SCIENCE AND TECHNOLOGY ORGANIZATION
CENTRE FOR MARITIME RESEARCH AND EXPERIMENTATION

Reprint Series **CMRE-PR-2019-079**

Learning with privacy in consensus + obfuscation

Paolo Braca, Riccardo Lazzeretti, Stefano Marano,
Vincenzo Matta

June 2019

Originally published in:

IEEE Signal Processing Letters, volume 23, issue 9, September 2016,
pp. 1174-1178, doi: 10.1109/LSP.2016.2587327

About CMRE

The Centre for Maritime Research and Experimentation (CMRE) is a world-class NATO scientific research and experimentation facility located in La Spezia, Italy.

The CMRE was established by the North Atlantic Council on 1 July 2012 as part of the NATO Science & Technology Organization. The CMRE and its predecessors have served NATO for over 50 years as the SACLANT Anti-Submarine Warfare Centre, SACLANT Undersea Research Centre, NATO Undersea Research Centre (NURC) and now as part of the Science & Technology Organization.

CMRE conducts state-of-the-art scientific research and experimentation ranging from concept development to prototype demonstration in an operational environment and has produced leaders in ocean science, modelling and simulation, acoustics and other disciplines, as well as producing critical results and understanding that have been built into the operational concepts of NATO and the nations.

CMRE conducts hands-on scientific and engineering research for the direct benefit of its NATO Customers. It operates two research vessels that enable science and technology solutions to be explored and exploited at sea. The largest of these vessels, the NRV Alliance, is a global class vessel that is acoustically extremely quiet.

CMRE is a leading example of enabling nations to work more effectively and efficiently together by prioritizing national needs, focusing on research and technology challenges, both in and out of the maritime environment, through the collective Power of its world-class scientists, engineers, and specialized laboratories in collaboration with the many partners in and out of the scientific domain.

Copyright © IEEE, 2016. NATO member nations have unlimited rights to use, modify, reproduce, release, perform, display or disclose these materials, and to authorize others to do so for government purposes. Any reproductions marked with this legend must also reproduce these markings. All other rights and uses except those permitted by copyright law are reserved by the copyright owner.

NOTE: The CMRE Reprint series reprints papers and articles published by CMRE authors in the open literature as an effort to widely disseminate CMRE products. Users are encouraged to cite the original article where possible.

Learning With Privacy in Consensus + Obfuscation

Paolo Braca, Riccardo Lazzeretti, Stefano Marano, and Vincenzo Matta

***Abstract*—We examine the interplay between learning and privacy over multiagent consensus networks. The learning objective of each individual agent consists of computing some global network statistic, and is accomplished by means of a consensus protocol. The privacy objective consists of preventing inference of the individual agents' data from the information exchanged during the consensus stages, and is accomplished by adding some artificial noise to the observations (obfuscation). An analytical characterization of the learning and privacy performance is provided, with reference to a consensus perturbing and to a consensus-preserving obfuscation strategy.**

***Index Terms*—Consensus, Multi-agent systems, obfuscation, privacy.**

I. Introduction

THANKS to the ongoing progression of distributed processing, as well as to the increasing availability of shared resources, the computation of functions from dispersed pieces of information is becoming an essential building block of many modern information systems. In such a setting, a prominent role is held by *consensus* algorithms, which are employed in multi-agent networks to enable the distributed computation of linear functions of the agents' data, by means of cooperation between neighboring network nodes [1]–[10].

The distributed computation paradigm implies unique challenges in terms of privacy, so as to make the traditional protection techniques helpless [11]. Indeed, while the learning objective of each individual agent consists of computing some *global* network statistic, such an objective is accomplished by propagating across the network the *local* data owned by the agents, which one might want to keep private. Therefore, the privacy objective consists of preventing inference of the individual agents' data from the information exchanged during the consensus stages. In a nutshell, one would like to compute a (global) function of the (local, private) agents' data, without letting each agent make inference about the agents' data themselves. The correct management of these conflicting requirements is a crucial task in several applications, which include medical diagnosis, multimedia services, social networks, public safety, and defense.

In the realm of distributed consensus, some privacy-preserving approaches have been recently proposed [13]–[18]. These approaches mainly exploit the so-called *obfuscation* strategies, which are common practice in the well-established frameworks of signal processing in the encrypted domain (SPED), multiparty computation [12], and differential privacy [19]. At one extreme, independent (artificial) noise samples are added to the individual agents' data, and the consensus algorithm operates on such obfuscated dataset [14]–[16]. As a result, the noise increases the privacy and reduces the fidelity in the computation of the global function. At another extreme, one might introduce some dependence, in time (across subsequent transmissions) and/or in space (e.g., across agents) among the noise samples, in order to let the arithmetic average of the noise samples vanish [17], [18]. In this way, the final consensus is preserved (over an infinite time horizon), while the introduced dependence might reduce to some extent the achievable privacy.

In order to capture such a tradeoff, it is crucial to observe that the consensus algorithm cannot run forever, and must be *stopped* at a certain time. Even when the obfuscation noise averages to zero, the finite time-horizon implies a residual error in estimating the arithmetic average of the agents' data.

In this letter, we provide an analytical characterization of the interplay between learning and privacy for a consensus-perturbing (see Section III) and for a consensus-preserving (see Section IV) strategy. This characterization is summarized by the privacy/learning curves in (10), (12), and (15).

II. Problem Formulation

We consider a network of N agents that collect observations about a certain phenomenon. The ℓth agent's observation is denoted by θ_ℓ, and $\boldsymbol{\theta} = [\theta_1, \theta_2, \ldots, \theta_N]^T$ is the ensemble of observations, which is modeled as a *random* vector, so as to include the presence of *noisy* measurements. The goal of *each* agent is to learn the arithmetic mean of the observations, $\mu \triangleq \bar{\theta}$ (the notation $\bar{v}$ will denote the arithmetic average of the entries in vector v), in a fully decentralized fashion, leveraging local data exchange with its own neighbors. When such data are transmitted uncovered, the privacy of the agent's datum is clearly not guaranteed. One commonly adopted strategy to add some privacy is the *obfuscation* strategy, where the observations are masked by *artificially* adding some noise ω [14]–[18]. The obfuscated dataset will be accordingly $x = \theta + \omega$.

The network agents implement an averaging consensus protocol. Denoting by $s(\tau) = [s_1(\tau), s_2(\tau), \ldots, s_N(\tau)]^T$ the *state* vector of the network at step $\tau = 0, 1, \ldots$, such a protocol is [1]–[3]

$$s(\tau) = \boldsymbol{W}(\tau)s(\tau-1),\ s(0) = x,\ s_\ell(\tau) \overset{\tau\to\infty}{\longrightarrow} \bar{x} = \mu + \bar{\omega} \tag{1}$$

where $\boldsymbol{W}(\tau)$ is the (random) weighting matrix at time τ, and the convergence holds (in several different senses, see [1]–[3]) under suitable conditions on the network topology and weighting policy. Equation (1) reveals that, while convergence to the desired term μ is obfuscated by the noisy term $\bar{\omega}$, the consensus

Manuscript received May 03, 2016; revised June 27, 2016; accepted June 30, 2016. Date of publication July 07, 2016; date of current version July 20, 2016. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Steeve Zozor.

P. Braca is with the NATO STO CMRE, La Spezia, Italy (e-mail: paolo.braca@cmre.nato.int).

R. Lazzeretti is with the University of Siena, 55 Siena, Italy, (e-mail: lazzeretti2@unisi.it).

S. Marano and V. Matta are with the University of Salerno, 84084 Fisciano, Italy (e-mail: marano@unisa.it; vmatta@unisa.it).

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

Digital Object Identifier 10.1109/LSP.2016.2587327

algorithm can nevertheless be used to provide an estimate of μ. The estimate produced by the ℓth agent at step τ will be denoted by $\hat{\mu}_\ell(\tau)$.

Since we are interested in the tradeoff between learning (the ability of each agent to infer μ) and privacy (the protection of the individual agents' data), let us introduce the pertinent performance indicators. To this aim, it is necessary to identify a proper optimality criterion. In complying with the linear nature of consensus, in this study, we choose the linear minimum mean square error (LMMSE) criterion. The LMMSE in estimating a variable ξ based on a dataset $\boldsymbol{d}$ is denoted by $\mathsf{Lmmse}(\xi|\boldsymbol{d})$, and the corresponding optimal estimator by $\hat{\xi}^\star$. The learning ability of the ℓth agent at step τ will be measured in terms of the average squared distance between the true value, μ, and its estimate, $\hat{\mu}_\ell(\tau)$. We accordingly introduce as learning index the error ratio

$$\mathscr{E}_\ell \triangleq \frac{\mathbb{E}[(\hat{\mu}_\ell(\tau)-\mu)^2]}{\mathsf{Lmmse}(\mu|\emptyset)} = \frac{\mathbb{E}[(\hat{\mu}_\ell(\tau)-\hat{\mu}^\star)^2]+\mathsf{Lmmse}(\mu|\boldsymbol{x})}{\mathrm{VAR}[\mu]} \tag{2}$$

where: 1) the error is scaled to the prior error (LMMSE with no data); and 2) the second equality follows by the orthogonality principle for LMMSE estimators. A high learning ability corresponds to a small $\mathscr{E}_\ell$. In particular, the best achievable error ratio for a given obfuscation policy is $\mathsf{Lmmse}(\mu|\boldsymbol{x})/\mathrm{VAR}[\mu] \leq 1$, which will be achieved if the estimate $\mu_\ell(\tau)$ is able to reproduce the LMMSE estimator $\hat{\mu}^\star$ as $\tau \to \infty$. Clearly, a *meaningful* learning regime requires that the achievable mean square error is smaller than the prior error, implying $\mathscr{E}_\ell < 1$.

Let us switch to the privacy indicator. Consider the case that the observation θ_ℓ of the ℓth individual agent must be estimated by the kth agent. We consider the worst-case scenario that the estimate must be produced based upon the whole obfuscated dataset $\boldsymbol{x}$. Clearly, the kth agent has at its own disposal also the uncovered observation θ_k, as well as the obfuscation noise sample ω_k. We measure the privacy of an individual agent with the same indicator, the MSE, used to measure the global *learning* objective. While there are many alternative, well-established notions of privacy, (e.g., privacy based on equivocation, differential privacy), such a symmetric choice has gained increasing attention [18], [20]. The privacy indicator of the ℓth agent is accordingly defined by the following LMMSE ratio:

$$\mathscr{P}_{\ell,k} \triangleq \frac{\mathsf{Lmmse}(\theta_\ell|\{\boldsymbol{x}_{-k},\theta_k,\omega_k\})}{\mathrm{VAR}[\theta_l]} \tag{3}$$

where $\boldsymbol{x}_{-k}$ denotes the dataset $\boldsymbol{x}$, deprived of the kth sample x_k. It is seen that $0 \leq \mathscr{P}_{\ell,k} \leq 1$. In particular, null privacy ($\mathscr{P}_{\ell,k}=0$) corresponds to a perfect estimate of the underlying observation (the numerator is zero), while perfect privacy ($\mathscr{P}_{\ell,k}=1$) corresponds to the case that the error in estimating θ_ℓ is the same as if no information were available. Finally, joining (2) and (3), a curve $\mathscr{P}^\star_{\ell,k}(\mathscr{E})$ is obtained, which will represent the privacy achieved for a given learning error $\mathscr{E}$.

Assumptions. The vector $\boldsymbol{\theta}$ has covariance matrix $\boldsymbol{C}_\theta = \boldsymbol{I}$, where $\boldsymbol{I}$ is the identity matrix. The noise-vector $\boldsymbol{\omega}$ is zero-mean, uncorrelated with $\boldsymbol{\theta}$, with covariance matrix $\boldsymbol{C}_\omega$. The consensus matrices $\boldsymbol{W}(\tau)$'s are independent and identically distributed realizations of doubly stochastic matrices. The second largest eigenvalue of the average matrix $\mathbb{E}[\boldsymbol{W}(\tau)\boldsymbol{W}(\tau)^T]$ is strictly less than one, a condition that correspond to a strongly connected network, and that will, therefore, guarantee the convergence of the consensus protocol [1]–[3].

III. Consensus-Perturbing Strategy

The first strategy is a simple jamming strategy where $\boldsymbol{C}_\omega \triangleq \mathbb{E}[\boldsymbol{\omega}\boldsymbol{\omega}^T] = \sigma^2\boldsymbol{I}$ namely, where the entries of the obfuscation noise are uncorrelated with (common) variance σ^2. We observe that, in general, we shall have $\bar{\omega} \neq 0$, since $\mathrm{VAR}[\bar{\omega}] = (1/N)\sigma^2 > 0$. Therefore, we see from (1) that this obfuscation strategy is *consensus perturbing*, since the true arithmetic average μ is not recovered as time elapses. The privacy indicator is immediately obtained, since the entries of $\boldsymbol{\theta}$, as well as the entries of $\boldsymbol{\omega}$, are uncorrelated, and so are $\boldsymbol{\theta}$ and $\boldsymbol{\omega}$. Therefore, only the (obfuscated) datum x_ℓ is used by the kth agent to estimate the ℓth agent's observation, θ_ℓ, which implies [22]: $\mathscr{P}_{\ell,k} = \mathsf{Lmmse}(\theta_\ell|x_\ell) = \sigma^2/(1+\sigma^2)$. Concerning the learning behavior, in the forthcoming sections, we consider the scenarios with and without prior information.

A. $\hat{\mu}_\ell(\tau)$ With Prior Information

Assume that the prior information useful to design LMMSE estimators (i.e., the second-order characterization of $\boldsymbol{\theta}$) is available to the agents. Calculating the LMMSE estimator of μ amounts to minimizing, over all the possible choices of the real coefficients $\alpha_1, \alpha_2, \ldots, \alpha_N$, the error $\mathbb{E}[(\sum_{\ell=1}^N \alpha_\ell x_\ell - \mu)^2]$. From the latter formula, it is straightforward to show that [22]

$$\hat{\mu}^\star = \frac{\bar{x}}{1+\sigma^2} \Rightarrow \mathsf{Lmmse}(\mu|\boldsymbol{x}) = \frac{\sigma^2/N}{1+\sigma^2}. \tag{4}$$

Therefore, the consensus estimator in (1) must be scaled as

$$\hat{\mu}_\ell(\tau) = \frac{1}{1+\sigma^2}\, s_\ell(\tau) \overset{\tau\to\infty}{\longrightarrow} \hat{\mu}^\star \tag{5}$$

yielding $\mathbb{E}[(\hat{\mu}_\ell(\tau)-\hat{\mu}^\star)^2] = (1+\sigma^2)^{-2}\,[\boldsymbol{C}_\epsilon(\tau)]_{\ell\ell}$, where $\boldsymbol{C}_\epsilon(\tau) = \mathbb{E}[\boldsymbol{\epsilon}(\tau)\boldsymbol{\epsilon}(\tau)^T]$ is the covariance matrix of the consensus error-vector $\boldsymbol{\epsilon}(\tau) = \boldsymbol{s}(\tau) - \mathbb{1}\,\bar{x}$, with $\mathbb{1}$ being the $N\times 1$ vector with all entries equal to one. Introducing now the matrix $\boldsymbol{A}(\tau) = \prod_{j=\tau}^1 \boldsymbol{W}(j)$, from (1), we can write

$$\boldsymbol{s}(\tau) = \boldsymbol{A}(\tau)\,\boldsymbol{x} \Rightarrow \boldsymbol{\epsilon}(\tau) = \boldsymbol{A}(\tau)\left(\boldsymbol{I} - \frac{\mathbb{1}\mathbb{1}^T}{N}\right)\boldsymbol{x} \tag{6}$$

where the latter equality follows because $\boldsymbol{A}(\tau)\mathbb{1} = \mathbb{1}$. Therefore, noting that $\mathbb{E}[\boldsymbol{x}\boldsymbol{x}^T] = (1+\sigma^2)\boldsymbol{I}$, the consensus-error covariance is readily evaluated as

$$\boldsymbol{C}_\epsilon(\tau) = (1+\sigma^2)\underbrace{\mathbb{E}\left[\boldsymbol{A}(\tau)\boldsymbol{A}(\tau)^T - \frac{\mathbb{1}\mathbb{1}^T}{N}\right]}_{\boldsymbol{\Phi}(\tau)} \tag{7}$$

which yields $\mathbb{E}[(\hat{\mu}_\ell(\tau)-\hat{\mu}^\star)^2] = (1+\sigma^2)^{-1}\,\Phi_{\ell\ell}(\tau)$. Using the latter result and (4) in (2), and recalling the result obtained for the privacy indicator, we get the following summary for the consensus-perturbing strategy with prior information:

$$\mathscr{E}_\ell = \frac{N\Phi_{\ell\ell}(\tau)+\sigma^2}{1+\sigma^2}, \quad \mathscr{P}_{\ell,k} = \frac{\sigma^2}{1+\sigma^2}. \tag{8}$$

We see that the privacy indicator is an increasing function of σ^2, that is zero in the absence of obfuscation, and that tends to its maximum value of unity when the obfuscation power increases without bound. With regard to the inference performance, we see that $\mathscr{E}_\ell \to \sigma^2/(1+\sigma^2)$ as $\tau \to \infty$. In view of (4), the limiting error represents the performance of the LMMSE estimator, which is expected since the estimator $\hat{\mu}_\ell(\tau)$ has been just chosen so as to converge to the LMMSE. We note also that, for finite τ, even in the absence of obfuscation, an error ratio $N\Phi_{\ell\ell}(\tau)$ is unavoidable. On the other hand, we know that an error ratio greater than one makes little sense in terms of learning. We conclude that: 1) the learning time, τ, must ensure that $N\Phi_{\ell\ell}(\tau)$ is smaller than one and 2) the obfuscation power must ensure that $\mathscr{E}_\ell < 1$, yielding

$$\boxed{N\Phi_{\ell\ell}(\tau) < \mathscr{E}_\ell < 1.} \tag{9}$$

Let us finally summarize the learning/privacy tradeoff by computing $\mathscr{P}^\star_{\ell,k}(\mathscr{E})$. Enforcing $\mathscr{E}_\ell = \mathscr{E}$ into the first relationship in (8) implies $\sigma^2 = (\mathscr{E} - N\Phi_{\ell\ell}(\tau))/(1-\mathscr{E})$. Inserting such a value into the second relationship in (8) yields

$$\boxed{\mathscr{P}^\star_{\ell,k}(\mathscr{E}) = \frac{\mathscr{E} - N\Phi_{\ell\ell}(\tau)}{1 - N\Phi_{\ell\ell}(\tau)}.} \tag{10}$$

B. $\hat{\mu}_\ell(\tau)$ *Without Prior Information*

As already stated, when no prior information is available, we set $\hat{\mu}_\ell(\tau) = s_\ell(\tau)$. Reasoning as in the previous section

$$\boxed{\mathscr{E}_\ell = (1+\sigma^2)N\Phi_{\ell\ell}(\tau) + \sigma^2, \quad \mathscr{P}_{\ell,k} = \frac{\sigma^2}{1+\sigma^2}.} \tag{11}$$

The comments about the latter two formulas are similar to those made for the scheme with prior information, but for an important aspect. Considering the error ratio, we here see that $\mathscr{E}_\ell \to \sigma^2$ as $\tau \to \infty$. This means that, as the obfuscation power increases, the inference error increases without bound. Such a behavior matches perfectly the *universal* nature of the estimator adopted in the scenario without prior information. Indeed, with extremely noisy data (large σ^2), the prior information would become dominant with respect to the information contained in the data. Thus, when prior information is available, the worst-case inference performance can be achieved ($\mathscr{E}_\ell \to 1$). In contrast, when it is not, the presence of extremely noisy observations leads to an *unbounded* error. From (11), we can also obtain the privacy/learning curve

$$\boxed{\mathscr{P}^\star_{\ell,k}(\mathscr{E}) = \frac{\mathscr{E} - N\Phi_{\ell\ell}(\tau)}{1 + \mathscr{E}}.} \tag{12}$$

IV. Consensus-Preserving Strategy

The consensus-preserving strategy proposed in this letter employs a *a single* noise sample per each agent. Thus, the dependence required to preserve consensus implies some coordination among the agents, which can be guaranteed in several ways, depending on the particular application. For instance, in a sensor network, sensors can be equipped with their obfuscation noise at the factory stage, or periodically, by a system manager; in multiparty computation, the obfuscation noise can be generated by a trusted third party. Similar operations could be performed by means of a key-distribution algorithm during an initialization stage, see, e.g., [23].

Consider a random vector $\boldsymbol{z}$, and set: $\mathbb{E}[\boldsymbol{z}] = 0$, $\mathbb{E}[\boldsymbol{z}\boldsymbol{z}^T] = \sigma^2\boldsymbol{I}$, $\boldsymbol{\omega} = \boldsymbol{z} - \bar{\boldsymbol{z}}$. Basically, the obfuscation noise is generated, from the auxiliary vector $\boldsymbol{z}$, by subtracting from $\boldsymbol{z}$ its arithmetic average. As a result, the arithmetic average $\bar{\omega}$ of the obfuscation-noise vector is forced to be zero, which allows preservation of consensus. The correlation properties of $\boldsymbol{z}$ reveal that the noise covariance is $\boldsymbol{C}_\omega = \sigma^2(\boldsymbol{I} - \frac{\mathbb{1}\mathbb{1}^T}{N})$. Moreover, from the condition $\bar{\omega} = 0$, we reach immediately the important conclusion: $\hat{\mu}^\star = \bar{x} = \mu \Rightarrow \mathsf{Lmmse}(\mu|\boldsymbol{x}) = 0$, which in turn motivates the choice $\hat{\mu}_\ell(\tau) = s_\ell(\tau)$. Accordingly, we introduce the error $\boldsymbol{\epsilon}(\tau) = \boldsymbol{s}(\tau) - \mathbb{1}\mu$, and using the noise covariance $\boldsymbol{C}_\omega$, we get $\boldsymbol{C}_\epsilon(\tau) = (1+\sigma^2)\boldsymbol{\Phi}(\tau)$, yielding

$$\mathscr{E}_\ell = \frac{\mathbb{E}[(\hat{\mu}_\ell(\tau) - \mu)^2]}{\mathsf{VAR}[\mu]} = \frac{[\boldsymbol{C}_\epsilon(\tau)]_{\ell\ell}}{1/N} = (1+\sigma^2)N\Phi_{\ell\ell}(\tau). \tag{13}$$

Let us move on examining the privacy indicator. To this aim, it is expedient to address first the case that one has to estimate $\boldsymbol{\theta}$ based upon the whole dataset $\boldsymbol{x}$, *without knowing any of the local observations contained in* $\boldsymbol{\theta}$. For this particular case, it is known that the covariance matrix of the LMMSE error is: $\mathbb{E}\left[(\hat{\boldsymbol{\theta}}^\star - \boldsymbol{\theta})(\hat{\boldsymbol{\theta}}^\star - \boldsymbol{\theta})^T\right] = \boldsymbol{C}_\theta - \boldsymbol{C}_\theta(\boldsymbol{C}_\theta + \boldsymbol{C}_\omega)^{-1}\boldsymbol{C}_\theta = \boldsymbol{I} - (\boldsymbol{I} + \boldsymbol{C}_\omega)^{-1} = \sigma^2/(1+\sigma^2)(\boldsymbol{I} - \mathbb{1}\mathbb{1}^T/N)$, where the matrix inversion has been performed by using the Sherman–Morrison–Woodbury identity [21]. As a result, for this case, we can write: $\mathbb{E}[(\hat{\theta}^\star_\ell - \theta_\ell)^2] = \sigma^2/(1+\sigma^2)(1-1/N)$. Let us come back to the evaluation of the LMMSE errors relevant to the privacy indicator in (3). Without loss of generality, assume that agent 1 must estimate the observations of the remaining agents, and accordingly, its dataset is: $\boldsymbol{y} = [\omega_1, x_2, \ldots, x_N]^T$, where θ_1 has been excluded since it is uncorrelated from θ_ℓ and x_ℓ, for $\ell \neq 1$. It is convenient to introduce the alternative dataset $\boldsymbol{y}' = \left[\omega_1, x_2 + \frac{\omega_1}{N-1}, \ldots, x_N + \frac{\omega_1}{N-1}\right]^T$. Clearly, any linear estimator that can be constructed from $\boldsymbol{y}$ can be constructed also from $\boldsymbol{y}'$. On the other hand, exploiting the form of $\boldsymbol{C}_\omega$ for all $k \neq 1$ we have: $\mathbb{E}\left[\omega_1\left(\omega_k + \frac{\omega_1}{N-1}\right)\right] = \sigma^2(-1/N + 1/N) = 0$, which implies (since the obfuscation is independent of $\boldsymbol{\theta}$) that the $\boldsymbol{y}'$ can be expurgated of ω_1, obtaining the equivalent dataset $\boldsymbol{y}'' = \left[x_2 + \frac{\omega_1}{N-1}, \ldots, x_N + \frac{\omega_1}{N-1}\right]^T$. Now, note that, for $k \neq 1$: $\mathbb{E}\left[(\omega_k + \frac{\omega_1}{N-1})^2\right] = 1 - 1/(N-1)$, while for $\ell \neq k$ and $\ell, k \neq 1$ $\mathbb{E}\left[(\omega_\ell + \frac{\omega_1}{N-1})(\omega_k + \frac{\omega_1}{N-1})\right] = -1/(N-1)$. The latter two equations reveal that the problem of estimating $\theta_2, \theta_3, \ldots, \theta_N$ from the dataset $\boldsymbol{y}''$ has, in terms of covariance, the same structure as the problem of estimating $\boldsymbol{\theta}$ from the dataset $\boldsymbol{x}$, but for the fact that the dimension of the problem is $N-1$ instead of N. The privacy indicator can be therefore evaluated as explained before, finally yielding

$$\boxed{\mathscr{E}_\ell = (1+\sigma^2)N\Phi_{\ell\ell}(\tau), \quad \mathscr{P}_{\ell,k} = \frac{\sigma^2}{1+\sigma^2}\left(1 - \frac{1}{N-1}\right).} \tag{14}$$

With regard to the learning index, different from the consensus-perturbing strategy, $\mathscr{E}_\ell \to 0$ as $\tau \to \infty$. This matches perfectly our intuition, since the consensus-preserving strategy must allow the estimator $\hat{\mu}_\ell(\tau)$ to reach the *true* arithmetic average of the agents' observations as time elapses. The pri-

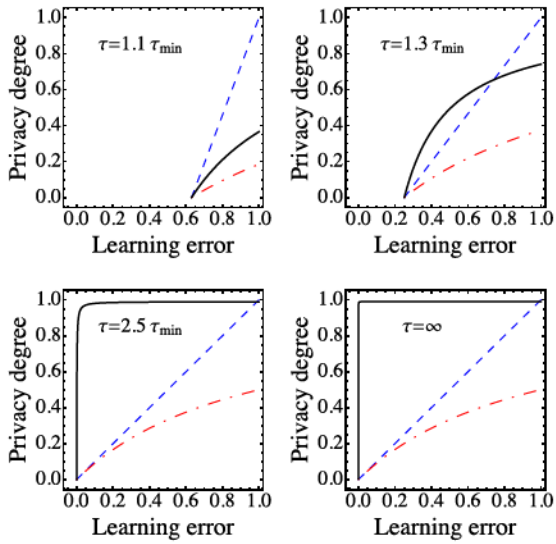

Fig. 1. Curve $\mathscr{P}^{\star}_{\ell,k}(\mathscr{E})$. Different panels correspond to different τ. Each panel displays the performance of the three strategies examined in the work: consensus-perturbing strategy with [(10), dashed] and without [(12), dotted-dashed] prior information, and the consensus-preserving strategy [(15), solid]. The relevant parameters are $\lambda = 0.9$ and $N = 100$.

vacy index increases monotonically with the obfuscation power. However, the limiting privacy value is now less than unity, and is equal to $1 - 1/(N-1)$. Interestingly, this limiting value corresponds to the LMMSE obtained when the kth agent tries to estimate the ℓth agent's datum θ_ℓ by observing, along with its own data θ_k and ω_k, the arithmetic average of the observations. Such a result makes perfect sense because irrespectively of the (large) value of the obfuscation-noise power, from the obfuscated vector x one can always compute $\bar{x} = \bar{\theta} + \bar{w} = \bar{\theta}$, namely, the arithmetic average is always accessible since it represents the information that we want to preserve through the proposed strategy. Finally, joining the two equations appearing in (14), we get

$$\boxed{\mathscr{P}^{\star}_{\ell,k}(\mathscr{E}) = \left(1 - \frac{1}{N-1}\right) \frac{\mathscr{E} - N\Phi_{\ell\ell}(\tau)}{\mathscr{E}}.} \qquad (15)$$

V. Visualizing the Learning / Privacy Tradeoff

In Fig. 1, we display the privacy/learning curve for the three strategies addressed in this study. Each panel corresponds to a different number of consensus steps, τ, which increases from left-to-right, top-to-bottom. In order to consider a meaningful learning regime, we focused on the range in (9). We observe from (10), (12), and (15) that the functional form of $\mathscr{P}^{\star}_{\ell,k}(\mathscr{E})$ will be independent of the specific value chosen for the term $N\Phi_{\ell\ell}(\tau)$. Accordingly, the privacy/learning curves could be in principle described by setting a certain value $N\Phi_{\ell\ell}(\tau) < 1$, without worrying about the physical meaning of such a term in connection to the underlying network structure. On the other hand, capturing the behavior of $\Phi_{\ell\ell}(\tau)$ as regards its dependence on the number of consensus steps is important to capturing the dynamic evolution of the system performance. A simple though revealing way to capture the consensus dynamics is offered by the known upper bound $\Phi_{\ell\ell}(\tau) \leq \lambda^{\tau}$, where λ is the second largest eigenvalue of the matrix $\mathbb{E}[\boldsymbol{W}(\tau)\boldsymbol{W}(\tau)^T]$, which is strictly less than one by assumption [1]–[3]. Accordingly, we replace $\Phi_{\ell\ell}(\tau)$ with the upper bound λ^{τ} to highlight the essence of the privacy learning tradeoff. Using this bound, we can now set the *minimum* number of consensus steps so as to cope with (9). Enforcing the condition $N\lambda^{\tau_{\min}} = 1$ we get, but for round-off errors, $\tau_{\min} = -\ln N / \ln \lambda$. We are now ready to examine in detail Fig. 1. By joint inspection of the four panels, we see that, in agreement with (9), all curves move leftward as τ increases. In particular, we observe that the consensus-perturbing strategy with prior information is always a line that joins the minimum-error/minimum-privacy point, $(N\Phi_{\ell\ell}(\tau), 0)$, with the maximum-error/maximum-privacy point, $(1, 1)$. As another general feature, we see in all panels that the consensus-perturbing strategy without prior information performs uniformly (across the error axis) worse than the other strategies. In fact, from (12) and (15), we see that consensus-perturbation with no prior outperforms consensus-preservation only if $\mathscr{E} > N - 2$. This inequality is verified only for the trivial cases $N = 1$ and $N = 2$, since we focus on the regime $\mathscr{E} < 1$.

Let us start by examining the leftmost-and-uppermost panel in Fig. 1, where the number of consensus steps, $\tau = 1.1\,\tau_{\min}$, is very close to the minimum value $\tau_{\min}$, namely, to the obfuscation-without-learning zone. We see that the consensus-perturbing strategy with prior information outperforms the consensus-preserving strategy. Such a behavior will be observed as long as the derivative of the consensus-perturbing curve at $\mathscr{E} = N\Phi_{\ell\ell}(\tau)$ is greater than the corresponding derivative for the consensus-preserving strategy. Using (10) and (15), simple algebra reveals that the latter situation corresponds to: $N\Phi_{\ell\ell}(\tau) \geq (N-2)/(2N-3)$, a condition that identifies the maximum learning time for which the consensus-perturbing strategy is always superior. Moving on to the rightmost-and-uppermost panel, the consensus-preserving strategy initially outperforms the consensus-perturbing strategy with prior information. For a certain error, the two pertinent curves cross each other, and the situation is reversed. As τ increases (lowermost panels), the intersection moves rightward. As $\tau \to \infty$ (rightmost-and-lowermost panel), the consensus-preserving strategy exhibits a uniform privacy equal to $1 - 1/(N-1)$, while the consensus-perturbing strategy with prior information is a line joining the points 0 and 1.

VI. Summary

We have provided an analytical characterization of the learning/privacy tradeoff over distributed networks, for consensus-perturbing as opposed to consensus-preserving algorithms. Our analysis revealed that consensus preservation is not necessarily beneficial, since its possible advantages depend on the time available for learning. In the accurate learning regime, a sharp difference is observed between the aforementioned approaches: the consensus-perturbing strategy exhibits *zero privacy*, while the consensus-preserving strategy achieves a privacy value close (with a slight loss) to the maximum theoretical value of unit privacy. The slight loss is ascribed to the fact that, irrespectively of the obfuscation power, the information preserved to consensus purposes allows making some inference about the individual agents' observations.

REFERENCES

[1] L. Xiao and S. Boyd, "Fast linear iterations for distributed averaging," *Syst. Control Lett.*, vol. 53, no. 1, pp. 65–78, Sep. 2004.

[2] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized gossip algorithms," *IEEE Trans. Inf. Theory*, vol. 52, no. 6, pp. 2508–2530, Jun. 2006.

[3] R. Olfati-Saber, A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proc. IEEE*, vol. 95, no. 1, pp. 215–233, Jan. 2007.

[4] A. G. Dimakis, S. Kar, J. M. F. Moura, M. G. Rabbat, and A. Scaglione, "Gossip algorithms for distributed signal processing," *Proc. IEEE*, vol. 98, no. 11, pp. 1847–1864, Nov. 2010.

[5] P. Braca, S. Marano, and V. Matta, "Enforcing consensus while monitoring the environment in wireless sensor networks," *IEEE Trans. Signal Process.*, vol. 56, no. 7, pp. 3375–3380, Jul. 2008.

[6] P. Braca, S. Marano, V. Matta, and P. Willett, "Asymptotic optimality of running consensus in testing statistical hypotheses," *IEEE Trans. Signal Process.*, vol. 58, no. 2, pp. 814–825, Feb. 2010.

[7] P. Braca, S. Marano, V. Matta, and P. Willett, "Consensus-based Page's test in sensor networks," *Signal Process.*, vol. 91, no. 4, pp. 919–930, Apr. 2011.

[8] D. Bajovic, D. Jakovetic, J. Xavier, B. Sinopoli, and J. M. F. Moura, "Distributed detection via Gaussian running consensus: Large deviations asymptotic analysis," *IEEE Trans. Signal Process.*, vol. 59, no. 9, pp. 4381–4396, Sep. 2011.

[9] D. Bajovic, D. Jakovetic, J. M. F. Moura, J. Xavier, and B. Sinopoli, "Large deviations performance of consensus + innovations distributed detection with non-Gaussian observations," *IEEE Trans. Signal Process.*, vol. 60, no. 11, pp. 5987–6002, Nov. 2012.

[10] S. Kar and J. M. F. Moura, "Convergence rate analysis of distributed gossip (linear parameter) estimation: Fundamental limits and tradeoffs," *IEEE J. Sel. Topics Signal Process.*, vol. 5, no. 4, pp. 674–690, Aug. 2011.

[11] A. Perrig, J. Stankovic, and D. Wagner, "Security in wireless sensor networks," *Commun. ACM*, vol. 47, no. 6, pp. 53–57, 2004.

[12] R. Lagendijk, Z. Erkin, and M. Barni, "Encrypted signal processing for privacy protection: Conveying the utility of homomorphic encryption and multiparty computation," *IEEE Signal Process. Mag.*, vol. 30, no. 1, pp. 82–105, 2013.

[13] R. Lazzeretti, S. Horn, P. Braca, and P. Willett, "Secure multiparty consensus gossip algorithms," in *Proc. IEEE Int. Conf. Acoustics, Speech Signal Process.*, Florence, Italy, May 4–9, 2014, pp. 7406–7410.

[14] M. Kefayati, M. S. Talebi, B. H. Khalaj, and H. R. Rabiee, "Secure consensus averaging in sensor networks using random offsets," in *Proc. IEEE Int. Conf. Telecommun. Malaysia Int. Conf. Commun.*, Penang, Malaysia, May 14–17, 2007, pp. 556–560.

[15] Z. Huang, S. Mitra, and G. Dullerud, "Differentially private iterative synchronous consensus," in *Proc. ACM Workshop Privacy Electron. Soc.*, Raileigh, NC, USA, Oct. 15, 2012, pp. 81–90.

[16] E. Nozari, P. Tallapragada, and J. Cortés, "Differentially private average consensus with optimal noise selection," in *Proc. 5th IFAC Workshop Distrib. Estimation Control Networked Syst.*, Philadelphia, PA, USA, Sep. 10–11, 2015, vol. 48, no. 22, pp. 203–208.

[17] N. E. Manitara and C. N. Hadjicostis, "Privacy-preserving asymptotic average consensus," in *Proc. Eur. Control Conf.*, Zürich, Switzerland, Jul. 17–19, 2013, pp. 760–765.

[18] Y. Mo and R. Murray, "Privacy preserving average consensus," in *Proc. IEEE 53rd Annu. Conf. Decision Control*, Los Angeles, CA, USA, Dec. 15–17, 2014, pp. 2154–2159.

[19] C. Dwork, "Differential privacy," in *Automata, Languages and Programming*. Berlin, Germany: Springer, 2006, pp. 1–12.

[20] C. Schieler and P. Cuff, "Rate-Distortion Theory for Secrecy Systems," *IEEE Trans. Inf. Theory*, vol. 60, no. 12, pp. 7584–7605, Dec. 2014.

[21] R. Horn and C. Johnson, *Matrix Analysis*. Cambridge, U.K.: Cambridge Univ. Press, 1985.

[22] S. M. Kay, *Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory*. Englewood Cliffs, NJ, USA: Prentice Hall, 1993.

[23] S. Camtepe and B. Yener, "Key distribution mechanisms for wireless sensor networks: A survey," Rensselaer Polytechnic Inst., Troy, NY, USA, Tech. Rep. 05–07, 2005.

Document Data Sheet

Security Classification		Project No.
Document Serial No.	**Date of Issue**	**Total Pages**
CMRE-PR-2019-079	June 2019	5 pp.

Author(s)

Paolo Braca, Riccardo Lazzeretti, Stefano Marano, Vincenzo Matta

Title

Learning with privacy in consensus + obfuscation

Abstract

We examine the interplay between learning and privacy over multiagent consensus networks. The learning objective of each individual agent consists of computing some global network statistic, and is accomplished by means of a consensus protocol. The privacy objective consists of preventing inference of the individual agents' data from the information exchanged during the consensus stages, and is accomplished by adding some artificial noise to the observations (obfuscation). An analytical characterization of the learning and privacy performance is provided, with reference to a consensus perturbing and to a consensus-preserving obfuscation strategy.

Keywords

Consensus, multi-agent systems, obfuscation, privacy

Issuing Organization

NATO Science and Technology Organization
Centre for Maritime Research and Experimentation

Viale San Bartolomeo 400, 19126 La Spezia, Italy

[From N. America:
STO CMRE
Unit 31318, Box 19, APO AE 09613-1318]

Tel: +39 0187 527 361
Fax:+39 0187 527 700

E-mail: library@cmre.nato.int