



Enabling the Next- Generation of 5G Platforms

WP029

2022.09.01

Copyright and Trademarks

Copyright © 2022 Achronix Semiconductor Corporation. All rights reserved. Achronix, Speedster and VectorPath are registered trademarks, and Speedcore and Speedchip are trademarks of Achronix Semiconductor Corporation. All other trademarks are the property of their prospective owners. All specifications subject to change without notice. WP029 V1.1

Notice of Disclaimer

The information given in this document is believed to be accurate and reliable. However, Achronix Semiconductor Corporation does not give any representations or warranties as to the completeness or accuracy of such information and shall have no liability for the use of the information contained herein. Achronix Semiconductor Corporation reserves the right to make changes to this document and the information contained herein at any time and without notice. All Achronix trademarks, registered trademarks, disclaimers and patents are listed at <http://www.achronix.com/legal>.

Achronix[®]
Data Acceleration

Achronix Semiconductor Corporation

2903 Bunker Hill Lane
Santa Clara, CA 95054
USA

Website: www.achronix.com
E-mail : info@achronix.com

Abstract

Radio access networks (RANs) and associated core network hierarchy, which link end-user equipment to both the central telecom network and the cloud, are essential in building ubiquitous cellular connectivity to expand the number and breadth of use cases that can be supported by the technology. Having a high-level understanding of what 5G entails and understanding where, how and when the technology will evolve can help to manage expectations when forming a strategy for developing and implementing a 5G RAN and core equipment. This paper outlines the current status of 5G standards and rollout, summarizes the new use cases 5G RANs need to support and examines the standards evolution to support higher bandwidth and additional use cases. Finally, the paper also explains how Achronix FPGA technology can be utilized by developers to meet the fundamental challenge facing them — supporting the optimization of 5G RAN architectures by offloading some of the processing from CPUs onto FPGA-based accelerators in a cost-, power- and area-efficient manner.

5G Deployments and Macro Trends

Clearly, 5G is now much more than just the next cellular technology for handset connectivity. 5G and the evolution of cellular connectivity enable multiple new use cases and open up a new commercial opportunity for companies not previously having cellular connectivity as part of their product mix. 5G is not just about providing telecoms connectivity anymore, but about enabling connectivity for a variety of other use cases such as industrial IoT, automotive, smart city and other applications. 5G is about enabling the connection of the billions of new devices such as cameras and a variety of other sensors in homes, cities and factories, providing remote medical support for doctors and patients, supporting convergence with IT technologies, and general replacement of tethered, wireline connectivity.

Fundamentally, 5G is a more spectrally efficient implementation of cellular technology than previous generations, with significantly additional air interface capacity combined with beam-forming/steering technology and the ability to aggregate 4G and 5G channels, all being utilized to good effect.

5G infrastructure deployments are beginning to ramp up with estimates showing a faster adoption than 4G, with mobile network operators (MNOs) rolling out 5G to reach one billion subscribers two years ahead of where it took 4G to achieve similar levels.

The following table describes the macro trends steering the development and evolution of the technology in the future.

Table 1 – Macro Trends Affecting 5G Evolution

Macro Trend	Description	Impact
Geopolitical pressures	Restriction on the western MNOs from using the China-based vendors has placed more of an emphasis on trying to generate a wider supplier base other than the current market-leading Tier 1 suppliers. This effort to generate a wider supply base aligns with the introduction of a series of new disruptive initiatives such as Open RAN and the ORAN Alliance.	<ul style="list-style-type: none"> • Reduced number of approved suppliers necessitates standardized interfaces to enable interchangeability between vendors • ORAN reference implementations to enable T2 and emerging OEMs to provide alternative, better solutions

Macro Trend	Description	Impact
Expanded 5G use cases	Enhancements to specifications to support ultra-reliable, low-latency communications (URLLC) and massive machine-type communications (MMTC) use cases are really the targets of 3rd Generation Partnership Project (3GPP) Release 17 and 18, where more efficient use of radio resources is combined with machine learning techniques to support high connectivity densities and lower-latency decision making.	Heavy increase in processing load on the air interface traffic for mid-band deployments means that new architectures need to be employed with acceleration offload from CPU subsystems
More stakeholders driving 5G evolution	<p>There is a blurring of the historical boundaries that existed in 4G, with an increasing diversity of stakeholders with an interest in the technology, including:</p> <ul style="list-style-type: none"> • Companies who did not historically have any cellular capability or knowledge, for example cloud providers looking at how they can utilize their cloud expertise in hosting workloads on 5G • Companies such as Microsoft with Azure and AWS with their Outpost Edge deployments • Industrial/automotive companies who have a specific problem to solve utilizing 5G 	New entrants influencing the direction of the evolution of 5G technology need to be considered. Consideration should be given to how cloud and radio technologies can be applied to vertical markets (e.g., industrial and automotive) with different players bringing their relative areas of expertise.

Building Blocks to Transform 5G

Previous RAN architectures (2G, 3G and 4G) were based on monolithic building blocks, where few interactions happened between logical nodes. Since the earliest phases of the new radio (NR) study, however, it was felt that splitting up the gNodeB base station (gNB) (the NR logical node) between centralized units (CUs), distributed units (DUs) and radio units (RUs) would bring added flexibility. Flexible hardware and software implementations allow for more scalable and cost-effective network deployments — but only if the hardware and software components are interoperable and can be mixed and matched from different vendors. A split architecture (between central and distributed units) allows for coordination for performance features, load management, real-time performance optimization and enables adaptation to various use cases. This split architecture also delivers the quality of service (QoS) needed by various applications (i.e., gaming, voice and video), which have variable latency tolerance and dependency on transport as well as different deployment scenarios, such as rural or urban, that have different access to transport, such as fiber versus wireless. **Figure 1** introduces the main building blocks required for a 5G deployment.

5G is no longer just about a RAN but requires technology that encompasses the full network connectivity from client to data center. Historically, intelligence was located at either end of the cellular network, in the client, base station and core of the network. As we head towards a trillion connected devices, MNOs can no longer add more and more capacity for data backhaul from the radio, into data centers for application processing and then back to the

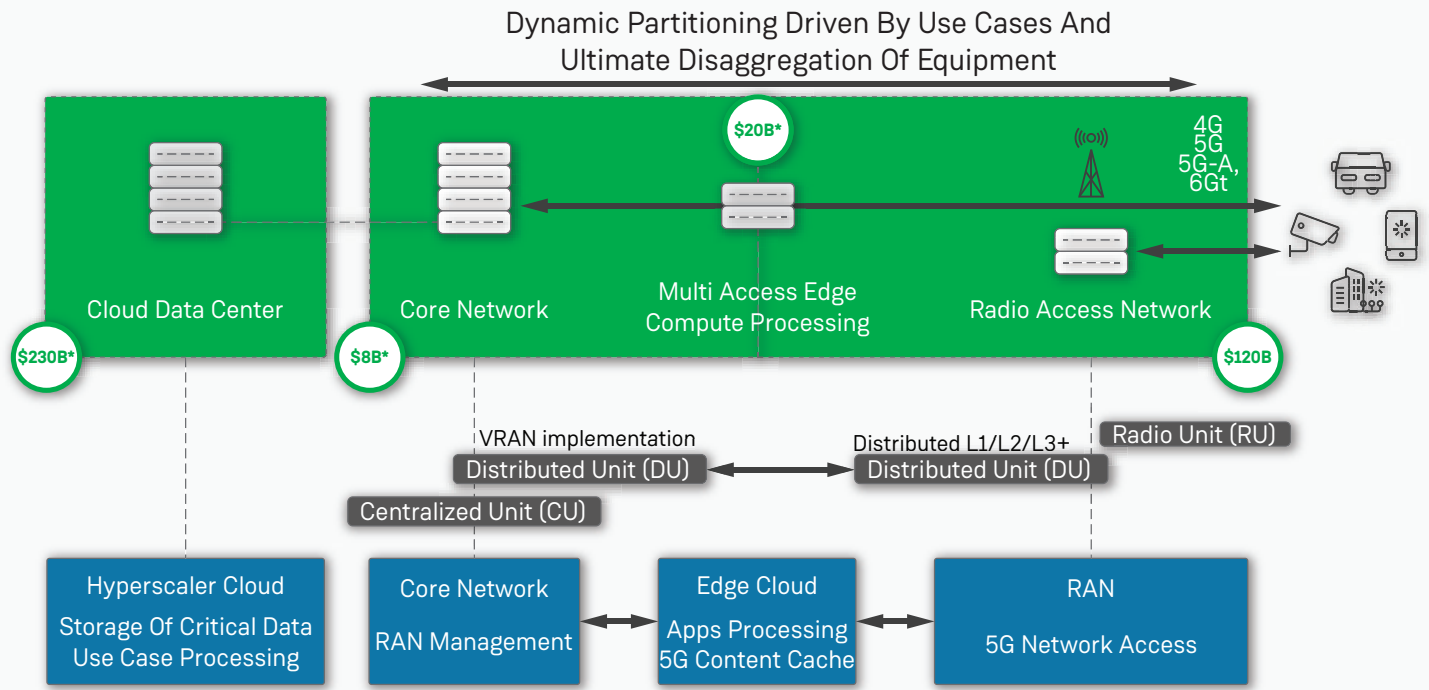


Figure 1 - Building Blocks to Transform 5G

client devices. For example, an increase from 400M connected image sensors today to 1B connected image sensors alone would increase network traffic to 400 Exabytes from today’s around 150 Exabyte mark.

One way to address this capital expenditure issue is to distribute intelligence more evenly across the network. This change necessitates the distribution of more compute power to enable making quicker and more efficient decisions. For example, the block in the diagram above labeled Multi Access Edge Compute Processing represents the type of addition supporting this distribution of intelligence.

The circled dollar values in the figure above illustrate the estimated equipment spend in the RAN and network hierarchy over a four-year period. The spend on the radio network itself is very significant at \$120B targeting a significant R&D spend.

The figure above represents the different elements that comprise a 5G radio network. To support the range of different 5G use cases from enhanced mobile broadband (eMBB) and massive machine type communications (mMTC) to ultra-reliable, low-latency communications (URLLC) requires flexibility in where these elements are located physically in the network. For example, the diagram represents how the distributed unit (DU) may be facilitated as a standalone unit close to the radio unit (RU) to support low-latency, more real-time requirements for 5G, whereas for non-latency intensive applications such as eMBB, the DU may be co-located with the CU in a vRAN-like deployment.

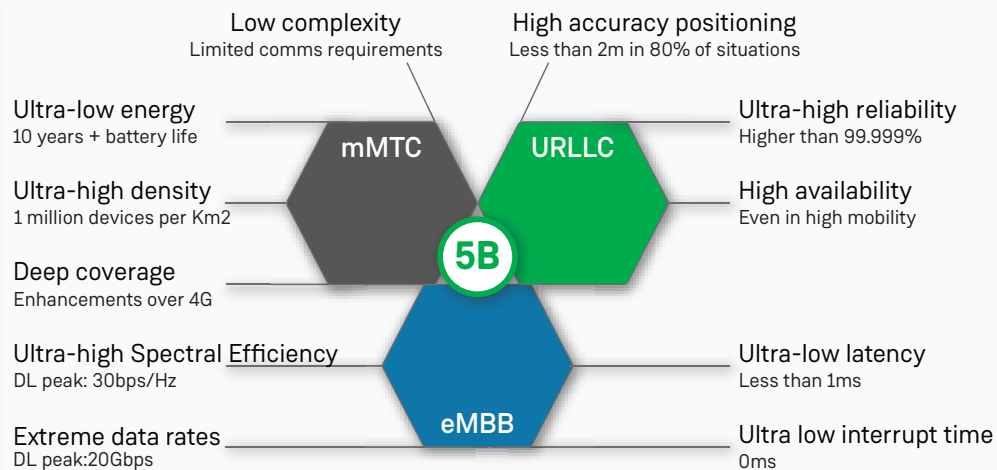
The need for flexibility drives a need for the building blocks used for these designs to be equally flexible and allow these designs to be partitioned in a number of ways with common elements. Diversity of SoC designs and how accelerator capabilities are implemented are important factors in meeting these challenges.

What Use Cases Will 5G RANs Need to Support?

As one of their first steps in defining 5G, the International Telecommunication Union Telecommunication Standardization Sector (ITU-T) identified the ways that cellular is and will be used by consumers, enterprises, and industry. 3GPP then set about implementing the needed standards. As part of what the 3GPP called their SMARTER project (Study on New Services and Markets Technology Enablers), they identified high-level current and future use cases for cellular as well as the needed features and functionality.

In addition to a category called fixed broadband, the body defined three categories of mobile use cases: mMTC, eMBB, and URLLC (Figure 2). While these category names are not particularly catchy, they have become industry-standard terms:

- mMTC** - Massive machine type communications introduces support for machine-to-machine interactions happening on a very large scale and includes battery-operated IoT devices. By and large, these devices need relatively low latency, highly reliable connectivity with energy efficiency. The challenge is to provide scalability and consistent connectivity for the billions of IoT devices that communicate relatively infrequently and in short bursts. Wide coverage and deep indoor penetration are important, as are low cost.
- eMBB** - If mMTC mostly addresses how machines use cellular, eMBB mostly addresses how people use cellular. This category includes activities such as 8K video streaming, immersive augmented reality/virtual reality (AR/VR), connected transportation infotainment, and enterprises supporting mobile broadband connectivity. The key requirements in this category are ultra-high spectral efficiency, extreme data rates, and ultra-low interrupt time. All of these requirements are addressed by the 5G NR defined in Release 15. As infrastructure support for the 5G NR begins to expand, these use cases become more widely available. This category can be considered a mix of evolution and revolution, since laptops using cellular for connectivity are not exactly new, while immersive AR/VR and other data-intensive applications have not really been possible with previous generations of cellular.
- URLLC** - Providing support for ultra-reliable, low-latency communications as a class of service is the truly revolutionary aspect of 5G, since it delivers a level of performance not yet seen in real-world applications. Adding support for URLLC enables applications such as intelligent transportation, including vehicles that can navigate complex road situations and avoid collisions by cooperating with each other, and use cases



98246942-02.2022.05.08

Figure 2 – 5G Use Case Categories

associated with the fourth industrial revolution, including time-critical factory automation. It also includes remote healthcare, which includes devices that measure vital signs and either automatically or semi-automatically respond as needed, as well as remote treatment, including surgeries performed in ambulances, during disaster situations, or in remote areas while using real-time guidance from an off-site doctor. In all these situations, the connection needs to be exceptionally stable and needs to operate with end-to-end latency rates on the order of a millisecond or less. The features needed to support URLLC are largely being defined in Release 16 and 17 of the 3GPP specifications. In other words, URLLC represents the future of 5G, even if that future is only a few years away.

The various features added to each 3GPP specification release are intended to address different aspects of these three categories. The specific use cases that are already active today or close to arriving are addressed in earlier specification releases, while use cases that are farther in the future are addressed in later releases — all part of the ongoing evolution of 5G.

Evolution of Requirements to Meet 3GPP Releases 17 & 18

The evolution of 5G brings a series of new standards that are agreed to by multiple companies participating in the ETSI 3GPP organization. But what are the likely technology requirements that result from the evolution of 5G standards?

The diagram below represents the current state of the 3GPP release process. Equipment being deployed in today's 5G networks is largely comprised of technology specified in Release 15 and 16 of the 3GPP specifications. The more advanced use cases and resultant network requirements are met with future releases of 3GPP specifications (Release 17 and 18).

Today, 3GPP has passed the midpoint of its work on Release 17 (Rel-17) with plans to publish in the mid-2022 time frame. Meanwhile, the discussions around the scope of Release 18 (Rel-18) are well underway. Rel-18 and subsequent releases have been termed 5G Advanced by 3GPP to recognize the evolution of the technology.

Features in Rel-17 are intended to enhance network performance for existing and new use cases. These new features are classified in the diagram below into three categories:

Air interface and management features:

- Upper and lower split L1 processing and offload — L1 kernels acceleration for uplink and downlink channels
- Complex L1 MAC scheduling acceleration
- Spectral efficiency, beam management and dynamic spectrum sharing
- Flexible DFE processing/offload

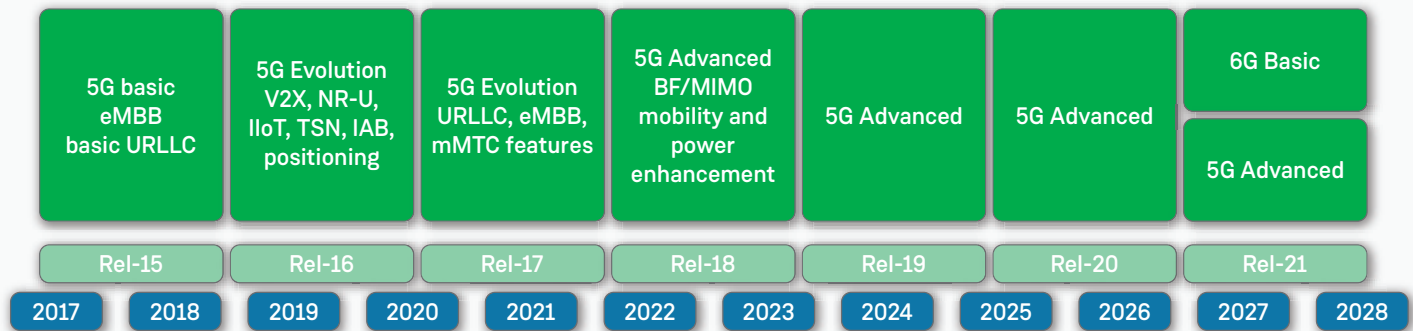
Connectivity and security:

- eCPRI offload and processing (Split 7.x DU/RU flexibility)
- Backhaul and security offload
- Network processing and balancing, including buffer and queue management

Compute and Application acceleration:

- C and U plane management: application of ML/AI to user route selection policy
- Network data analytics
- Edge compute hosting closer to radio
- ML radio and application-based processing

These categories and features are discussed in more detail later in this paper.



98246942-03.2022.05.08

Figure 3 – 3GPP Specification Release Timelines

Rel-18 or 5G Advanced (5G-A) takes Rel-17 a stage further, providing more intelligent network solutions by embedding ML technology in the radio and network hierarchy in order to support new additional use cases and increase efficiency in the network. Looking at changes specifically in the radio, Rel-18 (Advanced Antenna Systems) is the main tool for supporting increases in spectral efficiency, bringing further enhancements in beam forming and massive multiple input/multiple output (MIMO) specifically for the mid-band, sub-6 GHz spectrum.

In terms of new use cases for 5G-A, aside from automotive and industrial, there are national security and public safety applications where these new features could be used to support remote control of drones and rogue drone detection for example.

5G Network Hierarchy and Radio — Next-Generation Networks Drive the Need for Diverse Solutions

There are numerous drivers influencing the need for platform diversity. There has been a desire from mobile network operators to base their networks on network function virtualization (NFV) and software-defined networking (SDN) based technology running on commercial off-the-shelf (COTS) servers. However, Achronix believes that one homogeneous design cannot meet all the requirements for 5G moving forward. Diverse workloads place different pressures on the network, driving differing solutions to meet these requirements.

Benefits of an architecture with the flexibility to split and move 5G NR functions between centralized and distributed units allow:

- More flexible hardware implementations, allowing more scalable, cost-effective solutions.

- Coordination of performance features, load management, real-time performance optimization, and enabling NFV/SDN techniques according to use case.
- Different deployment scenarios to enable different use cases (eMBB, uRLLC, mMTC). These different deployment scenarios in turn support evolution of radio technologies by accommodating changes in network hierarchy/architecture (e.g., ORAN) and dynamic allocation of network resource by new features such as network slicing.

New network/function splitting is likely to influence the need for diverse equipment and SoC options.

Figure 4 shows the different option splits specified in 3GPP standards to support emerging use cases and different resulting traffic types. This diagram shows the different L1, L2 and L3 splits with functions running on the CU, DU and RU. Two of the most popular options are:

- L2 Option 6 split, where the upper layer functions are centralized in the network but the radio-specific scheduling of traffic and radio link control are pushed closer to the radio network.
- L1 Option 7.x split where the upper layer 1 processing is centralized with the L2 and L3 functions and only the lower layer 1 Phy is populated in the RU.

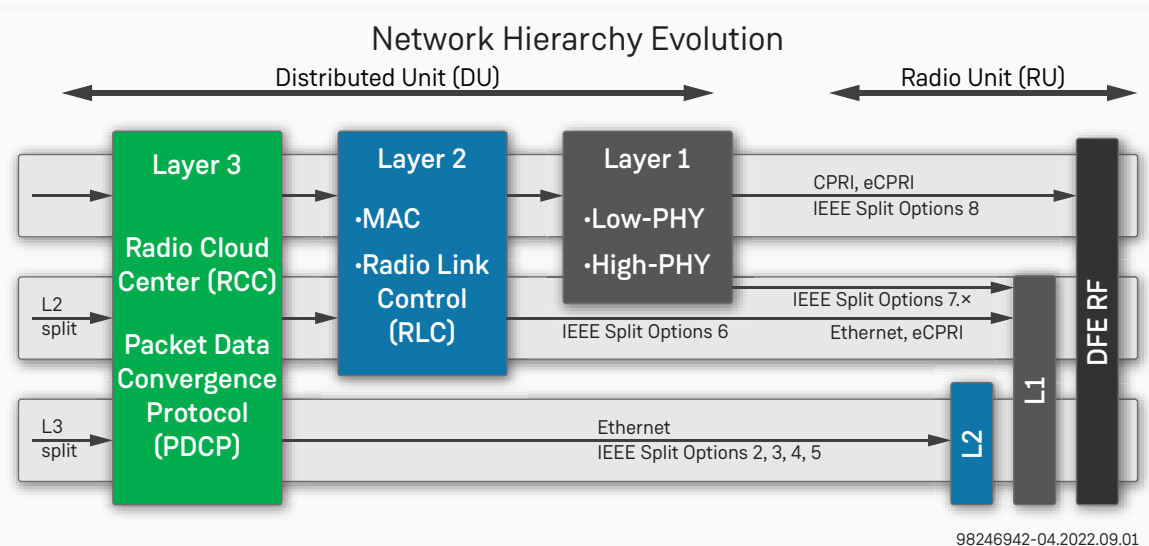
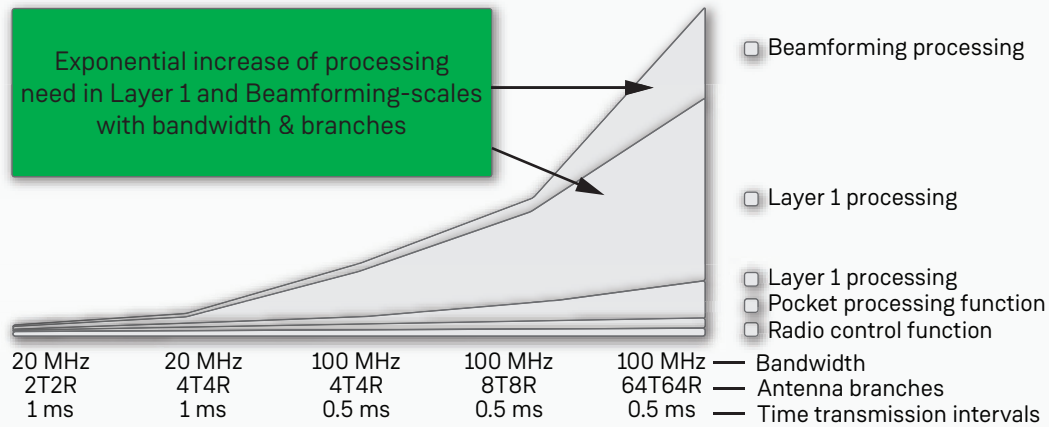


Figure 4 – Use Cases, Splits and Diversity

Figure 5 below graphically represents the challenges 5G NR brings, with respect to the amount of processing performance required to support some of the new antenna configurations. Low-band (20 MHz) MIMO antennas with two transmit and two receive paths (2T2R) are represented on the left of the diagram, and mid-band (100 MHz) antennas with 64 transmit and 64 receive paths (64T64R) are on the right. The evolution from low to mid band allows for higher spectrum channels with potentially spectrum sharing, dual connectivity and carrier aggregation with 4G. These mid-band requirements also come with the need to support much lower 0.5 ms transmission intervals and the need for significant amounts of processing for beam forming and steering.

Therefore, as represented in the figure below, the amount of compute power required, for the L1 processing in particular, starts to scale exponentially with these higher bandwidths. Layer 1 processing on the air interface, and managing beam forming and steering on mid-band spectrum, requires significantly higher processing requirements than low-band deployments.



98246942-05.2022.05.08

Figure 5 – Processing Loads Required for 5G Low- and Mid-band Spectrum
(Source: Ericsson Blog)

In order to meet the processing loads for L1, the industry must consider different heterogeneous solutions that can handle processing needs efficiently (both from the perspective of performance and power). Together with the new network/function splitting, these new solutions might introduce diverse equipment and, therefore, SoC options — one homogeneous solution cannot fit all the RAN needs.

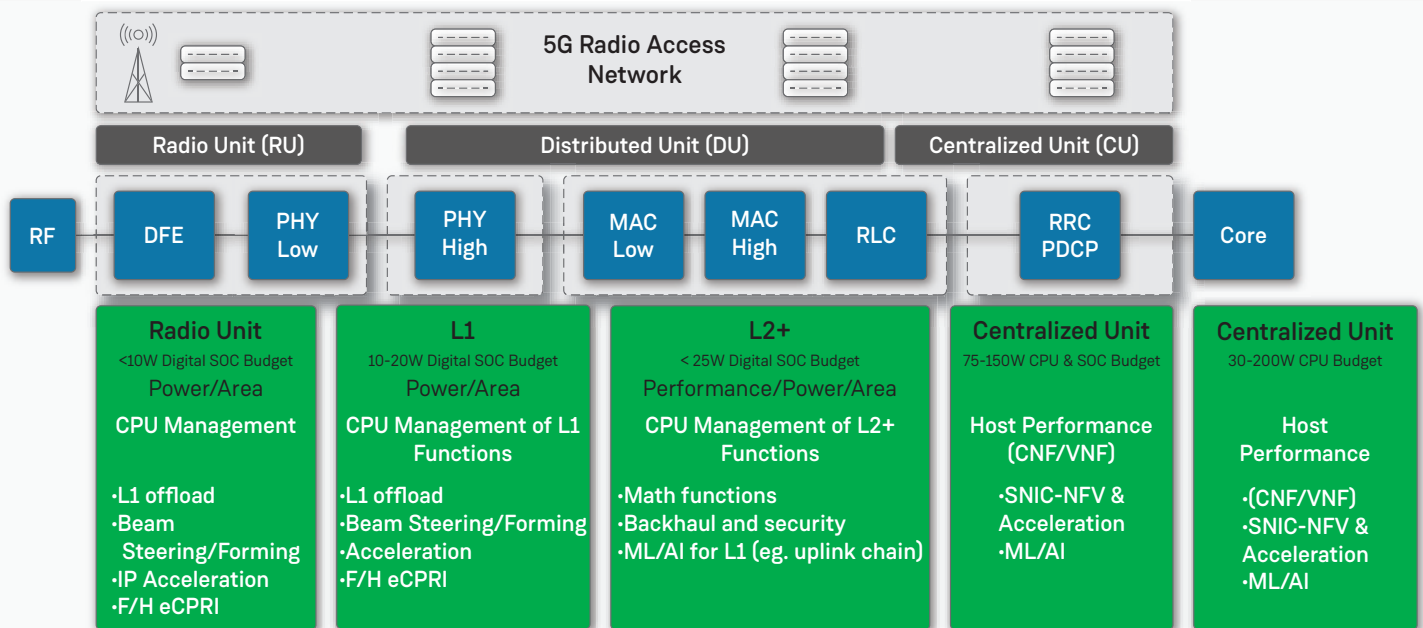
Dis-aggregation of 5G Equipment Drives the Need for Flexibility and Acceleration

New requirements in Rel-17 and Rel-18 drive the need for additional flexibility and acceleration offload from a homogeneous CPU only subsystem. The figure below shows the main elements in a 5G network: RU, DU and CU. For each of these elements, there is a requirement to consider how heterogeneous architectures — a mix of CPUs, DSPs and accelerators (such as, GPUs, FPGAs and eFPGAs) — can be utilized to meet the latency, power, area and cost targets for each of these new designs.

There has been a desire from network operators to use cloud-native, software-based technologies as much as possible for all RAN functions (centralized RAN-based deployments), with the assumption that solutions running on x86- or Arm-based CPU platforms maximizes flexibility. Studies have shown that for low-band deployments (roughly 600-700 MHz serving 50-25 Mbps), baseband and control can be serviced on a CPU platform with minimal acceleration offload. The result is centralized DU and CU functionality, using fiber to connect to the RU with only minimal processing capability in the radio.

Deployments could utilize a COTS server for all processing of a low-band cell with a single CPU core. For these types of deployments, running everything in software as virtualized or containerized workloads, the performance, cost, and power requirements could be viable. In that case, looking at the diagram, L2+ functions in the DU and possibly most of the L1 processing could be co-located in mini-servers together with the core networking functions in the CU.

However, moving towards mid-band deployments in the sub 6 GHz range, roughly the 3.5 to 3.6 GHz range, as seen in one of the previous diagrams, the processing for the radio (including the baseband functions in the L1 block and most of the functions in the L2 block) grows almost exponentially. In this case, there is a 20–40× increase in the downlink and uplink processing load. Running a fully loaded mid-band cell with no acceleration would require in



98246942-06.2022.05.08

Figure 6 – Dis-aggregation of 5G Equipment Drives Need For Flexibility and Acceleration

excess of 16 x86 cores. However, the cost and power consumption for such a system is not commercially viable, driving the importance of continuing to offload certain Layer 1 and Layer 2 functions into specialized hardware — hardware accelerators either co-located in the CU, or distributed closer to the radio interface in remote DU and RU.

The bulleted items listed here, with the exception of CNF/VNF, are ideal candidates to be considered for offloading into accelerator hardware from x86, Arm or R5 CPU subsystems. Some examples are:

- Network processing and classification management on the interfaces between boxes (transport/backhaul/security interface, eCPRI front-haul interface, or where traffic managers, classifiers, etc., are needed)
- L1 processing and beamforming is another area where acceleration is essential for maximizing throughput and optimizing power consumption with DSPs or eFPGA technology, or both.

In addition, a likely default requirement in almost every RAN SoC by 2025 is acceleration for machine learning — a feature that could be applied not only in learning and inference functions for use cases running over 5G, but also for the RAN L1 physical layer enhancements. Studies show that AI/ML can significantly improve L1 PHY performance, with one of the first areas for investigation being where AI/ML enhancement could be applied to beam management, channel estimation and prediction.

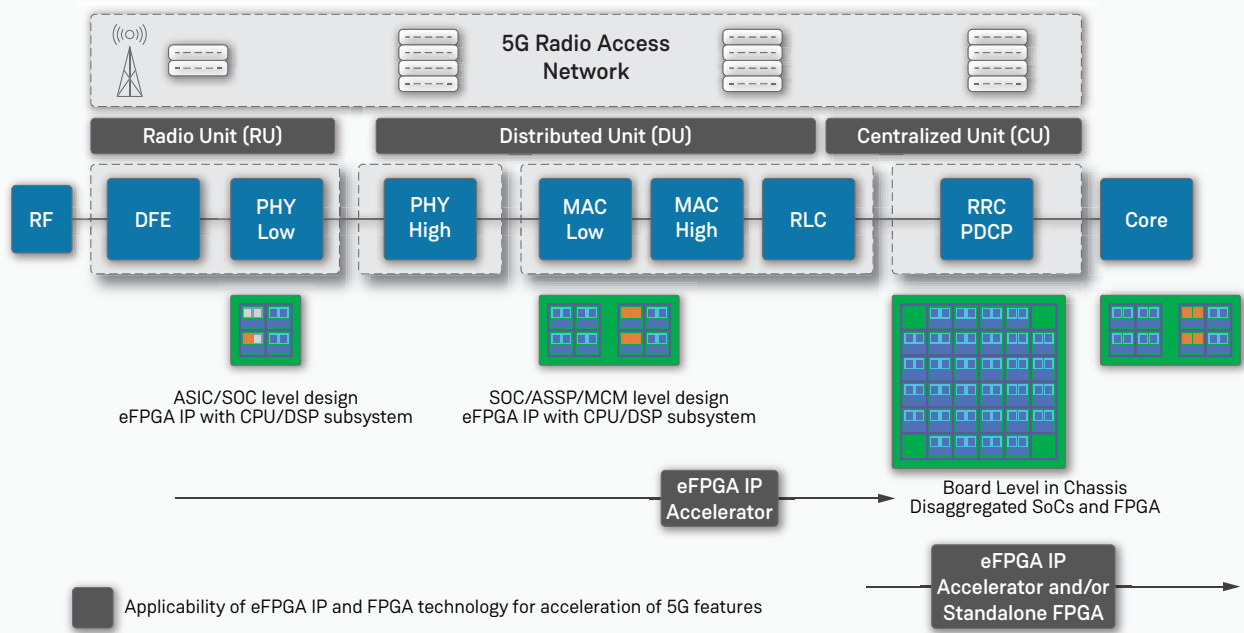
5G Advanced eFPGA and FPGA Acceleration

FPGA and eFPGA technology could be used across the range of 5G designs in the future. As discussed earlier, there will always be a tradeoff between programmability and compute efficiency. While CPUs offer the ultimate programmability, a hard solution based on GPUs, FPGAs and ASICs always offers the benefits of less power but with a lot less flexibility.

Historically, FPGAs have been used in many of the previous generations of cellular designs. In 3G and 4G designs, significant parts of the systems were designed around standalone FPGAs. These FPGAs were used to accelerate certain functions on the air interface, closely coupled with DSPs used for the air interface processing on the baseband unit. FPGAs were also used for transport and security interfaces from CPRI connectivity, chassis interface and backhaul and security interfaces.

The ability to embed FPGA capability within ASICs allows some of the challenges 5G designs present to be addressed. Integration of eFPGA capability in an SoC provides a lower cost route compared to standalone FPGAs since the designer only includes the resources needed while reducing board area, added packaging and I/O. Integration on a SoC with tight coupling to CPU and DSP resources allows higher bandwidth, lower latency and lower power while still offering the ability to increase flexibility by enabling real-time field upgrade to deployed equipment as specifications change.

In the **Figure 7** above, the boxes in red illustrate how Achronix eFPGA and FPGA technology could be used to incorporate flexibility into new designs for the RU, DU and CU either as a standalone device, a monolithic SoC or as a chiplet design packaged in a multi-chip module.



98246942-07.2022.05.08

Figure 7 – 5G Advanced: eFPGA IP and FPGA For Heterogeneous Compute Acceleration

For CU and core RAN applications, either one or multiple FPGAs could be used to provide support for very high data rates and compute densities for a variety of networking and radio specific offload from servers.

Achronix is working with a number of partners in this space who are developing targeted solutions. Companies, such as Napatech and Accolade, are developing FPGA IP targeting SmartNICs. These SmartNICs can be used for multiple different 5G needs including the DU for vRAN-based deployments. The resultant designs include technology for networking, PDCP, security (both air interface and backhaul), OVS and L1 offload. In the future, it is likely that these solutions will also be used for ML inferencing for both multi access edge compute and specifically for radio applications.

The elements shown in red above in the RU and DU show eFPGA capability and how one or multiple tiles of embedded FPGA can be integrated into SoC designs together with CPU, DSP and memory subsystems.

eFPGA Integration on SoC

An eFPGA is a core that is embedded within a custom SoC or ASIC. The IP can be licensed for use similar to that of other IP used in semiconductor designs. Unlike the design process for a standalone FPGA, eFPGA designers can select the exact amount of logic, DSP and memory resources required for their customer's application. An eFPGA can also be used to lower system cost, power and board space by eliminating the standalone FPGAs when moving into high-volume production.

Speedcore™ eFPGA IP architecture incorporates many architectural enhancements that dramatically increase performance, reduce power consumption, and shrink die area. When selecting a Speedcore eFPGA, designers can select the optimal mix of architectural elements including:

- Logic – 6-input look-up-tables (LUTs) plus integrated wide MUX functions and fast adders
- Logic RAM – 2 kb per memory block for LRAM2k, and 4kb per memory block for LRAM4k
- Block RAM – 72 kb per memory block for BRAM72k, and 20kb per memory block for BRAM20k
- DSP64 – 18 × 27 multiplier, 64-bit accumulator and 27-bit pre-adder per block
- Machine learning processors (MLP) – 32 multiplier/accumulators (MACs) per block, supporting integer and floating point formats

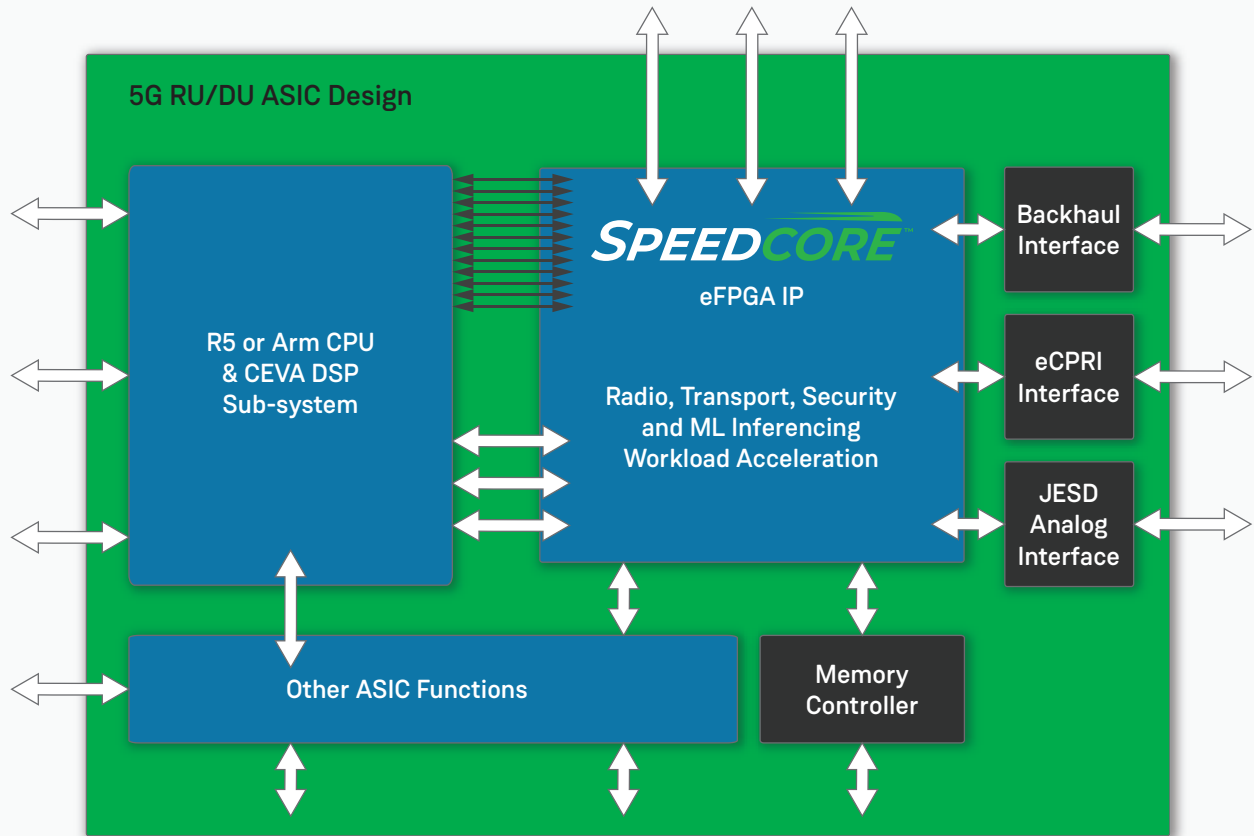
Integrating eFPGA capability in SoC-based designs is an ideal way to provide a flexible, scalable platform to maximize performance in a RAN design while still meeting stringent power targets for these new designs. Integrating eFPGA technology can bring the benefits of a standalone FPGA with several additional benefits:

- These designs offer much less power for the same compute capacity versus a CPU or GPU with the flexibility to add and change capabilities
- The re-configurability of eFPGAs deliver the flexibility to meet evolving standards and deploy updates in the field
- A low-latency, power-efficient and highly flexible eFPGA IP tile can be replicated in multiple SoC designs

Closely coupling FPGA capability with the CPU, DSP and memory subsystem also brings advantages. Standalone FPGAs consume power with integrated high-speed SerDesS/PHYs that connect to other chips. Integrating an eFPGA into an SoC eliminates the need for SerDes interfaces on both sides of the design and only utilize the functionality that you actually need. There are also an inherent savings in silicon area.

The designer can choose to integrate single or multiple eFPGA instances, which can be scaled anywhere from a few thousand LUTs up to a few hundred thousands of LUTs in an SoC. These eFPGA instances can be closely coupled with the CPU subsystem to make efficient use of shared cache and memory subsystems for high-performance, low-latency tasks. For example, Arm offers the CHI-E bus as part of their architecture supporting a coherent mesh interconnect, allowing for some applications that would place a high load on the CPU to be offloaded into the eFPGA block for dedicated processing.

Speedcore eFPGA technology is production proven. Our customers have shipped in excess of 10 million units into these types of applications for a variety of functions including support for eCPRI connectivity, backhaul and security interfaces, radio digital front end (DFE) algorithmic function offload for digital pre-distortion adaption, beamforming offload and baseband repartitioning with Split L1 (I/FFT, RACH, LDPC, etc.).



98246942-08.2022.05.08

Figure 8 – Meeting 5G Advanced Features in ASIC/SoC with eFPGA: RU, DU (and CU) Implementation

An eFPGA as an Accelerator for 5G NR Functions

Achronix is targeting Speedster® and Speedcore technology to address the requirements of 5G-A and 6G. Along with partners, Achronix is working to develop solutions that can be used to target the current and long-term trends influencing 5G development requirements. Some of the benefits Achronix technology can deliver include:

- High-performance architecture for acceleration of various 5G workloads – Achronix offers solutions with high levels of performance for each power/area budget and supports FPGA and eFPGA technologies to accelerate workloads with exceptional power efficiency.
- Diverse solutions and ecosystem – Achronix gives designers the freedom to closely couple custom accelerators and to supplement operation for both eFPGA and FPGA based environments. The Achronix ecosystem includes a broad range of partners to push innovation, with 5G features like eCPRI, radio offload and and chip-to-chip (C2C) interconnects.
- Scalable from cloud to the radio interface – Achronix solutions offer the performance needed to offload servers with FPGA SmartNIC designs adapted for 5G applications and performance that scales with eFPGA for throughput and power needs in the RAN. Moreover, this architecture is scalable across all the points in between.

This white paper has highlighted the following areas that represent the main challenges around the evolution of 5G:

- **Data processing** – To achieve higher spectrum efficiency and meet end-to-end latency requirements, a 5G RAN needs to perform more complex algorithms for data processing. When considering the requirements for these algorithms, it is important to find the right balance between hardware and software tasks, so that the system meets its goals for performance, power, and cost. The eFPGA is an ideal candidate for workloads offloaded from a CPU subsystem.
- **Deployment scenarios** – The specific use cases to be supported by a given RAN have a strong influence on the system as a whole, since each use case (mMTC, eMBB, URLLC) has its own unique characteristics. One size probably will not fit all. Deciding how to divide network functionality across different equipment to support a given set of use cases can impact the RAN design.
- **Radio and spectrum** – 5G uses more of the spectrum, with devices operating in low bands (below 1 GHz), mid bands (between 1 GHz and 2.6 GHz or between 3.5 GHz and 8 GHz), and high bands (between 24 GHz and 40 GHz). Each band has its own set of requirements for edge performance, capacity, speed, and latency. As new spectrum assets become available, these various requirements need to be addressed by the RAN system.
- **Supply chain and ecosystem** – 5G is disrupting the supply chain in several ways. There are initiatives that aim to reduce vendor dependence along with growing availability of proprietary and open software platforms. The level of infrastructure support also varies from region to region. OEMs might need to reassess and revise their ecosystem partnerships.
- **Emerging standards** – Significant investment is being made into evolving the 5G standards to support new use cases and additional features. Specifically Rel-17 and Rel-18 will enable many new use cases. In addition to the 3GPP, there are independent industry organizations, such as the Telecom Infra Project (TIP) and the Open RAN Alliance (O-RAN), that are working on aspects of 5G operation and deployment. There is a growing convergence towards the O-RAN alliance as the key industry body that drives interface specifications.

Summary

The radio access network and the 5G network hierarchy will change. Dis-aggregation of equipment from today's baseband and radio functions into separate boxes will require that functions need to reside potentially in multiple different parts of the network to support different option splits. In the future, MNOs will be required to dynamically partition network functions using slicing techniques. Use of containerized and virtualization functions running on COTS servers will become widespread with virtualization of functionality throughout the network.

However, the success of 5G is dependent on enabling flexible, scalable platforms with power, throughput and latency key in supporting L1 and massive MIMO in antennas especially in the RAN. In the network hierarchy, new functions, such as edge compute, will necessitate machine learning functions being pushed closer to the radio interface. Scalable, heterogeneous SoC architectures with CPU and DSP capability with acceleration offload on to FPGA and eFPGAs (based in ASICs, SoCs, ASSPs) will prevail in order to meet changes in specifications in the near and medium term.

In summary, eFPGA IP is a critical element of meeting these new design challenges as it provides the ability to scale functionality to meet new specifications and as-yet unknown features in 3GPP R'17 and R'18 — 5G Advanced and 6G.

Achronix[®]

Data Acceleration

Achronix Semiconductor Corporation