Resiliency in Open-Source Solutions for Disaggregated 5G Cloud Radio Access and Transport Networks

Shunmugapriya Ramanathan[#], Koteswararao Kondepu (Supervisor)[•], Andrea Fumagalli (Supervisor)[#]
[#]Open Network Advanced Research lab, The University of Texas at Dallas
[•]Indian Institute of Technology Dharwad, Dharwad, India
e-mail:sxr173131@utdallas.edu

***Abstract*—Cloud Radio Access Network (C-RAN) has emerged as an advantageous architecture to address the needs of future wireless networks. Leveraging both Network Functions Virtualization and Software-Defined Networking (NFV-SDN) technologies, the C-RAN components can be disaggregated and potentially run in a virtualized environment where either virtual machines (VMs) or containers are efficiently applied. C-RAN solutions require highly reliable fronthaul, midhaul, and backhaul network designs. This work first describes a programmable optical software-defined network testbed that offers fronthaul, midhaul, and backhaul transport capabilities supporting C-RAN functionalities with increased reliability.**

Optical transport programmable (SDN) capabilities also facilitate live-migration of C-RAN functions between sites in order to achieve improved compute load-balancing and fault-tolerance of the overall C-RAN system. Container-based frameworks are increasing gaining momentum due to their smaller footprint size compared to VM's. However, the extant container migration software frameworks do not support mobile network protocol stacks. The second contribution of our work focuses on the live-migration of both containerized core network and RAN central unit virtual functions, which is achieved through a proof-of-concept implementation based on open-source software.

As briefly outlined in this paper our studies demonstrate the achievable high-reliability in the fronthual network along with a guaranteed end-user service continuity without permanent interruption during live-migration of specific mobile network virtual functions. In addition, the temporary disruption time experienced by the end-user service during the function live-migration is reduced by more than 50% when using our designed container platform compared to conventional VM solutions.

***Index Terms*—5G resiliency, virtual RAN, docker container, container migration, CRIU.**

I. Introduction

Fifth-generation mobile networks are expected to experience significant growth in wireless data traffic and services in the upcoming years. According to the Cisco Internet Report for 2018-2023 white paper [1], there will be 5.3 billion total Internet users by 2023, up from 3.9 billion in 2018. To contain both capital (CAPEX) and operating expenses (OPEX), operators have increasingly relied on open standard solutions that leverage Software-Defined Networking (SDN), Network Function Virtualization (NFV), and Cloud Radio Access Network (C-RAN). For example, in the C-RAN architecture, baseband units (BBU) are relocated away from individual base stations to reside in centralized control and processing stations. The radio network elements are implemented as Virtual Network Functions (VNFs) to simplify network programmability, service flexibility, and reconfiguration.

Several functional split options are available to design C-RAN architectures by separating and hosting the next-generation Node B (gNB) functionalities across two distinct elements, i.e., the Central Unit (CU) and the Distributed Unit (DU). The resulting geographically distributed C-RAN architecture requires high-speed and low-latency transport networks. The high-capacity and low-latency characteristics of optical fiber communication, along with the already deployed fiber cables in the field, make Dense Wavelength Division Multiplexing (DWDM) a most desirable solutions for the C-RAN transport network. Several proprietary and open optical network solutions are beginning to offer *physical layer (SDN) programmability*, which can be readily leveraged to achieve highly reliable transport network connectivity. Based on the applied functional split option, the fronthaul bandwidth and latency requirements may vary significantly and are generally more demanding than the backhaul requirements in the same C-RAN solution. Therefore, in terms of resiliency, network link fault tolerance mechanisms must consider both the transport network's latency and restoration time constraints as they are dictated by the supported mobile applications.

Unlike traditional LTE networks that rely on proprietary turnkey solutions, the C-RAN architecture uses virtualized mobile network modules and runs as software on inexpensive commodity hardware. Virtualization platforms enable efficient resource utilization by providing resource isolation, multi-tenancy support, and flexible resource allocation. However, carrier-grade Quality of Service (QoS) is not naturally guaranteed by the deployed virtualization platform. For example, dedicated and proprietary hardware reliability is inherently superior to commodity Cloud servers, and this disadvantage must be compensated by an intelligent design and implementation of the virtualized software modules. In simple words, guaranteeing the high availability of mobile VNFs is essential to achieve carrier-grade reliable mobile network services. What's more, Container Network Function (CNF) continues to gain momentum at this point as it is believed to reduce management overhead and be easily managed through open-source Kubernetes or Docker [2].

Besides automatic fault handling, live-migration of both

VNFs and CNFs is instrumental in achieving efficient resource utilization in mobile network communication [3] and many other applications. A typical mobile network is expected to connect multiple gNB-vDUs to a single gNB-vCU, and multiple gNB-vCUs to a single core network. These modules communicate with one another using connection-oriented protocols which must therefore rely on efficient fault-tolerant mechanisms.

To address the much-needed resiliency improvement of the disaggregated mobile communication networks — inclusive of the optical transport network — this paper is organized as follows:

- Definition of the problem statements
- A survey of network and device resiliency of the mobile network
- A developed and validated Proof-of-Concept (POC) implementation of 1+1 network link restoration on the open-source RAN software (OpenAirInterface (OAI) repository) [4]
- A developed and validated implementation of the live migration of the Core Network (CN) and RAN central unit components [5], [6]
- Experimental evaluation of designed strategies using VM and *stateful* container migration techniques in the federated testbed [7]

II. Problem Statements

In the 5G New Radio (NR) architectural model, 3GPP defines a new, flexible architecture for the 5G RAN in Release 15, where the gNB is split into three logical nodes: the Central Unit (CU), the Distributed Unit (DU) and the Radio Unit (RU), each capable of hosting distinct functions of the 5G NR stack. In that, 3GPP specifies eight functional split options for distributing the radio functionality across the fronthaul, midhaul, and backhaul network. These options offer operators various trade-offs between latency requirements and bandwidth availability as detailed in [8]. Of these eight splits, the industry has widely adopted two split options.

- Option 2 – a high-level CU and DU split, which is primarily the separation of the control from the user plane. Commercial Off-The-Shelf (COTS) hardware combined with NFV is highly preferred for CU and DU in the cloud.
- Option 7 (or 7.2) - a low-level split for ultra-reliable low-latency communication (URLLC) and near-edge deployment. By enabling centralized scheduling and low-cost remote radio unit (RRU) designs for wide adoption, Option 7 split is an ideal solution for the distributed RAN deployment.

Problem Statement #1: Resiliency is essential in open transport optical networks. How to maintain a resilient fronthaul network with stringent latency requirements?

The transport network is expected to carry a converged mix of fronthaul, midhaul, backhaul, and non-5G traffic, thus requiring large capacity. In addition, guaranteeing reliability is of paramount importance. Disruptions or even the failure of a single component in the transport network might lead to cascading failures in other components. Many open and proprietary approaches are available to handle fault handling in the backhaul network. The widely used 1:1 fast failover table protection mechanism experiences packet loss during the protection procedure, as the switch takes time to detect the link outage and forward the flow over the secondary path. The radio interface needs to maintain a higher quality of service to ensure priority for time-critical data.

Problem Statement #2: COTS hardware gives better CAPEX, OPEX, and flexibility options in the data centers. How to ensure carrier-grade reliability while using COTS hardware in 5G applications?

While numerous methods have been discussed in the literature to handle element failures when 5G functions are virtualized, there is no detailed analysis of container-based live migration for the core network and radio access network components. For example, in [9], a two-step resiliency scheme is described that overcomes soft failures of optical circuits (lightpaths) carrying the fronthaul traffic. Containers are emerging as a lightweight alternative to virtual machines and are becoming part of the service lifecycle [10]. In [11], the authors evaluated the container-based live-migration of the blank and video applications to meet the multi-access edge computing requirements during radio handover.

III. Background and Related Work

A. Fronthaul Interface Resiliency Challenges

C-RAN initially proposes the centralization of radio intelligence to enable a model resembling the data centers, which has shown tremendous success. Then the virtualization of these centralized radio functions allows running them in commercial off-the-shelf hardware, removing the proprietary hardware dependency and enabling ease of management. All base station functionalities are split into central, distributed, and radio units. The term "fronthaul" refers to the network segment that connects the distributed unit with the radio unit. The split can occur at different radio protocol layers, thus resulting in different bandwidth and delay requirements of the mobile fronthaul. The split at the physical layer — called the low-level split — enables effective virtualization of radio functionalities with simple, low-cost radio unit design for wide adoption. A popular interface carried on these fronthaul links is Ethernet over Common Public Radio Interface (eCPRI), which is used to transmit high-resolution digitization of the radio signals. As reported in [12], when eCPRI is utilized to transport the option 7 split, the network delay must be within 100 s because of the time-sensitive I/Q radio signal data being transmitted. The key constraint in this approach is the latency-stringent requirement because of the subframe-level timing interactions between the part of the PHY layer in DU and part of the PHY layer in RUs.

Reliability is particularly crucial in this latency-stringent fronthaul network. Several methods have been proposed to recover from hardware and software failures. In the technical specification (TS) 23.007, 3GPP specifies a number of failure detection and recovery mechanisms for the CN [13], handling failure with the echo request/response messages. Several resilient schemes for recovering C-RAN failure are proposed in [14]. ETSI NFV defines the scalable architecture for reliability management [15], and part of it is implemented in open-source orchestration frameworks such as OpenStack [16]. Furthermore, to satisfy the fronthaul requirements, lightpath transmission adaptation is proposed in [17] to recover the radio functionality.

B. Container Migration Challenges

Recently, the container-related open-source community has been expanding containerization technology with the increasing need for agile architecture and innovation in digital transformation. A container live-migration of a generic application is relatively straight forward compared to doing the same for mobile network services. In fact, the generic application does not have strict timing and synchronization requirements. Conversely, the gateway of a mobile network must maintain the connection state of every end-user (UE). Therefore, it is essential to hold the connectivity when migrating one of the radio network virtual functions, which requires *stateful container migration*. Maintaining UE connectivity requires the migrated container application to retain its configuration, connection state, and IP address. In addition, the brief unavailability of service during the live migration creates a temporary call drop — referred to as downtime — in the network. If a UE experiences a connection problem due to migration, it typically waits for a specified time before it releases logical connection with the CN. Therefore, a critical study on downtime improvements is needed for the mobile network. Furthermore, the communication protocol in the 5G — Stream Control Transmission Protocol (SCTP) — is not yet supported by the most widely used open-source container migration software, i.e., CRIU [18]. For this reason perhaps studies on stateful migration approaches for the containerized 5G mobile network are limited at this time. Recently, several research efforts have started to analyze the radio component container migration [19]. The authors have examined the LTE baseband container live migration in [20], but because of the unavailability of SCTP support, they limited their analyses to the fronthaul only without accounting for the CN connectivity. The virtualized CU (vCU)/virtualized DU (vDU) VNF migration over DWDM network using CRIU is briefly discussed in [21].

IV. Our Contributions

A. Fronthaul Network Resiliency

We configured an option 7 split-based C-RAN (from OAI) testbed with a programmable DWDM fronthual and backhaul implemented using Cisco NCS 2000 to study network resiliency. Our testbed provides two built-in fault tolerant mechanisms, i.e., 1:1 + R and 1 + 1 + R, where both the 1:1 and 1 + 1 protection mechanisms are implemented at the Ethernet layer and the R restoration mechanism (optical circuit restoration) is implemented at the DWDM layer. The 1:1 protection data plane is implemented using the OpenFlow fast failover table function. We developed the 1 + 1 protection data plane (which is not readily available in OpenFlow/OVS) in kernel software to handle the tight latency requirements of the fronthaul network.

Shown in Fig. 1, the 1+1 protection mechanism (ingress switch OVS A) sends two copies of each data packet over two disjoint paths: one over the primary path and one over the secondary path [4]. At the other end of the two paths, the egress switch (OVS E) receives the duplicated flows removing the second copy at the kernel level to ensure that only one copy of each radio packet is delivered to the host. In option 7 split, the radio packets are sent in frame format. Each radio frame consists of 10 sub-frames, each of which contains 1 ms of signal samples. Each sub-frame has 2 slots with 14 OFDM symbols in it. The purpose of our software module is to filter the duplicate at the slot level of the radio subframe (0.5 ms). Each data packet is uniquely identified by combining $frame\_number$, sub_frame_number, and $symbol\_number$. The algorithm for filtering the duplicate radio packets is summirized in the next table.

Algorithm 1: Discard duplicate radio packets in the kernel

Input: sk **Output:** route_status
Function Name: pre_route_hook_function

if $sk->ip\_header->protocol$ is UDP **then**
    Read the $dest$ IP and port details
    if $dest\_IP$ is RAN_IP AND $dest\_port$ is RAN_port **then**
        Skip the udp and ip headers
        Read 32_bit fr_num
        Read 8_bit sub_fr_num
        Read 8_bit sym_num
        $unique\_pattern$ = {fr_num,sub_fr_num,sym_num}
        if $unique\_pattern$ in the list **then**
            Duplicate found
            Discard the packet from the list
            $route\_status$ is NF_DROP
            return
        else
            First copy of the packet
            Store $unique\_pattern$ in the list
            Forward the packet
            $route\_status$ is NF_ACCEPT
            return
        end if
    else
        $route\_status$ is NF_ACCEPT
        return
    end if
else
    $route\_status$ is NF_ACCEPT
    return
end if

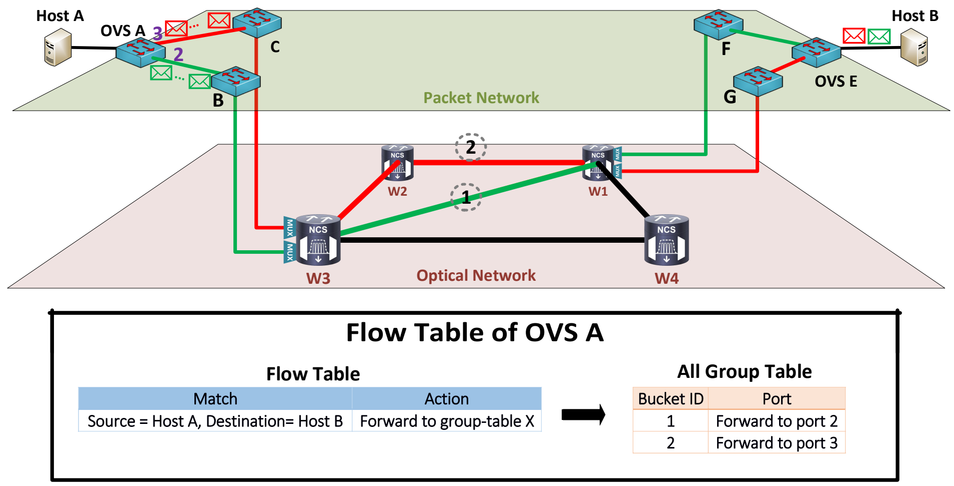

Fig. 1: 1+1 protection and restoration mechanism design [4]

The software removes the oldest received packet from the list when it exceeds a user-defined number of entries, N. This additional feature is required to prevent the list from growing unbounded in size when the second packet copy is not received, e.g., when one of the two paths is down.

This 1+1 protection mechanism is designed to overcome a single fiber failure in the network without causing any packet loss and therefore matching the stringent transmission requirements of the fronthaul network in the option 7 split. The companion R mechanism is implemented at the DWDM layer to automatically restore the lost optical circuit — for example the green path in Fig. 1 is restored using path W3-W4-W1 — thus promptly recreating the double flow of packets at the Ethernet layer to prepare the network to overcome a follow up outage. More details about this fronthaul and backhaul reliability study are available in [4]. The hook function software module is in BitBucket [1]. In addition, both network protection and restoration mechanisms have been demonstrated to overcome multiple fiber link outages [22].

B. Core Network Migration

For fault tolerance at the VNF level, live-migration of the virtual network components from one node to another achieves service connectivity with minimal interruption. Many studies have focused on the live-migration of generic applications and services in simulated or emulated network environments. However, experimental demonstrators of containerized mobile network components' live-migration have been limited due to the demanding protocol stack and timing requirements described in III-B. In our work, we have analyzed both the VM and container-based live-migration of the CN and the RAN CU in the soft real-time mobile network environment.

In the Non-Standalone (NSA) architecture, the 5G NR solution uses the existing LTE RAN and CN (EPC) as an anchor for mobility management and coverage to provide the 5G carrier. For the *CN migration study*, we started our investigation on the C-RAN mobile user service downtime in the presence of a failing mobility management entity (MME) when two alternative recovery mechanisms are applied: (i) *Proactive approach*: once the failure is predicted, live-migration of the virtual machine running the MME is performed from the affected data center to a backup data center, and (ii) *Reactive approach*: when a sudden unexpected failure occurs, the MME VM is replaced by another VM instance from a pool of geo-diverse MMEs, which is a functionality available in S1-flex with redundancy support [23].

We then explored the container-based CN live-migration [24], which presents a number of challenges in order to support the CN protocol stack using the open-source container migration software, CRIU. In [7], we proposed a framework that enables the migration of containers while running virtualized CN components. In addition, we did a comprehensive study on different failure scenarios in the virtualization environment, compared the migration and downtime between VM and container-based CN components in a soft real-time experimental setup, and also analyzed the system factors that influence the end-user service recovery time. Fig. 2 compares the migration time of the evolved packet core components — Home Subscriber Server (HSS), Mobility Management Entity (MME), and Serving and Packet Gateway (SPGW) — for three types of live-migration techniques: VM pre-copy, container stop-and-copy, and container pre-copy [7]

C. Radio Access Network Central Unit Migration

Live migration of virtualized Central Unit (vCU) using either VM or Container technologies is an additional critical functionality in the mobile network. With the 5G Standalone Architecture (SA) software modules available from the OAI, we considered the 5G SA option-2 split between the gNB-vDU and gNB-vCU, which in turn is connected to the containerized 5G CN components. In this configuration, multiple gNB-

[1] https://bitbucket.org/PriyaRamanathan/1-1-protection/src/master/

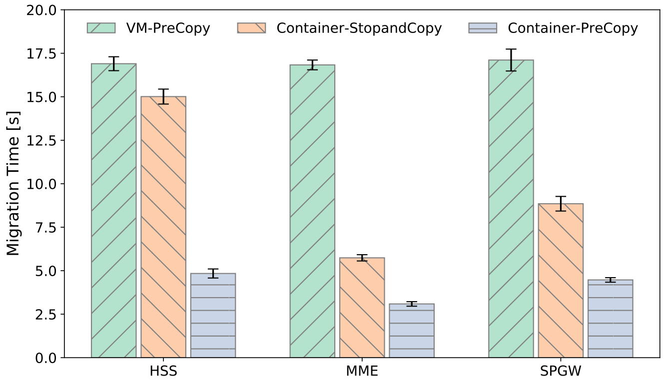

Fig. 2: Migration time comparison of core network components as a function of virtualization and migration type [7]

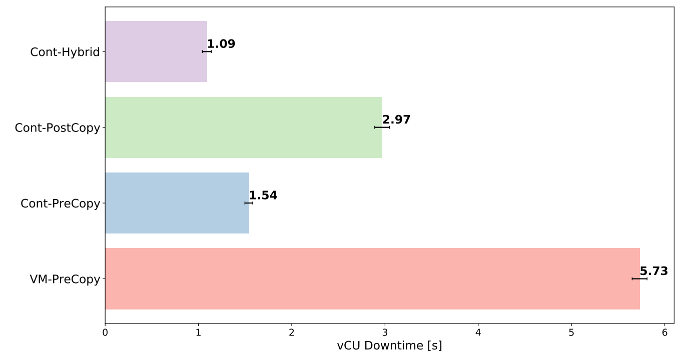

Fig. 3: Downtime comparison of gNB-vCU live-migration as a function of virtualization and migration type

vDUs are connected to a single gNB-vCU. Therefore a single CU outage would have a significant disruptive consequence on many connected UE devices. The challenge of this configuration is to handle the gNB-vCU live-migration without connectivity loss for the UE devices. Compared to VNF, CNF is more suitable to support the 5G NR access technology in the Next Generation Radio Access Network (NG-RAN) due to its shorter time to deploy and faster upgrade cycles [25].

In the 5G control plane stack, the gNB-vCU uses SCTP to exchange N2 control messages with the Access and Mobility Management Function (AMF) in the 5G CN. Similarly, the SCTP transport layer protocol is also adopted to exchange F1-C control messages with the gNB-DU. By so heavily relaying on SCTP (which is not currently supported by open-source migration software like CRIU), the containerized gNB-vCU live-migration cannot be easily achieved. We implemented a custom support for SCTP as a proof-of-concept in the container migration software and demonstrated cold copy-based gNB-vCU container migration without losing the end-user connectivity in a recent study [5]. The code base for the SCTP support is in BitBucket [2]. In addition, we explored different container migration approaches – pre-copy, post-copy, and hybrid – running under the runc hood. Some preliminary results of in-house lab testing for the gNB-vCU live-migration are shown in Fig. 3, which reports the UE service downtime caused by the gNB-vCU live-migration using a number of migration approaches. The Container hybrid-based migration yields the shortest downtime, approximately reducing it to 20% of that needed by the VM pre-copy approach. Both the migration time and UE service downtime can be reduced by fine tuning a number of system parameters [7].

V. Summary and Future Work

The growing user demand for real-time experiences has led operators to look for new ways of mobile service delivery at reduced CAPEX and OPEX. C-RAN and the complementing fields of NFV and SDN offer a bright future to communication systems. The availability, convergence, and economics of disaggregated fronthaul, midhaul and backhaul systems are also important factors to consider. Fast failovers are more suitable for the backhaul network where there are no strict real-time constraints. However, to meet the needs of the stringent 5G fronthaul system requirements, we have proposed and evaluated a mixed protection and restoration mechanism for achieving high service reliability. In addition to fault tolerance, flexible CU/DU placement may reduce overall power consumption and further increase reliability and availability of 5G mobile services. Demonstration of the first successful live-migration of containerized central unit and core network components with CRIU was achieved through custom software implementation. Overall, this research area is still in its infant stages, with our future work as detailed next.

A. Resilient Open Transport Network

The restoration mechanism reported in Section IV-A requires the use of a proprietary GMPLS implementation to control the DWDM transport network. However, latest efforts from the industry are increasingly pursuing open and all-photonic transport networks, e.g., IOWN [26]. Open standards and related open-source implementation efforts must be leveraged to implement restoration and protection mechanisms that offer reliability levels that are similar, if not superior, to the existing proprietary solutions. We propose to investigate the feasibility and scalability of protection and restoration mechanisms in open-source platforms like Transport-PCE [27], specifically designed to support multi-vendor OpenROADM compliant DWDM flex-grid transport equipment.

B. Reliability of Containerized Decoupled gNB-vCU

With the option 2 split, the gNB-vCU is decoupled into gNB-vCU Control Plane (CP) and gNB-vCU User Plane (UP). The gNB-vCU-UP holds mobile data traffic, and, as the number of connected UE devices increases, more data packets flow through the user plane. Implementing gNB-vCU-UP using data plane development kit (DPDK) [28], [29] is expected to improve the rate at which data packets are transmitted. In addition, the DPDK implementation in user space offers the unique advantage of making the solution kernel-independent, thus broadening the options available when choosing the target

[2]https://bitbucket.org/PriyaRamanathan/sctp_criu/

host for a given gNB-vCU-UP that must be migrated to in order to recover from an impending system failure or traffic congestion. Noting that open container migration software packages do not support DPDK, we propose to investigate the design and development of new software modules that will overcome this drawback.

C. Orchestration of Transport and Access Resources for Improved Reliability

As open solutions become available to automatically reprogram optical transport networks on the one hand, and improved container-based gNB-vCU-UP implementations through DPDK on the other, network designers must begin to look into automated solutions that can jointly leverage these two technologies. We propose to investigate the design and implementation of orchestrated (automated) procedures to further increase the reliability of the service offered to the UE by concurrently triggering protection and restoration mechanism in the transport network together with the instantiation of gNB-vCU-UP container migration. While intuition suggests an advantage when coordinately applying these two techniques — each designed to increase the reliability of the UE service — identifying a pragmatic approach and quantifying the achievable gains remain open challenges.

ACKNOWLEDGMENT

This work is supported in part by NSF grant CNS-1956357 and also received funding from DST SERB Startup Research Grant (SRG-2021-001522), and the SGNF project ("Reliability Evaluation of Virtualised 5G").

REFERENCES

[1] Cisco, "Cisco Internet Report White Paper," 2020. [Online]. Available: https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html

[2] C. Novaes, C. Nahum, I. Trindade, D. Cederholm, G. Patra, and A. Klautau, "Virtualized C-RAN Orchestration with Docker, Kubernetes and OpenAirInterface," 2020. [Online]. Available: https://arxiv.org/abs/2001.08992

[3] Y. Xiao, J. Zhang, and Y. Ji, "Energy efficient Placement of Baseband Functions and Mobile Edge Computing in 5G Networks," in *2018 Asia Communications and Photonics Conference (ACP)*, 2018, pp. 1–3.

[4] S. Ramanathan, M. Tacca, M. Razo, B. Mirkhanzadeh, K. Kondepu, F. Giannone, L. Valcarenghi, and A. Fumagalli, "A programmable optical network testbed in support of C-RAN: a reliability study," *Photonic Network Communications*, vol. 37, pp. 311–321, 2018.

[5] S. Ramanathan, A. Bhattacharyya, K. Kondepu, M. Razo, M. Tacca, L. Valcarenghi, and A. Fumagalli, "Demonstration of Containerized Central Unit Live Migration in 5G Radio Access Network," in *2022 IEEE 8th International Conference on Network Softwarization (NetSoft)*, 2022, pp. 225–227.

[6] S. Ramanathan, K. Kondepu, T. Zhang, B. Mirkhanzadeh, M. Razo, M. Tacca, L. Valcarenghi, and A. Fumagalli, "Orchestrating Virtualized Core Network Migration in OpenROADM SDN-Enabled Network," in *2021 International Conference on Optical Network Design and Modeling (ONDM)*, 2021, pp. 1–6.

[7] S. Ramanathan, K. Kondepu, M. Razo, M. Tacca, L. Valcarenghi, and A. Fumagalli, "Live Migration of Virtual Machine and Container Based Mobile Core Network Components: A Comprehensive Study," *IEEE Access*, vol. 9, pp. 105 082–105 100, 2021.

[8] Huber Shuner, "Cube-Optics-Functional-Split-Rev03." [Online]. Available: https://www.hubersuhner.com/en/documents-repository/technologies/pdf/fiber-optics-documents/5g-fundamentals-functional-split-overview

[9] K. Kondepu, A. Sgambelluri, N. Sambo, F. Giannone, P. Castoldi, and L. Valcarenghi, "Orchestrating Lightpath Recovery and Flexible Functional Split to Preserve Virtualized RAN Connectivity," *J. Opt. Commun. Netw.*, vol. 10, no. 11, pp. 843–851, Nov. 2018.

[10] Cisco, "Cloud-Native Network Functions." [Online]. Available: https://www.cisco.com/c/en/us/solutions/service-provider/industry/cable/cloud-native-network-functions.html#~introduction

[11] R. A. Addad, D. L. C. Dutra, M. Bagaa, T. Taleb, and H. Flinck, "Fast Service Migration in 5G Trends and Scenarios," *IEEE Network*, vol. 34, no. 2, pp. 92–98, 2020.

[12] D. Chitimalla, K. Kondepu, L. Valcarenghi, M. Tornatore, and B. Mukherjee, "5G fronthaul-latency and jitter studies of CPRI over ethernet," *Journal of Optical Communications and Networking*, vol. 9, no. 2, pp. 172–182, 2017.

[13] 3GPP TS 23.007, "Technical Specification Group Core Network and Terminals; Restoration procedures," Release 16, Mar. 2020.

[14] C. Colman-Meixner, G. B. Figueiredo, M. Fiorani, M. Tornatore, and B. Mukherjee, "Resilient cloud network mapping with virtualized BBU placement for cloud-RAN," in *2016 IEEE International Conference on Advanced Networks and Telecommunications Systems (ANTS)*, 2016, pp. 1–3.

[15] "Network Functions Virtualisation (nfv); Reliability; Report on Scalable Architectures for Reliability Management Group Specification."

[16] F. F. Moghaddam, A. Gherbi, and Y. Lemieux, "Self-Healing Redundancy for OpenStack Applications through Fault-Tolerant Multi-Agent Task Scheduling," in *2016 IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*, 2016, pp. 572–577.

[17] K. Kondepu, N. Sambo, F. Giannone, P. Castoldi, and L. Valcarenghi, "Orchestrating Lightpath Adaptation and Flexible Functional Split to Recover Virtualized RAN Connectivity," in *2018 Optical Fiber Communications Conference and Exposition (OFC)*, 2018, pp. 1–3.

[18] CRIU team, "CRIU What cannot be checkpointed," 2019. [Online]. Available: https://criu.org/What_cannot_be_checkpointed#Devices

[19] Jesutofunmi, 2019. [Online]. Available: http://rvs.unibe.ch/research/pub_files/Tofunmi_Thesis_.pdf

[20] E. Schiller, J. Ajayi, S. Weber, T. Braun, and B. Stiller, "Toward a Live BBU Container Migration in Wireless Networks," *IEEE Open Journal of the Communications Society*, vol. 3, pp. 301–321, 2022.

[21] J. Feng, J. Zhang, Y. Xiao, and Y. Ji, "Demonstration of Containerized vDU/vCU Migration in WDM Metro Optical Networks," in *2020 Optical Fiber Communications Conference and Exhibition (OFC)*, 2020, pp. 1–3.

[22] K. Kondepu, S. Ramanathan, M. Tacca, M. Razo, B. Mirkhanzadeh, F. Giannone, L. Valcarenghi, and A. Fumagalli, "Experimental Demonstration of a Packet-based Protection for Seamlessly Recovering from a Multi-layer Metro Network Fronthaul Failure," in *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 2019, pp. 906–908.

[23] S. Ramanathan, K. Kondepu, B. Mirkhanzadeh, M. Razo, M. Tacca, L. Valcarenghi, and A. Fumagalli, "Performance Evaluation of Two Service Recovery Strategies in Cloud-Native Radio Access Networks," in *2019 21st International Conference on Transparent Optical Networks (ICTON)*, 2019, pp. 1–5.

[24] S. Ramanathan, K. Kondepu, M. Tacca, L. Valcarenghi, M. Razo, and A. Fumagalli, "Container Migration of Core Network Component in Cloud-Native Radio Access Network," in *2020 22nd International Conference on Transparent Optical Networks (ICTON)*, 2020, pp. 1–5.

[25] 3rd Generation Partnership Project, "Study on new radio access technology; radio access architecture and interfaces," *Technical Specification Group Radio Access Network*, vol. 2.0.0, 2017.

[26] IOWN, "Open All-Photonic Network Functional Architecture." [Online]. Available: https://iowngf.org/wp-content/uploads/formidable/21/IOWN-GF-RD-Open-APN-Functional-Architecture-1.0-1.pdf

[27] TransportPCE, "UserGuide." [Online]. Available: https://docs.opendaylight.org/projects/transportpce/en/latest/user-guide.html

[28] Intel, "Data Plane Development Kit." [Online]. Available: https://www.intel.com/content/www/us/en/communications/data-plane-development-kit.html

[29] L. Lai, G. Ara, T. Cucinotta, K. Kondepu, and L. Valcarenghi, "Ultra-low latency nfv services using dpdk," in *2021 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN)*, 2021, pp. 8–14.