

HArMoNICS: High-Assurance Microgrid Network Infrastructure Case Study

Original

HArMoNICS: High-Assurance Microgrid Network Infrastructure Case Study / Roascio, G., Costa, G., Baccelli, E., Malina, L., Matulevicius, R., Momeu, M., Morkevicius, N., Russo, E., Stojanovic, B., Tasidou, A.. - In: IEEE ACCESS. - ISSN 2169-3536. - ELETTRONICO. - 10:(2022), pp. 115372-115383. [10.1109/ACCESS.2022.3218412]

Availability:

This version is available at: 11583/2974478 since: 2023-01-10T13:16:29Z

Publisher:

IEEE-INST ELECTRICAL ELECTRONICS ENGINEERS

Published

DOI:10.1109/ACCESS.2022.3218412

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

RESEARCH ARTICLE

HArMoNICS: High-Assurance Microgrid Network Infrastructure Case Study

GIANLUCA ROASCIO¹, GABRIELE COSTA², (Member, IEEE), EMMANUEL BACCELLI³,
LUKAS MALINA⁴, RAIMUNDAS MATULEVIČIUS⁵, MARIUS MOMEU⁶,
NERIJUS MORKEVIČIUS⁷, ENRICO RUSSO⁸, BRANKA STOJANOVIĆ⁹,
AND AIMILIA TASIDOU¹⁰

¹Department of Control and Computer Engineering, Polytechnic of Turin, 10129 Turin, Italy

²Systems Security Modeling and Analysis Group, IMT School for Advanced Studies, 55100 Lucca, Italy

³National Institute for Research in Computer Science and Automation, 75012 Paris, France

⁴Department of Telecommunications, Brno University of Technology, 60190 Brno, Czech Republic

⁵Institute of Computer Science, University of Tartu, 51009 Tartu, Estonia

⁶Department of Informatics, Technology University of Munich, 80333 Munich, Germany

⁷Department of Computer Science, Kaunas University of Technology, 44249 Kaunas, Lithuania

⁸Department of Computer Science, Bioengineering, Robotics and Systems Engineering, University of Genoa, 16145 Genoa, Italy

⁹Institute for Information and Communication Technologies, Joanneum Research Digital, A-8010 Graz, Austria

¹⁰Télécom SudParis, 91000 Courcouronnes, France

Corresponding author: Gianluca Roascio (gianluca.roascio@polito.it)

This work was supported by the European Union's Horizon 2020 Research and Innovation Programme under Grant 830892, Project SPARTA.

ABSTRACT Modern Intelligent Infrastructures (II) are highly complex, interconnected systems that are now emerging. For instance, II can integrate technologies and processes to provide citizens with faster services and better goods. An average II can include many technologies, e.g., Cloud applications and IoT devices, under different environments, e.g., industry 4.0 production plants and smart buildings. Although II bring concrete benefits to all of these contexts, they also carry security concerns. Reasoning about threats and security exposures that might affect II is non-trivial. This is only partially due to their inherent complexity. As a matter of fact, real II are typically in charge of some critical operations that cannot be interrupted or compromised for experimental purposes. An alternative solution is to rely on digital replicas which can provide a good trade-off between realism and usability. These assets represent a strategic and highly demanded resource for the security community. In this paper, we present HArMoNICS, a case study infrastructure meant to provide a playground for security experts interested in II security. HArMoNICS revolves around a digital replica of a real Smart Polygeneration Microgrid (SPM) located in Italy. Although most of the components are based on or inspired by the real system, HArMoNICS has been enriched with further security-relevant features. As a result, the case study includes vertical use cases focusing on specific security topics. Security researchers can use it to assess the effectiveness of new methodologies, to carry out security training activities, or even to extend it with new elements.

INDEX TERMS Security, cyber security, embedded systems, intelligent infrastructure, case study.

I. INTRODUCTION

The development of modern Intelligent Infrastructures (II) promises to raise the bar for several aspects of our everyday life. Nowadays, computationally-enabled devices include televisions, watches, alarm clocks, and many others surrounding us. These devices are progressively outnumbering other

types, e.g., personal computers and smartphones, and they form the well-known Internet of Things (IoT). The number of connected devices is expected to skyrocket from 8.74 billion (2020) to more than 25.4 billion (2030) [1]. IoT is an enabling technology for II, yet not the only one. As a matter of fact, other technological, e.g., Fog and Cloud computing [2], [3], and infrastructural pillars, e.g., 5G networks [4], are directly involved.

The associate editor coordinating the review of this manuscript and approving it for publication was Pedro R. M. Inácio¹⁰.

As often happens, along with opportunities, these technologies also bring new risks. Serious concerns exist that latent vulnerabilities may pave the way for attackers. Needless to say, security violations would have a dramatic impact, e.g., on the privacy of citizens and the continuity of critical services. Also, existing security mechanisms might not be applicable to II. A reason is that these infrastructures are extremely complex and heterogeneous, made of myriads of objects using hardware and software of many different manufacturers. Such extreme diversity, together with the quick development of modern technologies, requires appropriate countermeasures to ensure the security of the next-generation II.

A growing trend is that of using *Security-by-Design* (S×D) [5] as the leading principle. In short, S×D enriches the traditional development lifecycle with security-specific tasks. These tasks take place at every stage, from the very early design to the final deployment. By integrating these tasks in the development workflow, the S×D approach creates a security management process where every critical event triggers effects along the entire lifecycle. As a practical example, consider the case of a new vulnerability discovered in an existing II. Such a vulnerability may occur due to a design weakness, which one might want to evaluate against the initial security specifications. At the same time, the vulnerability may trigger the development of countermeasures, e.g., an emergency patch, as well as new security tests.

Although the ideas behind S×D and its potential benefits are clear, implementing it is still a major challenge. As a matter of fact, all the security tasks previously mentioned must be implemented and populated with actual security tools and procedures. For each of them, many alternatives exist and new ones appear over time. As a consequence, the entire security workflow must be continuously revised and maintained. Even worse, assessing the actual effectiveness of the involved security procedures is hardly feasible. For instance, most of these procedures require discontinuing the operations of the II, which is typically infeasible. Hence, the actual effectiveness of the processes is often only assessed when a real security event occurs and, if they fail, when it is too late.

A possible solution for the assessment of S×D frameworks is to rely on II replicas. For instance, computer simulations can accurately reproduce the behavior of a real system. However, virtual II are rarely available. A reason is that, although simplified, they typically bring part of the complexity of a real II. Thus, they are often considered valuable assets by private actors who may prefer not to be shared. Furthermore, when they are created to mimic a real II, they may lack the security aspects of interest, e.g., vulnerabilities, for assessing the effectiveness of security solutions.

In this paper, we present our *High-Assurance Microgrid Network Infrastructure Case Study* (HArMoNICS), developed within the EU project SPARTA.¹ Briefly, HArMoNICS

revolves around a zero-emission building scenario in the context of a smart microgrid. All the elements appearing in our case study are taken from or inspired by actual systems that reside inside the original infrastructure.

The main motivation behind our proposal is the strategic role of intentionally vulnerable environments, which provide a shared setting for research and development of security techniques. Indeed, such environments stimulate and favor innovation by providing a touchstone for new methodologies. The main goal of HArMoNICS is to be a valuable asset for the community interested in the security and privacy of critical and intelligent infrastructures. Researchers can use it both as a *playground*, e.g., for hosting training exercises, and as a *benchmark*, e.g., for systematically assessing the effectiveness of a certain security mechanism under different scenarios. These activities can be carried out by taking advantage of the built-in use cases. Furthermore, new use cases can be added to model specific security concerns, e.g., related to new technologies. Finally, through a VPN-based integration mechanism, physical devices can be plugged in, and dually, HArMoNICS can be used to extend existing environments, such as simulators or digital twins.

In summary, among the features of HArMoNICS, the following ones are those we consider highly relevant for the security community.

- The case study is designed and implemented in order to mimic a real II, and therefore it includes a number of technologies that may appear within the perimeter of a smart infrastructure.
- Our case study includes 8 security and privacy use cases revolving around specific threats and weaknesses common to most II, including but not limited to: (i) software integrity and updates for end-point devices; (ii) privacy-preserving data management and processing; (iii) intrusion detection; (iv) protocol verification, and; (v) fog computing orchestration and hardening.
- The case study blueprint and implementation are open-source, with the possibility of extending and modifying them.
- The entire case study can be executed inside a publicly-available virtual machine, with minimal computational resources.

The rest of this paper is organized as follows. Section II describes the general features of architecture and networking related to the case study infrastructure. Section III details every single scenario included, with the relative security and privacy problem statement. Section IV offers a glimpse of the current state of the art in such a kind of digital replicas. Finally, Section V concludes the paper.

II. ARCHITECTURE OVERVIEW

The HArMoNICS case study is based on a smart building scenario, composed of both IT and OT elements. The scenario is inspired by the *Zero-Emission Building* (ZEB), which is hosted inside the Genoa University Campus, located in Savona (Italy). Figure 1 shows a sky view of the Campus,

¹www.sparta.eu



FIGURE 1. The Savona campus and the zero-emission building (in red).

with the ZEB location highlighted in red. Among the research infrastructures and facilities that the Campus hosts, there is a smart grid, called *Savona Polygeneration Microgrid (SPM)*, consisting of several nodes for the generation of power. Power generation nodes rely on different sources, e.g., solar panels and gas turbines, and their production partially supplies the internal energy demand.

ZEB is a smart building where innovative technologies and materials are adopted in order to optimize energy consumption with the ultimate goal of nullifying the carbon footprint. It contains several laboratories, offices, and a gym. Some servers and networks reside in the building and they host the services which contribute to the IT infrastructure. Standard and Fog-enabled access points provide wireless connectivity to network devices, including IoT ones. Also, sensors and actuators have been deployed to monitor the environment (e.g., room temperature) and reconfigure it (e.g., by opening a window).

A. NETWORK INFRASTRUCTURE

The network infrastructure of HARMoNICS is here described in more detail. Figure 2 shows the reference network scheme of the case study. Briefly, the smart building hosts a three-segment network. The three segments are *dmz*, *intranet* and *iot*. Servers hosting public services, i.e., those that are accessible from outside the network perimeter, are connected to *dmz*. Other servers, instead, and hosts are connected to *intranet*. Finally, *iot* is used for connecting various field devices, e.g., sensors, actuators, and Fog nodes. The smart building network resides behind a router firewall that delimits the perimeter with the external network, i.e., the public Internet.

Networks and nodes are labeled with their *symbolic* names, e.g., *ns* and *www*, and IP addresses. Node names are managed by the DNS service running on node *ns*. Network address space represents the interval of IP addresses that can be assigned to connected devices. For brevity, statically assigned IPs are only represented by the last address segments. For instance, the IP address of node *www* is $198.51.100.3$. If a node has no address label, its IP address is dynamically assigned. Finally, when relevant, connections are labeled with a specification of the used channel, e.g., IEEE 802.15.4.

The colored and numbered circles indicate that the nearby devices belong to one of the 8 security and privacy-related scenarios described, which will be detailed in the next Section.

The network topology described above is implemented by means of Docker containers and networks [6]. Roughly speaking, Docker containers are lightweight virtual machines running inside isolated Linux processes. Container connectivity is granted by means of virtual networks which are emulated by the host machine.

Figure 3 highlights the implementation details of the network infrastructure of HARMoNICS. Below, the core aspects are discussed.

- *Public network simulation.* Several devices connect with the scenario infrastructure by means of the public Internet (see Figure 2). Some of them are actual, remote entities (e.g., web servers), while others must be deployed within the scenario. To support this hybrid structure, the infrastructure relies on a *simulated internet* (Figure 3). The simulated internet is implemented by means of a router (*rt-simint*) which connects three networks, i.e., *ext*, *simint* and *outside*. Briefly, *ext* is connected with a virtual interface of the host platform. By means of a virtual bridge, the host interface provides direct connectivity with the real Internet.
- *VPN access.* To make the infrastructure extensible, direct access to each network is provided, by means of a virtual private network (VPN) server. The VPN server is connected to *ext*. In this way, the VPN can be accessed from software running on the host platform and even from the Internet. The server accepts connections on different ports, in the range [8886 – 8889]. Each port is uniquely mapped into one of the networks in the infrastructure. In Figure 3, port numbers are used to label (in red) the connection with the corresponding network. For instance, establishing a session with $172.16.255.100:8886$ will connect the VPN client to the intranet. The main purpose of the VPN is to permit the integration of physical devices, e.g., think of a wireless access point or an IoT device. Further examples are provided by some of the scenarios described in Section III.

HARMoNICS has been designed with a standard *Infrastructure-as-Code (IaC)* framework called TOSCA [7]. Docker Compose [8], the default Docker orchestrator, has been chosen as the scenario deployment tool. Finally,

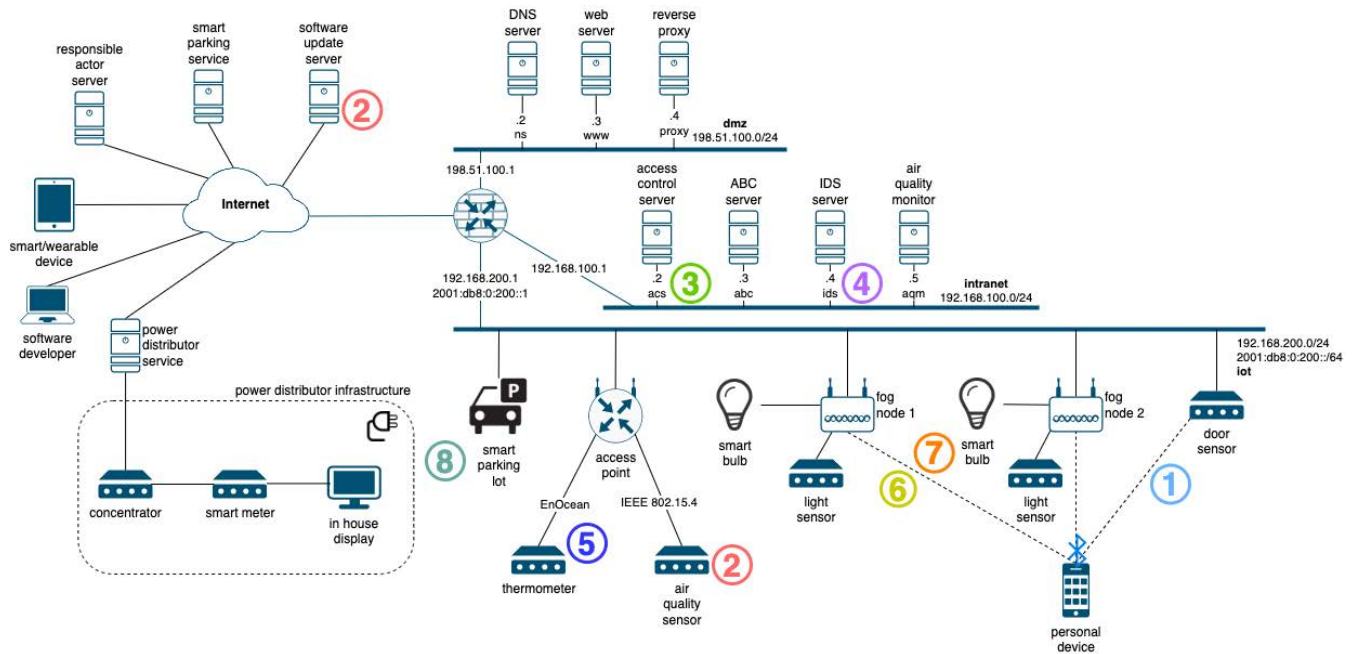


FIGURE 2. HArMoNICS general infrastructure overview with the 8 scenarios (described in Sections III.A–H).




FIGURE 3. Network infrastructure and VPN service overview.

the scenario code is hosted on a GitHub repository at <https://github.com/enricorusso/spartawp6>. In this way, by deploying HArMoNICS from the GitHub repository, continuous integration is also supported.

III. USE CASES

This Section provides a glimpse of the main use cases included in HArMoNICS. Each of the scenarios focuses on a different security concern and refers to a subsystem within the case study infrastructure. Furthermore, for each scenario, we carry out a discussion of open challenges and possible countermeasures that are being investigated in the context of the SPARTA project. The discussed techniques are relevant to highlight how security researchers can leverage on HArMoNICS for assessing their methodologies in a realistic setting.

A. SOFTWARE INTEGRITY FOR SMART SENSORS

One of the main threats to the security of II is the limited, or even absent, protection of the sensing devices that often form a significant part of the perimeter of the entire infrastructure. For example, the smart door sensor, placed within the *iot* network, has the task of monitoring the access of individuals inside the building (item 1 in Figure 2, in cyan). To do this, it pairs over Bluetooth with the mobile device of the person accessing it and requests the transmission of a unique identifier. An app is installed on the personal device, capable of interacting with the door sensor to pass it the person’s ID.

Given the reduced operation complexity and the low activation frequency, the door sensor is an embedded device on which firmware written in C language runs on bare metal.

C language permits low-level control without losing the advantages of high-level statements and data structures. Still, the manual management of data structures and memory pointers is often a source of vulnerabilities. The lack of memory safety capabilities (such as strong typing, present in other modern languages) enables attackers to exploit these flaws by maliciously altering the program behavior by partially or entirely hijacking control flow.

Reasonably, the most famous vulnerability for this scenario is *buffer overflow* [9], which is caused by increasing or decreasing a pointer without proper boundary checks on the data structure that is being accessed. This results in out-of-bounds writes that corrupt adjacent data areas, e.g., stack or heap. Similar problems may arise when indexing bugs are present in the code, i.e., boundary checks over an index for a given data structure are missing or incomplete. Indexing bugs are often caused by integer-related errors like an integer overflow, truncation or signedness bugs, or incorrect pointer casting.

In the present use case, the firmware may mistakenly consider the mobile app as a trusted actor. In particular, since the transmitted information has a fixed size in bytes, the sensor may not check the incoming messages, but just read bytes until the string termination symbol is received. An attacker can exploit this vulnerability in various ways, e.g., to steal data or enter the network. Moreover, this vulnerability is an enabling factor for, e.g., bypassing a possible *non-executable-stack* defense [10], and mounting a *Return-Oriented Programming (ROP)* attack [11], [12] [13].

Discussion: A possible solution to preserve control-flow integrity in unprotected devices is by using enforcement mechanisms, e.g., relying on security extensions that consist of adopting a runtime Policy Decision Point (PDP), also called *control-flow monitor*, and several Policy Enforcement Points (PEPs), inserted before runtime into the code. Because of the small capabilities of these devices, this extension must be as tiny as possible in terms of additional computational resources, and must not compromise the real-time nature of the execution. An example is given by [14], where a lightweight solution is presented to protect firmware running on a microcontroller through the use of an external FPGA, which implements a checker. Checks are triggered at critical branch points by a single instruction that invokes the FPGA. Here, the legitimacy of the control-flow transfer is checked through a completely parallel execution, and if necessary, CPU activity is interrupted if a violation is detected.

B. SOFTWARE UPDATE FOR IoT END DEVICES

Low-power indoor sensors (*AirMonitors*) continuously collect data about the air quality inside the ZEB, aiming to detect and prevent bad air quality situations, which could lead, e.g., to a higher risk of COVID transmission.

From the hardware point of view, the AirMonitor prototype present in HARMoNICS bundles a typical COTS System-on-Chip (SoC): an ARM Cortex-M microcontroller communicating with IEEE 802.15.4 low-power radio. This SoC is

connected via I2C/SPI bus on-board to a variety of sensors, for humidity, gas, and dust particles. The AirMonitor is networked via a low-power wireless access point, through which it communicates via IP protocols (6LoWPAN and CoAP) with a remote software update server (items 2 in Figure 2, in red). The software embedded in AirMonitors is based on RIOT [15], a popular open-source general-purpose operating system for low-power IoT devices.

The idea behind this scenario is to offer a test case for the security of software updates for low-power IoT devices. Over the last few years, the research community has been working on the definition of several IoT update processes [16], among which secure software updates for resource-constrained devices is a challenging research topic [17].

Discussion: The Internet Engineering Task Force (IETF) is currently defining a new standard for firmware updates called Software Updates for Internet of Things (SUIT) [18], [19]. The main goals of SUIT are interoperability and end-to-end security. The SUIT information model [20] defines a collection of security threats for the update process. Threats associated with this type of scenario are manifold. As discussed in [21], an attacker might update the IoT device with a modified and intentionally flawed firmware image. Alternatively, an attacker may replay a valid, but old (known-to-be-flawed) firmware, or a firmware update that is authentic, but for an incompatible device.

Although the SUIT model suggests a set of security requirements and countermeasures, it is worth noticing that all these threats are related to the integrity and the confidentiality of the update process only, while the content of the update is assumed to be trusted. Therefore, the SUIT workflow allows an Information System Management (ISM) service to upload a firmware image containing security vulnerabilities or malicious behaviors. Furthermore, as demonstrated by some recent work, the SUIT workflow is flexible in that it allows not only pre-quantum, but also post-quantum security [22], and does not only cater for full IoT firmware updates, but also for securing modular software updates on low-power IoT devices [23]. Last but not least, SUIT allows the ISM to transfer its authority to another entity, e.g., a third-party developer, that can deliver to the ISM some components of a software update (e.g., the executable of the application to be updated) or trigger the update process directly. In this case, the ISM has no mechanism to assess the content of the external software components and must trust the external entity.

C. PRIVACY-PRESERVING DATA PROCESSING

HARMoNICS includes an access control system that is managed by a dedicated server in the intranet subnetwork (item 3 in Figure 2, in light green). The system secures the access of persons inside the smart building. It also works as an access control system for drivers who want to use the parking lot of the campus. Access control is the process of mediating every request to data and resources owned by a system and determining whether the request should be granted or denied. In general, access control is a necessary condition to build

privacy in IoT solutions and to comply with the EU General Data Protection Regulation (GDPR) [24].

There are several risks associated with such a scenario, the following being the main ones.

- *Consent unawareness.* This risk relates to a user being unaware of the information disclosed to the system. The user could either provide too much information, e.g., allowing a malicious agent to retrieve her identity, or, on the contrary, inaccurate information, which can lead to wrong decisions or behaviors.
- *Policy and consent non-compliance.* This threat means that, even though the system shows its privacy policies to the users, there is no guarantee that the system actually complies with the advertised policies [25]. Therefore, although policies claim to prevent it, the user's personal data might still be revealed.
- *Information disclosure due to wrong design and/or implementation of access control.* The information disclosure threats expose personal information to unauthorized individuals [25]. This can happen in case the access control mechanisms in place are wrongly designed or implemented. Some modern approaches decompose access control in three main components: policy language, model, and enforcement [26]. Each of them may carry design/implementation errors.

A good privacy-preserving authentication system is an authenticated protocol that does not fully disclose the user's identity to a verifier. Only necessary pieces of the user identity (e.g., age, gender, membership, etc.) are provided during the verification phase. Furthermore, the authentication sessions should be mutually unlinkable, that is the protocol protects user identity and avoids profiling and tracking of the users.

Another aspect that needs to be taken into consideration within HArMoNICS is how data produced during the use of the II is being handled. The analysis of the usage and transaction logs of the II can provide valuable information on the characteristics and needs of the systems and their users, e.g. services demand, peak hours, and resource sufficiency. However, this type of analysis needs to be performed in a privacy-preserving way, so that the collected data is utilized while remaining protected both from data breaches and unauthorized uses.

Discussion: A promising privacy-enhancing technology that can be used for this purpose is searchable encryption (SE), supporting the storage of usage logs in encrypted form, while data remains available for processing [27]. This data processing system could be applied to the HArMoNICS smart building, as well as the smart parking lot and vehicle charging services. Desired properties of the SE service include:

- *Query expressiveness.* Support for complex, multi-keyword queries is required, in order to enable deriving useful conclusions from the data analysis.
- *Efficiency.* The query functionality needs to be efficient, in order to be applicable and practical in a real-world setting such as HArMoNICS.

- *Dynamic dataset support.* Dynamic updates of the encrypted dataset need to be supported in the system, to respond to the dynamic nature of the II.
- *Multi-client search support.* Enabling external entities to perform authorized queries on the encrypted dataset extends the utilization potential and value of the data processing service.

Searchable Symmetric Encryption (SSE) is a practical variation of searchable encryption providing a balance between efficiency, functionality, and security, while supporting the aforementioned requirements [28]. An open-source SSE implementation such as Clusion [29], based on the IEX SSE scheme [30], can be used for the implementation of the HArMoNICS privacy-preserving data processing system.

D. INTRUSION DETECTION

Intrusion detection is one of the main components of a global security strategy. In particular, it is typically considered the second line of defense against attacks. Reasonably, a complex II supported by a large-scale network, made of multiple and heterogeneous devices, is unlikely to be completely secure. Despite all the upstream efforts that are made during the design and development of a critical information system, malicious activities may succeed and compromise the confidentiality, availability, or integrity of the system during its life cycle. An Intrusion Detection System (IDS) aims at detecting such attacks against computer systems and networks. To deal with latent threats, an IDS continuously monitors the running system and analyses the gathered information to detect if an attack occurs or not. When the monitoring mechanism suspects that an attack has occurred (or is in progress), an alert is raised.

The presence of an IDS server within intranet network of HArMoNICS (item 4 in Figure 2, in purple) is useful to allow the investigation of currently trending challenges for the community. In particular, a crucial challenge is related to the availability of an appropriate dataset which is critical in the development of most IDSs. A bulk of state-of-the-art research does not provide reliable performance results since they rely on either the KDD99 or NSL-KDD benchmark datasets, which are concocted of traffic being over 20 years old. In this way, it does not represent recent attack scenarios and traffic behaviors. Obtaining traffic from simulated environments can help overcome this issue when merged with testing more recent datasets, such as the CICIDS 2017 [31]. Published datasets are available for different domains, such as industrial control systems (ICS) [32]. HArMoNICS can be used for assessing the effectiveness of datasets against real attack scenarios.

Discussion: An interesting direction for dealing with this scenario is to consider solutions using intrusion models that assume events to belong to a partially ordered set. For each participant of the distributed computation, the input trace describes the sequence of events that occurred locally in the participant's process. An event is a performed action, like sending or receiving a message, but also an internal

activity such as a system call. Clearly, for this model to apply, monitoring and logging code must be supported by all the participants, which is often the case.

Based on these traces, each distributed computation can be observed as a partially ordered set of events and be represented by a lattice of consistent cuts. This intermediate representation of learned normal behavior is potentially very large in size. Thus, it is used to infer smaller models that characterize the acceptable sequences of events. These models can take the form of an automaton or a list of temporal properties that have to be satisfied (likely invariants). Several types of models are constructed and used in parallel, as different models are often complementary during the detection phase. Also, using different types of models is key to reducing false negatives. Furthermore, to reduce false positives during the detection phase, the training phase must consider multiple distinct correct executions: the resulting models are obtained by combining the intermediate models defined during the learning of each normal execution.

E. IoT PROTOCOLS FLAWS

HARMoNICS includes two different scenarios related to the verification of IoT protocols. In the first scenario, the protocol EnOcean [33], mainly used in the smart building domain, has been considered (item 5 in Figure 2, in blue). Briefly, EnOcean is used to implement the communications between IoT devices interacting with a smart HVAC (Heating, Ventilation, and Air-Conditioning) system. Since IoT devices are provided by different manufacturers, design flaws in the EnOcean protocol may have dramatic effects on the correct behavior of the HVAC system. The HVAC hub server is located in the building area network (or local area network) and communicates through a gateway to the Internet and outside users. Via smartphone or tablet, the user can get access to the building network and she can monitor or configure the system. In case of an unusual event, the user will be notified immediately. Thus, an attacker could exploit flaws in the EnOcean protocol to carry out the following operations.

- Eavesdropping, i.e. spying on the system.
- Replay attack, where (parts) of messages are recorded to use it at a later stage.
- Man-in-the-middle attack, where the communication between two communication partners is intercepted, and potentially changed during transmission (modification attack).
- Denial-of-Service attack, i.e., preventing legitimate users from accessing the system.

The second scenario, instead, focuses on a risk analysis of two different smart building system configurations through threat modeling. This scenario carries three elements. The first one is modeling a system, e.g., with a threat modeling tool, in order to identify potential threats and vulnerabilities. The second element includes the attack scenario definition through the exploitation of selected vulnerabilities. The last one amounts to a risk analysis of the attack scenarios.

Discussion: Formal verification allows proving the correctness of a target protocol with respect to a certain specification or property *with mathematical rigor*. Possible checks include verification or falsification of security properties, functional correctness, qualitative and quantitative analysis of protocol specifications or implementations [34], [35], in presence of an attacker.

In the first scenario describe above, the formal verification of EnOcean can be based on ProVerif [36], [37], i.e., a protocol model checker which considers the well know Dolev-Yao attacker model [38]. The first step is the creation of an input model, e.g., based on the protocol specification. Since, in general, model checking the whole protocol specification is computationally expensive, often models include only the most critical parts in the protocol (e.g., initial authentication between the participants) to be verified against the relevant security goals (e.g., authentication happens correctly). Usually, the output of the formal verification process is either a proof of correctness or a potential vulnerability (a.k.a. a counterexample). Since ProVerif also handles unbounded protocol sessions, which are undecidable in general, it may return false positives, i.e., counterexamples/attacks that are not actually executable. For this reason, ProVerif also gives an attack derivation, which helps a human analyst to manually reconstruct an attack.

Instead, the second scenario can be treated via a probabilistic risk analysis of the two system configurations (see above) through model checking [39]. Since the problem to be modeled for this scenario is probabilistic, one may consider the Prism model checker [40]. Briefly, Prism models are specified through various types of Markov chains, where each transition occurs with a given probability. Assigned probabilities can correspond to the risk/likelihood score of a certain threat in the threat model.

F. FOG ORCHESTRATION SECURITY

The main idea behind this scenario is to check the placement of Fog and Edge devices and services for possible QoS and security-related issues and find the non-optimal distribution of services between Fog nodes. Two Fog nodes are physically placed in two different locations and both are connected to the IoT network (item 6 in Figure 2, in yellow). These Fog nodes are capable of running communication services to connect with related edge devices. Fog nodes host services that monitor the lighting characteristics in the rooms using light sensors, and are capable to adjust the lighting according to the preferences of the human by activating smart bulbs. Location and presence services are running on the Fog nodes to sense humans in the room and to monitor their exact position. A decision service “knows” the lighting preferences of the particular human, and controls smart bulbs according to the position of that human inside the room. For example, if the human sits on the couch near the TV, the lighting should be dimmed, etc.

The goal of this scenario is to provide a testbed for Fog orchestrators, and for measuring their ability to make

decisions on controlling the services according to the QoS and security requirements. Each decision of the orchestrator on starting/stopping/suspending/moving Fog services should be checked for the satisfaction of the minimal requirements imposed by various hardware and software restrictions of the involved physical devices, as well as requirements arising from the specifics of the area of application (e.g., health-related data should be protected better than environment monitoring data).

Orchestrators should follow three main steps:

- Each orchestrator should apply requirements on latency, bandwidth, security, and range imposed by the application area and hardware/software capabilities of each Fog node and decide if it is possible to start all required services without violating these requirements;
- Orchestrators should find the optimal distribution of the available services between different Fog nodes: this is useful for saving energy and computation resources in cases when some services may be stopped, suspended, or moved to another Fog nodes;
- Dynamic service allocation should be carried out. Dynamic allocation happens when the situation changes at runtime, and orchestrators must change the distribution of the services between available Fog nodes according to the new conditions.

Discussion: The technique for a Dynamic Service Orchestration, which addresses the issues discussed above, is presented in [41]. The control loop is developed on three main consecutive steps: Monitoring, Optimization, and Execution. Optimization aims to place find which placement of n available services in k Fog nodes makes a set of chosen QoS parameters optimal. The problem can be solved through a system of inequalities for parameters such that objective functions are minimized. Although, the objective functions are contradicting each other so there is no single solution to this multi-objective optimization problem that optimizes all the objective functions at the same time.

In fact, the process resolves into two further sub-steps. The integer Multi-Objective Particle Swarm Optimization (IMOPSO) method is used to find a set of Pareto optimal solutions. All service placements in this set are non-dominated (Pareto optimal), which means that each of them is better than all the other ones by at least one criterion. The second step is to choose the best solution from the Pareto optimal set by using the Analytical Hierarchy Process (AHP) [42]. AHP uses only a pairwise comparison of all alternatives by all objective functions, is easy to implement, and gives consistent results.

G. FOG HARDENING

As previously stated, Fog nodes play an important role in the IoT network of HARMoNICS. As a matter of fact, they serve edge computation for end-users in their proximity, sharing the load on the resources provided by cloud servers. However, this network topology optimization exposes the security of the user- and kernel-space software running on Fog nodes,

as they directly interact with the end-user device, which may be malicious (item 7 of Figure 2, in orange). Such exposure poses security risks that threaten the confidentiality of the data processed by user-space Fog applications, the integrity of the kernel-space Fog operating system, and, inherently, the whole Fog layer of the ZEB.

Attackers leverage memory corruption vulnerabilities to establish primitives for reading from or writing to the address space of a vulnerable application. These primitives form the foundation for code-reuse and data-oriented attacks [43]. The security enhancement should ensure the confidentiality of sensitive data, such as personal user information or user authentication material, that the user-space applications running on Fog nodes process. Moreover, the security extension should harden the underlying operating system kernel against data-oriented attacks, preventing an attacker from taking over the Fog node, which would allow her to have a foothold in the II network.

Discussion: With respect to the highlighted issues, virtualization extensions of modern CPUs could be leveraged to establish selective memory protection (xMP) primitives [44], that have the capability of thwarting data-oriented attacks. Such extensions, like Intel's Extended Page Table pointer (EPTP), offer the possibility to manage different views on guest-physical memory from inside a VM, without any interaction with a hypervisor. Therefore, selective protection of sensitive data in user or kernel space is obtained by isolating sensitive data in disjoint xMP domains.

In the specific use case scenario, such a solution can be used to harden the user-space decision service that the fog orchestrator deploys on fog nodes, which adjusts the light level based on user preferences received from their end device. Specifically, the end-user data that the service processes and stores in its address space can be isolated in a dedicated xMP domain. This way, the user information is prevented from being leaked in case an attacker exploits a memory corruption vulnerability that may emerge in the decision service.

H. MANAGING PERSONAL DATA IN VEHICLE RECHARGE PROCESS

HARMoNICS includes a vehicle recharge facility. Its main components are the user's personal device, the smart parking lot, and the power distributor infrastructure (item 8 of Figure 2, in dark green). The driver using her personal device initiates the charging process in the smart parking lot. Once the charging process is done, the power distributor (based on the smart meter) sends the charged energy amount to be paid.

The scenario focuses on the payment details and on the related security risks. The payment details may include (i) the driver's name, (ii) bank account and/or credit card information, and (iii) authorization to debit/credit the bank account for the service. The process is as follows: the driver submits her payment details to the parking lot, where the charging is happening. Once done, the parking lot initiates charging by sending an initiation command to the power distribution

infrastructure. Then, the infrastructure charges the car and sends information about the charged energy amount for the payment. Also, the power distribution infrastructure informs the parking lot that the charge is completed. After the parking lot sends the payment details to the smart building, the central management allows performing the payment transaction for the charged energy amount. Then, the lot is informed about the success of the payment and it sends the payment transaction receipt to the user's personal device.

A security issue related to the above scenario is that it is not clear how the driver's personal data, i.e., payment details, are handled during the transaction and whether the treatment of the personal data respects the principles of the GDPR. Verification tools and methodologies can help when designing and implementing data management processes like the one described above. For instance, the DPO tool [45] can help to assess how much the described process is compliant with the GDPR and recommend means to achieve this compliance [46]. For example, in the scenario above, it is possible to determine that the data owner is the driver and that her personal data, i.e., the payment details, are processed by the recharging infrastructure. Analysis of the process using the DPO tool results in a list of non-compliances, such as:

- Consent is missing (GDPR [24], Art. 7)
- Privacy policy is missing (GDPR [24], Art. 13, 14)
- No security measures are present (GDPR [24], Art. 25)
- Processing task is not being recorded (GDPR [24], Art. 30)

Discussion: The various non-compliances with regulations can be resolved through the adoption of several measures. For instance, to solve the issue of consent, the driver must provide consent for processing payment details to the infrastructure, and the infrastructure keeps consent for them. If it is not valid, the infrastructure informs the driver about the invalid consent; otherwise, it proceeds with sending permission to charge the vehicle.

The process can be made compliant with security and privacy requirements by adopting TLS protocol to send sensitive personal information. Then, public-key encryption must be applied so that payment details are encrypted before sending to the parking lot. The infrastructure decrypts them before performing the transaction.

Last, a specific task in the process must be allocated after each processing task to log details.

IV. RELATED WORK

HARMoNICS represents a case study that collects, within a single infrastructure, technologies that are related to some major security concerns in real environments. In fact, our proposal integrates both emulated technologies and real components, that replicate some attack vectors of interest. Furthermore, the assets of HARMoNICS are grouped within a virtual machine, open and downloadable by the community. For all these reasons, we believe HARMoNICS to be a distinguished and useful asset for the security community.

Although HARMoNICS is not meant to be a *digital twin* (DT) framework, it has a few similarities with this kind of system. Briefly, according to [47], [48] and [49], a DT is a digital replica of a real infrastructure, whose simulated execution is capable of generating the same amount of information as the original system [50]. In this respect, some of the scenarios of HARMoNICS can be seen as DTs of real systems. However, HARMoNICS is not designed to be generic and reconfigurable, but only for being extended with new scenarios.

Among the existing digital twins and security testbeds, EPICWTWIN [51] is a major proposal focusing on smart grids, where the users can deploy real-world attacks and countermeasures. It includes SCADA workstations, PLCs, end devices, and smart meters executed through a combination of simulation/emulation technologies, including VMs and Matlab-Simulink real-time models. Intuitively, even if EPICWTWIN is not an alternative to our proposal, it may be used to implement an infrastructure similar to that of HARMoNICS. To the best of our knowledge, such an implementation does not exist. Furthermore, DTs based on EPICWTWIN can be composed with HARMoNICS through the technologies discussed in Section II. Beyond EPICWTWIN, the reasoning discussed above also applies to other testbeds based on smart grids. Among them, some prominent examples are PRIME [52], the National SCADA testbed [53], and the infrastructure by the Mississippi State University SCADA Security Laboratory [54].

Some other authors have proposed systems for assessing smart infrastructures from a mainly functional point of view. As a consequence, security aspects are often neglected. Remarkably, [55] puts forward a demonstrator staged in a Campus of the West Cambridge university, which resembles the context of HARMoNICS. The work presents a detailed taxonomy for the layers on which an intelligent civil infrastructure must be based, sided by a data systematization model. Another example is [56], where a reference architecture for smart cities, also including a framework for responding to disastrous incidents, is given. Similar reference architectures are given in [57], [58], and [59]. Since they have a different target, none of these systems puts emphasis on IT security aspects and, thus, they cannot be directly compared with our work. In fact, they might even be composed with HARMoNICS to build larger case studies.

Finally, from the point of view of the formalization of security issues, [60] and [61] propose a systematic literature review on smart buildings and cities, respectively. Although these works do not provide an implementation, the security concerns gathered there partially overlap with those of HARMoNICS. Again, other security scenarios presented in these works can be integrated with HARMoNICS when implemented.

V. CONCLUSION

In this paper, we presented HARMoNICS, an open-source case study based on a virtual replica of a real intelligent

infrastructure located in the Savona Campus of the University of Genoa. HARMoNICS provides a series of vertical scenarios related to major security concerns of intelligent infrastructures, e.g., software integrity and upgradeability, privacy-preserving computing, and intrusion detection. The main goal of HARMoNICS is to become a strategic resource for the security community, which can rely on it for setting up experiments and benchmark activities. The infrastructure is currently used as an environment for the demonstration of security techniques (e.g., see [62]) developed in the context of the High-Assurance Intelligent Infrastructure Toolkit program of EU-funded project SPARTA.

In future work, we plan to leverage HARMoNICS for assessing the effectiveness of new methodologies aimed at increasing the security level of intelligent infrastructures. Moreover, we will consider further security scenarios to enrich HARMoNICS.

ACKNOWLEDGMENT

The authors would like to thank all colleagues not included in the list of authors but decisive for the activity: Katharina Hofer-Schmitz at Joanneum Research, Austria; Manon Knockaert and Jean-Marc Van Gyseghem at the University of Namur, Belgium; Petr Dzurenda at the Brno University of Technology, Czech Republic; Tewodros Beyene at the Fortiss Research Institute, Germany; Mari Seeba and Jake Tom at the University of Tartu, Estonia; Joaquin Garcia-Alfaro, Jean-Max Dutertre, Jean-Luc Danger, and Maryline Laurent at the Mines-Télécom Institute, France; Michel Hurfin and Ludovic Mé at the National Institute for Research in Computer Science and Automation, France; Paolo Prinetto and Alessandro Armando at the National Interuniversity Consortium for Informatics, Italy; Giorgio Bernardinetti, Francesco Mancini, and Gabriele Restuccia at the National Consortium for Telecommunications, Italy; Uwe Roth and Qiang Tang at the Luxembourg Institute of Science and Technology; and Marek Pawlicki at ITTI Institute, Poland.

REFERENCES

- [1] A. Holst. (2021). *Number of Internet of Things (IoT) Connected Devices Worldwide From 2019 to 2030*. Accessed: Oct. 27, 2022. [Online]. Available: <https://www.statista.com/statistics/1183457/iot-connected-devices-worldwide/>
- [2] N. Antonopoulos and L. Gillam, *Cloud Computing*. Berlin, Germany: Springer, 2010.
- [3] C. Matt, “Fog computing,” *Bus. Inf. Syst. Eng.*, vol. 60, no. 4, pp. 351–355, 2018.
- [4] A. Dogra, R. K. Jha, and S. Jain, “A survey on beyond 5G network with the advent of 6G: Architecture and emerging technologies,” *IEEE Access*, vol. 9, pp. 67512–67547, 2021.
- [5] J. C. S. Santos, K. Tarrit, and M. Mirakhorli, “A catalog of security architecture weaknesses,” in *Proc. IEEE Int. Conf. Softw. Archit. Workshops (ICSAW)*, Apr. 2017, pp. 220–223.
- [6] Docker. (2021). *What is a Container?—App Containerization*. Accessed: Oct. 27, 2022. [Online]. Available: <https://www.docker.com/resources/what-container>
- [7] P. Lipton, C. Lauwers, M. Rutkowski, C. Noshpitz, and C. Curescu, “TOSCA simple profile in YAML version 1.3,” OASIS, Burlington, MA, USA, Tech. Rep., Feb. 2020. [Online]. Available: <https://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.3/os/TOSCA-Simple-Profile-YAML-v1.3-os.pdf>
- [8] Docs Documentation. (2021). *Overview of Docker Compose*. Accessed: Oct. 27, 2022. [Online]. Available: <https://docs.docker.com/compose/>
- [9] (2019). *CWE-119: Improper Restriction of Operations Within the Bounds of a Memory Buffer*. Accessed: Oct. 27, 2022. [Online]. Available: <https://cwe.mitre.org/data/definitions/119.html>
- [10] M. Docs. (2021). *Data Execution Prevention—Win32 Apps*. Accessed: Oct. 27, 2022. [Online]. Available: <https://docs.microsoft.com/en-us/windows/win32/memory/data-execution-prevention>
- [11] R. Roemer, E. Buchanan, H. Shacham, and S. Savage, “Return-oriented programming: Systems, languages, and applications,” *ACM Trans. Inf. Syst. Secur.*, vol. 15, no. 1, pp. 1–34, 2012.
- [12] N. R. Weidler, D. Brown, S. A. Mitchel, J. Anderson, J. R. Williams, A. Costley, C. Kunz, C. Wilkinson, R. Wehbe, and R. Gerdes, “Return-oriented programming on a cortex-M processor,” in *Proc. IEEE Trust-com/BigDataSE/ICSS*, Aug. 2017, pp. 823–832.
- [13] N. R. Weidler, D. Brown, S. A. Mitchell, J. Anderson, J. R. Williams, A. Costley, C. Kunz, C. Wilkinson, R. Wehbe, and R. Gerdes, “Return-oriented programming on a resource constrained device,” *Sustain. Comput., Informat. Syst.*, vol. 22, pp. 244–256, Jun. 2019.
- [14] N. Maunero, P. Prinetto, G. Roascio, and A. Varriale, “A FPGA-based control-flow integrity solution for securing bare-metal embedded systems,” in *Proc. 15th Design Technol. Integr. Syst. Nanosc. Era (DTIS)*, Apr. 2020, pp. 1–10.
- [15] E. Baccelli, C. Gundogan, O. Hahm, P. Kietzmann, M. S. Lenders, H. Petersen, K. Schleiser, T. C. Schmidt, and M. Wahlisch, “RIOT: An open source operating system for low-end embedded devices in the IoT,” *IEEE Internet Things J.*, vol. 5, no. 6, pp. 4428–4440, Dec. 2018.
- [16] A. Kolehmainen, “Secure firmware updates for IoT: A survey,” in *Proc. IEEE Int. Conf. Internet Things (Things) IEEE Green Comput. Commun. (GreenCom) IEEE Cyber, Phys. Social Comput. (CPSCom) IEEE Smart Data (SmartData)*, Jul. 2018, pp. 112–117.
- [17] F. J. A. Padilla, E. Baccelli, T. Eichinger, and K. Schleiser, (2016). *The Future of IoT Software Must be Updated*. Accessed: Oct. 27, 2022. [Online]. Available: <https://hal.inria.fr/hal-01369681/document>
- [18] IETF. (2019). *Website of: IETF SUIT Draft Architecture*. Accessed: Oct. 27, 2022. [Online]. Available: <https://tools.ietf.org/html/draft-ietf-suit-architecture>
- [19] B. Moran, H. Tschofenig, H. Birkholz, and K. Zandberg, “A concise binary object representation (CBOR)-based serialization format for the software updates for Internet of Things (SUIT) manifest,” Internet Draft, French Nat. Inst. Res. Comput. Automat., Paris, France, Tech. Rep., 2021. [Online]. Available: <https://hal.inria.fr/hal-03137962/document>
- [20] IETF. (2019). *Website of: Firmware Updates for Internet of Things Devices—An Information Model for Manifests Draft-IETF-Suit-Information-Model-03*. Accessed: Oct. 27, 2022. [Online]. Available: <https://tools.ietf.org/html/draft-ietf-suit-information-model-03>
- [21] K. Zandberg, K. Schleiser, F. Acosta, H. Tschofenig, and E. Baccelli, “Secure firmware updates for constrained IoT devices using open standards: A reality check,” *IEEE Access*, vol. 7, pp. 71907–71920, 2019.
- [22] G. Banegas, K. Zandberg, E. Baccelli, A. Herrmann, and B. Smith, “Quantum-resistant software update security on low-power networked embedded devices,” in *Proc. Int. Conf. Appl. Cryptogr. Netw. Secur.* Cham, Switzerland: Springer, 2022, pp. 872–891.
- [23] K. Zandberg, E. Baccelli, S. Yuan, F. Besson, and B. Smith, “Femto-Containers: Lightweight virtualization and fault isolation for small software functions on low-power IoT microcontrollers,” in *Proc. Int. Middleware Conf. (MIDDLEWARE)*, 2022, pp. 1–12.
- [24] European Union. (2016). *Regulation (EU) 2016/679 of the European Parliament (General Data Protection Regulation)*. Accessed: Oct. 27, 2022. [Online]. Available: <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- [25] M. Deng, K. Wuyts, R. Scandariato, B. Preneel, and W. Joosen, “A privacy threat analysis framework: Supporting the elicitation and fulfillment of privacy requirements,” *Requirements Eng.*, vol. 16, no. 1, pp. 3–32, Mar. 2011.
- [26] S. D. C. D. Vimercati, S. Foresti, P. Samarati, and S. Jajodia, “Access control policies and languages,” *Int. J. Comput. Sci. Eng.*, vol. 3, pp. 94–102, Nov. 2007.
- [27] C. Bösch, P. Hartel, W. Jonker, and A. Peter, “A survey of provably secure searchable encryption,” *ACM Comput. Surveys*, vol. 47, no. 2, pp. 1–51, Jan. 2015.
- [28] Q. Gan, C. Zuo, J. Wang, S.-F. Sun, and X. Wang, “Dynamic searchable symmetric encryption with forward and backward privacy: A survey,” in *Network and System Security (Network and System Security)*, J. K. Liu and X. Huang, Eds. Cham, Switzerland: Springer, 2019, pp. 37–52.

- [29] (2021). *Clusion Library*. Accessed: Oct. 27, 2022. [Online]. Available: <https://github.com/encryptedsystems/Clusion>
- [30] S. Kamara and T. Moataz, "Boolean searchable symmetric encryption with worst-case sub-linear complexity," in *Proc. Annu. Int. Conf. Theory Appl. Cryptograph. Techn.* Cham, Switzerland: Springer, 2017, pp. 94–124.
- [31] S. Borah, R. Panigrahi, and A. Chakraborty, "An enhanced intrusion detection system based on clustering," in *Advances in Intelligent Systems and Computing*. Singapore: Springer, Dec. 2017, pp. 37–45.
- [32] R. C. B. Hink, J. M. Beaver, M. A. Buckner, T. Morris, U. Adhikari, and S. Pan, "Machine learning for power system disturbance and cyber-attack discrimination," in *Proc. 7th Int. Symp. Resilient Control Syst. (ISRCSS)*, Aug. 2014, pp. 1–8.
- [33] S. Marksteiner, V. J. E. Jimenez, H. Valiant, and H. Zeiner, "An overview of wireless IoT protocol security in the smart home domain," in *Proc. Internet Things Bus. Models, Users, Netw.*, 2017, pp. 1–8.
- [34] K. Hofer-Schmitz and B. Stojanovic, "Towards formal methods of IoT application layer protocols," in *Proc. 12th CMI Conf. Cybersecurity Privacy (CMI)*, Nov. 2019, pp. 1–6.
- [35] K. Hofer-Schmitz and B. Stojanović, "Towards formal verification of IoT protocols: A review," *Comput. Netw.*, vol. 174, Jun. 2020, Art. no. 107233.
- [36] B. Blanchet, "An efficient cryptographic protocol verifier based on prolog rules," in *Proc. 14th IEEE Comput. Secur. Found. Workshop*, Jun. 2001, pp. 82–96.
- [37] K. Hofer-Schmitz, "A formal analysis of EnOcean's teach-in and authentication," in *Proc. 16th Int. Conf. Availability, Rel. Secur.*, Aug. 2021, pp. 1–8.
- [38] D. Dolev and A. C. Yao, "On the security of public key protocols," *IEEE Trans. Inf. Theory*, vol. IT-29, no. 2, pp. 198–208, Mar. 1983.
- [39] H. Vallant, B. Stojanovic, J. Božić, and K. Hofer-Schmitz, "Threat modelling and beyond-novel approaches to cyber secure the smart energy system," *Appl. Sci.*, vol. 11, no. 11, p. 5149, Jun. 2021.
- [40] M. Kwiatkowska, G. Norman, and D. Parker, "PRISM: Probabilistic symbolic model checker," in *Proc. Int. Conf. Model. Techn. Tools Comput. Perform. Eval.* Cham, Switzerland: Springer, 2002, pp. 200–204.
- [41] N. Morkevicius, A. Venckauskas, N. Šatkauskas, and J. Toldinas, "Method for dynamic service orchestration in fog computing," *Electronics*, vol. 10, no. 15, p. 1796, Jul. 2021.
- [42] T. L. Saaty and L. G. Vargas, "The seven pillars of the analytic hierarchy process," in *Models, Methods, Concepts Applications of the Analytic Hierarchy Process*. Boston, MA, USA: Springer, 2012, pp. 23–40.
- [43] L. Szekeres, M. Payer, T. Wei, and D. Song, "SoK: Eternal war in memory," in *Proc. IEEE Symp. Secur. Privacy (SP)*, May 2013, pp. 48–62.
- [44] S. Proskurin, M. Momeu, S. Ghavamnia, V. P. Kemerlis, and M. Polychronakis, "XMP: Selective memory protection for kernel and user space," in *Proc. IEEE Symp. Secur. Privacy (SP)*, May 2020, pp. 563–577.
- [45] (2020). *DPO Tool—Evaluate Business Process*. Accessed: Oct. 27, 2022. [Online]. Available: <https://dpotool.cs.ut.ee/>
- [46] R. Matulevicius, J. Tom, K. Kala, and E. Sing, "A method for managing GDPR compliance in business processes," in *Proc. Int. Conf. Adv. Inf. Syst. Eng.* Cham, Switzerland: Springer, 2020, pp. 100–112.
- [47] A. Bolton, M. Enzer, and J. Schooling, *The Gemini Principles: Guiding Values for the National Digital Twin and Information Management Framework*. Cambridge, U.K.: Centre for Digital Built Britain.
- [48] GE Digital. (2017). *Digital Twins: The Bridge Between Industrial Assets and the Digital World*. Accessed: Oct. 27, 2022. [Online]. Available: <https://www.ge.com/digital/blog/digital-twins-bridge-between-industrial-assets-and-digital-world>
- [49] H. V. M. Catapult, "Feasibility of an immersive digital twin: The definition of a digital twin and discussions around the benefit of immersion," VR Forum, HVM Catapult Vis., U.K., 2018. [Online]. Available: https://www.amrc.co.uk/files/document/219/1536919984_HVM_CATAPULT_DIGITAL_TWIN_DL.pdf
- [50] M. Grieves and J. Vickers, "Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems," in *Transdisciplinary Perspectives on Complex Systems*. Berlin, Germany: Springer, 2017, pp. 85–113.
- [51] N. K. Kandasamy, S. Venugopalan, T. K. Wong, and L. J. Nicholas, "EPICTWIN: An electric power digital twin for cyber security testing, research and education," 2021, *arXiv:2105.04260*.
- [52] T. Becejac, C. Eppinger, A. Ashok, U. Agrawal, and J. O'Brien, "PRIME: A real-time cyber-physical systems testbed: From wide-area monitoring, protection, and control prototyping to operator training and beyond," *IET Cyber-Phys. Syst., Theory Appl.*, vol. 5, no. 2, pp. 186–195, Jun. 2020.
- [53] K. Barnes and B. Johnson, "National SCADA test bed substation automation evaluation report," Idaho National Lab., Daho Falls, ID, USA, Tech. Rep. INL/EXT-09-15321 TRN: US201001%345, 2009.
- [54] T. Morris, A. Srivastava, B. Reaves, W. Gao, K. Pavurapu, and R. Reddi, "A control system testbed to validate critical infrastructure protection concepts," *Int. J. Crit. Infrastruct. Protection*, vol. 4, no. 2, pp. 88–103, 2011.
- [55] Q. Lu, A. K. Parlakad, P. Woodall, G. D. Ranasinghe, X. Xie, Z. Liang, E. Konstantinou, J. Heaton, and J. Schooling, "Developing a digital twin at building and city levels: Case study of west Cambridge campus," *J. Manag. Eng.*, vol. 36, no. 3, May 2020, Art. no. 05020004.
- [56] D. N. Ford and C. M. Wolf, "Smart cities with digital twin systems for disaster management," *J. Manag. Eng.*, vol. 36, no. 4, Jul. 2020, Art. no. 04020027.
- [57] J. Dutta and S. Roy, "IoT-fog-cloud based architecture for smart city: Prototype of a smart building," in *Proc. 7th Int. Conf. Cloud Comput., Data Sci. Eng.*, Jan. 2017, pp. 237–242.
- [58] G. Alsuhli and A. Khattab, "A fog-based IoT platform for smart buildings," in *Proc. Int. Conf. Innov. Trends Comput. Eng. (ITCE)*, Feb. 2019, pp. 174–179.
- [59] Z. Chevallier, B. Finance, and B. C. Boulakia, "A reference architecture for smart building digital twin," in *Proc. ESWC*, 2020, pp. 1–12.
- [60] P. Ciholas, A. Lennie, P. Sadigova, and J. M. Such, "The security of smart buildings: A systematic literature review," 2019, *arXiv:1901.05837*.
- [61] J. Laufs, H. Borrión, and B. Bradford, "Security and the smart city: A systematic review," *Sustain. Cities Soc.*, vol. 55, Apr. 2020, Art. no. 102023.
- [62] L. Malina, P. Dzurenda, S. Ricci, J. Hajny, G. Srivastava, R. Matulevicius, A.-A.-O. Affia, M. Laurent, N. H. Sultan, and Q. Tang, "Post-quantum era privacy protection for intelligent infrastructures," *IEEE Access*, vol. 9, pp. 36038–36077, 2021.




GIANLUCA ROASCIO received the master's degree in embedded systems, in 2018. He is currently pursuing the Ph.D. degree with the Department of Control and Computer Engineering, Polytechnic of Turin. He is a member of the Cyber Security National Laboratory of CINI and the Italian National Consortium for Informatics. Since 2020, he is a Lecturer of Computer Security at the Italian Army Academy. His research interests include security of hardware devices and the

IoT, memory protection, and control-flow integrity.



GABRIELE COSTA (Member, IEEE) received the Laurea degree in computer science, in 2007 and the Ph.D. degree in computer science, in 2011. He is currently an Associate Professor at the SySMA Group, IMT School for Advanced Studies. He has been a Co-Founder of the Computer Security Laboratory (CSEC), DIBRIS, and CRO of Talos, a startup focusing on cyber security. He is a member of the Cyber Security National Laboratory of CINI and also a member of the Cyber Security Group of the Istituto di informatica e Telematica of the CNR. His appointments include a period as Visiting Researcher at ETH Zurich, from 2016 to 2017.



part in the project-team TRiBE. Since 2013, he is also a Co-Founder and a Coordinator of the Open Source Community Developing RIOT Operating System.

EMMANUEL BACCELLI received the Ph.D. degree from École Polytechnique, Paris, in 2006, the habilitation degree from Université Pierre et Marie Curie, in 2012. He is currently a Professor for “Open and Secure IoT Ecosystem” at Freie Universität Berlin, in partnership with the French National Research Institute for Digital Sciences (Inria) and the Einstein Center Digital Future (ECDF). Since 2007, he has been with Inria as a Scientific Researcher, where he currently takes



cyber security and information systems security.

NERIJUS MORKEVIČIUS received the Ph.D. degree in computer science from the Department of Applied Mathematics, Kaunas University of Technology (KTU), in 2002. He is currently an Associate Professor with KTU. He has taken part in several national and EU-funded research projects, has a number of publications in high-impact international conferences and journals. His main research interests include information management systems, IoT technologies, and



Tarragona, Spain, in 2015; St. Petersburg ITMO, Russia, in 2017; and KU Leuven, Belgium, in 2017. He has been leading Privacy-by-Design Task in SPARTA H2020 Project.

LUKAS MALINA received the M.Sc. (Hons.) and Ph.D. degrees from the Brno University of Technology (BUT), in 2010 and 2014, respectively. He earned the Dean Prize for the master’s thesis. He is currently a Senior Researcher with the Department of Telecommunications, BUT. He has published more than 70 articles in international journals and conferences and has provided several invited research and teaching lectures abroad, such as University of Tampere, Finland, in 2013; URV



at the University of Genova, and is personally involved in different live-fire cyber exercises as a member of the Green/Blue team.

ENRICO RUSSO received the M.Sc. and Ph.D. degrees in computer science, in 2001 and 2021, respectively. He is currently an Assistant Professor in computer engineering with the University of Genova. His work is focused on Cyber Range systems with particular emphasis on investigating techniques for simplifying the generation of the training environments and for automatizing the tasks of personnel involved in exercises. He also co-founded ZenHack, the Capture the Flag Team



CAiSE. He is an author of the book *Fundamentals of Secure System Modelling* (Springer, 2017). He has been involved in the SPARTA H2020 Project (task: Privacy-by-Design).

RAIMUNDAS MATULEVIČIUS received the Ph.D. degree in computer and information science from the Norwegian University of Science and Technology. He is currently a Professor of information security with the University of Tartu, Estonia. His publication record includes more than 100 articles published in peer-reviewed journals, conferences, and workshops. He has been a Program Committee Member at international conferences, such as NordSec, PoEM, REFSQ, and



interests include cyber security, artificial intelligence, pervasive computing, biometrics, and computer vision. Her work focuses on formal modeling of protocols for the IoT. With this topic, she has been included in the “Secure Orchestration” task of the SPARTA H2020 Project.

BRANKA STOJANOVIĆ received the bachelor’s degree in electrical engineering and the Ph.D. degree in electrical engineering and computer science from the School of Electrical Engineering (ETF), University of Belgrade, Serbia. She is currently a Key Researcher and the Deputy Head of the Cyber Security and Defence Research Group at Joanneum Research Digital—Institute for Information and Communications Technologies, Austria. She is CISSP certified. Her research



MARIUS MOMEU received the master’s degree, in 2020. He is currently pursuing the Ph.D. degree with the Technical University of Munich. His research interests include software memory corruption vulnerabilities, with particular emphasis on operating system kernel hardening, automated testing through symbolic execution and software mitigations for architectural flaws, and secure e-voting. He has been leading the “Hardening Legacy Components” task in SPARTA H2020 Project.



U.K., and France. Her research interests include work on anonymization techniques for real-world datasets, privacy-preserving utilization of personal data within personalized services, and context-aware systems.

AIMILIA TASIDOU received the B.Sc. degree in informatics from the University of Piraeus, Greece, the M.Sc. degree in artificial intelligence from the University of Edinburgh, U.K., and the Ph.D. degree in data privacy from the Department of Electrical and Computer Engineering, Democritus University of Thrace, Greece. She works as a Post-Doctoral Researcher at Télécom SudParis, Polytechnic Institute of Paris, France. She has worked on several research projects, in Greece,

...

Open Access funding provided by ‘Politecnico di Torino’ within the CRUI CARE Agreement