Large-Scale System Identification Using a Randomized SVD

Han Wang, James Anderson

Abstract—Learning a dynamical system from input/output data is a fundamental task in the control design pipeline. In the partially observed setting there are two components to identification: *parameter estimation* to learn the Markov parameters, and *system realization* to obtain a state space model. In both sub-problems it is implicitly assumed that standard numerical algorithms such as the singular value decomposition (SVD) can be easily and reliably computed. When trying to fit a high-dimensional model to data, even computing an SVD may be intractable. In this work we show that an approximate matrix factorization obtained using randomized methods can replace the standard SVD in the realization algorithm while maintaining the finite-sample performance and robustness guarantees of classical methods.

I. INTRODUCTION

We consider the problem of identifying a linear timeinvariant (LTI) system parameterized by the matrices $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, $C \in \mathbb{R}^{p \times n}$ and $D \in \mathbb{R}^{p \times m}$, that evolves according to

$$x_{t+1} = Ax_t + Bu_t + w_t$$

$$y_t = Cx_t + Du_t + v_t.$$
(1)

It is assumed that we have $N < \infty$ observations of the output signal $\{y_t^i\}_{t=0}^T$ and control signal $\{u_t^i\}_{t=0}^T$, where T denotes the length of the signal. The vectors $x_t \in \mathbb{R}^n, w_t \in \mathbb{R}^n$, and $v_t \in \mathbb{R}^p$ in (1) denote the system state, process noise, and measurement noise at time t, respectively.

In the fully observed setting, estimates for (A, B) can be obtained by solving ordinary least-squares (OLS) optimization problems. A series of recent papers [1–5] have derived non-asymptotic guarantees for ordinary least-squares (OLS) estimators. In the case of partially observed systems, which is conceptually more complicated than the fully observed case, OLS optimization can be used to estimate the *Markov parameters* associated with (1) from which the Ho-Kalman algorithm [6] can be employed to estimate the system parameters (A, B, C, D). The process of obtaining estimates of the system matrices from the Markov parameters is referred to as *system realization* which is the main focus of this paper. Using this framework, the authors of [3, 7–11] have derived non-asymptotic estimation error bounds for the system parameters which decay at a rate $O(\frac{1}{\sqrt{N}})$. Note that these papers make different assumptions about the stability, system order, and the number of required trajectories to excite the unknown system.

However, in contrast to the estimation error bounds, the computational complexity of system identification has received much less attention in the literature [12, 13]. Due to the fact that the OLS problem is convex, and the computational bulk of the Ho-Kalman Algorithm is a singular value decomposition (SVD), it is taken for granted that system identification can be carried out at scale. As mentioned in [12], with the increase of system dimension, the computational and storage costs of general control algorithms quickly become prohibitively large. This challenge motivates us to design algorithms that mitigate the "curse of dimensionality". In this paper, we aim to design an efficient and scalable system realization algorithm that can be deployed in the big data regime.

From the view of computational complexity, the system identification methods proposed in [3, 7, 10, 11] are not scalable since the size of the Hankel matrix increases quadratically with the length of output signal T and cubically with the system state dimension n. The result is the singular value decomposition used in the Ho-Kalman Algorithm cannot be computed. This quadratic/cubic dependence on the problem size greatly limits its application in large scale system identification problems.

Motivated by the limits of the scalability of numerical SVD computations, there has been a surge of work which has focussed on providing approximate, but more easily computable matrix factorizations. Thanks to advances in our understanding of random matrix theory and high dimensional probability (in particular, concentration of measure), *randomized methods* have been shown to provide an excellent balance between numerical implementation (in terms of storage requirements and computational cost) and accuracy of approximation (in theory and practice). Broadly speaking this field is referred to as randomized numerical linear algebra (RNLA), and we refer the reader to [14–16] and the references therein for an overview of the field.

The intuition is that randomized methods can produce efficient, unbiased approximations of deterministic operations while being numerically efficient to implement by exploiting modern computational architectures such as parallelization and streaming. The performance of the randomized SVD (RSVD) has been studied in many works [14–17] and has found applications in large-scale problems across machine

Han Wang and James Anderson are with the Department of Electrical Engineering, Columbia University in the City of New York (e-mail: {hw2786, james.anderson}@columbia.edu). HW is kindly supported by a Wei Family Foundation fellowship. JA is supported by DOE through award DE-SC0022234 and NSF through CAREER ECCS-2144634.

learning, statistics, and signal processing.

The main contribution of this work is a stochastic Ho-Kalman Algorithm, where the standard SVD (which constitutes the main computational bottleneck of the algorithm) is replaced with an RSVD algorithm, which trades off accuracy and robustness for speed. We show that the stochastic Ho-Kalman Algorithm achieves the same robustness guarantees as its deterministic, non-asymptotic version in expectation. However, it outperforms the deterministic algorithm in terms of speed/computational complexity, which is measured by the total number of required floating-point operations (flops) [[18], §C.1.1]. Compared with $O(pmn^3)$ flops required by the deterministic algorithm, the stochastic Ho-Kalman Algorithm only requires $O(pmn^2 \log n)$ flops.

II. PRELIMINARIES AND PROBLEM FORMULATION

Given a matrix $A \in \mathbb{C}^{m \times n}$, where \mathbb{C} is the set of complex numbers, ||A|| denotes the spectral norm and $||A||_F$ denotes the Frobenius norm, i.e., $||A|| = \sigma_1(A)$, where σ_1 is the maximum singular value of A, and $||A||_F = \sqrt{\text{Trace}(A^*A)}$, where A^* denotes the Hermitian transpose of A. The multivariate normal distribution with mean μ and covariance matrix Σ is denoted by $\mathcal{N}(\mu, \Sigma)$. A matrix is said to be *standard Gaussian* if every entry is drawn independently from $\mathcal{N}(0, 1)$.

A. Singular Value Decomposition

The singular value decomposition of the matrix $A \in \mathbb{C}^{m \times n}$, factors it as $A = U\Sigma V^*$, where $U \in \mathbb{C}^{m \times m}$ and $V \in \mathbb{C}^{n \times n}$ are orthonormal matrices, and Σ is an $m \times n$ real diagonal matrix with entries $\sigma_1, \sigma_2, \cdots, \sigma_n$ ordered such that $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_n \geq 0$. When A is real, so are U and V. The truncated SVD of A is given by $U_r \Sigma_r V_r^\top (r < \min\{m, n\})$, where the matrices U_r and V_r contain only the first r columns of U and V, and Σ_r contains only the first r singular values from Σ . According to the Eckart-Young theorem [19], the best rank-r approximation to A in the spectral norm or Frobenius norm is given by

$$A_{[r]} = \sum_{i=1}^{r} \sigma_i u_i v_i^{\top} \tag{2}$$

where u_i and v_i denote the i^{th} column of U and V, respectively. More precisely,

$$\underset{\operatorname{rank}(X) \le r}{\operatorname{minimize}} \|A - X\| = \sigma_{r+1}, \tag{3}$$

and a minimizer is given by $X^* = A_{[r]}$. The expression (3) concisely sums up the scalability issue we are concerned with: on the left hand side is non-convex optimization problem with no polynomial-time solution; on the right is a singular value which for large m and/or n cannot be computed. In the sequel we shall see how randomized methods use approximate factorizations to resolve these issues.

B. System Identification

We consider the problem of identifying a linear system model defined by (1) where $u_t \sim \mathcal{N}\left(0, \sigma_u^2 I_m\right), w_t \sim \mathcal{N}\left(0, \sigma_w^2 I_n\right)$, and $v_t \sim \mathcal{N}\left(0, \sigma_v^2 I_p\right)$. We further assume that the initial state variable $x_0 = 0_n$ (although the dimension, n, in unknown *a priori*). Under these assumptions, we generate N trajectories of length T. Data is recorded as

$$\mathcal{D}_{T}^{N} = \left\{ \left\{ y_{t}^{i}, u_{t}^{i} \right\} : 1 \le i \le N, 0 \le t \le T - 1 \right\},\$$

where *i* denotes *i*th trajectory and *t* denotes *t*th time-step in each trajectory. With the data \mathcal{D}_T^N , the system identification problem can be solved in two steps:

1) **Estimation:** Given \mathcal{D}_T^N , estimate the first T Markov parameters of the system which are defined as

 $G = [D, CB, CAB, \ldots, CA^{T-2}B] \in \mathbb{R}^{m \times Tp}.$

Ideally, the estimation algorithm will produce finite sample bounds of the form $||G - \hat{G}|| \le \epsilon(N, T)$, where \hat{G} is the estimate of G.

2) **Realization:** Given an estimated Markov parameter matrix \hat{G} , produce state-space matrices $(\hat{A}, \hat{B}, \hat{C}, \hat{D})$ with guarantees of the form $||A - \hat{A}|| \le \epsilon_A$, $||B - \hat{B}|| \le \epsilon_B$, etc. This is most commonly done using the Ho-Kalman algorithm.

The input/output trajectory $\{y_t, u_t\}_{t=0}^{T-1}$ is referred to as a *rollout*. There are two approaches to collecting data; single [7, 2, 20, 3, 21] and multi-rollout [22, 10, 5, 8]. As this paper focuses on the realization step, we can use either approach.

C. System realization via noise-free Markov matrix G

The Ho-Kalman Algorithm [6] produces a realization from \hat{G} . We first consider the noise-free setting. The main idea of the Ho-Kalman Algorithm is to construct and factorize a Hankel matrix derived from the *G*. Specifically, we generate the Hankel matrix:

$$\mathcal{H} = \begin{bmatrix} CB & CAB & \dots & CA^{T_2}B \\ CAB & CA^2B & \dots & CA^{T_2+1}B \\ CA^2B & CA^3B & \dots & CA^{T_2+2}B \\ \vdots & \vdots & \vdots & \vdots \\ CA^{T_1-1}B & CA^{T_1}B & \dots & CA^{T_1+T_2-1}B \end{bmatrix},$$

where $T = T_1 + T_2 + 1$. We use $\mathcal{H}^-(\mathcal{H}^+)$ to denote the $pT_1 \times mT_2$ Hankel matrix created by deleting the last (first) block column of \mathcal{H} . We assume that

1) the system (1) is observable and controllable, and

2) $n = \operatorname{rank}(\mathcal{H}) \le \min\{T_1, T_2\}.$

Under these assumptions, \mathcal{H} and \mathcal{H}^- are of rank n. We note that \mathcal{H}^- can be factorized as

$$\mathcal{H}^{-} = \begin{bmatrix} C \\ CA \\ \vdots \\ CA^{T_1-1} \end{bmatrix} \begin{bmatrix} B & AB & \dots & A^{T_2-1}B \end{bmatrix} = OQ,$$

2179

where O, Q denote the observability matrix and controllability matrix respectively. We can also factorize \mathcal{H}^- by computing its truncated SVD, i.e., $\mathcal{H}^- = U\Sigma_n V^T = (U\Sigma_n^{\frac{1}{2}})(\Sigma_n^{\frac{1}{2}}V^T)$. Therefore, the factorization of \mathcal{H}^- establishes $O = U\Sigma_n^{\frac{1}{2}}, Q = \Sigma_n^{\frac{1}{2}}V^T$. And doing so, we obtain the system parameter C by taking the first p rows of $U\Sigma_n^{\frac{1}{2}}$ and the system parameter B by taking the first m columns of $\Sigma_n^{\frac{1}{2}}V$. Then A matrix can be obtained by $A = O^{\dagger}\mathcal{H}^+Q^{\dagger}$, where $(\cdot)^{\dagger}$ denotes the Moore-Penrose inverse. Note that the state-space realization problem is a non-convex. There are multiple solutions yielding the same system input/output behavior and Markov matrix G: If (A, B, C, D) is a statespace realization obtained from G, then for any non-singular matrix S, $(SAS^{-1}, SB, CS^{-1}, D)$ is also a valid realization.

D. System realization via noisy Markov parameter G

In the setting with noise, the same algorithm is applied to the estimated Markov matrix \hat{G} instead of the true matrix G. In this case the Ho-Kalman Algorithm will produce estimates $\hat{A}, \hat{B}, \hat{C}$ and \hat{D} . The explicit algorithm is described in Alg 1 (deterministic). It was shown in [7] that the robustness of the Ho-kalman Algorithm provides an estimation error bounded by $O(\frac{1}{N^{1/4}})$, where N is the number of trajectories:

$$\max\left\{ \|\hat{A} - S^{-1}AS\|, \|\hat{B} - S^{-1}B\|, \|\hat{C} - CS\| \right\}$$

$$\lesssim \sqrt{\|G - \hat{G}\|} = O(\frac{1}{N^{1/4}}).$$
(4)

This result can be improved to $O(\frac{1}{\sqrt{N}})$ [3, 11, 9].

Note that the computational complexity of the Ho-Kalman Algorithm in Alg. 1 (deterministic) is dominated by the cost of computing the SVD (Step 7), which is $O(pT_1 \times mT_2 \times n)$ when using the Krylov method (see e.g. [23, 24]). Therefore, we want to use a small T to reduce the computational cost. However, to satisfy the second assumption that n =rank(\mathcal{H}) $\leq \min\{T_1, T_2\}$, where $T_1 + T_2 + 1 = T$, the smallest T we can choose is 2n + 1 with $T_1 = T_2 = n$. In summary, the lowest achievable computational cost for SVD is $O(n^3)$. Such dependency on the system dimension is prohibitive for large-scale systems (e.g. systems with n = 100 as we show in Section V). Motivated by the drawbacks of the existing method, we aim to answer the following question:

• Is there a system realization method which can significantly reduce the computational complexity without sacrificing robustness guarantees?

The main result of this paper is to answer this question in the affirmative: We design a stochastic version of the Ho-Kalman algorithm that is computationally efficient and produces competitive robustness guarantees. Algorithm 1 Stochastic/Deterministic Ho-Kalman Algorithm

- 1: Input: Length T, Estimated Markov parameters \hat{G} , system order n, (T_1, T_2) satisfying $T_1 + T_2 + 1 = T$
- 2: **Outputs**: State space realization $\hat{A}, \hat{B}, \hat{C}, \hat{D}$
- 3: Generate a Hankel matrix $\hat{H} \in \mathbb{R}^{pT_1 \times m(T_1+1)}$ from \hat{G} 4: $\hat{H}^- = \hat{H}(:, 1: mT_2)$ $\triangleright \dim(\hat{H}^-) = pT_1 \times mT_2$ 5: $\hat{H}^+ = \hat{H}(:, m+1: m(T_2+1))$ 6: **if** Deterministic **then** 7: $L = H^{-}_{[n]}$ \triangleright truncated SVD via (2) $\hat{U}, \hat{\Sigma}, \hat{V} = \texttt{SVD}(\hat{L})$ 8: 9: else if Stochastic then $\hat{U}, \hat{\Sigma}, \hat{V} = \texttt{RSVD}(\hat{H}^-, n, l)^1$ $\triangleright \tilde{L} = \hat{U}\hat{\Sigma}\hat{V} \approx \hat{L}$ 10: 11: end if 12: $\hat{O} = \hat{U}\hat{\Sigma}^{1/2}$ $\triangleright \dim(\hat{O}) = pT_1 \times n$ $\triangleright \dim(\hat{Q}) = n \times mT_2$ 13: $\hat{Q} = \hat{\Sigma}^{1/2} \hat{V}^*$ 14: $\hat{C} = \hat{O}(1:p,:), \hat{B} = \hat{Q}(:,1:m)$ 15: $\hat{A} = \hat{O}^{\dagger}\hat{H}^{\dagger}\hat{Q}^{\dagger}, \hat{D} = \hat{G}(:, 1:m)$ 16: **Return** $\hat{A} \in \mathbb{R}^{n \times n}, \hat{B} \in \mathbb{R}^{n \times m}, \hat{C} \in \mathbb{R}^{p \times n}, \hat{D} \in \mathbb{R}^{p \times m}$

III. RANDOMIZED SINGULAR VALUE DECOMPOSITION

The numerical computation of a singular value decomposition can be implemented in many ways. The structure of the matrix to be decomposed will likely play a role in determining which is the most efficient algorithm. We do not attempt to review methods here as the literature is vast. In the system realization problem, the Ho-Kalman Algorithm computes the SVD of \mathcal{H}^- , a dense truncated block Hankel matrix. To the best of our knowledge there are no specialized algorithms for this purpose. As such, we assume we are dealing with a general dense low rank matrix.

The objective of the RSVD it to produce matrices U, Σ, V , such that for a given matrix $A \in \mathbb{C}^{m \times n}$ with $\operatorname{rank}(A) = r < \min\{m, n\}$, and tolerance $\epsilon > 0$, the bound²

$$\|A - U\Sigma V^*\| \le \epsilon$$

is satisfied where U and V have orthornormal columns and $\Sigma \in \mathbb{R}^{k \times k}$ is diagonal with k < r.

Following [15], the RSVD of a matrix A with target rank k is computed in two stages (full implementation details are provided in Algorithm 2):

- Find a matrix P ∈ ℝ^{m×k} with orthonormal columns such that the range of P captures as much of the range of A as possible. In other words, A ≈ PP*A.
- 2) Form the matrix $M = P^*A \in \mathbb{R}^{k \times n}$ and apply the standard numerical linear algebra technique to compute the SVD of P.

Step 1 is the *range finding* problem. This is where randomization enters picture. Let $\omega^{(i)}$ be a standard Gaussian vector, and compute $y^{(i)} = A\omega^{(i)}$. This can be viewed as a sample of Range(A). Repeating this process k times and

¹In the following analysis, we will use $\tilde{U}, \tilde{\Sigma}, \tilde{V}, \tilde{L}$ to denote the variables used in the stochastic Ho-Kalman Algorithm.

 $^{^{2}}$ The A matrix in this section is to demonstrate RSVD. It is not related to the dynamic system in (1).

concatenating samples into matrices we have $Y = A\Omega$, an orthonormal basis for Y can then be computed using standard techniques, we use an economy QR decomposition. Again, we concatenate basis vectors q_i into a matrix P. Because k is selected to be small, this process is computationally tractable. When $||A - PP^*A||$ is small, PP^*A is a good rank-k approximation of A. In step 2, standard deterministic routines are called to compute the SVD of M. These routines are considered tractable as the the matrix Mhas dimension $k \times m$ where k is ideally much less than r. From the SVD of M, the matrices U, Σ, V can be easily constructed (lines 7-8 of RSVD).

Algorithm 2 Randomized SVD: RSVD

1: Input: an $m \times n$ matrix A, a target rank k, an oversampling parameter l2: **Output:** Approximate SVD s.t. $A \approx USV^T$ 3: $\Omega = \operatorname{randn}(n, k+l)$ 4: $P = \operatorname{orth}(A\Omega)$ \triangleright approx. basis for Range(A) 5: $M = P^{\mathrm{T}}A$ $\triangleright \dim(M) = (k+l) \times n$ 6: $[U, S, V] = \operatorname{svd}(M)$ 7: U = PU8: U = U(:, 1:k), S = S(1:k, 1:k), V = V(:, 1:k)9: Return $U \in \mathbb{R}^{m \times k}, S \in \mathbb{R}^{k \times k}, V \in \mathbb{R}^{n \times k}$

In practice, if the target rank is selected to be k, then one should sample the range of $A \ k+l$ times where l is a small integer. In RSVD, Ω is chosen to be a standard Gaussian matrix. Surprisingly, the computational bottleneck of RSVD is the matrix-vector multiplication in computing $A\Omega$ in step 4. To reduce the computational cost of this step, we can choose other types of random matrices such as the subsampled random Fourier transform (SRFT) matrix which reduces the flop count from O(mn(k+l)) to $O(mn\log(k+l))$ without incurring much loss in accuracy (we extend our results to this setting in Appendix A-B). It should be further noted that the computation of $A\Omega$ is trivially parallelizable.

The orth function called on line 4 of RSVD computes an orthonormal basis for the range of its argument. This can be done in many ways, here we use an economy QR decomposition.

A. Power Scheme for slowly decaying spectra

When the input matrix A has a flat spectrum, RSVD tends to struggle to find a good approximate basis. To improve the accuracy a power iteration scheme is employed [23, p. 332]. Loosely, the power iteration are based on the observation that the singular vectors of A and $(AA^*)^q A$ are the same, while the singular values with magnitude less than one will rapidly shrink. In other words, it can reduce the effect of noise. More precisely, we apply RSVD to the matrix $W = (AA^*)^q A$, and we have $\sigma_i(W) =$ $\sigma_i(A)^{2q+1}, \quad j = 1, 2, \cdots$ which shows that for $\sigma_j < 1$, the power iteration will provide singular values that decay more rapidly while the singular vectors remain unchanged. This will provide a more accurate approximation, however it will require 2q+1 times as many matrix vector multiplies. The following theorem provides a bound on the accuracy of the approximation that RSVD provides.

Theorem 1. [15] Suppose that A is a real $m \times n$ matrix with singular values $\sigma_1 \geq \sigma_2 \geq \sigma_3 \geq \cdots$. Choose a target rank $k \geq 2$ and an oversampling parameter $l \geq 2$, where $k+l \leq \min\{m,n\}$. Then RSVD called on $W = (AA^*)^q A$ with target rank k, and oversampling parameter l produces an orthonormal approximate basis P which satisfies

$$\mathbb{E} \| (A - PP^*A) \| \le \left[1 + \sqrt{\frac{k}{l-1}} + \frac{e\sqrt{k+l}}{l} \cdot \sqrt{\min\{m,n\} - k} \right]^{\frac{1}{2q+1}} \sigma_{k+1}$$

where e is Euler's number and \mathbb{E} denotes expectation with respect to the random matrix Ω .

Note that σ_{k+1} is the theoretically optimal value in the deterministic setting. Thus the price of randomization is the given by the contents of the square brackets with exponent $\frac{1}{2q+1}$.

RSVD is implemented based on the assumption that the target rank k is know a priori. In practice, we do not know the true rank k in advance. Therefore, it is desirable to design an algorithm that can find a matrix P with as few orthonormal columns as possible such that $||(I - PP^*)A|| \leq \varepsilon$ where ε denotes a given tolerance. The work in [15], based on results from [25] describes an adaptive randomized ranger finder that iteratively samples until the desired tolerance is obtained.

IV. MAIN RESULTS

The *stochastic Ho-Kalman* algorithm we propose replaces the deterministic singular value decomposition and truncation (lines 7-8, Algorithm 1) with a single approximate randomized SVD (line 10, Algorithm 1) obtained using RSVD. Proofs of all results are deferred to the appendix. For the remainder of the paper, symbols with a tilde denote that they were obtained from the stochastic Ho-Kalman Algorithm, while symbols with a hat denote that they were obtained from the deterministic Ho-Kalman Algorithm. Finally, symbols with neither have been obtained from the ground truth Markov matrix G. From [7], we have the following perturbation bounds:

Lemma 2. [7] The matrices H, \hat{H} and L, \hat{L} satisfy the following perturbation bounds:

- $\max\left\{ \|H^+ \hat{H}^+\|, \|H^- \hat{H}^-\| \right\} \le \|H \hat{H}\| \le \sqrt{\min\{T_1, T_2 + 1\}} \|G \hat{G}\|.$ $\|L \hat{L}\| \le 2\|H^- \hat{H}^-\| \le 2\sqrt{\min\{T_1, T_2\}} \|G \hat{G}\|.$

We will now use Theorem 1 and Lemma 2 to provide average and deviation bounds on the performance of the stochastic Ho-Kalman algorithm.

Lemma 3. (Average perturbation bound) Denote l > 2to be the oversampling parameter used in RSVD. Run the Stochastic Ho-Kalman Algorithm with a standard Gaussian matrix $\Omega \in \mathbb{R}^{mT_2 \times (n+l)}$ in line 3 of RSVD, where $n+l \leq \min\{pT_1, mT_2\}$. Then L and L satisfy the following perturbation bound:

$$\mathbb{E}\|L - \tilde{L}\| \le 2C_2 \Big(2 + \sqrt{\frac{n}{l-1}} + \frac{e\sqrt{n+l}}{l}C_1\Big)\|G - \hat{G}\| \quad (5)$$

where

$$C_1 = \sqrt{\min\{pT_1, mT_2\} - n}, \quad C_2 = \sqrt{\min\{T_1, T_2\}}$$

Furthermore, if we exploit the power scheme with RSVD, then the right-hand side of (5) can be improved to

$$4C_2 \left(1 + \frac{1}{2}\sqrt{\frac{n}{l-1}} + \frac{e\sqrt{n+l}}{2l}C_1\right)^{1/(2q+1)} \|G - \hat{G}\|.$$
 (6)

Proof. See appendix.

From (5), we have that the perturbation bound is determined by the ratio between the target rank n and the oversampling parameter l. The error is large if l is small. In practice, it is sufficient to use l = 5 or l = 10. And there is rarely any advantage to select l > 10 [15]. In addition, from (6), we know that the bound will decrease if we increase the power parameter q. The effect of l in terms of running time and realization error is studied further in Section V where we observe that the stochastic Ho-Kalman algorithm is robust to the choice of l. In case the average perturbation as characterized by Lemma 3 doesn't feel like a helpful quantity, a deterministic error bound is also achievable:

Lemma 4. (Deviation bound) Let the assumptions of Lemma 3 hold. Assume $l \ge 4$ and let C_1 and C_2 be defined as in Lemma 3. Then we have

$$\|L - \tilde{L}\| \le 2C_2 \left(2 + 16\sqrt{1 + \frac{n}{l-1}} + \frac{8\sqrt{n+l}}{l+1}C_1\right)\|G - \hat{G}\|$$
(7)

with failure probability at most $3e^{-l}$. Moreover,

$$\|L - \tilde{L}\| \le C_2 \Big(2 + 6\sqrt{(n+l)l\log l} + 3\sqrt{n+l} C_1 \Big) \|G - \hat{G}\|$$
(8)

with failure probability at most $3l^{-l}$.

Proof. See appendix.
$$\Box$$

Remark 1. Another way to implement the stochastic Ho-Kalman Algorithm is to use a structured random matrix like subsampled random Fourier transform, or SRFT to compute the RSVD. In contrast with Gaussian matrices, SRFTs have faster matrix-vector multiply profile. As a result RSVD computation time decreases. We present the bounds for SRFT random matrices in the appendix.

We are now ready to show the robustness of stochastic Ho-Kalman algorithm (valid up to a unitary transformation).

Theorem 5. Suppose the system A, B, C, D is observable and controllable. Let O,Q be order-n controllability/observability matrices associated with G and O, Q be approximate order-n controllability/observability matrices (computed by RSVD) associated with \hat{G} . Suppose $\sigma_{\min}(L) >$ 0 and the following robustness condition is satisfied:

$$\mathbb{E}\|L - L\| \le \sigma_{\min}(L)/2.$$

Then, there exists a unitary matrix $S \in \mathbb{R}^{n \times n}$ such that,

$$\mathbb{E} \|C - \tilde{C}S\|_F \le \mathbb{E} \|O - \tilde{O}S\|_F \le \mathbb{E} \sqrt{5n} \|L - \tilde{L}\|,$$

$$\mathbb{E} \|B - S^* \tilde{B}\|_F \le \mathbb{E} \|Q - S^* \tilde{Q}\|_F \le \mathbb{E} \sqrt{5n} \|L - \tilde{L}\|.$$

and \tilde{A} , A satisfy

$$\mathbb{E}||A - S^* \tilde{A}S||_F$$

$$\leq C_3 \left(\sqrt{\frac{\mathbb{E}||L - \tilde{L}||}{\sigma_{\min}(L)}} (||H^+|| + ||H^+ - \hat{H}^+||) + ||H^+ - \hat{H}^+|| \right),$$
where $C_3 = \frac{14\sqrt{n}}{\sigma_{\min}(L)}.$
Proof. See appendix.

Proof. See appendix.

As discussed in [7], $||H^+ - \hat{H}^+||, \mathbb{E}||L - \hat{L}||$ are perturbation terms that can be bounded in terms of $\|G - \hat{G}\|$ via Lemma 2 and Lemma 3. Theorem 5 shows that the stochastic Ho-Kalman Algorithm has the same error bounds as its deterministic counterpart, which says the estimation errors for system matrix decrease as fast as $O(\frac{1}{N^{1/4}})$. Our analysis framework can be easily extended to achieve the optimal error bounds $O(\frac{1}{\sqrt{N}})$ mentioned in [21, 3, 11, 9].

V. NUMERICAL EXPERIMENTS

A. Stochastic versus deterministic Ho-Kalman Algorithm

We begin the comparison between the stochastic and deterministic Ho-Kalman Algorithm on six randomly generated systems described by (1). Each entry of the system matrix is generated through a uniform distribution over a range of integers as follows: matrix A with random integers from 1 to 5, and matrices B, C, D with random integers from -2 to 2. The A matrix is re-scaled to make it Schur stable³, i.e., $|\lambda_{\max}(A)| < 1$. The standard deviations of the process and measurement noises are $\sigma_w = 1$ and $\sigma_v = 0.5$. The length of trajectory T is given in the second column in Table I with T_1 chosen to be the smallest integer not less than T/2 and $T_2 = T - 1 - T_1$. The third column in Table I denotes the matrix dimension of \hat{H}^- when we run Algorithm 1.

We denote the true system as $\mathcal{G}(A, B, C, D)$ and the estimated system returned by the stochastic/deterministic Ho-Kalman algorithm as $\tilde{\mathcal{G}}(\tilde{A}, \tilde{B}, \tilde{C}, \tilde{D})/\hat{\mathcal{G}}(\hat{A}, \hat{B}, \hat{C}, \hat{D})$. We will use $\mathcal{G}, \mathcal{G}, \mathcal{G}$ at times to reduce notational clutter. The realization error of the algorithm is measured by the normalized \mathcal{H}_{∞} error: $\frac{\|\tilde{\mathcal{G}}-\mathcal{G}\|_{\mathcal{G}_{\infty}}}{\|\mathcal{G}\|_{\mathcal{H}_{\infty}}}$ and $\frac{\|\hat{\mathcal{G}}-\mathcal{G}\|_{\mathcal{H}_{\infty}}}{\|\mathcal{G}\|_{\mathcal{G}_{\infty}}}$. The running time and the realization error of the deterministic and stochastic algorithms⁴ are reported in Table I where the results for the

³There is no requirement that the systems we work with be stable. However, we are using an \mathcal{H}_{∞} -norm metric to judge the approximation error, so such an assumption makes things more straight forward.

⁴We used the publicly available python package sklearn.utils.extmath. randomized_svd to compute the RSVD.

Eg	(n,m,p,T)	$\dim(\hat{H}^-)$	Running Time [s]		Realization Error	
			deterministic	stochastic	deterministic	stochastic
1	(30,20,10,90)	450×880	0.1079	0.0156	7.64e-04	7.70e-04
2	(40,30,20,100)	2000×2970	5.7456	0.0897	6.67e-04	1.19e-03
3	(60,50,40,360)	7200×8950	227.0116	0.9323	8.27e-04	1.75e-03
4	(100,80,50,500)	12500×19920	922.8428	4.4581	6.53e-04	1.66e-03
5	(120,110,90,600)	27000×32890	Inf	17.6603	N/A	1.96e-03
6	(200,150,100,600)	30000×44850	Inf	52.1762	N/A	1.45e-03

Table I: Comparison between the stochastic and deterministic Ho-Kalman Algorithm. The running time is in seconds. The approximate SVD is computed using RSVD with oversampling parameter l = 10. To benchmark the algorithm performance, a naive implementation of RSVD is used; we *do not* use power iterations and *do not* make use of parallelization. Inf and N/A indicates that the deterministic algorithm fail to realize the system.

stochastic algorithm are average over 10 independent trials. All experiments are done on a 2.6 GHz Intel Core i7 CPU.

The reported running time in the stochastic setting is highly conservative: we did not parallelize the sampling (i.e., constructing $A\Omega$ in line 4 of Algorithm 2). Furthermore, as noted earlier, standard Gaussian matrices are theoretically "nice" to work with but structured random matrices will offer superior running times. We observe that the stochastic Ho-Kalman algorithm consistently leads to a dramatic speed-up over the deterministic algorithm. The larger the system dimension is, the larger the run time gap is. Observe that the deterministic Ho-Kalman Algorithm fails to provide a result in the 5th and 6th examples where the system state dimensions are above 100. Meanwhile, the stochastic algorithm runs successfully and takes a fraction of the time the deterministic algorithm took to solve a 60 state realization problem.

B. Oversampling effects

To illustrate the influence of oversampling (parameter *l* in RSVD), we run the stochastic Ho-Kalman Algorithm on the 4th example (n = 100, m = 80, p = 50) in Table I and vary the oversampling parameter l from 1 to 10. In this experiment we use a power iteration parameter of q = 1. Running times (averaged over 10 runs) and realization errors (averaged over 10 runs) are shown via a boxplot in Figure 1a and graph in Figure 1b. In the boxplot, the central red mark indicates the median, and the bottom and top edges of the box indicate the 25th and 75th percentiles, respectively. Outliers are denoted by "+". We observe that in Fig 1b, the realization error tends to be larger when a small oversampling number is used, although the change is slight. The observed behavior is consistent with the theoretical analysis of Lemma 3. We can also observe from Fig 1a that the computational time is insensitive to the oversampling parameters, as such taking larger values of l = 10 is advantageous.

C. Power iteration effect

Based on the results of the previous subsection, we fix the oversampling parameter in RSVD as l = 10, and sweep q from 1 to 4. The results are shown in Figures 1c and 1d. We observe that the realization error decreases as the power parameter q increases as indicated in Eq (6). In contrast to the oversampling parameter l, the runtime demonstrably increases with q at an empirically linear rate. This trend is expected and analyzed in [15].

The power iteration method is most effective for problems where the spectrum of the matrix being approximated decays slowly. In the noise free setting, $\operatorname{rank}(\mathcal{H}^-) = n$, where n = 100 in this example. In contrast the dimensions of $\mathcal{H}^$ are 12500×19920 . When noise is introduced, \mathcal{H}^- becomes full rank and the spectral decay depends on σ_w and σ_v . For the values chosen, these results show that spectral decay appears to be sharp enough that the power iterations do not offer *significant* improvement in accuracy.

VI. CONCLUSION

We have introduced a scalable algorithm for system realization based on introducing randomized numerical linear algebra techniques into the Ho-Kalman algorithm. Theoretically it has been shown that our algorithm provides nonasymptotic performance guarantees that are competitive with deterministic approaches. Besides, without any algorithm optimization, we have shown that the stochastic algorithm easily handles problem instances of a size significantly beyond what classical deterministic algorithms can handle.

REFERENCES

- Y. Abbasi-Yadkori, D. Pál, and C. Szepesvári, "Online least squares estimation with self-normalized processes: An application to bandit problems," *arXiv preprint arXiv:1102.2670*, 2011.
- [2] M. Simchowitz, H. Mania, S. Tu, M. I. Jordan, and B. Recht, "Learning without mixing: Towards a sharp analysis of linear system identification," in *Conference On Learning Theory*. PMLR, 2018, pp. 439–473.
- [3] T. Sarkar, A. Rakhlin, and M. A. Dahleh, "Finite-time system identification for partially observed lti systems of unknown order," arXiv preprint arXiv:1902.01848, 2019.
- [4] M. K. S. Faradonbeh, A. Tewari, and G. Michailidis, "Finite time identification in unstable linear systems," *Automatica*, vol. 96, pp. 342–353, 2018.
- [5] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu, "On the sample complexity of the linear quadratic regulator," *Foundations of Computational Mathematics*, vol. 20, no. 4, pp. 633–679, 2020.
- [6] B. Ho and R. E. Kálmán, "Effective construction of linear state-variable models from input/output functions," *at-Automatisierungstechnik*, vol. 14, no. 1-12, pp. 545–548, 1966.



(a) Running time of stochastic Ho-Kalman Algorithm using RSVD with oversampling parameter l.



(c) Running time of stochastic Ho-Kalman Algorithm with varying power parameter q. The oversampling parameter l is 10.

Figure 1: Oversampling and power iteration effect



(b) Realization error of stochastic Ho-Kalman Algorithm using RSVD with oversampling parameter l.



(d) Realization error of the stochastic Ho-Kalman Algorithm with varying power parameter q. The oversampling parameter l is 10.

- [7] S. Oymak and N. Ozay, "Non-asymptotic identification of lti systems from a single trajectory," in 2019 American control conference (ACC). IEEE, 2019, pp. 5655–5661.
- [8] Y. Zheng and N. Li, "Non-asymptotic identification of linear dynamical systems using multiple trajectories," *IEEE Control Systems Letters*, vol. 5, no. 5, pp. 1693–1698, 2020.
- [9] H. Lee, "Improved rates for identification of partially observed linear dynamical systems," *arXiv preprint arXiv:2011.10006*, 2020.
- [10] Y. Sun, S. Oymak, and M. Fazel, "Finite sample system identification: Optimal rates and the role of regularization," in *Learning for Dynamics and Control.* PMLR, 2020, pp. 16–25.
- [11] A. Tsiamis and G. J. Pappas, "Finite sample analysis of stochastic system identification," in 2019 IEEE 58th Conference on Decision and Control (CDC). IEEE, 2019, pp. 3648–3654.
- [12] M. Sznaier, "Control oriented learning in the era of big data," *IEEE Control Systems Letters*, vol. 5, no. 6, pp. 1855–1867, 2020.
- [13] H. Wang and J. Anderson, "Learning linear models using distributed iterative hessian sketching," arXiv preprint arXiv:2112.04101, 2021.
- [14] S. Voronin and P.-G. Martinsson, "Rsvdpack: An implementation of randomized algorithms for computing the singular value, interpolative, and cur decompositions of matrices on multi-core and gpu architectures," *arXiv preprint arXiv:1502.05366*, 2015.

- [15] N. Halko, P.-G. Martinsson, and J. A. Tropp, "Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions," *SIAM review*, vol. 53, no. 2, pp. 217–288, 2011.
- [16] C. Musco and C. Musco, "Randomized block krylov methods for stronger and faster approximate singular value decomposition," Advances in Neural Information Processing Systems, vol. 28, pp. 1396–1404, 2015.
- [17] X. Feng, W. Yu, and Y. Li, "Faster matrix completion using randomized svd," in 2018 IEEE 30th International Conference on Tools with Artificial Intelligence (ICTAI). IEEE, 2018, pp. 608–615.
- [18] S. Boyd, S. P. Boyd, and L. Vandenberghe, *Convex optimiza*tion. Cambridge university press, 2004.
- [19] C. Eckart and G. Young, "The approximation of one matrix by another of lower rank," *Psychometrika*, vol. 1, no. 3, pp. 211–218, 1936.
- [20] M. Simchowitz, R. Boczar, and B. Recht, "Learning linear dynamical systems with semi-parametric least squares," in *Conference on Learning Theory*. PMLR, 2019, pp. 2714– 2802.
- [21] T. Sarkar and A. Rakhlin, "Near optimal finite time identification of arbitrary linear dynamical systems," in *International Conference on Machine Learning*. PMLR, 2019, pp. 5610– 5618.
- [22] S. Tu, R. Boczar, A. Packard, and B. Recht, "Non-asymptotic analysis of robust control from coarse-grained identification," *arXiv preprint arXiv*:1707.04791, 2017.

- [23] G. H. Golub and C. F. Van Loan, "Matrix computations. johns hopkins studies in the mathematical sciences," 1996.
- [24] L. N. Trefethen and D. Bau III, Numerical linear algebra. Siam, 1997, vol. 50.
- [25] F. Woolfe, E. Liberty, V. Rokhlin, and M. Tygert, "A fast randomized algorithm for the approximation of matrices," *Applied and Computational Harmonic Analysis*, vol. 25, no. 3, pp. 335–366, 2008.
- [26] —, "A fast randomized algorithm for the approximation of matrices," *Applied and Computational Harmonic Analysis*, vol. 25, no. 3, pp. 335–366, 2008.
- [27] S. Tu, R. Boczar, M. Simchowitz, M. Soltanolkotabi, and B. Recht, "Low-rank solutions of linear matrix equations via procrustes flow," in *International Conference on Machine Learning*. PMLR, 2016, pp. 964–973.

APPENDIX A

A. Proof of Lemma 3 and 4

~

~

Proof of Lemma 3. To prove (5), we first use the triangle inequality to get the following bound:

$$\mathbb{E}\|L - \tilde{L}\| \le \|H^{-} - \hat{H}^{-}\| + \mathbb{E}\|\hat{H}^{-} - \tilde{L}\|,$$
(9)

where $L = H^-$ and H^- is of rank n. Then we bound $\mathbb{E} \|\hat{H}^- - \tilde{L}\|$ by applying Theorem 1 to \hat{H}^- with q = 0, giving:

$$\begin{aligned} & \mathbb{E} \| H^{-} - L \| \leq \\ & \left(1 + \sqrt{\frac{n}{l-1}} + \frac{e\sqrt{n+l}}{l} \sqrt{\min\{pT_{1}, mT_{2}\} - n} \right) \| \hat{H}^{-} - \hat{H}_{[n]}^{-} \| \\ & \leq \left(1 + \sqrt{\frac{n}{l-1}} + \frac{e\sqrt{n+l}}{l} \sqrt{\min\{pT_{1}, mT_{2}\} - n} \right) \| \hat{H}^{-} - H \| . \end{aligned}$$

The first inequality follows from $\tilde{L} = PP^*\hat{H}^-$. The second inequality is due to the fact that $\hat{H}_{[k]}^-$ is the best rank-*k* approximation of \hat{H}^- . Plugging the inequality above into (9) and applying Lemma 2, we obtain the inequality (5).

Applying the bound in Theorem 1 to \hat{H}^- with a fixed positive integer q > 0 gives us

$$\mathbb{E} \| H^{-} - L \| \tag{10}$$

$$\leq \left(1 + \sqrt{\frac{n}{l-1}} + \frac{e\sqrt{n+l}}{l} \sqrt{\min\{pT_1, mT_2\}} - n \right)^{1/(2q+1)} \\
\times \| \hat{H}^{-} - H \| + 1^{1/(2q+1)} \| \hat{H}^{-} - H \| \stackrel{(a)}{\leq} \\
2 \left(1 + \frac{1}{2} \sqrt{\frac{n}{l-1}} + \frac{1}{2} \frac{e\sqrt{n+l}}{l} \sqrt{\min\{pT_1, mT_2\}} - n \right)^{1/(2q+1)} \\
\times \| \hat{H}^{-} - H \| \tag{11}$$

Inequality (a) holds because $x^{1/(2q+1)}$ is concave in x. Applying Lemma 2, we prove the inequality (6).

Proof of Lemma 4. Follow the same steps as Lemma 3 and use the bound given in Corollary 10.9 [15]. \Box

B. Perturbation bounds for Stochastic Ho-Kalman Algorithm with SRFT Test Matrices

The subsampled random Fourier transform (SRFT) is an $n \times \ell$ matrix of the form $\Omega = \sqrt{\frac{n}{\ell}}DFR$ where D is an $n \times n$ diagonal matrix whose entries are independent random variables uniformly distributed on the complex unit circle,

F is the $n \times n$ unitary discrete Fourier transform (DFT), whose entries take the values

$$f_{pq} = n^{-1/2} e^{-2\pi i (p-1)(q-1)/n}$$
 for $p, q = 1, 2, \dots, n$,

and R is an $n \times \ell$ matrix that samples ℓ coordinates from n uniformly at random [15]. When Ω is a SRFT matrix, we can calculate the matrix multiplication $Y = A\Omega$ using $O(mn \log(\ell))$ flops by applying a subsampled FFT [26].

Lemma 6. (Deviation bound) Denote $l \ge 2$ to be the oversampling parameter used in RSVD algorithm. Run the Stochastic Ho-Kalman Algorithm with a SRFT matrix $\Omega \in \mathcal{R}^{mT_2 \times (n+l)}$ in computing the RSVD step, where $4[\sqrt{n} + \sqrt{8\log(nmT_2)}]^2 \log(n) \le l + n \le mT_2$. Then L, \tilde{L} satisfy the following perturbation bound:

$$||L - \tilde{L}|| \le (1 + \sqrt{1 + \frac{7mT_2}{l+n}}) \times 2\sqrt{\min\{T_1, T_2\}} ||G - \hat{G}||$$

with failure probability at most $O(n^{-1})$.

Proof. We follow the same steps in proving Lemma 3 and use the bound in Theorem 11.2 of [15] to finish the proof. \Box

C. Proof of Theorem 5

To prove Theorem 5, we require two auxiliary lemmas.

Lemma 7. Suppose $\sigma_{\min}(L) \geq 2\mathbb{E}||L - \tilde{L}||$ where $\sigma_{\min}(L)$ is the smallest nonzero singular value (i.e. *n*-th largest singular value) of L. Let rank n matrices L, \tilde{L} have the singular value decomposition $U\Sigma V^*$ and $\tilde{U}\tilde{\Sigma}\tilde{V}^*$. There exists an $n \times n$ unitary matrix S so that

$$\mathbb{E} \| U\Sigma^{1/2} - \tilde{U}\tilde{\Sigma}^{1/2}S \|_F^2 + \mathbb{E} \| V\Sigma^{1/2} - \tilde{V}\tilde{\Sigma}^{1/2}S \|_F^2
\leq 5n\mathbb{E} \| L - \tilde{L} \|.$$
(12)

Proof. : Direct application of Theorem 5.14 of [27] guarantees the existence of a unitary S such that

LHS =
$$\mathbb{E} \| U \Sigma^{1/2} - \tilde{U} \tilde{\Sigma}^{1/2} S \|_F^2 + \mathbb{E} \| V \Sigma^{1/2} - \tilde{V} \tilde{\Sigma}^{1/2} S \|_F^2$$

 $\leq \frac{2}{\sqrt{2} - 1} \frac{\mathbb{E} \| L - \tilde{L} \|_F^2}{\sigma_{\min}(L)},$
(13)

where LHS refers to the left hand side of (12). To proceed, using $\mathbb{E} \operatorname{rank}(L - \tilde{L}) \leq 2n$ and by assumption $\sigma_{\min}(L) \geq 2\mathbb{E}||L - \tilde{L}|| \geq \sqrt{2/n}\mathbb{E}||L - \tilde{L}||_F$, we find LHS $\leq \frac{\sqrt{2n}}{\sqrt{2-1}}\mathbb{E}||L - \tilde{L}||_F \leq \frac{2n}{\sqrt{2-1}}\mathbb{E}||L - \tilde{L}|| \leq 5n\mathbb{E}||L - \tilde{L}||$. \Box

Lemma 8. Suppose $\sigma_{\min}(L) \geq 2 \geq 2\mathbb{E}||L - \tilde{L}||$. Then, $\mathbb{E}||\tilde{L}|| \leq 2||L||$ and $\sigma_{\min}(\mathbb{E}\tilde{L}) \geq \sigma_{\min}(L)/2$.

Proof. See Lemma 2.2 in [7].
$$\Box$$

Using these, we will prove the robustness of the stochastic Ho-Kalman Algorithm, which is stated in Theorem 5. The robustness will be up to a unitary transformation similar to Lemma 7.

Proof. The proof of Theorem 5 is obtained by following the proof of Theorem 5.3 in [7] and substitute \tilde{L} for \hat{L} .