

Data Sketching for Large-Scale Kalman Filtering

Dimitris Berberidis, *Student Member, IEEE*, and Georgios B. Giannakis, *Fellow, IEEE*

Abstract—In an age of exponentially increasing data generation, performing inference tasks by utilizing the available information in its entirety is not always an affordable option. This paper puts forth approaches to render tracking of large-scale dynamic processes via a Kalman filter affordable, by processing a reduced number of data. Three distinct methods are introduced for reducing the number of data involved in the correction step of the filter. Toward this goal, the first two methods employ random projections and innovation-based censoring to effect dimensionality reduction and measurement selection, respectively. The third method achieves reduced complexity by leveraging sequential processing of observations and selecting a few informative updates based on an information-theoretic metric. Simulations on synthetic data compare the proposed methods with competing alternatives, and corroborate their efficacy in terms of estimation accuracy over complexity reduction. Finally, monitoring large networks is considered as an application domain, with the proposed methods tested on Kronecker graphs to evaluate their efficiency in tracking traffic matrices and time-varying link costs.

Index Terms—Tracking, dimensionality reduction, censoring, random projections, Kalman filter, traffic matrix.

I. INTRODUCTION

TRACKING nonstationary dynamic processes is of paramount importance in various applications. In the context of big data, being able to perform accurate and economical state estimation may render problems of prohibitive scale feasible. Weather prediction is an example of tracking a slowly-varying dynamic process, from a massive volume of observations acquired from fast-sampling sensors per time interval; see e.g., [1]. Monitoring large and dynamically evolving networks, where nodes may join or leave and connections may be established or lost as time progresses, provides an exciting domain in which the acquisition and processing of network-wide performance metrics becomes challenging as the network size increases [2, Ch. 8]. For instance, monitoring path metrics such as delays or loss rates is challenging primarily because the number of paths generally grows as the square of the number of nodes in the network. Therefore, measuring and

storing the delays of all possible origin-destination pairs is hard in practice, even for moderate-size networks [2].

In this context, efforts to reduce the number of measurements used for tracking have pursued two different directions. The first is that of optimal experimental design (OED), where the goal is to perform *model-driven* sensor selection based on ensemble performance metrics (e.g. the trace of the error covariance). Channel-aware dimensionality reduction of observations was reported in [3] and [4] using distributed wireless sensor networks (WSNs). Optimal and near optimal sensor schedules for a finite time horizon estimation was dealt with in [5], while entropy- and mutual-information-based sensor selection were advocated in [6] and [7]. A posterior-CRLB-based method to select sensors for tracking was introduced in [8], via convex optimization in [9] and [10], while a greedy algorithm leveraging submodularity was developed in [11] for measurement selection in sequential estimation. The latter has also been advocated as a means of reducing the complexity of Kalman filters that operate with limited processing resources [12]. OED is nicely attuned for designing low-dimensional observation models, but it is *data-agnostic* and thus sub-optimal when observations become available and need to be reduced.

The second direction is that of *data-driven* methods that select available measurements for processing. Specifically, censoring has recently been employed to select data for distributed parameter estimation using resource-constrained WSNs, thus trading off performance for tractability [13], [14]. Furthermore, censoring has been proposed for signal estimation using WSNs, for tracking, and control of dynamical processes [15]–[17]. However, existing works on censoring mainly aim at reducing the rate at which sensors communicate their observations, and pertinent methods exhibit large computational complexity and storage requirements, which can be possibly afforded only at the fusion center.

The goal of this paper is to perform reliable tracking using the Kalman filter (KF), while reducing the amount of data and the computational complexity involved. To this end, the first two methods employ random projections and innovation-based censoring to respectively effect dimensionality reduction and measurement selection. The third method reduces complexity by leveraging sequential processing of observations when the noise is uncorrelated, and by selecting a few informative updates based on an information-theoretic metric. Finally, an efficient backward smoothing method is developed to mitigate the performance degradation caused by dimensionality reduction. Corroborating simulations compare with state-of-the-art greedy measurement selection algorithms, and illustrate the efficacy of the novel schemes. To demonstrate the applicability of the proposed

Manuscript received June 4, 2016; revised January 6, 2017; accepted March 28, 2017. Date of publication April 6, 2017; date of current version May 16, 2017. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Gesualdo Scutari. This work was supported in part by the Army Research Office under Grant W911NF-15-1-0492 and in part the National Science Foundation under Grant 1343860, Grant 1442686, and Grant 1500713. (*Corresponding author: Georgios B. Giannakis.*)

The authors are with the ECE Department, University of Minnesota, Minneapolis, MN 55455 USA (e-mail: bermpp001@umn.edu; georgios@umn.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2017.2691662

update selection approach on real-world problems, traffic matrix estimation and network link cost estimation is also considered.

The rest of the paper is organized as follows. Section II introduces the proposed model of reduced complexity KF. Sections III and IV present the two dimensionality reduction modules based on RPs and censoring, respectively. The proposed update selection method is introduced in Section V. Numerical experiments are in Section VII, while Section VIII includes experiments on network monitoring. Finally, concluding remarks are given in Section IX.

Notation. Lower- (upper-) case boldface letters denote column vectors (matrices). Calligraphic symbols are reserved for sets, while T stands for transposition. Vectors $\mathbf{0}$, $\mathbf{1}$, and \mathbf{e}_n denote the all-zeros, the all-ones, and the n -th canonical vector, respectively. Symbol $\mathbf{1}_E$ denotes the indicator for the event E . Notation $\mathcal{N}(\mathbf{m}, \mathbf{C})$ stands for the multivariate Gaussian distribution with mean \mathbf{m} and covariance matrix \mathbf{C} , while $\text{tr}(\mathbf{X})$, $\lambda_{\min}(\mathbf{X})$, and $\lambda_{\max}(\mathbf{X})$ are reserved for the trace, the minimum and maximum eigenvalues of matrix \mathbf{X} , respectively. Symbol $D \gg$ is used to denote that the number of observations D is “prohibitively large” relative to the problem at hand, as well as the computing platform.

II. PROBLEM STATEMENT AND PRELIMINARIES

Consider the following linear dynamical system model

$$\boldsymbol{\theta}_n = \mathbf{F}_n \boldsymbol{\theta}_{n-1} + \mathbf{G}_n \mathbf{u}_n + \mathbf{w}_n \quad (1)$$

$$\mathbf{y}_n = \mathbf{X}_n \boldsymbol{\theta}_n + \mathbf{v}_n \quad (2)$$

where $\boldsymbol{\theta}_n \in \mathbb{R}^p$ denotes the state vector at time n ; \mathbf{F}_n is the known state-transition matrix; \mathbf{G}_n and \mathbf{u}_n are known, deterministic control-input model and control-input vector respectively; $\mathbf{y}_n \in \mathbb{R}^D$ the measurement vector, and \mathbf{X}_n is the known $D \times p$ measurement matrix; while \mathbf{w}_n and \mathbf{v}_n are zero-mean, mutually uncorrelated and individually uncorrelated across time random noise vectors, with respective covariance matrices \mathbf{Q}_n and \mathbf{R}_n . The initial state $\boldsymbol{\theta}_0$ has mean \mathbf{m}_0 , and covariance \mathbf{P}_0 .

Given the information-bearing data $\mathcal{I}_n := \{\mathbf{y}_n, \mathbf{X}_n, \mathbf{R}_n\}$ of the measurement (2) at time n , the most recent estimate $\hat{\boldsymbol{\theta}}_{n-1|n-1}$ and its covariance matrix $\mathbf{P}_{n-1|n-1}$, the celebrated KF yields the minimum mean-square error (MMS-E) optimal estimate $\hat{\boldsymbol{\theta}}_{n|n}$ in two steps. First, the state prediction $\hat{\boldsymbol{\theta}}_{n|n-1}$ and its covariance matrix $\mathbf{P}_{n|n-1}$ are obtained using the model dynamics $\{\mathbf{F}_n, \mathbf{Q}_n\}$ as [cf. (1)]

$$\begin{aligned} \hat{\boldsymbol{\theta}}_{n|n-1} &= \mathbf{F}_n \hat{\boldsymbol{\theta}}_{n-1|n-1} + \mathbf{G}_n \mathbf{u}_n \\ \mathbf{P}_{n|n-1} &= \mathbf{F}_n \mathbf{P}_{n-1|n-1} \mathbf{F}_n^T + \mathbf{Q}_n. \end{aligned} \quad (3)$$

Subsequently, as \mathcal{I}_n becomes available, $\hat{\boldsymbol{\theta}}_{n|n}$ is obtained as

$$\hat{\boldsymbol{\theta}}_{n|n} = \arg \min_{\boldsymbol{\theta}} \|\mathbf{y}_n - \mathbf{X}_n \boldsymbol{\theta}\|_{\mathbf{R}_n^{-1}}^2 + \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_{n|n-1}\|_{\mathbf{P}_{n|n-1}^{-1}}^2. \quad (4)$$

The first term of the cost in (4) is a weighted least-squares term fitting the state $\boldsymbol{\theta}$ with \mathcal{I}_n that arises from the linear observation model in (2); while the second regularization term corresponds to treating $\hat{\boldsymbol{\theta}}_{n|n-1}$ as a prior of $\boldsymbol{\theta}_n$. Solving (4) and applying the matrix inversion lemma (MIL) yields the well known KF

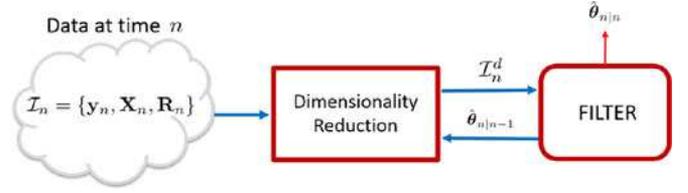


Fig. 1. Reduced-dimension filtering.

Algorithm 1: Reduced-dimension KF.

Initialization: $\hat{\boldsymbol{\theta}}_{0|0} = \mathbf{m}_0$, $\mathbf{P}_{0|0} = \mathbf{P}_0$

for $n = 1 : N$ **do**

Prediction Step

$$\hat{\boldsymbol{\theta}}_{n|n-1} = \mathbf{F}_n \hat{\boldsymbol{\theta}}_{n-1|n-1} + \mathbf{G}_n \mathbf{u}_n$$

$$\mathbf{P}_{n|n-1} = \mathbf{F}_n \mathbf{P}_{n-1|n-1} \mathbf{F}_n^T + \mathbf{Q}_n$$

Data Reduction

$$\{\check{\mathbf{y}}_n, \check{\mathbf{X}}_n, \check{\mathbf{R}}_n\} = \text{Sketching}(\{\mathbf{y}_n, \mathbf{X}_n, \mathbf{R}_n\}, \hat{\boldsymbol{\theta}}_{n|n-1})$$

Correction Step

$$\hat{\boldsymbol{\theta}}_{n|n} = \hat{\boldsymbol{\theta}}_{n|n-1} + \check{\mathbf{K}}_n (\check{\mathbf{y}}_n - \check{\mathbf{X}}_n \hat{\boldsymbol{\theta}}_{n|n-1})$$

$$\check{\mathbf{K}}_n = \mathbf{P}_{n|n-1} \check{\mathbf{X}}_n^T (\check{\mathbf{X}}_n \mathbf{P}_{n|n-1} \check{\mathbf{X}}_n^T + \check{\mathbf{R}}_n)^{-1}$$

$$\mathbf{P}_{n|n} = (\mathbf{I}_p - \check{\mathbf{K}}_n \check{\mathbf{X}}_n) \mathbf{P}_{n|n-1}$$

end for

correction step, e.g., [18, p. 205]

$$\hat{\boldsymbol{\theta}}_{n|n} = \hat{\boldsymbol{\theta}}_{n|n-1} + \mathbf{K}_n (\mathbf{y}_n - \mathbf{X}_n \hat{\boldsymbol{\theta}}_{n|n-1})$$

where the so-termed KF gain \mathbf{K}_n and the state covariance update are given by

$$\mathbf{K}_n = \mathbf{P}_{n|n-1} \mathbf{X}_n^T (\mathbf{X}_n \mathbf{P}_{n|n-1} \mathbf{X}_n^T + \mathbf{R}_n)^{-1}$$

$$\mathbf{P}_{n|n} = (\mathbf{I}_p - \mathbf{K}_n \mathbf{X}_n) \mathbf{P}_{n|n-1}.$$

A dual form of the KF known as the information filter (IF) relies on the MIL to offer a more efficient solver of (4) as D grows large [18, Ch. 7]. Nevertheless, even the low-complexity IF requires $\mathcal{O}(Dp^2)$ multiplications to solve (4) in the case of uncorrelated observations (\mathbf{R}_n diagonal), and $\mathcal{O}(D^2p)$ in general. Therefore, for large-scale KF problems where $D \gg$, dimensionality reduction of the datasets \mathcal{I}_n is well motivated for rendering the solution of (4) computationally tractable, while also reducing other data-related costs, such as storage and transmission.

Towards this goal, we introduce a reduced-complexity Kalman-like filter (see Algorithm 1) that extracts a reduced (size $d < D$), yet informative dataset $\mathcal{I}_n^d := \{\check{\mathbf{y}}_n, \check{\mathbf{X}}_n, \check{\mathbf{R}}_n\}$ from the original \mathcal{I}_n , where $\check{\mathbf{y}}_n \in \mathbb{R}^d$, $\check{\mathbf{X}}_n \in \mathbb{R}^{d \times p}$ and $\check{\mathbf{R}}_n \in \mathbb{R}^{d \times d}$ are the corresponding reduced-dimension observation vector, measurement matrix, and covariance matrix; see also Fig. 1. Consequently, the problem reduces to the design of low-complexity sketching modules for informative dimensionality reduction. In the ensuing two sections, a *data-agnostic* method based on RPs followed by a *data-adaptive* method based on censoring are developed.

III. RP-BASED KF

RP-based dimensionality reduction amounts to premultiplying measurements and regressors $\{\mathbf{y}_n, \mathbf{X}_n\}$ with a random

Algorithm 2: RP sketching module.

Dimensionality reduction with RPs

$$\begin{aligned}\check{\mathbf{y}}_n &= \mathbf{S}_d \mathbf{H} \mathbf{\Gamma} \mathbf{y}_n \\ \check{\mathbf{X}}_n &= \mathbf{S}_d \mathbf{H} \mathbf{\Gamma} \mathbf{X}_n \\ \check{\mathbf{R}}_n &= \mathbf{S}_d \mathbf{H} \mathbf{\Gamma} \mathbf{R}_n (\mathbf{S}_d \mathbf{H} \mathbf{\Gamma})^T\end{aligned}$$

matrix \mathbf{H} , and a diagonal matrix $\mathbf{\Gamma}$, whose entries take the values $\{+1/\sqrt{D}, -1/\sqrt{D}\}$ equiprobably. The net result is a linear transformation of the measurement equations so that all rows convey “comparable information”. A subset of d rows of the transformed system is then extracted by simple random sampling, implemented by left multiplication with a random $d \times D$ selection matrix \mathbf{S}_d .

Originally developed for linear regressions [19]–[21], the novelty here is RP-based reduced-dimensionality tracking of dynamical processes. Applying the Hadamard preconditioning and random sampling matrices on (2) yields the reduced-dimension observation model

$$\check{\mathbf{y}}_n := \mathbf{S}_d \mathbf{H} \mathbf{\Gamma} \mathbf{y}_n = \mathbf{S}_d \mathbf{H} \mathbf{\Gamma} (\mathbf{X}_n \boldsymbol{\theta}_n + \mathbf{v}_n) = \check{\mathbf{X}}_n \boldsymbol{\theta}_n + \check{\mathbf{v}}_n$$

where $\check{\mathbf{v}}_n := \mathbf{S}_d \mathbf{H} \mathbf{\Gamma} \mathbf{v}_n$ is zero mean with covariance $\check{\mathbf{R}}_n = \mathbf{S}_d \mathbf{H} \mathbf{\Gamma} \mathbf{R}_n (\mathbf{S}_d \mathbf{H} \mathbf{\Gamma})^T$. Given $\hat{\boldsymbol{\theta}}_{n|n-1}$ and the reduced data \mathcal{I}_n^d , state estimate $\hat{\boldsymbol{\theta}}_{n|n}$ can be obtained as [cf. (4)]

$$\hat{\boldsymbol{\theta}}_{n|n} = \arg \min_{\boldsymbol{\theta}} \|\check{\mathbf{y}}_n - \check{\mathbf{X}}_n \boldsymbol{\theta}\|_{\check{\mathbf{R}}_n}^2 + \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_{n|n-1}\|_{\mathbf{P}_{n|n-1}}^2. \quad (5)$$

Solving (5) and applying the MIL yields the novel RP-based KF, which is summarized as Algorithm 1 using Algorithm 2 as sketching module.

Implementing RPs can have affordable $\mathcal{O}(Dp \log d)$ complexity if \mathbf{H} is chosen to be a pseudo-random Hadamard matrix of size $D = 2^\nu$ for $\nu \in \mathbb{Z}_+$. Different from the more elaborate approaches in [3] and [4], the proposed RP-KF is an easy-to-implement, “one-size-fits-all” reduced-complexity tracker, using data-agnostic dimensionality reduction. Furthermore, RP-KF’s estimation performance can be guaranteed as asserted in the ensuing proposition, which provides a benchmark for the data-driven methods introduced in the following section.

Proposition 1: With $\mathbf{R}_n = \sigma_n^2 \mathbf{I}_D$, let $\mathbf{A}_n := [\mathbf{P}_{n|n-1}^{-1/2}, \sigma_n^{-1} \mathbf{X}_n^T]^T$, $\mathbf{b}_n := [\mathbf{P}_{n|n-1}^{-1/2} \hat{\boldsymbol{\theta}}_{n|n-1}, \sigma_n^{-1} \mathbf{y}_n^T]^T$, and $\mathbf{A}_n = \mathbf{U}_n \boldsymbol{\Lambda}_n \mathbf{V}_n^T$ the singular value decomposition of \mathbf{A}_n . If $\|\mathbf{U}_n \mathbf{U}_n^T \mathbf{b}_n\|_2 \geq \gamma \|\mathbf{b}_n\|_2$ for some $\gamma \in (0, 1]$, then by choosing $d = \mathcal{O}(p \ln(pD)/\epsilon)$ the following bound for the RP-KF estimates holds w.h.p.

$$\|\hat{\boldsymbol{\theta}}_{n|n} - \hat{\boldsymbol{\theta}}_{n|n}^*\|_2 \leq \sqrt{\epsilon} \left(\kappa(\mathbf{A}_n) \sqrt{\gamma^{-2} - 1} \right) \|\hat{\boldsymbol{\theta}}_{n|n}^*\|_2$$

where $\kappa(\mathbf{A}_n)$ denotes the condition number of \mathbf{A}_n , and $\hat{\boldsymbol{\theta}}_{n|n}^*$ is the full-data KF estimate.

Proof: See Appendix 1. ■

Proposition 1 asserts that, per time slot n , the estimate $\hat{\boldsymbol{\theta}}_{n|n}$ of the RP-KF can be guaranteed to be close enough, in the relative squared-error sense, to the estimate $\hat{\boldsymbol{\theta}}_{n|n}^*$ of the full KF, if the reduced dimension d is chosen to be large enough. Note that Proposition 1 only provides a *per-step* error guarantee,

meaning that the error between the estimates of the RP-KF and the full-data KF at slot n is bounded, given that the two filters share a common estimate at slot $n - 1$. Bounding the RP-KF error across multiple time slots is a more challenging task that goes beyond the scope and claims of the present paper. Naturally, the quality of the approximation also depends on other parameters such as the observation matrix, noise variance and covariance of prediction. Nevertheless, being data-agnostic and requiring storage and processing of \mathcal{I}_n in batch form per time slot renders the RP-KF less attractive in practice, and motivates the algorithms presented in the following two sections.

IV. CENSORING-BASED KF

Measurement censoring for estimating dynamical processes has been advocated as a means of reducing the inter-sensor transmission overhead when WSNs are deployed for distributed tracking [16], [17]; see also [15], [22], where censoring is employed for event-based estimation. Since the goal in the aforementioned applications is saving communication resources, censoring is performed solely on measurements \mathbf{y}_n , with \mathbf{X}_n and \mathbf{R}_n assumed known and used even for the censored entries of \mathbf{y}_n ; thus, [14]–[16], [22], and [17] rely on reducing the dimensionality of a dataset that only consists of observations; that is, $\mathcal{I}_n := \{\mathbf{y}_n\}$. A subset of d observations $\mathcal{I}_n^d := \{\mathbf{y}_{n, \mathcal{S}_n}\}$ is obtained, where $y_{n,i}$ is the i -th entry of \mathbf{y}_n , and $\mathcal{S}_n \subseteq \{1, \dots, D\}$ denotes a set collecting the indices of uncensored observations. Given $\mathbf{y}_{n, \mathcal{S}_n}$, \mathbf{X}_n and \mathbf{R}_n , [14]–[17], [22] develop sequential estimators to optimally estimate $\boldsymbol{\theta}_n$. Targeting reduction of communication load, optimal (in the maximum likelihood or MMSE sense) estimation from (un)censored observations comes with complexity comparable to that of using the full set of D measurements. For our big-data setups, this is not affordable.

Since the aim is dimensionality and complexity reduction, the starting point is on censoring entire rows of the full dataset $\mathcal{I}_n^D := \{\mathbf{y}_n, \mathbf{X}_n, \mathbf{R}_n\}$, in order to obtain a reduced set $\mathcal{I}_n^d := \{\mathbf{y}_{n, \mathcal{S}_n}, \mathbf{X}_{n, \mathcal{S}_n}, \mathbf{R}_{n, \mathcal{S}_n}\}$, where $\mathbf{x}_{n,i}^T$ denotes the i -th row of \mathbf{X}_n and $\mathbf{R}_{n, \mathcal{S}_n} := \text{cov}(\mathbf{v}_{n, \mathcal{S}_n})$. The goal here is to develop censoring rules in order to obtain \mathcal{S}_n , so that \mathcal{I}_n^d is an “informative” subset of \mathcal{I}_n . Most existing censoring schemes adopt the innovation $\tilde{\mathbf{y}}_n := \mathbf{y}_n - \mathbf{X}_n \hat{\boldsymbol{\theta}}_{n|n-1}$ as a measure of information contained in \mathbf{y}_n .

One approach—henceforth termed *block censoring* (BC)—is to censor the entire vector \mathbf{y}_n . From an information-theoretic viewpoint [17], the optimal BC rule relies on the magnitude of the prewhitened innovation $\boldsymbol{\Sigma}_n^{-1/2} \tilde{\mathbf{y}}_n$, where $\boldsymbol{\Sigma}_n := \text{cov}(\tilde{\mathbf{y}}_n) = \mathbf{X}_n \mathbf{P}_{n|n-1} \mathbf{X}_n^T + \mathbf{R}_n$; thus, \mathcal{S}_n is found as

$$\mathcal{S}_n := \begin{cases} \{1, \dots, D\}, & \|\boldsymbol{\Sigma}_n^{-1/2} \tilde{\mathbf{y}}_n\|_2 > \tau_n \\ \emptyset, & \text{otherwise} \end{cases}. \quad (6)$$

Clearly, having $\mathcal{S}_n = \emptyset$ corresponds to skipping the correction step of the KF. A major shortcoming of (6) is the cubic complexity $\mathcal{O}(D^3)$ associated with inverting $\boldsymbol{\Sigma}_n$. Furthermore, BC-KF can only reduce the data cost *on average* across iterations by entirely skipping correction steps.

Our idea of a more attractive alternative is to possibly censor separately each entry of \mathcal{I}_n . Such an entry-wise censoring rule

yields \mathcal{S}_n as

$$\mathcal{S}_n := \{1 \leq i \leq D \mid |\sigma_{n,i}^{-1} \tilde{y}_{n,i}| > \tau_n\} \quad (7)$$

where $\sigma_{n,i}^2 = [\mathbf{R}_n]_{ii}$, and τ_n can be tuned so that the set cardinality $|\mathcal{S}_n| \approx d$. Note that normalization in (6) includes pre-whitening of the innovation vector that is effected through left multiplication with $\Sigma_n^{-1/2}$. In contrast, any low-complexity *per-entry adaptive* censoring rule cannot explicitly consider the cross-correlation between different measurements, meaning that Σ_n is not utilized. Instead, (7) uses the diagonal entries of \mathbf{R}_n as an approximate measure of the per-entry innovation variance and therefore it serves as a normalization factor. Compared to BC-KF, the innovation-based entry-wise rule of (7) is more flexible in reducing the available data, since it can censor any subset of observations at slot n at much lower complexity. Nevertheless, to accurately perform measurement selection with (7), $|\tilde{y}_{n,i}|$ must reflect how informative $y_{n,i}$ is for the purpose of tracking θ_n . Using for this purpose the entry-wise predictor-based innovations $\tilde{y}_{n,i} := y_{n,i} - \mathbf{x}_{n,i}^T \hat{\theta}_{n|n-1}$ is a possibility, but turns out to be unsuitable for the proposed reduced-complexity KF, due to the fact that censoring rule (7) tends to yield “biased” observations for a given $\hat{\theta}_{n|n-1}$. Correction of this bias is possible through the incorporation of a maximum likelihood criterion (see, e.g. [23]). Since such an approach requires the additional knowledge of \mathbf{X}_n and incurs computational complexity at least as high as that of the full-data KF, it is only suitable for reducing the communication overhead.

Targeting a more suitable censoring rule, the *adaptive censoring* least mean-square (AC-LMS) algorithm we introduced in [23] for non-dynamical regressions can be employed to discard uninformative rows of \mathcal{I}_n . Within time slot n , rows of $(\mathbf{y}_n, \mathbf{X}_n)$ are processed sequentially; given a temporary estimate $\hat{\theta}_{n|n-1,i}$, the i -th row is discarded when indicated so by the censoring variable (1. denotes the indicator function)

$$c_i := 1\{|y_{n,i} - \mathbf{x}_{n,i}^T \hat{\theta}_{n|n-1,i-1}| \leq \tau_n \sigma_{n,i}^{-1}\}. \quad (8)$$

Given $\{\mathbf{x}_{n,i}^T, \hat{\theta}_{n|n-1,i-1}, \tau_n\}$, a “censoring slab” is specified in \mathbb{R}^{p+1} to determine whether $(\mathbf{x}_{n,i}, y_{n,i})$ will be censored (if inside this slab) or not (if outside this slab). If deemed informative enough ($c_i = 0$), the i -th row is added to \mathcal{S}_n , and subsequently involved in updating $\hat{\theta}_{n|n-1,i}$ as

$$\hat{\theta}_{n|n-1,i} = \hat{\theta}_{n|n-1,i-1} + (1 - c_i)\mu \mathbf{x}_{n,i} (y_{n,i} - \mathbf{x}_{n,i}^T \hat{\theta}_{n|n-1,i-1}). \quad (9)$$

The role of the first-order update in (9) is to perturb the censoring slab towards the direction of $(\mathbf{x}_{n,i}, y_{n,i})$, thus making it less likely for future measurements conveying information “close to” $(\mathbf{x}_{n,i}, y_{n,i})$ to be retained. Intuitively speaking, such updates eliminate measurement redundancies and reduce estimation error due to the bias of uncensored observations. The AC sketching module is summarized as Algorithm 3, and when plugged into Algorithm 1, it yields the proposed adaptive censoring (AC)-KF scheme. With regards to its performance, we have the following result.

TABLE I
PER TIME-SLOT SKETCHING MODULE COMPLEXITY FOR $d < D$ DATA

Sketching method	Complexity
Random sampling	$\mathcal{O}(D)$
Random projections (RP-KF)	$\mathcal{O}(Dp \log d)$
Adaptive censoring (AC-KF)	$\mathcal{O}(Dp)$

Algorithm 3: AC sketching module.

Measurement selection with AC-LMS

Input: $\hat{\theta}_{n|n-1}, \{\mathbf{y}_n, \mathbf{X}_n, \mathbf{R}_n\}$

Initialization: $\hat{\theta}_{n|n-1,0} = \hat{\theta}_{n|n-1}, \mathcal{S}_{n,0} = \emptyset$

for $i = 1 : D$ **do**

 Obtain c_i as in (8)

if $c_i = 0$, **then**

$\mathcal{S}_{n,i} = \mathcal{S}_{n,i-1} \cup \{i\}$

 Update $\hat{\theta}_{n|n-1,i-1}$ as in (9)

end if

end for

Return: $\{\check{\mathbf{y}}_n, \check{\mathbf{X}}_n, \check{\mathbf{R}}_n\} = \{\mathbf{y}_{n,\mathcal{S}_n}, \mathbf{X}_{n,\mathcal{S}_n}, \mathbf{R}_{n,\mathcal{S}_n}\}$

Proposition 2: If $\mathbf{R}_n = \sigma_n^2 \mathbf{I}, \sigma_n^2 \gg \text{tr}(\mathbf{P}_{n|n-1})$, and $\mathbf{w}_n, \mathbf{v}_n$ are zero-mean Gaussian, then the AC-KF with $\mu = 0$ yields unbiased estimates $\forall \tau$.

Proof: See Appendix 2. ■

The assumptions in Proposition 2 were made to simplify the proof and are not necessary. Extensive simulations indicate that the AC-KF remains unbiased even for low σ_n^2 and correlated noise, and also for $\mu > 0$. Nevertheless, the variance of AC-KF largely depends on the choice of μ . Tuning μ to optimize the MSE performance of AC-KF is a challenging task. Accurate rules for selecting μ is part of our ongoing research. However, even for possibly suboptimal values of μ , the proposed scheme yields promising results. Simulations in Section VII will demonstrate that the proposed AC-KF attains estimation accuracy close to that of the KF using the greedy measurement selection method in [11]. In addition, the proposed sketching module performs a single pass over the data, and requires $\mathcal{O}(Dp)$ computations, which is markedly lower than the $\mathcal{O}(Ddp^2)$ required to perform greedy selection. Furthermore, AC-KF is suitable for online implementation by processing rows of \mathcal{I}_n sequentially. Table I summarizes the per-slot computational complexity of applying the sketching modules corresponding to AC-KF, RP-KF and random sampling.

Remark 1: Note that (9) does not pertain to a filter update. Instead, it is an update of $\hat{\theta}_{n|n-1,i}$ within the AC sketching module (Algorithm 3) that leverages uncensored entries per slot, and it is only used for censoring entries within \mathcal{I}_n of slot n . Note also that for $\mu > 0$ Algorithm 3 performs *adaptive* censoring on \mathcal{I}_n . If $\mu = 0$, then $\hat{\theta}_{n|n-1,i} = \hat{\theta}_{n|n-1} \forall i \in 1, \dots, D$, and thus the censoring rule in (8) becomes *non-adaptive* across measurements at slot n . For AC-KF with *non-adaptive* censoring, scalar entries of the measurement vector \mathbf{y}_n can also be censored in a decentralized fashion across distributed sensors.

V. UPDATE-SELECTION KF

In the last two sections, dimensionality reduction schemes were proposed for KF, by reducing the number of observations processed. The resulting algorithms can be used to reduce the complexity of filtering as well as other data-related costs such as storage and transmission by reducing the dimensionality of \mathcal{I}_n . Specifically, if the observations need to be transmitted from a remote location, one could be interested in reducing the communication overhead as well computations; the RP-KF (where no feedback is required between the filter and the sensors) or the AC-KF (that requires $\hat{\boldsymbol{\theta}}_{n|n-1}$ as feedback) would be preferable in such cases, with their dimensionality-reduction modules reducing the amount of data that need to be transmitted. On the other hand, if D observations become available to the computing platform (block-by block or entry-by-entry) at each time-slot, reduction of computational complexity is the main concern, as well as the focus of the present section.

Suppose that the observation noise has diagonal covariance matrix with $[\mathbf{R}_n]_{ii} = \sigma_i^2 \forall n$. Then, the KF correction step at time n can be obtained by solving

$$\hat{\boldsymbol{\theta}}_{n|n} = \arg \min_{\boldsymbol{\theta}} \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_{n|n-1}\|_{\mathbf{P}_{n|n-1}^{-1}}^2 + \sum_{i=1}^D \frac{1}{\sigma_i^2} (y_{n,i} - \mathbf{x}_{n,i}^T \boldsymbol{\theta})^2. \quad (10)$$

A sequential (across entries of \mathbf{y}_n) solution of (10) can also be obtained using the recursive least-squares (RLS) algorithm with parameter estimate and error covariance matrix initialized at $\hat{\boldsymbol{\theta}}_{n|n-1}$ and $\mathbf{P}_{n|n-1}$, respectively. Specifically, let $\hat{\boldsymbol{\theta}}_{n|n,i} := \mathbb{E}[\boldsymbol{\theta}_n | \hat{\boldsymbol{\theta}}_{n|n-1}, \mathbf{y}_{n,1:i}]$ and $\mathbf{P}_{n|n,i} := \text{cov}(\hat{\boldsymbol{\theta}}_{n|n,i})$, for $i \in \{0, \dots, D\}$; clearly, $\hat{\boldsymbol{\theta}}_{n|n,0} = \hat{\boldsymbol{\theta}}_{n|n-1}$ and $\hat{\boldsymbol{\theta}}_{n|n,D} = \hat{\boldsymbol{\theta}}_{n|n}$. The following RLS-like iteration, corresponding to the i -th entry of \mathbf{y}_n , updates the state estimate as

$$\hat{\boldsymbol{\theta}}_{n|n,i} = \hat{\boldsymbol{\theta}}_{n|n,i-1} + \mathbf{k}_{n,i} e_{n,i} \quad (11)$$

where $e_{n,i} := y_{n,i} - \mathbf{x}_{n,i}^T \hat{\boldsymbol{\theta}}_{n|n,i-1}$ and

$$\mathbf{k}_{n,i} = \mathbf{P}_{n|n,i-1} \mathbf{x}_{n,i} s_{n,i}^{-1} \quad (12)$$

with

$$s_{n,i} := \mathbf{x}_{n,i}^T \mathbf{P}_{n|n,i-1} \mathbf{x}_{n,i} + \sigma_i^2. \quad (13)$$

The state covariance matrix is then updated as

$$\mathbf{P}_{n|n,i} = \mathbf{P}_{n|n,i-1} - \mathbf{P}_{n|n,i-1} \mathbf{x}_{n,i} \mathbf{x}_{n,i}^T \mathbf{P}_{n|n,i-1} s_{n,i}^{-1} \quad (14)$$

and the process is repeated until $i = D$, and all the measurements have been processed.

A common approach to dealing with $D \gg 1$ is to simply process as many data within time slot n as the available computational resources allow for; see, e.g. [18, Ch. 7]. In the present work however, to reduce computational complexity, we propose *judiciously skipping correction updates*. The criterion according to which the i -th row of $\{\mathbf{y}_n, \mathbf{X}_n\}$ will be used to update $\hat{\boldsymbol{\theta}}_{n|n,i-1}$ is based on how much the distribution $p(\boldsymbol{\theta}_n | \hat{\boldsymbol{\theta}}_{n|n-1}, [\mathbf{y}_n]_{1:i})$ after the update will diverge from

the posterior $p(\boldsymbol{\theta}_n | \hat{\boldsymbol{\theta}}_{n|n-1}, [\mathbf{y}_n]_{1:i-1})$ before the update. A commonly used measure of difference between probability density functions (pdfs) is the Kullback-Leibler (KL) divergence, also known as relative entropy (see, e.g. [24]). The KL divergence between two pdfs $p(x)$ and $q(x)$ is defined as

$$\mathcal{D}_{KL}(p||q) := \int p(x) \ln \frac{p(x)}{q(x)} dx = \mathbb{E}_p \left[\ln \frac{p(x)}{q(x)} \right]$$

and is not symmetric with respect to its arguments. In fact, one can interpret $p(x)$ as being the ‘‘true’’ pdf of x while $q(x)$ is an approximate one. Then, $\mathcal{D}_{KL}(p||q)$ is a measure of how far the approximation is from reality.

Aiming at carrying out only ‘‘useful’’ updates, $\mathcal{D}_{KL}(p_{n,i}||p_{n,i-1})$ can be used as an indicator of how informative the update that involves the i -th row of $\{\mathbf{y}_n, \mathbf{X}_n\}$ is, where

$$p_{n,i}(\boldsymbol{\theta}_n) := p(\boldsymbol{\theta}_n | \hat{\boldsymbol{\theta}}_{n|n-1}, [\mathbf{y}_n]_{1:i})$$

and

$$p_{n,i-1}(\boldsymbol{\theta}_n) := p(\boldsymbol{\theta}_n | \hat{\boldsymbol{\theta}}_{n|n-1}, [\mathbf{y}_n]_{1:i-1}).$$

Proposition 3: Let observations be generated according to (1)–(2) with \mathbf{w}_n and \mathbf{v}_n Gaussian. Then, for the sequential estimator in (11)–(14) it holds that

$$\mathcal{D}_{KL}(p_{n,i}||p_{n,i-1}) = \frac{1}{2} (\bar{e}_{n,i}^2 - 1) \frac{\gamma_{n,i}}{\gamma_{n,i} + \sigma_i^2} + \ln \sqrt{\frac{\gamma_{n,i} + \sigma_i^2}{\sigma_i^2}}$$

where $\gamma_{n,i} := \mathbf{x}_{n,i}^T \mathbf{P}_{n|n,i-1} \mathbf{x}_{n,i}$, and $\bar{e}_{n,i} := e_{n,i} s_{n,i}^{-1/2}$ is the per-entry normalized innovation.

Proof: See Appendix 3. ■

Proposition 3 offers a simple expression of $\mathcal{D}_{KL}(p_{n,i}||p_{n,i-1})$ that will come handy in performing informative updates. Consider first the quantities involved in $\mathcal{D}_{KL}(p_{n,i}||p_{n,i-1})$, namely the normalized innovation (residual) $\bar{e}_{n,i}$, which is a random variable and $\gamma_{n,i}$ that is deterministic. Interestingly, $\gamma_{n,i} := \|\mathbf{x}_{n,i}\|_{\mathbf{P}_{n|n,i-1}}^2$ is an *ensemble* quantity capturing the expected power of the i -th observation across the main directions of state uncertainty, while $\bar{e}_{n,i}$ is a random data-dependent variable that measures how important the i -th update is for a specific realization of the problem. Depicted in Fig. 2(a) is a simulated sequence of $\mathcal{D}_{KL}(p_{n,i}||p_{n,i-1})$ as a function of index i for an arbitrary time-slot n . Immediately noticeable is that the per-step divergence decreases with an approximate rate of $1/i$ following the rate of decrease of $\mathbf{P}_{n|n,i}$. One may also observe that certain updates yield higher KL divergence compared to the baseline. Fast and computationally efficient estimation may be achieved by isolating and performing only such informative updates.

While $\mathcal{D}_{KL}(p_{n,i}||p_{n,i-1})$ offers a measure of difference between the posteriors prior and after updating, it lacks symmetry and it is not conveniently interpreted as distance. Consequently, we considered the modified metric

$$\mathcal{D}(p_{n,i}, p_{n,i-1}) := \mathcal{D}_{KL}(p_{n,i}||p_{n,i-1}) + \mathcal{D}_{KL}(p_{n,i}||p_{n,i-1}) \quad (15)$$

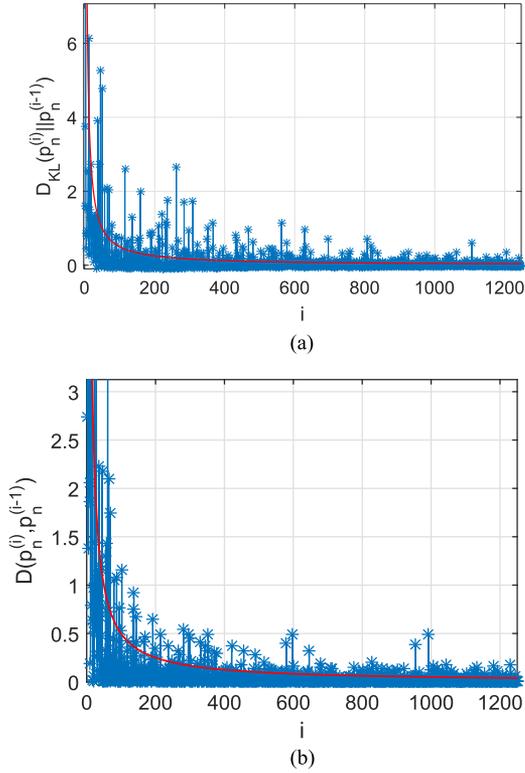


Fig. 2. Example of (a) KL divergence and (b) symmetric KL divergence evolution across sequential correction updates. Both metrics converge to 0 following a $\propto 1/i$ trend.

also known as the symmetric KL divergence. As seen in Figs. 2(a) and 2(b), both metrics follow a similar trend and converge to 0 as the state estimate converges in probability. Nevertheless, $\mathcal{D}(p_{n,i}, p_{n,i-1})$ enjoys symmetry as well as a more simple expression which as given in Proposition 4. Subsequently, the following rule is proposed for selecting informative updates

$$\mathcal{D}(p_{n,i} || p_{n,i-1}) \begin{cases} \geq \tau_{n,i}, & \text{Update } \hat{\boldsymbol{\theta}}_{n|i}, \mathbf{P}_{n|i}, \\ < \tau_{n,i}, & \text{Skip updates.} \end{cases} \quad (16)$$

Proposition 4: Let observations be generated according to (1)–(2) with \mathbf{w}_n and \mathbf{v}_n Gaussian. Then, for the sequential estimator in (11)–(14) it holds that

$$\mathcal{D}(p_{n,i}, p_{n,i-1}) = \frac{1}{2} \bar{c}_{n,i}^2 \left(2\gamma_{n,i} + \left(\frac{\gamma_{n,i}}{\sigma_i} \right)^2 \right) s_{n,i}^{-1}. \quad (17)$$

Proof: See Appendix 4. ■

Using (17), rule (16) can be readily implemented. Regarding the sequence of thresholds $\{\tau_{n,i}\}_{i=1}^D$, a judicious choice is

$$\tau_{n,i} = \tau_n \frac{1}{i} \quad (18)$$

which promotes updates with large informational value relative to the stage of the estimation process. The total number of updates per slot n can be tuned by τ_n . A couple of remarks are now in order.

Remark 2: KL divergence induced by a measurement was also employed by [17] to offer an alternative viewpoint

on a *distributed* censoring rule for reducing the *communication* load in WSNs. Specifically, it was shown that the KL divergence of $p(\boldsymbol{\theta}_n)$ with $p(\boldsymbol{\theta}_n | y_{n,i})$ as reference (i.e. $\mathcal{D}_{KL}(p(\boldsymbol{\theta}_n) || p(\boldsymbol{\theta}_n | y_{n,i}))$) is proportional to the magnitude of $e_{n,i}$, which implies that the latter is related to the informational value of a measurement. Apart from the different goals and context, a major difference of the present section's contribution relative to [17] is the explicit calculation, and use of the (symmetric) KL divergence in the proposed update selection rule.

Remark 3: Interestingly, our proposed *data-driven* update selection using $\mathcal{D}_{KL}(p_{n,i} || p_{n,i-1})$ is also related to OED-type sensor selection schemes that are based on the mutual information between a sensor and the model (e.g., [7]). This relation can be observed upon recalling that the mutual information $I(X; Y)$ between two random variables X and Y can be expressed as

$$I(X; Y) = \mathbb{E}_Y [\mathcal{D}_{KL}(p(X|Y) || p(X))]. \quad (19)$$

In the present context, (19) implies that the mutual information between the i -th “sensor” at time slot n and $\boldsymbol{\theta}_n$ in a sequential processing setting, equals its $\mathcal{D}_{KL}(p_{n,i} || p_{n,i-1})$ averaged over all possible measurements $y_{n,i}$.

A. Reduced-Complexity Censoring Rule

In the previous section, an update selection rule was introduced in (16) relying on the information metric in (17). Practical implementation of (16) requires careful consideration of the computational complexity needed to obtain $\mathcal{D}(p_{n,i}, p_{n,i-1})$. As seen in (17), to obtain the latter it suffices to compute $e_{n,i}$ and $\gamma_{n,i}$ (since $s_{n,i} = \gamma_{n,i} + \sigma_i^2$). While computing $e_{n,i}$ requires only $\mathcal{O}(p)$ products, obtaining $\gamma_{n,i} := \mathbf{x}_{n,i}^T \mathbf{P}_{n|i-1} \mathbf{x}_{n,i}$ requires a matrix-vector product that comes with $\mathcal{O}(p^2)$ complexity. Thus, even though checking whether an update is informative or not has smaller complexity than the update itself (cf. (11)–(14)), both tasks are of the same order of $\mathcal{O}(p^2)$ complexity. Ideally, checking the update should be less costly than performing the update by an order of magnitude.

For this purpose, a low-complexity approximation of $\gamma_{n,i}$ is highly desirable. One way to approximate $\gamma_{n,i}$ is to use the eigen-decomposition $\mathbf{P}_{n|i-1} = \mathbf{V}_{n,i-1} \boldsymbol{\Lambda}_{n,i-1} \mathbf{V}_{n,i-1}^T$ to produce the best k -rank approximation of $\mathbf{P}_{n|i-1}$ as

$$\hat{\mathbf{P}}_{n|i-1}^k = \mathbf{V}_{n,i-1}^k \boldsymbol{\Lambda}_{n,i-1}^k (\mathbf{V}_{n,i-1}^k)^T$$

where $\boldsymbol{\Lambda}_{n,i-1}^k$ is a $k \times k$ diagonal matrix containing the $k \leq p$ largest eigenvalues of $\mathbf{P}_{n|i-1}$, and $\mathbf{V}_{n,i-1}^k$ is the $p \times k$ matrix of the corresponding eigenvectors. Using $\hat{\mathbf{P}}_{n|i-1}^k$ to approximate $\gamma_{n,i}$, yields

$$\begin{aligned} g_{n,i}^k &:= \mathbf{x}_{n,i}^T \hat{\mathbf{P}}_{n|i-1}^k \mathbf{x}_{n,i} \\ &= \mathbf{x}_{n,i}^T \mathbf{V}_{n,i-1}^k \boldsymbol{\Lambda}_{n,i-1}^k (\mathbf{x}_{n,i}^T \mathbf{V}_{n,i-1}^k)^T \end{aligned} \quad (20)$$

which can be obtained with $\mathcal{O}(pk)$ complexity. Although $g_{n,i}^k$ exactly captures the power of $\mathbf{x}_{n,i}^T$ along the principal directions of $\mathbf{P}_{n|i-1}$, it is in general a poor estimate of $\gamma_{n,i}$. In fact, ignoring the power of the $p - k$ smallest eigenvalues of $\mathbf{P}_{n|i-1}$

leads to under-estimation; that is $g_{n,i}^k \leq \gamma_{n,i}^k$. To mitigate this problem, let us denote the power of $\mathbf{x}_{n,i}$ that is not captured by the first k principal directions as

$$g_{n,i}^{-k} := \|\mathbf{x}_{n,i}\|_2^2 - \|\mathbf{x}_{n,i}^T \mathbf{V}_{n,i-1}^k\|_2^2.$$

Then, if the the remaining power $g_{n,i}^{-k}$ is distributed evenly along the $p - k$ directions that correspond to the smallest eigenvalues of $\mathbf{P}_{n|i,i-1}$, an improved approximation of $\gamma_{n,i}^k$ is

$$\begin{aligned} \hat{\gamma}_{n,i}^k &= g_{n,i}^k + \frac{1}{p-k} g_{n,i}^{-k} \sum_{j=k+1}^p [\mathbf{\Lambda}_{n,i-1}]_{jj} \\ &= g_{n,i}^k + \frac{1}{p-k} g_{n,i}^{-k} [\text{tr}(\mathbf{P}_{n|i,i-1}) - \text{tr}(\mathbf{\Lambda}_{n,i-1}^k)] \end{aligned} \quad (21)$$

where for the second equality we used that $\text{tr}(\mathbf{P}_{n|i,i-1}) = \text{tr}(\mathbf{\Lambda}_{n|i,i-1})$. Essentially, by ignoring the angle of $\mathbf{x}_{n,i}^T$ along the $p - k$ least important directions of uncertainty, an estimate $\hat{\gamma}_{n,i}^k \approx \gamma_{n,i}^k$ can be found with $\mathcal{O}(pk)$ complexity. Simulations will demonstrate that for most cases k need not be very large for the purpose of obtaining a reliable approximation of $\gamma_{n,i}$, and thus of the update selection rule in (16).

Finally, let us consider the computational burden of re-computing the eigen-decomposition of $\mathbf{P}_{n|i,i-1}$ when an update is performed. Fortunately, the decomposition needs only be fully computed once for $\mathbf{P}_{n|n,0}$, after the prediction step. Then, exploiting that $\mathbf{P}_{n|i,i-1}$ is given as a sequence of rank-one updates of symmetric positive matrices (cf. (14)) allows for low-complexity $\mathcal{O}(p^2)$ updates of the eigen-decompositions, see e.g. [25] and [26], while the fact that only the first k eigen-pairs are required can further reduce the complexity of the updates. The need for tracking the principal eigen-pairs of $\mathbf{P}_{n|i,i-1}$ can be completely eliminated by setting $k = 0$, which yields the estimate

$$\hat{\gamma}_{n,i}^0 = \frac{1}{p} \|\mathbf{x}_{n,i}\|_2^2 \text{tr}(\mathbf{P}_{n|i,i-1}) \quad (22)$$

with $\mathcal{O}(p)$ complexity, at the cost of ignoring information given by the angle of $\mathbf{x}_{n,i}$. For cases where the eigenvalues of $\mathbf{P}_{n|i,i-1}$ are approximately uniform, (22) provides a practical and sufficiently accurate estimate of $\gamma_{n,i}$. Generally, obtaining (21) requires $\mathcal{O}(p(k+1))$ computations. Overall, the complexity of the correction step using the iterative method in (11)–(14) with the update selection rule in (16) and the approximation in (21) is $\mathcal{O}(dp^2) + \mathcal{O}(Dp(k+1))$, where $d \ll D$ is the number of updates. Depending on the size of p , d and D , the overall complexity of the proposed scheme can be considerably less than the standard $\mathcal{O}(Dp^2)$.

B. First-Order Updates

The fact that the update selection rule in (16) requires at least $\mathcal{O}(p)$ computations hints at possible modifications of the present scheme, that are considered in this section. Specifically, instead of using (16) to completely skip updates, one may incorporate a simple $\mathcal{O}(p)$ update without noticeably increasing the overall complexity of the algorithm. For instance, having computed $e_{n,i}$ which is required in (16), the following LMS-like parameter

update can readily be implemented

$$\hat{\boldsymbol{\theta}}_{n|i,i} = \hat{\boldsymbol{\theta}}_{n|i,i-1} + \mu_{n,i} \mathbf{x}_{n,i} e_{n,i} \quad (23)$$

where $\mu_{n,i}$ denotes a user selected stepsize. Given the update in (23), consider the difference

$$\Delta_n(\mu_{n,i}) := \text{MSE}(\hat{\boldsymbol{\theta}}_{n|i,i}) - \text{MSE}(\hat{\boldsymbol{\theta}}_{n|i,i-1}) \quad (24)$$

where $\text{MSE}(\boldsymbol{\theta}) := \mathbb{E}[\|\boldsymbol{\theta} - \boldsymbol{\theta}_n\|_2^2]$. Clearly, $\mu_{n,i}$ should be chosen such that $\Delta_n(\mu_{n,i}) \leq 0$, or, ideally such that $\Delta_n(\mu_{n,i})$ is minimized. But first, it is useful to derive an explicit expression for $\Delta_n(\mu_{n,i})$.

Proposition 5: For an update of the form (23) that follows an update of (11)–(14) it holds that

$$\Delta_n(\mu_{n,i}) = \|\mathbf{x}_{n,i}\|_2^2 (\gamma_{n,i} + \sigma_i^2) \mu_{n,i}^2 - 2\gamma_{n,i} \mu_{n,i}. \quad (25)$$

Proof: See Appendix 5. ■

To guarantee $\Delta_n(\mu_{n,i}) \leq 0$, it suffices to choose $\mu_{n,i}$ as

$$0 \leq \mu_{n,i} \leq \frac{2\gamma_{n,i}}{\|\mathbf{x}_{n,i}\|_2^2 (\gamma_{n,i} + \sigma_i^2)} \quad (26)$$

while for

$$\mu_{n,i}^* = \frac{\gamma_{n,i}}{\|\mathbf{x}_{n,i}\|_2^2 (\gamma_{n,i} + \sigma_i^2)} \quad (27)$$

the minimum of $\Delta_n(\cdot)$ is attained

$$\Delta_n(\mu_{n,i}^*) = -\frac{(\gamma_{n,i})^2}{\|\mathbf{x}_{n,i}\|_2^2 (\gamma_{n,i} + \sigma_i^2)}. \quad (28)$$

Although the LMS-like iteration (23) reduces the MSE by as much as $-\Delta_n(\mu_{n,i}^*)$, updating $\mathbf{P}_{n|i,i-1}$ incurs complexity $\mathcal{O}(p^2)$, and it is thus skipped. Skipping covariance updates for first-order updates is also well motivated by the fact that $-\Delta_n(\mu_{n,i}^*)$ is generally significantly smaller than the reduction achieved by the second-order updates (11)–(14). Nevertheless, one may in practice use $\mu_{n,i} \leq \mu_{n,i}^*$, to compensate for the (slow) decrease in estimation variance. Finally, while the exact value of $\gamma_{n,i}$ is generally not available, $\mu_{n,i}$ can be selected after using the estimate $\hat{\gamma}_{n,i}$ from (21) in (26) or (27). The overall proposed reduced-complexity update-selection (US) KF described in Section V is tabulated as Algorithm 4.

VI. BUDGETED FIXED-INTERVAL SMOOTHING

The methods introduced in Sections III, IV and V utilize dimensionality reduction, measurement selection, and update selection, in order to promote low-complexity correction updates of the KF. In the present section, we briefly explore another direction that allows for reliable tracking with smaller data usage and computational complexity. Specifically, consider the ‘‘smoothed’’ estimate $\hat{\boldsymbol{\theta}}_n^{\text{KS}} := \mathbb{E}[\boldsymbol{\theta}_n | \{\mathbf{y}_n\}_{n=1}^N]$, and let $\hat{\boldsymbol{\theta}}^{\text{KS}}$ be formed by concatenating all such smoothed estimates. This can also be written as

$$\begin{aligned} \hat{\boldsymbol{\theta}}^{\text{KS}} &= \arg \min_{\{\boldsymbol{\theta}_n\}_{n=1}^N} \frac{1}{2} \sum_{n=1}^N \|\mathbf{y}_n - \mathbf{X}_n \boldsymbol{\theta}_n\|_{\mathbf{R}_n}^2 \\ &\quad + \|\boldsymbol{\theta}_n - \mathbf{F}_n \boldsymbol{\theta}_{n-1}\|_{\mathbf{Q}_n}^2 + \|\boldsymbol{\theta}_0 - \mathbf{m}_0\|_{\mathbf{P}_0}^2 \end{aligned} \quad (29)$$

Algorithm 4: US-KF.

Initialization: $\hat{\theta}_{0|0} = \mathbf{m}_0$, $\mathbf{P}_{0|0} = \mathbf{P}_0$
for $n = 1 : N$ **do**
 Prediction Step:
 $\hat{\theta}_{n|n-1} = \mathbf{F}_n \hat{\theta}_{n-1|n-1} + \mathbf{G}_n \mathbf{u}_n$
 $\mathbf{P}_{n|n-1} = \mathbf{F}_n \mathbf{P}_{n-1|n-1} \mathbf{F}_n^T + \mathbf{Q}_n$
 Correction Step:
 Set parameters: k, τ_n
 Initialize: $\hat{\theta}_{n|n,0} = \hat{\theta}_{n|n-1}$, $\mathbf{P}_{n|n,0} = \mathbf{P}_{n|n-1}$
 Compute k first eigenpairs: $\{\mathbf{V}_{n,0}^k, \mathbf{\Lambda}_{n,0}^k\}$
 for $i = 1 : D$ **do**
 Obtain $\hat{\gamma}_{n,i}$ as in (21)
 Obtain $\mathcal{D}(p_{n,i}, p_{n,i-1})$ as in (17)
 if $\mathcal{D}(p_{n,i}, p_{n,i-1}) \geq \tau_n/i$ **then**
 Update $\{\hat{\theta}_{n|n,i}, \mathbf{P}_{n|n,i}\}$ as in (11)–(14)
 Update $\{\mathbf{V}_{n,i}^k, \mathbf{\Lambda}_{n,i}^k\}$ of rank-1 update of $\mathbf{P}_{n|n,i}$
 else
 Obtain $\mu_{n,i}$ by plugging $\hat{\gamma}_{n,i}$ in (26) or (27)
 Update $\hat{\theta}_{n|n,i}$ as in (23)
 end if
 end for
end for

which is optimal in the linear minimum mean-square error (LMMSE) sense.

Aiming at a recursive solver of (29), one can rely on the Rauch-Tung-Stribel (RTS) forward-backward KS algorithm [27]. In its forward pass, the RTS algorithm is identical to the KF. The KF estimates $\{\hat{\theta}_{n|n}\}_{n=1}^N$ are then stored and processed by the backward pass of the KS, while the error covariance matrices $\{\mathbf{P}_{n|n}\}_{n=1}^N$ are computed off-line.

Given $\hat{\theta}_{n+1|N}$, the backward iteration solves

$$\hat{\theta}_{n|N} := \arg \min_{\theta} \|\hat{\theta}_{n+1|N} - \mathbf{F}_n \theta\|_{\mathbf{Q}_n^{-1}}^2 + \|\theta - \hat{\theta}_{n|n}\|_{\mathbf{P}_{n|n}^{-1}}^2. \quad (30)$$

Similar to filtering, the minimizer of (30) is also given in closed form as

$$\hat{\theta}_{n|N} = \left(\mathbf{F}_n^T \mathbf{Q}_n^{-1} \mathbf{F}_n + \mathbf{P}_{n|n}^{-1} \right)^{-1} \left(\mathbf{F}_n^T \mathbf{Q}_n^{-1} \hat{\theta}_{n+1|N} + \mathbf{P}_{n|n}^{-1} \hat{\theta}_{n|n} \right).$$

After invoking the MIL and letting $\mathbf{B}_n := \mathbf{P}_{n|n} \mathbf{F}_n^T \mathbf{P}_{n+1|n}^{-1}$, the estimate $\hat{\theta}_{n|N}$ is given in the correction form of $\hat{\theta}_{n|n}$ as

$$\hat{\theta}_{n|N} = \hat{\theta}_{n|n} + \mathbf{B}_n \left(\hat{\theta}_{n+1|N} - \mathbf{F}_n \hat{\theta}_{n|n} \right) \quad (31)$$

with corresponding error covariance matrix

$$\mathbf{P}_{n|N} = \mathbf{P}_{n|n} + \mathbf{B}_n \left(\mathbf{P}_{n+1|N} - \mathbf{P}_{n+1|n} \right) \mathbf{B}_n^T. \quad (32)$$

A key property of the backward KS iteration, is that it improves KF performance using from $\{\mathcal{I}_n\}_{n=1}^N$ only the information encapsulated in the output $\hat{\theta}_{n|n}$ of the forward filter. Therefore, backward iterations can be readily applied on filtered estimates of RP-KF, AC-KF or the US-KF to limit the tracker's performance loss caused by the measurement reduction.

Algorithm 5: The budgeted Kalman smoother (Bud-KS).

for $n = N - 1 : 0$ **do**
 if $\hat{\theta}_{n|n} \in \Theta_n^S$ **then**
 $\hat{\theta}_{n|N} = \hat{\theta}_{n|n}$
 $\mathbf{P}_{n|N} = \mathbf{P}_{n|n}$
 else
 $\hat{\theta}_{n|N} = \hat{\theta}_{n|n} + \mathbf{B}_n \left(\hat{\theta}_{n+1|N} - \mathbf{F}_n \hat{\theta}_{n|n} \right)$
 $\mathbf{B}_n = \mathbf{P}_{n|n} \mathbf{F}_n^T \mathbf{P}_{n+1|n}^{-1}$
 $\mathbf{P}_{n|N} = \mathbf{P}_{n|n} + \mathbf{B}_n \left(\mathbf{P}_{n+1|N} - \mathbf{P}_{n+1|n} \right) \mathbf{B}_n^T$
 end if
end for

In addition, the backward iteration can also be modified to operate within a limited computational budget. Given the smoothed estimate at time $n + 1$, let us define the set

$$\Theta_n^b := \left\{ \theta \mid \|\hat{\theta}_{n+1|N} - \mathbf{F}_n \theta\|_{\mathbf{Q}_n^{-1}}^2 \leq \tau_b \right\} \quad (33)$$

of states at time n that are consistent enough with the transition model in the WLS sense. Based on (33), the Bud-KS estimate at time n is given as

$$\hat{\theta}_{n|N} = \begin{cases} \hat{\theta}_{n|n}, & \hat{\theta}_{n|n} \in \Theta_n^b \\ \hat{\theta}_{n|n} + \mathbf{B}_n \left(\hat{\theta}_{n+1|N} - \mathbf{F}_n \hat{\theta}_{n|n} \right), & \hat{\theta}_{n|n} \notin \Theta_n^b. \end{cases} \quad (34)$$

Clearly, for $\hat{\theta}_{n|n} \in \Theta_n^b$, it holds that $\mathbf{P}_{n|N} = \mathbf{P}_{n|n}$; while for $\hat{\theta}_{n|n} \notin \Theta_n^b$, the error covariance is given by (32). Essentially, KS estimates that are consistent enough with the system model are not smoothed, thus saving the computations required. Here, the threshold τ_b in (33) can be tuned to control the amount of ‘‘acceptable’’ deviation from the model. The novel economical, fixed-interval smoother on a budget, that we abbreviate as Bud-KS, is tabulated as Algorithm 5.

Regarding the computational complexity of Bud-KS, it is worth noting that implementing the rule (33) in the general case requires $\mathcal{O}(p^3)$ computations in order to invert \mathbf{Q}_n . The complexity of Bud-KS updates in (31) and (32) are on the same order of magnitude. Thus, Bud-KS is preferable when the covariance matrix of \mathbf{w}_n is time-invariant, meaning that $\mathbf{Q}_n = \mathbf{Q} \forall n$. In such cases, inversion of \mathbf{Q} is performed once offline, thus reducing complexity in (33) to $\mathcal{O}(p^2)$; likewise, when \mathbf{Q}_n is diagonal. In such scenarios, an update of $\mathcal{O}(p^3)$ complexity is skipped at the cost of an $\mathcal{O}(p^2)$ complexity rule, leading to computational savings that become more significant as p increases.

VII. NUMERICAL TESTS

The novel AC-KF, RP-KF, US-KF and Bud-KS algorithms are tested here on a simulated linear dynamical system. For this experiment, a simple state transition model that performs cyclical shifting of the entries of the state was implemented. The state transition matrix is

$$\mathbf{F}_{n,ij} \begin{cases} 1, & i = j - 1 \\ 0, & \text{otherwise} \end{cases}, \forall n$$

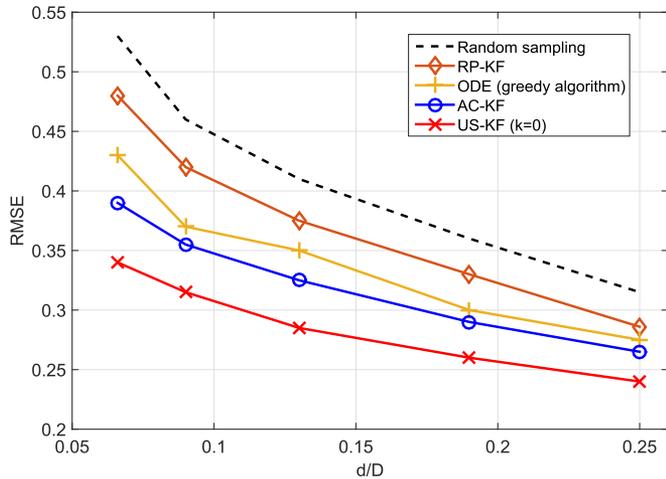


Fig. 3. Average RMSE for the US-KF, AC-KF, Greedy algorithm, RP-KF and random sampling as a function of d/D .

and $\mathbf{F}_{1p} = 1$, while the state dimension is set to $p = 50$. The state noise $\{\mathbf{w}_n\}_{n=1}^N$ was generated i.i.d. with $\mathbf{w}_n \sim \mathcal{N}(\mathbf{0}, \sigma_w^2 \mathbf{Q}_n)$, where $\mathbf{Q}_{n,i,j} = 0.5^{|i-j|}$ and $\sigma_w = 0.01$. Finally, the initial state is $\boldsymbol{\theta}_0 \sim \mathcal{N}(\mathbf{m}_0, \mathbf{P}_0)$, with \mathbf{m}_0 set to have two non-zero values 20 and -30 in its first and fifth entry, and $\mathbf{P}_0 = 0.04\mathbf{I}$. Per time instant $n \in \{1, \dots, N\}$ with $N = 100$, $D = 500$ measurements are obtained and concatenated in $\mathbf{y}_n = \mathbf{X}_n \boldsymbol{\theta}_n + \mathbf{v}_n$, where rows of \mathbf{X}_n are generated as i.i.d. standardized Gaussian vectors and then weighted independently by coefficients α drawn from $\alpha \sim \text{Unif}\{0.5, 1.5\}$. For this experiment, observations are correlated; thus, $\mathbf{v}_n \sim \mathcal{N}(\mathbf{0}, \sigma_v^2 \mathbf{R}_n)$, where $\mathbf{R}_{n,i,j} = 0.5^{|i-j|}$. For the following experiments, we set $\sigma_v^2 = 1$, upon observing that the results remain qualitatively similar for different noise levels.

A. AC-KF, RP-KF, and US-KF

To determine the average performance in terms of estimation error and computational complexity of AC-KF and RP-KF for different values of d/D , 20 Monte Carlo realizations were run on the same simulated linear dynamical system. The estimation performance was measured in terms of root mean-square error (RMSE) of the estimates across iterations; that is,

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{n=1}^N \|\hat{\boldsymbol{\theta}}_{n|n} - \boldsymbol{\theta}_n\|_2^2}.$$

AC-KF and US-KF were run first, with thresholds tuned such that a constant number of approximately d observations were selected per time slot; RP-KF and the greedy algorithm were then set to obtain d measurements per time slot. As a performance benchmark for the three algorithms, KF was also run with d randomly sampled observations per time step.

The average RMSE of the five methods as a function of d/D is plotted in Fig. 3. These plots confirm that the proposed data-agnostic RP-KF is useful for increasing the accuracy (compared to plain random sampling) when estimating dynamic processes. With regards to the more elaborate algorithms, AC-KF has comparable performance with the KF using greedy OED

TABLE II
AVERAGE RUNTIME OF ALGORITHMS FOR $D = 1000$

d/D	0.05	0.13	0.24
Random sampling	0.16 sec	0.31 sec	0.81 sec
RP-KF	0.26 sec	0.42 sec	1.3 sec
AC-KF	0.41 sec	0.51 sec	1.05 sec
US-KF	0.44 sec	0.64 sec	1.1 sec
Full-data KF	6.7 sec	6.7 sec	6.7 sec

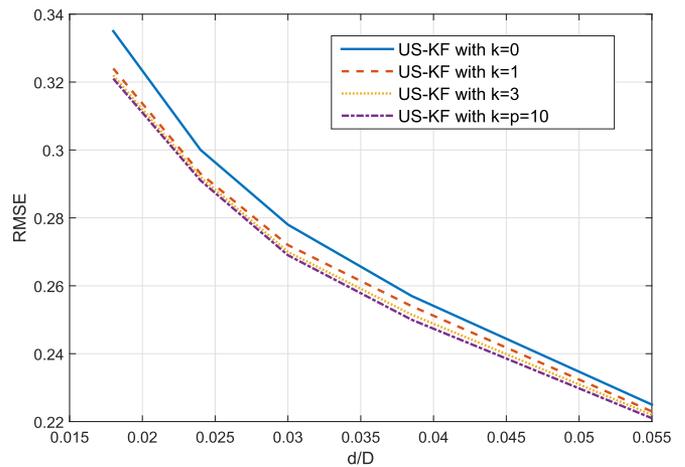


Fig. 4. Average RMSE for US-KF for different values of k .

measurement selection, while being orders of magnitude faster in terms of runtime. Last but not least, the US-KF with $k = 0$ outperforms the other methods while maintaining $\mathcal{O}(dp^2)$ complexity, even when the observation noise is correlated. Finally, the experiment was re-run with $D = 1000$ and for varying d/D , with the runtime of the algorithms listed in Table II. The greedy ODE algorithm is excluded from this experiment since it is an offline benchmark with runtime larger than that of the full-data KF. In comparison to random sampling, the proposed methods carry a certain computational overhead which becomes less significant as d/D (or D) increases. More importantly, the proposed algorithms enjoy a significantly lower runtime than the full-data KF.

Additional experiments were performed to assess sensitivity of the US-KF to the choice of parameter k . Recall that k determines the accuracy of the approximation of $\gamma_{n,i}$ (cf. (20)–(21)), and therefore how accurately the update selection rule in (16) is implemented; at the same time, the computational complexity of implementing (16) increases with k at a rate of $\mathcal{O}(p(k+1))$. Interestingly, experiments indicate that $k \ll p$ can be sufficient in practice, while sensitivity to k only manifests itself for relatively small values of the compression ratio d/D . As seen in Fig. 4, RMSE of US-KF with $k = 1$ is almost as low as the one achieved with $k = p$, while setting $k = 0$ still yields reliable estimates, with the gap becoming smaller as d/D increases. Recall that using $k = 0$ leads to the simple rule in (22), and bears the additional advantage that no eigenpairs of $\mathbf{P}_{n|n-1,i}$ need be tracked.

B. Bud-KS

In the last experiment, the extent to which backward smoothing iterations can improve reduced-observation filtering was examined. The AC-KF algorithm was first run with d/D

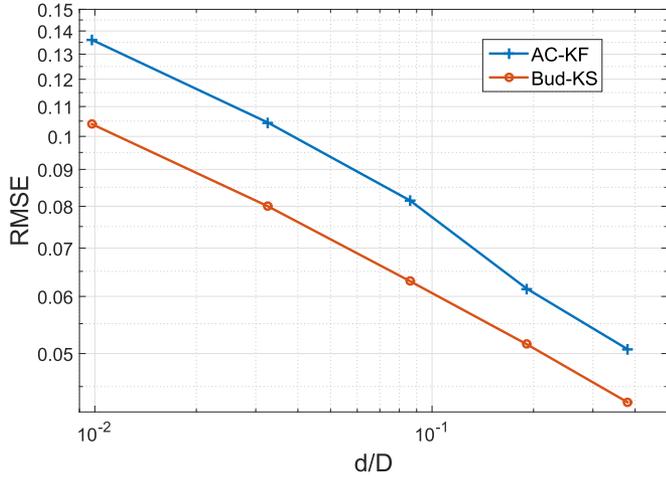


Fig. 5. RMSE of AC-KF versus Bud-KS, as a function of d/D .

ranging from 0.0095 up to 0.65; Bud-KF was then run with $\tau_b = 0$ in order to smooth all N filtered estimates. Fig. 5 depicts the average RMSE of the AC-KF with and without smoothing. Evidently, smoothing can significantly reduce RMSE over the entire range of dimensionality reduction, while its effect becomes more prominent as d/D decreases. Upon examining Fig. 5, the AC-KF using $< 1\%$ of the data followed by Bud-KS, attains the same RMSE as the AC-KF using 5% of the data; a surprising five-fold decrease. Thus, at the cost of introducing non-causality (or delay if a fixed-lag KS is used), smoothing offers room for significant decrease in the data requirements and complexity of tracking.

VIII. APPLICATION TO MONITORING DYNAMIC GRAPHS

Dynamically evolving graphs offer a promising application domain for our proposed algorithms. In this context, measurements are obtained from a graph of known and constant topology in order to infer a set of hidden time-varying properties. Specifically, *traffic matrix estimation* and *link cost estimation* are two tasks that involve tracking of large-scale dynamical processes from linearly obtained observations. To demonstrate the applicability of US-KF in reducing the complexity of such tasks, a Kronecker graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with $|\mathcal{V}| = 50$ vertices was generated. The adjacency matrix \mathbf{A}_k of a Kronecker graph can be generated recursively as $\mathbf{A}_k = \mathbf{A}_{k-1} \otimes \mathbf{A}_{k-1}$, and is completely determined by the initiator graph \mathbf{A}_1 . As shown in [28], Kronecker graphs exhibit many real-world graph properties such as power-law degree distributions, and are thus highly recommended for simulating algorithms. For our experiments, a Kronecker graph was generated with initiator

$$\mathbf{A}_1 = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{bmatrix}$$

until 100 nodes become available. Nodes adjacent to all other nodes were removed in order to decrease the connectivity of the graph to more realistic levels. The resulting adjacency matrix is depicted in Fig. 6.

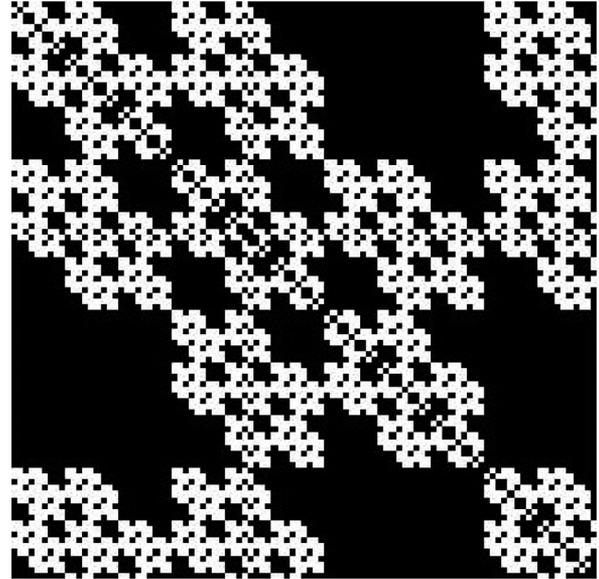


Fig. 6. Adjacency matrix of a Kronecker graph with 100 nodes.

A. Traffic Matrix Estimation

Consider the task of measuring the traffic volume at the links of a network, in order to estimate the volume of origin-to-destination (OD) flows, a very important task in many networks ranging from the Internet to transportation. Since OD flows are defined by a set of origins $\mathcal{O} \subseteq \mathcal{V}$ and a set of destinations $\mathcal{D} \subseteq \mathcal{V}$, they can be represented as the entries in an $|\mathcal{O}| \times |\mathcal{D}|$ traffic matrix \mathbf{F} . Similar to [29], [30] and [31], the following linear state-transition and observation models is considered

$$\mathbf{f}_n = \mathbf{f}_{n-1} + \mathbf{w}_n \quad (35)$$

$$\mathbf{l}_n = \mathbf{R}\mathbf{f}_n + \mathbf{v}_n \quad (36)$$

where $\mathbf{f}_n := \text{vec}(\mathbf{F}_n)$ is the vectorized traffic matrix at time slot n that is assumed to evolve according to a random walk with driving Gaussian noise \mathbf{w}_n with known covariance matrix $\sigma_f^2 \mathbf{Q}$ such that $\mathbf{Q}_{i,j} = 0.2^{|i-j|}$; \mathbf{l}_n contains the link measurements at time slot n ; and, \mathbf{v}_n is the observation noise with $\text{cov}(\mathbf{v}_n) = \sigma^2 \mathbf{I}$. The choice of a non-diagonal \mathbf{Q}_n was made to reflect the fact that flows tend to be highly correlated (see e.g. [2]). For this experiment, we set $\sigma_f = 0.02$, $\sigma = 0.5$, and generated the initial state as $\mathbf{f}_0 \sim \mathcal{N}(2 \cdot \mathbf{1}, \mathbf{Q})$. In this model, the role of the measurement matrix is played by the *routing matrix* $\mathbf{R} \in \{0, 1\}^{|\mathcal{E}| \times |\mathcal{O}| |\mathcal{D}|}$, each column of which corresponds to an OD flow with entries taking the value 1, if the corresponding links are part of the flow. Simply put, each column of \mathbf{R} describes the path that the corresponding OD flow takes through the graph. For this experiment, OD paths were chosen to be the shortest possible using Dijkstra's algorithm. To make this experiment even more challenging, flows with paths that consist of a single link were not considered; flows with no sampled links and irrelevant links were also removed from the model. Overall, 189 edges were sampled in order to track 689 OD flows.

Plotted in Fig. 7 is the MSE ($\mathbb{E}[\|\mathbf{f}_n - \hat{\mathbf{f}}_n\|_2^2]$) of the estimated traffic matrix across time, for the proposed US-KF (Alg. 4),

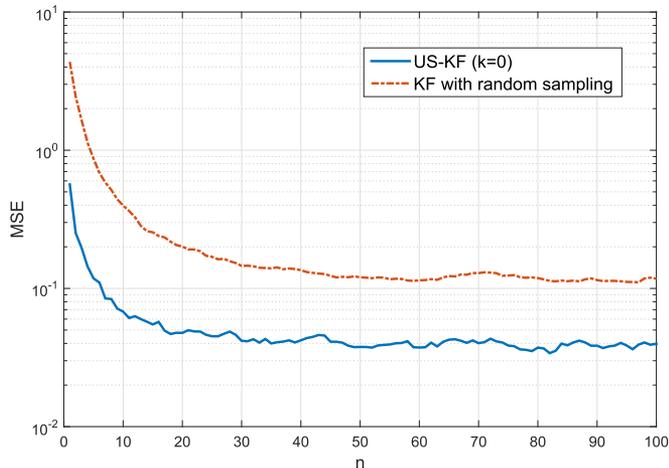


Fig. 7. Traffic matrix MSE vs time plot for the proposed Update selection KF (Alg.4) and the random sampling KF. Both algorithms were tuned to utilize 6% of edge measurements per time slot.

and the KF with random sampling. The algorithms were run for $N = 100$ time slots and the results were averaged across 100 runs. Both algorithms were tuned to utilize 6% of edge measurements per time slot, and require approximately the same runtime. As seen in the plot, the estimates $\hat{\mathbf{f}}_n$ of the proposed US-KF converge faster than those of the sub-sampled KF, and keep a closer track of the true traffic matrix \mathbf{f}_n . It should be noted that, due to the large state dimension, other methods such as the RP-KF or greedy OED become impractical.

B. Estimation of Link Costs From Path-Cost Measurements

Consider now that every edge ϵ of the graph is associated with a cost $c(\epsilon)$, and that the concatenation of all such costs forms the link cost vector \mathbf{c} . A common task associated with networks is inference of \mathbf{c} by measuring path costs p_{ij} , where p_{ij} is the total cost of a flow between nodes v_i and v_j (see e.g., [2, Ch. 9.4.1]). Since p_{ij} is the aggregation of all costs of the edges that the corresponding path crosses, it can be expressed as the inner product between \mathbf{c} and the corresponding row of the routing matrix. Consequently, path costs and link costs are linked through the linear observation model $\mathbf{p} = \mathbf{R}^T \mathbf{c}$, where \mathbf{p} is the vector with all the available path cost measurements. Considering dynamic graphs where the link costs \mathbf{c}_n and path costs \mathbf{p}_n evolve across time slots n , leads to the familiar linear state-transition and state-observation models

$$\mathbf{c}_n = \mathbf{c}_{n-1} + \mathbf{w}_n \quad (37)$$

$$\mathbf{p}_n = \mathbf{R}^T \mathbf{c}_n + \mathbf{v}_n \quad (38)$$

where $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \sigma_c^2 \mathbf{I})$, $\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$, and the initial state is $\mathbf{c}_0 \sim \mathcal{N}(\mathbf{m}, \sigma_0^2 \mathbf{I})$. For this experiment, we used the same graph and routing matrix as in the traffic estimation experiment, and generated \mathbf{c}_n and \mathbf{p}_n according to (37) and (38) correspondingly, with $\sigma_c = 0.04$, $\sigma = 0.1$, $\mathbf{m} = \mathbf{1}$ and $\sigma_0 = 0.1$.

Plotted in Fig. 8 is the MSE ($\mathbb{E} [\|\mathbf{c}_n - \hat{\mathbf{c}}_n\|_2^2]$) of the estimated link costs across time, for the proposed US-KF (Alg. 4) and the KF with random sampling, for $N = 100$ time slots and averaged

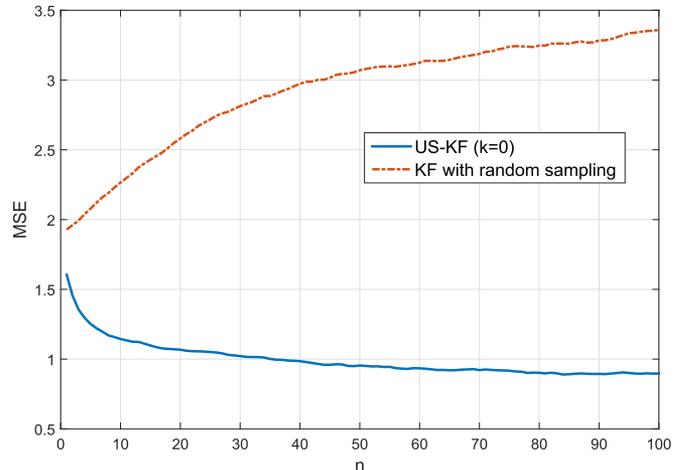


Fig. 8. Link-cost MSE vs time plot for US-KF (Alg.4) and the random sampling KF. Both algorithms were tuned to utilize 4% of flow measurements per time slot.

across 100 runs. Both algorithms were tuned to utilize 4% of path cost measurements per time slot, and require approximately the same runtime. As seen in the plot, the proposed US-KF successfully tracks the slowly evolving link costs by judiciously selecting and using a small fraction of the available path cost observations. Furthermore, it can be observed that if the same fraction (4%) of measurements is selected at random, then the KF fails to track the link costs, with its estimate diverging from the true value as time progresses. The divergence of the KF with random sampling is consistent with the results in [32], where it is shown that there exists a cut-off value for the data rate, below which the error covariance may become unbounded. Interestingly, the proposed reduced-complexity US-KF appears to be much more robust to divergence; as discussed in the following remark.

Remark 4: While KF based on random sampling (as well RP-KF) diverges when the compression ratio d/D becomes smaller than a certain threshold, this is not the case for the advocated censoring-based alternatives (AC-KF and US-KF) since diverging estimates prohibit censoring. This becomes evident upon realizing that a diverging estimate (i.e., $\|\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_n\|_2 \rightarrow \infty$) would imply infinitely large innovations that cannot be smaller than finite thresholds such as the ones used in censoring rules (7) and (16). This in turn implies that if AC-KF and US-KF were divergent, they would become equivalent to the full data KF. In a nutshell, if the ordinary KF is not divergent, the same holds for the proposed AC-KF and US-KF, since the latter will always obtain sets of observations that guarantee a bounded tracking error.

IX. CONCLUDING REMARKS

We introduced random projections and censoring as dimensionality reduction and measurement selection methods for tracking dynamical processes with generally time-varying parameters. The proposed methods are simple routines that can be used as dimensionality reduction modules coupled with an ordinary KF. Furthermore, we introduced a reduced-complexity

KF that processes measurements sequentially and performs updates that are deemed informative based on the information gain of corresponding measurements. Performance was not analytically performed, but simulations provide surprisingly strong evidence that the proposed methods perform close to the greedy measurement selection method in terms of estimation error. Furthermore, censoring-based measurement selection enjoys much lower computational complexity than greedy OED. To demonstrate applicability of the proposed update selection approach on real-world problems, we examined the network-related applications of traffic matrix estimation and network flow estimation.

APPENDIX

Proof of Proposition 1: Follows readily from [33, Th. 2]. ■

Proof of Proposition 2: From the assumption of large and uncorrelated noise $\mathbf{R}_n = \sigma_n^2 \mathbf{I}$, the inverse reduced innovation covariance matrix can be approximated as

$$(\tilde{\mathbf{X}}_n \mathbf{P}_{n|n-1} \tilde{\mathbf{X}}_n^T + \tilde{\mathbf{R}}_n)^{-1} \approx \sigma_n^{-2} \mathbf{I}$$

and hence the correction update as

$$\hat{\boldsymbol{\theta}}_{n|n} = \hat{\boldsymbol{\theta}}_{n|n-1} + \sigma_n^{-2} \mathbf{P}_{n|n-1} \sum_{i=1}^D \mathbf{x}_{n,i} \tilde{y}_{n,i} (1 - c_{n,i}) \quad (39)$$

where $\tilde{y}_{n,i} := \mathbf{x}_{n,i}^T (\boldsymbol{\theta}_n - \hat{\boldsymbol{\theta}}_{n|n-1}) + v_{n,i}$. Furthermore, for $\mu = 0$ the censoring rule in (8) simplifies to $1 - c_{n,i} = \mathbb{1}_{\{|\tilde{y}_{n,i}| \geq \tau_n \sigma_n\}}$, where $\tilde{y}_{n,i} \sim \mathcal{N}(0, \mathbf{x}_{n,i}^T \mathbf{P}_{n|n-1} \mathbf{x}_{n,i} + \sigma_n^2)$. If $\hat{\boldsymbol{\theta}}_{n-1|n-1}$ is unbiased, then it readily follows that $\hat{\boldsymbol{\theta}}_{n|n-1}$ is also unbiased, and (39) yields

$$\mathbb{E} [\hat{\boldsymbol{\theta}}_{n|n} - \boldsymbol{\theta}_n] = \sigma_n^{-2} \mathbf{P}_{n|n-1} \sum_{i=1}^D \mathbf{x}_{n,i} \mathbb{E} [\tilde{y}_{n,i} (1 - c_{n,i})]. \quad (40)$$

Since

$$\begin{aligned} \mathbb{E} [\tilde{y}_{n,i} (1 - c_{n,i})] &= \mathbb{E} [\tilde{y}_{n,i} \mathbb{1}_{\{|\tilde{y}_{n,i}| \geq \tau_n \sigma_n\}}] \\ &= \mathbb{E} [\tilde{y}_{n,i}] - \mathbb{E} [\tilde{y}_{n,i} \mathbb{1}_{\{|\tilde{y}_{n,i}| \leq \tau_n \sigma_n\}}] \\ &\propto - \int_{-\tau_n \sigma_n}^{\tau_n \sigma_n} \tilde{y}_{n,i} e^{-c(\tilde{y}_{n,i})^2} d(\tilde{y}_{n,i}) \\ &= \frac{1}{2c} \left(e^{-c\tau_n^2 \sigma_n^2} - e^{-c(-\tau_n \sigma_n)^2} \right) = 0 \end{aligned} \quad (41)$$

where $c := 0.5(\mathbf{x}_{n,i}^T \mathbf{P}_{n|n-1} \mathbf{x}_{n,i} + \sigma_n^2)^{-1}$, it follows from (41) and (40) that $\mathbb{E} [\hat{\boldsymbol{\theta}}_{n|n} - \boldsymbol{\theta}_n] = \mathbf{0}$, and the AC-KF is unbiased. ■

Proof of Proposition 3: For observations generated according to the linear Gaussian model, and since $\hat{\boldsymbol{\theta}}_{n|n,i}$ is the MMSE estimator of $\boldsymbol{\theta}_n$ given $\hat{\boldsymbol{\theta}}_{n|n-1}$ and $\mathbf{y}_{n,1:i}$, it follows that the posterior of $\boldsymbol{\theta}_n$ is also Gaussian with $p_{n,i}(\boldsymbol{\theta}_n) = \mathcal{N}(\hat{\boldsymbol{\theta}}_{n|n,i}, \mathbf{P}_{n|n,i})$. Similarly, one can obtain $p_{n,i-1}(\boldsymbol{\theta}_n) = \mathcal{N}(\hat{\boldsymbol{\theta}}_{n|n,i-1}, \mathbf{P}_{n|n,i-1})$. Using the closed-form identity for the KL divergence between

two multivariate normal pdfs, we arrive at

$$\begin{aligned} \mathcal{D}_{KL}(p_{n,i} || p_{n,i-1}) &= \frac{1}{2} \left[\text{tr} \left(\mathbf{P}_{n|n,i-1}^{-1} \mathbf{P}_{n|n,i} \right) \right. \\ &\quad + (\hat{\boldsymbol{\theta}}_{n|n,i-1} - \hat{\boldsymbol{\theta}}_{n|n,i})^T \mathbf{P}_{n|n,i-1}^{-1} \\ &\quad \times (\hat{\boldsymbol{\theta}}_{n|n,i-1} - \hat{\boldsymbol{\theta}}_{n|n,i}) \\ &\quad \left. - p + \ln \left(\frac{|\mathbf{P}_{n|n,i-1}|}{|\mathbf{P}_{n|n,i}|} \right) \right] \end{aligned} \quad (42)$$

where $\text{tr}(\mathbf{P})$ denotes the trace of matrix \mathbf{P} and $|\mathbf{P}|$ its determinant.

Using (14), the first summand in (42) can be expressed as

$$\begin{aligned} &\text{tr} \left(\mathbf{P}_{n|n,i-1}^{-1} \mathbf{P}_{n|n,i} \right) \\ &= \text{tr} \left(\mathbf{I}_p - \mathbf{x}_{n,i} \mathbf{x}_{n,i}^T \mathbf{P}_{n|n,i-1} \mathbf{s}_{n,i}^{-1} \right) \\ &= p - \mathbf{x}_{n,i}^T \mathbf{P}_{n|n,i-1} \mathbf{x}_{n,i} \mathbf{s}_{n,i}^{-1}. \end{aligned} \quad (43)$$

Upon observing that for the RLS-like iteration in (11) the inverse of the covariance matrix is updated as

$$\mathbf{P}_{n|n,i}^{-1} = \mathbf{P}_{n|n,i-1}^{-1} + \mathbf{x}_{n,i} \mathbf{x}_{n,i}^T \mathbf{s}_{n,i}^{-2} \quad (44)$$

the fourth summand in (42) can be expressed as

$$\begin{aligned} \ln \left(\frac{|\mathbf{P}_{n|n,i-1}|}{|\mathbf{P}_{n|n,i}|} \right) &= \ln (|\mathbf{P}_{n|n,i-1}|) + \ln \left(|\mathbf{P}_{n|n,i}^{-1}| \right) \\ &= \ln (|\mathbf{P}_{n|n,i-1}|) \\ &\quad + \ln \left(|\mathbf{P}_{n|n,i-1}^{-1} + \mathbf{x}_{n,i} \mathbf{x}_{n,i}^T \mathbf{s}_{n,i}^{-2}| \right) \\ &= \ln (|\mathbf{P}_{n|n,i-1}|) \\ &\quad + \ln \left(|\mathbf{P}_{n|n,i-1}^{-1} + \mathbf{x}_{n,i} \mathbf{x}_{n,i}^T \mathbf{s}_{n,i}^{-2}| \right) \\ &= \ln (|\mathbf{P}_{n|n,i-1}|) \\ &\quad + \ln \left(|\mathbf{P}_{n|n,i-1}^{-1}| (1 + \mathbf{x}_{n,i}^T \mathbf{P}_{n|n,i-1} \mathbf{x}_{n,i} \mathbf{s}_{n,i}^{-2}) \right) \\ &= \ln (1 + \mathbf{x}_{n,i}^T \mathbf{P}_{n|n,i-1} \mathbf{x}_{n,i} \mathbf{s}_{n,i}^{-2}) \\ &= \ln(s_{n,i}) - \ln(\mathbf{s}_{n,i}^2) \end{aligned} \quad (45)$$

where in the first equality we used the fact that $|\mathbf{P}^{-1}| = 1/|\mathbf{P}|$, and in the fourth one we applied the matrix determinant lemma for rank-one updates.

Finally, since $\boldsymbol{\theta}_{n|n,i-1} - \hat{\boldsymbol{\theta}}_{n|n,i} = -\mathbf{k}_{n,i} e_{n,i}$, the second summand in (42) becomes

$$\begin{aligned} &(\hat{\boldsymbol{\theta}}_{n|n,i-1} - \hat{\boldsymbol{\theta}}_{n|n,i})^T \mathbf{P}_{n|n,i-1}^{-1} (\hat{\boldsymbol{\theta}}_{n|n,i-1} - \hat{\boldsymbol{\theta}}_{n|n,i}) \\ &= (\mathbf{k}_{n,i} e_{n,i})^T \mathbf{P}_{n|n,i-1}^{-1} \mathbf{k}_{n,i} e_{n,i} \\ &= e_{n,i}^2 \mathbf{x}_{n,i}^T \mathbf{P}_{n|n,i-1} \mathbf{x}_{n,i} \mathbf{s}_{n,i}^{-2}. \end{aligned} \quad (46)$$

Substituting (43)–(46) into (42) and with $\gamma_{n,i} = \mathbf{x}_{n,i}^T \mathbf{P}_{n|n,i-1} \mathbf{x}_{n,i}$, we arrive at the result of Proposition 3. ■

Proof of Proposition 4: By the definition of $\mathcal{D}(p_{n,i}, p_{n,i-1})$ in (15) and expressing $\mathcal{D}_{KL}(p_{n,i-1} || p_{n,i})$ using arguments similar

to (43) and (45), it follows that

$$\begin{aligned} \mathcal{D}(p_{n,i}, p_{n,i-1}) &= \frac{1}{2} (\hat{\boldsymbol{\theta}}_{n|n,i-1} - \hat{\boldsymbol{\theta}}_{n|n,i})^T \\ &\quad \times \left(\mathbf{P}_{n|n,i-1}^{-1} + \mathbf{P}_{n|n,i}^{-1} \right) \\ &\quad \times (\hat{\boldsymbol{\theta}}_{n|n,i-1} - \hat{\boldsymbol{\theta}}_{n|n,i}). \end{aligned} \quad (47)$$

Utilizing (44) and that $\hat{\boldsymbol{\theta}}_{n|n,i-1} - \hat{\boldsymbol{\theta}}_{n|n,i} = -\mathbf{k}_{n,i} e_{n,i}$ yields

$$\begin{aligned} \mathcal{D}(p_{n,i}, p_{n,i-1}) &= \frac{1}{2} e_{n,i}^2 \mathbf{k}_{n,i}^T \\ &\quad \times \left(2\mathbf{P}_{n|n,i-1}^{-1} + \mathbf{x}_{n,i} \mathbf{x}_{n,i}^T \sigma_i^{-2} \right) \mathbf{k}_{n,i} \\ &= \frac{1}{2} \frac{e_{n,i}^2}{s_{n,i}^2} \left(\mathbf{x}_{n,i}^T \mathbf{P}_{n|n,i-1} \mathbf{x}_{n,i} \right. \\ &\quad \left. + (\mathbf{x}_{n,i}^T \mathbf{P}_{n|n,i-1} \mathbf{x}_{n,i})^2 \sigma_i^{-2} \right) \end{aligned}$$

and since $\gamma_{n,i} := \mathbf{x}_{n,i}^T \mathbf{P}_{n|n,i-1} \mathbf{x}_{n,i}$ and $\bar{e}_{n,i} := e_{n,i} s_{n,i}^{-1/2}$ the proposition holds. ■

Proof of Proposition 5: Recalling that $e_{n,i} := y_{n,i} - \mathbf{x}_{n,i}^T \hat{\boldsymbol{\theta}}_{n|n,i-1}$ and $y_{n,i} = \mathbf{x}_{n,i}^T \boldsymbol{\theta}_n + v_{n,i}$, (23) yields

$$\begin{aligned} \hat{\boldsymbol{\theta}}_{n|n,i} &= \hat{\boldsymbol{\theta}}_{n|n,i-1} - \mu_{n,i} \mathbf{x}_{n,i} \mathbf{x}_{n,i}^T (\hat{\boldsymbol{\theta}}_{n|n,i-1} - \boldsymbol{\theta}_n) \\ &\quad + \mu_{n,i} \mathbf{x}_{n,i} v_{n,i}. \end{aligned} \quad (48)$$

With $\tilde{\boldsymbol{\theta}}_{n,i} := \hat{\boldsymbol{\theta}}_{n|n,i} - \boldsymbol{\theta}_n$ denoting the error vector, (48) can be expressed as

$$\tilde{\boldsymbol{\theta}}_{n,i} = (\mathbf{I}_p - \mu_{n,i} \mathbf{x}_{n,i} \mathbf{x}_{n,i}^T) \tilde{\boldsymbol{\theta}}_{n,i-1} + \mu_{n,i} \mathbf{x}_{n,i} v_{n,i}. \quad (49)$$

The outer product of both sides in (49) yields

$$\begin{aligned} \tilde{\boldsymbol{\theta}}_{n,i} \tilde{\boldsymbol{\theta}}_{n,i}^T &= (\mathbf{I}_p - \mu_{n,i} \mathbf{x}_{n,i} \mathbf{x}_{n,i}^T) \tilde{\boldsymbol{\theta}}_{n,i-1} \\ &\quad \times \tilde{\boldsymbol{\theta}}_{n,i-1}^T (\mathbf{I}_p - \mu_{n,i} \mathbf{x}_{n,i} \mathbf{x}_{n,i}^T) \\ &\quad + 2 (\mathbf{I}_p - \mu_{n,i} \mathbf{x}_{n,i} \mathbf{x}_{n,i}^T) \tilde{\boldsymbol{\theta}}_{n,i-1} \mu_{n,i} \mathbf{x}_{n,i}^T v_{n,i} \\ &\quad + (\mu_{n,i})^2 \mathbf{x}_{n,i} \mathbf{x}_{n,i}^T (v_{n,i})^2. \end{aligned} \quad (50)$$

Since $\hat{\boldsymbol{\theta}}_{n|n,i-1}$ is unbiased, it follows that $\hat{\boldsymbol{\theta}}_{n|n,i}$ is unbiased too, and therefore the MSE equals the trace of the covariance matrix. Since the expected value of the second summand in (50) is zero, the trace of the expectation in (50) yields

$$\begin{aligned} \text{tr}(\bar{\mathbf{P}}_{n|n,i}) &= \text{tr} \left((\mathbf{I}_p - \mu_{n,i} \mathbf{x}_{n,i} \mathbf{x}_{n,i}^T)^2 \mathbf{P}_{n|n,i-1} \right) \\ &\quad + \mu_{n,i}^2 \|\mathbf{x}_{n,i}\|_2^2 \sigma_i^2 \\ &= \text{tr}(\mathbf{P}_{n|n,i-1}) + \mu_{n,i}^2 \text{tr}((\mathbf{x}_{n,i} \mathbf{x}_{n,i}^T)^2 \mathbf{P}_{n|n,i-1}) \\ &\quad - 2\mu_{n,i} \text{tr}(\mathbf{x}_{n,i} \mathbf{x}_{n,i}^T \mathbf{P}_{n|n,i-1}) + \mu_{n,i}^2 \|\mathbf{x}_{n,i}\|_2^2 \sigma_i^2. \end{aligned} \quad (51)$$

where $\bar{\mathbf{P}}_{n|n,i}$ is the covariance matrix after the first-order update in (23). Given that $\Delta_n(\mu_{n,i}) := \text{tr}(\bar{\mathbf{P}}_{n|n,i}) - \text{tr}(\mathbf{P}_{n|n,i-1})$, and upon observing that $\text{tr}(\mathbf{x}_{n,i} \mathbf{x}_{n,i}^T \mathbf{P}_{n|n,i-1}) = \gamma_{n,i}$ and $\text{tr}((\mathbf{x}_{n,i} \mathbf{x}_{n,i}^T)^2 \mathbf{P}_{n|n,i-1}) = \|\mathbf{x}_{n,i}\|_2^2 \gamma_{n,i}$, the proof is complete after using (51). ■

REFERENCES

- [1] P. Louka *et al.*, "Improvements in wind speed forecasts for wind power prediction purposes using Kalman filtering," *J. Wind Eng. Ind. Aero.*, vol. 96, no. 12, pp. 2348–2362, Nov. 2008.
- [2] E. D. Kolaczyk, *Statistical Analysis of Network Data*. Berlin, Germany: Springer, 2009.
- [3] H. Zhu, I. D. Schizas, and G. B. Giannakis, "Power-efficient dimensionality reduction for distributed channel-aware Kalman tracking using WSNs," *IEEE Trans. Signal Process.*, vol. 57, no. 8, pp. 3193–3207, Aug. 2009.
- [4] H. Ma, Y.-H. Yang, Y. Chen, K. R. Liu, and Q. Wang, "Distributed state estimation with dimension reduction preprocessing," *IEEE Trans. Signal Process.*, vol. 62, no. 12, pp. 3098–3110, Dec. 2014.
- [5] M. P. Vitus, W. Zhang, A. Abate, J. Hu, and C. J. Tomlin, "On efficient sensor scheduling for linear dynamical systems," *Automatica*, vol. 48, no. 10, pp. 2482–2493, 2012.
- [6] H. Wang, K. Yao, G. Pottie, and D. Estrin, "Entropy-based sensor selection heuristic for target localization," in *Proc. Int. Symp. Inf. Process. Sensor Netw.*, Berkeley, California, 2004, pp. 36–45.
- [7] E. Ertin, J. W. Fisher, and L. C. Potter, "Maximum mutual information principle for dynamic sensor query problems," in *Proc. Inf. Process. Sensor Netw.*, Springer, 2003, pp. 405–416.
- [8] L. Zuo, R. Niu, and P. K. Varshney, "Posterior CRLB based sensor selection for target tracking in sensor networks," in *Proc. Int. Conf. Acoust., Speech Signal Proc.*, vol. 2, Honolulu, Hawaii, 2007, pp. II–1041.
- [9] S. Joshi and S. Boyd, "Sensor selection via convex optimization," *IEEE Trans. Signal Process.*, vol. 57, no. 2, pp. 451–462, Feb. 2009.
- [10] F. Altenbach, S. Corroy, G. Bocherer, and R. Mathar, "Strategies for distributed sensor selection using convex optimization," in *Proc. GLOBECOM Conf.*, Anaheim, CA, 2012, pp. 2367–2372.
- [11] M. Shamaiah, S. Banerjee, and H. Vikalo, "Greedy sensor selection: Leveraging submodularity," in *Proc. 49th Conf. Decis. Control*, Atlanta, Dec. 2010, pp. 2572–2577.
- [12] N. Moshtagh, L. Chen, and R. Mehra, "Optimal measurement selection for any-time Kalman filtering with processing constraints," in *Proc. 48th Conf. Decis. Control*, Shanghai, China, 2009, pp. 5074–5079.
- [13] E. Msechu and G. B. Giannakis, "Sensor-centric data reduction for estimation with WSNs via censoring and quantization," *IEEE Trans. Signal Process.*, vol. 60, no. 1, pp. 400–414, Jan. 2012.
- [14] K. You, L. Xie, and S. Song, "Asymptotically optimal parameter estimation with scheduled measurements," *IEEE Trans. Signal Process.*, vol. 61, no. 14, pp. 3521–3531, Jul. 2013.
- [15] Q. Liu, Z. Wang, X. He, and D. Zhou, "A survey of event-based strategies on control and estimation," *Syst. Sci. Control Eng. Open Access J.*, vol. 2, no. 1, pp. 90–97, 2014.
- [16] G. Battistelli, A. Benavoli, and L. Chisci, "Data-driven strategies for selective data transmission in sensor networks," in *Proc. 51st Conf. Decis. Control*, Grand Wailea, Maui, 2012, pp. 800–805.
- [17] Y. Zheng, R. Niu, and P. K. Varshney, "Sequential Bayesian estimation with censored data for multi-sensor systems," *IEEE Trans. Signal Process.*, vol. 62, no. 10, pp. 2626–2641, Oct. 2014.
- [18] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation With Applications to Tracking and Navigation: Theory, Algorithms and Software*. Hoboken, NJ, USA: Wiley, 2004.
- [19] M. Mahoney, "Randomized algorithms for matrices and data," *Found. Trends. Mach. Learn.*, vol. 3, no. 2, pp. 123–224, 2011.
- [20] P. Drineas, M. W. Mahoney, and S. Muthukrishnan, "Sampling algorithms for 1 2 regression and applications," in *Proc. 17th SIAM-ACM Symp. Discrete Alg.*, Miami, Florida, 2006, pp. 1127–1136.
- [21] D. P. Woodruff, "Sketching as a tool for numerical linear algebra," *Found. Trends Theor. Comput. Sci.*, vol. 10, pp. 1–157, 2014.
- [22] Y. Wang, V. Krishnaswami, and G. Rizzoni, "Event-based estimation of indicated torque for IC engines using sliding-mode observers," *Control Eng. Pract.*, vol. 5, no. 8, pp. 1123–1129, 1997.
- [23] D. Berberidis, V. Kekatos, and G. B. Giannakis, "Online censoring for large-scale regressions with application to streaming big data," *IEEE Trans. Signal Process.*, vol. 64, no. 15, pp. 3854–3867, Aug. 2016.
- [24] S. Kullback and R. A. Leibler, "On information and sufficiency," *Annals Math. Stat.*, vol. 22, no. 1, pp. 79–86, 1951.
- [25] K.-B. Yu, "Recursive updating the eigenvalue decomposition of a covariance matrix," *IEEE Trans. Signal Process.*, vol. 39, no. 5, pp. 1136–1145, May 1991.
- [26] M. Gu and S. C. Eisenstat, "A stable and efficient algorithm for the rank-one modification of the symmetric eigenproblem," *SIAM J. Matrix Anal. Appl.*, vol. 15, no. 4, pp. 1266–1276, 1994.

- [27] H. E. Rauch, C. Striebel, and F. Tung, "Maximum likelihood estimates of linear dynamic systems," *AIAA J.*, vol. 3, no. 8, pp. 1445–1450, 1965.
- [28] J. Leskovec, D. Chakrabarti, J. Kleinberg, C. Faloutsos, and Z. Ghahramani, "Kronecker graphs: An approach to modeling networks," *J. Mach. Learn. Res.*, vol. 11, pp. 985–1042, 2010.
- [29] K. Rajawat, E. Dall'Anese, and G. B. Giannakis, "Dynamic network delay cartography," *IEEE Trans. Inf. Theory*, vol. 60, no. 5, pp. 2910–2920, May 2014.
- [30] A. Soule, K. Salamatian, A. Nucci, and N. Taft, "Traffic matrix tracking using Kalman filters," *SIGMETRICS Perform. Eval. Rev.*, vol. 33, no. 3, pp. 24–31, Dec. 2005.
- [31] P. Casas, S. Vaton, L. Fillatre, and L. Chonavel, "Efficient methods for traffic matrix modeling and on-line estimation in large-scale IP networks," in *Proc. Int. Teletraffic Congr.*, Paris, France, Sep. 2009, pp. 1–8.
- [32] B. Sinopoli, L. Schenato, M. Franceschetti, K. Poolla, M. I. Jordan, and S. S. Sastry, "Kalman filtering with intermittent observations," *IEEE Trans. Autom. Control*, vol. 49, no. 9, pp. 1453–1464, Sep. 2004.
- [33] P. Drineas, M. W. Mahoney, S. Muthukrishnan, and T. Sarlós, "Faster least squares approximation," *Numer. Math.*, vol. 117, no. 2, pp. 219–249, 2011.



Dimitris Berberidis (S'15) received the Diploma in electrical and computer engineering (ECE) from the University of Patras, Patras, Greece, in 2012 and the M.Sc. degree in ECE from the University of Minnesota, Minneapolis, MN, USA, where he is currently working towards the Ph.D. degree. His research interests include statistical signal processing, big data analytics, tracking, active learning, and adaptive sampling.



Georgios B. Giannakis (F'97) received the Diploma in electrical engineering from the National Technical University of Athens, Zografou, Greece, in 1981. From 1982 to 1986, he was with the University of Southern California (USC), where he received the M.Sc. degree in electrical engineering, in 1983, the M.Sc. degree in mathematics, in 1986, and the Ph.D. degree in electrical engineering, in 1986. He was with the University of Virginia from 1987 to 1998, and since 1999 he has been a professor with the University of Minnesota, where he holds an Endowed Chair in Wireless Telecommunications, a University of Minnesota McKnight Presidential Chair in ECE, and serves as director of the Digital Technology Center.

His general interests span the areas of communications, networking, and statistical signal processing on which he has published more than 400 journal papers, 700 conference papers, 25 book chapters, two edited books, and two research monographs (h-index 125). His current research focuses on learning from big data, wireless cognitive radios, and network science with applications to social, brain, and power networks with renewables. He is the (co-) inventor of 30 patents issued, and the (co-) recipient of 8 best paper awards from the IEEE Signal Processing (SP) and Communications Societies, including the G. Marconi Prize Paper Award in Wireless Communications. He also received Technical Achievement Awards from the SP Society (2000), from EURASIP (2005), a Young Faculty Teaching Award, the G. W. Taylor Award for Distinguished Research from the University of Minnesota, and the IEEE Fourier Technical Field Award (2015). He is a Fellow of EURASIP, and has served the IEEE in a number of posts, including that of a Distinguished Lecturer for the IEEE-SP Society.