

Charting the path to RAN virtualization: C-RAN, fronthaul and HetNets

By Monica Paolini, Senza Fili



In collaboration with



We thank these companies for sponsoring this report:



Table of contents

I. Market report 3

1. Introduction. Where does RAN virtualization fit in? 4
2. The emergence of a new RAN topology. From distributed RAN to C-RAN and virtualized RAN 6
3. From local to advanced C-RAN. Multiple types of C-RAN to target different environments 7
4. Picking the right C-RAN architecture. Assessment of multiple tradeoffs crucial to successful C-RAN deployments 9
5. The path to virtualization. The role of the RAN within the NFV platform 10
6. Fronthaul, the great enabler of C-RAN and vRAN architectures. A bottleneck or a requirement to meet? 11
7. Fronthaul beyond CPRI. CPRI alternatives attract operators' attention 13
8. Tradeoffs can ease the way to RAN virtualization. Functional split options to facilitate or enable the transition to C-RAN 14
9. Drivers to a virtualized RAN: cost. Cost dynamics in the short term and in the long term 17
10. Drivers to a virtualized RAN: mobility and interference management. Toward a more efficient use of network resources in densified networks 19
11. Drivers to a virtualized RAN: statistical multiplexing. Leveraging uneven traffic distribution to optimize resource utilization 21

12. Implications: The death of the cell? Moving baseband processing to the cloud erases RAN boundaries 22
13. Implications: A blurring demarcation between core and RAN. The importance of physical location in a virtualized world 24
14. Summary. 25

II. Vendor profiles and interviews 26

Airvana 27
ASOCS 34
Fujitsu 41
Radisys 47
SOLiD 54

III. Operator interviews 61

Oi 62
Orange 65
SoftBank 68
Telefonica 70

References 75

Glossary 76

I. Market report

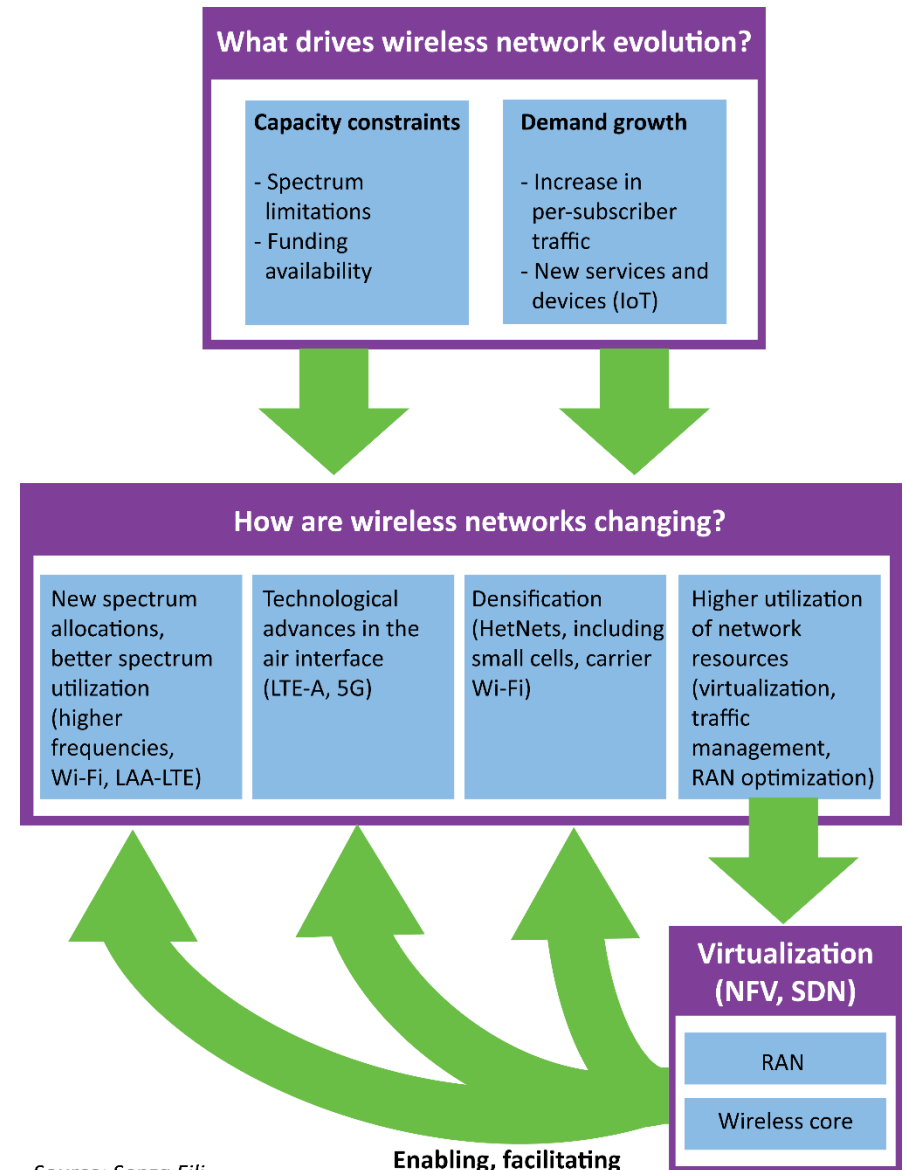
1. Introduction.

Where does RAN virtualization fit in?

RAN virtualization is gaining wider acceptance among mobile operators and vendors as part of the wider trend to virtualization in wireless networks that involves NFV and SDN. Yet, some still see C-RAN as a niche solution that will remain confined to specific environments (e.g., stadiums, malls or corporate buildings) or to markets where fiber is plentiful and cheap. Others believe that the target of a fully virtualized RAN is unattainable because of extremely tight processing requirements that preclude the use of VMs running on general purpose hardware.

There is still much uncertainty about the extent to which virtualized RAN topologies will spread among different countries and within operators' footprints, about the timeline for the transition from legacy RAN to virtualized RAN, and about the RAN functions that general purpose hardware can perform reliably and cost effectively. But the move to RAN virtualization has already started and will expand in the future – with wider commercial deployments and deeper functional virtualization.

Cost savings are frequently cited as the primary – or at least the initial – reason to move to virtualized RAN, and more generally to virtualized mobile networks. However, improvements in performance, ease of deployment and ease of maintenance will have a much stronger impact than short-term cost reduction. For many operators, especially if they do not have greenfield deployments, the transition costs of the move to vRAN may be higher than the ongoing cost savings. As is the case for the virtualization of the wireless core, the main motivation for RAN virtualization comes from a strategic, long-term evolution plan, in which short-term cost savings are often not the main goal. In the long term, virtualization should bring substantial cost savings at different levels, but it is extremely difficult, or impossible, to qualify and quantify them accurately now, because they will depend on the type of virtualization that the operator chooses, and its timeline and implementation.



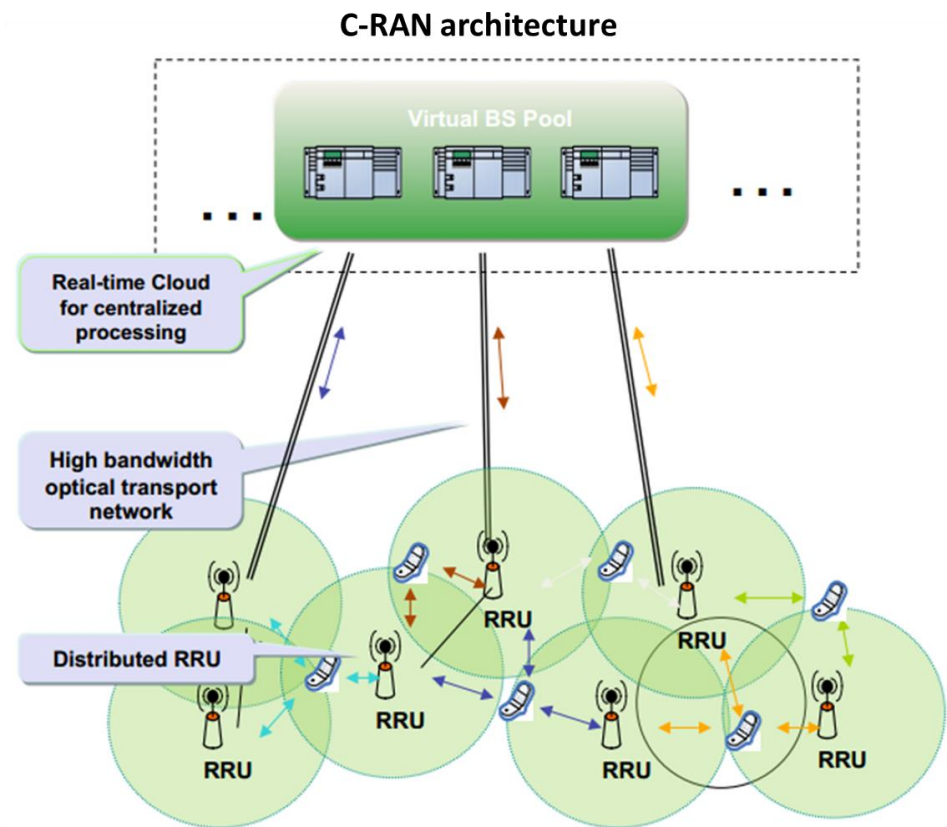
Source: Senza Fili

The trend toward network virtualization – and RAN virtualization, specifically – is driven by mobile networks’ need to accommodate fast traffic growth, driven by heavier usage by individual subscribers and by the spread of wireless connectivity to new classes of IoT devices. Network expansion is being constrained, though, by a combination of the limited spectrum and limited funding available. Operators are under intense pressure to deliver more services at a lower cost – i.e., minimize their per-bit costs – because subscribers are demanding faster access but strongly try to avoid higher monthly fees.

Within this framework, operators can – and typically do – follow multiple paths to increase network efficiency and performance. Traditionally, operators have focused on acquiring more spectrum assets and on ways to increase spectrum utilization – as the recent interest in higher frequencies and license-exempt bands attests.

But even though many spectrum assets are underutilized, the amount of spectrum available is necessarily limited and spectrum rights are expensive, so operators need to do more. Technological advances in the air interface – e.g., LTE-A or 5G – also help, but the performance enhancements they bring are not sufficient to meet the increase in traffic demand. Densification will provide a much needed increase in capacity concentration in areas of high traffic, but it requires substantial investment, plus careful interference and traffic management, to deliver the promised benefits. Virtualization and the use of advanced traffic management tools are relatively new approaches that operators are adopting for more efficient use of their network resources, crucially reinforcing their efforts in the spectrum, air interface and densification areas. For instance, RAN virtualization can facilitate the integration of unlicensed spectrum, the deployment of LTE-A tools like CoMP and CA, and the management of interference in HetNet deployments.

As part of the push for higher resource utilization, RAN virtualization acts as an enabler and facilitator within the evolution of mobile networks, rather than a new functionality. But this is a crucial role that – in conjunction with virtualization of the core – has the potential to fundamentally change the way mobile operators think of, plan and run their networks.



Source: China Mobile

2. The emergence of a new RAN topology.

From distributed RAN to C-RAN and virtualized RAN

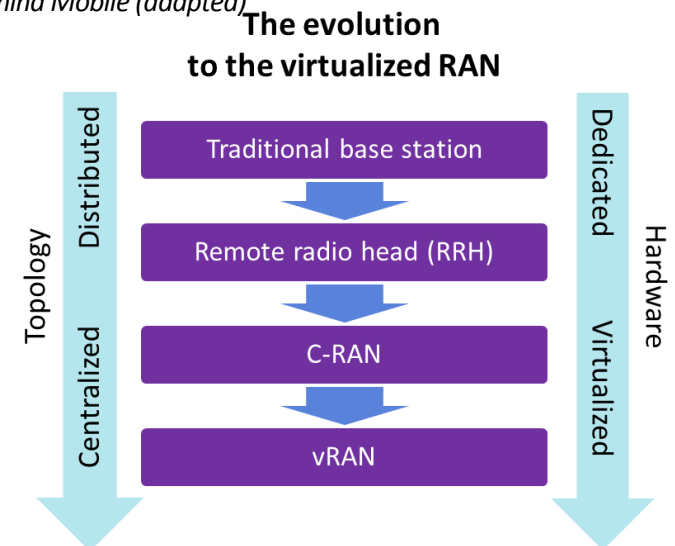
Long before the emergence of the C-RAN concept, the traditional RAN had started to evolve with the introduction of the RRH, which separated the radio and baseband processing elements that used to be in a single piece of equipment, the traditional base station. From this start, C-RAN and, eventually, vRAN architectures can be seen as the natural next steps, with the baseband processing progressively moved farther away in C-RAN, to a more centralized location, and in vRAN, becoming abstracted from a specific hardware implementation.

While this provides a first-pass description of the evolution from a distributed to a virtualized RAN, it combines multiple aspects of the evolution that coexist or compete with each other, and make this change much more than a cost-saving process or an alternative architecture. As such, the transition to vRAN is not going to follow a simple, well-demarcated path, but requires a complex evaluation of multiple tradeoffs.

The concept of C-RAN assumes different connotations among different industry players. While C-RAN architectures all assume remote baseband processing, the processing itself does not necessarily have to be pooled, and it may include only partial functionality, depending on an operator's goals. China Mobile, one of the strongest supporters of C-RAN, has identified four C-RAN concepts – four Cs: centralized, collaborative, cloud, clean – that define different flavors of C-RAN. In this report we take C-RAN to refer to cloud RAN, but assume that the other aspects – centralized, collaborative and clean – are also included in the concept. At the same time, we use a broad concept of cloud, to include different C-RAN types.

China Mobile's four concepts of C-RAN
Centralized RAN – Centralized processing resource pool that can support 10–10,000 cells.
Collaborative RAN – Multicell joint scheduling and processing to improve interference management.
Cloud RAN – Real-time cloud, leveraging open IT platform, consolidating the processing resource into a cloud. Enables flexible and multistandard operation and migration.
Clean RAN – Lower power consumption, lower opex, and fast system rollout.

Source: China Mobile (adapted)



3. From local to advanced C-RAN.

Multiple types of C-RAN to target different environments

The table below describes some C-RAN types that have been proposed. It is not an exhaustive list, but serves to convey the wide range of options available to operators as they plan to deploy C-RAN in different environments.

Local C-RAN is emerging as the initial stage in the evolution to C-RAN and vRAN. It fits smoothly into existing legacy networks, and it addresses the urgent need to manage the interference between the macro-cell and small-cell layers in HetNets, with a robust implementation of joint transmission with CoMP. Because of its limited scope, local C-RAN can be deployed in a targeted fashion – only where needed, when needed. Because it involves only a small number of small cells, all located at a short distance from the macro cell, fronthaul requirements can be met both by wireless and wireline fronthaul, facilitating deployments in areas where fiber is not available or is too expensive. Not only does local C-RAN offer a streamlined initial path that paves the way to more centralized architectures, it is also a crucial enabler for cost-effective small-cell

deployments. As a result, C-RAN and small cells mutually benefit each other: C-RAN improves HetNet performance, and small cells strengthen the business case for C-RAN deployment.

Another increasingly popular form of C-RAN that easily integrates into existing networks is micro C-RAN, which is deployed in indoor or tightly circumscribed environments. DAS deployments can be considered an instance of micro C-RAN, because baseband processing is separate from the antennas. DAS architecture, however, has been used for a long time, and can be seen as an ante litteram incarnation of C-RAN. An increasing number of small-cell solutions and hybrid small-cell and DAS architectures are also forms of micro C-RAN, because they keep RF and baseband components separate. Micro C-RAN's popularity comes from its ability to improve capacity and coverage in contained areas of very dense demand (e.g., stadiums) and to target specific areas, much as local C-RAN does.


Types of C-RAN

Types of C-RAN	
Local C-RAN	Located within a macro-cell site, local C-RAN includes (typically outdoor) small cells installed within the coverage area of the macro cell. This is an increasingly popular choice for legacy networks that need to support a growing number of small cells.
Micro C-RAN	Micro C-RAN can be private (e.g., enterprise) or public C-RAN (e.g., stadiums and other public venues). DASs are instances of micro C-RAN. Indoor deployments are mostly micro C-RANs.
Metro or wide C-RAN	Metro/wide C-RAN is deployed in high-traffic areas, with multilayer HetNet architectures that include outdoor small and macro cells.
Centralized or hyper C-RAN	Centralized/hyper C-RAN covers large metropolitan areas (up to 40 km radius) and is well suited to greenfield deployments.
Advanced C-RAN	Supported by DOCOMO, it covers a hotspot area and leverages CA to use different bands for macro cells (e.g., 2.1 GHz) and small cells (e.g., 3.5 GHz), with the macro layer mostly responsible for control plane functions.

Metro, or wide, C-RAN and centralized, or hyper, C-RAN are better suited to greenfield deployments and to networks in which large areas are rapidly moving to C-RAN. Their BBUs are typically hosted in a data center, with intra-BBU or inter-BBU pooling. The size of the BBU and the coverage area are much larger than in the local C-RAN and micro C-RAN cases, and they can include small cells, but they often don't. Cost savings as well as deployment and operational benefits are often the predominant drivers to adoption, at least during the initial stage.


Finally, the advanced C-RAN proposed by NTT DOCOMO is designed for high-traffic hotspots in which small cells use a different, typically higher, frequency from the macro cell. Data plane traffic is split among macro and small cells, but the control plane traffic is mostly managed by the macro cell. In the long term, this type of C-RAN configuration is likely to become increasingly common, because it allows operators to use spectrum more efficiently and remove the interference between the macro- and small-cell levels. More specifically, advanced C-RAN enables mobile operators to integrate transmission more effectively in higher bands, which can provide a high capacity density within a small coverage area within their networks (and hence it is well suited for small cells). Because such C-RAN configurations require the support of a new frequency in the handset and, in many cases, the acquisition of new spectrum, operators will need some time to plan for deployments. But when they do, they are likely to move forward fast with extensive rollouts.

Local C-RAN



Source: Orange

DOCOMO's advanced C-RAN architecture



Source: DOCOMO

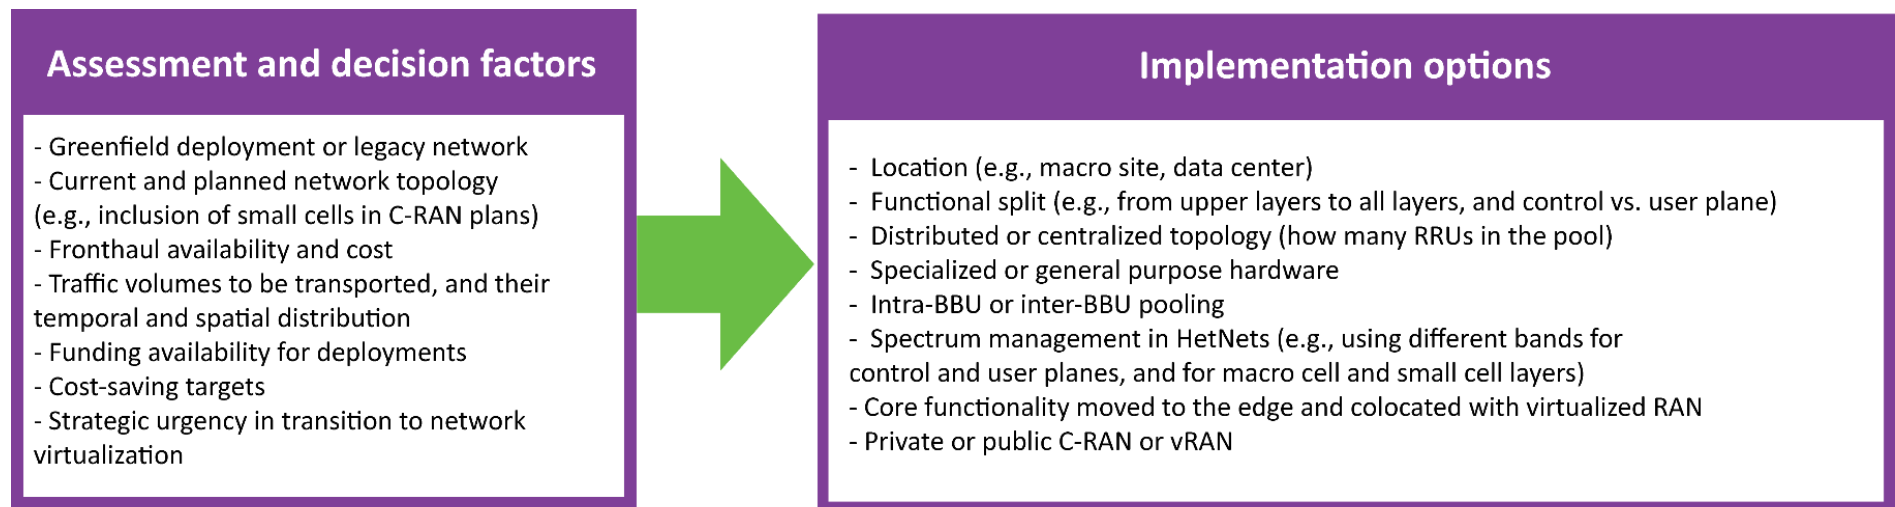
4. Picking the right C-RAN architecture.

Assessment of multiple tradeoffs crucial to successful C-RAN deployments

How will operators choose the C-RAN architecture that is best suited to their needs? In our talks with operators and vendors, we have identified factors that will guide their assessments and decisions. While the factors will remain the same, their role will change with time. Urgency may increase as fronthaul becomes more affordable. Or the inclusion of small cells may not be an issue initially if the operator has not yet deployed them. While the final target may be

a fully virtualized RAN, each operator is likely to chart a different path, depending on current network resources and requirements.

Based on the assessment and decision factors listed below, operators have to choose among multiple implementation options to determine which C-RAN type is best suited to their needs. In the rest of the report, we will discuss these options and the impact they will have on C-RAN deployments.



Source: Senza Fili

5. The path to virtualization.

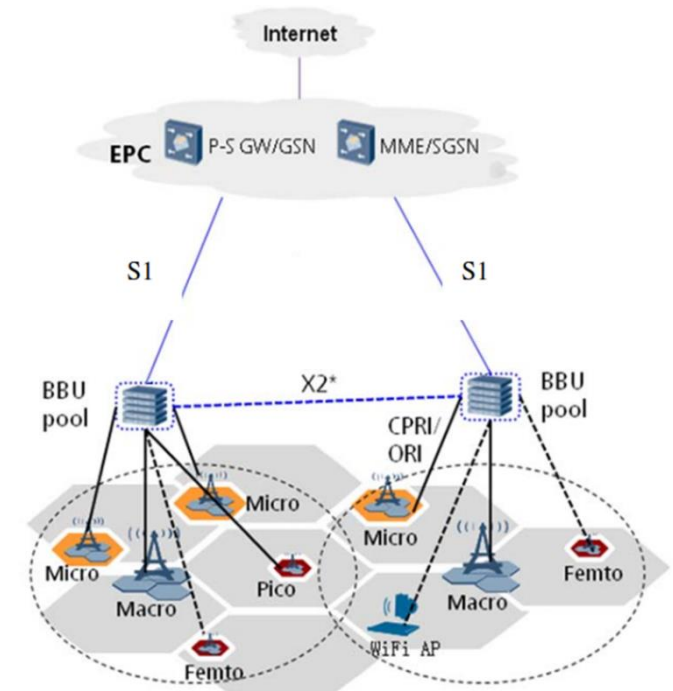
The role of the RAN within the NFV platform

Virtualization of the base station is one of the use cases defined by ETSI within the NFV platform. The virtualization targets that ETSI identifies are the baseband radio processing unit (PHY), MAC, RLC, PDPC, RRC, and CoMP. Two of the ETSI NFV PoCs are on RAN virtualization. In the first one, China Mobile, Alcatel-Lucent, Wind River Systems and Intel intend to test a C-RAN with general purpose hardware to support TD-LTE and GSM networks. In the second one, SK Telecom, Nokia and Intel will test a vRAN implementation in an LTE network.

In the China Mobile PoC, the hardware implementation includes modem servers (TD-LTE partial L1 and L2 protocols), a hardware accelerator (TD-LTE partial L1 functions such as FFT/iFFT and channel encoding and decoding), and control servers (TD-LTE L3 and GSM protocols). A detailed account of the PoC interim results is presented in the article “Recent Progress on C-RAN Centralization and Cloudification” by Chih-Lin I et al.; the authors demonstrated a successful implementation of CPRI compression, WDM for fronthaul, and CoMP. The uplink gain from CoMP was 20%–50% near the cell center, and 50%–100% at the cell edge.

The inclusion of RAN virtualization within the NFV platform has deep implications, and it increases the attractiveness of vRAN and its likelihood to succeed, because it becomes part of a platform that operators are already by and large committed to. More crucially, however, the inclusion of the RAN and the core within the shared NFV platform brings them closer to each other and removes the sharp demarcation line that has separated the two. Virtualizing the RAN outside the NFV framework could result in comparable cost savings and performance improvements, but would lack the long-term impact on the end-to-end network at the functional level. In a single-platform fully virtualized environment, the optimization of the use network resources can be more effective, because it can be done jointly across all network functions.

ETSI's LTE C-RAN architecture with a centralized BBU pool



Source: ETSI

6. Fronthaul, the great enabler of C-RAN and vRAN architectures.

A bottleneck or a requirement to meet?

Undoubtedly, fronthaul requirements are the major bottleneck in C-RAN adoption. The business case for C-RAN would be trivial if the fronthaul requirements were comparable to those of a distributed base station architecture. Invariably the argument against C-RAN rests on fronthaul: the capacity requirements are so high that only in areas with a good fiber infrastructure, such as many APAC countries, is C-RAN feasible. Even there it can be challenging, because fiber is not ubiquitous.

While fronthaul requirements pose a major challenge, the cost and performance benefits of C-RAN and eventually vRAN motivate operators – and vendors – to look for ways to meet them, rather than dismiss the opportunity. They are doing so in three ways. The first way is to reduce fronthaul requirements by using functional splits that enable some functions to be virtualized, while the others remain distributed. We will discuss this approach in the next section. The second way is to improve the efficiency of the fronthaul so that it can pack in more traffic. The third way is to use wireless alternatives where fiber is unavailable or too expensive.

Fronthaul efficiency can be improved by compression and by using solutions such as WDM instead of dark fiber (see table on the right). Dark fiber provides the best performance, but it is more expensive and, in many environments, its availability is limited. Solutions such as OTN, PON, CoE, WDM and wireless are less expensive, but they carry a performance penalty (see table), so their adoption has to be assessed against the operator's requirements.

Among operators that are more active in C-RAN trials and deployments, a combination of WDM and wireless is emerging as the preferred approach, as together they can provide cost-effective fronthaul connectivity to both small cells and macro cells. The ability to reach all cell sites within an area is a key benefit if operators deploy C-RAN to optimize performance and manage interference.

Fronthaul in China Mobile's PoC TD-LTE 20 MHz system, with 8 antennas

Requirements assuming 6 Gbps fiber connection using CPRI:

- 9.8 Gbps per carrier (2 fiber connections)
- 19.6 Gbps for UL and DL per carrier (4 fiber connections)
- 58.8 Gbps for a 3-sector macro cell (12 fiber connections)

Compression:

- CPRI compression (nonlinear quantization, I/Q data compression): 50%
- Single fiber bidirectional (uplink and downlink transmission over the same fiber connection): 50%
- With compression and bidirection, 3 fiber connections for a 3-sector site are needed, instead of 12.

Transport:

- WDM (multiple carriers on the same 6 Gbps fiber connection): in a network with 7 sites and 21 carriers, WDM requires 4 fiber cores, and dark fiber 42. WDM overhead latency was within the operator's accepted range.
- Microwave to reach sites without fiber.

Source: China Mobile (adapted)

While C-RAN can coexist with legacy distributed cell sites, the performance benefits vanish in an area where the C-RAN does not include all cell sites, because the gains from coordinated transmission and load balancing are compromised. To avoid partial C-RAN coverage, operators have to be able to reach all cell sites with fronthaul. Hence the combination of multiple solutions – and in particular the availability of wireless solutions – is crucial.

Fronthaul transport solutions

Dark fiber	Dark fiber is the native fronthaul solution that has the best performance. But because it requires a higher number of fiber connections, it is best suited to environments with abundant and affordable fiber.
Optical transport network (OTN), Passive optical network (PON), CPRI over Ethernet (CoE)	Attractive for outdoor small-cell fronthaul in dense urban areas, these solutions can reduce costs, but they have higher latency than dark fiber.
Wavelength-division multiplexing (WDM)	WDM enables an efficient use of fiber resources, and this is conducive to using fewer fiber connections compared to dark fiber, one way WDM can lower costs. It is also well-suited to multi-RAT networks. WDM is gaining traction among operators, and some consider it necessary to justify the business case for C-RAN and RAN virtualization. There are two types of WDM: CWDM and DWDM. CWDM is cheaper and easier to deploy, but most operators we talked to prefer DWDM, especially for the long term, because of its higher efficiency.
Wireless	Wireless is best suited to outdoor small-cell fronthaul, where distances are short and capacity requirements relatively low (small cells typically have only one sector). Wireless fronthaul brings cost savings, compared with fiber, for operators that do not have their own fiber transport network or in markets where fiber connectivity is expensive. Wireless links require LOS, which in some environments may not be available for all small cells. Relays can be used to compensate for lack of LOS, but they introduce cost and latency.

7. Fronthaul beyond CPRI.

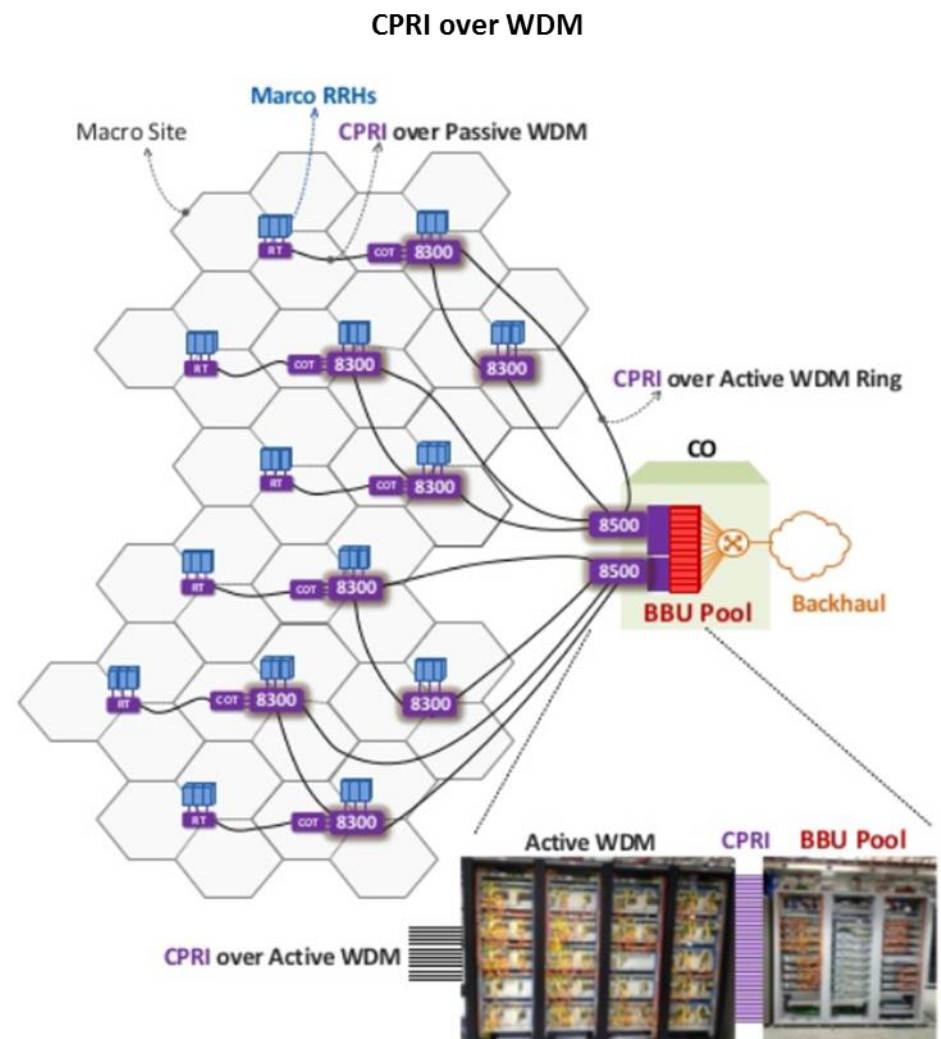
CPRI alternatives attract operators' attention

CPRI is the fronthaul interface that dominates today in both trials and commercial C-RAN deployments, because it meets the requirements for remote baseband processing with 0.4 ms latency, and very low jitter and bit error rate. It was originally developed by five vendors – Ericsson, Huawei, NEC, Nokia and Alcatel-Lucent – to connect the RRH and the BBU in a distributed base station; its use has naturally been extended to C-RAN.

However, there is some dissatisfaction with CPRI among both vendors and operators. They see it as a semi-proprietary vendor-defined interface. Although it meets the requirements, they do not consider CPRI a futureproof transport interface, and it is not a standards-based, open interfaces. In the short term, CPRI does provide a mature, well-established interface that has been crucial in proving the feasibility of C-RAN. Now, however, operators and vendors are exploring new interfaces and working on standard-based solutions.

One that is gaining ground with operators and vendors is ETSI's ORI, which is based on CPRI and designed to interoperate with it. As with CPRI, ORI is a digitized interface that connects the RRH and BBU in single-hop and multi-hop topologies and includes both user and data plane (L1 to L7), as well as the control and management plane and the synchronization plane.

Ethernet is another interface being considered, because it can leverage Ethernet equipment lower costs. However, getting Ethernet to match the level of CPRI's performance is not trivial, so Ethernet-based solutions have to be carefully managed end-to-end. As a result, at least initially, implementations will not be interoperable across vendors. This is also the case with other interfaces, though, so initial lack of interoperability is to be expected and does not limit the attractiveness of Ethernet for fronthaul.



Source: SK Telecom

8. Tradeoffs can ease the way to RAN virtualization.

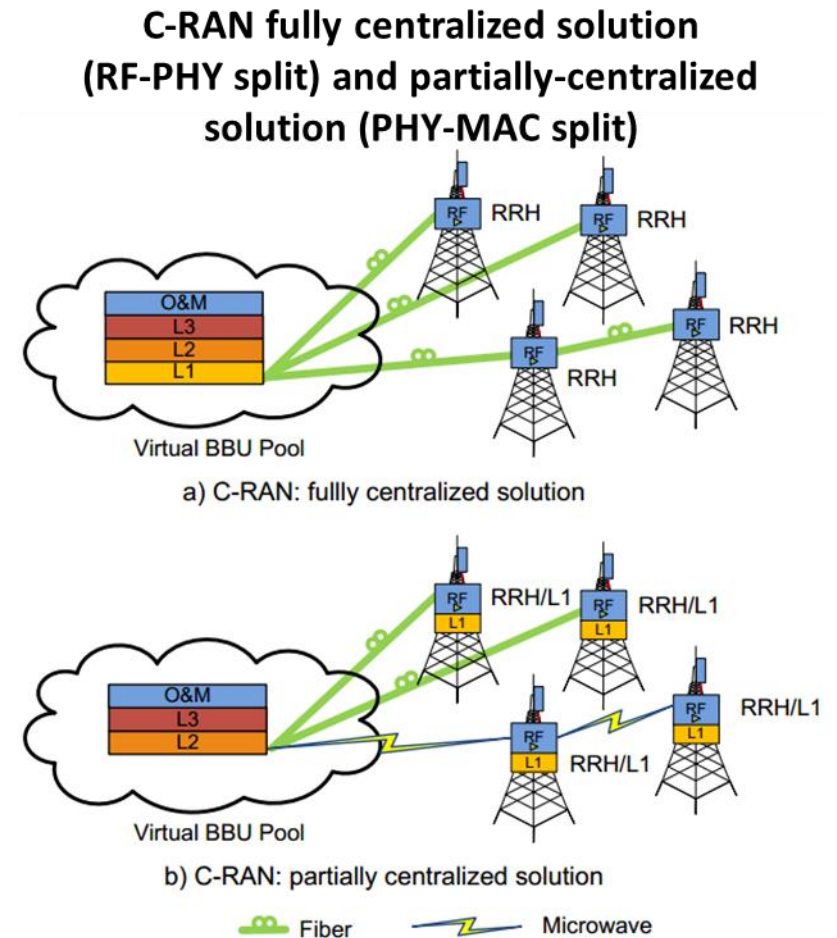
Functional split options to facilitate or enable the transition to C-RAN

The ultimate goal of RAN virtualization is to have all the BBU functionality pooled in a remote location using general purpose processors. This cannot be accomplished in commercial networks today, but is possible to contain the scope of virtualization efforts in order to get the transition to RAN virtualization under way. One of the crucial steps in this direction is to use functional splits that reduce the fronthaul capacity requirements – and, since fronthaul cost and availability are major bottlenecks, make C-RAN easier and cheaper to deploy.

When using a functional split, some of the BBU functions remain distributed and co-located with RRU at the cell site, while others (typically those in the upper layers) are moved to a data center or other remote location.

Multiple functional splits are possible (see graph below) and entail different tradeoffs, in terms of transport capacity needed and performance benefits. A RF-PHY split that allows remote processing of all BBU functions requires 20 to 50 times the bandwidth of a PHY-MAC split, according to China Mobile. In terms of fronthaul requirements, the difference is massive. But the performance gains vary as well. In the RF-PHY split, eICIC and CoMP can be implemented effectively, so the C-RAN can manage interference effectively. In the PHY-MAC split, coordinated transmission cannot be implemented in the pooled BBU, because they are L1 functions instantiated at the cell site. For an operator with a macro-only network, this can be an acceptable tradeoff. For an operator with a dense small-cell network, the reduced performance gains may weaken the overall business case for C-RAN.

While the choice of the appropriate functional split is dependent on the operator's assessment of its network conditions and performance requirements, the flexibility that functional splits allow is crucial to enabling a smooth and gradual transition to C-RAN, especially in legacy networks. The emerging trend is to implement, from the

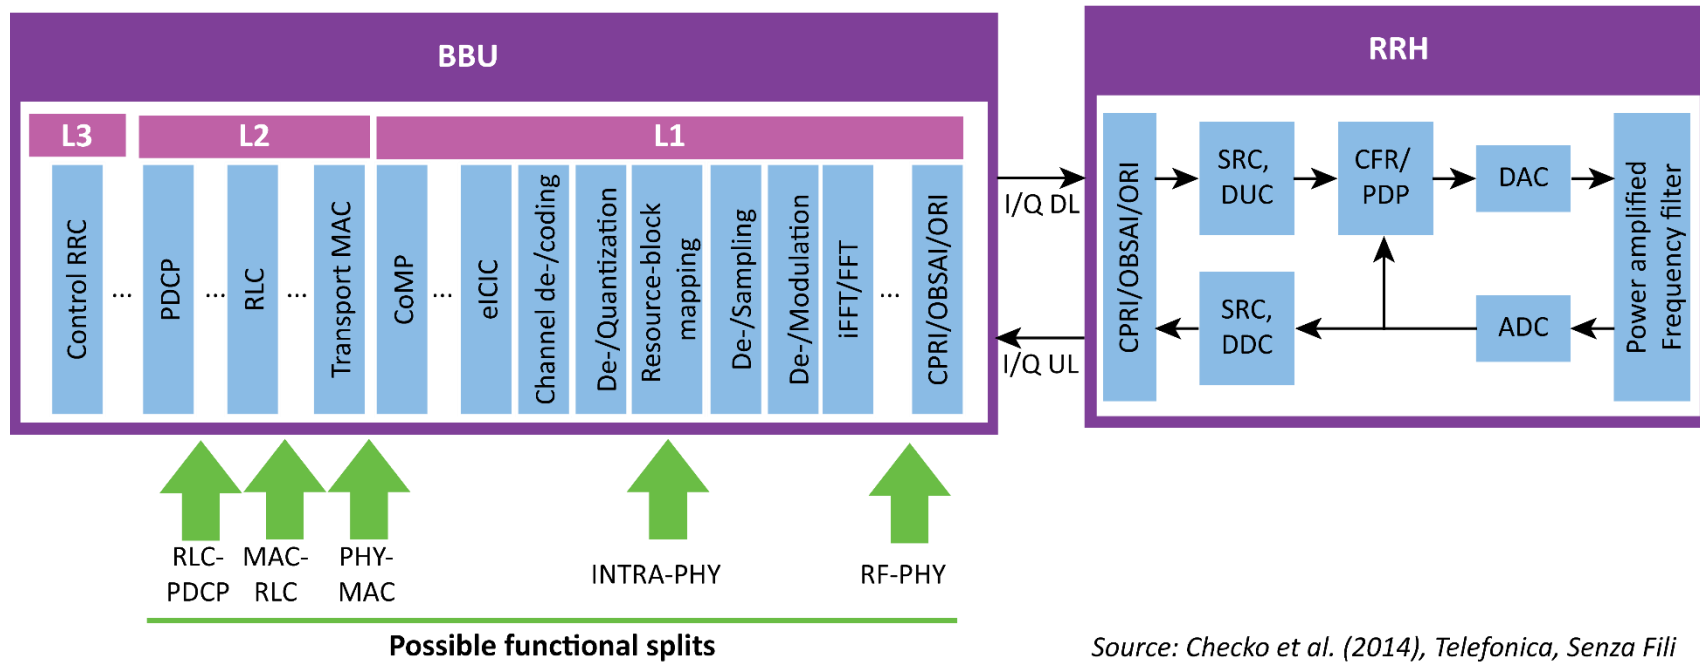


Source: Checko et al. (2014)

beginning, a fully centralized solution (RF-PHY split) for local C-RAN and for micro C-RAN – because the robust implementation of coordinated transmission requires L1 remote processing – and to select other functional splits for wider-area C-RAN.

An interesting variation at the intersection of C-RAN and V-RAN is used in the implementation at KDDI: most of the BBU processing remains at the cell site, but the scheduling is done remotely by a central scheduler over multiple cell sites. The MAC-level coordination among cells has lower latency and bandwidth requirements; hence Ethernet backhaul can be used effectively, instead of CPRI-based fronthaul, to improve resource utilization. This is an attractive solution in areas where fronthaul is not available at all cell sites. In this type of environment, the central scheduler across all cells may provide better performance improvement than a C-RAN with a RF-PHY split that includes only the subset of cells within the footprint with fronthaul connectivity.

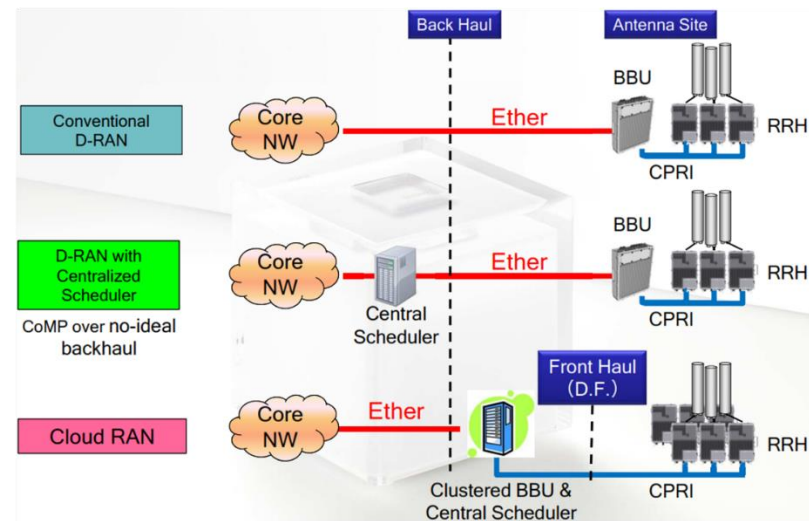
BBU and RRH functionalities and possible functional splits



Source: Checko et al. (2014), Telefonica, Senza Fili


Another interesting approach, promoted by DOCOMO, China Mobile and others, is to separate the control and user planes in HetNet deployments. Small cells, also called phantom cells, use a separate, higher-frequency band than the macro and transmit only user plane traffic. All the control functions (access, synchronization, RRM and handover control) are managed by the macro network. This C-RAN architecture lowers the cost and complexity of deploying small cells, and integrates them more tightly within the network, improving mobility management and reducing the signaling load on the network.

D-RAN with a centralized scheduler at KDDI



Source: KDDI

Phantom cell architecture with control and user plane split



Source: DOCOMO

9. Drivers to a virtualized RAN: cost.

Cost dynamics in the short term and in the long term

The previous sections have looked at how C-RAN and vRAN may be implemented. But what motivates the transition to a virtualized RAN – a long-term process that requires effort and investment from operators? Cost savings are clearly a primary driver. There is a growing body of data showing that C-RAN delivers cost savings and reduces power consumption in specific deployments or trials, but there is considerable variation across markets and operators.

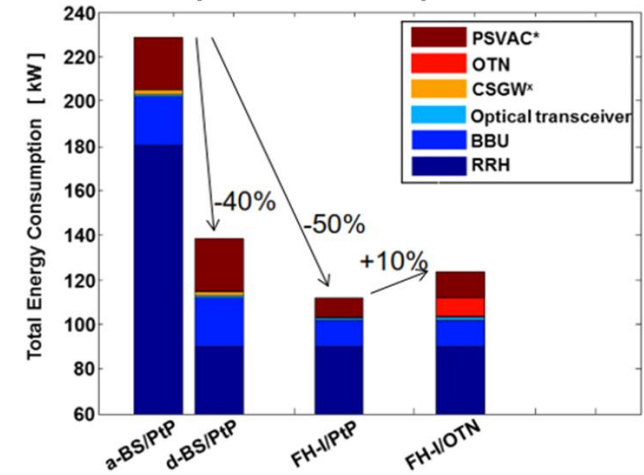
In its C-RAN trial (see figure on the right), Orange found a steep decrease in power consumption upon the transition to a distributed base station with RRH, and a further decrease when moving to a C-RAN architecture with wireless or wireline fronthaul.

China Mobile (see figure below) has shown a 30% reduction in capex and a 53% reduction in opex in its trials. Capex reduction mostly comes from equipment and installation costs. Most opex cost savings come from the lower site rents. C-RAN cost savings are made possible by multiple factors:


- Site selection: easier to identify suitable locations and negotiate a lease.
- Lower construction and maintenance costs, because the RRH site is easier to manage and because BBUs can be jointly maintained.
- Lower energy consumption.

In addition, C-RAN performance benefits and increased resource utilization add to the cost savings by lowering the per-bit costs. The capex and opex figures compare installation of the same infrastructure in a traditional RAN and in a C-RAN, so these results do not include the cost benefits derived from increased performance (e.g., the operator may need fewer cells if the throughput per cell is higher).


Reduction in power consumption with C-RAN




a-BS/PTP: Traditional BS + backhaul



d-BS/PTP: distributed BS + backhaul



fronthaul only



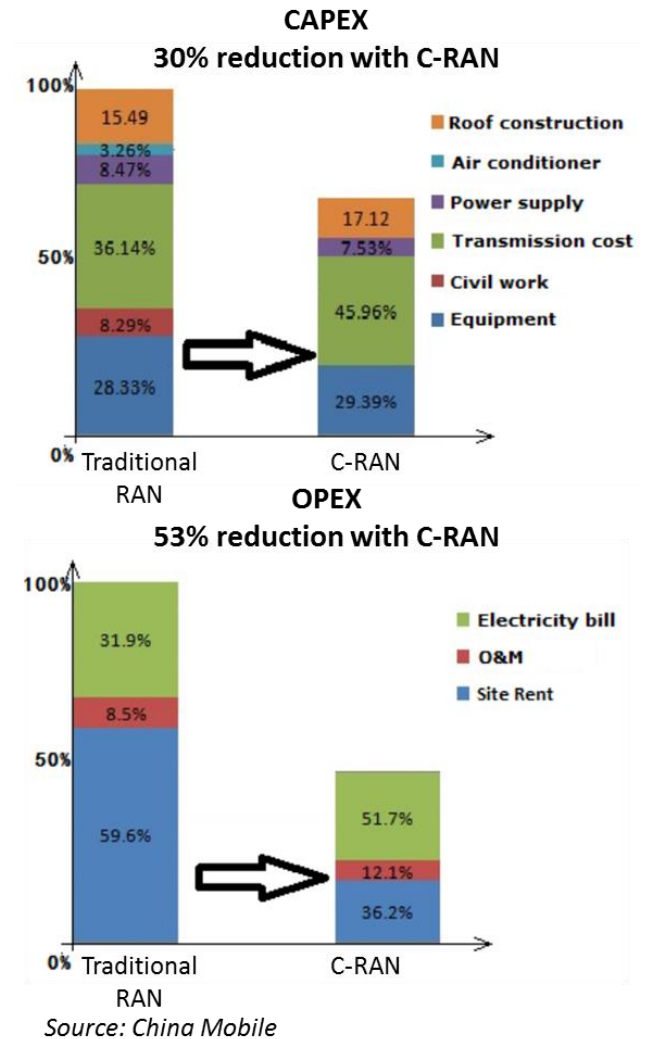
Source: Orange

The analysis of the per-bit costs is crucial to getting a comprehensive view of C-RAN adoption's financial impact, because per-bit costs factor in both the per-site cost savings and the performance-induced cost savings. For most deployments, both savings sources are needed to create a strong business case supporting C-RAN architectures. As operators move to a fully virtualized RAN, we expect the performance-induced cost savings to become dominant as the benefits from vRAN expand and resource utilization increases. As a consequence, it is likely that the cost savings will increase over time, because most initial C-RAN rollouts will benefit only from a subset of the vRAN features.

The cost savings reported by China Mobile are impressive, but it is difficult to assess how they can generalize across operators, geographies, and C-RAN or vRAN architectures. Among the sources of variability in cost savings are:

- **Fronthaul costs.** These vary greatly not only by market, but also by operator. Operators that own a fiber network enjoy a huge advantage, in both cost and access to fronthaul solutions. In addition, the selection of different combinations of fronthaul solutions (e.g., dark fiber versus WDM versus wireless) may result in further differences in cost savings.
- **Greenfield versus legacy deployments.** The business case for C-RAN is easier for a greenfield deployment, but in most cases C-RAN deployments have to coexist and be integrated with legacy RAN technologies. Multi-RAT C-RAN requires a higher initial expenditure to upgrade the existing infrastructure, but then delivers higher opex savings than an LTE C-RAN.
- **C-RAN type, functional split, hardware.** The type of C-RAN selected (e.g., local C-RAN versus metro C-RAN), the functional split (e.g., RF-PHY versus PHY-MAC), and the choice of dedicated versus general purpose hardware have a significant impact on the performance gain and cost savings.

Cost saving from C-RAN



10. Drivers to a virtualized RAN: mobility and interference management.

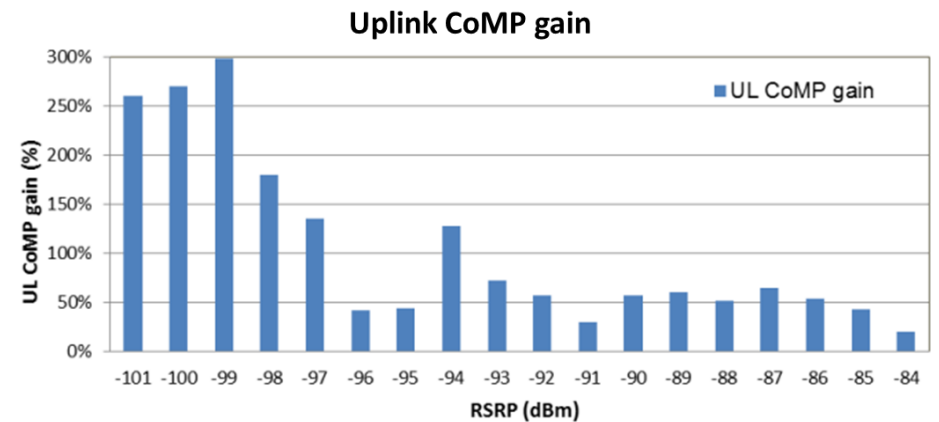
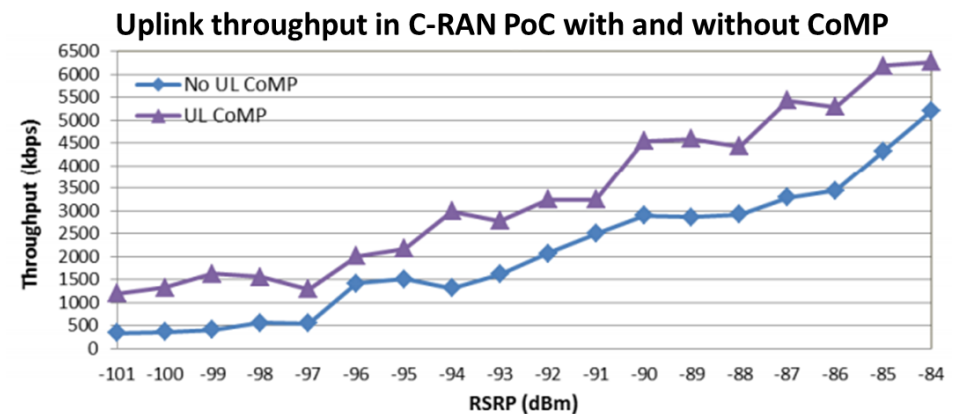
Toward a more efficient use of network resources in densified networks

C-RAN and vRAN deployments can lower the cost to deploy and operate cells, but the real and long-term financial benefits come from a higher utilization of network resources. The improved network performance that comes from this translates in financial gains, but they are more difficult to quantify than the cost savings at the single-cell level, and may accrue in a longer time frame. As a result, the virtualization of the RAN requires mobile operators to take a longer view of the transition, much as they need to do for the virtualization of the core. In fact, because of the demanding baseband processing requirements, full RAN virtualization is likely to take longer than core virtualization, and hence require an even more distant time horizon.

Improved resource utilization comes from different functionalities that are available to mobile operators even without C-RAN, but remote pooling enhances them because all the baseband processing is done at one location. All functions that require intercell coordination (e.g., eICIC, CoMP, mobility management) benefit from the extremely low latency at the BBU site.

Mobility management is a beneficiary of pooled processing, with faster and more robust handoffs. With the introduction of small cells and, more generally, with network densification, increasing the efficiency in mobility management becomes all the more important, because handoffs are bound to become more frequent with the subscribers' more frequent crossing of cell borders as they move around. In addition to improving handoffs, a C-RAN architecture can foster new ways to manage mobility that do not require handoffs – for instance, by treating an area with multiple cells as a logical single entity that has no cell edges, within which handoffs are not required.

Robust interference management is a main driver for C-RAN adoption, especially in local C-RAN deployments. The ability to effectively manage interference in co-



Source: China Mobile

channel outdoor small-cell deployments, where interference from small cells may cause an unacceptable degradation of macro capacity, will be many operators' primary motivation to install C-RAN.

At the same time, the ability to manage interference with tools like eICIC and CoMP in a C-RAN environment greatly enhances the business case for HetNet deployments with small cells, because the combined macro- and small-cell capacity is considerably higher with these tools – and, as a result, efficient implementation of eICIC and CoMP is a requirement for many operators planning small-cell deployments. To date, the inability to successfully implement CoMP in a traditional RAN architecture has slowed down the adoption of outdoor small cells. The graph on the previous page shows the gain from uplink CoMP in a C-RAN architecture in the China Mobile NFV PoC. Gains from CoMP are stronger (40%–100%) in areas with weak coverage (RSRP lower than –95 dBm), but still present (20%–50%) in areas with good coverage (RSRP higher than –90 dBm).

As with mobility management, C-RAN architectures will also facilitate or make possible entirely new approaches. The DOCOMO phantom-cell approach – in which small cells and macro cells use separate bands but transmission for both layers is managed jointly within the C-RAN – enables operators to integrate components in a HetNet and completely avoid co-channel interference. This makes the deployment of small cells much more attractive, requiring less RF planning and posing no threat to macro performance. Current plans to deploy LTE in license-exempt bands using LAA-LTE are expected to use a similar model. Transmission in LAA-LTE small cells can be managed by macro cells that – on the basis of real-time traffic demand and resource availability – select the traffic that has to stay on the macro layer and the traffic that can be diverted to the small cells.

This approach can be pushed further to include legacy interfaces – 2G and 3G – as well as Wi-Fi within the C-RAN, to optimize multi-RAT coordination by extending it beyond the single-cell boundaries. In a multilayer, multi-RAT HetNet architecture, the ability to manage transmission above the single-cell level can be beneficial to load balancing and can assist the network transition during spectrum refarming.

11. Drivers to a virtualized RAN: statistical multiplexing.

Leveraging uneven traffic distribution to optimize resource utilization


The traditional RAN architecture is unavoidably inefficient in network utilization. In Europe and North America, typically only 20% to 40% of the available network capacity is used. This is not because of bad network planning or operations, but mostly because traffic is not uniformly distributed in time or across the cell sites in the footprint. Subscribers tend to make calls or use their data connections during the day, and most do so during peak hours in places where many people congregate. Clearly it is not possible to eliminate the unevenness in traffic distribution across time and space, but it can be harnessed in C-RAN architectures, to some extent, by using statistical multiplexing.

The figure below illustrates how pooled processing can lead to higher resource utilization when individual cells have different temporal profiles, as is typically the case. Statistical multiplexing can greatly enhance network utilization if the

C-RAN includes RRUs from zones with different traffic profiles – e.g., if it combines business areas with residential areas. In this case BBU resources can shift to give more support to cells in one area or the other as demand changes throughout the day or the week.

The same approach works within a smaller area – e.g., within a metro hotspot – within which the daily traffic distribution is homogeneous. In this context, statistical multiplexing can increase BBU utilization and improve network performance by exploiting the real-time variation in traffic.

Statistical multiplexing has implications for the type of C-RAN topology chosen. If an operator intends, for instance, to include business and residential areas, it will likely choose a wide-area C-RAN strategically planned to include the appropriate mix of locations.



Source: Checko et al. (2014)

12. Implications: The death of the cell?

Moving baseband processing to the cloud erases RAN boundaries

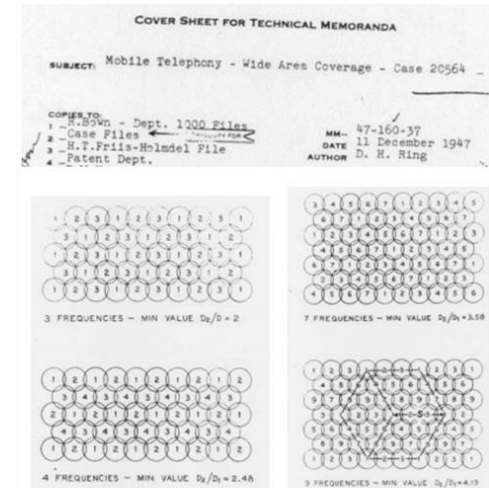
The virtualization of the RAN calls into question the concept of a cell in a mobile network – enshrined in the 1947 Bell Labs memo by D. H. Ring – as the foundational unit in which transmission originates and is managed. In a C-RAN, transmission remains in the cell, but baseband processing moves to a remote location and is pooled across multiple cells. In a vRAN architecture, not only is the processing done remotely, it is abstracted from the hardware – i.e., hardware resources are shared among cells. But more importantly, to reap the benefits of RAN virtualization, transmission has to be coordinated across cells, especially in multilayer, multi-RAN HetNets.

In this context, baseband processing becomes increasingly independent of the cell type and detached from the individual cell.


In a traditional D-RAN, the distinction between a macro cell and a small cell is clear – from the hardware form factor, to functionality, to power levels, to physical location. In a HetNet, even as small cells increasingly support the same functionality as macro cells do, they still have a role clearly separate from that of macro cells. This distinction quickly breaks down in a C-RAN environment, where small and macro cells share the hardware and have to jointly balance the use of spectrum resources.

Instead of maximizing the throughput for each cell, a C-RAN can optimize the throughput for the entire network, thus moving from a cell- to a user-centric perspective. The goal for the BBU is not to pack each cell with as many packets as possible, but to provide subscribers with the best possible experience, given the network resources available.

Douglas H. Ring & W. Rae Young, Bell Labs, 1947



China Mobile's "No more cells" approach



Source: China Mobile

The separation of control plane and user plane across cells in the advanced C-RAN architecture similarly challenges the concept of the individual cell as the fundamental unit for transmission processing.

China Mobile's vision for the "No more cells" approach includes four areas of change:

- The move from cell-centric to network-level resource management
- Elimination of cell physical ID, intercell interference, and handoffs
- Virtualization of infrastructure and improved spectrum utilization
- Decoupling of the signaling and data planes, and of uplink and downlink

The move to a user-centric network-based approach represents a fundamental shift in managing mobile traffic that is bound to have a huge impact, and in ways that it is too early to anticipate.

13. Implications: A blurring demarcation between core and RAN.

The importance of physical location in a virtualized world

The inclusion of RAN virtualization within the wider network virtualization heralded by NFV and SDN increases its attractiveness, cost efficiency and effectiveness for both mobile operators and vendors. In addition, it challenges the separation between the core and RAN. Traditionally, core and RAN have operated in physically separated environments, using different specialized hardware and addressing different requirements. Different teams within mobile operator organizations are in charge of these two network components, with distinct capabilities and approaches to network deployment and operations.

From the edge, RAN virtualization erases the borders between cells. From the center, it erases the demarcation line with the core. As BBUs separate from the RRU, they move closer to the core, both physically – as BBUs enter data centers or other locations where core functions may be co-located – and logically, as BBU processing shares the same NFV and SDN platforms. In a fully virtualized environment, core and RAN can share resources, which can lead to cost savings and increased efficiencies.

But with virtualization, there is more than a common platform. Core functionality can move farther out toward the edge and start merging with RAN functionality, as part of the trend (discussed in the previous section) toward a user-centric approach to optimizing the allocation of network resources. At the same time, RAN BBU processing may remain close to the RRU in a local C-RAN topology, or it may move closer to the core with a heavily centralized hyper C-RAN.

Virtualization frees mobile operators from restrictions about physical location – they can locate and distribute functionality across the network as they see fit. But this is not because physical location is no longer important. Quite the opposite, it is precisely because they can choose the physical location and leverage its advantages that physical location becomes more salient in a virtualized environment.

Distributed or centralized?

Examples of core functionality moving toward the edge

Real-time video / content optimization – Proximity to the edge lowers latency and enables services to adjust to network conditions and traffic load in real time.

Caching – Pushing caching to the edge may be more efficient and improve user experience by lowering latency.

TCP optimization, CDN, DSN, DPI, location-based services – A distributed approach for services that are tied to a network location may improve performance, reduce latency and increase resiliency.

How centralized should the virtualized RAN be?

The degree of centralization depends on the type of C-RAN selected. Local C-RAN includes only one macro site and the small cells in its footprint. A micro C-RAN consists of the small cells within the venue.

But how large should a metro or hyper C-RAN be? Fronthaul limitations set the maximum radius at 40 km, but most C-RANs are likely to be well within this limit.

According to Telefonica, C-RAN centralization obeys the law of diminishing returns. Most cost savings are accrued with 5 to 10 sectors. With 50 to 270 sectors, savings drop to less than 5%. As the size of the C-RAN grows, so do complexity and costs, so operators have to find the sweet spot that best meets their requirements.

14. Summary.

RAN virtualization, as an integral part of NFV and the SDN, plays a crucial role in driving a higher utilization of network resources within the evolution of wireless networks.

The transition from the traditional, distributed RAN to the fully virtualized vRAN will include multiple C-RAN types – local C-RAN, micro C-RAN, metro C-RAN, hyper C-RAN, and advanced C-RAN – among which operators will choose those that meet their network requirements and strategy preferences.

The choice of C-RAN topology is tied to the evaluation of multiple tradeoffs that give mobile operators wide flexibility in reaching their cost savings and/or performance targets.

Many industry players see fronthaul as an obstacle to C-RAN deployments, but there are many solutions – e.g., WDM and wireless fronthaul, CPRI compression – that can help operators reduce the financial impact of fronthaul and deploy C-RAN in areas without fiber connectivity.

CPRI is currently the dominant digital interface for fronthaul, but there is a growing dissatisfaction with CPRI among operators. Work is under way to develop ORI, a standards-based, open interface, and to use Ethernet-based solutions.

Functional splits that keep some of the baseband processing within the cell site make it possible to reduce fronthaul capacity and latency requirements, and facilitate the transition to C-RAN in legacy networks.

C-RAN deployments can reduce the capex and opex when deploying new cells, but high fronthaul costs may limit these cost savings. However, operators should look at the per-bit cost savings, which provide a more accurate way to assess the full financial implications of C-RAN.

The main driver to C-RAN deployment is improvement in performance, due to improved mobility and interference management and the ability to use statistical multiplexing in the allocation of network resources.

RAN virtualization challenges the concept of the cell as the basis of baseband processing, and opens the way for a shift from cell- to user-centric network optimization and resource allocation.

With virtualization, physical location assumes crucial relevance, with core functionality migrating toward the RAN and RAN functionality becoming more centralized, blurring the demarcation between wireless core and RAN.

II. Vendor profiles and interviews

Airvana

Airvana has worked on small cells for a long time, historically focusing on CDMA femtocells used in residential and small-business environments to improve coverage and capacity. With the OneCell solution, Airvana has moved into the enterprise LTE market using a single-cell C-RAN architecture to address crucial cost and interference challenges that DAS and small cells, respectively, face in indoor settings.

In OneCell deployments, Radio Points within an enterprise building or public venue are connected by an Ethernet LAN to a Baseband Controller, which is responsible for the baseband processing. All traffic from the Radio Points is relayed to the core network through the Baseband Controller.

All Radio Points act as a large single cell with transmissions coordinated by the Baseband Controller. This architecture departs from the dominant every-cell-for-itself model in which small cells interfere with each other along the overlapping borders. Instead, multiple Radio Points simultaneously receive transmissions from subscriber devices, and similarly, will jointly transmit to subscriber devices, thus eliminating cell-border interference and simplifying RF planning.

In addition, because the Radio Points all act as a single cell, there are no handoffs to manage as subscribers move within the coverage area and saves stationary subscribers from ping-ponging

between cells. Not only does this makes it easier to provide consistent LTE performance, it reduces overhead signaling activity needed to support mobility, which can be taxing in environments with a high density of small cells.

With Smart Reuse, the system simultaneously serves multiple users on the same frequency resources, optimizing the use of network resources and avoiding interference, according to Airvana.

Because the system runs over an Ethernet LAN, it reduces the cost of deployment versus DAS or RRH solutions that require fiber and/or coax cabling for in-building transmission. Like DAS, however, OneCell can support a multi-operator neutral-host business model, because the hardware supports multiple bands and each operator can select the one it wants to use.

By deploying the networks indoors, operators minimize interference with the macro network,

while providing the best possible support for indoor use, which accounts for the largest portion of data traffic (80% or more). As a result, macro network transmission does not need to be closely coordinated with the underlying small cells, as in an outdoor deployment, thus simplifying traffic and interference management.

OneCell may help operators protect their investment as they upgrade their network. Most OneCell functionality resides in the Baseband Controller, which uses a macro-cell grade chipset, which is software upgradable and can support carrier aggregation and other LTE-Advanced features. Because Radio Points are multi-band and the Ethernet infrastructure is frequency-independent, operators can simply re-configure OneCell's frequency band as they obtain new spectrum or re-farm spectrum channels in use, instead of having to upgrade each small cell in the network.

OneCell Radio Point and Baseband Controller



Source: Airvana



Taking C-RAN indoors to optimize small-cell performance

A conversation with Mike McFarland, Senior Director of Product Management and Marketing, Airvana

Monica Paolini: Our conversation today is with Mike McFarland, Senior Director of Product Management and Marketing at Airvana. Mike, how did Airvana get involved in the C-RAN space?

Mike McFarland: We've been in wireless since we were founded in 2000. We focused on macro-cell equipment, and established a very large and successful business there, OEM-ing products to both Nortel and Ericsson for EV-DO macro cells.

Back in 2006 we started development in small cells. Initially, we focused on stand-alone small cells that act completely independently as individual cells, and are targeted for residential and small business.

But with LTE, particularly in the enterprise environment, we saw an opportunity for C-RAN concepts to come into play and improve the performance and economics of a solution there. Towards the middle of last year, we announced a new product called OneCell, which is an enterprise LTE small-cell solution.

Monica: What is C-RAN for you? Everybody I talk to seems to have a different opinion.

Mike: In its broadest definition, we look at a C-RAN solution as having centralized baseband processing that is pooled across a variety of access points or remote radio heads, or RRHs. If you look in the macro environment, in a typical C-RAN deployment the baseband is centralized in a data center or exchange center. There are fiber lines running a CPRI protocol out to RRHs. Each RRH acts as an individual cell sector.

That's a very common definition for C-RAN. But when you're looking at small cells, particularly in a public venue or enterprise environment, there's a lot more that can be done with C-RAN. First and foremost, you can do coordination across those access points with joint transmission and reception to user devices, to boost the signal-to-interference and noise ratio (SINR) for those elements and get better throughput.

Another aspect is that operating these access points as a single cell eliminates handovers within the enterprise or public venue. In a typical deployment with stand-alone small cells, you have a lot of handovers – 20 to 25 per user per hour. That's because people are moving around in the enterprise. Some of them might be ping-ponging between individual cells, particularly if they're sitting between the cells. If you can create a single-cell solution in the enterprise that has no handovers, you can improve performance.

There is a big opportunity in the enterprise to take advantage of Ethernet rather than using the

traditional CPRI-over-fiber connectivity. By using Ethernet between the baseband and the remote access points, you reduce costs.

Monica: A lot of times CPRI is considered to be the requirement for C-RAN. Do you think that the C-RAN fronthaul could use Ethernet as well?

Mike: The fronthaul can definitely use Ethernet. You have to be careful about how you architect your solution in order to use Ethernet. It's not a trivial exercise, but we found a way to deal with some of the complexities of running over Ethernet versus just running over a dark-fiber line and using CPRI.

Because of that, our solution is very easy to deploy. Think of it as following the managed Wi-Fi deployment model. Anyone who understands Ethernet or has an IT background can put in the Ethernet cabling and switches, or can reuse existing infrastructure that's there. This reduces the cost and complexity of the deployment significantly.

Monica: What about the latency and the throughput?

Mike: You have to be careful about them. Our solution runs over Gigabit Ethernet, so it does require links of more than 100 Mbps to really get the full performance out of a 20 MHz LTE channel.

You obviously have to be careful about latency too, but we haven't found any issues to date. We typically test with up to four or five Ethernet hops. You don't usually go over that number in a

campus-enterprise deployment. So far, we're very pleased.

Monica: This raises the issue of interference. When you are indoors, you are much better off than if you're outdoors, because you don't have so much interference with the macro cells. But the small cells do interfere with each other, because they are in an environment where, as you mentioned, there are many handoffs.

Mike: One of the nice things about indoor small cells is that you do have less interaction with the macro than in an outdoor small cell. Outdoor small cells are challenging to deploy, because you have to get them installed, usually, on some sort of lamppost. That means you may have to get rights from the municipality to install them. And you need to provision backhaul and power to that location. All these things are complex, but then you have the additional complexity of people driving in cars and moving very quickly back and forth. You've got to do a lot of handoffs very frequently, very quickly.

When you're indoors, you're more insulated from those issues. There are typically defined entrances in and out of a building, or even a stadium or public venue. The people aren't usually driving in cars, so they're moving more slowly. The amount of interaction between the macro and the indoor small cells is much lower.

If you run the system inside the building like a single cell, like we do with our OneCell system, you have just a single cell that's interacting with the macro rather than having 20 or 50 or 100

stand-alone small cells all trying to communicate with the macro to handle handoffs and manage interference.

Consider for example the macro trying to perform eICIC with many individual small cells versus a single cell, which makes it a lot simpler in terms of interference management. It also makes it a lot easier for handsets to do interference cancellation. They can usually cancel out one or two interferers, but not more than that. In this case, they would just have the one macro-interferer, or vice versa.

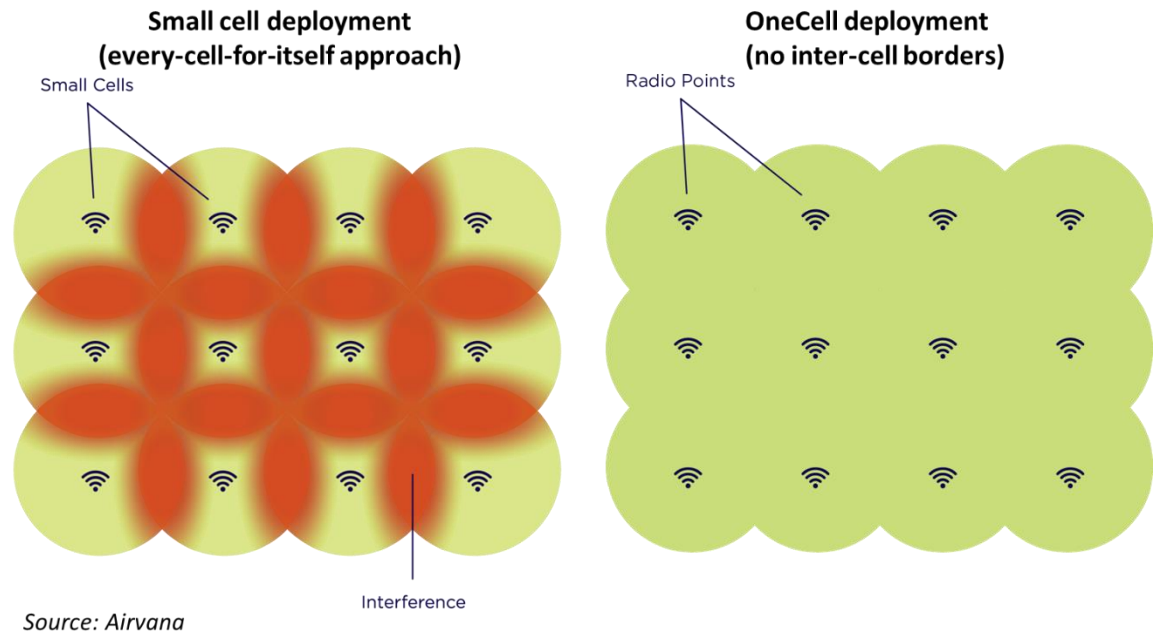
Monica: What led you to develop the OneCell system?

Mike: The initial idea of this approach started four years ago. That's when we started doing some

very in-depth, hardcore R&D into how can we split the baseband processing from the Radio Point, and be able to run this over Ethernet, but also create a single-cell architecture.

We developed the OneCell concept after looking at some of the challenges in an enterprise or public venue environment. You want to have consistent data rates throughout the building. When you have stand-alone cells, they create borders. At those borders, they create interference. They're not coordinated with each other. X2-based interfaces, or proprietary X3 interfaces, don't work fast enough to keep up with a system that is scheduling a multitude of users every millisecond.

Interference has tremendous impact on people's data rates. In these areas of interference we've



Source: Airvana

measured degraded performance by up to 90%. It also increases jitter, which has a significant impact on VoLTE. We've seen up to a 5X increase in jitter in these areas of interference on the borders.

The borders are not necessarily small. It can be 30%, 40%, maybe even 50% of the building, depending on how much overlap you have between cells. Of course, cells have to overlap in order to avoid coverage holes.

Having the system act as a single cell has significant advantages, because you don't have a decrease in data rates. You don't have the issue with the jitter. You have very good data performance. You have very good voice performance. There's the added benefit that you don't have any handovers. And you have just a single interface to the operator's core network.

Monica: How do you help the enterprise or the operators who deploy for the enterprise to go in and address all these issues?

Mike: Our solution has two main components.

There's a Baseband Controller. This is different from your typical controller that you might have seen on the market from other solutions. It does all of the baseband processing, all the scheduling for the users across all of the access points, which we call Radio Points, which is the second component.

We call them Radio Points because they're primarily just functioning as the radio. Those radio points are connected back to the baseband

controller over standard Ethernet LANs. Again, others talk about using Ethernet, but what they're really talking about using is the copper within the cat 5 cabling. In contrast, I'm talking about using regular Ethernet framing, and we connect through standard Ethernet switches. Using standard Ethernet makes our solution very easy to deploy.

We do have a third element, which is our Device Management System, or DMS. It sits in the operator's core network and manages the solution, and can manage a very large number of these deployments. It provides a plug-and-play capability. When the Baseband Controller initially boots up, it will automatically discover the Radio Points. They will connect to each other, and then the Baseband Controller will automatically connect


through a security gateway back to our management system.

The management system will then provision all of the appropriate configurations related to the core network and to the RF interfaces. We've tried to make the solution plug-and-play, with minimal configurations required from the installers. Of course, that enables operators to use lower-cost labor for most of the installation.

Monica: Is your solution better suited for larger enterprises?

Mike: It's definitely targeted for environments that are greater than 50,000 sq ft, or 5,000 sq m. If you're in a smaller environment than that, you're probably looking at only one or two stand-alone

Interference reduction with OneCell



SINR shows higher interference at the cell borders for standalone small cells, with an average of SINR advantage for OneCell, according to Airvana.

Source: Airvana, iBwave

small cells, maybe three, to cover the environment. In that case, it probably makes sense to use standalone small cells.

We have products in that space as well. We're not saying that standalone small cells are bad. They're a great solution for small or medium businesses, as well as for residential environments. But for larger enterprises, you need a C-RAN architecture.

Monica: How does your solution compare to DAS?

Mike: One of the great strengths of DAS is that it's targeted for multi-operator deployments. We specifically spent time architecting our solution to make it multi-operator. OneCell hardware is common across operators and it gets its operator personality from software. We also have a Multi-Radio Point Enclosure that holds up to four Radio Point modules.

It's fairly compact, about the size of a small toaster, because the Radio Points aren't very large. That allows us to feed the four Radio Point modules into a single antenna for aesthetic purposes. You can have a single antenna down below the ceiling, and have the Multi-Radio Point Enclosure above the ceiling, out of sight.

The Ethernet infrastructure can be shared across the four operators, so you don't have to duplicate it. The baseband controllers only take half of a rack unit, so you can fit two of them in one rack unit, and four of them in two rack units, so it's very compact.

We've also tried to make the solution multi-operator friendly from a business model perspective, similar to DAS, in the sense that the Radio Points are common across operators and regions. We have a North American Radio Point that covers all the major LTE bands for all of the North American operators and an International Radio Point that does likewise for other markets.

Typically, with DAS they'll purchase and deploy a lot of the cabling and antenna infrastructure separately. Then they'll have the operator bring in their base station. Similar business models can be used, if desired, in the case of our solution, because the Radio Points and the Ethernet infrastructure can be installed into the buildings, and then the operator can bring in our Baseband Controller to plug into it.

Monica: For the enterprise, VoLTE may be important to support voice services. It's not the bandwidth that is the issue with VoLTE, it's making sure that you have a reliable connection and low latency.

Mike: What has surprised us is how fast VoLTE is coming. We did not expect the market to move this quickly. I think it is somewhat due to the iPhone supporting it. We see over 50 operators, worldwide, in some stage of VoLTE deployment.

There's a lot of functionality that goes behind VoLTE in terms of the scheduler and providing the right guaranteed bit rates. What's unique to our solution is that we don't have any handovers. As users move around, they do not have to do frequent handovers between cells.

As I mentioned before, we don't have any jitter that gets introduced into the system at the cell borders, because we don't have any cell borders. That also helps improve our VoLTE performance.

Monica: We talked about multi-operator deployments. That is crucial because the business models for small cells in general require some sort of multi-operator environment to make the solution cost effective. How do you see that working out?

Mike: It's up to the operators to define how they want these business models to work. We leave it to them. What we try to do is provide an architecture that is flexible so they can implement the business model that they choose. Operators can go a traditional route, where they are deploying small cells into enterprises and public venues themselves, and they fund the equipment, so they own the network from start to finish. That's something that we can support.

Some operators may be interested in sharing the cost with enterprises, and enterprises may benefit from having better coverage inside their business. When someone comes to see a condominium or an apartment, one of the things that they are going to consider is, what's the wireless coverage here like? If the wireless coverage is fantastic in every single apartment and condominium, it's going to help retain their tenants as well as help fill vacancies. Thus, the owner of the building benefits. It's not unreasonable to expect the building owner to pay for some of that wireless infrastructure – in fact, they may be willing to fund a lot. In the past, the only alternative that building

owners had was DAS, but DAS has been prohibitively expensive. It wasn't viable economically. But with our OneCell solution, which is a fraction of the cost of DAS, building owners can fund most of the deployment and have the carrier perhaps only provide the Baseband Controller. The building owner would fund the Radio Points, the Ethernet infrastructure, and all of the deployment costs.

There are a variety of models here that can be leveraged by operators to reduce their cost and be able to service their customers very well by improving their coverage and capacity.

Monica: What do you hear when you talk with operators?

Mike: It varies. In some cases, they are looking to get more subscriber contracts from the enterprise. In other cases, they have a very strong interest in off-loading a nearby macro cell, particularly in dense urban environments, where macro cells are overloaded.

By offloading to an enterprise small-cell solution, the macro suddenly becomes less loaded. The users that are on the small cells will get much better performance, and the users outside the enterprise will get better performance from the macro cell, too.

Monica: Indoor small cells allow you to use the spectrum more efficiently. You increase spectrum reuse.

Mike: That's right. In the case of a new building with energy efficient windows very little RF signal goes in or out of the walls or windows, and so the indoor system becomes quite isolated from the outdoor system, except in the lobby environment.

There are a lot of advantages to using indoor small cells to improve the capacity of your overall network, versus outdoor small cells, with which you have to deal with a lot of interference challenges and handoffs.

Monica: In closing, what should we expect to see from C-RAN in the next few years? What new directions you are moving into with respect to C-RAN?

Mike: We are continuing to invest in our solutions, particularly in the management system. We have nearly a thousand small cells in the footprint of one single macro cell. That has not been easy to achieve. We've had to refine our management system and small cells to support advanced hybrid SON capabilities, so they are interacting in an appropriate way with the macro, and ensuring that all the handoffs are working.

That's been a big focus for us. I think you are going to see more news from us about our management system and how it can ensure that the small cells are not impacting the macro, because one of the biggest concerns operators we hear from operators is: "if I am going to go deploy hundreds of thousands, or millions, of small cells, I need to

make sure they're not going to impact my KPIs on the macro."

We've been able to achieve that, and we are bringing that technology now from 3G to 4G, in both our OneCell systems as well as our other 4G products.

About Airvana



Airvana enables mobile operators to deliver dramatically better service to subscribers in the most challenging and high-value environments: offices, public venues, and in the home. The company's award-winning 4G LTE and 3G small cell solutions ensure "five bar" voice coverage, superior data performance, and expanded capacity for wireless services. Independent research firms consistently rank Airvana among the small cell market share leaders. Headquartered in Chelmsford, Mass., USA, Airvana is a founding member of the Small Cell Forum and has 14 years of experience providing commercial wireless equipment to over 90 operators globally. For more information, visit www.airvana.com.

About Mike McFarland



With over 20 years of experience in the communications industry, Mr. McFarland currently leads Airvana's product management and marketing. Prior to Airvana, he was a Sr. Product Line Manager at Cisco Systems for several multi-billion dollar product lines. He was a founder of SetWave, and has had strategy and technology roles at Mercer Management Consulting, Ameritech, and Accenture. Mr. McFarland holds an MBA from the J. L. Kellogg School of Management and a BSME from Northwestern University, and graduated first in his class from both schools.

ASOCS

After the initial work on a programmable SDR platform for mobile handsets, back in 2012 ASOCS decided to shift its focus to the virtualization of the base station. While the ASOCS solution is designed to work in C-RAN centralized environments, the target of the company is to go beyond C-RAN to a cloud-based virtualized RAN, where all baseband resources and layers are pooled in a remote location and virtualized on general purpose hardware.

Central to the ASOCS vRAN solution is a software-based implementation of the traditional or distributed base station technology. The ASOCS vBS is aligned with the NFV requirements set forth by ETSI and supports SON and other LTE-Advanced features such as CoMP to increase the efficiency in the use of an operator's network and spectrum assets.

In the ASOCS vision, the centralized RAN is the initial phase towards virtualization. Centralized RAN delivers cost savings and lowers power consumption, but provides only some of the benefits of full virtualization. In the next stage, the vRAN brings the C-RAN into the SDN/NFV-based mobile network, widening the scope for cost savings and making the allocation of resources more flexible.

The vBS is treated like any VNF and co-located in data centers along with other VNFs. By extending the NFV platform to the RAN, operators create a

more homogeneous network. This enables operators to push closer to the edge such functions as EPC, IMS, CDN and security that today are centralized, and as they become VNFs as well, they can be flexibly deployed where needed.

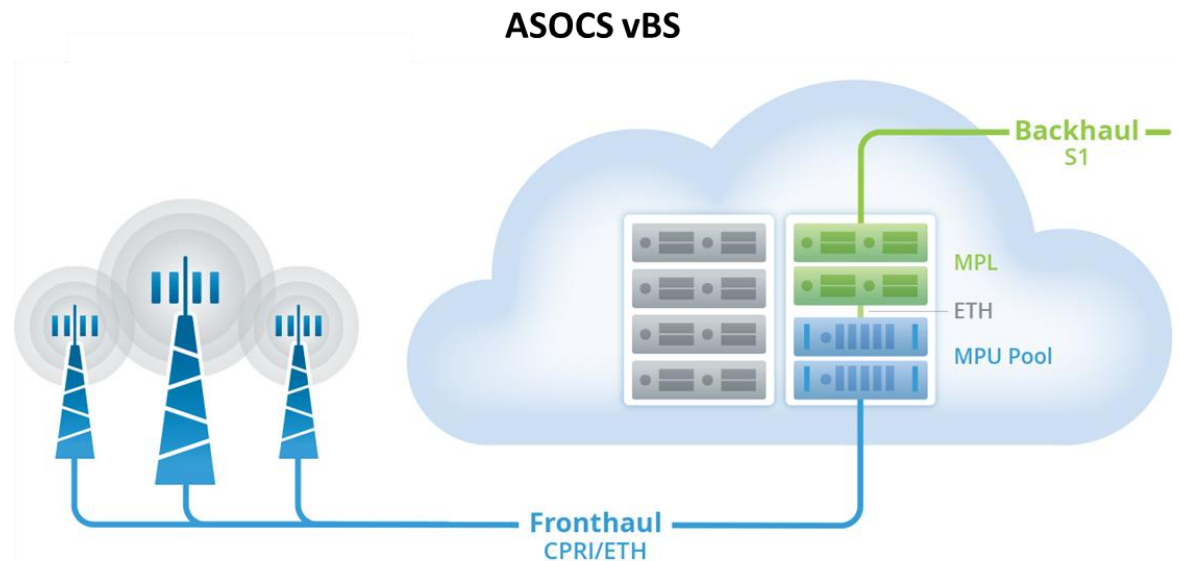
The ASOCS vBS targets three C-RAN segments:

- Mega-scale C-RAN, covering footprints that have a radius of 30–40 km. These are the C-RAN deployments that currently dominate in APAC.
- Metro C-RAN, covering a radius of 5–10 km, to be deployed in high-density environments, mostly in the US and Western Europe.
- Micro C-RAN, covering large venues (e.g., malls, airports, stadiums) and the densest urban areas.

The vBS solution works for both DAS and small-cell deployments, and supports neutral-host and communication as a service (CaaS) models.

The vBS covers all stack layers, including a virtualized L1. It is based on the ASOCS Modem Processing Unit (MPU) – a software-defined radio-programmable platform – and on the ASOCS Modem Programming Language (MPL).

In order to meet the real-time and deterministic low latency RAN requirements, ASOCS developed its Real-time-on-COTS (ROC) platform which is an open-source base, cross-platform software solution that transform massive scale cloud computing and SDN servers into a secure, reliable, real-time deterministic platform.



Source: ASOCS

Moving past C-RAN, to get the performance and cost gains of full RAN virtualization

A conversation with Gilad Garon, CEO and Founder, ASOCS

Monica Paolini: Today we are talking to Gilad Garon. He is the CEO and Founder of ASOCS. Gilad, how did you decide to start working on vRAN?

Gilad Garon: ASOCS has been around since 2003. In the first eight or nine years of the company, we focused on mobile handsets. In 2012, we made a dramatic shift and decided to switch the company's main operation into RAN virtualization. We started out working with the likes of China Mobile to transform base stations to Intel architecture in the early days of C-RAN.

As this industry matured and virtualization took hold over the industry, we found that there's an opportunity not just to sell silicon technology, but rather to offer a full NFV-compatible solution, end to end. This is when we decided to focus our efforts solely on building virtual base stations, rather than being silicon engines for the industry.

Monica: So you moved pretty early on, in the early days, towards the virtual transition path. What were the reasons? What made you make the decision? What were the drivers?

Gilad: We started with C-RAN a little bit before NFV started to become a mainstream ideology in the course of 2013. But through working with the Intel architecture, once virtualization became mainstream, we were NFV-ready.

In the early days we were thinking mainly about using virtualization technologies to solve capacity and cost challenges in the industry. But as this industry was maturing, we discovered that virtualization would not just be a cost reduction measure, it is the way that next-generation networks are going to be designed. This requires virtualizing the base station, which was a key component in the data path and in the cost structure of a network, as well as end-to-end services from the evolved packet core all the way to the edge.

So while we started with virtualization as a cost reduction measure, leveraging the intrinsic cost reduction of using IT equipment over proprietary telecom equipment, later on it became the springboard of moving into this NFV world.

Monica: What do you mean by C-RAN and vRAN, and how do they relate to each other? There seem to be a lot of different concepts out there in the industry.

Gilad: Again, as we started quite early on in the industry, the difference between 'C' as in

Cloud-RAN and vRAN was actually quite marginal. Later on, the term C-RAN was hijacked by traditional base station equipment vendors who said, "OK, the revolution would be that we would centralize existing base stations" and the term C-RAN became centralized RAN.

We were always, from day one, more focused on Cloud RAN, which is a combination of centralization and cloud. So now we simply use the word Cloud -RAN in order to differentiate from centralized RAN.

vRAN is just another term to try to emphasize the fact that we're talking about virtualization technologies. In centralized RAN, actually, the revolution or the transformation of the industry is quite limited. In places like Korea, centralized RAN was already deployed as early as two or three years ago, and the results were quite marginal.

Monica: Now we can go back to the cost and functional benefits issue. As you said, at the beginning, for a lot of players the first impulse was, "Well, we'll do C-RAN because of the cost savings," and then they realized there are also performance benefits.

Especially in terms of using your network resources efficiently, because your spectrum is limited as an operator. What the C-RAN allows you to do is to optimize that resource.

Gilad: You're absolutely right, Monica, that the cost of spectrum is a key component when you look at the total cost of deploying these new

virtualized or C-RAN networks. Spectrum is the scarcest resource, not the computation cost.

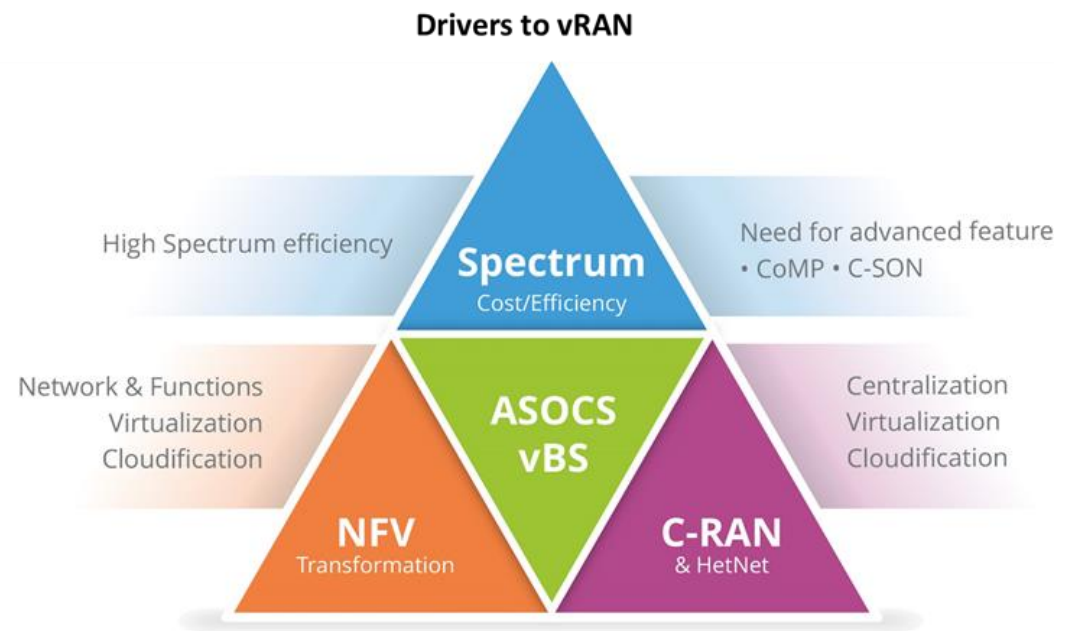
If you ask any radio manager on the planet, whether he would want a solution that brings more processing power to solve spectrum, he would tell you, "Absolutely, bring me as much processing as you can, because spectrum is scarce for me." Processing power, especially in the world of cloud computing, is limitless.

This is the core of the matter. C-RAN enables us to do exactly that. Every processing resource is limited, but for the problem at hand it's literally unlimited, because the scale of data centers and the amount of computing power available today on server technology enables us to throw this DSP processing at the problem.

For that we actually need also the centralization, because if you have a remote site where your antenna or the cell's tower is completely isolated, then the ability to bring unlimited processing to that tower would not make economic sense.

So in order to solve spectrum problems by using, shall we say, excessive processing power, you need to have that processing power located in some centralized location. And you want to have all these antennas connected.

One of two main mechanisms for solving spectrum limitations is CoMP, which stands for coordinated multipoint. Essentially CoMP uses the positive side of interference. This industry, for the last 30 to 40 years, has only focused on signal-to-noise ratio: how do you reduce interference of a single



Source: ASOCS

handset with a single tower, treating the rest of the towers as enemies, as the noise. Instead, CoMP uses all this additional signaling out there to enhance the signal.

The second approach, which is starting to happen and we're happy to be a part of, is C-SON, or cloud self-organizing networks, which are a relatively new field. Using cloud technologies for SON enables a real-time adjustment of existing networks and new C-RAN deployments. Whether the market moves more into CoMP or moves more into C-SON or moves together using both technologies, we would see significant spectrum utilization enhancement. Even the conservative estimates mention 20% to 35% increase in efficiency. Judging by cost of spectrum worldwide,

we're talking about tens of billions of dollars' worth of savings.

Monica: There has been a lot of trouble implementing CoMP and SON outside of a C-RAN environment. So it's a mutually reinforcing trend here. In this context, how important are small-cell deployments, which are where SON and CoMP are crucial? How important is the deployment of small cells in driving C-RAN, or of C-RAN in driving small cells?

Gilad: We're starting already to see trends, in the Small Cell Forum, of talking about virtualization. Once these small cells are virtualized, they're just remote radio units, as far as we're concerned. So whether they have on-board some signal

processing or off-board, or are they integrated with Wi-Fi or not, as far as we're concerned, they're just yet another resource for us to virtualize or to co-manage through SON.

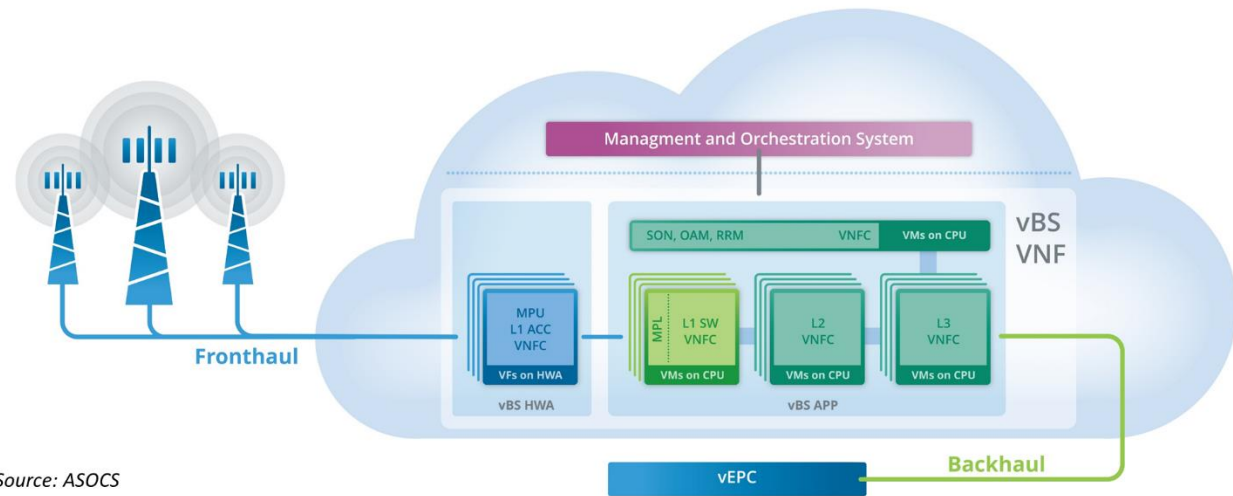
Monica: Once you have a C-RAN type of deployment, does it make any sense then to talk about macro cells, small cells, micro cells, pico cells – or are they just units within a C-RAN environment?

Gilad: I absolutely agree with that, with your assessment. One open issue is how do you manage legacy? Initially C-RAN was deployed as a greenfield technology. But as this industry is maturing, C-RAN has become a capacity enhancement layer over traditional legacy equipment. It wouldn't make sense for a carrier who is already in the second, third, or fourth year of LTE deployment to consider a greenfield C-RAN deployment.

Monica: The success of C-RAN and virtualization depends on the fact that you can adopt C-RAN in legacy networks. How can you help operators with a legacy network?

Gilad: The problem is complex. If we look at what we call the micro C-RAN – the in-building network in a campus, stadium, or airport – it could be deployed gradually. But the C-RAN is yet to prove itself as to whether it enhances the capacity in the already-deployed areas, and for that, as we mentioned, technologies such as C-SON will come into play.

ASOCS V-RAN network topology



Source: ASOCS

What we envision are hybrid cell sites, with a certain amount of on-premises DSP processing, which is already geared to and, in many cases, fused together with the antennas. Some of the virtual base station resources can be deployed at the cell site, or moved away, according to traffic. It's just another layer of coverage, which is cheaper and also mostly statistical, if we're again making it more economical to the carriers.

Monica: I guess that also works in the evolution towards vRAN, where you're not going to have everything virtualized on day one. You just have to go through different phases. Can you talk about how you see that developing through time?

Gilad: We all know eventually how it's going to end. One of the nice things about this industry is that it's clear that networks will be virtualized. Most of the economic benefits of applications such

as machine-to-machine and automation are not going to be realized, unless the network is virtualized.

The industry of today is focused more on the functions that are easier to virtualize, such as packet core processing, which was by definition more transaction oriented and which is going to be the first to move into virtualization.

This is definitely something that's already in early deployments. By the end of 2015, a little bit of 2016, we can have all of the packet core virtualized. That's done. But then there are other areas of the network we could virtualize as well. Here, the base station is critical. If I had a crystal ball, I would say that by 2018, we're going to see a great deal of the tier-one operators having a fair share of the network virtualized.

What these carriers are also looking at now is their strategy towards 5G. They do not want to repeat the same aching experiences they had with 4G, although it was better than 3G. One of the nice things that they like about vRAN is that it's going to be native in 5G.

Operators are looking at the problem and saying, "OK, we're getting some benefits in 2016 and 2017. But the bulk of the benefits will be in 2018 and beyond. What's good for us is that we already have the layout, which is going to take us all the way into 5G with a software upgrade."

Monica: Now let's talk about fronthaul. Some say that lack of fronthaul solutions will prevent C-RAN from happening. How do you deal with the fronthaul?

Gilad: An ideal C-RAN requires an order of magnitude – whether it's 6X or 8X – more capacity for fronthaul than backhaul for distributed RAN.

There's no denying that the limiting factor is the availability and cost of fronthaul; that is without argument. For this reason, many thought that C-RAN would be limited to Asia, where fronthaul appears to be much more available than in Europe and the United States.

When we entered this market, that was also our perception. What we have started to learn are two things. First, when you look at the in-building DAS or Distributed Radio Antenna Systems, or DRAS, market, fronthaul is not an issue, because these markets are moving into optical transport.

In the US, we have discovered that fronthaul is a non-issue, with the exception of specific cities or towns. Sometimes, fronthaul becomes an issue, not necessarily because of the technology, but because of the way operators buy backhaul or fronthaul equipment or service. So the answer again is that it's not all black and white. There are going to be areas in which fronthaul will delay the deployment of C-RAN.

If you look at the way that this industry is moving, I don't really see fronthaul as a significant barrier. The barriers are due to the fact that some carriers still have to make the decision that virtualization is the way they go. Most of them have – and at the end of the day, it also requires them to change their traditional business model somewhat.

That's a question mark that needs to be resolved at the management and board-room level around the globe. It's less the technology factor, the way we see it. Pretty much, other than very specific use cases that we've heard, the fronthaul problem is not a fundamental deal breaker, but rather something that can be resolved.

Monica: Let's talk a little bit about what you do at ASOCS. What sets you apart?

Gilad: I would say that what dramatically differentiates us from others is that we took the virtualization challenge from day one, and we aimed at solving it in a disruptive way.

We're not about taking legacy or existing architectures and redefining them at L3 or L2 or L4. We have gone all the way down to the core of

the virtualization of the baseband. If you want to present a truly disruptive solution, you have to solve the problem at the core, and this is what we have done.

We took the challenge of how to run communications and real-time applications, which by definition are more deterministic and can never fail. We are running them on IT equipment, which is by nature more statistical. We have managed to bridge that gap between the statistical advantages of IT and still being able to provide equipment that is reliable enough for telecom.

Putting it into more technical terms, we look at L1 resources as just another compute resource. In the classical NFV format, where the industry is looking at compute storage and networking, we added, if you will, another subsection into compute.

For us, monolithic L1 or monolithic PHY don't exist anymore. We treat the elements of radio processing just like any other compute resource. Using our core technologies that have been developed a decade before, and adding them to the classical Intel-architecture approach, we fuse these two together. Now you have a DSP or a physical layer resource which behaves like a compute. This is perhaps the dramatic gap, difference or differentiator that we have over other players in the industry.

As for the base station protocol software and applications, ASOCS present a unique open architecture solution approach, collaborating with RadiSys and virtualizing the proven Trillium LTE Software.

Monica: In closing, what should we expect from ASOCS, and maybe from the industry, in the next two, three years?

Gilad: 2015 is all about proof. This industry is looking for proof points that go beyond proofs of concept. Carriers are now saying, “We’ve got that. Let’s start playing with this a little more seriously. Let’s start building some limited-scale projects.”

There’s a lot of appetite in the carriers and other players to deploy small-scale C-RAN networks. All of them across the globe are looking at 2016 as the first year of commercial deployment, whether it’s going to be in this in-building or DAS topologies or in some layers of enhancement of capacity in urban areas.

There are also things that need to mature. In the NFV ecosystem, there are a lot of players who are

offering specific solutions at the orchestration level and at the component level. We’re still going to need to see a lot of software integration and maturity in the NFV industry.

Once these two parts mature, we see the later part of 2017 as a sweet spot of an industry takeoff. This is what we’re seeing right now, from all the discussions we’re having.

About ASOCS



ASOCS is a pioneer in development of virtual Base Station (vBS) solutions, enabled by its Modem Processing Unit (MPU) and Modem Programming Language (MPL), designed to meet current and future requirements, Cloud - Radio Access Networks (Cloud-RAN), Distributed Antenna System (DAS) and other wireless infrastructure cells. A vBS and its related resources will be managed as any other Network Function Virtualization (NFV) Virtual Network Function (VNF), virtual Infrastructure Manager (VIM) system and configured by an Orchestration platform, expanding NFV and SDN clouds to the edge.

About Gilad Garon



Gilad Garon, CEO, ASOCS, brings extensive knowledge in the telecommunications market with 20 years of hands-on experience in international sales, marketing and product management. Mr. Garon co-founded ASOCS in 2003 with the vision of software defined Communications removing all barriers to create a seamless holistic network. Today that vision is becoming a certainty. Mr. Garon led the transition of ASOCS from a semi-conductor company in the mobile space to a system software company pioneering virtual base station solutions in the NFV era. He holds a Master in Business Administration in Marketing and Information systems for the Hebrew University of Jerusalem.

Fujitsu

Experienced in providing transport and access solutions, Fujitsu offers a comprehensive set of C-RAN solutions to give operators the flexibility to meet their fronthaul requirements with resources they already have.

At the same time, Fujitsu has been very active on the access side, with its remote radio head and small-cell products. This gives the company an end-to-end perspective in the solutions for C-RAN that are available and being developed.

Fujitsu solutions support multiple interfaces – including CPRI and Ethernet – and transport options: dark fiber, OTN, PON, TWDM, CWDM and microwave. The array of solutions allows operators to balance cost, latency and performance tradeoffs. For instance, dark fiber is preferred where fiber is ubiquitous and affordable. But in many small-cell deployments, fiber is unavailable or too expensive, and operators need a wireless solution that meets fronthaul's requirements for latency and throughput and supports the chosen interface, which today is CPRI in most cases.

Fujitsu's BroadOne GX4000 is a wireless fronthaul solution to connect RRHs to a centralized BBU. This enables operators to include small cells in their C-RAN plans and, thus, more effectively manage interference using the X2 interface through tools like CoMP and eCIC.

The GX4000 uses the e-Band spectrum (71–76 GHz and 81–86 GHz); that spectrum is license-exempt

and provides high capacity, at very short distances. This is an advantage for tightly packed small-cell fronthaul, which is typically needed over short distances to bridge the link from an outdoor small cell to the closest fiber aggregation point (often the macro cell). In addition, the use of short, narrow links is conducive to low interference levels. Low power consumption, small footprint and low equipment costs are additional advantages.

Using Fujitsu's Impulse Radio technology, the GX4000 can carry up to 3 Gbps uncompressed per link, with per-link latency lower than 20 ms when using CPRI, according to the company. Operators

can use Fujitsu's NETSMART 1500 Management System for provisioning, operating and maintaining the wireless fronthaul links.

The FLASHWAVE® product line provides a set of complementary transport solutions that use packet optical networking, and also support the CPRI and Ethernet interfaces. These products can also be managed by the NETSMART 1500 Management System. The FLASHWAVE platform includes OTN, PON and CWDM, well suited for fronthaul. The choice among these alternatives has to be assessed in light of the impact that the adoption of the CPRI features may have on latency; that, in turn, affects C-RAN performance.

Fujitsu GX4000 E-Band Radio



Source: Fujitsu

Building end-to-end support for C-RAN and virtualized RAN

A conversation with Femi Adeyemi, LTE Solutions Architect, Fujitsu

Monica Paolini: As part of the report on C-RAN, today I'm talking to Femi Adeyemi, LTE Solutions Architect at Fujitsu.

Femi, could you give us an introduction to what you're doing with Fujitsu on the C-RAN front?

Femi Adeyemi: I've spent the greater part of my career in the wireless industry, building base stations and base station routers, and now I'm involved with Fujitsu on defining the next LTE architecture, for both small cells and fixed wireless.

Monica: There's really a lot of work to do there, because often we think of C-RAN as moving the baseband to a remote location, but there is much more to the C-RAN architecture than just having a remote baseband. At Fujitsu, how are you trying to capture the whole change that accompanies C-RAN?

Femi: We're looking at C-RAN as a total solution, both in the present and as we march into the future. For us, like you just described, it's not just remote radio heads and centralized basebands.

What we're finding out is that as you add more and more access technology, remote radio heads, small cells, the likelihood that you're going to introduce a lot of data into your core is very high.

We see this as a complete solution, where we understand what is going on at the access, understand the transport (CPRI or other fronthaul transport), understand the core, as well as all the individual chains that you need as you go from that access into the core. So, access, transport, aggregation points, as well as the core.

We're looking at this as a complete solution. For us that's very important, because we don't want to look at it as a disjointed technology on the access or the baseband centralization or virtualization. We see this as a total solution.

Monica: Let's start with the access. What do you see there?

Femi: On the access, we're very active in small-cell deployments. We build our own small cells solutions both in Asia – and deploy them in Asia – as well as in North America, both in the US and Canada.

Fujitsu has been very well renowned for building very reliable remote radio heads that have been deployed in many mobile networks today. On the access site, when we talk about remote radio heads, we are very prominent in that area, as well as small cells for enterprise and residential use. On the access side, we're very, very active.

Monica: That's step number one. The second step is the transport. This becomes more interesting

and challenging at the same time, when you talk about having a C-RAN architecture in conjunction with small cells.

Femi: Right.

Monica: Because of the requirements. What's going on there?

Femi: That is very correct. There is a stringent requirement for sending your I/Q samples from the baseband unit to the remote radio head.

In the industry, the technology you will have heard about is CPRI transport. On that fronthaul transport, we want to give operators flexibility and to allow them to choose, depending on cost and depending on deployment scenarios, which technology they want to deploy.

At Fujitsu, we're really looking at three different areas in fronthaul transport. Number one is what we call active CPRI transport – for example, where we are able to monitor what is happening with remote radio heads while we are still in this centralized location, which is a little bit more expensive but gives you visibility into what is happening between your remote radio head and a central location on your fiber.

For operators that are interested in cutting cost or not really concerned about what is happening at the remote radio heads, we have a passive solution as well: CWDM- and TWDM-based passive solutions.

Now, there are instances also where, for example, you have small cells that are outdoors and you

want to transport your data back to a centralized location, and you don't have access to fiber. For this scenario, we have a wireless transport solution that is E-band based.

We covered all the areas that will be very interesting to any operator. If you are interested in cost saving, we have the solutions for you. If you're interested in monitoring what is really happening among your fiber links, it's available from Fujitsu, as well as if you do not have access to fiber and have to transport your data back into the core.

Monica: Basically, an operator can decide how to mix and match. It's not necessary that one operator only have one solution for a whole network depending on where they are. It can get whatever works better for them.

Femi: That is correct. If you go back into what we've been discussing, we look at this as a complete solution. We don't want to limit the operator. We want to give you what is available based on cost, based on fiber availability, as well as based on network topology that the operator may have today.

Monica: Let's go to the next stages: the core and the aggregation. What are you doing in those parts?

Femi: As we look at the deployment that is going on today, going from the distributed RAN into what we call centralized RAN, as well as the next stage of virtualized RAN, we don't want to come into the industry with technologies that will not be deployable. We want to follow what the operators are really interested in.

Today we are moving from the distributed RAN into the centralized RAN. At the centralized RAN port, we have baseband aggregation happening. In the terminology in the industry today, we call that building baseband hotels.

We have a solution in that space, whether it's for small cells or for a macro cell. This will vary in the number of baseband units we are able to aggregate together, whether you want to aggregate a large number of basebands or a small number of basebands.

Then the next logical step for us is to virtualize that baseband. We've been able to operate this in the cloud along with some specialized devices, so that we move in these three steps that I'm describing: moving from the distributed RAN to a centralized RAN, and then into the terminology that I would call a cloud RAN, which will essentially virtualize the baseband hotels that you have.

With the baseband hotel, it's now easy to deploy logically, adaptively, sharing resources between the different baseband units that you have in the hotel.

Monica: So it's more than just a one-step process here. How do you see operators moving across this transition, especially those that already have a network and have to deal with legacy deployments? They have to look backward and forward. They need to look forward, but they also need to take into account what they already have.

Femi: That's a very good question, in the sense that we're running into scenarios – particularly

Fujitsu in-building LTE small cell



Source: Fujitsu

with established operators – where they already have a topology for their RAN architecture today.

Now, you can go in and ask them, "Oh, you want to deploy C-RAN architecture; let's pull apart what you already have and just start building from scratch." But that would lead to disruption in their service. What we're finding is that we have two layers that we're looking at, where we start with the existing architecture today, and when they plan for an expansion of the current topology, we build that out as cloud RAN, or C-RAN.

We keep what they already have in their network RAN architecture, but extend out our C-RAN to

their new deployment or attachments to their existing architecture. That's one scenario.

Another scenario is when they have a fresh greenfield deployment. Now we begin to think about complete C-RAN architecture for such greenfield deployment.

We want to make sure we work with the operator and get the best out of existing technology while we introduce this C-RAN architecture. There are issues that we come across as we do this. They have an existing transport layer. We want to make sure that the new C-RAN architecture works with the transport that they already have in place.

It's a two-stage process. We are not going in to break apart what they already have; we're building C-RAN on top of existing architecture and then, in the future, also think about pulling out what they have today, so that we can completely make it into C-RAN architecture.

Monica: How long will all this take?

Femi: It's ongoing, even as we are having this discussion. We are working through all these issues of transport. I can give you an example. We talk about CPRI fronthaul transport, which we are now going to integrate with, in most cases, fiber transport, as well as with Ethernet transport.

We are working through those issues, and within the next 12 to 18 months we are getting very, very close to fully deploying C-RAN architecture with most, if not all, operators in North America.

Monica: What about the rest of the world?

Femi: Well, C-RAN is actually very, very prominent today in Asia-Pacific, in South Korea, in Japan. Because of the dense population that we have in Tokyo, for example, C-RAN is already being deployed.

C-RAN is nothing new to the rest of the world, it's new to North America.


But we are also finding that in our metropolitan areas, we can replicate what we are currently doing with C-RAN in Far Asia in the US, for example.

Monica: In terms of the virtualization, the end target is to have all the network virtualized, and use generic hardware that can help support any processing.

Do you think it's possible to have COTS hardware for all the RAN functionality? There is a lot of debate as to whether we can go all the way there or at some point we need to keep specialized hardware. You still may need specialized hardware for some of the RAN functionality. What is your view on that?

Femi: You're absolutely right. Using generic servers to virtualize baseband functionality will not be completely possible, because of the kinds of

Centralized BBU and distributed RRH via CPRI interface over fiber and wireless links



Source: Fujitsu

data that need to be transported between the baseband and the remote radio heads.

My view is that we can virtualize much of the baseband, but we will still need specialized hardware when it comes to switching, for example, so that we can do high-speed switching of the I/Q samples that are going from the baseband into the remote radio heads.

It is not completely possible to virtualize the baseband without some form of specialized hardware. We will still need a little bit of specialized hardware to accomplish this.

Monica: Now, we talked a little bit about small cells before. Some people think that small cells and C-RAN don't mix very well. Others think that C-RAN is actually what small cells need to succeed. Where do you stand on that?

Femi: What is very important about small cells is that we are able to do away with all the large-scale network planning that needs to happen, which is what we have today in our current RAN deployment.

Today, in your RAN deployment, we do what we call network planning. You have to know where this RAN equipment is located, what it is doing, before you can introduce new equipment. If you don't do that planning properly, there will be all kinds of interference.

Now, with small cells and the way SON has been defined, the SON capability that small cells have is really going to help a lot. Because with SON, we arbitrarily place small cells anywhere we want to, and that can be managed from a centralized port.

My view is that C-RAN really needs a small cell to be more effective, because we can introduce remote radio heads, we can introduce small cells, such that they can work together interchangeably without having to worry about network planning as well as interference.

So you're right. A small cell is needed for effective C-RAN deployment as we go forward.

Monica: Let's go back to the operators. Short- to medium-term, what do you see that they're mostly focused on?

Femi: Operators today, they're mostly focused on integration. We know that C-RAN works. It has been deployed effectively in Asia Pacific. But the topology that we have in North America has a

unique mix of metropolitan areas as well as rural areas.

That integration of bringing in C-RAN into this topology that is peculiar to North America, where the challenge is, how do we migrate from what we have today into becoming C-RAN? Do we do it gradually, or go on and break what we have today and just start from a clear playing field? And the latter is not possible, by the way.

The third and final challenge that I see is, how much do we currently have in our transport to support what we are bringing in? A lot of studies have been done in terms of how do we do this for a new greenfield deployment. Do we marry it with our existing RAN architecture or RAN deployment, and is our core ready to support the kind of data that we're going to bring in?

Those are the things that the operators today are going through, fleshing them out as we get ready for the full-blown deployment within the next 18 months.

Monica: Then, as you say, transport is going to be a crucial part, because that's what enables C-RAN. If you don't have the fronthaul that performs at the right level, then everything else becomes a moot point. In that respect, there has been some debate in terms of using the CPRI interface: is it really necessary or not? What are your thoughts on that?

Femi: It depends on what is really good for the operator. In many of our metropolitan areas, even when fiber is available, it's very difficult to deploy –

whether it's a function of breaking new ground to allow this fiber to be placed or, in other areas, fiber is just not even available.

CPRI for transport is going to be very prominent in C-RAN. That's why, at Fujitsu, we don't want to limit the operator in terms of what can be used. In cases where fiber is not available, we have the option of either wireless transport, which is still CPRI based, and in places where fiber is available, the choice is whether you want to use active or passive fiber. For us, CPRI transport is crucial and is going to be an integral part of C-RAN deployment.

Monica: Now, in closing, may I ask you what you will be busy working on in the next two, three years at Fujitsu?

Femi: The future of C-RAN, as we go forward, is to allow resources to be shared among baseband units – shared effectively, and shared very intelligently. So, for us, the next logical step will be: how do we deploy software onto that baseband hotel effectively?

SDN will be very crucial, and NFV will be very crucial, because that will enable us to deploy access technology, aggregation technology, aggregation applications, that will really blow our minds as we look at the terrain of C-RAN going forward.

You can look for us to be very engaged in SDN and NFV, as well as the 5G technology as we go forward.

About Fujitsu



Fujitsu Network Communications, Inc. is a trusted partner to a broad spectrum of customers across all industries, enabling them to realize the maximum value from their communications networks. We are a market-leading US-based manufacturer of network equipment and a top US patent-holder in optical networking. Our solutions combine the best wireline, wireless, and software technology with extensive multivendor services expertise to deliver custom, end-to-end network integration and management solutions. For more information, please see: <http://us.fujitsu.com/telecom>.

About Femi Adeyemi



Dr. Femi Adeyemi is the Lead LTE Solutions Architect at Fujitsu. In this role, he is responsible for engaging wireless operators and partners in the Wireless and Small Cells Ecosystem. He is a 3GPP expert with design, development and deployment experience.

Femi has over 20 years of experience in the wireless and communications industry. Prior to working with Fujitsu, he led technology and product strategy for LTE and Voice in 3G/4G Macro and Femto product lines at Airvana Inc. He has also held various senior technical and management roles at Motorola, Engim and Elbera (a company he founded). Femi holds BS, MS, and Ph.D. degrees in Electrical Engineering.

Radisys

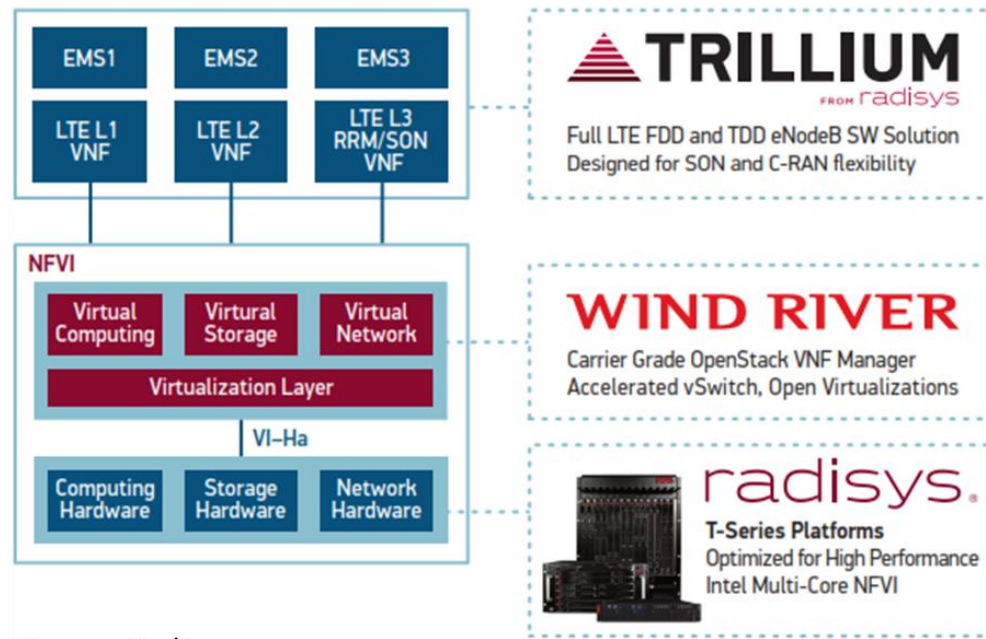
Radisys brings its expertise in developing application software and high-performance hardware to support C-RAN architectures. Radisys has been very active in supporting small-cell deployments and in promoting the nascent transition to NFV. The two activities that initially ran in parallel have now become the foundation for the company's C-RAN offerings.

In the small-cell space, Radisys offers the Trillium TOTALeNodeB, a turnkey solution for LTE small cells for SoC platforms that supports RRM, REM, SON and scheduler functionality. The TOTALeNodeB is the baseline for C-RAN solutions that will enable mobile operators to manage HetNets using remote baseband processing and more efficiently control interference and traffic loading.

Radisys has developed an edge router to help operators transition to SDN and NFV environments, by separating control plane and data plane, as an initial step towards end-to-end network virtualization.

On the hardware side, Radisys' T-Series Platforms offer telco-grade, high-performance software that is optimized for SDN and NFV environments, and targets C-RAN and virtualized RAN topologies. T-Series platforms are COTS, standards-based solutions that can use open-source software and comply with ETSI NFV requirements. The T-Series

Radisys solutions supporting C-RAN



Source: Radisys

platform ranges from the T-Series Compact rack mount server, to the T-100 Ultra solution with 100 Gbps interfaces and 2Tbps throughput.

Radisys collaborates with ASOCS, Intel and Wind River to optimize an application-ready SDN/NFV platform that can handle a frame within 1 ms, according to Radisys. The speed of network functions is critical to deploy C-RANs and to eventually achieve complete RAN virtualization.

The platform includes the following:

- Open Virtualization Profile (adaptive performance, interrupt delivery streamlining and management, system partitioning and tuning, and security management)

TRILLIUM
FROM radisys
Full LTE FDD and TDD eNodeB SW Solution
Designed for SON and C-RAN flexibility

WIND RIVER
Carrier Grade OpenStack VNF Manager
Accelerated vSwitch, Open Virtualizations

radisys
T-Series Platforms
Optimized for High Performance
Intel Multi-Core NFVI

- Wind River Titanium server, designed to support OpenStack for NFV deployments in legacy or greenfield networks
- Wind River accelerated vSwitch software
- Wind River carrier-grade profile for Wind River for Linux
- Intel processors
- Radisys' T-Series platform for C-RAN NFV
- Radisys' Trillium L2 and L3 protocol stacks, and Trillium Integrated eNodeB, a software solution that can be implemented in a variety of COTS hardware and that works with multiple vendors PHY.

Managing the C-RAN evolution, starting from small cells and DAS

A conversation with Renuka Bhalerao, Senior Product Line Manager, Radisys

Monica Paolini: Welcome to our conversation with Renuka Bhalerao, Senior Product Line Manager at Radisys. Renuka, could you tell us how Radisys got involved to work on C-RAN.

Renuka Bhalerao: Radisys has been active in the radio access network space for a long time. We have been providing software to the OEMs for the RAN, and lately we have been engaged in providing comprehensive small-cell solutions. We have been doing integrated small-cell solutions for a long time now.

With that as the context, we see C-RAN as a natural next step in the evolution of our product line. That's how we are getting into the C-RAN space and that is how we are seeing the small cells and the base stations evolving.

Monica: That's interesting, because some people think the C-RAN architecture is not very well suited for small cells. How do you see C-RAN and small cells related to each other?

Renuka: We think that the small cells actually do complement the C-RAN. The small cells have been providing the compact base station functionality near the edge; the C-RAN concept is about consolidating the resources to provide that same functionality in an efficient and more resourceful spending way.

Small cells actually transition into something like a C-RAN, where you have a cluster of small-cell functionality hooked together into a centralized, cloud-like environment. That's where you get to do the C-RAN.

We see them as technologies coming together nicely. We can give you certain examples here.

It is imperative to mitigate the interference in a network when you have a dense network with small cells. Algorithms used to mitigate interference need a lot of coordination at the L2 and also between the small cells themselves.

With the C-RAN architecture, we are actually putting that coordination into a central pool. Having to do that means that you're getting away from some backhaul restrictions. You are also coordinating better at the X2 and at the L2 scheduler levels.

You can also better implement a carrier aggregation feature for small cells in a cloud-like architecture.

We can also talk about a function called CoMP, or coordinated multipoint. The cloud architecture benefits small cells implementing CoMP.

Monica: You raise an important point, because small-cell deployments today are being held back by the need to manage interference. Operators are still trying to learn about this. It is becoming a more important issue, in terms of thinking about C-RAN, than the cost savings per se.

For instance, you can save on power with C-RAN, but the emphasis appears to be shifting towards the functionality benefits. Or towards the improvement in performance in terms of using the resources. Would you agree?

Renuka: Exactly. We would agree, in the sense that the end goal here is to provide the data usage to the end consumer at a cost – a capex and opex – that works for the operators.

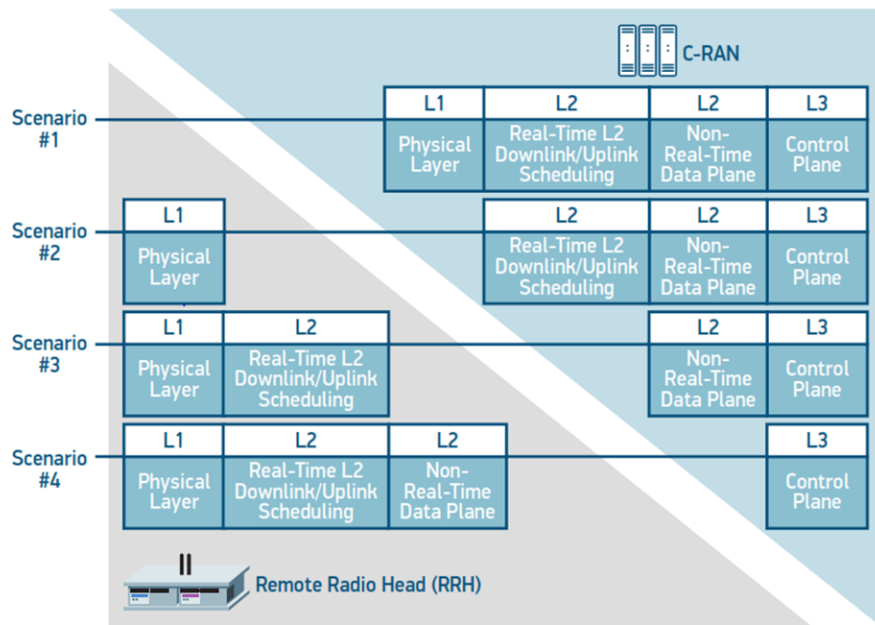
To meet that end goal, I think small cells, as well as C-RAN, will become part of that overall network transition. We have to see that the small cells will be something that will be part of HetNets. At the same time, the small cells will be evolving into what we can call the local C-RAN environment.

And when we have a multi-operator, multi-carrier environment, we will be able to drive more efficiency into small cells and also talk about DAS, and how it combines with the small cells.

With the small cells and DAS, the concept of distributed NFV eventually takes the form of a C-RAN.

Monica: How is C-RAN going to evolve? At the beginning, C-RAN might run on dedicated hardware, but the goal is to go to virtualized RAN

C-RAN implementation scenarios



Scenario #1

- All functions centralized

Scenario #2

- All functions centralized except PHY

Scenario #3

- All functions centralized except scheduling and PHY
- BBU, Time critical L2 co-located

Scenario #4

- Only L3 functions centralized
- L2 and PHY distributed
- BBU, L2 co-located

Source: Radisys

that runs on off-the-shelf hardware. In that case, do you think we're going to get to a point where all the hardware is going to be off-the-shelf, or are we still going to need specialized hardware for some RAN functions?

Renuka: Very interesting question here, Monica. You will see that the drive to C-RAN itself is coming from moving the base station into an NFV environment. And with that comes the fact that you have to separate out the software from the underlying hardware. And also, you should host almost all software functions onto a general purpose processor.

To that extent, we see that there will not be as much of a need for specialized hardware. This is a statement to be validated for technical feasibility, but a lot of R&D activity is going on to see how you can put most of that base station functionality near the core as much as you can.

Using the virtualization concept, you make VMs out of each of the components of the base station. These activities will be giving us some depth in the architecture that is coming together. The idea will be to not tie to specialized hardware but to merge RAN virtualization into the NFV concept truly.

Monica: Virtualization brings new flexibility. You don't have to virtualize all functions, and you don't have to have them perform in the remote locations. What are the options here for operators? What can they have in a remote location? What will be at the edge?

Renuka: When we talk about virtualization on the core network side, it's a different game than virtualization for the base station. The base station inherently – with the air interface on one side, specifically in LTE – has some definite latency and bandwidth restrictions. That's why there has been a debate about "Can we virtualize all of the base station functionality?"

The complete concept of C-RAN, however, is going to be at the end a complete virtualized base station, including L1 as well as the L2 and L3, and the application and management components. They all can sit as a VM.

Now, we can talk about certain deployment scenarios where not all the functions have to be virtualized. We can think of a case where all the base station is providing is the capacity and coverage. In certain cases, we can have virtualization only for the L3-and-above functionality.

Thus we can have the L2 and L1 stay near the baseband unit, functioning as in a small cell that is not completely virtualized but in a traditional environment. The L3 onwards can be virtualized in a cloud environment.

In fact, as operators move towards a complete virtualization, they are looking for a simpler transition initially, where they would like to test out the virtualization concept itself. In that scenario, this becomes more of a way to take the L3 and above into the cloud first, and then bring it slowly and gradually to where the complete function of the base station is now virtualized.

Monica: This is important, because operators need flexibility as they move towards a fully virtualized RAN. Also, a partially virtualized architecture reduces the requirements on the fronthaul. Will this facilitate and accelerate the adoption of C-RAN?

Renuka: When we consider separating out the L1 of the baseband from the radio unit, the fronthaul is needed to meet the LTE latency and bandwidth requirements.

In other cases, you can do away with that step. We have been doing some internal exploration of different ways of building the C-RAN architecture, looking at the tradeoffs of having L1 and L2 not virtualized, and the rest of the functionality being virtualized. And we compare this case with the one in which the fronthaul can take care of the required bandwidth, and we can do the complete L1 up to L3 in the cloud and rely on the fronthaul to provide the same type of performance that is driven today.

Monica: At Radisys, how do you help operators that want to transition gradually to the virtualized RAN?

Renuka: On one side, we are bringing the small-cell design into the DAS space, where we are helping operators see those C-RAN advantages at a smaller scale by combining the small cells with DAS.

On the other side, we are coming up with new partnerships in the C-RAN area, where we are taking on L2 and L3 software for the base station and putting it into the virtual environment, combining it with the L1 or the platform from partners such as ASOCS or Intel.

So with the combination of the platform and the L1 functionality that comes from our partners like ASOCS, we are bringing together a complete C-RAN environment. It is an ongoing effort at our end, and in 2015 that will be a big development that we will be coming in with. We will be testing it with some of our current customers in initial trials.

Monica: What this means is that operators can gradually add functionality into their location without having to do a major change.

Renuka: That's right. They need not look at it as a complex issue once they transition to the cloud.

C-RAN architecture scenarios features

Centralization – highest to lowest

	Scenario #1	Scenario #2	Scenario #3	Scenario #4
Complex fronthaul	Yes	Yes	No	No
Baseband pooling	Yes	No	No	No
Hardware accelerator pooling	Yes	Yes	Partial	No
L3 virtualization	Yes	Yes	Yes	Yes
Real-time L2 virtualization	Yes	Yes	No	No
Non real-time L2 virtualization	Yes	Yes	Yes	No

Source: Radisys

It's all about giving them different options which are tailored to their different deployments.

We can talk about outdoor hotspots, where you can use it as a local C-RAN. You can take the case of an enterprise environment, where you can use the C-RAN or a split base station architecture for deploying into the office environment.

And then you take it further and you have a complete cloud RAN that will be, not a replacement, but almost an equal functionality of a complete base station, at a significantly lower capex and opex.

Monica: What do you see from operators? What is it they are interested in doing right now when you get into trials?

Renuka: We see that they are very much interested in bringing that concept of virtualization

to the RAN side. We see that they are interested in actually trialing the virtualized core network as well as the virtualized base station together.

Operators are talking about cloud-based policy, cloud-based core management, cloud-based network resource management. It is a slightly new area for them, but they are open to seeing even the L2 and L1 functions of the base station move into the cloud. That will be part of these 2015 trials – not only in APAC, but also in North America and Europe.

Monica: In Asia Pacific, many C-RAN deployments have been mostly confined to greenfield deployments. I was wondering if you see any difference in operators depending on whether they have a greenfield deployment or if it's an existing network that they want to upgrade to C-RAN.

Renuka: We see action on both ends. We see some of our customers who are trying to bring the C-RAN into their existing deployments just because they are convinced of the advantages, because they are seeing that this is going to be a major way of bringing the evolution into the RAN.

At the same time, the greenfield deployments are always there to be built on new concepts. We see that in cases where we are trying to have the base station coverage put into new areas. The greenfield deployments will be benefiting from the C-RAN type of architecture.

Monica: Do you think that C-RAN is going to be a mixed solution – like DAS today – that works in

some environments but not in others, or is it going to be a more pervasive change in the platform and the topology used to run the overall network?

Renuka: In the long term, it is going to be your complete-evolution, next-step architecture for the RAN. In the near term, in the next two years or so, we will see variations. We will see that not every deployment has to be ripped apart and replaced with C-RAN. You'll have the capacity addition and the energy savings being taken into account, and the adoption of a smaller, local C-RAN concept. But eventually, it is going to be a new architecture for the base station.

Monica: Can we say that for 5G the C-RAN will be the dominating architecture?

Renuka: Nobody can say at this point what 5G will be. But the concept of 5G we agree on is about taking the network into that more efficient definition of it from an end-to-end perspective, be it the RAN, be it the spectrum management, the MIMO or the core network. In that definition, yes, we see that the C-RAN is what the 5G RAN is going to embrace.

Monica: If we were to look at the next few years, what do you see as being the major challenges that we need to address as an industry in order to accelerate the deployment of C-RAN?

Renuka: The challenges will be very similar to what we have with any new technology coming in. We have to find a way to make the deployments easier for the operators.

To that extent, the smaller steps of using the current small cells as the RRU – the remote radio unit – for the C-RAN, and then the general purpose processing unit taking care of the L2-and-above functionality, can show that the C-RAN architecture can fit into a current network without a major disruption. That will help build confidence within the operator community.

At the same time, we will have to come up with an alignment with HetNets. So we will have to structure the C-RAN architecture in a way that is not specific to LTE only, but takes Wi-Fi or 3G and other technologies, as well, into consideration.

Monica: In terms of timeline, it's not going to be a quick process. Can you give us some indication of what kind of timeframe we are looking at here?

Renuka: It is going to be a transition process. It's not an overnight shift, for sure. We are looking at the next two to five years. We could say that it's matching pretty much with how the 4.5G or the next-whatever LTE is coming into the network.

So you'll see a lot of activity going on in the next couple of years. The proofs of concept already demonstrated in the last year's trade shows are moving into pre-trials and limited trials in the networks. And then after this two years of proving the technology is over, we will see this expanding into the networks in the next three to four years.

Monica: In closing, can you tell us what we should expect from Radisys over the next two or three years?

Renuka: We are very excited, in fact, about 2015 from a C-RAN point of view. There are going to be certain trials that we will be part of, and major operators taking the concept to their networks. We will be banking on partners like Intel and ASOCS to bring the complementary technology to our Trillium software. Together, we are going to be

taking that from the lab into the field. We see it in the next two years evolving into a product line that we are proud of showing in an operator network and seeing deployed.

About Radisys

The Radisys logo consists of the word "radisys" in a white, lowercase, sans-serif font, centered within a dark red rectangular background.

Radisys helps communications and content providers, and their strategic partners, create new revenue streams and drive cost out of their services delivery infrastructure. Radisys' service aware traffic distribution platforms, real-time media processing engines and wireless access technologies enable its customers to maximize, virtualize and monetize their networks. For more information about Radisys please visit www.radisys.com.

About Renuka Bhalerao



Renuka is a Senior Product Line Manager at Radisys Corporation with her primary focus on Small cell Technologies for 3G and LTE and is responsible for the Trillium software portfolio. Prior to this, Renuka held the position of Principal Systems Architect in a customer-facing role specializing in Telecom Software and Systems. Renuka has 18+ years of telecom industry experience with expertise in wireless and VoIP solutions.

SOLiD

Over the years, SOLiD has developed RF amplifier, RF radio, and optical transport solutions designed to enable the densification of mobile networks and, hence, to support a massive but cost-effective increase in the capacity density through DAS and small-cell deployments.

SOLiD wants to ensure that backhaul and fronthaul do not become the bottlenecks in densification efforts. Instead, SOLiD aims to transform them into a scalable enablers that helps provide the capacity needed, where needed.

SOLiD solutions fit in organically to address the requirements of C-RAN topologies that include DAS and small cells, whether in indoor or outdoor locations. Working with SK Telecom in South Korea on the Smart Cloud Access Network (SCAN), SOLiD

has gained experience in the deployment of a commercial C-RAN deployment in a dense urban environment.

The premise at SOLiD is that densification, with its attendant need for a huge increase in backhaul and fronthaul capacity, requires not only a physical expansion of the optical transport network, but also a boost to its efficiency and a slash to its per-bit costs. Building out fiber is expensive and, as the density of the network infrastructure increases, scalability becomes an issue. As a result, SOLiD argues, a future-proof approach to densification requires operators to make the investment needed to optimize the use of the fiber assets available or being deployed.

SOLiD has been a leading proponent of DWDM, which splits a single fiber strand into multiple bidirectional channels to increase the capacity of fiber, multiplying the capacity of the link. Linear add/drops enable the operator to use the same

strand to serve multiple small cells, and to provide both fronthaul and backhaul.

The DWDM solution, INFINITY ACCESS, supports multiple protocols (e.g., CPRI, OBSAI and Ethernet) and can simultaneously support multiple access technologies (e.g., LTE and Wi-Fi) in the same strand. Operators no longer need to add a new fiber link when they add a new RRH or a small cell. And because DWDM allows operators to gradually add new links to the same strand, they can reduce their deployment and operating costs as they expand their networks.

The SK Telecom network includes 12,000 base stations and 80,000 RRHs, which are connected to BBUs in central offices or data centers. By using DWDM to increase the number of concurrent channels available, SK Telecom was able to increase the capacity of the existing fiber network without having to deploy new fiber where it was already available.

SOLiD DWDM INFINITY ACCESS



Source: SOLiD

Optical fronthaul as an enabler of scalability in C-RAN and vRAN architectures

A conversation with
Ken Sandfeld,
Executive Vice President,
SOLiD

Monica Paolini: Welcome to our conversation on C-RAN and fronthaul with Ken Sandfeld, Executive Vice President at SOLiD.

Ken, can you tell us what SOLiD is doing in the C-RAN and fronthaul space?

Ken Sandfeld: SOLiD is focused on the optical infrastructure for C-RAN, so the development and manufacture of unique optical network solutions at the physical layer that solve the problems that operators face when looking to build a C-RAN network topology.

Monica: What is it that you see, in terms of the C-RAN and fronthaul deployment? Because the end game is C-RAN. Fronthaul is what you really need to get there. What's the relationship between C-RAN and fronthaul?

Ken: The goal of C-RAN is to have the bulk of your processing done in a central location, and thereby increase the performance and the coordination of your assets right at the central office. The optical solutions provide the scale for the network, to be able to grow differently than previously was required.

With a C-RAN architecture, you're trying to push the assets back into the core, which means you have to change the way you deploy the topology on the access side of the network. This requires a different combination of network solutions, with the use of CWDM and DWDM. We have the capability of combining GPON, Ethernet and CPRI, all in the same fibers, thereby increasing the flexibility of the architecture.

We focused on the unique packaging of those solutions so that they better fit the network operators' requirements, because the traditional solutions just were not packaged and were not set up for that type of density. They were set up mostly for point-to-point links. This has enabled the operators to be able to design these networks and deploy them in this fashion. It really has to do with the number and type of services that you're trying to run across the cabling.

Monica: Why is it important to be able to mix and match different protocols?

Ken: It's very important, because no longer can you lease a dark fiber from point A to point B and say, "Well, that's all I need to do." The reality is there are going to be stops along the way. You might need to drop Ethernet for Wi-Fi. You might

drop Ethernet for a small cell on a pole, and that fiber asset may continue down that street and may also go into a building, where you might need to provide some broadband.

In the case of wireless operators in North America or Europe, that's not typically the case. In Asia, the same operator that supplies broadband to the building is supplying wireless, so it makes sense to have the ability to drop two different types of services.

However, in the end, the ability to mix and match those services on wavelengths is what's proving to be very powerful in the TCO of the solution.

Monica: This is a way to leverage your network resources in a more efficient way when it comes to C-RAN. Because what you often hear is that one of the major obstacles to C-RAN is that you're going to have cost savings, but the fronthaul is a big challenge because it's so expensive. How can you help with that part?

Ken: The upfront cost of the network – building new topology – is going to be higher than in the way we deploy today.

The reason it's going to be more expensive is that you don't have a ring. You're not trying to build a ring and then just add long access trunks off of that. You're essentially building multiple rings, and then you're building a whole new access network. That's expensive, because you have to put up new fiber in the ground, and you're going to go into locations where you've never been before. That's more expensive.

The benefit of doing so is that you have closer access to the nodes that you're going to use for future densification. You also have the physical-layer infrastructure, which we manufacture, that allows you to continue to add services without any additional fiber requirements.

By adding nodes out there – whether it be a small-cell node or whatever equipment is needed for the type of densification plans you have – the only requirement is a new card or the activation of additional ports into the ring. Wavelength services are what enable you to scale virtually on your fiber, and that's the most important benefit.

Not only can you scale in the quantity of ports, but in the type of service. Ethernet and CPRI are going to be interchangeable. You might even have to run both to a given node. You might have Wi-Fi at that same node as well. You may be supplying a radio head that's supplying service in the city for the Super Bowl, and at the same time, on that same pole, you're going to have Wi-Fi. You're going to need multiple services.

Monica: This sounds like a good TCO in the long term, although, on day one, you require substantial investment. Do you see that to be a challenge and something that is slowing down the deployment? Are operators ready to move on with this?

Ken: I've been watching this for three years, and it's just a very strategic investment. It's a big investment, and it requires a lot of change in the way things are done from a procedural standpoint.

I think those changes are still being made. Today, the procedures being used or the legacy procedures, legacy documentation, legacy steps for activating that particular device are one of the biggest problems. That's what's really slowing things down.

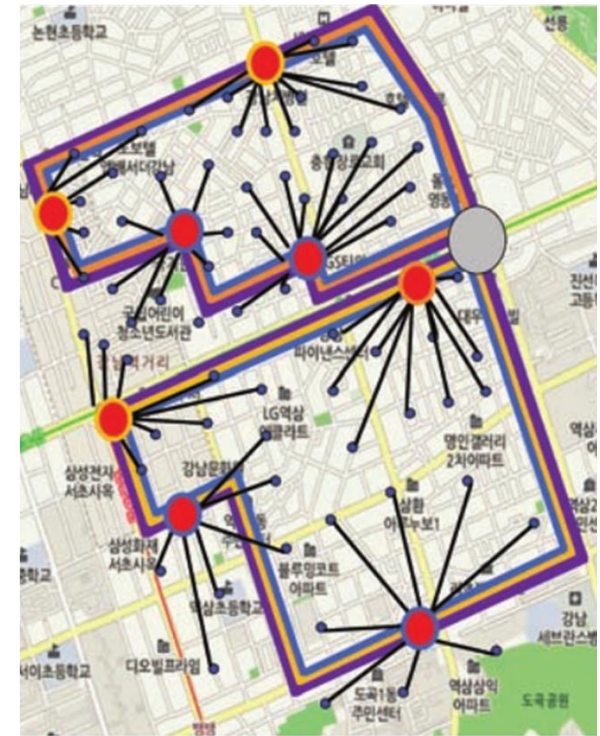
It's going to take some time for that to change, but I believe that this up-front investment is very strategic and requires a mind shift. It's just like the switch from 35-foot macro poles to 20-foot inner city poles. Even just by densification, you're effectively trying to cover less area. It takes time for everyone to learn how to adapt to those changes.

Plus we have a very robust macro network that we have to re-optimize. That's one of the steps in the densification with C-RAN: you have to reappportion the macro network to be able to accommodate all this densification. Actually, some of the macro locations will potentially even go away, because they're not suitable for that environment anymore. You're going to go from a big, tall pole and you're going to need lower locations.

Now, that doesn't mean that that physical location goes away. You may have a 20-foot shelter on it, you might actually repurpose it, and you might actually make that an aggregation point for all your fiber.

Monica: I wonder if you can comment on the differences you see in different regions. I know you do a lot of work in Asia, as well. What do you see that is different, in terms of the ability to have the strategic vision to move forward?

SOLiD deployment with SK Telecom



Source: SOLiD

Ken: Well, a lot of people – almost everyone in the US, the North American environment – will tell you, “It's different over there.” Some differences have to do with the rules on fiber and the access to new fiber, as well as the rules associated with that carrier providing multiple types of services on that. We have different rules here, and that causes things to be a little slower.

However, from an LTE deployment perspective, the problems are the same. That means the solutions are going to have to be similar. Regardless of how they deploy, and how hard or

easy it is compared to the North American market, the physical challenges are the same. The topology is the same. It has to go together, physically, the same way.

The investment required is going to be substantial. It may even change the entire model completely. It's like when the carriers started selling assets to the tower companies. They're saying, "You know what? We don't really need all those."

We may be going through a similar revolution with densification, because now you need to put up a 15-foot pole or a light pole on the street, and who wants to own that? Does the city own it? Does a third party own it? Does the carrier own it? Those are things that we're going through right now, but it does come back to the fiber usage. Fiber is going to be a very important commodity as we go here. Usage of the fiber is what we're focused on.

Monica: But in North America specifically, isn't there a difference in terms of, in the US, we are more spread out than in Asia and even in Western Europe?

Ken: Well, the investment is certainly bigger, because when you're densifying to that level, you have more square miles of coverage to densify. The US has always been your big, NFL cities that you're going to try to serve first. Where is your population? But the reality is the physics.

This isn't the coverage days. This is the capacity days. There is no way to deliver that capacity without having a pole near that user. A lot of my folks from Asia travel to the US, and they're always

amazed that LTE service is a matter of a couple of bars, whereas in Asia, there's five bars everywhere. The reason there's five bars isn't that they have more RF coverage blasting through the window. It's because they have more devices close by to serve the customers. That's why they have such strong RF coverage.

Unfortunately, there's no easy way around it, Monica. The reality is we're going to start at the city cores and we're going to have to build out from there. We first have to make the strategic investment in the cities, city by city, and changing the topology of our network. That's one of the reasons you see so many people trying to do wireless connections to these nodes, because they're like, "Well, jeez. I've got to do a gazillion of these things. I can't afford to run fiber to all these."

I think those solutions are going to be very critical in the build-out of these networks. However, all of these services ultimately have to aggregate to these points in the network. The aggregation points could either be the existing macro-node buildings or they could be the other CO offices in the city, but the reality is it's still going to be a big shift in the topology of the architecture.

There's just no way around it. If we don't, we're going to be here many years from now talking about this. There's just no way around it. You hear people say, "It's a physics exercise," and it really is. It's just like Wi-Fi in that perspective. If you're not close enough, you don't get enough bandwidth. It's the same thing.

Monica: With small cells, it becomes more challenging because, as you say, there is already the issue of ownership, and it's also the issue of where you put them. They need to be very close to the users. The closer you get to the users, the more complex it gets to get fiber there. Getting fiber on a lamppost is more difficult than getting it onto a cell tower. What can operators do?

Ken: Well, I think that's where in-building coverage is going to come in. Today you have, in the city, down-tilted antennas and all kinds of crazy things to try to provide some coverage inside buildings. This is going to change with in-building solutions coming back.

I believe the enterprise and the building owners are going to take some participation in that. Some buildings are going to utilize their bandwidth there to supply their services to their facility. Just like they supply Wi-Fi today, I believe there are going to be network solutions that allow them to provide future licensed-band coverage.

That's something that we're all working towards in this industry, in addition to the fiber networks. SOLiD is very focused on those solutions as well. The pipes to the buildings may be paid for by that enterprise to supply licensed and unlicensed services. There could be co-participation deals. Once all that IP is controlled by a gateway to the facilities in that building, you can decide who pays for what.

When you get back to the city street, that's when they require more ownership and more planning for the carrier to provide that. That's clearly an

outdoor, public space, but there's going to be fiber down the street. There's no other way around it, whether it's leased, owned, trenched or wireless, that connectivity has to come into play. I believe they're going to do it one block at a time. They're going to do it one building at a time.

Seventy percent of North America's buildings are not covered. The average facility that you and I travel to in North America has no coverage. If it does, it's an old system and it needs upgrades. It's designed for 3G. You may have a little service because they added an eNodeB to it, but there's not enough bandwidth or enough density on the RF to make it even work properly.

Regardless, we're talking about a virtualized core. That's really designed so that we can deploy additional nodes in the densification faster. It's all about lowering the costs and increasing the speeds to deployment. The fiber network equipment is getting cheaper.

If you have to buy the fiber, dark fiber, it'd be nice if you could do not just one thing. Maybe you can do 48 things with that fiber. That's what we're all about, is really trying to make more use of that infrastructure in a more efficient way to lower your TCO, given that, eventually, you're going to have to invest in that infrastructure, so let's make it more efficient.

Monica: What is your proposition for C-RAN, to operators going through the transition?

Ken: Well, in the beginning, the carriers are going to deploy basebands and radio heads for outdoor

solutions as a migration for the current macro topology. This has been going on for a long time, and it's only recently that the whole C-RAN concept has emerged. But that's already being done, at a more traditional level.

If we need fiber going from point A to B and you have these smart radios on top of the poles – that's not densification. That's just designing a more efficient macro pole. How do you duplicate that times 1,000 in that area? The reality is, you can't. You have to change the way you deploy.

Right now, they're going to continue to try all the electronics to make sure they work. They're going to be testing all the software. I think the software's probably one of the biggest reasons why it's been slow. They're trying to test all the software so that they can deploy these things quickly and have all the interference be managed automatically, so they don't have to do so much planning every time they make a little change. That's going on right now and has been for years.

Once you get to the next step, how we're going to be able to help them is, they're going to say, "Well, jeez. We don't need to deploy 10 sites here. We need to deploy 100 or 1,000. We know we have fiber in these areas." The fiber companies are going to become a real big part of that. They're going to take the carrier-owned fiber assets, they're going to take assets that they can lease or acquire some other way, and they're going to design new topologies from that perspective.

We're going to be able to provide the physical electronics and the optics to utilize all that

infrastructure. But it's not until they make the decision that they need to build their network that way that we really come into play.

One of the reasons we haven't really kicked off in a huge way in the North American market is that we haven't gotten to that problem point yet. Right now, carriers may tell you, "Well, I have enough fiber going to that pole. I have 12 strands going to that pole. What else do I need?" Well, it's not until they need to build 10 poles around that pole that they go, "Oh, my gosh, I don't have enough. I don't have any fiber." But they know they have the 12 strands going to that pole, so once they know they have that, they say, "Well, we know we have 12 strands going to this pole. There's a shelter next to it. We know we have connectivity at that point. Now, how do we get to the other 100 poles that we're going to put up around that?"

That was their way of looking at it. "Let's use what we have from a network. We've been investing in it for years. Let's try to see what we have and then kind of scale around that."

That's where the C-RAN optical network solutions really come into play, because you have 12 strands, so now how do you make better use of those 12 strands? The reality is, right now, those 12 strands are 100% full. Once we go to the C-RAN topology with the optics, we're probably only going to end up using 2 or 3, because we're going to take those 2 or 3 and we're going to create 48 virtual paths or more on every one of those fibers.

Now, all of a sudden, those 12 strands, we've got lots of excess capacity. That's where the TCO

model is just awesome. Because now, when the next year comes 5G and the next year comes whatever, now we have plenty of growth and scalability in our infrastructure to be able to add more channels into this architecture and grow however we need, up to a certain extent – up to the physical assets that you’ve acquired.

Monica: Let me ask, in closing, what should we expect from SOLiD over the next few years? What are your directions of development, in terms of products, trials?

Ken: As we’ve been working through all of the C-RAN possibilities in the Asia markets and actually learning with the carriers what’s working, what’s not, we’re modifying our products. We have a lot of new patents and new technologies that we’re looking to bring to market.

These aren’t just run-of-the-mill optic solutions. These are new, patented ideas and architectures that are totally different than the way things are deployed today. A lot of people commonly lump optics: “Well, it’s a piece of optic cable and you put stuff on both sides.” Yes, but how you package that and how you put that all together and how much you can do across a fiber is something of an art form and a science.

You’re going to see a lot of new, exciting solutions coming out of our factory. That’s obviously going to be combined with the coverage enhancement solutions on DAS, as well as other radio solutions that come from SOLiD.

We are actively looking at all aspects of the densification of the network. We’re focused on new technology as specifically geared towards making the C-RAN network deployments much more cost effective and much more powerful for the network operator. I think we’re right at the cusp of that starting. I think we’re right at that point.

About SOLiD



SOLiD (<http://www.solid.com>) helps keep people stay connected and safe in a rapidly-changing world through a portfolio of RF Amplifier, RF Radio and Optical Transport solutions. SOLiD enables indoor and outdoor cellular and public-safety communications at some of the world's best-known and most challenging venues including leading hospitals; professional, and college sports venues; government, university and Fortune 500 corporate buildings and campuses; international airports and metropolitan subways; and other high-profile sites. For further information on SOLiD DAS, Backhaul and Fronthaul solutions, go to www.solid.com or call 888-409-9997.

About Ken Sandfeld



As Executive Vice President, Ken Sandfeld leads the overall sales and product strategy activities for the SOLiD's portfolio of network densification solutions. Ken possesses over 17 years of experience in the wireless infrastructure industry and is passionate about bringing innovative technologies to market. Prior to his current leadership role, Ken held management positions at MobileAccess, Remec, Spectrian and Zyfer. Today Ken is focused on bringing SOLiD's leapfrog technologies out of incubation and into the market to solve some of the industry's biggest problems. Those areas include high-efficiency amplifiers for indoor and outdoor small cell applications as well as low-cost DWDM tunable optical solutions for the Enterprise and Wireless Operator markets.

III. Operator interviews

Oi

Optimizing the use of network resources with C-RAN

A conversation with Alberto Boaventura, Technical Consultant, Oi

Monica Paolini: Our conversation today is with Alberto Boaventura, a Technical Consultant at Oi in Brazil, and we are talking about Oi's work on C-RAN. Alberto, could you tell us a little bit about your role at Oi?

Alberto Boaventura: My main role here is related to strategy in technology. I'm focused on the access network, and I'm responsible for internalizing any kind of new technology.

Monica: At Oi, what are you working on with regards to C-RAN?

Alberto: We have been studying C-RAN for a couple years, and now we are talking with our suppliers in terms of identifying the state of the art. It's a technology that has huge potential to solve some challenges we are facing here, and eventually to improve the capacity or to minimize opex and capex. We

have some applications into which C-RAN will fit perfectly.

We are excited by the prospect that this technology will soon be available in mature products. Right now, in general, our vendors are aligned with our roadmap. Probably we will have something one year or a couple years from now.

Monica: Is C-RAN a niche application in a way that's similar to DAS and that works in only some environments, or does it represent a paradigm shift in how you run mobile networks?

Alberto: C-RAN can be applicable in several situations. DAS is doing the job right now in terms of improving the coverage, and C-RAN can do the same.

We need to address something that DAS does, for instance, in terms of neutral-host models. As soon as we have this, and the possibilities are RAN sharing (MOCN) and shared spectrum, C-RAN will be an alternative for DAS.

I think that the main contribution of C-RAN is to improve the capacity of the overall system, when we have situations in conjunction with macro-cell sites that would be a candidate to be a network data center to support the surrounding small cells. We use this kind of

site to improve the capacity in the neighborhood of the macro-cell site.

Monica: Where do you expect the increasing capacity to come from in a C-RAN environment?

Alberto: We think the C-RAN can act as a super base station, managing and coordinating resources in a huge number of cells. It can improve the development of functionalities to coordinate and mitigate the interference with technologies such as CoMP and eICIC. As known, both of them have specific synchronization and latency requirements, and with existing architecture (via the S1 and X2 interfaces), those are almost impossible to accomplish.

Monica: How important is it to have a C-RAN architecture to support small cells?

Alberto: C-RAN can minimize costs, or it can bring new suppliers. For operators, it can minimize the overall deployment costs by having a single type of hardware and using it with software from different suppliers. Is this some kind of a dream? This is a thing that I think we really need to have, and it is going to be feasible a few years from now.

The advantage we are expecting to recognize with capex and opex is the ability, eventually, to improve the network's elasticity. With C-RAN, we can balance traffic in different

regions in order to have the advantage of sharing available resources.

Today, in general, the deployment model that we have is for a fixed, high-rent license for a cell site. When we are talking about C-RAN, we can have a pool of resources that we can share with different regions. That's a view of functional virtualization that will increase the elasticity in general.

When we are talking about software-defined networks, or SDN, we have the possibility of using APIs to develop new algorithms or to improve operations, such as new SON functionalities, or eventually to minimize the interference, or to improve the capacity in general. As C-RAN can bring functional virtualization, and this in turn will bring network elasticity and minimize the capex, SDN will allow us to have more flexibility to develop and orchestrate virtual base stations.

Monica: How do you expect the move to C-RAN to take place? What are the intermediate steps? In a network that's already deployed, how would you transition from C-RAN to a fully virtualized RAN?

Alberto: At the beginning, rollouts will be done by niches, with specific cell sites. One of the first applications is indoor coverage, bringing an alternative for DAS in buildings, venues, stadiums and arenas. Thus, we can create in the same environment a network data center in order to take advantage of the existing

transmission resources (such as optical fiber, network connectivity, etc.) for further expanding to surrounding macro coverage.

Monica: For the long-term prospects of RAN virtualization, the plan is to move to off-the-shelf hardware for any function, but there are some questions as to whether that is going to be possible, or will we still need some specialized hardware for some of the functionality in the RAN. Is there a point at which virtualization of the RAN is no longer worth it or feasible?

Alberto: Today we don't have solutions – either COTS or ATCA – that can support this kind of open hardware.

The big challenge today is to migrate to open-source software tools, and have our traditional suppliers work with other hardware suppliers. A clear definition and standardization about a layer (hypervisor) to allow interoperability between base station software and hardware will be imperative.

Monica: Some people are really strong supporters of C-RAN, and others think C-RAN is more of a niche solution. The argument is that in C-RAN you need fronthaul and that is very difficult to secure, because of the latency and synchronization requirements. Is fronthaul going to be a challenge that may slow down the deployment of C-RAN?

Alberto: We need a new interface. Actually, CPRI, Common Public Radio Interface, has very tight requirements in terms of latency. The throughput requirements are huge, even for a MIMO 2x2. There is some good work that we are looking at, using compressed CPRI over IP, and GPON.

I think that the best way to do this, from an operator's perspective, is to have an open interface, as ETSI is standardizing in the Open Radio Interface, ORI. We would like to see a mix of the advantages of CPRI and OBSAI, and ORI can pick up these advantages.

Monica: A common assumption is that you need to have fiber and a CPRI interface, and that is being challenged in many ways. What is your view?

Alberto: When we are talking about the capacity, we are talking about fiber. For instance, in 5G, 10 Gbps – or at least 1 Gbps – will be a common requirement for the air interface.

As the incumbent for fixed line, we are in the right position now. Our footprint in terms of fiber is good. However, there are some solutions that will improve the capacity, and they will be very important to operators until we reach a point at which we have fiber everywhere.

One solution, multipoint-to-multipoint, will allow multiple connections instead of two,

and can work at the beginning, when throughput requirements are not too high. We could only use this kind of solution for some specific applications at this time. Links using the 28 GHz band and eventually e-Band may be an alternative for fiber or transmission.

Monica: In closing, let me ask you about the cost-saving question. Initially, it seemed like the major advantage of C-RAN is to reduce cost, but there are also, as you mentioned, improvements in capacity and resource utilization. If you were to look at what is going to drive the adoption of C-RAN, is it going to be mostly the cost savings or the improvement in the performance?

Alberto: The benefit of C-RAN is that it reduces capex and opex. But there are also functionality improvements that are feasible only with C-RAN implementation.

With C-RAN we can have a central scheduler handle several cell sites, which is different from traditional handovers. This will improve the capacity for cells in the C-RAN.

Another big advantage is having more elasticity in terms of network resources. We will be able to move resources to other applications that may not even be related to C-RAN, but that we can have in conjunction with C-RAN – for instance, to another network element such as RAN cache, which can

increase the customer experience by minimizing the latency in accessing network content.

Eventually we will simplify the RAN and core network architectures and avoid having too many interfaces. We are talking about simplified solutions in terms of maintenance, in terms of people we need to train, instead of having specialists for different functions. As the overall system will be simplified, we can improve the capacity, streamline the operations, and minimize opex and capex.

Orange

The interrelation of fronthaul availability and C-RAN requirements

A conversation with
Philippe Chanclou,
Team Manager Orange Labs,
Orange

Monica Paolini: In this conversation on C-RAN and fronthaul, we talk to Philippe Chanclou, Team Manager Orange Labs, from Orange in France.

Philippe, can you tell us about the work you're doing on C-RAN at Orange?

Philippe Chanclou: I work to understand the possibility of achieving centralized RAN in the mobile cloud. I'm focused on how the fiber and microwave fronthaul can establish the link between BBUs and RRHs.

Monica: Orange operates in many countries. Do you have different C-RAN activities in these markets?

Philippe: In Orange, we have several countries where we have deployed different radio access

networks from different vendors. We tried to get a general view of the evolution of C-RAN networks, including centralized RAN and cloud RAN.

We have made trials where we try to understand what would be the role of fronthaul to achieve these links and how we need to build it with the function of technical aspects and economical aspects.

Monica: What is it that you're trying to achieve with C-RAN in the long term?

Philippe: This is our learning phase of the trial. The learning phase allows us to understand what would be the impact of this cloud-RAN concept on our information system rules, and in the business case evolution.

Our short-term view has to be consolidated with a long-term view, and the fact that we have 5G arriving, we have CoMP arriving. We try to understand what could be the target cluster size of antenna sites with a BBU hotel and how that could be optimized in terms of performance.

Monica: Is C-RAN as a niche solution for some environments, or a new way of organizing and architecting your network?

Philippe: In our initial understanding, we consider C-RAN to be something that influences operators to build antenna sites that can reduce the time to build the network, the number of network operations, the expertise and knowledge that we need to have on the antenna site, and the footprint. All these things enable us also to deploy

data centers – or BBU hotels – in the long-term to operate the RAN in a clever way.

Monica: The demanding CPRI fronthaul requirements may slow down the deployment of C-RAN or limit its potential. What role do you see CPRI playing in fronthaul?

Philippe: We consider CPRI to have a valuable interface for now. All our RAN vendors build complete RAN solutions that include CPRI fronthaul. We try to achieve links over microwave, over fiber, using the CPRI interface, without affecting the software or hardware of our radio access technologies – a wonderful interface ready to deploy.

We are trying to have something that is operationally ready, supports any evolution of radio access technologies, and is futureproof to the coming 5G. We consider CPRI has the most relevant interface for now and that the considered optical and microwave transport solutions are compatible to any future network evolution of CPRI or something else.

Monica: If you can use wireless for fronthaul, you can have a C-RAN topology even where you do not have fiber. How important is this for small cells?

Philippe: Fiber and microwave are not mutually exclusive. In some cases we use only fiber, in some cases we use only microwave. Sometimes we could have the combination of fiber and microwave. They are two main solutions, and in different countries or markets you may use one or the other.

Monica: In the long term, do you see CPRI as dominating the interface for fronthaul, or do you see other interfaces becoming more important?

Philippe: We consider that CPRI is a very mature interface, and we are able to build fronthaul networks based on CPRI as it is.

Nevertheless, in studies at Orange Labs on the long-term, we try to understand if we need to study a new functional split between BBUs and RRHs, and define a new fronthaul interface. We also try to understand if Ethernet functions could be used inside the BBU and the RRH, as well as inside the CPRI switch element.

We do not think that a regular Ethernet network (including switch/router shared with existing Ethernet traffic as backhaul) is the adequate and optimum network for existing CPRI transport. But the dedicated CPRI switch for fronthaul that we have between the BBU pool and the different transport segments to connect the RRH pools could be based on some mechanism coming from Ethernet. The capacity to use 100%, or 10%, or 50% of all different functions coming from Ethernet is something that will have an impact on the evolution of the cloud RAN with BBU pooling.

In the long-term view, the evolved CPRI (based on an existing or new functional split) could be defined natively as an Ethernet interface to be compatible with switching and transport functions of an Ethernet network. For the short and medium term, we have to support CPRI as it is, without embedding it over some other protocol (e.g., CPRI-over-Ethernet), because we consider that several

flavors of CPRI exist, each with its own requirement constraints (e.g., latency, asymmetry, synchronization and bit error rate). The CPRI encapsulation protocols must be agnostic and robust to support any of these flavors and requirements. The evolution of the RAN, in terms of hardware or software, could have an impact on the CPRI frame, and so we have to be sure that our fronthaul segment is compatible for any evolution and multi-vendor compatibility.

We also have to consider the evolution of this fronthaul for 5G tomorrow. We don't know what this interface will be. It could be CPRI with a high bitrate or something completely different, but we want to be sure that our fronthaul network choice today will be compatible. We consider that our technology choices to achieve this network segment with fiber (based on passive infrastructure, with optionally low latency active equipment) and microwave are natively compatible with any future 5G implementation, and for that the solution is futureproof.

Monica: In the benefits that C-RAN brings you, there is clearly a cost component, but do you also expect improvements that derive from using your RAN resources better?

Philippe: The advantage for us is to create a cluster of BBUs connected to the RRHs, and the fact that we have a BBU hotel. We know that it will simplify the antenna site, because we will have a lighter antenna cabinet and no equipment that needs configuration and maintenance.

We've tried to study, also, what will be the size of this cluster of BBUs, and how many RRHs, and how many antenna sites we will connect to one location, where you host the different BBUs' ports.

We try to define this cluster's function in terms of security, in terms of performance, in terms of the latency that is acceptable in the link between BBUs and RRHs.

In terms of security, if we look at what we have done in fixed networks, where we try to consolidate the fiber to the home central office – the first aggregation point – we consider that the first aggregation point in our access will be no more than 50,000 customers. That will define the maximum size of a cluster and could be lower than this value.

Secondly, we know that CoMP is very efficient, when you connect the first cells, for suppressing the interference. That is for radio performance. Another performance measure is the capex efficiency, based on hardware savings. Some studies say that when you have about 100 cells together at the same BBU location, you reach the peak of gain in hardware saving.

The third thing is a function of the topology of your country and the topology of your urban area. The connection between BBUs and RRHs allows for 15 km of distance (75 μ s), so how many antennas could you connect to a local C-RAN?

Monica: How can you manage the transition to a C-RAN topology in a legacy network?

Philippe: The option we take in the evolution of the C-RAN is to define the centralized RAN for all radio access technologies in one antenna site. We don't focus on C-RAN for only 3G or 4G; we try to focus on centralized RAN on the antenna site.

If you have 2G, 3G and 4G at this antenna site, you try to define a fronthaul solution that supports the link between the BBU hotel and the RRHs for all mobile technology generations. This opens the way to building a solution that will also be compatible for 5G in terms of available resources using either fiber or microwave.

Monica: How do you plan to transition from a local RAN to a centralized RAN?

Philippe: We try to define some different ways to introduce centralized RAN. In the local C-RAN, there is a regular antenna site, where you have your BBU in the cabinet with power and the backhaul element.

When you want to add a fourth sector or connect a small cell to these BBUs, instead of introducing a new BBU and backhaul equipment, you install only RRHs to achieve these new cells. The local fronthaul allows you to share the existing BBU equipment with the existing macro cell.

In the local C-RAN, the microwave fronthaul could be an efficient solution to connect the small cell to the macro's sectors.

The second phase is no longer local C-RAN, but a more general, centralized RAN: you put all the BBUs not at the macro site, but inside the data center or central office. In this case, you use some fiber duct links to connect the several different antenna sites, which could be micro cells or macro cells. We study the way to reduce the number of fiber links needed by using wavelength-division multiplexing, or WDM, to achieve on the same fiber the parallel transmission of several CPRI links (i.e., colored links).

Monica: Is C-RAN going to be fundamental to 5G?

Philippe: We don't know exactly what 5G will be, but we know that in terms of capacity and in terms of the way to achieve an efficient and constant bitrate to the customers in the cells, we need to rely on CoMP. We know that efficient CoMP processing requires that all the BBUs are in the same place. And fronthaul is the solution to achieve this in an efficient 5G network.

Monica: What is the time frame for C-RAN adoption? This is a big transformation in the way

you run your networks. How long do you think the process is going to take for the industry as a whole?

Philippe: In terms of RAN, we have been ready for that for a long time, because the CPRI interface has existed for several years. We have been able to put some fiber between RRHs and BBUs for several years. It is a way both to transform and to accelerate the deployment by operators. All vendors are already ready to do that.

We need to have fronthaul solutions that are cost efficient, that have performance gains, and that make new antenna site construction and the modification of existing antenna sites easy. We also need to know what will be the impact in terms of RAN sharing. All these things have to be understood before we deploy C-RAN in our networks.

SoftBank

Cost and performance advantages in a C-RAN commercial network

A conversation with Hiroshi Kawai, Director, Technology Development Dept. and Wireless Technology Dept. Technology Division, SoftBank

Monica Paolini: Hiro, there is a lot of activity on C-RAN in Asia in general, and SoftBank has in many ways been leading the way. Could you give us an overview of how SoftBank has deployed C-RAN in Japan?

Hiroshi Kawai: There is a cable communications provider called NTT. Their fiber network is spreading radially from the big cities to small towns in every corner of Japan. We can rent their fiber at a reasonable cost, and also rent space in exchange stations and set up devices there.

A fundamental structure of Wireless City Network [SoftBank's TDD-LTE network in Japan] is putting BBUs together in an exchange station and connecting RRHs, which are set on the outside, through NTT's dark fiber. We have been making

the most of NTT network assets to build our C-RAN network.

Monica: What advantages did you see from deploying C-RAN architecture, from a financial, performance and operational perspective?

Hiro: We had to control interference carefully through cell RF planning design before, but cooperative control between cells now deals with it, so that we do not need to spend too much effort on cell planning. That also gives us flexibility in placing RRHs.

Furthermore, with C-RAN the wireless network can be optimized more efficiently compared to ordinary SON. This is because changing the configuration of base stations to upgrade quality in the wireless network is enabled on the TTI level, and the network can keep on top of sudden fluctuations in traffic.

From a financial perspective, C-RAN leads to cost reduction, just like introducing SON does.

Monica: What do you use for fronthaul for C-RAN where fiber is available? CPRI?

Hiro: That's right. As I explained earlier, Wireless City Planning's devices connect BBUs with RRHs through fiber, and we use CPRI to do so.

Monica: What do you use for fronthaul where fiber is not available or where it's too expensive?

Hiro: If the dark fiber is not available, we set up BBUs in the same site as the RRHs and install an L2 line, such as ADSL or satellite.

Monica: What is your view on small-cell topologies using a C-RAN approach (i.e., using a remote-baseband model for small cells)?

Hiro: The conceivable uses of small cells are to provide coverage in dead spots and to build HetNet topologies. The former has interference problems between adjacent cells, and the latter has the problems that come with micro cells. It is quite essential that the small cells work as a part of C-RAN, and that is why we connect the RRHs of small cells with BBUs through CPRI, like we do with macro cells.

Monica: How valuable is using a C-RAN approach with small cells to manage interference and do load balancing?

Hiro: We believe using C-RAN with micro cells is also absolutely valuable, for the reasons mentioned above.

Monica: How much of a challenge is it to ensure fronthaul for small cells?

Hiro: It is important for even small cells that the interference of adjacent cells is controlled, just as it is with regular cells. We were already thinking to use C-RAN, and dark fiber for fronthaul, when we were considering whether to implement small cells. Dark fiber is available almost everywhere here in Japan.

Monica: In your C-RAN deployments and trials, have you looked at selecting only some functions (e.g., only the control plane) for centralization/virtualization?

Hiro: No, we haven't.

Monica: What are the challenges that the industry still has to face to expand C-RAN deployments?

Hiro: When implementing CA with five-component carriers (5CCs) or massive MIMO in the future, the system bandwidth per site and the number of antennas will increase, and the necessary bandwidth of the link that connects RRHs with BBUs will increase accordingly. Then the current CPRI will need multiple fibers, which leads to rising costs for these reasons. A challenge we are facing now is increasing the bandwidth between BBUs and RRHs. And we have to reduce transmission delay, especially when using C-RAN in a wide area over BBUs in different municipalities.

Monica: What are your future plans for C-RAN?

Hiro: Our current business goals are to extend the power of cooperative control and to expand the network area for further improvement in spectral efficiency.

Telefonica RAN virtualization to achieve sustainable growth

A conversation with Javier Lorca Hernando, Project Manager for the RAN, Telefonica R&D

Monica Paolini: Good afternoon, and welcome to our conversation with Javier Lorca Hernando. He is the project manager for the Radio Access Network at Telefonica R&D in Spain.

Javier, can you tell us what you do at Telefonica?

Javier Lorca Hernando: I take part in all the innovation activities dealing with radio access networks in Telefonica R&D. We are a consultancy company within the Telefonica group, but since several years ago, we are more functionally integrated into the global CTO office in Telefonica, trying to coordinate all the network operations across the world, across the Telefonica footprint.

We conduct several activities related to innovation at the radio access network level.

The first activity will be all the things related to 5G: setting up all the future requirements and virtualization. Virtualization at the RAN access network level is an area of interest.

Our focus is on assessing the technology, because it's not yet mature. We are looking at changing the way in which the radio access networks are being deployed from a completely distributed and dedicated hardware point of view, like today, to a completely virtualized and general purpose hardware point of view in the future.

5G and virtualization are intimately linked. When you look at the different alternatives for our future 5G deployment, everything revolves around virtualization in some way or another. Even if the term virtualization itself for the RAN is not so clear right now – because we are looking at the first approaches, like cloud RAN, centralized RAN and so on – everything is tied to virtualization in some way or another. Running networks on general purpose hardware, virtualizing the RAN, making the network completely flexible, and splitting the network functions in different forms and different ways are among the most important activities for us at Telefonica.

Monica: At Telefonica, you are at the forefront for virtualization. How is the virtualization of the RAN inscribed in the overall virtualization process?

Javier: Up to now, I would say the RAN takes most of the opex, most of the capex – a great deal of the network expenses take place in the radio access part. If you look, for example, at the evolution of the different mobile networks, even if they are all 3GPP based, you change from 2G to 3G to 4G, and now we're looking at 5G. The model itself is not sustainable, because you have to increase your site capabilities to cover all the legacy infrastructure compatibility for even 2G terminals. For example, in Latin America, a lot of devices are still based exclusively on 2G, so you must support 2G, and you must support all the seamless operation, handovers and interoperation between 2G, 3G and 4G.

Clearly, the model is not sustainable anymore. Now we are seeing a very good opportunity that a completely different model may appear with 5G. With 5G, rather than looking at an increased throughput and reduction in latency, we can change completely the way in which networks are deployed. We are looking more closely at the economic sustainability, flexibility of deployments, not relying on a very rigid model for the radio access part, and changing the software model and the hardware model for this kind of network.

We are looking at moving from dedicated hardware to general purpose hardware. In many aspects – not in everything, but in many aspects of typical standards – RAN functionality is mainly based on software,

especially when looking at the MAC layer and above.

This is the major driver of RAN virtualization. People don't care about peak throughput as much as before. For example, if you had LTE-Advanced, you would have theoretically the capability to do 4 Gbps with 100 MHz of spectrum. Currently, no one delivers 4 Gbps, and it's not clearly the case that 5G needs, let's say, 100 Gbps.

The point is that the network should be adaptable so that it supports, for example, massive operation of sensors of M2M communications, and so that sections of your network in which delay should be minimized to give priority to reducing latency, as opposed to other areas of your network where the main driver is the reliability, for example.

Or being able to flexibly allocate the RAN functions in different locations, or in different ways compared to what we did in the past.

The main driver is the change in the model so that we can sustain the growth in network usage, in people using advanced services, and in creating new services completely different than before. Examples of these services can be smart-grid applications, health applications, security, or applications for disasters, emergencies or recovery.

A completely new model has to emerge. The only way to be sustainable from an economic point of view is to change the actual topology of the networks and make them much more flexible, especially at the RAN level, because the RAN is one of the most expensive parts of the network.

Monica: You're absolutely right. You want to optimize the network resources that you have, rather than just go improve your KPIs. But as you do that, you need to be much more sophisticated in the way you manage the functionality.

In this regard, you may choose to push some of the core functionality to the edge, or have the RAN functionality centralized. With virtualization, mobile operators have a choice between what they want centralized and what they want distributed.

Javier: Yes, exactly. For example, one of the main ideas behind the 5G architecture is that centralization should not be taken to the extreme, and distribution should also not be taken to the extreme. Everything in between can be valid, depending on your actual situation.

Your network should be able to flexibly change from one deployment to another without having to change all your network topology, network nodes and everything. That's the key point – centralization is not the best way to proceed, and a completely

distributed approach is also not the best way. Everything in between is, in general, applicable, depending on your situation.

For example, in 5G, there are very different requirements depending on your future applications. You may talk about sub-1 ms end-to-end delay for specific services such as tactile internet or even gaming. In that case, centralization doesn't make sense, because the fronthaul network introduces a long delay and the processing delay at the baseband unit is sometimes unacceptable.

Perhaps in these cases, it's better to have a completely distributed network, where your intelligence is at the edge, as you mentioned, and you have some part of the EPC or core network functions at the edge, and maybe even some applications or part of the OTT service residing there, very near to the user, so that operation is simplest.

But when trying to deploy a massive M2M deployment with a lot of sensors that requires a very low-cost access to the broadband services, the model can be completely different, because a sensor should not require a latency below 1 ms. Or perhaps, yes, it depends on the case: with smart-grid applications, it is definitely the case to have a very low delay.

So everything in between centralized and distributed topologies should be enabled by this kind of virtualization. You instantiate

virtual machines, rather than deploying specific nodes. Your RAN should be split into different atomic units, and you could activate or deactivate them upon demand without changing the basic topology chosen by the network operator and without requiring new investment.

The ultimate goal is to have a mixture of dedicated hardware, general purpose hardware, also virtual network functions that you could activate or deactivate upon demand, without having to change the topology. That creates new requirements for the hardware, and also for the network links themselves.

This also requires not relying on CPRI anymore, because we believe it's not a good concept for the future, it's not futureproof. But also not relying on the actual network transport being exclusively based on packets or transport blocks, like in 3GPP. We have to be much more flexible in terms of carrying a mix of digitized bits, packets, or other kinds of structures. It's a dramatic change in the way we deploy the network.

Monica: The idea of having generic hardware is very appealing. But does it apply to all the RAN functions? Or are there some RAN functions that you don't think will really lend themselves to full virtualization in the sense of using general purpose hardware?

Javier: So far, clearly it's not the case. General purpose hardware has a lot of limitations. One of them is the processing power. Another limitation is the power consumption, because it's a different model. It's a serial model where the CPU runs a different set of instructions, as opposed to a dedicated hardware model where everything runs in parallel, more or less.

In the future, there should be a gradual evolution from virtualizing only the higher layers, like could be today, to virtualizing almost everything that makes sense to virtualize. But there may be things that don't make sense to virtualize. For example, performing Fourier transforms, or some matrix operations, or very basic, physical layer activities that could be located perhaps at the remote site, or perhaps at the baseband unit, or anything in between. You get a hybrid model where you could have different kinds of hardware, but under the same operational model.

I would say, for example, you could instantiate a virtual machine that could be completely virtualized, or have part of it virtualized, part of it on dedicated hardware, without having to know it in advance. Because probably the amount of dedicated hardware will have to be reduced in the future, in favor of a more virtualized processing.

That same kind of operation, based on software, would take place irrespective of

whether your network is very much virtualized, or only specific parts of your MAC layer or your physical layer or anything in between. The model should be clearly the same model as is now in IT. You are not so much worried about the underlying hardware, and you could replace your hardware whenever it fails or you have to upgrade your capacity. You could change your hardware seamlessly, as you keep the same software level. For example, you instantiate more virtual machines, relying on more processing power or more memory, more storage or whatever, so that it is sustainable.

I would say that in ten years' time, it will be different than what happens today when you deploy a network initially, for the first two or three years, and after two or three years, it collapses. Everything is saturated, you have to upgrade your nodes, buy new software licenses, perhaps buy more spectrum, and increase the number of sites. That is not sustainable. In the future, this change should, ideally, be completely made on a software level.

Monica: Another thing you said is that you don't believe CPRI is the right interface. Can you tell me more about that?

Javier: CPRI is reused from the original scope, where it referred to a single base station where you wanted to separate your remote radio head from the baseband unit, to link them with several meters of digital link in

which the I/Q samples were digitized, and you wouldn't suffer from the cable losses in an analog deployment. It's from a single-base-station scope.

When you reuse CPRI in a wide-area-distribution scope, you have several limitations. The first limitation is that it is not IP traffic, it's just I/Q samples, which are completely, tightly synchronized. A reduced delay and jitter is required for that operation to be up and running, and the equipment is much more expensive because it doesn't rely on IP traffic.

The model is very good so far, but for future evolution of a virtualized RAN it doesn't make sense. The future evolution of a virtualized RAN should rely on a much more relaxed Ethernet fronthaul network, mainly based on packet switching, IP switching, making the necessary amendments, modifications and changes to carry the necessary traffic in the fronthaul, but relaxing the delay and throughput requirements. That way, on the one hand, the cost can be very much reduced, and on the other hand, you can do some kind of processing in your fronthaul network, depending on the splitting point that you want to have.

For example, if you want to have the complete physical layer, or you only want to have part of the physical layer running in this fronthaul network, or you want to deploy only higher-layer bits, like the MAC layer transport

bits, you can rely on the same basic Ethernet fronthaul network. You don't have to deal with multiple-Gbps constant bitrate traffic, which doesn't scale well when you increase the number of sites and deal with compression and multiplexing and so on.

A different model should be there. I recognize that there are several initiatives – most of them are now proprietary in the sense that vendors are promoting their own standards based on Ethernet – but they are not interoperable with other vendors.

That model does not work for us, because we believe that in the future we will have, for example, one vendor for the remote sites, others for the BBUs, others for the hardware, etcetera. Different standards should emerge from that, and we believe it's possible, because CPRI is the first approach, very basic, very important from the digital world from the inside of a base station.

Dealing with a completely distributed environment, you can relax those specifications and allow, for example, higher delays, lower throughputs, and some statistical multiplexing of traffic, which is not possible today in the CPRI links. I would say that in a few years' time, something should arrive in this sense.

Monica: This may make the implementation of certain architectures for small cells much more feasible. Because right now, fronthaul is

the bottleneck for C-RAN topologies in small-cell deployments.

Many people think, "You cannot do C-RAN in a small-cell dense environment, because you don't have fiber, or, if you have it, it's expensive." To have lower fronthaul requirements may help small cells as well. Is that right?

Javier: Yes, absolutely. For example, in Telefonica, the adoption of small cells has not been very wide, because of the problems that the small cells pose to the macro network in terms of interference when talking about co-channel deployments, which are usually the case in our operations where small cells use the same spectrum as the macro network. Network coordination between the macro layer and the small-cell layer is difficult, in general, and it relies on a complex interchange of information, a complex setup. It doesn't scale very well with the number of small cells.

This kind of centralized deployment at the macro site controlling a number of small cells would be a great idea, if it weren't for the need to deploy such an expensive fronthaul network connecting the small cells to the BBU or to the macro.

When you need not only fiber, but expensive dedicated fiber or dark fiber for the CPRI links between them, it doesn't make a lot of sense to deploy small cells, even though they

themselves can be quite cheap. With lower-cost fronthaul, small-cells adoption should be much higher, because you would coordinate transmission without having to do Almost Blank Subframes.

For example, if you want to do CoMP, C-RAN makes it much easier to centralize scheduling decisions or to do joint transmission between the macro and the small cells. But in any case, the hardware should change in such a way that, for example, not everything is centralized at the macro site, because that would require, again, the same CPRI fronthaul.

This architecture that is dominant in 5G is a driver for small-cell adoption. These are ultra-dense networks in which you could have small-cell sites as close together as every 20 meters or something like that. In that case, it doesn't make sense to deploy CPRI links every 20 meters to the macro cells.

Completely different models should be there, and I think in the future that will be an integral part of our RAN: a complete mixture of macros and low-power small cells, either coordinated or not, either distributed or centralized, with this kind of flexible fronthaul network, mainly dealing with Ethernet transport.

Monica: If you look at the big picture, the long-term prospects for the transition to virtualized

RAN, how do you see that happening? There is a lot of learning to do. How do you see this transition going over the next ten years or so?

Javier: I would say that the first thing we may expect to happen is to replace today's dedicated hardware with general purpose hardware, seeing how the software should be changed, should be developed in that sense – the drawbacks and the benefits of that, be it power consumption, the actual cost or whatever.

Simultaneously, we need to try to change the standard for the fronthaul network. This will probably take much longer, because it requires taking it to a standards body – the ETSI, 3GPP. I don't know what would be.

The first step, I would say, should be changing the hardware, and making some kind of original or unexplored split points in the future, that have not been tested so far, to see whether this brings benefits or not: What are the drawbacks? What is the fronthaul network required? And in parallel, the standardization of the fronthaul network, just to make things easier – not relying on CPRI anymore, but relying on the standard Ethernet fronthaul network.

In this sense, at first, the smallest vendors were the most active in trying to change the way that things are being done. Now, the

large vendors are very much into this idea too. Because even they, apart from seeing the need to be aligned with the other, smaller vendors to be part of the industry, they also see that relying on dedicated hardware and dedicated software and so on is not a good idea for the future.

Even their development costs and production costs could benefit from this kind of virtualization. Even if they set it as a black box, if they rely on general hardware and more or less general software – for example Linux-based or open source software – they will see the benefits for sure. Especially when we move into 2017 and 2018, when probably some kind of 5G standards will be decided in 3GPP and other forums, this kind of architecture should be much more concrete, and should be much more precisely detailed by the vendors themselves so that they can propose, for example, the separation of the hardware and the software, and the control plane from the data plane. Today this is a wish list from operators, but in five years' time it should be a reality.

References

Cheko, Aleksandra et al. (2014) Cloud RAN for Mobile Networks – a Technology Overview. IEEE Communucations Surveys & Tutorials

I, Chih-Lin et al. (2014) Recent Progress on C-RAN Centrifcation and Cloudification. IEEE Access

China Mobile (2013) C-RAN, The Road Towards Green RAN

DOCOMO (2014) DOCOMO 5G White Paper. 5G Radio Access: Requirements, Concepts and Technologies

Etemad, Kamran (2014) LTE-A Technology Enablers for Small Cell and CRAN Deployments

Ring, Douglas (1947) Mobile Telephony – Wide Area Coverage – Case 20564. Bell Labs

ETSI (2013) Network Functions Virtualization (NFV); Use Cases [ETSI GS NFV 001]

Fujitsu (2014) The Benefits of Cloud-RAN Architecture in Mobile Network Expansion

Liu, Jingchu (2014) On the Statistical Multiplexing Gain of Virtual Base Station Pools. asXiv

Namba, S. et al. (2012) Colony-RAN architecture for future cellular network. IEEE Future Network & Mobile Summit

Radisys (2014) Evaluating Cloud RAN Implementation Scenarios

Glossary

3G	Third generation	DAC	Digital to analog converter	ICIC	Inter-cell interference coordination
3GPP	Third Generation Partnership Project	DAS	Distributed antenna system	iFFT	Inverse Fast Fourier Transform
4G	Fourth generation	DDC	Digital downconversion	IMS	IP multimedia subsystem
5CC	Five-component carriers	DL	Downlink	IoT	Internet of things
5G	Fifth generation	DMS	Device Management System	IP	Internet Protocol
ADC	Analog-to-digital converter	DPI	Deep packet inspection	I/Q	In phase / quadrature
ADSL	Asymmetric digital subscriber line	D-RAN	Distributed RAN	IT	Information technology
API	Application programming interface	DRAS	Distributed Radio Antenna System	L1	[OSI] layer 1
ATCA	Advanced Telecommunications Computing Architecture	DSN	Distributed service network	L2	[OSI] layer 2
BBU	Baseband unit	DSP	Digital signal processing	L3	[OSI] layer 3
BS	Base station	DUC	Digital downconversion	LAA-LTE	Licensed-assisted access LTE
CA	Carrier aggregation	DWDM	Dense wavelength-division multiplexing	LAN	Local area network
CaaS	Communication as a service	eICIC	Enhanced ICIC	LOS	Line of sight
CDMA	Code division multiple access	EMS	Element management system	LTE	Long Term Evolution
CDN	Content delivery network	eNodeB	Evolved NodeB	LTE-A	LTE Advanced
CFR	Crest factor reduction	EPC	Evolved packet core	M2M	Machine to machine
CO	Central office	ETH	Ethernet	MAC	Media Access Control [layer]
CoE	CPRI over Ethernet	Ether	Ethernet	MIMO	Multiple input, multiple output
COGW	Central office gateway	ETSI	European Telecommunications Standards Institute	MME	Mobility management entity
CoMP	Coordinated multipoint	EV-DO	(CDMA) EVolution Data Optimized	MOCN	Multi-operator core network
COTS	Commercial off-the-shelf	FDD	Frequency division duplex	MPL	Modern Programming Language
C-plane	Control plane	FFT	Fast Fourier Transform	MPU	Modern Processing Unit
CPRI	Common public radio interface	GPON	Gigabit Passive Optical Networks	MS	Mobile subscriber
CPU	Central processing unit	GPRS	General packet radio service	NFV	Network Functions Virtualization
C-RAN	Cloud RAN	GSM	Global System for Mobile Communications	NFVI	NFV infrastructure
CSGW	Cell-site gateway	HetNet	Heterogeneous network	O&M	Operations and management
C-SON	Cloud SON	HWA	Hardware	OAM	Operations, administration and maintenance
CWDM	Coarse wavelength-division multiplexing			OBSAI	Open Base Station Architecture Initiative

OEM	Original equipment manufacturer	SRC	Sampling rate conversion
ORI	Open radio interface	SW	Software
OSI	Open Systems Interconnection	TCP	Transmission Control Protocol
OTN	Optical transport network	TDD	Time-division duplex
OTT	Over the top	TD-LTE	Time division LTE
PDP	Packet data protocol	TTI	Transmission Time Interval
PDPC	Packet data convergence protocol	TWDM	Time and wavelength-division multiplexing
PHY	Physical (layer)	UE	User equipment
PoC	Proof of concept	UL	Uplink
PON	Passive optical network	U-plane	User plane
P-S GW	Packet-switched gateway	vBS	Virtual base station
PSVAC	Power supply, ventilation, air conditioning	V-DAS	Virtualized DAS
PtP	Point to point	V-EPC	Virtual EPC
RAN	Radio access network	VM	Virtual machine
RAT	Radio access technology	VNF	Virtualized network function
REM	Radio environment map	VNFC	VNF component
RF	Radio frequency	VoLTE	Voice over LTE
RLC	Radio Link Control	vRAN	Virtualized RAN
ROCS	Real-time-on-COTS	V-SON	Virtualized SON
RRC	Radio resource control	vSwitch	Virtual switch
RRH	Remote radio head	WDM	Wavelength-division multiplexing
RRM	Radio resource management		
RRU	Remote radio unit		
RSRP	Reference signal received power		
SCAN	Smart Cloud Access Network		
SDN	Software-defined networking		
SDR	Software-defined radio		
SFR	Soft frequency reuse		
SGSN	Serving GPRS support node		
SGW	Serving gateway		
SINR	Signal-to-interference-plus-noise ratio		
SoC	System on a chip		
SON	Self-organizing network		

About RCR Wireless News



Since 1982, RCR Wireless News has been providing wireless and mobile industry news, insights, and analysis to industry and enterprise professionals, decision makers, policy makers, analysts and investors. Our mission is to connect, globally and locally, mobile technology professionals and companies online, in person, in print and now on video. Our dedication to editorial excellence coupled with one of the industry's most comprehensive industry databases and digital networks leads readers and advertisers to consistently choose RCR Wireless News over other industry publications.

About Senza Fili



Senza Fili provides advisory support on wireless data technologies and services. At Senza Fili we have in-depth expertise in financial modelling, market forecasts and research, white paper preparation, business plan support, RFP preparation and management, due diligence, and training. Our client base is international and spans the entire value chain: clients include wireline, fixed wireless and mobile operators, enterprises and other vertical players, vendors, system integrators, investors, regulators, and industry associations.

We provide a bridge between technologies and services, helping our clients assess established and emerging technologies, leverage these technologies to support new or existing services, and build solid, profitable business models. Independent advice, a strong quantitative orientation, and an international perspective are the hallmarks of our work. For additional information, visit www.senzafiliconsulting.com or contact us at info@senzafiliconsulting.com or +1 425 657 4991.

About the interviewer



Monica Paolini, PhD, is the founder and president of Senza Fili. She is an expert in wireless technologies and has helped clients worldwide to understand technology and customer requirements, evaluate business plan opportunities, market their services and products, and estimate the market size and revenue opportunity of new and established wireless technologies. She has frequently been invited to give presentations at conferences and has written several reports and articles on wireless broadband technologies. She has a PhD in cognitive science from the University of California, San Diego (US), an MBA from the University of Oxford (UK), and a BA/MA in philosophy from the University of Bologna (Italy). She can be contacted at monica.paolini@senzafiliconsulting.com.

© 2015 Senza Fili Consulting, LLC. All rights reserved. The views and statements expressed in this document are those of Senza Fili Consulting LLC, and they should not be inferred to reflect the position of the report sponsors. No selection of this material can be copied, photocopied, duplicated in any form or by any means, or redistributed without express written permission from Senza Fili Consulting. While the report is based upon information that we consider accurate and reliable, Senza Fili Consulting makes no warranty, express or implied, as to the accuracy of the information in this document. Senza Fili Consulting assumes no liability for any damage or loss arising from reliance on this information. Names of companies and products here mentioned may be the trademarks of their respective owners. Cover-page and last-page graphics from Neyro/Shutterstock.