Small-cell backhaul:
Industry trends and market overview

By Monica Paolini, Lance Hiley and Frank Rayal
Senza Fili

Table of contents

Senza Fili thanks the following companies for sponsoring and collaborating on this report:

I. Market trends report

1. Introduction: The evolving small-cell backhaul market.
A more precisely targeted approach to small-cell requirements

Backhaul is still one of the hottest topics in the small-cell community, and rightly so. This is an area of intense innovation, driven by many new entrants and an increased presence of established backhaul vendors, where much work is still needed to ensure that backhaul does not become the cost and performance bottleneck in small-cell deployments.

Since we published our first report a year ago, mobile operators have gained valuable experience through trials and initial limited deployments, which have led them to develop more detailed requirements. Vendors have listened and have devoted substantial effort to developing new solutions or adding new functionality to existing ones, while keeping equipment and operating costs down.

As a result, small-cell backhaul is emerging as a distinct segment within the backhaul market, with products specifically designed to meet a set of unique challenges. Small cells operate in continuously changing, cluttered environments over which operators have hardly any control, and which lack the physical and RF stability of cell towers or building roofs.

As we argued in our previous report, the emergence of small-cell backhaul as a separate segment is crucial to the success of small cells, because backhaul plays a crucial role in both the performance and the business case of the small-cell under-layer.

In this report, we provide an update on the latest developments in small-cell backhaul. The business-case tradeoffs and the advantages of different solution types have not changed much over the last year. But the vendor offerings have, and the report focuses on them. We let the vendors speak, and tried to get their inside story – what the drivers and targets behind their new products are, what challenges they faced, and where they want to go next.

In the first part of the report, we split our overview of the market between continuing and emerging trends – what has stayed the same, and what is new.

The second part gives a snapshot of what small-cell backhaul solutions are capable of today and how they meet the operators' requirements. We profiled 13 leading vendors that have agreed to sponsor this report and which have largely complementary solutions and approaches to the small-cell markets.

The third and final part of the report goes beyond features and specs, to tell us the vendors' stories and perspectives, which are all too often buried in the marketing materials and in the data sheets. We have recorded and transcribed conversations with each of the vendors to try to capture what drives the companies and the people working there. We found the conversations to be very engaging and hope that you will too.

2. A look at the last year: What stayed the same and what has changed.

A balancing act to improve performance and reduce equipment cost and size

These days support for small cells – broadly taken to be any base station that is smaller than a macro cell – is nearly universal. Both vendors and mobile operators see small cells as a necessary complement to the macro layer, to provide additional capacity and to improve spectrum utilization. There is much to learn, however, in how to deploy small cells, in terms of location, wireless interface, traffic and interference management, and, of course, backhaul.

Ultimately, many of the lessons will come only once small cells are commercially deployed in fully loaded networks. In the meantime, however, it has become clear that small cells are not going to be as cheap or easy to install as initially expected. To get good performance and reliability, to manage interference and to achieve the desired capacity density, operators have to invest in best-of-breed hardware, and plan their network and choose locations carefully. It is a process that is neither quick nor cheap. And the addition of a sub-layer is bound to increase overall network complexity.

As a corollary, large-scale small-cell deployments will take time. This will give mobile operators the opportunity to get a better understanding of how they have to evolve the network architecture and traffic management to integrate small cells, and it will give vendors time to introduce the solutions that meet performance, functionality and cost requirements.

The realization that small-cell deployments will require more time, cost and effort than initially expected has a strong impact on many aspects of backhaul requirements. We will review them in more detail in the following pages, but the overall push is for backhaul solutions that pack more performance, functionality and resiliency into a smaller form factor with a lower price tag.

Table 1. Evolution of the small-cell backhaul market

Confirmed trends
Small-cell backhaul requirements different from macro-cell backhaul
A higher percentage of backhaul TCO in small cells than in macro cells
Emergence of solutions specific to small-cell deployments
Coexistence of multiple solutions: Fiber and wireless, LOS and NLOS
LTE small-cell deployments still limited
Infrastructure sharing needed to lower TCO
Challenging network planning and installation
Complexity from a multi-vendor, multi-technology backhaul approach
Varying requirements across markets and operators
Emerging trends
Increased flexibility in available backhaul solutions
Partnerships among vendors
More Wi-Fi and 3G small cells
Indoor small cells gaining new support
More serious consideration of the 3.5 GHz band for small-cell access
A blurring line between LOS and NLOS
Vendor focus on shorter installation process
DAS as a complementary approach
Remote-baseband models creating demand for fronthaul solutions

3. Methodology

A common approach in analyst reports is to collect and filter information from vendors, and then distill a comprehensive overview of the competitive landscape that identifies main trends, differentiating approaches, and market targets. From detailed, complex and possibly confusing information, such reports strive to generalize an easy-to-grasp vision that helps the reader understand the market.

The first part of this report follows this approach and offers our assessment of the small-cell backhaul market based on consulting work, research and briefings with vendors and operators, as well as the research done for this report. We listened, thought about what heard, and reported. We expect our audience to be actively engaged in small-cell backhaul, so we have avoided broad overviews and concentrated on what we learned over the last year.

In the rest of the report we take a complementary approach. We take a deep dive into what vendors are doing, which is where most innovation is right now. Mobile operators are still primarily involved in trials and limited deployments and, hence, are busy evaluating vendors for their large-scale deployments. As a result, we believe this is a good time to present a report on what vendors are doing and make the results available to all.

To achieve these goals, we invited vendors in the small-cell backhaul market to share their perspectives and provide support for the report. The 13 participating vendors are listed in Table 2. They all received the same information about our process and the same request for information on their product. They all had the opportunity to send us background information on their product so we could identify good topics for the conversations. While we had a list of areas to cover during the conversations, they were unscripted and, to the extent possible, tailored to each vendor's approach. The report includes transcriptions of the conversations, and for most, a video version is also available online **here**.

Table 2. Report sponsors

Altobridge	www.altobridge.com
BLiNQ Networks	www.blinqnetworks.com
Bluwan	www.bluwan.com
CCS	www.ccsl.com
Cisco	www.cisco.com
DragonWave	www.dragonwaveinc.com
Intracom Telecom	www.intracom-telecom.com
NEC	www.nec.com
Proxim Wireless	www.proxim.com
Siklu	www.siklu.com
SOLiD Technologies	www.solid.com
Tellabs	www.tellabs.com
VubIQ	www.vubiq.com

4. Confirmed trends.

Backhaul is still a key element of small-cell deployments and TCO

- **Small-cell backhaul requirements different from macro-cell backhaul.** At first sight, small-cell backhaul can be perceived as a dressed-down version of its macro-cell counterpart. It has to be cheaper, but it also has lower capacity requirements, and many operators are willing to accept lower reliability and even to use license-exempt spectrum. But this is a deceptive view. In many ways, small-cell backhaul has a more stringent requirement set, which makes it a more challenging solution than macro-cell backhaul – even though one that may have lower profit margins for vendors. Small-cell backhaul solutions, regardless of technology or spectrum used (if wireless), have to meet multiple requirements:

  - Low cost
  - Small form factor
  - High capacity
  - Low power
  - Fast and easy installation that can be done by semiskilled employees
  - Tolerance of sway and ability to operate from precarious locations (wireless only)
  - Ability to cope with changes in the environment that can be unforeseen and are not under the operator's control (wireless only)
  - Scalability to accommodate the addition of new small cells within the same footprint
  - Low latency, to support LTE's X1 and S2 interfaces

  Most of the areas where small cells will be deployed, such as metro zones, are a much more challenging environment than macro cells' towers, where the equipment is securely installed in a protected location.

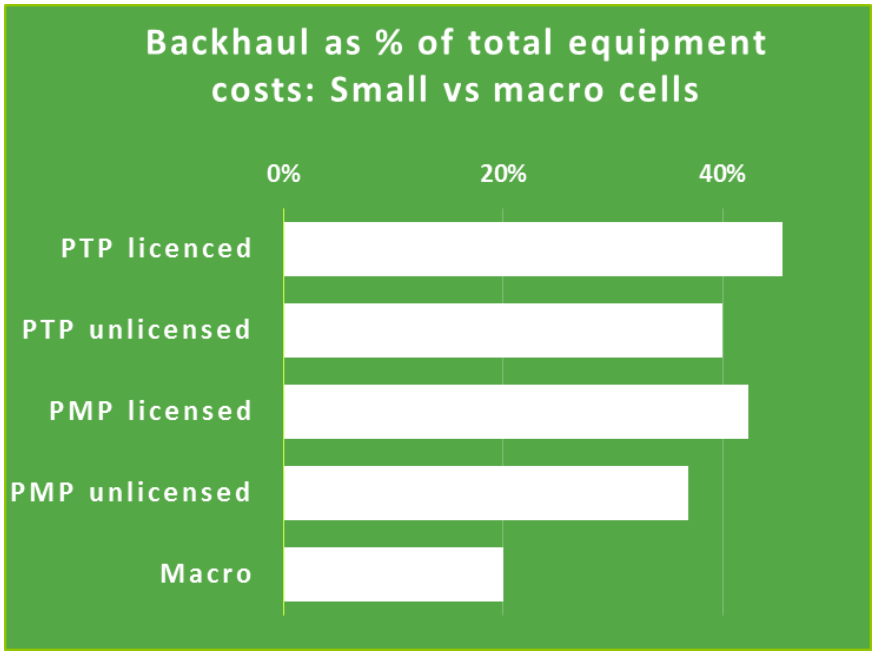

Figure 1. Backhaul capex and equipment as a percentage of TCO. Source: Senza Fili

- **A higher percentage of backhaul TCO in small cells than in macro cells.** Because of the differences in requirements, the percentage of RAN TCO accounted for by small-cell backhaul is approximately twice that accounted for by macro cells (Figure 1). The emergence of cost-effective solutions is going to be crucial to establishing a solid business case for small-cell backhaul. Without cost-effective solutions, it is difficult to see how extensive small-cell deployments can be rolled out, especially at a time when mobile operators face severe financial pressures.

- **Emergence of solutions specific to small-cell deployments.** The demanding environments in which small cells operate have driven efforts to create new types of backhaul products that provide a level of flexibility and specific functionality different from those of macro-cell backhaul solutions. While many small-cell backhaul vendors are also macro-cell and enterprise backhaul vendors, a number of new vendors have entered this market, with a tighter focus and new approaches to network management, leveraging previously underused spectrum bands and introducing advanced features.

- **Coexistence of multiple solutions.** Although many operators have a preference for one backhaul technology over others, they agree that no single technology can meet all the backhaul requirements. We have developed a decision tree (Figure 2) that summarizes how different technologies may be deployed. While it is not prescriptive, the graph illustrates the complexity of the decision process that the operator has to go through for each small cell.

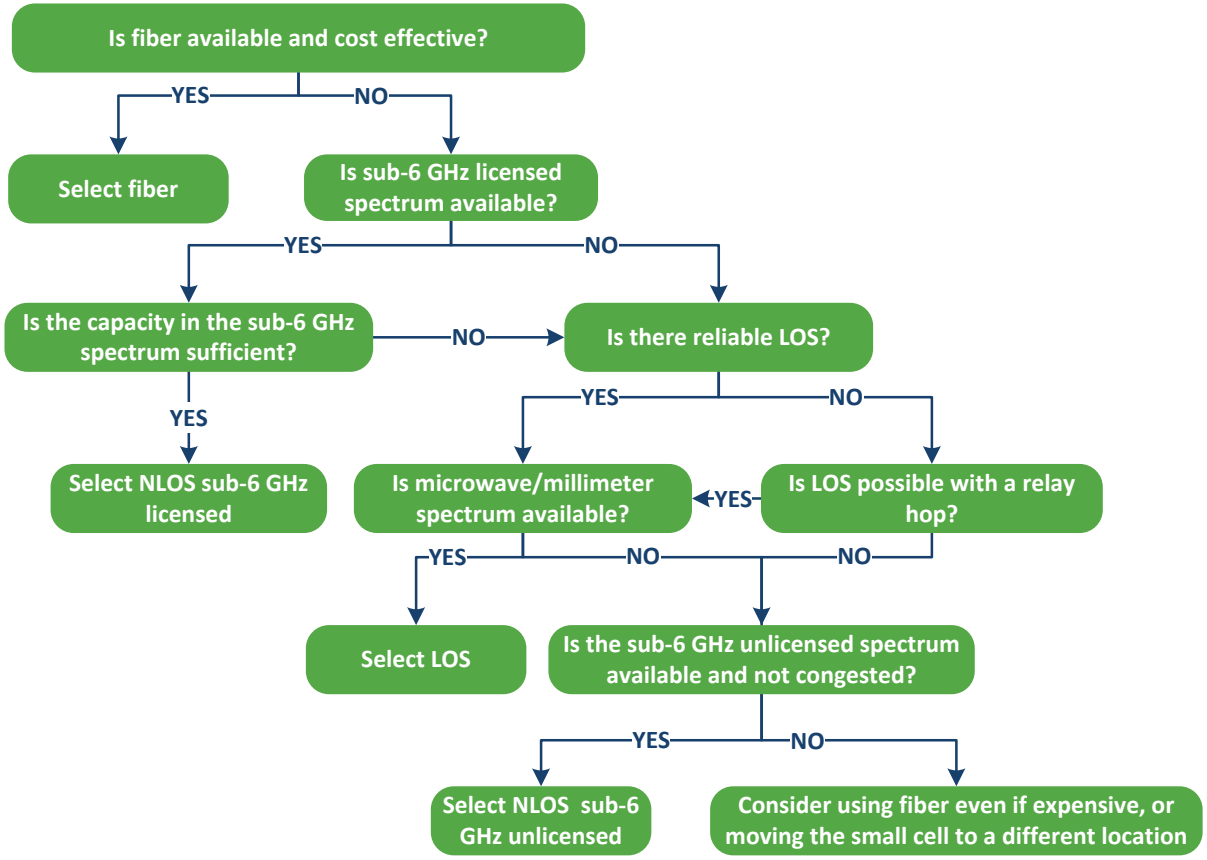

Figure 2. Small-cell backhaul decision tree. Source: Senza Fili

- **Fiber and wireless.** Unquestionably, fiber is the best solution for small-cell backhaul, so where it is available and cost effective, it typically wins over wireless solutions. But it is not always available, and when it is, often it is not cost effective, because either the installation costs (mostly due to trenching) or the operating costs (i.e., leasing costs) are too high. In many ways this is an easy decision point, mostly driven by how much an operator is willing to pay.

- **LOS and NLOS.** If the operator selects wireless backhaul, the choice among technologies becomes more complex, because it depends on a larger number of factors, including spectrum and LOS availability, and capacity requirements (Figure 3). Other things being equal, operators prefer NLOS solutions in licensed spectrum. But these solutions typically have less capacity and higher latency (especially in true NLOS conditions), and this makes them unsuitable for high-capacity small cells or for backhaul links that support multiple cells. In cases where 4G is combined with Wi-Fi or where operators share the backhaul link, capacity requirements grow quickly and operators are likely to select LOS solutions in the busiest locations, with the addition of relays where there is a need to compensate for the lack of direct LOS.

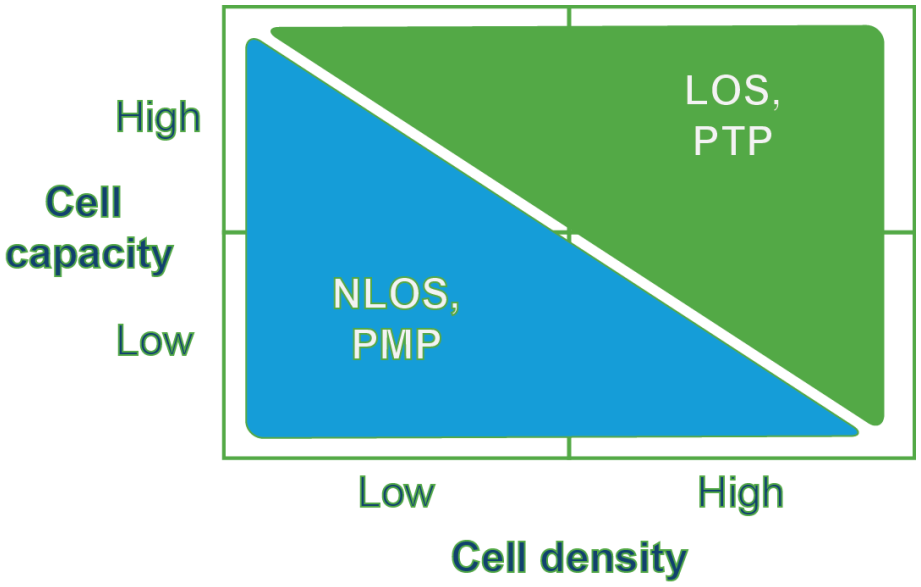

Figure 3. LOS or NLOS? Source: Senza Fili

- **LTE small-cell deployments still limited.** The most common question on small cells is when will they really take off – and, indeed, whether this will ever happen if, after so much attention, we have not yet seen wide commercial deployments, and even the rate of growth in mobile traffic shows signs of slowing down, thus reducing the urgency of small-cell rollouts.

  As we noted before, it will take time for small cells to become a primary traffic conduit. But more specifically, LTE small cells face a delay because most operators either do not have an LTE network yet, are still deploying the more basic macro infrastructure, or do not yet have congestion problems on their new LTE networks. Even operators like Verizon, which has

more than 50% of its data traffic on LTE, face most of their congestion in their 3G networks, not on LTE. Today these operators would need 3G small cells, not 4G ones. But the business case for 3G cells is not very good, so Wi-Fi fills the gap and will continue to do so for some time. The outcome of this is that operators are in no hurry to deploy LTE small cells, and this shifts the revenue opportunity for backhaul vendors forward into the future. This is especially true for those offering high-capacity links that are better suited for LTE than for 3G small cells.

- **Infrastructure sharing needed to lower TCO.** As operators better understand the complexity and costs involved in small-cell rollouts, they also try to identify ways in which costs can be kept under control. Possibly the most effective way is to share the infrastructure. Not only does infrastructure sharing lower costs, it also reduces the amount of hardware to be installed at street level, and this is beneficial from a planning and operational perspective.

  The main obstacles to infrastructure sharing are mobile operators' reluctance to give up some degree of control over the infrastructure, and the perception that infrastructure sharing will give competitors an advantage or reduce operators' ability to differentiate their service offerings. While many European operators see infrastructure sharing as necessary to justify the small-cell business case, in some markets, such as the US, there is strong resistance to it, possibly because operators there are under less financial pressure.

  Multiple approaches to infrastructure sharing are possible. Operators may share all the infrastructure including radios, or may simply co-locate their equipment. In most cases, however, the backhaul is shared, and this creates additional requirements for capacity and traffic management (e.g., QoS, OAM, and SLA management).

  Although they may have chosen a sharing arrangement to keep costs down, operators may become less cost sensitive once they choose to do so, because now performance, reliability and functionality assume increased relevance. They are more likely to gravitate toward high-end solutions and

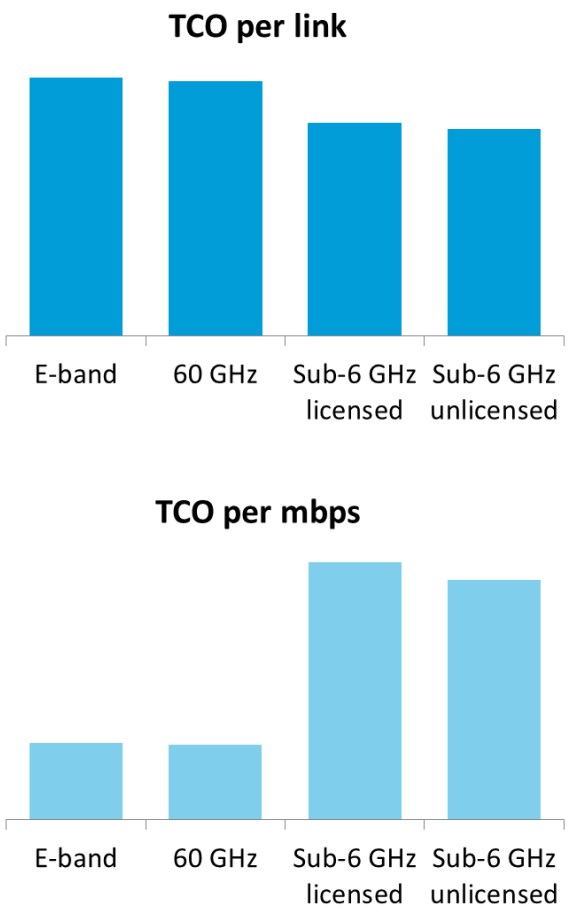

Figure 4. A comparison between millimeter wave and sub-6 GHz wireless backhaul solution. Per-link TCO is lower for sub-6 GHz solutions, but per-mbps TCO is lower for higher-capacity millimeter band solutions. Source: Senza Fili

may be more open to purchasing additional spectrum for backhaul if needed (and if available).

- **Challenging network planning and installation.** All vendors we talked to – both RAN and backhaul – acknowledge that ease of planning and installation are critical requirements and that they have been working to address them. The goal is to keep installation times and complexity at a minimum, with features such as zero-touch provisioning, self-organizing tools, or remote alignment.

  Much effort has gone into this, but the challenge remains, because of the variety of locations available to small cells (and the unpredictable conditions they pose) and because effective network planning depends on the interaction of multiple factors, such as macro-cell locations, topology and subscriber distribution within the coverage area. Improvement in the design of backhaul equipment certainly helps, but much of the planning and installation complexity stems from the environment rather than the equipment, and experience in commercial deployments will be essential to streamlining the deployment process further and developing best practices.

- **Complexity from a multi-vendor, multi-technology backhaul approach.** Mobile operators will have to deploy multiple backhaul solutions to reach all their small cells (Figure 5). To do so, they can choose a single vendor that offers multiple solutions and can manage all backhaul with tools provided by the vendor. In most cases, however, we expect operators to use more than one backhaul vendor within the same footprint. Within a metro area, an operator may have fiber plus a NLOS and a LOS wireless solutions, each with its own capacity, latency and functionality. Operators gain flexibility from adopting multiple solutions, but they pay a price in managing a backhaul architecture that has grown more complex. Interoperability and partnerships among vendors and their network management tools will also play a prominent role in enabling small-cell deployments and in coordinating transmission within the wider mobile network.

- **Varying requirements across markets and operators.** Small cells provide additional coverage and capacity in overloaded networks. Each operator,

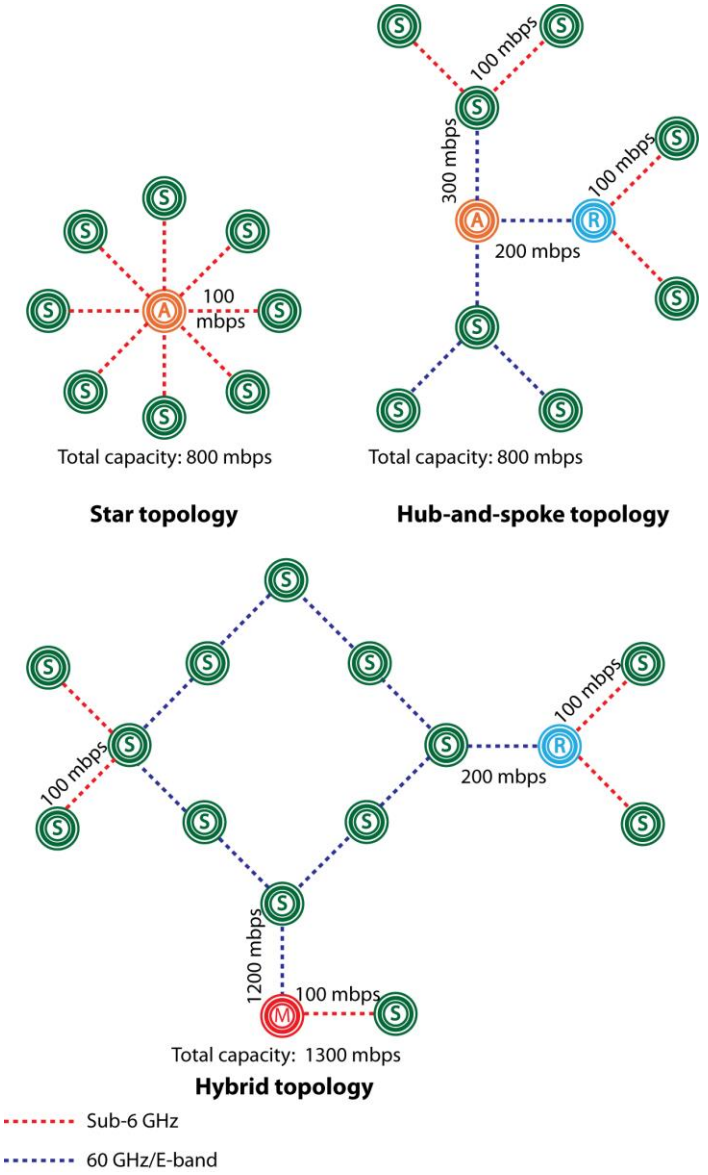

Figure 5. Examples of how different backhaul technologies may coexist in different topologies (S: small cell; M: macro-cell or other aggregation point; R; relay. Source: Senza Fili

depending on its deployed infrastructure, footprint characteristics and subscriber usage patterns, faces a unique set of needs to bring network performance up to the desired levels. Dense urban areas, high use of public transportation and more time spent outside the home and office create higher demand for capacity in metro locations in Asian and European countries, compared to North America. But the more generous traffic allowances in North America are responsible for higher volume generated by subscribers. Availability of (or plans for) LTE will also affect small-cell plans.

As a result, small-cell rollout strategies vary across mobile operators, and so do the backhaul strategies, as the capacity, functionality and reliability constraints change. In this context, the mobile operators' ability to choose from a wide array of solutions is critical to enabling all of them, regardless of their backhaul strategy and requirements, to find the right mix of solutions.

5. Emerging trends.

Variety in small-cell deployments requires flexibility in the backhaul

- **Increased flexibility in available backhaul solutions.** In parallel with the growth of small-cell backhaul as a separate segment within the backhaul market, we have seen an increase in differentiation and sophistication among vendors over the past year.

  Having a small-cell solution is no longer sufficient. Some vendors are expanding their existing solutions to new bands to give operators more choices and to be able to address markets with differing spectrum regulations. Other vendors are moving beyond terminal and hub equipment, expanding the scope of their solutions to include advanced management and self-configuration options. Still other vendors focus on transmission technologies such as beamforming to improve link capacity and reach.

  The net result is that there is healthy competition among backhaul vendors, but this makes mobile operators' vendor selection process more complex, because they have more tradeoffs to consider.

- **Partnerships among vendors.** In the short term, many operators may prefer to use the same vendor in the small-cell network that they do in the macro-cell network, to reduce complexity. In the long term, small cells have the potential to change the relationship with RAN vendors: operators are pushing more forcefully for interoperability across vendors within the same coverage area. In the choice of backhaul vendors, mobile operators feel less tied to their macro-cell backhaul vendors and hence are more open to working with new entrants or established vendors. At the same time, as they need to use multiple solutions to cover all their small-cell backhaul requirements, they are mindful of the complexity of integrating these

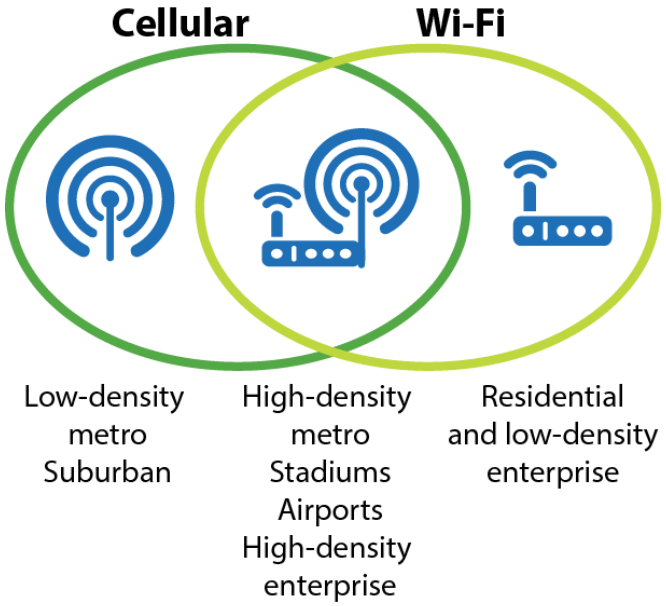

Figure 6. A complementary approach for cellular and Wi-Fi small cells. Source: Senza Fili

solutions, which have varying installation requirements, performance characteristics, and management functionality.

To address this concern, many vendors – and especially the new entrants with a tighter solution focus – have started to work together to build an ecosystem in which operators can choose the best-suited solutions but keep complexity at a minimum.

Partnerships, interoperability testing, and other marketing arrangements are becoming more common, and we expect to see more. The ability of new entrants and niche vendors to work together may turn out to be essential to their ability to penetrate the thick walls of mobile operators' procurement departments.

At the same time, successful collaboration among smaller vendors may put increasing pressure on tier-one vendors and accelerate movement along the inevitable path to consolidation. Indeed, the high number of new entrants in a market that can sustain only a limited number of them is a reason for concern – although it is also a sign that there is a need and a scope for the evolution of new approaches and solutions in this space.

- **More Wi-Fi and 3G small cells.** For a long time, the focus of the small-cell community was on LTE. Why deploy 3G small cells if LTE small cells are available? Doesn't LTE provide enough capacity to make Wi-Fi an unnecessary complement?

  It turns out that most subscribers are still on 3G, and operators need more capacity on their 3G networks and not (yet) on their LTE networks. And as we wait for LTE networks to reach capacity, Wi-Fi is doing a great job of offloading traffic from LTE (Figure 6). According to Cisco's VNI, a third of the traffic from devices with cellular connectivity goes over Wi-Fi – and other estimates are considerably higher. This is not bad for a technology that initially had no ambition to carry mobile traffic and that, in most cases, does so at no marginal cost to subscribers and no cost to mobile operators.

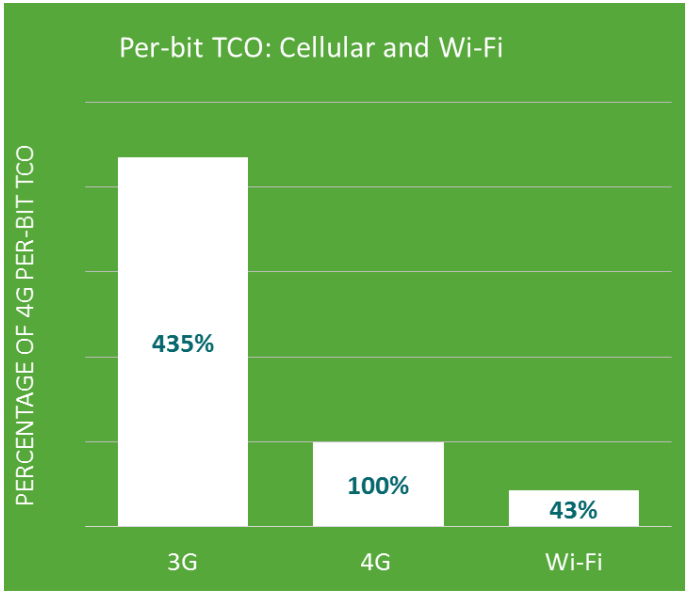

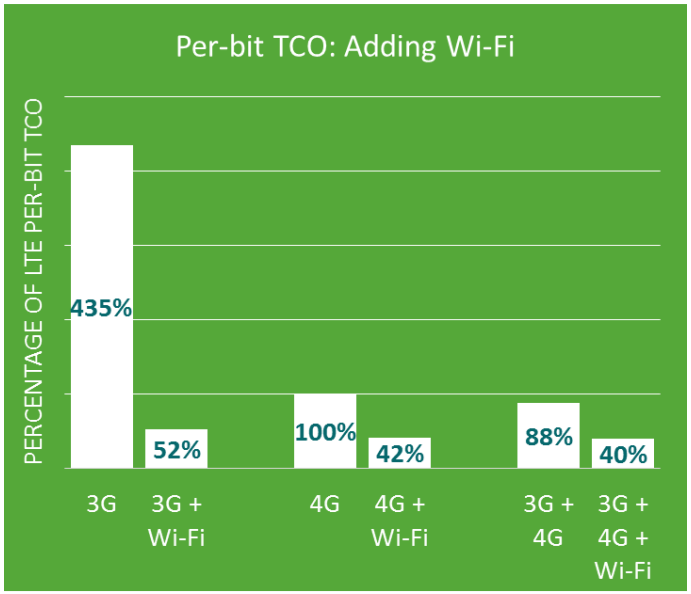

Figure 7. Cellular and Wi-Fi small-cell TCO over a five-year period. Source: Senza Fili

As we wait for demand for LTE small cells to grow, mobile operators eager to increase capacity are increasingly turning to 3G and Wi-Fi to meet their short-term requirements, planning to move to LTE small cells when appropriate – i.e., when LTE networks are overloaded and Wi-Fi cannot absorb the additional traffic. This may take some time, especially given that carrier Wi-Fi is emerging as much more attractive to mobile operators than the old-school, low-cost hotspots, which provide excellent value in a coffee shop but do not enable operators to control traffic and provide a consistent user experience. With carrier Wi-Fi, operators can integrate the Wi-Fi infrastructure with their cellular network, support roaming, and provide secure access. Wi-Fi is no longer used only for blind offload. It has acquired full rights as a new RAT alongside cellular technologies. In this perspective, Wi-Fi small cells become a natural extension of the concept of small cells – provided that they are carrier Wi-Fi small cells that support Hotspot 2.0, NGH and Passpoint functionality.

The economics behind Wi-Fi small cells and combined Wi-Fi and cellular small cells are compelling (Figure 7)[1], and we expect the move toward a heavier presence of 3G and Wi-Fi among small cells to have a substantial impact on backhaul as well. First, this trend may speed up small-cell deployments (no need to wait for LTE congestion), creating more demand for backhaul equipment. Second, and more importantly, while 3G backhaul requirements are substantially lower than LTE's, carrier Wi-Fi has capacity requirements comparable to LTE's. Because Wi-Fi does not create interference in the cellular network, it does not need the low latency that LTE requires for coordinating transmission with the macro network. Still, Wi-Fi small cells – owned and managed by the operator and an integral part of the mobile network – need a high-performance backhaul link, even though it is debatable whether it has to be carrier-grade.

- **Indoor small cells gaining new support.** Despite the fact that some operators report that up to 80% of data traffic comes from subscribers in indoor locations, most operators think of small cells as a mostly outdoor sub-layer. Venues such as stadiums, airports, or large malls may be natural targets for indoor small cells, but lampposts and other street furniture are the main targets of metro-zone deployments.

  Over the last year, however, operators have started to voice a stronger interest in indoor small-cell deployments that bring small cells even closer to their subscribers and move them farther away from macro cells. Closeness to subscriber means better modulation, and hence more capacity. Better separation from macro cells means less interference to manage – and, as a result, even more capacity. The downsides of indoor coverage – sites are more difficult to acquire, manage and operate – still limit the appeal of indoor small cells, but we expect to see an increased reliance on them in locations where operators can count on easy access to the infrastructure.

  The move to indoor cells can have substantial implications for backhaul, because wireline solutions are likely to dominate in indoor locations, and this may reduce the revenue opportunity for wireless vendors. At the same time, the assumption that indoor small-cells require wireline backhaul should be revisited, because in locations such as airports and stadiums, short wireless links may reduce the cost of deploying and

1. Senza Fili Consulting, Carrier Wi-Fi® for mobile operators (2013). White paper about per-bit TCO for Wi-Fi and cellular small cells, commissioned by the Wi-Fi Alliance.

operating an indoor network, and provide additional flexibility in the location of the indoor equipment.

- **More serious consideration of the 3.5 GHz band for small-cell access.** Most operators plan to deploy small cells in the same channel they use for macro access and to mitigate the resulting interference with tools such as CoMP and eICIC. But small cells add complexity to the network, and it is not yet clear how much benefit they will add. In any case, interference can be managed and mitigated, but not eliminated. Deploying small cells in a separate band is not feasible for most operators, because most operators do not have sufficient cellular spectrum.

  The 3.5 GHz band has emerged as a promising candidate, in the long term, for the dedicated use of small cells. The band's short range – a disadvantage in the macro-cell layer – enables operators to use the spectrum efficiently in the small-cell layer and keep interference down. Spectrum in the 3.5 GHz band is available and underutilized in many markets, even though often it is not in the hands of mobile operators. For a long time, mobile operators have been doubtful of their ability to use 3.5 GHz spectrum to support mobility, and of handset vendors' willingness to support the band. In the short term there are no major plans to use this band for small-cell deployments, and there is not sufficient demand (or spectrum scarcity) to justify the costs of adding 3.5 GHz in devices. But in the long term, the capacity that can be added with the 3.5 GHz band may become necessary.

  A move to use 3.5 GHz for access may limit the attractiveness of this band for backhaul, thus reducing the scope for sub-6 GHz licensed-spectrum solutions. At the same time, the necessary acquisition of the 3.5 GHz spectrum by operators may open the way for its use for in-band backhaul – which would allow mobile operators to concurrently use their spectrum assets for both access and backhaul.

- **A blurring line between LOS and NLOS.** The received wisdom in backhaul is that above 6 GHz, LOS is required. In sub-6 GHz bands, both LOS and NLOS are possible, and NLOS is the dominant architecture because it gives more flexibility.

  We are seeing two trends. While NLOS works in sub-6 GHz, it can be spectrally inefficient and lead to severe limitations in capacity that make many sub-6 GHz solutions not fit for LTE and Wi-Fi small cells, backhaul sharing, or multi-hop backhaul. As a result, there is increased interest in using sub-6 GHz spectrum in a PTP architecture or in using more advanced antenna technologies.

  At the same time, there is a promising exploration of using high-frequency bands for NLOS and near-line-of-sight scenarios, leveraging the signal reflection, diffraction and penetration in the environment. This approach may lead to innovative high-capacity backhaul solutions that use spectrum that is less expensive and more widely available than sub-6 GHz.

- **Vendor focus on shorter installation process.** A year ago, vendors were mostly concerned about capacity, and LOS versus NLOS requirements. Lately the emphasis has shifted to the installation process as operators have become increasingly aware of the complexity in deploying small cells.

  Vendors' efforts are aimed at three targets. The first is to simplify network planning – i.e., selecting the appropriate link for each cell, and determining the topology of the small-cell local network and location of relays. The second is to reduce the effort and time required to install a link during the initial rollout. The third is to facilitate network growth, as operators plan to incrementally add small cells to the existing footprint and want to seamlessly insert them into existing local backhaul networks that link the small cells to the nearest aggregation point.

- **DAS as a complementary approach.** DAS can be seen as an antecedent to small cells or as a competing solution. More interestingly, it may become one of the options used in small-cell deployments, in a continuum that also includes both standalone small cells and remote

baseband units. Small cells primarily address a capacity requirement, and secondarily a coverage requirement. At the opposite end, DAS's main benefit is improved coverage, with some capacity increase. In some environments operators may need capacity more urgently, in others coverage. Or they may need coverage initially, and capacity at a later stage.

The inclusion (and adoption) of DAS solutions in the mobile operator small-cell toolkit shifts the backhaul requirements toward fiber or fiber-equivalent solutions, limiting the choice among the available solutions.

- **Remote-baseband models creating demand for fronthaul solutions.** In addition to DAS-based solutions, small-cell deployments may use remote-baseband architectures that can reduce cost, use network resources more efficiently, and limit the amount of equipment that has to be installed on street furniture or at indoor locations open to the public.

  As in the DAS case, a shift to remote-baseband small cells will change the backhaul requirements and create the need for fronthaul solutions, which traditionally have meant fiber but may also be supported by wireless solutions over the short distances required in small-cell deployments.

II. Vendors' profiles

1. Altobridge

Overview. The Data-at-the-Edge (DatE) product by Altobridge minimizes backhaul bandwidth requirements by implementing byte-caching technology that operates on user data, while leaving signaling and voice data unaltered. Additionally, DatE incorporates a self-learning algorithm that learns the behavior of the most popular content being downloaded and pre-positions it at the access network to save backhaul capacity. Hence, byte caching operates in two dimensions: as personal cache and as location-specific cache. Altobridge estimates that DatE saves 50% on average in backhaul capacity, with a peak of as much as 70%. DatE preserves important operator requirements, such as legal intercept and E911, and is transparent to billing and content-filtering functions. It is compatible with 3G, LTE and Wi-Fi access technologies.

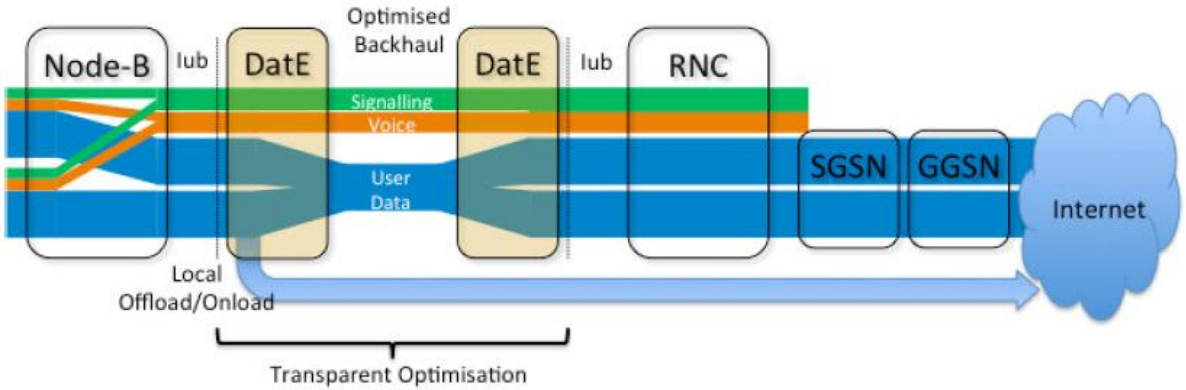

Figure 8. DatE byte caching is positioned on both sides of the transmission link. User data is minimized in terms of bandwidth requirements while voice and signaling are left unaltered. Source: Altobridge

Positioning. Altobridge positions DatE as the only byte-caching backhaul optimization solution in the market that supports the interface between the 3G base station and the RNC (Iub), in addition to other interfaces such as Iuh and S1. With types of traffic that lend themselves well to byte caching, such as video – which by the end of 2012 accounted for 51% of mobile data traffic[2] – DatE is well positioned to provide significant savings in backhaul capacity that should prompt operators to review their requirements and planned expenditures. Furthermore, DatE provides operators additional data on end-user behavior, which can strengthen their hand with respect to OTTs and which allows operators to launch new revenue-generating services.

Threats. Byte caching is relatively new in wireless networks. It relies on DPI techniques, which introduce some latency, although this is offset with large gains related to positioning the data close to the user, which results in potentially large savings in backhaul capacity. The effectiveness of the caching algorithm is a critical factor for the success of the caching solution. This is a matter that needs to be determined within the context of the type of subscriber data traffic and the demand for that traffic. The location of cache (core versus edge) is a tradeoff between cost and performance that must be considered. In the case of Altobridge, the approach is to use a book-end solution for the transport part of the network. This is in contrast to solutions that opt to place the caching engines at other points of the network, such as between the mobile device and base station.

Evolution. Altobridge plans to improve the hit rate of its byte-caching technology and to increase the utility of pre-positioned data in mobile networks by adding more intelligence through improved learning of user behaviors.

2. Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2012–2017. February 2013.

2. BLiNQ Networks

Overview. Although a recent entrant on the small-cell scene, BLiNQ has been working on the small-cell backhaul challenge for over five years. The company is a spin-off of Nortel's defunct backhaul transmission group that had been forming views on backhauling denser mobile networks since 2007. The X-100 self-organizing backhaul product operates in the licensed and license-exempt sub-6 GHz bands. BLiNQ focuses on ease of installation and configuration, seen as one of the main decision points for operators. BLiNQ says its solution can be installed in under 30 minutes, with network configuration taking place in the background as part of the self-organizing feature.

Positioning. The X-100 can be deployed in PTP LOS mode or as a PMP NLOS system with maximum link distances of 20 km and 1 km, respectively, and enabling BLiNQ to address small-cell backhaul deployment scenarios ranging from dense urban to rural using a common platform. The PMP hub is a two-box solution featuring all-outdoor access points with an external sector antenna. The remote terminals are a single-box unit that uses an integrated antenna.

Threats. BLiNQ relies on sub-6 GHz radio spectrum for its products. Licensed sub-6 GHz bands are convenient because they are area-licensed, but they are also in demand for RAN deployment, and that may increase the cost and limit the availability of the spectrum. License-exempt spectrum support relieves some of that spectrum cost and availability risk, but interference mitigation features may not be able to stem the challenge of a rising noise floor.

Evolution. BLiNQ is working to enhance interference mitigation with Managed Adaptive Resource Allocation (MARA). The company also expects that wider (20 MHz) channels will increase hub capacity to 240 mbps, and that dual-carrier 4 x 4 MIMO will further raise capacity to 1 gbps. The introduction of a self-install technology that automatically connects the remote terminal to the proper hub unit while simultaneously optimizing MIMO performance is also planned. In the longer term, BLiNQ is looking to integrate its backhaul know-how into third-party small-cell

Figure 9. BLiNQ X-100 small-cell backhaul system. Source: BLiNQ

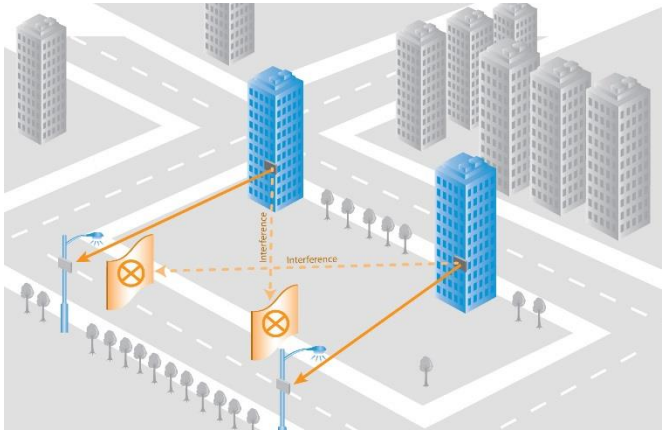

Figure 10. MARA is designed to increase backhaul capacity by reducing interference across the backhaul network. Source: BLiNQ

hardware at the physical or silicon level to drive cost, installation and physical dimensions down.

Features	X-100	X-1200
Market focus	Small-cell wireless backhaul	
Spectrum band	2–4 GHz	2–6 GHz
LOS/NLOS requirements	LOS to >20 km, NLOS to 1 km	
Channel sizes	5, 10 MHz	5, 10, 20, 40, 20, 10 MHz
TDD/FDD	TDD	TDD, FDD roadmap
Modulation	QPSK, 16 QAM, 64 QAM, 256 QAM	
Capacity	Up to 75 mbps	Up to 480 mbps
Latency	12 ms	5 ms
Antenna specs/configuration	Remote: integral 17 dBi antenna; Hub: external sector antenna	Remote: integral 16 dBi antenna; Hub: external sector antenna
Integrated/external antenna	Remote: integrated; Hub: external	
Maximum terminals per hub	4	8 (16 roadmap)
Hop length	PMP operates in a single hop	PMP operates in a single hop
Size	Hub and remote terminal: 31 x 21.8 x 8.3 cm	Remote terminal: 30 x 20 x 11.5 cm; Hub: 30 x 20 x 13 cm
Weight	Hub and remote terminal: 3.8 kg	
Form factor	Hub: separate antenna; Remote: single enclosure	
Integration with small cell	Small cell and remote hub can be combined within single radome enclosure	
Power consumption	35 W consumption for hub and remote	45 W consumption for hub and remote
Equipment cost	Approx. $3,500 per link in PMP, with 3 remote terminals	N/A
Installation	20-minute installation	5-minute installation
Architecture	PMP	
Topologies supported	Tree, star, ring	
Small cells supported by a link	2	6
X2 support	Interference mitigation technologies enable scalability of the solution	
Complementary technologies	Fiber, LOS microwave, E band, V band	

3. Bluwan

Overview. Building on technology that started inside the French multinational military and infrastructure group Thales, Bluwan was spun off in 2005 and has since developed PMP wireless access and backhaul solutions. Unlike other PMP architectures, Bluwan's does not rely on over-the-air statistical multiplexing to improve traffic management and capacity. Instead, Bluwan exploits an underutilized 3 GHz–wide millimeter-wave band at 42 GHz, and has developed an in-house radio to deliver flexible, high-capacity links. The radio can handle up to 1 GHz–wide spectrum blocks split into 40 MHz–wide channels. The channels can be channel-bonded or fused to create as much link capacity as required for individual customers or sites. The radio can operate in narrower blocks for deployment in smaller frequency allotments. The solution leverages the other architectural advantages of PMP with a central access point and only a single terminal required per connection.

Positioning. Bluwan targets both the macro-cell and small-cell backhaul. At 42 GHz the system is LOS, so Bluwan has developed a compact form factor and relay system to cope with difficult-to-reach sites. Bluwan's LinkFusion features an integrated antenna with a 6-degree beam width that reduces the effort required in installing and aligning the system in an urban environment. With a weight of 3 kg, the LinkFusion targets street furniture deployments. The hub is a two-box configuration that can be a challenge for some urban deployments.

Threats. Bluwan is relying on a single spectrum band for its solution – the availability of which is not guaranteed in every jurisdiction (the band was originally set aside for different purposes than backhaul). This is a scenario that besets many area-licensed spectrum bands required by PMP solutions.

Evolution. Bluwan has plans to introduce self-optimizing network features into its products, as well as increase capacity and efficiency through the introduction of 2 x 2 MIMO and beam-forming antennas.

Figure 11. LinkFusion NTE terminal, an area-licensed microwave backhaul system for small cells.

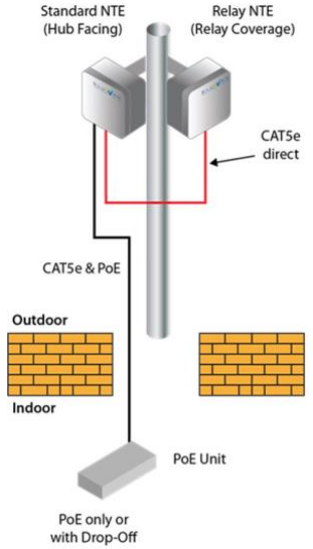

Figure 12. LinkFusion Relay extends the range of the standard terminal for hard-to-reach sites.

Features	LinkFusion
Market focus	HetNet/small-cell backhaul, enterprise access, PMP
Spectrum band	ETSI- and CEPT-harmonized 40.5–43.5 GHz
LOS/NLOS requirements	LOS with relay feature
Channel sizes	Up to 20 x 40 MHz channels on a single radio
TDD/FDD	Multichannel dynamic or static TDD
Modulation	BPSK, QPSK, 16 QAM, 64 QAM
Capacity	1 gbps capacity per LinkFusion IDU (8 x 40 MHz bonded), 2.4 gbps total radio capacity (20 x 40 MHz aggregated), 120 mbps per 40 MHz channel
Latency	Typical: 2 ms or less
Antenna specs/configuration	16 dBi (90 degree sector antenna), 19 dBi (45 degree sector antenna), 22 dBi (22.5 degree sector antenna) – all integrated into the ODU (no separate parabolic antenna)
Integrated/external antenna	Integrated patch antenna
Maximum terminals per hub	20 per sector in PTP mode; 100 per sector in PMP mode
Hop length	90 degree antenna: up to 2 km, can be extended to 2.5 km with narrower-beam antennas. In rain zone D, with 99.9% availability: 1.6 km. Relays can be used in 6, 22.5, 45 and 90 degree configurations to extend the range
Size	Small-cell end: 281 x 211 x 106 mm (NTE-120), or 180 x 180 x 150 Metro NTE; aggregation point ODU: 281 x 211 x 106 mm
Weight	LinkFusion ODU at the AP is 3.4 kg. LinkFusion NTE at the end point is 3 kg. Metro NTE is < 2 kg
Form factor	Single enclosure, PoE, all outdoor
Integration with small cell	No. Roadmap dependent on customer interest
Power consumption	Power at the small-cell end: < 25 W (PoE)
Installation	End-point installation, single person: < 2 hours. Sector installation: < 1 day
Equipment cost	Target of $3,000 per link, depending on volume and number of UE units in the network
Architecture	PMP low-latency TDD with channel aggregation; each channel can operate in PTP or PMP
Topologies supported	Hub and spoke, star, and relays
Small cells supported by a link	Single sector can support up to 100 4G/Wi-Fi small cells. Relay node can connect 5 small cells to an aggregation node
X2 support	Yes. LinkFusion provides enough peak capacity to meet backhaul requirements for X2, assuming 5–10% overhead
Complementary technologies	WWDM allows Bluwan's LinkFusion to aggregate multiple independent channels onto a single air interface Ultra-wide band (1 GHz) chipsets: Operating at 40.5–43.5 GHz, enables 2.3 gbps sector capacity Multi-gigabit PMP architecture: TDD architecture allows throughput to each terminal to be controlled
Competing technologies	PTP microwave products in the 40 GHz band; LOS PMP microwave solutions from CBNL, Intracom
Future product focus	Self-optimizing networks, 2 x 2 MIMO, multi-beam and beamforming high-gain antennas, radio over fiber, IPsec, Wi-Fi gateway

4. CCS

Overview. CCS is building on the experience of its two Cambridge University–educated founders, who have collaborated previously in two other wireless companies, Adaptive Broadband and Cambridge Broadband Networks. Refreshingly the new company does not feature broadband in its name, nor does it limit itself to mainstream technology. Instead, the company brings a new architecture for wireless backhaul – many point to many point (MPTMP) – that CSS says will combine the high capacity and low latency benefits of PTP microwave with the ability to reach difficult small-cell locations. The solution put forward by CCS is a multipoint network operating in the area-licensed 28 GHz band. It features identical nodes that self-organize via advanced algorithms that instruct each node to look for the best signal possible from another node. Once established, the nodes organize to maximize efficiency and performance. By having only one hardware component plus advanced configuration tools, CCS claims an electrician can install its node in 30 minutes.

Positioning. CCS is firmly focused on metropolitan small-cell backhaul. The terminal incorporates an integral 270 degree antenna and is intended for mounting on street furniture or corner structural locations. The system is designed to minimize installation and operating costs in evolving small-cell deployments in which operators gradually increase the density of small cells within the same footprint. Variable-channel bandwidth delivers up to 450 mbps in a 112 MHz channel. The system is available now at 28 GHz, and CCS plans variants at 26, 32 and 40 GHz.

Threats. CCS has a tight focus on what is still a niche market. Although its architecture and technology appear well suited for small-cell backhaul, availability and cost of spectrum and the LOS requirement may prove limiting.

Evolution. With many other PMP systems touting link capacities in the 1 gbps range, CCS will focus on capacity in the future. It also plans to add wider spectrum band support to address demand in countries where access to the 28 GHz band is not available.

Figure 13. CCS MPTMP area-licensed microwave backhaul terminal for small cells. Source: CCS

Feature	CCS Many point to many point system
Market focus	Metropolitan small-cell backhaul
Spectrum band	28, 26, 32, 42GHz
LOS/NLOS requirements	LOS
Channel sizes	28, 56, 112 MHz
TDD/FDD	TDD, FDD, and dual TDD
Modulation	64QAM 4/5 FEC, 64QAM 3/4, 64QAM 5/8, 16QAM 4/5, 16QAM 3/4, 16QAM 5/8, 16QAM 1/2, QPSK 4/5, QPSK 3/4, QPSK 5/8, QPSK 1/2
Capacity	112Mhz = 450mbps gross, 400mbps net Ethernet
Latency	Constrainable per node down to 125 µs per hop. Target of 1 ms for 8 hops
Antenna specs/configuration	270 degree x 20 degree 19dBi
Integrated/external antenna	Integrated antenna
Maximum terminals per hub	Each node supports up to 16 logical connections
Hop length	QPSK up to 1 km. 64Q AM up to 350 m 99.999% availability
Size	190 mm diameter, 130 mm height. Approx. 14 l volume
Weight	4.5 kg
Form factor	Single-enclosure cylinder with cut-out for the street furniture or wall corner
Integration with small cell	Backhaul only
Power consumption	36 W per node
Installation	30 min to 1 hr
Equipment cost	In line with market requirement for small-cell microwave equipment
Architecture	PMP
Topologies supported	Any topology, including tree, star, mesh, ring, linear
Small cells supported by a link	Each node has 2 GbE ports for connectivity to small cells or PoP backhaul connections
X2 support	Yes
Complementary technologies	Main competition comes from 60, 70/80 GHz PTP
Future product focus	Additional frequency bands. Additional capacity support

5. Cisco

Overview. Most small-cell solutions start with RAN requirements and, once they are met, move all the way to the mobile core to find a way to manage the new traffic created. Cisco's approach is complementary and, perhaps not too surprisingly, moves in the other direction, putting traffic management at the center as the unifying element to support multiple wireless interfaces and technologies within a single network. The ASR 901S router has been designed for small-cell deployments in which multiple vendors and multiple backhaul technologies coexist side by side.

As a necessary complement to the ASR 901S router, Cisco has launched the Unified MPLS for Mobile Transport (UMMT) architecture that optimizes MPLS for mobile backhaul, supporting 4G as well as legacy wireless interfaces. UMMT supports functionality that is required in 4G HetNets, including network synchronization, H-QoS, OAM, IPsec, and support for X1 and S2 interfaces.

With an end-to-end platform in place, Cisco has moved one step further to address the specific requirements of the small-cell market by creating an ecosystem of wireless backhaul vendors that support UMMT and have jointly tested the support for key features with Cisco. The initial group of vendors offers a wide range of solutions, including PTP and PMP, LOS and NLOS, and licensed and license-exempt options that collectively form a toolkit that mobile operators can use to select the best-suited solution for different environments.

Positioning. The ability to provide a common platform that can seamlessly support multiple backhaul vendors is valuable to mobile operators. They are increasingly aware of the additional complexity imposed by the necessary adoption of multiple backhaul solutions within the same footprint. In this context, small cells in a network may use multiple backhaul technologies – with different performance characteristics – to a single aggregation point. To manage traffic effectively, mobile operators need consistent functionality and tools available across all small cells.

Figure 14. Cisco ASR 901S small-cell router. Source: Cisco

Threats. Cisco's approach provides a boost to new entrants in the backhaul space: they can gain easier access to mobile operators that are traditionally inclined to select tier-one vendors. For their part, mobile operators can take advantage of the high level of innovation in small-cell backhaul that mostly comes from new entrants, while limiting their exposure to risk and increased complexity. At the same time this value proposition also has to face the competition of more entrenched, tier-one vendors that already provide the macro backhaul infrastructure to mobile operators. Cisco's success will depend on its ability to work with the most advanced backhaul vendors and present a compelling advantage over more established backhaul vendors.

Evolution. While the initial group of backhaul vendors encompasses a wide range of solutions, the addition of new vendors will be crucial to strengthening the ecosystem and driving mobile operators' support.

Cisco's small-cell backhaul ecosystem partners
BLiNQ: sub-6 GHz, NLOS
DragonWave: sub-6 GHz, microwave, 60 GHz, 80 GHz
Fastback: sub-6 GHz, NLOS
NEC: focus on 60 GHz; also sub-6 GHz and microwave
Radwin: sub-6 GHz, NLOS
Siklu: PTP, 60 GHz, 80 GHz

ASR 901S key features
Zero-touch provisioning, circuit validation, management tools
Wi-Fi interface to manage unit and limit need for physical access to the unit
Layer 2, Layer 3, MPLS deployment models
Support for OAM, IEEE 802.1ag/CFM, IEEE 1588, Y.1731, Y.1564
Installable on lampposts, walls and other street furniture in outdoor locations
Maximum power consumption 40 W. Support for POE+
Fanless, passive cooling design

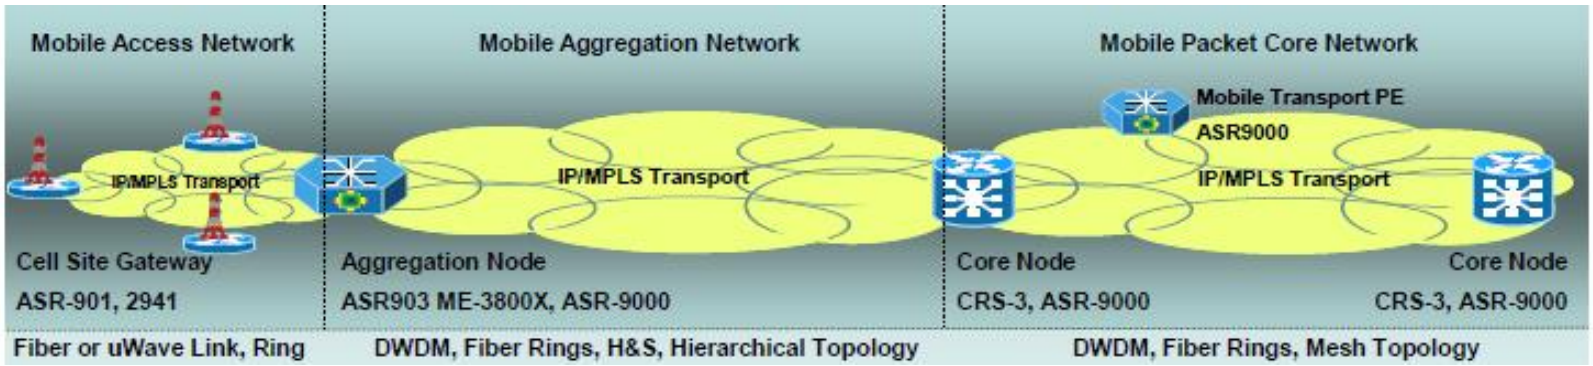

Figure 15. UMMT architecture. Source: Cisco

6. DragonWave

Overview. DragonWave's approach to small-cell backhaul is to provide a suite of solutions spanning different frequency bands and modes of operation. They include the Avenue Link PTP system operating in the 24–60 GHz band, and the Avenue Link Lite PMP system operating in the sub-6 GHz band for near-line-of-sight and NLOS application. DragonWave's solutions operate in both licensed and unlicensed spectrum bands. DragonWave designed the systems to maintain a small, integrated form factor suitable for deployments on public infrastructure at low elevation. For example, the Avenue Link includes a flat 5 in. antenna designed for small-cell backhaul. The microwave-band solution can operate at 2048 QAM modulation to provide high capacity over short distances.

Positioning. DragonWave recognizes that the deployment scenario for outdoor small-cell base stations is likely to lead to different backhaul requirements that are met by different solutions. Hence, the company provides solutions in sub-6 GHz bands that have better reach in the presence of obstructions, as well as in microwave and millimeter-wave bands that provide higher aggregate capacity. Both solutions share a single management system, DragonView, which may reduce efforts in the deployment and operation processes when different solutions are required.

Threats. Some of the bands for DragonWave's products rely on licensed spectrum that can be difficult to obtain (sub-6 GHz) or that requires licensing on a per-link basis (licensed microwave bands). The small, flat antenna in microwave-band systems would have lower gain and wider beamwidths than the typical high-gain parabolic antennas used in microwave links, which may lead to a greater need for frequency coordination to manage the higher risk of interference.

Evolution. DragonWave's focus is to continue to simplify the deployment aspects of small cells and to reduce costs. Integration of backhaul and access into a single unit is a way DragonWave is exploring to reduce the number of truck rolls, deployment time, and the operational costs.

Figure 16. Avenue Link Lite NLOS small-cell backhaul system in sub-6 GHz frequency bands. Source: DragonWave

Features	Avenue Link	Avenue Link Lite
Market focus	Outdoor metro, LOS, licensed and unlicensed	Outdoor metro, NLOS, licensed and unlicensed
Spectrum band	24–60 GHz	2–6 GHz
LOS/NLOS requirements	LOS	NLOS
Channel sizes	7–56 MHz	10–40 MHz
TDD/FDD	FDD	TDD
Modulation	Up to 2048 QAM	Up to 64 QAM
Capacity	> 500 mbps	230 mbps
Latency	0.1–0.2 ms	1–2 ms
Antenna specs/configuration	5 in. antenna	6 in. antenna
Integrated/external antenna	Integrated and external supported	
Hop length	Dependent on network, capacity and band Target of < 1 km	Dependent on network, capacity and band Target of <1 km in NLOS small-cell applications
Size	8 x 8 x 5.7 in	7.5 x 7.5 x 2 in
Weight	7.5 lb	< 4 lb
Form factor	Single enclosure	
Integration with small cell	In talks with small-cell vendors; currently 2 boxes	
Power consumption	< 50 W	< 25 W
Installation	< 15 min	
Architecture	PTP	PMP
Topologies supported	Ring, hub, daisy chain	
Small cells supported by a link	Dependent on capacity requirements, typically > 10	Dependent on capacity requirements, typically 1–3
X2 support	Yes, not tested	
Competing technologies	Fiber	
Future product focus	Improved size and cost	Expanded bands, synchronization capabilities
Small-cell vendor partners	Nokia Siemens Networks, Cisco	

7. Intracom

Overview. Intracom's approach to small-cell backhaul relies on area-licensed microwave bands, which are typically acquired at an upfront charge, to operate in PTP and PMP modes in a wide region. The licensing regime for these bands, typically in the 26, 28, 32, and 42 GHz frequencies, minimizes spectrum acquisition and coordination costs, especially when a relatively large number of links is required in an area. The StreetNode platform operates in PTP, PMP and relay modes to reach deep into the urban clutter where small cells are deployed. StreetNode features a dynamic bandwidth-allocation technology that is useful in PMP operation when serving multiple small cells with different capacity requirements. It also incorporates antenna auto-alignment to minimize deployment and installation costs.

Positioning. StreetNode seeks to provide wireless operators with a carrier-grade small-cell wireless backhaul solution that minimizes cost in four ways. First, it uses area-licensed microwave, which enables deployment of PMP links, in itself a cost-saving feature in that it avoids per-link licensing and coordination expenses. Second, the same hardware platform can be configured in different modes (PTP, PMP and relay), which streamlines operational processes. Third, the product incorporates antenna auto-alignment to automatically align two ends of the LOS microwave link, thereby simplifying the installation process by reducing reliance on highly skilled personnel and shortening the time required for deployment. Fourth, redundancy at the hub site increases the availability at the most critical node in a PMP deployment.

Threats. While the multimode capability of StreetNode is certainly advantageous, it also may require more units on a pole to reach deep into the urban clutter through the PTP and relay configuration. This is liable to increase the cost of deployment, especially when asset owners charge on a per-mounted-unit basis.

Evolution. Intracom plans to roll out StreetNode in additional frequency bands to cover the entire area-licensed microwave space. It also plans to provide a greater number of the features typically used in carrier-grade solutions, such as MPLS.

Figure 17. StreetNode area-licensed PMP microwave backhaul system for small cells featuring automatic antenna alignment. Source: Intracom

Features	StreetNode
Market focus	Microwave technology at street level
Spectrum band	Area-licensed spectrum at 26, 28, 32 and 42 GHz
LOS/NLOS requirements	LOS
Channel sizes	56 MHz
TDD/FDD	FDD
Modulation	Up to 1024 QAM
Capacity	540 mbps
Latency	0.3–0.6 ms, depending on operation mode
Antenna specs/configuration	ETSI EN 302 217-4-1 Class 2
Integrated/external antenna	Integrated
Maximum terminals per hub	30
Hop length	Dependent on network capacity and band, typically 1 km
Size	At the small cell and at the aggregation point: 263 x 143 x 166 mm (H x W x D) At a high-end aggregation point: 266 x 237 x 95 mm
Weight	At the small-cell end and at the aggregation point: 2.3 kg At a high-end aggregation point: 4.1 kg
Form factor	Single enclosure
Integration with small cell	Small-cell backhaul is a separate unit for an open and flexible approach in the first phase of small-cell deployments
Power consumption	< 25 W
Installation	> 30 min. Innovative antenna auto-alignment simplifies unit mounting, minimizes time and effort, and ensures optimum performance. Zero-touch provisioning reduces installation time and required personnel, while avoiding configuration faults. Real-time verification procedures ensure the installation is accomplished in a single visit
Architecture	Flexible, single hardware unit that is software defined for PMP hub, PMP terminal or PTP operation
Topologies supported	PTP, PMP and street-level relays
Small cells supported by a link	The number of small cells connected to an aggregation point depends on the expected mean busy-time traffic. Assuming mean traffic requirements of 50 mbps per small cell, a single StreetNode sector can aggregate 10 small cells. With dynamic bandwidth allocation, peak requirements of more than 500 mbps can be satisfied
X2 support	X2 interface is efficiently supported, allowing for traffic intra-switching at every node, with latency below 1 ms
Complementary technologies	StreetNode offers a complete carrier-class backhaul solution for small cells over a single area-licensed band, reaching any location at street level
Competing technologies	Sub-6 GHz NLOS, millimeter-wave PTP
Future product focus	Providing additional area-licensed spectrum bands; providing advanced networking features such as MPLS

8. Proxim Wireless

Overview. Proxim's main small-cell backhaul platform is the Tsunami 8200 system, which operates in the 4.9–5.9 GHz band with a channel bandwidth ranging between 5 and 40 MHz. While the system is based on the same IEEE 802.11n standard on which Wi-Fi is based, Proxim made a number of modifications to optimize the solution for backhaul. By reducing the overhead associated with the standard Wi-Fi MAC layer, the system's capacity reaches 250 mbps in PTP and 240 mbps in PMP.

Positioning. Proxim aims to provide low-cost systems for small-cell backhaul based on COTS IEEE 802.11n baseband modems that optimize the communication protocol stacks. The Tsunami 8200 operates in the 5 GHz unlicensed band, but extensions by Proxim enable it to serve adjacent frequency bands too. Branded as WORP, these enhancements of the MAC layer increase throughput efficiency by 50% to 80% of the PHY layer capacity, and lower latency from a few to 10 ms, according to Proxim estimates. Additional features such as automatic retransmission are designed to further improve the performance of Wi-Fi in backhaul. Furthermore, the solutions operate in nonstandard Wi-Fi channel widths such as 5 MHz, which is useful in markets where the 5 and 6 GHz bands can use only smaller channel bandwidths.

Threats. Because it is an unlicensed-spectrum solution, the Tsunami 8200 can be subject to interference from unknown sources. As more consumer devices integrate Wi-Fi n in the 5 GHz band, and with the advent of Wi-Fi ac, which brings greater capacity but also increases channel bandwidth, the specter of interference increases. Operators are divided on whether carrier-grade capability is necessary for small-cell deployments. The success of Tsunami will depend on where operators fall on the cost versus performance (carrier-grade capability) tradeoff.

Evolution. Proxim is planning the release of a smaller size that would blend better with the environment for low-height small-cell deployments. The company also plans products based on the IEEE 802.11ac standard, which would increase throughput capabilities by supporting higher modulation rates and wider channels.

Figure 18. Tsunami 8200 NLOS backhaul systems used in PTP and PMP small-cell backhaul applications, providing up to 250 mbps throughput. Source: Proxim

Features	Tsunami 8000 Series
Market focus	PMP equipment is predominantly used for the access; PTP equipment is typically used for backhaul
Spectrum band	5 GHz unlicensed bands, and 6–23 GHz licensed bands
LOS/NLOS requirements	2 x 2 and 3 x 3 MIMO solutions for NLOS environments. LOS environments pose no appreciable limits
Channel sizes	5, 10, 20 and 40 MHz in unlicensed bands; standard FCC and ETSI channel assignments for licensed bands
TDD/FDD	TDD for unlicensed products, FDD for licensed products
Modulation	BPSK, 64 QAM for existing products; up to 1024 QAM for IEEE 802.11ac planned products
Capacity	Unlicensed products: 240 mbps, with 40 MHz channel, 64 QAM; the throughput scales with smaller channel sizes
Latency	< 10 ms end-to-end with nominal network loading. At full capacity and > 50 subscribers, latency may increase to 40 ms.
Antenna specs/configuration	Connectorized versions for the radios allow for any size antenna to be connected to the radio. For integrated products, the antenna gain is in the range of 15–23 dBi
Integrated/external antenna	Both are available
Maximum terminals per hub	200
Hop length	This is a function of environment and antenna selection. Unlicensed PTP links operate reliably in the unlicensed band at more than 20 miles. Range for unlicensed PMP networks is 200 m to 4.5–6 km and is dependent upon the environment. Licensed links can extend much farther, especially when 6 GHz frequencies are used
Size	Varies by product; typically less than 12 x 12 x 5 in
Weight	4–15 lb
Form factor	Small (no larger than 12 x 12 x 5 in
Integration with small cell	Not at the present time. However, the radios are small enough to fit within most operators' enclosures with a cable run to an external antenna
Power consumption	Typically < 25 W for most products; most unlicensed-band products operate on PoE
Installation	Typically < 2 hours to install a link, assuming a tower climb is not necessary
Equipment cost	List price for unlicensed-band products: $469 to $3,999; for licensed-band products: in the range of $14,000
Architecture	PTP and PMP
Topologies supported	All are supported except mesh
Small cells supported by a link	4–5, assuming they require no more than 50–60 mpbs of useable TDD throughput per link
X2 support	As the transport layer, Tsunami can support and pass X2 information between nodes.
Competing technologies	Wire: fiber or coax, when installed, eliminate the need for wireless except for redundancy purposes
Future product focus	IEEE 802.11ac will allow for greater than 800 mbps of useable throughput
Small-cell vendor partners	As a manufacturer of transport products, Proxim can work with any manufacturer of small-cell equipment and any network operator deploying small-cell equipment

9. NEC

Overview. A well-established wireless backhaul solution provider, NEC has been working with key operators and trialing small-cell backhaul solutions with them for several years. The breadth of NEC's portfolio means that it is well prepared to present the toolkit approach that many operators are asking for, as it recognizes that no single technology will suffice. NEC has solutions ranging from conventional PTP microwave operating in the traditional LOS 6–42 GHz bands, to NLOS solutions in the sub-6 GHz band. But it is NEC's new 60 GHz iPASOLINK SX PTP millimeter-wave product that is the focus of its small-cell backhaul strategy. NEC believes it can deliver virtually future-proof capacity across the 9 GHz of bandwidth available in the 60 GHz band, and the light-licensing regime is also attractive to many operators.

Positioning. NEC is able to address virtually any backhaul requirement with its portfolio of products. But with choice can come indecision, and NEC is keen to make sure that its customers can assemble multifaceted systems easily with a range of network design and configuration tools, which include self-organizing features. NEC also offers the iPASOLINK GX, which is a miniature, outdoor router product that can support a wide range of topologies, giving its customers the full range of star, tree, mesh and partial meshing options. Finally, NEC offers an NMS platform that extends across both its wireless and its optical backhaul and network products.

Threats. A large installed base, broad wireless equipment portfolio, and mixed optical and wireless equipment strategy over a common NMS platform support NEC's small-cell products, but its size and multi-business unit structure could limit the company's agility to respond to smaller, more aggressive competitors.

Evolution. NEC plans to deliver continuous improvement in capacity, reduce the physical footprint, and expand the OAM&P integration with its own small-cell portfolio and the participation in operator-initiated partnerships.

Figure 19. iPASOLINK SX 60 GHz microwave backhaul system for small cells. Source: NEC

Features	iPASOLINK 100 / 100E / 200 / 400 / 1000	iPASOLINK SX	iPASOLINK EX	NLOS radio products	iPASOLINK GX	MS5000 NMS and resource optimization tools
Market focus	Urban small-cell backhaul aggregation (rooftop links), rural backhaul of small cells	Urban small-cell backhaul, street-level connectivity	Urban small-cell backhaul aggregation (rooftop links)	Urban small-cell backhaul, street-level connectivity, rural backhaul of small cells	Outdoor nodal aggregation and routing, branching and mesh topologies	Unified, multilayer network management platform provides a common OAM&P framework across the full toolkit
Spectrum band	6–42GHz	60 GHz	70–80 GHz	Sub-6 GHz	N/A	N/A
LOS/NLOS requirements	LOS	LOS	LOS	NLOS	N/A	N/A
Channel sizes	Up to 56 MHz	50 MHz	50, 250, 500 MHz	Up to 40 MHz	N/A	N/A
TDD/FDD	FDD	FDD	FDD	TDD	N/A	N/A
Modulation	Up to 2048 QAM (with AMR)	Up to 256 QAM (with AMR)	Up to 256 QAM (with AMR)	Up to 256 QAM	N/A	N/A
Capacity	Up to 500 mbps single channel; multi-gbps with spatial aggregation	330 mbps in 50 MHz channel; up to 1 gbps in 50+ MHz	Up to 10 gbps	More than 500 mbps	N/A	N/A
Hop length	Tens of km	Up to 1 km	Up to 4 km	Up to 4 km	N/A	N/A
Form factor	Split-mount IDU/ODU	All-in-one integrated	AOR and antenna	All-in-one integrated	All-in-one integrated	Software platform
Integration with small cell	No	Optional	No	Optional	Optional	N/A
Architecture	PTP	PTP	PTP	PTP/PMP	Router/nodal	N/A
Topologies supported	All	All	All	All	All	N/A
Small cells supported by a link	Depends on configuration of LTE (channel size and MIMO)			N/A		
X2 support	Yes	Yes	Yes	Yes	Yes	N/A
Complementary technologies	Full carrier-Ethernet feature set is implemented within the product; high-speed optical transceiver available as an option			End-to-end OAM&P NMS platform capabilities		

10. Siklu

Overview. Siklu is pioneering cost-effective, high-capacity, PTP millimeter-wave wireless backhaul solutions. It has been driving down the price of E-band millimeter-wave links with a silicon radio and baseband of its own design. Siklu has expanded its range with a 60 GHz product to add the millimeter-wave V-band. The 60 GHz band has attracted a lot of attention among operators due to its light-licensing or license-exempt regulations in most countries. Millimeter-wave products are attractive due to their massive capacity, made possible by the abundance of bandwidth available in the E- and V-bands. Siklu delivers 1 gbps capacity from its links using 500 MHz channels, with a latency of 350 µs. Siklu's patented antenna for its 60 GHz product, EtherHaul-600T, demonstrates the suitability of millimeter-wave radios for deployments on street furniture. EtherHaul-600T has been designed to be installed in street-level scenarios and to cope with the challenges of the twist, tilt and sway of poles. With automated alignment tools, Siklu claims that installation times can be less than 60 minutes.

Figure 20. EtherHaul-600T 60GHz backhaul system for small cells. Source: Siklu

Positioning. EtherHaul-600T is a palm-sized 60 GHz, all-outdoor small-cell backhaul product that enables rapid deployment anywhere, from street lamps to rooftops. It employs a number of networking features that enable it to act as a node in the network, supporting star, tree and mesh topologies, mostly targeted at dense urban environments. A second solution, the EtherHaul-1200F and 1200T radios, operate at the higher 71–86 GHz frequencies to provide wireless PTP gigabit Ethernet connectivity with MEF-compliant networking and QoS.

Threats. Siklu's exclusive focus on millimeter-wave technology gives a market advantage, but it also forces the company to partner with NLOS or PMP vendors to deliver a complete solution to operators.

Evolution. We expect Siklu will continue to put pressure on conventional PTP and the new small-cell backhaul entrants by pushing the link price of their high-capacity links lower, and by making them easier to install.

Feature	EtherHaul–600T (EH–600T)	EtherHaul–1200T (EH–1200T) / EtherHaul-1200F (EH-1200F)
Market focus	Small-cell backhaul, LOS	Macro-cell backhaul last mile, small-cell backhaul
Spectrum band	57–66 GHz (V-band)	71–76 GHz; 71–76 / 81–86 GHz (E-band)
LOS/NLOS requirements	LOS	LOS
Channel sizes	250 MHz, 500 MHz	250 MHz, 500 MHz
TDD/FDD	TDD	TDD/FDD
Modulation	QPSK, QAM 16, QAM 64	
Capacity	1 gbps in 500 MHz channel	
Latency	350 ms	250–350 ms
Antenna specs/configuration	Integrated Cassagrain; parabolic to cope with twist, tilt, sway; ETSI/FCC compliant	
Integrated/external antenna	200–500 meters, rain-zone dependent	Up to 3,000 meters, rain-zone dependant
Size	15 x 15 x 7 cm (H x W x D)	ODU (H x W x D): 24.5 x 22.5 x 7 cm ODU + antenna: 31 cm (dia. x depth): 31 x 13 cm
Weight	1.5 kg	ODU: 2.5 kg, ODU + antenna (31 cm): 4 kg
Form factor	All-outdoor, single-box, integrated antenna	
Integration with small cell	The EH-600T and EH-1200T are based on Siklu's all-silicon building blocks. This technology is modular, allowing very straightforward and beneficial integration with small-cell access equipment	
Power consumption	25 W max	35 W max / 45 W max
Installation	< 60 min, with auto-alignment tool. Mounting kit designed for both wall-mount and pole-mount scenarios. Wide range of azimuth and elevation fine alignment enabled. Local or remote configuration	
Equipment cost	Starting at $2,000; $1,000–$1,500 with volume	$3,000 (EH-1200T) / $5,000 (EH-1200F)
Architecture	PTP	
Topologies supported	Tree, daisy-chain, ring, mesh	
Small cells supported by a link	3 GbE ports per unit / small cell	3 GbE ports per unit / small cell
X2 support	Capacity and latency are well within the requirements of the X2	
Complementary technologies	E-band complements the EH-600T in long-distance links. NLOS backhaul complements it where there is no LOS	
Competing technologies	High-capacity PMP NLOS solutions	
Future product focus	Zero-touch installation and configuration capabilities, incl. installation tools such as the auto-alignment tool, auto configuration and SON capabilities. Multi-gigabit capacity and spectral efficiency, targeting 2 Gbps in 500 MHz channels	

12. SOLiD Technologies

Overview. SOLiD Technologies provides DAS solutions that expand the utility of fiber. With SOLiD's Infinity Access, up to 16 distinct channels can be multiplexed on a single fiber strand. The channels can carry different types of traffic, including CPRI, a high-data-rate protocol (up to 3 gbps) used to connect baseband modules to remote radio headends, and Ethernet. Infinity Access consists of three modules: First, the central Optical Line Terminal is the hub, which multiplexes multiple protocols onto a single fiber strand (GbE, E-PON, SONET and CPRI). The second module is the Optical Network Terminal, located close to the small cell (be it a Wi-Fi access node, an RRH, or a compact base station); it converts fiber-optic light signals to the desired protocol. The third module is the Access Passive Splitter, which is a bidirectional DWDM (de)multiplexor.

Positioning. The capacity that fiber provides is second to none. Hence it is the ideal solution in backhaul applications, because significant bandwidth can be provided for different services. In backhaul, Ethernet is the most common interface on base stations with data rate requirements ranging from the tens to the hundreds of mbps. High-data-rate protocols that extend throughput into the gbps range, like CPRI and OBSAI, are used in remote-radio head-end deployments, an application commonly referred to as fronthaul. SOLiD allows operators with fiber assets to deploy different types of small cells and leverage a common platform for transport. Furthermore, the low latency and tight jitter parameters facilitate the deployment of LTE-Advanced features like CoMP, which require tight synchronization.

Threats. Availability of fiber is the main barrier to deployment. Fiber can be very expensive to deploy in markets where trenching is mandated. It can also take a long time to deploy, thereby slowing down the operator's plans to increase network capacity. On the continuum of cost and performance tradeoffs, the Infinity Access is firmly positioned on the performance side in areas where fiber is available.

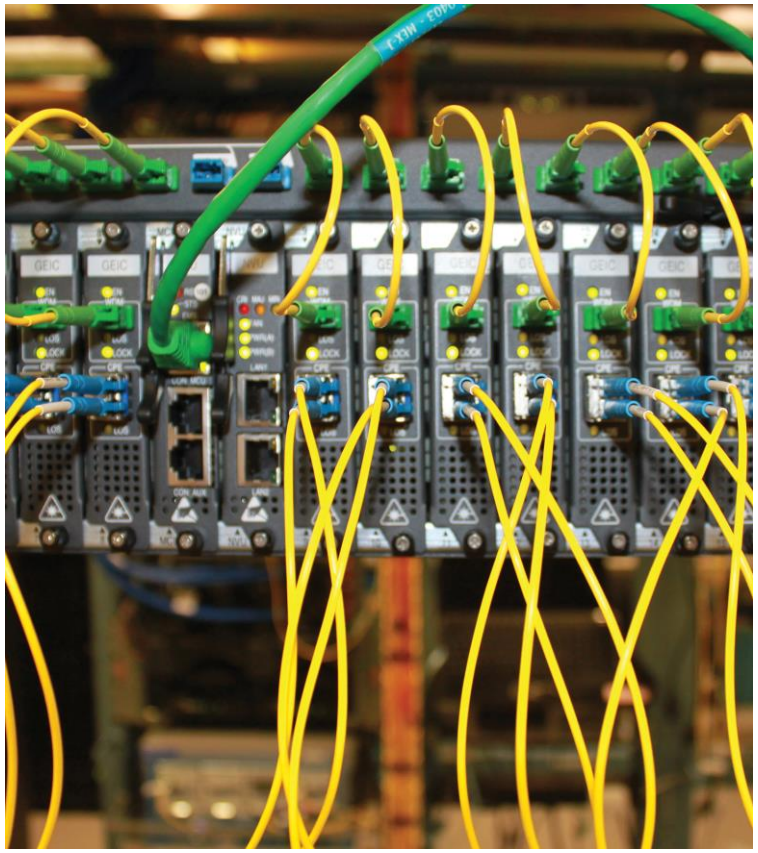

Figure 21. SOLiD Technologies Infinity Access DWDM fiber multiplexing system with capability to backhaul 16 small cells on a single fiber strand. Source: SOLiD Technologies

Evolution. SOLiD's future plans center on enhancing support for multiple transport protocols and increasing the number of supported channels on a single fiber strand.

Features	Infinity Access
Market focus	Metro fiber
Optical characteristics	C-band (uplink) and L-band (downlink)
Access technology	DWDM
Dispersion tolerance	1700 ps
Distance	Maximum 1 km between drops for 16 drops
System capacity	16 channels
Channel sizes	1.25–2.5 gbps
Capacity	20–40 gbps
Latency	121 ns
Maximum terminals per hub	16
Power consumption	12 W
Installation	Plug and play
Equipment cost	Link cost about $3,000
X2 support	Yes
Complementary technologies	Carrier Ethernet, CPRI
Competing technologies	Other Ethernet solutions
Future product focus	Simultaneous support for CPRI and carrier Ethernet on a single fiber strand
Small-cell vendor partners	Agnostic

13. Tellabs

Overview. Small cells have the potential to introduce a significant increase in complexity and cost. Tellabs strives to minimize costs and improve network performance and efficiency by leveraging its experience in supporting multi-vendor, multi-generation networks. To address the need for multiple technologies and multi-vendor coexistence, three elements are central to Tellabs' small-cell strategy: interoperability across vendors to give operators flexibility; backhaul pre-aggregation to manage backhaul traffic effectively; and partnerships with backhaul vendors for consistent functionality across equipment.

Positioning. At the core of its solution is Tellabs' outdoor, hardened 8602 Smart Router. It can be co-located with a cluster of small cells, and provide connectivity and pre-aggregation to wireless or optical backhaul connections. The 8602 Smart Router provides synchronization, SLA management, and SON and OAM&P tools that enable the third-party vendors to focus on their core competencies. The high level of integration and the availability of a consistent set of network management tools across multivendor, multi-RAT networks brings reduced complexity, more-compact equipment form factors and, in turn, lower costs. At the same time, Tellabs has developed a partnership program open to backhaul vendors that integrate Tellabs know-how into their products, either as a hardware module or as software.

Threats. Providing a unified platform to manage small cells within the broader macro networks is a valuable proposition that Tellabs offers to mobile operators – especially for those who already use Tellabs solutions in their existing network infrastructure. But as requirements change and vendor dynamics evolve with the introduction of small cells, mobile operators may become more selective in their choice of backhaul vendors. Tellabs' success in this area will depend on its ability to partner and deliver interoperable solutions with the key backhaul vendors.

Figure 22. Tellabs 8602 Smart Router for small-cell backhaul deployments. Source: Tellabs

Evolution. Tellabs plans to enhance the capabilities and performance of its automated network installation and configuration tools. It also intends to expand its backhaul-vendor partnership program.

Features	Tellabs 8602
Market focus	Intermediate and small-cell backhaul
Functionality	IP VPN (RFC 4364), VRF lite, Ethernet/VLAN pseudo-wires, Y.1731 frame loss, frame delay and frame-delay variation support, Ethernet MAC switching, BGP multipath for load balancing, IPv4 load balancing over RSVP-TE tunnels for IP VPN routes
Resiliency	1:1 RSVP-TE LSP protection, FRR, IP load balancing (ECMP), Ethernet link aggregation IEEE 802.1AX, Ethernet pseudo-wire redundancy
Synchronization	ITU-T G.813 option 1, ITU-T G.8262, Telcordia (GR-1244) Stratum-3, Synchronous Ethernet, SSM over Ethernet (G.8264), IEEE 1588v2 Precision Time Protocol, boundary clock for phase synch, GPS SFP support
Traffic management	DiffServ support for up to 7 traffic classes, 7 queues per interface with MPLS traffic engineering (DS-TE), IEEE 802.1P, IEEE 802.1Q mapping to IP or MPLS, policing and shaping, VLAN shaping
Security	RSVP-TE authentication, ACLs
Physical dimensions	320 x 130 x 50 mm
Power and cooling	100-253 VAC, power consumption: typical 20 W, maximum 30 W
Future product focus	Greater support for automated installation and configuration as well as high-performance network management

14. VubIQ

Overview. VubIQ's HaulPass systems provide a unique multicarrier capability with both 60 GHz and 5 GHz transceivers in one package. The 60 GHz link provides capacity up to 1.25 gbps by leveraging the very wide channel bandwidth available in that band (up to 2 GHz). The 5 GHz system is based on IEEE 802.n and optimized to improve performance in a PTP deployment scenario. VubIQ factored into its product design the cost and time associated with small-cell deployments, and built an antenna auto-alignment feature in the HaulPass SC to align the two ends of the wireless link without the need for highly skilled deployment personnel. Furthermore, the system offers a host of important capabilities for small-cell backhaul, including synchronization capabilities (e.g., IEEE 1588v2 and SynchE) and multiple Ethernet ports for different deployment configurations.

Positioning. VubIQ's main differentiation lies in combining millimeter-wave's high capacity with unlicensed sub-6 GHz support to increase the availability of the 60 GHz backhaul link in street-level deployments, where pole sway and path obstructions may reduce the availability of 60 GHz connectivity. The introduction of antenna auto-alignment simplifies the installation process, cutting the time required on site and reducing reliance on skilled personnel. The use of unlicensed spectrum eliminates associated licensing and coordination costs of licensed spectrum and works to lower the total cost of ownership. The high propagation losses in the 60 GHz band (about 20 dB/km) limit the potential of interference, particularly in short-distance deployments.

Threats. The 5 GHz failover mode would have a different system gain and range than the 60 GHz link, as well as higher susceptibility to interference from an increasing base of consumer devices with 5 GHz Wi-Fi capability. The capacity of the 5 GHz carrier would vary according to extraneous conditions and potentially allow only the most important traffic on the 60 GHz link to be carried in case of failover.

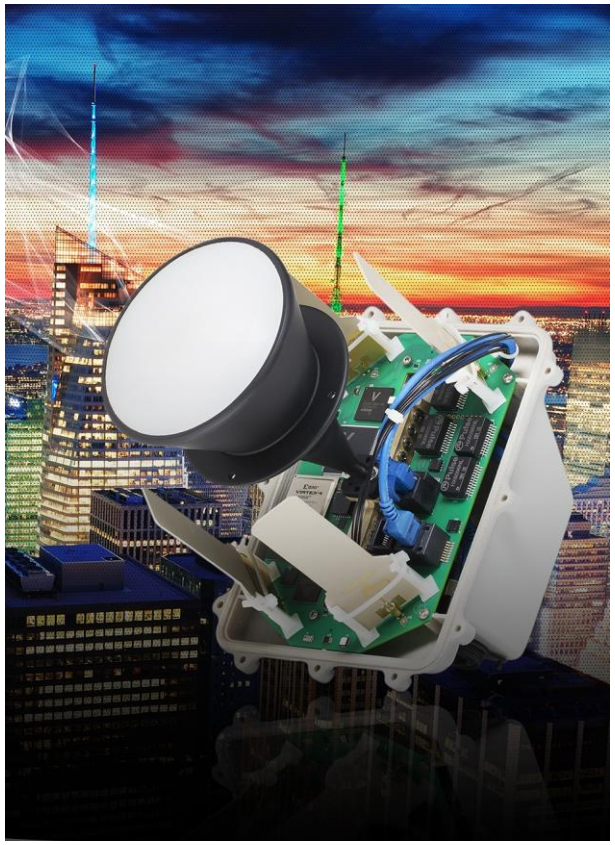

Figure 23. HaulPass SC small-cell backhaul system featuring dual-band operation in 60 GHz and 5 GHz bands, with automatic antenna alignment. Source: VubIQ

Evolution. VubIQ's focus is on providing greater bandwidth in both the 60 GHz and 5 GHz bands. With a target of 10 gbps aggregated capacity, the door should be wide open for the HaulPass to be deployed in additional applications such as fronthaul.

Features	HaulPass SC and microSC
Market focus	Outdoor metro, LOS and NLOS, unlicensed
Spectrum band	60 GHz and 5 GHz
LOS/NLOS requirements	LOS and NLOS
Channel sizes	60 GHz: up to 2 GHz; 5 GHz: up to 40 MHz (defined by IEEE 802.11n)
TDD/FDD	60 GHz: FDD; 5 GHz: TDD
Modulation	60 GHz: OOK; 5 GHz: up to 64 QAM
Capacity	60 GHz: 1.25 gbps; 5 GHz: 4 x 4 MIMO 600 mbps
Latency	60 GHz: 20 ms; 5 GHz: 1 ms
Antenna specs/configuration	60 GHz: 35 dBi gain; 5 GHz: 7 dBi gain per channel (x 4)
Integrated/external antenna	Integrated
Maximum terminals per hub	PMP makes use of multiple switched ports
Hop length	60 GHz: 500 m; 5 GHz: 1 km
Size	HaulPass SC: 9.25 in. diameter, 11 high (< 7.5 liters); HaulPass microSC: 8.6 x 6.8 x 5.4 in (< 4 liters)
Weight	HaulPass SC: 3.5 kg (7.6 lb); microSC: 1 kg (2.2 lb)
Form factor	HaulPass SC: cylindrical dome; microSC: small-footprint box
Integration with small cell	Interface is GbE (SC: 4 ports; microSC: 2 ports)
Power consumption	< 28 W
Installation	HaulPass SC: less than 15 minutes; microSC: 2 hours, with crew at each end of the link
Equipment cost	Cost per node anticipated to drop below $1,000 for volume orders
Architecture	PTP and PMP
Topologies supported	All network topologies can be supported with provided multiple GbE ports. Configurable Layer 2 switching
Small cells supported by a link	Depends on the operating bandwidth for each small cell. Up to 1.6 gbps aggregated bandwidth
X2 support	Depends upon the X2 interface requirements at the Ethernet level. Full Ethernet switching
Competing technologies	Other 60 GHz solutions, other NLOS solutions, fiber
Future product focus	Up to 10 gbps aggregated bandwidth

15. Acronyms

3G	Third generation
4G	Fourth generation
5G	Fifth generation
ACL	Access control lists
AMR	Adaptive multi-rate
ANDSF	Access network discovery and selection function
AOR	All outdoor radio
AP	Access point
AP	Access point
BGP	Border gateway protocol
BPSK	Binary phase-shift keying
CAGR	Compound average growth rate
CEPT	Conference of European Posts and Telegraphs
CoMP	Coordinated multi-point
COTS	Commercial off-the-shelf
CPRI	Common public radio interface
DAS	Distributed antenna system
DatE	Data at the Edge
Diffserv	Differentiated services
DPI	Deep packet inspection
DWDM	Dense wavelength division multiplexing
ECMP	Equal-cost multi-path
eICIC	Enhanced inter-cell interference coordination
E-PON	Ethernet passive optical network
ETSI	European Telecommunications Standards Institute
E-UTRAN	Evolved universal terrestrial radio access network
FCC	Federal Communications Commission
FDD	Frequency division duplex
FEC	Forward error correction
FRR	Fast reroute
GbE	Gigabit Ethernet
GGSN	Gateway GPRS support node
GigE	Gigabit Ethernet
GPS	Global positioning system
HetNet	Heterogeneous network
H-QoS	Hierarchical QoS
IDU	Indoor unit
IEEE	Institute of Electrical and Electronics Engineers
IP	Internet protocol
IPsec	Internet protocol security
IT	Information technology
Iub	The interface between NodeB and the RNC
Iuh	The interface between home NodeB and the home NodeB gateway
KPI	Key performance indicator
LOS	Line of sight
LSP	Label switched path
LTE	Long term evolution
MAC	Medium access control
MACsec	MAC security
MIMO	Multiple-input and multiple-output
MPLS	Multiprotocol label switching
MPTMP	Many point to many point
NGH	Next Generation Hotspot
NLOS	Non line of sight
NMS	Network management system
OAM	Operations, administration, and maintenance

OAM&P	Operations, administration, maintenance, and provisioning
OBSAI	Open base station architecture initiative
ODU	Outdoor unit
OTT	Over the top
PHY	Physical [layer]
PoE	Power over Ethernet
PoE+	PoE plus
PoP	Point of presence
PTP	Point to point
QAM	Quadrature amplitude modulation
QoS	Quality of service
QoS	Quality of service
QPSK	Quadrature phase shift keying
RAN	Radio access network
RFC	Request for comments
RNC	Radio network controller
RRH	Remote radio head
RSVP-TE	Resource reservation protocol - traffic engineering
S1	LTE interface between an eNodeB and a mobility management entity or a serving gateway
SAE_GW	System architecture evolution gateway
SFP	Small form-factor pluggable
SIM	Subscriber identity module
SoC	System on a chip
SONET	Synchronous optical network
SSM	Synchronization status message
TCO	Total cost of ownership
TDD	Time division duplex
VLAN	Virtual local area network
VNI	Visual Networking Index
VPN	Virtual private network
VRF	Virtual routing and forwarding
WWDM	Wireless wave division multiplexing
SLA	Service level agreement
SON	Self-organizing network
UMMT	Unified MPLS for mobile transport
URL	Uniform resource locator
uWave	Microwave
WORP	Wireless Outdoor Router Protocol
GPRS	General packet radio services
IEEE 1588	Precision time protocol (PTP)
IEEE 802.1ag	Connectivity fault management (CFM)
ITU	International Telecommunication Union
ITU-T	ITU Telecommunication Standardization Sector
Y.1731	ITU-T recommendation Y.1731
Y.1564	ITU-T recommendation Y.1564

About Senza Fili

Senza Fili provides advisory support on wireless data technologies and services. At Senza Fili we have in-depth expertise in financial modeling, 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.

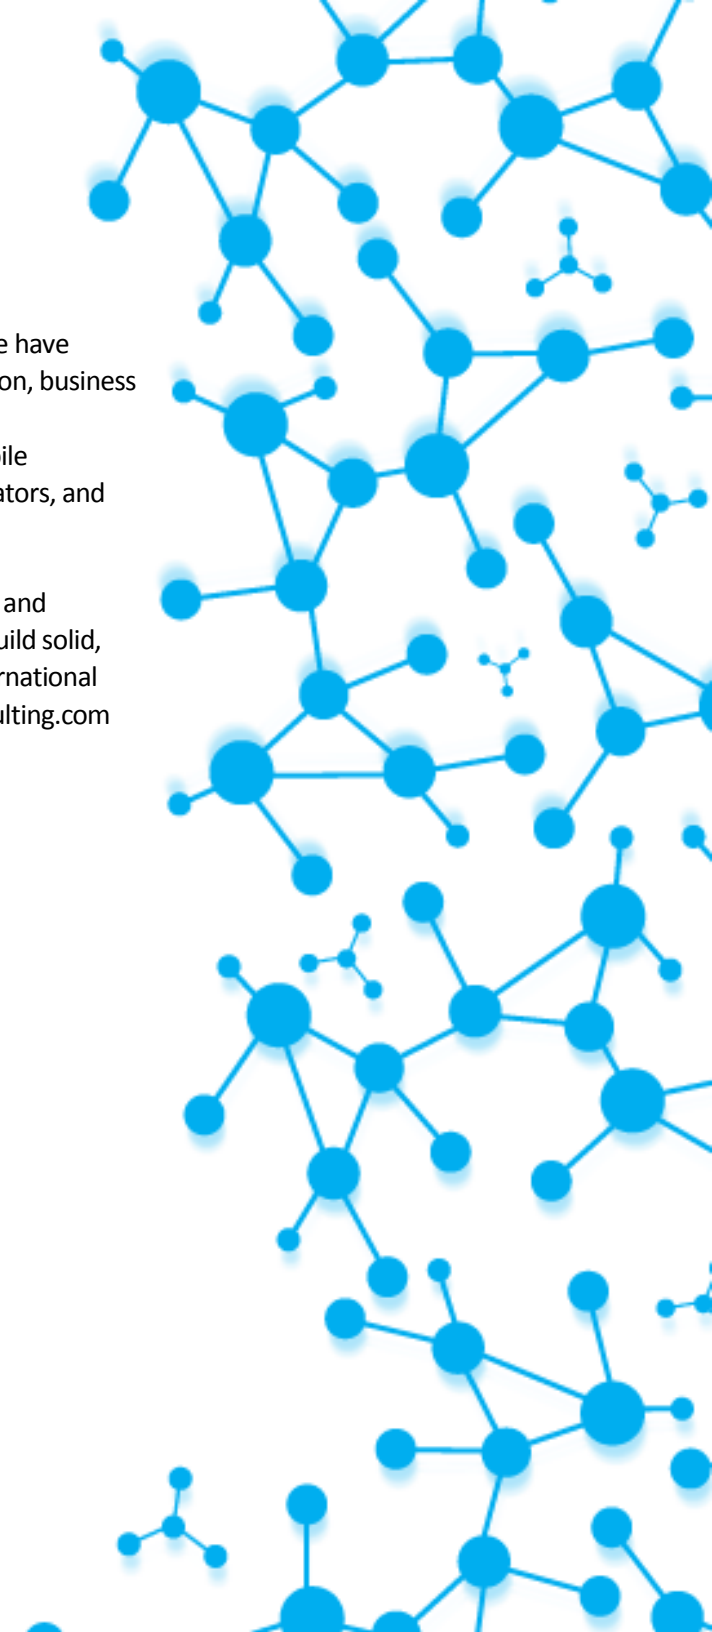

About the authors

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.

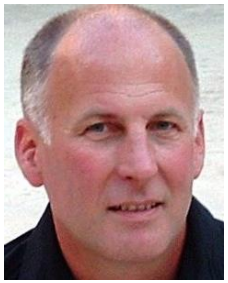

Lance Hiley is an independent wireless marketer and analyst working on behalf of Senza Fili. He is an expert in wireless backhaul technologies and worked with companies and industry organisations to bring clarity to backhaul products and issues. He has given numerous presentations at conferences, and has written articles and papers on wireless broadband technologies and products. He can be contacted at lance.hiley@gmail.com

Frank Rayal is a telecom industry professional with over 20 years of experience working with network operators and system vendors to develop and deploy innovative wireless solutions. He is a founding partner at Xona Partners a boutique management and technology advisory firm specialized in TMT and a founding member of small cell backhaul pioneer BLiNQ Networks. Frank held senior product management, marketing and business development positions at Ericsson, Redline, and Metawave. He holds a BS in Electrical Engineering from Case Western Reserve University, Cleveland, OH, and a MSc in Electrical Engineering and an MBA from the University of Toronto, Canada. Frank is a Senior Member of the IEEE, and a member of Professional Engineers Ontario.

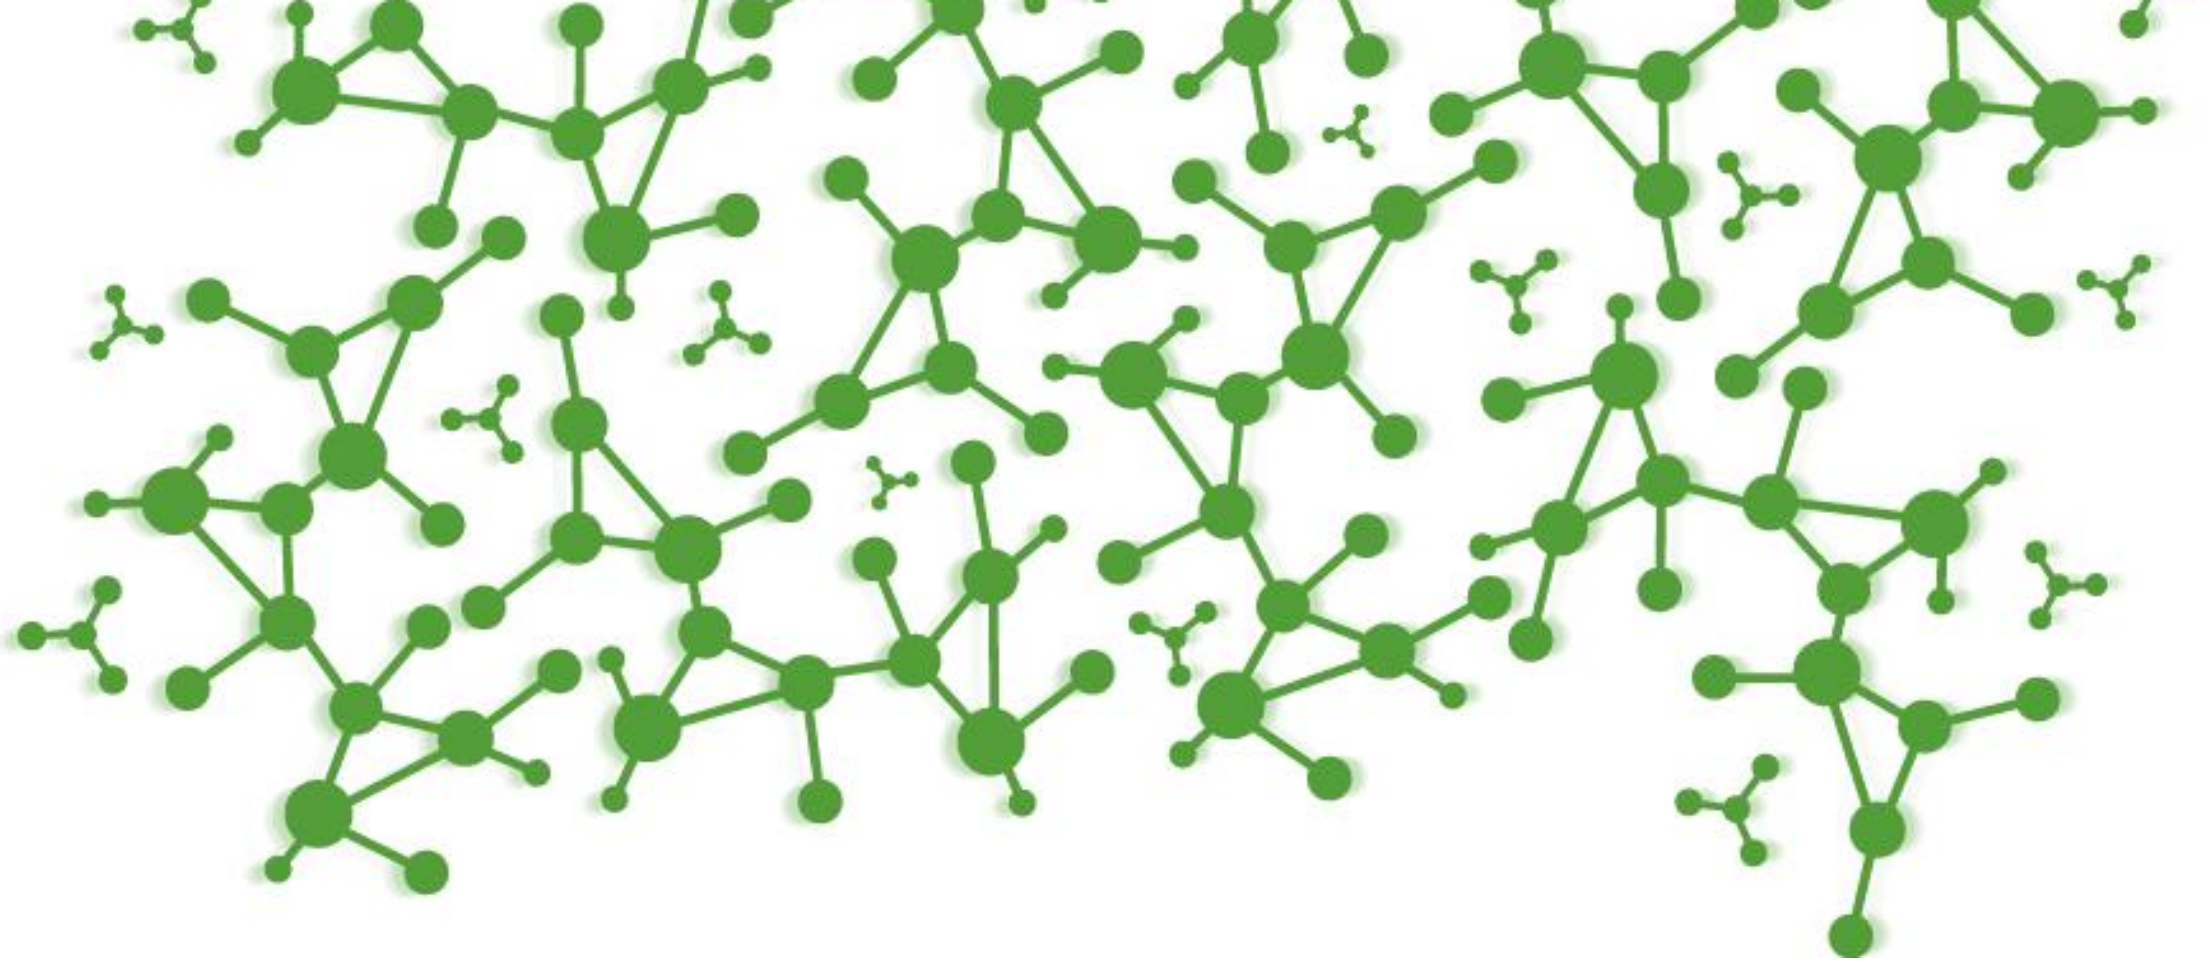

Small-cell backhaul:
Industry trends and market overview

Part III. Vendors' conversations

By Monica Paolini, Lance Hiley and Frank Rayal
Senza Fili

SENZA FILI CONSULTING

Table of contents

Senza Fili thanks the following companies for sponsoring and collaborating on this report:

III. Vendors' conversations

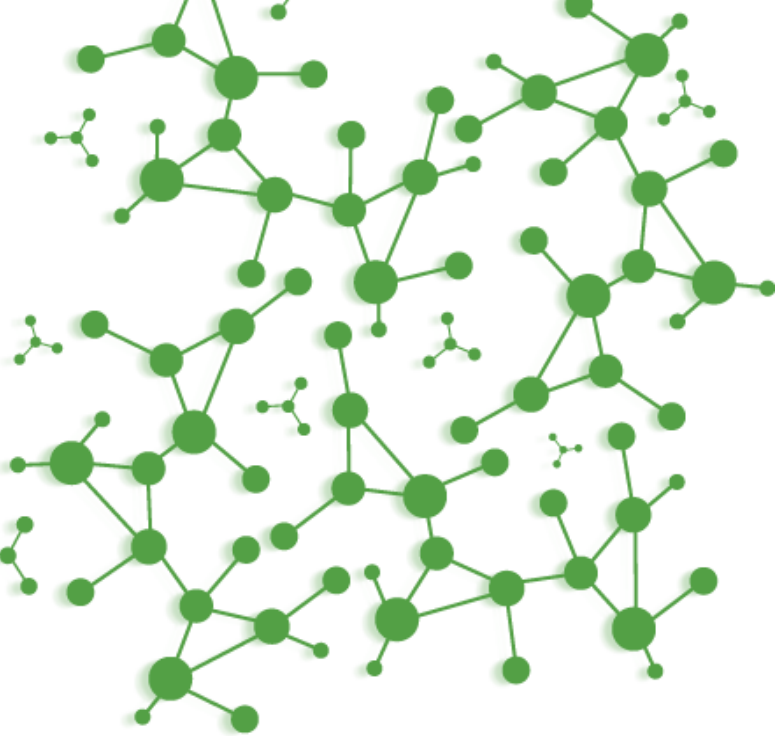

Optimizing backhaul capacity

A conversation with Richard Lord, CTO & Mats Vilander, VP EMEA Sales & Head of Global DatE, Altobridge

By Frank Rayal
Senza Fili Consulting

Frank Rayal. Today we are joined by Altobridge, and we have on the call with us Richard Lord, CTO of Altobridge, and Mats Vilander, VP of EMEA Sales and Head of Global DatE (Data-at-the-Edge) product. This conversation is part of Senza Fili's report on small-cell backhaul that gives an update on small-cell backhaul solutions and on the evolution of the mobile operator requirements for small-cell backhaul.

Altobridge is a leading developer of wireless network solutions that reduce network operating costs in delivering mobile voice, mobile broadband and big data services. The company developed a series of patented and patent-pending intelligent mobile broadband data optimization and delivery technologies that reduce backhaul costs across wireless networks. Altobridge is a privately held company headquartered in Ireland, with offices in the US, China, Malaysia and Indonesia.

Welcome, Richard and Mats.

Let me start by asking you to give us an introduction about your solution – and that would be the Data-at-the-Edge product – and how it works in the small-cell environment.

Mats Vilander. Thank you very much for that. Data-at-the-Edge, first of all, is a backhaul optimization product. It operates by saving backhaul cost in one environment, where it is driven out of satellite for rural backhaul, and secondly, in 3G and LTE congested networks, where you have a massive amount of content that is being distributed to the end users, and the down time is one of the most driving factors that prevents the end user from getting his content in time. So it is driven out of quality experience. We are also the only company that supports the local call interception, the LCI, which is already an existing 3GPP standard allowing the operators to overview the traffic over the Iuh and Iub interface. We are the only backhaul optimization company for byte caching that is supporting the Iub interface.

Frank Rayal. I take it that your product is a software solution – do you port it onto different vendors' hardware? Can you elaborate on exactly where the product is located in the mobile network?

Richard Lord. The solution is based on byte caching, which is a symmetrical caching technology – in other words, the caching engine required at both sides of the link to be optimized. In our solution we are optimizing the link between the remote base station and the core network.

As Mats indicated, we support all the major 3GPP cellular interfaces, so we are one of the few companies that have ported this solution onto an Iub interface. We also support the Iuh interface, which is for small cells, and the S1 interface, which is for LTE small cells and macro cells.

The solution is composed of two modules: one is called the central gateway and the other is called the remote gateway. The central gateway sits near the core network, and the remote gateway sits next to the base station. Basically, it is what people might refer to as a bookend solution, so you have a monitor on either side of the backhaul link that you are trying to optimize.

Frank Rayal. Backhaul is a major issue that is gating small-cell deployment. How does your product address the backhaul challenge, and how does it enable operators to deploy small cells?

Richard Lord. Backhaul is always, or very often, an issue with small cells. Small cells can be backhauled over a number of different media. In an urban environment, you will typically find DSL lines and you will find PTP or PTMP wireless backhaul solutions, microwave solutions, licensed and unlicensed spectrum. In remote areas, you will find small cells that are backhauled over DSL, over leased lines and also over satellite.

We came from a satellite background, therefore we have very restricted and a very expensive bandwidth, and that is where we learned our trade.

We have taken the same technology and adapted it to address the other types of backhaul – the microwave and DSL that you find in the urban environment. It allows the operator to effectively reduce their backhaul bandwidth requirements by approximately 50%.

You can look at that in a number of ways. You can say, "Well, I can get double the capacity for the same price."Or "I can pay less and cut my cost, reduce my cost for the same capacity, while at the same time reducing the congestion or the amount of time that my backhaul is congested, therefore giving the users a much better experience when they are watching videos or downloading internet content."

Mats Vilander. I would like to add is that congestion due to increase in data traffic is driven by companies such as Google, YouTube, Akamai, ChinaCache and Netflix, just to mention a few. They provide mobile applications over different wireless networks so we have a convergence scenario, where you have 3G, LTE and Wi-Fi all combined in this so-called HetNet environment in urban areas. The most important part is that we are the only one that can supply shared caching.

Frank Rayal. So you are saying that it applies to 3G and LTE, and it is the same solution that will be deployed and will be used by both networks at the same time.

Mats Vilander. We can currently do shared caching or mesh caching 3G and LTE and Wi-Fi, separately.

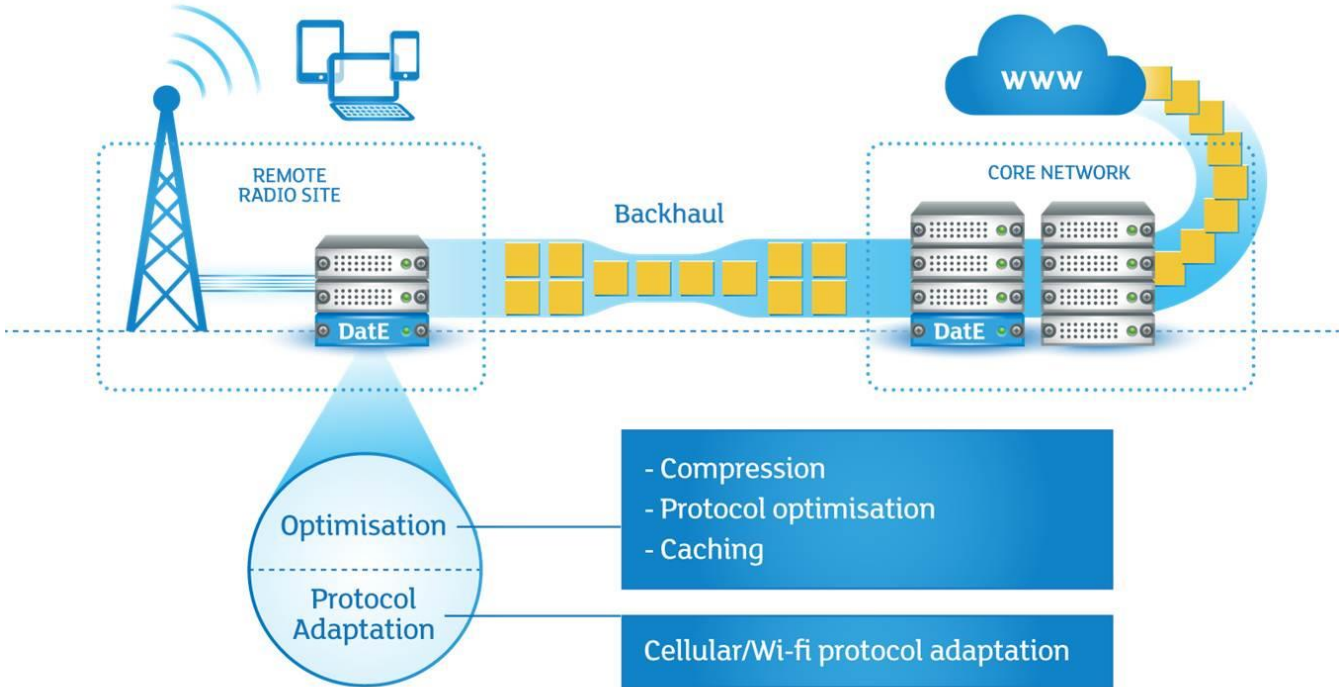

Figure 1 Data-at-the-Edge (DatE) backhaul optimization solution. Source: Altobridge

Frank Rayal. Richard and Mats, I have a question about the advantages of this technology regarding the business case of the operator and some of its features. You have already mentioned some of those, but can you expand on that in the small-cell context?

Mats Vilander. Sure, let's just start by adding on that in the business case for satellite backhaul, typically the current savings are 50% of the backhaul costs – that is one business case which is very specific. By saving 50% of the backhaul cost, from our experience, our solution is paid back within twelve months.

In a small cell urban environment, the cost savings are driven very much by two parameters: one is the cost for the backhaul, i.e., the leased lines, and the second is the entire capacity in traffic. By increasing traffic capacity, the backhaul cost savings will of course be realized.

The other very important aspect, moving in from the satellite transmission market to the urban market, is very much driven, as well, by the quality of experience. For operators, backhaul congestion is a big problem currently. And that is reflected in what was in the recent study of one of the transmission companies globally, announcing that over $9 billion is needed to cope with increased traffic of 3G and LTE.

Many of the operators are in situations where they have to upgrade their network and double its capacity to cope with increased traffic. The business case again lies with byte caching – transparent byte caching – that can save up to 70% in backhaul costs.

What we see in urban scenarios is that byte caching through pre-positioning of content optimizes bandwidth. At the data edge, there is an algorithm that can recognize traffic patterns by end user, so it can store the data of the most popular content close to the user's node. Thus the content will not be pushed back to its core network, where it would, of course, jam up the transmission.

Frank Rayal. I would like to ask you if there is intelligence in terms of distinguishing the different types of traffic.

Richard Lord. Absolutely, yes. We use deep packet inspection effectively to identify the type of traffic and decide, first, whether or not this traffic should be cached. We try and cache as much as we can.

Voice traffic, compressed traffic, or encrypted traffic are the types of traffic that do not lend themselves to caching– and therefore we don't cache it. Caching is most effective with streaming video, for example, or any form of large pictures, like JPEG, GIF or PNG. This is where caching is at its most effective, and so we focus on those protocols and those data types for caching.

Frank Rayal. I would like to also ask you about the competitive environment, in that you have a product that is fairly unique, but are their other ways of approaching this problem? Can you highlight some of the different approaches and how your solution would compare to them?

Richard Lord. I think that, from sort of a competitive point of view, there are two main areas, really. One is what we have always done in the past, is that we have always respected the philosophy that says the more intelligence you push out to the edge of the network, the more efficient it is from a backhaul perspective. I think that we were the first guys to put that sort of optimization intelligence and caching right out at the base station, and we have done that for all types of base stations: classic 3G, small-cell architecture, LTE architecture. We are the only company who has implemented successfully the Iub caching, which is something that is very complex from the technology perspective.

Then the second area, from a competitive point of view, is that rather than implement object cache or URL cache, we implement transparent byte caching, as Mats alluded to earlier, and byte caching has several unique features or aspects. One is that it is more effective as a caching technology than traditional URL caching or object caching.

That is because it works at a lower level: it is able to ignore a lot of the do-not-cache tags and indicators that you find on many web pages with dynamic content. So it gets around that dynamic content aspect by downloading the adapter – only the parts of the web page that have changed since the previous download.

It is also more effective because it is very good at caching partial content. This is a particular problem within applications like YouTube, for example, because in many cases people only download partial YouTube video. For example, if you download the first 30 seconds of a video, and I come along afterwards and download the first minute of that video, then the first 30 seconds comes out of the cache and the second 30-

seconds interval comes across the network – so it is a very effective caching solution.

Because the byte caching is a symmetrical solution that is implemented on both sides of the backhaul and it is only focusing on optimizing that backhaul, it is completely transparent of some of the key features in the core network. Those key features in the core network include billing, legal intercept, content filtering, and service management. Our solution is completely transparent to those key features for the operator, which I think causes a lot of our competitors' headaches, trying to meet requirements such as legal intercept and content filtering.

Frank Rayal. Latency is an important parameter on the backhaul side. How does caching impact backhaul latency?

Richard Lord. Clearly as soon as you start manipulating traffic on the backhaul link, you are going to introduce latency; even if you are doing something as simple as DPI, you introduce latency. Caching for sure introduces some latency.

Now, on the other hand, as soon as your link becomes congested – and this is what is happening all over the world, in small cell and in the macro cells – then, caching helps reduce that latency because you are not waiting for your content. Your content is effectively being delivered directly from the cache at the remote site to the base station, and then onto the user's device. Many parameters can affect the latency of the solution, but it is really a tradeoff. At the end of the day, if your backhaul is congested, then for sure the caching is going to reduce the latency on average, particularly during the busy times. And if your backhaul isn't congested, then you don't really need much caching, so you know the latency introduced isn't really relevant.

Frank Rayal. That is a very good point, Richard. Another question that I have is about the future roadmap, specifically talking about the DatE product (Data-at-the-Edge). Can you elaborate on where this product is now? And what are the unique features and capabilities that you plan to implement, and their timeline?

Richard Lord. Well, I think that the key thing, we have mentioned some numbers during the call: we have seen averages of 50% savings on the backhaul in certain peak times; as Mats said, we have seen up to 70% savings. Most of that is based on caching, and so there are two factors, I guess, which affect the effectiveness of your caching solution.

The first is how good are you at recognizing duplicate content, and we have chosen this byte caching technology, which is very good at recognizing duplicate content.

The second part of the question is: how much commonality is there across the content that users are looking at on a particular cell site? That obviously depends on the community, and then, depends on the size of the community, and the type of the content that they look at.

In order to improve the hit rate on the cache, which is what we are always striving to do as we move forward, we introduce more and more intelligence into the system, to allow the system to learn what the usage patterns are, what the popular data is and what type of content people are looking at, and what time of the day they are looking at it. We use that analytic-based information to pre-position the data and the remote caches.

You can take some very simple examples. You might discover in the morning that some particular news story is popular in New York. Well, you can be pretty sure that three hours later in California that news story is going to be popular. So you can take the initiative to push that content out to the remote caches out in California early in the morning while the system is still relatively lightly loaded.

There are many other examples of that, right down to the point where you can almost have a personal cache that follows a user around the network as he moves from one base station to the other.

There is really no limit to how intelligent you can get in terms of determining what content people are going to be looking at. All the more data that you can pre-position at the caches, the better the hit rate and the more effective the solution is in reducing the bandwidth required and delivering better-quality experience.

That is the main thrust of our development over the next couple of years. We already have a basic implementation of this pre-positioning, and it is something that we will develop over the next couple of years.

We are also fortunate to manage a number of the networks that we actually sell. We don't just sell boxes, but we are in a very good relationship with operators, and we help them manage their networks. This allows us to get to the firsthand feedback of what the problems they are facing are, and what the usage patterns are like. That drives our roadmap going forward.

Frank Rayal. And I take it that this roadmap would apply to both the 3G and the LTE technology, and even to Wi-Fi, is that correct?

Mats Vilander. All three of them.

Richard Lord. Absolutely. There is minimal difference to us between 3G and LTE. Obviously, the data volume in LTE is potentially much larger, but from a technology perspective, it is largely transparent to us. We have implemented a layered solution that can be shared across various access technologies

Frank Rayal. Well, Mats and Richard, is there anything else that you would like to add, on the company or on the solution?

Mats Vilander. I think we all – including Richard and myself – have been in this telecom industry for more than 20 years, in various roles and various RAN suppliers. I think that the biggest challenge right now for mobile operators– because in the end, it is an operator business case – is how to fund the additional investments needed to cope with the increased traffic.

We see an enormous number of new players that come in with entirely new content, but they are all struggling with the fact that many of the 3G networks are congested. And the downtime, if we compare with the fixed broadband services, is the biggest challenge.

Caching by itself, as we foresee it now, will be in every network globally: how do you optimize for more traffic on every node in the coming future, and how do you pre-store popular content close by its end users? Because we are talking about the fixed-mobile convergence that we have been talking about for fifteen years.

This will definitely happen in the 4G world. But it requires the caching functionality.

Frank Rayal. Now, you did mention 50% cost savings and twelve-month payback period. I am curious to understand: how much on average does your product save in terms of backhaul capacity?

Richard Lord. That is a key parameter, because that affects the whole business case to introducing a solution like this. And when we started out this whole sort of program a couple of years ago, we gave ourselves a target of saving 40% – in other words, taking 40% of the bandwidth requirement away from the backhaul – and what we are seeing with real deployments in the field is that we are doing better than that. We are taking 50% of the backhaul bandwidth out, and we are seeing during certain periods, particularly the busier periods, that number go up to 70%, and that is when it is important.

What we are working on is basically technology that will improve the hit rate of the cache and that will allow us to go beyond 70%. In other words, we would be able to take out more than 70% of the bandwidth requirement from the operator's perspective.

Frank Rayal. Richard, thanks a lot. It is very significant to be able to provide 70% or even over that in the future in capacity savings over the backhaul.

Well, Mats and Richard, thank you very much for the conversation today.

Richard Lord. Thank you.

Mats Vilander. Thank you.

Frank Rayal. This conversation is part of a Senza Fili report on small-cell backhaul that provides an overview of small-cell backhaul solutions along with an in-depth conversation like this from leading vendors who opted to participate in the report. The report can be downloaded from the Senza Fili website at **www.senza-fili.com**. Thank you for listening in, and, Richard and Mats, thanks again for this conversation.

Acronyms

3G	Third generation
4G	Fourth generation
3GPP	Third Generation Partnership Project
CTO	Chief technology officer
DatE	Data-at-the-Edge
DPI	Deep packet inspection
DSL	Digital subscriber line
HetNet	Heterogeneous network
Iub	The interface between NodeB (3G base station) and Radio Network Controller
Iuh	The interface between Home NodeB and the Home NodeB gateway
LCI	Local call interception
LTE	Long term evolution
PTMP	Point to multipoint
PTP	Point to point
RAN	Radio access network
S1	LTE interface between an eNodeB (LTE base station) and a mobility management entity or a serving gateway
URL	Uniform resource locator
JPEG	Joint photographic experts group
GIF	Graphics interchange format
PNG	Portable network graphics

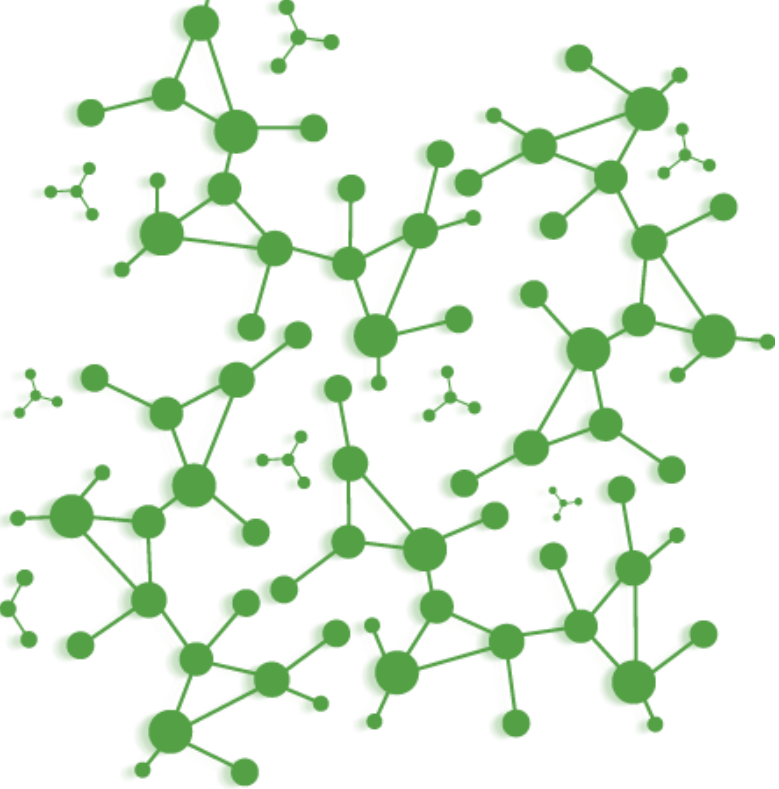

Small-cell backhaul with IQ

A conversation with Mickey Miller, President and CEO, BLiNQ Networks

By Lance Hiley
Senza Fili Consulting

Lance Hiley. Welcome to our audience with **BLiNQ** Networks. This conversation is part of a Senza Fili series on small-cell backhaul that gives an update on the small-cell backhaul solutions that are available today and the evolution of the requirements for small-cell backhaul that are coming directly from mobile operators.

My name is Lance Hiley and I am here today on behalf of Senza Fili. Today we are talking to Mickey Miller, President and CEO of BLiNQ Networks, a company that is pioneering backhaul self-organizing network solutions that offer another approach in how mobile operators can deliver mobile broadband services to urban areas.

With that I would like to introduce Mickey. Perhaps you could give us a high-level introduction to your company and what your product does.

Mickey Miller. Thank you, Lance, it is great to be here. BLiNQ Networks was founded based on assets acquired from Nortel Networks. Nortel had looked at this problem in the 2007–2009 timeframe and recognized that the biggest issues around scalability to small-cell networks would be backhaul. They looked at all of the various technologies they had available at the time and we have built on this foundation.

Our sole focus is solving the small-cell backhaul problem. The name of our company, BLiNQ, stands for "backhaul link with intelligence, IQ," and BLiNQ is completely focused on enabling carriers to deploy dense, scalable small-cell backhaul networks at half the cost and in half the time of traditional backhaul methods, while providing very cost-effective scalability for future growth.

We first started talking with carriers prior to forming the company, and a lot of those are very forward-looking in how they approach the small-cell problem. It was clear that they wanted to have a solution that was true non-line of sight so you didn't have multiple hops, but also had the reliability that the carriers expect, as well as the predictability.

When carriers deploy a link, they want to guarantee a consistent throughput and latency for the life of the link. We wanted to make it simple to plan, because the traditional approach that we use to deploy macros does not scale well for small cells. Given the amount of access points that would be deployed – we are talking about the need to be able to deploy an access point and a backhaul link in less than 30 minutes. Many times small cells are being deployed in high footfall areas like Times Square, you don't have the luxury to shut down traffic.

It also has to be scalable. We looked at Qualcomm's 1,000-times scaling forecast, and at what it would take to deploy 1,000 times the amount of access points that we have today. We need a backhaul solution that is able to scale with that in a very cost-effective manner. The backhaul, of course, has to have the

capacity for multimode 4G, 3G and Wi-Fi, so it has to have very high capacity and has to be delivered for a very low total cost of ownership. That is where we have been focused, Lance.

We have an intense focus, and we have developed a great deal of intellectual property around this and recently introduced our first devices. As we see it, it is a solution, it is not just a simple radio—it is a complete solution that solves the small-cell backhaul problem.

Lance Hiley. I heard you talking about the deployment of networks in dense urban areas, and clearly trying to get more capacity into mobile broadband networks that are deployed in hotspot areas and dense urban environments is a main focal point for small-cell networks. But can BLiNQ help with the notspots, as well as the rural networks and the suburban edge networks that operators are looking at using small cells to solve problems in, as well?

Mickey Miller. Absolutely, the notspot is a typical use case. Our solution works in dense urban, in suburban and even in rural areas. And we see carriers wanting to address both the hotspots, where there is overcapacity, as well as the notspots where they cannot get capacity.

Lance Hiley. That's good. And I also think I heard you talking about it being an RF solution, and I guess if it is RF, what bands are you using for your product? What is the spectrum availability like and what does it typically cost?

Mickey Miller. We have a very spectrally agile solution. It is agile and efficient, and agility is key,

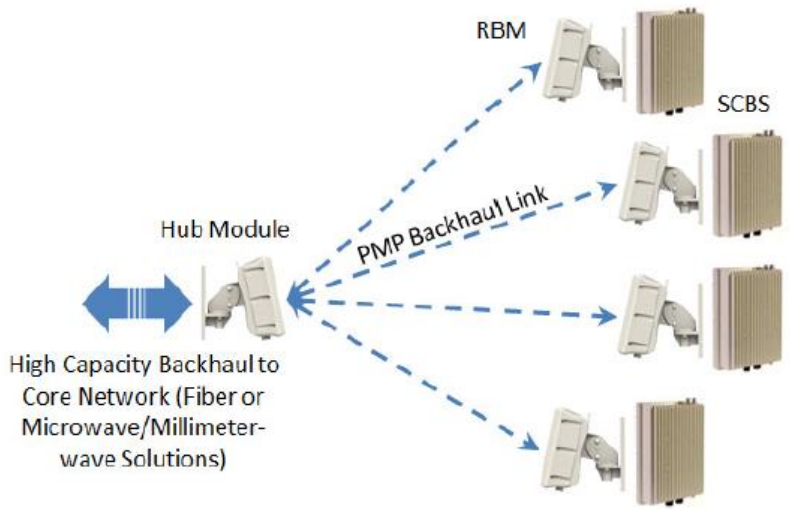

Figure 2. BLiNQ Point-to-multipoint NLOS wireless backhaul solution. Source: BLiNQ

because depending on where you are in the world, the availability of certain bands isn't consistent. We have both a licensed and an unlicensed – a dual-band – solution.

On the licensed sub-6 GHz side, we focus on 2.3, 2.6, 3.5 and 3.65 GHz, and on the unlicensed it is of course 5.8 GHz. As a system, we look to manage the traffic, and pair the type of traffic with the best available resource. We can put guaranteed-bit-rate traffic over licensed, and best-effort traffic over unlicensed, but it is all about identifying the best radio resource and actively managing traffic through the best radio link to be able to provide the end user experience that customers demand.

Lance Hiley. I think it is fair to say, though, that BLiNQ is not the only non–line of sight wireless backhaul proposition in the marketplace right now. What would you say would be the one or two key features that you have or advantages that you have that really distinguish your product from the others, to a mobile operator trying to make a decision right now?

Mickey Miller. I look at that question in two points. There are the traditional non-line-of-sight folks that have been deploying non-line-of-sight PTP or PMP links for some time. When I view those, that is a radio solution, and it doesn't necessarily lend itself to the self-organizing network that is required in small cells to be able to scale to the level these carriers want it scaled to. Also traditional vendors were not geared to provide dense high capacity point to multipoint networks in interference challenged deployments.

Then we have the new entrants like BLiNQ, which entered the market a year or a year and a half after we entered the market. We have not seen their products

yet, so I can't comment on what they do differently, but it is critical that people recognize that for a small-cell backhaul solution, the product has to be designed and developed with small cells in mind. The traditional ways of doing things just don't scale in the small-cell network world.

Lance Hiley. That is interesting, because I think one of the other competitive advantages that has been mentioned by BLiNQ in the past is the ability to integrate your solution into the small cell itself, and this has got to be of interest, because the operators are saying that they are really restricted as to how much equipment they could put up on different pieces of street furniture and other targets that they are thinking of using for the deployment of small cells.

Is that still a strategy or tactic that you are following?

Mickey Miller. Absolutely. We look at it this way: We are a software company and we deliver on a hardware platform. The hardware platform can be installed by itself, or it could be integrated in a mechanical package or at the base-band level – there is a wide variety of different ways in which it can be integrated.

Full integration will take some time, just because today there are massive challenges around multimode and multiband of what the current access points have, not to mention what the backhaul spectrum will be. But once those things are further defined, I think that there's a great opportunity to integrate into one single package.

But for us the hardware is just an element that is managed by the self-organizing functionality of the network, and this is where we focus our efforts. So whether the hardware is our product or a competitor's product, it could be seen either as an integrated access-backhaul point or a separate equipment part.

Lance Hiley. Another hot topic of discussion in outdoor deployed small cells is the topic of interference, and I believe that BLiNQ has something special in that area as well.

Mickey Miller. Absolutely: what we term our Managed Adaptive Resource Allocation. We can continuously manage and characterize the RF network. Based on that, we can understand what the foreign interferers are in unlicensed bands and we can also understand the interferers from our own network, and we are able to manage both. The time, space and frequency are all elements in which we can use to manage interference.

Our purpose was to create a powerful software tool set for interference management and one platform to successfully enable it. We have put together a variety of tools to cover multiple use cases starting from network-level power management and joint scheduling, to sophisticated space domain interference avoidance techniques.

By managing interference, we can optimize both the performance at a link level and at a network level. The idea is to provide optimal throughput at the network level – for instance through power management. We have taken many things that have been used in the access world for some time, and applied them to backhaul at both network and device level.

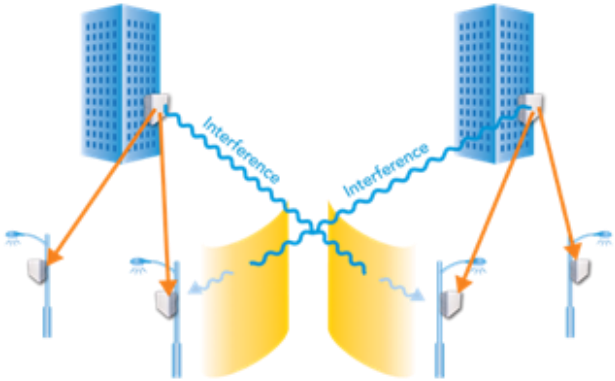

Figure 2. BLiNQ's MARA solution helps mitigate dense deployment interference issues. Source: BLiNQ

One simple example at device level is dynamic blanking where you turn the power down if you do not have traffic per frame-by-frame basis. Those are all capabilities that are baked into our solution, and we have innovations on both the hardware side and the software side to be able to provide an optimal format to deliver that.

Lance Hiley. That is really interesting, and we will come back to some of the cost ramifications that you just mentioned in a moment, but before I leave the area of hardware – what typically are your installation parameters, and how long does it take to install the product that you have visualized and you are bringing to market right now?

Mickey Miller. Less than 20 minutes; and that is a requirement. We expect to improve on that, because on average the goal for carriers is less than 30 minutes

for the access point and the backhaul point. We have been doing it in 20 minutes in many of the deployments that we have undertaken so far.

Lance Hiley. If you are installing in less than 20 minutes, then there can't be too long of a configuration time associated with that, so that literally you stick it up on a piece of street furniture or on the side of a building, and you switch it on and there is not much more involved to that. Or is there a lot going on in the background?

Mickey Miller. No, it is all baked into the software. We have software-defined radios that allow you to configure the link, depending on what bands and channels you want to operate on. It recognizes where it is and what hub to connect to, and then you are up and running, Plug and play.

Lance Hiley. So clearly the installation is going to be one of the main contributors to an improved cost model that your solution might offer to an operator. What are some of the other benefits that reduce the cost of deploying a network using a BLiNQ Network solution?

Mickey Miller. First, our solution is easy to plan and secondly is quick to install. Probably the most important is once you install this – and what we see is carriers will not go and deploy a thousand of these in a city – we will address their hotspots or their notspots.

Another important feature is that our devices can adapt to various deployment use cases to maximize their performance in any location where the link budget, capacity or interference present challenges. This can change anytime as network scales or customer needs it.

Always challenges come as network scales, and many spots can become interference challenged – and we can fix this on the run.

As the traffic patterns change and additional hotspots come on board, our solution allows operators to understand what the new RF network looks like, and our algorithms will do what-ifs. For instance, "If you put one access point here, here is how the network reconfigures and here are the types of throughputs that you can expect to get out of that point."

It becomes a real-time planning capability that the carriers can use as they deploy and scale their network.

Scalability is the biggest challenge here. We can talk today about a link here and a link there, but looking at the numbers that the carriers need to get to be able to give the user experience to their customers, we need significant densification. To be able to do that, the traditional RF planning tools that we have used with 2G and 3G just aren't going to allow them to cost effectively get to where they need to get to.

Lance Hiley. Okay, so with that in mind, then what should we expect to see from BLiNQ Networks this year and going into the future, in terms of product development or in the way that your business actually goes to market?

Mickey Miller. I think that this year is a critical year for us. We are introducing our X100 product, which is our 10 MHz, 82 mbps solution. We are also introducing our complete overall backhaul SON capability at CTIA. From there we will be introducing our future elements, which will be single band, multiband, higher bands, 20 MHz channels, multiband up to 40 MHz, and bring in LTE capabilities to the small-cell backhaul world.

It will be a continuum as we go from 4G to 5G. The requirements to densify the network will be there. Meeting the requirements of a dynamic, real-time backhaul network will be critical, and we see our self-organizing capability as having a key role in that.

We are bringing a solution that will allow operators to bring up to 500 mbps to one small-cell site, so hardware platforms mounted on the pole will stay there for a long time without additional maintenance or upgrade costs. All the rest is software upgrades or hub updates.

Lance Hiley. Well, thanks for that, Mickey, it sounds as though we have a lot to look forward to from BLiNQ Networks, even to the extent that you are thinking about 5G already, and certainly your comments and observations will be a great contribution to this report on small-cell backhaul.

Just to remind everyone: this interview is part of a report that provides an overview of the differences among small-cell backhaul solutions that are available. The report can be downloaded from the Senza Fili website, www.senzafiliconsulting.com.

And with that I would like to say thank you very much again to Mickey Miller, president and CEO of BLiNQ

Networks, and I hope that everyone has enjoyed the conversation that we have had today.

Acronyms

2G	Second generation
3G	Third generation
4G	Fourth generation
5G	Fifth generation
IP	Intellectual property
IQ	Intelligence quotient
LTE	Long term evolution
OEM	Original equipment manufacturer
PMP	Point to multipoint
PTP	Point to point
RAN	Radio access network
RBM	Remote backhaul module
RF	Radio frequency
SCBS	Small-cell base station
SON	Self-organizing network

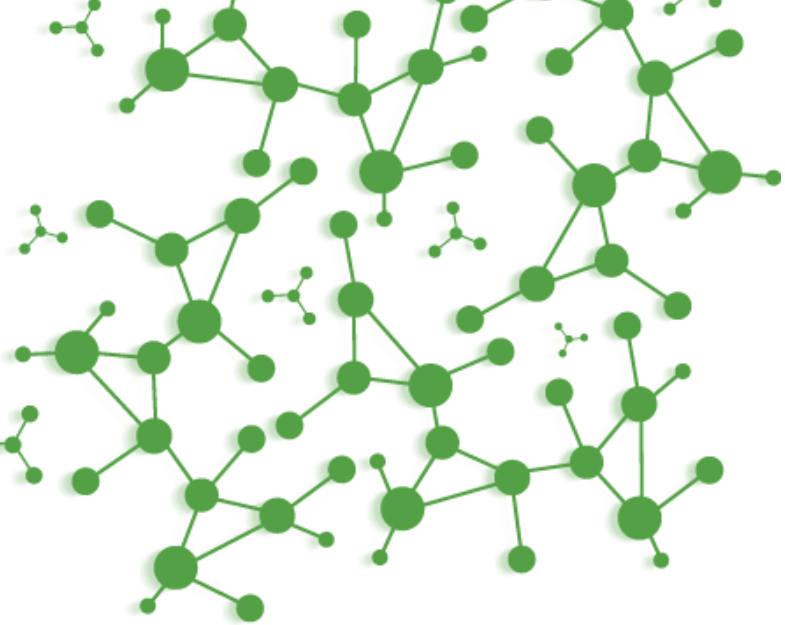

Bluwan's multi-gigabit point-to-multipoint HetNet backhaul technology

A conversation with Shayan Sanyal, CCO, Bluwan

By Lance Hiley
Senza Fili Consulting

Lance Hiley. Welcome to our audience with **Bluwan**. This conversation is part of a Senza Fili report looking at the different small-cell backhaul solutions and the evolution of mobile operators' requirements.

My name is Lance Hiley, I'm here today on behalf of Senza Fili, and I'm joined by Shayan Sanyal, Chief Commercial Officer of Bluwan, a company that is taking point-to-multipoint solutions to new spectrum bands and new levels of performance. Many operators are looking at point-to-multipoint wireless backhaul solutions as part of their small-cell backhaul strategy, so it is great to have Bluwan with us here today.

Shayan, could you give us an overview of your company and the products that you are targeting the small-cell backhaul market with?

Shayan Sanyal. Bluwan was spun off from Thales Communications, a French aerospace and defense company, in 2005. From there on, we developed a disruptive technology called Fibre Through The Air, which is really specifically targeted for the telecom industry.

Fibre Through The Air technology leverages four basic components. Our first one is a multi-gigabit wireless point-to-multipoint architecture, which enables operators to efficiently allocate bandwidth and reduce costs, while providing very significant network coverage and flexibility.

Second is the use of very wideband frequencies in the 42 GHz spectrum band that enables ultrahigh-bandwidth applications.

Third is a unique multiplexing technology, which allows operators to carry different data flows through multiple channels within our radios.

And last, a self-optimizing network that dynamically optimizes all of what I have just spoken about based on real-time network performance requirements and UE demand.

Specifically, for small-cell and multilayered heterogeneous network backhaul, we have developed a solution called LinkFusion. It's a millimeter-wave, point-to-multipoint wireless backhaul solution, specifically targeted for small-cell and HetNet backhaul, and delivering ultrahigh capacity over the last mile. One of the key differentiators is our ability to provide scalable capacity up to 10 gbps from a single transmission hub and to allow operators to easily address what we call the bandwidth density requirement, delivering peak point-to-point link performance to endpoints using a point-to-multipoint architecture.

Lance Hiley. It sounds like you have gone a long way to address a lot of the different topics and requirements that operators are saying are important to them when they are selecting small-cell backhaul solutions. Would you say that there is a specific positioning for your product? Is it a notspot solution or is it a hotspot solution? Are you targeting real metro line-of-sight–

challenged areas or non-line-of-sight–challenged areas?

Shayan Sanyal. We are focused on where the bandwidth demand is increasing exponentially, and where we are seeing that happen most of all is in dense urban metropolitan areas. However, the focus isn't just on any metropolitan area or any service provider operating in that environment, but operators that are specifically looking at expanding their network capacity through the rollout of either 3G small cells or 4G metro, or mobile network operators that are starting to de-saturate their existing networks by using carrier Wi-Fi–type solutions and potentially even rolling out 4G macro sites with peak requirements above 200 mbps.

What we are focusing on are those kinds of heterogeneous environments where operators are rolling out a diverse set of access technologies and where the forecasted capacity requirement in those areas is likely to be over a gigabit per square kilometer in the next 12 to 24 months.

Lance Hiley. It really sounds like you are trying to address the performance levels that a lot of operators are concerned about in terms of being able to deliver to not just their small cells but also to their macro-cell network as well. If you have that level of performance, you must be using a lot of radio spectrum. You mentioned the 42 GHz band earlier. Could you tell us a little bit about the spectrum that you're using and what it costs?

Shayan Sanyal. As you correctly pointed out, we are extremely focused on a particular spectrum band that is part of the Q band. A lot of operators will have heard of the V-band or E-band spectrum. The 42 GHz spectrum that we are using in the Q band is ETSI- and CEPT-harmonized and is located between 40.5 and 43.5 GHz. So in effect, it provides about 3 GHz of spectrum that has been, by and large, left untapped and unused. So we can actually call it virgin spectrum.

The advantage of having virgin spectrum at your disposal is that you can start looking at a contiguous-block allocation of that spectrum. And when you have contiguous-block allocation of spectrum, you are actually able to generate significantly more capacity using advanced radio equipment than if you were operating in a single channel, say a 56 MHz channel or a 28 MHz channel.

So as a result of that, when we saw this spectrum and identified its properties, compared it to other spectrums that are higher in terms of the 60, 70 and 80 GHz spectrum bands, or to other point-to-multipoint spectrum at 10.5, 28, or more recently, the 32 GHz bands, we saw that this was an ideal opportunity to build a multi-gigabit radio system that could leverage those contiguous spectrum blocks to generate multi-gigabit sectors that can in turn address some of the challenges that operators were facing with traditional point-to-multipoint microwave. Think of it as the best of both worlds: the capacity of millimeter wave with the efficiency of point-to-multipoint.

Lance Hiley. If there is 3 GHz of virgin spectrum available out there, how are you leveraging that with a technology advantage? What makes your solution better than the others in the line-of-sight, multipoint backhaul category?

Shayan Sanyal. The first thing would be the 1GHz ultra-wideband chipset that we have developed. We have designed and developed the MMICs, the power amplifiers and the synthesizers; they are protected and patented by us. This technology allows us to leverage 1 GHz contiguous spectrum blocks, but recognizing that not every operator may be in that position, we can support smaller 250, 500, 750 MHz blocks as well. The chipset is completely flexible in terms of what range of frequency band the operator can use.

The second thing that we did was develop an advanced multiplexing technology that allows us to take multiple 40 MHz channels, bond these channels together, and then feed the multiplex to our radio.

The third thing that we did was to develop a software and network optimization layer that allows operators to dynamically allocate traffic and throughput to endpoints based upon need. This is different from statistical multiplexing that analyzes peak versus mean. Our solution is more about looking at the type of demand that's at an endpoint, whether the download requirements are at 80% versus 20% for the upload. In certain cases, that's actually reversed: the upstream requirement is actually higher than the downstream requirement. This is part of the reason why we developed a low latency TDD-based system. It is because we wanted to efficiently use that spectrum at the endpoint and allow operators to adjust the bandwidth they were providing and the shape of that bandwidth based on UE demand.

Beyond that, because of the significant amount of capacity that we can provide to an endpoint, we can actually enable applications such as backhaul infrastructure sharing. We can deliver up to 240 mbps to an endpoint, and a single operator may not consume all of that. We know that there is a trend where operators may have to collaborate and potentially share backhaul infrastructure when putting in their actual small-cell units. So our system is able to essentially provide a multi-dimensional traffic profile to each of these operators. So say you have two operators sharing a lamppost. Operator One can have a downstream of 100 mbps and an upstream of 100 mbps, so a very symmetrical pattern. But another operator sharing that backhaul could have another. Each of them have their own QoS profiles, each of them can manage their policies individually. So they have complete freedom of movement in terms of customizing and tailoring their offer to their users, but the infrastructure is actually shared. And that can help a number of use cases, like providing managed service backhaul for small cells, something that MLL Telecom in the UK is trying to do.

Lance Hiley. So it sounds as though the product is not only adaptive and dynamic in terms of how it allows the operator to manage traffic, but it is also very flexible in how it can be configured for specific deployments. When it comes to deployment scenarios, what is this product like to install? How long does it take? What is involved in terms of skill sets, and are there any other installation parameters or aids to make it easier to install?

Shayan Sanyal. I think we are going to break this down into a couple of components. One is going to be physical. The other one is going to be more OAM related.

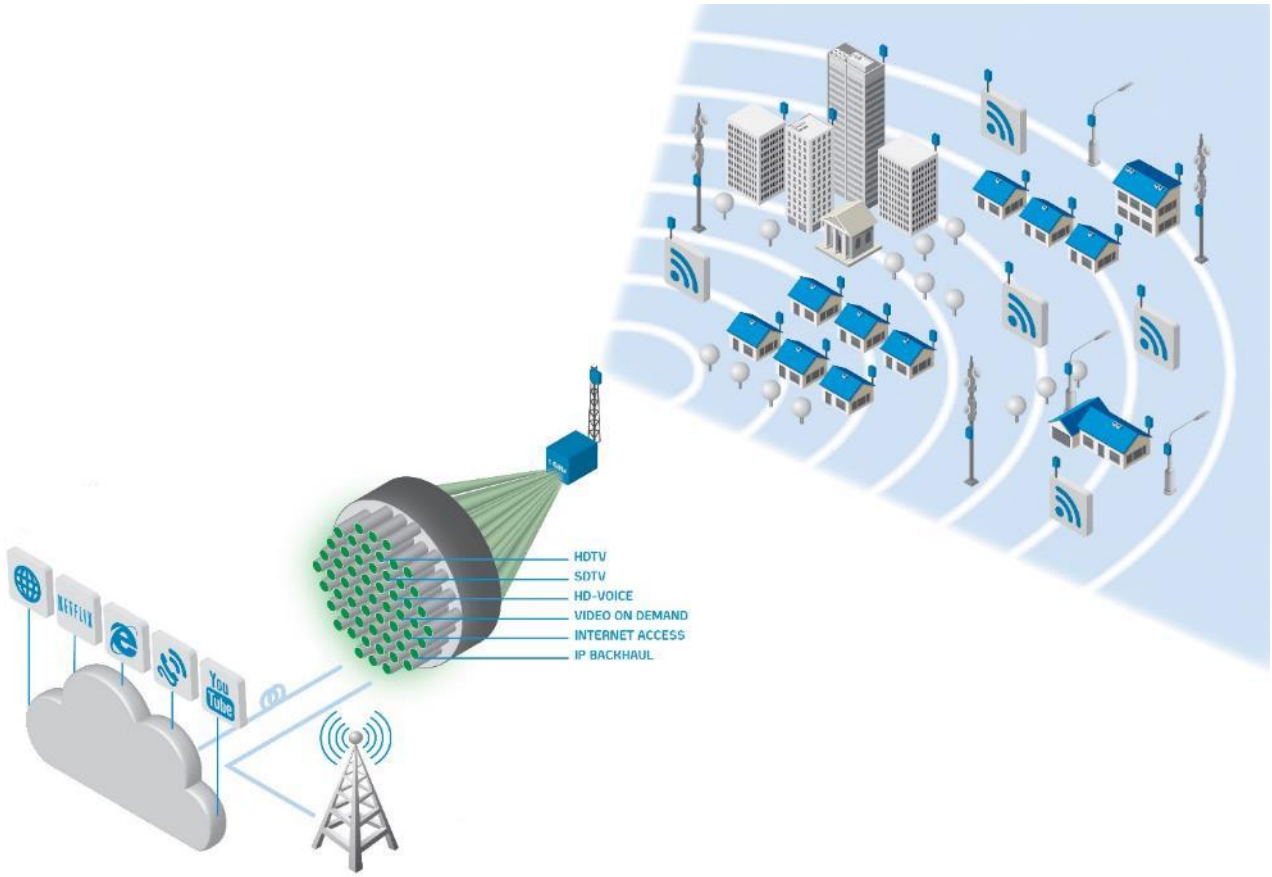

Figure 1. Bluwan's wireless wave-division multiplexer is one of the three components that enables scalable capacity in the LinkFusion system. Source: Bluwan

From the physical perspective, it is a point-to-multipoint system. So you are talking about an access point where you have a 90-degree sector that typically ranges around 2 km. That sector then provides up to 2.4 gbps of capacity. Once that is set up, and that typically takes a bit less than a day to do, and it is up and running, you can fine-tune your coverage.

So where do I actually get line of sight to the places I want to provide bandwidth to? Now, we are cognizant of the fact that line of sight can be a challenge in a dense urban environment. MNOs are faced with a high degree of clutter. There is also a lot of street

furniture on the ground. So having taken that into account, we also developed a relay capability. As you deploy terminals, some of these can be used as relays to light up shadowed areas. So essentially you have a 90-degree sector, but then within that sector, you may have relay nodes that are lighting up areas that don't have direct line-of-sight from the operator's aggregation node. The other thing we have developed in our radios is the integrated antenna that enables our very small form factor. Our access-point is a split-mount configuration, where the 1U indoor unit is actually inside the data center or the cabinet, and the outdoor radio unit is 3 kg, the size of a shoebox.

Once we have that in place, the actual network-terminating equipment, or the terminal equipment, is available in multiple configurations, depending on the usage the client needs to make in his network. For example, if you are providing capacity to a relay site, you might want to deliver 240 megabits of peak throughput to that site and then drop off maybe 120 mbps on that site and relay another 120 mbps down a street canyon. That street canyon would have 120 mbps of backhaul capacity that could be shared between four to five sites, using what we call our Metro NTE.

Our low cost metro NTE is our 15 x 15 cm small–form-factor PoE-powered device. It easily installs onto street furniture or the side of a building, and then, because of the 6-degree beamwidth of the NTE antenna, you don't need to spend a lot of time aligning the equipment.

It's not like the millimeter-wave point-to-point vendors that tout a 1-degree, pencil-beam width type of advantage. We actually view that as a disadvantage because it takes time to align. It requires significant resource to ensure that your street furniture is robust enough to ensure your equipment doesn't waiver, thereby degrading performance. We built in that flexibility. So that's the physical properties of the solution that make it easier to install. Typically an NTE takes less than three-quarters of an hour, more likely 20 to 30 minutes, to install. But maybe the first one takes a bit longer and then after that, you understand where your hub is and the alignment that comes along with it.

Our product requires a lot less skill to install as a result. We are working on software OAM upgrade features that allow automatic stress testing and auto-configuration. That means once you have the alignment done, the network controller provides all the configuration settings, so the technician can actually move on to the next lamppost, and then gets notified as and when things progress.

Lance Hiley. What other areas does your product contribute to for helping tight cost budgets that mobile operators have when they are planning their small-cell networks?

Shayan Sanyal. Questions about the total cost of ownership are among of the most important questions service providers ask. When we were speaking at Mobile World Congress earlier this year at the small-cell session, we actually had a debate on whether we should be talking about total cost of ownership or maximizing network value. Operators that are rolling out small cells are not reducing their TCO.

Figure 2. Bluwan's LinkFusion Metro NTE is a 15 x 15 cm high-capacity terminal for small cells. Source: Bluwan

What you are trying to do as an operator is maximize, as much as possible, the efficiency of that investment and maximize the value that you are creating by rolling out and investing in these types of technologies. That does not mean that you should not be looking at cost, but it is a very value-driven total five-year business case and making sure that that works. If that does not stack up, focusing just on capex alone isn't going to solve the equation. Capex is a very important point, but opex is one of the Achilles' heels of the successful execution of small-cell network deployments.

Of course, from a TCO vantage point, there are the traditional point-to-multipoint advantages. You do not have nearly as much equipment in point-to-multipoint

as you have in point-to-point. That has a ripple effect in your business case, because you end up paying less on site rental costs, especially for aggregation nodes or macro sites, where you need to provide the backhaul from. And typically these can be quite expensive, and they can go up to £1,500 to £2,000 a year for a 30 cm dish in the UK, $800 per dish or per antenna per month in Egypt. It can be very expensive.

The second thing is around the requirements to power all of this equipment. When you have half the equipment, you have a lower power overhead.

Finally, we have deployment, installation, and commissioning of this equipment. Point-to-point is always going to require a dual-ended installation. That means two teams lining up the link, and the planning requirement that goes into developing a point-to-point small-cell network is often not talked about. The fact that you can daisy chain is commendable, but the fact that you actually have to daisy chain and mesh your network, from a planning overhead perspective, requires so much more effort for an operator, and that adds costs.

Operators are not trying to scale up the number of network architects they have on staff at the present time. They are trying to outsource as much of that as possible. And, by deploying point-to-point architectures, they are creating a major planning overhead in terms of effectively having a point-to-point tree-branch or daisy-chain network. PMP is absolutely the opposite. You set up your sector and it is a fairly simple, hierarchical architecture. With a single-ended installation process, you end up with less people on the ground, and less overhead.

There is a second angle when it comes specifically to Bluwan LinkFusion: the spectrum allocation it uses. For example, in the UK, operators have had a very, very high cost advantage – meaning a low-cost investment in wide-area licensing. For example, UK Broadband in the UK was able to get 2 GHz of spectrum for £120,000 for a 15-year nation-wide non-revocable license. This cost advantage makes it so cheap that the spectrum has no capex impact on their business case whatsoever. You know that you are going to amortize that within a few months of deploying a solution.

The final thing around the total cost of ownership is the business case. We have done a number of studies with our trial customers, and we were able to demonstrate that from a cost-per-megabit perspective we can provide a 60% savings compared to the existing solutions that they have in their network today.

For an operator in the UAE, that can amount to nearly $100 million worth of savings over a four- or five-year business case across their high-density urban environments. That is a significant cost saving and one of the key things that we are focusing on.

Lance Hiley. That is great, Shayan, and thank you very much. Your insight and observations are going to be a great contribution to this report on small-cell bbackhaul. The report aims to provide an overview of small-cell backhaul solutions available, and like this interview, there will be in-depth interviews with all the leading vendors who have participated in the report. The report can be downloaded from the Senza Fili website, www.senzafiliconsulting.com.

Acronyms

3G	Third generation
4G	Fourth generation
ETSI	European Telecommunications Standards Institute
FCC	U.S. Federal Communications Commission
MMIC	Monolithic microwave integrated circuit
NMS	Network management system
NTE	Network terminal equipment
OAM	Operations administration maintenance
PMP	Point to multipoint
PoE	Power over ethernet
TCO	Total cost of ownership
TDD	Time division duplex
UE	User equipment

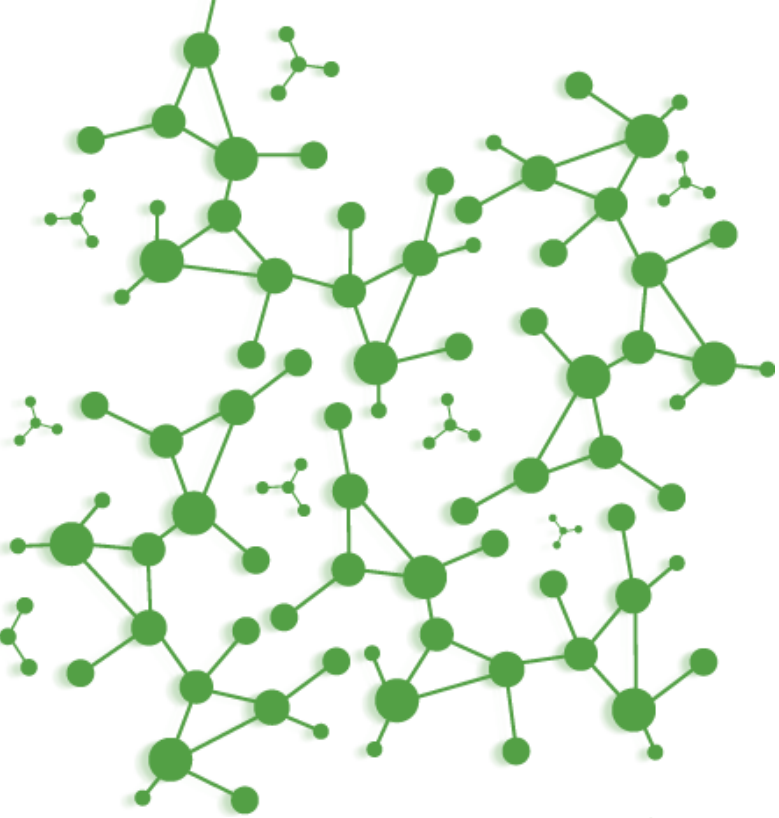

Smooth planning for small-cell backhaul growth

A conversation with Steve Greaves, CEO, CCS

By Lance Hiley
Senza Fili Consulting

Lance Hiley. Welcome to our audience with **CCS**. This conversation is part of a Senza Fili report on small-cell backhaul that gives an update on the solutions available today and on the evolution of mobile operators' requirements for small-cell backhaul.

My name is Lance Hiley, and I am here on behalf of Senza Fili. Today we are talking to Steve Greaves, Co-founder and CEO of CCS, founded by the same team that started Adaptive Broadband Corporation and Cambridge Broadband Networks Limited. CCS is looking to bring to market a small-form-factor backhaul solution that is self-organizing, addressing many of the challenges currently being discussed in small-cell backhaul. So with that I would like to introduce to you to Steve Greaves. Steve, could you perhaps start off with just explaining the background of your company and what CCS's product is all about?

Steve Greaves. Sure. Our other co-founder is John Porter, and we established this company in 2010. One of the first things that we did was go and talk to mobile operators about the problems that they foresaw in the future. One of the key issues that came up was outdoor small cells and the lack of suitable small-cell backhaul solutions. Then we went away and looked at the requirements.

The first requirement is link performance. We provide 450 mbps capacity and 125 µs latency for a single unit, 900 mbps for a dual unit.

The second requirement is form factor. It was very clear from the start that these small-cell backhaul units will be deployed at street level on street furniture and they have to be small. Form factor is very, very important.

The third requirement was a short installation time. The operational cost to an operator when considering conventional backhaul technologies to be used as small-cell backhaul was quite frightening for them. The operational cost of radio planning, installing, and going back to modify systems that didn't work was just too expensive. So we hit on the notion of a small-cell backhaul system that was self-organizing. In fact we have shown that an electrician can deploy our system in under half an hour.

In meeting these key requirements we have developed a carrier-class microwave backhaul system that can be deployed in a fraction of the time of conventional technologies – and by conventional I mean the PTP and the PMP technologies, and even NLOS systems, including sub-6 GHz ones.

Lance Hiley. You mentioned street furniture – are you specifically designing a solution for urban metro environments, or is that something that can be used across a wider range of small-cell deployments?

Steve Greaves. It is specifically focused on a dense urban deployment, so metro range would be correct. We made the decision not to try and do the WiMAX 75 mbps at 75 miles. We asked: what is a reasonable distribution of small cells? In answering this we took a lot of input from the mobile operators and designed our system to meet their requirements. Our system provides a standard range of about a kilometer, because we predominately see deployments in the range of 300 to 400 m for the many-point-to-many point (MPTMP) mode. However this range can be increased with the addition of a higher-gain antenna system.

The system architecture is generically multi-point. This means it can run in both PMP and MPTMP modes. Conceptually MPTMP looks like a PMP base station in communication with other PMP base stations. In MPTMP, we build multiple connections between base stations. To run the system as a PMP system, we can attach standard dish antennas to the units and get much longer range – in that case, standard PMP ranges.

When run in PMP he key advantage of our system over existing PMP is that you still maintain the self-organization with zero-frequency planning combined with up to 900 mbps and 125 µs link latency performance. With our system we give a single frequency band to the system and the system decides how to use that in a way that optimizes data transport performance in a dense urban and interference-limited environment.

Lance Hiley. I gathered from the way that you are describing the system is that it is a wireless system, but what spectrum do you use and why did you choose that spectrum, and what are the costs associated with that spectrum?

Steve Greaves. We are focused on what is termed the area-based microwave frequency bands. These bands are available worldwide, and they cover 26, 28, 32 and 42 GHz. Currently there is around 6.5 GHz of underutilized spectrum in these bands waiting for an application. These bands have historically been known as the PMP bands. For instance, the 28 GHz band was the very early LMDS band.

The advantage of those types of frequency bands is that they are geographically arranged. In a conventional PTP system, you have to go to the regulator and you will have to apply for a license to cover from point A to point B. This is clearly impractical in the small-cell deployment scenario.

The concept of having an area-based, geographic-based license is much better. We chose 28 GHz initially because that is the band that has the best component availability and projected cost reductions, meaning that we can hit our target price in volumes. So initially we started at 28 GHz and then we will look at 26, 32 and finally 42 GHz.

Lance Hiley. You have mentioned PTP and PMP, and all of these technologies are what we could loosely describe as falling into the LOS microwave market segment. Is your solution specifically LOS? And if it is, what advantages do you feel your approach has over the other solutions that are out there?

Figure 1: CCS Multipoint 28 GHz self-organizing small-cell backhaul node. Source: CCS

Steve Greaves. Our system operates predominantly as a LOS microwave system. It may also be possible to make use of passive reflectors to assist in routing and we are actively investigating this. Initial work has shown our system can detect reflected signals and

adjust network topology to suit this. That being said we expect that our nodes have a good connection to at least one other node, or if not a good reflector. As our nodes will be co-located with the small cells we expect the connectivity distances between 50 meters and anywhere up to 300 or 400 meters, with typical deployments at street level. However, even if there is no direct connection available between nodes the system can route along many different paths and build up routes around obstacles, creating a mesh connectivity architecture. This means that we can provide NLOS capability, by routing around obstacles if the potential path is there.

So it is very simple, because it takes away the process of large scale radio planning that is so difficult with these types of systems. We provide very high capacity at low latency akin to a line of LOS system, with the deployment simplicity of NLOS systems.

Lance Hiley. I would like to dig into the self-organizing versus network planning aspect of your solution a little bit. One of the benefits that the NLOS small-cell backhaul community are really pushing with their systems is that you can use fairly unskilled people to install it, switch it on, and it more or less aligns itself, and that is because in most cases you just have some kind of an omnidirectional antenna system on it.

Where I guess your system differs is that you actually have to plan where the nodes need to be located so that there is a reasonable opportunity or chance that they will be able to have a LOS to another node. Is that a correct way of describing it?

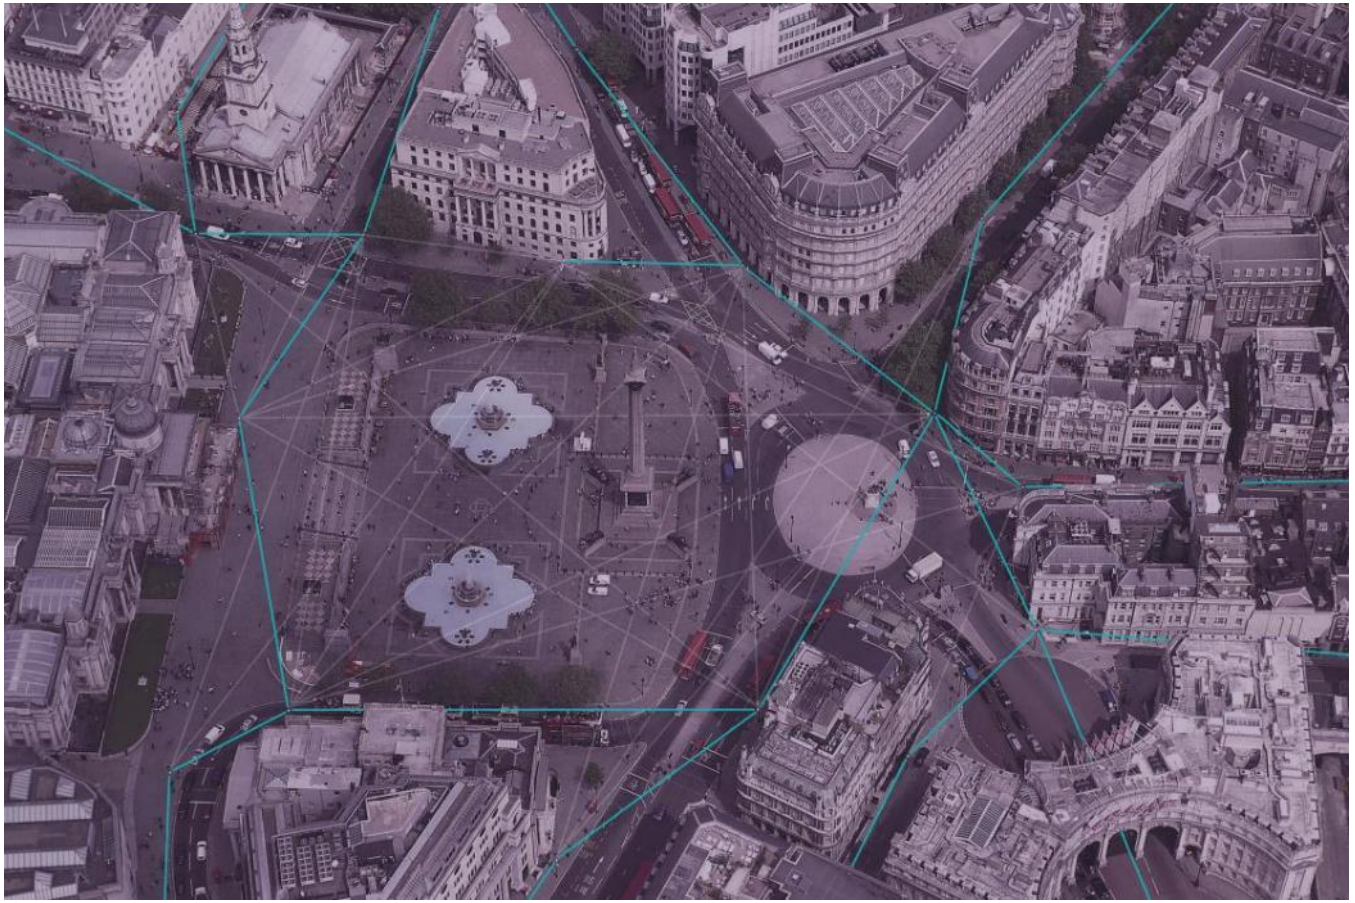

Figure 2: CCS multipoint nodes automatically seek out the best adjacent node to form MPTMP backhaul architectures that route around obstacles. Source: CCS

Steve Greaves. Ideally we would like to have LOS to another node, but stepping back to your analogy with the NLOS, we looked at sub-6 GHz systems and liked just turning up and switching the thing on. In reality, however, the capability and performance of a sub-6-GHz is ultimately determined by the degree LOS availability of the link, and interference.

With a practical power budget you can't have both things: a NLOS system and a guarantee of carrier-class availability, capacity and latency, as the variability of the NLOS channel and interference is so high. What we have seen in many NLOS deployments is that they are still running LOS. So you look at how those deployments have been configured, and they frequently have visibility of other nodes in the

network. So for NLOS systems to work at full throughput and the minimum latency that they are stating, they really need to have LOS, and somehow suppress self and external interference.

So the installation requirements of a sub-6 GHz system and our system turn out to be similar. As long as we can connect to another lamppost – and that lamppost may be 50, 100, 200 or 400 m away, or up to a km away – then we just turn that switch on and the system requires no aligning. You just put the unit up in the same way that you do with a sub-6 GHz system, you just switch it on. There is no difficult alignment. It finds any possible link, it determines if it is a good link or if it is a bad link and it effectively picks the optimal link to make the connection. However our system delivers a carrier class 450 mbps capacity and 125 µs latency for a single unit, 900 mbps for a dual unit.

The key to self-organization is network intelligence. All the high-capacity PTP and PMP systems out there are pretty dumb. They just sit there, and require external influence to plan and align them to make them work. We borrowed from the cognitive radio community and said: "Why can't our systems be intelligent and make intelligent decisions about where they want to connect, how they can connect, and how they want to change connections?"

The reason the system can achieve what it does is the intelligence of the individual nodes, which we build upon to provide overall, end-to-end system performance. In the CCS system we are not optimizing discrete link performance from a lamppost to a lamppost, or a small cell to a small cell. What we are doing is utilizing the overall interference information to create an end-to-end system that is optimal for the operator. This means that it delivers capacity where capacity is required whilst maximizing the spectrum utilization.

The self-organization process is dynamic so the whole system updates about every 200 ms, so we are continually monitoring these potential connections and we can change them as required.

We also looked at the 60 GHz systems and we wanted to emulate their performance. We really liked their high capacity, and we really liked their low latency, too. But what concerns us about 60 GHz systems is that they appear difficult to align, and to maintain alignment on a lot of the target infrastructure may be a challenge. Fundamentally they have very directional antennas: smallish dishes that need to be aligned. Some of these systems are going on moving structures. So we thought that was a difficult thing to do, and ultimately if you build a street-level connection, you are going to have to have multiple links. For instance, you may have to have one or more of these nodes on a particular lamppost pointing out in different directions. And, again, since very early on, we have been talking to people like city councils and infrastructure owners, and it became very clear that you could not have multiple backhaul units on a single piece of infrastructure. You have to have one.

We wanted to marry the capacity of the high-frequency systems with the easy installation of the sub-6 GHz NLOS systems, and that is what we really think we have achieved. We have brought the best aspects of both and brought them together in a single small package. We just put one of our units up there, and it stays there regardless how many links you connect. You don't have to keep going back and adding anything if you need to add more links.

We anticipate the deployment model for the small cells will be ad hoc: hotspot connections that eventually grow and grow, and then will migrate and mesh and build up that way. So we built the capability in our system to deal with that. So rather than having to revisit each site time and time again and re-plan everything, if I put a small cell in the middle of an existing planned network, the self-organization aspect enables you to turn up to that new site and backhaul it even if the system is quite established. So you can put a new network element in that will automatically become part of that system without any difficulty on your part.

If at any stage in the future somebody determines how to provide a 60 GHz system that enables us to produce the system solution we have created, we would utilize it – we are not knocking the technology and we do see a part for both 60 GHz and sub-6 GHz systems in small cell backhaul. Ultimately we are developing a solution, an end-to-end solution that optimizes the total cost of ownership for the operators. That is really what we are about, and we apply the technology that we have to enable that to happen.

Lance Hiley. I am glad that you made that last note about opex, because clearly one of the biggest stories in the whole small-cell debate is how operators can build networks with so many nodes. How does your solution fit into that whole cost aspect of it?

Steve Greaves. Looking from a capex point of view, we have an idea of where we need to be, and we are at quite a relatively early stage – low volume. But we made a conscious decision to design the system such that we could create what might be termed a volume backhaul product. That meant we couldn't put anything exotic in the physical part that would limit our cost reduction options going forward.

Fundamentally for the components that cost the most, the radio and the silicon, we discussed with our radio and chipset suppliers from day one where we felt we needed to get cost wise, if they could not convince us they could meet our cost expectations it didn't go into the design. We focus on capex because we know that there is a table-stakes argument that means that you can't have a ridiculous capex going forward.

But clearly our business is based on delivering big opex savings. We recognized that from day one, so how do you manage that? Well, we have one unit. We don't have an IDU, we don't have an access point or a CPE, and we don't have a remote terminal; we just have one unit. So one unit is deployed by the operator, and that adopts a certain software personality depending on how it is deployed. And once that unit is in place, it doesn't have to move. We don't have to keep visiting it, and we don't have to keep going back to the lamppost to add capability. If we want to build more links, you can ultimately scale your system to suit.

As mentioned, from an early stage we recognized the local authority planning departments could derail the whole small-cell deployment process so we made our system small so they didn't have to be involved in a large way. So by minimizing the number of visits and the deployment time, and by focusing on a system that we know will cost down with volume, we know we can minimize the TCO.

If you want a cheap tactical deployment that has got no real quality of service, that has high latency, then for sure you can put just about anything out there without much planning.

But if you want to ensure that you have a system that is scalable and that has sufficient bandwidth, low latency and good quality of service, historically you have to plan that in quite a complicated way. When we look at conventional PTP and PMP systems, you have to plan using extensive use of building data and planning tools. This is followed by a deployment and integration phase. That whole cycle takes just too long. To get your system set is very much like balancing a number of variables that give you no scope for growth. And here is what happens: if you want to grow that network and scale it, you have to go through the process again.

So removing the opex associated with ongoing network planning is a big part of what we do. Some technologies take so long to get in place, align antennas, and then make sure that the whole thing works, that it almost defeats the object. The key thing is to have a system that you turn up with and install with semi-skilled labor, switch it on and then you're finished. We are getting half-hour installs for our system with city council electricians. The self-organizing part of it is quite difficult to develop, but it is a fundamental requirement to ensure that you can deliver these capabilities in a cost-effective way.

Lance Hiley. Thank you very much, Steve. Your comments and observations will be a great contribution to this Senza Fili report on small-cell backhaul.

The report is going to provide an overview of the small-cell backhaul solutions from a variety of vendors and will be accompanied by in-depth interviews like this from all of the leading vendors that participated in the report. The report will be available to be downloaded from the Senza Fili website at **www.senzafiliconsulting.com**.

Acronyms

2G	Second generation
3G	Third generation
4G	Fourth generation
CPE	Customer premises equipment
eNB	Evolved NodeB
IDU	Indoor unit
LMDS	Local multipoint distribution system
LTE	Long term evolution
MPTMP	Many point to many point
ODU	Outdoor unit
PMP	Point to multipoint
PTP	Point to point
WiMAX	Worldwide Interoperability for Microwave Access

Building the ecosystem for small-cell backhaul

A conversation with Ed Chang, VP Product Management, and Eric Vallone, Senior Manager, Product Management, Cisco

By Monica Paolini
Senza Fili Consulting

Monica Paolini. Good afternoon and welcome to this conversation with Ed Chang, Vice President of Product Management at Cisco, and Eric Vallone, Senior Manager of Product Management at Cisco. I'm Monica Paolini and this conversation is part of a Senza Fili report on small-cell backhaul. The report provides an update on small-cell backhaul solutions and on the evolution of mobile operators' requirements for small-cell backhaul.

Today we will be talking about a new initiative on building the ecosystem for small-cell backhaul led by Cisco. Ed and Eric, thank you very much for joining us today.

I would like to start by asking you to give us an overview of the announcement you recently made on trying to create an ecosystem to move beyond a single solution.

Ed Chang. As background, we definitely see the large growth in data traffic as driven by the subscribers out there, both human subscribers and also machine to machine. And this data traffic growth is driving the need to deploy smaller and smaller cells in the rural, urban, or suburban environments.

Now, with the deployment of these small cells, clearly there needs to be a way to connect or to backhaul these cells to the rest of the mobile operator network. And what we're seeing is that there are many different locations where this small cell could be located. It could be on the side of a building, it could be on a lamppost, it could be on a bus stop, it could be in a rural area or in a small town somewhere. To drive connectivity to these new small cells, there is a need for multiple technologies to provide backhaul. The small-cell backhaul at a broader level can be broken up into a wired connection or a wireless connection.

In the wireless connection space, what we have announced is an ecosystem of partners that provide leading technologies in their own areas, which will be incorporated into the overall Cisco best-in-class mobile backhaul solution, the unified MPLS for mobile transport (UMMT) solution.

Monica Paolini. Can you tell us a little about what Cisco brings to this ecosystem? And who are the partners, and how do you select them?

Ed Chang. What we bring to the overall solution – and the ecosystem is part of the solution – is end-to-end connectivity for the mobile operators, from the small-cell site all the way back to their mobile core network. And that includes different types of connectivity, which includes wireless backhaul.

As part of that, we have formed an ecosystem of partners. They include BLiNQ Networks, DragonWave, Fastback Networks, NEC, RADWIN, and Siklu. And all of these bring their own unique technologies to wireless backhaul. What we use in our solution depends on the operator needs – where the small cell is located, what

Sponsored by:

CISCO

type of transmission speed they need, whether they need LOS or NLOS for each particular small-cell location. We would then work with any one of these ecosystem partners to deliver that end-to-end solution.

Monica Paolini. All the vendors you selected occupy different regions in the space of wireless backhaul solutions. Do you anticipate mobile operators will select different wireless solutions, in combination with wireline, and work with both you and the other vendors in the ecosystem to decide where to deploy what?

Ed Chang. Because there is going to be a large variety of installation locations for the small cells – as mentioned earlier, urban, suburban, or rural – many of these operators want to connect the small-cell locations, in the case of an urban environment, to the macro site itself. Between the macro-cell site and the small-cell site, both LOS and NLOS wireless technologies will be needed. So depending on how far the small cell is from the macro-cell site, what type of conditions are there between the two locations – in terms of buildings, trees, and other factors – we would then work with our ecosystem partners to come up with the best wireless solution to meet that particular installation type. And in some cases, we do envision working with multiple ecosystem partners at the same operator, because at different locations there may be different conditions, where one ecosystem partner's product may work and provide better characteristics than another ecosystem partner's.

Monica Paolini. In your vendor lineup, you have a wide range of solutions – PMP, PTP, licensed and unlicensed, different bands – and that gives mobile operators a lot of flexibility, but also you provide the ASR901 router, a common element.

What is it that vendors need to do in order to be able to be part of the ecosystem?

Eric Vallone. There's some fairly extensive certification and testing that we're putting them through before they can be a part of this ecosystem. They need to support unified MPLS for mobile transport architecture. There are some very specific tests that they need to perform, in terms of quality of service, of being able to fit into our management architecture, of how they handle timing across the network, of how they support IEEE 1588 and/or synchronous ethernet, and of how they support SLA management, including Y.1731 and IEEE 802.1ag. And to top it off, it's more than having a great technology; it's also having deployable technology. Not only do these vendors pass certification, they also have been actively deployed in concert with our Cisco technology in live customer networks.

Monica Paolini. It is support of the proper functionality and proving that it is actually working in real deployments.

Ed Chang. We are doing a two-part certification. The ecosystem partners all have their unique capabilities to provide the physical layer of connectivity, whether it is LOS or NLOS, PTP or PMP. They are the experts in the wireless transport connection end-to-end. And to what Eric said earlier, we need to make sure that the traffic, which is primarily ethernet traffic on an end-to-end basis over this wireless transport and over the routed infrastructure for overall backhaul, satisfies various requirements such as synchronization, packet timing, SLAs, Ethernet OAM, resiliency and scalability. So it is kind of a two-part technology certification: one is the wireless transport, and then on top of that is the ethernet or IP end-to-end features that we need.

Monica Paolini. Can you tell us a little more about the MPLS support they need to have? Why is it important for mobile operators?

Ed Chang. We feel that MPLS is important for mobile operators for a number of reasons.

From an overall scalability perspective, because mobile operators are going to have multiple numbers of nodes out there, and a node in this case is a small cell, you can imagine that it could be a 5x or 10x increase of small-cell versus macro-cell type of deployments. What operators want, and MPLS provides, is to have scalable connectivity to all these end points or small-cell sites, to provide the OAM and the SLA capabilities, to provide resiliency. For example, in the case of failover, if the path is broken, operators have an alternate path between the macro site and the small-cell site itself. All of these features are what operators have confirmed they need in a macro backhaul environment, and we do see many operators asking for the same functionality extended to their small-cell sites.

Monica Paolini. As mobile operators realize that they need to use different technologies, wireless or wireline, there is really a burden on their end to manage this diversity, because not only are they using different air interfaces, but also different features and

different capacities. So from a management point of view, it can be quite challenging. Is that what you are seeing in the market?

Ed Chang. Yes, we definitely see that. We're clearly seeing a multivendor environment out there for operators to choose from. And when you have a multivendor environment, each vendor technology and management system configuration is always going to be a bit different. So what we have been asked from operators is to provide an overall management system, which Cisco also provides, and it is called Prime. All this is integrated, to the extent that we can, with some of the ecosystem partners' management systems themselves – for example, their element management system for the device.

And the area that the operators are asking us to provide on an end-to-end basis starts with fault management – for example, if something breaks or they spot a fault. Provisioning: provisioning of the capacity on an end-to-end basis, provisioning of the SLAs, provisioning of the scalability and reliability features.

And then the third area they are asking us to do is on performance management. So that they can manage to their service level agreements, they can look at the OAM features, like Y.1731, to ensure the quality of the overall end-to-end circuit. An overall end-to-end management system is one thing we are working on.

Eric Vallone. The overall, end-to-end management system is designed to be a multivendor system from the ground up, to not only manage Cisco equipment but to manage other vendors' equipment.

To take it back to one part of your original question, I think what a lot of it is going to come down to is whether to deploy service at a location or not. I don't think carriers are going to have a choice but to be able to leverage multiple technologies. If not, they won't be able to reach their customers everywhere.

And that is, again, why we take the multivendor approach. It's to make sure that we can leverage these industry-leading solutions, so the carrier doesn't have to say no. If they need to increase capacity in a particular area, they have the ability to do so. Or if they hit a dead spot, they know they have the tools in their toolbox to be able to deal with that.

Monica Paolini. There is another dimension to that. You have a multivendor environment on the small-cell end. But you may also have another vendor in the macro layer. Operators that deploy small cells in the same frequency band they use for macro cells need to coordinate transmission. And that imposes even more stringent requirements on latency and reliability of the backhaul link. And I was wondering whether that has been a driver for you to step into this area, and what is it that you are doing to ensure that, for instance, latency is at an acceptable level?

Ed Chang. That is where our overall solution comes into play – the universal MPLS for mobile transport. What we envision as a deployment scenario is a small-cell radio and our ASR901 series router, then you have the wireless backhaul. On the other end, you have the wireless backhaul termination and then another ASR901 series router, which then connects the macro-cell site back to the operator.

And it is exactly between the macro- and the small-cell sites where SLAs' performance characteristics such as latency and error rate are measured, between the routers themselves. This is where all the intelligence is, and that is why we think we need to have MPLS all the way out to the small-cell site. And with that we are able to measure the physical layer characteristics and the latency between the two routers.

Monica Paolini. We have talked so far mostly about the performance of the backhaul link, but another

Figure 3. The ASR901 router. Source: Cisco

issue is installation. How are you going to ensure that installation is effortless, simple, and fast?

Ed Chang. Certainly the opex of deploying small-cell sites could grow as you send people out there. So our approach in the overall solution is to have zero-touch provisioning capabilities within the routers. We are working with the ecosystem partners also to enable this in their products. When you install the device, once the small cell lights up it finds its home-node location, which then provides its configuration down to the devices at the small-cell site. And this minimizes the amount of management capabilities you would need. You would not need to send a technician to the small-cell site to configure those devices, as this can be done remotely over the network.

Eric Vallone. Talking about the provisioning of the system itself, there is also the alignment of the antennas, and different technologies have different characteristics used to connect to end points, for instance in a microwave link. With NLOS we typically have easier operationalization than with LOS, but there are tradeoffs. Nonetheless the vendors we are choosing are absolutely leading in this industry to ensure the minimal operational requirements for antenna alignment and sync between units.

Monica Paolini. So there is going to be some variability depending on whether you will have LOS or NLOS, PTP or PMP links.

Eric Vallone. In general there is going to be a benefit of NLOS when you are looking at today's technology. It does not have to be a benefit forever, but it is definitely a benefit today.

Monica Paolini. Is there a particular challenge when you have multi-hop backhaul, with small-cells two or three steps removed from the aggregation point, and you may have different technologies – for instance, one step may be fiber, then two wireless technologies, with different performance characteristics?

Eric Vallone. That's one of the key attributes for any mobile backhaul system, with any of the varying technologies from the aggregation point in the mobile core, all the way out to the eNodeB or the small cell where the radio sits, and in the middle you have an IP-routed network, a microwave link, and the small-cell itself. There is a plethora of technology out there, and operators are looking very closely at the key performance characteristics across all these technologies.

One is packet timing or synchronization, on an end-to-end basis, between the radio and all the way to the aggregation site. In between you have a multitude of technologies that require key performance metrics that mobile operators are looking at.

The other technology they are looking at is OAM, which is used for SLA measurements such as delay and jitter on the link itself. Again, this link traverses multiple technologies from the cell site all the way to the aggregation point.

Monica Paolini. In addition to 4G small cells, you can also have carrier Wi-Fi small cells for access. Does your approach work with Wi-Fi small cells as well?

Eric Vallone. The connection to the user can be both through licensed (3G, LTE) and unlicensed spectrum (Wi-Fi). Once the connection is set up, all the backhaul characteristics that we mentioned earlier when talking primarily about licensed small cells apply to a Wi-Fi radio, if Wi-Fi is used to provide access to the subscriber.

Monica Paolini. What is the roadmap for your ecosystem? How is it going to evolve over time?

Eric Vallone. Increasing overall capacity per cell site is a main priority. Hand-in-hand is a reduction in the overhead and latency associated with that. We also see the ability to minimize the operational cost of the solution as a priority. Further, on the provisioning aspects that I talked about earlier, this may entail an easier installation of the wireless unit. One alignment technology, such as beamforming, may minimize the operational costs of turning up the solution. As I have alluded to earlier, the potential of deep integration within the router, the wireless unit, and potentially even the small-cell unit product that we have in our portfolio are other priorities for the coming years.

Monica Paolini. Ed and Eric, thank you very much, it was a pleasure talking to you today.

This conversation is part of a Senza Fili report on small-cell backhaul that provides an overview of small-cell backhaul solutions, along with in-depth conversations, like this one that we just had, from meeting with vendors who participated in the report. The report can be downloaded from the Senza Fili website at **www.senza-fili.com**.

Acronyms

3G	Third generation
4G	Fourth generation
eNodeB	Evolved NodeB
IEEE	Institute of Electrical and Electronics Engineers
IEEE 1588	Precision time protocol (PTP)
IEEE 802.1ag	Connectivity fault management (CFM)
IP	Internet Protocol
ITU	International Telecommunication Union
ITU-T	ITU Telecommunication Standardization Sector
LOS	Line of sight
LTE	Long term evolution
MPLS	Multiprotocol label switching
NLOS	Non line of sight
OAM	Operations administration and maintenance
PMP	Point to multipoint
PTP	Point to point
SLA	Service level agreement
UMMT	Unified MPLS for mobile transport
Y.1731	ITU-T recommendation Y.1731

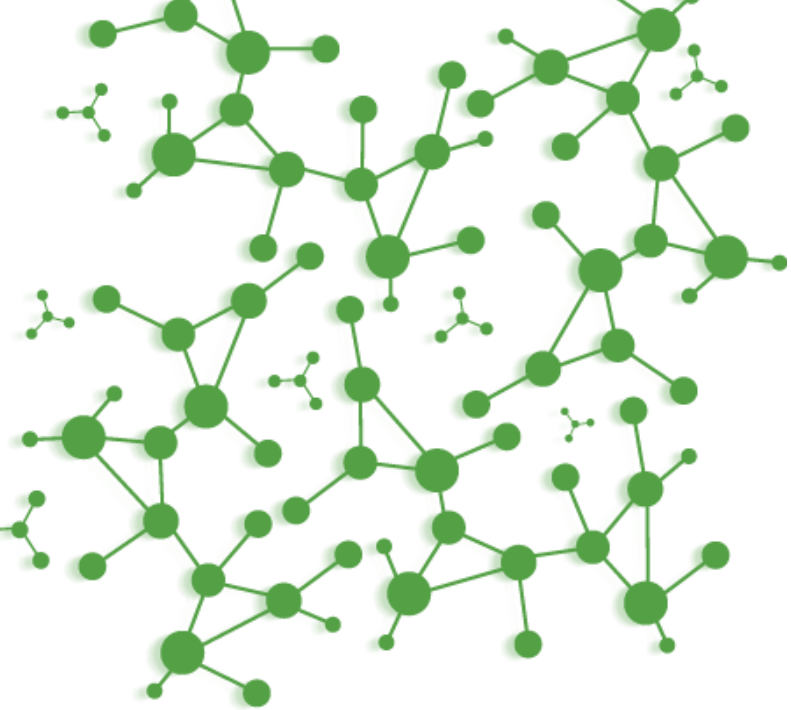

Building the small-cell backhaul toolkit for mobile operators

A conversation with Greg Friesen, Vice President of Product Management, DragonWave

By Frank Rayal
Senza Fili Consulting

Frank Rayal. Hello and welcome to this conversation with Greg Friesen, Vice President of Product Management at **DragonWave**. My name is Frank Rayal. This conversation is part of Senza Fili's report on small-cell backhaul that gives an update on small-cell backhaul solutions and on the evolution of the mobile operator requirements for small-cell backhaul.

Today we are talking with Greg Friesen at DragonWave, a leader in microwave backhaul solutions used by carriers and in other market verticals as well. Greg, thank you very much for joining us in this conversation on small-cell backhaul.

I would like to kick this off by asking you to give us an overview of DragonWave and more detail about your solution for small-cell backhaul.

Greg Friesen. Thank you very much, Frank. DragonWave, as you mentioned, is a leader in microwave backhaul, and we have products from the 2 GHz band all the way up to the 80 GHz band. We have historically focused on all-outdoor and packet microwave systems, and have expanded on those in recent years with hybrid capabilities and split-mount systems as well. We focus on mobile operators, but also have a large customer base in enterprises and vertical market, as well.

Small cells have been a strong focus of ours over the past year or two. We have a number of products that have been announced over the past two years. The first was our Avenue Link, which is a very small, compact product with an integrated antenna in the 24–60 GHz range, delivering over 500 mbps full duplex of capability. Next product, which we have just recently introduced, this year, is the Avenue Link Lite, which is our sub–6 GHz point-to-point product, also in a very small, integrated package with antenna, available in the 2.3 to 5.8 GHz, licensed and unlicensed bands. We view that as an important addition to the portfolio. Where the Avenue Link has a very high capacity for aggregation links and sometimes for access links, the Avenue Link Lite is a very good complement to that by providing non-line-of-sight and near-line-of-sight capabilities when obstructions are in the path, to help increase the network coverage.

Frank Rayal. Can you give us more details on the Avenue Link product in terms of performance metrics and other important parameters?

Greg Friesen. In terms of performance capabilities, the Avenue Link is in the 24–60 GHz band. Reach capabilities depend on the band, but can go anywhere from 500 meters to a couple of kilometers. Capacities are in excess of 500 mbps, with support for up to 2048 QAM, varied low delay at 0.1 milliseconds, full Ethernet OAM capabilities, and compression capabilities to further increase the traffic capacity beyond that 500 mbps. There is a bulk compression capability as part of the Avenue Link.

Frank Rayal. The licensing regime for spectrum in the 60 GHz band is very different from the lower bands, and maybe some features of the product may be different. Can you elaborate on that?

Greg Friesen. Yes, certainly the 60 GHz licensing regime in most places is an unlicensed band, so there are obviously benefits to that, in that it means that you can deploy very quickly and rapidly and you don't have to acquire spectrum. There is no annual cost, and it helps improve the economics. The downside, of course, is that there can be interference; however, the interference environment is much different than that in some of the more traditional unlicensed bands, like 5.8 GHz, in that in the 60 GHz band there isn't broad consumer use today. The beam width is very narrow and, in addition, the propagation distance is quite short because of the high air loss in this frequency band. So there is much less frequency and interference concern in this band than in other unlicensed bands, and then there is much more spectrum, but nonetheless it is still something that has to be engineered around.

Figure 4. Avenue Link (24 – 60 GHz). Source: DragonWave

Figure 5 Avenue Link Lite (sub-6 GHz). Source: DragonWave

Frank Rayal. How does your technology compare to solutions from other vendors in the 24–60 GHz bands?

Greg Friesen. I think that one of the key things in all of those bands is that we developed and integrated a 5-inch antenna to have a very small, compact footprint. We believe that it is very important in terms of deployability.

We support up to 2048 QAM at some of those lower bands, and 512 QAM at the 60 GHz band, which is industry-leading in this space. We also support the only bulk compression engine which can provide 30%-plus of capacity gain, so we feel that we are quite differentiated in terms of capacity capabilities that also can enable you to use lower modulations, because you've got the compressor, which can help you increase your reach capabilities.

Frank Rayal. I am curious, does the data compression work on a per-traffic type, or on all the data at the same time?

Greg Friesen. It is assignable per traffic. Basically, there are four compressors, and you can assign each of your four queues to a compressor to be compressed or not be compressed, or to different compress settings.

Frank Rayal. Can you address the different versions and capabilities you have on the Avenue Link Lite and the sub–6 GHz product?

Greg Friesen. The Avenue Link Lite is targeted more at the last mile of the microcell network. It has capacity capabilities up to 230 mbps and it supports near- and non-line-of-sight deployments. It is fairly low-latency as far as TDD systems go – 1 to 2 milliseconds latency for the system. Some of the things we have built in, to differentiate, and are focusing on are, of course, the delay and minimizing the packet delay variation. We do cut-through prioritization for timing packets, and we also support 1588 transparent clock and synch E, both of which we believe are quite important for deployment of LTE and LTE-TDD networks.

Frank Rayal. Now if you can address the frequency bands, because I think I heard you saying that you

have the product in the 2 GHz bands in addition to the 5GHz band.

Greg Friesen. That is right. We've just released the product. The first band we are releasing that is just becoming generally available is the 5 GHz band. So 4.9–5.8 GHz, and the next band, which we will be releasing early fall, is the 3.5 GHz, and we will be following that up with the 2.3–2.6 GHz bands at the end of the year.

Frank Rayal. What is the maximum channel bandwidth that you can do in that product?

Greg Friesen. We support user-selectable channel bandwidths from 5 to 40 MHz.

Frank Rayal. What are you hearing from the operators on deployment issues, and how do the products you have address the deployment issues?

Greg Friesen. Well, I think one of the things we are hearing is that it is not one-size-fits-all, at all, in terms of small-cell backhaul, and that is why you can see that we have pretty wide breadth of products and solutions. Some operators own 3.5 GHz spectrum, so that can be something very attractive to them. Some own 2 GHz spectrum, some are willing to deploy unlicensed and others are not, and that may be 5.8 unlicensed or perhaps 60 GHz unlicensed, and then some have bands in the 24–38 GHz bands. Some may have 26 GHz spectrum or 38 GHz spectrum, and a lot of that will vary by market and even location in the market.

Operators are looking to us to provide a complete solution that they can use in many areas of their network, and that means more than just a good 60 GHz product or just a 5.8 GHz product, but a wide tool kit and solution set that can help them deploy a wide network in many areas.

Frank Rayal. I think that is going to be one of the major challenges in terms of scaling small-cell deployments. Different solutions have to be deployed in different scenarios, so how do you help the operators in terms of minimizing the deployment timelines and then making the deployment and the design faster?

Greg Friesen. A couple of things we are doing there. One is all these products have very similar Ethernet capabilities, and they are all managed by a single management system, our DragonView management system, which enables end-to-end provisioning. We have made the systems have a fairly wide beamwidth, wider than some of the traditional systems to simplify alignment, and we have built in a number of alignment and commissioning tools to simplify deployment. Going forward, our future vision is to do more integration with base station vendors – and we have started that discussion with some of the key players – to be able to maximize the reuse of all resources and have a single install rather than two separate installs when you are mounting a base station and a backhaul link.

Frank Rayal. When it comes to integration, do you think all solutions will be considered, or will there be a focus on a specific band or specific type of solution, let's say line of sight or PTP, versus PMP?

Greg Friesen. I think, for the same reasons I discussed earlier, that it won't, although it would be ideal to be able to focus on a single band. We think that we will need to offer a pretty wide set of solutions. There may be some band that we can hit that are higher-volume, that we target more integration on, and that maybe are more suitable for integration, but we certainly see and believe that the base station vendors have a common vision that we will need a wide set of solutions and won't be able to just integrate with one product.

Frank Rayal. If I may ask a technical question: The 2048 QAM modulation, I hear a lot of people talking whether it is useful or not useful, and on the benefits of having it. What is your view on that?

Greg Friesen. Certainly you are right, 2048 QAM has a reduction in system gain, so that results in a reduction in link range. In small cells, though, you are not that constrained in the link range, especially in the 24–38 GHz bands when you have a 400-meter link. The capabilities of these bands are well beyond that – so it fits actually quite well in the small cell, and then 2048 QAM can provide about a 30% increase in throughput, even in longer links.

Because we have hitless adaptive modulation, we can allow the user to take the benefit of these higher modulations during clear periods when there is no rain, which is 99% of the time. And then the system will switch down during heavier rain back to what you would have been operating at regularly in a 256 QAM or 128 QAM, so it is really just taking advantage of those good weather conditions, which is most of the time, and increasing throughput during those periods.

Frank Rayal. When it comes to operator requirements for small cells, my feeling is that there has been a bit of confusion, or still a bit of searching, for really what we need to have for small cells in terms of availability, latency and throughput. From your point of view in discussing with the different operators, what do you see are the key requirements, and how have you addressed them in your product line?

Greg Friesen. I completely agree with you, I believe that there is still a lot of confusion – and I don't think that it is just confusion. I think that the spectrum varies a lot by operator.

Taking one of the parameters that you discussed as an example, the availability, so many operators have said, "Well, a small cell is part of my mobile network, and therefore I expect the same in backhaul availability for my small-cell network as my mobile network, and so I want four-9's or five-9's on this small-cell network," and that drives all sorts of things, like battery backup, redundancy, etc.

Other operators look at it differently and say, "Well, I have redundancy already, in that I already have coverage in this area, and this is really augmentation and capacity increase to my coverage, so my redundancy is the macro network, it doesn't need to be through the backhaul in the small-cell network, and therefore I would much rather optimize costs on this network, not have redundancy, and I accept the lower availability, maybe a three-9's, maybe even a two-9's and a five, so I am willing to look at different parameters."

Certainly you want availability. I don't think there is a right answer, I think that each operator has a different mentality and a different perspective, and we are working to support all of those.

In terms of delay, you do see ranges in delay, but typically from the operators we are seeing fairly low delay numbers. Typically they are targeting 5 to 10 millisecond end-to-end delay across the entire backhaul network, and so that may mean multiple hops through Ethernet switches, maybe even through some fiber as well. So certainly we are seeing that fairly commonly.

We also see a very wide range in capacity. You will see anywhere from someone saying, "I just want best effort, just to be able to have my cell working," or even, "5 to 10 mbps will work," others are saying, "I am using this as a capacity-increase strategy and I need at least 100-plus mbps per site." We do see a very wide range, and we are certainly working towards covering the higher end of that range. We believe that even if the capacity is not expected to be there short-term, in the long term the operator is going to grow into it.

Frank Rayal. These are tough issues, and I think it has a lot to do with the way small cells are deployed and the coordination aspects between the macro-cell layer and the small-cell layer, so we will see how these things will develop.

What are your thoughts on the total cost of ownership for small-cell deployments, and specifically the backhaul? We see the base stations have been reduced significantly in cost, going from a three-sector site to a single-sector site with much lower power, but backhaul solutions across the entire bands are pretty much the same thing still.

Do you think that there is a significant issue here related to cost, and how can the industry move forward and address that, and how are you trying to address that with your products?

Greg Friesen. First of all, I think if we look at DragonWave, we have seen the average selling price go down quite a bit, and our costs go down quite a bit, over the past ten years, definitely double digits per year.

That said, that certainly is not enough for a small cell. We are not looking for tens of percent here – we are looking for 3 to 5 times the difference, sometimes more. So a significant reduction.

There are three key areas to address this. One is the site cost – the cost of site leases, site acquisition, etc. We believe that the best way to address that is through a very small and low-weight solution, and that is what we are targeting and what we have built.

The second is the actual equipment cost, so we are working to optimize that, although that is one of the smaller portions of the total cost of ownership. In addition to that is the part which often gets locked in with that equipment cost, the capex, is the spectrum cost, so by having a flexible array of spectrum options, we are looking to allow the operator to minimize their spectrum cost.

Third and probably most important is the operation costs – the cost of deploying and maintaining the network. By simplifying the solution, doing as much integration as possible, and providing as much Ethernet OAM as possible. We are certainly trying to minimize that by offering end-to-end visibility in our management system, and very simple installations.

So all three areas we are going after. We believe the most important is the opex area. The second most important is the space and the site cost. The third is the capex costs.

Frank Rayal. In terms of future evolution of your product line, how do you see that evolving, and what are some of the key features that you plan to introduce in the future?

Greg Friesen. Obviously the one that everyone talks about is size, weight and cost, so we will be looking to optimize around that, as new technologies and new chips come out, and there is more integration on chips. We will be looking to leverage that, and to get our products smaller and minimize the cost. That is a pretty obvious step for everyone in the industry.

The other – a more important one but one for which there is not necessarily as clear an answer on, but I hinted at – is minimizing the opex and installation cost: making it very simple to deploy, improving, simplifying alignment, and automating alignment and realignment if it gets put out of line. So optimizing opex, and installation time and cost is a very important focus of ours, and what you will probably see us bringing on multiple features to minimize that cost.

Frank Rayal. Anything else that you would like to add in terms of final thoughts on your product line and how you see the small-cell market evolving?

Greg Friesen. We believe that we are at the infancy of this market. There have not been any large, wide-scale deployments to date. There have been some small deployments, and more of what I would call test deployments, than any wide-scale deployments. So we, as well as others, have all built products that are targeted at what we think the market requirement is, but I don't think that anyone really knows yet, because there hasn't been a wide-scale deployment to really see the challenges of street-level backhaul, which is really what we are facing here.

As we see the challenges of street-level backhaul on a wide scale, it is going to shape the operators' network architectures, as well as the product solutions that we need to offer to optimize their network architectures.

We strongly believe we are at the very early stage of this and that we have got a good solution for what we know about today, but that solution is probably going to have to evolve into things that we don't even know yet, and solve problems that we don't even know exist yet.

Frank Rayal. All right. Thank you, Greg, very much for this conversation. This was Greg Friesen with DragonWave. This conversation is part of a Senza Fili report on small-cell backhaul that provides an overview of small-cell backhaul solutions along with an in-depth conversation like this from leading vendors who opted to participate in the report. The report can be downloaded from the Senza Fili website at www.senza-fili.com. Thank you for listening in, and, Greg, thanks again for this conversation.

Greg Friesen. Thank you, Frank.

Acronyms

3G	Third generation
4G	Fourth generation
GSM	Global system for mobile communications
HD	High definition
IP	Internet protocol
LOS	Line of sight
LTE	Long term evolution
NLOS	Non line of sight
SON	Self-organizing network
VoIP	Voice over IP
VoLTE	Voice over LTE
PTP	Point to point
PMP	Point to multi point
OAM	Operations, administration and management

Carrier-grade microwave backhaul for street-level small-cell deployments

A conversation with Babis Papanastasiou, Product Line Manager, Intracom Telecom

By Frank Rayal
Senza Fili Consulting

Frank Rayal. Hello, and welcome to this conversation with **Intracom Telecom**. My name is Frank Rayal. This conversation is part of a Senza Fili report on small-cell backhaul that gives an update on small-cell backhaul solutions and on the evolution of mobile operators' requirements for small-cell backhaul.

Today we are talking with Babis Papanastasiou of Intracom Telecom, who is the Product Line Manager for StreetNode and the WiBAS point to multipoint platform. Intracom Telecom is a leading vendor of PTP and PTMP wireless backhaul solutions.

Thank you, Babis, for being here with us today. To start I would like to ask you to give us a brief introduction about Intracom and your product line.

Babis Papanastasiou. Intracom has more than 15 years of experience in wireless technologies and provides a wide solution portfolio of PTP and PTMP technologies, having the ability of combining these two qualities together and providing the optimum solution to our customers.

We have also recently addressed the needs of small-cell backhaul with a new and, I believe, a very interesting solution, which we named StreetNode.

Frank Rayal. Can you position that product for us, please – how it works, and what is the frequency band that it operates in?

Babis Papanastasiou. We are positioning our solution in the microwave band. It is a microwave system operating in area-licensed frequencies. These are the 26, 28, 32 and 42 GHz frequency bands, where the operator has the ability of getting a license for a full area and then has the ability to deploy the solution, without having to go back to an authority and get license for the backhaul per hop. This makes this solution very effective, to start with.

We also explored the benefits of the microwave technology. Microwave technology is a mature technology and has a lot of benefits. It is a carrier-grade technology, and we have taken all of the benefits of the technology, and now we have put them in a very attractive size, because StreetNode is a very small unit: it is only 4 liters in volume. It blends into the urban environment so that it is almost invisible, and it can be installed on a lamppost or on walls in an urban environment and provide a very efficient solution for the small-cell backhaul.

Because the solution is line of sight, it is very deterministic, which allows you to dimension your network. It is a carrier-grade solution, which we believe is mandatory for a backhaul solution.

Frank Rayal. Would you expand more on that? What do you mean by carrier-grade solution, and what are some of the features that the product has that would make it carrier-grade?

Babis Papanastasiou. First of all, the deterministic way you design the network makes you feel confident

about the performance of the network. StreetNode has a high capacity: we are offering 540 mbps per sector or per link, and we can have aggregation hops with two channels of 56 MHz that can give 2.2 gbps capacity per aggregation node. It has minimal latency, at below 1ms. It has the quality of service and all the mechanisms that are needed for a carrier-grade solution.

Frank Rayal. As a PTMP solution, how many multipoints will the solution be able to support?

Babis Papanastasiou. I would like to make a note here and correct a misconception. It is not only a PTMP solution; it is a very flexible solution that can accommodate any type of architecture. It can be PTP and it can be PTMP and it can do relays, so we can extend the connectivity within the urban clutter and create relays that literally reach any location. It is a line-of-sight solution with carrier-grade characteristics of a line-of-sight technology, but also we have all these mechanisms that bypass the problem of line of sight.

With our solution, our proposition, we can reach any location by implementing relays in the urban environment.

If you were talking about strictly PTMP, we can aggregate from a single hub 30 terminals, but it is not only a multipoint solution, it is a very flexible solution that can adapt to the needs of a specific network deployment.

Frank Rayal. The support for relays is interesting because, as you have mentioned, you will be able to go deep into the urban clutter. Can you expand on that to help the audience understand the relay functionality?

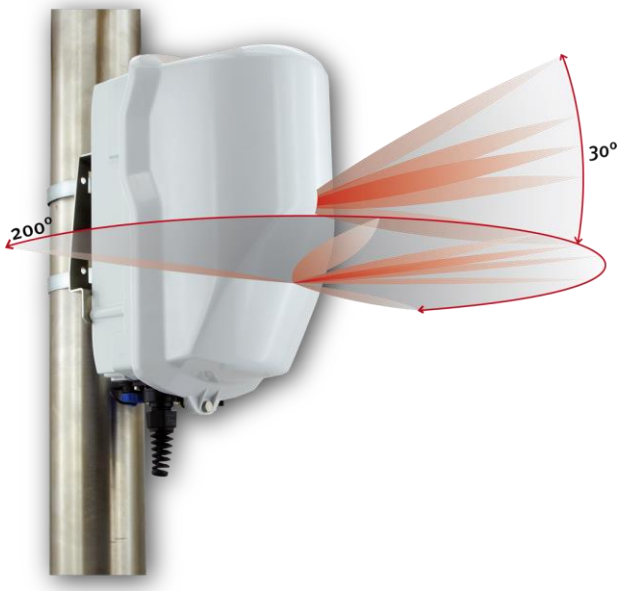

Figure 6. StreetNode PTMP microwave backhaul system. Source: Intracom

Babis Papanastasiou. First of all, we explore the benefit of the multipoint, where from a single hub you can have many terminals, and this in principle gives you less equipment, which then translates to less expense. But then you cannot really reach all points from a central aggregation point. You need to expand connectivity with a very flexible solution, and this is another key point to our proposition: we have a single hardware unit that is defined by software for different modes of operation.

The unit inside the urban clutter can operate as a terminal, as a relay PTP extension or as a relay PTMP extension in the clutter. By putting together one or two units on a lamppost, we can extend the connectivity and reach locations that do not have direct line of sight.

Frank Rayal. The product has a feature that supports a self-alignment. Can you expand on that?

Babis Papanastasiou. That is another key point about our proposition, because we need to bring intelligence into our solution. We are talking about a dense deployment of small cells, and, for a massive deployment, we need to have the intelligence in our solution that will simplify the installation, provisioning and, later, maintenance of the network.

Autoprovisioning is a key feature of our solution because we expect unskilled personnel to carry out the installation tasks. Aligning a PTP link can be very complex and time consuming. Imagine the locations where this is going to happen. This is going to happen in an urban environment with a lot of traffic, cars and people going by, so you need something that will be realized very quickly.

Autoalignment gives this option, because the system has the ability to find the optimum position and location, and to align to the other end automatically. You power up the system and it will find the best alignment.

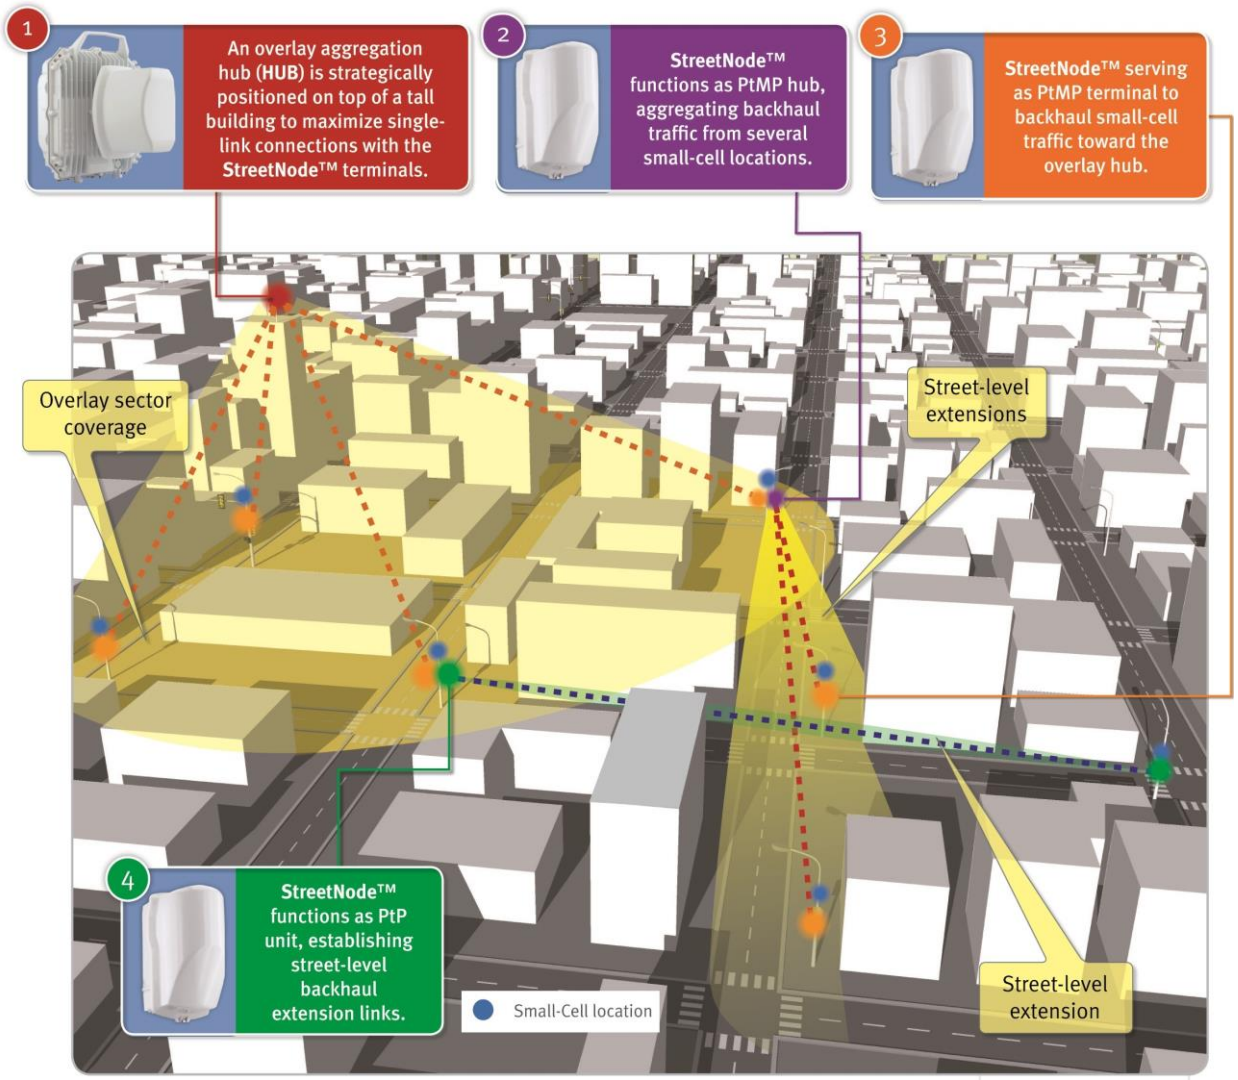

Figure 7. StreetNode combines overlay coverage and street-level extensions for optimum small-cell backhaul. Source: Intracom

Furthermore, the system will do the autoprovisioning. All of the configuration parameters, all of the profiles are created offline, and downloaded directly to the node. By the time you power up the device, very simply, in a few minutes, the node will automatically go alive.

Frank Rayal. How would you describe the competitive advantage of your solution and technology over others?

Babis Papanastasiou. First, we are talking about a frequency band that is already there. It is a block spectrum allocation, which can be reused very easily and provides a very efficient solution. With the relays we can reach everywhere with carrier-grade performance. It is a very flexible solution: we have a single hardware unit, which can be software defined in various operation modes, either PTP or PTMP.

Imagine how much logistics are simplified when you can deal with a single hardware unit that operates in different modes and provides reach anywhere in the clutter. Exploiting the benefits of the multipoint architecture with less equipment and expense is something that should be highly considered.

Frank Rayal. Is this spectrum widely available, and do you have any ideas on the cost of spectrum? How does it factor into the operator business case?

Babis Papanastasiou. It is available in England. The higher spectrum bands are less utilized, and recently we saw the higher bands open up in many countries. We also have seen a very attractive licensing scheme adopted by the authorities. That makes this spectrum quite attractive.

I would also like to mention that our solution is deployed at street level. The microwave transmission at street level is decoupled from whatever happens at the rooftop level. We could have traditional microwave links rooftop to rooftop and, at the same

time, without any interference, we can reuse the same spectrum inside the clutter and provide very efficient solutions. Many tier-one operators today may have this spectrum, which may be underutilized, so effectively they can deploy our solution without any additional cost.

Frank Rayal. When you are speaking with operators, what do you hear as the major issues they encounter in small-cell deployments and backhaul? How have you addressed those in your product line?

Babis Papanastasiou. Small cells are a very hot topic, and all the tier-one operators have been preparing for this. Providing a solution that is carrier grade is important, and carrier-grade backhaul is a requirement. Having the capacity and latency requirements to satisfy upcoming RAN technologies like LTE and LTE-Advanced is important.

The automated procedures that we bring in our solutions, which will simplify the deployment, are also a very important issue.

And of course we must have a unit that is not "telecom": a unit that will blend in with the urban environment, and it will be almost invisible, because at the end of the day all the operators must have authorization to install telecom equipment on a lamppost or on a wall in an urban area.

Frank Rayal. Can you also comment on redundancy? How does it apply to the small-cell environment?

Babis Papanastasiou. Having a multipoint solution, we can easily make redundant the most critical point: the hub. We can have hubs that provide coverage, and these are deployed in a redundant way. When one aggregation point fails, we will have a standby that will take over the process. Redundancy is critical in aggregation, and this is where we offer it.

Frank Rayal. What do you think are the major costs that the operator would encounter in a small-cell deployment, and how does your solution address the issue of maintaining a low total cost of ownership?

Babis Papanastasiou. The multipoint technology brings less equipment. Having one aggregation point connected to many terminals translates into less expense. In terms of spectrum license and fees, the area-licensed spectrum pricing is very attractive, and most importantly, it is fixed, so the operator knows exactly how much to pay. It is not an expense that will grow in time with more links. Regardless of the number of the deployed links, the operator knows exactly how much the spectrum license will cost.

This is very important, as it may be the case that they have already acquired the spectrum and it is currently underutilized, which means that it will effectively add no cost. Furthermore, the automated alignment process we have introduced will simplify installation, which means less install time. Making the product small and compact simplifies site acquisition and probably the cost.

Frank Rayal. How do you compare your solution to other backhaul solutions, let's say those that operate in the sub-6 GHz or in the 60 GHz spectrum? How do you see your solution fitting within the overall space in backhaul?

Babis Papanastasiou. The small-cell market is a new market, and we have very high expectations yet to be realized. Probably by next year we will see the market growing.

The major competitor for us is the sub-6 GHz, non-line-of-sight backhaul, where there are many issues. Probably the bandwidth is not enough. If the spectrum is licensed, it could be expensive. If it is unlicensed, then it is not a carrier-grade solution. If the spectrum is not there, it means that there is not enough capacity to address the requirements of small-cell backhaul.

Non-line-of-sight solutions give a statistical dimensioning of the network. They have a varying performance, and by no means they are a carrier-grade solution. In the long run, I believe that sub-6 GHz frequencies are going to be used for the access network to address the increasing need for bandwidth.

On the other end, we have the millimeter-wave technology. It is a very new technology, yet to be proven. Our solution is based on a proven technology. Millimeter-wave technology has very narrow pencilbeams, which could create some problems – for example, pole sway could be a problem – whereas our solution is based on microwave, a known and proven technology. The engineers know how to plan and design microwave networks. What we are proposing here is to transfer this knowledge to the street level.

Frank Rayal. How do you see your product evolving in the future, and what features do you plan to add?

Babis Papanastasiou. We would like the product to cover all the area-licensed bands we see coming – 26,

28, 32, and 42 GHz. We plan to address all these area-licensed bands in the product evolution, as well as carrier-grade networking, and to provide all the features for quality of service, OAM certification and MPLS. All of these issues are going to be addressed in our solution.

Ultimately, we would like to see our solution integrated with a small cell and effectively have a single unit that combines the small cell and the backhaul in one box.

Frank Rayal. Thanks, Babis, for being with us here and for this conversation about your products, the StreetNode and the PTMP wireless platform.

This conversation is part of a Senza Fili report on small-cell backhaul that provides an overview of small-cell backhaul solutions, along with in-depth conversations, like this one that we just had, from meeting with vendors who participated in the report. I would like to thank the audience for listening in and mention that this report can be downloaded from the Senza Fili website at **www.senza-fili.com**.

Acronyms

LTE	Long term evolution
MPLS	Multiprotocol label switching
OAM	Operations administration and maintenance
PTMP	Point to multipoint
PTP	Point to point
RAN	Radio access network

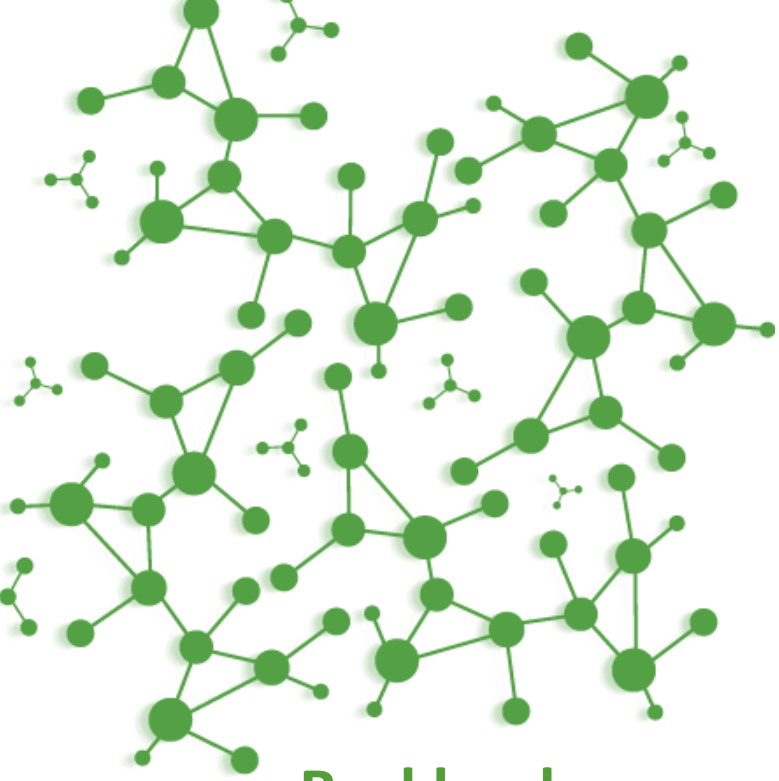

Backhaul as an integral component of small-cell networks

A conversation with Dejan Bojic, Director of Global Product Strategy, Mobile Backhaul, NEC

By Lance Hiley
Senza Fili Consulting

Lance Hiley. Welcome to our audience with **NEC**. This conversation is part of a Senza Fili report on small-cell backhaul that gives an update on the small-cell backhaul solutions available today and on the evolution in the mobile operator's requirements for small-cell backhaul. My name is Lance Hiley. I am here today on behalf of Senza Fili, and I am joined by Dejan Bojic, Director of Global Product Strategy, Mobile Backhaul at NEC, a leading wireless backhaul company. NEC has announced in the past 18 months some new perspectives on wireless backhaul, which we will be hearing more about today.

Dejan, could you give us an overview of your company and the product that targets the small-cell backhaul market?

Dejan Bojic. We have been supplying mobile backhaul solutions now for many years – several decades, in fact, from the early days of mobile networks. We hold a leading position in the wireless backhaul marketplace, having supplied solutions to 149 countries around the world, and we are bringing these years of expertise to a new focus on small-cell backhaul. We have been looking at what technological innovation is necessary to fulfill our customers' technical requirements and to deliver them with the right economics.

The way we like to work with our customers in the backhaul area is to be a one-stop shop for all their needs. What does that mean? First, it is about having the right technology tool kits. Probably the first thing I should explain is that from NEC's perspective, we provide a suite of technologies and products to address the various use-case scenarios for small-cell backhaul. Whether those small cells are there to address capacity needs in dense urban areas or perhaps to cover the areas of poor coverage in the rural communities, NEC aims to have solutions for both scenarios.

But beyond just the technology and product tool kit, it is also about providing the right sort of know-how, technical expertise and professional services to help with all the issues: securing the sites, planning the deployment, executing that deployment in as short a time as possible, and then, of course, helping our customers maintain the system with the lowest opex possible. That is our approach, and in terms of technology leadership we have focused on certain new technologies that perhaps some people would be surprised to see coming from NEC.

Lance Hiley. When you say that you have a tool kit, does that mean that you are in a position to address the entire spectrum of the small-cell marketplace, not only the hotspots, the dense urban deployments, but also the notspots, the rural applications that a lot of people, a lot of the operators are thinking of using small cells to serve as well?

Figure 1: iPASOLINK SX utilizes 60 GHz radio spectrum for small cell backhaul. Source: NEC

Dejan Bojic. Absolutely, that's exactly the case. And actually it is interesting to very briefly reflect on the recent history of small cells, and on the development of small cells in terms of adoption by mobile operators. If we look back at it, where small cells really started taking off is in the home environment, through what was called femtocells until not so long ago. NEC has a leading position in indoor small-cell solutions.

Over the last couple of years the industry has accepted more and more that we need to take the small cells outside, and that we also need to move from closed-mode to open-mode access. We look at where do these open-access small cells bring the most value, and we have quite polarized application cases.

There are two key cases. One is the dense urban areas where we need the capacity. People often talk about the areas where people congregate – sports arenas, train stations, shopping centers, all those kinds of places. And on the other hand, we are talking about rural communities, where perhaps it was cost-prohibitive until now to deliver macro-cellular solutions.

NEC provides small-cell solutions for both of those use cases, and by implication we also supply backhaul solutions for both. Now the requirements for the backhaul in those two use cases are significantly different, but then again, we are very well aware of that, and we provide the technologies accordingly.

Lance Hiley. Recently, Dejan, NEC has been talking about a different approach for wireless backhaul. NEC has been a leader for years in conventional point-to-point microwave links, operating in the 6 to 42 GHz range, but you have been talking about using different frequency bands.

Dejan Bojic. The great advantage of being a company like NEC, with such a heritage in microwave backhaul, is that we get early visibility of our customers' requirements. So it was as long as three years ago that we started a dialogue with mobile operators around the topic of outdoor small cells, focusing on the use case for dense urban areas.

Very quickly it became very apparent that if we are to translate this into its simplest, most direct form, the requirements are for high-capacity connections to small cells at an ultra low cost. So what we had to do is to look at how well we can innovate from a technology perspective to deliver that site capacity at that very low cost.

We started looking at different technology options, and came to a realization that we had to move to higher frequencies in order to deliver this. We started looking into millimeter wave and, specifically, found that at 60 GHz we could deliver what is effectively a future-proof capacity. The nature of 60 GHz propagation, the nature of the regulation and availability of the spectrum, also meant that we can actually deliver these products at very low cost, into a very small footprint, and at a very low cost.

Shortly afterwards, NEC started promoting the use of 60 GHz. We are happy to see that the industry has reached the consensus around this and that indeed 60 GHz is now recognized as one of the primary wireless technologies to deliver street-level connectivity to small cells.

Lance Hiley. What is the cost of using 60 GHz spectrum, and what is the availability like in different markets around the world?

Dejan Bojic. It is pretty good, and it is one of the reasons why we focused on 60 GHz as the key technology choice. If we look at the regulations in some of the key markets, like the United States and Europe, there is up to 9 GHz of spectrum bandwidth available at this frequency. And the industry recommendation for the use of this spectrum is to be license exempt, which basically means anybody can use it and it is free to use.

When we consider today's projections around quantities and high densities of small cells, the fact that operators don't need to pay for the backhaul spectrum and also that they don't need to incur the administrative costs of managing spectrum licenses is a benefit.

Of course, as with any spectrum regulation, the next level of regulatory implementation is at a national level. In many cases regulators have already followed through on general recommendations and have made this spectrum available for use. In some other countries there is still some way to go in recognizing the need for the use of 60 GHz and making it open for use. But we are quite confident that with the momentum behind small cells and the industry consensus for the need for 60 GHz, this should not be a major concern.

Lance Hiley. What would you say is the NEC technology advantage of your solution operating at this frequency over your competitors? What are the unique features? What makes it better than the other options that are out there?

Dejan Bojic. There are certain specifics around our feature sets based on operators' needs. One of those is the use of what we call narrow-band channels. The majority of products in this space use relatively wide channels – anywhere between 250 to 500 MHz, even in some cases 1,000 MHz – which is fine. NEC uses narrow-band transmission at 50 MHz, and that gives us much more room to maneuver in terms of the link design and propagation characteristics, the capacity and distance we can achieve, and the equipment design.

We actually designed the components for the new 60 GHz product from scratch to make it as efficient and low cost as possible. The use of the 60 GHz band is not new. It has been used in the past, but primarily for enterprise fixed access applications. But we have specifically targeted iPASOLINK SX as a brand new product with the right economics, right cost points, and right performance levels – to make sure we can reach several hundred megabits per second throughputs and make it cost effective.

Transmission at high frequencies is strength of NEC. This is not the first time that NEC has developed 60 GHz products as we have had products in the past available – for example, for the backhaul of high-definition television, which we had deployed in certain markets in Asia. This expertise has proven powerful for our customers and us because in our early trials of iPASOLINK SX the RF capability and RF performance were very solid right from the start.

With any new product you would expect teething problems at the beginning before the technology matures. But actually we are not seeing much of that at all. The performance of the iPASOLINK SX is already mature as proven through early evaluations with several of our customers, giving them the confidence that 60 GHz technology is ready to go.

Lance Hiley. That's fascinating, hearing about how you actually developed a product from scratch, and I guess the way I would summarize is that it sounds as though your radio design approach has been to deliver something that is very flexible, very robust, to your customers.

But at the end of the day, the radios have to be deployed in some sort of a topology. Can you talk about the different types of topologies that you can support with this radio, and do you have anything else that is contributing to the creation of or traffic management in those topologies?

Dejan Bojic. This is probably where we should explain our tool kit, because we have spent a few minutes talking specifically about 60 GHz, and rightfully so, because it is a new technology and it is a significant piece of innovation. But it is certainly not the only technology in our solution.

We recognize that in some cases customer requirements may be challenging for 60 GHz-only approach, so for that reason we have added other solutions as well. These are addressing street-level connectivity and are complementary in nature to 60 GHz for let's call them "typical lamppost deployments", and include non-line-of-sight wireless and optical wireline technologies.

However, that's not the only part of this story, because every time we roll out tens of hundreds of small cells we will have the knock-on impact on the rest of the backhaul network in terms of capacity and traffic management end-to-end. We cannot just treat the small-cell backhaul domain in isolation. You can almost see this is an extension of the capillaries of the existing backhaul network. And considering that NEC has a presence in the existing backhaul, we are in a good position to address that network-level design with the expansion into the small-cell domain.

For that reason, in addition to these street-level connectivity solutions, we have also introduced a new E-band radio product to address increased capacity in the aggregation in urban areas.

If you imagine traffic from several small cells at the street level being collected by 60 GHz radios then at the aggregation point there is a need to transmit this traffic further into the network – that is really where high-capacity E-band radio comes to the fore.

In addition, NEC have introduced a very interesting modular concept for integration of optical connectivity into our radio unit so we can provide a seamless extension from fiber connections onto radio links in the last mile.

We also have some very interesting designs around the nodal functions, product called iPASOLINK GX, which has been available in the market for almost a year and is effectively a miniaturized, all-outdoor router device. It acts as a router, but it also acts as a power supply unit for radio units, so actually one small outdoor device can power up to four radios.

Effectively iPASOLINK GX gives you that ability to branch in a very space-constrained environment, whether that is a street or a rooftop. A device of this type is a very low-cost way to expand into ring topologies, into meshing, into partial meshing, or just simply to extend out from the existing aggregation point.

All of this is a way of looking at the network element layer, the links, connectivity and topologies. But then again, all this needs to gel, and we need to manage the traffic to optimize the capacity through this backhaul infrastructure. With that in mind the whole issue of end-to-end quality-of-service management comes into mind.

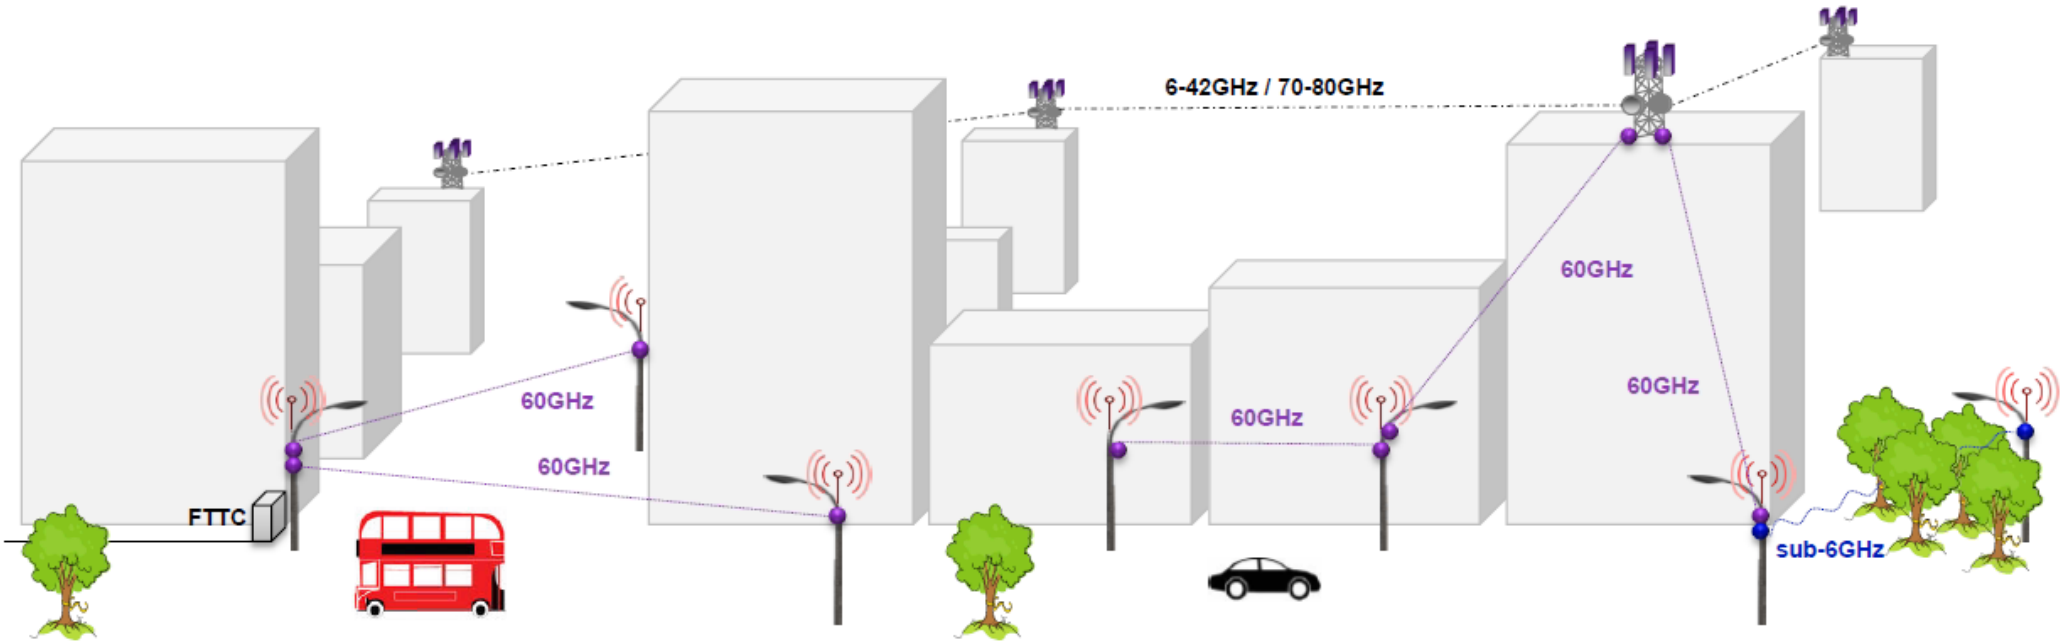

Figure 8. Technology toolkit and small-cell backhaul architecture. Source: NEC

At NEC we have put quite a lot of work into developing the network management solution that envelops all these technologies together, and we feel this is quite important, because what we are aiming to do is to give a common framework to our customers in terms of operation and maintenance of the network.

In addition to network management platform, we have also introduced tools for either static or dynamic optimization of capacity resources in backhaul. Sometimes we refer to these as backhaul resource manager applications which is something quite new. We are now in concept evaluation with several of our customers.

So that is really it: we build up our solution from that final few hundred yards to the small cells. We make sure that the knock-on effect on the existing backhaul is accommodated for and then we provide end-to-end common management framework for the complete backhaul. And in a sense that is what we feel gives us a good position to go to our customers and say, look, we are your one-stop shop for all your backhaul needs.

Lance Hiley. I would like to move the conversation to costs for a moment, because clearly the big issue with small cells is the fact that, one day, operators expect to have lots of them.

Clearly one of the focal points of a small-cell deployment is how much is it going to cost to connect each one of the small cells out to the rest of the network? What can you tell us about the installation and the cost parameters around your product?

Dejan Bojic. It is absolutely the question around the whole small-cell rollout. Cost is really what will determine the future successes of small cells and small-cell networks. There is little doubt that small cells are necessary in terms of overall network performance, but within the industry it still feels like the minutiae of TCO models and how will this all work out is very much up in the air. Our customers are grappling with this issue and that is why I think it's hard to answer your question in a very specific way.

What we know for a fact is that the cost of the equipment itself has to be minimized, as we have done at NEC, but that is only a fraction of the TCO. If we compare it to the past deployments of mobile backhaul for the macro-cellular environment, the cost of equipment itself is actually becoming less and less significant. It is all about costs of site acquisition, site engineering, installations, and commissioning.

One question you asked is about installation. Certainly we are aiming to minimize this to within an hour, perhaps a couple of hours per site. But there is a great dependency on which technology we use and what kind of environment we are working in.

When working at an aggregation site things are simpler with easier access. Working in busy streets presents its own challenges. This is why we are looking to do area-wide deployments in a single shot rather than site-by-site rollouts. For example we are working with our customers to deploy within a square mile at a time and basically optimize both people resources as well as engineering resources.

It is incredibly hard to pin down the numbers, and this is the focus of where we are with our customers. We are now trialing out all these technologies in the field and basically learning through experience how the economics of the installation and commissioning work. NEC is there with our customers to learn from trials and to tweak certain aspects of our equipment design in order to help them.

One example of this is automating of equipment provisioning. We are checking that have we put in the right steps in terms of our software design to make sure that when we mount the radios up we can download the configuration in a robust way. We are ensuring that the self-provisioning of equipment works as we are at the trial stages of real-life deployments.

Lance Hiley. If we were to look specifically at your 60 GHz product, the one that is going to be focused on the street-furniture type of deployments, what kind of estimates have you come up with, as far as installation times for the parameters involved?

Dejan Bojic. It is in the order of hours, I should say. I am struggling to give you a specific answer because it depends on the local installation and site procedures.

If the site is planned correctly, the power cabling is prepared and all the connectors are in place, and we do not have any hiccups, we can certainly mount the equipment, line up the link, and get the equipment commissioned and working within an hour, two hours at most. But in reality it is difficult, especially in the street environment. Sometimes we have to work on the fly. For example, you might arrive to the site and

realize that the power feed is not positioned correctly, and that will add time.

The real trick we see here is to create collaborative engagement with our customers. Not just with our customers' engineers but also with third-parties such as local installation companies. What we put a lot of focus on is planning and training, making sure that when we go into the field everything works smoothly. We are trying to focus on that pragmatic approach to real-life challenges as we look to grow small cell networks.

Lance Hiley. Well thank you, Dejan. Your comments and observations are going to be a great contribution to the Senza Fili report on small-cell backhaul.

Just a reminder, this report is aiming to provide an overview of all the different small-cell backhaul solutions available. And it is going to include an in-depth interview just like this one from all the leading vendors who have participated in the report. The report is available to be downloaded from the Senza Fili website, www.senzafiliconsulting.com.

Acronyms

FTTC	Fiber to the curb
RF	Radio frequency
TCO	Total cost of ownership

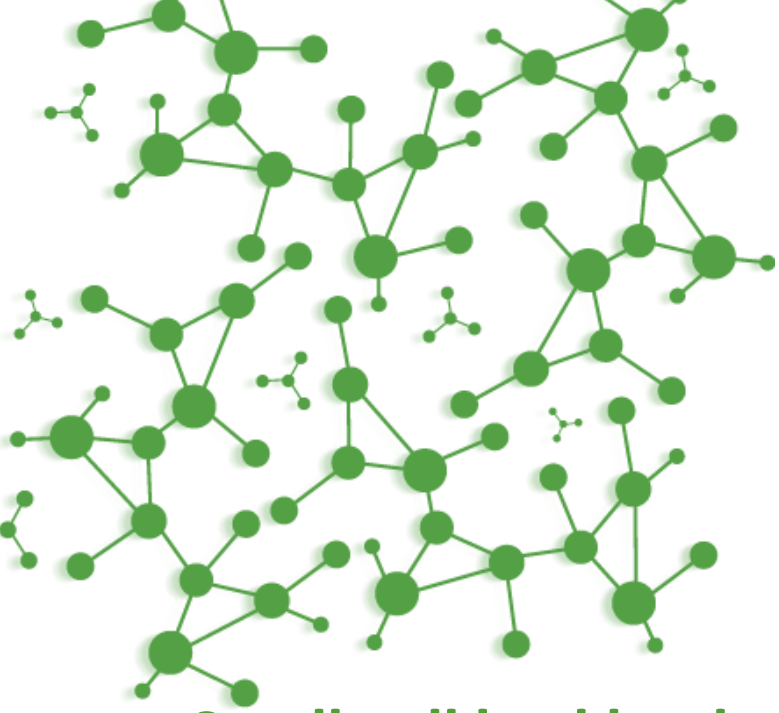

Small-cell backhaul with software-optimized commercial chipsets

A conversation with Lee Gopadze, CEO Proxim Wireless

By Frank Rayal
Senza Fili Consulting

Frank Rayal. Hello, and welcome to this conversation with **Proxim Wireless**. My name is Frank Rayal. This conversation is part of a Senza Fili report on small-cell backhaul that gives an update on small-cell backhaul solutions and on the evolution of mobile operator requirements for small-cell backhaul. Today we're talking with Lee Gopadze, CEO of Proxim Wireless. Proxim is a leading vendor of licensed- and unlicensed-spectrum backhaul systems used for video, mobile, backhaul and network applications.

Lee, tell us about the background of the company and its heritage.

Lee Gopadze. Proxim Wireless has been in the manufacturing of wireless communications equipment for about 25 years. We have, during that period of time, developed a cross-layer software capability called WORP, which stands for Wireless Outdoor Router Protocol, and to date have placed more than 2 million units into the marketplace with about 250,000 different companies around the world. We operate primarily out of our headquarters location in Milpitas, California, with an R&D center in Hyderabad, India, an operations center in Taipei, Taiwan, and sales offices around the world.

Frank Rayal. When it comes to small-cell backhauls, you have a new product. That's the Tsunami MP-8200 platform that you're targeting specifically for that market in addition to other applications. Tell us more about this product and how it meets the requirements of small-cell backhaul.

Lee Gopadze. First of all, let's talk a little bit about some of the things that we believe small-cell backhaul will need, and often these are competing requirements. Low equipment cost, low operational cost, guaranteed throughput, high capacity in terms of the number of units that can be serviced, deterministic service flows, RF signal propagation capability and the very physical size of the unit we think are all important characteristics of small-cell backhaul units, and ones that we intend to address through our 8200 platform and units that will come out about midyear of this year.

As we look at the 8200, it is a small-cell backhaul product that can be used both in a point-to-point and in a point-to-multipoint mode.

In a point-to-point mode, the unit will approach 250 mbps of throughput using the highest modulation capabilities.

In a point-to-multipoint mode the unit will deliver as much as 150 mbps to each one of the nodes to which a base station is attached. Coupled with the MAC-polling protocol that we use on the product and the class-of-service and quality-of-service characteristics of the product, and the cost, we think it is an able competitor for people who are interested in a small-cell backhaul to look at, either in the 5 GHz or in the 6 GHz band.

Frank Rayal. So this product mainly operates in the unlicensed-spectrum bands?

Lee Gopadze. It operates in the United States in the unlicensed-spectrum bands. In other areas of the world where both 5 and 6 GHz are licensed spectrum, the product operates within those same bands but they are licensed. So, for example, in Russia the product works at between 5.9 GHz and 6.4 GHz, which is licensed spectrum in Russia.

Frank Rayal. Lee, you mentioned software features that your company wraps under the headline of WORP. Can you tell us more about that?

Lee Gopadze. WORP is a cross-layer software product or protocol that we have developed over about the last 10 years. One way to understand it is to look at LTE and WiMAX protocols. What you would find is that WORP compares favorably in a throughput, quality-of-service, and class-of-service comparison with those particular protocols. What WORP has that is perhaps unique is its ability to deliver low latency, low jitter in a very flat network architecture.

Our point-to-point links typically operate at less than 3 ms latency; our point-to-multipoint links operate typically at less than 10 ms latency. Both of those are significantly below what you would find in a standard WiMAX offering or even in some of the earlier versions of LTE. We developed the WORP protocol primarily because we service the video surveillance market, where both latency and pixelization are very important to the agencies, many of whom are government agencies.

Frank Rayal. You have also mentioned flexible channel planning in your product. How does that work, and what does it do?

Lee Gopadze. If you look at most Wi-Fi products, they are available in 20 and 40 MHz channel sizes. Our channel sizes actually go down to 5 MHz, which allows the operator – be they a carrier or for other backhaul needs, such as video – to use small bits of spectrum that are available.

Outdoors is a very different environment than indoors, so much of what we do today is focused on allowing operators, both network operators as well as agencies, to use small bits of spectrum and still maintain very high throughput. In the outdoor world, finding a 20 MHz or a 40 MHz channel is often very difficult in urban areas, so we combine small channels with high throughput in order to deliver what we believe the market needs in the outdoor and urban areas.

Frank Rayal. When it comes to comparing your solution with other solutions that are in the same class, what are some of the differentiating features? You talked about WORP; any other things that you can share with us?

Lee Gopadze. I think it goes back to our ability to control the basic chipsets. We use what you might call a COTS, or a commercial off-the-shelf, chipset. Typically, that chipset will normally be sourced through someone like Atheros or one of the other major chipset vendors. When you look at those chipsets, typically what you find in standard 802.11n Wi-Fi where there is a relatively low utilization of available data rate versus throughput, in the 50% to 60% range.

What WORP does is actually increase the efficiency of the throughput. We regularly see WORP providing as much as 80% or 85% of the available data rate in relation to the channel size. If you have a 300 mbps data rate, you will see WORP ensuring that the throughput is as much as 250 mbps.

The other advantage is that WORP allows us to increase both the transmit power and the receive sensitivity, as well as expand the range of the product.

For example, most standard Wi-Fi chipsets work from about 5.15 to about 5.825 GHz. We are actually able to stretch that, using WORP, down into the 4.9 GHz range and as high as 6.4 GHz. We start out with a Wi-Fi chipset and end up with a software-defined radio with performance-monitoring asymmetric bandwidth control, where drops in packets due to interference are detected by WORP and then addressed by either retransmitting the packet or increasing the available signal strength to make sure that the link is secure and stable over time.

Figure 9 Tsunami MP-8200 Wireless Backhaul System.
Source: Proxim

Frank Rayal. You mentioned dropping packets, and that can be all due to interference. Using backhaul in unlicensed band has been a point of debate in the wireless community. Some operators see no problem with that; some of them are less amenable to deploying an unlicensed-band solution.

What is your take on that, and how do you strengthen the communication channel so that there is less interference or a more robust communication link?

Lee Gopadze. I think you have to look at it in terms of distance, throughput, and channel or packet stability. As you look at a particular environment, you want to maximize the modulation that you are able to push, to 64 QAM, or 16 QAM. Then you want to have a relatively short distance so that the signal strength is maintained relatively high, and then you want to control the latency and the other performance characteristics of the signal.

By using WORP, we increased the throughput, number one. Number two, by using a polling architecture rather than standard Wi-Fi, we created an environment very similar to what you find in a WiMAX product, where each station speaks when spoken to. Then by overlaying this with a high level of CoS and QoS, we are able to maintain a high-performance real-time environment.

I think a measure of this reliability is the fact that we are used by multiple law-enforcement agencies around the world to deliver video packets over secure links at secure installations. Video is a very demanding application, particularly high-definition video. We see pixelization happening very often. With our particular product – it is one of the reasons that the County of Los Angeles has bought it, among other customers – you do not get that pixelization. Our product is designed today to typically work out as far as five miles.

As we look at small-cell backhaul, we are seeing distances dramatically shorter than that, maybe as little as 300 to 500 meters. So we're very comfortable, within those parameters, that most of the time we'd be using 64 QAM or 16 QAM, which gives us a high available data rate and, therefore, addresses the need for significant throughput. The cost factor weighs in, as well.

Frank Rayal. When it comes to the business case of small-cell deployments, cost has been a major issue that prevents operators from going through this type of deployment. How does your product address the cost issue?

Lee Gopadze. It addresses it from a couple of standpoints. First of all, the product is built on a commercial off-the-shelf chipset, which means that we start out with a lower overall cost in general. We are not building specialized chipsets in order to address the need, and therefore, the cost of the internal electronics starts to come down dramatically as well.

The second way that we address cost is in total cost. When you look at many products on the market, they are using plastic housing and a variety of low-cost features. We do not think that is the way to do it. We have significant experience in building IP66 and IP67 enclosures, which will withstand the test of time, and that is why you see MTBF of our products is upwards of 200,000 hours. The last thing is the enterprise market, which is our primary market and is a very cost-conscious market around the world. We sell in India, China, Russia, as well as the United States, and carriers in all of those markets are very cost conscious. If you look at point-to-point microwave products, they often cost $5,000 to $15,000 per link. Typically, our products are less than $3,000 per link and will deliver as much as 240 to 250 mbps for a single location.

Frank Rayal. I am very curious about the spectrum issue, because you operate in unlicensed spectrum in many countries, and licensed in others. There have recently been moves by the FCC to provide more bandwidth in the 5 GHz band for Wi-Fi and other

unlicensed applications. You mentioned that there are some new rules in the 6 GHz band. Can you elaborate on that?

Lee Gopadze. There is a rulemaking before the FCC today to allocate about 200 more MHz in the 5 GHz spectrum band here in the United States. That obviously portends well for using 5 GHz, potentially, as an outdoor product for small-cell backhaul. At the same time, the 6 GHz spectrum here in the United States has typically been reserved for long-distance microwave products, where you have large dishes and split-mount installations with both an IDU and ODU.

The FCC is in the process of finalizing rules that will allow all-in-one enclosures so it is an all-outdoor unit. The minimum size of the antenna has come down to about 1 meter now, right around 3 feet, and they are even looking at potentially allowing point-to-multipoint operations. They have not created a rulemaking for that, but it is a subject of discussion by the technology group at the FCC.

We see all of these things as positive indicators that the 5 GHz band will continue to grow in terms of the spectrum that is made available to it, and that there is the potential to use the 6 GHz band in a very efficient manner outdoors, as well, for small-cell deployment.

Frank Rayal. Do you have any plans for developing licensed-band solutions? What do you think about having backhaul in the licensed bands?

Lee Gopadze. We think there is certainly an opportunity for that. We have products in the 18 to 23 GHz bands today. They are split-mount products in that they have an ODU and IDU. We do not think that that is an appropriate architecture for small-cell backhaul. We are actively looking at spectrum both in 3.5 GHz, and 28 to 31 GHz where, here in the United States, there is the ability to use that product.

Again, the very architecture of WORP means that, as long as we have a hardware platform, it is relatively easy for us to port WORP over to that hardware platform – for example, at 3.5 GHz – and to provide the same set of operating characteristics that we would at 5 or 6 GHz.

WORP is a cross-spectrum software capability. Our job becomes much easier: instead of developing both software and hardware for a particular spectrum band, all we are doing is looking at implementing a new chipset on the PCBA that we already have in place.

Frank Rayal. Looking at the evolution of your product, how do you see that going forward? What are the next features and plans for your product line?

Lee Gopadze. Later this year we will see a smaller-sized platform coming out. Frank, I think you mentioned that you had seen the 8160 on our website. That size of platform is where we are driving a lot of our product development towards. We believe that is absolutely critical for small cells.

The other area that we are driving is our software product. Today our software operates in mobile environments up to about 70 mph. We think it is important to increase that, so we are currently almost finished with a new revision of software that will drive that up to about 180 mph. This focuses on one of our primary market segments, which is intelligent transportation systems and high-speed trains and trams that are being constructed all over the world.

Frank Rayal. Any final thoughts on the small-cell market and the positioning of the Tsunami product line in that market?

Lee Gopadze. We are doing a couple of things in that area. We will actually be presenting to IWPC next month in their small-cell backhaul forum, which will address some of the small-cell backhaul requirements and how, from a performance standpoint, we think that we fit in well with the requirements of small-cell backhaul. We are also looking at 3.5 GHz, which is another spectrum area that we think has some value for small-cell backhaul.

So, a combination of expanding the spectrum bands that we operate in, reducing the size of the product, and ensuring that the existing class-of-service and quality-of-service characteristics are maintained. We think this gives us a protocol that is RF-spectrum agnostic and meets many of the needs of small-cell backhaul, whether you are here in the United States or in other parts of the world. Then, at some point in time, we probably need to be talking with some of the vendors who are providing the front-end, the radio-access network and the small cells, to see if incorporating our product with theirs makes sense over time.

Frank Rayal. There is a lot of work to be done in that space, and that's the good thing here. Lee, I would like to thank you very much for joining us in this conversation.

Lee Gopadze. I appreciate it Frank.

Frank Rayal. This conversation is part of Senza Fili's report on small-cell backhaul, which provides an overview of small-cell backhaul solutions along with in-depth conversation like this one from leading vendors who participated in the report. The report can be downloaded from the Senza Fili website at **www.senza-fili.com**.

Acronyms

CoS	Class of service
COTS	Commercial off-the-shelf
FCC	US Federal Communications Commission
IDU	Indoor unit
IP66	IP Code, Ingress Protection Rating: dust tight, powerful water jets
IP67	IP Code, Ingress Protection Rating: dust tight, immersion up to 1 m
IWPC	International Wireless Industry Consortium
LTE	Long term evolution
MAC	Medium Access Control [Layer]
MTBF	Mean time before failure
ODU	Outdoor unit
PCBA	Printed circuit board assembly
QAM	Quadrature amplitude modulation
QoS	Quality of service
R&D	Research and development
RF	Radio frequency
WORP	Wireless Outdoor Router Protocol

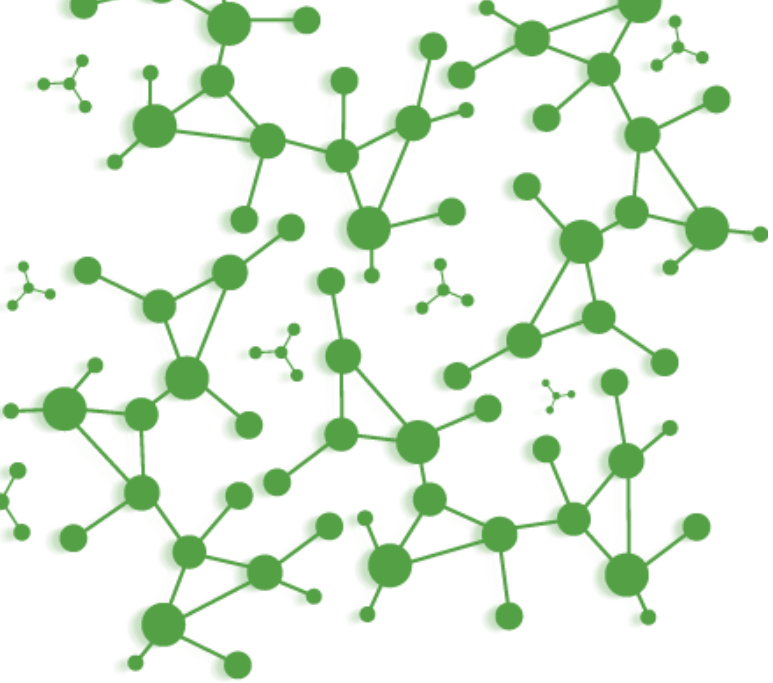

Siklu's all-silicon technology for small-cell backhaul

A conversation with Shahar Peleg, Director of Product Management, Siklu

By Lance Hiley
Senza Fili Consulting

Lance Hiley. Welcome to our conversation with Siklu. This conversation is part of the Senza Fili report on small-cell backhaul that looks at the current state of small-cell backhaul solutions available from different vendors and the evolution of mobile operators' requirements for small-cell backhaul. My name is Lance Hiley, and I am here on behalf of Senza Fili.

Today we're talking to Shahar Peleg, Director of Product Management at Siklu, a company that is pioneering cost-effective, high-capacity, point-to-point microwave wireless backhaul solutions operating in the E and V bands that are particularly well suited for small-cell-deployments. Shahar, how would you say your product fits into the small-cell backhaul marketplace?

Shahar Peleg. Siklu makes gigabit wireless backhaul solutions that operate in the 60, 70 and 80 GHz millimeter bands, which are also called, as you mentioned, the V band and the E band. We develop advanced, all-silicon designs that reduce the price and increase the reliability of these products. We have sold over 7,000 units to date, making us the leader in millimeter-wave backhaul.

Our systems are field-proven in all sorts of weather conditions around the world, from the monsoons of India to the cold weather of Russia. And our units also operated solidly during the recent Hurricane Sandy. Millimeter-wave E band and V band will play a significant role in small-cell backhaul, providing high-capacity backhaul solutions for small-cells and the microcell network.

Lance Hiley. Weather and environment may not affect your product, but what part of the small-cell backhaul market are you targeting? Would you say that you have a solution that's well suited for rural or for metro?

Shahar Peleg. I think the most interesting thing that we're doing in the small-cell backhaul ecosystem is that we have adapted our all-silicon technology for street-level urban deployment. We've developed a point-to-point gigabit wireless radio that operates in the 60 GHz band just for that purpose. Small-cell backhaul is going to be completely different than the existing mobile backhaul network. Small cells are going to be deployed at the street level, and the backhaul will have to be as well.

Deployment at the street level has to be extremely fast and easy. Physical installation needs to be quick because of the limited access that the installer will have to the street furniture where the backhaul will be installed.

Also, because of the sheer number of small-cells that will need to be deployed, the installation needs to be simple enough to be done by unskilled installers, while

the commissioning needs to be done automatically from a centralized location. All this and more makes the focus on the small-cell backhaul challenge very interesting; creativity plays an important role.

Lance Hiley. What would you say is your technological advantage over other solutions and in the line-of-sight wireless backhaul category?

Shahar Peleg. First of all, we think that the bulk of small-cell backhaul will be done in the 60 GHz V band, because 60 GHz offers an abundant chunk of spectrum, which supports high-gigabit capacities that will be required for small-cells. On top of that, the 60 GHz frequency band has unique propagation characteristics that support the ideal distances required for small-cell backhaul.

In addition to that, it is largely immune to interference, which is especially important for deployment in street canyons. Operators also get high spectrum reuse and the band is unlicensed, which allows even greater freedom when coming to choose the location for the installation of the backhaul. It gives you flexibility in choosing the location, even when on site, because it does not require a long, time-consuming, bureaucratic link-licensing coordination process.

From the vendor perspective, it allows us to design equipment – specifically an antenna – which is very small, enabling the simple installation we just mentioned, but which also copes with the street-level challenges such as pole sway – the tilt or movement of the street furniture itself.

Lance Hiley. One of the main reasons operators look to line–of-sight and specifically point-to-point line-of-sight technology when they're designing their backhaul networks and backhaul strategy is that they are looking for capacity. Traditionally point-to-point has delivered more in the way of capacity, more in the way of performance. How does your product fit into that situation against the other point-to-point solutions that are operating, perhaps, at lower frequencies?

Shahar Peleg. First of all, let me mention that we that we don't expect one single technology will be used for small-cell backhaul. It will always be a mix. Where fiber exists and is available, fiber will be used. 60 GHz, specifically our equipment, provides very high gigabit-per-second capacities and could be expanded to multi-gigabit-per-second capacities. In street canyon small-cell deployments, it will be very hard for the operators to stick to the traditional PTP technology, as it will be very hard to have line of sight to each and every small-cell from a centralized access point. This is why we expect small-cell backhaul networks will use daisy chain or ring topologies, where the capacities will need to be aggregated over a number of small cells.

Even though each cell site may only generate 100 mbps, once you aggregate a number of small cells, you very quickly reach multiple hundreds of megabits and even gigabit capacities. Supporting this magnitude of capacity using non-line-of-sight technologies is not achievable.

For the other traditional microwave frequencies, a few hundreds of megabits are achievable, but the form factor doesn't allow it to be deployed in small-cell scenarios. The traditional microwave frequencies are licensed frequencies, where there are requirements and standards dictating the size of the antenna and such, which limit the size of the equipment and make it unsuitable for small-cell deployment. This is where 60 GHz or V band has an advantage. It supports high-capacity, multi-gigabit traffic – and allows the vendors to design practically invisible backhaul systems that can be installed at street level.

Lance Hiley. You talked a little bit about the different deployment scenarios, because traditionally, as you said, point-to-point systems are deployed in tree-and-branch, daisy chain, ring or mesh configurations. Does the Siklu product have any other features that are integrated into it or as a standalone subsystem that support any or all of those different deployment scenarios?

Shahar Peleg. Siklu's product may be small in size, but it is packed with feature sets and capabilities. We have integrated a full suite of networking capabilities into the unit itself that allows you to deploy point-to-point links in daisy chain, in mesh, or in rings without the need of any additional demarcation equipment. We have also integrated advanced zero-touch installation capabilities based on SON principles to simplify the installation process. Of course, it pretty much goes without saying that you need to provide synchronization capabilities as well, so we have also integrated both SyncE and IEEE 1588. So not only do we provide gigabit-per-second capacities in this small form factor, we also provide networking capabilities and synchronization capabilities, all in a single box. And with total cost of ownership one of the major

Figure 10. Siklu EtherHaul PTP radio mounted on a street lamp. Source: Siklu

enablers for the mass deployment of small cells, we do all this at a very, very low cost.

Lance Hiley. We'll come back to costs in a moment. But I'm curious: what are some of your typical installation parameters? For example, how long does it take to install your product, and which features in the product assist with the installation process?

Shahar Peleg. When we designed our EtherHaul 600, which is our small-cell backhaul product, we had the installer in mind. And when I say the installer, I mean a non-telco installer.

First of all, our unit can be mounted on a variety of poles. It can also be installed on the side of a building. All this is done with a single mounting mechanism. The installer doesn't need to think too much about what he needs to take with him. Our mounting mechanism gives him the flexibility to decide on site where to install the unit.

One of the main challenges of 60 GHz V band being a line-of-sight technology is aligning the link. Once the installer installs the hardware, the units need to be aligned to maximize the reception at both ends. We understand that due the number of small-cells that will need to be installed, the installers may be people with little experience and expertise in the installation of telco equipment. To help, we are providing an auto-alignment tool that takes the effort out of the alignment process.

The installer installs the tool on each end of the link, and then they can go and drink their coffee. When they return the alignment will be done automatically for them.

By limiting the expertise required of the installer, you limit the human error. Our idea is that once the automatic tool does its work, that's it. You take the tool with you and you don't need to visit the site anymore. From there on in, our zero-touch system allows you to do all the commissioning and the integration of the link into the network from a centralized location by a person who has more experience and a higher level of expertise.

Our experience from trials that we've done, including those with leading mobile operators, shows that installation can be done in 10 to 20 minutes. You can have a link up and running, and continue on to the next link.

Lance Hiley. Cost is something that is really at the forefront of operators when it comes to small cells. They are very concerned with having so many operating nodes in the field and what it is going to take to actually deploy those. Can you give us a sense of why your solution may be more cost effective than others that are being put forward in the industry? And if relevant, what environment is it better suited to, from a cost perspective? I mean you mentioned that no one solution fits all, so clearly you must have a particular scenario or environments in mind that you're focusing your product on from a cost perspective.

Shahar Peleg. The main innovation of Siklu is that we've invested in designing and developing all-silicon technology. In a sense, we've taken all the silicon technologies and disciplines from consumer electronics and applied them to the carrier network, retaining the cost structure of consumer electronics on one hand while ensuring the carrier-rate performance that the operators require.

Obviously, in consumer electronics, cost is one of the key elements. We've developed a single radio-frequency integrated chip that integrates all the LNAs, the power amplifiers, the front-end low-noise amplifiers, the mixers, the filters – everything that was a discrete module – onto a single silicon chip, which reduces the cost.

We've also developed our own baseband chip that incorporates the modem, the D-A converters and a powerful machine that controls the quality of service over the link. Not only does this enable us to achieve optimal radio performance, it also eliminates our need for a third-party modem, which allows us to reduce the cost even more.

Figure 11. Siklu EtherHaul palm-sized outdoor unit. Source: Siklu

Another aspect is that we have designed – and we have a patent on – antennas that are not only very small in size, but that also comply to with all the FCC and ETSI standard requirements and are significantly lower cost compared to off-the-shelf antennas that we could have bought.

All these technologies put together result in a very low-cost product, but they also have some nice side effects which result in a TCO reduction. Products are more reliable, so you don't need to visit the sites for repairs and maintain a large stock for repairs. The power consumption is very, very low, so your power costs are greatly reduced. The small form factor also contributes to reducing your rental costs in many cases.

Of course, being an all-outdoor product, there are no indoor-associated costs, with space rental for indoor equipment being one of them and air conditioning being another. Operators are looking very closely at the TCO, and with backhaul being a large part of the ecosystem, EtherHaul is specifically designed with this in mind.

Lance Hiley. Okay, Shahar, so you have told us that your product is easier to install than many expect, and that is going to help with the cost, and you have explained to us how you are trying to drive down the manufacturing costs of the product. So what is the bottom line? Where is the cost per link of E-band and V-band products converging to, and is it comparable to what we have seen for traditional Ethernet gigabit point-to-point links?

Shahar Peleg. Siklu EtherHaul-600, which is a 60 GHz V-band product, costs in the area of $2,000 for a complete link, meaning two sides of the link. We do expect a significant price reduction as the small-cell market starts to take off. Once it does, the EtherHaul product may go down to as low as $1,000 per link.

Lance Hiley. So basically what you're saying is that the days of the $20,000 per link E-band product are long gone and this is a product that can really compete with all of the other conventional line-of-sight and non-line-of-sight small-cell solutions that are being put forward by the different manufacturers?

Shahar Peleg. Yes, the days of the $20,000 E-band link are long gone. When Siklu came into the market with our E-band products, we released it at a price of approximately $5,000 a link. Once again, this is all a result of our house silicon and self-designed technology and manufacturing capabilities.

Today we see other vendors coming down to Siklu's price range. But Siklu is reducing its price to the price of a traditional microwave link and even below that, yet providing the high capacity and high performance of the E and the V band.

Lance Hiley. That's great. Well, thank you very much, Shahar. It has been fantastic to hear your comments and observations. It is going to be a great contribution to the Senza Fili report. The report aims to provide an overview of the different small-cell backhaul solutions available. The report can be downloaded from the Senza Fili website, www.senzafiliconsulting.com.

Acronyms

D-A	Digital to analogue
ETSI	European Telecommunications Standards Institute
FCC	U.S. Federal Communications Commission
IEEE 1588	Standard for precision time protocol
LNA	Low-noise amplifier
PTP	Point to point
SON	Self-organizing network
SyncE	Synchronous Ethernet
TCO	Total cost of ownership

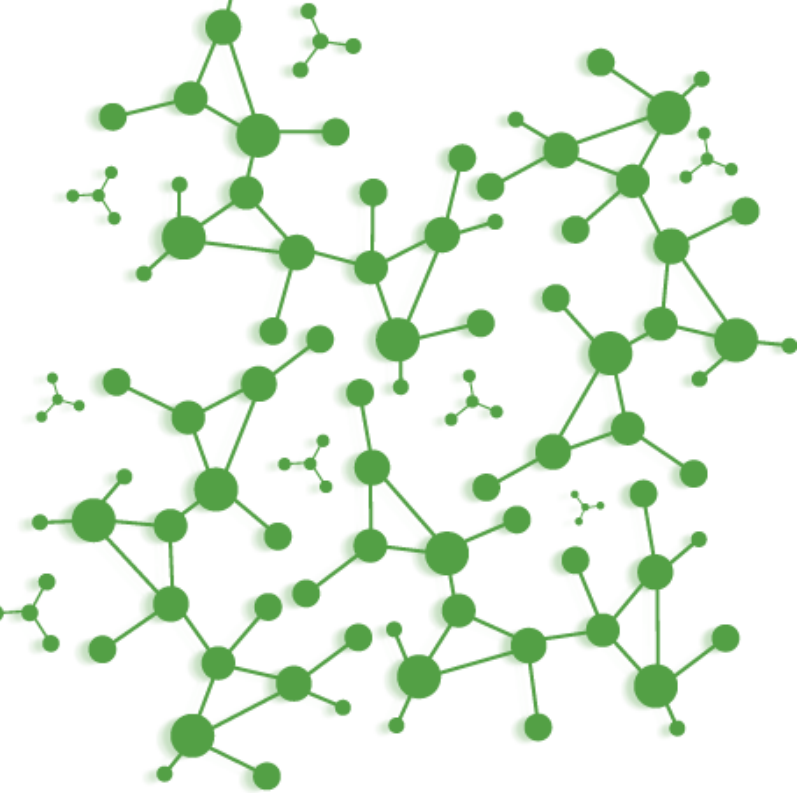

Expanding the utility of fiber in small-cell deployments

A conversation with Saeed Anwar, CTO SOLiD Technologies

By Frank Rayal
Senza Fili Consulting

Frank Rayal. Hello and welcome to this conversation with Saeed Anwar of **SOLiD Technologies**. My name is Frank Rayal. This conversation is part of a Senza Fili report on small-cell backhaul that gives an update on the small-cell backhaul solutions and on the evolution of mobile operator requirements for small-cell backhaul. Today we are talking with Saeed Anwar, CTO of SOLiD Technologies. SOLiD is a leading provider of advanced distributed antenna systems, commonly known as DAS. Saeed, welcome to this conversation. Thanks for being with us today.

Saeed Anwar. Good morning. Thank you.

Frank Rayal. I would like to start by asking you for a brief introduction about SOLiD Technologies.

Saeed Anwar. SOLiD is from South Korea. Our R&D and manufacturing is done in Seoul. The US operation is sales, marketing and logistics for distribution of the product into the North American markets. We are primarily selling DAS, both indoor and outdoor DAS, in North America, and we also have fiber optics for fronthaul and backhaul.

Frank Rayal. So the product that you are providing that is specifically targeted at the small-cell solutions, I believe is called Infinity Access.

Saeed Anwar. Yes.

Frank Rayal. Help us position this product. Tell us what it does and some of its features.

Saeed Anwar. The product is a DWDM (dense WDM) technology product. It has the capability of channelizing a single fiber strand into multiple bidirectional channels. The technology itself can handle 32 bidirectional channels. The form factor that we use allows us to do 16 bidirectional channels. With a single fiber strand from the hub site, where all the heavy equipment is, we can take the signaling and transmission over a distance and do 16 linear add/drops off that single fiber to wherever the small cells are, whether they are on a wall or lamppost or street furniture.

Frank Rayal. You mentioned fronthaul. Fronthaul is a relatively new name in the industry. Can you tell us more about what fronthaul is and how you define it?

Saeed Anwar. Let's start off with backhaul first. Backhaul is basically from the base station. The typical protocol that you use is Ethernet, linked back into the port.

For fronthaul, the base station itself is split in two parts. There is the baseband processing and there is the antenna system itself, which we call the remote radio head. The link between those is what is called fronthaul. It is traditionally designed for macro deployment, where you have the baseband unit at the

base of the cell tower and the remote radio heads at the top. It is a fiber-based system using a protocol called CPRI, and from there people have learned that you can actually take that remote radio head and distribute wider. Some people are deploying it over distances of tens of kilometers, and that is what we refer to as fronthaul in the industry.

Frank Rayal. So how about the applications of the Infinity Access, can you describe a few deployments in areas of how the system can be used?

Saeed Anwar. The small-cell deployment is one scenario. Using a single strand of fiber, you can do 16 linear add/drops on that fiber to street furniture, where the small cells are. That linear add/drop can simultaneously handle both backhaul and fronthaul.

Other deployments on where we use Infinity Access are in what I consider regular telecom and enterprise environments, where you have fiber capacity shortage and you want fast relief without retrenching or bringing more fiber. It is an easy way to add capacity on that fiber run.

Frank Rayal. So in terms of the deployment of your system, let's say I have a base station and I want to cover a certain area, whether it is indoor or outdoor. What are the interfaces, and how do you deploy your solution?

Saeed Anwar. Well, the protocol for that particular scenario, as I said, CPRI. There is another version called OBSAI. The scenario for that is really, at the current time, more outdoors.

Macro cells are split into baseband and remote radio heads. And we are seeing is a new form of micro cell: The one-watt systems that are considered small cell and also have the same fronthaul interface going back into a data center. So those two are the two primary scenarios that we see.

Frank Rayal. Where do you see most of the deployments of your system, the Infinity Access, and what type of environments do you see the operators deploying now?

Saeed Anwar. We see that in metro areas, where micro cells provide coverage in a concentrated area.

From a macro perspective, we also see a deployment scenario where you are deploying macro sites, but for efficiency you are trying to take the baseband processing and centralizing it into a small base-station hotel.

Frank Rayal. When it comes to small-cell backhaul, wireless has been widely promoted because the fiber infrastructure can be not as extensive as to reach every single point where the small cell is going to be deployed.

With your solution, do you see that being deployed mainly in indoor environments or outdoor? And how do you make the business case for it in the outdoor case, where the availability of fiber can be an issue?

Saeed Anwar. Certainly for indoor deployments we can use an Infinity Access.

For outdoors, in metro areas, we see a lot of fiber. Outside of the metro areas, yes, fiber is questionable. Because of the way Infinity Access is designed, it is a very efficient use of the existing fiber, so we can certainly use Infinity Access in a metro environment where the fiber does not exist, in combination with a wireless solution. For example, if you want a small cell on the one side of a street but the fiber is on the other side, then we would use wireless technology to bridge the gap.

Frank Rayal. You mentioned you can add or you can put 16 channels on one fiber strand. Is that the capacity of the card, or is it the hub capacity?

Saeed Anwar. That is the hub capacity.

Frank Rayal. And the 16 channels can be anything ranging from CPRI to Wi-Fi, Ethernet or any other type of channel, is that correct?

Figure 12. Infinity Access Optical Line Terminal (top) and Optical Network Terminal (bottom). Source: SOLiD Technologies.

Saeed Anwar. Yes. The primary channel protocols that we use on fiber are CPRI and Ethernet.

Frank Rayal. What are some of the other main features of your solution, and how do you compare your DAS solution to others?

Saeed Anwar. I think there are several factors. The biggest factor is the use of a single fiber strand for multiple drops. Second, within the Infinity Access, we have linear drops, which provide the capability for passive bypass of the site for maintenance or power outage.

So imagine if you have 16 drops down a street and one of your cell sites there is down because power to the building or lamppost is out. The other sites are totally unaffected because of the passive bypass.

Frank Rayal. When you look at DAS systems, what are some of the key differentiating parameters? Where does your solution differentiate on these parameters?

Saeed Anwar. From a DAS perspective, the parameters you are looking for are, first, the distribution coverage from a single point to a multipoint environment: What is the capability of the system to provide the widest coverage per RF? Efficient fiber usage is a second one. The ability to multiplex both cellular signals and public safety at the same time is a third one.

And then obviously there are several RF KPI metrics that you would have to consider. Another one that is significant is how you roll individual cell providers onto the system, and that is basically what we call our power-level settings for existing carriers after you bring on new carriers. This is a form of protection for existing carriers there.

Frank Rayal. There seems a debate on what is a small cell, and many people also consider remote radio heads as being a small cell, which in the pure definition of LTE-Advanced, it is. Where do you see operators deploying remote radio heads versus compact base stations, and how do you see the backhaul to these two things playing?

Saeed Anwar. Well, there are different scenarios. Obviously SOLiD has a very extensive remote-radio head deployment within South Korea, but in the US what we are seeing is remote radio heads first for hotspots, and others in areas with rapid increase in population like North and South Dakota.

The basic reason for deploying remote radio heads is an easier and a faster way to deploy RF without having the elaborate balancing issues you have with compact base stations. With compact base stations, you still need to make sure that all the interference issues are cleared up. With remote radio heads, you can deploy a lot of remote radio heads, especially if they are in the same baseband, without any interference, so it is an easier technology to deploy from that perspective.

Frank Rayal. In terms of the economics, do remote radio heads have an advantage, in certain types of deployment, over the compact base stations? How does the DAS system itself or the backhaul over fiber play in that type of a scenario?

Saeed Anwar. From an economics perspective, compact base stations would be prevalent from a capex perspective, but, from an opex perspective, the remote radio heads can be easier to manage, even more than compact base stations, primarily because of all the analysis that you would have to do for RF interference issues. For the backhaul, I think the compact base stations would be easier because you have multiple tools in your tool bag for compact base stations because you are carrying Ethernet. But for remote radio heads you will definitely need fiber.

Frank Rayal. I think that is one of the major issues. When it comes to remote radio heads, we are talking about gigabit per second in terms of throughputs, and when we talk about compact base stations, it is tens or hundreds of megabits per second.

In terms of evolution and roadmap of your product, where do you see this product heading?

Saeed Anwar. We certainly believe that Infinity Access is going to be deployed for small cells in high-density metro areas where there are existing fiber providers or a substantial fiber plant. We also plan on evolving the technology for our next-generation digital DAS systems that will combine the ability of carrying compact base stations and remote radio heads into what we call multiband, multiprotocol systems.

Frank Rayal. Can you detail for us the difference between the analog DAS and digital DAS?

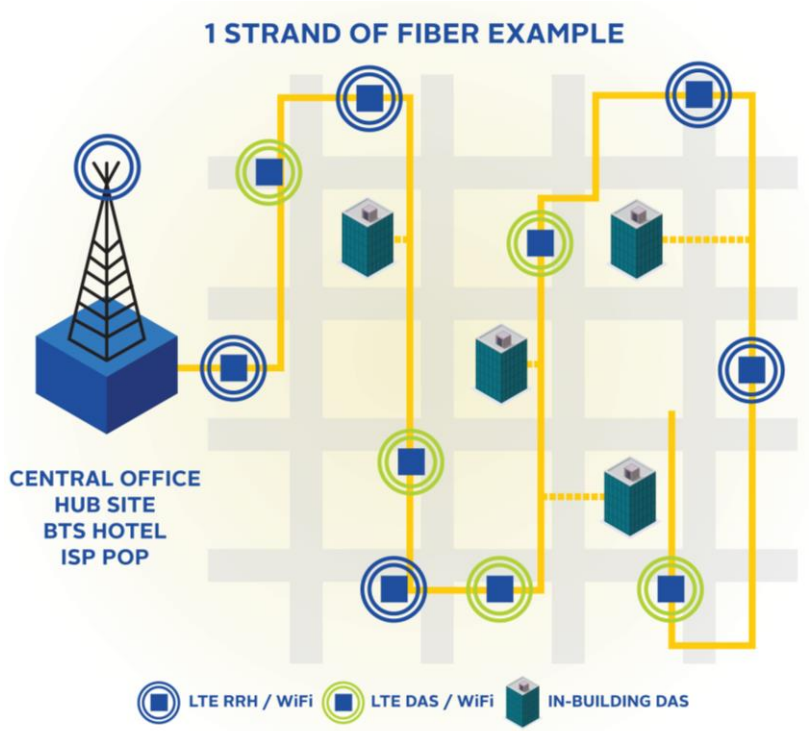

Figure 13. Infinity Access provides up to 16 multi-band, multi-technology add/drop capability over a single strand of fiber. Source: SOLiD Technologies.

Saeed Anwar. In the analog DAS system, we take all the band and protocol feeds from the base station, and combine them in our RF mixer. Then, we translate that RF from coaxial cable to fiber, we place that particular feed, once it hits the fiber, onto a WDM frequency, and finally carry that frequency over a distance. It is very similar to a cable plant. When you are receiving TV signals in your home from the hub site, you are taking fiber first, and then eventually you take coax right in front of your house.

The DAS fiber system is very similar. At the remote node, we convert the fiber to coax, which then goes into the antenna.

On the other hand, in digital DAS systems a remote radio headend is used which is connected to the baseband module through fiber. The interface between the RRH and baseband can be based on CPRI or OBSAI standards.

Frank Rayal. In terms of your view on the small-cell market, whether it is RRH or a compact base station, how do you see that evolving, and where do you see some of the main deployments happening right now?

Saeed Anwar. Obviously in heavier-density areas, the distance between cell sites is much smaller than in rural and suburban areas. With the deployment of LTE-Advanced, we see the distance is shrinking further. The primary reason for that is, if you consider cell sites today, that they do not talk to each other, they are talking back to core. In LTE-Advanced they will be able to talk to each other.

When two cell sites are talking to each other and the system is designed so that the connection goes back to the core and then comes back again, there are going to be substantial latency issues. This is where we see the advantage of our technology, because it is ultra-low latency. From an end-to-end connection basis, our latency, not including the fiber, would be around 120 nanoseconds. Our jitter is also extremely low, 6 nanoseconds. When you total that up, we provide a very good environment for LTE-Advanced compared to other technologies.

Frank Rayal. A low latency and jitter are important for some LTE-Advanced features, like coordinated multipoint, which requires very close synchronization between the different base stations or the transmitters.

Any other thoughts that you would like to add in terms of your product line and where you think you are heading with the solution?

Saeed Anwar. Small cells include Wi-Fi also. We are using the same technology to provide Wi-Fi within the same environment as DAS. For example, if you take a sporting venue deploying enhancements for cellular services, most of the venues are also asking for Wi-Fi. This is not just from an offload perspective, but because they are also providing other services within the venue, like video. There is a lot of traction on video, which is used for things like instant replay to keep the fans within their seats. We are taking Infinity Access as a first stage, even for digital DAS, and moving forward with Wi-Fi within stadiums.

Frank Rayal. And obviously with your type of system, you can put Wi-Fi and LTE, whether it is a remote radio head or a compact base station, and all that can be multiplexed on a single fiber strand and distributed to multiple locations.

Saeed Anwar. Correct, yes.

Frank Rayal. All right, Saeed, I would like to thank you very much for being part of this conversation.

This conversation is part of a Senza Fili report on small-cell backhaul that provides an overview on small-cell backhaul solutions, along with in-depth conversations like this from leading vendors who opted to participate in the report. The report can be downloaded from the Senza Fili website at **www.senza-fili.com**. I would like to thank the audience for joining us today, and Saeed, again, thank you very much for being part of this conversation.

Saeed Anwar. Thank you.

Acronyms

CPRI	Common Public Radio Interface
CTO	Chief technology officer
DAS	Distributed antenna systems
DWDM	Dense WDM
KPI	Key performance indicator
LTE	Long term evolution
NLOS	Non line of sight
OBSAI	Open Base Station Architecture Initiative
R&D	Research and development
RF	Radio frequency
RRH	Remote radio head
WDM	Wavelength division multiplexing

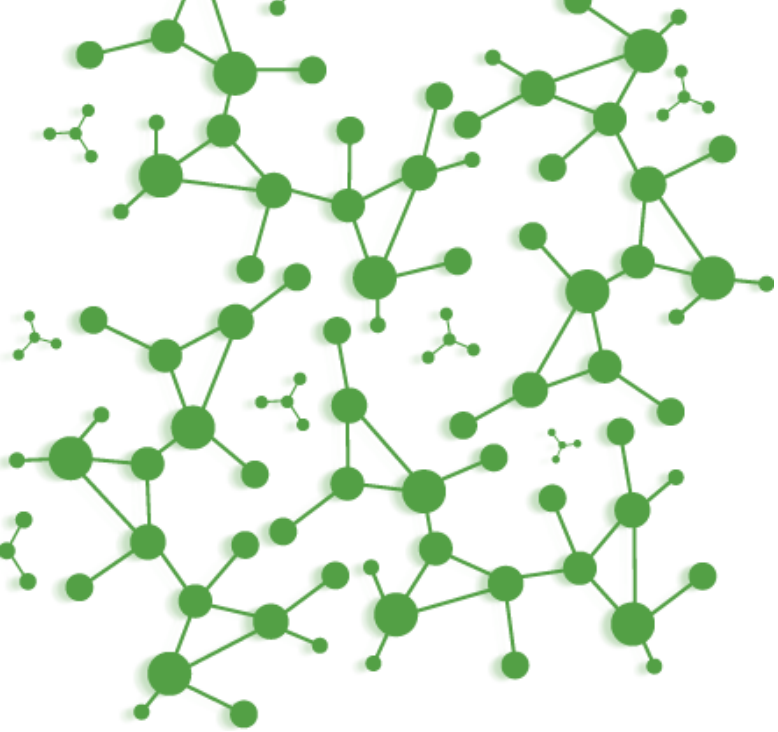

Addressing the scalability challenges of small-cell backhaul

A conversation with Tim Doiron, VP of Mobile Routing Product Development, Tellabs

By Lance Hiley
Senza Fili Consulting

Lance Hiley. Welcome to our audience with Tellabs. This conversation is part of a Senza Fili report on small-cell backhaul that gives an update on the small-cell backhaul solutions available and on the evolution of the mobile operators' requirements for small-cell backhaul.

My name is Lance Hiley, and I am here today on behalf of Senza Fili. We are talking to Tim Doiron, VP of Mobile Routing Product Development at **Tellabs**, which is a company that delivers a wide range of mobile internet solutions to mobile operators globally. So with that, Tim, perhaps you could explain to us, exactly what is the Tellabs position on small-cell backhaul products today?

Tim Doiron. Tellabs sees small-cell backhaul as a growing and evolving part of the mobile environment. Service providers are seeking to deploy solutions for locations we call hotspots and notspots – hotspots being denser urban areas where people may migrate to – the central business district of a city like Chicago, London or Helsinki — and notspots being at the edges of the network, where additional coverage may be needed. Tellabs has been in the macro mobile backhaul market for many, many years, and we see small-cell deployments and the backhaul associated with that as a logical extension of macro-cell backhaul.

Lance Hiley. You mentioned hotspots, obviously the dense urban metro areas that a lot of operators are challenged with today, as they try to figure out how to get capacity into city centers and into urban areas, but also notspots, the areas – for example in rural and some suburban situations – where today operators lack the mobile broadband performance that they would like to be able to deliver to their customers. Where does the Tellabs proposition fit into this? Are you positioning yourself across the full spectrum of the small-cell backhaul opportunities – notspots and hotspots – or are you focusing in on one of the segments?

Tim Doiron. Tellabs has a three-part development plan around supporting small-cell backhaul. We call the first part interoperability, and the second part is what we call pre-aggregation, and then integration.

From an interoperability perspective, Tellabs is interoperating with other ecosystem partners for things like next-generation microwave backhaul for small cells and macro cells. In addition to that, we are building a stand-alone mobile backhaul network element of a very small size and power for pre-aggregating the multiple small cells that you might find in an urban area or central business district. And then lastly, as part of our partner program, we are looking to integrate that same pre-aggregation technology into a set of partners' solutions for providing the layer 2 and the layer 3, and the backhaul component of the infrastructure.

Lance Hiley. That is really interesting, Tim, because certainly from my experience, Tellabs made their name in mobile, over the past few years, in providing a kind of universal interface at a lot of cell-site locations, that could integrate different technology – 2G, 3G and, now of course 4G eNodeB elements that are starting to come onto the platform – into a single interface. With your know-how to address that situation, it sounds as though you are trying to extend that expertise to help operators to cope with what might be a variety of different small-cell solutions that they are trying to aggregate at different points in their network, is that correct?

Tim Doiron. Yeah, that is very well stated, and that is exactly what we are attempting to do. If you take a look at small-cell deployments and small-cell backhaul, there will be more small cells than macro cells, and they will cover smaller geographies or portions of the network. And service providers, for cost-effectiveness reasons, will look to use whatever technology is available at the least cost, to provide the reliable transmission back to some other location, whether it is central office or macro-cell site location. Tellabs approach is to provide multiple physical-interface options via a combination of our standalone Tellabs 8602 Smart Router product and the integration of that technology into other ecosystem or partner products, as well.

Lance Hiley. Going back to your pre-aggregation product for a moment: It sounds as though that is going to be a key part of your strategy. What would you say your competitive advantage would be in providing that sort of network element when, listening to a lot of the small-cell backhaul and even some of the small-cell node vendors, they seem to think that they have a proposal for that technology as well?

Tim Doiron. The benefit Tellabs can bring to the overall solution is our years of experience in backhaul.

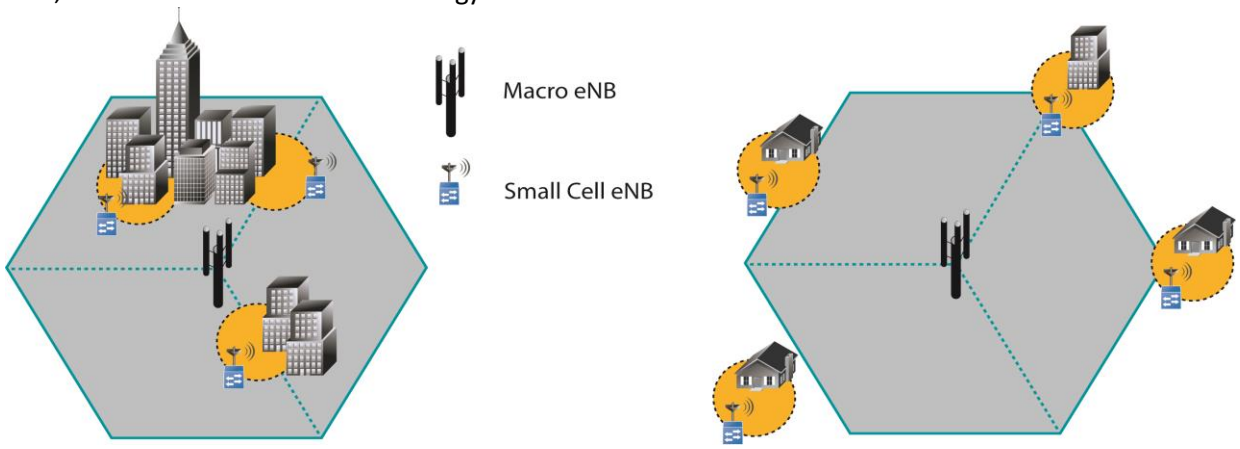

Hotspot capacity

Coverage extension near or beyond cell boundaries: notspots

Figure 14. Common use cases for small cells. Source: Tellabs

We provide the layer 2 and the layer 3 functionality, including the operations management and performance monitoring of the network components, along with industry-leading packet synchronization and timing capabilities that we spent many years evolving and adapting for packet-based networks and furthermore, getting tight timing on those networks for LTE, and LTE–Advanced.

The last piece is our universal network management called Tellabs 8000 Intelligent Network Manager. As you look at bringing a network management umbrella over the solution, you give the service provider the opportunity for easier deployment, easier troubleshooting, rapid turn-up and plug and play – required capabilities both for reliable deployment, and in ongoing support and maintenance of that network.

Lance Hiley. The really interesting thing for me in that proposal is that it sounds as though you want to make that technology available really in two ways. First, you can make it available as a stand-alone component that could be purchased by an operator and put in a small-cell location, and they can plug in whatever small cell device or other backhaul or network element into that solution of yours, and leverage that experience that you have got. But it also sounds like you want to possibly even integrate that into other vendors' equipment. How would that work?

Tim Doiron. That is correct. We have the stand-alone technology embedded into the product we call the Tellabs 8602 Smart Router, and we would take that same technology and make it available in a small-form-factor piece of hardware and software that we could integrate with a third-party supplier, and with that

provide the layer 2 and layer 3, as well as the performance monitoring and network management pieces of the network.

What exactly that is will certainly evolve over time and as we look at cost effectiveness and the need to continue to follow the cost. So maybe today it is a combination of multiple pieces of silicon and software; over time, you could have that evolving into just embedded software if there is enough processing power, capacity and capability in the third-party supplier equipment.

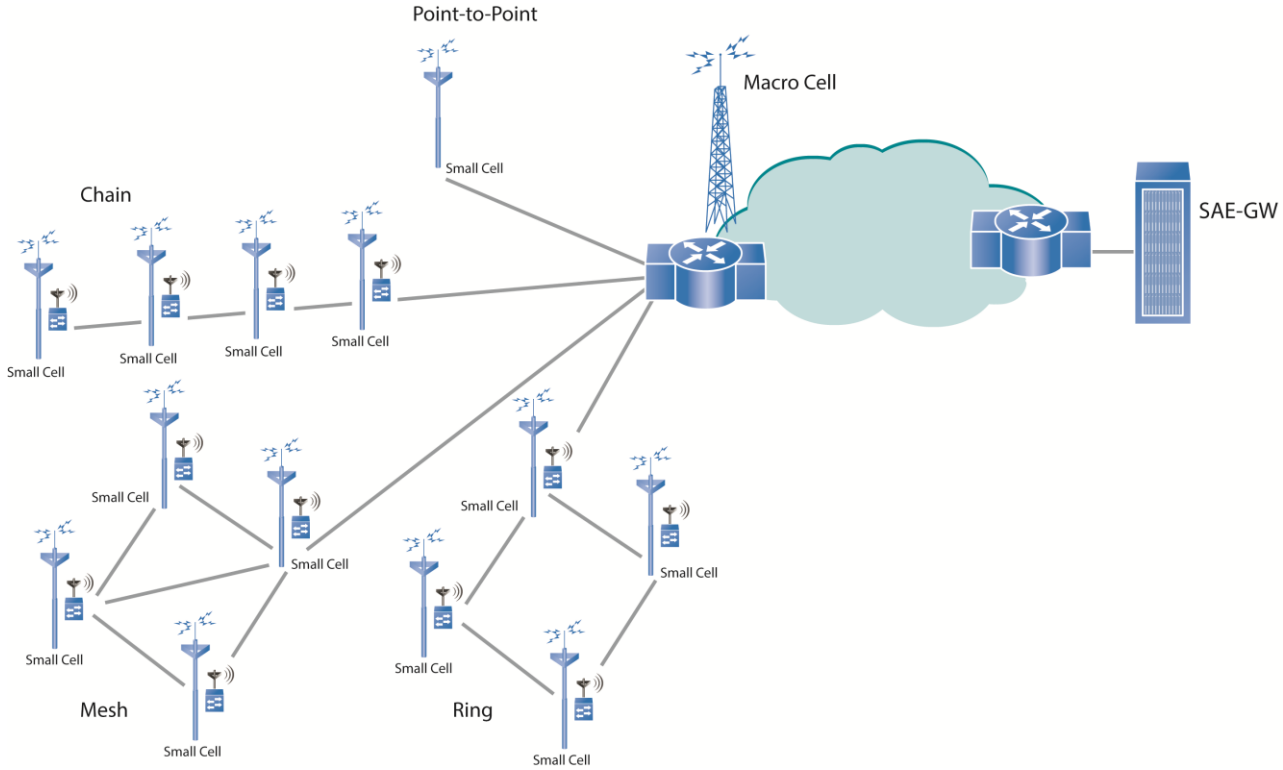

Figure 15. Different small-cell-site aggregation topologies. Source: Tellabs

Lance Hiley. In other words, when we look at the small-cell industry right now, especially when we look at small-cell backhaul, there are a lot of new entrants that are bringing innovative solutions to market. What you are essentially saying is that they can leverage the experience that Tellabs has with network integration, literally design that – whether it is through a subsystem or even to a silicon level – right into their products, and leverage the very best of your experience as well as their own.

Tim Doiron. That's right. And this gets back to what do various equipment suppliers believe their core competencies and their own differentiation attributes are for their products.

Tellabs believes that the experience and the years of deployment that we have gained and the intellectual property gained around deploying macro-cell and mobile networks has value, and can be brought to bear for equipment suppliers – whose core competencies or skills may be next-generation microwave, millimeter wave, or non-line-of-sight technology.

We could alleviate their need to focus on layer 2, layer 3 and the OA&M-type functionality, and the universal management aspect of that. Thus we let them differentiate and compete at what they are best at, while leveraging the skills, talents and capabilities that Tellabs has garnered over the last many years of mobile network deployments.

Lance Hiley. That sounds like a really compelling proposition to a lot of these new innovation entrants to the marketplace. How is that partnership realized? Is it a purely a technical partnership, or will you do co-marketing as well?

Tim Doiron. Well, we are in the early phases of identifying and laying out this program, but our vision is that this would be a development activity, but we are certainly open to co-marketing and co-selling discussions as well. Certainly we would begin with both a technology discussion and an integration-type discussion.

Lance Hiley. Clearly one of the main attractions of such a partnership program, to some of the other

vendors that you might be working with, is how it might save them money in terms of development. Ultimately, both you and your partners have to deliver a compelling cost proposition to your end customers, mobile operators. How do you see your program and the different elements of your small-cell backhaul strategy actually addressing the capex and opex concerns that operators have today, as they plan their small-cell networks?

Tim Doiron. You certainly raise the right two points that service providers are very much aware and concerned about: the capital expenditure associated with small-cell deployments, and the operational cost to do those deployments.

From a capital expenditure perspective, the numbers vary, but in general people think that you need to reach one-fifth, or one-tenth perhaps, the cost of the macro-cell deployment model, because of the sheer volume of small cells, and the small capacity range or geography that they are going to cover. So one of the things Tellabs has done, and one of the reasons it likes the pre-aggregation idea, is that a single network element can pre-aggregate multiple small cells, and so you don't need to deploy additional layer 2 or layer 3 functionality, either at the location of each one of those small cells or embedded in the small cell. The pre-aggregation function can add some statistical multiplexing value, which can help amortize the capital cost across multiple small cells deployed in an urban area.

In addition to that, we have invested time and energy in the operational side of deployment: how rapidly can I deploy and get the service turned up, the backhaul element configured, and move on to the next one? That contributes to some extent to your operational and deployment costs.

Tellabs has also been investing in a network management system and network elements to allow a plug-and-play, self-organizing, self-optimizing network approach wherever our small-cell pre-aggregation device is installed. It can let the network know it's present, and it can take an automatic configuration file download, and be up and running, active and in service with minimal human intervention.

Lance Hiley. Your hardware proposition touches on a number of different network management and traffic management – as well as just pure hardware and installation – concerns. It feels like a very powerful proposition. What sort of things can we expect to see in the future from Tellabs to further enhance what you are offering us today?

Tim Doiron. I think that you will see Tellabs continue to enhance performance monitoring and performance management, and being able to anticipate problems in the network before they happen. We are already doing that with some of our software tools, like our Insight Analytics that can detect when a path is down and your network is in a backup path state. That wouldn't be an outage, but it makes sure the issue is visible and highlights that it is a risk to the overall network. You will see us continuing to invest in that area, and to advance our automation, our rapid deployment and troubleshooting tools so that we can help continue to minimize those operational costs, in addition to the capital costs, associated with small-cell deployments.

Lance Hiley. That is great Tim, thank you very much. It sounds like we have a lot to look forward to from Tellabs and certainly your comments and observations will be a great contribution to this Senza Fili report on small-cell backhaul.

The report aims to provide an overview of the different small-cell backhaul solutions available, along with an in-depth interview like this one from all of the leading vendors who have participated in the report. The report can be downloaded from the Senza Fili website at **www.senzafiliconsulting.com**. I hope everyone enjoyed our interview today with Tim Doiron from Tellabs.

Acronyms

2G	Second generation
3G	Third generation
4G	Fourth generation
eNB	Evolved NodeB
LTE	Long term evolution
OEM	Original equipment manufacturer
SAE-GW	System architecture evolution gateway

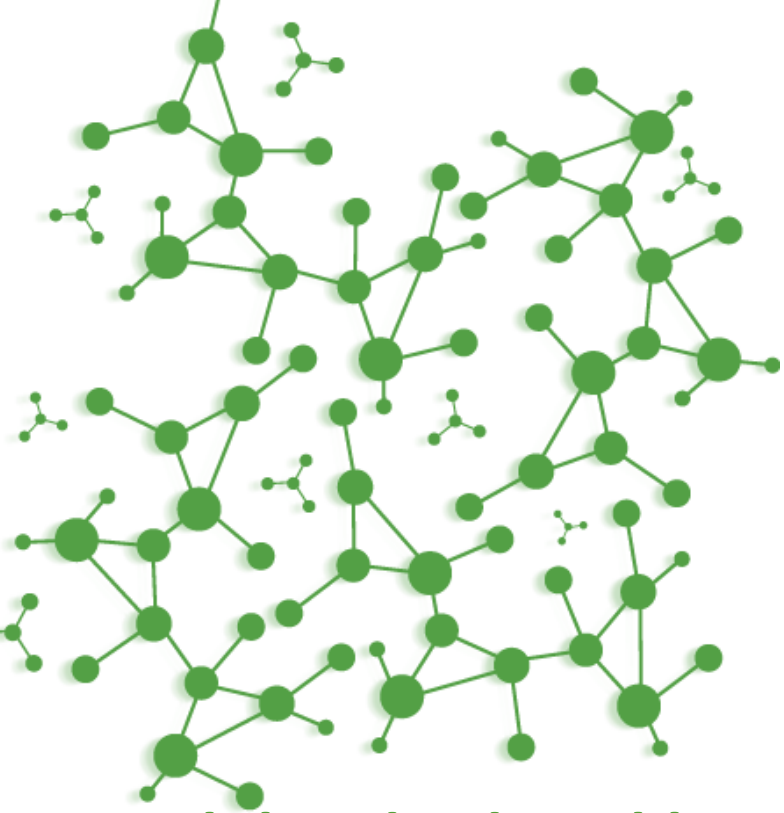

Multiband, adaptable small-cell backhaul

A conversation with Adam Button, CEO & Mike Pettus, CTO VubIQ

SENZA FILI CONSULTING

By Frank Rayal
Senza Fili Consulting

Frank Rayal. Hello, and welcome to this conversation with VubIQ. My name is Frank Rayal. This conversation is part of Senza Fili's report on small-cell backhaul, which gives an update on small-cell backhaul solutions and on the evolution of the mobile operator's requirements for small-cell backhaul. Today we are talking with Adam Button, CEO of VubIQ, and to Mike Pettus, CTO of VubIQ. VubIQ is an innovator in the small-cell backhaul space, particularly focusing on the 60 GHz band.

I would like to turn it over to you, Adam, and start off by asking you to give us a brief introduction about your company and your small-cell backhaul solution.

Adam Button. Sure, Frank. Thank you very much. VubIQ has been around for about 9 or 10 years now, exclusively focused on 60 GHz technology. We have developed products in the R&D market (for example, components and development systems), and we also build video links, creating products that transmit uncompressed high-definition video wirelessly at 1.5 gbps. These are for broadcast applications.

Our video link offerings led into a nice order with the National Hockey League to equip every net in the NHL – all thirty teams and all thirty venues – with GoalCams. These are high-definition cameras looking at the goal area to adjudicate whether or not a puck crossed over the line. The GoalCams are positioned in the back of each net, with the HD camera on one end of the unit focused on the goal line, and the other end wirelessly transmitting uncompressed video at 60 GHz up to the judging booth.

So we not only have theoretical and lab applications, but real-world applications of 60 GHz links. We are now moving into telecom, and are developing some very interesting products for small-cell backhaul utilizing both the 60 GHz band and also another unlicensed band, 5 GHz spectrum, for a dual-band solution.

Frank Rayal. This is a very interesting proposition, so the solution you have, the HaulPass, is a dual-band product that covers both the 60 GHz and the 5GHz bands. Can you elaborate on that?

Adam Button. Sure. We are unlicensed spectrum advocates, and I believe that there are tremendous advantages to using unlicensed spectrum versus paying for licensed spectrum. Both the 60 GHz band and the 5 GHz band we are using are unlicensed, and we are incorporating both of them into a very efficiently integrated unit that will be able to transmit backhaul for small cells, both line of sight over 60 GHz, and non line of sight over 5 GHz.

The 5 GHz band can act as a failover – for example, if you've got a line-of-sight connection at gbps speed and something happens, like a huge rainstorm, a truck

passing by, or tree growth over time – anything to block that line-of-sight path – our system will immediately and automatically fail over to the lower 5 GHz frequency to maintain availability.

Frank Rayal. What sort of capacity and throughput can the HaulPass provide in the 60 GHz and 5 GHz bands?

Adam Button. At 60 GHz we are at a full gbps speed. Actually, the overall data rate is 1.25 gbps in full duplex, but the payload is basically one gbps, full duplex. At 5 GHz, we are using a 4x4 MIMO implementation of Wi-Fi that allows us to achieve 600 mbps. So with both radios operating at the same time in a line-of-sight condition, the total aggregate throughput will be 1.6 gbps in both directions.

Figure 16. HaulPass SC 5/60 GHz small cell wireless backhaul system featuring automatic antenna alignment. Source: VubIQ

Frank Rayal. Achieving 600 mbps per second for the 5 GHz band is tough to do. How do you achieve that?

Mike Pettus. The interesting thing about what we are doing is that we are using 802.11n technology in the product and 4x4 MIMO. 600 mbps, of course, is the peak physical-layer bit rate that we can achieve. There are four antennas at each end of the link for the 5 GHz band surrounding the horn antenna for the 60 GHz band. There is good antenna diversity in terms of the 4x4 MIMO, with both polarization and spatial diversity.

The other advantage of using 802.11n in a point-to-point application is that we have the ability to optimize the controls. These systems are not being used as access equipment, of course; they are being used in point-to-point bridge mode, and in so doing, we will optimize for maximum throughput, maximum bandwidth, highest interference mitigation with the knowledge that the other end of the link is us.

In other words, it is a closed system, and because it is a closed system, we do not have to deal with 802.11n access point issues, and the system is basically in a point-to-point bridge mode. We have turned on all of the high-throughput, high-bandwidth controls in the 802.11n MAC, such that we maximize throughput for our benefit.

Frank Rayal. One issue that is associated with unlicensed-band solutions is the interference and latency, especially when it comes to 802.11n-type products. How do you address operator requirements? Can you comment on what you hear from operators on the latency requirements?

Mike Pettus. The expectation for millimeter wave is very low latency. Our latency, including forward error correction, is less than 20 µsec on the 60 GHz link, and the 802.11n latency is less than 1 msec. We achieve those very low latencies on 802.11n because we have full control over the MAC and we turn off a lot of the access point features on the MAC. So the interesting thing is that the 802.11n does have the ability to be a very low-latency link if you have control of both ends, and that is what is important.

Frank Rayal. What other requirements do you hear from operators on small-cell backhaul, and how does your product meet those requirements?

Figure 2. HaulPass micro SC 5/60 GHz small cell wireless backhaul system. Antennas are fixed. Source: VubIQ

Mike Pettus. The other requirements that are becoming very, very important in carrier-grade networks are not only the physical layer requirements – getting throughput and latency down to where they

need it for LTE – but also at the next layer, we will call it layer 2 or even layer 2.5: Synchronous Ethernet, IEEE 1588, PTP timing, for the new LTE-Advanced systems.

In order to achieve the capacities they need, they are moving to time-division duplex (TDD). When you use time-division duplex in a very dense network, you have to give very accurate timing and phase information to the small cells in order for them to achieve their capacity.

Our radios are fully compliant with 1588v2 and have GPS timing availability, and we are using a combinational algorithm in our synchronization and timing – I will call it 1588 with GPS assist – such that there is a fallback mechanism. Some carriers are saying GPS is great, and they are even using it in their small-cell base stations as a method for synchronization, but they do not count on it. There may be interference or a situation where small cells do not receive the GPS satellite signals, and in that situation we fall back to 1588 completely. In that mode, the version and the implementation of the 1588v2 that we are using gets us to less than 500 nanosecond phase error at the small-cell base station, and that is very important for synchronization.

Frank Rayal. Can you maintain the same type of synchronization over the 60 GHz as the 5 GHz band?

Mike Pettus. The 5 GHz will probably have a wider window of synchronization error than the 60 GHz, simply because of the packet delay variation numbers that you incur at 5 GHz versus 60 GHz. So, again, the theme behind this product is that we have multiple methods of delivery for both payload and synchronization.

We have two bands, and 1588 and GPS. We feel by putting all of the tools from the tool kit into a single product is how we achieve extremely high reliability in this system.

To answer your question further, we also handle full security. Not only we are transparently transporting IPsec, which operators may use in their network for security, but we also are implementing MACsec 802.1AE, MAC layer security so that every segment – if there are repeated segments of wireless backhaul – is protected and encrypted with AES 256-bit encryption, as well as being able to transport multiple streams that are isolated from each other.

What we see happening in the business model is that backhaul networks may actually transport competitive carrier streams in the future. Your equipment has to be ready to isolate and encrypt for each payload stream, and we are fully implementing MACsec at the layer 2 portion of the product.

The other features that we provide are full MPLS and full QoS capability. The customer can configure the QoS parameters and a full set of OAM capability for monitoring and configuration. So, more than just a physical radio link, we see these units being placed out into the backhaul network as very capable network nodes. And what is going to happen is an evolution from a point-to-point network to a partial ring, partial-mesh.

Adaptability is the key word here, because when operators get out into the field and they realize where they have to put their small-cell base stations, it is not as clean and as nice as it was when installing macro cells. They just have to adapt, and the products have to be ready to work in that kind of an adapted environment.

Adam Button. One reason we have some expertise in this area is because Mike ran engineering at Metricom for about eleven years. Metricom, as you may know, put out wide area networks, wireless networks for people to access the internet before Wi-Fi even existed. They were putting up radios on top of street lights and light poles outdoors in the city that were talking to each other in mesh configuration. This is remarkably similar to the current challenge for the small-cell infrastructure that we are looking at today – the major difference being that the radios today are about 1,000 times faster.

So we've got a lot of experience in real-world issues, and are developing our products knowing that when a truck rolls, what is listed on a plan or a map is not always the case, because there may be a building, a billboard or a tree or something in the way, and you have got to be able to adapt when you are out in the field.

Frank Rayal. This is very interesting, and there is actually a whole different area in terms of the logistics. What are your thoughts on this, having had this experience of doing it in the past? And how did you design your product in order to make it even easier to deploy and save cost from a deployment and implementation perspective?

Adam Button. I will make a quick comment on that before Mike answers, but – one of your questions before was: what are the carriers looking for, and what are we offering the carriers? Cost is clearly of huge importance to the carriers.

Now I am not talking just capital cost of the equipment itself, we are talking total cost of ownership. While our units will be very competitive on a cost basis alone, where they really shine is by minimizing the total cost of ownership because of the second major feature of this unit, which is automatic alignment and maintenance of a link.

So when you set up the link, it makes it very easy for a low-skilled installer to plug it in and walk away, without requiring the time, sometimes hours, needed for two installers, one on each end of the link, to get them set up and aligned, and to optimize the link and minimize the bit error rate. That takes a lot of time, and when you multiply that by the hundreds of thousands of small cells that are going to be deployed over the next few years, you see the installation cost is going to be tremendously high.

This is another feature that differentiates our product, not only from other categories, but from other 60 GHz solutions, and we think that it will be of great benefit to the carriers.

Mike Pettus. The auto-alignment and auto-stabilization feature of the HaulPass SC is key for carriers whose business models require a lower opex. In other words, if they want to minimize installation time and installation labor, the HaulPass SC is the product to go after which provides the auto-alignment. Because the installer can put it up, and he has no idea where it is pointed, he just installs it, connects it and goes away. It is literally a ten-minute install, and what happens is the system itself automatically provides alignment of the link.

In millimeter wave this is critical, because the antenna beamwidths at those frequencies are on the order of 2 to 3 degrees, so it requires very precise alignment.

Auto-alignment is a great thing for reducing the opex, because many of the platforms that they will be mounting on, like light poles, move in the wind, and for millimeter-wave links, this can be a disaster, because you can fall in and out of the beam as the pole is swaying. So the same mechanism that provides alignment also provides stabilization within a half degree. This is a big win in terms of the ability of the carriers and the installers to minimize their overhead and maintain a very reliable link.

To answer your second question, of how we provide an answer to an adaptive environment, you hear a lot these days about self-organizing networks, or SONs. Most of what they are talking about is at a logical layer, and of course we do that. It is a given that you have to be able to adapt addresses and new nodes that appear and disappear in the network, but because we have the auto-alignment feature, we can physically do self-organized networking.

One of the carriers said that they move things around from time to time. They want to move the positions of their base stations around, and this is a big win when you have an auto-alignment feature.

Frank Rayal. So in terms of availability of the link and also in terms of the ongoing maintenance, that is a major cost saving. How wide can the alignment be? Can it be done vertically or horizontally?

Mike Pettus. If you look at the HaulPass SC product, the auto-alignment range in the azimuth – we will call it the pan – the range is 360 degrees; and the auto-alignment range for the elevation – which is rotating up and through, on the vertical plane – is plus or minus 120 degrees, or 240 degrees total. So it has a lot of adaptability to point just about anywhere.

Adam Button. And it can be installed inverted, upside down. So you essentially have omnidirectional capability.

Mike Pettus. Yes, so if you are on a tall building and you have to point down at a steep angle, you can actually invert the mounting of the unit and it will point down that way.

Frank Rayal. Going back to the 60 GHz technology, is there anything that you are doing that would be differentiated from other 60 GHz solutions? Can you also give us a flavor of the 60 GHz market and how you fit into that?

Mike Pettus. The 60 GHz technology that we are using is now on the order of four years old, and we have a lot of history of developing 60 GHz. Our fundamental core technology is at 60 GHz, and the key issues with developing a reliable and manufacturable 60 GHz product are in getting around the problems of packaging.

You know about integrated circuits: they have little leads, and then you put them down on the board, and it is very simple. At 5 GHz it is a simple matter: you buy chips and you put them on a board and connect them to an antenna. At 60 GHz you cannot use the same techniques to get the energy on and off the die or the chips, because of such a high frequency.

Our waveguide module – you can see it on our website – is the core technology inside these radios, which is a very producible and reliable 60 GHz element in the radios. This points to the fact that our manufacturing cost is low for producing a millimeter-wave radio.

The other performance issue with these products is that some techniques at 60 GHz are lossy. If you do not do this right, you will lose energy or signal. We measure it in decibels, and if you have enough loss, you impact your range. Our packaging is patented – the entire packaging that we use for waveguide interfaces and 60 GHz die interfaces are all protected under our IP. The losses incurred in these packages are less than a decibel, and at 60 GHz this is very important. This points to performance, simple performance.

Adam Button. And, Frank, we have been shipping silicon-based 60 GHz products since May of 2008, just about five years next month. So we have a tremendous amount of experience and IP: whether it is actual patents or just trade secrets, we have a tremendous amount of experience and knowledge when it comes to building highly efficient, low-cost 60 GHz products.

Frank Rayal. You obviously made a conscious decision in going into unlicensed bands, 5 GHz and 60 GHz. How do you see that comparing to other solutions on the markets? Solutions that use, for example, microwave frequency, line of sight, point-to-point or point-to-multipoint, and also non line of sight.

Mike Pettus. Faster, better and cheaper. Adam said earlier, we are unlicensed-band advocates. We absolutely believe that technology solutions are going to be faster and more effective than policy-driven solutions, such as licensing.

That is the first thing, and if you look at all of the current technologies that are going to be deployed in LTE – such as OFDM, MIMO, coding – were developed in unlicensed technologies, specifically in IEEE standards, five to eight years before the cellular industry adopted them. So unlicensed technology usually gets more quickly to market with very high-technology solutions, very effective solutions, before the licensed users get to them, number one.

Number two, interference is a big question, and you asked that earlier, and I want to address that. Interference controlled by allocating fixed spectrum for users has been used traditionally in radio ever since the FCC came into existence – that is how interference has been controlled on a licensed basis. What has happened in the world of radio is the big word: digital. Digital radio techniques – OFDM, dynamic frequency selection, coding, and also MIMO, which is one of the biggest things during the last ten years – are the best solution to interference. Rather than saying, "I'm going to allocate a piece of spectrum, and no one else can use it but a certain entity or the licensee."

The third reason: Licensed spectrum by definition, especially below 5 GHz, is narrow-band. You might get a 10 MHz or a 20 MHz license; you might get it for an area, or you might get it nationwide at great expense. But the 5 GHz unlicensed band is on the order of a few hundred megahertz wide, and the FCC made a statement at CES in January that now another 195 MHz is being added. So it is not only a better technical solution – the policy makers are now realizing that unlicensed solutions are a good way to go, such that they are giving more and more spectrum to unlicensed operations. So we are big advocates of it.

On interference mitigation, at 60 GHz interference mitigation is basically the physical layer. 60 GHz has a very narrow beamwidth, and the falloff due to oxygen absorption, which at one time was considered a bug, is now a feature, especially in dense networks. The isolation, and therefore the network density, that you can provide at 60 GHz is unbelievable compared to other spectrum bands.

And at 5 GHz the impact of using MIMO for beam steering and also beam nulling, if there is an interfering source at a given angle from the radio MIMO, has the ability to null out that source dynamically on a packet by packet basis.

So our answer is: let technology figure out the interference solutions, rather than policy.

Adam Button. You also mentioned e-band and 70-80 GHz solutions that are also millimeter wave. These are light-licensed solutions, and people would say, "Okay, it is very simple: you get on the web and you fill out a small application for the link and pay a small fee and

no problems, it is a minor issue." Again, when we talk about rolling out tens of thousands, hundreds of thousands of small cells, each requiring backhaul, that becomes an entire huge process, another administrative process that is going to add cost to the carriers.

Frank Rayal. I've got one final question and that is about your roadmap and the evolution of your product. Where do you see taking the product next?

Mike Pettus. The roadmap for this product, which is a high-bandwidth backhaul product, simply points to one major area, which is bandwidth. We see the bandwidth going up over the next few years in the 60 GHz spectrum; at the end of our roadmap, we see 60 GHz carrying up to 10 gbps and we see more spectrum being added to the 5 GHz band, which gives us, again, more spectrum and more wide-band channels.

The first parameter that you will see change as products move along in the backhaul is bandwidth and, of course, the LTE-Advanced Release 10, and Release 12 coming around the corner. What is the major change that they need? More bandwidth, and that is at the top of the list.

Below that you will see a product lineup, with a variation of products that fit certain needs. There will be variations on the number of ports, the capability to cascade and mesh, and these types of things. So as the network complexity grows, the products have to adapt to be able to work in those highly complex networks.

Frank Rayal. Well, thank you very much, Mike and Adam, for a very insightful conversation. This conversation is part of a Senza Fili report on small-cell backhaul that provides an overview of small-cell backhaul solutions along with in-depth conversations, like this one that we just had, from leading vendors who participated in the report. The report can be downloaded from the Senza Fili website at www.senza-fili.com. Thanks again, Adam and Mike.

Mike Pettus. Thank you again for having us. We appreciate it.

Acronyms

AES	Advanced encryption standard
CEO	Chief executive officer
CES	Consumer Electronics Show
CTO	Chief technology officer
GPS	Global positioning system
IEEE	Institute of Electrical and Electronics Engineers
IP	Intellectual property
IPsec	Internet protocol security
LTE	Long term evolution
MAC	Medium access control [Layer]
MACsec	MAC security
MIMO	Multiple input, multiple output
MPLS	Multiprotocol label switching
NHL	National Hockey League
OAM	Operation Administration and Maintenance
OFDM	Orthogonal frequency division multiplexing
PTP	Precision time protocol
QoS	Quality of service
R&D	Research and development
SON	Self-organizing network
TDD	Time-division duplex

About Senza Fili

Senza Fili provides advisory support on wireless data technologies and services. At Senza Fili we have in-depth expertise in financial modeling, 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.

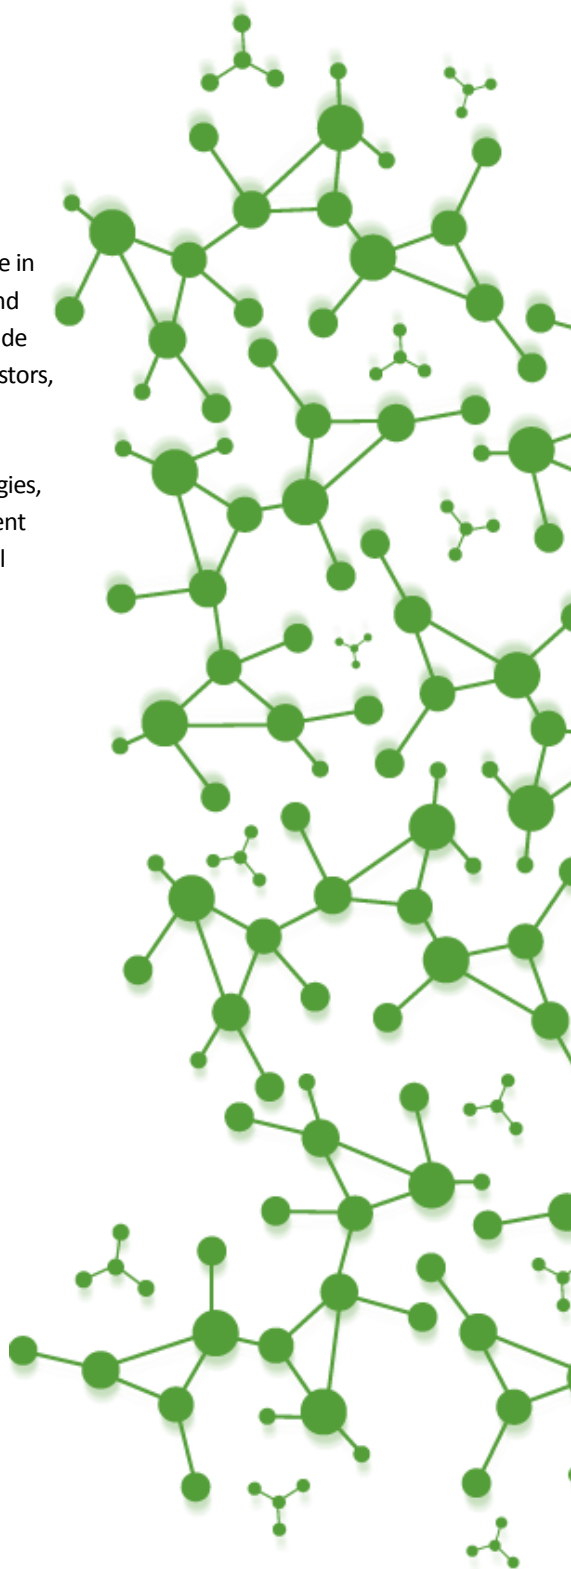

About the authors

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**.**

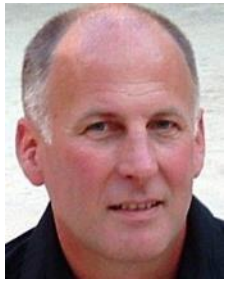

Lance Hiley is an independent wireless marketer and analyst working on behalf of Senza Fili. He is an expert in wireless backhaul technologies and worked with companies and industry organizations to bring clarity to backhaul products and issues. He has given numerous presentations at conferences, and has written articles and papers on wireless broadband technologies and products. He can be contacted at lance.hiley@gmail.com

Frank Rayal is a telecom industry professional with over 20 years of experience working with network operators and system vendors to develop and deploy innovative wireless solutions. He is a founding partner at Xona Partners a boutique management and technology advisory firm specialized in TMT and a founding member of small cell backhaul pioneer BLiNQ Networks. Frank held senior product management, marketing and business development positions at Ericsson, Redline, and Metawave. He holds a BS in Electrical Engineering from Case Western Reserve University, Cleveland, OH, and a MSc in Electrical Engineering and an MBA from the University of Toronto, Canada. Frank is a Senior Member of the IEEE, and a member of Professional Engineers Ontario.