The Invisible Threat: Interference and Wireless LANs

A Farpoint Group White Paper

Document FPG 2010-134.2
May 2010

Farpoint Group

Interference is a fact of life in the unlicensed bands used by wireless LANs (WLANs), and is an increasing challenge in all WLAN environments – the enterprise, outdoors (including metro-scale Wi-Fi meshes), and even in the residence. As the number of unlicensed devices grows and as ever more mission-critical applications are deployed on WLANs, interference represents a challenge that must be addressed.

Farpoint Group has been studying the impact of interference on wireless LANs for a number of years, and we have empirically measured the result in a variety of situations. We have learned that many forms of interference can have a detrimental and even *destructive* impact on WLAN traffic, degrading both data throughput and time-bounded traffic typified by voice and video. In addition, we have analyzed the impact of a variety of interferers often seen in the enterprise and have evaluated their effect on WLAN performance. The results of this work are available in a series of Farpoint Group Technical Notes listed in Appendix I of this document.

With the unlicensed bands available to many devices beyond WLANs, these potential sources of interference are, like WLANs themselves, increasing in number. Several vendors have responded with a new class of *spectral assurance (SA)* tools for dealing with this challenge, promising far-reaching benefits for WLAN systems and their users. Based on the concept of the spectrum analyzer long used by engineers, these new tools are WLAN-oriented and designed for use by IT staff who may not have engineering backgrounds and experience. And, most exciting, these tools are now moving into WLAN infrastructure with far-reaching benefits.

This White Paper discusses the threat that interference represents, and how spectral assurance tools can yield a significant improvement in reliability for network managers in enterprises of all sizes.

Radio-Frequency Interference and the Unlicensed Bands

Radio-Frequency Interference (RFI) is a major concern in the deployment and use of WLANs, and is often cited as a justification for avoiding their installation altogether. As we noted above, WLANs operate in the *unlicensed bands*, spectrum reserved by regulators worldwide for applications without a requirement for individual user or device licensing. A consequential challenge in utilizing these frequencies is that a potentially large number of wireless devices may be competing for the airwaves in a particular location, often resulting in interference and thus degraded user connectivity in terms of throughput, link quality, and range.

Regulations require unlicensed devices operating in these bands to *accept* any interference that may be present, and most interference in the unlicensed bands is in fact *unintentional* – the result of other devices operating legally in this spectrum. Interference may also originate from certain licensed services, including amateur radio equipment, RADAR systems, and a variety of other devices, operating at much higher power than is allowed for unlicensed products. These signals can be quite damaging indeed to unlicensed band transmissions. WLAN devices can also be subject to *intentional* interference, also known as *jamming*. While such is rarely encountered today, the potential for jamming exists and must be managed as any other risk to network integrity.

Regulators created the unlicensed bands to promote the use of low-power (and thus limited-range) radio-based equipment and to minimize both bureaucratic impediments to deployment and end-user requirements. The rules and regulations governing the unlicensed bands are similar (but not identical) throughout the world. The Federal Communications Commission (FCC) has jurisdiction in the United States; the applicable rules can be found in the *Code of Federal Regulations (CFR) Title 47, Telecommunications*, with specific rules for WLAN-based wireless LANs found in Parts 15.247 and 15.407 [http://www.access.gpo.gov/nara/cfr/waisidx_05/47cfr15_05.html] of these Regulations. The rules primarily specify the applicable frequency bands, power output limitations, and a wide variety of technical and other parameters including limitations on coordination of devices and requiring the use of spread-spectrum radio techniques in most cases.

The FCC's policy on sharing unlicensed spectrum effectively leaves it to industry to work out the details regarding interference. As was noted in an FCC Technical Advisory Committee report in December 2000 [http://www.fcc.gov/oet/tac/tac7report.pdf], "We are about to have an unplanned real-time experiment on the consequences of uncoordinated spectrum sharing by different services using incompatible etiquette rules." Thus far, the experiment has clearly been wildly successful, with billions of unlicensed devices in use today. But some network managers and IT analysts are rightfully concerned that there may soon be so many unlicensed devices operational that the unlicensed bands might no longer be useful in some locales - or at least not practical for mission-critical, time-bounded, or high-bandwidth applications. Indeed, the above FCC report even mentioned Yogi Berra's oft-quoted quip about a restaurant being so crowded that "no one goes there anymore." While the limited range (distance between endpoints) of unlicensed devices mitigates the possibility of severe interference to some degree, we are already seeing the effects of crowded airwaves in some venues today - and cumulative energy levels and transmit duty cycles continue to increase.

Understanding the Impact of Interference

Interference occurs when two radio signals with sufficient proximity to each other are transmitted on the same frequency at the same time. Interference can occur if the two (or more) simultaneous signals have similar relative transmit power, in which case they will likely mutually interfere, or if one signal has relatively greater power, in which case weaker signals will suffer (perhaps severe) interference from the stronger. Note that radio waves fade (lose power) exponentially with distance, an effect known as the *inverse power law*. As a given radio wave moves from transmitter to receiver, it can thus change from interferer to interferee - the signal might initially have enough power to damage another nearby in the same spectrum, and then, as it fades, it might for a time be at the same relative power level as another signal, with mutual interference the outcome. Finally, the signal might fade enough that a nearby stronger signal might present destructive interference to it.

Interference is moreover a function not only of *relative power*, but also *transmit duty cycle*, the percentage of time that a given device is actually transmitting, with a larger share here resulting in a greater probability of interference. It is possible for two otherwise potentially interfering signals to "timeshare" a given frequency (in an uncoordinated fashion, of course), resulting in relatively little mutual interference. But note also that except in the case of jamming, interference is (often maddeningly) intermittent in nature, making it very difficult to detect and analyze without the right tools.

In practice, interference in WLAN applications usually manifests itself as reduced data-traffic throughput, less effective range, and impaired quality of service (QoS) for voice and video applications, but can also include the complete - if temporary - failure of a given link. The cumulative effects of interference may be identifiable by analyzing network management logs, but diagnosing these symptoms in this manner can be very difficult because they can also result from other network-related problems. This situation further motivates the use of specialized PHY-layer tools for identifying and evaluating the sources and effects of interference.

With respect to WLANs, interference can come from a variety of sources. Interference from other WLAN networks is typically *co-channel* interference (CCI), usually between two access points on the same channel, or *adjacent-channel* interference (ACI) resulting from two access points operating in close proximity on abutting or overlapping channels. Since WLANs employ a “listen-before-talk” protocol, based on Carrier Sense Multiple Access with Collision Avoidance *(CSMA/CA)*, any interference between WLAN networks tends to work out somewhat cooperatively, with the two networks often sharing (albeit sub-optimally) channel capacity as noted above. In contrast, interference from non-WLAN sources, which use protocols different from those of WLANs, more often result in the severe degradation of WLAN transmissions. There are numerous non-WLAN devices that operate in the unlicensed bands, including Bluetooth products of many forms (some operating at the same power levels as WLANs), cordless phones, wireless video surveillance cameras, wireless security and energy management systems, proprietary wireless bridges, and computer peripherals such as cordless mice, keyboards, and game controllers. In addition, there are potentially-interfering emissions from commercial and industrial devices such as microwave ovens, certain RADAR systems, and even microwave-based lighting.

While market-research numbers vary, it is quite clear that the number of WLANs will continue to grow enormously well into the future. Farpoint Group estimates that less than half of all enterprises globally have deployed a WLAN for general office use. We further believe that the convenience of mobile computing, the low cost of WLAN technology (essentially free, in the case of clients, anyway), the constantly-improving price/performance of WLANs, a reduction in wired-network maintenance costs (via the use of wireless at the network edge, where wiring costs are higher), and the dramatically higher performance of new WLAN technologies (most notably 802.11n) are leading to WLANs becoming the primary and even *default* network connection, for both voice and data, over the next few years. Advances in VLSI implementations of 802.11 radios and related components will further spur WLAN deployments, especially in the form of dual-mode cellular/voice-over-IP-over-Wi-Fi *(VoFi)* handsets. We see these devices entirely *replacing* desktop phones, via *Fixed/Mobile* and *Mobile/Mobile Convergence (FMC/MMC)* and *mobile unified communications (MUC),* for most professionals and essentially for every worker not tied to a given location by the specifics of their work.

These factors, coupled with increasing deployments of metro-scale and other public and private WLANs and the lack of radio coordination inherent in these bands, create the opportunity for interference to become a *major* concern for enterprises, governments, and organizations of all sizes. Moreover, it is likely that residential WLAN deployments, now a practical vehicle for the real-time transfer of large, time-bounded data objects like video (and even HDTV video), will also begin to suffer from the effects of interference as the residential WLAN similarly becomes the default. The issue is ultimately not one of *security* - the traditional nightmare for network administrators - but rather of the fundamental *integrity* and *reliability* of the network itself. Fortunately, a number of tools and approaches are now available to help network administrators effectively manage this invisible threat.

Managing the Interference Threat

Regardless of frequency, even on licensed bands, no radio signal is entirely immune to interference. Farpoint Group believes that there are two key components to effective interference management: *continually monitoring for interference* (this monitoring includes identifying the source of any interference that threatens the integrity of a given WLAN) with spectral-assurance tools, and then *taking steps to mitigate any interference discovered*.

One approach to dealing with interference is to move to another channel or band, most obviously the 5 GHz. spectrum used by 802.11a and many 802.11n products. Farpoint Group often recommends deployment here, and not just because this spectrum is currently less likely to suffer from interference. There are 23 (20 MHz.) non-overlapping channels defined in this spectrum in the US (as compared to just three at 2.4 GHz.), offering significantly more uncongested capacity. The 5 GHz. bands have been underutilized primarily due to a lack of familiarity on the part of users, and a general belief that transmissions at 5 GHz. have less range than those at 2.4 GHz. While it is true that 5 GHz. signals do not propagate as far as signals at lower frequencies at any given transmit power level, we have found that the throughput of 5 GHz. networks to be as good as or better than that of any 2.4 GHz. network at any given operating range. Additionally, we recommend a strategy of *dense deployments* (see Farpoint Group White Papers 2004-193.1, *Rethinking the Access Point: Dense Deployments for Wireless LANs* and 2005-083.1, *Wireless LAN Dense Deployments: Practical Considerations* for more information on this topic), as opposed to attempting to optimize for maximum coverage for each AP. This strategy makes any reduced range at 5 GHz. inconsequential in enterprise settings. But just as WLAN products migrated from 900 MHz. to 2.4 GHz., so too will they move, driven by the need for more capacity, from 2.4 to 5 GHz. Interference monitoring and mitigation techniques will thus still be critical in the 5 GHz. spectrum, and, of course, at 2.4 GHz. as well - there will be many devices operating in the 2.4 GHz. bands for some time to come, including legacy devices, VoFi handsets, and Wi-Fi-based location and tracking tags, so it behooves us to address the interference challenges in this band regardless.

WLAN system vendors have long been cognizant of the issues surrounding interference, and have taken steps to attempt to deal with the problem, albeit in a coarse-grained and WLAN-traffic-specific manner. The most common approach has been to use *RF Spectrum Management (RFSM,* also called *Radio Resource Management)* tools, which are present in most contemporary enterprise-class WLAN systems. These tools enable the (in many cases, *automatic*) management of the PHY in much the same way that other networking equipment enables the management of the upper layers of the network protocol stack. While there are many possible functions in RFSM, the most important are the automatic setting of channel assignments and transmit power levels, and the re-configuration of these parameters as radio and network conditions change over time. All RSFM tools are useful, but only the most sophisticated RFSM implementations can make decisions based on non-WLAN traffic, mostly relating to a gross estimation of "noise". Because of this limitation, most RSFM tools turn out to be quite limited in scope for managing and mitigating interference. However, as we'll discuss below, RFSM tools are expected to broaden in scope over the next few years. More information on RFSM can be found in Farpoint Group White Paper 2003-201.1, *Beyond the Site Survey: RF Spectrum Management for Wireless LANs.*

But this brings up an important point - since a WLAN radio can only detect a WLAN signal, the radios used in WLAN APs and clients are usually not very useful for diagnosing non-WLAN interference. As we discussed earlier, there is an ever-growing list of non-WLAN devices, including cordless phones and Bluetooth devices that can create interference problems for WLANs. The device typically used today to identify these arbitrary wireless signals is a specialized radio receiver called a *spectrum analyzer*. These (usually quite expensive) pieces of test equipment look a lot like oscilloscopes and require an appropriate engineering background for effective use.

The core problems with most spectrum analyzers are their inherent difficult-to-use-for-non-engineers nature and their cost. Good spectrum analyzers are often priced at US$20,000 or more, as they are sensitive, calibrated test equipment designed primarily for component and product-engineering applications. Since interference can creep into a given facility at any time, it would be nice to be able to continually monitor for this eventuality - but the above two factors essentially eliminate this possibility with traditional spectrum analyzers. A third major issue is their lack of specificity to WLAN-related situations, which limits their practical application in the enterprise.

Fortunately, progress in VLSI, spectrum analyzer architecture, and associated software has resulted in a new class of WLAN assurance capability – what are known as *Spectral Assurance (SA)* tools, essentially spectrum analyzers designed for WLAN applications. These products combine spectrum analysis with the ability to determine if interference is causing a problem on the WLAN, identify and fingerprint specific interfering devices, and locate those devices. The first of these is *Spectrum Expert*™ from Cisco [http://www.cisco.com/en/US/products/ps9393/index.html], which can be seen in Figure 1. This is a simple yet very powerful PC Card-based product, frequently used with a clip-on external antenna, and based on a

Figure 1 - Cisco's *Spectrum Expert* is the first spectrum analysis product designed for WLAN applications. *Source:* Cisco Systems.

Spectral Assurance in the Infrastructure: Cisco's CleanAir Product Line

Cisco's recently-announced CleanAir product line is based on a new family of access points - the 3500 series - that integrate dedicated, specialized functionality just for spectral analysis. Coupled with additional capabilities in the company's Wireless Control Systems (WCS) management console (see Figure 2), enterprises can now gain the benefits of continuous, venue-wide, and even multi-site monitoring. The cost advantages here, as we have discussed, are mostly obvious - rather than sending out a trained (and expensive) engineer with a hand-held spectrum analysis tool when interference is suspected, operations staff at a central console can now be alerted to spectral issues, including the detection and classification of interference sources, quickly and automatically. A perhaps less-than-obvious benefit, though, is a reduction in a wide variety of related opportunity costs - including lost productivity, regulatory violations, and customer-facing failures - that could have very severe impacts indeed on the entire enterprise. The addition of Cisco's Mobility Services Engine (MSE) (see Figure 3, which also illustrates all members of the CleanAir product family) adds automated location of interferers, and a pervasive deployment of 3500s enables automatic remediation of spectrally-related challenges. We expect infrastructure-based spectral assurance to be one of the most important trends in wireless LANs over the next few years - the benefits (including the potential for much lower operational expense as discussed in this White Paper) are clear, and the challenges arising from not having this capability potentially quite damaging to many aspects of an enterprise's operations.

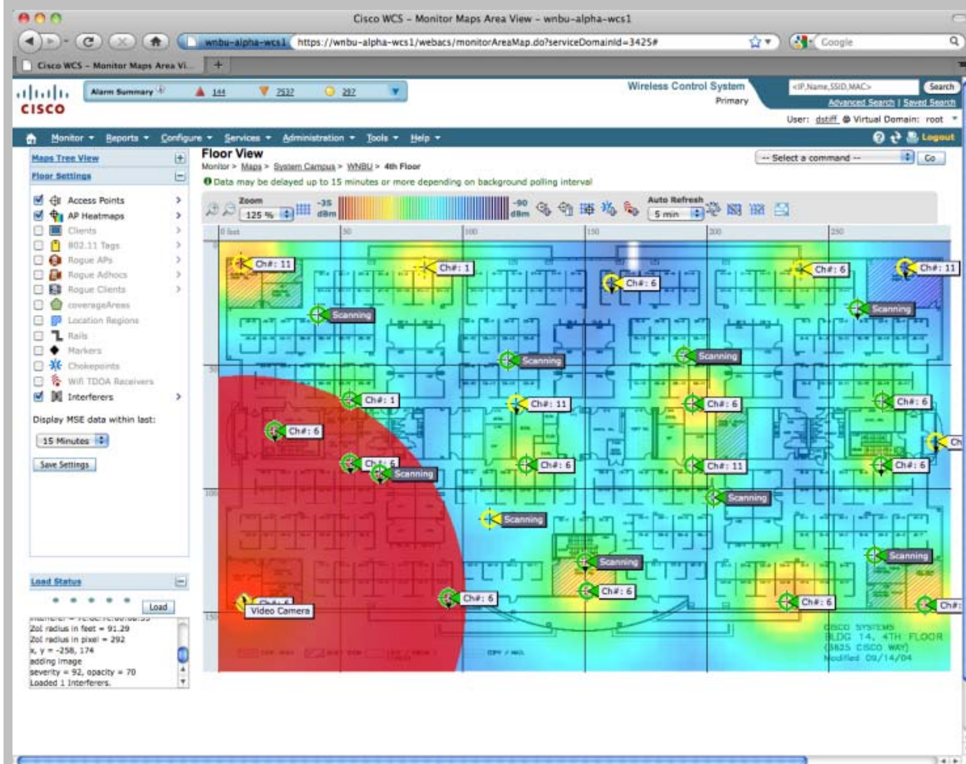

Figure 2 - Screen shot of Cisco's WCS showing the impact of an interferer (large red circle) on a particular installation. *Source:* Cisco Systems.

Figure 3 - The Cisco CleanAir product family includes the Aironet 3500i Access Point, Cisco 3500e Access Point, 5500 Series Wireless Controller, Catalyst 6500 Series Wireless Services Module (WiSM), and Mobility Services Engine (MSE). *Source:* Cisco Systems.

custom spectrum-analyzer-on-a-chip developed by Cisco. Coupled with a broad set of comprehensive and flexible software, Spectrum Expert defined a new and very cost-effective WLAN spectrum assurance solution that is very popular in enterprise settings. Farpoint Group regularly uses Spectrum Expert and highly recommends the product.

A final point - Farpoint Group believes that a “Spectrum Survey”, which is an “RF sweep” of a given location prior to the installation of a WLAN, is often very valuable in identifying possible sources of interference. This exercise involves sampling the spectrum at various locations using the spectrum assurance tool, looking for levels of energy that, irrespective of source, might prove detrimental in a production WLAN environment. Similarly, we will occasionally perform a post-installation RF sweep if interference is suspected at that time. We believe however, that continual monitoring with spectral assurance tools within the WLAN infrastructure (see Sidebar, *Spectral Assurance in the Infrastructure: Cisco’s CleanAir*) is going to become the norm over time – and, indeed, essential to the success of large-scale WLAN installations.

Conclusions and Current Directions

Cisco integrated Spectrum Expert into their Wireless Control System (WCS), under the label of *Spectrum Intelligence*, and now, with CleanAir, has introduced continual, automated spectral monitoring, interferer location, and remediation capabilities. This development enables much greater and more convenient monitoring, control, and interference-resolution possibilities, optimizing the benefits of having both protocol- and energy-based analysis within a single framework. The IEEE 802.11 is also active here, with work completed on the 802.11k (Radio Resource Measurement) standard, and continuing within the 802.11v (Management) Task Group. Spectrum analysis and assurance represent one of the most exciting and, we believe, ultimately beneficial areas of wireless LAN innovation today. We are beginning to see spectral assurance as an integral feature in network management systems, automatically working around interference challenges with little, if any, manual intervention. Regardless, other steps, such as identifying and moving interfering devices, replacing them with non-interfering equivalents, and similar measures, remain good practices, but automated awareness and remediation are clearly the most desirable route.

This White Paper has discussed radio interference in the unlicensed bands especially with respect to wireless LAN deployments. We have reviewed the tools and techniques available for addressing the challenge of RF interference, and we have outlined methodologies that will enable large-scale wireless LANs systems to continue to expand with all of the convenience and performance inherent in the promise of wireless networking. While radio-frequency interference will always remain a concern, we believe that we now have the tools to render this situation more than manageable. Thanks to new technologies like spectral assurance, RF interference will be handled effectively and often automatically in the course of normal enterprise network operations. We will, of course, continue to monitor developments in this space and will report new advances as they occur.

Appendix I: For Further Reading

Farpoint Group has spent significant time gathering empirical data on the nature and effects of interference in the unlicensed bands, particularly with respect to wireless LAN systems and applications. The following Technical Notes are available for those who want to explore this subject in more depth:

- Farpoint Group Technical Note 2010-135.1, *Evaluating Interference in Wireless LANs: Recommended Practice* (April 2010)
- Farpoint Group Technical Note 2006-328.3, *The Effects of Interference on General Wi-Fi Traffic* (January 2008)
- Farpoint Group Technical Note 2006-329.3, *The Effects of Interference on VoFi Traffic* (January 2008)
- Farpoint Group Technical Note 2006-330.3, *The Effects of Interference on Video Over Wi-Fi* (January 2008)
- Farpoint Group Technical Note 2006-373.3, *Interference and Metro-Scale Wi-Fi Mesh Networks* (January 2008)

Farpoint Group

Ashland MA 01721
508-881-6467
www.farpointgroup.com
info@farpointgroup.com