

Modular Multi-Domain Digital Twin Architecture: Sustainable Intent-Driven 6G Management

Berk Buzcu, *Student Member, IEEE*, Marcin Pakula, Gevher Yesevi Keskin, Laura Finarelli, *Student Member, IEEE*, Gianluca Rizzo, *Senior Member, IEEE*, Engin Zeydan, *Senior Member, IEEE*, Jorge Baranda, *Senior Member, IEEE*, Aitor Alcázar-Fernández, Javier Velázquez-Martínez, Luis M. Contreras, Gil Kedar, Paweł Kryszkiewicz, *Senior Member, IEEE*

Abstract—Future 6G networks will operate across distributed and heterogeneous domain infrastructures, making conventional single-domain management insufficient for proactive, trustworthy automation. Network Digital Twins (NDTs) enable what-if analysis, AI-assisted optimization, and risk-free validation of control actions before deployment, yet monolithic end-to-end twins remain impractical due to scalability, fidelity, and cross-domain coordination challenges. Accordingly, this paper proposes a Digital Twin-enabled 6G architecture that exposes NDT capabilities as a specialized service domain within a multi-domain orchestration framework built on a state-of-the-art service-based 6G architecture. A DT Orchestrator interprets *predictive* and *prescriptive* what-if queries and composes domain-specific DT modules and simulators on demand, while decision authority remains with the requesting entity. Furthermore, a generalized workflow covers telemetry synchronization, simulation-based decision support, and closed-loop execution. The framework is demonstrated through a green-networking use case that couples a system-level O-RAN cellular digital twin component with a two-stage solar-allocation simulator, evaluated over a 105-base-station deployment in Poznan using simulative datasets. Joint coverage and renewable optimization reduces daily grid consumption by 28.5% with 32 solar panels at the diminishing-returns threshold, with 17 base stations identified as both coverage-active and high-priority solar candidates as evidence that cross-domain NDT coordination enables sustainable, intent-driven 6G network management.

Index Terms—6G, Digital Twin, Edge Computing, Network Virtualization, Artificial Intelligence.

This work is supported by UNITY-6G project, co-funded from European Union's Horizon Europe Smart Networks and Services Joint Undertaking (SNS JU) research and innovation programme under the Grant Agreement No 101192650 and from the Swiss State Secretariat for Education, Research and Innovation (SERI).

B. Buzcu and L. Finarelli are with the University of Applied Sciences and Arts Western Switzerland (HES-SO Valais/Wallis), Switzerland, and with Technische Universität Berlin, Germany. (e-mail: berk.buzcu@hevs.ch, laura.finarelli@hevs.ch)

M. Pakula and P. Kryszkiewicz are with the Poznań University of Technology (PUT), Poland (e-mail: marcin.pakula@doctorate.put.poznan.pl, pawel.kryszkiewicz@put.poznan.pl).

G. Yesevi Keskin is with Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands, and with Eindhoven University of Technology, The Netherlands. (e-mail: g.yesevi.keskin@tue.nl)

E. Zeydan and J. Baranda are with the Centre Tecnològic de Telecomunicacions de Catalunya (CTTC), Castelldefels, Spain. (e-mail: ezeydan@cttc.es, jorge.baranda@cttc.es)

A. Alcázar-Fernández, J. Velázquez-Martínez, and L. M. Contreras are with Telefónica Innovación Digital, Madrid, Spain. (e-mail: aitor.alcazarfernandez@telefonica.com, j.velamar@telefonica.com, luismiguel.contrerasmurillo@telefonica.com)

G. Kedar is with Ceragon Networks, Israel (e-mail: gilke@ceragon.com).

G. Rizzo is with the University of Applied Sciences and Arts Western Switzerland (HES-SO Valais/Wallis), Switzerland, and with University of Turin, Italy. (e-mail: gianluca.rizzo@hevs.ch)

I. INTRODUCTION

The evolution toward new generations of cellular networks, and particularly Sixth Generation (6G), requires a shift from reactive troubleshooting to proactive, AI-driven network management. Future networks will operate in highly dynamic and heterogeneous environments, integrating radio access, transport, core, and edge-cloud infrastructures while meeting stringent requirements for latency, reliability, scalability, and sustainability. Traditional management frameworks, constrained by isolated domains, *telco-only* perspectives, static control policies, and offline optimization, are no longer sufficient.

To manage this complexity, 6G networks are expected to rely on the Network Digital Twin (NDT) as a synchronized virtual representation of network assets, services, and operational states. The NDT is not merely a monitoring tool, but a high-fidelity virtual environment that enables operators to simulate *what-if* scenarios, assess alternative configurations, and validate AI-driven decisions before live deployment. In this role, the NDT provides a risk-aware decision-support layer between network telemetry and operational control, improving the reliability of closed-loop automation. However, constructing a monolithic end-to-end NDT for future 6G ecosystems remains impractical due to the scale, heterogeneity, and computational complexity of multi-domain environments. A more feasible approach is to coordinate multiple domain-specific twins, simulators, and analytical models under a common orchestration framework. In this work, we adopt the inter-domain management and orchestration architecture proposed in [1], which enables coordinated service delivery across heterogeneous administrative and technological domains.

Beyond conventional telecom optimization, NDTs also create an opportunity to incorporate non-telecommunication information sources into network control loops. This is particularly relevant for sustainability-oriented 6G systems, where renewable energy availability, carbon-intensity indicators, and localized energy forecasts can influence network decisions. In this context, the NDT enables networks to evolve from being merely energy-efficient, i.e., consuming less power, to becoming carbon-aware, i.e., consuming energy when and where it is most sustainable. The need for this capability is reinforced by the evolution of the energy sector. Modern power grids are no longer passive distribution infrastructures; they increasingly rely on forecasting, demand-response mechanisms, and autonomous management of decentralized Renewable

Energy Sources (RES) [2]. Consequently, the 6G network should not treat the grid as an opaque energy source, but as an intelligent external domain whose state can be exploited for network optimization. The NDT acts as the mediator in this relationship, ingesting energy-domain information such as predicted carbon intensity or renewable energy surpluses and translating it into actionable network-management decisions.

Motivated by these challenges, this paper proposes a Digital Twin (DT)-enabled multi-domain 6G architecture in which the NDT acts as a specialized orchestration capability for scenario analysis, intent-driven optimization, and coordinated decision support. The framework integrates NDT functionalities within a multi-domain orchestration layer, enabling coordinated management of heterogeneous network segments through structured what-if analysis and closed-loop control. Its applicability is demonstrated through a sustainability-oriented use case that combines network simulation outputs with renewable energy information to optimize base-station operation under energy-aware objectives.

The paper is organized as follows. Section II reviews the state of the art on DTs and multi-domain NDT architectures, identifying the gap addressed in this work. Section III presents the proposed multi-domain 6G orchestration architecture and describes how NDT capabilities are integrated to support coordinated service management. Section IV details the corresponding service workflow, including DT instantiation, what-if analysis, KPI-driven optimization, and closed-loop execution. Section V introduces the sustainability use case, where the NDT Manager orchestrates the interaction between a system-level cellular RAN DT Module, and a energy simulator to resolve a cross-domain what-if scenario. Section VI elaborates on the evaluation method of the framework. Finally, Section VII concludes the paper.

II. RELATED WORK

This section reviews the evolution from general DT concepts toward NDT architectures for future wireless systems. First, it consolidates the foundational understanding of DTs as synchronized, purpose-driven virtual representations. Then, it analyses the role of NDTs in 6G networks, with emphasis on decision support, pre-validation, and trustworthy automation. Finally, it discusses the transition from scoped, domain-specific NDTs toward multi-domain NDT coordination, which directly motivates the architecture proposed in this work.

A. Foundations of Digital Twin Architectures

The concept of a DT originated as a virtual counterpart to a physical asset or system, continuously updated with operational data to support monitoring, analysis, prediction, and optimization. Over time, the literature has shifted from broad and sometimes inconsistent interpretations to a more operational understanding: a DT is not a static model, simulator, or visualization tool, but a synchronized and purpose-driven virtual representation that supports system-level reasoning. Its value lies not in reproducing every internal building block of the physical system, but in providing purpose-specific fidelity,

i.e., accurately capturing the behaviours, KPIs, or decision outcomes relevant to the intended use case.

For networking, this distinction is critical. Communication systems are distributed, dynamic, and multi-layered, making permanent full-fidelity replication impractical. Accordingly, recent NDT-oriented work increasingly frames the twin as a modular architectural capability composed of data collection and synchronization mechanisms, virtual models, analytics functions, lifecycle management procedures, and interfaces exposed to external consumers [3]–[5]. This interpretation is increasingly reflected in standardization efforts such as ITU-T Y.3090 [6], the IRTF NMRG reference architecture draft [7], and 3GPP TS 28.561 [8], which move the discussion from abstract definitions toward scenario-aware, lifecycle-managed, and automation-oriented NDT realizations.

From this perspective, an NDT analysis evaluates the operational consequences of candidate network decisions before they are applied to the production system. Such analyses may address planning, configuration validation, fault assessment, congestion mitigation, recovery strategies, energy optimization, or ML-enabled control. The common objective is to reduce operational risk and improve the reliability of automated decision-making by evaluating the expected impact of alternative actions under realistic assumptions.

B. Network Digital Twins in Future Wireless Networks (6G)

In 6G, the NDT evolves from a supporting management tool into a core enabler of trustworthy, proactive, and automation-ready network operations. Future wireless systems must support dense multi-RAT deployments, cloud-edge-native service execution, AI-native control loops, service-aware optimization, and stringent requirements for latency, reliability, and energy efficiency. In these conditions, validating decisions directly on production infrastructure becomes increasingly costly, risky, and often impractical. The NDT addresses this limitation by providing a synchronized and scenario-aware environment where network behavior, service performance, and control actions can be evaluated before implementation in the real system [9]–[11].

Recent studies therefore position NDTs as decision-support and pre-validation tools for 6G planning, optimization, troubleshooting, policy verification, and AI/ML lifecycle support [12]–[14]. In this role, the NDT acts as an intermediate reasoning layer between live telemetry and control. It allows alternative actions to be explored before affecting the production network, supports the comparison of AI-driven decisions, and improves the robustness of closed-loop management.

This role is especially relevant when network optimization depends on information that is not available through conventional telecom interfaces. Sustainability-oriented operation is a representative example: energy-aware policies can be derived from traffic and power-state information, but carbon-aware policies also require external indicators such as renewable energy availability, grid carbon intensity, and energy forecasts. This extends the NDT beyond domain-internal modelling and motivates architectures capable of combining telecom and non-telecom information sources in a consistent decision process.

TABLE I
COMPARISON OF CLASSICAL AND MULTI-DOMAIN NDT LITERATURE

Class	Representative works	Domains and segments	Positioning w.r.t. this paper
Classical NDTs	Nguyen <i>et al.</i> [15], Sanz Rodrigo <i>et al.</i> [16], Nardini and Stea [17], Vilà <i>et al.</i> [18], Apostolakis <i>et al.</i> [19], Vilalta <i>et al.</i> [20], Cherini <i>et al.</i> [21], Polverini <i>et al.</i> [22]	Mobile operator, slice mgmt., simulation/service across RAN, core, and edge/cloud segments, optical transport, and routing/control domains, and IP/BGP segments	NDT as support for 5G/B5G operation, synchronization, and simulation-based evaluation; bounded to single operator/system view, lacking inter-domain orchestration.
Multi-domain NDTs	Wang <i>et al.</i> [23], Robitzsch <i>et al.</i> [24], Giardina <i>et al.</i> [25], Faye <i>et al.</i> [26], Zaki-Hindi <i>et al.</i> [27], Raza <i>et al.</i> [28]	Service, logical, and resource domains across cross-layer 6G architectures and end-to-end service paths	Introduce cross-domain coordination across heterogeneous environments; partial definition of inter-twin orchestration and result consolidation.

C. Classical NDT vs Multi-Domain NDT

A useful way to position the literature is to distinguish between classical, domain-scoped NDTs and multi-domain NDT approaches. Classical NDTs typically serve as domain-specific capabilities limited to a particular operational segment, such as RAN, transport, core, edge/cloud, or service management. They provide mature mechanisms for modelling, synchronization, analytics, and what-if evaluation within bounded environments. Moreover, domain-specific NDTs can preserve provider privacy by exposing only abstracted models, capabilities, or simulation outcomes, without requiring each domain owner to disclose detailed network configurations or proprietary operational data. However, they usually do not address how decisions spanning several technological or administrative domains should be decomposed, distributed, evaluated, and consolidated.

Multi-domain NDT approaches address this broader scope by coordinating multiple models, simulators, or twin capabilities across heterogeneous environments rather than assuming a single monolithic twin. This direction is aligned with the operational reality of future 6G systems, where end-to-end services depend on the joint behavior of radio, transport, core, cloud-edge, and external domains. The key challenge is therefore not only to construct accurate domain-level twins, but also to define how an end-to-end intent or what-if request is translated into domain-specific analyses and how their outputs are aggregated into an actionable operational decision. In addition, coordination mechanisms must also account for privacy and confidentiality constraints, ensuring that domain providers can participate without exposing detailed network configurations, internal policies, or proprietary operational data.

Table I summarizes representative works from this perspective. Each entry is characterized by the domains it addresses, covering functional, technological, or administrative scopes across network segments such as Radio Access Network (RAN), transport, core, and edge/cloud. The comparison highlights a clear trend: while scoped NDTs provide mature mechanisms for modelling and evaluation within bounded environments, multi-domain approaches extend the architectural scope but still lack a systematic framework for inter-domain coordination, result consolidation, and end-to-end decision support.

Overall, the literature indicates that current NDT research is progressing from bounded, segment-specific realizations toward broader multi-domain architectures. However, this tran-

sition remains incomplete. Existing works do not yet fully specify how end-to-end decisions should be translated into distributed twin queries, how heterogeneous simulation outputs should be consolidated, or how the resulting recommendations should be exposed to orchestration and closed-loop control functions. This gap directly motivates our work, which focuses on the role of the Multi-Domain Manager as the coordination function responsible for interacting with domain-specific twin capabilities and consolidating their outputs for end-to-end decision support.

III. BACKGROUND: MULTI-DOMAIN ORCHESTRATION ARCHITECTURE FOR 6G NETWORKS

Future wireless network environments, especially 6G systems, require a shift from traditional single-domain management to highly federated, multi-domain ecosystems. These networks will integrate diverse access technologies, non-terrestrial components, and complex service chains distributed from deep edge deployments to centralized clouds. To manage the scale and heterogeneity of these cross-domain environments, the UNITY-6G project [29] proposes a distributed architecture based on the Service-Based Management Architecture (SBMA) approach presented in [30].

At the core of this distributed paradigm, a layered control structure decouples global service composition from localized infrastructure execution. As shown in Figure 1, the Inter-Domain Management Orchestrator (IDMO) entity is at the top of the architecture. The IDMO oversees multiple administrative and technological domains to establish end-to-end (E2E) network services at the highest level of hierarchical control. When an E2E service intent request is received, the IDMO is responsible for decomposing the high-level intent into functional components (e.g., RAN, Transport, Mobile Core). It then executes a multi-criteria decision-making process to select the optimal destination domains based on available service capabilities, resource availability, and latency or coverage suitability.

At each selected domain, the Domain Management Orchestrator (DMO) receives, via the Service Adapter module, the decomposed service component to perform its Life Cycle Management (LCM). The DMO optimizes intra-domain resources, deploys virtualized network functions (VNFs/CNFs) at selected Point of Presences (PoPs), monitors domain and service-level KPIs, and reports status to the IDMO through interaction with the underlying Infrastructure Domain Managers (IDMs), which operate the computing resources at the




Fig. 1. Multi-domain Orchestration architecture for unified 6G networks

infrastructure level. All communication related to management, orchestration, and policy synchronization between the IDMO and domains occurs over the Cross-Domain SBMA bus. This vertical approach enables the IDMO to maintain global coordination while delegating local orchestration and infrastructure control to the DMOs and IDMs, respectively, reducing the computational workload of decision making. In later sections, we explain how our DT architecture integrates into this broader network architecture and which components replace the general structure.

IV. PROPOSED APPROACH

Given the increasing complexity and heterogeneity of future wireless networks [31], accurately simulating the entire network continuum is computationally prohibitive. The diversity of user equipment, proprietary hardware, and domain-specific variables makes relying on a single, monolithic simulator impractical for actionable what-if analysis. As a result, the development of holistic DT is significantly hindered, since maintaining real-time, end-to-end fidelity across all network domains is extremely difficult to achieve in practice, especially in real time.

Nevertheless, because service delivery spans highly interdependent segments, including the RAN, transport layers, mobile core, and distributed edge-cloud infrastructures, isolated domain-specific testing is insufficient. Optimizing these networks requires accounting for cross-domain effects, such

as a RAN configuration shift directly impacting edge computational loads. Therefore, a multi-domain NDT is essential to provide the comprehensive network visibility necessary to optimize interdependencies and guarantee deterministic performance across the entire service path. Accordingly, in this section we describe our proposed architecture, which can be broken down by task, and our query-based workflow that uses this architecture to resolve DT-related queries.

A. DT Architecture

To overcome the mentioned constraints, the proposed architecture adopts a pragmatic, query-driven approach: rather than maintaining a monolithic continuous simulation, a DT Manager receives a what-if query, identifies which available simulators are relevant to that query's domain requirements, invokes them with the appropriate parameters, and consolidates their outputs into a unified result.

The composition of simulators is determined by the query, not pre-configured in advance, which allows the same architectural framework to serve radically different scenario types by selecting different functional subsets. This design is a direct architectural response to the impossibility claim above, because no single function can span all domains. An important detail of this design is that the orchestration layer can natively ingest external, non-telco contextual information (e.g. local renewable energy availability and smart grid telemetry) by treating simulators on equal footing with digital twin modules within the same composition framework.

Figure 2 illustrates the internal organisation of the DT domain and its interactions with the selected future wireless networks architecture (See Section III) multi-domain architecture. At the top, the Cross-Domain Service-Based Management Bus (SBMA) serves as the northbound communication plane through which the other domains forward what-if tasks and receive the DT's response. A task may be *predictive* (e.g., *what happens if I install 25 solar panels?*), to which the service replies with evaluation results, or *prescriptive* (e.g., *which configuration maximizes greenness under a fixed panel budget?*), to which the DT replies with a recommended configuration. In either case, enforcement authority remains with the requesting entity.

Within the DT domain, the *DT Orchestrator* acts as the control entry point: its *DT Resolver* receives incoming tasks from the SBMA, while the *DT Instantiator* creates a dedicated *DT Manager* instance scoped to the specific query so that concurrent, query-isolated twin instances without contention are instantiated. The DT Manager is the central coordination unit. Its *Data Adaptor* ingests both standard network telemetry (e.g. traffic flow, user association) and non-telco contextual feeds (e.g. real-time renewable generation availability, grid carbon intensity) through the data-continuum. The *ML/RL Sandbox* provides an isolated environment for offline training of reinforcement learning agents on synthesised DT instances, decoupling the model development cycle from live network operation. The *DT Executor* pipelines the external simulators (e.g. *Energy Simulator*) with the DT Modules (e.g. *RAN DT Module*) at query time to enable multi-domain what-if




Fig. 2. Overall digital architecture and corresponding components interactions. Northbound interfaces represent how the applications interact with the DT whereas Southbound interfaces represent where the twin data comes from [6].

reasoning without requiring a unified simulation model to answer a certain query as explained in further sections (See Section V-C) and consolidates their respective outputs into a unified, actionable result. Alongside the DT Manager, the *DT Repository* maintains versioned virtual models spanning energy, RAN, and transport domains, coordinated through a *DT Registry* that governs lifecycle and versioning across successive what-if evaluations. At the bottom of the architecture, the *Data Continuum* and its Southbound Interfaces expose a unification layer that continuously synchronises physical infrastructure state into the Repository, sustaining the fidelity of the virtual models over time.

B. Network Digital Twin Service Workflow Description

A central design point is that the DT operates in parallel with service establishment rather than as a precondition for it: an end-to-end service is first instantiated with a baseline configuration through the standard cross-domain SBMA procedure sketched in Section III, and the DT subsequently derives refined configurations while drawing on telemetry from the already running deployment. This avoids the impasse of having to identify a configuration from a system that is not yet operational and ensures that the DT reasons over a measurable, live state.

When an end-to-end intent expresses an analytical requirement (e.g., maximizing renewable utilization under a bounded outage), the inter-domain coordinator decomposes the intent into per-domain components. The RAN, Transport, and mobile Core components are dispatched to their respective appropriate domain, where DMOs bring up the corresponding service component, while a separate DT sub-request carrying only the analytical intent (target Key Performance Indicators (KPIs), scope, and time horizon) is forwarded through the SBMA to the DT domain. Domain selection reuses the capability-based mechanism (See Section III) where the DT domain advertises its modelling capabilities (e.g., RAN, Core, Transport), and is selected when those capabilities cover the requested scope. The coordinator does not specify which simulators to run or how to compose them, so the workflow is not bound to a specific inter-domain coordinator implementation and can equivalently be initiated by another domain holding that role.

Composition is therefore a responsibility of the DT Orchestrator. The DT Resolver interprets the analytical intent and consults the DT Registry to identify the virtual models and simulators in the DT Repository that match the requested scope; the DT Instantiator then allocates a query-scoped DT Manager. The DT Manager populates the selected models with topology, configuration, and live telemetry obtained from the running service through the Data Continuum, complemented

when needed by non-telco feeds such as renewable generation forecasts or grid carbon intensity. The DT Executor invokes the composed simulators and consolidates their outputs. This is also where cross-domain effects are captured: an individual DMO is by construction limited to the visibility of its own domain, whereas the DT Orchestrator selects simulators that jointly cover the relevant domains and the DT Manager combines their results into a single evaluation.

Candidate configurations are ranked against the requested KPIs and the best-performing one is returned to the DT Orchestrator, which forwards it through the SBMA to the execution DMOs that own the affected resources (RAN, Core, Transport, or any combination thereof). Each receiving DMO enforces the configuration through its associated IDMs. The resulting operational state is reflected back into the Data Continuum and feeds subsequent DT evaluations, closing the loop. DT Service Request box illustrates a concrete instance of this workflow.

C. Network Twinning

Network twinning denotes the process by which selected operational evidence from the deployed system is transformed into an executable representation that can be used by the DT environment. In the proposed architecture, this transformation is not limited to copying static inventory data, but relies on exposure mechanisms capable of presenting the real system in a structured and machine-consumable form. For instance, when an operational domain is controlled or observed through Software Defined Networking (SDN), an interface such as TerraFlow SDN [32] can serve as an abstraction layer between heterogeneous network resources and the DT environment. Instead of requiring the simulator to interact directly with individual devices, vendor-specific management systems, or low-level configuration procedures, the controller exposes a normalized view of the managed domain through programmable interfaces. This view captures the operational information required for the requested evaluation and presents it in a form that can be interpreted by the twinning function.

The twinning function then maps this exposed operational state into the modelling primitives required by the selected DT modules or simulators. In the case of an NS-3-based [33] component, this may involve translating the exposed system view into simulated entities, relationships, parameters, and runtime conditions that correspond to the relevant part of the real network. The same principle applies to other domain models: the objective is not to reproduce the entire deployed system, but to instantiate a scoped executable model whose structure and parameters are grounded in the current or recent operational state. In this sense, TerraFlow SDN and similar exposure functions make twinning possible by converting heterogeneous managed resources into a consistent control-plane representation that can be consumed, parameterized, and updated by the DT. The resulting twin is therefore not a passive database replica, but a purpose-specific operational model that supports simulation, analysis, and decision validation before actions are applied to the real system.

DT Service Request

Service Intent: Maximise renewable energy utilisation across the RAN within a fixed solar panel budget.

Target KPI:

- Minimise grid energy consumption while keeping network outage below **5%**

Context:

- 105 base stations operating under a cell on/off switching policy
- Local solar generation available; grid carbon intensity varies hourly

DT Task:

- Simulate network load and BS operational states under the active coverage policy
- Evaluate solar dispatch scenarios across candidate BS deployments and panel capacities

Expected Output:

- Optimal set of BSes for solar panel installation maximising grid displacement
- Projected grid energy reduction per deployment size

Action:

- Forward recommended BS configuration via SBMA to the RAN domain for execution at DMO

V. GREEN NETWORK USE CASE

Realizing the sustainability goals of modern telecommunication networks requires a paradigm shift from simply managing operational energy efficiency to actively expanding localized renewable generation. A fundamental problem in 6G service deployment methodologies is their highly siloed nature where they focus almost exclusively on optimizing coverage, capacity, and internal traffic loads. Consequently, they remain "blind" to the external spatial and environmental factors that present physical opportunities to make the infrastructure greener, specifically the untapped solar potential of existing base station (BS) sites and their surrounding topologies.

By strategically deploying solar panels as a supporting instrument for base station energy demand, the network can drastically decrease its dependency on the main energy grid. Not only it is advantageous for the base stations, but it also helps to decrease the stress on the main grid, especially when energy demand peaks during the day. However, planning this integration requires overcoming the divide between telecommunications engineering and environmental modeling. To address this, a dual-simulation approach is essential. First, a dedicated energy simulator is required to evaluate the greenness generation potential considering varying settings. Second, a module representing the network is needed to accurately model the traffic-driven power consumption and other operational constraints of the base stations.

A. Sustainability Aware Simulator

The Energy Simulator is a non-domain-specific simulation component that models the spatiotemporal availability of solar

energy and its allocation to a distributed set of demand nodes (e.g. base stations). Unlike a static data lookup, the simulator accepts a parameterized scenario (a candidate deployment or a set of active components, and a time horizon) and evaluates the energy outcomes of that scenario under varying solar energy generation capacities. For the given renewable configurations (i.e. generation capacity, number of panels), the simulator evaluates how it can support the changing energy needs of base stations during the day. Algorithm 1 outlines the methodology of this simulator, which operates in two primary stages: (i) candidate node selection and (ii) temporal resource dispatch.

The first stage (i) identifies target stations that contribute most significantly to system load during critical periods. We evaluate a set of historical system states, each showing different selected base stations to be turned on/off, and their corresponding effect on the energy grid on different times of the day. The goal of this stage is to select candidates that could benefit the most from the dispatched solar energy, therefore it is important to identify the critical system states that demonstrates strong correlation with the hourly patterns of the solar generation. To do this, each set is ranked based on the product of grid stress severity and concurrent solar resource availability. This step isolates periods characterized by high demand and viable resource generation. From the top N critical system states, we identify the active nodes. Each node is assigned a utility score, calculated by multiplying its maximum energy demand during these critical states by its observation frequency (the ratio of states in which the node is active). The nodes are then sorted in descending order based on this score, and we select the top m nodes for resource integration.

The second stage (ii) details the time-step dispatch of the potentially generated solar energy to the selected candidates. The goal of this step is to distribute the hourly generated energy to the different base stations to support the grid during peak energy hours (in other words; peak shaving). For a given solar capacity factor (representing the ratio of how close the actual output can be compared to the generation capacity), we calculate the total available resource for each time interval. Concurrently, the algorithm determines the maximum demand for each candidate node at that specific interval. We then sequentially allocate the available resource to the nodes. Nodes that demand the highest energy are supplied first until the interval's resource generation is fully utilized. We repeat this dispatch process across a range of generation capacities by increasing the number of solar panels to evaluate marginal gains of adding each, and to identify the ideal number of panels that the candidate stations can benefit based on the diminishing returns.

B. RAN DT Module

The module models a multi-cell 5G network, where base station locations are obtained using real-world data sources, using publicly available base station location databases and BTS search [34], to ensure a realistic spatial topology. Each base station is characterized by its transmission parameters, antenna configuration, and maximum power consumption,

Algorithm 1 Two-Stage Greedy Heuristic for Resource Dispatch

Require: Set of all system states $\mathcal{S}$, Set of all nodes $\mathcal{V}$, N (critical states limit), m (candidate limit), K (resource scaling array)

Ensure: Temporal allocation $a_v(t)$ for varying capacities $k \in K$

Stage 1: Candidate Selection

- 1: **for** each state $s \in \mathcal{S}$ **do**
- 2: Rank score $R_s \leftarrow \text{stress}_s \times \text{resource_avail}_s$
- 3: **end for**
- 4: $\mathcal{S}_{critical} \leftarrow$ top N states sorted by R_s (descending)
- 5: **for** each node $v \in \mathcal{V}$ **do**
- 6: $f_v \leftarrow \frac{\text{active_states}_v}{|\mathcal{S}|}$
- 7: $U_v \leftarrow 0$
- 8: **for** each $s \in \mathcal{S}_{critical}$ where v is active **do**
- 9: $U_v \leftarrow U_v + (D_{max,v,s} \times f_v)$
- 10: **end for**
- 11: **end for**
- 12: $\mathcal{V}_{cand} \leftarrow$ top m nodes sorted by U_v (descending)

Stage 2: Temporal Dispatch

- 13: **for** each capacity scalar $k \in K$ **do**
- 14: **for** each time step t **do**
- 15: $R_{total}(t) \leftarrow k \times \text{unit_capacity} \times \text{resource_avail}(t)$
- 16: Calculate worst-case demand $D_v(t)$ for all $v \in \mathcal{V}_{cand}$
- 17: Sort $\mathcal{V}_{cand}$ by $D_v(t)$ in descending order
- 18: $R_{remaining} \leftarrow R_{total}(t)$
- 19: **for** each $v \in \mathcal{V}_{cand}$ **do**
- 20: $a_v(t) \leftarrow \min(D_v(t), R_{remaining})$
- 21: $R_{remaining} \leftarrow R_{remaining} - a_v(t)$
- 22: **if** $R_{remaining} == 0$ **then**
- 23: **break**
- 24: **end if**
- 25: **end for**
- 26: **end for**
- 27: **end for**

which are later used in the energy-aware analysis. The User Equipments (UEs) are modeled using a stochastic mobility process based on a random walk model with dynamic spatial traffic distributions over time [35]. This approach captures temporal and spatial movements in network load, which are critical for evaluating the interaction between traffic demand and renewable energy availability. At each simulation step, UEs are associated with serving base stations based on signal strength and path loss conditions. The radio resource management follows a round-robin scheduling scheme, allowing the simulator to estimate per-user throughput, SINR, and resource utilization. Based on these metrics, the instantaneous and worst-case load of each base station is computed. The DART simulator produces time-resolved network states, including base station load, user distribution, and spectral efficiency.

C. Simulator Composition and Query Resolution

Figure 3 illustrates the interaction between the DT Manager, the RAN DT Module and the Energy Simulator for the green




Fig. 3. Process of the interaction between the DT modules, simulators and the DT Manager.

networking use case in this study. The process begins when the DT Orchestrator receives the what-if query from the SBMA and forwards it to the DT Manager, which then parses the query and determines the requirements: (i) RAN Simulator for network-layer state generation and (ii) the Energy Simulator for solar allocation as elaborated earlier. In the first phase, O-RAN Simulator is invoked with the base station topology and a stochastic user mobility model. It produces a set of time-resolved network states $\mathcal{S}$, where each state captures per-BS load, active or inactive status, and worst-case power consumption. These states represent the demand-side input that the Energy Simulator cannot determine independently. In the second phase, the DT Manager passes $\mathcal{S}$, the solar generation profile, and a panel budget parameter m to the simulator. It executes Algorithm 1 to identify the m candidate stations and compute temporal solar dispatch, returning the projected grid displacement for each candidate deployment size.

The DT Manager selects the deployment scenario that maximizes grid energy displacement within the panel budget and returns this recommendation as the actionable output to the IDMO. The value of this composition is that neither simulator alone can answer the query: O-RAN Simulator cannot evaluate solar allocation, and the Energy Simulator cannot determine which stations are operationally active without the network simulation.

VI. EVALUATION

We apply the generalized framework to localized 5G telecommunications to evaluate the viability of integrating solar photovoltaics (PV) to improve the greenness of the network. The generalized parameters of Algorithm 1 are directly mapped to our 5G deployment model. Specifically, the

nodes ($\mathcal{V}$) represent clustered 5G network base stations, while the system states ($\mathcal{S}$) include historical network configurations and localized traffic load profiles. We redefine system stress to reflect localized network performance degradation, quantified as the service failure rate (i.e., $1 - \text{service success rate}$, representing unserved traffic or dropped connections). Resource availability remains the hourly solar capacity factor. The product of these metrics identifies timeframes when a base station is under severe traffic load and solar generation is viable. Because nodes experiencing high service failure operate at maximum power draw, targeting them ensures the generated green energy is fully consumed locally by nodes estimated to be under the most duress, maximizing the displacement of grid-derived carbon emissions. Furthermore, the demand (D_v) represents the maximum power consumption of a specific base station during a given configuration or hour. Applying this context to the energy heuristic, our model dynamically routes simulated solar power to the most heavily burdened base stations. Iterating through the solar array sizes (N_{PV}) determines the marginal utility of expanding physical hardware deployment, ensuring that renewable investments are directed to the nodes where they achieve the greatest reduction in the network's overall carbon footprint.

The Cell On/Off Switching (COOS) xApp (described in [35]) is located in the Near-RT RIC in the RAN DT Module and is responsible for implementing the logic to turn off certain BSes when necessary. Specifically, each BS $v \in \mathcal{V}$ monitors its traffic load $L_v(t)$ and exchanges load information with its neighboring nodes. A BS is switched off when its load falls below a predefined threshold α_{off} , i.e., $L_v(t) < \alpha_{\text{off}}$, indicating underutilization. Conversely, a previously deactivated BS is reactivated when the average load




Fig. 4. Poznan base station placement and average power consumption during the day.

of its neighboring cells exceeds a threshold α_{on} , i.e., $L_{(v)}(t) > \alpha_{on}$, suggesting potential overload in the surrounding network. Typically, $\alpha_{on} > \alpha_{off}$ is assumed to prevent the ping-pong effect. From an energy perspective, this mechanism reduces total network power consumption by eliminating the static power component of underutilized BSes. Traffic consolidation increases the utilization of active nodes, which is more energy-efficient.

A. Assumptions

The network topology comprises 105 base stations. Their geographic coordinates, obtained from a real-world database [34], are projected onto a local Cartesian reference frame to enable spatial operations, as shown in Fig. 4. All BSes operate at a carrier frequency of 2100 MHz with a bandwidth of 20 MHz, and a transmit power of 37 dBm, while inter-cell interference is explicitly taken into account. The UEs are randomly distributed over the considered area, with their number varying between 200 and 600 over a 24-hour period to reflect realistic traffic dynamics. Each UE follows a random walk mobility model with an average speed of 5.4 km/h and generates traffic with a constant bitrate demand of 1 Mbps.

Ten validated BS set configurations (Sets 0-9) represent alternative network operating states with mean outage rates ranging from 8.5 % to 15.9 % (Table II) where all the sets have 25 random base stations on.

The coverage optimization algorithm (BS Set 10) achieves a mean outage of 2.46 %, meanwhile the energy optimization algorithm (BS Set 11) achieves a mean outage of 3.88 %, both activating 25 BSes and deactivating the remaining 80. We refer to this as the *Algorithm-25 scenario*. Finally, the candidate selection step of the energy simulation produces a set of BSes similarly comprised of 25 BSes, we refer to the utilization of these BSes in the RAN DT Module as the *Heuristic-25 scenario*.

The system integrates real-world grid generation data from the European Network of Transmission System Operators for Electricity (ENTSO-E) with location specific solar generation data obtained from the renewables.ninja platform using the NASA MERRA-2 reanalysis dataset [36], [37]. A single representative measurement point at the Poznan city centre was used, consistent with the homogeneous spatial scale of

TABLE II
MEAN NETWORK OUTAGE PER BS SET CONFIGURATION.

BS Set	Mean Outage (%)
Set 1	8.51
Set 4	8.89
Set 2	9.51
Set 8	11.16
Set 0	12.19
Set 9	12.21
Set 7	13.10
Set 6	13.94
Set 5	15.41
Set 3	15.87
Set 10 (Algorithm-25)	2.46
Set 11 (Heuristic-25)	3.88

the study area. The modelled PV output peaked at 0.608 kW per panel on the simulation date. Solar capacity factors were computed hourly as:

$$\text{solarCF}(t) = \frac{P_{\text{solar}}(t)}{P_{\text{rated}}} \in [0, 1] \quad (1)$$

where $P_{\text{rated}} = 1 \text{ kW}$ per panel (simulator default capacity). The energy dispatch calculations are tested across settings with different number of panels $k \in [20, 50]$.

B. Results

Figure 5 shows the geographic location of all the BSes, colour-coded by their joint classification under the two approaches¹.




Fig. 5. BS coverage map. Orange (17 BSes): active in the algorithm and selected by the energy simulator. Green (8 BSes): selected by the energy heuristic but inactive per the O-RAN simulator. Blue (8 BSes): active. Grey (72 BSes): inactive and not selected by either approach.

¹All relevant results and data can be found in <https://tinyurl.com/yt75257b>




Fig. 6. Energy demand peak shaving for Algorithm-25 (Poznan) and same configurations for North Africa (Sahara) use case. Blue rectangle shows the threshold for diminishing benefits of adding a panel for the system in Poznan, whereas orange rectangle shows the same for North Africa use case.

The orange BSeS represent the productive agreement between the two frameworks where 17 stations are both active in the coverage-optimal configuration and ranked as high-priority solar candidates by the energy heuristic. The solar panels installed at an orange BS offset real, ongoing grid consumption at a station the network genuinely needs to keep running. The green BSeS tell the opposite story where the energy heuristic selected them because they appear active and power-hungry in the random validation configurations, but the coverage algorithm has determined that the network performs better without them.

The blue BSeS represent these stations are active under the algorithm and consume grid energy throughout the day, yet the energy heuristic did not identify them as solar candidates. They represent the power drawn from the grid that the solar could have offset. The large gray cluster, comprising 72 BSeS, is ignored by both approaches so these stations are switched off by the algorithm and ignored by the energy heuristic when making either candidate list.

Once the energy simulator identifies the base stations to connect to solar energy, the next step is to allocate the expected energy generation to those base stations using varying numbers of solar panels. Since the main goal of this allocation is to support the grid during periods where energy demand is at the peak during the day (i.e., peak shaving), Figure 6 shows the total maximum peak energy saved as the number of panels (and consecutively, the total solar energy that can be injected) increases. When considering the total energy consumption needs of the selected base stations operating in the Algorithm-25 scenario, we observe that after 32 solar panels, each additional panel yields diminishing returns.

There are two factors that define the diminishing effect when adding a new panel: i) the total amount of energy that the given set of stations need to consume, and ii) potential solar generation amount depending on the location of the use case. If the same set of base stations with the same configurations had been located in a different region with more solar benefits, we would expect to observe a different number of panels that fits the requirements of the given region. To test this hypothesis, we assume the same network to be operating in the North Africa at the same time period of the original use case. By using our solar energy dispatcher, we allocate the updated amount of solar energy to the base stations with same

power consumption and outage values. Similar to the original observation, the results show as the number of panels increase, more peak energy per hour could be relieved from the main grid, until the panels don't contribute to the peak demand periods anymore. For the North Africa use case, we see this critical number of panels is 28, as shown in Figure 6.

The Algorithm-25 scenario establishes the upper bound: among 25 operating base stations, 17 benefit from dispatched solar energy. The stations selected by Algorithm-25 collectively require 458.30 kWh from the grid over the day, of which 32 panels supplying the selected base stations provide 130.4 kWh. Algorithm-25 selects 25 base stations based on network traffic distribution, balancing the need to keep the outage metric low with reducing power consumption. As a result, net grid consumption decreases to 327.9 kWh, corresponding to a 28.5% reduction. Alternatively, if all 25 base stations connected to the solar resources operate instead of 17, their total energy consumption over the day becomes 444.76 kWh. With the same solar panel assumptions, the load that can be supplied by solar remains 130.4 kWh, and the net grid consumption decreases by 29.3%.

It should be noted that if two distinct sets of operating base stations yield nearly the same total energy consumption, the region's solar potential remains unchanged. Although the reduction in net grid consumption is a useful metric for assessing the impact of solar potential, Cell On/Off switching decisions are driven by objectives that differ from those of the energy simulator. This complementary simulation helps network operators make long-term decisions to meet sustainability targets while maintaining network performance and reliability.

VII. CONCLUSIONS & FUTURE WORK

This paper presented a DT-enabled 6G framework in which NDT capabilities are integrated as a specialized service within a multi-domain orchestration architecture. The proposed approach supports what-if scenario analysis, KPI-driven optimization, and safe validation of candidate actions before deployment in operational networks. To demonstrate its practical value, a sustainability-oriented use case was developed where network operational data were combined with renewable energy availability to enable carbon-aware base-station operation. Results showed that aligning renewable resource placement with network-aware activation decisions can improve local green-energy utilization while reducing avoidable grid consumption up to 28.5%, highlighting the benefit of combining telecom and external data sources within future control loops. Overall, the proposed framework illustrates how NDTs can evolve from passive monitoring tools into active decision-support components for intelligent 6G orchestration.

A natural extension of this framework is its use in online, fault-driven settings, where choices such as which simulations to run are made on the fly. For instance, the DT Manager could be asked: *"How would network KPIs be affected if the solar panel connected to a base station were to fail?"*. This scenario should trigger a coupled simulation in which the energy simulator represents the reduction in generation

capacity at that node, while O-RAN Simulator assesses the resulting rise in grid reliance and any loss of coverage under the updated energy budget. Such disaster-response scenarios would demonstrate the online, closed-loop potential of the architecture and are a natural next step following the offline planning use case presented here.

REFERENCES

- [1] B. Buzcu *et al.*, “Energy-aware multi-domain orchestration architecture for sustainable 6g networks,” in *Proceedings of the 21st IFIP/IEEE Wireless On-Demand Network Systems and Services Conference (WONS)*. IEEE, 2026.
- [2] Y. Sanjalawe, S. Al-E’Mari, S. Fraihat, S. Naser Makhadmeh, and E. Alzubi, “Ai-powered smart grids in the 6g era: A comprehensive survey on security and intelligent energy systems,” *IEEE Open Journal of the Communications Society*, vol. 6, pp. 7677–7719, 2025.
- [3] L. U. Khan, Z. Han, W. Saad, E. Hossain, M. Guizani, and C. S. Hong, “Digital twin of wireless systems: Overview, taxonomy, challenges, and opportunities,” *IEEE Communications Surveys & Tutorials*, vol. 24, no. 4, pp. 2230–2254, 2022.
- [4] R. Poorzare, D. N. Kanellopoulos, V. K. Sharma, P. Dalapati, and O. P. Waldhorst, “Network Digital Twin Toward Networking, Telecommunications, and Traffic Engineering: A Survey,” *IEEE Access*, vol. 13, pp. 16 489–16 538, 2025.
- [5] B. Qin, H. Pan, Y. Dai, X. Si, X. Huang, C. Yuen, and Y. Zhang, “Machine and Deep Learning for Digital Twin Networks: A Survey,” *IEEE Internet of Things Journal*, vol. 11, no. 21, pp. 34 694–34 716, 2024.
- [6] *Digital twin network — Requirements and architecture*, International Telecommunication Union Std. ITU-T Y.3090, feb 2022.
- [7] C. Zhou, H. Yang, X. Duan, D. Lopez, A. Pastor, Q. Wu, M. Boucadair, and C. Jacquenet, “Network Digital Twin: Concepts and Reference Architecture,” Jul. 2025, version Number: draft-irtf-nmrg-network-digital-twin-arch-11.
- [8] *Management and orchestration; Management aspects of Network Digital Twins*, 3GPP Std. TS 28.561, Oct. 2025, release 19.
- [9] X. Lin, L. Kundu, C. Dick, E. Obiodu, T. Mostak, and M. Flaxman, “6G Digital Twin Networks: From Theory to Practice,” *IEEE Communications Magazine*, vol. 61, no. 11, pp. 72–78, Nov. 2023.
- [10] D. Campoy, I. Vilà, J. Pérez-Romero, O. Sallent, and X. Gelabert, “Digital Twin as a Service for 6G Radio Access Networks: Functional Model and Key Challenges,” in *Proceedings of the 31st International Conference on Telecommunications (ICT)*, Apr. 2025.
- [11] D.-H. Tran, N. Waheed, Y. M. Saputra, X. Lin, C. T. Nguyen, T. S. Abdu, V. N. Vo, V.-Q. Pham, M. Alsenwi, A. B. M. A. Adam, S. Chatzinotas, E. Lagunas, H. Tran, T. H. Dac, and N. V. Huynh, “Network Digital Twin for 6G and Beyond: An End-to-End View Across Multi-Domain Network Ecosystems,” *IEEE Open Journal of the Communications Society*, vol. 6, pp. 6866–6911, 2025.
- [12] A. Karamchandani, J. Nunez, L. de-la Cal, Y. Moreno, A. Mozo, and A. Pastor, “On the Applicability of Network Digital Twins in Generating Synthetic Data for Heavy Hitter Discrimination,” *IEEE Communications Magazine*, vol. 63, no. 10, pp. 78–84, Oct. 2025.
- [13] A. Calvillo-Fernandez, T. Dimitrovski, M. Groshev, A. Ganesh, C. Ayimba, and A. de la Oliva, “Attention to Virtualization: Making Network Digital Twins Aware of Network Slicing,” *IEEE Network*, vol. 39, no. 3, pp. 134–139, 2025.
- [14] Z. Xiao, T. Tao, C. Ye, Y. Hu, Y. Feng, T. Jiao, and L. Cai, “Towards native intelligence: 6g-llm trained with reinforcement learning from ndt feedback,” 2026. [Online]. Available: <https://arxiv.org/abs/2601.09992>
- [15] H. X. Nguyen, R. Trestian, D. To, and M. Tatipamula, “Digital Twin for 5G and Beyond,” *IEEE Commun. Mag.*, vol. 59, no. 2, pp. 10–15, Feb. 2021.
- [16] M. Sanz Rodrigo, D. Rivera, J. I. Moreno, M. Álvarez Campana, and D. R. López, “Digital Twins for 5G Networks: A Modeling and Deployment Methodology,” *IEEE Access*, vol. 11, pp. 38 112–38 126, 2023.
- [17] G. Nardini and G. Stea, “Enabling simulation services for digital twins of 5G/B5G mobile networks,” *Computer Communications*, vol. 213, pp. 33–48, Jan. 2024.
- [18] I. Vilà, O. Sallent, and J. Pérez-Romero, “On the Design of a Network Digital Twin for the Radio Access Network in 5G and Beyond,” *Sensors*, vol. 23, no. 3, p. 1197, Jan. 2023.
- [19] N. Apostolakis, L. E. Chatzieleftheriou, D. Bega, M. Gramaglia, and A. Banchs, “Digital Twins for Next-Generation Mobile Networks: Applications and Solutions,” *IEEE Commun. Mag.*, vol. 61, no. 11, pp. 80–86, Nov. 2023.
- [20] R. Vilalta, L. Gifre, R. Casellas, R. Muñoz, R. Martínez, A. Mozo, A. Pastor, D. López, and J. P. Fernández-Palacios, “Applying Digital Twins to Optical Networks with Cloud-Native SDN Controllers,” *IEEE Commun. Mag.*, vol. 61, no. 12, pp. 128–134, Dec. 2023.
- [21] R. Cherini, N. Gonzalez-Montoro, A. Gonzalez-Montoro, L. Compagnucci, and J. M. Finochietto, “Building operator-centric digital twins for optical transport networks: from data models to real applications,” *J. Opt. Commun. Netw.*, vol. 17, no. 6, p. B1, Jun. 2025.
- [22] M. Polverini, I. Germini, A. Cianfrani, F. G. Lavacca, and M. Listanti, “A Digital Twin based Framework to Enable “What-If” Analysis in BGP Optimization,” in *NOMS 2023-2023 IEEE/IFIP Network Operations and Management Symposium*. Miami, FL, USA: IEEE, May 2023, pp. 1–6.
- [23] X. Wang, C. Liu, J. He, N. Shi, T. Zhang, and X. Pan, “6G Network Architecture Based on Digital Twin: Modeling, Evaluation, and Optimization,” *IEEE Network*, vol. 38, no. 1, pp. 15–21, Jan. 2024.
- [24] S. Robitzsch, A. De La Oliva, C. F. Chiasserini, C. E. Casetti, F. Agraz, G. Landi, L. M. Contreras, P. Szilagy, P. G. Giardina, R. Rosales, S. Spadaro, and V. Frasca, “Standardisation Assessment of a Digital Twin-Based Multi-Domain Deterministic Communications System,” in *2024 IEEE Globecom Workshops (GC Wkshps)*. Cape Town, South Africa: IEEE, Dec. 2024, pp. 1–7.
- [25] P. G. Giardina, P. Szilagy, C. F. Chiasserini, J. L. Carcel, L. Velasco, S. Spadaro, F. Agraz, S. Robitzsch, R. Rosales, V. Frasca, R. Doostnejad, A. Calvillo, and G. Bernini, “A hierarchical AI-based control plane solution for multi-technology deterministic networks,” in *Proceedings of the Twenty-fourth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing*. Washington DC USA: ACM, Oct. 2023, pp. 316–321.
- [26] S. Faye, M. Camelo, J.-S. Sottet, C. Sommer, M. Franke, J. Baudouin, G. Castellanos, R. Decorme, M. P. Fanti, R. Fuladi, G. Kesik, B. Mendes, C. Murphy, S. Parker, S. Pryor, S. M. Senouci, and I. Turcanu, “Integrating Network Digital Twinning into Future AI-based 6G Systems: The 6G-TWIN Vision,” in *2024 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*. Antwerp, Belgium: IEEE, Jun. 2024, pp. 883–888.
- [27] A. Zaki-Hindi and others, “A Reference Functional Architecture for Network Digital Twins in 6G Systems,” *IEEE Open Journal of the Communications Society*, vol. 7, pp. 2068–2101, 2026.
- [28] S. M. Raza, R. Minerva, N. Crespi, M. Alvi, M. Herath, and H. Dutta, “A comprehensive survey of Network Digital Twin architecture, capabilities, challenges, and requirements for Edge-Cloud Continuum,” *Computer Communications*, vol. 236, p. 108144, Apr. 2025.
- [29] European Commission, “UNITY6G - UNified archITecture for Open RAN-enabled Distributed, Scalable and Sustainability-enhanced 6G Networks (UNITY6G),” 2025. [Online]. Available: <https://unity-6g.eu/>
- [30] 3GPP, “Management and Orchestration; Architecture framework (Release 19),” 3rd Generation Partnership Project (3GPP), Tech. Rep., 2025.
- [31] W. Tarneberg, E. Fitzgerald, M. Bhuyan, P. Townend, K.-E. Arzen, P.-O. Ostberg, E. Elmroth, J. Eker, F. Tufvesson, and M. Kihl, “The 6g computing continuum (6gcc): Meeting the 6g computing challenges,” in *2022 1st International Conference on 6G Networking (6GNet)*. IEEE, 2022, pp. 1–8.
- [32] ETSI Software Development Group TeraFlowSDN, “TeraFlowSDN: Open source cloud-native sdn controller and network automation framework,” ETSI SDG TFS Website, 2026, accessed: 2026-05-03. [Online]. Available: <https://www.etsi.org/technical-groups/tfs/>
- [33] G. F. Riley and T. R. Henderson, “The ns-3 network simulator,” in *Modeling and Tools for Network Simulation*, K. Wehrle, M. Güneş, and J. Gross, Eds. Berlin, Heidelberg: Springer, 2010, pp. 15–34.
- [34] “Bts search – base station database,” <https://beta.btsearch.pl/>, accessed: 2026.
- [35] P. Kryszkiewicz, Ł. Kułacz, M. Pakula, M. Dryjanski, M. Hoffmann, P. Skrzypczak, H. Lehmann, and M. Stahn, “Large-scale wireless network management via open-ran tandem apps: Cell on/off switching use case,” *arXiv preprint arXiv:2604.24465*, 2026.
- [36] S. Pfenninger and I. Staffell, “Long-term patterns of european pv output using 30 years of validated hourly reanalysis and satellite data,” *Energy*, vol. 114, pp. 1251–1265, 2016.
- [37] I. Staffell and S. Pfenninger, “Using bias-corrected reanalysis to simulate current and future wind power output,” *Energy*, vol. 114, pp. 1224–1239, 2016.