

B_oTM: Basestation-on-the-move, a Radio Access Network Management Primitive

Aashish Gottipati
School of Computing
University of Utah

Salt Lake City, United States of America
u1174579@utah.edu

Jacobus Van der Merwe
School of Computing
University of Utah

Salt Lake City, United States of America
kobus@cs.utah.edu

Abstract—Software-defined Radio Access Networks (SD-RANs) enable unparalleled flexibility and the opportunity to customize and/or optimize network operations. In particular, network function virtualization suggests that network functions are no longer tightly coupled to any specific hardware instance or location. In this paper we present our work on B_oTM, *basestation-on-the-move*, where we exploit network function virtualization in an SD-RAN environment to dynamically “move” a mobile base station from one location to another to realize a general RAN management primitive. Specifically, we show how an SD-RAN environment enables orchestration across both the virtual infrastructure layer (i.e., the base station) as well as the mobile network functional layer (i.e., the protocol interactions in a mobile network), allowing a base station *and* its associated endpoints to be moved, fostering new network management functionality. We present the design and implementation of this network management primitive using Open Air Interface and the FlexRAN framework, with experimental results and efficiency metrics.

I. INTRODUCTION

The “softwarization” wave has changed the way wired networks are realized, managed, and operated, emerging as a key paradigm for mobile and wireless networks as well. Specifically, the “softwarization” phenomenon has revolutionized the radio access network (RAN). These software-defined radio access networks (SD-RANs) include virtualization and programmability capabilities, promise to enable unparalleled flexibility, and provide the opportunity to customize and/or optimize network operations [1]–[9].

In this paper we present our work on B_oTM, *basestation-on-the-move*, where we exploit network function virtualization in an SD-RAN environment to dynamically “move” a mobile base station from one location to another to realize a general RAN management primitive. Specifically, we show how an SD-RAN environment enables orchestration across both the virtual infrastructure layer (i.e., the base station) as well as the mobile network functional layer (i.e., the protocol interactions in a mobile network), allowing a base station and its associated endpoints to be moved, fostering new network management

This work was supported by the National Science Foundation (NSF) through a Research Experience for Undergraduates (REU) program at the University of Utah as part of NSF award CNS-1827940, and by an Undergraduate Research Opportunities Program (UROP) award from the University of Utah.

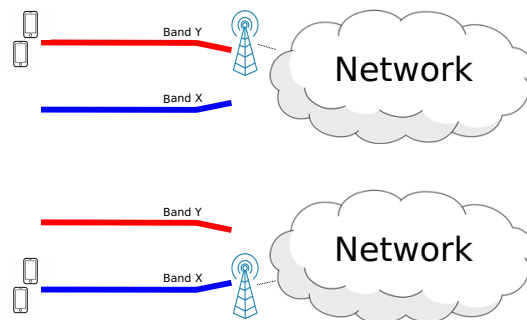


Fig. 1. B_oTM concept: In this case the basestation is “moved”, together with its associated endpoint devices, from using one frequency band to another.

functionality (Figure 1). We have validated our design using Open Air Interface (OAI) [10] and FlexRAN [8], and we have provided the flexibility for our design to generalize to more sophisticated frameworks such as O-RAN [2], [4].

Motivation. With 5G promising support for more devices and more types of communication, lower levels of latency, and optimized user experience, ensuring a high level of quality of service (QoS) and quality of experience (QoE) becomes absolutely essential [11]. Specifically, with the adoption of mmWave technology, maintaining adequate QoS and QoE levels becomes even more challenging due to the nature of attenuation and signal loss at the high frequency ranges [12]. To illustrate, a typical office worker may spend a large amount of time in-doors. This closed-indoor environment increases the difficulty for mmWaves to propagate to a gNodeB base station, which, depending on the environment, can lead to a complete loss in service [13].

In addition, with 5G integrating more and more types of devices and users into the mobile network, the need for efficient spectrum and RAN resource utilization becomes imperative [14], [15]. Challenges with enforcing fine-grained spectrum utilization arise from the chaotic fluctuations in spectrum usage; however, a subset of these fluctuations are directly dependent on time [16]. Given this time dependency, spectrum management can be fine-tuned based on the environment and the time of day, leading to more efficient spectrum utilization. In addition, by leveraging RAN virtualization techniques, base stations can be dynamically instantiated and released, allocating or freeing their resources respectively [10], [17].

For example, during morning work commutes, there is a high likelihood that more users will be concentrated along the main street roads. Given this assumption, more base stations along the main-street path can be instantiated during the most active commute hours, while base station instances with little to no traffic can be released. These base station instances can be tuned to match the optimum frequency for the time of day, enhancing utilization in both the RAN resource and spectrum domains.

Lastly, to foster high levels of QoE, QoS, and resource utilization, a high level of security within the 5G ecosystem is required [18]. With the integration of massive industrial machines, autonomous vehicles, and internet-of-things (IoT) devices, many time sensitive and critical communications will be transmitted over 5G, meaning adverse repercussions exist and are likely to happen, if the mobile network infrastructure becomes compromised. However, mobile infrastructure attacks that target specific patches of spectrum, e.g. spectrum jamming, can be mitigated by dynamically migrating to a new piece a spectrum [19], improving the overall security of the network.

Design principles. Given these motivations, we develop and adhere to the following design principles for BoTM :

- BoTM must be a general network management primitive and must be realizable via emerging RAN management frameworks.
- BoTM 's migration must be transparent, meaning that the migration results in little to no impact on the user equipment's (UE) QoS and QoE.

Contributions. We make the following contributions:

- The design of BoTM , a realizable RAN migration design that incorporates various network abstractions to increase generalizability, ensuring seamless integration with other RAN management frameworks.
- Prototype implementation and evaluation of BoTM using OAI [10] and FlexRAN [8] is shown to produce negligible overhead. From an application viewpoint, seamless communication with the RAN is ensured, meaning that BoTM is nearly imperceptible to user equipment.
- Quantitative results demonstrating the benefits of our procedure by highlighting its capabilities.

The rest of the paper is organized as follows: Section II provides the necessary background information on RAN programmability to understand the key enablers of our system, Section III details our system design and key implementation decisions, Section IV highlights our system's capability within the context of RAN interference management, Section V describes future work and we conclude with Section VI.

II. BACKGROUND AND RELATED WORK

RAN programmability has been an active area of research in recent years. RAN programmability builds on two key network paradigms— Network Function Virtualization (NFV) and Software-Defined Networking (SDN). NFV enables the decoupling of network functionality from hardware by defining

network functionality in software, which can then be run on commodity hardware [20]. On the other hand, SDN focuses on the separation of the control plane and data plane of the network, focusing on how packets flow through the network [21]. RAN programmability refers to the extension of these two key paradigms to the RAN. With these two enablers, the notion of RAN management emerges; i.e., the ability to supervise the life cycle of an SD-RAN instance.

Furthermore, dynamic RAN management frameworks have garnered more and more attention over the past decade. Static RAN management system implementations such as FLARE [1] became the precursor to more efficient dynamic RAN management system proposals [5]. Early iterations of dynamic RAN management systems focused mostly on near real-time control functions such as media access control (MAC) scheduling, meaning that early SD-RAN implementations were limited to configuring and controlling the RAN [8]. Within recent years, modern RAN management systems have evolved to enable RAN orchestration, supporting non-real time control functions such as instantiating a RAN topology and configuring RAN policies [2], [3], [6], [7]. The BoTM design aims to extend and exploit the functionality of these modern RAN management frameworks.

State of the art RAN management systems implement various techniques to ensure the efficient and dynamic sharing of RAN resources. While current RAN management systems may distribute resources efficiently, these systems may not maximize resource utilization. For example, Heterogeneous Networks (HetNets) have recently been proposed to increase cell coverage and meet increased data rate demands in next generation networks. HetNets are composed of a variety of radio access technologies (RAT) that operate within both large cells and small cells. By adopting a variety of RATs and cell sizes, cell coverage and data rates can be improved by mitigating potential RAT failures with the ability to fall back to other nearby RATs within the network [22]; however, by utilizing more RATs and more cells, the potential for communication interference increases. Since interference scales relative to the density of HetNet deployments, the need for improved interference management becomes essential [22]. HetNet interference management centers around coordinating individual cells to minimize conflicting radio-frequency (RF) transmissions [22]. Another approach to minimizing HetNet interference is to adopt an orthogonal frequency-division multiplexing (OFDM) approach, in which each cell transmits on different sub-bands [23]. Our design aims to build upon these approaches by developing a general RAN management primitive through cross-layer RAN orchestration.

The flexibility of the BoTM RAN management primitive enables a variety of operational use cases. For example, in the context of carrier aggregation, the RAN topology may be instantiated in such a way that many UEs may be unable to utilize carrier aggregation. Carrier aggregation allows UEs to utilize RAN resource blocks from both the telephone carrier that the UE is subscribed to and the other adjacent telephone carriers e.g. an AT&T UE utilizes resources from AT&T and

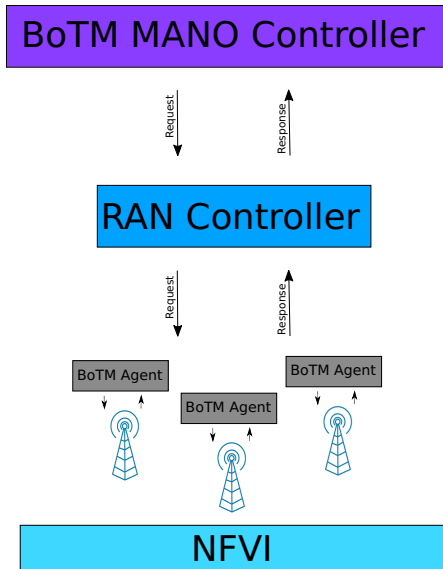


Fig. 2. BoTM Architecture

T-Mobile. Depending on the RAN topology, UEs may require more than one transceiver to utilize carrier aggregation [24], preventing access to nearby radio resources; however, by shifting the topology, spectrum resources can possibly be conjoined into a single, extended piece of spectrum, requiring only a single transceiver. Related work in this space utilized inter-frequency handover to move between telephone carriers within a simulated environment [25]. Our work emphasizes a realistic implementation, and our experimentation are conducted in a controlled RF environment.

III. BoTM SYSTEM DESIGN & IMPLEMENTATION

This section details the BoTM design and implementation.

A. BoTM Network Management Primitive

BoTM's migration primitive consists of four fundamental steps (Fig. 3). First, a migration event is triggered by either the RAN, the RAN controller, or the MANO controller to determine whether the base station should migrate to a new position and/or piece of spectrum. A migration event relays crucial information to the BoTM MANO controller, indicating diminishing QoS or QoE, sub-optimal use of resources, or potential attacks on the mobile network infrastructure. Second, upon detecting a migration event, the MANO controller triggers the migration procedure, instantiating a new base station. Depending on the type of event, the newly instantiated base station's parameters may be tailored accordingly. For example, in the case of diminishing QoS or QoE, the base station may be instantiated with more resource blocks, an increased transmit gain, and carrier aggregation enabled, while, in the case of inadequate utilization of spectrum or potential jamming attacks, a base station may be instantiated on a new transmit frequency. Base station instantiation may be automated through pre-configured configuration files or dynamically allocated based on real-time parameters relayed to the BoTM agent.

Third, after instantiating the new base station instance, a handover between the source base station and the new target base station is then initiated by the MANO controller and executed by the RAN controller. If a direct link exists between the source and target base station, then the handover may be executed over the X2 interface, with the S1 interface being utilized in the absence of a direct link, which, in turn, reduces locality constraints on the target base station. Furthermore, handovers may range from intra-frequency, inter-frequency, inter-RAT, or a combination of the latter. For example, intra-frequency handover may be utilized in a scenario where spectrum usage is constrained and load balancing is warranted, while inter-frequency handover may be employed in a situation where spectrum is being utilized inefficiently. Furthermore, in the case of diminishing 5G services, a combination of inter-frequency and inter-RAT handover may be utilized, e.g., the target base station may constitute an LTE eNodeB base station within the sub 6 GHz frequency range. Fourth, after the handover succeeds, the MANO controller then removes the source base station instance from the network, freeing the associated resources. BoTM effectively emulates the movement of a base station and its associated endpoints to a new location and/or piece of spectrum, in an effort, to expand the toolkit of existing RAN management systems.

B. MANO Controller

The Management and Orchestration (MANO) controller is primarily responsible for overseeing the life cycle of base station instances, managing the RAN topology, and carrying out non-real time control functions. We mainly utilize the MANO controller to signal the RAN to initiate BoTM's migration procedure and to instantiate or remove base stations. The MANO controller runs "above" the RAN and the RAN controller, providing direct access to a global view of the RAN. This allows the MANO controller to aggregate data to influence and enforce RAN policy decisions such as the frequencies being utilized, the number of base stations deployed, and the types of RATs present.

Policy decisions and life cycle updates are delivered directly to BoTM agents running "above" the individual base stations, signaling instantiation, deletion, or configuration. Base station configuration is crucial to enforcing compatibility and collaboration between neighboring base stations; however, configuration messages can be relayed to multiple base stations, in order, to shift and mold the RAN topology itself. For example, frequencies can be fine-tuned to minimize transmission interference between base stations, while a variety of RATs can be adopted to minimize network failure through RAT fallback mechanisms. In addition, based on an existing cell's occupancy, the given cell can potentially be "reshaped," divided, or expanded on-the-fly by instantiating or removing base station instances, allowing BoTM to support cell-level orchestration and load-balancing within the RAN. Specifically, BoTM's MANO controller enables us to leverage RAN orchestration as a primitive for flexible RAN management.

C. RAN Controller

The RAN controller is mainly responsible for carrying out real-time control functions, integrating network applications, and maintaining statistical data. Particularly, we leverage the RAN controller to dynamically trigger a network-initiated handover between a source base station and a target base station. We also utilize the RAN controller to enforce RAN policy updates such as altering the distribution of base station resources, coordinating base station transmission times, and maintaining minimum levels of network QoS and QoE, in an effort, to improve RAN resource efficiency and mitigate inter-base station interference. In conjunction with RAN orchestration, we utilize the RAN controller's near real-time protocol configuration and policy management as another building block for flexible RAN management.

D. Virtualized RAN

Base station instances are instantiated over a virtualized infrastructure (NFVI). The virtualized nature of the underlying infrastructure promotes RAN orchestration through the dynamic instantiation and removal of custom, lightweight base station instances. Due to the low cost of deploying and removing SD-RAN instances, base stations can be rapidly deployed or removed, which enables base station parameters, such as the number of resource blocks, the transmit frequency, the receive gain, and the transmit gain, to be fine-tuned based on the deployed environment and the time of day, fostering a more diverse set of RAN management policies, e.g., RAN policies can potentially be optimized and enforced based on specific times and days.

E. Implementation

BoTM was implemented and evaluated with OAI [10] and FlexRAN [8]. We implemented a basic BoTM MANO controller, which was deployed "above" the FlexRAN controller, allowing the BoTM MANO controller to access a global view of the RAN, in addition to FlexRAN's aggregated statistics. We introduced BoTM base station agents to run over the top of OAI's LTE eNodeB framework, enabling direct communication with the BoTM MANO controller. This communication link facilitated dynamic instantiation and deletion of base station instances, extending the OAI framework to support RAN orchestration. The BoTM MANO controller leveraged fluctuations in the Channel Quality Indicator (CQI) for the migration decision process, i.e., if the CQI dropped below a predefined threshold, the migration process was triggered.

Upon the migration trigger, the BoTM MANO controller issued a create command to the target BoTM agent, which then instantiated the target base station with a pre-configured OAI configuration file, ensuring compatibility with the source base station. Inter-base station communication was managed through OAI's X2 application protocol (X2AP) implementation, which enabled the newly instantiated target base station to peer directly with the source base station. After the target base station successfully peered with the source base station, the BoTM MANO controller notified the FlexRAN controller

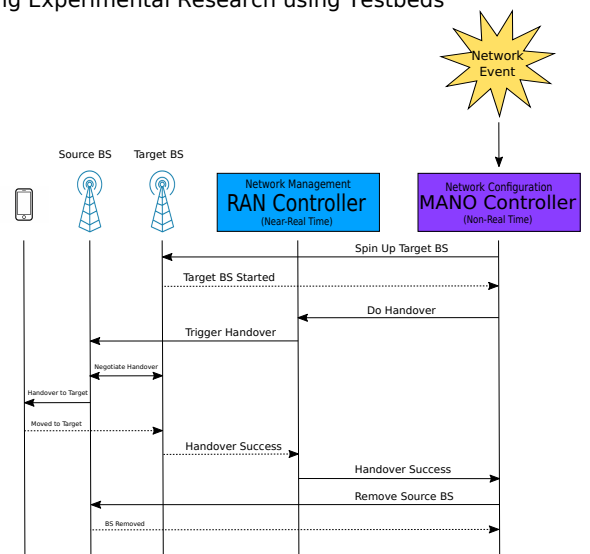


Fig. 3. BoTM Network Management Primitive

to initiate an intra-frequency X2 handover from the source base station to the target base station, which was negotiated over the X2 interface. Upon a successful handover, the BoTM MANO controller issued a removal command to the source BoTM agent, which then passed a kill command to the source base station, removing it from the network. Lastly, to eliminate potential functionality limitations that arise from the use of software-defined UEs, we chose to utilize a commercial off the shelf (COTS) UE.

IV. EVALUATION

In this section, we evaluate our prototype realization of the BoTM network management primitive.

A. Experimental Setup

All necessary resources were provided through the POWDER wireless testbed [26]. The experiment was conducted within POWDER's controlled-RF environment. For the experiment, the MANO controller and RAN controller [8] were deployed on a POWDER d430 compute node. The d430 was equipped with an 8-core E5-2630v3 processor and 64 gigabytes of RAM, running Ubuntu 18.04 Standard. Two agent-enabled eNodeBs were deployed on Intel-NUC5300s. The eNodeBs utilized a dual-core intel i5-5300U processor, 16 gigabytes of RAM, and a software-defined radio (Ettus B210 USRP), running the latest OAI release [10] on Ubuntu 16.04 low-latency. The eNodeBs were deployed with a total of 25 physical resource blocks (5 MHz) on band 4. An Evolved Packet Core (EPC), based on NextEPC [27], was deployed on a POWDER d430 compute node. The testbed utilized a Nexus 5 COTS UE running Android KitKat (4.4.4 standard). The starting network topology was constructed with the two agent-enabled eNodeBs connected to the MANO controller, RAN controller, and the EPC over a local area network, with the UE connected to a single eNodeB over an RF-link.

B. Interference Management

We explore the practicality of our prototype as a network management primitive by analyzing its performance within the context of a specific use case, namely, RF interference




Fig. 4. B₀TM Channel Quality Indicator Over Time. Interference was emulated in both the uplink and downlink directions at the 9 second mark. At the 10 second mark, upon detecting the drop in signal quality, the MANO controller initiated the migration process. The target base station was instantiated near the 18 second mark. After instantiating the new base station, the MANO controller then initiated a UE handover from the source base station to the new base station, which concluded at the 20 second mark. After the UE moved to the new base station, the MANO controller removed the source base station instance from the network.

management. We evaluated this B₀TM use case in terms of the observed CQI, uplink bitrate, and downlink bitrate. We began our experiment by generating simultaneous uplink and downlink UDP traffic, from the UE and EPC respectively, via the iPerf2 network tool. We then simulated interference within our controlled RF environment in both the uplink and downlink directions. Upon detecting a drop in signal quality, the MANO controller initiated the migration procedure. During the experiment, we sampled the observed CQI with OAI’s built in trace mechanism, while recording the observed bitrates produced by iPerf.

As shown in Figure 4, after the migration procedure, we can see an increase in the perceived CQI, with the uplink and downlink bitrates (Fig. 5) remaining unchanged. It is important to note that, during B₀TM’s migration procedure, the observed CQI remains low for a relatively long period of time, which can be attributed to the start up overhead required to spin up a brand new base station instance.

As for the uplink and downlink bitrates, the data indicates that B₀TM’s migration procedure introduces little to no impact. When the UE attaches to the base station, B₀TM’s migration procedure experiences a large spike in the perceived uplink bitrate. We attribute this large spike to a Packet Data Convergence Protocol (PDCP) buffer buildup of the uplink UE traffic, which, momentarily, emulates an increase in throughput, when the buffer state is processed at the target base station [28]. The data also indicates that B₀TM’s orchestrated migration procedure introduces a small, relatively similar, impact on the UE’s QoS and QoE as a standalone OAI X2 handover, with an average of 2.1% and 0.16% packet loss for the standalone X2 handover in the downlink and uplink traffic respectively, and 1.9% and 0.19% packet loss for B₀TM in the downlink and uplink directions (Table I). While B₀TM’s migration procedure and the standalone X2 handover produce similar impacts on the UE’s QoS and QoE, we emphasize the orchestration aspect




Fig. 5. B₀TM Downlink and Uplink Bitrate Over Time

of B₀TM, since an X2 handover cannot be executed without an existing, neighboring base station, while B₀TM’s migration procedure can. We attribute the small impact on the user QoS and QoE to the overhead required to execute a handover between the two base stations. Despite the small overhead, we believe the increased CQI and flexibility of B₀TM, within interference prone environments, to be significant, demonstrating the RAN management enhancements offered by B₀TM: a new RAN management primitive.

V. FUTURE WORK

Spectrum Migration. B₀TM currently leverages intra-frequency X2 handover; however, in the future, we plan to extend its functionality to utilize a form of inter-frequency handover to mimic spectrum migration, i.e., migrating a base station to a new piece of spectrum, which, ultimately, may lead to a more efficient use of spectrum.

Load Balancing. B₀TM’s current functionality is limited to migrating a single base station instance. We posit that B₀TM can be enhanced to not simply migrate a base station but to “split” a base station instance, in order, to decrease network load. For example, instead of instantiating a single target base station, multiple target base stations can be instantiated, allowing the UEs to be optimally distributed, which, in turn, may lower the average load per base station.

Intelligence. Our prototype’s migration decision process was constituted by a simple, preset CQI threshold. With the growing complexities of next generation networking, human engineered heuristics may not achieve optimal results within the growing network ecosystem. The shift to minimize human involvement by incorporating artificial intelligence has delivered promising results for the future of networking by dealing with the added complexities [29]. Furthermore, recent trends posit that artificial intelligence will become a key enabler of next generation networking [30]–[32]. Investigating the possibility of integrating intelligence into B₀TM to create a more dynamic and reactive management system is a topic of future interest.

VI. CONCLUSION

We have presented B₀TM, a novel RAN management primitive, and validated its implementation through a proof-of-

TABLE I
DOWNLINK AND UPLINK MEASUREMENT SUMMARY

	<i>Average Bitrate</i>	<i>Average Packet Loss</i>	<i>Total Dropped Packets</i>	<i>Total Packets Sent</i>
BoTM Downlink	1.03 Mbps	1.9%	206	10701
Standalone Downlink	1.03 Mbps	2.1%	223	10701
BoTM Uplink	1.03 Mbps	0.19%	20	10564
Standalone Uplink	1.03 Mbps	0.16%	17	10565

concept prototype. BoTM illustrates the flexibility of emerging SD-RAN systems to enable powerful RAN management primitives through cross-layer orchestration. We argue that primitives such as BoTM will be important to realize efficient RAN resource management, amidst the growing complexity and sophistication of next generation mobile and wireless networks. We have packaged our work into a POWDER profile that enables others to replicate our results and to serve as a starting point for related research efforts [33].

REFERENCES

- [1] "FLARE: Open Deeply Programmable Network Node Architecture," 2012. [Online]. Available: http://netseminar.stanford.edu/seminars/10_18_12.pdf
- [2] "O-RAN: Towards an Open and Smart RAN," <https://www.o-ran.org/resources>, 2018.
- [3] "xRAN Fronthaul Working Group White Paper v1.0," 2018.
- [4] "O-RAN: Use Cases and Deployment Scenarios," <https://www.o-ran.org/resources>, 2020.
- [5] I. Akyildiz, p. Wang, and S.-C. Lin, "Softair: A software defined networking architecture for 5G wireless systems," *Computer Networks*, vol. 85, 05 2015.
- [6] E. Coronado Calero, S. N. Khan, and R. Riggio, "5g-empower: A software-defined networking platform for 5G radio access networks," *IEEE Transactions on Network and Service Management*, vol. PP, pp. 1–1, 04 2019.
- [7] X. Foukas, M. Marina, and K. Kontovasilis, "Orion: Ran slicing for a flexible and cost-effective multi-service mobile network architecture," 10 2017, pp. 127–140.
- [8] X. Foukas, N. Nikaein, M. M. Kassem, M. K. Marina, and K. Kontovasilis, "Flexran: A flexible and programmable platform for software-defined radio access networks," in *Proceedings of the 12th International Conference on Emerging Networking EXperiments and Technologies*, ser. CoNEXT '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 427–441. [Online]. Available: <https://doi.org/10.1145/2999572.2999599>
- [9] K. Koutlia, R. Ferrus, E. Coronado Calero, R. Riggio, F. Palacio, A. Umberto, and J. Pérez-Romero, "Design and experimental validation of a software-defined radio access network testbed with slicing support," *Wireless Communications and Mobile Computing*, vol. 2019, pp. 1–17, 06 2019.
- [10] N. Nikaein, M. K. Marina, S. Manickam, A. Dawson, R. Knopp, and C. Bonnet, "Openairinterface: A flexible platform for 5g research," *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 5, p. 33–38, Oct. 2014. [Online]. Available: <https://doi.org/10.1145/2677046.2677053>
- [11] NEC Corporation, "Making 5G a Reality," https://www.nec.com/en/press/201802/global_20180222_03.html.
- [12] H. Zhao, R. Mayzus, S. Sun, M. Samimi, J. K. Schulz, Y. Azar, K. Wang, G. N. Wong, F. Gutierrez, and T. S. Rappaport, "28 ghz millimeter wave cellular communication measurements for reflection and penetration loss in and around buildings in new york city," in *2013 IEEE International Conference on Communications (ICC)*. IEEE, 2013, pp. 5163–5167.
- [13] A. Narayanan, E. Ramadan, J. Carpenter, Q. Liu, Y. Liu, F. Qian, and Z.-L. Zhang, "A first look at commercial 5g performance on smartphones," in *Proceedings of The Web Conference 2020*, ser. WWW '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 894–905. [Online]. Available: <https://doi.org/10.1145/3366423.3380169>
- [14] I. F. Akyildiz, W. Lee, M. C. Vuran, and S. Mohanty, "A survey on spectrum management in cognitive radio networks," *IEEE Communications Magazine*, vol. 46, no. 4, pp. 40–48, 2008.
- [15] T. O. Olwal, K. Djouani, and A. M. Kurien, "A survey of resource management toward 5g radio access networks," *IEEE Communications Surveys Tutorials*, vol. 18, no. 3, pp. 1656–1686, 2016.
- [16] A. Ghosh, "Spectrum usage analysis and prediction using lstm networks," M.S. thesis, Dept. School of Computing, Univ. of Utah, 2020.
- [17] I. Gomez-Miguel, A. Garcia-Saavedra, P. D. Sutton, P. Serrano, C. Cano, and D. J. Leith, "srslte: An open-source platform for LTE evolution and experimentation," *CoRR*, vol. abs/1602.04629, 2016. [Online]. Available: <http://arxiv.org/abs/1602.04629>
- [18] "5G PPP Phase 1 Security Landscape," <https://5g-ppp.eu/>, 2017.
- [19] R. v. Nee and R. Prasad, *OFDM for Wireless Multimedia Communications*, 1st ed. USA: Artech House, Inc., 2000.
- [20] B. Han, V. Gopalakrishnan, L. Ji, and S. Lee, "Network function virtualization: Challenges and opportunities for innovations," *IEEE Communications Magazine*, vol. 53, no. 2, pp. 90–97, 2015.
- [21] D. Kreutz, F. M. V. Ramos, P. E. Verissimo, C. E. Rothenberg, S. Azodolmolky, and S. Uhlig, "Software-defined networking: A comprehensive survey," *Proceedings of the IEEE*, vol. 103, no. 1, pp. 14–76, 2015.
- [22] D. Lopez-Perez, I. Guvenc, G. de la Roche, M. Kountouris, T. Q. S. Quek, and J. Zhang, "Enhanced intercell interference coordination challenges in heterogeneous networks," *IEEE Wireless Communications*, vol. 18, no. 3, pp. 22–30, 2011.
- [23] J. G. Andrews, "Seven ways that hetnets are a cellular paradigm shift," *IEEE Communications Magazine*, vol. 51, no. 3, pp. 136–144, 2013.
- [24] R. Ratasuk, D. Tolli, and A. Ghosh, "Carrier aggregation in lte-advanced," in *2010 IEEE 71st Vehicular Technology Conference*, 2010, pp. 1–5.
- [25] A. Krotov and B. Bojovic, "Enabling LTE CA handover to secondary cell," <https://www.nsnam.org/wiki/GSOC2017Lte>, 2017.
- [26] J. Breen, A. Buffmire, J. Duerig, K. Dutt, E. Eide, M. Hibler, D. Johnson, S. K. Kasera, E. Lewis, D. Maas, A. Orange, N. Patwari, D. Reading, R. Ricci, D. Schurig, L. B. Stoller, J. Van der Merwe, K. Webb, and G. Wong, "POWDER: Platform for open wireless data-driven experimental research," in *Proceedings of the 14th International Workshop on Wireless Network Testbeds, Experimental Evaluation and Characterization (WiNTECH)*, Sep. 2020.
- [27] N. Inc., "Open Source Implementation of LTE EPC," <https://nextepc.org/>, 2019.
- [28] B. Nguyen, A. Banerjee, V. Gopalakrishnan, S. Kasera, S. Lee, A. Shaikh, and J. Van der Merwe, "Towards understanding tcp performance on lte/epc mobile networks," in *Proceedings of the 4th Workshop on All Things Cellular: Operations, Applications, & Challenges*, ser. AllThingsCellular '14. New York, NY, USA: Association for Computing Machinery, 2014, p. 41–46. [Online]. Available: <https://doi.org/10.1145/2627585.2627594>
- [29] S. Niknam, A. Roy, H. S. Dhillon, S. Singh, R. Banerji, J. H. Reed, N. Saxena, and S. Yoon, "Intelligent o-ran for beyond 5g and 6g wireless networks," 2020.
- [30] K. B. Letaief, W. Chen, Y. Shi, J. Zhang, and Y. A. Zhang, "The roadmap to 6g: Ai empowered wireless networks," *IEEE Communications Magazine*, vol. 57, no. 8, pp. 84–90, 2019.
- [31] W. Saad, M. Bennis, and M. Chen, "A vision of 6g wireless systems: Applications, trends, technologies, and open research problems," *IEEE Network*, vol. 34, no. 3, pp. 134–142, 2020.
- [32] C. Wang, M. D. Renzo, S. Stanczak, S. Wang, and E. G. Larsson, "Artificial intelligence enabled wireless networking for 5g and beyond: Recent advances and future challenges," *IEEE Wireless Communications*, vol. 27, no. 1, pp. 16–23, 2020.
- [33] "POWDER BoTM profile," <https://www.powderwireless.net/PowderProfiles/BoTM>.