Truong, Linh; Magoutis, Kostas

A Multi-Continuum View for Swarm-Edge-Cloud Service-based Applications

Published: 01/01/2025

Document Version
Early version, also known as pre-print

Please cite the original version:
Truong, L., & Magoutis, K. (2025). *A Multi-Continuum View for Swarm-Edge-Cloud Service-based Applications*.

A Multi-Continuum View for Swarm-Edge-Cloud Service-based Applications

Hong-Linh Truong, *Aalto University, Finland*

Kostas Magoutis, *University of Crete & FORTH-ICS, Greece*

Abstract—We leverage various enablers of network connectivity, IoT, and AI inference, to develop service-based applications, deployed atop swarm-edge-cloud infrastructures, as solutions in key domains. A majority of works focus on distributing functionality across swarm-edge-cloud computing resources, on the so-called computing continuum. However, opportunities to explore the computing continuum must be matched with capabilities to achieve time continuum and intelligence continuum, which are not formally acknowledged and included in the contemporary discussion. This paper highlights the new direction of supporting time, intelligence and computing capabilities (i.e., the multi-continuum view) for achieving robustness and reliability of swarm-edge-cloud applications. This will allow computing-continuum applications to also operate continuously over long horizons (the time continuum) and to improve their automation and manageability via use of AI in concert with human experts (the intelligence continuum). We discuss how to strengthen existing enablers for implementing multi-continuum computing with a focus on a few coordination-centric components.

Introduction

Tools and concepts of edge-cloud computing continuum are nowadays well known. They aim at delivering a seamless integration of edge and cloud computing resources with the support of network connectivity among the edge's and the cloud's compute nodes. A developer can utilize Internet of Things (IoT) data sources, messaging systems with NATS, MQTT or Apache Kafka, and containerized and Kubernetes-managed computing resources for launching distributed (near real time) data workflows, IoT-based machine learning (ML) pipelines, or distributed microservice-based applications in a cross-infrastructure of edge and cloud compute nodes; the so-called computing continuum [1] as we know it today. However, acquiring diverse types of computing capabilities dynamically, when needed, does not necessarily achieve the needs of continuum for the mission carried out by the application execution, even with the current vision of computing continuum [2]. For example, a long-running critical mission, such as the very common search and rescue (S&R), requires adaptability during long-lasting execution (or the ability to operate in a *time continuum*), maintaining the application's state until its completion, whenever that might be. We are still far from achieving smooth operations in terms of the time continuum however: S&R systems and human operators still have to spend great efforts to (often manually) save, load and link snapshots/chunks of collected data from multiple sessions of a long mission carried out by a complex application, due to the interruption of workflows within the long mission. Or, for a long running mission, we are forced to break down the design into separate applications, due to the lack of supporting continuum mechanisms. There are other examples of the need to heavily involve human operators, besides the increasing use of AI/GenAI to provide additional knowledge and smart decisions within a mission, such as to interpret the rescue time series sensing measurements and log records, and near real time images and videos captured by the drones via AI detection and analytics. Consider the recent development of GenAI/LLMs, AI Agents and Agentic AI systems [3], which are not only available via cloud APIs but also on edge computing nodes, such as using LiteLLM [4] and Ollama [5]. We now have the opportunity to implement a *continuum of intelligence* utilizing intelligence capabilities from GenAI/LLMs and AI Agent services, deployed in the edge and the cloud, and from human (expert) operators via easy-to-use API and human tasks integration.

Clearly, we cannot achieve the time and intelligence continuum without new technology enablers driven by the foundational computing continuum, to support dynamic, seamless integration and invocation of required capabilities. Industry and research results achieved so far in the computing continuum pave the way for us to build-up required support for these two additional types of continuums for many important classes of real-world applications. Supporting time and intelligence continuum will improve mission quality in terms of time-to-resolution, quality of results, and automation for robustness and reliability.

The need for multi-continuum in integrated swarm, edge and cloud capabilities

Let us consider a long survey scenario in a disaster zone (like in a flood in EU [6]). A long running application needs several (a swarm of) Unmanned Autonomous Vehicles (UAVs), edge and cloud storage and computing resources and services, and human operators and analysts. We call such an application a swarm-edge-cloud service-based application (*SEC_App*). Here the swarm part refers to (micro)services deployed in UAV nodes; these services typically work and coordinate their tasks and the movement of UAV nodes as a collective, carrying out part of the mission. The edge and cloud parts provide various data analytics, AI/ML, and mission-related knowledge services working with the swarm part; all together form and support SEC_App for the mission. Using enablers for computing continuum, capabilities from edge-cloud compute nodes and UAV nodes establish a continuum infrastructure from which we can acquire various types of computing resources for the application. We call a set of such computing resources an application-specific resource ensemble (ASR_Ensemble) [7], which includes storage and database, messaging systems, ML/LLM serving service, to name just a few. Utilizing an ASR_Ensemble, the search and rescue (S&R) application consists of multiple service-based components deployed across swarm-edge-cloud computing infrastructures; an ASR_Ensemble can be changed during runtime due to the elasticity and resilience needs of the application and is guaranteed by the computing continuum infrastructure.

Consider, within a SEC_App, an application-level (micro)service being deployed and active within a UAV node. The service can be (self)controlled, due to the energy capabilities of the UAV, to save its states and transfer its unprocessed data and its sensing and analytics tasks to other services (executed in other UAVs or edge/cloud computing nodes) before controlling the underlying UAV hardware to fly back for recharges. Overall, many application states within the application's distributed services must be saved and maintained over time, while one part of the infrastructure of the computing continuum – UAV nodes – can stop working and new UAV nodes may be provisioned for launching replacement or new services to continue the work automatically. With that situation, we avoid the interruption of the infrastructure causing many hard-to-link snapshots/chunks of data and analytic insights that lowers the mission objectives (e.g., a longer time to resolution and a fewer number of problems detected) or requires manual effort from humans to launch new nodes and services. This is one example of the requirements for *time continuum*.

Due to the high complexity of such a mission, tasks carried out by human operators today may be delegated to GenAI/LLMs for interpreting multi-modal data captured from services in UAVs and other devices. Hence, intelligence can be taken from human operators and AI Agents, in a replacement or complementary mechanism, depending on the availability, capability, and design patterns of humans and AI services. When an application needs to strengthen its intelligence, it requires support to build up an intelligence continuum in the sense that AI abilities delivered must be paired with human intelligence to accomplish complex functions without interruption, regardless of the (partially) failure of AI software services or the unavailability of human operators during the execution of the application.

It is our contention that, with the ability to achieve seamless use of computing resources, connectivities, and data (the computing continuum), the time and intelligence continuum can also be supported. In fact, as a developer, we wish to express application needs in terms of time and intelligence capabilities, rather than of the computing resource capabilities. The current state of the art has not developed such a combined view and ways to achieve it. We thus aim to raise awareness of the need for a view on multi-continuum and to analyze the relationships between these different continuum dimensions. We will outline some concrete actions for the implementation of multi-continuum from the service management viewpoint.

Definition of different types of continuum

We first provide some fundamentals by defining the three dimensions of the multi-continuum: computing, time and intelligence.

Computing continuum: reflects the view that applications should be distributable over computing resources, storage, and network, from swarms of IoT-capable nodes to edge and cloud resources. This view has been discussed intensively and several enabling technologies have been introduced [1,2]. The computing continuum takes an infrastructure and connectivity perspective and does not deal with, guarantee or define mission time or intelligence goals related to the applications, although the requests for computing resources may indicate related intelligence capabilities (types of services) and time (SLAs about resource usage).

Time continuum: reflects the continuous execution of swarm-edge-cloud applications without state interruption over the computing continuum. From the application viewpoint, an execution of the application serves a mission. Due to the complexity to achieve the mission goal, technically, the application execution may be split into multiple phases and there may be disruptions among phases that require reconciliation of states, including data, analytics results, insights, etc., to have a complete view of outcomes of the mission (sometimes also seen in post processing). Such disruptions are due to the limitation of the infrastructures and resource allocation, including required intelligence capabilities. For example, a UAV runs out of battery, the data space for the analytics is over the quota due to the allocated infrastructure (such a specific compute node with a powerful LLM deployment), or a human operator is not available to steer the task; all of these will lead to disruptions and data incoherence. Practically, time continuum means the application can be designed for undisrupted

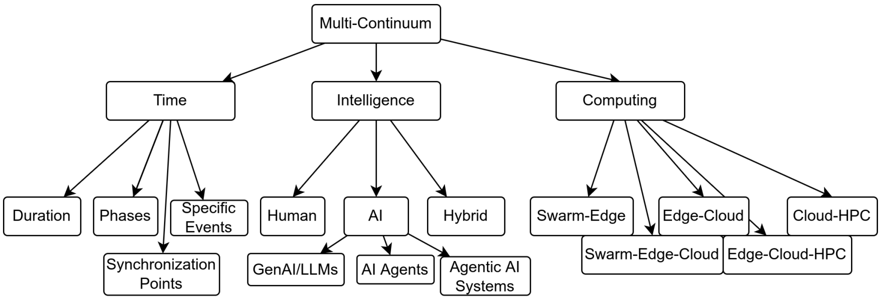

FIGURE 1. Main categories in the Continuum Attribute Space

execution over a long time, and intermediate states and data can be saved and reloaded without changing the design logic of the applications, regardless of changes to the underlying swarm, edge or cloud computing resources. Time continuum is crucial as, without proper support, complex long running applications have to be designed as separate (sub)executions and various types of aggregation and manual processes must be done to reconcile these (sub)executions into a coherent whole mission. This is expensive, such as in data analytics and ML training (with separate phases of data collection, feature engineering, and training), a lot of connectors and information are missing, creating big problems (e.g. changes, merging, and quality of data assurance). We shall note that specifying and guaranteeing time continuum for an application are not the same for a computing resource used by the application (which can be achieved using computing continuum).

Intelligence continuum: Intelligence continuum reflects the ability to switch, combine and utilize different types of intelligence from AI software services and human intelligence without interruption. Intelligence will be stored and maintained over the course of the applications and be provided from edge to cloud to human. Given the computing and time continuums, it is possible to achieve intelligence because the network, edge and cloud allow the seamless connections while AI Agents could act as a bridge to connect various intelligence sources like tools, LLMs, humans and other agents. Furthermore, intelligence abilities can be combined in different compositions, such as sequence workflows or collectives.

While the expectation and engineering of time and intelligence continuum may be implicitly defined when building SEC_App, we propose to make them explicit to support the requirements and implementation. In the past, an explicit specification of cloud elasticity in existing work allowed us to present techniques/tools to control computing resources across edge-cloud in computing continuum. This shall be done also for the time and intelligence continuums, highlighting the importance of all three dimensions, or the *multi-continuum computing*.

Characteristics and Modeling of Multi-Continuum

To support the multi-continuum concept, we propose to characterize computing continuum, time continuum and intelligence continuum via a Continuum Attribute Space (CAS) as shown in Figure 1.

The attribute space is proposed based on key attribute categories associated with entities in each continuum that concern application requirements:

- *Time:* key entities in a mission from a time viewpoint are tasks and resources. Duration, phases, specific events and synchronization points are common notions when carrying tasks in the view of time.

Time continuum allows us to specify requirements for such entities that characterize important points of time in the mission goal. Thus, they allow the continuum provisioning and management techniques to prepare resources for long missions.

- *Intelligence:* main forms of intelligence capabilities are provided by human, AI or hybrids (collaborations between human and AI). Human intelligence can take the form of an individual, team of (subject matter) experts or a crowd, and hybrid intelligence can be provided as a sequence pattern or a collective one of human and AI capabilities.
- *Computing:* main entities are different combinations of computing resources. Thus, main attributes reflect the typical combination of resources, such as swarm-edge, edge-cloud, cloud-HPC, or edge-cloud-HPC. A combined configuration is usually suitable for provisioning resources for a part of the application execution and is carried out in an elastic manner. Attributes can help to select suitable resources that fulfill data, privacy and AI regulations as well.

Interdependencies among different types of continuums

To implement multi-continuum applications, we need to understand dependencies among continuum dimensions. Figure 2 provides a view on dependencies of time and intelligence continuum on swarm-edge-cloud resources and services provided by computing continuum. Time and intelligence continuum requirements are captured by the developer/operator or corresponding provisioning/configuration component/service that will specify the needs, usages and configurations of time and intelligence continuum, which rely on the computing continuum. The process of development and management of such requirements, needs, usages and configurations are continuous through the development and operations of the application.

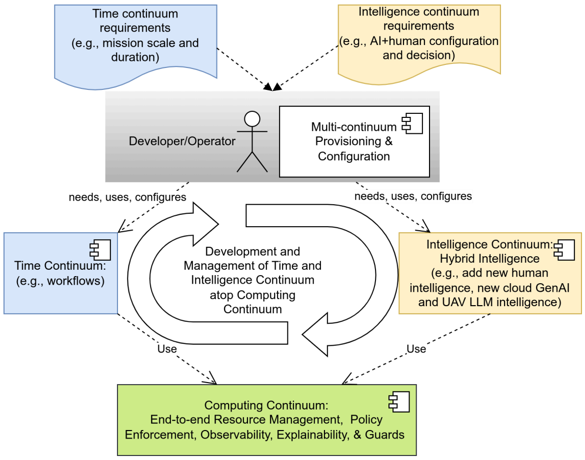

FIGURE 2. Relationships among requirements and specific continuum dimensions.

Given a mission, requirements along the time and intelligence continuum dimensions can be expressed for an application. In terms of time, it is important to support mission operations without a predefined execution duration. In terms of intelligence, it is expected that AI can play the decisive role, making decisions like human operators, besides typical AI capabilities like providing inputs for other tasks and humans. Thus, hybrid intelligence requirements from AI and humans allow the implementation of the intelligence continuum. The requirements are translated into the appropriate configurations, enabling scalability along both dimensions. To allow the required level of scalability in the time dimension, we support sequences of reconfiguration actions during the long lifecycle of the application stages and workflows by saving intermediate state that is accessible anywhere in the computing continuum and by supporting long-running data and analytics workflows. To allow the required scale in the intelligence dimension, our approach identifies the necessary human expertise and augments it with AI/GenAI services in the edge and cloud so as to be able to handle higher operational complexity in long-running missions with fewer human operators. Intelligence scaling can, among others, be interpreted in terms of types (AI, human or hybrid), interactions (as structured workflows or collectives), and performance.

Multi-continuum requirements are associated with SEC_App, whereas abilities to support the requirements are based on ASR_Ensemble and how we select, configure and control services in ASR_Ensemble to deliver required capabilities. Therefore, we need a systematic way to utilize the requirements of SEC_App to meet the multi-continuum requirements. For this, we consider available components or already substantially developed that we can use or extend for multi-continuum and new components for accelerating the development of multi-continuum. The requirements are translated into the appropriate configurations, enabling time and intelligence dimensions over the (existing and new) computing continuum.

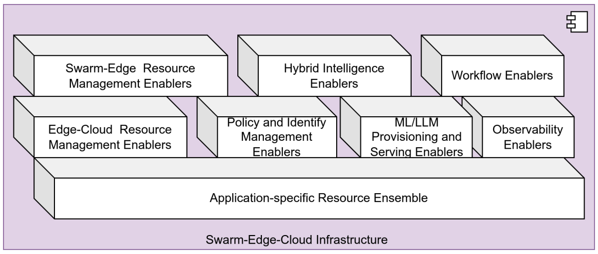

FIGURE 3. Example of key enablers.

Roadmap for Multi-Continuum - a Focus on Coordination Aspects

Since multi-continuum involves a great deal of computing resources and intelligence capabilities across time and space, **coordination** is one of the key aspects. Novel coordination techniques are needed for both time and intelligence continuum. In terms of time continuum, the focus is mainly about coordinating states and resources (such as, saving the state of services in a UVA and planning for a new UAV that takes over the run-out-of-battery UVA). In terms of intelligence continuum, we foresee the connection between ML/AI and human intelligence. We discuss new coordination-centric enablers expecting the enablers to accelerate the development of multi-continuum.

Multi-continuum Computing Resource Coordination

State-of-the -art: Many key enablers for edge-cloud computing resources management have been developed, such as container- and Kubernetes-based edge cloud resource orchestration which allows an application to obtain an elastic, reliable and resilient container ensemble and deploy application microservices, which can also be scaled, saved, and recovered. Existing, serverless models and orchestration tools can help to launch analytics functions in swarm-edge-cloud environments. Adding UAV resources, quite similar to the edge counterpart, creates a lot of challenges that we must support. One aspect is to think of swarm resources as Kubernetes entities to easily deploy services, but due to connectivity configurations [12] it will not work easily. Although common techniques like virtualization and deployments will be needed, the swarm part will interact with the other edge and cloud in a different way. In addition, edge storage and cloud storage essentially enable us to acquire an elastic data store across edge and cloud for storing persistent data, especially for supporting reliable state and data saving, required for time and intelligence continuum. Various reliable messaging systems are also ready for message exchanges across the computing continuum.

Towards multi-continuum support: Resource management will encompass both swarm-edge-cloud and human-assisted activities, covering from planning and design of SEC_App to its operation and maintenance. A workflow can help in the allocation of tasks to services in individual UAVs within the swarm, track their progress, and provide feedback on their performance. One type of workflow is to coordinate services to process swarm events in near real time, controlling SEC_App services to update path plans, and reconfigure the swarm and/or edge-cloud resources. One of the challenging questions is the coordination implementation models. Due to the structure of inter-connected infrastructures in the computing continuum, a reactive choreography using messages could be a suitable choice for both inter- and intra-infrastructures, but would be too complex for the entire computing continuum. Thus, we believe that new research to reexamine existing basic, distributed coordination schemes [16] should be carried out. In terms of engineering, while messaging systems and reactiveness using function-as-a-service or containers are popular in the edge and cloud parts, integrating them for the swarm part is challenging.

End-to-end Policy and Identity Management and Enforcement

State-of-the-art: distributed services provide a virtual edge-cloud secret-as-a-service (vaults) for applications to store and manage digital identities for information security and governance across swarm-edge-cloud. Currently services in the edge can utilize centralized secret-as-a-service or vaults hosted in the cloud but an extension to services in UAVs must be made/strengthened. Similarly, policy management for the edge subsystem or the cloud subsystem have been supported widely. With a swarm, it is also possible to dynamically control policies. But we need to strengthen these supports in the cross swarm-edge-cloud computing continuum.

Towards multi-continuum support: policy languages and policy enforcement techniques for multi-continuum must be investigated. A question of whether a single language, such as Rego [13], is suitable or not, or what needs to be done to have a holistic policy language. To enforce the policy, obviously we see a complex model of coordination combinations, such as centralized and decentralized and complex pluggable enforcements atop existing computing continuum. For example, edge-cloud computing resources can be enforced with Kubernetes or underlying mechanisms, whereas swarms can be based on ROS2 [14]. One open issue is how to integrate and enforce policies and manage identities for humans and AI services when scaling intelligence capabilities. In addition, to enable intelligence continuum in dealing with policies, GenAI and AI Agents could be investigated for implementing policy enforcement in the context of security orchestration and incident responses.

Observability, Explainability and Risk Guards

State-of-the-art: with plenty of observability tools and systems available and due to the heterogeneity of continuum, it is hard to think that a single, even powerful, observability platform can be used. For example, Prometheus [4] and OpenTelemetry [5] may be very powerful but they are not quite suited or explored for observability of services deployed in UAVs nodes. Furthermore, services in SEC_App usually also have their own application monitoring to deal with states of the application that shall be reported for multi-continuum analysis.

Towards multi-continuum support: we will oversee all service actions and interactions within and among services and workflows that will facilitate human-AI interactions to support AI risks, explainability and responsible AI. On the one hand, both SEC_App and ASR_Ensemble behaviours must be observed using a combination of tracing and monitoring tools. On the other hand, actions within SEC_App, including AI inferences, must be explained and guarded against potential risks. It is tricky as many aspects of SEC_App from an application-specific view must be captured, requiring a strong integration with the application development. The explainability requires protocols and specification to specify what needs to be explained both runtime and offline explainability. The explainability need is far beyond common AI model explainability [17] that must be considered for interactions among various types of services spanning over the time. Another feature is that it will continuously collect feedback from services, aggregating the feedback and putting it into context to explain decisions, service composition, AI-human composition, mission status and possible failures in a human-understandable way. Thus, it helps the human operators to be aware of the system's current and future state to also intervene in interactions related to the intelligence continuum.

Hybrid Intelligence Coordination

State-of-the-art: Concepts of hybrid intelligence and collective intelligences via human-service and multi-agent have been well-developed for many years [15]. However, considering them in the context of edge-cloud computing with contemporary ML/LLMs, they are still not available for SEC_App. In summary we can see that existing ML/LLMs provisioning and serving in the edge and cloud which are static deployed can be used.

- Traditional ML services can be provisioned in different computing resources.
- LLMs in edge nodes are starting to be popular. With a combination of cloud-based GenAI, we have multiple ML/LLMs capabilities that can be invoked via service APIs.
- Distributed ML serving allows an application to deploy and distribute its ML models to different serving nodes in UAV, edge and cloud resources, allowing complex end-to-end and composable ML pipelines.
- ML/LLM orchestration can automate the management of AI workflows, deciding where to run an ML model based on a set of criteria or rules.

Towards multi-continuum support: The research focus for hybrid intelligence includes *engineering* techniques and tools to establish hybrid intelligence service-based software with humans and AI capabilities. Such hybrid intelligence is delivered via hybrid human-agent systems to serve for different purposes: analytics and decision making. One focused direction is to develop suitable patterns and adaptation techniques for hybrid intelligence. For example, common ML and AI Agents can be used to interpret data provided by SEC_App to address a typical need to properly schedule or react to different types of findings from SEC_App. Task planning at an abstract level helps generate quickly feasible actions for the swarm as a whole to perform tasks, without processing the technical details of how each action can be implemented by the individual AI or humans.

Sidebar: Further related work

Many complex applications require dedicated computing resources for tasks execution; such resources conceptually establish a form of ASR_Ensembles from distributed infrastructures. A simple example is the case of Flink application clusters as dedicated computing resources (edge or cloud) for a specific application [1]. Although work still has been evolving, especially for provisioning UAVs as computing resources for swarm-edge-cloud continuum, existing UAVs as a service [2] and virtualization techniques will address a seamless swarm-edge-cloud computing continuum. While integrating swarm, edge and cloud for computing continuum has been in the research focus, achieving scalability in the time and intelligence continuum (resilient and elastic systems for highly available, durable missions with AI-enabled capabilities) is still unexplored. Several tools have been developed to enable elasticity within a single space, either at the cloud or at the edge, and at a single layer, such as virtual containers at the infrastructures. But, there is a lack of tools for managing resilience and elasticity across swarm-edge-cloud infrastructures to support time continuum.

Control via elasticity modeling and other policy specifications has been explored but currently does not include specifications for multi-continuum requirements in terms of time and intelligence. Elasticity controls for single attributes have been developed as well as optimization techniques targeting individual layers (infrastructure, platform, and application level in infrastructural computing continuum). For instance, Baresi et. all addressed management of applications deployed in the edge-cloud continuum [3], but such approaches need to be significantly extended to achieve end-to-end optimization for resilience and elasticity of swarm-edge-cloud applications. Additionally, none of these prior works considers human-software elasticity for achieving time and intelligence continuums.

[1] Flink Application Cluster, https://nightlies.apache.org/flink/flink-docs-master/docs/concepts/flink-architecture/
[2] B. Alkouz, A. Bouguettaya, and S. Mistry, "Swarm-based drone-as-a-service (sdaas) for

delivery," in 2020 IEEE International Conference on Web Services (ICWS), pp. 441–448, 2020.
[3] L. Baresi, D. F. Mendonça, M. Garriga, S. Guinea, and G. Quattrocchi. 2019. A Unified Model for the Mobile-Edge-Cloud Continuum. ACM Trans. Internet Technol. 19, 2, Article 29 (May 2019), 21 pages. https://doi.org/10.1145/3226644

Conclusion

We advocate a new research focus on applications that operate on the combined capabilities from the computing continuum, time continuum and intelligence continuum, the concepts we collectively refer to as multi-continuum in this work. We provide definitions and key aspects for enabling this vision. We discuss the need to strengthen current enablers to support coordination of multi-continuum capabilities for swarm-edge-cloud service-based applications.

Acknowledgements

We thank our colleagues in the HyperScaleOps proposal for fruitful discussions that helped us cement the views of the multi-continuum concept.

References

1. M. Jansen, A. Al-Dulaimy, A. V. Papadopoulos, A. Trivedi and A. Iosup, "The SPEC-RG Reference Architecture for The Compute Continuum," in 2023 IEEE/ACM 23rd International Symposium on Cluster, Cloud and Internet Computing (CCGrid), Bangalore, India, 2023, pp. 469-484
2. M. Tortonesi, "The Compute Continuum: Trends and Challenges," in Computer, vol. 58, no. 3, pp. 105-108, March 2025
3. Sapkota, R., Roumeliotis, K. I., and Karkee, M., "AI Agents vs. Agentic AI: A Conceptual Taxonomy, Applications and Challenges", arXiv e-prints, Art. no. arXiv:2505.10468, 2025.
4. LiteLLM. https://github.com/BerriAI/litellm, Last access: May 28, 2025
5. *Ollama Framework*. https://www.ollama.com/. Last access: Mar 9, 2025, from
6. https://www.theguardian.com/commentisfree/2024/oct/02/floods-havoc-vienna-austria-climate-crisis
7. H. -L. Truong, "ASRE — Towards Application-specific Resource Ensembles across Edges and Clouds," *2021 17th International Conference on Network and Service Management (CNSM)*, Izmir, Turkey, 2021, pp. 239-243
8. Calavaro, C., Cardellini, V., Lo Presti, F. *et al.* Beyond Cloud: Serverless Functions in the Compute Continuum. *SN COMPUT. SCI.* 6, 194 (2025).
9. *OpenTelemetry*. Retrieved June 23, 2025, from https://opentelemetry.io/
10. *Prometheus: Monitoring system & time series database*. Retrieved June 23, 2025, from https://prometheus.io/
11. Chalima Dimitra Nassar Kyriakidou, Athanasia Maria Papathanasiou, Vasilios A. Siris, Nikos Fotiou, George C. Polyzos, Eduardo Cánovas Martínez, and Antonio Skarmeta. 2025. Identity and Access Management for the Computing Continuum. In Proceedings of the 2nd International Workshop on MetaOS for the Cloud-Edge-IoT Continuum (MECC '25). Association for Computing Machinery, New York, NY, USA, 33–39.
12. Pietro Boccadoro, Domenico Striccoli, Luigi Alfredo Grieco, An extensive survey on the Internet of Drones, Ad Hoc Networks, Volume 122, 2021, 102600, ISSN 1570-8705
13. OPA, https://www.openpolicyagent.org/docs/policy-language, Retrieved June 23, 2025.
14. S. Macenski, T. Foote, B. Gerkey, C. Lalancette, W. Woodall, "Robot Operating System 2: Design, architecture, and uses in the wild," Science Robotics vol. 7, May 2022.
15. O. Scekic et al., "A Programming Model for Hybrid Collaborative Adaptive Systems," in IEEE Transactions on Emerging Topics in Computing, vol. 8, no. 1, pp. 6-19, 1 Jan.-March 2020
16. Gabriele Mencagli. 2016. A Game-Theoretic Approach for Elastic Distributed Data Stream Processing. ACM Trans. Auton. Adapt. Syst. 11, 2, Article 13 (July 2016)
17. Joao Marques-Silva and Xuanxiang Huang. 2024. Explainability Is Not a Game. Commun. ACM 67, 7 (July 2024), 66–75.

Hong-Linh Truong is an associate professor at Aalto University, Finland. Contact him at linh.truong@aalto.fi

Kostas Magoutis is a professor at University of Crete and FORTH-ICS, Greece. Contact him at magoutis@csd.uoc.gr.