Next Generation Edge-Cloud Continuum Architecture for Structural Health Monitoring

Lorenzo Gigli, *Graduate Student Member, IEEE*, Ivan Zyrianoff, *Graduate Student Member, IEEE*, Federica Zonzini, *Member, IEEE*, Denis Bogomolov, Nicola Testoni, *Member, IEEE*, Marco Di Felice, *Member, IEEE*, Luca De Marchi, *Member, IEEE*, Giuseppe Augugliaro, Canio Mennuti, and Alessandro Marzani, *Member, IEEE*

***Abstract*—Assessing the integrity of industrial and civil appliances has become a priority worldwide. Noteworthy, this goal requires a strong synergy between multiple tools, disciplines, and approaches to be attained via a joint hardware-software co-design of the different Structural Health Monitoring (SHM) system components. This work proposes the** MAC4PRO **architecture, a sensor-to-cloud monitoring platform that seamlessly integrates sensing and software technologies for accurate data measurement, transmission, and analysis. The developed solution stands out for its interoperability and versatility, making it a promising candidate for integration in the next generation of smart structures. Our platform was validated during extensive experimental campaigns targeted at various industrial scenarios. The results show that the** MAC4PRO **architecture can identify subtle changes, such as 1mm size leakage events in pipeline circuits, or less than 1% frequency drifts in civil buildings after seismic excitation, while ensuring more than 90% reduction in the edge-to-cloud data transfer process.**

***Index Terms*—Edge-cloud continuum, Internet of Things, interoperability, structural health monitoring.**

Manuscript received 28 September 2023; accepted 20 November 2023. Date of publication 25 December 2023; date of current version 4 April 2024. This work was supported by INAIL within the BRIC/2018, ID = 11 framework under Grant MAC4PRO "Smart maintenance of industrial plants and civil structures via innovative monitoring technologies and prognostic approaches." Paper no. TII-23-3773. *(Lorenzo Gigli and Ivan Zyrianoff contributed equally to this work.) (Corresponding author: Ivan Zyrianoff.)*

Lorenzo Gigli, Ivan Zyrianoff, and Marco Di Felice are with the Advanced Research Center on Electronic Systems for Information and Communication Technologies "Ercole De Castro, University of Bologna, 40136 Bologna, Italy, and also with the Department of Computer Science and Engineering, University of Bologna, 40136 Bologna, Italy (e-mail: lorenzo.gigli@unibo.it; ivandimitry.ribeiro@unibo.it; marco.difelice3@unibo.it).

Federica Zonzini, Nicola Testoni, and Luca De Marchi are with the Advanced Research Center on Electronic Systems for Information and Communication Technologies "Ercole De Castro, University of Bologna, 40136 Bologna, Italy, and also with the Department of Electrical, Electronic, and Information Engineering, University of Bologna, 40136 Bologna, Italy (e-mail: federica.zonzini@unibo.it; nicola.testoni@unibo.it; l.demarchi@unibo.it).

Denis Bogomolov is with the Advanced Research Center on Electronic Systems for Information and Communication Technologies "Ercole De Castro, University of Bologna, 40136 Bologna, Italy (e-mail: denis.bogomolov2@unibo.it).

Giuseppe Augugliaro and Canio Mennuti are with INAIL Technological Innovations Department, 00078 Rome, Italy (e-mail: g.augugliaro@inail.it; c.mennuti@inail.it).

Alessandro Marzani is with the Department of Civil, Chemical, Environmental, and Materials Engineering, University of Bologna, 40136 Bologna, Italy (e-mail: alessandro.marzani@unibo.it).

Color versions of one or more figures in this article are available at https://doi.org/10.1109/TII.2023.3337391.

Digital Object Identifier 10.1109/TII.2023.3337391

I. Introduction

In recent years, the *smart structures* paradigm [1] has emerged as a novel and practical approach to assessing the integrity of industrial and civil assets. The definition comes from the fact that, in the next generation, engineered structures will be equipped with *intelligent sensor systems* featuring on-board and advanced *decision-making functionalities*. Hence, implementing such structural health monitoring (SHM) architectures requires perfect coordination among the sensing, communication, and decision subsystems to achieve a timely and reliable diagnosis [2]. This progress has been made possible by the vertical coalescence between the research contributions in sensing technologies, data science, and signal processing [3].

More in detail, the effectiveness of the SHM systems is based on the optimal integration between the required hardware resources for signal acquisition, conditioning, and digitalization, and the associated software infrastructure in charge of data management, data analytics, and visualization. Such integration must consider the unique requirements posed by the SHM application context. On the one side, resilient monitoring strategies are needed to ensure operational serviceability in the presence of faulty devices. On the other hand, the developed architectures must handle the heterogeneity of sensing units, which may differ based on the type of sensed signals (e.g., accelerations, strains, displacements), data formats, and acquisition protocols [4]. Besides, managing SHM data is another pivotal challenge since—in multiple deployment cases—the acquired information may exhibit all the four dimensions of big data (volume, variety, velocity, and veracity) [5]. Depending on the application context, data management policies could be implemented at any stage of the computational continuum, such as on *the extreme edge*, when damage-sensitive features are processed directly on the sensors' board, on *the edge* (i.e., on computational units nearby the sensors' board) or on *the cloud* for scenarios involving burdensome postprocessing phases [3], [6], [7]. Deployment scenarios often entail noncompatible infrastructure, posing challenges regarding adaptability, and the distribution of software components across the edge-cloud continuum presents tradeoffs. Executing processing at the edge reduces the workload for subsequent components, resulting in bandwidth and processing efficiencies. However, edge devices may have lower robustness compared to cloud resources.

To meet these requirements, some research studies on smart structures propose multilayered, Internet of Things (IoT)-based architectures involving both smart sensor devices in charge of measuring, preprocessing, and forwarding physical data and remote processing units, which merge and handle the huge data volume, finally executing structural assessment algorithms [8], [9], [10], [11], [12]. While the referred solutions have achieved considerable results, two open points still need to be tackled. First, some research works fail to consistently deploy all the cyber–physical components, as is the case of [8], [10], [12], or vice versa; they fail in test-fielding the software components on real-world SHM scenarios [11]. Another issue is the lack of generality since most of the proposed architectures fit a specific SHM use case, hence being barely extendable to support different sensing, processing, and monitoring tasks.

In this article, we advance the state-of-the-art of multilayered SHM-IoT architectures by presenting the results of the MAC4PRO[1] project, a research effort aimed at developing an infrastructure-agnostic and general-purpose monitoring platform for the condition assessment of industrial and civil infrastructures leveraging the ultimate technologies delivered by the information, software, and industrial engineering communities. Three main contributions of the MAC4PRO project are presented in this article as follows.

1) We illustrate a generic and modular IoT multilayered architecture for SHM applications. It includes four layers (sensing, interoperability, data management, and service) to handle generic SHM scenarios. Pursuing this goal is achieved by abstracting from the structure's characteristics and supporting the sensing units' heterogeneity.
2) We present a complete implementation of the abstract architecture, which meets the aforementioned requirements. As a sensing layer, we designed a distributed monitoring network based on redundant, low-cost sensing units with embedded signal processing support. The other layers are addressed by the MODRON software, a general-purpose SHM platform for heterogeneous data acquisition, storage, visualization, and third-party software integration.
3) We validate both the architecture and the platform versatility in two distinct experimental campaigns related to the condition monitoring of real facilities in their operative environment. The first testbed refers to the monitoring of a concrete building during seismic events simulated through a shaking table. The second use case pertains to identifying leakage in hydraulic circuits via acoustic emissions (AEs). For both campaigns, we discuss the deployment of the architectural components in the edge-cloud continuum and present diagnostic results. In addition, we conducted a comprehensive performance evaluation to assess the architecture capabilities.

The rest of this article is organized as follows. An overview of the current sensing solutions and preliminary works on software functionalities of IoT-SHM systems are reviewed in Section II. Section III encloses a thorough description of the MAC4PRO architecture, introducing the specificity of each layer and providing a general description of its defining features, which is complemented by Section IV that provides the implementation details of the software components describes. Section V presents the performance evaluation conducted for the proposed architecture. Section VI is dedicated to an extensive experimental validation phase on two representative benchmark scenarios for condition monitoring. Finally, Section VII concludes this article.

II. Related Works

An extensive analysis of the sensor technologies and the practical issues to be faced in the deployment process of SHM networks is conducted in the critical work by Pengfei et al. [13]. In general terms, the physical parameters to be acquired vary largely depending on the nature of the structural response to be monitored. In particular, the dynamic behavior is primarily dictated by vibrations, such as those induced by operational or environmental actions [14]. Consequently, sensing technologies measuring vibration-related quantities (e.g., accelerations, angular velocities) are usually employed to accurately extract and characterize dynamic-dependent properties, typically consisting of frequency of vibrations, damping coefficients, and mode shapes. To this end, microelectromechanical system (MEMS) sensors have recently gained increasing attention for the successful installation of low-cost instrumentation [15]. At the same time, it is well proven that industrial and civil structures are subjected to important degradation processes over time [16], such as corrosion, cracks, and delaminations, which manifest slowly; therefore, they are considered static mechanisms and involve entirely different measuring technologies.

Apart from the mere physical principle at the basis of each sensing solution, a paramount aspect that needs to be considered in practical scenarios concerns the ease of deployment they allow for, namely the possibility to interface them with wired or wireless communication links. At the same time, the effective implementation of SHM systems also depends on other critical aspects, which include the sensor placement, i.e., the locations at which sensors are installed based on the local or global information they can and should capture [17].

In addition to sensor selection and placement, SHM systems must implement other specialized data processing and management components. To this aim, several recent studies investigate the integration of the IoT paradigm for asset monitoring. Besides allowing for connecting smart sensors to the Internet, such integration consists of deploying software architectures in the continuum (i.e., edge and cloud environments) to collect, store, and analyze SHM sensor data. A reference architecture is discussed in [8], including four different components: i) smart objects, ii)

[1][Online]. Available: https://site.unibo.it/mac4pro/

TABLE I
COMPARISON OF MAC4PRO ARCHITECTURE TO THE LITERATURE

Solution	F1	F2	F3		F4
			F3.1	F3.2	
[11]	✗	✓	✗	✓	✗
[20]	✗	✓	✓	✗	✓
[21]	✗	✓	✗	✓	✗
[22]	✗	✗	✓	✗	✗
[23]	✗	✓	✗	✗	✓
MAC4PRO	✓	✓	✓	✓	✓

gateway, iii) cloud, and iv) remote station for data access and visualization. Two IoT-SHM use cases related to safeguarding and protecting masonry buildings are presented, although the development of the software component is still in a preliminary stage. As previously observed, structural assessment through nondestructive techniques may involve collecting a large amount of data due to high-frequency sensor sampling and long measuring periods necessary for high-quality information retrieval [18]. For this reason, cloud infrastructures often provide the storage and computational resources of IoT-SHM platforms. In [9], the authors employed the AWS IoT cloud platform to manage smart sensors installed on a single-line railway bridge; the raw sensor data are stored with DynamoDB and displayable through a custom web interface. Similar sensor-to-cloud workflows are proposed and test-fielded in [10] and [19]. Finally, in [12] the authors proposed a digital twin (DT) framework for SHM that leverages edge nodes for data cleaning and preprocessing tasks. However, this article focuses on the DT's modeling rather than on the underlying platform supporting it.

We highlight that the mentioned approaches propose architectures tailor-made for a specific use case; moreover, the application components are usually bound to a specific processing location—e.g., the cloud. Much less attention has been dedicated to designing general SHM software architectures that can abstract from domain-specific industrial and civil engineering information and support heterogeneous sensor devices. We review the most promising solutions available in the literature (see Table I) and highlight the differences concerning our work by focusing on the following features.

1) Interoperability (F1): the support to easily integrate dissonant interfaces and data structures from sensors and software components into the system.
2) Modularity (F2): the ability to divide the system into independent modules that can be developed, maintained, and managed separately.
3) Agnostic-design (F3):
   - Infrastructure (F3.1): the capacity to remain independent of the underlying hardware infrastructure in various deployment scenarios, accommodating different edge, and cloud computing node configurations.
   - Application (F3.2): the flexibility to work with various software configurations without being tied to specific technologies or implementations.
4) Edge-cloud continuum support (F4): the capability to seamlessly operate in edge and cloud computing environments.

In [11], the authors investigated the advantages of service-oriented architectures for SHM systems, given the fact that independent and reusable software components can be composed for the specific scenario; based on these premises, the MAC4PRO architecture developed in the this article follows a similar approach; however, we consider the whole edge-cloud continuum and do not centralize the computation solely in the cloud. In [20], it is proposed an SHM architecture where software components are distributed across the edge-cloud continuum. They adopt a modular design capable of accommodating various infrastructures. Our architecture differs by prioritizing interoperability and maintaining flexibility in the software components, which are not bound to particular technology implementations. A collection and classification of several SHM systems deployed on bridges was presented in [21]. The authors analyze the various sensing and communication technologies and the most common data processing algorithms for early warning. Although many aspects were covered, the replicability of software architectures was not considered, nor was the potential of the edge-cloud continuum. The work in [22] presented the integration of a bridge SHM system with a BIM model, enabling a bridge 3-D model to directly access bridge health data in real-time. IoT sensors were deployed, communicating through WiFi to an IoT web platform. The limitation of this work is that it focuses on the BIM aspects of the system and the specifics of bridge modeling rather than on the software architecture to support it. In the work presented in [23], an architecture was designed and implemented to support optical fiber sensors for SHM. This architecture enables real-time data processing using an Apache Kafka-based stream processing system. Their system was custom-built to meet stringent time constraints while computing high data volumes. Consequently, the architecture is tightly coupled with its specific technological implementation and underlying infrastructure, with limited consideration for interoperability. In contrast to other architectures, our approach prioritizes interoperability, simplifying the integration of new sensors and software components and enabling efficient updates. Therefore, this design significantly enhances the deployment of our system across various monitoring structures, which frequently demand different hardware and software components.

III. ARCHITECTURE

Fig. 1 depicts the MAC4PRO IoT-SHM architecture proposed in this article. It consists of four layers as follows.

1) *Sensing* layer, consisting of the sensing units necessary to acquire the physical phenomena to be monitored.
2) *Interoperability* layer, offering a uniform application programming interface (API) for two-way interaction with the *sensing* layer devices and related feature-extraction tasks.
3) *Data management* layer, working as data lake and enabling SHM data acquisition, aggregation, storage, and processing via SHM anomaly detection algorithms.
4) *Service* layer, including user applications able to query the data lake for custom data visualization and processing, being them built-in services or external, and third-party applications.

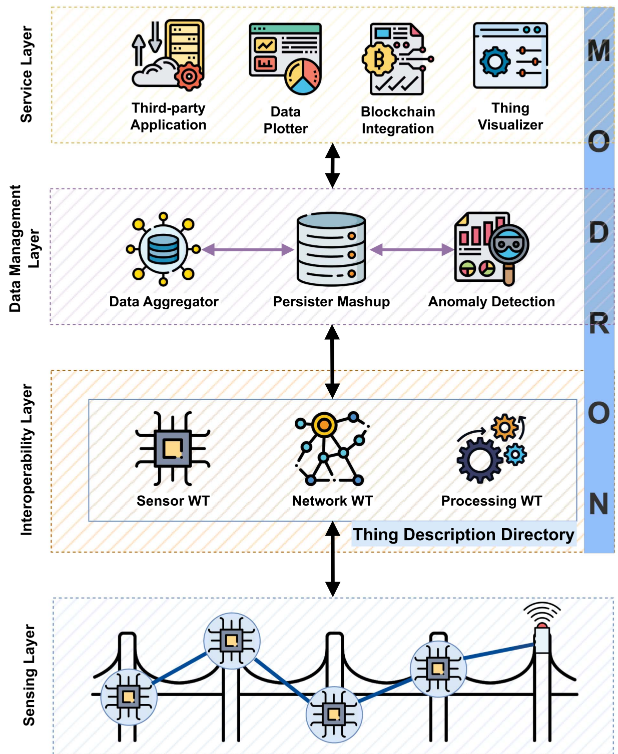

Fig. 1. MAC4PRO abstract architecture.

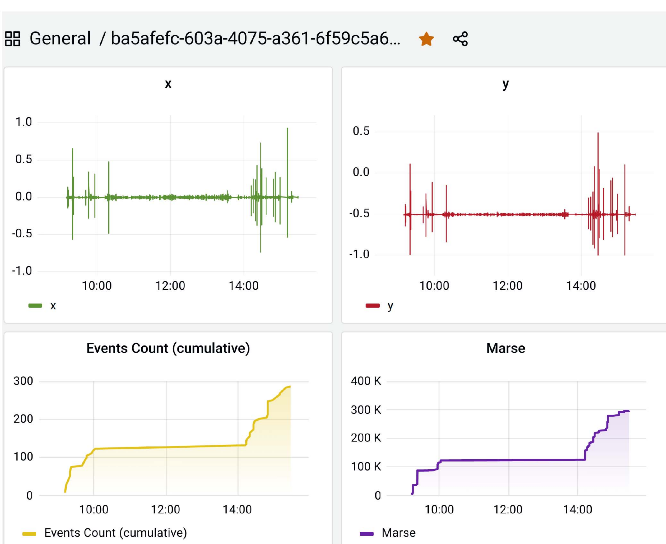

Fig. 2. MODRON *data plotter* depicting vibration and AE sensor data from the experimental campaigns.

The functionalities of the last three layers are provided by a software platform, called MODRON [24] and further detailed in the following sections. The architectural design presented so far is decoupled from the deployment plan. Based on the requirements and the available resources, the software components can be variously configured and deployed through the edge-cloud continuum. For instance, they can be assigned entirely to edge nodes near the monitored structure or distributed among the cloud and the edge nodes. We emphasize that the flexibility in distributing software components through the continuum pertains to deployment time rather than runtime. Nevertheless, our architecture inherently supports runtime migration since the migration of services between edge and cloud was already explored in literature [25], [26], [27]. We stress that our innovations rely not on specific algorithms but on the architectural system combining hardware and software components. We will provide evidence of such decoupling in Section VI, in which we deploy the architecture in different configurations.

A. Sensing Layer

The *sensing* layer corresponds to the sensor networks (SNs) in charge of data gathering. As a general observation, for a large-scale structure, we expect that the deployment of a single, multihop SN will introduce some pitfalls in terms of both end-to-end performance and reliability, as largely discussed in the literature [28]. For this reason, we assume that the same structure can be instrumented with multiple SNs consisting of extreme edge nodes (EENs), geographically isolated or with some spatial redundancy. These EENs may be heterogeneous regarding sensing and computational capabilities, communication protocols, and generated data formats. This heterogeneity is important for full-scale structural inspection to overcome the limitations of individual sensing technologies and their operative ranges. To capitalize on that, we have designed a distributed SN consisting of small-footprint, low-power, and light-weight EENs, i.e., peripheral devices integrating, in a thin form factor, all the circuitry necessary for heterogeneous data sampling, conditioning, and pre-/postprocessing [29]. Indeed, each EEN can acquire, at the same time, vibration data (i.e., triaxis accelerations and triaxis angular velocities via an inertial measurement unit having an output data rate as high as 6.66 kHz and up to three AE signals, which can be acquired with a maximum sampling frequency of 2 MHz.

Despite the advantages in terms of electrical characteristics [30], the designed EEN is unique in that it offers computing functionalities implementing sensor-near feature extraction for structural diagnostics. The embedded algorithms comprise, among others, an exhaustive list of parameters (e.g., amplitude, energy, count) for acoustic and vibration data processing.

In general terms, the number of EEN (and consequent amount of SNs) depends on the nature of the information to be processed, as well as on the specific diagnostic parameter to be estimated. In case vibration features are considered, it is indicated that, independently from the size and scale of the monitored target, the minimum number of sensors should be equal to the number of natural frequencies to be tracked [31]: this condition is necessary in case spatial-dependent information, such as mode shapes, have to be reconstructed. Preferable positions for vibration sensor installation are those located in correspondence of antinodal points, i.e., locations in which the energy of the identified modes is maximal. Beside, AE sensors are sensitive to stress-related phenomena, which typically manifest nearby joints or welding. Even if no practical boundary exist on their density, at least three sensors have to be installed in case AE source localization tasks have to be fulfilled [32].

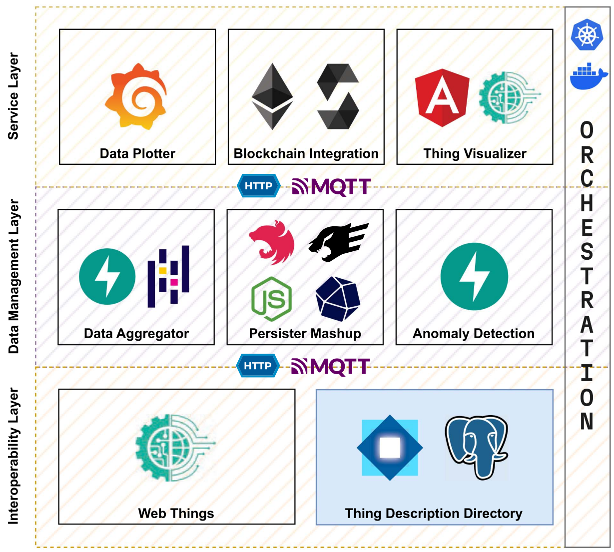

Fig. 3. MAC4PRO implementation.

B. Interoperability Layer

The *Interoperability* layer allows the MODRON platform to abstract as much as possible from the characteristics of the sensing units. The interoperability support takes advantage of the web of things (WoT) standard proposed by the W3C [33]. As per definition, *"The goal of the WoT is to preserve and complement existing IoT standards and solutions"* by providing strategies to describe what already exists rather than prescribing new mechanisms. A core component of the WoT standard is the web thing (WT), a physical or virtual entity whose metadata and interfaces can be uniformly and well-defined by a WoT thing description (TD). The latter represents a collection of standardized, machine-understandable metadata that allows consumers to discover and interpret the capabilities of a WT [33]. Among other fields, the TD includes: i) the *affordances*, providing an abstract model of the WT interface in terms of properties (i.e., the state variables of the WT), actions (i.e., commands that can be invoked on the WT) and events (i.e., notifications sent by the WT); ii) the *protocol bindings*, defining the mapping between the abstract affordances and the network strategies (e.g., the protocols) used to interact with the WT; iii) the *security configuration*, defining the mechanisms to control the accesses to the affordances. In MAC4PRO, the *interoperability* layer is composed of three classes of WTs, as shown in Fig. 3.

1) *Sensor-related WTs:* We associate a WT to each sensing unit of the SN, exposing the data produced from that EEN as readable properties, the configuration settings as writable properties, and supported commands as actions. For instance, for the case of triaxis accelerometers illustrated in Section VI, the properties include the raw signal values in each direction and the sampling frequency, while the actions include the possibility to turn ON/OFF the data acquisition on a specific axis. Thanks to the WT abstraction, the MODRON platform can establish a bidirectional, logical communication channel with each device of the *sensing* layer through a uniform and well-defined software interface.
2) *Network-related WTs:* We assign a WT to each SN, modeled as a whole. In such case, the WT includes links to the *sensor* WTs composing that SN. In addition, it may expose aggregated properties (e.g., the average network performance) and global commands (e.g., turning ON/OFF the SN by issuing the same command on each sensor WT).
3) *Processing-related WTs:* We associate WTs with software tasks in charge of processing the sensor data, extracting second-layer information from the monitored structure, and enabling error-handling capabilities. In such case, the WT is not connected to any physical device but acquires data from multiple sensor WTs, acting as virtual sensors. However, since it exposes a TD with its own properties, actions, and events, it can be displayed and controlled by the MODRON platform as a real device. For instance, the implementation of vibration data analysis tasks, such as natural frequency identification, can be provided with a dedicated WT, acquiring data from sensor WT (i.e., the accelerometers) and providing peak spectral values in output as TD properties.

The users and applications must be able to discover the WTs of an SHM scenario in order to interact with them. To this aim, the *interoperability layer* includes a thing discovery directory (TDD), which is a register of available WTs that provide search capabilities upon the metadata description of the devices (i.e., their TD) to the upper layer.

C. Data Management Layer

The *data management* layer collects the SHM data, aggregates it, and supports the analytics. To this aim, it leverages the facilities of the underlying *interoperability* layer, specifically the WT abstractions, to gather data from heterogeneous sensing units. The data acquisition is performed via the *persister mashup* through a three-stage procedure called *task*. First, a task retrieves the TDs of the WTs of interest from the TDD. Then, it establishes a direct connection to each WT. Finally, it issues a sampling sequence of actions and saves the returned data to the designated databases. The user can fully configure each task through a REST API; for instance, it is possible to specify the start time and frequency of execution, as well as its type (one-shot or periodic). In addition, the user must indicate the data source(s) to be queried (i.e., the WTs), the interaction affordance necessary for sampling the sensory value (e.g., a property), and a proper mapping function if the received data must fit a predefined database scheme.

Finally, SHM data are stored in the target database(s): we highlight that MODRON supports several adapters that enable different types of storage (e.g., time-series databases) to be managed. Generally speaking, the storage system can likely become the bottleneck of any IoT-SHM architecture due to the potential massive flow of data to be handled in real-time and the latency of persistence operations. To cope with such issue, our *persister* module is designed to scale vertically and horizontally.

During execution, the process of persisting SHM data into the target database(s) is queued, and the queue is processed using the maximum number of available threads. This approach ensures that the application can fully utilize the available resources on a given hardware platform, enabling efficient and low-latency data storage. It can also be distributed by splitting the workload between multiple instances, enabling task parallelization, increasing the system throughput, and reducing processing time. With this approach, the application can easily handle traffic and data volume spikes without compromising performance or reliability.

The collected data are immediately available to the *service* layer above and, simultaneously, are also used by the other two components on the *data management* layer shown in Fig. 3. The *data aggregator* extracts features from the time-series data retrieved from the *persister*. It includes standard statistical aggregation methods (e.g., mean, maximum, and minimum of a time-series), allowing for easy integration of context- and sensor-dependent aggregation methods through an extendable interface. The extracted features are stored to be quickly accessible for later processing or visualization steps. The data management pipeline includes collaboration between the data aggregator and the anomaly detection module: the first extracts features that the second utilizes to perform its computation. Both components use raw and feature data to assess the condition of the monitored structure, detect anomalies, and provide insights for maintenance operations. Different technical appliances may require specific analysis strategies and models to identify defects accurately; for this reason, the component is designed to be highly modular, allowing for the dynamic loading of new algorithms. This need for modularity is justified because the resource utilization of the involved diagnostic algorithms may vary depending on multiple causes: the nature of processed waveforms (vibrations versus AEs) and the global/local value of the computed damage-sensitive features. Indeed, the characterization process is conventionally performed in the frequency domain for appliances in a dynamic regime, as is the case of industrial rotors, wind turbines, and civil buildings that vibrate due to external stimuli. In particular, anomalies are identified by observing how the position and amplitude of the peak spectral values contained within the cross-power spectral density (CPSD) of the gathered signals, also termed natural modes of vibration, evolve over time [34]. Global damage indexes are employed in this case, which can be computed only in a postprocessing phase upon aggregating EEN-related information. Therefore, vibration data's data aggregation and anomaly detection components work asynchronously, with periodic queries extracting data from the databases, processing it on the cloud, and storing the computed outputs.

Conversely, AE techniques are recommended when the primary source of defects is intrinsically related to energy release. In this case, the trend of the compute AE features can provide insightful information about the intensity of long-aging stress. The key AE parameters for real-time assessment are signal peak amplitude, signal energy and AE count, which are strictly EEN-dependent. Time analysis of these parameters allows early detection and localization of incipient faults, such as growing cracks in reinforced concrete (RC), corrosion processes in metal structures, or leaks from pipelines [35]. Hence, when handling AE data, performing feature extraction at the EEN level is suitable.

D. Service Layer

The *service* layer uses the data access APIs of the *data management* layer to provide user functionalities. Third parties can develop these or be custom implemented directly by the MODRON platform. The built-in services include the *data plotter*, an advanced dashboard for visualizing the stored data (raw and processed), the *thing visualizer*, a graphic interface to manage devices, and the *blockchain integration*, which guarantees data transparency and immutability.

The *data plotter* offers a wide range of functionalities allowing users to create custom charts, filter data based on the WTs time interval, and apply various transformations to the displayed data. A representative screenshot of the *data plotter* is reported in Fig. 2, considering real-world data of the two experimental campaigns presented in Section VI. In addition, it supports exporting charts and data in multiple formats, making it easy to embed it into other applications and share it with different collaborators.

The *thing visualizer* allows users to manage and interact with the available WTs. The service can render the TD into a web interface, i.e., automatically generating a graphical panel that reflects the properties, actions, and events defined in the TD. Through the *thing visualizer*, end users can observe the value of a property, issue an action for modifying WT configurations, and subscribe to events to receive real-time updates from the WT.

Finally, the *blockchain integration* service provides an additional layer of security and trust to the MODRON platform in SHM scenarios of critical facilities (e.g., buildings, bridges, industrial plants) [36], [37]. The data logging system exploits the blockchain's feature of maintaining a permanent and unalterable history of transactions to guarantee the immutability and transparency of the SHM data. Specifically, events related to detecting an anomaly in the structure are stored on the chain, making them easily verifiable by external auditors.

IV. Software Implementation

This section describes the MAC4PRO software implementation. It details the technologies and methodologies utilized to develop the architectural components. Some of these components have been developed from scratch to meet the specific requirements of our system. In contrast, others are industry-adopted solutions that have been integrated into the MAC4PRO software. The interoperability layer is implemented using the Eclipse Thingweb node-wot runtime,[2] which serves as the execution environment for the JavaScript applications responsible for instantiating the WTs described in Section III-B. We develop the TDD to meet the system requirements for scalability and advanced querying capabilities. Our current solution is developed

[2] [Online]. Available: https://www.thingweb.io

using Node.js and utilizes a Postgres database for storage. One of its standout features is the full support for JSONPath queries, which provides the upper layers with a mechanism to search for WTs within and across monitored structures.

We custom-built all the components in the data management layer since we identified very specific requirements regarding APIs and performance constraints during our architectural analysis. The persister mashup, developed in Node.js, leverages the Fastify[3] web framework for its REST API and uses InfluxDB[4] as our primary time-series database. In addition, it incorporates the node-cron package to schedule and run user tasks. These technologies have been chosen for adequate reliability and scalability even under high workloads without risking service unavailability and data loss. The persister mashup is highly modular and configurable, supporting multiple databases through an abstraction layer implemented using the adapter pattern, which enables the component deployment into constrained edge environments where a full-fledged database may not be supported. We implemented the data aggregator using the FastAPI[5] framework, capitalizing on its rapid performance and asynchronous capabilities. To manage the concurrency of requests, we utilized Bull[6] as our queuing system. The data aggregator exposes an API that allows other system components to request specific data manipulations, like calculating the mean of a designated time-series or undertaking other statistical analyses. As requests arrive, they are placed into the Bull queue and processed asynchronously, optimizing system performance. Once a manipulation task concludes, the data aggregator updates the database with the results and immediately notifies the calling service via dedicated webhooks. Our anomaly detection component is also built on the FastAPI framework and designed for modularity and adaptability. Users can easily integrate specific detection algorithms, addressing the diverse needs of SHM environments. The system supports concurrent algorithms and configurable alarms to notify users of detected anomalies. Some examples of these algorithms can be found in Section VI.

The data plotter component is a Grafana[7] dashboard that has been modified to integrate our authentication flow while retaining its core functionalities. The thing visualizer is a custom application developed in Angular that can manage and render WTs of different users and organizations. In order to communicate with the WTs, we used Eclipse Thingweb node-wot as a browser-side JavaScript Library with both the protocol bindings for HTTP/HTTPS and MQTT. Finally, the blockchain integration component is implemented through a dedicated Python application; it directly routes specific, critical events into an Ethereum virtual machine[8]—compatible blockchain configurable by the user. The data available on-chain cannot only be utilized for auditing but also be cross-verified. Dedicated smart contracts can be leveraged to retrieve external IoT data via oracle systems [37], enabling a comprehensive comparison with the data persisted on the blockchain.

Orchestrating these diverse and numerous components is crucial for achieving a robust and scalable system. Given the microservices nature of our implementation, we adopted an orchestration platform that could handle the management, deployment, and scaling of these services across the edge-cloud continuum. To this end, we employed Docker[9] and Kubernetes.[10] By using these tools, we ensure seamless service discovery, load balancing, and self-healing across both edge devices and centralized cloud infrastructure. All our components are containerized using only Alpine-based images to reduce disk space, build time, and spare bandwidth when there is need to move containers in our stack or update them with new releases. Kubernetes handles the orchestration, ensuring that the desired state of our services is consistently maintained. Furthermore, it optimizes resource allocation and automatically restarts or replaces containers during failures, ensuring high availability. This is a critical aspect for SHM use cases since if the system (especially the WTs and the persister mashup) experiences downtime during critical events, such as earthquakes or other unpredictable structural stresses, it would result in the loss of invaluable data.

V. Performance Analysis

This section aims to comprehensively characterize the proposed architecture and its implementation by examining its key variables and identifying potential bottlenecks. We conducted a thorough evaluation encompassing the most significant metrics at each infrastructure level—i.e., EEN, edge, and cloud. We examined the effects of task execution under different edge-cloud continuum configurations, focusing on the tradeoffs associated with the feature extraction task detailed in Section III-C. This task holds particular significance as it reduces the data dimensionality. To evaluate its impact, we analyzed the deployment of this task on both the EEN and in the cloud.

Concerning the EEN, we analyzed the impact on energy consumption when performing feature extraction onboard or not. To this end, the number of collected samples per sensor (single acquisition) was varied (1024, 4096, and 16 384), and the energy spent to compute these features and transmit them (or raw data) has been analyzed, encompassing different wireless transmission technologies [38]. Fig. 4 summarizes the results and highlights that performing feature extraction on the EEN is more efficient: this is coherent with the fact that the energy expenditure in data processing is minimal if compared to the transmission costs of wireless modulus. This pattern characterizes all the evaluated communication technologies and is enhanced when the data payload increases. An IoT analyzer [39] has been used for energy profiling, considering the hardware and software components of our specific architecture. In particular, in computational terms, we have assumed the electrical properties of the STM32F3 family of microprocessors (maximum clock frequency of 72 MHz, 40 mA, and 10 μA in run and sleep mode, respectively).

[3][Online]. Available: https://fastify.dev
[4][Online]. Available: https://influxdata.com
[5][Online]. Available: https://fastapi.tiangolo.com
[6][Online]. Available: https://optimalbits.github.io/bull
[7][Online]. Available: https://grafana.com
[8][Online]. Available: https://ethereum.org/en/developers/docs/evm
[9][Online]. Available: https://docker.com
[10][Online]. Available: https://kubernetes.io

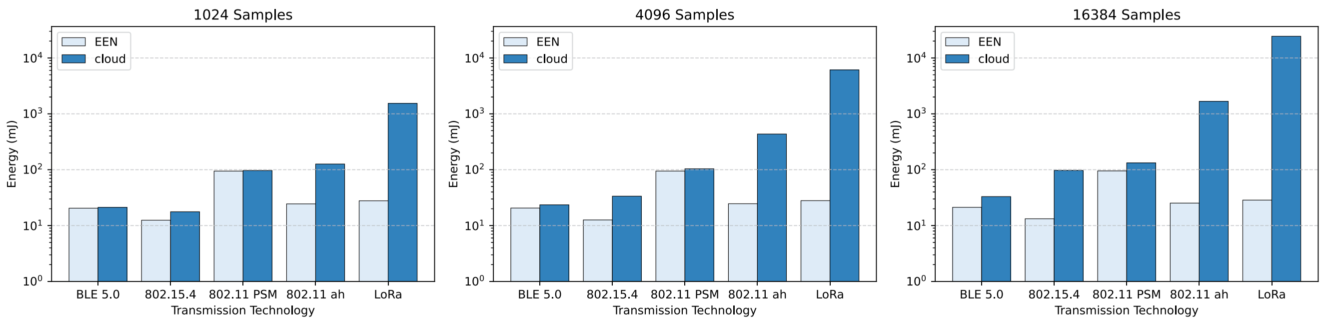

Fig. 4. Energy consumption analysis of feature extraction on EEN.

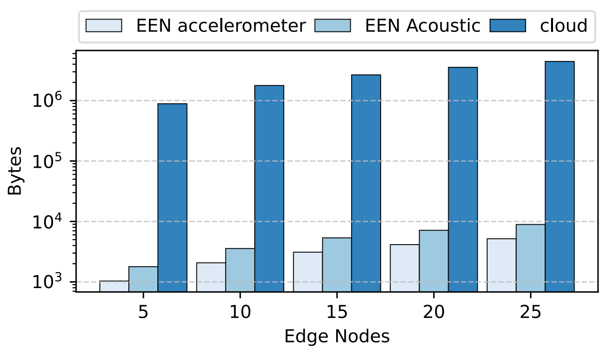

Fig. 5. Data payload size comparison when performing feature extraction in EEN versus the cloud.

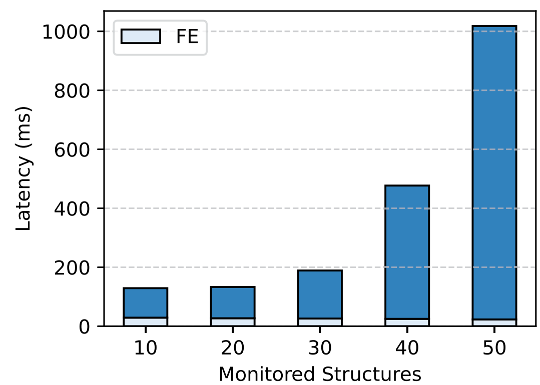

Fig. 6. Cloud application scalability.

The placement of feature extraction within the edge-cloud continuum directly impacts the amount of data transferred between the edge and the cloud. Edge-cloud data transfer is a key variable in SHM scenarios, where poor network connectivity is a common challenge. Structures under monitoring often lack dedicated networking infrastructure and are exposed to various environmental hazards, including adverse weather conditions. Therefore, we evaluated data payload sizes when performing feature extraction at the EEN or in the cloud. We assume that each data acquisition generates 2500 samples, which are transformed into human- and machine-readable formats through the WT abstractions. Each edge device is equipped with a triaxial acceleration sensor and with three-channel AE SNs. In the AE case, we extract eleven features (summing all the energy and time-related features), while in the accelerometer case, we extract six features—the number of frequencies of interest. Consequently, the feature extraction payload size varies between these two types of SNs. The results of this evaluation, depicted in Fig. 5, illustrate the byte size of the payload for different configurations. Notably, performing feature extraction in the cloud results in a higher workload on the network.

Finally, we aim to understand how the increasing workload impacts the cloud, especially when receiving raw data and performing feature extraction. As the same cloud application can monitor several structures, we scaled up the number of monitored structures to analyze how the cloud application performs under these conditions. We assume that each monitored structure is equipped with ten edge nodes. As such, we are modeling dense and realistically complicated geometries, which require a fine sensor installation plan to capture different damaging phenomena. We realize that in real use cases, this number is different depending on the physical properties of each structure. For our experiments, we simulated data transmission frequencies based on real-world scenarios. Each edge node transmits an accelerometer payload every hour and an AE payload every minute. This higher AE sampling rate is typically necessary when monitoring degraded structures, we opt for evaluating the system in this configuration as the deployed cloud applications needs to able to support the system in critical scenarios. To scale the data generation process, we developed a workload generator that emulates the edge transmission of accelerometer and AE data. The cloud application in this scenario comprised two components: the data aggregator and the persister mashup, consisting of a NodeJS and InfluxDB Docker container, respectively, as detailed in Section IV. Each container had access to one logical CPU (Intel(R) Xeon(R) Gold 6238R CPU @ 2.20 GHz) and 4 GB of RAM. The workload generator is connected via LAN to accurately assess the performance of the cloud application under controlled and consistent network conditions. Each experiment represents a time window of 10 min of continuing transmitting data. Fig. 6 shows the processing latency from data generation to its inclusion in the database and it demonstrates that the feature extraction component is scalable, as increasing the workload does not significantly impact its execution time. In addition, the system bottleneck is associate to managing multiple connections and efficiently transforming and storing data in the database. Notably, the system showcases high scalability, as even with constrained resources, it can handle multiple monitored structures. To further enhance the system scalability, scaling the cloud computational nodes horizontally or vertically is a viable option.

Summarizing the conducted evaluation, performing feature extraction in the EEN decreases its energy consumption and decreases the the amount of bytes transferred between the edge and the the cloud. On the other hand, the additional computation imposed by performing the feature extraction in the cloud is minimal, and the cloud offers greater stability compared to the edge. Moreover, conducting feature extraction in the edge increases the complexity of the EEN, which may require more expensive equipment. In conclusion, there is no universally optimal component placement within the continuum. Instead, a careful analysis of tradeoffs is essential to address each scenario's unique characteristics. Fortunately, our architecture versatility enables the exploration of various deployment configurations, empowering users to adapt it to their specific needs.

VI. Real Field Deployment

We deployed the MAC4PRO platform in two distinct testbeds to validate its feasibility and showcase its versatility. The goal is not to quantify MAC4PRO performance but to demonstrate its practicability and applicability in real-world deployments. The first scenario involved monitoring a concrete frame structure during mechanically simulated seismic events utilizing a large-scale shaking table. We built and instrumented a hydraulic circuit in the second scenario to identify water leakages. Both scenarios were monitored in real-time. We emphasize that the scenarios evaluated differ not only in their respective domains but also in the nature of the raw and processed data. In the first one, data bursts are generated and transmitted in each seismic excitation, whereas in the second, a data stream is transmitted constantly through the entire system uptime.

A. RC Frame Under Seismic Excitation

1) Structure: The specimen used within this experimental campaign is a two-story RC frame, 3×3 m (x and y in-plane directions), 4 m tall (z-direction), having columns and beams with cross-sections 20×20 cm and 20×30 cm, respectively. The frame was purposely designed according to nonseismic codes, namely without joint resources for ductility as requested by recent seismic provisions, to represent a bulk of existing non-recent buildings. The structure is built upon two one-way ribbed floor slabs lightened by hollow clay blocks, with the possibility of applying additional masses by steel plates at different points: a picture is shown at the bottom of Fig. 7.

Testing was conducted at the seismic hall of the ENEA Casaccia Research Center, which is equipped with a 4×4 m shaking table capable of applying seismic inputs on large mockups of structures up to 30 t of weight. To this end, the seismic acceleration recorded in Norcia at Savelli Station on October 30, 2016 was selected as the input shaking force since it provoked a disastrous earthquake in Italy. Signals were applied with increasing levels of maximum peak ground acceleration (PGA), from 0.1 to 0.8 g, to damage the frame progressively. These inputs were spaced by white noise (WN) excitations of

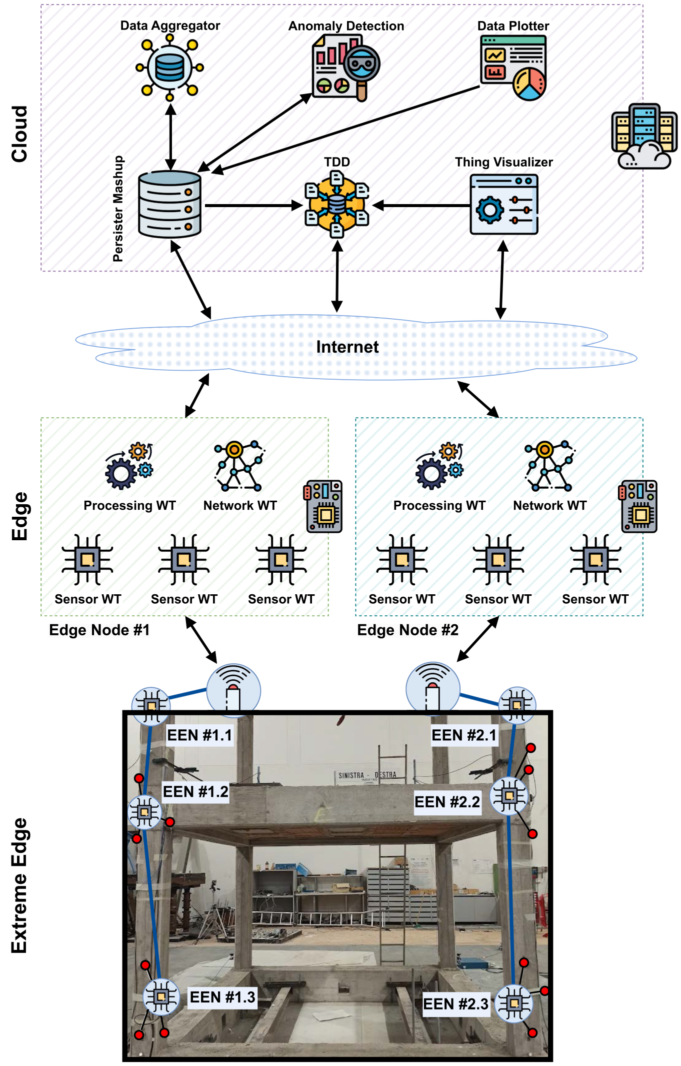

Fig. 7. MAC4PRO deployment plan on the RC frame (red dots indicate the position of the AE transducers).

constant amplitude and duration for later use in characterizing the structure's frequency parameters.

2) Sensor Network: Two SNs comprising three EEN devices (one per floor) were installed on two opposite columns of the frame: the positions were selected after a preliminary numerical simulation of the elastomechanical properties of the structure. Splitting the EENs in two different SNs has been preferred over the deployment of one single network in order to minimize the length and the number of cables to be deployed, which would have otherwise been affected by high electromagnetic noise and interference. Two EENs (one per SN) were connected to three G150 AE transducers (working frequency up to 400 kHz): the bonding to the structure was realized via metal platforms for landing magnetic sensor holders. All EENs were programmed to acquire, at the same time, triaxial accelerations; for those connected to AE transducers, the sensing and on-board processing of AE features was also enabled.

3) MAC4PRO Deployment Plan: Fig. 7 illustrates the deployment plan of MAC4PRO to enable the monitoring of the concrete structure. In this scenario, each SN was connected to an edge device (i.e., Raspberry Pi), which abstracts the particularities of the SN and EENs, instantiating them as WTs. In addition, a processing WT was utilized to perform edge computation tasks. Specifically, it reads raw sensor data—as a binary data stream—and converts it to a well-structured human- and machine-understandable format, including additional metadata (i.e., collected data timestamp, unique universal sensor id). It performs the first step of data cleaning and error handling. The processing WT handles exceptions and ill-formed data (generated by sensor or communication errors) not to jeopardize the whole processing pipeline. It is strategically placed at the edge for a twofold reason: i) there is intense data communication between the processing WT and the other WTs that could occupy a large portion of the available bandwidth between cloud and edge; ii) some error handling strategies trigger device commands, which need to be executed with low latency not to propagate errors and, thus, minimize the data loss. Finally, the edge is the first infrastructure component of our platform connected to the Internet. For this reason, it implements security mechanisms. The data transmission from the cloud to the edge is encrypted through the HTTPS protocol. In addition, we leverage the usage of unique identifiers assigned to authenticated applications and users to control and monitor the access for the exposed WTs.

Regarding the continuum configuration, the computation-intensive tasks were deployed in the cloud server—namely, the *data management* and *service* layers components. The TDD is the only component of the *interoperability* layer (see Section III-A) deployed in the cloud since it indexes the WTs from all edges nodes, and it is the entry point enabling the discovery of the MAC4PRO WTs. In this scenario, the vibration data analysis—performed by the data aggregator and the anomaly detection—was not performed in real-time since it was not the testbed goal. The vibration data were first acquired through the SNs, then stored, and finally processed to investigate the effects of the emulated earthquake. On the other hand, feature extraction of AE data was computed directly at acquisition time by exploiting the unique on-sensor computing capabilities of the developed EEN devices. The EEN processing is crucial to diminish the burden of the subsequent applications in the computing pipeline, saving considerable bandwidth and storage. Considering only AE data, the performed EEN computation imposes a **reduction of 99.8%** in the transmitted payload [40].

4) Vibration-Based Monitoring: This analysis aimed to determine the variation in the spectral signature of the structure due to the deterioration induced by increasing seismic activity. To this end, the waterfall in the first singular value of the CPSD matrix, which is a cumulative measure of the spectral profile of the structure taking into consideration all the deployed sensors, has been estimated on the acceleration signals recorded along the y-direction and reported in the left panel of Fig. 8. From these plots, it is possible to observe that all the natural frequencies undergo a consistent reduction while increasing the extent of the PGA over time, moving from less than 1% variation for PGA $= 0.1$ g to more than 3% at the end of the testing campaign. Furthermore, a structural mechanics-based damage index

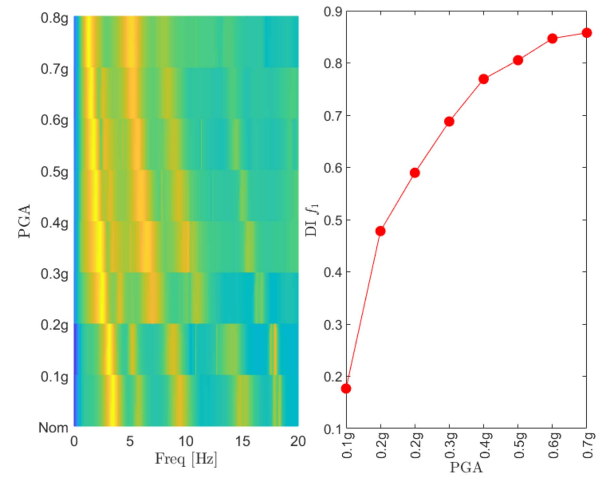

Fig. 8. Waterfall in the first singular value of the CPSD along the y-direction (left). DI_1 for the first natural mode of vibration showing a significant increment as the severity of the seismic activity increases (higher PGA) (right).

$$DI_i = 1 - \left(\frac{f_i}{\hat{f}_i}\right)^2 \tag{1}$$

has been implemented to identify anomalies related to variations in high-mass participating modes, alike the first mode of the RC frame [41]. DI_i expresses, each ith mode of vibration, the square relative ratio between the mode-related frequency value $\hat{f}_i$ computed in reference conditions and the one estimated during on-condition maintenance. $DI_i = 0$ indicates absence of relevant structural variation (namely, $f_i = \hat{f}_i$), while higher values manifest significant changes in the spectral pattern of vibration. In particular, DI_1 (associated to the first frequency of vibration identified along the y-axis) has been computed and reported in the right panel of Fig. 8. As can be observed from the reported trend, DI_1 exhibits a significant increment while moving from low seismic activity ($DI_1 = 0.17$ for PGA $= 0.10$ g) to very consistent earthquake-like excitation ($DI_1 = 0.85$ for PGA $= 0.7$ g).

5) AE-Based Monitoring: In AE testing of concrete structures under incremental cyclic loading, AE energy is widely used for damage characterization [42]. Consistently, in our case, we investigated the pattern in the cumulative AE energy acquired during WN tests to spot the occurrence of structural changes in response to seismic activity. As it can be observed from Fig. 9, in which the trends for the different AE transducers connected to the two EENs are reported, all curves follow a similar pattern: initially (PGA from 0.1 to 0.4 g), the AE activity increases, then a gradual reduction happens for 0.5 g until it minimizes at 0.8 g. In accordance with the experimental study of RC beams under cyclic load [43], the decrease in the total AE energy released can be explained by brittleness in the test frame due to the high percentage of reinforcing steel inside the considered element.

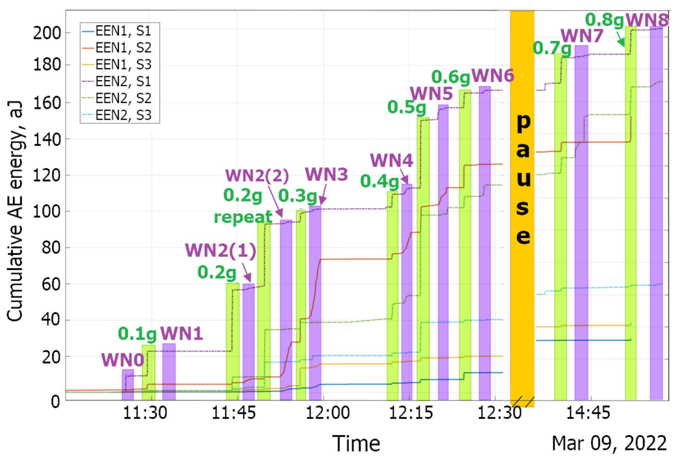

Fig. 9. AE energy trend during testing from 0.1 to 0.8 g of PGA (EEN = extreme edge node, S1–S3 = AE sensors, WN = white noise, 0.x g = seismic motion).

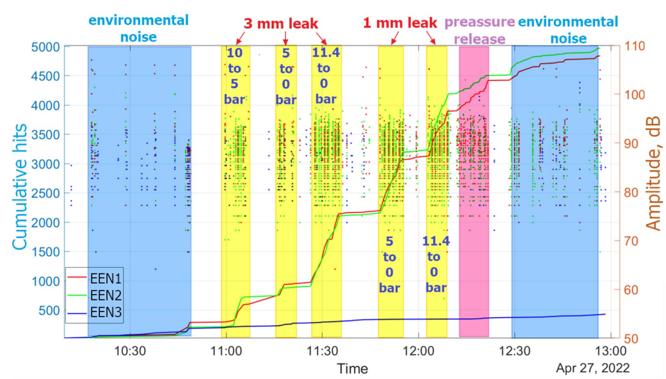

Fig. 10. Three-hour monitoring of the hydraulic circuit: trend in cumulative AE hits and peak amplitude in various pressure/leakage conditions.

B. Hydraulic Circuit Under AE Leakage

1) Structure: The second experimental campaign aims to validate the ability of MAC4PRO architecture to detect fluid leakage that might arise during the pressurization of industrial facilities, such as vessels. To this purpose, the hydraulic circuit in Fig. 11 has been built and exploited as a target structure. The circuit comprises a pipe loop that can independently pressurize two 1000-L vessels; moreover, it can be controlled from a dedicated control room or remotely through a dashboard.

2) Sensor Network: As depicted in Fig. 11, one SN consisting of three EENs, each connected to one AE transducer (i.e., S1, S2, and S3) installed by magnetic connections along the pipeline of the pressure circuit. One sensor (S2) was attached in proximity to the opening valve (inset in the center of the figure) used to simulate leakage, while the remaining two are far apart along the pipe components.

3) MAC4PRO Deployment Plan: Fig. 11 depicts the placement of the MAC4PRO architectural components in the hydraulic circuit. Compared to the previous testbed presented in Section VI-A, we introduced the following changes in the deployment plan. First, we deployed a single edge node since

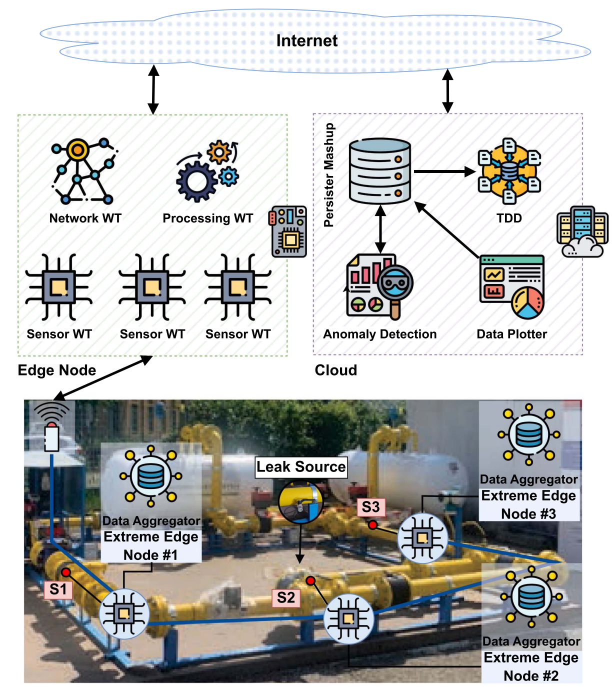

Fig. 11. MAC4PRO hydraulic circuit deployment plan.

only one SN exists in the scenario. Second, the *data aggregator* component was deployed only in each EEN, while in the previous testbed the feature extraction was performed in the cloud and the EEN. Moreover, the *thing visualizer* component was not deployed in this testbed since updating sensor configurations and metadata at runtime was unnecessary. We highlight that such changes to the deployment plan were possible thanks to the modularity and versatility of the MAC4PRO architecture can be easily customized to support many SHM scenarios.

4) Leakage Detection: AE testing is widely used to inspect the condition of pipelines [44] and pressure vessels [45]. Therefore, we focus in this paragraph on AE-based diagnostics. For the sake of clarity, the objective pursued in this experimental campaign was to deliver a new monitoring framework that, thanks to the seamless integration of SW and HW tools, could be compatible with long-term AE installation. By doing so, we deviate from the conventional approaches adopted in AE testing for pipe and pressure vessel monitoring, which involve halting normal plant operations to perform nondestructive inspections in controlled conditions. Accordingly, a continuous monitoring system was developed and validated by recording several experimental tests during which leakage was simulated by opening/closing a valve fitted with a nozzle of different diameters (from 3 to 1 mm) and under different pressure levels. Fig. 10 presents the results regarding cumulative AE hits and peak amplitude over time: the areas marked in blue refer to a temporary increase in environmental noise, as the hydraulic circuit is located close to a carriageway. The areas marked in yellow present several cycles of pressure release by opening the valve and the leakage activation. As shown, both the peak amplitude in time distribution and the cumulative count responded with a notable trend change, which confirms the system's high sensitivity to low-frequency and small-size AE events, such as water leakages. The red region demonstrated the system's activity when the circuit was depressurized due to the imperfect closure of the valve. In addition, comparing the outcome from the three EENs, it is reasonable to observe a more intense acoustic activity in correspondence with EEN2 (the one connected to S2) since it was closer to the leakage source and, thus, more sensitive to the induced defect. EEN3 featured the lowest leakage sensitivity due to the longest distance from the AE source.

VII. Conclusion

In this work, the SHM-IoT MAC4PRO architecture has been thoroughly described and showcased for the condition assessment of two structural targets representative of industrial and civil appliances. The architecture is designed to abstract from the specific use case and the underlying infrastructure, which is achieved by integrating advanced SN solutions with the state-of-the-art SHM data acquisition, modeling, and processing techniques. The proposed architecture comprises four layers: Sensing, interoperability, data management, and service. These layers can be deployed adaptively across the edge-cloud continuum based on the requirements and characteristics of the SHM applications. We have demonstrated the versatility of our framework in two experimental campaigns, in which we discussed the deployment plans in the continuum and presented diagnostic results. In future works, we plan to integrate advanced anomaly detection strategies empowered by artificial intelligence techniques at either the edge or cloud level. We also intend to validate our architecture on additional real-world SHM scenarios to further demonstrate its effectiveness and robustness.

References

[1] B. Natesha and R. M. R. Guddeti, "Fog-based intelligent machine malfunction monitoring system for industry 4.0," *IEEE Trans. Ind. Informat.*, vol. 17, no. 12, pp. 7923–7932, Dec. 2021.

[2] M. Z. Alam Bhuiyan, J. Wu, G. Wang, and J. Cao, "Sensing and decision making in cyber-physical systems: The case of structural event monitoring," *IEEE Trans. Ind. Informat.*, vol. 12, no. 6, pp. 2103–2114, Dec. 2016.

[3] M. Mishra, P. B. Lourenço, and G. V. Ramana, "Structural health monitoring of civil engineering structures by using the Internet of Things: A review," *J. Building Eng.*, vol. 48, 2022, Art. no. 1103954.

[4] R. Li, T. Mo, J. Yang, S. Jiang, T. Li, and Y. Liu, "Ontologies-based domain knowledge modeling and heterogeneous sensor data integration for bridge health monitoring systems," *IEEE Trans. Ind. Informat.*, vol. 17, no. 1, pp. 321–332, Jan. 2021.

[5] J. Fan, F. Han, and H. Liu, "Challenges of big data analysis," *Nat. Sci. Rev.*, vol. 1, no. 2, pp. 293–314, 2014.

[6] D. Mourtzis, *Design and Operation of Production Networks for Mass Personalization in the Era of Cloud Technology*. Amsterdam, The Netherlands: Elsevier, 2021.

[7] M. Cheraghzade and M. Roohi, "Deep learning for seismic structural monitoring by accounting for mechanics-based model uncertainty," *J. Building Eng.*, vol. 57, 2022, Art. no. 104837.

[8] C. Scuro, F. Lamonaca, S. Porzio, G. Milani, and R. Olivito, "Internet of things (IoT) for masonry structural health monitoring (SHM): Overview and examples of innovative systems," *Construction Building Mater.*, vol. 290, 2021, Art. no. 123092.

[9] V. Naraharisetty, V. S. Talari, S. Neridu, P. Kalapatapu, and V. D. K. Pasupuleti, "Cloud architecture for IoT based bridge monitoring applications," in *Proc. Int. Conf. Emerg. Techn. Comput. Intell.*, 2021, pp. 39–42.

[10] M. C. De Simone, A. Lorusso, and D. Santaniello, "Predictive maintenance and structural health monitoring via IoT system," in *Proc. IEEE Workshop Complexity Eng.*, 2022, pp. 1–4.

[11] S. Jeong and K. Law, "An IoT platform for civil infrastructure monitoring," in *Proc. IEEE 42nd Annu. Comput. Softw. Appl. Conf.*, 2018, pp. 746–754.

[12] H. V. Dang, M. Tatipamula, and H. X. Nguyen, "Cloud-based digital twinning for structural health monitoring using deep learning," *IEEE Trans. Ind. Informat.*, vol. 18, no. 6, pp. 3820–3830, Jun. 2022.

[13] Z. Pengfei, T. K. Boon, and W. Yixin, "Sensor selection in wireless sensor networks for structural health monitoring," in *Proc. IEEE Sensors*, 2019, pp. 1–4.

[14] F. Zonzini, M. Zauli, M. Mangia, N. Testoni, and L. De Marchi, "Model-assisted compressed sensing for vibration-based structural health monitoring," *IEEE Trans. Ind. Informat.*, vol. 17, no. 11, pp. 7338–7347, Nov. 2021.

[15] B. Qiu and J. Tan, "Real-time, automatic and wireless bridge monitoring system based on MEMS technology," *J. Civil Eng. Archit.*, vol. 10, pp. 1027–1031, 2016.

[16] H. Wenzel, *Health Monitoring of Bridges*. Hoboken, NJ, USA: Wiley, 2008.

[17] C. Bigoni, Z. Zhang, and J. S. Hesthaven, "Systematic sensor placement for structural anomaly detection in the absence of damaged states," *Comput. Methods Appl. Mechanics Eng.*, vol. 371, 2020, Art. no. 113315.

[18] P. D'Antuono, W. Weijtjens, and C. Devriendt, "On the minimum required sampling frequency for reliable fatigue lifetime estimation in structural health monitoring. how much is enough?," in *Proc. Eur. Workshop Struct. Health Monit.*, Springer, 2022, pp. 133–142.

[19] J. Wang, J. Yang, and Z. Zhang, "Design of cloud computing platform based accurate measurement for structure monitoring using fiber Bragg grating sensors," in *Proc. IEEE 2nd Int. Conf. Big Data, Artif. Intell. Internet Things Eng.*, 2021, pp. 807–811.

[20] C. Martín, D. Garrido, L. Llopis, B. Rubio, and M. Díaz, "Facilitating the monitoring and management of structural health in civil infrastructures with an edge/fog/cloud architecture," *Comput. Standards Interfaces*, vol. 81, 2022, Art. no. 103600.

[21] Z. He, W. Li, H. Salehi, H. Zhang, H. Zhou, and P. Jiao, "Integrated structural health monitoring in bridge engineering," *Automat. Construction*, vol. 136, 2022, Art. no. 104168.

[22] M. Fawad et al., "Automation of structural health monitoring (SHM) system of a bridge using BIMification approach and BIM-based finite element model development," *Sci. Rep.*, vol. 13, no. 1, Aug. 2023, Art. no. 13215.

[23] A. G. Mohapatra et al., "Fiber Bragg grating sensors driven structural health monitoring by using multimedia-enabled IoT and Big Data technology," *Multimedia Tools Appl.*, vol. 81, no. 24, pp. 34573–34593, Oct. 2022.

[24] C. Aguzzi et al., "MODRON: A scalable and interoperable web of things platform for structural health monitoring," in *Proc. IEEE 18th Annu. Consum. Commun. Netw. Conf.*, 2021, pp. 1–7.

[25] C. Aguzzi, L. Gigli, L. Sciullo, A. Trotta, and M. Di Felice, "From cloud to edge: Seamless software migration at the era of the web of things," *IEEE Access*, vol. 8, pp. 228118–228135, 2020.

[26] K. Fu, W. Zhang, Q. Chen, D. Zeng, and M. Guo, "Adaptive resource efficient microservice deployment in cloud-edge continuum," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 8, pp. 1825–1840, Aug. 2022.

[27] Y. Chen, Y. Sun, C. Wang, and T. Taleb, "Dynamic task allocation and service migration in edge-cloud IoT system based on deep reinforcement learning," *IEEE Internet Things J.*, vol. 9, no. 18, pp. 16742–16757, Sep. 2022.

[28] S. Sharma and A. Kaur, "Survey on wireless sensor network, its applications and issues," in *J. Phys.: Conf. Ser.*, vol. 1969, no. 1, 2021, Art. no. 012042.

[29] N. Testoni et al., "A sensor network with embedded data processing and data-to-cloud capabilities for vibration-based real-time SHM," *J. Sensors*, vol. 2018, 2018, Art. no. 2107679.

[30] F. Zonzini, M. M. Malatesta, D. Bogomolov, N. Testoni, A. Marzani, and L. De Marchi, "Vibration-based SHM with upscalable and low-cost sensor networks," *IEEE Trans. Instrum. Meas.*, vol. 69, no. 10, pp. 7990–7998, Oct. 2020.

[31] R. Brincker et al., "Modal identification from ambient responses using frequency domain decomposition," in *Proc. 18th Int. Modal Anal. Conf.*, San Antonio, TX, USA, 2000, pp. 625–630.

[32] K. M. Holford and D. Carter, "Acoustic emission source location," *Key Eng. Mater.*, vol. 167, pp. 162–171, 1999.

[33] L. Sciullo, L. Gigli, F. Montori, A. Trotta, and M. D. Felice, "A survey on the web of things," *IEEE Access*, vol. 10, pp. 47570–47596, 2022.

[34] W. Fan and P. Qiao, "Vibration-based damage identification methods: A review and comparative study," *Struct. Health Monit.*, vol. 10, no. 1, pp. 83–111, 2011.
[35] M. Wevers and K. Lambrighs, *Applications of Acoustic Emission for SHM: A Review*. Hoboken, NJ, USA: Wiley, 2009, pp. 9–15.
[36] L. Gigli, L. Sciullo, F. Montori, A. Marzani, and M. Di Felice, "Blockchain and web of things for structural health monitoring applications: A proof of concept," in *Proc. IEEE 19th Annu. Consum. Commun. Netw. Conf.*, 2022, pp. 699–702.
[37] L. Gigli, I. Zyrianoff, F. Montori, C. Aguzzi, L. Roffia, and M. Di Felice, "A decentralized oracle architecture for a blockchain-based IoT global market," *IEEE Commun. Mag.*, vol. 61, no. 8, pp. 86–92, Aug. 2023.
[38] F. Zonzini, V. Dertimanis, E. Chatzi, and L. D. Marchi, "System identification at the extreme edge for network load reduction in vibration-based monitoring," *IEEE Internet Things J.*, vol. 9, no. 20, pp. 20467–20478, Oct. 2022.
[39] A. Morin, M. Maman, R. Guizzetti, and A. Duda, "Comparison of the device lifetime in wireless networks for the Internet of Things," *IEEE Access*, vol. 5, pp. 7097–7114, 2017.
[40] D. Bogomolov, N. Testoni, F. Zonzini, M. Malatesta, L. de Marchi, and A. Marzani, "Acoustic emission structural monitoring through low-cost sensor nodes," in *Proc. 10th Int. Conf. Struct. Health Monit. Intell. Infrastructure*, Porto, Portugal, 2021, pp. 383–388.
[41] J. R. Casas and J. J. Moughty, "Bridge damage detection based on vibration data: Past and new developments," *Front. Built Environ.*, vol. 3, 2017, Art. no. 4.
[42] M. Desa, M. Ibrahim, S. Shahidan, N. Ghadzali, and Z. Misri, "Fundamental and assessment of concrete structure monitoring by using acoustic emission technique testing: A review," in *Proc. IOP Conf. Series: Earth Environ. Sci.*, IOP Publishing, 2018, Art. no. 012142.
[43] R. V. Sagar, "Acoustic emission characteristics of reinforced concrete beams with varying percentage of tension steel reinforcement under flexural loading," *Case Stud. Construction Mater.*, vol. 6, pp. 162–176, 2017.
[44] M. Nicola, C. Nicola, A. Vintilă, I. Hurezeanu, and M. Duţă, "Pipeline leakage detection by means of acoustic emission technique using cross-correlation function," *J. Mech. Eng. Auto*, vol. 8, pp. 59–67, 2018.
[45] S. Ramesh, "Acoustic emission testing (AET)," in *Book: Applied Welding Engineering*. Boston, MA, USA: Butterworth-Heinemann, 2012.

Lorenzo Gigli (Graduate Student Member, IEEE) received the master's degree (*summa cum Laude*) in computer science from the University of Bologna, Bologna, Italy, in 2019. He is currently working toward the Ph.D. degree in engineering and information technology for structural and environmental monitoring and risk management from University of Bologna, Bologna, Italy.

He was a Research Fellow on the MAC4PRO project with the University of Bologna. He is part of the IoT PRISM laboratory directed by Mr. Marco Di Felice. His research interests include the Internet of Things, blockchains, and distributed systems.

Ivan Zyrianoff (Graduate Student Member, IEEE) received the B.S. degree in computer science and the M.S. degree in information engineering from the Federal University of ABC, Santo André, Brazil. He is currently working toward the Ph.D. degree in engineering and information technology for structural and environmental monitoring and risk management with the University of Bologna, Bologna, Italy.

His research interests include interoperability for the Internet of Things and edge computing.

Mr. Zyrianoff is a Member of the IoT-Prism Lab. He was involved in the SWAMP Project, an EU-Brazil collaborative research project that developed IoT-based methods for smart water management. He also participated in the Arrowhead Tools project, which aims for the digitalization and automation solutions for the European industry.

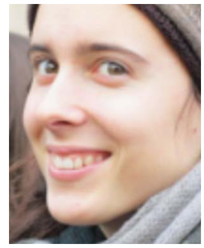

Federica Zonzini (Member, IEEE) received the B.S. and M.S. degrees in electronic engineering and the Ph.D. degree in structural and environmental health monitoring and management (SEHM2) from the University of Bologna, Bologna, Italy, in 2016, 2018, and 2022, respectively.

She is Junior Research Assistant in electronics with the University of Bologna, Bologna, Italy. Her research interests include the design of intelligent sensor systems and edge computing in the context of structural health monitoring, encompassing advanced signal processing, and tiny machine learning.

Denis Bogomolov received the M.Sc. degree in devices, methods of quality control and diagnostics from Omsk State Technical University, Omsk, Russia, in 2017, and the Ph.D. degree in EIT4SEMM (Engineering and Information Technology for Structural and Environmental Monitoring and Risk Management) program from the University of Bologna, Bologna, Italy, 2022.

He is currently a Postdoctoral Researcher with the University of Bologna, Bologna, Italy. The general area of investigation looks at evaluation and monitoring structural integrity of various structures in industrial and civil engineering facilities. His research interests include nondestructive testing, signal processing, acoustic wave theory, and metrology of electrical devices. Particular attention is being paid to how acoustic emission technique can be applied for long-term monitoring and real-time damage assessment.

Nicola Testoni (Member, IEEE) received the M.Sc. degree in microelectronics and the Ph.D. degree in information technology from Bologna University, Bologna, Italy, in 2004 and 2008, respectively.

He is currently an Adjunct Professor with the Department of Electrical, Electronic, and Information Engineering "Guglielmo Marconi," Bologna University. His research interests include guided waves, analog circuit design, nonlinear signal processing, wavelet theory and applications, neural signal denoising, and event sorting.

Marco Di Felice (Member, IEEE) received the *Laurea* and Ph.D. degrees in computer science from the University of Bologna, Bologna, Italy, in 2004 and 2008, respectively.

He is currently a Full Professor of computer science with the University of Bologna, and the Co-Director of the IoT PRISM Laboratory. He has held Visiting Research positions with the Georgia Institute of Technology, Atlanta, Georgia, and Northeastern University, Boston, MA, USA. He has authored or coauthored more than 120 papers. His research focuses on wireless and mobile systems, including self-organizing networks, unmanned aerial systems, IoT, WoT, and context-aware computing.

Luca De Marchi (Member, IEEE) received the M.Sc. and Ph.D. degrees in electronics engineering from the University of Bologna, Bologna, Italy, in 2002 and 2006, respectively.

He is currently an Associate Professor in electronics with the University of Bologna, Bologna, Italy. He has authored more than 200 articles in international journals or the proceedings of international conferences. He holds two patents. His research interests include multiresolution and adaptive signal processing, with a particular emphasis on SHM applications.

Giuseppe Augugliaro received the bachelor's degree in mechanical engineer and the Ph.D. degree in engineering of processes and interaction between energy and environment from the University of Rome Tor Vergata, Rome, Italy, in 1995 and 1999, respectively.

He was the first Technologist at INAIL, Rome, Italy, in the IX laboratory, "Innovative Technologies for Safety," DIT, conducting research in nondestructive testing. Engaged in developing advanced sensor systems for monitoring equipment and infrastructures, and creating procedures for nondestructive testing on metallic and plastic materials. Brings extensive experience in applying the acoustic emission method for pressure equipment verification. Involved in research for periodic requalification procedures for LPG tanks. Leads the working group developing a procedure for visual examination of external surfaces (PEVA) using remotely piloted aircraft systems (drones).

Mr. Augugliaro holds a Level-3 certification in acoustic emission, ultrasonics, radiography, and liquid penetrants according to the international standard UNI EN ISO 9712.

Canio Mennuti received the bachelor's degree in civil engineering from the University of Naples Federico II, Naples, Italy, in 1989.

He is a Technological Manager with INAIL, Rome, Italy, where he leads the Innovative Technologies for Safety Lab. His role involves planning, coordinating, and reviewing research activities in diagnostic technologies for structural integrity.

Mr. Mennuti is the President of the UNI PND Technical Committee, actively contributing to nondestructive testing standards. He is a Board Member of the Italian Center for Non-Destructive Testing and oversees the INAIL Examination Center at the Monte Porzio Catone Research Center. He holds a Level-3 certification in acoustic emission, ultrasonics, radiography, and liquid penetrants, in accordance with UNI EN ISO 9712.

Alessandro Marzani (Member, IEEE) received the M.Sc. (*Laurea*) degree in civil engineering from the University of Bologna, Bologna, Italy, in 2001, the M.Sc. degree in structural engineering from the University of California San Diego, San Diego, CA, USA, in 2004, and the Ph.D. degree in engineering of materials and structures from the University of Calabria, Arcavacata, Italy, in 2005.

He is currently a Professor of structural mechanics with the Department of Civil, Chemical, Environmental and Material Engineering, University of Bologna. His research interests include nondestructive evaluation techniques of materials and structures, structural monitoring, linear and nonlinear ultrasonic guided wave propagation, structural optimization and identification strategies, and structured materials for wave propagation control metamaterials.