IEEE Open Journal of Control Systems

Received 30 April 2024; revised 17 July 2024 and 21 August 2024; accepted 2 September 2024. Date of publication 6 September 2024; date of current version 14 October 2024. Recommended by Guest Editor Hideaki Ishii.

Digital Object Identifier 10.1109/OJCSYS.2024.3455899

A Control-Theoretical Zero-Knowledge Proof Scheme for Networked Control Systems

CAMILLA FIORAVANTI [1] (Member, IEEE), CHRISTOFOROS N. HADJICOSTIS [2] (Fellow, IEEE), AND GABRIELE OLIVA [1] (Senior Member, IEEE)

(Resilient and Safe Control in Multi-Agent Systems)

[1]University Campus Bio-Medico of Rome, 00128 Rome, Italy
[2]University of Cyprus, 1678 Nicosia, Cyprus

CORRESPONDING AUTHOR: GABRIELE OLIVA (e-mail: g.oliva@unicampus.it)

ABSTRACT Networked Control Systems (NCS) are pivotal for sectors like industrial automation, autonomous vehicles, and smart grids. However, merging communication networks with control loops brings complexities and security vulnerabilities, necessitating strong protection and authentication measures. This paper introduces an innovative Zero-Knowledge Proof (ZKP) scheme tailored for NCSs, enabling a networked controller to prove its knowledge of the dynamical model and its ability to control a discrete-time linear time-invariant (LTI) system to a sensor, without revealing the model. This verification is done through the controller's capacity to produce suitable control signals in response to the sensor's output demands. The completeness, soundness, and zero-knowledge properties of the proposed approach are demonstrated. The scheme is subsequently extended by considering the presence of delays and output noise. Additionally, a dual scenario where the sensor proves its model knowledge to the controller is explored, enhancing the method's versatility. Effectiveness is shown through numerical simulations and a case study on distributed agreement in multi-agent systems.

INDEX TERMS Computer/network security, control applications, networked control systems, resilient control systems, zero knowledge proof.

I. INTRODUCTION

Networked Control Systems (NCS) [1], [2], which can be defined as control systems wherein the control loops are closed through a real-time communication network, are emerging as disruptive technologies in many fields, such as industrial automation [3], [4], autonomous vehicles [5], and smart grids [6]. These kinds of systems allow information flows (e.g., reference input, plant output, control input) to be exchanged via a network between control system components, such as sensors, controllers, and actuators. This allows efficient fusion of global information to make intelligent decisions over a large physical space, improves scalability and interoperability, and reduces complexity, with significant practical and economic benefits [7], [8].

The above advantages, however, are not exempt from risks. This is due to the exposure of the data on public networks and third-party platforms, which foster cyber threats [9], [10], [11]. In particular, the most harmful attacks include, on one side, eavesdropping and spoofing of sensitive information and, on the other side, false data injection and data manipulation, which are often perpetrated on industrial control systems and critical infrastructures [12], [13], [14], with significant economic and social effects. Moreover, despite great advantages and wide applicability, the insertion of communication networks in feedback control loops makes the analysis and design of NCSs complicated [15], as the network itself is a dynamic system subject to time delays and measurement noise, which may degrade the performance of the control systems [16], [17]. For this reason, it is imperative to develop protection and authentication strategies that can make systems resilient to possible third-party intrusion into the public networks of the infrastructure, while also taking into account the effects of delays and noise on the system.

Over the past decade, significant research has been conducted in the field of secure control of networked systems, with a focus on encrypted control to protect the privacy and confidentiality of critical system states and model parameters [9], [18]. To achieve this, some cryptographic methods have been developed to ensure the privacy of the exchanged data and their resistance to potential eavesdroppers [19], [20]. One class of such techniques involves *homomorphic encryption* schemes, which enable mathematical operations to be completed on encrypted data, i.e., multiplication [21], [22], addition [23], [24], or both (fully homomorphic) [25], [26]. However, these schemes introduce quantization errors and impose a high computational burden that hinders their applicability to encrypted control on power- and memory-constrained devices, unless restricting encryption to a small number of operations allowed on ciphertexts [27]. Another technique to protect data confidentiality is based on the *differential privacy* concept [28], [29], [30], which obfuscates unstructured data by injecting layers of noise. This method has the merit of being less computationally onerous but it is known to provide results with scarce accuracy.

Although the previously mentioned techniques are effective in protecting the confidentiality of exchanged data, they are not intended to protect data integrity, since they cannot recognize or block injection attacks or data manipulation attacks [31], [32]. For this purpose, authentication mechanisms like *hash-based message authentication code* (HMAC) [33], [34] have been exploited in the context of control systems. However, these methods may jeopardize the temporal behavior of the network data communication because of the computational and communication overhead, and require the prior exchange of a secret key to apply the algorithm to the received messages. The exchange of information necessary for encryption, although minimal, places numerous constraints on networked systems, where the risk of spoofing is especially high.

To address this, a powerful technique, namely *Zero-Knowledge Proof* (ZKP) [35], [36], has recently found application in several domains such as machine learning [37], Internet of Things [38], or Blockchain [39], as a mechanism by which one party (the prover) can prove to another party (the verifier) that it knows a secret, without revealing the secret itself. Notably, the ZKP concept is particularly useful in the context of authentication, where one party wants to prove its identity via some secret information, but without sharing it [40], [41], [42].

Starting from this concept, some theoretical mechanisms were developed to implement exchanges between a prover and a verifier in a way that maintains the zero-knowledge property, such as those based on discrete logarithms [43], [44] or Hamiltonian cycles for a large graph [45]; several protocols have been developed over the years, starting from *Pinocchio* [46] and *Geppetto* [47] and culminating in more recent schemes, such as *Aurora* [48] and *Zilch* [49], which are considered plausibly post-quantum protocols, i.e., not susceptible to known attacks that may leverage on quantum computing.

To date, to the best of our knowledge, existing ZKP applications involve problems related to information technology and computer science at large [37], [38], [39], [40], [41], [42], although the concept has been applied in the context of embedded systems [50]. In particular, existing approaches are typically computer-theoretical in nature, in that both the prover and the verifier are able to exchange messages and perform computations. Notable examples in this sense include proving knowledge of the discrete logarithm of an integer value [43] or possession of information related to a graph (e.g., knowledge of a Hamiltonian cycle [45] or other properties of a graph [51]). Also, computer-theoretical ZKP approaches typically require iterating the same procedure multiple times (either sequentially, affecting the completion time of the procedure, or in parallel, which may require the transmission of large messages).

Conversely, applications of ZKPs to networked control system problems are lacking. Yet, NCS would largely benefit from an effective authentication mechanism; in doing so, an interesting feature would be the ability to actively leverage on the dynamical system, on the sensors measuring the system's output, and on the controller injecting the control signals.

A. CONTRIBUTION

In this paper, we introduce an innovative ZKP scheme specifically designed for the domain of NCSs. The proposed approach focuses on a scenario in which a networked controller, functioning as the prover, aims to demonstrate its knowledge of a discrete-time linear time-invariant (LTI) system's dynamical model by manipulating input signals to affect the system. The verifier, in this context represented by a sensor, selects target outputs to be generated by the system and challenges the prover to select adequate control signals to fulfill this requirement. The proposed scheme is shown to possess properties of completeness (i.e., the prover succeeds only provided that it knows the system's model), perfect soundness (the prover fails if it does not know the model), and zero-knowledge (i.e., the verifier gains no insights on the system's model). Moreover, differently from traditional computer-theoretical approaches, the proposed scheme actively leverages on the presence of an actual system and on an asymmetric setting where the prover can excite the system via a suitably chosen input, while the verifier can assess the effect of the injection in terms of the resulting measured outputs. Notably, the proposed approach is single shot, i.e., a single execution of the protocol is sufficient to ascertain whether or not the prover knows the secret. Moreover, since the prover testifies its knowledge of the system model via the injection of a control signal and the measurement of the resulting outputs, the proposed method keeps the number of exchanged messages to a minimum. The scheme is also extended to the cases of delayed effect of the input on the output as well as in the presence of measurement noise. Additionally, a dual scheme is provided where the roles are reversed, and the sensor takes on the role of the prover, seeking to verify its knowledge of the system model to the controller. To illustrate the practicality and applicability of our

proposed schemes, we have included numerical examples and a case study.

B. MOTIVATIONAL EXAMPLES

Consider a multi-agent system, interacting over a graph topology, that aims to reach distributed agreement on the average of some value. The agents are provided with different initial values and execute a discrete-time consensus process with the aim to compute the average of these initial values. In particular, we assume that the links of the graph, which connect the agents, represent respective knowledge of the agents' presence (e.g., an IP address) in a computer network. In this scenario, Peggy (the prover) is a higher-level entity that knows the topology of the graph and aims at recommending topological changes (e.g., providing additional IP addresses to some of the agents and/or requesting that some connection is dropped) in order to increase the overall network connectedness or resilience. In Section VII, we present a case study based on the above motivational example. Specifically, we assume that an agent plays the role of Victor (the verifier); in other words, before accepting topological changes, Victor challenges Peggy in order to verify its knowledge of the graph. In more detail, Victor selects arbitrary values for its own state; Peggy addresses this challenge by injecting a carefully chosen exogenous signal at some other agent in order to force Victor's state to its desired values. This scenario is considered in the case study in Section VII.

Another example would be the case of integrated power networks. In this scenario, different companies manage distinct segments of the network. For instance, consider that Company A is responsible for Subnetwork A, while Company B manages Subnetwork B. These two subnetworks are connected, and at their juncture there is a sensor that is owned by Company B. In order for Company A to effectively monitor and manage the network's performance, there might be the need that the sensor provides its information to Company A; however, before providing such measurements, the sensor must be convinced that it should provide data to Company A. To this end, Company B might issue a request, which is accepted by the sensor only if Company A succeeds in the ZKP challenge, proving its knowledge of the subnetwork model without revealing it to others.

C. PAPER OUTLINE

The outline of the paper is as follows: Section II collects some preliminary concepts and definitions, including the zero-knowledge proof (ZKP) model; Section III details the proposed approach for authentication in NCSs, while in Section IV we demonstrate the validity of the three fundamental properties of ZKP applied to the proposed approach. Section V collects useful alternative schemes that extend the proposed ZKP scheme in the presence of delays and noise, and to the dual case problem. Section VI provides some illustrative examples to numerically evaluate the effectiveness of the protocol and extensions, while Section VII presents the case study. Finally, Section VIII draws some conclusions and discusses possible directions for future work.

II. PRELIMINARIES

A. NOTATION

We denote vectors with boldface lowercase letters and matrices with uppercase letters. We refer to the (i, j)-th entry of a matrix A by A_{ij}. We represent by $\mathbf{0}_n$ and $\mathbf{1}_n$ vectors with n entries, all equal to zero and to one, respectively, and we use $0_{n\times m}$ to denote the $n \times m$ matrix with all entries equal to zero, while I_n denotes the $n \times n$ identity matrix. We use $\|\cdot\|$ to denote the Euclidean norm.

B. ZERO KNOWLEDGE PROOFS

Suppose Peggy has a password and wants to authenticate or access a Web site run by Victor, but they do not trust the computer Victor is using to verify the password. Hence, Peggy would aim to convince Victor that they know something without Victor finding out exactly what Peggy knows. This seemingly contradictory situation is addressed with Zero Knowledge Proofs (ZKP) [52], [53]. In the literature on ZKPs, Peggy's role is called the *prover*, since they wish to prove something, while Victor's role is called the *verifier*, since they wish to verify that the prover actually knows something. The main insight behind zero-knowledge proofs is that it is easy to prove possession of specific information by revealing it, but the real challenge lies in proving this possession without disclosing any details about the information itself. In a realistic scenario, nontrivial ZKPs require an interaction between the prover and the verifier, based on one or more challenges submitted by the verifier to the prover.

A trivial example that is used to explain the mechanism of ZKP is the Ali Baba cave [54]. Peggy has discovered the secret word to open a magic door in a cave in the shape of a ring, with the entrance on one side and the magic door blocking the opposite side. Peggy wants to prove to Victor that they know the secret word without revealing it. The two characters label the left and right paths of the entrance as A and B. First, Victor waits outside the cave while Peggy enters one of the two paths without being seen. Then, Victor enters the cave and shouts out the name of the path they want Peggy to use to get back, A or B, chosen at random. If Peggy really knows the magic word, they are able to open the door, if necessary, and go back along the desired path. Since Victor chooses A or B at random, Peggy has a 50% chance of guessing. If they were to repeat this challenge many times, the chance of successfully anticipating all of Victor's requests would be greatly reduced, and Victor might conclude that it is extremely likely that Peggy actually knows the secret word.

Some definitions are now in order to better define the main properties of ZKPs.

Definition 1 (Completeness): A Zero-Knowledge proof is said to be *complete* if Victor accepts the proof with probability one whenever Peggy knows the secret.

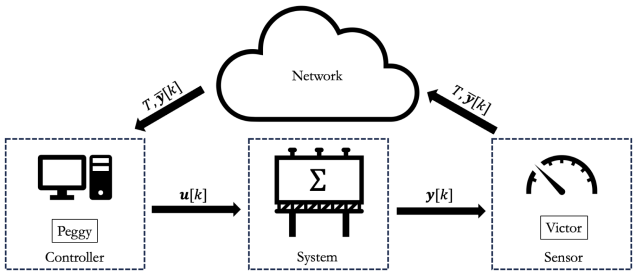

FIGURE 1. Scheme of the proposed protocol.

This property is related to the assumption that an honest verifier will always accept the answers of an honest provider.

Definition 2 (Soundness): A Zero-Knowledge proof is said to be *sound* if the probability that Victor accepts the proof when Peggy does not know the secret is small.

In other words, this property ensures that a cheating prover cannot convince an honest verifier, except with some small probability. From this property, it is possible to derive an even more stringent one, in which the probability is set to zero.

Definition 3 (Perfect Soundness): A Zero-Knowledge proof is said to be *perfectly sound* if the probability that Victor accepts the proof when Peggy does not know the secret is zero.

To conclude the last definition is about the "Zero-Knowledge" meaning the lack of new knowledge transmitted from the prover to the verifier.

Definition 4 (Zero-Knowledge): A proof is said to be *zero-knowledge* if Victor does not acquire any new knowledge from Peggy during the verification process.

This ensures security guarantees for honest provers since malicious verifiers cannot obtain any new information from what is sent to them.

III. PROBLEM STATEMENT

Let us consider a discrete-time LTI system as follows

$$\begin{cases} \boldsymbol{x}[k+1] = A\boldsymbol{x}[k] + B\boldsymbol{u}[k] \\ \boldsymbol{y}[k] = C\boldsymbol{x}[k], \end{cases} \tag{1}$$

with $A \in \mathbb{R}^{n\times n}$, $B \in \mathbb{R}^{n\times p}$, $C \in \mathbb{R}^{q\times n}$, and $\boldsymbol{x}[k]$, $\boldsymbol{u}[k]$, $\boldsymbol{y}[k]$ being vectors of dimensions n, p and q, respectively. Moreover, let us suppose that Peggy (which we can think of as a control unit) knows A, B, C and is able to apply $\boldsymbol{u}[k]$ to the system; however, Peggy has no direct way to measure $\boldsymbol{y}[k]$. Moreover, let us assume that Victor (which can play the role of a sensor) is able to measure $\boldsymbol{y}[k]$, but has no knowledge of the system's model, i.e., Victor does not know[1] A, B, C. Moreover, Peggy and Victor are able to exchange messages, e.g., via the Internet (see Fig. 1).

Let us now make the following assumptions.

Assumption 1: Peggy knows the initial conditions $\boldsymbol{x}[0]$.

Remark 1: The above assumption is required in order to be able to use the proposed ZKP scheme also when the system is not at rest. In fact, as discussed later in the paper, when $\boldsymbol{x}[0] \neq \boldsymbol{0}_n$, to be able to implement the proposed ZKP scheme Peggy must compensate for the system's free evolution, which depends on $\boldsymbol{x}[0]$. Notably, when $\boldsymbol{x}[0] = \boldsymbol{0}_n$, the ZKP scheme greatly simplifies.

Assumption 2: It holds $p = q = r$ and $CB \in \mathbb{R}^{r\times r}$ is nonsingular.

Remark 2: As discussed later in this section, the above assumption guarantees that Victor can select arbitrary values for the outputs and that such values correspond to a choice of the inputs. A simple case where Assumption 2 is verified is when there are just one input and one output (i.e., $p = q = 1$) and $CB \neq 0$. Another case is when there are sets $\mathcal{J} \subseteq \{1, \ldots, p\}$ and $\mathcal{M} \subseteq \{1, \ldots, q\}$ with $|\mathcal{M}| = |\mathcal{J}| = r$ such that the matrix $C^{(\mathcal{M})}B^{(\mathcal{J})} \in \mathbb{R}^{r\times r}$ is nonsingular, where $C^{(\mathcal{M})} \in \mathbb{R}^{r\times n}$ and $B^{(\mathcal{J})} \in \mathbb{R}^{n\times r}$ feature the rows of C indexed by $\mathcal{M}$ and the columns of B indexed by $\mathcal{J}$, respectively. In this latter case, Peggy will only use the inputs that correspond to the indices in $\mathcal{J}$, and Victor will only take into account the outputs indexed by $\mathcal{M}$.

Let us now state the main problem in this paper.

Problem 1: How can Peggy prove to Victor knowledge of A, B and C without revealing them to Victor?

In this paper we propose a ZKP interactive protocol to address Problem 1; the protocol relies on the fact that Peggy is able to apply input signals to the system while Victor is able to measure the system's output.

In particular, we consider a scenario where Victor challenges Peggy to apply a suitable control signal $\boldsymbol{u}[k]$ to the system, for $k \in \{0, \ldots, T-1\}$ so that $\boldsymbol{y}[1], \ldots, \boldsymbol{y}[T]$ assume arbitrarily selected values $\bar{\boldsymbol{y}}[1], \ldots, \bar{\boldsymbol{y}}[T]$. In other words, the protocol amounts to the following three steps:

I) Victor selects arbitrary $T \geq 1$ and $\bar{\boldsymbol{y}}[k] \in \mathbb{R}^q$ for $k \in \{1, \ldots, T\}$ and sends them to Peggy.
II) Peggy applies[2] a suitable control signal $\boldsymbol{u}[k]$ to the system for $k \in \{0, \ldots, T-1\}$.
III) Victor verifies whether or not $\boldsymbol{y}[k] = \bar{\boldsymbol{y}}[k]$ for $k \in \{1, \ldots, T\}$.

Matrix $\Omega[T]$ below will play a pivotal role within the proposed ZKP scheme:

$$\Omega[T] = \begin{bmatrix} CB & 0 & \ldots & \ldots & \ldots & 0 \\ CAB & CB & 0 & \ldots & \ldots & 0 \\ CA^2B & CAB & CB & 0 & \ldots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ CA^{T-2}B & \ddots & \ddots & \ddots & \ddots & 0 \\ CA^{T-1}B & \ldots & \ldots & \ldots & CAB & CB \end{bmatrix}.$$

[1] Interestingly, this implies that Victor (who is able to measure $\boldsymbol{y}[k]$), knows the number of outputs, but does not know the number of state variables.

[2] For the sake of simplicity, and without loss of generality, we assume the time at which Peggy begins to apply the signal is $k = 0$.

Moreover, let

$$\Gamma[T] = \begin{bmatrix} CA \\ \vdots \\ CA^T \end{bmatrix}.$$

In view of the later developments in this paper, we need the following technical proposition.

Proposition 1: Under Assumption 2, $\Omega[T]$ is nonsingular for all $T \geq 1$.

Proof: In order to prove the result we observe that, $\Omega[T]$ is block triangular with diagonal blocks equal to CB. It is well known that, for block triangular matrices, the rank is lower bounded by the sum of the ranks of the diagonal blocks. Therefore, we have that

$$\texttt{rank}(\Omega[T]) \geq T\,\texttt{rank}(CB).$$

Since, by Assumption 2, CB is nonsingular, we conclude that $\texttt{rank}(CB) = r$. This implies that

$$\texttt{rank}(\Omega[T]) \geq Tr,$$

and since $\Omega[T] \in \mathbb{R}^{Tr \times Tr}$ we conclude that $\texttt{rank}(\Omega[T]) = Tr$ and thus $\Omega[T]$ is nonsingular. This completes our proof. ■

Since $\Omega[T]$ is nonsingular, the system of linear equations

$$\boldsymbol{Y} = \Omega[T]\boldsymbol{U} + \boldsymbol{F}$$

with $\boldsymbol{Y}, \boldsymbol{F} \in \mathbb{R}^{Tr}$ and $\boldsymbol{U} \in \mathbb{R}^{Tr}$ admits a unique solution $\boldsymbol{U} = \Omega^{-1}[T](\boldsymbol{Y} - \boldsymbol{F})$. At this point, let us define

$$\bar{\boldsymbol{Y}}[T] = \begin{bmatrix} \bar{\boldsymbol{y}}^\top[1] & \dots & \bar{\boldsymbol{y}}^\top[T] \end{bmatrix}^\top$$

and

$$\begin{aligned} \bar{\boldsymbol{U}}[T] &= \Omega^{-1}[T]\left(\bar{\boldsymbol{Y}}[T] - \Gamma[T]\boldsymbol{x}[0]\right) \\ &= \begin{bmatrix} \bar{\boldsymbol{u}}^\top[0] & \dots & \bar{\boldsymbol{u}}^\top[T-1] \end{bmatrix}^\top. \end{aligned} \tag{2}$$

Within the proposed ZKP protocol, Peggy reacts to the challenge posed by Victor by applying the input $\boldsymbol{u}[k] = \bar{\boldsymbol{u}}[k]$, for all $k \in \{0, \dots, T-1\}$.

IV. PROTOCOL ANALYSIS

In this Section, we aim to show that the proposed protocol belongs to the class of zero-knowledge proofs, by demonstrating the three fundamental properties: completeness, soundness, and zero-knowledge (see Section II-B).

Let us now first establish the completeness of the proposed control-theoretical ZKP protocol.

Theorem 2: Under Assumptions 1–2, the proposed control-theoretical ZKP scheme is complete.

Proof: In order to prove completeness, let us assume Peggy knows A, B, C. Based on the proposed scheme, Victor transmits $T, \bar{\boldsymbol{Y}}[T]$ to Peggy. Peggy reacts by injecting the input $\boldsymbol{u}[k] = \bar{\boldsymbol{u}}[k]$ to the system for $k \in \{0, \dots, T-1\}$. We recall that, by Assumption 1, Peggy knows $\boldsymbol{x}[0]$. Therefore, for $k \geq 1$, the system's output corresponds to

$$\begin{aligned} \boldsymbol{y}[k] &= CA^k\boldsymbol{x}[0] + C\sum_{h=0}^{k-1} A^{k-h-1}B\boldsymbol{u}[h] \\ &= CA^k\boldsymbol{x}[0] + C\sum_{h=0}^{k-1} A^{k-h-1}B\bar{\boldsymbol{u}}[h]. \end{aligned}$$

Stacking the above equation for all $k \in \{1, \dots, T\}$, we have that

$$\begin{aligned} \boldsymbol{Y}[T] &= \begin{bmatrix} \boldsymbol{y}^\top[1] & \dots & \boldsymbol{y}^\top[T] \end{bmatrix}^\top \\ &= \Gamma[T]\boldsymbol{x}[0] + \Omega[T]\bar{\boldsymbol{U}}[T] \\ &= \Gamma[T]\boldsymbol{x}[0] + \Omega[T]\underbrace{\Omega^{-1}[T]\left(\bar{\boldsymbol{Y}}[k] - \Gamma[T]\boldsymbol{x}[0]\right)}_{\bar{\boldsymbol{U}}[T]} \\ &= \bar{\boldsymbol{Y}}[T]. \end{aligned}$$

where $\Omega^{-1}[T]$ is guaranteed to exist by Proposition 1. In conclusion, when A, B, and C are known by Peggy, Peggy is able to fulfill Victor's request. This completes our proof. ■

We now prove the perfect soundness of the proposed ZKP scheme.

Theorem 3: Let Assumptions 1 and 2 hold true. Then, the proposed control-theoretical ZKP scheme is perfectly sound.

Proof: In order to prove the result, let us first show that, if Peggy does not know A, B, and C, then they have zero probability to set $\boldsymbol{Y}[k] = \bar{\boldsymbol{Y}}[k]$.

In a worst-case scenario for Victor, let us assume that Peggy knows n and $\boldsymbol{x}(0)$ and is able to inject the input $\boldsymbol{u}[k]$. We have that Peggy must solve for $A, B, C, \bar{\boldsymbol{U}}[T]$ a system of equations in the form

$$\bar{\boldsymbol{Y}}[T] = \Omega[T]\bar{\boldsymbol{U}}[T] + \Gamma[T]\boldsymbol{x}[0]. \tag{3}$$

In particular, Peggy must solve Tr polynomial equations in $n^2 + 2nr + Tr$ unknowns. Hence, the system of polynomial equations is underdetermined. Since by Proposition 1, for any given choice of A, B, C that satisfies Assumption 2, the system has a unique solution

$$\bar{\boldsymbol{U}}[T] = \Omega^{-1}[T]\left(\bar{\boldsymbol{Y}}[T] - \Gamma[T]\boldsymbol{x}[0]\right), \tag{4}$$

we conclude that the system of polynomial equations has infinitely many solutions. Moreover, for given A, B, C that satisfy Assumption 2, since $\bar{\boldsymbol{U}}[T]$ is unique, we have that the set of solutions has zero measure. Therefore, Peggy has probability zero to select any such solution. This completes our proof. ■

Using an argument similar to the one used above, we now establish that the proposed ZKP scheme is zero-knowledge.

Theorem 4: Let Assumptions 1 and 2 hold true. Then, the proposed control-theoretical ZKP scheme is zero-knowledge.

Proof: In a worst-case scenario, let us assume that Victor knows n and $\boldsymbol{x}[0]$. Similarly to Theorem 3, to correctly identify $A, B, C, \bar{\boldsymbol{U}}[T]$ Victor must solve a system of Tr polynomial equations in $n^2 + 2nr + Tr$ unknowns, which is

underdetermined and has an infinity of solutions. Moreover, the set of solutions have zero measure. Therefore, Victor has zero probability of identifying A, B, C and, thus, attains zero knowledge from the proposed protocol. This completes the proof. ■

Remark 3: Notice that, being a sensor that measures some linear combination of the state variables, Victor could be able to detect if some of the outputs are blowing up, which could imply the presence of at least one unstable eigenvalue. However, in the presence of nonzero initial conditions, the system would exhibit a divergent dynamics anyway, independently of the signal injection undertaken by Peggy. Moreover, we would like to point out that the input injected by Peggy is designed in order to have the system's dynamics assume arbitrarily selected outputs, independently from the possible stability or instability of the system. For instance, Peggy could mask the instability of the system by fulfilling the requirement of the ZKP scheme and then by continuing to inject carefully chosen inputs that completely override the natural stability/instability of the system. Thus, by observing the outputs after time T, Victor would not be able to reach a conclusion on the stability or instability of the system.

V. EXTENSIONS AND ALTERNATIVE SCHEMES

This section collects useful alternative schemes that extend the proposed ZKP scheme.

A. INPUTS WITH DELAYED EFFECT ON THE OUTPUT

The proposed ZKP scheme requires that CB is full row rank. However, our scheme can be applied also in the case characterized by the following assumption.

Assumption 3: There is a positive integer s such that $CA^hB = 0_{r\times r}$ for all $h \in \{0, \dots, s-1\}$, while CA^sB is nonsingular.

The above assumption models a scenario where there is a delay between the application of an input and its effect on the output. Before discussing the proposed extension, let us show that s must be smaller than or equal to n.

Lemma 1: Let Assumption 3 hold true. Then, it must hold that $s < n$.

Proof: Let us prove our claim by induction. To this end we observe that, by the Cayley-Hamilton Theorem, for all $g \geq 0$ it holds

$$A^{n+g} = -\alpha_0 A^g - \dots - \alpha_{n-1}A^{n+g-1},$$

where $\alpha_0, \dots, \alpha_{n-1}$ are the coefficients of the characteristic polynomial of A. At this point, let us consider the case $g = 0$. We have that either $s < n$ or $CA^hB = 0_{r\times r}$ for all $h \in \{1, \dots, n-1\}$. Assuming that the latter holds true (otherwise our claim is trivially verified), we have that

$$CA^nB = -\alpha_0 CI_nB - \dots - \alpha_{n-1}CA^{n-1}B = 0.$$

Now, let us assume that for some $g \geq 0$ it holds $CA^hB = 0_{r\times r}$ for all $h \in \{1, \dots, n+g\}$ and let us show that, then, also $CA^{n+g+1}B = 0_{r\times r}$. To this end, we observe that

$$CA^{n+g+1}B = -\alpha_0 CA^{g+1}B - \dots - \alpha_{n-1}CA^{n+g}B = 0.$$

This completes our inductive proof. ■

Let us now characterize the proposed extended ZKP scheme. To this end, we observe that, under Assumption 3, it is sufficient that Peggy sends s to Victor before the ZKP procedure begins, and that Victor selects arbitrary

$$\bar{\boldsymbol{Y}}[s, T] = \left[\bar{\boldsymbol{y}}^\top[s+1] \quad \dots \quad \bar{\boldsymbol{y}}^\top[s+T]\right]^\top,$$

ignoring the values assumed by the output before time $s+1$. To fulfill the challenge posed by Victor, Peggy must solve for $\bar{\boldsymbol{U}}[T]$ the following equation

$$\bar{\boldsymbol{Y}}[s, T] = \Gamma[s, T]\boldsymbol{x}[0] + \Omega[s, T]\bar{\boldsymbol{U}}[T],$$

where

$$\Omega[s, T] = \begin{bmatrix} CHB & 0 & \dots & \dots & \dots & 0 \\ CHAB & CHB & 0 & \dots & \dots & 0 \\ CHA^2B & CHAB & CHB & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ CHA^{T-2}B & \ddots & \ddots & \ddots & \ddots & 0 \\ CHA^{T-1}B & \dots & \dots & \dots & CHAB & CHB \end{bmatrix},$$

with $H = A^s$, and

$$\Gamma[s, T] = \begin{bmatrix} CA^{s+1} \\ \vdots \\ CA^{s+T} \end{bmatrix}.$$

Let us now characterize the properties of this variation of the proposed ZKP scheme. Notice that, since Peggy provides s to Victor and since $s \leq n$, the price to pay is that, although still being unable to determine A, B, C, Victor gains insights on a lower bound on the number of state variables.

Theorem 5: Under Assumption 3, the proposed ZKP scheme with delayed effect of the input on the output is complete and perfectly sound. Moreover, although being able to obtain a lower bound on n, Victor is unable to ascertain A, B, C.

Proof: In order to prove the statement we observe that, if Peggy knows A, B, C, then they are able to choose the correct $\bar{\boldsymbol{U}}[T]$. In fact, following the same argument as in Proposition 1, $\Omega[s, T]$ is nonsingular and thus the unique solution is

$$\bar{\boldsymbol{U}}[T] = \Omega^{-1}[s, T]\left(\bar{\boldsymbol{Y}}[s, T] - \Gamma[s, T]\boldsymbol{x}[0]\right). \tag{5}$$

Thus, the ZKP scheme is complete. Let us now prove perfect soundness and zero-knowledge. To this end we observe that, to reconstruct $A, B, C, \bar{\boldsymbol{U}}[T]$, there is a need to solve a system of $(T+s)r$ polynomial equations with $n^2 + 2nr + Tr$

unknowns, which is underdetermined if

$$s < \frac{n^2}{r} + 2n. \tag{6}$$

Therefore, following the same argument as in Theorems 3 and 4, the ZKP scheme is perfectly sound. Moreover, under the additional requirement that s satisfies (6), Victor is unable to ascertain A, B, C. However, by Lemma 1, we have that $s < n$, hence (6) is always satisfied. In any case, Victor obtains knowledge of s, which is a lower bound on the number n of state variables. This completes our proof. ■

Remark 4: Notice that Peggy could communicate any value $\bar{s}$, with

$$s < \bar{s} < \frac{n^2}{r} + 2n,$$

to Victor and set all the outputs $\boldsymbol{y}[s], \boldsymbol{y}[s+1], \ldots, \boldsymbol{y}[\bar{s}-1]$ to zero. Even in the case $s = 0$, Peggy could communicate a nonzero $\bar{s}$ and prevent Victor from understanding if there is a real delay of $\bar{s}$ steps or if $\bar{s}$ is artificial.

B. EXTENSION TO IMPRECISE KNOWLEDGE OF A, B, C

In this subsection, we consider the case where Peggy has imprecise knowledge regarding the matrices A, B, C of the system. Specifically, let us assume that Peggy knows A', B', C' which are in general different from A, B, C, respectively. As a result, Peggy is only able to compute $\Omega'[T] = \Omega[T] + \Delta\Omega[T]$, but has no way to compute the exact $\Omega[T]$. The next theorem characterizes the resulting relative output error based on the magnitude of $\Delta\Omega[T]$.

Theorem 6: Suppose Peggy takes part to our proposed ZKP scheme selecting the input based on $\Omega'[T]$. Moreover, assume that $C'B'$ is nonsingular and let

$$\phi[T] = \frac{\|\Delta\Omega[T]\|}{\sigma_{\min}(\Omega[T] + \Delta\Omega[T])},$$

where $\sigma_{\min}(\Omega[T] + \Delta\Omega[T])$ is the minimum singular value of $\Omega[T] + \Delta\Omega[T]$ and satisfies $\sigma_{\min}(\Omega[T] + \Delta\Omega[T]) > 0$. Then, under Assumption 2, it holds

$$\frac{\|Y[T] - \overline{Y}[T]\|}{\|\overline{Y}[T]\|} \leq \phi[T],$$

with $\phi[T] < \infty$ and $\lim_{\|\Delta A\| \to 0} \phi[T] = 0$.

Proof: In order to prove the statement, let us consider a scenario where Victor requests a given $\overline{Y}[T]$ and, based on an imperfect knowledge on A, B, C, Peggy injects an input

$$U[T] = (\Omega[T] + \Delta\Omega[T])^{-1}\overline{Y}[T].$$

Notice that $\Omega[T]$ is nonsingular by Assumption 2. Similarly, we observe that also $\Omega[T] + \Delta\Omega[T]$ is lower block triangular with diagonal blocks equal to $C'B'$. Hence, since we assumed $C'B'$ is nonsingular, also $\Omega[T] + \Delta\Omega[T]$ is nonsingular. As a consequence of the above input, the output becomes

$$Y[T] = \Omega[T]U[T] = \Omega[T](\Omega[T] + \Delta\Omega[T])^{-1}\overline{Y}[T].$$

This implies that

$$(\Omega[T] + \Delta\Omega[T])\Omega^{-1}[T]Y[T] = \overline{Y}[T],$$

i.e.,

$$Y[T] = (I + \Delta\Omega[T]\Omega^{-1}[T])^{-1}\overline{Y}[T].$$

At this point, we use the Woodbury formula [55] to expand

$$(I + \Delta\Omega[T]\Omega^{-1}[T])^{-1} = I - \Delta\Omega[T](I + \Omega^{-1}[T]\Delta\Omega[T])^{-1}\Omega^{-1}[T]$$

obtaining

$$\begin{aligned} Y[T] - \overline{Y}[T] &= -\Delta\Omega[T](I + \Omega^{-1}[T]\Delta\Omega[T])^{-1}\Omega^{-1}[T]\overline{Y}[T] \\ &= -\Delta\Omega[T](\Omega[T] + \Delta\Omega[T])^{-1}\overline{Y}[T]. \end{aligned}$$

As a consequence

$$\|Y[T] - \overline{Y}[T]\| \leq \|\Delta\Omega[T]\|\|(\Omega[T] + \Delta\Omega[T])^{-1}\|\|\overline{Y}[T]\|$$

i.e.,

$$\begin{aligned} \frac{\|Y[T] - \overline{Y}[T]\|}{\|\overline{Y}[T]\|} &\leq \|\Delta\Omega[T]\|\|(\Omega[T] + \Delta\Omega[T])^{-1}\| \\ &= \frac{\|\Delta\Omega[T]\|}{\sigma_{\min}(\Omega[T] + \Delta\Omega[T])} = \phi[T]. \end{aligned} \tag{7}$$

Notably, since $\Omega[T] + \Delta\Omega[T]$ is nonsingular, we have that $\sigma_{\min}(\Omega[T] + \Delta\Omega[T]) > 0$ and thus $\phi[T] < \infty$. To conclude, we observe that, as $\|\Delta\Omega[T]\|$ approaches zero $\sigma_{\min}(\Omega[T] + \Delta\Omega[T])$ approaches $\sigma_{\min}(\Omega[T])$ and $\phi[T]$ approaches zero. This completes our proof. ■

Theorem 6 states that, when Peggy has incorrect knowledge of $\Omega[T]$ because of a small perturbation $\Delta\Omega[T]$, the resulting output will also have a small relative difference with respect to the one Victor is expecting. In this case, the proposed scheme is no longer perfectly sound, but this extension allows the overall ZKP scheme to be more flexible. For instance, Victor could accept results that are within a given threshold for the relative output error, and Peggy succeeds provided that the imperfectly known matrices A, B, C are such that $\|\Delta\Omega[T]\|$ is small.

C. EXTENSION TO MEASUREMENT NOISE

In this subsection, we extend our original ZKP protocol to the case where the outputs measured by Victor are affected by measurement noise. To this end, for the sake of simplicity, we assume $\boldsymbol{x}[0] = \boldsymbol{0}_n$.

Suppose that the measurements $\boldsymbol{y}[k]$ are affected by independent zero mean Gaussian noises, i.e.,

$$\boldsymbol{y}[k] = C\boldsymbol{x}[k] + \boldsymbol{w}[k], \tag{8}$$

with $\boldsymbol{w}[k] \sim \mathcal{N}(\boldsymbol{0}_r, W)$ for some covariance matrix $W \in \mathbb{R}^{r \times r}$, which we assume to be unknown to Victor.

Notably, since Victor challenges Peggy to generate outputs $\bar{\boldsymbol{y}}[k]$ and measures noisy $\boldsymbol{y}[k]$, we have that Victor is able to compute

$$\overline{\boldsymbol{w}}[k] = \boldsymbol{y}[k] - \bar{\boldsymbol{y}}[k].$$

Let us now consider the null hypothesis

$$\mathtt{H}_0: \text{ Peggy knows } A, B, C$$

and the alternative hypothesis

$$\mathtt{H}_1: \text{ Peggy does not know } A, B, C.$$

In the first case, since by Theorem 2 the proposed ZKP scheme in the absence of noise is complete, Peggy successfully sets states $\boldsymbol{x}[k]$ such that $C\boldsymbol{x}[k] = \bar{\boldsymbol{y}}[k]$. Therefore, the vectors $\overline{\boldsymbol{w}}[k] = \boldsymbol{w}[k]$ have zero mean, i.e., for $k \in \{1, \dots, T\}$, it holds

$$\overline{\boldsymbol{w}}[k] \sim \mathcal{N}(\mathbf{0}_r, W),$$

and, thus

$$\widetilde{\boldsymbol{w}} \sim \mathcal{N}\left(\mathbf{0}_r, \frac{1}{T}W\right),$$

where

$$\widetilde{\boldsymbol{w}} = \frac{1}{T}\sum_{k=1}^{T} \overline{\boldsymbol{w}}[k]$$

is the experimental average of the vectors $\overline{\boldsymbol{w}}[k]$.

Translating $\mathtt{H}_1$ into mathematical terms is more challenging. In fact, when Peggy does not know the system's model, in principle she can select any control signal. In the following, in order to characterize in probabilistic terms the effect of the injection of a wrong input signal, we assume that (as a result of the injection) the outputs are $\boldsymbol{y}[k] = \boldsymbol{y}^{\dagger}[k] + \boldsymbol{w}[k]$, with $\boldsymbol{y}^{\dagger}[k] \neq \bar{\boldsymbol{y}}[k]$. In this case, with $\boldsymbol{x}[0] = \mathbf{0}_n$, we have that $\overline{\boldsymbol{w}}[k] = \boldsymbol{y}^{\dagger}[k] + \boldsymbol{w}[k] - \bar{\boldsymbol{y}}[k]$. In other words, since $\boldsymbol{w}[k]$ is zero-mean, it holds

$$\overline{\boldsymbol{w}}[k] \sim \mathcal{N}(\boldsymbol{y}^{\dagger}[k] - \bar{\boldsymbol{y}}[k], W)$$

and thus

$$\widetilde{\boldsymbol{w}} \sim \mathcal{N}\left(\boldsymbol{m}_1, \frac{1}{T}W\right), \quad \boldsymbol{m}_1 = \frac{1}{T}\sum_{k=1}^{T}\left(\boldsymbol{y}^{\dagger}[k] - \bar{\boldsymbol{y}}[k]\right).$$

Based on the above scenarios, we can restate our hypotheses as follows. Specifically, assume a null hypothesis

$$\mathtt{H}_0: \widetilde{\boldsymbol{w}} \sim \mathcal{N}\left(\boldsymbol{m}_0, \frac{1}{T}W\right),$$

where $\boldsymbol{m}_0 = \mathbf{0}_r$, versus the alternative hypothesis

$$\mathtt{H}_1: \widetilde{\boldsymbol{w}} \sim \mathcal{N}\left(\boldsymbol{m}_1, \frac{1}{T}W\right).$$

In order to decide, let us assume that Victor resorts to the *Hotelling's* $\mathtt{T}^2$ *test* (e.g., see [56]). Assuming W is unknown to Victor, let us consider the statistic

$$\mathtt{T}^2 = T\left(\widetilde{\boldsymbol{w}} - \boldsymbol{m}_0\right)^{\top} S^{-1}\left(\widetilde{\boldsymbol{w}} - \boldsymbol{m}_0\right), \tag{9}$$

where

$$S = \frac{1}{T-1}\sum_{k=1}^{T}\left(\overline{\boldsymbol{w}}[k] - \widetilde{\boldsymbol{w}}\right)\left(\overline{\boldsymbol{w}}[k] - \widetilde{\boldsymbol{w}}\right)^{\top}$$

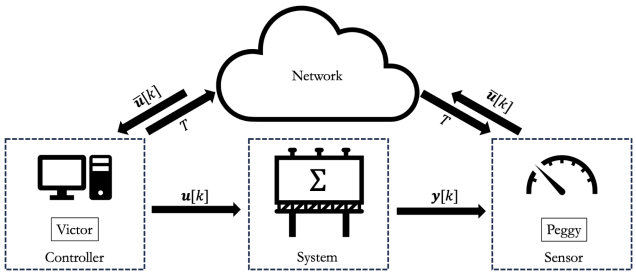

FIGURE 2. Dual scheme of the proposed protocol.

is the experimental covariance matrix.

It is well known (e.g., see [56]) that

$$\mathtt{T}^2 \sim \frac{(T-1)r}{T-r}\mathcal{F}_{r,T-r},$$

where $\mathcal{F}_{r,T-r}$ denotes the $\mathtt{F}$-distribution with r and $T-r$ degrees of freedom. Therefore, Victor rejects $\mathtt{H}_0$ with a significance level α (typically 0.05 or 0.01) if

$$\mathtt{T}^2 > \frac{(T-1)r}{T-r}\mathcal{F}_{r,T-r}(\alpha), \tag{10}$$

where $\mathcal{F}_{r,T-r}(\alpha)$ is the upper (100α)th percentile of the $\mathcal{F}_{r,T-r}$ distribution, i.e., the value $\mathcal{F}_{r,T-r}(\alpha)$ is such that

$$\int_0^{\mathcal{F}_{r,T-r}(\alpha)} \mathcal{F}_{r,T-r}(z)dz = 1-\alpha.$$

This implies that the proposed ZKP is no longer complete: in fact, when Peggy knows the secret, Victor accepts the result only with probability $1-\alpha$.

Notably, due to the noise, the ZKP protocol is also not perfectly sound. In fact, there is a probability β that Victor does not reject $\mathtt{H}_0$ when $\mathtt{H}_1$ is true. In particular, it is well known that $1-\beta$, the *power* of the test (i.e., the probability to reject $\mathtt{H}_0$ when $\mathtt{H}_1$ is true), is the area to the right of $\mathcal{F}_{r,T-r}(\alpha)$ for the noncentral $\mathtt{F}$-distribution $\mathcal{F}_{r,T-r,\lambda}(\cdot)$ with r and $T-r$ degrees of freedom and noncentrality parameter

$$\lambda = T(\boldsymbol{m}_0 - \boldsymbol{m}_1)^{\top} S^{-1}(\boldsymbol{m}_0 - \boldsymbol{m}_1).$$

In other words, we have that

$$1-\beta = \int_{\mathcal{F}_{r,T-r}(\alpha)}^{\infty} \mathcal{F}_{r,T-r,\lambda}(z)dz,$$

which corresponds to

$$\beta = 1 - \int_{\mathcal{F}_{r,T-r}(\alpha)}^{\infty} \mathcal{F}_{r,T-r,\lambda}(z)dz$$
$$= \int_0^{\mathcal{F}_{r,T-r}(\alpha)} \mathcal{F}_{r,T-r,\lambda}(z)dz.$$

D. A DUAL ZKP SCHEME

Let us consider a dual problem where Peggy and Victor are swapped (see Fig. 2): Peggy knows $A, B, C, \boldsymbol{x}(0)$ and is able to measure the outputs, while Victor does not know the system's model, but has the ability to inject a control signal. In this dual

problem, the aim of Peggy is to prove to Victor that they know A, B, C. In this view, let us consider a ZKP protocol consisting of the following steps:

I) Victor selects arbitrary $T \geq 1$ and sends it to Peggy.
II) Victor applies an arbitrary control signal $\boldsymbol{u}[k]$ to the system for $k \in \{0, \dots, T-1\}$.
III) Peggy measures the outputs $\boldsymbol{y}[k]$ for $k \in \{1, \dots, T\}$, computes the input values $\bar{\boldsymbol{u}}[k]$ for $k \in \{0, \dots, T-1\}$ and sends them to Victor.
IV) Victor verifies whether or not $\boldsymbol{u}[k] = \bar{\boldsymbol{u}}[k]$ for $k \in \{0, \dots, T-1\}$.

Based on Assumptions 1–2, with a similar argument as in our original ZKP protocol, if Peggy knows A, B, C then they are able to compute the unique solution given in (4), hence the dual ZKP protocol is complete. If, conversely, they do not know A, B, C, following the same argument as in Theorems 3 and 4, the ZKP scheme is perfectly sound and zero-knowledge.

This dual scheme can also be extended to the case discussed in Section V-A, where the inputs affect the output with a delay s. In this case, Peggy computes the input values according to (5). and the scheme is complete, perfectly sound, and zero-knowledge.

VI. ILLUSTRATIVE EXAMPLES

In this section, we provide some illustrative examples to numerically show the effectiveness of the protocol.

Example 1: Let us consider a dynamic system characterized by the following matrices

$$A = \begin{bmatrix} 0.466 & -0.002 & -0.140 & -0.013 \\ 0.167 & 0.530 & -0.136 & -0.141 \\ -0.077 & 0.068 & 0.348 & 0.060 \\ 0.264 & 0.005 & 0.105 & 0.454 \end{bmatrix},$$

$$B = \begin{bmatrix} 0.913 & -1.888 \\ 0.976 & 0.736 \\ -1.577 & 1.117 \\ 0.204 & 0.743 \end{bmatrix}, \quad C = \begin{bmatrix} 1.616 & 1.667 \\ -0.103 & 0.623 \\ -0.193 & 0.795 \\ 0.161 & 0.539 \end{bmatrix}^\top.$$

Moreover, let us assume that

$$\boldsymbol{x}[0] = \begin{bmatrix} -0.140 & -0.983 & 0.408 & 0.243 \end{bmatrix}^\top.$$

Notice that

$$CB = \begin{bmatrix} 1.712 & -3.223 \\ 0.986 & -1.400 \end{bmatrix}$$

is nonsingular and thus satisfies Assumption 2.

Let us assume that Victor wants to verify that Peggy knows A, B, C (according to the scheme in Fig. 1). To this aim, Victor selects $T = 6$ and challenges Peggy to generate the outputs

$$\bar{\boldsymbol{y}}[1] = \begin{bmatrix} 0.600 \\ 0.266 \end{bmatrix}, \quad \bar{\boldsymbol{y}}[2] = \begin{bmatrix} 0.285 \\ 0.254 \end{bmatrix}, \quad \bar{\boldsymbol{y}}[3] = \begin{bmatrix} 0.328 \\ 0.144 \end{bmatrix},$$

$$\bar{\boldsymbol{y}}[4] = \begin{bmatrix} 0.165 \\ 0.964 \end{bmatrix}, \quad \bar{\boldsymbol{y}}[5] = \begin{bmatrix} 0.960 \\ 0.188 \end{bmatrix}, \quad \bar{\boldsymbol{y}}[6] = \begin{bmatrix} 0.024 \\ 0.205 \end{bmatrix}.$$

In response to this challenge, Peggy selects adequate inputs $\bar{\boldsymbol{u}}[\cdot]$ according to (2), i.e.,

$$\bar{\boldsymbol{u}}[0] = \begin{bmatrix} 1.681 \\ 0.665 \end{bmatrix}, \quad \bar{\boldsymbol{u}}[1] = \begin{bmatrix} -2.075 \\ -1.021 \end{bmatrix}, \quad \bar{\boldsymbol{u}}[2] = \begin{bmatrix} 1.988 \\ 0.930 \end{bmatrix},$$

$$\bar{\boldsymbol{u}}[3] = \begin{bmatrix} 1.227 \\ 0.724 \end{bmatrix}, \quad \bar{\boldsymbol{u}}[4] = \begin{bmatrix} -5.338 \\ -3.052 \end{bmatrix}, \quad \bar{\boldsymbol{u}}[5] = \begin{bmatrix} 7.737 \\ 4.063 \end{bmatrix}.$$

Finally, Victor verifies the output provided by the system in response to the output sequence, comparing it with $\bar{\boldsymbol{y}}$. Notice that, this example is similarly applicable to the dual problem discussed in Section V-D. In fact, it can be shown that the above choices for the inputs and the outputs satisfy (3).

Example 2: Let us now analyze the case described in Section V-A, which features a delay between the application of the input and its effect on the output. To model it, let us consider the following input and output matrices

$$B = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad C = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix},$$

while A is the same as Example 1. We notice that

$$CB = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix},$$

while

$$CAB = \begin{bmatrix} -0.140 & -0.013 \\ -0.136 & -0.141 \end{bmatrix}$$

is nonsingular. Therefore, Assumption 3 is satisfied for $s = 1$. For simplicity of exposition, let us assume that Victor selects $T = 6$ and the same outputs as the previous example. In this case, the inputs provided by Peggy to obtain the desired outputs with a finite delay $s = 1$ are

$$\bar{\boldsymbol{u}}[0] = \begin{bmatrix} -4.799 \\ 0.004 \end{bmatrix}, \quad \bar{\boldsymbol{u}}[1] = \begin{bmatrix} 1.625 \\ 0.422 \end{bmatrix}, \quad \bar{\boldsymbol{u}}[2] = \begin{bmatrix} -1.519 \\ 1.648 \end{bmatrix},$$

$$\bar{\boldsymbol{u}}[3] = \begin{bmatrix} 0.947 \\ -7.114 \end{bmatrix}, \quad \bar{\boldsymbol{u}}[4] = \begin{bmatrix} -6.974 \\ 12.190 \end{bmatrix}, \quad \bar{\boldsymbol{u}}[5] = \begin{bmatrix} 5.165 \\ -6.346 \end{bmatrix}.$$

Let us now provide an example where Peggy has incorrect knowledge of A.

Example 3: Let us consider the same matrices A, B, C as in Example 1 but, for simplicity, let $\boldsymbol{x}[0] = \boldsymbol{0}_4$. Moreover, we assume that Peggy knows the nominal B, C matrices, but has imperfect knowledge on A. Specifically, Peggy knows $A' = A + \Delta A$ where, for the sake of simplicity, $\Delta A = \alpha I_4$. In Fig. 3 we plot the results obtained for the relative output error and for the upper bound $\phi[T]$ for different values of the parameter

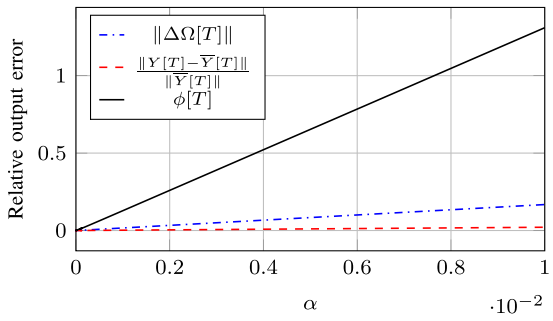

FIGURE 3. Relative output error (red dashed line) and upper bound $\phi[T]$ (black solid line) experienced by Victor when Peggy does not perfectly know A, based on the magnitude α of the perturbation. The plot also reports $\|\Delta\Omega[T]\|$ (dashed–dotted blue line).

α, along with the corresponding magnitude $\|\Delta\Omega[T]\|$ of the perturbation of $\Omega[T]$. According to the plot, the three curves are essentially linear in α. In particular, by linear fitting, we experimentally obtain a slope of 2.11 for the relative output error, a slope of 16.77 for $\|\Delta\Omega[T]\|$ and a slope of 130.97 for $\phi[T]$. Overall, this analysis suggests that, when the magnitude of $\|\Delta\Omega[T]\| \ll 1$, also the resulting relative output error will be small, meaning that our approach can be extended to a scenario where Victor accepts using a given error threshold (at the cost of losing perfect soundness of the ZKP scheme).

Let us conclude this section by considering an example where the output measured by Victor is affected by noise.

Example 4: Let us consider the same matrices A, B, C as in Example 1 but, for simplicity, let $\boldsymbol{x}[0] = \mathbf{0}_4$. Moreover, let $T = 50$ and assume that Victor selects

$$\bar{\boldsymbol{y}}[k] = [-25 + k, 26 - k]^\top, \quad k \in \{1, \dots, 50\}$$

while, for the sake of brevity, we omit the values of $\bar{\boldsymbol{u}}[k]$, which are computed according to (4). Let us assume that the noise affecting the output is zero-mean and has covariance

$$W = \begin{bmatrix} 0.1 & 0.05 \\ 0.05 & 0.1 \end{bmatrix}.$$

Let us first assume that Peggy knows A, B, C and is able to correctly select the input. In this case, the Hotelling test yields a score $\mathtt{T}^2 = 0.046$ while, for $\alpha = 0.05$, the threshold is 6.514. Therefore, there is not enough evidence to reject $\mathtt{H}_0$ and Victor concludes that Peggy knows the secret with high confidence. Let us now assume that Peggy does not know the system's model and, for simplicity, assume $\boldsymbol{u}[k] = \mathbf{0}_2$, which results in $\boldsymbol{y}^\dagger = \mathbf{0}_2$ and, consequently, in

$$\boldsymbol{m}_1 = -\frac{1}{T}\sum_{k=1}^{T} \bar{\boldsymbol{y}}[k].$$

Interestingly, in this example, we have that the probability β that Victor does not reject $\mathtt{H}_0$ when $\mathtt{H}_1$ is true is $\beta = 0.103$. Conversely, if Peggy does not know A, B, C, assuming that no input is injected, we have that $\mathtt{T}^2 = 12.636$ is almost twice the threshold, hence in this case $\mathtt{H}_0$ is rejected and Victor concludes that Peggy does not know the secret.

VII. CASE STUDY

Let us consider a scenario where a network of n agents interact over a connected graph topology $G = \{V, E\}$, where V is the node set and E is the set of edges. Each agent is represented by a node $v_i \in V$, while an edge $(v_i, v_j) \in E$ models the fact agent v_i and agent v_j interact. For the sake of simplicity, let us assume that the graph is undirected (i.e., $(v_i, v_j) \in E$ whenever $(v_j, v_i) \in E$) and connected (i.e., each agent is able to reach each other agent via a path that is constructed using the edges in E). The agents interact according to a discrete-time consensus process with an exogenous input, i.e., each agent is characterized by a dynamics in the form

$$x_i[k+1] = \sum_{v_j \in \mathcal{N}_i} P_{ij} x_j[k] + b_i u_i[k],$$

where $\mathcal{N}_i$ is the *neighborhood* of agent v_i, i.e., the set of agents that are connected to agent v_i via an edge. In a compact form for all agents, the above dynamics reads

$$\boldsymbol{x}[k+1] = P\boldsymbol{x}[k] + B\boldsymbol{u}[k], \tag{11}$$

where

$$\boldsymbol{x}[k] = \begin{bmatrix} x_1[k] & x_2[k] & \dots & x_n[k] \end{bmatrix}^\top,$$
$$\boldsymbol{u}[k] = \begin{bmatrix} u_1[k] & u_2[k] & \dots & u_n[k] \end{bmatrix}^\top,$$

and P_{ij} (i.e., the entry of P at the i-th row and j-th column position) is zero whenever i and j are not connected by an edge. Furthermore, P is a doubly stochastic matrix, i.e., $P_{ij} \in [0, 1]$, $P\mathbf{1}_n = \mathbf{1}_n$, and $\mathbf{1}_n^\top P = \mathbf{1}_n^\top$. Regarding B, we have that $B = I_n$ is the identity matrix. It is well known that, in the absence of an exogenous input (i.e., when $u_i[k] = 0$ for all k and all i), the above dynamics converges to $x_{\mathtt{ave}}\mathbf{1}_n$, where $x_{\mathtt{ave}}$ is the average of the agents' initial states. Let us now assume that there is an entity i (e.g., the agent v_i itself or another entity), Victor, and let us assume that Victor is provided with an output

$$y_i[k] = C_i\boldsymbol{x}[k],$$

where $C_i^\top \in \mathbb{R}^n$ and C_i has all entries equal to zero, except the i-th one, which is equal to one. This choice of C_i models a scenario where Victor knows the state of agent v_i.

Suppose further that there is another entity, Peggy, claiming to know P, and able to inject some $u_j[k]$ into the system. Peggy is interested in demonstrating this claim in order to authenticate itself as a regulatory entity. For instance, in a scenario where an edge models the knowledge of the IP address of an agent, Peggy might want to prove knowledge of the topology before recommending local topological changes in Victor's neighborhood (e.g., to improve the overall connectivity or resilience of the network). In this view, the proposed ZKP scheme could be effectively leveraged upon to avoid an adversary that pretends to be the higher-level entity and modifies the topology to cause harm to the network.

The system dynamics, in this case, are

$$\begin{cases} \boldsymbol{x}[k+1] = P\boldsymbol{x}[k] + B_j u_j[k], \\ \boldsymbol{y}_i[k] = C_i\boldsymbol{x}[k], \end{cases}$$

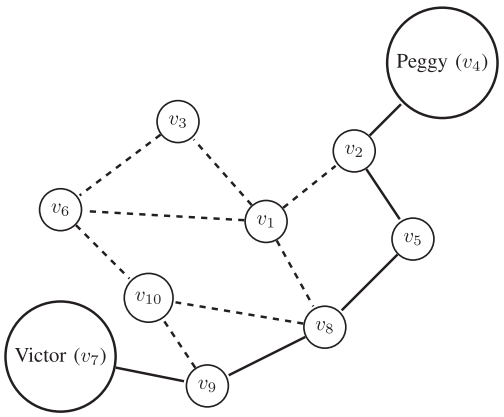

FIGURE 4. Graph topology considered for the simulations.

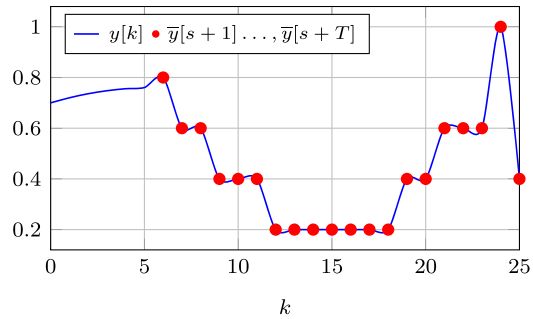

FIGURE 5. Output measured by Victor. The red dots represent the desired outputs $\bar{y}_7[k]$ with $k \in \{6, \ldots, 25\}$, selected by Victor.

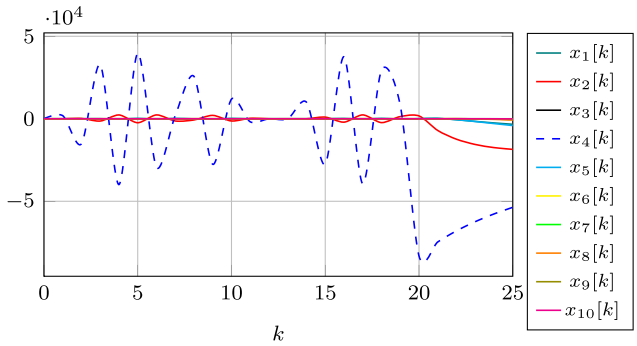

FIGURE 6. Evolution of the states of the agents as a result of the exogenous signal $u_4[k]$ injected by Peggy. The state of node v_4 is shown via the blue dashed line.

where B_j is the j-th column of B. In particular, suppose that Assumption 3 holds true for some s, i.e., suppose that it takes s steps for the input to have an effect on the output. Suppose further that the agents have, in general, nonzero initial conditions and are bound to follow the consensus dynamics in (11).

Notice that, for the sake of this application, we can assume that C_i (i.e., the identifier i of the node whose state is available to Victor) is known to Peggy (e.g., because Victor tells it to Peggy). This implies that, in this scenario, the number of unknowns for Victor is slightly reduced. In particular, here $r = 1$ and Victor is provided with $T + s$ equations and $n^2 + n + T$ unknowns. By Lemma 1, we are guaranteed that $n^2 + n > s$, hence the system of equations is still underdetermined. This implies that the ZKP scheme is still perfectly sound and, even though Victor gets to know a lower bound s on the number n of state variables, there is no way to determine A, B, C. Moreover, the fact Victor knows C_i has no negative effect on the completeness of the protocol, which still holds true.

In order to verify Peggy's claims, Victor chooses arbitrarily a positive integer T and the output values $\bar{y}_i[k]$ with $k \in \{s+1, \ldots, s+T\}$ and transmits them to Peggy. In turn, Peggy computes suitable values for $u_j[0], \ldots, u_j[T-1]$ and injects them into the system. By verifying whether or not $y_i[k]$ corresponds to $\bar{y}_i[k]$ for $k \in \{s+1, \ldots, s+T\}$, Victor is able to ascertain if Peggy's claim is justified, but gains zero knowledge about the entry point of the input (described by B_j) and the topology of the network (captured by P).

To numerically prove the effectiveness of the proposed protocol in this case, let us consider a graph topology G (reported in Fig. 4) with 10 nodes, each holding a scalar initial condition expressed by

$$\boldsymbol{x}[0] = \begin{bmatrix} 0.1 & 0.2 & 0.3 & 0.4 & 0.5 & 0.6 & 0.7 & 0.8 & 0.9 & 1 \end{bmatrix}^\top.$$

In this framework, we assume that node v_4 (Peggy) claims to know the matrix $P = I - 0.1L$, where L is the Laplacian matrix[3] associated with the graph G, and to be able to inject a signal $u_4[k]$ onto itself; in other words, B_4 is a column with all entries equal to zero, except for the fourth one, which is equal to one. To verify Peggy's claim, node v_7 (Victor) challenges Peggy to modify their state; in other words, the output matrix C_7 is a row with all entries equal to zero, except for the seventh one, which is equal to one. Notice that, for this example, Assumption 3 is satisfied for $s = 5$, which corresponds to the length of the minimum path from node v_4 to node v_7. In particular, Victor selects $T = 20$ and the following output sequence

$$\begin{aligned}
&\bar{y}[6] = 0.8 \quad \bar{y}[7] = 0.6 \quad \bar{y}[8] = 0.6 \quad \bar{y}[9] = 0.4 \\
&\bar{y}[10] = 0.4 \quad \bar{y}[11] = 0.4 \quad \bar{y}[12] = 0.2 \quad \bar{y}[13] = 0.2 \\
&\bar{y}[14] = 0.2 \quad \bar{y}[15] = 0.2 \quad \bar{y}[16] = 0.2 \quad \bar{y}[17] = 0.2 \\
&\bar{y}[18] = 0.2 \quad \bar{y}[19] = 0.4 \quad \bar{y}[20] = 0.4 \quad \bar{y}[21] = 0.6 \\
&\bar{y}[22] = 0.6 \quad \bar{y}[23] = 0.6 \quad \bar{y}[24] = 1 \quad \bar{y}[25] = 0.4.
\end{aligned}$$

To prove the claim, Peggy injects suitable input signals (i.e., computed via (5)) into the system, and the effect of this injection is evident in Fig. 6, which reports the output of the system from the perspective of node v_7. According to the figure, the output follows the behavior indicated by Victor, with a delay equal to the number of communication steps $s = 5$ required to transfer the information from Peggy to Victor. Fig. 5 shows the evolution of the states $\boldsymbol{x}[k]$ of all the agents in the network; it is

[3]Given an undirected graph $G = \{V, E\}$ with n nodes, we define the Laplacian matrix L as the $n \times n$ matrix such that $L_{i,j} = -1$ if $(v_j, v_i) \in E$, $L_{i,j} = d_i$ if $i = j$, and $L_{i,j} = 0$, otherwise.

observed that node 4 which injects the input signal undergoes the most pronounced state changes.

VIII. CONCLUSIONS AND FUTURE WORK

In this paper, we presented a novel Zero-Knowledge Proof scheme tailored to networked control systems. In particular, we considered a scenario where a networked controller aims to prove to a sensor that it knows the dynamical model of a system and, to this end, is able to inject input signals to the system. The sensor, acting as a verifier, selects desired outputs and transmits them to the prover which, in turn, reacts by carefully selecting adequate control signals. The proposed scheme is proved to be complete, perfectly sound and zero knowledge. Moreover, extensions to cases where there is a delayed effect of the input on the output or the measurements are noisy are presented. Finally, a dual scheme is discussed where the sensor knows the model and aims to prove this knowledge to a controller. Numerical examples and a case study in the context of distributed agreement complete the paper.

Future work will follow four main directions: (1) extend the approach to nonlinear, delayed, and uncertain systems (both linear and nonlinear), while also considering the definition of system-theoretical cryptosystems; (2) apply the scheme as an authentication procedure in hybrid contexts involving both local, distributed agents and higher level centralized entities (e.g., optimization, estimation, load balancing or localization problems with a combination of local and global computations); (3) devise approaches to simultaneously stabilize a system while fulfilling a ZKP challenge (e.g., considering systems with multiple inputs, where some degrees of freedom can be used to stabilize the system, while some others could be used to convince a sensor about the fact that the controller knows the dynamical model); (4) develop system-theoretical approaches that are grounded on existing ZKP and cryptographic primitives and/or on computationally challenging control-theoretical problems.

REFERENCES

[1] A. Bemporad et al., *Networked Control Systems*. vol. 406, Berlin, Germany: Springer, 2010.

[2] A. Rajagopal and S. Chitraganti, "State estimation and control for networked control systems in the presence of correlated packet drops," *Int. J. Syst. Sci.*, vol. 54, no. 11, pp. 2352–2365, 2023.

[3] B. Rahmani and A. H. D. Markazi, "Networked control of industrial automation systems-a new predictive method," *Int. J. Adv. Manuf. Technol.*, vol. 58, pp. 803–815, 2012.

[4] H. Yang, C. Peng, and Z. Cao, "Attack-model-independent stabilization of networked control systems under a jump-like TOD scheduling protocol," *Automatica*, vol. 152, 2023, Art. no. 110982.

[5] D. Plöger, L. Krüger, and A. Timm-Giel, "Analysis of communication demands of networked control systems for autonomous platooning," in *Proc. IEEE 19th Int. Symp. "A World Wireless, Mobile Multimedia Netw."*, 2018, pp. 14–19.

[6] A. K. Singh, R. Singh, and B. C. Pal, "Stability analysis of networked control in smart grids," *IEEE Trans. Smart Grid*, vol. 6, no. 1, pp. 381–390, Jan. 2015.

[7] R. A. Gupta and M.-Y. Chow, "Overview of networked control systems," *Networked Control Syst.: Theory Appl.*, pp. 1–23, 2008.

[8] G. C. Walsh, H. Ye, and L. G. Bushnell, "Stability analysis of networked control systems," *IEEE Trans. Control Syst. Technol.*, vol. 10, no. 3, pp. 438–446, May 2002.

[9] M. S. Darup, A. B. Alexandru, D. E. Quevedo, and G. J. Pappas, "Encrypted control for networked systems: An illustrative introduction and current challenges," *IEEE Control Syst. Mag.*, vol. 41, no. 3, pp. 58–78, Jun. 2021.

[10] R. A. Nafea and M. A. Almaiah, "Cyber security threats in cloud: Literature review," in *Proc. IEEE 2021 Int. Conf. Inf. Technol.*, 2021, pp. 779–786.

[11] Y. Tan, Y. Yuan, X. Xie, E. Tian, and J. Liu, "Observer-based event-triggered control for interval type-2 fuzzy networked system with network attacks," *IEEE Trans. Fuzzy Syst.*, vol. 31, pp. 2788–2798, Aug. 2023.

[12] Y. Zou and G. Wang, "Intercept behavior analysis of industrial wireless sensor networks in the presence of eavesdropping attack," *IEEE Trans. Ind. Informat.*, vol. 12, no. 2, pp. 780–787, Apr. 2016.

[13] G. Liang, J. Zhao, F. Luo, S. R. Weller, and Z. Y. Dong, "A review of false data injection attacks against modern power systems," *IEEE Trans. Smart Grid*, vol. 8, no. 4, pp. 1630–1638, Jul. 2017.

[14] A. S. Mohammed, E. Anthi, O. Rana, N. Saxena, and P. Burnap, "Detection and mitigation of field flooding attacks on oil and gas critical infrastructure communication," *Comput. Secur.*, vol. 124, 2023, Art. no. 103007.

[15] X. Ge, F. Yang, and Q.-L. Han, "Distributed networked control systems: A brief overview," *Inf. Sci.*, vol. 380, pp. 117–131, 2017.

[16] X. Luan, P. Shi, and F. Liu, "Stabilization of networked control systems with random delays," *IEEE Trans. Ind. Electron.*, vol. 58, no. 9, pp. 4323–4330, Sep. 2011.

[17] X.-S. Zhan, J. Wu, T. Jiang, and X.-W. Jiang, "Optimal performance of networked control systems under the packet dropouts and channel noise," *ISA Trans.*, vol. 58, pp. 214–221, 2015.

[18] A. Tsiamis, K. Gatsis, and G. J. Pappas, "State estimation codes for perfect secrecy," in *Proc. IEEE 56th Annu. Conf. Decis. Control*, 2017, pp. 176–181.

[19] K. Kogiso and T. Fujita, "Cyber-security enhancement of networked control systems using homomorphic encryption," in *Proc. IEEE 54th Conf. Decis. Control*, 2015, pp. 6836–6843.

[20] C. Fioravanti, V. Bonagura, G. Oliva, C. N. Hadjicostis, and S. Panzieri, "Exploiting the synchronization of nonlinear dynamics to secure distributed consensus," *IEEE Open J. Control Syst.*, vol. 2, pp. 249–262, 2023.

[21] T. ElGamal, "A public key cryptosystem and a signature scheme based on discrete logarithms," *IEEE Trans. Inf. Theory*, vol. 31, no. 4, pp. 469–472, Jul. 1985.

[22] K. Kogiso, "Encrypted control using multiplicative homomorphic encryption," *Privacy Dyn. Syst.*, pp. 267–286, 2020.

[23] P. Paillier, "Public-key cryptosystems based on composite degree residuosity classes," in *Proc. Int. Conf. Theory Appl. Cryptographic Techn.*, 1999, pp. 223–238.

[24] T. B. Ogunseyi and T. Bo, "Fast decryption algorithm for paillier homomorphic cryptosystem," in *Proc. IEEE 2020 Int. Conf. Power, Intell. Comput. Syst.*, 2020, pp. 803–806.

[25] K. Teranishi, T. Sadamoto, and K. Kogiso, "Input-output history feedback controller for encrypted control with leveled fully homomorphic encryption," *IEEE Trans. Control Netw. Syst.*, vol. 11, no. 1, pp. 271–283, Mar. 2024.

[26] P. Stobbe, T. Keijzer, and R. M. Ferrari, "A fully homomorphic encryption scheme for real-time safe control," in *Proc. IEEE 61st Conf. Decis. Control*, 2022, pp. 2911–2916.

[27] J. H. Cheon, A. Kim, M. Kim, and Y. Song, "Homomorphic encryption for arithmetic of approximate numbers," in *Proc. Adv. Cryptology–ASIACRYPT 2017: 23rd Int. Conf. on Theory Appl. Cryptology Inf. Secur.*, 2017, pp. 409–437.

[28] Y. Zhao and J. Chen, "A survey on differential privacy for unstructured data content," *ACM Comput. Surv.*, vol. 54, no. 10s, pp. 1–28, 2022.

[29] M. U. Hassan, M. H. Rehmani, and J. Chen, "Differential privacy techniques for cyber physical systems: A survey," *IEEE Commun. Surv. Tut.*, vol. 22, no. 1, pp. 746–789, Firstquarter 2020.

[30] J. Cortés, G. E. Dullerud, S. Han, J. Le Ny, S. Mitra, and G. J. Pappas, "Differential privacy in control and network systems," in *Proc. IEEE 55th Conf. Decis. Control*, 2016, pp. 4252–4272.

[31] A. Sargolzaei, K. Yazdani, A. Abbaspour, C. D. Crane III, and W. E. Dixon, "Detection and mitigation of false data injection attacks in networked control systems," *IEEE Trans. Ind. Informat.*, vol. 16, no. 6, pp. 4281–4292, Jun. 2020.
[32] A. Abdallah and X. S. Shen, "Efficient prevention technique for false data injection attack in smart grid," in *Proc. IEEE 2016 Int. Conf. Commun.*, 2016, pp. 1–6.
[33] G. Martins, A. Bhattacharjee, A. Dubey, and X. D. Koutsoukos, "Performance evaluation of an authentication mechanism in time-triggered networked control systems," in *Proc. IEEE 7th Int. Symp. Resilient Control Syst.*, 2014, pp. 1–6.
[34] G. Martins, A. Moondra, A. Dubey, A. Bhattacharjee, and X. D. Koutsoukos, "Computation and communication evaluation of an authentication mechanism for time-triggered networked control systems," *Sensors*, vol. 16, no. 8, 2016, Art. no. 1166.
[35] I. Aad, "Zero-knowledge proof," in *Trends in Data Protection and Encryption Technologies*. Berlin, Germany: Springer, 2023, pp. 25–30.
[36] J. Kurmi and A. Sodhi, "A survey of zero-knowledge proof for authentication," *Int. J. Adv. Res. Comput. Sci. Softw. Eng.*, vol. 5, no. 1, 2015, Art. no. 1145.
[37] Y. Fan, B. Xu, L. Zhang, J. Song, A. Zomaya, and K.-C. Li, "Validating the integrity of convolutional neural network predictions based on zero-knowledge proof," *Inf. Sci.*, vol. 625, pp. 125–140, 2023.
[38] R. Singh, A. D. Dwivedi, G. Srivastava, P. Chatterjee, and J. C.-W. Lin, "A privacy preserving Internet of Things smart healthcare financial system," *IEEE Internet Things J.*, pp. 18452–18460, Nov. 2023.
[39] L. Zhou, A. Diro, A. Saini, S. Kaisar, and P. C. Hiep, "Leveraging zero knowledge proofs for blockchain-based identity sharing: A survey of advancements, challenges and opportunities," *J. Inf. Secur. Appl.*, vol. 80, 2024, Art. no. 103678.
[40] D. Gabay, K. Akkaya, and M. Cebe, "Privacy-preserving authentication scheme for connected electric vehicles using blockchain and zero knowledge proofs," *IEEE Trans. Veh. Technol.*, vol. 69, no. 6, pp. 5760–5772, Jun. 2020.
[41] W. Li, C. Meese, H. Guo, and M. Nejad, "Aggregated zero-knowledge proof and blockchain-empowered authentication for autonomous truck platooning," *IEEE Trans. Intell. Transp. Syst.*, pp. 9309–9323, Sep. 2023.
[42] J. K. Shahrouz and M. Analoui, "An anonymous authentication scheme with conditional privacy-preserving for vehicular ad hoc networks based on zero-knowledge proof and blockchain," *Ad Hoc Netw.*, vol. 154, 2024, Art. no. 103349.
[43] D. Chaum, J. H. Evertse, and J. V. D. Graaf, "An improved protocol for demonstrating possession of discrete logarithms and some generalizations," in *Proc. Adv. Cryptol.-EUROCRYPT'87: Workshop Theory Appl. Cryptographic Techn.*, 1988, pp. 127–141.
[44] S. Kim, H. Lee, and J. H. Seo, "Efficient zero-knowledge arguments in discrete logarithm setting: Sublogarithmic proof or sublinear verifier," in *Proc. Int. Conf. Theory Appl. Cryptol. Inf. Secur.*, 2022, pp. 403–433.
[45] M. Blum, "How to prove a theorem so no one else can claim it," in *Proc. Int. Congr. Mathematicians*, 1986, pp. 1444–1451.
[46] B. Parno, J. Howell, C. Gentry, and M. Raykova, "Pinocchio: Nearly practical verifiable computation," *Commun. ACM*, vol. 59, no. 2, pp. 103–112, 2016.
[47] C. Costello et al., "Geppetto: Versatile verifiable computation," in *Proc. 2015 IEEE Symp. Secur. Privacy*, 2015, pp. 253–270.
[48] E. Ben-Sasson, A. Chiesa, M. Riabzev, N. Spooner, M. Virza, and N. P. Ward, "Aurora: Transparent succinct arguments for R1CS," in *Proc. Adv. Cryptology–EUROCRYPT 2019: 38th Annu. Int. Conf. Theory Appl. Cryptographic Techn., Part I. 38*, 2019, pp. 103–128.
[49] D. Mouris and N. G. Tsoutsos, "Zilch: A framework for deploying transparent zero-knowledge proofs," *IEEE Trans. Inf. Forensics Secur.*, vol. 16, pp. 3269–3284, 2021.
[50] X. Salleras and V. Daza, "ZPiE: Zero-knowledge proofs in embedded systems," *Math.*, vol. 9, no. 20, 2021, Art. no. 2569.
[51] M. F. Esgin, R. Steinfeld, J. K. Liu, and D. Liu, "Lattice-based zero-knowledge proofs: New techniques for shorter and faster constructions and applications," in *Proc. 39th Annu. Int. Cryptol. Conf.*, 2019, pp. 115–146.
[52] N. P. Smart, *Cryptography Made Simple*. Berlin, Germany: Springer, 2016.
[53] D. Firsov and D. Unruh, "Zero-knowledge in EasyCrtypt," in *Proc. IEEE 36th Comput. Secur. Foundations Symp.*, 2023, pp. 1–16.
[54] J.-J. Quisquater et al., "How to explain zero-knowledge protocols to your children," in *Proc. Conf. Theory Appl. Cryptol.*, 1989, pp. 628–631.
[55] M. A. Woodbury, "Inverting modified matrices," in *Memorandum Rept. 42, Statistical Research Group*. Princeton, NJ, USA: Princeton University, 1950, p. 4.
[56] R. A. Johnson et al., *Applied Multivariate Statistical Analysis*, New York, NY, USA: Taylor & Francis, 2002.

CAMILLA FIORAVANTI (Member, IEEE) received the M.Sc. degree in biomedical engineering and the Ph.D. degree in science and engineering for humans and the environment from the University Campus Bio-Medico of Rome, Roma, Italy, in 2020 and 2024, respectively. She is currently a Postdoc Research Fellow with the University Campus Bio-Medico of Rome. She spent a visiting period in 2022 and 2023 with the University of Cyprus. Her research interests include distributed systems, distributed estimation, security and privacy-preserving approaches, and fault detection.

CHRISTOFOROS N. HADJICOSTIS (Fellow, IEEE) received the S.B. degrees in electrical engineering, computer science and engineering, and mathematics, the M.Eng. degree in electrical engineering and computer science in 1995, and the Ph.D. degree in electrical engineering and computer science in 1999, from the Massachusetts Institute of Technology, Cambridge, MA, USA. Since 2007, he has been with the Department of Electrical and Computer Engineering, University of Cyprus, Nicosia, Cyprus, where he is currently a Professor and Interim Director of the Daedalus Research Center. His research interests include fault diagnosis and tolerance in distributed dynamic systems, error control coding, monitoring, diagnosis and control of large-scale discrete-event systems, and applications to network security, anomaly detection, energy distribution systems, and medical diagnosis. He is the Editor-in-Chief of the *Journal of Discrete Event Dynamic Systems* and as Senior Editor of IEEE TRANSACTIONS ON AUTOMATIC CONTROL. He was also an Associate Editor for *Automatica*, IEEE TRANSACTIONS ON AUTOMATION SCIENCE AND ENGINEERING, IEEE TRANSACTIONS ON CONTROL SYSTEMS TECHNOLOGY, IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS I, and *Hybrid Systems: Nonlinear Analysis*.

GABRIELE OLIVA (Senior Member, IEEE) received the M.Sc. and Ph.D. degrees in computer science and automation engineering from the University Roma Tre of Rome, Rome, Italy, in 2008 and 2012, respectively. He is currently an Associate Professor of automatic control with the University Campus Bio-Medico of Rome, Italy, where he directs the Complex Systems & Security Laboratory (CoserityLab). His main research interests include distributed multiagent systems, optimization, estimation, decision-making, and critical infrastructure protection. Since 2019, he has been an Associate Editor on the Conference Editorial Board of the IEEE Control Systems Society. Since 2022, he has been an Associate Editor for IEEE CONTROL SYSTEMS LETTERS.

Open Access funding provided by 'Università "Campus Bio-Medico" di Roma' within the CRUI CARE Agreement