Attachment A

Analysis of the Interference Temperature Concept to Support Spectrum Sharing Between Licensed Services and Unlicensed Devices

Dr. Jay E. Padgett
Dr. Robert A. Ziegler

Telcordia Technologies, Inc.
Applied Research
Wireless Systems and Networks

April 5, 2004

Executive Summary

In its recent Notice of Inquiry/Notice of Proposed Rule Making (NOI/NPRM), the FCC has proposed to use an "Interference Temperature" (ITemp) metric to facilitate the sharing of licensed spectrum by unlicensed devices. The basic concept is that the unlicensed devices would be allowed to generate, into licensed receivers, an aggregate interference level up to some threshold. Hence, there are two main components associated with implementing ITemp: (1) setting the appropriate interference threshold; and (2) regulating the aggregate interference so that it does not exceed the threshold. How these two things are done will depend on the nature of the affected licensed service; there is no single universal ITemp implementation.

This paper explores ITemp implementation in general and provides detailed analyses for two specific cases: mobile radio and fixed point-to-point services.[1] The general conclusions, based on the analysis of interference statistics, show that monitoring by the unlicensed device itself, or a network of arbitrarily-placed ITemp "thermometers," would be ineffective. The mobile radio analysis shows that ITemp is unworkable on the downlink and impractical on the uplink, and that even if a perfect theoretical implementation on the uplink is assumed, the addition of the unlicensed devices would cause an overall degradation in spectrum efficiency, accounting for both the added capacity of the unlicensed devices and the reduction in licensed capacity. For point-to-point services, the transmit power control (TPC) and dynamic frequency selection (DFS) approaches proposed in the NPRM will not protect the licensed receiver. The unlicensed devices must be required to operate beyond the radio horizon of any cochannel fixed receive stations.

[1] The analyses apply to system types rather than to specific frequency bands. For example, the mobile radio analysis is relevant to MDS and PLRMS as well as CMRS.

ÓTelcordia Technologies 2004

To implement the ITemp concept, it would be necessary to control the aggregate unlicensed interference *as seen at the licensed receiver.* One way to do this is to use a closed-loop approach, whereby a monitoring receiver tracks the aggregate interference level. That level would be broadcast over a feedback channel to the unlicensed devices, which then would decrease their transmit power levels as necessary to maintain the total interference below the threshold. However, for this closed loop approach to work, the interference power seen by the monitoring receiver must be highly correlated with the interference into the victim licensed receiver. This requirement is satisfied if the monitoring receiver is co-located with the victim receiver (and has the same antenna pattern). There may be some flexibility in locating monitoring receivers if an "exclusion zone" (an area within which unlicensed devices may not transmit) surrounds the victim licensed receiver. In that case, the monitoring receiver need not be exactly co-located with the victim receiver, but must be near it relative to the exclusion zone radius (*i.e.*, well away from the exclusion zone boundary).

There are three significant conclusions that follow from the foregoing observations. First, if there is no exclusion zone, then monitoring the interference at a location other than that of the victim receiver is not useful, because the measured interference will be uncorrelated with the interference to the victim receiver. Second, "self monitoring" by the unlicensed transmitter itself cannot be effective; to measure the interference accurately, the unlicensed transmitter would need to be near the center of the exclusion zone, but by definition, the unlicensed would not be allowed to operate there. Third, for any monitoring to be effective, the locations of the victim receivers must be known (and generally stationary as well). Otherwise, the monitors cannot be located sufficiently close to the victim receivers to be effective.

From these conclusions, it follows that there are some types of licensed receivers for which ITemp monitoring-and-feedback cannot be reasonably implemented under any circumstances. Prominent among these are most types of mobile transceivers as well as broadcast receivers, for which an exclusion zone cannot practically be implemented. The locations of such receivers are typically unknown. Broadcast receivers have no associated transmission capability, and mobile radio units often are idle (not transmitting), which means the proximity of a licensed receiver cannot reliably be estimated by measuring its associated transmitted signal.

Therefore, short of integrating the ITemp monitoring-and-feedback capability within the affected licensed devices themselves (which is often not feasible), the closed-loop form of ITemp is limited to licensed receivers for which the monitoring receiver can be either co-located or located well within the boundary of any exclusion zone that might exist. This means that a network of arbitrarily-placed monitoring receivers will generally not be an effective way to implement closed-loop ITemp.

Significant technical challenges exist even in cases for which the nature of the licensed receivers will accommodate co-location of a monitoring receiver. For example, the monitoring receiver would need to distinguish the aggregate unlicensed interference from

the co-channel licensed signals as well as background noise, both are which are likely to be significantly higher than the allowed ITemp interference level.

Despite such implementation design problems, this paper also examines the tradeoffs between the benefits gained from the additional unlicensed usage of the spectrum and the degradation experienced by the "host" licensed service. This paper analyzes in detail the impact of ITemp on the reverse link (uplink) of a commercial mobile radio services (CMRS) system that uses code-division multiple access (CDMA) technology. In such a case, the ITemp monitoring receiver would be co-located with the base station antenna. For purposes of the tradeoff analysis, it is assumed that the idealized closed-loop ITemp system is able to regulate perfectly the total unlicensed interference into the CDMA base station receiver by the monitoring/feedback process.

Not surprisingly, the capacity of the CDMA uplink is reduced by the unlicensed interference, and the degree of the reduction depends on the allowed ITemp threshold. In return, unlicensed capacity is gained, but the performance of the unlicensed devices is limited by the fact that they would sustain interference from the CMRS handsets. It is clear from the results of the analysis that based on overall efficiency of spectrum utilization, ITemp sharing is a losing proposition in this case, even assuming that the implementation mechanism could operate perfectly and at zero monetary cost (and disregarding any additional spectrum that might be needed for feedback signals to the unlicensed devices). The value lost in terms of CMRS capacity exceeds what is gained in unlicensed value. Fundamentally, this inefficiency is due to the mixing of unlike systems. The conclusion is that the ITemp concept is neither useful nor workable for mobile services.

The advantage of the closed-loop form of ITemp is that, at least in principle, it should be able to guarantee that the interference threshold is not exceeded. However, under certain conditions, other forms of ITemp might be feasible. The NPRM gives two examples: the fixed satellite service (FSS) and fixed point-to-point terrestrial services. In the FSS case, the satellite receiver tends to see the aggregate interference from many unlicensed devices over a large area, but there is a large guaranteed physical separation and the path loss is very high. Because of the large number of devices involved, the statistical variation of the interference relative to its average will be very small, and the average can be used to calculate the interference impact on the satellite receiver, as was done in the NPRM. Provided that the total number of transmitting devices within view of the satellite does not exceed some limit, the total interference should not exceed its threshold.

For a fixed point-to-point one-way link, the NPRM proposes an open-loop approach whereby the unlicensed device would measure the power received from the licensed transmitter and then add a fixed increment to it (in dB) to calculate the allowed transmit power. This is intended to ensure that the interference received by the licensed receiver (at the other end of the transmission path) will not exceed some predetermined level, and the analysis in this paper shows that this technique will work *for a single unlicensed transmitter*. However, as the NPRM notes, this open-loop approach does not account for interference aggregation, which can be a significant factor in this case. Moreover, setting

the increment conservatively, to guarantee that the signal-to-interference ratio does not drop below a certain threshold (*e.g.*, 50 dB), will, even for a single unlicensed transmitter, result in a very low operating range (several meters) for the unlicensed link, due to the interference from the licensed transmitter. A non-measurement based method, such as requiring unlicensed devices to use GPS and consult a data base of licensed receiver locations, may be a more reliable approach for allowing unlicensed devices to exploit unused space in frequency bands licensed to microwave fixed services. In that case, the rules should require that the unlicensed devices be beyond the radio horizon of the fixed receiver in order to operate in its band.

In sum, based on the technical analyses performed to this point, the ITemp concept does not appear to be promising except perhaps in very special cases (primarily those with large exclusion zones). The closed-loop (monitoring and feedback) form of ITemp suffers from some fairly restrictive constraints on the locations of the monitoring receivers, and there are some types of receivers that are not practical candidates, such as mobile terminals and broadcast receivers. Even for cases in which the licensed system topology might permit closed-loop ITemp (a monitor collocated with a stationary receiver at a known location, and a path loss from the unlicensed transmitters to the licensed receiver that can be calculated by the unlicensed transmitters), the implementation challenges are significant. Finally, for the CMRS case examined in detail here, even with perfect implementation, net spectrum efficiency would be reduced by addition of the ITemp unlicensed devices.

More generally, these results also show the importance of subjecting any potential ITemp application to detailed technical analysis, including consideration of the effect on the *unlicensed* devices of the *licensed* transmitters, in order to quantify the tradeoff between licensed functionality lost and unlicensed value gained.

CONTENTS

1. Introduction 7
2. Overview of ITemp Implementation Principles 10
2.1 Scope of ITemp 10
2.2 Interference Temperature Control Mechanisms 11
2.3 Power Control using Monitoring and Feedback 11
2.4 The Ineffectiveness of Self Monitoring for ITemp 13
2.5 Bandwidth and Waveform Sensitivity Factors 17
2.6 Implementation Challenges 18
2.7 Spectrum Utilization Cost/Benefit Tradeoffs 18
3. Analysis of the Aggregate Unlicensed Interference 19
3.1 Introduction 19
3.2 Assumptions and Model 20
3.3 Mean and Variance of the Aggregate Interference Power 21
3.4 Interference Mean and Standard Deviation Versus Distance from Exclusion Zone Center 26
3.5 Correlation Between Interference Levels at Different Locations 29
3.6 Implications for ITemp Monitoring 32
3.7 Cumulative Distribution Function of the Aggregate Interference 33
3.8 Single-Interferer vs. Multiple Interferer Models 35
3.9 Minimum Distance Calculations 36
3.10 Discussion 37
4. The Impact of ITemp Applied to a Mobile Services CDMA Network 39
4.1 Forward Link 39
4.2 Reverse Link 40
4.2.1 SINR and Jamming Margin 40
4.2.2 Basic Uplink Capacity Relationships 41
4.2.3 Uplink Capacity Reduction due to External Interference 44
4.2.4 Uplink Throughput Reduction 46
4.2.5 EIRP of Unlicensed Devices 47
4.2.6 Implementation Issues 50
4.2.7 Summary 51
5. Cost/Benefit Analysis for ITemp Sharing of the CDMA Uplink 52
5.1 Carrier-to-Interference Ratio as Seen by the Unlicensed Receivers 52
5.2 Unlicensed Receiver Outage Probability 54
5.3 Local Propagation and Unlicensed Rate vs. Range 55
5.4 The Throughput-Coverage Product for the Unlicensed Devices 58
5.5 Licensed Cost vs. Unlicensed Benefit 61
5.5.1 Mobile and Portable Communication Spectrum Efficiency 61
5.5.2 ITemp Cost/Benefit Analysis for CMRS 63
5.6 Conclusions 66
6. The Throughput vs. Coverage Tradeoff 67
6.1 Introduction 67
6.2 Average Throughput over a Coverage Area 68

6.3 Cellular System Analysis with Discrete Modulation Rates 70
6.4 Average Downlink Throughput per Cell 73
6.5 Outage, Minimum Rate, and Shadow Fading 76
6.6 Summary 77
7. Interference Temperature Analysis for Fixed Point-to-Point Services 79
7.1 Introductory comments on fixed point-to-point services 79
7.2 Explicit consideration in NPRM 80
7.2.1 Initial FCC Analysis Supporting Application of Dynamic Frequency Selection (DFS) and Transmit Power Control (TPC) 81
7.2.2 Specific Scenario Analysis 85
7.2.3 Interference Aggregation Effects 94
7.2.4 Dynamic Frequency Selection Thresholds 96
7.3 Two-Way Links 97
7.4 Feedback-based ITemp Management 98
7.5 Unlicensed Sharing without Measurements: Exclusion Zone Assessment 98
7.6 Summary 105
8. Annex A: CDF of the Aggregate Interference 108
9. Annex B: Interference Correction Factors 115

1. Introduction

The FCC has released a Notice of Inquiry (NOI) and Notice of Proposed Rule Making (NPRM) in ET Docket 03-327 regarding the possibility of using an "interference temperature" concept to manage interference into licensed services from unlicensed devices sharing the same spectrum.[2] The concept is to allow unlicensed devices to be introduced into selected licensed bands in a controlled way such that the aggregate interference from the unlicensed devices will not raise the noise floor seen by the incumbent licensed service by more than some threshold amount. Because thermal noise power in a receiver can be calculated as kTB (where k is Boltzmann's constant, T is the effective noise temperature in degrees Kelvin, and B is the bandwidth), the effect of the added interference can be viewed as increasing the noise temperature by some amount ΔT, which is the effective "interference temperature" (ITemp) seen by the victim receiver. The basic premise underlying the ITemp concept is that if a mechanism can be devised to limit ΔT to some predetermined threshold, then it may be technically possible for unlicensed devices to coexist with licensed systems, thereby providing a means of increasing utilization of the spectrum while limiting the impact on the licensed systems to predictable levels.

The purpose of this paper is to explore, from a purely technical perspective (without regard to economic, legal or regulatory policy issues) the implementation aspects of the ITemp concept. There are a number of dimensions to the implementation challenge, including:

1. Quantifying the impact on the incumbent licensed system; *i.e.*, determining the correct measure of the effect of adding the interference (*e.g.*, decrease in system availability, reductions in coverage, capacity, service quality). The appropriate metric will depend on the type of incumbent system and the nature of the service it is providing.
2. Determining the relationship between impact and $\Delta T/T$ (the effective increase in the noise floor). This depends on (a) the nature of the primary system, and (b) the characteristics of the interference, including waveform properties and transmit timing characteristics.[3]
3. Setting a threshold for maximum impact, which can be translated to an upper limit on $\Delta T/T$.
4. Devising a mechanism for regulating $\Delta T/T$ and ensuring that it does not exceed the limit. In some cases, it may not be possible to effectively regulate $\Delta T/T$.

---

[2] *See* Establishment of an Interference Temperature Metric to Quantify and Manage Interference and to Expand Available Unlicensed Operation in Certain Fixed, Mobile and Satellite Frequency Bands, ET Docket 03-237, Notice of Inquiry and Notice of Proposed Rulemaking, FCC 03-289 (Nov. 28, 2003) ("NOI/NPRM").

[3] Not all waveforms will affect a given receiver in the same way as noise. This point is discussed in more detail in Annex B.

The paper is organized as follows.

Section 2 provides a technical overview of ITemp implementation principles, defining the scope of the concept and outlining the two basic types of interference control mechanisms. This section also explains why self monitoring of the intended transmit channel by an unlicensed device is ineffective for regulating the interference into a victim receiver, and briefly summarizes other key topics, including the need to account for the sensitivity of receivers to specific waveforms (rather than assuming that interference affects the receiver in the same way as noise), and the usefulness of analyzing spectrum usage cost/benefit tradeoffs.

Section 3 develops a mathematical model of the aggregate interference from unlicensed devices, which often are randomly-located relative to the random receiver. The section first derives expressions for the mean and standard deviation of the aggregate interference if there is an "exclusion zone" (a guaranteed minimum separation between the victim receiver and any interfering transmitters). Monte Carlo computations are then used to show the mean and standard deviation of the interference seen by a monitoring receiver versus its distance from the center of the exclusion zone. Following that, the correlation coefficient between interference power levels at different locations is computed. Finally, the actual cumulative distribution function (CDF) is derived for the aggregate interference, and the CDF for the aggregate interference is compared to that of the interference power received from the single nearest transmitter. Section 3 shows that without an exclusion zone, the upper tail of the interference CDF (corresponding to strong interference levels) is determined by the location of the single nearest interferer, which tends to be random. The conclusion is that to regulate the total aggregate $\Delta T/T$, there must be either (1) a real-time feedback control mechanism; or (2) a guaranteed separation distance that is adequate, given the transmit power levels and spatial densities of the interfering transmitters. Otherwise, $\Delta T/T$ is essentially a random variable, which would seem to defeat the purpose of the ITemp concept. Further, ITemp monitoring appears to be practical only when the locations of the victim receivers are fixed and known.

Section 4 analyzes the impact of applying ITemp to bands used by the uplink of mobile radio services, and in particular, mobile services using code division multiple access (CDMA) technology. The section provides a detailed analysis of the relationship between $\Delta T/T$ and the capacity impact on the CDMA uplink, and discusses some of the implementation problems.

Section 5 continues with the mobile/CDMA analysis, viewing spectrum-sharing from the perspective of the unlicensed devices themselves, which will experience interference from the CDMA handsets. A mathematical model is developed that shows the tradeoff between impact on the CDMA uplink capacity (via $\Delta T/T$) and the performance of the unlicensed devices themselves (expressed as a combination of aggregate data rate and per-link coverage). This model is used to demonstrate the spectrum utilization cost/benefit relationships associated with ITemp in this particular case, and to show that

overall spectrum efficiency is lost in this case if licensed capacity is sacrificed to gain unlicensed capacity. The total value equation includes both the loss in licensed mobile capacity and the gain in unlicensed capacity.

Section 6 discusses the inherent tradeoff between coverage and capacity, which is important to understand in evaluating the cost/benefit relationships associated with deployment of ITemp with frequency reuse systems. This material supports the analysis of section 5.

Section 7 analyzes the sharing of fixed point-to-point spectrum with unlicensed services, as proposed in the NPRM. It is shown that the approach the NPRM proposes, whereby the licensed device transmits at a power level that is a fixed number of decibels above the power received from the licensed transmitter, results in extremely short operating range for the unlicensed devices, even if aggregation effects are disregarded.

Annex A derives the cumulative distribution function (CDF) of the aggregate interference from randomly-located transmitters and relates to the material in section 3. Annex B briefly discusses interference correction factors, which can be used to adjust for the fact that not all types of interfering signals will affect the victim receiver in the same way as random noise.

2. Overview of ITemp Implementation Principles

2.1 Scope of ITemp

Although a review of the table of frequency allocations suggests there is virtually no unallocated spectrum in the U. S., a frequency scan at any given location and time would show that there are large portions of the radio spectrum with no detectable activity. There are various reasons for this: some bands are receive-only (radio astronomy); some bands are deliberately unused in certain locations to prevent interference (broadcast); and some bands are used intermittently (public safety, military communications). It could be argued that such cases represent unused/underutilized spectrum capacity and, in turn, offer unrealized opportunities to increase spectrum utilization.

The "interference temperature" ("ITemp") concept is part of an FCC effort to find ways of applying advanced radio technologies to increase spectrum utilization by allowing unlicensed devices to share spectrum that is already used by licensed services. Such an effort can be divided into two broad categories:

1. *The Total Isolation Approach:* Use of licensed spectrum in areas where it is actually "unused", meaning that any receivers associated with the primary spectrum allocation are beyond the radio horizon of any newly-introduced transmitters. An example might be television broadcasting, for which the frequencies are allocated to specific geographic areas and are relatively static. Unlicensed devices might include built-in GPS receivers and store the frequency allocations and transmitter locations of licensed services using the bands of interest. The unlicensed devices would be prohibited from operating within some radius of the transmitter location, and that radius would be computed such that it would beyond the radio horizon of any licensed receivers that could be served by the broadcast transmitter. The goal would be to cause no degradation to the licensed service. In this paper, this approach is termed "total isolation."
2. *The ITemp Concept:* Sharing of licensed spectrum by unlicensed devices within the geographic operating area of the licensed service by allowing a *limited* and *controlled* degradation of the licensed service by the unlicensed transmissions. This is the ITemp concept as discussed in the NOI/NPRM and this paper.

The two key elements for ITemp are: (a) a *limit* on the allowed impact of the unlicensed devices to the licensed service (quantified in some way that is meaningful to the specific licensed service under consideration), and (b) a mechanism to *control* the transmissions of unlicensed devices to enforce that limit. It is assumed throughout this paper that these two conditions must be met for a licensed/unlicensed sharing scenario to fall within the scope of ITemp. This means that cases in where unlicensed devices are deployed that intentionally radiate within the passband of a licensed receiver, with no restrictions on device density or proximity to a victim receiver, are beyond the scope of this paper and assumed to be beyond the scope of ITemp. An example of such a case is the ultra wide band (UWB) proceeding, ET Docket 98-153.

It should be emphasized that for purposes of assessing and controlling the impact of unlicensed devices on a licensed service, what is important is the interference *at the receivers* associated with the licensed service – not the interference threshold at the unlicensed transmitter.

2.2 Interference Temperature Control Mechanisms

Within this scope of ITemp, there seem to be two basic mechanisms for controlling interference generated by unlicensed devices, which can be used individually, or in some cases combined.

- A form of *power control*, to limit the aggregate interference from unlicensed devices into a victim receiver. In general, this mechanism limits some combination of the received interference per device and the number of interfering devices. In the most literal realization, doing so requires monitoring and active feedback because of the potential for aggregate interference. The monitoring must be implemented in such a way as to "see" the radio environment from the perspective of the victim receiver; self monitoring of the channel by the prospective unlicensed transmitting device is ineffective, as explained in detail below.
- A physical or electromagnetic *exclusion zone* between the victim receiver and any unlicensed transmitter, enforced using either geographical information (*e.g.*, GPS and knowledge of the exclusion zone boundaries) or the monitoring of an available signal and calculation of path loss based on the received signal strength. This differs from "total isolation" as mentioned above, in that there will be some interference into the victim receiver in this case, but it will be controlled by the exclusion zone.[4] A satellite uplink is an example of a natural candidate for this approach, because of the guaranteed distance limit. One problem with this approach is that there is no feedback mechanism to regulate aggregate interference. Therefore, when setting the size of the exclusion zone and the power limits for the unlicensed devices, some maximum value must be assumed for the density of active unlicensed devices. If the actual device density exceeds this value, then the interference from the unlicensed devices will exceed the established limit (*i.e.*, cause harmful interference).

It should be noted that dynamic frequency selection (DFS) is not actually an ITemp control technique, but rather a means for the unlicensed device to expand its set of options. If the unlicensed device is unable to transmit at adequate power on a particular frequency, then with DFS it can attempt to use another frequency, for which the restrictions may be different.

2.3 Power Control using Monitoring and Feedback

A strict interpretation of the ITemp principle requires the power control approach. Indeed, the NOI/NPRM states:

[4] Clearly, total isolation can be viewed as an extreme case of the exclusion zone approach.

> The interference temperature model could represent a fundamental paradigm shift in the Commission's approach to spectrum management by specifying a potentially more accurate measure of interference that takes into account the cumulative effects of all undesired RF energy; *i.e.*, energy that may result in interference from both transmitters and noise sources, that is present at a receiver at any instant of time (¶ 1).

Strict adherence to this definition requires that every ITemp-compliant unlicensed transmitter know two things:

1. How much additional interference can be applied to the victim licensed receiver before the limit is reached; and
2. How much power the ITemp device itself can transmit to introduce a given level of interference into the licensed receiver.

Condition (1) requires that there be (i) a monitoring capability at the victim receiver site, and (ii) a feedback mechanism to report the monitored interference level to the unlicensed ITemp devices. Condition (2) requires that each ITemp device be able to measure the electromagnetic path loss between itself and the victim receiver(s), which in turn requires that there be a signal transmitted from the victim receiver site that can be received by the ITemp device and used to compute the path loss (by comparing the signal transmit power, which must be known, to the received signal power). This signal typically will be on a frequency that is different from the intended ITemp device transmit frequency. It may be part of the normal operation of the licensed system. Examples include a cellular radio downlink, and the transmission of a frequency-duplexed fixed point-to-point link.

It is immediately clear from condition (2) that this strict form of ITemp cannot be used with licensed receivers that are sometimes or always passive (have no co-site transmitter active), because there is no way for the ITemp device to be aware of the receiver location or measure the path loss between itself and the receiver. Examples of such receivers include commercial broadcast receivers (which have no associated transmit signal), and commercial mobile radio services (CMRS) mobile units, which often are idle (not transmitting) but locked to the system (receiving broadcast channels and awaiting incoming messages or calls).

Condition (1) states that monitoring must be done at the site of the victim receiver. Moreover, the same antenna configuration must be used for the monitoring receiver as for the victim receiver, so that the monitoring receiver sees the same interference as the victim. Monitoring of the intended transmit channel by the ITemp device at its own location is generally not useful, as discussed below. Similarly, networks of "monitoring stations" that are not collocated with victim receiver sites are not likely to be very useful in implementing ITemp either. This is because interference due to point-source transmitters (the unlicensed devices in this case) is highly location-dependent. If the transmitters are randomly located, the interference power measured at one location is

generally a poor predictor of the interference at another location. This is a important point, and the next subsection discusses it in detail.

2.4 The Ineffectiveness of Self Monitoring for ITemp

It might be assumed that ITemp monitoring could be implemented by requiring an unlicensed device to monitor its intended transmit channel. However, as explained below, such "self monitoring" is not effective. With the simplest form of self monitoring, if the observed signal level is below some threshold, the device would transmit; otherwise, transmission would be inhibited. There are existing unlicensed devices that use this approach. These include wireless local area network (WLAN) devices using the IEEE 802.11 protocols, which use carrier sense multiple access with collision avoidance (CSMA/CA), as well as the unlicensed PCS devices operating under Subpart 15D of the FCC Rules, which require "listen before transmit" (LBT). LBT and CSMA are essentially the same concept, and differ only in the detailed procedures associated with a specific application. For example, CSMA uses a random backoff if a channel is busy or if a collision occurs, to reduce the chances of initial and recurring collisions.

While LBT and CSMA are adequate solutions for contention-based channel access among peer unlicensed devices, they are imperfect because the transmitter views the state of the transmission medium from its own perspective rather than that of the receiver, or potential victim receivers. This leads to the well-known "hidden station" and "exposed station" problems as illustrated in Figure 1 and Figure 2, respectively.

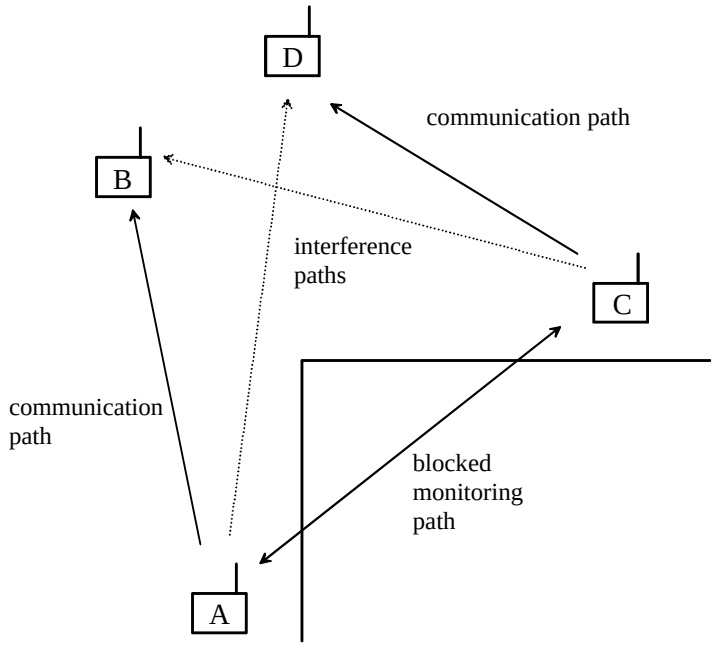

Figure 1: *Illustration of the hidden station problem.*

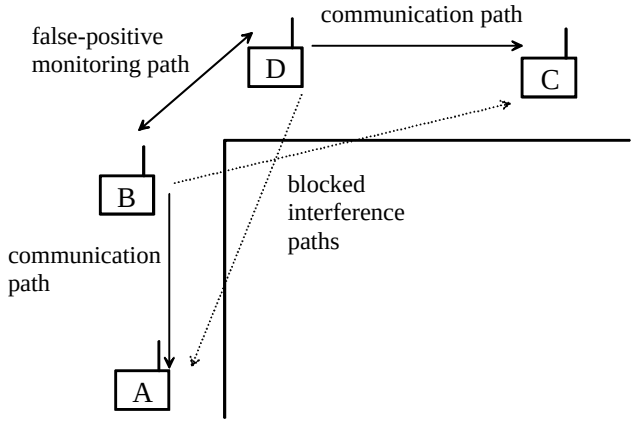

Figure 2: *Illustration of the exposed station problem*

In Figure 1, station **C** is transmitting to **D** and **A** has a packet to transmit to **B.** Station **A** monitors and finds the channel clear because the signal from **C** is blocked as shown. Once **A** transmits, it interferes with reception at **D**. With the geometry shown, the transmission from **C** also happens to interfere with reception at **B**, but there could be other examples for which that is not the case. The net result is that interference occurs because the transmitter **A**, when monitoring, cannot see the channel from the perspective of a potential victim receiver **D**. Figure 2 illustrates the opposite problem, which is that transmission can be inhibited unnecessarily. In this case, both interfering paths are blocked as shown, but when monitoring, **B** detects the transmission from **D** and inhibits its transmission. However, in the context of IEEE 802.11, this is not necessarily bad, because after receiving the packet from **D**, station **C** must send back an acknowledgment to **D**. Reception of this acknowledgment requires that **D** have a clear channel. Thus, the more troublesome problem is that of hidden stations, which leads to interference and consequently repeated transmissions and the resulting wasted resources.

The hidden station problem arises because the victim receiver is invisible to the interfering transmitter, since at the time that the transmitter monitors the channel, the victim receiver is not transmitting. The receiving station can be made "visible" using a request-to-send, clear-to-send (RTS/CTS) exchange, whereby the sending station transmits an RTS and the receiving station responds with a CTS (which makes it visible to all stations within reception range). This procedure alerts stations within range of the transmitting and/or receiving station that a message is being transmitted and the other stations need to be silent. The total length of the message is embedded in a dedicated field in the RTS and CTS packets, so that all stations that can hear either the transmitting or receiving station will know the length of the transmission interval and inhibit transmission for that length of time. For long message packets, it is worth incurring the overhead of the RTS/CTS exchange, since the longer the message packet, the more likely it is to suffer a collision, and if a collision occurs, the higher the resource cost (of retransmission). Of course, RTS/CTS requires that all stations involved use the same air interface, because they must be able to decode the RTS/CTS (not simply sense energy), and must comply with the transmit-inhibition rules.

Self monitoring of the transmit channel is a workable if imperfect vehicle for communication among the elements of an unlicensed wireless network, and collisions can be reduced with protocol enhancements such as RTS/CTS. Even so, collisions still occur, and the lack of a response from the intended receiving station provides a natural feedback mechanism. With IEEE 802.11, a station responds to a correctly received message with an acknowledgment packet (ACK). If the ACK is not received by the transmitting station, the message either was not received (*i.e.*, recipient out of range) or was corrupted by interference (detected by a checksum failure). Upon failure to receive the ACK, the transmitting station waits for a random time interval and then retransmits. The point is that in this case, the penalty for incorrect decisions due to the inaccuracy of self monitoring is a message delivery failure on the part of at least one of the stations involved in the collision, and that failure is remedied by the retransmission protocol.

Controlling the total interference into a licensed receiver as required for implementing ITemp is a different and more demanding problem. For a situation in which unlicensed devices are allowed to contribute to the total interference, the noise plus interference (N+I) experienced by a receiver will be highly dependent on the location of the receiver. This is because the interfering (unlicensed) transmitters are point sources. Figure 3 reproduces Figure 1 from the NOI/NPRM. It could be inferred from this Figure that N+I, after implementation of the ITemp concept, is unvarying with location. Figure 4 shows N+I (vertical axis) over area for an example in which there are four arbitrarily located transmitters, and as can be clearly seen, the N+I depends heavily on the location of the receiver. This means that a receiver "monitoring" a channel at one location generally cannot draw an accurate conclusion about the interference at another location. When interference is due to a number of discrete transmitters, the N+I cannot be smooth or uniform as suggested by Figure 3. Section 3 provides a detailed mathematical framework for the aggregate interference from randomly-located devices.

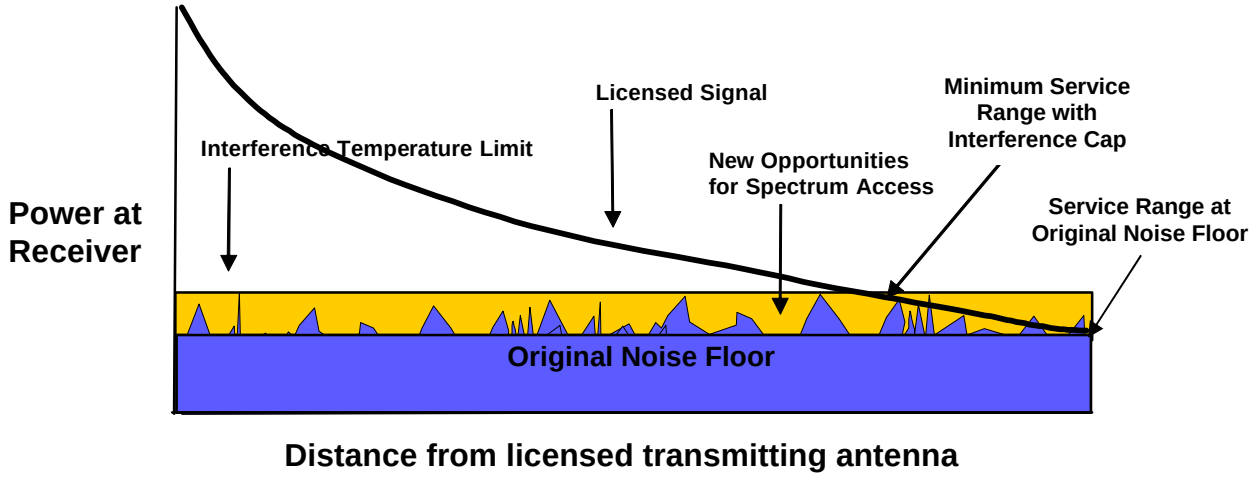

Figure 3: *Noise temperature concept illustration, reproduced from Figure 1 of the NOI/NPRM.*

, nor does it give any indication of the path loss It is clear that self monitoring of the intended transmit channel by the unlicensed device is essentially useless for implementation of the ITemp concept using the power control approach. Self monitoring does not provide a reliable measure of the level of interference seen by the victim receiver between the ITemp device and the victim receiver.

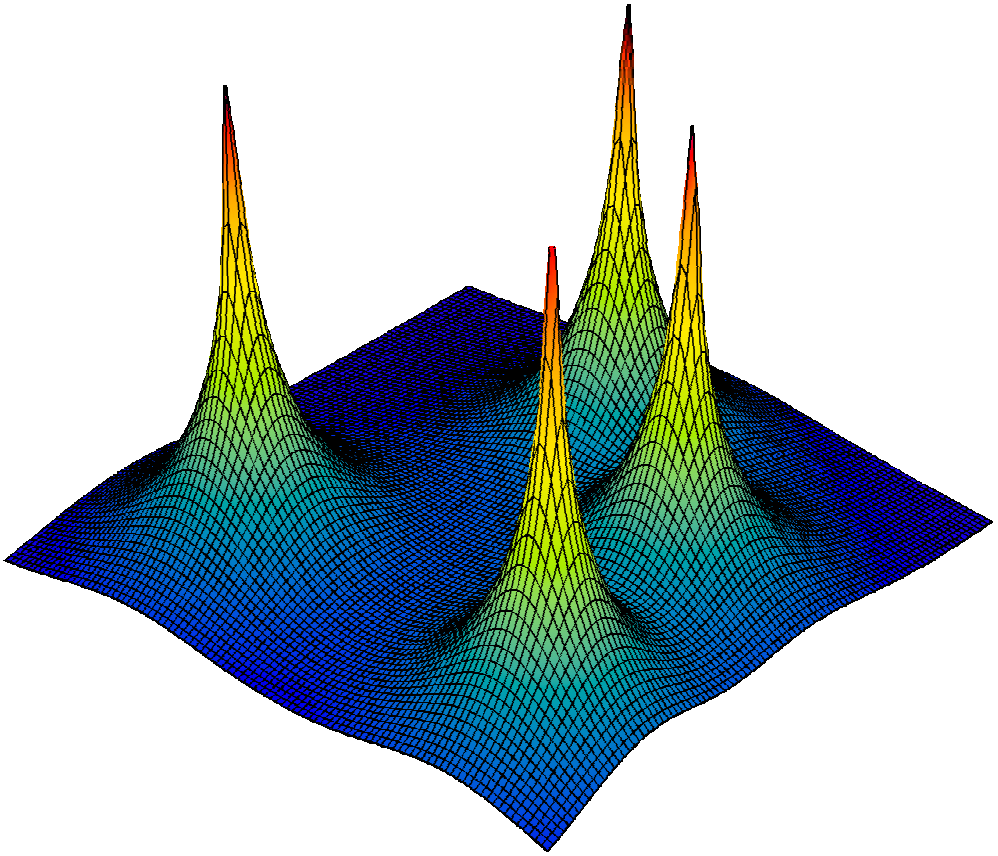

Figure 4: *Illustration of total interference power over a plane due to four point sources.*

Self monitoring might also be contemplated as a way of determining whether the unlicensed device is within a particular service area, to support the exclusion zone approach to ITemp. For example, assume that an unlicensed device is allowed to use a television broadcast channel outside of the broadcast coverage area, as defined by some signal strength coverage contour. It is tempting to assume that the unlicensed device could monitor the broadcast TV channel, with the appropriate intermediate frequency (IF) bandwidth and detection parameters, and conclude that if the signal level is below some threshold, the unlicensed device is outside the service area and can use the channel. However, this is not a reliable approach, as shown in Figure 5. In this case, a receive antenna is optimized for reception of the broadcast signal (elevated, outdoors, and with directive gain). The unlicensed device generally does not have these advantages and may, in fact, be partially blocked from the broadcast signal by local clutter as shown, while having a strong interference path to the victim receive antenna. This situation is particularly likely to occur near the edge of the coverage area, where the broadcast signal

is marginal (but still adequate for the well-engineered antenna installation). The unlicensed device may not receive a detectable signal and conclude that it is outside the service area. Upon transmission, it renders the channel unusable for the receiver connected to the rooftop antenna.

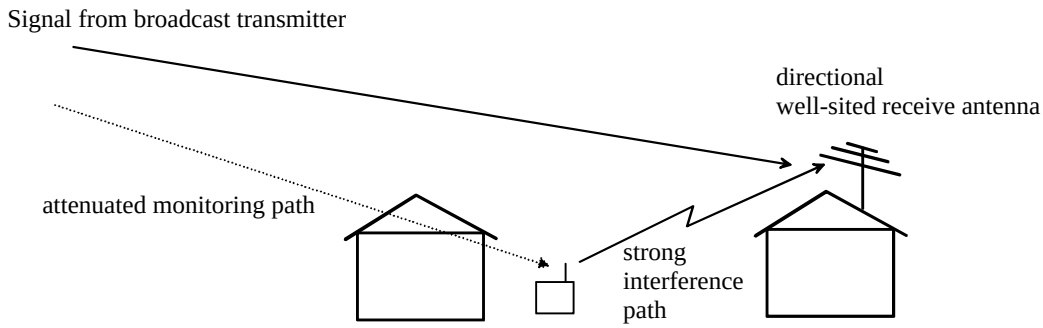

Figure 5: *Example of unreliable service area monitoring*

In sum, self monitoring of its own transmit frequency by an unlicensed device prior to transmission is not an effective means of controlling interference into the victim receiver. Such monitoring cannot be reliably used by the unlicensed device to determine whether or not it is within the usable service area of the victim receiver. It also cannot be used to determine the state of the victim receiver relative to the interference threshold, or the effect that a transmission would have on the victim receiver, because if the victim receiver is passive (not transmitting) there is no way for the unlicensed device to know the path loss between itself and the victim receiver.

2.5 Bandwidth and Waveform Sensitivity Factors

Design and analysis of potential ITemp implementations must account for differences in bandwidth between the interfering unlicensed signal and the victim licensed receiver, as well as the sensitivity of the victim receiver to the specific waveform transmitted by the unlicensed device. The first point is fairly obvious. If the bandwidth of the unlicensed device is B_{ul} and that of the victim licensed receiver is B_{vlr}, and $B_{ul} > B_{vlr}$, then the licensed receiver will tend to see a fraction B_{vlr}/B_{ul} of the power transmitted by the unlicensed device. Of course, this is somewhat simplified; the actual average interference power captured by the victim receiver will depend on the shape of the power spectral density (PSD) of the unlicensed transmission across the victim receiver passband and the frequency response of the victim receiver.

The second point, waveform sensitivity, is more subtle. Receivers are usually characterized in terms of their performance vs. signal-to-noise ratio (SNR) where the "noise" is the average power within the receiver passband of a Gaussian noise. If the interference affects the receiver differently than Gaussian noise, then either more or less average interference power may be required to produce a given degradation level,

compared to noise. The spectral, temporal, and statistical characteristics of the interference must be taken into account in evaluating specific potential ITemp implementations.

2.6 Implementation Challenges

The strict (power control) form of ITemp requires monitoring of the aggregate unlicensed interference at the licensed receiver site and feedback of that information to the unlicensed devices. The monitoring tends to be problematic, since the same channel (e.g., frequency band) is being used by both licensed and unlicensed transmitters, and the ITemp monitoring receiver will need to be able to distinguish the aggregate power received from the unlicensed devices from the signals due to licensed transmitters and background noise. If the impact of the unlicensed devices on the performance of the licensed system is to be kept at a modest level, the power received from the licensed transmitters will tend to be much higher than that from the unlicensed devices. This will make it difficult for the ITemp monitoring receiver to measure accurately the total unlicensed interference within the licensed band.

One possible solution would be to have a component of the unlicensed signal that is outside the licensed band, so that the total power from the unlicensed devices can be measured more easily. Alternatively, there could be a narrowband "pilot" tone in a separate band, and the strength of the transmitted pilot tone would be made proportional to that of the main (interfering) signal. In this way, the monitoring receiver could track the total power received from unlicensed devices as seen at the site of the victim receiver, but without needing to separate the unlicensed transmissions from cochannel licensed signals. The drawback of this "pilot tone" approach is that it would require allocation of additional bandwidth that is unused within the operating area of the unlicensed devices, which would appear to undermine the ITemp sharing concept.

2.7 Spectrum Utilization Cost/Benefit Tradeoffs

Although there are significant technical, practical, economic and regulatory challenges involved in implementing ITemp, it is worthwhile to assume a perfectly-operating implementation and examine the tradeoffs inherent in using ITemp to support sharing between licensed and unlicensed devices – that is, to quantify the balance between what is gained by introducing the unlicensed devices compared to what is lost by degrading the performance of the licensed services.

If possible, it is useful to develop a value metric that can be applied to both the licensed and unlicensed systems, to facilitate a direct cost/benefit comparison. This tends to be difficult unless the applications are similar. For example, a CMRS system provides a certain throughput over a certain coverage area. An unlicensed system that provides mobile/portable communication therefore can be compared directly to CMRS. Section 5 provides an example of a detailed cost/benefit analysis for the CMRS case. Conversely, a licensed point-to-point link provides some fixed throughput over a given path length, with a specified reliability objective. In that case, cost/benefit analysis *vis-à-vis* unlicensed portable communication is less straightforward.

3. Analysis of the Aggregate Unlicensed Interference

3.1 Introduction

The NOI suggests that one possible means of monitoring and controlling the interference temperature is to deploy a system of monitoring receivers at various locations. These "thermometers" would somehow be networked together to provide data on the current interference temperature within a frequency band. For this interference monitoring approach to be useful, the monitored levels must be highly correlated with the interference into the licensed receivers that are being protected. Therefore, to understand the requirements for effective interference monitoring, it is necessary to understand the statistics of the aggregate unlicensed interference, which is the topic of this section.

The interference due to randomly-located unlicensed devices, with no restrictions on the proximity of unlicensed transmitters to the victim receiver, will be highly location-dependent. In that case, "thermometers" at locations other than that of the victim receiver are of limited value. Figure 6 shows an example of the aggregate interference from four transmitters. The vertical coordinate represents total power in decibel units at the corresponding point on the plane. The transmitters are easily located by the signal power peaks.

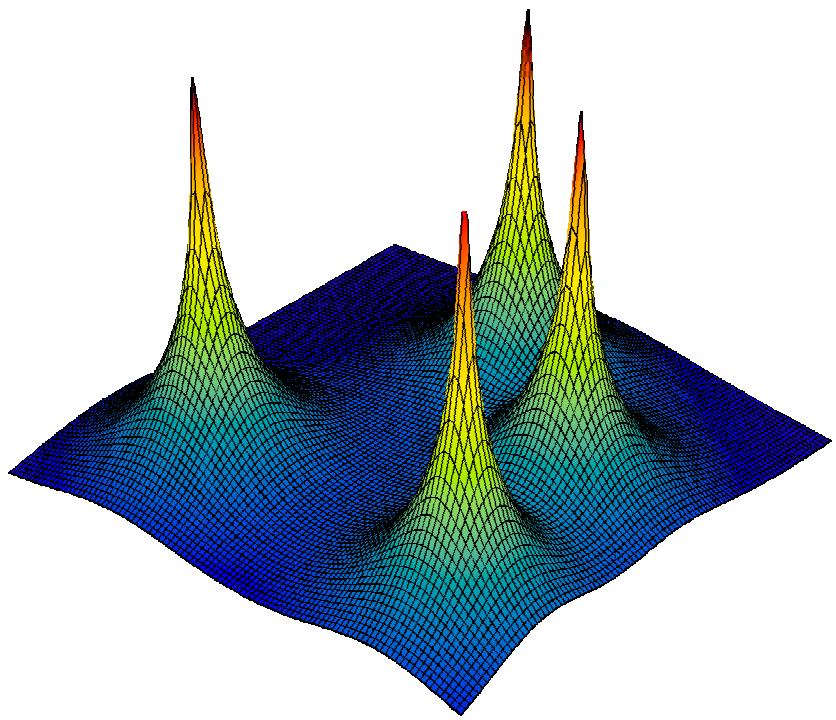

Figure 6: *Illustration of total interference power over a plane due to four point sources.*

It is clear even from this illustrative example that among nearby interfering transmitters, the received interference power can vary significantly with location. In that situation, a

monitoring receiver at one location cannot accurately gauge the interference power at another location. However, as will be shown here, if there is a large "exclusion zone" surrounding the victim receiver, and within this zone no unlicensed transmitters can operate, the interference to the victim receiver would be less variable. In addition, it would not be necessary that the monitoring receiver be at the same location as the victim receiver, only that the distance between them be small relative to the exclusion zone radius.

To understand the limitations on the usefulness of interference monitoring, it is necessary to analyze the statistics of the aggregate interference from multiple unlicensed devices, which is the subject of this section. After the assumptions and model are summarized, the mean and standard deviation of the aggregate interference are calculated. Monte Carlo computations are then used to show the mean and standard deviation of the interference seen by a monitoring receiver versus its distance from the center of the exclusion zone. Following that, the correlation coefficient between interference power levels at different locations is computed. Finally, the actual cumulative distribution function (CDF) is derived for the aggregate interference, and the CDF for the aggregate interference is compared to that of the interference power received from the single nearest transmitter. It is shown that without an exclusion zone, the upper tail of the interference CDF (corresponding to strong interference levels) is determined by the power from the single nearest interferer.

3.2 Assumptions and Model

The assumptions made here are:

- Unlicensed transmitters are uniformly randomly distributed over area (in two dimensions) with an average density of r_u.
- There may be some minimum distance $r_{\min}$ between the unlicensed transmitter and the victim receiver. The model allows the possibility that $r_{\min} = 0$.
- If *d* is the distance between the unlicensed transmitter and the narrowband receiver, the path loss is proportional to d^g where $g > 2$ (i.e., path loss exceeds free-space loss).
- The average power levels from multiple unlicensed transmitters seen by the victim receiver are additive; that is, the interference contributions of the unlicensed transmitters add non-coherently as power, rather than coherently as voltage.

What is modeled here is the *total power* received from a group of randomly-distributed unlicensed transmitters surrounding the victim receiver. That is, if there are *J* active unlicensed transmitters, all transmitting the same average power within the passband of the victim receiver, the total receive interference is

$$I = a \sum_{j=1}^{J} d_j^{-g} \tag{1}$$

where d_j is the distance between the victim receiver and the jth unlicensed transmitter.

3.3 Mean and Variance of the Aggregate Interference Power

It is assumed that in general, there is some minimum distance $r_{\min}$ between the unlicensed transmitter and the victim receiver; see Figure 7. In some cases $r_{\min} = 0$; that is, the unlicensed transmitter is permitted to be arbitrarily close to the victim receiver. In other cases, $r_{\min} \neq 0$, either because of an enforced "exclusion zone", or for physical reasons (e.g., there may be a barrier such as a fence around the licensed receiver). The outer circle has a radius of *D* and is centered on the victim receiver (*D* might be the radio horizon in some cases). Unlicensed transmitters are assumed to be randomly-distributed over the shaded area between the two circles, which has an area of $p\left(D^2 - r_{\min}^2\right)$. Figure 8 shows an example of a set of randomly-generated interfering transmitter positions with an exclusion zone.

The unlicensed transmitter is a random distance d_u from the victim receiver. For a uniform random distribution of unlicensed transmitters over the shaded area, the probability density function (PDF) of d_u is:

$$f_{d_u}(x) = \frac{2x}{D^2 - r_{\min}^2} \qquad r_{\min} \le x \le D \tag{2}$$

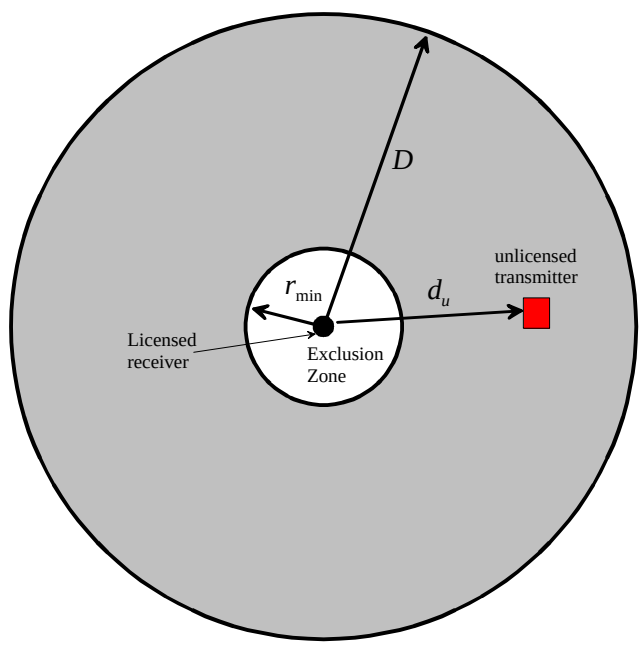

Figure 7: *Geometry for interference from unlicensed transmitters to licensed receiver*

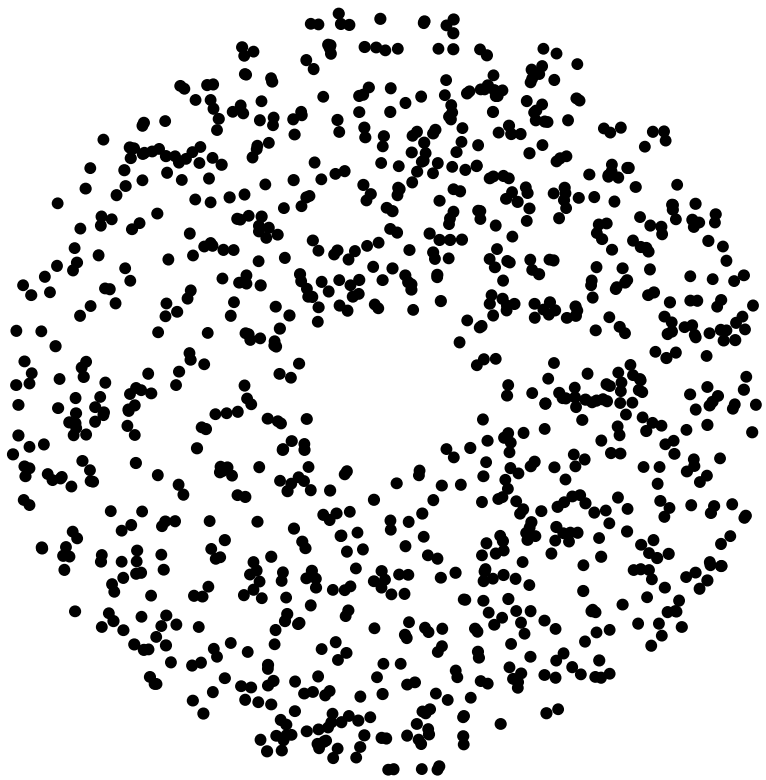

Figure 8: *Example of randomly-located interfering transmitters with an exclusion zone*

The interference from the unlicensed transmitter is $I(d_u) = I_{\max}(d_u / r_{\min})^{-g}$, where $I_{\max} = a r_{\min}^{-g}$ is the interference that would be received from an unlicensed transmitter with $d_u = r_{\min}$.

The average interference power from a single unlicensed transmitter that is randomly-located in the shaded region therefore is:

$$\bar{I}_1 = I_{\max} \int_{r_{\min}}^{D} \left(\frac{x}{r_{\min}} \right)^{-g} f_{d_u}(x)\, dx = \begin{cases} \dfrac{I_{\max}}{g/2-1} \cdot \dfrac{1}{D^2 - r_{\min}^2} \cdot \left[r_{\min}^2 - D^2 \left(\dfrac{r_{\min}}{D} \right)^g \right] & g > 2 \\ 2 I_{\max} \ln\left(\dfrac{D}{r_{\min}} \right) \dfrac{r_{\min}^2}{D^2 - r_{\min}^2} & g = 2 \end{cases} . \quad (3)$$

If the average density of active unlicensed transmitters is r_u, then the average number of unlicensed transmitters in the shaded area is

$$\bar{J} = p r_u \left(D^2 - r_{\min}^2 \right), \quad (4)$$

and the mean interference from unlicensed transmitters in the shaded area is:

$$\bar{I} = \bar{J} \cdot \bar{I}_1 = \begin{cases} \dfrac{p r_u I_{\max}}{g/2-1}\left[r_{\min}^2 - D^2 \left(\dfrac{r_{\min}}{D} \right)^g \right] & g > 2 \\ 2 p r_u I_{\max} r_{\min}^2 \ln\left(\dfrac{D}{r_{\min}} \right) & g = 2 \end{cases} . \tag{5}$$

For $D >> r_{\min}$, the right hand term of the expression in the brackets is negligible for g in the range of interest here, and the upper bound for $g > 2$ is:

$$\bar{I} < I_{\max} \frac{p\, r_{\min}^2\, r_u}{g/2-1} \qquad g > 2 \tag{6}$$

The standard deviation of the aggregate interference, denoted s_I, is also of interest, and can be found using

$$s_I^2 = \left\langle I^2 \right\rangle - \left(\bar{I} \right)^2 \tag{7}$$

The mean squared interference is

$$\left\langle I^2 \right\rangle = \left\langle \left(\sum_{j=1}^{J} I_j \right)^2 \right\rangle \tag{8}$$

where I_j is the interference received from the jth transmitter. It is assumed here that the interference power levels received from two different transmitters are uncorrelated; that is:

$$\left\langle I_j I_k \right\rangle = \left(\bar{I} \right)^2 + s_{I_1}^2 d\left[j-k \right] \tag{9}$$

where $d[n]$ is the Kronecker delta function, defined as

$$d[n] = \begin{cases} 1 & n = 0 \\ 0 & n \neq 0 \end{cases} \tag{10}$$

and $s_{I_1}^2$ is the variance of the interference power from a single transmitter.

Assuming a fixed value of *J* and taking the expectation over the $\{I_j\}$ therefore gives

$$\left\langle I^2 \middle| J \right\rangle = J s_{I_1}^2 + J^2 \left(\bar{I}_1\right)^2 \tag{11}$$

Next, expectation must be taken over *J*, which is assumed to be a Poisson-distributed random variable, in which case $s_J^2 = \bar{J}$, so $\overline{J^2} = \left(\bar{J}\right)^2 + \bar{J}$ and

$$\overline{I^2} = \bar{J} s_{I_1}^2 + \left[\left(\bar{J}\right)^2 + \bar{J}\right]\left(\bar{I}_1\right)^2 = \bar{J}\,\overline{I_1^2} + \left(\bar{J}\,\bar{I}_1\right)^2 = \bar{J}\,\overline{I_1^2} + \left(\bar{I}\right)^2 \tag{12}$$

Thus,

$$s_I^2 = \bar{J}\,\overline{I_1^2}\,. \tag{13}$$

The mean-square interference from a single unlicensed transmitter within the shaded area is

$$\begin{aligned}\overline{I_1^2} &= I_{\max}^2 \int_{r_{\min}}^{D} \left(\frac{x}{r_{\min}}\right)^{-2g} f_{d_{uwb}}(x)\,dx = \frac{I_{\max}^2}{g-1} \cdot \frac{1}{D^2 - r_{\min}^2} \cdot \left[r_{\min}^2 - D^2\left(\frac{r_{\min}}{D}\right)^{2g}\right] \\ &< \frac{I_{\max}^2}{g-1} \cdot \frac{r_{\min}^2}{D^2 - r_{\min}^2}\end{aligned} \tag{14}$$

giving the variance of the total interference as:

$$s_I^2 = p r_u \frac{I_{\max}^2}{g-1}\left[r_{\min}^2 - D^2\left(\frac{r_{\min}}{D}\right)^{2g}\right] \tag{15}$$

with $I_{\max} = a r_{\min}^{-g}$, this becomes

$$s_I^2 = p r_u \frac{a^2}{g-1}\left(r_{\min}^{2(1-g)} - D^{2(1-g)}\right) \tag{16}$$

For $D >> r_{\min}$ and $g \geq 2$, the variance is tightly upper-bounded by:

$$s_I^2 < \frac{p r_u r_{\min}^2 I_{\max}^2}{g-1}, \tag{17}$$

and for $D >> r_{\min}$, the ratio of the standard deviation to the mean is approximately

$$\begin{aligned} \frac{s_I}{\bar{I}} &\cong \frac{g/2-1}{\sqrt{g-1}} \frac{1}{\sqrt{p r_{\min}^2 r_u}} && g > 2 \\ \frac{s_I}{\bar{I}} &\cong \frac{1}{2\sqrt{p r_u r_{\min}^2} \ln(D/r_{\min})} && g = 2 \end{aligned} . \tag{18}$$

The radius $r_{\min}$ describes a circular "exclusion zone" surrounding the victim receiver. The average number of interfering transmitters that would be within the exclusion zone if they were permitted is

$$N_{xz} = p r_{\min}^2 r_u, \tag{19}$$

which essentially normalizes the size of the exclusion zone to the transmitter density.

Hence,

$$\begin{aligned} \frac{s_I}{\bar{I}} &\cong \frac{g/2-1}{\sqrt{g-1}} \frac{1}{\sqrt{N_{xz}}} && g > 2 \\ \frac{s_I}{\bar{I}} &\cong \frac{1}{2\sqrt{N_{xz}} \ln(D/r_{\min})} && g = 2 \end{aligned} \tag{20}$$

and for $D >> r_{\min}$,

$$\begin{aligned} \bar{I} &\cong I_{\max} \frac{N_{xz}}{g/2-1} && g > 2 \\ \bar{I} &\cong 2 I_{\max} N_{xz} \ln\left(\frac{D}{r_{\min}}\right) && g = 2 \end{aligned} \tag{21}$$

with

$$I_{\max} = a r_{\min}^{-g} = a\left(\frac{p r_u}{N_{xz}}\right)^{g/2}. \tag{22}$$

Thus,

$$\begin{aligned} \bar{I} &= \frac{a}{g/2-1}(p r_u)^{g/2} N_{xz}^{1-g/2} && g > 2 \\ \bar{I} &= 2 p r_u a \ln\left(\frac{D}{r_{\min}}\right) && g = 2 \end{aligned} \tag{23}$$

$$s_I = \frac{a}{\sqrt{g-1}}(p r_u)^{g/2} N_{xz}^{(1-g)/2} \tag{24}$$

It can be seen that as $N_{xz} \to 0$, $s_I / \bar{I} \to \infty$, and also $\bar{I} \to \infty$. That is, for vanishingly small exclusion zones, the average interference becomes very large and the standard deviation becomes larger, which quantitatively supports the observation made in the introduction to this section – namely, that the interference is highly variable as location changes.

3.4 Interference Mean and Standard Deviation Versus Distance from Exclusion Zone Center

The values of $\bar{I}$ and s_I derived above pertain to a single receiver located at the center of the exclusion zone, be it the victim receiver or a monitoring receiver. It is interesting to determine the effect of an offset from the center of the exclusion zone on $\bar{I}$ and s_I, which might represent a monitoring receiver somewhere in the exclusion zone but not at the center. Figure 9 shows the mean and standard deviation of the monitored interference (determined from Monte Carlo simulations) vs. $d_{mon}/r_{\min}$ where d_{mon} is the distance of the receiver from the exclusion zone center, for $N_{xz} = 50$. Figure 10 and Figure 11 show similar curves for $N_{xz} = 10$ and $N_{xz} = 1$, respectively. Levels are shown in dB relative to $\bar{I}$, computed using (21).

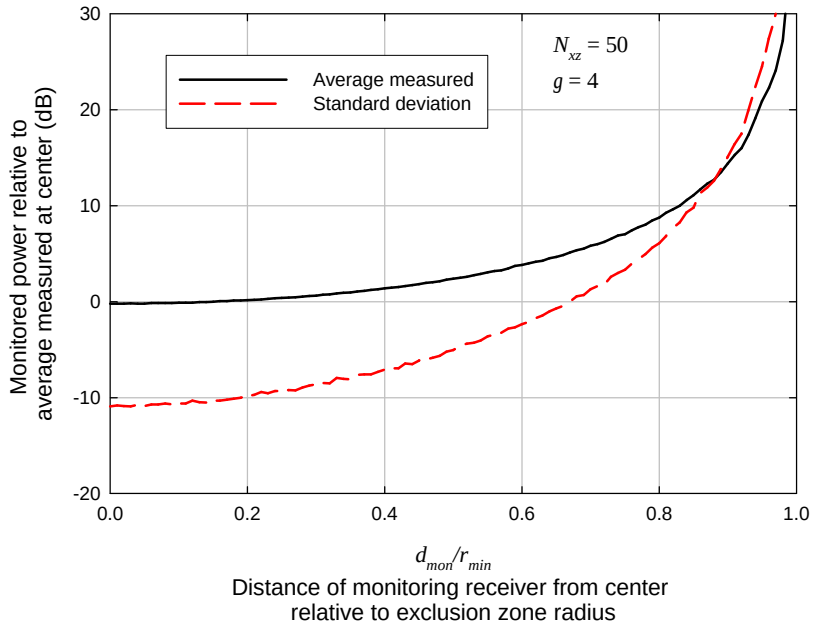

Figure 9: *Mean and standard deviation of monitored aggregate interference vs. offset of monitor from exclusion zone center for* $N_{xz} = 50$*; Monte Carlo, 1000 samples/point.*

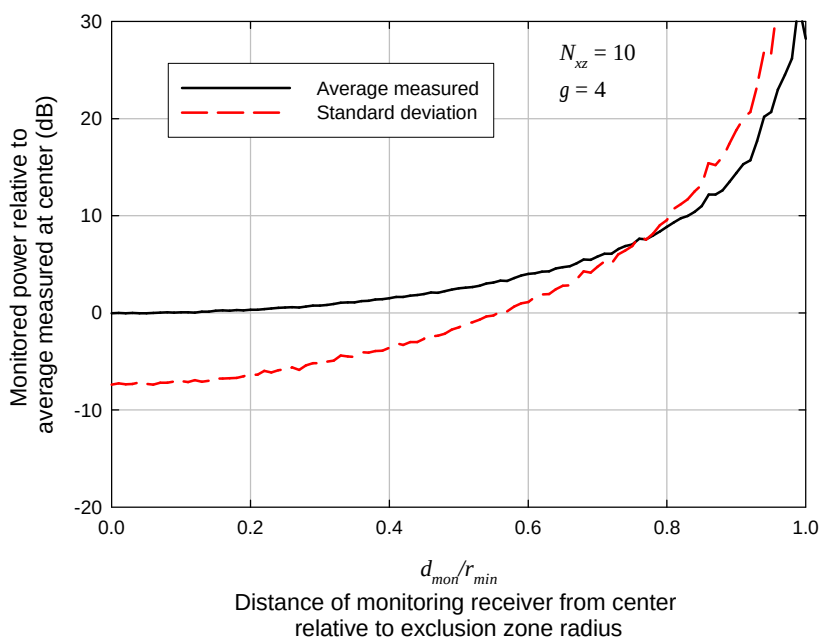

Figure 10: *Mean and standard deviation of monitored aggregate interference vs. offset of monitor from exclusion zone center for* $N_{xz} = 10$*; Monte Carlo, 1000 samples/point.*

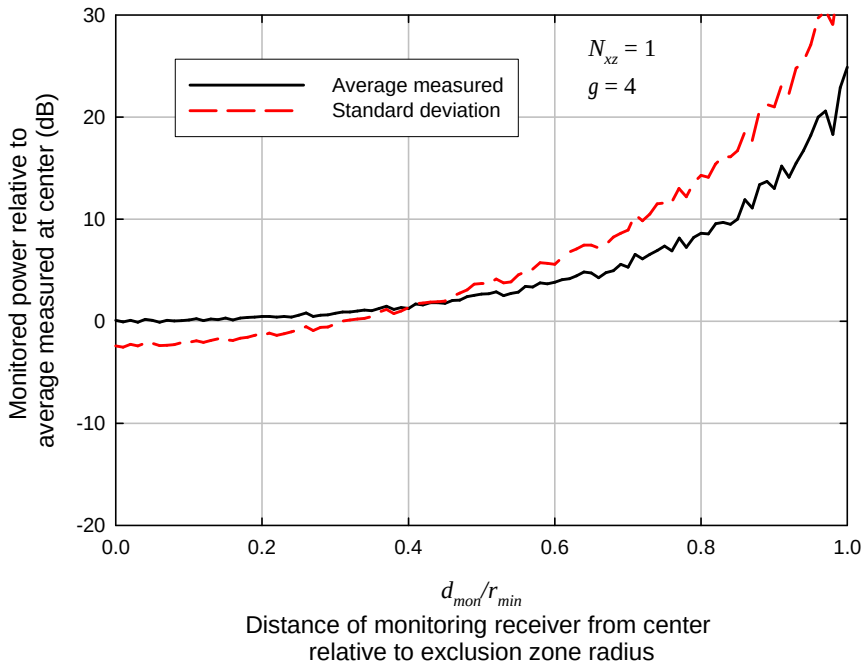

Figure 11: *Mean and standard deviation of monitored aggregate interference vs. offset of monitor from exclusion zone center for* $N_{xz} = 1$*; Monte Carlo, 10,000 samples/point.*

The x-axis increment was 0.01 (100 equally-spaced values of $d_{mon}/r_{\min}$). For each of these points, 1000 samples were used in the cases of Figure 9 and Figure 10 to calculate the average and standard deviation. For Figure 11, 10,000 samples per point were used, and despite that, the curves in Figure 11 are noticeably less smooth than the others. The explanation is useful, because it bears on the problem of using monitoring to regulate interference temperature.

To calculate the average interference shown in the curve, the Monte Carlo program simply generates *M* independent samples of the aggregate interference (for a given value of $d_{mon}/r_{\min}$) and computes the estimate as:

$$\hat{\bar{I}} = \frac{1}{M}\sum_{m=1}^{M} I_m \tag{25}$$

As is well known, the standard deviation of the estimate is

$$s_{\hat{\bar{I}}} = \frac{1}{\sqrt{M}} s_{I_m} \ . \tag{26}$$

Since $s_{I_m} \propto N_{xz}^{(1-g)/2}$,[5] it is necessary that $N_{xz}^{(1-g)/2} \big/ \sqrt{M}$ remain constant as N_{xz} changes if $s_{\hat{I}}$ is to be held constant. For the cases shown, $g = 4$, which means that $N_{xz}^3 M$ must be held constant, or $M \propto 1 \big/ N_{xz}^3$. Thus, to maintain the same "smoothness" as the $N_{xz} = 10$ case with 1000 samples per point, one million samples per point would be needed for the $N_{xz} = 1$ case.

3.5 Correlation Between Interference Levels at Different Locations

The correlation among interference samples from monitoring stations in different locations will (not surprisingly) depend on how far apart they are relative to the exclusion zone radius. The correlation coefficient for interference power levels received at two different locations is calculated by first computing the covariance:

$$C = \langle I_1 I_2 \rangle - \langle I_1 \rangle \langle I_2 \rangle \tag{27}$$

The correlation coefficient is then:

$$r = \frac{C}{s_{I_1} s_{I_2}} \tag{28}$$

The correlation coefficient for the power levels received from two monitoring stations, at $\pm d_{mon} \big/ r_{\min}$ was computed via Monte Carlo and is shown in Figure 12 for $N_{xz} = 50$, Figure 13 for $N_{xz} = 10$, and Figure 14 for $N_{xz} = 1$ (10,000 samples per point in each case). As can be seen, there does not seem to be significant difference among the cases. Stations close to each other and near the center of the exclusion zone are highly correlated, while stations separated by about $1.2 r_{\min}$ or more are essentially uncorrelated.

Figure 15 and Figure 16 show, for $N_{xz} = 50$ and $N_{xz} = 1$, respectively, the correlation coefficient for interference at the exclusion zone center and interference at a point d_{mon} from the center. Note that the correlation degrees from 1 to 0 in a nearly linear fashion as d_{mon} goes from 0 to $r_{\min}$.

[5] This is true for a receiver at the center of the exclusion zone, and apparently roughly true for other locations, judging from the Monte Carlo results.

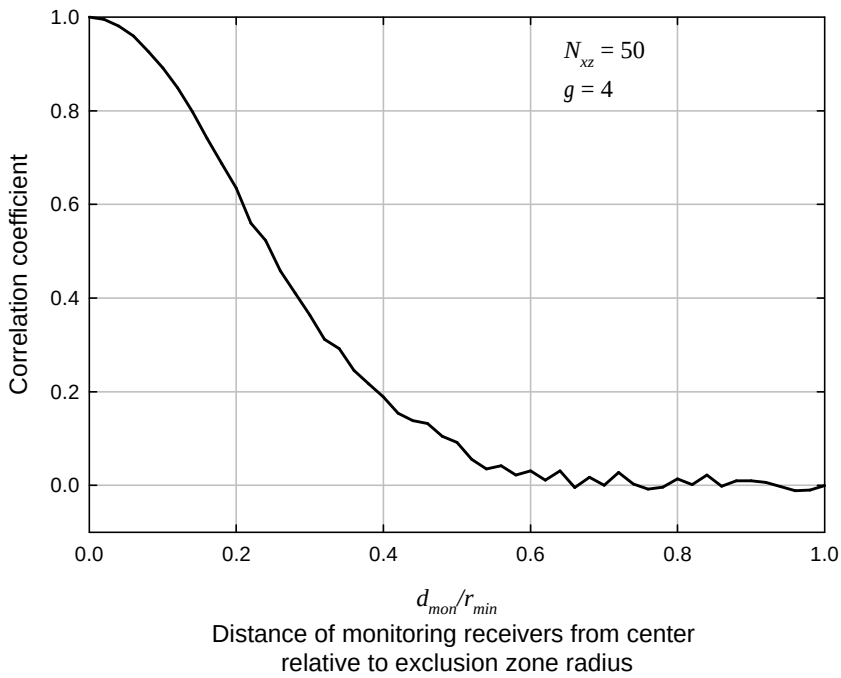

Figure 12: *Correlation coefficient from Monte Carlo computations with* $N_{xz} = 50$.

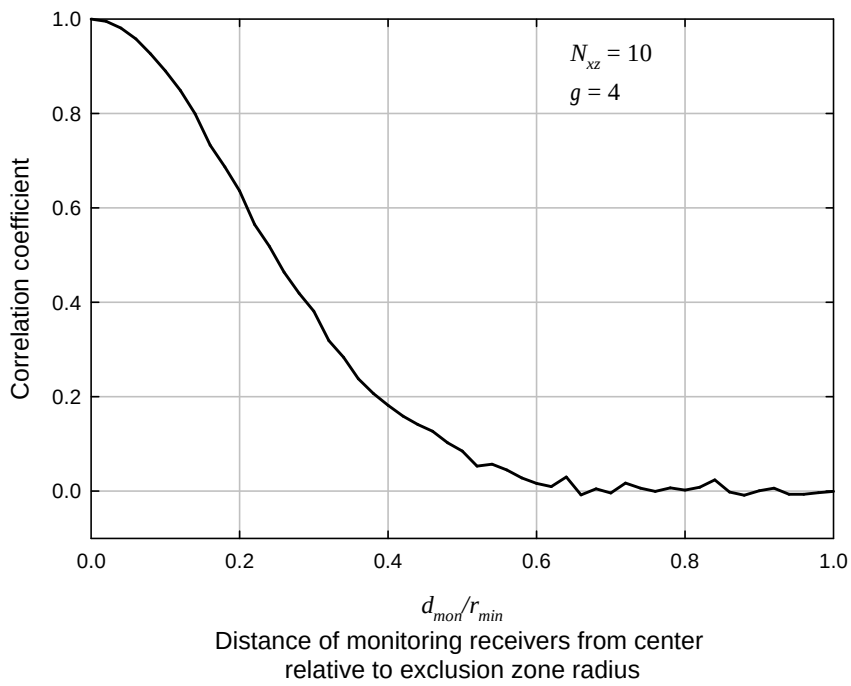

Figure 13: *Correlation coefficient from Monte Carlo computations with* $N_{xz} = 10$.

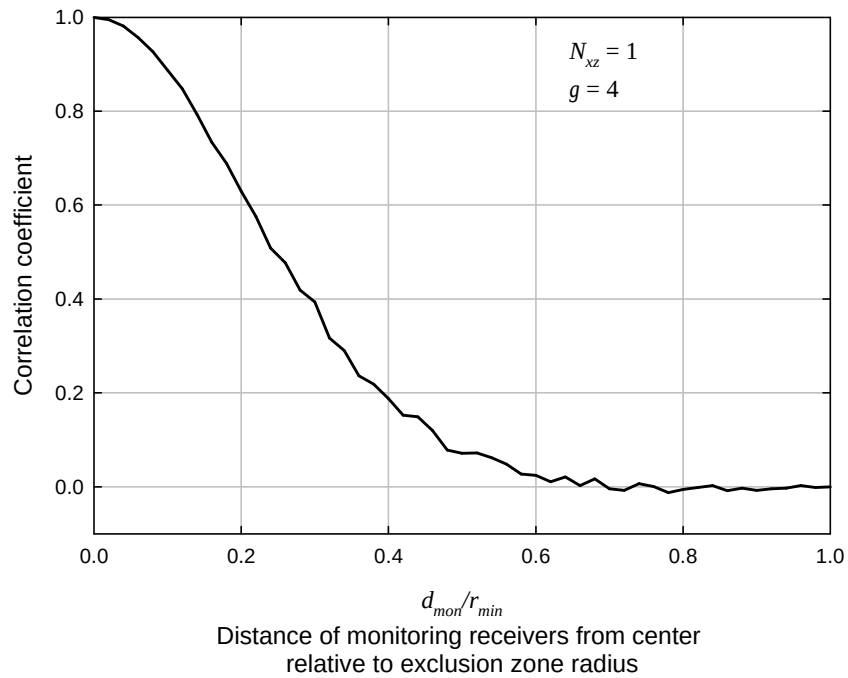

Figure 14: *Correlation coefficient from Monte Carlo computations with* $N_{xz} = 1$.

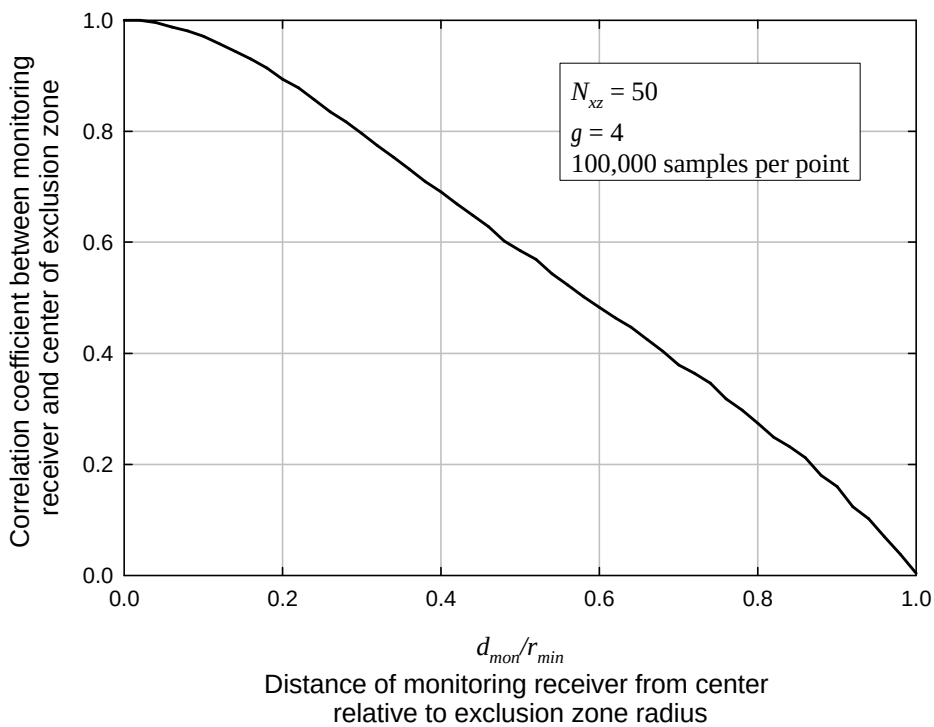

Figure 15: *Correlation between interference at exclusion zone center and a point* d_{mon} *from the center, for* $N_{xz} = 50$ *(100,000 samples per point)* .

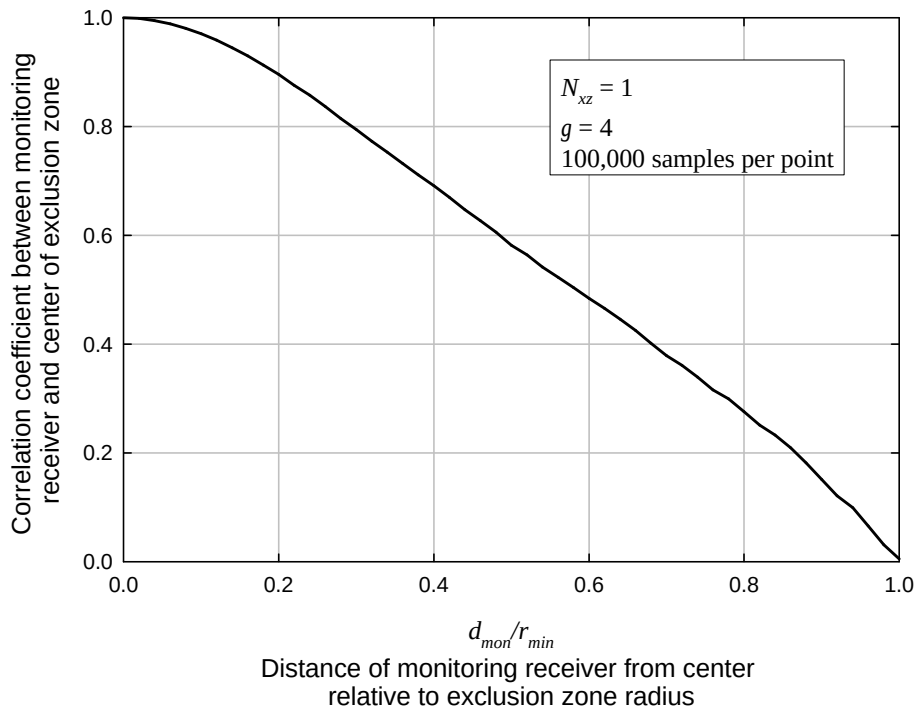

Figure 16: *Correlation between interference at exclusion zone center and a point* d_{mon} *from the center, for* $N_{xz} = 1$ *(100,000 samples per point).*

3.6 Implications for ITemp Monitoring

It is worthwhile at this point to summarize the key points from the above analyses:

- At the center of the exclusion zone, the standard deviation of the interference relative to the mean varies inversely with the exclusion zone radius.
- The mean and standard deviation of the aggregate interference is determined by the size of the exclusion zone relative to the average density of the interfering devices. That is, the exclusion zone size is best expressed by $N_{xz} = p r_u r_{\min}^2$.
- For small exclusion zones, interference is highly variable.
- As a receiver moves away from the center of the exclusion zone toward its edge, the mean and standard deviation of the interference increase. However, within $\pm 0.2 r_{\min}$ of the center, the mean and standard deviation are very near their values at the center.
- The correlation coefficient between the aggregate interference at the exclusion zone center and that at some point d_{mon} from the center decreases roughly linearly from 1 to 0 as d_{mon} goes from 0 to $r_{\min}$. However, near the center the decrease is less than linear; the correlation coefficient is about 0.9 for $d_{mon}/r_{\min} = 0.2$.
- The temporal dimension was not included in the analysis (*i.e.*, autocorrelation or autocovariance of the interference at a particular point), and would depend on the timing characteristics of the interfering transmitters. However, samples spaced far

enough in time will be uncorrelated, assuming that each licensed transmitter has a limited overall duty cycle.

- Although omnidirectional antennas were assumed in the above, the results apply equally to directional antennas, with the constraint that the monitoring receiver have the same antenna pattern and orientation as the victim receiver.

These points seem to suggest the following conclusions for ITemp monitoring:

- For real time monitoring (a station that tracks the ITemp continuously and triggers some feedback mechanism when the ITemp exceeds some threshold), the victim receivers and the monitoring station should all be near the exclusion zone center and the monitoring station should be within about $\pm 0.2 r_{\min}$ of the victim receivers it is protecting, where $r_{\min}$ is the exclusion zone radius.
- Multiple monitors within a single exclusion zone would be of little value, since they must be fairly near the victim receiver to effectively protect it, and therefore near each other. As a result, their measured power levels would be highly correlated and would not be useful for averaging.
- Non-real time monitoring with temporal averaging (a station collects a time-average of the interference over some time window, and uses the result to exert control over the unlicensed devices) could be used in cases where the standard deviation of the interference is low relative to the mean (*e.g.*, a very large exclusion zone). Some margin in the threshold for this time-averaged interference would have to be allowed for the standard deviation. Allowance in the averaging interval would need to be made for any longer-timescale trends (variation with time of day, day of the week, *etc.*).

Two points follow from the foregoing: (a) ITemp monitoring seems to be practical only when the victim receiver is fixed and known, and (b) the larger the exclusion zone surrounding the victim receiver, the more flexibility in locating the monitoring receiver.

3.7 Cumulative Distribution Function of the Aggregate Interference

If the area of the exclusion zone surrounding the victim receiver is small relative to the density of unlicensed transmitters; that is, $r_u p r_{\min}^2 << 1$, then $s_I / \bar{I} >> 1$, and the *average* unlicensed interference therefore is not a reliable measure of the interference impact. A much more useful statistic is the probability that the interference exceeds some specified level; i.e., the cumulative distribution function (CDF) of the interference. The CDF is fairly straightforward to derive if the total interference is approximated as the interference from the unlicensed transmitter nearest the victim receiver. Accounting for the aggregate power from all unlicensed transmitters is more complicated. However, as will be demonstrated, the single-source CDF is adequate in many cases.

Consistent with the assumption of a uniform planar distribution of unlicensed transmitters, the transmitter locations are modeled using a Poisson point process. The average number of unlicensed transmitters within some region of total area *A* is

$$\overline{K_A} = r_u A \tag{29}$$

The probability that the region does not contain a unlicensed transmitter therefore is:

$$P_0 = e^{-\overline{K_A}} \tag{30}$$

The area of the ring bounded on the outside by a circle of radius d and on the inside by a circle of radius $r_{\min}$, both centered on the victim receiver, is $A = p\left(d^2 - r_{\min}^2\right)$. Therefore, the probability that there are no unlicensed transmitters within a distance *d* of the victim receiver is

$$P_0(d) = e^{-pr_u\left(d^2 - r_{\min}^2\right)} = \frac{e^{-pr_u d^2}}{e^{-pr_u r_{\min}^2}}, \qquad d \geq r_{\min} \tag{31}$$

Since the interference from a unlicensed transmitter at distance *d* is $I(d) = I_{\max}\left(d/r_{\min}\right)^{-g}$, $P_0(d) = \Pr\{I_u < I(d)\}$, giving the desired CDF:

$$\Pr(I_u < I) = \exp\left\{-pr_u r_{\min}^2\left(\left(I/I_{\max}\right)^{-2/g} - 1\right)\right\} \quad I \leq I_{\max} \tag{32}$$

It is useful to normalize the interference using

$$Z \equiv \left(pr_u r_{\min}^2\right)^{-g/2} \frac{I_u}{I_{\max}}. \tag{33}$$

Clearly, $Z_{\max} = \left(pr_u r_{\min}^2\right)^{-g/2}$, and the CDF in (32) can be written as

$$\Pr(I_u < I) = \Pr(Z < z) = e^{Z_{\max}^{-2/g}} e^{-z^{-2/g}} \qquad z \leq Z_{\max}. \tag{34}$$

Also,

$$\Pr(Z < z) = \Pr\left(\frac{I_u}{I_{\max}} < z\left(pr_u r_{\min}^2\right)^{g/2}\right) \tag{35}$$

The quantity $pr_u r_{\min}^2$ represents the average number of unlicensed transmitters that would be within an area $pr_{\min}^2$ with a uniformly-distributed field of unlicensed transmitters of density r_u unlicensed transmitters/km^2.

3.8 Single-Interferer vs. Multiple Interferer Models

The CDF given in (32) is based on the interference from only the nearest unlicensed transmitter, rather than the combined interference from all co-channel unlicensed transmitters. However, at the upper tail of the CDF (high probability values on the ordinate), this "single-interferer" model gives essentially the same results as a model which accounts for the combine interference from multiple sources.

Clearly, the CDF of Z is tightly upper-bounded by

$$\Pr(Z < z) < F_Z(z) = e^{-z^{-2/g}} \tag{36}$$

where $F_Z(z)$ would be the CDF if unlicensed transmitters were not restricted in the model to a distance greater than $r_{\min}$ from the victim receiver.

If the total power from all unlicensed transmitters is taken into account, as shown in Annex A, $F_Z(z)$ becomes:

$$F_Z(z) = 1 - \frac{1}{p}\sum_{k=1}^{\infty}\frac{\Gamma(kn)}{k!}\left[\frac{\Gamma(1-n)}{z^n}\right]^k \sin kp(1-n), \quad z > 0 \tag{37}$$

where $\Gamma(\cdot)$ is the Gamma function. Note that $g > 2$ is a necessary condition for convergence. For the special case of $g = 4$, the distribution reduces to the closed form:

$$\Pr\{Z < z\} = \operatorname{erfc}\left(\frac{\sqrt{p}}{2\sqrt{z}}\right), \quad z > 0 \quad \text{for } g = 4. \tag{38}$$

where erfc(·) is the complementary error function, defined as:

$$\operatorname{erfc}(x) \equiv \frac{2}{\sqrt{p}}\int_x^{\infty} e^{-x^2}\,dx \tag{39}$$

Figure 17 shows $F_Z(z)$ for this case, along with the single-interferer model used in the unlicensed transmitter interference calculations. As can be seen, for probability levels

greater than 90%, there is no significant difference in the results. The reason is that the upper tail of the CDF corresponds to strong interference, which is dominated by a single strong (nearby) source. At lower levels on the CDF, the combined effect of multiple sources becomes more significant, and the curves diverge.

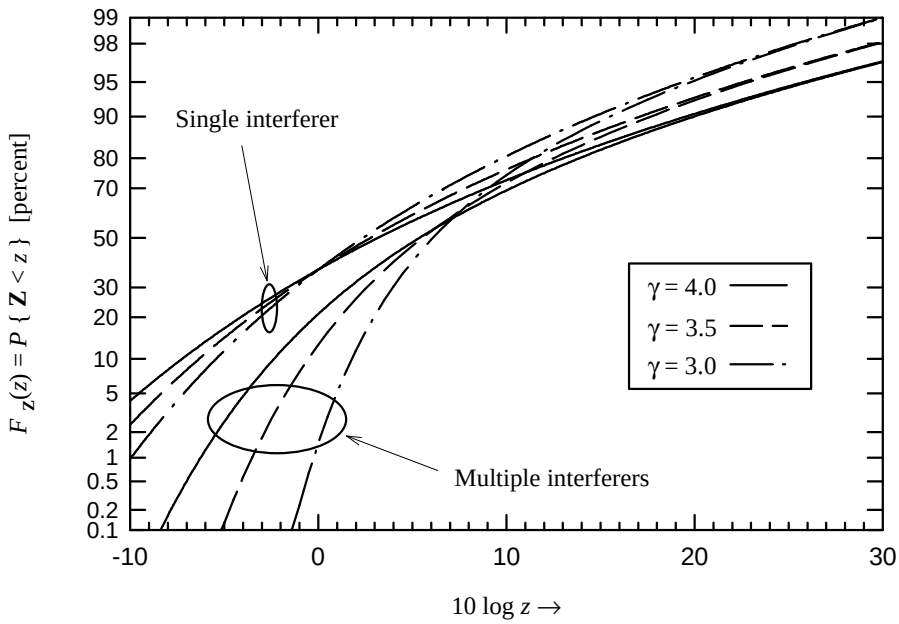

Figure 17: *CDF of Z for single-interferer and multiple-interferer models*

3.9 Minimum Distance Calculations

With the single-interferer model, the distance to the nearest interferer (unlicensed device) is of interest. The probability that this distance, denoted d_0, exceeds some distance r is:

$$P_0 = \Pr\{d_0 > r\} = e^{-pr^2 r_u} \tag{40}$$

giving

$$r\sqrt{r_u} = \sqrt{\frac{-\ln P_0}{p}}\,. \tag{41}$$

From this, tables such as the following can be developed:

P_0	$r\sqrt{r_u}$	r_u (unlicensed devices / km^2)			
		100	1000	10,000	100,000
0.99	0.056	r =5.6 m	r = 1.8 m	r = 0.56 m	r = 0.18 m
0.95	0.128	12.8 m	4.1 m	1.28 m	0.40 m
0.90	0.183	18.3 m	5.8 m	1.83 m	0.58 m
0.50	0.47	47 m	14.9 m	4.7 m	1.49 m

The meaning of these numbers is best illustrated with an example. Suppose that there are on average 1000 unlicensed devices per square kilometer simultaneously active (*roughly* one active device per quarter-acre). There is a 90% probability that the nearest one to the victim receiver will be more than 5.8 m away, a 95% probability that it will be more than 4.1 m away, *etc.*, assuming that the stated assumptions apply (uniform distribution of unlicensed devices, no inherent constraint on minimum separation).

3.10 Discussion

This section has provided an introduction to the analysis of aggregate unlicensed interference. The model developed here represents at statistical "snapshot" of the total average interference seen at the victim receiver at a given point in time. The model does not provide information about how the average interference varies with time (*e.g.*, as unlicensed transmitters turn on and off).

Duty cycle effects, as well as propagation shadowing and multipath effects, are easily taken into account with the model. The density parameter r_u represents the average density of unlicensed devices that are active at a particular time. If the total density is r_{tot} and the average duty cycle is h, then $r_u = h r_{tot}$. The effects of multipath and shadow fading can be added by introducing the random variables v_{Sj} (typically lognormal) and v_{Mj} (exponentially-distributed for Rayleigh fading), giving:

$$I = a \sum_{j=1}^{J} v_{Sj} v_{Mj} d_j^{-g} \tag{42}$$

With the inclusion of multipath and shadowing effects, there does not appear to be a closed form solution for the CDF of the aggregate interference, and the Monte Carlo

technique must be used (although the mean and variance of the interference with shadowing and multipath can be expressed in closed form).

The model given here suggests that in some cases, the CDF of the aggregate interference can be reasonably approximated by the CDF of the interference from the nearest transmitter. This approximation might apply when the exclusion distance $r_{\min}$ is small relative to the average density of active unlicensed transmitters; i.e., $pr_{\min}^2 r_u << 1$. In that case, on the upper tail of the interference CDF (high interference levels), there will tend to be a single dominant interferer (one unlicensed transmitter much closer to the receiver than any others). However, as $pr_{\min}^2 r_u$ becomes larger, there will tend to be multiple unlicensed transmitters approximately the same distance from the victim receiver, and aggregation effects become more significant.