MOTOROLA

Controlling Interference: The Motorola® Canopy™ Approach

August 2002

White Paper

Introduction

Market Demand for Broadband Drives Broadband Wireless Access (BWA)

Living in New York, Chicago, San Francisco or London, business moves at Internet speed, and a broadband connection to the world at large is as necessary as it is taken for granted. Cable modems and multiple DSL providers are ready and eager to provide companies with high-speed connections. Yet the vast majority of the world's population and businesses are not located in large metropolises. This does not mean that the need and demand for broadband is any less; on the contrary, living and working in a tier-two or tier-three city often means these connections are even more important.

Unfortunately for many businesses, domestically and, especially, internationally, reasonably priced broadband is not readily available, sometimes not at any price. The expense of building out new DSL networks, re-working or conditioning the lines that exist, or converting existing cable plants to carry two-way traffic is very expensive. This means a solid business case must exist before broadband can even be contemplated.

Broadband wireless access has been touted as a possible solution to this dilemma, able to bridge the "digital divide" for those suffering from bandwidth starvation. This claim has been heard for the past several years, yet the industry as a whole has not been successful on a large scale. The reasons for this have as much to do with perception as reality. The truth is that until now, equipment lacked the features and price points necessary to support a viable business model.

Even though there are over 30 million broadband subscribers today (Source: In-Stat/MDR, July 2002), the majority of the world is still unable to receive reliable high-speed data and/or voice connections. The most promising access medium, broadband wireless access, accounts for less than five percent of the total broadband access connections (Source: In-Stat/MDR, July 2002). This is not surprising, since the leading broadband access technologies (namely, xDSL and cable) had an enormous head-start on BWA. Nevertheless, the current BWA market is seeing a rebirth in new approaches to solving the issues that had previously stalled its potential growth. An insidious issue for service providers, for example, has been the lack of ability to handle RF interference, which has resulted in higher costs due to additional equipment and an inability to meet service-level agreements (SLAs) contracted with their customers.

Times change and technology moves forward, and today products are emerging that address this critical issue. Once again BWA's star is rising as a potentially powerful answer to the broadband dilemma.

In this paper we will examine one of the key challenges to making BWA ubiquitous as the third leg of broadband access - interference. Customers must be assured that the technology chosen is hassle-free and always available. With BWA the number one threat to this is interference. Specifically, we will cover:

- What interference is and where it is prevalent
- Methods for dealing with interference at the RF level
  - Robust modulation
  - Carrier to interference ratios (C/I) and what they mean
  - Antenna performance
  - Synchronization
- Methods for dealing with interference in the media access layer
  - How the choice of frame size can have an impact
  - Automatic Retransmission reQuest (ARQ)
  - Central transmission control

And, finally, we will take a look at what all these technical issues mean when it comes down to the practical task of deployment.

BWA Exists in Licensed and License-Exempt Bands

Many of the concerns over the demise of BWA were centered on companies deploying high-frequency, 38 GHz point-to-point (PTP) networks. Point-to-multipoint (PMP) systems at these high frequencies were and are available, but the price points and limitations on range and coverage generally limit their effectiveness.

All of these networks were deployed in a licensed band, meaning the company deploying the equipment first had to purchase the spectrum from the government. While there were several reasons for the demise of BWA companies, the extra burden of paying high dollar for licenses made operating in the black difficult if not impossible.

In addition, when licensed bands are designated, typically only two to three licenses per region are granted. On the surface this means that BWA will only be deployed in those places where the license fee can be recouped, and only by a few players. Such a situation effectively reduces the number of potential competitors and, hence, options available to the end customer, freezing out competing BWA options.

Enter BWA in the license-exempt bands. In the US and internationally, spectrum has been set aside in which anyone may deploy and operate a network as long as the equipment adheres to the rules for those bands. These rules are designed to allow multiple networks to co-exist with minimal interference, enabling multiple operators to serve a given geographic region These bands are generally accepted around the globe, providing vendors with the ability to develop a single product for a large market. Figure 1 below shows both sets of frequency bands.

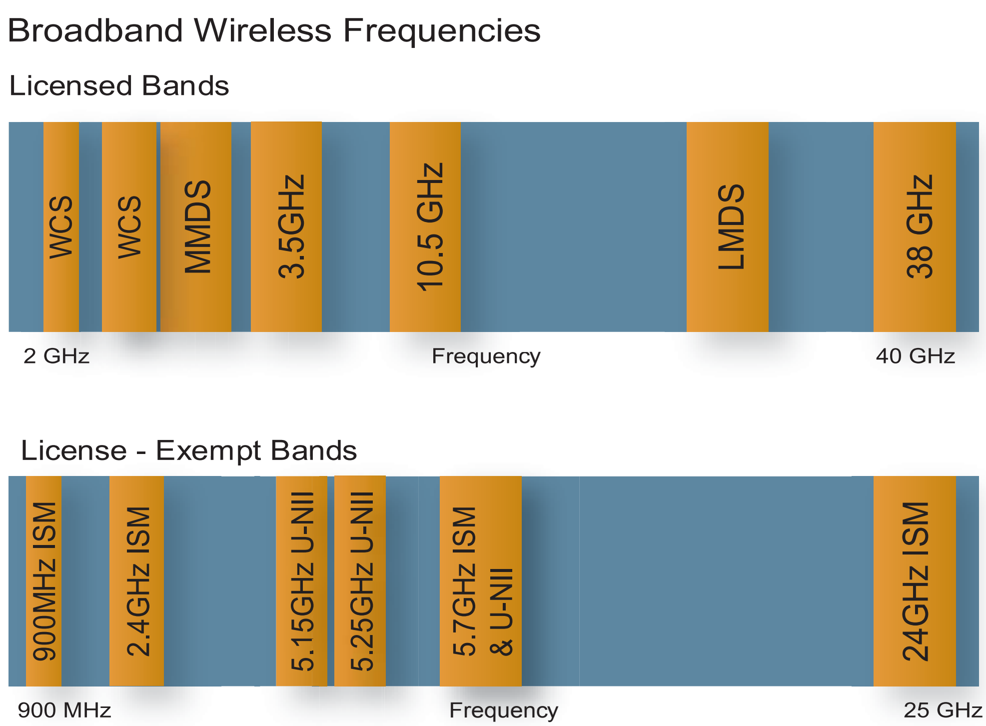

Figure 1. Global Frequency Bands

Interference Must Be Addressed

But -- there are always trade-offs. While operating in a licensed band adds an initial cost to the bottom line, the license itself means that no one else can operate in that frequency. As far as interference goes, the network design must only account for self-induced interference. When deploying a PMP network in the license-exempt bands, interference from your own network is only one issue to be considered; interference from other networks must also be taken into account. These interfering sources may be present at the time of initial network deployment, or they may be deployed a year or more down the road.

In order to ensure that a BWA system will last for a significant period of time providing an adequate return on investment, operations in these bands must take into account all forms of interference. Equally important, customers and end users must have confidence at a high level that the technology driving the network which is delivering business-critical broadband connections is reliable and stable.

When setting up a PTP network in license-exempt bands, interference is still an issue. As will be discussed later, these problems are not as significant as those encountered in PMP networks. However, interruption to these links can often have an even greater impact, as they typically carry traffic for many customers at a time.

The ability of a network deployed in license-exempt bands to tolerate interference - both self-induced and from outside sources - is largely a function of the equipment being used. It is important that the core product itself be designed from the ground up with stability and RF robustness as a key design goal.

The Motorola Canopy Solution

One of the newest entrants into the BWA world is Motorola, who is addressing the problem of delivering a cost-effective BWA system in the license-exempt bands with the Canopy product line. Motorola is well known in the wireless sector, delivering best-of-breed products to the cellular, government, commercial and industrial markets. Motorola has taken the expertise gained in these challenging large-scale deployment models and applied it to the BWA market.

From the start, the overriding design goal of the Motorola Canopy team has been to deliver the most robust simple-to-use BWA system on the market. The Canopy system achieves this goal.

In the following sections, we will review the methods available for combating interference and the Canopy approach, leaving the reader with a deeper understanding of and appreciation for the complexity of the problem - as well as the solution provided by the Canopy product line for both PTP and PMP networks.

What is Interference?

In the wireless world, interference, by definition, is a situation where unwanted radio signals operate in the same frequency channels or bands - i.e. they mutually "interfere," disrupt or add to the overall noise level in the intended transmission.

Interference can be divided into two forms based on whether it comes from your own network(s) or from an outside source. If the interfering RF signals emanate from a network under your control, whether it is on the same cell tower or several miles away, it is termed "self-interference." If the opposing signals come from a network, device, or other source that is not under your control, it is termed "outside interference." Thus, the definition of what type of interference is being combated is not based on technology, but ownership.

In licensed bands, it is self-interference alone that must be taken into account; however given a more or less known operating environment (the radio spectrum will only have signals transmitting that are under control by a single entity) proper product design and network deployment can reduce these interferers to a level where they do not impact network performance.

Self-interference is not a phenomenon that is confined to licensed band operations; license-exempt bands must address the same issues. The techniques and design elements of a given product that serve to reduce and tame self-interference in licensed band operations can be applied directly to the license-exempt systems. Taking advantage of the experience gained in cellular network product design, the Motorola Canopy BWA system embodies many of these same features and deployment guidelines in delivering reliable, self-interference-free operation.

The License-Exempt Challenge

In the license-exempt bands, not only must self-interference be accounted for, but, given the nature of the regulations governing these bands, external interference must be designed for as well. This can be extremely challenging, as there is no way of knowing in advance where these outside signals may be or will be sourced from, or even how strong the interfering transmissions will be relative to the desired transmission. This aspect of the license-exempt bands represents the possible "downside" to license-exempt network operation.

Yet as potentially damaging and unpredictable as external interference can be in license-exempt networks, a properly designed BWA product can make a significant difference in the performance of a network under siege from unwanted external radio transmissions.

In the Canopy system, these issues are well understood and great care was taken when designing this product to make it the most robust BWA system on the market today. This has been accomplished not just through solid RF design, but in more subtle areas such as MAC layer structure and error correction techniques.

License-Exempt Bands

There are several frequency bands that have been allocated as license-exempt in the US. Internationally, these bands generally cover the same frequency spectrum but with some variations in allowed power limits, channel size, etc. Delving slightly deeper, there are two sets of rules applied to the license-exempt bands:

- Industrial, Scientific and Medical (ISM) - three frequency allocations that operate under these rules: 902 MHz to 928 MHz, 2.40 GHz to 2.4835 GHz and 5.725 GHz to 5.85 GHz

- Unlicensed National Information Infrastructure (U-NII) - 5.15 GHz to 5.25 GHz, 5.25 GHz to 5.35 GHz and 5.725 GHz to 5.825 GHz.

The primary difference between the rules governing these bands has to do with whether or not the signal must be spread. In the U-NII band spreading is not required, which generally allows greater data rates than equipment in the ISM bands.

FCC protection – The Unlicensed Bands.
While these are the “un-licensed” bands, there are some FCC rules that provide protection from the most egregious interference. These can be divided into three areas: certification, spreading (ISM only), and power limitations.

- Certification – The intent here is to allow unfettered deployment of new and better technology while preventing the garage ham radio operator from widely deploying nonconforming equipment. ISM devices are governed under Part 15.247 while U-NII devices follow the rules in Part 15.407. The rules do not allow equipment to be sold for revenue until they have been certified by the FCC. This also includes providing a list of tested antennas to ensure power limitations are met.
- Spreading – For devices operating under ISM FCC rules (900MHz, 2.4GHz and 5.7GHz) the FCC requires that the signal be spread either with Direct Sequence (code spreading) or Frequency Hopping techniques. By spreading the signals, it is possible for several different networks to operate in the same area without undue interference. In the U-NII band spreading is not required.
- Power Limitations – The rules here limit both the amount of transmitted power (what comes out the RF analog port) and Effective Isotropic Radiated Power (EIRP) which is the resultant power emitted from the antenna. This includes power from the RF port plus the antenna gain, minus any cable and connector losses. For ISM devices in point-to-multipoint operation, the power limit is one watt and the EIRP is four watts. U-NII power has different levels depending on the specific band, and is tied to the amount of bandwidth used. Only those devices using 20MHz of bandwidth are allowed to use the full power regulated. If a device only operates in 10MHz of spectrum, then the power allowed will be pro-rated accordingly, or one half.

The U-NII spectrum has only recently (within the last four years) been allocated in the US (with international adoption slightly behind), while the ISM bands have been accepted for years. Because of this, in many countries around the world, including the US, operation in the 2.4 GHz band can be very difficult, since a great deal of equipment has been deployed, already crowding the air waves. In addition, numerous consumer devices, such as cordless phones, WLANs, garage door openers and even newer fluorescent light bulbs, are designed to operate in this band, adding to the amount of external interference.

For these and other reasons, most of the newer systems being developed and offered in the BWA market are designed to operate in the U-NII band. The Canopy product is one such system, and in its initial form, is offered in the middle and upper U-NII bands.

Where It Hits and What It Means

Interference in a PMP BWA network, either self-induced or external, is generally an issue more for the Access Point (AP) than for the CPE or Subscriber Modules (SM).

At the AP, typically antennas with much wider angles are used so that they may communicate with many SMs spread over a given geography. Beam widths for these devices can range from 45 to 360 degrees. The wider the angle, the more potential there exists for either self-interference or external interference.

Because a single AP supports dozens, if not hundreds, of end users or customers, interference at this stage in the network deployment can have a large impact. The issue in BWA networks designed to support data or IP-based traffic can be even more insidious. In this instance, a very small amount of RF interference can have a huge performance impact on the network throughput: in some instances three to four percent RF interference can result in a 40 percent reduction in actual end-to-end data rates.

The problem of interference in PTP networks operating in the license-exempt bands, while not as severe as that encountered in PMP networks (due to the use of highly directional antennas at both ends), must still be addressed.

Thus, in a BWA system in license exempt bands, it is critical that the best design possible be used to combat what can be a network-crippling factor. In the ensuing sections, several design factors that impact BWA system performance under the stress of interference are examined in comparison to the Canopy solution.

Methods of Dealing With Interference- The Physical Layer

Modulation and the C/I Ratio

At the most fundamental level, an interfering RF source disrupts the digital transmission by making it too difficult for the receiving station to "decode" the signal. How much noise or interference a digital RF transmission can tolerate depends on the modulation used.

Fundamentally, modulation is the method whereby zeros and ones are communicated by varying one of three aspects of a radio signal. The three portions of an RF signal that can be changed or modulated are phase, frequency, and amplitude. Shifting the properties of any of these parameters can be used to communicate different "states." These states, in turn, are translated to zeros and ones for binary communications.
For example, with frequency modulation, if the sine wave is at frequency one, it is decoded as a zero. If the sine wave is shifted slight-

ly to frequency two, this is decoded as a one. This type of modulation is referred to as Binary FSK (BFSK), or Frequency Shift Keying. In this example, a system must only be able to tell the difference between one of two states or phases. More complex modulations, such as 16QAM (quadrature amplitude modulation), attempt to differentiate among 16 different possible states of an incoming signal.

The advantage to 16QAM is that it conveys more information per bandwidth or more bits/Hz. The disadvantage lies in the fact that, in order to distinguish among the 16 different states, the signal must be very clean and very strong relative to background noise or, even more important, interference.

The ability of a receiving station to decode an incoming signal at the most basic physical layer is dependent on a factor called the "carrier to interference ratio," or C/I. This fancy-sounding term means exactly what it says: how strong the desired signal (the carrier) is relative to the unwanted signals (the interference). C/I ratios are based primarily on the modulation used, with more complex modulations requiring higher C/I numbers than more robust modulations, such as BFSK.

The Canopy product employs BFSK for modulation. With this modulation the C/I ratio necessary to operate properly with an error rate of 1x 10-4 bits per second is only 3dB; i.e. the wanted signal need be only 3dB higher in power than the unwanted interferers. A system operating with 16 QAM at these levels would require a C/I ratio of roughly 12 to 14dB.

Putting this into perspective, with every 3dB of additional signal strength, the power of a signal is doubled. This means that the Canopy system, with its C/I ratio of 3dB, can tolerate an interfering signal that is many times more powerful than a 16QAM system and still operate at the specified error rate. Whether the interference is from another cell site on the network or another network completely, the Canopy system employing BFSK modulation will tolerate substantially higher levels of interference before the communication stream becomes impacted. All other PHY layer techniques are designed to improve this most fundamental measurement of network robustness and operational effectiveness by sustaining the necessary C/I level.

Antennas

When a BWA signal is followed from end to end, it leaves the radio and travels first through a transmitting antenna, over the air to a receiving antenna, and into the radio. The antenna, an important component in the RF chain, can also have an impact on how well the network tolerates interference, both internal and external.

Antenna performance is specified in a variety of ways, but for the purpose of this discussion, the one that is most important is the front- to-back ratio. The front-to-back ratio of an AP antenna indicates how much of an incoming signal will be absorbed coming into the front of the antenna as compared to how much of a signal arriving at the back of the antenna is absorbed.

When deploying networks in a cellular topology, the performance of the antenna in rejecting unwanted signals from behind is an important feature. The Canopy system, with its integrated antennas at the AP, has a front-to-back ratio of 20dB. Coupled with the excellent C/I ratio, this means a Canopy AP receiving a signal at threshold (the weakest signal it can still detect) can be hit with an interfering signal from behind, either internal or external, on the order of -60dBm and still support connections at an acceptable error rate.

TDD Synchronization

BWA networks that use Time Division Duplexing for separating upstream and downstream communications are ideally suited for asymmetric traffic, such as data. The ability to adjust the amount of bandwidth dedicated for upstream and downstream communications without changing hardware is a powerful feature.

TDD systems operate by transmitting downstream (from the AP to the SM) for a period of time -- 1ms for example. Following a short guard time, the SMs then transmit on the same frequency in the upstream. For a cell site with more than one radio operating in TDD mode, it is important that all the sectors of the cell transmit and receive at precisely the same time. Otherwise, if sector 1 is transmitting when sector 2 is receiving, sector 2's incoming transmission can be interfered with even if they are on different frequency channels because the sector 1 signal is so close it is strong enough to "flood" or overwhelm the electronics in sector 2.

When deploying a TDD system in a cellular topology, it is desirable to be able to use the same frequency in each cell site even though those cell sites are possibly several miles away. This means sector 1 from AP A may interfere with sector 1 of AP B. The frequency planning used for the Canopy product is displayed in Figure 2, showing how these signals can interfere. In this case, inter-cellular synchronization is required, making sure that all the sectors in all the cell sites are properly timed and synchronized in terms of downstream and upstream communications.

Delivering tight synchronization across potentially hundreds of square miles can be a challenge. With the Canopy system, designed for large scale, dense network deployments, TDD synchronization is a critical requirement. This has been solved with the use of a GPS signal. These precise satellite signals are used for timing and, ultimately, transmit/receive synchronization, thus tying all sectors in a Canopy network to the same "clock."

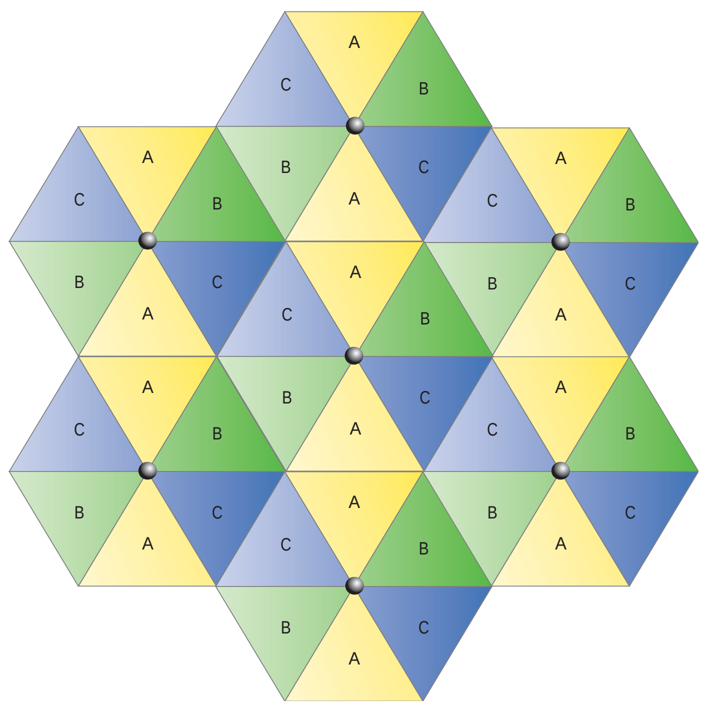

Figure 2. Frequency planning in a cellular deployment.

Methods of Dealing With Interference- The MAC Layer

Up until this point the discussion has centered on various physical layer techniques for addressing interference. Often overlooked, but as important in combating unwanted signals, is the design approach taken in the MAC layer.

Frame/Slot size

A typical MAC frame for a TDD system such as Canopy is shown in Figure 3. As can be seen, the upstream and downstream portions of the frame are divided into slots, each slot carrying what can be termed a "radio data packet," or RDP. The original data, an IP packet datagram, for example, is segmented into packets that fit into a RDP.

Despite all the best deployment design and use of the extremely robust Canopy system, there will be instances where interference will overcome these measures and corrupt a MAC frame or a portion of a MAC frame. When this happens, the corrupted data must be sent

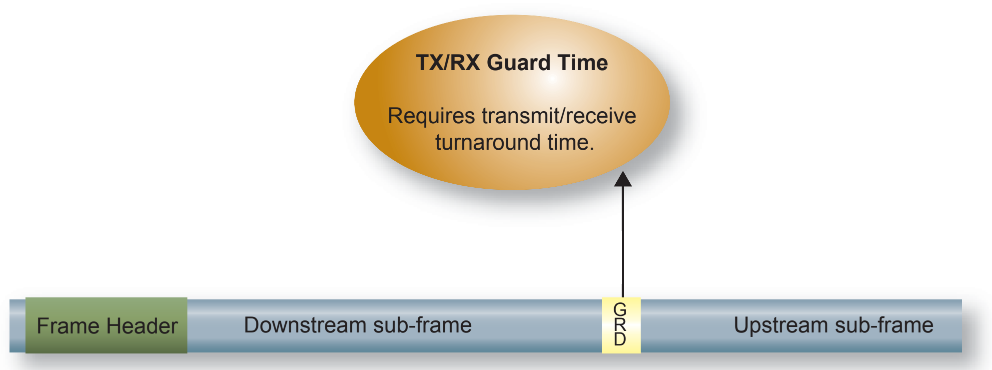

Figure 3. TDD MAC Frame.

again. If the MAC frame is designed for large RDPs on the order of several hundred bytes, the entire slot must be re-transmitted even if only a small amount of this packet is damaged.

The impact on network throughput as a result can be large, with a few bytes in error causing hundreds of bytes to be re-sent. Canopy solves this problem by using RDPs sized at 64 bytes. With this smaller RDP size, the re-transmission can be contained to only those bytes that were damaged, thus avoiding the re-send of large chunks of valid data.

The 64-byte slot could have been made even smaller, but as RDP size decreases, the slot header which is fixed becomes a more significant portion of the packet data, hence increasing the MAC layer overhead. In addition the 64-byte slot is ideally sized for handling the TCP acknowledgements sent for most IP packets.

Automatic Retransmission Request

As discussed earlier, small amounts of interference can have large impacts on end-to-end network performance. This is tied to the way TCP/IP networks were designed to operate in the wired world.

This protocol was designed to operate over wire, where interference was assumed to be negligible. The protocol design calls for a positive acknowledgement sent from the receiving station to the sending station for every IP packet sent out. If the sending station does not receive the TCP ACK in a certain amount of time, it is assumed that the cause was congestion of the network - not an error resulting from transmission impediments. When encountering congestion, TCP responds by dramatically slowing down the transmission and then increasing transmission speed slowly.

In a BWA network, a lost or corrupted packet is not due to congestion but interference. But the TCP protocol has no knowledge or ability to account for higher error rates, and responds by slowing down the end- to-end data rates. This is the phenomenon that can multiply a small amount of RF interference into significant network degradation.

In the Canopy networks, this is not a factor. Recognizing the dilemma of combining TCP/IP with wireless networks and the attendant error rates, the Canopy system solves the problem with a feature called Automatic Retransmission reQuest or ARQ. ARQ actually inspects the RDPs that come into the receiving SM and looks for errors. If an error is detected, the SM (or AP) will send a request to the sending entity to re-send the RDP. All of this is accomplished two layers below TCP in the protocol stack. The net effect is that as far as TCP is concerned, it never receives a packet of data with an error as a result of the wireless portion of the network.

This prevents TCP from invoking the slow start algorithm, keeping the end-to-end data rates at the peak or just slightly below peak operational rates.

Centralized Transmission Control

Some BWA MAC protocols, such as that used in the IEEE 802.11 standard, operate in what is referred to as a distributed control manner. This means that each SM has the ability to send a packet at its own discretion. Typically in this scenario the SM will "listen," and if it does not hear any transmissions, it will assume the channel is clear and send its data.

The problem arises if the sending SM cannot hear other SMs. In this instance, two or more SMs may send a packet at the same time, corrupting both and causing a retransmission. Interference is also a culprit in blocking SMs from hearing each other with the same effect.

Canopy solves this problem by implementing a demand contention access scheme where the AP controls all transmissions in the sector, both upstream and downstream. An SM will only send its data when allowed. If an SM's request to send data is interfered with, it will wait and try again, but at no point will it ever transmit into the uplink data channel until it is permitted by the AP.

Net Effect

When deciding to deliver BWA service to an area or region, the first issue an operator must decide is how they want to deploy the network. Based on the number of potential customers expected, one of two methods is selected - cellular or spot deployment.

PMP Applications – Cellular Deployments

Cellular deployments are chosen when a service provider decides that they want to deliver ubiquitous BWA service to a large area. The benefit of a cellular approach lies in the fact that all of the area will be covered uniformly. The downside lies in the increase in up-front planning and design required.

As depicted in Figure 2, in order to avoid self-interference, the frequencies used in each sector and their orientation from cell to cell must follow strict deployment guidelines. The fewer frequencies that can be used, the easier the planning. Ideally, a good single carrier system, such as Canopy, can re-use the same frequency channels from one cell to another. Note: in Figure 2 depicting a typical Canopy six sector per AP cellular deployment, there are actually only three channels being used in the entire network.

For systems with less robust modulation schemes, the C/I requirements often drive complex frequency plans. This is because before a given frequency channel can be used again at a second cell site, it has to be far enough away to satisfy the C/I requirement. [As noted in sidebar below, every time the distance is doubled, the signal fades by 6dB.] Thus, to ensure the "interfering signal" from cell site A does have not an impact the next time that channel is used, its signal level strength must be reduced. This is accomplished by ensuring it is several cell radii away (actual distance also depends on antennas and other factors such as line of sight) from any of the other sites using that channel.

The range of a given system with a clear Line of Sight path can be calculated as follows. First determine the “link budget” available, than compare to the charts below. Chart one is for operation in the 2.4GHz band and Chart 2 is for the 5.8GHz band. Note the return path is typically the limiting factor and it is suggested that this be used for determining the range.

Link Budget (dB) = Tx power (dBm)
+ Transmitting antenna gain (dBi)
+ Receive antenna gain (dBi)
- Receive Sensitivity (-xxdBm @BER of 10-6)
- Antenna cable losses
- RF Fade Margin (15dB recommended)

Chart 1 – 2.4GHz

Link Budget	101	104	107	110	113	116	119	122	125	128	131
Distance (KM)	1	1.5	2	3	4	6	8	11	16	23	32

Chart 2 – 5.8GHz

Link Budget	101	104	107	110	113	116	119	122	125	128	131	134	137	139
Distance (KM)	0.4	0.6	0.8	1	1.7	2.5	3.5	5	7	10	14	20	27	32

The net effect for a cellular deployment with a system using higher-order modulations typically means more channels are required in order to satisfy the C/I ratio.

Thus the ability to deploy efficiently in a cellular format depends on many factors, such as the C/I ratio, antenna performance, and TDD design and synchronization. All of these factors, as well as ease of design for the network, are clearly understood by Motorola and the Canopy design team, and have been addressed in the Canopy product.

PMP Applications - Spot Deployments

Many BWA deployments start out being installed in what is referred to as a "spot deployment" model. This topology refers to a single cell site, or possibly several, that are not geographically contiguous but are chosen to serve specific areas of need. This is contrasted with a cellular deployment approach in which the goal is to provide BWA coverage across an entire region and hence the cells or AP sites are deployed such that there are no LOS gaps in coverage.

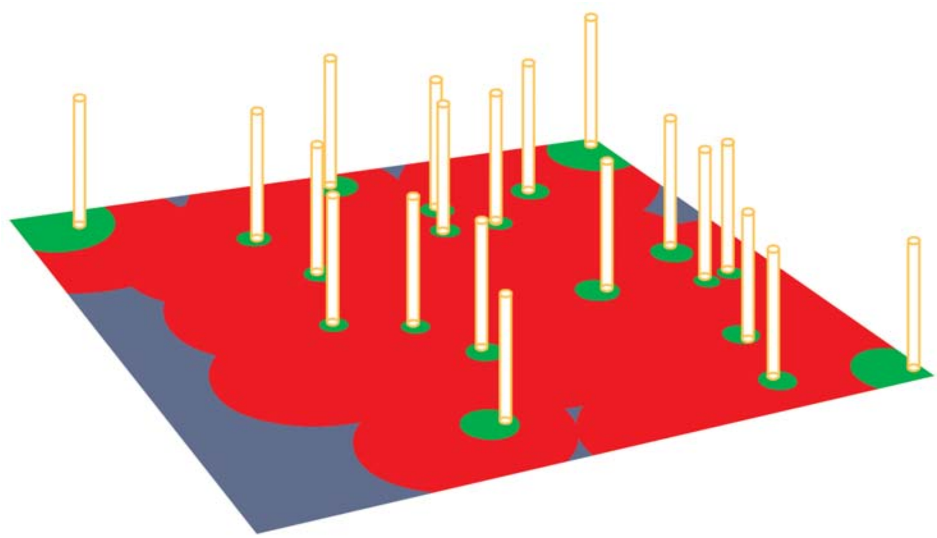

Figure 4a. Poor coverage in a dense deployment.

When a BWA system is deployed in a spot method, frequency coordination and planning are not usually an issue, and each cell site is installed with what is best for that area of coverage alone as a deployment guide.

For spot deployments, a problem arises when an operator starts adding many AP central locations to an area. As is explained in the

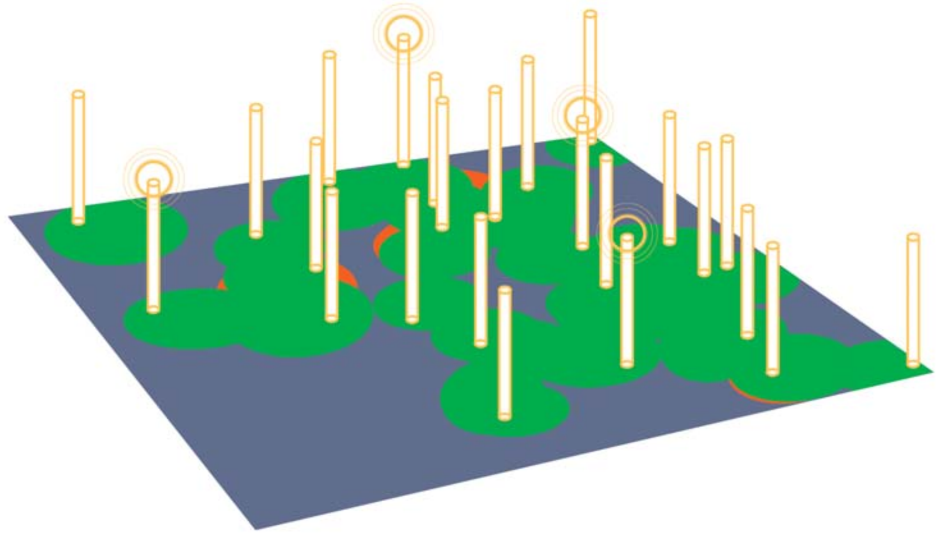

Figure 4b. Canopy coverage in a dense deployment.

cellular deployment model above, self-interference from neighboring AP's or SMs on the same frequency channel rapidly becomes a problem. Because little if any frequency planning was performed as the sites were rolled out, the problem with self-interference can become very serious very quickly.

The result is that operators will be severely constrained in where and how close they can install additional AP sites. As is shown in Figure 4a, this will result in large coverage gaps. The Green shaded areas can be served by the AP sites shown. Red represents areas that can't be served due to interference between AP sites. Large gaps in coverage not only prevent a service provider from serving potentially lucrative areas of interest, but also leave the public at large with a poor perception of the service. It only takes a few potential customers to hear, "I'm sorry, we cannot serve your business" to get the word out that the service is "not real."

The Canopy system's design minimizes these impacts, allowing the service provider to grow the network far beyond less-robust BWA systems. With Canopy, an operator utilizing a spot deployment methodology will be able to significantly increase the number of AP sites without the coverage gaps of other BWA networks. Fundamentally, more AP sites means greater coverage, which means more customers per AP. The net effect is simple -- more revenue per square mile. Figure 4b shows a dense deployment using the Canopy system. Note that vast increase in coverage area (shown in green).

PTP/Backhaul Applications

While most of the talk in BWA surrounds PMP networks and products, a large percentage of these systems will be located in areas where the infrastructure is not well developed. The location of the AP site is chosen based on where potential customers are located, where a good high tower or building can be used, etc. What does not factor into this decision is the question of where is there a good connection to the Internet or a private network.

Most times the AP will be located such that there is no significant copper or fiber connection into the WAN. It almost seems ludicrous that BWA networks can be hindered by the lack of wire-line infrastructure to connect the data aggregated at the AP to the WAN.

Nonetheless, it is a reality that for BWA systems to be effective, they need to have a PTP solution in the arsenal to provide the AP-to-WAN connection when no wire-line options exist. Figure 5 depicts the Canopy solution. The issue with BWA PTP links lies in the fact that they carry all the traffic for a AP site, or data from possibly hundreds of customers. The net effect is that this link becomes both mission critical and a single point of failure - a dangerous combination.

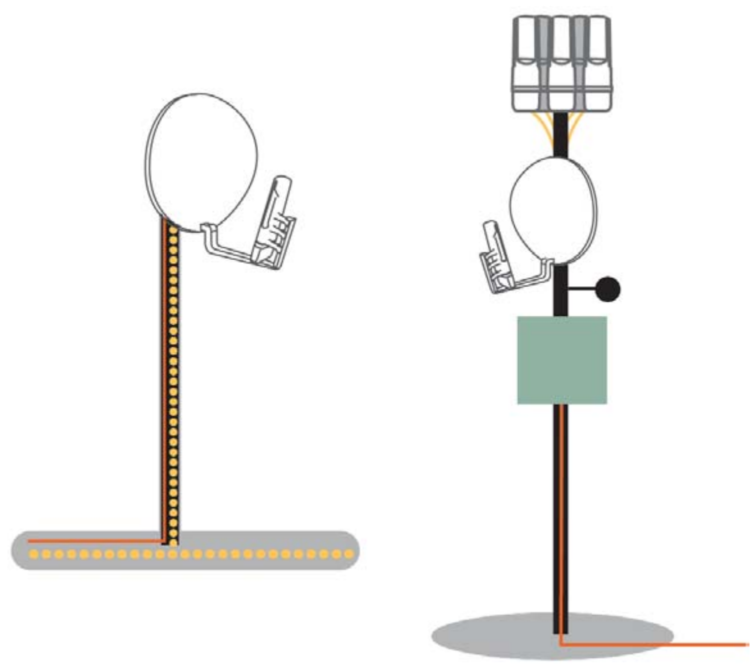

Figure 5. Canopy AP site with backhaul.

For the same reasons a service provider may have chosen license-exempt equipment for their BWA deployment, they will likely need a license-exempt solution for the backhaul or PTP link. This places a burden of reliability and robustness on the PTP connection.

All of the features described previously that are implemented in Canopy to "harden" it against interference in PMP applications are applicable to the PTP application. While the likelihood of interference in PTP links is substantially less (due to two narrowly focused antennas, as discussed above) the impact is substantially greater. When the Canopy solution is used in a PTP application, the system employs a slightly different MAC, reducing overhead and delivering more useable or sellable traffic on the link.

The Canopy PTP solution is not limited to just backhauling traffic from Canopy cell sites, although they have been designed to work together interference free. With such a robust approach to wireless communications, the Canopy PTP system has the reliability necessary for connecting cellular telephony as well. Indeed, any time a high capacity dedicated connection that has increased reliability requirements occurs, the Canopy PTP solution can serve the need. Schools, government, enterprise campuses separated by a road - all are candidates for this approach. No longer will leased lines be needed to support mission-critical applications between two points. Not only are PTP wireless links more economical, they are also more scalable, allowing for users to deploy the amount of bandwidth needed.

Summary

The world of broadband everywhere at this point is more of a dream today than a reality. While wire-line initiatives are being deployed rapidly, they alone cannot bring broadband to everyone, for many reasons. Cost, technology and, fundamentally, sound business judgment means that if wire-line alone was the broadband solution, huge numbers of users would be left out for years to come.

BWA, with all of its inherent promise in rapid deployment and cost-effective delivery of broadband, is a reality. However, BWA systems have their own hurdles to overcome, with reliability, or perception of reliability, being at the top of the list. Interference lies at the heart of the reliability challenge, and interference in the license-exempt bands can be a much greater factor than that faced by licensed band systems.

It is a fact that license-exempt solutions have an opportunity - and a challenge - in becoming accepted as a broadband option. In order to address perception and reality, these networks must deliver on the promise of long-term, robust operation, bringing revenue assurance to the operator and reliable service to the end user.

To that end, it is critical that BWA solutions designed for the license-exempt bands address this issue head on. It is also clear that in order to do so, proper design at a very detailed level must be accommodated in the core of the product. Solid, reliable BWA networks do not happen by chance; they are a result of keeping a focus on the issues and delivering the right solutions. The Canopy product is an industry leader in this area. Whether one's need is mission-critical high volume traffic (Canopy PTP) or serving large numbers of customers in a geographic region (Canopy PMP), operators can be assured their networks will last for years.

MOTOROLA

WHITE PAPER
August 2002