DEMO: SPARC: Adaptive Resource Control for Multi-site Spectrum Management in NextG Networks

Ushasi Ghosh[1], Azuka Chiejina[2], Nathan Stephenson[2], Vijay Shah[3], Srinivas Shakkottai[4], Dinesh Bharadia[1]
[1]UC San Diego, [2]George Mason University, [3]NC State University, [4]Texas A&M University

ABSTRACT

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-real-time 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%.

CCS CONCEPTS

• **Networks** → **Resource management**; • **Interference** → **Spectrum management**;

KEYWORDS

RAN Intelligent Control, Interference detection, traffic

ACM Reference Format:
Ushasi Ghosh[1], Azuka Chiejina[2], Nathan Stephenson[2], Vijay Shah[3], Srinivas Shakkottai[4], Dinesh Bharadia[1] [1]UC San Diego, [2]George Mason University, [3]NC State University, [4]Texas A&M University. 2024. DEMO: SPARC: Adaptive Resource Control for Multi-site Spectrum Management in NextG Networks. In *The 30th Annual International Conference on Mobile Computing and Networking (ACM MobiCom '24), November 18–22, 2024, Washington D.C., DC, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3636534.3701547

ACM MobiCom '24, November 18–22, 2024, Washington D.C., DC, USA

ACM ISBN 979-8-4007-0489-5/24/11.
https://doi.org/10.1145/3636534.3701547

1 INTRODUCTION

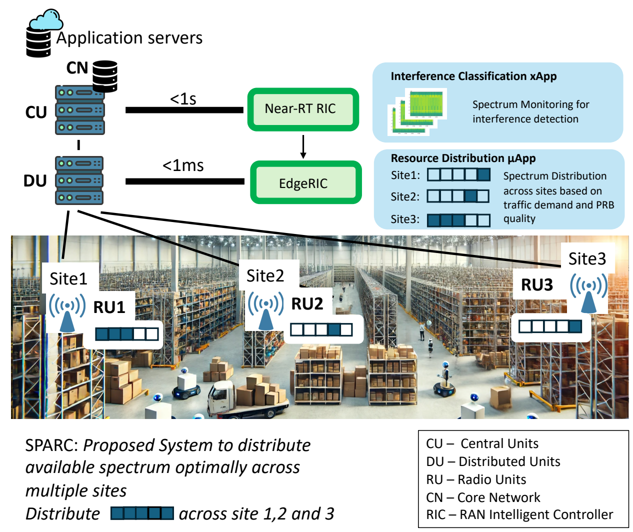

Figure 1: System and Network Overview

The rapid growth of next-generation applications, such as industrial IoT and AR/VR, places significant demands on data throughput, latency, and reliability in modern communication networks. These applications increasingly rely on 5G technology, extending beyond traditional wide area networks to private 5G networks in spaces like large enterprises and warehouses. A critical challenge in these deployments is the limited availability of licensed spectrum, requiring efficient frequency reuse strategies across cell sites, as shown in Figure 1. This paper introduces SPARC (Spatio-Temporal Adaptive Resource Control), a novel approach to multi-site spectrum management in NextG cellular networks. SPARC leverages the Open Radio Access Network (O-RAN) architecture to address spectrum limitations by enabling dynamic monitoring and control of radio access networks across different timescales. Our approach centers on developing a multi-timescale RAN Intelligent Controller (RIC) framework, featuring an xApp for near-real-time (< 1s) interference detection and localization and a μApp for real-time (<1 ms) intelligent resource allocation to the sites. 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, thus enhancing signal-to-noise ratio (SNR), and overall system throughput, making it a robust solution for optimizing resource usage efficiency and network performance.

SPARC employs multiple Radio Units (RUs) across different sites, leveraging disaggregated cellular architectures. This, 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.

2 SPARC: SYSTEM OVERVIEW

SPARC introduces several key innovations in multi-site spectrum management. It employs a multi-timescale RIC approach that enhances spectrum management by improving information sharing and joint optimization between near-real-time and real-time RICs, leading to an SNR gain of up to 7dB. Additionally, SPARC uses base stations as spectrum sensors, with an xApp that utilizes object detection to identify and localize interference, providing critical insights into affected Physical Resource Blocks (PRBs). The framework also includes a μApp for real-time intelligent resource allocation across RAN sites, enhancing spectral efficiency by up to 15% through resource block blanking.

In summary, SPARC offers a comprehensive solution for improving energy efficiency and network performance in next-gen cellular networks by integrating innovative spectrum management and intelligent resource control, adopting a multi-timescale monitoring and control strategy for optimized system behavior.

■ **Spectrum Sensing: A Near Real-time Approach**: We implement spectrum sensing for base stations as an xApp—a microservice for monitoring, detecting, and localizing spectrum within the near-RT RIC. Raw I/Q samples from RAN sites are collected by the E2 agent via the standardized E2 interface and forwarded to the RIC's policy controller, stored in separate buffers per site. Periodically, the last 10ms (one LTE/5G frame) of I/Q samples, reducible to 5ms for lower latency, are sent to a data processing microservice that converts them into spectrograms. These spectrograms are stored in a Redis-based database within the near-RT RIC, accessible to any xApp. The interference detection and localization xApp retrieves these spectrograms and uses a machine learning model, trained on a spectrograms dataset generated in [2], to detect interference signals and identify affected Physical Resource Blocks (PRBs). The inference results are then stored in the Redis-based database for access by other RICs, microservices, or xApps.

■ **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. The real time RIC - EdgeRIC [3] 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 (Real time 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-RT RIC database, in conjunction with the information regarding each RAN site's pending data and channel quality to determine which and how many PRBs to allocate to each site.

■ **Summary of Results**: We provide evaluations to address our 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 2, which illustrates how our system performs under various traffic profiles in the presence of interference. We introduce frequency-hopping interferers in our system, 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-RT RIC, we can accurately detect the presence of the interfering signals. This detection allows us to strategically avoid the PRBs affected by the interfering signals at each site. EdgeRIC is then updated about the interfered PRBs at each site, enabling it to judiciously select the parts of the spectrum to allocate per site.

Figure 2(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 2(b) corroborates these benefits, showing the improved average uplink SINR achieved

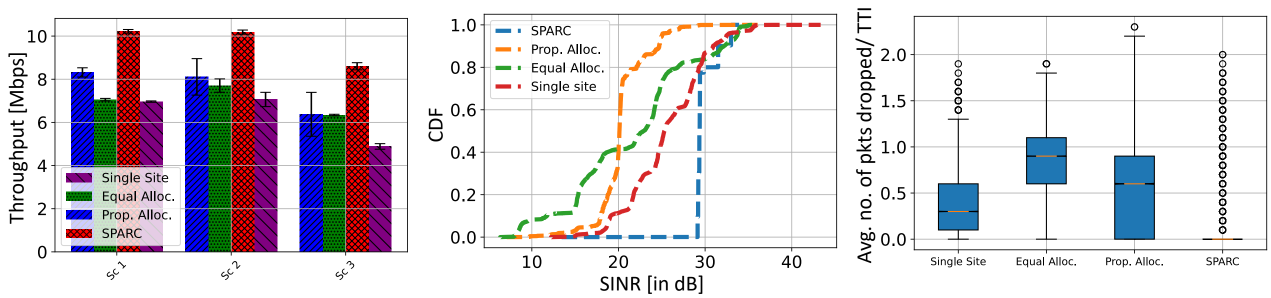

(a) SPARC supports a higher throughput (b) SPARC improves overall SINR (c) SPARC reduces packet drops

Figure 2: Summary of system benefits realized by SPARC under interference

when the system efficiently avoids the compromised frequencies. Finally, Figure 2(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.

3 SETUP DETAILS

Our setup was designed using the Open AI Cellular (OAIC) platform [1] which is a platform developed for prototyping and testing AI based solutions for next-generation wireless networks. This platform is built on top of the srsRAN [4] codebase version 21.10 hosted on the different desktop computers for the UEs and base stations. Each desktop is equipped with an ubuntu release 20.04 OS and running on an intel core i7-8700 having 6 CPU cores, 16GB RAM, 12 threads and running at a clock speed of 3.2GHz. Each desktop also has USRP B210 Software Defined Radios (SDRs) connected to them. We also have the real-time RIC (EdgeRIC) co-located with the edge distributed unit (DU), a near-RT RIC, and SDR based jammers connected to two different laptop computers running GNU radio. This system setup is shown in Figure 3.

The near-RT RIC is hosted on a rack server and has the capacity to serve multiple RANs as shown in Figure 3. The server hosting the near-RT 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-RT 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.

For our experiement which is done in an indoor lab setting, we are operating in the Frequency Division Duplex (FDD) mode and considering the uplink traffic direction from the UEs to the base station operating on the 2.56GHz carrier frequency. We utilize a total of 5MHz bandwidth configuration for analysis which is equivalent to having 25 PRBs. Both base stations and UEs are all stationary for simplicity reasons. For the traffic, we generate different uplink traffic load in the uplink direction for the UEs at different sites using iperf. The UE connected to site one generates 2Mbps of traffic in the uplink direction while the UE connected to the second site generates 4Mbps of traffic in the uplink direction. This is to show the effect of performing resource distribution to different sites based on traffic demand. We would use this setup for our demonstration.

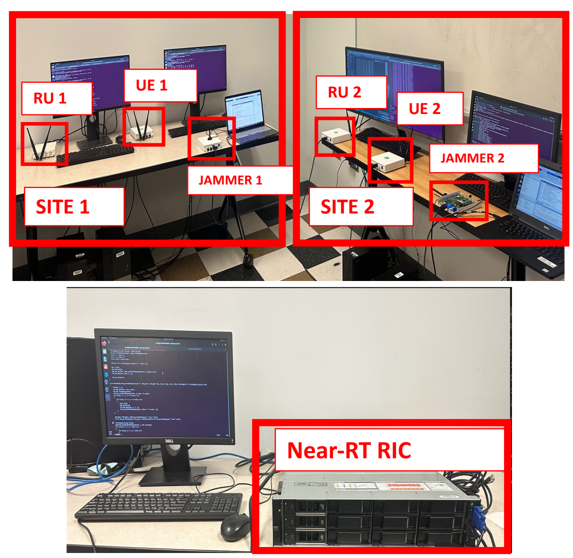

Figure 3: System Implementation on over the air setup

REFERENCES

[1] 2023. Open AI Cellular. (2023). https://github.com/openaicellular
[2] 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.
[3] 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
[4] srsRAN, Inc. 2023. srsRAN, Inc. https://www.srslte.com/. (2023).