

No-Regret Prediction in Marginally Stable Systems

Udaya Ghai

Princeton University

UGHAI@CS.PRINCETON.EDU

Holden Lee

Duke University

HOLEE@MATH.DUKE.EDU

Karan Singh

Princeton University

KARANS@PRINCETON.EDU

Cyril Zhang

Princeton University

CYRIL.ZHANG@PRINCETON.EDU

Yi Zhang

Princeton University

Y.ZHANG@PRINCETON.EDU

Editors: Jacob Abernethy and Shivani Agarwal

Abstract

We consider the problem of online prediction in a marginally stable linear dynamical system subject to bounded adversarial or (non-isotropic) stochastic perturbations. This poses two challenges. Firstly, the system is in general unidentifiable, so recent and classical results on parameter recovery do not apply. Secondly, because we allow the system to be marginally stable, the state can grow polynomially with time; this causes standard regret bounds in online convex optimization to be vacuous. In spite of these challenges, we show that the online least-squares algorithm achieves sublinear regret (improvable to polylogarithmic in the stochastic setting), with polynomial dependence on the system’s parameters. This requires a refined regret analysis, including a structural lemma showing the current state of the system to be a small linear combination of past states, even if the state grows polynomially. By applying our techniques to learning an autoregressive filter, we also achieve logarithmic regret in the partially observed setting under Gaussian noise, with polynomial dependence on the memory of the associated Kalman filter.

Keywords: Dynamical systems, marginally stable systems, online least squares.

1. Introduction

We consider the problem of sequential state prediction in a linear time-invariant dynamical system, subject to perturbations:

$$x_t = Ax_{t-1} + Bu_{t-1} + \xi_t.$$

This is a central object of study in control theory and time-series analysis, dating back to the foundational work of Kalman (Kalman, 1960), and has recently received considerable attention from the machine learning community. In a typical learning setting, the system parameters A, B are unknown, and only the past states x_t and exogenous inputs u_t are observed. Sometimes, another layer of difficulty is imposed: the *latent state* x_t can only be observed noisily or through a low-rank transformation. These models serve as an abstraction for learning from correlated data in stateful environments, and have helped to understand empirical successes in reinforcement learning and of recurrent neural networks.

Many recent works (Simchowitz et al., 2018; Sarkar et al., 2019) are concerned with finite-sample *system identification*, in which the matrices A, B can be recovered due to structure in the perturbations or inputs. These results rely on matrix concentration inequalities and careful error propagation applied to classic primitives in linear system identification, and have settled some important statistical questions about these methods. However, as in the classical control theory literature, these theorems require unrealistic assumptions of i.i.d. isotropic random perturbations and recoverability of the system, and often require the user to select the “exploration” inputs u_t (Simchowitz et al., 2019; Dean et al., 2017). Furthermore, under model misspecification, the guarantees of parameter identification pipelines break down.

Another line of work seeks to obtain more flexible guarantees via the online learning framework (Hazan et al., 2016). Here, the goal of parameter recovery is replaced with regret minimization, the excess prediction loss compared to the best-fit system parameters in hindsight (Kozdoba et al., 2019; Hazan et al., 2017, 2018). This approach gives rise to algorithms which adapt to adversarially perturbed data and model misspecification, and can be extended beyond prediction to obtain new methods for robust control (Arora et al., 2018). However, these algorithms can diverge significantly from the classical parameter identification pipeline. In particular, they can be *improper*, in that they may use an intentionally misspecified (e.g. overparameterized) model. Thus, these algorithms can be incompatible with parameter recovery and downstream methods.

In this work, we show that the same algorithm used for parameter identification (online least squares) has a no-regret guarantee, even in the challenging setting of prediction under marginal stability and adversarial perturbations, where recovery is impossible. More precisely, in this setting, where the state is allowed to grow polynomially with time, we show that the regret of this algorithm is sublinear, with a polynomial dependence on the system’s parameters. This does not follow from the usual analysis of online least squares: the magnitude of loss functions (and associated gradient bounds) can scale polynomially with time, causing standard regret bounds to become vacuous. Instead, we conduct a refined regret analysis, including a structural *volume doubling* lemma showing x_t to be a small linear combination of past states.

By replacing the worst-case structural lemma with a stronger martingale analysis, we also show a polylogarithmic regret bound for least squares in the stochastic setting. Again, this analysis does not go through parameter convergence, and thus applies in the setting of unidentifiable systems and non-isotropic noise. The same techniques allow us to prove a logarithmic regret bound in the partially observed setting under Gaussian noise, with polynomial dependence on the memory of the associated Kalman filter.

Paper structure. In Section 2, we formally introduce the problem and the natural *online least squares* algorithm. In Section 3, we give an overview of related work. In Section 4, we state our main results. In Section 5, we sketch the proofs. In Section B, we show that a structural condition on a time series gives a regret bound for online least squares. We then prove our main theorems for fully observed LDS in the adversarial setting (Section C), fully observed LDS in the stochastic setting (Section D), and for partially observed LDS in the stochastic setting (Section E), through establishing this structural condition. Section G provides an alternative approach.

2. Problem setting and algorithm

The problem of online state prediction for a LDS falls within the framework of online least squares. We first introduce the general problem of online least squares (Section 2.1), and then specialize to

the prediction problem for a fully observed (Section 2.2) or partially observed (Section 2.3) LDS. Because the observations come from a sequential process, they have extra structure that we will leverage to obtain better guarantees than for black-box online least squares. In Section 2.4, we describe the challenges associated with marginally stable systems.

2.1. The online least squares problem

In the problem of online least squares, at each time t we are given $x_t \in \mathbb{R}^m$, and asked to predict $y_t \in \mathbb{R}^n$. We choose a matrix $A_t \in \mathbb{R}^{n \times m}$ and predict $\hat{y}_t = A_t x_t \in \mathbb{R}^n$. The desired output y_t is then revealed and we suffer the squared loss $\|\hat{y}_t - y_t\|_2^2$. A natural goal in this setting is to predict as well as if we had known the best matrix in hindsight; hence, the performance metric is given by the *regret* with respect to A , defined by

$$R_T(A) = \sum_{t=1}^T \|A_t x_t - y_t\|_2^2 - \sum_{t=1}^T \|A x_t - y_t\|_2^2.$$

Define the regret with respect to a given set $\mathcal{K} \subseteq \mathbb{R}^{n \times m}$ by $R_T = \sup_{A \in \mathcal{K}} R_T(A)$. In general, we would like to achieve R_T that is sublinear in T , or equivalently, average regret $\frac{R_T}{T}$ that converges to 0. In some cases, we can do better, and achieve regret R_T that is polylogarithmic in T .

A natural algorithm for online least squares is to choose A_t that minimizes the total squared prediction error for all the pairs (x_s, y_s) , $1 \leq s \leq t-1$ seen so far, plus a regularization term:

Algorithm 1 Online Least Squares Regression

Input: Regularization parameter μ .

for $t = 1$ to T **do**

Let $A_t = \operatorname{argmin}_{A \in \mathbb{R}^{n \times m}} (\mu \|A\|_F^2 + \sum_{s=0}^{t-1} \|A x_s - y_s\|_2^2)$.

Predict $\hat{y}_t := A_t x_t$ and observe cost $\|\hat{y}_t - y_t\|_2^2$.

end for

We state the standard regret bound for Algorithm 1. Note that [Cesa-Bianchi and Lugosi \(2006\)](#) show the $n = 1$ case, but the proof goes through in the multi-dimensional case (Appendix F).

Theorem 1 (OLS regret bound; Thm. 11.7, [Cesa-Bianchi and Lugosi 2006](#)) *In the online least squares setting, suppose that $\|x_t\| \leq M$ for all t . Then, Algorithm 1 incurs regret*

$$R_T(A) \leq \mu \|A\|_F^2 + \max_{0 \leq t \leq T-1} \|y_t - A_t x_t\|_2^2 m \ln \left(1 + \frac{TM^2}{m} \right).$$

Thus, if there is a uniform bound on the prediction errors $\|y_t - A_t x_t\|$, online least squares achieves logarithmic regret. This follows immediately in the usual OLS setting, where x_t, y_t , are bounded and A_t is restricted to a bounded set. However, in the case where they can grow with time, as in marginally stable systems, a more sophisticated analysis will be necessary to get sublinear regret.

2.2. Prediction in fully-observed linear dynamical systems

A special case of online least squares is state prediction in a time-invariant linear dynamical system (LDS), defined as follows. Given an initial state $x_0 \in \mathbb{R}^d$, matrices $A \in \mathbb{R}^{d \times d}$ and $B \in \mathbb{R}^{d \times m}$,

inputs $u_0, \dots, u_{T-1} \in \mathbb{R}^m$ and a sequence of perturbations $\xi_1, \dots, \xi_T \in \mathbb{R}^d$, the LDS produces a time series of states $x_1, \dots, x_T \in \mathbb{R}^d$ according to the following dynamics:

$$x_t = Ax_{t-1} + Bu_{t-1} + \xi_t, \quad 1 \leq t \leq T. \quad (1)$$

This setting generalizes the *linear Gaussian model* from control theory and time-series analysis, in which each ξ_t is drawn i.i.d. from a Gaussian distribution. Aside from modeling disturbances, ξ_t can also represent model uncertainty or misspecification.

In the prediction problem for LDS, we are asked to predict x_{t+1} as a linear function of the current x_t and the input u_t . We can treat this as an online least squares problem, by casting (x_t, u_t) as the input at time t , and x_{t+1} as the desired output. At each step, the learner produces A_t and B_t and predicts $\hat{x}_{t+1} = A_t x_t + B_t u_t$. Thus, we can adapt Algorithm 1 to this setting with the substitution $x_t \leftarrow \begin{pmatrix} x_t \\ u_t \end{pmatrix}$, $y_t \leftarrow x_{t+1}$, $A \leftarrow (A, B)$, $(n, m) \leftarrow (d, d + m)$, and obtain Algorithm 2. Translated to this setting, the goal of regret minimization becomes that of predicting as accurately as if one had known the system's underlying matrices A and B .

Algorithm 2 Online Least Squares Regression (LDS setting)

- 1: **Input:** Regularization parameter μ .
- 2: **for** $t = 0, \dots, T - 1$ **do**
- 3: Estimate dynamics as

$$(A_t, B_t) := \underset{(A, B)}{\operatorname{argmin}} \left(\mu \|(A, B)\|_F^2 + \sum_{s=0}^{t-1} \|Ax_s + Bu_s - x_{s+1}\|^2 \right).$$

- 4: Predict state: $\hat{x}_{t+1} := A_t x_t + B_t u_t$ and suffer loss $\|\hat{x}_{t+1} - x_{t+1}\|^2$.
 - 5: **end for**
-

Note that in the stochastic setting, when the covariance of the noise is lower-bounded in each direction, OLS gives a consistent estimator for A and B ; convergence rates for recovery are analyzed in Sarkar and Rakhlin (2019); Simchowitz et al. (2018). However, in the adversarial setting, recovery of A is an ill-posed problem. The perturbations can be biased or rank-deficient, causing the recovery problem to be underdetermined in general, and the optimal A may change as time.

2.3. Prediction in partially-observed linear dynamical systems

A partially-observed linear dynamical system is defined by

$$x_t = Ax_{t-1} + Bu_{t-1} + \xi_t \quad (2)$$

$$y_t = Ch_t + \eta_t, \quad (3)$$

where $u_t \in \mathbb{R}^m$ are inputs, $x_t \in \mathbb{R}^d$ are hidden states, $y_t \in \mathbb{R}^n$ are observations, $A \in \mathbb{R}^{d \times d}$, $B \in \mathbb{R}^{d \times m}$ and $C \in \mathbb{R}^{n \times d}$ are matrices, and $\xi_t \in \mathbb{R}^d$ and $\eta_t \in \mathbb{R}^n$ are perturbations. We consider the stochastic setting, so that ξ_t and η_t are independent zero-mean noise terms. Crucially, only the y_t , and not the x_t , are observed.

For prediction in this setting, we use Algorithm 3, regressing with the previous ℓ observations and inputs, so we slightly modify the definition of the regret in (118) to start accruing from $t = \ell + 1$:

$$R_T(A, B, C) = \sum_{t=\ell+1}^T \|\hat{y}_t - y_t\|^2 - \sum_{t=\ell+1}^T \|\hat{y}_{\text{KF},t} - y_t\|^2,$$

where $\hat{y}_{\text{KF},t}$ is the prediction of the *steady-state Kalman filter* for the system (A, B, C) ; see Appendix E for a review and formal definitions. Note that we will learn the system in an *improper* manner: that is, we will predict \hat{y}_t using a general autoregressive filter, rather than the Kalman filter of some system.

Algorithm 3 Online Least Squares Autoregression for LDS

- 1: **Input:** Regularization parameter μ , rollout length ℓ .
- 2: **for** $t = 0$ to $T - 1$ **do**
- 3: Estimate the autoregressive filter:

$$(F_t, G_t) := \underset{F \in \mathbb{R}^{n \times \ell m}, G \in \mathbb{R}^{n \times \ell n}}{\operatorname{argmin}} \left(\mu \|(F, G)\|_F^2 + \sum_{s=\ell-1}^{t-1} \|Fu_{s:s-\ell+1} + Gy_{s:s-\ell+1} - y_{s+1}\|^2 \right).$$

- 4: Predict state: $\hat{y}_{t+1} := F_t u_{t:t-\ell+1} + G_t y_{t:t-\ell+1}$ and suffer cost $\|\hat{y}_{t+1} - y_{t+1}\|^2$.
 - 5: **end for**
-

2.4. Marginally stable systems

In this work, we are interested in prediction in marginally stable systems. In both the fully-observed and partially-observed cases, the *spectral radius* $\rho(A)$ of the system is defined to be the magnitude of the largest eigenvalue of the transition matrix A , as in Equations 1 and 2. An LDS is *marginally stable* if $\rho(A) = 1$.

As opposed to *strictly* stable systems (ones for which $\rho(A) < 1$), these systems model phenomena where the state does not reset itself over time, often representing physical systems which experience little or no dissipation. As discussed in Section 3, their capacity to represent long-term dependences presents algorithmic and statistical challenges. An *inverse spectral gap* factor $\frac{1}{1-\rho(A)}$ appears in the computational and statistical guarantees for many learning algorithms in these settings (see, e.g. [Hardt et al. \(2016\)](#)), as a finite-impulse truncation length or mixing time, rendering those results inapplicable.

Among marginally stable systems, the hardest cases are those with large Jordan blocks corresponding to large eigenvalues. Defining the *Jordan matrix*

$$J_{\lambda,r} := \begin{pmatrix} \lambda & 1 & 0 & \cdots & 0 \\ 0 & \lambda & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \lambda & 1 \\ 0 & 0 & 0 & 0 & \lambda \end{pmatrix} \in \mathbb{R}^{r \times r},$$

we see that for marginally stable systems, $\|J_{\lambda,r}^t\|_2 = \Omega(\lambda^{t-r+1}t^{r-1})$ can grow polynomially in t . These occur naturally in physical systems as discrete-time integrators of r -th degree ordinary differential equations. The primary challenge we overcome in this work is to show the sublinear regret of least squares, even when the state grows polynomially. As is also the case in our work, recent advances in parameter recovery of marginally stable systems (Simchowitz et al., 2018; Sarkar and Rakhlin, 2019) exhibit exponential dependences on the largest Jordan block order r .

3. Related work

Linear dynamical systems have been studied in a number of disciplines, including control theory (Ghahramani and Hinton, 1996; Kalman, 1963), astronomy (Chiappori et al., 1992), econometrics (Hendry and Doornik, 1994), biology (Saunders, 1994), and chemical kinetics. They capture many popular models in statistics and machine learning (Ghahramani and Hinton, 1996). We first describe the results on parameter estimation and prediction in fully observed systems, and then describe results more broadly applicable to partially observed systems. Unless noted otherwise, the results hold under the assumption that the noise is i.i.d.; some results also require that it be Gaussian.

Fully-observed LDS. Dean et al. (2017) show that when given independent rollouts of a LDS, the least-squares estimator of the parameters is sample-efficient. Using this, they obtain sub-optimality bounds for control. Simchowitz et al. (2018) consider the more challenging case when only a single trajectory is given, and show that the least-squares estimator is still efficient, despite correlations across timesteps. Their results hold for marginally stable systems. Improving over Simchowitz et al. (2018) and Faradonbeh et al. (2018), Sarkar and Rakhlin (2019) offer bounds applicable even to explosive systems, with the restriction that explosive eigenvalues have unit geometric multiplicity. In order to obtain results for parameter recovery, all these results assume that the covariance of the noise is lower bounded. Because we are concerned with prediction, this requirement will not be necessary for our results.

Partially-observed LDS. For a system with known parameters, the celebrated Kalman filter (Kalman, 1960) provides an analytic solution for the posterior distribution of the latent states and future observations given a series of observations. When the underlying parameters are unknown, the Expectation Maximization (EM) algorithm can be used to learn them (Ghahramani and Hinton, 1996). However, due to nonconvexity of the problem, EM is only guaranteed to converge to local optima. In the absence of process noise ξ_t , Hardt et al. (2016) show that gradient descent on the maximum likelihood objective converges to the global minimum; however, they require the roots of the associated characteristic polynomial to be well-separated and the system to be strictly stable.

Subspace identification methods circumvent the nonconvexity of maximum likelihood estimation. For strictly stable systems, Oymak and Ozay (2018) demonstrate that the Ho-Kalman algorithm learns the Markov parameters of the underlying system at an optimal rate (in T), and identifies the parameters approximately up to an equivalent realization (at a rate of $T^{-1/4}$) under further assumptions of observability¹ and controllability. Simchowitz et al. (2019) showed that a *prefiltered* variant of least-squares offers stronger guarantees that apply even to marginally stable systems and systems with adversarial noise. For strictly stable stochastic systems, Sarkar et al. (2019) improve upon previous works to give an optimal rate for parameter identification. We note that these works require the control inputs, and often the noise, to be Gaussian. This may not hold when the control

1. This notion of “observable” is not to be confused with the system being partially observable.

inputs are exogenous (not under user control). In contrast, our results can handle arbitrary (bounded) control inputs.

In a notable departure from this trend, [Tsiamis and Pappas \(2019\)](#) demonstrate optimal recovery of system parameters in the absence of control inputs for marginally stable systems. Under similar conditions, [Tsiamis et al. \(2019\)](#) prove the first result that integrates former system identification results with a perturbation analysis for the Kalman filter to obtain error bounds on prediction. These results only apply to stochastic systems subject to persistent excitation, which our stochastic-case result does not require.

Prediction via improper learning. For strictly stable partially observed systems without process noise, it is sufficient to learn a finite impulse response (FIR) filter on the inputs, as observed by e.g., [Hardt et al. \(2016\)](#). [Tu et al. \(2017\)](#) give near-optimal sample complexity bounds for learning a FIR filter under design inputs. [Hazan et al. \(2017, 2018\)](#); [Arora et al. \(2018\)](#) instead use *spectral filtering* on the inputs to achieve regret bounds that apply in the presence of adversarial dynamics and marginally stable systems, much like the present work. However, in the presence of process noise, the regret compared to the optimal filter can grow linearly. This is an inherent limitation of any FIR-based approach; see [Lee and Zhang \(2019\)](#) for a discussion.

If the LDS is observable, then the associated Kalman filter is strictly stable ([Anderson and Moore, 2012](#)). Using this, [Kozdoba et al. \(2019\)](#) show that in this case, the Kalman filter can be arbitrarily well-approximated by an autoregressive (AR) model. [Anava et al. \(2013\)](#) give algorithms for prediction in ARMA models with adversarial noise. However, their results hold under conditions more stringent than even strict stability. [Kozdoba et al. \(2019\)](#) show that online gradient descent on the AR model gives regret bounds scaling with the size of the observations. As discussed in Section 2.4, in the marginally stable case, this could be polynomial in the time T . [Lee and Zhang \(2019\)](#) give guarantees for learning an autoregressive filter in a stricter notion of \mathcal{H}_∞ norm, but require the system to be strictly stable. In concurrent work, [Tsiamis and Pappas \(2020\)](#) establish polylogarithmic regret in the marginally stable case, but their guarantee depends on the coefficients of the characteristic polynomial, which may be exponentially large (see Section 5.3 and C.1).

Online learning. We use tools from online learning (see [Hazan et al. \(2016\)](#); [Shalev-Shwartz et al. \(2012\)](#); [Cesa-Bianchi and Lugosi \(2006\)](#) for a survey). The standard regret bounds for online least-squares scale with an upper bound on the maximum instantaneous loss, through the gradient norm or the exp-concavity factor [Zinkevich \(2003\)](#); [Hazan et al. \(2007\)](#). Our core argument shows that this quantity is sublinear in T in the LDS setting. This cannot be true for online least-squares for arbitrary polynomially growing x_t , so black-box results cannot apply; see Appendix A. We note the similarity of our approach to [Rakhlin and Sridharan \(2012, 2013\)](#) where the authors show that approximate knowledge of cost functions or gradients revealed one step in advance can give “beyond worst-case” regret bounds.

4. Our results

4.1. Fully-observed LDS

For prediction in a fully-observed LDS, we show that we can achieve sublinear regret in the adversarial setting and polylogarithmic regret in the stochastic setting.

We will make the following assumptions for both theorems:

Assumption 2 *The linear dynamical system*

$$x_t = Ax_{t-1} + Bu_{t-1} + \xi_t$$

with $u_t \in \mathbb{R}^m$, $x_t, \xi_t \in \mathbb{R}^d$, $A \in \mathbb{R}^{d \times d}$, $B \in \mathbb{R}^{d \times m}$ satisfies the following:

- *The initial state is bounded:* $\|x_0\| \leq C_0$.
- *The inputs are bounded:* $\|u_t\| \leq C_u$.
- *The perturbations are bounded:* $\|\xi_t\| \leq C_\xi$ for $1 \leq t \leq T$.
- $\|(A, B)\|_F \leq R$, $\rho(A) \leq 1$, and A can be written in Jordan form as $A = SJS^{-1}$ where J has Jordan blocks of size $\leq r$ and $\|S\|_2 \|S^{-1}\|_2 \leq C_A$.
- B satisfies $\|B\|_2 \leq C_B$.

We will let $C_{\text{sys}} = \max\{C_0, C_A, C_B, C_\xi\}$.

We note that the bound on the perturbations C_ξ is necessary. This prevents, for example, the pathological case when the system switches between two very different linear dynamical systems $x_t = x_{t-1} + u_{t-1}$ and $x_t = -(x_{t-1} + u_{t-1})$ and linear regret is unavoidable. We also note that $\|S\|_2 \|S^{-1}\|_2$ (the condition number of S) is a standard quantity that often appears in learning guarantees.

Our main theorem in the adversarial setting is the following; see the appendix more precise dependences on individual constants.

Theorem 3 (Sublinear regret in the adversarial setting) *Suppose Assumption 2 holds. Then, there is an explicit choice of regularizer μ such that Algorithm 1 achieves regret*

$$\begin{aligned} R_T(A, B) &\leq T^{\frac{2r+1}{2r+2}} \text{poly}(C_{\text{sys}}, R, (d+m), (\ln T), (\ln C_{\text{sys}})) \\ &\quad + T^{\frac{1}{2r+2}} \text{poly}(C_{\text{sys}}, R, (d+m)^r, (\ln T)^r, (\ln C_{\text{sys}})^r). \end{aligned} \quad (4)$$

If $C_{\text{sys}} = O(1)$ and $\|(A, B)\|_F = O(d+m)$,

$$R_T(A, B) = O\left((d+m)^4 d^2 r^2 \cdot T^{\frac{2r+1}{2r+2}} \ln^3 T\right)$$

as $T \rightarrow \infty$. The dependence of μ on T is $T^{\frac{2r+1}{2r+2}}$.

Remark *This is a pessimistic bound. The worst case is when the eigenvalues of the large Jordan blocks are close to 1. If $\|A^k\| \leq C'k^{r'}$, then we can replace the dependence on r with r' , and instead suffer a $\text{poly}(C')$ dependence.*

In the case where A is diagonalizable, Theorem 3 implies $\tilde{O}(T^{3/4})$ regret:

Corollary 4 *Suppose Assumption 2 holds, and further suppose A is diagonalizable. There is an explicit choice of regularizer μ such that Algorithm 2 achieves regret*

$$R_T(A, B) \leq T^{3/4} \text{poly}(C_{\text{sys}}, R, d, m, \ln T). \quad (5)$$

When $C_{\text{sys}} = O(1)$ and $R = O(d+m)$, we have $R_T(A, B) = O((d+m)^{3.5} d T^{3/4} \ln^3 T)$ as $T \rightarrow \infty$. The dependence of μ on T is $T^{3/4}$.

Our main theorem in the stochastic setting is the following.

Theorem 5 (Polylogarithmic regret in stochastic setting) *Suppose Assumption 2 holds, and further that ξ_t is a random variable satisfying $\mathbb{E}[\xi_t | \xi_{t-1}, \dots, \xi_1] = 0$. Then with probability at least $1 - \delta$, Algorithm 2 with $\mu = 1$ achieves regret*

$$R_T(A, B) \leq \text{poly}(C_{\text{sys}}, R, d^r, (\ln T)^r, (\ln C_{\text{sys}})^r, \ln(1/\delta)). \quad (6)$$

Note that there is no requirement that the noise be i.i.d., nor that their covariance is greater than some multiple of the identity. At the expense of a $\ln T$ factor, the theorem can be applied to subgaussian random variables, by first conditioning on the event that $\|\xi_t\| \leq C_\xi$ for all $1 \leq t \leq T$.

4.2. Partially-observed LDS

Our assumptions in the partially observed setting are the following. We will assume that the noise is i.i.d. Gaussian; this is the analogue of the linear-quadratic estimation (LQE) setting where we only care about predicting the observation.

Assumption 6 *The partially-observed LDS defined by (2)–(3) satisfies the following: The perturbations are i.i.d. Gaussian with $\xi_t \sim N(0, \Sigma_x)$ and $\eta_t \sim N(0, \Sigma_y)$. Moreover, the system matrix pair (A, C) is observable (i.e., $O_n = [C; CA; \dots; CA^{n-1}]$ has rank n)², the pair $(A, \Sigma_x^{1/2})$ is controllable (i.e., $[\Sigma_x^{1/2}, A\Sigma_x^{1/2}, \dots, A^{n-1}\Sigma_x^{1/2}]$ has rank n), and $\Sigma_y \succ 0$.*

Assumption 7 *The partially-observed LDS defined by (2)–(3) satisfies the following:*

- *The initial state has steady-state covariance: $x_0 \sim N(x_0^-, \Sigma_0)$ with $\|x_0^-\| \leq C_0$.*
- *The inputs are bounded: $\|u_t\| \leq C_u$.*
- *The perturbations are Gaussian: $\xi_t \sim N(0, \Sigma_x)$ and $\eta_t \sim N(0, \Sigma_y)$.*
- *$\|A\|_F \leq R$, $\rho(A) \leq 1$, and A can be written in Jordan form as $A = SJS^{-1}$ where J has Jordan blocks of size $\leq r$ and $\|S\|_2 \|S^{-1}\|_2 \leq C_A$.*
- *B and C satisfy $\|B\|_2 \leq C_B$ and $\|C\|_2 \leq C_C$.*

We will let $C_{\text{sys}} = \max\{C_0, C_A, C_B, C_C, \|\Sigma_0\|, \|\Sigma_x\|, \|\Sigma_y\|\}$.

For simplicity, our result assumes that x_0 has steady-state covariance. If this is not the case, then one would need quantitative bounds on how quickly the time-varying Kalman filter converges to the steady-state Kalman filter, to bound the additional regret incurred by using a fixed filter.

The theorem also depends on the *sufficient length* $R(\varepsilon)$ of the Kalman filter, which is roughly the length at which we can truncate the unrolled filter to incur an error of at most ε ; see Definition 33 for a precise account. If the filter decays exponentially, then $R(\varepsilon) = O(\ln(1/\varepsilon))$. It remains an interesting problem to handle the case where the filter is *also* marginally stable.

2. We use $[A_1; A_2; \dots]$ to denote a block-column matrix.

Theorem 8 (Polylogarithmic regret for LQE) *Assume Assumptions 6 and 7. Suppose the corresponding Kalman filter is given by A_{KF} , B_{KF} , and C_{KF} . Let $R(\cdot)$ denote the sufficient length of the Kalman filter system, and choose $\ell = R\left(\delta \left[T^r C_C C_A \left(C_B C_u + 10 \|\Sigma_x\| \sqrt{d \ln\left(\frac{T}{\delta}\right)}\right)\right]^{-1}\right)$. Let $(F_t, G_t) = C_{\text{KF}} A_{\text{KF}}^t B_{\text{KF}}$ be the unrolled Kalman filter, and suppose $\sum_{t=0}^{\ell-1} \|F_t\|^2 + \|G_t\|^2 \leq R^2$. Then with probability $1 - \delta$, Algorithm 3 with $\mu = 1$ achieves regret*

$$R_T(A, B, C) \leq \text{poly}(\ell, C_{\text{sys}}, R, (d+m)^r, (\ln T)^r, (\ln C)^r, \ln(1/\delta)). \quad (7)$$

5. Outlines of main proofs

In this section, we explain the key ideas behind our results, using a fully observable LDS with adversarial noise as an example. For simplicity, we sketch the proof in the case that the matrix A is diagonalizable, matrix B is zero, and $\|\xi_s\|_2 = O(1)$, which already captures the core difficulty of the problem. In the proof sketch, we assume that all relevant parameters for the LDS are $O(1)$. Full proofs are in Sections B–E.

5.1. Regret bounds for online least squares with large inputs

Our starting point is the regret bound for online least squares, Theorem 1, which depends on the maximum prediction error $\max_{0 \leq t \leq T-1} \|y_t - A_t x_t\|^2$. To obtain sublinear regret using this bound, we must show that the maximum prediction error is $o(T/\log T)$. We show in Theorem 12 that this holds as long as the following structural condition (formally defined in Definition 11) on the regressor sequence holds:

Definition 9 (Anomaly-free sequences; informal) *A sequence $(x_t)_{t < T}$ is **anomaly-free** if whenever the projection of any x_t onto a unit vector w is large, then there must have been $\Omega(|w^\top x_t|)$ indices $s < t$ for which the projection of x_s to w has norm at least $\Omega(|w^\top x_t|)$.*

Intuitively, the inputs are anomaly-free if no input x_t is large in a direction where we have not already seen many inputs. Note this does not hold in the general case of polynomially-bounded x_t ; see the counterexample in Appendix A. To prove Theorem 12, we first express $A_t x_t - y_t$ in terms of the preceding states and errors $\{x_s, \xi_s\}_{s < t}$ (Lemma 13). Next, we show this expression is bounded in the 1-dimensional case (Lemma 14). Finally, we reduce the general d -dimensional case to the 1-dimensional case by diagonalizing the sample covariance matrix $\sum_{s=0}^{t-1} x_s x_s^\top$. In the reduction, we project onto the eigendirections; this is why we want there to be many large inputs when projected to any direction w .

Note that Theorem 12 is stated more generally, allowing $\Omega(|w^\top x_t|^\alpha)$ indices where $|w^\top x_s|$ is large. This allows superlinear growth in the x_t , and hence can be applied to dynamical systems with matrices having Jordan blocks.

5.2. Proving LDS states are anomaly-free

Our main result (Theorem 3) follows by verifying that LDS states are anomaly-free. The main idea is that the evolution of the LDS ensures that x_t is always approximately a linear combination of past states with small coefficients. More precisely, we need $x_t = \left(\sum_{s=0}^{t-1} a_s x_s\right) + v$ with small $\sum_{s=0}^{t-1} |a_s|$ and $\|v\|_2$ (Lemma 20). Once we have this, projecting onto w gives $w^\top x_t = \left(\sum_{s=0}^{t-1} a_s w^\top x_s\right) + w^\top v$,

showing that one of the projections $|w^\top x_s|$ is large. To obtain many indices s for which this is large, we apply the same argument to the k -step dynamical systems defined by A^2 , A^3 , and so forth, keeping track of how many times an index can be overcounted. This shows the states are anomaly-free and finishes the proof of Theorem 3.

We provide two approaches to decompose x_t into previous states as needed. The simpler approach provides coefficients of size $\exp(d)$, while a more involved approach provides $\text{poly}(d)$ sized coefficients. In order to have a $\text{poly}(d)$ dependence in the final regret bound, the latter approach is necessary.

5.3. $\exp(d)$ -sized coefficients using the Cayley-Hamilton theorem

To show that x_t is always a linear combination of previous states with small coefficients, a first idea is to use the Cayley-Hamilton theorem. In the noiseless case, the theorem implies that the x_t satisfy a recurrence $x_t = \sum_{i=1}^d a_i x_{t-i}$, where a_i are the coefficients of the characteristic polynomial of A . Adding the noise back, we may get an error term v of size $\sum_{i=1}^d |a_i|$ by inspecting how the noise propagates through this recurrence. Even though $\sum_{i=1}^d |a_i|$ can be as large as 2^d , this suffices to get a bound that is sublinear in T while exponential in d . For ease of reading, we first present this weaker result in Lemma 17 in Section C.1.³

5.4. $\text{poly}(d)$ -sized coefficients via volume doubling

As an alternative, we now present a novel volume-doubling argument leading to a recurrence with only $\text{poly}(d)$ coefficient size, which may be of independent interest. For the ease of presentation, we introduce some notation below:

- ℓ_1 -span of S . $\Delta(S) := \{\sum_{u \in S} a_u u : a_u \in \mathbb{R}, \sum_{u \in S} |a_u| \leq 1\}$. For short, we denote $\Delta_t := \Delta(\{x_s\}_{s=0}^t)$.
- Norm with respect to $\Delta(S)$. $\|x\|_{\Delta(S)} := \min_{a_s, v} \sum_{s=0}^t |a_s| + \|v\|_2$ s.t. $x = (\sum_{s=0}^t a_s x_s) + v$ for $x_s \in S$.
- Set of “outlier” indices. $I_t := \{0 \leq s \leq t : \|x_s\|_{\Delta_{s-1}} \geq \frac{2}{\ln 2} d\}$.

Now, our goal is summarized as bounding $\|x_t\|_{\Delta_{t-1}}$. To this end, we first prove an upper bound on the size of $|I_t|$ using a general potential-based argument, and then relate it to the norm $\|x_t\|_{\Delta_{t-1}}$ by unrolling the dynamics appropriately.

Bounding the number of outliers. We have to show that $|I_t|$ is at most polynomial in d and logarithmic in t . We prove a general lemma that if $\|x\|_{\Delta(S)}$ is large enough ($\|x\|_{\Delta(S)} \geq \frac{2}{\ln 2} d$), then adding x to the S increases its volume significantly: $\text{Vol}(\Delta(S \cup \{x\})) \geq 2 \text{Vol}(\Delta(S))$ (Lemma 22). Applied to our situation, this shows that if $\|x_t\|_{\Delta_{t-1}} \geq \frac{2}{\ln 2} d$, then $\text{Vol}(\Delta_t) \geq 2 \text{Vol}(\Delta_{t-1})$. The total volume is bounded by $(\max_{1 \leq s \leq t} \|x_s\|)^d$, so the number of outliers at or before time t is bounded by $\log_2 \left[(\max_{1 \leq s \leq t} \|x_s\|)^d \right]$, which is polynomial in d and logarithmic in t (Lemma 23).

3. We note that an alternate approach to the one we present below is to find a multiple of the characteristic polynomial with small coefficients; see Appendix G.

Bounding $\sum_{s=0}^{t-1} |a_s|$ and $\|v\|_2$. We obtain an inequality showing that $\|x_t\|_{\Delta_{t-1}}$ is not much larger than $\|x_{t-k}\|_{\Delta_{t-k-1}}$ for small delays k . To see this, note that x_t is generated from x_{t-k} by evolving the LDS k times, keeping track of noise. The contribution from the noise here is at most $\text{poly}(k)$. Now, it suffices to find a small k such that $\|x_{t-k}\|_{\Delta_{t-k-1}}$ is small, or in other words $t-k$ is not an outlier. Because the number of outliers is $O(d \log t)$, we can choose $k = O(d \log T)$, and we conclude $\|x_t\|_{\Delta_{t-1}}$ is at most $\text{poly}(d \log t)$. This argument is formalized in the proof of Lemma 20.

Controlling overcounting using prime index gaps. One technicality is that we have to apply the same argument to the dynamical systems $x_t \approx A^p x_{t-p}$ for different values of p . For each value of p , we get an index $t-pk$ such that $w^\top x_{t-pk}$ is large. To make sure we obtain enough indices this way, in the proof of Lemma 26 we only take p to be prime, and we use a lower bound on primorials $\prod_{\text{prime } p \leq X} p$ to show that we can collect enough distinct indices.

5.5. Stochastic cases

Finally, we provide a brief comment on how to prove Theorems 5 and 8. In the fully-observed setting, we can use a martingale argument, rather than the structural result, to obtain $\text{poly} \log(T)$ regret. To do this, we show that with high probability, $\max_{0 \leq t \leq T-1} \|A_t v_t - v_{t+1}\|$ is bounded by $\text{poly} \log(T)$ (Lemma 28).

Let A_t be the matrix predicted by online least squares with regularization parameter μ . By Lemma 13, $A_t x_t - x_{t+1} = \sum_{s=0}^{t-1} \xi_{s+1} v_s^\top \Sigma_t^{-1} v_t - \mu A \Sigma_t^{-1} x_t - \xi_{t+1}$. The main term we need to bound is the first one. Let $b_s = v_s^\top \Sigma_t^{-1} v_t$. Pretending for a moment that b_s is $\mathcal{F}_s = \sigma(\xi_1, \dots, \xi_s)$ -valued, by Azuma's inequality it suffices to bound the variation $\sum_{s=0}^{t-1} b_s^2$. We have already shown that $v_t = \sum_{s=0}^{t-1} a_s v_s$ with $\sum_{s=0}^{t-1} a_s^2 \leq \left(\sum_{s=0}^{t-1} |a_s|\right)^2 \leq O(1)$, so we can bound the variation $\sum_{s=0}^{t-1} b_s^2$ in terms of $\sum_{s=0}^{t-1} \left\| \Sigma_t^{-\frac{1}{2}} v_s \right\|^2$. By definition $\Sigma_t = \mu I + \sum_{s=0}^{t-1} v_s v_s^\top$, so $\sum_{s=0}^{t-1} \left\| \Sigma_t^{-\frac{1}{2}} v_s \right\|^2 \leq d$.

However, we can't apply Azuma's inequality directly, since b_s depends on v_t , and thus is not \mathcal{F}_s -valued. However, the dependence on non- \mathcal{F}_s -valued random variables is only through $z_t = \Sigma_t^{-1} v_t$, which does not depend on s , so we can use Azuma's inequality on an ε -net of possible values for $z = \Sigma_t^{-1} v_t$. More precisely, note that z_t has the property that $\sum_{s=0}^{t-1} (v_s^\top z_t)^2$ is small, so it suffices to bound $\sum_{s=0}^{t-1} \xi_{s+1} (v_s^\top z)$ for a ε -net of z such that $\sum_{s=0}^{t-1} (v_s^\top z)^2$ is small. These z 's live in a d -dimensional space, so we only incur factors of d .

For the partially-observed setting, we reduce to Theorem 5 by lifting the state: we choose a large enough horizon ℓ such that truncating the unrolled Kalman filter to length ℓ incurs an approximation error of at most ε . Then, by letting the state space be the past ℓ observations and inputs, the partially observed LDS is approximately described by a fully-observed LDS. This incurs an additional polynomial factor in the length ℓ .

6. Conclusion

We have shown that online least-squares, with a carefully chosen regularization and a refined analysis, has a sublinear regret guarantee in marginally stable linear dynamical systems, even in the most difficult cases when the state can grow polynomially. In the stochastic setting, adopting the same view of low-regret prediction as opposed to parameter recovery, we have shown logarithmic regret bounds and bypassed usual isotropic noise assumptions. We list a few plausible lines of further inquiry in appendix H.

References

- Oren Anava, Elad Hazan, Shie Mannor, and Ohad Shamir. Online learning for time series prediction. In *COLT 2013 - The 26th Annual Conference on Learning Theory, June 12-14, 2013, Princeton University, NJ, USA*, pages 172–184, 2013.
- Brian DO Anderson and John B Moore. *Optimal filtering*. Courier Corporation, 2012.
- Sanjeev Arora, Elad Hazan, Holden Lee, Karan Singh, Cyril Zhang, and Yi Zhang. Towards provable control for unknown linear dynamical systems. *International Conference on Learning Representations, Workshop track*, 2018. URL <https://openreview.net/forum?id=BygpQlbA->.
- Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- Xue Chen, Daniel M Kane, Eric Price, and Zhao Song. Fourier-sparse interpolation without a frequency gap. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 741–750. IEEE, 2016.
- Pierre-André Chiappori, Pierre-Yves Geoffard, and Roger Guesnerie. Sunspot fluctuations around a steady state: the case of multidimensional one-step forward looking economic models. *Econometrica: Journal of the Econometric Society*, pages 1097–1126, 1992.
- Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the sample complexity of the linear quadratic regulator. *arXiv preprint arXiv:1710.01688*, 2017.
- Mohamad Kazem Shirani Faradonbeh, Ambuj Tewari, and George Michailidis. Finite time identification in unstable linear systems. *Automatica*, 96:342–353, 2018.
- Zoubin Ghahramani and Geoffrey E Hinton. Parameter estimation for linear dynamical systems. Technical report, Technical Report CRG-TR-96-2, University of Totronto, Dept. of Computer Science, 1996.
- Moritz Hardt, Tengyu Ma, and Benjamin Recht. Gradient descent learns linear dynamical systems. *arXiv preprint arXiv:1609.05191*, 2016.
- P Jeff Harrison. Convergence and the constant dynamic linear model. *Journal of Forecasting*, 16(5):287–292, 1997.
- Elad Hazan and Zohar Karnin. Volumetric spanners: an efficient exploration basis for learning. *The Journal of Machine Learning Research*, 17(1):4062–4095, 2016.
- Elad Hazan, Amit Agarwal, and Satyen Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.
- Elad Hazan, Karan Singh, and Cyril Zhang. Learning linear dynamical systems via spectral filtering. In *Advances in Neural Information Processing Systems*, pages 1–2, 2017.
- Elad Hazan, Holden Lee, Karan Singh, Cyril Zhang, and Yi Zhang. Spectral filtering for general linear dynamical systems. *arXiv preprint arXiv:1802.03981*, 2018.

- Elad Hazan et al. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- David F Hendry and Jurgen A Doornik. Modelling linear dynamic econometric systems. *Scottish Journal of Political Economy*, 41(1):1–33, 1994.
- Rudolf Emil Kalman. Mathematical description of linear dynamical systems. *Journal of the Society for Industrial and Applied Mathematics, Series A: Control*, 1(2):152–192, 1963.
- Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82.1:35–45, 1960.
- Edward W Kamen and Jonathan K Su. *Introduction to optimal estimation*. Springer Science & Business Media, 1999.
- Mark Kozdoba, Jakub Marecek, Tigran T. Tchakian, and Shie Mannor. On-line learning of linear dynamical systems: Exponential forgetting in kalman filters. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4098–4105, 2019.
- Holden Lee and Cyril Zhang. Robust guarantees for learning an autoregressive filter. *arXiv preprint arXiv:1905.09897*, 2019.
- Melvyn B Nathanson. *Additive Number Theory The Classical Bases*, volume 164. Springer Science & Business Media, 2013.
- Samet Oymak and Necmiye Ozay. Non-asymptotic identification of lti systems from a single trajectory. *arXiv preprint arXiv:1806.05722*, 2018.
- S Joe Qin. An overview of subspace identification. *Computers & chemical engineering*, 30(10-12):1502–1513, 2006.
- Alexander Rakhlin and Karthik Sridharan. Online learning with predictable sequences. *arXiv preprint arXiv:1208.3728*, 2012.
- Sasha Rakhlin and Karthik Sridharan. Optimization, learning, and games with predictable sequences. In *Advances in Neural Information Processing Systems*, pages 3066–3074, 2013.
- Tuhin Sarkar, Alexander Rakhlin, and Munther A Dahleh. Finite-time system identification for partially observed lti systems of unknown order. *arXiv preprint arXiv:1902.01848*, 2019.
- Tuhin Subhra Sarkar and Alexander Rakhlin. Near optimal finite time identification of arbitrary linear dynamical systems. In *International Conference on Machine Learning*, pages 5610–5618, 2019.
- Peter T Saunders. Evolution without natural selection: further implications of the daisyworld parable. *Journal of Theoretical Biology*, 166(4):365–373, 1994.
- Shai Shalev-Shwartz et al. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194, 2012.

Max Simchowitz, Horia Mania, Stephen Tu, Michael I Jordan, and Benjamin Recht. Learning without mixing: Towards a sharp analysis of linear system identification. *arXiv preprint arXiv:1802.08334*, 2018.

Max Simchowitz, Ross Boczar, and Benjamin Recht. Learning linear dynamical systems with semi-parametric least squares. *arXiv preprint arXiv:1902.00768*, 2019.

Anastasios Tsiamis and George Pappas. Online learning of the kalman filter with logarithmic regret. *arXiv preprint arXiv:2002.05141*, 2020.

Anastasios Tsiamis and George J. Pappas. Finite sample analysis of stochastic system identification. *ArXiv*, abs/1903.09122, 2019.

Anastasios Tsiamis, Nikolai Matni, and George J. Pappas. Sample complexity of kalman filtering for unknown systems, 2019.

Stephen Tu, Ross Boczar, Andrew Packard, and Benjamin Recht. Non-asymptotic analysis of robust control from coarse-grained identification. *arXiv preprint arXiv:1707.04791*, 2017.

Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.

Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning*, pages 928–936, 2003.

Appendix A. Impossibility of sublinear regret without the LDS

For sake of completeness, we include a simple one-dimensional counterexample showing the insufficiency of black-box regret bounds for online least-squares when it is only assumed that an upper bound for the regressors grows with time. Even without adversarial perturbations, this information-theoretic lower bound shows that we require a refined notion of gradual growth of regressors x_t , as in the structural results of Section 5.2, to achieve sublinear regret.

Proposition 10 *There is a joint distribution over $a \in [-1, 1]$ and length- T sequences $x_t, y_t \in \mathbb{R}$, for which $y_t = ax_t$ and $|x_t| \leq t$, but any online algorithm incurs at least T^2 expected regret.*

Proof We construct this distribution, choosing $a = \pm 1$ with equal probability. We choose $x_t = y_t = 0$ for all $1 \leq t \leq T - 1$, then choose $x_T = T$, so that $y_T = ax_T$. In this example, at time T , all previous feedback (x_t, y_t) is independent of a , so the best prediction at time T is $\hat{y}_T = 0$, which suffers expected least-squares loss T^2 . ■

Appendix B. Anomaly-free inputs imply sublinear OLS regret

We begin by defining a structural condition on OLS inputs, whereby if any time x_t is large in a direction then many previous x_s 's are also large in that direction.

Definition 11 (Anomaly-free sequences) A sequence (x_t) is (c, c_1, c_2, α) -anomaly-free if for any t and any unit vector $w \in \mathbb{R}^d$, if $|w^\top x_t| = M > c$, there exist at least $c_1 M^\alpha$ indices $s < t$ such that $|w^\top x_s| \geq c_2 M$.

We show that when every y_t is obtained from anomaly-free x_t from a fixed linear transformation, plus a bounded perturbation, then OLS attains sublinear regret.

Theorem 12 For constants $C_\xi, c, c_1, c_2, \alpha \in (0, 1]$, $c_1 \leq \frac{1}{2}$, $c \geq \frac{1}{c_1}$, suppose an online least-squares problem satisfies the following conditions:

1. There exists $A^* \in \mathbb{R}^{n \times m}$ such that for all t , $y_t = A^* x_t + \xi_t$ with $\|\xi_t\| \leq C_\xi$.
2. The input vectors are bounded: $\|x_t\| \leq B$.
3. $(x_s)_{s < t}$ is (c, c_1, c_2, α) -anomaly-free.

Then online least squares with regularization parameter μ incurs regret

$$R_t(A) \leq \mu \|A\|_F^2 + O\left(m^3 \max\left\{\frac{C_\xi c^2 t}{\mu} + \|A\|_2 c, \frac{C_\xi}{c_2^2}, \frac{1}{c_1^{\frac{1}{\alpha+2}} c_2^{\frac{2}{\alpha+2}}}\left(\|A\|_2 \mu^{\frac{1}{\alpha+2}} + \frac{C_\xi t^{\frac{1}{2}}}{\mu^{\frac{1}{2} - \frac{1}{\alpha+2}}}\right)\right\} \cdot \ln\left(1 + \frac{tB^2}{m}\right)\right).$$

If (1a) $\mu \geq c \frac{4(\alpha+2)}{\alpha+4} c_1^{\frac{2}{\alpha+4}} c_2^{\frac{4}{\alpha+4}} t^{\frac{\alpha+2}{\alpha+4}}$, (1b) $\|A\|_2 \leq \frac{C_\xi c t}{\mu}$, (2) $\mu \leq \frac{c_2^{\frac{4(\alpha+2)}{\alpha}} t^{\frac{\alpha+2}{\alpha}}}{c_1^{\frac{2}{\alpha}}}$, and (3) $\mu \leq \frac{t}{C_\xi^2 \|A\|^2}$,

then $R_t(A) \leq \mu \|A\|_F^2 + O\left(\frac{m^3 C_\xi^2 t}{c_1^{\frac{2}{\alpha+2}} c_2^{\frac{4}{\alpha+2}} \mu^{1 - \frac{2}{\alpha+2}}}\right)$. If $\|A\|_F = O(m)$, $C_\xi = O(1)$, and $\mu = \frac{m^{\frac{\alpha+2}{\alpha+1}} t^{\frac{\alpha+2}{2(\alpha+1)}}}{c_1^{\frac{1}{\alpha+1}} c_2^{\frac{2}{\alpha+1}}}$ satisfies these conditions, then

$$R_t(A) = O\left(\frac{m^{2 + \frac{1}{\alpha+1}} t^{\frac{\alpha+2}{2(\alpha+1)}} \ln\left(1 + \frac{tB^2}{m}\right)}{c_1^{\frac{1}{\alpha+1}} c_2^{\frac{2}{\alpha+1}}}\right) \quad (8)$$

Note that if $t \rightarrow \infty$ while the other constants are fixed, the inequalities do hold for the choice of μ . The only inequality that is not immediately clear is (1a); it follows from comparing the exponents of t : $\frac{\alpha+2}{2(\alpha+1)} \geq \frac{\alpha+2}{\alpha+4}$ for $\alpha \leq 1$.

By Theorem 1, it suffices to bound $\|A_t x_t - y_t\|$, which is done in Lemma 15. We start with the following calculation.

Lemma 13 Let $x_t \in \mathbb{R}^m$, $A \in \mathbb{R}^{n \times m}$, $\xi_t \in \mathbb{R}^n$ for $0 \leq s \leq t-1$. Suppose $y_s = Ax_s + \xi_s$ for each $0 \leq s \leq t-1$. Let $A_t = \operatorname{argmin}_{A \in \mathbb{R}^{n \times m}} [\mu \|A\|_F^2 + \sum_{s=0}^{t-1} \|y_t - Ax_t\|_2^2]$. Let

$$\Sigma_t = \mu I_m + \sum_{s=0}^{t-1} x_s x_s^\top.$$

Then

$$A_t x_t - y_t = \sum_{s=0}^{t-1} \xi_s x_s^\top \Sigma_t^{-1} x_t - \mu A \Sigma_t^{-1} x_t - \xi_t. \quad (9)$$

Proof We calculate

$$\begin{aligned}
 A_t &= \operatorname{argmin}_A \left[\mu \|A\|_F^2 + \sum_{s=0}^{t-1} \|y_t - Ax_t\|_2^2 \right] \\
 &= \sum_{s=0}^{t-1} y_s x_s^\top \Sigma_t^{-1} = \sum_{s=0}^{t-1} (Ax_s x_s^\top \Sigma_t^{-1} + \xi_s x_s^\top \Sigma_t^{-1}) \\
 A &= A \left(\sum_{s=0}^{t-1} x_s x_s^\top + \mu I_m \right) \Sigma_t^{-1} = \sum_{s=0}^{t-1} Ax_s x_s^\top \Sigma_t^{-1} + \mu A \Sigma_t^{-1} \\
 A_t - A &= \sum_{s=0}^{t-1} \xi_s x_s^\top \Sigma_t^{-1} - \mu A \Sigma_t^{-1} \\
 A_t x_t - y_t &= (A_t - A)x_t - \xi_t \\
 &= \sum_{s=0}^{t-1} \xi_s x_s^\top \Sigma_t^{-1} x_t - \mu A \Sigma_t^{-1} x_t - \xi_t.
 \end{aligned}$$

■

We now bound the size of the residual when condition 3 from Theorem 12 holds. We focus on the one-dimensional case in Lemma 14 where the scale of the unknown parameter is a . Afterwards, we bound the residuals in the general case by diagonalization and reduction to the one-dimensional result in Lemma 15.

Lemma 14 *Let c, c_1, c_2 be constants, $c_1 \leq \frac{1}{2}$, $c \geq \frac{1}{c_1}$. Let $0 < \alpha \leq 1$. Suppose that $(x_s)_{s < t}$ is (c, c_1, c_2, α) -anomaly-free and suppose $|\xi_s| \leq C_\xi$ for each s , then for $a > 0$,*

$$\begin{aligned}
 &\sum_{s=0}^{t-1} \left| \xi_s x_s \left(\mu + \sum_{r=0}^{t-1} x_r^2 \right)^{-1} x_t \right| + a \left| \mu \left(\mu + \sum_{s=0}^{t-1} x_s^2 \right)^{-1} x_t \right| \\
 &\leq O \left(\max \left\{ \frac{C_\xi c^2 t}{\mu} + ac, \frac{C_\xi}{c_2^2}, \frac{1}{c_1^{\frac{1}{\alpha+2}} c_2^{\frac{2}{\alpha+2}}} \left(a \mu^{\frac{1}{\alpha+2}} + \frac{C_\xi t^{\frac{1}{2}}}{\mu^{\frac{1}{2} - \frac{1}{\alpha+2}}} \right) \right\} \right).
 \end{aligned}$$

Recall, from Theorem 1, the univariate regret then scales as the square of this quantity plus a regularization term scaling linearly in μ . For the case $\alpha = 1$, to balance the terms, we set $\left(\frac{t^{\frac{1}{2}}}{\mu^{\frac{1}{6}}} \right)^2 \sim \mu$, so $\mu = t^{\frac{3}{4}}$ with regret $\tilde{O}(t^{\frac{3}{4}})$.

Proof Let $M = \max_{0 \leq s \leq t-1} |x_s|$.

First suppose $M \leq c$. Then

$$\left| \sum_{s=0}^t \xi_s x_s \left(\mu + \sum_{r=0}^{t-1} x_r^2 \right)^{-1} x_t \right| \leq \frac{C_\xi c^2 t}{\mu}$$

and the second term satisfies

$$\left| \mu \left(\mu + \sum_{s=0}^{t-1} x_s^2 \right)^{-1} x_t \right| \leq c.$$

Now suppose $M > c$. Then there exists a subsequence x_{i_1}, \dots, x_{i_n} where $n = \lceil c_1 M^\alpha \rceil > 1$ and $i_j < t$, $|x_{i_j}| \geq c_2 M$ for each j . Let $S = \{i_1, \dots, i_n\}$. We have

$$\sum_{k=1}^n |x_{i_k}| \leq \lceil c_1 M^\alpha \rceil M = O(c_1 M^{\alpha+1}) \quad \sum_{k=1}^n x_{i_k}^2 \geq \lceil c_1 M^\alpha \rceil (c_2 M)^2 = \Omega(c_1 c_2^2 M^{\alpha+2}).$$

Thus

$$\begin{aligned} & \sum_{s=0}^{t-1} \left| \xi_{s+1} x_s \left(\mu + \sum_{r=0}^{t-1} x_r^2 \right)^{-1} \right| + a \left| \mu \left(\mu + \sum_{s=0}^{t-1} x_s^2 \right)^{-1} \right| \leq \frac{a\mu + C_\xi \sum_{s=1}^{t-1} |x_s|}{\mu + \sum_{s=1}^{t-1} x_s^2} \\ & \lesssim \frac{a\mu + C_\xi (c_1 M^{\alpha+1} + \sum_{i \in [t-1] \setminus S} |x_i|)}{\mu + c_1 c_2^2 M^{\alpha+2} + \sum_{i \in [t-1] \setminus S} x_i^2} \\ & \leq \max_{0 \leq x'_i \leq M \text{ for } 1 \leq i \leq t-1-n} \frac{a\mu + C_\xi (c_1 M^{\alpha+1} + \sum_{i=1}^{t-1-n} x'_i)}{\mu + c_1 c_2^2 M^{\alpha+2} + \sum_{i=1}^{t-1-n} x_i'^2} \\ & \leq \max_{0 \leq x \leq M} \frac{a\mu + C_\xi (c_1 M^{\alpha+1} + (t-1-n)x)}{\mu + c_1 c_2^2 M^{\alpha+2} + (t-1-n)x^2} \end{aligned}$$

where the last inequality follows by Jensen's inequality since $x \mapsto x^2$ is convex (holding $\sum_{i=1}^{t-1-n} x'_i$ constant, the smallest value of $\sum_{i=1}^{t-1-n} x_i'^2$ is attained when they are equal). Continuing,

$$\begin{aligned} & \leq \max_{0 \leq x \leq M} \left(\frac{C_\xi c_1 M^{\alpha+1}}{\mu + c_1 c_2^2 M^{\alpha+2}} + \frac{a\mu + C_\xi (t-1-n)x}{\mu + c_1 c_2^2 M^{\alpha+2} + (t-1-n)x^2} \right) \\ & \leq \max_{0 \leq x \leq M} \left(\frac{C_\xi c_1 M^{\alpha+1}}{\mu + c_1 c_2^2 M^{\alpha+2}} + \frac{a\mu + C_\xi t x}{\mu + c_1 c_2^2 M^{\alpha+2} + t x^2} \right) \end{aligned}$$

because $x \mapsto \frac{ax}{b+cx}$ is increasing for $a, b, c > 0$ and $x \geq 0$. Thus

$$\begin{aligned} & \sum_{s=0}^{t-1} \left| \xi_s x_s \left(\mu + \sum_{r=0}^{t-1} x_r^2 \right)^{-1} x_t \right| + a \left| \mu \left(\mu + \sum_{s=0}^{t-1} x_s^2 \right)^{-1} x_t \right| \\ & \lesssim M \cdot \max_{0 \leq x \leq M} \left(\frac{C_\xi c_1 M^{\alpha+1}}{\mu + c_1 c_2^2 M^{\alpha+2}} + \frac{a\mu + C_\xi t x}{\mu + c_1 c_2^2 M^{\alpha+2} + t x^2} \right) \\ & \leq \max_{0 \leq x \leq M} \left(\frac{C_\xi}{c_2^2} + \frac{aM\mu + C_\xi M t x}{\mu + c_1 c_2^2 M^{\alpha+2} + t x^2} \right) \end{aligned}$$

Applying the weighted arithmetic-geometric mean inequality, the second term is bounded by

$$\begin{aligned} & aM\mu \left(\frac{1}{\mu} \right)^{\frac{\alpha+1}{\alpha+2}} \left(\frac{1}{c_1 c_2^2 M^{\alpha+2}} \right)^{\frac{1}{\alpha+2}} + C_\xi M t x \left(\frac{1}{\mu} \right)^{\frac{1}{2} - \frac{1}{\alpha+2}} \left(\frac{1}{c_1 c_2^2 M^{\alpha+2}} \right)^{\frac{1}{\alpha+2}} \left(\frac{1}{t x^2} \right)^{\frac{1}{2}} \\ & \leq \frac{1}{c_1^{\frac{1}{\alpha+2}} c_2^{\frac{2}{\alpha+2}}} \left(a\mu^{\frac{1}{\alpha+2}} + \frac{C_\xi t^{\frac{1}{2}}}{\mu^{\frac{1}{2} - \frac{1}{\alpha+2}}} \right). \end{aligned}$$

■

Now we are ready to bound (9) in the multidimensional case.

Lemma 15 Suppose the conditions specified in Theorem 12 are satisfied, then for each $0 \leq t \leq T - 1$,

$$\|A_t x_t - y_t\| \leq O \left(m \max \left\{ \frac{C_\xi c^2 t}{\mu} + \|A\|_2 c, \frac{C_\xi}{c_2^2}, \frac{1}{c_1^{\alpha+2} c_2^{\alpha+2}} \left(\|A\|_2 \mu^{\frac{1}{\alpha+2}} + \frac{C_\xi t^{\frac{1}{2}}}{\mu^{\frac{1}{2} - \frac{1}{\alpha+2}}} \right) \right\} \right).$$

Proof We bound each of the terms in (9).

Because $\sum_{s=0}^{t-1} x_s x_s^\top$ is symmetric, we can find an orthogonal matrix U such that $U \left(\sum_{s=0}^{t-1} x_s x_s^\top \right) U^\top$ is diagonal. Let $z_s = U x_s$. Then

$$\begin{aligned} \sum_{s=0}^{t-1} \xi_s x_s^\top \left(\mu I_m + \sum_{r=0}^{t-1} x_r x_r^\top \right)^{-1} x_t &= \sum_{s=0}^{t-1} \xi_s x_s^\top U^\top \left(\mu I_m + \sum_{r=0}^{t-1} U x_r x_r^\top U^\top \right)^{-1} U x_t \\ &= \sum_{s=0}^{t-1} \xi_s z_s^\top \left(\mu I_m + \sum_{r=0}^{t-1} z_r z_r^\top \right)^{-1} z_t \\ &= \sum_{i=1}^m \sum_{s=0}^{t-1} \xi_s z_{si} \left(\mu + \sum_{r=0}^{t-1} z_{ri}^2 \right)^{-1} z_{ti} \end{aligned} \quad (10)$$

where we can “decouple” the coordinates because $\mu I + \sum_{s=1}^{t-1} z_s z_s^\top$ is diagonal. Similarly,

$$\begin{aligned} \mu A \left(\mu I_m + \sum_{s=0}^{t-1} x_s x_s^\top \right)^{-1} x_t &= \mu A U^\top \left(\mu I + \sum_{s=0}^{t-1} z_s z_s^\top \right)^{-1} z_t \\ &= \sum_{i=1}^m (A U^\top)_{\cdot i} \mu \left(\mu I_m + \sum_{s=0}^{t-1} z_{si}^2 \right)^{-1} z_{ti} \end{aligned} \quad (11)$$

$$\begin{aligned} \left\| \mu A \left(\mu I_m + \sum_{s=0}^{t-1} x_s x_s^\top \right)^{-1} x_t \right\| &\leq \max_{1 \leq i \leq m} \left\| (A U^\top)_{\cdot i} \right\| \sum_{i=1}^m \left| \mu \left(\mu I + \sum_{s=0}^{t-1} z_{si}^2 \right)^{-1} z_{ti} \right| \\ &\leq \|A\|_2 \sum_{i=1}^m \left| \mu \left(\mu I_m + \sum_{s=0}^{t-1} z_{si}^2 \right)^{-1} z_{ti} \right| \end{aligned}$$

From the hypothesis of the lemma and the fact that $z_{ti} = U_{i \cdot} x_t$, we conclude that z_{ti} satisfies the conditions of Lemma 14. Apply Lemma 14 to the z_{ti} to obtain bounds for

$$\sum_{s=0}^{t-1} \|\xi_s\| \left| z_{si} \left(\mu I + \sum_{r=0}^{t-1} z_{ri}^2 \right)^{-1} z_{ti} \right| + \|A\|_2 \left| \mu \left(\mu I + \sum_{s=0}^{t-1} z_{si}^2 \right)^{-1} z_{ti} \right|$$

for each i . Summing over over i gives a factor of m , and gives a bound for (10) + (11). Finally, the term $-\xi_t$ contributes at most C_ξ , which can be absorbed into the bound. \blacksquare

Proof [Proof of Theorem 12] The theorem follows from plugging in the bound on $\|A_t x_t - y_t\|$ in Lemma 15 into the regret bound of Theorem 1.

For the second statement, we check that (1a) implies $\frac{C_\xi^2 c^2 t}{\mu} \leq \frac{1}{c_1^{\frac{1}{\alpha+2}} c_2^{\frac{2}{\alpha+2}}} \cdot \frac{C_\xi t^{\frac{1}{2}}}{\mu^{\frac{1}{2}-\frac{1}{\alpha+2}}}$, (1b) implies that $\|A\|_2 c \leq \frac{C_\xi c^2 t}{\mu}$, (2) implies that $\frac{C_\xi}{c_2} \leq \frac{1}{c_1^{\frac{1}{\alpha+2}} c_2^{\frac{2}{\alpha+2}}} \cdot \frac{C_\xi t^{\frac{1}{2}}}{\mu^{\frac{1}{2}-\frac{1}{\alpha+2}}}$, and (3) implies that $\|A\|_2 \mu^{\frac{1}{\alpha+2}} \leq \frac{C_\xi t^{\frac{1}{2}}}{\mu^{\frac{1}{2}-\frac{1}{\alpha+2}}}$. Then the maximum term is $O\left(\frac{1}{c_1^{\frac{1}{\alpha+2}} c_2^{\frac{2}{\alpha+2}}} \cdot \frac{C_\xi t^{\frac{1}{2}}}{\mu^{\frac{1}{2}-\frac{1}{\alpha+2}}}\right)$.

For the third statement, we note μ was chosen to equate the two terms. ■

Appendix C. Proof of Theorem 3 (fully observed system, adversarial noise)

We complete the details for the proof outline from Section 5. We start with the following simple observation: if x_T is a linear combination of previous x_t 's with small coefficients, and x_T has large projection in some direction, then at least one of the x_t 's also has large projection in that direction.

Lemma 16 *Suppose there exist $a_t \in \mathbb{R}$ and $v \in \mathbb{R}^d$ such that*

$$x_T = \sum_{t=0}^{T-1} a_t x_t + v \quad (12)$$

and $\sum_{t=0}^{T-1} |a_t| \leq L_a$, $\|v\|_2 \leq L_v$.

Then for any unit vector $w \in \mathbb{R}^d$, there exists $0 \leq t \leq T-1$ such that

$$|w^\top x_t| \geq \frac{|w^\top x_T| - L_v}{L_a}. \quad (13)$$

Proof We have

$$w^\top x_T = \sum_{t=0}^{T-1} a_t w^\top x_t + w^\top v. \quad (14)$$

Applying Hölder's inequality, there exists $0 \leq t \leq T-1$ such that

$$|w^\top x_t| \geq \frac{|w^\top x_T - w^\top v|}{\sum_{t=0}^{T-1} |a_t|} \geq \frac{|w^\top x_T| - |w^\top v|}{L_a}. \quad (15)$$

We then note that $|w^\top v| \leq \|v\|$ as w is a unit vector, and the result follows. ■

C.1. State is a small ℓ_1 combination of previous states (exp(d) version)

As a warm-up, we first give a simpler proof that obtains a bound on $\sum_{t=1}^{T-1} |a_t|$ that is exponential in d . This subsection is for exposition purposes only and may be skipped. In the next subsection we obtain a poly(d) bound. (We note that an alternate way of obtaining a poly(d) bound is by using Lemma 17 with multiple of the characteristic polynomial with small coefficients; see Appendix G.)

Lemma 17 (Large past x 's via Cayley-Hamilton) *Given Assumptions 2 with $B = O$, suppose that A is diagonalizable as $A = VDV^{-1}$ where $\|V\| \|V^{-1}\| \leq C_A$. Then for any unit vector $w \in \mathbb{R}^d$, there exist $\min \left\{ \left\lfloor \frac{|w^\top x_t|}{d^2 2^{d+1} C_A C_\xi} \right\rfloor, \left\lfloor \frac{t}{d^2} \right\rfloor \right\}$ values of s , $0 \leq s \leq t - 1$, such that*

$$|w^\top x_s| \geq \frac{|w^\top x_t|}{2^{d+1}}.$$

Note that in fact the d in the bound can be replaced with the degree of the minimal polynomial.

Proof By the Cayley-Hamilton Theorem, if $p(x) = \sum_{i=0}^d a_i^{(1)} x^i$ is the characteristic polynomial of A , then $\sum_{i=0}^d a_i^{(1)} A^i = O$. We can bound the size of the coefficients as follows. Let r_1, \dots, r_d be the roots of $p(x) = 0$. Then $p(x) = \prod_{i=1}^d (x - r_i)$. Because $\rho(A) \leq 1$, every root satisfies $|r_i| \leq 1$. By the triangle inequality, the coefficients of $p(x)$ are at most the corresponding coefficients of $q(x) = \prod_{i=1}^d (x + |r_i|)$ in absolute value. The sum of coefficients of $q(x)$ is $q(1) \leq 2^d$, so $\sum_{i=0}^d |a_i^{(1)}| \leq 2^d$.

We now proceed to bound the error term due to the noise. By unfolding the recurrence we obtain

$$\begin{aligned} x_t &= A^d x_{t-d} + \sum_{\tau=1}^d A^{d-\tau} \xi_{t-d+\tau} \\ &= \sum_{i=0}^{d-1} -a_i^{(1)} A^i x_{t-d} + \sum_{\tau=1}^d A^{d-\tau} \xi_{t-d+\tau} \\ &= \sum_{i=0}^{d-1} -a_i^{(1)} \left(x_{t-d+i} - \sum_{\tau=1}^i A^{i-\tau} \xi_{t-d+\tau} \right) + \sum_{\tau=1}^d A^{d-\tau} \xi_{t-d+\tau} \\ &= \sum_{i=0}^{d-1} -a_i^{(1)} x_{t-d+i} + \underbrace{\sum_{i=0}^{d-1} a_i^{(1)} \sum_{\tau=1}^i A^{i-\tau} \xi_{t-d+\tau}}_{=: v^{(1)}} + \sum_{\tau=1}^d A^{d-\tau} \xi_{t-d+\tau} \end{aligned} \quad (16)$$

Note that $\|A^i \xi\| \leq \|VD^i V^{-1}\| \|\xi\| \leq C_A \|\xi\| \leq C_A C_\xi$. We write $x_t = -\sum_{i=1}^{d-1} a_i^{(1)} x_{t-d+i} + v^{(1)}$ where

$$\begin{aligned} \|v^{(1)}\| &= \left\| \sum_{i=0}^{d-1} a_i^{(1)} \sum_{\tau=1}^i A^{i-\tau} \xi_{t-d+\tau} + \sum_{\tau=1}^d A^{d-\tau} \xi_{t-d+\tau} \right\| \\ &\leq \sum_{i=0}^{d-1} |a_i^{(1)}| \sum_{\tau=1}^i C_A \|\xi_{t-d+\tau}\| + \sum_{\tau=1}^d C_A \|\xi_{t-d+\tau}\| \\ &\leq 2^d (d-1) C_A C_\xi + d C_A C_\xi \leq d 2^d C_A C_\xi \end{aligned}$$

If $w \in \mathbb{R}^d$ is a unit vector such that $|w^\top x_t| \geq d 2^{d+1} C_A C_\xi$, we have $|w^\top x_t| \geq 2 \|v^{(1)}\|$. Noting that $w^\top x_t = \sum_{i=0}^{d-1} -a_i^{(1)} x_{t-d+i} + v^{(1)}$, by Lemma 16 there exists an index $0 \leq i \leq d-1$ such that

$$|w^\top x_{t-d+i}| \geq \frac{|w^\top x_t| - \|v^{(1)}\|}{\sum_{i=0}^{d-1} |a_i^{(1)}|} \geq \frac{|w^\top x_t|}{2 \sum_{i=0}^{d-1} |a_i^{(1)}|} \geq \frac{|w^\top x_t|}{2^{d+1}}$$

In order to obtain many large past x 's, we apply the same argument on sequences $x_t, x_{t-k}, \dots, x_{t-kd}$ by considering the recurrence

$$x_t = A'x_{t-k} + \xi_t^{(k)} \quad (17)$$

$$A' = A^k \quad (18)$$

$$\xi_t^{(k)} = \sum_{j=1}^k A^{k-j} \xi_{t-k+j}. \quad (19)$$

Let $p_k(x) = \sum_{i=0}^d a_i^{(k)} x^i$ be the characteristic polynomial of A^k . Defining $v^{(k)}$ similarly to $v^{(1)}$, we have

$$\|v^{(k)}\| = \left\| \sum_{i=0}^{d-1} a_i^{(k)} \sum_{\tau=1}^i A^{i-\tau} \xi_{t-dk+\tau k}^{(k)} + \sum_{\tau=1}^d A^{d-\tau} \xi_{t-dk+\tau k}^{(k)} \right\| \quad (20)$$

$$\leq k2^d(d-1)C_A C_\xi + kdC_A C_\xi \leq kd2^d C_A C_\xi \quad (21)$$

so that we obtain $x_t = -\sum_{i=1}^{d-1} a_i^{(k)} x_{t-k(d-i)} + v^{(k)}$ with $\sum_{i=0}^d |a_i^{(k)}| \leq 2^d$ and $\|v^{(k)}\| \leq kd2^d C_A C_\xi$. We then pick $k = 1, 2, \dots, \min \left\{ \left\lfloor \frac{|w^\top x_t|}{2d2^d C_A C_\xi} \right\rfloor, \left\lfloor \frac{t}{d} \right\rfloor \right\}$. For each choice of k , we know by design that $|w^\top x_t| \geq 2\|v^{(k)}\|_2$, and therefore there must exist an x in the sequence x_{t-k}, \dots, x_{t-kd} such that $|w^\top x| \geq \frac{|w^\top x_t|}{2^{d+1}}$. In this way, we are able to collect in total $L = \min \left\{ \left\lfloor \frac{|w^\top x_t|}{2d2^d C_A C_\xi} \right\rfloor, \left\lfloor \frac{t}{d} \right\rfloor \right\}$ many such x 's. To finish the argument, we note that out of the L collected x 's, there are at least $\lfloor \frac{L}{d} \rfloor$ distinct ones, since one x can appear in at most d different sequences. \blacksquare

Plugging in the conclusion of Lemma 17 into Theorem 12, we can obtain the following theorem.

Theorem 18 *Assume Assumptions 2, and furthermore suppose A is diagonalizable and $u_t = 0$ for each t . Then online least squares (Algorithm 1) with $\mu = T^{\frac{3}{4}}$ achieves regret*

$$R_T(A, B) \leq T^{\frac{3}{4}} \exp(O(d)) \text{poly}(C, R, \ln T).$$

We omit the details, as we will prove the stronger Theorem 3.

C.2. State is a small ℓ_1 combination of previous states (poly(d) version)

We will use the following notation to keep track of the growth of A^k .

Definition 19 *Given a matrix $A \in \mathbb{R}^{d \times d}$, define $f_A(k) := \|A^k\|$.*

Given a set $K \subseteq \mathbb{R}^d$, define $f_{A,K}(k) := \max_{\xi \in K} \|A^k \xi\|$.

For a rougher bound, note we can bound $f_{A,K}(k) \leq f_A(k) \max_{\xi \in K} \|\xi\|$, but keeping $f_{A,K}$ separate allows us to separately bound the dependence on the size of the starting state and on the perturbations.

The first step of the proof is to show the following. Denote the L^2 ball of radius r in d dimensions by $B_r^d = \{\|x\|_2 \leq r : x \in \mathbb{R}^d\}$.

Lemma 20 *Suppose that $x_0 \in K_0 \subset \mathbb{R}^d$ and $x_t = Ax_{t-1} + \xi_t$ with $A \in \mathbb{R}^{d \times d}$ and $\xi_t \in K$ for $1 \leq t \leq T$. Suppose $M \geq 1$ and $\|x_t\| \leq M$ for $0 \leq t \leq T$. Then there exist $a_t \in \mathbb{R}$ and $v \in \mathbb{R}^d$ such that*

$$x_T = \sum_{t=0}^{T-1} a_t x_t + v \quad (22)$$

and letting $k' = \lfloor d \log_2(M) \rfloor$,

$$\sum_{t=0}^{T-1} |a_t| \leq \frac{2}{\ln 2} d \quad (23)$$

$$\|v\|_2 \leq \max_{0 \leq k \leq k'} f_{A, K_0 \cup B_{\frac{2}{\ln 2} d}}(k) + \left(\sum_{k=0}^{k'-1} f_{A, K}(k) \right) \left(\frac{2}{\ln 2} d + 1 \right). \quad (24)$$

We emphasize that we will only need the existence of a_t , x_t , and v ; the algorithm will not need to compute the linear decomposition. We mention that the related notion of volumetric spanners based on the L^2 norm that has been applied to online convex optimization (Hazan and Karnin, 2016).

Note that we can bound $f(k)$ in terms of the size of the Jordan blocks (Lemma 24). To prove Lemma 20, we will use two lemmas (Lemma 22 and Lemma 23).

C.2.1. BOUNDING I_t

Define the convex set spanned by a set $S \subset \mathbb{R}^d$ by

$$\Delta(S) := \left\{ \sum_{u \in S} a_u u : a_u \in \mathbb{R}, \sum_{u \in S} |a_u| \leq 1, a_u = 0 \text{ for all but finitely many } u \right\}. \quad (25)$$

Suppose that S is compact and spans \mathbb{R}^d . The centrally symmetric convex set $\Delta(S)$ defines the norm (called the Minkowski functional of $\Delta(S)$)

$$\|x\|_{\Delta(S)} = \inf_{\substack{x = \sum_{u \in S} a_u u \\ a_u = 0 \text{ f.a.b.f.m. } u}} \sum_{u \in S} |a_u|, \quad (26)$$

where f.a.b.f.m. abbreviates ‘‘for all but finitely many.’’ First we note that we can in fact replace the inf with the min over linear combinations of $d + 1$ terms.

Proposition 21 *Suppose $S \subset \mathbb{R}^d$ is compact and spans \mathbb{R}^d . The following hold:*

$$\Delta(S) = \left\{ \sum_{i=1}^{d+1} a_i u_i : a_i \in \mathbb{R}, u_i \in S, \sum_{i=1}^{d+1} |a_i| \leq 1 \right\}$$

$$\|x\|_{\Delta(S)} = \min_{x = \sum_{i=1}^{d+1} a_i u_i, u_i \in S} \sum_{i=1}^{d+1} |a_i|.$$

Proof If $v \in \Delta(S)$ has a representation $v = \sum_{u \in S} a_u u$ in the form of (25), then by Carathéodory's Theorem on $\frac{v}{\sum_{u \in S} |a_u|}$, it can be written also as $v = \sum_{i=1}^{d+1} b_i u_i$ for $\sum_{i=1}^{d+1} |b_i| \leq \sum_{u \in S} |a_u|$, $u_i \in S$.

Next, we claim that $\Delta(S)$ is closed. To see this, note that $\Delta(S)$ is the image of the compact set $S^{d+1} \times \{(a_i)_{i=1}^{d+1} \in \mathbb{R}^{d+1} : \sum_{i=1}^{d+1} |a_i| \leq 1\}$ under the continuous map

$$((a_i)_{i=1}^{d+1}, (u_i)_{i=1}^{d+1}) \mapsto \sum_{i=1}^{d+1} a_i u_i.$$

The image of a compact set is compact, and a compact set in \mathbb{R}^d is closed and bounded, by the Heine-Borel Theorem.

Finally, this implies $\|x\|_{\Delta(S)} = \inf_{x = \sum_{u \in S} a_u u} \sum_{u \in S} |a_u| = \min_{x = \sum_{i=1}^{d+1} a_i u_i, u_i \in S} \sum_{i=1}^{d+1} |a_i|$.

■

Define

$$\Delta'(S) = \Delta(S \cup \{v : \|v\|_2 \leq 1\}) = \left\{ \left(\sum_{u \in S} a_u u \right) + v : a_u \in \mathbb{R}, a_u = 0 \text{ f.a.b.f.m. } u, v \in \mathbb{R}^d, \sum_{u \in S} |a_u| + \|v\|_2 \leq 1 \right\} \quad (27)$$

(note that the L^2 norm is used with v). Then

$$\|x\|_{\Delta'(S)} = \min_{y = (\sum_{u \in S} a_u u) + v} \left(\sum_{u \in S} |a_u| + \|v\|_2 \right). \quad (28)$$

Now define

$$\Delta_t = \Delta'(\{x_s : 0 \leq s \leq t\}) \quad (29)$$

for $-1 \leq t \leq T$. We need to keep track of the number of times when x_t is not a small linear combination of previous x_s 's, so let

$$I_t = \left\{ 0 \leq s \leq t : \|x_s\|_{\Delta_{s-1}} \geq \frac{2}{\ln 2} d \right\}. \quad (30)$$

We first show that if $s \in I_t$, then there is a large increase in volume between Δ_{s-1} and Δ_s . Then, we use a bound on the volume of Δ_T to bound $|I_T|$.

Lemma 22 *Let w be such that $\|w\|_{\Delta(S)} \geq 2Cd$. Then*

$$\text{Vol}(\Delta(S \cup \{w\})) \geq (1 + 2e^{-1/C}) \text{Vol}(\Delta(S)). \quad (31)$$

Proof Consider the set $\Delta(S)$ slightly shrunk, and translated along the direction of w :

$$\pm \frac{1}{Cd} w + \left(1 - \frac{1}{Cd}\right) \Delta(S) \subseteq \Delta(S \cup \{w\}). \quad (32)$$

Now for $u \in \Delta(S)$, $\left\| \pm \frac{1}{Cd} w + \left(1 - \frac{1}{Cd}\right) u \right\|_{\Delta(S)} \geq \frac{1}{Cd} \|w\|_{\Delta(S)} - \|u\|_{\Delta(S)} > \frac{1}{Cd}(2Cd) - 1 \geq 1$. Hence, $\pm \frac{1}{Cd} w + \left(1 - \frac{1}{Cd}\right) \Delta(S)$ does not intersect $\Delta(S)$. This means that

$$\text{Vol}(\Delta(S \cup \{w\})) \geq \text{Vol}(\Delta(S)) + \text{Vol}\left(\pm \frac{1}{Cd} w + \left(1 - \frac{1}{Cd}\right) \Delta(S)\right) \quad (33)$$

$$\geq \text{Vol}(\Delta(S)) + 2\left(1 - \frac{1}{Cd}\right)^d \text{Vol}(\Delta(S)) \quad (34)$$

$$\geq (1 + 2e^{-1/C}) \text{Vol}(\Delta(S)). \quad (35)$$

■

Lemma 23 *Let $M = \max_{0 \leq t \leq T} \|x_t\|$. Then $|I_T| \leq d \log_2 M$.*

Proof If $|I_t| > |I_{t-1}|$, then by definition $\|x_t\|_{\Delta_{t-1}} \geq \frac{2}{\ln 2} d$. By Lemma 22, $\text{Vol}(\Delta_t) \geq 2 \text{Vol}(\Delta_{t-1})$. Hence

$$2^{|I_T|} \text{Vol}(B_1^d) = 2^{|I_T|} \text{Vol}(\Delta_{-1}) \leq \text{Vol}(\Delta_T) \leq M^d \text{Vol}(B_1^d). \quad (36)$$

Taking logarithms gives $|I_T| \leq d \log_2 M$. ■

C.2.2. BOUNDING $\|x_t\|_{\Delta_{t-1}}$

Proof [Proof of Lemma 20] First we show that if x_t is a linear combination of previous x_s 's with small coefficients, then so is x_{t+k} for small k . Suppose that

$$x_t = \sum_{s=0}^{t-1} a_s x_s + v. \quad (37)$$

We claim by induction that for any $k \geq 0$,

$$x_{t+k} = \sum_{s=0}^{t-1} a_s x_{s+k} + A^k v + \sum_{j=1}^k A^{k-j} \left(\sum_{s=0}^{t-1} -a_s \xi_{s+j} + \xi_{t+j} \right). \quad (38)$$

Indeed, if this holds for k , then

$$x_{t+k+1} = A x_{t+k} + \xi_{t+k+1} \quad (39)$$

$$= A \left(\sum_{s=0}^{t-1} a_s x_{s+k} + A^k v + \sum_{j=1}^k A^{k-j} \left(\sum_{s=0}^{t-1} -a_s \xi_{s+j} + \xi_{t+j} \right) \right) + \xi_{t+k+1} \quad (40)$$

$$= \left(\sum_{s=0}^{t-1} a_s A x_{s+k} + A^{k+1} v + \sum_{j=1}^k A^{k+1-j} \left(\sum_{s=0}^{t-1} -a_s \xi_{s+j} + \xi_{t+j} \right) \right) + \xi_{t+k+1} \quad (41)$$

$$= \left(\sum_{s=0}^{t-1} a_s (x_{s+k+1} - \xi_{s+k+1}) + A^{k+1} v + \sum_{j=1}^k A^{k+1-j} \left(\sum_{s=0}^{t-1} -a_s \xi_{s+j} + \xi_{t+j} \right) \right) + \xi_{t+k+1} \quad (42)$$

$$= \sum_{s=0}^{t-1} a_s x_{s+k+1} + A^{k+1} v + \sum_{j=1}^{k+1} A^{k+1-j} \left(\sum_{s=0}^{t-1} -a_s \xi_{s+j+1} + \xi_{t+j+1} \right). \quad (43)$$

Now we consider the sizes of the coefficients in (38). If $\sum_{s=0}^{t-1} |a_s| = L_a$ and $\|v\|_2 = L_v$, then (38) expresses x_{t+k+1} as a linear combination with

$$\sum_{s=0}^{t-1} |a_s| = L_a \quad (44)$$

$$\left\| A^k v + \sum_{j=1}^k A^{k-j} \left(\sum_{s=0}^{t-1} -a_s \xi_{s+j} + \xi_{t+j} \right) \right\| \leq \|A^k v\| + \left(\sum_{s=0}^{k-1} f_{A,K}(s) \right) \left(\left(\sum_{s=0}^{t-1} |a_s| \right) + 1 \right) \quad (45)$$

$$\leq \|A^k v\| + \left(\sum_{s=0}^{k-1} f_{A,K}(s) \right) (L_a + 1). \quad (46)$$

By Lemma 23, $|I_T| \leq d \log_2 M$, so for $1 \leq t \leq T$, there exists $k \leq |I_T| \leq k' = \lfloor d \log_2(M) \rfloor$ such that either

- $t - k \notin I_T$, so $\|x_{t-k}\|_{\Delta_{t-k-1}} \leq \frac{2}{\ln 2} d$, or
- $t - k = 0$, so $x_{t-k} = x_0 \in K_0$.

In either case, we can write $x_{t-k} = \sum_{s=0}^{t-k-1} a_s x_s + v$ with $\sum_{s=0}^{t-k-1} |a_s| \leq L_a := \frac{2}{\ln 2} d$ and $v \in K_0 \cup B_{\frac{2}{\ln 2} d}^d$. Then by (44) and (46), we can write x_t as

$$x_t = \sum_{s=0}^{t-1} a'_s x_s + v' \quad (47)$$

with

$$\sum_{s=0}^{t-1} |a'_s| \leq \frac{2}{\ln 2} d \quad (48)$$

$$\|v\|_2 \leq \|A^k v\| + \left(\sum_{s=0}^{k'-1} f_{A,K}(s) \right) (L_a + 1) \quad (49)$$

$$\leq \max_{0 \leq k \leq k'} f_{A, K_0 \cup B_{\frac{2}{\ln 2} d}^d}(k) + \left(\sum_{k=0}^{k'-1} f_{A,K}(k) \right) \left(\frac{2}{\ln 2} d + 1 \right). \quad (50)$$

■

C.2.3. BOUND IN TERMS OF JORDAN FORM

Lemma 24 *Suppose that A has spectral radius ≤ 1 and $A = SJS^{-1}$ with $\|S\|_2 \|S^{-1}\|_2 \leq C_A$ and J has Jordan blocks of rank $\leq r$. Then*

$$\|A^k\|_2 \leq \binom{k+r}{r-1} C_A \leq (k+1)^{r-1} C_A \quad (51)$$

$$\sum_{s=0}^k \|A^s\|_2 \leq \binom{k+r+1}{r} C_A \leq (k+1)^r C_A \quad (52)$$

Proof Let $J_{\lambda,r} = \begin{pmatrix} \lambda & 1 & & \\ & \lambda & \ddots & \\ & & \ddots & 1 \\ & & & \lambda \end{pmatrix} \in \mathbb{R}^{r \times r}$. Then entrywise, $J_{\lambda,r}^k$ has absolute value at most $J_{1,r}^k = \begin{pmatrix} 1 & \binom{k}{1} & \cdots & \binom{k}{r-1} \\ & 1 & \ddots & \vdots \\ & & \ddots & \binom{k}{1} \\ & & & 1 \end{pmatrix}$. We have for $k \geq 1$ that

$$\|J_{\lambda,r}^k\|_2 \leq \|J_{\lambda,r}^k\|_{1,1} \leq \|J_{1,r}^k\|_{1,1} \quad (53)$$

$$\leq \sum_{q=0}^{r-1} \sum_{j=0}^q \binom{k}{j} \leq \sum_{q=0}^{r-1} \sum_{j=0}^q \binom{k+j-1}{j} \leq \sum_{q=0}^{r-1} \binom{k+q}{q} \quad (54)$$

$$\leq \binom{k+r}{r-1} \leq \frac{(k+1)^r}{(r-1)!} \cdot \frac{(k+r) \cdots (k+2)}{(k+1)^r} \leq (k+1)^{r-1} \quad (55)$$

By decomposing into Jordan blocks, we find that

$$\|A^k\|_2 \leq \|S\|_2 \|J_{\lambda,r}^k\|_2 \|S^{-1}\|_2 \leq C_A \|J_{\lambda,r}^k\|_2. \quad (56)$$

Together (55) and (56) give (51). Summing over k gives (52). \blacksquare

C.3. Concluding the states are anomaly-free

We need the following number-theoretic lemma.

Lemma 25 (Nathanson 2013, Theorem 6.3) *There exists a constant $C > 0$ such that for all $x \geq 2$,*

$$\prod_{1 < p \leq x, p \text{ prime}} p \geq x^{Cx}.$$

The next lemma, Lemma 26, is technical to state, so we first explain the idea. Recall that we want the sequence (x_t) to be anomaly-free in order to apply the bound on regret in Theorem 12, i.e., we want a lot of x_s 's such that $|w^\top x_s|$ is large relative to $w^\top x_T$. Lemma 20 together with Lemma 16 only gives a single large $|w^\top x_s|$, so the idea is to apply the lemmas to the dynamical systems $x_t \approx A^p x_{t-p}$ for different values of p . We choose prime p 's and use Lemma 25 to control overcounting.

We separate out the core argument from the numerical bounds, so Lemma 26 is stated in terms of the functions L_v and Q . In the lemma, $L_v(p)$ tracks how large the residual term in (24) can get. This is the term left over when writing x_T as a linear combination of previous terms. We will instantiate $L_v(p) = O((pk')^r(d+m))$ in the proof of Theorem 3 (ignoring constants related to the LDS). In order to apply Lemma 16, we would like to only consider p small enough so that $2L_v(p) \leq |w^\top x_T|$, or $(pk')^r(d+m) = O(|w^\top x_T|)$. Roughly, this is true for $p = O\left(\frac{1}{k'} \left(\frac{|w^\top x_T|}{d+m}\right)^{1/r}\right)$. This RHS is approximately the Q appearing in the lemma, $Q(x) \approx \frac{1}{k'} \left(\frac{x}{d+m}\right)^{1/r}$. For each of these primes $p \leq Q(|w^\top x_T|)$, we obtain some x_s such that $|w^\top x_s|$ is large. This gives almost $Q(|w^\top x_T|)$ many large x_s 's. We lose logarithmic factors, since we restrict to primes and account for overcounting.

The key dependence to note here is that on r , the size of the largest Jordan block: the residual $L_v(p)$ has r th power dependence, so the number of large $|w^\top x_s|$'s has $\frac{1}{r}$ -power dependence. Any constant power gives a sublinear regret bound in Theorem 12, with a smaller value of $\alpha = \frac{1}{r}$ giving a worse regret bound.

Lemma 26 *Let $T > 0$. Suppose that $\|x_0\| \leq C_0$ and $x_t = Ax_{t-1} + \xi_t$ with $\xi_t \in K$ for $1 \leq t \leq T$. Let $k' = \lfloor d \log_2(M) \rfloor$ and let*

$$L_v(p) = \max_{0 \leq k \leq p^{(k'+1)-1}} f_A(k) \max \left\{ \frac{2}{\ln 2} d, C_0 \right\} + \left(\sum_{k=0}^{p^{(k'+1)-2}} f_{A,K}(k) \right) \left(\frac{2}{\ln 2} d + 2 \right) \quad (57)$$

Suppose the following hold.

- (Bound on x_t) $M \geq 1$, $\|x_t\| \leq M$ for $1 \leq t \leq T$.
- ($|w^\top x_T|$ is large enough) $c \geq \frac{4}{\ln 2} d C_0$, and $|w^\top x_T| > c$.
- (Q gives the range of p that are “good” for Lemma 16) $Q(x)$ is a function such that whenever $x > c$, then $2 \leq Q(x)$ and for all $p \leq Q(x)$, we have $x \geq 2L_v(p)$.

Then there are at least $\Omega\left(\frac{Q(|w^\top x_T|)}{\ln T}\right)$ values of s , $0 \leq s < t$, such that

$$|w^\top x_s| \geq \frac{\ln 2}{4d} |w^\top x_T|. \quad (58)$$

Proof For $p \leq T$, we apply Lemma 16 with the bounds from Lemma 20 with the sequence $x_{T \bmod p}, \dots, x_{T-p}, x_T$. This satisfies the conditions of Lemma 20 with

$$x_t = A' x_{t-p} + \xi_t^{(p)} \quad (59)$$

$$A' = A^p \quad (60)$$

$$x_{T \bmod p} \in K'_0 := A^{T \bmod p} B_{C_0}^d + A^{(T \bmod p)-1} K + \dots + K \quad (61)$$

$$\xi_t^{(p)} = \sum_{j=1}^p A^{p-j} \xi_{t-p+j} \in K' := \sum_{j=1}^p A^{p-j} K. \quad (62)$$

Using this, we write $f_{A', K'_0 \cup B_{\frac{2}{\ln 2} d}}(k)$ in terms of f_A and $f_{A,K}$. We upper bound by taking the maximum inside the sum, to get the following.

$$\max_{0 \leq k \leq k'} f_{A', K'_0 \cup B_{\frac{2}{\ln 2} d}}(k) \quad (63)$$

$$= \max \left\{ \max_{\substack{0 \leq k \leq k' \\ v \in B_{C_0}^d, \xi_1, \dots, \xi_{T \bmod p} \in K}} \left(A^{(T \bmod p)+kp} v + \sum_{j=1}^{T \bmod p} A^{(T \bmod p)-j+kp} \xi_j \right), \quad (64) \right.$$

$$\left. \max_{0 \leq k \leq k'} \left\| A^{kp} \right\|_2 \frac{2}{\ln 2} d \right\}$$

$$\leq \max_{0 \leq k \leq p^{(k'+1)-1}} f_A(k) \max \left\{ \frac{2}{\ln 2} d, C_0 \right\} + \sum_{k=0}^{p^{(k'+1)-2}} f_{A,K}(k). \quad (65)$$

Similarly,

$$\sum_{k=0}^{k'-1} f_{A',K'}(k) = \sum_{k=0}^{k'-1} \max_{\xi_0, \dots, \xi_{p-1} \in K} \left\| \sum_{j=1}^p A^{pk+p-j} \xi_j \right\| \leq \sum_{k=0}^{k'-1} \max_{\xi \in K} \|A^k \xi\| = \sum_{k=0}^{p k'-1} f_{A,K}(k). \quad (66)$$

Then for the sequence $x_{T \bmod p}, \dots, x_T$, the L_v in Lemma 20 equals

$$\max_{0 \leq k \leq k'} f_{A',K'_0 \cup B_{\frac{2}{\ln 2}d}}(k) + \left(\sum_{k=0}^{k'-1} f_{A',K'}(k) \right) \left(\frac{2}{\ln 2}d + 1 \right) \quad (67)$$

$$\leq \max_{0 \leq k \leq p(k'+1)-1} f_A(k) \max \left\{ \frac{2}{\ln 2}d, C_0 \right\} + \left(\sum_{k=0}^{p(k'+1)-2} f_{A,K}(k) \right) \left(\frac{2}{\ln 2}d + 2 \right) =: L_v(p) \quad (68)$$

using (65) and (66). Apply Lemma 16 to get that there exists $k \in \mathbb{N}$ such that

$$|w^\top x_{T-kp}| \geq \frac{|w^\top x_T| - L_v(p)}{\frac{2}{\ln 2}d}. \quad (69)$$

When $|w^\top x_T| > c$, and $p \leq Q(|w^\top x_T|)$, by assumption $|w^\top x_T| \geq 2L_v(p)$ and hence $\frac{|w^\top x_T| - L_v(p)}{\frac{2}{\ln 2}d} \geq \frac{|w^\top x_T|}{\frac{4}{\ln 2}d}$. Then equation (69) becomes: there exists $k \in \mathbb{N}$ such that

$$|w^\top x_{T-kp}| \geq \frac{\ln 2}{4d} |w^\top x_T| \quad (70)$$

Put another way, there exists s such that $|w^\top x_s| \geq \frac{\ln 2}{4d} |w^\top x_T|$ and p divides $T - s$. Let $Q = Q(|w^\top x_T|)$. This means that $\prod_{|w^\top x_s| \geq \frac{\ln 2}{4d} |w^\top x_T|} (T - s)$ has to be divisible by all primes $p \leq \min\{Q, T\} = Q$. (Note $Q \leq T$ because otherwise, (70) would hold for $p = T < Q$ and we have $C_0 \geq |w^\top x_0| > \frac{\ln 2}{4d} |w^\top x_T| \geq \frac{\ln 2}{4d} c$, contradiction.) Let $N = \left| \left\{ s : |w^\top x_s| \geq \frac{\ln 2}{4d} |w^\top x_T| \right\} \right|$. By Lemma 25, for some $C' > 0$,

$$T^N \geq \prod_{|w^\top x_s| \geq \frac{\ln 2}{4d} |w^\top x_T|} (T - s) \geq \prod_{\text{prime } p \leq Q} p \geq Q^{C'Q}. \quad (71)$$

Thus

$$N \geq \frac{C'Q \ln Q}{\ln T} \geq \Omega \left(\frac{Q}{\ln T} \right). \quad (72)$$

■

C.4. Finishing the proof

Lemma 27 *Assume Assumption 2. Then*

$$\max_{0 \leq k \leq T} \|x_k\| \leq (T+1)^{r-1} C_A C_0 + T^r C_A (C_B C_u + C_\xi). \quad (73)$$

Proof Using Lemma 24 we have

$$\begin{aligned} x_t &= A^t v + \sum_{k=0}^{t-1} A^k B u_{t-1-k} + \sum_{k=0}^{t-1} A^k \xi_{t-k} \\ \max_{0 \leq k \leq T} \|x_k\| &\leq \max_{0 \leq k \leq T} \|A^k\| C_0 + \sum_{k=0}^{T-1} \|A^k B\| C_u + \sum_{k=0}^{T-1} \|A^k\| C_\xi \\ &\leq (T+1)^{r-1} C_A C_0 + C_A C_B C_u T^r + C_A C_\xi T^r. \end{aligned}$$

■

Proof [Proof of Theorem 3] Let $A' = \begin{pmatrix} A & B \\ 0 & B \end{pmatrix}$. We apply Lemma 26 to the system

$$\begin{pmatrix} x_t \\ u_t \end{pmatrix} = A' \begin{pmatrix} x_{t-1} \\ u_{t-1} \end{pmatrix} + \begin{pmatrix} \xi_t \\ u_t \end{pmatrix} \quad (74)$$

with $\begin{pmatrix} \xi_t \\ u_t \end{pmatrix} \in K := B_{C_\xi}^d \oplus B_{C_u}^d$ for $1 \leq t \leq T$. Let $k' = \lfloor (d+m) \log_2(M) \rfloor$ and $C = C_A(C_\xi + C_B C_u)$. Note that

$$\begin{aligned} f_{A',K}(k) &= \max_{\begin{pmatrix} \xi \\ u \end{pmatrix} \in K} \left(\|A^k \xi\|_2 + \begin{cases} \|A^{k-1} B u\|_2, & k \geq 1 \\ C_u, & k = 0 \end{cases} \right) \\ \sum_{k=0}^{p(k'+1)-2} f_{A',K}(k) &\leq \sum_{k=0}^{p(k'+1)-2} (f_A(k) C_\xi + f_A(k) C_B C_u) + C_u \end{aligned} \quad (75)$$

Then using Lemma 24, $L_v(p)$ in Lemma 26 equals

$$\begin{aligned} L_v(p) &= \max_{0 \leq k \leq p(k'+1)-1} f_{A'}(k) \max \left\{ \frac{2}{\ln 2} (d+m), C_0 \right\} + \left(\sum_{k=0}^{p(k'+1)-2} f_{A',K}(k) \right) \left(\frac{2}{\ln 2} (d+m) + 2 \right) \\ &\leq \max_{0 \leq k \leq p(k'+1)-1} f_{A'}(k) \max \left\{ \frac{2}{\ln 2} (d+m), C_0 \right\} \\ &\quad + \left(\sum_{k=0}^{p(k'+1)-2} (f_A(k) C_\xi + f_A(k) C_B C_u) + C_u \right) \left(\frac{2}{\ln 2} (d+m) + 2 \right) \\ &\leq (p(k'+1))^{r-1} \max \left\{ \frac{2}{\ln 2} (d+m), C_0 \right\} + [C(p(k'+1))^r + C_u] \left(\frac{2}{\ln 2} (d+m) + 2 \right) \\ &\leq p^r (k'+1)^r \left(\frac{2}{\ln 2} (d+m)(C+1) + C_0 + 2 \right) + C_u \left(\frac{2}{\ln 2} (d+m) + 2 \right). \end{aligned} \quad (76)$$

Let

$$\begin{aligned} Q(x) &:= \left(\frac{x - 2C_u \left(\frac{2}{\ln 2} (d+m) + 2 \right)}{2C(k'+1)^r \left(\frac{2}{\ln 2} (d+m)(C+1) + C_0 + 2 \right)} \right)^{\frac{1}{r}} \\ c &:= \max \left\{ 2^{r+1} C(k'+1)^r \left(\frac{2}{\ln 2} (d+m)(C+1) + C_0 + 2 \right) + 4C_u \left(\frac{2}{\ln 2} (d+m) + 2 \right), \frac{4dC_0}{\ln 2} \right\}. \end{aligned}$$

It is easy to check that when $x \geq c$, then $2 \leq Q(x)$ and for all $p \leq Q(x)$, $x \geq 2L_v(p)$. Moreover, for $x \geq c$, we have $x \geq 4C_u \left(\frac{2}{\ln 2}(d+m) + 2 \right)$, so $Q(x) \geq \left(\frac{x}{4C(k'+1)^r \left(\frac{2}{\ln 2}(d+m)(C+1) + C_0 + 2 \right)} \right)^{\frac{1}{r}}$. By Lemma 26, the states are (c, c_1, c_2, α) -anomaly-free, with

$$\begin{aligned} \alpha &= \frac{1}{r} \\ c &= \max \left\{ 2^{r+1} C(k'+1)^r \left(\frac{2}{\ln 2}(d+m)(C+1) + C_0 + 2 \right) + 4C_u \left(\frac{2}{\ln 2}(d+m) + 2 \right), \frac{4dC_0}{\ln 2} \right\} \\ c_1 &= \Theta \left(\frac{1}{C^{\frac{1}{r}}(k'+1) \left(\frac{2}{\ln 2}(d+m)(C+1) + C_0 \right)^{\frac{1}{r}} \ln T} \right) \\ c_2 &= \frac{\ln 2}{4d}, \end{aligned}$$

satisfying the hypothesis of Theorem 12. By Theorem 12, the regret is

$$\begin{aligned} R_T \leq & \mu R^2 + O \left((d+m)^3 \max \left\{ \frac{C_\xi c^2 T}{\mu} + \|(A, B)\|_2 c, \frac{C_\xi}{c_2^2}, \frac{1}{c_1^{\frac{1}{\alpha+2}} c_2^{\frac{2}{\alpha+2}}} \left(\|(A, B)\|_2 \mu^{\frac{1}{\alpha+2}} + \frac{C_\xi T^{\frac{1}{2}}}{\mu^{\frac{1}{2} - \frac{1}{\alpha+2}}} \right) \right\} \right. \\ & \left. \cdot \ln \left(1 + \frac{TM^2}{d+m} \right) \right). \end{aligned}$$

Plugging in $\mu = T^{\frac{2r+1}{2r+2}}$, and using the bound $M \leq (T+1)^{r-1} C_A C_0 + T^r C_A (C_B C_u + C_\xi)$ from Lemma 27, we obtain the theorem. Note that the only terms with an r th power are the two terms involving c . The dependence on T of the larger such term is $\frac{T}{\mu} = T^{\frac{1}{2r+2}}$, hence the additive term in (4).

We now simplify the bound when all the constants are $O(1)$ and $\|(A, B)\|_F = O(d+m)$. We have $c_1 = \Omega \left(\frac{1}{(d+m)^r \ln T (d+m)^{\frac{1}{r}} \ln T} \right)$ and $c_2 = \Omega \left(\frac{1}{d} \right)$. Using (8) in Theorem 12, with $\mu = \frac{(d+m)^{\frac{\alpha+2}{\alpha+1}} T^{\frac{\alpha+2}{2(\alpha+1)}}}{c_1^{\frac{1}{\alpha+1}} c_2^{\frac{2}{\alpha+1}}}$, we obtain the bound (as $T \rightarrow \infty$)

$$\begin{aligned} R_T &= O \left(\frac{(d+m)^2 (d+m)^{\frac{r}{r+1}} T^{\frac{2r+1}{2r+2}} r \ln T}{c_1^{\frac{r}{r+1}} c_2^{\frac{2r}{r+1}}} \right) \\ &= O \left((d+m)^{2+\frac{r}{r+1}} T^{\frac{2r+1}{2r+2}} r \ln T \cdot \left((d+m)^{\frac{r+1}{r}} r (\ln T)^2 \right)^{\frac{r}{r+1}} d^{\frac{2r}{r+1}} \right) \\ &= O \left((d+m)^{2+\frac{r}{r+1}} d^{1+\frac{r-1}{r+1}} (d+m) T^{\frac{2r+1}{2r+2}} r^2 (\ln T)^3 \right) \\ &= O \left((d+m)^4 d^2 r^2 T^{\frac{2r+1}{2r+2}} (\ln T)^3 \right). \end{aligned} \tag{77}$$

Note that Corollary 4 follows from plugging $r = 1$ into (77). ■

Appendix D. Proof of Theorem 5 (fully observed system, stochastic noise)

Lemma 28 *Let $v_t, w_t \in \mathbb{R}^d$ for $0 \leq t \leq T - 1$. Suppose that the following hold.*

- For any $0 \leq t \leq T - 1$, there exist $a_s \in \mathbb{R}$ such that

$$v_t = \sum_{s=0}^{t-1} a_s v_s + v, \quad \sqrt{\sum_{s=0}^{t-1} a_s^2} \leq L_a, \quad \|v\|_2 \leq L_v. \quad (78)$$

- $w_t = Av_t + \varepsilon_{t+1} + \xi_{t+1}$ where $\|A\|_2 \leq R$, $\|\varepsilon_t\| \leq C_\varepsilon$, and ξ_t , $1 \leq t \leq T$ are random variables such that $\xi_{t+1} | \xi_1, \dots, \xi_t$ is mean 0 and C_ξ -subgaussian for $0 \leq t \leq T - 1$.
- $\|v_t\| \leq g(t)$ for $0 \leq t \leq T - 1$.

Let $L' > 0$, and let

$$\varepsilon_{\text{net}} = \frac{1}{\sum_{s=0}^{T-1} g(s)} \quad (79)$$

$$R_z = \mu^{-\frac{1}{2}} \left(L_a^2 + \frac{L_v^2}{\mu} \right)^{\frac{1}{2}} d^{\frac{1}{2}} \quad (80)$$

$$S_b = \sqrt{2} \left(\left(L_a^2 + \frac{L_v^2}{\mu} \right) d^2 + 1 \right)^{\frac{1}{2}} \quad (81)$$

$$L = \sqrt{2} C_\xi S_b \left(d \ln \left(1 + \frac{2R_z}{\varepsilon_{\text{net}}} \right) + \ln \left(\frac{T}{\delta} \right) \right)^{\frac{1}{2}}. \quad (82)$$

Then

$$\mathbb{P} \left(\max_{0 \leq t \leq T-1} \|A_t v_t - v_{t+1}\|_2 \leq L + (\mu^{\frac{1}{2}} R + C_\varepsilon d^{\frac{1}{2}} \sqrt{T}) \left(L_a^2 + \frac{L_v^2}{\mu} \right)^{\frac{1}{2}} d^{\frac{1}{2}} + 2C_\xi + C_\varepsilon \right) \geq 1 - \delta. \quad (83)$$

The reason for the choice of values of ε_{net} , R_z , S_b , and L can be seen in (103), (97), (99), and (106), respectively.

We note several lemmas we will need.

Lemma 29 *Let $v_s \in \mathbb{R}^d$ for $1 \leq s \leq t$. Suppose that $\sum_{s=1}^t v_s v_s^\top \preceq \Sigma$. Then*

$$\sum_{s=1}^t v_s^\top \Sigma^{-1} v_s \leq d. \quad (84)$$

Proof We can choose $v'_s \in \mathbb{R}^d$, $1 \leq s \leq t'$ such that $\sum_{s=1}^t v_s v_s^\top + \sum_{s=1}^{t'} v'_s v'_s{}^\top = \Sigma$. Then

$$\sum_{s=1}^t v_s^\top \Sigma^{-1} v_s + \sum_{s=1}^{t'} v'_s{}^\top \Sigma^{-1} v'_s = \text{Tr} \left(\Sigma^{-1} \left(\sum_{s=1}^t v_s v_s^\top + \sum_{s=1}^{t'} v'_s v'_s{}^\top \right) \right) = \text{Tr}(I_d) = d \quad (85)$$

$$\implies \sum_{s=1}^t v_s^\top \Sigma^{-1} v_s = d - \sum_{s=1}^{t'} v'_s{}^\top \Sigma^{-1} v'_s \leq d. \quad (86)$$

■

Lemma 30 (Vershynin 2010, Lemma 5.2) *There is an ε -net of B_1^d of size $(1 + \frac{2}{\varepsilon})^d$.*

Lemma 31 (Generalization of Azuma's Inequality (Simchowitz et al., 2018, Lemma 4.2)) *Let $\{\mathcal{F}_t\}_{t \geq 0}$ be a filtration, and $\{Z_t\}_{t \geq 1}$ and $\{W_t\}_{t \geq 1}$ be real-valued processes adapted to \mathcal{F}_t and \mathcal{F}_{t+1} respectively. Moreover, assume $W_t | \mathcal{F}_t$ is mean 0 and σ^2 -sub-Gaussian. Then for any positive real numbers L and β ,*

$$\mathbb{P} \left[\left\{ \sum_{t=1}^T W_t Z_t \geq L \right\} \cap \left\{ \sum_{t=1}^T Z_t^2 \leq \beta \right\} \right] \leq \exp \left(-\frac{L^2}{2\sigma^2\beta} \right)$$

Proof [Proof of Lemma 28] Let

$$\Sigma_t = \mu I_d + \sum_{s=0}^{t-1} v_s v_s^\top. \quad (87)$$

By Lemma 13,

$$A_t v_t - w_t = \sum_{s=0}^{t-1} (\varepsilon_{s+1} + \xi_{s+1}) v_s^\top \Sigma_t^{-1} v_t - \mu \Sigma_t^{-1} v_t - \varepsilon_{t+1} - \xi_{t+1}. \quad (88)$$

Bounding $\sum_{s=0}^{t-1} \xi_{s+1} v_s^\top \Sigma_t^{-1} v_t$. Let $b_s = v_s^\top \Sigma_t^{-1} v_t$.

To use Azuma's inequality, we need to bound the following. Using Cauchy-Schwarz ($\langle v, w \rangle^2 \leq \|v\|^2 \|w\|^2$),

$$\sum_{s=0}^{t-1} b_s^2 \leq \sum_{s=0}^{t-1} \left\| v_s^\top \Sigma_t^{-\frac{1}{2}} \right\|_2^2 \left\| \Sigma_t^{-\frac{1}{2}} v_t \right\|_2^2 \quad (89)$$

$$= \left(\sum_{s=0}^{t-1} v_s^\top \Sigma_t^{-1} v_s \right) \left\| \Sigma_t^{-\frac{1}{2}} v_t \right\|_2^2 \quad (90)$$

Choosing a_s and v as in (78),

$$\left\| \Sigma_t^{-\frac{1}{2}} v_t \right\|_2^2 = \left\| \Sigma_t^{-\frac{1}{2}} \left(\sum_{s=0}^{t-1} a_s v_s + v \right) \right\|_2^2 = \left\| \sum_{s=0}^{t-1} a_s \Sigma_t^{-\frac{1}{2}} v_s + \Sigma_t^{-\frac{1}{2}} v \right\|_2^2 \quad (91)$$

$$\leq \left(\sum_{s=0}^{t-1} a_s^2 + \frac{\|v\|^2}{\mu} \right) \left(\sum_{s=0}^{t-1} \left\| \Sigma_t^{-\frac{1}{2}} v_s \right\|_2^2 + \frac{\mu}{\|v\|^2} \left\| \Sigma_t^{-\frac{1}{2}} v \right\|_2^2 \right) \quad \text{by Cauchy-Schwarz} \quad (92)$$

$$\leq \left(L_a^2 + \frac{L_v^2}{\mu} \right) \left(\sum_{s=0}^{t-1} v_s^\top \Sigma_t^{-1} v_s + \mu \frac{v^\top}{\|v\|} \Sigma_t^{-1} \frac{v}{\|v\|} \right). \quad (93)$$

Now $v_s v_s^\top + \mu \frac{v v^\top}{\|v\|^2} \preceq \Sigma_t$ satisfies the hypothesis of Lemma 29, so

$$\sum_{s=0}^{t-1} v_s^\top \Sigma_t^{-1} v_s + \mu \frac{v^\top}{\|v\|} \Sigma_t^{-1} \frac{v}{\|v\|} \leq d \quad (94)$$

and so (93) gives

$$\left\| \Sigma_t^{-\frac{1}{2}} v_t \right\|^2 \leq \left(L_a^2 + \frac{L_v^2}{\mu} \right) d. \quad (95)$$

From (90), (94), and (95), we get

$$\sum_{s=0}^{t-1} b_s^2 \leq \left(L_a^2 + \frac{L_v^2}{\mu} \right) d^2. \quad (96)$$

Let $z_t = \Sigma_t^{-1} v_t$. By (93) and (95),

$$\|z_t\| = \left\| \Sigma_t^{-\frac{1}{2}} \right\|_2 \left\| \Sigma_t^{-\frac{1}{2}} v_t \right\|_2 \leq \mu^{-\frac{1}{2}} \left(L_a^2 + \frac{L_v^2}{\mu} \right)^{\frac{1}{2}} d^{\frac{1}{2}} = R_z \quad (97)$$

so $z_t \in B_{R_z}^d$. By Lemma 30, there is an ε_{net} -net \mathcal{N} of $B_{R_z}^d$ of size $(1 + \frac{2R_z}{\varepsilon_{\text{net}}})^d$. Since $z_t \in B_{R_z}^d$, there exists $\|\Delta z\| \leq \varepsilon_{\text{net}}$ such that $z_t + \Delta z \in \mathcal{N}$. Note that

$$\sum_{s=0}^{t-1} |v_s^\top (z_t + \Delta z)|^2 \leq 2 \left(\sum_{s=0}^{t-1} |v_s^\top z_t|^2 + |v_s^\top \Delta z|^2 \right) \leq 2 \left(\left(\sum_{s=0}^{t-1} b_s^2 \right) d^2 + \sum_{s=0}^{t-1} g(s)^2 \varepsilon_{\text{net}}^2 \right) \quad (98)$$

$$\leq 2 \left(\left(L_a^2 + \frac{L_v^2}{\mu} \right) d^2 + 1 \right) =: S_b^2. \quad (99)$$

For a fixed z , consider the event $E_z = \{ \sum_{s=0}^{t-1} |y_s^\top z|^2 > S_b^2 \text{ or } \sum_{s=0}^{t-1} \xi_{s+1} y_s^\top z \leq L \}$. Then by the generalization of Azuma's inequality 31,

$$\mathbb{P}(E_z^c) \leq \mathbb{P} \left(\sum_{s=0}^{t-1} |y_s^\top z|^2 \leq S_b^2 \text{ and } \sum_{s=0}^{t-1} \xi_{s+1} y_s^\top z > L \right) \quad (100)$$

$$\leq \exp \left(-\frac{L^2}{2C_\xi^2 S_b^2} \right). \quad (101)$$

Now we use the triangle inequality to bound the sum by the maximum value on the ε_{net} -net.

$$\left\| \sum_{s=0}^{t-1} \xi_{s+1} v_s^\top z \right\| \leq \left\| \sum_{s=0}^{t-1} \xi_{s+1} v_s^\top (z + \Delta z) \right\| + \left\| \sum_{s=0}^{t-1} \xi_{s+1} v_s^\top (\Delta z) \right\| \quad (102)$$

$$\leq \max_{\substack{z' \in \mathcal{N} \\ \sum_{s=0}^{t-1} |v_s^\top z'|^2 \leq S_b^2}} \left| \sum_{s=0}^{t-1} \xi_{s+1} v_s^\top z' \right| + C_\xi \sum_{s=0}^{t-1} g(s) \varepsilon_{\text{net}} \quad (103)$$

$$= \max_{\substack{z' \in \mathcal{N} \\ \sum_{s=0}^{t-1} |v_s^\top z'|^2 \leq S_b^2}} \left| \sum_{s=0}^{t-1} \xi_{s+1} v_s^\top z' \right| + C_\xi \quad (104)$$

Under the event $\bigcap_{z' \in \mathcal{N}} E_{z'}$, we have that the maximum above is $\leq L$. Thus

$$\mathbb{P} \left(\left\| \sum_{s=0}^{t-1} \xi_{s+1} v_s^\top z \right\| > C_\xi + L \right) \leq \mathbb{P} \left(\bigcup_{z' \in \mathcal{N}} E_{z'}^c \right) \quad (105)$$

$$\leq \left(1 + \frac{2R_z}{\varepsilon_{\text{net}}} \right)^d \exp \left(-\frac{L^2}{2C_\xi^2 S_b^2} \right) \leq \frac{\varepsilon}{T}. \quad (106)$$

by the choice of L .

Bounding $\sum_{s=0}^{t-1} \varepsilon_{s+1} v_s^\top \Sigma_t^{-1} v_t$. A crude bound suffices here. We have by (96) that

$$\sum_{s=0}^{t-1} \varepsilon_{s+1} v_s^\top \Sigma_t^{-1} v_t \leq C_\varepsilon \sum_{s=0}^{t-1} |v_s^\top (\Sigma_t)^{-1} v_t| \quad (107)$$

$$\leq C_\varepsilon \sqrt{t} \sqrt{\sum_{s=0}^{t-1} |v_s^\top (\Sigma_t)^{-1} v_t|^2} \quad (108)$$

$$\leq C_\varepsilon \sqrt{T} \left(L_a^2 + \frac{L_v^2}{\mu} \right)^{\frac{1}{2}} d. \quad (109)$$

Bounding $\mu A \Sigma_t^{-1} v_t$. We have by (97) that

$$\left\| \mu A \Sigma_t^{-1} v_t \right\| \leq \mu A R_z = \mu^{\frac{1}{2}} R \left(L_a^2 + \frac{L_v^2}{\mu} \right)^{\frac{1}{2}} d^{\frac{1}{2}} \quad (110)$$

From (88), (106), (109), and (110), and noting that $\|\varepsilon_{t+1} + \xi_{t+1}\| \leq C_\xi + C_\varepsilon$, we get that

$$\mathbb{P} \left(\|A_t v_t - w_t\| > L + (\mu^{\frac{1}{2}} R + C_\varepsilon d^{\frac{1}{2}} \sqrt{T}) \left(L_a^2 + \frac{L_v^2}{\mu} \right)^{\frac{1}{2}} d^{\frac{1}{2}} + 2C_\xi + C_\varepsilon \right) \leq \frac{\varepsilon}{T}. \quad (111)$$

Union-bounding over $0 \leq t \leq T-1$ finishes the proof. \blacksquare

Proof [Proof of Theorem 5] We apply Lemma 28 to the system (74). Let $k' = \lfloor (d+m) \log_2(M) \rfloor$ and $C = C_A(C_B C_u + C_\xi)$. By (76) with $p = 1$, we can write $x_t = \sum_{s=0}^t a_s x_s + v$ with $\sum_{s=0}^{t-1} |a_s| \leq L_a$ and $\|v\|_2 \leq L_v$, where

$$L_a = \frac{2}{\ln 2} (d+m) \quad (112)$$

$$L_v = (k' + 1)^r \left(\frac{2}{\ln 2} (d+m)(C+1) + C_0 + 2 \right) + C_u \left(\frac{2}{\ln 2} (d+m) + 2 \right). \quad (113)$$

using (75) in the last step.

Then Lemma 28 is satisfied with these values of L_a , L_v , and (by Lemma 27), $g(t) = (t+1)^{r-1} C_A C_0 + t^r C_A (C_B C_u + C_\xi)$. Defining ε_{net} , R_z , S_b , L as in Lemma 28, we get that

$$\mathbb{P} \left(\max_{0 \leq t \leq T-1} \|A_t v_t - v_{t+1}\|_2 \leq L + \mu^{\frac{1}{2}} \left(L_a^2 + \frac{L_v^2}{\mu} \right)^{\frac{1}{2}} d + 2C_\xi \right) \geq 1 - \varepsilon. \quad (114)$$

Then with probability $1 - \varepsilon$, we have by Theorem 1 that

$$R_T(A, B) \leq \mu R^2 + \left(L + \mu^{\frac{1}{2}} \left(L_a^2 + \frac{L_v^2}{\mu} \right)^{\frac{1}{2}} d + 2C_\xi \right)^2 (d + m) \ln \left(1 + \frac{TM^2}{d + m} \right) \quad (115)$$

where L_a, L_v, L, M have the desired parameter dependences; take $\mu = 1$. ■

Appendix E. Proof: Partially observable system, stochastic setting

E.1. Learning the steady-state Kalman filter

In the prediction problem, at each time step t , we have observed y_0, \dots, y_t and u_0, \dots, u_t , and are asked to predict y_{t+1} . Note that this does not immediately fit in the framework for online least squares, because y_{t+1} is a linear function of the unobserved x_t (the latent state) plus noise. However, as we will see, we can still place it in this framework if we use a linear autoregressive estimator.

The Kalman filter (Kalman, 1960; Kamen and Su, 1999) gives the optimal linear estimator in the case that the parameters of the LDS, the noises, and the initial state are drawn from known Gaussian distributions, i.e. $x_0 \sim N(x_0^-, \Sigma_0)$, $\xi_t \sim N(0, \Sigma_x)$, $\eta_t \sim N(0, \Sigma_y)$. We can compute matrices $A_{\text{KF}}^{(t)}$, $B_{\text{KF}}^{(t)}$, and $C_{\text{KF}}^{(t)} = C$ such that the optimal linear estimate of the latent state \hat{h}_t and the observation \hat{y}_t are given by a time-varying LDS (taking the y_t as feedback) with those matrices :

$$x_t^- = A_{\text{KF}}^{(t)} x_{t-1}^- + B_{\text{KF}}^{(t)} \begin{pmatrix} u_{t-1} \\ y_{t-1} \end{pmatrix} \quad (116)$$

$$\hat{y}_t = C_{\text{KF}}^{(t)} x_t^- \quad (117)$$

This is known as the predictor form of the system Qin (2006). We will denote $B_{\text{KF}} = (B_{\text{KF},u} \ B_{\text{KF},y})$, where $B_{\text{KF},u}$ and $B_{\text{KF},y}$ are the submatrices acting on u_{t-1} and y_{t-1} , respectively. In this case, x_t^- and \hat{y}_t are the maximum a posteriori (MAP) estimators, and the actual hidden state x_t and the observation y_t are Gaussians when conditioned on $\mathcal{F}_{t-1} = \sigma(y_0, \dots, y_{t-1})$ (the observations up to time $t - 1$): $x_t | \mathcal{F}_{t-1} \sim N(x_{t-1}^-, \Sigma_{\text{KF},x}^{(t)})$ and $y_t | \mathcal{F}_{t-1} \sim N(\hat{y}_t, \Sigma_{\text{KF},y}^{(t)})$ for some covariance matrices $\Sigma_{\text{KF},x}^{(t)}$, $\Sigma_{\text{KF},y}^{(t)}$. Our goal is to predict as well as the Kalman filter without knowing the parameters of the original LDS.

If the original system satisfies Assumption 6 (most notably, it is observable and the noise is i.i.d.), taking $t \rightarrow \infty$, the matrices $A_{\text{KF}}^{(t)}$, and $B_{\text{KF}}^{(t)}$ approach certain fixed matrices A_{KF} and B_{KF} , and the covariance matrices $\Sigma_{\text{KF},x}^{(t)}$ and $\Sigma_{\text{KF},y}^{(t)}$ approach fixed matrices $\Sigma_{\text{KF},x}$ and $\Sigma_{\text{KF},y}$ (Harrison, 1997; Anderson and Moore, 2012). In the Gaussian case, at steady-state, the actual hidden state x_t and observation y_t will be distributed as $x_t | \mathcal{F}_{t-1} \sim N(x_t^-, \Sigma_x)$ and $y_t | \mathcal{F}_{t-1} \sim N(\hat{y}_t, \Sigma_y)$. To simplify the problem, we will assume that the LDS starts with the steady-state covariance⁴, so that the steady-state Kalman filter is the optimal filter for all time.

4. Note that steady-state refers to the covariance of the $x_t | \mathcal{F}_{t-1}$ and $y_t | \mathcal{F}_{t-1}$ being constant, rather than the distribution of $x_t | \mathcal{F}_{t-1}$ and $y_t | \mathcal{F}_{t-1}$ being constant. If the system is not strictly stable, it does not have a steady state in the open-loop setting, because x_t will diverge.

As before, our task is to predict \hat{y}_{t+1} at time step t . The regret is now defined by

$$R_T(A, B, C) = \mathbb{E} \left[\sum_{t=0}^{T-1} \|\hat{y}_{t+1} - y_{t+1}\|_2^2 - \sum_{t=0}^{T-1} \|\hat{y}_{t+1, \text{KF}} - y_{t+1}\|_2^2 \right] \quad (118)$$

where $\hat{y}_{t+1, \text{KF}}$ is the prediction given by (117). The challenge to competing with the Kalman filter prediction is that the Kalman filter has memory: its prediction depends on a state estimate x_t^- kept in memory. We can remove this dependence by “unrolling” the Kalman filter and then truncating. Then we find that $\hat{y}_{t+1, \text{KF}}$ is approximately a linear function of $u_{t-\ell+1}, \dots, u_t$ and $y_{t-\ell+1}, \dots, y_t$ for large enough ℓ : letting $C_{\text{KF}} = C$,

$$\begin{aligned} \hat{y}_{t+1, \text{KF}} &= F u_{t:t-\ell+1} + G y_{t:t-\ell+1} \\ \text{where } F &= (C_{\text{KF}} B_{\text{KF}, u}, C_{\text{KF}} A_{\text{KF}} B_{\text{KF}, u}, \dots, C_{\text{KF}} A_{\text{KF}}^{\ell-1} B_{\text{KF}, u}) \\ \text{and } G &= (C_{\text{KF}} B_{\text{KF}, y}, C_{\text{KF}} A_{\text{KF}} B_{\text{KF}, y}, \dots, C_{\text{KF}} A_{\text{KF}}^{\ell-1} B_{\text{KF}, y}). \end{aligned}$$

In other words, we can approximate the Kalman filter with an autoregressive filter of length ℓ .

The framework of online least-squares (Algorithm 1) now applies with $x_t \leftarrow \begin{pmatrix} u_{t:t-\ell+1} \\ y_{t:t-\ell+1} \end{pmatrix}$, $y_t \leftarrow y_{t+1}$, $A \leftarrow (F, G)$, $n \leftarrow n$, and $m \leftarrow \ell(d+m)$, giving Algorithm 3. We let $u_s = 0$ and $y_s = 0$ for $s < 0$.

E.2. Norms and sufficient length

First we define the *sufficient length* of a system. Roughly speaking, the sufficient length $R(\varepsilon)$ is the length at which we can truncate a finite impulse response (FIR) filter, so that when inputs are bounded by 1, we incur at most ε prediction error at any time step. This notion was introduced by Tu et al. (2017) in the one-dimensional setting.

We first recall some concepts from control theory. In particular, the definition of sufficient length depends on the \mathcal{H}_∞ norm.

Definition 32 Let F be a stable, linear time-invariant (LTI) system, represented as the transfer function $F(z) = \sum_{j=0}^{\infty} F_j z^{-j} \in \mathbb{R}^{n \times m}[[z^{-1}]]$. (This is a matrix-valued Laurent series whose coefficients (F_0, F_1, \dots) form the impulse response function, that is, the response to input $v \in \mathbb{R}^m$ is $(F_0 v, F_1 v, \dots)$.) Define the \mathcal{H}_∞ norm of F to be $\|F\|_{\mathcal{H}_\infty} := \max_{|z|=1} \|F(z)\|_2$.

Definition 33 (Sufficient length condition, (Tu et al., 2017, Definition 1)) We say that a Laurent series (or LTI system) F has stability radius $\rho \in (0, 1)$ if F converges for $\{x \in \mathbb{C} : |x| > \rho\}$. Let F be stable with stability radius $\rho \in (0, 1)$. Fix $\varepsilon > 0$. Define the sufficient length

$$R(\varepsilon) = \left\lceil \inf_{\rho < \gamma < 1} \frac{1}{1-\gamma} \ln \left(\frac{\|F(\gamma z)\|_\infty}{\varepsilon(1-\gamma)} \right) \right\rceil. \quad (119)$$

Note that having a dependence on sufficient length is analogous to having a $\frac{1}{1-\rho(A)}$ dependence on the spectral radius of A , for learning a LDS. Roughly, if we ignore factors depending on condition numbers, for a LDS with dynamics given by A , the sufficient length $R(\varepsilon)$ is on the order of $O\left(\frac{1}{1-\rho(A)} \cdot \ln\left(\frac{1}{\varepsilon}\right)\right)$.

The following lemma says that if we are content with an error of ε , we can safely truncate the impulse response function at length $R(\varepsilon)$. Note that Tu et al. (2017) give the proof for the one-dimensional case, but the same proof works in the multi-dimensional setting.

Lemma 34 (Tu et al. 2017, Lemma 4.1) *Suppose F is stable with stability radius $\rho \in (0, 1)$. Then $\|F_{\geq L}\|_1 := \sum_{k \geq L} \|F(k)\| \leq \max_{\rho < \gamma < 1} \frac{\|F(\gamma z)\|_{\infty} \gamma^L}{1 - \gamma}$. Hence, if $L \geq R(\varepsilon)$, then $\|F_{\geq L}\|_1 \leq \varepsilon$.*

If the LDS is not stable, then we cannot truncate the dependence on past inputs. The key observation is that even if the original LDS is not stable, the LDS defined by the Kalman filter is stable (Anderson and Moore, 2012). Hence, there is some sufficient length at which we can truncate the unrolled Kalman filter.

Theorem 35 ((Anderson and Moore, 2012, §4.4)) *When the LDS (2)–(3) satisfies Assumption 6, the associated Kalman filter (as defined in Section 2.3) is strictly stable: $\rho(A_{KF}) < 1$.*

E.3. Proof of Theorem 8

Proof [Proof of Theorem 8] We first condition on all the noise terms $\|\xi_t\| \leq \|\Sigma_x\| O\left(\sqrt{d \ln\left(\frac{T}{\delta}\right)}\right) =: C_\xi$, $\|\eta_t\| \leq \|\Sigma_y\| O\left(\sqrt{n \ln\left(\frac{T}{\delta}\right)}\right) =: C_\eta$. Choosing the constants large enough, we can ensure this happens with probability $\geq 1 - \frac{\delta}{3}$.

Let $v_t = (h_t; \dots; h_{t-\ell+1}; u_t; \dots; u_{t-\ell+1})$. For convenience set $h_t = 0$, $u_t = 0$ for $t < 0$. Then $(v'_t)_{t=0}^T$ satisfies the following:

$$v'_t = A'v'_{t-1} + \xi'_t \quad (120)$$

$$A' = \left(\begin{array}{ccc|ccc|ccc} A & & & O & & & B & O & \dots & O \\ I & & & O & & & O & \ddots & & \vdots \\ & \ddots & & \vdots & & & \vdots & & & \\ & & I & O & & & O & \dots & & O \\ \hline & & & & & & O & \dots & O & O \\ & & & & & & \hline & & & & & & I & & & O \end{array} \right) \quad (121)$$

$$v'_0 = \begin{pmatrix} h_0 \\ \mathbf{0} \\ u_0 \\ \mathbf{0} \end{pmatrix} \in K'_0 := B_{C_0}^d \oplus \{0\}^{(\ell-1)d} \oplus B_{C_u}^m \oplus \{0\}^{(\ell-1)m} \quad (122)$$

$$\xi'_t = \begin{pmatrix} \xi_t \\ \mathbf{0} \\ u_t \\ \mathbf{0} \end{pmatrix} \in K'_t := B_{C_\xi}^d \oplus \{0\}^{(\ell-1)d} \oplus B_{C_u}^m \oplus \{0\}^{(\ell-1)m}. \quad (123)$$

If $h_0 = h$, $u_0 = u$, and $u_t = 0$ for $t \geq 1$, and there is no noise, then $h_t = A^t h + A^{t-1} B u$ for $t \geq 1$. Hence

$$\left\| (A')^{k'} \begin{pmatrix} h \\ \mathbf{0} \\ u \\ \mathbf{0} \end{pmatrix} \right\|_2 \leq \sqrt{\ell} \left[\max_{0 \leq k \leq k'} f_A(k) \|h\| + \left(\max_{0 \leq k \leq k'-1} f_A(k) C_B + 1 \right) \|u\| \right]. \quad (124)$$

Note also that

$$\max_{v \in B_r^{\ell(d+m)}} \|(A')^{k'} v\| = \max_{(h;u) \in B_r^{d+m}} \left\| (A')^{k'} \begin{pmatrix} h \\ \mathbf{0} \\ u \\ \mathbf{0} \end{pmatrix} \right\| \quad (125)$$

because we can check that

$$\left\| (A')^{k'} \begin{pmatrix} h_{\ell-1} \\ \vdots \\ h_0 \\ u_{\ell-1} \\ \vdots \\ u_0 \end{pmatrix} \right\| \leq \left\| (A')^{k'} \begin{pmatrix} \frac{\|h_{\ell-1;0}\|}{\|h_{\ell-1}\|} h_{\ell-1} \\ \mathbf{0} \\ \frac{\|u_{\ell-1;0}\|}{\|u_{\ell-1}\|} u_{\ell-1} \\ \mathbf{0} \end{pmatrix} \right\|$$

We will apply Lemma 20. Let $k' = \lfloor \ell(d+m) \log_2(M) \rfloor$ and $C' = C_A C_0 + C_A C_B C_u + C_u$. We bound (24). By (124), (125), and Lemma 24,

$$\begin{aligned} f_{A', K'_0 \cup B_{\frac{2}{\ln 2} \ell(d+m)}^{\ell(d+m)}}(k) &= \sqrt{\ell} \left[\max_{0 \leq k \leq k'} f_A(k) \left(C_0 + \frac{2}{\ln 2} \ell(d+m) \right) \right. \\ &\quad \left. + \left(\max_{0 \leq k \leq k'-1} f_A(k) + 1 \right) \left(C_u + \frac{2}{\ln 2} \ell(d+m) \right) \right] \\ &\leq \sqrt{\ell} \left[C_A (k'+1)^{r-1} \left(C_0 + \frac{2}{\ln 2} \ell(d+m) \right) \right. \\ &\quad \left. + (C_A C_B (k')^{r-1} + 1) \left(C_u + \frac{2}{\ln 2} \ell(d+m) \right) \right] \\ &\leq \sqrt{\ell} \left[C' (k'+1)^{r-1} + (C_A (C_B + 1) + 1) \frac{2}{\ln 2} \ell(d+m) \right] \\ \sum_{k=0}^{k'-1} f_{A', K'}(k) &\leq \sqrt{\ell} \left[\sum_{k=0}^{k'-1} \max_{0 \leq j \leq k} f_A(k) C_\xi + \left(\sum_{k=0}^{k'-2} \max_{0 \leq j \leq k} f_A(k) C_B + 1 \right) C_u \right] \\ &\leq \sqrt{\ell} [(k')^r C_\xi + ((k'-1)^r C_B + 1) C_u]. \end{aligned}$$

By Lemma 20, for $0 \leq t \leq T$, there exist $a_s \in \mathbb{R}$ and $v \in \mathbb{R}^{\ell(d+m)}$ such that $v_t = \sum_{s=0}^{t-1} a_s v_s + v$ with

$$\begin{aligned} \sum_{t=0}^{T-1} |a_t| &\leq \frac{2}{\ln 2} \ell(d+m) =: L'_a \\ \|v\| &\leq \max_{0 \leq k \leq k'} f_{A', K'_0 \cup B_{\frac{2}{\ln 2} \ell(d+m)}^{\ell(d+m)}}(k) + \left(\sum_{k=0}^{k'-1} f_{A', K'}(k) \right) \left(\frac{2}{\ln 2} \ell(d+m) + 1 \right) \\ &= \sqrt{\ell} \left[C' (k'+1)^{r-1} + (C_A (C_B + 1) + 1 + (k')^r C_\xi \right. \\ &\quad \left. + ((k'-1)^r C_B + 1) C_u \right] \left(\frac{2}{\ln 2} \ell(d+m) + 1 \right) =: L'_v \end{aligned}$$

Now let $v_t = (y_t; \dots; y_{t-\ell+1}; u_t; \dots; u_{t-\ell+1})$. Then

$$v_t = C'v'_t + \begin{pmatrix} \eta_{t:t-\ell+1} \\ \mathbf{0} \end{pmatrix}$$

where $C' = \left(\begin{array}{c|c} I_\ell \otimes C & O \\ \hline O & I_{\ell m} \end{array} \right)$.

Now if $v'_t = \sum_{s=0}^{t-1} a_s v'_s + v$ with $\sum_{s=0}^{t-1} a_s \leq L_a$ and $\|v\| \leq L_v$, then

$$\begin{aligned} C'v'_t &= \sum_{s=0}^{t-1} a_s C'v'_s + C'v \\ \implies v_t - \begin{pmatrix} \eta_{t:t-\ell+1} \\ \mathbf{0} \end{pmatrix} &= \sum_{s=0}^{t-1} a_s \left(v_s - \begin{pmatrix} \eta_{s:s-\ell+1} \\ \mathbf{0} \end{pmatrix} \right) + C'v \\ \implies v_t &= \sum_{s=0}^{t-1} a_s v_s + \begin{pmatrix} \eta_{s:s-\ell+1} \\ \mathbf{0} \end{pmatrix} - \sum_{s=0}^{t-1} a_s \begin{pmatrix} \eta_{s:s-\ell+1} \\ \mathbf{0} \end{pmatrix} + C'v \end{aligned}$$

so it can be written as a linear combination of previous v_s 's with

$$\begin{aligned} L_a &= L'_a \\ L_v &= C_C L'_v + (L_a + 1)\sqrt{\ell}C_\eta \end{aligned}$$

If $x_0 \sim N(0, \Sigma_{\text{KF},x})$, then $x_0^- = 0$, the steady-state Kalman filter applies, and unfolding the Kalman filter recurrence gives

$$y_{t+1} = \sum_{s=0}^t F_s u_{t-s} + \sum_{s=0}^t G_s y_{t-s} + \cancel{C_{\text{KF}} A_{\text{KF}}^{t+1} x_0^-} + \zeta_{t+1}$$

where $\zeta_{t+1} | \mathcal{F}_t \sim N(0, \Sigma_{\text{KF},y})$ is $\|\Sigma_{\text{KF},y}\|^2$ -subgaussian. Then (y_t) satisfies Lemma 28 with L_a, L_v , and with $w_t = y_{t+1}$ where

$$\begin{aligned} y_{t+1} &= (F, G) \begin{pmatrix} u_{t:t-\ell+1} \\ y_{t:t-\ell+1} \end{pmatrix} + \varepsilon_{t+1} + \zeta_{t+1} \\ \varepsilon_{t+1} &= \sum_{s=\ell}^t (F_s u_{t-s} + G_s y_{t-s}). \end{aligned}$$

By Lemma 27,

$$\max_{0 \leq s \leq t-\ell} \max\{\|u_t\|, \|y_t\|\} \leq T^r C_C C_A (C_B C_u + C_\xi) =: K.$$

By choice of $\ell = R(\varepsilon')$, where $\varepsilon' = \frac{\varepsilon}{K}$, we get

$$\|\varepsilon_{t+1}\| = \left\| \sum_{s=\ell}^t (F_s u_{t-s} + G_s y_{t-s}) \right\| \leq \varepsilon' \max_{0 \leq s \leq t-\ell} \max\{\|u_t\|, \|y_t\|\} \leq \varepsilon.$$

We also know $\xi_{t+1} | \mathcal{F}_t \sim N(0, \Sigma_{\text{KF},y})$. Apply Lemma 28 to get a polynomial bound on $\max_{0 \leq t \leq T-1} \|F_t u_{t:t-\ell+1} + G_t y_{t:t-\ell+1} - y_{t+1}\|$ and Theorem 1 to finish. ■

Appendix F. OLS regret bound

We show Theorem 1.

We note that [Cesa-Bianchi and Lugosi \(2006\)](#) show the case where $A \in \mathbb{R}^{1 \times m}$; the case for $A \in \mathbb{R}^{n \times m}$ essentially follows the same proof. Note that the objective function for A decomposes as a sum of objective functions for each row A_i :

$$\mu \|A\|_F^2 + \sum_{s=0}^{t-1} \|Ax_s - y_s\|^2 = \sum_{i=1}^n \left(\mu \|A_i\|^2 + \sum_{s=0}^{t-1} \|A_i x_s - (y_s)_i\|^2 \right).$$

Hence, running Algorithm 1 is equivalent to running the algorithm on each coordinate of $y_t \in \mathbb{R}^n$ separately.

Note, however, that if we use ([Cesa-Bianchi and Lugosi, 2006](#), Thm. 11.7) as a black box, and apply it to every row of A_t , we get the bound

$$\begin{aligned} R_T(A) &\leq \mu \|A\|_F^2 + \sum_{i=1}^n \max_{0 \leq t \leq T-1} ((y_t)_i - (A_t)_i x_t)_2^2 m \ln \left(1 + \frac{TM^2}{m} \right) \\ &\leq \mu \|A\|_F^2 + \max_{0 \leq t \leq T-1} \|y_t - A_t x_t\|_2^2 mn \ln \left(1 + \frac{TM^2}{m} \right) \end{aligned}$$

where the second inequality follows from using the naive bound $((y_t)_i - (A_t)_i x_t)_2^2 \leq \|y_t - A_t x_t\|_2^2$ and has an extra factor of m .

We refer to [Cesa-Bianchi and Lugosi \(2006\)](#) for the notation we will use.

Proof [Proof of Theorem 1] For consistency with ([Cesa-Bianchi and Lugosi, 2006](#), Thm. 11.7), we index starting from $t = 1$. Define the objective function $\Phi_t^* : \mathbb{R}^{n \times m} \rightarrow \mathbb{R}$ by

$$\Phi_t^*(A) = \mu \|A\|_F^2 + \sum_{s=1}^t \|Ax_s - y_s\|^2$$

and the component objective functions $\Phi_{t,i}^* : \mathbb{R}^{1 \times m} \rightarrow \mathbb{R}$ (for $1 \leq i \leq n$) by

$$\Phi_{t,i}^*(a) = \mu \|a\|^2 + \sum_{s=1}^t \|ax_s - (y_s)_i\|^2$$

so that

$$\Phi_t^*(A) = \sum_{i=1}^n \Phi_{t,i}^*(A_i).$$

Hence the Bregman divergence also decomposes. Let $\Sigma_t = \mu I + \sum_{s=1}^t x_s x_s^\top$. Using the calculation in ([Cesa-Bianchi and Lugosi, 2006](#), pg. 319)

$$\begin{aligned} D_{\Phi_t^*}(A_{t-1}, A_t) &= \sum_{i=1}^n D_{\Phi_{t,i}^*}((A_{t-1})_i, (A_t)_i) \\ &= \sum_{i=1}^n ((A_{t-1})_i x_t - (y_t)_i)^2 x_t^\top \Sigma_t^{-1} x_t \\ &= \|A_{t-1} x_t - y_t\|^2 x_t^\top \Sigma_t^{-1} x_t. \end{aligned}$$

The rest of the proof then follows the proof of Theorem 11.7, together with the remark following the proof. The only difference is that $\|w_{t-1}^\top x_t - y_t\|^2$ has been replaced by $\|A_{t-1}x_t - y_t\|^2$. Note also that they take $\mu = 1$, but the calculations go through with arbitrary $\mu > 0$. ■

Appendix G. Alternate approach to proving anomaly-freeness

In this section we give an alternate proof of anomaly-freeness, which we discovered after the first draft of the paper. This approach does not require the number-theoretic lemma, Lemma 25. However, the polynomial bounds are worse than the proof via the volume doubling argument.

The key fact is that the characteristic polynomial has a multiple of not-too-large degree with small coefficients. Then we can apply the same approach as Lemma 17. The following lemma is an adaptation of (Chen et al., 2016, Lemma 5.4).

Lemma 36 *For any $k \geq 1$ and any $z_1, \dots, z_k \in \mathbb{C}$ of absolute value at most 1, there exists a degree $n = O(k^2 \ln k)$ polynomial $P(z) = \sum_{j=0}^n c_j z^j$ with the following properties:*

$$\begin{aligned} \prod_{i=1}^k (z - z_i) \mid P(z) \\ c_n = 1 \\ |c_j| \leq 11, \quad \forall j \in \{1, \dots, n\}. \end{aligned}$$

Lemma 37 (Claim A.4, Chen et al. (2016)) *Let $z_1, \dots, z_k \in \mathbb{C}$ have absolute value at most 1, and let $Q(z) = \prod_{i=1}^k (z - z_i)$. For any $m = \Omega(k^2 \ln k)$ and $P^*(z) = \sum_{i=0}^m \alpha_i z^i$ with coefficients $|\alpha_i| \leq 10$ for any $i \in \{0, 1, \dots, m\}$, such that every coefficient of $P^*(z) \bmod Q(z)$ is bounded by $2^{-m/k}$.*

Note this is a corrected statement of Claim A.4. (The original statement erroneously had 2^{-m} instead of $2^{-m/k}$.) The proof is an elegant and delightful application of the pigeonhole principle. Note that Chen et al. (2016) state the theorem for z_1, \dots, z_k on the unit circle, but the exact same proof goes through if they are allowed to be inside the unit circle.

Proof [Proof of Lemma 36] Take $m = \Theta(k^2 \ln k)$, so that $2^{-m/k} \leq \varepsilon := \frac{1}{12}$, and so that Lemma 37 applies. Let $P^*(z) = \sum_{i=0}^n \alpha_i z^i$ be as in Lemma 37, with $n \leq m$ and nonzero leading coefficient $\alpha_n \neq 0$, and let $r(z) = P^*(z) \bmod Q(z) = \sum_{i=0}^{k-1} \gamma_i z^i$. Set $\gamma_i = 0$ for $i \geq k$. Because $r(z)$ is the residue, $P^*(z) - r(z)$ is divisible by $\prod_{i=1}^k (z - z_i)$.

Let

$$P(z) = \frac{P^*(z) - r(z)}{\alpha_n - \gamma_n}.$$

Note that $P(z)$ is monic, and every coefficient is bounded in absolute value by $\frac{10+\varepsilon}{1-\varepsilon} \leq 11$. This is the desired polynomial. ■

Using the multiple of the characteristic polynomial given by Lemma 36 instead of the characteristic polynomial in Lemma 17 gives the following result on anomaly-freeness.

Lemma 38 *Given Assumptions 2 with $B = O$, suppose that $A \in \mathbb{R}^{d \times d}$ is diagonalizable as $A = VDV^{-1}$ where $\|V\| \|V^{-1}\| \leq C_A$. Suppose that $|w^\top x_t| = \Omega(d^2(\ln d)C_A C_\xi)$ and $t = \Omega(d^2 \ln d)$. Then for any unit vector $w \in \mathbb{R}^d$, there exist $O\left(\min\left\{\frac{|w^\top x_t|}{d^4(\ln d)^2 C_A C_\xi}, \frac{t}{d^4(\ln d)^2}\right\}\right)$ values of s , $0 \leq s \leq t-1$, such that*

$$|w^\top x_s| \geq \Omega\left(\frac{|w^\top x_t|}{d^2 \ln d}\right).$$

Proof Because $\rho(A) \leq 1$, all zeros of the characteristic polynomial of A have absolute value ≤ 1 . By Lemma 36, there exists a multiple $p(x) = \sum_{i=0}^n a_i^{(1)} x^i$ of the characteristic polynomial of A such that:

- p has degree $n = O(k \ln k)$.
- p is monic ($a_n^{(1)} = 1$).
- All coefficients are bounded: $|a_i^{(1)}| \leq 11$.

Then by unfolding the recurrence as in (16), we obtain

$$x_t = \sum_{i=0}^{n-1} -a_i^{(1)} x_{t-n+i} + \underbrace{\sum_{i=0}^{n-1} a_i^{(1)} \sum_{\tau=1}^i A^{i-\tau} \xi_{t-n+\tau} + \sum_{\tau=1}^n A^{n-\tau} \xi_{t-n+\tau}}_{=:v^{(1)}}$$

Note that $\|A^i \xi\| \leq \|VD^i V^{-1}\| \|\xi\| \leq C_A \|\xi\| \leq C_A C_\xi$. We write $x_t = -\sum_{i=1}^{n-1} a_i^{(1)} x_{t-n+i} + v^{(1)}$ where

$$\begin{aligned} \|v^{(1)}\| &= \left\| \sum_{i=0}^{n-1} a_i^{(1)} \sum_{\tau=1}^i A^{i-\tau} \xi_{t-n+\tau} + \sum_{\tau=1}^n A^{n-\tau} \xi_{t-n+\tau} \right\| \\ &\leq \sum_{i=0}^{n-1} |a_i^{(1)}| \sum_{\tau=1}^i C_A \|\xi_{t-n+\tau}\| + \sum_{\tau=1}^n C_A \|\xi_{t-n+\tau}\| \\ &\leq 11(n-1)C_A C_\xi + nC_A C_\xi \leq 12nC_A C_\xi \end{aligned}$$

If $w \in \mathbb{R}^d$ is a unit vector such that $|w^\top x_t| \geq 24nC_A C_\xi$, we have $|w^\top x_t| \geq 2\|v^{(1)}\|$. Noting that $w^\top x_t = \sum_{i=0}^{n-1} -a_i^{(1)} x_{t-n+i} + v^{(1)}$, by Lemma 16 there exists an index $0 \leq i \leq n-1$ such that

$$|w^\top x_{t-n+i}| \geq \frac{|w^\top x_t| - \|v^{(1)}\|}{\sum_{i=0}^{n-1} |a_i^{(1)}|} \geq \frac{|w^\top x_t|}{2 \sum_{i=0}^{n-1} |a_i^{(1)}|} \geq \frac{|w^\top x_t|}{22n}$$

In order to obtain many large past x 's, we apply the same argument on sequences $x_t, x_{t-k}, \dots, x_{t-kd}$ by considering the recurrence

$$\begin{aligned} x_t &= A' x_{t-k} + \xi_t^{(k)} \\ A' &= A^k \\ \xi_t^{(k)} &= \sum_{j=1}^k A^{k-j} \xi_{t-k+j}. \end{aligned}$$

Let $p_k(x) = \sum_{i=0}^n a_i^{(k)} x^i$ be a multiple of the characteristic polynomial of A^k given by Lemma 36, for some $n = O(d^2 \ln d)$. We obtain $x_t = -\sum_{i=1}^{n-1} a_i^{(k)} x_{t-k(n-i)} + v^{(k)}$ with

$$\|v^{(k)}\| = \left\| \sum_{i=0}^{n-1} a_i^{(k)} \sum_{\tau=1}^i A^{i-\tau} \xi_{t-nk+\tau k}^{(k)} + \sum_{\tau=1}^n A^{n-\tau} \xi_{t-nk+\tau k}^{(k)} \right\| \quad (126)$$

$$\leq 11k(n-1)C_A C_\xi + knC_A C_\xi \leq 12knC_A C_\xi \quad (127)$$

and $\sum_{i=0}^{n-1} |a_i^{(k)}| \leq 11n$. We then pick $k = 1, 2, \dots, O\left(\min\left\{\frac{|w^\top x_t|}{d^2(\ln d)C_A C_\xi}, \frac{t}{d^2 \ln d}\right\}\right)$. For each choice of k , we know by design that $|w^\top x_t| \geq 2\|v^{(k)}\|_2$, and therefore there must exist an x in the sequence x_{t-k}, \dots, x_{t-kd} such that $|w^\top x| = \Omega\left(\frac{|w^\top x_t|}{d^2 \ln d}\right)$. In this way, we are able to collect in total $L = O\left(\min\left\{\frac{|w^\top x_t|}{d^2(\ln d)C_A C_\xi}, \frac{t}{d^2 \ln d}\right\}\right)$ many such x 's. To finish the argument, we note that out of the L collected x 's, there are at least $O\left(\frac{L}{d^2 \ln d}\right)$ distinct ones, since one x can appear in at most $O(d^2 \ln d)$ different sequences. \blacksquare

Appendix H. Open Questions

Several fundamental questions come to mind:

1. *Is the $T^{\frac{2r+1}{2r+2}}$ rate optimal?* Even in the diagonalizable ($r = 1$) case, this is unresolved.
2. *What is the rate for partially observed LDS when the noise is not Gaussian?* We stated Theorem 8 for Gaussian noise only, but a similar result will hold as long as at steady state, $\mathbb{E}[y_t | \mathcal{F}_{t-1}]$ is given by a linear function of $y_{t-1:0}, u_{t-1:0}$, and the estimated state x_0^- . This is required in order for the random variable $y_t | \mathcal{F}_{t-1}$ to be a linear function of past observations and inputs, plus a random variable ζ_t with zero mean. In general, if the noise ξ_t is not Gaussian, then ζ_t is not zero-mean (even if ξ_t is zero-mean). We use the same machinery as in the proof of Theorem 5 to conclude Theorem 8, so our proof strategy cannot handle arbitrary zero-mean noise ξ_t .

For non-Gaussian zero-mean noise ξ_t , we can instead treat the ζ_t as adversarial noise, and use the machinery behind Theorem 3 to obtain a $T^{\frac{2r+1}{2r+2}}$ regret bound. It is an interesting question whether we can obtain polylogarithmic regret with respect to the best linear filter in this case.

3. *Can we obtain bounds depending on system order d rather than rollout length ℓ ?* Theorem 8 depends polynomially on the sufficient rollout length ℓ , rather than the intrinsic dimensionality of the problem given by d and m . This seems to be a limitation of using the improper autoregressive approach; can we do better using techniques from system identification?

We leave these for future work.