A dynamical approach to privacy preserving average consensus

Claudio Altafini

Altafini, C., (2019), A dynamical approach to privacy preserving average consensus, *2019 IEEE 58th Conference on Decision and Control (CDC)*. https://doi.org/10.1109/CDC40024.2019.9029712

A dynamical approach to privacy preserving average consensus

Claudio Altafini

***Abstract*— In this paper we propose a novel method for achieving average consensus in a continuous-time multiagent network while avoiding to disclose the initial states of the individual agents. In order to achieve privacy protection of the state variables, we introduce maps, called output masks, which alter the value of the states before transmitting them. These output masks are local (i.e., implemented independently by each agent), deterministic, time-varying and converging asymptotically to the true state. The resulting masked system is also time-varying and has the original (unmasked) system as its limit system. It is shown in the paper that the masked system has the original average consensus value as its only attractor. However, in order to preserve privacy, it cannot share an equilibrium point with the unmasked system, meaning that in the masked system the attractor cannot be also stable.**

I. Introduction

Preserving privacy in a multiagent system means performing a computation in a distributed manner among the agents of a network without revealing the individual values that the agents contribute to the computation process. For instance a privacy preserving average consensus problem consists in computing the mean of the state variables of the agents without disclosing the initial state values to the other agents or to external observers. Several approaches have been proposed in recent years for this task. One of them relies on differential privacy [6], [7], which consists in adding to the state being transmitted by an agent a noise from an appropriate source. In this way, even if the value is publicly broadcasted, the knowledge that an observing agent can acquire of the true state is limited to a predetermined precision. In the average consensus problem, this method has been investigated for instance in [5], [8], [9], [17]. Another approach relies on cryptography. Encrypted messages can be exchanged among the agents in various ways, e.g. through trusted third parties [10], obfuscation [3], or through distributed cryptography schemes [12], [22]. For instance in [13], [15], [20] the encryption is realized as a perturbation with zero sum (or integral) over time. A third approach is based on understanding what is observable and what is not at a node [1], [19], and on trying to guarantee privacy as a loss of observability.

The scope of this paper is to propose a novel approach for privacy-preserving average consensus. Our approach is inspired by system-theoretical considerations and relies crucially on interpreting a distributed computation problem like consensus as a dynamical system, hence we refer to it as *dynamical privacy*. We push the idea of "lack of observability as a form of privacy" to its extreme, by defining output maps that we call *output masks* which are altering (or "masking") the internal state of an agent before it is transmitted to the neighboring agents. As these output masks remain private to each agent, they do not allow for any form of observability of the state of the agents. Our output masks are deterministic, time-varying transformations that each agent can implement independently, and that asymptotically converge to the true internal state. The resulting masked system is also a time-varying system. Its characteristic feature is that its limit system [4] corresponds to the original (unmasked) system, i.e., to the average consensus dynamics. If the average consensus problem being investigated is on a static graph, then the masked system is a case of asymptotically autonomous time-varying system [4], [14].

We show in the paper that for suitably chosen output masks, in the consensus problem privacy protection can always be guaranteed even in the worst cases, for instance in the degenerate situation of all agents having the same initial state (and hence being already at the consensus value). In order to do so, an output mask must be able to escape neighborhoods of any point (i.e., it cannot be stable), feature which can be achieved by using output masks which are inhomogeneous in the state they are hiding. This is reminiscent of the additive noise term used in differential privacy. A consequence of having inhomogeneous output maps is that the masked system lacks equilibria (again, fixed points in the masked dynamics would lead to breaching of privacy for certain initial conditions). In spite of the absence of equilibria, we shown in the paper that our time-varying masked system has the average consensus value as a uniform attractor for its dynamics. In fact, the masked system converges asymptotically to the original consensus problem, and for all times a conservation law like preservation of the average of the true states remains valid also on the masked system.

Technically, global attractivity (on the complement of the agreement subspace) can be shown using the same Lyapunov function of a standard consensus problem, but without requiring its derivative along the trajectories of the masked system to remain nonpositive. As already mentioned, the lack of stability around the average consensus value is needed to guarantee indiscernibility of the initial conditions from the masked outputs. Global attractivity is obtained if the derivative of the Lyapunov function is upper bounded by terms that decay to 0 asymptotically.

In this paper, the continuous-time version of the average consensus problem is considered. Unlike [20], we do not

Work supported in part by a grant from the Swedish Research Council (grant n. 2015-04390).

C. Altafini is with the Division of Automatic Control, Department of Electrical Engineering, Linköping University, SE-58183 Linköping, Sweden. E-mail: claudio.altafini@liu.se

require that our perturbations have time integral that asymptotically tend to 0. However, as in [20], we must impose that in-neighbors of the agents are not completely contained into each other, assumption also considered in [15]. The other conditions under which our dynamical privacy guarantees privacy protection are mild and reasonable. Namely, we require that only the internal state and the parameters of the output mask are kept private to each agent, while the masked output is communicated to the first neighbors on the interaction graph, and the Laplacian of the problem can be publicly available. It is also worth observing that breaching the privacy of an agent does not compromise that of the remaining agents, as each output mask is created locally by each agent.

II. Preliminaries

A continuous function $\alpha : [0, \infty) \to [0, \infty)$ is said to belong to class $\mathcal{K}_\infty$ if it is strictly increasing and $\alpha(0) = 0$. Subclasses of $\mathcal{K}_\infty$ which are homogeneous polynomials of order i will be denoted $\mathcal{K}^i_\infty$: $\alpha(r) = ar^i$ for some constant $a > 0$. A continuous function $\zeta : [0, \infty) \to [0, \infty)$ is said to belong to class $\mathcal{L}$ if it is decreasing and $\lim_{t\to\infty} \zeta(t) = 0$. In particular, we are interested in $\mathcal{L}$ functions that are exponentially decreasing: $\zeta(t) = ae^{-\delta t}$ for some $a > 0$ and $\delta > 0$. We shall denote such subclass $\mathcal{L}^e \subset \mathcal{L}$. A continuous function $\beta : [0, \infty) \times [0, \infty) \to [0, \infty)$ is said to belong to class $\mathcal{KL}^{i,e}_\infty$ if the mapping $\beta(r, t)$ belongs to class $\mathcal{K}^i_\infty$ for each fixed t and to class $\mathcal{L}^e$ for each fixed r, i.e., $\beta(r, t) = ar^i e^{-\delta t}$ for some $a > 0$ and $\delta > 0$.

Consider

$$\dot{x} = g(t, x), \qquad x(t_o) = x_o \tag{1}$$

where $g : \mathbb{R}_+ \times \mathbb{R}^n \to \mathbb{R}^n$ is Lipschitz continuous in x, measurable in t, and such that for each $x_o \in \mathbb{R}^n$ and each $t_o \in \mathbb{R}_+$ the solution of (1), $x(t, x_o)$, exists in $[0, \infty)$. A point $x^* \in \mathbb{R}^n$ is an equilibrium point of (1) if $g(t, x^*) = 0$ for a.e. [1] $t \geq t_o$.

A point $x^* \in \mathbb{R}^n$ is *uniformly globally attractive* for (1) if for each $\nu > 0$ there exists $T = T(\nu) > 0$ such that for each solution $x(t, x_o)$ of (1) it holds that $\|x(t, x_o) - x^*\| < \nu$ for each $t > t_o + T$, each $x_o \in \mathbb{R}^n$ and each $t_o \geq 0$. In particular, if x^* is a uniformly global attractor for (1), then as $t \to \infty$ all trajectories $x(t, x_o)$ converge to x^* uniformly in t for all $t_o \geq 0$ and x_o. A point x^* can be attractive for (1) without being an equilibrium of (1) (we will use this fact extensively in the paper).

Given (1), denote $g_s(t, x)$ the translate of $g(t, x)$: $g_s(t, x) = g(t + s, x)$. A (possibly time-dependent) system $\dot{x} = \tilde{g}(t, x)$ is called a *limit system* of (1) if there exists a sequence $\{s_k\}$, $s_k \to \infty$ as $k \to \infty$, such that $g_{s_k}(t, x)$ converges to $\tilde{g}(t, x)$ [4]. An existence condition for a limit system $\tilde{g}(t, x)$ is given in Lemma 1 of [11]: when $g(t, x)$ is a uniformly continuous and bounded function, then there exists increasing and diverging sequences $\{s_k\}$ such that on compact subsets of $\mathbb{R}^n$ $g_{s_k}(t, x)$ converges uniformly to a continuous limit function $\tilde{g}(t, x)$ on every compact of $[0, \infty)$. In general the limit system may not be unique nor time-invariant. However, when it exists unique, then it must be autonomous [4], [21] because all translates $g_{s+s'}(t, x)$ must have themselves a limit system hence the latter cannot depend on time. The time-varying system (1) is called *asymptotically autonomous* in this case.

The ω-limit set of $x(t, x_o)$, denoted Ω_{x_o}, consists of all points x^* such that a sequence $\{t_k\}$, with $t_k \to \infty$ when $k \to \infty$, exists for which $\lim_{k\to\infty} x(t_k, x_o) = x^*$. For time-varying systems, if a solution is bounded then the corresponding Ω_{x_o} is nonempty, compact and approached by $x(t, x_o)$. However, it need not be invariant. Only for limit systems the invariance property may hold, although not necessarily (it may fail even for asymptotically autonomous systems, see [4]).

The following lemma is inspired by [16], Thm 2.1 and [23], Prop. 5, and provides us with a suitable comparison function to be used later in the paper. The proof of this and other results is omitted due to lack of space. It can be found on the preprint [2] available in the arXiv.

Lemma 1 *Consider the scalar system*

$$\dot{v} = -\alpha(v) + \beta(v, t) + \zeta(t), \qquad v(t_o) = v_o \geq 0. \tag{2}$$

If $\alpha(v) \in \mathcal{K}^2_\infty$, $\beta \in \mathcal{KL}^{1,e}_\infty$ and $\zeta \in \mathcal{L}^e$, then the solutions of (2) *are all prolongable to ∞ and bounded $\forall\, v_o \geq 0$ and $\forall t_o \geq 0$. Furthermore,*

$$\lim_{t\to\infty} v(t) = 0 \qquad \forall\, v_o \geq 0, \quad \forall\, t_o \geq 0.$$

III. Problem formulation

Consider a distributed dynamical system on a graph with n nodes:

$$\dot{x} = f(x), \qquad x(0) = x_o, \tag{3}$$

where $x = \begin{bmatrix} x_1 & \dots & x_n \end{bmatrix}^T \in \mathbb{R}^n$ is a state vector and $f = [f_1 \dots f_n]^T : \mathbb{R}^n \to \mathbb{R}^n$ is a Lipschitz continuous vector field. Standing assumptions in this paper are that (3) possesses a unique solution continuable on $[0, \infty)$ for all $x_o \in \mathbb{R}^n$ and that information can be exchanged only between first neighbors on the graph, i.e.,

$$\dot{x}_i = f_i(x_i, x_j, j \in \mathcal{N}_i), \qquad i = 1, \dots, n \tag{4}$$

with $\mathcal{N}_i$ the in-neighborhood of node i. Furthermore, to avoid trivial situations, we impose that $\mathcal{N}_i$ is the "essential neighborhood" of agent i [15], i.e., $\forall\, i = 1 \dots, n$

$$f_i(x_i, x_j, j \in \tilde{\mathcal{N}}_i) \neq f_i(x_i, x_j, j \in \mathcal{N}_i) \;\; \forall \tilde{\mathcal{N}}_i \subsetneq \mathcal{N}_i. \tag{5}$$

We are interested in cases in which the presence of a conservation law (as in the consensus problem) leads to exponential stability on some submanifold depending on the initial conditions, i.e., $\lim_{t\to\infty} x(t) = x^*(x_o)$.

The *privacy preservation problem* consists in using a system like (3) to perform the computation of x^* in a distributed manner, while avoiding to divulgate the initial condition x_o to the other nodes. Clearly this cannot be achieved directly

[1] almost every, i.e., except for at most a set of Lebesgue measure 0.

on the system (3) which is based on exchanging the values x_i between the nodes. It can however be achieved if we insert a mask on the value $x(t)$ which preserves convergence to x^*, at least asymptotically. The masks we propose in this paper have the form of time-varying output maps.

A. Output masks

Consider a continuously differentiable time-varying output map

$$\begin{aligned} h : \mathbb{R}_+ \times \mathbb{R}^n \times \mathbb{R}^m &\to \mathbb{R}^n \\ (t,\, x,\, \pi) &\mapsto y(t) = h(t, x(t), \pi) \end{aligned} \tag{6}$$

where $y = \begin{bmatrix} y_1 & \dots & y_n \end{bmatrix}^T \in \mathbb{R}^n$ is an output vector of the same size as x, and $\pi \in \mathbb{R}^m$ is a vector of parameters splittable into n subvectors (not necessarily of the same dimension), one for each node of the network: $\pi = \{\pi_1, \dots, \pi_n\}$.

In the following we refer to $h(t, x(t), \pi)$ as an *output mask* and to y as a *masked output*. The state x of the system is first masked into y and then sent to the first out-neighbors on the graph. The original system (3) can therefore be modified into the following *masked system:*

$$\dot{x} = f(y) \tag{7a}$$
$$y = h(t, x, \pi). \tag{7b}$$

We assume in what follows that the vector field $f(\cdot)$ is publicly known and that each node knows the output trajectories $y_i(t)$ of its in-neighbors. The state x and the output mask $h(t, x, \pi)$ (functional form plus values of the parameters π) are instead private to each agent, as explained more in detail next.

Definition 1 *A C^1 output map h is said a* local mask *if it has components that are local, i.e.,*

P1: $h_i(t, x, \pi) = h_i(t, x_i, \pi_i) \qquad i = 1, \dots, n.$

The property of locality guarantees that the output map h_i can be independently decided by node i. Both the functional form chosen for $h_i(\cdot)$ and the numerical value of the parameters π_i can therefore remain hidden to the other agents.

In order to confound an agent monitoring the communications, the output map needs also to avoid mapping neighborhoods of a point x^* of (3) (typically an equilibrium point) into themselves.

Definition 2 *A C^1 output map h is said* not to preserve neighborhoods *of a point x^* if for all small $\epsilon > 0$, $\|x_o - x^*\| < \epsilon$ does not imply $\|h(0, x_o, \pi) - x^*\| < \epsilon$.*

These notions are used in the following definition.

Definition 3 *A C^1 output map h is said a* privacy mask *if it is a local mask and in addition*

P2: $h_i(0, x_i, \pi_i) \neq x_i \;\forall\, x_i \in \mathbb{R}^n,\; i = 1, \dots, n;$

P3: $h(t, x, \pi)$ *does not preserve neighborhoods of any* $x \in \mathbb{R}^n$*;*

P4: $h_i(t, x_i, \pi_i)$ *strictly increasing in* x_i *for each fixed* t *and* $\pi_i,\; i = 1, \dots, n.$

Property P4 resembles a definition of $\mathcal{K}_\infty$ function, but it is in fact more general: $x = 0$ is not a fixed point of h for any finite t, and h need not be nonnegative in x. It follows from Property P4 and locality that h is a bijection in x for each fixed t and π, although one that does not preserve the origin.

In many cases, it will be necessary to impose that the privacy mask converges asymptotically to the true state, i.e., that the perturbation induced by the mask is vanishing.

Definition 4 *The output map h is said a* vanishing privacy mask *if it is a privacy mask and in addition*

P5: $|h_i(t, x_i, \pi_i) - x_i|$ *is decreasing in* t *for each fixed* x_i *and* π_i*, and* $\lim_{t\to\infty} h_i(t, x_i, \pi_i) = x_i,\; i = 1, \dots, n.$

B. Examples of output masks

The following are examples of output masks.

a) Linear mask:

$$h_i(t, x_i, \pi_i) = (1 + \phi_i e^{-\sigma_i t}) x_i, \qquad \phi_i \geq 0, \quad \sigma_i > 0 \tag{8}$$

(i.e., $\pi_i = \{\phi_i,\, \sigma_i\}$). This local vanishing mask is not a proper privacy mask since $h_i(0, 0, \pi_i) = 0$ i.e. the origin is not masked. Notice that all homogeneous maps have this problem (and they fail to escape neighborhoods of x_i).

b) Additive mask:

$$h_i(t, x_i, \pi_i) = x_i + \gamma_i e^{-\delta_i t}, \qquad \delta_i > 0, \quad \gamma_i \neq 0 \tag{9}$$

(i.e., $\pi_i = \{\delta_i,\, \gamma_i\}$) is a vanishing privacy mask.

c) Affine mask:

$$h_i(t, x_i, \pi_i) = c_i(x_i + \gamma_i e^{-\delta_i t}), \quad c_i > 1, \quad \delta_i > 0, \quad \gamma_i \neq 0 \tag{10}$$

(i.e., $\pi_i = \{c_i,\, \delta_i,\, \gamma_i\}$) is also a privacy mask. Since $\lim_{t\to\infty} h_i(t,\, x_i,\, \pi_i) = c_i x_i$, it is however not vanishing.

d) Vanishing affine mask:

$$\begin{aligned} h_i(t, x_i, \pi_i) =& (1 + \phi_i e^{-\sigma_i t})(x_i + \gamma_i e^{-\delta_i t}), \\ & \phi_i > 0, \quad \sigma_i > 0, \quad \delta_i > 0, \quad \gamma_i \neq 0 \end{aligned} \tag{11}$$

(i.e., $\pi_i = \{\phi_i,\, \sigma_i,\, \delta_i,\, \gamma_i\}$). This privacy mask is also vanishing. Notice that in vector form, assuming all nodes adopt it, the vanishing affine mask can be expressed as

$$h(t,\, x, \pi) = (I + \Phi e^{-\Sigma t})(x + e^{-\Delta t}\gamma) \tag{12}$$

where $\Phi = \operatorname{diag}(\phi_1, \dots, \phi_n)$, $\Sigma = \operatorname{diag}(\sigma_1, \dots, \sigma_n)$, $\Delta = \operatorname{diag}(\delta_1, \dots, \delta_n)$, and $\gamma = \begin{bmatrix} \gamma_1 & \dots & \gamma_n \end{bmatrix}^T$.

C. Dynamically private systems

The problem of privacy preserving as it is formulated here cannot be cast as an observability problem, as each $h_i(\cdot)$ is unknown to the other agents. However, even if privacy masks are "privacy enabling", by themselves they do not guarantee that in a system like (7) x_o cannot be reconstructed asymptotically.

To make things more precise, we introduce the following definition. Consider the system (7). Denote $y(t, x_o)$, of

components $y_i(t, x_o)$, $i = 1, \ldots, n$, the output trajectory of (7) from the initial state x_o, of components $x_{o,i}$.

Definition 5 *The initial condition of agent i, $x_{o,i}$, is said* indiscernible for agent j *($j \neq i$) if knowledge of the output trajectories $y_k(t, x_o)$, $k \in \mathcal{N}_j$, $t \in [t_o, \infty)$, and of the vector field $f(\cdot)$ is not enough to reconstruct $x_{o,i}$ in* (7)*. An initial condition x_o is said* indiscernible *if all its components $x_{o,i}$ are indiscernible for all agents $j \in \{1, \ldots, n\} \setminus \{i\}$.*

When an initial condition is not indiscernible, then it must be discernible or partially discernible (with obvious meaning of these terms) for at least one agent.

If we consider observations $y(t, x_o)$ over a finite horizon, i.e., $t \in [t_o, t_1]$, $t_1 < \infty$, then in order to have discernible initial states, the following three conditions must all be satisfied:

(i) The exact functional form of the output mask $h(\cdot)$ must be known;
(ii) The parameters π must be identifiable given the trajectory $y(t, x_o)$ and the vector field $f(\cdot)$;
(iii) The system (7) must be observable.

However, encoding indiscernibility simply as a systematic violation of these conditions at all nodes is not enough, because when $t_1 \to \infty$ more subtle forms of disclosure of x_o may appear if $\lim_{t\to\infty} y(t) = x(t)$ (case normally considered in the following). In fact, assuming convergence and denoting $\lim_{t\to\infty} y_i(t) = y_i^* = x_i^*$, from

$$x_i^* = x_{o,i} + \int_0^\infty f_i(y)dt, \tag{13}$$

if $f_i(y)$ and y_i^* are both known, then $x_{o,i}$ can be estimated from (13). In order to guarantee that indiscernibility does not get lost asymptotically, we have therefore to impose that no agent $j \neq i$ can estimate the integral in (13). Since $\int_0^\infty f_i(y)dt = \int_0^\infty f_i(y_i, y_k, k \in \mathcal{N}_i)dt$, this is achieved if we make the following assumption [20] (see also [15], Corollary 1).

Assumption 1 (No overlapping neighborhoods) *The system* (7) *is such that $\{\mathcal{N}_i \cup \{i\}\} \nsubseteq \{\mathcal{N}_j \cup \{j\}\}$, $\forall\, i, j = 1, \ldots, n$, $i \neq j$.*

Assumption 1 guarantees that no node has complete information of what is going on at the other nodes. This is a condition on the topology of the graph, and therefore a *system* property, rather than simply a property of well-conceived output maps.

Combining with privacy of the output masks, we can formulate the following definition.

Definition 6 *The system* (7) *is called a* dynamically private *version of* (3) *if*

1) h is a privacy mask;
2) the solution of (7) *exists unique in* $[0, \infty)$ *and is bounded $\forall\, x_o \in \mathbb{R}^n$;*
3) $\lim_{t\to\infty} y(t) = x(t)$;
4) indiscernibility of the initial condition is guaranteed.

The next proposition relates indiscernibility to Assumption 1.

Proposition 1 *If the system* (7) *satisfies conditions 1-3 of Definition 6 and Assumption 1, then it is a dynamically private version of* (3)*.*

The privacy property of $h(\cdot)$ suggests that in a dynamically private system we cannot have equilibrium points and therefore we cannot talk about stability (of equilibria), while convergence of $y(t)$ to $x(t)$ suggests that as long as $f(\cdot)$ is autonomous, a dynamically private system is asymptotically autonomous with the unmasked system as limit system. This can be shown to be always true if the output mask is vanishing.

Proposition 2 *If* (7) *is a dynamically private version of* (3)*, then it cannot have equilibrium points. Furthermore, if $h(\cdot)$ is a vanishing privacy mask, then the system* (7) *is asymptotically autonomous with limit system* (3)*.*

IV. Privacy-preserving average consensus

In the average consensus problem, $f(x) = -Lx$, with L a weight-balanced Laplacian matrix: $L\mathbf{1} = L^T\mathbf{1} = 0$, with $\mathbf{1} = \begin{bmatrix}1 & \ldots & 1\end{bmatrix}^T \in \mathbb{R}^n$. When L is irreducible, the equilibrium point is $x^*(x_o) = \mathbf{1}^T x_o/n$. The system has a continuum of equilibria, described by $\mathrm{span}(\mathbf{1})$, and each $x^*(x_o)$ is globally asymptotically stable in $\mathrm{span}(\mathbf{1})^\perp$, see [18].

Theorem 1 *Consider the system*

$$\dot{x} = -Lx, \qquad x(0) = x_o \tag{14}$$

where L is an irreducible, weight-balanced Laplacian matrix, and denote $\eta = \boldsymbol{1}^T x_o/n$ its average consensus value. Then $x^* = \eta\boldsymbol{1}$ is a global uniform attractor on $\mathrm{span}(\boldsymbol{1})^\perp$ for the masked system

$$\begin{aligned} \dot{x} &= -Ly \\ y &= h(t,\, x,\, \pi) = \left(I + \Phi e^{-\Sigma t}\right)\left(x - e^{-\Delta t}\gamma\right). \end{aligned} \tag{15}$$

Furthermore, if Assumption 1 holds, then (15) *is a dynamically private version of* (14)*.*

Proof: Notice first that the system (15) can be written as

$$\dot{x} = -L\left(I + \Phi e^{-\Sigma t}\right)\left(x + e^{-\Delta t}\gamma\right), \tag{16}$$

from which it is clear that the system (15) cannot have equilibrium points. It is also clear from (16) that $\mathbf{1}^T\dot{x} = 0$ i.e., also (15) obeys to the conservation law $\mathbf{1}^T x(t) = \mathbf{1}^T x_o = \eta\mathbf{1}$. As in the standard consensus problem [18], we can therefore work on the $n-1$ dimensional projection subspace $\mathrm{span}(\mathbf{1})^\perp$ and consider the time-varying Lyapunov function for the "displacement vector" $x - \eta\mathbf{1} \in \mathrm{span}(\mathbf{1})^\perp$:

$$V(t, x) = (x - \eta\mathbf{1})^T\left(I + \Phi e^{-\Sigma t}\right)(x - \eta\mathbf{1}).$$

From now on we assume that all calculations are restricted to $\mathrm{span}(\mathbf{1})^\perp$. The derivative of V along the solutions of (15) is

$$\begin{aligned}
&\dot V(t,\,x) = \frac{\partial V}{\partial x}\dot x + \frac{\partial V}{\partial t}\\
&= -\,2(x-\eta\mathbf{1})^T\left(I+\Phi e^{-\Sigma t}\right)L\left(I+\Phi e^{-\Sigma t}\right)\left(x+e^{-\Delta t}\gamma\right)\\
&\quad -\,(x-\eta\mathbf{1})^T\left(\Sigma\Phi e^{-\Sigma t}\right)(x-\eta\mathbf{1})\\
&= -\,(x-\eta\mathbf{1})^T\left(I+\Phi e^{-\Sigma t}\right)(L+L^T)\left(I+\Phi e^{-\Sigma t}\right)(x-\eta\mathbf{1})\\
&\quad -\,\eta(x-\eta\mathbf{1})^T\left(I+\Phi e^{-\Sigma t}\right)(L+L^T)\left(I+\Phi e^{-\Sigma t}\right)\mathbf{1}\\
&\quad -\,(x-\eta\mathbf{1})^T\left(I+\Phi e^{-\Sigma t}\right)(L+L^T)\left(I+\Phi e^{-\Sigma t}\right)e^{-\Delta t}\gamma\\
&\quad -\,(x-\eta\mathbf{1})^T\left(\Sigma\Phi e^{-\Sigma t}\right)(x-\eta\mathbf{1}).
\end{aligned}\tag{17}$$

Since $\phi_i > 0$, it is $1+\phi_i e^{-\sigma_i t} \geq 1\ \forall\, t \geq 0$, and $I+\Phi e^{-\Sigma t}$ is a positive definite diagonal matrix, for the first term of (17) we have

$$\begin{aligned}
&(x-\eta\mathbf{1})^T\left(I+\Phi e^{-\Sigma t}\right)(L+L^T)\left(I+\Phi e^{-\Sigma t}\right)(x-\eta\mathbf{1})\\
&\geq (x-\eta\mathbf{1})^T(L+L^T)(x-\eta\mathbf{1}) \geq \alpha_1(\|x-\eta\mathbf{1}\|) > 0
\end{aligned}$$

for some function $\alpha_1 \in \mathcal{K}_\infty^2$. The second term of (17) is linear in $\|x-\eta\mathbf{1}\|$, and from $L\mathbf{1} = L^T\mathbf{1} = 0$, we have

$$\begin{aligned}
&-\,\eta(x-\eta\mathbf{1})^T\left(I+\Phi e^{-\Sigma t}\right)(L+L^T)\left(I+\Phi e^{-\Sigma t}\right)\mathbf{1}\\
&= -\eta(x-\eta\mathbf{1})^T\left(I+\Phi e^{-\Sigma t}\right)(L+L^T)\Phi e^{-\Sigma t}\mathbf{1}\\
&\leq \beta_1(\|x-\eta\mathbf{1}\|,\,t)
\end{aligned}$$

for some function $\beta_1 \in \mathcal{KL}_\infty^{1,e}$. Similarly, for the third term of (17),

$$\begin{aligned}
&-\,(x-\eta\mathbf{1})^T\left(I+\Phi e^{-\Sigma t}\right)(L+L^T)\left(I+\Phi e^{-\Sigma t}\right)e^{-\Delta t}\gamma\\
&\leq \beta_2(\|x-\eta\mathbf{1}\|,\,t)
\end{aligned}$$

for some $\beta_2 \in \mathcal{KL}_\infty^{1,e}$. Finally, the fourth term of (17) is

$$(x-\eta\mathbf{1})^T\left(\Sigma\Phi e^{-\Sigma t}\right)(x-\eta\mathbf{1}) = \alpha_2(\|x-\eta\mathbf{1}\|,\,t)$$

for some $\alpha_2 \in \mathcal{KL}_\infty^{2,e}$, i.e., it is positive definite for all finite t, and vanishes as $t \to \infty$. Hence there exists $\alpha \in \mathcal{K}_\infty^2$ such that

$$\alpha(v) \geq \alpha_1(v) + \alpha_2(v,t) > 0 \quad \forall\, v \in \mathbb{R}^+.$$

Denote $\beta(\|x-\eta\mathbf{1}\|,t) \in \mathcal{KL}_\infty^{1,e}$ a proper majorization of $\beta_j(\|x-\eta\mathbf{1}\|,t)$, $j=1,\,2$. Since, for all t, V is quadratic, positive definite, radially unbounded and vanishing in $x = \eta\mathbf{1}$, there exists two class $\mathcal{K}_\infty^2$ functions α_3 and α_4 such that

$$\alpha_3(\|x-\eta\mathbf{1}\|) \leq V(t,x) \leq \alpha_4(\|x-\eta\mathbf{1}\|). \tag{18}$$

We can therefore apply the comparison lemma, using (2) with initial condition $v(0) = V(0,x_o)$, where x_o such that $\mathbf{1}^T x_o/n = \eta$. From Lemma 1, it follows that it must be $\lim_{t\to\infty} V(t,x(t)) = 0$ for all x_o such that $\mathbf{1}^T x_o/n = \eta$, hence from (18) $\lim_{t\to\infty}\alpha_3(\|x-\eta\mathbf{1}\|) = 0$ or $\lim_{t\to\infty} x(t) = \eta\mathbf{1}$ for all x_o such that $\mathbf{1}^T x_o/n = \eta$. Since $h(t,x,\pi) = \left(I+\Phi e^{-\Sigma t}\right)\left(x+e^{-\Delta t}\gamma\right)$ is a privacy mask, $\lim_{t\to\infty} y(t) = \lim_{t\to\infty} x(t) = \eta\mathbf{1}$ and Assumption 1 holds, from Proposition 1 (15) is a dynamically private version of (14). ∎

Corollary 1 *The masked system* (15) *is asymptotically autonomous with* (14) *as limit system. The ω-limit set of* (15) *is given by $\{(\mathbf{1}^T x_o/n)\mathbf{1}\}$ for each x_o.*

Remark 1 Even if (14) has $x^* = \eta\mathbf{1}$ as a globally asymptotically stable equilibrium point in $\mathrm{span}(\mathbf{1})^\perp$, the masked system (7) does not have equilibria and, because of the extra inhomogeneous term in the right hand side of (16), not even stability of $\eta\mathbf{1}$ is guaranteed. Nevertheless, $x^* = \eta\mathbf{1}$ remains a global attractor for all trajectories of the system in $\mathrm{span}(\mathbf{1})^\perp$.

Remark 2 Since the evolution of the masked system (14) is restricted to the $n-1$ dimensional subspace $\mathrm{span}(\mathbf{1})^\perp$, our masked consensus problem (as any exact privacy preserving consensus scheme) make sense only when $n > 2$. The case $n = 2$ never satisfies Assumption 1 when L is irreducible.

Example 1 In Fig. 1 a private consensus problem is run among $n = 100$ agents. Both $x(t)$ (private) and $y(t)$ (public) converge to the same consensus value $\eta = \mathbf{1}^T x(0)/n$, but the initial condition $y(0)$ does not reflect $x(0)$, not even when $x_i(0)$ is already near η ($h(\cdot)$ does not preserve neighborhoods, see panel (c) of Fig. 1). Notice that $\mathbf{1}^T x(t)/n$ is constant over t, while $\mathbf{1}^T y(t)/n$ is not, which confirms that the output masks indeed act as confounding factors. Notice further that a standard Lyapunov function used for consensus, like $V_{mm}(t) = \max_i(x_i(t)) - \min_i(x_i(t))$, does not work in our privacy-preserving scheme (see panel (d) of Fig. 1), which reflects the fact that the system (15) is not asymptotically stable in $\mathrm{span}(\mathbf{1})^\perp$. Violation of $\dot V_{mm} \leq 0$ is however not systematic but depending on the initial conditions, see Fig. 2. The convergence speed of the time-dependent part can be manipulated by selecting the factors σ_i and δ_i appropriately.

V. Conclusions

When a distributed computation can be thought of as a multiagent dynamical system, then the problem of protecting the initial states of the agents can be formulated using classical tools from systems and control theory. For cases like average consensus, such system-theoretical framework provides a full solution: it exploits the naturally converging character of the original dynamics while at the same time hiding the original initial states through output masks which render the system non-observable, and hence the true state non-estimable from the masked output. The framework, here applied only to average consensus for lack of space, is fairly general, and valid for a broad range of multiagent problems, see [2]. Notice how it is crucial for our method to deal with multiagent *dynamics*. Only with a dynamical system, in fact, can the extra layer introduced by the output mask decay and disappear over time, allowing convergence to the true state. In "static" contexts such as database query, where problems like privacy were originally formulated, our method is unlikely to provide any benefit.

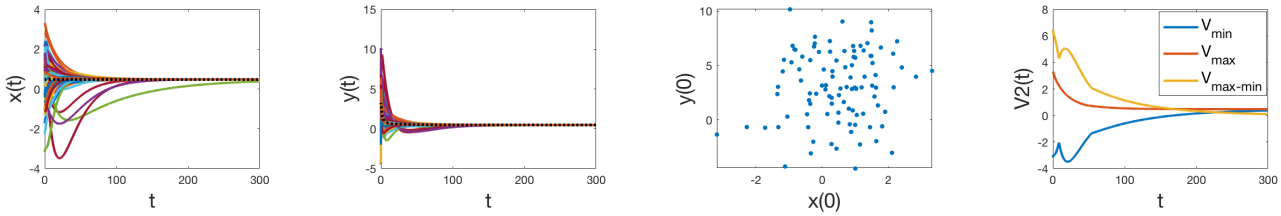

(a) (b) (c) (d)

Fig. 1. Privacy-preserving consensus of Example 1. (a): $x(t)$; (b): $y(t)$; (c): $x(0)$ vs. $y(0)$; (d): $V_{mm}(t) = \max_i(x_i(t)) - \min_i(x_i(t))$. The black dotted line in (a) resp. (b) represent $\mathbf{1}^T x(t)/n$, resp. $\mathbf{1}^T y(t)/n$.

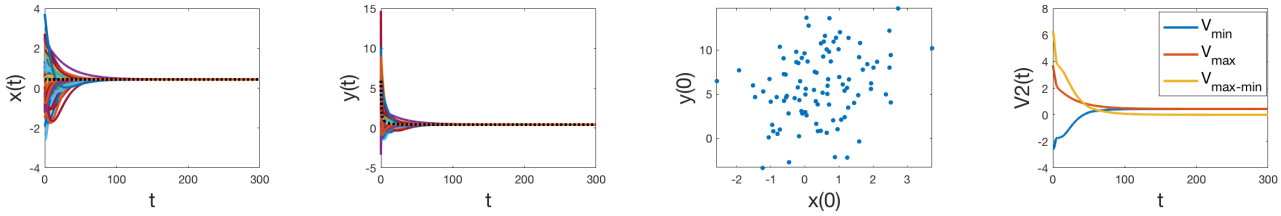

(a) (b) (c) (d)

Fig. 2. Another case of the privacy-preserving consensus of Example 1. (a): $x(t)$; (b): $y(t)$; (c): $x(0)$ vs. $y(0)$; (d): $V_{mm}(t) = \max_i(x_i(t)) - \min_i(x_i(t))$. The black dotted line in (a) resp. (b) represent $\mathbf{1}^T x(t)/n$, resp. $\mathbf{1}^T y(t)/n$.

REFERENCES

[1] A. Alaeddini, K. Morgansen, and M. Mesbahi. Adaptive communication networks with privacy guarantees. In *2017 American Control Conference (ACC)*, pages 4460–4465, May 2017.

[2] C. Altafini. A system-theoretic framework for privacy preservation in multiagent dynamics. *arXiv preprint arXiv:1904.11246*.

[3] M. Ambrosin, P. Braca, M. Conti, and R. Lazzeretti. Odin: Obfuscation-based privacy-preserving consensus algorithm for decentralized information fusion in smart device networks. *ACM Trans. Internet Technol.*, 18(1):6:1–6:22, Oct. 2017.

[4] Z. Artstein. Limiting equations and stability of nonautonomous ordinary differential equations. In J. LaSalle, editor, *The stability of dynamical systems*, CBMS Regional Conference Series in Applied Mathematics. SIAM, Philadelphia, 1976.

[5] J. Cortés, G. E. Dullerud, S. Han, J. L. Ny, S. Mitra, and G. J. Pappas. Differential privacy in control and network systems. In *IEEE 55th Conf. on Decision and Control*, pages 4252–4272, Dec 2016.

[6] C. Dwork. Differential privacy. In *Proceedings of the 33rd International Conference on Automata, Languages and Programming - Volume Part II*, ICALP'06, pages 1–12, Berlin, Heidelberg, 2006. Springer-Verlag.

[7] C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4):211–407, Aug. 2014.

[8] N. Gupta, J. Katz, and N. Chopra. Privacy in distributed average consensus. *IFAC-PapersOnLine*, 50(1):9515 – 9520, 2017. 20th IFAC World Congress.

[9] Z. Huang, S. Mitra, and G. Dullerud. Differentially private iterative synchronous consensus. In *Proceedings of the 2012 ACM Workshop on Privacy in the Electronic Society*, WPES '12, pages 81–90, New York, NY, USA, 2012. ACM.

[10] R. Lazzeretti, S. Horn, P. Braca, and P. Willett. Secure multi-party consensus gossip algorithms. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7406–7410, May 2014.

[11] T.-C. Lee, D.-C. Liaw, and B.-S. Chen. A general invariance principle for nonlinear time-varying systems and its applications. *IEEE Transactions on Automatic Control*, 46(12):1989–1993, Dec 2001.

[12] Y. Liu, J. Wu, I. R. Manchester, and G. Shi. Dynamical Privacy in Distributed Computing – Part I: Privacy Loss and PPSC Mechanism. *arXiv e-prints*, page arXiv:1902.06966, Feb 2019.

[13] N. E. Manitara and C. N. Hadjicostis. Privacy-preserving asymptotic average consensus. In *2013 European Control Conference (ECC)*, pages 760–765, July 2013.

[14] L. Markus. Asymptotically autonomous differential systems. In S. Lefschetz, editor, *Contribution to the theory of nonlinear oscillations*, Annals of Mathematical Studies. Princeton Univ. Press, Princeton, 1956.

[15] Y. Mo and R. M. Murray. Privacy preserving average consensus. *IEEE Transactions on Automatic Control*, 62(2):753–765, Feb 2017.

[16] X. Mu and D. Cheng. On the stability and stabilization of time-varying nonlinear control systems. *Asian Journal of Control*, 7(3):244–255, 2005.

[17] E. Nozari, P. Tallapragada, and J. Cortés. Differentially private average consensus: Obstructions, trade-offs, and optimal algorithm design. *Automatica*, 81:221 – 231, 2017.

[18] R. Olfati-Saber and R. Murray. Consensus problems in networks of agents with switching topology and time-delays. *Automatic Control, IEEE Transactions on*, 49(9):1520 – 1533, sept. 2004.

[19] S. Pequito, S. Kar, S. Sundaram, and A. P. Aguiar. Design of communication networks for distributed computation with privacy guarantees. In *53rd IEEE Conference on Decision and Control*, pages 1370–1376, Dec 2014.

[20] N. Rezazadeh and S. Kia. Privacy preservation in a continuous-time static average consensus algorithm over directed graphs. In *American Control Conference*, pages 5890–5895, 06 2018.

[21] N. Rouche, P. Habets, and M. Laloy. *Stability Theory by Liapunov's Direct Method*. Applied Mathematical Sciences. Springer New York, 2012.

[22] M. Ruan and Y. Wang. Secure and privacy-preserving average consensus. *CoRR*, abs/1703.09364, 2017.

[23] A. Saberi, P. Kokotovic, and H. Sussmann. Global stabilization of partially linear composite systems. *SIAM Journal on Control and Optimization*, 28(6):1491–1503, 1990.