

Build More Cost-Effective and More Efficient 5G Radios with Intel® Agilex™ FPGAs

With Intel® Xeon®-D CPUs, Intel® Agilex™ FPGAs, Intel® eASIC™ devices, and ASIC technologies, Intel is the only company on the planet that has an end-to-end silicon solution for 5G radio access networks.

Why Intel® Agilex™ FPGAs?

- ~2X better fabric performance per watt compared to competing 7 nm FPGAs*
- 15% to 20% faster on average than competing 7 nm FPGAs for 5G O-RU functions*
- 45% higher fabric performance on average compared to Intel® Stratix® 10 FPGAs*
- Up to 40% lower total power compared to Intel® Stratix® 10 FPGAs*
- TCO reduction via migration to lower cost, lower power Intel® eASIC™ devices

*See configuration information at [intel.com/performanceindex](https://www.intel.com/performanceindex)

In the context of telecommunications, the term '5G' refers to the fifth-generation technology standard for broadband cellular networks. 5G is the planned successor to existing 4G networks, which currently provide connectivity to the vast majority of the world's cellphones. 5G also represents a new "network of networks," bringing the wireless, computing, and cloud worlds together.

There is a lot of terminology in this area, and the functions associated with different portions of networks are changing and evolving, so it is worth establishing a few foundational definitions. For example, a radio access network (RAN) is part of a mobile telecommunications system. Conceptually, the RAN resides between the core network (CN) and wireless connected devices like mobile phones, which may be referred to as user equipment (UE). Over the years, RAN technology has evolved, and 5G RANs (along with some recently deployed 4G RANs) now feature a suite of virtualized networking technologies whose capabilities are significantly more advanced as compared to earlier RAN standards.

An Open RAN (O-RAN) is a nonproprietary implementation of a RAN that allows interoperability between cellular network equipment provided by different vendors. To support the O-RAN concept, the O-RAN Alliance is a world-wide community of mobile network operators, vendors, and research and academic institutions operating in the RAN industry.

A centralized-RAN (C-RAN) is a centralized, coordinated, cloud computing-based architecture for RANs that supports 4G, 5G, and future wireless communication standards.

The 3rd Generation Partnership Project (3GPP) is an umbrella term for a group of standards organizations that develop protocols for mobile telecommunications. 3GPP is a consortium with seven national or regional telecommunication standards organizations as primary members ("organizational partners") and a variety of other organizations as associate members ("market representation partners"). 5G NR (New Radio) is a new radio access technology (RAT) developed by 3GPP for 5G that is designed to be the global standard for the air interface of 5G networks.

Authors

Richard Maiden

Director, Wireless Systems Solution Engineering

Christian Lanzani

Director, Head of Wireless Business Unit

Ankur Vora

Wireless Business Developer

Network Business Division
Intel Programmable Solutions Group

4G to 5G RAN Evolution

In the case of Evolved Universal Terrestrial Radio Access Network (E-UTRAN) RANs—that is, 4G or LTE RANs—one or more radio towers, known as remote radio heads (RRHs) are connected to a baseband unit (BBU), which is itself connected to the core network. In this context, the term 'fronthaul' refers to the fiber-based connection between the RRH and the BBU, while the term 'backhaul' refers to the connection between the BBU and the core network (Figure 1a).

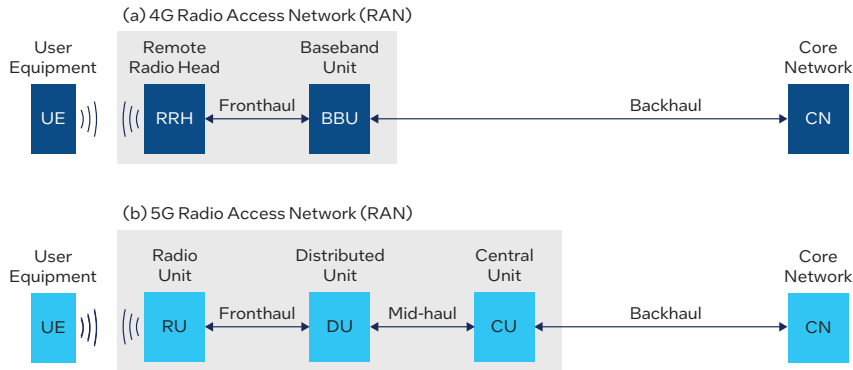


Figure 1. In the 5G RAN architecture, the base station is split into three logical nodes.

Many Next-Generation Radio Access Network (NG-RAN)—that is, 5G RAN—use cases demand a significant increase in fronthaul bandwidth. One way to mitigate this and reduce fronthaul bandwidth is to move functions from the BBU into the radio. In the case of the 5G RAN architecture, the functionality of the 4G BBU is split into three logical nodes (Figure 1b). Some of the functionality moves up into the radio unit (RU), while the remaining functionality is divided between the distributed unit (DU) and the central unit (CU). In this scenario, ‘fronthaul’ refers to the typically fiber-based connection between the RU and the DU, ‘backhaul’ refers to the connection between the CU and the core network, and a new ‘mid-haul’ term is introduced to reflect the connection between the DU and the CU.

It should be noted that Figure 1 is a gross simplification and there are many possible topologies. In the case of a 4G RAN, for example, there may be multiple BBUs and multiple RRHs feeding each BBU. Similarly, in the case of a 5G RAN, there may be multiple DUs and multiple RUs feeding each DU. In addition to reflecting the connection between a DU and a CU, the term mid-haul may also be employed for DU-to-DU connections.

The 5G NR RAN architecture also addresses the fact that different use cases may be better served by different network configurations. The phrase “3GPP Functional Splits” refers to the fact that 3GPP allows for a variety of options for distributing (“splitting”) the functionality of the 5G NR RAN stack between the DU and the RU where Split 7.2 is most frequently used to provide the best option for supporting 5G O-RANs.

End-to-End 5G Silicon Solutions

Intel provides silicon technologies that address every aspect of the 5G O-RAN architecture, including O-RAN RUs (O-RUs), O-RAN DUs (O-DUs), O-RAN CUs (O-CUs), and O-DRUs, which are formed from the combination of an RU and DU. These technologies include Intel Agilex FPGAs, Intel Xeon-D CPUs, Intel Optane™ persistent memory (high-capacity, high-speed, non-volatile), and Ethernet solutions (Intel’s Agilex FPGA-based programmable acceleration card (PAC) network adapters, cards, controllers, and accessories). Intel augments and enhances these hardware technologies with complementary software solutions, including the FlexRAN software stack, Intel’s virtual RAN (vRAN) enablement package (Figure 2).

The Intel FlexRAN software stack is a 4G and 5G L1 reference design that can run on any Intel Xeon CPU, including the latest generation Ice Lake (the code name for the 10th generation Intel Core mobile and 3rd generation Intel Xeon Scalable server processors based on the new Sunny Cove microarchitecture) and Sapphire Rapids (the code name for Intel’s next-generation Intel Xeon server processors based on the Intel 7 nm process). FlexRAN enables the creation of end-to-end O-RAN solutions with Intel Xeon processor, FPGA, Intel eASIC devices, and Ethernet technologies. FlexRAN is complemented by OneAPI, which is an open standard for a unified application programming interface (API) intended to be used across different compute accelerator (coprocessor) architectures, including graphics processing units (GPUs), artificial intelligence (AI) accelerators, and FPGAs.

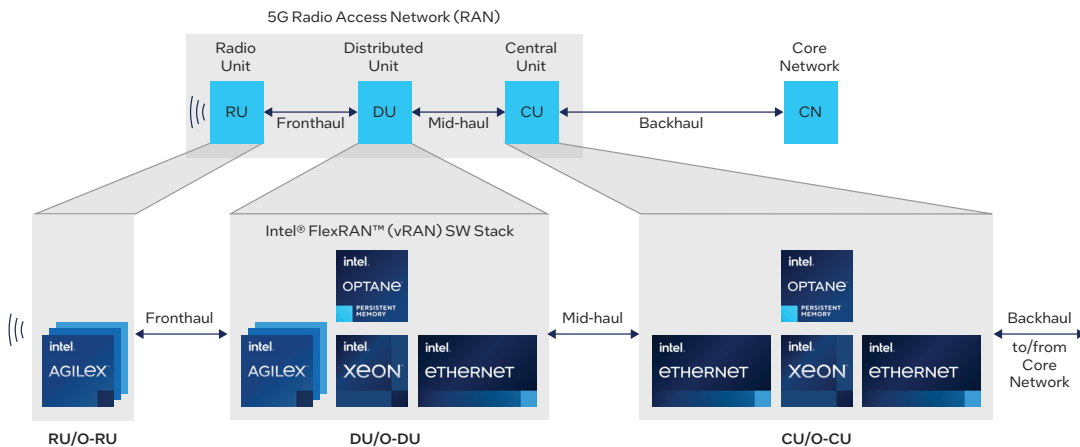


Figure 2. Intel offers end-to-end silicon solutions for 5G RANs.

Intel Agilex FPGAs offer the ability to accelerate the FlexRAN (vRAN) 5G RAN algorithms in a massively parallel fashion, thereby providing extreme performance while consuming relatively low power. Also, the programmable fabric in Intel Agilex FPGAs allows developers to respond to rapidly changing standards and evolving RAN protocols, including remote updates after systems have been deployed into the field. By comparison, in the case of an ASIC implementation, updating the design to accommodate changes in protocols may require a whole new spin of the design, verification, and tapeout process, which is resource-intensive, time-consuming, and costly.

To accommodate various business and development models, Intel also provides market-leading Intel eASIC device and ASIC options. Intel eASIC devices are structured ASICs, which are an intermediate technology between FPGAs and standard-cell ASICs. Alternatively, standard-cell ASICs can be created using Intel Foundry Services, which is a fully vertical, stand-alone foundry business built to help meet the growing global demand for semiconductors.

Intel eASIC device performance falls between FPGA and ASIC implementations. Using an Intel eASIC device, you can keep the same clock rate as the FPGA implementation and reduce power, or you can increase performance while staying inside the same thermal/power budget. Intel eASIC devices also provide faster time to market (TTM) and lower non-recurring engineering (NRE) costs compared to standard-cell ASICs. By comparison, standard-cell ASICs offer the highest performance with the lowest power consumption. Once an Intel Agilex FPGA implementation of a design has been proven, that design can be hardened into a lower cost and lower power Intel eASIC device or into a higher performing even lower power ASIC. In some cases, developers might start with an Intel Agilex FPGA implementation of the design, migrate to an Intel eASIC device, and then migrate once again to a full ASIC (Figure 3).




Figure 3. Intel offers paths from FPGA to eASIC and ASIC to reduce costs, lower power, and increase performance.

As will be demonstrated later in this paper, with respect to FPGA implementations, Intel Agilex FPGAs offer the best solution for 5G RANs on the market. This is augmented by FPGA-to-eASIC-to-ASIC cost reduction, power reduction, and performance enhancement paths that no other company offers. With its Intel Xeon-D CPUs, Intel Agilex FPGAs, Intel eASIC devices, and ASIC technologies, Intel is the only company on the planet that has an end-to-end silicon solution for 5G RANs.

Intel Agilex FPGAs for 5G RANs

Intel Agilex FPGAs can provide tremendous performance benefits from the edge to the cloud; that is, from edge and endpoint devices to hyperscale data centers. In the case of 5G RANs, Intel Agilex FPGAs can be deployed as raw devices laid down on custom printed circuit boards (PCBs) in the RU, to Intel FPGA-based programmable acceleration cards (PACs)—such as the Intel FPGA PAC N3000 and N6000—which can be used to accelerate network traffic to support low-latency, high-bandwidth 5G applications. These PAC cards, which support PCI Express (PCIe) Gen4 and Gen5, also allow developers to create custom-tailored solutions for vRAN and core network workloads, and to achieve faster time to market with the support of industry-standard orchestration and open-source tools.

Intel Agilex FPGAs leverage Intel's heterogeneous 3D system-in-package (SiP) technology. This features Intel's Embedded Multi-die Interconnect Bridge (EMIB)¹, which is an elegant and cost-effective approach to the in-package high density interconnect of heterogeneous chips. In this case, EMIBs are used to connect the core FPGA die with other dies (a.k.a. tiles or chiplets) such as transceivers, thereby allowing each function to be created using the most appropriate process technology.

Intel Agilex FPGAs are the first FPGA fabric built on 10 nm SuperFin technology. This is Intel's third generation FinFET technology. These devices also use the second generation of the Intel® Hyperflex™ FPGA Architecture, which delivers up to 45% higher performance or up to 40% lower power for applications in the data center, networking, and edge compute as compared to its predecessor Intel Stratix 10 FPGA. When compared to our competition's 7 nm FPGA portfolio, Intel Agilex delivers ~2X better fabric performance per watt. Intel Agilex SoC FPGAs also integrate 64-bit quad-core Arm Cortex-A53 processors—also known as the hard processor system (HPS)—to provide high system integration.

Intel® Quartus® Prime Software is Intel's FPGA design software suite. The Intel Quartus Prime Software enables the analysis and synthesis of hardware description language (HDL) designs, enabling developers to compile their designs, perform timing analysis, examine register transfer language (RTL) diagrams, simulate a design's reaction to different stimuli, and configure the target device. The Intel Quartus Prime Software supports Verilog and VHDL, the visual editing of logic circuits, and vector waveform simulation.

5G RANs involve a lot of digital signal processing (DSP). DSP Builder for Intel® FPGAs is a DSP design tool that enables HDL generation of DSP algorithms directly from the MathWorks Simulink environment. This tool allows developers to describe the design at a high-level of abstraction, and it then generates highly efficient structures in the form of high-quality, automatically pipelined, synthesizable VHDL/Verilog code from MATLAB functions and Simulink models. This register transfer level (RTL), which facilitates rapid design space exploration and speeds time to market, can be combined with other RTL and intellectual property (IP) blocks before being fed to the Intel Quartus Prime Software.

It is important to note that the DSP Builder for Intel FPGAs can generate optimized RTL for both FPGA and Intel eASIC devices. The functionality of the RTL for the two targets will be bit-identical for the same DSPB design.

Intel Agilex FPGAs are supported by a wide range of hardware and software IP blocks and functions. In the case of 5G RANs, some key IP offerings are as follows:

- **Ethernet:** Ethernet is a family of wired computer networking technologies commonly used in local area networks (LANs), metropolitan area networks (MANs), and wide area networks (WANs). As was previously noted, Intel Agilex FPGAs leverage heterogeneous 3D SiP technology. In addition to the main FPGA silicon die, these SiPs include a number of smaller dies, which are also known as chipllets or tiles. The Intel Agilex FPGA F-Tile, which is an evolution of the E-Tile, offers hardened Ethernet IP that incorporates a fracturable, configurable, hardened Ethernet protocol stack for supporting rates from 10G to 400G, compatible with IEEE 802.3 specification, and other related Ethernet Consortium specifications. The IP core is available in multiple variants providing different combinations of Ethernet channels and features. These include optional Reed-Solomon Forward Error Correction (RSFEC) and optional IEEE 1588v2 Precision Time Protocol (PTP). The user can choose a media access control (MAC) and a physical coding sublayer (PCS) variation, a PCS-only variation, a Flexible Ethernet (FlexE) variation, or an Optical Transport Network (OTN) variation.
- **JESD204C:** JESD204C is a standard of the Joint Electron Devices Engineering Council (JEDEC). It is a high-speed interface designed to interconnect fast analog-to-digital converters (ADCs) and digital-to-analog converters (DACs) to high-speed processors, FPGAs, and ASICs. The JESD204C Intel® FPGA IP addresses multidevice synchronization using Subclass 1 to achieve deterministic latency. It also supports TX-only, RX-only, and duplex (TX and RX) modes.
- **eCPRI:** The Common Public Radio Interface (CPRI) is a legacy interface that is used to implement the fronthaul connection between RRHs and BBUs in 4G RANs. Enhanced or evolved CPRI (eCPRI) provides a way of splitting up baseband functions to reduce traffic strain on the system and implement more flexible and more efficient communications between O-RUs and O-DUs in 5G RANs. The Intel eCPRI IP implements the eCPRI specification version 2.0. It can support 10G and 25G Ethernet ports while also supporting one-way delay measurement similar to IEEE standard 1588 PTP hardware-based timestamping.
- **O-RAN Fronthaul Interface:** The O-RAN Alliance Workgroup 4 (O-RAN WG4) fronthaul interface defines the interface between a O-RAN radio unit (O-RU) and a lower-layer split O-RAN distributed unit (O-DU) in 4G and 5G implementations with a lower layer functional split-7.2 based architecture. The current Intel O-RAN Fronthaul Interface IP supports CAT-A and CAT-B type radios.

Intel Agilex FPGAs for 5G O-RUs

The Groupe Speciale Mobile Association (GSMA) is an association representing the interests of mobile operators and the broader mobile industry worldwide. According to the GSMA², the total cost of ownership (TCO) of 5G RAN infrastructure can increase by 65% compared to current 4G RAN costs. Energy costs can increase by up to 140%. Minimizing these costs is critical to the operators' businesses.

A key aspect to cost is the power consumption of any FPGAs used in the system. Almost all radios are convection cooled (no fans), so the thermal efficiency of the design is critical. Lower FPGA power requires less cooling, which means less metal, reduced size, reduced weight, reduced difficulty of installation, and reduced wind loading on the tower. As was previously noted, Intel Agilex FPGAs are based on the second generation of the Intel Hyperflex FPGA Architecture, which allows them to consume up to 40% lower power as compared to their Intel Stratix 10 FPGA predecessor.

Today's wireless infrastructure is built around heterogeneous networks composed of many different sized radios, ranging from femtocells to picocells to microcells to macrocells. Each of these cell types supports different features and functions, including the following:

- The number of bands, including sub-6 GHz (actually 600 MHz to 7.125 GHz) and millimeter wave (a.k.a. mmWave) bands (24.25 GHz to 52.6 GHz).
- RAT technology (GSM, CDMA, 4G LTE/LTE-A, 5G, NB-IoT).
- RF output power level, which typically ranges from 125 milliwatts up to 80 watts per antenna.
- The number of antenna elements, which is typically up to eight in a macrocell, 16-to-64 in a massive multiple-input and multiple-output (MIMO) array, and hundreds in a mmWave deployment depending on the desired equivalent isotropic radiated power (EIRP) coverage.
- The number of component carriers required to cover 4G, 5G, NB-IoT, and—possibly—legacy standards like 2G and 3G.
- The number of users per cell and the types of services these users are employing.
- The configuration (integrated antenna, traditional active antenna arrays, remote radio heads, etc.)

In order to be able to provide solutions for all of these scenarios, network operators need to be able to draw on a broad range of radio hardware. Flexible platform solutions are desirable because they allow operators to quickly adapt to the changing standards and evolving performance requirement characteristics that are the hallmarks of modern wireless communication systems. Also, scalable platform solutions are desirable because they minimize design effort, resources, and costs; they reduce test time and inventory; and they maximize design reuse. Intel Agilex FPGAs provide flexibility and support scalability.

Key IP components of an O-RU Digital Front End (DFE) include frontal interface processing (CPRI, eCPRI, O-RAN), low layer 1 (LL1), digital up conversion (DUC), digital down conversion (DDC), crest factor reduction (CFR), and digital pre-distortion (DPD).

Earlier in this paper, it was noted the 3GPP Split 7.2 is considered to provide the best option for supporting 5G O-RANs. The O-RAN WG4 fronthaul features a 7.2x split interface with low PHY in the O-RU and high PHY in the O-DU. The key O-RAN low PHY modules for an Intel Agilex FPGA-based split 7.2 O-RU implementation are shown in Figure 4.

These optimized, production-ready functions include the key processing IP elements and the software framework bundled together as an integrated radio solution. This solution is a resource-efficient implementation of the O-RAN standard optimized for the Intel Agilex FPGA. It allows easy reconfiguration of the functionality for various applications while still maintaining an ultra-small resource footprint.

The IP elements shown here include finite impulse response (FIR), fast Fourier transform (FFT), inverse FFT (iFFT), cyclic prefix (CP) add/remove (+/-), radio timing, physical random-access channel (PRACH), layer mapping, IQ compression/decompression, user plane framer/deframer, control plane multiplexer/demultiplexer, and TDD switching. The software environment runs on the Intel Agilex FPGA's hard processor system (HPS).

This 5G O-RAN solution includes integration with JESD204B/C, IEEE1588, and Ethernet components, along with other modules to form a complete radio sub-system. The solution is also scalable to support different antenna configurations, bandwidths, RATs, and carrier configurations. This O-RU implementation can be connected to an Intel FlexRAN-based O-DU server by means of an O-RAN compliant fronthaul interface using eCPRI links in accord with 3GPP Split 7.2X.

Intel Agilex FPGAs vs. Xilinx Versal FPGAs for 5G O-RUs

For the purposes of these evaluations, we picked a pair of devices that were as close as possible in terms of resources. We used an Intel Agilex AGF 014 Series FPGA and a Xilinx Versal VM1802 Prime Series FPGA. In both cases we used the slowest speed grade and an industrial temperature grade (Table 1).

The key modules required to implement a 5G O-RU design are shown in Figure 5.

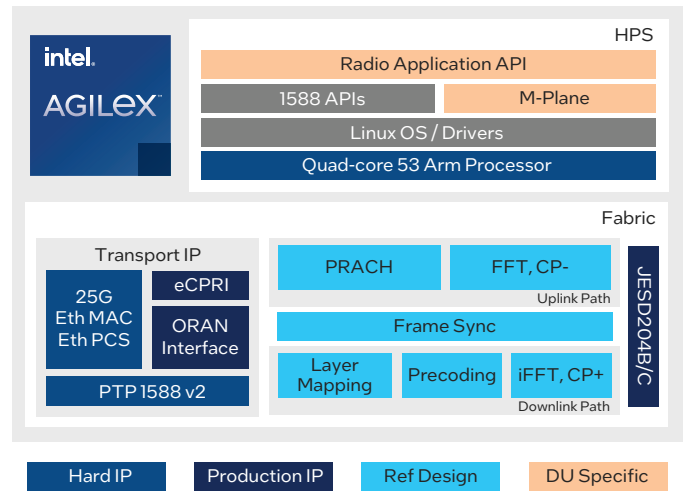


Figure 4. Key O-RAN low PHY modules for an Intel Agilex FPGA-based split 7.2 O-RU implementation.

Resource	Intel® Agilex™ FPGA	Xilinx Versal FPGA
Series	AGF 014 Series	VM1802 Prime Series
Part Number	AGFB014R24A3I3V	XCVM1802-VFVC1760-1LHP-i-S
Speed Grade	Slowest	Slowest
Temperature Grade	Industrial	Industrial
Logic Capacity (ALMs / Slices)	487,200 ALMs	112,480 Slices
DSP (Blocks / Slices)	4,510 DSP Blocks	1,968 DSP Slices
RAM Blocks (M20ks / Block RAM Tiles)	7110 M20ks	967 Block RAM Tiles

Table 1. High-level comparison of devices used in evaluations

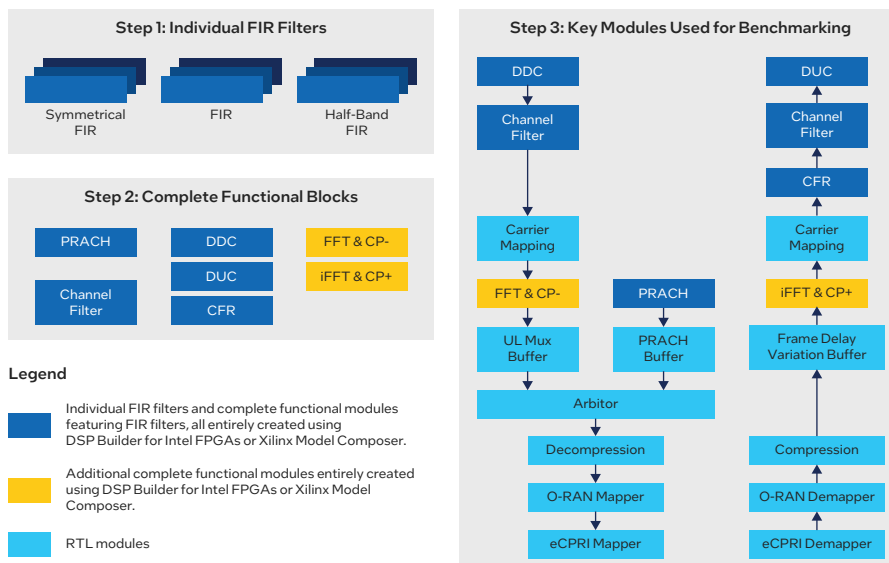


Figure 5. The 5G O-RU design used as a basis for the Intel Agilex FPGA vs Xilinx Versal evaluations.

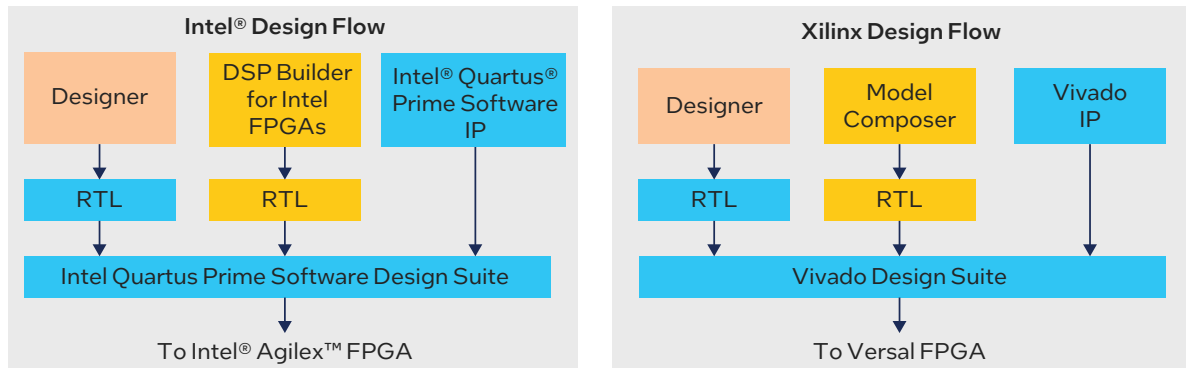


Figure 6. The design flows used to perform the Intel Agilex vs Xilinx Versal evaluations.

In the case of tests performed using the Intel Agilex FPGA, we used our DSP Builder for Intel FPGAs tool, Intel Quartus Prime Software Design Suite, and Intel Quartus Prime Software IP. By comparison, with respect to the corresponding tests performed using the Xilinx Versal FPGA, we used their Model Composer tool, Vivado Design Suite, and Vivado IP (Figure 6).

In the case of the Intel design flow, the tools used were Intel Quartus Prime Software version 21.3 and DSPBA version 21.3. In the case of the Xilinx design flow, the tools used were Vivado version 2021.1 and Model Composer version 2021.1. Both flows used MATLAB version R2020b 64 bit. The server configuration used in both cases was as follows: Server = PowerEdge R630, Processor = Intel Xeon processor E5-2699 v4 product family, operating system = Cent OS Linux 7 (Core), Memory = 256 GB. More details on configurations, optimizations, definitions, tools, and hardware can be found on Intel’s Performance Index web pages³.

Step 1: For our first suite of tests, we focused on FIR filters, which are composed of chains of delays, multipliers, and adders. These are fundamental building blocks found in larger functions. IP modules like DUCs, DDCs, and CFRs all contain some form of FIR filter or multiple FIR filters chained together.

We created almost 60 designs featuring different flavors of FIR filters. Some have more channels, others have more taps, some are symmetrical, while others are half-band symmetrical. Essentially, we performed a “sweep” across this entire design space.

Using these small designs allowed us to perform a true “apples-to-apples” comparison between the two FPGAs. The target frequency for these designs was set to 614.4 MHz. The results were then ranked on the performance of the Intel Agilex FPGA (Figure 7). The X-axis shows the different Design Configuration (DCG) and Y-axis shows FMAX in MHz. The respecting DCG configurations are as shown in Table 2. The slowest design, shown on the left, runs at around 620 MHz, while the fastest, shown on the right, runs at almost 800 MHz. Also shown are the relative performance values for the same designs running on the Xilinx Versal device. Of particular interest are the two tests that failed with Versal when Xilinx Model Composer attempted to use block RAM. No such issue was seen with the DSP Builder for Intel FPGAs.

	Filter Type	Coefficients	Number of channels	Input Sample Rate	Number of Taps	CFG_01	CFG_02	CFG_03	CFG_04	CFG_05
DES1 DES7	Channel Filter Halfband Filter	Programmable	See CFG list	15.36 MSps	87	1	2	4	8	16
DES2 DES8	Channel Filter Halfband Filter	Fixed								
DES3 DES9	Channel Filter Halfband Filter	Programmable	4	See CFG list	107	7.68 MSps	15.76 MSps	30.72 MSps	61.44 MSps	122.88 MSps
DES4 DES10	Channel Filter Halfband Filter	Fixed								
DES5 DES11	Channel Filter Halfband Filter	Programmable	8	30.72 MSps	See CFG list	35	53	87	107	125
DES6 DES12	Channel Filter Halfband Filter	Fixed								

Table 2. Design configuration (DCG)




Figure 7. FIR benchmark results for Intel Agilex and Xilinx Versal FPGAs.

FPGAs are binned (sorted) into different speed grades, where each speed grade roughly equates to around 15% to 20% increase in speed*, and each increase in speed grade leads to a higher price point (*Performance varies by use, configuration, and other factors. Learn more on Intel's Performance Index web pages³). Also, in these times of supply chain disruptions and shortages, faster speed grades are typically less available (slower speed grades are usually slower moving inventory).

From Figure 7 we see that, on average, the Intel Agilex FPGA closed timing at a frequency 15% to 20% faster than the Xilinx Versal device (it was also an average of 5% smaller in terms of logic footprint (*Performance varies by use, configuration, and other factors. Learn more on Intel's Performance Index web pages³)). This difference is equivalent to an increment in speed grade so to match these results with a Xilinx FPGA, customers would have to purchase a higher speed grade device, resulting in a higher cost.

Step 2: For our second suite of tests, we moved up the design hierarchy to consider some core functional elements that are more significant in terms of complexity and size: FFT & CP-, iFFT & CP+, DDC, DUC & CFR, and PRACH (Figure 8).

For the purposes of these evaluations, it's important to note that there are some fundamental clock frequencies that designers typically use to satisfy the latency requirements. Two of these frequencies are 491.52 MHz and 614.40 MHz (see sidebar).

Figure 8 shows that the Intel Agilex FPGA meets the 614.40 MHz timing with the FFT & CP- and the IFFT & CP+, and it meets timing at 491.52 MHz for all of the other designs. It is important to note that these results were achieved with a "Push Button" use of the tools (that is, with no optimization steps).

By comparison, the Xilinx Versal FPGA failed to meet the 614.40 MHz timing for all of the functions, and it even failed to meet the 491.52 MHz timing for the "DUC & CFR" module (once again, details on configurations, optimizations, definitions, tools, and hardware can be found on Intel's Performance Index web pages³).

To be fair, it should be noted that things grow more complex as we move up the design hierarchy. As a result, a developer might try one approach with one FPGA architecture and a different approach with an alternative architecture. In the case of this module, we did our best to make the Xilinx Versal FPGA look good by employing various optimization strategies (Table 3), but it still failed to meet its timing goals.

Magic Frequency Numbers

Many newcomers to the industry have questions regarding the origin of "magic frequency numbers" like 614.40 MHz and 491.52 MHz (dubbed as being "magic" because they seem to appear from nowhere). In fact, the W-CDMA spread-spectrum modulation technique used in 3G RANs employs 5 MHz carriers. Also, its pulse-shaping filter uses a root-raised cosine FIR filter, which creates a 'chipping' rate of 3.84 MHz. Based on this, the CPRI interface adopted a basic frame rate of 3.84 MHz. LTE introduced wider carrier frequencies and abandoned W-CDMA for OFDM. The industry wanted to keep CPRI compatible between 3G and 4G/LTE, so 3GPP tweaked the cyclic prefix lengths so that 20 MHz LTE uses a 30.72 MSps sample rate, which is $2^3 \times 3.84$ MSps = 8×3.84 MSps. 5G is also OFDM and the trend continued. A 5G 100 MHz carrier (with 15 kHz sub-carrier spacing) uses a 122.88 MSps sampling rate which is $2^5 \times 3.84$ MHz. Thus, 491.52 MHz and 614.4 MHz are "magical" frequencies because they are 4x and 5x 122.88 MHz, respectively.

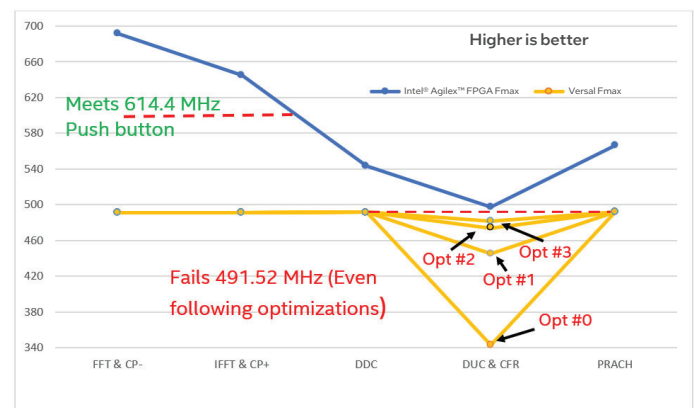


Figure 8. IP module benchmark results.

Optimization Level	Achieved FMAX (MHz)	Versal DUC/CFR Optimization Efforts
Opt #0 (No optimization)	343	<ol style="list-style-type: none"> Initial design Vivado default Implementation strategy
Opt #1	445	<ol style="list-style-type: none"> Pipeline logic added in the design Channel filters and half-band filter structure implemented with four parallel paths each handling one channel. Performance_ExploreWithRemap setting enabled in Vivado (Implementation Strategy)
Opt #2	474	<ol style="list-style-type: none"> CORDIC: # of add-sub-iterations set to 25 (Advanced Configurations) Performance_ExploreWithRemap setting enabled in Vivado
Opt #3	482	<ol style="list-style-type: none"> Channel Filter Memory settings set to implement as Distributed Memory Performance_SpreadSLLS setting enabled in Vivado (Implementation Strategy)

Table 3. Optimization strategies attempted to make the Xilinx Versal FPGA meet timing (it didn't).

Step 3: This step involves the 5G O-RU design shown in Figure 5. In the case of the Intel Agilex FPGA, both the FFT & CP- and iFFT & CP+ maintained the target frequency of 614.40 MHz. Given that this clock speed is not achievable for these modules in the Xilinx Versal FPGA, the target was dropped to 491.52 MHz. In a real-world implementation, this would force a complete redesign of these modules and a corresponding 25% increase in the module size (this re-design effort was not undertaken for the purposes of these benchmarks).

The remainder of the project continued to target 491.52 MHz. A summary of FPGA resources used is shown in Table 4.

In the case of the Intel Agilex FPGA, frequency/timing targets were easily met (maximum placement effort was used). By comparison, in the case of the Xilinx Versal FPGA, even with all the optimizations described in Table 3, the frequency for the full design failed to hit the 491.52 MHz target. In fact, it managed only 372.2 MHz. This is no surprise since the underlying DUC/CFR module could not hit 491.52 MHz. So, not only was the 614.4 MHz target abandoned for the FFT & CP- and iFFT & CP+ modules, the Versal implementation of the complete design still could not achieve 491.52 MHz. As a last-ditch effort, a mid-speed grade device was tried, and this finally managed to meet the 491.52 MHz target (the actual value achieved was 499.62 MHz).

	Utilization		Available		% Utilization	
	Intel Agilex FPGA	Versal	Intel Agilex FPGA	Versal	Intel Agilex FPGA	Versal
Multiplier count	1153 (601 DSP blocks)	1156 DSP Slices	9020 (4510 DSP Blocks)	1968 DSP slices	13% Multipliers	59% DSP Engines
ALM/Slices	87k ALMs	21k CLB Slices	487K ALM	112k CLB Slices	18% ALMs	19% CLB Slices
RAM Blocks	569 M20Ks	357 RAMB36E5 & 67 RAMB18E5	7110 M20K	967 BRAM36E5	8% M20Ks	40% BlockRAM Tiles

Table 4. Summary of FPGA resources used to implement the 5G O-RU design.

Conclusion

Although 5G RAN deployments commenced in 2019 and 2020, and started to gain traction in 2021, we are still in the early days of RAN deployment and the very early days of 5G adoption by end users. Having said this, things are starting to ramp up quickly, and it's anticipated that there will be 3 billion 5G subscribers by 2025⁴.

The worldwide deployment of 5G will require hundreds of thousands of radio units, so carriers are extremely cost conscious. As has been demonstrated in this white paper, with regard to the functions required to implement a 5G O-RU, Intel Agilex FPGAs are, on average, 15% to 20% faster than their Xilinx Versal counterparts while also consuming an average of 5% less logic*. (*Performance varies by use, configuration and other factors. Learn more on Intel's Performance Index web pages³)

The bottom line is that, if we assume that a carrier has a certain 5G O-RU bandwidth requirement that can just be achieved using a single Intel Agilex FPGA, then to match these results using members of the Xilinx Versal family, the carrier will either have to use two Xilinx Versal FPGAs or a higher-speed grade device with both options costing a significant amount more.

References

- ¹ <https://www.intel.com/content/www/us/en/silicon-innovations/6-pillars/emib.html>
- ² <https://www.gsma.com/futurenetworks/wiki/5g-era-mobile-network-cost-evolution/>
- ³ <https://edc.intel.com/content/www/us/en/products/performance/benchmarks/overview/>
- ⁴ <https://www.statista.com/statistics/760275/5g-mobile-subscriptions-worldwide/>

Additional Resources

- Intel Agilex SoC FPGAs
www.intel.com/content/www/us/en/products/details/fpga/agilex.html
- Intel eASIC devices
www.intel.com/content/www/us/en/products/details/easic.html
- Intel Xeon processor
www.intel.com/content/www/us/en/products/details/processors/xeon/scalable.html
- Intel Quartus Prime Software
www.intel.com/content/www/us/en/software/programmable/quartus-prime/overview.html
- DSP Builder for Intel FPGAs
www.intel.com/content/www/us/en/software/programmable/quartus-prime/dsp-builder.html



Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel does not control or audit third party data. You should consult other sources for accuracy.

Intel technologies may require enabled hardware, software or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.