

Point-to-Point Broadband Wireless for the Enterprise:

Licensed Performance in the Unlicensed Band

Introduction

As enterprises extend their networks beyond the confines of walls and buildings, point-to-point, wireless Ethernet systems have become more widely adopted to bridge IP traffic – especially across spaces that are too difficult or too expensive to bridge with wires. Whether the requirement is to link separate loops within individual buildings, communicate between buildings or link networks in a campus setting, enterprises increasingly turn to wireless Ethernet as the preferred solution.

This paper provides:

- Multiple approaches and technologies for point-to-point wireless
- Challenges and advantages of those approaches for enterprises that seek reliable, high-throughput, economical wireless solutions for their high-performance networks

Wireless Ethernet Versus Microwave

When enterprises need to bridge IP traffic, there are a number of internal and external factors that influence decision makers to choose Ethernet over traditional microwave or analog technologies. Those factors can include the inability to lay cable or fiber, infrastructure complexities, cost, high bandwidth requirements, path obstructions, applications and throughput requirements. Increasingly IT professionals choose wireless Ethernet over conventional microwave technology because they prefer a native IP infrastructure end-to-end, rather than converting to circuit-switched technology just to cross an air gap. Plus, they favor the speed of digital, multi-megabit performance over T1 data rates on analog networks.

While the reasons to adopt wireless Ethernet are fairly clear-cut, the decision about which type of wireless Ethernet solution to adopt is less obvious. Organizations essentially have two options: licensed and unlicensed. Licensed Ethernet operates within the part of the radio spectrum (e.g., 6.0 GHz in the U.S.) designated by government regulators to be reserved for individual license holders. Licensed operators are guaranteed exclusive use of that part of the band over an assigned geographic area. That means a low likelihood of interference but more complexities from a set-up and management perspective.

The other option is to use the part of the spectrum regulators have set aside for unlicensed operation (currently 5.8 GHz in the U.S.). You may not have exclusive use of the band but neither do you have the delay and expense of obtaining a license. Unlicensed links also tend to be much less expensive – by a ratio of five to one in many cases. In addition, unlicensed links are not restricted to a specific coverage area. Typically, they can be deployed at the owner's discretion, offering greater flexibility for today's mobile and virtual enterprises.

Unlicensed Ethernet systems are typically less expensive, easier to deploy and offer a lower cost of ownership than licensed radios.

Written by Randy Cronk
greatwriting.com

Interference

As unlicensed operation has become increasingly popular, unused unlicensed availability has become increasingly scarce in many areas. As a result, there is a greater likelihood that operator systems will interfere with one another, leading to lower system reliability and level of service. And, should there be a continual need to adjust transmission frequencies, links will become more difficult to deploy and manage.

The ideal solution would be to obtain an Ethernet system that operates in the unlicensed spectrum but effectively delivers the throughput, reliability, availability and management simplicity that enterprise customers require. It would be even better if that system could deliver those benefits without the restrictions, delays and higher costs associated with licensed operations.

Fortunately for business, technology has a way of transforming itself into ideal solutions at the point when enterprises begin to need them. Recently a number of products have emerged that take wireless technology to the next level, performing at high-speeds in Non-Line-of-Sight applications. One such advanced wireless product was developed by Orthogon Systems, whose initial, point-to-point wireless Ethernet bridge, OS-Gemini, received a “Best Products” designation from NetworkWorldFusion in early 2004.

Orthogon Systems designed its latest product, the OS-Spectra point-to-point, wireless Ethernet bridge with this ideal in mind. First, it's a very fast bridge – with up to 300 Mbps aggregate throughput. However, by itself, speed is not necessarily a solution for spectrum crowding. As an analogy, a runner's ability to run fast may not help much if other racers block the lane. With OS-Spectra, many of the same features that boost performance also reduce crowding or make any crowding that does remain easier to manage. Of course, any feature that reduces crowding will also boost performance since there will be less interference to decrease the data rate.

It's this recognition of the relationship between speed and spectrum that makes OS-Spectra unique in its ability to achieve licensed performance in the unlicensed space. The following sections describe the various technologies employed in wireless Ethernet bridging and explain how those technologies impact reliability and performance.

An Orthogonal Approach

To understand what an “orthogonal approach” does, it helps to see it as applied to an actual Ethernet bridge such as Orthogon Systems' OS-Spectra. The name of the company provides some insight into the key architectural features underlying the bridge. Specifically, each Orthogon-enabled link employs four signals between two bridges on opposite ends of the air gap, each of which is connected to an antenna. One antenna on each end is oriented to polarize signals in the vertical plane while the other antenna polarizes signals in the horizontal plane.

On the receiving end, each antenna receives two orthogonal signals, i.e., two signals oscillating at right angles to each other. This basic architecture creates a number of mutually reinforcing advantages with respect to performance, reliability and spectrum management:

Multi-Beam Coding

Simply having four paths from which to sample data greatly increases the probability that at least one of these paths will succeed. This is like saying that instead of flipping one coin to get heads once (1:2 probability), we will flip four coins to get heads once (11:15 probability). Conversely, the probability of failure is reduced by the same degree, so this technique allows a higher data rate than if only a single path was used. It assumes that the sampling algorithm only makes a simple decision to sample whichever of the four signals happens to be the strongest. An even lower failure rate – and proportionately higher data rate – is achieved if data from all paths are sampled (even the weakest) and the signal reconstituted. This is what happens in the OS-Spectra.

However, multi-beam transmission does more than just provide more data to sample. When combined with orthogonal signal propagation, it is also the basis for more advanced performance-boosting techniques like Space Time Coding. The real point is that it does so in a way that simultaneously makes the link more spectrum efficient.

Space Time Coding (STC)

Transmitting mutually orthogonal signals can only happen if you allow for higher data rates (i.e., by sending multi-beam signals). Similarly, the benefits of STC are maximized only if you oscillate these individual beams on orthogonal planes. The reason is that orthogonal signals are less likely to fade – for example, by canceling each other out at various points in the waveform. Fading can occur when correlated signals (i.e., same timing, modulation and frequency) become de-correlated due to atmospheric disturbances or when bouncing off obstructions. Under these conditions a signal may even degrade itself if slightly different variations of its waveform reach the receiver at multiple, but slightly varying, times.

Because orthogonal beams don't "see" each other, they don't interact and are much less likely to fade. This is key to the design of non-line-of-sight wireless bridges that have to constantly contend with highly de-correlated signals arriving at the receiver from many different reflection points. Multi-beam Space Time Coding is a technique that can take advantage of this "fade immunity" to create signals that are even more tolerant of fading as well as other types of signal degradation such as distance and interference.

In OS-Spectra, Multi-beam Space Time Coding does this by deliberately de-correlating the four beams of the orthogonal multi-beam coding with respect to time, modulation or frequency. The fact that the beams behave differently reduces the chances that any condition that degrades one of them will have the same effect on all of them. Whatever data does reach the other end of the link is then available to the Multi-beam Space Time Coding algorithms for signal reconstitution. Not surprisingly, the weight of a beam's individual contribution to the final output depends on its clarity, among other variables.

A Legal View

The terms "license-exempt" (also called "license-free" or "unlicensed") and "licensed" refer to the radio frequency spectrum rules defined by the U.S. Federal Communications Commission (FCC) or equivalent national government regulatory body. In the United States, FCC Rules Part 15 governs the license-exempt frequency spectrum, and Rules Part 101 governs the licensed frequency spectrum. Licensed products require regulatory approval before deployment while license-exempt products can be deployed without any regulatory approval.

A "license-exempt" system can be installed virtually anywhere within a given country without obtaining a license to operate from the regulatory authorities. Such a system must already be certified to operate as license-exempt in that country. Manufacturers desiring license-exempt certification must apply to the FCC or equivalent national authority for approval to operate the particular product in specific radio frequency bands. FCC rules encourage efficient use of RF bandwidth and harmonious co-existence of different systems using the same radio spectrum. The FCC (or equivalent authority outside of the US) process helps to isolate the opportunity for interference by requiring all license-exempt devices to conform to the same standards. Once a product adheres to specific national regulations and the manufacturer obtains certification by the appropriate governing authority, anyone can deploy that manufacturer's equipment anywhere in the country without further regulatory approvals.

The fact that Multi-beam Space Time Coding de-correlates beams in the frequency domain implies the need for spectrum separation between beams, and, consequently, the need for more spectrum in order to carry a wider signal. This opposes the objective of squeezing more data into less spectrum. Thanks to orthogonal propagation, it is possible to both shrink band utilization and achieve spectrum separation at the same time. In fact, OS-Spectra's "spectral footprint" is almost 1/3 less than that of conventional wireless bridges. Because the orthogonal beams don't see each other, the frequencies of their de-correlated waveforms can overlap to a far greater extent than is possible in most STC implementations. That allows many more channels in the same band and less crowding. In many applications, almost no actual antenna separation is required on the tower to maintain signal separation, permitting integration of antennas within a single enclosure for easy and unobtrusive mounting.

This assumes that Multi-beam Space Time Coding multi-beam coding is needed to overcome some adverse condition. But what if it's not? In a clear-line-of-sight application (i.e., the type of application where you might have considered a licensed link), you might not need to send the same data on four beams simultaneously. If you were to transmit different data on each antenna, the data rate would double within an extremely narrow band (30 MHz for OS-Spectra). With the OS-Spectra, the decision of whether to employ multi-beam coding is a configurable option.

Intelligent OFDM

There are two basic ways wireless radios encode data – single carrier or multi-carrier. Single carrier encodes symbols as states of amplitude and/or phase per unit of time on a single carrier. Multi-carrier does the same thing, except it decomposes symbols and spreads the “pieces” in parallel across the multiple carriers and reconstitutes them on the other end. Because no one carrier is responsible for the entire symbol, parity checking can recover information on lost or degraded channels. This recovery feature makes multi-carrier encoding more fault tolerant than single carrier.

There are also two basic methods of multi-carrier encoding currently in wide use: CDMA (code division multiple access) and OFDM (orthogonal frequency division modulation). As its name implies, CDMA uses a mathematical code (called a Walsh code) to separate and allocate data across the band. OFDM also separates data into channels, but these channels overlap in frequency – that is, the channels are orthogonal to each other, so they do not interfere with each other. OFDM-based wireless bridges (like OS-Spectra) are inherently more spectrum efficient while benefiting from the fault tolerance of multi-carrier encoding, resulting in higher data throughput.

In addition, OFDM has qualities that better exploit the attributes of the Multi-beam Space Time Coding environment. In particular, it allows for the bridge to apply a uniform phase correction (say 1/30th KHz), if needed, to all channels simultaneously to compensate for an environmental condition. That’s a result of how OFDM spreads data across the band. Instead of using a predetermined encryption code, it uses a derived parametric value (a sine wave) which can be modified on the fly in response to external events, such as cross-interference between multiple beams.

This correction technique, called channel equalization, can be applied intelligently and dynamically provided that adequate information exists with which to do so. Orthogon provides this information by transmitting pilot tones of known characteristics in advance of the data and dynamically adjusts the coding value until the signal is optimum. The result is virtually instant recovery from even the deepest of fading situations.

More Spectrum Efficiency

Although not fundamental to the architecture like multi beam coding, orthogonal propagation, Multi-beam Space Time Coding and OFDM, there are some additional technologies that enable users to achieve even higher data rates with greater spectral efficiency. They include:

256 QAM Modulation

As described earlier, wireless digital devices transmit information as phase modulations. To send more information, you must increase either the amount of digital information (i.e., number of bits sent) or the amount of analog information (i.e., channel frequencies used). It’s a direct tradeoff. 256 QAM (the default in OS-Spectra) is a modulation scheme that employs eight bits per symbol versus the six bits employed in 64 QAM (the industry norm). That reduces the number of channels needed (and resulting spectrum footprint) by 25%. Because it is an inherently higher-performing bridge, it can trade channel size for bits while still providing high data throughput.

Adaptive Modulation

Within a user-specified range the operator can set the bridge to operate within a specified frequency band and (if possible) at a specified bit rate. Alternatively, the operator can set the bridge to operate at a specified bit rate and (if possible) within a specified frequency band. The bridge will then dynamically shift modulation modes as conditions warrant to satisfy the request. Available modes are 256QAM, 64QAM, 16QAM, QPSK, BPSK, BPSK Multiple FEC rates, single and dual payload.

ISM Band Operation

Operators of radio transmitters must abide by one of two sets of rules governing radio wave emissions: ISM (Instrumentation & Monitoring) or UNII (Universal Network Information Infrastructure). The ISM rules allow for a channel that is 25 MHz wider than UNII rules, giving the equipment that employ these rules (like OS-Spectra) a larger spectrum in which to find interference-free channels.

Dynamic Frequency Selection

The wireless bridge monitors all available radio channels and dynamically selects those over which it can sustain the highest data rate. This means the bridge is very likely to find a clear channel (without operator intervention) even in a crowded space. From the user's point of view, this experience is equivalent to having exclusive rights to use a licensed channel.

Security

Most operators of licensed links care about security. They wouldn't make the considerable investment in infrastructure if their data wasn't important. Most network managers are especially concerned about over-the-air transmissions where physical interception is an obvious threat.

If the phrase "licensed performance in the unlicensed space" is to have meaning, then security must be addressed. Whether a device is licensed or unlicensed does not, by itself, reveal much about the level of security provided. Unlicensed wireless bridges that provide robust encoding – using techniques like Multi-beam Space Time Coding and intelligent OFDM – have a headstart on security versus licensed wireless bridges.

To boost performance and spectral efficiency, Orthogon applies a proprietary Space Time Coding algorithm. To decode the data, a hacker would not only have to know Space Time Coding, the hacker would also have to know Orthogon's unique software. Part of what makes that software exclusive is that one bridge will communicate only with its matched counterpart at the other end of the link, and with none other. That communication is also encoded (within the OFDM encoding) using a proprietary encryption algorithm. A fourth layer of encryption is provided (layered on top of STC, iOFDM and proprietary encryption) with the optional addition of 128-bit AES encryption. The result is a highly secure transmission.

Management Ties It All Together

Just as they expect industrial-strength security, licensed users also expect industrial-strength management. The two go hand-in-hand given that end-to-end visibility of the entire network, both wireless and wired, is a requirement of both security and management. This requires that the wireless link publish a MIB (management information block) in an industry-standard format like SNMP (Simple Network Management Protocol) accessible to third-party network management systems. Optimally, it would also provide an interface compatible with popular web browsers to allow anywhere-anytime, hands-free configuration, status and performance monitoring as well as configuration setup.

Beyond the fundamentals, however, unlicensed wireless bridges can also offer management capabilities expressly designed to help maximize throughput and minimize interference from other unlicensed users. For example, in addition to self-selecting the least congested channel, the OS-Spectra also provides a real time and historical view (up to 25 hours) of activity across the wireless band to help operators pinpoint which channels offer the clearest data paths. The fact that the radio is more spectral efficient means that the channels required to broadcast can be smaller, which makes the job of finding clear paths that much easier. In other words, better spectral efficiency fosters better spectrum management.

Interference is mitigated by the ability to synchronize the transmissions of multiple bridges on the same mast so that some bridges are not receiving data while others are sending. This is a feature that can be enabled with the connection of an optional GPS receiver as a source of universal time.

Interference is also minimized by the ability to determine in which direction (upload or download) users receive the greatest allocation of available bandwidth. Most IP users download much more data than they upload, so allocating more bandwidth for downloads not only improves the user's experience it also improves return on assets. Of course, some operators will want to allocate the same amount of available bandwidth in both directions equally and have that flexibility as well.

These are the types of features that help enterprises manage their businesses, while enjoying the benefits of technological advances.

What Licensed Performance Really Means

Achieving 100% interference-free transmission is not possible even with licensed systems. The primary difference between licensed and unlicensed systems is that licensed radio users have a regulatory body that will assist them in overcoming any interference issues that may arise, while unlicensed users must resolve interference issues without governmental assistance. In both cases, proper selection of the frequencies and methodical engineering of the path are key to reducing the potential for interference.

Ultimately, what determines licensed performance is not the license, but whether the wireless bridge meets a certain expectation for performance, reliability, security and manageability. When an unlicensed radio achieves a high standard in any of these criteria, it enjoys a head start toward achieving a high standard on the others, as we saw with performance and security as well as security and management. That's especially true if the wireless bridge manufacturer exploits the potential synergies inherent between high performance and high spectral efficiency.

Licensed radios are not built to be spectral-efficient because they don't have to be. Darwinian logic would say that – given this lower threshold – they might not be able to do other things well either. That could mean that sometimes the best path to true licensed performance is on an unlicensed link.

Get Ready For WiMAX

As WiMAX standards get sorted out and momentum for the standard builds, the terms "WiMAX Compatible" or "WiMAX Ready" will be sung louder and more often from the marketing materials of wireless bridge manufacturers. Yet, what does "WiMAX Ready" mean – especially in the context of a point-to-point solution as described in this paper?

WiMAX is a multipoint-to-multipoint solution like WiFi except over a larger geographic area. The concept is for multiple users who walk into a WiMAX hotspot with their laptops or other wireless devices and go online. However, a point-to-point link is set up for the purpose of allowing just two devices, on opposite ends of a link, to communicate, so they don't – and probably shouldn't – "speak" WiMAX.

Does that mean there's no place for "WiMAX compatibility" in a point-to-point wireless solution? Of course there is, especially where a point-to-point link will co-exist with a WiMAX network. In that context, a wireless point-to-point link is WiMAX-compatible to the extent that it meets three tests:

1. **Integrated management.** The manager of a WiMAX network should not have to worry about "islands of connectivity" when managing an entire infrastructure. The point-to-point equipment should be completely visible and manageable on the operator's console just as if it were one of the WiMAX devices. That means publishing a WiMAX-compliant MIB.
2. **High throughput.** Today's high-performance networks generate great quantities of diverse traffic, and demand significant data rates, even in adverse environmental conditions. A point-to-point solution should be able to extend connection distances while decreasing the number of hops and increasing data rate.

**Orthogon Systems Solutions
Succeed More Often**

3. **High spectral efficiency.** By establishing multiple transmission paths, a point-to-point solution increases the probability that data will successfully transmit. In addition, by sampling each transmission on each path and automatically adjusting to the optimum signal, the probability increases even further. The combination of the two provides highly efficient spectrum management.

Orthogon Systems manufactures a series of point-to-point 5.8 GHz wireless Ethernet bridges under the brand names OS-Spectra and OS-Gemini. The success of these products results from Orthogon's unique combination of technologies, the overall effect of which is much more powerful than using any of them individually. You can acquire this combination in a surprisingly small form factor that is easy to install and maintain, even in aesthetically or physically restrictive environments. These key technologies include:

- Dynamic Spectrum Management
- Multi-Beam Space/Time Encoding
- Best-in-Class Radios
- OFDM (Orthogonal Frequency Division Modulation)
- Dynamic Frequency Selection
- Adaptive Modulation

By any measure, Orthogon's solutions succeed where more conventional products disappoint ... and achieve far more reliable connections, much more often, at higher data rates.

Configuring OS-Spectra

OS-Spectra creates a point-to-point, wireless Ethernet link across an air gap of any distance up to several miles. The two endpoints communicate via radio waves on one of 18 different channels. Having so many channels is only possible because each channel uses only 30 MHz of bandwidth. This greatly increases the probability that users can find at least one usable channel in an air space already crowded with other radio communications.

Each end of an OS-Spectra link consists of an integrated outdoor unit (ODU) and a small, powered indoor unit (PIDU), with the required mounting equipment. The PIDU is a lightweight, wall-mountable box about the size of a pocket dictionary and takes up no rack space. It connects to the ODU via a single RJ-45 (CAT5) cable and supplies the ODU with both power and the Ethernet data to be communicated over the link.

The ODU is a small, lightweight transceiver that contains all the required radio and networking elements. It comes with a mounting bracket and can be positioned easily and quickly on a pole or mast. The unit's small size and light weight make the ODU ideal for setup in space-constrained and aesthetically challenging environments. OS-Spectra includes an embedded web server to manage the link either directly or remotely, and setup is easy. Simply install the units on their respective mountings and align the antennas, using an audible signal.

Put OS-Spectra To Work For You

In enterprises, OS-Spectra supports high-bandwidth applications in environments where wired networks are either too expensive or impossible to implement. Offering high speed and reliability over long distances, OS-Spectra efficiently uses the frequency spectra to reduce interference and boost performance for business-critical applications.

Whether the requirement is to migrate from an analog to a digital network, link separate loops within individual buildings, communicate between buildings or link networks in a campus setting, OS-Spectra delivers faster data streams for networked applications in a variety of vertical markets, including:

- Utilities
- Transportation
- Healthcare
- Government
- Education

HEADQUARTERS

Orthogon Systems
Unit A1, Linhay Business Park
Eastern Road, Ashburton
Devon, TQ13 7UP, UK
Tel +44 1364 655500

USA OFFICE

Orthogon Systems LLC
890 Winter Street, Suite 320
Waltham, MA 02451

Sales and Technical Support in
North America:
+1 877 515-0400

Outside of North America:
+44 1364 655569

www.orthogonsystems.com

©Copyright 2005 Orthogon Systems LLC. All rights reserved. All trademarks are the property of their respective owners. All statements of fact contained herein are provided for informational purposes only and are subject to change without notice. No warranty of accuracy is expressed or implied and the user of this information assumes all liability.