

ABSTRACT

Title of Document: USING COMMERCIAL RAY TRACING
SOFTWARE TO DRIVE AN ATTENUATOR-
BASED MOBILE WIRELESS TESTBED

Keith Richard Taylor, Master of Science, 2012

Directed By: Professor Joseph Jaja
Department of Electrical and Computer
Engineering

We propose and build a prototype architecture for a laboratory-based mobile wireless testbed that uses highly detailed, site-specific channel models to dynamically configure a many-to-many analog channel emulator. Unlike similar systems that have used abstract channel models with few details from the physical environment, we take advantage of commercial ray tracing software and high-performance hardware to make realistic signal power and characteristics predictions in a highly detailed environment. The ray tracing results are used to program a many-to-many analog channel emulator. Using this system, we can conveniently, repeatedly, and realistically subject real wireless nodes to the effects of mobility. We use our prototype system and a detailed CAD model of the University of Maryland campus to compare field test measurements to measurements made from the same devices in the same physical scenario in the testbed. This thesis presents the design, implementation, and validation phases of the proposed mobile wireless testbed.

USING COMMERCIAL RAY TRACING SOFTWARE TO DRIVE AN
ATTENUATOR-BASED MOBILE WIRELESS TESTBED

By

Keith Richard Taylor

Thesis submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Master of Science
2012

Advisory Committee:
Professor Joseph Jaja, Chair/Advisor
Professor Richard La
Professor Gang Qu

© Copyright by
Keith Richard Taylor
2012

Acknowledgements

Foremost, I would like to express my sincere gratitude to my advisor Dr. Joseph Jaja for his unwavering support and patience throughout this whole process. Taking on such a large project wouldn't have been possible if it weren't for his faith in my abilities.

I would like to thank Dr. Padma Mundur and Brenton Walker for their motivation and encouragement to pursue my thesis at LTS, working with one of the most intelligent groups of people that I've been around. Without this support, I wouldn't have had the opportunity to work in such a new and exciting research field.

I am extremely grateful for the many hours of assistance from Richard Graham, for without his wealth of knowledge and insightful suggestions this project would not have turned out the way it did.

I dedicate this thesis to my family. I cannot thank them enough for their love and support throughout my life that has enabled me to reach this point. Their encouragement to continue my education has only led to unforgettable experiences that otherwise wouldn't have been possible.

Lastly, I would like to thank my roommates for putting up with the giant mess that has accumulated in the basement while living at my computer desk. Don't worry; it's all cleaned up now.

Table of Contents

Acknowledgements.....	ii
Table of Contents.....	iii
List of Tables.....	v
List of Figures.....	vi
List of Acronyms & Abbreviations.....	viii
Chapter 1: Introduction.....	1
1.1 Motivation.....	1
1.2 Related Work.....	2
1.2.1 Wireless Simulators.....	2
1.2.2 Wireless Channel Emulators.....	4
1.2.3 Wireless Testbeds.....	6
1.3 Previous Work.....	7
1.3.1 MeshTest.....	8
1.3.2 Virtual MeshTest.....	11
1.4 BounceHaus Overview.....	13
1.5 Thesis Outline.....	14
Chapter 2: Testbed Design.....	16
2.1 Hardware Architecture.....	16
2.2 Testbed Implementation.....	20
2.2.1 Precomputation.....	21
2.2.2 Testbed Operation.....	23
Chapter 3: Wireless Channel Modeling.....	24
3.1 Basics of RF Propagation.....	24
3.2 Path Loss Models.....	29
3.2.1 Categorization.....	29
3.2.2 Performance.....	31
3.3 Ray Tracing.....	33
3.4 Wireless InSite.....	36
3.4.1 Performance.....	36
3.4.2 Terrain & Building Models.....	39
3.4.3 Materials.....	40
3.4.4 Transmitter & Receiver Sets.....	41
3.4.5 Antennas.....	42
3.4.6 Waveforms.....	43
3.4.7 Study Areas.....	43
3.4.8 Propagation Models.....	44
3.4.9 Output & Filters.....	46
Chapter 4: Measurement Methodology.....	49
4.1 Hardware Setup.....	50
4.2 Hardware Modifications.....	52
4.2.1 Continuous Beacon Transmission.....	52
4.2.2 Reliable RSS Measurements.....	53
4.3 Data Collection Procedure.....	55

4.4 Calibration & Performance	58
Chapter 5: Integration of Wireless InSite	64
5.1 Terrain Model	64
5.1.1 Terrain Detail	64
5.1.2 Material Selection	66
5.2 Building Models.....	68
5.2.1 Model Creation	70
5.2.2 Building Simplification.....	72
5.2.3 Placement in WI.....	75
5.2.4 Allowed Interactions.....	76
5.3 Transmitters & Receivers	77
5.3.1 Placement.....	77
5.3.2 Equalizer & Correlator Modeling.....	79
5.3.3 Antennas	81
5.3.4 Waveforms.....	84
5.3.5 Collection Radius & Ray Spacing	84
Chapter 6: Field Test Descriptions	86
6.1 AVW	86
6.1.1 Description.....	86
6.1.2 Data Analysis.....	88
6.1.3 Representation in WI	91
6.2 Stadium Drive	92
6.2.1 Description.....	92
6.2.2 Data Analysis.....	94
6.2.3 Representation in WI	96
Chapter 7: Results & Discussion	97
7.1 Results.....	97
7.1.1 AVW.....	99
7.1.2 Stadium Drive	105
7.2 Transmitter & Receiver Modeling.....	107
7.2.1 ETOA & Number of Paths.....	107
7.2.2 Summing Complex Electric Fields	109
7.2.3 Ideal vs. Realistic Antenna Patterns.....	110
7.3 Building Models.....	110
7.3.1 Allowed Interactions.....	110
7.3.2 Building Detail.....	111
7.4 Recommendations.....	112
7.5 Running WI Predictions on BounceHaus	114
Chapter 8: Conclusion & Future Work.....	116
8.1 Conclusion	116
8.2 Future Work	117
Appendix A: Building Models.....	120
Appendix B: List of Contributors	146
Bibliography	147

List of Tables

Table 3.1: Database of Materials Built-in to Wireless InSite	41
Table 5.1: Path Loss (dB) due to a Single Reflection off Different Materials	67
Table 5.2: Modeled Buildings with Identification Numbers	69
Table 5.3: Number of Faces in each Building Model as well as the Full Model.....	75
Table 7.1: RMSE for AVW Line 1	99
Table 7.2: Mean Error for AVW Line 1	100
Table 7.3: RMSE for AVW Line 2	101
Table 7.4: Mean Error for AVW Line 2	102
Table 7.5: RMSE for AVW Line 3	103
Table 7.6: Mean Error for AVW Line 3	104
Table 7.7: RMSE for Stadium Drive	105
Table 7.8: Mean Error for Stadium Drive.....	106

List of Figures

Figure 1.1: MeshTest RF Switch Diagram	10
Figure 1.2: The Premise Behind Virtual MeshTest	12
Figure 2.1: Hardware Architecture	16
Figure 2.2: Node Placed in Shielded Enclosure.....	17
Figure 2.3: BounceHaus RF Switch Diagram.....	18
Figure 2.4: Phases of BounceHaus	21
Figure 2.5: Example of an XML-Encoded Scenario File	22
Figure 2.6: Dynamic Attenuator Settings for the Testbed Controller.....	23
Figure 3.1: Contributions to Noise in the Wireless Channel	24
Figure 3.2: Vertical and Horizontal Antenna Gain Patterns for a 3 dBi Half-Wave Dipole.....	25
Figure 3.3: Mapping of Received Power Produced by Wireless InSite	27
Figure 3.4: Section of Ottawa Showing Transmitter and Receivers.....	37
Figure 3.5: Path Gain of Albert Street	38
Figure 3.6: Path Gain of Queen Street.....	39
Figure 4.1: Dell Latitude Showing Antenna Mount	50
Figure 4.2: RSS Drops due to Wireless Link Adaptation Mechanisms.....	54
Figure 4.3: RSS Correction by Filtering.....	55
Figure 4.4: View of Field Test GUI.....	57
Figure 4.5: Accuracy and Consistency of Nodes over RF Switch.....	61
Figure 4.6: Snapshot of Walking Hallway Test.....	62
Figure 4.7: Results of LOS Field Test	63
Figure 5.1: Color-Coded Topological map of UMD	65
Figure 5.2: WI Simulation Testing Material Reflection Coefficients.....	67
Figure 5.3: UMD Campus map Showing Section Modeled	68
Figure 5.4: Modeled Section of UMD Campus	69
Figure 5.5: View of V1 Buildings Models from the North East.....	72
Figure 5.6: View of V1 Buildings Models from the South East.....	72
Figure 5.7: Example of Window Simplification.....	73
Figure 5.8: Example of External Structure Simplification	73
Figure 5.9: Example of Roof & Recess Simplification	74
Figure 5.10: Vertical and Horizontal Gain Patterns of 3 dBi Half-wave Dipole.....	82
Figure 5.11: Anechoic Chamber used to Measure Antenna Gain Patterns.....	83
Figure 5.12: Vertical and Horizontal Gain Patterns of 3.127 dBi Omnidirectional Antenna.....	83
Figure 6.1: 2D Overview of the AVW Field Test Scenario	88
Figure 6.2: Moving Average on Line1 Mitigating Effects of Fast Fading	89
Figure 6.3: Line1 Motion Points after Moving Average vs. Stationary Points	90
Figure 6.4: Line2 Motion Points after Moving Average vs. Stationary Points	90
Figure 6.5: Line3 Motion Points after Moving Average vs. Stationary Points	91
Figure 6.6: Transmitter and Receiver Placement in WI for AVW Field Test	92
Figure 6.7: 2D Overview of the Stadium Drive Field Test Scenario.....	93
Figure 6.8: Antenna Mount used During Stadium Drive Field Tests.....	94

Figure 6.9: Moving Average on Stadium Drive Mitigating Effects of Fast Fading ...	95
Figure 6.10: Stadium Drive vs. Stadium Drive Reverse at Measured Positions	95
Figure 6.11: Transmitter and Receiver Placement in WI for Stadium Drive Field Test	96
Figure 7.1: Running AVW Line1 Predictions on BounceHaus	115

List of Acronyms & Abbreviations

AWGN	Additive White Gaussian Noise
CIR	Channel Impulse response
DSP	Digital Signal Processing
DTN	Delay-Tolerant Network
EM	Electromagnetic
ETOA	Excess Time of Arrival
FDTD	Finite-Difference Time-Domain
FSPL	Free Space Path Loss
GTD	Geometric Theory of Diffraction
GO	Geometric Optics
GUI	Graphical User Interface
ISM	Industrial, Scientific, and Medical
I/O	Input/Output
IEEE	Institute of Electrical and Electronics Engineers
KFTH	Keith's Field Test Helper
MAC	Media Access Control
MANET	Mobile Ad-Hoc Network
MIMO	Multiple-Input/Multiple-Output
NTP	Network Time Protocol
PAPD	Peak-to-Average Power Detection
RF	Radio Frequency
SBR	Shooting and Bouncing Ray

SNR	Signal to Noise Ratio
UMD	University of Maryland
VMT	Virtual MeshTest Testbed
WI	Wireless InSite

Chapter 1: Introduction

1.1 Motivation

With the growing ubiquity of wireless communication, Mobile Ad-Hoc Networks (MANETs) and Delay-Tolerant Networks (DTNs) have emerged as important and active areas of research and development. These networks feature mobile, multi-hop, and disconnected topologies and provide services in harsh conditions (high latency, opportunistic connections, variable link states, lack of end-to-end connectivity, etc.), alleviating the need for fixed infrastructures. However, evaluating new routing protocols, applications, and hardware implementations requires a testing environment capable of providing repeatable and controlled experiments without sacrificing realism; an environment that is difficult to come by.

Conducting live field tests using actual hardware in the environment of interest is an experimental technique that offers a high level of realism, but has many shortcomings. Representing a topology of mobile nodes requires a large number of subjects and can be difficult to choreograph and control over large a geographical area, only getting worse as the number of nodes scales. Also, measurements are susceptible to drastic changes in the environment such as moving vehicles, people walking by, and other interfering radio devices, limiting experimental control and making field tests only marginally reproducible. Trying to mitigate the undesirable effects of real world environments makes the already time consuming task of conducting field tests a challenging endeavor.

Due to the time, difficulty, and expense of running live field tests, researchers have developed a wide variety of test environments, including wireless network simulators, wireless channel emulators, and recently emerging wireless testbeds. The remainder of this chapter will describe popular implementations of these testing environments, highlighting their strengths and weaknesses, followed by a description of our wireless testbed prototype and an outline of this thesis.

1.2 Related Work

1.2.1 Wireless Simulators

Even with the growing interest in wireless testbeds, wireless network simulators, such as NS-2 [26] and the Opportunistic Network Environment (ONE) [21], have remained the preferred method for evaluating the performance of MANET and DTN protocols and applications. Simulators provide repeatable experiments in a controlled environment, without the high costs of hardware and labor associated with real field deployments and testbeds. Moreover, scalability in simulators is not limited by the number of hardware devices, but the increased runtime of as a result of simulating large network topologies. The popular simulators mentioned above are often criticized for not keeping up with 802.11 standards and for a lack of documentation, stemming from their open source nature. This has led to the development of commercial simulators such as OPNET [27] and QualNet [37]. These simulators claim to have better support for current standards and offer added features such as modeling antenna patterns and weather. While simulators possess many desired qualities for performance evaluation, the difficult task of modeling both the

system at all layers of the network protocol stack as well as the interactions in the physical environment has reduced their experimental realism.

Most simulators contain detailed protocol models of the IEEE 802.11 specifications, but then make vast oversimplifications of electromagnetic wave theory, resulting in a fairly unrealistic model of the physical layer. Kotz et al. [24] conducted a survey of several publications, enumerating six common assumptions in which simulators either explicitly or implicitly relied upon: the world is two-dimensional; a radio's transmission area is circular; all radios have equal range; if I can hear you, you can hear me; if I can hear you at all, I can hear you perfectly; signal strength is a simple function of distance. The authors then conduct a set of field tests, collecting data in a real world outdoor environment. From their measurements, they concluded that the underlying assumptions all contradicted behaviors observed in the field test. Other research [40] indicates that not only are there discrepancies between simulators and the real world at the physical layer, but at the MAC layer as well.

Calvin et al. [6] present a set of measures collected during the simulation of a flooding router, chosen for its simplicity, across three popular simulators (OPNET, NS-2, GloMoSim). The results show that significant divergences exist between the simulators, both numerical and behavioral, in some cases making them barely comparable; mostly resulting from different implementations of the physical layer. When it comes to the accuracy of simulators, these divergences demonstrate how the performance of ad-hoc network protocols deeply depends on the modeling of the lower layers of the protocol stack.

1.2.2 Wireless Channel Emulators

The lack of experimental realism present in today's wireless network simulators has caused some researchers to adopt emulation as a means for evaluation. Emulation embraces the use of real hardware devices, simulating network communication without having to modify their protocol stacks. As the name suggests, wireless channel emulators focus primarily on the physical layer, emulating signal propagation characteristics between each device over the wireless transmission medium. The efforts discussed below both use Digital Signal Processing (DSP) to support dynamic channel conditions such as fading and multipath, allowing for fine grained control of the physical medium.

Azimuth Systems produces a line of Multiple-Input/Multiple-Output (MIMO) wireless channel emulators [2], providing solutions for testing WiFi, WiMAX, and 2G/3G/4G UMTS/LTE products. They use DSP to model RF channels as a Channel Impulse Response (CIR), implementing the response as a tapped delay line. Each tap represents a different path taken by a propagating signal between the transmitter and receiver. The tapped delay line allows users to modify attributes of RF signals such as time delay and power, representing the distance the signal traveled along with the power due to interactions in the physical medium. These devices also support changing signal phase, applying Doppler effects, fast fading models, as well as an Additive White Gaussian Noise (AWGN) generator to simulate real world background noise in the environment.

The wireless channel emulators developed by Azimuth Systems provide fine grained control of the physical medium allowing for the manipulation of several

signal characteristics. However, these devices are severely limited by their internal connections between the I/O ports. The ports are laid out in two sets, bank A and bank B, where no links exist between ports of the same bank. This design was intended for the performance characterization of MIMO devices, not supporting the cross-connections necessary to simulate signal characteristics between arbitrarily arranged nodes. In addition, a method for controlling this device is still required.

The CMU Wireless Emulator uses a Field Programmable Gate Array (FPGA) to implement DSP, digitally emulating wireless signal propagation between nodes. RF nodes are placed in shielded enclosures for isolation, preventing inadvertent cross-talk, and are connected to the emulator through shielded cables. On transmission, the output signal from each node is attenuated, converted to a lower frequency, and then digitized using an A/D converter. The digitized signal is then sent to the DSP Engine residing in the FPGA. The DSP Engine implements a tapped delay line, where the signal is pulled off after a programmed amount of delay. The output signals are then combined with the scaled signals from other RF nodes before being converted back to both an analog signal using a D/A converter and the original frequency.

This design allows for complete control of multipath signal characteristics, simulating real world impairments such as path loss and path delay using multi-tapped delay lines and dynamic path gains. Moreover, these impairments are applied in real time, avoiding long simulation times. The current implementation supports the full 2.4 GHz Industrial, Scientific, and Medical (ISM) band and can handle up to 15 nodes. The next version of this emulator is claimed to support up to 32 nodes. The use of an FPGA for DSP enables all-to-all communication, allowing for representations

of arbitrary network topologies. The main drawback is that allowing for complete control of multipath characteristics requires a method to provide the necessary path loss and path delays. The current platform uses abstract channel models and doesn't take full advantage of its capabilities. This emulator is an attractive option and future implementations of the work in this thesis may focus on integration of this device.

1.2.3 Wireless Testbeds

Over the last decade, a wide variety of wireless testbeds have been developed and it is important understand each classification and their key differences. One class of testbeds involves full-scale deployments in real world environments, allowing for wireless testing in realistic conditions. In MIT's Roofnet [7], wireless nodes are spread across Cambridge, MA on volunteers' rooftops. These nodes form a large-scale static mesh topology and provide broadband internet service to its users. The network is mainly used for studying wireless network characteristics, varying transmit powers in order to form or break links. A similar deployment, the UMass DOME Testbed [4], introduces mobility by placing its wireless nodes on 35 buses driving on the University of Massachusetts campus in Amherst, MA. The added node mobility makes this testbed more flexible, but is still limited to bus routes, not allowing for arbitrary topologies. Unfortunately, these deployments suffer from the drawbacks of real world environments discussed in section 1.1.

Another class of testbeds introduces greater control and repeatability by emulating large-scale wireless networks in a much smaller area, usually indoors. Testbeds of this nature include ORBIT [35] and MiNT-m [10]. The ORBIT testbed is

comprised of 400 custom-designed nodes in an indoor 20m x 20m grid. Since the wireless node grid is static, different network topologies are created by selecting which nodes are to be active, transferring the state of one node to another to emulate mobility. Signal reception between nodes is controlled by injecting noise into the environment, manipulating the noise floor and changing the SNR according to a mapping algorithm. Instead of using a static set of nodes, the MiNT testbed makes use of robots to introduce mobility and topology reconfiguration. To emulate large environments they add an attenuator between a nodes antenna connector and antenna, thereby shrinking the physical space. Both of these testbeds offer an interesting approach, but it is unclear how precisely they can control individual links between node pairs, an essential component in emulating multi-hop large-scale topologies. Moreover, the indoor environment introduces its own multipath effects that are not representative of the multipath that would exist in an outdoor environment, casting doubt on how realistic links may be. In order to scale these networks, another physical node needs to be introduced. This can be bound by the size of the room as well as the cost-per-node, which is more of an issue for ORBIT since it uses custom-designed hardware.

1.3 Previous Work

To avoid the complications of conducting field tests, many researchers have relied on wireless network simulators which allow for a high level of control at almost every layer of the protocol stack. Even though simulators are an attractive option, their highly controlled, repeatable experiments come at the cost of accuracy as

previously discussed in section 1.2.1. Wireless channel emulators such as the CMU Emulator are a viable option, but require complete control of each propagation path between node pairs (attenuation and time delay). This provides an interesting option that may integrate well with the work presented in this thesis and will be discussed in the future work section at the end of this thesis. The development of testbeds like the ones discussed in section 1.2.3 use real RF for their experiments in attempt to control a more realistic environment in a simulator fashion, but they too have their own shortcomings. The testbeds presented in the following sections are a hybrid between using real RF and simulation, providing a more accurate wireless test environment. Both are part of a collaborative effort between the Laboratory for Telecommunication Sciences (LTS) and the University of Maryland. A list of current and past contributors can be found in appendix B.

1.3.1 MeshTest

The MeshTest testbed [9] was designed to provide the qualities of a simulator such as control, manageability, and repeatability, without sacrificing the realism of an RF environment. The basic setup consists of wireless nodes placed in shielded enclosures. The RF energy from each node is wired through shielded cables from its antenna connector to an RF switch of programmable attenuators. Control software is then used to dynamically configure the attenuations between each node pair, allowing us to effectively simulate mobility and arbitrary topologies. This testbed uses real hardware, real operating systems, and real applications, while manipulating a real RF environment.

The original testbed contained 12 nodes obtained from the ORBIT testbed through a partnership with Rutgers. The ORBIT nodes were placed in shielded enclosures offering roughly 80 dB of signal isolation, preventing inadvertent cross-talk. The wireless antennas were removed and the cards wired to SMA bulkhead connectors. A shielded cable then connects the antenna connector to one of 16 input ports of the RF switch. The internal design of the $n \times b$ switch is shown in figure 1.1, n being the number of inputs that connect through nb Ethernet-controlled digital attenuators to b buses. Since there are b buses, there are b paths between any pair of inputs. A signal on any path must pass through two attenuators to simulate loss between node pairs. This architecture makes computing attenuator setting more difficult than one may think; the algorithm is thoroughly detailed in [9]. The RF switch simulates channel loss between node pairs, offering an attenuation range of 0-127dB. However, the RF switch is unable to simulate propagation delays. When a signal passes into the switch, it experiences about 45 dB of insertion loss. This loss is due to passing through the switch's cables, connectors, and mainly the energy division in the splitters. The loss of 45 dB may seem concerning, but in free space it's only equivalent to the first 1.8m of path loss.

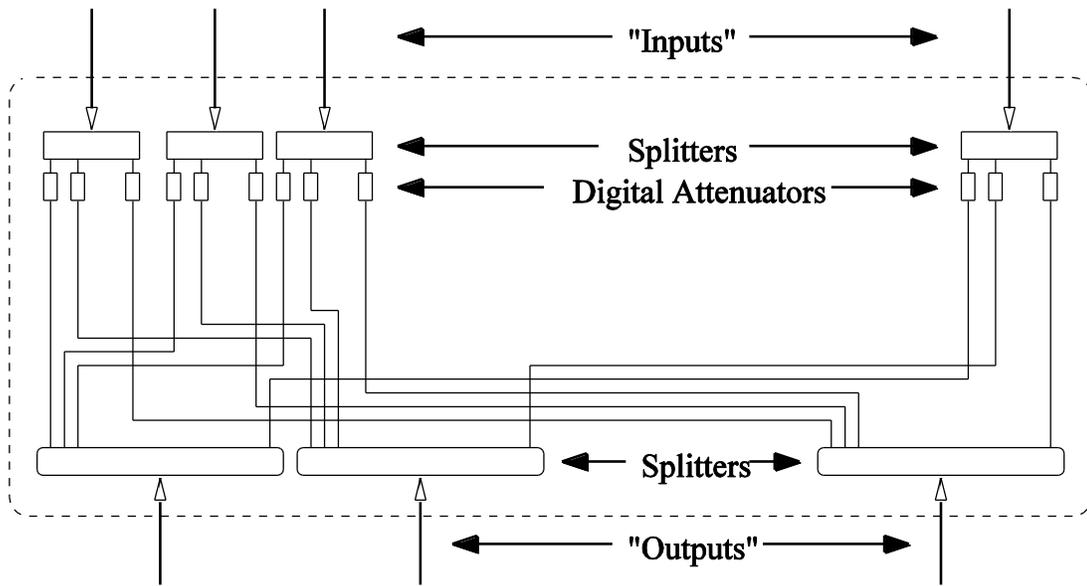


Figure 1.1: MeshTest RF Switch Diagram

For configurability and manageability, MeshTest takes advantage of the ORBIT testbeds management software to control attached nodes. This gives the ability to supply an input script to the control software describing the mobile network at each step in time. Software is added to this framework to first map physical arrangements of nodes to attenuator settings, and then update the RF switch through a TCP socket, simulating experiments in real time.

Unlike simulators, this setup doesn't require any modeling of the protocol stack and subjects real hardware to a real RF environment. Experiments can be conducted with a more diverse set of multi-hop network topologies without being limited by the physical medium. It is also important to note that any device of interest may be placed inside the shielded enclosures as long as they are within the frequency range of the RF switch. Further analysis of this testbed can be found in [9] and [44].

The MeshTest testbed offers an entirely new class of wireless testbeds, but comes with a few limitations. First, the scalability of experiments is limited to the 16 input ports of the RF Switch. This issue is addressed in the next section by the successor of MeshTest. Also, the testbed only controls attenuation between node pairs. An attractive option is to incorporate the CMU Wireless Emulator, allowing for direct control of individual paths between node pairs. Most importantly, MeshTest uses Free Space Path Loss (FSPL) to calculate the attenuation between nodes and the architecture of the switch makes for more difficult attenuation calculations. Both of these issues are resolved by the testbed presented in this thesis.

1.3.2 Virtual MeshTest

To address the issue of scalability present in the MeshTest testbed, our research group designed the Virtual MeshTest testbed (VMT) [14]. The setup is similar to that of MeshTest; wireless nodes are placed in shielded enclosures with their RF energy wired through shielded cables to an input of one of two RF switches containing programmable attenuators. Control software is then used to dynamically configure the attenuations between node pairs on each switch. VMT expands the capacity of the MeshTest by introducing virtualization of nodes and live migrations, activating clusters of nodes across two 8-port RF switches.

VMT consists of 16 physical nodes placed in shielded enclosures, similar to MeshTest, except the nodes are split into two subsets with each subset connected to its own eight port RF switch. The current nodes are running Debian Linux with the 2.6.32-5-xen-amd64 kernel and have Broadcom BCM4321 802.11 a/b/g/n wireless

cards. Each RF switch follows the same design as in MeshTest (figure 1.1) and attenuator settings are calculated using the same technique. Both RF switches apply 45 dB of insertion loss to input signals and each programmable attenuator has a range of 0-127dB.

The Virtualization of nodes and live migration set this testbed apart from its predecessor. The testbed controller takes a scenario file as an input which describes the topologies and mobility of nodes over time. Each node in a scenario is given a virtual machine, a *vnode*, initially running on a virtualization server. When a group of *vnodes* are about to come into communication range, they are migrated onto physical nodes, called *pnodes*. *Vnodes* that are geographically isolated will remain running on the virtualization server until future interactions are feasible. This concept of housing nodes until they become active is depicted in figure 1.2.

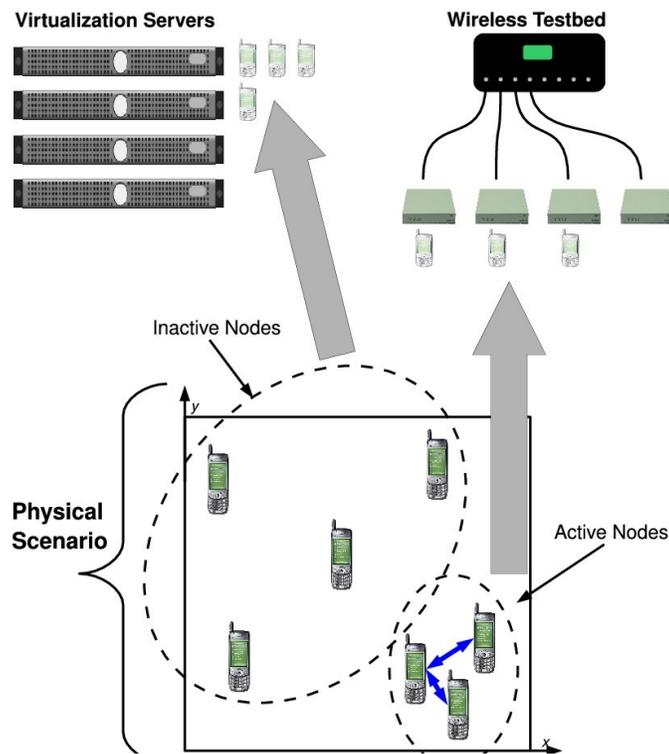


Figure 1.2: The Premise Behind Virtual MeshTest

VMT preserves the wireless test environment presented in MeshTest, but addresses scalability by introducing virtualization and live migration. Experiments have shown that VMT is able to handle scenarios of up to 34 nodes [23], but this requires that nodes exhibit cluster like behavior and that the size and number of communication clusters can be properly allocated to the ports of the RF switch. Another issue lies within the migration of vnodes. Each migration can take three to four seconds, increasing experiment runtimes, and within that migration time communication is suspended across the switch. It is because of these minor issues that the presented testbed will focus on expanding the MeshTest testbed, which will simplify initial validations. VMT provides a viable option to address scalability, but will be further explored after initial validation of the presented testbed has been completed.

1.4 BounceHaus Overview

In this thesis, we present a prototype of a laboratory-based mobile wireless testbed extending the concepts of MeshTest, called *BounceHaus*. BounceHaus is a hybrid between using real RF and simulation, offering researchers the desirable experimental qualities of control, manageability, and repeatability, without sacrificing realism. Nodes are placed in shielded enclosures, having their RF energy wired through shielded cables from their antenna connectors to an RF switch of programmable attenuators. Control software then dynamically configures the attenuators between node pairs in real time. This allows us to effectively simulate mobility and arbitrary topologies while subjecting real hardware running real

operating systems and applications to a real RF environment. Our testbed is freed from the task of modeling the protocol stack of each node, only manipulating the physical medium connecting them. The key feature of BounceHaus setting itself apart from MeshTest is its use of commercial ray tracing software and high-performance hardware to make realistic signal power and characteristics predictions, circumventing the use of abstract channel models.

Unlike abstract channel models, ray tracing software looks at the specific interactions between propagating electromagnetic waves and objects in a highly detailed 3D environment. It calculates the total path loss between a transmitter and receiver by constructively or destructively adding the multiple paths at the receiver that result from these interactions. This physics based approach accurately predicts the effects of large-scale fading, shadowing, and fast fading that occur due to terrain, buildings, foliage, and other objects in any 3D environment, whether it be outdoor, rural, indoor, or mixed path. With the integration of ray tracing software, the testbed is able to control the attenuations between node pairs with a higher level of experimental realism. The testbed design and architecture are described in Chapter 2.

1.5 Thesis Outline

The remainder of this thesis proceeds as follows. Chapter 2 presents testbed design including the hardware architecture and testbed implementation. Chapter 3 provides a discussion of wireless channel modeling starting with the basics and the discussing path loss models, ray tracing, and a description of the commercial ray tracing software used. Chapter 4 presents our measurement methodology describing

the hardware used and general procedures when conducting field tests. Chapter 5 then discusses the integration of Wireless InSite into our testbed prototype as well as the various parameters to be studied. Chapter 6 gives a description of the conducted field tests before Chapter 7 presents a discussion of the results. Chapter 8 concludes this thesis.

Chapter 2: Testbed Design

2.1 Hardware Architecture

The high-level hardware architecture of the BounceHaus testbed is shown in figure 2.1. The *Ray Tracing Server* is a Colfax CX-1250-XF2 and has been equipped with two NVIDIA Tesla M2090 GPUs for reduced calculation times. This server is connected to a gigabit Ethernet switch, allowing it to communicate with the *Switch Control Server*, a Penguin Relion 1600SC containing the testbed control software.

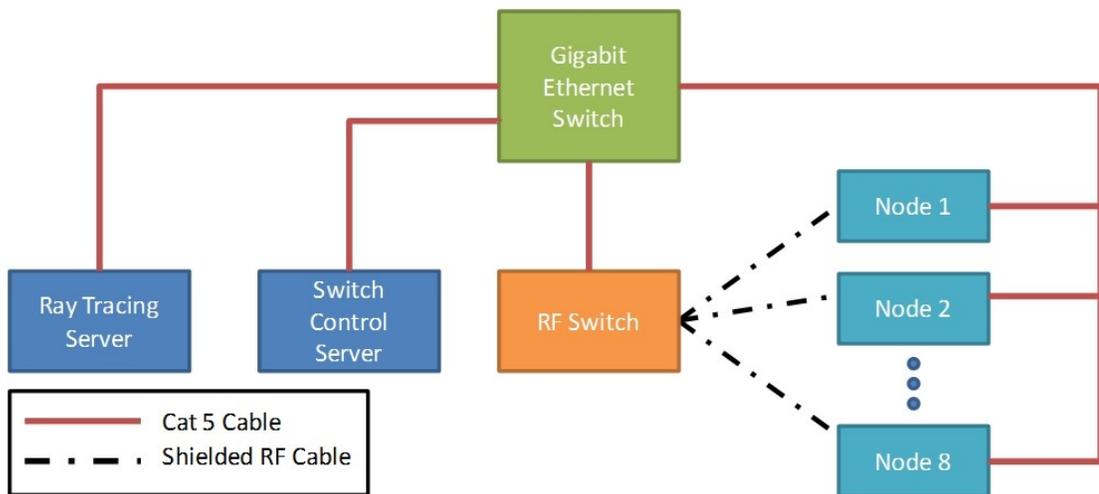


Figure 2.1: Hardware Architecture

The testbed supports up to eight nodes, but in this thesis we're using three Dell Latitude E6500 laptops, each equipped with an 802.11a/b/g Atheros AR5424 wireless chipset. The field tests used for evaluating testbed performance will be conducted using these same laptops. Each node is placed in a shielded enclosure that offers roughly 80 dB of signal isolation, preventing inadvertent cross-talk. The enclosures also offer filtered Ethernet, USB, and DC power for the devices placed inside. The

wireless antenna of the device is removed and the antenna connector is wired to an SMA bulkhead connector. A shielded cable then connects the wireless card to an RF switch of programmable attenuators. Each node is also connected to the gigabit Ethernet switch for remote access. While the current setup contains 802.11 devices, any device, from cellular phones to software-defined radios, may be placed into the shielded enclosures. The BouceHaus testbed does not require the use of custom hardware. Figure 2.2 shows one of our nodes placed in a shielded enclosure.



Figure 2.2: Node Placed in Shielded Enclosure

Figure 2.3 (obtained from [16]) shows the logical construction of the RF switch, a JFW 50PMA-030 8-port RF Transceiver Test System. The design consists of a power divider for each of the eight inputs, as well as 28 Ethernet-controlled digital attenuators, one between each input pair providing 0-127 dB of attenuation with a resolution of 1 dB, also having an accuracy of ± 0.5 dB. Individual attenuators

are controlled via TCP socket and take approximately 3ms to process a request and apply the new attenuation setting. The RF switch has an operational frequency range of 200-2500 MHz, which is one of few restrictions for devices that can be connected to the testbed.

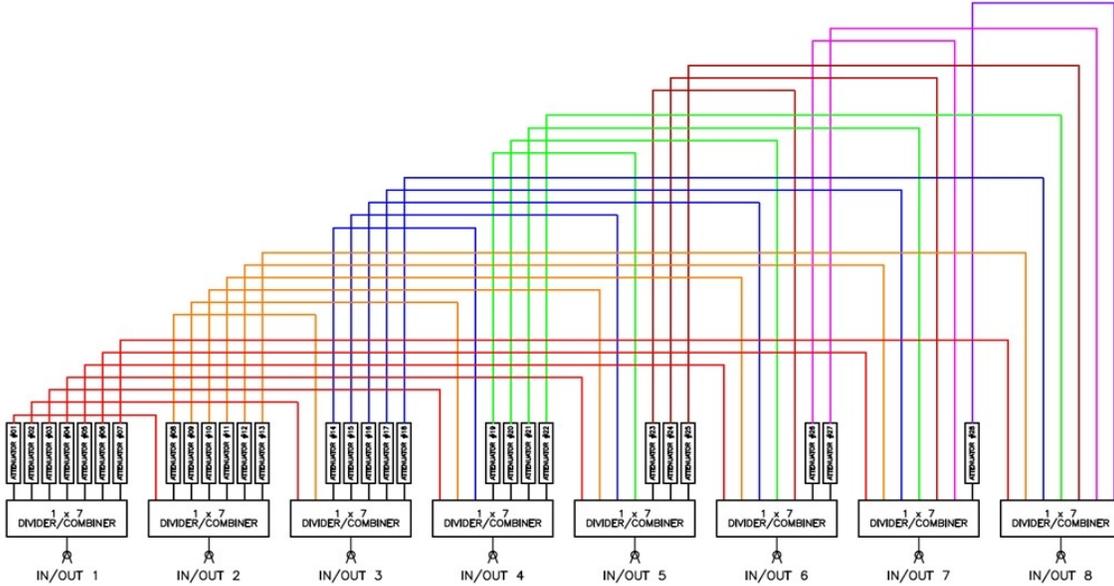


Figure 2.3: BounceHaus RF Switch Diagram

When propagating an RF signal between inputs, the power undergoes an insertion loss. Insertion loss can be defined as the ratio of power supplied on input i to the power propagated out of input j when the attenuator on the path from i to j is set to 0 dB, much like what is described in [9]. This insertion loss is a result of a signal passing through the switch's cables and connectors, mainly from the energy division of the power dividers.

We measured the insertion loss between each port pair on the RF switch, isolating all individual paths. This was done by connecting a signal generator to input i , producing an unmodulated 2.4GHz sine wave and then attaching a spectrum

analyzer to another input, j , and setting the attenuation between those ports to zero. The difference between the transmitted and received power is the insertion loss for that pair of inputs. We found that the insertion loss between all pairs of input ports was approximately 37 dB. Although an insertion loss of 37 dB may seem severe, it is equivalent to the first 0.69m of free space path loss at 2.4 GHz and will not significantly limit the scenarios which can be simulated on the BounceHaus testbed.

Since the testbed is continually changing the attenuations applied by the digital attenuators, it is important to understand the effects of switching attenuation during signal transmission. Even for small changes, changing the attenuation during a transmission will almost always result in a phase discontinuity, potentially disrupting modulation schemes such as OFDM. As done in [44], we used iperf to transmit packets between two nodes connected to the RF switch, we found that 31000 Kbps to be the highest rate at which we could maintain a median packet drop rate of zero with no attenuator switching. The median packet drop rate in this case was taken over a 10s period. To test the effects of switching attenuation, we repeatedly switched the attenuation between 15dB and 16dB at regular intervals. It was found that for switching intervals less than 0.01s, the number packets dropped was related to the switching rate. The only times where the switch attenuations may or could be updated at this frequency is if were replaying a trace captured during a field test or by using a path loss model. The use of ray tracing to produce signal predictions of this frequency carries a computational requirement that will be undesirable to users. But overall, the discontinuities will not interfere with testbed experiments.

The use of this RF switch as a shared physical medium provides many desirable advantages. Due to the simplified internal structure of this RF switch compared to ones used in MeshTest and VMT, algorithms to calculate attenuator settings are simply replaced by a mapping of attenuators to input port pairs, simplifying the control software. One of the main advantages over wireless simulators and other testbeds is that devices experience actual interference from one another. Tests conducted in [44] were able to show that the previously used RF switch acted as a realistic shared medium with several pairs of nodes communicating at once. Their results illustrated that the 802.11 MAC protocol efficiently divided available bandwidth and behaved well even for a group of nodes placed in a small simulated area. Through observations, the same is true of the RF switch in BounceHaus. Because the switch behaves as a realistic shared medium, it is possible to inject noise into the simulated environment with a high level of control.

2.2 Testbed Implementation

The BounceHaus testbed implementation consists of two stages, *Precomputation* and *Testbed Operation*, both of which are shown in figure 2.4 and discussed in the following sections.

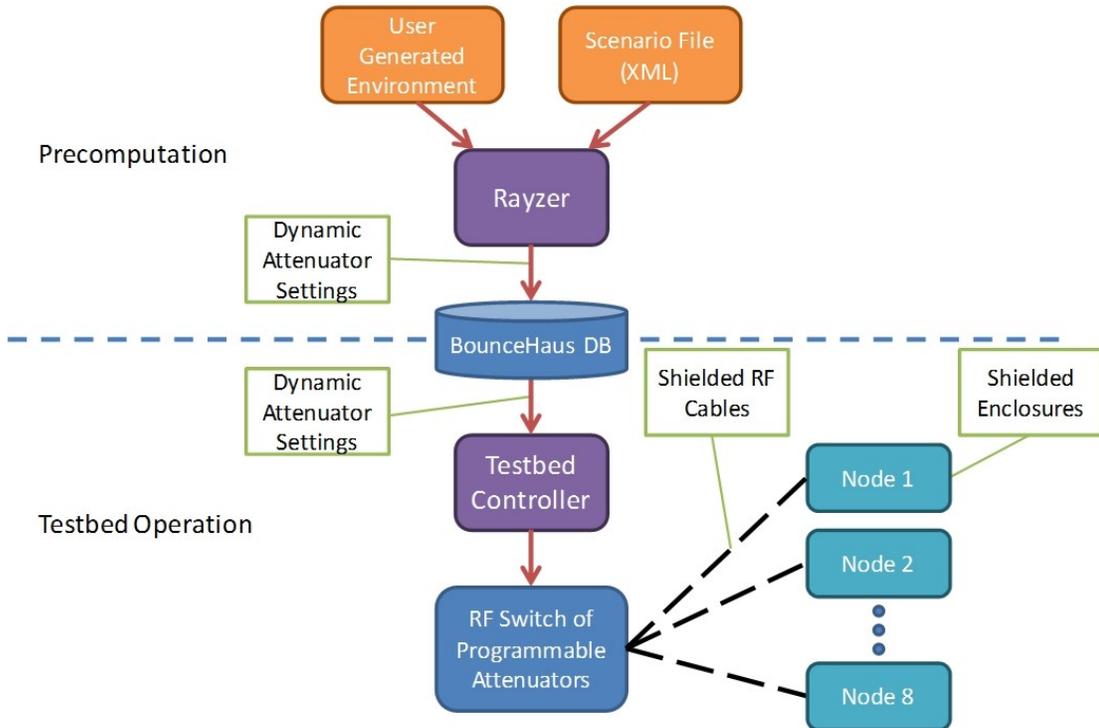


Figure 2.4: Phases of BounceHaus

2.2.1 Precomputation

The primary function of the Precomputation phase is to integrate the ray tracing software, calculating all attenuator settings before a mobility scenario is emulated on the testbed. This is done using the ray tracing control software, called *Rayzer*. *Rayzer* takes two inputs, a scenario file and a user generated environment. The scenario file is an XML-encoded list of node locations for all nodes in the experiment per timestep. An example of a scenario file is shown in figure 2.5. The file gives the GPS origin for the scenario, as well as a location per timestep, given as an offset from the origin in Cartesian coordinates. The second input is the user generated environment and is the primary focus of this thesis. The environment comes to *Rayzer* in the form of an ASCII file which is generated by the ray tracing software. This file contains details of the 3D environment including building models,

terrain, foliage, and material properties. It also contains the properties for the selected waveforms, antenna patterns, transmitters, receivers, and the requested output. All of these properties are introduced in chapter 3, but described in detail in chapter 5 along with how to model a 3D environment.

```
<BOUNCEHAUS>
<!-- Input files:
  1 car_one.dat
  2 car_two.dat
-->
<PHYSICAL TIME="1312218000" UNITS="meters" GPS_ORIGIN="38.8921341770206,-77.0764374672897">
  <NODE INPUT="1" X="374.81921517048" Y="323.92573035602" DESC="node1" />
  <NODE INPUT="2" X="484.04549028418" Y="344.784451749338" DESC="node2" />
</PHYSICAL>
<PHYSICAL TIME="1312218001" UNITS="meters" GPS_ORIGIN="38.8921341770206,-77.0764374672897">
  <NODE INPUT="1" X="373.676606554257" Y="353.384657071641" DESC="node1" />
  <NODE INPUT="2" X="482.957790179183" Y="375.669017618092" DESC="node2" />
</PHYSICAL>
<PHYSICAL TIME="1312218002" UNITS="meters" GPS_ORIGIN="38.8921341770206,-77.0764374672897">
  <NODE INPUT="1" X="370.740472776634" Y="381.12081367136" DESC="node1" />
  <NODE INPUT="2" X="482.95811660471" Y="398.426032522244" DESC="node2" />
</PHYSICAL>
<PHYSICAL TIME="1312218003" UNITS="meters" GPS_ORIGIN="38.8921341770206,-77.0764374672897">
  <NODE INPUT="1" X="368.403205370129" Y="413.683053813811" DESC="node1" />
  <NODE INPUT="2" X="481.326167988173" Y="417.390228879105" DESC="node2" />
</PHYSICAL>
<PHYSICAL TIME="1312218004" UNITS="meters" GPS_ORIGIN="38.8921341770206,-77.0764374672897">
  <NODE INPUT="1" X="366.506545627052" Y="440.150826021722" DESC="node1" />
  <NODE INPUT="2" X="481.326424178987" Y="435.270733059479" DESC="node2" />
</PHYSICAL>
</BOUNCEHAUS>
```

Figure 2.5: Example of an XML-Encoded Scenario File

During operation, Rayzer takes the scenario file and models each individual timestep in the given environment, managing the calculations by sending calculation requests to the ray tracing software on the Colfax server. The output from these calculations is then parsed, sending the attenuator settings for each timestep to the BounceHaus database. This process may be time consuming, depending on the level of detail of the environment and selected parameters, but only needs to be done once for a scenario file. Note that Rayzer is only used for the placement of nodes in the

supplied ray tracing environment and the managing of several calculations. The details of node placement and environment generation are discussed in chapter 5.

2.2.2 Testbed Operation

Unlike MeshTest, BounceHaus performs all signal characteristics calculations in the Precomputation phase, avoiding any calculations during Testbed Operation. Upon user request, the *Testbed Controller* retrieves the dynamic attenuator settings for a scenario from the BounceHaus database. Operating on the *Switch Control Server*, it then initiates the experiment in real time, sending attenuator commands to the RF switch for each timestep in the scenario. The dynamic attenuator settings retrieved from the database are contained in an XML-encoded file. An example is shown in figure 2.6. In this file, attenuator settings are specified for each timestep in the scenario. If there is no attenuation specified for an attenuator, the Testbed Controller sets its attenuation to 127 dB. This applies a total of 164dB of attenuation due to insertion loss, blocking any communication between the intended node pair.

```
<BOUNCEHAUS>
<!-- Dynamic Attenuator Settings -->
<!-- Attenuators Used:
  1
-->
<PATHLOSS TIME="1312218000" UNITS="meters" DIMENSION="3" MODEL="Full13D" >
  <ATTENUATOR INPUT="1" ATT="20" DESC="node1->node2" />
</PATHLOSS>
<PATHLOSS TIME="1312218001" UNITS="meters" DIMENSION="3" MODEL="Full13D" >
  <ATTENUATOR INPUT="1" ATT="27" DESC="node1->node2" />
</PATHLOSS>
<PATHLOSS TIME="1312218002" UNITS="meters" DIMENSION="3" MODEL="Full13D" >
  <ATTENUATOR INPUT="1" ATT="45" DESC="node1->node2" />
</PATHLOSS>
<PATHLOSS TIME="1312218003" UNITS="meters" DIMENSION="3" MODEL="Full13D" >
  <ATTENUATOR INPUT="1" ATT="60" DESC="node1->node2" />
</PATHLOSS>
<PATHLOSS TIME="1312218004" UNITS="meters" DIMENSION="3" MODEL="Full13D" >
  <ATTENUATOR INPUT="1" ATT="55" DESC="node1->node2" />
</PATHLOSS>
</BOUNCEHAUS>
```

Figure 2.6: Dynamic Attenuator Settings for the Testbed Controller

Chapter 3: Wireless Channel Modeling

3.1 Basics of RF Propagation

When transmitting data over a wireless channel, the transmitter first modulates the information onto a carrier frequency in the form of an electromagnetic plane wave. The signal is then radiated into the environment, undergoing attenuation, phase shifts, and frequency shifts due to interactions in the environment before reaching the receiving antenna. The receiver's ability to then correctly demodulate the received signal is dependent on the received signal power in relation to the interference and noise on the channel, known as the Signal to Interference and Noise Ratio (SINR). This process of sending information between a transmitter and receiver in a wireless medium is more thoroughly discussed below and is depicted in figure 3.1 (obtained from [5]). This figure subdivides the noise sources of a wireless channel into multiplicative and additive effects.

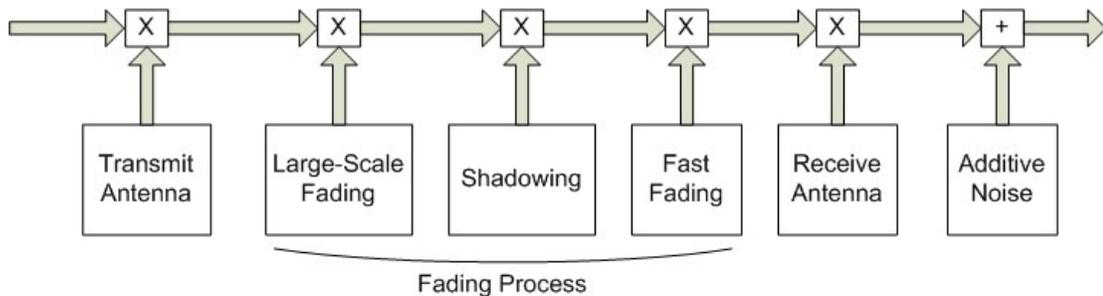


Figure 3.1: Contributions to Noise in the Wireless Channel

The initial radiation and propagation of a wireless signal through an environment is determined by the transmitting antenna's gain pattern. The gain of an

antenna describes how effectively the antenna converts an input power into electromagnetic waves in a specified direction, the opposite being true for the gain at the receiving end. Antenna gain is usually defined as the ratio of the power produced by the antenna to the power produced by an isotropic antenna, expressed in units of dBi. A positive gain corresponds to an increased radiated or received power for the specified direction, while a negative gain corresponds to a decreased radiated or received power for the specified direction. When a direction is not specified, the gain refers to the maximum gain over all directions. Figure 3.2 shows the vertical and horizontal antenna gain patterns of a 3dBi half-wave dipole antenna. The patterns show two main *lobes*, emphasizing signal transmission and reception in those directions.

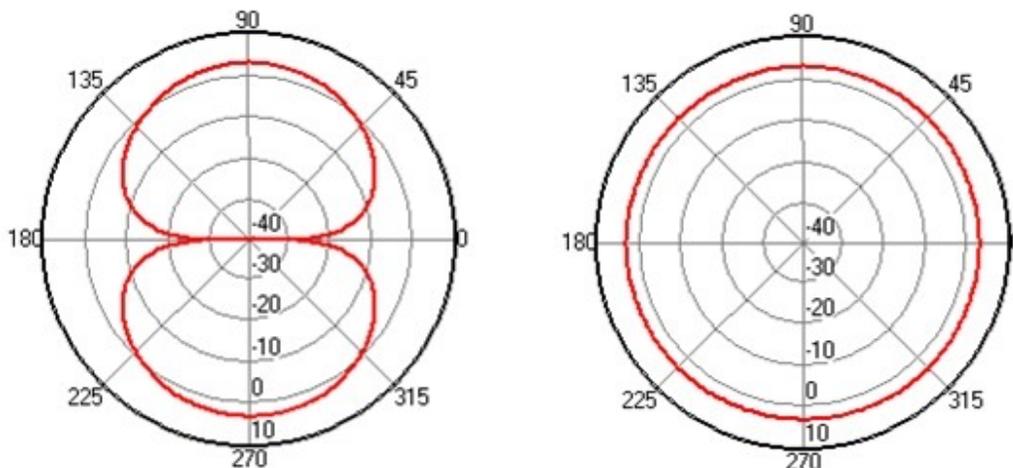


Figure 3.2: Vertical and Horizontal Antenna Gain Patterns for a 3 dBi Half-Wave Dipole

The interactions between propagating electromagnetic waves and objects in an environment play a primary role in the magnitude of the received power. These interactions include ground bounces, reflections, diffractions, transmissions, and absorption. How much an individual interaction with an object affects the received

power depends on the dielectric properties of the material as well as the object size in relation to a wavelength. Mobile transceivers can also introduce frequency shifts due to *doppler effects*. Because an antenna radiates a signal in several directions, there exist several paths from the transmitter to the receiver, known as *multipath*. Each path is attenuated and phase altered by its interactions with the environment, arriving at the receiver at different times. The spread of arrival times is called the *delay spread*. The contribution of each path to the total received power is determined by its phase upon arrival in comparison to the other paths. Signals that are in phase will result in constructive interference, while signals that are out of phase will result in destructive interference. How the receiver decides which paths to combine is further discussed in chapter 5.

The multiplicative environmental effects on signals mentioned above can be split into three types of fading: large-scale fading, medium-scale fading (shadowing), and small-scale fading (fast fading), all of which are time-varying processes between the transmitting and receiving antennas [5]. Large-scale fading is defined as the overall field strength as a function of frequency and distance. This type of fading is due to waves traversing through free space, a deterministic effect. When fading is caused by objects that are static to the environment, it is called shadowing. Shadowing introduces significant variations over large distances based on how the objects are positioned in the environment. Lastly, fast-fading consists of sharp variations over short distances, usually a fraction of a wavelength, caused by small moving objects and the constructive or destructive interference of multiple wave fronts reaching the receiver. Figure 3.3 shows a mapping of received power in an area

between two buildings (shown in black), based on a fixed transmitter location (shown in white and positioned in the middle of the red). The red shows an area of higher received power while purple shows an area of lower received power. The gradual drop in received power radiating from the transmitter shows large-scale fading. Obvious drops in received power occur due to the obstructions of the smaller building, an example of shadowing. Many examples of fast fading can be seen by the sudden drops in received power, some being near the transmitter.

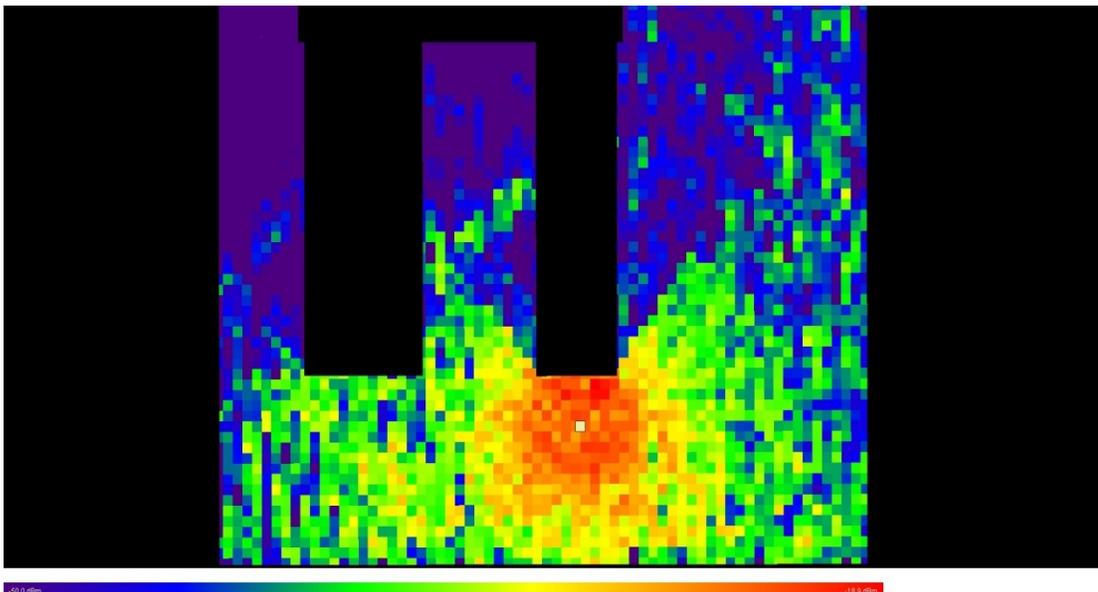


Figure 3.3: Mapping of Received Power Produced by Wireless InSite

Additive noise present in the environment comes from sources both internal and external to the receiver. The receiver itself produces both thermal radiation and noise generated from active and passive devices. This noise is considered to be constant and can be modeled as AWGN. External to the receiver, other similar devices, in this thesis other 802.11b devices, operating on the channel of interest or on other overlapping channels can produce unwanted interference. Since each

802.11b channel spans 22 MHz, only channels 1, 6, and 11 are isolated from one another. Other research [13][19] has shown that non-802.11 devices such as Bluetooth and microwaves can be sources of bursty interference.

After converting the transmit power to the log domain (dBm), the received power of a wireless signal can be expressed as:

$$P_{rx} = P_{tx} + G_{tx} + G_{rx} - PL \quad (3.1)$$

where P_{rx} is the received power, P_{tx} is the transmitted power, G_{tx} and G_{rx} are the antenna gain at the transmitter and receiver, and PL is the overall path loss over the wireless channel [34]. The overall path loss PL can be computed by:

$$PL = L_l + L_s + L_f \quad (3.2)$$

where L_l is the path loss due to large-scale fading, L_s is the path loss due to shadowing, and L_f is the path loss due to fast fading. Fast fading proves to be the hardest to model since it would require complete knowledge of the environment, including details of transient objects and other radio interference. Hence, L_f is modeled separately from L_l and L_s by treating fast fading as a random variable and using statistical models such as Rayleigh or Rician distributions [39][31]. Most path loss models focus on determining L_l and L_s as they capture the overall behavior of the channel.

3.2 Path Loss Models

Decades of development has led to a plethora of path loss models of a wide variety. This is a result of the complexities of different RF environments and the difficulty of accurately predicting field strength within them. In this section, we will first give a brief overview of the categories of path loss models that have been proposed, followed by a discussion on the performance of these models.

3.2.1 Categorization

In [29], Phillips *et al.* conduct an extensive survey of more than 50 proposed path loss models from the last 60 years. They describe 30 of these models in detail, highlighting the key differences between the various types of models. This work provides a rare update to the research community and can be used as a guideline for choosing a path loss model for a preferred environment. The authors also proposed a new taxonomy for path loss models, grouping them into seven major categories and fourteen subcategories. The major categories consist of:

- 1) Theoretical/Foundational Models
- 2) Basic Models
- 3) Terrain Models
- 4) Supplementary Models
- 5) Stochastic Fading Models
- 6) Many-Ray Models
- 7) Active Measurements Models

Theoretical/Foundational models are purely analytical and idealize the theory behind electromagnetic wave propagation. These models are quite simple and are usually a function of distance, frequency, and antenna height of the transmitter and receiver pair. Basic models also use these parameters as inputs, but make corrections

based on measurements taken in one or more environments. These models compute path loss over a single path and are the most numerous. Terrain models behave like Basic models, but try to account for additional losses due to diffractions, such as Fresnel zone blockage. These models require knowledge of the terrain, usually in the form of a Digital Elevation Model (DEM). Supplementary models are not new approaches, but make corrections to existing models, correcting for phenomenon such as frequency coverage, obstructions, and directivity. Stochastic fading models focus on adding the effects of fast fading by adding a random variable to the path loss model. Popular fast fading models include Raleigh or Ricean distributions. Many-Ray models consist of ray tracing methods that look at the multiple paths that exist between a transmitter and receiver in a highly detailed environment, and are able to calculate signal characteristics other than path loss, such as delay spread and phase information. Ray Tracing methods will be further discussed in sections 3.3 and 3.4. The categories previously mentioned are *a priori*, using previously obtained knowledge of an environment, whether that be empirical data or known details of the physical environment. Active measurement models use no prior knowledge of an environment, but provide a method for collecting sample measurements and predicting values for specified locations. These methods usually detail how measurements should be made and how to aggregate data in order to make predictions.

3.2.2 Performance

Even with today's overwhelming pool of path loss models, the differences between them and their performance in a particular environment are difficult to assess. In most published works, authors collect data from their environment of interest to develop and evaluate their proposed model or improvement on an existing model, usually providing a comparison with a select few competing models. The data used for comparison is seldom available to the research community making it difficult to conduct comparative evaluations. Since it is unclear how model performance will translate from one domain to another, it is best to look at comparative studies involving a large number of path loss models over a common realistic dataset from a diverse set of environments (urban, rural, indoor).

In this thesis, we're interested in the path loss prediction in urban environments between devices operating in the 2.4 GHz ISM band. For our performance evaluation in comparison to other path loss models, we'll focus on a study that conducted extensive comparative evaluation between a diverse set of path loss models using realistic data set from a production network, CU WART (University of Colorado Wide Area Radio Testbed) [28]. This is the first study of its kind looking at results in the widely used 2.4 GHz and 5.8 GHz bands, analyzing 30 path loss models spanning 65 years of publications. Not all models analyzed were designed exactly for the sort of environment being studied. The authors state that because many researchers use models well outside their coverage requirements, all models would be given equal chance for making predictions, eliminating any bias. CU WART consists of six 8-element uniform circular phased array antennas mounted

on rooftops around the CU campus and in Boulder, Colorado. The devices can electrically change their antenna pattern, and are changed to represent both directional and omnidirectional antennas for the collected data sets.

The study used four sets of measurements to best capture the behavior of a general urban environment, not just one in particular. The first set was collected from CU WART, looking at RSS measurements in rooftop-to-rooftop scenarios. The second data set came from a municipal wireless mesh network in Portland, where 70 access points were installed on utility poles in a 2km by 2km grid. This data set is representative of ground-to-ground scenarios in urban environments. The third data set represents a wide area infrastructure and was conducted using CU WART. The scenario consisted of one transmitter communicating with several mobile ground-based nodes. Finally, the fourth data set was a reference set used for comparison. The data was collected by the COST-231 group at 900MHz in Munich, Germany.

From the four data sets, the authors concluded that no single model was able to accurately predict path loss consistently, showing that models that performed well in one scenario had poor performance in others. They found that only after tuning the models with measurements from the collected data sets, a best-case performance of 8-9 dB Root Mean Square Error (RMSE) could be obtained. Models that could not be tuned achieved a best-case performance of 12-15 dB RMSE. Both of these RMSE bounds agree with other publications in the field, as pointed out in [28]. We will use these RMSE values as a basis for comparison while exploring ray tracing techniques.

3.3 Ray Tracing

Ray Tracing falls under the Many-Ray category of path loss models and consists of a much different approach. This approach is purely physics-based, treating electromagnetic waves as rays that propagate through a site-specific environment according to the laws of Geometric Optics (GO). Individual rays are cast from a specified transmitter, undergoing reflections, transmissions, and diffractions found using the Geometric Theory of Diffraction (GTD) [20] before reaching a specified receiver. This requires a detailed description of the structures that exist within an environment, including positions, shapes, and material properties. Path loss is calculated by summing the contributions of each individual ray that reaches the receiver, more accurately modeling RF wave propagation and hardware functionality. But this approach goes beyond calculating path loss, providing more detailed information of each path such as phase, direction of arrival, and time of arrival. These methods are used to evaluate the electric field strength and make predictions of signal characteristics in a site-specific environment, allowing for immense customization of environment detail. Signal predictions can be done in both real world representations and hypothetical environments. Moreover, ray tracing methods are not constrained by the limitations of path loss models caused by how they were derived. For example, empirical path loss models contain antenna height and frequency range restrictions as a result of the heuristics based on statistical analysis in the environment the model was initially developed.

While overcoming the restrictions of path loss models, ray tracing techniques have long faced criticism for their accuracy-sensitive coverage mapping due to large

computation and data requirements [29]. The inclusion of reflected, transmitted, and diffracted ray interactions during propagation in realistic 3D environments often results in complex and slow calculation procedures, especially compared to 2D environments. A majority of recent work hasn't been focused on accuracy, but instead optimization and preprocessing while maintaining the accuracy. A survey of various ray tracing methods (optimization) and ray tracing acceleration techniques (preprocessing) can be found in [15]. Along with advances in ray tracing methods, computational times have been greatly reduced by the progression of high performance hardware including today's multi-core processors and GPU capabilities.

Although computation requirements have become less of a factor, the accuracy of ray tracing is tightly coupled to the amount of detail in the modeled environment. Accurate predictions require a precise description of the buildings, terrain, and other objects including their shape, position, elevation, and material properties (permittivity and conductivity). General building footprints, positions, and elevations can be obtained from aerial photography or Light Detection and Ranging (LIDAR) and can be expensive, but many digital databases are available providing this data. However, because it is uncertain how precise this data is and the material properties of each surface may not be known, using this data in a ray tracing environment can only provide approximations for RF signal characteristics. Near perfect accuracy would require complete knowledge of an environment down to the nearest wavelength, detail which is hard to come by and time consuming to create. This would include the impossible task of modeling all transient objects and other sources of interference.

Despite having large data requirements, many researchers have shown that accurate results can be obtained even with generalized environment details. In [36], measurements are taken in two cities in Switzerland at a frequency of 1890 MHz and compared with 2D ray tracing results. In this study, the authors used only a general building layout, estimated electrical characteristics of building walls, base-station location, antenna patterns, and frequency. They found that even with their assumptions and generalizations, they were able to obtain mean errors of 1-8 dB. In [30], measurements were compared in two residential areas in Trenton, NJ. Two base-stations were used to transmit from different heights, at and above rooftop, while a van containing the receiver drove down predefined routes. The average and standard deviation of the error between measured and predicted results were less than 1 and 5 dB, respectively.

In [22], AT&T Bell Laboratories made comparisons between a collected set of measurements from Rosslyn, VA and compared them to predictions from their 3D ray tracing tool. In these measurements, two vans acted as the transmitter and receiver, driving down predefined routes in the streets of Rosslyn collecting path loss measurements at both 908 MHz and 1.9 GHz. The building models were created from extensive aerial photography and complex detail was removed for faster computations. The environment studied consisted of both low residential buildings as well as the high rise structures in the downtown area. It was found that that their tool predicted the propagation loss with a mean error of less than 7dB and a standard deviation of less than 8dB.

3.4 Wireless InSite

In this thesis we have chosen Remcom's EM wave propagation software suite, Wireless InSite® (WI) [32]. Wireless InSite is a suite of ray tracing models and high-fidelity EM solvers for predicting the effects of buildings, terrain, and other objects based on the placement of transmitters and receivers in a 3D environment. Urban, indoor, rural, or mixed path environments can be built within WI using its editing tools, or imported through a number of popular file formats. The calculations are made by shooting rays from the transmitters, propagating them throughout the defined environment and finding interactions with environmental features before reaching the receivers. These interactions include transmissions through features, reflections off feature faces, and diffractions around features. The propagation paths and signal strength coverage can be viewed using the project view, a 3D environment viewer. This GUI allows for easy configuration of a simulation and provides access to other requested outputs. All calculations are done by a separate executable, the calculation engine. The calculation engine writes all output to specific files (ASCII format) from which the GUI reads. The rest of this section is dedicated to describing the main features of WI after first showing the performance of WI in an urban environment. A full description of how WI is integrated into the BounceHaus testbed is given in chapter 5.

3.4.1 Performance

This subsection presents one of many validations made by Remcom for WI. For this validation, Remcom used measurements taken by the authors of [45] in the

city of Ottawa, CA and compared them to predictions made within WI. All measurements were made at a frequency of 910 MHz between a stationary transmitter and a mobile receiver. The transmitter had a height of 8.5m while the receiver had a height of 3.65m. Figure 3.4 shows the 1000m x 1600m area in which the measurements were taken. The bright green dot represents the transmitter location while the red line is the path of the mobile receiver (Albert Street).

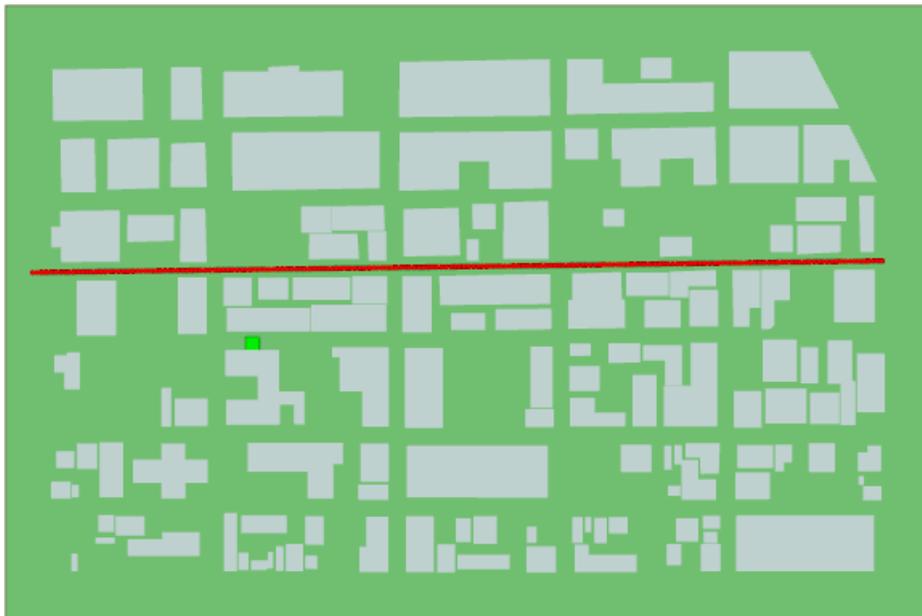


Figure 3.4: Section of Ottawa Showing Transmitter and Receivers

Footprints of the buildings were obtained from maps in [41] since no building data was given. Due to a lack of building height information, the predictions in WI all use the 2D Urban Canyon propagation model which assumes infinitely tall structures. The terrain in WI was modeled as a flat terrain since no terrain was included in the original paper. Also, no information about building materials was reported. Material properties of both the building materials and terrain were set to numbers suggested by the authors of [41].

Figure 3.5 shows the results of the receivers on Albert Street. The figure plots the measurements against both the Urban Canyon predictions and free space path loss. The free space path loss was included to show the loss of attenuation due to buildings. Figure 3.6 shows another set of measurements and predictions from placing the receiver route on the next parallel street to the north (Queen Street). From both of these figures, it can be seen that the predicted path loss was mostly within 10 dB of the measured values. The few regions with a considerable amount of error were attributed to objects in the environment that were not modeled in WI. Also, it is important to note that these results were obtained with minimal effort in environment modeling. Validations using measurements in other cities and environments can be found in [33].

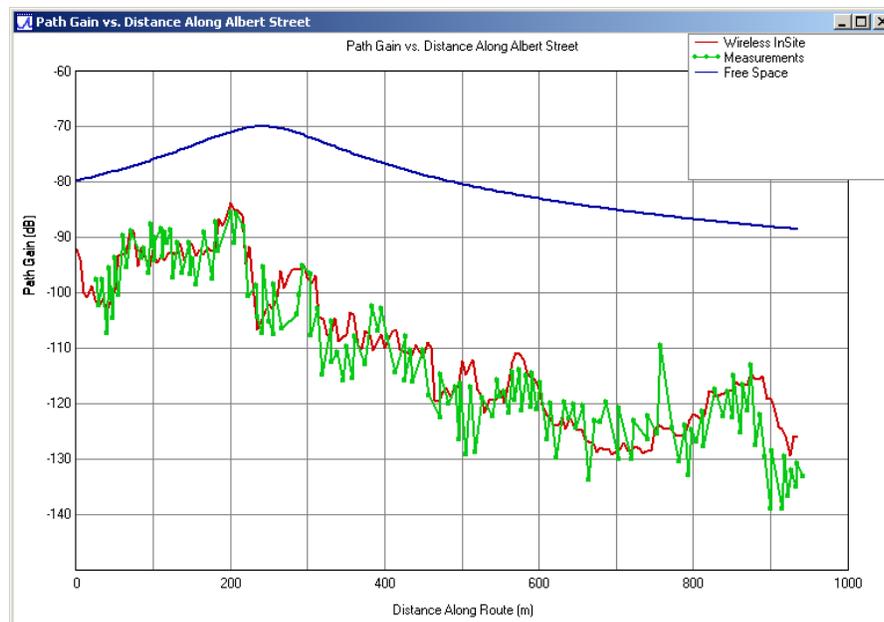


Figure 3.5: Path Gain of Albert Street

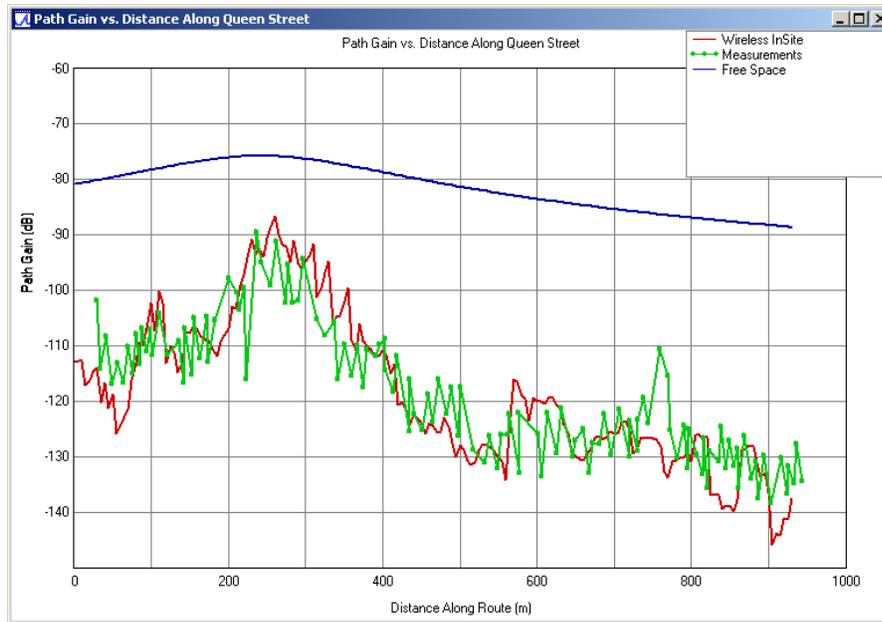


Figure 3.6: Path Gain of Queen Street

3.4.2 Terrain & Building Models

When creating a terrain, users can either use the built-in terrain editor or import a terrain profile from one of the supported file formats, including Digital Terrain Elevation Data (DTED) or Digital Elevation Model (DEM). The built-in editor can be used to define the elevation of individual vertices or materials of individual faces that form the terrain. WI also supports the modeling of foliage, represented by a polygonal shape and a material specifying the type of vegetation.

WI uses three built-in formats to model buildings and objects, these include city, floor plan, and object. The city format is designed to represent the layout of buildings and structures that make up a city and is edited using the city editor, defining the heights and footprints of individual structures. This format can support building models of any shape, but was intended for rectangular vertical walls and a horizontal roof to support running simulations on a large collection of simplified

buildings (an urban environment). The floor plan format allows for the analysis of indoor environments, specifying the location and size of walls, windows, doors, floors, ceilings, and hallways. The floor plan editor is used to place features and apply material properties. The object format is intended for items such as vehicles, posts, tables, and other various items that exist in the environment. An editor is available to create basic polygonal shapes, but it is best to import these objects. All three of these formats (city, floor plan, and object) can be imported using the supported file formats which include Drawing Exchange Format (DXF), ESRI Shapefile, or raster data of building heights commonly obtained by aerial photography or Light Detection and Ranging (LIDAR).

3.4.3 Materials

All terrain and defined structures are represented by a collection of faces, each having its own material properties. Table 3.1 shows the built-in database of materials that was collected from various sources. Each material type has several defined properties such as roughness, thickness, relative permittivity, electrical conductivity, and reflection and transmission coefficients. Users can modify any of these material properties and save the new material type as a new material, or create the new material supplying the material properties listed above. The default properties for the built-in materials can be found in [32].

Building Materials	Terrain Materials	Foliage Materials	Miscellaneous
Concrete Wall Brick Wall Wood Plate Glass Asphalt Layered Drywall	Wet Earth Dry Earth Dry Sand Fresh Water Sea Water	Dense Deciduous Forest In Leaf Sparse Deciduous Forest In Leaf Dense Deciduous Forest Out Leaf Sparse Deciduous Forest Out Leaf Dense Pine Forest Sparse Pine Forest Grass	Free Space Perfect Absorber

Table 3.1: Database of Materials Built-in to Wireless InSite.

3.4.4 Transmitter & Receiver Sets

Wireless InSite provides several types of transmitter/receiver sets that can be used to analyze the characteristics of an environment. Some of these types include points, routes, trajectory, XY grid, vertical surface, arc, and vertical arc. The points set is the most straight-forward method of placement, specifying the location of each transmitter/receiver independently. A route is composed of a series of evenly spaced points connected by line segments. A trajectory can model moving points along a path and is made up of several connected routes. An XY grid allows a large area to be covered with evenly spaced points on the horizontal plane, exactly what the name implies. A vertical surface is the same as an XY grid except in the vertical plane. Horizontal and vertical arc sets can be used, specifying the radius and degree of separation. Other sets include cylinder, sphere, polygon, points on face, and plane waves.

Transmitter and receiver sets have several properties that can be defined. Properties in common with one another include location, height, point spacing (depending on set type), antenna type (see section 3.4.4), antenna orientation, and waveform type (see section 3.4.5). For transmitters, a specific input power may be

given in dBm. Additional properties unique to a receiver include the noise figure and collection radius. The noise figure (dB) is simply the contribution of the receiver itself to thermal noise at its output. The collection radius defines an area around each receiver point where if reached by a propagating path, that path is considered to have reached the receiver point.

3.4.5 Antennas

Each transmitter and receiver point is associated with an antenna as required for propagation calculations. WI offers several built-in antenna models, including short dipole, short monopole, linear dipole, half-wave dipole, linear monopole, quarter-wave monopole, axial mode helix, circular and rectangular loops, circular and rectangular apertures, circular and rectangular patch, horn, and parabolic reflector. WI also offers three general, synthesized patterns, omnidirectional, directional, and isotropic. Most antenna properties available depend on which antenna is selected. Properties common among all patterns include maximum gain, receiver threshold, and noise temperature. The receiver threshold is a specified value (dBm) that determines which ray paths to ignore at the receiver point.

WI also supports user-defined antenna patterns in a format specified in [32]. These patterns can be easily produced and imported by specifying gain and phase at various degrees in a spherical coordinate system. Other formats such as National Spectrum Managers Association (NSMA), Odyssey, and Planet are also supported.

3.4.6 Waveforms

Each transmitter and receiver point is associated with a defined waveform. WI offers a selection of built-in waveforms such as blackman envelope, chirp, Gaussian, Gaussian derivative, hamming envelope, hanning envelope, raised cosine, root raised cosine, sinusoid, and turkey envelope. Each waveform carries its own set of defined properties. Properties common among all waveforms include the carrier frequency, effective bandwidth, and phase. WI can also import a user-defined waveform using a format specified in [32].

3.4.7 Study Areas

A study area defines a region in which to perform a simulation. Only transmitters, receivers, and objects that lie within the study area boundary are considered during calculations. Any ray paths outside of this boundary are simply discarded. Study areas also offer users control over various parameters during a simulation. These parameters include the propagation model, ray spacing (degree at which rays are launched), ray tracing method, allowed interactions, and the number of each type of interaction allowed. The suite of propagation models and ray tracing methods used are discussed in section 3.4.8. Another important parameter is how to sum complex electric fields which has three options. The first option, 'None', sums the powers of each path while ignoring phase. In the second option, 'All', fields are summed with all phase information. Finally, the 'Correlated' option combines the phase of paths that follow nearly the same path in the environment, then summing the powers of the correlated groups. Many of the mentioned parameters will be evaluated in this thesis.

3.4.8 Propagation Models

Originally developed for outdoor urban radio propagation predictions, Wireless InSite has been extended to offer various propagation models for irregular terrain, foliage, indoor, and mixed environments. Available ray-based models include Full 3D, X3D, Urban Canyon, and Vertical Plane. These models combine ray-tracing algorithms with the UTD [3], evaluating the complex electric field of each ray path between a transmitter and receiver. WI also offers additional models such as Free Space, HATA, and COST-HATA, as well as full-wave Urban Canyon FDTD and Moving Window FDTD. These full-wave models simulate radio wave propagation by using the Finite-Difference Time-Domain (FDTD) [42] method, directly solving Maxwell's equations. Although integrating WI allows for the use of any of these models, we'll focus our validations on the most comprehensive model, Full 3D.

The Full 3D model is the most general of the available ray-based models and can be applied in any environment or terrain. It is the only model that places no restriction on object shape and includes transmissions through surfaces. Two available ray tracing methods are available with the Full 3D model: Eigenray and the Shooting and Bouncing Ray (SBR) [38] method. These methods are used to propagate rays through the environment including the effects of reflections, transmissions, and diffractions on the electric field.

The Eigenray method involves the construction of ray paths between a transmitter and receiver that satisfy Fermat's Principle of least time, excluding refractions at transmission points. This method allows a maximum of three reflections and diffractions, and 30 transmissions for each ray path. The computation time is

proportional to $(N_F)^X$ where N_F is the number of faces in the environment and X is the number of reflections and diffractions allowed for any path, the max being three. The limitation of the number of allowed reflections and diffractions as well as the computation time required by using image theory has encouraged us to use the more popular and widely used SBR method.

The SBR method works by first tracing ray paths through the 2D geometry without regard for the location of specific field points. Rays are launched from each transmitter at a set angular spacing, the ray spacing, to find specular reflections off building walls up to the number of allowed reflections defined by the user. It is highly unlikely that rays launched at discrete angles will reach the exact point of the receiver, so a defined area is created to represent the receiver called the collection radius. Ray paths are determined to have reached the receiver if they pass through the receiver's collection radius, while ray paths that reach the study area boundary are simply discarded. After the rays have been shot and bounced from the transmitters, they are then shot from all of the diffracting edges that were found. Diffractions are found by using UTD, identifying discontinuities in the GO fields and determining the diffraction coefficients that will be used to calculate the field strength and phase for each direction away from the diffracting edge.

This procedure is repeated until the allowed number diffractions is reached. From the resulting database of 2D paths, the 3D paths are then constructed and used to evaluate the electric field strength. During this procedure, rays are sorted according to their interactions with feature faces in order to remove rays that follow essentially the same path. These rays represent the same wave front and only one of these rays

will be used to avoid over predicting field strength. The details of electric field evaluation can be found in [32].

When using SBR, the Full 3D propagation model has a maximum of 30 reflections, 4 diffractions, and 30 transmissions. Computation time is proportional to the number of faces in the project when no diffractions are requested. Requesting one diffraction results in a computation time proportional to the number of faces squared. Increasing the number of diffractions above one will further increase computation time, but not exponentially since SBR limits further diffractions to coplanar edges.

3.4.9 Output & Filters

Wireless InSite offers several output types that describe the channel characteristics of a simulated environment. Some of these outputs being received power, path loss, path gain, propagation paths, time of arrival, delay spread, electric field magnitude, electric field phase, complex electric field, direction of arrival, direction of departure, complex impulse response, power delay profile, electric field vs. frequency, electric field vs. time, doppler shift, and diagnostic data. The outputs we're most interested in this thesis are received power and path loss, which will be used for validation against field tests and as attenuator settings for the RF switch.

The way in which received power is calculated is determined by the 'Sum complex electric fields' option discussed in section 3.4.7. A more complete description of the following received power equations can be found in [32]. If no phase information is to be used, the time averaged receiver power (Watts) is given by:

$$P_R = \sum_{i=1}^{N_p} P_i \quad (3.3)$$

where N_p is the number of paths, and P_i is the time averaged power (Watts) of the i^{th} path. P_i is given by:

$$P_i = \frac{\lambda^2 \beta}{8\pi\eta_0} |E_{\theta,i} g_{\theta}(\theta_i, \Phi_i) + E_{\phi,i} g_{\phi}(\theta_i, \Phi_i)|^2 \quad (3.4)$$

where λ is the wavelength, η_0 is the impedance of free space (377Ω), $E_{\theta,i}$ and $E_{\phi,i}$ are the theta and phi components of the electric field of the i^{th} path calculated by the requested propagation model, $g_{\theta}(\theta_i, \Phi_i)$ and $g_{\phi}(\theta_i, \Phi_i)$ are the direction of arrival components, and β is the overlap of the frequency spectrum of the transmitted waveform and the spectrum of the frequency sensitivity of the receiver.

If fields are to be combined with all phase information, the received power is given by:

$$P_R = \frac{\lambda^2 \beta}{8\pi\eta_0} \left| \sum_{i=1}^{N_p} [E_{\theta,i} g_{\theta}(\theta_i, \Phi_i) + E_{\phi,i} g_{\phi}(\theta_i, \Phi_i)] \right|^2 \quad (3.5)$$

If fields are to be combined using the 'Correlated' option, the received power is given by:

$$P_R = \sum_{i=1}^{N_c} Q_i \quad (3.6)$$

where N_c is the number of groups. Q_j is the time average power of each group of correlated paths, given by:

$$Q_j = \frac{\lambda^2 \beta}{8\pi\eta_0} \left| \sum_{i=1}^{M_j} [E_{\theta,i} g_{\theta}(\theta_i, \Phi_i) + E_{\phi,i} g_{\phi}(\theta_i, \Phi_i)] \right|^2 \quad (3.7)$$

where M_j is the number of paths in the j^{th} group.

The path loss output calculation is similar to equation 3.1 and is given by:

$$L_{Path} = P_T(dBm) - P_R(dBm) + G_T(dBi) + G_R(dBi) - L_S(dB) \quad (3.8)$$

where $P_T(dBm)$ and $P_R(dBm)$ are transmitted and received power, $G_T(dBi)$ and $G_R(dBi)$ are the gains of the transmitter and receiver with respect to antenna orientation, and $L_S(dB)$ is the sum of other losses in the system. A description of the other output types can be found in [32].

Output filters provide a convenient way to isolate ray paths with specific interactions, defining a subset of the paths found during simulation. Filters can be applied defining the min and max values for properties for each ray path. These properties include power, time of arrival, and excess time of arrival. Excess time of arrival (ETOA) is the delay of each path after the first arriving path at the receiver. The use of excess time of arrival will be further discussed in chapter 5.

Chapter 4: Measurement Methodology

In the following chapters, field test measurements will be used for guidance on environment and receiver modeling, as well as validation of modeling techniques. Field test measurements need to be collected in a manner that best represents a MANET or DTN. This type of network requires measurements to be collected between each node pair, requiring every node be independent from others and capable of both transmitting and receiving simultaneously throughout the duration of a field test. Equipment such as channel sounders could be used to capture the impulse response of the channel providing detailed multipath information. However, their cost can be quite prohibitive, not to mention the trouble of managing multiple channel sounders in our topology of interest.

In this thesis, we take advantage of inexpensive commodity hardware and use 802.11b packet-based RSS measurements to calculate the path loss between node pairs. Unlike other research using a similar approach [17][11], we apply modifications that alter 802.11b behaviors, enabling our hardware to reliably collect accurate RSS measurements without interruption. Using commodity hardware comes with a few disadvantages including the lack of RF interference characterization, the resolution of measurements which is limited by packet transmit rates, and having to cope with channel congestion. Measurements cannot be collected RSS values for packets that fail to demodulate. Also, packet-based methods are not able to obtain specific multipath information and are limited to RSS. The current testbed implementation uses an RF switch that only controls attenuation between nodes, so

RSS will suffice. In the following sections we describe the basic hardware setup, hardware modifications, data collection procedure, and then address any reservations of using commodity hardware by conducting calibration and performance analysis.

4.1 Hardware Setup

RSS measurements are captured using three Dell Latitude R6500 laptops, each equipped with the 802.11a/b/g Atheros AR5424 wireless chipset. The wireless card is connected to an external omnidirectional antenna mounted on the side of the laptop using a U.FI to RP-SMA pigtail adapter, shown in figure 4.1. Each antenna was tested in an anechoic chamber and was found to have approximately 3 dBi of gain, as opposed to the 5 dBi gain claimed by the manufacturer (details in section 5.3.3). With our hardware setup, nodes can be configured as a transmitter or a receiver, or both. As previously mentioned, all field tests presented in this thesis have nodes configured as both in order to represent a MANET or DTN topology.



Figure 4.1: Dell Latitude Showing Antenna Mount

For transmitting, a node is configured as an Access Point (AP) by running *hostapd*. This allows for a continuous stream of 802.11b beacons to be broadcast at a low modulation rate of 1 Mbps from which receivers will take RSS measurements. The beacon interval can be set within the *hostapd* configuration file and will be changed depending on the field test conducted. The ability to associate with this AP has been disabled by modifying the *hostapd* acceptance list, assuring that 802.11 devices external to the experiment are not connecting to our nodes. The transmit power of every node is set at 20 dBm and is kept constant during every field test, transmitting the 802.11b beacons on the least congested channel. Beacons are tagged with sequence numbers, enabling us to detect when packet reception fails due to interference, poor RSS, or deep fades.

For receiving, we create a virtual interface and put it into monitor mode on the same channel as the transmitter. Received packets can then be logged by running *tcpdump* on that interface. From this log, we're able to retrieve the RSS and MAC timestamps for individual packets, as well as other information about transmissions on the channel. The Atheros AR5424 based cards measure RSSI at the beginning of packet acquisition, so RSSI readings are quick samples and not an average of the whole packet [17]. The RSS measurements resulting from RSSI to RSS vendor mappings are given with 1 dBm granularity. In addition, the receiver has a sensitivity of -95 dBm which will be taken into account when receiver modeling within Wireless InSite.

4.2 Hardware Modifications

The hardware setup described in the previous section enables a node to act as both a transmitter and receiver. But in order to enable commodity hardware as a reliable measurement tool, certain hardware modifications are required to alter the intended behavior of our wireless cards. The following modifications are aimed at achieving two goals: broadcasting a continuous stream of beacons without any interruptions, and attain accurate and reliable RSS measurements.

4.2.1 Continuous Beacon Transmission

The first modification ensures continuous beacon transmissions by disabling network discovery. Virtually all Linux distributions use Network Manager to manage all available nearby networks. Users can either select their network or Network Manager can automatically associate with an access point of the strongest signal. Keeping an updated list of available networks involves periodic scans of all available channels, each scan lasting up to a few seconds. From our observations, any beacon transmitted within the time frame of the frequency scan would in turn be transmitted at the frequency of the current channel being scanned. Since we desire a continuous stream of beacon transmissions on the same channel as the receiver, we simply disable the Network Manager before any field test.

While the first modification prevents a node from disrupting network traffic for network discovery, the second modification prevents disruptions in beacon transmissions by responding to requests from other devices (e.g. probe and association requests). Any packet in the transmit queue other than a beacon could

potentially delay the transmission of the next beacon. This issue becomes more of a concern when using short beacon intervals in an environment comprised of numerous 802.11 devices. To allow only for the transmission of beacons, we modified the *mac80211* kernel module to drop unwanted packets from the transmit queue before transmission.

4.2.2 Reliable RSS Measurements

Figure 4.2 shows two raw signal traces collected after applying the modifications described in the previous section. The traces were both collected by standing in a parking with the distance between the transmitter and receiver being 5m and 60m. While the traces show the continuous beacon stream resulting from the application of the proposed modifications, they also demonstrate that periodically, and only for a single packet, the RSS drops by a significant amount compared to neighboring packets. These sudden drops in RSS are a result of two wireless link adaptation mechanisms, the first being antenna diversity. The antenna diversity mechanism is found in most commodity hardware devices and allows for the switching of antennas to improve the quality and reliability of a wireless link. Giustiniano et al. [11] discovered that Atheros based cards employ the following diversity control mechanisms:

- 1) antenna switching performed for data retransmission

- 2) antenna switching performed on the loss of two consecutive packets

Since we are only using one antenna in our field tests, any RSS reading from or transmission on the second antenna interface will result in lower measurements, provided the packet is not lost. To prevent us from having to disable antenna diversity

before every experiment, we modified the *ath5k* kernel module to only use the main antenna interface.

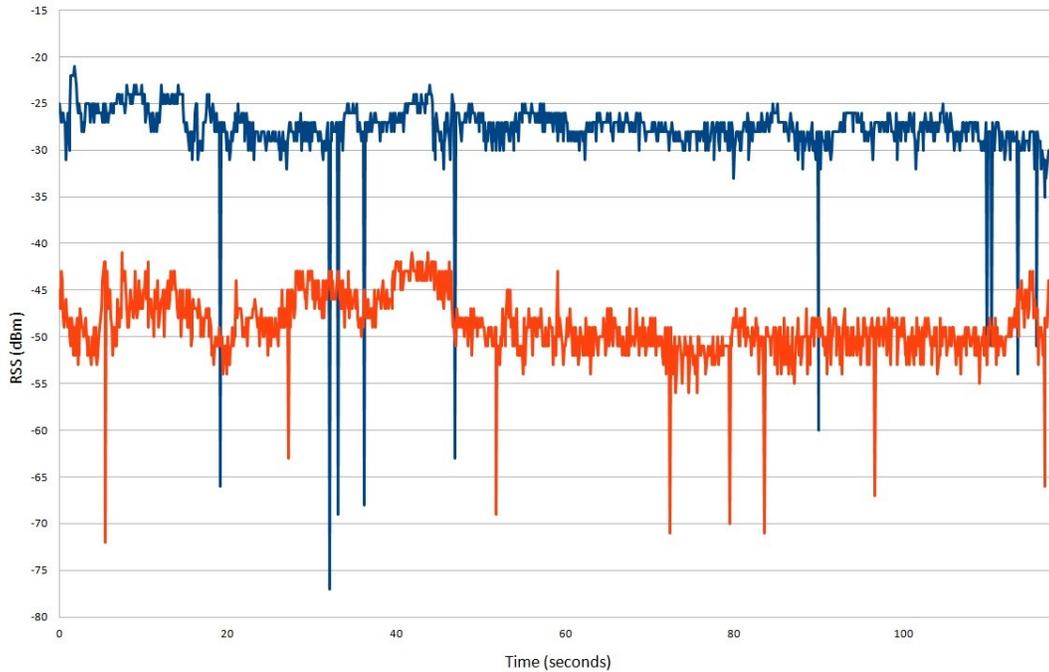


Figure 4.2: RSS Drops due to Wireless Link Adaptation Mechanisms

Even when disabling antenna diversity, drops in RSS could still be observed similar to those shown in figure 4.2. Discovered in [25], these power drops were caused by a calibration procedure which was invoking a peak-to-average power detection mechanism (PAPD), causing the wireless card to transmit the next packet with the lowest possible power while monitoring transmit power to compute gain. These packets are easily isolated and filtered out of the trace data for all of our field tests. The resulting traces are shown in figure 4.3. Because of their low frequency, filtering out these packets will not affect our analysis. Further calibration for RSS accuracy is conducted in section 4.4.

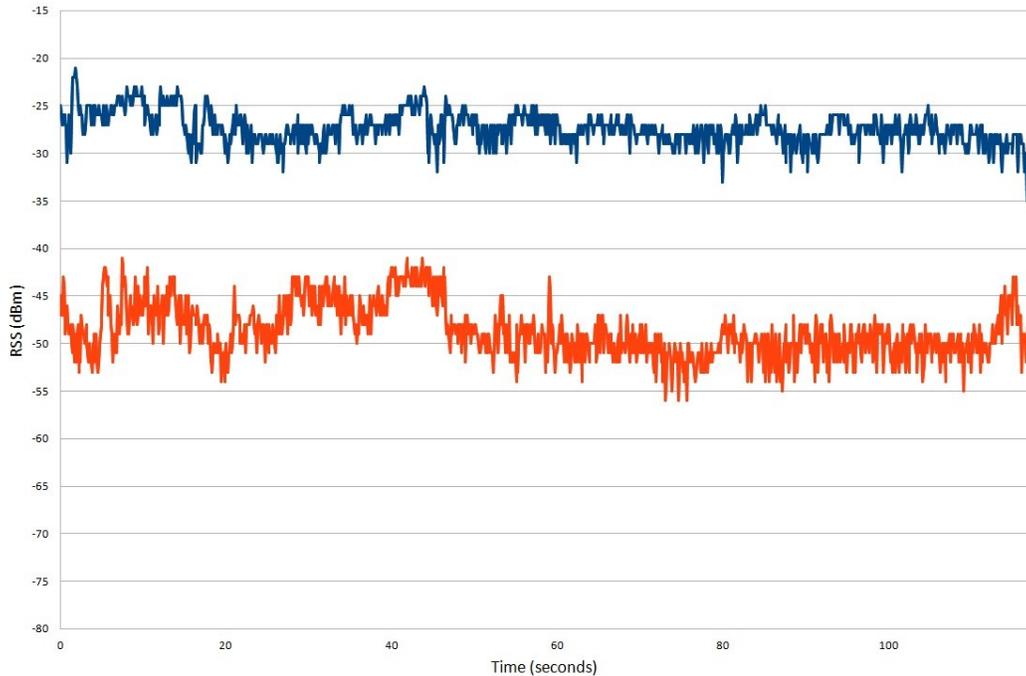


Figure 4.3: RSS Correction by Filtering

4.3 Data Collection Procedure

This section describes the general data collection procedure common among all field tests in this thesis. Details unique to each field test will be presented in chapter 6.

Before conducting a field test, all nodes used are plugged into a power source to charge their batteries. All of our mobility is generated by walking over pre-defined paths, but if cars were used one could use a power adapter to keep a node charged during the field test. Also, all nodes are connected to the internet and use the Network Time Protocol (NTP) for clock synchronization. When powered off or disconnected for a long period of time, the nodes rely on their hardware clocks and experience a considerable amount of clock drift.

To combat the difficulties of conducting outdoor field tests, a program with an interactive Graphical User Interface (GUI) called "Keith's Field Test Helper" (KFTH) was developed to simplify the data collection process. KFTH can be configured through the GUI or through an input configuration file. Settings in the configuration file include project name, node number, wireless interface, channel, enabling as a transmitter or receiver, and MACs of interest. The project name and node number are used in the program's file system to better organize field test data. The wireless interface and channel are used to configure `hostapd` and `tcpdump`, only being initiated based on the functionality of that node specified during configuration. The MACs of interest are a list of MAC addresses of the other nodes in the field test.

When all of the configuration settings are specified, KFTH first checks that the correct kernel modules are running and turns off the network manager. It then creates a virtual interface running in monitor mode on the specified channel. This interface is used by `tcpdump` to log received packets in a pcap file and also by *tshark*, running a live capture to update the display with the devices in communication range along with their RSSI, shown in figure 4.4. The user can either view all devices in range or just the devices with their MAC addresses listed in the configuration. The display also shows a log of events, including the creation and manipulation of interfaces, initialization of programs, and any errors that occur during operation. Field tests are conducted by using the start and stop buttons controlling `hostapd` and `tcpdump`. The GUI also shows an activity bar making it easier for the user to tell when the program is transmitting or collecting data. KFTH greatly simplifies the data

collecting process, taking much of the difficulty away from the user and providing a debugging tool allowing for quick adjustments.

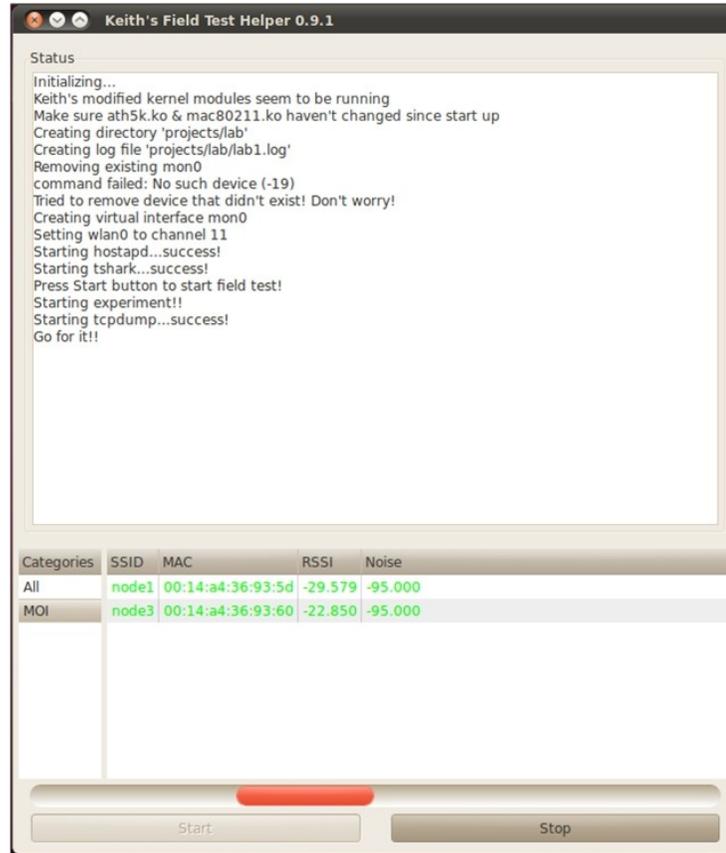


Figure 4.4: View of Field Test GUI

While running KFTH during an experiment, nodes are held so that the omnidirectional antennas are pointing toward the zenith sky, maintaining their vertical polarization. If the position of the nodes results in obvious RF propagation through the body of the conductor, a mount will be used to elevate the antenna above the conductor, limiting undesirable attenuation in the near field. Each conductor uses a clock script to record the UNIX timestamp of when the node reaches each point on a pre-defined path.

The original plans for the field tests conducted in this thesis involved using GPS devices to keep track of node locations. Unfortunately, the devices available to us were only accurate to approximately 10 feet at any given time. After conducting several traces on the UMD campus, we observed large discrepancies between repeated traces. For our desired level of accuracy, it was decided to contend with human error of marking times to predefined locations and limit our field tests to smaller areas. Future work will involve more accurate GPS devices and field test scenarios taking advantage of the entire modeled area.

Although an impossible feat, our best effort was given to try and control environmental conditions. All tests were conducted at times when few people and cars were around, trying to mitigate fast fading due to transient objects in the environment. Field test data was discarded if a car drove into the field test area.

4.4 Calibration & Performance

This section is focused on the performance of our modified commodity hardware as a reliable measurement tool. Even with the previous modifications, there are issues common amongst commodity devices that can lead to variations in RSS measurements. The researchers in [25] found that even though wireless cards allow for per-packet power control, they also exhibit anomalous fluctuations in transmitted power most likely caused by adaptation/diversity/calibration mechanisms beyond what we have already corrected for using the previously mentioned modifications. Judd et al. [18] investigated the causes behind asymmetric link behavior observed by several research groups. They found that the dominating factor in asymmetric links

was the transmit power variation present across devices of the same model. Eleven seemingly identical wireless cards were tested and it was found that the cards fell into two groups. While one group showed similar average transmit powers, the other group exhibited significant variations. While wireless cards may be marketed and labeled as identical, it is possible they may come from different vendors. Judd et al. also concluded that link asymmetry is also a cause of additive effects of several causes usually local to a device in the environment.

Even when using homogeneous devices, RSS characterization should be performed for each wireless card to determine whether or not each device provides the desired level of accuracy and if any calibration is needed. Ideally, we would use an external source with a known transmit power like an Agilent Vector Signal Generator to calibrate our RSS measurements, removing transmit power uncertainties. However, the costs of such devices are quite prohibitive and while we did have access to one, we did not have the correct modules to transmit Wi-Fi packets.

For the RSS characterization of our devices, we first directly connect two of our nodes through a shielded RF cable while attaching a 10 dB attenuator. One node is configured to transmit beacons at 20 dBm while the other node is set to take RSS measurements. This is done with each node configured as a receiver and between all node pairs of our three nodes. It was found that in this simple case, all nodes had a mean error of less than 0.65 dB from the expected value. Note that all tests in this section use a frequency of 2462 MHz (channel 11).

Even though our wireless devices show great measurement accuracy for high RSS, it is imperative to look at both consistency and accuracy of RSS measurements at varying applied attenuation. To apply varying levels of attenuation, we place two of the nodes in shielded enclosures and connect them to our RF switch. We then varied the attenuation from 0 dB to 75 dB in 5 dB increments. This should bring the RSS measurements close to the receiver sensitivity of -95 dBm taking into consideration the insertion loss of the switch and added cable loss. Figure 4.5 shows the RSS measurements for the applied attenuations with the error bars representing one standard deviation. Looking at the standard deviation, we see two areas with increased variation. The attenuation range of 25 dB to 40 dB exhibits a standard deviation of approximately 1.1 dB, increasing from the 0.6 dB standard deviation at lower attenuations. We also see the region of 60 dB to 75 dB increase the standard deviation to approximately 1.5dB, although it was expected to see larger variations at lower RSS readings. In terms of accuracy, the mean error of these readings is less than 0.5 dBm except for the applied attenuations of 60dB to 75dB, which experience a mean error between 1.48dB and 1.95dB. Considering the typical fade margin is 10dB, a mean error less than 2dB for lower RSS measurements is very reasonable. These tests were done for each node pair and all nodes showed similar behavior.

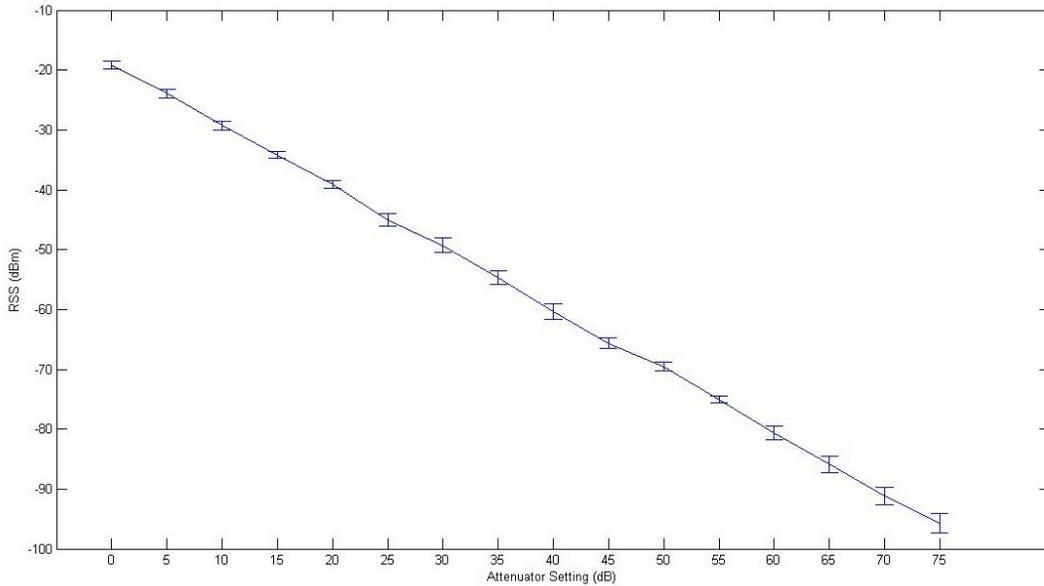


Figure 4.5: Accuracy and Consistency of Nodes over RF Switch

To better display the agreement between node measurements in a more realistic environment, we have conducted two field tests. The first field test was conducted in the hallways of our building. Two nodes were placed at the end of two different hallways, each 17m from the perpendicular intersection of two hallways resulting in a 34m route between them. During this test, nodes were set to both transmit and receive with one node stationary and the other walking towards the stationary node. A 30 second snapshot of approximately 300 RSS measurements within the 3 minute trace is shown in figure 4.6. Although both traces display many similar characteristics, a few regions of variation do exist. Even with these variations the two traces have a RMSE of only 2.8 dB.

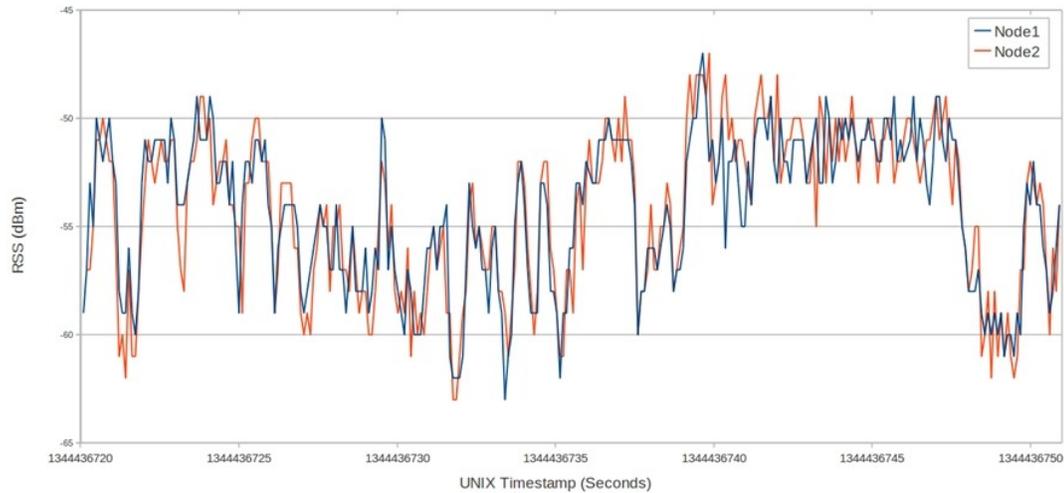


Figure 4.6: Snapshot of Walking Hallway Test

The second field test was conducted between two nodes in an open field, varying the distance by first 2.5m and 5m, then in 10m intervals from 10m to 120m. Each node was set to both transmit and receive with a beacon interval of 100ms, taking measurements at each position for approximately two minutes. Figure 4.7 shows the average RSS measurements taken by each node at each distance, along with the free space path loss and predictions from Wireless InSite. Readings from both nodes were in good agreement, having a RMSE of only 0.75dB. When comparing measurements taken by both nodes to the predictions from WI, we get a worst case RMSE of only 1.76dB. Most notably, the measurements from both the field test and WI show a reduction of RSS at a distance of 30m. Given that the height of the nodes was approximately 1.3m, this distance represents the boundary of the first Fresnel zone. In this case, the radio waves reflecting off of the ground are arriving at the receiver out of phase, effectively reducing the RSS. By using WI, we can take advantage of phase information and hopefully increase the accuracy of predictions for real world scenarios.

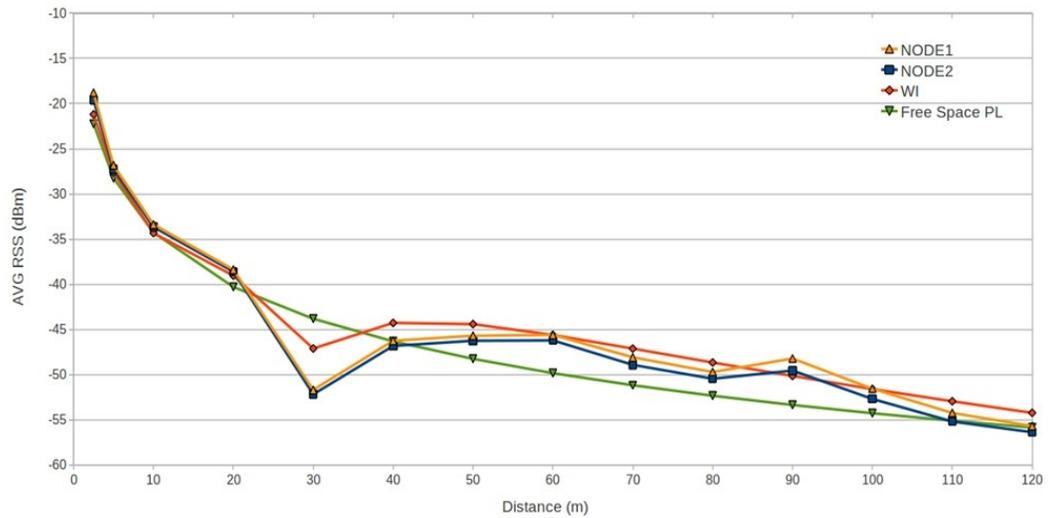


Figure 4.7: Results of LOS Field Test

Overall, we've been able to obtain consistent and accurate RSS measurements in both ideal and realistic environments with each of our wireless cards. Moreover, the readings from each card has shown good agreement with the others; the worst case RMSE from our initial tests being 2.8 dB in a noisy office building. In later field tests, the effects of these variations will be mitigated using aggregation, limiting the effects of fast fading during our comparisons. Based on these findings, the inaccuracies produced by each device are not considered a significant source of error and thus no calibration will be performed.

Chapter 5: Integration of Wireless InSite

In this chapter, we discuss the integration of ray tracing software into the BounceHaus testbed. We explore the various parameters within WI, tuning them in attempt to properly model both the physical 3D environment and the hardware communicating over the physical medium. To do this, we have created 3D models of a selection of buildings from the University of Maryland (UMD) campus. Using field test measurements conducted at UMD (chapter 6), we will evaluate the performance of the BounceHaus testbed based on how we set parameters within WI (chapter 7). We will first focus on terrain and building modeling, discussing the levels of detail that will be analyzed, and then the parameter settings to properly model transmitter and receiver hardware.

5.1 Terrain Model

5.1.1 Terrain Detail

Figure 5.1 shows a color-coded topographical map of the section of UMD that we have modeled and will use for evaluation. This data was obtained from contacts at UMD who had the elevation of campus professionally mapped. From north to south, the elevation increases by 1.31m over approximately 200m. From east to west, the elevation increases by 4.67m over approximately 400m, mostly within the last 100m. All of the changes in elevation are gradual and there exist few drastic changes that would have a significant impact on our collected measurements. Because our terrain is relatively flat, we have decided to use a flat terrain model in WI for initial

evaluation. The field tests used in this thesis were conducted in the center and eastern end of the modeled area, removing most uncertainty due to changes in elevation.

Using a flat terrain in our WI environment will cause an elevation bias in the UMD model based on how the building heights are set and will be discussed in section 5.2.

Future work will involve realistic terrains using either the built-in terrain editor to create the sloped terrain, or by using the terrain data obtained from UMD. Applying terrain elevations will prove more valuable when conducting field tests that take full advantage of the area in which we have modeled. But, given the location and scale of our field tests, the impact of using a flat terrain will be less significant than the details of our building models.

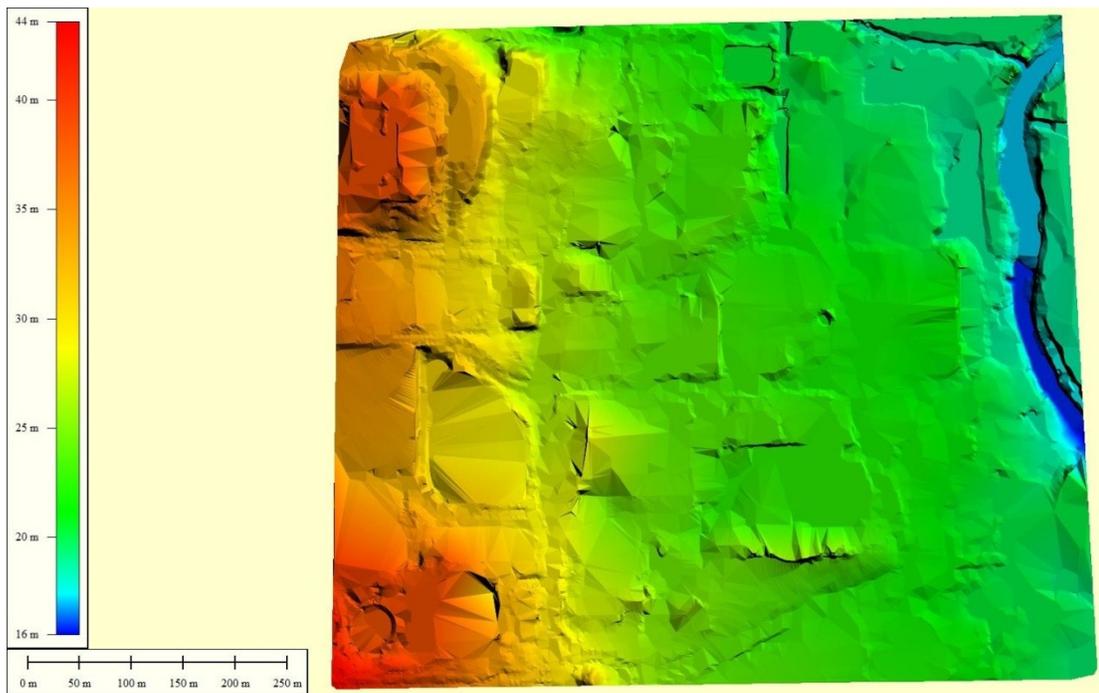


Figure 5.1: Color-Coded Topological map of UMD

5.1.2 Material Selection

The characteristics of the reflections and transmissions depend heavily on the permittivity and conductivity of the materials that make up these surfaces [5]. As pointed out in [36], one reflection could have a variation of loss up to 6 dB, depending on the electrical parameters applied to the surface. Due to these variations, accurate predictions in site-specific propagation modeling require not only the details of structures in the environment, but a reasonable selection of permittivity and conductivity for building materials.

To further demonstrate the variations in loss due to reflections, we have provided a simulation from WI which is shown in figure 5.2. In this simulation, a hexagon structure with a circumradius of 10m is placed directly between a transmitter and a receiver. Each of the six faces on the hexagon is made up a different material from the WI database. These materials include brick, concrete, asphalt, dry sand, dry earth, and wet earth; their material properties can be found in [32]. A smaller hexagon structure given the material properties of a perfect absorber is placed between the transmitter and receiver in order to prevent a LOS path. The result of this setup will be a specular reflection off each face, one for each material type. We vary the angle (θ) from the removed LOS path between 15, 30, 45, 60, and 75 degrees by changing the distance between the transmitter and receiver, recording the path loss (dB) from each path. The transmitter was set to transmit with 20 dBm at a center frequency of 2462 MHz. In table 5.1, we can see that a variation of 6 dB between material types is the best case given the angles of reflection.

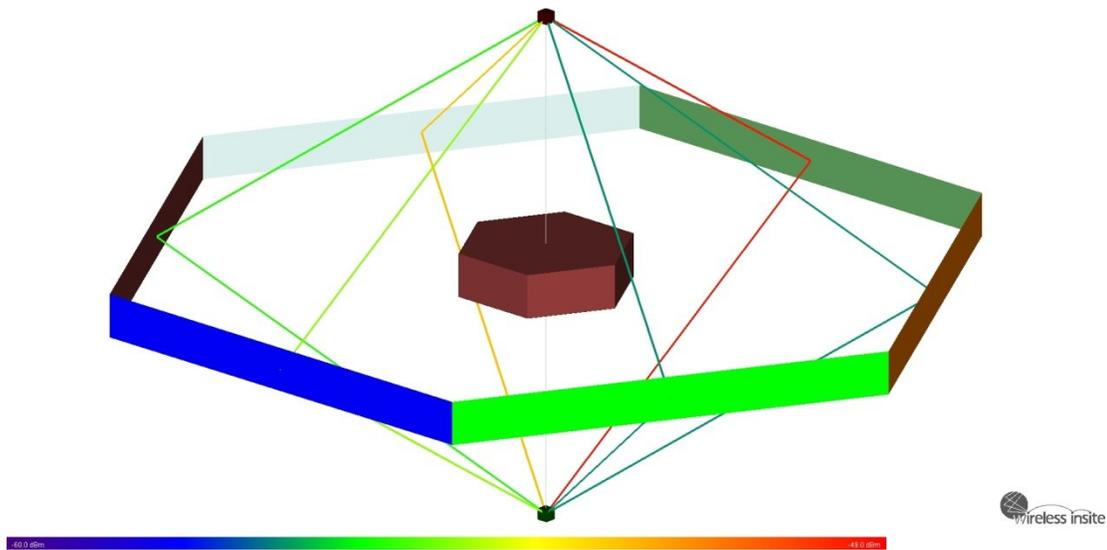


Figure 5.2: WI Simulation Testing Material Reflection Coefficients

Material \ θ	15	30	45	60	75
Brick	-92.05	-103.3	-78.88	-55.30	-46.47
Concrete	-97.67	-80.81	-60.04	-52.32	-52.68
Asphalt	-93.59	-86.30	-64.65	-54.38	-48.68
Dry Sand	-91.72	-135.9	-67.77	-56.53	-50.55
Dry Earth	-91.72	-131.9	-68.77	-56.53	-50.55
Wet Earth	-104.8	-71.56	-57.86	-49.32	-44.28

Table 5.1: Path Loss (dB) due to a Single Reflection off Different Materials

By observation, a low percentage of interactions in the WI representations of our field tests include a ground bounce. Because of this, and since the modeling of buildings is of greater focus in this thesis, the flat terrain will be comprised of a single material instead of a collection of several materials. As suggested in [36], we use measurements from our field tests to make preliminary comparisons, selecting the material that provides the closest match between predictions and measurements. After comparing the built-in terrain materials with our field test results, we found that the *Wet Earth* material was the closest match and will be used during our analysis.

5.2 Building Models

A map of the University of Maryland campus [43] showing the area that we modeled is shown in figure 5.3. Besides being relatively flat, the selected area offered the largest building density with a street layout best suited for an urban environment field test. This selection also prevented us from having to model a larger section of campus, which would have increased the difficulty of organizing and conducting field tests. A map of the selected area of campus is shown in figure 5.4, along with the buildings names and identification numbers in table 5.2.

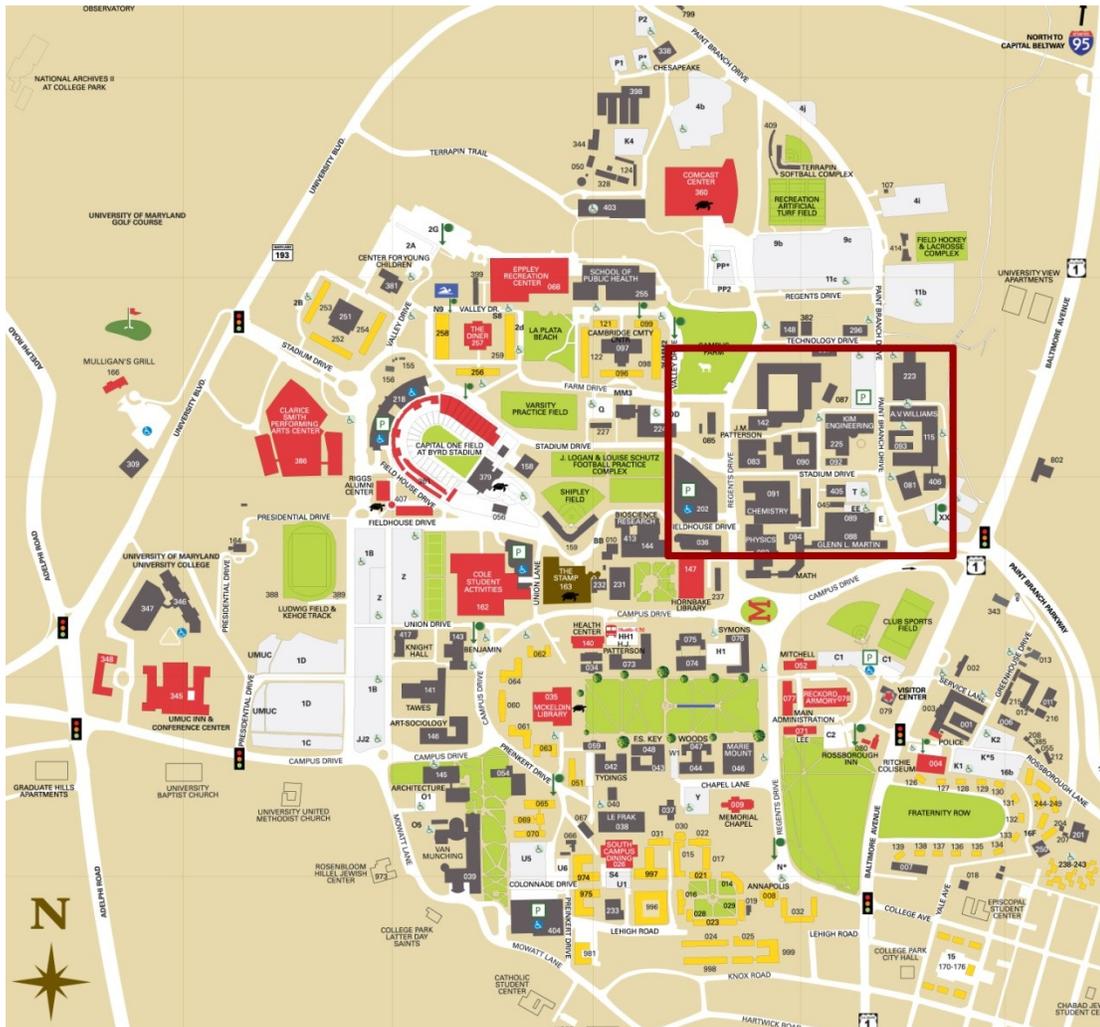


Figure 5.3: UMD Campus map Showing Section Modeled

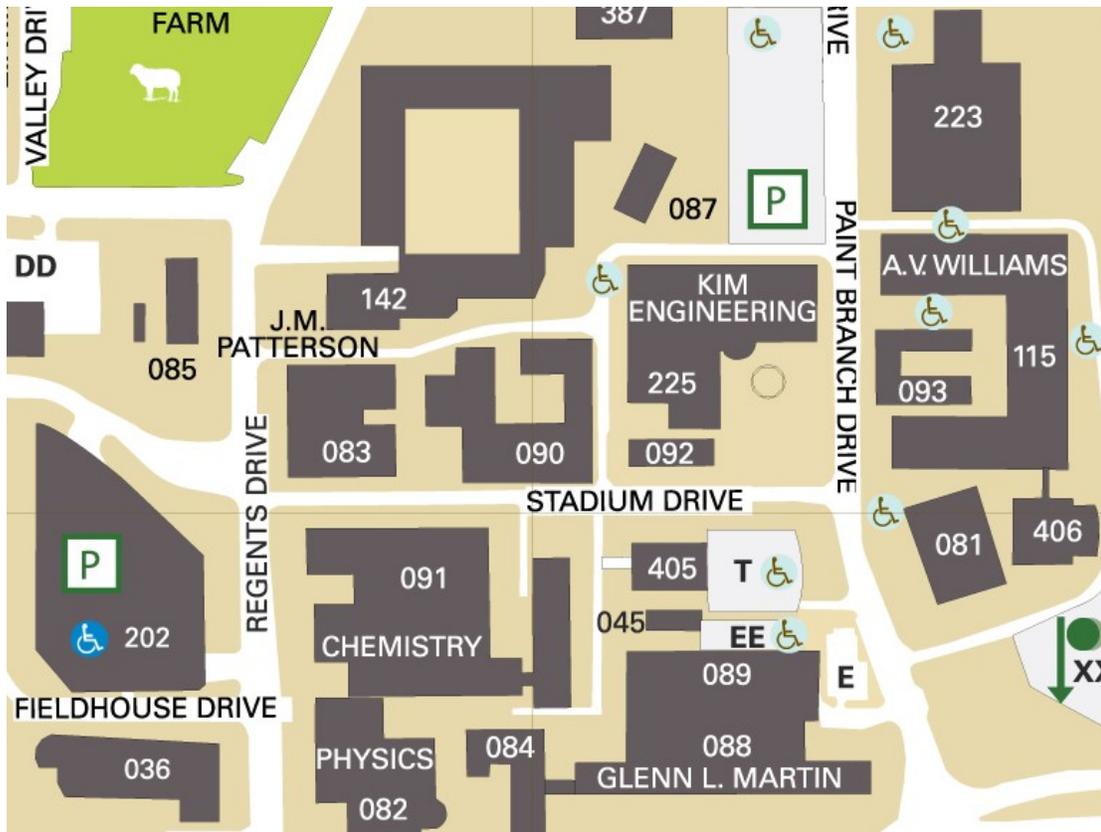


Figure 5.4: Modeled Section of UMD Campus

Name	Identification Number
A.V. Williams	115
Animal Science/Agricultural Engineering	142
Central Animal Resources Facility	087
Chemical and Nuclear Engineering	090
Chemistry	091
Computer Science Instructional Center	406
Energy Research Facility	223
Engineering Annex	093
Engineering Laboratory	089
Glenn L. Martin	088
Institute for Physical Science and Tech.	085
Instructional Television Facility	045
J.M. Patterson	083
Jeong H. Kim Engineering	225
Potomac	092
Regents Parking Garage	202
Satelite Central Utility Building 4	405
Wind Tunnel	081

Table 5.2: Modeled Buildings with Identification Numbers

One goal of this thesis is to give insight into how much detailed is needed in a modeled environment for a certain level of realism. To do this, we have created three models of varying detail for each building. Subsection 5.2.1 discusses the creation of the most detailed building models, V1, which emphasize both material separation and the depth of structures. Subsection 5.2.2 outlines the simplification process of V1, creating two less detailed models, V2 and V3, emphasizing building shells with material separation and building shells made of brick, the dominating material on the UMD campus. How these building models are placed in WI as well as the allowed interactions in our simulations are then discussed in subsections 5.2.3 and 5.2.4.

5.2.1 Model Creation

For the creation of the building models, we were able to obtain 3D building models in the Google Sketchup format from contacts at UMD. We were also able to obtain building floor plans, accurate building positions, and elevation data giving elevations of both the terrain and individual buildings. Although helpful, it was quickly recognized that the 3D models were accurate in two dimensions, but not three.

As stated in [32], the WI DXF converter can only convert DXF objects containing polylines, polyface meshes, and 3D faces using AutoCAD *grip points* so that the DXF converter can identify faces that are touching. The 3D building models obtained from UMD were first imported into AutoCAD to create the required DXF. All the building surfaces were then changed into 3D faces, each of which were then associated with AutoCAD *layers* according to their material type. Converting all

building surfaces into 3D faces instead of polylines and polyface meshes seemed to be the most compatible with WI.

To create the building models for our simulations we took a top-down approach, creating the models of highest detail and then applying a common simplification technique to all models. The models of highest detail, V1, emphasize material separation as well as the depth of building features, including windows, doors, columns, overhangs, stairs, ramps, and any structure unique to a building. Using the floor plans, we were able to determine the width and depth of these features. The elevation data was used to determine the height of structures at different points and the accurate placement of the buildings on the terrain. To determine the heights of individual features on each building, independent measurements had to be made. Since brick is the dominant surface material for almost all buildings on campus, it was relatively straightforward to approximate feature heights by counting bricks and knowing the measurements of the brick used. Although we went through the painstaking process of creating buildings of such detail, we hope to show that this level of detail is not required. It is also important to note that all buildings were created with a realistic terrain in mind. The terrain and building elevation data was used to add padding to each building, extending the model into the terrain enough so that if elevation is applied to the terrain model, no section of the building would be hovering above the terrain. Images of all the V1 building models placed in WI are shown in figures 5.5 and 5.6. In these images, the different colors represent the various building materials. For pictures of individual buildings, as well as the simplified models described in the next subsection, please refer to appendix A.



Figure 5.5: View of V1 Buildings Models from the North East

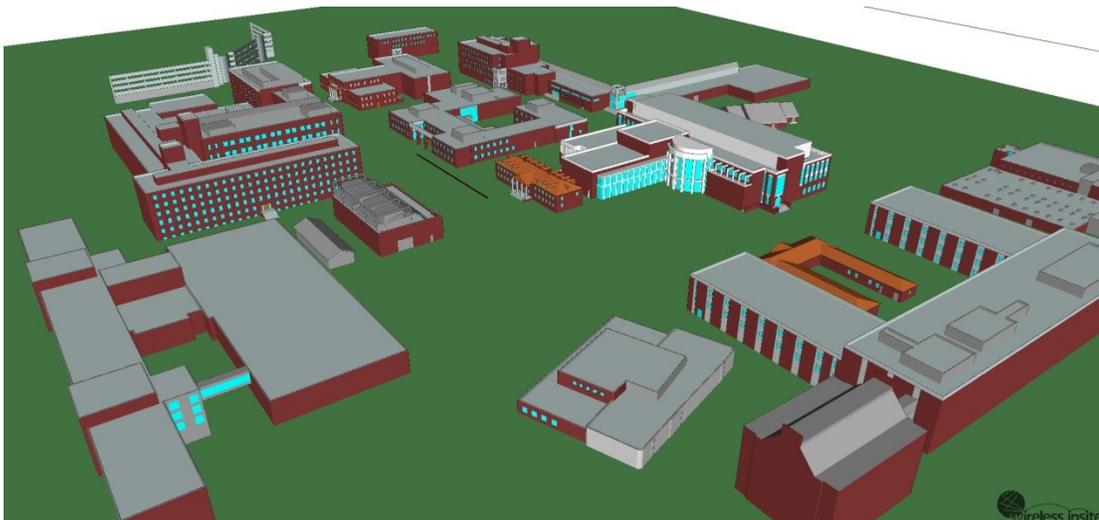


Figure 5.6: View of V1 Buildings Models from the South East

5.2.2 Building Simplification

This section describes the simplification techniques used to create two simplified versions (V2 and V3) of the building models created in the previous subsection. For V2, emphasis is placed in material separation, taking away the depth

of the building and its features created in V1. We first removed any depth due to the presence of windows or doors, bringing the feature materials flush with the main wall, shown in figure 5.7. Then details unique to a building are removed including stairs, ramps, columns, or structures external to the general makeup of the building, shown in figure 5.8. Finally, to make the buildings more box-like, we then simplified both the roof and the external wall structure. Any details such as slopes, ledges, overhangs, or other roof structures were removed unless the structure contributed to the general makeup of the building, resulting in a flat roof. For the walls of each side of the building, any small recessions were removed. An example of this simplification is shown in figure 5.9.

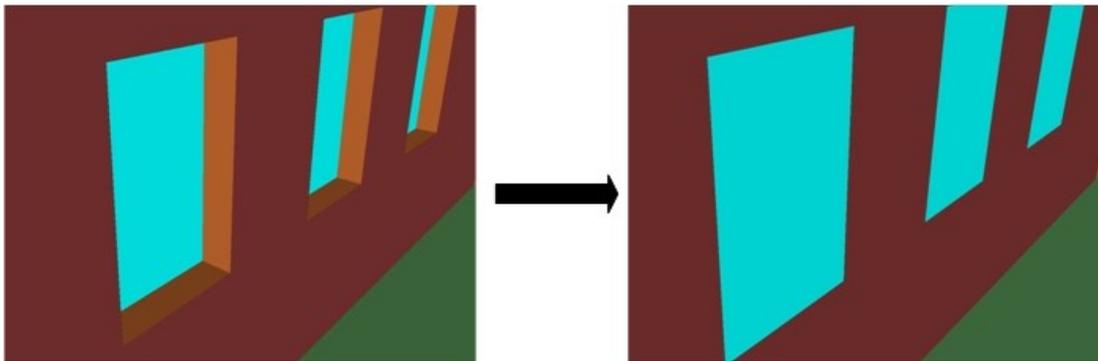


Figure 5.7: Example of Window Simplification

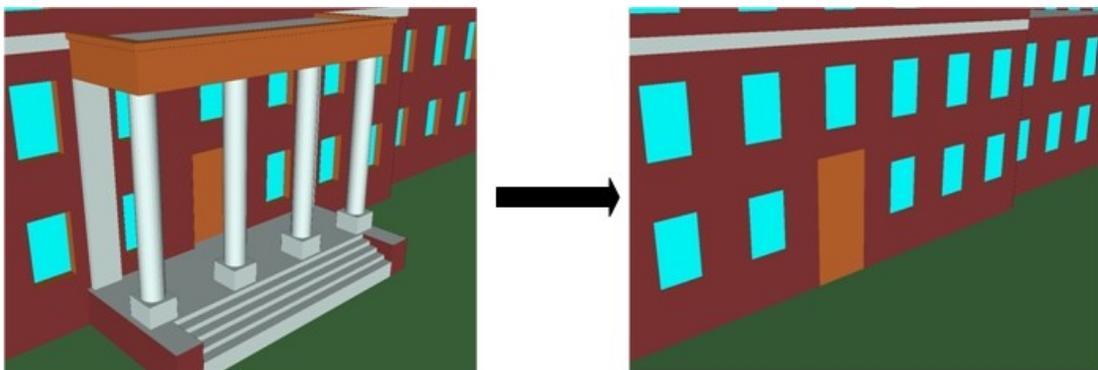


Figure 5.8: Example of External Structure Simplification

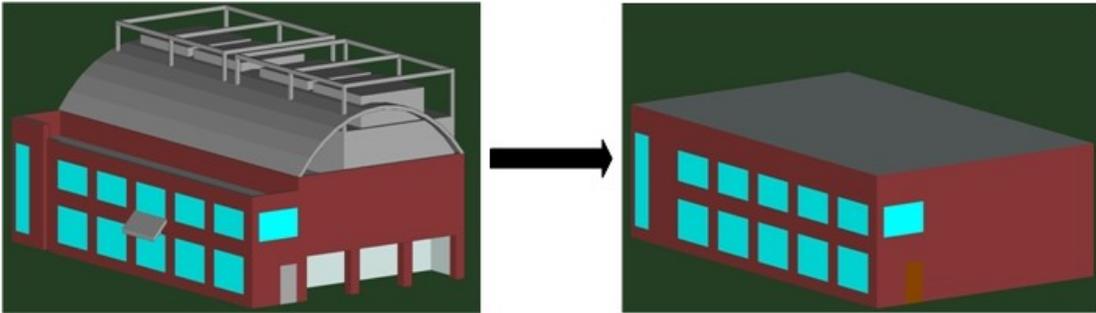


Figure 5.9: Example of Roof & Recess Simplification

The end result of the previous simplification techniques are basic building shells representing the general structures of buildings, while maintaining material separation. For V3, we take V2 and remove material separation by replacing all materials with the material that dominates the structure, in most cases brick. This produces a basic shell of a building consisting of one material, a model that can be easily created using the tools built-in to WI. Appendix A provides images of all buildings and their different versions, first showing V1 in comparison to the actual building and then the simplification process of V2 and V3. Table 5.3 provides the number of faces that make up each building model across each version, as well as the total number of faces when all buildings are placed in WI.

Building ID	Version1	Version2	Version3
115-406	2494	1034	20
142	1001	340	75
087	113	41	10
090	1069	309	31
091	3102	1099	81
223	480	90	29
093	270	93	11
088-089	1866	206	47
085	207	75	5
045	18	5	5
083	825	249	23
225	1217	674	88
092	866	162	5
405	284	55	5
081	300	112	26
202	280	45	44
Full Model	14392	4589	505

Table 5.3: Number of Faces in each Building Model as well as the Full Model

5.2.3 Placement in WI

Using the DXF converter, the building models are imported into WI as floor plans since the city format was intended for simplified building structures. Floor plans were intended for indoor simulations, but are also used for mixed environments and allow us to represent our highly detailed models. When importing each building model, each layer created in AutoCAD is associated with a default building material from the WI material database. Initial analysis will be done using these default materials. If we find during our analysis that one material is dominating the interactions in the environment, future work will fit our field test measurements to more suitable permittivity and conductivity parameters for that material.

To position the buildings in the ray tracing environment, we use the data obtained from our contacts at UMD. But our decision to use a flat terrain creates a

bias in the model that depends on how the building heights are set. Because there is little difference in elevation between the center and east side of the selected area (the sites of the field tests), we start by placing the eastern most building, 115, working our way west. Building 115 is placed at the proper elevation according to the side closest to Paint Branch Drive (facing west). The buildings to the north and south of 115 are placed by their elevation in relation to 115. From there, the building directly west of 115 is placed by its elevation in relation to 115, then placing the buildings to the north and south of that building. This process is repeated until every building is placed in the environment. Because of this placement, buildings along Regents Drive are placed higher above the terrain than they would be in a realistic terrain since these buildings undergo the most change in elevation. The change in elevation from north to south is 1.31m and is fairly constant, creating little bias in those directions.

5.2.4 Allowed Interactions

For all of our simulations, no transmissions are allowed through any of our building models. Almost all buildings are predominantly brick and rather large making it very unlikely that a radio wave would propagate through a building and still have a significant impact in the power at the receiver. Moreover, it is impossible to know what interactions are occurring between entry and exit of a building due to several unknowns such as moving people and knowledge of the details in the environment. For these reasons, we focus on reflections and diffractions in the outdoor urban environment.

In this thesis, we will evaluate the accuracy of WI compared to our field test measurements by first limiting the interactions to reflections, and then adding a single

diffraction. One criticism of ray tracing is the large amount of time spent doing computations. While there are many other factors that contribute to computational complexity, much time is needed to search for and cast rays from diffraction points. By conducting our simulations with and without diffractions, we hope to explore the benefits of adding such interactions. We limit the number of diffractions to one with the assumption that since all our nodes are near earth, reflections will contain larger power contributions than two or more diffractions due to the close proximity of buildings. Multiple reflections should be able to reach most areas in our field test scenarios. We also limit the number of reflections to 6, which by observation produces a sufficient number of ray paths given the size of our scenarios. The allowed number of reflections should be increased for larger field tests.

5.3 Transmitters & Receivers

5.3.1 Placement

Due to uncertainties in both the field tests and WI, it is unclear how the placement of transmitters and receivers will affect the accuracy of WI predictions. In the conducted field tests, mobility scenarios were generated from UNIX timestamps recorded by the user when walking over a predetermined location. All field tests are prone to human error and it is difficult to determine how accurately these locations were marked by each participant. Vast variations and deep fades in measurements can result from changing the position of the transmitter or receiver by even a wavelength. Despite careful planning, field tests will always contain inaccuracies that need to be accounted for.

WI is a deterministic method of calculating received power and other signal characteristics. Received power is calculated either using all phase information or by correlating ray paths into groups, combining those correlated paths with phase before adding the powers of each group. This method of calculating power can, to some extent, model the fast fading due to the constructive and destructive interference of waves in the environment, resulting in large variations and deep fades within predictions. Some form of aggregation is needed for both field test data and WI predictions to mitigate these effects.

To remove the variations of fast fading from our field test data while retaining the effects of shadowing, we use a form of low-pass filter that is simple to implement, a moving average. The number of measurements to average (window size) is determined by expanding the window size until most sharp variations are removed without changing the overall shape of the measurements over time. Despite its simplicity, this method proves to be very effective.

In order to perform the same aggregation to both field test measurements and the predictions of WI, receivers are placed according to the velocity of the node as well as the transmitted beacon interval while making sure that each predetermined location from the field test is represented by one receiver. Since the transmitters in our field tests are stationary, they are simply modeled as a single point, only performing the moving window aggregation around each main receiver point using the added receiver points.

It is important to note that the proposed placement method based on node movement and beacon interval will only work for our simplified test scenarios. For

tests involving two nodes and a mobile transmitter, slight modifications to this method may still work. But when more than two nodes make up a field test scenario, a different placement method must be used to better represent the mobile scenario. Proposed techniques will be discussed in the future work section of chapter 8.

5.3.2 Equalizer & Correlator Modeling

As a consequence of multipath, the contributions of ray paths at the receiver vary in time of arrival due to their differences in path lengths. Because of these delays, part of the transmitted symbol spreads into subsequent symbols creating noise and making correct detection of those symbols increasingly difficult. This type of distortion is called Intersymbol Interference (ISI). To mitigate the effects of ISI, filters are put in place at the receiver to attempt canceling the ISI introduced by multipath, also known as an *equalizer*.

While many forms of equalizers are used, manufacturers of commodity hardware rarely release the details of their hardware implementations. For our wireless chipset, the Atheros AR5424, we were able to find a technical overview document claiming our receiver contained an "advanced wideband receiver with best path sequencer for better range and multipath resistance than conventional equalizer-based designs". Without any technical specifications, it is unclear how exactly the "best path sequencer" functions. For our analysis, we make the assumption that the "best path sequencer" uses some method to keep track of multipath peaks in the correlator and selects the strongest signal within a time window set by the modulation scheme. To test our assumption, we will vary the number of paths being combined at

the receiver in WI between one and fifty and compare the results with our collected field test measurements.

The nodes in our field tests are transmitting 802.11b beacons at a rate of 1 Mbps spread using a Direct-Sequence Spread Spectrum (DSSS) modulation technique. DSSS uses a pseudonoise spreading code to spread transmitted data over wide bandwidth. This pseudonoise code is the same for every user in the network and is made up of an 11 bit barker code, each bit represented by a chip. Transmitting 1 Mbps using an 11 bit barker code allows for the transmission of 1 Msymbol/s giving a symbol period of 1 microsecond. To more accurately model the DSSS correlator in the receiver we use the output filter in WI to filter out any paths arriving after the symbol period of 1 microsecond from the first path that arrives (excess time of arrival). In reality, the period of time in which the correlator can distinguish a signal from multipath noise is probably much less than the symbol period. Applying the excess time of arrival filter of 1 microsecond is a loose restriction that will be re-examined after initial evaluation.

In our analysis, we will also look at the 'summing complex electric fields' parameter in WI, describing how the receiver is adding the individual contributions of each ray path. This setting will affect our received power when increasing the number of paths to combine. As described in section 3.4.7, the 'all phase' option sums the powers of each ray path using all phase information. The 'correlated phase' option combines the phase of paths that follow nearly the same path in the environment, then summing the powers of the correlated groups. Both of these options will be evaluated,

ignoring the 'none' option. This option would best represent a rake receiver, a receiver that is better fit to coherently add multiple multipath components.

5.3.3 Antennas

In urban environments, there exist multiple reflected and diffracted paths other than LOS between a transmitter and a receiver. The signal characteristics of reflected paths vary depending on the surface material properties, polarization, and incidence angle. The multipath components arrive at the receiver at different angles, depending on where the signal was radiated from, the interactions undergone, and the beamwidths of the used antennas. Therefore, the antenna pattern can have a significant impact on the transmission and reception of radio waves between the transmitter and receiver.

Along with the parameters already being tested, we'll explore the accuracy of WI predictions using both ideal and realistic omnidirectional antennas. For the ideal omnidirectional antenna, we use the built-in vertically polarized half-wave dipole. Its vertical and horizontal gain patterns are shown in figure 5.10.

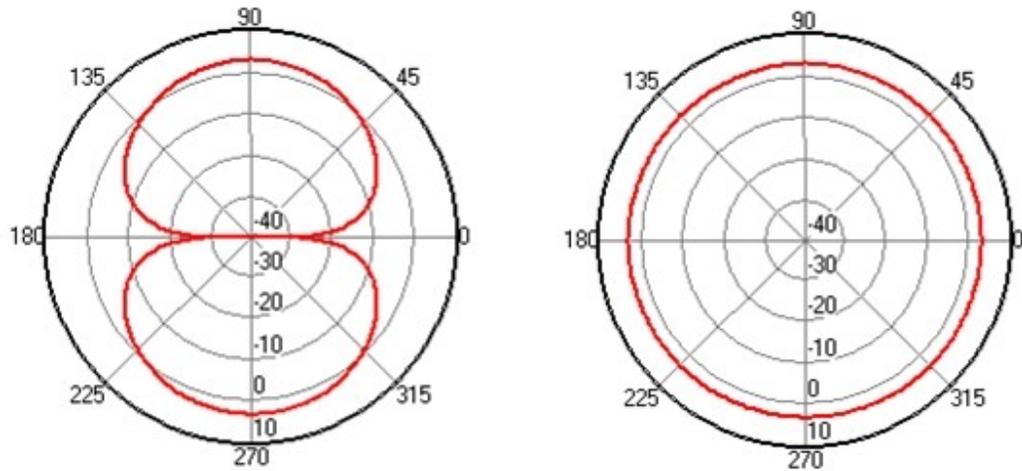


Figure 5.10: Vertical and Horizontal Gain Patterns of 3 dBi Half-wave Dipole

To determine the antenna radiation patterns of the antennas used in our field tests, we used measurements collected in an anechoic chamber. An anechoic chamber is a room designed to provide a controlled environment, insulated from exterior sources of noise. The walls are coated with RF absorbers, preventing any multipath between a transmitter and receiver caused by reflections off the walls leaving only the LOS path. This environment allows us to precisely compute the antenna gain pattern of our antennas without any interference. Figure 5.11 show a picture of the anechoic chamber used. In this scenario, a simple diffraction test was being conducted between a transmitter and a receiver using a sheet of metal. The resulting vertical and horizontal antenna gain patterns of testing one antenna are shown in figure 5.12. The two other antennas were also tested, producing near identical results. The maximum gain of the main lobe was measured to be 3.127 dBi as opposed to the 5dBi claimed by the manufacturer. Because of this discovery, we gave our ideal antenna a 3.127 dBi gain to focus our analysis on the differences in radiation patterns.



Figure 5.11: Anechoic Chamber used to Measure Antenna Gain Patterns

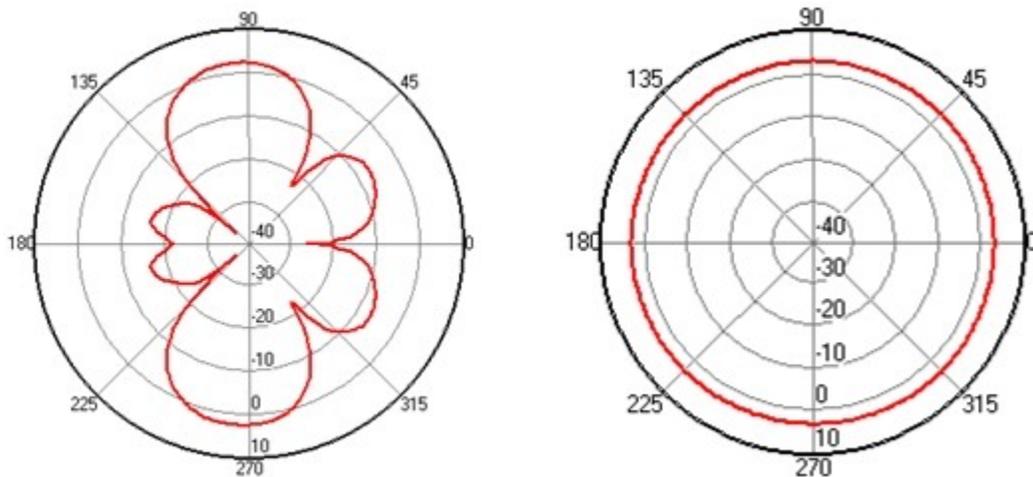


Figure 5.12: Vertical and Horizontal Gain Patterns of 3.127 dBi Omnidirectional Antenna

To continue our goal of accurate receiver modeling, we set the received power threshold on the receiver's antenna, disregarding any arriving paths with a power beyond the receiver's capabilities. While it's tempting to just set this power to the

receiver sensitivity, we must take into the consideration multiple waves combining at the antenna. When RF signals arrive at nearly the same time they undergo constructive and destructive interference at the antenna with the receiver only seeing the combined wave. It is possible that two signals of lower power, beyond the receiver's capabilities, combine into a signal that the receiver can demodulate. Because of this phenomenon, we set the received power threshold to -115 dBm. This value is a rough approximation and is not claimed to be the correct value. Since we are only looking at the top so many paths that arrive at the receiver within a certain time frame, this will only have an effect when increasing the total number of paths being combined.

5.3.4 Waveforms

Since DSSS essentially phase modulates a continuous sine wave instead of a pulse-type waveform, we select a sinusoidal waveform within WI. For our simulations we set the effective bandwidth to 22 MHz and the center frequency to 2462 MHz (channel 11) since this was the least congested channel that was used during our field tests.

5.3.5 Collection Radius & Ray Spacing

The SBR ray tracing algorithm used by the full 3D propagation model will find a number of common rays that have followed essentially the same path, given a reasonable collection radius and ray spacing have been selected. These rays represent

the same wave front and are removed as duplicate paths to avoid over predicting the received power. But, there is no analytical solution for selecting these parameters.

To select a reasonable collection radius, several considerations must be made. If the collection radius is too small, the ray spacing will need to be set at a small angle, greatly increasing the ray tracing run time. Making the collection radius too large will include ray paths that never would have hit the receiver in a real world environment. According to [32], convergence tests are the best way to develop an intuition for selecting these parameters in different urban environments. In order to eliminate unwanted ray paths with the suggested 2.5m collection radius, we use a radius of 1m. We then tested to see any differences between the default ray spacing of 0.25° and the recommended 0.2° . After noticing no difference in results, we reduced the ray spacing to 0.3° and notice a reduction of paths arriving at the receiver. For our ray tracing simulations, a collection radius of 1m with a ray spacing of 0.25° was determined to be a reasonable selection.

Chapter 6: Field Test Descriptions

In this chapter, we describe two field tests we have conducted on the UMD campus to evaluate the performance of the BounceHaus testbed. The first field test was conducted between the A.V. Williams building (#115) and the Engineering Annex (#093) and is referred to as the *AVW* field test. The second test was conducted on Stadium Drive and the off street between the Kim building (#225) and the Chemical and Nuclear Engineering building (#090) and is referred to as the *Stadium Drive* field test.

6.1 AVW

6.1.1 Description

A two-dimensional overview of the AVW field test scenario is shown in figure 6.1 (obtained from Google Maps). In this field test, a stationary transmitter, represented by the black circle, was positioned between the Engineering Annex and the A.V. Williams buildings. We then created three lines spanning 25m parallel to the Annex building and moving away from the transmitter, each consisting of 26 measured positions from which to record measurements from. A separation of 1m was kept between lines. The purpose of this scenario is to create both LOS and NLOS paths between the transmitter and receiver in close proximity.

The transmitter in this field test was set to transmit on channel 11 (the least congested channel) at 20 dBm with a beacon interval of 100ms, or 10 beacons/second. Both laptops were held just above the waist with an approximate

antenna height of 1.3m which was maintained to the best of our ability throughout the entirety of the field test. Both participants faced each other and held the laptops so the antenna experienced no obstructions from the laptop monitor. This was done to mitigate near field effects of the human body and laptop monitor. While standing in this position and given this scenario, it is very unlikely that a path between the transmitter and receiver was attenuated by these near field effects. In addition, this test was conducted on a weekend when there was little activity on the UMD campus. If any cars pulled into the testing area, the measurements were repeated; the same being true for people walking by.

We conducted several measurements using the positions on all three lines. First, we took stationary measurements from all 26 points on each line. Measurements at each point were taken for approximately two minutes, giving us enough data to average out any undesired interference. After the stationary measurements, we conducted three motion tests, one for each line. The participant holding the receiver started at the point on the line furthest away from the transmitter, then walking towards the transmitter at approximately 0.5m/s and marking the UNIX timestamps when crossing over each of the 26 predefined points by using a clock script. Moving the receiver towards the transmitter eliminated any near field effects due to the human body.

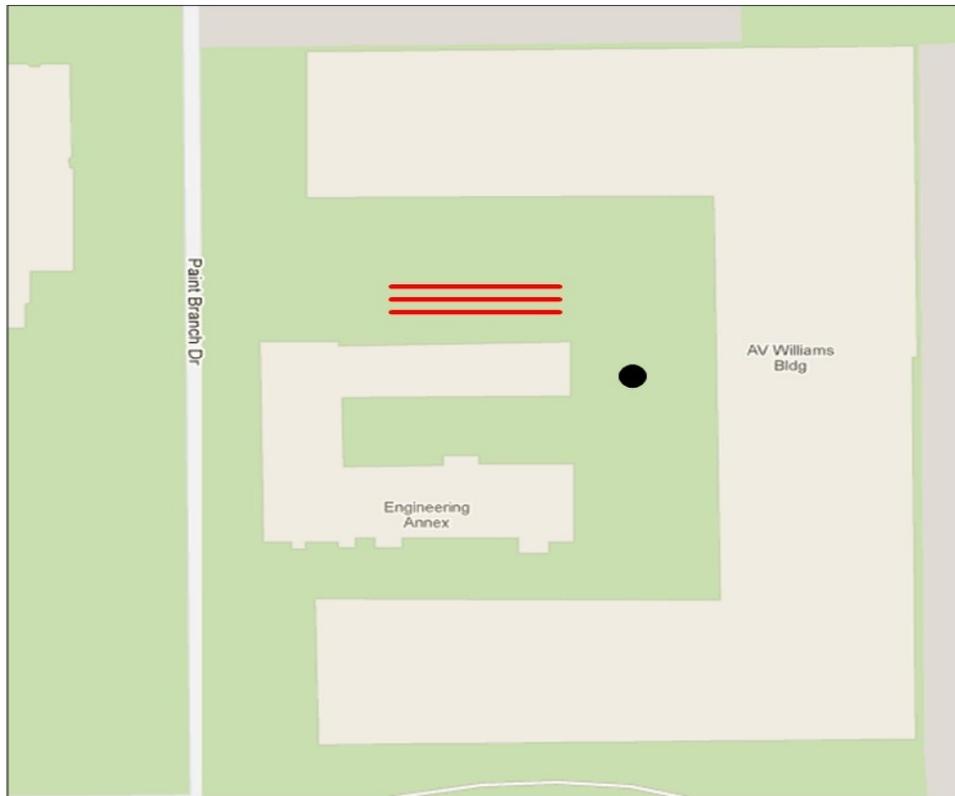


Figure 6.1: 2D Overview of the AVW Field Test Scenario

6.1.2 Data Analysis

To first mitigate the effects of fast fading, we perform a moving average with a window size of 21 on the motion data for each line. Figure 6.2 shows both the original measurements from line1 as well as the moving average. It can be seen that the overall shape of the measurements over time is kept intact, preserving the effects of shadowing. It is important to note that before any averaging of RSS is done, the measurements are first converted back to milliwatts, and then converted back to dBm after averaging. Averaging the RSS in dBm is actually performing a geometric mean instead of an arithmetic mean which can alter averages in the presence of high variance.

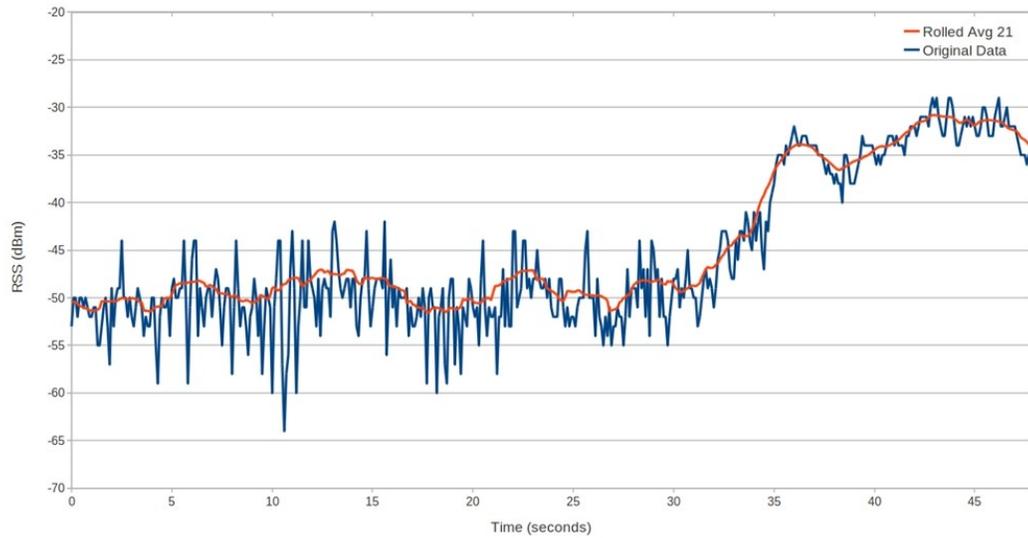


Figure 6.2: Moving Average on Line1 Mitigating Effects of Fast Fading

For all comparisons, we first compute this moving average and then use the points associated with each of the 26 marked locations during the field test. Figures 6.3, 6.4, and 6.5 show these points taken from the moving averages of line1, line2, and line3, respectively, and plot them against the stationary measurements that were also taken for each line. While the motion measurements match the stationary measurements for line1 very well, there are two areas within both line2 and line3 that show major differences. After looking at the original data it was seen that these stationary measurements were within the variance of the original motion measurements. Even with these slight variances, the RMSE of the line1, line2, and line3 comparisons were 2.11 dB, 3.33 dB, and 2.05 dB, well within the margin of error for repeating field test measurements [36].

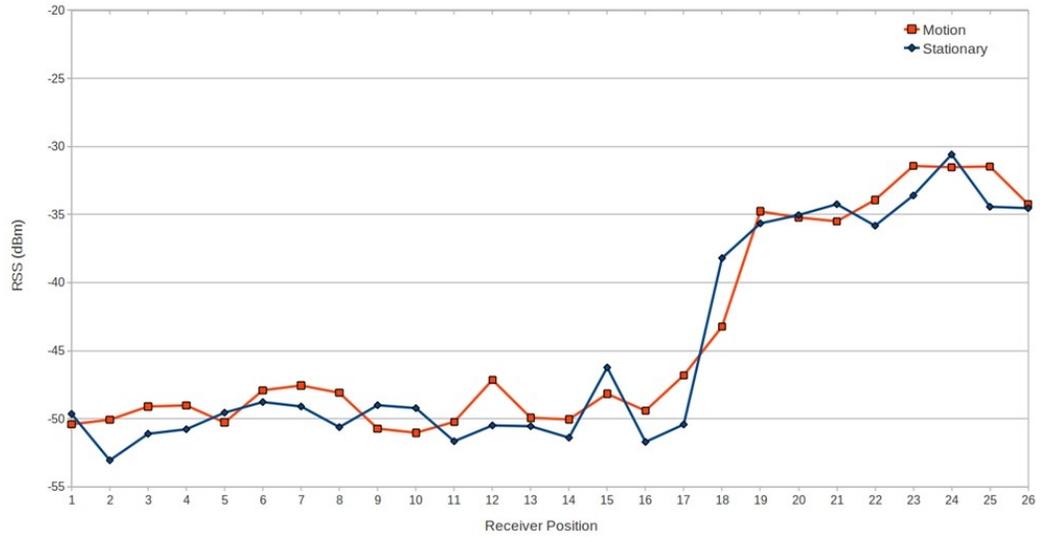


Figure 6.3: Line1 Motion Points after Moving Average vs. Stationary Points

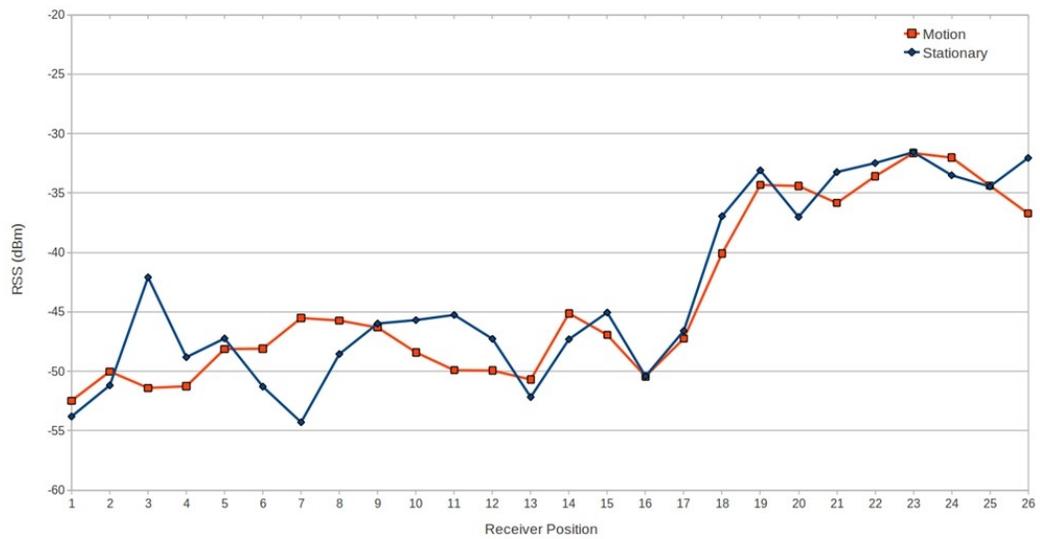


Figure 6.4: Line2 Motion Points after Moving Average vs. Stationary Points

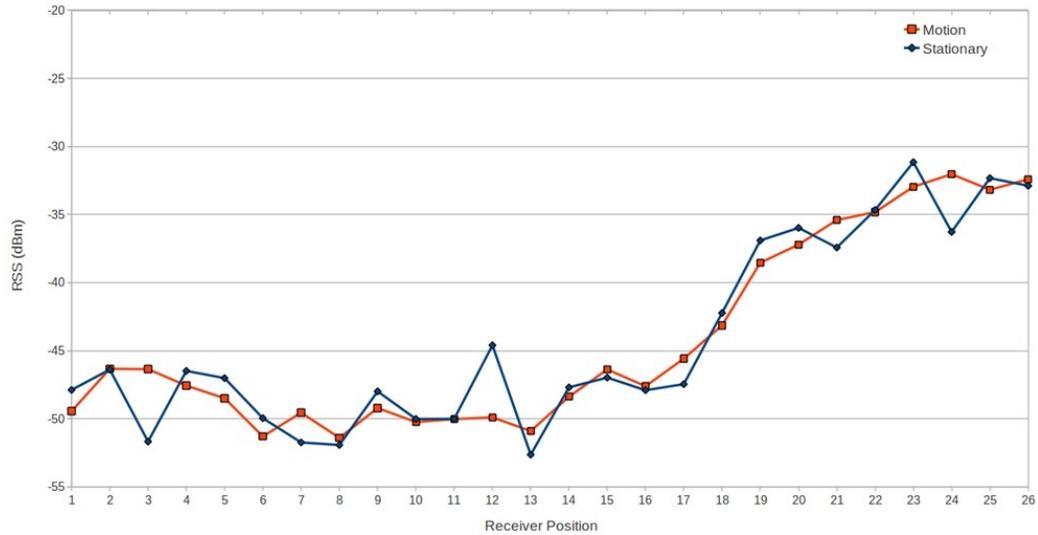


Figure 6.5: Line3 Motion Points after Moving Average vs. Stationary Points

6.1.3 Representation in WI

As mentioned in section 5.3.1, to compare the field test measurements with WI predictions we use the node velocity and beacon interval to approximate the positions upon beacon reception. This allows us to place receivers in the approximated locations in order to perform a moving average on the WI data, mitigating the effects of fast fading before making direct comparisons. In this field test, nodes had an approximate velocity of 0.5m/s and the beacon interval was set at 100ms, therefore there were approximately 20 beacon readings per meter. In WI, each line was modeled as a receiver route with a spacing of 0.05m resulting in lines containing 501 receivers. Given that we used a window of 21 measurements for our field test averaging, the same window was applied to the WI predictions. Figure 6.6 shows the placement of the transmitter and receiver routes in WI.

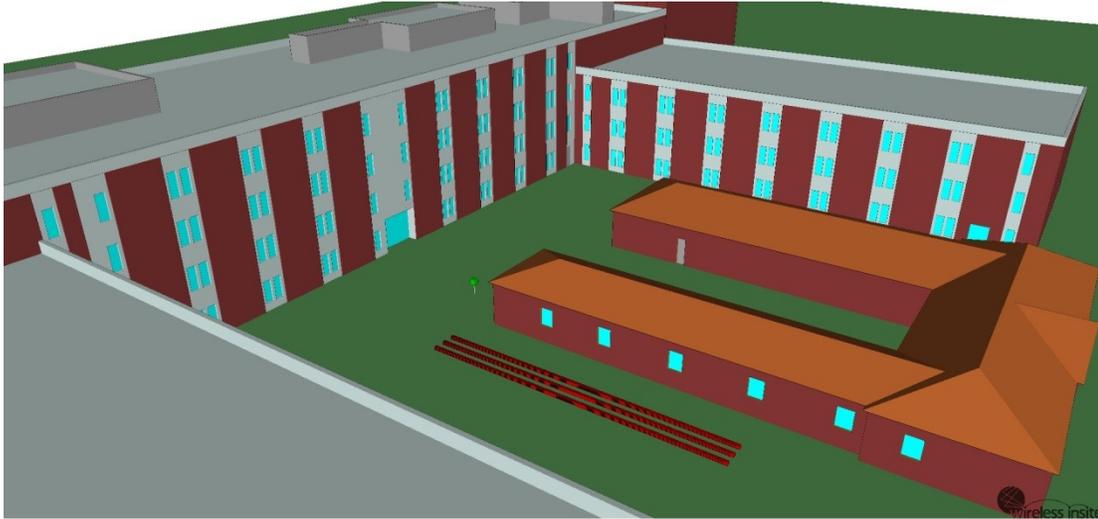


Figure 6.6: Transmitter and Receiver Placement in WI for AVW Field Test

6.2 Stadium Drive

6.2.1 Description

A two-dimensional overview of the Stadium Drive field test scenario is shown in figure 6.7 (obtained from Google Maps). Points A,C, and D represent the ends of the paths (in purple) while point B represents the intersection between them. Each path segment (B to A, B to C, and B to D) is 30m long with points measured every 2m in each direction away from B. In this field test, a stationary transmitter is positioned at point D while a mobile receiver is positioned at node A, walking from A to B to C while collecting signal measurements from the transmitter. The receiver will be moving at approximately 1m/s over a distance of 60m and marking the UNIX timestamp when the node crosses over each of the 31 predefined points. This scenario was also conducted in reverse order with the receiver moving from C to B to A. A comparison between the two tests will be done in the next section. These test were

designed to subject our nodes to a more realistic urban environment, unlike the simple two building scenario in the AVW field test.

The transmitter in this field test was set to transmit on channel 11 (the least congested channel) at 20 dBm with a beacon interval of 50ms, or 20 beacons/second. Due to the positions of the transmitter and receiver during this scenario, we created a mount to raise the antenna height to approximately 2m (figure 6.8). With the antennas mounted above field test participants, the participant is most likely to be in the path of any signal heading toward the null below the main lobe of our antenna pattern, shown in figure 5.12. This reduces the chance of the participant's body attenuating any path of significance. When analyzing the field test results, the loss due to the added cables is taken into consideration. As done in the previous field test, tests were discarded if any cars or people came into the area of interest.

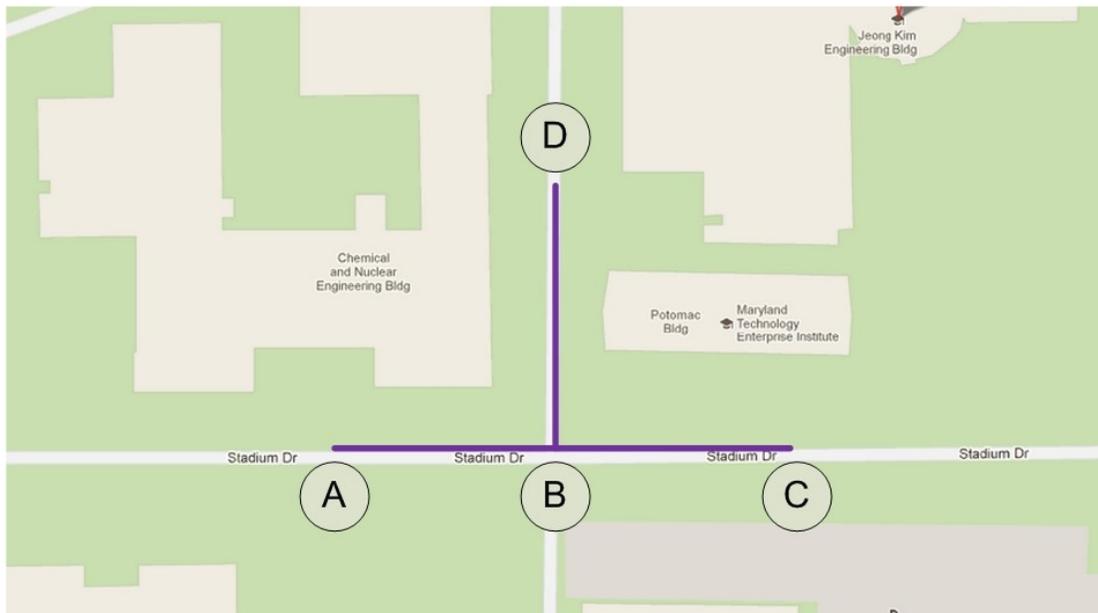


Figure 6.7: 2D Overview of the Stadium Drive Field Test Scenario



Figure 6.8: Antenna Mount used During Stadium Drive Field Tests

6.2.2 Data Analysis

To first mitigate the effects of fast fading, we performed a moving average with a window size of 21 on the RSS measurements. Figure 6.9 shows both the original measurements as well as the moving average. It can be seen that the overall shape of the measurements over time is kept intact, preserving the effects of shadowing. For all comparisons, we first compute this moving average and then use the points associated with each of the 31 marked locations during the field test. Figure 6.10 shows the averaged points compared to the moving average of the same field test

conducted in reverse order. Both traces are near identical with a RMSE of 1.49 dB.

Because of the accuracy between both traces, we will select the original trace to use

in our comparisons with WI.

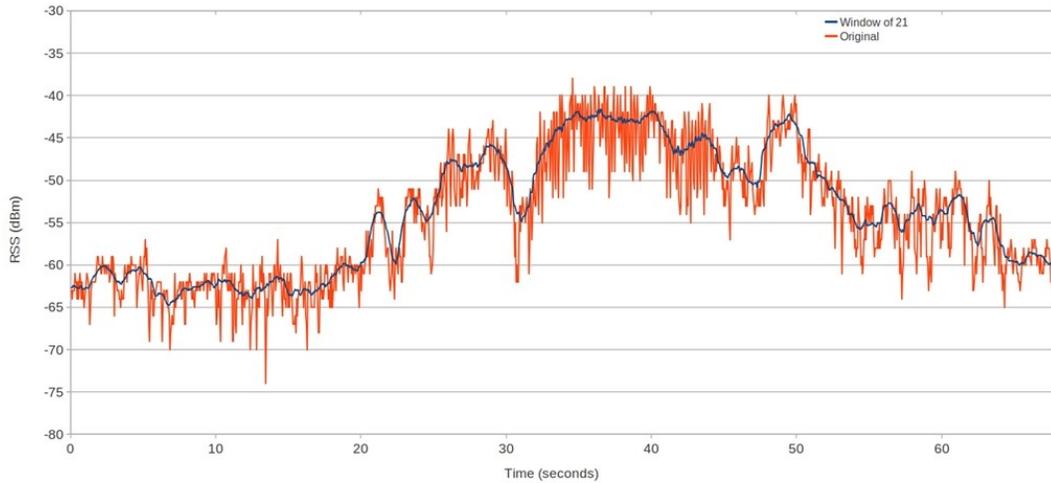


Figure 6.9: Moving Average on Stadium Drive Mitigating Effects of Fast Fading

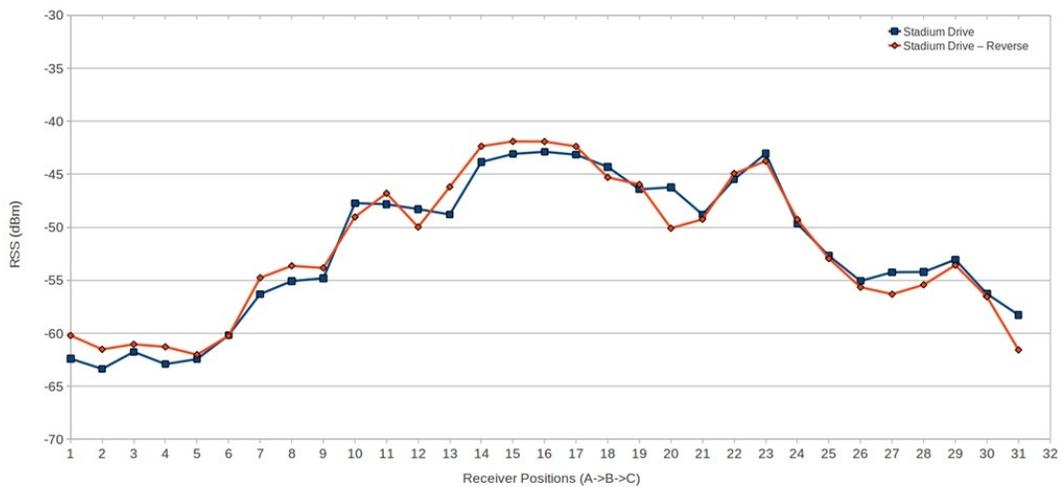


Figure 6.10: Stadium Drive vs. Stadium Drive Reverse at Measured Positions

6.2.3 Representation in WI

Receiver positions are determined with the same method used in the AVW field test. In this field test, the receiver had an approximate velocity of 1m/s with a beacon interval of 50ms, therefore there were around 20 beacon readings per meter, the same as in AVW. In WI, the path of the receiver was modeled as a receiver route with a spacing of 0.05m resulting in a line containing 1201 receivers. Given that we used a window of 21 measurements for our field test averaging, the same window being applied to the WI predictions. Figure 6.11 shows the placement of the transmitter and receiver routes in WI.



Figure 6.11: Transmitter and Receiver Placement in WI for Stadium Drive Field Test

Chapter 7: Results & Discussion

7.1 Results

To analyze the performance of WI predictions based on various parameter combinations, we propose the use of two metrics. The first metric is the overall RMSE and is defined by:

$$RMSE \text{ (dBm)} = \sqrt{\frac{\sum_{i=1}^n (x_{1,i} - x_{2,i})^2}{n}} \quad (7.1)$$

where n is the number of samples and $x_{1,i} - x_{2,i}$ represents the error between the field test measurement and predicted value from WI. Expressing this formula in words, differences between the observed and predicted values are each squared and then averaged over the sample, before taking the square root of the average. The RMSE provides a non-linear fit and gives the distance, on average, of a data point from this fitted line. Since errors are squared before being averaged, the RMSE applies a higher weight to large errors which is useful when large errors are undesirable.

The second metric is the *mean error*, which is the overall sum of the errors between observed and predicted values scaled by the number of predictions. Other research [28] has shown that some path loss models may have a cost/benefit for under or over-predictions. Over-predictions (under-predicted received signal strength) are represented by higher positive values for this metric. Under-predictions (over-predicted received signal strength) are represented by large negative values. A score of zero is given for models that have an equal amount of under and over-predictions.

Overall, using RMSE and mean error for performance evaluation will give us a sense of both the bound of the error as well as the tendency to over or under-predict.

In the following subsections, we present the RMSE and mean error results from the AVW and Stadium Drive field tests. We will refer back to these tables in the following sections. Note that the best value for each column is highlighted in red for each table.

7.1.1 AVW

AVW Line1 – RMSE								
Building V1								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	4.04	4.04	4.17	4.17	4.04	4.04	4.17	4.17
2 Paths	2.38	3.22	2.44	3.29	2.38	3.22	2.44	3.29
3 Paths	2.18	2.31	2.23	2.39	2.18	2.31	2.23	2.39
4 Paths	3.00	2.08	3.04	2.15	3.00	2.08	3.04	2.15
5 Paths	3.87	2.02	3.92	2.07	3.87	2.02	3.92	2.07
10 Paths	4.50	2.06	4.53	2.10	4.50	2.07	4.53	2.11
50 Paths	4.67	2.21	4.70	2.24	4.64	2.25	4.69	2.27
Building V2								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	4.04	4.04	4.17	4.17	4.04	4.04	4.17	4.17
2 Paths	2.38	3.22	2.44	3.29	2.38	3.22	2.44	3.29
3 Paths	2.21	2.44	2.25	2.51	2.21	2.44	2.25	2.51
4 Paths	2.95	2.26	2.99	2.32	2.95	2.26	2.99	2.32
5 Paths	3.80	2.19	3.85	2.25	3.80	2.19	3.85	2.25
10 Paths	4.60	2.24	4.65	2.28	4.77	2.92	4.79	2.88
50 Paths	4.89	2.32	4.92	2.36	4.99	3.17	5.01	3.12
Building V3								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	4.04	4.04	4.17	4.17	4.04	4.04	4.17	4.17
2 Paths	2.38	3.22	2.44	3.29	2.38	3.22	2.44	3.29
3 Paths	2.26	2.26	2.29	2.34	2.26	2.26	2.29	2.34
4 Paths	3.10	2.09	3.14	2.15	3.10	2.09	3.14	2.15
5 Paths	4.06	2.04	4.11	2.09	4.06	2.04	4.11	2.09
10 Paths	4.60	2.11	4.64	2.15	4.70	2.24	4.72	2.27
50 Paths	4.85	2.21	4.90	2.24	4.99	2.43	5.03	2.45

Table 7.1: RMSE for AVW Line 1

AVW Line1 – Mean Error								
Building V1								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	3.68	3.68	3.82	3.82	3.68	3.50	3.82	3.65
2 Paths	1.32	2.83	1.37	2.91	1.32	2.83	1.37	2.91
3 Paths	0.06	1.65	0.13	1.75	0.06	1.65	0.13	1.75
4 Paths	-0.92	1.17	-0.86	1.25	-0.92	1.17	-0.86	1.25
5 Paths	-1.81	0.87	-1.77	0.92	-1.81	0.87	-1.77	0.92
10 Paths	-2.31	0.16	-2.28	0.22	-2.33	0.14	-2.31	0.20
50 Paths	-2.45	-0.27	-2.42	-0.21	-2.37	-0.40	-2.35	-0.33
Building V2								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	3.68	3.68	3.82	3.82	3.29	2.35	3.43	2.53
2 Paths	1.32	2.83	1.37	2.91	1.32	2.83	1.37	2.91
3 Paths	0.15	1.72	0.22	1.80	0.15	1.72	0.22	1.80
4 Paths	-0.89	1.27	-0.84	1.34	-0.89	1.27	-0.84	1.34
5 Paths	-1.77	0.97	-1.73	1.02	-1.77	0.97	-1.73	1.02
10 Paths	-2.43	0.27	-2.40	0.33	-2.55	-0.71	-2.51	-0.61
50 Paths	-2.64	0.01	-2.61	0.06	-2.68	-1.17	-2.63	-1.06
Building V3								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	3.68	3.68	3.82	3.82	3.29	3.26	3.43	3.40
2 Paths	1.33	2.83	1.37	2.91	1.33	2.83	1.37	2.91
3 Paths	-0.01	1.53	0.08	1.63	-0.01	1.53	0.08	1.63
4 Paths	-1.11	1.10	-1.06	1.17	-1.11	1.10	-1.06	1.17
5 Paths	-2.08	0.76	-2.05	0.81	-2.08	0.76	-2.05	0.81
10 Paths	-2.41	0.10	-2.38	0.15	-2.45	-0.19	-2.42	-0.13
50 Paths	-2.69	-0.16	-2.69	-0.10	-2.72	-0.54	-2.70	-0.48

Table 7.2: Mean Error for AVW Line 1

AVW Line2 – RMSE								
Building V1								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	4.56	4.56	4.66	4.66	4.56	4.56	4.66	4.66
2 Paths	3.11	3.87	3.14	3.91	3.11	3.87	3.14	3.91
3 Paths	3.18	3.21	3.19	3.24	3.18	3.21	3.19	3.24
4 Paths	3.61	3.01	3.62	3.04	3.61	3.01	3.62	3.04
5 Paths	4.20	2.97	4.23	3.00	4.20	2.97	4.23	3.00
10 Paths	4.86	2.96	4.89	2.97	4.88	2.07	4.91	2.97
50 Paths	5.00	2.99	5.02	3.00	5.05	2.25	5.08	3.10
Building V2								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	4.56	4.56	4.66	4.66	4.56	4.56	4.66	4.66
2 Paths	3.11	3.87	3.14	3.91	3.11	3.87	3.14	3.91
3 Paths	3.19	3.24	3.20	3.27	3.19	3.24	3.20	3.27
4 Paths	3.55	3.12	3.56	3.14	3.55	3.12	3.56	3.14
5 Paths	4.19	3.09	4.22	3.10	4.19	3.09	4.22	3.10
10 Paths	4.74	3.04	4.77	3.05	4.86	3.45	4.87	3.43
50 Paths	4.89	3.05	4.92	3.06	5.19	3.61	5.18	3.57
Building V3								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	4.56	4.56	4.66	4.66	4.56	4.56	4.66	4.66
2 Paths	3.11	3.87	3.14	3.91	3.11	3.87	3.14	3.91
3 Paths	3.12	3.10	3.13	3.13	3.12	3.10	3.13	3.13
4 Paths	3.54	2.96	3.56	2.98	3.54	2.96	3.56	2.98
5 Paths	4.25	2.92	4.29	2.94	4.25	2.92	4.29	2.94
10 Paths	4.70	2.95	4.72	2.96	4.62	2.81	4.65	2.82
50 Paths	4.81	2.98	4.82	2.99	4.76	2.92	4.78	2.93

Table 7.3: RMSE for AVW Line 2

AVW Line2 – Mean Error								
Building V1								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	3.14	3.14	3.29	3.29	3.14	2.97	3.29	3.12
2 Paths	0.89	2.38	0.94	2.45	0.89	2.38	0.94	2.45
3 Paths	-0.02	1.25	0.05	1.34	-0.02	1.25	0.05	1.34
4 Paths	-0.89	0.80	-0.83	0.88	-0.89	0.80	-0.83	0.88
5 Paths	-1.75	0.49	-1.71	0.56	-1.75	0.49	-1.71	0.56
10 Paths	-2.70	-0.16	-2.68	-0.09	-2.72	-0.16	-2.69	-0.10
50 Paths	-2.95	-0.54	-2.92	-0.48	-2.89	-0.72	-2.86	-0.65
Building V2								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	3.14	3.14	3.29	3.29	2.92	2.13	3.06	2.30
2 Paths	0.89	2.38	0.94	2.45	0.89	2.38	0.94	2.45
3 Paths	-0.03	1.27	0.05	1.36	-0.03	1.27	0.05	1.36
4 Paths	-1.00	0.86	-0.94	0.94	-1.00	0.86	-0.94	0.94
5 Paths	-1.87	0.58	-1.84	0.64	-1.87	0.58	-1.84	0.64
10 Paths	-2.73	-0.05	-2.70	0.02	-2.99	-0.64	-2.94	-0.54
50 Paths	-3.09	-0.29	-3.07	-0.22	-3.45	-1.25	-3.40	-1.15
Building V3								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	3.15	3.15	3.29	3.29	2.92	2.99	3.06	3.13
2 Paths	0.89	2.38	0.94	2.46	0.89	2.38	0.94	2.46
3 Paths	-0.12	1.15	-0.03	1.25	-0.12	1.15	-0.03	1.25
4 Paths	-1.12	0.74	-1.06	0.82	-1.12	0.74	-1.06	0.82
5 Paths	-2.04	0.43	-2.01	0.49	-2.04	0.43	-2.01	0.49
10 Paths	-2.74	-0.17	-2.72	-0.11	-2.81	-0.27	-2.79	-0.21
50 Paths	-3.01	-0.38	-2.97	-0.32	-3.13	-0.62	-3.09	-0.55

Table 7.4: Mean Error for AVW Line 2

AVW Line3 - RMSE								
Building V1								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	4.51	4.51	4.60	4.60	4.51	4.51	4.60	4.60
2 Paths	3.39	3.88	3.40	3.90	3.39	3.88	3.40	3.90
3 Paths	3.56	3.24	3.56	3.26	3.56	3.24	3.56	3.26
4 Paths	4.03	3.13	4.03	3.14	4.03	3.13	4.03	3.14
5 Paths	4.61	3.13	4.62	3.14	4.61	3.13	4.62	3.14
10 Paths	5.33	3.19	5.35	3.18	5.28	3.18	5.30	3.18
50 Paths	5.47	3.31	5.48	3.31	5.38	3.36	5.38	3.35
Building V2								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	4.51	4.51	4.60	4.60	4.51	4.51	4.60	4.60
2 Paths	3.39	3.88	3.40	3.90	3.39	3.88	3.40	3.90
3 Paths	3.52	3.25	3.52	3.27	3.52	3.25	3.52	3.27
4 Paths	3.99	3.11	4.00	3.12	3.99	3.11	4.00	3.12
5 Paths	4.56	3.09	4.58	3.10	4.56	3.09	4.58	3.10
10 Paths	5.43	3.12	5.45	3.12	5.78	3.79	5.77	3.73
50 Paths	5.87	3.19	5.88	3.18	6.19	4.04	6.17	3.98
Building V3								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	4.51	4.51	4.60	4.60	4.51	4.51	4.60	4.60
2 Paths	3.39	3.88	3.40	3.90	3.39	3.88	3.40	3.90
3 Paths	3.62	3.23	3.62	3.25	3.62	3.23	3.62	3.25
4 Paths	4.19	3.14	4.20	3.14	4.19	3.14	4.20	3.14
5 Paths	4.85	3.15	4.88	3.15	4.85	3.15	4.88	3.15
10 Paths	5.19	3.21	5.21	3.21	5.10	3.15	5.13	3.14
50 Paths	5.42	3.26	5.42	3.26	5.48	3.28	5.50	3.27

Table 7.5: RMSE for AVW Line 3

AVW Line3 – Mean Error								
Building V1								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	2.59	2.59	2.73	2.73	2.59	2.39	2.73	2.53
2 Paths	0.31	1.85	0.36	1.93	0.31	1.85	0.36	1.93
3 Paths	-0.80	0.73	-0.73	0.84	-0.80	0.73	-0.73	0.84
4 Paths	-1.72	0.29	-1.66	0.38	-1.72	0.29	-1.66	0.38
5 Paths	-2.46	0.01	-2.41	0.07	-2.46	0.01	-2.41	0.07
10 Paths	-3.49	-0.58	-3.46	-0.51	-3.41	-0.62	-3.37	-0.56
50 Paths	-3.79	-0.94	-3.75	-0.88	-3.63	-1.18	-3.59	-1.11
Building V2								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	2.59	2.59	2.73	2.73	2.26	1.19	2.40	1.36
2 Paths	0.31	1.85	0.36	1.93	0.31	1.85	0.36	1.93
3 Paths	-0.75	0.76	-0.68	0.86	-0.75	0.76	-0.68	0.86
4 Paths	-1.65	0.34	-1.60	0.42	-1.65	0.34	-1.60	0.42
5 Paths	-2.45	0.06	-2.41	0.12	-2.45	0.06	-2.41	0.12
10 Paths	-3.48	-0.45	-3.45	-0.39	-3.80	-1.30	-3.75	-1.20
50 Paths	-3.81	-0.66	-3.77	-0.60	-4.09	-1.79	-4.03	-1.69
Building V3								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	2.59	2.59	2.73	2.73	2.26	2.28	2.40	2.42
2 Paths	0.31	1.85	0.37	1.93	0.31	1.85	0.37	1.93
3 Paths	-0.90	0.65	-0.82	0.75	-0.90	0.65	-0.82	0.75
4 Paths	-1.88	0.24	-1.82	0.32	-1.88	0.24	-1.82	0.32
5 Paths	-2.65	-0.06	-2.62	0.01	-2.65	-0.06	-2.62	0.01
10 Paths	-3.22	-0.63	-3.19	-0.57	-3.21	-0.83	-3.19	-0.76
50 Paths	-3.48	-0.80	-3.43	-0.74	-3.59	-1.13	-3.56	-1.06

Table 7.6: Mean Error for AVW Line 3

7.1.2 Stadium Drive

Stadium Drive – RMSE								
Building V1								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	2.85	2.85	2.85	2.85	2.85	3.67	2.85	3.65
2 Paths	3.82	3.05	3.84	3.03	3.72	2.97	3.74	2.95
3 Paths	4.00	3.34	4.00	3.30	4.05	3.23	4.07	3.20
4 Paths	4.41	3.47	4.40	3.44	4.39	3.56	4.41	3.53
5 Paths	4.55	3.63	4.57	3.60	4.57	3.75	4.59	3.72
10 Paths	5.06	3.95	5.03	3.92	4.99	4.47	4.96	4.43
50 Paths	5.66	4.09	5.63	4.06	5.70	5.02	5.67	4.98
Building V2								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	2.85	2.85	2.85	2.85	2.85	3.50	2.85	3.48
2 Paths	3.82	3.05	3.84	3.03	3.69	2.97	3.71	2.95
3 Paths	3.96	3.32	3.95	3.29	4.01	3.22	4.03	3.19
4 Paths	4.12	3.47	4.12	3.44	4.14	3.54	4.16	3.51
5 Paths	4.40	3.58	4.40	3.56	4.39	3.85	4.38	3.82
10 Paths	5.00	3.92	4.99	3.90	5.04	4.34	5.02	4.30
50 Paths	5.24	4.09	5.22	4.06	5.36	4.84	5.35	4.80
Building V3								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	2.85	2.85	2.85	2.85	2.85	3.14	2.85	3.12
2 Paths	3.23	3.05	3.27	3.02	2.93	2.99	2.95	2.97
3 Paths	3.07	3.26	3.08	3.23	3.02	3.22	3.04	3.20
4 Paths	3.04	3.38	3.05	3.34	3.07	3.48	3.09	3.45
5 Paths	3.14	3.46	3.16	3.42	3.13	3.70	3.15	3.66
10 Paths	3.42	3.65	3.42	3.61	3.42	3.79	3.42	3.75
50 Paths	3.51	3.69	3.50	3.65	3.65	3.95	3.66	3.91

Table 7.7: RMSE for Stadium Drive

Stadium Drive – Mean Error								
Building V1								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	-0.17	-0.17	-0.03	-0.03	-0.17	-0.63	-0.03	-0.51
2 Paths	-0.81	-1.04	-0.65	-0.93	-0.79	-0.84	-0.67	-0.61
3 Paths	-1.30	-1.43	-1.15	-1.33	-1.26	-1.32	-1.16	-0.98
4 Paths	-1.92	-1.62	-1.79	-1.52	-1.87	-1.79	-1.77	-1.24
5 Paths	-2.16	-1.81	-2.06	-1.72	-2.18	-2.01	-2.07	-1.52
10 Paths	-2.95	-2.17	-2.84	-2.08	-2.85	-2.60	-2.73	-2.12
50 Paths	-3.58	-2.35	-3.47	-2.26	-3.54	-3.11	-3.43	-2.62
Building V2								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	-0.17	-0.17	-0.03	-0.03	-0.17	-0.69	-0.03	-0.55
2 Paths	-0.81	-1.04	-0.65	-0.93	-0.72	-0.89	-0.68	-0.66
3 Paths	-1.27	-1.41	-1.12	-1.30	-1.28	-1.33	-1.18	-1.02
4 Paths	-1.48	-1.60	-1.34	-1.50	-1.35	-1.81	-1.29	-1.42
5 Paths	-1.73	-1.75	-1.60	-1.65	-1.73	-2.00	-1.58	-1.90
10 Paths	-2.27	-2.11	-2.15	-2.02	-2.47	-2.47	-2.35	-2.37
50 Paths	-2.61	-2.28	-2.48	-2.19	-2.61	-2.93	-2.49	-2.83
Building V3								
	Reflections				1 Diffraction			
	Ideal Antenna		Realistic Antenna		Ideal Antenna		Realistic Antenna	
	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase	All Phase	Corr Phase
Top Path	-0.17	-0.17	-0.02	-0.02	-0.17	-0.39	-0.02	-0.24
2 Paths	0.33	-0.70	0.50	-0.70	-0.22	-0.79	0.52	-0.66
3 Paths	-0.21	-1.03	-0.06	-1.03	-0.21	-1.12	-0.07	-0.98
4 Paths	-0.15	-1.21	0.02	-1.21	-0.14	-1.29	0.04	-1.25
5 Paths	-0.22	-1.32	-0.07	-1.32	0.21	-1.54	-0.03	-1.42
10 Paths	-0.62	-1.54	-0.51	-1.54	-0.21	-1.71	-0.47	-1.60
50 Paths	-0.86	-1.62	-0.73	-1.62	-0.57	-1.93	-0.84	-1.82

Table 7.8: Mean Error for Stadium Drive

7.2 Transmitter & Receiver Modeling

7.2.1 ETOA & Number of Paths

Before learning that our receiver hardware has a "best path sequencer", a variation in predictive accuracy was seen between LOS and NLOS situations. When there is LOS between a transmitter and receiver, the LOS ray tends to dominate the others being added due to the power differential, resulting in accurate predictions no matter how many paths were being added. In NLOS situations, it was noticed that due to a lack of a 'dominant path', prediction accuracy was related to the number of paths being added. We also observed that adding a large number of paths resulted in over-prediction and that by reducing the number of paths, the predictions would drop to a level more accurately matching observed values before under-predicting. Although it's difficult to see in the mean error results of the Stadium Drive field test due to the accuracy of the top path, the mean error results of the AVW comparison clearly demonstrate this over/under-prediction pattern.

Initially, output filters were applied to restrict the ETOA in attempt to limit the number of paths. Although this filter produced fairly accurate results with extreme restrictions, it was less effective than simply limiting the number of paths which more accurately models the discovered "best path sequencer". The ETOA filter of one symbol time was kept throughout the analysis as we believe it is a reasonable assumption.

When limiting the number of paths, the RMSE values for both field tests clearly indicate that the highest accuracy is obtained when limiting the number of

paths to five or less. In AVW, obtaining the highest accuracy depends on how phase information is dealt with at the receiver, but in both cases the top path is not optimal. If summing with all phase information, the best RMSE is found when combining two or three paths giving a RMSE between 2.18 dBm and 3.40 dBm. Little variation is seen by varying building detail or other parameters. Selecting the number of paths on either side of this range results in the RMSE nearly doubling. When correlating phase, the best RMSE is found when selecting five or ten paths giving a RMSE between 2.02 dBm and 3.14 dBm. It turns out that correlating phase using more than two paths is quite forgiving and results in a RMSE within a few tenths of a dBm. Again little variation is seen by varying building detail or other parameters. The mean error seems to be in good agreement with the RMSE in that it centers around the best selection for RMSE.

In Stadium Drive, obtaining the best RMSE can be achieved by only using the top path, except for when correlating phase and including a diffraction (but still very close). The RMSE gets increasingly worse as the number of paths increases, no matter how phase is combined or any other parameter is varied. Although the RMSE increases with the number of paths, the spread of the RMSE decreases as the building detail is simplified. In V1, V2, and V3, the RMSE when summing all phase ranges from approximately 2.85 dBm to 5.7 dBm, 2.85 dBm to 5.3 dBm, and 2.85 dBm to 3.6 dBm. This same pattern occurs when correlating phase. In terms of the mean error, V1 and V2 are best around the top path and increasingly under-predict as the number of paths combined increases. This is true when varying all other parameters. When using V3, using all phase information gives a mean error between dBm and -1

dBm, while correlating phase gives a mean error between 0 dBm and -2 dBm. Both phase options increasingly under-predicts as the number of paths combined is increased.

7.2.2 Summing Complex Electric Fields

Since the method of summing complex electric fields was tied into the discussion in the previous section, we will only give a general outline in this section. When comparing the results from either combining all phase information or only within correlated groups, we see patterns unique to each field test. The RMSE results from AVW show that we can include more paths in the received power calculation as long as the correlated phase option is used. If using all phase, combining only the top two or three paths yields the best RMSE, getting increasingly worse as more paths are added. Both of these patterns are true no matter what level of detail or other parameters were used. The mean error seems to be in good agreement with the RMSE in that it centers around the best selection for RMSE.

The Stadium Drive field test showed that correlating phase was better overall for V1 and V2 having a smaller spread in RMSE, but using all phase became just as accurate, if not more as the building details were simplified in V3. But in almost all cases, the RMSE gets worse as we include more than just the top path. As described in the previous section, the mean error for this field test is better when combining all phase and always increasingly under-predicts as the number of paths increases.

7.2.3 Ideal vs. Realistic Antenna Patterns

When looking at the differences in RMSE and mean error in the cases for both field tests, only slight differences can be seen when comparing the use of ideal and realistic antenna patterns. Using a realistic antenna pattern will result in a limited selection of paths arriving within the main lobe compared to the use of the ideal antenna pattern. The minute differences seen across the board suggest that the restrictions placed on direction of departure and arrival result in a nearly similar selection of paths, and that a majority of paths are in fact arriving within the main lobe region. After analyzing the arriving paths in WI, this was seen to be the case.

Given that the ideal antenna pattern performed nearly as well as the obtained realistic pattern, in some cases better, it's hard to justify the use of a realistic pattern. We're fortunate enough to have access to an anechoic chamber at no cost, but this is not the case for most researchers. The cost of obtaining access to an anechoic chamber can be quite prohibitive. Although we were able to show that the ideal pattern performed just as well as the realistic pattern, this performance is limited to the environment we have created.

7.3 Building Models

7.3.1 Allowed Interactions

Along with the amount of detail in the environment, the allowed interactions can have a significant impact on the time spent doing computations. In AVW, little difference, if any, can be seen in RMSE and mean error when using only reflections

or including a single diffraction. This is because diffractions are not present in the top five paths and are therefore not used in the calculation of received power. Due to the close proximity of the two buildings, a group of reflections are seen to have greater power at the receiver than the simple diffractions around the corner of the annex building. In Stadium Drive, diffractions represent the top few paths and are therefore included when calculating the received power. Even though this scenario includes more diffractions in the top five paths, we still see little difference when using only reflections. In fact, using reflections produces results, both RMSE and mean error, which are approximately the same, if not better than the inclusion of a single diffraction. This can be seen when varying any of the other parameters and building detail. Due to the time required to find and shoot and bounce rays from diffraction points, we cannot yet justify the use of diffractions based on our current results.

Based on observation and analysis of the paths in both AVW and Stadium Drive, it is believed that both the close proximity of buildings and the dominance of brick on the UMD campus contribute to the performance of using only reflections. All areas of both field tests can be illuminated using only a few reflections. Other research [36] has indicated that the inclusion of a single diffraction yields dramatic performance improvements over using only reflections, but was done over much longer distances. Future field tests will continue to explore the use of diffractions as they expand to include larger portions of the UMD model.

7.3.2 Building Detail

The amount of detail in the modeled environment can significantly impact the amount of time spent doing computations. In AVW, we see virtually no difference in

RMSE or mean error across all three versions of building models. In the Stadium Drive field test, we see the RMSE spread as we vary the number of paths decrease as buildings details are simplified. This is especially true for summing complex electric fields using all phase information. Moreover, the mean error shows less under-prediction as details are simplified, bringing the mean error much closer to zero. It turns out that simplifying the environment reduced the variance of RMSE and mean error since the simplified environment resulted in less paths between the transmitter and receiver. The top 50 paths turned into the top 20-30 paths depending on LOS or NLOS situations. As mentioned earlier, limiting the number of paths usually produces predictions that better match our field test measurements. Less variance was not seen in the AVW scenario because there were many interactions with complex features that were less dominant, resulting in few differences.

7.4 Recommendations

In this section, recommendations are made for the tested WI parameters based on the tables presented in section 7.1 and the previous discussions. When it comes to selecting an antenna pattern for an omnidirectional antenna, the built-in generic half-wave dipole produced similar results to a realistic pattern. Given the field test comparisons in this thesis, the cost of using an anechoic chamber cannot be justified. However, we were able to find that our antenna gain was approximately 3 dBi instead of the 5 dBi claimed by the manufacturer. Testing in an anechoic chamber can be avoided as long as other tests can be done to determine the actual gain of the antennas being used.

Since we saw few advantages to using just reflections over a single diffraction, it is recommended to use reflections in near earth, close proximity scenarios where the buildings are dominated by few materials. Not only will this lead to less time running computations, but will also allow for further accuracy through extended tuning. As noted in [36], one of the limitations of a model based on purely reflections is that it's near impossible to tune material properties due to the complex nature of urban environments. Luckily, the UMD campus model used in this thesis is predominantly brick and field test measurements can be used to find the properties for brick that provide the best accuracy. Future work involves finding the 'magic brick' for our environment model.

Since we don't know details of environment as well as the position during RSS measurements with sub-wavelength accuracy, it may not be accurate or even reasonable to combine paths at the receiver using all phase information. Some form of aggregation is needed for both transmitter/receiver position (discussed in future work section) and phase information (correlated). As mentioned in previous discussions, correlating phase can be much more forgiving in terms of accuracy when selecting more than two paths and combined with other parameter selections. Looking at the results for both AVW and Stadium Drive, correlating phase while using the V3 building models has a higher probability for accurate predictions due to little variation in both RMSE and mean error. Using these parameter selections and by looking at the field test results, the RMSE had a range of 2.04 dBm to 3.7 dBm and the mean error had a range of -1.62 dBm to 1.53 dBm.

7.5 Running WI Predictions on BounceHaus

After obtaining the predicted path loss from WI and performing the desired method of aggregation, the resulting path loss values (minus insertion loss) can be combined with timestamp information in an XML-encoded scenario file and then run on the BounceHaus testbed. Translating the accuracy seen in the WI predictions to accuracy on the testbed depends on the frequency of attenuation updates between node pairs. More frequent updates require more predictions from WI which can increase the total amount of time running computations. The time spent predicting path loss values will depend on environment detail, whether or not diffractions are being used, the number of nodes modeled, type of aggregation, ray spacing, and collection radius. Given that we've been able to achieve a high level of accuracy with both a simplified environment and by only using reflections, providing updates every second can be done in a reasonable amount of time. For more frequent updates, users will have to tend with longer computation times in the Precomputation phase of the testbed.

Figure 7.1 shows the original trace of AVW line1 after aggregation compared to running the predictions from WI on the BounceHaus testbed. The selected parameters used an ideal antenna pattern and correlating phase within the V3 building models only using the top three paths. This selection in WI gave a RMSE of 2.26 dBm when comparing the measurements at the predefined points. Using the predicted values from WI, adjusted for cable loss and antenna gain, the RF switch was updated approximately every two seconds and attenuator settings were rounded to the nearest integer since the attenuators have a resolution of 1 dB. Most notably, the most error

occurs around the transition from NLOS to LOS. This sharp drop was predicted in WI no matter what parameters were selected. It is believed that a more gradual transition exists in the field test due to environmental objects not modeled within WI.

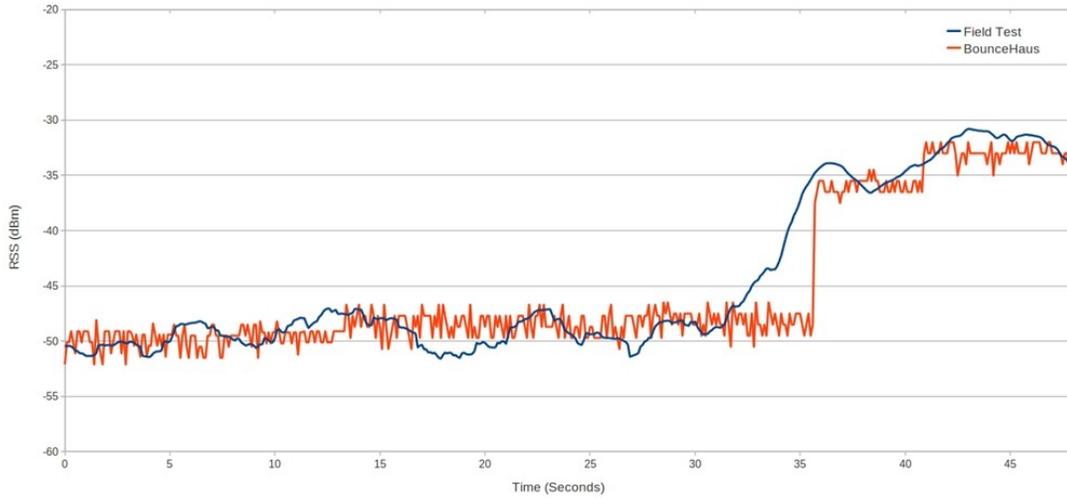


Figure 7.1: Running AVW Line1 Predictions on BounceHaus

Chapter 8: Conclusion & Future Work

8.1 Conclusion

In this thesis, we presented a prototype of the BounceHaus testbed offering researchers the desirable experimental qualities of control, manageability, and repeatability, without sacrificing realism. BounceHaus leverages ray tracing techniques to implement highly detailed, site-specific channel models to dynamically configure a many-to-many analog channel emulator. To validate the use of ray tracing techniques in a mobile wireless testbed, we first modified commodity hardware to make an accurate and reliable RSS measurement tool. These devices were then used to conduct field tests on the University of Maryland campus, collecting RSS measurements to be compared to predictions from Wireless InSite. Before making any comparisons, we created 3D models of several buildings on the University of Maryland campus and built a 3D environment within Wireless InSite. Using the field test measurements and the created 3D environment, we then studied various parameters within the ray tracing software to evaluate its performance as a path loss prediction tool.

After extensive testing of the various parameters involved in both hardware and environment modeling, it was determined that a high level of accuracy could be achieved using generalized parameters and avoiding extensive tuning. Also, it was found that simplified building models provided nearly as much accuracy as high detailed models due to the characteristics of the modeled environment. However, the BounceHaus testbed is not a complete replacement of simulation. Simulation serves

better for large-scale experiments and for initial development of applications or protocols. Due to the amount of work involved in creating a 3D environment, the BounceHaus testbed may be best for developers wanting to move past theoretical testing and evaluate performance in a realistic wireless environment.

8.2 Future Work

The work in this thesis provides significant contributions to the studying the viability of ray tracing technique leverage within mobile wireless testbeds. However, this work only provides initial validations and leaves much work to be done before the true potential of the BounceHaus testbed can be unlocked.

Future work in terrain modeling includes studying the effects of foliage, mixed material terrains, and adding elevation data. The building models were created to handle the addition of elevation data and this data is readily available to us from contacts at the University of Maryland. One hope is that using elevation data will boost the performance realistic antenna patterns over the built-in generalized patterns. Also, we were able to achieve a high level of accuracy with the V3 building models, but a few anomalies exist when analyzing the propagation paths between transmitter and receiver in some of our tests. Future work looks as using V3 but adding certain environmental objects such as metal containers, chain link fence, and any other objects external to building structures that may act as a reflector or obscure measurements. Extensive studies will also be done to find the material properties for brick that best match our field test measurements, finding the 'magic brick'.

Much work has yet to be done to correctly model transmitter/receiver hardware within Wireless InSite. In this thesis, we made loose assumptions and were attempting to model hardware that we didn't know the details of. Thorough testing needs to be done to better understand hardware behaviors so that efforts to model hardware aren't as blind.

Since we were able to achieve high levels of accuracy without extensive tuning with the full3d propagation model, other propagation models within the Wireless InSite suite should be examined to see if they can achieve similar levels of accuracy. Using other propagation models by reduce the overall amount of work required to use full3d.

The field tests used for initial validation are fairly basic LOS and NLOS situations and only involve two nodes. Future work includes conducting more extensive field tests to look at more complex situations and adding multiple nodes using GPS and cars. Adding more nodes will change the way we aggregate power measurements within Wireless InSite. Possible aggregation methods include a 2D square of receivers spanning several wavelengths at the fixed receiver height, or extending this idea into 3D by varying the receiver height over multiple wavelengths. Note that with multiple nodes each node will have to be modeled as both a transmitter and a receiver to represent all of the possible links.

There are several improvements to the BounceHaus architecture that will soon be implemented. After the initial creation of the building models used in this thesis, Remcom released a new version of Wireless InSite which took advantage of GPU hardware for increased performance. However, we found several issues when trying

to run the new X3D propagation model using our building models. We are currently working with Remcom to try and resolve these issues. Using GPUs will open up new possibilities as the time spent doing computations will become even less of a factor.

Another improvement to the testbed architecture involves the physical medium. One issue is that using one switch limits scalability to eight nodes. Possible solutions involve both implementing a virtual environment like VMT and replacing the RF switch with a 16 port RF switch. Another issue lies in the fact that by using an RF switch, we are limited to controlling only attenuation between node pairs. We are currently working with Carnegie Mellon University and have future plans to replace the RF switch with the CMU wireless emulator. Their emulator provides an attractive option as it allows us to control signal strength as well as multipath and propagation delays. Wireless InSite provides all of the necessary output to implement this high level of control.

Appendix A: Building Models

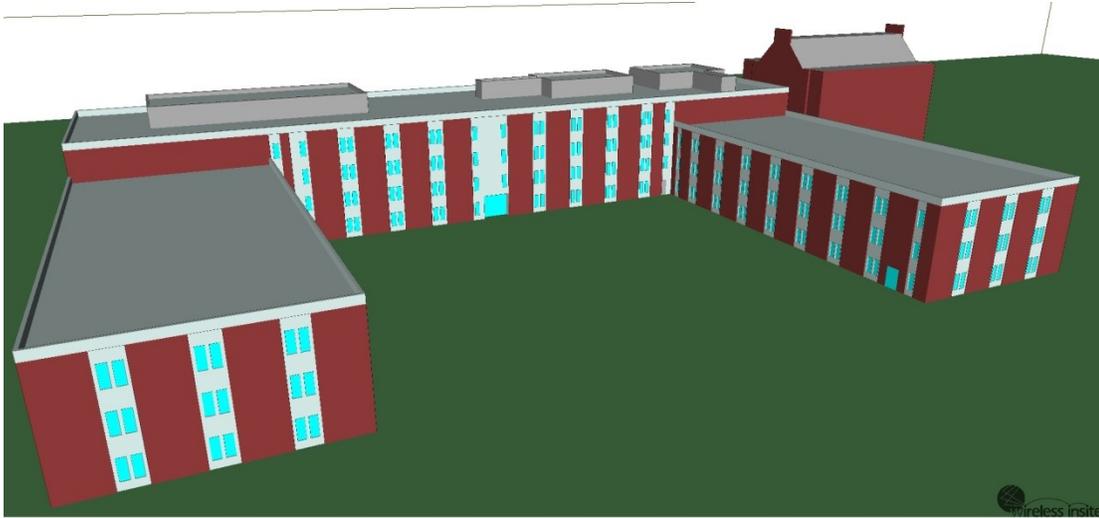
In this appendix, you will find images of all buildings modeled for the work in this thesis along with their simplified versions. Unless otherwise stated, the images use the material color scheme shown in the following table.

Material	Color
Brick	Red
Concrete	Gray
Glass	Light Blue
Metal	Dark Brown
Wet Earth	Green
Wood	Orange

Before presenting the building models, either ground or aerial photos of the actual buildings are shown.

A.V. Williams / Computer Science Instructional Center (115-406)

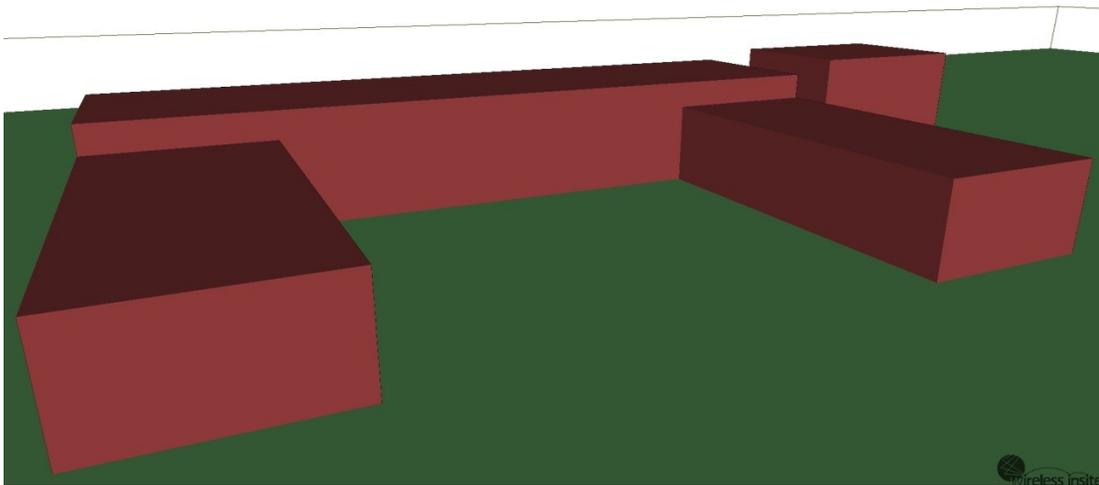




115-406 Version 1

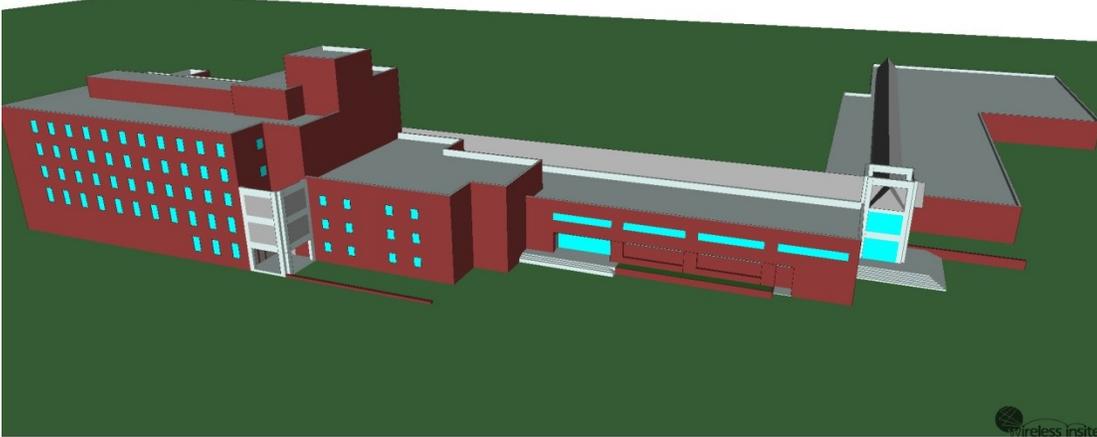


115-406 Version 2

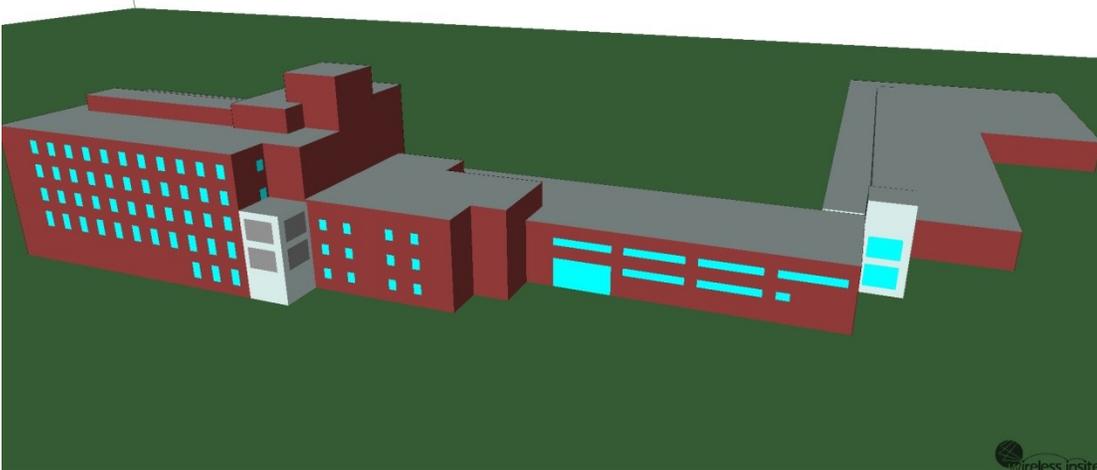


115-406 Version 3

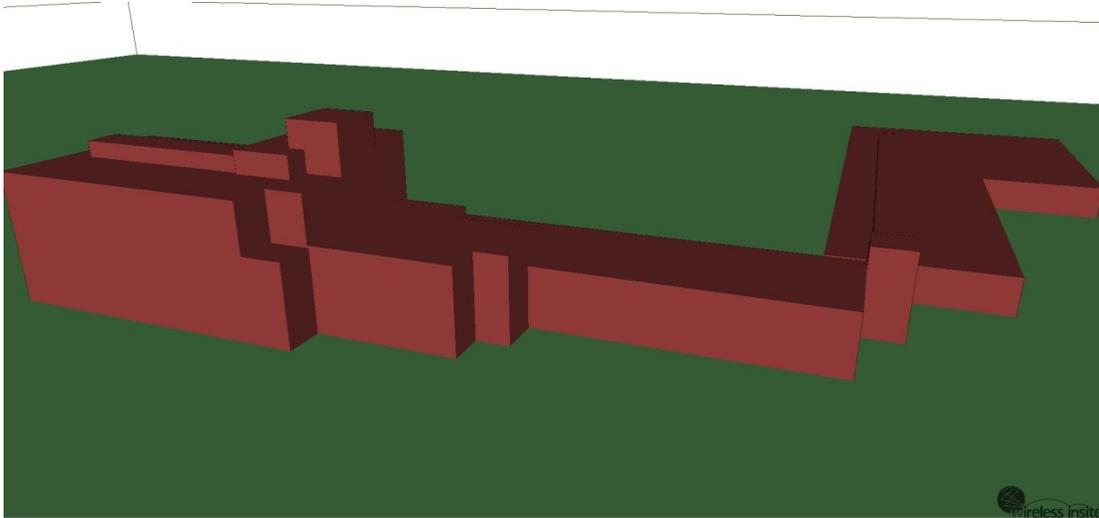
Animal Science/Agricultural Engineering (142)



142 Version 1



142 Version 2



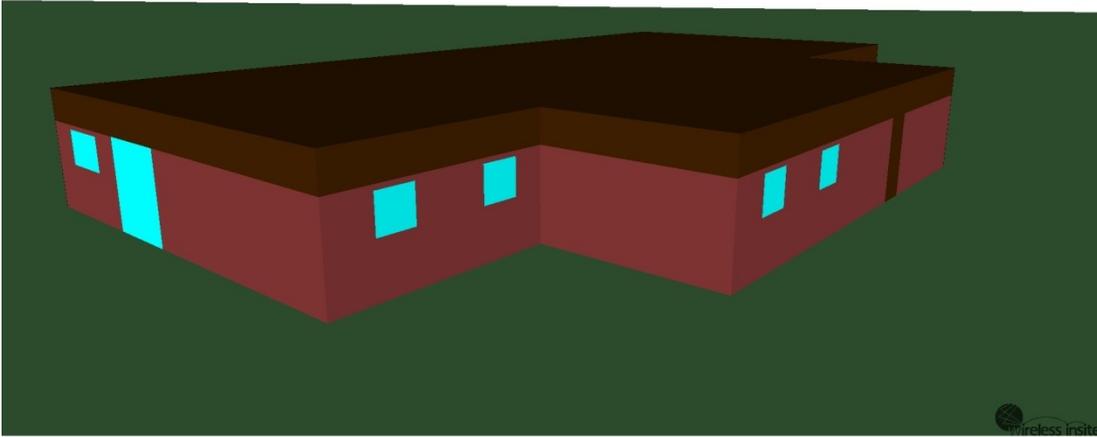
142 Version 3

Central Animal Resources Facility (087)

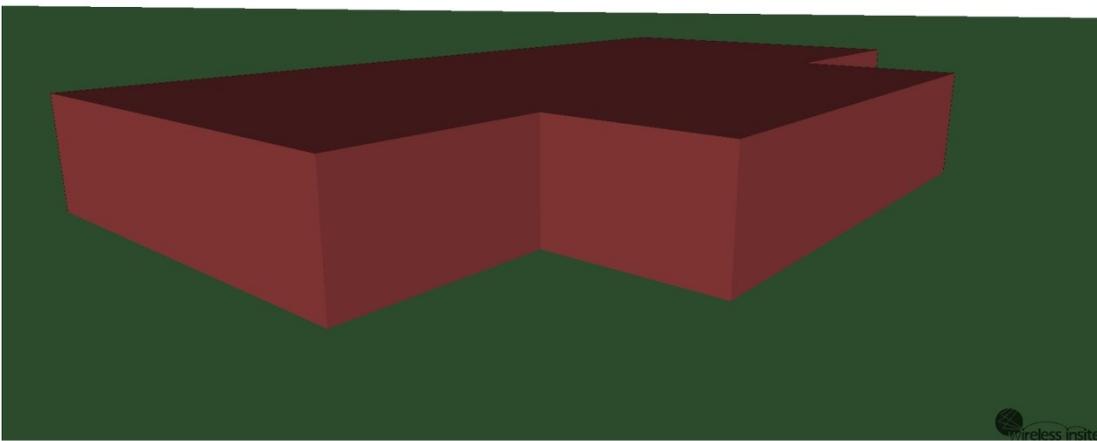




087 Version 1

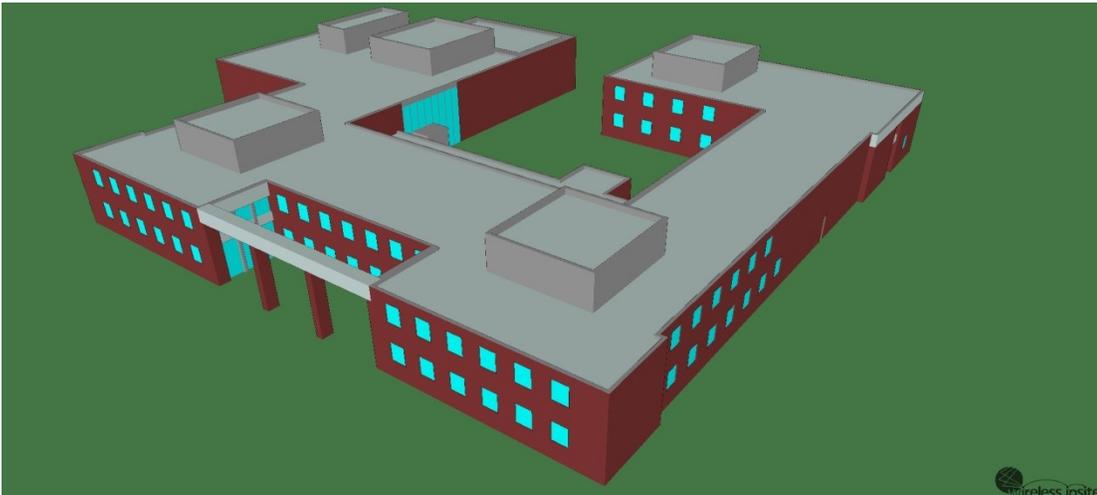


087 Version 2

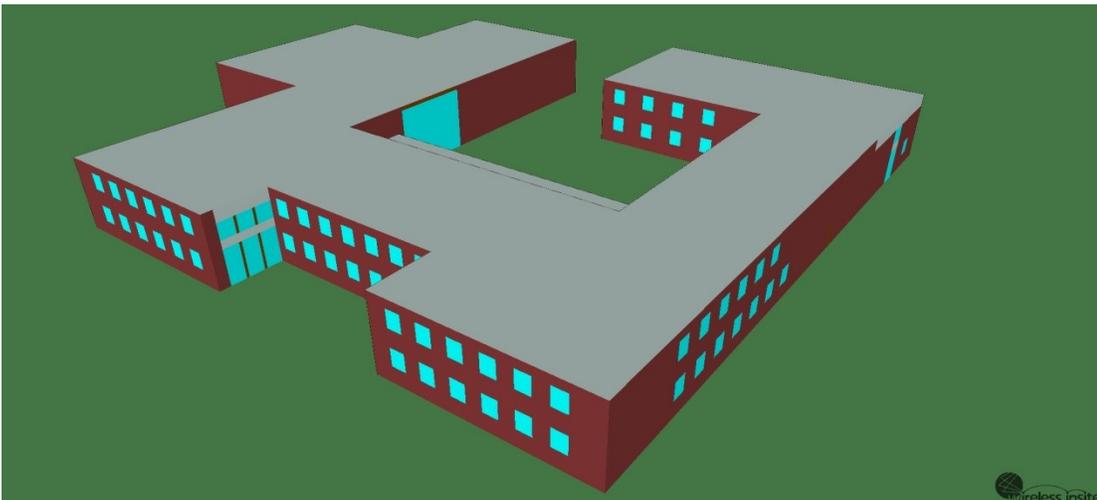


087 Version 3

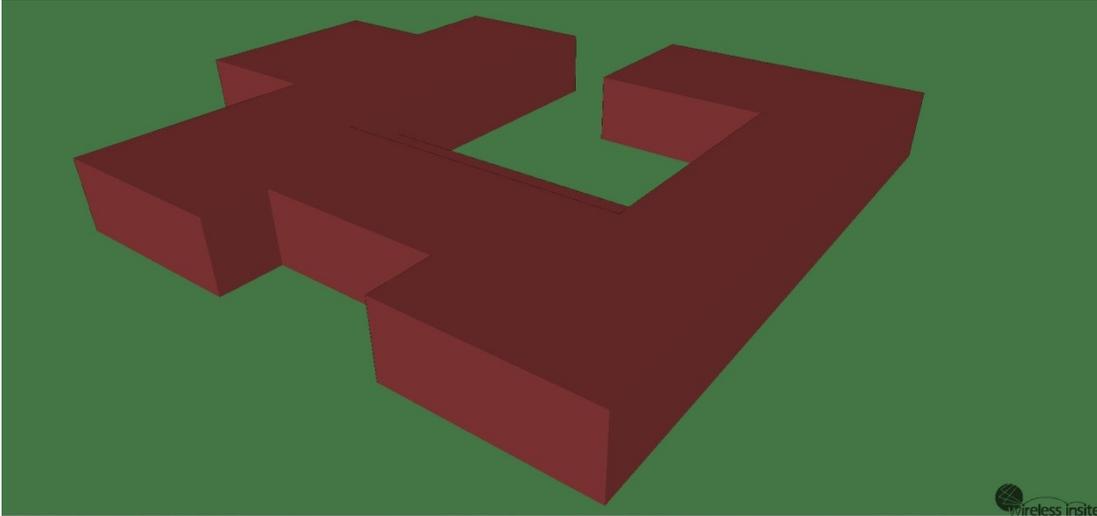
Chemical and Nuclear Engineering (090)



090 Version 1



090 Version 2



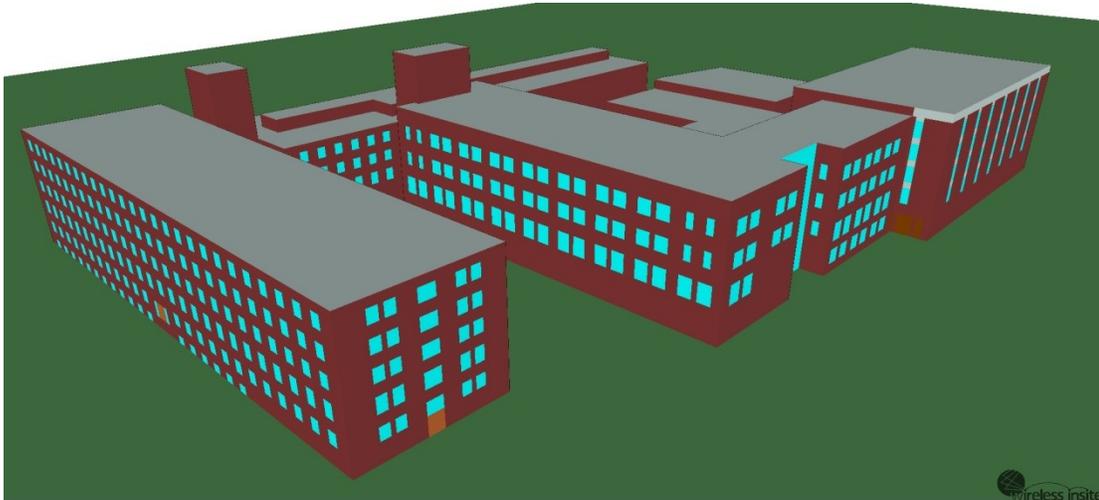
090 Version 3

Chemistry (091)

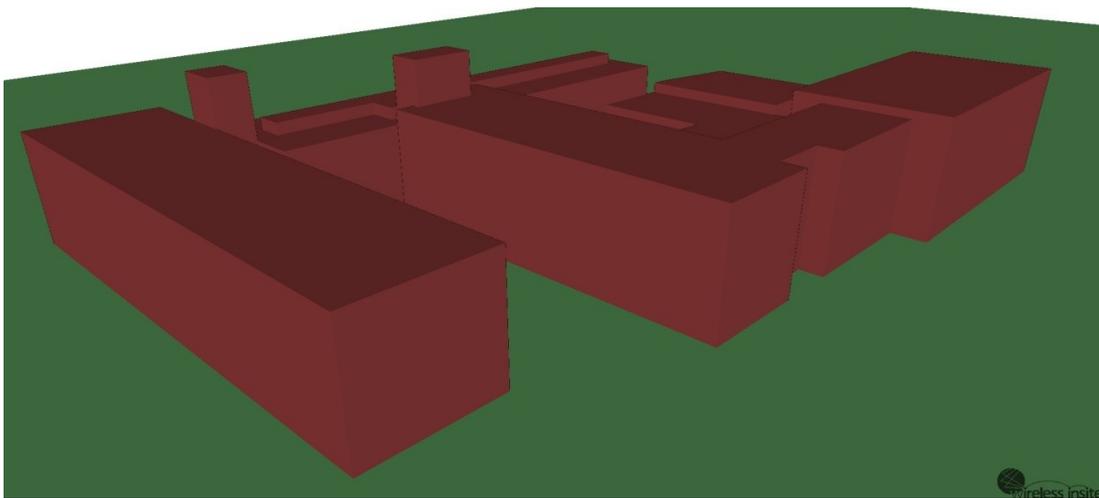




091 Version 1

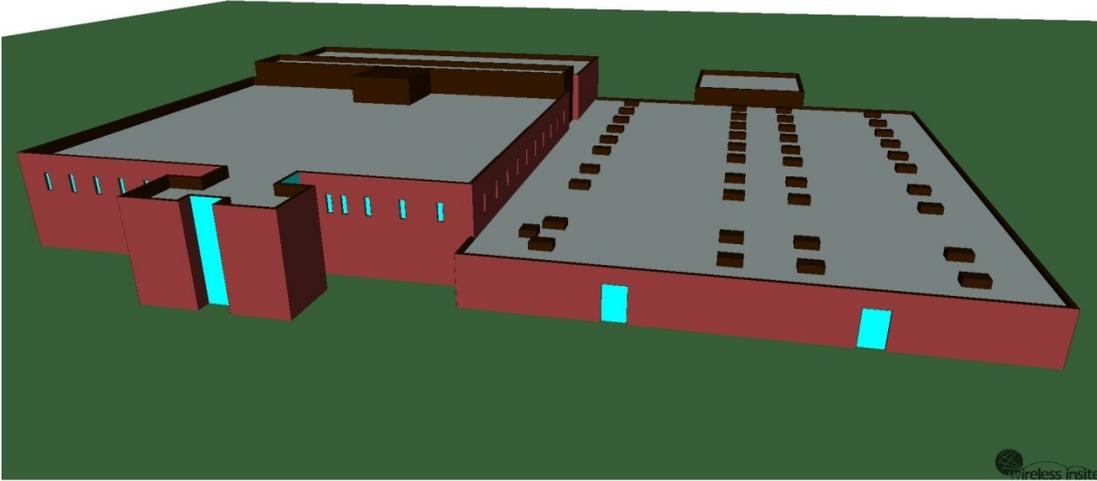


091 Version 2

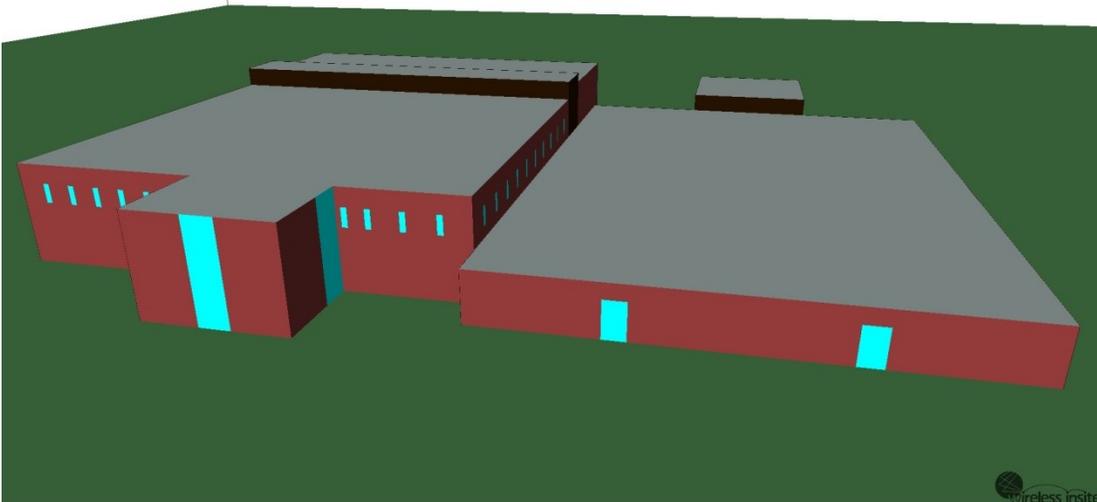


091 Version 3

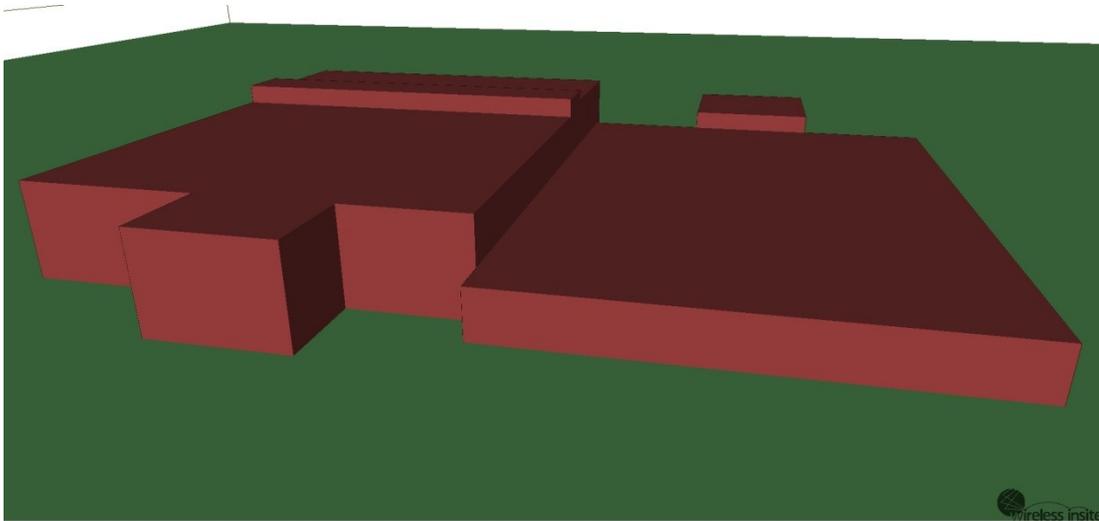
Energy Research Facility (223)



223 Version 1

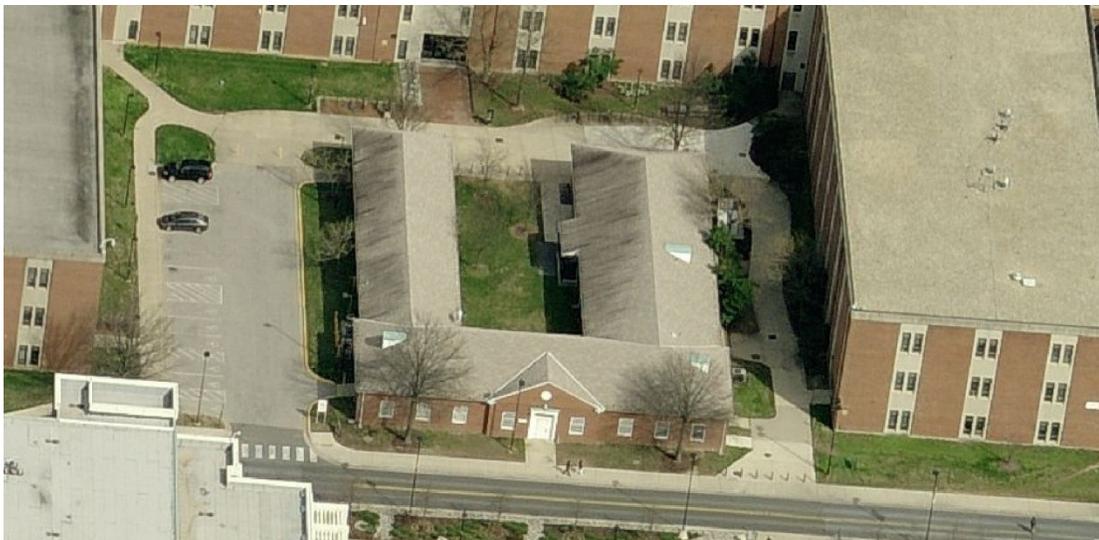


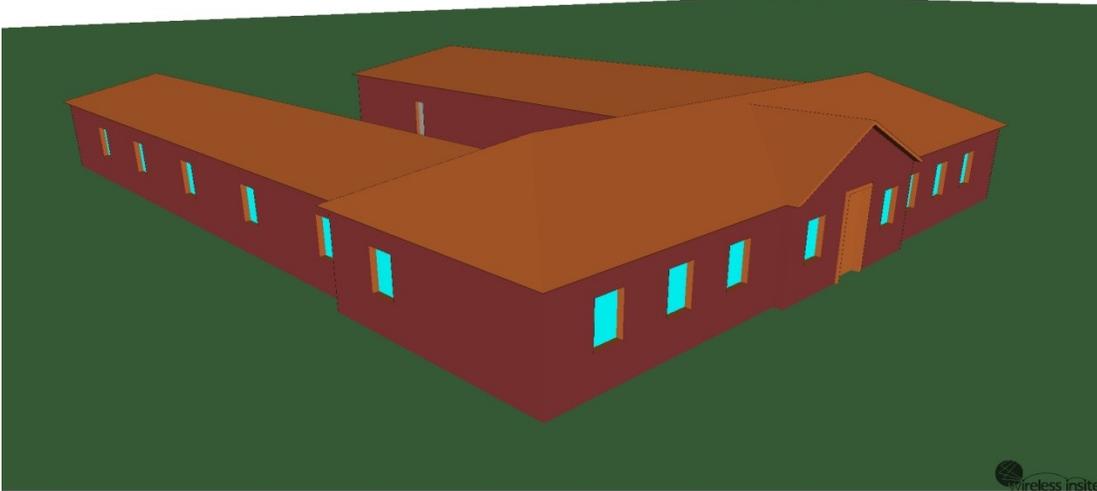
223 Version 2



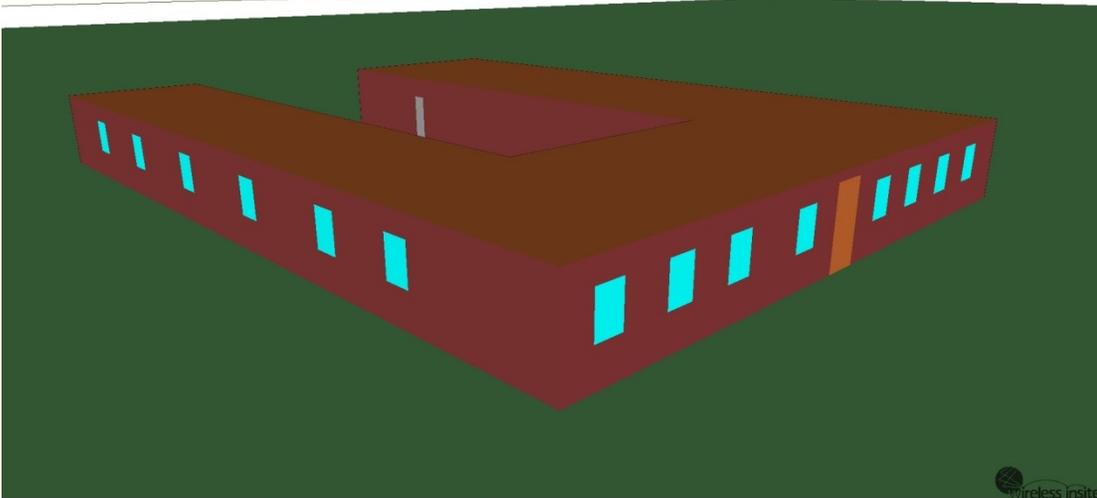
223 Version 3

Engineering Annex (093)

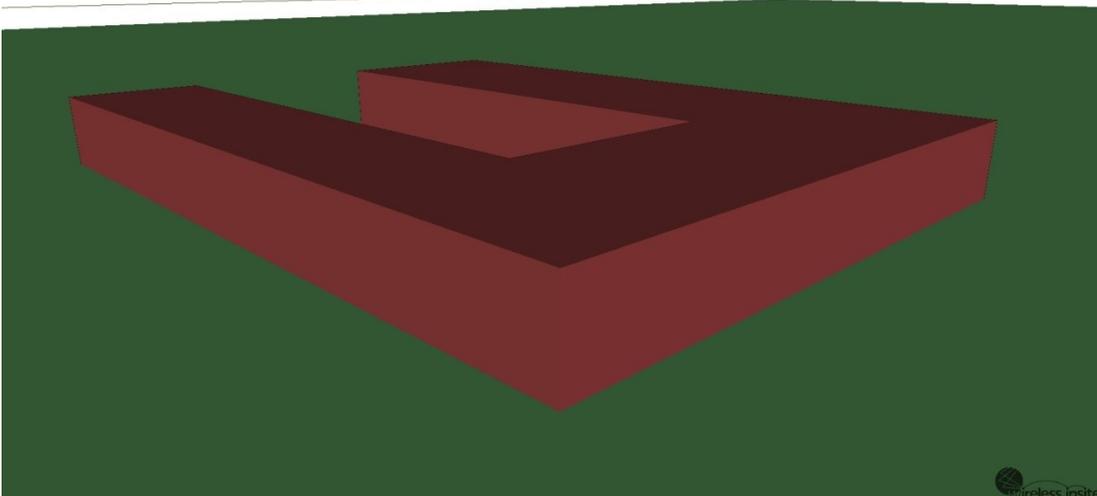




093 Version 1

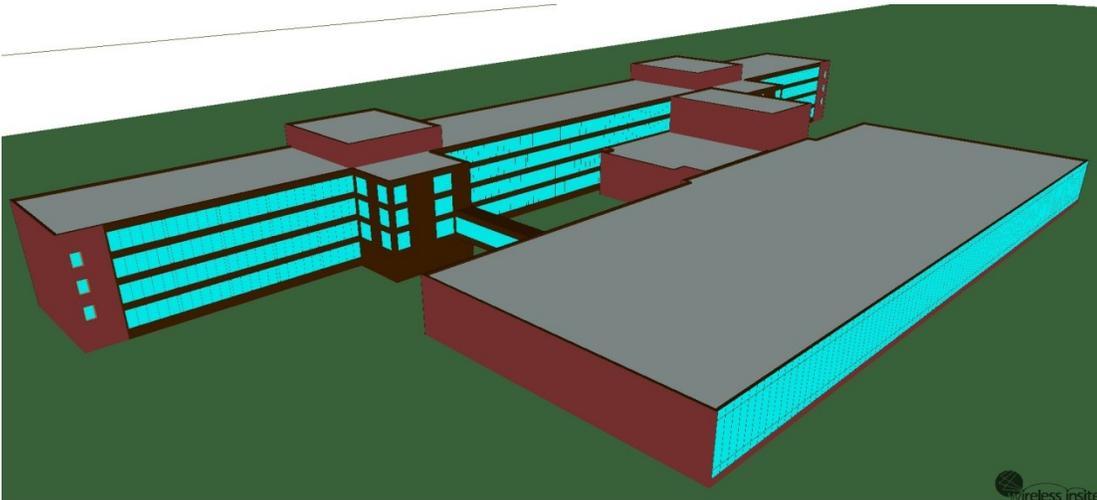


093 Version 2

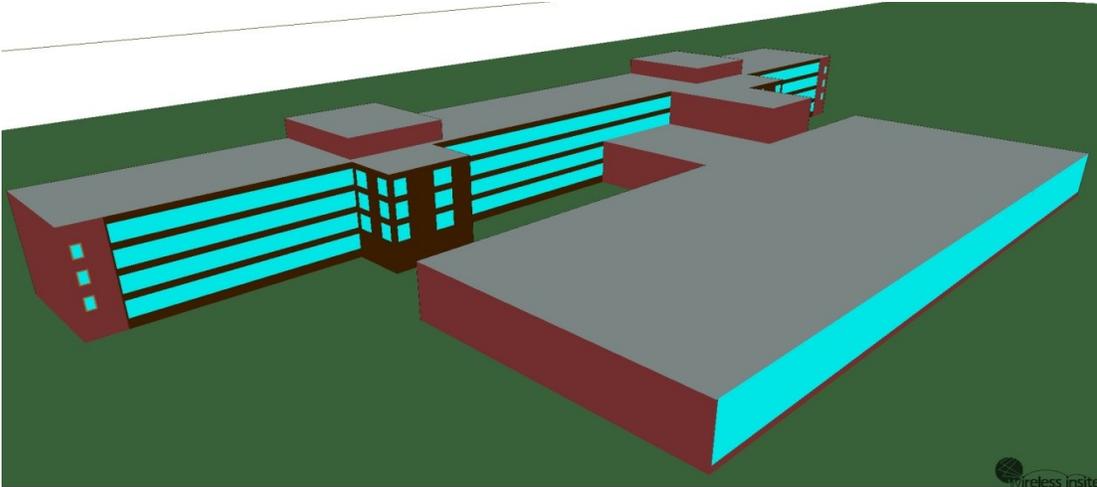


093 Version 3

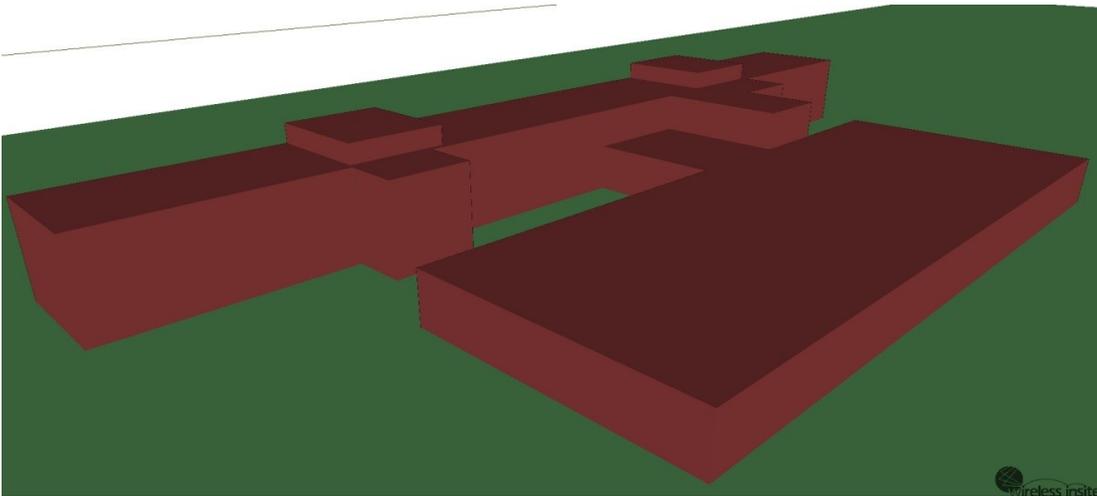
Glenn L. Martin / Engineering Laboratory (088-089)



088-089 Version 1



088-089 Version 2



088-089 Version 3

Institute for Physical Science and Tech. (085)

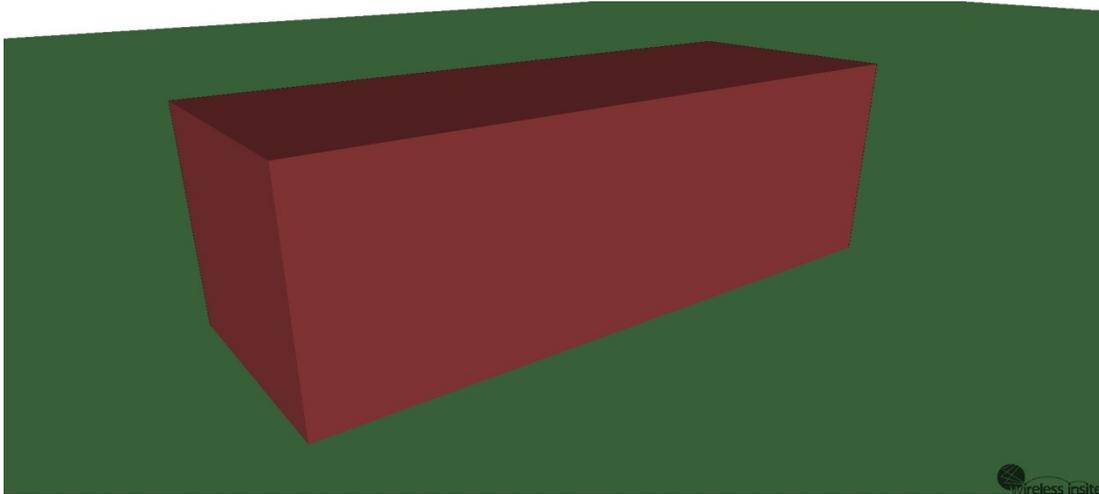




085 Version 1

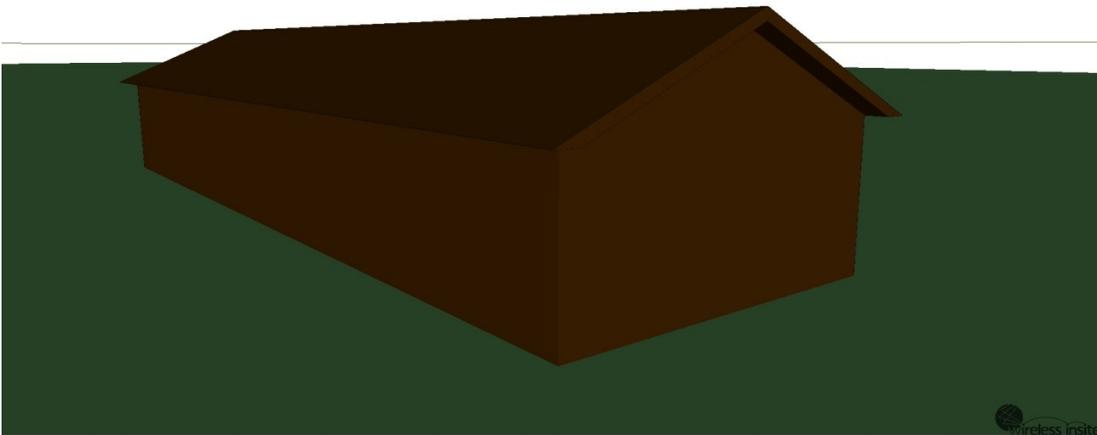


085 Version 2

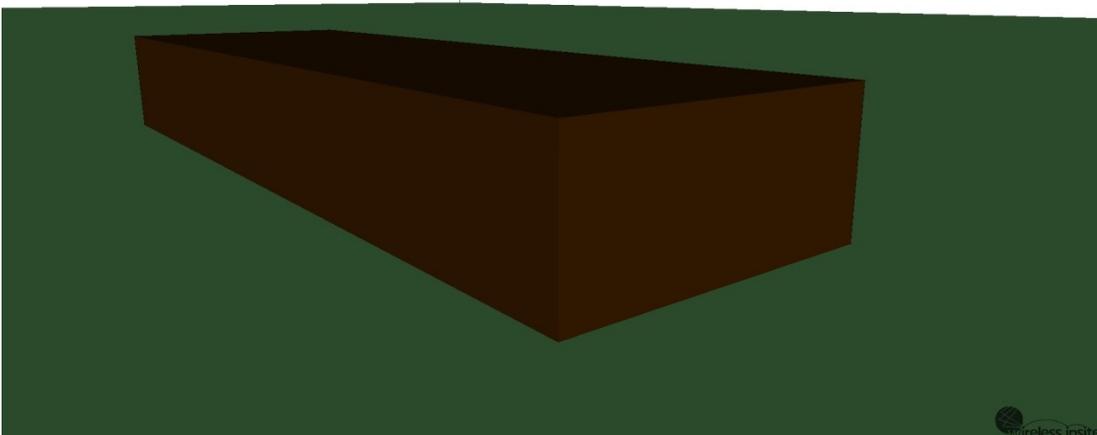


085 Version 3

Instructional Television Facility (045)

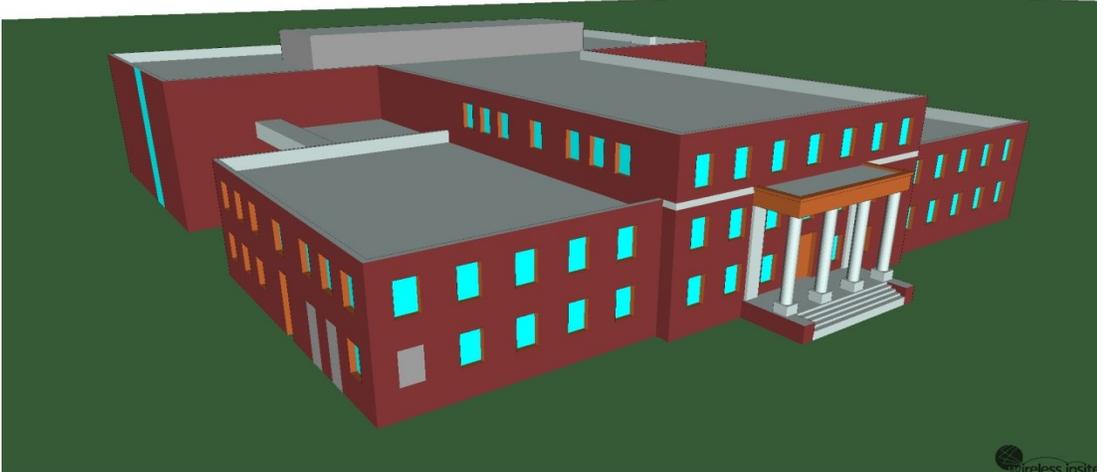


045 Version 1

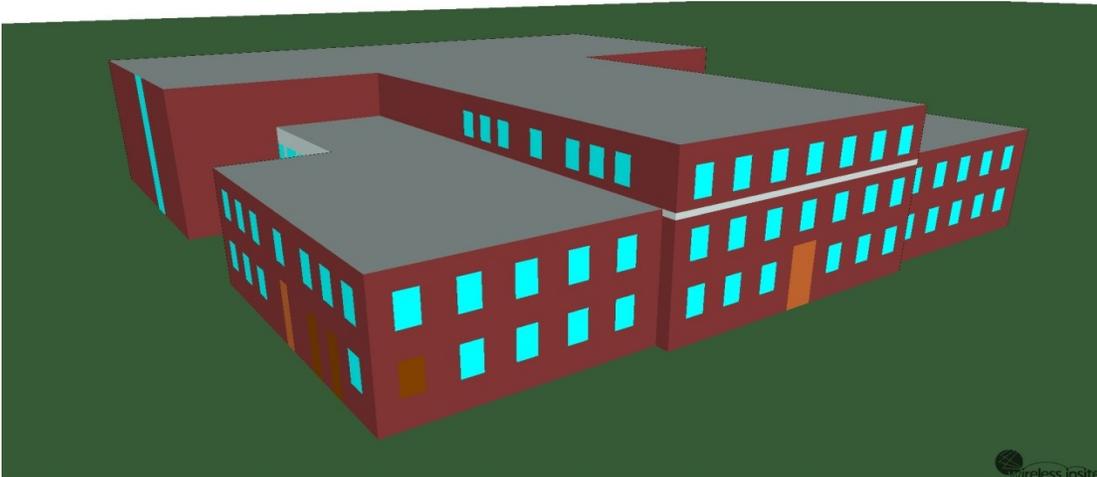


045 Versions 2&3

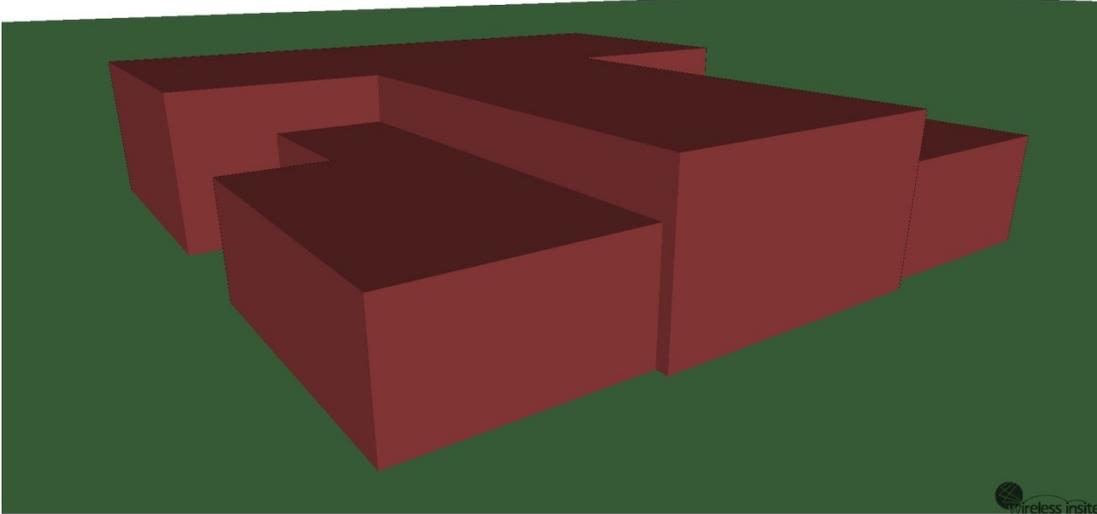
J.M. Patterson (083)



083 Version 1



083 Version 2



083 Version 3

Jeong H. Kim Engineering 225

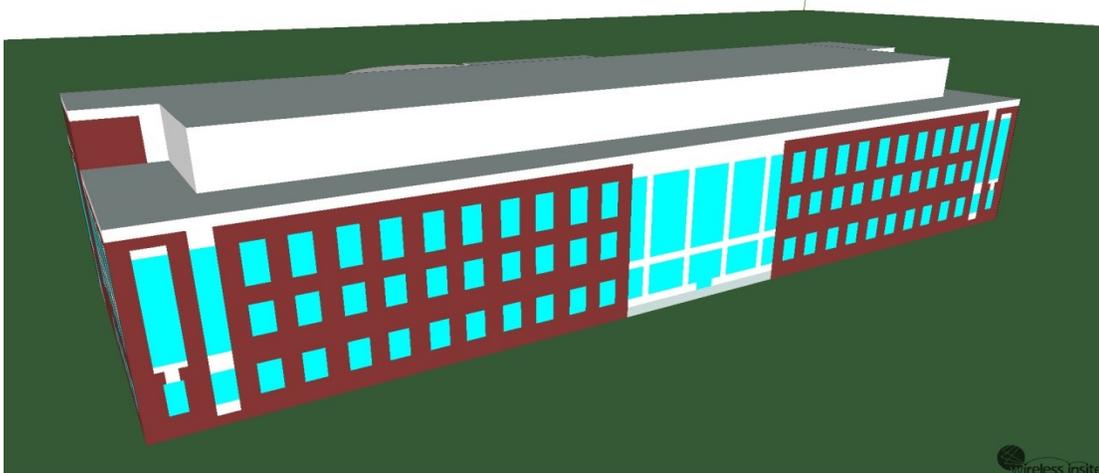
For the building models of building 225, metal surfaces are shown in white. This is simply done for visual purposes.



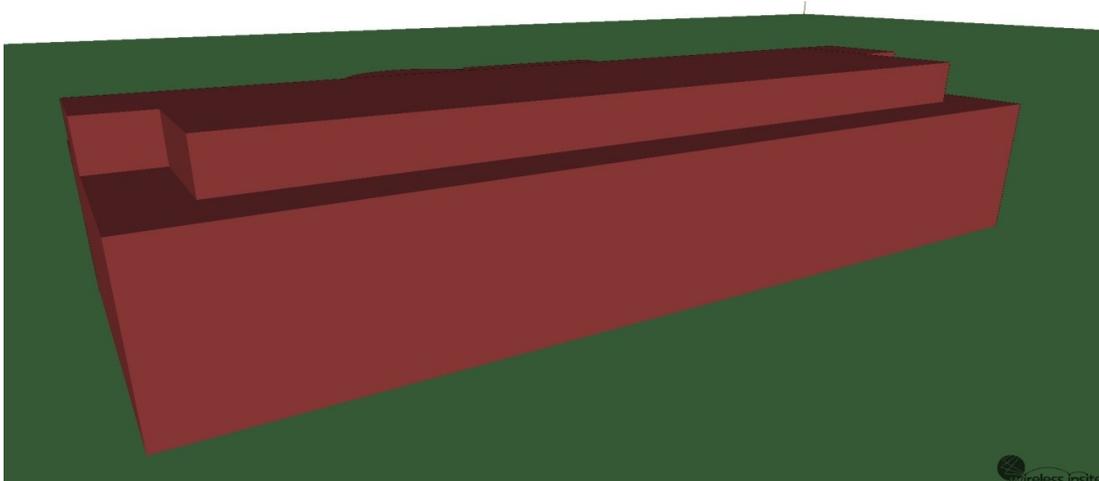
225 Northern Aerial View



225 Version 1 North Side



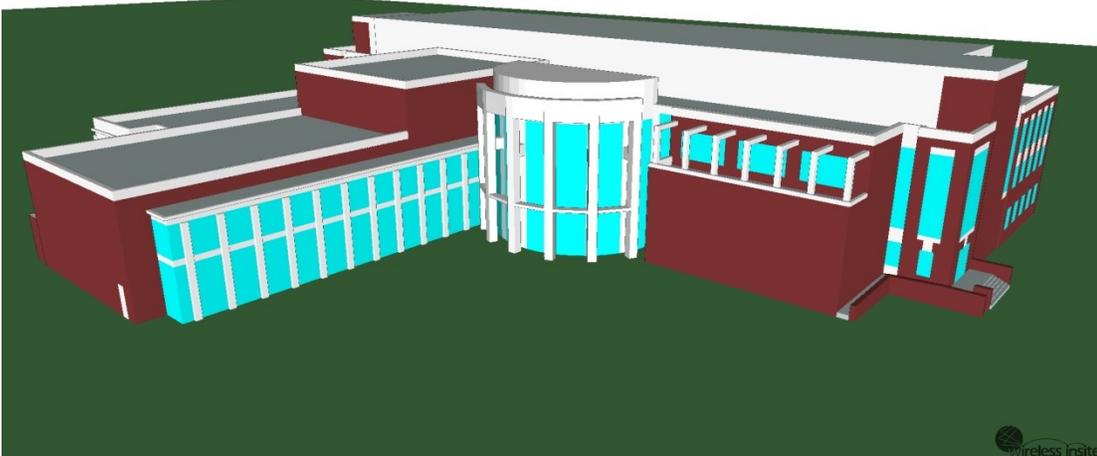
225 Version 2 North Side



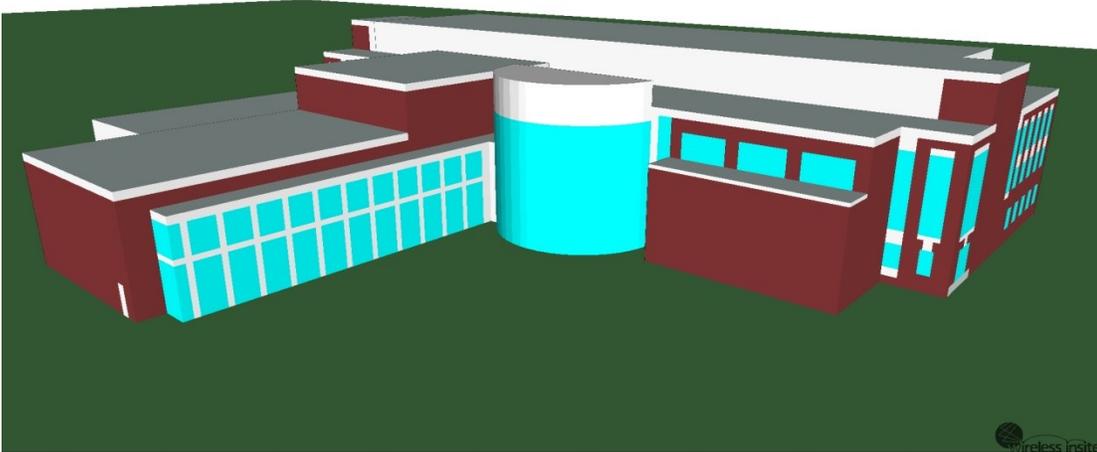
225 Version 3 North Side



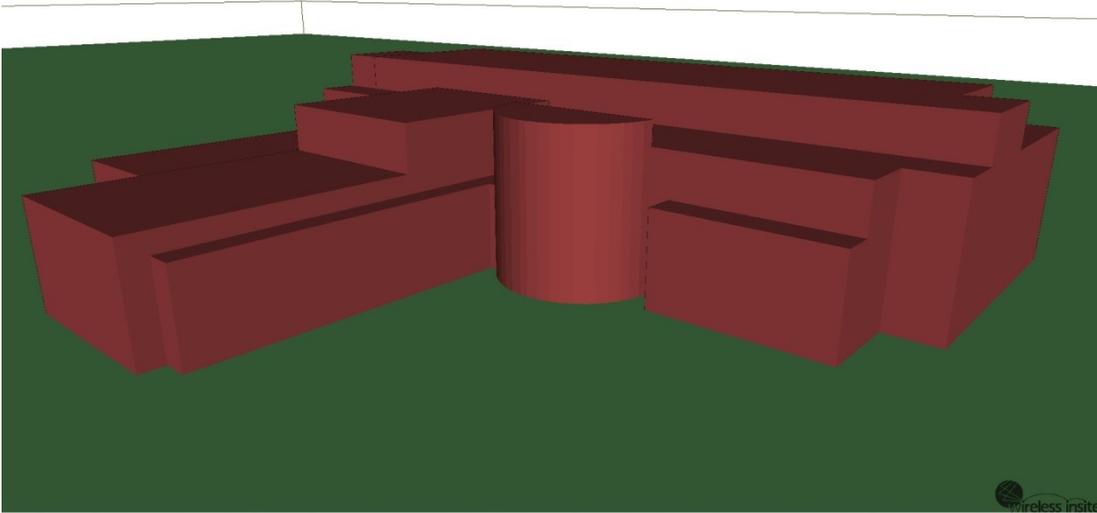
225 Southern Aerial View



225 Version 1 South Side



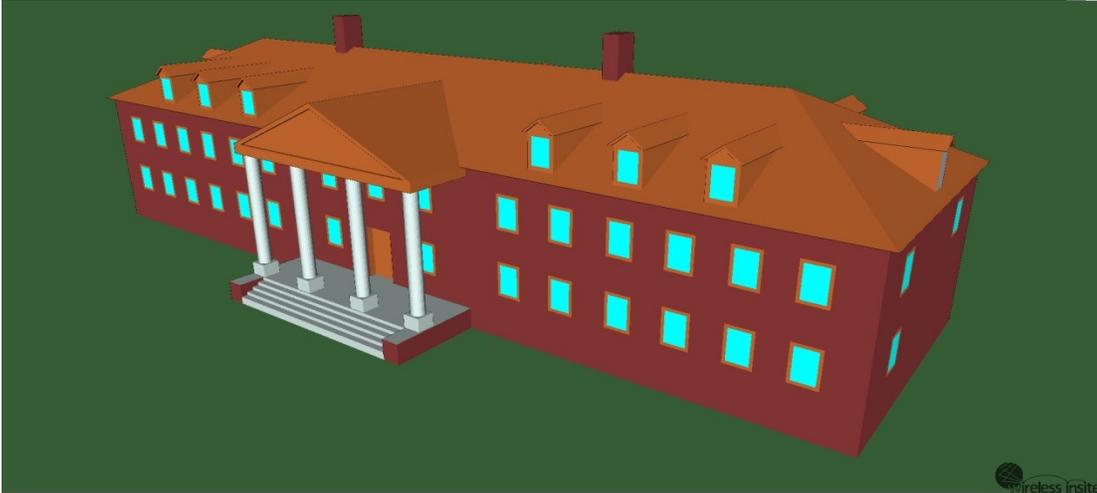
225 Version 2 South Side



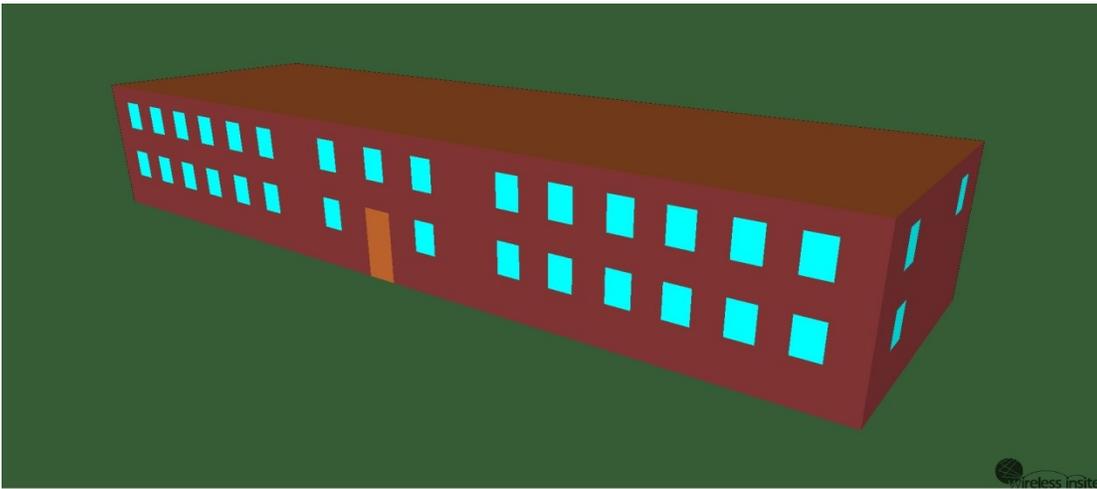
225 Version 3 South Side

Potomac (092)

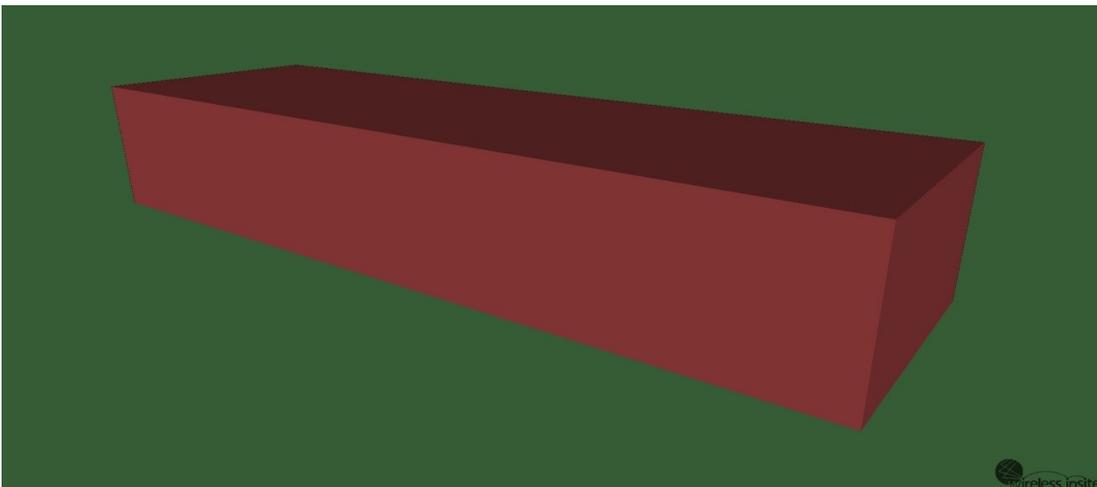




092 Version 1



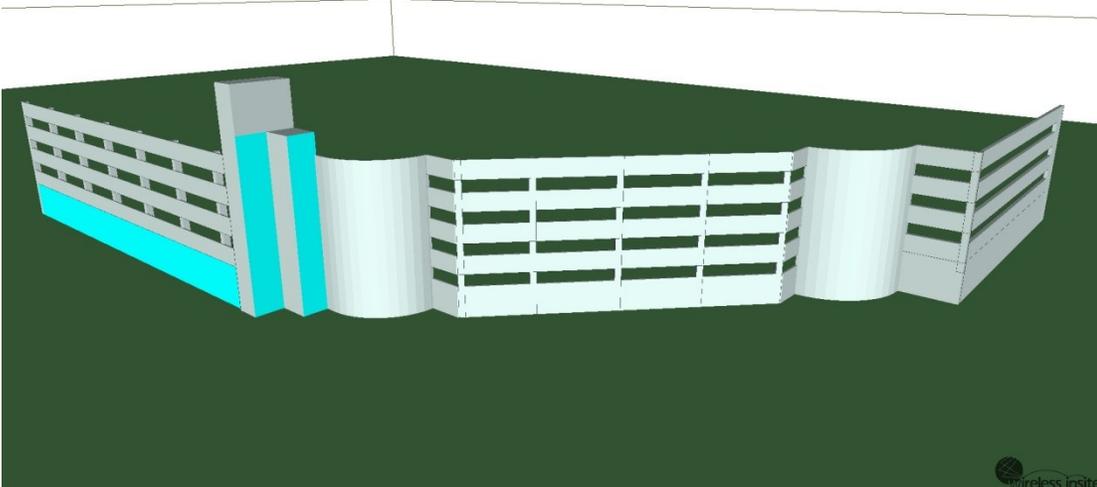
092 Version 2



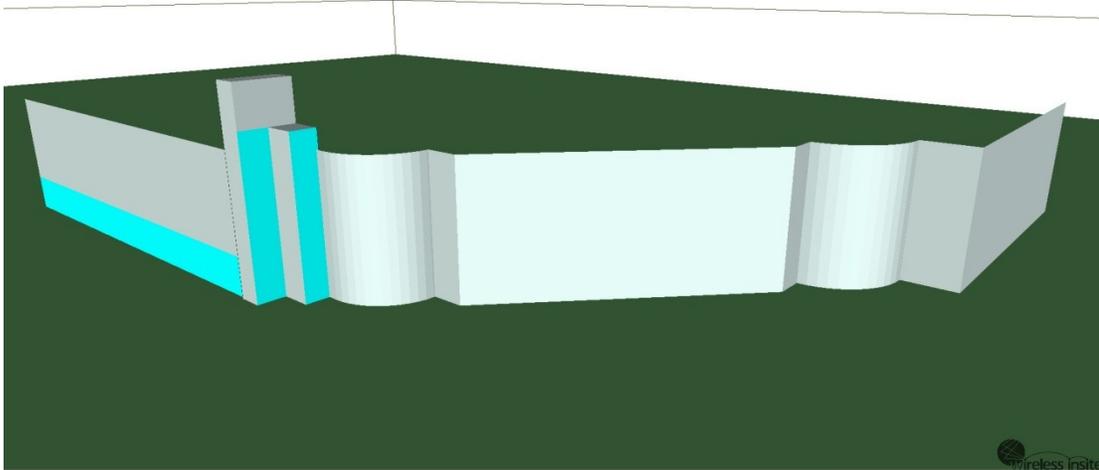
092 Version 3

Regents Parking Garage (202)

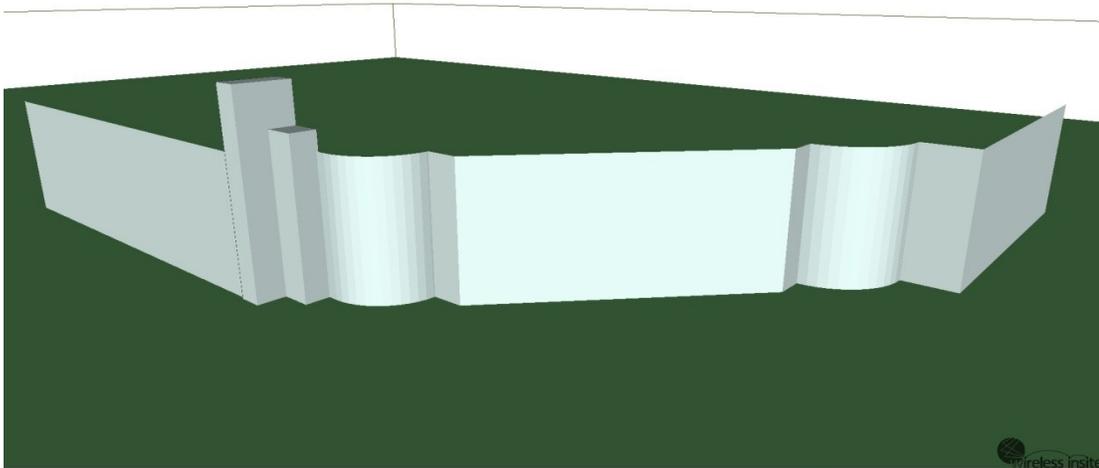
This building lies on the edge of the area where future field tests will be conducted. We make the assumption that any electromagnetic waves passing into this structure will be lost and only model the surfaces facing our field test area to capture reflections off of the structure.



202 Version 1



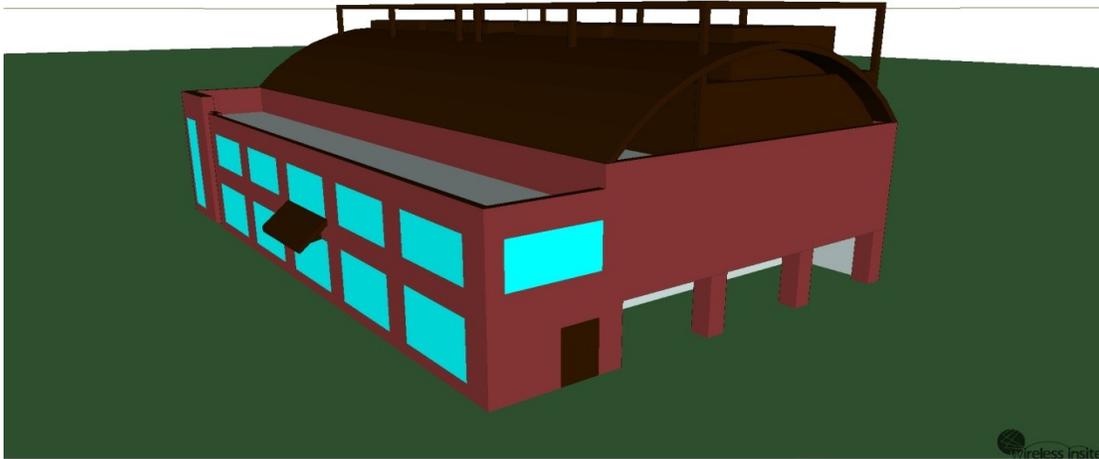
202 Version 2



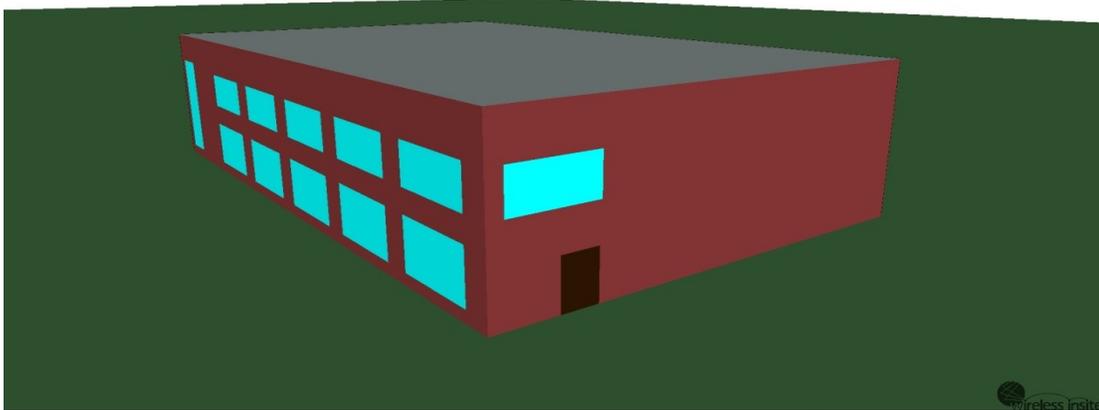
202 Version 3

Satellite Central Utility Building 4 (405)

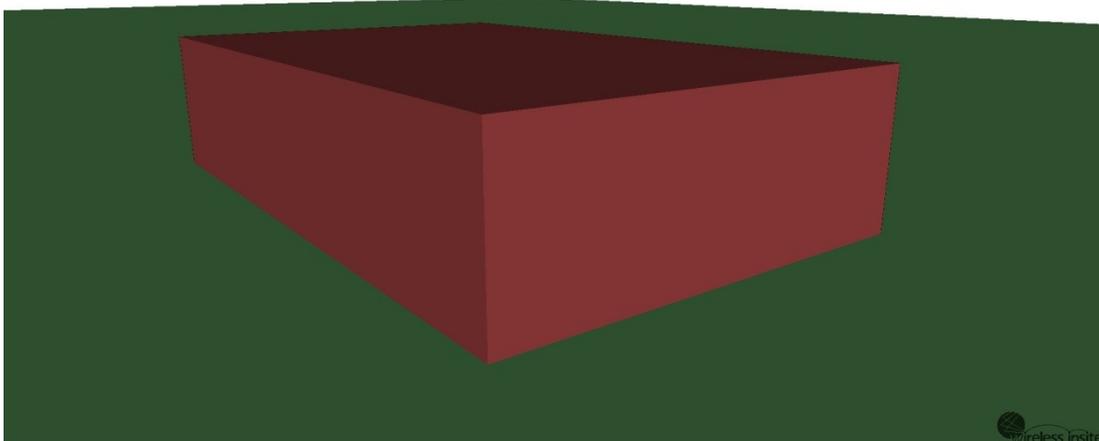




405 Version 1



405 Version 2

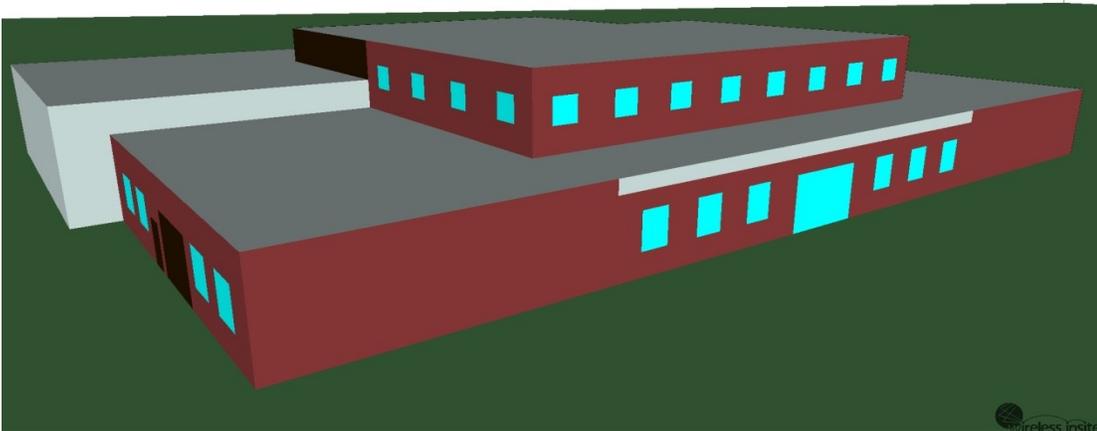


405 Version 3

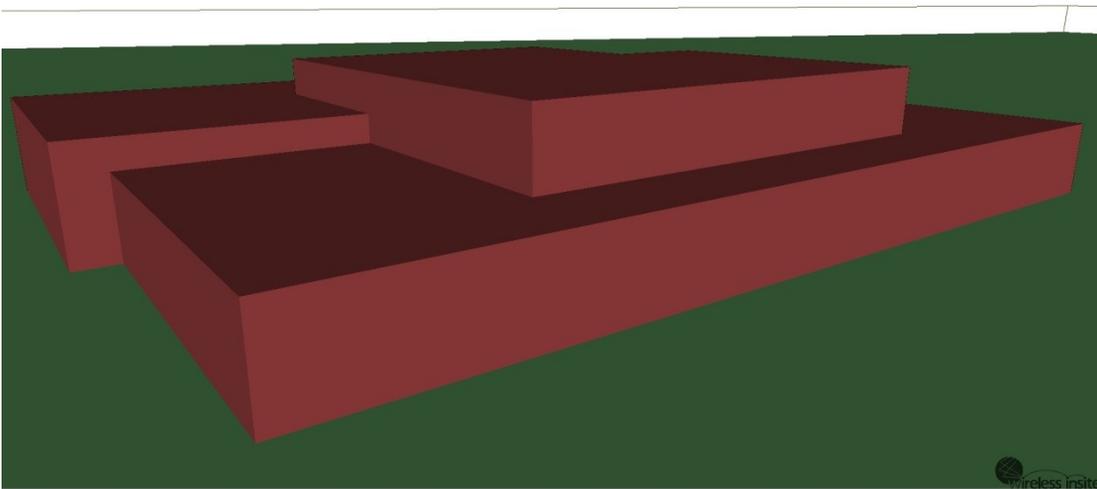
Wind Tunnel (081)



081 Version 1



081 Version 2



081 Version 3

Appendix B: List of Contributors

Name (last, first)	Organization
Vo, Ian	Columbia University
Kennedy, Chris	Gardetto Engineering
Matusiak, Jason	Gardetto Engineering
Beecher, Matt	Laboratory for Telecommunications Sciences
Clancy, Charles	Laboratory for Telecommunications Sciences
Graham, Richard	Laboratory for Telecommunications Sciences
Seligman, Matt	Laboratory for Telecommunications Sciences
Walker, Brenton	Laboratory for Telecommunications Sciences
Devillar, Justin	Naval Postgraduate School
Harley, Peter	Naval Postgraduate School
Cho, Inkeun	University of Maryland
Dunbar, Carson	University of Maryland
Hahn, Dongwoon	University of Maryland
Jesuraj, Jonathan	University of Maryland
Kim, Youngil	University of Maryland
Lee, Ginnah	University of Maryland
Lin, Kun	University of Maryland
Mundur, Padma	University of Maryland
Schulman, Aaron	University of Maryland
Seastrom, Jessica	University of Maryland
Taylor, Keith	University of Maryland

Bibliography

- [1] Atheros AR5006EXS. AR5006EXS Technology Overview.
[http://www.wikidevi.com/files/Atheros/specsheets/AR5006EXS_\(AR5424\).pdf](http://www.wikidevi.com/files/Atheros/specsheets/AR5006EXS_(AR5424).pdf)
- [2] Azimuth Systems. Azimuth ACE series. <http://www.azimuthsystems.com/>
- [3] C. Balanis. *Advanced Engineering Electromagnetics*. Wiley, New York, 1989.
- [4] N. Banerjee, M. Corner, D. Towsley, and B. Levine. Relays, base stations, and meshes: enhancing mobile networks with infrastructure. In *MobiCom '08*, San Francisco, CA, September, 2008. ACM.
- [5] L. Barclay. *Propagation of Radiowaves*. 2nd Edition. Institute of Engineering and Technology, London, 2003.
- [6] D. Calvin, Y. Sasson, and A. Schiper. On the accuracy of MANET simulators. In *POMC'02*, Toulouse, France, October 2002. ACM.
- [7] B.A. Chambers. "The grid roofnet: A rooftop ad hoc wireless network." MIT Master's Thesis, June 2002.
- [8] D. Cichon and T. Kurner. Digital mobile radio towards future generation systems: cost 231 final report. *COST European Cooperation in the Field of Scientific and Technical Research*. 1993.
- [9] C. Clancy, B. Walker. MeshTest: laboratory-based testbed for large topologies. In *TridentCom '07*, Orlando, FL, May 2007.
- [10] P. De, A. Raniwala, S. Sharma, T. Chiueh. MiNT: a miniaturized network testbed for mobile wireless research. In *InfoCom '09*, Miami, FL, March 2009.
- [11] D. Giustiniano, I. Tinnirello, L. Scalia, and A. Levanti. Revealing transmit diversity mechanisms and their side-effects in commercial IEEE 802.11 cards. *Telecommunication Networking Workshop on QoS in Multiservice IP Networks, 2008. IT-NEWS 2008. 4th International* , vol., no., pp.135-141, Feb. 2008
- [12] D. Green and A. Obaidat. An accurate line of sight propagation performance model for ad-hoc 802.11 wireless LAN (WLAN) devices. In *IEEE International Conference on Communications, 2002. ICC 2002*. Vol. 5, pg 3424-3428, 2002.
- [13] R. Gummadi, D. Wetherall, B. Greenstein, and S. Seshan. Understanding and mitigating the impact of RF interference on 802.11 networks. In *SIGCOMM '07*, Kyoto, Japan, 2007.

- [14] D. Hahn, G. Lee, B. Walker, M. Beecher, and P. Mundur. Using virtualization and live migration in a scalable mobile wireless testbed. *SIGMETRICS Perform. Eval. Rev.* 38, December 2010. pg 21-25. ACM.
- [15] M. Iskander and Z. Yun. Propagation prediction models for wireless communication systems. *IEEE Trans. Veh. Tech.*, 50:662-672, 2002.
- [16] JFW Industries. RF Transceiver Test Systems. <http://www.jfwindustries.com>
- [17] G. Judd and P. Steenkiste. A simple mechanism for capturing and replaying wireless channels. In *E-Wind'05*, Philadelphia, PA. August 2005. ACM
- [18] G. Judd and P. Steenkiste. Characterizing 802.11 wireless link behavior. *Wireless Networks*, 16(1), January 2010.
- [19] A. Kamerman and N. Erkocevic. Microwave oven interference on wireless lans operating in the 2.4 GHz ISM band. In *Proc. PIMRC '97*, Helsinki, Finland, Sept 1997.
- [20] J. Keller. "Geometric theory of diffraction," *Journal of the Optical Society of America*, vol. 62, pp. 1448-1461, 1974.
- [21] A. Keränen, J. Ott, and T. Kärkkäinen. The one simulator for DTN protocol evaluation. In *Simutools '09*, Rome, Italy, March 2009. ICST.
- [22] S. Kim, B. Guarino, T. Willis, V. Erceg, S. Fortune, R. Valenzuela, L. Thomas, J. Ling, and J. Moore. Radio propagation measurements and prediction using three-dimensional ray tracing in urban environments at 908 MHz and 1.9 GHz. *IEEE Trans. Veh. Tech.*, 48:931-946, 1999.
- [23] Y. Kim, K. Taylor, C. Dunbar, B. Walker, P. Mundur. Reality vs emulation: running real mobility traces on a mobile wireless testbed. In *HotPlanet'11*, Washington DC, USA, June 2011. ACM.
- [24] D. Kotz, C. Newport, R. Gray, J. Liu, Y. Yuan, and C. Elliot. Experimental evaluation of wireless simulation assumptions. In *MSWiM '04*, Venice, Italy, October 2004. ACM.
- [25] K. Kowalik, M. Bykowski, B. Keegan, and M. Davis. Practical issues of power control in IEEE 802.11 wireless devices. In *ICT'08*, St. Petersburg, Russia, June 2008. IEEE.
- [26] S. McCanne and S. Floyd. The network simulator NS-2. <http://isi.edu/nsnam/ns/>, April 1999.
- [27] OPNET Tech. Opnet, <http://www.opnet.com>.

- [28] C. Phillips, D. Sicker, D. Grunwald. Bounding the error of path loss models. In *IEEE Dynamic Spectrum Access Networks (DySPAN)*, May 2011. IEEE.
- [29] C. Phillips, D. Sicker, D. Grunwald. A survey of wireless path loss prediction and coverage mapping methods. *Communications Surveys & Tutorials, IEEE* , vol.PP, no.99, March 2012, pp.1-16.
- [30] L. Piazzzi and H. Bertoni. Achievable accuracy of site-specific path-loss predictions in residential environments. *IEEE Trans. Veh. Tech.*, 48:922-930, 1999.
- [31] R. Punnoose, P. Nikitin, and D. Stancil. Efficient simulation of ricean fading within a packet simulator. In *Proc. of the IEEE Conference on Vehicular Technology*. Boston, MA, September 2000. IEEE.
- [32] REMCOM. Wireless InSite, Version 2.6.2. <http://www.remcom.com/wireless-insite>
- [33] REMCOM. Wireless InSite Application Examples. <http://www.remcom.com/examples/category/wireless-insite>
- [34] T. Rappaport. *Wireless Communications: Principles and Practice*. Prentice-Hall, Englewood Cliffs, NJ, 2002.
- [35] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Lui, and M. Singh. Overview of the ORBIT radio grid testbed for evaluation of next-generation wireless network protocols. In *Proc. of WCNC'05*, New Orleans, LA, March 2005. IEEE.
- [36] K. Rizk, J. Wagen, and F. Gardiol. Two-dimensional ray-tracing modeling for propagation prediction in microcellular environments. *IEEE Transactions on Vehicular Technology*, Vol. 46, no.2, pg 508-518. May 1997. IEEE.
- [37] Scalable Network Tech. Qualnet. <http://www.scalable-networks.com/>.
- [38] J. Schuster and R. Luebbers. Hybrid SBR/GTD radio propagation model for site specific predictions in an urban environment. *12th Annual Rev. of Progress in Applied Computational Electromagnetics*, Vol.1, pg 84-92, 1996.
- [39] B. Sklar. Rayleigh fading channels in mobile digital communications systems part I: characterization. *IEEE Commun. Mag.*, 35(7):90-100, July 1997. IEEE.
- [40] K. Tan, D. Wu, A. Chan, and P. Mohapatra. Comparing simulation tools and experimental testbeds for wireless mesh networks. In *WoWMoM'10*, Montreal, Canada, June 2010. IEEE.

- [41] S. Tan and H. Tan. Propagation model for microcellular communications applied to path loss measurements in Ottawa city streets, *IEEE Trans. Veh. Technol.*, vol. 44, no. 2, pp. 313-317, Aug. 1995
- [42] A. Toflove, S. Hagness. *Computational Electromagnetics: the Finite Difference Time Domain Method*. Artech House Publishers, 2000.
- [43] University of Maryland. Campus Map. <http://www.umd.edu>
- [44] B. Walker, C. Clancy. A quantitative evaluation of the meshtest wireless testbed. In *TridentCom '08*, Innsbruck, Austria, March 2008.
- [45] J.H. Whitteker. Measurements of path loss at 910 MHz for proposed microcell urban mobile systems, *IEEE Trans. Veh. Technol.*, vol. 37, no. 3, pp. 125-129, Aug. 1998