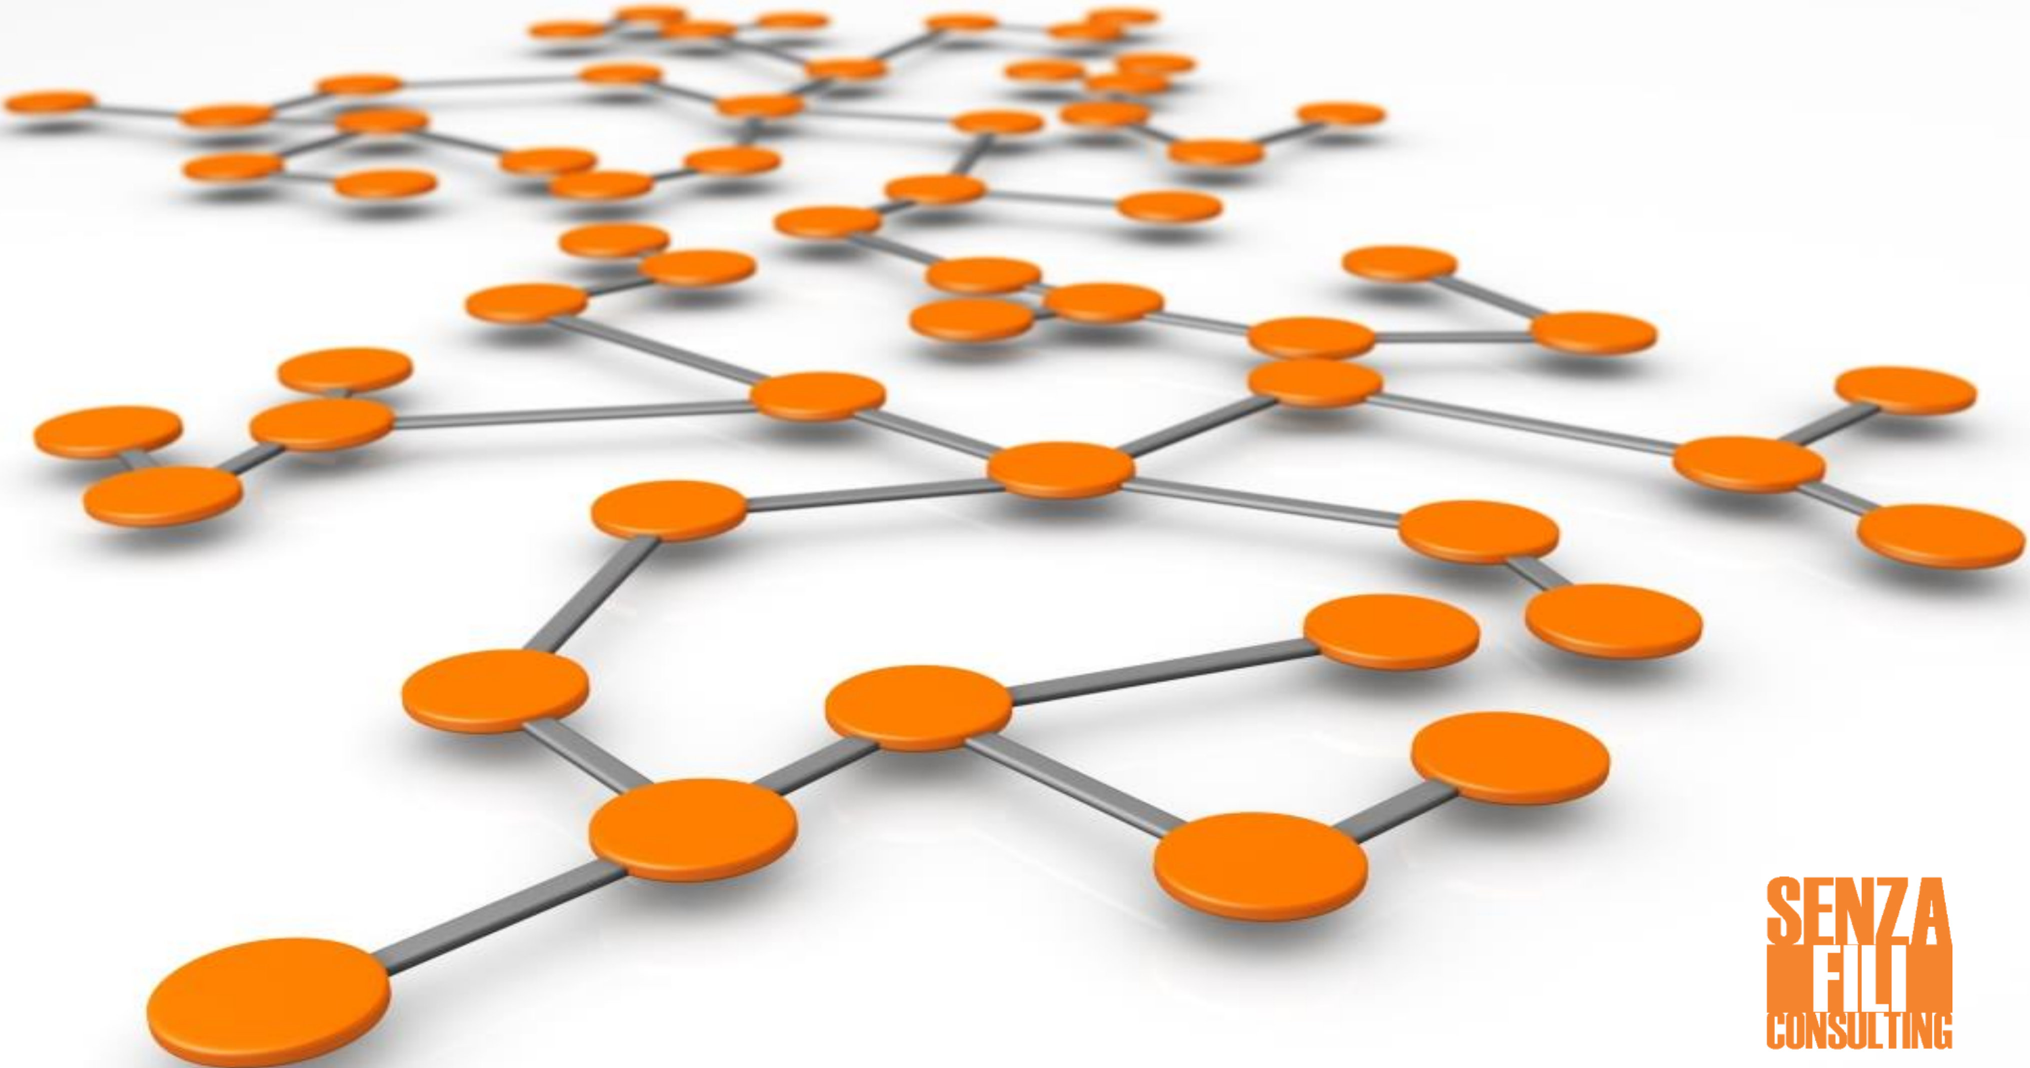

SENZA FILI CONSULTING

A backhaul solution for organic small-cell growth

A TCO analysis

By Monica Paolini, Senza Fili

Sponsored by

CCS

1. Introduction.
Small cells: A new type of backhaul

Densification of the mobile infrastructure has become a prominent – and unavoidable – way to address steep traffic growth that strains against limited expansion in spectrum availability. Operators need more capacity (i.e., Mbps) in their network, but because traffic distribution is not uniform across their footprint, the challenge they face is to increase the capacity density (i.e., Mbps per km^2) in areas of high traffic. RAN densification is necessary to meet this challenge, and small cells play an increasingly important role in it.

Small cells differ from their macro cousins in more than form factor and power consumption. They also change the way operators plan and deploy their networks. In their small-cell trials, operators quickly realized they could not simply put the small cells on just any lamppost or building corner and hope they would work. Most operators mention three specific challenges they face in deploying small cells that are markedly different from the challenges for macro cells: cell-site acquisition and planning, interference management, and backhaul.

Operators can choose among many mature macro-cell backhaul solutions. Yet most of these solutions are not well suited to small cells, because the small-cell environment is much more challenging: closer to the clutter, impossible for operators to control, and inherently dynamic. To complicate matters further, operators plan for a continuously growing and evolving small-cell layer that expands as traffic grows in an organic fashion. With small cells, the goal of operators is to increase the per-km^2 capacity – i.e., increasing depth of capacity rather than width of coverage, targeting those locations where traffic growth leads to congestion. Most mobile operators cannot afford to build massive small-cell coverage from the beginning, and instead plan to add small cells where needed, when needed.

Fiber remains an important element in the small-cell backhaul aggregation, but in most cases it is neither cost effective nor well suited to meeting the ever-changing

Requirements that make small-cell backhaul challenging

- Small form factor
- Low cost
- Low energy consumption
- Easy, fast install
- Low operation and maintenance costs
- Operating in a changing environment, largely outside the operator's control
- Flexible, adapts to organic network growth and automated topology optimization
- Compliant with municipal council planning requirements, e.g., for design aesthetics, weight limits for lampposts, health and safety regulations
- Protected against vandalism, security and hacking threats

Requirements in common with macro-cell backhaul

- High capacity and low latency
- Reliability
- Standards compliance (ETSI, FCC)
- Network management requirements

network and environment requirements of small-cell deployments. Operators need equipment that is low cost and fast to install, and that has high capacity, low cost of operation and maintenance, and a small form factor acceptable to municipal planners. It is a tall order to meet all these requirements at the same time, and in the last few years, this has been one of the most active areas of innovation and startup activity in the backhaul area.

While the split between wireless and wireline backhaul will vary across operators, we expect wireless backhaul to dominate in small-cell sites, both for cost considerations (e.g., high installation costs and/or high recurring costs of fiber), and for operational considerations (e.g., difficulties in bringing fiber to structures like lampposts, and lack of flexibility). For small cells with wireless backhaul, two solutions have gained the most support to date:

- NLOS solutions in the sub-6 GHz bands, which mostly use PMP architectures. Cheaper and easier to deploy, most NLOS solutions have limited capacity and require licensed spectrum that is difficult to obtain.
- LOS solutions in the V/E millimeter bands (60 GHz and 80 GHz), using PTP links. These can be more expensive, but they provide more capacity, and spectrum is available in many markets. The need to add relays to compensate for lack of LOS can increase deployment costs. Deployment is slower and more expensive due to alignment requirements and frequent visits to sites for realignment, needed when the operator adds new small cells to the network.

To address the specific challenges of small-cell backhaul, new solutions are gaining traction. CCS, one of the new players in this space, has developed a new approach to small-cell backhaul that gives operators more flexibility as they increase the density of the small cells in their deployments.

With a self-organizing mesh solution that uses point-to-multipoint (PMP) spectrum in a multipoint-to-multipoint (MP2MP) architecture, CCS enables operators to easily add new small cells within their HetNet deployments. Nodes in CCS's architecture use a mix of PTP and PMP links that auto-configure and change dynamically depending on backhaul terminal location and traffic load. The backhaul unit in a new

CCS: A multipoint-to-multipoint microwave small-cell backhaul system with full self-organizational capability

The CCS MP2MP solution combines PTP and PMP functionality to move beyond the tradeoffs that these two types of solutions impose:

- While it uses LOS microwave spectrum (e.g., 26, 28, 32, 38 and 42 GHz), it uses only a single channel, with dynamic allocation of spectrum to optimize capacity, giving operators more flexibility in establishing and expanding their small-cell networks.
- Self-organizing mesh features simplify installation and maintenance, lowering capex and opex. As operators add new small cells, the backhaul network can automatically reconfigure itself. As the traffic patterns change, the network can automatically optimize the topology of backhaul connections to balance traffic loads.

Form factor and equipment costs per link for the CCS solution are comparable to those for V/E millimeter-wave PTP solutions. However, because of the 270-degree visibility, only one unit is required per lamppost, whereas a V/E-band solution in a daisy chain or ring topology will require multiple radios. Capacity ranges from 480 to 960 Mbps in a 112 MHz channel, with a latency of under 125 µs.

CCS has designed its solution for fast and simple installation (15 minutes per terminal) that does not require telecom-engineering skills. Each local network partition includes multiple cell sites that are connected to a controlling node, either directly, through other sites or relays – and, typically, through more than one link, to provide some redundancy. Each cell site needs only one unit, with the same form factor across site types (i.e., controlling node, relay or edge cell site). Terminals do not require alignment during installation or subsequent network changes or optimization, because the units self-configure and self-heal.

small cell automatically connects to the backhaul network, which reconfigures itself to optimize transmission each time a new small cell is installed.

The auto-configuration capabilities and flexibility in topology can create cost savings for operators, because they reduce effort for both installation and operations, and because they do not require backhaul RF planning or manual realignment to accommodate new small cells within the footprint.

In this paper, we look at whether MP2MP solutions have a cost advantage compared to PTP links in the 60 GHz and 80 GHz millimeter bands. We present the results of a TCO analysis, based on data assumptions from CCS, to compare the cost drivers for PTP and MP2MP solutions for small-cell backhaul and, hence, to understand how these solutions may address the needs of mobile operators.

2. TCO model assumptions.
Deployment models for MP2MP and V/E-band PTP

The TCO model we present in this paper assumes a gradual deployment (shown in the figure at right) in which an operator expands the number and density of small cells over five years, starting with a limited number of small cells in a new market. The model, however, can also estimate the TCO for larger deployments and produce directionally consistent results, because the TCO dynamics scale linearly with the number of small cells deployed.

The figure below illustrates how an operator may add new small-cell sites over three years, and how this organic growth may change the role of backhaul nodes, and the links among them, within the backhaul network.

We track the deployment's growth within the same footprint over the five-year period, as the operator increases the density of small cells to meet the increase in traffic load. The deployment model has three elements:

Small-cell and backhaul equipment

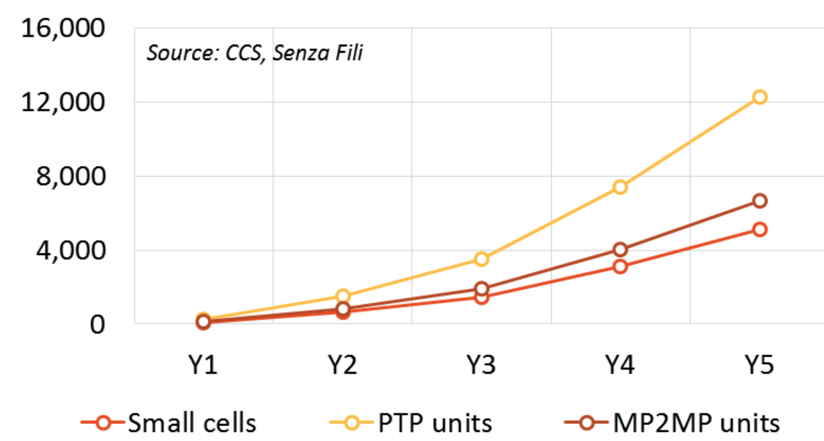

- **Controlling nodes**, which aggregate the traffic from the local backhaul and from which a fiber (or leased-line) connection carries the traffic to the core network.
- **Relays**, to connect small-cell sites that are not in LOS with any other element in the network. They are in a location that has LOS to both the unconnected small-cell sites and a connected element (controlling node or backhaul site) in the backhaul network.
- **Backhaul sites**, each co-located with a small cell. Backhaul sites send their traffic to another backhaul site, a relay, or a controlling node, with the choice depending on availability of LOS and distance. Some backhaul sites aggregate traffic from the co-located small cell and other small cells connected to them by a backhaul link.

As the network grows, the density of small cells – and hence of backhaul elements – increases, and their role and position within the network changes. A backhaul site at the edge may receive additional traffic from a new small cell. Another site may switch its connection to a newly installed controlling node.

The ability of the network to reconfigure itself in response to the installation of new small cells within the same footprint is not strictly necessary, but it makes it possible to optimize the local backhaul network in response to traffic loads and cell location. In the figure at right showing the network evolution over three years, if there were no network reconfiguration, links would be longer (and hence less likely to have LOS or to operate at a maximum throughput), and the addition of new controlling nodes would be less effective.

In a PTP backhaul network, changes in topology are possible, but come at an additional cost. The network needs to be redesigned for frequency-channel allocations and interference management, the equipment needs to be manually realigned based on the updated RF plan, and, if using licensed spectrum, the rights to more PTP links must be purchased.

In an MP2MP network, the addition of new elements does not require any change to the existing elements. In most cases, there is no need to acquire new spectrum, because the license grants spectrum use over a geographic area. The new element

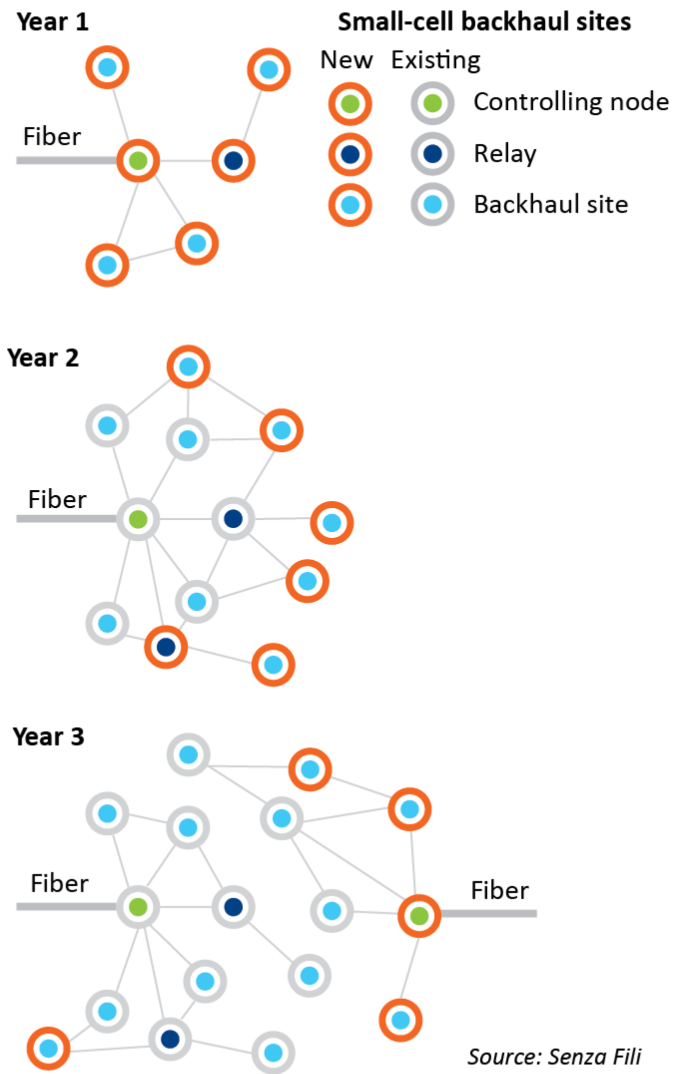

Source: Senza Fili

automatically detects the other network elements that it can connect to and picks the best-suited one, avoiding disruption in network performance.

The deployment model's assumptions for PTP and MP2MP are the same: a local network with 10 small cells, one of which is co-located with the controlling node, and three relay nodes, assuming that 30% of the small-cell locations are not in LOS with another network element. The controlling node is the only element connected to the rest of the network via fiber or leased line.

As a result, in the PTP scenario, the operator needs 24 backhaul terminals, because each of the 12 links requires two terminals. In the MP2MP scenario, the operator needs 13 terminals, one per link plus one in the controlling node. The graph at the beginning of this section shows the equipment requirements for the network we use as the base case in this paper. The different number of backhaul terminals does not noticeably affect the capex for the equipment, because we assume MP2MP terminals to be more expensive, but fewer are required. As a result, the equipment cost per link is approximately the same for both the PTP and MP2MP cases. In the PTP scenario, we include a PoP aggregation point to aggregate traffic and manage QoS. The higher number of terminals in the PTP case, however, results in higher installation costs and opex, because the operator has to install, maintain and operate more equipment than in the MP2MP case.

The number of deployed small cells grows from 100 to 5,110 by Year 5, in line with rates estimated by operators and small-cell market forecasts. Newly deployed small cells account for 65% of the total of small cells by Year 5.

The cost assumptions for the deployment model are from CCS and based on data reviewed with mobile operators. For simplicity, the model calculates the average cost on a per-small-cell basis.

The model assumes the average capex to install the backhaul in a small cell to be $3,401 for MP2MP and $5,820 for PTP, a cost saving of 42% with MP2MP. We assume equipment costs of $2,000 per link for the MP2MP solution, and $1,000 per terminal (or $2,000 per link) for the PTP solution. In a 10-site network, this translates

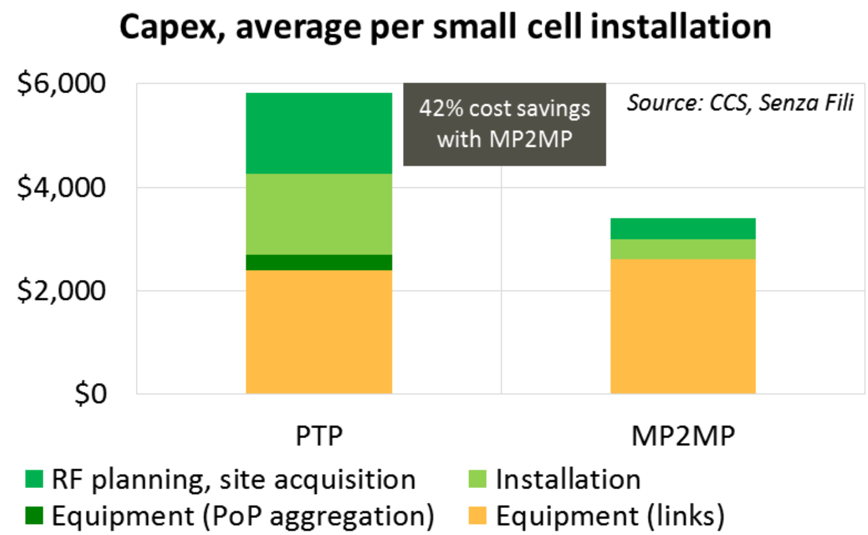

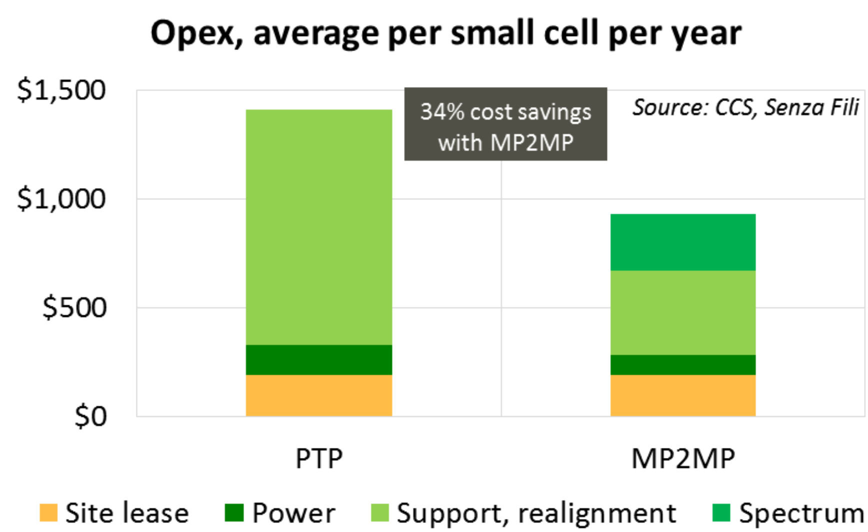

to a per-site equipment cost of $2,600 for MP2MP and $2,700 for PTP (this figure includes PoP aggregation equipment costs estimated at $300 per small cell).

The bulk of the savings come from the lower cost of planning and installing the small cells with MP2MP, due to the faster installation enabled by network self-organizing functionality.

The annual average opex per small cell is projected to be $932 for MP2MP and $1,408 for PTP, a 34% cost saving with MP2MP. The backhaul costs and site lease costs are the same. Power is approximately a third more expensive for PTP (because of the higher number of equipment elements in a PTP architecture).

No spectrum costs are included for PTP. We assumed that PTP spectrum in the E-band and V-band does not require any fee (although use of the spectrum may require a license). In the MP2MP scenario, we include spectrum costs as an opex item, with a per-link rental fee of $200 per year (or $236 per small cell, to account for the additional spectrum cost incurred by relays). In the model, spectrum fees account for 19% of the small-cell opex, which is a conservative figure likely to overestimate spectrum fees in most countries; it represents the high-end cost assumptions provided by mobile operators.

We could have treated MP2MP spectrum costs as capex. In most markets, spectrum rights in microwave bands reserved for PMP are auctioned on a per-region basis (instead of a per-link basis) relatively inexpensively. However, operators have to amortize the cost over time and over a reasonably large deployment. The model described here looks only at a single market within a wider network deployment, and it is not appropriate to allocate the entire spectrum costs to a small part of the network. The per-link spectrum fees enable us to incorporate spectrum costs more accurately. This approach, however, assumes that the operator is indeed deploying a larger network over which it can amortize spectrum costs effectively, or that the PMP spectrum holder leases the spectrum on a per-link or per-subregion basis, as is the case in markets such as the US. Per-link leases of PMP spectrum are beneficial to the MP2MP business case, because an MP2MP solution is typically less cost-effective

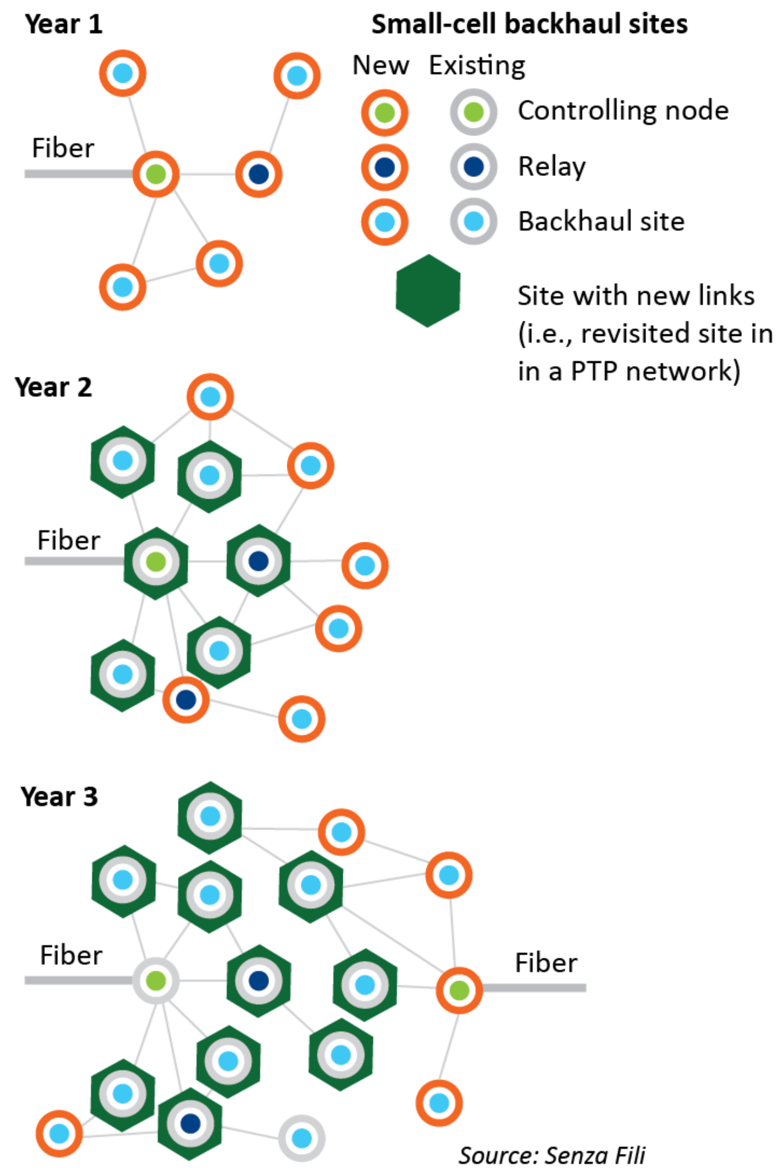

Source: Senza Fili

for an operator that plans a limited small-cell deployment if it has to acquire spectrum rights on a regional basis.

The opex in the TCO model also includes the costs associated with the organic growth and evolution of the network. As the number and density of small cells increases within the same footprint, the operator has to reconfigure the local backhaul network to link new small cells to a controlling node, and change the links among existing cells to optimize the networks and to accommodate the introduction of new controlling nodes. In a PTP architecture, this requires a revision of the RF plan and a realignment of the backhaul links, estimated to cost on average $480 per cell site per year. In an MP2MP architecture, these adjustments come for free, as the network automatically optimizes the local backhaul connectivity and realigns the newly established links accordingly. The figure above shows the cell sites with new links – i.e., those revisited in a PTP deployment – as a result of the expansion of the small-cell network in the example used before.

Five-year TCO

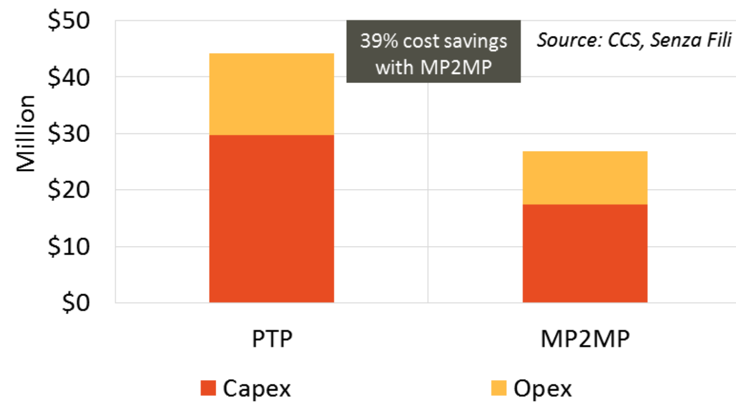

Capex: MP2MP cost savings breakdown

Source: CCS, Senza Fili

3. Five-year TCO analysis. MP2MP is cost effective where small-cell deployments grow organically

The financial model estimates that the TCO to deploy and operate an organically evolving backhaul network – growing from 100 small cells in Year 1 to 5,110 cells in Year 5 – is $44.1 million with a PTP solution in the V or E band, and $29.9 million with an MP2MP solution in the 28 GHz band. This means a cost saving of 39% for the MP2MP scenario, or $17.2 million over five years. By Year 5, the annual cost to operate and expand the network is $18.8 million for the PTP solution, and $11.5 million for the MP2MP solution.

The MP2MP savings over PTP are greater for capex (42%) than for opex (34%). Since capex is nearly the same percentage of TCO in both scenarios (67% of TCO in the PTP scenario, and 65% in the MP2MP scenario), capex saving accounts for most (72%) of the overall cost savings ($12.4 million out of $17.2 million).

Capex savings are due to MP2MP's lower planning and installation costs. Equipment costs account for 41% of PTP capex and 76% of MP2MP capex, but the total amount is approximately the same in the two cases, as discussed above.

Self-configuration in MP2MP networks removes the need to align backhaul terminals and leads to a 74% cost reduction for RF planning, site acquisition and installation over the PTP scenario, mostly due to the faster installation. As a result, installation costs represent 11% of the MP2MP capex, but 27% of the PTP capex.

Among opex costs, site lease costs are the same for both solutions. Power is a limited source of cost savings: it accounts for only 7% (MP2MP) and 8% (PTP) of opex. A more substantial source of savings is maintenance and realignment costs: they contribute 41% (MP2MP) and 77% (PTP) of opex. The TCO model estimates a 34% cost savings for MP2MP maintenance and realignment.

Spectrum costs lower the overall opex savings in the MP2MP case. We anticipate V and E spectrum to be available in a license-exempt or lightly licensed environment, in which spectrum can be treated as free. In the model, we treat spectrum costs as an opex item that accounts for 28% of MP2MP opex.

If the operator does not own a spectrum license and has a lease arrangement with the spectrum owner, spectrum costs are an opex item. If the operator has a regional or countrywide license, spectrum costs are a capex item whose impact on the TCO will depend on the number of backhaul links that use the licensed band and on the cost of the license. In our model, we treat spectrum costs as opex items so that their impact on the TCO linearly scales with the number of links. This allows us to distribute spectrum costs fairly, regardless of the network size. However, this approach does not capture the likely cost savings that derive from deploying a large number of small cells using MP2MP backhaul.

The TCO model's assumptions represent median values for a developed market, and they are not specifically based on a mobile operator. However, there are substantial

Sensitivity analysis: Spectrum annual costs

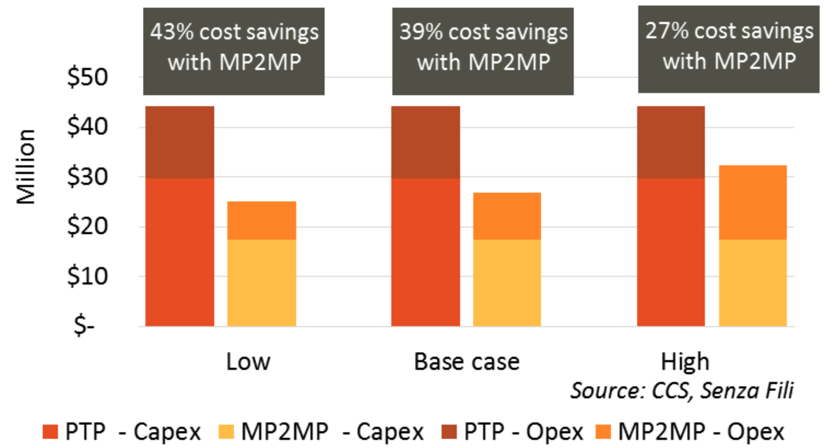

Source: CCS, Senza Fili

Sensitivity analysis: CCS estimates of MP2MP savings

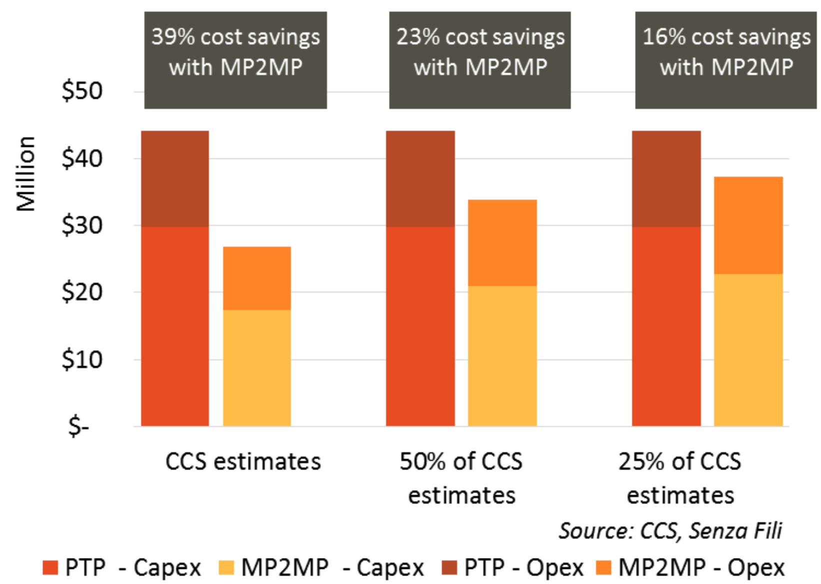

Source: CCS, Senza Fili

differences across markets and, within the same market, across operators – mostly due to differences in spectrum regulation and in staff and operating costs. Some changes across markets have a comparable impact on both PTP and MP2MP solutions.

As described above, however, spectrum costs affect the two solutions differently, because only in the MP2MP case do we include spectrum fees. To assess the impact of spectrum costs, we compared scenarios in which spectrum annual costs vary from $67 to $600, compared to $200 per year per cell site in the MP2MP base case. Spectrum fees are highly variable across countries, or change as a result of selecting a different band or getting spectrum access from a wholesale provider. In the low-cost scenario, the cost savings with MP2MP are higher (43%) than in the base case (39%), because the lower spectrum fees of the MP2MP scenario decrease the opex. Conversely, high spectrum fees result in a smaller (27%) cost advantage for MP2MP, but demonstrate that the MP2MP solution is cost effective even if operators face higher spectrum prices than anticipated by the TCO model's base case.

The cost savings afforded by MP2MP depend on cost data for installation, operation and maintenance from CCS's operator trials. What is the impact on the TCO if we lower the estimated cost advantage? In a second sensitivity analysis, we looked at the impact on cost savings if we reduced the CCS advantage to a half (50%) and a fourth (25%) of the CCS estimates used in the base case. The cost savings for MP2MP dropped from 39%, in the base case, to 23% if we assume half of CCS cost savings, and to 16% if we assume a fourth of CCS cost savings. Consequently, on all tested scenarios there are significant cost savings resulting from the implementation of a MP2MP backhaul solution that is fully self-organizing.

Implications: Cost dynamics of backhaul in evolving small-cell networks

Small cells will be crucial to enabling the increase in capacity density in mobile networks that operators will need over the coming years, but they will do so as part of a well-paced evolution. Operators plan to add small cells gradually, and only in those locations where they see traffic congestion.

As the number and density of small cells grows, operators will have to reshape the backhaul networks serving small cells. The ability to dynamically and cost effectively change the topology of backhaul links is a key benefit for operators deploying multilayer HetNets with small cells.

CCS has developed its MP2MP solution to address this challenge by using a mesh architecture that optimizes the backhaul connectivity in a local small-cell network that continues to evolve and expand through time.

The TCO analysis shows that the lower costs to install, operate, and expand a backhaul network that automatically configures and heals itself may result in a 32% cost savings in capex and opex over a five-year period. The cost savings depend on the manpower-driven capex and opex items – the reduced need to operate equipment, and the ability to install it in a shorter time and to reconfigure it automatically.

Glossary

ETSI	European Telecommunications Standards Institute
FCC	Federal Communications Commission
HetNet	Heterogeneous network
LOS	Line of sight
MP2MP	Multipoint to multipoint
NLOS	Non line of sight
PMP	Point to multipoint
PoP	Point of presence
PTP	Point to point
QoS	Quality of service
RAN	Radio access network
RF	Radio frequency
TCO	Total cost of ownership

About Senza Fili

Senza Fili provides advisory support on wireless data technologies and services. At Senza Fili we have in-depth expertise in financial modelling, market forecasts and research, white paper preparation, business plan support, RFP preparation and management, due diligence, and training. Our client base is international and spans the entire value chain: clients include wireline, fixed wireless, and mobile operators, enterprises and other vertical players, vendors, system integrators, investors, regulators, and industry associations. We provide a bridge between technologies and services, helping our clients assess established and emerging technologies, leverage these technologies to support new or existing services, and build solid, profitable business models. Independent advice, a strong quantitative orientation, and an international perspective are the hallmarks of our work. For additional information, visit www.senzafiliconsulting.com or contact us at info@senzafiliconsulting.com or +1 425 657 4991.

About the author

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.

CCS – The first self-organizing small cell microwave backhaul

A new approach: self-organizing and optimizing

CCS has developed the world's first small cell microwave backhaul system with self-organizing capabilities. It is the only solution that meets all key requirements for high capacity, reliability and low latency, with rapid deployment and simple, low-cost operation in a small, low-impact design.

Self-organizing, self-healing links automatically reconfigure themselves to optimize performance across a resilient, multipoint-to-multipoint topology. Easy to scale, with no radio planning or re-alignment required as nodes are added, it delivers the lowest TCO for any small cell backhaul technology.

CCS's system enables small cell deployment in a flexible, organic way as additional capacity is required, maintaining a quality of service that meets customer expectations.

Commercial deployments

ETSI-compliant and CE-marked, the first commercial implementation of the technology is now taking place with China Mobile in the first deployment of small cells in China. CCS has also selected China for its manufacturing base to meet demand in what is expected to be the world's largest market for small cell networks.

CCS is also conducting commercial pilots with other global Tier 1 mobile operators, which are expected to lead to phased commercial deployment over the coming months.

About CCS (www.ccsl.com)

Established in 2010, Cambridge Communication Systems (CCS) is a specialist in small cell microwave backhaul systems. The team has a successful track record in creating innovative wireless solutions and its founders – CEO Steve Greaves and CTO John Porter – previously co-founded Adaptive Broadband and Cambridge Broadband Networks.

Based in Cambridge, Europe's most successful technology center, the company draws on the city's rich pool of world-class engineering talent. In 2012, it won the Cambridge Business Weekly award for Best Start-Up of the Year.