Network Slicing made easy with Precise-ITC Flex-E and Multi-channel MAC Cores

1 Overview

The soon to be ratified Fifth Generation (5G) wireless standard is going to change the way people and machines use the network. It's going to allow not only high bandwidth applications like augmented reality but also real-time applications like remote surgery. There's a lot of excitement at the edge, as evident with the proposed Broadcom & Qualcomm merger, but at Precise-ITC we're looking at the impact in the core.

5G along with Software Defined Networking (SDN) are ushering-in the concept of Network Slicing. 5G is going to create the capability for very-low latency networks. The promise of SDN is a more agile network that will unleash a variety of functions such as bandwidth-on-demand and the ability scale up (and down) capacity in near-real-time. Together it allows Mobile Network Operators (MNOs) the ability to manage and operate multiple virtual networks over common physical network infrastructure. It enables the virtual partitioning of the Radio Access Network (RAN), core network components and switching and aggregation networks all the way to the data centers, where content and applications are hosted.

This all makes sense at the edge and it isn't hard to visualize the relationship for deterministic network slices in applications like vehicle-to-infrastructure applications. It's significantly foggier in the core where channels are aggregated into "hard-pipes".

2 Introduction to Network Slicing

An end-to-end service orchestration engine "slices" virtual (logical) networks out of physical network resources. The slices can be carved-out using dedicated QoS assignment or committed bandwidth assignments. The same functionality can be achieved when network endpoints apply priority tagging and the network is able to honor tagging from end to end, for example in remote medicine the end equipment can apply high priority tagging or time-sensitive-network tagging to packets at the surgeon's robotic hand and the patient's robotic surgeon. However, this capability can only be achieve if the network operator is able to assign priority and map performance for the packets as they traverse across the network from end-point to end-point.

At the edge it's easy to tag packets. Indeed the RAN already does this for voice versus data packets at the RAN. The challenge has been the ability to maintain priority of the packets as they move through the aggregation network, and better yet to manipulate performance of the network to honor priority tagging (per 802.1Q).

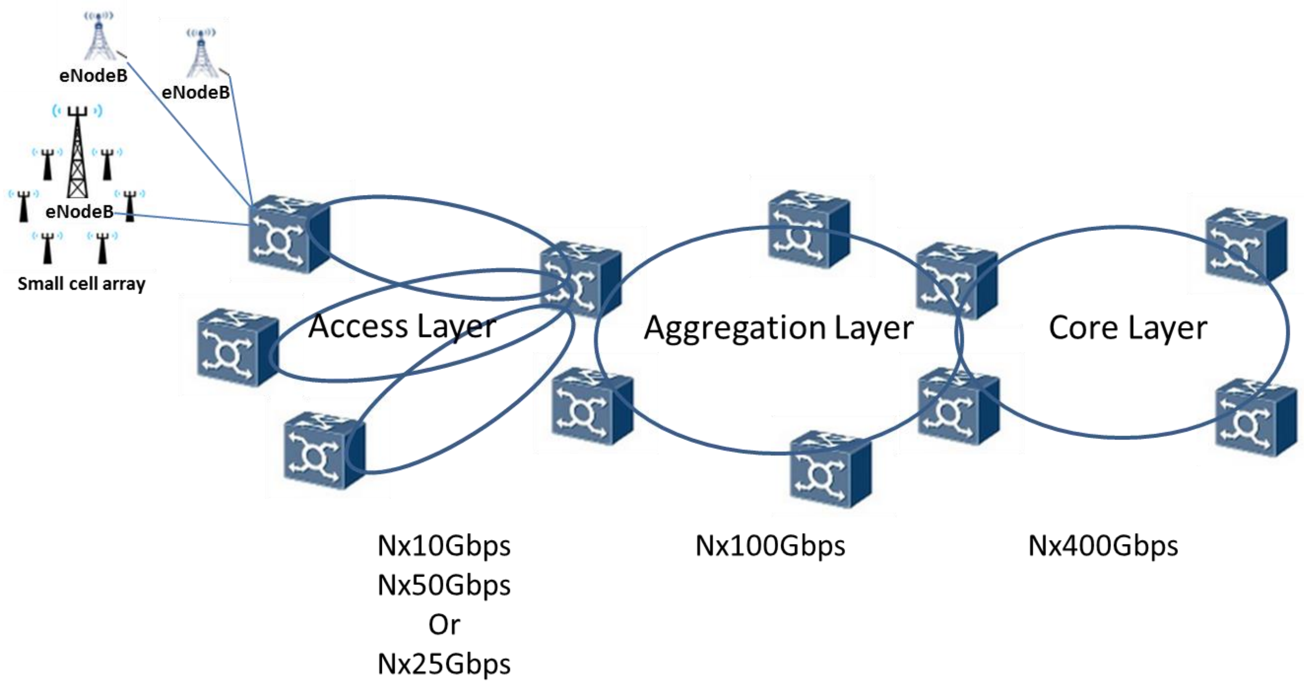

Figure 1: Network View

3 FlexE 2.0 Overview

The Optical Internetworking Forum (OIF) recently ratified FlexE 2.0 which provides part of the solution for network slicing in the aggregated network. FlexE is a shim within the Layer 1 PHY. The first iteration of the specification created the ability to aggregate asynchronous lower-rate MACs onto a higher rate PHY, specifically up to ten 10GE signals, each running on its own timing domain into a single 100GE PHY. This technology was rapidly adopted at the aggregation layer. Functionality is achieved through Time Division Multiplexing (TDM) using traditional bit-stuffing to map low-bandwidth pipes to high bandwidth pipes. The opposite function at the far-end restores original timing. Any intermediary hops are oblivious to the presence of the lower bandwidth pipes unless the channels are demultiplexed.

The FlexE Core is compliant to OIF FlexE 2.0 IA. It supports 20 calendar slots and each calendar slot is 5Gb/s or 4 calendar slots each at 25Gb/s per 100G instance. It can support up to 256 logical MAC clients. FlexE 2.0 IA supports 200GE and 400GE PHY with multiple 100G FlexE instance. This greatly addresses the need for higher bandwidth demand for today's application.

FlexE 2.0 enhances the original interoperability specification with the ability to aggregate much higher data-rates, but more importantly, a richer calendar provides more granularity for increasingly lower bandwidth channels, which are necessary ingredients for aggregating 5G data streams. This provides flexibility since network operators can configure the inputs at 25Gbps, for example with high bandwidth edge processing, or at 5Gbps, for example to aggregate multiple RANs, or any combination thereof. This flexibility means that the same equipment can target different applications.

Amongst other challenges associated with the adoption of 5G is the uncertainty over network architecture. While MSOs and OEMs have envisioned different deployment scenarios, it's anybody's guess what will be the killer app. There is uncertainty how much processing will occur at the edge or in the fog. Knowing this will dictate the bandwidth requirements at the aggregation points, however, in absence of certainty, a flexible architecture, using Precise-ITC cores, provides the capability to adapt equipment to meet the emerging requirements. In fact using Precise-ITC's new proprietary SuperFlexE Extension tm operators can have visibility down to 1G per channel.

Precise-ITC recognized that FlexE2.0 with its dense calendar was necessary for network slicing but insufficient for time-sensitive or deterministic applications. They developed a multi-channel MAC (MCM) controller that sits on top of the FlexE shim (logically). Refer to Figure 2: Multi Channel MAC and FlexE 2.0 Shim in ISO Stack.

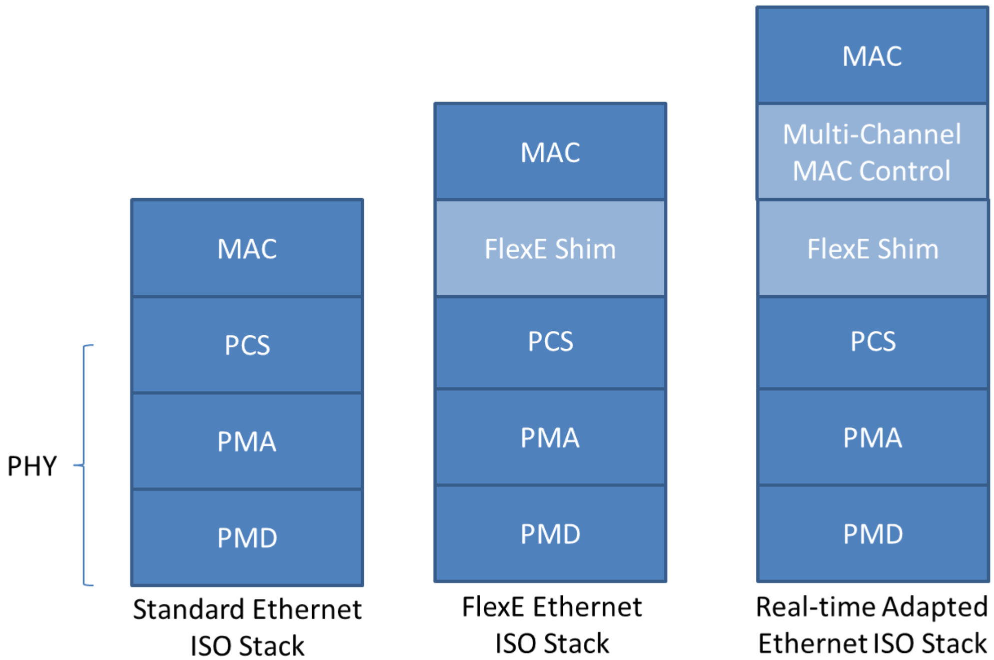

Figure 2: Multi Channel MAC and FlexE 2.0 Shim in ISO Stack

4 Proprietary Multi-channel MAC

The proprietary multi-channel MAC (MCM) from Precise-ITC provides the ability to respond to priority flows as defined by 802.1Qbb, and in conjunction with 802.1Qbu (draft) standard for time sensitive networking, adds the ability to pre-empt flows at layer 2 with no impact to the Layer 1 PHY.

Whereas transport systems may have a single MAC source and destination, there is more granularity closer to the edge. Here again, the flexibility provided by the MCM to match the data sources future-proofs today's implementation. The MCM is used to evaluate tagging in the incoming data and can be scaled up or down correspond to the type of traffic being aggregated.

In the example of Network Slicing, the multi-channel MAC allows control of the user experience based on each individual channel. This can allow the operator to tag packets and see a corresponding dynamic change in bandwidth or latency or add features such as security to a telemedicine application. The network operator can manage a specific SLA (service level agreement) on a per channel or frame-by-frame basis.

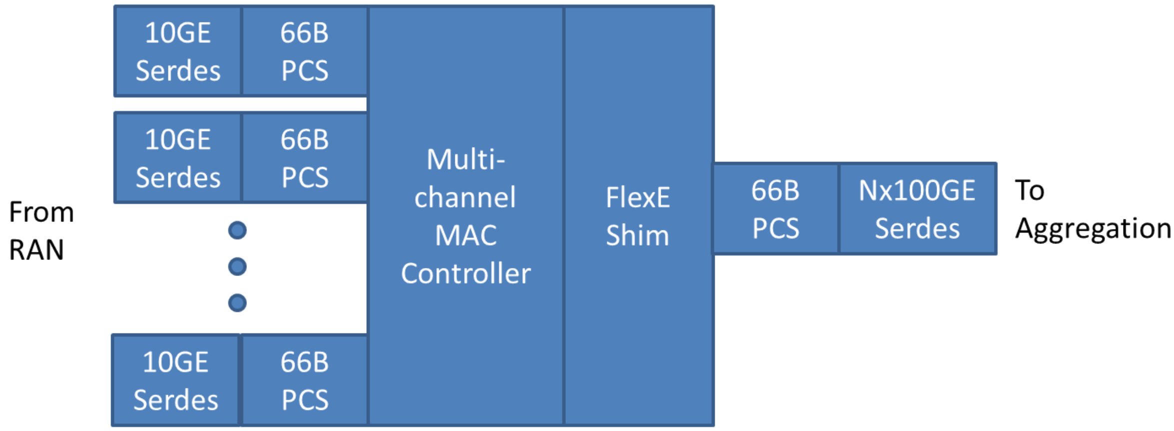

Figure 3: Block Diagram of Multi-Channel MAC and FlexE IP Core

5 Prioritizing Tagged Packets

The multi-channel MAC (MCM) sits logically above the FlexE shim. The MCM doesn't terminate the incoming frames but provides a bridging function. For Network Slicing, the MCM has the role of unpacking the priority tags on the incoming data stream. The MCM can prioritize the higher priority frames over the regular frames within a MAC channel (802.3br/1Qbu). It also supports multi-priority flow control (802.1Qbb).

As shown in Figure 4: Functional Block Diagram of Pre-emption Mechanism, the MCM parses MAC frames as they arrive for aggregation. Priority tagged frames are advanced in priority prior to loading into the FlexE calendar for transmission.

Within a "fat-pipe" of aggregated data, the MCM can re-prioritize the MAC frames within its internal buffering and queuing space. Prioritized frames are presented to the FlexE shim in advance of regular traffic, providing they are part of the same channel. This provides efficient aggregation to respect time-sensitive traffic.

6 Dynamic Bandwidth Allocation

The FlexE core can bond multiple lower bandwidth MAC channels into a high bandwidth PHY (i.e. multiples 5Gs to 100G/200G/400GE) or break up a large bandwith MAC into smaller PHYs (i.e. 400G MAC carried over 4 x 100GE PHY).

The MCM can also bond multiple input channels to a single higher bandwidth channel in conjunction with the FlexE shim. This is the traditional role of aggregation equipment. In this case, the same core can be placed closer to the edge to aggregate multiple RANs. For example, a football stadium may have dozens 5Gbps antenna that are dynamically aggregated to several pairs of 25Gbps links within the MCM and FlexE core.

The MCM is the network slicer that operates seamlessly with the FlexE core to provide hitless, dynamic scaling such as using 25G calendar slots with the MCM to honor priorities queues with granularity to 1Gbps. It can also dynamically adjust bandwidth at the edge on a per MAC basis and make the translation to the FlexE calendar.

.

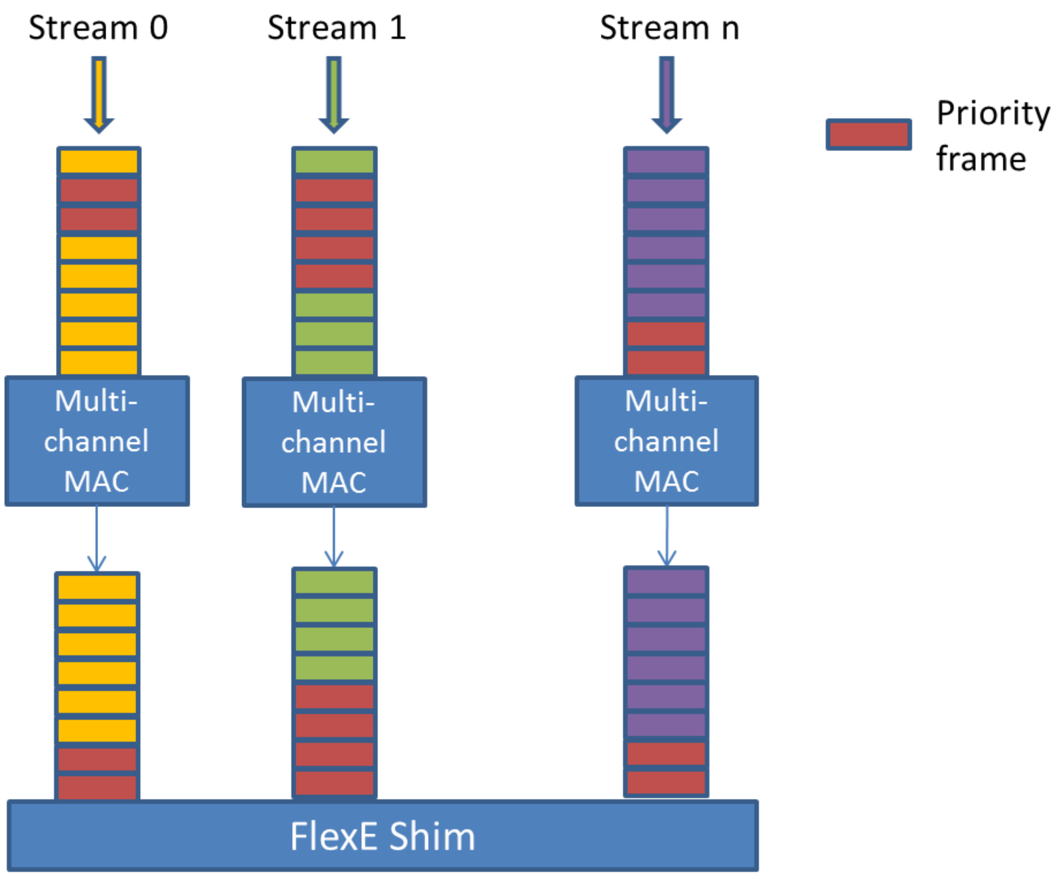

Figure 4: Functional Block Diagram of Pre-emption Mechanism

The International Telecommunication Union states that technologies designed for 5G must reach a peak data transfer rate of up to 20Gbps in under 1ms latency which is a significant technical hurdle at the edge (base-station to handset), but any latency savings at the edge are often lost in the metro network. For instance when the 20Gbps tributary is multiplexed into a 100Gbps pipe. The combination of FlexE and Multi-channel MAC allow MNOs to achieve these objectives, not only at the edge, but also through the metro. This kind of breakthrough is required to enable real-time applications over moderate distances, such as remote surgery.

Certainly the MNOs and carriers would be able to monetize technology that allows them to provide on-demand Class of Service (CoS), were it available and so there is demand for the technology. But FlexE 2.0 is brand new and like most technologies the uptake is slow; mostly because the technology doesn't exist for OEMs to begin experimenting.

The Precise-ITC FlexE SHIM Wrapper Core works seamlessly with the Multi-Channel MAC and Ethernet PCS Cores. FlexE IP also supports lower rate PCS Cores, which is supported through the Super FlexE Extension. More importantly, the cores are available in both FPGA and ASIC versions. This provides OEMs with a low-risk strategy to begin experimentation and address this emerging need. Precise-ITC is first to market with IP cores that support network slicing in the core, supporting connections from 10Gbps through multi-100Gbps implemented through a combination of a Flex-E shim and a multi-channel MAC controller. These technologies allow operators to provision bandwidth and priority at the MAC level even when these channels are aggregated at very high density.

About the authors

Silas Li is the Director of Engineering at Precise-ITC. He has over 20 years' experience designing time division multiplexed and Ethernet systems which maintain timing across complex networks. He is currently developing 800Gbps IP for datacenter interconnect for major companies in North America and China.

Murray Baldock is VP of Sales and Marketing for Precise-ITC. A 30-year veteran of the industry he estimates that 70% of the world's IP traffic traverses a product with which he's been involved.