2025 IEEE 25th International Symposium on Cluster, Cloud and Internet Computing Workshops (CCGridW) | 979-8-3315-0938-5/25/$31.00 ©2025 IEEE | DOI: 10.1109/CCGRIDW65158.2025.00022

Data-Driven Operational Artificial Intelligence for Computing Continuum: a Natural Disaster Management Use Case

Serena Sebbio*, Lorenzo Carnevale*, *IEEE Member*, Daniel Balouek†,
Manish Parashar‡, Massimo Villari*, *IEEE Member*
*Department of Mathematics, Physics, Informatic and Earth Science, University of Messina, Italy
{serena.sebbio, lcarnevale, mvillari}@unime.it
†IMT Atlantique, Inria, LS2N, UMR CNRS 6004, F-44307 Nantes, France
{daniel.balouek}@inria.fr
‡Scientific Computing and Imaging Institute, University of Utah, USA
{manish.parashar}@utah.edu

***Abstract*—The increasing frequency of documented natural disasters can be attributed to advances in communication technologies, such as satellites, the Internet, and smart devices that facilitate better disaster reporting. This is coupled with an actual rise in the occurrence of such events and improved documentation of their impacts. These trends underscore the pressing need for scalable and intelligent technological solutions to efficiently process large datasets, allowing informed decision-making and effective disaster response. This study presents a Computing Continuum framework that integrates intelligence across cloud, edge and deep edge tiers for efficient disaster data processing. A significant characteristic is the incorporation of Artificial Intelligence for IT Operations (AIOps), which leverages machine learning and analytics to facilitate dynamic resource management and adaptive system modeling, thereby addressing the intricate challenges posed by disaster scenarios. The architecture encompasses an AI-driven framework for monitoring and managing service, network, and infrastructure layers, tailoring policies to specific disaster needs. The proposed framework is applied to wildfire management, leveraging an AI Operation Manager to coordinate sensor-equipped drones for real-time data acquisition and processing. Operating at the deep edge tier, these drones transmit environmental data to edge and cloud infrastructures for analysis. This multi-tiered approach improves situational awareness, disaster response, and resource utilization.**

***Index Terms*—Natural Disaster Management, Artificial Intelligence, Computing Continuum, Software Architecture, AIOps**

I. Introduction

The frequency and impact of natural disasters have escalated significantly in recent decades. According to the Institute for Economics and Peace [1], the number of recorded disasters increased from 39 in 1960 to 396 in 2019, with peak events in 2005 resulting in 90,000 fatalities and 160 million people requiring urgent assistance. The economic burden has also risen sharply, with annual damage costs growing from $50 billion in the 1980s to $200 billion in the last decade. These trends highlight the necessity of enhancing natural disaster management (NDM) through computational approaches that improve real-time response efficiency, optimize resource allocation, and support decision-making under uncertainty. The main issue in NDM is the inefficiency of traditional computing architectures in handling real-time, large-scale, and dynamically evolving disaster scenarios, preventing the timely analysis and dissemination of critical data collected from heterogeneous sources, essential for interactive decision support systems. Existing solutions rely on centralized cloud infrastructures, introducing latency and bandwidth limitations that delay first responders' access to actionable insights. Additionally, the lack of Data-driven intelligence across the computing continuum hinders efficient information sharing and adaptive resource allocation, while limiting the ability to integrate and visualize extreme data sources in real time. These shortcomings compromise situational awareness and reduce the effectiveness of disaster response operations.

To address these limitations, this paper proposes a novel software architecture that integrates AI-driven decision-making across the computing continuum, distributing computational workloads efficiently. By leveraging deep edge, edge, and cloud resources, the framework enhances responsiveness in dynamic and resource-constrained environments, improving situational awareness and real-time decision support. Furthermore, it facilitates data fusion, visualization, and autonomous system reconfiguration, enabling first responders and coordination centers to process and act upon extreme data streams more effectively. This work presents the conceptual design of a three-tier hybrid architecture, comprising deep edge, edge, and cloud layers, along with service, network, and infrastructure layers orchestrated by an operation manager.

The rest of the paper is organized as follows. Section II details the motivation for this architectural design. Related works are discussed in Section III, followed by an overview of the proposed architecture in Section IV. Although a comprehensive implementation and evaluation extend beyond the scope of this preliminary design paper, Section V provides a use case illustrating its application in an NDM scenario. This includes a subsection that elaborates on the microservices,

979-8-3315-0938-5/25/$31.00 ©2025 IEEE
DOI 10.1109/CCGridW65158.2025.00022

delineates a potential implementation scenario, and outlines key elements for future evaluation in accordance with this architectural blueprint. Finally, Section VI concludes the paper, highlighting ongoing efforts and future research direction, including the potential implementation strategies and rigorous evaluation of the proposed architecture.

II. Motivation

In the current era of climate change, the NDM is becoming increasingly important at a global level to address the new challenges that are appearing in Nature. In fact, first responders are in urgent need of technological support [2].

End-users are, typically, highlighting a pool of needs, such as (a) the ability to collect and share vital information, (b) the lack of an enhanced sharing information tool, (c) the fusion of many heterogeneous extreme data sources and (d) advanced and interactive visualization systems, combining data fusion capabilities for enabling advanced real-time decision support features. On the other hand, a natural disaster scenario is typically characterized by energy- and Internet-disconnected areas, where battery-powered devices (e.g., drones, weather stations) are used for supporting first responders (e.g., civil protection, firefighters) during the exploration and rescue missions. At the same time, a coordination center monitors the operations from remote, often facing the issue of transmitting data over a satellite network, which increases the latency and reduces the possibility to make decisions in real-time. Our solution meets those requirements, by developing a cloud-edge intelligence platform. Furthermore, the management of NDM tasks in the aforementioned scenario is facilitated by pipelines that handle data in near real-time across various devices, libraries, and end-user entry points. Balancing local data collection and aggregation decisions with global optimizations that incorporate learning for system operations necessitates a software architecture that operates both horizontally to distribute data processing across the Computing Continuum and vertically by utilizing Artificial Intelligence capabilities for cross-optimizations between the infrastructure and application services. The following technical requirements must be integrated to achieve a cohesive vision of NDM:

- **Integrating learning for local decisions and global optimizations** for gaining insights and knowledge of processes while reconfiguring applications under new events or constraints.
- **Support geographically distributed deployment and device mobility** to handle dynamically different sizes of systems and fragmented applications.

III. State of the Art

The Internet of Things (IoT) is commonly linked with utilizing Machine Learning algorithms to analyze data, derive insights, and predict patterns future occurrences [3]–[5]. The progressive convergence of cloud computing and IoT devices is resulting in the computing continuum [6]–[8]. The concept of edge computing first became to represent a middle layer, where data and computation were migrated in runtime, and later was fused with the IoT devices, because these increasingly acquired computational capacity. In the context of natural disasters, selected scientific works are available in literature to leverage the Computing Continuum for natural disasters. Løvholt et al. proposed a prototype for utilizing real-time seismic parameters and HPC in tsunami early warning and rapid post disaster assessments, faster than the physical propagation time of a tsunami [9]. Balouek et al. has proposed a solution for harnessing the computing continuum for urgent science. An Early Earthquake Warning (EEW) workflow was used as a driver for proposing a system stack that can enable the fluid integration of distributed analytics across a dynamic infrastructure spanning the computing continuum [10]. They extend this work to introduced a hierarchical approach for modeling distributed stream-based applications on a large set of heterogeneous data processing and management frameworks across the network. [11]. Babu et al. focused on using IoT home units to design a Flood and earthquake Observatory System. Results presented encouraging performance, along with the integration of solar energy for emergency situations [12]. Balouek et al. discussed research directions for tackling a fire science scenario that includes sensors at the edge of the network for smoke detection, and computational models launched in the cloud for wildfire simulation and air quality assessment [13].

The proposed software architecture differs from the ones discussed in this Section because of the exploiting of operational data. The intent is re-configuring the system in runtime exploiting AI for operational data by building configuration according to the monitor of resources and running services.

IV. Architecture

In the following, we describe the software architecture designed to distribute data processing across the Computing Continuum by leveraging Artificial Intelligence capabilities. The architecture meets the requirements described in Section II.

A. Cloud-Edge-Deep Edge Tiers

Figure 1 shows the interactions in the Computing Continuum, composed of two edge tiers and one cloud tier. This hierarchical structure supports real-time decision-making by addressing constraints such as latency, connectivity, and energy efficiency.

- **Deep Edge**: This tier consists of autonomous and/or mobile devices responsible for data collection and local inference of machine learning models. These devices, though commonly identified as IoT (e.g., Raspberry Pi, Jetson Nano), possess significant computational capacity.
- **Edge Tier**: It comprises one or more base stations, such as portable computers or small-scale data centers, facilitating real-time data aggregation and intelligent processing while communicating with mobile devices (e.g., unmanned aerial vehicles) via wireless telemetry.

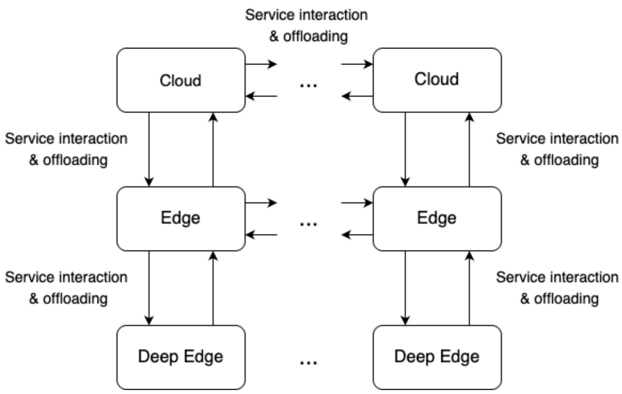

Fig. 1. The Cloud-Edge-Deep Edge tiers are thought to distribute the computation at different distances from the data source. The offloading may happen by following a horizontal or vertical logic.

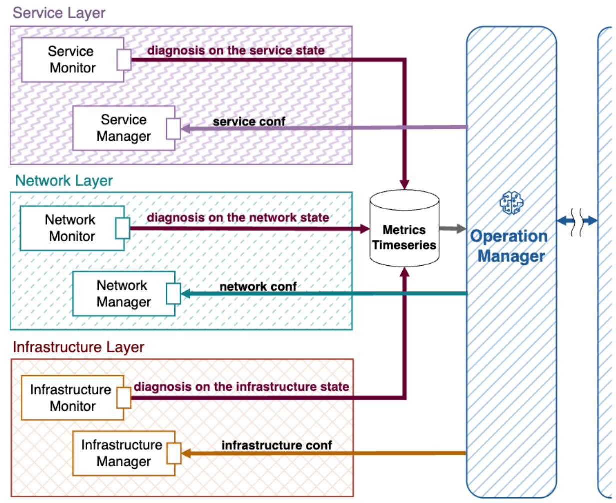

Fig. 2. Layered architecture within a single node, composed of Infrastructure, Network, and Service layers. The AI-driven Operation Manager enables autonomous optimization and interaction with other nodes across the system.

- **Cloud Tier**: This tier provides remote support and resources for massive computation requirements, ensuring scalability and resilience.

Computation is handled by distributed computational units (e.g., agents) deployed across all tiers. By ensuring that critical processing occurs as close as possible to data sources, the architecture minimizes delays and enhances situational awareness. Tasks are executed following a deep edge-to-edge-to-cloud offloading strategy, prioritizing local execution unless resource constraints (e.g., battery limitations, network failures) necessitate offloading to higher tiers.

B. Architecture for Intelligent Intra-Node Management

This subsection details the layered architecture within a single node, regardless of whether it is deployed at the Deep Edge, Edge, or Cloud tier. Each node follows a structured model that includes the following layers:

- **Infrastructure Layer**: Manages computational resources such as processing power, memory, and storage.
- **Network Layer**: Governs communication, connectivity, and data transmission.
- **Service Layer**: Executes application logic and and manages workloads based on system demands.

They are all orchestrated by an intelligent cross-layer component called the Operation Manager. This design ensures efficient local resource allocation while enabling coordination with other nodes in the distributed system, as illustrated in Figure 2.

Each layer integrates:

- **Resource Monitor** – Tracks real-time operational data within the layer, providing essential metrics for performance assessment.
- **Layer Manager** – Implements configuration adjustments generated by the Operation Manager to refine deployment strategies.

The intelligent cross-layer component, the **Operation Manager**, processes operational data from all three layers to optimize configurations, ensure efficient workload management, and enable dynamic adaptation. This feedback loop allows for continuous adaptation based on real-time data, ensuring optimal resource utilization and system resilience. AI-driven policy learning plays a key role in optimizing configurations. By analyzing key performance indicators (KPIs) such as latency, resource availability, and network congestion, the system dynamically adjusts deployments to match operational demands. Additionally, computational units within the Edge and Cloud tiers interact through global AI cross-layers, enabling efficient workload distribution across nodes.

C. Operation Manager

The **Operation Manager** is the core intelligence of the framework, consisting of multiple AI agents that process time-series operational data from each layer. By leveraging real-time analytics and machine learning models, it facilitates adaptive decision-making for resource allocation, workload offloading, and system reconfiguration, as illustrated in Figure 3.

Each AI agent is structured into two primary components: a Predictive Neural Network **(Predictive NN)** and an **Anomaly Detection** Module. The Predictive NN is responsible for forecasting system behavior based on historical and real-time data, allowing for proactive adjustments in resource allocation. Meanwhile, the Anomaly Detection Module continuously monitors deviations from expected performance, detecting potential faults or inefficiencies in system operations. These two components work in tandem, ensuring both predictive optimization and rapid response to unexpected conditions. Each AI agent focuses on optimizing a specific subset of resources (e.g., CPU, memory, network bandwidth, energy

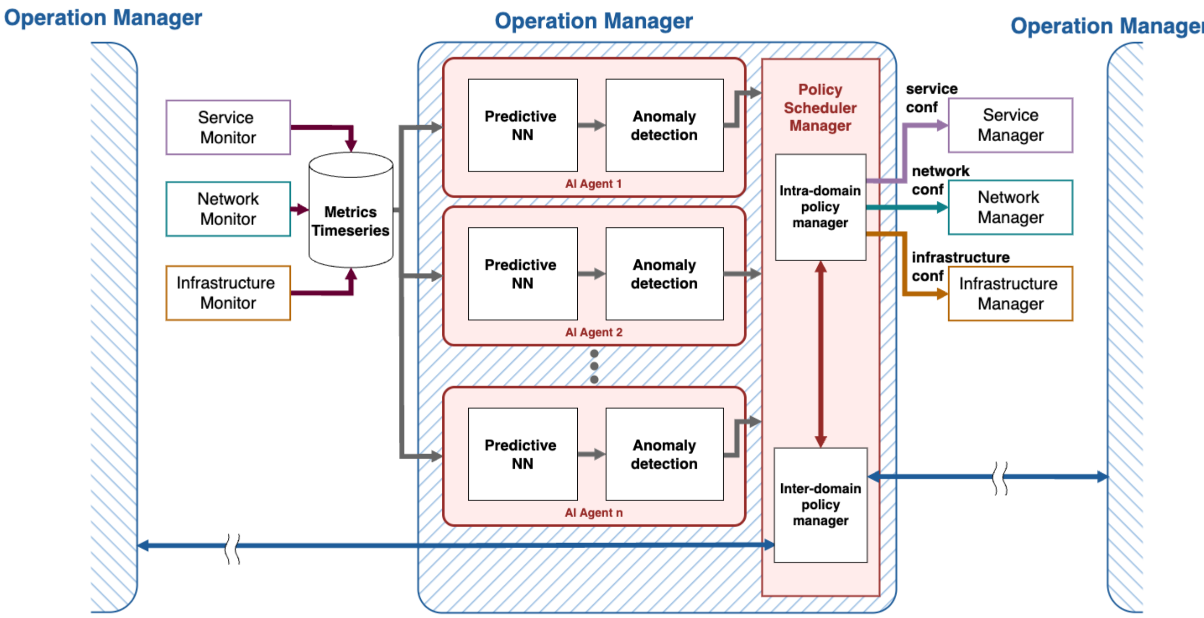

Fig. 3. The Operation Managers are cross-layers components that interact with each other for evolving the architecture behavior as if it is a complex system.

consumption). The agents collaborate through an ensemble learning approach, ensuring that system-wide decisions align with operational priorities.

The **Policy Scheduler Manager** responsible for synthesizing AI-generated insights and generating configuration updates that align with the system's objectives, operates at two levels:

- **Intra-domain Policy Management:** Regulates resources and task allocation within a specific layer (Deep Edge, Edge, or Cloud), ensuring local optimization of computing nodes.
- **Inter-domain Policy Management:** Coordinates resource-sharing and workload distribution across different layers, enabling dynamic task offloading and cross-layer cooperation.

The Operation Manager is not limited to local optimization. Each node in the system—whether deployed in the Deep Edge, Edge, or Cloud—can exchange information with the Operation Managers of other nodes, forming a distributed intelligence network. This capability allows nodes to collaboratively manage workloads, offload tasks dynamically, and balance system-wide resource utilization, ensuring efficient decision-making across the entire computing continuum. By leveraging this intelligent intra-node management architecture, the system allows real-time adaptability, energy-efficient operation, and resilient performance, making it particularly suitable for scenarios requiring autonomous decision-making in resource-constrained environments, such as emergency response and large-scale monitoring applications.

D. AI-Driven Resource Optimization

AI-driven decision-making in resource optimization relies on key operational data:

- **Service Performance Metrics**: Metrics such as latency, throughput, and availability guide intent-based optimization, ensuring user experience enhancements [14]. Instead of defining resource requirements explicitly, high-level intents are translated into policies based on continuous KPI monitoring [15].
- **Network Performance Metrics**: Connectivity parameters like bandwidth, latency, congestion, and packet loss enable real-time adjustments for efficient data transmission [16]. Intent-driven orchestration dynamically optimizes routing and congestion control, ensuring alignment with service-level objectives [17].
- **Infrastructure Performance Metrics**: Insights into CPU, memory, storage, and energy consumption support dynamic resource allocation and energy-aware provisioning [15]. Proactive fault detection minimizes downtime, enhancing reliability across cloud and edge infrastructures [17]–[19].

By correlating these data sources, the Operation Manager continuously refines decision-making, adapting resource provisioning and workload distribution to changing conditions. Optimization strategies depend on the specific operational data considered, particularly in application-specific scenarios such as natural disaster management.

E. Intent-Based Scheduling

The **Policy Scheduler Manager** translates high-level operational goals (e.g., minimizing response time, maximizing energy efficiency) into concrete resource allocation strategies. This intent-based approach ensures that system adaptations align with real-time mission objectives.

- **Predictive Analytics**: AI models simulate scheduling scenarios to anticipate bottlenecks and optimize workload distribution before issues arise.
- **Adaptive Security**: Dynamic policy adjustments enhance system resilience against cyber threats and ensure secure data sharing.
- **Continuous Learning**: A feedback loop refines scheduling policies over time, integrating insights from both AI models and real-world operational data.

By integrating AI-driven intelligence at every layer, the proposed architecture empowers real-time, data-driven decision support, improving operational efficiency and situational awareness in complex and dynamic environments.

The operation data are not only used standalone. They are often considered in combination because the state of a layer may affect the state of another layer. For example, Service Operation may be correlated with Network Operation in terms of latency or bandwidth constraints, optimizing the data transfer protocol. Network Operation may be correlated with the Infrastructure Operation to coordinate network resource provisioning based on service-level demands and infrastructure constraints. Infrastructure Operation may be correlated with the Service Operation in terms of provisioning recommendations, planning or container orchestration. Additionally, Infrastructure Operation may be correlated with the Network Operation in terms of alignment between network resource provisioning and infrastructure capacity planning.

By understanding the capabilities and limitations of the layers, the *Operation Manager* adapts its decision-making processes and optimization strategies to ensure alignment with the available resources, the constraints (e.g., service level agreement).

The Policy Scheduler Manager is a key component within the Operation Manager, responsible for optimizing resource allocation and scheduling across the infrastructure, network, and service layers. It synthesizes insights from the three layers—Service, Network, and Infrastructure—by processing operational data and intent-based directives, coordinating the actions of multiple AI agents, ensuring that their decisions are aligned and do not lead to resource contention or inefficiencies [20].

Using an intent-based approach, the scheduler translates high-level user goals (e.g., reducing latency, maximizing availability) into concrete resource allocation strategies [17]. By integrating real-time operational data, it dynamically adjusts scaling policies and resource distribution based on factors like CPU, memory, energy consumption, disk, and network I/O.

The combination of intent-based networking and digital twin technology supports predictive analysis by simulating scheduling scenarios, allowing for proactive adjustments that optimize performance and mitigate bottlenecks [19]. Additionally, intent-based scheduling applies security policies by dynamically adjusting network segmentation, reducing attack surfaces, and minimizing conflicts between cluster and application administrators.

A continuous feedback loop refines the system's decision-making, integrating both AI-driven subjective assessments and objective evaluations from digital twins. This ensures that scheduling remains adaptive to evolving operational demands while aligning with user-defined intent policies [17], [19].

V. Use Case and Discussion

In the context of NDM, the deployment of an Operation Manager provides a structured framework for facilitating the deployment of edge devices, such as drones equipped with sensors for real-time fire detection and monitoring, in wildfire management scenarios. In fact, leveraging an AI Operation Manager can significantly enhance the coordination, data exchange and decision-making processes across different layers of the architecture, enhancing responsiveness in fire detection and containment efforts.

A. Case Study: Wildfire Management

Wildfires represent a significant global threat, particularly in regions prone to extreme heat and dry weather conditions. Early detection and real-time monitoring are essential to mitigate damage, protect human life, and allocate firefighting resources effectively. Authorities typically deploy a network of drones equipped with advanced sensors to provide first-responding support in detecting and monitoring fire outbreaks. The proposed framework is well-suited to the challenges of wildfire management, enabling real-time data acquisition, processing, and decision-making across multiple layers of computational infrastructure.

Drones operate as deep-edge devices, collecting data from their surroundings and transmitting them to edge and cloud infrastructure for analysis and decision-making. Computational resources are deployed both on the edge for first-response actions and in the cloud for management and information fusion with, e.g., social network posts and satellite images. Table I summarizes the hardware components deployed across the Deep Edge, Edge, and Cloud tiers in this wildfire management scenario.

TABLE I
HARDWARE COMPONENTS BY COMPUTATIONAL TIER

Tier	Device Type	Purpose
Deep Edge	UAV Drones	Image acquisition, environmental sensing
	Ground Sensors	Temperature, humidity, wind speed/direction
Edge	Base Station (laptop)	Fire detection, smoke segmentation
	Raspberry Pi rack	Fire detection, sensor statistics
Cloud	Data Center VMs	Fire simulation, information fusion, decision support

B. Microservice-Based Modular Approach

To ensure flexibility and efficient resource allocation in the proposed framework, a modular microservice-based approach is adopted. Each computational process, spanning from deep-edge devices to cloud infrastructure, is structured as an independent service that communicates through standardized interfaces.

This modularization enables dynamic workload distribution and efficient orchestration of tasks, ensuring real-time responsiveness. Each task is isolated and can be scaled independently depending on demand: for example, if a particular wildfire incident escalates, additional containers running fire detection microservices can be instantiated in the edge or cloud tier to handle the increased workload. The core services include:

- **Data Acquisition Services:** Responsible for collecting data from drones, ground sensors, and satellites. These services preprocess raw data at the deep edge before transmitting it to higher layers.
- **Fire and Smoke Detection Services:** Deployed on both edge and cloud tiers, these services analyze sensor and image data to identify potential fire outbreaks.
- **Decision Support Services:** Hosted in the cloud, these services integrate multiple data sources, including social network reports and environmental simulations, to provide actionable insights for disaster management authorities.
- **Communication and Coordination Services:** Facilitate interaction between different layers, ensuring synchronization and secure data exchange across the infrastructure.

Table II outlines the primary microservices deployed in each tier, along with their respective functionalities.

TABLE II
MICROSERVICES AND THEIR MAPPING TO COMPUTATIONAL TIERS

Microservice	Deep Edge	Edge	Cloud
Drone-based Image Acquisition	✓		
Drone-based Data Acquisition	✓		
Ground Sensor Data Acquisition	✓		
Fire Detection		✓	✓
Smoke Detection		✓	✓
Person Detection		✓	✓
Car Detection		✓	✓
Fire Segmentation		✓	✓
Smoke Segmentation		✓	✓
Person Re-identification		✓	✓
Sensor Statistics		✓	✓
Drone Planning	✓	✓	✓
Visualizer (Edge)		✓	
Fire Simulation			✓
Information Fusion			✓
Satellite-based Fire Detection			✓
Decision Support			✓
Visualizer (Cloud)			✓

C. Multi-Tier AI-Driven Coordination for Wildfire Response

1) Deep Edge Tier: Devices and Services: The *Deep Edge tier* consists of Unmanned Aerial Vehicles (UAVs) and ground sensor networks responsible for real-time data acquisition and initial processing.

- *Devices*: UAVs, equipped with cameras and AI kits (e.g., Raspberry Pi, Jetson Nano Orin), capture images and videos. Ground sensors monitor temperature, humidity, wind speed, and wind direction, detecting anomalies such as sudden temperature spikes or wind pattern changes.
- *Services*: Despite their limited computational power, Deep Edge devices perform low-level AI-based inference and execute microservices for early smoke and fire detection. When necessary, they offload computationally demanding tasks to Edge base stations.

2) Edge Tier: Devices and Services: The *Edge tier* acts as an intermediary layer for real-time data processing and decision-making.

- *Devices*: Edge base stations receive images from UAVs and run fire and smoke detection algorithms. These stations may consist of laptops, NVIDIA Jetson Orin devices, or small ARM/Intel-based clusters.
- *Services*: This tier performs fire segmentation, sensor data aggregation, and UAV mission planning. If local processing is insufficient, tasks are offloaded to the Cloud or horizontally distributed among other Edge nodes.

3) Cloud Tier: Devices and Services: The *Cloud tier* provides high-performance computing for advanced analytics and decision support.

- *Devices*: Cloud data centers house GPU-accelerated virtual machines (VMs) for intensive computations.
- *Services*: Tasks include fire spread simulations, multi-source data fusion (satellite imagery, UAV footage, sensor data), and decision support for authorities. The Cloud also serves as a repository for long-term storage and AI model training.

D. Role of the Layers

- *Infrastructure Layer*: Evaluates energy-performance trade-offs for UAVs and other resource-constrained devices, adapting deployment strategies based on environmental conditions.
- *Network Layer*: Manages communication between Edge and Cloud components, ensuring efficient bandwidth allocation and minimizing latency.
- *Service Layer*: Hosts applications for fire detection, monitoring, and emergency response, leveraging insights from the Infrastructure and Network layers.

When talking about the Infrastructure Layer, it must fit, e.g., the trade-off between energy consumption and performance to maximize the effectiveness of the deployed drones. Therefore, it evaluates factors such as energy availability and response time requirements to, e.g., optimize and scale drone deployment. That is, the evolution of the system is based, e.g., on environmental conditions and operational demand to ensure optimal allocation and scalability [21].

E. Intelligent Task Offloading

To maintain system efficiency, AI-driven *predictive models* (e.g., LSTM networks) monitor and forecast CPU usage, memory consumption, and network bandwidth. These models are updated in real-time to account for changing environmental conditions and workloads. The AI system uses this information to decide when to offload tasks to other tiers, ensuring that devices in the deep edge and edge tiers do not become overburdened by high computational loads. For instance, if the Network Monitor reports high latency or low bandwidth, the *Operation Manager* might offload tasks to a local edge base station rather than the cloud. Similarly, if Infrastructure Monitor detects low battery levels in a UAV, the *Operation Manager* may reduce the frequency of data transmissions or offload more processing to the edge or cloud tiers. The integration of predictive AI models with real-time resource monitoring ensures that the system operates efficiently, even in highly dynamic and resource-constrained environments such as wildfire management.

F. Achievements

The proposal of such architecture holds promise in improving coordination among the various stakeholders involved in disaster response and mitigation. The integration of edge devices, such as drones, with sensors and a cloud-edge intelligence platform is envisioned to enable real-time data collection, transmission and analysis, potentially facilitating timely decision-making during emergencies. Furthermore, the conceptualization of resource allocation strategies aims to optimize performance while considering energy constraints, potentially enhancing the effectiveness of deployed resources. Addressing critical needs in NDM, such as efficient information sharing and interactive visualization for decision support, underscores the potential progress that could be made to improve the overall response capabilities.

G. Challenges

Despite the strengths of our proposed architecture, several challenges must be addressed to ensure effective deployment and long-term performance. Key challenges include:

- **AI Agent Coordination and Synchronization**: Coordinating multiple AI agents across Deep Edge, Edge, and Cloud layers introduces complexity. Each agent may act autonomously, leading to conflicting resource allocation strategies. To mitigate this,adopting *Multi-Agent Reinforcement Learning (MARL)* can enable agents to collaboratively learn optimal workload distribution strategies, improving synchronization across layers. Prior research has shown that MARL techniques can significantly enhance multi-agent coordination in complex environments [22], [23]. Additionally, federated learning could support decentralized AI model training, enhancing decision-making capabilities while reducing data transfer overhead [24].
- **Real-Time Data Synchronization**: Disaster scenarios are highly dynamic, requiring rapid propagation of data from sensors to the Cloud. By integrating Prometheus for real-time metric collection and MQTT for low-latency messaging, we aim to minimize delays and improve synchronization.
- **Workload Adaptability in Disaster Scenarios:** Disaster scenarios are highly unpredictable, causing sudden surges in data volume and computational demands. Efficient resource management requires rapid adaptation strategies. Implementing *hierarchical orchestration models*—where edge nodes manage immediate processing while cloud infrastructure handles high-level analytics—has proven effective in reducing latency during peak loads ([25]). Developing predictive scheduling algorithms that anticipate resource demands can further improve responsiveness.
- **Real-World Validation and Adoption:** The successful deployment of our framework requires collaboration among AI researchers, emergency responders, and policymakers. Establishing pilot programs in wildfire-prone regions can provide practical insights into system effectiveness. Additionally, incorporating *human-in-the-loop interfaces* can improve trust and usability by allowing emergency personnel to interact seamlessly with AI-driven recommendations. Aligning AI decision-making with regulatory frameworks and ethical guidelines is essential for ensuring transparency and accountability.

VI. Conclusion

In conclusion, the application of a software architecture equipped with AI Operational Managers in NDM presents a promising approach to address critical challenges in disaster response and mitigation. By leveraging advanced technologies and intelligent AI Operation managers across different layers, our proposed solution enhances coordination, efficiency and adaptability in managing natural disasters.

The deployment of edge devices, such as drones equipped with sensors, coupled with a cloud-edge intelligence platform, facilitates real-time data collection, transmission and analysis, essential for timely decision-making during emergencies. The integration of adaptive resource allocation strategies ensures optimal performance while considering the energy constraints inherent in battery-powered devices.

In addition, our solution addresses key NDM needs, including efficient information sharing, data fusion from heterogeneous sources and interactive visualization for decision support. By providing near-real-time data processing capabilities and enabling cross-layer optimizations, our architecture improves the resilience and effectiveness of disaster management systems.

As for future work, further exploration could involve the development of more detailed design specifications using the Unified Modeling Language (UML), facilitating comprehensive system modeling and refinement. Furthermore, ongoing research could focus on enhancing the scalability and robustness of the proposed architecture to handle larger-scale disaster scenarios and accommodate evolving technological advancements. Implementing the proposed framework would also be a crucial future step in evaluating its effectiveness in real-world disaster management scenarios.

Overall, our study underscores the importance of adopting innovative approaches in disaster management, leveraging a data operational driven AI-based architectures to improve responsiveness and ultimately mitigate the impact of natural disasters on communities and infrastructure.

ACKNOWLEDGMENT

This work is partially supported by the Horizon Europe "Trusted Extremely Precise Mapping and Prediction for Emergency Management" project (Grant Agreement, project 101093003 - TEMA), the Horizon Europe NEUROKIT2E project (Grant Agreement 101112268), the "SEcurity and RIghts in the CyberSpace (SERICS)" project (PE00000014), under the MUR National Recovery and Resilience Plan funded by the European Union - NextGenerationEU, and the Italian Ministry of University and Research (MUR) "Research projects of National Interest (PRIN-PNRR)" through the project "Cloud Continuum aimed at On-Demand Services in Smart Sustainable Environments" (CUP: J53D23015080001-IC: P2022YNBHP). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union - European Commission. Neither the European Commission nor the European Union can be held responsible for them.

REFERENCES

[1] Institute for Economics Peace. Ecological threat register 2020: Understanding ecological threats. Resilience and Peace, Sydney, September 2020.

[2] Christian Sicari, Alessio Catalfamo, Lorenzo Carnevale, Antonino Galletta, Daniel Balouek-Thomert, Manish Parashar, and Massimo Villari. Tema: Event driven serverless workflows platform for natural disaster management. In *2023 IEEE Symposium on Computers and Communications (ISCC)*, pages 1–6. IEEE, 2023.

[3] Mehdi Mohammadi, Ala Al-Fuqaha, Sameh Sorour, and Mohsen Guizani. Deep learning for iot big data and streaming analytics: A survey. *IEEE Communications Surveys & Tutorials*, 20(4):2923–2960, 2018.

[4] Yogesh S. Lonkar, Abhinav Sudhakar Bhagat, and Sd Aasif Sd Manjur. Smart disaster management and prevention using reinforcement learning in iot environment. In *2019 3rd International Conference on Trends in Electronics and Informatics (ICOEI)*, pages 35–38, 2019.

[5] Tausifa Jan Saleem and Mohammad Ahsan Chishti. Deep learning for the internet of things: Potential benefits and use-cases. *Digital Communications and Networks*, 7(4):526–542, 2021.

[6] Daniel Balouek-Thomert, Eduard Gibert Renart, Ali Reza Zamani, Anthony Simonet, and Manish Parashar. Towards a computing continuum: Enabling edge-to-cloud integration for data-driven workflows. *International Journal of High Performance Computing Applications*, 33:1159–1174, 11 2019.

[7] Pete Beckman, Jack Dongarra, Nicola Ferrier, Geoffrey Fox, Terry Moore, Dan Reed, and Micah Beck. Harnessing the computing continuum for programming our world. *Fog Computing: Theory and Practice*, pages 215–230, 2020.

[8] Schahram Dustdar, Victor Casamayor Pujol, and Praveen Kumar Donta. On distributed computing continuum systems. *IEEE Transactions on Knowledge and Data Engineering*, 35(4):4092–4105, 2022.

[9] Finn Løvholt, Stefano Lorito, Jorge Macias, Manula Volpe, Jacopo Selva, and Steven Gibbons. Urgent tsunami computing. In *2019 IEEE/ACM HPC for Urgent Decision Making (UrgentHPC)*, pages 45–50, 2019.

[10] Daniel Balouek-Thomert, Ivan Rodero, and Manish Parashar. Harnessing the computing continuum for urgent science. *ACM SIGMETRICS Performance Evaluation Review*, 48:41–46, 11 2020.

[11] Daniel Balouek-Thomert, Pedro Silva, Kevin Fauvel, Alexandru Costan, Gabriel Antoniu, and Manish Parashar. Mdsc: Modelling distributed stream processing across the edge-to-cloud continuum. *ACM International Conference Proceeding Series*, 12 2021.

[12] Venita Babu and Vishnu Rajan. Flood and earthquake detection and rescue using iot technology. In *2019 International Conference on Communication and Electronics Systems (ICCES)*, pages 1256–1260, 2019.

[13] Daniel Balouek-Thomert, Eddy Caron, Laurent Lefevre, and Manish Parashar. Towards a methodology for building dynamic urgent applications on continuum computing platforms. *Proceedings of CIW-IUS 2022: Combined International Workshop on Interactive Urgent Supercomputing, Held in conjunction with SC 2022: The International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 25–30, 2022.

[14] Thijs Metsch, Magdalena Viktorsson, Adrian Hoban, Monica Vitali, Ravi Iyer, and Erik Elmroth. Intent-driven orchestration: Enforcing service level objectives for cloud native deployments. *SN Computer Science*, 4(3), 2023.

[15] Mohammad S. Aslanpour, Sukhpal Singh Gill, and Adel N. Toosi. Performance evaluation metrics for cloud, fog and edge computing: A review, taxonomy, benchmarks and standards for future research. *Internet of Things (Netherlands)*, 12, 2020.

[16] Gabriel Ioan Arcas, Tudor Cioara, Ionut Anghel, Dragos Lazea, and Anca Hangan. Edge offloading in smart grid. *Smart Cities*, 7(1):680 – 711, 2024.

[17] Negin Akbari, Adel N. Toosi, John Grundy, Hourieh Khalajzadeh, Mohammad S. Aslanpour, and Shashikant Ilager. Icontinuum: An emulation toolkit for intent-based computing across the edge-to-cloud continuum. In *Proceedings of the IEEE International Conference on Cloud Computing (CLOUD)*, page 468 – 474, 2024.

[18] Hui Song, Ahmet Soylu, and Dumitru Roman. Towards cognitive self-management of iot-edge-cloud continuum based on user intents. In *Proceedings of the 2022 IEEE/ACM 15th International Conference on Utility and Cloud Computing (UCC)*, page 313 – 316, 2022.

[19] Jose Santos, Eddy Truyen, Christoph Baumann, Filip De Turck, Gerald Budigiri, and Wouter Joosen. Towards intent-based scheduling for performance and security in edge-to-cloud networks. In *Proceedings of the 27th Conference on Innovation in Clouds, Internet and Networks (ICIN)*, page 222 – 227, 2024.

[20] Xiangning Li, Yuqian Cai, and Jingyue Tian. Computing force network self-intelligence based on intent driven and digital twin. In *Proceedings of the 2023 3rd International Conference on Electronic Information Engineering and Computer Science (EIECS 2023)*, page 929 – 936, 2023.

[21] Marc Gamell, Ivan Rodero, Manish Parashar, and Rajeev Muralidhar. Exploring cross-layer power management for pgas applications on the scc platform. In *Proceedings of the 21st International Symposium on High-Performance Parallel and Distributed Computing*, HPDC '12, page 235–246, New York, NY, USA, 2012. Association for Computing Machinery.

[22] Chao Zeng, Xingwei Wang, Rongfei Zeng, Ying Li, Jianzhi Shi, and Min Huang. Joint optimization of multi-dimensional resource allocation and task offloading for qoe enhancement in cloud-edge-end collaboration. *Future Generation Computer Systems*, 155:121 – 131, 2024.

[23] Amanda Jayanetti, Saman Halgamuge, and Rajkumar Buyya. Multi-agent deep reinforcement learning framework for renewable energy-aware workflow scheduling on distributed cloud data centers. *IEEE Transactions on Parallel and Distributed Systems*, 35(4):604 – 615, 2024.

[24] Zhenguo Ma, Yang Xu, Hongli Xu, Jianchun Liu, and Yinxing Xue. Like attracts like: Personalized federated learning in decentralized edge computing. *IEEE Transactions on Mobile Computing*, 23(2):1080 – 1096, 2024.

[25] Santanu Ghosh and Pratyay Kuila. Efficient offloading in disaster-affected areas using unmanned aerial vehicle-assisted mobile edge computing: A gravitational search algorithm-based approach. *International Journal of Disaster Risk Reduction*, 97, 2023.