Received: 18 February 2018 | Revised: 23 July 2018 | Accepted: 1 October 2018

DOI: 10.1002/rnc.4392

SPECIAL ISSUE ARTICLE

WILEY

Differentially private nonlinear observer design using contraction analysis

Jerome Le Ny

Department of Electrical Engineering and GERAD, Polytechnique Montreal, Montréal, Canada

Correspondence
Jerome Le Ny, Department of Electrical Engineering and GERAD, Polytechnique Montreal, C.P. 6079, Succursale Centre-ville, Montréal, QC H3C 3A7, Canada.
Email: jerome.le-ny@polymtl.ca

Funding information
Natural Sciences and Engineering Research Council of Canada, Grant/Award Number: RGPIN-5287-2018 and RGPAS-2018-522686

Summary

Real-time information processing applications such as those enabling a more intelligent infrastructure are increasingly focused on analyzing privacy-sensitive data obtained from individuals. To produce accurate statistics about the habits of a population of users of a system, this data might need to be processed through model-based estimators. Moreover, models of population dynamics, originating for example from epidemiology or the social sciences, are often necessarily nonlinear. Motivated by these trends, this paper presents an approach to design nonlinear privacy-preserving model-based observers, relying on additive input or output noise to give differential privacy guarantees to the individuals providing the input data. For the case of output perturbation, contraction analysis allows us to design convergent observers as well as set the level of privacy-preserving noise appropriately. Two examples illustrate the proposed approach: estimating the edge formation probabilities in a social network using a dynamic stochastic block model, and syndromic surveillance relying on an epidemiological model.

KEYWORDS

differential privacy, nonlinear filtering, nonlinear observer design, privacy-preserving data analysis

1 | INTRODUCTION

The possibility to analyze vast amounts of personal data capturing information about the activities of private individuals is a foundational principle behind* many current technology-driven trends such as the "Internet of Things," electronic biosurveillance systems, or developing an intelligent infrastructure enabling smart cities. In many respects, however, the data collection practices envisioned to operate these systems often go against basic privacy rights.[2] Concerns about the acquisition and use of personal data by companies and governments, eg, for potential price and service discrimination, are rising[3-5] and could lead to people rejecting these technologies despite their suggested benefits. Rigorous privacy-preserving data analysis methodologies are needed to support regulations and allow people to appropriately trade off the privacy risks they increasingly incur with the benefits they can expect in return.

Typically, large-scale monitoring and control systems only require aggregate statistics computed from personal data streams, eg, a dynamic map showing road traffic conditions built from location traces sent by smartphones or an estimate of power consumption in a neighborhood updated using smart meter data from individual homes. Aggregation is beneficial to privacy, but past examples have shown that this is not sufficient to a priori rule out the possibility of significant

*A preliminary version of this paper was presented at CDC 2015.[1]

privacy breaches.[6-8] Privacy attacks are often *linkage attacks*, where some newly published information is combined with other available data in order to make new inferences about specific individuals, and predicting at system design time how any such attack could be carried out is very difficult. Yet, as explained later, it is still possible to compute aggregate statistics with formal privacy guarantees for the individuals from whom the data originates, which could help alleviate some of the justified concerns and encourage wider adoption of certain pervasive sensing and control systems.

Various information theoretic definitions have been proposed to capture quantitatively the concept of privacy and are potentially applicable to the processing of data streams in real time.[9] In this paper, we focus on the notion of *differential privacy*, which originates from the database and cryptography literature.[10] Intuitively, a differentially private mechanism publishes information about a data set in a way that is not too sensitive to a single individual's data. As a result, an individual receives a guarantee that whether or not she decides to provide her data will not dramatically change the ability of a third party to make new inferences about her. Previous work has considered the design of linear filters processing sensitive time series data with differential privacy guarantees.[11-17] The problem studied in this paper is that of designing privacy-preserving nonlinear model-based estimators, which, to the best of our knowledge, has not been studied in a general setting before. A convenient way of achieving differential privacy for an estimator is to bound its so-called *sensitivity*,[10] a form of incremental system gain between the private input signals and the published output.[15] Various tools could be used for this purpose, and here, we rely on contraction analysis.[18-21] This idea was proposed in a preliminary version of this paper.[1] Here, we extend our previous analysis in particular to a more general type of observer and to contraction with respect to more general metrics, including Riemannian metrics, enlarging the class of systems for which a contraction property holds and for which we can bound the sensitivity as a result. Moreover, by using different proof strategies, we can provide tighter sensitivity bounds.

The rest of this paper is divided as follows. Section 2 presents the problem statement formally, provides a brief introduction to the notion of differential privacy, and describes and compares privacy-preserving data analysis mechanisms with input and output perturbations. In Section 3, we discuss some fundamental results in contraction analysis and present a type of "input-to-state stability" property of contracting systems similar to the one proved in the work of Sontag[19] but stated here for discrete-time systems. This property is used in Section 4 to design differentially private observers with output perturbation. The methodology is illustrated via two examples involving the analysis of dynamic data originating from private individuals. In Section 5.1, we consider the problem of estimating link formation probabilities in a social network using a dynamic version of the classical stochastic block model,[22] which involves a nonlinear measurement model. In Section 5.2, we consider a nonlinear epidemiological model and design a differentially private estimator of the proportion of susceptible and infectious people in a population, assuming a syndromic data source.

Notation. The expressions "if and only if" and "independent and identically distributed" are abbreviated as iff and iid, respectively. $\mathbb{N} := \{0, 1, \dots\}$ denotes the set of nonnegative integers, C^1 is the set of continuously differentiable functions, and a class $\mathcal{K}$ function $\beta : \mathbb{R}_+ \to \mathbb{R}_+$ is a strictly increasing continuous function such that $\beta(0) = 0$. For $H : \mathsf{X} \to \mathsf{Y}$ a linear map between finite dimensional vector spaces X and Y equipped with the norms $|\cdot|_{\mathsf{X}}$ and $|\cdot|_{\mathsf{Y}}$ respectively, we denote by $\|H\|_{\mathsf{Y}}^{\mathsf{X}}$ its induced norm, so that $|Hx|_{\mathsf{Y}} \leq \|H\|_{\mathsf{Y}}^{\mathsf{X}} |x|_{\mathsf{X}}$, for all x in X. If $\mathsf{X} = \mathsf{Y}$ and both spaces are equipped with the same norm $|\cdot|_{\mathsf{X}}$, we simply write $\|\cdot\|_{\mathsf{X}}$. For $1 \leq p < \infty$, the p-norm on $\mathbb{R}^n$, denoted $|\cdot|_p$, is defined as $|v|_p := (\sum_{i=1}^n |v_i|^p)^{1/p}$, and $|v|_\infty := \max_{1\leq i\leq n} |v_i|$. For $v = \{v_k\}_{k\in\mathbb{N}}$ a vector-valued discrete-time signal, where $v_k \in \mathbb{R}^n$ has components $\{v_{k,i}\}_{i=1}^n$, the ℓ_p signal norm is $\|v\|_p = (\sum_{k=0}^\infty \sum_{i=1}^n |v_{k,i}|^p)^{1/p} = (\sum_{k=0}^\infty |v|_p^p)^{1/p}$ for $1 \leq p < \infty$, and $\|v\|_\infty = \sup_{k\geq 0} |v_k|_\infty$. We use $\mathrm{diag}(v)$ to denote a diagonal matrix with the components of the vector v on the diagonal. For P a symmetric matrix, P positive definite is denoted $P \succ 0$ and P positive semidefinite is denoted $P \succeq 0$. For $P \succeq 0$, we denote its (unique) positive semidefinite square root as $P^{1/2}$, ie, $P = P^{1/2}P^{1/2}$. For P, Q symmetric matrices, $P \succeq Q$ means $P - Q \succeq 0$, and $P \preceq 0$ means $-P \succeq 0$.

2 | PROBLEM STATEMENT

2.1 | Observer design

Consider the problem of estimating a discrete-time signal denoted $x := \{x_k\}_{k\in\mathbb{N}}$, with $x_k \in \mathsf{X} = \mathbb{R}^n$ for some positive integer n, which represents an aggregate state for a population of privacy-sensitive individuals. For example, x_k could be the density at period k of drivers or pedestrians at a finite number of spatial locations, the proportion of individuals infected by a disease in a population, etc. We assume that x_k cannot be perfectly observed but that we can measure instead

a privacy-sensitive discrete-time signal $\{y_k\}_{k\in\mathbb{N}}$, with $y_k \in \mathsf{Y} = \mathbb{R}^m$ for some positive integer m, for which we have a state-space model of the form

$$x_{k+1} = f_k(x_k) + w_k, \tag{1}$$

$$y_k = g_k(x_k) + v_k, \tag{2}$$

where w_k, v_k are noise signals representing modeling and measurement errors, and f_k and g_k are $\mathcal{C}^1$ functions. Note that the case where $y_k = x_k + v_k$ is possible, when the privacy-sensitive signal is a direct measurement of the underlying state. Our aim is to publish an estimate z_k of x_k, computed from y_k by an observer of the following form[23]:

$$z_{k+1} = f_k(z_k) + h_k(z_k, y_k - g_k(z_k)), \tag{3}$$

with, for each k in $\mathbb{N}$, $h_k : \mathsf{X} \times \mathsf{Y} \to \mathsf{X}$ a $\mathcal{C}^1$ function such that $h_k(x, 0) = 0$. We initialize (3) with some estimate z_0 of x_0. Note that (3) could describe an observer for a model (1)-(2) that has already been transformed under a suitable change of coordinates to a form that facilitates observer design, eg, an observability canonical form.[24,25] With straightforward modifications to our arguments, the "prediction" form (3) could also be replaced by an observer using the most recent observations

$$\begin{aligned} z_0 &= \bar{z}_0 + h_0(\bar{z}_0, y_0 - g_k(\bar{z}_0)), \quad \text{for some estimate } \bar{z}_0 \text{ of } x_0, \\ z_{k+1} &= f_k(z_k) + h_{k+1}(z_k, y_{k+1} - g_k(f_k(z_k))), \text{ for } k \geq 0. \end{aligned} \tag{4}$$

In the applications discussed later in the paper, the signal y_k is collected from privacy-sensitive individuals, hence needs to be protected, in a sense defined below. For the examples of states x_k mentioned above, y_k could consist of location traces or be the number of people presenting certain symptoms visiting emergency rooms for instance. On the other hand, model (1)-(3), ie, the functions f_k, g_k, and h_k, is assumed to be publicly available or at least could be potentially known to an adversary trying to make inferences about y based on z. The data aggregator wishes to publicly release the signal z produced by (3). However, since z depends on the sensitive signal y, we only allow the release of an approximate version of z carrying certain privacy guarantees, which are presented formally in the next section. As a result, it will emerge that the functions h_k need to be carefully chosen to balance accuracy or convergence speed of the observer with the level of privacy offered.

Remark 1. We do not provide here nor use any model of the noise signals w and v in (1), (2), which are simply introduced as a device to explain the discrepancy between any measured signal y and the signals that can be predicted by a deterministic model $x_{k+1} = f_k(x_k), y_k = g_k(x_k)$.

Remark 2. More generally, we might just want to publish an output $\chi_k(x_k)$, function of the state x_k. As explained below, this can be done by first obtaining a privacy-preserving estimate $\hat{x}$ of the signal x and then publishing $\chi_k(\hat{x}_k)$, relying on the fact that sound privacy guarantees such as differential privacy are preserved by the final transformation through χ_k.

2.2 | Differential privacy

The published signal should provide an accurate estimate of x under an additional constraint that is not satisfied a priori by z from (3), aiming at preserving the privacy of the individuals from which the measured signal y originates. More precisely, we impose that the published signal be differentially private,[10] which requires adding artificial noise somewhere in the signal processing system to randomize the published output. A differentially private version of observer (3) should produce a randomized output signal whose distribution is not too sensitive to certain variations associated with the effect of any individual's data on the signal y, input of the observer. The formal definition of differential privacy is given in Definition 1 and requires that we specify the type of variations in y that should be hard to detect from the published output. This is done by defining a symmetric binary relation, called adjacency and denoted Adj, on the space of data sets D of interest, here the space of signals y, so that two adjacent input signals y and $\tilde{y}$ should produce (randomized) output signals with similar distributions. It is possible to define different adjacency relations[15] to model different data analysis scenarios. In this paper, y is assumed to represent a (possibly multidimensional) signal that already aggregates the data obtained from multiple users, eg, y_k at a particular time period k could be the number of people waiting in a hospital

emergency room, the total power consumption of a group of homes during that period, etc. We then consider in particular the following adjacency relations between signals:

$$\text{Adj}(y, \tilde{y}) \text{ iff } \|y - \tilde{y}\|_p \leq B_p, \tag{5}$$

for $p = 1$ or $p = 2$ and some given fixed constant $B_p > 0$, as well as the more restrictive adjacency relation

$$\text{Adj}(y, \tilde{y}) \text{ iff } \exists k_0 \geq 0 \text{ s.t. } \begin{cases} y_k = \tilde{y}_k, & k < k_0 \\ |y_k - \tilde{y}_k|_p \leq K\alpha^{k-k_0}, & k \geq k_0, \end{cases} \tag{6}$$

where again $p = 1$ or $p = 2$ and $K > 0, 0 \leq \alpha < 1$ are given fixed constants. In other words, we aim at hiding deviations in the signal y (eg, due to the contribution of one individual to the signal) that are bounded in p-norm (relation (5)) or more explicitly that can start at any time k_0 but then subsequently decrease geometrically (relation (6)). Note that even the more restrictive condition (6) is much more general than the adjacency relation considered in some previous work on the design of a differentially private counter,[11,12,14] where adjacent (scalar) signals can differ at a single time period by at most one. In comparison, the adjacency condition (6) greatly enlarges the set of signal deviations that can result from the presence of any individual and for which we provide guarantees (deviation at a single period is obtained for $\alpha = 0$). We can now state the definition of a differentially private mechanism, ie, of a randomized map from input to output signals.

Definition 1. Let D be a space equipped with a symmetric binary relation denoted Adj, and let $(\mathsf{R}, \mathcal{M})$ be a measurable space, where $\mathcal{M}$ is a given σ-algebra over R. Let $\epsilon, \delta \geq 0$. A randomized mechanism M from D to R is (ϵ, δ)-differentially private (for Adj) if for all $d, d' \in \mathsf{D}$ such that $\text{Adj}(d, d')$, we have

$$\mathbb{P}(M(d) \in S) \leq e^{\epsilon}\mathbb{P}(M(d') \in S) + \delta, \quad \forall S \in \mathcal{M}. \tag{7}$$

If $\delta = 0$, the mechanism is said to be ϵ-differentially private.

This definition quantifies the admissible deviations for the output distribution of a differentially private mechanism, when a variation at the input satisfies the adjacency relation. Smaller values of ϵ and δ correspond to stronger privacy guarantees. In this paper, the space D was defined as the space of input signals y, the adjacency relation considered is (5) or (6), and the output space R is the space of output signals for the observer, here $\mathsf{X}^{\mathbb{N}}$ since we wish to estimate x. The problem is to publish an accurate (but randomized) estimate of the state x while satisfying the property of Definition 1 for specified values of ϵ and δ.

Remark 3. Definition 1 depends on the choice of σ-algebra $\mathcal{M}$, which must contain enough sets S to provide a meaningful differential privacy guarantee. The interested reader can find a discussion of measurability issues in a previous paper.[15]

2.3 | Sensitivity and basic differentially private mechanisms

Enforcing differential privacy can be done by randomly perturbing the published output of a system[10,15] at the expense of its quality or utility. Hence, we are interested in evaluating as precisely as possible the amount of noise necessary to make a mechanism differentially private. For this purpose, the following quantity plays an important role.

Definition 2. Let $q \geq 1$. The ℓ_q-sensitivity of a system G with m inputs and n outputs, with respect to an adjacency relation Adj on the input signals, is defined by $\Delta_q G = \sup_{\text{Adj(u,u')}} \|Gu - Gu'\|_q$.

In practice, we are interested in the sensitivity of a system for the cases $q = 1$ and $q = 2$. The basic mechanisms of Theorem 1 (with proofs and references in the previous paper[15]) can be used to produce differentially private signals. First, we need the following definitions. A zero-mean Laplace random variable with parameter b has the probability density function $\exp(-|x|/b)/2b$, and its variance is $2b^2$. The Q-function is defined as $Q(x) := \frac{1}{\sqrt{2\pi}} \int_x^{\infty} e^{-\frac{u^2}{2}} du$. Then, for $\epsilon > 0$, $0.5 \geq \delta > 0$, define $\kappa_{\delta,\epsilon} = \frac{1}{2\epsilon}(Q^{-1}(\delta) + \sqrt{(Q^{-1}(\delta))^2 + 2\epsilon})$, which can be shown to behave roughly as $O(\sqrt{\ln(1/\delta)}/\epsilon)$.

Theorem 1. *Let G be a system with m inputs and n outputs and fix a relation Adj in Definition 2. The mechanism $Mu = Gu + w$, where all $w_{k,i}, k \in \mathbb{N}, 1 \leq i \leq n$, are independent Laplace random variables with parameter $b \geq (\Delta_1 G)/\epsilon$, is ϵ-differentially private for Adj. If w is instead a white Gaussian noise such that the covariance matrix of each sample w_k is $\sigma^2 I_n$ with $\sigma \geq \kappa_{\delta,\epsilon}\Delta_2 G$, then the mechanism is (ϵ, δ)-differentially private.*

The mechanisms of Theorem 1 are called the Laplace and the Gaussian mechanism. One reason for introducing the Gaussian mechanism is that typically the ℓ_2-sensitivity is smaller than its ℓ_1 counterpart, which leads to lower noise levels if one can tolerate $\delta > 0$ in the privacy guarantee (7).

2.4 | Input and output perturbation

Theorem 1 says that we can obtain a differentially private signal at the output of a system G by adding noise with standard deviation proportional to $\Delta_1 G/\epsilon$ or to $\kappa_{\delta,\epsilon}\Delta_2 G$. A very useful additional result stated here informally says that postprocessing a differentially private signal without re-accessing the privacy-sensitive input signal does not change the differential privacy guarantee.[15] Now, system G in Theorem 1 can simply be the identity with ℓ_1- and ℓ_2- sensitivity for the adjacency relation (6) equal to $K/(1-\alpha)$ and $K/\sqrt{1-\alpha^2}$ respectively (or B_1 and B_2 for (5)). This immediately gives a first possible design approach for our privacy-preserving observer by simply adding Laplace or Gaussian noise directly to the input signal y (see Figure 1A). The observer can then be designed according to any desired methodology and should try to mitigate the effect of the artificial input noise, whose distribution is known, in addition to the usual measurement error. We call this design an input perturbation mechanism. Note that for α close to 1, $1/\sqrt{1-\alpha^2}$ is significantly smaller than $1/(1-\alpha)$, so that if we are willing to accept some $\delta > 0$ in the privacy guarantee and to use the 2-norm on Y in the adjacency relation (6), we can obtain much better accuracy by using the ℓ_2-sensitivity.

The input perturbation mechanism is attractive for its simplicity and might perform well, especially with low privacy level requirements (relatively high ϵ, δ). In particular, the sensitive data can be made differentially private at the source before sending it to any third party for processing. However, it can also potentially exhibit the following drawbacks. First, the noise added to y might be unnecessarily large because it is not tailored to the task of estimating the state x of model (1)-(2) and does not take into account the temporal correlations between samples of the signal y captured by this model. Significant noise at the input of the observer can also lead to poor performance, ie, slow convergence and large errors in the state estimate, or even perhaps divergence of the estimate from the true state trajectory, since the convergence of nonlinear observers is often local. Second, characterizing the output error (state estimation error) due to the privacy-preserving noise requires understanding how this noise is transformed when passing through the nonlinear observer. In general, for nonlinear systems, the noise distribution at the output can become multimodal and nonzero mean, and hence, the observer could produce a systematically biased estimate that could be hard to correct.

An alternative to input perturbation is the output perturbation mechanism shown on Figure 1B. In this case, following Theorem 1, a privacy-preserving noise signal proportional to the sensitivity of observer G is added at its output. Computing the sensitivity of G, or in practice upper bounding it, should be done as accurately as possible to reduce the conservatism of the approach. On the other hand, the output noise does not impact any stability or bias analysis of observer G. As discussed in more details in the following sections, we should then try to design an observer that has both good tracking performance for the state trajectory and low sensitivity in order to minimize the level of privacy-preserving noise necessary at the output. These two desired properties are essentially in conflict. Figure 1B shows that we can also add a terminal filter to smooth out the Laplace or Gaussian noise,[26] although this can generally affect the transient performance of the overall system (eg, its convergence speed). We do not discuss the design of a potential smoothing filter in this paper, except briefly in Section 5.1.

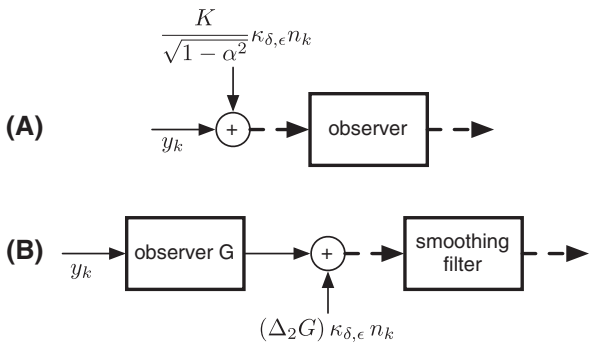

FIGURE 1 Gaussian mechanisms with input (A) and output (B) perturbation for the adjacency relation (6). n_k represents a zero-mean standard white Gaussian noise with identity covariance matrix. Dashed lines represent a differentially private signal

Example 1. Consider the memoryless system $y_k \mapsto \phi(y_k) := y_k^2$, which could be a simple state estimator for a measurement model $y_k = \sqrt{x_k}$ in (2), not taking the dynamics (1) into account. Consider the adjacency relation (6) for $\alpha = 0$, so that we have a deviation at some (unknown) single time period k_0 of at most K between adjacent signals y_k and $\tilde{y}_k$. For the input perturbation scheme and the Gaussian mechanism, assuming for simplicity that $\kappa_{\delta,\epsilon} = 1$, the signal $z_k = (y_k + K\xi_k)^2 = y_k^2 + 2Ky_k\xi_k + K^2\xi_k^2$ is differentially private when ξ is a standard Gaussian white noise. The privacy-preserving noise at the input induces a systematic bias at the output between z_k and y_k^2 equal to $\mathbb{E}[2Ky_k\xi_k + K^2\xi_k^2] = K^2$. Since K is assumed publicly known, in this case, the bias can be compensated and a possibly better approximation of ϕ that is still differentially private is $z'_k = (y_k + K\xi_k)^2 - K^2$. One can verify that the variance of the remaining error is $e'_k = \mathbb{E}[(z'_k - y_k^2)^2] = 4K^2y_k^2 + 2K^4$.

Suppose we know in addition that $y_k \in [0, 1]$ for all $k \geq 0$. Then, we can bound the sensitivity of the memoryless system as

$$\Delta_2\phi = \left|y_{k_0}^2 - \tilde{y}_{k_0}^2\right| = |y_{k_0} - \tilde{y}_{k_0}||y_{k_0} + \tilde{y}_{k_0}| \leq 2|y_{k_0} - \tilde{y}_{k_0}| \leq 2K. \tag{8}$$

Hence, the signal $z''_k = y_k^2 + 2K\xi_k$ is also differentially private and unbiased, with ξ a standard white Gaussian noise as before. The variance of the error is $e''_k = \mathbb{E}[(z''_k - y_k^2)^2] = 4K^2$, which is smaller than the worst-case value $4K^2 + 2K^4$ for e'_k However, e''_k is larger than e'_k as soon as $y_k < \sqrt{1 - K^2/2}$, the typical case since K should be much less than one, otherwise both the input and output mechanisms essentially destroy the signal. The upper bound (8) on the sensitivity is conservative in order to be independent of the actual values of the sensitive signal y, which is necessary when Theorem 1 is used to provide a differential privacy guarantee.

In the rest of this paper, we focus on the output perturbation mechanism of Figure 1B. There are two aspects to the differentially private observer design problem in this case. First, we need to enforce appropriate convergence of z toward x, which is the observer design problem itself. Second, we also need to control and bound explicitly the magnitude of the changes in z when the observer input changes from y to an adjacent signal $\tilde{y}$, in order to apply Theorem 1 and set the output noise level providing the differential privacy guarantee. In this paper, both aspects of the problem are treated by using contraction analysis to design the observer as well as quantify its sensitivity to variations in the measured signal y. A motivation for this approach is the exponential convergence of trajectories of contractive systems toward each other, which provides a degree of robustness against input disturbances[18,19,27,28] and, as a consequence, sensitivity bounds for variations in input data streams y. The next section provides some background on contraction analysis that is necessary to describe our design approach in Section 4.

3 | CONTRACTING SYSTEMS

Contraction analysis is an "incremental" stability analysis methodology for dynamical systems emphasizing convergence of trajectories toward each other, popularized in particular by the work of Lohmiller and Slotine.[18] Earlier related work can also be found in the mathematics literature.[29,30] Contraction and incremental stability analysis have seen significant developments in the past two decades,[18-21,27,31-33] and we refer the reader to the recent paper by Forni and Sepulchre[20] for a comparison of different variations that have emerged and additional references. The purpose of this section is to review some aspects of this methodology for discrete-time systems, which are not emphasized as much as continuous-time systems in the literature and to state and prove some results that we rely on to design differentially private observers with output perturbation. Although these results could potentially be derived from the ideas presented in some of the references cited above, we provide here a self-contained discussion and, in particular, *explicit* bounds on distances between trajectories that are necessary to precisely set the level of privacy-preserving noise since qualitative guarantees of incremental convergence are insufficient.

3.1 | Basic results

Consider a discrete-time system

$$x_{k+1} = f_k(x_k), \tag{9}$$

with $f_k : \mathsf{X} \to \mathsf{X}$ a C^1 function, for all $k \in \mathbb{N}$. Let us denote by $\phi(k; k_0, x_0)$ the value at time $k \geq k_0$ of the solution of (9) taking the value x_0 at time k_0. A forward invariant set for system (9) is a set $C \subset \mathsf{X}$ such that if $x_0 \in C$, then for all k_0 and

all $k \geq k_0$, $\phi(k;k_0,x_0) \in C$. Although we assume in this paper $\mathsf{X} = \mathbb{R}^n$, it is useful to introduce here some language from differential geometry and view X more generally as an n-dimensional differentiable manifold.[20,34] For each point $x \in \mathsf{X}$, the tangent space to X at x, ie, informally, the n-dimensional vector space of all tangent vectors to curves on X passing through x, is denoted $\mathsf{T}_x\mathsf{X}$. The tangent bundle of X is denoted $\mathsf{TX} := \cup_{x\in\mathsf{X}}\{x\} \times \mathsf{T}_x\mathsf{X}$ and is equipped with a time-varying family of norms $|\cdot|_{[x,k]}$, smoothly varying with x for each k, so that $|\cdot|_{[x,k]}$ is a norm on $\mathsf{T}_x\mathsf{X}$, for all $k \in \mathbb{N}$. For each $x, \tilde{x} \in \mathsf{X}$, let $\Gamma(x,\tilde{x})$ be the set of piecewise C^1 curves joining x and $\tilde{x}$, ie, functions $\gamma : [0,1] \to \mathsf{X}$ with $\gamma(0) = x$, $\gamma(1) = \tilde{x}$. We define the (time-varying) length of such a curve γ by[20,35]

$$L_k(\gamma) = \int_0^1 |\gamma'(r)|_{[\gamma(r),k]} dr,$$

where $\gamma'(r) := \frac{d\gamma}{dr}(r) \in \mathsf{T}_{\gamma(r)}\mathsf{X}$. We then have a notion of (time-varying) geodesic distance on X defined as

$$d_k(x,\tilde{x}) = \inf_{\gamma\in\Gamma(x,\tilde{x})} L_k(\gamma), \ \forall x, \tilde{x} \in \mathsf{X}. \tag{10}$$

Moreover, if the norms $|\cdot|_{[x,k]}$ are in fact independent of x, thus denoted $|\cdot|_{[k]}$, and if X is a convex set in $\mathbb{R}^n$ (possibly equal to $\mathbb{R}^n$), then the infimum in (10) is achieved by straight lines $\gamma(r) = x + r(\tilde{x} - x)$ and $d_k(x,\tilde{x}) = |\tilde{x} - x|_{[k]}$ in (10). Finally, each function f_k in (9) is associated to a Jacobian $F_k(x) := \frac{\partial f_k}{\partial x}(x)$, which defines a linear map from $\mathsf{T}_x\mathsf{X}$ at time k to $\mathsf{T}_{f_k(x)}X$ at time $k+1$. As a result, for all vectors $v \in \mathsf{T}_x\mathsf{X}$,

$$|F_k(x)v|_{[f_k(x),k+1]} \leq \|F_k(x)\|^{[x.k]}_{[f_k(x),k+1]} |v|_{[x,k]}, \tag{11}$$

where $\|\cdot\|^{[x.k]}_{[f(x),k+1]}$ denotes the norm induced by $|\cdot|_{[x,k]}$ and $|\cdot|_{[f_k(x),k+1]}$.

Remark 4. The discussion could be carried out in a slightly more general framework by allowing asymmetric norms on the tangent spaces[20] rather than standard norms, but we will not need this level of generality.

Definition 3. Let ρ be a nonnegative constant. System (9) is said to be *ρ-contracting* for the norms $|\cdot|_{[x,k]}$ on a forward invariant set $C \subset \mathsf{X}$ if, for any $k_0 \in \mathbb{N}$ and any two initial conditions $x_0, \tilde{x}_0 \in C$, we have, for all $k \geq k_0$,

$$d_k(\phi(k;k_0,x_0), \phi(k;k_0,\tilde{x}_0)) \leq \rho^{k-k_0} d_{k_0}(x_0,\tilde{x}_0). \tag{12}$$

Let $\gamma_k \in \Gamma(x,\tilde{x})$ be a curve joining two points x and $\tilde{x}$ in X at a fixed time k. Let $\gamma'_k(r)$ be the tangent vector to γ_k at the point $\gamma_k(r)$, for $r \in [0,1]$. The curve γ_k is transported at time k by (9) to a curve γ_{k+1} joining $f_k(x)$ and $f_k(\tilde{x})$. Taking the derivative with respect to r in the equation $\gamma_{k+1}(r) = f_k(\gamma_k(r))$, we obtain the important *linear* relation between tangent vectors

$$\gamma'_{k+1}(r) = F_k(\gamma_k(r))\gamma'_k(r), \ \forall r \in [0,1], \forall k \geq 0. \tag{13}$$

The following fundamental theorem of contraction analysis is then a consequence of (13).

Theorem 2. *Let $F_k = \frac{\partial f_k}{\partial x}$ be the Jacobian of f_k, for all $k \geq 0$. A sufficient condition for system (9) to be ρ-contracting for the norms $|\cdot|_{[x,k]}$ on a forward invariant set $C \subset \mathsf{X}$ is that*

$$\|F_k(x)\|^{[x.k]}_{[f_k(x),k+1]} \leq \rho, \quad \forall x \in C, \forall k \in \mathbb{N}. \tag{14}$$

Proof. Consider a curve $\gamma_{k_0} : [0,1] \to \mathsf{X}$ in $\Gamma(x_0,\tilde{x}_0)$. This curve is transported by (9) to a sequence of curves $\gamma_{k_0+1}, \gamma_{k_0+2}, \ldots$, ie, $\gamma_{k+1}(r) = f_k(\gamma_k(r))$, for all $r \in [0,1]$, with γ_k joining $\phi(k;k_0,x_0)$ and $\phi(k;k_0,\tilde{x}_0)$. We have, for all $k \geq k_0$, using (13)

$$L_{k+1}(\gamma_{k+1}) = \int_0^1 \left|\gamma'_{k+1}(r)\right|_{[\gamma_{k+1}(r),k+1]} dr = \int_0^1 \left|F_k(\gamma_k(r))\gamma'_k(r)\right|_{[\gamma_{k+1}(r),k+1]} dr.$$

Now, using (11) and then assumption (14)

$$L_{k+1}(\gamma_{k+1}) \leq \int_0^1 \|F_k(\gamma_k(r))\|_{[\gamma_{k+1}(r),k+1]}^{[\gamma_k(r),k]} \left|\gamma_k'(r)\right|_{[\gamma_k(r),k]} dr \leq \rho \int_0^1 \left|\gamma_k'(r)\right|_{[\gamma_k(r),k]} dr = \rho L_k(\gamma_k), \tag{15}$$

and hence by immediate recursion, $L_k(\gamma_k) \leq \rho^{k-k_0} L_{k_0}(\gamma_{k_0})$. To conclude, let $\epsilon > 0$ and take the curve γ_{k_0} above to satisfy

$$L_{k_0}(\gamma_{k_0}) \leq (1+\epsilon) d_{k_0}(x_0, \tilde{x}_0).$$

Then, since $\gamma_k \in \Gamma(\phi(k; k_0, x_0), \phi(k; k_0, \tilde{x}_0))$, we have

$$d_k(\phi(k; k_0, x_0), \phi(k; k_0, \tilde{x}_0)) \leq L_k(\gamma_k) \leq \rho^{k-k_0} L_{k_0}(\gamma_{k_0}) \leq (1+\epsilon)\rho^{k-k_0} d_{k_0}(x_0, \tilde{x}_0). \tag{16}$$

Since this inequality is true for all $\epsilon > 0$, (12) holds. □

Remark 5. Note that to obtain useful results in continuous time (in particular, to detect convergent dynamics), it is crucial to use a tighter inequality replacing the first inequality of (15) by Coppel's inequality[36] to bound the solutions of linear differential equations. This leads to a sufficient condition for continuous-time systems similar to (14) stated in terms of matrix measures instead of induced norms.[18,19,33] However, this does not apply to discrete-time systems.

Corollary 1. *With the notation defined as in Theorem 2, suppose that C is a convex forward invariant subset of $\mathbb{R}^n$ and that the norms $|\cdot|_{[x,k]}$ on the tangent spaces are independent of x and denoted $|\cdot|_k$. Let $\|\cdot\|_{k+1}^k$ be the matrix norm induced by $|\cdot|_k$ and $|\cdot|_{k+1}$. Then, if $\|F_k(x)\|_{k+1}^k \leq \rho$ for all $x \in C$ and for all $k \in \mathbb{N}$, we have*

$$|\phi(k; k_0, x_0) - \phi(k; k_0, \tilde{x}_0)|_k \leq \rho^{k-k_0} |x_0 - \tilde{x}_0|_{k_0}, \quad \forall x_0, \tilde{x}_0 \in C, \forall k \geq k_0.$$

Proof. The result follows immediately from Theorem 2 and the remarks on geodesic distances preceding Definition 3. □

Corollary 2. *With the notation defined as in Theorem 2, suppose that the norms on the tangent spaces are defined for all x and k by $|v|_{[x,k]} = |P_{[x,k]} v|_1$, where $P_{[x,k]} = diag(p_{[x,k]})$, with $p_{[x,k]}$ a vector with positive components $p_{[x,k],i}$. Hence, $|v|_{[x,k]} = \sum_{i=1}^n p_{[x,k],i} |v_i|$. Then, the system is ρ-contracting for the associated distances on X if the following linear programs are feasible, for all $x \in C$ and $k \in \mathbb{N}$*

$$\sum_{i=1}^n p_{[f_k(x),k+1],i} |F_{k,ij}(x)| \leq \rho p_{[x,k],j}, \quad \forall 1 \leq j \leq n, \tag{17}$$

$$p_{[x,k],i}, p_{[f_k(x),k+1],i} > 0, \ \forall 1 \leq i \leq n. \tag{18}$$

In particular, if C is convex and if there exist positive vectors $p_{[k]}$ independent of x satisfying the above inequalities (17), (18) for all x, k, then, with $P_{[k]} := diag(p_{[k]})$, $x_k := \phi(k; k_0, x_0)$, $\tilde{x}_k := \phi(k; k_0, \tilde{x}_0)$, we have

$$\left|P_{[k]}(x_k - \tilde{x}_k)\right|_1 \leq \rho^{k-k_0} \left|P_{[k_0]}(x_0 - \tilde{x}_0)\right|_1, \quad \forall x_0, \tilde{x}_0 \in C, \forall k \geq k_0. \tag{19}$$

Proof. Inequalities (17), (18) come from satisfying (14) for the 1-norm weighted by $R := P_{[x,k]}$ and $S := P_{[f_k(x),k+1]}$. Condition (14) is equivalent to the induced 1-norm of the matrix $SF_k(x)R^{-1}$ being less than ρ, and this matrix has entries $p_{[f_k(x),k+1],i} F_{k,ij}(x) / p_{[x,k],j}$. The induced 1-norm of an $n \times m$ matrix $A = [a_{ij}]_{i,j}$ is $\max_{1 \leq j \leq m} \sum_{i=1}^n |a_{ij}|$. The result follows from these facts. □

Corollary 3. *With the notation defined as in Theorem 2, suppose that the norms on the tangent spaces are defined by $|v|_{[x,k]} = (v^T P_{[x,k]} v)^{1/2} = |P_{[x,k]}^{1/2} v|_2$, where $P_{[x,k]} \succ 0$, for all x and k. Then, the system is ρ-contracting for the associated distances on X if the following linear matrix inequalities (LMIs) are satisfied*

$$F_k(x)^T P_{[f(x),k+1]} F_k(x) \leq \rho^2 P_{[x,k]}, \ \forall x \in C, \forall k \in \mathbb{N}. \tag{20}$$

Suppose C is convex. If there exist matrices $P_{[k]} \succ 0$, $k \in \mathbb{N}$, independent of x, satisfying these LMIs, then we have

$$\left|P_{[k]}^{1/2}(x_k - \tilde{x}_k)\right|_2 \leq \rho^{k-k_0} \left|P_{[k_0]}^{1/2}(x_0 - \tilde{x}_0)\right|_2, \quad \forall x_0, \tilde{x}_0 \in C, \forall k \geq k_0, \tag{21}$$

where $x_k := \phi(k;k_0,x_0)$, $\tilde{x}_k := \phi(k;k_0,\tilde{x}_0)$. *If there exist matrices* $P_{[x,k]}$ *satisfying (20) and if there exist 2 matrices* $P_{\min} \succ 0$ *with minimum eigenvalue* $\lambda_{\min} > 0$ *and* $P_{\max} \succ 0$ *with maximum eigenvalue* $\lambda_{\max} > 0$ *such that we have* $\lambda_{min} I \preceq P_{\min} \preceq P_{[x,k]} \preceq P_{\max} \preceq \lambda_{\max} I$, *for all* x, k, *then*

$$\left|P_{\min}^{1/2}(x_k-\tilde{x}_k)\right|_2 \leq \rho^{k-k_0}\left|P_{\max}^{1/2}(x_0-\tilde{x}_0)\right|_2, \quad \forall x_0,\tilde{x}_0 \in C, \forall k \geq k_0,$$

and hence,

$$|x_k-\tilde{x}_k|_2 \leq \rho^{k-k_0}\sqrt{\frac{\lambda_{\max}}{\lambda_{\min}}}|x_0-\tilde{x}_0|_2.$$

Proof. This is a corollary of Theorem 2 since satisfying (14) for the norm induced by the weighted 2-norms with matrices $P_{[x,k]}$ and $P_{[f(x),k+1]}$ can be written $v^T F_k(x)^T P_{[f(x),k+1]} F_k(x) v \leq \rho^2 v^T P_{[x,k]} v$, for all v in $\mathbb{R}^n$. The second part uses the fact

$$\int_0^1 \sqrt{\gamma'(r)P_{\max}\gamma'(r)}dr \geq L_k(\gamma) = \int_0^1 \sqrt{\gamma'(r)P_{[\gamma(r),k]}\gamma'(r)}dr \geq \int_0^1 \sqrt{\gamma'(r)P_{\min}\gamma'(r)}dr,$$

and moreover, $\int_0^1 \sqrt{\gamma'(r)P_{\min}\gamma'(r)}dr \geq |P_{min}^{1/2}(x-\tilde{x})|_2$ if $\gamma \in \Gamma(x,\tilde{x})$ since, for a constant norm on $\mathbb{R}^n$, the geodesic curves are straight lines. Finally, referring to the argument leading to (16), we get

$$\left|P_{min}^{1/2}(x_k-\tilde{x}_k)\right|_2 \leq L_k(\gamma_k) \leq \rho^{k-k_0}L_{k_0}(\gamma_{k_0}) \leq \rho^{k-k_0}\left|P_{\max}^{1/2}(x_0-\tilde{x}_0)\right|_2.$$

□

Remark 6. The first part of Corollary 3 is the classical contraction result,[18] in discrete time, for norms associated with an inner product (Riemannian structure on X). Using state-dependent P matrices enlarges the set of systems for which we can prove contraction, but in our case, we also need to explicitly bound the Euclidean distances $|x_k-\tilde{x}_k|_2$, not just general geodesic distances, to be able to evaluate the level of noise necessary for the Gaussian mechanism of Theorem 1.

3.2 | Effect of disturbances

For the computation of ℓ^1 and ℓ^2-sensitivities, we need to study the trajectory deviations of contracting systems subject to disturbances. Qualitatively, the exponential convergence of trajectories of a contracting system provides some robustness against disturbances.[18,19,28] However, to precisely set the level of privacy-preserving noise, quantitative worst-case bounds on the ℓ^1 or ℓ^2-norms of the trajectory deviations are needed. Hence, consider a system

$$x_{k+1} = f_k(x_k, \pi_k(x_k)), \tag{22}$$

where $\pi_k : \mathsf{X} \to \mathsf{P} := \mathbb{R}^p$, for some p, represents a C^1 disturbance signal, and for all $k \geq 0$, $f_k : \mathsf{X}\times\mathsf{P} \to \mathsf{X}$ is C^1. We equip the tangent spaces of the product manifold $\mathsf{X}\times\mathsf{P}$ with time-varying norms assumed for simplicity to be fixed for the disturbance part, ie, $|(v,w)|_{[(x,\pi),k]} = |v|_{[x,k]} + |w|_{\mathsf{P}}$, for a fixed norm $|\cdot|_{\mathsf{P}}$. The nominal system under zero disturbance is

$$\bar{x}_{k+1} = f_k(\bar{x}_k, 0). \tag{23}$$

We denote $\frac{\partial f_k}{\partial x}$ and $\frac{\partial f_k}{\partial \pi}$ the Jacobian matrices of $f_k(x,\pi)$ with respect to the components of x and π, respectively.
For $r \in [0,1]$, denote by $\phi(k;r,k_0,x_0)$ the iterates of

$$x_{k+1} = f_k(x_k, r\,\pi_k(x_k)), \tag{24}$$

starting from x_0 at time k_0. Note that (22) corresponds to $r = 1$ and (23) to $r = 0$. Let us also define

$$J_k(x;r) := \frac{\partial f_k}{\partial x}(x, r\,\pi_k(x)) + r\frac{\partial f_k}{\partial \pi}(x, r\,\pi_k(x))\frac{\partial \pi_k}{\partial x}(x), \ \forall x \in \mathsf{X}, \forall r \in [0,1]. \tag{25}$$

For all x in X, denote $x_+^{k,r} := f_k(x, r\,\pi_k(x))$. Formally, the "differential" maps (25) are from $\mathsf{T}_{[x,k]}\mathsf{X}$ to $\mathsf{T}_{[x_+^{k,r},k+1]}\mathsf{X}$, with the corresponding induced norms $\|\cdot\|_{[x_+^{k,r},k+1]}^{[x,k]}$. We then have the following result.

Theorem 3. *Consider a trajectory $\bar{x}_k := \phi(k;0,k_0,\bar{x}_0)$ for (23) starting from $\bar{x}_0$ and a trajectory $x_k := \phi(k;1,k_0,x_0)$ for the perturbed system (22) starting from x_0. Suppose that there exists a sequence $\{M_k\}_{k\geq 0}$ such that*

$$\left|\frac{\partial f_k}{\partial \pi}(x, r\,\pi_k(x))\,\pi_k(x)\right|_{\left[x_+^{k,r},k+1\right]} \leq M_k, \quad \forall r\in[0,1], \forall x\in C, \forall k\geq k_0, \tag{26}$$

and that

$$\|J_k(x;r)\|_{\left[x_+^{k,r},k+1\right]}^{[x,k]} \leq \rho, \quad \forall r\in[0,1], \forall x\in C, \forall k\geq k_0, \tag{27}$$

where C is a forward invariant set for (24), for all $r\in[0,1]$. Then, we have, for all $k\geq k_0$ and the distances d_k defined in (10),

$$d_k(\bar{x}_k,x_k) \leq \rho^{k-k_0} d_{k_0}(\bar{x}_{k_0},x_{k_0}) + \sum_{l=0}^{k-k_0-1} \rho^l M_{k-1-l}.$$

Remark 7. As an example, in the case of additive disturbances on $\mathsf{X}=\mathsf{P}=\mathbb{R}^n$, ie,

$$f_k(x,\pi_k(x)) = \tilde{f}_k(x) + \pi_k(x), \tag{28}$$

with a fixed norm $|\cdot|$ on $\mathbb{R}^n$, condition (26) can be written more simply $\sup_{x\in\mathsf{C}}|\pi_k(x)|\leq M_k$, ie, M_k is a bound on the disturbance term.

Remark 8. Note that if the disturbance π_k does not depend on x, then (25) reads $J_k(x;r) := \frac{\partial f_k}{\partial x}(x, r\,\pi_k)$ and (27) is a type of contraction condition on the perturbed system. If moreover the perturbation is in fact additive as in (28), then (27) simply asks that the Jacobian of the nominal system $\tilde{f}_k$ satisfy the contraction assumption.

Proof. Consider a curve $\gamma_{k_0}\in\Gamma(\bar{x}_0,x_0)$, ie, such that $\gamma_{k_0}(0)=\bar{x}_0$ and $\gamma_{k_0}(1)=x_0$, transported by (24) to the sequence

$$\gamma_k(r) = \phi(k;r,k_0,\gamma_{k_0}(r)), \ \forall r\in[0,1], \forall k\geq k_0.$$

Then, for $k\geq k_0$, we have $\gamma_k\in\Gamma(\bar{x}_k,x_k)$, where $\bar{x}_k := \phi(k;0,k_0,\bar{x}_0)$ and $x_k := \phi(k;1,k_0,x_0)$. Following the idea of the proof of Theorem 2, define $\gamma_k'(r) := \frac{d}{dr}\phi(k;r,k_0,\gamma_{k_0}(r))$, so that we have, for all k and all $r\in[0,1]$,

$$\gamma_{k+1}'(r) = J_k(\gamma_k(r);r)\,\gamma_k'(r) + \frac{\partial f_k}{\partial \pi}(\gamma_k(r), r\,\pi_k(\gamma_k(r)))\,\pi_k(\gamma_k(r)),$$

which implies, by (27) and (26),

$$\left|\gamma_{k+1}'(r)\right|_{[\gamma_{k+1}(r),k+1]} \leq \rho\left|\gamma_k'(r)\right|_{[\gamma_k(r),k]} + M_k, \quad \forall r\in[0,1], \forall k\geq k_0,$$

and by integration over $r\in[0,1]$

$$L_{k+1}(\gamma_{k+1}) \leq \rho L_k(\gamma_k) + M_k, \ \forall k\geq k_0.$$

By the comparison lemma,[37] we then have that $L(\gamma_k)\leq u_k$ for u_k satisfying the linear scalar dynamics

$$u_{k_0} = L_{k_0}(\gamma_{k_0}), \ u_{k+1} = \rho\,u_k + M_k, \forall k\geq k_0.$$

Hence, $L_k(\gamma_k)\leq \rho^{k-k_0}u_{k_0} + \sum_{l=0}^{k-k_0-1}\rho^l M_{k-1-l}$. As in the end of the proof of Theorem 2, we can then choose γ_{k_0} so that $L_{k_0}(\gamma_{k_0})$ is arbitrarily close to $d_{k_0}(\bar{x}_0,x_0)$, and then use $d_k(\bar{x}_k,x_k)\leq L_k(\gamma_k)$ to conclude. □

We can now make convergence assumptions on the bounding sequence $\{M_k\}_{k\geq 0}$ in (26) to state more concrete results. The following corollaries follow by standard calculations[28] on the sequence u_k introduced at the end of the proof of Theorem 3.

Corollary 4. *Let $1\leq p\leq\infty$ be an integer. Suppose that $\{M_k\}_{k\geq 0}$ in (26) is a sequence in ℓ^p, with norm $\|M\|_p$. Then, with the notation and assumptions of Theorem (3), if $\rho<1$, there exists a class $\mathcal{K}$ function $\beta:\mathbb{R}_+\to\mathbb{R}_+$ such that*

$$\left(\sum_{k=k_0}^{\infty} d_k(\bar{x}_k,x_k)^p\right)^{1/p} \leq \beta(d_{k_0}(x_0,\bar{x}_0)) + \frac{\|M\|_p}{1-\rho}, \tag{29}$$

where, for $p=\infty$, the left-hand side of the inequality is interpreted as usual as $\sup_{k\geq k_0} d_k(\bar{x}_k,x_k)$.

By further restricting the class of disturbances, we get slightly tighter bounds on the deviations for $p\geq 2$.

Corollary 5. *Let $1 \leq p \leq \infty$ be an integer. Suppose that $\{M_k\}_{k\geq 0}$ in (26) satisfies the following condition:*

$$\exists K \geq 0, 1 > \alpha \geq 0, and \;\; k_0 \in \mathbb{N} s.t. M_k = \begin{cases} 0, & if\, k < k_0, \\ K\alpha^{k-k_0}, & if\, k \geq k_0. \end{cases} \tag{30}$$

Then, with the notation and assumptions of Theorem 3, for $k \geq k_0$,

$$d_k(\bar{x}_k, x_k) \leq \rho^{k-k_0} d_{k_0}(\bar{x}_0, x_0) + K\frac{\rho^{k-k_0} - \alpha^{k-k_0}}{\rho - \alpha}.$$

Hence, if $\rho < 1$,

$$\sum_{k=k_0}^{\infty} d_k(\bar{x}_k, x_k) \leq \frac{1}{1-\rho} d_{k_0}(\bar{x}_0, x_0) + \frac{K}{(1-\rho)(1-\alpha)},$$

and for any $p \geq 2$, there exists a class $\mathcal{K}$ function $\beta : \mathbb{R}_+ \to \mathbb{R}_+$ such that

$$\left(\sum_{k=k_0}^{\infty} d_k(\bar{x}_k, x_k)^p\right)^{1/p} \leq \beta(d_{k_0}(\bar{x}_0, x_0)) + \frac{K}{|\rho - \alpha|}\left(\sum_{k=0}^{\infty} |\rho^k - \alpha^k|^p\right)^{1/p}. \tag{31}$$

Remark 9. If the norms on TX are given by weighted 1- and 2-norms as in Corollaries 2 and 3, then condition (27) corresponds to the feasibility of a family of linear programs or LMIs, and if moreover C is convex and the weight matrices in these norms are independent of x, then we can replace the distances $d_k(\bar{x}_k, x_k)$ in (29), (31) by $|P_{[k]}(\bar{x}_k - x_k)|_1$ or $|P_{[k]}^{1/2}(\bar{x}_k - x_k)|_2$ as in (19), (21).

4 | DIFFERENTIALLY PRIVATE OBSERVERS WITH OUTPUT PERTURBATION

Let us now return to our initial differentially private observer design problem with output perturbation. Two adjacent measured signals y and $\tilde{y}$ produce distinct observer state trajectories z and $\tilde{z}$ by (3), such that

$$z_{k+1} = f_k(z_k) + h_k(z_k, y_k - g_k(z_k)), \tag{32}$$

$$\tilde{z}_{k+1} = f_k(\tilde{z}_k) + h_k(\tilde{z}_k, y_k - g_k(\tilde{z}_k) + \pi_k), \tag{33}$$

where $\pi_k = \tilde{y}_k - y_k$. We can now attempt to choose the functions h_k to design a contractive observer while at the same time minimizing the "gain" of the map $\pi \to z$. First, contraction provides a notion of convergence for the observer. Namely, if model (1), (2) were valid under no modeling noise assumptions (zero v, w), then any the sequence x satisfying (1), (2) would also satisfy the dynamics (32) (since $y_k = g(x_k)$), and the trajectories x, z would converge exponentially toward each other, so that any initial difference between z_0 and x_0 would eventually be forgotten. Second, the results of Section 3.2 give us tools to bound the sensitivity of contractive observers, ie, the deviations between z and $\tilde{z}$ above and, hence, a mean to set the level of privacy-preserving noise using Theorem 1.

Given two measured signals y and $\tilde{y}$, define the notation $v_k^{y,\tilde{y}}(x; r) := y_k - g_k(x) + r\pi_k = (1-r)y_k + r\tilde{y}_k - g_k(x)$ and

$$J_k^{y,\tilde{y}}(x; r) = \frac{\partial f_k}{\partial x}(x) + \frac{\partial h_k}{\partial x}\left(x, v_k^{y,\tilde{y}}(x; r)\right) - \frac{\partial h_k}{\partial y}\left(x, v_k^{y,\tilde{y}}(x; r)\right)\frac{\partial g_k}{\partial x}(x). \tag{34}$$

The proof of the following proposition follows immediately from Theorem 3 and Remark 8.

Proposition 1. *Consider observer (3), and two measured signals y, $\tilde{y}$ producing respectively the trajectories z, $\tilde{z}$, assuming the same initial condition $z_0 = \tilde{z}_0$ to initialize the observer. Suppose that we have the bound*

$$\left\| J_k^{y,\tilde{y}}(x; r) \right\|_{\left[x_+^{k,r}, k+1\right]}^{[x,k]} \leq \rho, \quad \forall r \in [0,1], \forall x \in C, \forall k \in \mathbb{N}, \tag{35}$$

where $J_k^{y,\tilde{y}}$ *is defined by (34),* $x_+^{k,r} := f_k(x) + h_k(x, \nu_k^{y,\tilde{y}}(x;r))$ *and* C *is a set containing* z_0*, which is forward invariant for observer (32) for any input signal* $(1-r)y + r\tilde{y}$*,* $r \in [0,1]$*. Suppose moreover that*

$$\sup_{x\in C, r\in[0,1]} \left| \frac{\partial h_k}{\partial y}\left(x, \nu_k^{y,\tilde{y}}(x;r)\right)(\tilde{y}_k - y_k)\right|_{\left[x_+^{k,r}, k+1\right]} \leq M_k, \quad \forall k \in \mathbb{N}. \tag{36}$$

Then, we have, for the distances d_k *associated to the norms* $|\cdot|_{[x,k]}$*,*

$$d_k(z_k, \tilde{z}_k) \leq \sum_{l=0}^{k-1} \rho^l M_{k-1-l}.$$

The result of Proposition 1 is still quite general. To illustrate how it can be applied and to simplify the following discussion, we now focus on the simpler situation where a Luenberger-type observer can be used to estimate the state[38]

$$z_{k+1} = f_k(z_k) + H_k \times (y_k - g_k(z_k)), \tag{37}$$

where H_k represents an $n \times m$ matrix to design. In other words, we set $h_k(x,y) = H_k y$. Then, expression (34) reads simply $\frac{\partial f_k}{\partial x}(x) - H_k \frac{\partial g_k}{\partial x}(x)$ and becomes, in particular, independent of r and $y, \tilde{y}$. Next, fix a norm $|\cdot|_{\mathsf{X}}$ on TX, independent of x, k, and a p-norm $|\cdot|_p$ on Y, and let $\bar{H}_{\mathsf{X}}^p := \sup_k \|H_k\|_{\mathsf{X}}^{\mathsf{Y}}$. Then, in (36), we can take $M_k = \bar{H}_{\mathsf{X}}^p |y_k - \tilde{y}_k|_p$. This leads to the following corollary of Proposition 1 similar to the Corollaries 4 and 5, which we will use next in the illustrative examples. We introduce the notation $\|v\|_{p,\mathsf{X}} := (\sum_{k=0}^{\infty} |v_k|_{\mathsf{X}}^p)^{1/p}$, for $1 \leq p \leq \infty$.

Corollary 6. *Consider observer (37), and two measured signals* $y, \tilde{y}$ *producing respectively the trajectories* $z, \tilde{z}$*, assuming the same initial condition* $z_0 = \tilde{z}_0$ *to initialize the observer. Fix the norms* $|\cdot|_{\mathsf{X}}$*, on* TX*, independent of* x, k*. Suppose that we have the bound*

$$\left\| \frac{\partial f_k}{\partial x}(x) - H_k \frac{\partial g_k}{\partial x}(x) \right\|_{\mathsf{X}} \leq \rho, \quad \forall x \in C, k \in \mathbb{N} \tag{38}$$

for some constant $\rho < 1$*, where* C *is a set containing* z_0 *and forward invariant for (32) for any input signal* $y + (1-r)\tilde{y}$*,* $r \in [0,1]$*. Then, if the signals* $y, \tilde{y}$ *are adjacent according to (5), we have, for the same value of* p*,*

$$\|z - \tilde{z}\|_{p,\mathsf{X}} \leq \frac{B_p \bar{H}_{\mathsf{X}}^p}{1-\rho}. \tag{39}$$

Moreover, if the signals $y, \tilde{y}$ *are in fact adjacent according to (6), we have more precisely, for the same value of* p*,*

$$\|z - \tilde{z}\|_{p,\mathsf{X}} \leq \frac{K \bar{H}_{\mathsf{X}}^p}{|\rho - \alpha|} \left(\sum_{k=0}^{\infty} |\rho^k - \alpha^k|^p \right)^{1/p}. \tag{40}$$

Remark 10. For the adjacency relation (6) with $p = 1$, both (40) and (39) give the same upper bound $\frac{K \bar{H}_{\mathsf{X}}^p}{(1-\rho)(1-\alpha)}$.

In Corollary 6, the choice of H_k has an impact both on ρ and on the ℓ^p-sensitivity bound. Increasing the gains H_k can help decrease the contraction rate ρ to obtain a more rapidly converging observer, but at the same time, it increases the sensitivity, in the sense of Section 2.3, and thus the level of noise necessary for differential privacy. Hence, in general, we should try to achieve a reasonable contraction rate ρ with the smallest gain possible. We conclude this section with two more corollaries, describing differentially private observers with output perturbation.

Corollary 7. *Let* $P = \mathrm{diag}(p)$*, with* $p_i > 0, 1 \leq i \leq n$*, and assume that the conditions of Corollary 6 are satisfied for the weighted 1-norm* $|Pv|_1 = \sum_{i=1}^n p_i |v_i|$ *on* X*. Consider the signal* $\hat{x}_k = z_k + \xi_k$*, where* z_k *is computed from (37), and* $\xi_{k,i}$ *are iid Laplace random variables with parameters* b/p_i*, for* $1 \leq i \leq n$*, where*

$$b = \frac{B_1 \sup_k \|PH_k\|_1}{\epsilon(1-\rho)}. \tag{41}$$

Then, this signal $\hat{x}_k$ *is* ϵ*-differentially private for the adjacency relation (5) with* $p = 1$ *and for (6) with* $p = 1$ *when* $B_1 = \frac{K}{1-\alpha}$.

Proof. From bound (39) for $p = 1$, since $\|z - z\|_{1,\mathsf{X}} = \sum_{k=0}^{\infty} |P(z_k - \tilde{z}_k)|_1$, we deduce by Theorem 1 that $Pz_k + \zeta_k$ is a differentially private signal, where ζ_k has Laplace distributed iid components with the parameter b. Hence, $P^{-1}(Pz_k + \zeta_k)$ is also differentially private (by resilience to postprocessing[15]) and we define $\xi_k = P^{-1}\zeta_k$ in the corollary. □

Corollary 8. *Let P be a positive definite matrix, and assume that the conditions of Corollary 6 are satisfied for the weighted 2-norm $|P^{1/2}v|_2$ on X. Consider the signal $\hat{x}_k = z_k + \xi_k$, where z_k is computed from (37), and ξ_k is a Gaussian white noise with covariance matrix $\sigma^2 P^{-1}$, where $\sigma = \kappa_{\delta,\epsilon} K_2 \sup_k \|P^{1/2}H_k\|_2$. Then, this signal $\hat{x}_k$ is (ϵ, δ)-differentially private for the adjacency relation (5) with $p = 2$ if $K_2 = B_2/(1-\rho)$, and for the adjacency relation (6) with $p = 2$ if $K_2 = \frac{K}{|\rho-\alpha|}(\sum_{k\geq 0}(\rho^k - \alpha^k)^2)^{1/2}$.*

Proof. From bounds (39) or (40), we deduce by Theorem 1 that $P^{1/2}z_k + \zeta_k$ is a differentially private signal, where ζ_k is a Gaussian white noise with covariance matrix $\sigma^2 I$. Hence, $P^{-1/2}(P^{1/2}z_k + \zeta_k)$ is also differentially private (by resilience to postprocessing[15]) and we define $\xi_k = P^{-1/2}\zeta_k$ in the corollary. □

Corollaries 7 and 8 give two differentially private mechanisms with output perturbation, provided that we can design the matrices $\bar{H}_k$ to verify the assumptions of Corollary 6 with the (weighted) 1- or 2-norm on X. The next section discusses application examples for the privacy-preserving observer design methodology.

5 | APPLICATION EXAMPLES

5.1 | Estimating link formation preferences in dynamic social networks

Statistical studies of networks have intensified tremendously in recent years, with one motivating application being the emergence of online social networking communities. In this section, we focus on a recently proposed state-space model[39] describing the dynamics of link formation in networks, called the dynamic stochastic block model. It combines a linear state-space model for the underlying dynamics of the network and the classical stochastic block model of Holland et al,[22] resulting in a nonlinear measurement equation. Examples of applications of this model include mining email and cell phone databases,[39] which obviously contain privacy-sensitive data.

Consider a set of n nodes. Each node corresponds to an individual and can belong to one of N classes. Let θ_k^{ab} be the probability of forming an edge at time k between a node in class a and a node in class b, and let θ_k denote the vector of probabilities $[\theta_k^{ab}]_{1\leq a,b\leq N}$. For example, edges could represent email exchanges or phone conversations. Edges are assumed to be formed independently of each other according to θ_k. Let $y_k^{ab} = \frac{m_k^{ab}}{n^{ab}}$ be the observed density of edges between classes a and b, where m_k^{ab} is the number of observed edges between classes a and b at time k, and n^{ab} is the maximum possible number of edges between these two classes. For simplicity, we assume that the quantities n^{ab} are publicly known (this is the case, for example, if the class of each node is public information), and we focus on the problem of estimating the parameters θ_k^{ab} by using the signals y_k^{ab}. This corresponds to the "a priori" block modeling setting.[22,39] The links formed between specific nodes constitute private information however, so directly releasing m_k^{ab} or y_k^{ab} or an estimate of θ_k based on these quantities is not allowed.

If n^{ab} is large enough, previous work has argued[39] using the central limit theorem that an approximate model where y_k^{ab} is Gaussian is justified, so that

$$y_k = \theta_k + v_k, \tag{42}$$

where v_k is a Gaussian noise vector with diagonal covariance matrix V_k (whose entries theoretically should depend on θ_k, but this aspect is neglected in the model). Rather than defining a dynamic model for θ_k, whose entries are constrained to be between 0 and 1, let us redefine the state vector to be the so-called logit of θ_k, denoted ψ_k, with entries $\psi_k^{ab} = \ln \frac{\theta_k^{ab}}{1-\theta_k^{ab}}$, which are well defined for $0 < \theta_k^{ab} < 1$. The dynamics of ψ_k is assumed to be linear

$$\psi_{k+1} = F\psi_k + w_k, \tag{43}$$

for some known matrix F, and for noise vectors w_k assumed to be iid Gaussian with known covariance matrix W.[39] The observation model (42) now becomes

$$y_k = g(\psi_k) + v_k, \tag{44}$$

where the components of g are given by the logistic function applied to each entry of ψ, ie,

$$g^{ab}(\psi_k) = \frac{1}{\left(1 + e^{-\psi_k^{ab}}\right)}.$$

An extended Kalman filter is proposed in the work of Xu and Hero[39] to estimate ψ, but we pursue here a deterministic observer design to illustrate the ideas discussed in the previous sections. Hence, for simplicity, we consider an observer of the form

$$\hat{\psi}_{k+1} = F\hat{\psi}_k + H(y_k - g(\hat{\psi}_k)) = (F\hat{\psi}_k - Hg(\hat{\psi}_k)) + Hy_k,$$

with H a constant square gain matrix. To enforce contraction as in Corollary 6, we should choose H so that $\|F - HG(\psi)\| \leq \rho$, where $G(\psi)$ is the Jacobian of g at ψ. Note that $G(\psi)$ is a square and diagonal matrix with entries $G_{ii}(\psi) = \frac{e^{-\psi^i}}{(1+e^{-\psi^i})^2}$, with i indexing pairs (a, b). The only nonlinearity in model (43), (44) comes from the observation model (44).

To further simplify the following discussion, let us assume that F is also diagonal (an assumption also made in the previous work,[39] where the coupling between components occurs only through the nondiagonal covariance matrix W). In this case, the systems completely decouple into scalar systems, and it is natural to choose H to be diagonal as well. The observer for one of these scalar system takes the form

$$z_{k+1} = fz_k + h \times \left(y_k - \frac{1}{1 + e^{-z_k}}\right) = fz_k - \frac{h}{1 + e^{-z_k}} + hy_k, \tag{45}$$

where $h \in \mathbb{R}$ is the observer gain to set, $f \in \mathbb{R}_+$, $z_k \in \mathbb{R}$ is one component (a, b) of $\hat{\psi}_k$ and y_k now represents just the corresponding scalar component of the measurement vector as well. Since the state space X is now $\mathbb{R}$, the norm $|\cdot|_{\mathsf{X}}$ is simply the absolute value. The contraction condition (38) reads, for some $0 < \rho < 1$,

$$-\rho \leq f - \frac{he^{-z}}{(1 + e^{-z})^2} \leq \rho \tag{46}$$

$$\text{ie, } f - \rho \leq \frac{he^{-z}}{(1 + e^{-z})^2} \leq f + \rho. \tag{47}$$

Now, note that $0 \leq \frac{e^{-z}}{(1+e^{-z})^2} \leq \frac{1}{4}$ for all z. Hence, by taking $h \leq 4(f + \rho)$, the right inequality (47) is satisfied. To satisfy the left inequality, if $f < 1$, we could potentially take $\rho \geq f$, although the estimation performance might not necessarily be satisfying in this case. Alternatively, if $f \geq 1$ or if we want to achieve a smaller contraction parameter ρ than the value of f, we can enforce the left inequality on a subset of the state space. Namely, for $-a \leq z \leq a$, we have $\frac{e^{-z}}{(1+e^{-z})^2} \geq \frac{e^{-a}}{(1+e^{-a})^2}$. In this case, for $\rho < f$, by taking $h \geq (f - \rho)e^a(1 + e^{-a})^2$, the left-hand side of (47) is also satisfied.

Suppose for example that $f = 1$ in the dynamics (45), so that (43) describes a Gaussian random walk and that the adjacency relation considered is (6). By Corollary 7, we can publish an ϵ-differentially private estimate of ψ by computing z_k using (45) and adding Laplace noise to it with parameter $b = Kh/(\epsilon(1 - \rho)(1 - \alpha))$. Small noise requires small values of h and of ρ. Since we must take $\rho < 1$, we cannot enforce the left inequality of (47) for all values of z. Suppose then that we want to design a privacy-preserving observer assuming that θ remains in the interval $[0.1, 0.9]$ or equivalently $\psi \in [-2.197, 2.197]$ approximately. In this interval, we have $0.09 \leq \frac{e^{-\psi}}{(1+e^{-\psi})^2} \leq \frac{1}{4}$, and so ρ and h must also satisfy

$$\frac{f - \rho}{0.09} \leq h \leq 4(f + \rho), \quad \text{ie, } \frac{1 - \rho}{0.09} \leq h \leq 4(1 + \rho). \tag{48}$$

Note in particular that the factor $h/(1 - \rho)$ also appearing in parameter b is lower bounded by $1/0.09 \approx 11.1$. We should then set $h = (1 - \rho)/0.09$, satisfying the left inequality in (48) with equality, for the value of the contraction parameter ρ that we want to achieve. For example, for faster observer convergence we should try to achieve the lowest possible value of ρ, although this might amplify the steady-state variance due to measurement noise. The inequalities (48) can only be satisfied for $\rho \gtrapprox 0.47$, a contraction parameter that can then be achieved by taking $h \approx 5.88$.

Figure 2 illustrates the behavior of the privacy-preserving observer, when the privacy parameters are $\epsilon = \ln(3), \delta = 0$ and $K = 3 \times 10^{-3}$ and $\alpha = 0.25$ in (6). That is, for the pair of classes (a, b) under consideration, we want to provide a differential privacy guarantee making it hard to detect a transient variation in the number of edges, as long as this variation represents initially at most 0.3% of all the edges between classes a and b, and subsequently decreases at least geometrically with rate 1/4. Concretely, if edges represent phone conversations for example, this means that if an individual in class a suddenly increases his call volume with class b but by an amount representing less than a proportion K of all calls between

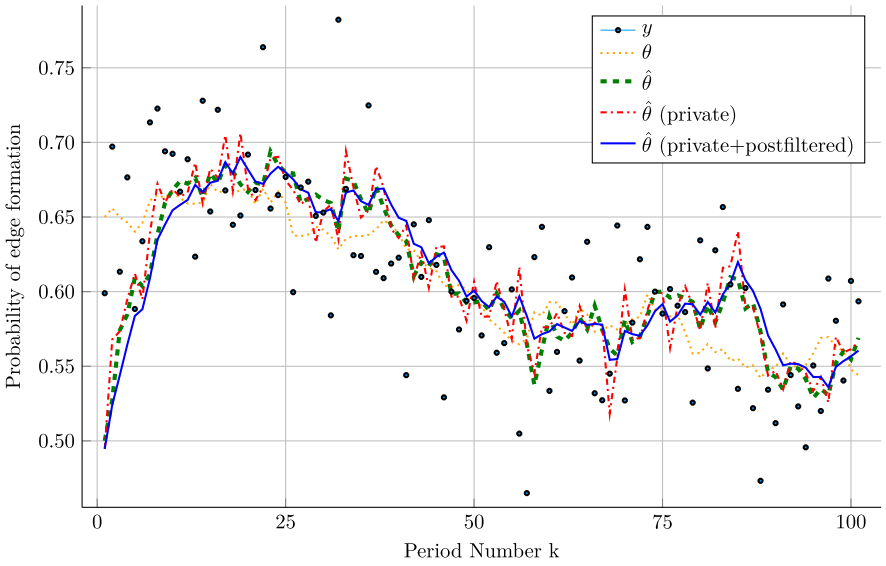

FIGURE 2 Sample path of the estimate of the edge formation probability θ_k^{ab}, for some classes (a, b). The measured edge density is generated from one component of model (42), (43) with $f = 1$ and w_k, v_k iid Gaussian random variables with zero mean and standard deviation 0.03 and 0.04, respectively. The trajectory θ (dotted line) starts at the value 0.65, and the observers are all initialized at the value 0.50. The upper bound ρ on the contraction rate of observer (45) is set to $\rho = 0.9$, providing a good trade-off between convergence speed and steady-state variance (see green dashed curve for the nonprivate observer), with corresponding gain $h = 1.11$. The dot-dashed line shows $1/(1+\exp(-\tilde{z}_k))$ as our private estimate of θ_k, where $\tilde{z}_k$ is a ln(3)-differentially private estimate of ψ_k (hence, the estimate of θ_k is also ln(3)-differentially private), obtained by the Laplace mechanism, for the adjacency relation (6) with parameter values $K = 3 \times 10^{-3}$, $\alpha = 0.25$. We also show a differentially private estimate obtained after further postfiltering, as explained in the main text [Colour figure can be viewed at wileyonlinelibrary.com]

a and b, and subsequently reduces this temporary activity at rate α, then an adversary having access to a differentially private estimate of θ_k^{ab} can only achieve a low probability of correctly detecting this event.[40]

As explained in Figure 1, it can be useful to further filter the differentially private signal produced above since this signal exposes directly the privacy-preserving noise. In this case, one can interpret the private estimate $\tilde{z}_k = z_k + \xi_k$, with ξ the Laplace noise as in Corollary (7), as a noisy measurement of ψ, now with a trivial linear measurement model, in contrast to (44). A possible simple postfilter smoothing $\tilde{z}_k$ can then be the linear observer

$$\hat{\psi}_{k+1} = f\hat{\psi}_k + k_{\text{post}}(\tilde{z}_k - f\hat{\psi}_k),$$

and Figure 2 also represents $\hat{\theta}_k = g(\hat{\psi}_k)$ for the gain value $k_{\text{post}} = 0.4$.

5.2 | Syndromic surveillance

Syndromic surveillance systems monitor health-related data in real-time in a population to facilitate early detection of epidemic outbreaks.[41] In particular, recent studies have shown the correlation between certain nonmedical data, eg, search engine queries related to a specific disease, and the proportion of individuals infected by this disease in the population.[42] Although time series analysis can be used to detect abnormal patterns in the collected data,[41] here, we focus on a model-based filtering approach[43] and develop a differentially private observer using a two-dimensional epidemiological model.

The following SIR model of Kermack and McKendrick[44,45] models the evolution of an epidemic in a population by dividing individuals into 3 categories: susceptible (S), ie, individuals who might become infected if exposed; infectious (I), ie, currently infected individuals who can transmit the infection; and recovered (R) individuals, who are immune to the infection. A simple version of the model in continuous-time includes bilinear terms and reads

$$\frac{ds}{dt} = -\mu\mathcal{R}_o is$$
$$\frac{di}{dt} = \mu\mathcal{R}_o is - \mu i.$$

Here, i and s represent the proportion of the total population in the classes I and S. The last class R need not be included in this model because we have the constraint $i + s + r = 1$. The parameter $\mathcal{R}_o$ is called the basic reproduction number and represents the average number of individuals infected by a sick person. The epidemic can propagate when $\mathcal{R}_o > 1$. The parameter μ represents the rate at which infectious people recover and move to class R. More details about this model can be found in the monograph of Brauer et al.[45]

Discretizing this model with sampling period τ, we get the discrete-time model

$$s_{k+1} = s_k - \tau\mu\mathcal{R}_o i_k s_k + w_{1,k} = f_1(s_k, i_k) + w_{1,k} \tag{49}$$

$$i_{k+1} = i_k + \tau\mu i_k(\mathcal{R}_o s_k - 1) + w_{2,k} = f_2(s_k, i_k) + w_{2,k}, \tag{50}$$

where we have also introduced noise signals w_1 and w_2 in the dynamics. We assume here for simplicity that we can collect syndromic data providing a noisy measurement of the proportion of infected individuals, ie,

$$y_k = i_k + v_k,$$

where v_k is a noise signal. We can then consider the design of an observer of the form

$$\hat{s}_{k+1} = f_1(\hat{s}_k, \hat{i}_k) + h_1(y_k - \hat{i}_k)$$
$$\hat{i}_{k+1} = f_2(\hat{s}_k, \hat{i}_k) + h_2(y_k - \hat{i}_k).$$

We define the Jacobian matrix of system (49), (50)

$$F(s,i) = I_2 + \tau\mu\mathcal{R}_o \begin{bmatrix} -i & -s \\ i & s - 1/\mathcal{R}_o \end{bmatrix},$$

as well as the gain matrix $H = [h_1, h_2]^T$ and observation matrix $C = [0, 1]$. Here, we design a differentially private observer with Gaussian noise using Corollary 8, for the adjacency relation (6) with $p = 2$.

Following Corollary 3, the contraction rate constraint (38) for a 2-norm on $\mathbb{R}^2$ weighted by a matrix $P \succ 0$ is equivalent to the family of inequalities, for all (s, i) in the region of $[0, 1]^2$ where we want to show contraction

$$(F(s,i) - HC)^T P (F(s,i) - HC) \preceq \rho^2 P$$
$$F_x^T P F_x - F_x^T P H C - C^T H^T P F_x + C^T H^T P H C \preceq \rho^2 P,$$

where we used $F_x := F(s, i)$ to simplify the notation. Defining the new variable $X = PH$, this can be rewritten as

$$F_x^T P F_x - F_x^T X C - C^T X^T F_x + C^T X^T P^{-1} X C \preceq \rho^2 P,$$

which, using the Schur complement, is equivalent to the family of LMIs

$$\begin{bmatrix} \rho^2 P - F_x^T P F_x + F_x^T X C + C^T X^T F_x & C^T X^T \\ XC & P \end{bmatrix} \succeq 0, \tag{51}$$

for all $x = (s, i)$ in the region where we want to prove contraction. If we can find P, X satisfying these inequalities, we recover the observer gain vector simply as $H = P^{-1}X$.

For a given value of ρ, the covariance matrix of the Gaussian noise in Corollary 8 is proportional to $\|P^{1/2}H\|_2^2 P^{-1} = (H^T P H)P^{-1} = (X^T P^{-1} X)P^{-1}$, and hence, it is advantageous to minimize a function of this matrix. Note that $X^T P^{-1} X$ is a scalar. Minimizing $(X^T P^{-1} X)\,\mathrm{Tr}(P^{-1})$ does not appear to directly lead to an efficiently solvable optimization problem, but as a proxy, we can choose to minimize instead the sum $X^T P^{-1} X + \nu \mathrm{Tr}(P^{-1})$, for some tuning parameter ν. After taking Schur complements, this leads to the following semidefinite program, for a given value of the contraction parameter ρ:

$$\min_{\Sigma \succeq 0, \lambda \geq 0, P \succeq 0, X} \quad \lambda + \nu\,\mathrm{Tr}(\Sigma)$$
$$\text{subject to} \begin{bmatrix} \lambda & X^T \\ X & P \end{bmatrix} \succeq 0, \begin{bmatrix} \Sigma & I_2 \\ I_2 & P \end{bmatrix} \succeq 0, \text{and (51)}.$$

Alternatively, one can minimize $\lambda \mathrm{Tr}(\Sigma)$ for fixed values of λ subject to the constraints above and perform a one-dimensional search for a minimizing value of λ.

Example 2. Let us assume $\mu = 0.1$, $\mathcal{R}_o = 2$, $\tau = 0.1$, $K = 10^{-3}$, $\alpha = 0.25$ in (6), and $\epsilon = 2$, $\delta = 0.05$. That is, we wish to provide a $(2, 0.05)$-differential privacy guarantee for maximum deviations of 0.1% (see the discussion in the previous section). Although not a perfectly rigorous contraction certificate, we sample the continuous set of constraints (51) by sampling the set $\{(s, i) | 0.01 \leq i \leq 0.25, 0.01 \leq s \leq 1 - i\}$ at the values of s, i multiple of

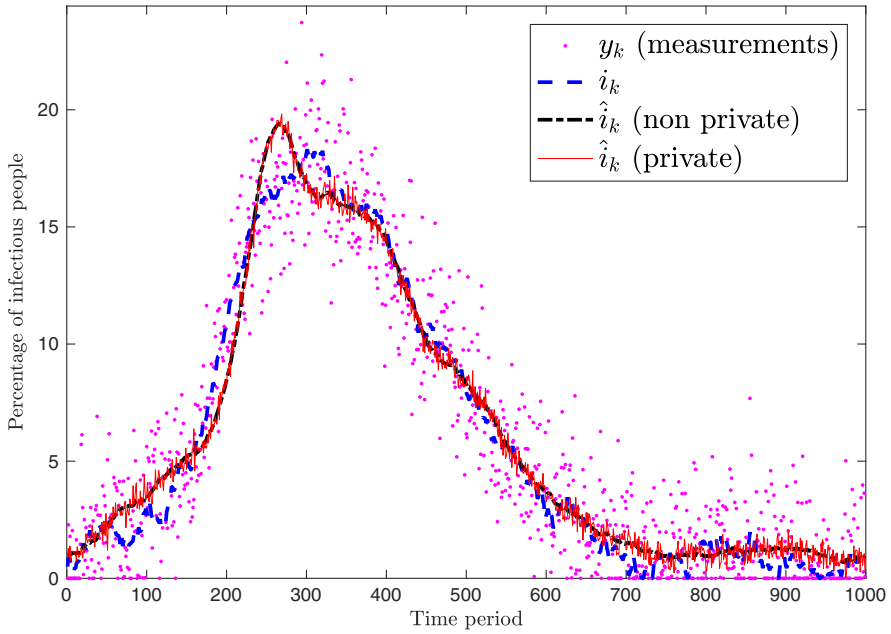

FIGURE 3 Sample path of the estimate of the percentage of infectious people over time produced by the observer. The standard deviations for the dynamics and measurement noise were set to $\sigma_{w_k} I_2 = 0.005\sqrt{\tau} I_2$ and $\sigma_{v_k} = 0.02$, respectively. The signals were truncated to maintain positive values for i, s, y in the simulation. The true proportion of infectious people starts at 0.5%, whereas the estimate used to initialize the observer is 1%. The output of the differentially private observer is not filtered [Colour figure can be viewed at wileyonlinelibrary.com]

0.01, to obtain a finite number of LMIs. A more rigorous approach to enforce these constraints could make use of sum-of-squares programming.[46] Following the procedure above, for the choice $\rho = 0.996$, we obtain the observer gain $H = [3.9304; 0.2003]$ and the covariance matrix Σ with $\Sigma^{1/2} = \begin{bmatrix} 691 & 22 \\ 22 & 17 \end{bmatrix} \times 10^{-4}$ for the privacy-preserving Gaussian noise. Sample trajectories of the nonprivate and private (nonsmoothed) estimates of i are shown on Figure 3.

Remark 11. If model (1)-(2) is linear and time-invariant

$$x_{k+1} = Fx_k + w_k,$$
$$y_k = Cx_k + v_k,$$

then the preceding discussion, replacing F_x in the LMI (51) with the constant matrix F, also provides a method to design a time-invariant observer gain H based on the result of Corollary 8. If moreover we have stochastic Gaussian models for w_k, v_k, a more complex set of LMIs was provided in our previous work[15] to design a type of differentially private Kalman filter using output perturbation.

6 | CONCLUSION

This paper introduces a design methodology for nonlinear observer design, which provides differential privacy guarantees when the measured signals are privacy sensitive, by perturbing the published output signal of the observer. Tools from contraction analysis are used both to enforce convergence of the observer and to set the level of output noise necessary in order to provide the differential privacy guarantee. More concretely, we bound the sensitivity of the observers by leveraging a robustness property of contractive systems. The observer design methodology is illustrated through two examples where we construct estimators for models with nonlinear dynamics or measurements.

ACKNOWLEDGEMENTS

This work was supported by the Natural Sciences and Engineering Research Council of Canada under grants RGPIN-5287-2018 and RGPAS-2018-522686.

ORCID

Jerome Le Ny https://orcid.org/0000-0002-3417-4135

REFERENCES

1. Le Ny J. Privacy-preserving nonlinear observer design using contraction analysis. In: Proceedings of the 54th Annual Conference on Decision and Control (CDC); 2015; Osaka, Japan. http://doi.org/10.1109/CDC.2015.7402922
2. Warren SD, Brandeis LD. The right to privacy. *Harv Law Rev*. 1890;4(5):193-220. http://doi.org/10.2307/1321160
3. McDaniel P, McLaughlin S. Security and privacy challenges in the smart grid. *IEEE Secur Priv*. 2009;7(3):75-77.
4. President's Council of Advisors on Science and Technology. Big data and privacy: a technological perspective. Report to the president. Washington, DC: Executive Office of the President of the United States; 2014.
5. Markey EJ. Tracking and hacking: security and privacy gaps put american drivers at risk. US senator's report. 2015.
6. Narayanan A, Shmatikov V. Robust De-anonymization of large sparse datasets (how to break anonymity of the Netflix prize dataset). In: Proceedings of the IEEE Symposium on Security and Privacy; 2008; Oakland, CA.
7. Calandrino JA, Kilzer A, Narayanan A, Felten EW, Shmatikov V. "You Might Also Like": privacy risks of collaborative filtering. In: Proceedings of the 2011 IEEE Symposium on Security and Privacy; 2011; Berkeley, CA.
8. Wilson DH, Atkeson C. Simultaneous tracking and activity recognition (STAR) using many anonymous, binary sensors. In: *Pervasive Computing: Third International Conference, PERVASIVE 2005, Munich, Germany, May 8-13, 2005. Proceedings*. Berlin, Germany: Springer-Verlag Berlin Heidelberg; 2005:62-79. *Lecture Notes in Computer Science*; vol 3468.
9. Sankar L, Trappe W, Ramchandran K, Poor HV, Debbah M. The role of signal processing in meeting privacy challenges: an overview. *IEEE Signal Process Mag*. 2013;30(5):95-106. Special Issue on Signal Processing for Cybersecurity and Privacy.
10. Dwork C, McSherry F, Nissim K, Smith A. Calibrating noise to sensitivity in private data analysis. In: Proceedings of the Third Theory of Cryptography Conference; 2006; New York, NY.
11. Dwork C, Naor M, Pitassi T, Rothblum GN. Differential privacy under continual observations. In: Proceedings of the Forty-Second ACM Symposium on the Theory of Computing (STOC); 2010; Cambridge, MA.
12. Chan T-HH, Shi E, Song D. Private and continual release of statistics. *ACM Trans Inf Syst Sec*. 2011;14(3):26:1-26:24.
13. Le Ny J, Pappas GJ. Differentially private Kalman filtering. In: Proceedings of the 2012 50th Annual Allerton Conference on Communication, Control, and Computing; 2012; Monticello, IL.
14. Bolot J, Fawaz N, Muthukrishnan S, Nikolov A, Taft N. Private decayed predicate sums on streams. In: Proceedings of the 16th International Conference on Database Theory (ICDT); 2013; Genoa, Italy.
15. Le Ny J, Pappas GJ. Differentially private filtering. *IEEE Trans Autom Control*. 2014;59(2):341-354.
16. Le Ny J, Mohammady M. Differentially private MIMO filtering for event streams. *IEEE Trans Autom Control*. 2018;63(1):145-157.
17. McGlinchey A, Mason O. Differential privacy and the $l1$ sensitivity of positive linear observers. Paper presented at: 20th IFAC World Congress; 2017; Toulouse, France.
18. Lohmiller W, Slotine J-J. On contraction analysis for non-linear systems. *Automatica*. 1998;6:683-696.
19. Sontag ED. Contractive systems with inputs. In: *Perspectives in Mathematical System Theory, Control, and Signal Processing: A Festschrift in Honor of Yutaka Yamamoto on the Occasion of his 60th Birthday*. Berlin, Germany: Springer-Verlag Berlin Heidelberg; 2010:217-228.
20. Forni F, Sepulchre R. A differential Lyapunov framework for contraction analysis. *IEEE Trans Autom Control*. 2014;59(3):614-628.
21. Angeli D. A Lyapunov approach to incremental stability properties. *IEEE Trans Autom Control*. 2000;47(3):410-421.
22. Holland PW, Laskey KB, Leinhardt S. Stochastic blockmodels: first steps. *Soc Networks*. 1983;5(2):109-137.
23. Sontag ED, Wang Y. Output-to-state stability and detectability of nonlinear systems. *Syst Control Lett*. 1997;29(5):279-290.
24. Gauthier J-P, Kupka I. *Deterministic Observation Theory and Applications*. Cambridge, UK: Cambridge University Press; 2001.
25. Isidori A. *Lectures in Feedback Design for Multivariable Systems*. Basel, Switzerland: Springer International Publishing; 2017.
26. Cortés J, Dullerud GE, Han S, Le Ny J, Mitra S, Pappas GJ. Differential privacy in control and network systems. In: Proceedings of the 2016 IEEE 55th Conference on Decision and Control (CDC); 2016; Las Vegas, NV.
27. Chung S-J, Bandyopadhyay S, Chang I, Hadaegh FY. Phase synchronization control of complex networks of Lagrangian systems on adaptive digraphs. *Automatica*. 2013;49(5):1148-1161.
28. Khalil HK. *Nonlinear Systems*. Upper Saddle River, NJ: Prentice Hall; 2002.
29. Lewis DC. Metric properties of differential equations. *Am J Math*. 1949;71(2):294-312.
30. Hartman P. *Ordinary Differential Equations*. 2nd ed. Cambridge, MA: Birkhäuser; 1982.
31. Aghannan N, Rouchon P. An intrinsic observer for a class of Lagrangian systems. *IEEE Trans Autom Control*. 2003;48(6):936-945.
32. Pavloc A, van de Wouw N, Nijmeijer H. *Uniform Output Regulation of Nonlinear Systems: A Convergent Dynamics Approach*. New York, NY: Birkhäuser Boston; 2006.
33. Russo G, di Bernardo M, Sontag ED. Global entrainment of transcriptional systems to periodic inputs. *PLoS Comput Biol*. 2010;6(4).
34. do Carmo MP. *Riemannian Geometry*. Boston, MA: Birkhäuser; 1992.
35. Shen Z. *Lectures on Finsler Geometry*. Singapore: World Scientific Publishing; 2001.

36. Vidyasagar M. *Nonlinear Systems Analysis*. 2nd. Englewood Cliffs, NJ: Prentice Hall; 1993.
37. Lakshmikantham V, Trigiante D. *Theory of Difference Equations: Numerical Methods and Applications*. New York, NY: Marcel Dekker; 2002.
38. Besançon G, ed. *Nonlinear Observers and Applications*. New York, NY: Springer Science+Business Media; 2007.
39. Xu KS, Hero III AS. Dynamic stochastic blockmodels for time-evolving social networks. *J Sel Topics Signal Proc*. 2014;8(4):552-562. Special Issue on Signal Processing for Social Networks.
40. Wasserman L, Zhou S. A statistical framework for differential privacy. *J Am Statist Assoc*. 2010;105(489):375-389.
41. Lawson AB, Kleinman K. *Spatial & Syndromic Surveillance For Public Health*. Chichester, UK: John Wiley & Sons; 2005.
42. Ginsberg J, Mohebbi MH, Patel RS, Brammer L, Smolinski MS, Brilliant L. Detecting influenza epidemics using search engine query data. *Nature*. 2009;457:1012-1014.
43. Skvortsov A, Ristic B. Monitoring and prediction of an epidemic outbreak using syndromic observations. *Math Biosci*. 2012;240(1):12-19.
44. Kermack WO, McKendrick AG. A contribution to the mathematical theory of epidemics. *Proc Royal Soc Lond Ser A*. 1927;115(772):700-721.
45. Brauer F, van den Driessche P, Wu J. *Mathematical Epidemiology*. Berlin, Germany: Springer-Verlag Berlin Heidelberg; 2008. *Lecture Notes in Mathematics*; vol 1945.
46. Aylward EM, Parrilo PA, Slotine J-JE. Stability and robustness analysis of nonlinear systems via contraction metrics and SOS programming. *Automatica*. 2008;44(8):2163-2170.

How to cite this article: Le Ny J. Differentially private nonlinear observer design using contraction analysis. *Int J Robust Nonlinear Control*. 2018;1–19. https://doi.org/10.1002/rnc.4392