

Scaling Physics-to-Protocol Digital Twins for 5G/6G Networks

Salim El Ghalbzouri^{*†}, Daoud Burghal[†], Russell Ford[‡], Pranav Madadi[†], Ardavan Tehrani[†]
Charlie Jianzhong Zhang[†], and Adlen Ksentini^{*}

^{*}Communication Systems Department, EURECOM, France

[†]Standards and Mobility Innovation, Samsung Research America, USA

[‡]NVIDIA, USA

Abstract—Digital Twins (DTs) are central to the 6G vision, serving as the definitive proving ground for pervasive network intelligence. The ultimate value of a DT hinges on closing the sim-to-real gap, which requires faithfully modeling dynamic, dense multi-gNB (next-generation Node B)/multi-UE (User Equipment) behaviors. However, current platforms can emulate high-fidelity but isolated slices of the protocol stack; scaling this fidelity horizontally across dense environments encounters prohibitive computational bottlenecks. This creates a binary trade-off: scale broadly using abstracted simulators that sacrifice protocol realism, or achieve genuine protocol execution within the narrow bounds of costly hardware testbeds. This paper introduces a Physics-Based Digital Twin (PBDT) architecture that breaks this deadlock. We deploy a centralized proxy that interconnects fully 3GPP-compliant OpenAirInterface (OAI) gNB and UE stack instances, replacing per-node baseband processing with slot-level message routing. The proxy conditions each gNB–UE link with channel metrics (SINR, RSS, RSRP) derived from site-specific ray tracing, so every node perceives a radio environment shaped by 3-D urban geometry, enabling direct MAC-to-MAC communication at deployment scale on commodity infrastructure. Experimental validation confirms the physics fidelity of the ray-traced channel conditioning along a vehicular mobility trace and demonstrates a $3.4\times$ per-slot cost reduction over baseband-coupled baselines, extending the feasible real-time operating range by $3\times$ on commodity hardware. The platform provides a native environment for training, validating, and stress-testing AI-driven RAN optimization against real 3GPP protocol behavior, closing the loop between algorithm development and deployment-grade network dynamics.

Index Terms—Digital Twin, OpenAirInterface, Ray tracing, nFAPI, Scalability.

I. INTRODUCTION

Autonomous network intelligence requires training environments that reproduce the protocol-state trajectories and radio conditions it will encounter in production, under load, mobility, and contention. The 6G roadmap positions AI across the Radio Access Network (RAN) and designates the Digital Twin (DT) as the canonical environment in which these controllers are developed, validated, and stress-tested before deployment [1], [2]. These dynamics emerge from the concurrent operation of many gNBs and UEs under correlated spatial propagation.

The scalability ceiling of existing platforms stems from a fundamental architectural trade-off between protocol fidelity and deployment scale. Simulators achieve topological density by abstracting protocol execution away from 3GPP-compliant

software, a trade-off that enables scale but forfeits the procedural realism. Production-grade open-source stacks eliminate this compromise, but the real-time constraint of symbol-level baseband processing, where every scheduled subframe must be encoded, transmitted, and decoded within strict timing deadlines, binds per-instance compute to the physical layer and prevents horizontal scaling. Facility-scale testbeds resolve this through dedicated RF infrastructure, an investment profile that places this capability beyond the reach of the broader research community. Orthogonal to the scaling challenge, the channel models underlying these platforms limit spatial realism: statistical models approximate propagation through parametric distributions, discarding the site-specific geometry that shapes real deployments. Ray tracing addresses this by deriving channel conditions directly from the physical environment. Yet when ray-traced channels are integrated with protocol stacks at the physical layer, they inherit the full computational weight of baseband processing while adding GPU-intensive channel computation on top. Spatial realism and horizontal scalability remain structurally incompatible.

This paper eliminates this trade-off. We present a Physics-Based Digital Twin (PBDT)¹ that deploys a centralized proxy acting as a virtual physical layer for all gNB and UE instances in the deployment, fundamentally eliminating the need for per-instance baseband processing, which is computationally prohibitive. By replacing this *complex and expensive* per-node baseband processing with lightweight slot-level message routing, our approach establishes direct MAC-to-MAC communication across the network while maintaining full 3GPP compliance. The proxy employs a Block Error Rate (BLER)-based PHY abstraction, conditioned on site-specific channel metrics (SINR, RSS, RSRP) derived from ray tracing over real 3D environmental geometry, to determine link-level behavior for each gNB–UE pair. The PHY abstraction captures the channel at the granularity that drives higher-layer protocol decisions: whether a transport block is successfully decoded, at what Modulation and Coding Scheme (MCS), and with what Hybrid Automatic Repeat Request (HARQ) feedback. This architectural transformation enables *horizontal scalability* by replacing the entire sub-symbol waveform processing chain, leaving the full protocol stack driven by deterministic

¹We release the PBDT implementation as open source. *Archive:* <https://doi.org/10.5281/zenodo.21915631>. *Source:* <https://gitlab.eurecom.fr/oai/openairinterface5g/-/tree/l2-sim-new>.

spatial physics at a per-instance cost that scales with *upper-layer* processing rather than baseband computation.

Dense multi-gNB/multi-UE topologies that today demand facility-grade testbed infrastructure become reproducible on standard research infrastructure, democratizing access to deployment-scale protocol experimentation. The same architectural property that enables this scale produces a deterministic execution environment: AI/ML agents can train against real protocol behavior at deployment density in closed loop, evaluate and iterate on the protocol-level impact of their actions. The platform delivers what the autonomous, AI-native networks envisioned for 6G require: 3GPP-compliant protocol execution at network scale, driven by deterministic spatial physics.

The core contributions of this work are:

- An architecture that decouples protocol-stack scaling from baseband processing overhead, enabling deployment-scale populations of independent, 3GPP-compliant gNB and UE protocol stacks to execute under physics-derived channel conditions through a virtual PHY with link-level abstraction.
- A channel-conditioning methodology that integrates site-specific, ray-traced propagation directly at the PHY abstraction interface, so that real schedulers and mobility procedures respond to environmental geometry without waveform synthesis.
- Experimental validation demonstrating that proxy-mediated multi-gNB/multi-UE topologies produce spatially coherent channel dynamics under urban vehicular mobility, with a $3.4\times$ per-slot cost reduction and $3\times$ increase in feasible real-time operating range over baseband-coupled baselines.

The remainder of this paper is organized as follows. Section II positions the work within the systems landscape. Section III presents the architecture and channel-integration methodology. Section IV reports experimental results, and Section V concludes.

II. RELATED WORK

The systems landscape for 5G digital twinning is organized by a structural variable: where the channel model meets the protocol stack, and at what level of fidelity and scalability each is represented. This boundary determines which properties a platform can offer and which it precludes. We survey the limits of existing architectures through the paradigms that define the design space our work occupies.

Discrete-event simulators such as ns-3 [3] achieve network-scale topologies by implementing behavioral models of 3GPP procedures rather than executing 3GPP-compliant protocol software, and DTs built on these simulators inherit the gap. Recent work has substantially improved the channel-modeling half of this paradigm. SioLENA [4] couples Sionna RT [5] into the ns-3 NR module for site-specific channel evaluation, replacing statistical 3GPP propagation with geometry-aware ray-traced channels. These integrations meaningfully advance spatial realism within the simulator paradigm. However, the

TABLE I
CAPABILITY COMPARISON OF 5G/6G DT PLATFORMS

	SioLENA [4]	Colosseum [8]	OWDT [7]	PBDT
Physics fidelity	✓	×	✓	✓
Protocol fidelity	×	✓	✓	✓
Multi-gNB/multi-UE scalability	✓	✓	×	✓
Commodity hardware	✓	×	✓	✓

protocol stacks that consume the ray-traced channels remain behavioral models. AI policies trained in these environments face a protocol-level gap that compromises their portability to production deployments.

Production-grade stacks such as OpenAirInterface (OAI) [6] eliminate this abstraction entirely, executing 3GPP-compliant protocol software. The Open Wireless Digital Twin (OWDT) [7] integrates Sionna RT with OAI through the `rfsimulator`, OAI's built-in RF emulator that transports I/Q samples over TCP between gNB and UE processes, applying pre-computed ray-traced channel impulse responses to baseband I/Q samples. This achieves both site-specific channel fidelity and full protocol execution. However, at this coupling point, per-link baseband processing must still execute within real-time constraints, a cost inherited from OAI's physical layer that scales with every additional gNB-UE link and prevents horizontal scaling.

Facility-scale testbeds such as Colosseum [8] absorb this per-instance cost through dedicated infrastructure, connecting software-defined radio nodes at scale through channel emulators to sustain real-time multi-gNB experiments with production protocol stacks. The infrastructure cost of this approach trades software accessibility for hardware-dependent scale.

The architecture presented in this paper replaces per-node waveform processing with slot-level message routing through a centralized proxy built on OAI, injecting physics-derived channel metrics from Sionna RT into a PHY abstraction layer. This yields site-specific channel fidelity, production protocol execution, and horizontal multi-gNB/multi-UE scalability on commodity hardware. Table I summarizes the capability landscape.

III. SYSTEM ARCHITECTURE

A. Architecture Overview

PBDT integrates two complementary planes (Fig. 1): a *protocol plane* and a *physics plane*. The protocol plane executes N_{gNB} independent 3GPP-compliant gNB stacks and N_{UE} UE stacks interconnected through a centralized Layer-2 (L2) network Functional Application Platform Interface (nFAPI) proxy, with a shared 5G core network. Each gNB instance runs the OAI NR protocol stack (RRC, PDCP, SDAP, RLC, and the MAC scheduler) as an independent Virtual Network Function (VNF) process. Each UE instance executes the OAI NR UE stack from NAS and RRC downward through MAC, with the physical layer replaced by a PHY abstraction model




Fig. 1. PBDT (Physics-Based Digital Twin) architecture.

(Section III-B). This decomposition follows the nFAPI functional split [9], where the control-plane (P5) interface handles per-gNB configuration and the data-plane (P7) interface carries slot-level scheduling messages between each gNB-VNF and the centralized Physical Network Function (PNF). In our architecture, the PNF is a software proxy shared across all gNB-VNFs, eliminating per-node baseband computation by exploiting the nFAPI boundary between MAC and PHY in the OAI stack. Each gNB-VNF additionally exposes an E2 agent interface, enabling integration with O-RAN near-RT RIC frameworks such as FlexRIC [10] for closed-loop policy evaluation.

The physics plane derives per-gNB path gains along UE trajectories from site-specific ray tracing (Section III-C). The resulting per-UE, per-gNB channel traces are loaded by the proxy and delivered at slot granularity.

B. L2 nFAPI Proxy

The proxy builds on the open-source EpiSci proxy [11], enhanced, integrated with the physics plane, and extended to multi-gNB operation. It operates at the nFAPI boundary, maintaining a P7 connection per gNB-VNF and one dedicated bidirectional UDP channel per UE.

At each slot boundary ($T_{\text{slot}} = 500 \mu\text{s}$ for NR numerology $\mu=1$, 30 kHz subcarrier spacing), the proxy generates a `SLOT.indication` that synchronizes all gNB-VNFs to a frame and slot counter and carries per-UE channel information from the current trace entry. Each gNB-VNF then emits a `DL_TTI.request` carrying per-UE PDSCH descriptors and a `TX_DATA.request` with transport block payloads. The proxy filters each downlink message by RNTI and forwards the PDUs destined for each UE.

In the uplink direction, each UE returns `UCI.indication` (HARQ ACK/NACK, CQI, scheduling request) and `CRC.indication` with transport block decode status. The proxy routes these indications to the corresponding gNB-VNF and injects the CQI derived from the current trace entry, so that the gNB's link-adaptation module operates on physics-consistent channel quality at every slot.

On the UE side, the physical layer is replaced by a PHY abstraction that maps the proxy-injected SINR to a BLER via pre-computed AWGN lookup tables covering each supported MCS index. For each transport block, the CRC outcome is determined by a Bernoulli trial with failure probability equal to the BLER: if the block fails, the UE generates a HARQ NACK; otherwise, the transport block is delivered upward through MAC, RLC, PDCP, and SDAP to the application. In the reverse direction, each trace entry's SINR is mapped to a CQI index via a monotone lookup table calibrated at a target BLER of 10%, following the CQI tables (TS 38.214 [12], Table 5.2.2.1-2). The resulting CQI is injected into UCI at slot granularity, closing the link-adaptation loop: the gNB's MAC observes physics-derived CQI, selects MCS accordingly, and the PHY abstraction evaluates BLER at the resulting operating point. This abstraction captures the channel at the granularity that drives protocol-layer decisions: whether a transport block is decoded, at what MCS, and with what HARQ outcome, while eliminating the entire sub-symbol signal processing chain.

C. Physics Plane

Sionna RT's [5] ray-tracing engine computes site-specific propagation over a 3-D scene of the target deployment area. Let s index the deployed gNBs and u index the UEs. For

a transmitter gNB_s and a receiver position $\mathbf{r} = (x, y, z)$, the engine launches rays into the scene and traces their interactions with the 3-D geometry, resolving line-of-sight propagation, specular reflections, diffraction, and scattering. Each resolved path carries a complex amplitude $a_{s,i}$. The per-gNB path gain is obtained by coherently summing all path contributions:

$$G_s(\mathbf{r}) = \left| \sum_{i=1}^{N_{\text{paths}}} a_{s,i}(\mathbf{r}) \right|^2 \quad (1)$$

The received signal strength follows as

$$\text{RSS}_{u,s}(t) = P_{\text{tx},s} \cdot G_s(\mathbf{r}_u(t)) \quad (2)$$

where $P_{\text{tx},s}$ is the configured transmit power of gNB_s and $\mathbf{r}_u(t)$ is the position of UE_u at time step t along its mobility trajectory. For multi-gNB deployments, the SINR at UE_u served by gNB_{s_0} is computed by:

$$\text{SINR}_{u,s_0}(t) = \frac{\text{RSS}_{u,s_0}(t)}{N_0 + \sum_{s \neq s_0} \text{RSS}_{u,s}(t)} \quad (3)$$

where N_0 is the thermal noise power. RSRP is derived from RSS following TS 38.215 [13]:

$$\text{RSRP}_{u,s}(t) = \text{RSS}_{u,s}(t) \Big|_{\text{dBm}} - 10 \log_{10}(12 \cdot N_{\text{PRB}}) \quad (4)$$

where N_{PRB} is the number of PRBs determined by the configured bandwidth and subcarrier spacing.

This two-plane decomposition decouples the cost of protocol execution from the cost of channel computation. Scaling the network requires only additional protocol processes, without replicating the baseband processing chain that conventionally binds each node.

IV. EXPERIMENTAL EVALUATION

This section validates the architecture along two axes: spatial fidelity of the physics-plane channel conditioning, and computational scalability as network density increases.

A. Physics-Plane Validation

Four gNBs are deployed over the Munich scene, with per-gNB transmit power of 44–46 dBm at 3.5 GHz (40 MHz, 106 PRBs, 30 kHz SCS). Sionna RT computes per-gNB path gain across a spatial grid at street level, resolving LoS propagation, specular reflections, and diffraction through the 3-D building mesh.

Fig. 2 presents the best-server path gain across the deployment area. Each gNB's coverage footprint is shaped by the surrounding geometry, with path gain spanning from -60 dB in near-field LoS to below -120 dB in deep NLoS, values consistent with standard 3.5 GHz urban macro measurements.

A vehicular UE (50 km/h) traverses a road-following trajectory through the deployment, passing through three successive serving-gNB regions (Fig. 3). Per-gNB RSRP is derived from the ray-traced path gain following Eq. (4) and L3-filtered per




Fig. 2. Best-server path gain across the Munich deployment area.




Fig. 3. Vehicular UE trajectory (50 km/h).

TS 38.331 [14] to suppress fast-fading fluctuations. Serving-gNB selection applies 3 dB hysteresis: a neighbor gNB must sustain a filtered RSRP advantage of at least 3 dB over the current serving gNB before re-selection occurs, preventing rapid ping-pong transitions at coverage boundaries.

Fig. 4 presents the resulting RSRP traces. The UE begins near gNB_3 with serving RSRP of approximately -62 dBm, which decays to -74 dBm as the UE moves away. After the first serving transition, gNB_0 provides -68 to -85 dBm as the UE traverses the central urban area. The second transition hands the UE to gNB_1 , whose RSRP rises from -80 dBm to -65 dBm as the UE approaches. gNB_2 remains consistently weak throughout, positioned in NLoS relative to the entire trajectory. The lower heatmap confirms that all four per-gNB RSRP profiles evolve with the ray-traced geometry.

These per-gNB, per-slot channel metrics are the values the




Fig. 4. Per-gNB RSRP traces along the vehicular trajectory.

TABLE II
MEASUREMENT PARAMETERS

Parameter	Value
Carrier frequency	3.5 GHz (band n78)
System bandwidth	40 MHz
Subcarrier spacing	30 kHz ($\mu = 1$)
Number of PRBs	106
Slot duration	500 μ s
UE counts	2, 4, 8, 16, 32, 64
Platform	Intel Core Ultra 5 135U, 12 cores / 14 threads, 32 GB

proxy injects into each UE's channel state (Section III-B), driving CQI reporting and MCS selection across independent gNB protocol-stack instances.

B. Computational Scalability

This section quantifies the per-slot processing cost of the PBDT proxy relative to OAI's I/Q baseband pipeline under the *rf*simulator, which OWDT [7] relies on for waveform-level channel coupling and thus inherits full baseband processing overhead. The comparison uses identical upper-layer configurations; deployment and measurement parameters are listed in Table II. Both architectures run on a single commodity workstation, and absolute values will shift with CPU frequency, core count, and memory bandwidth; the comparison isolates the structural cost difference at a fixed hardware point.

The I/Q baseline executes OAI's complete downlink L1 chain (DLSCH encoding, scrambling, modulation, precoding, and OFDM modulation) with sample-level transport between gNB and UE processes over TCP. The proxy replaces this chain with slot-level nFAPI message routing at the MAC-PHY boundary. For each configuration, per-slot processing time is collected over approximately 100s of steady-state operation after initial UE attachment.

Fig. 5 decomposes the per-slot cost at a single UE. Under the *rf*simulator, the gNB's per-slot L1 pipeline consumes 154 μ s, distributed across DL encoding (52 μ s), OFDM modulation (46 μ s), OFDM demodulation (48 μ s), UL decoding (5 μ s), and control overhead (3 μ s). The proxy replaces this with three nFAPI message-routing operations: DL forwarding (21 μ s), slot indication dispatch (6 μ s), and UL message aggregation (18 μ s), totaling 45 μ s, a 3.4 $\times$ reduction.




Fig. 5. Per-slot CPU-work breakdown at a single UE.




Fig. 6. Per-slot processing time vs. number of UEs.

This reduction is structural: it follows from replacing per-PRB waveform operations with per-UE message copies, and persists regardless of hardware. The proxy consumes 9% of the 500 μ s slot budget; the baseline consumes 31%.

Fig. 6 presents empirical per-slot processing time from 2 to 64 UEs. The proxy grows from 49 μ s to 442 μ s, remaining within the 500 μ s real-time deadline across the full range, with a consistent marginal cost of approximately 4–6 μ s per additional UE, since each UE adds one per-RNTI message copy. The *rf*simulator grows from 123 μ s to 286 μ s through 16 UEs, then accelerates sharply, reaching 790 μ s at 32 UEs, 1149 μ s at 64 UEs, as per-UE LDPC encoding and I/Q sample buffering saturate the serial L1 pipeline. Once the *rf*simulator exceeds the slot deadline, the system remains operational but can no longer sustain the NR slot grid: frames are dropped, the TCP-based I/Q transport becomes I/O-bound, and real-time scheduling fidelity degrades. On this hardware, the proxy extends the feasible real-time operation beyond 64 UEs whereas the baseline exceeds the slot deadline at approximately 23 UEs.

Multi-gNB/multi-UE topologies, where use cases such as load balancing and mobility management arise, are the deployment scenarios that AI-driven RAN controllers must train




Fig. 7. Per-slot processing time over the (UEs $\times$ gNBs) plane.

against. Our architecture enables these configurations without dedicated RF infrastructure, placing them within reach of the research community with access to a standard workstation. The proxy processes each gNB's slot messages sequentially within a single timing loop; Fig. 7 maps the resulting per-slot cost over the (UEs $\times$ gNBs) plane, where the bottom row (1 gNB) contains the empirical measurements from Fig. 6 and each subsequent row is computed analytically. The feasible operating region scales directly with the available CPU frequency, core count, and memory bandwidth; the heatmap captures the structural cost profile at a fixed hardware point, and higher-end platforms shift the real-time boundary outward proportionally.

V. CONCLUSION AND FUTURE WORK

This paper presented an architecture that decouples protocol-stack scaling from baseband processing, enabling dense multi-gNB/multi-UE DT deployments on commodity infrastructure. By replacing per-node waveform processing with slot-level message routing at the nFAPI MAC-PHY boundary and conditioning each link with physics-derived channel metrics from site-specific ray tracing, the platform establishes direct MAC-to-MAC communication between production 3GPP protocol stacks under deterministic, geometry-aware channel conditions. Experimental validation confirmed two results: ray-traced propagation over an urban deployment produces spatially coherent per-gNB channel conditions that evolve consistently with urban geometry along a vehicular mobility trace, and the PBDT achieves a $3.4\times$ per-slot cost reduction over baseband-coupled baselines, extending feasible real-time operation from approximately 23 to beyond 64 UEs.

By closing the sim-to-real gap, the PBDT provides a faithful training ground for autonomous network intelligence. Each gNB in the platform exposes standard E2 interfaces, enabling xApps to interact with production protocol stacks through the O-RAN Near-RT RIC at deployment density.

Beyond those workflows, the platform's scalable and physics-grounded execution environment is well suited to emerging agentic AI, where autonomous agents explore, learn, and adapt network policies against production protocol behavior in closed loop.

Building on this foundation, we are currently extending verification from the physics plane to the protocol plane, characterizing mobility, handover, and scheduling behavior across the denser multi-gNB/multi-UE deployments mapped in Fig. 7, and evaluating the xApp and rApp policies that drive them, which in turn provides the substrate for the agentic workflows outlined above.

ACKNOWLEDGMENT

The work of the EURECOM authors was funded by PEPR 5G NF-NAI Project and EU SNS JU 6G-FLECON (Grant No. 101192462).

REFERENCES

- [1] L. U. Khan, W. Saad, D. Niyato, Z. Han, and C. S. Hong, "Digital-twin-enabled 6G: Vision, architectural trends, and future directions," *IEEE Communications Magazine*, vol. 60, no. 1, pp. 74–80, 2022.
- [2] Z. Tao, W. Xu, Y. Huang, X. Wang, and X. You, "Wireless network digital twin for 6G: Generative AI as a key enabler," *IEEE Wireless Communications*, vol. 31, no. 4, pp. 24–31, 2024.
- [3] N. Patriciello, S. Lagen, B. Bojovic, and L. Giupponi, "An E2E simulator for 5G NR networks," *Simulation Modelling Practice and Theory*, vol. 96, p. 101933, 2019.
- [4] T. Ropitault, M. Bordin, P. Testolina, M. Polese, P. Johari, N. Golmie, and T. Melodia, "Enabling site-specific cellular network simulation through ray-tracing-driven ns-3," in *Proc. IEEE Consumer Communications and Networking Conference (CCNC)*, 2026, pp. 1–6.
- [5] J. Hoydis, F. Ait Aoudia, S. Cammerer, M. Nimier-David, N. Binder, G. Marcus, and A. Keller, "Sionna RT: Differentiable ray tracing for radio propagation modeling," in *Proc. IEEE Globecom Workshops (GC Wkshps)*, 2023, pp. 317–321.
- [6] F. Kaltenberger, A. P. Silva, A. Gosain, L. Wang, and T.-T. Nguyen, "OpenAirInterface: Democratizing innovation in the 5G era," *Computer Networks*, vol. 176, p. 107284, 2020.
- [7] T. Iye, M. Sakamoto, S. Takaya, E. Sato, Y. Susukida, Y. Nagaoka, K. Maruta, and J. Nakazato, "Open wireless digital twin: End-to-end 5G mobility emulation with OpenAirInterface and ray tracing," *IEEE Access*, vol. 13, pp. 175 109–175 122, 2025.
- [8] L. Bonati, P. Johari, M. Polese, S. D'Oro, S. Mohanti, M. Tehrani-Moayyed, D. Villa, S. Shrivastava, C. Tassie, K. Yoder, A. Gosain, and T. Melodia, "Colosseum: Large-scale wireless experimentation through hardware-in-the-loop network emulation," in *Proc. IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, 2021, pp. 105–113.
- [9] Small Cell Forum (SCF), "5G FAPI: PHY API specification," SCF222, v4.0, 2021, <https://www.smallcellforum.org>.
- [10] R. Schmidt, M. Irazabal, and N. Nikaein, "FlexRIC: An SDK for next-generation SD-RANs," in *Proc. ACM International Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, 2021, pp. 411–425.
- [11] EpiSci, "OAI LTE/5G Multi-UE Proxy," <https://github.com/EpiSci/oai-lte-5g-multi-ue-proxy>.
- [12] 3GPP, "TS 38.214: NR; physical layer procedures for data," 3rd Generation Partnership Project, Technical Specification 38.214, 2024, release 18.
- [13] —, "TS 38.215: NR; physical layer measurements," 3rd Generation Partnership Project, Technical Specification 38.215, 2024, release 18.
- [14] —, "TS 38.331: NR; radio resource control (RRC); protocol specification," 3rd Generation Partnership Project, Technical Specification 38.331, 2024, release 18.