WHITE PAPER

The Evolution of Mobile Fronthaul to Support RAN Architecture Migration to 5G

Why the Architecture of the Radio Access Network (RAN) Is Changing

It is widely understood that there is rapid bandwidth growth in wireless networks and there is plenty of literature detailing the likely path that radio architecture in the RAN will take to accommodate this growth. Technologies like coordinated multi-point (CoMP), carrier aggregation, and more complex and dynamic multiple-input multiple-output (MIMO) schemes will allow radio access networks to scale to accommodate immense amounts of bandwidth per handset as rollouts of 4G and 4.5 Long Term Evolution Advanced (LTE-A) networks are completed and operators ultimately move to 5G.

What seems to be discussed less is how the transport network in the RAN will mold itself to accommodate these new technologies. There are a number of aspects of these technologies that will force the RAN transport network to totally reinvent itself over the next 10 years. Key drivers will include:

1. The need for immense amounts of bandwidth in the RAN (moving from less than 1 gigabit per second (Gb/s) of actual data transmission per tower to multiples of 10 Gb/s)
2. Energy efficiency
3. The requirement for extremely low latency to provide for a better user experience
4. The addition of different classes of service for different service types
5. Centralization of control to accommodate more complex network planning and optimization
6. A push toward virtualization and standardization to simplify network hardware.

One of the most significant changes in the architecture of the RAN transport network is the move to mobile fronthaul. Fronthaul deployments are currently in their early stages with varying levels of adoption across the globe. Initial deployments have started in the Asia-Pacific region, supported by a **business case*** that proves better network economics for all wireless network technologies from 2G to 4G by moving to a centralized-RAN (C-RAN) architecture. These economics vary from carrier to carrier and from region to region. At present, fronthaul is in its infancy outside the APAC region, wi th many operators evaluating the technology and in some cases running initial lab or live network trials. In these cases, network operators are typically looking at 4.5G and 5G as the primary driver for migration to fronthaul.

Today's Decisions Affect Future Advantages

Decisions that are made today concerning fronthaul choices will affect the ability and cost required to take advantage of emerging wireless technologies in the future. In addition, one must consider that not doing anything now, while waiting for fronthaul to mature, will limit a network's ability to take advantage of advanced technologies that are already being rolled out today, such as baseband unit (BBU) pooling and coordination. So, the trick will be for network operators to figure out how to roll out fronthaul technology today while not limiting the network's ability to scale and adapt as fronthaul matures. Of course, using fronthaul alone does not address and solve all of the drivers and requirements listed above. The move to fronthaul helps with energy efficiency and latency primarily. To help with some of the other considerations, many network operators are looking to other technologies like software-defined networking (SDN)/network function virtualization (NFV) and time-sensitive Ethernet. How these technologies will play in future architectures will be discussed later.

The purpose of this paper is to provide a view of the direction that RAN transport architecture will take and how this will enable wireless service providers to make smarter decisions as they move down the path of re-architecting their RAN. This paper will also discuss some of the technologies that Infinera will use to make the move to this next-generation architecture more seamless and scalable.

Please note, this paper assumes a degree of understanding of current RAN architectures and the use of mobile fronthaul networks. Background reading suggestions can be found in the Further Reading section at the end of the document.

What is the Target Architecture and Why?

As previously mentioned, there are a number of drivers that will lead to a dramatic change in the way that RAN transport architectures are built in the future. The current view of the net result of these drivers will be a RAN network architecture that utilizes a mix of fronthaul and backhaul with full virtualization of resources. In addition, there will be a push to move content closer and closer to the subscriber in order to provide for a better user experience and meet 5G latency targets. This architecture is shown in Figure 1.

In the below figure, the controller/orchestrator provides orchestration and control plane for all network and compute resources in the RAN, including those located at the tower. Many resources will eventually move to generic off-the-shelf hardware architectures. This is likely for radio resources and a possibility for network resources as well.

Now this paper will focus on the previously mentioned drivers that specifically affect how the RAN transport network will evolve one at a time. The first was the growth in bandwidth as we move from 3G and 4G to 4.5G and ultimately 5G. It is not yet determined how much bandwidth will be required to be delivered by the transport network to each tower but it is likely to be on the order of 10 Gb/s or multiples of this. It is largely accepted that future RAN transport architectures will be built utilizing fronthaul or a mix of fronthaul and backhaul, often called crosshaul.

Today, most networks are based on a backhaul architecture supporting a global average of less than 1 Gb/s of traffic per tower. Most networks that are moving to fronthaul today are utilizing 2.5 Gb/s Common Public Radio Interface (CPRI) for 4G LTE deployments, with most service providers looking at 9.8 Gb/s CPRI as a next step for fronthaul. It is anticipated that 9.8 Gb/s CPRI will be utilized for 4.5G LTE-A (20 MHz, 8x8 MIMO) and it is likely that this will take us through to 5G. If CPRI were to be utilized for 5G fronthaul deployments, it is likely that it would also utilize 9.8G, or might be more dependent on factors such as the flavor of MIMO that is selected and ultimately what functionality the remote radio heads (RRHs) themselves support.

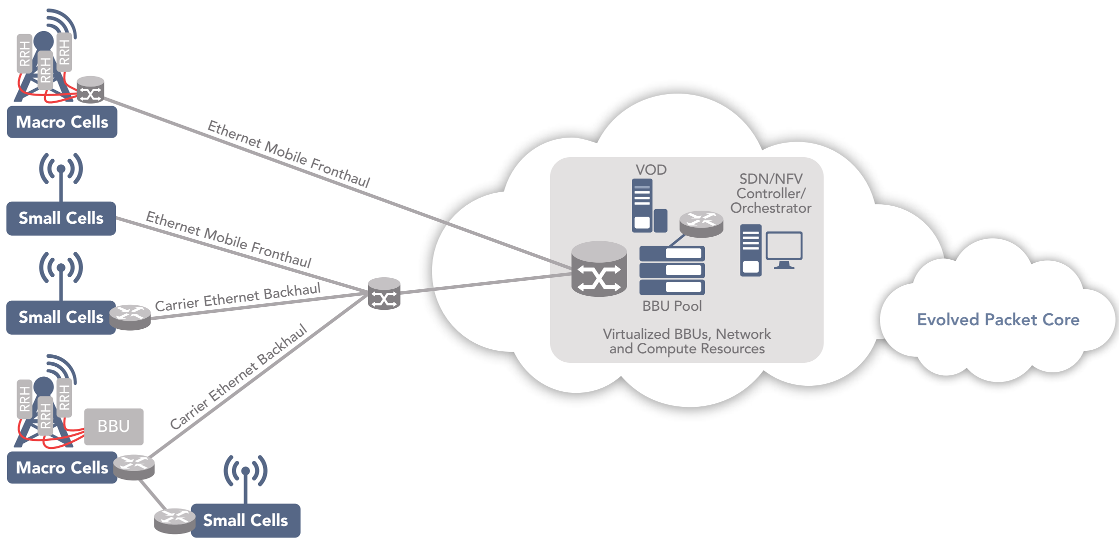

Figure 1: Future 5G Transport Network Architecture

Future Ethernet-based Fronthaul Architectures

In Next Generation Fronthaul Interface (NGFI) discussions within the Institute of Electrical and Electronics Engineers (IEEE), Ethernet seems to be the target for fronthaul in the longer term, with an Ethernet-based fronthaul architecture as a goal for 5G*. If the industry goes down this path, initial line rates will likely be either 10 Gb/s or multiples of 10 Gb/s, although this has not yet been determined. If it is deemed that rates of more than 10 Gb/s rates are necessary, it is likely that the target will be Ethernet rates that line up with those that are designed for and utilized in the wide area network (WAN), which would point to 10 Gb/s, multiples of 10 Gb/s (possibly multiplexed up into 100 Gb/s), and eventually 100 Gb/s. Of course, 25 Gb/s and 50 Gb/s may be considered, but these rates are more accepted and utilized inside the data center and not over the distances required in the WAN environment. Ethernet will also have to be enhanced to support time-sensitive applications to support the demanding latency and synchronization requirements of fronthaul networks. Mobile backhaul is typically Ethernet-based already, and if it continues to be utilized out to the towers with 5G, even if capacity per subscriber increases by a factor of ten or more, operators are still likely to deploy 10 Gb/s or low multiples of 10 Gb/s per tower for backhaul. In any case, Ethernet will be utilized to overcome capacity requirements in the RAN transport network.

Fronthaul as a Way of Reducing Energy Consumption in the RAN

Fronthaul plays an important part in increasing capacity in the RAN and in reducing energy consumption and cost. Fronthaul helps increase capacity by enabling technologies like CoMP and enhanced inter-cell interference coordination (eICIC) to work more efficiently by reducing communications delay between BBUs. As larger amounts of capacity are carried through the RAN, energy utilization and often cost will increase. Because of this, it is an important target of 5G to decrease energy utilization in wireless networks. Fronthaul contributes to this goal by moving energy consumption out of towers, where energy is sometimes expensive, particularly if the wireless carrier does not own the tower, and by allowing for capabilities brought about by BBU pooling. An example would be load sharing that allows a pool of BBUs to serve a pool of RRHs. In this architecture, BBU power can be reallocated based on traffic patterns, allowing for lower overall utilization of power as compared to BBUs acting independently.

Reducing Latency with Fronthaul

Fronthaul also helps to reduce latency in the network. Today, by utilizing a C-RAN architecture with CPRI for fronthaul implementations, high-latency functions (like routing and switching) that typically comprise the backhaul network can be moved away from the towers and deeper into the network. By doing this, the number of aggregation hops that are necessary in the backhaul network is reduced and hence the number of high-latency network devices between the tower and the Evolved Packet Core (EPC) is reduced too. In the future, if it is the goal of the industry to utilize Ethernet for fronthaul instead of CPRI, steps will have to be taken to handle time-sensitive

*In addition to the NGFI initiative, there is work ongoing in 3GPP, ITU, ETSI and other organizations related to 5G standardization.

traffic within the Ethernet standards. Basically, new standards will need to focus on how Ethernet is carried, switched, and routed through the network in order to deal with the time sensitivity of radio traffic.

Classes and Qualities of Service Enabled by Ethernet-based Fronthaul

The longer-term move to all Ethernet for both fronthaul and backhaul not only brings us ubiquity in protocol and operations but it also helps us solve one of our market drivers: the need for different classes of service for different service types. Given the move to hybrid fronthaul and backhaul architectures, once both fronthaul and backhaul are transported using Ethernet, we can take advantage of Carrier Ethernet quality of service (QOS)/class of service (COS) to provide for the different classes and qualities of service necessary for different service types in future 5G networks. In 5G, not only will there be different classes and qualities of service, but the network will also be "sliced" by service type, allowing for each slice to be managed by a carrier client using that service. This is called "network slicing." Because the classes of service and bandwidth profiles for different service types will vary so dramatically, it will be necessary to treat varying types of traffic differently as they traverse the RAN. For instance, traffic utilized for telemetry will require only small amounts of bandwidth and will have very little dependency on low latency. On the other hand, traffic utilized for video streaming will demand both large amounts of bandwidth and very low latency. Because Ethernet helps solve all of the aforementioned challenges, the market may push toward Ethernet-based fronthaul rather than CPRI in the longer term once the necessary standardization work is complete.

Finally, to help with centralization of control and virtualization of the network, all resources in the RAN, such as network, BBU, and compute, will migrate over time to SDN initially and then eventually NFV-based platforms. What used to be switching offices will then become data centers.

The Evolution of Ethernet to Accommodate Mobile Fronthaul

There are a number of efforts in the industry to evolve Ethernet for fronthaul. One of the first and most important is the 802.1 Time-Sensitive Networking Task Group within IEEE. Time-sensitive Ethernet is under study by this group and is a new standard for Ethernet that will allow time-sensitive traffic to be transported over an Ethernet network. Today, when Ethernet traffic is switched through a network, there is no accounting for its sensitivity to time delays or latency. Traffic is only categorized by its desired quality of service and its priority. Time-sensitive Ethernet will also add time sensitivity to the equation for prioritizing and processing Ethernet traffic, making it more deterministic as a transport protocol and therefore more feasible as a protocol for fronthaul networks. Time-sensitive Ethernet will allow for streams to be reserved through the network to transport time-sensitive traffic, and will also provide for special forwarding, queueing and policing of this traffic as it traverses the network. Finally, time-sensitive Ethernet will have a dependency on the existing 1588v2 standard for time synchronization of network resources in order to process this traffic.

Time-sensitive Ethernet will allow for digitized radio frequency (RF) traffic to be deterministically carried through the network as it is with CPRI, but there is also a need for standardization of

how RF gets digitized. There were previously a couple of working groups within IEEE looking at this challenge. The first was the 1904.3 standard, which standardized RF over Ethernet, and the second was the previously mentioned 1914.1 NGFI Working Group, which is specifically looking at next-generation fronthaul. More recently, the two groups have combined into the NGFI Working Group. It is likely that two standards will come out of this group – one for the migration of existing CPRI traffic to a next generation architecture, CPRI over Ethernet, and the second for carrying future fronthaul traffic directly over Ethernet without the need for CPRI at all. Ultimately, of course, it is desirable to get to Ethernet directly in order to utilize bandwidth more effectively, and this group will take advantage of time-sensitive Ethernet. It is through the work of this group that we may see Ethernet evolve to where it needs to be to support the future of fronthaul networks. The NGFI group has a lot of support from within the industry and may emerge over the alternative of continuing to evolve CPRI to higher and higher rates.

As previously mentioned, the Ethernet rates that will be supported within these new standards are not yet solidified, but there will likely be a focus on 10G and multiples of 10G in order to align with Ethernet rates commonly used in WAN transport.

The Evolution of Transport Equipment to Accommodate Next-generation Fronthaul

How Is It Done Today?

Today, fronthaul networks are planned and built utilizing CPRI. There are a number of ways to transport CPRI over fiber. Operators can always transport a single CPRI service per fiber or per fiber pair. This can be done by utilizing grey optics at the RRH and BBU sites and simply connecting each RRH directly to a BBU using a single fiber or fiber pair for each RRH to BBU connection.

This deployment option of using grey optics and dedicated fiber has the advantage that it does not require colored optics at the BBU and RRH, it costs less than putting active equipment out at the tower and at the BBU site, and it will scale regardless of the rate or protocol that is being utilized between the RRU and BBU in the future. The disadvantages of this architecture are that the fiber is not utilized efficiently and the cost of fiber is high. In addition, this architecture will not allow network operators to take advantage of aggregation of multiple fronthaul connections from multiple towers, and therefore operators cannot aggregate traffic with the same priority and carry it back to the BBU with the same QoS. Also, fault isolation for issues that can happen

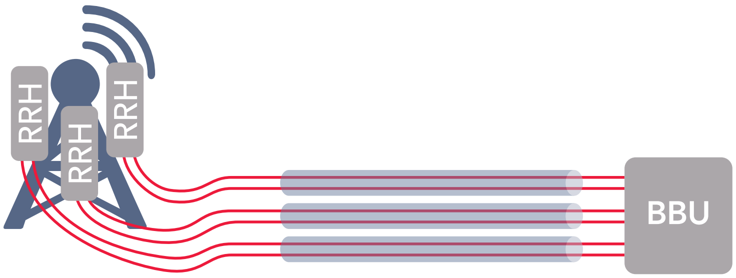

Figure 2: Fronthaul Utilizing One Fiber per CPRI

in the transport network, such as a fiber break or the effects of the weather, versus faults that are attributed to the radio or tower, is not provided.

Finally, the architecture is not sliceable in that all services that originate in a sector will be treated the same, at least until they hit the BBU. In general, this architecture has challenges today and does not allow the advantage that Ethernet brings once Ethernet becomes the fronthaul transport of choice. In order to do that, equipment will have to be added at a later date. Since this architecture does not have an add-drop capability, fiber has to be dedicated to each tower, pole, etc. even if they are deployed as a chain in a row or in a ring, further compounding the fiber cost and inefficiency.

Passive WDM Options

To alleviate the issues described above, many carriers are turning to utilizing wavelength-division multiplexing (WDM), allowing the multiplexing of multiple fronthaul connections over a single fiber or fiber pair.

In the figure below, passive WDM is used. Only passive mux/demux (MDU) filter units are placed at the RRH and BBU locations, carrying one fronthaul connection per wavelength over a fiber and thus allowing a fiber pair to carry many fronthaul connections. This architecture still has the advantages that it costs less than putting active equipment out at the tower and at the BBU site, and that it will scale regardless of the rate or protocol that is being utilized between the RRH and BBU in the future. It also fixes the fiber efficiency issue and adds add-drop capability, which is particularly important in dense urban areas where fronthaul is likely to be used and where cells are at the tops of buildings or other locations where other revenue-generating services could be delivered. Using WDM elevates the value of fiber. This architecture also allows for aggregation of multiple fronthaul connections at the wavelength level, and operators can use a band splitter in the RAN to aggregate multiple fronthaul connections/wavelengths from multiple fibers onto a single fiber.

The limitations of the passive option are that it still does not provide management capabilities beyond simple inventory management, such as fault isolation capabilities, and the deployed network is not sliceable. While it allows the advantage of aggregation at the wavelength level, it does not allow aggregation of services with similar priority from each cell over the same wavelength

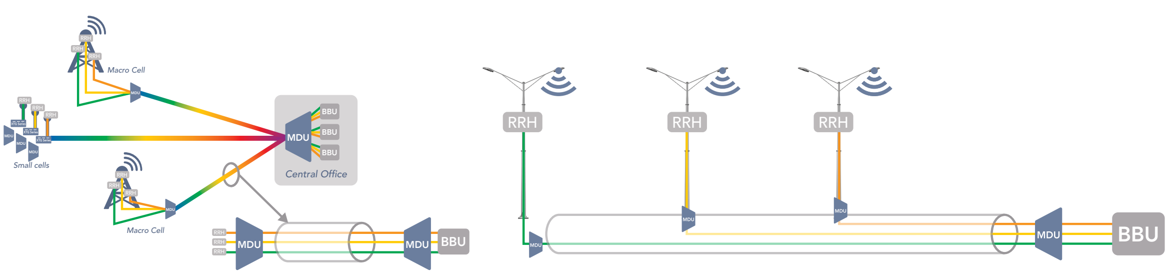

Figure 3: Fronthaul Using WDM (Passive) Using WDM to Chain Small Cells

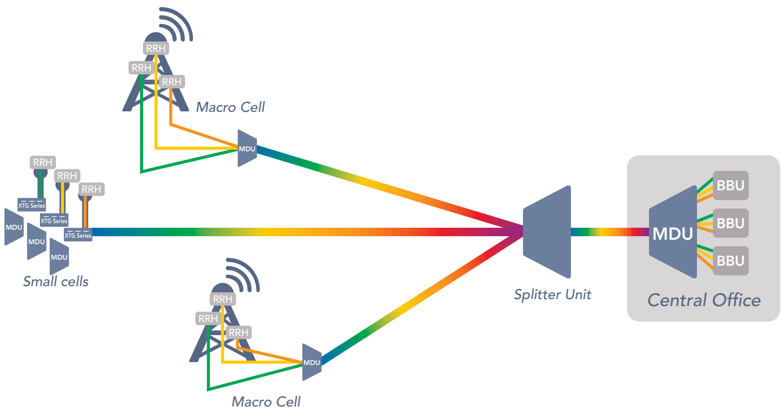

Figure 4: Passive Configuration Using a Band Splitter to Aggregate Fronthaul Connections/Wavelengths

and transportation of those services together and at the same QOS/priority back to the BBU site. Like the dark fiber scenario, in general, this architecture does not allow the advantage that Ethernet brings once Ethernet becomes the fronthaul transport of choice. Note also that this architecture requires the use of colored optics in the RRH and BBU. Not all radio solutions support colored optics in the RRH and BBU, which would eliminate passive WDM as an option.

Semi-passive WDM Options

To address the issue of a lack of management capabilities, some operators are evaluating and deploying semi-passive architectures that solve issues such as fault isolation between radio/tower failures and issues with the fiber plant. Two of the more pertinent semi-passive options - one using optical channel monitoring and link status monitoring (LSM), and the other using digital diagnostics monitoring (DDM), will be discussed here.

The two semi-passive options above both provide additional management capabilities over purely passive WDM, although the two are a little different in their approaches. Both retain passive infrastructure at the RRH site and transport one fronthaul connection per wavelength. The first option, shown on the left, adds an optical channel monitor (OCM) and fiber break detection card at the BBU site. These active components monitor per-wavelength power levels as they leave and arrive at the BBU site as well as detect fiber breaks. This allows for fault isolation between issues caused by the fiber plant and issues caused by the radio or failures at the tower

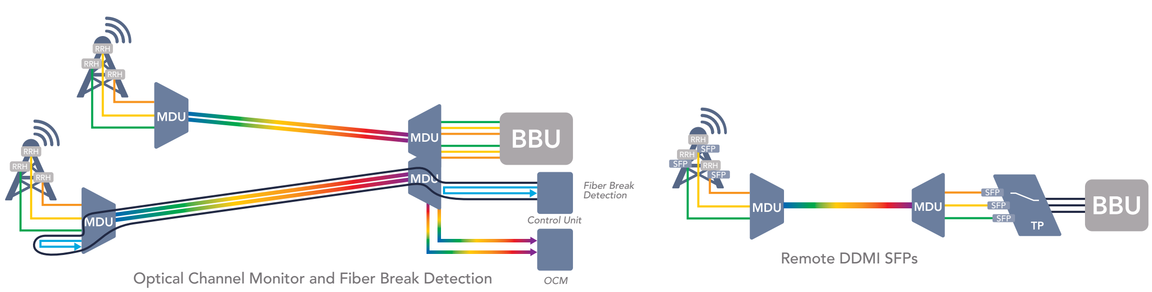

Figure 5: Semi-passive Options for Fronthaul: OCM/ LSM option and DDM option

itself. This option has all the advantages that passive WDM does, plus additional management capabilities. Since the OCM is a standard dense WDM (DWDM) or coarse WDM (CWDM) OCM, it will not have to be upgraded for future protocols and rates unless those rates or protocols are not aligned with the CWDM and DWDM grids. This is not likely to happen anytime soon since all rates up to 200 Gb/s can be supported in the DWDM C-band grid (50 GHz) and even higher rates could be supported over CWDM (100 GHz). The fiber break detection is independent of fronthaul traffic and will not be effected by the move to new rates and protocols in the future. However, this option still has all of the limitations that passive DWDM solutions do in terms of not supporting the advantages that Ethernet will ultimately bring to fronthaul. While it allows the advantage of aggregation at the wavelength level, it does not allow the aggregation of services with similar priority from each cell to transport those services together and at the same QOS/ priority back to the BBU site.

The second option, to the right, utilizes special small form-factor pluggable (SFP) optical modules that plug into the BBU and the RRH. In addition to providing the necessary WDM wavelength, these SFPs capture diagnostic information as the fronthaul connection leaves or enters the BBU and RRH sites. The remote SFPs at the RRH site are capable of transmitting their diagnostic information back to the BBU site on an out-of-band carrier that rides over the fronthaul traffic so no data connection network (DCN) is required out to the tower. This option shares the same benefits and drawbacks as the previous semi-passive option and adds the consideration that these special SFPs will have to be available for new rates, protocols, and form factors as technology moves forward, as opposed to for instance, just needing readily available standard Ethernet SFPs/C form-factor pluggables (CFPs)/quad SFPs (QSFPs).

Active WDM Options

Operators that need to address the challenges that neither passive nor semi-passive architectures can resolve can look at options based on active WDM. These are predominantly either transponder-based, supporting a single fronthaul signal per wavelength, or muxponder-based, supporting multiple lower-speed fronthaul signals over a higher-speed wavelength.

Currently available solutions on the market vary considerably in terms of performance and suitability in fronthaul networks. Active WDM solutions add electronics into the network, and as the CPRI protocol is extremely sensitive to both latency and synchronization, care should be taken to ensure that the selected solution supports tight fronthaul performance requirements. This subject and Infinera's full range of fronthaul options is explained further in Infinera's application note entitled "Enabling Cloud-RAN with Mobile Fronthaul" (see further reading section below).

Mobile operators are selecting active fronthaul solutions with extremely low latency due to the huge impact this has on network performance, but there are also other factors taken into consideration. These include the ability to support Ethernet and CPRI over muxponders and the support of

multiple options for building practices. Whereas typical WDM deployments are in controlled environments such as central offices or customer premises, fronthaul requires deployments at cell towers and therefore mandates a range of deployment building practices including standard plug-in units for deployment in standard telecoms environments, hardened fanless pizza boxes for more rigorous environments such as street or pole cabinets, or fully hardened "clamshell" options for outdoor environments such as direct mounting on poles or towers.

Some WDM vendors have added additional fronthaul-specific features such as ring protection with real-time latency monitoring and compensation. This feature may be important in circumstances where network operators wish to take advantage of dual fiber routes to protect fronthaul links, and the BBU/RRH solution requires constant latency, even in protection switch scenarios.

Service providers often choose to use these active options for a number of reasons. First, the active option does not require a colored optic in the RRH or BBU, which has two advantages: it enables operators to support radio solutions that do not support colored optics today and it supports the upgrade of existing RRH deployments that already have grey optics from the tower without having to replace the RRH optics. This is particularly useful when operators are upgrading existing non-fronthaul architectures to fronthaul architectures.

The other advantage that some active options have is that they may satisfy the one challenge on the list of drivers that the passive and semi-passive solutions did not: in the future, they could have the ability to be upgraded to play more ubiquitously in an all Ethernet NGFI environment. This means that active options could take advantage of some of the strengths that Ethernet brings to the RAN fronthaul environment: the ability to mux, slice, and prioritize traffic with QOS, which are all requirements for the migration to 5G.

Active options also provide the highest level of management and fault isolation. Because there is a transponder or muxponder on both ends of a fronthaul connection, very detailed statistics

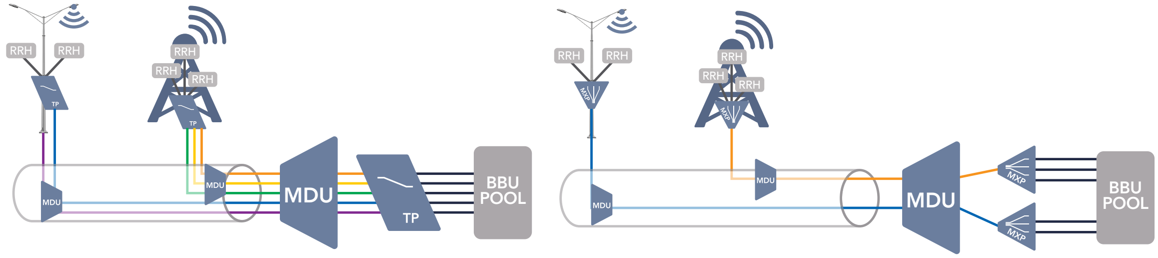

Figure 6: Active Configurations Adding and Dropping Fronthaul Connections to Sequential Cells

showing what is going on between the two devices are provided. These statistics typically include:

- Performance monitoring such as errored seconds (ES) and severely errored seconds (SES)
- Digital diagnostics monitoring interface (DDMI) statistics such as input power, output power and laser current
- Loss of signal isolated to the ends of the fiber plant

Active options provide good demarcation points for wholesale fronthaul services. Active options also utilize WDM over the fiber plant, solving all the same issues that the passive and semi-passive options solve, such as fiber utilization.

Migrating Active Fronthaul to NGFI

There are two likely capabilities that will come out of the NGFI standards with regard to Ethernet and fronthaul. One is the ability to carry CPRI over Ethernet, which will allow for the migration of existing fronthaul installations to an RF-over-Ethernet NGFI environment. The other important development expected out of the IEEE NGFI group is running RF directly over Ethernet without the need for CPRI. Some CPRI fronthaul hardware is built on field-programmable gate array (FPGA) technology that will allow future software upgrades to support new protocols and mapping as the standards evolve and as line rates increase in fronthaul networks.

In the above example, a fanless pizza box node containing a CPRI transponder/muxponder with an integrated WDM filter was initially installed to take CPRI signals from the RRHs and transport them as CPRI back to the BBU pool. Then it may be later upgraded to provide CPRI over Ethernet on the line interface rather than native CPRI. This allows the network to support a next-generation fronthaul architecture network even if the radio technology in the towers is not upgraded. Notice in the example above that there are multiple NGFI Time Sensitive Ethernet

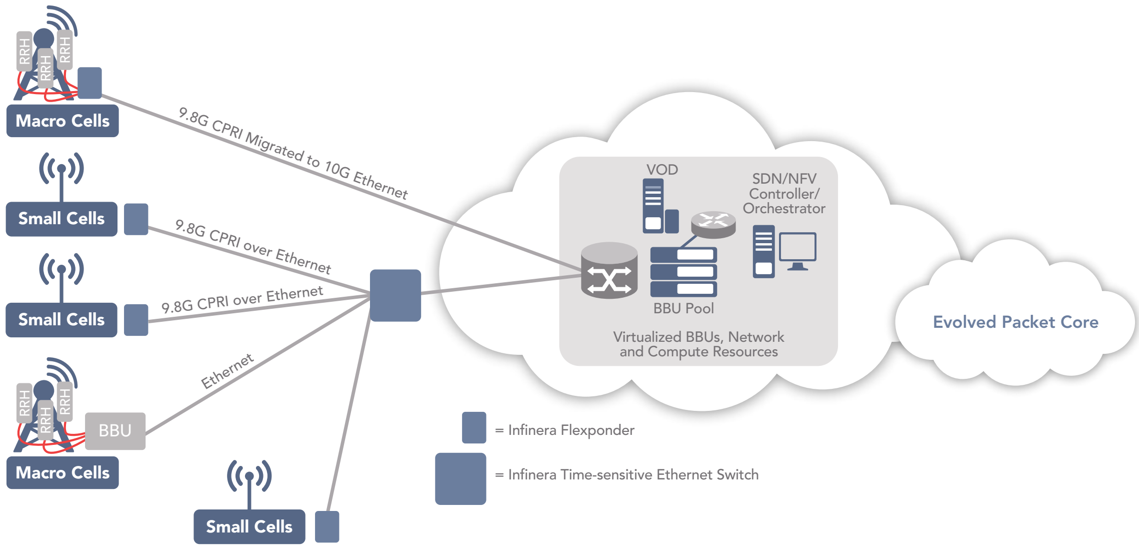

Figure 7: Hybrid Legacy Fronthaul and NGFI Network

Switches that aggregate NGFI streams from both converted fronthaul sites and NGFI sites. At the switch sites, fronthaul data streams that have similar priority and COS can be grouped together and carried back to the BBU pool. This allows for prioritization of individual traffic types as is required by 5G whether the fronthaul traffic is coming from NGFI or legacy CPRI sites.

Finally, no 5G paper would be complete without a discussion of the SDN and NFV requirements of the fronthaul network. Networks that are deployed today should be SDN capable and roadmaps should include Netconf interfaces to allow operators to interface directly to a standards-based SDN controller/orchestrator. This will allow networks that are being deployed today to fit into the virtualized environment being proposed for next-generation RANs.

Conclusion

Mobile fronthaul is still in an early evolutionary stage. At this point in the evolution of fronthaul there are at least two potential paths that next-generation fronthaul networks may go down: continued use of CPRI or the new NGFI approach. As of today, NGFI appears to be gaining traction within the industry, which melds well with the goals of 5G networks.

The use of Ethernet within NGFI will provide ubiquity of protocol and operations throughout the whole mobile transport network, across both fronthaul and backhaul, and will also provide a means by which different traffic types can be provided class of service and quality of service. The use of Ethernet combined with SDN additionally enables a means by which the network can be sliced and the control of different slices can be given to different end users.

Most service providers may take advantage of WDM in the fronthaul network in order to better utilize fiber plant and to allow for more diverse configurations. Passive, semi-passive, and active architectures for mobile fronthaul all have their advantages and limitations, but no matter what choice an operator makes, Infinera's diverse portfolio and market-unique technologies such as low-latency multiplexing, programmability and the upgradable active flexponder will provide them alternatives that will make the move to 5G and beyond simpler and more cost-effective.

Further reading:
China Mobile C-RAN White Paper 3.0
Infinera Application Note: Enabling Cloud-RAN with Mobile Fronthaul

Infinera

Infinera Corporation
140 Caspian Court
Sunnyvale, CA 94089 USA
Telephone: +1 408 572 5200
Fax: +1 408 572 5454
www.infinera.com

Have a question about Infinera's products or services?
Please **contact us**.

Specifications subject to change without notice.

Document Number: 0187-WP-RevA-0519