

Challenging computer architectures with an intensive vRAN deployment

1st Romain Beurdouche
Communication Systems dept.
EURECOM
Biot, France
romain.beurdouche@eurecom.fr

2nd Raymond Knopp
Communication Systems dept.
EURECOM
Biot, France
raymond.knopp@eurecom.fr

Abstract—Several computer architectures were offered by the industry to provide at the same time the performance and flexibility necessary for implementing virtualized Radio Access Networks (vRAN). Assessing these solutions requires to challenge them in complete full-stack setups at scale. In this demo, we will show how we implement a deployment of several O-RAN Distributed Units (O-DU) on two different systems to observe and compare their respective behavior and performance.

Index Terms—5G, vRAN, O-RAN, Distributed Unit, Baseband Processing, GPU, OpenAirInterface, Open Source

I. CONTEXT

Several silicon industrials offered competing solutions for implementing O-DUs from the O-RAN standard architecture [1]. O-DUs are network components seating a few kilometers from cell-sites that include especially the upper part of the physical layer of the RAN protocol stack which includes demanding computing blocks such as channel coding, precoding or channel estimation. It is also the first network component in the RAN up from the cell that can be virtualized on top of a shared edge infrastructure. Then, Industrials often offer systems relying on combinations of general-purpose processors and Hardware Accelerators (HA) to implement this component. For instance, a combination of an AMD EPYC processor and an RFSoc can be used to implement an O-DU with the OpenAirInterface RAN software stack [2]. NVIDIA's Grace Hopper 200 system, a combination of a powerful Arm Neoverse processor and of a GPU, can as well be used to implement O-DUs with the OpenAirInterface stack [3].

Evaluating the performance and fitness for given use-cases of different computer architectures is key to the design of

future generation of mobile networks. Solutions should be at some point evaluated in operational conditions, with full-stack operation and in conditions in which computer architectures make a difference, in intensive usage conditions. A framework enabling experiments with several O-DUs sharing compute would enable these evaluations on existing and future computer architectures for vRANs. Such a framework is what we want to show in this demo.

II. DEMO CONTENT

In this demo, we will implement a deployment of several 5G small-cell O-DUs on the two systems from AMD and NVIDIA that we introduced above. We will show that this deployment enables us to make relevant observations on the behavior of networks on the two systems under load. This is a protocol with which we already carried out several studies on computer architectures from the industry [2], [3]. The setups we will demonstrate are actually the same that we described with more details and used in [3].

As illustrated in Figure 1, we implement up to 8 O-DUs with the OpenAirInterface software stack on the different systems. Each O-DU implements a 100MHz 4T4R small-cell. The setup we present requires limited real radio appliances. Only one O-DU instance connects to a real radio appliance in order to monitor the end-to-end behavior of a network on the system. Traffic is emulated on the other O-DUs in order that they create the typical load of an O-DU on the system. The O-DUs share the system including the host processor, the Network Interface Card (NIC) and the HA. The sharing of resources and the behavior it induces are key concerns when operating this setup. Sharing the HAs especially can be made in a variety of ways including Single-Root I/O Virtualization (SR-IOV) for the RFSoc of the AMD solution while, on the NVIDIA system, the GPU can be shared by using the Multi-Process Service (MPS) or Multi-Instance GPU (MIG).

REFERENCES

- [1] *O-RAN Architecture Description*, Technical Specification, O-RAN Alliance, October 2025. [Online]. Available: <https://specifications.o-ran.org/download?id=932>
- [2] R. Beurdouche and R. Knopp, "A framework for experimental dense vRAN deployments," in *WiNTECH '26*. Austin, TX, USA: ACM, 2026.
- [3] "Implementing vRANs on an Arm and GPU system," in *WiMob 2026 (Submitted)*, 2026.




Fig. 1: Schematic of the intended vRAN deployment