

# Understanding Small-Cell Wireless Backhaul

Bernard Prkić

Article

April 3, 2013

The cellular industry is buzzing with small-cell fever. Within the general thrust of implementing small-cell outdoor cellular networks, backhaul is seen as one of the technical and economic barriers to deployment. Wireless backhaul technologies seemingly have the capabilities to cost-effectively overcome this barrier.

## Background

Small cells are not a new phenomenon. In 1991, analog advanced mobile phone service (AMPS) operators (1G operators) in North America successfully deployed small cells to combat debilitating congestion that was averaging 15% and severely testing customer loyalty.

In those early days, voice capacity demand—although tiny according to today's standards—was already outstripping supply. Traditional capacity enhancement methods were either simply not available (more spectrum and carriers) or already taken to their interference limits (macro cell densification, cell splitting).

The solution that did work was to provide blanket coverage at all corners of nearly all major streets in the grid in dense urban areas, using 200-m spacing between microcells and a 100-m microcell radius. The 1G AMPS small cells were rather costly, yet the solution paid off by improving customer satisfaction and reducing churn.

With the advent of 2G/GSM, the need for small cells diminished. GSM was able to carry significantly more traffic at higher quality than AMPS and at much lower cost due to extra spectrum, higher

spectrum efficiency (e.g., through efficient voice codecs)<sup>1</sup> and superior frequency reuse capability. Despite the lesser need for small cells, several 2G vendors offered them.

As early as 2000, a well-known GSM infrastructure vendor launched its ultra-compact picocellular indoor single-carrier 2G basestation, marketing it as “laptop-size, laptop-price.” Both claims were true. This product was rather revolutionary, both in terms of price (at least one order of magnitude lower than macro basestations) and size (several orders of magnitude smaller), so it was a truly unique value proposition.

The product was very similar to modern microcellular basestations in terms of look and feel and area coverage capability. Despite all its obvious merits, it was only deployed sparingly, mostly at the carrier’s own headquarters and equipment room (e.g., switch) premises, primarily to provide premium indoor coverage to employees and guests. In some instances, 2G microcells were deployed to plug coverage gaps.

The conclusion is that small cells are neither new nor revolutionary. However, they have not been deployed in large numbers since the advent of second-generation (2G) mobile standards.

#### Four Obstacles In Adoption

Why were early micro/picocellular 2G basestations not successful, and why aren’t 3G small cells ubiquitous? Knowing the answer to those questions helps us assess whether small cells are here to stay or just hype.

First, the very notable increase in the capacity of macrocellular networks since 2000 has pushed the necessity for small-cell deployment further into the future. Technological advances like higher-order modulations and multiple-input multiple-output (MIMO) have boosted spectral efficiency, while remote radio heads improve the link budget, achieving lower noise figure receivers and more transmit power by eliminating feeder losses and noise.

Also, regulators have opened up more spectrum (e.g., digital dividend, new bands) for mobile broadband, responding to the growing importance of the industry. Highly efficient yet robust and adaptive voice codecs reduce the bandwidth required for a single voice channel to a mere 8 or 16 kbits/s as well.

Second, small cells require excessively high operational expenditures (OPEX), including but not limited to mobile backhaul. About the year 2000, E1/T1/STM/SONET (leased) lines were most widely used to backhaul cell site traffic. Physically, they were carried by copper using symmetrical high-density digital

subscriber line (SHDSL), by plesiochronous digital hierarchy (PDH) or synchronous digital hierarchy (SDH) microwave radio, or by fiber through SDH or PDH multiplexers.

A typical macro site would boast six to 18 carriers, covering a large geographical area, with a 3-km radius not being exceptional. A single E1 or T1 would support up to 12 carriers. This carrier-to-E1/T1 ratio made good economic sense at the time. With a single/dual-carrier microcellular or picocellular basestation, the very same expensive E1/T1 would typically serve just one or two carriers, raising the transport cost per carrier (per unit effective on-air capacity) by a factor of six to 12, or roughly one order of magnitude.

Additionally, the geographical coverage of a microcell basestation (up to 10,000 m<sup>2</sup>) would be in the order of 1/3000<sup>th</sup> of the coverage of a traditional macro cell site, providing roughly three orders of magnitude less coverage per unit transport capacity. Consequently, small-cell backhaul OPEX alone would have killed the business case for small cells back in 2000.

Third, there was no ecosystem for mobile broadband driving the need for even higher capacity densities in mobile networks that would necessitate small-cell deployment. Primarily, there was a lack of consumer demand for mobile broadband. Suitable video content was simply not available on the Internet. Remember that YouTube came into being in 2005. Video and video on demand are the key drivers for (mobile) broadband today. Next, there was a lack of powerful, easy-to-use smartphones for consumers to deal with rich multimedia Internet content while on the move. The first successful consumer smartphone was the iPhone, launched as recently as June 2007.

And fourth, 3G small cells require an entire carrier to operate. Operators typically have just one to three broadband 3G carriers (5, 10, or 20 MHz each) at their disposal (against dozens of narrowband 200-kHz carriers in 2G) and are not willing to give up one (macro network) carrier for localized small-cell deployment. Co-channel operation of small cells in 3G is problematic, certainly at the cell edge. Due to the low operator demand, mainstream network vendors simply do not offer 3G small-cell systems.

## Small-Cell Deployment

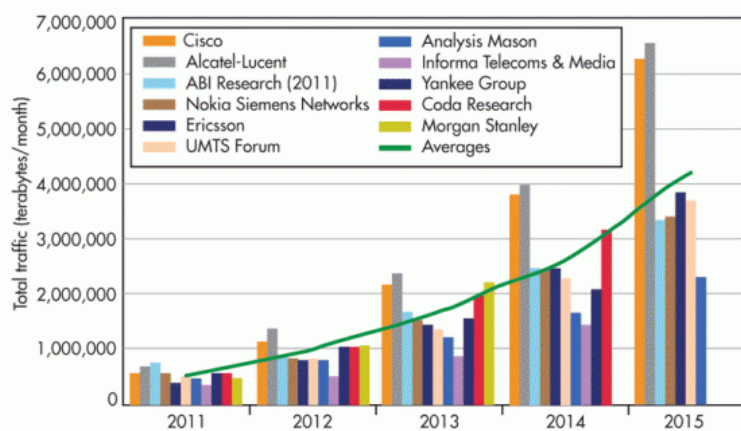
Developed markets, where LTE was widespread in 2013, will lead small-cell deployment. LTE is an indispensable driver for small cells as it represents the ultimate in terms of spectral efficiency today and tomorrow (LTE-Advanced). More specifically, on a network level, small cells serve three main and overlapping purposes:

- Increasing capacity density in areas with high user densities, mostly in (indoor) public places in dense urban areas and inside buildings with a high subscriber density
- Improving coverage and thus available data rates and service quality in important, user-dense, marginal-coverage areas that cannot be adequately reached by macro cells such as at cell edges or in “black spots,” both outdoors and indoors
- Extending handset battery life since higher-order modulations, broader bandwidth, MIMO, and enhanced processor power all drive up handset power consumption while smaller cells reduce power consumption

Whether we’ll see small cells deployed has proven quite trivial, at least on a conceptual level. Yet the main and difficult questions of when we’ll see widespread deployment of small cells and how these deployments will take shape remain.

The mobile broadband landscape has changed beyond recognition since 2005. Today, most people in advanced markets have developed the latent need to access any content, anytime and anywhere, at a high speed. Ubiquitous broadband access to the Internet has become a necessity of life. This is why customers are forcing carriers to deploy 4G and, ultimately, small cells to meet their expectations, particularly where user density is highest and spectrum availability is limited.

One of the main four impediments to small cell deployment, then—the lack of demand for mobile broadband—is definitely no more, at least in developed markets (*Fig. 1*). We also know that the capacity (density) of current macrocellular 4G networks will continue to increase in the foreseeable future since there’s still spectrum available around the world that could be used or reused for mobile broadband.



1. Projections vary, but driven by video and other bandwidth-heavy applications, mobile traffic definitely will see significant increases through 2015.

If there's a strong need to free up that spectrum, and there is, regulators will be inclined to do so, as they have done in the past, because of the sheer importance of mobile broadband for societies. At the same time, spectrum is by default a finite resource and much of the most suitable spectrum for macro cells, featuring low propagation losses and good obstacle/building penetration, has already been allocated to mobile broadband.

Even if more sub-6-GHz spectrum is eventually allocated for mobile broadband use, it won't happen overnight. It won't be cheap either. Therefore, we cannot expect miracles in terms of capacity gain through new spectrum allocations. Most short-term (efficiency) gains will probably be made through refarming current 2G and 3G spectrum for 4G use.

Fortunately, there are ample well-known and more experimental ways to further increase spectral efficiency, like 4x4 or 8x8 MIMO, adaptive beamforming, higher modulations, and coordinated multi-point. It's hard to gauge where this will lead us in terms of spectral efficiency, but it can take us pretty close to the theoretical Shannon limit. Some of those promising techniques look less promising, though, when subjected to closer scrutiny.

Using higher modulations is a proven, reliable, and well-understood method to increase capacity in a given communication channel, but it has clear limits. With each and every modulation step that we take, such as from 64-state quadrature amplitude modulation (QAM) to 128QAM or from 1024QAM to 2048QAM, the available link budget decreases more or less linearly by 2 to 4 dB, limiting range and thereby cell size and, ironically, capacity. To add insult to injury, the higher the modulation, the lower the differential gain of switching to the next-higher modulation. The relative capacity gain decreases as a function of modulation order.

Next, 4x4 MIMO is very hard to implement on our diminutive "pocketable" mobile devices. 2x2 MIMO appears to be the practical limit, simply due to the small dimensions of portable mobile equipment that users expect. There is insufficient room on the device for proper antenna spacing between the four MIMO antennas in the favorable sub-6-GHz bands used by macro cells.

Surely, using higher frequencies like millimeter waves could fix this MIMO antenna placement challenge, yet it would at the same time increase path loss significantly and be very ineffective if not dysfunctional in non-line-of-sight (NLOS) scenarios that are rather the rule in macrocellular settings. High-order MIMO could in the end be confined to (near) line-of-sight (LOS) scenarios, such as small cells (*Fig. 2*).



2. A three-band, six-sector LTE micro site with single-band antennas using 8x8 MIMO would be very complicated. Now imagine your eight-antenna smartphone or tablet.

Adaptive beam forming and beam steering look great on paper. They also work superbly for military radar applications and perhaps in the countryside where (near) LOS between eNodeB and the mobile device is more readily available. But in the common NLOS situation in the urban macro-cell jungle, where we'd need the additional capacity most, trying to hit and consistently track an "invisible" and (fast) moving target in a totally unpredictable and dynamic electromagnetic environment can be a vexing challenge.

Add to that the sheer cost and complexity of large phased array antennas required for "pencil" beamforming, and it doesn't take much thought to dispel this approach—one that has been around for almost two decades—as a viable solution to capacity density challenges downtown, except again perhaps for millimeter-wave LOS (small cell) scenarios. Another challenge with beam steering is that space division multiple access (SDMA) has to be added into the media access controller (MAC), and this is very complicated.

The availability of new spectrum and spectral efficiency increases will continue to work against small-cell deployment, but the efficacy and cost-efficiency of those measures (their incremental gain) will decrease. In the end, everything boils down to “bang for the buck.”

So depending on the local regulatory framework and technological advances in spectral efficiency, in terms of total cost of ownership, small cells will outrun advances in macrocellular spectral efficiency, the cost of additional spectrum and/or “re-farming” 2G and 3G spectrum, and distributed antenna systems.

The OPEX component, including the backhaul OPEX, also deserves closer scrutiny as an impediment to small-cell success. It will be governed by the cost of site acquisition and site rental with a readily available power supply and availability of wired backhaul. It will be governed by mounting height as well. In case mounting height is limited to 2 m above ground level, the small-cell output power is legally limited, materially reducing cell size and coverage.

With wireless backhaul, OPEX will be governed by LOS to a suitable chain, the hub or fiber point of presence (PoP) site, or the non/near-LoS capability (nLoS/NLoS) of the wireless backhaul solution. OPEX also will be governed by installation and maintenance costs. For example, do we need a union crew with a bucket truck and traffic police, rerouting traffic for all installations or maintenance like in New York City if a traffic light is acquired?

Spectrum fees (when wireless backhaul occurs in a licensed band) and depreciation are factors as well, though they are well known and readily quantifiable, and they don't pose a challenge for widespread small-cell deployment. The costs of site acquisition, installation, and maintenance, though, remain true and open challenges.

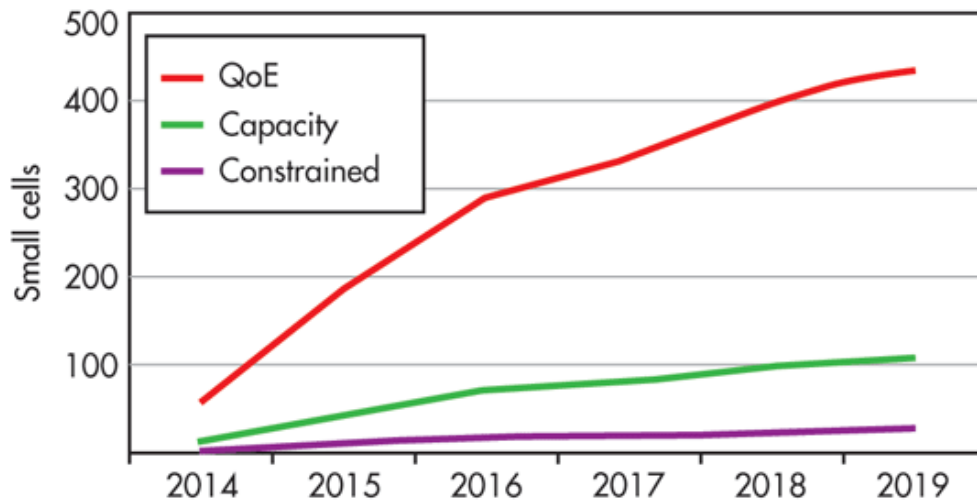
A satisfactory solution for the site acquisition and installation challenges does not exist—at least not one that will at the same time combine a low cost with the right technical attributes for broadband small-cell backhaul like low latency, frequency and phase synchronisation, reasonable capacity, and, in some instances, nLoS/NLoS capability for wireless backhaul.

Due to the diversity of the small-cell site acquisition challenges—negotiating with owners of street furniture, local councils, individual building or structure owners, etc.—it's rather unlikely that a single “universal” small-cell solution will emerge to become the small-cell solution of choice (*see the table*). Several different technical approaches will coexist in different markets, cities, and countries, depending on the local regulatory and legal frameworks.

Small-cell installation and maintenance costs also need to be further optimized to allow for installation and maintenance by less skilful subcontractors than highly skilled professionals who are used to working on macrocellular sites. The subcontractors who will deal with small cells in the future will likely be those tasked with maintenance of lampposts and street furniture. The lower installation and maintenance costs then will have to be enabled by easy-to-use, foolproof, self-configuring, self-diagnosing devices. We expect vendors to rise to this challenge, and it's only a matter of time until it will be addressed.

In terms of OPEX, then, site acquisition and installation primarily still stand in the way of widespread small-cell deployment. Operators will have to be smart and develop novel mass site acquisition and low-cost deployment models to surmount these challenges. A secondary yet temporary challenge involves self-configuring small-cell equipment.

As the first small-cell trials are being conducted, and there will be a strong need to deploy them in more developed markets, the first small-scale deployments this year will grow into widespread deployments in 2016 (Fig. 3). Depending on the small-cell deployment strategy, capacity-based or quality-of-experience (QoE) based, 100 to 400 cells per average city will be required. Obviously, large cities will need more, small towns less.



3. Average cities will require more than 400 small cells each by 2019, driven by three scenarios. Constrained deployment is limited by barriers of sites and backhaul, and it's only used to resolve critical issues in most congested macrocells. Capacity-only deployment uses small cells to ease the congestion of the busiest macrocells. And, wide QoE deployment enhances the QoE for most macrocells.

### Key Technical Attributes



Besides its low cost, a small-cell backhaul solution must offer low latency of 2 ms or better. It also must support accurate frequency and (in the future) phase synchronization. It must operate on low power, preferably Power over Ethernet Plus (PoE+). It must come in a small form factor, possibly integrated with the basestation, and be outdoor-capable. Its capacity must range from 40 Mbits/s up to 400 Mbits/s. Its Ethernet feature set must include virtual local area network (VLAN), quality of service (QoS), traffic shaping, and policing with an embedded L2 switch and at least two Ethernet ports for chaining. And, it must display a high degree of self-configuration.

Wireless small-cell backhaul presents its own requirements, including nLoS/NLoS capability as well as support of both unlicensed (low OPEX) and licensed (high dependability) bands. Of course, wireless small-cell backhaul comes with five key challenges.

First, wireless small-cell backhaul must provide sub-2-ms latency in case sub-6-GHz time-division duplex (TDD) solutions are used. Some products have addressed this challenge, such as DragonWave's Avenue Link Lite, but not most. More specifically, point-to-multipoint (PmP) solutions tend to seriously underperform in this area, which is a real issue for LTE quality of experience.

Second, synchronization must be suitable for LTE in case (sub-6 GHz) TDD solutions are used. Again, many TDD products on the market do not provide a proper solution.

Third, basestation integration is a challenge. Integrating both the antennas and the electronics in a non-adaptable way is not viable, as the optimal antenna alignment for the basestation and backhaul solution will rarely align in a pre-determined way. Two approaches are likely to emerge:

- Integration of microwave electronics into the basestation, with a separate MWR antenna (or vice versa, with an integrated microwave antenna and separate basestation antenna)
- Active beamforming antennas electronic (beam steering), most likely to appear in the millimeter-wave bands

Fourth, the capacity in the sub-6-GHz bands must suit nLoS/NLoS applications. In the millimeter-wave bands, capacity is not an issue, but LOS definitely is. As the bandwidth in the sub-6-GHz band is by default scarce, channels tend to be limited to 10, 20, or 40 MHz. Capacity in a single channel currently therefore peaks at about 230 Mbits/s aggregate (uplink + downlink). Furthermore, a limited number of channels is available. Several approaches are pursued to address this issue.

Moving from 64QAM to 256QAM modulation is the best way to increase capacity in this part of the spectrum, as there's ample link budget to make higher-order modulation work on the short hops

required for small-cell backhaul. Also, moving from 2x2 to 4x4 MIMO can double spectral efficiency. The downside will be the larger form factor for the antennas and steeply increased hardware cost. A bandwidth accelerator can provide low-cost lossless wire-speed data compression with an average capacity gain (actual gain depends on traffic mix) of 30% to over 100%.

Broader channels, such as 80-MHz channels, can help as well, though this may not always be feasible due to spectrum scarcity. Dual-band or multi-band microwave radio also can help. This may not always be feasible due to spectrum scarcity, but it enables the use of a licensed and an unlicensed band simultaneously, which renders the link more dependable in the face of possible (spurious) interferers in the unlicensed band.

And fifth, wireless small-cell backhaul must achieve a high degree of self-configuring capability, sometimes called self-optimized networks and plug-and-play. Possible features are:

- Automatic assignment of Internet protocol (IP) addresses for management purposes and automatic connection to a centralized network management system
- Automatic configuring based on intelligent algorithms and/or fetching configuration data from a central network management system database
- Self-alignment of the antenna, radio channel selection, modulation scheme selection, etc., and automated RF configuration
- Automated Layer 2 (Ethernet) configuration, such as VLANs and QoS

Some leading vendors have addressed sub-2-ms latency and synchronization suitable for LTE. The remaining challenges are still part of ongoing development activities and will be properly addressed over the next three years. Wireless backhaul technology will be an important enabler, certainly not an impediment to small-cell success.

## Conclusions

The concept of small cells is neither new nor revolutionary. Small cells have been around for decades. Unlike some 1G networks, though, they weren't really required in 2G and 3G networks.

The proliferation of bandwidth-hungry broadband multimedia devices like tablets and smartphones plus the popularity of bandwidth-hungry applications like photo sharing, video sharing, and video on demand is straining LTE networks and forcing operators to contemplate alternatives to traditional

macro-cell capacity extensions to cope with the data tornado. This is fuelling small-cell demand in areas with a high user density and limited spectrum assets.

There are still several technical challenges facing small-cell mobile backhaul, though none of them looks insurmountable. Based on current information and trends, then, we expect early, low-volume 4G small-cell deployments to commence in 2014—the year of trials—and a serious liftoff around 2016.

**Reference**

A voice codec (voice coder/decoder) is an algorithm employing data compression and data reduction techniques optimized for human voice data patterns that can, for instance, reduce a 64-kbit/s data stream fourfold to eightfold to 16 (full-rate codec) or 8 kbits/s (half-rate codec) without severely degrading voice quality in terms of mean opinion score (MOS).

THE CHALLENGES OF SMALL-CELL ACQUISITION			
	Access to location	Power availability	Backhaul
Sides of buildings	Fragmented ownership	Available from real-estate owners	Fixed broadband from real-estate owners
Street furniture	City halls, electricity companies	Available in lampposts, bus shelters	Fiber not naturally present; wireless required
FTTx infrastructure	Own or shared infrastructure	May not be available (PON avoids it)	Integrated (fiber)

**View this Article Online**

<http://electronicdesign.com/communications/understanding-small-cell-wireless-backhaul>