Plotting the performance landscape for 5G RAN Packet Networks

WHITE PAPER

How Ethernet impairment testing can enable efficient and cost-effective testing of dynamic 5G-RAN packet networks

Contents

Executive Summary 2
5G RAN presents new challenges 3
Broader service support 3
From 4G macro-cells to 5G micro-cells 4
Open, packet-based RAN 5
The transition to packet based fronthaul networks with eCPRI 6
Edge computing 7
Virtual RANs 7
The performance assurance challenge in dynamic 5G networks 8
Plotting the performance landscape with impairment testing 9
What is Ethernet impairment testing? 10
Key focus areas for 5G RAN packet network impairment testing 11
The 5G RAN fronthaul network 11
5G CU/DU placement 12
Time Synchronization 13
Network Sharing and Connectivity Services 13
Impairment testing adds a critical dimension to performance assurance 14
Manage Ethernet uncertainty 15
Cost-effective alternative to field-testing 15
Emulate observed conditions to understand root-causes 15
Plan new features and deployments 15
Plot the 5G RAN performance landscape with Chimera 16
Controlled testing at scale 16
Ideal for 5G RAN emulation 17

EXECUTIVE SUMMARY

5G deployments are now underway across the globe but are merely the first steps in a multi-year journey for the mobile industry. Many practical challenges still remain, not least of which is the challenge of cost-effectively deploying 5G Radio Access Networks (RANs) at scale and assuring the performance of supported 5G services.

The 5G RAN architecture provides an open, virtual, packet-based network that extends from the core to the radio antenna. The virtual architecture of 5G RAN with network slicing and new functional elements, such as the Central Unit (CU) and Distributed Unit (DU), enables multiple demanding services to share the same infrastructure without compromising on their specific performance requirements. Virtual, packet-based RANs also enable network sharing over open interfaces as well as multi-vendor implementations, both of which are important to the 5G business case.

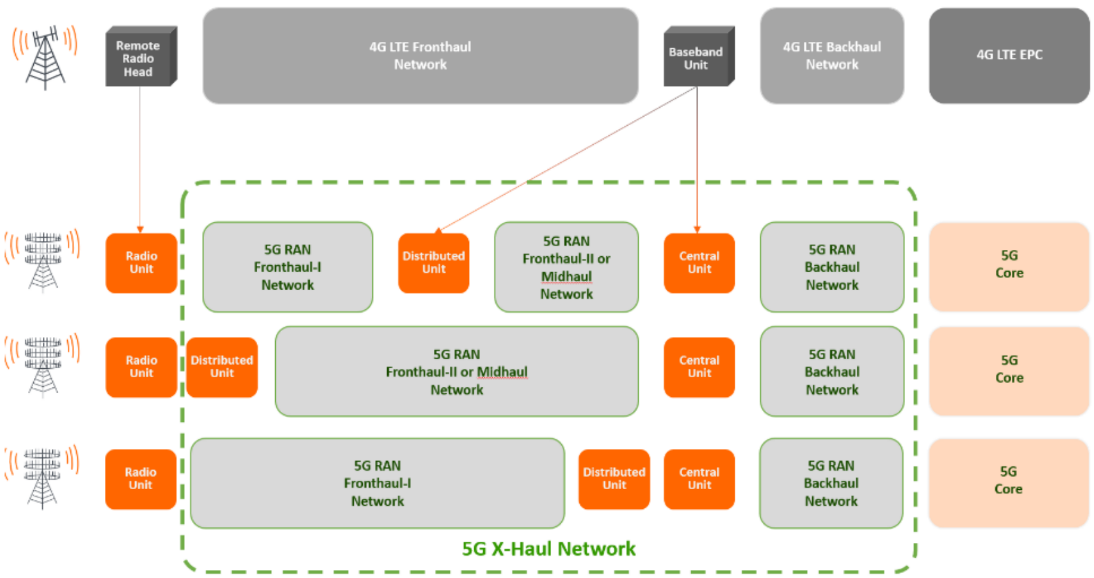

Figure 1: 5G introduces new virtual functions and a more dynamic X-Haul Network

As can be seen in Figure 1, hardware appliances, such as the Remote Radio Head and Baseband Unit of 4G LTE are now replaced with virtual software running on whitebox or Commercial Off-The-Shelf (COTS) hardware. In 5G, the Baseband Unit is split into two new functions, namely the CU and DU. Both these functions can be located either close to the 5G core or close to the Radio Unit (RU) to meet latency and backhaul requirements. This leads to a much more dynamic fronthaul, midhaul and backhaul network, which is now collectively referred to as the 5G cross-haul or “X-haul” network.

It also means that the traditional aggregation architecture where higher capacity is required closer to the core is no longer the norm. As cloud computing and other real-time services move closer to the edge of the network, more traffic is likely to remain close to the subscriber requiring higher capacity X-haul networks with 10G, 25G and even 100G Ethernet connectivity.

The flexibility and dynamic nature of 5G RANs poses new challenges for mobile solution vendors, network operators and service providers. One of those challenges is assuring performance in a highly dynamic environment. Traditional functional, application, integration and protocol testing can be used to test whether 5G equipment is implemented properly and performs as expected. However, these tests are not designed to determine how resilient or robust the equipment will be in the face of network issues that are common in Ethernet networks. 5G RAN packet networks are dynamic with many network configuration possibilities that can change instantaneously. There is, therefore, an additional need to test the limits of network equipment performance under various conditions to ensure dynamic changes in the Ethernet packet network do not lead to service performance issues or outages.

In other words, there is a need to plot the “performance landscape” of the 5G RAN, contouring the performance limits of network equipment and software functions and the 5G RAN itself. Network emulation of the Ethernet-based 5G RAN packet network with impairment testing can be used to stress test network equipment and virtual software in various configurations. This enables vendors to assure performance of their network equipment and software before delivery to network operators. It also enables network operators and service providers to plot their 5G RAN performance landscape at a fraction of the current cost and time and with greater accuracy. This is because testing can now be performed in the lab where a broad variety of scenarios can be tested that are often difficult, if not impossible, to perform in the field.

5G RAN PRESENTS NEW CHALLENGES

5G RAN presents unique challenges that have not been seen before and require new testing methodologies and solutions. To understand these challenges, it is instructive to examine the changes that 5G introduces compared to existing 4G LTE networks with specific focus on the RAN.

Broader service support

Previous generations of mobile networks were designed with specific services in mind. Telephony services have always been the primary offering and Internet connectivity has grown in importance from 3G to 4G. 5G is the first mobile generation designed to support a wide range of services with extremely demanding and very different requirements.

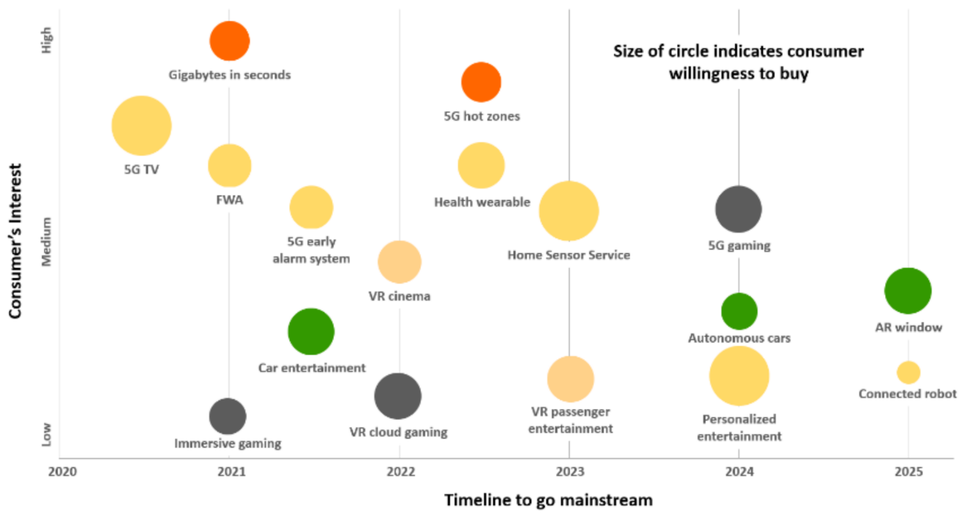

Figure 2: Forecast adoption of new consumer 5G services (adapted from source: https://www.idtechex.com/en/research-report/5g-technology-market-and-forecasts-2020-2030/753)

Enhanced Mobile Broadband (eMBB) services, such as high-definition video streaming lasting for an hour or more, are being supported at the same time as massive Machine Type Communications(mMTC) for industrial Internet of Things services, which can burst small amounts of traffic for short periods of time. Some services will also be critical, such as tele-medicine or autonomous vehicle connectivity, and require Ultra Reliable Low Latency Communication (URLLC). The 5G RAN needs to be able to accommodate the needs of all these services as well as the un-predictability of when the service demand will arise. As can be seen in Figure 2, this multi-service reality is not far in the future and we will see these challenges intensify within the next 5 years.

From 4G macro-cells to 5G micro-cells

The concept of micro-cells was introduced in 3G networks to supplement the capacity of existing macro-cell deployments and provide coverage in rural areas and blind-spots. Micro-cells can also be used to meet specific service needs, such as high-bandwidth or ultra-low latency for a specific application and are therefore a useful tool in meeting multi-service requirements.

In 5G, higher radio-spectrum frequencies are being used that will naturally reduce the range and coverage of 5G cells, but also means that 5G cells are more prone to interference. A typical 4G LTE cell tower can cover a 10 km radius, while a mid-band 5G cell covers up to 4 km and a 5G mmWave cell provides only about 150 m of coverage. 5G RAN deployments will therefore require an order of magnitude more cells than 4G networks, which in turn means an order of magnitude increase in the number of Radio Units (RUs), DUs and even CUs. As we can see in Figure 3, deployments of small cells in

4G networks will continue to meet growing bandwidth demands, but 5G small cells will quickly account for the majority of deployments and continue to grow.

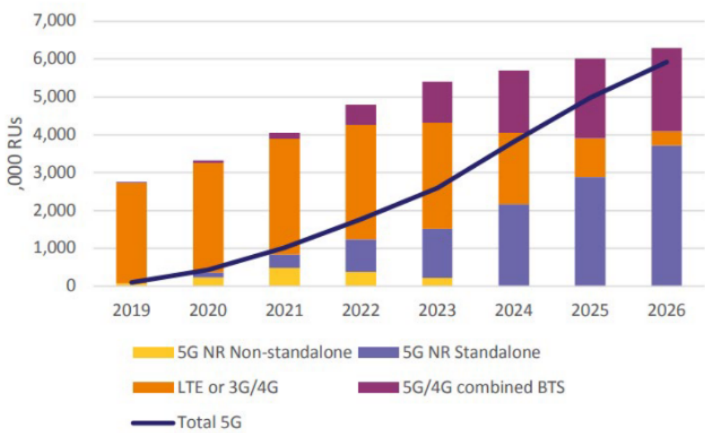

Figure 3: Deployments and upgrades of small cells by technology (Source: Small Cell Forum, Small cells market status report July 2020)

This is expected to increase the cost of 5G RAN deployment. 4G RANs currently account for 57.5% of total mobile network operational costs (OPEX), but with 5G these costs are expected to increase to an average of 67%[1] and possibly more for carriers that heavily rely on mmWave radio spectrum. Sharing the RAN and the cost of deploying 5G cells can reduce RAN operating costs by up to 45%. This is the reason why opening the 5G RAN is so important.

Open, packet-based RAN

4G LTE is already a packet-based network from the core to the antenna. The introduction of the enhanced Common Public Radio Interface (eCPRI) enabled migration to a packet-based front-haul interface in 4G LTE. What is different about 5G is the opening of this interface enabling multiple vendors to deliver RU, DU and CU solutions.

Mobile carriers are driving new consortia focused on defining specifications for open RANs. The O-RAN Alliance has published a number of specifications and driven whitebox implementations of RUs, DUs and CUs. The Telecom Intra-Project (TIP) hosts the OpenRAN project group who are also focused on producing whitebox implementations. The Open Networking Foundation (ONF) has several projects focused on producing open-source implementations of functionality. The Open RAN Policy Coalition is advocating for government policies to drive open RAN adoption.

These organizations are also cooperating. The O-RAN Alliance signed a liaison agreement with TIP OpenRAN in early 2020 to prevent duplication of effort. ONF is also cooperating

[1] Source: Typical network cost splits of active RAN sharing vs no RAN sharing (chart) - STL Partners

with the O-RAN Alliance and using the O-RAN reference architecture for their SD-RAN project. The GSMA has also partnered with the O-RAN Alliance. These agreements underline the importance of open RANs and carrier commitment to making open RANs a reality. One of the technologies that makes this possible is virtualization.

The transition to packet based fronthaul networks with eCPRI

Until recently, the interface between the Remote Radio Head (RRH) and Baseband Unit (BBU) in 4G was based on the Common Public Radio Interface (CPRI), which is still widely deployed. CPRI provides digital radio over fiber (D-ROF) transport of radio signals between the RRH and BBU.

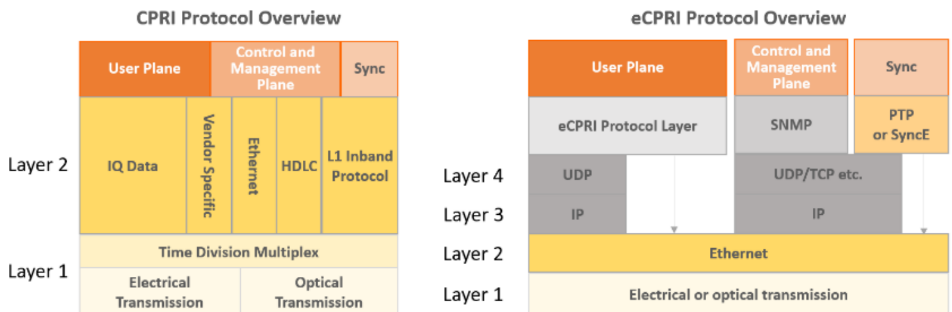

Even though the CPRI signal is digitized, it is not packetized, but carried directly over the fiber interface using the CPRI protocol. This also means that a dedicated fiber is needed from the BBU to the RRH. While CPRI is a standard, there are many options enabling vendor-specific implementations, which effectively means that the RRH and BBU CPRI interface implementations need to be provided by the same vendor.

Enhanced CPRI (eCPRI) was introduced to address these two issues, but also to meet the demands of 5G. eCPRI relies on Ethernet for transport rather than the dedicated CPRI physical layer. This allows much more flexibility in how traffic is transported to the RRH using different types of Ethernet-based connectivity and fronthaul networks. eCPRI is also an open interface allowing any vendor to offer the RRH or BBU.

While Ethernet provides great benefits, it also introduces new issues that need to be managed. Packets can be buffered and delayed in switches. Packets can be delivered out of order or even dropped if there is congestion. The built-in synchronization of the CPRI protocol is now provided separately by the Ethernet-based Precision Time Protocol (PTP) or Synchronous Ethernet, which can also suffer from some of the same packet-based network issues.

These characteristics can seriously affect the operation of the eCPRI protocol as well as the RRH and BBU. It is therefore important to test and characterize the performance of the Ethernet connectivity supporting the eCPRI interface to ensure that issues that can occur in packet networks do not adversely affect services.

An alternative to eCPRI is IEEE 1914.3 Radio over Ethernet (RoE), where digitized radio IQ data is encapsulated in Ethernet frames, which will face similar challenges and require similar solutions.

Edge computing

The broad adoption of cloud services and high-definition streaming is leading to more computing resources being moved to the edge of the network as close to the subscriber as possible. This is challenging traditional network architectures based on bandwidth aggregation, which has been the norm in telecom networks to date. In 5G, it is no longer the case that subscriber traffic needs to be sent to the core of the network for processing. Multi-Access Edge Computing (MEC) components can be located at the edge of the network, which reduces the need to backhaul to the core.

However, this also means that more traffic at higher speeds remains in the X-haul network, which in turn requires higher bandwidth connections. 10G, 25G and 100G connections in the X-haul network are not uncommon and will most likely become more prevalent as consumption of 5G services increase over the next 5 years.

Virtual RANs

One of the major differences between 4G LTE and 5G is the introduction of virtual functions. 4G networks are appliance-based with software functionality tightly integrated with the supporting hardware platforms. By virtualizing 5G, the software-based functions can now be instantiated on any hardware platform meeting the required performance specification. This, in effect, means that 5G functions can be instantiated and moved to any location in the RAN where there is suitable hardware.

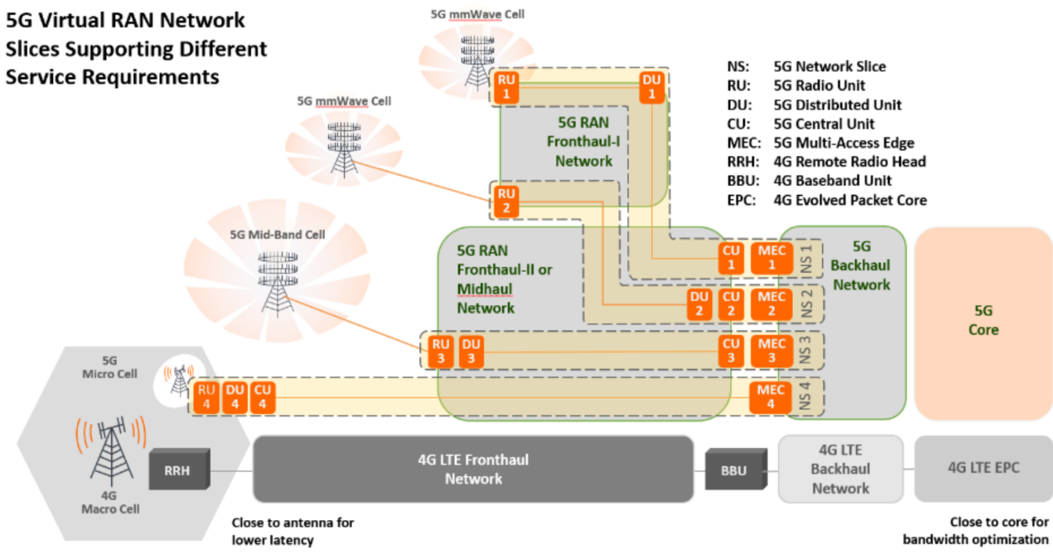

Figure 4: 5G Network slices and virtual function placement options

When this capability is coupled with network-service chaining and the concept of "network slices", then it is possible to instantiate a chain of 5G virtual functions supporting a specific service instance. The individual 5G virtual functions can be deployed in locations that provide the right balance between latency and bandwidth optimization. This enables the performance of the network slice to be tailored to the needs of the specific service. It also means that multiple network slices with vastly different performance requirements can share the same 5G virtual RAN infrastructure.

THE PERFORMANCE ASSURANCE CHALLENGE IN DYNAMIC 5G NETWORKS

The new capabilities provided by 5G RANs also introduce performance assurance challenges.

An order of magnitude increase in the number of 5G cells leads to a corresponding increase in the number of elements that need to be tested. Because the RAN is open, virtual functions can be delivered by different vendors on equipment delivered by yet another vendor. Services might also have to be delivered over another carrier's network through open APIs. Because the functions are virtual, they can be instantiated anywhere, at any time, but must still operate in a network slice that must adhere to stringent end-to-end Service Level Agreement (SLA) requirements. Finally, the service running on this network slice must be delivered while sharing the same network resources as multiple other services with very different performance requirements and at very high speeds, such as 100G.

This leads to a very dynamic network where the configuration can change at any time and where network equipment and functions must be robust enough to adapt and still meet their service delivery obligations. One can no longer assume that if the components of the network meet individual performance specifications then the network itself can meet all of its obligations. It is therefore important to determine how the entire system including supporting functions and equipment from different vendors will perform as one solution under extreme conditions. In other words, it is important to determine the "performance landscape" for the network.

PLOTTING THE PERFORMANCE LANDSCAPE WITH IMPAIRMENT TESTING

Plotting the performance landscape means contouring the limits of network performance. It can be considered a stress test where network conditions are emulated to see how the network functions and equipment, as well as the services they support, respond to different events. The best approach for achieving this is to use Ethernet impairment testing.

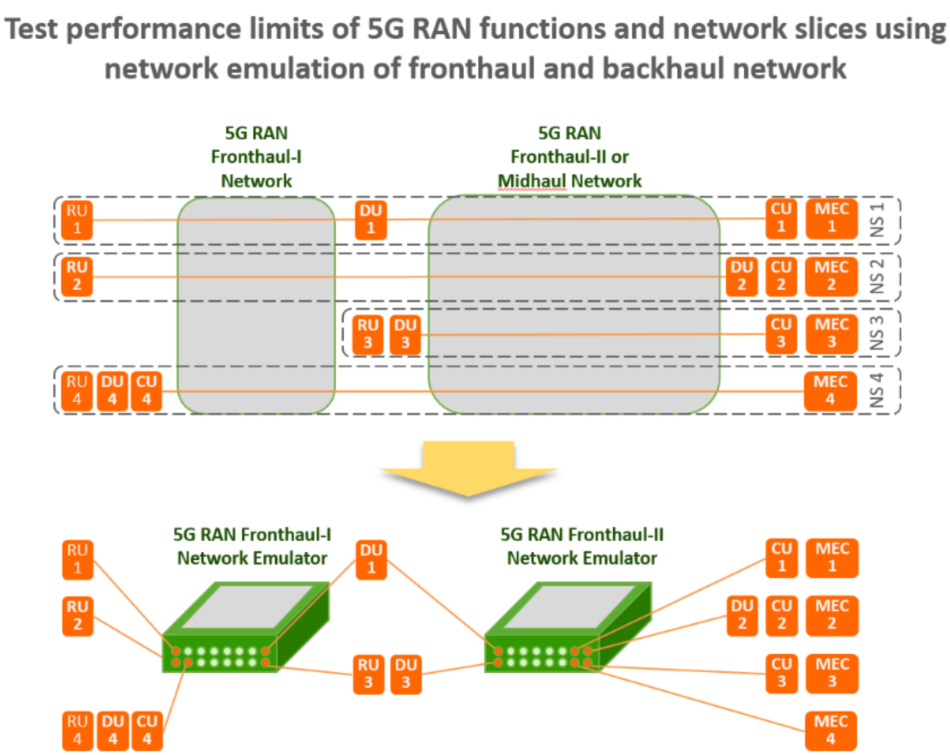

Figure 5: Emulating the 5G fronthaul networks

A network emulator is a device that emulates network behavior so that it appears to attached devices as if they are connected to a network. With a network emulator, it is possible to change how the "network" is behaving using configurable parameters. It is also possible to introduce error conditions or impairments, such as high latency, jitter, packet loss or other errors that emulate what could happen in real networks.

What is Ethernet impairment testing?

In Ethernet packet networks, issues can occur that hinder the timely and orderly delivery of data. These are "impairments" and can be issues like packet drops, packet duplication, packet latency and jitter or packets arriving out-of-order. In real-world packet networks like Ethernet, impairments can occur under network congestion or network switch overload. The issues occur spontaneously, often in very short periods of time, but can have significant impact on network performance and service delivery.

To test the resilience of network equipment in such situations, a mechanism is needed to introduce impairments into the network. However, this can be difficult to control in a live network. A network emulator simulates the behavior of an Ethernet network, but since it is a self-contained appliance, it provides predictable performance without any outside influences. This allows impairments to be introduced in a controlled manner and their affects on devices under test to be examined.

The impairments can be introduced individually as a single event or based on a schedule and distribution profile that emulates how a specific Ethernet network behaves. For example, introducing a packet drop impairment at regular intervals over a specific period of time and for specific types of Ethernet traffic. Impairments can also be combined to emulate complex error conditions. This can be used to reproduce behavior observed in real networks, but which is hard to isolate in a live network.

Ethernet impairment testing is a cost-effective alternative to field testing. Tests can be performed in the lab at fraction of the cost. In addition, there are no limits to the range of tests that can be performed including the ability to automate allowing continuous repeated testing. This allows specific scenarios to be tested that can be difficult to test in the field.

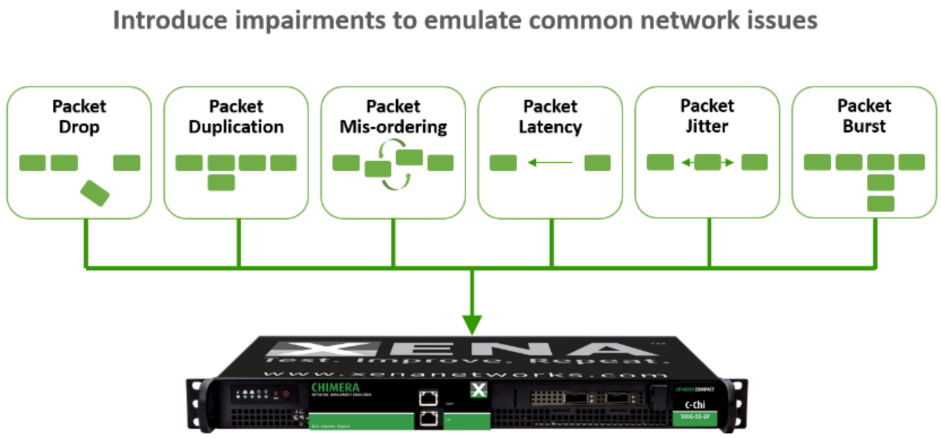

Figure 6: Network emulator impairments

Using a network emulator and impairment testing, it is possible to introduce error conditions in the emulated 5G RAN packet network and thereby test how attached network functions and equipment respond to these error conditions. By plotting the results for different types of impairment test, it is possible to develop an overview, or landscape, of the performance limits of the network functions and equipment under test as well as the individual services supported by the RAN. The technique could also be used to plot the landscape for the 5G packet network itself when it relies on other carrier's networks for fronthaul or backhaul connectivity and can form the basis for SLAs between carriers sharing network infrastructure or providing connectivity services.

KEY FOCUS AREAS FOR 5G RAN PACKET NETWORK IMPAIRMENT TESTING

There are key areas of the 5G RAN packet network that need special attention with respect to the new capabilities offered by 5G RANs:

- The 5G RAN fronthaul network
- 5G CU/DU placement
- Time synchronization
- Network sharing and connectivity services

The 5G RAN fronthaul network

The 5G RAN fronthaul network connecting RUs and DUs relies on the Ethernet-based eCPRI or IEEE 1914.3 Radio over Ethernet (RoE) protocols. This enables the RU and DU to be placed up to 50 km apart or co-located, and any location between these two extremes. However, the eCPRI protocol stack is split between the RU and DU with 5G implementations based on the so-called “lower-layer split” or “split-PHY” model, as shown in Figure 7. This means that the Ethernet-based fronthaul network connecting the upper and lower parts of the PHY in eCPRI needs to meet stringent requirements.

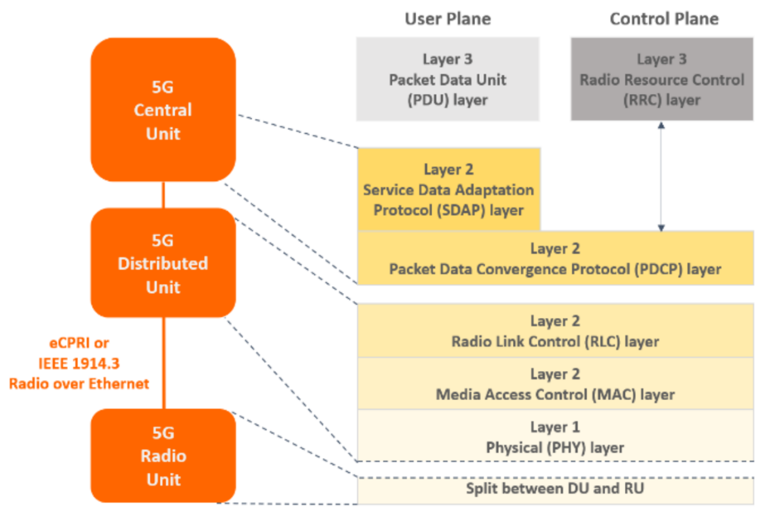

Figure 7: The lower layer split-PHY and PDCP-RLC upper-layer splits

Using a network emulator and impairment testing, it is possible to emulate the underlying packet-based fronthaul network and plot the performance landscape for the eCPRI implementation in both the RU and the DU. It is possible to see how the RU and DU will respond to packet loss, jitter, latency and time synchronization issues and thereby plot the contour of performance limits for these implementations. This can provide insight that can be used as deployment guidance and requirements for underlying fronthaul packet networks.

5G CU/DU placement

The 5G RAN fronthaul network connects the CUs and DUs over the so-called “upper layer split” between the Packet Data Convergence Protocol (PDCP) and Radio Link Control (RLC) layers, as shown in Figure 7. This enables the CUs and DUs to be co-located or separated. The placement of CU’s and DUs is dictated by the performance requirements of the service to be supported.

For example, if ultra-low latency is important, then the CU and DU can be co-located with the RU. If, on the other hand, bandwidth capacity is more important, then the CU and DU can be co-located centrally with the Multi-Access Edge Cloud (MEC). The placement of the CU and DU can thus be tailored to meet trade-offs in latency and bandwidth performance.

Ethernet impairment testing can not only be used to emulate the underlying Ethernet network to determine how CU and DU implementations perform under various conditions, but can also be used to emulate where CUs and DUs are physically instantiated in the network and how that affects the performance profile. For example, the network emulator can introduce delay and jitter that emulates the effect of placing a CU or DU in different locations. This can be used to determine the optimal placement of CUs and DUs depending on network conditions as input to network orchestration and control policies.

Time Synchronization

One of the strictest requirements for 5G networks, and 5G RAN in particular, is end-to-end time synchronization requirements. End-to-end time error budgets have been defined of +/- 1.5 microseconds from the core to the antenna. Between RU antennae, the timing error should be under +/- 65 nanoseconds.

The 5G RAN relies on the Ethernet-based Precise Time Protocol (PTP) telecom profiles defined by ITU-T in G.8275.1 and G.8275.2. For network emulation of Ethernet packet networks, it is important not to interfere with the operation of PTP and therefore network emulators should act as Telecom Transparent Clocks (T-TCs) by measuring their own contribution to delay and update packet time stamps in order to remain invisible to PTP. As a T-TC, the emulator must also be able to forward Synchronous Ethernet timing, which is required in the PTP telecom profiles.

Network Sharing and Connectivity Services

Considering that 5G RANs can account for as much as 67% of total network operational costs or higher, it is no surprise that carriers are considering sharing RANs or availing of connectivity services provided by other carriers. However, it can be difficult to determine an appropriate SLA for dynamic 5G networks.

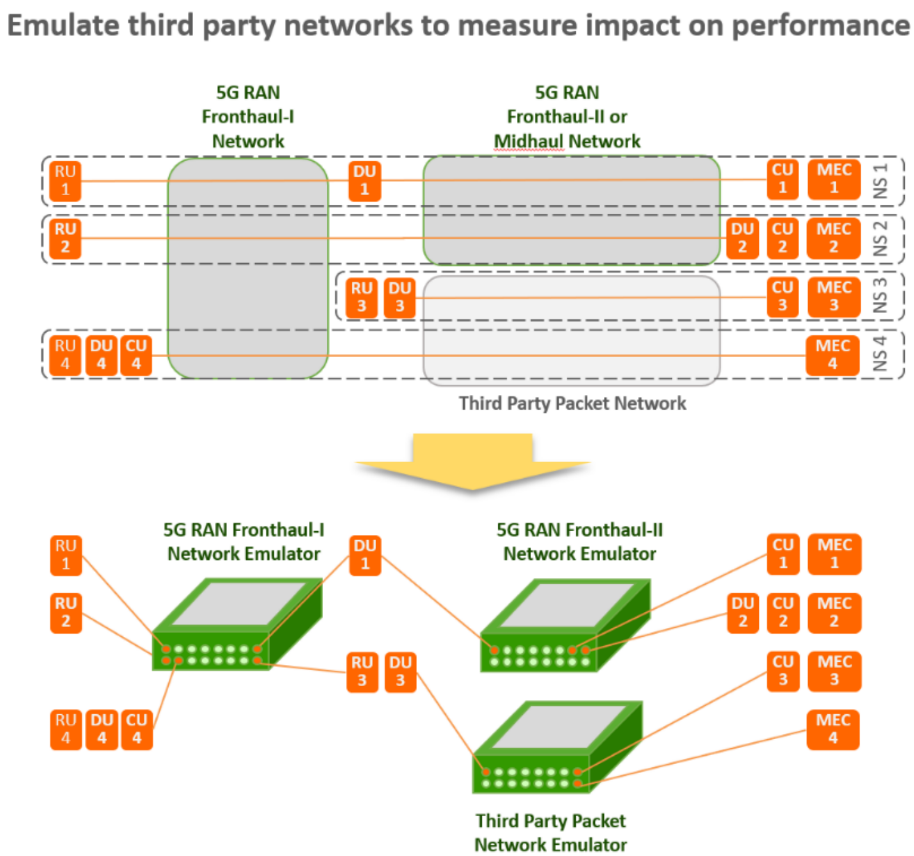

Figure 8: Emulating third party packet networks in the RAN

Ethernet impairment testing can be used to emulate the performance of a third-party network and to plot the performance landscape as input to SLA requirements. Using impairment tests, the effect of maximum packet loss, latency, jitter, time synchronization time error and other parameters on performance can be examined. These results can be used to determine what can be tolerated before service performance is affected.

IMPAIRMENT TESTING ADDS A CRITICAL DIMENSION TO PERFORMANCE ASSURANCE

Stateful and stateless testing are critical and necessary testing methodologies that are highly relevant for 5G RANs. Nevertheless, Ethernet impairment testing adds a new, critical test dimension that provides invaluable insight and tangible benefits.

Manage Ethernet uncertainty

The adoption of eCPRI and RoE in the 5G RAN introduces Ethernet networks from antennae to the core. This provides a great deal of flexibility and openness, but also introduces extra uncertainty as packet networks like Ethernet are not deterministic. Issues such as packet loss, delay, jitter, re-ordering and duplication can arise and it is important to be prepared for these unpredictable real-world events.

With network emulation, it is possible to test the resilience and robustness of 5G function and equipment implementation to changes in the Ethernet packet-based network providing X-haul connectivity.

Cost-effective alternative to field-testing

Plotting the performance landscape is not new. However, it is usually performed in the real-world during field trials and live deployments where feedback is used to improve new versions of network equipment and functions. This is costly, risky and takes years.

Ethernet impairment testing can be performed in a lab environment and address the vast majority of potential conditions, events and other performance cases that could be encountered in live deployments. In fact, a meticulous impairment test can provide more accurate and complete information on performance under various conditions. Feedback from field trials and live deployments can be incomplete, as extreme conditions might never be encountered. This is, in reality, a worst-case scenario, as failure when extreme conditions are encountered could have catastrophic consequences.

The ability to automate these tests with precise schedules and test distributions allows specific scenarios and conditions to be emulated over a period of time that provides invaluable insight into the longer-term effects of Ethernet network configurations that can eventually lead to major issues.

Emulate observed conditions to understand root-causes

5G RAN architectures can be extremely complex and when a performance issue is encountered, it can be difficult to determine the root-cause. Ethernet impairment testing can be used to reproduce the behavior observed and provide insight into what the root-cause could be. This has the advantage of not interfering with the existing network that is also supporting multiple other services that are not affected. It can also be conducted in a lab together with relevant vendors and other network providers collaborating on resolving the issue.

Plan new features and deployments

Network emulation allows new features in network functions and supporting equipment to be tested accurately before deployment as the existing network conditions can be emulated accurately. New network deployment configurations can also be tested in the same way allowing network architects and engineers to experiment with different deployment strategies. “What-if” test scenarios can be defined to anticipate potential issues. For example, the maximum number of hops in a supporting network can be emulated to determine how extra hops affect overall performance.

PLOT THE 5G RAN PERFORMANCE LANDSCAPE WITH CHIMERA

Ethernet impairment testing provides a new dimension in testing for 5G RANs that can be used to plot the performance landscape of 5G RAN virtual functions and their supporting equipment.

TEST MODULE (Chi-100G-5S-2P)

Chimera is currently available as a 2-slot test module for a ValkyrieBay chassis, the Val-C12-2400G, where ValkyrieManager is used to emulate impairment as part of it standard user-interface.

STANDALONE CHASSIS (C-Chi-100G-5S-2P):

ChimeraCompact is a standalone version of Chimera where the test module is installed in a quiet desktop-sized chassis.

Xena Networks provides a comprehensive Ethernet impairment testing product named Chimera. Chimera can be used to plot the performance landscape of 5G functions and supporting hardware by emulating 5G Ethernet-based X-Haul packet networks and issues, such as packet drops, latency and jitter. This provides insight into the resilience and robustness of 5G functions and their supporting hardware in the face of real-world Ethernet network issues.

Controlled testing at scale

Chimera supports multiple ports of 10G, 25G, 40G, 50G and 100G allowing a number of devices under test simultaneously. A comprehensive range of Ethernet impairments can be introduced on a per port and per-flow basis. Each impairment can be configured for a specific behavior in a specific time period that can be repeated. This allows complex combinations of tests that can accurately even the most complex network behavior. The network emulator provides an easy-to-use and understand Graphical User Interface (GUI). Nevertheless, the solution is entirely script driven enabling configuration and reporting to be automated and controlled using Chimera’s strong scripting interface.

Ideal for 5G RAN emulation

Chimera is fully integrated with the Xena Networks Valkyrie stateless traffic generator allowing a complete testing solution for 5G RAN emulation. Chimera is designed for 5G networks and is implemented as a PTP T-TC with Synchronous Ethernet forwarding that allows Chimera to emulate Ethernet network connectivity without affecting 5G RAN time synchronization performance. All 5G X-Haul networks can be tested including the effect of Ethenet impairments on eCPRI and RoE fronthaul networks. Custom Ethernet impairment distributions allow tailored impairments to be defined that emulate specific Ethernet network behavior all of which can be automated with configurable durations and repetition.

For more information see https://xenanetworks.com/chimera/.

>> Book a consultation with a Xena tech expert