

Campus5G: A Campus Scale Private 5G Open RAN Testbed

Andrew E. Ferguson*
 The University of Edinburgh
 United Kingdom
 Andrew.E.Ferguson@ed.ac.uk

Tianxin Wang
 The University of Edinburgh
 United Kingdom
 twang4@ed.ac.uk

Ujjwal Pawar*
 The University of Edinburgh
 United Kingdom
 u.pawar@sms.ed.ac.uk

Mahesh K. Marina
 The University of Edinburgh
 United Kingdom
 mahesh@ed.ac.uk

ABSTRACT

Mobile networks are embracing disaggregation, reflected by the industry trend towards Open RAN. Private 5G networks are viewed as particularly suitable contenders as early adopters of Open RAN, owing to their setting, high degree of control, and opportunity for innovation they present. Motivated by this, we have recently deployed CAMPUS5G, the first of its kind campus-wide, O-RAN-compliant private 5G testbed across the central campus of the University of Edinburgh. We present in detail our process developing the testbed, from planning, to architecting, to deployment, and measuring the testbed performance. We then discuss the lessons learned from building the testbed, and highlight some research opportunities that emerged from our deployment experience.

CCS CONCEPTS

• Networks → Mobile networks; Network experimentation;

KEYWORDS

Open RAN, Private 5G, Testbeds

1 INTRODUCTION

Mobile networks in the recent past have embraced *disaggregation* at different levels, driven by the need for cost-effective and flexible system architecture that accelerates innovation and enables new services. Complimentarily, they are headed towards data-driven and AI powered operation to unlock significant energy, operational and spectral efficiency gains. These trends are most prominently evident with *Open Radio Access Network (RAN)* [84], being standardized by the O-RAN Alliance [61].

The O-RAN architecture [12] disaggregates the traditional monolithic base stations into Centralized (CU), Distributed Unit (DU) and Radio Unit (RU) components that communicate over open standardized interfaces. Of these components, CU and DU are intended to be realized as virtualized network functions (VNFs) over commodity compute hardware, thus reduce CAPEX for operators. This disaggregated architecture also promises to diversify the mobile telecoms ecosystem, as it offers flexibility for operators to mix and match solutions from different vendors. In addition, the O-RAN architecture features RAN Intelligent Controllers (RICs) operating at different timescales and with corresponding set of “Apps” (for energy saving, load balancing, anomaly detection, etc.) that can leverage AI/ML towards efficient RAN operation.

While Open RAN is expected to become an integral part of future (6G) mobile network system architecture generally [70], its deployment in national scale public mobile networks today is mostly limited to trials [8, 60]. This could be partly attributed to Open RAN not making it to the last investment cycle for public mobile network operators (MNOs) when they transitioned to 5G [16].

It is in this context that *private 5G mobile networks* [68] are seen as the early adopters of Open RAN. Unlike national scale public mobile networks, private networks are “local” scale deployments that target industrial (e.g., manufacturing, mining, warehouses) and enterprise (e.g., business parks, universities, hospitals) settings, transport hubs, venues and such [42]. Private mobile networks are appealing for these environments because they not only allow high degree of control and customization but also offer greater opportunity for innovation (e.g., edge AI applications [9, 29, 52]). Private 5G deployments are rapidly growing [4, 6, 10, 11, 44, 63]. Private 5G is a key use case for Open RAN [48]. New breed of operators and vendors target private 5G with Open RAN as a default [1, 2, 14, 46].

University campuses are a representative and compelling setting for Open RAN based private 5G networks. Moreover, as history has shown with software-defined networking (SDN), campus networks can be excellent experimentation grounds for innovative technological solutions and services that have the potential to shape our future connectivity [51].

Motivated by the above, we have recently deployed CAMPUS5G, the first of its kind campus-wide, O-RAN-compliant private 5G testbed across the central campus of the University of Edinburgh. Our deployment setting is a typical dense urban outdoor environment with high footfall and mobile traffic demand. We designed CAMPUS5G with blanket coverage across the campus in mind while at the same time allowing ample flexibility for experimentally evaluating innovative designs across the whole O-RAN architecture in a real-world setting. This includes the ability to use commercial and open source software components from different vendors in the Open RAN ecosystem as well as research prototypes. We also embed AI compute to facilitate research on AI for RAN and edge AI applications. From a user perspective, our testbed allows the use of unmodified commodity mobile devices through custom physical or eSIMs for our network.

CAMPUS5G is intended to enable broad ranging experimental research on next generation wireless access networks, including experimentally exploring systems research challenges in the Open

*Equal contributors to this article.

Table 1: CAMPUS5G versus representative set of existing private 5G Open RAN testbeds.

Testbed	Size	Indoor or Outdoor	O-RAN Compatibility	Commercial RUs or SDRs	Software Flexibility
Arena [18]	24	Indoor	Partially	SDRs	Open Source
Powder [20]	8	Outdoor	Partially	SDRs	Open Source
Colosseum [83]	128	Emulated	Partially	SDRs	Open Source
X5G [82]	8	Indoor	Yes	Commercial RUs	Open Source or Commercial
Microsoft Cambridge [17]	20	Indoor	Yes	Commercial RUs	Open Source or Commercial
SIT-FCT [71]	5	Indoor and Outdoor	Yes	Commercial RUs	Commercial
OpenRAN@Brazil [50]	8	Indoor and Outdoor	Yes	Commercial RUs	Commercial
CAMPUS5G	20	Outdoor	Yes	Commercial RUs	Open Source or Commercial

RAN context; developing and validating innovative data-driven solutions for optimized RAN operation and beyond; and investigating issues unique to private 5G networks and their role in the future mobile network system architecture. In the longer term, we envision this testbed to become a living lab routinely used by university staff and students.

In this paper, we present our end to end process with the CAMPUS5G testbed deployment from the planning stage and architecture design to component selection, configuration and measurement based coverage analysis. Compared to existing private 5G or Open RAN testbeds (discussed in §2 and summarized in Table 1), our testbed is unique not only in its campus level scale, high degree of flexibility and fully O-RAN compliant nature but also due to its deployment in an urban outdoor environment. We share some of the key lessons we learned from the testbed deployment as well as potential research opportunities we identified in the process. We hope this paper can serve as a blueprint to inform other researchers intending to deploy similar testbeds. Equally, we believe our experience reported in this paper will benefit Open RAN research and experimentation in the context of private 5G and beyond.

2 RELATED WORK

By way of positioning our CAMPUS5G testbed with respect to existing private 5G and Open RAN testbeds, here we briefly discuss those prior efforts. There exist a few private 5G deployments made up of commercial components that are being used as testbeds. Examples include NICT Local 5G testbed in Japan [49], South Bend CBRS network in the US [81] and Fraunhofer IIS Industrial Testbed in Germany [34]. Some of these like NICT and South Bend testbeds are not O-RAN compatible while others like the Fraunhofer testbed are primarily meant for interoperability testing and so offer very limited flexibility for research use.

Many testbeds currently exist to support O-RAN research. These testbeds by their nature are private (4G/5G) networks. Majority of them rely on USRP software defined radio (SDR) based radio frontends for the RUs, and are for the most part deployed indoors. Arena testbed [18] is an early example. NEC testbed [72] is another. NSF PAWR platforms [28] also broadly fall under this category, though each with a different focus (e.g., mmWave in COSMOS [22], UAVs for AERPAW [55], rural applications for ARA [40]). Of these, Powder [20] features an 8-node outdoor testbed with SDRs deployed across the campus of University of Utah. Similar SDR based Open RAN testbeds exist outside the US, including the Open Ireland testbed [3], testbeds in Romania and Norway [48] and various testbeds under Fed4Fire and follow on projects [23]. Note that

these SDR based testbeds do not fully support the O-RAN fronthaul interface between DU and RU based on the 3GPP 7.2x split that puts High-PHY on the DU side and Low-PHY on the RU side [69]. As such, they are only partially compliant with the O-RAN standard [48] and consequently also are not compatible with O-RAN compliant commercial virtualized RAN (vRAN) solutions. Colosseum facility at Northeastern University [83] is also based on SDRs but relies instead on hardware based channel emulation, and so can support experimentation based on replaying real-world channel measurement traces.

In contrast to the above, other testbeds have emerged recently that use commercial RUs and are fully O-RAN compliant. The testbed at Microsoft Research, Cambridge [17] is a building scale indoor deployment with 4 O-RAN compliant Foxconn RPQN-7800 RUs per floor across 5 floors, and uses an Intel FlexRAN based Layer 1 (L1), i.e., PHY [38]. X5G testbed at Northeastern is also an indoor testbed that consists of 8 Foxconn RPQN7801 RUs but instead uses NVIDIA ARC based GPU accelerated L1. Recently, a nation wide Open RAN testbed deployment effort was reported from Brazil [50]. While this testbed is large scale in its geographical span covering multiple interconnected sites across Brazil, each site has a small scale deployment with around handful of commercial RUs (from Foxconn, Benetel and VVDN), mostly indoor. The FCT O-RAN testbed at Singapore Institute of Technology [71] mixes indoor (3 Foxconn RPQN780) and outdoor (2 Benetel RAN650) RUs with a FlexRAN L1 and Radisys CU / DU. Our CAMPUS5G, while broadly in the same category as the above three testbeds, is deployed outdoors at a campus scale in a dense urban outdoor environment. It is currently based on Intel FlexRAN based L1 but is expected to also feature GPU accelerated L1 in the near future. Table 1 summarizes the comparison of CAMPUS5G with representative existing testbeds.

3 NETWORK PLANNING

3.1 Radio Placement and Spectrum

The goal of the planning process was to find a suitable placement of the radios that would provide both *blanket coverage* across the entire campus and surrounding areas, as well as *sufficient capacity* for known high-footfall areas where we expected larger numbers of users to congregate or traverse on a regular basis. We also preferred to over-provision rather than the opposite, as over-provisioning allows us to slice the network on a radio basis; for example one slice can serve our production users while others can be used for experimental and integration purposes.

An additional constraint was that we did not have unlimited flexibility for where we could deploy the RUs. Rather, we had to work within certain constraints placed by the heritage status of several of the university buildings and the feasibility of installing new fiber to buildings to support the O-RAN fronthaul requirements. This exercise resulted in identifying six university building sites where we could place the RUs.

From a spectrum perspective, the only viable option was to use the shared access spectrum made available through Ofcom (the UK communications regulator) for private 5G networks [64]. Specifically, there are 4 potential shared access spectrum bands to choose from – 1800 MHz, 2300 MHz, 3.8-4.2 GHz and 26 GHz. Of

these, very limited amount of spectrum was available to access for 1800 MHz and 2300 MHz bands – 6.6 MHz and 30 MHz, respectively. The 2300 MHz band is also restricted to indoor use. On the other hand, 26 GHz band provided abundant amount of spectrum – 3.05 GHz of spectrum – but would have required 100s of radios to be deployed to achieve campus wide coverage.

The 3.8-4.2 GHz (n77) band turned out to be a suitable middle ground with up to 390 MHz of spectrum and so we opted for that. Within this band, we had two different types of licenses to choose from with different power level restrictions but both allowed channel bandwidths up to 100 MHz. The ‘low-power’ license allowed power levels up to 21 dBm / 5 MHz for carriers > 20 MHz (EIRP per cell) for multiple cell sites within a 50m radius, whereas the ‘medium-power’ alternative required per-site license while allowing power levels up to 36 dBm / 5 MHz for carriers > 20 MHz.

Given the above constraints, we leveraged our knowledge of the campus environment and footfall in different areas to come up with a radio network plan that would achieve blanket coverage across campus, while ensuring sufficient capacity in all areas. Moreover, we optimized for reducing licensing cost where possible, by using medium-power licenses to reach most of the campus, with the remaining areas served by low-power licenses. This plan consisted of 20 RUs distributed across 6 building sites on campus along with the choice of power level, antenna type (omni or directional) and orientation for each of them. Fig. 1 (left) illustrates this plan overlaid on the campus map.

3.2 Pre-Deployment Coverage Map Estimation

Here we estimate the radio signal coverage map with our above outlined radio network plan. A coverage map [56] associates signal quality metrics to each discrete bin on a 2D plane, enabling mobile operators to evaluate signal reception at arbitrary locations. We mainly focus on the received signal strength indicator (RSSI) metric. To achieve a proper trade-off between coverage map resolution and computational complexity, the estimated coverage maps are obtained at the 1m×1m bin granularity within the campus area.

For the coverage map estimation, we employ the Sionna ray tracing (RT) based tool (SionnaRT) [56] because it supports accurate ray-tracing-based channel modeling, customized antenna pattern designs, is fully open-source and also GPU-accelerated. Before using SionnaRT, we first create a real-world campus scene model consisting of three types of objects, i.e., buildings, foliage, and ground (the terrain) with Blender 3.6 [19]. Specifically, we first import the footprints of buildings on campus from OpenStreetMap [67] into Blender. The foliage is then manually modeled by 3D cuboids, and the terrain information is based on the public LiDAR point cloud data for Scotland [33]. Each object in Blender is constituted by a mesh of triangles that is acceptable to SionnaRT. Furthermore, the radio-frequency materials with the parameters from the ITU standards [41] are assigned to each object.

With the above scene model in hand, we then setup the ray tracer in SionnaRT as follows. The maximum number of interactions (including reflection, diffraction, and scattering) between a ray and a scene object is set to 5. In total, 2×10^6 random rays are traced. The 4-port omni-directional array pattern and the 4-port directional array pattern are carefully customized based on the respective

antenna data sheets. The receiver antenna (reflecting the mobile device) is configured as a single isotropic antenna.

We now estimate the RSSI coverage maps on the receiver plane at 1m×1m bin granularity for three distinct cases: 1) switching on all the radios; 2) only switching on medium-power radios; and 3) only switching on low-power radios. Each of those coverage maps is obtained using the highest RSSI among all the transmitters computed at each bin. As shown in Fig. 2(a), when all the 20 radios are switched on, the estimated coverage map reflects blanket coverage across the campus (between -65 and -50 dBm), with particularly strong coverage in the center of the campus area. When only medium-power radios are activated, most regions across the campus still experience satisfactory signal coverage (above -80 dBm) as illustrated in Figure 2(b). But certain areas highlighted by red circles show significant coverage deterioration compared with the scenario with all radios active. This result validates our radio network planning strategy, emphasizing the essential role of low-power directional radios in significantly enhancing signal quality in certain parts of the campus.

4 TESTBED DESIGN AND DEPLOYMENT

4.1 Architecture and Design Choices

From a testbed design perspective, besides the number of radios and their placement covered in the previous section, we have few further design decisions to make: (1) choosing an O-RAN deployment scenario; (2) fronthaul design; (3) handling time synchronization. For (1), we opt for the O-RAN deployment scenario “A” [7] where CU, DU and near real-time (RT) RIC are collocated at an edge cloud while RUs are realized as physical network functions that communicate with DU over the O-RAN fronthaul.

Regarding (2), we initially considered leveraging existing university fiber infrastructure across the campus. But O-RAN fronthaul has stringent latency and bandwidth requirements. Specifically, fronthaul latency must be below 100 μ s whereas multiple Gbps of fronthaul bandwidth per RU is needed – the exact breakdown between upstream/downstream dependent on the TDD configuration used [5, 26]. As we found this to be infeasible while sharing our university fiber infrastructure, we had to go with custom fiber installation between the edge cloud location and the RUs.

For (3), for O-RAN to reliably function, RUs must be time synchronized at a fine-grained level (sub-microsecond level) with the functions hosted at the edge cloud, particularly the DU [57]. This would only be feasible to achieve with the Precision Time Protocol (PTP) [25]. PTP operates in a master-slave setup, where the master is responsible for serving as the timing source for the slaves and master itself uses a highly-accurate source like GPS as reference. In our deployment, we had two options for handling time synchronization: (i) to have a single PTP master for the whole network hosted at the edge cloud; (ii) have an independent PTP master at each RU and the edge cloud. Given our decision to have custom fiber installed for the fronthaul, (i) was sufficient to ensure network-wide precise time synchronization. Specifically, we used Viavi Qg 2 [36] as the PTP clock in our testbed.

Our CAMPUS5G testbed architecture following the above design decisions is illustrated in Fig. 1 (right). It consists of eight key components: the radio units (RUs); two different network switches

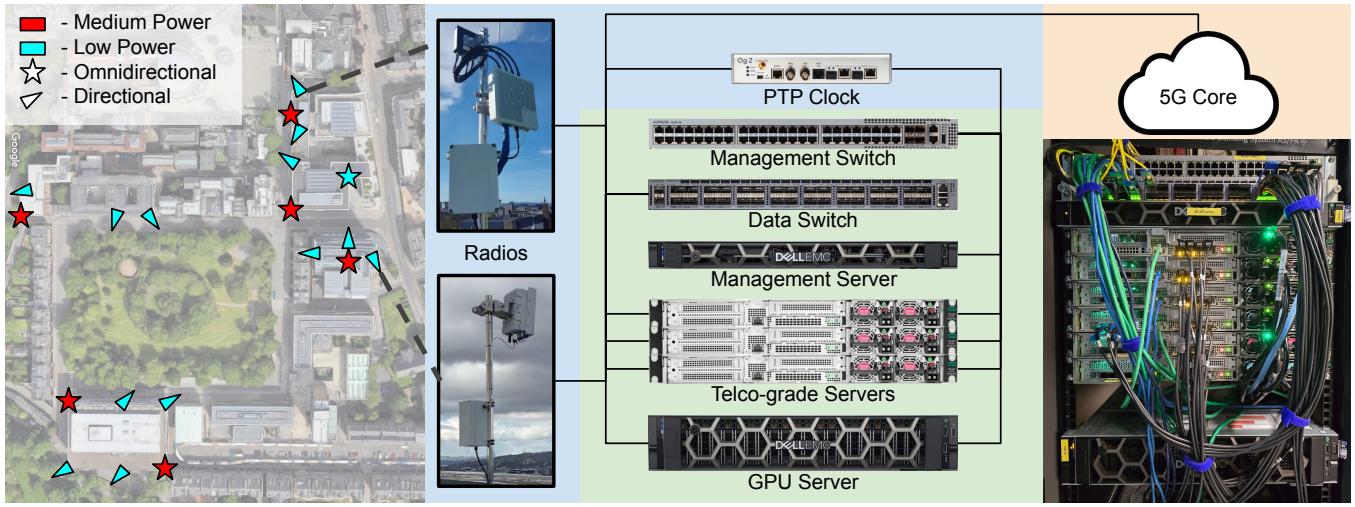


Figure 1: CAMPUS5G physical radio network plan (left) and network architecture (right).

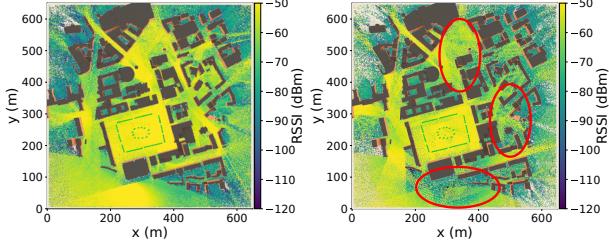


Figure 2: Estimated RSSI coverage maps for two different cases: (a) all radios switched on; (b) only the medium-power radios (with omni-directional antennas) switched on.

– one for data traffic, another for management traffic; a computing cluster consisting of a set of telco-grade servers to run the vRAN NFs; a GPU node to process RAN related AI workloads (e.g., AI based Apps); a management node to tie the cluster together; a PTP clock for precise time synchronization; and a cloud-based 5G core.

4.2 Testbed Components

Radios. We selected two O-RAN compliant RUs from two different vendors: the Benetel RAN650 [47] and the VVDN Mid-Power RU (MPRU) [79]. These were the only radios we found that meet our criteria: O-RAN 7.2x split compatibility, work in our desired n77 spectrum band, support for 100MHz and 4x4 MIMO, support from both commercial and open source vRANs, and outdoor-rated. Our deployment specifically used 12 Benetel and 8 VVDN RUs. Note that SDR option is not viable for our outdoor deployment as it did not readily support the O-RAN fronthaul.

Servers. The chosen telco-grade servers were the HP DL110 G10 Plus Telco servers, selected due to their track-record of supporting O-RAN workloads [17]. Each server was fitted with an Intel ACC100 vRAN Accelerator card to provide hardware-acceleration for the High-PHY related signal processing required in the O-RAN DU; these accelerators enabled flexibility through their wide support by commercial RANs. To ensure adequate network capacity and redundancy, each server was fitted with dual quad-port 25Gbps network cards, giving a total capacity of 200Gbps per server. Our

current deployment consists of six of these telco servers, the number carefully chosen to meet the processing needs across 20 RUs.

The management and GPU servers were more straightforward: the management server could be any machine capable of acting as the controller in a six-node cluster, whereas the GPU server needed only to support modern graphics cards and high-speed network interconnect to the telco nodes.

Switches. Our architecture uses two network switches. The data switch is the key one that connects the radios to the servers. Due to the demanding fronthaul bandwidth and latency requirements, this switch must support 10 Gbps links from each RU, and eight 25 Gbps links from each server. To fulfill this requirement, we chose the Arista 7050X3 switch. The second switch is a management switch, used for keeping management and data traffic separate. Unlike the data switch, there are no strict throughput or latency requirements, so we opted for a lower spec Arista 7010X.

Platform. The role of the platform is to provide a software layer atop the hardware, for all the upper layers to run on. The platform is a key enabler of flexibility, providing a base for running a diverse set of possible RANs, RICs and Apps. For the base layer, i.e., the operating system (OS), we chose Ubuntu [21], optimizing it to support real-time processing – a necessity for the RAN functions. Atop the OS, Ansible [35] was used to provision basic utilities such as networking, some drivers, and the Kubernetes orchestrator [32]. Kubernetes is the main component of the platform. All of the workloads intended to be deployed (RANs, RICs and Apps) were containerized and deployed atop it. To manage this orchestration, we made extensive use of Helm charts [31], which allowed ‘one-click’ deployment of a specified RAN setup.

Mobile Network Functions. Our platform accommodates a variety of mobile network functions, emphasizing flexibility and interoperability in line with the spirit of Open RAN. The platform is designed to seamlessly integrate any O-RAN-compliant vRAN function. We have extensively tested open-source vRAN solutions, specifically srsRAN [77] and OpenAirInterface (OAI) [66]. Both srsRAN and OAI can be flexibly deployed either as monolithic

network functions, or separated into distinct CU and DU components. We have also assessed the feasibility of deploying commercial vRAN solutions (e.g., from CapGemini). In general, with our servers equipped with Intel ACC100 accelerator cards [39], our infrastructure is capable of running any Intel FlexRAN [38] based vRAN implementation from vendors like CapGemini, Mavenir, Radisys, Parallel Wireless and Rakuten.

We use srsRAN as the default vRAN software in our deployment and use a split CU-DU architecture with a single CU mapped to 20 DUs (one per RU). This setup eases mobility management, specifically enabling intra-CU inter-DU handovers and ensuring seamless user transitions between cells. For the core network, we currently use an Open5GS [65] based cloud-hosted solution. Our testbed supports access from commodity devices through custom eSIM and physical SIM cards.

RICs and Apps. Our platform supports the deployment of standard O-RAN E2 [58] based RICs, enabling real-time network optimization through intelligent control loops via xApps. Specifically, we have successfully integrated O-RAN SC RIC [75] and FlexRIC [73] along with standard applications such as the KPIMON xApp, facilitating basic monitoring and control functionalities. Given the rising interest in non-standard RICs, which offer more granular control over RAN operations, we have additionally integrated two research RIC solutions: EdgeRIC [43] and Microsoft Janus/jbpf [30, 53]. Our testbed also features advanced AI/ML based Open RAN Apps (e.g., SpotLight [78]).

4.3 RAN Configuration

RAN systems expose hundreds of configurable parameters [62, 76] for which vendors typically ship default values for each. However, certain parameters must be explicitly updated for the network to function at all. These critical settings depend on the specific RU, DU, platform (including OS and server hardware), and the deployment topology. In our deployment, the initial setup focused on key platform-specific parameters such as network interface bindings, IP addressing, MCC/MNC values, physical cell IDs, and PCI device mappings for DPDK. These were combined with DU and RU vendor-specific configurations that were automated using Kubernetes and Helm, streamlining repeatable deployments across multiple nodes.

However, we found that not all vendor defaults were viable in our environment, and several needed to be manually tuned or overridden. For instance, we deployed 20 radios overall from two different RU vendors, VVDN and Benetel, each supporting different feature subsets. VVDN RUs, for example, support only a fixed TDD configuration and short PRACH format, while Benetel supports seven TDD configurations along with both short and long PRACH formats. Additionally, each RU vendor specifies its own fronthaul delay profile [59], but these defaults often do not reflect the real-world deployments. In one case, migrating a VVDN RU from a lab testbed using short direct-attach-copper cables to a rooftop site over fiber introduced additional propagation delay, pushing eCPRI timestamps outside the DU’s expected window. This resulted in packet loss until we manually adjusted the delay profile to align with the physical link characteristics.

With 20 radios deployed across our campus, most locations observe coverage from 2–3 cells simultaneously. Such a deployment density required careful PRACH planning [86]. We allocated

Table 2: Comparison of various interpolation methods.

Category	Method	Metrics on five folds (Avg. \pm Std.)		
		RMSE	NMSE	MAPE (%)
Deterministic	IDW	8.21 ± 0.14	0.24 ± 0.01	7.06 ± 0.15
	RBF	7.25 ± 0.14	0.19 ± 0.01	5.85 ± 0.10
Stochastic	OK	7.37 ± 0.10	0.19 ± 0.00	5.95 ± 0.04
ML-based	RF	7.46 ± 0.06	0.20 ± 0.00	6.03 ± 0.02
	XGBoost	7.77 ± 0.09	0.21 ± 0.00	6.44 ± 0.05
Physics-aware	MRI	7.40 ± 0.15	0.19 ± 0.00	6.05 ± 0.06

a unique *prach_root_sequence_index* to each cell site to prevent random access (RACH) collisions or failures. These configurations were sufficient to bring up a stable and functional RAN across our deployment. However, they are not yet fully optimized. Many additional parameters require further tuning and per-site calibration to reach optimal performance.

With our current deployment and default TDD configuration, we have achieved max downlink (uplink) throughput of 800 Mbps (around 80Mbps). This is lower than the absolute maximum possible throughput of 1.4Gbps with 100MHz and 4x4 MIMO, suggesting room for optimization through configuration parameter tuning. To put these numbers in context, what we already achieve is more than 3.5 times the max downlink throughput measured with public 5G mobile networks deployed in our campus area.

5 MEASUREMENT BASED COVERAGE ANALYSIS

In this section, we present measurement data from CAMPUS5G collected through various walking paths across the campus. Since measurements cannot cover every possible location within the target area, interpolation is required to generate a complete coverage map. To determine an efficient interpolation method, we first evaluate candidate schemes using simulated coverage data produced by SionnaRT (§3.2). Then, we introduce the collected real-world measurement dataset, and the selected interpolation approach is applied to the measured data to create complete coverage maps.

5.1 Comparison of Interpolation Methods

The goal of spatial interpolation is to obtain the values at the unsampled locations. Let $\hat{Z}(s_0)$ be the estimate at an unsampled location s_0 from observations $\{Z(s_i)\}_{i=1}^N$ where s_i denotes an observed location. To determine the best interpolation method for the coverage map, we group approaches into four distinct families (outlined in the Appendix A) based on how they construct $\hat{Z}(s_0)$ and benchmark one or two representative algorithms from each family:

- **Deterministic:** *Inverse Distance Weighting* (IDW) and *Radial Basis Function* (RBF).
- **Geostatistical:** *Ordinary Kriging* (OK).
- **Machine learning based:** *Random Forest* (RF) and *XGBoost*.
- **Physics aware:** *Model-based radio interpolation* (MRI) [74].

To evaluate these algorithms, we perform 5-fold cross-validation on the RSSI dataset generated by SionnaRT, using 80% of the data for training and the remaining 20% for testing in each fold. The

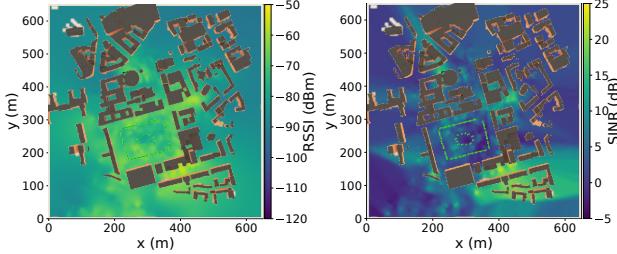


Figure 3: Measurement based RSSI coverage (left) and SINR (right) maps.

average and standard deviation of three metrics, i.e., Root Mean Square Error (RMSE), Normalized Mean Square Error (NMSE), and Mean Absolute Percentage Error (MAPE), are computed across folds, and reported in Table 2. These results show that RBF achieves the best RMSE and MAPE performance, while Ordinary Kriging and MRI are marginally better in terms of NMSE and its standard deviation. So, we select RBF to interpolate measurement based coverage maps, specifically with a multi-quadratic kernel (shape parameter $\epsilon = 1$ and smoothing parameter $\delta = 0.1$).

5.2 Measurement Based Coverage Maps

We conducted a 15 Km walking test across the campus to collect measurements using a commodity 5G smartphone, specifically Nothing Phone 2 [45], including GPS locations and UE reported metrics (RSRP, RSRQ and SINR). Around 11K samples in the form of [GPS location, reported metrics] are collected. To align the metrics obtained from Sionna and field measurements, we compute RSSI from the measured data based on the following relationship [13]:

$$\text{RSSI}_{\text{dBm}} = \text{RSRP}_{\text{dBm}} + 10 \log_{10} N - \text{RSRQ}_{\text{dB}}, \quad (1)$$

where N is the number of physical resource blocks (PRBs) in the measurement bandwidth configured in SS/PBCH block ($N = 20$ for our testbed).

The measurement based RSSI coverage map via RBF interpolation is shown in Fig. 3(a). As with the earlier ray-tracing based estimated coverage map in Fig. 2, measured map also confirms blanket coverage across the campus area.

To analyze the tradeoff between coverage and interference, we also interpolate the SINR metric to obtain the SINR based coverage map, as shown in Fig. 3(b). While our testbed delivers high signal strength in the central campus area, it also leads to severe interference (resulting in low SINR levels between -5 dB and 5 dB) when all radios are operating at full power. This highlights the critical trade-off between signal coverage and interference in dense deployments, necessitating careful interference management to optimize network performance.

6 LESSONS LEARNED

Cellular is attractive for campus wide coverage. Our post-deployment measurement based analysis of CAMPUS5G coverage revealed that it was remarkably resilient when compared to the existing WiFi network on campus. In scenarios where the WiFi network would drop out, our testbed (with the equivalent signal strength) worked well. This demonstrates the advantages of deploying cellular over WiFi, aligned with observation made in other

recent work [29]. This is particularly the case in large areas such as open spaces in university campuses, which would otherwise require a large number of WiFi access points to be deployed [24]. **First, make the platform real-time.** Ensuring real-time behavior is essential to run latency-sensitive DU workloads. Installing a real-time kernel is only the first step. Meaningful performance requires careful tuning of both hardware and software. This includes configuring BIOS settings to disable power-saving features, isolating CPU cores for RAN functions, removing interrupts from those cores, and disabling low-power C-states. We used tools like tuned [80] along with x86 processor-specific guidelines from Intel [37] and AMD [15] to apply these changes. Without these steps, even optimized RAN stacks showed inconsistent performance under load.

No two vendors are the same. We observed considerable variation between vendors, both in vRAN NFs (DU, CU) and O-RUs. Features like TDD configurations, PRACH support, and delay handling differ widely. In many cases, vendor defaults were incompatible with our environment, and given the early maturity of open RAN ecosystems, support teams were not always able to help. These differences meant we often had to debug and adapt configurations ourselves. Seamless multi-vendor interoperability remains a challenge, requiring significant deployment-specific tuning.

Performance vs. stability trade-off. Configuring RAN for maximum performance can introduce system-level issues. Aggressively prioritizing RAN threads may yield higher throughput but risks starving kernel processes or other workloads of shared resources, such as CPU caches. On systems using a real-time kernel, this can make the platform unresponsive, as even essential kernel threads may be blocked. In our experience, slightly trading off from peak performance in favor of system stability yields more reliable and maintainable deployments.

Deployment costs are insightful. Examining costs of deployment reveals that civil engineering costs (fiber, electrical and antenna installation) make up a substantial fraction (at least a third) of the overall CAPEX costs. However, overall deployment costs for a campus scale private 5G network like our CAMPUS5G testbed are surprisingly in the same ballpark as the costs for deploying a single macrocell in a dense urban area [85].

7 RESEARCH OPPORTUNITIES

Our experience deploying the CAMPUS5G testbed brought to the fore several key research opportunities that we outline below.

RAN configuration and optimization is a broad topic for future work. Even small changes in RAN configuration parameters can lead to significant behavioral shifts, motivating research on their optimization and dynamic adaptation. An illustrative example from our deployment is the impact of tuning the measurement reporting interval on handovers. In a controlled walk-around experiment we conducted on campus, we evaluated how different reporting interval settings affect the number of handovers triggered for the same physical route. Table 3 summarizes the results from this experiment.

The above example conveys how parameter tuning can influence not only the signaling overhead and mobility behavior but also user experience and system load. However, the optimal configuration is often context-dependent, influenced by mobility patterns, network density, and cell overlap.

Table 3: Impact of measurement reporting interval on handover frequency.

Report Interval (ms)	#Handovers
480	133
1024	39
5120	23

Systematic exploration of these parameter spaces, especially in conjunction with RIC based closed-loop control, represents a rich area for future research. Interference management as alluded to in §5.2 and Fig. 3(b) is another case demanding further work. More issues to look at along these lines include: (1) intelligent and dynamic use of spectrum within a private 5G network and across networks; (2) continual optimization of resource and energy usage, which requires a holistic understanding of energy consumption and its interplay with the use of compute and radio resources.

Complementarily, monitoring and optimizing user experience is another topic with several research opportunities. Our measurement effort in §5.2 is a case in point. To obtain data for measurement based coverage maps, we conducted a 15 Km walk test to evaluate RAN behavior across our deployment. While effective, this process was very manual and required constant supervision. Automating such measurements towards zero-touch management for private 5G and Open RAN presents a strong research opportunity. Another significant one concerns optimizing user experience while moving both within the coverage area of a private network as well as at its edges with other networks.

8 CONCLUSIONS

In this paper, we have captured our journey planning, deploying and measuring CAMPUS5G, the first-of-its kind campus scale O-RAN testbed. Urban, outdoor and university campus setting that our deployment represents is a perfect blend for innovation on future wireless access in a real-world environment. We have reflected on our experience deploying CAMPUS5G and shared the key lessons. We have also outlined a number of research opportunities that came to the surface through our deployment experience.

ACKNOWLEDGMENTS

This work was mainly supported by a project funded by the UK Department for Science, Innovation and Technology (DSIT). The work of T. Wang and M. Marina was also supported by the UKRI/EPSRC grant UKRI860.

REFERENCES

- [1] [n. d.]. JET Connectivity. <https://jet-eng.co.uk/>. ([n. d.]).
- [2] [n. d.]. Neutral Wireless. <https://www.neutralwireless.com/>. ([n. d.]).
- [3] [n. d.]. Open Ireland Testbed. <https://open-ireland.atlassian.net/wiki/spaces/OP/overview>. ([n. d.]).
- [4] 2019. Worcester Bosch launches first 5G factory trial. <https://www.bosch.co.uk/internet-of-things/leading-in-5g-technology/>. (Apr 2019).
- [5] 2020. TSN Frame Preemption Meets Stringent Latency Requirements of 5G Fronthaul. <https://www.keysight.com/blogs/en/tech/traf-gen/2020/07/31/tsn-frame-preemption-meets-stringent-latency-requirements-of-5g-fronthaul>. (Jul 2020).
- [6] 2022. BMW builds open RAN 5G testbed. <https://www.fierce-network.com/private-wireless/bmw-builds-open-ran-5g-testbed>. (Jul 2022).
- [7] 2023. O-RAN Cloud Deployment Scenarios. <https://rimedolabs.com/blog/o-ran-deployment-scenarios/>. (Jul 2023).
- [8] 2024. Current State of Open RAN - Countries & Operators deploying & trialing Open RAN. <https://tecknexus.com/5g-network/5g-magazine/open-ran-june-2021/current-state-of-open-ran-countries-operators-deploying-trialing-open-ran/> an/. (Nov 2024).
- [9] 2024. Ericsson Mobility Report. <https://www.ericsson.com/en/reports-and-papers/mobility-report>. (Nov 2024).
- [10] 2024. Innovative private network deployments from 2023. <https://stlpartners.com/articles/private-cellular/private-network-deployments-from-2023/>. (Apr 2024).
- [11] 2025. A (baker's) dozen key UK private 5G projects to make Industry 4.0 rise. <https://www.rcrwireless.com/20250122/fundamentals/uk-private-5g-industry>. (Jan 2025).
- [12] 2025. O-RAN Architecture Overview. <https://docs.o-ran-sc.org/en/latest/architecture/architecture.html>. (2025).
- [13] 3rd Generation Partnership Project (3GPP). 2023. *NR; Physical layer procedures for control*. Technical Specification (TS) 38.213. 3rd Generation Partnership Project (3GPP). https://www.etsi.org/deliver/etsi_ts/138200_138299/138213/17.07.00_60_ts_138213v170700p.pdf Release 17.
- [14] Accelleran. [n. d.]. Open Private 5G, Today. <https://accelleran.com/>. ([n. d.]).
- [15] AMD. [n. d.]. Low Latency Tuning for AMD EPYC CPU-Powered Servers. https://www.amd.com/content/dam/amd/en/documents/epyc-technical-docs/white-papers/58649_amd-epyc-tg-low-latency.pdf. ([n. d.]).
- [16] Amit Nagpal. 2025. What is happening to Open RAN? <https://cms.law/en/int/expert-guides/cms-expert-guide-to-5g-regulation-and-law/what-is-happening-to-open-ran>. (Mar 2025).
- [17] Paramvir Bahl, Matthew Balkwill, Xenofon Foukas, Anuj Kalia, Daehyeok Kim, Manikanta Kotaru, Zhihua Lai, Sanjeev Mehrotra, Bozidar Radunovic, Stefan Saroiu, Connor Settle, Ankit Verma, Alec Wolman, Francis Y. Yan, and Yongguang Zhang. 2023. Accelerating Open RAN Research Through an Enterprise-scale 5G Testbed. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking (ACM MobiCom '23)*. Association for Computing Machinery, New York, NY, USA, Article 138, 3 pages. <https://doi.org/10.1145/3570361.3615745>
- [18] Lorenzo Bertizzolo, Leonardo Bonati, Emrecan Demirors, Amani Al-Shawabka, Salvatore D’Oro, Francesco Restuccia, and Tommaso Melodia. 2020. Arena: A 64-antenna SDR-based ceiling grid testing platform for sub-6 GHz 5G-and-Beyond radio spectrum research. *Comput. Networks* 181 (2020), 107436. <https://doi.org/10.1016/j.comnet.2020.107436>
- [19] Blender Foundation. 2023. *Blender (Version 3.6)*. Blender Foundation. <https://www.blender.org> Free and open-source 3D creation suite.
- [20] Joe Breen, Andrew Buffmire, Jonathon Duerig, Kevin Dutt, Eric Eide, Annesa Ghosh, Mike Hibler, David Johnson, Sneha Kumar Kasera, Earl Lewis, Dustin Maas, Caleb Martin, Alex Orange, Neal Patwari, Daniel Reading, Robert Ricci, David Schurig, Leigh B. Stoller, Allison Todd, Jacobus Van der Merwe, Naren Viswanathan, Kirk Webb, and Gary Wong. 2021. Powder: Platform for Open Wireless Data-driven Experimental Research. *Computer Networks* 197 (2021), 108281. <https://doi.org/10.1016/j.comnet.2021.108281>
- [21] Canonical. [n. d.]. Ubuntu Operating System. <https://ubuntu.com/>. ([n. d.]).
- [22] Tingjun Chen, Prasanthi Maddala, Panagiotis Skrimponis, Jakub Kolodziejczyk, Abhishek Adhikari, Hang Hu, Zhihui Gao, Arun Paidimari, Alberto Valdes-Garcia, Myung Lee, Sundeep Rangan, Gil Zussman, and Ivan Seskar. 2023. Open-access millimeter-wave software-defined radios in the PAWR COSMOS testbed: Design, deployment, and experimentation. *Comput. Netw.* 234, C (Oct. 2023), 12. <https://doi.org/10.1016/j.comnet.2023.109922>
- [23] Fed4FIRE+ Consortium. [n. d.]. Fed4FIRE+. <https://www.fed4fire.eu/>. ([n. d.]).
- [24] Sude Dogan-Tusha, Armed Tusha, Muhammad Iqbal Rochman, Hossein Nasiri, Joshua Roy Palathinkal, Mike Atkins, and Monisha Ghosh. 2025. Evaluation of Indoor/Outdoor Sharing in the Unlicensed 6 GHz Band. In *2025 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*.
- [25] John C Eidson. 2006. *Measurement, control, and communication using IEEE 1588*. Springer Science & Business Media.
- [26] O-RAN ALLIANCE e.V. 2025. *O-RAN Control, User and Synchronization Plane Specification 17.01 (O-RAN.WG4.TS.CUS.0-R004-v17.01)*. Technical Report.
- [27] Mah-Rukh Fida, Andra Lutu, Mahesh K. Marina, and Özgür Alay. 2017. ZipWeave: Towards efficient and reliable measurement based mobile coverage maps. In *IEEE INFOCOM 2017 - IEEE Conference on Computer Communications*. 1–9. <https://doi.org/10.1109/INFOCOM.2017.8057098>
- [28] Platform for Advanced Wireless Research Project Office. [n. d.]. PAWR Wireless Platforms. <https://advancedwireless.org/platforms/>. ([n. d.]).
- [29] Xenofon Foukas and Bozidar Radunovic. 2025. The future of the industrial AI edge is cellular. In *Proceedings of the 26th International Workshop on Mobile Computing Systems and Applications (HotMobile '25)*. Association for Computing Machinery, New York, NY, USA, 61–66. <https://doi.org/10.1145/3708468.3711887>
- [30] Xenofon Foukas, Bozidar Radunovic, Matthew Balkwill, and Zhihua Lai. 2023. Taking 5G RAN Analytics and Control to a New Level. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking (ACM MobiCom '23)*. Association for Computing Machinery, New York, NY, USA, Article 1, 16 pages. <https://doi.org/10.1145/3570361.3592493>
- [31] The Linux Foundation. [n. d.]. Helm - The package manager for Kubernetes. <https://helm.sh/>. ([n. d.]).

[32] The Linux Foundation. [n. d.]. Kubernetes: Production-Grade Container Orchestration. <https://kubernetes.io/>. ([n. d.]).

[33] Scottish Government. 2024. Scotland LiDAR Public Elevation Dataset. [https://data.gov.scot.](https://data.gov.scot/) (2024). Accessed December 23, 2025.

[34] Thomas Hoeschle, Florian Kaltenberger, Andreas Ingo Grohmann, Elif Tasdemir, Martin Reisslein, and Frank H. P. Fitzek. 2022. 5G InterOPERAbility of Open RAN Components in Large Testbed Ecosystem: Towards 6G Flexibility. In *European Wireless 2022; 27th European Wireless Conference*. 1–6.

[35] Red Hat Inc. [n. d.]. Ansible Collaborative: A gathering space to build automation skills and success. <https://www.redhat.com/en/ansible-collaborative>. ([n. d.]).

[36] VIAVI Solutions Inc. 2025. Qg 2 Multi-Sync Gateway and PTP Grandmaster. <https://www.viavisolutions.com/en-us/products/qg-2-multi-sync-gateway-and-ptp-grandmaster>. (2025).

[37] Intel. [n. d.]. Real-Time Host Installation. <https://github.com/intel/FlexRAN?tab=readme-ov-file#real-time-host-installation>. ([n. d.]).

[38] Intel. 2025. FlexRAN™ Reference Architecture for Wireless Access. <https://www.intel.com/content/www/us/en/developer/topic-technology/edge-5g/tools/flexran.html>. (2025). Accessed December 23, 2025.

[39] Intel. 2025. Intel® vRAN Accelerator ACC100 Adapter. <https://www.intel.com/content/www/us/en/products/sku/211796/intel-vran-accelerator-acc100-adapter/specifications.html>. (2025). Accessed December 23, 2025.

[40] Taimoor Ul Islam, Joshua Ofori Boateng, Md Nadim, Guoying Zu, Mukaram Shahid, Xun Li, Tianyi Zhang, Sail Reddy, Wei Xu, Ataberk Atalar, Vincent Lee, Yung-Fu Chen, Evan Gossling, Elisabeth Permatasari, Christ Somiah, Owen Perrin, Zhibo Meng, Reshal Afzal, Sarath Babu, Mohammed Soliman, Ali Hussain, Daji Qiao, Mai Zheng, Ozdal Boyraz, Yong Guan, Anish Arora, Mohamed Y. Selim, Arsalan Ahmad, Myra B. Cohen, Mike Luby, Ranveer Chandra, James Gross, Kate Keahey, and Hongwei Zhang. 2025. Design and implementation of ARA wireless living lab for rural broadband and applications. *Computer Networks* 263 (2025), 111188. <https://doi.org/10.1016/j.comnet.2025.111188>

[41] ITU-R. 2015. *Recommendation ITU-R P.2040: Effects of Building Materials and Structures on Radio-wave Propagation*. Standard P.2040. International Telecommunication Union.

[42] Ibraheem Kasujeer and Tom Rebbeck. 2023. Unlock the Power of Private 5G: Top Applications and Key Business Outcomes. (2023). Analysys Mason Report.

[43] Woo-Hyun Ko, Ushasi Ghosh, Ujwal Dinesha, Raini Wu, Srinivas Shakkottai, and Dinesh Bharadia. 2024. EdgeRIC: Empowering Real-time Intelligent Optimization and Control in NextG Cellular Networks. In *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24)*. USENIX Association, Santa Clara, CA, 1315–1330. <https://www.usenix.org/conference/nsdi24/presentation/ko>

[44] Weaver Labs. [n. d.]. Project Arana: Boosting Fan Experience in Stadiums. <https://weaverlabs.io/case-study/project-arana-introducing-cell-stack-as-a-naas-platform-accelerating-network-monetisation-through-new-business-models>. ([n. d.]).

[45] Nothing Technology Limited. [n. d.]. Nothing Phone 2. <https://nothing.tech/pages/phone-2>. ([n. d.]).

[46] Telet Research Limited. [n. d.]. Telet. <https://teletresearch.com/what-we-do/private-5g-networks>. ([n. d.]).

[47] Benetel Ltd. [n. d.]. RAN650 5G Split 7.2x Outdoor O-RU. <https://benetel.com/ran650>. ([n. d.]).

[48] Alexandru Martian, Razvan-Florentin Trifan, Teodora-Cristina Stoian, Marius-Constantin Vochin, and Frank Y. Li. 2025. Towards Open RAN in beyond 5G networks: Evolution, architectures, deployments, spectrum, prototypes, and performance assessment. *Comput. Netw.* 259, C (April 2025), 19. <https://doi.org/10.1016/j.comnet.2025.111087>

[49] Takeshi Matsumura. 2025. NICT Outdoor Local 5G Testbed. (May 2025). Keynote at IEEE DySPAN 2025 Workshop on Spectrum Access for Local Networks: Current Status and Future Challenges.

[50] Weskley Mauricio, Francisco Hugo Costa Neto, Maykon Silva, Rodrigo Kenji Yaly Aoki, Eduardo Melão, Sérgio Barros, Fernando Farias, Lucas Bondan, and Fuad M. Abinader. 2025. Assessing Nationwide Distributed Open RAN Deployments with the OpenRAN@Brasil Testbed. *IEEE Communications Magazine* 63, 2 (2025), 46–52. <https://doi.org/10.1109/MCOM.001.2400317>

[51] N. McKeown et al. 2008. OpenFlow: Enabling Innovation in Campus Networks. *ACM SIGCOMM Computer Communication Review* 38, 2 (March 2008), 69–74. <https://doi.org/10.1145/1355734.1355746>

[52] Fraser McLean, Leyang Xue, Chris Xiaoxuan Lu, and Mahesh Marina. 2022. Towards edge-assisted real-time 3D segmentation of large scale LiDAR point clouds. In *Proceedings of the 6th International Workshop on Embedded and Mobile Deep Learning (EMDL '22)*. Association for Computing Machinery, New York, NY, USA, 1–6. <https://doi.org/10.1145/3539491.3539591>

[53] Microsoft. [n. d.]. jbpf. <https://github.com/microsoft/jbpf>. ([n. d.]).

[54] Sanaz Mohammadjafari, Sophie Roginsky, Emir Kavurmacioglu, Mucahit Cevik, Jonathan Ethier, and Ayse Basar Bener. 2020. Machine Learning-Based Radio Coverage Prediction in Urban Environments. *IEEE Transactions on Network and Service Management* 17, 4 (2020), 2117–2130. <https://doi.org/10.1109/TNSM.2020.3035442>

[55] Magreth Mushi, Yuchen Liu, Shreyas Sreenivas, Ozgur Ozdemir, Ismail Guvenc, Mihail Sichitiu, Rudra Dutta, and Russ Gyurek. 2024. Open RAN testbeds with controlled air mobility. *Comput. Commun.* 228, C (Dec. 2024), 20. <https://doi.org/10.1016/j.comcom.2024.107955>

[56] NVIDIA. 2025. SionnaRT – Ray-Tracing Simulator. <https://nvlabs.github.io/sionna/>. (2025). Accessed December 23, 2025.

[57] O-RAN. [n. d.]. O-RAN Control, User and Synchronization Plane Specification 17.01. <https://specifications.o-ran.org/download?id=856>. ([n. d.]).

[58] O-RAN. 2022. Near-Real-time RAN Intelligent Controller E2 Service Model (E2SM) KPM. <https://orandownloadsweb.azurewebsites.net/download?id=393>. (2022).

[59] O-RAN. 2024. O-RAN Fronthaul Interoperability Test Specification (IOT) 12.0. <https://specifications.o-ran.org/download?id=742>. (2024). Accessed December 23, 2025.

[60] O-RAN Alliance. 2025. O-RAN Map. <https://map.o-ran.org/>. (2025). Accessed December 23, 2025.

[61] O-RAN ALLIANCE e.V. [n. d.]. O-RAN Alliance. <https://www.o-ran.org/>. ([n. d.]).

[62] OAI. 2025. OpenAirInterface documentation overview. https://gitlab.eurecom.fr/oai/openairinterface5g/-/tree/develop/doc?ref_type=heads. (2025). Accessed December 23, 2025.

[63] 5G Observatory. [n. d.]. 5G private networks. <https://5gobservatory.eu/5g-private-networks/>. ([n. d.]).

[64] Office of Communications (Ofcom). [n. d.]. *Shared Access Licence Guidance Document*. Guidance document. Office of Communications (Ofcom). <https://www.ofcom.org.uk/siteassets/resources/documents/consultations/category-1-1-0-weeks/consultation-supporting-increased-use-of-shared-spectrum/associated-documents/shared-access-licence-guidance-document-2024.pdf?v=388817> Version 1.5; updated 14 January 2025.

[65] Open5GS. 2025. Introduction to Open5GS. <https://open5gs.org/open5gs/docs/guide/01-quickstart/>. (2025). Accessed December 23, 2025.

[66] OpenAirInterface. 2025. OAI 5G RAN PROJECT GROUP. <https://openairinterface.org/oai-5g-ran-project/>. (2025). Accessed December 23, 2025.

[67] OpenStreetMap contributors. 2025. Building Footprints. <https://openstreetmap.org>. (2025). Accessed December 23, 2025.

[68] L. Peterson, O. Sunay, and B. Davie. 2023. *Private 5G: A Systems Approach*. <https://5g.systemsapproach.org/>

[69] Michele Polese, Leonardo Bonati, Salvatore D’Oro, Stefano Basagni, and Tommaso Melodia. 2023. Understanding O-RAN: Architecture, Interfaces, Algorithms, Security, and Research Challenges. *IEEE Communications Surveys & Tutorials* 25, 2 (April 2023), 1376–1411. <https://doi.org/10.1109/COMST.2023.3239220>

[70] Michele Polese, Mischa Dohler, Falko Dressler, Melike Erol-Kantarci, Rittwik Jana, Raymond Knopp, and Tommaso Melodia. 2024. Empowering the 6G Cellular Architecture With Open RAN. *IEEE Journal on Selected Areas in Communications* 42, 2 (2024), 245–262. <https://doi.org/10.1109/JSAC.2023.3334610>

[71] Pedro Rezende, Shyam Mahato, Amogh PC, Aris Risdianto, Vignesh Nagamuthu, Yiyang Pei, and Sumei Sun. 2025. Comprehensive End-to-End Performance Evaluation of a Multi-Vendor 5G O-RAN-based Testbed. In *Proceedings of the 2nd ACM Workshop on Open and AI RAN (OpenRan ’25)*. Association for Computing Machinery, New York, NY, USA, 8–14. <https://doi.org/10.1145/3737900.3770165>

[72] J. Xavier Salvat, Jose A. Ayala-Romero, Lanfranco Zanzi, Andres Garcia-Saavedra, and Xavier Costa-Perez. 2023. Open Radio Access Networks (O-RAN) Experimentation Platform: Design and Datasets. *IEEE Communications Magazine* 61, 9 (2023), 138–144. <https://doi.org/10.1109/MCOM.003.2200711>

[73] Robert Schmidt, Mikel Irazabal, and Navid Nikaein. 2021. FlexRIC: an SDK for next-generation SD-RANs. In *Proceedings of the 17th International Conference on Emerging Networking Experiments and Technologies (CoNEXT ’21)*. Association for Computing Machinery, New York, NY, USA, 411–425. <https://doi.org/10.1145/3485983.3494870>

[74] Hyojeong Shin, Yohan Chon, Yungeun Kim, and Hojung Cha. 2015. MRI: Model-Based Radio Interpolation for Indoor War-Walking. *IEEE Transactions on Mobile Computing* 14, 6 (2015), 1231–1244.

[75] O-RAN Software Community. [n. d.]. O-RAN Software Community. <https://o-ran-sc.org/>. ([n. d.]).

[76] SRSRAN. 2025. Configuration Reference. https://docs.srsran.com/projects/project/en/latest/user_manuals/source/config_ref.html. (2025). Accessed December 23, 2025.

[77] SRSRAN. 2025. srsRAN Project Documentation. <https://https://docs.srsran.com/projects/project/en/latest/>. (2025). Accessed December 23, 2025.

[78] Chuanhai Sun, Ujjwal Pawar, Molham Khoja, Xenofon Foukas, Mahesh K. Marina, and Bozidar Radunovic. 2024. SpotLight: Accurate, Explainable and Efficient Anomaly Detection for Open RAN. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking (ACM MobiCom ’24)*. Association for Computing Machinery, New York, NY, USA, 923–937. <https://doi.org/10.1145/3636534.3649380>

[79] VVDN Technologies. [n. d.]. 5G NR Enterprise Radio Units. <https://www.vvdn.com/adaptive-compute-and-comms/enterprise-radio-unit>. ([n. d.]).

[80] tuned. [n. d.]. TuneD : Tuning Profile Delivery Mechanism for Linux. <https://tuned-project.org/>. ([n. d.]).

[81] Armed Tusha, Seda Dogan-Tusha, Hossein Nasiri, Muhammad Iqbal Rochman, Patrick McGuire, and Monisha Ghosh. 2024. A Comprehensive Analysis of Secondary Coexistence in a Real-World CBRS Deployment. In *2024 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*. 79–87. <https://doi.org/10.1109/DySPAN60163.2024.10632859>

[82] Davide Villa, Imran Khan, Florian Kaltenberger, Nicholas Hedberg, Rúben Soares da Silva, Stefano Maxenti, Leonardo Bonati, Anupa Kelkar, Chris Dick, Eduardo Baena, Josep M. Jornet, Tommaso Melodia, Michele Polese, and Dimitrios Koutsoukolas. 2025. X5G: An Open, Programmable, Multi-vendor, End-to-end, Private 5G O-RAN Testbed with NVIDIA ARC and OpenAirInterface. (2025). arXiv:cs.NI/2406.15935. <https://arxiv.org/abs/2406.15935>

[83] Davide Villa, Miead Tehrani-Moayyed, Clifton Paul Robinson, Leonardo Bonati, Pedram Johari, Michele Polese, and Tommaso Melodia. 2024. Colosseum as a Digital Twin: Bridging Real-World Experimentation and Wireless Network Emulation. *IEEE Transactions on Mobile Computing* 23, 10 (2024), 9150–9166. <https://doi.org/10.1109/TMC.2024.3359596>

[84] M. Wani et al. 2025. Open RAN: A Concise Overview. *IEEE Open Journal of the Communications Society* 6 (2025), 13–28. <https://doi.org/10.1109/OJCOMS.2024.3430823>

[85] David Wisely, Ning Wang, and Rahim Tafazolli. 2018. Capacity and costs for 5G networks in dense urban areas. *IET Communications* 12, 19 (2018), 2502–2510. <https://doi.org/10.1049/iet-com.2018.5505> arXiv:<https://ietresearch.onlinelibrary.wiley.com/doi/pdf/10.1049/iet-com.2018.5505>

[86] Xincheng Zhang. 2017. *PRACH Optimization*. 663–701. <https://doi.org/10.1002/9781119158981.ch12>

A OVERVIEW OF INTERPOLATION METHODS

- **Deterministic.** The spatial field is assumed deterministic in this category. The interpolator is a linear/non-linear function of known data points. *Inverse Distance Weighting* (IDW) is the canonical linear method that embodies pure distance-decay weighting, while *Radial Basis Function* (RBF) non-linear interpolation adds global smoothness via a radial kernel. We specifically consider a multi-quadratic kernel for RBF.
- **Geostatistical.** The spatial field is modeled by a random process with a specified covariance (variogram) structure. Among the algorithm in this category, *Ordinary Kriging* (OK) is the canonical case with an unknown but constant mean and it serves as the benchmark for kriging variants and Gaussian process regression methods [27].
- **Machine-learning-based.** In this category, location coordinates (and any covariates) form a feature vector \mathbf{x} , and then a flexible function $\hat{Z}(\mathbf{s}_0) = f(\mathbf{x}_0; \theta)$ is learned by minimizing an interpolation loss [54], where θ denotes unknown parameters. We consider two effective ML-based algorithms that are widely used, namely *Random Forest* (RF) and *XG-Boost*. Both algorithms scale well to large spatial datasets and require much lighter computation than the neural network based methods.
- **Physics-aware.** Instead of relying solely on data, physics-aware methods incorporate physical laws/constraints into the interpolation process. A representative method in this category is *Model-based radio interpolation* (MRI) [74], which assumes a log-distance path-loss model for each transmitter and fits its parameters using linear regression.