White Paper

Design Strategies to Reduce TCO of Wireless Backhaul: The DragonWave Example

DragonWave

www.dragonwaveinc.com

NOTICE

This document contains DragonWave proprietary information. Use, disclosure, copying or distribution of any part of the information contained herein, beyond that for which it was originally furnished, requires the written permission of DragonWave Inc.

The information in this document is subject to change without notice and relates only to the product defined in the introduction of this document. DragonWave intends that information contained herein is, to the best of its knowledge, correct and accurate. However, any/all liabilities associated with the use or accuracy of the information contained herein must be defined in a separate agreement between DragonWave and the customer/user.

TABLE OF CONTENTS

This paper discusses a combination of innovative design strategies that can help reduce the TCO of wireless back by as much as 60-80%, depending on the network circumstances. The paper illustrates how these design strategies have been implemented in the DragonWave Horizon Line of products, enabling short ROI and sustained profitability for cellular carriers and network operators.

1. Introduction

In the previous paper, “Understanding the TCO of Wireless Backhaul” we illustrated that wireless microwave is the most cost-effective solution for backhaul when compared to fiber and copper. However, high capacity wireless backhaul is nonetheless still a serious cost challenge facing wireless service providers today, taking up a large portion of the total network lifecycle costs.

We argued that when planning a network expansion or a new build it is important to analyze all aspects of the network cost in detail rather than focusing on the popular, but diminishing, importance of the capital costs (CAPEX) number. We discussed that albeit being more economical than copper and fiber, wireless backhaul is posed with certain technical and logistical challenges and without a proper design and planning it will not be as cost-effective as its potential. We analyzed a breakout of the backhaul TCO based on a ten-year network model and illustrated that the CAPEX costs of the network architecture take up only a small part of the total cost whereas the operating costs (OPEX), combined together, comprise the majority of all costs in the backhaul TCO. We further demonstrated that in order to get a significant improvement in the TCO, the antenna/tower and site leasing costs need to be addressed first of all, as these costs comprise more half of the TCO.

It follows then that as more and more capacity is added, amid rising costs per bit and declining Average Revenue per User (ARPU), the profit margins of cellular carriers are not going to increase if backhaul costs are not restrained. In this paper, we will discuss a set of specific design strategies that help reduce the TCO of wireless backhaul by a considerable margin, helping operators turn backhaul revenues intro profits.

We will demonstrate that by deploying mesh network architecture, reducing antenna size and maximizing spectral efficiency via advanced modulation techniques along with the integrated antenna approach it is possible to reduce the backhaul burden on the TCO by more than half.We will also demonstrate that software-controlled scalability and flexible bandwidth pricing help decrease the backhaul costs even further.

Additionally, we will demonstrate that the dramatically reduced TCO can go hand in hand with optimized performance and outstanding capacity, comprising a highly scalable backhaul architecture that can optimally carry all services and support unpredictable user demands. For example, if one is to use Dragon Wave’s Horizon Duo on a Dual Polarity Radio Mount (DPRM), it is possible to support an astounding 1.6 Gbps of full duplex bandwidth per link, which is twice as much as what can be supported in a typical wireless link.

We will demonstrate how these and other innovative design strategies have been implemented in the DragonWave Horizon Line of products, paving the way to a short ROI and sustained profitability for cellular carriers and network operators.

2. Reducing OPEX

Before we begin discussing specific design strategies to reduce the wireless backhaul TCO, let us briefly review the TCO model, which was analyzed in the first paper.

The TCO of wireless backhaul consists of OPEX and CAPEX and includes the following components: monthly tower/antenna lease costs, monthly space/site and power costs, spectrum license costs, installation, maintenance and management costs, and money costs (OPEX); microwave equipment and Ethernet switching costs (CAPEX).

With the exception of the tower/antenna as well as space/site and power costs that are incurred on a monthly basis, all other costs are either a one-time final fee or a one-time fee to be incurred every 10 years (e.g.

spectrum license). Figure 1 demonstrates the breakout of the TCO components for a typical network in North America and Western Europe over 10 years.

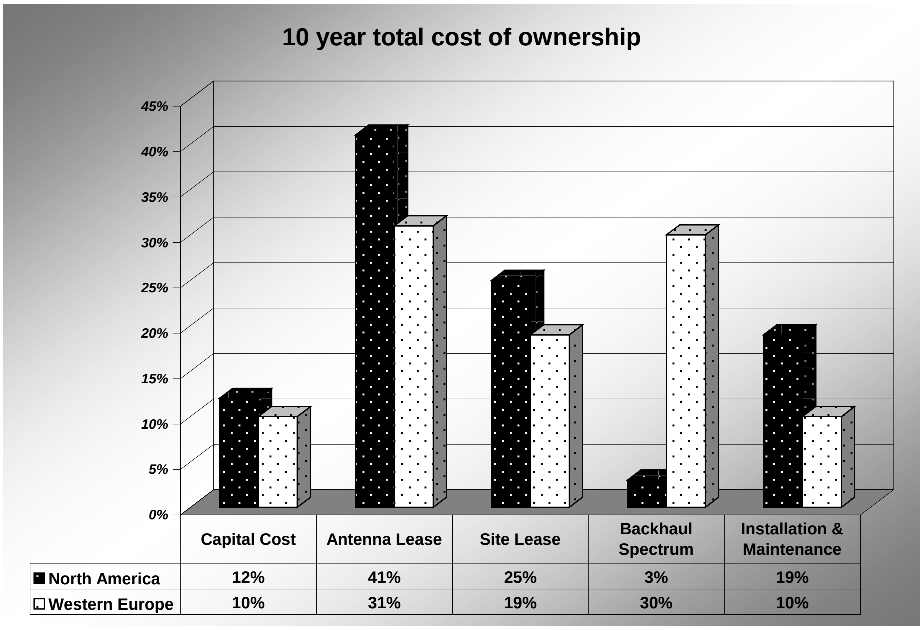

Figure 1 – 10-year Cost of Ownership Components in North America and Western Europe

It becomes very clear that if we want to reduce the overall network cost addressing equipment cost will only have a small effect on the total costs. To get a significant reduction in the total costs, the antenna/tower and site/space monthly leasing costs need to be addressed.

2.1 Reducing Monthly Lease Costs

Reducing monthly lease costs is easier said than done, thanks to the growing demand for wireless deployments and the associated rise in leasing costs. The only way to reduce these costs is to plan on reducing the amount of leases that are required within the network. There are two monthly lease areas that constitute more than half of the TCO: antenna/tower leases and space/site leases.

2.1.1 Reducing Antenna/Tower Lease Costs

Reducing Antenna Costs via Adaptive Modulation

The antenna lease costs are proportional to the total real estate consumed on the tower or building and comprise approximately 40% of the TCO, being the largest individual expense in a network deployment. Achieving considerable reduction in the antenna lease costs will help drive the overall lifecycle costs down by a large margin.

Antenna lease costs can be cut in half by reducing antenna sizes by 50%. Antenna sizes can be reduced, without sacrificing availability, by using high power systems or by using adaptive modulation techniques.

Automatic adaptive modulation (AAM) allows maximizing spectrum usage and increasing the capacity of a given link budget in wireless backhaul links, which allows for a minimization of antenna sizes and spectrum consumption. AAM automatically provisions a bandwidth on a given link according to link conditions, while maintaining connectivity during rain fade events (at reduced capacity). Varying the modulation varies the amount of bits transferred per signal, which enables higher capacity and better spectral efficiencies. For example, 256-QAM modulation can deliver approximately four times the capacity of QPSK.

The DragonWave Horizon products support modulation schemes ranging from QPSK to 256 QAM, which are software selectable at any time. With DragonWave's AAM, operators can use up to 50% smaller antenna sizes while achieving higher capacities per link and significantly maximizing performance of existing infrastructure.

Reducing Antenna Costs via Mesh/Ring Network Architecture

Supplementing AAM, mesh/ring architectures also facilitate reduction in antenna sizes. Wireless mesh/ring network can allow an operator to engineer a network to 99.999%, while engineering the individual links to a slightly lower availability, e.g. 99.99-99.992%. This is enabled by the alternate path capabilities of the ring. The reduction in per-link availability requirements allows using smaller antenna sizes while maintaining a "five-nine" end-to-end network availability. In addition, given that rings distribute equipment and path through the network, the average path length is shorter than in a hub-and-spoke configuration, which also helps reduce the required antenna sizes.

A combination of adaptive modulation with mesh/ring architecture allows the operator to significantly reduce the average antenna size, which results in a corresponding reduction in the antenna lease costs.

In this context it is important to note that mesh/ring architecture offers a more robust network protection mechanism than hub-and-spoke topology. The operator can take advantage of the fact that equipment failures and rain fades are not correlated to effectively increase the available bandwidth (or reduce the antenna size at the same bandwidth and availability targets). During normal operation traffic can be routed both ways around the ring with the full link capacity enabled by AAM delivering double the deign bandwidth for best effort services. During a rain fade, the capacity on one direction is reduced, but the alternate ring direction is still available, and the operator can deliver approximately 1.5 times the design bandwidth. During equipment failure the alternate path still has the full link capacity available and 100% of the design bandwidth is available. The impact of a mesh/ring architecture deserves an individual discussion, we will talk about it in a separate section called Using Mesh/Ring Architecture, pp. 9.

Integrating Access and Backhaul Antenna Panels

Another way to reduce antenna lease costs is to integrate RAN antenna panels together with backhaul antennas to form a single multi-beam antenna. The resulted integrated antenna, though larger in size than the original RAN antenna, eliminates the need for a second antenna lease. The small form factor of the DragonWave product line enables a smooth integration of backhaul antennas with RAN antennas, which helps to reduce the total lifecycle costs by a large margin.

By successfully deploying AAM, using mesh/ring architecture as well as adopting integrated antenna solutions, DragonWave has significantly reduced monthly antenna lease costs while at the same time significantly increasing system capacity and reliability.

2.1.2 Reducing Space Lease Costs

Next in significance to monthly antenna leases stand monthly site/space lease costs. The site lease costs can be significantly reduced or even eliminated by utilizing all-outdoor equipment.

Zero Footprint Impact

All-outdoor equipment eliminates the requirement for indoor microwave equipment, which reduces footprint and decreases monthly site lease costs.

Having microwave equipment permanently placed outdoors, i.e. on a roof or on a tower, offers an additional benefit of simplified cabinet wiring, which helps reduce spending even further.

Outdoor switch and pseudowire housings combined with battery backup and lightning protection can eliminate the need for any indoor rack space at all sites. With energy prices climbing to record heights, if power costs were to increase further, using indoor equipment will lead to higher Heating Ventilating Air Conditioning (HVAC) costs which would lead to an increase in the total monthly site/space lease costs.

DragonWave has been recognized in the industry for its all-outdoor licensed microwave equipment, which has been enjoying successful sales worldwide, making it easier for operators to reduce costs. The table below illustrates the range of outdoor and indoor options available with the DragonWave products.

Product	All Indoor	Indoor/ Outdoor	All Outdoor
Horizon DUO	x	x	
Horizon Compact	x		x
Air Pair	x	x	x
Service Delivery Unit (SDU)	x		x

Deploying an all-outdoor microwave backhaul system provides 20% reduction to the TCO while deploying outdoor switches provides another 5% reduction.

Graphical Illustration of Antenna/Tower and Site Cost Reductions

Let us graphically illustrate the effect of reducing antenna lease requirements and site lease requirements on the network costs. As shown in Figure 2, by reducing antenna sizes and providing RAN antenna extensions to include backhaul antennas, network cost can be reduced by 35%. Deploying an all-outdoor microwave backhaul system provides an additional 20% reduction, and deploying switches outside provide another 5% reduction, for a total of a 60% reduction compared to the current network cost.

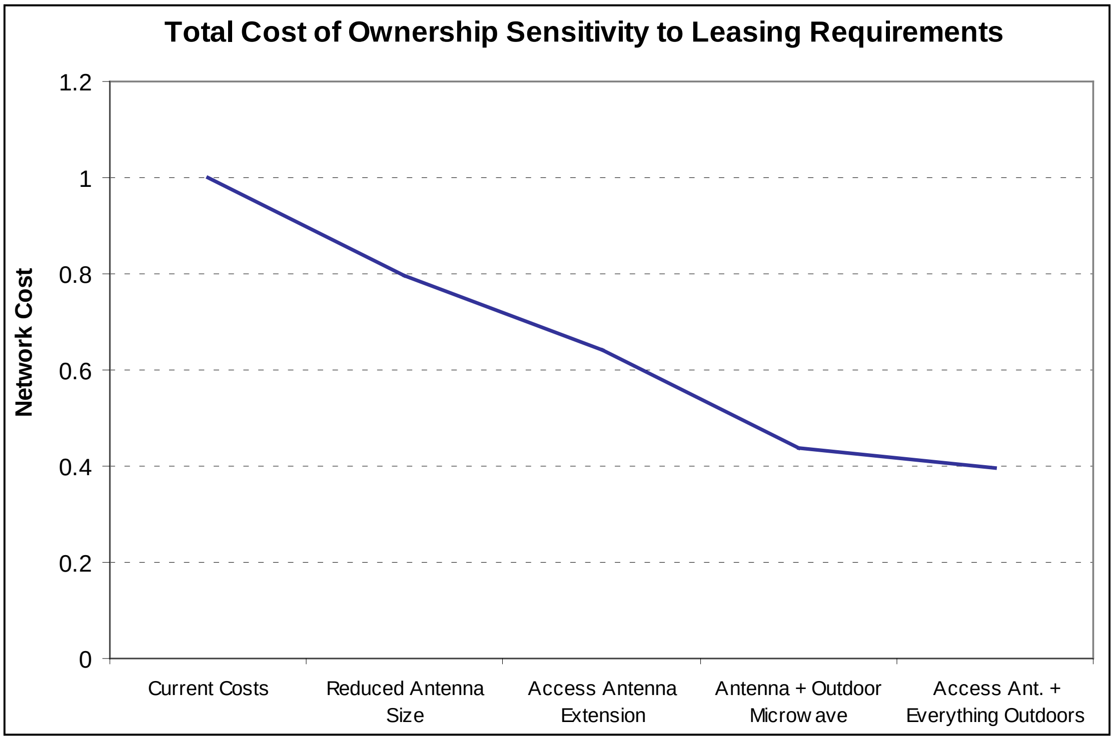

Figure 2 – Cost of Ownership Sensitivity to Leasing Requirements

2.2 Reducing Spectrum Lease Costs

Spectrum licenses authorize a licensee to use specified radio frequencies within a defined geographic area. The spectrum lease costs vary greatly by geographical regions. In North America where the spectrum lease costs represent a one-time flat fee they are less of an issue than in Western Europe and other locales where the spectrum costs are calculated by use and can be as much as an equivalent of $30,000 for a period of 10 years. Furthermore, the spectrum costs are also affected by whether the pricing is set on a per link or area channel basis. In North America (except Mexico), the spectrum costs are calculated link by link with no regard to channel size. In Europe, spectrum lease costs are calculate on per link basis, but with direct correlation to the amount of spectrum used on that link, so it becomes very important to minimize the spectrum requirements in order to minimize the licensing costs. In many other regions, such as South Asia and Latin America, the spectrum is leased on a channel-by-channel basis across a given region. In these regions, minimizing the number of channels used can significantly reduce the spectrum costs. In these cases, increasing spectral efficiency in smaller and cheaper channels becomes of critical importance.

2.2.1 Increasing Spectral Efficiency in Smaller Channels

The principal way of increasing spectral efficiency is by using higher modulation techniques. The higher the order of modulation the larger the cost savings. Traditional systems have supported a maximum of 16 QAM in 7/10 MHz channels sizes. Even new Ethernet systems support high modulation in large channel sizes, but are still limited to 16 QAM in small channel sizes. DragonWave's Horizon Compact and Horizon Duo address this issue by supporting up to 128 QAM in 7/10 MHz channels and 256 QAM in 14 MHz Channels.

Figure 3 below shows the throughput versus channel size of the various technologies.

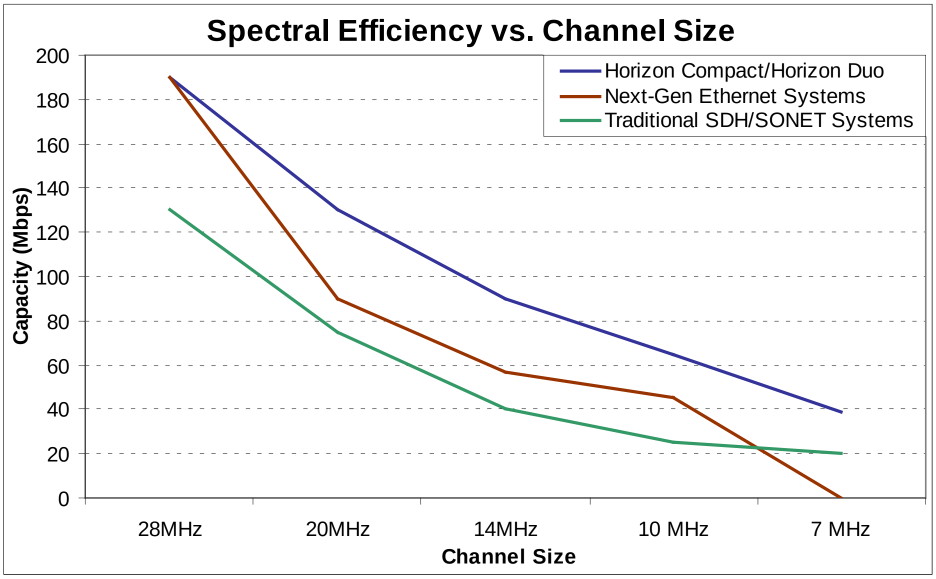

Figure 3

As shown in the figure above, Horizon Compact and Horizon Duo provide about a doubling in capacity at low channel sizes, due to the support of 128 QAM in 7 MHz and 256 QAM in 14 MHz channels.

2.2.2 Using Mesh Ring Architecture

Furthermore, the choice of network architecture also affects spectrum lease costs. For example, the monthly spectrum lease costs are adversely affected by the hub-and-spoke architecture in those regions where these costs are established on a channel-by-channel basis: Pakistan, India, Mexico, Latin America, and others. One of the downfalls of the hub-and-spoke architecture is its propensity to congestion. To avoid self-interference more channels would be required which would increase the licensing costs and potentially make the network un-deployable. One of the main advantages of the mesh/ring architecture is that it significantly reduces congestion, thus decreasing the number of unique channels required. This is achieved by distributing spectrum throughout a market, rather than concentrating it at a hub site.

Furthermore, spectrum costs can be reduced by using equipment designed to operate in specific area-licensed spectrum options. This will eliminate the need to lease additional spectrum. Unlike other vendors whose equipment does not support specific area-licensed spectrum options, DragonWave's backhaul equipment is designed to operate in 24 and 28 GHz, 24 UL, and >40GHz.

High modulation in small channels, coupled with mesh/ring architecture and targeted area-licensed equipment has allowed DragonWave to solve the spectrum scarcity problem very elegantly, eliminating the need for channel overlays and reducing monthly spectrum lease costs by a large margin.

2.3 Impact of Mesh/Ring Architecture

Network architecture does not only affect the antenna and spectrum lease costs. The choice of network architecture affects the overall TCO of the wireless backhaul and it does so in the most profound way.

Conventionally, backhaul networks had been built predominately on a hub-and-spoke architecture. While this architecture may be well suited for access networks and voice-only networks, for high capacity backhaul networks it proves to be very limiting and cost prohibitive. With network expansion, the hub-and-spoke architecture will lead to scalability problems, because this architecture requires a tremendous aggregation of frequencies and antennas at a single hub site. Additionally, due to availability constraints the hub-and-spoke architecture can only span 2-4 links from the fiber site. This increases fiber requirements and also limits the number of sites that can be served by the wireless architecture. In order to correct the situation, the backhaul portion of the network needs to be built using a more flexible and scalable mesh/ring architecture.

Mesh/ring networking routes communications between nodes following a mesh/ring topology whereby continuous connections and reconfiguration around broken or blocked paths are accomplished by “hopping” from node to node until the required destination is reached. As opposed to hub-and-spoke (tree) networks, in mesh networks connectivity is distributed across multiple sites, avoiding the site congestion of a hub-and spoke architecture. Furthermore, because the mesh/ring network provides path redundancy, network availability targets can be achieved while going 8-10 hops away from the fiber PoP. This leads to reduced fiber use and increased site coverage.

Finally, in order to achieve an acceptable availability in a hub-and-spoke architecture, redundant systems are required to harden the links, leading to a significant increase in CAPEX. The mesh/ring network, on the other hand, has an inherent equipment protection which is enabled via the possibility of alternate routes. This inherent redundancy of the mesh configuration eliminates the requirement for the doubling of hardware of the hub-and-spoke architecture, leading to significant savings in CAPEX.

2.3.1 Connectivity of Fiber PoPs

In a hub-and-spoke architecture, in order to establish a connection to the fiber line, each node needs to be connected to a fiber Point of Presence (PoP). In a mesh/ring network, as long as there is least one node present as a PoP, the connectivity to the fiber line is enabled for all nodes in the network. This peer-to-peer connectivity opens opportunities for multiple choices for a fiber location, which takes the fiber connectivity burden out of the network design.

2.3.2 Line of Site Flexibility

In a mesh/ring architecture, the site needs to be connected to another ring or mesh site. Because there are many more ring nodes than there are fiber nodes, there are significantly more line of site options to a fiber node than it is possible in a hub-and-spoke architecture. This is especially important as the network becomes field verified and line of sight candidates are eliminated. Given that there are many line of sight options in the mesh/ring architecture, a new path can quickly be chosen without drastically changing the design.

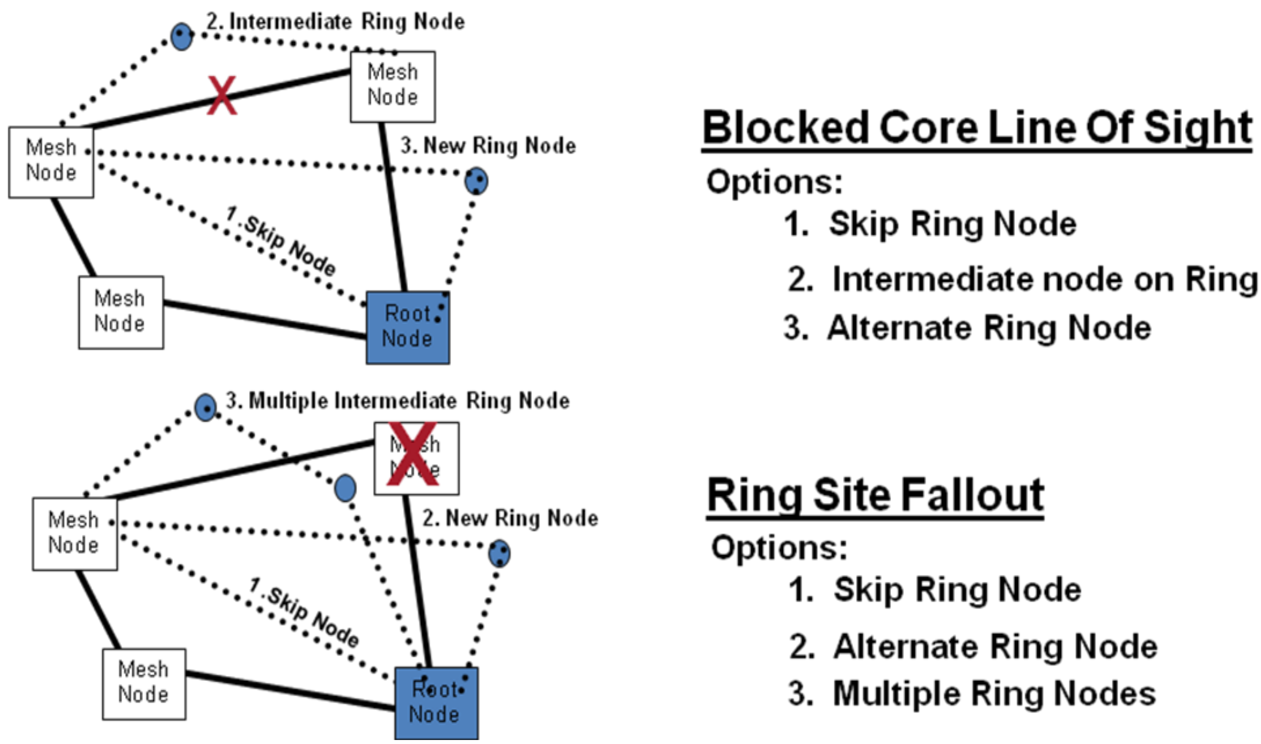

2.3.3 Ease of Deployment, Scalability and Expansion

Unlike hub-and-spoke architecture, mesh/ring architecture makes it easy to scale capacity of an existing network as well as add more coverage. With a ring network, as a capacity increases are required, the ring capacity can be scaled up to its maximum capacity. After that, the ring can split in the middle to double the total capacity. If further scale is required, the rings can be overlaid, using the second polarization and same antennas to provide a further doubling in capacity. Using Horizon Duo's 1.6 Gbps capacity, this methodology can deliver 100 Mbps per site to a 48 site cluster, or 200 Mbps per site to a 24 site cluster.

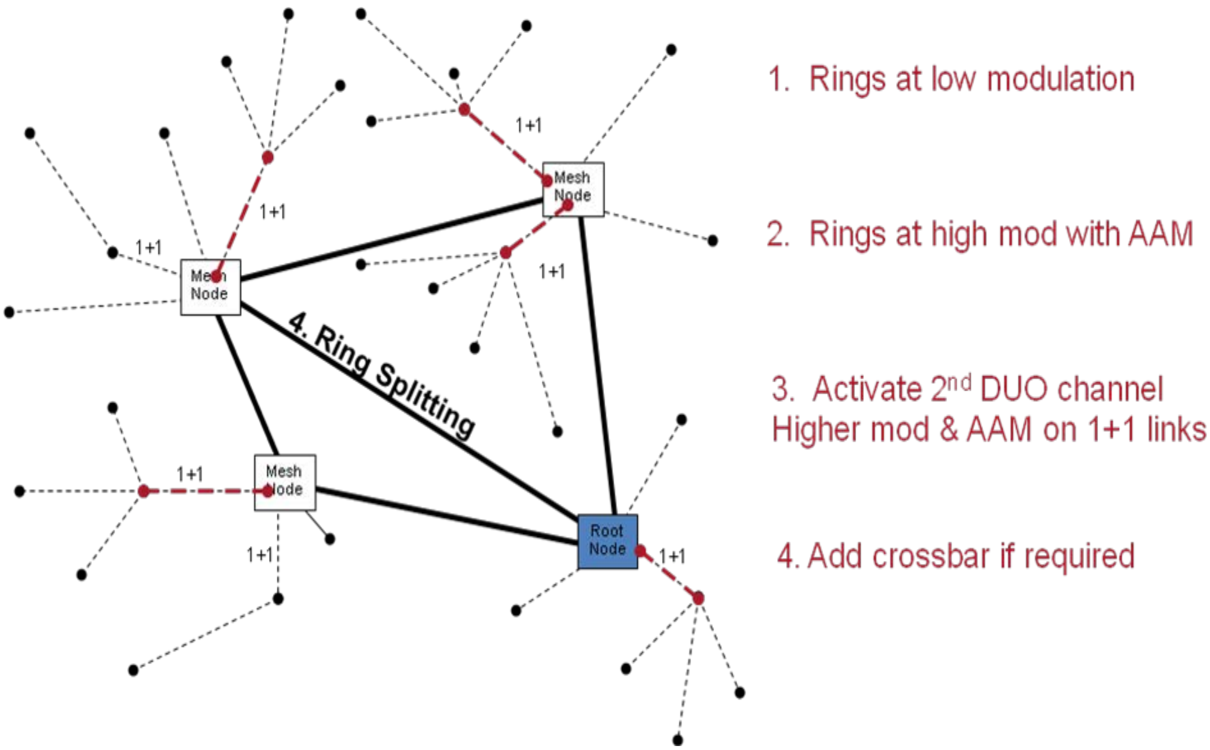

2.3.4 Redundancy at Lower Costs

Unlike mesh, hub-and-spoke also requires additional resources to enable redundancy. The mesh/ring N+1 architecture provides redundancy while requiring only one more link than there are nodes, which helps minimize network costs. This is significantly lower than in a 1+1 hub-and-spoke architecture, where at least two or more links are needed per node, depending on the diversity required.

2.3.5 Shortened Radio Links

Apart from many other important benefits, the mesh/ring architecture allows for shortened radio links, which means reduced antenna size and repeater requirements. This shortening of links, is caused, because each link goes back to a ring site, and then multiple other ring sites before going to the fiber pop, as opposed to a hub-and-spoke architecture, where long links are used to connect back direct to the fiber PoP.

2.3.6 Simplified Zoning

Furthermore, mesh/ring networks allow dealing with various municipal zoning requirements in a more effective way. The distribution of resources in a rign network avoids site overloading. In addition, if a site becomes difficult to obtain from a zoning or leasing perspective, the ring design can quickly be modofied to include an alternate site, or to skip that site altogether, without having to redesign the rest of the ring.

2.3.7 The DragonWave Mesh Networks

Unlike most other vendors of wireless backhaul equipment whose equipment is bound by a hub-and-spoke architecture, DragonWave's technology had been designed to enable ring/mesh networks. The mesh/ring is deployed by using DragonWave links, in tandem with an approved Ethernet switch. The Ethernet switch provides protection switching in less than 50 ms when coupled with DragonWave's patented Rapid Link Shutdown feature.

The mesh/ring can be deployed with any of the following licensing options:

- Area License (i.e. 24 GHz DEMS, 28 GHz LMDS, 26 GHz LMDS);
- Link-by-Link licensing (i.e. 6, 7, 8, 11, 13, 15, 18, 23, 26 GHz); and
- 24 GHz Unlicensed

Implementing mesh/ring architecture is the first most important step towards achieving quick profitability on wireless backhaul.

2.4 Reducing Installation, Maintenance and Management Costs

As discussed in the previous paper ,the installation costs of a microwave link is a one-time expense, including line of site survey and site preparation. It is assumed to be $7,500 per link in North American and an equivalent expense in Western Europe. The installation costs can be reduced by consolidating backhaul installation with base station installation. This consolidation allows reusing the installation climbing resources, which helps to reduce the backhaul installation cost by a large margin.

Maintenance and management costs are reduced by DragonWave by including an integrated RF loopback, simplifying troubleshooting and halving the number of technicians required to go to the field during a fault scenario. In addition, commissioning timing can be reduced by DragonWave's network Ethernet loopback.

Given that the installation, maintenance and management costs comprise a considerable part of the TCO, a wisely implemented consolidation approach can help reduce the overall costs by a noticeable margin.

3. Reducing CAPEX

CAPEX costs are usually a fixed upfront investment required to purchase equipment. Depreciated and amortized over time, in the early stages of deployment the CAPEX costs put a serious financial pressure on carriers and operators. The good news is that the CAPEX costs can be reduced.

One way of reducing CAPEX is to offer **remote software-enabled bandwidth scalability with a flexible pricing scheme**. DragonWave's IP-based radios offer flexible bandwidth pricing to mitigate upfront CAPEX costs. DragonWave's Horizon Duo provides a configurable throughput speed from 10 Mbps up to 800 Mbps in 10 Mbps increments. Depending on the customer requirements, the Horizon Duo link can be engineered and licensed at the maximum throughput speeds, but deployed at current demanded throughput levels. This can allow operators to license for a target bandwidth and set the system to that target modulation, but start at a lower initial capacity. DragonWave believes that this capability to"flex" the bandwidth of the network and map capital outlays to growth requirements to revenues is crucial for operators. The flexible bandwidth scaling enables operators to have a backhaul network that is affordable day one while being able to gracefully expand and meet incremental bandwidth requirements when and if they occur. With Horizon Duo, operators can double the bandwidth to 1.6 Gbps simply by adding a second radio via the Dual Polarization Radio Mount onto the existing antenna.

Furthermore, the upfront backhaul CAPEX costs on equipment can be reduced by purchasing microwave equipment with **a highly integrated form factor in the physical design**. DragonWave's Horizon Compact is one such system with a unique all-outdoor compact design. Featuring zero-footprint, the radio and the modem are integrated into one compact all-outdoor unit. Designed as an Ethernet platform from the ground up, the DragonWave Horizon Compact meets the critical needs demanded by carrier class customers delivering a wireless GigE/100bT connection of up to 800 Mbps full duplex over licensed or unlicensed frequency allocations. With a native Ethernet design and ultra-low latency, the Horizon Compact is optimized for next generation services. When compared to equivalent leased TDM line operations, Horizon Compact offers significant cost benefits that are paralleled with unsurpassed industry-leading performance along with remarkably simple management and operations.

4. Conclusion

Combined together the design strategies analyzed in this paper reduce the burden of the backhaul network on the overall business case. By selecting the appropriate backhaul technology for each situation and by paying attention to the overall lifecycle cost, rather than simply the capital cost of the equipment, and by implementing a combination of proactive design strategies, it is possible to deliver services at significantly lower cost per bit than today's networks, enabling the evolution from a voice centric mobile network to one that delivers truly mobile personal broadband – profitably.

With the understanding that the backhaul portion of the network eats away the profitability of the entire network, we calculated in the previous paper that 12% of the TCO is taken up by CAPEX while the staggering 88% is consumed by OPEX, wherein 75% of OPEX (and 66% of the entire TCO) is consumed by monthly antenna and lease costs.

This paper illustrated that DragonWave's technology has successfully reduced the backhaul burden by astonishing 60-80%, thus removing obstacles to its quick profitability and enabling operators to expand their market coverage and meet increased bandwidth requirements rapidly and affordably.

The DragonWave Horizon products provide software controlled scalability, ultra-low latency, wire-speed connectivity up to 1.6 Gbps full duplex, 99.999% service availability and a full suite of network management options. Ring/mesh network architecture coupled with adaptive modulation, and integrated antenna solutions enable a >50% reduction in the total life cycle cost compared to conventional designs. Offering an option of a zero footprint solution that integrates the IDU and ODU functionality of a conventional radio into a single high performance all outdoor unit, DragonWave helps reduce space/site leasing costs. By integrating pseudowire functionality as a convenient plug-in option on an IDU, DragonWave enables operators to leverage and maximize investments in existing TDM infrastructure.

Building upon its reputation as the world leader in high capacity broadband wireless networking systems, DragonWave has solved the wireless backhaul cost challenge, enabling carriers and network operators to turn backhaul revenues derived from new services and expanded capacities into profits.