SPARC: Spatio-Temporal Adaptive Resource Control for Multi-site Spectrum Management in NextG Cellular Networks

USHASI GHOSH*, UC San Diego, USA
AZUKA CHIEJINA*, George Mason University, USA
NATHAN STEPHENSON, George Mason University, USA
VIJAY K. SHAH, NC State University, USA
SRINIVAS SHAKKOTTAI, Texas A&M University, USA
DINESH BHARADIA, UC San Diego, USA

This work presents SPARC (Spatio-Temporal Adaptive Resource Control), a novel approach for multi-site spectrum management in NextG cellular networks. SPARC addresses the challenge of limited licensed spectrum in dynamic environments. We leverage the O-RAN architecture to develop a multi-timescale RAN Intelligent Controller (RIC) framework, featuring an xApp for near-realtime interference detection and localization, and a μApp for real-time intelligent resource allocation. By utilizing base stations as spectrum sensors, SPARC enables efficient and fine-grained dynamic resource allocation across multiple sites, enhancing signal-to-noise ratio (SNR) by up to 7dB, spectral efficiency by up to 15%, and overall system throughput by up to 20%. Comprehensive evaluations, including emulations and over-the-air experiments, demonstrate the significant performance gains achieved through SPARC, showcasing it as a promising solution for optimizing resource efficiency and network performance in NextG cellular networks.

CCS Concepts: • **Networks** → **Resource management**; • **Interference** → **Spectrum management**.

Additional Key Words and Phrases: RAN Intelligent Control, Interference detection, traffic

ACM Reference Format:
Ushasi Ghosh*, Azuka Chiejina*, Nathan Stephenson, Vijay K. Shah, Srinivas Shakkottai, and Dinesh Bharadia. 2024. SPARC: Spatio-Temporal Adaptive Resource Control for Multi-site Spectrum Management in NextG Cellular Networks. *Proc. ACM Netw.* 2, CoNEXT4, Article 35 (December 2024), 18 pages. https://doi.org/10.1145/3696405

1 Introduction

The rapid growth of next-generation applications places substantial demands on data throughput, latency, and reliability in modern communication networks. These applications increasingly rely on 5G technology to connect them to cloud and edge computing services, extending beyond traditional wide area networks to private 5G networks designed for enclosed spaces such as large enterprises, warehouses and office buildings. Network-sensitive applications like industrial IoT, which demand high reliability, and AR/VR, which are extremely data-intensive and require low latency, further emphasize the need for advanced network management solutions.

* Equal Contribution.

Authors' Contact Information: Ushasi Ghosh*, UC San Diego, USA, ughosh@ucsd.edu; Azuka Chiejina*, George Mason University, USA, achiejina@gmu.edu; Nathan Stephenson, George Mason University, USA, nstephenson@gmu.edu; Vijay K. Shah, NC State University, USA, vkshah@ncsu.edu; Srinivas Shakkottai, Texas A&M University, USA, sshakkottai@tamu.edu; Dinesh Bharadia, UC San Diego, USA, dineshb@ucsd.edu.

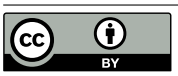

ACM 2834-5509/2024/12-ART35
https://doi.org/10.1145/3696405

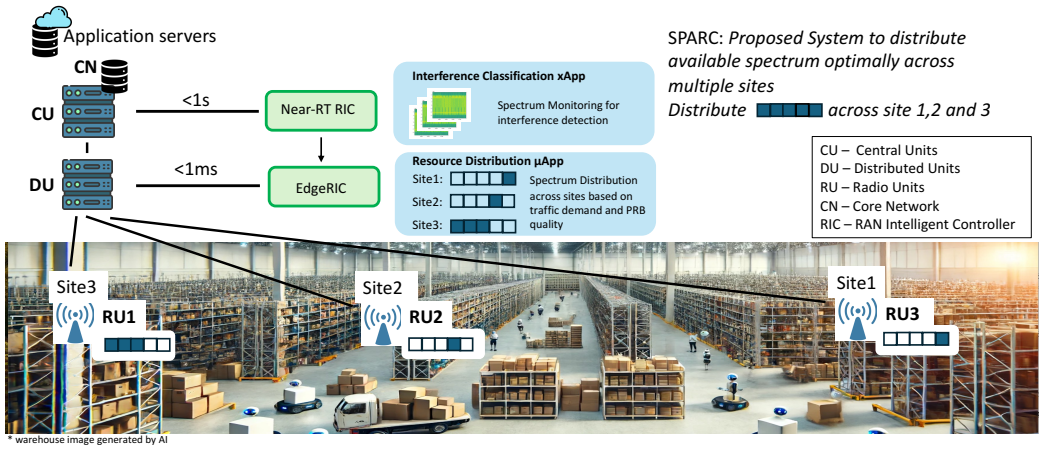

Fig. 1. System and Network Overview

A critical challenge in these deployments is the limited availability of licensed spectrum, necessitating efficient frequency reuse strategies across cell sites that jointly provide coverage over the space as shown in Figure 1. Here, three cell sites jointly provide coverage over a limited area and spectrum resources must be allocated across them. This paper introduces SPARC (Spatio-Temporal Adaptive Resource Control), a novel approach for multi-site spectrum management in NextG cellular networks. SPARC addresses the challenge of limited licensed spectrum in dynamic environments by leveraging the Open Radio Access Network (O-RAN) architecture that enables monitoring and control of radio access networks at different timescales.

Our approach centers on developing a multi-timescale RAN Intelligent Controller (RIC) framework, featuring an xApp for near-realtime (< 1s) interference detection and localization and a μApp for real-time (<1 ms) intelligent resource allocation. The two RICs operate in concert, sharing information and taking control actions to determine the appropriate spectrum resources to allocate at each site as shown in Figure 1. By using base stations as spectrum sensors, SPARC enables efficient and fine-grained dynamic resource allocation across multiple sites. This approach enhances signal-to-noise ratio (SNR), spectral efficiency, and overall system throughput, making it a robust solution for optimizing resource usage efficiency and network performance.

To address the dual requirements of resource usage efficiency and throughput maximization, SPARC employs multiple Radio Units (RUs) across different sites, leveraging disaggregated cellular architectures such as the O-RAN Split 7.2. This configuration effectively distributes the processing load and enhances local signal strength by permitting the deployment of many relatively simpler/cheaper RUs over a given area. However, the challenge of limited spectrum availability persists, requiring dynamic reallocation to maximize system throughput, especially under variable traffic levels. Efficient spectrum management necessitates sophisticated intelligence capable of operating at very low time granularity to determine and allocate the optimal spectrum parts to specific sites.

SPARC introduces several key contributions:

- **Multi-timescale RIC Approach**: We develop and demonstrate a multi-timescale RIC approach for efficient spectrum management, enabling information sharing and joint optimization using both a near-realtime RIC and a real-time RIC in multi-site scenarios. Our evaluations show a significant enhancement in SNR by up to 7dB.
- **Base Station as Spectrum Sensor**: We leverage the base station as a spectrum sensor, designing an xApp capable of detecting and localizing interference using object detection

techniques. This provides critical information about the physical resource blocks (PRBs) affected by interference.
- **Intelligent Resource Distribution**: Utilizing information from the xApp, we design a μApp for intelligent resource distribution in real-time across different RAN sites. This introduces the notion of resource block blanking to optimally redistribute limited spectrum, improving spectral efficiency by up to 15%.
- **Comprehensive Evaluations**: Through simulations and over-the-air experiments, we validate the benefits of SPARC, observing performance gains in terms of throughput (up to 20%), SNR, and spectral efficiency. These results highlight the effectiveness of the joint capabilities of near-realtime RIC and EdgeRIC for multi-site resource sharing.

In summary, SPARC offers a promising solution for enhancing energy efficiency and network performance in next-generation cellular networks through innovative spectrum management and intelligent resource control mechanisms.

2 Motivation and Background

Recent incidents have underscored the critical importance of spectrum awareness in modern communication networks. In one notable instance, a village in Wales with 400 residents experienced daily DSL Internet outages for 18 months due to electrical interference from an old TV [25]. Network operators were baffled until they traced the issue to a single household appliance emitting electrical noise. In another case, a truck driver's GPS jammer disrupted satellite systems at Newark airport [24], interfering with an advanced system designed to improve airport operations. Additionally, concerns have arisen regarding the susceptibility of 5G networks to jamming attacks, which could lead to denial of service (DoS) in critical applications, with severe impacts on both individuals and infrastructure [23]. These examples show how external interference from stray devices can infiltrate communication bands, causing widespread disruptions. They highlight the necessity for robust spectrum awareness to identify and mitigate such interference, ensuring network integrity. Without it, operators may struggle to pinpoint the source of failures, risking prolonged outages. Therefore, spectrum awareness is vital for maintaining reliable and resilient communication networks.

In this work, our primary goal is to leverage the O-RAN architecture to augment a cellular network with spectrum sensing capabilities. We use base stations as spectrum sensors to collect I/Q samples and deploy an interference detection and localization module as an xApp on the RAN Intelligent Controller (RIC) to identify and locate interference in the spectrum. This capability enables the xApp to suggest blanking out interfered bands, improving the signal-to-interference/noise ratio (SINR), enhancing system throughput, and reducing block error rates (BLER). The concept of Resource Block (RB) blanking leads to valuable use cases, such as optimizing resource redistribution across multiple cell sites. Our second goal is to support multiple radio units (RUs) operating over the same spectrum, extending RB blanking to optimize resource allocation in real-time, driven by traffic demand and spectrum sensing. This approach enables deploying multiple RUs over the same bandwidth, bringing transmitters closer to user equipment for improved performance.

An added advantage of this method addresses the important issue of sustainability. Research by [17] shows that base-station densification can create more sustainable wireless networks that scale effectively with user demand. Instead of relying on a single powerful base station to reach distant clients, multiple smaller base stations with simpler hardware and reduced signal levels can achieve the same task more efficiently and with lower power consumption. Further support for this approach comes from [16], who highlight that the main contributors to the increased carbon footprint in wireless networks are smartphone batteries and base stations consuming energy for

last-mile connectivity. Base-station densification mitigates these issues, reducing both sources of the carbon footprint.

Our system design promotes sustainability through densification by integrating spectrum sensing and real-time resource allocation within the O-RAN architecture. In this approach, a dense deployment of radio units (RUs) operating on the same frequency band is connected to a Distributed Unit (DU), with intelligence provided by RICs. The near-realtime RIC offers intelligence through xApps, which conduct sensing and optimization of network parameters at near-real-time granularity (<1s). In our system, spectrum monitoring is embedded within an xApp, enabling the network to detect interference and dynamically reconfigure its parameters accordingly.

At a finer time granularity, we use EdgeRIC [20], which operates in real-time and hosts μApps for intelligent control. In our system, this control involves resource allocation based on interference conditions observed by the xApp. This combined approach enables efficient spectrum management and resource redistribution. Figure 2 formalizes our problem: Case 1 shows a traditional monolithic cellular stack without spectrum awareness, Case 2 adds spectrum awareness with Resource Block (RB) blanking, and Case 3 illustrates our system, where resources are distributed across multiple sites based on local interference. This approach brings transmitters closer to user equipment. By integrating multi-timescale monitoring and control, we offer a framework for interference avoidance and resource efficiency. To our knowledge, we are the first to explore multi-site management within the O-RAN and RIC framework.

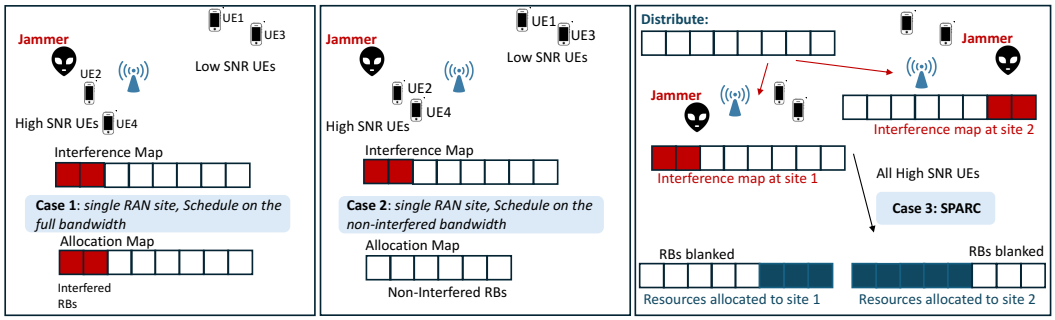

Fig. 2. Problem Definition illustrated

2.1 Background on O-RAN

The cellular network infrastructure is shifting towards Open Radio Access Networks (Open RAN) [29], promoting vendor diversity and interoperability while moving away from traditional monolithic architectures. This transition involves the softwarization and disaggregation of the 5G cellular stack, with higher layers (Central Unit, CU) hosted in data centers near the core and edge servers, and lower layers (Distributed Units, DUs) handling signal processing near the Radio Units (RUs). The O-RAN Alliance, comprising industry and academic experts, is standardizing these networks and defining key use cases [27, 28]. Open-source 5G stacks like Open Air Interface [26] and srsRAN [35], along with contributions from companies like Mavenir [22], Radisys [30], Nvidia [2], Microsoft [5], and Intel [12], are advancing ORAN-compliant infrastructures, marking a new era in the evolution of radio access networks.

The O-RAN movement is reshaping network architecture by opening interfaces for efficient metric collection and integrating AI/ML solutions. It introduces standardized, programmable RAN Intelligent Controllers (RICs), supporting near-realtime RIC and non-realtime RIC components. The near-realtime RIC runs xApps that use ML algorithms to optimize the RAN within milliseconds

to seconds, leveraging RAN data. It also hosts databases and an Internal Messaging Infrastructure (IMI) for data routing via the standardized E2 interface, which collects data and communicates control decisions to the RAN. Meanwhile, the non-realtime RIC, part of the Service Management Orchestration (SMO) framework, supports rApps for longer timescale control. Open-source options like OSC RIC [8] and FlexRIC [33] enhance both non-realtime and near-realtime needs. Real-time RICs (sub-5ms RAN control) are still emerging, with recent research highlighting their potential in the ORAN framework [11, 13, 14, 20, 21].

In this work, we leverage the O-RAN architecture to demonstrate the advantages of a multi-timescale monitoring and control approach, with a focus on enhancing spatial diversity across multiple sites. By integrating the near-realtime RIC and the real-time RIC, we aim to optimize resource distribution across sites, enabling more efficient resource management within the RAN. This collaborative framework highlights the potential of a synchronized multi-RIC system to improve network performance and efficiency.

2.2 Related Work

Spectrum sensing using spectrograms is a widely adopted method for detecting energy levels across frequency bands [7, 9]. This approach has advanced with the integration of machine learning (ML) techniques to analyze these time-frequency images generated from detected energy levels [39]. In this work, we use an ML-based technique to detect external interference within our band of interest, leveraging spectrograms to infer the spectrum environment and perform interference localization.

While 3GPP standards use reference signals like Sounding Reference Signals (SRS) and Demodulation Reference Signals (DMRS) to estimate channel characteristics per UE [1, 18, 40], they lack a system-wide view. In contrast, spectrograms observed at the base station offer a comprehensive perspective of the spectral environment, detecting external signals even without active UE traffic. Advanced spectrogram-based methods [37] provide deeper insights into spectrum usage compared to traditional KPI monitoring methods, which focus on metrics like packet error rate (PER), bit error rate (BER), and SINR [38].

While KPIs like CQI, SINR, and error rates help inform MCS selection, they don't reveal the spatial distribution or interference patterns across the spectrum, limiting their role in network optimization. Spectrograms offer granular visibility into interference, enabling intelligent resource allocation by blanking affected PRBs and utilizing unaffected ones, thus improving efficiency. In shared spectrum environments like CBRS, spectrograms help differentiate between radar signals and external interference, supporting better coexistence strategies.

Studies [31, 32, 34] highlight the potential of spectrograms to improve network performance and regulatory compliance, with one study achieving 100% interference detection accuracy under SINR conditions >=12dB. However, spectrograms may struggle with low-probability intercept signals. In this work, SPARC uses spectrograms from I/Q data collected at each RU to detect and map interference, optimizing frequency allocation and improving network management.

In the cellular network domain, significant advancements have been made in spectrum sensing to identify and avoid compromised frequency bands. A notable work within the O-RAN framework is ChARM [6], which enhances network management through spectrum awareness. ChARM integrates an additional radio into the base station that processes raw IQ samples to detect interference. Upon identifying suspicious activity, ChARM suggests shifting the entire network's center frequency of operation at the affected site.

Other works targeting spectrum sensing and sharing for optimized resource utilization include [15], which introduces ProSAS, a data-driven spectrum management solution. ProSAS predicts radio resource demand and manages spectrum to minimize surplus or deficit in RANs. Additionally, [31] presents SenseORAN, an enhancement to cellular communications and spectrum sensing

through Open RAN (O-RAN). SenseORAN uses a YOLO-based machine learning framework in the near-realtime RIC to detect radar pulses in the CBRS band, significantly improving radar interference management response times.

Our work takes a more nuanced approach by leveraging the spatio-temporal diversity of frequency bands. Interference impacts can vary spatially and temporally, meaning different areas may experience affected frequency bands at different times. To exploit this, we propose deploying multiple transmitters with spectrum sensing capabilities across various locations. Each node can independently detect compromised spectrum regions, specifically the affected Physical Resource Block (PRB) areas. Based on these localized sensing results, spectrum can be dynamically redistributed, allowing each transmitter to operate on the clearest available frequencies. This is especially beneficial for enterprises with limited licensed spectrum, as it maximizes utilization across the bandwidth. By adjusting frequency usage in real time, our system enhances network performance and efficiency, ensuring optimal spectrum use.

3 System Design

Figure 3 provides an overview of our system's main components within the O-RAN framework for multi-site resource sharing. The network is O-RAN compatible, with radio units distributed across various sites. EdgeRIC [20] handles the real-time (~1ms) monitoring and control of DU functions, communicating with the lower RAN layers via the RT-E2 interface. The near-realtime RIC operates at a coarser timescale (< 1 s) and communicates with the RAN over the E2 interface. We summarize the system architecture in the next subsection.

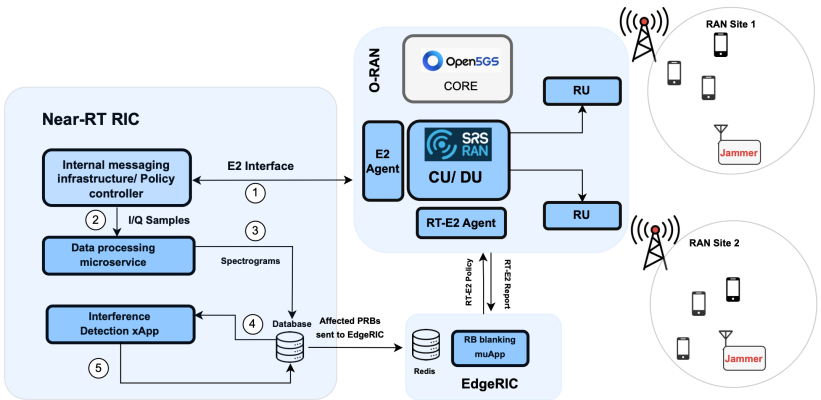

Fig. 3. Overview of System

3.1 System Architecture

We adopt a multi-timescale monitoring and control approach to optimize system behavior, focusing primarily on the uplink spectrum. Uplink scenarios are particularly challenging due to power constraints at the User Equipment (UE), especially for small sensor nodes with limited power. These devices are highly sensitive to interference, which can increase retransmissions, power consumption, and reduce battery life. Addressing uplink interference is crucial to ensure these devices operate efficiently without depleting their energy reserves. In contrast, downlink interference is less critical, as base stations have sufficient power to transmit effectively despite interference. Our focus on uplink scenarios aims to improve reliability and energy efficiency for power-constrained UEs, which is vital for the sustainability of mobile devices, sensor networks, and IoT applications. In Time Division Duplex (TDD) systems, channel reciprocity allows interference patterns observed

in the uplink to be inferred for the downlink spectrum as well. Additionally, CSI-RS reports help identify downlink interference. For completeness, we will also present a downlink scenario from our simulation, highlighting the benefits of real-time resource distribution across multiple sites.

■ **Spectrum Sensing: A Near Real-time Approach** We implement spectrum sensing capabilities for the base station as an xApp, running as a spectrum monitoring, detection and localization microservice in the near-realtime RIC. Raw I/Q samples from the RAN sites are collected by the E2 agent through the E2 interface into the RIC for further processing and analysis. An elaborate breakdown of the overall system is as follows:

1 After connection establishment between the near-realtime RIC and E2 agent which serves as entry point to the RAN, I/Q samples from both RF frontends are forwarded to the policy controller within the near-realtime RIC via the E2 interface. These I/Q samples are initially stored in separate buffers to distinguish different RAN sites. For experimental purposes, we periodically collect the last 10 ms segments of I/Q samples from each RAN site which is equivalent to the length of one LTE/5G frame. To reduce round trip time, there can be a trade-off between collecting a full frame of 10 ms or reducing it to 5 ms segments.

2 These I/Q samples for the two RAN sites are then forwarded to a data processing microservice which is used to process and convert the raw I/Q samples into spectrograms.

3 The data processing microservice then forwards the computed spectrograms to a Redis-based database hosted within the near-realtime RIC. This is to ensure that the data is accessible by any xApp that may want to utilize.

4 The interference detection/Localization xApp queries the database to get the latest spectrograms for the two RAN sites. These spectrograms give a good picture of the spectrum at the two RAN sites.

5 Using the ML model deployed within the xApp, we first detect the presence of interference signal at the two RAN sites using the spectrograms, then we go a step further to determine which Physical Resource Blocks (PRBs) are affected by the interference. These inference results are then forwarded into the database to be used by any other microservice/xApp.

6 Finally, the latest information of the affected PRBs is made accessible to EdgeRIC.

■ **Resource Distribution: A Real-time Approach**

The allocation of resources per site is determined every transmission time interval (TTI, 1ms) based on the traffic requirements at each site. EdgeRIC communicates with the MAC layer of the DU to impart control decisions regarding resource allocation by indicating which RBs to blank out for each RU site. Blanked RBs at a site mean those RBs will not be available for use at that site, thereby making them available for use at another site. Essentially, the unblanked RBs are the ones available for use at a particular site or RU.

The communication between EdgeRIC and the RAN occurs over the RT-E2 interface. The RT Report carries information on the RAN state, including pending data and channel quality. The RT E2 policy message consists of the control information, specifically the range of RBs to blank out at a RAN site. The number of RBs to blank depends on the total pending data waiting to be transmitted at each site, which is indirectly a function of the traffic load at the site.

Additionally, situational awareness is crucial for deciding which RBs to allocate to a site. If there is an interfered PRB at a site, it is preferable to avoid transmitting on that PRB. Therefore, we combine the decision on affected PRBs at each site which was derived from the near-realtime RIC database, in conjunction with the information regarding each RAN site pending data and channel quality to determine which how many PRBs to allocate to each site.

3.2 Interference Detection/Localization xApp

To fully leverage the capabilities and benefits of the O-RAN, particularly the near-realtime RIC platform, for multi-site resource sharing, we developed an interference detection and localization algorithm that utilizes spectrograms for both training and inference. This algorithm operates in two steps: first, it processes an image of shape (NxM), where N represents the frequency axis (height) and M represents the time axis (width). Using this image, the model outputs the presence or absence of a jammer or interferring signal, along with the jammer's location dimensions in the spectrum, i.e., $J^{NxM} \rightarrow \mathcal{J}[f_L, f_H, t_L, t_H], \overline{\mathcal{J}}[F_L, F_H, T_L, T_H]$, where $\mathcal{J}$ and $\overline{\mathcal{J}}$ indicate the presence and absence of a jammer signal, respectively. When a jammer is detected, the model returns a list $[f_L, f_H, t_L, t_H]$, where f_L and f_H denote the lower and higher frequencies occupied by the jammer as shown in Figure 4a, and t_L and t_H represent the time axis dimensions. For our purposes, we focus on the values of f_L and f_H to estimate the bandwidth covered by the jammer. Similarly, as illustrated in Figure 4a, F_L and F_H indicate the lower and higher frequencies occupied by the LTE/5G signal.

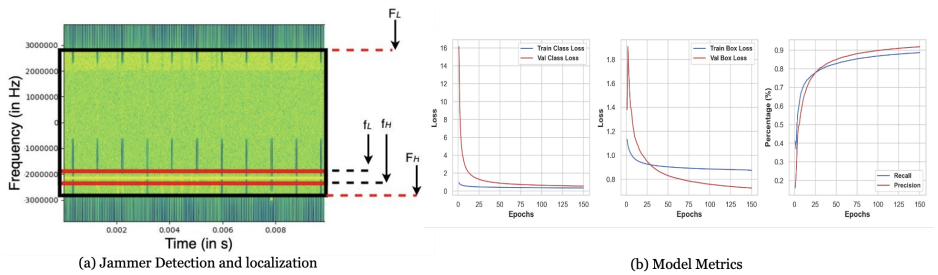

(a) Jammer Detection and localization (b) Model Metrics

Fig. 4. Object Detection for Interference Detection/Localization and Model Training/Validation Metrics.

In the first step, we leverage state-of-the-art deep learning algorithms such as Convolutional Neural Networks (CNNs) and YOLOV8 [19], using the open-source dataset from over-the-air experiments in [10] and [36]. We utilize [4], an open-source data annotation tool, to label a few samples. To increase dataset diversity, we apply data augmentation techniques such as adding Gaussian noise and creating contrast. These augmentations introduce complexity, ensuring our model can detect interfering signals in spectrograms under varying conditions. For the training dataset, we used 2000 images/spectrograms containing information on both the Signal of Interest (SOI) (i.e., the LTE/5G signal) and the jammer in the form of Continuous Wave Interference (CWI) (N.B: Other interference categories can be considered having wider bandwidths and variation in the spectrum). The validation dataset on the other hand consists of 1000 spectrograms.

We use YOLOv8l, a pretrained model with 43.7 million parameters, initially trained on the COCO dataset. For our dataset, we trained the model for 150 epochs with a batch size of 16 and a learning rate of 0.01, which provided optimal accuracy and minimized loss for detecting jammer signals, even with narrow bandwidths. This task is more challenging than typical object detection tasks involving larger objects. Figure 4b shows the training and validation metrics from our dataset.

For detecting and localizing interference, such as Continuous Wave Interference (CWI), it is essential to identify each jammer in the spectrogram while minimizing false detections to avoid unnecessary resource blanking by the μapp. After 150 epochs of training, our model achieves up to 95% precision and 90% recall. Additionally, we focus on the model's accuracy in detecting interference and estimating bounding boxes to correctly identify interference locations. Figure 4b also shows box and class losses for both training and validation, indicating proper learning. There is still potential for further optimization to enhance these results.

After obtaining the values of f_L, f_H, F_L, and F_H, the next step is to estimate which PRBs are affected by the interference signal. This involves considering the original signal parameters, such as channel bandwidth, the number of PRBs for the chosen numerology, and the guard bandwidth. We first determine the dimensions of the SOI from F_L and F_H and then use the guard bands to calculate the actual signal bandwidth in terms of PRBs. Finally, we map the jammer dimensions (f_L and f_H) to the corresponding affected PRBs.

3.3 Multi-site resource distribution on UL Spectrum

3.3.1 Primer on UL scheduling. In LTE, uplink transmissions use Single Carrier Frequency-Division Multiple Access (SC-FDMA), similar to Discrete Fourier Transform - Orthogonal Frequency-Division Multiplexing (DFT-OFDM), to reduce Peak-to-Average Power Ratio (PAPR). SC-FDMA is preferred for uplink due to its lower energy requirements but requires contiguous PRBs. Uplink scheduling starts by searching for a contiguous block of PRBs within the mask. If none are found, it returns an empty set. If a valid interval is found but too short, the system tries to extend it by adding PRBs, or reduces it if necessary, ensuring energy-efficient, compliant transmissions, though it restricts resource allocation to contiguous PRB chunks per site.

3.3.2 Resource blanking. To illustrate the concept of resource blanking, we first consider a scenario without interference. In this case, the weight of a site i (w_i) is determined by the traffic demands at each site, algorithm 1 elaborates on this. The demand is computed from the pending UL data buffers of the UEs at that site, while also taking into consideration the desired service requirements. The total resources allocated to a site are proportional to w_i relative to the total available Resource Blocks (RBs). This can be expressed as: $\texttt{rbs\_site}_i$ = int(round ($w_i \times n\_prb$)) where n_prb is the total RBs available. Based on the required number of resources, we blank out (n_prb - $\texttt{rbs\_site}_i$) RBs at each site i

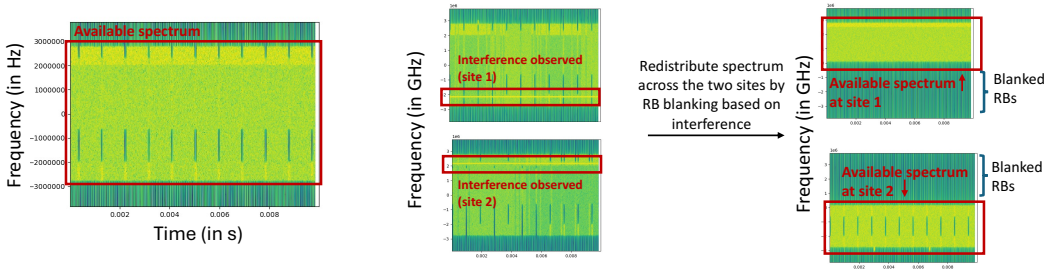

a) Spectrogram over full bandwidth

b) Example Scenario: Spectrogram – in presence of interference

c) Spectrogram – with RB blanking

Fig. 5. Spectrogram visualizations: The light green portion indicates available spectrum

Figure 5 illustrates the concept of Resource Block (RB) blanking for efficient spectrum management across multiple sites in the presence of interference, using spectrograms. In Figure 5(a), the full available spectrum is shown without any interference, providing a baseline view of the total bandwidth that can be used for communication. Figure 5(b) introduces interference at the sites, indicated by red boxes marking the affected frequency bands. This interference reduces the effective bandwidth at each site, limiting the available spectrum without compromising signal quality. Figure 5(c) demonstrates how RB blanking mitigates this interference. The top plot shows site 1 after RB blanking, where the interfered RBs are blanked out and unavailable. The bottom plot shows site 2, where the RBs blanked at site 1 are available, allowing the system to make efficient use of non-interfered frequencies across both sites. This dynamic allocation ensures overall system efficiency despite the interference.

3.3.3 Resource Block Allocation Scheme. In this section, we detail our method for allocating Resource Blocks (RBs) across multiple sites, guided by the interference maps of each site. Our proposed allocation scheme operates in real-time (every TTI) and follows a two-stage approach to efficiently distribute resources while minimizing interference and ensuring demand-aware fairness across sites. Given the complexity of wide-area deployments with multiple small sites, we assume that not all sites can be allocated spectrum during each Transmission Time Interval (TTI). To manage computational resources and maintain real-time performance, we limit the number of sites served in each TTI to three, accounting for the time required to infer the optimal allocation.
Step 1: Computing Site Priority: At each TTI, a priority index is computed for every site to determine the urgency for resource allocation. This priority index is based on the resource demands ('$buffer$' - buffer size) and service violations ('$violation$' - throughput deviation) of the site's UEs (User Equipments). For each UE, the deviation between desired and actual throughput is calculated, and the UE's contribution to the site's priority is derived as the product of its buffer size and the throughput violation. The overall site priority index is determined by the highest UE index at the site, representing the UE with the most critical resource demand. If no UEs at the site have a non-zero demand, the site's index is set to zero. This method ensures that resource allocation is guided by both current traffic conditions and any unmet service requirements. Algorithm 1 outlines the detailed process for computing the priority index at each site.

Algorithm 1 Generate Site Weights Based on UE Demand

procedure GenerateSiteWeights($sites$, $time\_step$)
  for each $site\_id$ in $sites$ **do**
    $indices \leftarrow []$ ▷ List to store UE indices
    $ues \leftarrow sites[site\_id]$ ▷ Get UEs for the site
    for each ue_id in ues **do**
      $buffer \leftarrow sites[site\_id][ue\_id]["buffer"]$
      $actual\_throughput \leftarrow \frac{sites[site\_id][ue\_id]["total\_bytes\_sent"]}{1000 \times time\_step}$
      $desired\_throughput \leftarrow sites[site\_id][ue\_id]["desired\_thrpt"]$
      $violation \leftarrow desired\_throughput - actual\_throughput$
      $index \leftarrow buffer \times violation$ ▷ Normalize the index
      Append $index$ to $indices$
      Log: UE ue_id in site $site\_id$: Index = $index$, Violation = $violation$, Buffer = $buffer$
    end for
    $site\_index \leftarrow \max(indices)$ if $indices \neq \emptyset$ else 0
    Log: Site $site\_id$ has index: $site\_index$
    Output $site\_index$ for site $site\_id$
  end for
end procedure

Step 2: Optimal Resource Distribution: Once the priority indices for each site are computed (as described in Step 1), we proceed with the resource block (RB) allocation. In each TTI, we select the top three sites with the highest priority indices for resource allocation. These indices represent the relative urgency and demand of each site, where higher indices correspond to greater resource needs. Algorithm 2 outlines the RB allocation scheme, which takes two primary inputs: the total number of available Physical Resource Blocks (n_prb) and $sites\_info$. The $sites\_info$ is a dictionary where each site contains the following key attributes: '$weight$' indicates the weight of each site, which is the priority index computed in Step 1, 'bad_rbs' is list of PRBs at the site that are adversely

affected by interference, and '*rbs*' stores the required number of RBs at each site ($\texttt{rbs\_site}_i$). Using the priority weights of the top three selected sites, the algorithm proportionally distributes the available RBs among them based on their weights. Next, the algorithm explores different permutations of possible RB allocations to these sites. For each permutation, it calculates the ranges of PRBs assigned to each site and evaluates the impact of interference by identifying the number of affected PRBs (*bad_rbs*) within these ranges. The optimal allocation is the permutation that results in the least interference, i.e., the one with the fewest affected PRBs. This ensures that the selected sites receive their proportional share of resources while minimizing the performance degradation due to interference. Given that we are dealing with only three sites, we can afford to perform a "brute-force search" with a computational complexity of $O(3!)$. This exhaustive search evaluates all possible RB allocation permutations and ensures optimal resource distribution in real time.

Algorithm 2 Allocation of Physical Resource Blocks (PRBs) to Sites

procedure OptimalAllocateRBs(n_prb, $sites\_info$)
  $sites\_info[site\_key]['weight'] = w_i \leftarrow$ extract weight of each site from $sites\_info$
  $sites\_info[site\_key]['rbs'] = \texttt{rbs\_site}_i \leftarrow$ int(round ($w_i \times n\_prb$))
  $allocations \leftarrow$ all permutations of site keys
  $min\_affected \leftarrow \infty$
  for each *allocation* in *allocations* **do**
    $alloc\_ranges \leftarrow$ GetRBRanges($allocation, sites\_info$)
    $affected \leftarrow$ CalculateAffectedPRBs($alloc\_ranges, sites\_info$)
    if $affected < min\_affected$ **then**
      $min\_affected \leftarrow affected$
      $best\_allocation \leftarrow \{key \rightarrow range$ for each key in $allocation\}$
    end if
  end for
  return *best_allocation*
end procedure
function GetRBRanges($allocation$, $sites\_info$)
  $start \leftarrow 0$
  for each $site\_key$ in $allocation$ **do**
    $end \leftarrow start + sites\_info[site\_key][`rbs`]$
    $sites\_info[site\_key][`alloc\_range`] \leftarrow (start, end - 1)$
    $start \leftarrow end$
  end for
  return $sites\_info$
end function
function CalculateAffectedPRBs($alloc\_ranges$, $sites\_info$)
  $total\_affected \leftarrow$ sum affected PRBs in $alloc\_ranges$
  return $total\_affected$
end function

3.4 Microbenchmarks

In this section, we present microbenchmarks that validate the benefits of our system. Figure 6 summarizes the performance improvements observed with context-aware resource block (RB) blanking. Cases 1 and 2 in Figure 2 represent single RAN site scenarios, while Case 3 demonstrates our proposed system.

Figure 6(a) shows the Cumulative Distribution Function (CDF) of the Signal-to-Noise and Interference Ratio (SINR) for UE1 (a low SNR UE) and UE2 (a high SNR UE located closer to the transmitter) with and without RB blanking in the presence of interference. The results indicate an SNR improvement of 5dB to 7dB when affected PRBs are appropriately blanked, improving the average signal quality. Figure 6(b) illustrates the CDF of the total uplink throughput for three cases: no blanking in a single site (Case 1), blanking applied in a single site (Case 2), and blanking with resource distribution across multiple sites (Case 3). Cases 2 and 3 show a significant throughput improvement (20%) compared to Case 1, with Case 3 achieving the highest throughput, highlighting the advantages of our system. Figure 6(c) shows the average number of packet drops per Transmission Time Interval (TTI) over time. Case 1 has the highest packet drop rate, while Case 2 shows a notable reduction in errors. Case 3 maintains consistently low packet drops, reducing errors by up to 30% compared to the single-site scenario in Case 2.

Overall, the microbenchmarks demonstrate that interference-aware RB blanking across sites improves SNR, system throughput, and packet drop rates, leading to enhanced network performance.

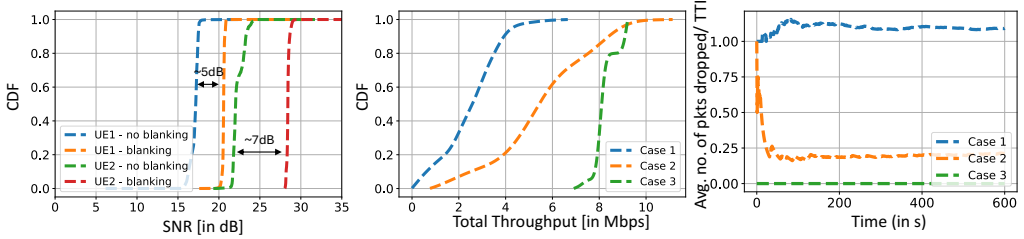

(a) SNR improvements observed (b) Total system throughput on uplink (c) Packet errors observed on uplink

Fig. 6. Microbenchmarks: Interference aware resource distribution across multiple sites (Case 3) does improve network reliability and performance

4 System Implementation

For our over-the-air experiments, we utilized the Open AI Cellular (OAIC) platform [3], designed for prototyping AI-based solutions for next-generation wireless networks. Built on the srsRAN [35] codebase version 21.10, the platform is deployed across desktop computers acting as UEs and base stations. Each desktop runs Ubuntu 20.04, with Intel Core i7-8700 CPUs (6 cores, 12 threads, 3.2GHz), 16GB RAM, and USRP B210 SDRs. The real-time RIC (EdgeRIC) is co-located with the edge DU, and near-realtime RIC, while SDR-based jammers are connected to laptops running GNU Radio. This setup is shown in Figure 7(a).

The near-realtime RIC is hosted on a rack server and has the capacity to serve multiple RANs as shown in Figure 7(a). The server hosting the near-realtime RIC is an AMD EPYC™ 7443P with 24 CPU cores, 48 threads, 64GB RAM and a base clock speed of 2.85GHz. It acts as an intelligent controller for the RAN. The near-realtime RIC interfaces with the RAN via an E2 interface, allowing it to make decisions and control RAN functions based on real-time data and network conditions. The table in Figure 7(b) presents the overall timing for each step described in section 3.1 explaining our system architecture. This clearly shows that spectrum monitoring occurs on a (<1s) timescale. It can be observed that a chunk of the time here is due to the interference detection and localization step which is due to the fact that the model used for this task has 43.7M parameters. The real-time PRB allocation on the other hand are done within sub-millisecond granularity.

For our indoor lab experiment, we operate in Frequency Division Duplex (FDD) mode, focusing on uplink traffic at a 2.56GHz carrier frequency. We use a 5MHz bandwidth (25 PRBs) and stationary UEs and base stations. Traffic is generated via iperf, with the UE at site one transmitting 2Mbps

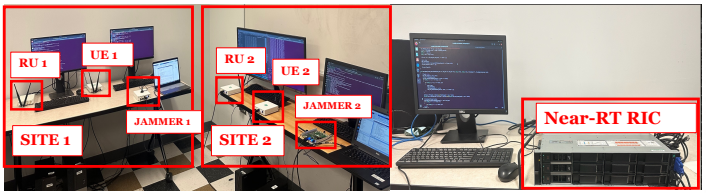

(a) Experiment Setup

Steps	Timings
Near Real-time Monitoring	
Receive I/Q samples	**272.0ms**
Time to forward data for processing	**132.68ms**
Data processing and storing in RIC database	**97.02ms**
Interference detection and Localization	**300.88ms**
Total time	**801.7ms (<1s)**
Real-time Control	
Resource distribution policy compute time	**50μs**
Control decision to RAN	**40μs**
Total time	**<1ms**

(b) Breakdown of system timing of overall processes

Fig. 7. System Implementation on over the air setup

and the UE at site two transmitting 4Mbps, demonstrating resource distribution based on traffic demand, as explained in Section 3.3.

4.1 Over the Air System Benchmarks

In Figure 8, we highlight key observations from our real-world setup to address the question: *Can we improve network/signal quality with our proposed system?*

Figure 8(a) shows the observed SINR under various scenarios. The "vanilla system w/o IF" (blue curve) represents a setup where both sites fully utilize the entire bandwidth, resulting in significant inter-site interference due to the proximity of the RAN sites, leading to lower SINR performance compared to our proposed method (orange curve). By strategically blanking resources, our method improves SINR by 10dB. The "vanilla system w/ IF" (green curve), further degraded by interference, shows even poorer SINR, while our system (red curve) effectively mitigates interference, achieving a 12dB gain in SINR. Though not as high as the orange curve (due to remaining interference at other sites), it significantly improves signal quality compared to the green curve. Figure 8(b) presents packet drop data, inversely correlated with the SINR results. Lower SINR leads to higher packet drops, reinforcing the negative impact of poor signal quality on network reliability.

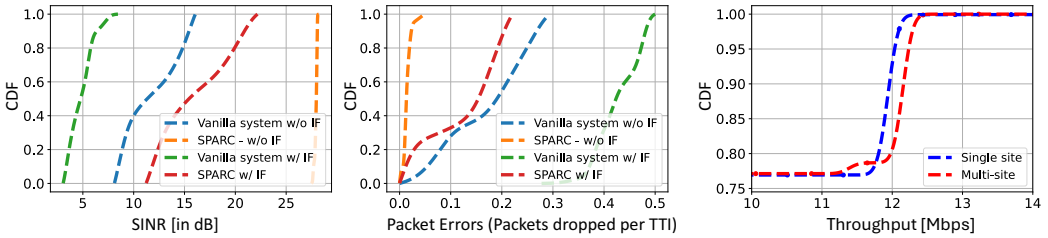

(a) SINR improvements observed

(b) Packet errors observed on uplink

(c) Spectral Efficiency retained with RB blanking

Fig. 8. System Benchmarks: SPARC significantly enhances the network quality

Furthermore, Figure 8(c) addresses the question, *Do we compromise uplink spectral efficiency by resource blanking?* When operating with an equal load on all connected UEs and assuming each UE has a very high uplink SNR, we observe that resource blanking, which redistributes resources across multiple sites, achieves spectral efficiency comparable to that of the traditional single-site scenario. In fact, by increasing the number of transmitters through a multi-site approach, we effectively enhance the uplink SNR for all UEs, thereby potentially increasing spectral efficiency.

5 System Evaluations

In this section we provide a comprehensive evaluation of our system performance in a wide range of traffic and interference scenarios, trying to specifically answer the following questions, *(i) Is demand based real-time resource distribution across multiple sites useful?* and *(ii) Does spectrum*

aware resource distribution offer enhanced system behaviour? . We specifically compare our proposed system against the following allocation schemes to establish the benefits of **real-time** resource distribution across **multiple sites**:

Equal Allocation: A fixed and equal number of resources is reserved for each site.

Proportional Allocation: Operating on a coarse granularity timescale (approximately 500ms), this scheme calculates the total bitrate observed at each site over the last 500 slots. Based on this data, it proportionally distributes the spectrum across each site, updating the number of available PRBs every 500 Transmission Time Intervals (TTIs).

Single RAN site: We also compare our proposed multi site system performance with a traditional single RAN environment.

Fairness based Allocation: In Section 5.3, we present an additional comparison between our SPARC resource distribution strategy and a fairness-based allocation scheme. In the fairness-based scheme, the weight of each site, $w_i = \sum_{j \in \text{UEs in site } i} \texttt{buffer}_j \; / \; \sum_{j \in \text{UEs in site } i} (\texttt{total\_bytes\_sent}_i)$.

In Table 1, we detail various traffic scenarios to assess our system, which comprises four UEs connected to two RUs within the cellular network. Specifically, UE1 and UE2 are connected to RU1, while UE3 and UE4 connect to RU2. It is presumed that all UEs are proximate to their transmitters, thereby benefitting from strong uplink channels. Scenarios 1 through 5 (Sc 1-5) employ iperf as the traffic generator. In contrast, Scenario 6 (Sc 6) utilizes a custom traffic generator designed to simulate different traffic profiles. For instance, the Enhanced Mobile Broadband (eMBB) and Extended Reality (XR) scenarios involve periodic traffic, where packets are generated at fixed intervals—akin to video frame rates. The eMBB scenario is configured for a traffic load of 3 Mbps, while the XR scenario supports 5 Mbps. Additionally, the Ultra-Reliable Low-Latency Communications (URLLC) traffic flow is characterized by bursty patterns, with random bursts of 3-5 MB occurring sporadically.

Table 1. Summary of all scenarios

Scenarios	**Traffic profile for the connected UEs**
Static Traffic	
Sc 1	iperf - UE1: 3Mbps, UE2: 8Mbps, UE3: 2Mbps, UE4: 9Mbps
Sc 2	iperf - UE1: 7Mbps, UE2: 1Mbps, UE3: 7Mbps, UE4: 1Mbps
Sc 3	iperf - UE1: 7Mbps, UE2: 7Mbps, UE3: 1Mbps, UE4: 1Mbps
Sc 4	iperf - UE1: 3Mbps, UE2: 6Mbps, UE3: 0.01Mbps, UE4: 0.01Mbps
Dynamic Traffic	
Sc 5	iperf traffic: offered load randomly changes b/w 2-9 Mbps every 1s for all UEs
Sc 6	Custom traffic: UE1: embb, UE2: urllc, UE3: XR, UE4: urllc

5.1 Impact of real-time demand-based multi site resource distribution

In this subsection, we present evaluations (Figure 9) to address our first research question: Is demand-based, real-time resource distribution across multiple sites beneficial?

Figure 9(a) illustrates scenarios with static traffic profiles, where the traffic offered to each UE remains constant throughout our experiments. Scenarios 1 and 2 (Sc1 and Sc2) involve cases where the total traffic demand at both sites is equivalent. In these scenarios, we observe that performance under the Equal Allocation (Equal Alloc.) and Proportional Allocation (Prop. Alloc.) schemes is comparable. However, we observe that the total amount of data pending for transmission can vary significantly at any given moment. By adopting a real-time approach that adjusts resource distribution based on the instantaneous total data pending at each site, we can surpass the performance of both the Equal and Proportional Alloc. schemes. Similar patterns are evident in Scenarios 3 and 4

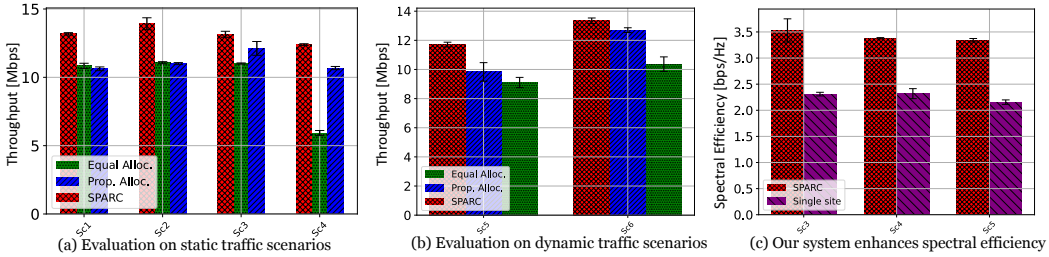

Fig. 9. Real-time resource distribution across multiple sites via SPARC enhances the total system throughput

(Sc3 and Sc4), where the traffic demand differs between the sites—site1 requires more resources than site2. In these cases, Equal alloc results in poorer performance, while PF moderately improves resource allocation by adjusting to the traffic at a coarser timescale. Nonetheless, our approach, which makes instantaneous resource distribution decisions, consistently achieves superior throughput, thereby validating its effectiveness in enhancing system performance. SPARC is able to support at least 20% higher system throughput.

Figure 9(b) displays results for scenarios (Sc 5 and Sc 6) where traffic flow varies among the UEs throughout the duration of our experiments. These scenarios further confirm that real-time decision-making enhances system performance, providing optimal outcomes even under fluctuating traffic conditions.

In Figure 9(c), we explore how our multi-site system stacks up against traditional single-site configurations. By strategically positioning transmitters or RUs closer and more densely around UEs, we significantly improve the uplink SNR. This setup increases the achievable bitrate and the Modulation and Coding Scheme (MCS), thus substantially boosting overall system throughput.

5.2 Impact of interference-aware demand-based multi site resource distribution

In this subsection, we provide evaluations to address our second question: Does spectrum-aware resource distribution enhance system behavior? Given the critical role of spectrum awareness in identifying external interference—such as jammers that can severely disrupt operations—we present our results in Figure 10, which illustrates how our system performs under various traffic profiles in the presence of interference.

We introduce frequency-hopping interferers in our system, specifically single-tone jammers that transmit randomly across various frequencies within our spectrum of interest. Leveraging the interference detection and localization xApp, which operates within the near-realtime RIC, we can accurately detect these interfering frequencies. This detection allows us to strategically avoid these frequencies at each site. EdgeRIC is then updated about the interfered or compromised PRBs at each site, enabling it to judiciously select the parts of the spectrum to allocate per site.

Figure 10(a) highlights the throughput benefits of our proposed multi-site system in the presence of interference across various traffic profiles. SPARC support at least 25% higher throughput in all scenarios. Figure 10(b) corroborates these benefits by showing the improved uplink average SINR achieved when the system efficiently avoids the compromised frequencies. Finally, Figure 10(c) demonstrates how our system significantly reduces packet drops by steering clear of the bad channel RBs, thereby potentially lowering overall latency by eliminating the need for retransmissions. SPARC is able to offer near to zero percent packet drops.

5.3 Scalability Case study

In this section, we present results from a scaled-up evaluation of our system. Due to hardware limitations, which prevent us from scaling beyond two sites, we rely on simulation-based assessments for bench marking. Specifically, we evaluate SPARC's real-time resource distribution scheme

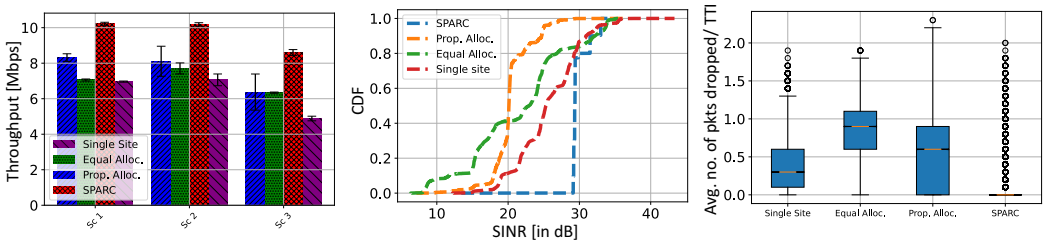

(a) SPARC supports a higher throughput (b) SPARC improves overall SINR (c) SPARC reduces packet drops

Fig. 10. Summary of system benefits realized by SPARC under interference

across multiple sites, dynamically adjusting based on demand. Our experiments show that SPARC consistently delivers throughput aligned with system capacity across all scenarios by employing an adaptive, real-time approach. In contrast, other schemes fail to reach the achievable system throughput. Figure 11(a) outlines the different scenarios tested, while Figure 11(b) illustrates the total system throughput attained in each case. Additionally, Figure 11(c) highlights SPARC's ability to minimize service demand violations at each site, focusing on Scenario 2, where the total system demand is within feasible limits. SPARC successfully meets site-specific demands through real-time resource reallocation, outperforming alternative approaches that either distribute resources equally or rely on coarser timescale redistribution schemes.

Scenario	Site	#UEs	Traffic Demand
Sc 1	site 1	2UEs	14Mbps
	site 2	2UEs	16Mbps
	site 3	2UEs	3Mbps
	site 4	2UEs	3Mbps
	site 5	2UEs	4Mbps
Sc 2	site 1	1UEs	10Mbps
	site 2	4UEs	12Mbps
	site 3	2UEs	9Mbps
	site 4	2UEs	5Mbps
	site 5	2UEs	5Mbps
Sc 3	site 1	3UEs	22Mbps
	site 2	3UEs	26Mbps
	site 3	3UEs	9Mbps
	site 4	3UEs	7Mbps
	site 5	3UEs	6Mbps

(a) Summary of all scenarios

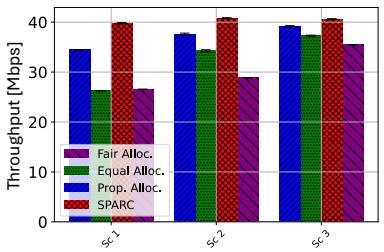

(b) SPARC can achieve the system capacity

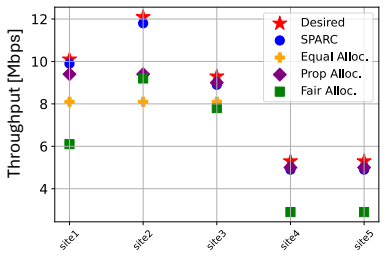

(c) SPARC can offer the least service violations - Results demonstrated for scenario 2

Fig. 11. Summary of evaluations in a 5-site system

6 Conclusion

In this paper, we presented a multi-site cellular network system that integrates near real-time spectrum monitoring and spectrum-aware resource distribution. Our experiments addressed two key questions: First, we showed that real-time, demand-based resource distribution significantly boosts system throughput by dynamically adjusting resource allocation based on traffic demands across multiple sites. Second, spectrum-aware resource management proved effective in mitigating external interference, such as jammers. By leveraging near-realtime RIC for detection using an xApp and EdgeRIC for spectrum control using a μApp, our system maintained high throughput and improved SINR, reducing packet drops and enhancing network reliability and latency.

Acknowledgement: Authors acknowledge the funding support from NSF CCRI 2120411, and NSF Grants CNS 2312978 and CNS 2312979. All finding and opinions are of the authors.

Ethical concerns: This work does not raise any ethical issues.

References

[1] 2014. Review of channel quality indicator estimation schemes for multi-user MIMO in 3GPP LTE/LTE-A systems. *KSII Trans. Internet Inf. Syst.* 8, 6 (June 2014), 1848–1868.

[2] 2023. Nvidia Aerial. (2023). https://developer.nvidia.com/aerial
[3] 2023. Open AI Cellular. (2023). https://github.com/openaicellular
[4] 2024. Open Data Annotation Platform. https://www.cvat.ai/. (2024).
[5] 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.* Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3570361.3615745
[6] Luca Baldesi, Francesco Restuccia, and Tommaso Melodia. 2022. ChARM: NextG Spectrum Sharing Through Data-Driven Real-Time O-RAN Dynamic Control. In *IEEE INFOCOM 2022 - IEEE Conference on Computer Communications*. IEEE Press, 240–249. https://doi.org/10.1109/INFOCOM48880.2022.9796985
[7] Richard Bell, Kyle Watson, Tianyi Hu, Isamu Poy, fred harris, and Dinesh Bharadia. 2023. Searchlight: An accurate, sensitive, and fast radio frequency energy detection system. 397–404. https://doi.org/10.1109/MILCOM58377.2023.10356230
[8] Fransiscus Asisi Bimo, Ferlinda Feliana, Shu-Hua Liao, Chih-Wei Lin, David F. Kinsey, James Li, Rittwik Jana, Richard Wright, and Ray-Guang Cheng. 2022. OSC Community Lab: The Integration Test Bed for O-RAN Software Community. (2022). arXiv:cs.NI/2208.14885
[9] Lianning Cai, Kaitian Cao, Yongpeng Wu, and Zhou Yuan. 2022. Spectrum Sensing Based on Spectrogram-Aware CNN for Cognitive Radio Network. *IEEE Wireless Communications Letters* 11 (10 2022), 1–1. https://doi.org/10.1109/LWC.2022.3194735
[10] Azuka Chiejina, Brian Kim, Kaushik Chowhdury, and Vijay K Shah. 2024. System-level Analysis of Adversarial Attacks and Defenses on Intelligence in O-RAN based Cellular Networks. In *Proceedings of the 17th ACM Conference on Security and Privacy in Wireless and Mobile Networks*. 237–247.
[11] Salvatore D'Oro, Michele Polese, Leonardo Bonati, Hai Cheng, and Tommaso Melodia. 2022. dApps: Distributed Applications for Real-Time Inference and Control in O-RAN. *IEEE Communications Magazine* 60, 11 (2022), 52–58. https://doi.org/10.1109/MCOM.002.2200079
[12] Xenofon Foukas, Navid Nikaein, Mohamed M. Kassem, Mahesh K. Marina, and Kimon Kontovasilis. 2016. FlexRAN: A Flexible and Programmable Platform for Software-Defined Radio Access Networks. In *Proceedings of the 12th International on Conference on Emerging Networking EXperiments and Technologies (CoNEXT '16)*. Association for Computing Machinery, New York, NY, USA, 427–441. https://doi.org/10.1145/2999572.2999599
[13] Xenofon Foukas, Bozidar Radunovic, Matthew Balkwill, and Zhihua Lai. 2023. Taking 5G RAN Analytics and Control to a New Level *(ACM MobiCom '23)*. Association for Computing Machinery, New York, NY, USA, Article 1, 16 pages. https://doi.org/10.1145/3570361.3592493
[14] Xenofon Foukas, Bozidar Radunovic, Matthew Balkwill, Zhihua Lai, and Connor Settle. 2023. *Programmable RAN Platform for Flexible Real-Time Control and Telemetry.* Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3570361.3614065
[15] Sneihil Gopal, David Griffith, Richard A Rouil, and Chunmei Liu. 2024. ProSAS: An O-RAN Approach to Spectrum Sharing between NR and LTE. *arXiv preprint arXiv:2404.09110* (2024).
[16] Agrim Gupta, Adel Heidari, Jiaming Jin, and Dinesh Bharadia. 2024. Densify & Conquer: Densified, smaller base-stations can conquer the increasing carbon footprint problem in nextG wireless. (2024). arXiv:cs.NI/2403.13611
[17] Agrim Gupta, Ish Jain, and Dinesh Bharadia. 2022. Multiple smaller base stations are greener than a single powerful one: Densification of wireless cellular networks. In *1st Workshop on Sustainable Computer Systems Design and Implementation (HotCarbon)*.
[18] Xiaolin Hou and Hidetoshi Kayama. 2011. *Demodulation reference signal design and channel estimation for LTE-Advanced uplink.* INTECH Open Access Publisher.
[19] Jacob Solawetz and Francesco. 2024. What is YOLOv8? The Ultimate Guide.[2024]. https://blog.roboflow.com/whats-new-in-yolov8/. (2024).
[20] 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
[21] Chang Liu, Gopalasingham Aravinthan, Ahan Kak, and Nakjung Choi. 2023. *TinyRIC: Supercharging O-RAN Base Stations with Real-time Control.* Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3570361.3615743
[22] Mavenir, Inc. 2022. Mavenir, Inc. https://www.mavenir.com/. (2022).
[23] Jorge Navarro-Ortiz, Pablo Romero-Diaz, Sandra Sendra, Pablo Ameigeiras, Juan J. Ramos-Munoz, and Juan M. Lopez-Soler. 2020. A Survey on 5G Usage Scenarios and Traffic Models. *IEEE Communications Surveys Tutorials* 22, 2 (2020), 905–929. https://doi.org/10.1109/COMST.2020.2971781

[24] News Article 2013. Intentional jamming by truck driver at new york airport. https://www.cnet.com/culture/truck-driver-has-gps-jammer-accidentally-jams-newark-airport/. (2013).
[25] News Article 2020. old tv set interferes with connectivity:. https://arstechnica.com/information-technology/2020/09/old-tv-set-interfered-with-villages-dsl-internet-each-day-for-18-months/. (2020).
[26] Navid Nikaein, Mahesh K. Marina, Saravana Manickam, Alex Dawson, Raymond Knopp, and Christian Bonnet. 2014. OpenAirInterface: A Flexible Platform for 5G Research. *SIGCOMM Comput. Commun. Rev.* 44, 5 (oct 2014), 33–38. https://doi.org/10.1145/2677046.2677053
[27] ORAN Alliance Feb, 2019. O-RAN use cases and deployment scenarios. Tech Republic.. (Feb, 2019).
[28] ORAN Alliance March, 2019. O-RAN working group 2: AI/ML workflow description and requirements. Tech Republic.. (March, 2019).
[29] Michele Polese, Leonardo Bonati, Salvatore D'Oro, Stefano Basagni, and Tommaso Melodia. 2023. Understanding O-RAN: Architecture, Interfaces, Algorithms, Security, and Research Challenges. *Commun. Surveys Tuts.* 25, 2 (apr 2023), 1376–1411. https://doi.org/10.1109/COMST.2023.3239220
[30] Radisys, Inc. 2022. Radisys, Inc. https://www.radisys.com/. (2022).
[31] Guillem Reus-Muns, Pratheek S Upadhyaya, Utku Demir, Nathan Stephenson, Nasim Soltani, Vijay K Shah, and Kaushik R Chowdhury. 2023. Senseoran: O-RAN based radar detection in the cbrs band. *IEEE Journal on Selected Areas in Communications* (2023).
[32] Shamik Sarkar, Milind Buddhikot, Aniqua Baset, and Sneha Kumar Kasera. 2021. DeepRadar: a deep-learning-based environmental sensing capability sensor design for CBRS. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking (MobiCom '21)*. Association for Computing Machinery, New York, NY, USA, 56–68. https://doi.org/10.1145/3447993.3448632
[33] 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
[34] Nasim Soltani, Vini Chaudhary, Debashri Roy, and Kaushik Chowdhury. 2022. Finding Waldo in the CBRS Band: Signal Detection and Localization in the 3.5 GHz Spectrum. In *GLOBECOM 2022 - 2022 IEEE Global Communications Conference*. 4570–4575. https://doi.org/10.1109/GLOBECOM48099.2022.10001638
[35] srsRAN, Inc. 2023. srsRAN, Inc. https://www.srslte.com/. (2023).
[36] Nathan H. Stephenson, Azuka J. Chiejina, Nathaniel B. Kabigting, and Vijay K. Shah. 2023. Demonstration of Closed Loop AI-Driven RAN Controllers Using O-RAN SDR Testbed. In *MILCOM 2023 - 2023 IEEE Military Communications Conference (MILCOM)*. 241–242. https://doi.org/10.1109/MILCOM58377.2023.10356330
[37] Matteo Varotto, Stefan Valentin, and Stefano Tomasin. 2024. Detecting 5G Signal Jammers Using Spectrograms with Supervised and Unsupervised Learning. (2024). arXiv:2405.10331
[38] Yujia Zhang, Guanlin Chen, Wenyong Weng, and Zebing Wang. 2010. An overview of wireless intrusion prevention systems. In *2010 Second International Conference on Communication Systems, Networks and Applications*, Vol. 1. 147–150. https://doi.org/10.1109/ICCSNA.2010.5588671
[39] Yixuan Zhang and Zhongqiang Luo. 2023. A Review of Research on Spectrum Sensing Based on Deep Learning. *Electronics* 12, 21 (2023). https://doi.org/10.3390/electronics12214514
[40] Jim Zyren and Wes McCoy. 2007. Overview of the 3GPP long term evolution physical layer. *Freescale Semiconductor, Inc., white paper* 7 (2007), 2–22.

Received June 2024; revised September 2024; accepted October 2024