

Towards a Hybrid Hierarchical Digital Twin Architecture for the 6G Compute Continuum

José Santos*, Javad Sameri*[‡], Sam Van Damme*, Susanna Schwarzmann[†], Qing Wei[†],
Riccardo Trivisonno[†], Maria Torres Vega[‡], Filip De Turck*,

* IDLab, Department of Information Technology, Ghent University - imec, Ghent, Belgium

[†] Huawei Technologies, Germany

[‡] eMedia Research Lab, Department of Electrical Engineering (ESAT), KU Leuven, Belgium

Email: josepedro.pereiradossantos@UGent.be

Abstract—The emergence of the 6G era demands seamless orchestration across an increasingly heterogeneous and distributed compute continuum—spanning edge, fog, and cloud resources. Moreover, the next generation of the mobile network is poised to redefine the digital landscape by enabling pervasive intelligence, ultra-low latency communication, and extreme heterogeneity across the entire network infrastructure. This transformation introduces unprecedented orchestration challenges due to the dynamic, multi-domain, and resource-constrained nature of emerging workloads such as Generative Artificial Intelligence (GenAI) inference, immersive eXtended Reality (XR), and autonomous systems. To tackle this complexity, we advocate for a Hybrid Hierarchical Digital Twin (DT) architecture that serves as a foundation for intelligent, adaptive, and real-time orchestration in 6G environments. We present a comprehensive vision for integrating DTs as enablers of intelligent, context-aware, and adaptive orchestration mechanisms that span across multiple domains. The proposed architecture introduces a multi-layered DT hierarchy combining local and global views, enabling scalable coordination and real-time decision-making. We highlight key architectural enhancements required to realize this vision, including inter-twin interoperability and behavioral modeling for QoE estimation. This work aims to guide researchers and practitioners in shaping the foundations of resilient and efficient orchestration frameworks for 6G systems.

Index Terms—Digital Twin, Reinforcement Learning, Orchestration, Cloud Computing, Compute Continuum

I. INTRODUCTION

The advent of 6G networks is promising transformative capabilities that will reshape the digital landscape [1]. Unlike previous generations, 6G is expected to deliver pervasive intelligence, ultra-low latency, and seamless integration of heterogeneous resources across a widely distributed compute continuum, including multiple edge, fog, and cloud layers. However, the increased complexity of emerging applications such as Generative Artificial Intelligence (GenAI) inference, immersive eXtended Reality (XR), and the diverse nature of the underlying infrastructure pose significant orchestration challenges. The Compute Continuum (CC) in 6G is characterized by a dynamic, multi-domain, and resource-constrained environment where efficient management of resources and workloads is critical to meet stringent performance requirements. Traditional orchestration methods, often centralized, are inadequate to cope with the scale, heterogeneity, and real-time demands of next-generation mobile networks.

These orchestration challenges highlight the potential of integrating Digital Twins (DTs) [2] as complementary, simulation-driven platforms for training and refining orchestration policies. DTs provide dynamic, real-time digital replicas of the network and its users, accurately mirroring and forecasting of physical system behavior. By acting as virtual testbeds, DTs facilitate safe exploration and evaluation of network adaptations prior to real-world deployment. In our vision, DTs extend beyond modeling traditional infrastructure-level metrics, such as Quality of Service (QoS), to also capture behavioral patterns and perceptual feedback, i.e., Quality of Experience (QoE). This comprehensive modeling empowers Reinforcement Learning (RL) agents to train and operate within simulated environments that mimic actual network conditions. However, realizing this vision demands addressing critical challenges, including achieving high-fidelity, real-time synchronization between the DT and the physical 6G system; preserving model accuracy that encompasses both technical parameters and human-centric factors; and ensuring seamless integration of DT-generated insights with the orchestration layer to deliver timely, actionable feedback to RL agents.

This paper presents our envisioned hybrid hierarchical orchestration framework for 6G networks, addressing key challenges in multi-tier decision-making and the integration of DTs as foundational elements for context-aware and adaptive management across the CC. By combining localized telemetry with aggregated regional and global system views, our approach aims to enable a more intelligent, resilient, and scalable orchestration paradigm that spans from edge resources to centralized cloud infrastructures. Building further on previous works, including subjective studies and hybrid QoE models [3], [4], and RL-based orchestration methods [5], [6], our approach leverages simulation-driven optimizations to bridge the gap between network-level Key Performance Indicators (KPIs) and actual user-perceived experience in dynamic and heterogeneous environments. In particular, this paper addresses critical limitations in RL-driven orchestration, such as sample inefficiency, and limited policy generalization, by introducing DTs as high-fidelity virtual replicas of the 6G system. These enable safe, accelerated, and context-rich training and testing of orchestration policies. At the same time, the use of DTs introduces its own architectural and operational challenges,

including real-time synchronization, modeling of user behavior for accurate QoE estimation, and tight integration with the orchestration layer. By addressing these challenges, our orchestration framework supports proactive, intelligent management strategies capable of adapting in real-time to user and network dynamics—laying the groundwork for scalable, and user-aware 6G services.

The remainder of this paper is organized as follows: Sec. II briefly discusses the background and related works on RL and DT. Sec. III introduces an overview of the proposed architecture, followed by an analysis of different orchestration approaches in Sec. IV. Sec. V outlines the architectural enhancements in the proposed hybrid hierarchical DT architecture within the 6G CC. Lastly, Sec. VI concludes this paper.

II. BACKGROUND & RELATED WORK

Cloud-Native (CN) eXtended Reality (XR) [7] represents the evolution of immersive technologies through their integration with scalable and distributed cloud infrastructures. Unlike traditional XR systems that rely heavily on co-located, high-performance hardware, CN XR leverages the elasticity and ubiquity of cloud services to deliver immersive content seamlessly across heterogeneous devices and networks. Offloading compute-intensive operations—such as rendering, encoding, and scene reconstruction—to remote cloud or edge resources allows for lightweight client devices while maintaining high-fidelity experiences. This shift not only enhances scalability and accessibility but also aligns XR service delivery with emerging CC paradigms [7].

Kubernetes (K8s) [8] has become a cornerstone technology in the orchestration of containerized applications, including latency-sensitive workloads such as CN XR. As a de facto standard in cloud-native environments, K8s offers robust primitives for automated deployment, scaling, and resource management across distributed infrastructures. For XR workloads, K8s enables dynamic allocation of compute and memory resources, supports real-time scaling based on user activity, and facilitates seamless application updates—crucial for maintaining low-latency, high-quality experiences under variable network and compute conditions.

The Compute Continuum (CC) [9] represents a new computing paradigm where resources span a unified, hierarchical infrastructure from the far edge to centralized clouds. Rather than treating edge and cloud as separate silos, the CC emphasizes continuity and fluidity in resource allocation and data flow. This paradigm is especially relevant in the context of 6G networks, which are expected to support demanding applications such as collaborative XR, real-time DTs, and distributed GenAI. By enabling workload placement and migration across the continuum in response to latency constraints, energy considerations, and user mobility, the CC provides the infrastructure backbone needed for future-ready, adaptive XR systems.

Reinforcement Learning (RL) [10] is a powerful Machine Learning (ML)-based technique that enables agents to learn optimal behaviors through trial and error, guided by rewards

and penalties from their environment. Unlike traditional supervised learning, where a model is trained on a fixed dataset, RL involves an agent interacting with its environment to discover the best strategies for achieving specific goals. Thus, RL is well-suited for dynamic and complex decision-making tasks where actions must be determined in near real-time. In the context of 6G, RL can be applied to optimize QoE for end users by dynamically adjusting network configurations and allocating computational resources based on real-time information, ensuring that performance metrics such as latency, bandwidth, and user satisfaction are continuously optimized, steering orchestration decisions toward an enhanced user experience. In addition, the versatility of RL methods to handle multi-objective optimization problems makes it ideal for scenarios where multiple, often conflicting, goals must be balanced, such as maximizing user experience while minimizing energy consumption of the infrastructure.

Digital Twin (DT) [2] refers to a digital replica of a physical system, continuously synchronized through data exchange to reflect the real-time state, behavior, and evolution of the system it represents. In contrast to static models, DTs provide a dynamic and context-aware abstraction of network infrastructure, user behavior, and service demands. This real-time mirroring capability allows for advanced monitoring, simulation, and predictive control, making DTs a cornerstone technology for enabling intelligent orchestration in 6G. By integrating DTs with AI-driven decision-making mechanisms such as RL, it becomes possible to simulate and evaluate various actions before deployment, reducing operational risks and enabling proactive, closed-loop management of network resources. Ultimately, DTs enhance the adaptability, efficiency, and resilience of 6G orchestration frameworks by bridging the gap between physical and virtual network environments.

Recent works have proposed DT architectures for 5G and beyond networks [11]–[14], and the CC [15]–[17]. These approaches demonstrate the potential of DT to model, monitor, and optimize complex networked environments by bridging physical and digital layers. However, existing solutions are typically tailored to either centralized control paradigms or domain-specific optimizations, limiting their scalability and adaptability in dynamic, multi-tier infrastructures. Moreover, current DT frameworks often assume static or cloud-centric deployments, which fail to meet the performance and resilience requirements of emerging 6G workloads such as immersive XR. These applications demand context-aware orchestration that spans across edge and cloud boundaries, where decisions must be made rapidly and locally, while still aligned with global optimization goals.

In contrast, this paper proposes a hybrid hierarchical DT architecture that integrates localized, regional, and global DTs into a cohesive orchestration fabric. This approach enables fine-grained responsiveness at the edge, situational awareness at the regional level, and policy coherence at the global tier. By distributing intelligence and decision-making capabilities across the CC, our architecture supports predictive, adaptive, and scalable orchestration mechanisms

suit for the dynamic conditions and heterogeneous resources of 6G environments.

III. SYSTEM OVERVIEW

This section describes a modular architecture for real-time QoE-driven orchestration in 6G CC environments. The proposed system integrates QoE-aware Network Functions (NFs), AI-based orchestration policies, and runtime feedback loops powered by DTs and RL. This work aims to dynamically adapt the deployment and operation of latency-sensitive XR services based on both real-time user experience metrics and network and infrastructure information.

Quality of Experience (QoE) Network Function (NF).

To bridge the gap between infrastructure-level metrics and end-user satisfaction, we propose a dedicated QoE NF deployed at the network edge. This NF ingests application-level telemetry, such as video resolution, frame rate, packet loss, and latency, alongside contextual data including user location and device capabilities. Leveraging lightweight ML models, such as decision trees or compact Neural Networks (NNs), the NF provides real-time QoE score estimations that reflect each user's subjective experience. Our envisioned QoE model integrates both Technological Influencing Factors (TIF) and Human Influencing Factors (HIF): conventional network KPIs, such as throughput, packet size, and inter-arrival times, are enriched with physiological signals, including heart rate and skin conductivity monitored via wearable sensors. By applying modeling techniques such as Support Vector Classifier (SVC) or XGBoost (XGB) on these features, the NF generates accurate, near real-time QoE predictions tailored to various dimensions of user experience, including perception of jerkiness, latency, and synchronization [4]. These QoE scores are fed to the RL orchestrator to inform and guide system-level decision-making. Unlike traditional QoS metrics, this user-centric abstraction enables the system to prioritize perceptual outcomes, resulting in more intelligent and adaptive orchestration strategies. Moreover, the QoE NF opens the door to novel Service Level Agreement (SLA) definitions and pricing models grounded in user satisfaction rather than rigid performance guarantees.

Digital Twin (DT). A key innovation of our approach is the integration of a QoE-centric DT that mirrors the current and predicted future states of the network, compute infrastructure, and service topology. The DT ingests real-time metrics and orchestration actions, and supports the simulation of hypothetical scenarios, allowing the orchestrator to assess the impact of decisions before enactment. This predictive modeling capability is particularly valuable in the presence of conflicting objectives, such as minimizing latency while conserving energy, and helps the system anticipate performance degradation before it occurs. In contrast to traditional methods, which focus on physical or network behavior alone, our DT maintains a model of QoE dynamics and supports multi-objective optimization at runtime.

Reinforcement Learning (RL) Orchestrator. At the core of our architecture is an intelligent orchestrator powered by

RL. This agent observes system states from both the QoE NF and the DT and learns to take optimal actions such as migrating services, adjusting resource limits, or rerouting flows based on a multi-objective reward function. The reward function combines real-time QoE metrics with energy consumption and service-level compliance indicators. Importantly, the RL agent does not require prior knowledge of network or service models. Instead, it learns optimal behavior through interaction and feedback. This allows the system to adapt to novel workloads and operating conditions, including unforeseen user mobility or traffic spikes.

IV. ANALYSIS: MAIN DIFFERENCES BETWEEN POTENTIAL ORCHESTRATION STRATEGIES

In the 6G CC, the delivery of XR immersive services demands stringent QoE guarantees. To meet these requirements, the orchestration layer must be capable of dynamically managing heterogeneous and geographically distributed resources under strict latency, reliability, and bandwidth constraints [18]. Therefore, the choice of orchestration strategy—whether **Centralized**, **Distributed**, or **Hybrid**—is a critical design consideration. Each approach presents distinct trade-offs in terms of latency (i.e., decision-making response time), computational overhead, scalability across network domains, resilience to failures, and overall deployment complexity. This paper investigates these dimensions to identify which orchestration paradigm is best suited to support real-time, QoE-sensitive applications in next-generation mobile networks, highlighting how architectural choices directly impact service quality, operational cost, and system robustness. Table I provides a qualitative comparison of centralized, distributed, and hybrid orchestration strategies, considering key architectural and operational dimensions relevant to 6G networks and QoE-aware service provisioning.

Centralized Orchestration (CO) offers the advantage of strong policy consistency due to its global view and control of the system. All decisions are made at a single control point, which ensures uniformity in resource allocation and policy enforcement. However, this comes at the cost of lower responsiveness, as control messages must traverse the network, resulting in higher latency and potential bottlenecks. Moreover, centralized systems face limitations in scalability; as the number of nodes or services increases, the central controller may become overwhelmed, leading to degraded performance and increased failures. In such systems, fault tolerance is inherently low because a failure in the central controller can impact the entire network. In contrast, **Distributed Orchestration (DO)** exhibits high responsiveness and scalability. By decentralizing control, decisions can be made closer to where services are consumed, resulting in lower latency and more adaptive resource management. This model is also more fault-tolerant, as failures are typically contained within localized domains. Furthermore, distributed strategies enhance data locality and improve security and privacy by keeping sensitive data at the edge. However, these benefits are offset by higher coordination overhead and increased deployment complexity, as consistent

global behavior must emerge from interactions among many autonomous controllers.

Lastly, **Hybrid Orchestration (HO)** attempts to combine the best of both. It leverages centralized coordination for global objectives and policy dissemination, while allowing local controllers to make fast, context-aware decisions. This leads to a medium-to-high score across most categories. Policy consistency is maintained within regions, responsiveness remains high due to local decision-making, and the architecture is generally resilient to failures. Nonetheless, hybrid models still inherit some deployment complexity and coordination challenges from distributed systems, albeit to a lesser extent.

Table II provides a quantitative comparison of centralized, distributed, and hybrid orchestration strategies across several key metrics relevant to 6G service provisioning. These values are derived from prior studies [4], [7], and system design considerations. *Decision latency* is highest in centralized architectures (80–150 ms), due to the round-trip delays between endpoints and the centralized controller. DO significantly reduces this latency to 10–30 ms by enabling localized decision-making. Hybrid solutions offer a trade-off, with latencies in the range of 20–60 ms, benefiting from hierarchical or regional control layers. The *model convergence time*, indicative of how long learning-based components (e.g., for QoE estimation or resource prediction) take to stabilize, is longest in centralized systems (2–3 hours). DO allows much faster convergence (30–45 minutes) by leveraging parallelized and localized training, while hybrid systems fall in between (45–90 minutes). *Control plane traffic* is notably high in centralized architectures (200 Mbps), as all nodes must communicate frequently with the central controller. Distributed systems limit signaling to local domains, reducing traffic to around 50 Mbps. HO incurs moderate signaling overhead, between 80 and 120 Mbps, by balancing global awareness with local autonomy.

When considering *instantiation time*—the time to deploy network slices or services—centralized systems are the slowest (500–1000 ms), while distributed models achieve much faster deployment times (100–200 ms). Hybrid approaches again offer a compromise, requiring 200–400 ms on average. In terms of *failure domain size*, centralized models are the most vulnerable; a controller failure can affect 100% of the system. Distributed systems are more resilient, typically limiting the impact to just 5–10% of the infrastructure. Hybrid approaches localize failures regionally, with affected areas generally ranging from 5–25%. Finally, concerning *scalability*, DO supports the largest scale, managing more than 1000 domains effectively. Centralized approaches are constrained by bottlenecks and can handle approximately 200 domains. Hybrid architectures can support between 500 and 1000 domains, depending on the coordination overhead and design optimizations employed.

In summary, each orchestration strategy entails specific benefits and drawbacks, particularly when evaluated against performance, scalability, and system complexity (Table III). Based on the comparative analysis of orchestration strategies, we argue for the adoption of a **Hybrid Hierarchical DT**

TABLE I: Qualitative Comparison of Orchestration Strategies.

Aspect	Centralized	Distributed	Hybrid
Policy Consistency	High	Low	Medium
Responsiveness	Low	High	High
Scalability	Medium	High	High
Fault Tolerance	Low	High	High
Data Locality	Low	High	High
Security/Privacy	High	Low	Medium
Coordination Overhead	Low	High	Medium
Deployment Complexity	Low	High	Medium–High

TABLE II: Quantitative comparison of Orchestration Approaches.

Metric	Centralized	Distributed	Hybrid
Decision Latency	80–150 ms	10–30 ms	20–60 ms
Convergence Time (Model)	2–3 hrs	30–45 min	45–90 min
Control Plane Traffic	200 Mbps	50 Mbps	80–120 Mbps
Instantiation Time	500–1000 ms	100–200 ms	200–400 ms
Failure Domain Size	100%	5–10%	5–25% (regional)
Scalability	~200 domains	>1000 domains	500–1000 domains


Architecture as the foundational model for 6G CC orchestration. This architectural paradigm merges the strengths of both centralized and distributed approaches while mitigating their inherent drawbacks. As shown in our evaluation, CO offers high policy consistency and efficient global optimization but suffers from high decision latency, longer model convergence time, and limited scalability. In contrast, DO approaches are responsive, scalable, and privacy-preserving, yet incur significant coordination overhead and risk policy fragmentation. The HO approach strikes an effective middle ground: it enables regional autonomy for rapid, localized decisions, while also maintaining global coherence through hierarchical coordination and DT representations of the network state. The envisioned architecture introduces a layered control model. At the lower tier, *regional edge domains* are empowered with localized inference and adaptation capabilities, enabling real-time decision-making based on local observations. At the higher tier, a *global controller* equipped with a DT monitors macro-level trends and enforces long-term optimization goals.

V. TOWARDS A HYBRID HIERARCHICAL DIGITAL TWIN ARCHITECTURE FOR 6G

To address the trade-offs identified in our analysis, we propose a **Hybrid Hierarchical DT Architecture** tailored for 6G CC orchestration to support immersive XR services. This architecture integrates multi-level control with DT-based network state representations to optimize resource allocation, QoE management, and resilience while maintaining scalability and low-latency responsiveness. Fig. 1 illustrates the proposed hybrid hierarchical DT architecture consisting of three primary tiers: the **Global Tier**, the **Regional Tier**, and the **Local Tier**, each responsible for distinct levels of control and decision-making. At the top of the hierarchy, the **Global RL Orchestrator** acts as the central coordination entity overseeing macro-level network management and long-term optimization goals. It interacts closely with the **Global DT**, a virtual representation maintaining a comprehensive and up-to-date model of the entire network state. This Global DT enables policy dissemination, system health monitoring, and strategic planning decisions.

TABLE III: The advantages and disadvantages of each orchestration approach.

Approach	Advantages	Disadvantages
Centralized	Global data leads to accurate and generalized model training. Simplified policy enforcement due to unified control logic.	High decision latency (80–150 ms) Slice instantiation time (500–1000 ms) Limited suitability for real-time services. Large failure domain (up to 100%) poses resilience concerns.
Distributed	Low decision latency (10–30 ms) Fast slice instantiation (100–200 ms). High scalability (>1000 domains).	Increased coordination and synchronization overhead. Potential inconsistencies in policy enforcement. Requires more complex orchestration logic and state management.
Hybrid	Balances local responsiveness with global/regional optimization. Convergence time and control overhead are moderate. Scalability supports 500–1000 domains (large deployments).	Increased architectural complexity due to multi-tier coordination. Sophisticated mechanisms for coordination and policy alignment. Deployment is challenging, demanding fine-tuned control hierarchies.

**Fig. 1:** An overview of the Hybrid Hierarchical DT Architecture. A Global RL orchestrator coordinates multiple Regional RL Orchestrators and Regional DTs, managing local resources and XR user devices.

Beneath the global tier, multiple **Regional RL Orchestrators** operate within geographically distributed edge domains. Each is paired with a **Regional DT** modeling the localized state and behavior of its domain. These Regional DTs provide real-time insights, enabling low-latency, context-aware decision-making and rapid adaptation to dynamic edge conditions. The regional orchestrators communicate bidirectionally with their DTs to synchronize network state and control decisions. They also maintain hierarchical communication channels with the Global Orchestrator and Local DTs in their region, ensuring alignment between regional autonomy and global objectives through coordination messages and policy updates. At the edge, the architecture interfaces directly with XR devices and end users. These devices interact primarily with their local orchestrator, which manages resources and service delivery with minimal latency, meeting the strict QoE demands of immersive XR applications. The architecture is composed of three main control tiers:

- **Local Tier:** This tier operates at the network edge close to end devices, such as base stations, access points, or on-device compute nodes, and serves as the first line of orchestration. Each local zone hosts a lightweight *local orchestrator* that handles immediate, latency-sensitive

tasks, leveraging minimal compute and storage. These orchestrators interact directly with small-scale DTs.

- **Regional Domains (Intermediate Tier):** These domains manage geographically localized clusters of compute, storage, and network resources near end users. Each regional domain contains a *Regional Orchestrator* empowered with embedded AI inference engines for fast, context-aware decisions. The regional orchestrator relies on near real-time telemetry and local DTs mirroring physical resources and active network slices within the domain. This enables rapid adaptation to dynamic conditions such as user mobility, workload fluctuations, and transient faults.
- **Global Controller (Central Tier):** At the top level, a *Global Orchestrator* maintains a high-level DT synthesizing data from Regional DTs and other global metrics. This controller focuses on long-term optimization goals like load balancing across domains, enforcing cross-domain policies, and coordinating AI model training to improve global decision accuracy. By aggregating multi-domain telemetry, the global controller enables efficient resource provisioning and policy consistency at scale.

DTs play a pivotal role in the envisioned hybrid architecture

by providing virtual representations of the state of the network at both local and global levels:

- **State Synchronization:** Continuous synchronization between physical resources and their virtual counterparts enables accurate monitoring and prediction of resource availability, network congestion, and QoE metrics.
- **Predictive Analytics:** DTs leverage historical data and ML to forecast future system states, enabling proactive resource orchestration, congestion avoidance, and fault anticipation.
- **Model Training and Adaptation:** Regional DTs perform localized model updates using domain-specific data to enhance inference accuracy and reduce convergence times. Periodically, the global DT aggregates these updates to retrain global models that guide long-term policy refinement.

The main benefits of the envisioned approach are the following. First, it **reduces decision latency**, as local orchestrators empowered with DTs can minimize latency by confining time-critical decisions within regional boundaries, achieving sub-50 ms responsiveness suitable for XR QoE demands. Second, it **enhances scalability and resilience** by hierarchically partitioning the orchestration logic, allowing the system to scale beyond 1000 domains without overwhelming a single controller. This approach also reduces the blast radius of faults, enabling graceful degradation by limiting failure domains to regional scopes. Third, it **balances coordination overhead** through hybrid control, where local DTs report summarized states to the global DT, thus reducing control plane traffic compared to flat distributed systems while preserving enough global awareness to prevent policy inconsistencies. Finally, it **improves policy consistency** by allowing global orchestration to enforce high-level policies across regions, while local orchestrators adapt these policies to domain-specific contexts. The DT framework ensures synchronized state views, mitigating the policy fragmentation risks typically found in purely distributed designs.

In summary, the proposed hybrid architecture synthesizes the advantages of CO and DO paradigms to meet the stringent QoE and scalability requirements of 6G CC. It offers a flexible, resilient framework for the delivery of next-generation XR services, enabling near real-time local decisions, supported by coordinated global oversight.

VI. CONCLUSIONS

As 6G networks continue to evolve towards pervasive intelligence and extreme heterogeneity, the need for scalable, context-aware, and adaptive orchestration frameworks becomes increasingly critical. In this paper, we proposed a hybrid hierarchical DT architecture as a foundational pillar for managing the complexity of the 6G CC. Our framework integrates local, regional, and global orchestration layers, each empowered by high-fidelity DTs and AI-driven decision-making, to enable real-time, resilient, and intelligent resource coordination across edge, fog, and cloud domains. By leveraging DTs not only as passive monitors but as active enablers

of safe and accelerated learning, the proposed architecture addresses key limitations of current orchestration methods, such as poor generalization, sample inefficiency, and limited context awareness. We outlined essential architectural enhancements and discussed open research challenges, including inter-twin interoperability, real-time synchronization, behavioral modeling for QoE estimation, and cross-tier policy alignment. This work aims to lay a foundational step toward proactive, sustainable, and user-centric orchestration in next-generation 6G systems, ultimately empowering future services such as immersive XR, and large-scale GenAI applications.

ACKNOWLEDGMENT

José Santos is funded by the Research Foundation Flanders (FWO), grant nrs. 1299323N and 1253226N. This work is partly funded by FWO WaveVR project, grant nr. G034322N.

REFERENCES

- [1] L. F. Bittencourt *et al.*, “The computing continuum: Past, present, and future,” *Computer Science Review*, vol. 58, p. 100782, 2025.
- [2] L. U. Khan *et al.*, “Digital-twin-enabled 6g: Vision, architectural trends, and future directions,” *IEEE Communications Magazine*, vol. 60, no. 1, pp. 74–80, 2022.
- [3] J. Sameri *et al.*, “Collaborative cooking in vr: Effects of network distortion in multi-user virtual environments,” in *Proceedings of the 15th ACM Multimedia Systems Conference*, 2024, pp. 509–515.
- [4] J. Santos *et al.*, “Leveraging user perception for 6g edge-cloud orchestration of networked extended reality,” *IEEE Communications Magazine*, 2025.
- [5] A. F. Ocampo *et al.*, “Reinforcement learning-driven service placement in 6g networks across the compute continuum,” in *2024 20th International Conference on Network and Service Management (CNSM)*. IEEE, 2024, pp. 1–9.
- [6] J. Santos *et al.*, “Hephaestusforge: Optimal microservice deployment across the compute continuum via reinforcement learning,” *Future Generation Computer Systems*, p. 107680, 2025.
- [7] J. Santos, “Towards cloud-native virtual reality applications: State-of-the-art and open challenges,” in *NGMSE2024, the Next-Generation Multimedia Services at the Edge: Leveraging 5G and Beyond workshop, co-located with ISCC2024*, 2024.
- [8] C. Carrión, “Kubernetes as a standard container orchestrator-a bibliometric analysis,” *Journal of Grid Computing*, vol. 20, no. 4, p. 42, 2022.
- [9] A. Al-Dulaimy *et al.*, “The computing continuum: From iot to the cloud,” *Internet of Things*, vol. 27, p. 101272, 2024.
- [10] X. Wang *et al.*, “Deep reinforcement learning: A survey,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 35, no. 4, pp. 5064–5078, 2022.
- [11] H. X. Nguyen *et al.*, “Digital twin for 5g and beyond,” *IEEE Communications Magazine*, vol. 59, no. 2, pp. 10–15, 2021.
- [12] N. P. Kuruvatti *et al.*, “Empowering 6g communication systems with digital twin technology: A comprehensive survey,” *IEEE access*, vol. 10, pp. 112 158–112 186, 2022.
- [13] H. Ahmadi *et al.*, “Networked twins and twins of networks: An overview on the relationship between digital twins and 6g,” *IEEE Communications Standards Magazine*, vol. 5, no. 4, pp. 154–160, 2022.
- [14] X. Lin *et al.*, “6g digital twin networks: From theory to practice,” *IEEE Communications Magazine*, vol. 61, no. 11, pp. 72–78, 2023.
- [15] D. Borsatti *et al.*, “Kubetwin: A digital twin framework for kubernetes deployments at scale,” *IEEE Transactions on Network and Service Management*, 2024.
- [16] P. Bellavista *et al.*, “Exploiting microservices and serverless for digital twins in the cloud-to-edge continuum,” *Future Generation Computer Systems*, vol. 157, pp. 275–287, 2024.
- [17] A. G. Wermann *et al.*, “Ktwin: A serverless kubernetes-based digital twin platform,” *Computer Networks*, p. 111095, 2025.
- [18] J. Santos *et al.*, “Efficient orchestration of service chains in fog computing for immersive media,” in *2021 17th International Conference on network and service management (CNSM)*. Ieee, 2021, pp. 139–145.