Centralized RAN and Fronthaul

Kevin Murphy
Solutions Manager
kevin.murphy@ericsson.com

1 Introduction

A Centralized Radio Access Network (C-RAN) is an architecture where the Digital Units (or Baseband Units) are placed at a centralized location, and the Remote Radio Units (RRU's) are placed at distances up to several kilometers away from the baseband site. The connection from a baseband unit to a remote radio unit is typically a fiber facility which is referred to as "fronthaul" to differentiate it from backhaul which connects the baseband unit to the network. This document will discuss the background and origin of the C-RAN architecture and the cost savings that can be achieved using this approach.

2 Baseband Units and Radio Units

Traditionally, radio base stations (RBS) consisted of two functional units, the baseband unit and the radio unit. The baseband unit connects to the network via the backhaul, which is typically an Ethernet circuit. The Radio Unit transmits and receives the carrier signal that is transmitted over the air to the end user equipment (UE).

The function of the baseband unit is to translate the data stream coming from the network into a form that is suitable for transmission over the air, and in the other direction, to take the data stream from the radio units, and transform it into a form suitable for transport back to the network.

Cellular radios use a form of carrier modulation called QAM, Quadrature Amplitude Modulation which modulates, not just the amplitude of the carrier, but also the phase. This is achieved by creating two waveforms called the ***in phase*** waveform (I) and the ***quadrature*** waveform (Q). The quadrature waveform is out of phase with the in phase waveform by one quarter of a wavelength. Modulating the amplitude of both of these waveforms and then adding them together results is a signal that has both amplitude and phase variations. This waveform is used to modulate the carrier that is transmitted to the user equipment.

In traditional base stations, coaxial RF cables were used to carry the RF signal up the tower to the antennas. These cables tended to be bulky and expensive and also introduced a lot of RF power loss into the signal path.

To overcome this problem, the newer base stations digitize the baseband signal and transmit it over fiber optic cable to the radio unit mounted up on the tower. This avoids the long runs of co-ax cable and the associated RF losses. This digitized baseband data stream is known as **IQ data** and it is encapsulated in a protocol known as ***CPRI*** (Common Public Radio Interface).

One advantage of using CPRI over fiber is that it enables radios to be placed quite far away from the baseband unit (up to several miles). This turns out to be a useful feature as will be discussed in the next section, but first we will take a more in- depth look at the CPRI protocol.

3 CPRI Basics

The CPRI specification can be found at www.cpri.info . The parties cooperating to define the CPRI specification are Ericsson AB, Huawei Technologies Co. Ltd, NEC Corporation, Alcatel Lucent and Nokia Networks.

The CPRI protocol consists of a User Plane that carries the IQ data, a Control Plane that provides control signalling to and from the remote radio, and a Sync Plane, that keeps the remote radio and baseband units in sync.

The standardization was done to enable chipset vendors to create components that can be used by any radio equipment vendor. However, the standard is not completely open. Certain aspects of the signalling are vendor specific, so mixing radios and baseband units from different vendors is not guaranteed to work.

The CPRI protocol is a full duplex serial data stream. In other words there is a transmit fiber and a receive fiber. It is also a synchronous protocol, which means the baseband must establish sync with the remote radio before communication can be established.

One CPRI stream can support several IQ data flows to several different radios. Operators typically use multiple radios at each cell site (with different carrier frequencies), and the radios are usually divided into multiple sectors. So a three band, three sector cell site, would have nine radios each with a separate IQ data flow.

Within the CPRI protocol there are multiple data blocks called antenna carriers (AxC). One AxC carries the IQ data necessary to serve one carrier at one independent antenna element. This is why there are several different CPRI line rates (Option 1 to Option 9) that range from 614Mbps to 12Gbps. The most commonly used CPRI rates are Option 3 which is 2.5Gbps and Option 7 which is 10Gbps. Perhaps one of the reasons for this is that a standard SFP can handle up to 2.5 Gbps whereas the more expensive SFP+ can handle up to 10 Gbps.

CPRI has a very high data rate because it is a digital representation of a very high frequency waveform. An LTE baseband signal with a 20MHz channel bandwidth has to be sampled 30.72 million times per second to be accurately digitized. The stream is then duplicated for 2x2 MIMO resulting in a CPRI rate of 2.4576 Gbps which is the Option 3 CPRI rate. A 20Mhz LTE channel can deliver up to 150 Mbps in the downlink direction and requires a CPRI data rate of 2.5 Gbps. So CPRI is inherently quite inefficient from a data transport perspective.

These kinds of data rates usually require fiber optic transmission, although microwave can be used for some of the lower CPRI rates.

4 Cell Densification and Coordination

One of the primary reasons to use a C-RAN architecture with CPRI fronthaul is to enable tight coordination between cells. This becomes especially critical when more cells are introduced into the existing cell structure in a process known as densification, or when small cells are deployed in the same frequency band as the macro cells.

When you consider the tiny antenna in a smart phone, it is quite impressive that this weak signal can make it through building walls and windows and all the way to the nearest cell tower. As the signal loses power on the way to the cell tower, the data rate must be stepped down to reduce data errors. Increasing the power at the cell site end would only solve the problem in one direction, but would also create more interference in the adjacent cells. So we say that coverage and capacity in a cellular network is 'uplink constrained'.

The only way to significantly increase user data throughput in both directions is to move the cell site closer to the end user. This is achieved by densification (placing new cell sites between the existing ones) and by deploying small cells in coverage holes or hot spots that have a lot of data demand.

Traditionally, a macro cell site has an Ethernet backhaul facility back to the network. CPRI provides a way of extending remote radios from an existing macro site, out into the cell's coverage area to address coverage holes, indoor locations or areas with heavy demand. This CPRI circuit is referred to as "fronthaul" to differentiate it from the backhaul to the network.

Using fronthaul has a number of benefits over connecting the small cells directly back to the network using Ethernet. These benefits are especially significant when introducing LTE Advanced into the network. One of the most important advantages is coordination.

When small cells are deployed in a macro coverage area, there will be areas, where the signal strength of the macro and the signal strength of the small cell are about equal (see Figure 1). In these areas, it is possible to have some users communicating with the macro cell, while other users are communicating with the small cell. In a densely populated area with lots of simultaneous users, like a campus or a train station, this can result in substantial interference. It is in these areas that we get the most benefit from using an LTE Advanced feature called CoMP (Coordinated Multipoint).

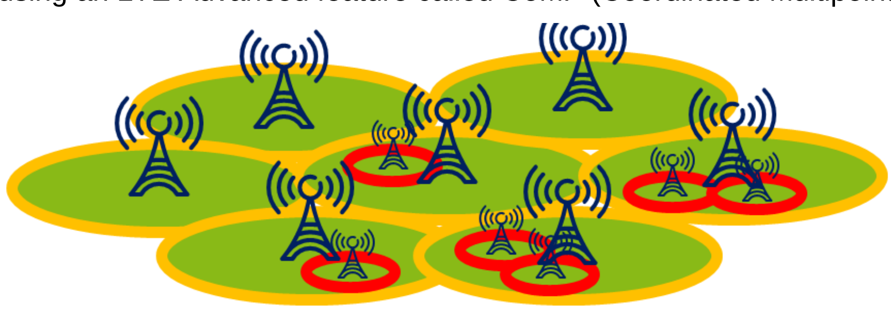

Figure 1: Macro/Small Cell Interference

CoMP works by using peer to peer communications between the cells over what is known as the X2 interface. This allows the cells to coordinate the communication with end users so that interference is minimized. One approach is to use scheduling where cells take turns signalling to users so as not to “talk over” each other.

A “Coordination Set” is a set of cells that may interfere with each other and hence utilize the CoMP feature to implement coordination. Tight coordination using CoMP enables much more efficient use of the access spectrum in the cell edge, increasing both the capacity and the speed. Considering the billions of dollars that operators invest in spectrum, increasing spectral efficiency, and hence end user experience, is a worthwhile investment.

5 CPRI latency

For a C-RAN architecture to perform well it is important that the overall CPRI transport network is designed to keep the latency as low as possible. A guideline is to design the network so that the one way delay is below 75 microseconds. Light travels approximately 1 kilometer in 5 microseconds, so 15 km of fiber would produce a one way latency of 75 microseconds. Transport equipment that injects extra data into the stream will add additional delay.

The baseband unit is processing many control signals to and from the radio units and end user equipment. If the baseband and radio unit are co-located then there is little delay in sending signals between the two. If the radio is placed 5 kilometers away, the baseband must wait at least 25 microseconds for the signal to get out to the radio, and 25 microseconds for the response signal to come back. So every transaction will take 50 microseconds longer than it would take if the baseband unit and radio unit were co-located. The net effect is that the efficiency of the baseband unit is reduced because tasks now take longer to execute. Heavy traffic loads, with a large amount of signalling over a network with high latency may result in performance degredation.

When the the coordination feature CoMP is activated it requires very fast signalling between radios within a Coordiantion Set. If each signal has to travel a long distance back to the baseband unit, the timing limits may be exceeded.

Voice over LTE (VoLTE) is another feature that requires low latency and the network must be planned to accomodate the future demands that will be placed on the network as these and other new features proliferate.

6 CPRI Transport

A physical CPRI link will typically be no more than 15km end to end (for latency reasons) and preferably each CPRI stream will have a data rate of 10 Gbps. As discussed above, one cell site typically requires several AxC’s, one for each radio, so using the higher order Option 7 CPRI (10 Gbps) is more efficient.

The best facility for CPRI transport is dark fiber, because encapsulating CPRI into another protocol introduces too much latency and jitter. Since the monthly lease of a dark fiber pair can be up to $1,500 per mile in a large city, the operator needs to put as many CPRI streams as possible on a single fiber pair. As the number of streams on a single pair increases it becomes more important to have geographical redundancy and ring protection in case of a fiber cut.

The two best options for CPRI transport are Coarse Wave Division Multiplexing (CWDM) and Dense Wave Division Multiplexing (DWDM).

CWDM systems are passive and allow multiple wavelengths to be transported on a single fiber by using matched “colored” SFP’s in the baseband and remote radio end. Up to 16 CPRI wavelengths can be combined on the same fiber using a passive filter. The downside of CWDM is that it doesn’t support ring protection and scaling to hundreds of CPRI links becomes an operational challenge. Also, since CWDM is a passive technology there are no active devices that can perform supervision functions such as link surveillance and alarm generation.

Another option is DWDM. However, traditional DWDM systems were designed to address a different kind of transport problem. Unlike CPRI links that are quite short, DWDM systems are designed to transport wavelengths over long distances. To achieve this, a G.709 “digital wrapper” is inserted into the bit stream to provide FEC (Forward Error Correction) and OAM (Operations Administration and Management) functions, and optical amplifiers are used to boost the signal. This can add up to 40 microseconds of latency to the link, and since CPRI already has its own built in FEC and control plane, adding this extra content is to some extent redundant. Some P-OTN systems can disable the FEC function but still create about 8 micro seconds of latency which is equivalent to adding a mile of fiber to the fiber route.

Systems that are custom made for CPRI transport, like the Ericsson PAU 6000, minimize latency by not inserting anything into the CPRI data stream and using an external dedicated wavelength as the management channel to support supervision and ring protection. This keeps added latency under 0.5 microseconds and allows the system to be completely agnostic to the protocol that is being transported. The transported signal can be any vendor’s CPRI, Ethernet (GigE or 10 GigE) or any other optical signal. Another advantage of short CPRI links is that they don’t need expensive optical amplifiers, and they can use less expensive thin film filters instead of Arrayed Wave Guides (AWG).

A key feature for CPRI DWDM systems is “drop and continue” where a subset of wavelengths can be dropped at a site and the rest of the wavelengths continue on to the next site. This capability is useful because a single fiber pair can transport enough CPRI streams for many cell sites. The cumulative loss in this type of chain is quite high when you consider the fiber losses plus the add drop multiplexor losses. But since the CPRI link will never be longer than 15km, four sites can easily be served without optical amplifiers.

Creating a chain or a ring that drops 6 to 12 CPRI at each site is a very economical way to serve many sites using only one fiber pair.

So in summary, the key to designing a fiber optic CPRI distribution network is to consider that every microsecond of latency reduces the efficiency of the baseband. Distances to the remote radios will be fixed and the operator has little choice when it comes to the fiber route. So it is critical that the transport network uses as few fiber pairs as possible (preferably only one) and the added latency is as low as possible.

CPRI over microwave is also technically possible but only in the higher 70 and 80 GHz frequency bands where there spectrum is plentiful. High frequency microwave does not have very long reach, but considering the cell radius of a macro in a dense urban environment is often less the one kilometer, the CPRI signal to a small cell does not have to go that far. A 2.5 Gbps CPRI link over microwave would typically reach about 500m maximum. This is useful in campus or stadium environments or to access areas were fiber is difficult to deploy. Latency over microwave is very low and these radios can use small antennas that are less than 12 inches in diameter.

7 Total Cost of Ownership

Using CPRI fronthaul as opposed to Ethernet backhaul has a number of cost benefits which may, in part, offset the high cost of leasing or deploying dark fiber.

The following are some of the benefits compared to Ethernet backhaul:

- CPRI provides its own synchronization. No need for GPS at the site
- No baseband unit at the site. Lower power consumption and footprint
- No cell site router required
- Tight coordination possible with the macro layer
- Dark fiber, leased per mile instead of per megabit, can be less expensive.

Lowering the equipment footprint and power requirements at the cell site can be a major savings especially if the large cabinets that house the baseband units and rectifiers are no longer needed.

The cost of leasing a 100 Mbps facility from a carrier Ethernet provider is currently about $1200 per month, and 500 Mbps is about $2,000, so as the data in the network grows so does the backhaul cost. Dark fiber is leased per mile and extra wavelengths can be added without affecting the monthly recurring charge. So as the cell density increases and the need for coordination becomes more critical, there comes a point where dark fiber using CPRI fronthaul is more economical.

Centralizing the baseband units may also provide what is known as a pooling gain. This means that baseband units whose CPRI ports were only partly filled at a cell site can be fully utilized in a centralized baseband site. Therefore, the same number of cells can be supported by a fewer number of baseband units.

8 Scaling the C-RAN Network

The next issue to be considered is the optimum number of cell sites that should be connect to one baseband site. The maximum size of a baseband site is constrained, both by CPRI reach (the 15km limit), and the single point of failure risk. There is also the cost associated with acquiring and maintaining the baseband site.

To get a rough idea of how big to scale a baseband site, here is an example of a possible configuration.

If we assume that macro sites are roughly 1 kilometer apart, we could envisage 16 sites in a 3 x 3km square as shown below. In this configuration the fiber distance to every site is 12km or less, whether we go clockwise or anti-clockwise around the ring.

Assuming very little latency is added by the transport equipment, we still have 3km of latency to spare for irregularities in the fiber route.

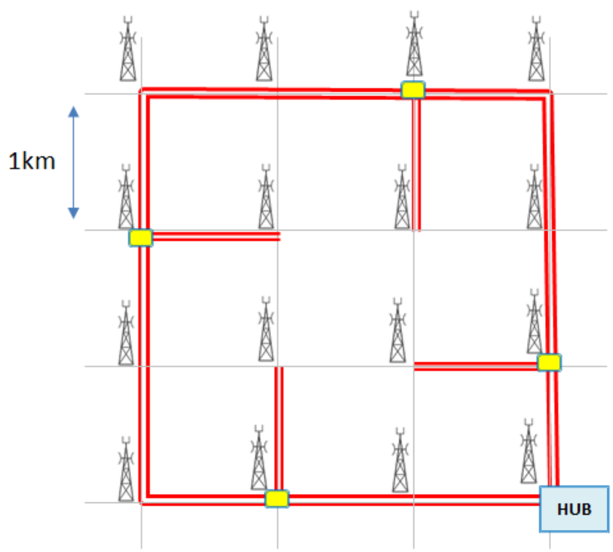

Figure 2: Baseband Site Maximum Coverage Area

If we had four of these rings radiating out from the baseband hub site, we could reach 64 sites (4 x 16). So 64 sites connected to a baseband site is an approximate upper limit from a latency point of view.

9 Summary

In summary, the following are the main considerations when building a centralized RAN:

- Centralized RAN is best suited for a dense cellular network with many users

- A baseband site could, in theory, serve up to 64 cell sites in a 6 x 6km area, but in practice it would probably be less than that
- Leasing dark fiber becomes cost effective when the distances are short and the amount of data (e.g. CPRI) is high. Pay by the mile not by the megabit
- One of C-RAN's biggest benefits is the ability to introduce coordination to improve RF performance
- The CPRI transport network should be designed for as low latency as possible
- Cost savings may be realized at the cell site due to smaller footprint, lower power consumption, and potentially lower site lease costs.

For more information, please contact:

Kevin Murphy

Solutions Manager

kevin.murphy@ericsson.com

