

Optimized FIR-Kalman architecture for differentially private event stream filtering

M. Jadidi, J. Le Ny, S. Gao

G-2026-31

July 2026

La collection *Les Cahiers du GERAD* est constituée des travaux de recherche menés par nos membres. La plupart de ces documents de travail a été soumis à des revues avec comité de révision. Lorsqu'un document est accepté et publié, le pdf original est retiré si c'est nécessaire et un lien vers l'article publié est ajouté.

Citation suggérée : M. Jadidi, J. Le Ny, S. Gao (Juillet 2026). Optimized FIR-Kalman architecture for differentially private event stream filtering, Rapport technique, Les Cahiers du GERAD G-2026-31, GERAD, HEC Montréal, Canada.

Avant de citer ce rapport technique, veuillez visiter notre site Web (<https://www.gerad.ca/fr/papers/G-2026-31>) afin de mettre à jour vos données de référence, s'il a été publié dans une revue scientifique.

The series *Les Cahiers du GERAD* consists of working papers carried out by our members. Most of these pre-prints have been submitted to peer-reviewed journals. When accepted and published, if necessary, the original pdf is removed and a link to the published article is added.

Suggested citation: M. Jadidi, J. Le Ny, S. Gao (July 2026). Optimized FIR-Kalman architecture for differentially private event stream filtering, Technical report, Les Cahiers du GERAD G-2026-31, GERAD, HEC Montréal, Canada.

Before citing this technical report, please visit our website (<https://www.gerad.ca/en/papers/G-2026-31>) to update your reference data, if it has been published in a scientific journal.

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2026
– Bibliothèque et Archives Canada, 2026

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Legal deposit – Bibliothèque et Archives nationales du Québec, 2026
– Library and Archives Canada, 2026

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Télec. : 514 340-5665
info@gerad.ca
www.gerad.ca

Optimized FIR-Kalman architecture for differentially private event stream filtering

Mohadese Jadidi

Jérôme Le Ny

Shuang Gao

*GERAD & Polytechnique Montréal, Montréal (QC),
Canada, H3T 1J4*

`mohadese.jadidi@polymtl.ca`

`jerome.le-ny@polymtl.ca`

`shuang.gao@polymtl.ca`

July 2026

Les Cahiers du GERAD

G–2026–31

Copyright © 2026 Jadidi, Le Ny, Gao

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Abstract : We investigate a state estimation problem for a linear time-invariant Gaussian system, based on a measured scalar privacy-sensitive signal. This is formulated as a Kalman filtering problem under an “event-level” differential privacy constraint, for which existing approaches, directly perturbing the measured signal, tend to degrade the estimation accuracy significantly. Here, we propose a two-stage differentially private architecture combining a finite impulse response (FIR) pre-filter with the Kalman filter. We cast the joint design of the FIR and Kalman filters as a (bilinear) optimization problem to improve the trade-off between privacy and estimation accuracy. Simulations on an epidemiological example demonstrate that the proposed method preserves estimation accuracy while ensuring event-level privacy.

Keywords: Differential privacy; Kalman filtering; finite impulse response filtering; state estimation; bilinear optimization.

1 Introduction

Large-scale monitoring and control systems, such as smart grids or intelligent transportation systems, often require individuals to continuously share their personal information. Hence, their potential benefits must be balanced against privacy concerns. For instance, sharing location data to obtain accurate navigation information can expose users to third-party tracking [1]. Even aggregate data, such as the number of people at each location, may reveal individual mobility patterns [2].

Various approaches have been proposed for privacy-preserving data analysis, e.g., using cryptography, measurement transmission scheduling, data perturbation [3], etc., with different types of privacy guarantees. In particular, when publishing aggregate statistics about a population based on a sensitive dataset, Differential Privacy (DP) [4] provides a framework to adjust the level of perturbation and limit one's ability to use the published outputs to make inferences about the data of specific individuals. For dynamic data taking the form of temporal signals, DP has been used to provide *event-level privacy* [5], when processing an aggregated signal impacted by a given individual at a single or a few time steps (e.g., a presence detector), or *user-level privacy*, when processing multiple signals, each one coming from a different individual [6]. In systems and control, DP has found applications to estimation and filtering [7], consensus [8], distributed optimization [9], or optimal control [10].

Differentially private (dp) estimation, filtering and control schemes often rely on adding privacy-preserving noise directly on the sensitive input signals [10, 11] or on the published outputs [7, 12]. When noise is added at the input, post-processing that exploits any publicly available model of the data can mitigate the impact of the noise [7, 13]. Further performance improvements can be obtained with two-stage architectures [6, 14], where the first stage reduces the amount of required privacy-preserving noise by removing information that is potentially redundant or irrelevant to the query, while the second stage again attenuates the effect of the privacy-preserving noise. For user-level privacy, [15] implements the first stage by linearly combining the participants' signals, and the second stage with a Kalman filter. For event-level privacy, [16, 17] develop two-stage dp filtering architectures when the private input signals are wide-sense stationary with known second-order statistics.

In this paper, we investigate the problem of Kalman filtering under event-level privacy to process an aggregate scalar-valued private signal for which a linear time-invariant (LTI) state-space model is known. As illustrated in Section 4, the sensitive signal could be for instance the number of people recovering from a virus each day (data reported by a hospital), from which one would like to publish a dp estimate of the total number of exposed and infectious individuals in the population, leveraging an epidemiological state-space model. We propose a two-stage architecture combining a finite impulse response (FIR) pre-filter as the first stage and a Kalman filter exploiting the LTI model as the second stage. We then develop an optimization scheme, involving a bilinear matrix inequality (BMI) [18], to optimize the architecture's estimation performance at any given DP level. Simulation results demonstrate that adding and optimizing the FIR pre-filter can provide significant performance improvements.

The rest of the paper is organized as follows. Section 2 outlines the necessary background on DP and formulates the problem. Section 3 presents the methodology for the design of the dp estimation architecture, including the optimization of the FIR pre-filter. Section 4 illustrates through simulations the effectiveness of the proposed method, and Section 5 concludes the paper.

2 Preliminaries and problem statement

2.1 Notation

We fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where $\mathcal{F}$ is a σ -algebra on the sample space Ω , and $\mathbb{P}$ is a probability measure on $\mathcal{F}$. The p -norm of a vector $v \in \mathbb{R}^n$ for $p \in [1, \infty)$ is defined as $|v|_p = (\sum_{i=1}^n |v_i|^p)^{1/p}$. We denote $\mathcal{S}^n := (\mathbb{R}^n)^{\mathbb{N}}$ the space of n -dimensional discrete-time signals. For $p \in [1, \infty)$, the ℓ_p -norm of a

signal $x = \{x_k\}_{k \geq 0} \in \mathcal{S}^n$ is $\|x\|_p = (\sum_{k=0}^{\infty} \sum_{i=1}^n |x_{k,i}|^p)^{1/p}$. We denote the discrete-time scalar impulse signal at time step l by δ^l , i.e., $\delta_k^l = 1$ if $k = l$ and 0 otherwise. The $\mathcal{H}_2$ -norm of a discrete-time system with transfer function $G(z)$ is given by $\|G\|_{\mathcal{H}_2} = \sqrt{\frac{1}{2\pi} \int_{-\pi}^{\pi} \text{Tr}(G(e^{j\omega})^* G(e^{j\omega})) d\omega}$. A random vector v following a Gaussian distribution with mean μ and covariance matrix Σ is denoted by $v \sim \mathcal{N}(\mu, \Sigma)$. A discrete-time white Gaussian noise (WGN) signal $\{w_k\}_{k \geq 0}$ is a sequence of independent and identically distributed zero-mean Gaussian random vectors. It is called *standard* when the covariance of each vector is the identity matrix. $\mathbb{S}^n$ denotes the set of $n \times n$ symmetric matrices. For a $P \in \mathbb{S}^n$, the notation $P \succeq 0$ (respectively $P \succ 0$) means that P is positive semi-definite (respectively positive definite).

2.2 Background on differential privacy

An algorithm, also called *mechanism*, processing private inputs in some space Y is dp if it publishes randomized outputs whose distribution is not very sensitive to variations in the input data due to the presence of any single individual. These possible variations are captured by a binary symmetric relation defined on Y , called adjacency.

Definition 1. Let $\epsilon \geq 0$, $\delta \in [0, 1]$. Let Y be a space equipped with a binary symmetric relation Adj, called adjacency, and $(R, \mathcal{R})$ be a measurable space with $\mathcal{R}$ a given σ -algebra over R . A randomized mechanism M from Y to R is said to be (ϵ, δ) -dp if for all $y, y' \in Y$ such that $\text{Adj}(y, y')$,

$$\mathbb{P}(M(y) \in S) \leq e^\epsilon \mathbb{P}(M(y') \in S) + \delta, \quad \forall S \in \mathcal{R}.$$

In Definition 1, smaller values of ϵ and δ give stronger privacy guarantees, by forcing the distributions of $M(y)$ and $M(y')$ to be closer, for adjacent inputs y and y' .

Next, we introduce the Gaussian mechanism [7, 19], providing a simple method to produce dp outputs taking values in a vector space. It relies on the following quantity.

Definition 2 (ℓ_p -sensitivity). Let Y be a space equipped with an adjacency relation Adj. The ℓ_p -sensitivity of a mapping $G : Y \rightarrow \mathcal{S}^n$ is defined as

$$\Delta_p G = \sup_{y, y' : \text{Adj}(y, y')} \|G(y) - G(y')\|_p,$$

where $\|\cdot\|_p$ denotes the ℓ_p -norm on $\mathcal{S}^n$.

The basic Gaussian mechanism produces (ϵ, δ) -dp outputs by adding Gaussian noise proportional to the ℓ_2 -sensitivity of a mapping of interest.

Theorem 1 ([7]). Let $\epsilon > 0$, $\delta \in (0, 1)$, and $G : Y \rightarrow \mathcal{S}^n$. Define $k_{\epsilon, \delta} := \frac{1}{2\epsilon} \left(Q^{-1}(\delta) + \sqrt{(Q^{-1}(\delta))^2 + 2\epsilon} \right)$, where $Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-u^2/2} du$. Then, the mechanism $M(y) = G(y) + (\sigma k_{\epsilon, \delta}) \zeta$, where ζ is a standard WGN signal and σ is a scalar such that $\sigma \geq \Delta_2 G$ is (ϵ, δ) -dp.

2.3 Privacy-preserving state estimation problem

Consider now a privacy-sensitive discrete-time scalar-valued signal $\{y_k\}_{k \geq 0} \in \mathcal{S}^1$, which can be modeled as the output of an LTI state-space system

$$\begin{aligned} x_{k+1} &= Ax_k + Bw_k, \\ y_k &= c^T x_k + dv_k, \end{aligned} \tag{1}$$

where $x_k \in \mathbb{R}^n$ is the state at period $k \geq 0$, $x_0 \sim \mathcal{N}(\bar{x}_0, \Sigma_0^-)$ with $\Sigma_0^- \succ 0$, $w_k \in \mathbb{R}^p$ and $v_k \in \mathbb{R}$ for $k \geq 0$ are discrete-time standard WGN signals, independent of each other and of x_0 .

The modeling parameters $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times p}$, $c \in \mathbb{R}^n$, $d \in \mathbb{R}$, $\bar{x}_0 \in \mathbb{R}^n$, $\Sigma_0^- \in \mathbb{S}^n$ are assumed to be publicly known. We make the following assumption to ensure the well-posedness of steady-state state estimation problem.

Assumption 1. The pair (A, B) is stabilizable and the pair (A, c^T) is detectable.

Based on the measured signal y , we aim to publish at each time step $k \geq 0$, a causal estimate $\hat{z}_k \in \mathbb{R}^q$ of $z_k = Lx_k$, a linear transformation of the system state, where $L \in \mathbb{R}^{q \times n}$ is a publicly known matrix. In terms of accuracy, this estimate should minimize the steady-state mean-squared error (MSE)

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=0}^{T-1} \mathbb{E}[|z_k - \hat{z}_k|_2^2]. \quad (2)$$

In a standard setting, the optimal estimator would be $\hat{z}_k = L\hat{x}_k$ with $\hat{x}_k$ provided by the Kalman filter taking y as an input. However, because y is assumed to be privacy-sensitive, we also impose a DP constraint on the published signal $\hat{z}$.

In this paper, the space Y of Section 2.2 is taken to be $\mathcal{S}^1$ and we work with the following adjacency relation.

Definition 3. Let $r > 0$. The adjacency relation Adj_r is defined on $Y = \mathcal{S}^1$ as

$$\begin{aligned} &\text{Adj}_r(y, y') \text{ if and only if } \exists l \in \mathbb{N}, \rho \in \mathbb{R} \text{ with } |\rho| \leq r, \\ &\text{such that } y - y' = \rho \delta^l. \end{aligned} \quad (3)$$

An interpretation of Definition 3, which slightly generalizes the notion of event-level DP introduced in [5], is that we are interested in situations where a single individual contributes to the signal y at a single time instant by a maximum amount r . For instance, y could be event counts from a motion detector in an area where an individual is expected to pass at most once during the time frame of the data analysis.

2.4 A two-stage architecture for dp estimation

Theorem 1 can be used to design a class of dp estimators, following the architecture of Figure 1. There, the signal y is processed by an LTI pre-filter $G : \mathcal{S}^1 \rightarrow \mathcal{S}^1$, producing a new signal $\psi = G(y)$. Noise is then added to ψ to ensure that the resulting signal s is dp. Given the adjacency relation (3), the ℓ_2 -sensitivity of the pre-filter G is bounded by

$$\begin{aligned} \Delta_2 G &= \sup_{y, y' : \text{Adj}(y, y')} \|G(y) - G(y')\|_2 \\ &= \sup_{l \in \mathbb{N}, |\rho| \leq r} |\rho| \|G(\delta^l)\|_2 \leq r \|G\|_{\mathcal{H}_2}, \end{aligned}$$

because, by the Plancherel theorem, $\|G\|_{\mathcal{H}_2}^2$ represents the energy of the impulse response of G . Hence, by Theorem 1, adding WGN with variance $k_{\epsilon, \delta}^2 r^2 \|G\|_{\mathcal{H}_2}^2$ ensures that s is dp. Furthermore, transforming a dp signal does not alter the DP guarantee [6, Section 1.3.3]. Because the signal $s = c^T x + dv + (k_{\epsilon, \delta} r \|G\|_{\mathcal{H}_2}) \zeta$ with ζ a standard WGN, can still be viewed as the output of a linear Gaussian system, one should then use a Kalman filter to compute a dp estimate $\hat{z} = L\hat{x}$ from s , minimizing the MSE.

A special case is to take the pre-filter G to be the identity filter (with $\|G\|_{\mathcal{H}_2} = 1$) and perturb directly the private signal y . However, this so-called *input perturbation mechanism* [7] may fail to leverage the potential temporal redundancies in y when setting the noise level, and tends to produce estimators with poor accuracy, i.e., high MSE, see Section 4. In the rest of this paper we develop a methodology for designing pre-filters G within the architecture of Figure 1, generalizing the input perturbation scheme to produce dp estimates with improved accuracy.




Figure 1: Illustration of the proposed dp mechanism with two filters. ζ is a scalar standard WGN, which needs to be scaled by $k_{\delta, \epsilon} r \|G\|_{\mathcal{H}_2}$ according to Theorem 1 for the signal s to be dp. The dashed lines represent dp signals. Note that the input perturbation mechanism [7] is obtained by taking the pre-filter G to be the identity filter, i.e., $\psi = y$.

2.5 A motivating example

To motivate the pre-filtering approach, consider a system whose output signal y in the frequency domain exhibits two dominant peaks at ω_1 and ω_2 , see Figure 2. Suppose a data aggregator is interested in publishing an estimate of a quantity z associated with the component at frequency ω_1 .




Figure 2: Magnitude of the frequency response of the low-pass filter $|G(e^{j\omega})|$, and magnitude spectrum of the system output $|Y(e^{j\omega})|$.

Under input perturbation, the pre-filter is $G = 1$ and the privacy-preserving noise is distributed as $\zeta_k \sim \mathcal{N}(0, k_{\epsilon, \delta}^2 r^2)$. On the other hand, an ideal low-pass filter G with cut-off frequency ω_c satisfying $\omega_1 < \omega_c < \omega_2$, with unit magnitude below ω_c , preserves all the information needed for the data aggregator to estimate z , and requires a smaller amount of privacy-preserving noise because

$$\|G\|_{\mathcal{H}_2}^2 = \frac{1}{2\pi} \int_{-\omega_c}^{\omega_c} 1 d\omega = \frac{\omega_c}{\pi} < 1,$$

so that now $\zeta_k \sim \mathcal{N}(0, k_{\epsilon, \delta}^2 r^2 \omega_c / \pi)$.

3 FIR pre-filter optimization

Intuitively, to optimize the performance of the architecture of Figure 1, the signal y should be passed through a pre-filter G that both exhibits low sensitivity (to inject a minimal amount of noise) and preserves sufficient information for estimating z . In this section we develop a method to optimize this pre-filter, constraining it for tractability to be an FIR filter.

3.1 FIR pre-filter model

The output $\psi_k \in \mathbb{R}$ for $k \geq 0$ of an FIR filter of order m is a weighted sum of the current and past measurements

$$\psi_k = \xi y_k + \sum_{l=1}^m \lambda_l y_{k-l},$$

where $\xi, \lambda_1, \lambda_2, \dots, \lambda_m$ are the filter coefficients. Defining

$$\Gamma = \begin{bmatrix} 0 & 0 & \cdots & 0 & 0 \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \end{bmatrix}_{m \times m}, \quad \theta = \begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}_{m \times 1}, \quad (4)$$

$$\lambda^T = [\lambda_1 \quad \lambda_2 \quad \cdots \quad \lambda_m],$$

a state-space model of the FIR filter is given by

$$\begin{aligned} \eta_{k+1} &= \Gamma \eta_k + \theta y_k, \\ \psi_k &= \lambda^T \eta_k + \xi y_k, \end{aligned} \quad (5)$$

where $\eta_k \in \mathbb{R}^m$ represents the state vector with $\eta_0 = 0$, and $\xi \in \mathbb{R}$, $\lambda \in \mathbb{R}^m$, and $m \in \mathbb{N}$ are treated as design variables.

3.2 Kalman filter design and performance objective

To derive the equations for the Kalman filter on Figure 1, we define the augmented state vector $\phi_k = [x_k^T \quad \eta_k^T]^T$. With

$$\begin{aligned} E &= \begin{bmatrix} A & 0_{n \times m} \\ \theta c^T & \Gamma \end{bmatrix}, & F &= \begin{bmatrix} B & 0_{n \times 1} & 0_{n \times 1} \\ 0_{m \times p} & \theta d & 0_{m \times 1} \end{bmatrix}, \\ g^T &= [\xi c^T \quad \lambda^T], & h^T &= \begin{bmatrix} 0_{1 \times p} & \xi d & k_{\epsilon, \delta} r \|G\|_{\mathcal{H}_2} \end{bmatrix}, \end{aligned} \quad (6)$$

the augmented state-space model becomes

$$\begin{aligned} \phi_{k+1} &= E \phi_k + F \nu_k, \\ s_k &= g^T \phi_k + h^T \nu_k, \end{aligned} \quad (7)$$

with $\nu_k = [w_k^T \quad v_k \quad \zeta_k]^T$ and ζ a scalar standard WGN.

For the model (7), the state estimate $\hat{\phi}_k$ minimizing the steady-state MSE is given by the Kalman filter as

$$\hat{\phi}_{k+1} = E \hat{\phi}_k + \chi (s_k - g^T \hat{\phi}_k), \quad (8)$$

where $\chi \in \mathbb{R}^{n+m}$ is the steady-state Kalman gain

$$\chi = EP_{\infty}g(g^T P_{\infty}g + h^T h)^{-1}. \quad (9)$$

Here, P_{∞} denotes the steady-state error covariance matrix, which satisfies the discrete-time algebraic Riccati equation

$$\begin{aligned} P_{\infty} &= EP_{\infty}E^T + F^T F \\ &\quad - EP_{\infty}g(g^T P_{\infty}g + h^T h)^{-1}g^T P_{\infty}E^T. \end{aligned} \quad (10)$$

The corresponding estimation error $e_k = \phi_k - \hat{\phi}_k$ evolves according to

$$e_{k+1} = (E - \chi g^T) e_k + (F - \chi h^T) \nu_k. \quad (11)$$

The transfer function (parametrized by ξ, λ) from ν to e is

$$W(z; \xi, \lambda) = (zI - (E - \chi g^T))^{-1} (F - \chi h^T). \quad (12)$$

Let $\bar{L} = [L \ 0_{q \times m}]$. Given that $\hat{z}_k = L \hat{x}_k$ for the minimum MSE estimate of z_k , from (11) and the fact that ν is a standard WGN, the steady-state MSE in (2) is

$$\begin{aligned} \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=0}^{T-1} \mathbb{E}[\|\bar{L} \phi_k - \bar{L} \hat{\phi}_k\|_2^2] &= \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=0}^{T-1} \mathbb{E}[\|\bar{L} e_k\|_2^2] \\ &= \|\bar{L} W\|_{\mathcal{H}_2}^2, \end{aligned} \quad (13)$$

by a standard property of the $\mathcal{H}_2$ norm [20, Lemma 4.3.1]. Accordingly, the optimization problem of interest is to design the pre-filter parameters ξ and λ so as to minimize

$$\inf_{\xi \in \mathbb{R}, \lambda \in \mathbb{R}^m} \|\bar{L} W(\cdot; \xi, \lambda)\|_{\mathcal{H}_2}^2, \quad (14)$$

where choosing larger values of m provides more design freedom but increases the computational difficulty. Next, we give conditions under which a time-invariant Kalman filter stabilizing the error (11) can be designed [21].

Lemma 1. Under Assumption 1, the pair (E, F) is stabilizable. Moreover, if $\xi + \sum_{i=1}^m \frac{\lambda_i}{\alpha^i} \neq 0$ for any eigenvalue α of A such that $|\alpha| \geq 1$, then the pair (E, g^T) is detectable.

Proof. Since the eigenvalues of E consist of the eigenvalues of Γ , which are all 0s, and those of A , any eigenvalue of E that has magnitude greater or equal to 1 must correspond to an eigenvalue of A . In the case where all the eigenvalues of A have magnitude strictly less than 1, (E, F) is clearly stabilizable. Now, consider the case where A has some eigenvalue α with $|\alpha| \geq 1$. Then, any left eigenvector of E associated with α denoted by $u^T = [u_1^T, u_2^T] \neq 0$ necessarily satisfies

$$u^T E = [u_1^T A + u_2^T \theta c^T \quad u_2^T \Gamma] = [\alpha u_1^T \quad \alpha u_2^T], \quad (15)$$

and in particular $u_2^T \Gamma = \alpha u_2^T$. This implies that $u_2 = 0$, since $\alpha \neq 0$ cannot be an eigenvalue of Γ , and hence u_1^T is a left eigenvector of A based on (15). Furthermore, since (A, B) is stabilizable, $u^T F = [u_1^T B \ 0 \ 0]$ cannot be a zero vector. We conclude that (E, F) is stabilizable.

The detectability of the pair (E, g^T) when all the eigenvalues of A (and hence E) have magnitude strictly less than 1 is immediate. Thus, we only need to consider the case where A (and hence E) has some eigenvalue α with $|\alpha| \geq 1$. Let $u = [u_1^T, u_2^T]^T$ be an associated right eigenvector for E . Then $Eu = \alpha u$, that is,

$$Au_1 = \alpha u_1, \quad \theta c^T u_1 + \Gamma u_2 = \alpha u_2. \quad (16)$$

This implies that u_1 is an eigenvector of A for α . Hence, $c^T u_1 \neq 0$ under Assumption 1. The second part of (16) can be explicitly written as follows:

$$c^T u_1 = \alpha u_{2,1}, \quad u_{2,1} = \alpha u_{2,2}, \quad \dots, \quad u_{2,m-1} = \alpha u_{2,m},$$

which implies $u_2 = (c^T u_1) \begin{bmatrix} \frac{1}{\alpha} & \dots & \frac{1}{\alpha^m} \end{bmatrix}$. Now, for the eigenvector test for detectability, we consider

$$g^T u = \xi c^T u_1 + \lambda^T u_2 = (c^T u_1) \left(\xi + \sum_{i=1}^m \frac{\lambda_i}{\alpha^i} \right).$$

Since we already know that $c^T u_1 \neq 0$, the additional condition $(\xi + \sum_{i=1}^m \frac{\lambda_i}{\alpha^i}) \neq 0$ ensures that $g^T u \neq 0$, and so the pair (E, g^T) is detectable. $\square$

Remark 1. A consequence of Lemma 1 is that (E, g^T) is detectable for $\xi = 1$ and $\lambda = 0$, i.e., when the FIR filter is the identity filter and the scheme reduces to input perturbation.

3.3 Filter design optimization

We now focus on solving the optimization problem (14) numerically. Note that the FIR filter coefficients ξ and λ determine the value $\|G\|_{\mathcal{H}_2}$ appearing in h in (6). Specifically,

$$\|G\|_{\mathcal{H}_2} = \sqrt{\xi^2 + \lambda^T \lambda}. \quad (17)$$

Lemma 2. Multiplying a pre-filter G in Figure 1 by a nonzero scalar factor does not change the objective function in (14). Hence, without loss of generality, we can assume that the pre-filter satisfies $\|G\|_{\mathcal{H}_2} = 1$.

Proof. Let G be an LTI pre-filter with state-space realization $(\Gamma, \theta, \lambda^T, \xi)$ in (5). Let $\iota \neq 0$ and $\tilde{G} = \iota G$. Then, a state-space realization of $\tilde{G}$ is $(\Gamma, \theta, \iota \lambda^T, \iota \xi)$. The vectors g and h in (6) scale as $\tilde{g}^T = [\iota \xi c^T \quad \iota \lambda^T] = \iota g^T$ and $\tilde{h}^T = [0_{1 \times n} \quad \iota \xi d \quad k_{\epsilon, \delta} r^T \|\tilde{G}\|_{\mathcal{H}_2}] = \iota h^T$. Substituting $\tilde{g}$ and $\tilde{h}$ into (10) implies that the steady-state error covariance matrix does not depend on ι . Using (9), the Kalman gains corresponding to $\tilde{G}$ and G are related by $\tilde{\chi} = \frac{1}{\iota} \chi$. Hence, we have $\tilde{\chi} \tilde{g}^T = \chi g^T$ and $\tilde{\chi} \tilde{h}^T = \chi h^T$, and the transfer function in (12) and objective function in (14) do not depend on ι . As a result, we can replace any pre-filter G by $G/\|G\|_{\mathcal{H}_2}$, or equivalently, assume that $\|G\|_{\mathcal{H}_2} = 1$. $\square$

The next proposition formulates the design of the pre-filter and Kalman filter as a numerical optimization problem.

Proposition 1. Problem (14), together with the stability requirement of the error system (11), is equivalent to

$$\min_{\substack{Q \succ 0, Z \succ 0, \kappa \in \mathbb{R}^{n+m}, \\ \tau > 0, \xi, \lambda \in \mathbb{R}^m}} \tau, \quad \text{such that} \quad (18)$$

$$\text{Tr}(Z) < \tau, \quad (19)$$

$$\begin{bmatrix} Z & \bar{L} \\ * & Q \end{bmatrix} \succ 0, \quad (20)$$

$$\begin{bmatrix} Q & QE - \kappa g^T & QF - \kappa h^T \\ * & Q & 0_{(n+m) \times (p+2)} \\ * & * & I_{p+2} \end{bmatrix} \succ 0, \quad (21)$$

$$\begin{bmatrix} 1 & \xi & \lambda^T \\ * & I_{m+1} & \\ * & & \end{bmatrix} \succeq 0, \quad (22)$$

$$\text{with } g^T = [\xi c^T \quad \lambda^T], \quad h^T = [0_{1 \times p} \quad \xi d \quad k_{\epsilon, \delta} r^T],$$

and $*$ denoting symmetric terms in the matrices. We then set $\chi = (Q^*)^{-1} \kappa^*$ in (8), where Q^* and κ^* are part of an optimal solution for the problem above.

Proof. (*sketch*) According to Lemma 2, we can first impose the constraint $\|G\|_{\mathcal{H}_2} = 1$, hence removing the term $\|G\|_{\mathcal{H}_2}$ from h^T in (6). Next, minimizing the objective function in (14) is equivalent to minimizing the variable τ subject to the constraint $\|\bar{L}W\|_{\mathcal{H}_2}^2 < \tau$, where W is the transfer function given in (12). This constraint together with the internal stability of the dynamics (11) can be written equivalently as (19), (20), and (21), see [22, Section 4.3.2], where we defined $\kappa = Q\chi$.

Finally, we relax the constraint $\|G\|_{\mathcal{H}_2} = 1$ to $\|G\|_{\mathcal{H}_2} \leq 1$, which, given (17), can be written as (22) after taking a Schur complement. Then, it remains to show that an optimal solution G^* , corresponding to optimal parameters ξ^* and λ^* , necessarily satisfies $\|G^*\|_{\mathcal{H}_2} = 1$, i.e., that the relaxation is tight. This can be shown by contradiction. Indeed, if $\|G^*\|_{\mathcal{H}_2} < 1$, then there exists $\iota > 1$ such that the parameters $\tilde{\xi} = \iota \xi^*$ and $\tilde{\lambda} = \iota \lambda^*$, corresponding to the scaled filter $\tilde{G} = \iota G^*$, still satisfy (22), i.e., $\|\tilde{G}\|_{\mathcal{H}_2} < 1$. Hence, using $\tilde{G}$ together with WGN with the same variance $k_{\epsilon, \delta}^2 r^2$ on Figure 1 still ensures

DP. Intuitively, however, using $\tilde{G}$ instead of G^* strictly improves the signal-to-noise ratio, i.e., strictly decreases the MSE, which contradicts the fact that G^* is optimal. The formal proof of this statement is deferred to the full version of this paper. $\square$

3.4 Optimization algorithm

Due to the bilinear terms κg^T and κh^T , the optimization problem (18)-(22) involves a BMI constraint [18], which renders it non-convex. To address this issue, we apply an iterative alternating convex minimization approach. Although this method is not guaranteed to reach a globally optimal solution, it still serves as a useful heuristic in practice, see Section 4. It proceeds by iteratively solving convex subproblems through the following steps:

1. Variable splitting: Group the decision variables into two subsets $S_1 = \{\tau, Q, Z, \kappa\}$ and $S_2 = \{\tau, Q, Z, \xi, \lambda\}$ such that fixing one subset transforms the BMI into a linear matrix inequality (LMI).
2. Initialization: Set as initial values $\xi_0 = 1$ and $\lambda_0 = 0$ (corresponding to the input perturbation mechanism), and set $i = 0$.
3. Step 1 (optimize over S_1): Increment $i \leftarrow i + 1/2$. Minimize the objective τ_i over $(\tau_i, Q_i, Z_i, \kappa_i)$, using the fixed values $\xi_i = \xi_{i-1/2}$ and $\lambda_i = \lambda_{i-1/2}$ from the previous iteration.
4. Step 2 (optimize over S_2): Increment $i \leftarrow i + 1/2$. Minimize the objective τ_i over $(\tau_i, Q_i, Z_i, \xi_i, \lambda_i)$, using the fixed value $\kappa_i = \kappa_{i-1/2}$ from the previous iteration.
5. Repeat: Alternate between Steps 1 and 2 until the objective τ converges (in practice, until the change from one iteration to the next is small enough).

We have the following convergence result.

Proposition 2. The sequence of objective values $\{\tau_i\}_i$ generated by the iterative alternating minimization approach above is non-increasing, hence guaranteed to converge.

Proof. According to Remark 1, after initialization with $\xi = 1$ and $\lambda = 0$, the pair (E, g^T) is detectable. Minimization in Step 1 over the variables in S_1 corresponds to a standard semi-definite program for $\mathcal{H}_2$ -norm minimization and internal stabilization, which admits a solution when (E, g^T) is detectable. After this first step, at each optimization step, there is a feasible solution and the objective cannot increase, since the feasible solution from the previous step could be left unchanged. Since the sequence $\{\tau_i\}$ is also bounded below by zero, it is then guaranteed to converge. $\square$

4 Simulation example

Consider a scenario in which epidemiologists collect over time the number of individuals who recover from a disease, reported by a hospital. We consider a continuous-time SEIR epidemiological model [23] given by

$$\begin{aligned} \frac{dS}{dt} &= \Lambda - \mu S - \frac{\beta IS}{N}, & \frac{dE}{dt} &= \frac{\beta IS}{N} - (\mu + \alpha)E, \\ \frac{dI}{dt} &= \alpha E - (\gamma + \mu)I, & \frac{dR}{dt} &= \gamma I - \mu R, \end{aligned}$$

where S , E , I , and R denote the number of susceptible, exposed, infectious, and recovered individuals, respectively. Moreover, Λ , μ , β , α , and γ are the birth, natural death, transmission, exposed-to-infectious, and recovery rates, respectively. $N = S + E + I + R$ is the total population, which is assumed to be constant (i.e., $\Lambda = \mu N$). After linearization around the equilibrium point $[N \ 0 \ 0 \ 0]^T$, discretization using the forward Euler method with step size $\Delta t = 1$, and eliminating the state S

to avoid linear dependency in the state vector, the remaining states $\mathcal{X}_k = [E_k \ I_k \ R_k]^T$ evolve according to the dynamics (without process noise)

$$\mathcal{X}_{k+1} = \begin{bmatrix} 1 - (\mu + \alpha) & \beta & 0 \\ \alpha & 1 - (\gamma + \mu) & 0 \\ 0 & \gamma & 1 - \mu \end{bmatrix} \mathcal{X}_k.$$

The privacy-sensitive signal y of Section 2 is the data stream reported by the hospital to the epidemiologists. According to the adjacency relation (3), the difference between two adjacent data streams, i.e., the impact of adding or removing a single individual, should be limited to a single time step. Hence, the hospital should not release the total number of recovered individuals R_k (because one recovering individual then contributes a step function), but can instead report the number of *newly* recovered individuals $R_k - R_{k-1}$ at each step k . Assuming that each individual gets sick and recovers at most once, (3) is then satisfied for data streams differing by one individual with $r = 1$.

Due to modeling uncertainty and counting errors, we introduce process and measurement noise. By defining the state vector as $x_k = [E_k \ I_k \ R_k \ R_{k-1}]^T$, the system matrices become

$$A = \begin{bmatrix} 1 - (\mu + \alpha) & \beta & 0 & 0 \\ \alpha & 1 - (\gamma + \mu) & 0 & 0 \\ 0 & \gamma & 1 - \mu & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix},$$

$$c^T = [0 \ 0 \ 1 \ -1].$$

We use the parameter values $\mu = 1/(80 \times 365)$, $\beta = 0.06$, $\alpha = 1/7$, $\gamma = 1/14$, and choosing $B = \text{diag}(0.3162, 0.3162, 0.3162, 0.01)$ and $d = 0.3162$. It can be verified that (A, c^T) is detectable and (A, B) is stabilizable. The FIR filter order is set to $m = 10$. The iterative alternating optimization algorithm of Section 3.4 is implemented using CVX [24, 25] with SeDuMi [26] as the solver. The algorithm converges within a few iterations with the relative change in the objective falling below a prescribed tolerance 10^{-4} . The epidemiologists want to provide at each period k an estimate of the number of infectious individuals, hence $L = [0 \ 1 \ 0 \ 0]$. Figure 3 shows that the proposed FIR-Kalman architecture achieves significantly lower steady-state MSE than the input perturbation mechanism for different values of ϵ , with $\delta = 0.05$. Finally, Figure 4 illustrates the transient evolution of the estimation error for the infectious population under the input perturbation mechanism and the proposed FIR-Kalman architecture.




Figure 3: Steady-state MSE of input perturbation and the FIR-Kalman architecture as a function of the privacy parameter ϵ . Here $\delta = 0.05$.




Figure 4: Estimation errors of the infectious population with ± 2 standard deviation bounds for the dp approaches based on input perturbation and the FIR-Kalman architecture. The privacy parameters are $\epsilon = 1$ and $\delta = 0.05$.

5 Conclusions

This paper introduced a two-stage FIR-Kalman architecture for achieving event-level DP in state estimation based on a privacy-sensitive data stream. By combining an FIR filter with a Kalman filter, our approach generalizes the input perturbation mechanism. We formulated the joint filter design problem as an optimization problem and demonstrated in a simulated example an improved trade-off between privacy level and estimation accuracy.

References

- [1] X. Feng and C. Zhang, MPLDP: Multi-level personalized local differential privacy method, *IEEE Access*, 12, 99, 739–99 754, 2024.
- [2] Y. Cao, M. Yoshikawa, Y. Xiao, and L. Xiong, Quantifying differential privacy in continuous data release under temporal correlations, *IEEE Transactions on Knowledge and Data Engineering*, 31(7), 1281–1295, 2019.
- [3] X. Yan, G. Zhou, D. E. Quevedo, C. Murguia, B. Chen, and H. Huang, Privacy-preserving state estimation in the presence of eavesdroppers: A survey, *IEEE Transactions on Automation Science and Engineering*, 22, 6190–6207, 2025.
- [4] C. Dwork and A. Roth, The algorithmic foundations of differential privacy, *Foundations and Trends in Theoretical Computer Science*, 9(3–4), 211–407, 2014.
- [5] C. Dwork, M. Naor, T. Pitassi, and G. N. Rothblum, Differential privacy under continual observations, in *Proceedings of the ACM Symposium on the Theory of Computing (STOC)*, Cambridge, MA, June 2010.
- [6] J. Le Ny, *Differential Privacy for Dynamic Data*. Springer, 2020.
- [7] J. Le Ny and G. J. Pappas, Differentially private filtering, *IEEE Transactions on Automatic Control*, 59(2), 341–354, February 2014.
- [8] E. Nozari, P. Tallapragada, and J. Cortés, Differentially private average consensus: Obstructions, trade-offs, and optimal algorithm design, *Automatica*, 81, 221–231, 2017.
- [9] Y. Wang and A. Nedić, Differentially-private distributed optimization with guaranteed optimality, in *IEEE Conference on Decision and Control (CDC)*, Singapore, 2023.
- [10] K. Yazdani, A. Jones, K. Leahy, and M. Hale, Differentially private LQ control, *IEEE Transactions on Automatic Control*, 68(2), 1061–1068, 2023.
- [11] Y. Wang, Z. Huang, S. Mitra, and G. E. Dullerud, Differential privacy in linear distributed control systems: Entropy minimizing mechanisms and performance tradeoffs, *IEEE Transactions on Control of Network Systems*, 4(1), 118–130, 2017.
- [12] Y. Kawano and M. Cao, Design of privacy-preserving dynamic controllers, *IEEE Transactions on Automatic Control*, 65(9), 3863–3878, 2020.

- [13] X. Yan, B. Chen, Y. Zhang, and L. Yu, Guaranteeing differential privacy in distributed fusion estimation, *IEEE Transactions on Aerospace and Electronic Systems*, 59(3), 3416–3423, 2023.
- [14] C. Li, M. Hay, V. Rastogi, G. Miklau, and A. McGregor, Optimizing linear counting queries under differential privacy, in *Principles of Database Systems (PODS)*, 2010.
- [15] K. H. Degue and J. Le Ny, Differentially private Kalman filtering with signal aggregation, *IEEE Transactions on Automatic Control*, 68(10), 6240–6246, 2023.
- [16] J. Le Ny, On differentially private filtering for event streams, in *Proceedings of the Conference on Decision and Control (CDC)*, Florence, Italy, December 2013.
- [17] J. Le Ny and M. Mohammady, Differentially private MIMO filtering for event streams, *IEEE Transactions on Automatic Control*, 63(1), 145–157, 2018.
- [18] J. G. VanAntwerp and R. D. Braatz, A tutorial on linear and bilinear matrix inequalities, *Journal of process control*, 10(4), 363–385, 2000.
- [19] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, Our data, ourselves: Privacy via distributed noise generation, in *Proceedings of the 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, St. Petersburg, Russia, May 2006, pp. 486–503.
- [20] T. Chen and B. A. Francis, *Optimal Sampled-Data Control Systems*. Springer-Verlag, 1995.
- [21] B. Anderson and J. B. Moore, *Optimal filtering*. Prentice-Hall, 1979.
- [22] R. J. Caverly and J. R. Forbes, LMI properties and applications in systems, stability, and control theory, *arXiv:1903.08599*, 2019.
- [23] L. J. Allen, F. Brauer, P. Van den Driessche, and J. Wu, *Mathematical Epidemiology*. Springer, 1945, 2008.
- [24] CVX Research, Inc., CVX: Matlab software for disciplined convex programming, version 2.0, Available: <https://cvxr.com/cvx>, 2012.
- [25] M. Grant and S. Boyd, Graph implementations for nonsmooth convex programs, in *Recent Advances in Learning and Control*. Springer, 2008, pp. 95–110.
- [26] J. F. Sturm, Using sedumi 1.02, a MATLAB toolbox for optimization over symmetric cones, *Optimization Methods and Software*, 11(1–4), 625–653, 1999.