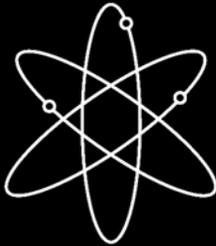


NUREG/CR-6939
ORNL/TM-2006/86

Coexistence Assessment of Industrial Wireless Protocols in the Nuclear Facility Environment



Oak Ridge National Laboratory



**U.S. Nuclear Regulatory Commission
Office of Nuclear Regulatory Research
Washington, DC 20555-0001**

NUREG/CR-6939
ORNL/TM-2006/86

Coexistence Assessment of Industrial Wireless Protocols in the Nuclear Facility Environment

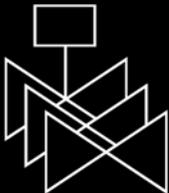
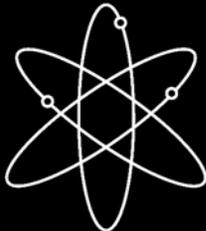
Manuscript Completed: May 2007
Date Published: July 2007

Prepared by
M. Howlader, C.J. Kiger and P.D. Ewing

Oak Ridge National Laboratory
Managed by UT-Battelle, LLC
P.O. Box 2008
Oak Ridge, TN 37831-6283

T. V. Govan, NRC Project Manager

Prepared for
Division of Fuel, Engineering and Radiological Research
Office of Nuclear Regulatory Research
U.S. Nuclear Regulatory Commission
Washington, DC 20555-0001
NRC Job Code Y6475



Coexistence Assessment of Industrial Wireless Protocols in the Nuclear Facility Environment

Manuscript Completed: May 2007
Date Published: July 2007

Prepared by
M. Howlader
C. J. Kiger
P. D. Ewing

Oak Ridge National Laboratory
Oak Ridge, TN 37831-6283
Managed by UT-Battelle, LLC

T. V. Govan, NRC Project Manager

Prepared for
Division of Fuel, Engineering, and Radiological Research
Office of Nuclear Regulatory Research
U.S. Nuclear Regulatory Commission
Washington, DC 20555-0001
NRC Job Code Y6475

ABSTRACT

This report details an interference study of the three most prominent wireless devices in use today, using computer models and simulations. The goal is to determine whether Bluetooth, Zigbee, and Wireless Fidelity (WiFi) wireless devices can coexist in an industrial environment. All three wireless devices operate in the 2.4-GHz industrial, scientific, and medical (ISM) frequency band. Simulations are conducted because of the amount of time that it would take to physically collect measurements for a plausible coexistence study. Numerous possible combinations of transmitters, receivers, and interferers are simulated. Both general channel models and site-specific channel models, incorporating the physical layout of an industrial building, are created and computed to simulate the effects of interference for certain combinations of different wireless devices. The considered channel models are basic additive white Gaussian noise (AWGN), general Rayleigh fading, and site-specific Ricean and Rayleigh fading. The results of the simulations demonstrate the performance of the three wireless devices in a practical wireless environment and their influence on other wireless devices.

Related work performed to date under JCN Y6475 includes a companion study that identified and assessed wireless technologies, both current and emerging, that have the potential for deployment in nuclear facilities. The companion study explored the technology differentiators that need to be considered before deploying a wireless system. In addition, deployment issues were investigated and current wireless deployments in nuclear facilities were examined. Results of that study are reported in NUREG/CR-6882, *Assessment of Wireless Technologies and Their Application at Nuclear Facilities* (May 2006). Security issues are also briefly explored in NUREG/CR-6882, but it was not the intent of that study to comprehensively address wireless security issues. The need for a comprehensive treatment of wireless security has been recognized and additional work will be performed in this area. The results of the planned wireless security study are anticipated to augment the findings from this study and the work reported in NUREG/CR-6882.

PAPERWORK REDUCTION ACT STATEMENT

This NUREG/CR report does not contain any information collections and, therefore, is not subject to the requirements of the Paperwork Reduction Act of 1995 (44 U.S.C. 3501 et seq.).

FOREWORD

This NUREG-series report discusses research, sponsored by the U.S. Nuclear Regulatory Commission (NRC) and conducted by Oak Ridge National Laboratory, to document computer models and simulations that can be used as an auxiliary resource to simulate the operation of wireless systems in nuclear facilities and confirm potential deployment issues. This research is relevant to an important regulatory issue, in that it identifies and assesses the potential impact on safety-related systems and functions that may arise with the increasing implementation of wireless systems in nuclear facilities. The contributions of this project to the mission of the NRC are to determine the potential safety issues posed by wireless technology, and contribute to the technical basis for establishing regulatory guidance on wireless system implementation.

Wireless technology is not currently used as an integral element of safety-related systems in nuclear facilities. However, benefitting from advances in the telecommunications industry, this technology is increasingly finding its way into the nuclear industry. The most prevalent introductory uses are for (1) in-facility communications among personnel and (2) supplemental information transmission.

The goal of this study was to determine whether Bluetooth, Zigbee, and Wireless Fidelity (the three most prominent wireless technologies) can coexist effectively in an industrial environment. The results of this research will support the NRC's review of anticipated applications of wireless technology in the Nation's nuclear facilities. The regulatory products from this project also include identification of deployment issues and acceptance criteria that should be considered in regulatory reviews, as well as a computer-based assessment tool that can be used to investigate those deployment issues. Ultimately, this research will be used to develop a technical basis for guidance to address safety-related issues associated with the implementation of wireless systems in the nuclear industry. Additionally, where necessary, acceptance criteria will be developed to address the use of wireless technologies in the nuclear industry.

Brian W. Sheron, Director
Office of Nuclear Regulatory Research
U.S. Nuclear Regulatory Commission

CONTENTS

ABSTRACT.....	iii
FOREWORD.....	v
CONTENTS.....	vii
FIGURES.....	ix
TABLES.....	xi
EXECUTIVE SUMMARY.....	xiii
ACKNOWLEDGEMENTS.....	xv
ACRONYMS.....	xvii
1. WIRELESS COMMUNICATION.....	1
1.1 <i>Wireless Specifications</i>	1
1.2 <i>WiFi</i>	4
1.3 <i>ZigBee</i>	4
1.4 <i>Bluetooth</i>	4
2. ISM BAND STANDARDS.....	5
2.1 <i>ZigBee</i>	5
2.1.1 <i>ZigBee Applications</i>	5
2.1.2 <i>Physical Layer</i>	7
2.1.3 <i>MAC Layer</i>	9
2.2 <i>WiFi</i>	10
2.2.1 <i>Applications</i>	10
2.2.2 <i>802.11 Standards</i>	11
2.3 <i>Bluetooth</i>	16
2.3.1 <i>Applications</i>	16
2.3.2 <i>Physical Layer</i>	17
2.3.3 <i>MAC Layer</i>	19
3. SIMULATION OF PHYSICAL LAYER SYSTEM MODELS.....	21
3.1 <i>Monte Carlo Simulation</i>	21
3.2 <i>ZigBee Simulation</i>	21
3.2.1 <i>ZigBee Transmitter</i>	22
3.2.2 <i>ZigBee Receiver</i>	23
3.3 <i>WiFi Simulation</i>	24
3.3.1 <i>WiFi Transmitter</i>	24
3.3.2 <i>WiFi Receiver</i>	25
3.4 <i>Bluetooth Simulation</i>	26
3.4.1 <i>Bluetooth Transmitter</i>	26
3.4.2 <i>Bluetooth Receiver</i>	27
3.5 <i>Channel</i>	27
3.5.1 <i>AWGN</i>	27
3.5.2 <i>Fading</i>	28
3.5.3 <i>Ricean</i>	29
3.5.4 <i>Rayleigh</i>	29
3.5.5 <i>Simulation Design—Power</i>	30
3.5.6 <i>Simulation Design—Phase</i>	30
3.5.7 <i>Simulation Design—Delay</i>	31
3.5.8 <i>Bluetooth Delay</i>	31
3.5.9 <i>ZigBee Delay</i>	32
3.5.10 <i>WiFi Delay</i>	33
3.6 <i>Interference</i>	35
3.6.1 <i>Site-Specific Channel Model</i>	35

3.6.2	Chip Rate.....	36
3.6.3	Bandwidth.....	38
4.	RESULTS FROM GENERAL CHANNEL MODELS.....	41
4.1	<i>Coding Gain</i>	41
4.2	<i>Interference over an AWGN Channel</i>	42
4.2.1	ZigBee.....	42
4.2.2	WiFi.....	45
4.2.3	Bluetooth.....	46
4.2.4	AWGN Conclusion.....	48
4.3	<i>Interference over a Rayleigh Flat-Faded Channel</i>	50
4.3.1	ZigBee.....	50
4.3.2	WiFi.....	52
4.3.3	Bluetooth.....	54
4.3.4	Rayleigh Fading Conclusions.....	56
5.	SITE-SPECIFIC CHANNEL MODEL.....	59
5.1	<i>Wireless InSite Environment Simulation</i>	59
5.2	<i>Site-Specific Room Model</i>	59
5.2.1	Transmitter/Receiver Placement.....	60
5.2.2	Signal Properties.....	60
5.2.3	Wireless InSite Output.....	60
5.3	<i>Results</i>	65
5.3.1	Transmitter/Receiver Performance.....	65
5.3.2	BER Curves.....	68
5.3.3	ZigBee LOS.....	68
5.3.4	ZigBee NLOS.....	70
5.3.5	WiFi LOS.....	73
5.3.6	WiFi NLOS.....	75
5.3.7	Bluetooth LOS.....	79
5.3.8	Bluetooth NLOS.....	79
6.	SUMMARY.....	85
7.	REFERENCES.....	89

FIGURES

Figure	Page
Figure 1.1. United States frequency allocation.....	2
Figure 1.2. Wireless protocol coverage.....	3
Figure 1.3. Wireless protocol coverage vs bit rate.....	3
Figure 2.1. O-QPSK chip offsets.....	9
Figure 3.1. ZigBee transmitter.....	22
Figure 3.2. Chip sequence separation with half-sine pulse shaping.....	23
Figure 3.3. ZigBee receiver.....	23
Figure 3.4. WiFi transmitter.....	24
Figure 3.5. WiFi receiver.....	25
Figure 3.6. Bluetooth transmitter.....	27
Figure 3.7. Bluetooth receiver.....	27
Figure 3.8. Rayleigh simulator.....	30
Figure 3.9. Bluetooth delay amplitude.....	32
Figure 3.10. ZigBee delay amplitude.....	33
Figure 3.11. WiFi delay amplitude.....	34
Figure 3.12. Channel path simulator.....	36
Figure 3.13. Chip cycles.....	37
Figure 3.14. Channel assignments.....	39
Figure 4.1. Coding gain for ZigBee and WiFi.....	42
Figure 4.2. ZigBee signal with a Bluetooth interferer–AWGN channel.....	43
Figure 4.3. ZigBee signal with a WiFi interferer–AWGN channel.....	44
Figure 4.4. ZigBee signal with a ZigBee interferer–AWGN channel.....	45
Figure 4.5. WiFi signal with a Bluetooth interferer – AWGN channel.....	46
Figure 4.6. WiFi signal with a WiFi interferer–AWGN channel.....	47
Figure 4.7. WiFi signal with a ZigBee interferer–AWGN channel.....	47
Figure 4.8. Bluetooth signal with a Bluetooth interferer–AWGN channel.....	48
Figure 4.9. Bluetooth signal with a Bluetooth interferer–AWGN channel.....	49
Figure 4.10. Bluetooth signal with a ZigBee interferer–AWGN channel.....	49
Figure 4.11. ZigBee signal with a Bluetooth interferer – Rayleigh flat fading.....	51
Figure 4.12. ZigBee signal with a WiFi interferer – Rayleigh flat fading.....	51
Figure 4.13. ZigBee signal with a ZigBee interferer – Rayleigh flat fading.....	52
Figure 4.14. WiFi signal with a Bluetooth interferer–Rayleigh flat fading.....	53
Figure 4.15. WiFi signal with a WiFi interferer–Rayleigh flat fading.....	53
Figure 4.16. WiFi signal with a ZigBee interferer – Rayleigh flat fading.....	54
Figure 4.17. Bluetooth signal with a Bluetooth interferer–Rayleigh flat fading.....	55
Figure 4.18. Bluetooth signal with a WiFi interferer–Rayleigh flat fading.....	55
Figure 4.19. Bluetooth signal with a ZigBee interferer–Rayleigh flat fading.....	56
Figure 5.1. Transmitter and receiver locations.....	61
Figure 5.2. Waveform properties.....	61
Figure 5.3. Transmitter properties.....	62
Figure 5.4. Propagation paths for WiFi transmitter #5.....	63
Figure 5.5. Propagation paths for WiFi transmitter #7.....	63
Figure 5.6. Power vs delay–WiFi transmitter #5.....	64
Figure 5.7. Power vs delay–WiFi transmitter #7.....	64
Figure 5.8. ZigBee transmitter #3 with ZigBee interferers.....	68
Figure 5.9. ZigBee transmitter #6 with Bluetooth interferers.....	69

Figure 5.10. ZigBee transmitter #6 with WiFi interferers.	70
Figure 5.11. ZigBee transmitter #7 with Bluetooth interferers.	71
Figure 5.12. ZigBee transmitter #7 with WiFi interferers.	72
Figure 5.13. ZigBee transmitter #10 with WiFi interferers.	72
Figure 5.14. WiFi transmitter #5 with ZigBee interferers.	74
Figure 5.15. WiFi transmitter #6 with WiFi interferers.	74
Figure 5.16. WiFi transmitter #2 with Bluetooth interferers.	76
Figure 5.17. WiFi transmitter #9 with Bluetooth interferers.	77
Figure 5.18. WiFi transmitter #10 with Bluetooth interferers.	77
Figure 5.19. WiFi transmitter #11 with WiFi interferers.	78
Figure 5.20. Bluetooth transmitter #1 with ZigBee interferers.	80
Figure 5.21. Bluetooth transmitter #3 with WiFi interferers.	80
Figure 5.22. Bluetooth transmitter #6 with Bluetooth interferer.	81
Figure 5.23. Bluetooth transmitter #7 with WiFi interferers.	82
Figure 5.24. Bluetooth transmitter #10 with ZigBee interferers.	83
Figure 5.25. Bluetooth transmitter #11 with Bluetooth interferer.	84

TABLES

Table	Page
2.1 ZigBee symbol-to-chip mapping sequences.....	8
2.2 Bit pattern.....	13
3.1 Bluetooth delay.....	32
3.2 ZigBee delay.....	33
3.3 WiFi delay.....	35
5.1 WiFi transmitter #5 propagation paths.....	65
5.2 WiFi transmitter #7 propagation paths.....	65
5.3 Bluetooth propagation paths.....	66
5.4 WiFi propagation paths.....	66
5.5 ZigBee propagation paths.....	67

EXECUTIVE SUMMARY

This report details an interference study of the three most prominent wireless devices in use today, using computer models and simulations. The goal is to determine whether Bluetooth, Zigbee, and Wireless Fidelity (WiFi) wireless devices can coexist in an industrial environment. All three wireless devices operate in the 2.4-GHz industrial, scientific, and medical (ISM) frequency band. Simulations are conducted because of the amount of time that it would take to physically collect measurements for a plausible coexistence study. Numerous possible combinations of transmitters, receivers, and interferers are simulated. Both general channel models and site-specific channel models, incorporating the physical layout of an industrial building, are created and computed to simulate the effects of interference for certain combinations of different wireless devices. The considered channel models are basic additive white Gaussian noise (AWGN), general Rayleigh fading, and site-specific Ricean and Rayleigh fading.¹

The effects of Bluetooth as an interferer are considered first, followed by the response of Bluetooth to interfering devices. For the most part, throughout all three of the interference models, Bluetooth was shown to intrude upon the performance of the other devices the least. This means that Bluetooth is a “good neighbor” and allows for the coexistence of devices within the 2.4-GHz ISM frequency band. This fact can be attributed not only to Bluetooth’s transmitting at a relatively low power level, 4dBm, but also to the frequency-hopping nature of the scheme Bluetooth employs. On average, most hops will not lie within the bandwidth of another signal. Even when these hops do occur within the wide bandwidth of WiFi, the nature of the narrowband signal that Bluetooth encompasses compared with the wideband of WiFi minimizes the effects when the WiFi signal is decoded.

The overall performance of ZigBee in the roles of both the interferer and the transmitter is considered and found to be very similar to that of Bluetooth, although for different reasons. ZigBee does not seem to interfere with other devices because of its low radiated power level, which is 4 dBm below that of Bluetooth and 17 dBm lower than that of WiFi. ZigBee also has a relatively small bandwidth, only 2 MHz; therefore the effects of ZigBee on a wide signal such as WiFi are minimized. When ZigBee is assumed to be the transmitting device, it is able to defend itself from the effects of both the Bluetooth and ZigBee interferers successfully through the processing gain associated with the spreading of its signal until the point at which the interfering signal’s power completely dominates the signal. Alternately, because WiFi begins with such a large power advantage over ZigBee, it inadvertently hinders the performance of the ZigBee signal. Therefore, deployment of ZigBee devices in the presence of WiFi must be undertaken with caution.

As an interfering device, WiFi tends to limit the performance of other devices and does not allow for coexistence. In the general channel models, WiFi limits the performance of other transmitters as a result of its wide bandwidth and the likelihood that these devices will be located within the same frequency space. In the site-specific scenario, this dominance is taken one step further as a result of the increased power advantage that WiFi has over the other wireless protocols. WiFi limits the performance of a transmitting device long before the other interferers do. One reason is that WiFi was one of the first

¹ The AWGN channel model is one in which the only impairment is the linear addition of wideband or white noise with a constant spectral density and a Gaussian distribution of amplitude. The model does not account for the phenomena of fading, frequency selectivity, interference, nonlinearity or dispersion. Rayleigh fading is a statistical model that assumes that the power of a signal passing through a transmission medium will vary randomly, or fade, according to a Rayleigh distribution — the radial component of the sum of two uncorrelated Gaussian random variables. Rayleigh fading occurs where there are multiple indirect paths between transmitter and receiver, with no distinct dominant path. The idea is that there is no clear desired signal. The signals arriving at the receiver, instead, represent a sum of multiple, independent, random variables. Ricean fading is similar to Rayleigh fading, with the exception that there is a direct, or at least dominant, component in the mix of signals that reach the receiver. It effectively “biases” the Rayleigh distribution, generating a stochastic distribution about a more firmly characterized mean amplitude value.

standards developed for products within the 2.4-GHz band; therefore, the presence of other devices operating in the same band was not considered. Considering the effects that other devices have upon a transmitted WiFi signal, WiFi appears to be the least well equipped to handle the effects of interference. WiFi seems to be sensitive to any small changes in its signal caused by the Bluetooth and Zigbee interferers.

This study covers only a coexistence assessment of Bluetooth, Zigbee, and WiFi wireless devices. Related work performed to date under JCN Y6475 includes a companion study that identified and assessed wireless technologies, both current and emerging, that have the potential for deployment in nuclear facilities. The companion study explored the technology differentiators that need to be considered before deploying a wireless system. In addition, deployment issues were investigated and current wireless deployments in nuclear facilities were examined. Results of that study are reported in NUREG/CR-6882, *Assessment of Wireless Technologies and Their Application at Nuclear Facilities* (May 2006). Security issues are also briefly explored in NUREG/CR-6882, but it was not the intent of that study to comprehensively address wireless security issues. The need for a comprehensive treatment of wireless security has been recognized and additional work will be performed in this area. The results of the planned wireless security study are anticipated to augment the findings from this study and the work reported in NUREG/CR-6882.

ACKNOWLEDGEMENTS

The authors wish to thank Bill Kemper, Branch Chief, and Tekia Govan, JCN Y6475 Project Manager, of the U.S. Nuclear Regulatory Commission Office of Nuclear Regulatory Research for their help in initiating, planning, and implementing this research effort.

ACRONYMS

ACL	asynchronous connectionless
ASK	amplitude-shift keying
AWGN	additive white Gaussian noise
BER	bit error rate
BPSK	binary phase-shift keying
CCK	complementary code keying
CRC	cyclic redundancy check
CSMA/CA	carrier sense multiple access with collision avoidance
DC	direct current
DQPSK	differential quadrature phase-shift keying
DSSS	direct-sequence spread spectrum
FCC	Federal Communications Commission
FEC	forward error correction
FHSS	frequency hopping spread spectrum
FM	frequency modulation
FSK	frequency-shift keying
GFSK	Gaussian frequency-shift keying
GMSK	Gaussian minimum-shift keying
GSM	global system for mobile communication
HVAC	heating, ventilation, and air conditioning
IEEE	Institute of Electrical and Electronic Engineers
IR	infrared
ISI	inter-symbol interference
ISM	industrial, scientific, and medical
LAN	local area network
LOS	line of sight
MAC	medium access control
MAN	metropolitan area network
MSK	minimum-shift keying
NLOS	non-line of sight
OFDM	orthogonal frequency division multiplexing
O-QPSK	offset quadrature phase-shift keying
PAN	personal area network
PDA	personal data assistant
PHR	physical header
PHY	physical
PLCP	physical layer convergence protocol
PN	pseudonoise
PSDU	physical layer convergence protocol service data unit
PSK	phase-shift keying
QAM	quadrature amplitude modulation
QPSK	quadrature phase-shift keying
RF	radio frequency
SCO	synchronous connection-oriented
SFD	start frame delimiter
SHR	synchronization header
SIR	signal-to-interference ratio
SNR	signal-to-noise ratio
spc	samples per chip

UNII	unlicensed national information infrastructure
WAN	wide area network
WiFi	wireless fidelity
WiMax	worldwide interoperability for microwave access
XOR	exclusive or

1. WIRELESS COMMUNICATION

Wireless communication is accomplished through the use of electromagnetic waves. A radio wave can “oscillate at frequencies between about 3 kHz and 100 GHz” [1]. This spectrum is broken down into multiple sections designed for private use, reserved for government use, or considered to be unlicensed and free to all users so long as they abide by government regulations. The spectrum allocation is pictured in Figure 1.1. The lower portion of the spectrum, any frequency below 2.4 GHz, is used for a wide variety of devices and purposes, but it is primarily concerned with voice communication.

Entertainment uses of this lower spectrum include amplitude modulation and frequency modulation (FM) radio, from 535 kHz to 1.7 MHz and from 88 to 108 MHz, respectively. Television stations are also located in this spectrum, ranging from 54 to 220 MHz (very high frequency) and 470 to 890 MHz (ultra high frequency), subtracting the band of FM radio frequencies. Voice communication in this lower spectrum can be used in cell phones, which use the frequencies 824 to 849 MHz, or in Citizens’ Band radios from 26.96 to 27.41 MHz. Other devices also operate within the frequencies of this spectrum, including remote-controlled toys, global positioning system devices, and even garage door openers.

1.1 Wireless Specifications

Standards that deal with wireless communication primarily consist of the family of Institute of Electrical and Electronic Engineers (IEEE) 802 standards. Several of these standards, along with their intended coverage areas and corresponding bit rates, can be found in Figure 1.2, which shows that there are four main types of coverage areas. A personal area network (PAN) is defined as the immediate space surrounding a device, usually confined to a single room, and has a range only on the order of 10 m (~33 ft). A local area network (LAN) is an expansion of a PAN and can include multiple rooms; this type of network usually delivers service to a number of devices, whereas a PAN is designed as a point-to-point connection. Expanding upon a LAN, a metropolitan area network (MAN) can deliver point-to-multipoint communication among devices within a business building or an entire block of business buildings; a MAN is typically referred to when dealing with an urban environment. On the other hand, in a rural environment in which there are very few obstructions, a wide area network (WAN) is commonly considered. A WAN can serve an entire community.

A representation of how the types of networks are arranged so that the coverage area for one does not overlap with other coverage areas can be found in Figure 1.3. Notice that when the coverage area in which a set of devices is intended to deliver service increases, the maximum bit rate for each network has a tendency to drop. This tendency can be attributed both to the limits that distance places on bandwidth and to the fact that as a network needs to provide service to more users, the bandwidth available to each user decreases. Only devices complying with the Wireless Fidelity (WiFi), Zigbee, and Bluetooth standards are explored in the research documented in this report, as they operate in the 2.4 GHz band, the frequency range of interest. Detailed discussions of these standards, as well as standards that do not operate in the 2.4 GHz band (WiMax, WiMedia, and MobileFi) can be found in NUREG/CR-6882, *Assessment of Wireless Technologies and Their Application at Nuclear Facilities* [2].

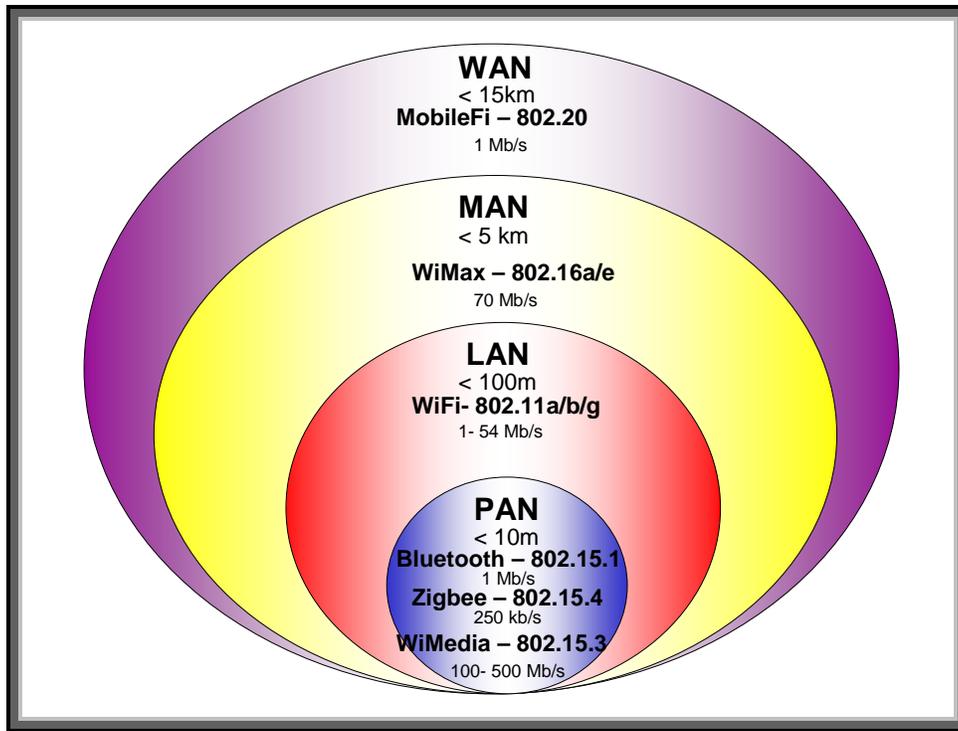


Figure 1.2. Wireless protocol coverage.

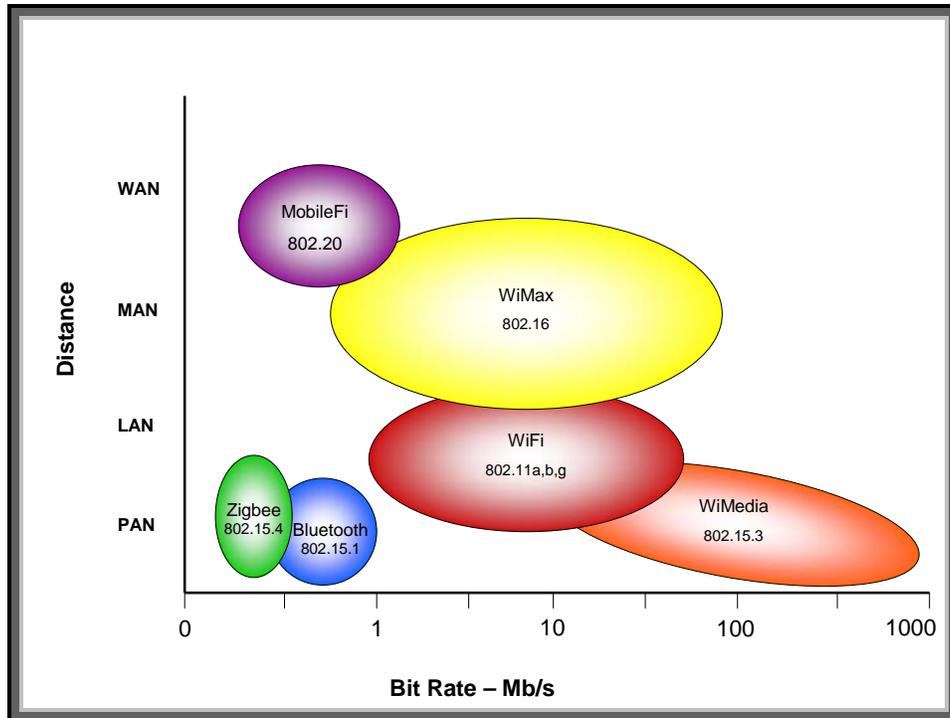


Figure 1.3. Wireless protocol coverage vs bit rate.

1.2 WiFi

WiFi is perhaps the most widely known of the six standards mentioned. It is used in routers as a link between a computer and the Internet. WiFi comes in three different forms: 802.11a, 802.11b, and 802.11g. With these three different types of devices, data rates of between 1 and 54 Mbit/s are possible. A typical range in the area of 100 m (~328 ft) can be expected for all devices. WiFi is one of the three standards selected for our research; it will be explored in detail in Section 2.

1.3 ZigBee

Another standard that will be analyzed in depth in Section 2 is ZigBee. ZigBee is geared toward low-power, low-rate communication techniques used in functions such as home automation and sensors. ZigBee achieves a data rate of only 250 kbit/s, but because it serves an area of only 10 to 70 m (~33 to 230 ft), it can use more power-efficient methods of transmission.

1.4 Bluetooth

The third protocol under investigation in this report is Bluetooth. Bluetooth is a cable-replacement device used mainly in conjunction with computers but also finding applications in cell phones. It was developed to be a lower-power, lower-cost alternative to WiFi, much as ZigBee was developed to be a lower-power and lower-cost solution than Bluetooth. Bluetooth can provide 1 Mbit/s data rates for coverage from a few meters to a hundred meters, depending on its three different transmitted power levels. Bluetooth will also be explored in detail in Section 2.

2. ISM BAND STANDARDS

The industrial, scientific, and medical (ISM) band is a span of frequencies in the 2.4-GHz range, more specifically 2.4 to 2.4835 GHz. It is a free and unlicensed band that can be used by anyone to transmit information wirelessly. The U. S. Government created this range so that no one entity could hold the rights to use these frequencies of interest. It was created in hopes that any device acting wirelessly would have the opportunity to exist at the frequency along with other wireless devices. Since the creation of the ISM band, the Federal Communications Commission (FCC) has put numerous regulations on its use, and IEEE has adopted several standards for devices that can be used in the ISM band. Three such standards, along with their most common protocols, are IEEE Std 802.11b, under which part of the protocol of WiFi falls; IEEE Std 802.15.4, for which ZigBee is the rising protocol; and IEEE Std 802.15.1, also referred to as Bluetooth. In some cases, these protocols are not the only ones being used under a given standard, but because they are the most common, they will be treated as one and the same.

Explanations of the Zigbee, WiFi, and Bluetooth standards are given below. An overview, possible applications, and descriptions of the PHY and MAC layers are provided. The intent is to make the reader aware of how the ISM-band devices operate and how the devices differ in their characteristics. This information should be helpful in understanding the research results.

2.1 ZigBee

ZigBee has come about out of convenience more than anything else. A collection of major corporations—the most significant eight being Ember, Freescale, Honeywell, Invensys, Mitsubishi, Motorola, Philips, and Samsung—are all committed to standardizing cost-effective, low-power, wirelessly networked monitoring and control products based on an open global standard [3]. Consequently, these companies are looking for a protocol that does not use a large amount of bandwidth and is not very complex, because both higher bandwidth and complexity lead to increased cost and power consumption. Many products today fall into this category, opening up a large market; since its inception, more than a hundred more companies have joined the ZigBee Alliance.

2.1.1 ZigBee Applications

The market base for ZigBee consists of three main categories: personal, business, and industrial. The personal category deals with individual consumers; the main focus for ZigBee in this area is home automation, creating ways to make everyday activities easier through the use of wireless devices. The business market deals with companies that are not producing any type of product on-site, and the main focus for ZigBee is office applications. The industrial category entails a broad range of services, from a nuclear power plant producing electricity to a FedEx warehouse distributing packages. Because the industrial environment is of the greatest concern, the focus of this report will be on these types of venues.

2.1.1.1 Sensory Devices

Wireless sensors are probably the biggest market available to ZigBee products within an industrial environment simply because they encompass such a vast array of applications. Within a nuclear power plant, the purposes of these sensors can range from controlling the environment of the power plant itself, including heating and lighting, to protecting the safety of the workers, and even security protection for the plant.

The environment of the power plant can include both the lighting and the heating, ventilation, and air conditioning (HVAC), which can be controlled through the use of either light or temperature sensors [4].

For instance, if sensors placed on the windows in a room in the plant detect enough light entering the window, then the lights within the room can be dimmed because of the decreased need for lighting; the result would be less use of electricity, meaning lower overhead costs. The HVAC can be controlled in a similar way; if temperature sensors placed throughout a large control room maintain a temperature below a certain threshold, then the heat will be turned on, and vice versa for the cooling. In offices located within a power plant, if motion sensors do not detect someone is occupying an office, then both the lights and the HVAC can be turned off to the room until someone reenters the office, the temperature goes beyond a second threshold, or a certain time of the day is reached. (In the morning, for instance, the lights and HVAC could be turned on in anticipation of someone entering the room.)

2.1.1.2 Safety

The safety of the workers could also be protected through the use of ZigBee sensors. This protection could range from making sure that machinery is operating properly, to monitoring reactor coolant temperature levels, to maintaining healthful conditions within the work environment.

For instance, sensors could be placed on water pumps and generators to ensure that they are performing properly, maintaining their appropriate revolutions-per-minute speeds, and not drawing too much energy. Sensors could also be used to monitor the wear that various parts within the machines experience in everyday use, and sound sensors could observe noises the machines are producing. If the sensors recognize abnormal noise or that parts are close to failure, then an alarm could be triggered.

Another potential application for sensors is the monitoring of various processes throughout a nuclear power plant. Temperature gauges and other sensors could be placed within a coolant chamber to report not only whether the coolant is at an acceptable temperature but also the coolant level itself and whether a leak of any kind has occurred. These types of monitoring applications could be extended to having radiation sensors and other types of warning sensors placed throughout the plant to warn of contamination in the air or within the cooling water systems.

Certain traits of the ZigBee specification could be exploited by these tasks. For instance, through the use of on-chip intelligence, the sensors could analyze the information they are accumulating on their own and then relay information to a main terminal only if a given threshold has been exceeded. Such a system would allow for slower data rates because of the decrease in the amount of data needing to be exchanged, thereby saving battery consumption. The reduction in the duty cycle of these devices from transmitting data only intermittently would increase their battery life expectancy to a range of 1 to 3 years, an appealing figure compared with the hours and/or days of Bluetooth and WiFi. The battery life extension would allow users to place the sensors without concern that they would fail as a result of battery failure. If the devices were not so power-efficient, the battery in each device would constantly need replacement, defeating the convenience of wireless devices.

Another attribute of ZigBee products that helps them keep power consumption to a minimum is that they can enter a sleep mode; in this state they consume almost no power but can be awakened at any time. There is typically a 15-ms delay for a device to change from sleep mode to being awake, and then there is another 15-ms delay for the active slave to access the channel. This is comparable to the 30 ms it would take for enumeration (i.e., the time required for a new device to access a given network) [3]. For these types of sensors, a 15- or 30-ms delay is well within the latency requirements because of the polling nature of their applications, which means that most devices will spend much of the time in sleep mode, only awakening to send data about their current state at a given time.

2.1.1.3 Security

Wireless sensors (ZigBee devices) could also serve to aid the functionality of various security devices. Whether used with motion sensors on the ceiling or pressure sensors within the floor, they could be used to detect whether a restricted area has been accessed and then alert the central security system, which could then relay information to other security features (e.g., controls for lights, alarms, door locks, and cameras).

One topic not yet mentioned is the distance across which a wireless device must be able to transmit. The range of most ZigBee applications will fall within 10 m (~33 ft), although in some applications it can exceed 70 m (~230 ft). For more coverage area, a higher transmitted power is required, causing the device to draw more energy from the battery and creating the need to change or recharge the battery more often. A larger coverage area is typically not the aim for a device using ZigBee, so it is not customary to use ZigBee in this way, although it could be applicable in certain situations.

One way ZigBee gets around the distance dilemma is to relay information between several devices until it reaches the desired device. ZigBee can conform to various topologies, two of which are star and peer-to-peer networks. Within a star network, there is only one coordinator, and the rest of the devices are considered the slaves. Within this configuration, the slaves may talk to only the coordinator and not to each other. In a peer-to-peer network, also considered a cluster, there is only one coordinator, but the slaves may communicate with each other [4]. To accommodate the previous situation in which the information obtained by a ZigBee device needs to travel a long distance, the total transmission length may be broken up between several devices or clusters, allowing for less transmitted power and thus longer battery life per device.

Although sensors represent a broad range of applications, they are not the only applications in which ZigBee could excel. ZigBee could be used in radio frequency (RF) tagging—either tagging of employees to allow them access to buildings and other areas or tagging of inventory to track packages and equipment by allowing ZigBee transmitters to give updates of their locations at regular intervals. Therefore, ZigBee should be considered in the design of any type of wireless device.

2.1.2 Physical Layer

To better understand why ZigBee is so useful for the previously explored applications, a background of the actual specifications of the protocol itself is needed. ZigBee incorporates the use of a direct-sequence spread spectrum (DSSS) system to help make it more robust and less susceptible to interference. In the 2.4-GHz range, ZigBee uses the frequencies 2.405 to 2.480 GHz. This range is subdivided into 16 different channels, each with an equal spacing of 5 MHz. Allocating the available bandwidth in this fashion allows for signal quality improvement due to less inter-symbol interference (ISI), because while the channel has an available bandwidth of 5 MHz, the signal occupies a spectrum of only 2 MHz. This also allows for the implementation of more channels, if the need ever arises for such an improvement, because of the extra 3 MHz of available channel bandwidth [5].

ZigBee has a basic bit rate of 250 kbps for the 2.4-GHz frequency range. To spread the signal and make it become DSSS, the signal is mapped into a 32-chip-length pseudonoise (PN) sequence. Unlike most other DSSS systems, ZigBee does not multiply input bits by a PN sequence; it maps just the input bits to a pre-defined PN sequence. ZigBee has a databank of 16 different 32-chip sequences. These 16 different chip sequences can represent four information bits; therefore, four bits represent a ZigBee symbol. A 250-kbps bit rate, when divided by the four bits per symbol, results in a 62.5 ksymbol/s symbol rate. Taken one step further, each symbol represents the 32 chips in a PN sequence, so the chip rate becomes 2.0 Mchip/s.

The set of sixteen 32-chip PN sequences are quasi-orthogonal to each other and come from cyclic shifts and/or conjugation. The first 8 sets of the 16 simply cyclically shift the length by 4 chips each time, so after the first sequence, the 4 chips at the end are put in the front, and the rest of the chips are pushed back by 4. For the next sequence, the last four chips of the second sequence are placed in the front, and all the bits are again pushed back by four. At the ninth sequence, a new sequence is introduced that takes the first original sequence and inverts the odd indexed chips (starting with the first chip indexed as the zero point). Once the ninth sequence is created, the following sequences are then found by shifting the ninth sequence in the same manner as were the first eight. This process continues until the full set of 16 different PN sequences is made. The set of 16 sequences and the corresponding data symbols can be found in Table 2.1 [6].

Table 2.1. ZigBee symbol-to-chip mapping sequences

Data symbol (decimal)	Data symbol (binary) (b_0, b_1, b_2, b_3)	Chip values ($c_0, c_1, \dots, c_{30}, c_{31}$)
0	0000	11011001110000110101001000101110
1	0001	11101101100111000011010100100010
2	0010	00101110110110011100001101010010
3	0011	00100010111011011001110000110101
4	0100	01010010001011101101100111000011
5	0101	00110101001000101110110110011100
6	0110	11000011010100100010111011011001
7	0111	10011100001101010010001011101101
8	1000	10001100100101100000011101111011
9	1001	10111000110010010110000001110111
10	1010	01111011100011001001011000000111
11	1011	01110111101110001100100101100000
12	1100	00000111011110111000110010010110
13	1101	01100000011101111011100011001001
14	1110	10010110000001110111101110001100
15	1111	11001001011000000111011110111000

Once the appropriate PN sequence has been chosen for the input symbol, successive chip sequences are concatenated, and the chips are modulated using offset quadrature phase-shift keying (O-QPSK) [6]. Because the 802.15.4 standard specifies that half-sine pulse shaping must be used, the O-QPSK modulation is equivalent to minimum-shift keying (MSK), which can be defined as continuous-phase frequency-shift keying (FSK) with a minimum modulation index ($h=0.5$) that will produce orthogonal signaling [7].

O-QPSK is a form of quadrature phase-shift keying (QPSK), which can send two bits of information per symbol, but O-QPSK employs a technique of delaying the Q phase of transmission by one bit period. QPSK is formed by separating a signal into its I phase (in phase or direct phase) and Q phase (quadrature-phase). It can be thought of as the real and imaginary parts of a complex number, with the I phase being the real part and the Q phase being the imaginary part. Delaying the Q phase by a bit period allows for only one zero crossing to occur at a time, with a zero crossing of the phases representing a change in the data bit. With only one zero crossing occurring at a time, the phase transition for O-QPSK is only 90° instead of the 180° that can occur for QPSK. For half-sine pulse shaping, this results in one phase of the signal being at its peak of the sine wave while the other is at a zero crossing and vice versa for the other phase, allowing for a much more reliable demodulation of the signal.

For ZigBee, to separate the signal into its I and Q phases, the PN sequence is broken down into 2 sets of 16 different chips. From the original sequence, the even-indexed chips are placed in the I phase, and the odd-indexed chips are placed in the Q phase. For the offset found in O-QPSK, the Q phase is delayed by half of a chip or double the inverse of the chip rate. It was found earlier that the chip rate was equal to 2 Mchip/s ; T_c in Figure 2.1 corresponds to the inverse of that value [6]. The figure also shows that the chips in the individual phases have duration times of twice the chip period; thus, the per-phase chip rate is half the overall chip rate and is equal to 1 Mchip/s .

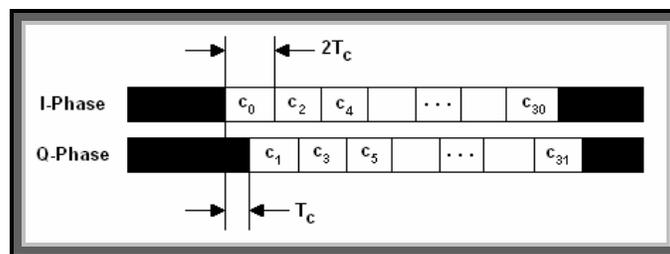


Figure 2.1. O-QPSK chip offsets.

2.1.3 MAC Layer

The medium access control (MAC) layer deals with how information is sent at the packet level. As stated earlier, a data symbol consists of four information bits. To know where these four bits come from, a more abstract look will be taken. But first it is important to note that for IEEE Std 802.15.4, transmission is done in octets, or groups of eight bits, as will be shown in the following explanation. All the following information about how the packets are formed is taken from the 802.15.4 standard. A packet, also known as a physical (PHY) beacon packet, consists of three fundamental elements, which are the synchronization header (SHR), the physical header (PHR), and the payload.

The SHR consists of the synchronization components, namely the preamble and the start frame delimiter (SFD), and is a total of five octets, or 40 bits, long. The preamble requires four of those octets, and each octet is composed of all binary zeros, or 32 zeros. It is used to synchronize the receiver with the incoming

signal. The fifth and final octet is called the SFD field. The SFD field is used to designate when the incoming data are about to begin. After the 32 zeros go through, the SFD is composed of the bits 11100101, and once the receiver recognizes these bits, it knows that the preamble is over and it must prepare itself for the next stage, which is the PHR.

The PHR is only one octet long but is vital to receiving the information that was originally sent. It specifies how long the PHY layer convergence protocol (PLCP) service data unit (PSDU) is formed in the MAC sub-layer; it is not within the scope of this report, but it is essentially the sent information itself and will be in octets. A maximum of 2^7 , or 128, octets can be sent for one packet, corresponding to 1024 bits of total information. The final bit of the PHR octet is reserved for later use.

The final component of the packet is the PHY payload or PSDU. It can be of varying lengths, as specified by the PHR above. One small detail to note is that when the octets are grouped into data symbols, each data symbol is composed of four bits, while the packet information is grouped into eight bits. The way this discrepancy is resolved is to place the first four bits (b_0 , b_1 , b_2 , and b_3) into one data symbol and the second group of four bits (b_4 , b_5 , b_6 , and b_7) into a second data symbol [6].

Now that the fundamentals for the ZigBee standard have been set forth, the next standard, WiFi, followed by Bluetooth, will be presented so that the differing standards can be compared.

2.2 WiFi

It is hard to find a place nowadays in which wireless Internet is unavailable. Whether in hotel rooms, local cafés, or business offices, people everywhere are demanding wireless Internet access points because of the surge in the number of laptop computers and PDAs.

2.2.1 Applications

Within an industrial environment, laptops and PDAs are being introduced throughout the workplace. With machines becoming less dependent on a human interface and processes being converted to become computerized, there is a need to be able to upgrade and test equipment. This need has depended on allowing the equipment to be connected to a laptop so a diagnostics test can be performed and newer software downloaded. Rather than loading all of the necessary software on the laptop, it can be placed on a main server and simply downloaded as it is needed through the use of a wireless network.

Wireless networks also could be exploited by using devices that are accessible to the Internet through WiFi. PDAs could allow technicians to communicate with troubleshooters while inspecting faulty equipment; they could use picture and text messaging to obtain instant feedback and avoid being delayed by poor communication. The use of PDAs could also allow for better methods of ordering supplies and equipment needed to maintain proper working conditions. Rather than taking inventory, checking the needed supplies, and then entering orders into a computer, a more efficient approach would be to order materials directly through a PDA as the inventory is taken. This approach would increase efficiency and decrease the chances for a mistake.

More and more laptops are replacing desktop computers because of their portability. The use of wireless Internet would allow for legacy systems to be upgraded without the added cost of running Ethernet cord throughout a building and maintaining hundreds of access ports. If WiFi routers were used, a minimal amount of cord would need to be placed within walls, floors, and ceilings to connect the routers, cutting down on the added overhead from remodeling. As long as the networks were secure and did not allow access to unwanted users, these networks would work just as well as wired ones; and they would provide users the added flexibility of not needing to work near an access port.

2.2.2 802.11 Standards

Within the 802.11 family of standards, the three that have found prominence today are 802.11a, 802.11b, and 802.11g. IEEE Stds 802.11a and 802.11b can be considered distinct protocols within themselves; IEEE Std 802.11g is a fusion of those two standards into one, because 802.11g encompasses the more attractive trait of 802.11a, which is the speed, and the broad compatibility of 802.11b. The relevant functional aspects of each of the three standards and a more in-depth explanation of their operations are provided in the following sections. An interesting similarity to note is that all three protocols instantiate the same MAC layer defined by the 802.11 standard, and it is only the way in which the PHY layer is implemented that distinguishes the protocols from one another. Because the MAC layers are the same for all three and secondary to the modeling effort, information pertaining to the MAC layer will be presented only as needed. For completeness, explanations are given for all three protocols, even though only IEEE Std 802.11b will be part of the modeling effort. IEEE Std 802.11a will not be pursued within this study because it operates in the 5-GHZ UNII band that is unsuitable for the interference studies at the 2.4-GHZ band. Because 802.11g uses the same basic protocol as WiFi and operates in the same frequency band, the interference simulation results for 802.11g will be similar to the WiFi results. Hence, 802.11g will not be a protocol pursued at this time.

2.2.2.1 802.11b

The most prominent of the three protocols for IEEE Std 802.11 is 802.11b (also referred to as WiFi), which has found a wide market. Almost all wireless routers today are WiFi compliant, even though numerous 802.11g-compliant devices are coming on the market. The universal switch from WiFi to 802.11g has not yet fully occurred; thus, WiFi will be the main topic of discussion in the following sections. Unlike for ZigBee, the aim for WiFi is not to implement it in wireless sensors or as a simple cable replacement for computer devices but to connect devices through the use of the Internet. WiFi was made with data speed in mind, not low power consumption or low complexity. Therefore, there is no real comparison between ZigBee and WiFi. In most ZigBee applications, WiFi would not only consume too much power but also would also be overkill; none of ZigBee's applications require high-speed data rates. In a similar manner, ZigBee is not suited to perform WiFi tasks because WiFi has the potential to perform a given task 44 times faster than ZigBee. Consequently, comparing the two systems would be like comparing dial-up with broadband access. With such a decrease in throughput through the use of ZigBee, the efficiency within the work environment would also be drastically decreased; therefore, WiFi would be the more desirable choice for Internet connection.

2.2.2.2 802.11b Physical Layer

WiFi can be broken down into four different data rates (1, 2, 5.5, and 11 Mbps), but before exploring them, one must understand some common requirements set forth for all data rates. First and most important, WiFi operates in the same ISM band as ZigBee, specifically from 2.4 to 2.4835 GHz. In the United States, out of a total of 14 channels, only the first 11 are allowed to be used; and their center frequencies range from 2.412 to 2.462 GHz, with each channel occupying approximately 22 MHz. Thus, only three channels do not overlap each other—channels 1, 6, and 11 [8]. Second, the aggregate chip rate for all four data rates is 11 Mchip/s, which corresponds to occupying a total bandwidth of 22 MHz. Finally, while the 802.11 standard supports three types of PHY layers—DSSS, frequency hopping spread spectrum (FHSS), and infrared (IR)—it is recognized that all four data rates encompass the DSSS system, but only the two slower data rates are used in FHSS and IR systems.

A brief introduction is needed into the workings of the 1- and 2-Mbps schemes to understand how the 5.5- and 11-Mbps rates were later achieved, although all of the testing done for this study was conducted

assuming an 11-Mbps rate. The main reason it is important to understand the slower data rates is that all compliant WiFi devices that operate at any given data rate must also be able to operate at the 1-Mbps rate. The requirement was set into place because in order for a device to send or receive data, it must know specifics about the message itself; therefore, all vital information about a message, including its data rate and its payload, is sent in the headers (specifically called the PLCP headers), which are transmitted at 1-Mbps modulation. This is one way of ensuring that not only will the receiver know which modulation scheme to use, whether it is for the slower or faster data rates, but also the other transmitters will know how long to refrain from using the channel before attempting to access it for themselves. Because the header contains information about the rate and the payload of a signal (which in this context signifies the time duration of the packet transmission), the other transmitters can decipher this information from the header and know to wait at least that long before trying to access the channel.

The 1-Mbps data rate is realized using a combination of a Barker code DSSS spreading function along with a binary phase-shift keying (BPSK) modulation scheme. The Barker code is the result of an “exclusive or” (XOR) operation between two operands, an 11-bit sequence (10110111000) and the input data stream [9]. XOR is a logic-based action on two operands that produces a logical value of “true” if and only if one, but not both, of its operands is true. The XOR operation is done by concatenating the successive 11-chip sequences resulting from the XOR between the 11-bit Barker code and the individual message bits. Thus, each message bit is encoded by 11 chips—the 1 bit at a 1-Mbps rate multiplied by the 11 chips that represent that 1 bit, yielding a chip rate of 11 Mchip/s as previously specified. The resulting chip sequence is then modulated using BPSK modulation, which represents one bit per symbol of transmission.

On the contrary, the 2-Mbps data rate incorporates a QPSK modulation scheme that can represent two bits of information per transmitted symbol. Thus, twice the information can be sent using QPSK in the same bandwidth as in the 1-Mbps BPSK scheme. Therefore, the information is encoded in the same way as before, but now instead of BPSK modulation, differential QPSK (DQPSK) is used. DQPSK modulates sequential symbols by a phase rotation. This increase in the bit rate occurs at the expense of either a need for a higher transmitted power or a diminished range of effectiveness. Because the FCC has put regulations on the maximum effective transmission power in the ISM band, which is 1000 mW, the only factor left to control is the effective range. Thus, as the distance between the transmitter and receiver increases, the modulation scheme used adjusts to one of the slower rates to maintain a tolerable signal level [9].

The 11- and 5.5-Mbps data rates can be thought of as an extension of the 2-Mbps data rate previously discussed. Both schemes maintain the 11-Mchip/s chip rate, and they both are modulated using DQPSK. The difference is that the two higher data-rate formats incorporate a different and more complex DSSS technique that will change the way bits are grouped and the way they are spread.

The technique implemented by the higher data rates is a design first conceived by Marcel J. E. Golay in 1951. Golay had been doing some work with uses of spread-spectrum models pertaining to light emitting through slits. In doing this work, he stumbled across complementary sequences that proved to contain valuable mathematical properties. He later published a paper about the binary sequences he had discovered, mainly about what made them so appealing and how they were created. It was an extension of work similar to this that helped bring about the evolution of complementary code keying (CCK) codes, which are types of polyphase complementary codes. The codes that Golay helped discover are types of polyphase codes called binary complementary codes. The difference between the two is that binary complementary codes take on binary values (ones or zeros), while polyphase complementary codes can take on a number of different values so long as they maintain complementary properties. For the case at hand, CCK uses codes containing four different phase values that take on complex values, namely the values $\{1, -1, j, -j\}$ [10].

Because of its superior coding properties compared with the Barker sequence, CCK was implemented to make the data transmission of WiFi more efficient and robust. The efficiency comes from the increase in data rate within the same signal bandwidth, and the robustness comes from the improved coding capability of incorporating multiple sets of possible transmitted code words rather than just one Barker sequence implemented by the slower data rates.

To increase the speed of the data being transmitted, CCK transmits eight complex chips for every eight information bits, yielding an 8:8 or 1:1 ratio rather than the 11:1 ratio of chips to bits for 1-Mbps transmission. This accounts for an 11-fold increase in data throughput. The increase in data is accomplished in a multistage process. First, groups of eight bits must be gathered to create an information symbol. The individual groups of eight bits are then separated into two unequal partitions, one portion being the first two bits of the symbol and the second portion the last six bits of the symbol. The first portion will be used later to modulate the signal but will be ignored for now. Because the second portion contains a group of 6 bits, there is a possibility of 2^6 , or 64, potential combinations. These 6 bits will determine which one of the 64 possible 8-chip code words will be output. This information will be more useful after the discussion of the two-step method for determining the code word.

The different eight-chip code words are created in either of two ways; both ways are essentially the same, but one way uses a direct method and the other uses a two-step method. Before the two methods are explained, a nomenclature common to both methods needs to be introduced. The eight bits that make up a message symbol can be broken down into four groups of two bits each. As an example, if the eight-bit message symbol $(b_0, b_1, b_2, b_3, b_4, b_5, b_6, b_7)$ sent was to be 01110110, then the four groups of two bits $\{(b_0, b_1), (b_2, b_3), (b_4, b_5), (b_6, b_7)\}$ would be $\{(01), (11), (01), (10)\}$. By examining Table 2.2 from the 802.11b standard, the corresponding phase values can be found.

Table 2.2. Bit pattern

Bit pattern $\{d_i, d_{(i+1)}\}$ $\left(\begin{array}{l} d_i \text{ is first} \\ \text{in time} \end{array} \right)$	Phase values φ_1	Phase values $\varphi_2, \varphi_3, \text{ and } \varphi_4$
0 0	0	0
0 1	$\frac{\pi}{2}$	$\frac{\pi}{2}$
1 0	$3\pi/2 \text{ or } \left(-\frac{\pi}{2}\right)$	π
1 1	π	$3\pi/2 \text{ or } \left(-\frac{\pi}{2}\right)$

It can be seen that the phase values for the example octet (01, 11, 01, 10) can be determined as shown in Eq. (1) below.

$$\varphi_1 = 01 = \frac{\pi}{2}; \varphi_2 = 11 = \frac{3\pi}{2}; \varphi_3 = 01 = \frac{\pi}{2}; \varphi_4 = 10 = \pi \quad (1)$$

These four phase values will be used in both the direct and two-step methods that will be presented.

The direct method simply takes the previous four phase values and plugs them directly into an equation to determine what the sent coded bits will be. Even though this is not the conventional way of creating the code word, it is the intuitive way and merits explanation. The four phase values are entered into Eq. (2) below.

$$c = \{e^{j(\varphi_1+\varphi_2+\varphi_3+\varphi_4)}, e^{j(\varphi_1+\varphi_3+\varphi_4)}, e^{j(\varphi_1+\varphi_2+\varphi_4)}, \\ -e^{j(\varphi_1+\varphi_4)}, e^{j(\varphi_1+\varphi_2+\varphi_3)}, e^{j(\varphi_1+\varphi_3)}, -e^{j(\varphi_1+\varphi_2)}, e^{j\varphi_1}\} \quad (2)$$

By solving Eq. (2), the sent complex code word is obtained and can take on the values of $\{1, -1, j, -j\}$. From the previous example, using the input bits 01110110, the following output code word would be $\{-j, 1, -1, j, j, -1, -1, j\}$. As was stated earlier, this is the intuitive method and the one most often used to determine the code word. However, this is not the way the code word is found in modulating the signal in a real-world system. To determine this, the two-step method must be presented.

The two-step method uses the same four phase values as before, but instead of plugging the four phase values directly into the equation, a different phase equation is used. As can be seen from Eq. (2), the first phase value, φ_1 , can be found in each term. Consequently, the term $e^{j\varphi_1}$ can be factored out of the expression and simply used as a multiplier at a later step to rotate all of the terms. When factored out of Eq. (2), it leaves the equation as shown in Eq. (3) below.

$$c = \{e^{j(\varphi_2+\varphi_3+\varphi_4)}, e^{j(\varphi_3+\varphi_4)}, e^{j(\varphi_2+\varphi_4)}, -e^{j(\varphi_4)}, e^{j(\varphi_2+\varphi_3)}, e^{j(\varphi_3)}, -e^{j(\varphi_2)}, 1\}. \quad (3)$$

This equation will output an eight-chip code word, but this code word is not rotated by the φ_1 value. This is the equation that will be used to create the databank of 64 different 8-chip code word sequences previously mentioned. When the bits are actually modulated, only the last six bits of the message symbol are sent to the code word decision block. This means that the second equation must be used because it does not account for the first two bits of the message symbol. The needed code word value could be determined on the fly using Eq. (3), but it is more likely that all 64 different possible code word values will be predetermined and stored in an accessible memory bank. This way, as the six bits are read in, they will correspond to a certain place in memory, and those eight code word chips in memory will be output. Using the same example as shown for the direct method, if the same phase values were used for the new equation, the output chips would be $\{-1, -j, j, 1, 1, j, j, 1\}$. These chips obviously are not equal to the sent chips found in the direct method. This is why the second step of the method is required; if all of the values in the chip sequence above were rotated 90° or a value of $\pi/2$ [i.e., multiply each chip by $e^{j(\pi/2)}$ or by the value $\{j\}$], the resulting chip sequence would be equal to the chip sequence for the direct method, which was $\{-j, 1, -1, j, j, -1, -1, j\}$.

DQPSK modulation is accomplished by using the first two bits of the symbol, which were previously set aside and are used to differentially modulate the code word so that each chip maintains the same phase rotation. This phase rotation can be detected in the receiver and add two more bits of coding. The question might arise why DQPSK modulation should be used instead of QPSK with the direct method of determining the code word because they have already been differentially encoded. The two methods are identical in theory, but they are actually much different when the decoding within the receiver is considered. The main problem occurs in the receiving and decoding of the data. To determine the sent message symbol, the direct method would require a bank of 256 correlators. Correlators tend to make receiver design more complex and thus more expensive. Therefore, in using the two-step method of DQPSK modulation, only 64 correlators would be required along with a phase detector. Although both implementations are theoretically equivalent, the latter has been shown to be more cost-efficient.

2.2.2.3 802.11a

Several different PHY layer specifications have evolved out of IEEE 802.11, with the most prominent being 802.11b, or WiFi, although 802.11a is more appealing in many ways. The frequency range for 802.11a does not lie within the 2.4-GHz ISM band. This is attractive because this ISM band has become overcrowded with ZigBee, 802.11b and 802.11g, Bluetooth, and even microwave ovens, for a few examples. With the congestion comes a need for a system less susceptible to interference; to maintain this robustness, a system must sustain a degradation in efficiency and will experience a decrease in overall data throughput. To increase its output bit rate, 802.11a takes advantage of the 5-GHz unlicensed national information infrastructure (UNII) band.

2.2.2.4 802.11a Physical Layer

The 802.11a scheme has an improvement in output data rates on the order of a five-fold increase over 802.11b. From IEEE Std 802.11a, this PHY layer can support eight different data rates, which are 6, 9, 12, 18, 24, 36, 48, and 54 Mbps, although the mandatory rates are 6, 12, and 24 Mbps [11]. These rates are realized through the use of 52 different subcarriers, as required by the Orthogonal Frequency Division Multiplexing (OFDM) system used by 802.11a. A more detailed look into how the 52 subcarriers are used can be found in a Linksys Group white paper entitled “A Comparison of 802.11a and 802.11b Wireless LAN Standards” [12]. OFDM is implemented as pipelining. For example, instead of transmitting a 24-Mbps data rate on one carrier, the 20-MHz channel is broken down into 52 different subcarriers, with 48 of the 52 each carrying a data rate of 0.5 Mbps in parallel with the others. The other four subcarriers do not carry data and are “pilot” tones [12]. These 52 subcarriers are modulated using BPSK, QPSK, 16-quadrature amplitude modulation (QAM) or 64-QAM, depending on the desired data rate [11]. Sending data using multiple carriers has several advantages over single-carrier modulation that are not within the scope of this report but can be found within the Linksys white paper [12].

Within the 5-GHz UNII band, 802.11a is subdivided into three different channels of 100 MHz each, occupying a total of 300 MHz. This means that 802.11a and 802.11b both have approximately 100 MHz of bandwidth to use. The three different channels for 802.11a and their respective bandwidths are the UNII lower band (5.150 to 5.250 GHz), the UNII middle band (5.250 GHz to 5.350 GHz), and the UNII upper band (5.725 to 5.825 GHz). The need for the three different bands results from 802.11a offering 12 separate subchannels, each with a bandwidth of 20 MHz. This 20-MHz channel is divided into the 52 separate subcarriers. The use of the three bands allows for no overlap of the subchannels; this arrangement is very different from WiFi, where only three channels do not have overlap. A more important reason for the three main UNII bands of 802.11a is to distinguish among different transmission output powers. It allows for a maximum of 40 mW in the lower UNII band. This band is mainly for indoor use, and the lower power can be used because it does not have to span a long distance. For use outdoors, where distances could be greater than they would be indoors, the upper UNII band allows for a maximum output power of 800 mW. For cases in between, such as in a large warehouse or an application that needs to span short distances between an outdoor and indoor transceiver, the middle UNII band allows for a maximum output power of 200 mW. For 802.11a-compliant devices, it is not a requirement that they be able to transmit and receive in all three bands [12].

2.2.2.5 802.11g

The relatively new 802.11g is a later version of WiFi with backward compatibility and capable of maintaining 802.11a-type data rates—up to 54 Mbps. It is essentially another version of 802.11a placed in the ISM band but with a few slight differences. Therefore, the similarities and differences between 802.11g and both 802.11a and WiFi need to be explained.

Because 802.11a uses 100 MHz of bandwidth in the 5-GHz UNII band, and WiFi occupies 82 MHz of bandwidth in the ISM band, the task of translating one to the other seems feasible. The protocol 802.11g specifically does so; it incorporates the same OFDM carrier modulation as 802.11a and can obtain the same data rates. IEEE Std 802.11g also uses a packet binary convolutional code modulation scheme that can be found in the 802.11b standard, which is a different method of coding other than CCK that can achieve higher data rates. It will not be discussed at length in this report.

IEEE Std 802.11g must be backward-compatible with WiFi devices; thus it must be able to operate at the same data rates as WiFi and using the same modulation schemes. This restriction was put in place so that the existing networks using WiFi, mainly the wireless devices placed in laptops and PDAs, would still operate in the new 802.11g environments. This attribute would alleviate the problem of having a full-scale switch from one protocol to the other. It would allow network routers within a building to be switched to 802.11g, while the devices connected through the wireless network could be operated using either of the two protocols. Over time, the WiFi devices will be phased out and replaced with the updated and faster 802.11g devices.

2.3 Bluetooth

Bluetooth is another option that has emerged as a popular choice for wireless connectivity. Bluetooth-compliant devices communicate with the modern PC and this capability opens Bluetooth to a vast array of potential applications because computers are essential in any type of research and in controlling processes such as nuclear reactions and other extremely delicate processes that require precise monitoring. To aid in human control of the computers, there has been a demand for more wireless products, including keyboards, mice, and printers.

2.3.1 Applications

Certain Bluetooth qualities make it desirable for applications that use such computer accompaniments. While Bluetooth does not offer the speed of WiFi, which operates at 11 Mbps compared with 1 Mbps for Bluetooth, 1 Mbps is still well suited for use in conjunction with a computer when it deals with a human interface. Because of the nature of applications such as entering words on a keyboard or scrolling with a mouse, these activities do not require a large amount of bandwidth. Even ZigBee, with a maximum bit rate of 250 kbps has started to be used for some of these devices.

Another desirable quality of Bluetooth (although ZigBee might be shown to have a bigger advantage) is the amount of energy that it consumes. Normally, a Bluetooth mouse has built-in rechargeable batteries with the charger located within the receiver. Whenever the battery runs low, docking the mouse into the receiver unit allows it to charge. A typical charge will last from one to three days depending on usage. A keyboard, on the other hand, is typically powered by off-the-shelf disposable batteries and can be used for anywhere from six months to a year before the batteries must be replaced. This type of keyboard is much more efficient than one that incorporates the WiFi standard, which would have to be plugged into a wall outlet because of its power consumption, thereby defeating the purpose of being wireless.

Because keyboards and mice are typically used within close proximity to a PC, there is no need to consider the range of such devices. Distance does become an issue, however, when devices such as printers or scanners are considered. For instance, printers could be located on the opposite side of a room or even in a completely different room altogether. The distance might or might not be a problem, depending on the conditions of the channel through which the signal must travel. Typically Bluetooth is operational up to a range of about 10 to 100 m (~33 to 328 ft), depending on the selected power level. Bluetooth offers transceivers with a range of different power levels. Power Class 1 transceivers can

transmit with a maximum output power of 20 dBm, allowing a range up to 100 m. Power Class 2 devices can transmit in a range of up to approximately 10 m (~33 ft), with a maximum output power of 4 dBm. The most common applications use Power Class 2 chips. The third and final power-level transceivers, Power Class 3, can transmit a maximum of only 0 dBm and are for very short-range applications, typically 1 m (~3.3 ft) [13]. Therefore, if printers were being used in conjunction with Power Class 1 transceivers, then in a wide-open area, a range of 100 m (~328 ft) could be achieved. This figure would drop drastically, however, within an indoor environment because of walls and objects causing reflections and absorbing power, so the overall range could be considerably less than 100 m (~328 ft). The additional power consumption could be tolerated because it would occur in printers, desktop computers, or laptops, which are typically supplied by wall sources and do not need battery recharge.

Bluetooth is a cable-replacement protocol. For instance, with PDAs and cell phones becoming increasingly more sophisticated, there is a demand to have synchronization between these devices and PCs. For example, it might be desirable to share calendars, e-mail, and files between two devices. If, for example, a PDA were used for repair work or for ordering materials, and wireless Internet were not available, then the user could download information from the PDA to a computer through a wireless Bluetooth connection rather than manually copying everything from one machine to another. Also, when various processes are measured throughout a power plant, the results could be recorded using one device and transferred to an off-site computer at another facility for processing. For instance, if an outside company were using a channel sounder to characterize the fading parameters within a nuclear facility, and by rule no outside devices with memory were allowed into the facility, then secure laptops owned by the nuclear plant could be used to take the measurements, the contents could be analyzed to ensure that no critical information had been compromised, and then the results could be transferred from the laptop at a later time through a Bluetooth connection to the consulting firms' device for data analysis.

2.3.2 Physical Layer

To help determine what is within a Bluetooth transceiver chip that allows for wireless connectivity, a look into the standard itself is required. Much like WiFi and ZigBee, Bluetooth can be dismantled into two separate partitions, the MAC and PHY layers. For the purpose of this study, the PHY layer will be explained in some depth, whereas the MAC layer will be only briefly touched upon. A more extensive look into both layers can be found in the Bluetooth core specification [14].

First, the Bluetooth approach to combating interference is very different from the way both ZigBee and WiFi try to accomplish this task. ZigBee and WiFi use DSSS techniques in which the narrow signal bandwidth is extended into a wider bandwidth, thereby making the chance of an interferer depleting the entire signal minimal. In the time domain, this result is accomplished by taking a signal bit and multiplying it by a PN sequence. If a 63-length PN sequence is used, then for every bit of information, 63 corresponding chips will be sent. At the receiver, therefore, there is a greater chance that many of the 63 chips can be received without error than that a single bit can be received. Bluetooth, on the other hand, uses an FHSS technique. Rather than spreading the entire signal over a portion of the allotted frequency band, Bluetooth keeps the same narrowband signal and changes only the carrier frequency of the transmitted signal, thus hopping from one frequency range to another. This hopping minimizes the likelihood that an interferer will be located on several hop sequences in a row.

Multiple variations of hopping schemes are used for FHSS systems. The main two methods are fast-hop, which incorporates multiple hops per bit, and slow-hop, which sends multiple bits per hop. Bluetooth is of the latter type and transmits one complete packet per frequency hop. A packet can contain anywhere from 126 to 2971 bits [14]. The maximum time duration of any one packet is 625 μ s based on the rate of 1600 hops/s that Bluetooth employs.

In the 2.4-GHz unlicensed ISM frequency band, Bluetooth uses 79 MHz of the 83.5 MHz bandwidth available. This usage allows for 79 channels, which are 1-MHz wide and correspond to the data rate of 1 Mbps, to be used as determined by Eq. (4) [15].

$$(2402 + k) \text{ MHz}, k = 0, 1, 2, \dots, 78. \quad (4)$$

(e.g., 2402, 2403, 2404, 2405, 2406, 2480)

Because of regulations set forth by FCC, Part 15.247, a device may not transmit for longer than 0.4 s on any particular channel within a given 30-s timeframe. This means that at least 75 of the 79 channels must be used in the hopping sequence. The hopping sequence is a predetermined sequence that combines the 79 channels in pseudo-random order [16].

Bluetooth incorporates a modulation scheme similar to the cellular standard for the global system for mobile communication (GSM). GSM uses a Gaussian minimum-shift keying (GMSK) technique, and Bluetooth uses a very similar Gaussian frequency-shift keying (GFSK). Both techniques are modeled from FSK modulation schemes. To better understand GFSK and GMSK, a conceptual foundation in FSK and MSK must be explored.

FSK is used so as to change the frequency of a signal when either a binary one or zero is sent, unlike PSK, which changes the phase of the signal, or amplitude-shift keying (ASK), which changes the amplitude of the outgoing signal. For Bluetooth, a positive frequency deviation from the carrier corresponds to a binary one being transmitted; conversely, a negative frequency deviation is transmitted when a binary zero is sent. The minimum frequency deviation acceptable according to the standard has been set to 115 kHz. FSK has two distinct advantages over both ASK and PSK modulation techniques. First, the additive thermal noise in the receiver directly affects the amplitude and phase messages contained within the signal, whereas the noise will not affect the message of the FSK signal in a direct manner. The second advantage is contained within the complexity that must go into the transmitter and receiver design itself. The varying envelope of the ASK and PSK signals makes the design of the RF circuitry very complex, causing the footprint of the design to be bigger, both of which lead to a higher cost. FSK, on the other hand, is a constant envelope modulation, creating less complexity, a smaller footprint, and a lower cost, all desirable for Bluetooth products.

There are two different kinds of FSK modulators; one is a continuous-phase modulation scheme and the other is discontinuous. The difference comes from the type of transmitter used. If two different oscillators are used, one to send the higher frequency and one to send the lower frequency, then when the signals switch from one oscillator to the other, the phase of the outgoing signal will instantaneously change, causing it to be discontinuous. If a frequency modulator is used in which changes to the signal are accomplished with the aid of an integration technique, then the outgoing phase will be continuous. This leads into MSK, which is a form of continuous-phase FSK.

MSK is very similar to the O-QPSK modulation scheme with half-sine pulse shaping used for ZigBee. MSK is a continuous-phase FSK signal with a minimum modulation index ($h=0.5$) that will produce orthogonal signaling. The similarities between O-QPSK and MSK can be found in *Digital and Analog Communication Systems* [7]. When the sinusoidal pulse shaping of MSK is replaced with Gaussian pulse shaping, the resulting modulation is GMSK. GMSK improves the spectral efficiency of the MSK signal and helps to stabilize the frequency variations over time [17]. When dealing with GMSK, it is useful to know the value of BT , which is defined as the product of the parameters B , the 3-dB baseband bandwidth of the Gaussian filter, and T , the baseband symbol duration [17].

To understand why GFSK and GMSK are similar but not equal, it will be best to look at an example between the parameters of the two wireless standards, Bluetooth and GSM. GSM uses GMSK with a

modulation index $h=0.5$ and a BT product of the Gaussian filter $BT = 0.3$. Bluetooth, on the other hand, uses a modulation index h to be between 0.28 and 0.35 with a BT product of $BT=0.5$. Thus, it can be seen that because the modulation index is less than 0.5, the signaling is not MSK and therefore is considered GFSK. Similarities can be seen between the values of h and BT for Bluetooth and GSM. This is because both systems made tradeoffs between the two variables. As both h and BT increase, the bandwidth also increases. But as h and BT decrease, the eye of the signal becomes wider; and as h decreases, the bit error rate (BER) increases. This means that if it is assumed that the system is modeled as an MSK rather than an FSK system, the result will be an overestimate of the true performance. The results obtained for Bluetooth will be better than if the system were modeled after an FSK system.

2.3.3 MAC Layer

The standard covers the PHY layer, as previously discussed, and the MAC layer, which deals with how the data are organized and sent. Each Bluetooth device is given a 48-bit device address, which is used for authenticity when creating a connection between devices. If two or more devices are trying to communicate with each other, then a piconet is created. A piconet consists of a master device and its accompanying slave devices. There is nothing about a device that makes it either a master or a slave; all devices are built equally and being a master depends only on who initiates the contact. A piconet can contain up to 255 devices, although at a given time only 8 can be active—one master with seven slaves. The reason behind this will be explained later in discussing the different modes of operation. When several piconets are located within the same area and some devices belong to several piconets, a scatternet evolves. Within a scatternet, masters and slaves can belong to different piconets. A master in one piconet can be a slave in another and vice versa. However, a master of one piconet cannot be a master of a different piconet; if that were the case, the two piconets would be considered one big piconet because the timing and hopping sequence is controlled off the master's clock.

The timing within a Bluetooth network is critical because the protocol follows a time division multiple access system in which devices are given certain time increments during which they are allowed to transmit data. The communication is broken down into time slots, and each time slot has a duration of $625 \mu\text{s}$, which derives from 1600 hops/s. The master transmits on the odd time slots, and the slaves are allowed to transmit on the even time slots as determined by the master. Therefore, only one time slot is allotted for each packet, although in some cases a packet might need to transmit more bits than are allowed in a single time slot. In that case, a master can allow a given device to transmit for more than one time slot, as long as the total number of slots used is odd (e.g., 1, 3, 5, and so on). The practice of allowing a packet to last longer than one time slot is used only in pure data transmission [18].

Bluetooth allows for two different types of communication links, synchronous (SCO) and asynchronous connectionless (ACL). SCO links are mainly used for voice transmissions in which an application must have forward and reverse communication at regular intervals on dedicated time slots. SCO packets are rarely transmitted with coding and are never retransmitted because the voice transmission cannot be delayed for retransmission [19].

ACL links are pure data links that can allow for retransmission if a packet is received in error. These packets are coded with a cyclic redundancy check (CRC) or forward error correction (FEC) or both. ACL links also allot time slots as they are needed for certain devices. Rather than having dedicated time slots, as is done with SCO links, it allows a packet to last for more than one time slot. If an SCO link is being used, however, the ACL link must wait for time slots to become available because the real-time nature of the voice SCO link takes priority [19].

Before discussing the details of the FEC and CRC procedures, it might be useful to describe the structure of a general Bluetooth packet. The packet can be broken down into three parts: the access code, the

header, and the payload. The access code comes at the beginning of every packet and consists of 72 bits that are used for synchronization, direct current (dc) offset compensation, and identification [14]. The header consists of 54 bits that are used to determine which device is transmitting, the type of packet being sent, the length of the packet being sent, whether the previous data were successfully received, and the header integrity [14]. Finally, the payload can consist of 0 to 2745 bits that contain the intended data to be transmitted along with any CRC or FEC that was incorporated into the data. A more in-depth detailing of the packet can be found with the Bluetooth standard, but it is outside of the scope of this report [14].

Four basic types of error correction are incorporated into the Bluetooth standard. They might or might not be implemented, depending on the application. Two FEC methods, the 1/3 rate and the 2/3 rate, can be instantiated. The 1/3-rate FEC is a redundancy technique that simply repeats each bit three times, so only a third of the payload of the packet is actually filled with information. The 2/3-rate FEC integrates a (15, 10) shortened Hamming code with a generator polynomial equal to 65 in octal representation [14]. This is equal to taking 10 information bits and generating 15 coded bits that can be sent. The third error correction scheme is an automatic repeat request system in which the packet will simply be retransmitted until an acknowledgement is sent saying that the packet was successfully received. This acknowledgement is incorporated into the packet header as a single bit. If no acknowledgement is sent, then the repetition will continue until a timeout value is reached. Once it is exceeded, the packet will be discarded, and the next packet will be transmitted. The fourth and final error correction system is the CRC. The CRC is composed of 16 bits generated by the CRC-CCITT polynomial 210041, which is given in octal representation [14].

While in a piconet, a device can be in one of four different modes. The first mode is an active mode in which the device is interacting with the other devices within a piconet, sending and receiving data. The next three modes all deal with power saving. The sniff mode uses the most power of the three and occurs when the device listens to the communication traffic within the piconet at a reduced rate but does not transmit data of its own. Another mode, which is a little more power-efficient, is the hold mode. In the hold mode the device neither transmits nor receives data. With this downtime, the device can perform other operations or simply go into a power-saving mode. The fourth mode is the park mode. The only difference between the hold and park modes is that in the park mode the device gives up its member address, whereas in the hold mode it maintains its member address. This is where a piconet can contain either 8 or 255 devices. The member address is a three-bit number, which can maintain up to eight active members of a piconet. Each device is also given an eight-bit inactive member address in which it can stay synchronized with the master unit but does not actively participate. Therefore, if a piconet contains 34 devices, 8 of these devices are in active mode and 26 are in park mode. If the master tells an active device to enter park mode, then the master can signal one of the parked devices to enter into active mode. While in parked mode, the device is in the most power-efficient mode.

3. SIMULATION OF PHYSICAL LAYER SYSTEM MODELS

Why is there a need to simulate various ISM band devices and perform an interference study? Why not just perform testing on-site? When an on-site study is executed, there is no doubt about the results because the tests have been performed and the system tested. Subsequently, the performance of measurements will yield a near 100% accuracy for the stimuli tested. Because the overall goal is to determine whether the three protocols can coexist in an industrial environment, and because the measurements did precisely this (i.e., measured the interference between the different protocols for the assumed parameters), then why is there a need for simulation?

Simulation is needed because the amount of time that it would take to physically collect measurements in all areas for which a coexistence study has been deemed necessary would be daunting. The different possible combinations of transmitters, receivers, and interferers would result in an endless number of measurements to be taken. If a computer program could be created that would be able to take a physical layout of a building and compute the interference for certain combinations of different signals, then an educated guess could be made from these calculations, and only a minimal number of measurements would need to be taken on-site to verify the functionality of the computer simulator. Therefore, modeling the three protocols and simulating their PHY layers could save a large amount of both time and cost.

Wireless InSite™, developed by Remcom, Inc., is modeling software that predicts the propagation path characteristics of electromagnetic waves. It was originally designed for outdoor urban environments, but was later extended to include irregular terrain, foliage, and indoor and indoor/outdoor environments. This study takes advantage of the indoor propagation prediction techniques. The indoor modeling software includes a full three-dimensional vector ray-tracing model that can have various features set for either an imported floor-plan layout or one created through the floor-plan editor provided. These different features can include placement of objects within a room that are made of different materials, and the tool can take into account the material and thickness of walls and whether they are concrete or drywall. It also allows for the inclusion of doors and windows. Wireless Insite can also be used to determine the theoretical energy of wireless signals.

3.1 Monte Carlo Simulation

The next question to arise might be what sort of testing should be conducted. How are the models to be built? The first step is usually to conduct Monte Carlo simulations. Monte Carlo simulations are based on the probability of a random event occurring. As pertaining to a simulation of a communication device, Monte Carlo simulations generally assume rectangular pulse shaping; additive white Gaussian noise (AWGN); or fading channel, independent and equally probable data symbols, and no filtering in the system [20]. Incorporating all of these assumptions into an experiment might seem to oversimplify the simulation, but it is an excellent starting point. Not only does it get to the basis of the simulations and aid a full understanding of the operability of a given protocol, but it also provides a fast and efficient simulation as a result of the no-filtering assumption. It allows the user to identify the results to be compared with known results and to ascertain that the models are being properly created.

3.2 ZigBee Simulation

For ZigBee, the Monte Carlo simulation was set up using the procedure described in the following sections. The procedure can be broken down into three different parts, which involve the transmitter, channel, and receiver. The transmitter generates the sent message bits and then encodes them so that they can be transmitted. The channel adds both the fading and the AWGN to the transmitted signal. The receiver then decodes the sent signal and determines what the sent signal was.

3.2.1 ZigBee Transmitter

The transmitter for the simulation can be broken down into two separate parts, message generation and message encryption or spreading. A block diagram of the transmitter can be found in Figure 3.1. The first part of the transmitter is the message generation portion, and because the basic building block of a ZigBee information packet is a data symbol composed of four message bits, a random variable between 0 and 15 needs to be created. This figure comes from the fact that four message bits can have 16 different possible values. The resulting symbol that has a value of 0 to 15 is then fed into the spreading function of the transmitter.

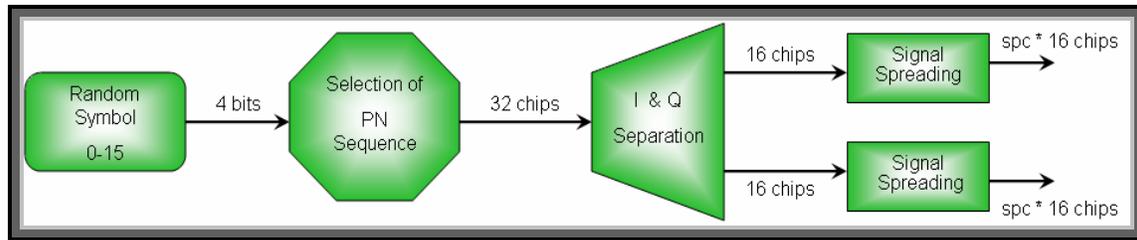


Figure 3.1. ZigBee transmitter.

As shown earlier, ZigBee has its own set of 16 quasi-orthogonal PN sequences, with each sequence corresponding to a symbol between 0 and 15 as determined by its binary equivalents. So once the symbol is passed to the spreading function, the spreader will output the corresponding PN sequence. For instance, if a symbol of value 9 (1001) was sent to the input of the spreader, the 32-chip PN sequence corresponding to the 9 would be output.

Once the proper chip sequence has been selected, it will pass through a few more steps before being ready to be entered into the channel. The next two steps are basically signal conditioning steps and involve simply a change of value and a change of indexing. The first step is to change the chip values from 1s and 0s, as shown in Table 3.1, to values of 1s and -1 s. Values of 1 stay as 1, and values of 0 change to values of -1 . This is done for convention and to refrain from multiplying by a zero because the zero becomes a -1 . This would also be used so that if pulse shaping were later added, the signals of 1s and -1 s could be multiplied by the half-sine pulse shaping, and the correct output would be realized. As opposed to a 1 and 0 representation in which the half-sine pulse shaping was multiplied by the message, the portion corresponding to a sent binary 1 would look like the positive portion of a sine wave; but the portion corresponding to a sent binary zero would look like a flat line at zero rather than a negative portion of a sine wave.

The next step of the signal conditioning is referred to as a change-of-indexing step. Another way of putting it is that the signal would be separated into a direct (I) phase and a Q phase so that it could be QPSK modulated. If the 32-chip sequence was thought to contain place values from 0 to 31, the first value of the chip sequence would be thought to hold the zero place, while the last value of the chip sequence would be thought to hold the 31st place. When the chip sequence was separated into its I- and Q-phase equivalents, the even-indexed chips would be placed in the I phase, and the odd indexed chips would be placed in the Q phase. This can more easily be seen by an example from the 802.15.4 standard. Figure 3.2 shows how the chips would be grouped had the input sequence been 32 chips long, beginning with the chips 1101100 and so forth. If separated as specified, the I phase includes the values $\{1\ 0\ 1\ 0\}$, and the Q phase includes the values $\{1\ 1\ 0\ 1\}$. This figure also shows how half-sine pulse shaping would be implemented. Once the chip sequence is separated into its I and Q phases, the signal is almost ready to be introduced to the channel.

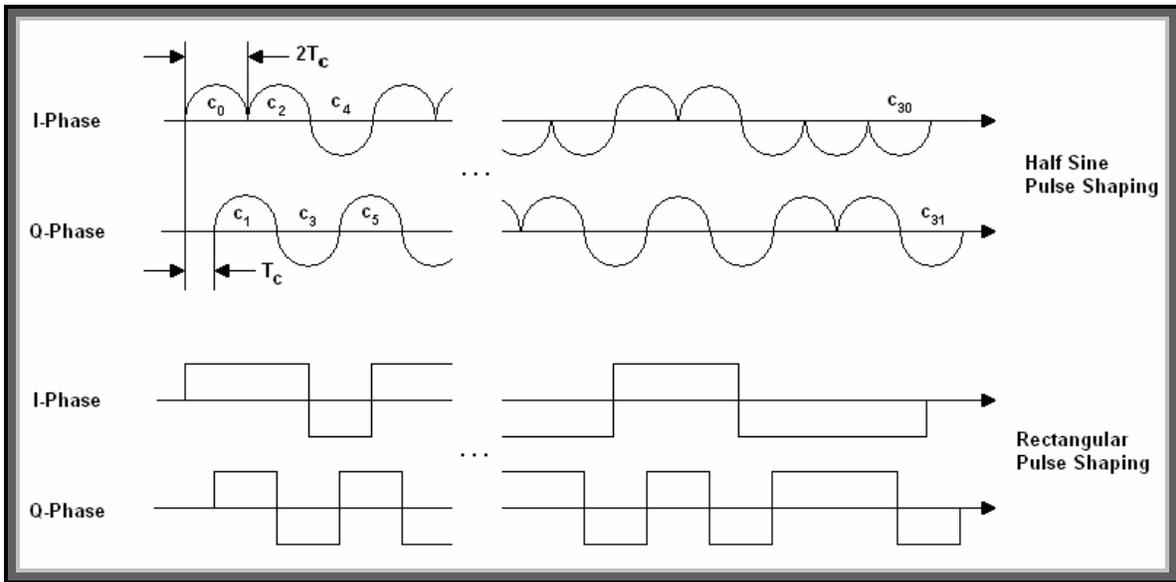


Figure 3.2. Chip sequence separation with half-sine pulse shaping.

3.2.2 ZigBee Receiver

Upon reception of the transmitted signal that has either been faded or had noise added to it, the signal must be down-sampled or, more specifically, integrated over the chip interval. This means that when the chip is up-sampled, if the samples per chip (spc) are equal to two, the two values would be summed for the down-sampling. This will be done for every set of chips values, resulting in a signal half as long as the transmitted signal and equal to the length of the original signal before the up-sampling. The resulting values can span the entire range of positive and negative values. Now that the signal has been down-sampled, it is ready to be entered into the receiver design.

For the ZigBee receiver, which is shown in Figure 3.3, after the summation over the chip interval, the result is a soft decision that can take on any range of values and not just a solid 1 or -1 . This allows for a more reliable result to be obtained from the correlation. It also allows for a value near a strong 1 to have more of an influence on the value's being a 1 rather than a value more near 0, which could have resulted from a -1 having noise added to it.

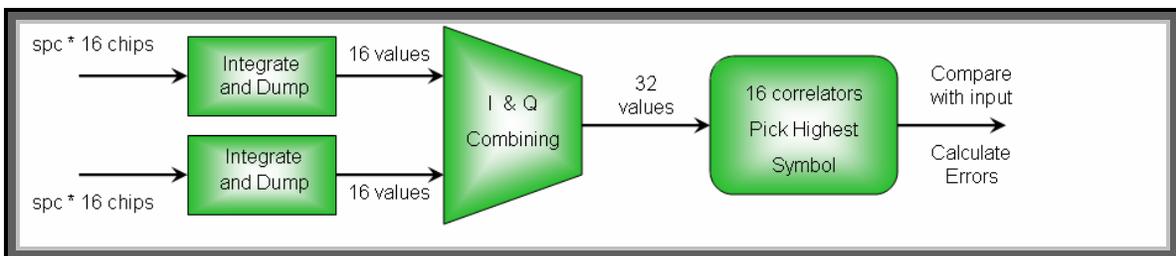


Figure 3.3. ZigBee receiver.

Once the soft decision has been made, the received signal is separated into blocks of 16 chips. There are still separate I and Q phases; therefore, the grouping of 16 chips in each phase will result in being equivalent to the 32 chips in one PN sequence. Therefore, each of the possible PN sequences is grouped into its respective I- and Q-phase representations, as set forth earlier, in which the even chips are placed in

the I phase, and the odd chips are placed in the Q phase. Rather than the PN sequence containing 1s and 0s, they are transformed to contain values of 1 and -1 . Once separated, the I and Q phases of both the received signal and the transmitted PN sequence are correlated with each other. This results in the need for 32 correlators—16 for the I phase and 16 for the Q phase. The correlators consist of multiplying each of the 16 possible transmitted sequences with the received sequence for both the I- and Q-phase sequences. The results will be that if two chips have the same sign, they will add to the correlation, and if the two chips have different signs, then the value will decrease the correlation. Once all 16 values are summed, the highest value will be chosen, corresponding to the PN sequence assumed to have been sent. If the PN sequences for the I and Q phases agree, then that sequence will be chosen as being transmitted and the corresponding symbol obtained. However, if the two phases do not agree on the same sequence, then the two sequences will have their correlation values added from each phase, and the largest value will be chosen. Once the data symbol is found, it is converted to a binary number and compared with the transmitted binary message. For instance, if a symbol 4 {0 1 0 0} was sent, but the received symbol was a 12 {1 1 0 0}, then the result is a total of 1 error because only one bit is different between the two symbols.

3.3 WiFi Simulation

For WiFi, the Monte Carlo simulation is set up using the procedure described in the following sections. The procedure for WiFi can be broken down into the same three parts as that for ZigBee: the transmitter, channel, and receiver. The transmitter generates the sent message bits and then encodes them so that they can be transmitted. The channel adds both the fading and the noise to the transmitted signal. The receiver then decodes the received signal and compares it with the sent signal to obtain the number of errors.

3.3.1 WiFi Transmitter

The transmission technique for WiFi varies from the transmission of ZigBee, shown in Figure 3.4, even though both claim to use a DSSS system. Both contain the signal generation and the system encoding, but there is a slight difference in the former and a big difference in the latter. ZigBee has a chip rate eight times greater than the bit rate, while for WiFi, the chip rate is equal to the bit rate.

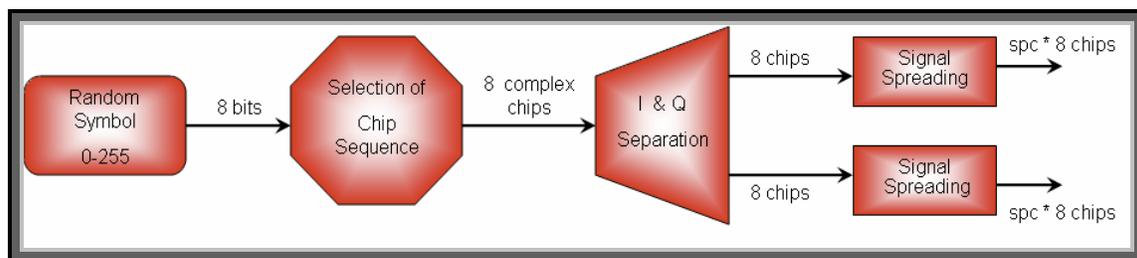


Figure 3.4. WiFi transmitter.

The first step for the WiFi transmitter is the generation of a random signal. Because WiFi uses a base symbol size of eight bits, it is convenient to create eight random variables that have given values of 1 or 0 and place them into one symbol. The next task for the transmitter is to encode the data bits. This process takes three steps to complete and consists of breaking the eight data bits into the four phase values, determining the code word, and then doing some signal conditioning to make the signal easier to transmit. The first step is fairly straightforward and is determined in the background information; it takes the eight data bits and groups them into four sets of two. Each group of two bits represents a phase value.

The following step to encoding the data involves the use of Eq. (1). When using WiFi, Eq. (2) is normally used; but because correlators have not yet been introduced into the system and because the computation can be minimized by not taking the extra step of creating the code word and then rotating the code word,

it can all be done at once using Eq. (1) without harming the integrity of the simulation at all. The effect is to create eight coded chips from the eight data bits.

The eight coded chips can have the values $\{1, -1, j, -j\}$. To represent these values in I and Q representation, three symbols are needed, 1, -1, and 0, but it is desirable to require only two symbols, 1 and -1. To make this transformation, the coded chip values must first be rotated and then scaled. This is accomplished by adding a rotation of $\pi/4$ to the coded chip values. This results in values of the form $\{\pm 0.707 \pm j0.707\}$, which would work when trying to transmit either a 0.707 or a -0.707 because they are two distinct values; but the convention is to use a 1 or -1 so that the values are simply swapped, 1 for 0.707 and -1 for -0.707. This is a valid switch because the increase in signal power is accounted for when the noise variance is calculated.

To allow for greater efficiency, the simulation execution runtime needs to be sped up; therefore, the above transmitter process is performed in only one single instance at the very beginning of the simulation. Every possible eight-bit input is fed into the process, and the resulting eight-chip outputs are stored in a lookup. This allows for a possibility of 256 different combinations. Therefore, all that is needed during the execution of the simulation is for a random variable to be chosen from the values of 0 to 255 and, using that value to scan through the lookup table, for the corresponding eight-chip sequence to be output. This saves time because it is not necessary to group the bits, encode them, and then rotate them.

The only step left before transmission is to up-sample the chips to produce one or more copies of each chip, which is not necessary for a Monte Carlo simulation. This step would have a bigger impact had pulse shaping been used, and therefore it is incorporated so that the simulation can be built upon at a later time.

3.3.2 WiFi Receiver

Phase detection is the method used within the decoder found in Figure 3.5. But before going into phase detection, the receiver must down-sample the received signal by accounting for the up-sampling of the chips. This is done by adding the adjacent chips that were copies of each other and making the decision that if the sum is positive, a 1 was sent, but if the sum is negative, then a zero was sent.

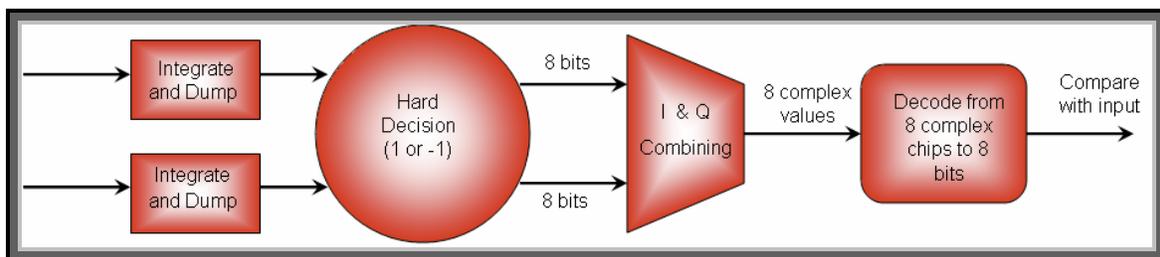


Figure 3.5. WiFi receiver.

Next the two phases, I and Q, are added back together by making the I phase the real part and the Q phase the imaginary part of the received code word. This results in a code word that has values of $\{\pm 1 \pm j1\}$, which are not equal to the desired values, which are $\{1, -1, j, -j\}$. Therefore, the values must be rotated back by a value of $\pi/4$. In this simulation, the rotation is done by assignment. The program checks to see if the real part is a 1. If it is, then it checks for the value of the imaginary part. If the imaginary part is 1, then the chip value is a 1; otherwise, the chip value is a -j. On the other hand, if the real part was a -1, then the routine would check to see if the imaginary part was a 1. If it was, then the chip value would be a j; otherwise, the chip value would be set to be -1. Once completed for all the chips, the simulation must

obtain the phases. This is done by inverting the equation for the chip value of the eighth chip in the code word. As can be seen, this chip value contains only the value of phase 1. From this value, similar inverting can be done to the fourth, sixth, and seventh chips in the code word to determine their phases.

Now that the phases are known, which can have values of $\{0, \pi/2, \pi, 3\pi/2\}$, the corresponding binary values can be found from Table 3.2. The binary values can then be concatenated, and the resulting eight-bit received data bits will then be known. The data bits can be compared with the original sent data bits, and a count can be kept for the number of errors that occur. Much as with ZigBee, the WiFi simulation is conducted for different values of signal-to-noise ratio (SNR), and the corresponding number of errors is tabulated. Once completed, the number of bits in error is divided by the total number of bits sent, and the result is plotted versus SNR.

3.4 Bluetooth Simulation

In the final simulation model for Bluetooth, the Monte Carlo simulation is created in a manner very similar to those used for the previous two models. For Bluetooth, however, there are some assumptions that need to be made before a model can be created that can be joined with the previous two models so that an interference study can be conducted. There are two main groups of assumptions that must be made. One deals with the fact that Bluetooth is based on an FSK modulation scheme, which must be dealt with in the transmitter; the second stems from the FHSS nature of Bluetooth and will be dealt with within the interference study when the models must be joined.

3.4.1 Bluetooth Transmitter

The first assumption that can be made is that the GFSK system of Bluetooth will be modeled as a regular FSK system because in a Monte Carlo simulation, pulse shaping is not considered. This means that the Gaussian pulse shapes can be disregarded and not considered. This assumption goes along with the other two models' assumptions in that their pulse-shaping techniques were not included, such as the half-sine pulse shaping of ZigBee. The main reason not to consider the Gaussian filter to shape the outgoing signal is to decrease the runtime of the model. To include pulse shaping, each bit must be spread to resolve a Gaussian pulse out of it, which in turn increases the amount of time needed for the simulation to run by a factor of however many samples are taken—upward of 50 to 100 times.

The second major consideration that must be taken into account is that the other two simulations are broken down into their I- and Q-phase components. To ease the combining of all of these models, therefore, it would be helpful for them all to be similarly broken down, so it is logical to create an FSK model that will use I and Q components rather than frequency components. One way to do this is to model an FSK system using an MSK system. GMSK and GFSK are very similar modulation techniques. If the assumption is made that h , the modulation index for Bluetooth, is equal to 0.5, then the model becomes a GMSK system, whereby after pulse shaping has been discarded, it becomes an MSK system. An MSK signal can be represented using I and Q components. Therefore, the assumption will be made that Bluetooth's modulation can be approximated using MSK techniques.

To create a transmitter for Bluetooth modeled after an MSK system involves having a data source to create a data stream consisting of 1s and -1s. For MSK, the I and Q components are created in much the same way as for ZigBee. The even-indexed bits are placed in the I phase, and the odd-indexed chips are placed in the Q phase, with the first bit indexed as the 0th bit. Because Bluetooth uses FHSS instead of DSSS, the created data bits are the sent bits; they are not spread using any type of PN sequence. Now that the data are in the I and Q phases, they can be sampled, so there could be more than one sample per chip. This occurrence will become more prevalent when the models for the interference study are combined.

The Bluetooth transmitter can be found in Figure 3.6. Note that there is no type of PN sequence involved because of the FHSS system.

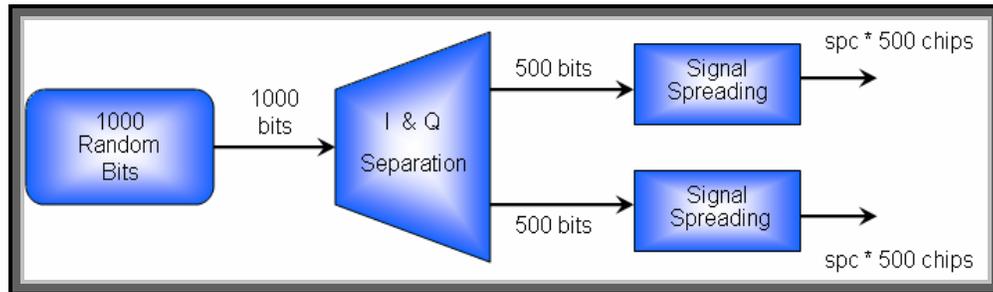


Figure 3.6. Bluetooth transmitter.

3.4.2 Bluetooth Receiver

The receiver for the Bluetooth model will be modeled to perform much like the integrate-and-dump receiver design shown in the block diagram of Figure 3.7. The incoming data are summed over the length of a bit period, whether it consists of a single chip or a number of chips, depending on the value of spc . After being summed, if the value is greater than zero, a value of one is assumed to have been sent; if the value is less than zero, then a binary zero is assumed to have been sent. This is done in both the I and Q branches of the receiver, and once completed, the results are intermixed in the same way that they were separated and become one long data stream. This data stream is compared with the sent data, and the number of errors is determined. Exactly as in the previous two models, the Bluetooth simulation is conducted for different values of SNR, and the corresponding number of errors is tabulated.

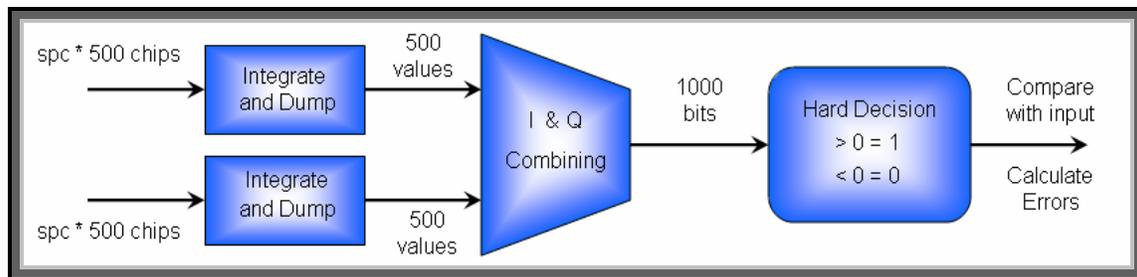


Figure 3.7. Bluetooth receiver

3.5 Channel

Now that the transmitted signal has been created, it is ready to be passed through the wireless channel model. The model will consist of two separate parts, the fading model and the AWGN model. The fading model is used only when fading is considered and is not used for the ideal channel case in which only AWGN is considered. Both cases will be considered within the results portion of this report.

3.5.1 AWGN

The AWGN model is perhaps the trickiest portion of the whole simulation in terms of subtlety. It is something that does not draw a lot of attention to itself, but it can be difficult to implement, mainly because of the scaling factor. The noise is created from two separate noise generation sources—one source for the I phase and another source for the Q phase. The two phases must have independent sources; otherwise, the AWGN assumption might not hold true anymore because the two phases could become

cross-coupled. Therefore, for each phase, the procedure described below must be followed. The noise consists of generating a random variable with a zero mean and a standard deviation and variance both equal to one. The random variable is then scaled by the noise variance.

The noise variance is dependent upon the amplitude of the signal, which for the Monte Carlo simulation for ZigBee is $\sqrt{2}$; the number of chips per bit, which is equal to eight (due to four bits being represented by 32 chips); the number of spc, which is arbitrarily set to two (more important for pulse shaping, which does not change the simulation because as there are more spc, the noise variance also increases); and the SNR. The main variable is the SNR because as it increases, the noise variance decreases, making the noise level lower.

Once the noise variance is multiplied by the random variable, the result is then added to a single chip. For example, if the random variable was 0.25 and the noise variance was 0.4, then the noise value would be equal to {0.1}. If, for instance, the transmitted chip was a 1, then the noise added to the signal would yield a value of 1.1. As seen later, this will result in the chip not being in error. If, for example, the next chip was sent, and this time the new random variable had a value of -2.8 (because a new random variable needs to be created for each chip) and the noise variance was still 0.4 (the noise variance changes only for a new SNR), then the resulting noise value would be equal to {-1.12}. If that next chip was also a 1, then when the noise was added to the chip, the result would be a chip with the value of {-0.12} that will be detected as a 0, and an error will have occurred.

3.5.2 Fading

To improve the accuracy of the system model for the on-site floor plan design, fading must be incorporated into the channel. There are two assumed channel types that are widely used to model indoor wireless channels. In the first case, there is a line-of-sight (LOS) path between the transmitter and the receiver. In this situation, a Rician fading envelope is assumed. In the second case, an object is obstructing the LOS path between the transmitter and receiver, such as a wall, desk, pipes, or a person walking. In each case, there becomes no direct path between the transmitter and receiver, creating an NLOS scenario. This scenario is often modeled as a Rayleigh fading distribution over the wireless channel. These distributions will be explored further, and their similarities and differences will become apparent.

There are two different varieties of fading—large and small scale. Large-scale fading occurs as a result of signal attenuation caused by path loss. It fluctuates as the transmitter is moved away from the receiver at distances much greater than a few wavelengths of the signal, and it is the average power of the signal over a local distance. Small-scale fading, on the other hand, is not caused by path loss because small-scale fading is the instantaneous fluctuations in the power of the signal. Because over short distances, on the order of wavelengths, the path loss will remain approximately constant, there must be some other phenomenon affecting the signal.

The phenomenon can be explained in two parts. The first part of the explanation deals with the phase of the signal. When two signals intersect in space, the overall effect of the resulting signal can range from the summation of the two signals to the subtraction of the two signals. This is because two signals that are out of phase with each other will add destructively, and conversely, two signals that are in phase with each other will add constructively. The second part can be explained by imagining any number of signals combining at a single point. The overall result will be one single value, as happens within a receiver, and this value will change as the environment in which the signal is propagating changes. However, because our environment is assumed to be static to an extent, this value will not vary a great deal; therefore, the receiver is thought to be receiving signals in a ray form. The ray form is composed of numerous signals

added to give the overall effect of one signal. Because this ray is formed from a wide variety of signals, its value will fluctuate as if it were the receiver point. The simulation assumes ten of these rays are being received. Each of the ten rays will be Rayleigh faded unless there is a LOS ray, and in that case the LOS ray will not fluctuate and will maintain a constant value.

3.5.3 Ricean

The first fading channel to be studied is the Ricean distribution, in which a LOS path is present and unfaded. In this situation, the individual paths are not Ricean faded. Ricean fading is a collection of the entire ten paths and signifies the distribution after which the overall effect can best be modeled. Ricean fading is a collection of a number of Rayleigh-faded paths that can be superimposed upon a signal with a constant value. This constant value can be used to determine the Ricean K factor, which signifies the energy contained within the dominant path with respect to the spectral energy of the secondary paths; however, for the present simulation, the K factor is not used to produce the fading because the theoretical energy of each signal has already been determined through the use of Wireless InSite. For this simulation, the nine Rayleigh-faded rays are added to one ray, which is given a normalized energy value. The K factor can be used to explain the connection between Ricean and Rayleigh fading.

3.5.4 Rayleigh

Rayleigh fading is a special case of Ricean fading, meaning that when the Ricean K factor is equal to 0, Ricean fading degenerates to Rayleigh fading. This is another way of saying that the stationary dc value corresponding to the LOS path is no longer present. For the simulation, this means that the ten rays are all Rayleigh faded, and none are composed of a normalized energy value.

The Rayleigh fading simulator is shown in Figure 3.8. The diagram is a frequency domain implementation for creating a Rayleigh envelope. The Rayleigh envelope consists of creating two independent random Gaussian sequences. The sequences are then filtered through a Jakes Doppler spectrum. The spectrum follows a bathtub curve and is represented as the square root of Eq. (5) [17] as follows:

$$S_{E_z}(f) = \frac{1.5}{\pi f_m \sqrt{1 - \left(\frac{f - f_c}{f_m}\right)^2}}, \quad (5)$$

where f_c is the carrier frequency, set as 2.4 GHz, and f_m is the Doppler frequency. In this case, the Doppler frequency is approximately 14 Hz. The Doppler frequency is proportional to the velocity of objects in the path of propagation. For the indoor environment, the maximum speeds that need to be considered are those of people walking, which are less than 4 mph, or 6 ft/s.

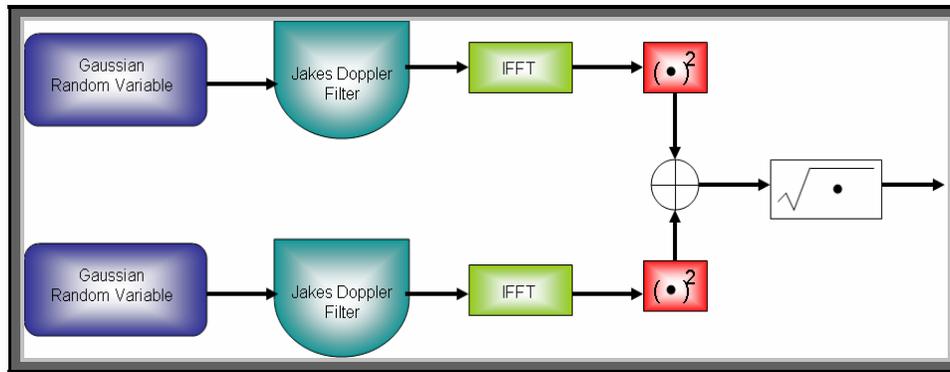


Figure 3.8. Rayleigh simulator.

After the random variables have been filtered, an inverse fast Fourier transform is performed on the two separate sequences. By adding the squares of each of the two signals and then taking the square root of the result, a real signal will be generated. This signal can be multiplied by the transmitted signal to create a Rayleigh-faded signal. Because of the two phases of the transmitted signal, the direct and quadrature phases, a separate Rayleigh envelope must be generated for each phase.

3.5.5 Simulation Design—Power

The first step in determining the values of the ten path gains to be used in the simulation is to take the ten values of received power from Wireless InSite and normalize the path gains to the greatest value of path gain, normally the first arriving path. In doing this, the first arriving path obtains a value equal to 0 dB, and the subsequent values of the nine other paths will all be less than 0 dB. This vector-of-ten path gain value is then used to scale the value of the Rayleigh channel. For instance, the first path gain of 0 dB means a value of 1 in linear units. Therefore, multiplying the simulated Rayleigh-faded signal by the path gain of the first path will result in a multiplication by a value of 1, so there will be no change. All of the subsequent paths, however, will have path gains that are less than one, so the value of the resulting Rayleigh-faded paths will be less than that of the simulated Rayleigh signal.

3.5.6 Simulation Design—Phase

Nearly as important as the power in each path, and maybe even more important considering that the power is normalized, is the value of the phase of the signal. The phase signifies the phase of the sine wave with which the signal is being transmitted, referring to the phase of the message signal, not that of the carrier. The phase is important because two signals with different phases can potentially cancel each other out if they are 180° out of phase with equal power.

The phase is dealt with in a manner similar to that in which the power was handled. Because the first arriving path has its power normalized to 0 dB or a value of 1, the phase of that signal is also normalized to 0° because that is where the phase detector will lock on and compare the other signals with the phase of the first arriving path. Therefore, because the first path will have the phase of its signal subtracted from it to have a phase equal to 0, all subsequent paths will also have the phase value of the first path subtracted from them.

To translate the phase from degrees to a viable number, the cosine of the phase must be taken. The resulting values, when multiplied by the normalized power values, yield the effect of the multipath. For instance, because the power of the first path is normalized to 1 and the phase has a value of 0, which yields a value of $\cos(0) = 1$, meaning that the power multiplied by the phase is $1(1) = 1$, the first path will

always have a value equal to 1. Subsequent paths will have values that fall in the range of 1 down to -1 , meaning that some of the paths could take away information from the message. One example would be if the second path had a power value of -3 dB, which refers to a value of 0.5 , and if the phase were 135° , which yields $\cos(135) = -0.707$, then the resulting effect of the second path would be $(0.5)(-0.707) = -0.35$, which means it diminishes the normalized signal by a value of $1/3$.

3.5.7 Simulation Design—Delay

A third consideration for improving the accuracy of the simulation is also provided through the output of Wireless InSite; this is the time delay of each signal relative to the time of the signal being transmitted. The delays of the signals are important because if the relative delays of two signals are significant compared with the chip interval, then the two signals could add with chips being intermixed. A chip from one signal will not be superimposed with itself in the second signal; it will combine with the chip that was received in that timeframe. In this situation being presented, ISI is not the merging of pulses within a single path; it is the combining of pulses from different paths. Depending on the delay, the distortion of the symbol caused by ISI will change. If the delay is on the order of half-chip periods, then the distortion will be very noticeable. If the distortion is a multiple of a chip period, then the distortion might not be noticeable, but information could be lost as a result of the canceling of signals.

For this simulation, because the relative delays were not considerable when the chip periods were considered, they were not implemented in the simulation. To accurately model the signals with their delays, instead of the chip period's being broken into two sections or samples per chip, the number of slices would need to be on the order of nanoseconds. For the protocols being simulated, this corresponds to taking anywhere from 100 samples per chip for WiFi and up to 1000 samples per chip for Bluetooth. Because the runtime of the simulation is directly proportional to the number of samples per chip, this would result in the runtime of each simulation being increased by a magnitude of 10^2 or 10^3 , corresponding to additional days of computer processing. Because the delays calculated for the three separate protocols are not significant compared with the chip period, the distortion caused by the addition of the ISI would not affect the overall signal in a substantial way. This means that the increase in accuracy of the simulation would change only slightly, while the complexity of the system coupled with the computer processing required for the simulations would increase exponentially. Because the tradeoffs do not offset each other, the delay does not warrant being included and thus will be left out of the simulations; but the simulations can be modified later.

To illustrate the assertion that incorporating the delay into the simulation would have a diminishing return—considering the amount of accuracy gained through increased complexity versus the escalation in runtime caused by the number of samples needed per chip—transmitter files were chosen for each protocol under different propagating characteristics. The propagation parameters will include a LOS case when the propagation is in the range of only 1 m (~ 3.3 ft), an NLOS in which the wave must travel on the order of 5 m, and finally a case in which the NLOS travel is stretched farther and must propagate a distance exceeding 15 m (~ 49 ft).

3.5.8 Bluetooth Delay

Because Bluetooth has the longest chip period, $1\mu\text{s}$, corresponding to the chip rate of 1 Mchip/s, it will be the first scenario considered. For Bluetooth, the LOS situation will be observed to determine how the reflections affect the direct path and whether the shifting of the information bits because of the delay is significant enough that it should be modeled within the system. Parameters and similar data were generated using Wireless InSite and are shown in Table 3.1. From the table, it can be calculated that the first path, the LOS path, contains 90% of the total power found within all ten paths. Therefore, a

significant amount of the energy contained within the signal is received at the instant the first path is detected, which does not indicate a need to delay each given path. Also, upon inspection of the delays of each path found in Table 3.1, all signals are received within 38.1 ns of the first arriving signal, which corresponds to all signals being received within the first 4% of the chip period. Figure 3.9 demonstrates that characteristics of the propagation path due to the delay approximate a rectangle, which is the assumed shape used within the simulations.

Table 3.1. Bluetooth delay

Path #	Power (dBm)	Delay (ns)	Normalized power	Normalized Delay	% Path Power of Total Power	% Delay of Total Chip Period
1	-37.35	5.22	1	0.00	89.85	0
2	-51.47	8.85	0.039	3.60	3.50	0.36
3	-53.56	12.24	0.024	7.00	2.16	0.7
4	-54.22	14.32	0.021	9.10	1.89	0.91
5	-57.06	16.62	0.011	11.40	0.99	1.14
6	-59.19	21.92	0.007	16.70	0.63	1.67
7	-62.43	23.05	0.004	17.80	0.36	1.78
8	-61.62	29.05	0.003	23.80	0.27	2.38
9	-63.9	29.92	0.002	24.70	0.18	2.47
10	-65.1	43.31	0.002	38.10	0.18	3.81

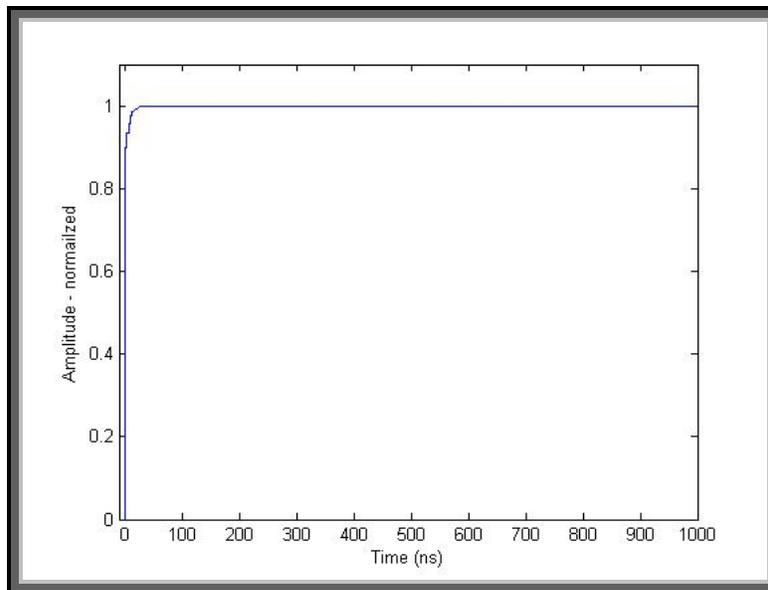


Figure 3.9. Bluetooth delay amplitude.

3.5.9 ZigBee Delay

With a chip rate of 2 Mchip/s, double that of Bluetooth, ZigBee has an equivalent chip period of 500 ns. The ZigBee transmitter and receiver combination was chosen so that there is an NLOS scenario, in which the signal will have to propagate through a wall with a minimum path distance of 5 m (~16 ft) between

the transmitter and receiver. As can be seen in Figure 3.10, the results are very similar to those for Bluetooth. The shape approximates a rectangle with minor imperfections caused by both the power and delay of the secondary paths, two through ten. In looking at Table 3.2, which shows the delay corresponding to ZigBee, it can be observed that 93% of the total energy is contained within the first five paths, and that these five paths occur within the first 1.12% of the chip period, or the first 5.6 ns. The total delay spread is only 27.7 ns, which relates to only 5.6% of the chip period, making the ZigBee delay not worth being included in the simulation.

Table 3.2. ZigBee delay

Path #	Power (dBm)	Delay (ns)	Normalized power	Normalized Delay	% Path Power of Total Power	% Delay of Total Chip Period
1	-47.84	12.94	1	0.00	73.58	0
2	-57.79	16.06	0.101	3.10	7.43	0.62
3	-62.16	17.03	0.037	4.10	2.72	0.82
4	-58.36	17.55	0.089	4.60	6.55	0.92
5	-61.59	18.58	0.042	5.60	3.09	1.12
6	-62.49	23.27	0.034	10.30	2.50	2.06
7	-67.03	24.12	0.012	11.20	0.88	2.24
8	-67.34	25.15	0.011	12.20	0.81	2.44
9	-64.37	29.14	0.022	16.20	1.62	3.24
10	-67.32	40.66	0.011	27.70	0.81	5.54

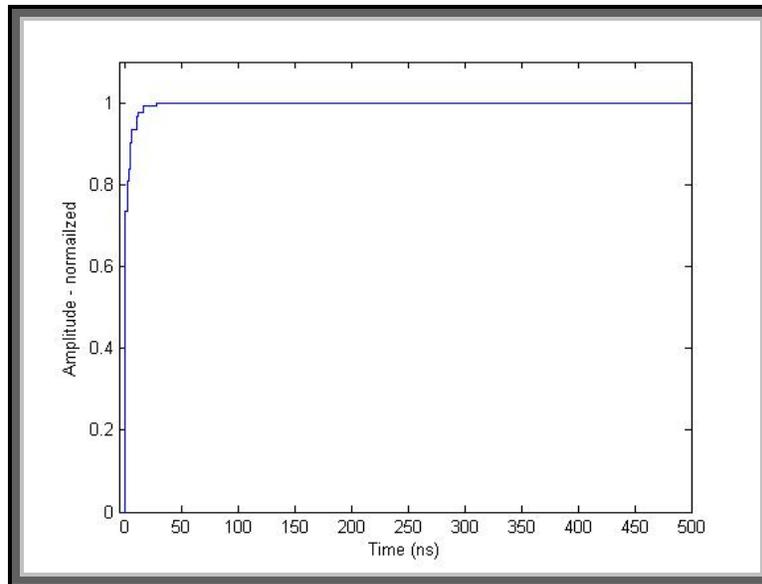


Figure 3.10. ZigBee delay amplitude.

3.5.10 WiFi Delay

Compared with Zigbee and Bluetooth, WiFi has a much shorter chip duration of only 90 ns, making it more than five times faster than ZigBee. The longer duration is caused by the 11 Mchip/s assumed from

the 802.11b standard. It means that the delay associated with each path will have a greater effect on the overall signal. Coupled with the fact that the propagation path implemented for WiFi was for the signal to undergo several transmissions and/or reflections through or off walls during its travel, the distance is more than 15 m, increasing the delay spread of the signal. As can be seen from both the figure representing the waveform caused by the delay in Figure 3.11 and the numerical form in Figure 3.3, the

overall signal is affected much more by the delay than in the previous two examples. The question that remains, therefore, is whether the deformation of the signal is significant enough to merit replicating it within the simulation. From Table 3.3, it can be seen that the last signal arrives at 26.8 ns, which is a third of the total chip period; but this path contributes less than half of 1% to the overall signal energy. The first five arriving paths account for more than 92% of the total energy of the signal and occur within the first 5% of the chip period, or 4.4ns. Because the paths approximate a rectangle, the delay can be disregarded.

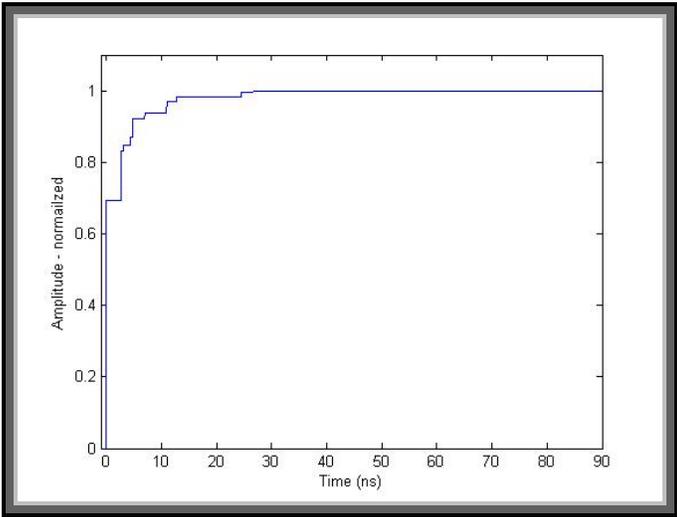


Figure 3.11. WiFi delay amplitude.

Table 3.3. WiFi delay

Path #	Power (dBm)	Delay (ns)	Normalized power	Normalized delay	Percent path power of total power	Percent delay of total chip period
1	-34.69	18.08	1	0.00	69.11	0.00
2	-41.64	20.81	0.202	2.70	13.96	3.00
3	-51.33	21.2	0.022	3.10	1.52	3.44
4	-49.36	22.47	0.034	4.40	2.35	4.88
5	-45.98	22.86	0.074	4.80	5.11	5.33
6	-51.33	25.08	0.022	7.00	1.52	7.77
7	-48	29.06	0.047	11.00	3.25	12.21
8	-52	30.83	0.019	12.80	1.31	14.21
9	-51.58	42.61	0.021	24.50	1.45	27.19
10	-56.76	44.86	0.006	26.80	0.41	29.74

3.6 Interference

The interference between the three different protocols was modeled in a block fashion. The first step was to take the transmitter portions of the three different models and place them into one big transmitter file. Then, by adding the three separate receivers to the one big transmitter file, three different models were created: (1) ZigBee as the intended signal and WiFi and Bluetooth as the interferers; (2) WiFi as the intended signal and the other two as interferers; and (3) Bluetooth as the intended signal and the other two as interferers. Once these three models had been developed, another interfering transmitter model was added to them, each one duplicating the model of the intended signal for the specific model. (For example, in the ZigBee model, another ZigBee transmitter was added that would act as a ZigBee interferer, and the same held true for both the WiFi with a WiFi interferer and a Bluetooth model with a Bluetooth interferer.) Once the foundation for an interference model had been constructed, the model could be manipulated so that it would simulate a real-world environment showing how the interference would take place.

3.6.1 Site-Specific Channel Model

For the site-specific model, the channel is formulated from Figure 3.12. This figure shows the direct-phase branch of the transmitted signal. An identical method is used to construct the fading channel for the quadrature phase of the transmitted signal. The flow of the diagram is that each individual path is faded and scaled, and then the total signal is scaled. For the NLOS situation, all ten faded paths are created in a similar way. The first step is to pass the transmitted signal through the Rayleigh simulator depicted in Figure 3.8. This results in a path consisting of values greater than zero, showing that the path is only increased or decreased in amplitude and has not been negated in any way. The Rayleigh-faded signal is then scaled by a factor of both the path gain and the path phase. The path gain will have a value between 0 and 1; therefore, the signal is not losing any information, and only the amplitude is decreased. The path phase, on the other hand, can have a value between 1 and -1, making it possible for information to be taken away from the overall signal. After each path is scaled to the properties calculated from Wireless InSite, the ten paths are summed and divided by the total energy contained within the signal. The total energy comes from the multiplication of the path gain and path phase. Because some of the values could be negative, the absolute value of multiplication is needed because energy can only have a positive value. Even though the effect of a negative signal is that it takes information away from the signal, the energy contained within the signal must be added to the overall signal.

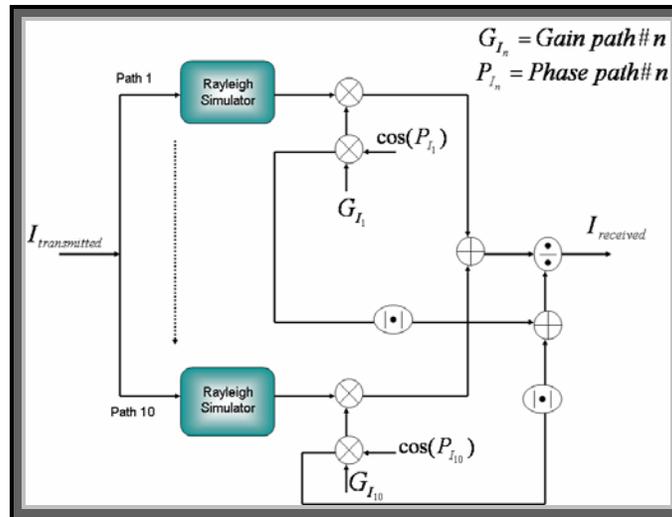


Figure 3.12. Channel path simulator.

For the LOS case within the site-specific channel model, the only variance from the NLOS model shown in Figure 3.12 is that instead of the first path being Rayleigh faded, the block is passed over. The Rayleigh simulator is not present because the LOS path is not faded because of the absence of scatterers in its path. This phenomenon is the cause of the difference between Ricean and Rayleigh fading for the LOS and NLOS cases.

3.6.2 Chip Rate

To make the interference as realistic as possible, the chip rate had to be considered when combining the three different protocols because ZigBee has a chip rate of 2 Mchip/s, WiFi a chip rate of 11 Mchip/s, and Bluetooth a chip rate of 1 Mchip/s. The relative chip rates of the three protocols are modeled in Figure 3.13 along with their respective overlaps. Therefore, if ZigBee is considered to be the intended signal, then during 1 s, 2 million ZigBee chips will be sent, 11 million WiFi chips, and only 1 million Bluetooth chips. This correlates to one ZigBee chip being sent during the amount of time it would take for

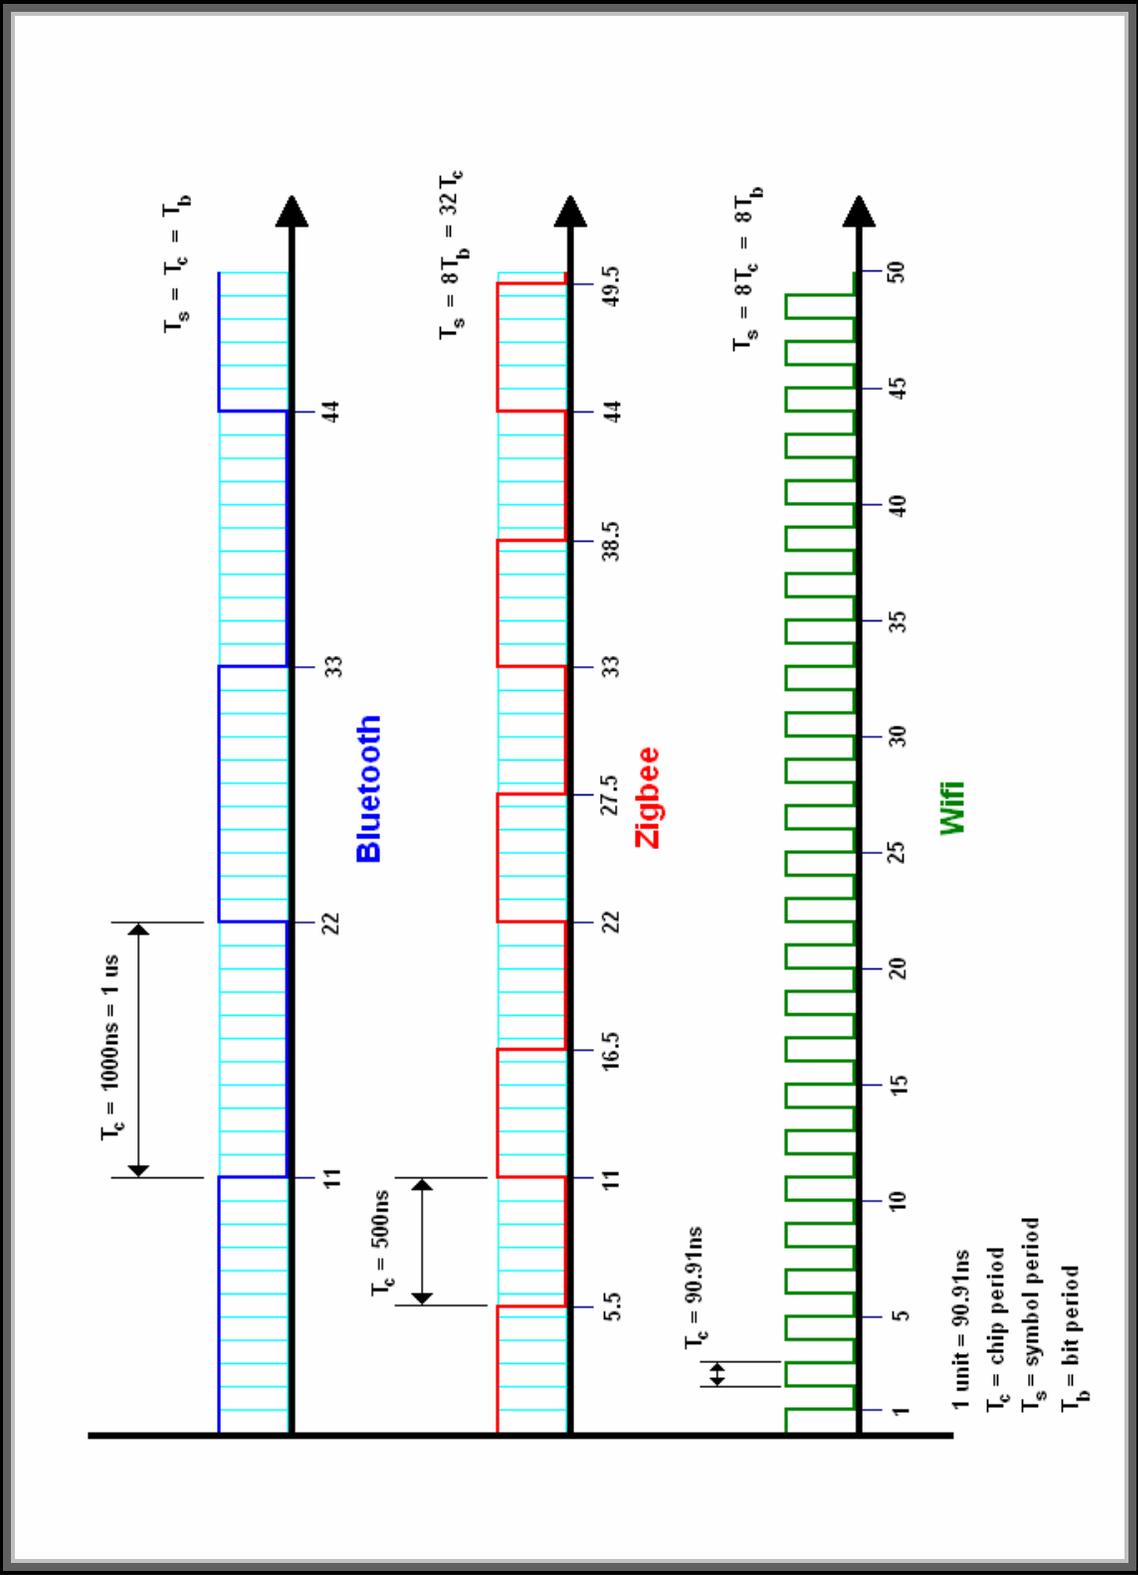


Figure 3.13. Chip cycles.

five and a half WiFi chips or one-half of a Bluetooth chip to be transmitted. To combat this discrepancy, the following courses of action were taken.

It is fairly straightforward to account for the Bluetooth interferer because only half of a chip is causing interference. The other half of the chip would be occurring during the next chip, meaning that if the Bluetooth signal is stretched by a factor of two, then the two signals would match up perfectly. It is necessary only to spread the Bluetooth signal by copying each chip. The WiFi interferer, however, is troublesome. Instead of trying to combine only five and a half chips, as long as ZigBee is sampled at an even rate (i.e., the number of samples per chip is an even number), then 11 WiFi chips could be combined and averaged, allowing that number to be added to two of the ZigBee chips. Thus for every two ZigBee chips being sent, 11 WiFi chips would also be sent. The interference caused by the WiFi signal would be, therefore, the average value of the WiFi signal over those 11 chips. This average value would be added to the intended ZigBee signal that is sent. The effect of the interfering signal would depend on its power level and the resistance of ZigBee to the interference through its spread spectrum system.

To conduct the WiFi interference as it pertains to the signals' chip rates, the signals are manipulated in much the same way as for ZigBee. This time, Bluetooth is spread to cover a value of 11 chips; therefore, instead of one chip being sampled once, it is sampled a total of 11 times. This number is proportional to the number of samples per chip specified in the WiFi program, typically two samples per chip, meaning one Bluetooth chip will be sampled 22 times. ZigBee creates a problem because two is not a factor of 11. Therefore, each sample of a ZigBee chip is spread to cover 11 chips; consequently, WiFi must then have an even number of samples per chip so that, when combined, 11 ZigBee chips will affect the two WiFi chips, which is the correct proportion.

Perhaps the only straightforward interference model can be found in Bluetooth, in which ZigBee is summed over two chips and WiFi summed over 11 chips and then averaged and added to the Bluetooth signal. The only other interferers to be accounted for in these simulations are the transmitters for the cases when the interferer was of the same type as the intended signal. Therefore, no chip-rate modifications needed to be applied to the signal; they were assumed to be at the same chip rate, so the signals were already matched.

3.6.3 Bandwidth

The second means in which the models must be manipulated to better model actual interference deals with bandwidth and just how much of a signal can be considered to interfere within the same amount of space as the intended signal. In the ISM band, there is a total of 83.5 MHz of bandwidth; of that space, ZigBee occupies sixteen 5-MHz channels; WiFi occupies 3 non-overlapping, 22-MHz-wide channels; and Bluetooth uses 79 channels, each with a bandwidth of 1 MHz. The total 83.5 MHz—along with the channel assignments for ZigBee, Bluetooth, and WiFi—can be found in Figure 3.14.

Starting again with the ZigBee model, assuming that each signal occupies only 2 MHz in the 5-MHz channel, if a Bluetooth signal were introduced that occupied only 1 MHz, then within the 80-MHz range there would be a $2/80$ or $1/40$ chance that the Bluetooth signal would fall into the same given bandwidth as ZigBee. So on average, Bluetooth would interfere in a total of 1 in 40 hops, allowing the equivalent Bluetooth signal to be scaled down by a factor of 40. The same is true for the WiFi signal; on average, the WiFi signal would occupy approximately $1/4$ of the allotted bandwidth and would interfere with the ZigBee signal on average $1/4$ of the time. Therefore, on average, the WiFi signal could be scaled back by a factor of four before being added to the ZigBee signal. Finally, the ZigBee interferer must also be modified. Because ZigBee has 16 channels, there would be a 1 in 16 chance that the interferer would be located on the same channel as the intended ZigBee signal; therefore, ZigBee could be scaled by a factor of 16.

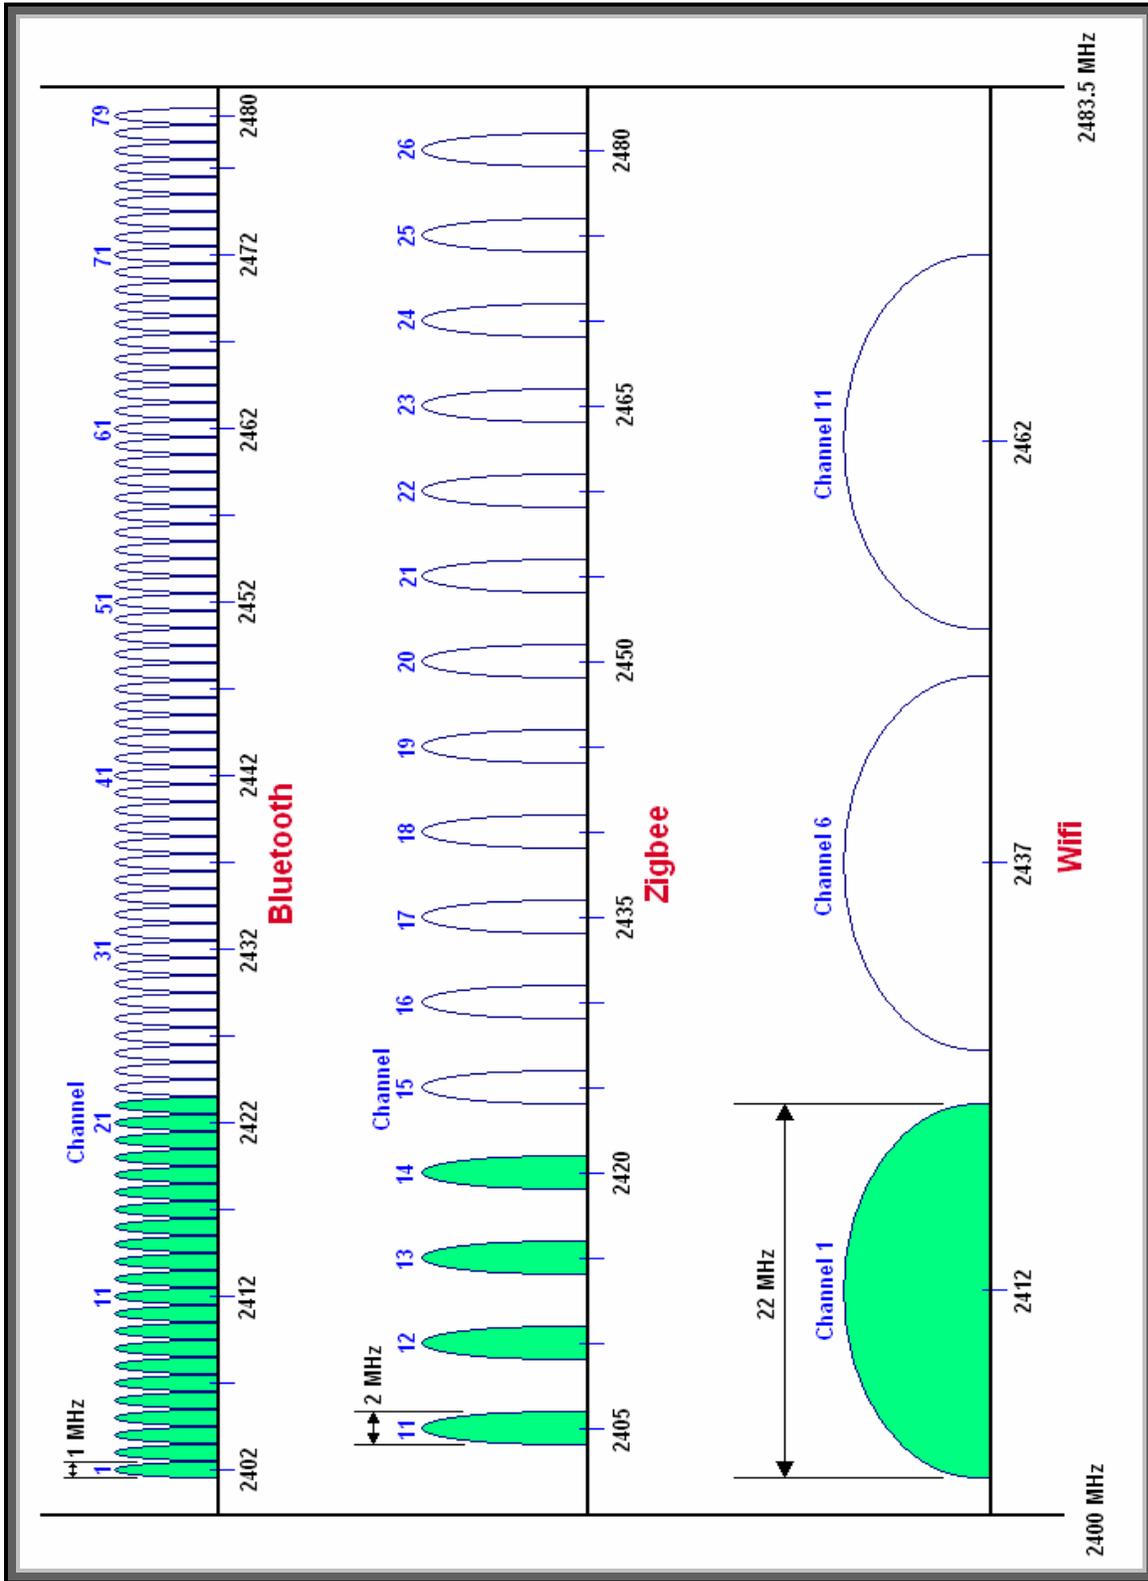


Figure 3.14. Channel assignments.

The same manipulations to the interferers can be applied for the other two models as for the ZigBee model. For instance, the WiFi signal would receive interference a fourth of the time from the ZigBee signal and a third of the time from the other WiFi interfering signal. Because WiFi covers 22 MHz, Bluetooth would interfere approximately 22 times within its 79-hop sequence, allowing for roughly a scale factor of 22/80, on average, to be multiplied by the Bluetooth interferer. The case for the Bluetooth model can be derived in a similar way. The WiFi would need to be scaled down by a factor of 22/80 and the ZigBee interferer by a factor of 1/40; and the Bluetooth interfering signal, assuming they are not on the same hopping sequence, would interfere approximately 1 in 80 times. Now that the two main modifications have been made, accounting for both the chip rate and bandwidths of the signals, the interference models are complete.

The results obtained within Sections 4 and 5 are valid only for the assumptions discussed above and might contradict results from previous coexistence studies. This is because the assumptions intend to consider the performance of the system as a whole rather than on a case-by-case basis. The case-by-case basis would involve simulating the performance when two devices are located within the same channel bandwidth. Even in trying to simulate this case, both devices would not always be transmitting at the same time; consequently, these results would be valid only under a different set of assumptions, unless a more complex system were developed.

The system being modeled considers on-average interference. Bluetooth is inherently a burst error interferer, meaning that Bluetooth will cause a block of errors when the interferer is located on the same channel as the system. As Bluetooth hops to another channel, the system will not detect any errors. Over a short time window, the errors depend upon whether the interferer is co-channel located. As the reference window is expanded, however, there will still be burst errors, but over time a statistical average can be taken. This average is modeled as causing errors over the entire range of frequencies rather than only during certain channels. The same averaging is done to the DSSS systems for both ZigBee and WiFi. The devices will interfere with each other only when they are located within the same channel space because the band-pass filter on the front end of the receiver will reject the out-of-band frequencies not within the receiving channel. In considering a device with an interferer that might or might not be located within the same channel bandwidth, the effects of the interference are averaged over a substantial time period, so they can be modeled as a statistical average affecting the interferers' signal amplitude.

4. RESULTS FROM GENERAL CHANNEL MODELS

Simulations were performed for the three interference models, one for each protocol, using different parameters and under separate conditions. The first case is to carry out the simulations with a general AWGN channel. This entails varying both the signal's SNR and its signal-to-interference ratio (SIR). Each protocol produces three separate BER plots, one for each interferer. Note that BER values worse (higher) than 10^{-3} are not acceptable for any application and this is noted on the plots as the BER threshold. However, some BER requirements are application specific. For example, for video applications the BER value can not be worse (higher) than 10^{-5} . In addition, the SNR requirement for a particular BER is also application specific. For an AWGN channel, generally if the SNR requirement for a certain BER value goes above 15 dB, that is considered as not acceptable. For a fading channel, that value can be as high as 30 dB. This case will aid in determining what the signal's average power range would need to be for the receiver to be able to detect a usable signal. Before the AWGN channels results are presented, the coding gain found within both ZigBee and WiFi will be shown to illustrate why the theoretical curves for both QPSK and DQPSK modulation are different from the curves for ZigBee and WiFi, respectively.

The second section of the results takes the AWGN system models and incorporates the Rayleigh-fading model. The same resulting BER plots are determined as before for the pure AWGN case. The fading plots are constructed through a single-pulse Rayleigh-fading simulator, which then creates flat, slow-fading signals for both the desired signal and for the power of the interferer, unlike the results in Section 5, which are obtained by using ten faded paths.

4.1 Coding Gain

ZigBee and WiFi both incorporate spreading techniques with intrinsic coding gains. These are separate from the processing gains found within DSSS systems. The coding gain helps to combat the effects of a noisy environment, whereas processing gain is used to minimize the effects of interference. Therefore, in the AWGN channel, the coding gain can be found from the performance of the system against the theoretical probability of error for a given modulation technique. The processing gain, however, does not change the performance in an AWGN channel because the noise is proportional to the number of chips per bit. As the code is spread by higher factors, the noise variance is also increased. Figure 4.1 shows the performance of ZigBee versus QPSK modulation and WiFi versus DQPSK modulation; the performance of the two protocols results in a decrease in BER from the theoretical calculation. This exhibits the effects that coding gain has on the overall system. The improvement for ZigBee is calculated as 2.5 dB at a BER of 10^{-5} , illustrating that the performance of ZigBee at 10^{-5} is approximately 2.5 dB better than a QPSK system. At the same 10^{-5} threshold, WiFi has an SNR that is 1 dB less than that of DQPSK for the same BER; therefore, CCK incorporates 1 dB of coding gain into WiFi.

Since the establishment of ZigBee, the FCC has suspended the 10-dB requirement of processing gain for devices operating within the 2.4-GHz ISM band. Processing gain is defined as the chip rate divided by the data or bit rate. The ZigBee signal is spread by a factor of eight, going from four information bits to 32 chips, resulting in a processing gain of 9 dB. The CCK characteristics of WiFi inflict an 11-dB total processing gain [21, 22]. Bluetooth, on the other hand, does not contain any spreading of the signal; it is a completely narrowband signal. So from an instantaneous standpoint, Bluetooth has 0 dB processing gain; but if taken statistically over an entire set of hopping sequences, the effects of one particular channel will not be considered for any of the other 78 channels. This results in a processing gain, which can be thought of as a hopping gain, of 19 dB.

4.2 Interference over an AWGN Channel

The results for the AWGN channel will be presented with the ZigBee protocol the first to be subjected to the interferers. The interferers will be introduced individually, first the Bluetooth interferer, then a WiFi interferer, and then a ZigBee device being interfered with by another ZigBee device. The resulting BER curves for the other two protocols, WiFi followed by Bluetooth, will be introduced to interferers in the same order as for ZigBee.

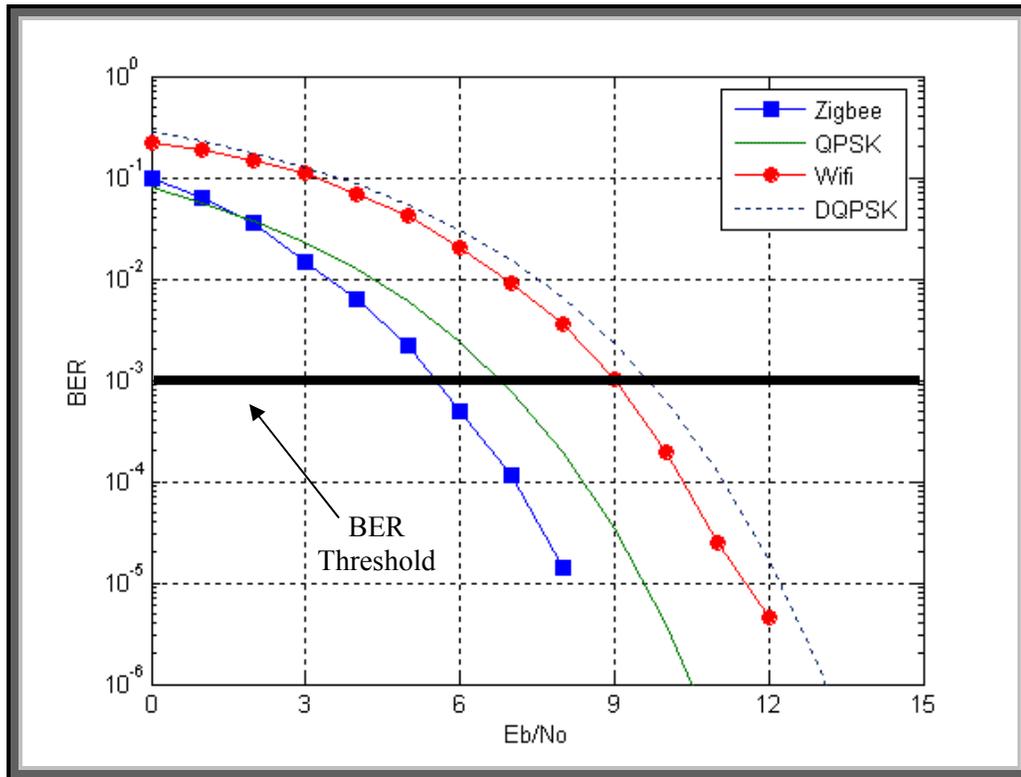


Figure 4.1. Coding gain for ZigBee and WiFi.

4.2.1 ZigBee

The first model simulated is the ZigBee model, and the first interferer considered is a Bluetooth interferer. The results of this simulation can be found in Figure 4.2. It can be seen that for values greater than -16.75 dB, meaning that the ZigBee signal is 16.75 dB below the Bluetooth signal, the system appears to be interference-limited. Here, by ‘interference-limited’ it is meant that the BER value cannot be improved by increasing the SNR because the sum of all interference causes so much signal distortion that the receiver noise is not relevant. This causes the BER curve to flatten out, and the probability of an error will never decrease. For values of SIR greater than -16.5 dB, the system turns into a noise-limited system, meaning that as the SNR increases, the system’s probability of error will decrease and approach zero. This means that the noise is what is hurting the system; the interfering signal level is low enough that it does not affect the system as much when the SIR increases. To maintain a BER of at least 10^{-3} , the SNR must be above 20 dB, and the interfering signal cannot have a value greater than 16.75 dB above the ZigBee signal power. To maintain results typical of a system with no interferers, such as the AWGN case, ZigBee’s SIR power cannot be lower than 10 dB.

In the next case considered, a WiFi interferer replaces the Bluetooth interferer. These results can be found in Figure 4.3. This case with a WiFi interferer appears to be more interference-limited than was the previous case because there is a distinct interference-limited level at least down to a probability of error of 10^{-5} . The SIR at this point is -10.5 dB; if the WiFi power level is any greater than 10.5 dB above the ZigBee signal power, the BER will flatten out at a probability of error greater than 10^{-5} . To achieve a noise-limited system with a BER lower than this, the SIR must be greater than -10.5 dB. The model does not achieve AWGN-type results until the SIR is at least -5 dB, which is the point at which the interference does not affect the signal compared with the noise.

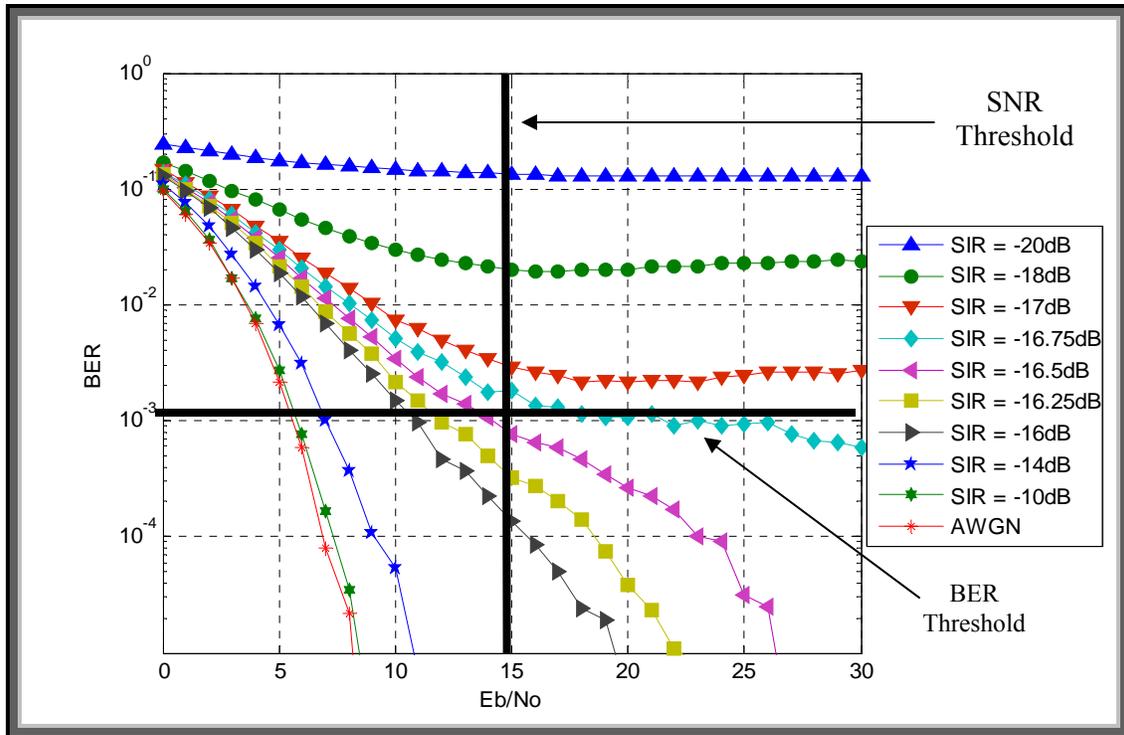


Figure 4.2. ZigBee signal with a Bluetooth interferer-AWGN channel.

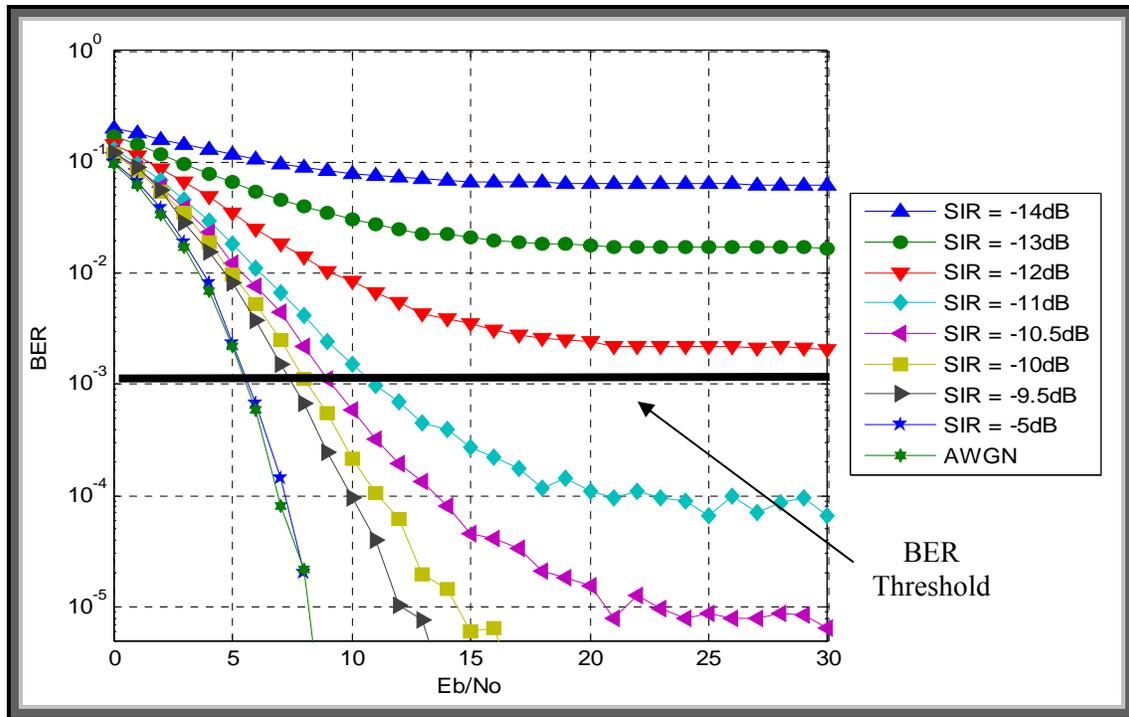


Figure 4.3. ZigBee signal with a WiFi interferer–AWGN channel.

The third and final scenario for the ZigBee signal is a second ZigBee signal interfering with the first. If the two devices were located on the same network, then the ZigBee MAC layer would thwart the interference through a listen-before-talk system known as “carrier sense multiple access with collision avoidance” (CSMA/CA). However, assuming the two devices are not part of the same network and had CSMA/CA turned off, then the results shown in Figure 4.4 are obtained when the signal is sent through an AWGN channel. Unlike in the previous two examples, when a ZigBee interferer is introduced to the ZigBee signal, the result is no longer an interference-limited system; it becomes a noise-limited system until the interference completely overwhelms the intended signal. The system does not level off at a given probability as in the previous two cases. Either the system’s BER approaches zero errors as the SNR increases, or it hovers near 10^{-1} , which occurs when the SIR falls below -12 dB. The BER curve then decreases, approaching the AWGN situation, which occurs when the SIR is above -5 dB.

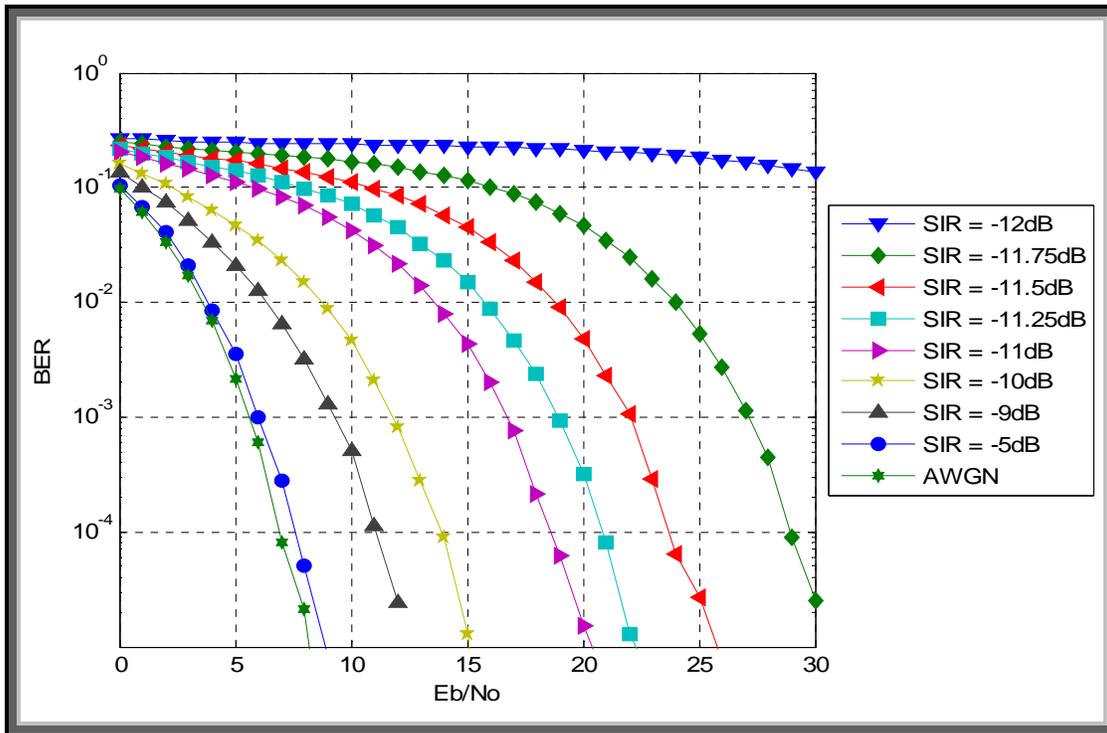


Figure 4.4. ZigBee signal with a ZigBee interferer–AWGN channel.

4.2.2 WiFi

The next simulation is the WiFi model through an AWGN channel to help determine what types of SIR values can be tolerated within this system for the three different types of interferers. Starting with Bluetooth as the initial interferer, the results are examined in Figure 4.5. Much as when the a ZigBee interfered with another ZigBee signal, it appears that the results in Figure 4.5 are also noise-limited only up to the point that an interfering signal completely cancels out the intended WiFi signal. The signal is completely destroyed with an SIR of -6 dB. The BER rapidly improves and approaches the AWGN case as the SIR becomes 0 dB, but this improvement slows after it reaches 0 dB and does not fully reach the AWGN case until the power of the interferer is more than 5 dB below that of the intended WiFi signal.

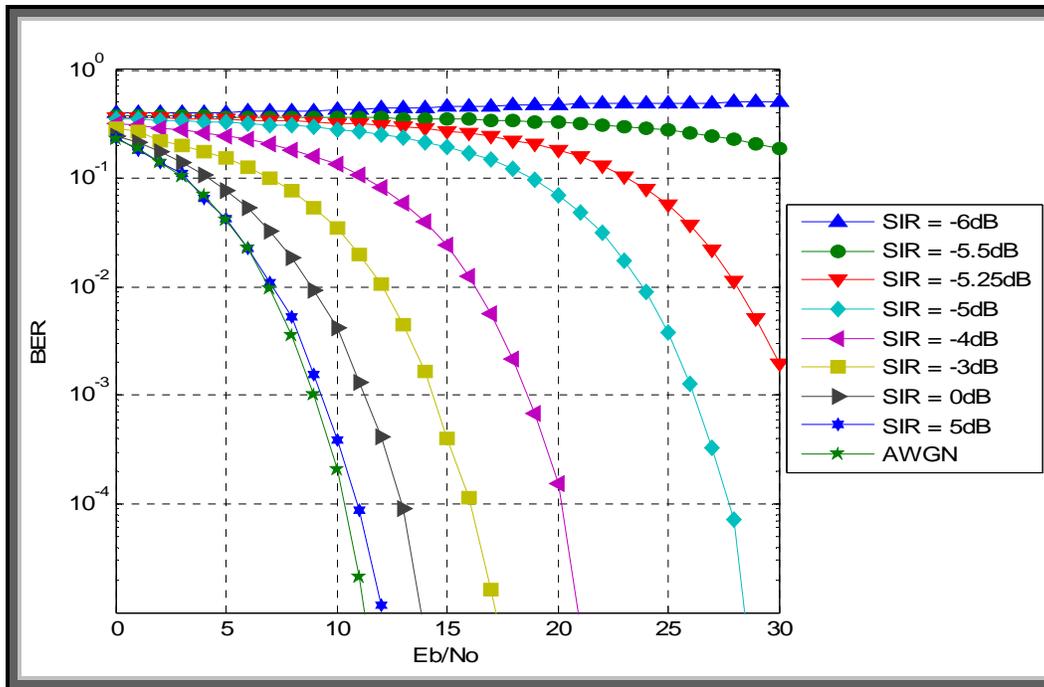


Figure 4.5. WiFi signal with a Bluetooth interferer – AWGN channel.

Introducing the interferers in the same order as for the scenario in which ZigBee was the intended signal, the results of the case in which a WiFi signal is interfering with another WiFi signal can be found in Figure 4.6. As for ZigBee, this can happen only when both devices are located on separate networks, and even then, WiFi uses a CSMA/CA system to avoid interference between WiFi devices. The figure is almost a mirror image of the case in which a Bluetooth interferer was considered. It appears to be noise-limited up to the point at which the interference completely floods the intended signal. For this case, the signal is completely overtaken when the SIR is less than -5 dB, and the BER improves rapidly until it reaches 0 dB. It then continues to improve slowly until 5 dB, at which point it is approximately equal to the AWGN case in which no interference is included.

The last interferer to be implemented with a WiFi signal is the ZigBee interferer. Figure 4.7 is very close to being a replica of Figure 4.6, incorporating even the same values of SIR. The threshold for not being able to boost the SNR so that a signal can be detected is when an SIR of less than -5 dB. For all values of SIR above that, the system is noise-limited. Much as in the previous example, the BER improves rapidly from an SIR of -5 to 0 dB. Then the improvement slows until 5 dB, at which point it reaches the AWGN curve. WiFi appears to be able to resist the effects of interference better than ZigBee, meaning that the system never really becomes interference-limited; but WiFi cannot tolerate a value of SIR as high as the value ZigBee tolerates because of the coding gain of ZigBee.

4.2.3 Bluetooth

The last simulated protocol is the Bluetooth model, which will help illustrate the effects that an interfering signal has on an FHSS system. From Figure 4.8 it can be seen that the effects are very similar to those of the previous two systems, except that Bluetooth has more resilience to interference than does WiFi. Figure 4.8 shows the effects of a Bluetooth signal with a Bluetooth interferer. For this scenario, the system is noise-limited for values of SIR above -19 dB. The effects of the interferer are diminished once the value of the interferer is less than 8 dB higher than the signal's power. There is an extreme

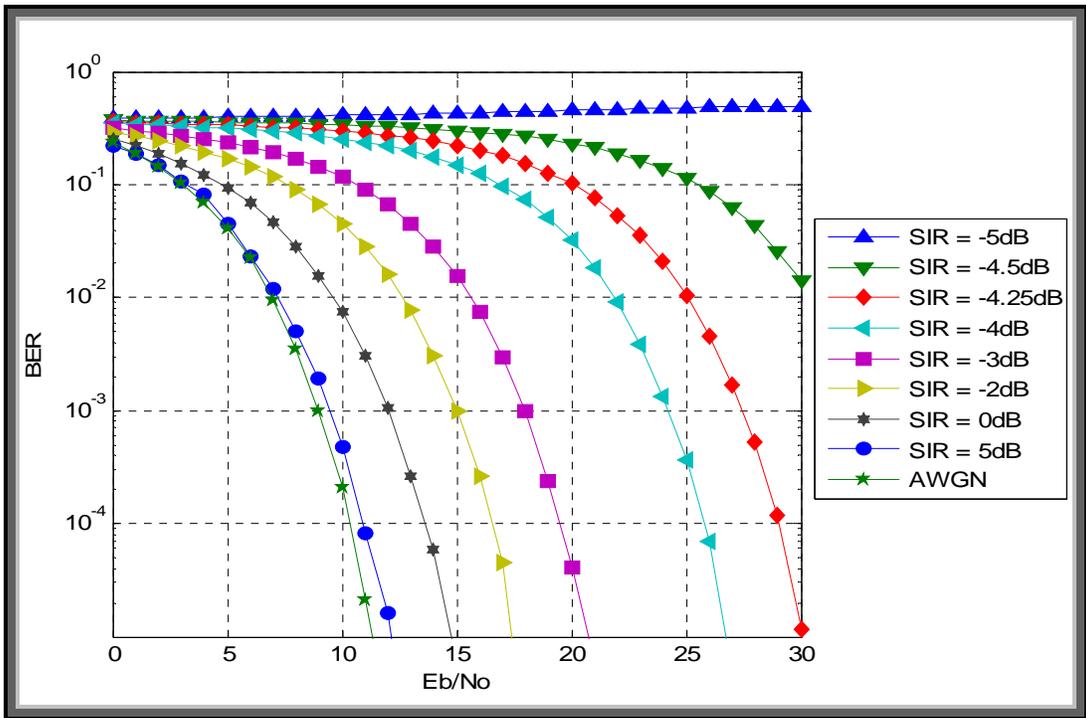


Figure 4.6. WiFi signal with a WiFi interferer-AWGN channel.

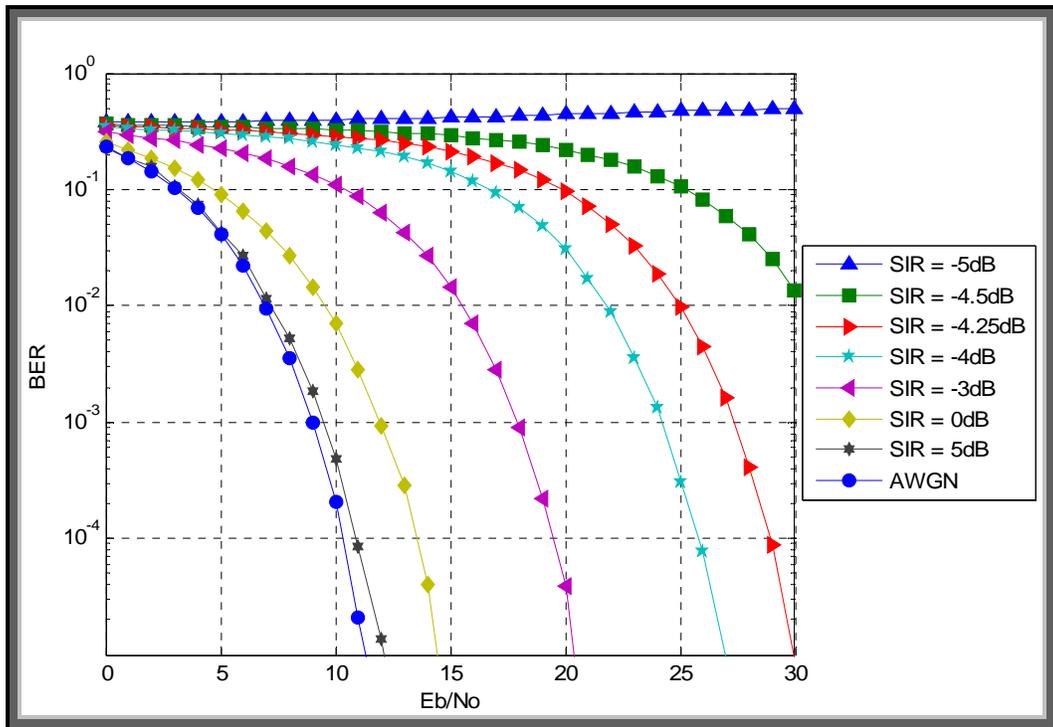


Figure 4.7. WiFi signal with a ZigBee interferer-AWGN channel.

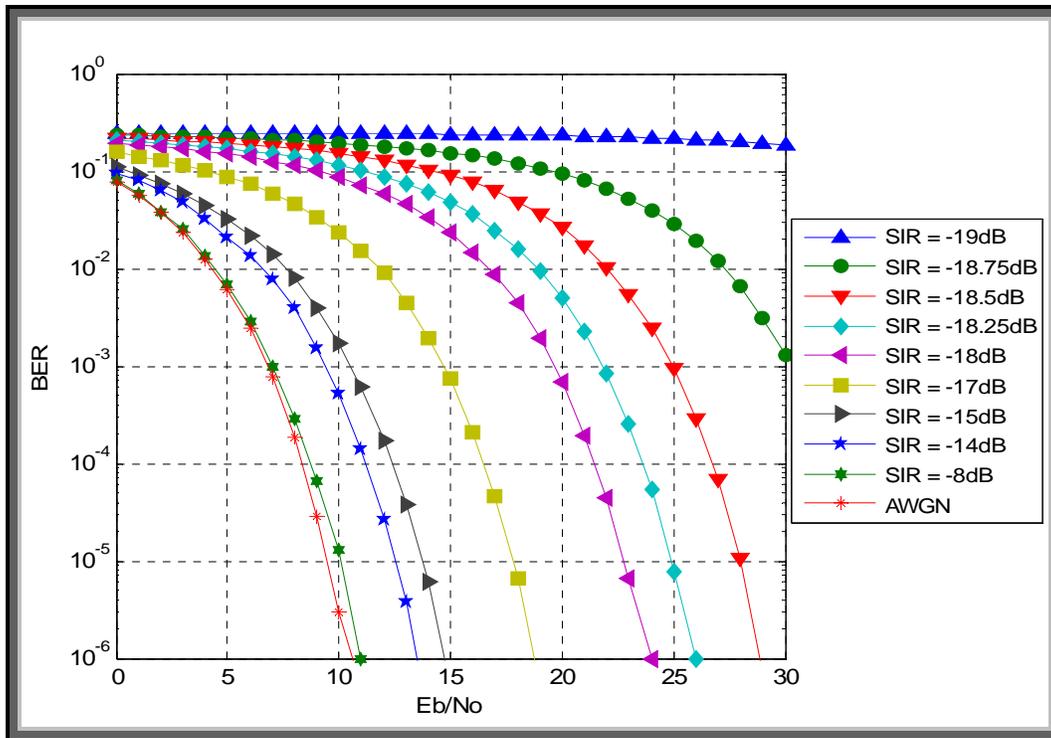


Figure 4.8. Bluetooth signal with a Bluetooth interferer–AWGN channel.

improvement in the BER over this 11-dB range, where it sharply falls off until it reaches -15 dB and then the improvement starts to slow down.

When a WiFi interferer is placed with the Bluetooth signal, the results more resemble the interference-limited systems that are found with ZigBee. The BER values level off for different values of a given SIR value, as can be seen in Figure 4.9. When the Bluetooth signal is more than 10 dB below the interfering WiFi’s signal power, the Bluetooth signal is unrecoverable. To obtain a probability of error of at least 10^{-3} for SNR values greater than 20 dB, then an SIR value greater than -6.5 dB must be used. For the BER to be less than 10^{-5} with an SNR of at least 20 dB, the signal can be no more than 5.5 dB below the interfering WiFi’s signal power. Once the SIR reaches the 0-dB threshold, the interference no longer affects the signal, and only the noise is a factor in causing errors within the system because the system changes from interference-limited to noise-limited near an SIR of -5 dB.

The final AWGN channel scenario to be simulated is the Bluetooth signal with a ZigBee interferer. Much like the Bluetooth signal with a Bluetooth interferer, this scenario also appears to be noise-limited. In Figure 4.10 it can be seen that if the SIR is less than -16 dB, the interference drowns out the signal; while if the SIR is greater than -16 dB, then the system will approach zero errors as the SNR is increased. The performance rapidly changes from having an unrecognizable signal at -16 dB to having a BER of less than 10^{-3} at an SNR value of 25 dB, for an SIR value of -15.5 dB. The performance continues to rapidly increase toward the AWGN case, and the interfering signal no longer affects the signal for values of SIR greater than -5 dB.

4.2.4 AWGN Conclusion

Even though the obtained results are only theoretical, especially when assuming a channel composed solely of noise and no fading, the results are a good indicator of the type of performance that can be

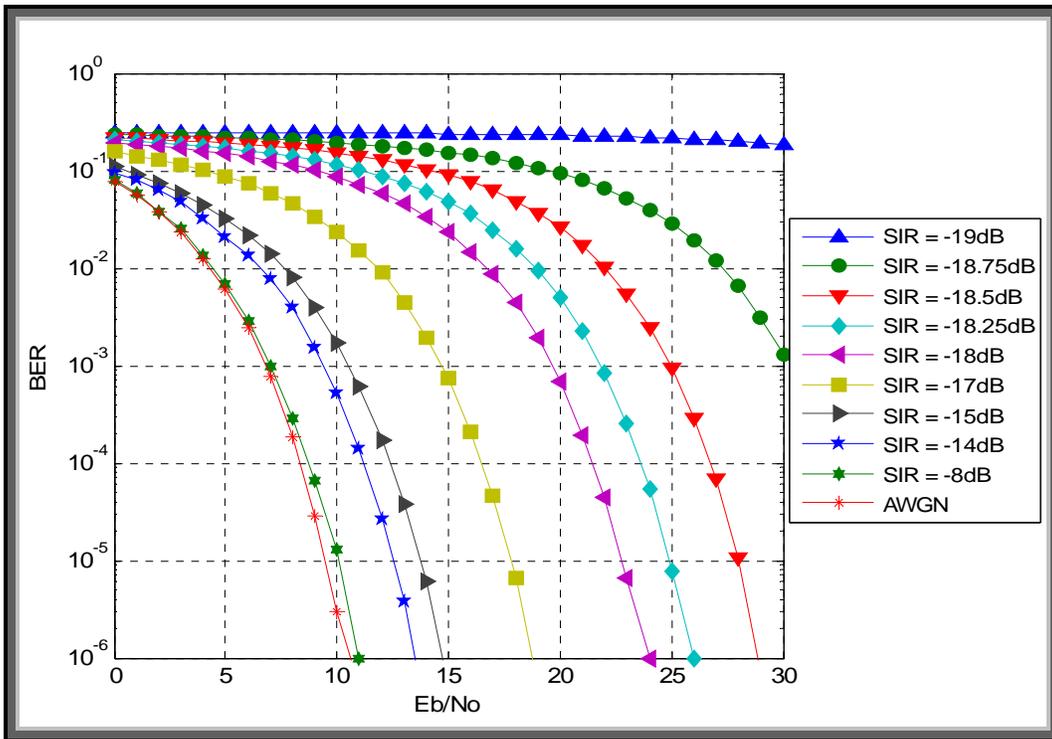


Figure 4.9. Bluetooth signal with a Bluetooth interferer-AWGN channel.

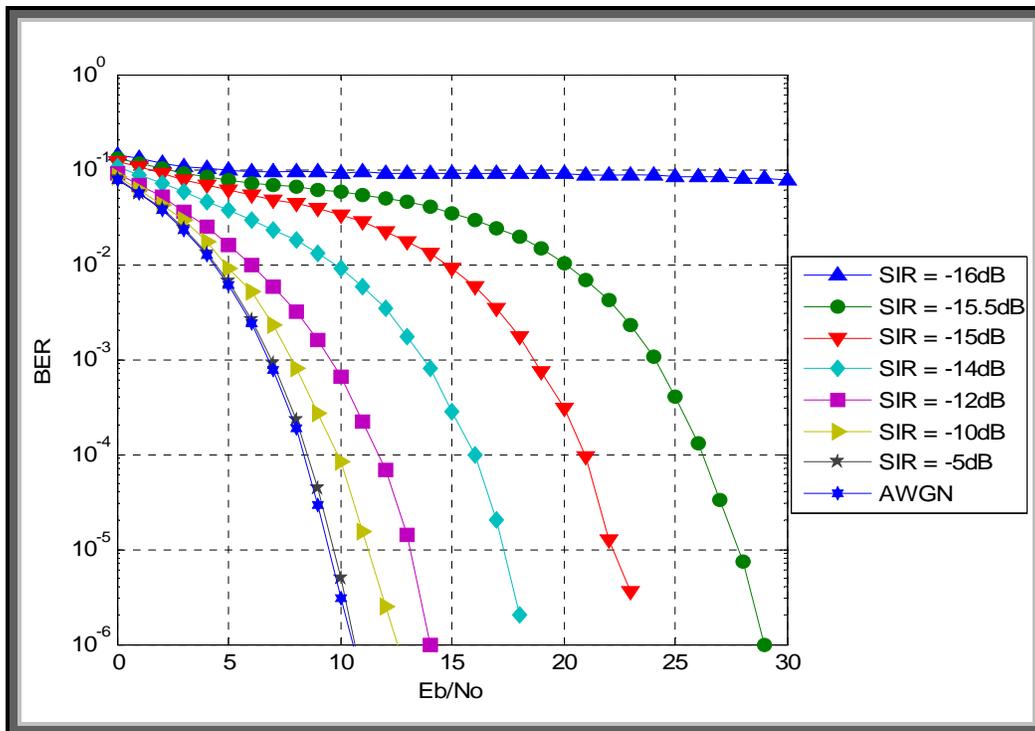


Figure 4.10. Bluetooth signal with a ZigBee interferer-AWGN channel.

expected. Based on the results, ZigBee and Bluetooth appear to be protocols that can survive in an environment in which they are not the dominant devices. This is because they are able to completely disregard interferers with a power level that is at least 5 dB above their own power level. The only exception is Bluetooth with a WiFi interferer, in which case the power levels must be equal. These devices can still tolerate an interferer when the SIR is less than -10 dB; however, Bluetooth with a WiFi interferer can tolerate an SIR of only -6.5 dB. Here the ability of a signal to tolerate an interferer means that in the presence of the interferer, the signal can still reach a BER of at least 10^{-3} at an SNR of 20 dB.

The effects of interference on WiFi differ from those obtained for Bluetooth and Zigbee. The main difference is that WiFi cannot function at full capacity when an interferer is present with a power level higher than that of WiFi. For WiFi to completely ignore the effects of the interferers, it must have an SIR of 5 dB. On the other hand, it can tolerate interferers that have higher power levels, up to 3 dB higher than the WiFi signal. This is why WiFi needs to be transmitted at such a higher power level than either Bluetooth or ZigBee. A second difference between WiFi and the other protocols is that WiFi is affected in the same way no matter which interfering protocol is present, whereas the performances of ZigBee and Bluetooth fluctuate from interferer to interferer.

4.3 Interference over a Rayleigh Flat-Faded Channel

To better understand whether the three different protocols can coexist and what power levels of interference they can tolerate, more has to be considered than just a simple AWGN channel. Within a real-world environment, the signal will experience some sort of fading. The fading is typically modeled as having a Rayleigh distribution. To simulate this situation, a general Rayleigh fading path will be considered for both the signal and the interferer. Because these paths will each consist of only a single path, the incurred fading will be flat fading in which only the signal amplitude is affected and the signal itself is not distorted. The results will be presented exactly as they were for the AWGN scenarios, beginning with ZigBee and ending with Bluetooth. All of the results appear to represent interference-limited systems, and the probabilities of error will be flat for a certain range of SIR values.

4.3.1 ZigBee

The first case considered for ZigBee is the case of a Bluetooth interferer. Figure 4.11 provides the BER results for the varying SIR values. There is a 5- to 10-dB decrease in performance from the AWGN channel, which can be attributed to the Rayleigh flat fading. The system does not perform like a system without a Bluetooth interferer until the SIR ratio is 0 dB. To maintain a probability of error below 10^{-3} for an SNR value of 20, the SIR ratio must now be above -10 dB. The signal becomes unusable for values of SIR below -15 dB.

The second fading example for the case of a ZigBee signal involves a WiFi interferer. Similarly to the previous plot, Figure 4.12 depicts an interference-limited system because the BER flattens out to a certain probability rather than approaching zero. This time the degradation in performance from the AWGN case drops by a value of only 2 dB. The signal now becomes clouded with interference for values lower than -12 dB, whereas before this clouding took place at -14 dB. The system also performs as if no interferers were present at -3 dB, an increase from -5 dB. For this system, values of SIR need to be -7 and -3 dB to maintain BERs of less than 10^{-3} and 10^{-4} , respectively, for values of SNR above 25 dB. The line showing when no interferers are present has approximately a constant slope, showing that the system is flat fading. If it were frequency-selective fading, then the plot would suffer in BER and have a downward curve to it.

Much as in the previous two cases, when a ZigBee interferer is introduced to a ZigBee signal, the flat fading causes an increase in BER from the AWGN case. The results in Figure 4.13 demonstrate the

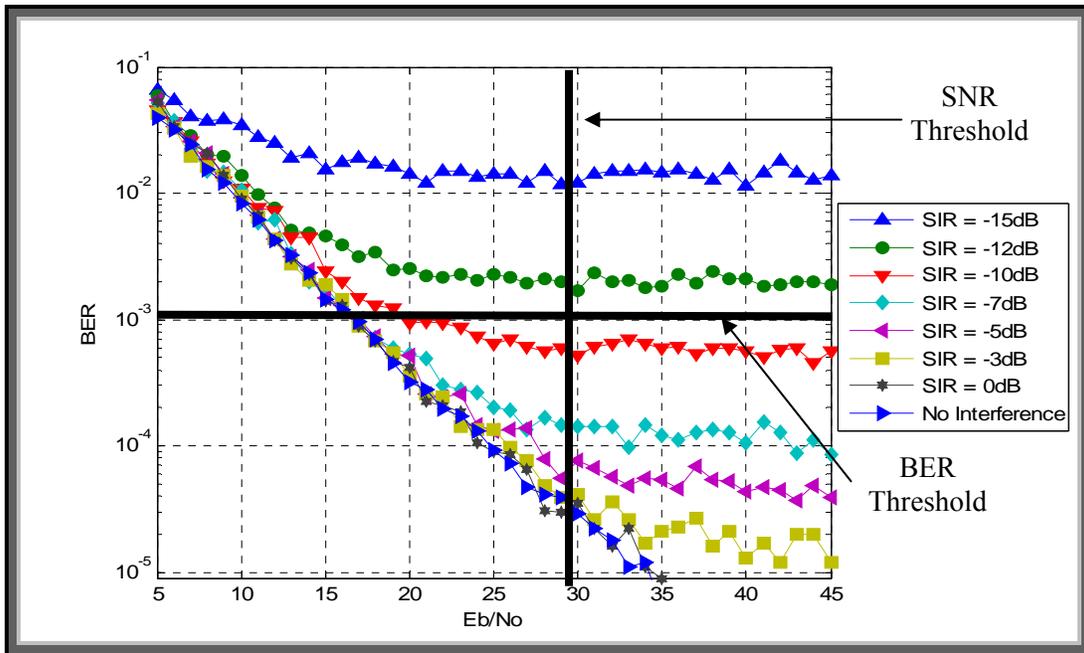


Figure 4.11. ZigBee signal with a Bluetooth interferer – Rayleigh flat fading.

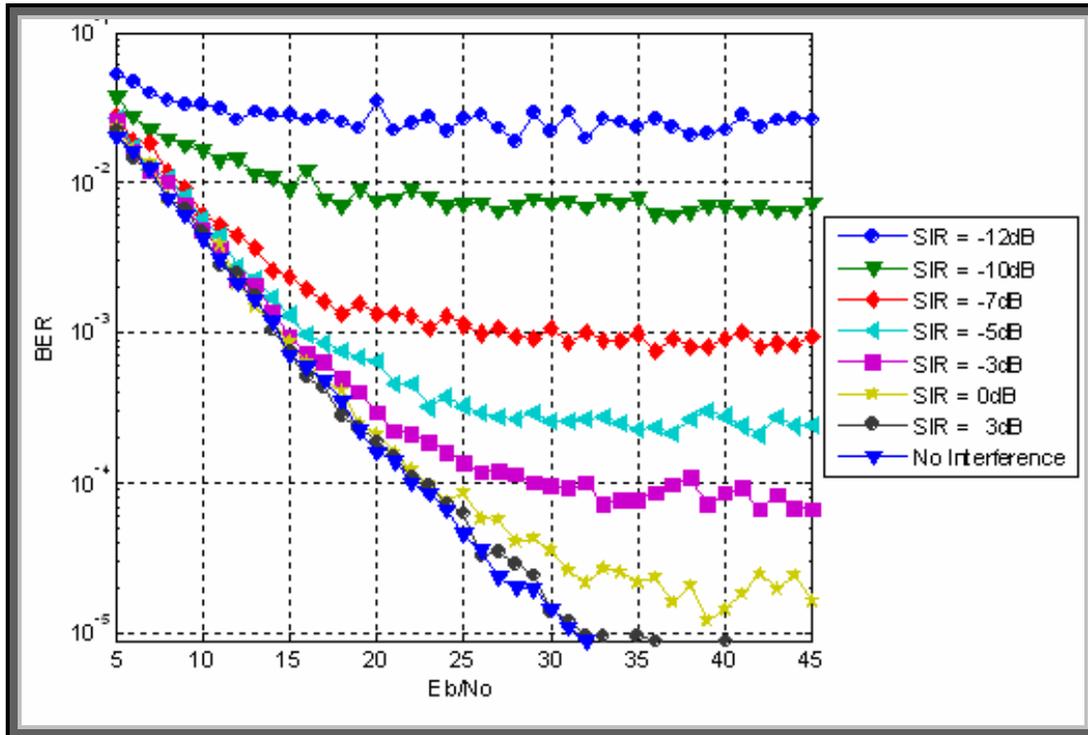


Figure 4.12. ZigBee signal with a WiFi interferer – Rayleigh flat fading.

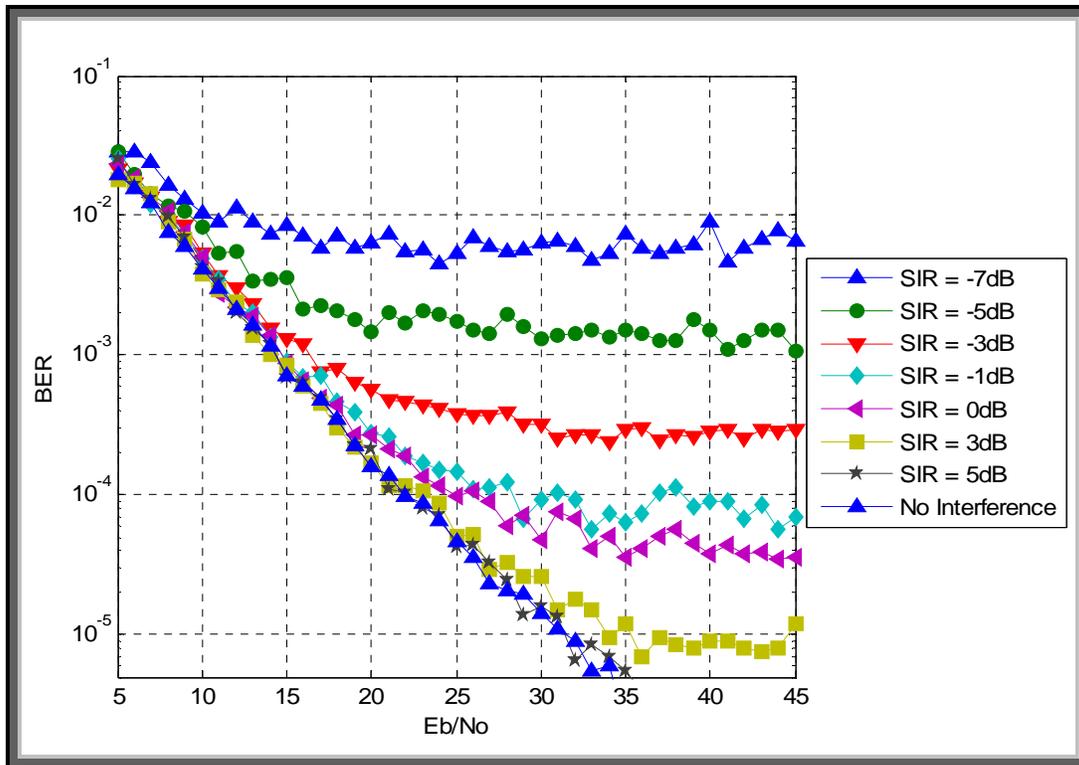


Figure 4.13. ZigBee signal with a ZigBee interferer – Rayleigh flat fading.

effects of the ZigBee interferer. This time the increase in SIR is again between 5 and 10 dB. The system completely gives in to the interference at -7 dB, whereas in the AWGN case in which it gave in at -12 dB. To maintain a BER below 10^{-3} for an SNR above 20 dB, an SIR of approximately -4 dB must be used. The system does not completely block out the other ZigBee signal's interference until the intended signal has a signal power at least 5 dB greater than that of the interferer, a change of 10 dB from the previous case in which it could be 5 dB below the interfering signal's power.

4.3.2 WiFi

The simulations explored how well the WiFi system model performs when subjected to both fading and AWGN environments along with interferers. The first interferer is again Bluetooth, and the results can be found in Figure 4.14. A significant difference can now be seen from the other cases in that the SIR must now be a positive number, meaning the signal power must be greater than the interferer's power. A decrease of 6 to 15 dB in performance from the previous WiFi example with a Bluetooth interferer can be seen. When the signal and interferer have equal signal power, the interferer dominates the signal. As the signal power increases to an SIR of 10 dB, the probability of error hovers under 10^{-3} for SNR values above 35 dB. It takes another 10-dB increase in SIR before the signal can completely overcome the interfering signal and perform as if no interferers are present at 20 dB.

The second interferer applied to the WiFi system is the entrance of a second WiFi signal. Figure 4.15 shows that the SIR falls between the same values as those of Figure 4.14. It ranges from 0 to 20 dB, with 0 dB indicating when the signal is completely distorted by the interferer and 20 dB correlating to when the interferer is completely dominated by the signal. Again, an SIR of at least 10 dB must be achieved for a probability of error of 10^{-3} to be achieved with an SNR of at least 35 dB. For the scenario in which only noise is considered and not fading, there is approximately 10 dB of degradation in performance.

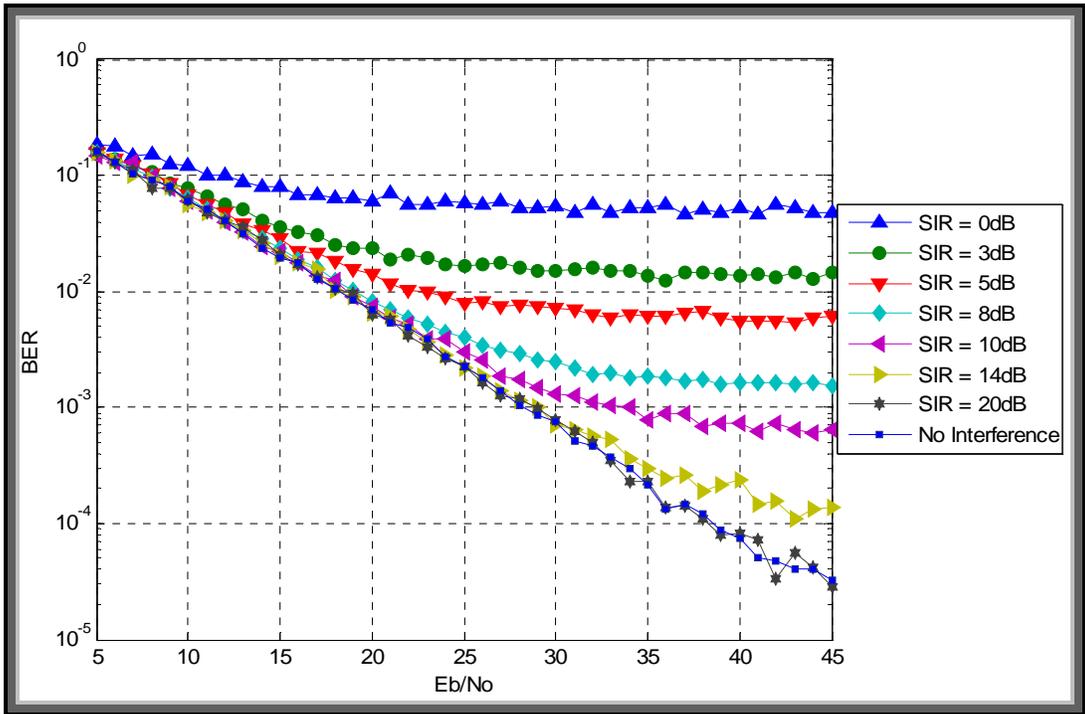


Figure 4.14. WiFi signal with a Bluetooth interferer–Rayleigh flat fading.

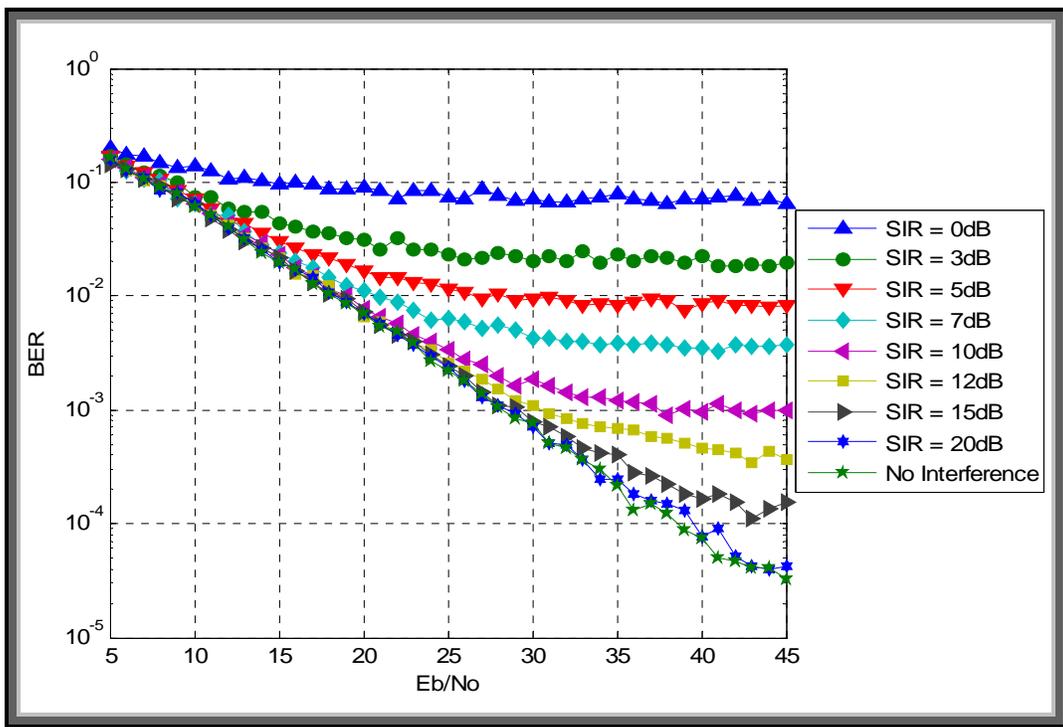


Figure 4.15. WiFi signal with a WiFi interferer–Rayleigh flat fading.

The final interferer to which the WiFi model is subjected is a ZigBee interfering signal. For the first AWGN case, the SIR ranges from values of -5 to 5 dB, from signals that are swamped by interferers to signals that completely block the interfering ZigBee signal. According to the results shown in Figure 4.16, when the effects of fading are included, the SIR must be increased to 20 dB before it can thwart the effects of the interferer. As in the two previous examples, the interfering signal overtakes the signal when the two signals have equal powers; and for the system to achieve a BER of 10^{-3} for values of SNR above 35 dB, the SIR must be equal to or greater than 10 dB. Based on the results in Figures 4.14, 4.15, and 4.16, it appears that no matter what type of interfering signal is subjected to WiFi, the performance stays roughly the same.

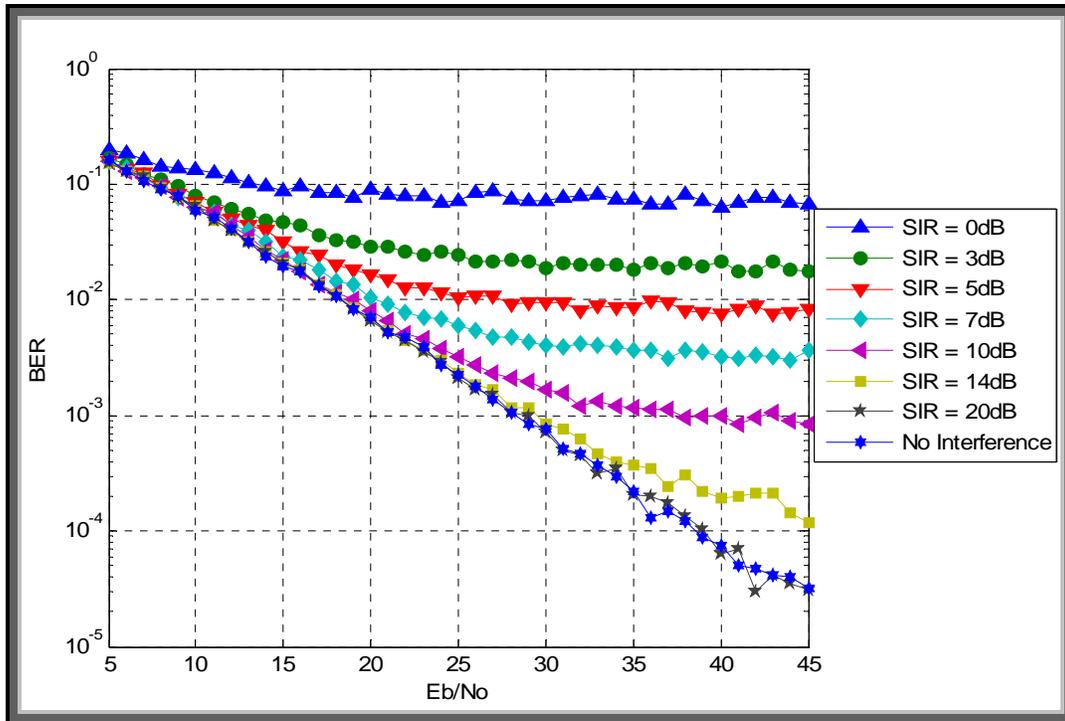


Figure 4.16. WiFi signal with a ZigBee interferer – Rayleigh flat fading.

4.3.3 Bluetooth

The final model in which fading needs to be considered for is Bluetooth. The three different cases all still experiencing flat fading and the performance of all three have been hindered by the addition of fading to the AWGN channel. The Bluetooth interferer creates a 20 -dB range of SIR, from -15 to 5 dB, in which the performance will range from being unable to extract any data from the received signal to receiving the signal as if no interference is involved. These results can be seen in Figure 4.17. From the graph it can be seen that for an SIR of -5 dB, the BER stays below 10^{-3} for values of SNR greater than 30 dB.

When Bluetooth is combined with a WiFi interferer, the performance is considerably worse than when a Bluetooth interferer is considered. The value of SIR must be roughly 10 dB higher to achieve the same performance. Also, the system suffers about a 5 -dB loss in signal power from the case in which only an AWGN case is considered. In Figure 4.18 it can be seen that a WiFi signal completely saturates the Bluetooth signal for values of SIR lower than -5 dB. For SNR values greater than 25 dB, the system

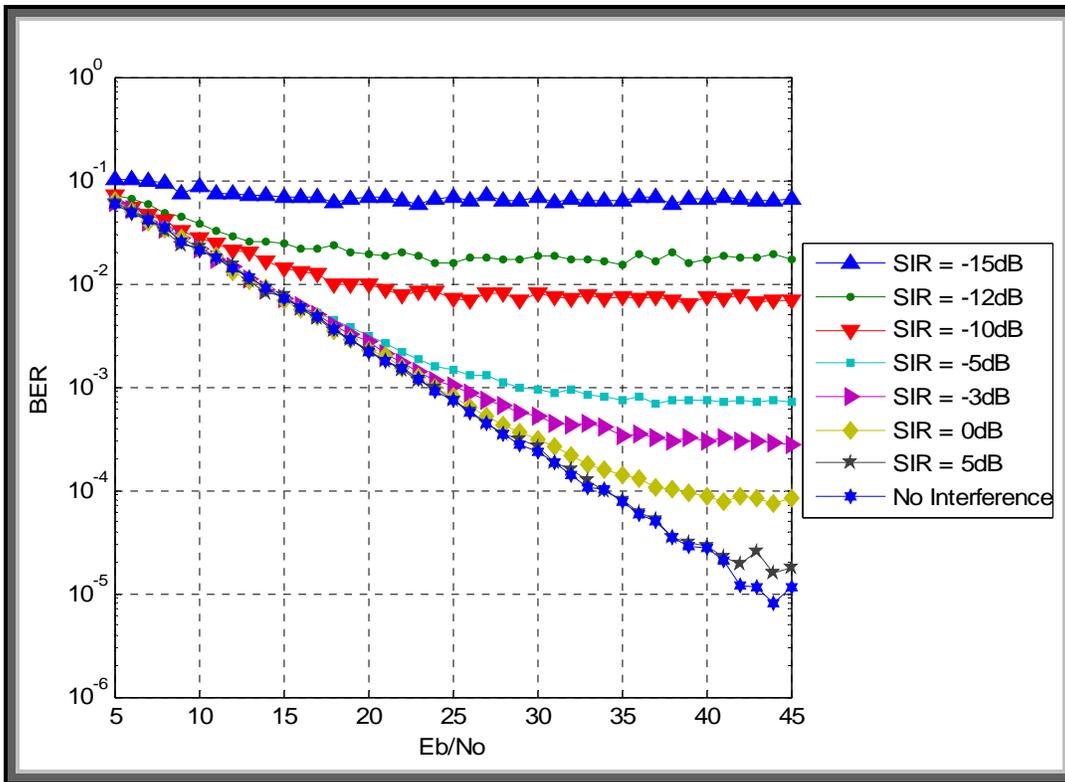


Figure 4.17. Bluetooth signal with a Bluetooth interferer–Rayleigh flat fading.

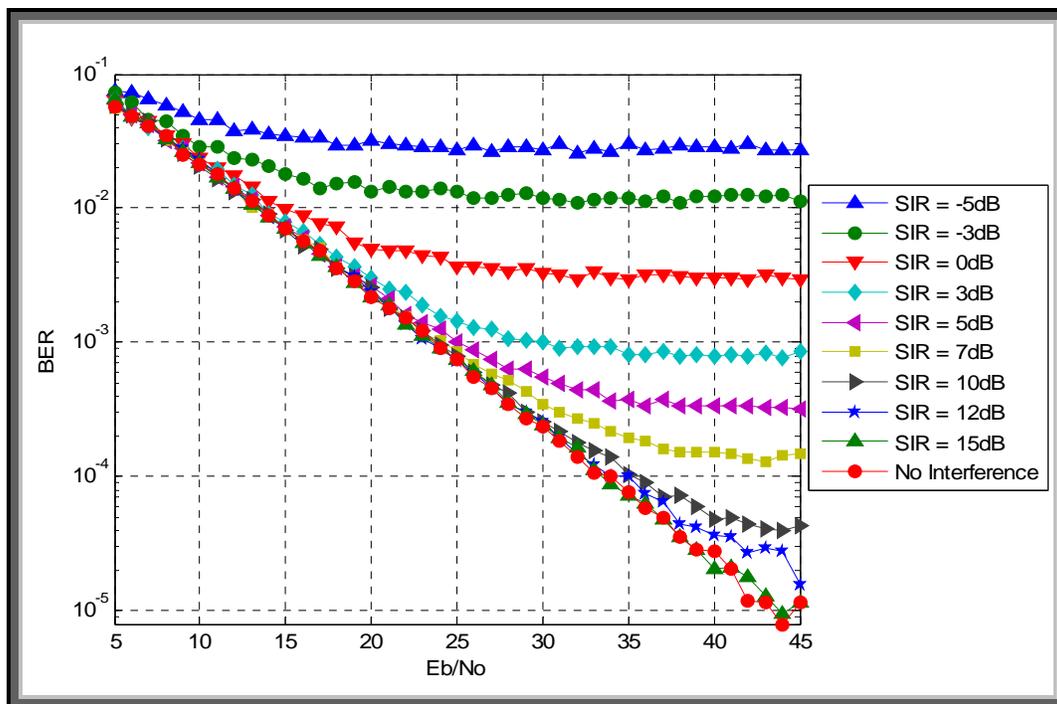


Figure 4.18. Bluetooth signal with a WiFi interferer–Rayleigh flat fading.

can sustain a BER of 10^{-3} as long as the Bluetooth signal is 3 dB greater than the interfering WiFi signal. Bluetooth cannot block out the WiFi interferer until the SIR is at least 15 dB.

The final fading case for the Bluetooth signal is when a ZigBee interferer is combined with it. For the AWGN-only case, Bluetooth is fairly resilient against the ZigBee signal and is able to maintain a signal power 5 dB below ZigBee's power level. For other cases, the Bluetooth power level must increase to at least 5 dB above the interfering signal of ZigBee, as demonstrated in Figure 4.19. For Bluetooth to maintain the 10^{-3} BER for SNR values above 30 dB, the SIR value must be in the neighborhood of -4 dB. The Bluetooth signal does not become saturated with the ZigBee interfering signal until the signal falls to more than 12 dB below the ZigBee signal.

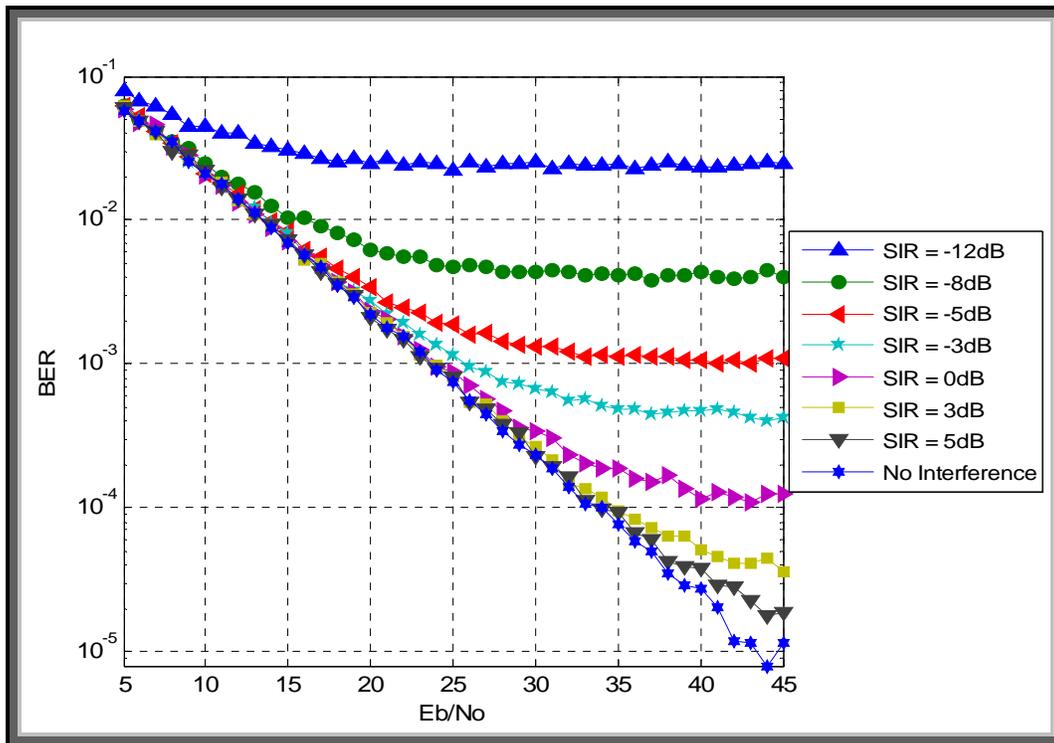


Figure 4.19. Bluetooth signal with a ZigBee interferer–Rayleigh flat fading

4.3.4 Rayleigh Fading Conclusions

As expected, the performance for the simulations incorporating Rayleigh fading are substantially lower than for those incorporating only an AWGN channel. On average, the effects of fading cause a 10- to 15-dB drop in performance. As a result of the 10- to 15-dB decrease, the level at which a signal can tolerate another signal has been extended to at least 10^{-3} at an SNR of 30 dB. WiFi operates along this 15-dB drop from the AWGN case. All interferers still affect the signal in the same way; however, rather than an SIR of 5 dB being needed to completely block out the interferers, the value now must be equal to 20 dB. Also, the signal can now tolerate only interferers with an energy at least 12 dB below its own energy, whereas in the AWGN case the signal could still be received when the interferer had an energy greater than that of the WiFi signal.

ZigBee is less affected by the fading than WiFi. ZigBee follows a 10-dB drop in performance. However, for ZigBee to operate as if no interferers are present, the SIR must be at least 5 dB in the presence of

another ZigBee signal, but it can extend up to 0 dB when a Bluetooth signal is incorporated. When ZigBee only must be able to tolerate the signals and can withstand needing a BER of at least 10^{-3} , the ZigBee interferer can have a higher power level than the ZigBee signal and approach an SIR of -3 dB. This value can be decreased to -10 dB in the presence of a Bluetooth interferer; the performance from the WiFi interferer falls between that of the other two.

Bluetooth is affected much more than ZigBee and encompasses the 15-dB decrease in performance. For Bluetooth to operate at a level equal to operation with no interferers, it must have an energy level at least 5 to 15 dB greater than that of the interferers. For a WiFi interferer, an SIR of 15 dB is required. Bluetooth can tolerate signals with values between 3 and -3 dB of its transmitted power, approximately 10 dB less than that of the previous AWGN situation. The tolerance of Bluetooth for interferers is less affected by fading than is the ability to block them.

5. SITE-SPECIFIC CHANNEL MODEL

5.1 Wireless InSite Environment Simulation

To take the simulation one step further, rather than incorporating only one faded path into the simulation, multiple paths should be considered to account for the various paths a signal can travel between the transmitter and receiver. In reality, the presence of only one propagation path is likely only during communication with a satellite in space or with a ship on the ocean because of the absence of reflectors and/or scatterers. For the purpose of this study, which aims to predict the performance of wireless devices within a nuclear power plant, numerous scatterers and reflectors will be present that need to be accounted for, especially when a direct path between the transmitter and receiver does not exist. Therefore, a tool is needed that not only can calculate the different paths between a transmitter and receiver and their respective propagation properties, but also can include objects such as desks, shelves, pipes, air ducts, equipment, and other types of items expected to be present within an industrial environment.

A software tool with the capability of performing both required tasks is the Wireless InSite modeling package developed by Remcom and discussed in Sect. 3. Available materials include concrete, wood, metal, and glass, among others, and the user can also create new materials and save them for future use. The materials are used to determine various coefficient values such as reflection, transmission, and diffraction, which will all help to create a dependable simulation. To complement selection of various materials, Wireless InSite also permits the selection of several other parameters and features. Location of transmitters and receivers is allowed through the floor-plan editor, as is selection of the type of antenna they are using, such as isotropic or dipole. The software also allows the selection of several different signal characteristics, such as waveform, bandwidth, and carrier frequency. The specific values incorporated into this simulation will be discussed later.

5.2 Site-Specific Room Model

For security reasons, detailed drawings of a nuclear power plant are not used in the site-specific room model. Instead, the test environment used for the simulations will consist of a partial representation of an actual building that would be representative of an office/laboratory-type environment. Both the placement and properties of the transmitters and receivers will be considered.

5.2.1 Transmitter/Receiver Placement

As a starting point for the simulation, only one receiver is considered to be present, allowing for a minimization of variables contained within the simulation. Figure 5.1 shows that the receiver is placed near the perimeter of the room above a metal workspace. The receiver height, along with all the transmitter heights, was set at 2 m (~6.6 ft), just above eye level, to ensure the objects placed within the room will have a minimal effect on the results. Because all the objects are less than 2 m (~6.6 ft) tall, transmitters placed within the main room will be LOS transmitters, and ones placed outside the main room will be NLOS transmitters. After the simulation has been validated, a more complex placement of the receiver will be possible.

Figure 5.1 shows that the transmitters are placed so as to maximize coverage. The main room is split into six sections with a transmitter placed in each partition. This placement allows for an analysis of the effects of the performance of the primary LOS path because of the presence of the secondary propagation paths. Transmitters #7 through #11 were placed outside the main room so that some transmitters would need to transverse through only one wall to reach the receiver, while others would need multiple reflections and/or transmissions to arrive at the receiver. Conducting this trial will not determine the range of the protocols because all powers are normalized at the receiver, but it will demonstrate how the path on which the signal propagates affects the performance.

5.2.2 Signal Properties

Next, the properties of the signals that the transmitters and receivers are using will be explored. In Figure 5.1 note that there are actually three different transmitters or receivers at each transmitter and receiver location, one for each protocol. This is because the different signals will not have the same propagation characteristics; therefore, separate transmitters and receivers are needed. The properties specific to the waveform for each transmitter and receiver can be found in Figure 5.2, which shows a window taken from the Wireless InSite program. This figure shows that the Bluetooth signal uses a Gaussian-type signal with a bandwidth of 1 MHz located at a carrier frequency of 2.405 GHz. This value corresponds to channel #4 within the Bluetooth spectrum. The second waveform is a sinusoidal waveform representing ZigBee, with a bandwidth of 2 MHz and also located at 2.4 GHz, which is channel #1 for ZigBee. WiFi is assumed to contain a raised cosine pulse shape, occupies a bandwidth of 22 MHz and is centered at 2.412 GHz, relating to channel #1.

After the waveform for each transmitter and receiver is specified, several other parameters must also be chosen. A second window from Wireless InSite can be seen in Figure 5.3. This figure shows the specifics for a Bluetooth transmitter. One difference between selecting the properties for the transmitter and selecting those for the receiver is the selection of the transmitter power level. As seen in the figure, the transmitter uses the Bluetooth waveform along with an isotropic antenna; in fact, all antennae in this simulation are isotropic. Another consideration specific to the transmitter is the radiated power. In Figure 5.3, the power is set to 4 dBm, as specified for Bluetooth. The other two protocols use different values; ZigBee transmits at 0 dBm, while WiFi radiates at a much higher 17 dBm.

5.2.3 Wireless InSite Output

After various parameters for the simulation of the propagation paths are selected, a proper understanding of how to identify and interpret the results is needed. The output of two different transmitters, WiFi transmitters #5 and #7, which are in close proximity to each other, is presented to show the difference caused by having to deal with obstructions located in the path of propagation, in this case a wall.

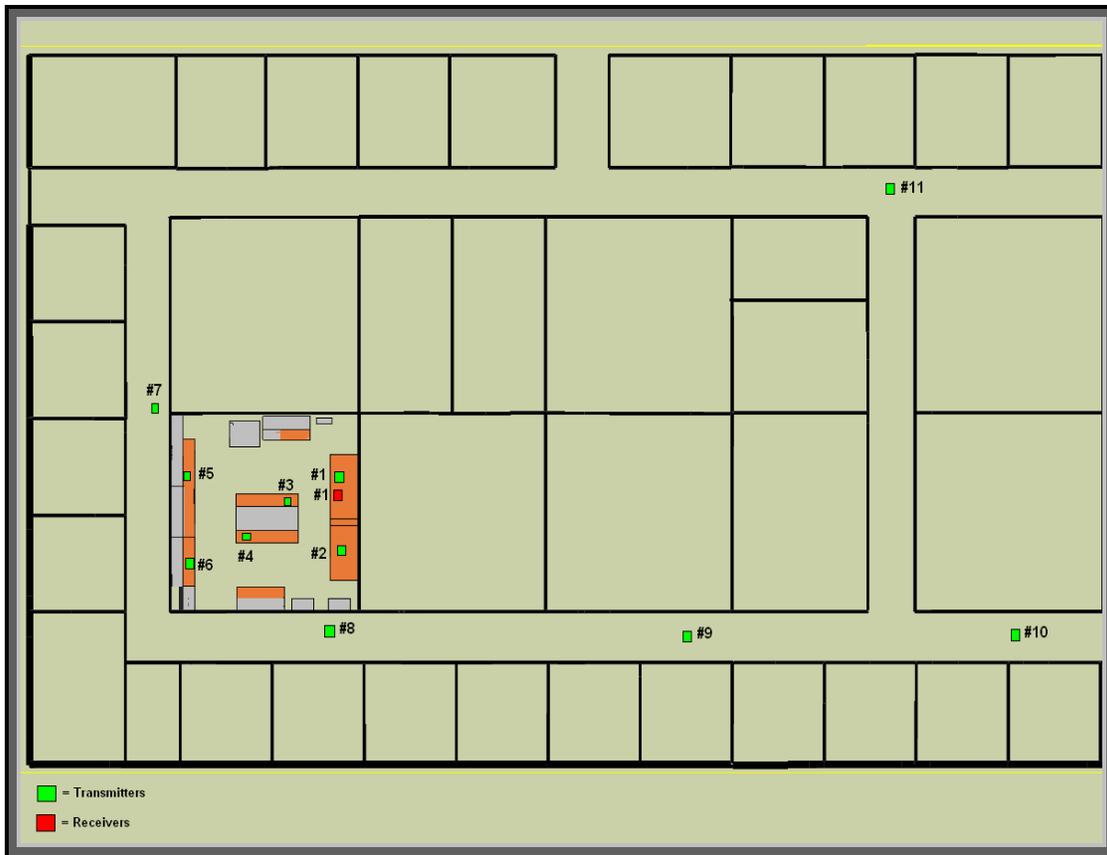


Figure 5.1. Transmitter and receiver locations.

The screenshot shows the 'Wireless InSite' software interface. The title bar reads 'Wireless InSite - Main: (B1_furniture_with 2nd floor) [C:\...\3500 cases\B1_furniture_with 2nd floor_Chad.se...'. The menu bar includes 'Project', 'Edit', 'View', and 'Help'. The toolbar contains icons for 'Study areas', 'Transmitters', 'Receivers', 'Materials', 'Antennas', 'Waveforms', and 'Requested out'. Below the toolbar is a table with the following data:

In use	Type	Description	Frequency	Bandwidth	Dispersive
Yes	Gaussian	Gaussian_BL	2405.	1.000	No
Yes	Sinusoid	ZigBee	2405.	2.000	No
Yes	Raised cosine	Raised Cosine_WiFi	2412.	22.00	No

Below this table is another table with columns: 'Type', 'Short description', 'Frequency', 'Location', and 'Ld'. The status bar at the bottom indicates 'Selection: ZigBee [Sinusoid]'.

Figure 5.2. Waveform properties.

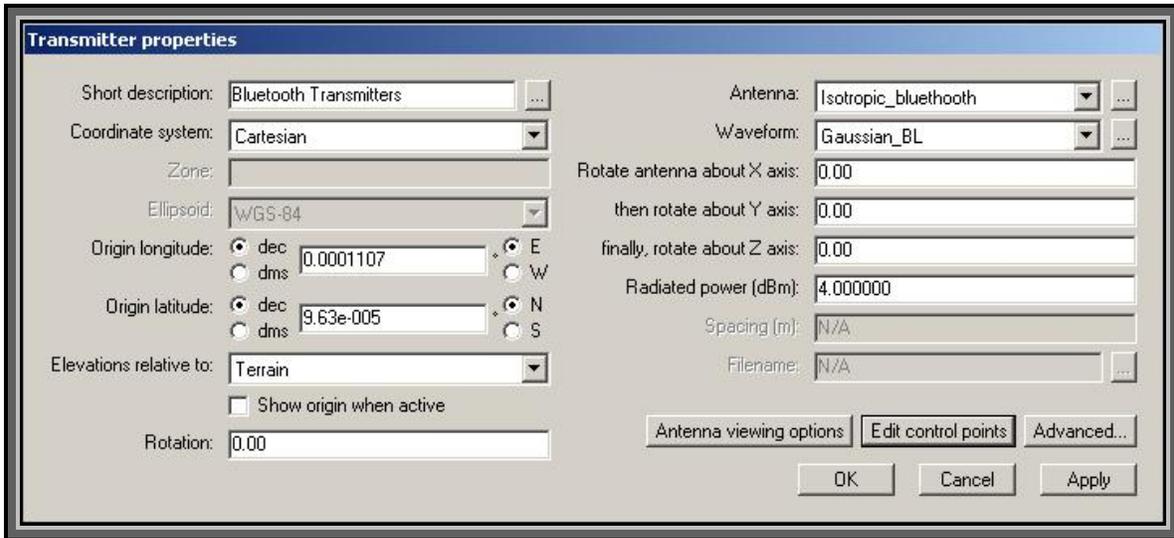


Figure 5.3. Transmitter properties.

WiFi transmitter #5 is a LOS transmitter located in the upper left-hand corner of the room shown in Figure 5.1. The resulting propagation paths associated with the transmitter will be composed of one direct path with nine supporting paths, each of which contains one or more reflections. NLOS WiFi transmitter #7 is located within the same area as transmitter #5 except that it is outside the room; thus, the direct paths must transmit through the wall to reach the receiver. Therefore, all secondary paths will also contain a transmission through a wall along with their multiple reflections. The propagation paths can be found in Figure 5.4 for transmitter #5 and in Figure 5.5 for transmitter #7. These figures also show the relative power contained within each path. The lighter-colored lines indicate more power contained within the path, while the darker lines correspond to a lower received power, as specified by the bar indicator along the bottom of the figures.

The power contained within each path can also be shown in graph form. Figures 5.6 and 5.7 show the relative received power of each path versus the delay associated with each path for transmitters #5 and #7, respectively. These graphs are used to reinforce the notion of the advantage in power that a LOS transmitter has over an NLOS transmitter, as well as the power advantage that the first arriving path has over all secondary paths, which undergo reflections. The results for transmitters #5 and #7 are provided in Tables 5.1 and 5.2, respectively. The results indicate that the LOS path of transmitter #5 has a 7-dB advantage over the performance of the single transmission path of transmitter #7; therefore, a transmission can cause a decrease in power of approximately 7 dB. Comparing the first path with the second path for each set of results, this shows that a reflection has occurred. Thus it appears that a reflection inflicts a 10-dB drop in power. These results are only approximate and are valid for only concrete walls and in considering a WiFi signal, but the trend of a transmission harming the power of a signal less than a reflection is a valid observation within this study.

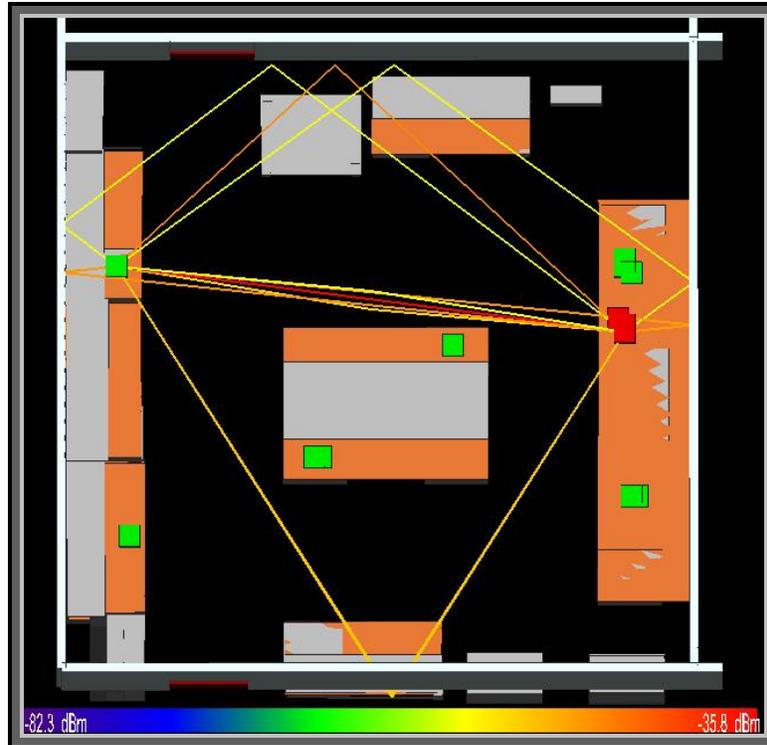


Figure 5.4. Propagation paths for WiFi transmitter #5.

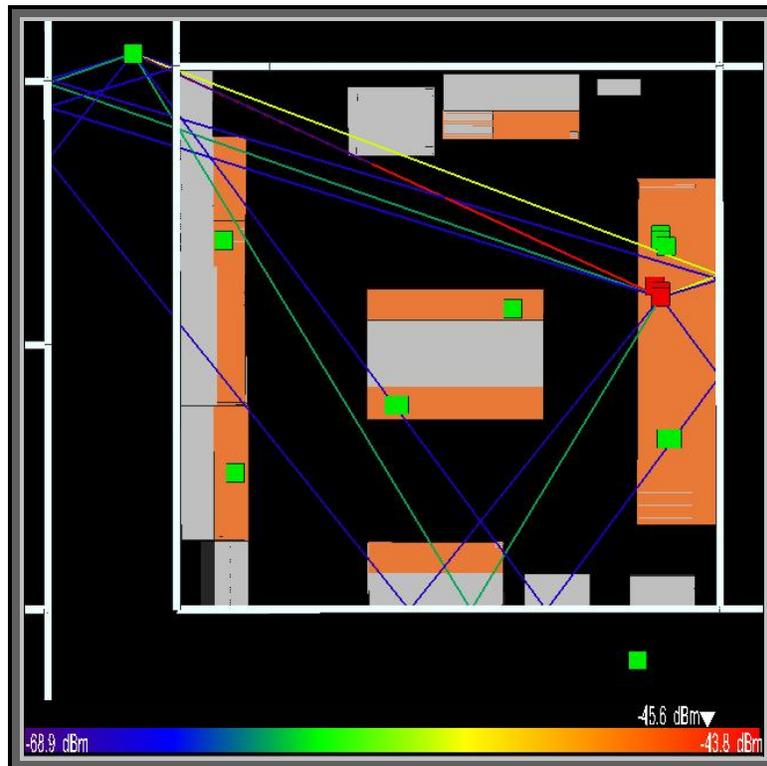


Figure 5.5. Propagation paths for WiFi transmitter #7.

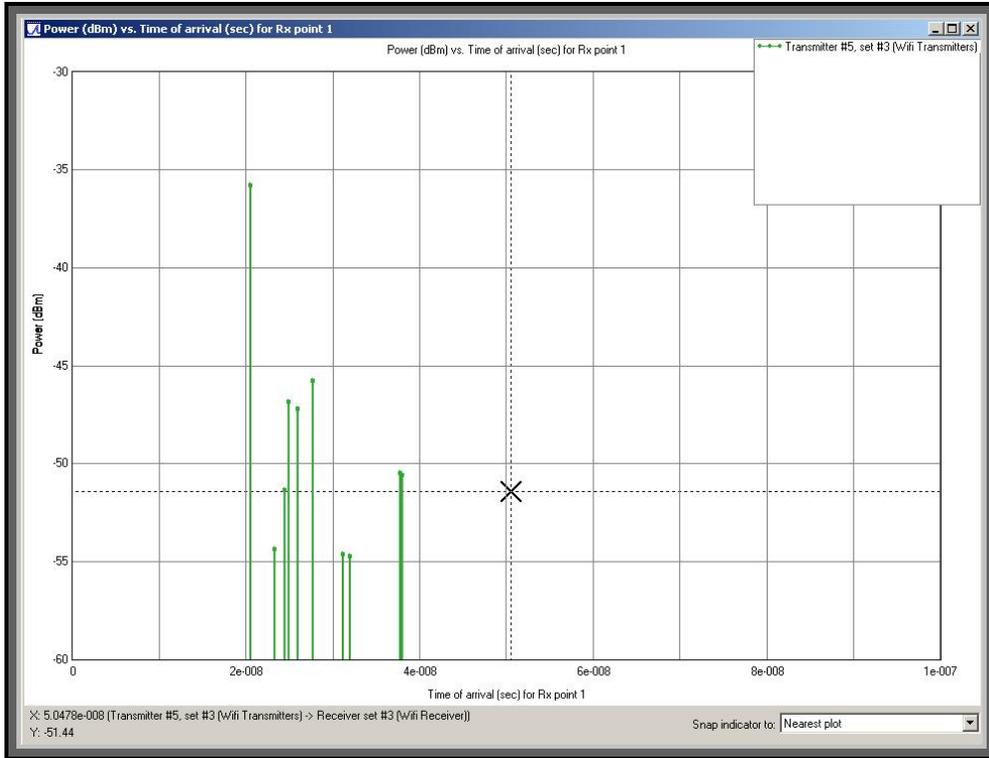


Figure 5.6. Power vs delay–WiFi transmitter #5.

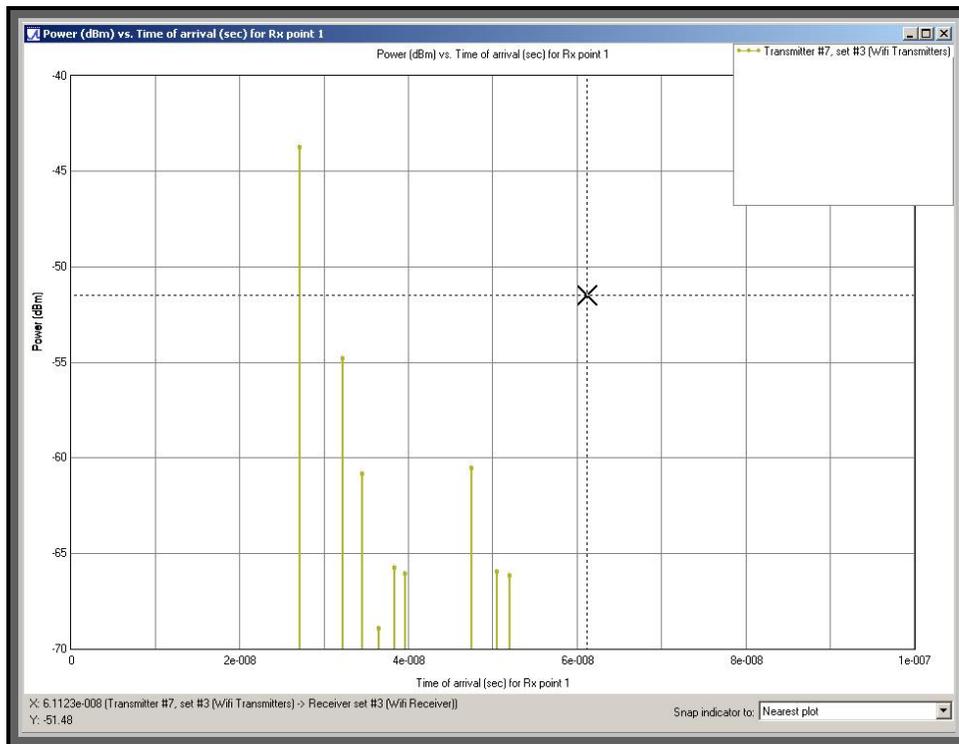


Figure 5.7. Power vs delay–WiFi transmitter #7.

Table 5.1. WiFi transmitter #5 propagation paths

Path number	Phase (degrees)	Time (s)	Power (dBm)
1	-151.852	0.205496E-07	-35.804
2	-104.970	0.277587E-07	-45.809
3	-173.928	0.249576E-07	-46.864
4	-15.263	0.260218E-07	-47.225
5	-148.987	0.378032E-07	-50.521
6	27.667	0.380144E-07	-50.603
7	19.466	0.244994E-07	-51.385
8	-64.001	0.233424E-07	-54.371
9	1.848	0.311631E-07	-54.649
10	-18.962	0.320198E-07	-54.746

Table 5.2. WiFi transmitter #7 propagation paths

Path number	Phase (degrees)	Time (s)	Power (dBm)
1	173.318	0.271163E-07	-43.755
2	-71.842	0.321982E-07	-54.803
3	-44.896	0.474554E-07	-60.570
4	-21.160	0.345179E-07	-60.840
5	-157.689	0.383456E-07	-65.787
6	-15.598	0.505357E-07	-65.973
7	113.873	0.396962E-07	-66.063
8	125.829	0.520404E-07	-66.180
9	125.829	0.520404E-07	-66.180
10	-77.880	0.364853E-07	-68.938

5.3 Results

The resulting BER curves can be grouped into three different categories, one for each protocol, showing the effects that an interfering signal has on a given protocol. Furthermore, each separate protocol can be broken down into its LOS and NLOS situations. Eleven different transmitter points were chosen with three different devices located at each position, and there are three separate interference cases for each protocol; that corresponds to the production of 99 different BER curves.

5.3.1 Transmitter/Receiver Performance

Before presenting the results for each device, it is important to look at the propagation characteristics from Wireless InSite for each transmitter location because the results will help determine whether the protocol should be able to perform at an acceptable level even before the interference has been introduced. Depending on the powers and phases of the subsequent paths after the first arriving path, these secondary paths might decrease the overall performance of the system. Therefore, the propagation path properties of the three protocols from Wireless InSite are summarized in Table 5.3 for Bluetooth, Table 5.4 for WiFi, and Table 5.5 for ZigBee. The three tables show the three different devices with each

of the 11 different transmitters for each protocol. Along the top of the table are the ten different paths along with the contribution that each path makes to the overall signal. The contribution of each signal is found by taking the normalized path power of the individual path and multiplying it by the cosine of the arriving phase of the path. Therefore, the number can be positive or negative depending on the relative phase of each path with the first arriving path. The total of all individual path contributions gives the total power contained within the received signal.

Table 5.3. Bluetooth propagation paths

		Propagation path number										
		1	2	3	4	5	6	7	8	9	10	Total
Bluetooth transmitter number	1	1.000	0.008	-0.008	-0.005	0.003	-0.002	0.001	0.001	0.001	-0.002	0.997
	2	1.000	0.253	-0.120	0.108	0.006	0.017	0.004	0.004	-0.005	-0.007	1.260
	3	1.000	-0.020	0.000	-0.029	0.000	-0.003	-0.003	-0.003	0.004	-0.003	0.943
	4	1.000	-0.725	0.056	-0.042	-0.035	0.022	0.014	0.000	-0.006	0.006	0.290
	5	1.000	-0.102	0.070	-0.047	-0.026	0.025	-0.027	0.001	-0.002	-0.007	0.885
	6	1.000	0.724	-0.077	0.038	0.048	-0.063	-0.058	0.043	-0.001	-0.001	1.653
	7	1.000	-0.014	0.037	-0.018	0.006	-0.002	0.000	0.004	0.004	0.004	1.021
	8	1.000	-0.364	-0.048	0.037	0.028	0.025	0.013	-0.008	0.009	-0.001	0.691
	9	1.000	-0.011	-0.011	0.350	-0.182	-0.152	-0.068	-0.043	-0.043	-0.003	0.837
	10	1.000	0.242	-0.069	0.082	-0.075	-0.009	-0.020	0.012	-0.011	0.010	1.162
	11	1.000	0.053	0.012	0.000	-0.012	0.003	0.001	0.002	-0.001	0.002	1.060

Total path powers of transmitters whose signals are able to be received and demodulated are highlighted.

Table 5.4. WiFi propagation paths

		Propagation path number										
		1	2	3	4	5	6	7	8	9	10	Total
WiFi transmitter number	1	1.000	-0.017	-0.008	-0.005	0.003	-0.001	0.001	-0.000	-0.000	0.000	0.973
	2	1.000	-0.296	-0.156	0.104	0.026	-0.018	0.004	-0.004	0.005	0.003	0.668
	3	1.000	-0.014	0.006	-0.029	-0.004	-0.007	0.001	0.001	0.004	0.001	0.959
	4	1.000	-0.828	0.057	0.004	-0.004	0.019	0.022	0.013	-0.005	0.006	0.284
	5	1.000	0.068	0.073	-0.052	0.034	-0.033	-0.027	0.001	-0.012	-0.009	1.043
	6	1.000	0.741	-0.049	0.054	0.062	-0.052	-0.044	0.002	-0.011	-0.011	1.712
	7	1.000	-0.033	-0.016	-0.019	0.006	-0.006	0.003	0.004	0.004	-0.001	0.942
	8	1.000	-0.411	-0.047	0.035	0.029	0.029	0.018	0.015	0.012	-0.002	0.678
	9	1.000	0.208	0.448	0.267	-0.143	-0.069	-0.044	-0.019	0.031	-0.005	1.674
	10	1.000	0.241	0.060	0.083	-0.094	-0.017	-0.016	0.012	0.009	0.010	1.288
	11	1.000	0.165	0.052	0.007	0.010	-0.001	-0.004	0.002	-0.001	-0.001	1.229

Total path powers of transmitters whose signals are able to be received and demodulated are highlighted.

Table 5.5. ZigBee propagation paths

		Propagation path number										
		1	2	3	4	5	6	7	8	9	10	Total
ZigBee transmitter number	1	1.000	-0.002	-0.008	-0.005	0.003	0.002	0.000	0.000	0.000	0.000	0.990
	2	1.000	0.501	0.148	0.106	-0.032	-0.010	0.004	-0.007	0.005	-0.003	1.712
	3	1.000	0.024	0.03	0.006	-0.013	-0.009	-0.003	-0.003	0.003	0.002	1.037
	4	1.000	-0.769	0.048	-0.027	-0.026	-0.011	-0.018	-0.003	0.002	0.003	0.199
	5	1.000	-0.045	0.067	0.060	-0.032	0.031	-0.017	0.009	0.012	0.005	1.090
	6	1.000	0.746	0.072	0.010	0.069	-0.067	-0.015	-0.039	0.012	0.011	1.799
	7	1.000	0.071	0.024	-0.019	0.006	-0.006	0.006	0.005	0.003	0.003	1.093
	8	1.000	-0.299	-0.050	-0.009	0.025	0.013	-0.021	-0.017	0.010	-0.002	0.650
	9	1.000	-0.127	0.227	-0.173	-0.164	0.103	-0.065	-0.042	0.043	-0.035	0.767
	10	1.000	0.250	-0.155	0.072	-0.058	-0.003	-0.022	-0.004	0.008	0.010	1.098
	11	1.000	-0.380	-0.039	-0.054	0.084	-0.005	-0.008	-0.001	-0.002	-0.14	0.455

Total path powers of transmitters whose signals are able to be received and demodulated are highlighted.

In Table 5.3, which shows the summation of the Bluetooth paths, the total path powers of the transmitters whose signals can be received and demodulated are highlighted. All transmitters with a total contained power greater than 0.885 can be recovered, while all transmitters with values less than that are unusable because the secondary paths took too much information away from the first arriving path.

The results in Table 5.4 for WiFi are somewhat different than those for Bluetooth. Even though the total received power for transmitter #2 is equal to 0.668, the signal still can be received. However, although transmitter #7 has a total value of 0.941, it yields unacceptable performance. This discrepancy can be attributed to the fact that the dominant LOS path of transmitter #2 carries a strong unfaded path between the transmitter and receiver, and the secondary paths that take information away from the signal are faded and have varying fluctuations within the signal. Therefore, their effect on the overall signal can be minimized. But in considering the NLOS path of transmitter #7, the faded secondary signals have a greater effect on the overall signal because the dominant signal is also faded and cannot maintain a solid power level. This causes the performance of transmitter #7 to be unacceptable. The results for the ZigBee transmitters, found in Table 5.5, are on a par with the results for the Bluetooth transmitters in that all transmitters with a total combined power of more than 0.8 generate acceptable results, even though the same transmitters between Bluetooth and ZigBee did not perform in the same manner.

An interesting phenomenon to note regarding the propagation paths common to all three of the protocols involves the second paths of transmitters #2, #4, #6, and #8. As can be seen from the tables, these transmitters all have a secondary path with an absolute power greater than 0.25, even though the first path is either a LOS path or an NLOS path consisting of just one transmission. This occurrence can be credited to the metal desk off which the second path bounces before entering the receiver. In the cases of transmitters #4 and #8, the energy caused by the second path, coupled with the phase shift associated with striking the metal surface, causes the transmitters to become inoperable. The second path cancels out too much of the power of the overall signal. The effect is totally opposite for transmitter #6. In this case, the signal always constructively adds with the primary path and increases the total energy by almost 75% of the energy found within the first path. For transmitter #2, the performance is enhanced for both the Bluetooth and ZigBee simulations, but the effects in considering the WiFi protocol tend to harm the

signal. This illustrates that in harsh environments in which there are many metal objects, the performance is highly dependent upon how the signal interacts with those objects.

5.3.2 BER Curves

Next to be considered are the transmitters that yield more acceptable propagation paths and are more likely to result in successful outcomes. Beginning with the ZigBee LOS scenario, the performance should follow a Ricean behavior because of the dominant LOS component. Furthermore, flat fading should be expected as a result of disregarding the delay aspect of the simulation because of its minimal presence. After careful inspection of the LOS cases, which include transmitters #1 through #6, the results can be represented by three different curves, as described in the following sections.

5.3.3 ZigBee LOS

A reasonable representation of the BER curves of all interferers for ZigBee transmitters #1 and #3 can be found in Figure 5.8, which shows the plot of ZigBee transmitter #3 with ZigBee interferers. In this BER curve, the results of the separate interferers are indistinguishable from the case in which no interferers are considered. All results reach a BER of 10^{-5} before an SNR value of 9 dB. As a result of the Ricean LOS component, the results can be approximated as coming from an AWGN environment because the LOS path was dominant.

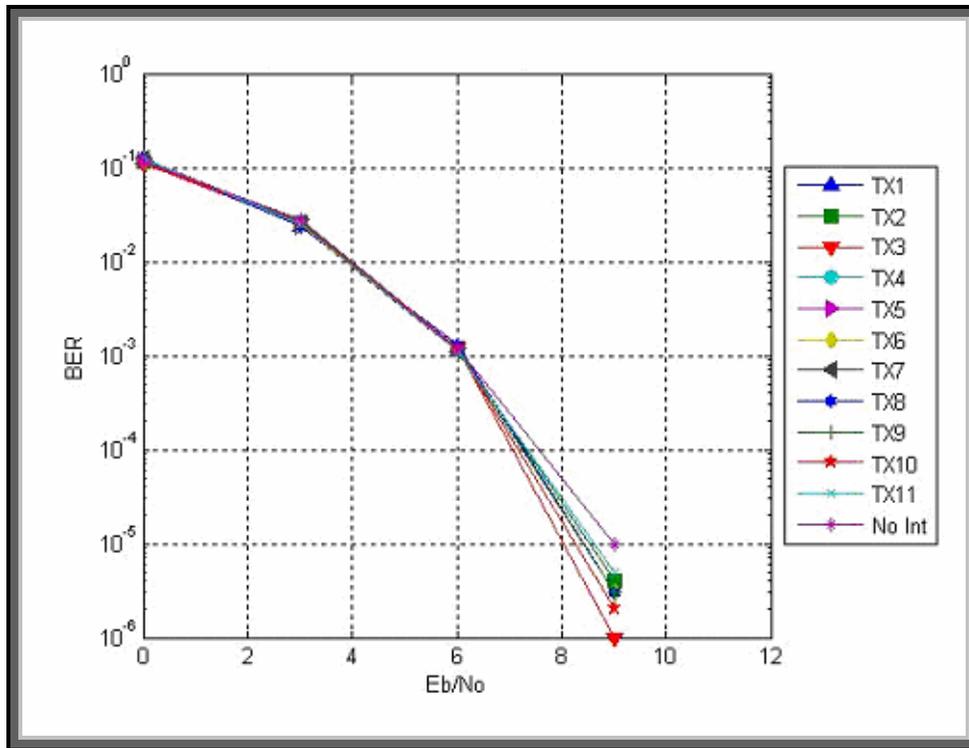


Figure 5.8. ZigBee transmitter #3 with ZigBee interferers.

The previous example illustrated that the interferers had little or no effect on the overall signal. Thus, the message can be received as if no interference were present. Even though the interferers might not have an effect on the signal, there are other factors that can limit the performance. For transmitter #2 with any interferer and transmitters #5 and #6 for the cases in which either a Bluetooth or ZigBee device was presented as an interferer, the ability of the signal to resist the effects of the interferers produces the same

results as in the previous example in Figure 5.8: there is little deviation from the no-interference setting. The difference, however, can be seen in Figure 5.9, which reflects the performance of Bluetooth interferers on ZigBee transmitter #6. In comparing the two figures, the only variation is the curve around which all the results are concentrated. For Figure 5.9, the performance at 9 dB is only approximately 5×10^{-4} , whereas it was 10^{-5} for the previous figure. For the performance to reach 10^{-5} for Figure 5.9, the SNR would have to be equal to 11 dB. This change in performance can be attributed to the relative characteristics of the propagation paths for the different transmitters. Compared with the first arriving path, it is obvious that transmitters #1 and #3 have more favorable secondary paths than do transmitters #2, #5, and #6.

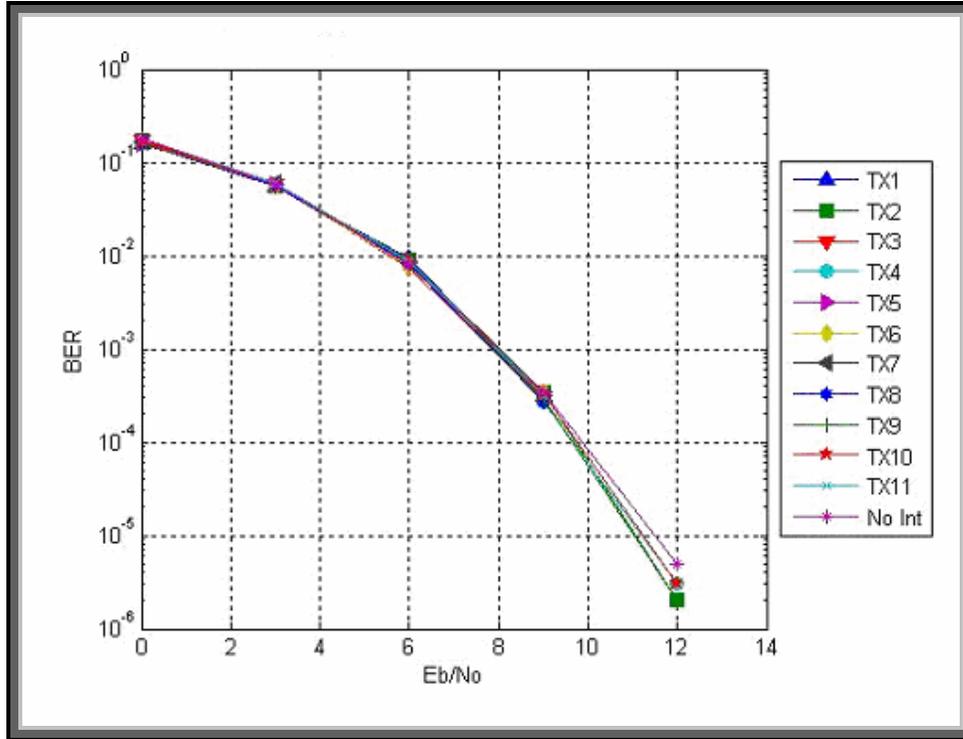


Figure 5.9. ZigBee transmitter #6 with Bluetooth interferers.

Because of the power advantage that WiFi has over ZigBee, radiating at an average power of 17 dBm as opposed to 0 dBm, it follows that WiFi would affect the performance much more than either a Bluetooth or ZigBee interferer. Unlike the previous examples in which the interferers did not have any effect on the transmitted signal, as shown in Figure 5.10, the WiFi interferers do affect the operation of the protocol. The effects of the WiFi interferer on transmitter # 6, as shown in Figure 5.10, and transmitter #5 are very similar. Only interferers #1, #2, and #4 cause the performance to diminish from the no-interference case. As a rule, any performance that results in a BER greater than 10^{-3} will be considered to be undetectable. Therefore, based on the results shown in Figure 5.10, interferers #2 and #4 cripple the performance and make those situations interference-limited. On the other hand, even though interferer #1 limits the performance, the BER still maintains a probability of error less than 10^{-4} , so the signal can be received and demodulated. Interferer #1 does not affect the performance as much as the other two interferers, even though it is closer to the receiver, because of the relative phases of the two signals. Because the phase relation between the received signal of interferer #1 and that of transmitter #6 is on the order of 90° , it minimizes the impact that interferer #1 has on the signal more than it minimizes the impact on the other two interferers. This means that even though interferer #1 has more signal energy than the other two interferers, it is less correlated with the transmitted ZigBee signal, and its effect is limited.

Comparing the results for ZigBee transmitter #5 when WiFi interferers are present with the results shown in Figure 5.10 for transmitter #6, the performances are very similar. When the interferers do not affect the signal, the BER reaches 10^{-5} between SNR values of 10 to 15 dB. However, for transmitter #5, only interferers #2 and #4 cause the performance to be interference-limited and to flatten out. Interferer #2 causes the signal to be unrecognizable by generating too many errors, but interferer #4 does not completely distort the signal, and the performance settles around 10^{-4} , which can be received.

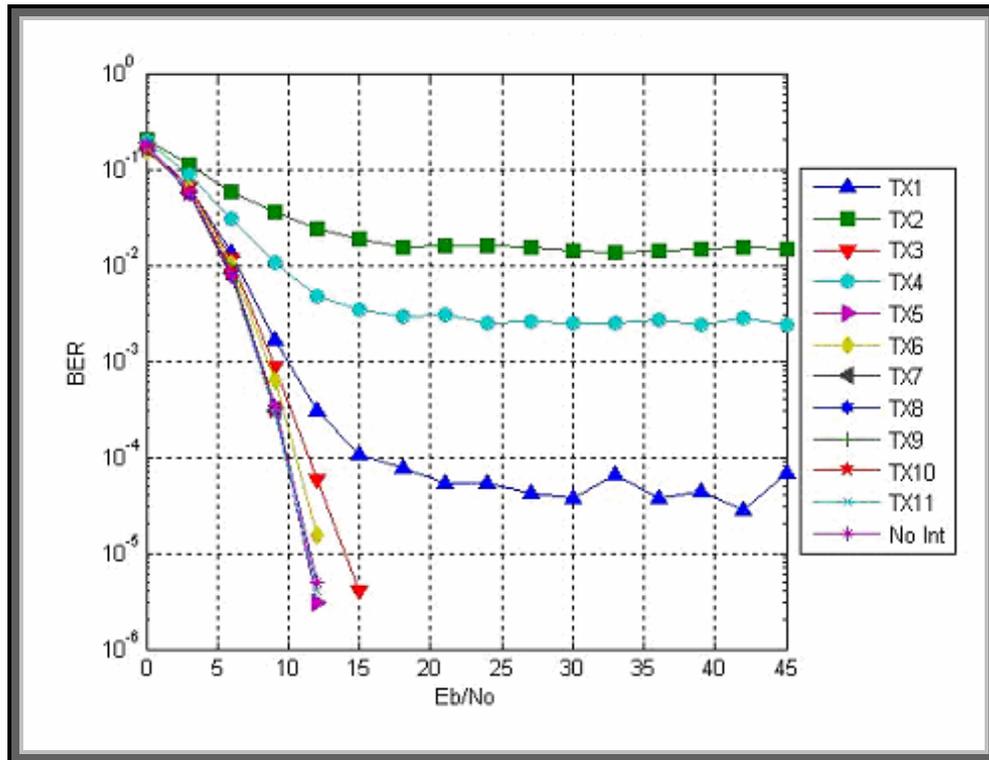


Figure 5.10. ZigBee transmitter #6 with WiFi interferers.

5.3.4 ZigBee NLOS

The performances within the NLOS cases, ZigBee transmitters #7 through #11, differ greatly from those when a LOS path is present. First, the performance of the system when only fading is considered (without interferers) is degraded. For the LOS scenario, the system would reach a BER of 10^{-5} at 9 dB; but now that crossing position has been pushed farther out, and the best performer is ZigBee transmitter #7, in which the crossing point is at 20 dB. Second, the performance has been corrupted because, while the message signal has moved into the NLOS scenario, some of the interfering transmitters are still located within the LOS of the receiver. Thus greater received powers affect the performance.

Looking at ZigBee transmitter #7, the performance when a Bluetooth or a ZigBee interferer is considered stays relatively the same. Figure 5.11—which provides the results obtained for ZigBee transmitter #7 with Bluetooth interferers—shows that the NLOS interferers #7 through #11 do not affect the signal in a substantial manner, and the results appear to be within close proximity to the no-interference case. This is also true when ZigBee interferers are considered. All NLOS interferers except #7 and #8 follow the no-interference case and reach a BER of 10^{-5} at 20 dB. The performance of interferers #7 and #8 appears to settle to a value of 2×10^{-6} . As also shown in Figure 5.11, the performance when any of the interferers are

present still reaches a point less than 10^{-3} . Therefore, all of the signals will be detectable. Interferers #1 and #2 keep the performance above the 10^{-4} level, while interferers #3, #4, and #6 can all maintain a BER nearer to 10^{-5} . Interferer #5 does not cause the signal to be interference-limited and allows the BER to perform as if the NLOS interferers were present. Also, when ZigBee is assumed as the interferer, interferer #5 does not affect the performance, but the outcome of the other LOS interferers is a great deal different. First, interferers #1 and #2 cause the signal to be no longer detectable because they cause the BER to rise to 10^{-2} . Similarly, the performance with interferers #3, #4, and #6 also rises but still maintains a value near but less than 10^{-3} .

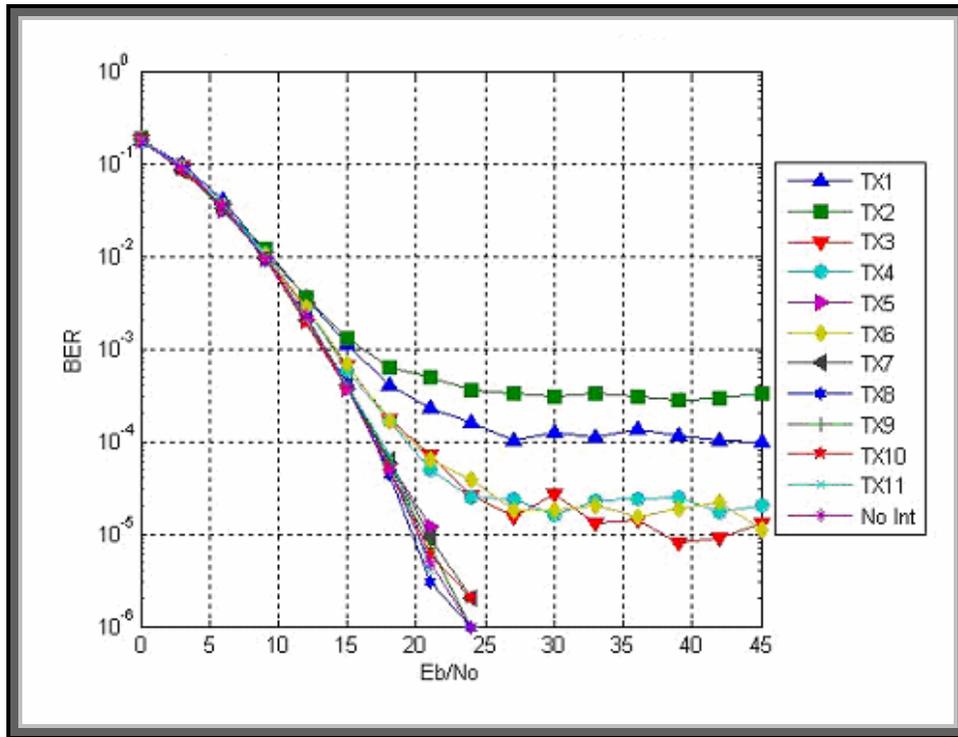


Figure 5.11. ZigBee transmitter #7 with Bluetooth interferers.

In looking at WiFi as an interferer for ZigBee transmitter #7 (Figure 5.12), it can be seen that as a result of the substantial increase in transmitted power between the two protocols and the characteristics of the LOS path, each of the interferers located within the B1 room has a considerable effect on the ability of the ZigBee receiver to demodulate the transmitted signal within an acceptable probability of error. The BER ranges from completely destroyed for interferers #1, #2 #3, #4, and #6, with nearly half of all received bits in error, to leveling out slightly above 10^{-4} whenever interferer #5 is present. It should also be noted that both the #7 and #8 interferers cause the BER to drop to near 10^{-2} , meaning that they destroy the signal. Therefore, it can be seen that among Bluetooth, WiFi, and ZigBee, WiFi has the most substantial effect on the interference, pushing the number of devastating interferers from zero for Bluetooth to two for ZigBee and finally to seven for WiFi.

The final ZigBee example that warrants further examination is that of ZigBee transmitter #10. For each of the three groups of interferers, the BER curves are nearly identical to the ones shown in Figure 5.13, which include the presence of ZigBee interferers. This curve is interesting because as a result of all the interferers in and surrounding the B1 room (#1, #2, #3, #4, #5, #6, #7, and #8), the message is unrecognizable, and approximately half of the bits are received in error.

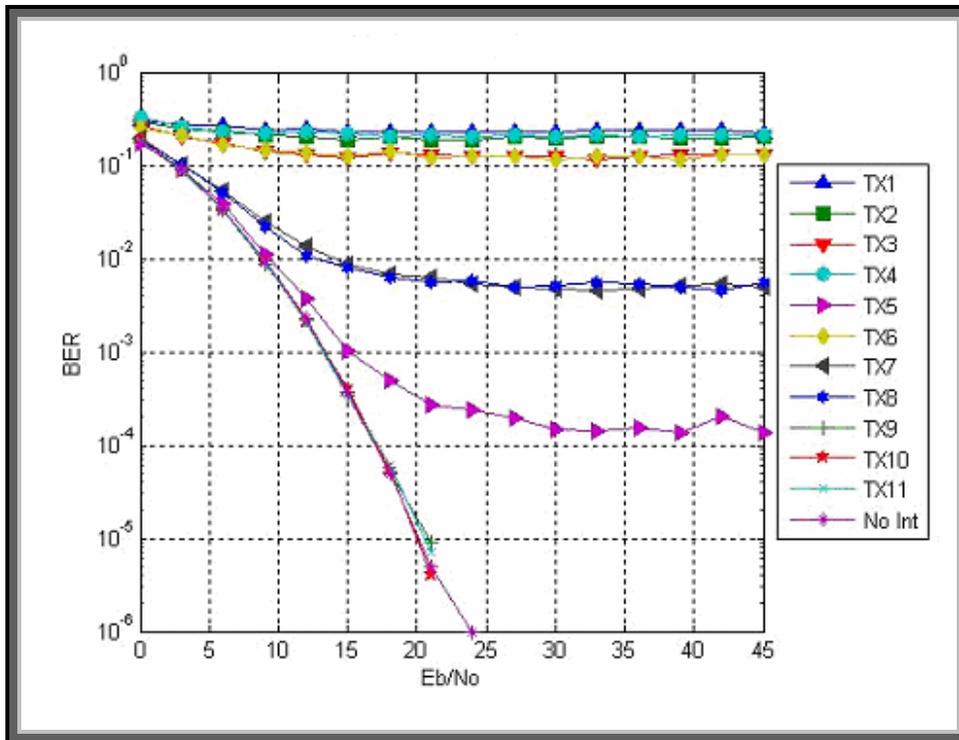


Figure 5.12. ZigBee transmitter #7 with WiFi interferers.

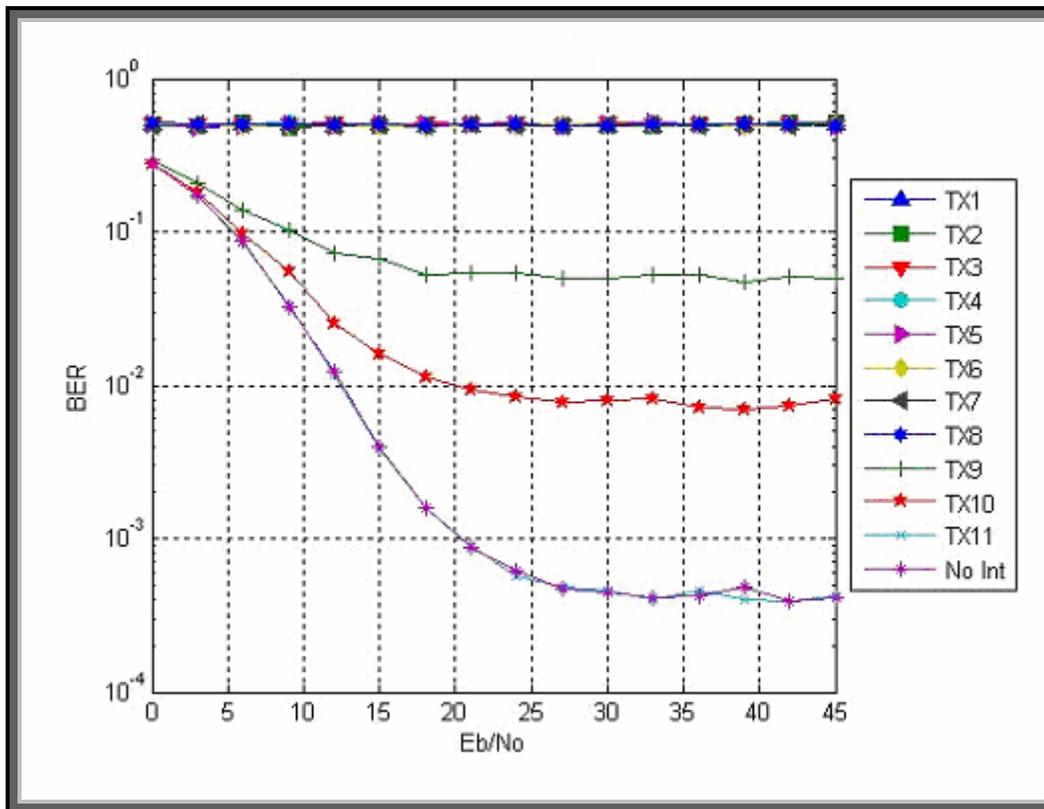


Figure 5.13. ZigBee transmitter #10 with WiFi interferers.

Conversely, for the interferers that are located a considerable distance away from the receiver (#9, #10, and #11), more on the order of the distance of the ZigBee transmitter, the performance corresponds to that when no interferers were used, meaning that those interferers have little effect on the reception of the transmitted signal. There are two important details to observe. The first is that under the no-interference scenario, the performance never dropped below 4×10^{-4} . The second is the locations of interferers #9 and #10. The first feature can be attributed to the characteristics of the propagation paths for the transmitter. The locations of the ninth and tenth interferers should be noted because, for the cases in which Bluetooth or ZigBee interferers are used, their performance is comparable to cases in which no interference is considered. However, their locations change when the WiFi interferers are present, as seen in Figure 5.10. The level degenerates to an undetectable message for both interferers to a level near or below 10^{-2} , which can be attributed to the much higher transmitter power of WiFi.

As a result, it appears that for ZigBee, Bluetooth is the least disturbing protocol. The reason is that, except for the NLOS transmitters with the LOS interferers, in the presence of Bluetooth, the performance of the ZigBee transmitters does not tend to vary from cases when no interferers are considered. Much like Bluetooth, ZigBee does not infringe upon the performance of another ZigBee device, with a few exceptions. However, based on the results of transmitter #7, it appears that ZigBee interferers affected the performance slightly more than did the Bluetooth interferers; therefore, ZigBee is slightly more destructive than Bluetooth. WiFi, on the other hand, seemed to have the most influence on the performance of the ZigBee devices; therefore, WiFi appears to be the most damaging of the three protocols its use with a ZigBee device is considered. This fact can be attributed to WiFi's having a high radiated power and occupying a large bandwidth.

5.3.5 WiFi LOS

For the WiFi transmitter, the effects from the interferers should decrease dramatically because of the increase of radiated power from 0 dBm for ZigBee to an aggregate 17 dBm for WiFi. The improved performance is most prevalent for the LOS transmitters, especially in dealing with either the Bluetooth or ZigBee interferers. Because of the 13- to 17-dBm power advantage, the ZigBee and Bluetooth interferers do not have a profound impact on the operation of the WiFi receiver. The cases in which the WiFi signal can be detected are for transmitters #1, #3, #5, and #6. When Bluetooth and ZigBee interferers are involved, and the only transmitters WiFi interferers are #1 and #3, the performance looks very similar to that shown in Figure 5.14. This figure shows the effects of ZigBee interferers on WiFi transmitter #5. For each interferer, the performance crosses the 10^{-4} threshold between 10 and 15 dB, with the performance of the nearer transmitters pushing toward 10 dB and the transmitter farther from the receiver reaching the threshold nearer to 15 dB. In some cases the performance from interferers #1 and #3 pushes the curve out slightly farther.

The results are considerably different when a WiFi interferer is introduced into the system for transmitters #5 and #6. The change can be seen in Figure 5.15, which shows the effects of the WiFi interferers on the #6 WiFi transmitter. While interferers #3, #5, and #7 through #11 behave in a manner similar to cases when a Bluetooth or ZigBee interferer is used, with the results approximating those of the no-interferer scenario, the #1, #2, and #4 interferers greatly affect the performance. The #2 interferer completely disrupts the WiFi signal by making half of the signal unrecognizable, while the interference from #4 causes the BER to lie between 10^{-2} and 10^{-3} . The effect of interferer #1 is that it produces a signal that might or might not be acceptable because the BER settles near the 10^{-3} threshold. Interferer #6 pushes the BER curve so that the probability of error does not consistently stay below 10^{-5} until the transmitter reaches an SNR value of 40 dB, an increase of more than 20 dB from the situation in which no interference is considered.

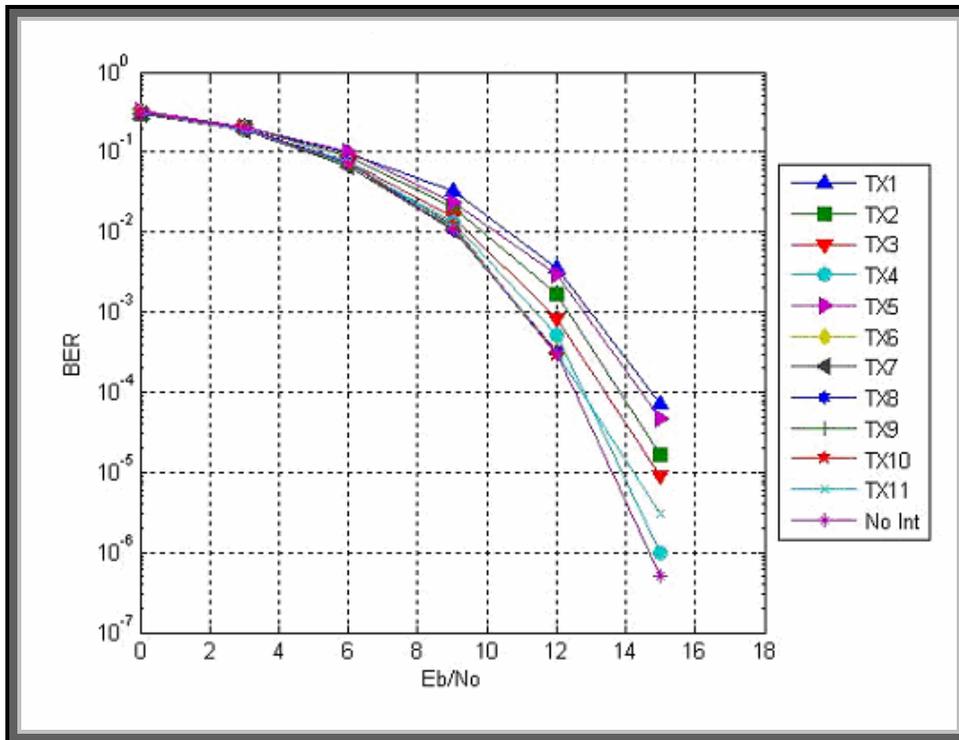


Figure 5.14. WiFi transmitter #5 with ZigBee interferers.

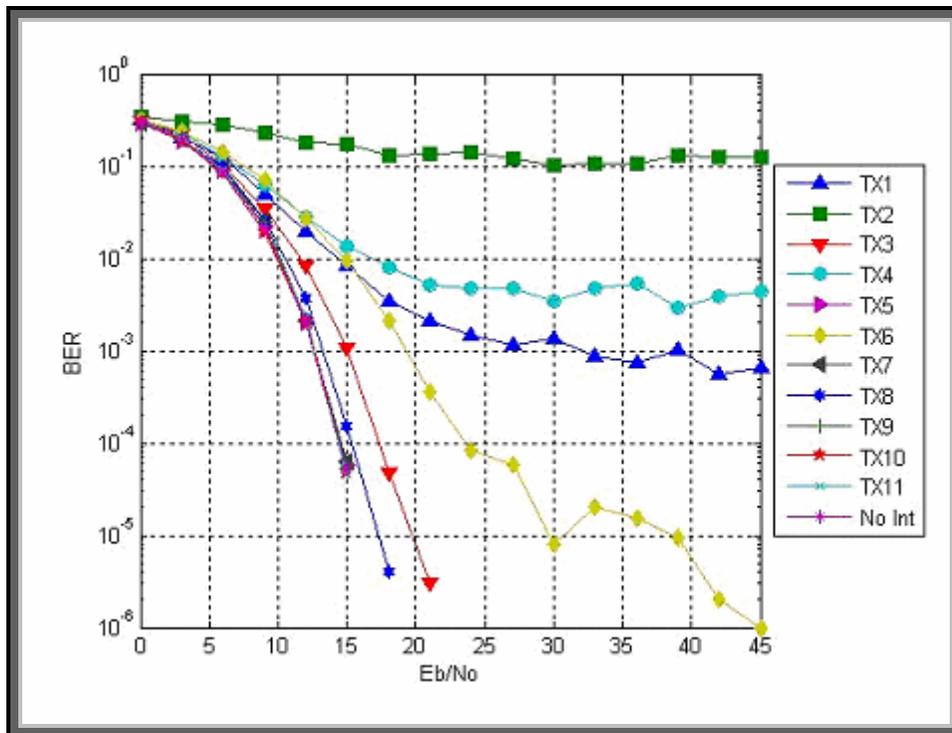


Figure 5.15. WiFi transmitter #6 with WiFi interferers.

Results similar to those found in Figure 5.15, which deals with WiFi transmitter #6, are prevalent for WiFi transmitter #5 when the presence of WiFi interferers is considered. The difference is that because WiFi transmitter #5 is closer to the receiver, not as many interferers affect the performance. In this case, only the first two interferers affect the performance in a significant way, but they both cause the BER to stay above the acceptance threshold and cause the signal to be undetectable. In the presence of all other interferers, the performance maintains a probability of error below 10^{-6} for an SNR above 20 dB. The change in performance caused by the WiFi interferer is attributable to the fact that WiFi has a wide bandwidth, equal to 22 MHz, nearly a third of the 80-MHz-wide ISM band. On average both WiFi transmitters could be transmitting on the same channel if the selection of the channel was not based on availability and was completely random. The change is also a result of the fact that even though WiFi has a power advantage over the other two protocols, when a WiFi interferer is considered, the transmitter and interferer are transmitting at the same radiated power. Therefore, when two WiFi transmitters are placed near each other, without the enhancements made through the upper network layers such as CSMA/CA and channel assignment, the two devices might not be able to coexist.

Unlike ZigBee transmitter #2, WiFi transmitter #2 does not produce an AWGN-type BER curve. As can be seen in Figure 5.16, which shows the performance of transmitter #2 with Bluetooth interferers, the system's performance levels off near a BER of 10^{-5} . This performance can be attributed to the propagation characteristics of the path between the transmitter and receiver. ZigBee, with its superior coding gain, is able to overcome the effects of fading. WiFi could not combat the fading as effectively. As shown in the figure, all LOS interferers except #4 and #6 cause the performance to be unacceptable and reside near a BER of 10^{-2} . All other interferers slightly affect the performance, and thus they are concentrated between the 10^{-4} and 10^{-5} error levels.

When a switch is made from Bluetooth interferers to ZigBee interferers, the transmitter performance stays relatively uniform. This also holds true for the case in which WiFi interferers are used, except that a third trend line emerges with performance hovering just below the 10^{-3} threshold. The two interferers located on this new line are #1 and #6; therefore, the performance of WiFi transmitter #2 is better considering interferer #1 of WiFi than when considering interferer #1 on the previous two standards. The results yield a performance that is acceptable in the presence of the WiFi interferer but not in the presence of a Bluetooth or ZigBee interferer. The opposite is true the effects of the #6 interferer are considered. Even though the #6 interferer does not cause the performance to become unacceptable when the WiFi interferer is presented, it does cause the performance to be inferior to that produced when a Bluetooth or ZigBee interferer is considered.

5.3.6 WiFi NLOS

The performance of transmitter #2 is not the only difference between the performance of the two protocols, ZigBee and WiFi. When the NLOS scenarios for ZigBee are considered, the only two transmitters that produce acceptable results are #7 and #10. Upon inspection of the NLOS results for WiFi, the #7 transmitter no longer can be used. However, transmitters #9, #10, and #11 all produce acceptable results that could be detected in a receiver. The trends among the three transmitters are fairly universal for the cases of WiFi and ZigBee; the interferers either cause the receiver to be unable to demodulate the signal, or the performance roughly mirrors that of the no-interference situation. The same cannot be said about the performance with the Bluetooth interferers. In that case, the BER curves are spread out between the no-interference case and a point at which the signal is unrecognizable.

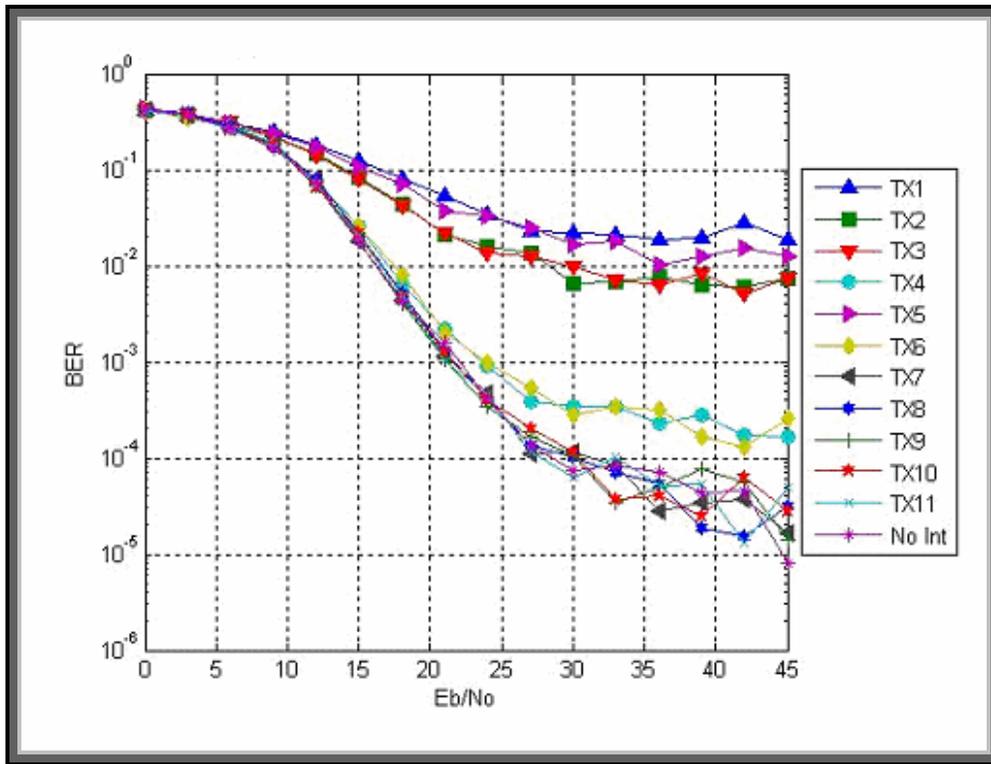


Figure 5.16. WiFi transmitter #2 with Bluetooth interferers.

This spreading of the BER curves can be seen by looking at WiFi transmitter #9 in Figure 5.17, which illustrates the effects resulting from the Bluetooth interferer. The performance varies from the NLOS interferers being situated near the no-interference case and maintaining a BER near 3×10^{-5} at 35 dB, to the LOS interferers being spread from interferer #5 and having a probability of error of 10^{-4} at 35 dB. From there, interferers #1, #2, and #4 all settle near the acceptable performance threshold for SNR values above 30 dB. The only two interferers that cause the performance to be unacceptable are the #3 and #6 interferers.

The trends found in Figure 5.17 are very different from what can be seen for the other two interferers of WiFi and ZigBee. For the condition in which ZigBee is the interferer, the only interferers that yield an acceptable performance are #9, #10, and #11; they all mirror the no-interference case. All other interferers distort the signal and cause nearly half of the bits to be in error. The same trend is found when WiFi interferers are being used; the only interferers not causing half of the bits to be in error are #9 through #11. For the case of WiFi, these interferers do not follow the no-interference case because the #10 and #11 interferers settle to a BER slightly above 10^{-4} , and the #9 interferer does not allow the BER to reach 10^{-2} ; therefore, the effects of this interferer destroy the signal.

Nearly identical performance to that of transmitter #9 can be seen in Figure 5.18, which depicts WiFi transmitter #10 with Bluetooth interferers. The difference is that instead of only two interferers causing the signal to become distorted beyond recognition, five of the LOS interferers damage the signal. The lone LOS interferer that does not significantly corrupt the signal is #4. However, its performance is limited to 3×10^{-4} at 30 dB, which is greater than the values between 5×10^{-6} and 5×10^{-5} at an SNR of 35 dB for the NLOS interferers and the situation in which no interference is present.

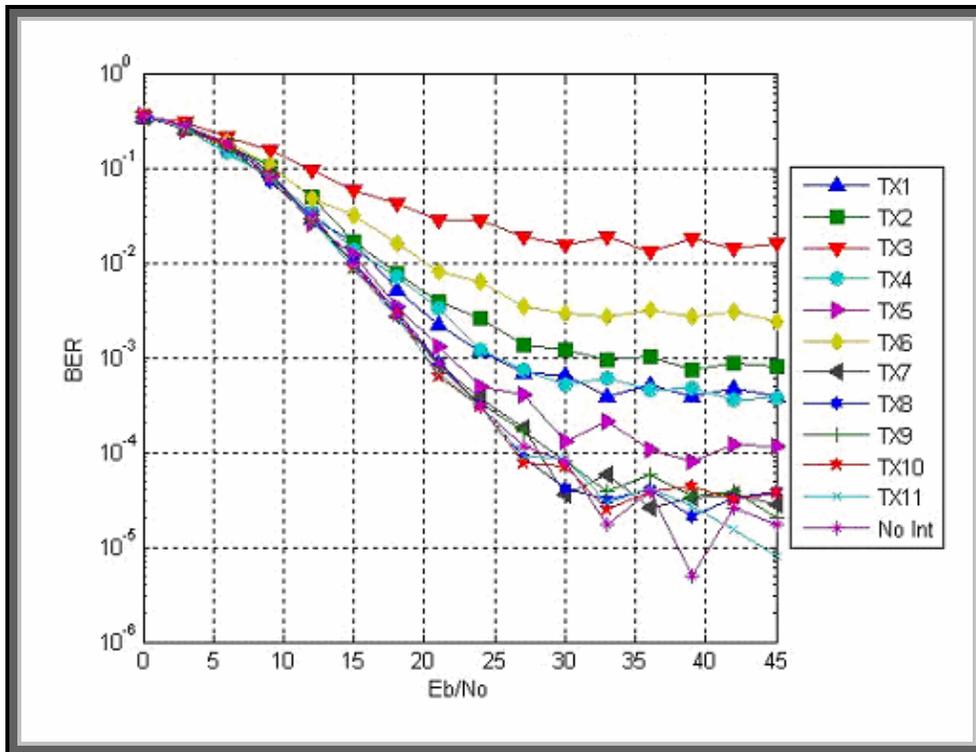


Figure 5.17. WiFi transmitter #9 with Bluetooth interferers.

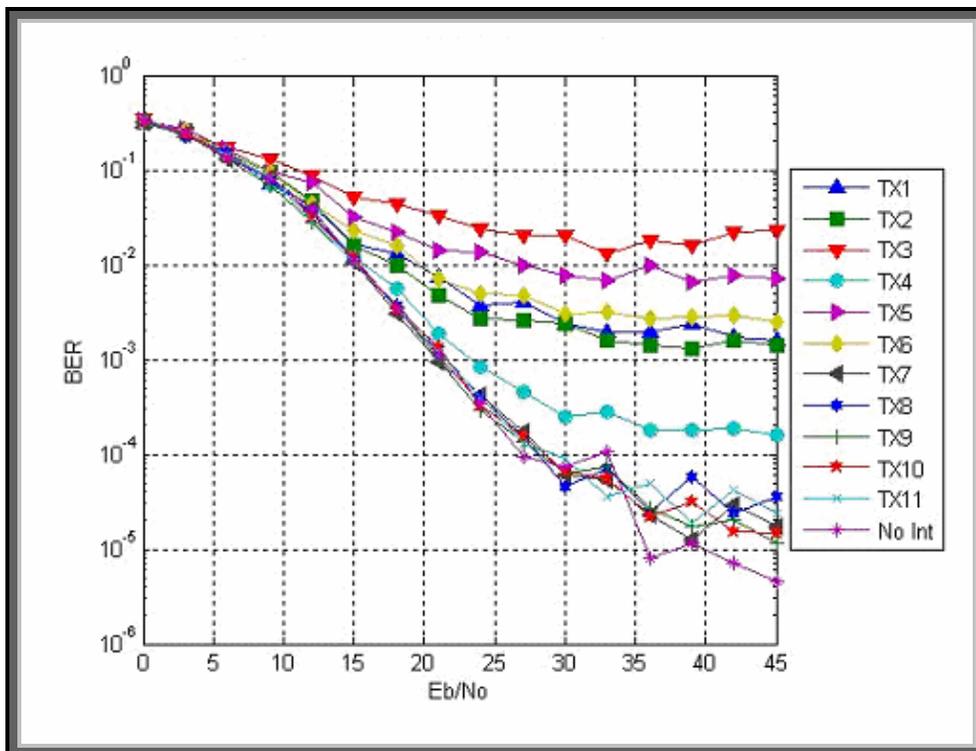


Figure 5.18. WiFi transmitter #10 with Bluetooth interferers.

Again, taking into account the effects of the ZigBee and WiFi interferers, the performances are very similar. For ZigBee, the LOS interferers all completely distort the signal, as do the #7 and #8 NLOS interferers. Just as before, the #9 through #11 interferers do not affect the signal, and their performances are comparable to cases when no interferers are used. Differences appear, however, in moving toward the WiFi interferers. When transmitter #9 is being used, two of the interferers still allow an acceptable signal to be received. When transmitter #10 is used, however, none of the interferers allow for signal to be detected, and they all completely damage the signal.

The results for the final WiFi transmitter, #11, can be found in Figure 5.19, which shows the effects of including a WiFi interferer. The first noticeable difference between these results and those for the previous two transmitters is that the line depicting no interference shows that the BER curve resembles the performance of a noise-limited system and continues to slope downward, whereas the previous two transmitters appeared to flatten out to a constant level. This phenomenon can be attributed to the unique propagation characteristics of this particular transmitter. From the figure, it can also be seen that the only interferers that do not significantly affect the signal are #7, #8, and #9. When transmitter #11 is the transmitter farthest from the receiver, the effects of interferers #10 and #11 can be noticed because their power level is now within the same range as the radiated power from transmitter #11. Therefore, when coupled with the relative phases, it appears that interferers #10 and #11 impair the signal. Conversely, interferers #7 and #8 do not impair the performance; they follow the same trend line as the no-interference scenario. Interferer #9, however, causes the BER to become interference-limited and to level off near a BER slightly less than 10^{-5} .

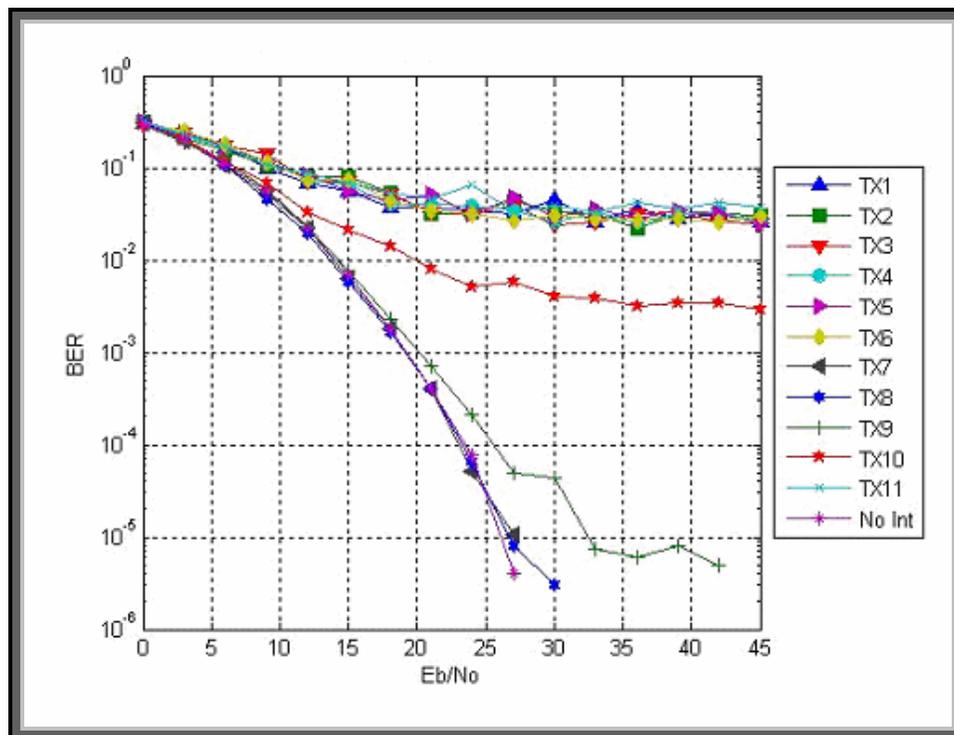


Figure 5.19. WiFi transmitter #11 with WiFi interferers.

For transmitter #11 with a Bluetooth interferer, the performance is spread between the LOS interferer's performance and that of the NLOS interferers. As before, the NLOS interferers do not cause the performance to vary considerably from the no-interference scenario. The LOS interferers, however, cause

the performance to fluctuate near the 10^{-3} threshold for the #1, #2, #4, and #5 interferers, while the #3 and #6 interferers cause the signal to be somewhat destroyed and unacceptable. When a ZigBee interferer is considered, the results are pushed to either one extreme or the other, and half of all bits are received in error for the LOS interferers as well as for two of the NLOS interferers, #7 and #8. On the other hand, the performance mirrors that of the no-interference situation whenever the other NLOS interferers, #9, #10, and #11, are considered.

Based on these results, the overall trend for WiFi is very similar to that for ZigBee when the effects of the interferers are considered. As before, Bluetooth appears to be less intrusive upon WiFi than the other two interferers. The difference is that for the NLOS transmitters of ZigBee, the effects of Bluetooth typically do not cause the signal to be completely distorted.

When the NLOS transmitters of WiFi are considered, however, that is not the case. The LOS interferers considerably affect the performance, sometimes allowing the signal to still be detectable and at other times causing the performance to go slightly above the threshold of acceptable performance. The NLOS interferers, on the other hand, do not affect the performance at all, allowing the signal to be received as if no interferers were present.

5.3.7 Bluetooth LOS

Unlike the previous two protocols, which used the advantage of spreading a signal over a wide bandwidth, Bluetooth uses a small bandwidth and jumps from one frequency to another. This approach aids in its ability to minimize interference by simply avoiding the interferers, because the chances are small that the same interferer will interfere with consecutive hops. Whenever Bluetooth has a clear LOS path to the receiver, the other protocols do not have an effect on the system at all. This fact is evident in all usable Bluetooth transmitters, #1, #2, #3, #5, and #6, for all cases when the interference associated with Bluetooth, WiFi and ZigBee interferers is considered. The only differences in performance can be found in the changes in the no-interference case across the five different transmitters.

Figure 5.20 shows the results obtained from Bluetooth transmitter #1 with ZigBee interferers. From this plot, it can be seen that the interferers do not affect the signal in a profound way; but inspection of the SNR when the BER crosses the 10^{-5} error level shows the SNR is equal to 10 dB. In going from Figure 5.20 to Figure 5.21, which shows the effects that WiFi interferers have on the signal of Bluetooth transmitter #3, the interferers have no impact. The curve does not reach 10^{-5} until an SNR of 11 dB. This performance is the same as that for all the other interferers of both transmitters #3 and #2; but for this transmitter, the SNR must exceed 12 dB to exhibit the same performance. When the transmitters farther from the receiver are used, the performance maintains the trend of pushing the curve to the right, which in turn means more energy is required to maintain the same performance. An inspection of the results of transmitter #6 with Bluetooth interferers in Figure 5.22 shows that the SNR must now be equal to 14 dB to achieve the 10^{-5} error rate, which is true for the other interferers as well. The results for transmitter #5 are very similar, except the SNR needs to reach only 13 dB for the 10^{-5} BER.

5.3.8 Bluetooth NLOS

Once the Bluetooth transmitters are placed outside the B1 room, the performance varies greatly because of the interferers. Three different transmitters perform to at least a minimum acceptable standard, which is that the instance in which no interferers are considered must at least surpass the 10^{-3} probability-of-error threshold. The three transmitters that qualify under this standard are the #7, #10, and #11. Of these, the #7 and #11 transmitters reach a BER of less than 10^{-5} , while the #10 transmitter's performance goes beyond only the 10^{-3} level. Even though the actual values vary for these three transmitters, the tendencies for each can still be classified into three different categories when the interferers are considered. These categories

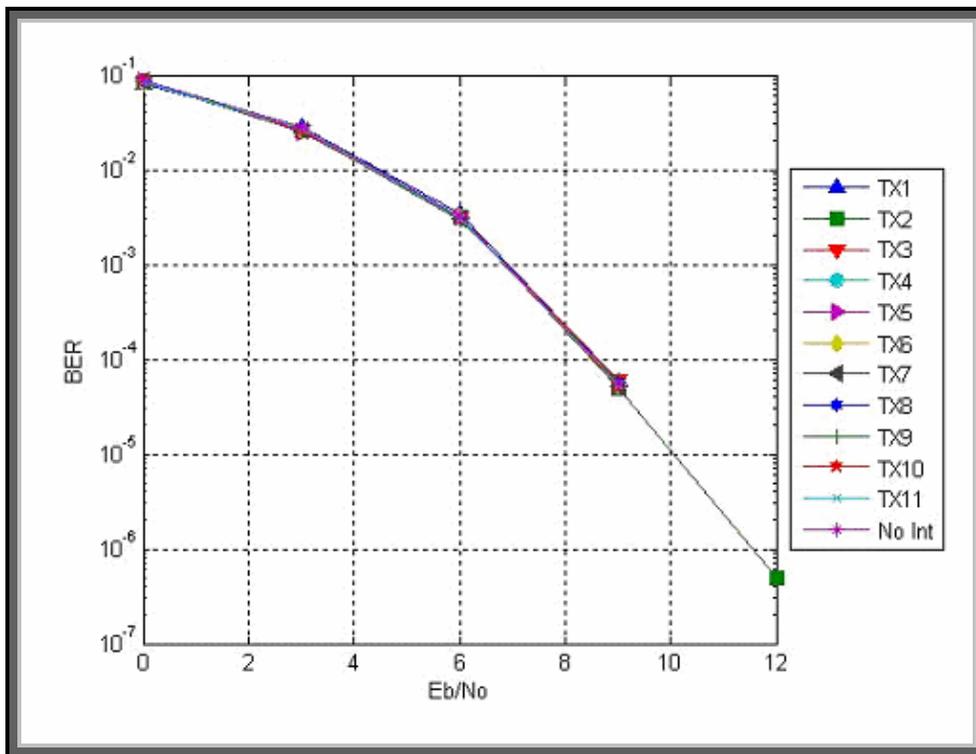


Figure 5.20. Bluetooth transmitter #1 with ZigBee interferers.

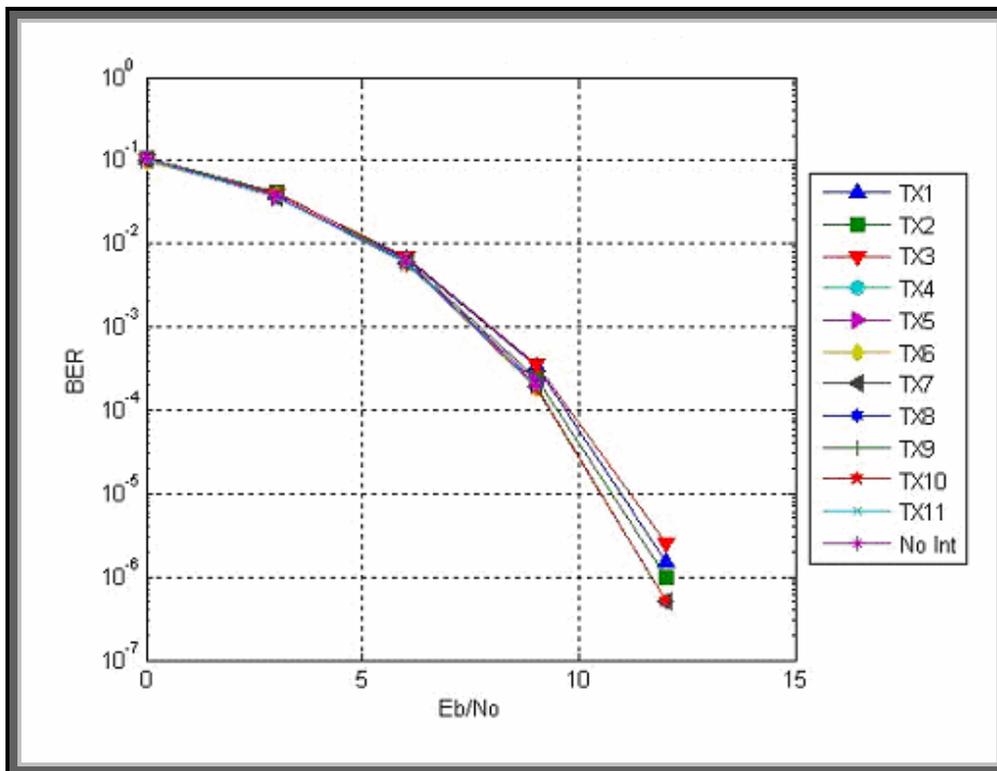


Figure 5.21. Bluetooth transmitter #3 with WiFi interferers.

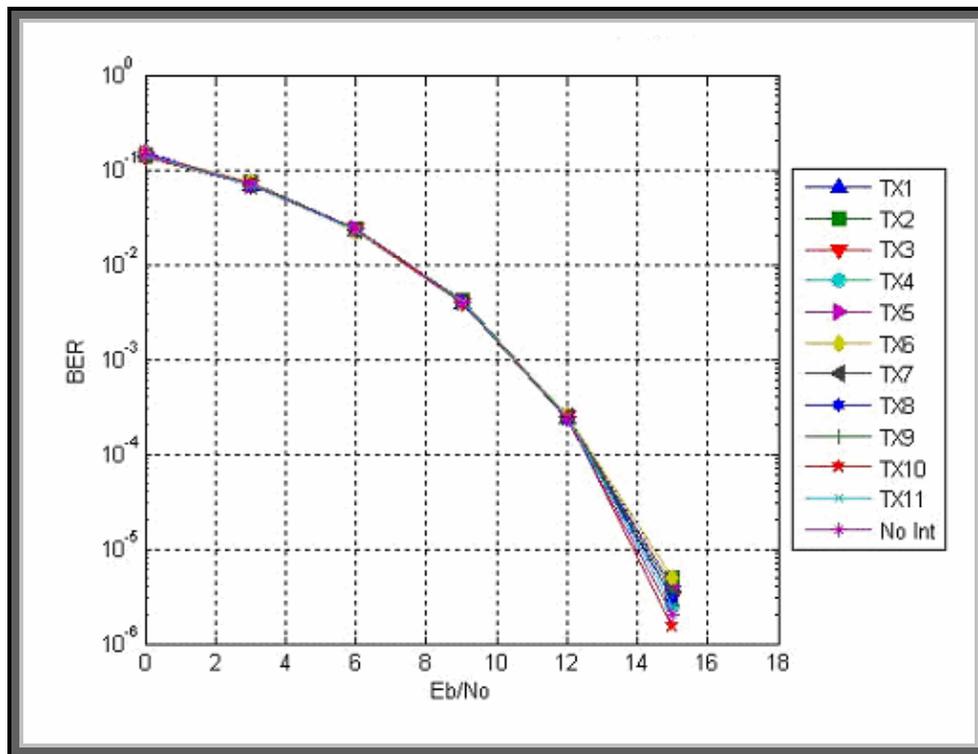


Figure 5.22. Bluetooth transmitter #6 with Bluetooth interferer.

are when the interferers (1) cause the signal to be undetectable, (2) do not affect the signal, and (3) are somewhere between the previous two levels. In general, the interferers with a short LOS path to the receiver, #1 through #4, can be placed in the first category because those interferers have a dominant path, and that path is received with a much higher total power than the signals received from outside the room. The second category of interferers consists of those that are not placed within close range of the receiver. These interferers do not contain signal powers that can overwhelm the transmitted power of the Bluetooth signal because the propagation loss is too great. This category, therefore, consists of interferers #9, #10, and #11.

The preceding two scenarios represent the possible extremes. There is a middle ground, however, that includes the effects of the interferers that are neither close to nor too far away from the receiver. The performance of these interferers is fairly unpredictable. This group consists of both LOS and NLOS interferers, #5 through #8. The phase of the signals becomes more critical than the power because these interferers have power levels similar to those of the NLOS Bluetooth transmitters. Therefore, it matters if the signal and interferer have relatively close phases or if they are 180° out of phase.

Figure 5.23 illustrates these three categories by showing how the different WiFi interferers affect Bluetooth transmitter #7. As expected, interferers #1 through #4 keep the BER near 10⁻³, the threshold between detectable and undetectable signals. Also, interferers #9, #10, and #11 do not affect the signal in a significant way, as was predicted. Moving on to the four interferers that were questionable, it appears that the two LOS interferers, #5 and #6, affect the signal, but still allow the BER to fluctuate around 10⁻⁴. The #7 and #8 interferers did not impinge on the signal, and the performance parallels that of the other NLOS interferers, #9 through #11.

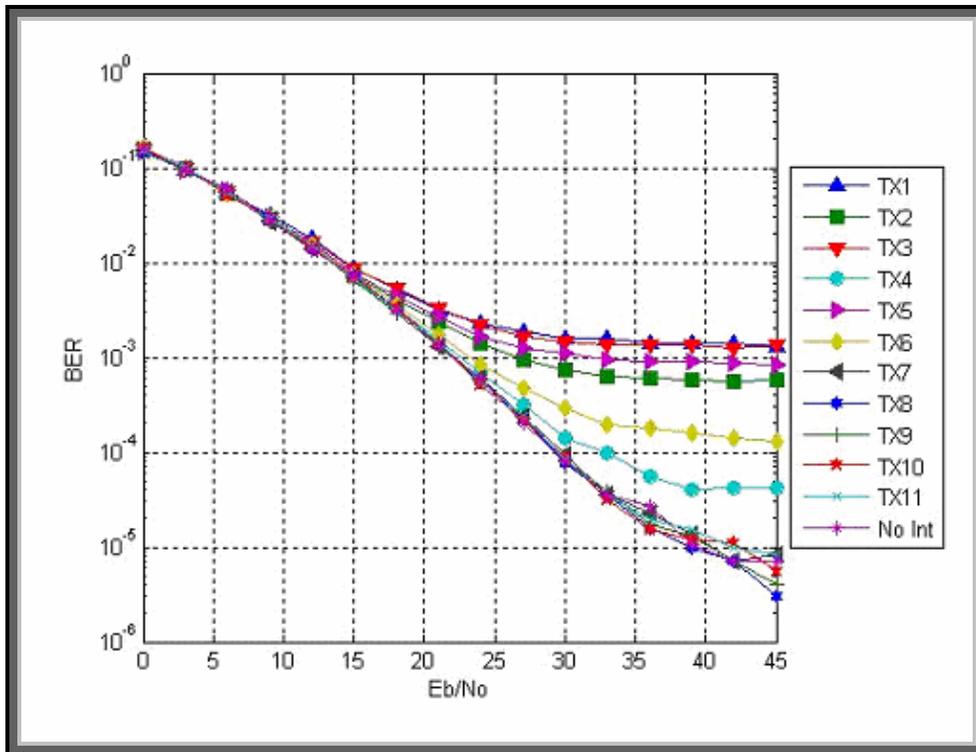


Figure 5.23. Bluetooth transmitter #7 with WiFi interferers.

Comparing the results of the WiFi interference case of Bluetooth transmitter #7 with those of the other interferers, similar results are obtained for the Bluetooth interferers, although the results for ZigBee are very different. For the ZigBee interferers, both the second- and third-category interferers perform as they should. The second group, the #9, #10, and #11 interferers, does not affect the signal. The third group is unpredictable, as expected, because none of them appears to significantly influence the signal. The LOS interferers, however, do not perform as expected. The only two interferers that change the BER of the signal are #1 and #2. Although their effects are not drastic, they cause the BER to reach only 2×10^{-4} rather than 10^{-5} ; the other signals reached the latter at SNR values above 40 dB. When the effects of a Bluetooth interferer upon transmitter #7 are considered, the results are as expected. Interferers #9 through #11 do not affect the performance, interferers #1, #3, #4, and #6 all cause the BER to settle to a value around 5×10^{-4} , interferer #2 completely destroys the signal, and interferers #7 and #8 do not allow the BER to reach 10^{-5} but keep it near 3×10^{-5} .

When a transmitter farther away is considered, the performance of the system as a whole changes, but the contribution due to each interferer is approximately the same. Looking at the results obtained in Figure 5.24—which illustrates a case when a ZigBee interferer attacks the signal of Bluetooth transmitter #10—the results are almost an exact replica of those for the other two interferers. Interferers #9 through #11 do not affect the performance, and they settle to a value of 4×10^{-4} for SNR values above 30 dB, as does the no-interference case. All other interferers cause the BER to be above the 10^{-3} threshold, making the signals meaningless. The only difference is that when interferer #5 is used as a WiFi interferer, it has a BER hovering around the 10^{-3} threshold; therefore, the signal might or might not be acceptable depending on the application.

Disregarding the effects of interferer #5 in the single instance from the previous results, the effects of the interferers on the results for transmitter #11 are the same as for transmitter #10. The no-interference situations are completely different because transmitter #11 continues to decrease, while transmitter #10

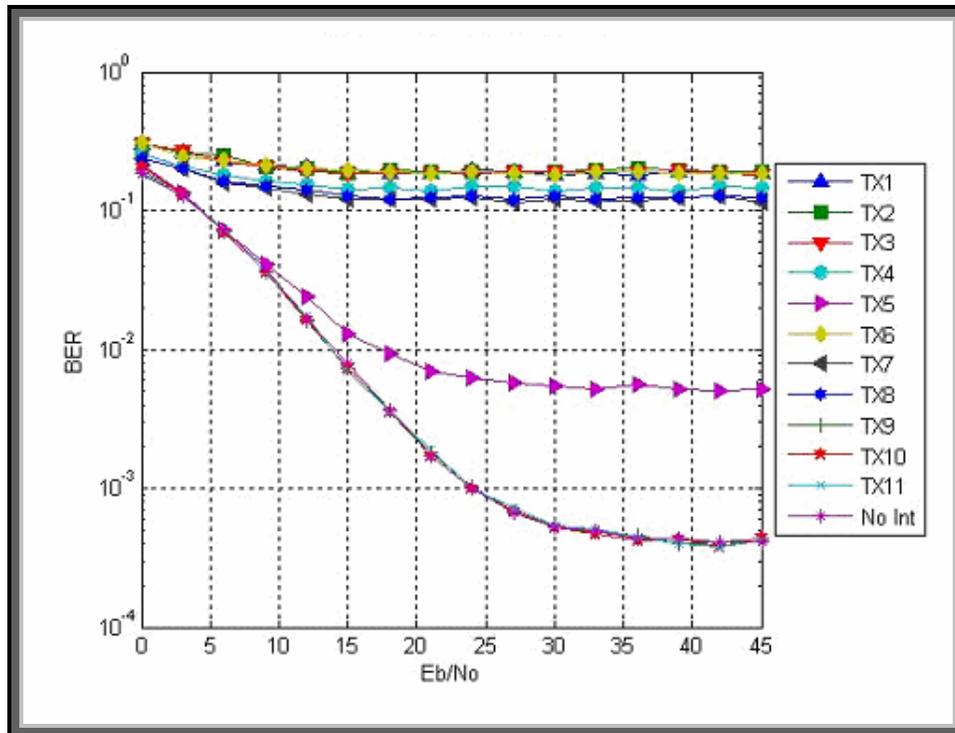


Figure 5.24. Bluetooth transmitter #10 with ZigBee interferers.

flattens out. All of the interferers, #1 through #8, cause the signal to be irreconcilable. For interferers #9 through #11, the performance mirrors that of the no-interference case when the interferers are considered to be WiFi and ZigBee interferers. The performance, as depicted in Figure 5.25, changes, however, when Bluetooth interferers are used. As can be seen in this figure, the three interferers have different effects on the transmitted signal of Bluetooth transmitter #11. Bluetooth interferer #9 causes the performance to be equal to 10^{-4} for SNR values above 35 dB. Interferer #10 appears to level off to a BER of 10^{-6} at an SNR of 45 dB, while the effects of interferer #11 allow the performance to continue to decrease past a probability of error of 10^{-6} , very near the case of no interference.

Based on the results of the case when Bluetooth is the transmitted signal, Bluetooth appears to be the most well-equipped protocol for resisting interference from other devices. This is at least true for the LOS case because none of the interferers were able to affect the performance of Bluetooth. However, the performance in an NLOS case fell in line with the performance of the previous two devices. When the interfering transmitter was located in proximity to the receiver, the signal was damaged; but as long as the interferers were located at a distance comparable to that of the Bluetooth transmitter, they did not affect the signal.

One distinction from the previous two interferers, though, is that Bluetooth in the presence of another Bluetooth device seems to be a much harsher interferer than is the case for the other two protocols. Now instead of the Bluetooth interferer hardly affecting the performance, it affects the performance as much as, if not more than, the other two interferers.

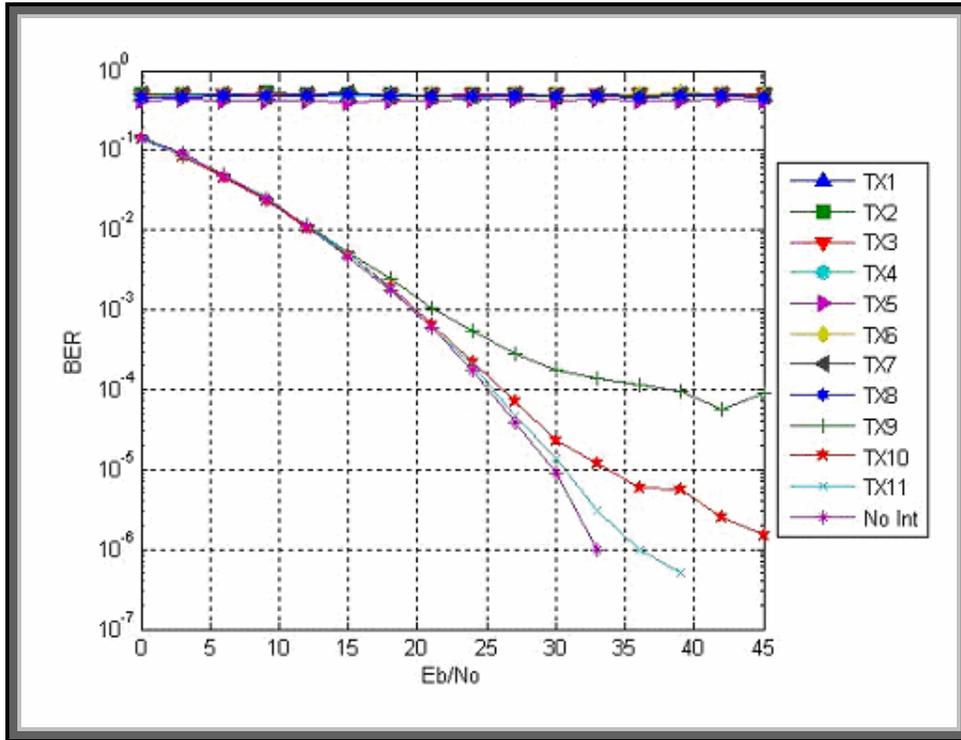


Figure 5.25. Bluetooth transmitter #11 with Bluetooth interferer.

6. SUMMARY

For this study, the primary goal was to develop a baseline simulation environment to examine the coexistence of the WiFi, Bluetooth, and Zigbee wireless standards. We focused on the physical layer issues, where the BER was presented as a function of E_b/N_0 for various interference conditions from other devices. We adopted related parameters from the standards. Note that our simulation models can be adapted, integrated or upgraded for present or future systems.

For the AWGN and Rayleigh fading of the general channel models, the relative received powers of the three protocols were assumed to be equal, allowing for variation of the interferer power level to determine the range of operability. On the other hand, in the site-specific channel model, the transmitted power levels of both the signal and the interferer were fixed according to their respective operating points set by the standard. For the general channel models, therefore, the interferer power is the variable. In the site-specific model, however, the propagation paths between the transmitter and receiver are the variables, so the characteristics of the physical environment influence the performance. Under these settings, the following inference can be summarized from our simulation results.

The effects of Bluetooth as an interferer are considered first, and then followed by the response of Bluetooth to interfering devices. Throughout all three channel models, the two generalized forms and the site-specific model, Bluetooth intrudes upon the performance of the other devices the least. This means that Bluetooth is a “good neighbor” and allows for the coexistence of devices within this frequency band. This fact can be attributed not only to Bluetooth’s transmitting at a relatively low power level, 4 dBm, but also to the frequency-hopping nature of the scheme Bluetooth employs. On average, most hops will not lie within the bandwidth of another signal. Even when these hops do occur within the wide bandwidth of WiFi, the nature of the narrowband signal that Bluetooth encompasses compared with the wideband of WiFi minimizes the effects when the WiFi signal is decoded.

For the AWGN channel, the performance of the Bluetooth system in the presence of interferers is as good as that of ZigBee and better than that of WiFi because of the minimal chance that the Bluetooth transmitter and the interferer are located within the same frequency space. However, the performance dips when the general Rayleigh channel is applied, more so than for ZigBee. Since the same flat-fading effects are assumed to occur to each signal, this difference in performance can be attributed to the coding gain associated with the spreading of the ZigBee signal. In the site-specific environment, Bluetooth shows the highest performance in the presence of a LOS signal or a very strong single-transmission NLOS signal. However, in the other NLOS situations, Bluetooth is unable to overcome the effects of the interferers, and performance is severely limited.

The overall performance of ZigBee in the roles of both interferer and transmitter is very similar to that of Bluetooth, although for different reasons. ZigBee does not seem to interfere with other devices because of its low radiated power level, which is 4 dB below that of Bluetooth and 17 dB lower than that of WiFi. ZigBee also has a relatively small bandwidth, only 2 MHz, as a result of its low data rate; therefore, the effects of ZigBee on a wide signal such as WiFi are also minimized and occur on only 2 of the 79 channels of Bluetooth. As a result, ZigBee does not cause a significant disturbance to the other devices within the site-specific model.

When ZigBee is assumed to be the transmitting device, it is able to defend itself from the effects of both the Bluetooth and ZigBee interferers successfully through the processing gain associated with the spreading of its signal until the point at which the interfering signal powers completely dominate the signal. Alternately, because WiFi begins with such a large power advantage over ZigBee, it inadvertently hinders the performance of the ZigBee signal even when ZigBee has a LOS path to the receiver. Therefore, deployment of ZigBee devices in the presence of WiFi must be undertaken with caution,

which means Zigbee device should not be located near WiFi access points. The actual distance will be case specific, depending on the propagation path characteristics of the environment, but simply moving away from the WiFi access points might be the solution if the interference is caused by the WiFi.

The final protocol, WiFi, has similar results between the general and specific cases in which it is considered as an interfering device. WiFi tends to limit the performance of other devices and does not allow for coexistence. In the general channel models, WiFi limits the performance of other transmitters as a result of its wide bandwidth and the likelihood that these devices will be located within the same frequency space. Moving to the site-specific scenario, this dominance is taken one step further as a result of the increased power advantage that WiFi has over the other protocols. WiFi limits the performance of a transmitting device long before the other interferers do. One reason for this is that WiFi was one of the first standards developed for products within the 2.4-GHz ISM band; therefore, the presence of other devices was not considered. The only goal was to deliver the highest speed, along with the widest possible range.

Comparing the BER results between an AWGN channel and the fading channel, the performance for the simulations incorporating Rayleigh fading are substantially lower than for those incorporating only an AWGN channel. On average, the effects of fading cause a 10- to 15-dB drop in performance. As a result of the 10- to 15-dB decrease, at a 10^{-3} BER the level at which a signal can tolerate another signal has been extended to a 30 dB SNR. WiFi continues to operate along this 15-dB drop from the AWGN case. All interferers still affect the signal in the same way; however, rather than an SIR of 5 dB being needed to completely block out the interferers, the value now must be equal to 20 dB. Also, the signal can now tolerate only interferers with an energy at least 12 dB below its own energy, whereas in the AWGN case the signal could still be received when the interferer had an energy greater than that of the WiFi signal.

ZigBee is less affected by the fading than WiFi. ZigBee follows a 10-dB drop in performance. However, for ZigBee to operate as if no interferers are present, the SIR must be at least 5 dB in the presence of another ZigBee signal, but it can extend up to 0 dB when a Bluetooth signal is incorporated. At BER of at least 10^{-3} , the ZigBee interferer can have a higher power level than the ZigBee signal up to SIR value of -3 dB. This value can be decreased to -10 dB in the presence of a Bluetooth interferer and the performance from the WiFi interferer falls between that of the Zigbee and Bluetooth interferers.

Bluetooth is affected much more than ZigBee and encompasses a 15-dB decrease in performance due to fading. For Bluetooth to operate at a level equal to operation with no interferers, it must have an energy level at least 5 to 15 dB greater than that of the interferers. For a WiFi interferer, an SIR of 15 dB is required. Bluetooth can tolerate signals with values between 3 and -3 dB of its transmitted power, approximately 10 dB less than that of the AWGN situation.

Under the the general channel models, for the same level of interfering signal power for either Bluetooth or Zigbee, WiFi is more vulnerable to interferers. This result is shadowed in the results of the site-specific channel model because of the advantage in power that WiFi has over the interferers. WiFi uses a brute-force technique to obtain data communication and just overpowers all other devices. The only interferer that can hinder the performance of a LOS WiFi transmitter is another WiFi device. WiFi cannot be completely blamed for this fault because at the time WiFi was being developed, mass production of other devices occupying the same frequency range as WiFi had not begun. To achieve a more efficient use of the ISM band, when other devices are produced, they must make a conscious effort to minimize the interference that these devices will cause to other devices, much as Bluetooth and ZigBee do. Otherwise, the entire band will suffer as a result of degradation in the amount of throughput, which is the total bits transmitting through the channel per unit time measured as bits/sec between two points/nodes, because of the amount of time a device must wait before being able to access a clear channel.

Our simulation shows that both Zigbee and Bluetooth can tolerate interference of higher magnitude than their own power. For an AWGN channel, a 5 dB higher value of the interference from the desired signal is acceptable. On the other hand, WiFi can not function well if the interference power is more than the desired user signal. These power requirements can be translated to distance for any particular environment. The relationship between the received power and the tolerable distance depends on the path loss exponent of an environment. For example, assuming the value of the path loss exponent is 4, if the distance doubles, the power drops 16 times (12 dB). Typical value of the path loss exponent varies from 2 to 6. For free space, the path loss exponent is 2 and for the outdoor cellular environment the value is 4. For WLAN applications, the value of the path loss exponent is about 3.2 and which varies with the nature of the environment.

Since the fading characteristics of a channel are random in nature and the influence of the interference is also random, the simulation results show the average BER performance for various setting environments. Our simulation results provide a qualitative guiding for the deployment of these three standards in terms the BER performance. Of course, the simulation tool is capable of providing a quantitative value of the BER for any particular fading environment.

7. REFERENCES

1. K. Werbach, "Radio Revolutions – The Coming Age of Unlicensed Wireless," New America Foundation, Public Knowledge. Washington, D.C.
2. NUREG/CR-6882, Assessment of Wireless Technologies and Their Application at Nuclear Facilities, Oak Ridge National Laboratory, July 2005.
3. Homepage of ZigBee™ Alliance, <http://www.ZigBee.org/>.
4. E. Callaway, P. Gorday, L. Hester, J.A. Gutierrez, M. Neave, B. Heile, and V. Bahl, "Home networking with IEEE 802.15.4: a developing standard for low-rate wireless personal area networks," *IEEE Communication Magazine*, **40**(8), pp. 70–77 (August 2002).
5. J. Notor, A. Caviglia, and G. Levy, *CMOS RFIC Architectures for IEEE 802.15.4 Networks*, white paper, Institute of Electrical and Electronics Engineers, 2003.
6. IEEE Std 802.15.4, *IEEE Standard for Information Technology—Telecommunications and Information Exchange Between Systems—Local and Metropolitan Area Networks—Specific Requirements. Part 15.4: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (LR-WPANs)*, Institute of Electrical and Electronics Engineers, 2003.
7. L. W. Couch, II. *Digital and Analog Communication Systems*, 6th edition, Prentice Hall PTR, Indianapolis, IN, 2001.
8. IEEE Std 802.11b, *Supplement to IEEE Standard for Information Technology—Telecommunications and Information Exchange Between Systems—Local and Metropolitan Area Networks—Specific Requirements. Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: High Speed Physical Layer Extension in the 2.4 GHz Band*, Institute of Electrical and Electronics Engineers, 1999.
9. Vocal Technologies, Ltd., "IEEE 80211b White Paper," October, 27 2003, available at http://www.vocal.com/white_paper/80211b_wp1pdf.pdf.
10. B. Pearson, "Complementary Code Keying Made Simple," white paper, May 2000, available at http://www.eetasia.com/ARTICLES/2001MAY/2001MAY25_NTEK_DSP_AN.PDF.
11. IEEE Std 802.11a, *Supplement to IEEE Standard for Information Technology— Telecommunications and Information Exchange Between Systems—Local and Metropolitan Area Networks—Specific Requirements. Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: High Speed Physical Layer in the 5 GHz Band*, Institute of Electrical and Electronics Engineers, 1999.
12. The Linksys Group, Inc., "A Comparison of 802.11a and 802.11b Wireless LAN Standards," white paper, October 6, 2004.
13. Tech Lab Review Staff, "Bluetooth, ready for the big leagues?" *Tech-Edge*. October 31, 2002, available at http://homepage.mac.com/techedgeezine/networking_bluetooth1.htm.
14. IEEE Std 802.15.1, *IEEE Standard for Information Technology—Telecommunications and Information Exchange Between Systems—Local and Metropolitan Area Networks—Specific Requirements. Part 15.1: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (LR-WPANs)*, Institute of Electrical and Electronics Engineers, 2002.
15. A. C. Davies, "An overview of Bluetooth Wireless Technology™ and some competing LAN Standards," Invited Plenary Lecture, in *Proceedings of the 1st IEEE International Conference on Circuits and Systems for Communications*, St. Petersburg, Russia, pp. 206–211.
16. P. McDermott-Wells, "What Is Bluetooth?" *IEEE Potentials*, **23**(5), pp. 33–35 (December 2004/January 2005).
17. T. Rappaport, *Wireless Communications: Principles & Practices*, Prentice Hall, Inc., Upper Saddle River, NJ, 1996.

18. T. Cooklev., *Wireless Communications Standards, A Study of IEEE 802.11TM, 802.15TM and 802.16TM*, IEEE Press, New York, NY, 2004, p. 225.
19. P. Bhagwat, "Bluetooth: Technology for Short-Range Wireless Apps," *IEEE Internet Computing*, May/June 2001, pp. 96–103.
20. W. H. Tranter, K. S. Shanmugan, T. S. Rappaport, and K. L. Kosbar, *Principles of Communication Systems Simulation with Wireless Applications*, Prentice Hall PTR, Indianapolis, IN, 2004.
21. T. Karhima, A. Silvennoinen, M. Hall, and S.G. Haggman, "IEEE 802.11B/G WLAN Tolerance to Jamming," MILCOM Conference, October 31–November 3, 2004.
22. C. Andren and M. Webster, "CCK Modulation Delivers 11 Mbps for High Rate IEEE 802.11 Extension," Wireless Symposium/Port by Design Conference, Spring 1999, available at http://www.csse.monash.edu.au/courseware/cse5501/reading/CCK_Mod_Delivers_11Mbps.pdf