

Nuberu: Reliable RAN Virtualization in Shared Platforms

GARCIA-SAAVEDRA Andres, COSTA PEREZ Xavier

Abstract

RAN virtualization will become a key technology for next-generation mobile networks. However, due to the computing fluctuations inherent to wireless dynamics and resource contention in shared computing infrastructure, the price to migrate from dedicated to shared platforms may be too high. We present Nuberu, a novel pipeline architecture for 4G/5G DUs specifically engineered for shared platforms. Nuberu has one objective to attain reliability: to guarantee a minimum set of signals that preserve synchronization between the DU and its users during computing capacity shortages and, provided this, maximize network throughput. To this end, we use techniques such as tight deadline control, jitter-absorbing buffers, predictive HARQ, and congestion control.

Keywords



O-RAN, vRAN, DU, Shared Platforms, computational schedule

1. Introduction

The virtualization of radio access networks (RANs), based hitherto on hardwired ASICs, will become the spearhead of next-generation mobile systems beyond 5G. Initiatives such as the carrier-led O-RAN ALLIANCE have spurred the market and the research community to find novel solutions that import the flexibility and cost-efficiency of network function virtualization (NFV) into the very far edge of mobile networks.

Fig. 1 shows the architecture of a vRAN, with base stations (BSs) split into a central unit (CU), which hosts the highest layers of the stack; a distributed unit (DU), which hosts the physical layer (PHY); and a radio unit

(RU), which hosts basic radio functions such as amplification or sampling. To reduce costs, vRANs may rely on cloud platforms comprised of pools of shared computing resources (mostly CPUs, but also hardware accelerators brokered by an abstraction layer), to host virtualized functions such as the PHY.

However, shared computing platforms provide a harsh environment for DUs because they trade off the predictability supplied by dedicated platforms for higher flexibility and cost-efficiency. While CUs are amenable to virtualization in regional clouds, virtualized DUs (vDUs) — namely, the vPHY therein — require fast and predictable computation in the edge.

To ease the explanation, we focus on frequency division duplex where uplink (UL) and downlink (DL) transmissions occur concurrently in different frequency bands, and on 5G's baseline numerology ($\mu = 0$ in 3GPP TS 38.211), which yields one transmission time interval (TTI) per subframe (SF), and a SF has a duration of 1 ms. **Fig. 2** illustrates the basic operation of a typical 4G/5G DU processor. Every TTI n , a worker initiates a DU job comprised of a pipeline of tasks (hereafter referred to as DU tasks).

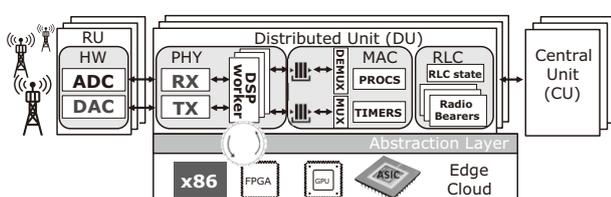


Fig. 1 vRAN architecture.

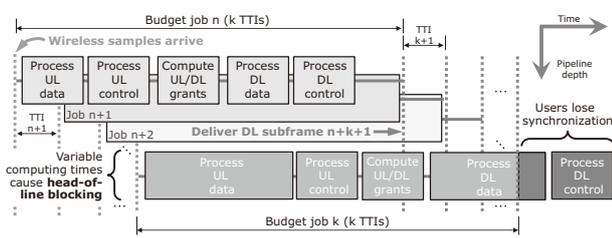


Fig. 2 Every TTI (=1 ms), a worker must execute a DU job, comprised of a pipeline of interdependent DU tasks to process UL SF n and DL SF $n + M$, within $M-1$ ms.

- (1) Process data
- (2) Control channels carried by UL SF n
- (3) Schedule UL/DL radio grants to be transported by DL SF $n + M$
- (4) Process data
- (5) Control channels for DL SF $n + M$

A worker executes a DU job in a thread, using computing resources allocated by a task scheduler; and multiple workers perform DU jobs in parallel to handle one DL SF and one UL SF every TTI, as shown in Fig. 2. 3GPP establishes a 4-ms one-way latency budget between UEs and CUs for eMBB traffic¹. Consequently, there is a hard constraint on M that imposes a computing time budget of roughly $M-1$ ms to process each DU job (usually, $M = 4$).

Indeed, completing a DU job every TTI is vital to preserve synchronization between the BS and its users and thus attain reliability. However, this is challenged by some compute-intensive operations within DU tasks such as forward error correction (FEC). These operations require substantial processing time and the solutions applied today on the market, namely, dedicated hardware acceleration, diminish the very reasons that make virtualization appealing for the RAN in the first place: flexibility and cost-efficiency.

2. Nuberu design

We propose Nuberu, a novel pipeline architecture for 4G/5G DUs that is suitable for shared computing platforms. Our design follows one objective: to build a minimum viable subframe (MVSF) with critical signals for synchronization and control every TTI first to provide reliability first during moments of computing capacity shortage and, provided this, maximize network throughput.

To this end, we set up a deadline within every DU job to begin building an MVSF even if data processing tasks are unfinished. This deadline, depicted in black in Fig. 3, is set such that there is enough time to process an MVSF

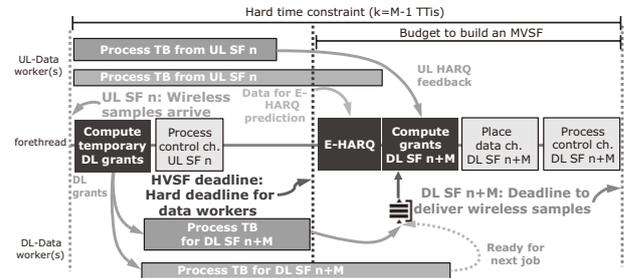


Fig. 3 To provide reliability, Nuberu decouples UL/DL data processing tasks from the rest of the pipeline by integrating a 2-stage radio scheduler, a novel E-HARQ mechanism, and a congestion controller.

before the final job completion deadline (in dark gray in the figure). This is viable because, different from data processing tasks, the tasks involved in building an MVSF require little and roughly deterministic time. To do this efficiently, we need to decouple data processing tasks such that the information required to build an MVSF is ready on time and network throughput is maximized during computing capacity fluctuations. Consequently, we apply the following techniques.

(1) To process DL data channel tasks

- 1) We adopt a two-stage DL radio scheduling approach
 - We issue temporary DL grants as early as possible in the DU pipeline, as shown in Fig. 3. Dedicated workers process (encode, modulate, etc.) these grants in separated threads and store the resulting data in a buffer.
 - Upon the MVSF deadline, final DL data grants are computed based on those already processed successfully that are available in the buffer by that time. Grants generated in a job n that are not processed on time are hence delayed for a later job.
- 2) To mitigate the number of delayed DL data grants, the amount of DL data granted by the temporary scheduler is regulated by a congestion controller that adapts the flow of DL data grants to the availability of computing resources. To this end, Nuberu's radio schedulers use a DL congestion window ($cwnd_{DL}$) that regulate the flow of DL grants. We adopt an additive-increase / multiplicative-decrease (AIMD) algorithm where $cwnd_{DL}$ increases by a PRBs every DU job n ($cwnd_{DL}^{(n+1)} = cwnd_{DL}^{(n)} + \alpha$) as long as congestion is not detected or the maximum PRB capacity is reached, and multiplicatively decreases by $\beta \leq 1$ ($cwnd_{DL}^{(n+1)} = cwnd_{DL}^{(n)} \cdot \beta$) if conges-

tion is detected. Nuberu infers congestion if the buffer of encoded TBs contains $\lambda > 0$ times the vDU's PRB capacity or more.

(2) To process UL data channel tasks

- 1) Dedicated workers process (demodulate, decode, etc.) UL data carried by each UL SF in separated threads.
- 2) Upon the MVSF deadline, an early HARQ (E-HARQ) mechanism infers the decodability of UL data based on feedback from the workers, as shown in Fig. 3. This enables us to estimate the radio information that is required to build an MVSF even if UL data processing tasks have not finished on time. The key idea behind our E-HARQ approach rests upon the concept of extrinsic information, which spawns organically by belief propagation algorithms used by both turbo and LDPC codes. We refer the reader to the literature²⁾ for detailed information about these coding techniques. In a nutshell, belief information is encoded into log-likelihood ratios (LLRs), $L_b := \ln \frac{\text{Prob}(b=+1|\text{input})}{\text{Prob}(b=-1|\text{input})}$, where "input" refers to all the inputs of each decoding node i in a decoder, and b represents the information symbol (bit). The key to iterative decoding is the sequence of a posteriori LLRs of the information symbols, which is exchanged every iteration k between the decoding nodes of the decoder so each node takes advantage of the information computed by the others. To improve the bit estimations every iteration, the different nodes need to exchange belief information that do not originate from themselves. The original concept of extrinsic information was in fact conceived to identify the information components that depend on redundant information introduced by the incumbent code. Such extrinsic LLRs are used to transform a posteriori LLRs into a priori LLRs used as an input in the next iteration. The evolution of the average magnitude of extrinsic information over decoding iterations is easily distinguishable for decodable data. Hence, by observing the pattern of extrinsic information we build a simple predictor that classifies the input data into three types: UNKNOWN, if the uncertainty is too high; DECODABLE, if the predictor is certain the data will be decoded successfully; and UNDECODABLE, if the predictor is certain the data will not be decoded successfully.
- 3) To maximize the predictive performance of E-HARQ, which depends on the amount of work done (iterations) before the MVSF deadline, another congestion controller adapts the allocation

of UL radio resources to the available computing capacity. Like its downlink counterpart, we define a congestion window $cwnd_{UL}^{(n)}$, which bounds the number of UL PRBs that can be allocated to the DU's users. Again, we adopt a simple AIMD algorithm that increases $cwnd_{UL}^{(n)}$ additively ($cwnd_{UL}^{(n+1)} = cwnd_{UL}^{(n)} + \alpha$) when congestion is not inferred, and decreases multiplicatively, ($cwnd_{UL}^{(n+1)} = cwnd_{UL}^{(n)} \cdot \beta$) with $\beta < 1$, when congestion is inferred. The obvious approach to estimate congestion from UL workload is to signal so every time a UL-Data worker does not finish before the MVSF deadline. However, this method does not fully exploit the predictive capability of our E-HARQ mechanism, which can infer the decodability of a TB well before explicit CRC confirmation. Conversely, our approach infers congestion every time E-HARQ cannot provide a prediction with certainty, i.e., outputs UNKNOWN, which occurs every time a UL-Data worker is unable to run sufficient decoding iterations before its deadline, and hence the uncertainty over the prediction is too large.

3. Evaluation

We next evaluate an experimental prototype of Nuberu.

We first set up an experiment with two DUs implemented with vanilla srsRAN ($M = 4$), an open-source implementation of a fully-fledged 3GPP-compliant base station. We associate each DU with one user implemented with srsUE and virtualized over Linux containers sharing 5 Intel Xeon x86 cores @ 1.9GHz.

In this experiment, one DU (vDU 1) transmits and receives as much data as possible. Conversely, the second DU (vDU 2) transmits and receives traffic following a random process with different parameters, which generate normally-distributed computing workload with the mean (line) and variance (shaded area) shown at the bottom of Fig. 4: the higher the load variance of vDU 2, the larger the fluctuations of the computing capacity available for vDU 1.

Fig. 4 (top) depicts vDU 1's relative network throughput in light gray ("Baseline") as a function of the workload produced by vDU 2. The figure shows that the performance of vDU 1 quickly deteriorates. The reason is that, because both vDUs share the same CPU pool, vDU 1 occasionally suffers from CPU resource deficit when vDU 2 produces a peak in demand. As a result, vDU 1 workers executing DU jobs violate their deadline to send out the corresponding DL SF, as illustrated at the bottom of Fig. 2, which causes the user to lose synchronization

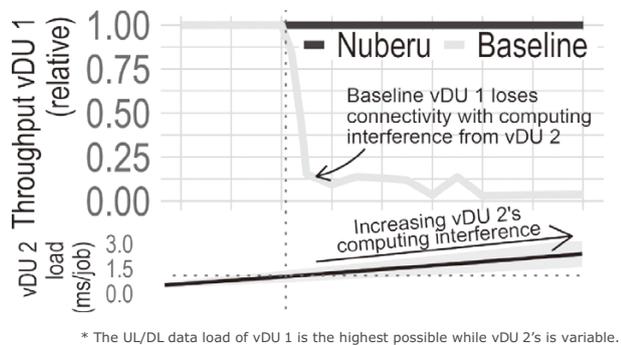


Fig. 4 Two vDUs competing for computing resources.

and throughput to drop.

Conversely, as shown by the dark gray line in Fig. 4, Nuberu can sustain maximum throughput despite severe fluctuations in computing capacity.

References

- 1) 3rd Generation Partnership Project (3GPP): 3GPP TR 38.913; Technical Specification Group Radio Access Network; Study on Scenarios and Requirements for Next Generation Access Technologies; (Release 17). Technical Report, 2022
- 2) Yang Sun and Joseph R Cavallaro: A flexible LDPC/turbo decoder architecture, Journal of Signal Processing Systems 64, 1–16, 2011

Authors' Profiles

GARCIA-SAAVEDRA Andres

Principal Research Scientist
6G Networks
NEC Laboratories Europe

COSTA PEREZ Xavier

Manager
6G Networks
NEC Laboratories Europe

Information about the NEC Technical Journal

Thank you for reading the paper.

If you are interested in the NEC Technical Journal, you can also read other papers on our website.

Link to NEC Technical Journal website

Japanese

English

Vol.17 No.1 Special Issue on Open Network Technologies

– Network Technologies and Advanced Solutions at the Heart of an Open and Green Society

Remarks for Special Issue on Open Network Technologies
NEC's Technological Developments and Solutions for Open Networks

Papers for Special Issue

Open RAN and Supporting Virtualization Technologies

Innovations Brought by Open RAN
Reducing Energy Consumption in Mobile Networks
Self-configuring Smart Surfaces
Nuberu: Reliable RAN Virtualization in Shared Platforms
vrAIn: Deep Learning based Orchestration for Computing and Radio Resources in vRANs

Wireless Technologies for 5G/Beyond 5G

NEC's Energy Efficient Technologies Development for 5G and Beyond Base Stations toward Green Society
Millimeter-wave Beamforming IC and Antenna Modules with Bi-directional Transceiver Architecture
Radio-over-Fiber Systems with 1-bit Outphasing Modulation for 5G/6G Indoor Wireless Communication
28 GHz Multi-User Massive Distributed-MIMO with Spatial Division Multiplexing
28 GHz Over-the-Air Measurements Using an OTFS Multi-User Distributed MIMO System
Comprehensive Digital Predistortion for improving Nonlinear Affection and Transceivers Calibration to Maximize Spatial Multiplexing Performance in Massive MIMO with Sub6 GHz Band Active Antenna System
Black-Box Doherty Amplifier Design Method Without using Transistor Models
39 GHz 256 Element Hybrid Beam-forming Massive MIMO for 8 Multi-users Multiplexing

Initiatives in Open APN (Open Optical/All Optical)

NEC's Approach to APN Realization — Towards the Creation of Open Optical Networks
NEC's Approach to APN Realization — Features of APN Devices (WX Series)
NEC's Approach to APN Realization — Field Trials
Wavelength Conversion Technology Using Laser Sources with Silicon Photonics for All Photonics Network
Optical Device Technology Supporting NEC Open Networks — Optical Transmission Technology for 800G and Beyond

Initiatives in Core & Value Networks

Technologies Supporting Data Plane Control for a Carbon-Neutral Society
NEC's Network Slicing Supports People's Lives in the 5G Era
Application-Aware ICT Control Technology to Support DX Promotion with Active Use of Beyond 5G, IoT, and AI
Using Public Cloud for 5G Core Networks for Telecom Operators

Enhancing Network Services through Initiatives in Network Automation and Security

NEC's Approach to Full Automation of Network Operations in OSS
Autonomous Network Operation Based on User Requirements and Security Response Initiatives
Enhancing Information and Communications Networks Safety through Security Transparency Assurance Technology
Enhancing Supply Chain Management for Network Equipment and Its Operation

Network Utilization Solutions and Supporting Technologies

Positioning Solutions for Communication Service Providers
The Key to Unlocking the Full Potential of 5G with the Traffic Management Solution (TMS)
Introducing the UNIVERGE RV1200, All-in-one Integrated Compact Base Station, and Managed Services for Private 5G
Vertical Services Leveraging Private 5G to Support Industrial DX
Integrated Solution Combining Private 5G and LAN/RAN

Global 5G xHaul Transport Solutions

xHaul Solution Suite for Advanced Transport Networks
xHaul Transformation Services
xHaul Transport Automation Solutions
Fixed Wireless Transport Technologies in the 5G and Beyond 5G Eras
SDN/Automation for Beyond 5G
OAM Mode-Multiplexing Transmission System for High-Efficiency and High-Capacity Wireless Transmission

Toward Beyond 5G/6G

NEC's Vision and Initiatives towards the Beyond 5G Era

NEC Information

2022 C&C Prize Ceremony



Vol.17 No.1
September 2023

Special Issue TOP