Solution Brief

Intel Agilex 7 FPGAs

Implementing Advanced Networking Solutions with F-Tiles in Intel Agilex® 7 FPGAs

F-Tile transceivers in Intel Agilex® 7 FPGAs and SoC FPGAs power extreme data bandwidth applications, including radio access networks key protocols (JESD204x, Ethernet, CPRI, eCPRI , PCIe), passive optical networks (PON), and more.

Authors

Graham Baker
Agilex Product Line Manager
Intel Corporation

Christian Lanzani
Director, Head of Wireless Business Unit
Intel Corporation

Ankur Vora
Wireless Business Developer
Intel Corporation

Bret Gustafson
Wireless Business Developer
Intel Corporation

Table of Contents

F-Tiles in Intel Agilex 7 FPGAs

Intel Agilex 7 FPGAs and SoC FPGAs leverage the full breadth of Intel innovation and manufacturing capability (Figure 1). The Intel Agilex 7 FPGA family consists of three variants: the F-Series FPGA, which is targeted at a wide range of applications, the I-Series FPGA, which is targeted at bandwidth-intensive applications, and the M-Series FPGA, which is targeted at compute-intensive and memory intensive applications.

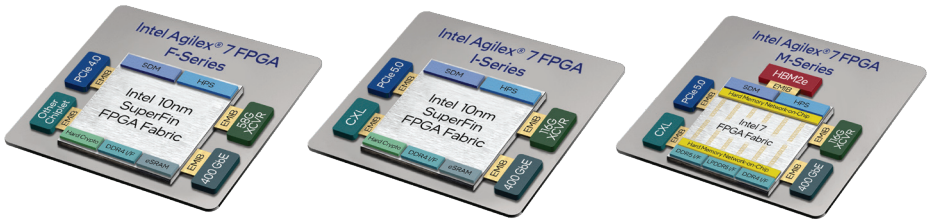

Figure 1. Intel Agilex 7 SoC FPGA Family Members

A chiplet or tile is a small integrated circuit die containing a well-defined subset of hardened functionality. In addition to the main FPGA die, Intel Agilex 7 devices can contain between one and six hardened transceiver (XCVR) tiles. These tiles and the HMB2E stacks in the case of M-Series devices—are connected to the main FPGA die using Intel embedded multi-die interconnect bridge (EMIB) technology. EMIB is an elegant and cost-effective approach to the in-package high density interconnect of heterogeneous chips. The result is that all of these chiplets function as if they are a single large die. Key features for each of these families are shown in Table 1.

	F-Series	I-Series	M-Series
Logic Capacity	573K-2.7M LE	up to 4.0M LE	up to 3.9M LE
On-chip Memory	5-35 MB	up to 53 MB	up to 46 MB
DSP blocks	up to 8k	up to 12k	up to 12k
XCVR speeds	up to 58G	up to 116G	up to 116G
PCIe	PCIe 4.0	PCIe 5.0	PCIe 5.0
Off-chip Memory	DDR4	DDR4	LPDDR5 / DDR5/ DDR4
Hard Crypto / Ethernet	200G / 400GbE	200G / 400GbE	400GbE
Tiles	E & P or F	F & R or F	F & R or F
Security	Triple-modular redundant hard processor, encryption, boot order, tamper detection		
Arm SoC	Quad-core Arm Cortex-A53 up to 1.41 GHz, NEON co-processor, DMA, cache, etc		
Coherency Option		Compute Express Link (CXL)	Compute Express Link (CXL)
High-Bandwidth Memory Option			HBM2E

Table 1. Key features of the Intel Agilex 7 FPGA families.

To complement their extreme processing capability, Intel Agilex 7 devices support multiple types of XCVR tiles, including E-Tiles, F-Tiles, P-Tiles, and R-Tiles. Different members of the Intel Agilex FPGA family provide different numbers and combinations of these tiles. Facilitated by Intel's EMIB and chiplet technology system, integrators have the possibility to devise a broad variety of XCVR capabilities utilizing different tile combinations. This allows customers to employ hardened transceivers while keeping the full capacity and flexibility of the FPGA soft logic. F-Tiles can support up to 400 Gbps Ethernet, while R-Tiles can support PCIe 5.0 and CXL interfaces.

The Intel Agilex 7 FPGA and SoC FPGA F-Tile feature set is as follows:

Feature	Description
Number of available PMAs	• Up to 20 PMAs • FHT (High speed Transceivers): up to four per tile • FGT (General purpose Transceivers): up to 16 per tile
Data rate range	FHT: • 24-29 Gbps NRZ • 48-58 Gbps NRZ and PAM4 • 96-116 Gbps PAM4 FGT: • 1-32 Gbps NRZ • 20-58.125 PAM4 Refer to PMA Data Rates
PCIe hard IP modes Ethernet hard IP modes with number of supported PMAs for each, where 10GbE-1 is 10GbE mode supporting one PMA	Up to one 4.0 x16, two 4.0 x8, or four 4.0 x4. • From 10GbE-1, 25GbE-1, up to 400GbE-8, and 400GbE-4, with these optional features: • Auto-negotiation • Link training • IEEE 1588 precision time protocol (PTP) Includes Ethernet PCS and MAC for all data rates. Not all features are supported for all data rates. Refer to F-Tile Ethernet Intel FPGA Hard IP User Guide
Forward error correction (FEC) and Reed-Solomon FEC (RS-FEC) modes	• IEEE 802.3 BASE-R Firecode (CL 74) • Ethernet Technology Consortium (ETC) RS(272, 258) • IEEE 802.3 RS(528, 514) (CL91) • IEEE 802.3 RS(544, 514) (CL 134) Refer to F-Tile Supported FEC Modes and Compliance Specifications

The Role of Transceivers in O-RAN Architectures

As connectivity and connected devices continue to experience exponential growth, this drives the demand for a flexible radio access network (RAN) architecture. The Open RAN Alliance[1] is a worldwide community of mobile network operators, vendors, researchers, and academic institutions operating and working in the RAN industry. This organization is dedicated to evolving RANs towards being smart and open. The goal is to migrate legacy proprietary RANs towards an ecosystem of innovative, multi-vendor, interoperable autonomy, thereby reducing costs while increasing performance and agility.

The term Open RAN (O-RAN) refers to a non-proprietary implementation of a RAN that allows interoperability between hardware and software provided by different vendors. In the case of a 5G RAN architecture, the base station is split into three logical nodes: the radio unit (RU), the distributed unit (DU), and the central unit (CU). The equivalent O-RAN compliant versions are referred to as the O-RAN RU (O-RU), O-RAN DU (O-DU), and O-RAN CU (O-CU).

As illustrated in Figure 2, Intel provides end-to-end (E2E) silicon technologies that address every aspect of the 5G O-RAN architecture, including O-RUs, O-DUs, and O-CUs. These technologies include Intel Agilex 7 FPGAs, Intel® Xeon® CPUs, and a wide variety of Ethernet solutions (network adapters, cards, controllers, and accessories). Intel augments and enhances these hardware technologies with software counterparts, including a library of intellectual property (IP), LowPHY and FHGW enablement packages, mMIMO white box, FlexRAN™ software stack, and the Virtual RAN (vRAN) Enablement Package.

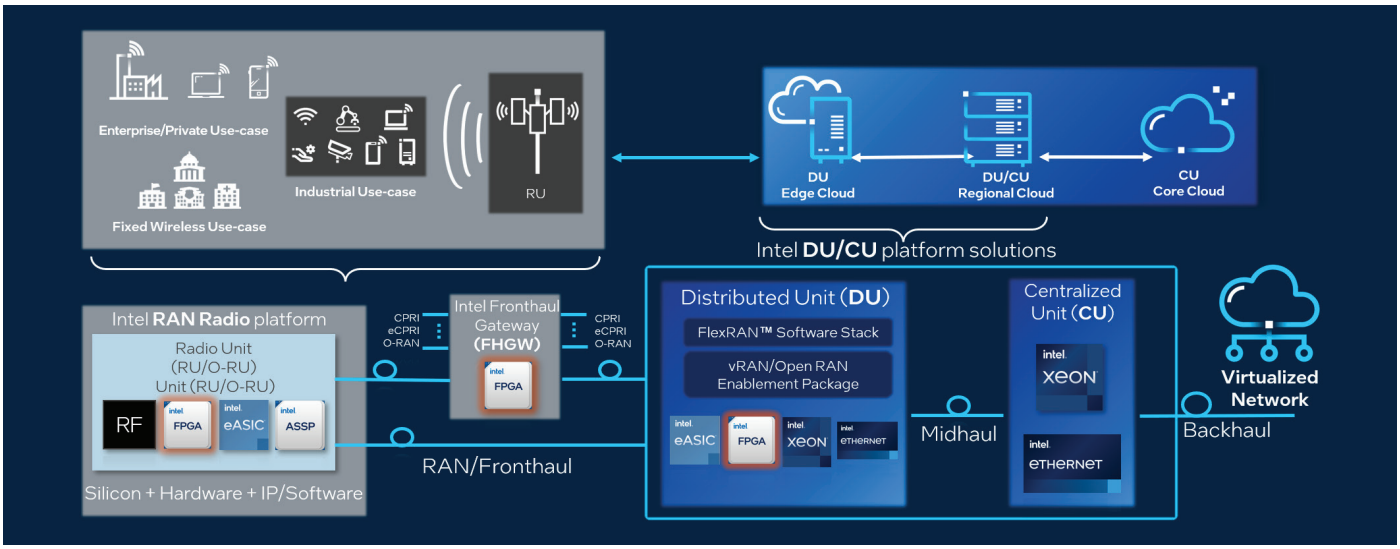

Figure 2. Intel offers end-to-end silicon solutions for 5G O-RANs.

Digital transformation leads to the virtualizing of RAN functions to bring more agile and multifunctional software-defined RANs in the form of virtual RANs (vRANs). The majority of vRAN functions can be implemented on Intel Xeon CPUs. However, certain high-bandwidth, low-latency workloads are better served when implemented in hardware powered by Intel Agilex 7 SoC FPGAs.

Functions such as encryption/decryption, timing synchronization, channel state estimation, error correction, beamforming vector calculations, fronthaul interfaces (e.g., CPRI, eCPRI, RoE), and protocol conversion (e.g., CPRI <-> eCPRI and RoE <-> eCPRI) can all benefit from the pipelining and hardware algorithmic implementation afforded by the soft logic in FPGAs.

In a competitive multi-vendor landscape, O-RANs are seeing the introduction of unique, feature-rich implementations that differentiate one offering from another. In turn, this new marketplace is dramatically furthering the technological advances of RANs, pushing the bounds of what is possible. None of this could be achieved without FPGAs, which enable developers to quickly and easily design and implement their differentiating technologies (their "secret sauce").

Due to the unique technological advantages afforded by Intel's EMIB and chiplet technologies, system integrators have the possibility to leverage a broad variety of XCVR capabilities utilizing Intel chiplets. As a result, customers not only can utilize the FPGA's soft logic flexibility and time-to-market (TTM) advantages, they can also leverage Intel's EMIB technology and hardened transceivers to choose from a variety of configurations tailored to their given application requirements. In turn, customers are able to take on many unique and differentiating feature sets.

Intel Agilex 7 FPGAs allow all these workloads to be implemented in a massively parallel fashion, thereby providing extreme performance while consuming relatively lower power. Offloading these tasks to the FPGA frees up cores on the host CPU, thereby allowing service providers to monetize those cores for other critical workloads. To complement the extreme compute performance afforded by the Intel Agilex 7 FPGA programmable fabric, F-tile transceivers support the extreme data bandwidth required to feed this compute engine.

F-Tiles for Radio Applications

The Groupe Special Mobile Association (GSMA) is a global organization that advocates for the interests of mobile operators and the broader mobile industry. According to the GSMA, a 5G RAN infrastructure's total cost of ownership (TCO) could rise by 65% compared to 4G RAN costs, and energy expenses may increase up to 140%.[2] As a result, minimizing these costs is crucial for the operator's business.

The power consumption of FPGAs used in the system is a critical factor with respect to operational costs. Since nearly all radios are convection cooled, the design's thermal efficiency becomes definitive, emphasizing the importance of low-power FPGA design implementations. By reducing the design's power, the need for cooling is decreased. Smaller passive heat sinks mean a reduction in size, weight, and wind loading on the tower, making installation easier. The Intel® Hyperflex™ FPGA Architecture allows Intel Agilex 7 devices to consume up to 40% lower power than their Intel® Stratix® 10 FPGA predecessors[3], helping system integrators meet or exceed their power budgets and reducing the operator's cost of ownership.

Today's wireless infrastructure comprises heterogeneous networks composed of various radios of different sizes, such as femtocells, picocells, microcells, and macrocells. Each cell type and unique variant supports different features and functions, including the number of bands, RF output power level, RAT technology, the number of antenna elements, component carriers, the number of users per cell, and configuration.

With new standards and communication protocols emerging at a rapid pace, radio designs are in a constant state of flux. To remain competitive, operators expect enhanced feature sets and greater computation and bandwidth capabilities, which put more pressure on power and thermal requirements. System integrators and operators alike demand flexible and scalable platforms to quickly adopt changing standards and evolving performance requirements, minimizing design effort, resources, and costs, cutting test time and inventory, and maximizing design reuse. The combination of FPGAs' soft logic programmability, EMIB packaging technology, and the use of hardened F-tiles allows for unparalleled flexibility, scalability, and the ability to upgrade products—requirements that are often difficult to reconcile. Meanwhile the quick time-to-market and affordability of FPGAs complement these advantages.

Radios Leveraging JESD204

Data-intensive wireless applications continue to push the boundaries for an interface protocol standard that can deliver data fast and efficiently. The JESD committee created the JESD204 data converter serial interface standard to standardize and reduce the number of data inputs/outputs between high-speed data converters, such as analog-to-digital converters (ADCs) and digital-to-analog converters (DACs), and other devices, such as FPGAs.

The JESD204 serial interface has several versions: JESD204, JESD204A, JESD204B, and JESD204C. In conjunction with JESD204C Intel soft FPGA IP, the F-Tile transceivers in Intel Agilex 7 FPGAs and SoC FPGAs support JESD204x communications between DACs and ADCs targeting 5G O-RAN deployments. Current IIntel JESD204C IP and F-tile transceivers support data rates up to 32.45 GHz to facilitate wider bandwidth and higher sampling rates support.

Radios Leveraging Ethernet

Historically, most FPGA-based Ethernet IP solutions have been based on soft IP (i.e., logic built using the FPGA's programmable fabric). In these cases, although the SERDES blocks were implemented as hard IP, the media access controller (MAC), forward error correction (FEC), and physical coding sublayer (PCS) functions were implemented as soft IP. In addition to consuming the FPGA's programmable logic resources, soft IP implementations also consume higher power. With the bandwidth requirements increasing, coupled with the need for higher transfer rates, the soft IP size grows larger, making a bigger impact on area and power requirements.

Ethernet is being used more than ever at higher data rates in transformative markets like automotive V2X infrastructure, data centers, and—of course-RANs. Hence, using a hardened IP solution for Ethernet is highly beneficial because of the area savings, power savings, reduction in latency, and faster time to market. A hard IP tile-based solution also affords users the benefit of lowered cost because they can target smaller FPGA devices and/or use the freed-up programmable logic resources for implementing their value-added or proprietary IP.

The F-tile supports transceivers with a maximum rate of 4x116 Gbps PAM4 and 16x32G with NRZ modulation. Together with Intel's IP for CPRI, eCPRI, and ORAN IP, the F-Tile can support the 3GPP and ORAN compliant CAT-A and CAT-B O-RUs. These IPs supports optional KR-FEC and KP-FEC. The Ethernet IP also supports auto-negotiation, link training, and forward error correction, along with optional support for the 1588 Precision Time Protocol (PTP). When using the F-Tile, there is also soft IP support for Triple Speed Ethernet, and 10GE MAC with multi-rate support.

F-Tiles for Fronthaul Applications

RAN deployments are divided into two main architectures. In C-RAN (Cloud RAN) deployments, the DUs are centralized and may be located as far away as 5/10/20 km or more from the RUs. This is a typical deployment for dense urban scenarios and allows for baseband pooling, which is a technique to optimize processing needs. By comparison, in D-RAN (Distributed RAN) deployments, the DUs are distributed at the cell sites. This is a typical deployment scenario for rural environments.

A 5G D-RAN involves a combination of the RUs (radios and antennas), the DU, and any ancillary equipment (Figure 3). Here, a cell site router (CSR) connects the RUs with the DU. The CSR also connects the DU with the rest of the mobile infrastructure where the CU resides. The fronthaul connection between the RU and the CSR/DU is usually fiber. The midhaul connection between the DU/CSR and the aggregation router/CU may be implemented using fiber, microwave, or satellite link using IP over multi-protocol label switching (MPLS).

Handling all of these diverse fronthaul applications requires multiple instances of Ethernet MAC and PCS functions while, at the same time, leaving enough headroom in the FPGA's programmable fabric to process the incoming data. If the I/O interfaces were implemented in the programmable fabric, it would be difficult to process the large workload while meeting thermal, latency, and throughput requirements. Having hardened XCVR functionality allows more headroom in the FPGA fabric to process these computationally expensive workloads while maintaining a low thermal envelope.

Fronthaul Gateways

Today's 5G O-RAN use cases demand a significant increase in fronthaul bandwidth. In addition to 5G, many networks also leverage the legacy generation of 4G/LTE devices and protocols. Combining 4G and 5G increases fronthaul traffic and connectivity challenges.

Fronthaul deployments typically employ a combination of the common public radio interface (CPRI) for legacy 4G applications and enhanced CPRI (eCPRI) for 5G applications. 5G fronthaul traffic protocol conversion is very demanding as it requires high throughput computation with low latency. Such conversions must also deliver extremely accurate clock synchronization using precision time protocol (PTP), which is a timing-over-packet protocol defined in the IEEE 1588v2 standard. A device called a fronthaul gateway (FHGW) is often used to make conversions between fronthaul protocols (e.g., CPRI <-> eCPRI).

F-Tile functionality in an Intel Agilex 7 FPGA allows the device to import and export Ethernet data up to 400 Gbps. The combination of hard IP in F-Tile transceivers and soft IP implemented in the FPGAs programmable fabric facilitates the support of multiple fronthaul formats. The Intel fronthaul enablement package makes CPRI <-> eCPRI conversion possible to accommodate 5G NR O-RUs and legacy generation radios.

Virtual CSRs

A key part of an O-RAN solution is a cell site router (CSR). One aspect of virtualizing the RAN is to virtualize the CSR. Intel's accelerated virtual CSR (vCSR) solution is implemented on an Intel® FPGA SmartNIC N6000-PL Platform with O-RAN-compliant precision timing.

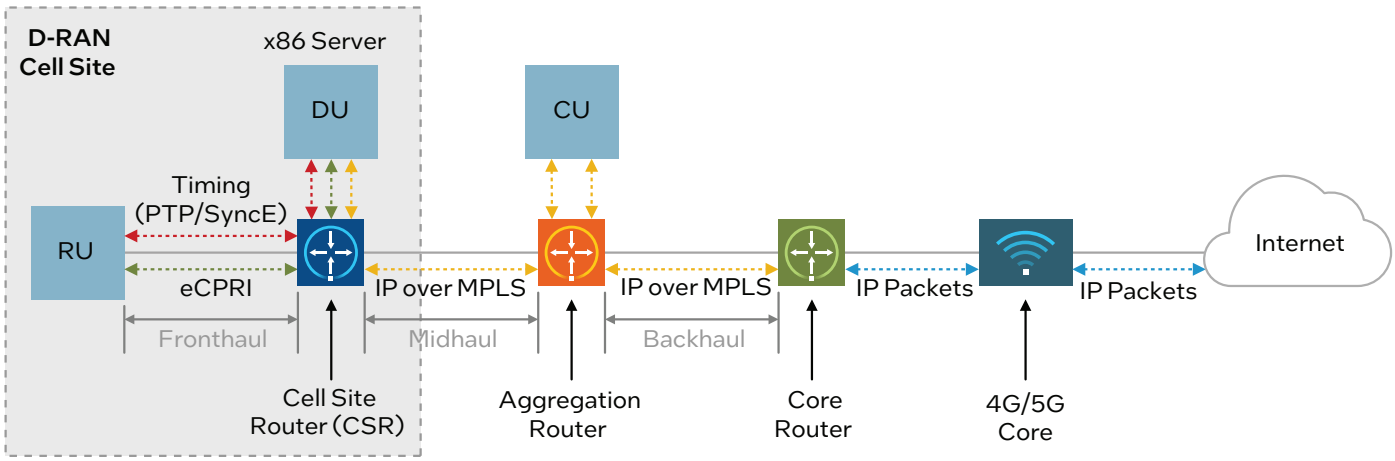

Figure 3. High-level representation of a cell site in the context of a 5G RAN.

At the time of this writing, most CSRs are implemented as stand-alone devices external to the DU. However, as the industry moves from 4G to 5G, there is an opportunity to use an Intel Agilex FPGA-based acceleration card to create a virtual CSR (vCSR). This vCSR implementation can also support an optional FHGW function for 4G radios along with optional future proofing with 6G Numerology acceleration. Furthermore, this card could be combined with a DU implemented in commercial off the shelf (COTS) server, thereby providing a complete DU solution with vCSR integrated using optional FHGW and Baseband Acceleration as a single COTS platform. (If you want to learn more about this single COTS solution, check out Intel's vCSR Solution Brief.)[4]

Powered by F-tiles in Intel Agilex 7 devices, the integrated vRouter function handles L2 traffic management and L3 routing, and also provides 1588/PTP support. This solution is designed to support the Juniper Cloud-Native Router stack, which provides a commercial-grade, high-performance, scalable routing solution. Open-source options such as the FRRouting Project can also be used. These vCSR tasks are run on the Intel Agilex 7 FPGA's hard processor subsystem (HPS) cores, thereby freeing up the host server's processor cores for revenue-generating functions.

The Intel Accelerated vCSR can support the midhaul (F1 interface) over the integrated MPLS or SR-MPLS data-planes. Support for 1588/PTP, especially clock accuracy, is key for any CSR. The ability to generate these clocks adds value to Intel's Accelerated vCSR solution as it removes the need to have an additional special switch for timing generation. Furthermore, Intel's vCSR supports the clock accuracy required for Class B and C requirements in 5G.

F-Tiles for Passive Optical Network Applications

A passive optical network (PON) is a point-to-multi-point telecommunications technology for delivering broadband network access to end-customers via fiber-optics. Using this technology, internet service providers (ISPs) can deliver high-bandwidth "last-mile" service to their customers. A high-level representation of a PON is illustrated in Figure 4.

The starting point for a PON is the optical line terminal (OLT), which resides at the ISP's central office. A typical OLT will be presented in the form of a rack. In addition to the PON card that communicates with the outside world by means of a single optical fiber, this OLT rack will contain several other cards/modules to provide control and switching functions, redundancy protection, power supply units (PSUs), etc.

At the customer end of the PON are optical network units (ONUs). The optical traffic being communicated back and forth between the OLT and the ONUs passes through a hierarchy of passive (unpowered) optical splitters, hence the "passive" portion of the PON's name.

One of the more recent PON developments is 25G PON (a.k.a. 25GS-PON), which operates at 25 Gbps downstream and both 10 and 25 Gbps upstream. Communications industry operators, system vendors, component vendors have come together in the form of the 25GS-PON MSA Group,[5] whose mission is to define and develop 25G PON technology.

Achieving the extreme data rates demanded by today's PONs involves many challenges, the most significant of which is handling the burst and random nature of upstream traffic. In the upstream direction, any data from the customers comes in random bursts. Data from multiple customers is combined using time-division multiple access (TDMA) for delivery back to the central office's OLT. TDMA allows multiple users (ONUs) to share the same frequency channel by dividing the signal into different time slots. The users transmit in rapid succession, one after the other, each using its own time slot. However, since the ONUs reside in different physical locations, the phase relationship between their data as it appears at the OLT is random.

Due to the "burstiness" of the data as it appears at the OLT, coupled with the randomness of the phase relationship of each burst, a special type of clock data recovery (CDR) implementation is required. This implementation, which is called burst-mode CDR (BCDR), allows for obtaining symbol lock in a relatively short period of time, but achieving it requires sophisticated transceiver technology. The F-Tiles in Intel Agilex 7 devices provide a wide range of features to address the challenges posed by PONs. Some examples are as follows:

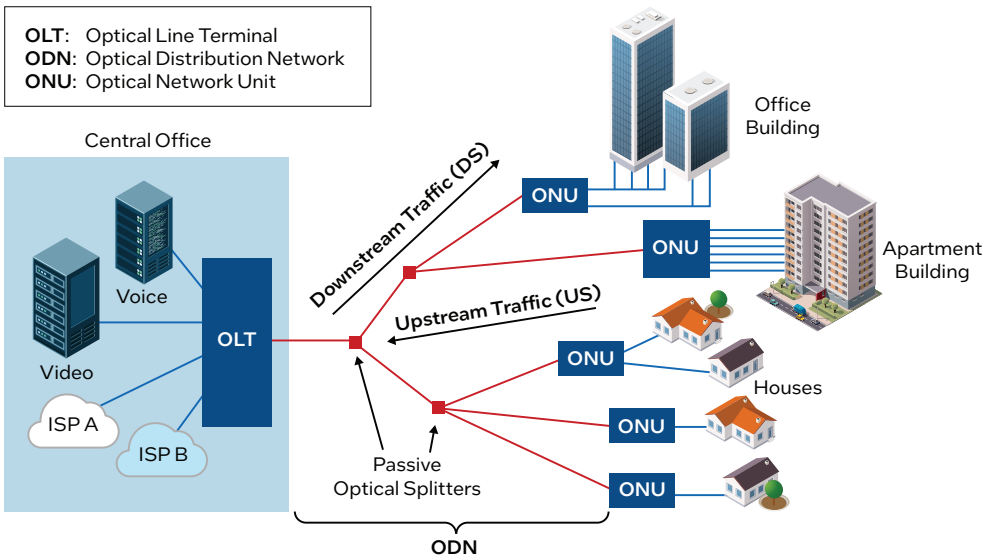

Figure 4. High-level representation of a passive optical network.

- CDR Freeze Mode: The F-Tile transceiver contains a circuit to lock the CDR's frequency in-place to keep it from drifting in between PON bursts. In this way, the CDR needs to perform phase acquisition only on the next data burst. This phase acquisition can be done in a short enough time by using the F-Tile phase-locked loop (PLL)-based CDR to recover the proper payload in the burst.
- CDR Lock-to-Reference (LTR) Mode: The F-Tile transceiver contains a circuit to force the CDR to lock to its incoming reference clock instead of trying to lock to the incoming data stream. This is beneficial during more catastrophic conditions by forcing any frequency deviation back to 0 ppm.
- Deterministic Latency: As was previously noted, PONs require deterministic latency during regular operational modes and during the ranging process. The F-Tile transceiver provides this capability even when the BCDR loses and regains lock.
- RX Equalization: To ensure that the highest quality signal is obtained for data and clock recovery, the F-Tile transceiver provides for static receiver equalization. These equalization settings can be hand-selected by the designer to compensate

Intel Agilex 7 FPGAs and CXL

One of the problems associated with systems employing multiple processors in the form of CPUs and XPUs (GPUs, NPUs, and FPGAs) is that each processor may maintain its own local high-speed cache memory. When multiple cores maintain caches of a common memory resource, problems may arise with incoherent data. The term "cache coherence" refers to the process of managing such conflicts by ensuring that a coherent view of the data values is maintained in the multiple caches.

The processors may communicate directly via a shared bus in conjunction with a proprietary or industry-standard protocol. One very common standard protocol that is used for this purpose is Peripheral Component Interconnect express (PCI Express or PCIe).

Unfortunately, PCIe is less-than-optimal with respect to maintaining cache, memory, and storage coherence. By comparison, a more recent communications protocol called Compute Express Link (CXL), which is built on the PCIe physical and electrical interface, features cache-coherent capabilities for accessing device memory and system memory.

As was previously noted, in addition to F-Tiles, Intel Agilex 7 FPGAs I-Series and M-Series also support R-tiles, which feature Intel CXL protocol hard IP. This hard IP allows for full bandwidth 5.0 x16 configuration support for CXL 1.1 and 2.0 with minimal use of FPGA fabric resources. CXL is also supported by 4th Gen Intel Xeon Scalable processors (previously codenamed Sapphire Rapids (SPR)). CXL will improve the coherency, latency, and speed of CPU-to-CPU and CPU-to-XPU communications.

CPUs do not have the capability to take on high bandwidth, low latency 400 Gbps Ethernet data by themselves. However, a system utilizing the combination of one or more 4th Gen Intel Xeon Scalable processors coupled with one or more Intel Agilex 7 FPGAs equipped with both CXL (via their R-Tiles) and 400 Gbps Ethernet (via their F-Tiles) enables CPUs to take full advantage of high bandwidth, low latency data.

O-DUs have to perform the complicated task of taking in large amounts of data of a variety of protocols with little packet loss and to perform a multitude of both hardware and software-oriented tasks. By leveraging the strengths of both CPU and FPGA devices, system integrators are capable of taking on more radios supporting higher bandwidths while offering a more diverse set of complex functions, coupled with the ability to quickly and easily implement unique functionality using software and soft logic, all enhanced by the low latency and cache coherency afforded by CXL.

Conclusion

O-RAN standards-compliant O-RUs, O-DUs, and O-CUs can be powered by the vast range of Intel's hardware silicon portfolio, which includes Intel Agilex 7 SoC FPGAs, Intel Xeon CPUs, and a variety of Ethernet solutions. Intel Agilex 7 FPGAs equipped with F-Tile transceivers have proven themselves to be an essential component of 5G O-RAN deployments. These devices can accelerate various workloads, providing extreme low latency, high bandwidth performance while consuming relatively little power.

In the case of radio applications, the F-Tile allows designers to take full advantage of protocols like JESD204C up to 33.45 Gbps and 400 Gbps Ethernet. With respect to fronthaul applications, the F-Tile allows developers to create sophisticated applications like fronthaul gateways and virtual CSRs. In the case of PON, the F-Tile provides a wide range of features that address the challenges associated with the latest 10G, 25G, and 50G implementations of these networks. Furthermore, the combination of the F-Tile (which supports 400 Gbps Ethernet) and the R-Tile (which supports PCIe 5.0 and CXL) allows 4th Gen Intel Xeon Scalable processors (which also support PCIe 5.0 and CXL) to take full advantage of the high bandwidth, low latency data being processed by the Intel Agilex 7 devices.

In conclusion, the F-Tile transceiver in Intel Agilex 7 SoC FPGAs is a vital component that enhances the transmission of data and accelerates a wide variety of workloads, providing extreme performance while consuming relatively little power. The F-Tile is an essential element of wireless applications, especially in 5G O-RAN deployments, which require a flexible radio access network architecture that allows interoperability between hardware and software provided by different vendors. Moreover, Intel's FPGA IP in conjunction with Intel Agilex 7 devices supports the efficient execution of 5G RAN workloads, making the F-Tile an indispensable technology for developers serving the wireless industry.

If you want to unleash the power of F-Tile transceivers in Intel Agilex 7 FPGAs and SoC FPGAs to satisfy the extreme data demands of your own projects, contact your local Intel representative today to learn more.

Learn More

- Intel Agilex 7 FPGAs
- Intel FPGA Wireless Page

References

1 O-RAN Alliance

2 GSMA: 5G-era Mobile Network Cost Evolution

3 Intel Performance Claims

4 vCSR solution brief

5 25GS-PON MSA

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

No product or component can be absolutely secure.

Your costs and results may vary.

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

© 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.

SS-1161-1.0