



Point-to-Point Broadband Wireless for Enterprises

Licensed performance in the unlicensed space



CONTENTS

Pg **Section**

3	Introduction
3	Analog versus IP
3	Licensed and Unlicensed Wireless
4	Speeding without a License
4	Orthogonal Architecture
5	Multiple-Input Multiple-Output (MIMO)
6	<i>Intelligent</i> OFDM
6	More Spectrum Efficiency
7	Security
7	System Management
8	Reliability and Performance
8	WiMAX
9	Configuring a Link
9	Summary

Introduction

As enterprises extend their networks beyond the confines of walls and buildings, point-to-point, wireless Ethernet systems have become more widely utilized 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

Analog Versus IP

Increasingly, enterprises choose to bridge IP traffic with microwave rather than to lay new cables or incur monthly leased-line charges. They also 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. This decision to remain digital end-to-end offers several clear-cut advantages:

- Reduced infrastructure complexities
- Lower equipment cost
- Ability to overcome path obstructions
- Higher throughput

Wireless Ethernet in particular offers enterprises reliable, secure multi-megabit speeds with significantly better performance than the T1/E1 data rates of analog networks.

Licensed and Unlicensed Wireless

The decision to adopt a wireless Ethernet bridge brings with it another decision: whether to operate in the licensed or unlicensed space. Both have their advantages. However, the ideal solution would be to enjoy the advantages of both, or what might be called, *“licensed performance in the unlicensed space.”*

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 permitted exclusive use of part of the band over an assigned geographic area. With exclusive rights, an operator should be able to operate without interference or spectrum crowding caused by other operators transmitting in the same frequency band, in the same geographic area. Less interference should translate into higher throughput – i.e., better link performance.

Although exclusive rights are granted, that doesn't mean that they are enforced, or that they are enforced quickly. The burden of protecting one's rights often falls to the rights holder – and legal remedies can take time. In addition, obtaining a license is expensive and can require significant time. A license also restricts the operator to the geographic area covered by the license.

The unlicensed part of the spectrum (e.g., 5.4 and 5.8 GHz in the U.S.) does not promise exclusive use of the band. However, what it does do is eliminate the delay and expense of obtaining a license. Unlicensed equipment (radios and antennas) also tends to be much less expensive to buy and install. In addition, unlicensed links are not restricted to a licensed coverage area. Typically, they can be deployed at the owner's discretion, offering greater flexibility for today's providers and telecoms to serve mobile and virtual enterprises.

Therefore, the statement, *“licensed performance in the unlicensed space,”* suggests that an unlicensed wireless solution can achieve the performance implied by exclusive rights without the delays, higher costs and geographic restrictions of a license.

Speeding without a License

Motorola's Point-to-Point Wireless Ethernet Bridges – PTP 400 and PTP 600 Series – were designed with performance and reliability in mind. First, they are very fast bridges with models that can deliver 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.

Since 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 levels of service.

The Motorola PTP 400 and PTP 600 Series systems have features that squeeze more signals into a limited amount of bandwidth to 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 spectral efficiency that makes Motorola's point-to-point wireless systems unique in their ability to achieve high performance in the unlicensed space.

Orthogonal Architecture

Each Motorola-enabled point-to-point link employs multiple signals between two bridges on opposite ends of the air gap, each of which is connected to two antennas. 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 architecture, combined with the following technologies, creates a number of mutually reinforcing advantages with respect to performance, reliability and spectrum management:

- Multiple-Input Multiple-Output (MIMO)
- *Intelligent* Orthogonal Frequency Division Multiplexing (*i*-OFDM)

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 U.S.) process helps to isolate the opportunity for interference by requiring all license-exempt devices to adhere to at least the same minimum 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. Of course, a manufacturer can go beyond the minimum standards and achieve even more efficient spectrum usage than governments mandate or than other manufacturers (licensed or unlicensed) offer.

Multiple-Input Multiple-Output (MIMO)

Through the use of MIMO technology, the PTP 400 and PTP 600 radios radiate multiple beams from the antenna. Having multiple beams from which to sample data greatly increases the probability that at least one of these beams will succeed. 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 correspondingly higher data rate – is achieved if data from all beams are sampled (even the weakest) and the signal reconstituted. This is what happens in the Motorola PTP 400 and PTP 600 Series bridges.

The benefits of MIMO are maximized when the individual beams in the multi-beam architecture oscillate on orthogonal planes because 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 critical 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. MIMO is a technology that can take advantage of this "fade immunity" to create signals that are even more tolerant of fading and other types of signal degradation such as distance and interference.

With MIMO technology, the PTP 400 and PTP 600 Series bridges deliberately de-correlate the beams 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 MIMO 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.

De-correlating 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, the "spectral footprint" of Motorola's point-to-point bridges is significantly less than that of comparable 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 MIMO 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.

These techniques assume that MIMO technology is needed to overcome some adverse condition. But what if it's not? In a clear-line-of-sight application, it might not be necessary to send the same data on four beams. In such cases, the PTP 600 Series radios will intelligently switch to "Dual Payload" mode if RF conditions will support it. In this mode, different data can be transmitted in parallel on each transmitter, effectively doubling the bandwidth at the higher modulation rates.

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 multiplexing). 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 Motorola’s point-to-point bridges) 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 MIMO. In particular, it allows 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. Motorola uses an enhanced version of OFDM, called *intelligent* OFDM, and 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 orthogonal propagation, MIMO and *μ*-OFDM, there are some additional technologies that enable users to achieve even higher data rates with greater spectral efficiency. They include:

QAM Modulation

Wireless digital devices transmit information as phase modulations. To send more information, either the amount of digital information (i.e., number of bits sent) or the amount of analog information (i.e., channel frequencies used) must increase. It’s a direct tradeoff. In 64 QAM (the industry norm), the modulation scheme employs six bits per symbol. However, 256 QAM (the default in Motorola’s PTP 600 Series bridges) employs a modulation scheme that uses eight bits per symbol, reducing the number of channels needed and resulting spectrum footprint in these bridges. For PTP 600 Series systems, available modes are 256 QAM, 64 QAM, 16 QAM, QPSK and BPSK, multiple FEC rates, single and dual payload. For PTP 400 Series systems, available modes are 64 QAM, 16 QAM, QPSK, BPSK and multiple FEC rates.

Adaptive Modulation (AMOD)

AMOD enables the transmitter and receiver to negotiate the highest mutually sustainable data rates, then dynamically “upshifts” and “downshifts” the rate as RF conditions change. In addition, within a user-specified range, the operator can set a PTP 400 or PTP 600 Series 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.

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 the PTP 400 and PTP 600 Series systems) a larger spectrum in which to find interference-free channels.

Advanced Spectrum Management with *i*-DFS

Intelligent Dynamic Frequency Selection (*i*-DFS) is at the heart of Motorola's exceptional spectrum management capabilities. At power-up and throughout operation, each point-to-point wireless bridge scans the band – 500 times a second – and automatically switches to the clearest channel. Plus, the 30-day, time-stamped database alerts network operators to any interference that does exist and provides statistics that help pinpoint which channels offer the clearest data paths. From the user's point of view, this experience is equivalent to having exclusive rights to use a licensed channel.

Security

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

In its point-to-point systems, Motorola offers robust, multi-layered security techniques that provide excellent security for data transmissions. To begin, each PTP 400 and PTP 600 Series radio is pre-programmed to communicate only with a matched radio. At installation, each link is programmed with the MAC and IP address of its partner. Then the two ends of the link will communicate only with each other, eliminating "man-in-the-middle" attacks. The pre-pairing also allows fast deployment as all that is needed is power for the modules to start searching for each other.

Over-the-air security is achieved through a proprietary scrambling mechanism that cannot be disabled or spoofed by commercial tools.

On transmission, the signal passes through the following processes:

- Reed Solomon forward error correction where added bits are applied
- Scrambling with a code that repeats every eight Reed Solomon code words
- Interleaver where the signal is then changed in order
- Convolutional Encoding where the signal is scrambled into two streams and then sent serially with some bits unscrambled
- The signal is coded onto one of BPSK, QPSK, 16 QAM, 64 QAM or 256 QAM waveforms
- Then the signal is interleaved across a 1024-carrier OFDM waveform

Plus, an additional layer of security can be applied with FIPS 197 compliant, 128- or 256-bit AES Encryption (optional). Motorola also encourages encryption of data before it is transmitted by using the security measures built into routers, network devices and web sites in order to ensure end-to-end protection of data.

System Management

Just as they seek industrial-strength security, network managers also look for industrial-strength management. The two go hand in hand, given that end-to-end visibility of the entire network (wireless and wired) is a requirement of both security and management. This means that the wireless link must 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.

However, wireless bridges can also offer management capabilities expressly designed to help maximize throughput and minimize interference from other users. For example, in addition to self-selecting the least congested channel, the Motorola PTP 400 and PTP 600 Series bridges also provide a real time and historical view of activity across the wireless band to help operators pinpoint which channels offer the clearest data paths. The fact that the radios are more spectrally 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 network operators will want to allocate the same amount of available bandwidth in both directions equally and have that flexibility as well.

Reliability and Performance

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 crucial to reducing the potential for interference.

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

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 point-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 PTP 600 system is completely visible and manageable on the operator's console just as if it were one of the WiMAX devices through its 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.

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.



Motorola PTP 400 Series Radio



Motorola PTP 600 Series Radio



Motorola PTP 400 PIDU Plus Unit



Motorola PTP 600 PIDU Plus Unit

Configuring a Link

Each PTP 400 and PTP 600 Series solution creates a point-to-point, wireless Ethernet link across an air gap of any distance up to 124 miles (200 km). 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 12 MHz or 30 MHz of bandwidth (PTP 400 and PTP 600 respectively). This greatly increases the probability that users can find at least one usable channel in an air space which may be crowded with other radio communications.

Each end of a PTP 400 or PTP 600 Series link consists of an integrated outdoor unit (ODU), a small, powered indoor unit, called the PIDU Plus, and the required mounting equipment. A lightweight, wall-mountable box, the PIDU Plus is 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. Each PTP 400 or PTP 600 Series system 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.

In Summary

Motorola provides a comprehensive portfolio of best-of-breed wireless broadband solutions that enable enterprises to deploy the most advantageous system to meet their individual connectivity requirements. Available as stand-alone solutions or integrated with Motorola's point-to-multipoint and Mesh products, the PTP 400 and PTP 600 point-to-point wireless Ethernet bridges deliver the reliable, secure and high-bandwidth data, voice and video traffic that enterprises require for applications, such as:

- Building-to-building connectivity
- Replacing or augmenting T1/E1 lines
- Network redundancy
- High-bandwidth connectivity for Voice-over-IP, distance learning and video surveillance

The success of these products results from Motorola's unique combination of technologies, the overall effect of which is much more powerful than using any of them individually. The key advantages of these technologies include:

- Ability to overcome path obstructions
- Reliable, high-throughput communications
- Multi-layered security
- Reduced infrastructure complexities
- Choice of configurations - 10/100/1000 Base T, 1000 Base SX, single or dual T1/E1, DS3, OC-3
- Fast, easy installation and management
- Fast return on investment – typically less than one year



The Motorola Point-to-Point Wireless Ethernet Bridges – PTP 400 and PTP 600 Series – are part of Motorola's MOTOwi4 portfolio of innovative wireless broadband solutions that create, complement and complete IP networks. Delivering IP coverage to virtually all spaces, the MOTOwi4 portfolio includes Fixed Broadband, WiMAX, Mesh and Broadband-over-Powerline solutions for private and public networks.



800-377-2929
info@ersdatasolutions.com



Motorola, Inc.

MOTOROLA, the stylized M Logo and all other trademarks indicated as such herein are trademarks of Motorola, Inc. ® Reg. US Pat & Tm. Office. All other product or service names are the property of their respective owners. © 2007 Motorola, Inc. All rights reserved.