

# Approximate Kalman filtering for large-scale systems with an application to hyperthermia cancer treatments

S.A.N. Nouwens<sup>1</sup>, B. de Jager<sup>1</sup>, M.M. Paulides<sup>2,3</sup>, W.P.M.H. Heemels<sup>1</sup>

**Abstract**—Accurate state estimates are required for increasingly complex systems, to enable, for example, feedback control. However, available state estimation schemes are not necessarily real-time feasible for certain large-scale systems. Therefore, we develop in this paper, a real-time feasible state-estimation scheme for a class of large-scale systems that approximates the steady state Kalman filter. In particular, we focus on systems where the state-vector is the result of discretizing the spatial domain, as typically seen in Partial Differential Equations. In such cases, the correlation between states in the state-vector often have an intuitive interpretation on the spatial domain, which can be exploited to obtain a significant reduction in computational complexity, while still providing accurate state estimates. We illustrate these strengths of our method through a hyperthermia cancer treatment case study. The results of the case study show significant improvements in the computation time, while simultaneously obtaining good state estimates, when compared to Ensemble Kalman filters and Kalman filters using reduced-order models.

## I. INTRODUCTION

State estimation for large-scale systems is often characterized by a trade-off between computational complexity of the estimation algorithm and the quality of the resulting state estimate. Indeed many schemes, included the ones in [1]–[3] strive to approximate the Kalman filter, given its optimality as the minimum variance estimator in a linear stochastic setting [4]. However, the computational complexity of the estimator is limited when the state estimates are used in real-time, e.g., for state-based controllers, such as, MPC [5]. As a result, the state estimate (and control action) must be computed within the sample time, which restricts the applicability of certain state estimation schemes for large-scale systems.

In this paper, we, therefore, propose a state estimator for a class of large-scale systems that strives to approximate the steady state Kalman filter, while simultaneously being computationally tractable. In particular, we focus on systems derived from Partial Differential Equations (PDEs). For these systems, the state-vector typically has an intuitive interpretation on the spatial domain. This strong link to the spatial domain will be extensively exploited to improve the computational aspects of our estimator. More specifically, we will exploit the (often strong) spatial correlation between the state elements on the spatial domain.

This research is supported by KWF Kankerbestrijding and NWO Domain AES, as part of their joint strategic research programme: Technology for Oncology II. The collaboration project is co-funded by the PPP Allowance made available by Health~Holland, Top Sector Life Sciences & Health, to stimulate public-private partnerships.

<sup>1</sup>Control Systems Technology, Department of Mechanical Engineering, Eindhoven University of Technology, Eindhoven, The Netherlands

<sup>2</sup>Electromagnetics for Care & Cure, Department of Electrical Engineering, Eindhoven University of Technology, Eindhoven, The Netherlands

<sup>3</sup>Department of Radiotherapy, Erasmus University Medical Center Cancer Institute, Rotterdam, The Netherlands

We will show the strengths of our method using a hyperthermia cancer treatment case study [6]–[9], which inspired this work, as we observed the computational complexity of standard Kalman filter-based estimators in this setting. Employing model-based estimators to improve Magnetic Resonance Thermometry (MRT) has been investigated already in literature. For example, in [10], a Kalman filter was used for models of small tumors. In [8], [9] a Kalman filter using a reduced-order model (ROM) was proposed. However, ROM-based Kalman filters are not easy to design and can suffer from long range spurious correlations [1]. In [11], a simple Luenberger observer is used. The full state measurements in [11] enable an easy to design diagonal innovation gain. However, this simple design does not exploit the strong spatial correlations in the temperature field. In this paper, we will compare our method to several well-known Kalman filter approximations and the Luenberger observer. Our method strives to accurately estimate states for large models, e.g., with  $10^6$  states and a computation time in the order of 1 – 10 seconds by bridging the gap between computationally intensive Kalman filter and the simple Luenberger observer. We achieve this by interpreting the steady state Kalman filter as a kernel regularized optimization problem. Additionally, we will experimentally show the strengths of our method on an actual hyperthermia treatment setup.

The remainder of this paper is structured as follows. In Section II, we introduce the problem and several well-known state estimators. In Section III, we present our proposed state estimation scheme. In Section IV, we compare our method against the introduced state estimators using a hyperthermia cancer treatment case-study. Last, in Section V, we provide our conclusions and give an outlook for future work.

## II. PRELIMINARIES

In this section, we start by introducing the system. Hereafter, we summarize the Ensemble Kalman filter, reduced-order model Kalman filter, and Luenberger observer, as these form benchmarks used to compare our newly proposed estimation scheme to.

### A. System class

In this paper, we consider the discrete-time linear time-invariant (LTI) system

$$\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k + \mathbf{B}\mathbf{u}_k + \mathbf{w}_k, \quad \mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}), \quad (1a)$$

$$\mathbf{y}_k = \mathbf{C}\mathbf{x}_k + \boldsymbol{\eta}_k, \quad \boldsymbol{\eta}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}), \quad (1b)$$

where  $\mathbf{x}_k \in \mathbb{R}^{n_x}$ ,  $\mathbf{y}_k \in \mathbb{R}^{n_y}$ ,  $\mathbf{u}_k \in \mathbb{R}^{n_u}$  denote the state, output, and input, respectively, at time  $k \in \mathbb{N}$ . Additionally,  $\mathbf{w}_k \in \mathbb{R}^{n_x}$  and  $\boldsymbol{\eta}_k \in \mathbb{R}^{n_y}$  denote the zero mean Gaussian distributed process and measurement noise with covariance

$\mathbf{Q}$  and  $\mathbf{R}$ , respectively. Moreover,  $\mathbf{A}$ ,  $\mathbf{B}$ , and  $\mathbf{C}$  are matrices of appropriate dimensions. Recall that we focus on the case where  $\mathbf{x}_k \in \mathbb{R}^{n_x}$  with  $n_x \gg 1$  and where the elements of  $\mathbf{x}_k$  are the result of a spatial discretization, i.e., the  $i$ -th element is given by  $\mathbf{x}_k^{(i)} = x_k(\mathbf{r}_i)$  with  $x_k : \mathbb{D} \rightarrow \mathbb{R}$  and  $\mathbf{r}_i \in \mathbb{D}$  for  $i = 1, \dots, n_x$ , where  $\mathbb{D} \subset \mathbb{R}^D$  denotes the spatial domain.

Below, we review existing observer schemes. In particular for large-scale systems, the observers in Sections II-B, II-C, and III require two assumptions.

- $\mathbf{R}^{-1}\mathbf{v}$  and  $\mathbf{R}^{-\top}\mathbf{v}$  are computationally cheap.
- $\mathbf{C}\mathbf{v}$  and  $\mathbf{C}^\top\mathbf{v}$  are computationally cheap.

A typical  $\mathbf{R}$ , that satisfies the first assumption results from uncorrelated measurement noise, i.e.,  $\mathbf{R}$  is diagonal.

### B. Ensemble Kalman filter

The Ensemble Kalman filter (EnKF) [1] strives to represent the state estimate covariance by an ensemble of state estimates, i.e.,  $\text{Var}(\hat{\mathbf{x}}_k) \approx \bar{\mathbf{X}}_k \bar{\mathbf{X}}_k^\top$  with  $\bar{\mathbf{X}}_k := \frac{1}{\sqrt{N-1}} \left( \begin{bmatrix} \mathbf{x}_k^{(1)} & \dots & \mathbf{x}_k^{(N)} \end{bmatrix} - \frac{1}{N} \sum_{j=1}^N \mathbf{x}_k^{(j)} \right)$ . Here,  $\hat{\mathbf{x}}_k$  denotes the state estimate and  $\mathbf{x}_k^{(j)} \in \mathbb{R}^{n_x}$  refers to the  $j$ -th ensemble member and its evolution is given by

$$\mathbf{x}_k^{(j)} = \mathbf{A}\hat{\mathbf{x}}_{k-1}^{(j)} + \mathbf{B}\mathbf{u}_{k-1} + \mathbf{w}_k^{(j)}, \quad (2)$$

where  $\mathbf{w}_k^{(j)}$  is a particular process noise realization. When the ensemble is sufficiently large, the statistical properties of  $\bar{\mathbf{X}}_k \bar{\mathbf{X}}_k^\top$  approximate the true covariance well.

The EnKF updates its ensemble using

$$\hat{\mathbf{x}}_k^{(j)} = \mathbf{x}_k^{(j)} + \bar{\mathbf{X}}_k \bar{\mathbf{Y}}_k^\top \mathbf{R}^{-1} (\mathbf{I} - \bar{\mathbf{Y}}_k (\mathbf{I} + \bar{\mathbf{Y}}_k^\top \mathbf{R}^{-1} \bar{\mathbf{Y}}_k)^{-1} \bar{\mathbf{Y}}_k^\top \mathbf{R}^{-1}) (\mathbf{y}_k - \bar{\mathbf{Y}}_k), \quad (3)$$

where  $\bar{\mathbf{Y}}_k = \mathbf{C} \bar{\mathbf{X}}_k$ . Observe that (3) uses the assumptions in Section II-A in order to reduce the computational complexity of the matrix inverse using the Woodbury identity [12]. To obtain a single state estimate from the ensemble, one can simply compute the mean, i.e.,  $\hat{\mathbf{x}}_k = \frac{1}{N} \sum_{j=1}^N \hat{\mathbf{x}}_k^{(j)}$ . Note that there are many variations on the EnKF tailored to specific systems, see [1] and the references therein.

We summarized the computational complexity of each step for the EnKF in Table I. Observe that the computational complexity is dominated by  $\mathcal{O}(Nn_x^2)$  for large-scale systems (typically  $N \ll n_x$ ). Indeed, propagating the ensemble using (2) proved to be computationally intensive, see Section IV.

### C. Reduced-order model Kalman filter

Large-scale systems are often well-described by smaller ROMs [13], [14]. To this end, it is possible to rewrite the full-order Kalman filter using the ROM given by

$$\mathbf{z}_{k+1} = \mathbf{A}_r \mathbf{z}_k + \mathbf{B}_r \mathbf{u}_k + \mathbf{V} \mathbf{w}_k, \quad (4a)$$

$$\mathbf{y}_k = \mathbf{C}_r \mathbf{z}_k + \boldsymbol{\eta}_k, \quad (4b)$$

where  $\mathbf{z}_k := \mathbf{V} \mathbf{x}_k \in \mathbb{R}^{n_r}$  is the reduced-order state vector with  $n_r \ll n_x$  and  $\mathbf{V}$  is the projection matrix that satisfies  $\mathbf{V}^\top \mathbf{V} = \mathbf{I}$ . The Kalman filter using the ROM is given by

$$\bar{\mathbf{z}}_k = \mathbf{A}_r \bar{\mathbf{z}}_{k-1} + \mathbf{B}_r \mathbf{u}_{k-1}, \quad (5a)$$

$$\bar{\mathbf{P}}_k = \mathbf{A}_r \bar{\mathbf{P}}_{k-1} \mathbf{A}_r^\top + \mathbf{V}^\top \mathbf{Q} \mathbf{V}, \quad (5b)$$

$$\mathbf{K}_k = (\mathbf{C}_r \bar{\mathbf{P}}_k \mathbf{C}_r^\top + \bar{\mathbf{P}}_k^{-1})^{-1} \mathbf{C}_r^\top \bar{\mathbf{P}}_k^{-1}, \quad (5c)$$

$$\hat{\mathbf{z}}_k = \bar{\mathbf{z}}_k + \mathbf{K}_k (\bar{\mathbf{y}}_k - \mathbf{y}_k), \quad (5d)$$

$$\bar{\mathbf{P}}_k = (\mathbf{I} - \mathbf{K}_k \mathbf{C}_r) \bar{\mathbf{P}}_k. \quad (5e)$$

Here,  $\bar{\mathbf{z}}_k$  and  $\hat{\mathbf{z}}_k$  denote the predicted and estimated state, respectively. Note that we reduce the computational complexity of (5c) using the Woodbury identity [12]. Crucially, by projecting  $\mathbf{Q}$  to the reduced-order state space, it is likely that part of the process noise is not well described by the reduced-order Kalman filter (ROM-KF).

The computational complexity for the ROM-KF is provided in Table I. Note that for  $n_r \ll n_x$ , the computational complexity is small, even when the original system is large.

### D. Luenberger observer

Last, we introduce the Luenberger observer

$$\bar{\mathbf{x}}_k = \mathbf{A} \hat{\mathbf{x}}_{k-1} + \mathbf{B} \mathbf{u}_{k-1}, \quad (6a)$$

$$\hat{\mathbf{x}}_k = \bar{\mathbf{x}}_k + \mathbf{K}_{\text{Luenberger}} (\mathbf{C} \bar{\mathbf{x}}_k - \mathbf{y}_k), \quad (6b)$$

where  $\mathbf{K}_{\text{Luenberger}}$  is the user-designed innovation gain.

For large-scale systems with full state measurements, i.e.,  $\mathbf{C} = \mathbf{I}$ , as seen in some hyperthermia cancer treatments [11], [15], an easy-to-design Luenberger observer uses a diagonal innovation gain, i.e.,  $\mathbf{K}_{\text{Luenberger}}$  is diagonal. We summarized the computational complexity of (6) when  $\mathbf{C} = \mathbf{I}$  and  $\mathbf{K}_{\text{Luenberger}}$  is diagonal in Table I. Observe that the computational complexity is dominated by propagating the state forward in time when  $\mathbf{K}_{\text{Luenberger}}$  is sufficiently sparse.

TABLE I: Computational complexity of the observers in Sections II-B to II-D

Equation	(2)	(3)	(5a)	(5b)	(5c)
Complexity	$\mathcal{O}(Nn_x^2)$	$\mathcal{O}(N^3)$	$\mathcal{O}(n_r^2)$	$\mathcal{O}(n_r^3)$	$\mathcal{O}(n_r^3)$
Equation	(5d)	(5e)	(6a)	(6b)	
Complexity	$\mathcal{O}(n_r n_y)$	$\mathcal{O}(n_r^2 n_y)$	$\mathcal{O}(n_x^2)$	$\mathcal{O}(n_x)$	

## III. APPROXIMATE KALMAN FILTERING FOR LARGE-SCALE SYSTEMS

In this section, we present our proposed approximate Kalman filter that is computationally efficient, intuitive to design, and can provide accurate state estimates, which we will refer to as the least-squares kernel Kalman filter (LSK-KF). To this end, we first introduce an equivalent least-squares form of the Kalman filter [16]. Hereafter, we present our key insights that enable the efficient computation of the least-squares minimization problem for large-scale systems.

### A. Least-squares Kalman filter

The least-squares minimization problem that is equivalent to the Kalman filter is

$$\bar{\mathbf{x}}_k = \mathbf{A} \hat{\mathbf{x}}_{k-1} + \mathbf{B} \mathbf{u}_{k-1}, \quad (7a)$$

$$\hat{\mathbf{x}}_k = \bar{\mathbf{x}}_k + \mathbf{d}_k, \quad (7b)$$

$$\mathbf{d}_k := \arg \min_{\mathbf{d}} \|\mathbf{d}\|_{\mathbf{P}_k}^2 + \|\mathbf{y}_k - \mathbf{C} \bar{\mathbf{x}}_k - \mathbf{C} \mathbf{d}\|_{\mathbf{R}^{-1}}^2 \quad (7c)$$

with  $\|\mathbf{v}\|_{\mathbf{M}}^2 := \mathbf{v}^\top \mathbf{M} \mathbf{v}$ . Note that (7) describes the update for the state estimate  $\hat{\mathbf{x}}_k$  but not for the state estimate covariance  $\mathbf{P}_k$ . As a result, (7) is often used with a static  $\mathbf{P}_k$ , i.e.,  $\mathbf{P}_k = \mathbf{P}$  for all  $k \in \mathbb{N}$ , as is the case in the steady state Kalman filter, see, e.g., [3]. Despite a static state estimate covariance, storing and computing products with  $\mathbf{P}$  is not tractable when  $n_x \gg 1$ . Moreover,  $\mathbf{P}$  is often ill-conditioned. Hence, solving (7) can be extremely challenging, especially when the estimator has to be used in real-time applications.

## B. Key concept

To solve the issue that (7) is computationally challenging, we can decompose  $\mathbf{P}$  as  $\mathbf{P} = \mathbf{P}^{\frac{1}{2}} \mathbf{P}^{\frac{1}{2}}$  [2]. This decomposition facilitates the coordinate change  $\mathbf{d} = \mathbf{P}^{\frac{1}{2}} \mathbf{f}$ , which can be used to pre-condition (7). Note this change of coordinates simultaneously avoids the inversion of  $\mathbf{P}$  and improves the conditioning of the optimization problem, as we will see.

In our work, we generalize this idea to the decomposition  $\mathbf{P} = \mathbf{L}\mathbf{L}^\top$ , where  $\mathbf{L}$  can be any matrix for which matrix-vector products  $\mathbf{L}\mathbf{v}$  and  $\mathbf{L}^\top\mathbf{v}$  can be computed efficiently. Intuitively, directly designing  $\mathbf{P}$  is analogous to kernel regularization methods, where the kernel specifies desired properties of  $\mathbf{d}$  in (7) [17]. Hence, we will refer to this method as the *least-squares kernel Kalman filter* (LSK-KF), as it utilizes a kernel to efficiently approximate the least-squares form of the Kalman filter. Substituting the  $\mathbf{d}_k = \mathbf{L}\mathbf{f}_k$  and  $\mathbf{P} = \mathbf{L}\mathbf{L}^\top$  (with  $\mathbf{P}_k = \mathbf{P}$ ) into (7b) and (7c) yields

$$\hat{\mathbf{x}}_k = \bar{\mathbf{x}}_k + \mathbf{L}\mathbf{f}_k, \quad (8a)$$

$$\mathbf{f}_k := \arg \min_{\mathbf{f}} \|\mathbf{f}\|_2^2 + \|\mathbf{y}_k - \mathbf{C}\bar{\mathbf{x}}_k - \mathbf{C}\mathbf{L}\mathbf{f}\|_{\mathbf{R}^{-1}}^2. \quad (8b)$$

As we will see, interpreting (8b) (and also (7)) as a kernel-regularized least-squares problem will provide insight into the design of  $\mathbf{L}\mathbf{L}^\top$ . We compute (8b) by solving for the first-order optimality conditions, i.e., the gradient of (8b) with respect to  $\mathbf{f}$  being zero. This yields the equality

$$\mathbf{f} + \mathbf{L}^\top \mathbf{C}^\top \mathbf{R}^{-1} \mathbf{C} \mathbf{L} \mathbf{f} = \mathbf{L}^\top \mathbf{C}^\top \mathbf{R}^{-1} (\mathbf{y}_k - \mathbf{C}\bar{\mathbf{x}}_k), \quad (9)$$

which can be efficiently solved with iterative linear solvers, such as, for example, the Conjugate Gradient (CG) method [18]. Solving (9) efficiently with iterative linear solvers requires that the matrix-vector products  $\mathbf{R}^{-1}\mathbf{v}$ ,  $\mathbf{C}\mathbf{v}$ ,  $\mathbf{C}^\top\mathbf{v}$ ,  $\mathbf{L}\mathbf{v}$ , and  $\mathbf{L}^\top\mathbf{v}$  are computationally cheap (for a  $\mathbf{v}$  of appropriate dimension). This observation also follows from the computational complexity of CG, which is  $\mathcal{O}(c\sqrt{\kappa})$ , where  $c$  denotes the complexity of the matrix-vector products with  $\mathbf{f}$  in (9) and  $\kappa$  the condition number of (8b) [18].

## C. Choosing $\mathbf{L}$

The insight that (7) can be efficiently solved by decomposing a static  $\mathbf{P}$  as  $\mathbf{L}\mathbf{L}^\top$ , does not solve the problem of choosing an appropriate  $\mathbf{L}$ . In order for LSK-KF to approximate the steady state Kalman-filter, we identified two crucial properties

- $\mathbf{L}\mathbf{L}^\top$  approximates the asymptotic state estimate covariance of the steady state Kalman filter.
- Matrix-vector products  $\mathbf{L}\mathbf{v}$  and  $\mathbf{L}^\top\mathbf{v}$  must be computationally cheap.

In order to satisfy these two assumptions, we propose to design  $\mathbf{L}$  as compositions and linear combinations of so-called ‘‘building blocks’’. A crucial property of a ‘‘building block’’ is that matrix-vector products are efficient and have an interpretation on the spatial domain, e.g., spatial smoothing. The interpretation on the spatial domain will be used at the end of this section as a design tool to visualize the covariance  $\mathbf{L}\mathbf{L}^\top$ , and thus the properties of the resulting state estimator.

In Table II, we provide a selection of building blocks that can be used to construct  $\mathbf{L}$  using linear combinations or compositions. We will start by introducing one building block in detail. Suppose  $\mathbf{v} \in \mathbb{R}^n$  is the result of spatially discretizing

$v : \mathbb{D} \rightarrow \mathbb{R}$  on  $\{\mathbf{r}_1, \dots, \mathbf{r}_n\} \subset \mathbb{D}$ , i.e.,  $\mathbf{v}^{(i)} := v(\mathbf{r}_i)$  for  $i = 1, \dots, n$ . Then, we can approximate the convolution  $(l * v)(\mathbf{r}) := \int_{\mathbb{D}} l(\mathbf{r} - \mathbf{r}')v(\mathbf{r}')d\mathbf{r}'$ , with  $l : \mathbb{D} \rightarrow \mathbb{R}$ , by the matrix-vector product,  $\mathbf{L}\mathbf{v}$  where  $\mathbf{L} \in \mathbb{R}^{n \times n}$  and  $\mathbf{L}^{(i,j)} = l(\mathbf{r}_i - \mathbf{r}_j)$ . The key insight is to evaluate  $\mathbf{L}\mathbf{v}$  on the spatial domain using (with some abuse of notation)  $(l * v)(\mathbf{r}) = \mathcal{F}^{-1}(\mathcal{F}(l)\mathcal{F}(v))(\mathbf{r})$ , where  $\mathcal{F}$  denotes the Fourier transform. In fact, the computational complexity of the matrix-vector product  $\mathbf{L}\mathbf{v}$  is generally  $\mathcal{O}(n^2)$ , but when  $\mathbf{L}\mathbf{v}$  can be computed using the Fast Fourier Transform (FFT), it is reduced to  $\mathcal{O}(n \log(n))$  [19]. Clearly, some matrix-vector products  $\mathbf{L}\mathbf{v}$  allow for efficient implementation without explicitly constructing  $\mathbf{L}$ . We will refer to such matrix-vector products, as matrix-free matrix-vector products.

The remaining building blocks are the following. Multiplying two functions on  $\mathbb{D}$ , which is equivalent to a matrix-vector product with a diagonal matrix, and can be thus computed with element-wise multiplications. The convolution with a separable kernel  $\int_{\mathbb{D}} l(\mathbf{r}, \mathbf{r}')v(\mathbf{r}')d\mathbf{r}'$ , with  $l(\mathbf{r}_i, \mathbf{r}_j) = \prod_{d=1}^D l_d(\mathbf{r}_i^{(d)}, \mathbf{r}_j^{(d)})$ , which is equivalent to the matrix-vector product where  $\mathbf{L}$  is the result of a Kronecker product, i.e.,  $\mathbf{L} = \bigotimes_{d=1}^D \mathbf{L}_d$  with  $\mathbf{L}_d \in \mathbb{R}^{n_d}$  for  $d = 1, \dots, D$ . This matrix-vector product can be efficiently computed using fast tensor products [20]. Observe that all building blocks have an intuitive interpretation on the spatial domain, which helps in designing an appropriate  $\mathbf{L}$ , as we will see.

TABLE II: A selection of matrix-vector products  $\mathbf{L}\mathbf{v}$  and  $\mathbf{F}\mathbf{v} := \mathbf{L}^\top\mathbf{v}$  with efficient matrix-free implementations.

Matrix $\mathbf{L}$	Interpretation of $\mathbf{L}\mathbf{v}$ on $\mathbb{D}$	Complexity
$\mathbf{L}^{(i,i)} = l(\mathbf{r}_i)$	$l(\mathbf{r})v(\mathbf{r})$	$\mathcal{O}(n)$
$\mathbf{L}^{(i,j)} = l(\mathbf{r}_i - \mathbf{r}_j)$	$l(\mathbf{r}) * v(\mathbf{r})$	$\mathcal{O}(n \log(n))$
$\mathbf{L} = \bigotimes_{d=1}^D \mathbf{L}_d$	$\prod_{d=1}^D l_d(\mathbf{r}_i^{(d)}, \mathbf{r}_j^{(d)}) * v(\mathbf{r})$	$\mathcal{O}(n \sum_{d=1}^D n_d)$
Matrix $\mathbf{F}$	Interpretation of $\mathbf{F}\mathbf{v}$ on $\mathbb{D}$	Complexity
$\mathbf{F}^{(i,i)} = l(\mathbf{r}_i)$	$l(\mathbf{r})v(\mathbf{r})$	$\mathcal{O}(n)$
$\mathbf{F}^{(i,j)} = l(\mathbf{r}_j - \mathbf{r}_i)$	$l(-\mathbf{r}) * v(\mathbf{r})$	$\mathcal{O}(n \log(n))$
$\mathbf{F} = \bigotimes_{d=1}^D \mathbf{L}_d^\top$	$\prod_{d=1}^D l_d(\mathbf{r}_j^{(d)}, \mathbf{r}_i^{(d)}) * v(\mathbf{r})$	$\mathcal{O}(n \sum_{d=1}^D n_d)$

Besides computational efficient matrix-vector products, we need to design an appropriate  $\mathbf{L}$  such that  $\mathbf{L}\mathbf{L}^\top$  approximates the state estimate covariance of the steady state Kalman filter. Clearly, due to the size of  $\mathbf{L}\mathbf{L}^\top$ , inspecting the matrix itself is prohibitively difficult. Therefore, we propose using the conditional expected value to evaluate  $\mathbf{L}\mathbf{L}^\top$ . More specifically, when

$$\begin{bmatrix} \mathbf{v}_a \\ \mathbf{v}_b \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} \mathbf{P}_a & \mathbf{P}_c \\ \mathbf{P}_c^\top & \mathbf{P}_b \end{bmatrix}\right), \quad (10)$$

the conditional expected value is given by

$$\mathbb{E}(\mathbf{v}_a | \mathbf{v}_b) = \mathbf{P}_c \mathbf{P}_b^{-1} \mathbf{v}_b \quad (11)$$

and expresses the expected value of  $\mathbf{v}_a$  conditioned on  $\mathbf{v}_b$ . First of all, visualizing  $\mathbb{E}(\mathbf{v}_a | \mathbf{v}_b)$  on the spatial domain will show the correlation of  $\mathbf{v}_a$  with  $\mathbf{v}_b$ , which is useful as  $\mathbf{L}\mathbf{L}^\top$  can be interpreted as a regularization kernel. Hence, visualizing  $\mathbb{E}(\mathbf{v}_a | \mathbf{v}_b)$  will show state distributions that are associated with a low cost, i.e., state distributions that are promoted by (8b). As a result, the conditional expected value is a useful design tool in designing an appropriate  $\mathbf{L}$  by, for example, mimicking the steady state Kalman filter from

a small-scale model. We will show how to utilize  $\mathbb{E}(v_a | v_b)$  in designing  $\mathbf{L}$  using our hyperthermia case study in Section IV. Last, observe that (11) is easy to compute with matrix-vector products when  $v_b \in \mathbb{R}$  by using  $[\mathbf{P}_c^\top \ \mathbf{P}_b]^\top = \mathbf{L}\mathbf{L}^\top [0 \ 0 \ \dots \ 1]^\top$ .

#### IV. EXPERIMENTAL RESULTS

In this section, we demonstrate the proposed LSK-KF observer by comparing the state estimates with the observers introduced in Section II using a hyperthermia cancer treatment case-study. Hereto, we first introduce the experimental setup used in actual hyperthermia treatments. Hereafter, we define the state observers, and last, the performance of each observer is quantified by means of the computation time and the state estimation error.

##### A. Experiment setup

We experimentally validate the observers in Sections II and III using an anthropomorphic phantom of the pelvic region to mimic a typical deep hyperthermia treatment, see [9]. A voxel-based model of the phantom is shown in Figure 1. The phantom and hyperthermia applicator can be modeled using a Partial Differential Equation (PDE)

$$\rho(\mathbf{r})c(\mathbf{r})\dot{T}(\mathbf{r}, t) = \nabla \cdot (k(\mathbf{r})\nabla T(\mathbf{r}, t)) + q(\mathbf{r}, t), \quad (12a)$$

for  $\mathbf{r} \in \Omega \subset \mathbb{R}^3$ , with  $q(\mathbf{r}, t) = b_1(\mathbf{r})u_1(t) + b_2(\mathbf{r})u_2(t)$  and convective boundary condition

$$k(\mathbf{r})\nabla T(\mathbf{r}, t) \cdot \mathbf{n} = hT(\mathbf{r}, t), \quad (12b)$$

for  $\mathbf{r} \in \partial\Omega$ , where  $\Omega = \Omega_{\text{soft-tissue}} \cup \Omega_{\text{shell}} \cup \Omega_{\text{spine}} \cup \Omega_{\text{pelvis}}$  denotes the phantom domain. Additionally,  $T$ ,  $k$ ,  $c$ ,  $\rho$ ,  $b_i$ ,  $u_i$  denote the temperature field with respect to room temperature, thermal conductivity, heat capacity, density,  $i$ -th heat load, and  $i$ -th input, respectively. We have two inputs  $u_i$  as we consider two RF-antenna inputs, i.e., right focus (corresponding to  $b_1(\mathbf{r})$ ) and left focus (corresponding to  $b_2(\mathbf{r})$ ), see Figure 2. The relevant material properties for (12) are given in [9]. Note that we consider temperatures with respect to a steady state, hence, the temperature of the boundary condition is not relevant for this demonstration.

After discretizing (12) in space on a well-chosen nodal grid  $\{\mathbf{r}_1, \dots, \mathbf{r}_{n_x}\} \subset \Omega$  ( $n_x \approx 10^6$ ) and in time with a sample time of 93 seconds (the time to acquire an MRI scan), we obtain a discrete-time system in the form of (1) with state vector  $\mathbf{x}_k := [T(\mathbf{r}_1, t_k) \ \dots \ T(\mathbf{r}_{n_x}, t_k)]^\top$  and  $t_k = 93k$ ,  $k \in \mathbb{N}$ . The output of (1) results from temperature measurements using Magnetic Resonance Imaging (MRI) by exploiting the temperature-dependent Proton Resonance Frequency Shift (PRFS), see [9] and the references therein. As a result, we define the  $i$ -th element of the output as  $\mathbf{y}_k^{(i)} := T(\mathbf{r}_j, t_k)$  for all  $\mathbf{r}_j \in \Omega_{\text{soft-tissue}}$ . Note that only the soft tissue has a temperature-dependent PRFS. The variance associated with the  $i$ -th output, i.e.,  $\mathbf{R}^{(i,i)}$ , is estimated from the measurement data. We refer the interested reader to [8], [9] for a more in-depth description of the phantom, corresponding thermal model, and the temperature measurement using MRI.

To validate the estimated temperature (using any observer), six temperature probes are inserted in the center transversal slice of the phantom using catheter tracks, see Figure 3 for

the probe locations. These probes serve as a spatially sparse ground truth temperature with a resolution of  $0.1^\circ\text{C}$  and accuracy  $\leq 0.1^\circ\text{C}$ . During an experiment of approximately 30 minutes, two sets of inputs are applied to the phantom. The experiment sequence is given by (13). Both inputs apply approximately 200 W to the phantom.

$$u_1(t_k) = 1 \text{ if } (2 \leq k \leq 7), \text{ and } 0, \text{ otherwise,} \quad (13a)$$

$$u_2(t_k) = 1 \text{ if } (8 \leq k \leq 15), \text{ and } 0, \text{ otherwise.} \quad (13b)$$

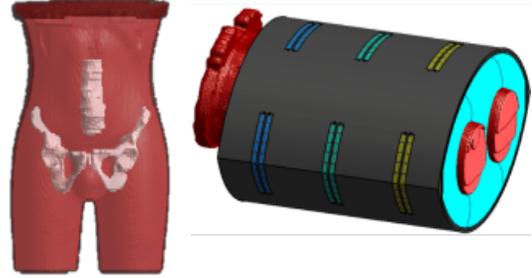


Fig. 1: Transparent front view of the phantom model (left) and a phantom model placed in the applicator (right).

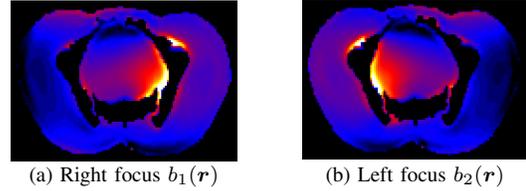


Fig. 2: Model-based heat loads  $b_1(\mathbf{r})$  and  $b_2(\mathbf{r})$  on the central transversal slice.

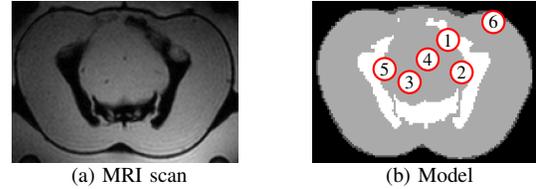


Fig. 3: Temperature probe locations on the central transversal slice of the phantom and model. The white and gray regions in Figure 3b denote the pelvis and soft-tissue, respectively.

##### B. LSK-KF design

In this section, we present the design of the LSK-KF observer tailored to the anthropomorphic phantom model (12). For this particular model, we design  $\mathbf{L}$  such that we obtain strong spatial correlations between nearby states within the *same* material and weak correlation between *different* materials. This insight was obtained by solving the discrete-time algebraic Riccati equations to obtain the steady state Kalman filter for a *small-scale* thermal model analogous to (12). Hereto, we choose  $\mathbf{L}$  such that the matrix-vector product  $\mathbf{d} = \mathbf{L}\mathbf{v}$  with  $\mathbf{d}^{(i)} = d(\mathbf{r}_i)$  for  $i = 1, \dots, n_x$  is given by

$$d(\mathbf{r}) = \sum_{i=1}^2 \varphi_i(\mathbf{r})(k(\mathbf{r}) * (\varphi_i(\mathbf{r})v)), \quad (14a)$$

$$\varphi_i(\mathbf{r}) = \begin{cases} 1, & \text{if } \mathbf{r} \in \Omega_i \\ 0, & \text{otherwise} \end{cases}, \quad k(\mathbf{r}) = \gamma e^{-\sigma^{-2}\|\mathbf{r}\|_2^2}. \quad (14b)$$

Here,  $\Omega_1 = \Omega_{\text{soft-tissue}}$  and  $\Omega_2 = \Omega_{\text{shell}} \cup \Omega_{\text{spine}} \cup \Omega_{\text{pelvis}}$ . Note that for this particular choice  $\mathbf{L}$  in (14), we have

$Lv = L^\top v$ . In (14), we designed the parameters  $\gamma$  and  $\sigma$  to match the small-scale steady state Kalman filter by using the conditional expected value (11). Observe that the computational complexity of the matrix-vector product (14) scales with  $\mathcal{O}(n_x \log(n_x))$ , when the convolution with  $k(r)$  is implemented using the FFT [19].

In Figure 4, we demonstrate the conditional expected value for two different state elements, one in the soft-tissue material and one in the plastic pelvis. From Figure 4, we conclude that our proposed design (14) indeed achieves strong, localized, spatial correlations within a *single* material and no correlation between *different* materials, as was observed from the steady state Kalman filter using the small-scale model.

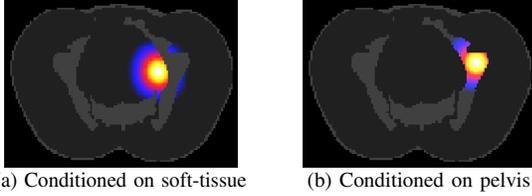


Fig. 4: Conditional expected value conditioned on a state in the soft-tissue (a) and a state in the pelvis (b).

### C. Observer definitions

In this section, we briefly present the designs for the observers from Section II. First, for the EnKF, we compute the process noise realization as  $w = Lv$  with  $v \sim \mathcal{N}(0, I)$  and  $L$  as in (14). To evaluate the EnKF we use two ensemble sizes of  $N = 20$  and  $N = 100$ , respectively. As we will see, the computational tractability of the EnKF will be limited by the ensemble size.

Second, the model (12) and corresponding  $V$  used in the ROM-KF (5a) is computed using data-driven balanced truncation, and truncating the singular values to include 99.9% of the energy [14]. The resulting model is of dimension  $n_r = 10$ . This significant reduction in model order is largely the result of the fact that the model has only two inputs. The reduced-order process noise is estimated by projecting  $LL^\top$  onto the reduced-order subspace, i.e.,  $Q_r := (LV)^\top (LV)$ , where  $LV$  can be efficiently computed as  $V$  has few columns.

Last, the diagonal gain for the Luenberger observer is designed as  $K_{\text{Luenberger}} := C^\top D$  with  $D^{(i,i)} = \frac{d}{d + R^{(i,i)}}$ . Here,  $d \in \mathbb{R}$  is chosen as  $d = \frac{1}{n_x} \text{trace}(WW^\top)$  and  $W := [w^{(1)} \dots w^{(500)}]$ , where  $w^{(i)}$  is the same as for the EnKF. As a result,  $d$  models the approximate uncertainty of the state estimate for all states.

### D. Results

Figure 5 presents the estimated temperature for each observer at the probe locations together with the probe measurements. For a visual comparison between the estimators and the measured temperature, a transversal slice is shown in Figure 6. Here, it is clear that both EnKF estimators have significant temperature irregularities, which are not physically motivated. In the Luenberger observer, the spatially high frequent noise is slightly attenuated, which shows that the spatial correlation between states is not exploited. Both the ROM-KF and the LSK-KF show spatially smooth temperature fields, indicating that the spatial correlation is exploited.

The results from Figure 5 are summarized in Table III using the Root Means Square (RMS) error, the computation time, and the approximate standard deviation of the estimator error for the last two time-steps. The approximate standard deviation of the estimator error assumes the state estimate is given by  $\hat{x}_k = x_k + e_k$  where  $x_k$  is the true state and  $e_k$  is the estimator error, which is assumed IID. Given this assumption, when  $x_k \approx x_{k+1}$ , we can approximate the standard deviation of the elements in  $e_k$  using

$$\text{STD}(e_k) \approx \frac{\sqrt{2}}{2} \left( \frac{1}{n_x - 1} \sum_{i=1}^{n_x} (\hat{x}_k^{(i)} - \hat{x}_{k+1}^{(i)} - \mu)^2 \right)^{1/2} \quad (15)$$

with  $\mu = \frac{1}{n_x} \sum_{i=1}^{n_x} \hat{x}_k^{(i)} - \hat{x}_{k+1}^{(i)}$ . Note that  $\frac{\sqrt{2}}{2}$  corrects the standard deviation as both state estimates  $x_k$  and  $x_{k+1}$  have an estimation error.

From Table III, it is clear that our LSK-KF observer agrees well with the probe data that serve as a ground truth while simultaneously being computationally attractive, even for a model with  $10^6$  states. Based on the approximate standard deviation in Table III, we observe that two consecutive state estimates show little variation when the temperature field is approximately constant between the two time-steps.

With regard to the other state estimation schemes, the EnKF proved to be computationally intractable, even for modest ensemble sizes. As a result of the modest ensemble size, the state estimates do not agree well with the probes. The ROM-KF, agrees reasonably well with the probes and is computationally cheap. However, we observed that the mode shapes in the projection matrix  $V$  are not “rich” enough to describe the required temperature field as there are significant low frequency spatial correlations in the difference between the measurement and the estimated state. Including more reduced-order states to ROM-KF might solve this problem, but finding relevant modes shapes to extend the ROM-KF is not trivial. The Luenberger observer showed good agreement with the probes, is computationally attractive, but has a higher standard deviation, which is also seen by the spatially high frequent noise in Figure 6.

TABLE III: RMS error of the state estimate relative to the probe measurement, mean computation time, and standard deviation for each observer.

	LSK-KF	EnKF $N = 20$	EnKF $N = 100$	ROM-KF	Luenberger
RMS probe 1	0.33	0.74	0.49	0.54	0.28
RMS probe 2	0.22	0.41	0.33	0.19	0.28
RMS probe 3	0.39	0.60	0.26	0.40	0.31
RMS probe 4	0.21	1.04	0.45	0.56	0.25
RMS probe 5	0.31	1.56	0.93	0.47	0.38
RMS probe 6	0.81	1.29	1.13	1.34	0.70
RMS Total	0.43	1.02	0.65	0.69	0.41
Time [s]	7.21	40.7	200.8	0.05	2.13
STD( $e_{17}$ )	0.05	0.13	0.17	0.03	0.12

## V. CONCLUSION AND OUTLOOK

In this paper, we presented a computationally attractive state estimation framework for large-scale systems, that is particularly useful for systems derived from partial differential equations. In fact, the proposed LSK-KF can approximate the steady state Kalman filter by approximating the state

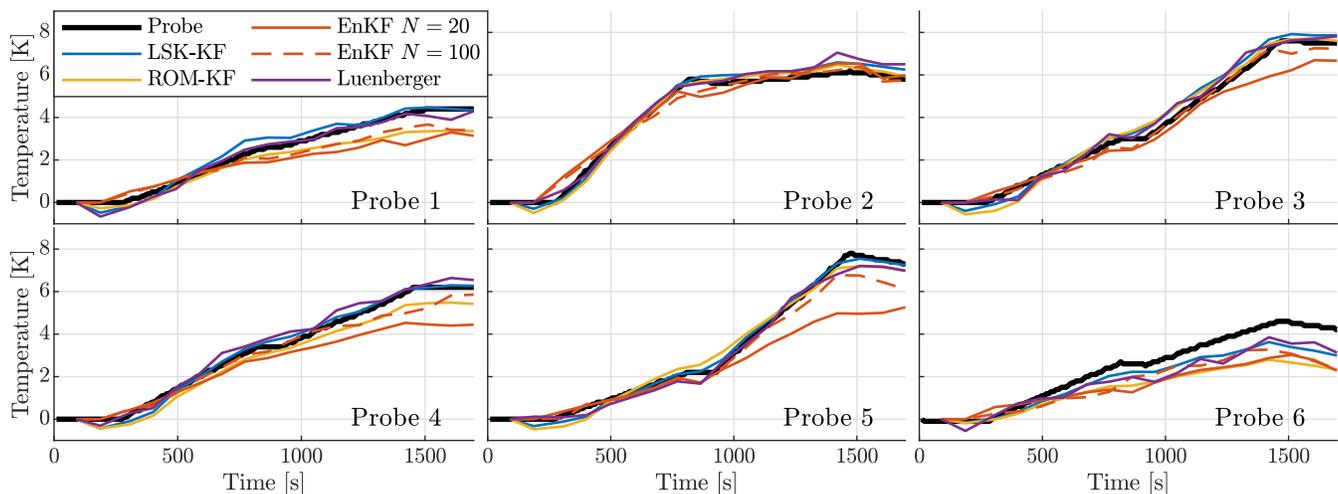


Fig. 5: Probe temperature and estimated temperature at the probe locations for different observers.

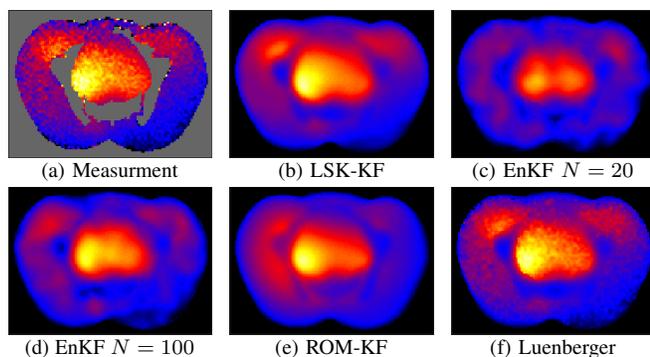


Fig. 6: Temperature measurement and estimates at  $t = 1674$  seconds on the transversal slice centered in the phantom. In Figure 6a, gray areas denote regions with no measurement.

estimate covariance using computationally attractive building blocks. These building blocks exploit the observation that some matrix-vector products can be efficiently evaluated on the spatial domain. We also proposed the use of the conditional expected value as a computationally efficient tool to design the observer.

A hyperthermia cancer treatment case study demonstrated the strengths of our estimator. We observed a good correlation to thermal probes, while simultaneously having a low computational complexity. Moreover, we showed how insights obtained from a small-scale model can be used to design the state estimator using the conditional expected value.

This work can be extended in various directions. For instance, expanding the library of building blocks, i.e., matrices with computationally attractive matrix-vector products, is of interest. Additionally, further automating the observer design procedure including stability guarantees for large-scale systems can even further improve the applicability of this work.

## REFERENCES

- [1] G. Evensen, "The Ensemble Kalman Filter: theoretical formulation and practical implementation," *Ocean Dyn.*, pp. 343–367, 2003.
- [2] S. A. Haben, A. Lawless, and N. Nichols, "Conditioning of incremental variational data assimilation, with application to the Met Office system," *Tellus A Dyn. Meteorol. Oceanogr.*, pp. 782–792, 2011.
- [3] D. M. Barker, W. Huang, Y. R. Guo, A. J. Bourgeois, and Q. N. Xiao, "A three-dimensional variational data assimilation system for MM5: Implementation and initial results," *Mon. Weather Rev.*, pp. 897–914, 2004.
- [4] R. E. Kalman, "A New Approach to Linear Filtering and Prediction Problems," *J. Basic Eng.*, pp. 35–45, 1960.
- [5] D. Q. Mayne, "Model predictive control: Recent developments and future promise," *Automatica*, 2014.
- [6] J. van der Zee, "Heating the patient: A promising approach?" *Ann. Oncol.*, pp. 1173–1184, 2002.
- [7] M. Kroesen, *et al.*, "Confirmation of thermal dose as a predictor of local control in cervical carcinoma patients treated with state-of-the-art radiation therapy and hyperthermia," *Radiother. Oncol.*, pp. 150–158, 2019.
- [8] R. Hendrikx, *et al.*, "POD-Based Recursive Temperature Estimation for MR-Guided RF Hyperthermia Cancer Treatment: A Pilot Study," in *Conf. Decis. Control*, 2018, pp. 5201–5208.
- [9] I. VilasBoas-Ribeiro, *et al.*, "POD-Kalman filtering for improving non-invasive 3D temperature monitoring in MR-guided hyperthermia," *Med Phys.*, pp. 1–16, 2022.
- [10] S. Roujol, B. D. de Senneville, S. Hey, C. Moonen, and M. Ries, "Robust Adaptive Extended Kalman Filtering for Real Time MR-Thermometry Guided HIFU Interventions," *IEEE Trans. Med. Imaging*, pp. 533–542, 2012.
- [11] D. A. Deenen, *et al.*, "Offset-Free Model Predictive Temperature Control for Ultrasound-Based Hyperthermia Cancer Treatments," *IEEE Trans. Control Syst. Technol.*, pp. 2351–2365, 2021.
- [12] W. W. Hager, "Updating the Inverse of a Matrix," *SIAM Rev.*, pp. 221–239, 1989.
- [13] A. C. Antoulas, *Approximation of Large-Scale Dynamical Systems*. SIAM, 2005.
- [14] C. W. Rowley, "Model Reduction for Fluids, using Balanced Proper Orthogonal Decomposition," *Int. J. Bifurc. Chaos*, pp. 997–1013, 2005.
- [15] D. Hensley, R. Orendorff, E. Yu, C. Danielson, V. Salgaonkar, and C. Diederich, "Model predictive control for treating cancer with ultrasonic heating," in *Am. Control Conf.*, 2015, pp. 220–225.
- [16] R. Nash, A. Gelb, and J. Kasper, *Applied Optimal Estimation*, 4th ed. Cambridge (Mass.): MIT press, 1974.
- [17] G. Pilonetto, F. Dinuzzo, T. Chen, G. De Nicolao, and L. Ljung, "Kernel methods in system identification, machine learning and function estimation: A survey," *Automatica*, pp. 657–682, 2014.
- [18] Y. Saad, *Iterative Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics, 2003.
- [19] J. W. Cooley and J. W. Tukey, "An algorithm for the machine calculation of complex Fourier series," *Math. Comput.*, pp. 297–301, 1965.
- [20] P. Fernandes, B. Plateau, and W. J. Stewart, "Efficient descriptor-vector multiplications in stochastic automata networks," *J. ACM*, pp. 381–414, 1998.