Preprints of the 21st IFAC World Congress (Virtual)
Berlin, Germany, July 12-17, 2020

Sparsity Preserving Discretization With Error Bounds ⋆

James Anderson * **Nikolai Matni** ** **Yuxiao Chen** *

* *California Institute of Technology, Pasadena, CA 91125 USA*
james@caltech.edu, chenyx@caltech.edu
** *University of Pennsylvania, Philadelphia, PA 19104 USA*
nmatni@seas.upenn.edu

Abstract: Typically when designing distributed controllers it is assumed that the state-space model of the plant consists of sparse matrices. However, in the discrete-time setting, if one begins with a continuous-time model, the discretization process annihilates any sparsity in the model. In this work we propose a discretization procedure that maintains the sparsity of the continuous-time model. We show that this discretization out-performs a simple truncation method in terms of its ability to approximate the "ground truth" model. Leveraging results from numerical analysis we are also be able to upper-bound the error between the dense discretization and our method. Furthermore, we show that in a robust control setting we can design a distributed controller on the approximate (sparse) model that stabilizes the dense model.

Keywords: Decentralized and distributed control; Discretization; Numerical methods

1. INTRODUCTION

Consider the continuous-time LTI model

$$\dot{x}(t) = \hat{A}x(t) + \hat{B}_1 w(t) + \hat{B}_2 u(t), \tag{1}$$

where $x(t) \in \mathbb{R}^n, w(t) \in \mathbb{R}^{n_w}$, and $u(t) \in \mathbb{R}^{n_u}$ are the state, disturbance, and control vectors at time t, and the task of discretizing it to take the form

$$x_{k+1} = Ax_k + B_1 w_k + B_2 u_k, \tag{2}$$

where we assume k is the integer sequence $k = 0, 1, 2, \ldots$ of sampling points and that a zero-order-hold scheme is used.

Discretization, the process of converting (1) to (2), is a well studied topic and numerous methods have been proposed, however all methods we have encountered destroy the sparsity patterns in $(\hat{A}, \hat{B}_1, \hat{B}_2)$ when constructing (A, B_1, B_2). This is unfortunate as the sparsity patterns typically encode some sort of graph or network structure in the physical system. In this work, motivated by distributed control synthesis, we seek to construct sparse discretizations of (1) that respect the network structure of the continuous-time model and are close (in norm) to the "true" discrete models.

2. BACKGROUND

2.1 Discretization

The zero-order-hold sampling method (see for example Chen and Francis (2012)) for discretization maps $(\hat{A}, \hat{B}_1, \hat{B}_2)$ to (A, B_1, B_2) by specifying a sample rate $\tau > 0$ and setting

$$A = e^{\hat{A}\tau}, \quad B_i = \int_0^\tau e^{\hat{A}\lambda} \mathrm{d}\lambda \hat{B}_i, \quad i \in \{1, 2\}, \tag{3}$$

where e^X defines the exponential

$$e^X = \sum_{k=0}^{\infty} \frac{1}{k!} X^k$$

of a matrix $X \in \mathbb{C}^{n \times n}$. When $\hat{A}$ is non-singular the expression for B_i reduces to $\hat{A}^{-1}(e^{\hat{A}\tau} - I)\hat{B}_i$. Define $B = [B_1, B_2]$, a simple method for computing (A, B) (which is applicable when the $\hat{A}$ is non-singular) was derived by Van Loan (1978) and proceeds as follows; Define the matrix

$$\Psi = \begin{bmatrix} -\hat{A}^T & I & 0 & 0 \\ 0 & -\hat{A}^T & Q & 0 \\ 0 & 0 & \hat{A} & \hat{B} \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

where $Q = Q^T$ is an $n \times n$ real matrix and compute the exponential $e^{\Psi\tau}$ using an appropriate method c.f. Moler and Loan (2003). The exponential takes the form

$$e^{\Psi\tau} = \begin{bmatrix} F_1(\tau) & G_1(\tau) & \star & \star \\ 0 & F_2(\tau) & G_2(\tau) & \star \\ 0 & 0 & F_3(\tau) & G_3(\tau) \\ 0 & 0 & 0 & F_4(\tau) \end{bmatrix},$$

where $F_3(\tau) = e^{\hat{A}\tau}$ and $G_3(\tau) = \int_0^\tau e^{\hat{A}(\tau-\lambda)} \mathrm{d}\lambda \hat{B}$. Elements denoted by $\star$ have analytic expressions but are not required here, the remaining F_i and G_i functions are structurally similar to the case of $i = 3$ given above. The reader is referred to Van Loan (1978) for the full details.

An alternative method to the sample-and-hold approach is to take a bi-linear transformation (often referred to as

⋆ J. Anderson and Y. Chen are supported by PNNL on grant 424858. J. Anderson is additionally supported by ARPA-E through the GRID DATA program. N. Matni is supported in part by ONR awards N00014- 17-1-2191 and N00014-18-1-2833 and the DARPA Assured Autonomy (FA8750-18-C-0101) and Lagrange (W911NF-16-1-0552) programs.

Tustin's method) Chen and Francis (2012); Jiang et al. (2014). In this case $(\hat{A}, \hat{B})$ are mapped to (A, B) by

$$A = \left(I - \frac{\tau}{2}\hat{A}\right)^{-1}\left(I + \frac{\tau}{2}\hat{A}\right), \quad B = \frac{\tau}{2}\left(I - \frac{\tau}{2}\hat{A}\right)^{-1}\hat{B},$$

provided that the necessary inverse exists.

It should be clear that any non-trivial sparsity patterns in the system matrices in (1) will be lost if (2) is obtained from either (3) or the bilinear-transformation method as both involve computing a matrix exponential or taking an inverse.

We focus on approximate discretization based on the transformation (3). Formally, we would like to construct a matrix F such that for a sparse matrix X, $e^X \approx F$ with $\|e^X - F\| < \delta$, where δ is known a priori and F is sparse.

It should be pointed out that computing approximations of the matrix exponential is a mature topic in numerical analysis. In general these methods look to compute the exponential in a computationally efficient and stable manner rather than preserve any sparsity pattern (Moler and Loan, 2003). Rational approximations seek to replace e^X with $p(X)/q(X)$ where p and q are matrix-valued polynomial functions (Higham, 2005). Spectral methods are perhaps intuitively the simplest methods; Let $X = VDV^{-1}$ where D is a diagonal matrix and V is a matrix formed from the eigenvectors of X, then $e^X = Ve^DV^{-1}$. For matrices with dependent eigenvectors, other factorizations can be used. Clearly neither of these methods attempts to produce a sparse exponential. A related problem that arises in numerical solutions of linear ordinary differential equations is that of computing the *action* $v \mapsto e^{At}v$ (Al-Mohy and Higham, 2011). Krylov subspace methods (Saad, 1992) attempt to approximate this action but again it is not clear how one could incorporate sparsity constraints and obtain error bounds using such a method.

In this work we will make use of the fact that the matrices we care most about are likely banded. Such matrices are well known to have favourable localization properties, i.e. if X is banded, then for many functions, $f(X)$, the entries far from the diagonal decay exponentially fast (Benzi and Razouk, 2007). In particular, we will make use of the work of Iserles (2000) that deals with the matrix exponential function.

2.2 Distributed Control

In this section we briefly review the System Level Synthesis (SLS) framework (Wang et al., 2019; Doyle et al., 2017; Anderson et al., 2019) for solving distributed control problems. We consider the state-feedback problem where the plant $\mathbf{P}$ models the state dynamics given by (2) augmented with the error signal $\bar{z}_k = C_1x_k + D_{11}w_k + D_{12}u_k$. Compactly the system is written as

$$\mathbf{P}(z) = \left[\begin{array}{c|cc} A & B_1 & B_2 \\ \hline C_1 & D_{11} & D_{12} \\ I & 0 & 0 \end{array}\right],$$

which defines the map

$$\begin{bmatrix} \bar{\mathbf{z}}(z) \\ \mathbf{y}(z) \end{bmatrix} = \mathbf{P}(z)\begin{bmatrix} \mathbf{w}(z) \\ \mathbf{u}(z) \end{bmatrix}.$$

We seek to design a controller $\mathbf{u}(z) = \mathbf{K}(z)\mathbf{y}(z) = \mathbf{K}(z)\mathbf{x}(z)$. For the remainder of the paper we drop the dependence on z from our notation and simply use boldface symbols to denote signals in the z-domain.

Unlike classical control synthesis methods which seek to design controllers that minimize the norm of the map from $\mathbf{w}$ to $\bar{\mathbf{z}}$, SLS controllers work with the closed-loop system response which maps $\boldsymbol{\delta}_x$ to $(\mathbf{x}, \mathbf{u})$ where $\boldsymbol{\delta}_x = B_1\mathbf{w}$. The system response is described by

$$\begin{bmatrix} \mathbf{x} \\ \mathbf{u} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\Phi}_x \\ \boldsymbol{\Phi}_u \end{bmatrix}\boldsymbol{\delta}_{\boldsymbol{x}},$$

where

$$\boldsymbol{\Phi}_x = (zI - A - B_2\mathbf{K})^{-1}, \ \boldsymbol{\Phi}_u = \mathbf{K}\boldsymbol{\Phi}_x^{-1}. \tag{4}$$

The proceeding theorem parameterizes all achievable closed-loop system responses and provides a realization of an internally stabilizing controller Wang et al. (2019).

Theorem 2.1. Consider the LTI system (2), evolving under a dynamic state-feedback control policy $\mathbf{u} = \mathbf{K}\mathbf{x}$. The following statements are true:

(1) The affine subspace defined by

$$[zI - A \quad -B_2]\begin{bmatrix} \boldsymbol{\Phi}_x \\ \boldsymbol{\Phi}_u \end{bmatrix} = I, \quad \boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u \in \frac{1}{z}\mathcal{RH}_\infty \tag{5}$$

(where $z^{-1}\mathcal{RH}_\infty$ denotes the space of strictly proper and stable transfer matrice) parameterizes all system responses from $\boldsymbol{\delta}_{\boldsymbol{x}}$ to $(\mathbf{x}, \mathbf{u})$ as defined in (4), achievable by an internally stabilizing state feedback controller $\mathbf{K}$.

(2) For any transfer matrices $\{\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u\}$ satisfying (5), the controller $\mathbf{K} = \boldsymbol{\Phi}_u\boldsymbol{\Phi}_x^{-1}$ is internally stabilizing and achieves the desired system response (4).

The significance of Theorem 2.1 is that (5) provides an *affine characterization* of *all* achievable system responses. In recent work it was shown that if the affine expression in (5) is not satisfied, it is still possible to construct a stabilizing controller based on an *approximate* system response (Matni et al., 2017).

Theorem 2.2. Let $(\hat{\boldsymbol{\Phi}}_x, \hat{\boldsymbol{\Phi}}_u, \boldsymbol{\Delta})$ with $\hat{\boldsymbol{\Phi}}_x, \hat{\boldsymbol{\Phi}}_x \in \frac{1}{z}\mathcal{RH}_\infty$ be a solution to

$$[zI - A \ -B_2]\begin{bmatrix} \hat{\boldsymbol{\Phi}}_x \\ \hat{\boldsymbol{\Phi}}_u \end{bmatrix} = I + \boldsymbol{\Delta}. \tag{6}$$

Then, the controller $\mathbf{K} = \hat{\boldsymbol{\Phi}}_u\hat{\boldsymbol{\Phi}}_x^{-1}$ internally stabilizes the system (A, B_2) if and only if $(I + \boldsymbol{\Delta})^{-1}$ is stable. Furthermore, the actual system responses achieved are given by

$$\begin{bmatrix} \mathbf{x} \\ \mathbf{u} \end{bmatrix} = \begin{bmatrix} \hat{\boldsymbol{\Phi}}_x \\ \hat{\boldsymbol{\Phi}}_u \end{bmatrix}(I + \boldsymbol{\Delta})^{-1}\boldsymbol{\delta}_{\boldsymbol{x}}.$$

Theorem 2.2 forms the basis of what we term a system level synthesis problem, i.e., a mathematical program that returns a distributed optimal controller. The standard SLS problem is

$$\begin{aligned} \text{minimize}_{\gamma\in[0,1)} \ \underset{\hat{\boldsymbol{\Phi}}_x, \hat{\boldsymbol{\Phi}}_u, \boldsymbol{\Delta}}{\text{minimize}} \quad & g(\hat{\boldsymbol{\Phi}}_x, \hat{\boldsymbol{\Phi}}_u) \\ \text{subject to} \quad & \hat{\boldsymbol{\Phi}}_x, \hat{\boldsymbol{\Phi}}_x \in \frac{1}{z}\mathcal{RH}_\infty, \\ & \|\boldsymbol{\Delta}\| < \gamma, \\ & (6), \quad \begin{bmatrix} \hat{\boldsymbol{\Phi}}_x \\ \hat{\boldsymbol{\Phi}}_u \end{bmatrix} \in \mathcal{S}. \end{aligned} \tag{7}$$

The SLS problem (7) is quasi-convex; the inner problem is convex, and so a simple bisection on γ suffices. Note that the choice of norm on $\mathbf{\Delta}$ must be an induced-norm, this is a sufficient condition for $(I+\mathbf{\Delta})^{-1}$ to be stable. The cost functional g is chosen to be

$$\left\| \begin{bmatrix} C_1 & D_{12} \end{bmatrix} \begin{bmatrix} \hat{\mathbf{\Phi}}_x \\ \hat{\mathbf{\Phi}}_u \end{bmatrix} \right\|_{\mathcal{H}_2}^2.$$

Finally the constraint $\hat{\mathbf{\Phi}} \in \mathcal{S}$ encodes locality constraints on $\{\hat{\mathbf{\Phi}}_x, \hat{\mathbf{\Phi}}_u\}$. Further details on locality constraints can be found in Anderson et al. (2019); broadly speaking these constraints ensure that when a disturbance hits the system, the effect does not propagate to parts of the system far away. Similarly, the affect of the disturbance will not persist for long periods of time. In both cases, locality can be imposed on $\{\hat{\mathbf{\Phi}}_x, \hat{\mathbf{\Phi}}_u\}$ by constraining their spectral elements. For example, we can write $\hat{\mathbf{\Phi}}_x = \sum_{k=0}^{\infty} z^{-k} \hat{\Phi}_x[k]$, where $\Phi_x[k]$ is a real matrix and is known as the k^{th} spectral element. Localizing $\hat{\mathbf{\Phi}}_x$ (and/or $\hat{\mathbf{\Phi}}_u$) in time corresponds to ensuring the spectral elements are zero matrices for all $k > T$, for some positive integer T. Spatial locality is enforced by controlling the sparsity of the elements of each of the spectral components. In this context, constraining the spectral elements of $\{\hat{\mathbf{\Phi}}_x, \hat{\mathbf{\Phi}}_u\}$ to have a bandwidth of at most d, (where d is a positive integer) ensures that if a disturbance hits node i then only the nodes $i \pm d$ feel the disturbance. By defining the bandwidth of the spectral elements over $k = 1, \ldots, T$ we can effectively constrain the closed-loop system response.

One of the advantages of constraining the closed-loop response is that the constraints on the spectral elements typically take the form of convex constraints. Moreover, objective functionals and locality constraints that satisfy a separability property allow for the control synthesis problem (7) to be decomposed into smaller finite-dimensional convex sub-problems and solved in parallel (Wang et al., 2018). In contrast, classical methods scale poorly with system size as the underlying optimization problem grows rapidly with the magnitude of the state of the system.

Norms that we will consider are the $\mathcal{H}_2, L_1$, and the $\mathcal{E}_1$ norms. The L_1-norm of a transfer matrix $\mathbf{G}(z)$ is the induced $\ell_\infty \to \ell_\infty$ norm which can be computed as

$$\|\mathbf{G}\|_{L_1} = \max_{1 \le i \le m} \sum_{j=1}^{n} \sum_{k=0}^{\infty} |G_{ij}[k]|, \quad \|\mathbf{G}\|_{\mathcal{E}} = \|\mathbf{G}^T\|_{L_1}.$$

3. RESULTS

The $\mathbf{\Delta}$ block that appears in Theorem 2.2 allows us to formulate robust control problems, i.e. the design of a controller that stabilizes a the plant over all realizations of an uncertainty set. In particular, we will consider the ground-truth model to be of the form

$$A = A_n + \mathbf{\Delta}_A, \quad B_2 = B_n + \mathbf{\Delta}_B, \tag{8}$$

where (A_n, B_n) represent the nominal system data and $\mathbf{\Delta}_A, \mathbf{\Delta}_B$, represent perturbation matrices. The robust problem we are interested in is: given upper-bounds on $\|\mathbf{\Delta}_A\|, \|\mathbf{\Delta}_B\|$, can we synthesize a robustly stabilizing controller from the nominal system matrices? Consider the case where the nominal system satisfies

$$\begin{bmatrix} zI - A_n & -B_n \end{bmatrix} \begin{bmatrix} \mathbf{\Phi}_x \\ \mathbf{\Phi}_u \end{bmatrix} = I, \tag{9}$$

then

$$\begin{bmatrix} zI - A & -B_2 \end{bmatrix} \begin{bmatrix} \mathbf{\Phi}_x \\ \mathbf{\Phi}_u \end{bmatrix} = I + \underbrace{\begin{bmatrix} \mathbf{\Delta_A} & \mathbf{\Delta}_B \end{bmatrix} \begin{bmatrix} \mathbf{\Phi}_x \\ \mathbf{\Phi}_u \end{bmatrix}}_{:= \mathbf{\Delta}},$$

where $\mathbf{\Delta}$ above is as defined in Theorem 2.2. Dean et al. (2017) show that when when $\|\mathbf{\Delta}_A\|_2 \le \rho_{\boldsymbol{A}}, \|\mathbf{\Delta}_B\|_2 \le \rho_{\boldsymbol{B}}$ then a tractable upper-bound for $\|\mathbf{\Delta}\|_{\mathcal{H}_\infty}$ is possible.

Lemma 3.1. (Dean et al. (2017)). For $\mathbf{\Delta}$ as defined above and with the bounds $\rho_{\boldsymbol{A}}, \rho_{\boldsymbol{B}}$, for any $\alpha \in (0,1)$

$$\|\mathbf{\Delta}\|_{\mathcal{H}_\infty} \le \left\| \begin{bmatrix} \frac{\rho_{\boldsymbol{A}}}{\sqrt{\alpha}} \mathbf{\Phi}_x \\ \frac{\rho_{\boldsymbol{B}}}{\sqrt{1-\alpha}} \mathbf{\Phi}_u \end{bmatrix} \right\|_{\mathcal{H}_\infty}.$$

The upper-bound above is then used in place of the constraint $\|\mathbf{\Delta}\|_{\mathcal{H}_\infty} < \gamma$ in the SLS problem (7). The lemmas that follow are in the same spirit as the bound above, but are for the L_1 and $\mathcal{E}_1$-norms. Assume that $\|\mathbf{\Delta}_A\|_\infty \le \epsilon_{\boldsymbol{A}}$ and $\|\mathbf{\Delta}_B\|_\infty \le \epsilon_{\boldsymbol{B}}$.

Lemma 3.2. Given the scalar bounds $\epsilon_{\boldsymbol{A}}$ and $\epsilon_{\boldsymbol{B}}$, then for all $\alpha \in (0,1)$

$$\begin{aligned} \|\mathbf{\Delta}\|_{L_1} &\le \left\| \begin{bmatrix} \frac{\epsilon_{\boldsymbol{A}}}{\alpha} \mathbf{\Phi}_x \\ \frac{\epsilon_{\boldsymbol{B}}}{1-\alpha} \mathbf{\Phi}_u \end{bmatrix} \right\|_{L_1} \\ &= \max\left\{ \frac{\epsilon_{\boldsymbol{A}}}{\alpha} \|\mathbf{\Phi}_x\|_{L_1}, \frac{\epsilon_{\boldsymbol{B}}}{1-\alpha} \|\mathbf{\Phi}_u\|_{L_1} \right\}. \end{aligned}$$

Proof. From the definition of $\mathbf{\Delta}$ we have

$$\begin{aligned} \|\mathbf{\Delta}\|_{L_1} &= \left\| \begin{bmatrix} \frac{\alpha}{\epsilon_{\boldsymbol{A}}} \mathbf{\Delta}_A & \frac{(1-\alpha)}{\epsilon_{\boldsymbol{B}}} \mathbf{\Delta}_B \end{bmatrix} \begin{bmatrix} \frac{\epsilon_{\boldsymbol{A}}}{\alpha} \mathbf{\Phi}_x \\ \frac{\epsilon_{\boldsymbol{B}}}{1-\alpha} \mathbf{\Phi}_u \end{bmatrix} \right\|_{L_1} \\ &\le \left\| \begin{bmatrix} \frac{\alpha}{\epsilon_{\boldsymbol{A}}} \mathbf{\Delta}_A & \frac{(1-\alpha)}{\epsilon_{\boldsymbol{B}}} \mathbf{\Delta}_B \end{bmatrix} \right\|_{L_1} \left\| \begin{bmatrix} \frac{\epsilon_{\boldsymbol{A}}}{\alpha} \mathbf{\Phi}_x \\ \frac{\epsilon_{\boldsymbol{B}}}{1-\alpha} \mathbf{\Phi}_u \end{bmatrix} \right\|_{L_1} \\ &\le \left(\left\| \frac{\alpha}{\epsilon_{\boldsymbol{A}}} \mathbf{\Delta}_A \right\|_{L_1} + \left\| \frac{(1-\alpha)}{\epsilon_{\boldsymbol{B}}} \mathbf{\Delta}_B \right\|_{L_1} \right) \left\| \begin{bmatrix} \frac{\epsilon_{\boldsymbol{A}}}{\alpha} \mathbf{\Phi}_x \\ \frac{\epsilon_{\boldsymbol{B}}}{1-\alpha} \mathbf{\Phi}_u \end{bmatrix} \right\|_{L_1} \\ &= \left(\frac{\alpha}{\epsilon_{\boldsymbol{A}}} \|\mathbf{\Delta}_A\|_{L_1} + \frac{(1-\alpha)}{\epsilon_{\boldsymbol{B}}} \|\mathbf{\Delta}_B\|_{L_1} \right) \left\| \begin{bmatrix} \frac{\epsilon_{\boldsymbol{A}}}{\alpha} \mathbf{\Phi}_x \\ \frac{\epsilon_{\boldsymbol{B}}}{1-\alpha} \mathbf{\Phi}_u \end{bmatrix} \right\|_{L_1} \\ &\le \left(\frac{\alpha}{\epsilon_{\boldsymbol{A}}} \epsilon_{\boldsymbol{A}} + \frac{(1-\alpha)}{\epsilon_{\boldsymbol{B}}} \epsilon_{\boldsymbol{B}} \right) \left\| \begin{bmatrix} \frac{\epsilon_{\boldsymbol{A}}}{\alpha} \mathbf{\Phi}_x \\ \frac{\epsilon_{\boldsymbol{B}}}{1-\alpha} \mathbf{\Phi}_u \end{bmatrix} \right\|_{L_1} \\ &= \left\| \begin{bmatrix} \frac{\epsilon_{\boldsymbol{A}}}{\alpha} \mathbf{\Phi}_x \\ \frac{\epsilon_{\boldsymbol{B}}}{1-\alpha} \mathbf{\Phi}_u \end{bmatrix} \right\|_{L_1} \\ &= \max\left\{ \frac{\epsilon_{\boldsymbol{A}}}{\alpha} \|\mathbf{\Phi}_x\|_{L_1}, \frac{\epsilon_{\boldsymbol{B}}}{1-\alpha} \|\mathbf{\Phi}_u\|_{L_1} \right\}. \end{aligned}$$

The first inequality results from applying the triangle inequality, the second comes from the fact that for $B \in \mathbb{C}^{m \times n_1}$ and $C \in \mathbb{C}^{m \times n_2}$

$$\|[B\ C]\|_\infty \le \|B\|_\infty + \|C\|_\infty,$$

(and noting that L_1-norm of a constant matrix is simply the standard matrix ∞-norm.) and the final inequality comes from substituting in the upper bounds. The last equality follows since

$$\left\| \begin{bmatrix} X \\ Y \end{bmatrix} \right\|_{L_1} = \max\left\{ \|X\|_{L_1}, \|Y\|_{L_1} \right\}.$$

Now, assume instead that we have the bounds $\|\boldsymbol{\Delta}_A\|_1 \leq \nu_{\boldsymbol{A}}$, $\|\boldsymbol{\Delta}_B\|_1 \leq \nu_{\boldsymbol{B}}$.

Lemma 3.3. Given the scalar bounds $\nu_{\mathbf{A}}, \nu_{\mathbf{B}}$, the following bound holds:

$$\|\boldsymbol{\Delta}\|_{\mathcal{E}_1} \leq \left\| \begin{bmatrix} \nu_{\mathbf{A}} \boldsymbol{\Phi}_x \\ \nu_{\mathbf{B}} \boldsymbol{\Phi}_u \end{bmatrix} \right\|_{\mathcal{E}_1}.$$

Proof. Follows similar arguments to the proof provided for Lemma 3.2.

Thus Lemmas 3.1–3.3 provide upper-bounds that incorporate the perturbation magnitude (in terms of three different norms of $\boldsymbol{\Delta}_A, \boldsymbol{\Delta}_B$) into the SLS problem (7).

From the definition of the matrix exponential, the most obvious way of constructing a sparse approximation is via truncation. Given a sample parameter $\tau > 0$, truncating the exponential after $k = 1$ terms we have

$$A_\tau^{trunc} := I + A\tau \approx e^{\hat{A}\tau}.$$

Clearly A_τ^{trunc} has the same sparsity as $\hat{A}$. Indeed, this approximation is the basis for first-order Euler methods for solving initial point problems. It is well known that truncation methods do not preserve stability. To reconcile notation, we have $A_\tau^{trunc} = A_n$, where A_n is defined in (8).

One advantage of this approximation method is that there is a clean bound for the error.

Theorem 3.4. (Moler and Loan (2003)). Given a matrix $\hat{A} \in \mathbb{C}^{n\times n}$ and a constant $\tau > 0$, then

$$\|A_\tau^{trunc} - e^{\hat{A}\tau}\|_2 \leq \left(\frac{\|\hat{A}\|_2^2 \tau^2}{2} \right) \left(\frac{1}{1 - \frac{\tau}{3}\|\hat{A}\|_2} \right).$$

This upper-bound would take the form of $\rho_{\boldsymbol{A}}$ in our robust control problem.

In practice, the bound Theorem 3.4 is sharp for small τ and useless for large values. We now propose a second more accurate method for computing an estimate of the matrix exponential. The method is simple; compute the full matrix exponential and then project it onto the support of $\hat{A}$. Formally, we define support matrix H as

$$[\mathbf{supp}(H)]_{ij} = \begin{cases} 1 \text{ if } H_{ij} \neq 0 \\ 0 \text{ otherwise} \end{cases}.$$

For a fixed constant τ, the *projected exponential* is given by

$$A_\tau^{proj} := \mathbf{supp}(|\hat{A}| + |I|) \circ e^{\hat{A}\tau} \approx e^{\hat{A}\tau}$$

where $\circ$ denotes the Hadamard (element-wise) product between two matrices of equal dimension. Note that the **supp** operation binds before the Hadamard product.

For the projection of $\hat{B}_i$, we use $\mathbf{supp}(|\hat{A}| + |I|)$ to approximate $\mathbf{supp}(\int_0^\tau e^{\hat{A}\lambda} \mathrm{d}\lambda)$ and have:

$$B_{i,\tau}^{proj} := \mathbf{supp}((|\hat{A}| + |I|) \cdot |\hat{B}_i|) \circ \int_0^\tau e^{\hat{A}\lambda} \mathrm{d}\lambda \hat{B}_i.$$

In order to obtain bounds on the approximation error we will need to impose some structure on the continuous drift-matrix $\hat{A}$.

Definition 1. Given a matrix $Y \in \mathbb{C}^{n\times n}$, the bandwidth of Y is the smallest integer s such that $Y_{jk} = 0$ for all $|j - k| \geq s + 1$. We use $Y_{(s)}$ to denote that Y has has bandwidth s.

According to Definition 1 a diagonal matrix has a bandwidth of zero, a tridiagonal matrix has bandwidth $s = 1$, etc.

Let $E_{ij} = e_i e_j^T$ where e_i is the standard i^{th} basis vector for $\mathbb{R}^n$. Then a bandwidth s matrix can be extracted from a dense matrix via $Y_{(s)} = \sum_{|i-j|\leq s} E_{ii} Y E_{jj}$, where the summation is taken over all pairs $\{i, j\}$ that satisfy $|i - j| \leq s$.

Assumption 1. The $n \times n$ drift matrix $\hat{A}$ from (1) is a banded matrix, or, there exists a permutation matrix Π such that $\Pi\hat{A}\Pi^{-1}$ is banded.

We note that this is not a major assumption, many applications produce matrices naturally in this form, examples can be found in the online catalogue described in Davis and Hu (2011) and in the recent work Vo and Sidje (2017). While the exponential of a banded matrix is formally dense, numerical analysts have noted that elements away from the diagonal decay rapidly. Specifically, "*provided that A is a banded matrix, e^A is itself within an exceedingly small distance from a banded matrix.*", Iserles (2000). This observation is illustrated qualitatively in Figure 1.

Recall the (non sub-multiplicative) max-norm of a matrix X defined as

$$\|X\|_{\max} := \max_{k,l} \quad |X_{kl}|.$$

The following result provides a bound on the elements of the matrix exponential obtained from a banded matrix.

Theorem 3.5. (Iserles (2000)). Let $A = e^{\hat{A}\tau}$ for some $\tau > 0$, where $\hat{A}$ is a banded matrix with bandwidth $s \geq 1$. Let $\alpha = \|\hat{A}\tau\|_{\max}$.Then for $|i - j| \gg 1$,

$$|A_{ij}| \leq \underbrace{\left(\frac{\alpha s}{|i-j|} \right)^{\frac{|i-j|}{s}} \left[e^{\frac{|i-j|}{s}} - \sum_{m=0}^{|i-j|-1} \frac{(|i-j|/s)^m}{m!} \right]}_{\mathcal{B}_{ij}(\alpha, s)}.$$

Under the assumption that $\hat{A}$ from (1) is banded (with bandwidth s), we define $A = A_n + \boldsymbol{\Delta}_A$ where $A_n = A_\tau^{proj}$, thus A_n is banded and $\boldsymbol{\Delta}_A$ is is the complement of a banded matrix. Appealing to Theorem 3.5 we can derive upper-bounds on $\|\boldsymbol{\Delta}_A\|$ for various norms. Moreover, the upper-bounds can be computed using only scalar operations. For any choice of norm it follows that

$$\begin{aligned} \|\boldsymbol{\Delta}_A\| &= \left\| \sum_{|i-j|>s} E_{ii} e^{\hat{A}\tau} E_{jj} \right\| \leq \sum_{|i-j|>s} \left\| E_{ii} e^{\hat{A}\tau} E_{jj} \right\| \\ &= \sum_{|i-j|>s} |[e^{\hat{A}\tau}]_{ij}| \leq \sum_{|i-j|>s} \mathcal{B}_{ij}(\alpha, s). \end{aligned}$$

Define $\rho_{\boldsymbol{A}}^\star, \epsilon_{\boldsymbol{A}}^\star$, and $\nu_{\boldsymbol{A}}^\star$ to be $\|\boldsymbol{\Delta}_A\|$ for the $2, \infty$, and 1-norm respectively.

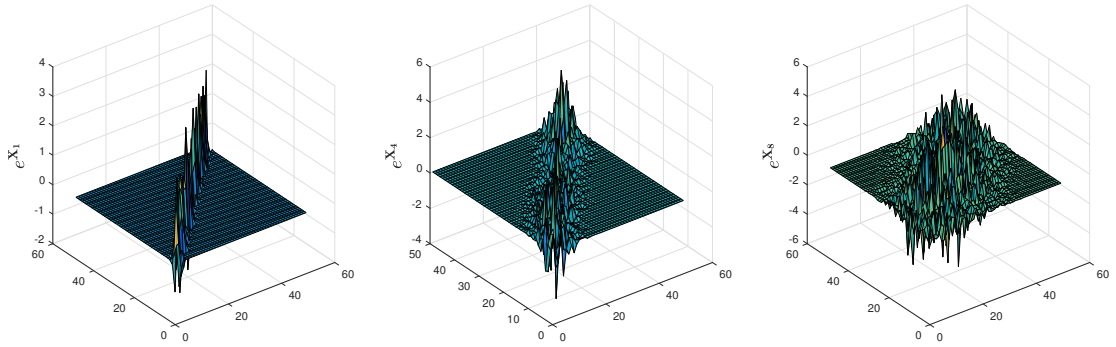

Fig. 1. The matrix exponential of three banded matrices X_1, X_4, and X_8 with bandwidth $1, 4$, and 8 are shown from left to right respectively.

Theorem 3.6. Let $\rho^\star_{\boldsymbol{A}}, \epsilon^\star_{\boldsymbol{A}}$, and $\nu^\star_{\boldsymbol{A}}$ be defined as above. Given a matrix $\hat{A} \in \mathbb{R}^{n\times n}$ of bandwidth s. Then $\rho^\star_{\boldsymbol{A}}, \epsilon^\star_{\boldsymbol{A}}$, and $\nu^\star_{\boldsymbol{A}}$ can be upper-bounded by computations that involve only scalar operations on the indices i, j and parameters α and s. Specifically,

$$\rho^\star_{\boldsymbol{A}} \leq \sum_{|i-j|>s} \mathcal{B}_{ij}(\alpha, s), \quad \epsilon^\star_{\boldsymbol{A}} \leq \max_{1\leq i\leq n} \sum_{\substack{j=1 \\ |i-j|>s}}^{n} \mathcal{B}_{ij}(\alpha, s),$$

$$\text{and } \nu^\star_{\boldsymbol{A}} \leq \max_{1\leq j\leq n} \sum_{\substack{i=1 \\ |i-j|>s}}^{n} \mathcal{B}_{ij}(\alpha, s).$$

Proof. The proof follows by applying Theorem 3.5, the inequalities derived above, and the definition of the 1 and ∞ norms.

Similar bounds for the norms of $\|\boldsymbol{\Delta}_B\|$ are easily derived. In Figure 2 the upper-bound on $\|\boldsymbol{\Delta}_A\|_2$ from Theorem 3.6 is compared to the true value for matrices of bandwidth 4, i.e. $\hat{A}_{(4)} \in \mathbb{R}^{n\times n}$. We show how the upper-bound changes as a function of the dimension of the matrix where $n \in \{20, 40, 60, 100, 200, 500, 1000\}$. It can be clearly seen, that even for large matrices, the estimates $\epsilon_{\boldsymbol{A}}$ are easily small enough to be useful.

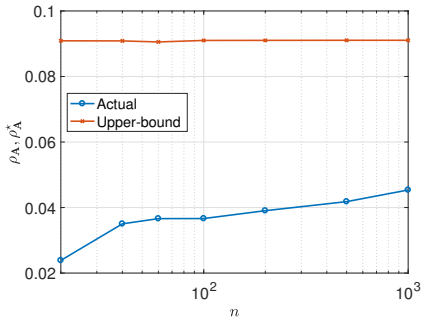

Fig. 2. The error magnitude $\|\boldsymbol{\Delta}_A\|_2 = \epsilon^\star_{\boldsymbol{A}}$ and the upper-bound for $\|\boldsymbol{\Delta}_A\|_2$ given by $\epsilon_{\boldsymbol{A}}$ from Theorem 3.6.

4. EXAMPLES

We demonstrate the proposed method on a power grid control example. The model comes from Zimmerman et al. (2011) and has 57 buses with 7 generator buses. The network topology is shown in Figure 3.

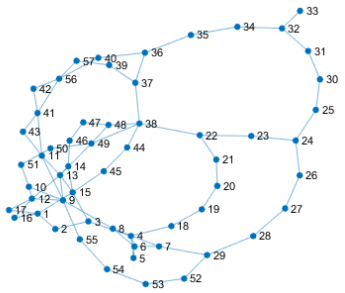

Fig. 3. Topology of the power network. A disturbance hits the frequency state of node 3.

The power grid dynamics are described by the coupled system of ordinary differential equations:

$$\begin{aligned} \dot{\theta}_i &= \omega_i, \\ M_i\dot{\omega}_i &= -D_i\omega_i - d_i - u_i - \sum_{j\in\mathcal{N}_i} H_{ij}(\theta_i - \theta_j), i \in \mathcal{G} \\ 0 &= -D_i\omega_i - d_i - u_i - \sum_{j\in\mathcal{N}_i} H_{ij}(\theta_i - \theta_j), i \in \mathcal{L}, \end{aligned} \tag{10}$$

where θ_i and ω_i are the phase angle and frequency of the voltage at bus i, d_i is the uncontrollable load at bus i which is treated as a disturbance. u_i is the controllable load, which is used to regulate bus i. $\mathcal{G}$ and $\mathcal{L}$ represent the set of generator buses and the set of pure load buses. In this example $\mathcal{G} = \{1, 2, 3, 6, 8, 9, 12\}$ and $\mathcal{L} = \{1, \ldots, 57\} \setminus \mathcal{G}$. For a generator bus, M_i is the inertia and D_i is the "damping coefficient"; for a load bus, there is zero inertia and ω_i is determined by an algebraic equation. A generator bus is modeled with 2 states ($x_i = [\theta_i, \omega_i]^T$); and a load bus is modeled with 1 state ($x_i = \theta_i$). H_{ij} represents the sensitivity of the power flow to phase variations, it is nonzero when bus i and bus j are neighbors. We assume an impulsive disturbance hits bus 3 and affects its frequency.

Using the sparse discretized model, the distributed controller derived from (7) is feasible for an FIR horizon $T \geq 5$ and the locality radius of $d \geq 4$. Informally this means that if a single impulsive disturbance hits a node, then only nodes within 3 "hops" of the affected node feel

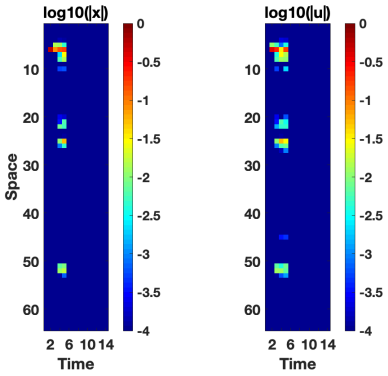

Fig. 4. The localized system response implemented on the sparse model. Clearly the response to the disturbance is localized in time and space.

the disturbance. In this example we have that $\|\mathbf{\Delta}_A\| = \{0.455, 0.394, 0.873\}$ and $\|\mathbf{\Delta}_B\| = \{0.001, 0.006, 0.001\}$ for the $1, 2$, and ∞ norms respectively (to 3 s.f.).

When using the exact discretized model obtained using (3), the SLS program (7) is only feasible if all nodes in the network area able to respond to the disturbance, this corresponds to a locality radius of which requires $d = 12$. The result is that localized distributed control is not possible (i.e. the SLS problem has no localization and is thus a centralized controller) if the underlying model (10) is converted to discrete time using standard methods. In Figure 4 the closed-loop response of the controller on the sparse model is plotted. In Figure 5, we show the system response when the controller is designed on the sparse model and implemented on the dense discrete model. Despite the model mismatch the robust controller still manages to localize the disturbance in time and space.

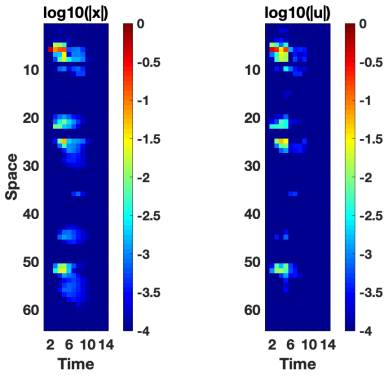

Fig. 5. The controller is designed on the sparse model and simulated on the dense model (3).

5. CONCLUSION

We have presented a simple projection based method for sparsity-preserving discretization of a continuous-time dynamical system. For the special case of banded matrices, bounds on the approximation error were derived and shown to perform well in practice. For non-banded systems, the existing bound is useful in certain sampling parameter regimes. In the banded case, the bounds were incorporated into the $\mathbf{\Delta}$ uncertainty parameter. The results were illustrated on a 57-bus power network where it was shown that an SLS distributed controller designed on the sparse approximation performs well when implemented on the dense "ground truth" model.

REFERENCES

Al-Mohy, A.H. and Higham, N.J. (2011). Computing the action of the matrix exponential, with an application to exponential integrators. *SIAM journal on scientific computing*, 33(2), 488–511.

Anderson, J., Doyle, J.C., Low, S.H., and Matni, N. (2019). System level synthesis. *Annual Reviews in Control*.

Benzi, M. and Razouk, N. (2007). Decay bounds and $o(n)$ algorithms for approximating functions of sparse matrices. *Electron. Trans. Numer. Anal*, 28, 16–39.

Chen, T. and Francis, B.A. (2012). *Optimal sampled-data control systems*. Springer Science & Business Media.

Davis, T.A. and Hu, Y. (2011). The University of Florida sparse matrix collection. *ACM Transactions on Mathematical Software (TOMS)*, 38(1), 1.

Dean, S. et al. (2017). On the sample complexity of the linear quadratic regulator. *arXiv preprint arXiv:1710.01688*.

Doyle, J.C., Matni, N., Wang, Y.S., Anderson, J., and Low, S.H. (2017). System level synthesis: A tutorial. In *2017 IEEE 56th Annual Conference on Decision and Control*, 2856–2867. IEEE.

Higham, N.J. (2005). The scaling and squaring method for the matrix exponential revisited. *SIAM Journal on Matrix Analysis and Applications*, 26(4), 1179–1193.

Iserles, A. (2000). How large is the exponential of a banded matrix? *New Zealand J. Math*, 29, 177–192.

Jiang, Y., Hu, X., and Wu, S. (2014). Transformation matrix for time discretization based on Tustin's method. *Mathematical Problems in Engineering*, 2014.

Loan, C.V. (1978). Computing integrals involving the matrix exponential. *IEEE transactions on automatic control*, 23(3), 395–404.

Matni, N., Wang, Y.S., and Anderson, J. (2017). Scalable system level synthesis for virtually localizable systems. In *2017 IEEE 56th Annual Conference on Decision and Control*, 3473–3480.

Moler, C. and Loan, C.V. (2003). Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM review*, 45(1), 3–49.

Saad, Y. (1992). Analysis of some Krylov subspace approximations to the matrix exponential operator. *SIAM Journal on Numerical Analysis*, 29(1), 209–228.

Vo, H.D. and Sidje, R.B. (2017). Approximating the large sparse matrix exponential using incomplete orthogonalization and Krylov subspaces of variable dimension. *Numerical Linear Algebra with Applications*, 24(3), e2090.

Wang, Y.S., Matni, N., and Doyle, J.C. (2018). Separable and localized system-level synthesis for large-scale systems. *IEEE Transactions on Automatic Control*, 63(12), 4234–4249.

Wang, Y.S., Matni, N., and Doyle, J.C. (2019). A system level approach to controller synthesis. *IEEE Transactions on Automatic Control*.

Zimmerman, R.D., Murillo-Sánchez, C.E., and Thomas, R.J. (2011). Matpower: Steady-state operations, planning, and analysis tools for power systems research and education. *IEEE Transactions on power systems*, 26(1), 12–19.