

REPORT DOCUMENTATION PAGE					Form Approved OMB No. 0704-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>						
1. REPORT DATE (DD-MM-YYYY) 01/01/2006		2. REPORT TYPE Technical Report - Doctoral Thesis			3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE Models and Methodologies for Realistic Propagation Simulation for Urban Mesh Networks				5a. CONTRACT NUMBER		
				5b. GRANT NUMBER		
				5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)				5d. PROJECT NUMBER		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Delaware Newark DE United States				8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Army Research Laboratory Adelphi MD United States				10. SPONSOR/MONITOR'S ACRONYM(S)		
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT A = Approved For Public Release 12/4/2015 No						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT By providing connectivity to mobile users, mesh networks are poised to become a major extension of the Internet. More than 300 cities and towns have plans to deploy mesh networks, and several dozen cities have already deployed these networks. These mesh networks are meant to enhance city and emergency services communication as well as to provide city-wide, low-cost, ubiquitous Internet access for residents and visitors. Such networks promise to bring dramatic changes to data						
15. SUBJECT TERMS						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code)	

**MODELS AND METHODOLOGIES FOR REALISTIC
PROPAGATION SIMULATION FOR URBAN MESH
NETWORKS**

by

Vinay Sridhara

A dissertation submitted to the Faculty of the University of Delaware in
partial fulfillment of the requirements for the degree of Doctor of Philosophy in
Electrical and Computer Engineering

Spring 2007

© 2007 Vinay Sridhara
All Rights Reserved

**MODELS AND METHODOLOGIES FOR REALISTIC
PROPAGATION SIMULATION FOR URBAN MESH
NETWORKS**

by

Vinay Sridhara

Approved: _____
Gonzalo R. Arce, Ph.D.
Chair of the Department of Electrical and Computer Engineering

Approved: _____
Eric W. Kaler, Ph.D.
Dean of the College of Engineering

Approved: _____
Conrado M. Gempesaw II, Ph.D.
Vice Provost for Academic and International Programs

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed: _____

Stephan K. Bohacek, Ph.D.
Professor in charge of dissertation

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed: _____

Charles Boncelet, Ph.D.
Member of dissertation committee

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed: _____

Chien-Chung Shen, Ph.D.
Member of dissertation committee

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed: _____

Saishankar Nandagopalan, Ph.D.
Member of dissertation committee

To

Horanadu Shree Annpoorneshwari

and

Amma and Anna

ACKNOWLEDGEMENTS

I first thank my Advisor Prof. Stephan Bohacek for all the guidance and encouragement during the course of my research. He was always able to make time for me no matter how busy he was. I would like to thank him for inducing the never give up attitude in me. He never ceases to surprise me with his approach and attitude towards a problem. I would always appreciate and cherish my association with Prof. Bohacek from the USC days where I started working with him as a masters student. I would like to express my thanks and deep appreciation towards him. I feel lucky to have an understanding and supporting advisor like him.

I thank my dissertation committee members: Prof. Chien-Chung Shen, Prof, Charles Boncelet and Dr. Saishankar Nandagopalan. I would like to thank them for providing valuable advise and investing support in me. I would like to thank Prof. Shen for giving me an opportunity to work with him and his group. I would like to thank Dr. Saishankar Nandagopalan for his time and effort to be on my dissertation committee. I would like to thank all my Professors and staff at Department of Electrical Engineering for their guidance and support. I also thank Dr. Sanjiv Nanda and Arnaud Meylan for their support and guidance during my internship at Qualcomm Research and Development.

I am fortunate to have interacted and worked with my colleagues in our research group. Jonghyun, Hweechul, Gaurav, Andjela, Peng, Andreas and Damian have all made it a very enjoyable time working in our research lab. I would like to thank Jonghyun for working with me on several projects which were very productive.

I would also like to thank Hweechul for working with me on several measurement based projects which were the hardest due to weather, time and place constraints.

It is an understatement to say that I am blessed with really understanding and loving parents. They have supported and embraced all the decisions that I have taken in my life. I thank my dad (anna) for being there for me and having guided me through my formative years as a father, mentor and more as a great friend. I thank my mother (amma) for providing me support by listening to me at times of excitement and despair alike. I love to thank them for everything that they have provided for me. I thank my grandparents (anni, sathu ajji, mandya ajji) for their unconditional love and support. My grandmother (sathu ajji) has been one of my greatest mentors and friend.

I thank my closest and dearest set of friends. I thank each one of them Nagendra, Karthik and Vaibhav for being there. I have enjoyed and continue to enjoy every moment spent with you guys. Vaibhav Dutta, my friend and philosopher has always been there for me at every moment no matter what. I especially thank him for being such a great friend and a great human being. This thesis is a dedication to my parents and God, the entities that form the core of my life. I thank all my family and friends who have supported me in this journey.

My doctoral work was funded by the U.S. Army Research Laboratory under the the Collaborative Technology Alliance Program, and the University of Delaware. My thanks to all for their generous funding that made this research possible.

TABLE OF CONTENTS

LIST OF FIGURES	xi
LIST OF TABLES	xix
ABSTRACT	xx

Chapter

1 INTRODUCTION	1
1.1 Urban Mesh Networks	1
1.2 Motivation	3
1.3 Problem statement	4
1.4 Key contributions	4
1.5 Thesis outline	6
2 REALISTIC PROPAGATION SIMULATION OF URBAN MESH NETWORKS	8
2.1 Stochastic Models of Propagation	10
2.2 Characteristics of Urban Propagation	14
2.2.1 Overview	14
2.2.2 Transmissions	15
2.2.3 Reflections	18
2.2.4 Multipath Fading	20
2.2.5 Diffraction	21
2.2.6 Scattering	24
2.2.7 Time-Varying Channel Gain	24

2.2.8	Delay Spread	25
2.3	Computational Techniques	27
2.3.1	Overview	27
2.3.2	Outdoor Propagation	28
2.3.3	Indoor Propagation	33
2.3.4	Computational Complexity	33
2.3.5	Computational complexity for packet simulation	35
2.3.6	Validation	37
2.4	Impact of Realistic Propagation	41
2.5	Related Work	44
3	THE GRAPHICAL PROPERTIES OF URBAN MESH NETWORKS	46
3.1	Related work	48
3.2	Simulations	49
3.3	Connectivity	49
3.4	Duration in and out of the largest connected component	52
3.5	Degree distribution of nodes in an Urban Mesh Network	53
3.6	Centrality distribution of node in an urban mesh network	54
3.7	Spatial aspects of the graph	55
4	OBSERVATIONS AND MODELS OF TIME-VARYING CHANNEL GAIN IN CROWDED AREAS	60
4.1	Introduction	60
4.2	Measurement setup	62
4.3	Diffusion process	63
4.3.1	Parameter estimation methodology	65
4.4	Data analysis	66
4.4.1	Stationary distribution	67
4.4.2	Transition probabilities	70
5	PERFORMANCE OF 802.11B/G IN THE INTERFERENCE	

LIMITED REGIME	73
5.1 A brief primer on 802.11 b/g physical layer	75
5.2 Experiment Description	77
5.2.1 Experimental Setup and Protocol	77
5.2.2 Experimental results	80
5.3 Probability of Transmission Error in the Interference Limited Regime	82
5.3.1 Analysis	82
5.3.2 Discussion	87
5.4 Bit-rate Selection in the Interference Limited Regime	88
5.5 Related work	92
6 SIMULATION METHODOLOGY	94
6.1 Simulation overview	98
6.2 City maps	98
6.3 A brief primer on UDelModels urban mobility simulator	99
6.3.1 Activity model	101
6.3.2 Task model	102
6.3.3 Agent Model - Node Dynamics and Interactions	103
6.4 Urban Population Size	104
7 INSIGHTS GAINED FROM REALISTIC SIMULATIONS	106
7.1 Coverage in urban mesh networks	107
7.2 Achievable bit rates in urban mesh networks	112
7.3 Scalable and efficient paging protocol utilizing the urban mesh topology	114
7.3.1 The Topology of LUMNets	116
7.3.2 Paging in Wireless Networks	118
7.3.3 Scalable Paging for LUMNets	121
7.3.3.1 Overview	121

7.3.3.2	Similarity and Gratuitous Registration	123
7.3.3.2.1	Active TGW	124
7.3.3.2.2	Gratuitous Registration	125
7.3.3.3	Detecting a Change of Tendril and Appropriation of Neighbors' Active TGW Lists	127
7.3.3.3.1	Detecting a Change of Tendril	127
7.3.3.3.2	Detecting Entry into a Tendril	132
7.3.3.4	Other issues	133
7.3.4	Analysis of Performance	135
7.3.5	Simulation results	139
7.3.5.1	Client information maintenance	140
7.3.5.2	Route Search/Paging Overhead	142
8	CONCLUSIONS AND FUTURE WORK	144
8.1	Propagation simulator	144
8.2	Graphical properties of urban mesh and mobile ad hoc networks . . .	145
8.3	Variation in propagation characteristics	145
8.4	Effect of interference on performance of 802.11b/g	146
8.5	Simulation framework	146
8.6	Insights gained from realistic simulations	147
8.7	Future work	148
	BIBLIOGRAPHY	151

LIST OF FIGURES

1.1	Urban mesh networks are composed of wireless links between infrastructure nodes and wireless links between infrastructure and mobile clients. Both types of nodes may be indoors or outdoors and may be elevated or at street level. Communication between mobile clients and the Internet is provided via multiple wireless hops over the infrastructure (dotted lines). Many types of signal propagation are relevant. In the figure above, the signals that reach the car are impacted by diffraction around the corner of the building, propagation through the interior of the building, transmission through an exterior building wall, and reflections off of a building and off of the ground.	3
2.1	Left, channel gain generated by a correlated lognormal model. Right, channel gains generated by ray-tracing in an urban area. The rectangle shapes are buildings and the shaded boxes inside the rectangles represent channel gain to locations on the first floor of the building. In both the left and right frames, the star marks the location of the transmitter.	13
2.2	Routing in an Urban ad hoc Network. The source and destination are inside of two different buildings. The shortest hop path travels along the road, it does not follow the shortest geometric route (i.e., a straight line).	13
2.3	Transmission through a wall. The left-hand frame shows a side-view of the simulated scenario. The right-hand frames depict the attenuation of the signal when transmitted through walls of different kinds of materials. It also depicts the nature of attenuation due to relative heights of the transmitter and the receiver.	17

2.4	Propagation down a street. The left-hand frame shows a top view of the simulated scenario. The right-hand frames depict the effect of propagation down an urban canyon for different wall materials, different attenuation exponents, and for the free space.	19
2.5	Fast fading in narrow band and wide band communication scenarios. Left: the signal strength as a function of frequency for a narrow band signal propagating down an urban canyon. Middle: signal strength of a narrow and wide band signal as a function of receiver position perturbation when propagating down an urban canyon. Right: variance in signal strength due to change in receiver position as a function of bandwidth for when propagating down a urban canyon as shown in Figure 2.4.	21
2.6	Diffraction around a corner (left) and over a building (right). The left-hand frame and second from the right frame show side views of the simulated scenario. The other frames show the attenuation as a function of h . As h increases, the signal strength decreases. The dashed lines show the signal strength if the loss was only due to free space propagation while the solid line shows the signal strength when the loss due to diffraction is also included.	22
2.7	Left, experimental setup for measuring the time varying nature of propagation in a busy street in Philadelphia. Right, an observed time-series of the channel gain.	25
2.8	Bit-error probabilities as a function of delay-spread and SNR. Simulation of 802.11b at 1 Mbps (without a RAKE receiver) and 802.11a at 6 Mbps. The different curves for 802.11b correspond to different RMS-delay spreads. For a delay-spread of 7ns to 220 ns, 802.11a gave the same relationship between SNR and bit-error rate, hence this relationship is shown with a single curve.	28
2.9	Left: a top-view of the scene on the right. The lines shown are found from 2-D ray-tracing. Right:: Two vertical plane paths (found from 2-D ray-tracing) and 5 ray paths.	30

2.10	Beam Tracing. Suppose a beam strikes the wall tile marked A. The reflection of this beam is modeled as a beam emanating from the virtual source as shown. This beam is between and is defined by the two rays marked 1 and 2. These rays are used to select the floor tile ray-neighbors that this beam strikes (since ray-neighbors are precomputed, this selection is computationally efficient). Similarly, the rays 1 and 2 are used to select the set of A's wall tile ray-neighbors that the beam strikes. One such wall tile is shown in the upper right and is marked B. The beam that strikes tile B is defined by the rays marked 3 and 4. These rays define the beam that is transmitted through tile B and can be used to find the beam that is reflected off of B. This process of reflection and transmission is repeated until the beam's strength is small enough that it can be neglected (e.g., -110 dB).	32
2.11	Simulation Execution Time. The left-hand frame shows the simulation execution time for three simulation techniques. The right-hand plot shows a break-down of the execution time when realistic propagation is included. The propagation computation time includes the time to identify which floor tile a node is nearest and the time to look-up the propagation between nodes from the propagation matrix.	37
2.12	Observed and estimated path loss in the campus environment. . . .	39
2.13	Observed and estimated path loss at an intersection in Philadelphia.	40
2.14	Observed and estimated path loss in an indoor environment. The measurement points shown on the x-axis correspond to the numbered locations in the map. The map is of the third floor of Evans Hall in University of Delaware campus.	40
2.15	Statistics of degree distribution in random way-point with open-space propagation, urban mobility with open-space propagation and urban mobility with realistic urban propagation. The mesh infrastructure nodes are fixed on top of streetlights, while the client nodes are mobile. The coefficient of variance (CoV) is the standard deviation divided by the mean.	43
2.16	Performance of a network under different types of mobility and propagation.	43

3.1	Signal Strength in an Urban Environment. The colors indicate the signal strength from a transmitter on the right-hand side of the modeled area. While the figures show a large variation in coverage, the only difference between the two figures is that the position of the transmitter has moved 10 meters.	47
3.2	Cities used for Simulation. the rectangles are office buildings with evenly distributed offices and hallways (not shown). Outside sidewalks are shown. The mobile nodes are restricted to sidewalks, offices and hallways.	50
3.3	Connectivity Ratio as a Function of the Number of Node in the Population. The upper curve show the connectivity ratio for just the nodes that are outside, the lower curve is of nodes inside buildings and the middle curve is for all nodes (both inside and outside). . .	51
3.4	Left-hand figure. The Complimentary Distribution of the Duration that a Node Spends in the LCC. Right-hand figure. The Complimentary Distribution of the Duration that a Node Spends out of the LCC.	53
3.5	Degree Distribution.	54
3.6	Centrality Density (Histogram).	56
3.7	Spatial Graph Properties of the University of Delaware. The upper right figure indicates the mean degree at each location. If a node never entered a particular region, the region is left uncolored. The upper left plot shows the probability that a node in a location is in the largest connect component. However, two color scales are applied, one for inside locations and one for outside locations. The lower left shows the centrality of nodes in each position. Again, two different color scales are used, one for inside and one for outside. The lower right figure indicates the frequency or proportion of time a node was in each location.	57
3.8	Spatial Graph Properties of the Paddington Environment	58
4.1	Measurement Configuration. T-L and T-H indicate transmitter at lower and elevated positions respectively. R indicates the receiver position.	63

4.2	Observed Time Series and Synthetic Time Series of Channel Gains in Different Scenarios. The synthetic time series were generated with a diffusion model. Note that difference in scales.	67
4.3	Observed and Fitted Stationary Distribution for Three Different Pedestrian Densities and Two Different Antenna Heights	68
4.4	The above table shows the description of different experimental scenarios and the parameters and errors for stationary distribution and probability transition functions	69
4.5	The plot shows the least squares fit for relationship between the variance and the mean channel gain	70
4.6	Transition Probabilities. The upper plots show the transition probabilities from the high density pedestrian scenario with a low antenna height, while lower set of plots show the transitions probabilities from the lower pedestrian density with the elevated antenna. In each case three transition probabilities are shown corresponding to three different initial conditions (x_0), namely, x_0 is the 25th percentile, the 50th percentile, and the 75th percentile. . .	71
4.7	Relationship between variance and the model parameters γ and σ . .	71
5.1	Structure of PPDU for 802.11 b/g standards. The top two figures show the PPDU for 802.11 b with long and short preambles respectively. The bottom most figure shows the PPDU for the ERP OFDM PLCP used in 802.11g.	76
5.2	The figure shows the block diagram of the experimental setup for all sets of measurements collected for analysis. The red lines indicate the RF cables used and the green thick line indicate the Ethernet cable connections used for controlling the experiment.	78
5.3	The observed probability of loss during the experiments described in Section 5.2.2. Note that not every frame experiences a collision, and hence, even at very low SIR, not all frames are lost. The fraction of frames lost depends on the probability that a frame experiences interference, the synchronization scheme used by the physical layer and the bit-rate used. For example, the probability of observing a	

	transmission error at low SIR is equal to the probability of a collision occurring, i.e., every frame that experiences a collision is lost. . . .	82
5.4	Possibilities of a sender's PLCP transmission overlapping with an interferer's transmission	83
5.5	The possibilities of a sender's payload overlapping with the interferer's transmission. This can lead to CRC errors or loss of lock errors.	84
5.6	Probability of synchronization failing as a function of SIR.	85
5.7	Probability of CRC error observed during the experiments described in Section 5.2 as a function of frame size. The left-hand size is for 2Mbps while the right-hand size is for 11 and 12 Mbps, as indicated. Note that the probability of CRC error occurring is linear in the frame size.	86
5.8	Probability of error as a function of SIR for different types of modulation.	87
5.9	Average time require to complete a transmission of a frame when $\rho = 0.3$ and $\lambda = 1/10 \text{ m sec}^{-1}$ and when the frame is 40B (left-hand side) and 1400B (right-hand side).	91
5.10	Optimal Rate Regions. For a given combination of channel utilization, ρ , frame size, and SIR, there exists a particular bit-rate that results in the smallest delay. The above shows the optimal bit-rate in each region of the packet size/SIR plane. The region where a particular bit-rate is optimal is marked with a particular color. In these plots, $\lambda = 1/10 \text{ m sec}^{-1}$	92
6.1	Different stages involved in realistic simulation of urban mesh networks. The output of each stage is represented in italics.	97
6.2	The figure illustrates the three tier model employed in the mobility simulation.	101

7.1	(a) The figure illustrates the disparity in coverage of indoor and outdoor locations due to infrastructure nodes placed outdoors at various channel loss thresholds. (b) The figure illustrates the coverage on different floors. The results obtained are from a simulation carried out for a section of downtown Chicago using UDel propagation simulator. 27 infrastructure were placed outdoors on lamp posts (3.5m) with a separation of 50m between them. . . .	108
7.2	The figure shows the coverage inside of buildings due to the infrastructure nodes placed outdoors with a separation of 50m between them. It can be observed as the floors number increases the coverage becomes almost nil. (The darker color indicates better coverage)	109
7.3	Coverage. These plots show the fraction of people that can communicate with the base station in different scenarios. The left hand column considers only nodes that are outside, the middle column only considers nodes that are inside, while the right hand column includes all nodes. The upper row is for coverage at 54 Mbps with 802.11g (-69 dBm signal strength). The middle row is for 11 Mbps (-85 dBm), while the bottom row is for 1 Mbps (-93 dBm). The x-axis indicates the distance between the infrastructure nodes. Within each infrastructure scenario, the right most bar indicates the coverage when mobile nodes are not allowed to act as relays, while the left column shows the coverage when mobile nodes may act as relays.	113
7.4	Achievable bit rates.	114
7.5	The figure depicts the structure of multihop communication in a building. This topology resembles that of a tendril.	117
7.6	A conceptual visualization of a mesh networks. The picture above highlights the different elements of multihop wireless mesh networks.	119
7.7	Heterogeneous Propagation. The 5th and 95th percentile of the distances from the WFIN to nodes that are exactly 1,2, 3, and 4 hops away from the WFIN.	121
7.8	Flowchart describing the functionality of a WFIN on reception of a registration from a client.	128

7.9	Flow chart depicting the functionality of a Tendril Gateway or Direct Client.	130
7.10	Flow chart depicting the functionality of a Tendril Client.	131
7.11	(a). Mean time to receive three neighborhood broadcasts when a new node broadcasts its TGW history as a function of the total number of neighbors and T_{NB} , the time between broadcasting neighborhood information. (b) Mean time to receive a TGW beacon when there are different numbers of TGWs. The x-axis is the time between TGW beacon transmissions.	139
7.12	Mean time that a node is unreachable after it changes from a direct client to a tendril client. It is assumed that the $T_{INBeacon} = 0.5\text{sec.}$, $\text{MAX_IN_MISSED_BEACONS} = 6$, $T_{NB} = 15 \text{ sec.}$, and the other parameters are as indicated.	140
7.13	Rates and types of packets that are sent and received. These packet types are used to maintain the client information.	141
7.14	Mean time for a client node to register after entering a tendril. In these simulations, $\text{MAX_IN_MISSED_BEACONS} = 6$ and $k = 3$. The values of the other parameters are as shown. The two building scenario was used in these simulations.	142
7.15	Comparison of overhead for performing a single page using the efficient paging algorithm and route search technique used in AODV. Each building has 10 floors and is populated with 100 nodes.	143

LIST OF TABLES

2.1	Coverage and computation time	23
2.2	Computational complexity	35
5.1	SIR and SNR required for probability of error of $1/2$	88
7.1	Mesh Network Scenarios	114

ABSTRACT

By providing connectivity to mobile users, mesh networks are poised to become a major extension of the Internet. More than 300 cities and towns have plans to deploy mesh networks, and several dozen cities have already deployed these networks. These mesh networks are meant to enhance city and emergency services communication as well as to provide city-wide, low-cost, ubiquitous Internet access for residents and visitors. Such networks promise to bring dramatic changes to data accessibility and hence have a major impact on society.

Simulation is the integral step in the validation of mesh networking protocols. However, currently, most simulations use the trivial disk propagation model (i.e., the signal propagates exactly R meters and no further) or a highly idealized propagation model such as the free-space or the two-ray model. Due to the presence of buildings, propagation in urban environments is far more complicated than the propagation presented by these simple models. Consequently, channel variability, which is a key aspect of wireless networking, is not well modeled in today's simulations. The result is that many insights gained from the simplistic environment do not necessarily hold in the urban environment, casting doubt on the applicability of the conclusions drawn from these simulations. While the reasons that realistic propagation models are typically not included into network simulations is not well documented, it seems that one of the main reasons include the belief that propagation modeling is computationally intractable and that propagation cannot be realistically modeled due to small-scale fading and delay spread.

This dissertation focuses on developing models for realistic propagation simulation and the implications of these models on simulation of urban mesh networks. The dissertation also provides measurements and models for time-varying nature of propagation and characterization of the effect of interference on packet success probability.

To this effect a realistic propagation simulator has been developed. This propagation simulator can estimate channel gain from any given point to any other given point in the simulated region (e.g., city downtown). In addition to estimating channel gain, other key metrics such as delay spread, angle of arrival and departure can be computed. It is shown that the simulation time scales linearly with the number of processors and also this computation needs to be carried out only once for a simulated region.

Channel variability that results in temporally varying topologies is one of the biggest challenges faced by wireless networks. This dissertation presents measurements and models of propagation between stationary transmitters and receivers in a dynamic environment. A diffusion-based stochastic model is presented that can be used to accurately model the time-varying propagation.

Propagation model only provides the signal strength and other propagation metrics. However, realistic simulation requires the packet success probability (PSP). The signal to Interference plus noise ratio impacts the probability with which the receiver will be able to successfully decode the signal. This dissertation aims at characterizing the effect of varying interference and signal strengths on the packet success probability. As a result of this characterization a model for bit-rate selection is provided.

This dissertation provides a framework for realistic simulation of urban mesh networks. The models developed are available online and are compatible with widely adopted discrete event packet simulators such as NS2 and QualNet.

Chapter 1

INTRODUCTION

Wireless networks bring about a paradigm shift in the way of information dissemination. The advent of wireless networks has provided the users with the flexibility of "*access on the go*". In recent years wireless networks has seen tremendous growth in the name of Mobile Ad-Hoc networks, Sensor Networks and more recently, the city wide deployment of Wireless Mesh Networks to provide ubiquitous internet access to the mobile and stationary users alike. The growth of wireless networks is also fueled by the IEEE standards [45] development such as 802.11 for wireless local area networks, the low complexity and low power standard 802.15 for personal area networks, 802.16 for fixed broadband wireless networks and P1415.5 for standardizing the wireless sensor networks.

Mesh networking is a promising technology for providing robust broadband wireless access for various multimedia applications and for various other public safety applications. Recently, there has been interest in developing protocols for urban mesh networks such as the ones to be deployed in Philadelphia [75], San Francisco [22], Taipei [96], Minneapolis [23], Anaheim, California, and Tempe, Arizona.

1.1 Urban Mesh Networks

A distinguishing feature of urban mesh networks is the presence of a large number of infrastructure nodes, each with relatively small coverage. This contrasts cellular networks that have far fewer infrastructure nodes, each with large regions of coverage. Urban mesh networks also differ from traditional mobile ad hoc networks

in that they have a fixed infrastructure that includes some wired base stations. On the other hand, mesh networks are similar to ad hoc networks in that paths may cross multiple wireless links. However, a path in a mesh network typically includes multiple infrastructure to infrastructure hops and a single infrastructure to mobile hop (See Figure 1.1). Another difference between these networks is that while cellular networks are carefully planned and mobile ad hoc networks are not planned at all, urban mesh networks occupy the middle ground in that they are somewhat planned. This lack of careful planning may result in infrastructure nodes having highly overlapped coverage areas, in contrast to cellular networks.

Besides architectural aspects, these networks differ in terms of propagation. In cellular networks, the network performance is largely controlled by propagation between the outdoor base station and the mobile client. As shown in Figure 1.1 propagation between the infrastructure and the mobile client is not the only concern. For example, communication between infrastructure nodes is critical. Also, in cellular networks, transmissions by nearby clients are controlled with CDMA, FDMA, or TDMA, and hence, client-to-client interference is not important (e.g., indoor propagation has little impact on performance). However, in urban mesh networks, the coverage areas of infrastructure nodes may greatly overlap, and hence clients communicating with different infrastructure nodes may experience considerable interference. Consequently, indoor-to-indoor propagation is important as well as the propagation between street level nodes. In summary, the performance of urban mesh networks is impacted by

- propagation between mobile clients, which may be indoors or outdoor and, if outdoors, are at street level,
- propagation between mobile client and infrastructure nodes, where both types of nodes could be either indoors or outdoors, but outdoor infrastructure nodes are typically elevated, and

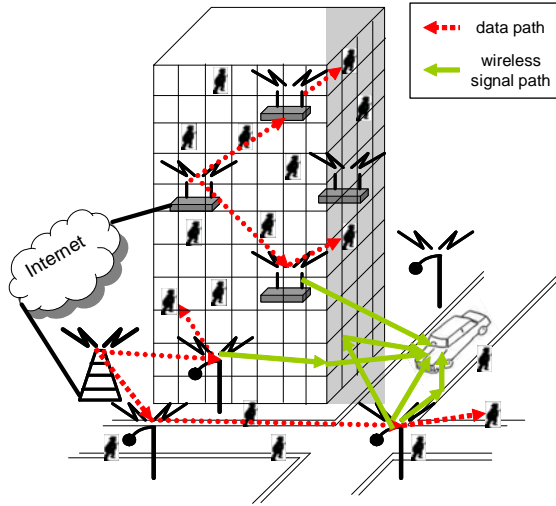


Figure 1.1: Urban mesh networks are composed of wireless links between infrastructure nodes and wireless links between infrastructure and mobile clients. Both types of nodes may be indoors or outdoors and may be elevated or at street level. Communication between mobile clients and the Internet is provided via multiple wireless hops over the infrastructure (dotted lines). Many types of signal propagation are relevant. In the figure above, the signals that reach the car are impacted by diffraction around the corner of the building, propagation through the interior of the building, transmission through an exterior building wall, and reflections off of a building and off of the ground.

- propagation between infrastructure nodes, where, again, these nodes can be indoors or outdoors.

1.2 Motivation

To simulate a wired protocol it is essential to model the link (bandwidth, delay, loss), the desired quality of service of the user, amount and type of the competing traffic etc. The situation is even more complex when wireless protocols are considered. Not only can the network traffic vary as in the wired case, but the wireless links and the resulting topology add a new dimension to the heterogeneity.

Currently, most simulations use the trivial disk propagation model (i.e., the signal propagates exactly R meters and no further) or a highly idealized propagation model such as the free-space or the two-ray model. Due to the presence of buildings, propagation in urban environments is far more complicated than the propagation presented by these simple models. Consequently, channel variability, which is a key aspect of wireless networking, is not well modeled in today's simulations. The result is that many insights gained from the free-space environment and other simplistic models do not necessarily hold in the urban environment, casting doubt on the applicability of the conclusions drawn from today's simulations.

1.3 Problem statement

Simulation is the integral step in the validation of mesh networking protocols. Simulation of a protocol involves simulation of multiple underlying aspects of the network. Wireless propagation is one such aspect that has a predominant effect on the nature of communication links. Specifically propagation has a significant impact on the performance of protocols due to the variation induced by it in terms of link bandwidth and error rate. Hence it becomes important to model the effects of propagation to evaluate the performance of any wireless protocol.

This dissertation focuses on providing models and methodologies for simulating wireless propagation in urban environments. This work proposes and investigates various models for realistically modeling various aspects of propagation to enhance the quality of simulations.

1.4 Key contributions

This dissertation provides models and mechanisms for realistic simulation of wireless propagation in urban environments. The Following are the key contributions of this dissertation.

- At the frequencies used in today's wide band communication, wireless signals may undergo reflections off of buildings and ground, transmissions through walls, and diffractions over and around buildings. Wireless communication extends far beyond what line-of-sight (LOS) communication will offer. Indeed, our simulations show that majority of a node's neighbors are not within LOS. Goals of realistic propagation simulation include simulating realistic coverage and realistic variation of the signal strength. To this effect a propagation simulator capable of estimating key metrics such as signal strength, delay spread etc. between any two given nodes in the network has been developed.
- A study of graphical properties of the topology formed by the nodes using realistic propagation is presented. The study gives an overview of the statistical properties of the graph formed by the nodes in mobile networks in urban environments. This study is aimed as a tool toward development of better protocols for urban mesh networks.
- It is common to see various models for estimating the channel variability between a transmitter and a receiver where one of the nodes is mobile. These models attribute the channel variability to the mobility of the transmitter or receiver. But in practicality it is true that the channel can also vary between a pair of stationary nodes due to the objects moving in the vicinity. This dissertation examines the channel variability between a stationary pair of nodes in a dynamic environment. To this effect a diffusion model has been developed that models the channel variability as a function of the mobile object density in the environment.
- Signal to interference plus noise ratio, SINR, is one of the main factors that affects the quality of wireless communication. While the impact of white

Gaussian noise on a wireless channel is well understood, impact of interference remains one of the less explored areas. With the deployment of dense mesh networks, the interference will be a dominant factor that affects the transmission errors. This work explores the performance of 802.11b/g when subject to interference. The findings are based on various controlled experiments in the laboratory setting. One finding of this work is that in contrast to communication over links where the noise is Gaussian, in 802.11b/g, the probability of successfully transmitting a packet is dominated by the ability of the receiver to synchronize with the carrier. As a result, changing to a lower bit-rate with same synchronization scheme will not make the transmission more resilient to interference. The significance of this result on bit-rate selection is also explored.

- The dissertation also aims at developing a general framework for protocol simulation using the developed models. To this effect an entire chapter is dedicated to discussing various steps involved and overall simulation methodology for realistic simulation of urban mesh networks.

1.5 Thesis outline

The rest of the thesis is organized as follows. The chapter 2 presents a detailed overview of the propagation simulator design and development. This chapter also provides some discussion on the impact that realistic simulation has on the performance of networking protocols. As one might expect, traditional simulations that use random way-point mobility and free-space propagation yield very different results than the simulations carried out using realistic mobility and propagation. It is important to note that realistic simulation of wireless networks is an ongoing effort and that as the computational capabilities increase, more detailed and accurate

simulations are possible. Chapter 3 presents the statistical analysis of the graphical properties of mobile networks. The modeling and characterization of channel variability between two stationary nodes in a dynamic environment is presented in chapter 4. Chapter 5 presents a study of relationship between the SIR, SNR and packet success probability. This chapter also provides a model for bit rate selection for varying SIR conditions. Chapter 6 provides a detailed simulation overview using the propagation simulator. A brief description of the mobility model used for the simulations is also provided in this chapter. Various insights gained from the simulations using these realistic models are presented in chapter 7. Issues such as Coverage, bit rates achievable and topological formations are discussed in this chapter. An efficient paging protocol that makes use of the unique topological formation found in indoor mesh nodes is developed and discussed. Working details of the protocol and simulation results are also provided. Chapter 8 provides concluding remarks and the directions for future work.

Chapter 2

REALISTIC PROPAGATION SIMULATION OF URBAN MESH NETWORKS

Recently, there has been interest in developing protocols for urban mesh networks such as the ones to be deployed in Philadelphia [76], San Francisco [22], Taipei [96], Minneapolis [23], Anaheim, California, and Tempe, Arizona. Simulation is the common technique to determine the performance of these protocols. However, today, most simulations use the trivial disk propagation model (i.e., the signal propagates exactly R meters and no further) or a highly idealized propagation model such as the free-space or the two-ray model. Due to the presence of buildings, propagation in urban environments is far more complicated than the propagation presented by these simple models. Consequently, channel variability, which is a key aspect of wireless networking, is not well modeled in today's simulations. The result is that many insights gained from the free-space environment do not necessarily hold in the urban environment, casting doubt on the applicability of the conclusions drawn from today's simulations.

While the reasons that realistic propagation models are typically not included into network simulations is not well documented, it seems that the typical reasons include the belief that propagation modeling is computationally intractable and that propagation cannot be realistically modeled due to small-scale fading and delay spread. To the contrary, this thesis demonstrates the feasibility of including realistic propagation into network simulation. To this end, a simulator has been developed

in accordance with the guidelines set forth in this chapter and in [92] and [17]. This simulator is validated in three scenarios and it is shown that when used with packet level simulation, realistic propagation has a dramatic impact on simulation results. Furthermore, while realistic propagation is computationally intensive, the computational complexity of the approach taken here results in the propagation aspects of the typical simulation taking approximately as long as the processing of the discrete events by the packet simulator. Consequently, this thesis provides techniques and methods for greatly improving the quality of simulations.

In order to justify the simulation strategy used, an overview of the key features of propagation in an urban environment is presented. This non-mathematical discussion of propagation also provides networking researchers with an intuition of propagation in urban environments. For example, some of the myths that can be found in the networking literature are dispelled, and issues that are neglected in the networking literature, such as the importance of building materials, are discussed.

The remainder of this chapter proceeds as follows. In the next section, a brief discussion of urban mesh networks is provided. Section 2.1 discussed stochastic models of propagation. While stochastic models are often used in communication theory, their applicability to networking is limited, and, as will be discussed, they should be used with care. In contrast, the propagation simulator discussed here is deterministic. Section 2.2 discusses several characteristics of propagation. Specifically, the subsections of Section 2.2 discuss transmissions through walls, reflections off of walls, multipath fading, diffraction, scattering, time-varying channels, and delay spread. Section 2.3 discusses computational issues, discusses details of the UDelModels' propagation simulator, and presents validation of the simulator. The section 2.4 intends to demonstrate the impact of using realistic propagation on the simulation of mesh networks. Specifically, the difference between using simplistic models such as random way-point mobility model with open space propagation and

using the above realistic simulation models is examined. This section also highlights the difference in the performance of AODV [74] routing protocol under realistic simulation scenario and under simplistic simulation scenarios using open space propagation model and random way point mobility model. Finally section 2.5 discusses related work.

2.1 Stochastic Models of Propagation

One of the most important aspects of propagation is the channel gain (other characteristics are discussed below). The objective of propagation simulation is to determine the channel gain (and other channel characteristics). The physical layer model then uses these channel characteristics to determine the probability of transmission error. The relationship between the received signal power P_{received} and transmitter signal power $P_{\text{transmitted}}$ is $P_{\text{received}} = K \times H \times P_{\text{transmitter}}$, where K is a constant that depends on the wavelength and the antennas, and H is the channel gain. It is common for this relationship to be specified in dB¹ and dBm (dB milliwatts), i.e.,

$$P_{\text{received}}[\text{dBm}] = 10 \log_{10} K + H[\text{dB}] + P_{\text{transmitter}}[\text{dBm}].$$

To get an idea of common values, consider the typical 802.11b transmissions. In this case, the transmission power is 15 dBm and $10 \log_{10} K = -41$ dB. Furthermore, the sensitivity of commercially available 802.11b receiver cards is approximately -93 dBm when receiving at 1Mbps. Thus, if the channel gain is lower than -67 dB, then it is not possible to decode the transmission with *marginal reliability*. It is also common to require an extra *gain margin* to allow for the signal to be decoded reliably (not simply marginally reliability), to account for losses due to connectors and other analogue circuits, and to allow for overly optimistic sensitivity specifications. The gain margin depends on the transmitter/receiver manufacturers and system design,

¹ For the remainder of the chapter, we specify channel gains in dB.

and can range from 3 dB up to 9 dB, which require the 802.11b communication channel gains to be above -64 dB and -58 dB respectively.

The propagation of wireless signals can be accurately modeled by Maxwell's Equations. However, it is beyond today's computational abilities to simulate large urban regions using these equations. For this reason, there has been extensive research effort on developing techniques that provide computationally tractable estimates of the channel gain and other characteristics. In general, the techniques can be classified into deterministic and stochastic models. Furthermore, they can be divided into site specific approaches (i.e., detailed information about the environment is included into the propagation model) and non-site specific approaches. As discussed next, for simulation of urban mesh networks, deterministic site specific propagation models are appropriate.

We can easily rule-out deterministic non-site specific approaches such as the Okumura-Hata model [42] or COST 231 [34]. Such models provide a relationship between the channel gain and the distance between the receiver and transmitter. However, these models only provide an average behavior and do not model the variability of the channel.

On the other hand, non-site specific stochastic models can account for channel variability and hence are widely used in communication theory. However, there are many pitfalls to stochastic models which often make them not applicable to urban mesh network simulation. Consider the well established lognormal shadowing model [80]. This model specifies that the received signal power obeys

$$P_{\text{received}} = \frac{K}{d^\alpha} \times Y \times P_{\text{transmission}}, \quad (2.1)$$

where K is a constant, d is the distance between the transmitter and receiver, α is the attenuation factor, $P_{\text{transmission}}$ is the transmitted power, and Y is a lognormally distributed random variable that accounts for shadowing. This model dictates that if a transmitter and receiver are placed at randomly selected locations, then the

relationship between distance, transmitted power and received power obeys (2.1). While the model has been verified through experimentation, the model focuses only on a *single* link. On the other hand, in networking, the interest is on a large number of links, specifically, the focus is on the graph formed by links. If the relationship between received and transmitted signal strength merely obeys (2.1), then the graph formed by links would be a type of a random graph, whereas an urban setting induces a graph with deterministic components. Consider the left and right frames in Figure 2.1. The left-hand frame shows the channel gains (including correlations as described in [37] and [40]) when a lognormal shadow fading model is used, while the right-hand shows the channel gains in an urban area that is found through the techniques described in the next sections. While it is reasonable to assume that even in the urban setting, if two arbitrary locations are selected, then the signal strength will follow the lognormal distribution (i.e., (2.1) is valid), it is quite clear that channel gain is not random. Rather, the channel gain is greatly influenced by the map of the urban area. For example, the signal propagates far down the streets, while it propagates poorly into or through/over buildings. To see the implication of this, consider the least-hop path shown in Figure 2.2. In this example, only the first and last hop are able to propagate between indoor and outdoor, while other hops follow streets (outdoor to outdoor propagation). Thus, we see that the propagation results in routes in urban areas that are not random, but are dependant on the structure of the city. For this reason, we do not use non-site specific models, and focus on site-specific models where the layout of the city and buildings influences the propagation. Currently, the UDelModels' propagation simulator is deterministic. However, future work will include appropriate stochastic elements.

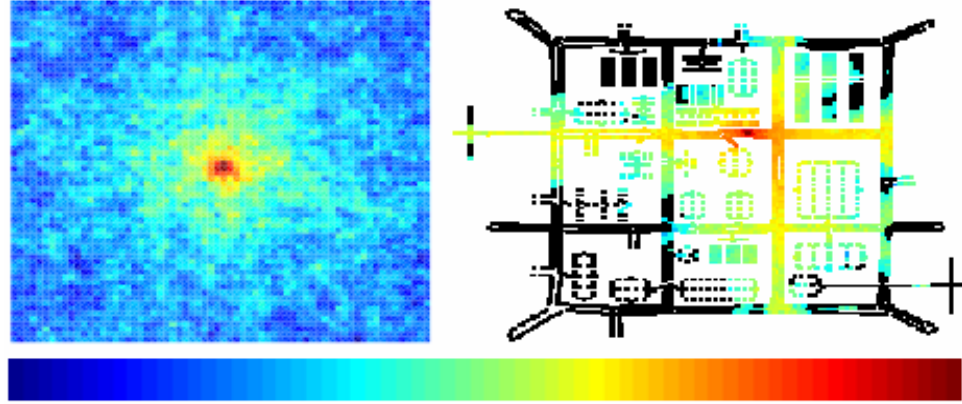


Figure 2.1: Left, channel gain generated by a correlated lognormal model. Right, channel gains generated by ray-tracing in an urban area. The rectangle shapes are buildings and the shaded boxes inside the rectangles represent channel gain to locations on the first floor of the building. In both the left and right frames, the star marks the location of the transmitter.




Figure 2.2: Routing in an Urban ad hoc Network. The source and destination are inside of two different buildings. The shortest hop path travels along the road, it does not follow the shortest geometric route (i.e., a straight line).

2.2 Characteristics of Urban Propagation

2.2.1 Overview

As mentioned above and depicted in Figure 1.1, there are several aspects of propagation that are important for UMNets. Specifically, the relevant types of propagation include line-of-sight, reflections off of buildings and the ground, transmissions through exterior walls of buildings, indoor propagation, and diffraction around and over buildings. While particular types of communication might be more impacted by particular types of propagation (e.g., indoor-to-indoor communication is mostly impacted by indoor propagation), these aspects of propagation can impact any type of communication (e.g., indoor signals may exit the building, experience a reflection, reenter the building, and impact the communication between two indoor nodes). In this section, these types of propagation are briefly discussed. The objective is to provide intuition of propagation in urban environments. While there are mathematical models of propagation, we do not provide any mathematical details. An excellent reference that contains mathematical details is [10].

Remark 1 *We typically focus on the propagation for a particular transmitter and receiver. However, due to reciprocity (e.g. \cite{reciprocity}), the channel between two antennas does not depend on which one is the transmitter and which is the receiver. To some degree, the symmetry of channels contradicts the observed asymmetry of links (e.g., [5] and [105]). It is important to note that the observed asymmetry of links is not due to the channel but due to difference in the analogue electronics in the transmitters and receivers. These differences may result in unequal transmit powers and unequal receiver attenuation. In some cases, these differences can be reduced by calibrating the transmit powers. It is also possible that one node experiences more interference than another node. In this case, the probability of error-free packet transmission depends on which node is the transmitter.*

2.2.2 Transmissions

In urban mesh networks, reflections off of objects such as buildings and the ground and transmission into buildings play a primary role in the magnitude of the received signal strength. Diffraction and scattering are also important, but from numerical experiments it can be shown that they play a secondary role. When the object that the signal is hitting is much larger than the wavelength, then the behavior of the signal can be accurately modeled as a ray that is partially reflected off of the object, partially transmitted through the object, and partially absorbed by the object.

If this ray model is accurate, then there are well known formulas that describe the behavior (e.g., see [10]). In general, the reflection can be written in terms of six components,

- $H_{\text{Transmit}}^{\perp}(\theta_{\text{Incidence}}, \varepsilon_r, \varepsilon_r'', w, f)$
- $H_{\text{Reflect}}^{\perp}(\theta_{\text{Incidence}}, \varepsilon_r, \varepsilon_r'', w, f)$
- $H_{\text{Absorb}}^{\perp}(\theta_{\text{Incidence}}, \varepsilon_r, \varepsilon_r'', w, f)$
- $H_{\text{Transmit}}^{\parallel}(\theta_{\text{Incidence}}, \varepsilon_r, \varepsilon_r'', w, f)$
- $H_{\text{Reflect}}^{\parallel}(\theta_{\text{Incidence}}, \varepsilon_r, \varepsilon_r'', w, f)$
- $H_{\text{Absorb}}^{\parallel}(\theta_{\text{Incidence}}, \varepsilon_r, \varepsilon_r'', w, f)$

where $\theta_{\text{Incidence}}$ is the angle between the ray and the material, ε_r' and ε_r'' are the real and imaginary part of the relative dielectric constant of the material, w is the width of the material, f is the frequency of the signal, and the superscript denotes the polarization.

Instead of detailing how each of these factors impacts the propagation, we will consider an example. We examine the received signal strength on one side of a wall

with a transmitter on the other side as shown in left-hand frame of Figure 2.3. We consider the received signal strength in two locations, specifically, one at the same height as the transmitter and one elevated from the transmitter. We assume that the signal strength decays according to free-space propagation and then intersects the wall and suffers insertion loss $H_{\text{Transmit}}^{\parallel}$ (where we assume that the wireless signal is vertically polarized, hence the \parallel superscript). We take the frequency to be 2.4 GHz, as in 802.11b/g. Three types of wall materials are examined, namely, brick ($\epsilon'_r = 3.9$, $\epsilon''_r = 0.001$, $w = 10$ cm.), concrete ($\epsilon'_r = 5$, $\epsilon''_r = 0.2$, $w = 20$ cm.), and glass ($\epsilon'_r = 5.8$, $\epsilon''_r = 0.003$, $w = 2.5$ cm.). The behavior for these materials is shown in the right three frames of Figure 2.3.

Many aspects of transmission through material can be observed in Figure 2.3. First, notice how in the case where the receiver and transmitter are at the same height, the channel gain increases as the distance between the transmitter and receiver decreases. This agrees with the intuition from free-space propagation. The impact of the material can be seen by noticing the variation in the received signal strength for different materials. For reference purposes, consider that without the wall, when the transmitter is about 5 m from the wall, the channel gain is -14 dB from the wall and a receiver that is at the same height as the transmitter. Thus, when the angle of incidence is 0, the wall reduces the channel gain by 0 dB (for glass) to 6 dB (for concrete). The impact of the angle of incidence can be emphasized by considering the channel gain when the transmitter and receiver are not at the same height. As can be seen, this results in the signal strength first increasing and then decreasing as the distance between the transmitter and receiver increases; this is in conflict with the behavior in free-space. The key characteristic of propagation at work here is that when the ray hits the wall at a grazing angle, only a small amount of signal power is transmitted through the wall. The point where the signal strength switches from increasing with the transmitter-receiver distance to decreasing with




Figure 2.3: Transmission through a wall. The left-hand frame shows a side-view of the simulated scenario. The right-hand frames depict the attenuation of the signal when transmitted through walls of different kinds of materials. It also depicts the nature of attenuation due to relative heights of the transmitter and the receiver.

transmitter-receiver distance depends on the material, the thickness of the material, and the height of the receiver.

Remark 2 *One interesting implication of Figure 2.3 is that if an urban mesh network uses relays mounted on lampposts, then the lamppost are quite close to buildings. Thus, relays mounted on a lamppost on the far side of the street will provide better coverage of a building than a relay mounted on a lamppost on the near side of the street. We concluded that, in some cases, the angle of incidence can be more important than the transmitter-receiver distance.*

Figure 2.3 shows variation of the channel gain due to the different building materials. The dependence on material poses a serious challenge for simulation. First, walls are typically not made of a single material, but of layers of material. It is, however, possible to extend the method to accommodate layers of different material (e.g., see [88]) at the expense of computational complexity. A second and more serious problem is that it is not realistic to know the construction material for each building. This dependence on construction materials emphasizes the difficulty

of accurate prediction of channel gains and demonstrates why realistic simulation is more reasonable than accurate simulation. In the UDelModels, it is possible to specify different building materials, but it is assumed that the material is homogeneous for each building.

2.2.3 Reflections

Reflections are complementary to transmissions. However, it is important to note that some power is absorbed by the material, hence, the power transmitted through a wall and the power reflected off of a wall do not necessarily sum to the power of the incident signal and, depending on the material and angle of incidence, a significant amount of energy can be absorbed.

To get an understanding of the impact of reflections, we consider propagation down a building lined street, as shown in Figure 2.4. Here the signal is repeatedly reflected off of walls and, in a sense, focused down the street. This effect is sometimes referred to as propagation down an urban canyon. A similar effect can also arise in hallways and tunnels. The result is that the received signal strength is stronger when propagating down a street than it would be with the same transmitter-receiver distance but in free-space.

Figure 2.4 shows the channel gain down an urban canyon for the same set of materials considered above and the free-space approximation. Note that free-space predicts substantially smaller received signal strength than the simulated received signal. For example, when the receiver-transmitter distance is 300 m., the difference between the free-space and the model that accounts for reflection ranges from 13 dB to 5 dB, depending on the material and the distance between the walls. Figure 2.4 also shows approximations of the channel gain given by $1/d^\alpha$, where, for a canyon width of 7 m., $\alpha = 1.38$ and a canyon width of 35 m., $\alpha = 1.47$. It should be pointed out that it is often assumed that $\alpha \in [2, 4]$, however, as this simple example

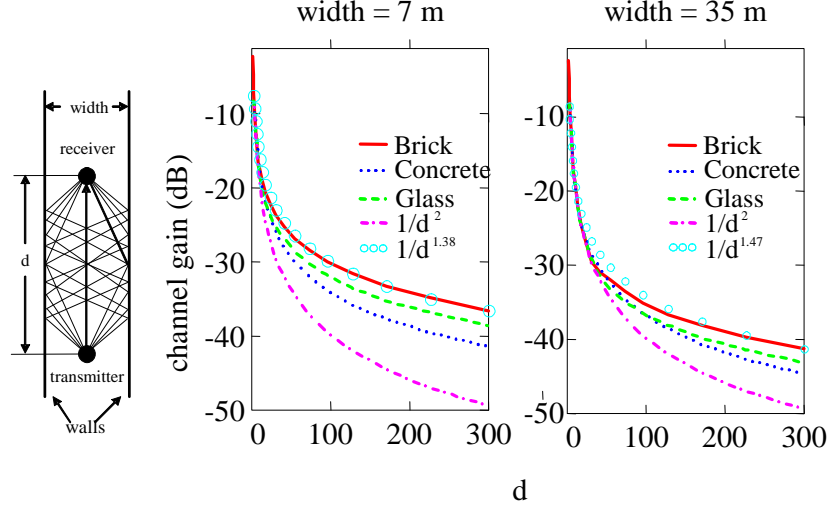


Figure 2.4: Propagation down a street. The left-hand frame shows a top view of the simulated scenario. The right-hand frames depict the effect of propagation down an urban canyon for different wall materials, different attenuation exponents, and for the free space.

demonstrates, it is possible that for some paths we have $\alpha < 2$. That is, while buildings may block communication, they may also enhance communication.

Note the lack of smoothness when the walls are made of brick and the width is 35 m. (e.g. $d = 50$). This is due to the wide variation in the reflection coefficient as the angle of incidence varies. Indeed, depending on the material and the width of the material, there may be some angles of incidence where no signal is reflected. Such angles are called *Brewster angles*. In some settings, the wide variation of the reflected signals strength as a function of incident angle results in large fluctuations in the total received signal strength for small movements of the receiver or transmitter antenna. This amounts to further sensitivity of communication on building materials.

2.2.4 Multipath Fading

It is well known that wireless signals can experience small-scale fading or multipath fading. Such fading results from the constructive and destructive interference of signals that follow different paths from the transmitter to the receiver. As a result, when the wavelength is small, a small displacement of the transmitter or the receiver will cause a change in the interference and hence a change in the received signal strength. While fast-fading is a significant problem in many wireless systems, it is typically not relevant in wide bandwidth communications such as those often used for data communications (e.g., 802.11). The reason is that the received signal strength is essentially averaged over the bandwidth, which, by definition, is wide. The left-hand frame in Figure 2.5 shows an example of the signal strength at various frequencies. Note that at some frequencies, the signal strength is quite low, even 40 dB less than the mean signal strength of 0 dB. A narrow bandwidth communication will experience such degradations in signal strength. However, when averaged over a sufficiently wide bandwidth, the average signal does not experience degradation. For example, the middle frame in Figure 2.5 shows how the signal strength varies for small changes in position. Note that the narrow bandwidth signal experiences rapid and large fluctuations, while the wide bandwidth case experiences much smaller variations.

In general, the wider the bandwidth, the less susceptible to fast fading. However, variation in signal strength also depends on the environment. The right-hand frame of Figure 2.5 shows the variance of the received signal strength over a circle with one meter radius as a function of the bandwidth. Each curve is for a slightly different environment, but in all cases it is for propagating down the urban canyon as shown in the left-hand frame of Figure 2.4. Note that in general, the variance decreases as the bandwidth increases. However, the variance of the narrow bandwidth case and the rate that the variation decreases with the bandwidth depends on the




Figure 2.5: Fast fading in narrow band and wide band communication scenarios. Left: the signal strength as a function of frequency for a narrow band signal propagating down an urban canyon. Middle: signal strength of a narrow and wide band signal as a function of receiver position perturbation when propagating down an urban canyon. Right: variance in signal strength due to change in receiver position as a function of bandwidth for when propagating down a urban canyon as shown in Figure 2.4.

environment. In all cases, the variance is quite small when the bandwidth exceeds 10 MHz, while 802.11b has a bandwidth of 22 MHz. It should be noted that there are other factors besides multipath reflections that could cause large changes in signal strength over small changes in position, e.g. changes in antenna orientation.

2.2.5 Diffraction

While reflections and transmission play an important role in propagation, diffraction is also significant and should not be neglected. Diffraction allows the signal to "bend" around corners and over/around buildings. However, the diffracted signal will not be as strong as the original signal; the sharper the bend, the weaker the signal. Consider the examples in Figure 2.6, which use the Uniform Diffraction Theory [59], [7]. The left-hand frames show the diffraction around a corner. The second from the left frame compares free-space propagation around the corner to free-space propagation with the added loss due to diffraction. The right-hand frames




Figure 2.6: Diffraction around a corner (left) and over a building (right). The left-hand frame and second from the right frame show side views of the simulated scenario. The other frames show the attenuation as a function of h . As h increases, the signal strength decreases. The dashed lines show the signal strength if the loss was only due to free space propagation while the solid line shows the signal strength when the loss due to diffraction is also included.

show diffraction over a building where two diffractions are required. In both cases, as h increases, the signal must make a sharper bend (i.e., the diffraction angle increases) and, hence, more loss is incurred. It is important to note that the signal strength decreases quickly as h increases. From the right-hand frames in Figure 2.6, it can be seen that diffracting over a building that is only 5 m. higher than the transmitter and receiver results in a channel gain that is too small for typical 802.11 communication (recall the discussion in Section 2.1), while diffracting around a single corner (left hand plots of Figure 2.6) might result in a sufficiently large channel gains for most 802.11 communications. Indeed, our simulations indicate that for 802.11, diffracting around two corners results in loss that is typically too great for 802.11.

While diffracting around two or more corners leads to large losses, 802.11 might be able to communicate effectively around a single corner. Consequently, a significant portion of the coverage area of 802.11 is due to diffraction. Using the

Table 2.1: Coverage and computation time

Experiment	Coverage	Time (sec./transmit location)
Line of sight	937	56
1 iteration	2623	59
1 iteration without diffraction	1960	56
2 iterations	3927	62
2 iterations without diffraction	2616	57
3 iterations	4243	67
3 iterations without diffraction	2862	58
4 iterations	4265	85
4 iterations without diffraction	3065	63
All iterations with diffraction	4265	122

UDelModels' propagation simulator, it is possible to quantify the impact of diffraction. Table 2.1 shows the number of locations to which a particular transmitter is able to communicate when only line-of-sight is used, when line-of-sight and reflections/transmissions are used, and when all the components are used, namely, line-of-sight, reflections/transmissions and diffraction. The table also shows the difference in the coverage when different numbers of iterations are used. By the number of iterations, we mean the maximum number of reflections, transmissions, or diffractions that each ray may experience. We see that about 22% of the locations are missed if diffraction is not used. However, if only line-of-sight is considered, then 78% of the locations are missed. For these calculations, a map of the Paddington area of London was used and the transmitter and receivers were located 1.5 m. above the ground or above the floor. If the transmitter was higher (e.g., a mobile phone base station), or transmissions with smaller channel gains can be received, then it is possible that diffraction could play a larger role than is illustrated in this example.

2.2.6 Scattering

Scattering is often used to refer to the impact of smaller unmodeled objects on the propagation, e.g., lampposts, trees, vehicles, people, and office furniture. Scattering also accounts for the unevenness of building walls (e.g., windows, doors, or facades). Without detailed knowledge of the location and dimension of small objects and without details of building walls, it is difficult to include the effects of these types of scatterers into propagation simulation. Furthermore, the inclusion of such details greatly increases the computational complexity of propagation simulation. As a result, relatively little work has focused on including scattering into propagation prediction. One exception is [84] where the impact of scatterers was examined. It was found that in dense urban environments, such scattering does not dominate. More specifically, at short distances, scatterers have little impact, but may have some impact at greater distances (e.g., in some cases the difference between prediction and measurements decreased from 6 dB to 3 dB when scatterers were included). On the other hand, in areas where building density is low (e.g., suburban areas), scatterers such as lampposts, trees, and vehicles may dominate the propagation. This is one reason why this thesis focuses on the urban setting.

Scattering can also occur when a wireless signal propagates through vegetation. Indeed, if the vegetation is large and dense, the scattered signal dominates over the direct, non-scattered signal. Even when the vegetation is thin, as it usually is in urban areas, the vegetation can cause loss. At 2.4 GHz, vegetation causes a loss of approximately 0.2 dB per meter of vegetation [100]. It is also common to model diffraction over vegetation [60].

2.2.7 Time-Varying Channel Gain

While the variability of channel gain is greatly affected by the movement of the transmitter or receiver, the channel gain can also change when the transmitter and receiver are fixed, but objects in the environment move. The right-hand frame of

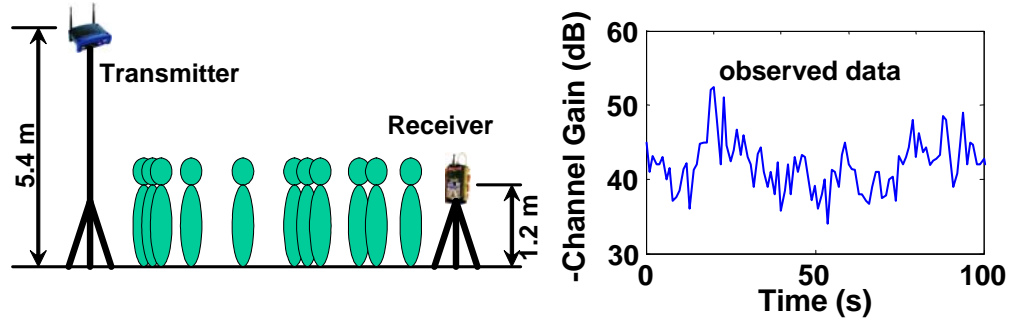


Figure 2.7: Left, experimental setup for measuring the time varying nature of propagation in a busy street in Philadelphia. Right, an observed time-series of the channel gain.

Figure 2.7 shows the channel gain as a function of time for a receiver and transmitter as observed along a sidewalk during rush-hour in the city center of Philadelphia. As shown in the left-hand frame of Figure 2.7, the transmitter was 5.4 m. above the sidewalk, while the receiver was only 1.2 m. above the sidewalk. The transmitter was a Linksys BEFW11S4 and the receiver was the Yellowjacket [2] receiver made by Berkeley Varitronics Systems. While time-varying propagation is difficult to model, in [94], a diffusion-based model was developed that accurately models the variability shown in Figure 2.7. Chapter 4 deals with this model in more detail.

2.2.8 Delay Spread

As noted in sections 2.2.2, 2.2.3 and 2.2.5, the wireless signal may follow several different paths from the transmitter to the receiver. While one impact of multiple paths is that the signal may experience multipath fading, another result is that the multiple copies of the wireless transmission will be received at different points of time. Essentially, these multiple copies will interfere with each other. When this self-interference is considered as noise, it is clear that the effective SNR is decreased by the presence of these delayed copies. Since the traditional calculation

of SNR considers delayed copies of the transmission as useful signal power, the traditional SNR might not be a good indicator of probability of transmission error.

A channel where multiple copies of the signal arrive at different times is said to have delay-spread. While there is no ideal metric of delay-spread, two common measures are the mean delay spread and the RMS delay spread (see [80] page 199 for definitions). To explore the impact of delay-spread, consider Figure 2.8, which shows the relationship between SNR and bit-error probability for different amounts of delay spread for a *simple* 802.11b receiver and for an 802.11a receiver. Note that the bit-error probability of 802.11b is greatly affected by the delay-spread. Also, it can be seen that for an RMS delay-spread of above 60 ns., which is quite common in an urban setting, the impact of the delay is quite severe. From measurements, it has been found that indoors, the RMS delay-spread is typically less than 50 ns., but outdoors, it can exceed 500 ns. The dependence of delay-spread on the communication environment has been widely observed (e.g., [5] and [95]). It should be noted that the delay-spread typically increases with the distance between the transmitter and receiver [38]. And therefore, roof-top configurations such as those studied in [5], which are able to propagate considerable distances due to the lack of obstructions at higher heights, are likely to experience large delay-spread values. As pointed out in [5], to predict the transmission error probability, in some cases, both delay-spread and SNR are required.

Incorporating the effects of delay-spread into wireless simulation poses significant difficulties. First, the mapping from RMS or mean delay-spread to transmission error probability is not well defined. It is possible to have a channel with high RMS delay-spread that yields lower transmission error probability than the one that has a lower RMS delay-spread and conversely. Second, RAKE receivers can be incorporated into receivers that can greatly mitigate the effects of delay-spread. However,

the implementation of the RAKE receiver has a significant impact on the performance. For example, [25] shows that high performance RAKE receivers can mitigate the impact of delay-spread, while less effective RAKE receivers will perform poorly when the delay-spread is large. Thus, it is difficult to make any conclusions about the effect of delay-spread on *future* mesh networks based on *current* receivers.

In anticipation of receivers more suitable for outdoor use, another alternative is to neglect delay-spread. Not only is such an approach reasonable when the inclusion of algorithms that mitigate the effect of delay-spread into 802.11b receivers is anticipated, but also because other physical layers based on OFDM greatly mitigate the impact of delay-spread. Consider Figure 2.8, which shows the relationship between SNR and bit-error probability for different values of RMS delay-spread when 802.11a is used. Note that the impact of delay-spread is not discernible. Schemes such as 802.11g and 802.16 use OFDM and hence are also relatively immune to delay-spread. In summary, while delay-spread can have a significant impact on transmission error, there are a number of physical layer schemes that can mitigate these effects, and, considering the rapid migration between physical layers (e.g., from 802.11b to 802.11g), it seems likely that when there is a need (i.e., when outdoor mesh networks are finally deployed) such schemes will become wide-spread. Therefore, depending on the assumptions made about the network being simulated, realistic simulation of future urban mesh networks does not necessarily require consideration of delay-spread.

2.3 Computational Techniques

2.3.1 Overview

While much is known about propagation, a major obstacle to simulation of propagation is the computational complexity. This complexity forces a trade-off between computational complexity and accuracy. The difficulty to specify the details of the environment (e.g., wall materials and exact dimensions and structure




Figure 2.8: Bit-error probabilities as a function of delay-spread and SNR. Simulation of 802.11b at 1 Mbps (without a RAKE receiver) and 802.11a at 6 Mbps. The different curves for 802.11b correspond to different RMS-delay spreads. For a delay-spread of 7ns to 220 ns, 802.11a gave the same relationship between SNR and bit-error rate, hence this relationship is shown with a single curve.

of buildings) also forces a trade-off with accuracy. It is important to note that the propagation must be determined from every location that a transmitter/receiver could be located to every other possible location that a transmitter/receiver could be located. The simulation of a $1\text{km} \times 1\text{km}$ urban region often requires hundreds of thousands of locations. Hence, the computation to determine the signal strength from a single location, which in itself is computationally difficult, must be repeated hundred of thousands of times. In this section, techniques to efficiently simulate propagation are discussed. The techniques described are part of the UDeIModels simulation package. In the next subsection, outdoor propagation is discussed, and is followed by a discussion on indoor propagation simulation.

2.3.2 Outdoor Propagation

Once the map of buildings, bandwidth, and building materials have been defined, propagation can be estimated. As noted above, extreme care must be taken

to reduce the computation. A significant computational savings can be achieved if it is assumed that all walls are vertical. In this case, the 3-D ray tracing problem reduces to a two stage problem, each of which are much smaller than the 3-D problem. The first stage is illustrated in the left-hand plot in Figure 2.9, where two vertical plane paths are shown (the path that diffracts around the smaller building is not shown). The first stage is a 2-D raytracing. Once the vertical plane paths are found, the 3-D ray paths restricted to the vertical plane paths can easily be computed. The right-hand frame in Figure 2.9 shows the paths of rays in the vertical planes. One vertical plane path has three ray paths, one that diffracts over a building, one that reflects off of the ground and passes through a building, and one that passes straight through a building. The other vertical plane path has two ray paths, one reflecting off of the wall of a building and one reflecting off of the wall of building and undergoing a ground reflection. While Figure 2.9 shows each vertical plane paths having two or three ray paths, in general, in one vertical plane path there are many ray paths that may include reflection off of the ground, and transmissions through buildings or diffractions over buildings. The UDelModels' propagation simulator only considers three types of ray paths, direct paths (no ground reflections and no diffractions over buildings), ground reflected paths with no diffraction, and paths that include no transmissions through buildings but rather diffract over buildings. Note that diffractions and transmissions through buildings result in significant loss, and hence neglecting ray paths with multiple diffractions or diffractions and transmissions into buildings has little impact on accuracy. This method of 3-D raytracing is known as vertical plane launching [62], and is known to yield accurate propagation estimates in real cities, where all walls are not necessarily vertical [82].

A straightforward implementation of even 2-D raytracing is computationally difficult. Instead, a technique that is more appropriately called beam tracing can




Figure 2.9: Left: a top-view of the scene on the right. The lines shown are found from 2-D ray-tracing. Right: Two vertical plane paths (found from 2-D ray-tracing) and 5 ray paths.

be performed. Like ray tracing, the goal of beam tracing is to determine the paths from the transmitter to receiver. Beam tracing begins with the source broadcasting the signal in all directions. This transmission is not modeled as a large number of rays, but as a small number of beams. When a beam intersects a building, two beams are generated, one is reflected off of the building and one is transmitted into the building. If only a part of the beam intersects the building, then multiple beams are generated, in particular, one that reflects off of the building, one that transmits into the building, one that passes by the building, as well as several beams that account for the diffracted signal. Each diffracted beam models a particular range of angles of diffraction. The strength of the diffracted beam is the average strength for the diffraction angles that the beam models. Note that smaller the range of angles for each beam, the more accurate the diffraction model is. But this generates more beams, increasing the computational complexity of the process. Finally, if the receiver is found to be included within the span of a beam, the vertical plane path is found and the 3-D rays from transmitter to receiver are determined. Once a 3-D ray is found, the *3-D ray information*, specifically, signal strength, phase, delay, angle of departure from the transmitter, and angle of arrival at the receiver are recorded.

The beam tracing computation can be further simplified by dividing the 2-D

space into a grid and the determining the propagation between the center points of each square. Each square of the grid is called a *floor-tile*. Outdoors and indoors are discretized in this manner. To reduce the number of floor tiles, the entire space is not discretized. Rather, floor-tiles are placed only along the center-lines of sidewalks, roads, and hallways. For rooms, floor tiles are placed in every location that a mobile node can be located. Similarly, the walls of buildings are also divided into *wall-tiles*. Since the beam tracing is in 2-D, the wall-tiles are segments (1-D tiles). The smaller the size of these tiles (floor-tiles and wall-tiles) more accurate the simulation results, but also the more computationally expensive.

The computation is divided into two parts, preprocessing and beam tracing. During preprocessing, *ray neighbors* for each tile are found. A tile's ray neighbors are all the tiles that could be directly reached (i.e., without reflection, transmission through a wall, or diffraction) by a ray emanating from the tile. Once the ray neighbors are found, beam tracing can be performed efficiently. Figure 2.10 illustrates how the beam tracing with wall tiles is performed.

The process of beam tracing Figure 2.10 is carried out in a breadth first manner² with each beam continued to be reflected, transmitted until either the beam exits the modeled area or until the estimated channel gain of the beam falls below a threshold. Each time a beam is generated, the floor-tiles it covers are determined and the 3-D ray information is recorded. Note that the 3-D ray information allows one to interpolate the signal strength between grid points as well as determine the delay-spread, the Doppler spread (once velocity is known), as well as simulate directional antennas. It is possible to determine the receiver power over a particular bandwidth by integrating the signal strengths in a way that accounts for phase (i.e., accounting for constructive/destructive interference). The 3-D ray information

² Here, breadth first means that all the k th reflections are found before the $(k + 1)$ th reflections are found.

requires a significant amount of information be saved. A considerable amount of savings is achieved by simply recording the received power over a bandwidth and RMS delay-spread (which is merely two integers). Of course, this approach cannot be used if directional antennas are simulated or Doppler spread included.

Another important computational saving is achieved by recognizing that the communication is not possible between most locations, hence the propagation matrix is sparse. Computational savings are possible by using techniques to store sparse matrices. However, since the resulting data is large, and since simulation requires the determination of a large number of channel gains, it is important that the sparse matrix techniques support fast access.

2.3.3 Indoor Propagation

Beam tracing can be performed indoors as well as outdoors. However, the computational complexity depends on the number of walls. Since building interiors have a large number of walls, beam tracing inside all the buildings within a large region of a city exceeds today's computational abilities. Fortunately, it has been found that a realistic estimate of indoor propagation can be performed without using beam tracing. Specifically, the *attenuation factor* (AF) model has been shown to provide realistic channel gain estimates, with the error found to be within 4dB [80]. The AF model assumes that communication indoors takes a straight line path (i.e., no reflections off of interior walls). Furthermore, transmissions through each interior wall and transmissions through each floor result in attenuation. While the amount of attenuation depends on the building, a value of 4dB per wall (for an office building) has been shown to work well [80]. Realistic attenuation for passing through 1 floor is 30dB, passing through 2 floors is 35dB, passing through 3 floors is 39dB and passing through 4 or more floors is 40dB [80]. In summary, outdoors, rays make reflections off of buildings, diffractions over and around buildings, and transmissions into buildings. Once inside a building, the ray will continue in the




Figure 2.10: Beam Tracing. Suppose a beam strikes the wall tile marked A. The reflection of this beam is modeled as a beam emanating from the virtual source as shown. This beam is between and is defined by the two rays marked 1 and 2. These rays are used to select the floor tile ray-neighbors that this beam strikes (since ray-neighbors are precomputed, this selection is computationally efficient). Similarly, the rays 1 and 2 are used to select the set of A's wall tile ray-neighbors that the beam strikes. One such wall tile is shown in the upper right and is marked B. The beam that strikes tile B is defined by the rays marked 3 and 4. These rays define the beam that is transmitted through tile B and can be used to find the beam that is reflected off of B. This process of reflection and transmission is repeated until the beam's strength is small enough that it can be neglected (e.g., -110 dB).

same direction, experiencing further attenuation for any interior wall or floor that it passes through. When a ray strikes an exterior wall from the inside, it is both reflected back inside and transmitted outside.

2.3.4 Computational Complexity

Determining the propagation matrix of a region of an urban area is a feasible but highly computationally complex task. The complexity is both in terms of memory usage and processing time. Processing times for a $1km \times 1km$ urban region is often on the order of tens of processor days. But the process is highly parallelizable and nearly scales with the number of processors used (i.e., 75 processor days takes about 5 days on 15 processors). Of course, the entire channel gain matrix for each city only needs to be found once. Table 2.2 outlines the memory requirements and the processing time for two cities. These two cities represent two ends of the spectrum. Paddington is a dense city with relatively small buildings, whereas the University of Delaware campus has much more open space and larger buildings. Note that the campus map is significantly larger than Paddington, but only has slightly more buildings.

The memory requirement is dominated by the lists of ray neighbors. In Table 2.2, the wall tiles were 2 meters long and floor tiles were $1m \times 1m$. While the number of wall and floor tiles scales linearly with the reciprocal of the size of the tiles (i.e., if the wall tiles are twice as long, there are half as many)³, the total ray neighbors scale quadratically. For the simulations shown in Table 2.2, there are between 20,000 to 40,000 exterior floor tiles, and 80,000 to 100,000 interior floor tiles. As expected, sidewalks and roads utilize little area outdoors, but indoors, hallways and offices fill large areas.

³ Recall that the floor tiles cover sidewalks, roads, hallways, and offices only.

Table 2.2: Computational complexity

City	Size (m.)	Num buildings	Num floor tiles	Num wall tiles	Total num. ray neighbors
Paddington	857×836	131	108K	17K	15M
UD Campus	1768×1597	151	147K	19K	67M
City	Size (m.)	Total num. ray neighbors	Ave. num reflections per source	Ave. num vertical plane paths/source	Total comp. time (processor days)
Paddington	857×836	194K	367K	1.3M	53
UD Campus	1768×1597	84K	332K	946K	72

As shown in Table 2.2, there are on the order of 10 to 100 million ray neighbors for a city of size about 1km×1km. Since each list entry requires 40B, the memory requirement approaches 4GB. Assuming sufficient memory resources, the computation time for the preprocessing stage is relatively short; the cities shown in Table 2.2 took a single day on an AMD Athlon 64 FX 55 with 8GB RAM.

Once the ray neighbors are found, the propagation characteristics between each pair of floor tiles can be found. From a single source, the propagation characteristics to all destinations can be found at the same time. Table 2.2 shows that each source produces vertical planes that make approximately 500,000 reflections, diffractions, or transmissions. These reflections, diffractions, or transmissions are shared among all destinations. However, for each destination, all the vertical plane paths between the source floor tile and destination floor tile are found. Hence, the total number of vertical plane paths found greatly exceeds the total number of reflections, diffractions, and transmissions.

For each source it is necessary to find the all reflections. Thus, in the campus map, the total number of reflections found was around 263 billion and 12 billion for Paddington. Several optimizations and tricks to efficiently distribute the program results in the total processing time shown.

2.3.5 Computational complexity for packet simulation

As discussed in Section 2.3.4, determining the propagation matrix is computationally intensive, however, it only must be done once for each city. Once the propagation matrix is found and a mobility model determines the locations of nodes, it is necessary to select the entries from the propagation matrix corresponding to each pair of node locations. This channel information is saved into a *propagation trace file*, which the protocol simulator reads during packet simulation. Figure 2.11 shows the execution time, which includes the time to construct the propagation trace file and the time to run the packet simulation for a one minute of simulated time. Figure 2.11 demonstrates that realistic propagation does increase the execution time. This increase is less than a factor of two for high node densities, but is quite large for low node densities. On the other hand, the right-hand frame in Figure 2.11 shows that for 250 mobile nodes, the packet simulation time makes up the majority of the execution time, and at all node densities, the time to construct the propagation trace file is similar to the packet simulation time. Furthermore, since the time to perform a packet simulation depends on the number of discrete events that must be processed, the propagation environment will have a significant impact on the packet simulation time. Hence, comparing packet simulation times of open-space propagation and realistic propagation is difficult.

While Figure 2.11 shows that the time it takes to perform packet simulation is similar to the time it takes to generate the propagation trace file, it is possible to generate the propagation trace file once and reuse the same trace file for multiple simulations. In fact, a large set of propagation trace files are available online [17]. If this approach is used, then the only computation related to realistic propagation is loading the propagation trace file. However, since the trace file specifies the propagation between all pairs of nodes, if a large number of nodes are simulated,

then the trace file can be quite large. On the other hand, the trace file only specifies the propagation between node pairs when the propagation changes; hence, if nodes move slowly, then the propagation trace file is quite small. Moreover, the trace file only provides the propagation once every second⁴. Therefore, if high data rates are achieved, then the computational complexity of processing packet transmission/reception events is higher than the computational complexity of loading the propagation matrix. Thus, when the propagation trace file is precomputed, the impact of including realistic propagation might be relatively small. In the right-hand frame of Figure 2.11, it can be seen that the simulation time grows slowly with the number of nodes. More specifically, in the worst-case, the propagation trace file grows quadratically with the number of nodes. However, Figure 2.11 shows that the simulation time does not grow quadratically except for 1500 nodes, which shows a large jump in simulation time. Interestingly, at this node density, the simulation with realistic propagation takes approximately as long as the simulations with open-space propagation.

2.3.6 Validation

The goal of the propagation simulation for simulating mobile wireless networks is not to predict the signal strength, but to produce signal strength that behaves in a realistic fashion. Nonetheless, it is useful to understand the accuracy of the propagation simulation. To this end, three experiments were performed, two outside and one inside. In all cases, an 802.11b access point was placed on a 1.5 meter tripod and the Berkeley Varitronics Yellowjacket wireless receiver [2] was placed on a second 1.5 meter tripod. The access point was placed at a fixed location and the average receiver signal strength was determined by receiving 600 802.11b beacons. Figure 2.12 (b) shows a part of the University of Delaware campus and

⁴ Currently, linear interpolation is used to determine the propagation between updates.

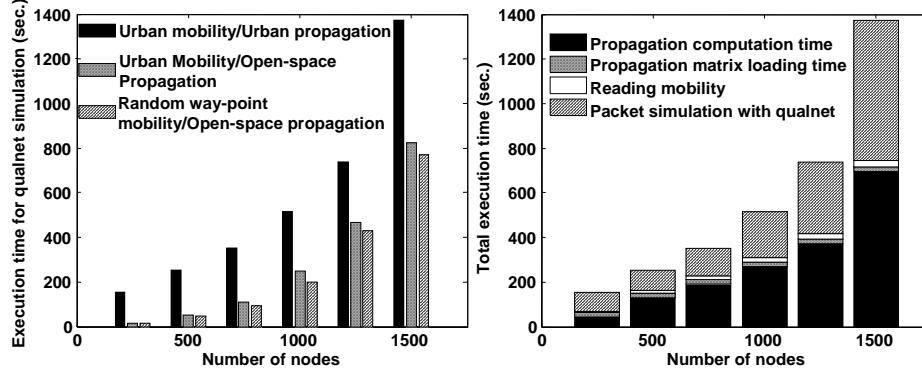


Figure 2.11: Simulation Execution Time. The left-hand frame shows the simulation execution time for three simulation techniques. The right-hand plot shows a break-down of the execution time when realistic propagation is included. The propagation computation time includes the time to identify which floor tile a node is nearest and the time to look-up the propagation between nodes from the propagation matrix.

Figure 2.13 (b) shows a street intersection in Philadelphia. In Figure 2.12 (b), the buildings were approximately 14 meters high while in Figure 2.13 (b) the buildings were at least 40 meters high. In both cases, the X-mark denotes the location of the transmitter while the receiver is moved along the indicated path. Figure 2.12 (a) shows the observed and modeled channel gain corresponding to the path starting from the transmission point and moving along the path in the counter-clockwise direction. Figure 2.13 (a) shows the model and observed channel gain starting near to the transmitter, moving down and then turning the corner.

Figures 2.12 (a) and 2.13 (a) show that the model and observations match well both qualitatively and quantitatively (within 5dB in most locations). To gain more insight into propagation modeling we examine the propagation prediction quality at different locations, especially the location where the prediction quality is low. In the area between B and C, there is an unmodeled archway that is indicated in Figure 2.12 (b). Similarly at location F, there is a bridge as depicted with the rectangle.

Ignoring these objects impacts the accuracy of the propagation prediction. In the locations marked E and G, there are several moderate sized unmodeled objects (large air conditioners at E and trees at G) that partially blocked the signal. As mentioned above, sometimes such objects are called scatterers. We see that scatterers can slightly decrease the received signal strength. On the other hand, in the areas where there is purely line-of-sight (marked As), line-of-sight with reflections (marked Bs) and reflections with diffraction (marked D), there is good agreement between the model and the observations. See [85] for further discussion on the relationship between accuracy of propagation prediction and accuracy of the map. Furthermore, it is important to note that due to mobility of objects, propagation measurements made at different times will be different. For example, in [83] it was found that identical measurements made at different times can differ by 5 dB.

Figure 2.13 (a) also shows a good fit. Again, the influence of scatters can be observed. In this case the scatterers includes things such as mailboxes, parked cars, and irregularity of the walls (e.g., doors that are set back from the wall). Nonetheless, the model and observations are within a few dB. Note that the signal strength decreases quite slowly as the distance from the corner increases. This is due to the urban canyon effect discussed in Section 2.2.3 and depicted in Figure 2.4.

Finally, Figure 2.14 (b) shows the layout of a building interior and Figure 2.14 (a) compares the modeled and observed signal strength for the points indicated in Figure 2.14 (b). Again, we see that there is reasonable good agreement between the model and the observations.

In summary, it is clear that accurate propagation prediction requires more detailed knowledge of the environment. However, coarse knowledge (e.g., location of buildings alone) provides realistic propagation, both qualitatively and quantitatively.

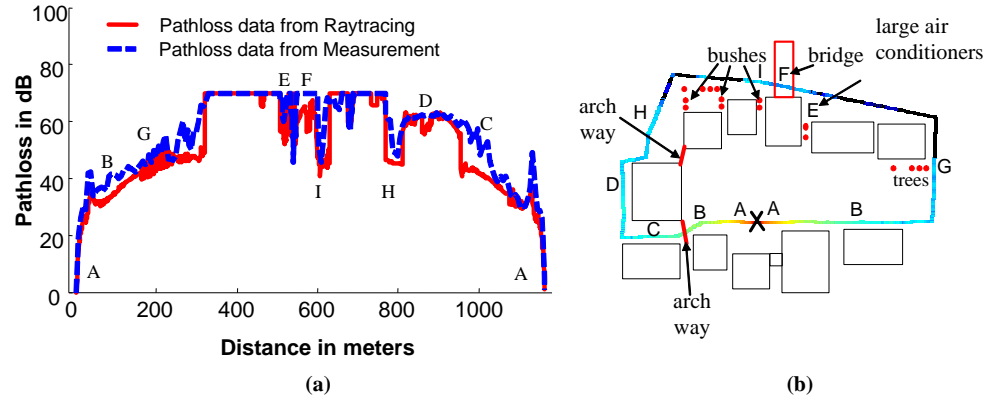


Figure 2.12: Observed and estimated path loss in the campus environment.

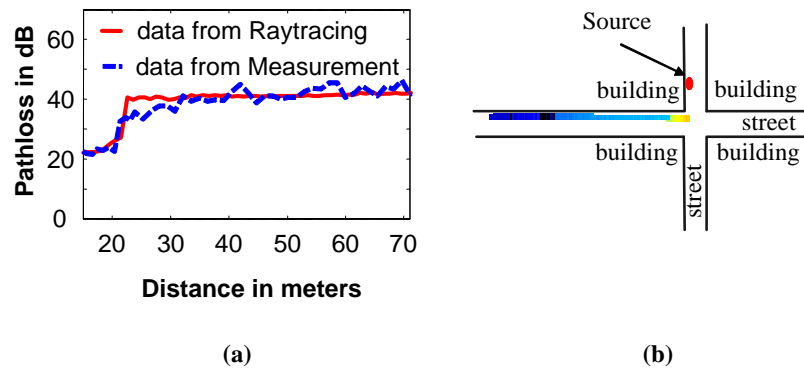


Figure 2.13: Observed and estimated path loss at an intersection in Philadelphia.




Figure 2.14: Observed and estimated path loss in an indoor environment. The measurement points shown on the x-axis correspond to the numbered locations in the map. The map is of the third floor of Evans Hall in University of Delaware campus.

2.4 Impact of Realistic Propagation

In much of the mobile wireless networking literature, random way-point with open-space propagation⁵ (RWP/OS) are used. It is clear that this simple propagation models is quite different from the one discussed here. A natural question is how the simulation strategy impacts the conclusions drawn from the simulations? However, a direct comparison of open-space propagation and realistic urban propagation is difficult. One problem is that random way-point mobility makes little sense in urban area. For example, since urban areas are in three dimensions, a random way-point model would have nodes flying through the air between buildings. Hence, in order to perform sensible simulation, an urban mobility model is necessary. To this end, the UDel Urban Mobility Model is used [55]. Briefly, the UDel Models Urban Mobility model realistic urban mobility model. Nearly all aspects of the model are based on surveys. For example, the model includes the time people (nodes) arrive

⁵ By open space, we mean the 2-ray model, where the received signal power decays as C_2^2/d^2 for $d < 200m$ and as C_4/d^4 for $d \geq 200m$, where C_2 and C_4 are constants.

at work, take lunch breaks, etc. The data for this part of the model is from the US Bureau of Labor and Statistics Time Use Study [71]. The mobility of office workers is based on surveys of office meetings (including two-person meetings) (e.g., [86]). And vehicle traffic is from the City of San Francisco [24]. In all, over twenty-five different types of surveys are incorporated in to the model.

Another problem with comparing open-space propagation to realistic urban propagation is that the resulting networks are very different and a hence meaningful comparison is difficult. For example, in an open-space environment, distant nodes are able to communicate whereas in an urban setting, the nodes separated by very small distances may not be able to communicate (e.g. when one of the communicating nodes is indoors and the other is outdoors), or, as shown in Figure 2.4, nodes that are far apart may be able to communicate (e.g., when communicating down an urban canyon). To see the impact of the variability in propagation, consider Figure 2.15. These plots show the mean number of nodes that a node can communicate with (i.e., the mean degree) and the coefficient of variance (CoV) of the degree. These plots show are from three scenarios, namely, urban mobility with realistic urban propagation, urban mobility with open-space propagation, and random way-point mobility with open-space propagation. In all cases, the simulated region is generated from a 9 block ($\sim 0.5km \times 0.5km$) region of Chicago. The case of random way-point also uses the Chicago map, but nodes move according to the random way-point mobility and propagation is given by the 2-ray model. While the random way-point mobility is distinct from urban mobility, we selected the node speeds to be similar, i.e., Gaussian with mean 1.34 m/s and standard deviation 0.26 [55]. In all cases, a mesh infrastructure was placed on 29 uniformly distributed streetlights. The statistics of the degree distribution of the infrastructure of the mesh and the mobile client nodes are shown for several different numbers of mobile nodes in the network. It can be seen that the degree distributions are different in a number of




Figure 2.15: Statistics of degree distribution in random way-point with open-space propagation, urban mobility with open-space propagation and urban mobility with realistic urban propagation. The mesh infrastructure nodes are fixed on top of streetlights, while the client nodes are mobile. The coefficient of variance (CoV) is the standard deviation divided by the mean.

ways. The most obvious difference is that the open-space propagation results in a much larger mean degree. On the other hand, the coefficient of variance is smaller, indicating less variation in the degree, i.e., the degree is more stable in open-space propagation environment than in the urban environment. Furthermore, comparing the statistics of the degree distribution of the infrastructure nodes to the client nodes, we see that the urban setting shows a larger difference between these types of nodes than does the open-space.

To further examine the differences between the realistic propagation and open-space propagation, consider the plots shown in Figure 2.16. These performance statistics are also generated from the 9-block region of Chicago as described above. Twenty UDP connections were initiated from a centrally located base station to randomly selected client nodes. Each connection attempted to send a packet once every 250 milli seconds. An adapted version of AODV [74] was used. Specifically, when a mobile node receives a route request (RREQ) from an infrastructure node, the mobile node will not relay the RREQ until 100 milli seconds. has past. In this




Figure 2.16: Performance of a network under different types of mobility and propagation.

way, routes tend to follow the infrastructure, but mobile-to-mobile relaying is still possible. Three environments were considered, namely, urban mobility with urban propagation, urban mobility with open-space propagation, and random way-point with open-space propagation.

As expected, Figure 2.16 shows dramatic difference between urban propagation and open-space propagation. As alluded to above, realistic node densities for urban environments result in very high densities when open-space propagation is used. Such high densities result in a large number of collisions which lead to route searches failing and even route failures. Indeed, for 750 or more mobile nodes, very few packets are successfully delivered in the open-space propagation setting. Note that there is little difference between urban mobility with open-space propagation and random way-point mobility with open-space propagation. However, this does not indicate that mobility is unimportant, but indicates that mobility is unimportant when used with open-space propagation.

2.5 Related Work

Currently, free-space and two-ray propagation are the most popular propagation models for MANETs research. For example, ns-2 [69], [70] only supports free-space and 2-ray propagation models. On the other hand, QualNet [89] supports open-space propagation as well as stochastic propagation models such as Rayleigh, Rician and Lognormal fading. QualNet also supports channel gain trace files. The UDelModels simulation package generates trace files that are compatible with QualNet. Furthermore, OPNET [4] supports open-space propagation models as well as an enhanced open-space model that accounts for hills, foliage and atmospheric affects.

There has been little effort on integrating realistic propagation into mobile wireless network simulation. Some examples where propagation modeling is considered include [51], [49], and [66] where obstacles were included in the simulated environment and propagation was limited to line-of-sight. In [49], the obstacles were randomly placed buildings. But these have limited applicability because, as is shown in the Table 2.1, most of the communication in an urban area is non line-of-sight. Other examples of propagation simulation in MANETs are [33], where raytracing is used to enhance ns-2's propagation model, and [20], where indoor MANET simulations make use of the attenuation factor model similar to what is discussed in Section 2.3.3.

Within the communications area, propagation prediction continues to be an active area of research. Raytracing (e.g., [98], [90]) and the vertical plane method [62] discussed in Section 2.3.2 are often used. See [46] for a detailed review of work in propagation prediction.

There are several commercial packages that can be used to predict coverage of a single or a small number of mobile phone base stations or wireless access points (examples include [103], [36] and [101]). These tools have limited applicability to

LUMNet simulation. Specifically, due to different goals (realistic vs. prediction), propagation predictions are typically more computationally intensive than is required for realistic network simulation. Furthermore, many tools focus on outdoor coverage for mobile phones, or indoor coverage for wireless base stations; they neglect mixed indoor/outdoor simulation scenarios.

Chapter 3

THE GRAPHICAL PROPERTIES OF URBAN MESH NETWORKS

Perhaps one of the most difficult challenges of designing network protocols is that networks are heterogeneous. A protocol that performs well in some environments might perform quite poorly in other environments. In the Internet, the environment includes the link speed, the amount and type of competing traffic, as well as the desires of the end user. In urban mesh networks, the situation is even more complex. Not only can the environment vary in the same way as it does in the Internet, but the wireless links and ad hoc topology add a new dimension to the heterogeneity. This new dimension has not been explored, rather it is typically assumed that each node is homogeneous in that it can successfully transmit to nodes within a fixed radius¹. However, in many realistic settings the wireless link is not homogeneous and a successful transmission in a fixed radius is not possible. The transmission region is heterogeneous and depends on the exact location of the node and the environment.

In mixed indoor and outdoor urban environments, this heterogeneity is extreme. The region of successful transmission does not resemble a disk and rapidly

¹ When power control is used, the radius is not fixed, when directional antennas are used, the region of transmission is not a disk. However, in both cases, a protocol is responsible for defining the region of successful transmission and it is often assumed that the protocol has unrealistic latitude when selecting a transmission region.




Figure 3.1: Signal Strength in an Urban Environment. The colors indicate the signal strength from a transmitter on the right-hand side of the modeled area. While the figures show a large variation in coverage, the only difference between the two figures is that the position of the transmitter has moved 10 meters.

varies with node movement. Figure 3.1 shows the transmission range in an urban environment with a small displacement of the transmitter. The figure shows how the radio transmission is both blocked and reflected by the buildings. Also, it shows the difficulty that radio transmissions have penetrating into buildings. This work studies mesh networks in urban environments, in particular it focuses on the graph properties of the ad hoc network. This includes an examination of connectivity, the degree distribution, and centrality. Of special interest is the spatial aspects of these properties. It will be shown that there are particular areas where nodes in that area assume certain characteristics. For example, there are areas where nodes in that area will typically have a high degree and other areas where nodes will typically have low degree. These spatial characteristics remain even when the nodes move. It will be seen that the graph from an urban mesh network displays far richer properties than the graph produced by the popular free-space/random waypoint model. It is hoped that features of this graph can be exploited to design protocols specially suited for the difficult propagation environment that face an urban mesh network.

This chapter is organized as follows. In the next section, some related work is discussed. Section 3.2 gives a brief explanation of the simulations used. Section 3.3 begins the analysis of the graph properties of mesh networks in urban environments with an examination of the dependence of connectivity on the number of nodes in the population. In Section 3.4, the duration that a node spends in and out of the largest connected cluster is examined. In Section 3.5, the degree distribution of nodes is investigated. In Section 3.6, the distribution of the centrality is found. In Section 3.7, the spatial aspects of the graph are investigated.

3.1 Related work

There has been a considerable amount of work focused on graph theoretic properties of MANETs. However, the vast majority of this work assumes a free-space environment and neglects the effect of realistic channel gain [32], [11], [31]. One exception is [12], where the connectivity is found assuming a random shadowing model. The idea behind random shadowing is that the channel gain between each link is modeled with an independent random variable. However, this random variable depends on the distance between the two nodes, but not on the channel gain between other nodes. We have found that there is a strong dependence between channel gains and thus such models are not appropriate. For example, consider three nodes, node A, node B and node C. And suppose we seek a probabilistic model for whether node A is connected to node C. A dependence between channel gains means that the connectivity between A and C depends on whether node A and B are connected. This dependence can easily be understood by examining Figures 3.7 or 3.8. For example, if a node A is in the street, its degree is typically quite high and hence has a high probability of being connected to some other node. On the other hand, if a node is inside a building, then it is less likely to be connected to other nodes. The fact that node A is connected to node B increases the probability that node A is outside and hence increases the probability that node A can communicate with

node C. Thus, random channel gain models such as Rayleigh fading models are not appropriate when *networks* are considered (they are only good for point to point communication models, as they are typically used).

There has been some previous work on other channel gain models for MANETs in urban areas. In [51] and more recently in [49], MANET simulations included some consideration of propagation in the presence of obstacles. The idea in above mentioned citations was to model buildings not as reflective, but as perfect absorbing obstacles. In [49] this model is called obstruction cone. Of course, in reality, walls reflect, absorb and transmit wireless transmissions.

3.2 Simulations

Several types of simulations are presented. These simulations use the urban propagation model and urban mobility model described in [92], [56], [55] and [91]. Briefly, the propagation characteristics are found using mechanisms described in chapter 2 and mobility model used is based on different surveys conducted by various urban planning institutions, transportation departments and Department of U. D. Bureau of Labor Statistics.

Figure 3.2 shows some of the cities investigated. As indicated, the left-hand urban area is a model of the Paddington area in London. This is a tightly packed urban area and provides a good contrast to the University of Delaware shown on the right. In comparison to Paddington, the university consists of large open areas and large buildings. The images follow different scales, the university covers nearly twice the area as that of Paddington.

3.3 Connectivity

Connectivity is perhaps the most basic property of a graph. In the case of mesh networks in urban environments, the network is often not connected. The reason for lack of connectivity is that an urban area provides many locations that

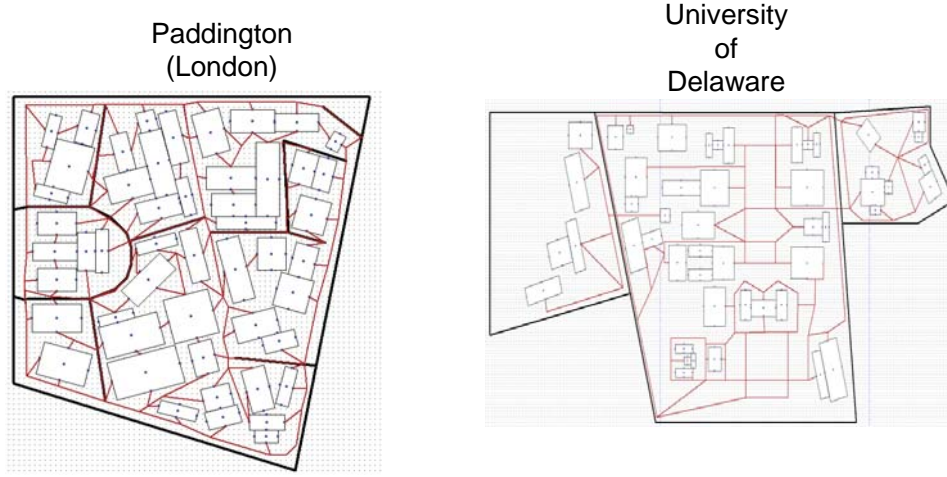


Figure 3.2: Cities used for Simulation. the rectangles are office buildings with evenly distributed offices and hallways (not shown). Outside sidewalks are shown. The mobile nodes are restricted to sidewalks, offices and hallways.

are well shielded from wireless transmissions. The connectivity problem does not necessarily change quickly when the size of the population is increased. The reason for this slow variation is that as the node density increases, the probability of a node residing in a very difficult to reach location (from the propagation perspective) also increases. For example, nodes in upper floors of corner offices where the corner does not face a busy thoroughfare are difficult to reach. Also as noted in the previous chapter, without indoor infrastructure nodes, these indoor nodes on the upper floors are hard to connect with the outdoor infrastructure nodes. In order to guarantee connectivity with such locations, there must be a high node density within the building as well as around the building.

Since the graph is so often disconnected, there is little point in measuring how often the graph is disconnected. Rather, we compute the connectivity ratio, which is the ratio of the number of nodes in the largest connected component (LCC) to the number of nodes in the entire population. Figure 3.3 shows the connectivity ratio




Figure 3.3: Connectivity Ratio as a Function of the Number of Node in the Population. The upper curve show the connectivity ratio for just the nodes that are outside, the lower curve is of nodes inside buildings and the middle curve is for all nodes (both inside and outside).

as a function of the number of nodes in the population. Note that the connectivity ratio for nodes that are inside refers to the ratio of nodes that are inside and in the LCC to the total number of nodes inside. Here the LCC contains both nodes inside and outside. Similarly, for the case of outdoor nodes, the connectivity ratio with nodes that are outside only is computed and in the case of all nodes connectivity ration with both inside and outside nodes is computed. We immediately see that the position of the node plays an important role; nearly all nodes that are outside are in the LCC, while only a small fraction of nodes in buildings are in the LCC. However, as more nodes are added to the population, the connectivity ratio increases. The figure shows that the rate of this increase depends on environment. Nodes within the large buildings in the University of Delaware campus are more difficult to connect to than the nodes in the smaller and densely packed buildings in Paddington.

3.4 Duration in and out of the largest connected component

While the main focus of this work is to examine the graph properties, the dynamic properties of the nodes and the graph are of interest as well. Here we examine the duration that a node spends in the LCC before leaving, as well as the duration spent outside of the LCC before reentering. Figure 3.4 shows the complementary distribution of these durations. Two distinct behaviors can be observed in the left-hand figure. In the case of Paddington, the distribution of the time within the LCC only depends weakly on the number of nodes in the population. Also, the figure shows that the node is likely to quickly leave the LCC (within the first 50 seconds), but as can be detected by observing the "flatness" of the distribution, if it lasts longer than 50 seconds, it will likely remain for more time (i.e., the duration suffers from infant mortality). The longer the node stays in the LCC, the more the population size matters, with the largest population leading to the longest lifetime. On the other hand, in the case of the University of Delaware, the distribution strongly depends on the size of the population with the smallest population enjoying the longest durations in the LCC. In the case of large populations on the university campus, the duration in the LCC suffers from a high infant mortality rate. The duration out of the LCC is quite different with both environments showing similar behavior and the larger the population, the shorter the time out of the LCC. 64 nodes on campus has a longest duration not in the LCC and the longest duration in the LCC. 256 nodes in Paddington have the shortest duration not in the LCC and the shortest duration in the LCC. Note that the duration not in the LCC is so short that the nodes, on average, send a significant amount of time in the LCC. This exiting and reentering of the LCC bodes poorly for mesh performance since a node not in the LCC will not be able to have connectivity.

Note that in all cases the probability decays exponentially. This can be easily seen by the "straight line" of the semilog plot. Analysis shows that these durations




Figure 3.4: Left-hand figure. The Complimentary Distribution of the Duration that a Node Spends in the LCC. Right-hand figure. The Complimentary Distribution of the Duration that a Node Spends out of the LCC.

can be modeled as a mixture of two exponential distributions. The first exponential distribution for the initial infant mortality and the second exponential distribution for longer lifetimes.

3.5 Degree distribution of nodes in an Urban Mesh Network

It is well known in wired networks that the degree distribution plays an important role in the topology. Specifically, nodes with large degree often play an important role in forwarding packets. Also, nodes with a large degree are likely to be well connected to the other nodes. Figure 3.5 shows the degree distribution for the two different urban environments and different population sizes. Here a degree of one implies that the node can only communicate with itself. Again, there is an obvious difference between nodes that are indoors and those that are outdoors; the degree of outdoor nodes is much higher. However, in both cases, the mean degree increases as the population size grows. But in the case of outdoor nodes, the degree




Figure 3.5: Degree Distribution.

distribution shows that there is a lower tail with a significant number of nodes with small degree. The reason for this is that while outdoor does provide good connectivity, there are small regions that are not well connected. As mentioned in Section 3.3, as more nodes are added, it is more likely that a node is in a hard to reach location. Note that the behavior as more nodes are added depends on the environment. Outdoors, the degree is larger for the University of Delaware campus, while indoors, the degree is larger for Paddington. The reason is that the campus buildings are large and difficult to penetrate, but the outdoor region is spacious with few buildings to block the signal.

The result of a mixed topology (i.e., one that has nodes both inside and outside) is a diverse degree distribution. Specifically, the degree distribution is bimodal with the modes sometimes a factor of 10 apart (e.g., the university campus with 256 nodes). While today's protocols do not make use of such information, an intelligent network could treat nodes with very high degree different than it treats the nodes with low degree.

3.6 Centrality distribution of node in an urban mesh network

Centrality is a useful measure to determine the extent to which a node is connected to the other nodes in the population. While there are many possible definitions of centrality, a common definition is the mean distance, in hops, to all

other nodes. A well connected node would have a low centrality and a node on the "edge" of the network would have a larger centrality. Since the graph is not connected, we speak of the centrality to nodes within the largest connected component and neglect the impact of nodes not in the LCC. Figure 3.6 shows the distribution of the centrality. The plot labeled "outside" is the distribution of the centrality from the nodes that are outdoors while the plot label "inside" is the centrality from nodes indoors. In both cases, the centrality measures the distance to all nodes (both indoors and outdoor nodes). Again the distinction between the indoor and outdoor nodes is clear. The centrality of nodes indoors is roughly the same, independent of the environment. Specifically, the distribution of the centrality of indoor nodes when there are 64 nodes in the environments is roughly the same, as the distribution when there are 128 nodes. There is a larger difference between the campus and Paddington for 256 nodes. On the other hand, the centrality of nodes outside show a clear distinction between the environments. Specifically, as nodes are added to the population, the centrality in the Paddington environment increases much faster than it does in the university campus environment. The reason that the centrality grows as nodes are added is that as more nodes are added, the nodes within buildings are included into the LCC. However, it often takes a few hops to reach nodes within buildings (typically one hop just to pass from just outside of the building to just inside). Thus, as the graph becomes more connected, the centrality of the outside nodes grows. This growth is faster for the Paddington environment as is shown in Section 3.3.

3.7 Spatial aspects of the graph

While the preceding sections clearly show the heterogeneity of the urban mesh networks in terms of indoor versus outdoor nodes, urban environments also have heterogeneity among the outdoor nodes as well as among the indoor nodes. Figure 3.7 and 3.8 show the spatial properties of urban mesh networks. Four metrics

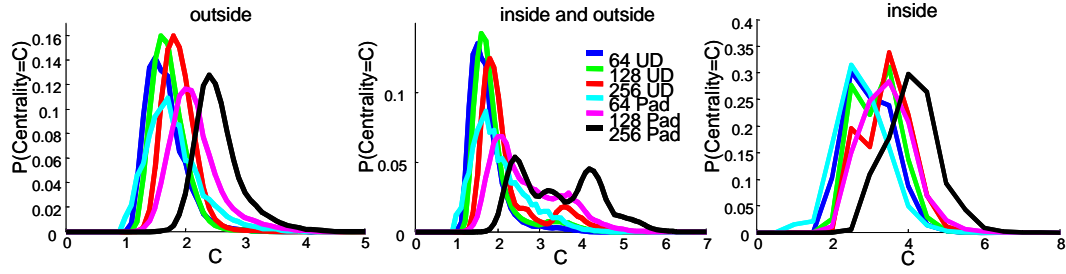


Figure 3.6: Centrality Density (Histogram).

are investigated. Three metrics, probability of being in the LCC, the mean degree, and the mean centrality have already been discussed in the previous sections. The fourth metric shows how often nodes occupy each spatial location. Specifically, the frequency plot shows the probability that a node would be in each location. These figures use colors to indicate the value of metric at each spatial location. For example, a dark area in the mean degree plot indicates that when a node is in that location, the degree of the node is typically high. The color bar along the right indicates the correspondence between the color and the value of the metric. The values of the probability of being in the LCC and the mean centrality greatly varies when the location is inside as against when it is outside. Thus, two different correspondences between color and the value of the metric are used, one for inside and one for outside. In these cases, the color bar shows both correspondences. Note that a white area indicates that no nodes were in the location and hence no value of the metric is given.

The spatial heterogeneity of urban mesh networks is clear. The figures show that along major roads or open areas, the nodes are well connected (high degree and low centrality). However, in regions where there is no major roads, nodes have poor connectivity. For example, in Paddington, consider the region near the point (300,200), or more generally, the regions along the major roads that pass




Figure 3.7: Spatial Graph Properties of the University of Delaware. The upper right figure indicates the mean degree at each location. If a node never entered a particular region, the region is left uncolored. The upper left plot shows the probability that a node in a location is in the largest connect component. However, two color scales are applied, one for inside locations and one for outside locations. The lower left shows the centrality of nodes in each position. Again, two different color scales are used, one for inside and one for outside. The lower right figure indicates the frequency or proportion of time a node was in each location.

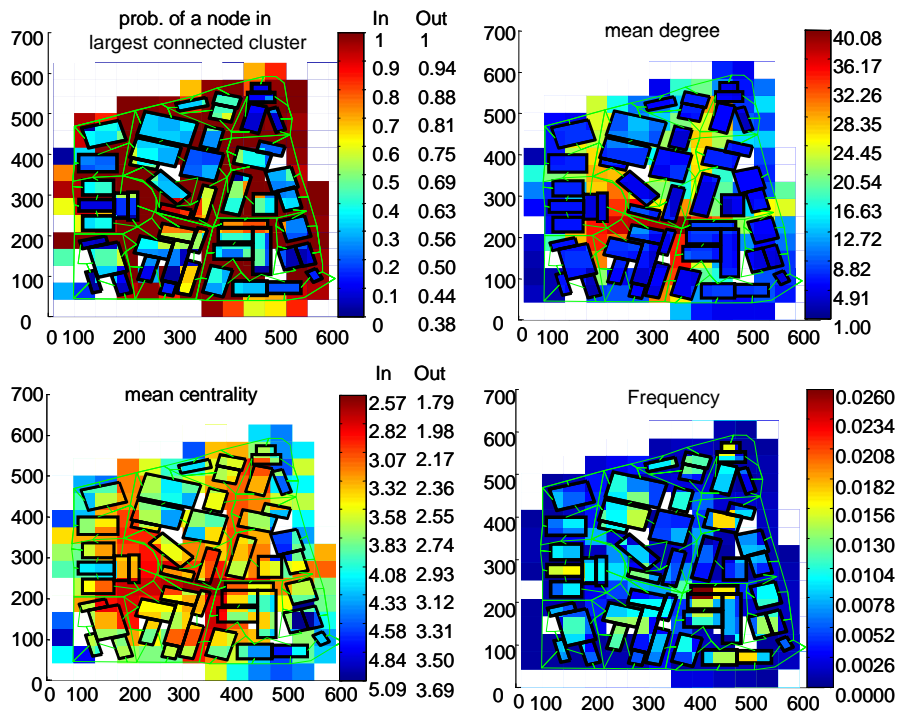


Figure 3.8: Spatial Graph Properties of the Paddington Environment

through Paddington along points $(50, 200) \rightarrow (500, 250)$, $(375, 550) \rightarrow (275, 50)$, and $(150, 500) \rightarrow (150, 350) \rightarrow (200, 250) \rightarrow (150, 200) \rightarrow (150, 50)$. Nodes along these roads are well connected. As a comparison, consider the region to the North of $(300, 200)$ where there are fewer roads and no major roadways. Nodes that find themselves in this region would be poorly connected. The lower right figure shows the fraction of time that a region is occupied by a node. Note, as can be observed by examining the occupancy, these major roadways also carry a significant amount of pedestrian traffic. However, it is not true that a region where nodes are well connected always has high occupancy. For example, the region near $(375, 500)$ is not particularly highly occupied, but nodes in this area would be well connected. The reason is that this region is on a roadway and has line-of-sight with the major intersection at $(300, 200)$. Note that this intersection has high occupancy. It should be noted that the roadways in the Paddington area are so narrow that they act like a waveguide (refer to section 2.2.3), thus the signal is transmitted down these roads very efficiently.

In these figures, the upper left figure shows the probability of a node in each location being in the LCC. As mentioned, the color scales for inside and outside are different. The connectivity into a building is affected by several factors. We see that small buildings near well connected and highly occupied regions are relatively likely to be in the largest connected component. The parts of the buildings near $(300, 450)$ are far from any major roadway and are hence unlikely to have nodes in the LCC. Recall that the communication through buildings is difficult, so it is difficult to communicate via several inside hops. Rather, in order for an inside node to be in the LCC, it must be near an outside node, or the population of nodes must be extremely large. This gives some indication of the difficulty involved in constructing an operational mesh network in an urban environment.

Chapter 4

OBSERVATIONS AND MODELS OF TIME-VARYING CHANNEL GAIN IN CROWDED AREAS

4.1 Introduction

One of the main challenges of mobile wireless networking is time-varying channels. This variability results in nodes being able to communicate at a high bit-rate and with high quality at one moment, and then, a moment later, not able to communicate at all or be able to only communicate at a low bit-rate. In order to understand the performance of mobile wireless networks in realistic environments, high-fidelity simulators or test-beds are necessary to capture the behaviors that are not amendable to analytic methods. In case of simulators and test-beds, it is important that the channels display the full range of realistic variations. In the case of test-beds, either the test-bed must include the physical mechanisms that result in realistic channel variation, or, if the physical means are not available, then some portion of the channel must be realized through artificial means. Thus, in order to verify the behavior of mobile wireless networks, the behavior of realistic channels variations must be well understood and perhaps modeled and simulated.

In the context of communication theory, variability of the channel gains is well known. In order to understand the performance of communication techniques, a number of models of the variability of channels have been developed. For example, a fast fading channel is assumed to be the result of the receiver and/or the transmitter moving. The resulting channel variations are often modeled with Jake's Model

[48]. To account for the variability induced by the mobile node moving among large objects (e.g., buildings), a time-varying shadowing model can be used. One approach is to model the channel as a correlated Gaussian process [40], [99], [14].

It is important to note that the commonly used channel models used in communication theory assume that it is the mobility of the receiver or/and transmitter that causes the channels to vary. On the other hand, it is widely known that the mobility of objects in the environment can also result in time-varying channels. However, there has been little effort focused on quantifying the impact of mobility in these dynamic environments. Furthermore, to the best of our knowledge, there are no models for the channel variation in such environments. This chapter reports on the characteristics of the time-varying channels that arise when the receiver and transmitter are stationary, but the environment contains mobile objects. We specifically focus on the impact of pedestrian mobility on the variability of the channel gain.

In the context of mesh networks, the infrastructure nodes are fixed. However, the environment could contain moving pedestrians. Furthermore, even when the receiver and/or transmitter are moving (e.g., a client node in a mesh network), the mobility of pedestrians may be responsible for a significant contribution to the channel variability. It is found that the channel in such settings displays large variability, with variation in channel gain exceeding 10 dB. Recall that the dynamic range¹ of 802.11b at 11 Mbps is roughly 50 dB [1], so, in a sense, a 10 dB variation is 20% of the entire dynamic range.

Considering the significance of the channel variation due to movement of objects in the environment, it is important that test-beds and simulators account for

¹ Typically the transmission power of 802.11b is -10dBm (100mW). Considering the *first meter* loss to be \equiv 40dB the pathloss has to be \leq 50dB in-order for the received signal be above the receiver sensitivity.

this type of channel variation. To support this need and to gain further understanding of channel variability, a continuous-time diffusion model of the variability of the channel gain is developed. It is important to note that the motivation for the diffusion-based model is that diffusion processes can provide a compact way to completely describe the observed process. A one-dimensional, four-parameter diffusion process is found to provide a good approximation. Specifically, the stationary distribution of this process is the Gamma distribution, which closely matches with the observations (Section 4.4.1). And furthermore, we find that the transition probabilities of the diffusion process approximate the empirical transition probabilities (Section 4.4.2). While the four-parameter model is fairly compact, it appears that the number of parameters could be further reduced. As will be shown, the process's parameters are strongly dependent on the variance, and the variance is dependent on the pedestrian density. This points to the possibility of developing a single parameter model, where the parameter is a function of the pedestrian density. Future work will verify this possibility.

The remainder of this chapter is as follows. In the next section a brief description of the measurement setup and the measured scenarios is provided. Then the class of diffusion models used is briefly described in Section 4.3. An overview of how the model parameters are estimated is provided in Section 4.3.1. In the data analysis section several feature are examined. Section 4.4.1 focuses on the stationary distributions, while Section 4.4.2 focuses on the transition probabilities.

4.2 Measurement setup

The central goal of this effort was to measure the *channel gain* when the channel might be subjected to blockage due to mobile pedestrians. Furthermore, the impact of the antenna height was investigated. Two Linksys BEFW11S4 routers were used. The parameters of the BEFW11S4s were set such that the data rates were fixed at 1Mbps. They were also outfitted with a 2 dBi antenna.




Figure 4.1: Measurement Configuration. T-L and T-H indicate transmitter at lower and elevated positions respectively. R indicates the receiver position.

To determine the received signal strength, the Yellowjacket [2] receiver made by Berkeley Varitronics Systems was used. The Yellowjacket 1.40 firmware was used with the system. The base stations and the receiver were configured to support 20 packets per second.

The configuration of the equipment is illustrated in Figure 4.1. Note that two transmitters are used at the same time. We refer to the higher antenna as the *elevated antenna*, and refer to the other antenna as the *lower antenna*.

Three scenarios were investigated. The *high pedestrian node density scenario* was measured between 4 PM and 5 PM on Walnut Street between 15th and 16th Streets in Philadelphia. The other two scenarios are from Trabant University Center at the University of Delaware. The *low pedestrian density* measurements were made from 4 to 5 PM, while the *moderate pedestrian density* measurements were made from 12 to 1 PM, when the building is fairly crowded due to lunchtime.

4.3 Diffusion process

Diffusions processes have been used to model a large number of continuous time (and continuous space) processes [50]. While diffusion processes have many

properties that make them especially amendable to various calculations, our purpose for employing diffusion is limited. Specifically, a diffusion process allows for a few parameters to specify the invariant distribution as well as the transition probabilities. Furthermore, the transition probability is specified for all times. That is, let $p(x, t|x_0)$ denote the probability density of the process taking the value x at time t , given that the process took the value x_0 at time 0. A diffusion model provides the function p for all the values taken by the variables x , t , and x_0 . Finally, it is straightforward to simulate the process. This last property makes it convenient and important for the hi-fidelity simulations and test-bed experiments of wireless networks. The challenge in using diffusion models is to find a model that has minimum number of parameters. We will utilize a process with four parameters.

In general, a diffusion process can be described through a stochastic differential equation, i.e.,

$$dX_t = \mu(X_t) dt + s(X_t) dW_t, \quad (4.1)$$

where X_t is the value of the process at time t , μ and s are one-dimensional functions and W_t is Brownian motion. Note that if $s \equiv 0$, then (4.1) is an ordinary differential equation. Here we are interested in the case where

$$\begin{aligned} \mu_{\lambda, \phi, \gamma, \sigma}(x) &= \frac{\sigma^2}{2} ((\gamma + \lambda - 1)x^{\gamma-1} - \phi x^\gamma) \\ s_{\gamma, \sigma}^2(x) &= \sigma^2 x^\gamma. \end{aligned} \quad (4.2)$$

In this case, it can be shown that $X_t \in [0, \infty)$ and that the probability density of the invariant (or stationary) distribution of X_t is given by $h_{\lambda, \phi}(x) = \frac{\phi^\lambda}{\Gamma(\lambda)} x^{\lambda-1} \exp(-\phi x)$ [39]. That is, X_t is Gamma distributed with parameters λ and ϕ . Note that if $\gamma = 1$, then (4.1) is the widely used CIR model [27].

For stochastic calculus [50], it can be shown that the transition probability

of X_t obeys

$$\begin{aligned} \frac{\partial p_{\lambda,\phi,\gamma,\sigma}(x,t|x_0)}{\partial t} = & \\ & - \frac{\partial}{\partial x} (\mu_{\lambda,\phi,\gamma,\sigma}(x) p_{\lambda,\phi,\gamma,\sigma}(x,t|x_0)) \\ & + \frac{1}{2} \frac{\partial^2}{\partial x^2} (s_{\gamma,\sigma}^2(x) p_{\lambda,\phi,\gamma,\sigma}(x,t|x_0)), \end{aligned}$$

where $p_{\lambda,\phi,\gamma,\sigma}(x,t|x_0)$ is the probability density of X_t given that $X_0 = x_0$ and given the parameters λ , ϕ , γ and σ . The above partial differential equation can be solved numerically, or Monte Carlo simulations of 4.1 can be used to determine the transition probability.

4.3.1 Parameter estimation methodology

The parameters, γ , λ , ϕ , and σ , are estimated in a two-stage process. First, all the data samples are used to estimate the stationary distribution, from which λ and ϕ , the parameters of the Gamma distribution, are found. Once λ and ϕ are found, σ and γ can be estimated. While it is possible to estimate the model parameters via maximum likelihood estimation, here we elected to select the parameters that minimize the L^1 norm between the observed probability density function (i.e., the histogram) and the modeled probability distribution. The motivation for using the L^1 norm is that it can be shown that if f and g are probability densities, then $\int |f(x) - g(x)| dx = 2 \sup_A |\int_A f(x) dx - \int_A g(x) dx|$. In other words, the error from using f to estimate the probability of event A as oppose to using g , is bounded by one-half of the L^1 difference between f and g [29].

The parameters are estimated as follows. Let $u_\Delta(x)$ be the fraction of observations between $x - \Delta$ and $x + \Delta$. Similarly, let $v_\Delta(x, t|x_0)$ be the fraction the observations such that $x_0 - \Delta < X_\tau \leq x_0 + \Delta$ and $x - \Delta < X_{\tau+t} \leq x + \Delta$. Then, the λ and ϕ are selected such that

$$\min_{\lambda,\phi} \sum_{k=1}^{\infty} \left| u_\Delta(k2\Delta) - \int_{k2\Delta-\Delta}^{k2\Delta+\Delta} h_{\lambda,\phi}(y) dy \right|. \quad (4.3)$$

Once λ and ϕ are determined, γ and σ can be found in a similar fashion, by solving

$$\min_{\gamma, \sigma} \sum_{k=1}^{\infty} h_{\lambda, \phi}(k2\Delta) \sum_{j=1}^{\infty} |v_{\Delta}(2j\Delta, T|2\Delta k) - \int_{j2\Delta-\Delta}^{j2\Delta+\Delta} \int_{k2\Delta-\Delta}^{k2\Delta+\Delta} p_{\lambda, \phi, \gamma, \sigma}(x, T|y) dy dx|. \quad (4.4)$$

Note that the above is the weighted L^1 norm of the transition probabilities where the weighting is given by the invariant probability distribution. The rational behind this is to force the parameters to provide the best fit for the most likely initial conditions x_0 .

In the computations below, $\Delta = 1/2$ and $T = 1$ sec.

4.4 Data analysis

The data collected represents the channel gain in dBm. We transform the data as follows. Denote the collected data at time t as Y_t dBm. Then, define $X_t := -Y_t + \max_t(Y_t)$. Thus, $X_t \geq 0$. Note that the smaller X_t , the louder the received power. It is important to note that the received signal power has been normalized by $\max_t(Y_t)$. Hence statistics such as the mean of X_t cannot be used to directly determine the probability of transmission error. However, the offset $\max_t(Y_t)$ does not impact the variability, which is of primary focus here.

Figure 4.2 shows several time series of the channel gains collected in the different scenarios described in Section 4.2. This figure also shows a sample time series plot generated by the diffusion process 4.1 and 4.2 with model parameters derived from the observations and using the techniques discussed in Section 4.3.1. As can readily be seen, the observed time series and the generated data are qualitatively the same. Hence, the diffusion model can be used to generate channel gains that are, in some ways, realistic. In this section, the observed data and quality of the diffusion process fit is analyzed in terms the stationary probability function (Section 4.4.1) and in terms of the transition probability functions (Section 4.4.2).




Figure 4.2: Observed Time Series and Synthetic Time Series of Channel Gains in Different Scenarios. The synthetic time series were generated with a diffusion model. Note that difference in scales.

4.4.1 Stationary distribution

Utilizing (4.3), the parameters of the Gamma distribution were found that minimize the L^1 norm between the histogram of the experimental data and the modeled distribution. Figure 4.3 shows the resulting Gamma distribution along with the observed histogram for various measurement scenarios. As can be seen, the Gamma distribution provides a good fit to the observed behavior for all pedestrian densities observed.

As can be observed in Figure 4.3 and in Figure 4.4, as the pedestrian density increases, the variance of the channel gain increases. This conclusion is reasonable. In a low pedestrian density setting, the received signal is a multipath signal that has bounced off of several stationary objects (e.g., the ground, walls, etc.). However, when the pedestrian density is high, there are a larger number of objects for the signal to be bounced off. The result of the increase in the number of reflectors results

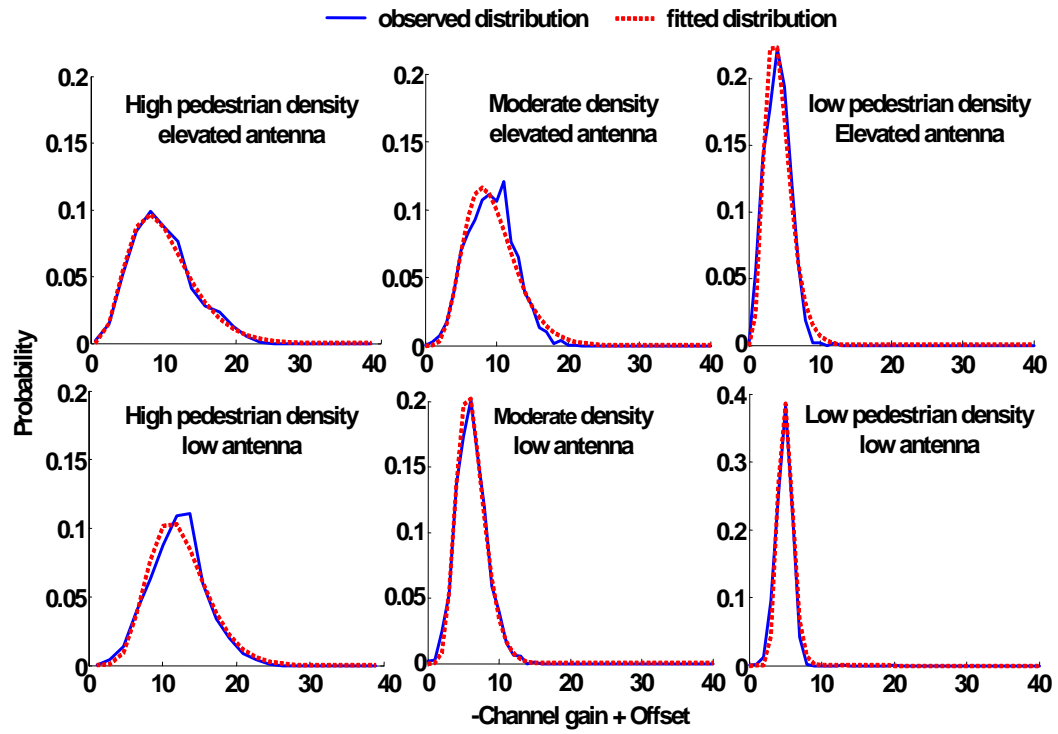


Figure 4.3: Observed and Fitted Stationary Distribution for Three Different Pedestrian Densities and Two Different Antenna Heights

Experimental Scenario	Mean	variance	λ	Φ	Stationary distribution L ¹ error	\mathbf{S}	γ	Probability transition L ¹ error
Philadelphia (Peak hour) Elevated – high density	13.9401	17.4602	4.9440	0.4920	0.1288	0.8400	1.2700	0.0584
Philadelphia (Peak hour) Low – high density	12.1627	13.8924	9.6395	0.7817	0.2302	0.8600	1.3000	0.0604
Trabant (Lunch time) Elevated – mod density	9.2259	10.7039	6.5718	0.6989	0.1467	0.8900	1.3500	0.0899
Trabant (Lunch time) Low – mod density	6.1820	4.3546	9.4812	1.5267	0.0777	0.9200	1.3800	0.1532
Trabant–(late afternoon) Elevated – low density	4.1192	2.8668	5.2893	1.2332	0.1633	0.9300	1.5300	0.1922
Trabant–(late afternoon) Low – low density	4.8699	1.2866	23.176	4.6211	0.1269	0.9800	1.9300	0.4132

Figure 4.4: The above table shows the description of different experimental scenarios and the parameters and errors for stationary distribution and probability transition functions

in a higher variability of the received signal power.

Comparing the distribution of the channel gain for elevated antennas to lower antennas, there is a relatively small difference. It can be seen that the variance is larger when the antenna is elevated. One possible explanation for this behavior is that the wireless signal transmitted from the elevated antenna is sometimes able reach the receiver via a strong line-of-sight path. When a line-of-sight path is not available, then the propagation environment facing the elevated antenna is similar to the environment facing the lower antenna, and hence, has similar channel gain variation. As a result, the received signal strength ranges from a line-of-sight strength to a rather weak signal that has experienced many reflections

Figure 4.5 shows the relationship between the mean and variance for different pedestrian densities. As expected, as the pedestrian density increases, the channel gain shows more variability. We also see that the mean increases with the pedestrian density. Furthermore, the mean approximately obeys an affine relationship with the variance, with slope of 0.68 and y-intercept 2.67. Hence, while the Gamma




Figure 4.5: The plot shows the least squares fit for relationship between the variance and the mean channel gain

distribution requires two parameters, if the variance is specified, then the mean is also approximately specified. Future work will examine the relationship between the pedestrian flow rate and the variance.

4.4.2 Transition probabilities

Now we turn to the transition probabilities. While the transition probabilities are important for determining channel behaviors such as the outage duration, we investigate these probabilities for the simple reason that the combination of the stationary probability distribution and transition probability completely characterizes a stochastic process. Hence, if the transition probabilities observed match those derived from the diffusion process given by (4.1) and (4.2), then the diffusion model can be used to generate synthetic channel gains that realistically mimic the channel gains that arise in crowded areas.

Following the approach discussed in Section 4.3.1, the parameters γ and σ were found. These parameters were used to numerically determine the transition probabilities. Figure 4.6 shows several observed and fitted transition probabilities for the high pedestrian density and low antenna scenario and the low pedestrian density




Figure 4.6: Transition Probabilities. The upper plots show the transition probabilities from the high density pedestrian scenario with a low antenna height, while lower set of plots show the transitions probabilities from the lower pedestrian density with the elevated antenna. In each case three transition probabilities are shown corresponding to three different initial conditions (x_0), namely, x_0 is the 25th percentile, the 50th percentile, and the 75th percentile.




Figure 4.7: Relationship between variance and the model parameters γ and σ .

and high antenna scenario. The other scenarios 4.4 resulted in similar quality of fit.

In Section 4.4.1 it was shown that the mean and variance approximately obey an affine function. Here we determine if any simple relationship exists for the parameters γ and σ . To this end, Figure 4.7 shows the relationship between the variance of the channel gain and these parameters. Besides the low pedestrian density scenario, both γ and σ follow an affine relationship with the channel gain variance. While more work is required to understand the low pedestrian density case, it appears that if the variance is known, then γ and σ can be found.

Chapter 5

PERFORMANCE OF 802.11B/G IN THE INTERFERENCE LIMITED REGIME

One of the key advantages of mesh networks over cellular-based networks is that mesh nodes are relatively inexpensive. Hence, it is economically feasible to spread mesh nodes at a high enough density to provide high data rates to a dense user population. Specifically, the high node density results in a short distance between receivers and transmitters, and hence high SNR channels and high bit-rates are possible. However, an important issue of high density networks is that transmissions are significantly impacted by interference. Thus, in low density networks, transmission errors are due to low SNR, while in high density networks, transmission errors are due to low SIR. The first scenario is referred to as the *noise limited regime* and the second scenario is referred to as the *interference limited regime*. The behavior of 802.11 is well understood when transmissions are noise limited. However, performance in the interference limited case is considerably less well understood. This work explores 802.11b/g in the interference limited regime.

The findings presented here are the result of a large number of laboratory experiments. Thus, the channels between transmitters and receivers were controlled and all transmissions from external sources were eliminated with EMF shielding. The experiments lead to two key findings. First, in the interference limited regime, the behavior of 802.11b/g is dominated by the ability of the receiver to synchronize

to the sender's carrier. Recall that transmissions begin with the broadcast of a synchronizing preamble followed by the data. Thus, in order to decode the data, the receiver must first successfully synchronize. The experiments discussed here show that for a large number 802.11b/g bit-rates, transmission errors are mostly caused by synchronization errors. In 802.11b, there are two similarly performing synchronization schemes, and 802.11g provides a third scheme. Thus, while 802.11b/g provides a large number of bit-rates, it essentially only provides two synchronization schemes. As a result, in the interference regime, most of the 802.11g bit-rates have the same tolerance to interference. Therefore, in contrast to the noise limited regime, in the interference limited regime decreasing the data rate provides no added ability to decode frames. Furthermore, since lowering the data rate increases the duration of the transmission, it increases the possibility that a collision will occur. As a result, in many cases, lowering the data rate decreases the performance. This has important implications for selecting the bit-rate that minimizes the time until a transmission is successful. This issue is further discussed in Section 5.4.

A second finding of the experiments described here is that if synchronization is successful, then the packet error is independent of the packet size. This significantly differs from the noise limited case where the probability of successful transmission obeys $(1 - p_{BE})^Z$, where p_{BE} is the probability of bit-error and Z is the frame size. Thus, in the noise limited case, the probability of packet loss is exponential in Z .

The remainder of this chapter proceeds as follows. The next section provides a brief discussion on the physical layer convergence procedure for 802.11b and 802.11g. In the section 5.2, the experiments and experimental set-up is described. Section 5.2.2 presents some experimental results. Section 5.3 discusses computation of transmission error probability from the data collected. Section 5.4 describes how the models presented in Section 5.3 can be used to determine the bit-rate that minimizes the expected time to successfully transmit a frame.

5.1 A brief primer on 802.11 b/g physical layer

Physical layer protocol data unit (PPDU) is the basic protocol data unit that is transmitted by any physical layer standard. For 802.11 b/g PPDU primarily consists of physical layer convergence procedure (PLCP) and the data (MPDU - MAC protocol data unit) is followed by it. The PLCP can be subdivided into two parts, the PLCP preamble and the PLCP header. PLCP preamble is as a set of symbols that are used by the receiver to train the demodulator to acquire the incoming BPSK, CCK or the OFDM signals. The PLCP header comes after the preamble and contains information such as the modulation scheme used and the length of the frame. Figure 5.1 shows the PPDU for 802.11b and 802.11g standards.

The top two frames show the PLCP preamble and the PLCP header for the 802.11b standard. Two types of preambles, short and long are defined for this standard. Most of the commercially available 802.11b cards and access points implement both these types of preambles. The short preamble is 72 bits long and the long preamble is 144 bits long. The synchronization field for the short preamble is 56 bits as against 128 bits for the long preamble. The start frame delimiter (SFD) is 16 bits long in either case. PLCP header primarily consists of 48 bits. First 8 bits of the PLCP header indicates the modulation scheme used for the MPDU that is to follow. The next 8 bits indicate the service field which is reserved for future use. The length field which is 16 bits indicate the length of the MPDU. By making use of the length and the type of modulation used, it is possible to determine the amount of time required for the current transmission to complete. Final 16 bits are reserved for CRC. 802.11b PLCP is always performed at 1Mbps (Barker code spreading with DBPSK modulation) regardless of the data rate (modulation scheme) used for transmitting MPDUs. An important thing to note about the short preamble is that after the preamble is transmitted the PLCP header is transmitted at 2Mbps (Barker code spreading with DQPSK). The 802.11b PLCP with long preamble takes $192\mu s$ and

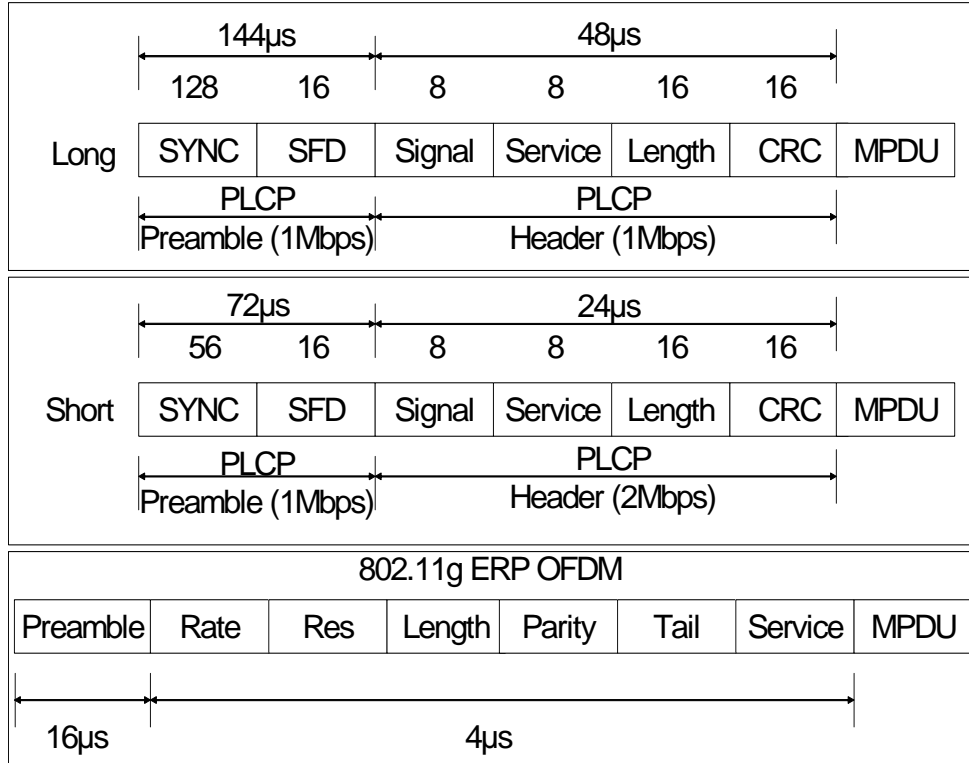


Figure 5.1: Structure of PPDU for 802.11 b/g standards. The top two figures show the PPDU for 802.11 b with long and short preambles respectively. The bottom most figure shows the PPDU for the ERP OFDM PLCP used in 802.11g.

the PLCP with short preamble takes $96\mu s$ to transmit.

Bottom frame of the figure 5.1 depicts the PLCP preamble and header for the 802.11g standard. Four different kinds of PLCP, ERP-DSSS/CCK, ERP-OFDM, ERP-PBCC and ERP-DSSS can be used for the 802.11g standard. When the service set contains clients that operate at 802.11b rates, ERP-DSSS/CCK PLCP is used. In this case either the short or the long preambles can be used. Of the other three mechanisms ERP-OFDM (extended rate phy - OFDM) is the one that is most widely implemented. ERP-OFDM allows for 7 bit rates (6,9,12,24,36 and 54Mbps) to be used. The PLCP header and the PLCP preamble of the 802.11g are always

transmitted at 6Mbps regardless of what physical rate is used for the data (MPDU) transmission. The 802.11g PLCP preamble is $16\mu s$ long with 12 symbols. Ten of these symbols are short and are used for AGC (automatic gain control), diversity selection and frequency offset estimation. The two symbols are long and are used for fine tuning of the frequency estimation. The way the signal is divided into groups of 1, 2, 4 or 6 bits depends on the data rate chosen. PLCP maps the data bits to BPSK, 16QAM or 64QAM depending on the data rate. The header of the PLCP consists primarily of fields that indicate data rate (modulation scheme), length of the MPDU frame and parity. There are other fields such as service and reserved which are left for future use. The PLCP header is $4\mu s$ long and hence the total time for PLCP procedure is $20\mu s$ in case of 802.11g.

5.2 Experiment Description

5.2.1 Experimental Setup and Protocol

Figure 5.2 depicts the block diagram of the experimental setup. The main aim of the setup is to precisely control of the channels between the transmitters and receivers and eliminate external interference.

The setup included of two access points, three laptops, and a controller computer. The two access points were the Cisco 1240 a/b/g [21] with Broadcom chipset. Prior work has shown that unlike some PCMCIA-based transmitters, the Cisco 1240 provides good transmit power stability. The sender and interferer controllers were used to adjust transmission bit-rate of the corresponding APs, transmit packets to the AP via the Ethernet (the AP then broadcasted these packets via the wireless transmitter), and to receive frame via the wireless transmitter (which was used for reference purposes). The receiver laptop was used to log all received frames. A modified MADWifi [65] driver was used to collect all frames received, including those received with bit-errors. The sender and interferer controllers and the receiver

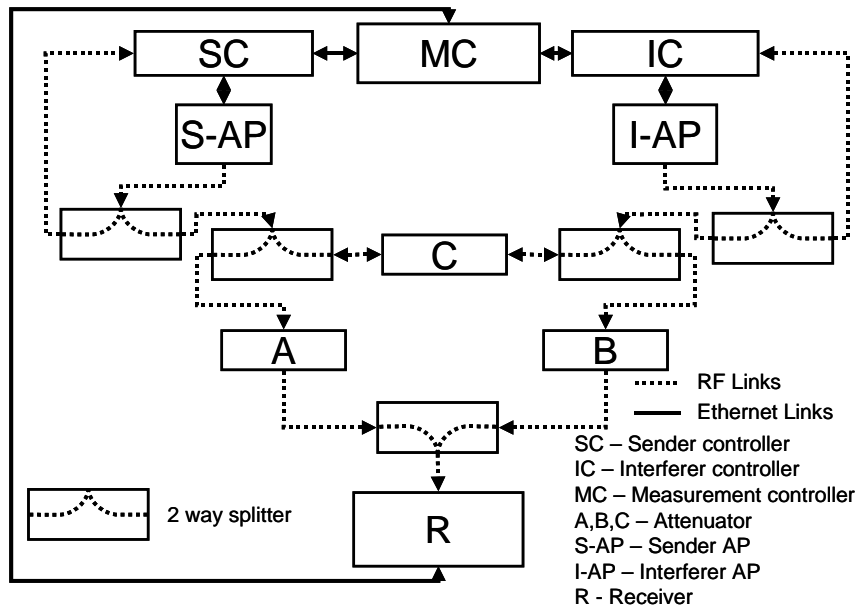


Figure 5.2: The figure shows the block diagram of the experimental setup for all sets of measurements collected for analysis. The red lines indicate the RF cables used and the green thick line indicate the Ethernet cable connections used for controlling the experiment.

were equipped with Proxim Orinoco b/g Gold Cards with Atheros AR5212 chip-set [78, 26].

In order to conduct the experiments in a controlled and repeatable fashion, the receiver, the transmitter, and the interferer were isolated from each other. The isolation is achieved by using shielded wires to carry the "wireless" signals and by using attenuators between them. In order to prevent the RF leakage from the devices, all of the access points, laptops, splitters and connectors are wrapped with an RF resistant cloth. It was found that a single layer of the RF resistant cloth provides at least 40dB of attenuation, which was found to be sufficient to keep the weak RF leakage from affecting the experiments. Furthermore, the wireless transmissions used channel 1, while there were no nearby transmitters on channels 1-10.

Each attenuator was composed of calibrated Agilent 8495A and 8494A attenuators, providing repeatability within 0.3dB. These attenuators were used to control the received signal strengths, hence, the AP's transmit power was not used. In all cases, the combined attenuation of attenuator A and B exceeded 130 dB. This is critical since the signal transmitted by the sender will partially reflect off of the receiver and be transmitted to the interferer. Thus, a combined attenuation of 130 dB ensured that the interferer is unable to detect the sender and vice versa. The level of attenuation and the inability to communicate between sender and interferer when $C = -\infty$ dB was verified through experiments. Thus, the experimental setup resulted in the hidden node topology.

The objective of the experiment was to determine the probability of receiving a frame when the transmission is subject to interference. Since transmitting frames at precise times is difficult and error prone, frames were transmitted at random times so that collision occurred at random. Specifically, the sender transmitted packets at a fixed interval of approximately 12.4 msec between transmissions, while

the interferer transmitted packets randomly with the time between transmissions exponentially distributed with mean 31 msec. The number of interferer transmissions was recorded, and hence the total duration that the channel was occupied by the interferer could be determined. With this duration and the transmission duration, the probability that the frame experienced interference can be determined. The details of this are provided in Section 5.3.

Finally, each trial consisted of the sender broadcasting 10,000 frames with RTS/CTS disabled. The frames were broadcasted, and hence, there were no retransmission nor ACKs. It was found that 10,000 frames resulted in a suitably small confidence interval.

5.2.2 Experimental results

As discussed above, the performance in the interference limited regime was investigated by transmitting packets so that collisions occurred at random. Some of the results of these experiments are shown in Figure 5.3. In this case, the average time between the beginning of the interferer's frames was approximately 31 msec. The interferer frames were 576B and sent at 1Mbps. The sender's frames were 1464B.

At high SIR, transmission errors occur with low probability. As expected, at low SIR, transmission errors occur regularly. It is important to note that the x-axis Figure 5.3 is the SIR when a collision occurs, but collisions do not always occur. Therefore, the probability of observing transmission error at low SIR is the probability of a collision occurring, i.e., the SIR is so low that if a collision occurs, then there is always an error. Since the SNR is high, when a collision does not occur, the frame is decoded with a high probability.

For 11, 12, and 24 Mbps, there is a plateau between the low SIR and the high SIR regions. Although not shown, this plateau also occurs at 5.5, 6, 9, and 18 Mbps. In the case of 6, 9, 12, 18, and 24 Mbps, this plateau ends at around 12

dB. For SIR above 12 dB, the observed probability of error is nearly the same for these bit-rates. 5.5 and 11 behave similarly, but the plateau ends at 7 dB. While not shown in Figure 5.3, 1 and 2 Mbps are similar to 11 Mbps in that they transition between a non-zero probability of transmission error and (nearly) zero probability of transmission error at 11 dB.

The reason that 6, 9, 12, 18, and 24 all have a plateau that ends at the same SIR is that they all use the same synchronization scheme and, apparently, this synchronization performs poorly when subjected to interference with SIR less than 12 dB. Similarly, 1, 2, 5.5, and 11 all use similar synchronization schemes¹ that perform poorly when the SIR is less than 7 dB. Indeed, the height of the plateaus is the probability that the sender's synchronization phase overlaps with the interferer's transmission (See the next section for details). Thus, for 802.11g (802.11b), if the SIR during synchronization is below 12 dB (7dB), then synchronization will fail nearly every time. If the sender's synchronization phase does not overlap with the interferer's transmission, then there still is a possibility that the sender's data transmission will overlap with the interferer's transmission. The probability of incorrectly decoding the data part of the frame depends on the bit-rate. Figure 5.3 indicates that for many bit-rates, the SIR that results in errors in decoding the data part of the packet is considerably smaller than the SIR required to synchronize. For example, at 12Mbps, few data errors occur if the SIR exceeds 5 dB, but synchronization will always fail unless the SIR exceeds 12 dB. In summary, the region where synchronization always fails and data decoding always succeeds is exactly the plateau. Thus, the plateaus end at similar points when the synchronization schemes are the same. Since different modulation schemes have different tolerance to interference, the plateaus begin at different points. The height of each plateau

¹ Our experiments have found no performance difference between 802.11b long and short preamble.




Figure 5.3: The observed probability of loss during the experiments described in Section 5.2.2. Note that not every frame experiences a collision, and hence, even at very low SIR, not all frames are lost. The fraction of frames lost depends on the probability that a frame experiences interference, the synchronization scheme used by the physical layer and the bit-rate used. For example, the probability of observing a transmission error at low SIR is equal to the probability of a collision occurring, i.e., every frame that experiences a collision is lost.

depends on the duration of the sender's transmissions, which, of course, depends on the bit-rate used. Note that in the case of 36Mbps (and also 48Mbps and 54Mbps), synchronization is at least as tolerant to interference as decoding data is. Thus, in these cases, there is no plateau.

5.3 Probability of Transmission Error in the Interference Limited Regime

5.3.1 Analysis

When the received signal strength is sufficiently high, a transmission error can only occur when the sender's transmission overlaps with an interferer's transmission. In this case, there are three ways that the frame is lost. First, it is possible that synchronization will fail; we refer to this type of loss as a *sync error* (See Figure 5.4). The event where synchronization fails is denoted with *SE*. Second, a frame can be lost if there are bit-errors that cannot be recovered from FEC (if FEC is

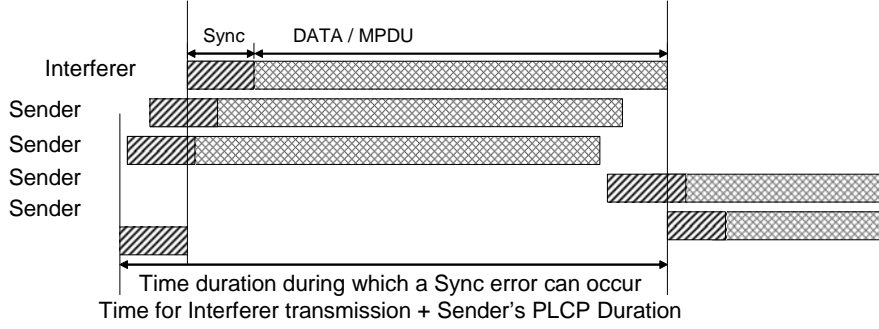


Figure 5.4: Possibilities of a sender's PLCP transmission overlapping with an interferer's transmission

used) . Since these bit-errors are detected with CRC errors, the event that a frame is lost due to a CRC error is denoted with $CRCE$ (See Figure 5.5). The third way that packets are lost is that during the decoding of the packet header and payload, synchronization lock is lost (See Figure 5.5). A loss of lock event is denoted with LL . Note that $CRCE$ can only occur if synchronization succeeded. Similarly, a loss of lock can only occur if the initial lock occurs. Furthermore, bit errors are irrelevant if there was a loss of lock. Thus, the events, $CRCE$, SE and LL are mutually exclusive. The objective of this section is to determine $P(CRCE|SIR)$ and $P(SE|SIR)$, where these probabilities denote the probability of the event occurring when a collision occurs. Prior work has found that when the SIR is low enough that $P(LL|SIR) > 0$, bit errors are common and hence $P(CRCE|SIR) \approx 1$. Thus, $P(LL|SIR)$ has no impact on the probability of frame error. For this reason, $P(LL|SIR)$ is not investigated.

Since the focus is on the interference limited regime and not on the noise limited regime, the received signal strength is high. Thus, an error can occur only when the sender is transmitting at the same time as the interferer. Specifically, a SE can

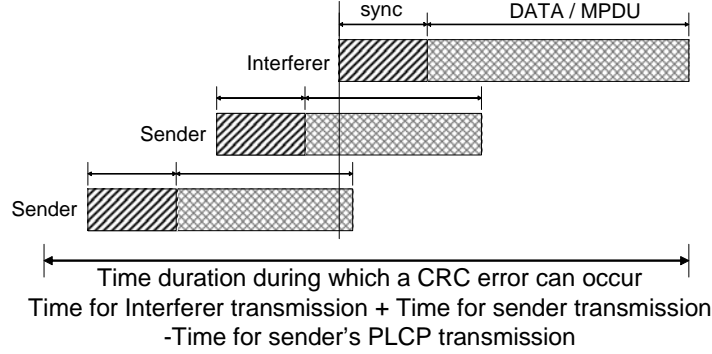


Figure 5.5: The possibilities of a sender's payload overlapping with the interferer's transmission. This can lead to CRC errors or loss of lock errors.

only occur if the sender's synchronization phase overlaps with the interferer's transmission. There are different types of overlaps. Specifically, the sender's synchronization phase could completely or partially overlap with the interferer's transmission. In the case of the experiments described in Section 5.2, the probability that the sender's synchronization phase is entirely overlapped by the interferer's transmission is $(T_{\text{Sync,Interferer}} + T_{\text{Data,Interferer}} - T_{\text{Sync,Sender}}) / T_{\text{InterfererInterval}}$, where $T_{\text{Sync,Interferer}}$ is the duration of the interfere's synchronization, $T_{\text{Data,Interferer}}$ is the duration that the interferer's transmits data, $T_{\text{InterfererInterval}}$ is the time between interferer's transmissions. A partial overlap occurs when n synchronization symbols overlap with the interferer's transmissions. In the experiments described in Section 5.2, the probability that exactly n symbols of the synchronization are overlapped with the interferer's transmissions is $T_{\text{SyncSymbol,Sender}} / T_{\text{InterfererInterval}}$, where $T_{\text{SyncSymbol,Sender}}$ is the duration of a synchronization symbol. Thus, let $P(SE|SIR, n)$ be the probability of a sync error when n of the synchronization symbols are overlapped with the interferer's transmission. Then, the probability of observing a sync error in the experiments

described in Section 5.2 is

$$\begin{aligned}
P(\text{observed SE}|\text{SIR}) &= P(SE|\text{SIR}, N_{\text{SyncSymbol,Sender}}) \\
&\times \frac{T_{\text{Sync,Interferer}} + T_{\text{Data,Interferer}} - T_{\text{Sync,Sender}}}{T_{\text{InterfererInterval}}} \\
&+ \sum_{n=1}^{N_{\text{SyncSymbol,Sender}}-1} P(SE|\text{SIR}, n) \frac{T_{\text{SyncSymbol, Sender}}}{T_{\text{InterfererInterval}}},
\end{aligned}$$

where $P(\text{observed SE}|\text{SIR})$ is the fraction of frames transmitted in the experiment that resulted in synchronization error and $N_{\text{SyncSymbol,Sender}}$ is the number of symbols in the sender's synchronization. Since the synchronization phase has a short duration, the probability of a partial overlap is much smaller than the probability of complete overlap. Thus,

$$\begin{aligned}
&P(\text{observed SE}|\text{SIR}) \\
&\approx P(SE|\text{SIR}) \frac{T_{\text{Sync,Interferer}} + T_{\text{Data,Interferer}} - T_{\text{Sync,Sender}}}{T_{\text{InterfererInterval}}},
\end{aligned} \tag{5.1}$$

where we drop the $N_{\text{SyncSymbol,Sender}}$ from $P(SE|\text{SIR}, N_{\text{SyncSymbol,Sender}})$ since a complete overlap is the only type of collision that is significant.

$P(SE|\text{SIR})$ was estimated for 11 Mbps with long and short preamble and for 12 Mbps. For verification purposes, $P(SE|\text{SIR})$ was estimated for several other bit-rates. Figure 5.6 shows $P(SE|\text{SIR})$ for various types of synchronization types. It can be seen that 802.11b short preamble and long preamble behave nearly the same. This is surprising since the long preamble has 128 bits, while the short preamble has only 56. On the other hand, this behavior was also detected in the noise limited case (an analysis of the noise limited case is not included due to space limitations).

Next, the probability of bit-errors occurring is examined. When the channel is noise limited, the probability of successfully transmitting a packet is given by $(1 - p_{BE})^Z$, where p_{BE} is the probability of bit error and Z is the frame size. However, the conclusion of extensive measurements is that when the transmission is interference limited, the probability of successfully transmitted a packet does not




Figure 5.6: Probability of synchronization failing as a function of SIR.

obey $(1 - p_{BE})^Z$. Rather, $P(CRCE|SIR)$ is independent of the frame length. On the other hand, if interference occurs randomly, then a longer frame is more likely to experience interference. Specifically, in the experiments described in Section 5.2, the probability that a frame experiences interference is $(Z_{\text{Sender}}/R_{\text{Sender}})/T_{\text{InterfererInterval}}$, where Z_{Sender} is size of the sender's frame and R_{Sender} is the sender's transmitted bit-rate. Thus, if $P(CRCE|SIR)$ is independent of frame size, then

$$P(\text{Observed CRCE}|SIR) = P(CRCE|SIR) \frac{Z_{\text{Sender}}/R_{\text{Sender}}}{T_{\text{InterfererInterval}}}, \quad (5.2)$$

which is linear in Z . Figure 5.8 shows $P(\text{Observed CRCE}|SIR)$ as a function of Z for several bit-rates and confirms that $P(\text{Observed CRCE}|SIR)$ is linear in Z and hence $P(CRCE|SIR)$ is independent of the frame length.

Figure 5.8 shows the $P(CRCE|SIR)$ for $Z_{\text{Sender}} = 1464\text{B}$ and several modulation schemes. For reference, Figure 5.8 also show $P(SE|SIR)$.

5.3.2 Discussion

Perhaps the most significant implications of Figure 5.8 is that when subject to interference, synchronization is considerably less robust to interference than decoding data. If the SNR is high and the short preamble is used, then 2, 5.5, and 11 Mbps all use the same synchronization. Thus, when subjected to interference, these modulation schemes will all have the same probability of error. Similarly,




Figure 5.7: Probability of CRC error observed during the experiments described in Section 5.2 as a function of frame size. The left-hand size is for 2Mbps while the right-hand size is for 11 and 12 Mbps, as indicated. Note that the probability of CRC error occurring is linear in the frame size.




Figure 5.8: Probability of error as a function of SIR for different types of modulation.

Bit rate (Mbps)	SNR	SIR
2	0	-5
11	5	4.75
12	5	4
24	7	8
36	12	13
48	15	18
54	16	19

Table 5.1: SIR and SNR required for probability of error of 1/2.

when subject to interference, 6, 9, 12, 18, 24, and 36 Mbps will all have the same probability of error. This behavior has implications for selecting a bit-rate. This issue is examined in more detail in the next section.

It is often assumed that the relationship between bit error and SNR is the same as the relationship between bit error and SIR or SNIR. To test this assumption, a larger number of experiments were performed to determine the noise limited performance of 802.11b/g. Table 5.1 used the results of these experiments to compare the SIR such that $P(CRCE|SIR) = 0.5$ to the SNR such that $P(CRCE|SNR) = 0.5$ under the assumption that the noise factor is 7dB, and hence the total noise (i.e., thermal noise plus the noise factor) is -93dB. Thus, excluding 2Mbps, 48Mbps, and 54Mbps, the assumption that relationship between SNR and bit-error is the same as the relationship between SIR and bit-error appears to approximately hold. Therefore, on the one hand, SIR can be treated as SNR when considering decoding the data. On the other hand, since synchronization is sensitive to interference, the performance of 802.11 frame decoding significantly depends on whether the noise is Gaussian or is interference. Thus, the relationship between SNR and transmissions error is not the same as the relationship between SIR and transmissions error.

The behavior at 2Mbps is intriguing and required a huge number of experiments to verify. On the one hand, at high SIR (e.g., 7dB), 2Mbps performed worst than 11Mbps. However, at -5dB, 2Mbps significantly outperforms 11Mbps. The

causes of this behavior are currently under investigation.

5.4 Bit-rate Selection in the Interference Limited Regime

The poor performance of 802.11b/g synchronization in the interference limited regime has significant implications for bit-rate selection. For example, if a channel has a SNR that can support 36Mbps, then reducing the bit-rate to 6, 9, 12, 18, or 24 will not impact the ability to synchronize and not increase robustness to interference. Furthermore, decreasing the bit-rate increases the transmission time and hence increases the possibility of experiencing interference. Hence, arguably, the highest bit-rate that the channel can support should be used regardless of the frequency of transmission errors. This contradicts the common approach to ARF that reduces bit-rate when packet losses are detected [52]. This section examines bit-rate selection in more detail.

Since interference can result in transmission errors, the behavior of the 802.11 backoff algorithm must be included. Specifically, for transmission at bit-rate R of a frame of size Z with probability of success p , the expected time to transmit is

$$\begin{aligned}
 ETx(Z, R, p) & \tag{5.3} \\
 &= T_{slot} \left(\sum_{j=1}^6 \frac{(2^{j+4} - 1)}{2} (1 - p)^j + \frac{1023}{2} \sum_{j=7}^{\infty} (1 - p)^j \right) \\
 &+ \frac{1}{1 - p} (DIFS + SIFS + T_{ACK} \\
 &+ T_{Sync} + T_{PLCPHeader} + \frac{Z}{R}) ,
 \end{aligned}$$

where T_{slot} is the duration of a time-slot, DIFS and SIFS are the durations of the DCF and short interface frame spacings, respectively, T_{Sync} is the time it takes to synchronize, $T_{PLCPHeader}$ is the time to transmit the PCLP header, T_{ACK} is the duration required to send an ACK or, in case that the transmission failed,

$T_{ACK} + DIFS$ is the time that the transmitter waits before beginning to decrement the backoff timer. Here, it is assumed that the ACK is transmitted at 2Mbps, which is the default value used 802.11 APs such as the Cisco 1240. The constants T_{Sync} and $T_{PLCPHeader}$ depend on whether 802.11b or g is used and in the case of 802.11b, they depend on whether the long preamble or short preamble is used. Here, it is assumed that the short preamble is used with 802.11b and the 802.11g PLCP header is used with 802.11g bit-rates. In (5.3), it is assumed that the initial value of the contention window is 31 (i.e., $2^5 - 1$) and the maximum value is 1023, i.e., $CW_{\min} = 31$ and $CW_{\max} = 1023$.

Under the assumption that RTS/CTS eliminates the majority of interference, the time to transmit a frame when RTS/CTS is used is

$$ETx(Z, R, 1) + 2 \times SIFS + T_{RTS} + T_{CTS}, \quad (5.4)$$

where T_{RTS} and T_{CTS} are the times to transmit the RTS and CTS, respectively. Again, in the results that follow, it is assumed that control packets are sent at 2Mbps.

When the transmission is interference limited, the probability of successful transmission depends on how often a collision occurs. Suppose that the interference is such that the fraction of time that the channel is occupied is ρ and the duration of the silent times between interferer transmissions is exponentially distributed with mean λ . Thus, we have the following.

Proposition 1 *The probability of a transmission failure is*

$$P(\text{frame error}|SIR) = \quad (5.5)$$

$$(\rho + (1 - \rho)(1 - \exp(-\lambda T_{Sync}))) P(SE|SIR) \quad (5.6)$$

$$+ (1 - \rho) \exp(-\lambda T_{Sync}) \quad (5.7)$$

$$\times (1 - \exp(-\lambda(Z/R))) P(CRCE|SIR) \quad (5.8)$$

$$+ (1 - (1 - \rho) \exp(-\lambda T_{Sync})) \quad (5.9)$$

$$\times (1 - P(SE|SIR)) P(CRCE|SIR) \quad (5.10)$$

Proof. There are three ways in which a failure can occur. First, there could be a sync error. A sync error can only occur if the channel is busy during synchronization, which occurs with probability $(\rho + (1 - \rho)(1 - \exp(-\lambda T_{Sync})))$. To see this, note that the channel is busy at the beginning of the synchronization phase with probability ρ . If the channel is not busy at the beginning of the synchronization phase (which occurs with probability $1 - \rho$), then it will become busy during synchronization with probability $1 - \exp(-\lambda T_{Sync})$. Thus, (5.6) is the probability of synchronization failure.

The second way a transmission failure can occur is if the channel is not busy during synchronization, but becomes busy during the transmission of the data, and then a bit-errors occur. The probability of this type of error occurring is given in (5.7-5.8).

The final way that a transmission error can occur is when the channel is busy during synchronization, but the synchronization succeeds. However, the bit-errors occur during data transmission phase. The probability of this type of error is given in (5.9-5.10). ■

With (5.3) and (5.5), the time required to transmit a frame in the face of occasional interference can be determined. Figure 5.9 shows the average time to transmit a 40B and 1400B frame when the channel is purely interference limited




Figure 5.9: Average time require to complete a transmission of a frame when $\rho = 0.3$ and $\lambda = 1/10 \text{ m sec}^{-1}$ and when the frame is 40B (left-hand side) and 1400B (right-hand side).

and the channel utilization is $\rho = 0.3$. As expected, the fastest bit-rates are either 11Mbps, 36Mbps, 48Mbps, or 54Mbps. Which of these rates is best depends on the packet size and SIR. Note that 48Mbps is nearly the same as 54Mbps, hence, little performance is lost if only 11Mbps, 36Mbps and 54Mbps are considered. We see that for small size frames, and $SIR = 10 \text{ dB}$, 11 Mbps is the best. Furthermore, at this SIR, 11Mbps is faster than 36 and 54Mbps by a factor of five, or $300\mu\text{s}$. Similarly, when the frame is 1400B and the $SIR = 16 \text{ dB}$, 36 Mbps is the fastest and is faster than 54Mbps by a factor of 2 or $400\mu\text{s}$.

Figure 5.9 shows that in the interference limited case, the bit-rate that results in the smallest expected time to successful transmission depends on the frame size, the SIR, and the channel utilization. Figure 5.10 shows the bit-rate for a wide range of frame sizes, SIRs, and channel utilizations. Figure 5.10 also shows where RTS/CTS leads to the smallest expected transmission time. In 802.11, frames larger than $RTSThreshold$ use RTS/CTS. However, the optimal value the $RTSThreshold$ is unknown. Figure 5.10 shows that the optimal value of $RTSThreshold$ depends on the channel utilization (at $\rho \leq 0.5$, $RTSThreshold > 1500\text{B}$; $\rho = 0.2$, $RTSThreshold = 500\text{B}$ and at $\rho = 0.3$, $RTSThreshold = 0$). It is possible to compute the optimal

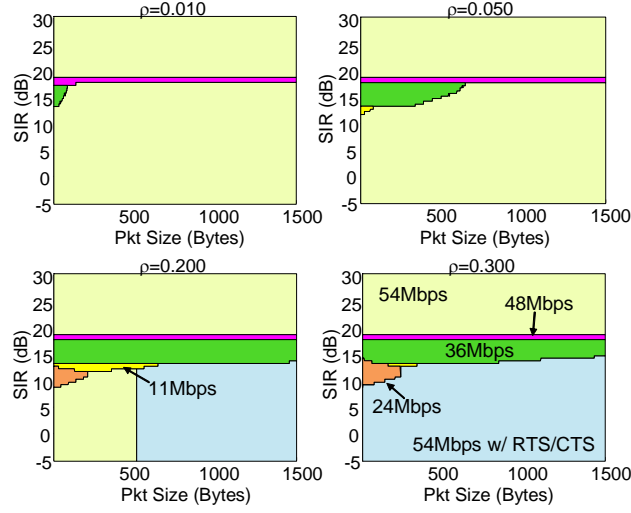


Figure 5.10: Optimal Rate Regions. For a given combination of channel utilization, ρ , frame size, and SIR, there exists a particular bit-rate that results in the smallest delay. The above shows the optimal bit-rate in each region of the packet size/SIR plane. The region where a particular bit-rate is optimal is marked with a particular color. In these plots, $\lambda = 1/10 \text{ m sec}^{-1}$.

value of RTSThreshold as a function of channel utilization. It is also possible to explore the optimal bit-rate and RTSThreshold when the channel is slightly SNR limited. For example, if the channel can only support bit-rates of 24Mbps or less.

5.5 Related work

There has been considerable effort focused on understanding the behavior of 802.11. In [13], measurements are used to explore the types of transmission errors at various bit-rates. Furthermore, [13] provides a useful explanation of many subtleties of transmission and decoding in 802.11. However, [13] does not examine the interference limited regime. In [5], the 802.11b nodes were examined in a rooftop network setting where it was found that packet error was not closely correlated with SNR. One possible explanation of this behavior is that the channel suffered from

delay spread. However, the transmissions could have also suffered from interference. One important drawback of the work presented in this paper is that the combined impacts of delay spread and interference was not studied. Further investigation in this issue is required. In [81], a PHY receiver model for decoding frames in the presence of interference was developed from measurement. While the developed model had some predictive power, it did not include the impact of synchronization, but rather relied on the replacing SNR with SNIR in the relationship between SNR and bit-error (See Section 5.3.2 above). In [58], capture was studied and a simple model for capture was developed from measurements. Contrary to findings presented here, [58] found that synchronization and decoding could as long as $SIR > 0$ dB.

Chapter 6

SIMULATION METHODOLOGY

It is well known that the variability of node-to-node communication is a major challenge facing mobile wireless networks (mobile ad hoc and mesh networks). For example at one moment, high quality communication between two nodes may be possible while a short time later, communication between the nodes may not be possible. In the case of wide bandwidth communication, such drastic changes in link quality are typically the result of node mobility. For example, if a node moves around a corner of a building, then, since the signal is not easily able to penetrate through buildings; the communication between the two nodes may be severed. Thus, a combination of node mobility and complex propagation due to the environment results in rapid variability of communication links. However, while great strides have been made in protocols, there has been very little effort devoted to understanding how to best simulate them, specifically, how to best simulate the node mobility and signal propagation. This lack of effort contrasts the simulation of wired networks where there has been extensive work focused on simulation issues such as background traffic and topology (e.g., [35], [63], [43], [61]). This chapter focuses on the simulation methodology for realistic simulation of urban mesh networks.

In urban mesh networks, the study of topology is complicated by the dependence on the propagation characteristics of the environment and location of the nodes. Propagation and the location of nodes is not random, but is dominated by structure. For example, streets, especially well traveled, wide, and straight, have

high node density and excellent propagation properties. Thus, nodes on a major street will have a large number of other nodes within communication range. However, these nodes will not be able to directly communicate with nodes on parallel streets since such communication requires transmissions through buildings or over them; something that is difficult if the buildings are large. Hence, the topology in an urban environment with large buildings will consist of well-connected nodes along the streets. Nodes near intersections will provide connectivity between two streets. Hence, the topology of outdoor nodes looks like a street map of the city. Within buildings, nodes have a smaller propagation range. Thus, the local topology of indoors and outdoors is very different. Realistic topologies can be simulated only if the propagation and mobility simulations are realistic.

There is little doubt that the mobility models currently used in the simulation of mobile ad hoc networks are not realistic. To some extent, since open-space (i.e., free-space and two-ray) propagation models that neglect the impact of objects have been used in the past, there has been little reason to use mobility models where nodes avoid or interact with objects. However, when propagation in urban environments is considered, mobility must also be addressed. Specifically, the mobility model must take into account the structure of the urban environment such as streets, sidewalks and buildings.

One of the reasons that mobility models must not be overly simplified is that in reality pedestrians and vehicles tend to move in clusters [79], [97]. That is, the locations of nodes are correlated. Furthermore, there is a well-studied relationship between node density and node speed (e.g., recall the "two-second rule" that specifies the safe driving distance between cars). Since the spatial distribution of nodes has an important impact on the behavior of mesh networking protocols, mobility models must be realistic.

In summary, the objectives of the simulation approach discussed here is to

provide realistic simulation of mobility and propagation. Specifically, for mobility, the goal is to provide realistic

- node distribution,
- node clustering (i.e., correlation in node location),
- trips including trip lengths, paths, and generation rates,
- and node speeds.

For propagation simulation, the goal is to provide realistic

- propagation range,
- signal strength,
- and spatial variation of the link quality.

Together, the mobility and propagation simulators provide realistic

- topologies,
- and variations of topologies.

The mobility simulation objectives can be achieved by employing models and model parameters that have been developed and verified by the urban planning and traffic engineering research communities. The propagation simulation objectives can be achieved by verifying and by comparing the propagation model to observations. It would be a reasonable assumption that if the propagation and mobility models are realistic then the topology and the dynamics of the topology would be realistic.

It is important to note that the objective is realistic simulation, not accurate simulation. By this we mean that the simulation should provide mobility and propagation similar to what *could* occur in an urban environment, not necessarily what *would* occur in a particular urban environment. Accurate prediction requires in depth knowledge of the modeled urban environment. For example, accurate prediction requires precise knowledge of location and dimensions of buildings and other large to moderate sized structures, as well as knowledge of the building materials used and the layout of building interiors. Furthermore, accurate mobility simulation requires knowledge of details such as the types of establishments within each building (e.g., restaurant, office, shopping, etc.) and origin-destination flow matrices for

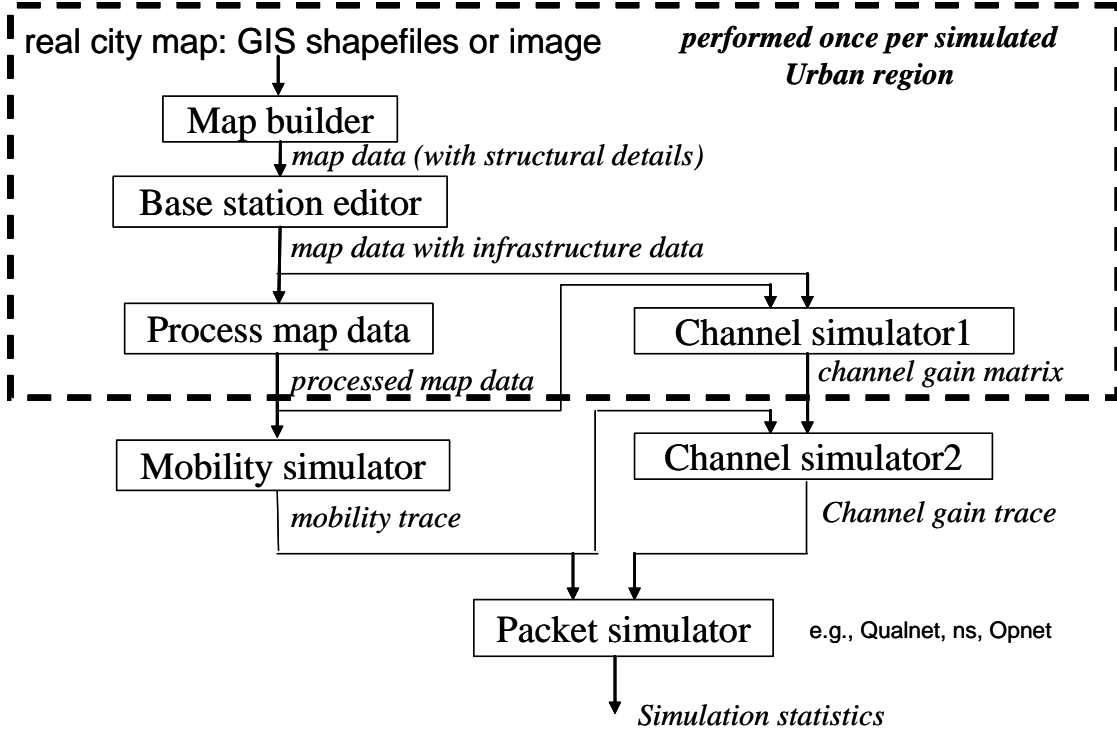


Figure 6.1: Different stages involved in realistic simulation of urban mesh networks. The output of each stage is represented in italics.

vehicle traffic. Realistic simulation, on the other hand, merely needs realistic dimensions and locations of buildings, building materials, layout of buildings interiors, and realistic trip generation for vehicles and pedestrians.

This first section of this chapter provides the overview of the different steps involved in the realistic simulation. The section 6.2 provides a brief overview of urban maps used for the propagation and the mobility simulation. Section 6.3 provides a brief primer on the UDeModels urban mobility simulator. Finally section 6.4 details the use of realistic population size (number of nodes) when conducting the simulation.

6.1 Simulation overview

There are several stages involved in realistic simulation of mesh networks. The first step detailed in 6.2 is to define the map of the urban region that is to be simulated. The stage involves editing the map to place the infrastructure nodes. The third step is to determine the propagation matrix for the city. The propagation matrix includes channel characteristics such as path loss, delay spread and angle of arrival for each source-destination in the city. This step has been dealt in detail in Chapters 2 through 5. Processed map information is used to generate mobility trace. The section 6.3 deals with this step. From the mobility trace file obtained from the stage 4 and the propagation matrix obtained from stage 3, propagation trace file containing the propagation characteristics between all pairs of nodes at every given time in the simulation can be generated. Finally the mobility and the propagation trace files can be used with any packet simulator such as QualNet, NS2, etc. to simulate protocols. The Figure 6.1 graphically depicts these different stages of simulation.

6.2 City maps

In order to model urban mesh networks, it is necessary to have a model of the urban area. There are several ways that maps suitable for simulation can be developed. First, a random city can be built as was done in [49]. In this case buildings are placed at random and a Voronoi diagram is used to construct sidewalks between the buildings. One drawback of such an approach is that important aspects of cities such as long thoroughfares and big intersections are neglected. It is well known that streets play an important role in mobile phone communication and it has been shown that streets play an important role in urban mesh networks [16].

A more realistic way to generate cities is to utilize the detailed GIS datasets [3]. These datasets include 3-dimensional building map information that provides enough detail for realistic simulation. There is an abundant number of such datasets.

For example, there are GIS datasets for most, if not all, American cities. Our map building suite of tools converts GIS datasets into format suitable for a specialized graphical editor. The graphical editor is used to "touch-up" the GIS map (e.g., remove spurious buildings, add roads, sidewalks, traffic lights, and fixed base stations). The graphical editor is also used to define locations where vehicles enter and exit the modeled area.

A third way to generate city maps is to develop a map directly in the editor. For example, idealized grid city could be generated within the editor. And finally, there has been some work on generating random, yet realistic cities [73]. Often, random cities produce GIS datasets, and hence can easily be used for propagation and mobility simulation. These realistic random cities are often generated to meet certain aesthetic requirements. It is unclear if these random cities would span a relevant range of mobility and propagation.

While GIS datasets have details of building heights and position, they typically do not provide any details about the interiors of the building. In lieu of actual interiors, they must be automatically generated. Our suite of tools assumes that all buildings are office buildings with offices that are 3.5 meters wide and $3/8$ of the building depth deep and the width of hallways is $1/4$ of the depth of the building. The hallway runs in the center of the building and stairs are on each end of the building. Incorporating automatic generation of heterogeneous building interiors will be left for future work.

6.3 A brief primer on UDelModels urban mobility simulator

The performance of mesh networks is clearly impacted by the distribution of nodes. The majority of mobility simulators assume that the nodes are uniformly spread or at least distributed according to a smooth distribution. For example, the popular random way-point mobility model leads to a smooth distribution where

nodes tend to be at the center of the modeled area [68]. Such distributions differ significantly from those that arise in realistic mobility in two ways. First, nodes are restricted to sidewalks, buildings, or roads, and second, the positions of nodes are correlated, specifically, nodes often move in groups (i.e., node arrivals are bursty). Such groups of nodes are called platoons and are well known to have an impact on the capacity of roads and sidewalks [97]. Platoons of vehicles and pedestrians can arise from traffic lights and from faster nodes catching-up, but not passing slower nodes. In the case of pedestrians, the second cause is increased by nodes that are in groups by choice. Such groups move slower than solitary nodes and limit the ability of faster nodes to pass, thus expanding the size of the group.

In this section a mobility model of people in urban areas for mesh network simulation is presented. Unlike most mobility models found in the literature, this model attempts to provide realistic mobility of people in an urban setting. This simulator is based on the data collected by various surveys including the time-use study conducted by the US Department of Bureau of Labor Statistics and several transportation departments (San Francisco, Connecticut, etc.). The simulator is primarily organized as a three tier model as shown in Figure 6.2. The highest layer is an activity model that determines the high level activity that the node is performing (e.g., working). The second layer is a task model that determines the specific task within an activity (e.g., meeting with three people). And the third layer is an agent model that determines how the person moves from one location to another (e.g., how a node navigates down a crowded hallway). The activity model is based on a recent US Department of Bureau of Labor Statistics time-use study. The task model focuses on mobility of office workers and is based on the current findings by the meetings analysis research community. The agent model is based on the work from urban planning that has extensive knowledge of pedestrian flow. This section provides a very brief introduction to the mobility model. More about

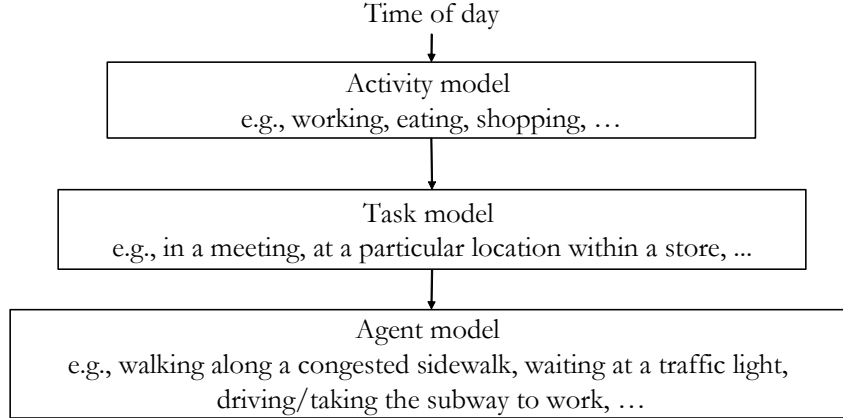


Figure 6.2: The figure illustrates the three tier model employed in the mobility simulation.

the simulation overview can be found at [92] and more about the mobility model can be found at [55] [53] and [56].

6.3.1 Activity model

This part of the mobility model is based on the US Bureau of Labor Statistics 2003 time-use study [71]. This study identifies a large number of activities. We focus on those activities that indicate location and group activities together that are performed in the same location (e.g., all activities performed at home are grouped together into the *at home* activity). While the BLS study also collected coarse location information, both activity and location information were used to determine the location used in the modeling effort.

This model effort focuses on the work day which consists of being at home, going to work, working, and perhaps taking a break and returning to work and then leaving work and returning home. The model neglects activities before and after work. Future work will include the rest of the day. For each person, the following steps are taken to determine the activities that they perform.

1. Select a home and office.
2. Determine the arrival time at work.
3. Determine the duration at work.
4. Determine if a break from work is taken. (The next 5 steps assume a break is taken.)
5. Determine the break start time.
6. Determine the number of activities performed during a break
7. Determine which activities are performed during the break.
8. Determine the duration of each activity.
9. Determine the arrival time back at work and determined if a break is taken again. If so, steps 5-9 are repeated.

6.3.2 Task model

Some activities consist of a single task. For example, eating consists of going to a restaurant. However, shopping and working consist of multiple tasks. We model shopping as a simple random walk inside the store. However, work is modeled in a more complicated manner that focuses on modeling meetings. Specifically, [86], [72], [67] have collected data on the frequency, size, and durations of meetings; [72] includes two person meetings. These studies allow the model to include worker interactions. Thus, we model mobility while at work as a sequence of meetings followed by working in the node's office. This process repeats until the work activity is complete.

More specifically, meetings are simulated as follows. The time between meetings is assumed to be exponentially distributed. When a meeting begins, a random

number of people are selected to attend the meeting. Based on the number of people attending, the duration of the meeting is determined. The duration is assumed to be exponentially distributed.

The model parameters of the model are the mean time between meetings, the distribution of the size of meetings, and the relationship between number of meeting participants and the mean meeting duration. These parameters are determined from [86], [72], [67].

6.3.3 Agent Model - Node Dynamics and Interactions

The agent model focuses on the dynamics and interaction between moving people. More specifically, the agent model consists of enforcing a distance-speed relationship between nodes and lane changing rules. The following aspects are considered in the agent model

1. Inter-node distance speed relationship: When a node with a higher desired speed catches up with a slower moving node, it will either follow or pass. To understand the dynamics of catching up, it is necessary to understand the distance-speed relationship. The impact of this relationship is that nodes can and will be tightly packed (i.e. high density) if their speed is low (congestion), but if the speed is higher, then the nodes must be further apart (low density). Since the density of nodes plays an important role in mesh networks and MANET performance, the distance-speed relationship must be understood and realistically modeled.
2. Probability of a node changing lanes: While traffic lights are an important cause of platooning, lane changing also plays an important role [79]. A node will certainly not pass if there is no room (e.g., if the other lanes are full). Even if there is room, both pedestrian and vehicle nodes might not pass out of choice and select to slow down and follow the node ahead [108]. Such decisions

lead to platooning. Platooning can change the node density at a particular place which in turn affects the communication pattern of the nodes in or near the platoon.

Another important aspect of mobility is the mode of transportation and the arrival rates in relation with the time of day. Detailed explanation of mobility model is out of scope of this thesis. More on mobility model could be found at [55], [53] and [56].

6.4 Urban Population Size

It is well known that the number of users has a major impact on the performance of the network. Thus, realistic node population sizes are an important part of realistic simulation. While the number of nodes in a network depend on the number of people in the simulated region, it also depends on the fraction of people that subscribe to the network. Today, mobile phone penetration in Europe exceeds 80% while in the US the fraction of subscribers is approximately 60%. Of course, in the early period of mobile phone deployment, the fraction of subscribers was much smaller. Hence, a wide range of penetration rates are realistic.

As expected, realistic population size in an urban region can be quite large. For example, 1 km² of Manhattan may contain 10,000 people outdoors [79], a number that is far larger than most simulations currently found in the literature. However, in a less dense city, if 10% of the population participates in the network, then a nine block region of Chicago would contain about 4000 nodes, a number that can be supported by protocol simulators such as QualNet [89]. The following presents guidelines for determining the population size in an urban region.

In the urban core, most of the indoor space is used for commercial purposes, including offices, stores, and restaurants, with office space being the most prevalent. As one moves away from the core, a larger fraction of the indoor space is used for

residences. However, it is assumed that the map specifies which buildings are office, retail, residential, or mixed usage. For office space, a survey of office use in the UK found that typical densities are approximately 15 m.² per person [104]. Thus the total working population can be determined by computing the total area of office space and dividing by 15.

The US Census American Housing Survey finds that in urban areas there is approximately 1 person per 65 m.² of residential space. The size of the residential population can be found by determining the total area of residential space and dividing by 65. However, in simulation, we assume that 92% of the people that live in the city will also work within the city, and hence are counted in the working population.

The simulator sets the population as follows

$$\begin{aligned}
\text{Number of office workers} &= \frac{\text{total office area}}{15}, \\
\text{Number of people living locally} &= \min \left(\frac{\text{total residential area}}{65}, \frac{\text{number of office workers}}{0.92} \right), \\
\text{Number of people simulated} &= \text{number of office workers} + \\
&\quad \text{number of people living locally} \times 0.08 + \\
&\quad \text{number of nonworking visitors}, \\
\text{Number of people} &= (\text{Number of office workers} - \\
\text{who commute via subway} &\quad \text{Number of people living locally} \times 0.92) \\
&\quad \times \text{MassTransitRatio}, \\
\text{Number of people who} &= (\text{Number of office workers} - \\
\text{commute via car} &\quad \text{Number of people living locally} \times 0.92) \\
&\quad \times (1 - \text{MassTransitRatio})^1
\end{aligned}$$

where the values are such that the office worker density is maintained even if there is an abundance of residential space. Note that we allow for some nonworking visitors. However, further work is required to determine realistic sizes of the nonworking visitor populations.

¹ The MassTransitRatio is the fraction of commuters that take the subway.

Chapter 7

INSIGHTS GAINED FROM REALISTIC SIMULATIONS

This chapter discusses the insights gained from simulations using realistic propagation model. Specifically the performance in terms of coverage and achievable bit rates, and topologies formed are dealt with. It is shown that the topologies formed in mesh networks pave way for development of new protocols. Specifically, an efficient paging protocol for urban mesh network setting is developed that makes use of tendril like structure formed by the mobile nodes inside a building.

The rest of this chapter is organized as follows. The section 7.1 briefly explains the indoor and outdoor coverage in an urban mesh network. The main aim of this section is to highlight the importance of differentiating the nodes with respect to indoor and outdoor environments. One of the important aspects of performance in wireless networks is the maximum data rate achievable. Section 7.2 deals with this aspect. Section 7.3 depicts the topology that can arise in the urban mesh networks that pave way for development of new protocols that utilize these topological formations. This section presents the details of an efficient paging protocol based on the topology of nodes inside buildings for the urban mesh network setting. The details of this protocols and simulation results have been presented in this section.

There have been several other topics that make use of these realistic propagation models. Some of them include [54], [15] [14], [93], [64], [18] etc.

7.1 Coverage in urban mesh networks

Coverage is the first issue that arises when considering the performance of a mesh network. Here we examine the fraction of nodes that are able to communicate with the wired network. We say that a node is able to communicate with the wired network if there is a sequence of communication channels with sufficiently high signal to noise ratio that form a path from the node to a wired base station.

The propagation environment of urban areas imposes a specific structure on the urban mesh network. Specifically, the outdoor communication occurs over good propagation environments, while indoor communication suffers from poor propagation environments. Thus, the outdoor area is well connected. For example, a relatively small number of hops are required to travel great distances. As a result, urban mesh networks are expected to provide very good coverage outdoors. For example, in Philadelphia, the requirement is that 90% of all outdoor locations will be covered. On the other hand, since the infrastructure is restricted to be outdoor, client nodes within buildings are not well connected. They are mostly only connected to nodes on the same floor or at most a few floors away. For example, consider Figure 7.1. This figure is based on propagation found from a 3-D raytracing experiment on a nine block region of downtown Chicago where the infrastructure nodes are approximately 50 m. apart. The left-hand frame shows the fraction of all outdoor and indoor location that are reachable by the infrastructure at different channel gains. As expected, outdoor locations are well connected and can support high bit-rate (i.e., low channel gain threshold) communication, while even if low bit-rate communication is used (e.g., 70 dB path loss results in 1 Mbps when 802.11b/g is used), only a small fraction of indoor locations are reachable from the infrastructure. The right-hand frame of Figure 7.1 provides further insight into the problem with indoor coverage. This figure shows how the coverage varies as a function of the floor (here it is assumed that a path loss of 70 or less will support communication).




Figure 7.1: (a) The figure illustrates the disparity in coverage of indoor and outdoor locations due to infrastructure nodes placed outdoors at various channel loss thresholds. (b) The figure illustrates the coverage on different floors. The results obtained are from a simulation carried out for a section of downtown Chicago using UDel propagation simulator. 27 infrastructure were placed outdoors on lamp posts (3.5m) with a separation of 50m between them.

Clearly, while good coverage is provided at the lower floors, the outdoor radios is not able to penetrate the upper floors. If the mesh network is composed of roof top infrastructure, then the figure would be flipped, with good coverage on the upper floors, but poor coverage on the lower floors.

To understand the topology further, consider Figure 7.2. This figure indicates the coverage of in an urban area. The area under consideration in Figure 7.2 is based on a section of downtown Chicago and the propagation is modeled with a 3-D raytracing based propagation model. The coverage of a building is indicated by the color. Specifically, the darker the color, the higher the fraction of the locations on that floor of the building is covered by the infrastructure. In Figure 7.2, the infrastructure is designed so that each location outdoors is covered by at least one infrastructure node. However, as clearly shown, in Figures 7.2 and 7.1 a significant part of the indoor area is not within range to the infrastructure. Specifically, the white colored floors on the upper floors indicate that no locations on these floors




Figure 7.2: The figure shows the coverage inside of buildings due to the infrastructure nodes placed outdoors with a separation of 50m between them. It can be observed as the floors number increases the coverage becomes almost nil. (The darker color indicates better coverage)

can communicate directly with the infrastructure.

From the bare minimum coverage point of view it can be stated that the coverage in an urban mesh network is impacted by two primary issues

1. Propagation environment: From the chapter 2 and 3, it is evident that the structure or the environment in which the mesh network is operating has a very high impact on the quality of communication. It was seen in sections 3.4, 3.5 and 3.6 that the duration in and out of the largest connected component, degree distribution and centrality of the nodes are quite different in case of outdoor and indoor nodes for simulated sections of Paddington and University of Delaware campus. For outdoor nodes, degree distribution is higher for the University of Delaware campus than that of Paddington area due to the reason that the campus environment is more spacious than that of the Paddington area. On the contrary the degree distribution of the indoor nodes for the

Paddington is higher than that of the Campus due to the reason that the campus buildings are bigger and hence are more difficult to penetrate. Similar arguments can be made about the centrality distribution and the duration of the nodes in and out of the largest connected component.

2. Indoor and outdoor nodes: In urban areas, coverage is complicated by the fact that buildings can reflect wireless signals while allowing small amount of energy to penetrate into or out of the building. As a result, communication from within a building to outside is severely impacted to the point that nodes that are indoors may not be able to directly communicate with nodes that are outdoors even when the nodes are relatively close. Similarly, wireless propagation indoors must pass through many interior walls. While interior walls typically result in less loss than exterior walls, interior propagation is also impacted to the point that communicating nodes must be closer when they are inside as compared to when they are outside. Thus, we expect that an outdoor mesh network will provide significantly better coverage to nodes that are outside than to those that are inside.

Further analysis shows that the other three issues that impact the coverage are

3. Physical layer used: The physical layer used also impacts the coverage of a mesh network. Typically, high bit-rate physical layers require less loss (stronger received signal power) than lower bit-rate physical layers. Thus, a mesh network that uses high bit-rate physical layer schemes will typically provide worst coverage than one that allows lower bit-rates.
4. Density of the infrastructure nodes: The density of the infrastructure plays an important role in the coverage; the more base stations, the more nodes that will be in range with at least one base station. On the other hand, if each fixed

relay is able to communicate with at least one base station directly or indirectly (i.e., via other fixed relays), then the fraction of the infrastructure nodes that are wired base stations does not impact the coverage. Thus, when considering coverage, there is no need to consider the fraction of the infrastructure that is made up of base stations or fixed relays.

5. Mobile nodes acting/not acting as relays: Another feature of a mesh network that may affect the coverage is whether mobile nodes may act as relays.

The degree to which these five issues impact coverage is shown in Figure 7.3. This figure is based on the behavior averaged over 20 observations over the period from 11:30 AM to 11:35 AM.

The difference between the coverage of nodes inside versus those that are outside is obvious and as expected; an outdoor infrastructure provides better coverage of nodes outside. However, note that when all nodes are considered, the coverage is nearly the same as the coverage of nodes that are inside. This is due to the fact that people are mostly inside. Even when considering the lunch time, e.g., 12:30 PM, we find that most nodes are inside (either in their work place, a restaurant, a shop, etc.).

Figure 7.3 shows that nodes that are indoors or nodes in general (i.e., when a node is selected at random) can expect "spotty" coverage from networks built with infrastructure nodes spaced at 300 m., even when only low bit-rates are desired. While, if infrastructure nodes are placed every 150 m., then 75% of nodes will be able to connect to the infrastructure at 1 Mbps. At high infrastructure densities, low bit-rate coverage of indoors is possible. Bit-rates of 54Mbps to nodes indoors appears to be not possible even when there is only 50 m. between base stations.

The impact of allowing mobile nodes to act as relays depends on the scenario. In general, the higher the density of the infrastructure, the lower the impact. Thus, when base stations are 300 m. apart, the impact of allowing mobile nodes to act

as relays is significant. For example, for general nodes, allowing mobile nodes to act as relays increases the coverage from 30% to 65%. On the other hand, when the base stations are less than 150 m. apart, allowing mobile nodes to act as relays provides only moderate improvements in coverage. It should be noted that in these simulations we assumed that about 12% of the population was participating in the network. Future work will examine the impact of a larger and smaller fraction of users.

7.2 Achievable bit rates in urban mesh networks

It is common to use achievable bit-rates as a performance metric for wireless networks. While total capacity is computationally difficult, a simple estimate of the achievable bit-rate can be found as follows. We assume that data originates at the base stations and flows to the mobile nodes. Furthermore, the data that flows to a mobile originates only at its nearest base station. The nearest base station and the routing to the mobile nodes is found using least cost routing where the cost is the path loss (i.e., the reciprocal of the channel gain). Hence, data flows along links with the highest quality. No load balancing is attempted.

It is assumed that a node receives and transmits data at the maximum bit-rate possible with the 802.11b/g physical layer. Transmissions are assumed to be at 15 dBm.

An upper bound on the bit-rate achieved by a single flow is given by the lowest link bit-rate along its path from base station to destination. Figure 7.4 shows the bit-rate averaged over all mobile nodes that are able to communicate with a base station. These bit-rates were further averaged over 20 time points. The x-axis denotes the mesh network scenario where the scenarios are defined in Table 7.1.

Figure 7.4 shows that when a node is able to communicate with the base station, the achievable data rate are quite high. For example, even when the base

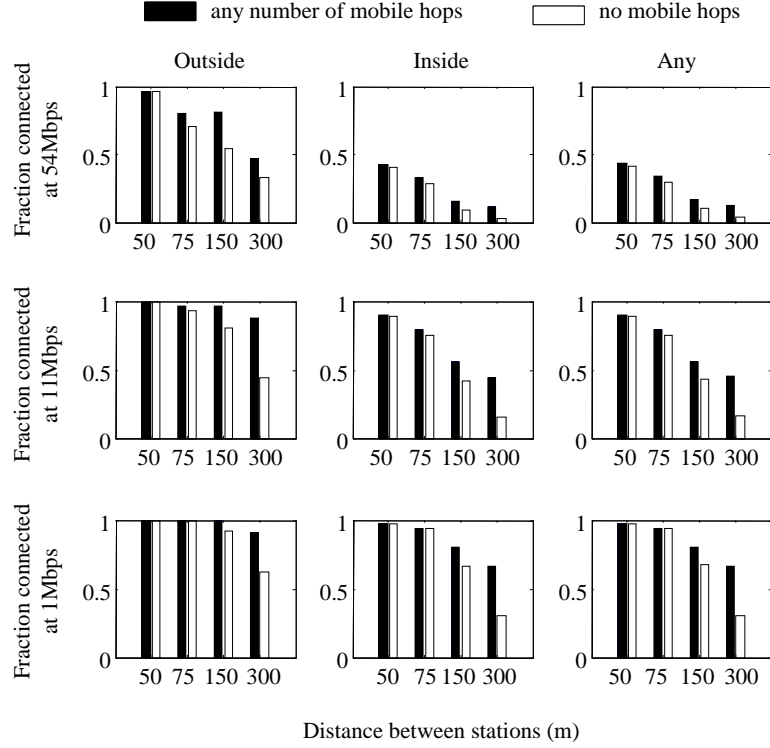



Figure 7.3: Coverage. These plots show the fraction of people that can communicate with the base station in different scenarios. The left hand column considers only nodes that are outside, the middle column only considers nodes that are inside, while the right hand column includes all nodes. The upper row is for coverage at 54 Mbps with 802.11g (-69 dBm signal strength). The middle row is for 11 Mbps (-85 dBm), while the bottom row is for 1 Mbps (-93 dBm). The x-axis indicates the distance between the infrastructure nodes. Within each infrastructure scenario, the right most bar indicates the coverage when mobile nodes are not allowed to act as relays, while the left column shows the coverage when mobile nodes may act as relays.

Table 7.1: Mesh Network Scenarios

Scenario	Distance between stations	Fraction of stations that are wired ¹
1	50 m.	1.0
2	50 m.	0.50
3	50 m.	0.25
4	75 m.	1.0
5	75 m.	0.50
6	75 m.	0.25
7	150 m.	1.0
8	150 m.	0.50
9	300 m.	1.0

**Figure 7.4:** Achievable bit rates.

stations are 300 m. apart, if communication is possible, then average data rate is approximately 4 Mbps. As more base stations and fixed relays are added, average bit-rates reach 10 Mbps. However, achievable bit-rates should be treated with care. As it was examined in [93], the high achievable bit-rates shown in Figure 7.3 may not translate into good performance of an application as experienced by the average user. Further work is required to ascertain this fact.

7.3 Scalable and efficient paging protocol utilizing the urban mesh topology

While large scale urban mesh networks (LUMNets) are being deployed, there are several challenges that remain unsolved. Perhaps the most pressing issues are related to scalability and coverage. Consider, for example, that during the lunch time, downtown Manhattan has approximately 10,000 people outdoors in 1 km² [77]. Similarly, combination of office usage surveys [104] and maps of downtown areas [3] [19] indicates that urban core may have office worker densities that exceed 40,000 people per km² in certain city sections.

Another important issue is coverage. In the case of Philadelphia, the initial plan was that there would be only one hop from the infrastructure nodes (either wired or wireless infrastructure) to the client host (or subscriber host). Due to the administrative and financial difficulties with the city owned networking including infrastructure nodes that are inside of private buildings, the infrastructure nodes are restricted to be outdoors. However, the 802.11 physical layer is not able to easily communicate from indoors to outdoors. As a result, it is not possible for the mesh network to provide one hop connectivity to all locations inside the buildings. For this reason, the initial requirements in Philadelphia and San Francisco are that the network should have at least 90% coverage outdoors and that there will be *some* location in 90% of the buildings that can connect to the wireless network. Recognizing the limitation that this puts on the utility of the network, the planners in Philadelphia considered consumer owned relays that could extend the coverage to inside of buildings. That is, the planners of these networks recognized the possibility that the urban mesh network might include not only the infrastructure and client nodes, but also client nodes that act as relays. In terms of topology, such LUMNets then resemble hybrid networks, which are essentially cellular networks that allow mobile to mobile relaying to enhance performance. While there has been some

attention to developing protocols for hybrid networks, hybrid network topologies for LUMNets is an area that has not been well explored.

The task of forming a route can be broken into two parts, finding any route to the desired host and refining that route so that it is suitable for the connection. This first task is referred to as paging and consists of simply delivering a single packet to the destination. Traditional MANET routing protocols perform this task via flooding. However, when the number of nodes ranges into the thousands, flooding becomes highly inefficient. In the case of cellular networks, paging is performed with one of more base stations broadcasting a page message. However, in large-scale urban networks, nodes are not necessarily in direct contact with the infrastructure, hence such techniques are not applicable. In this chapter, the topology and architecture of a large-scale urban mesh network is examined and a paging scheme is developed. This scheme is designed for and takes advantage of the topology, mobility, and architecture found in urban mesh networks. The resulting scheme is highly efficient in that it generates little overhead and is scalable.

While LUMNets such as those described above can be viewed as a mixture of infrastructure and ad hoc networks, the flooding techniques used in MANET routing protocols are not appropriate. For example, in the case of Philadelphia, it is expected that the network will have 200,000 users [76]. It is unreasonable to flood a network of 200,000 nodes every time a path to a destination is required. This section presents a technique that is suitable for the types of LUMNet planned for Philadelphia, San Francisco, etc. It is important to note that since these LUMNets include both infrastructure and mobility relaying, traditional paging techniques are not applicable. A key aspect of the scheme developed here is that it takes advantage of the unique topologies that arise in LUMNets.




Figure 7.5: The figure depicts the structure of multihop communication in a building. This topology resembles that of a tendril.

7.3.1 The Topology of LUMNets

Section 7.1 examined the coverage issues that arise in mesh networks for single hop communication from infrastructure nodes to the client nodes. This section deals with the topological and coverage issues that arise when multi-hop communication is considered.

If multi-hop communication is considered, then the coverage of the LUMNet is greatly extended. Indeed, if the subscriber density is high enough, then all regions of the city will be able to communicate with the infrastructure. However, due to the structure of propagation (i.e., good outdoor propagation, and bad indoor propagation), the topology of the LUMNet takes a particular form with a highly connected core that is composed of outdoor nodes and weakly connected tendrils that are composed of client nodes within buildings. Figure 7.5 shows an example of a tendril that resulted from a raytracing experiment.

We use the following nomenclature to describe the clients within the LUM-Net topology (also, see Figure 7.6). If a client node is directly reachable by the infrastructure, then the client is referred to as a *direct client* (DC). While most, or even all, outdoor nodes will be DCs, some DCs will be indoors as well. We refer to tendrils clients (TC) as client nodes that are not directly reachable from the infrastructure, but are reachable when client nodes are used as relays. Of the DCs, some can act as a gateway to the tendrils in that they have infrastructure nodes as neighbors and TCs as neighbors. Such DCs are referred to as tendrils gateways (TGWs). Of course, as client nodes move, their status may change.

Besides the client nodes, there are two types of infrastructure nodes, namely base stations that act as gateways between the wireless network and wired Internet, and fixed wireless relays. We make no distinction between these types of infrastructure nodes, and simply refer to them as WFIN (Wi-Fi infrastructure nodes).

One type of node not discussed here is the permanent consumer owned devices. Since these nodes are not owned by the mesh network provider, they cannot be full-fledged WFINs. However, it may be possible to have such nodes play an elevated role in relaying as compared to the typical mobile clients. This can be done in several ways. For example, as described below, efficient flooding within a tendrils is used to find client nodes. It is straightforward to implement efficient flooding to increase the role of some nodes. Furthermore, if the permanent client node is a TGW, then, as described below, it would advertise a long duration connection with the WFIN. The mobile clients make use of such information, and hence, are more apt to rely on such permanent clients as TGWs. However the details of taking advantage of fixed consumer owned devices is left for future work.

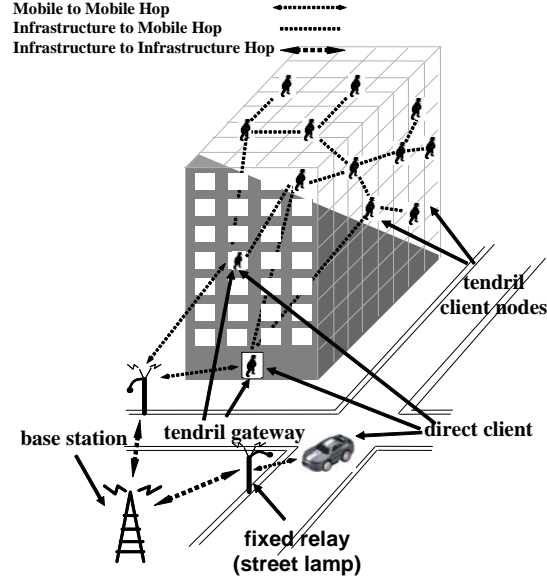


Figure 7.6: A conceptual visualization of a mesh networks. The picture above highlights the different elements of multihop wireless mesh networks.

7.3.2 Paging in Wireless Networks

As mentioned earlier, this work focuses on paging, that is, the problem of finding a destination. Earlier in this section, we distinguished paging from constructing a route. Thus, the only objective here is to efficiently find client nodes. It is important to note that we only consider the problem of the infrastructure finding a client, and not the problem of the client finding the infrastructure. However, since the infrastructure is large (it is made up of many nodes) and a particular client is small (it is only one node), this second problem is considerably more simple than the one examined here.

Paging is an integral part of mobility management and there has been extensive work on paging, especially in the cellular phone scenario [57], [87], [6], [41], [8], [9]. Paging has also been integrated into mobility management for data networks (e.g., [107], [28]). However, in these cases, the clients are within one hop of

the infrastructure, and hence these techniques cannot be applied for paging TCs. Nonetheless, the basic ideas are common to all paging algorithms, including the one developed here. Specifically, each client node will occasionally register with the infrastructure. This registration will contain various information and allows the infrastructure to learn the current location of the client. One approach is for clients to periodically register with the infrastructure, while another is for nodes to register when an event occurs. For example, a client might register when it can no longer "hear" its last WFIN or it can not hear any node from a set of WFINs that was determined during the last registration. Alternatively, a client might register when it has changed WFINs a specific number of times [6]. Besides periodic registration, these techniques require that the client is within communication range of a WFIN. While periodic registration could be applied to the LUMNets considered here, this approach has the drawback that clients that are not moving register as frequently as clients that are not moving. Hence, event-triggered registration where the events are related to mobility is preferable to periodic registration. The scheme presented here utilizes event triggered registration.

The goal of client registration is to reduce the size of the region that must be searched when the client is paged. For example, if a DC registers whenever it moves out of range of the last WFIN with which it registered, it is clear which WFIN it is reachable through. However, when a TC registers, the information about which WFIN the client is reachable from is not sufficient to reduce the region searched. To see this, one must consider the structure of the propagation discussed above, i.e., the good² propagation outdoors and bad propagation indoors. This structure means that distant parts of the city can be reached in a few hops, while the upper floors of a building requires a large number of hops. Consider, for example, a TC

² Due to the straight streets and buildings lining the streets, it is common that when propagating down a street, the signal strength decays like $d^{-\alpha}$, where $\alpha < 2$. This effect is known as propagation in urban canyons.

that has recently registered with a WFIN and has provided further information that it is 10 hops from the WFIN (such values are common when the TC is on the upper floors of a building). Thus, if the WFIN desires to reach the TC, it could, perhaps, broadcast a paging message that is to be relayed no more than $10 + K$ times, where K allows for a few extra hops in case the TC has moved. However, due to the good propagation conditions outdoors, such a message would be propagated to the far reaches of the city. Furthermore, the message needs to be relayed into all buildings near the WFIN. Consequently, the fixed radius flood would utilize bandwidth in a large part of the city.

This effect can also be seen in Figure 7.7. This figure shows the 5th and 95th percentile of the distance from the WFIN to nodes that are exactly 1, 2, 3, and 4 hops away from the WFIN. Note that as the number of hops increases, the 5th percentile does not change. These nodes are within buildings. However, the 95th percentile increases with the number of hops. As a result, specifying that a flood only spread K hops will have little impact of the geographic area the flood covers. The scheme presented next does not use hop count to control the region where the searched is performed

7.3.3 Scalable Paging for LUMNets

7.3.3.1 Overview

The basic idea of the scalable paging approach developed here is that WFIN should limit the search of the TC to the tendril where TC is located. Thus, the infrastructure must keep track of which tendril a node resides in and TCs within the tendril must also know which tendril they reside in. If these objectives are met, then the WFIN simply broadcasts a page message with the correct tendril identification and all nodes in the identified tendril forward the page message according to an efficient flooding scheme (a few possible efficient flooding schemes can be found in [102] and [106]). Hence, the overhead for the search is the overhead to efficiently




Figure 7.7: Heterogeneous Propagation. The 5th and 95th percentile of the distances from the WFIN to nodes that are exactly 1,2, 3, and 4 hops away from the WFIN.

flood a single tendril. A tendril may be a single building as shown in Figure 7.5, or may have some other structure. Regardless, the overhead required to efficiently flood a tendril is far less than the overhead required to efficiently flood the entire network. Furthermore, overhead is independent of the size of the network and only depends on the size of the tendrils within the network.

A tendril is identified by the TGW nodes that provide connectivity between the tendril and the infrastructure. In order for the TC to know which TGWs serve the tendril, the TGW floods beacons through the tendril (here, and in all situations where flooding is used, it is assumed that some sort of efficient flooding scheme is used). By examining how many hops a beacon has traveled, and by comparing this hop count to the hops traveled by beacons from other recently heard TGWs, the TC decides whether to propagate the TGW beacon. In this way, TCs can determine a set of TGWs that are between it and the WFIN. When a TC sends a registration packet to the infrastructure, it includes a list of all TGWs that it has recently heard and that provide high quality communication between the WFIN and the TC (these TGWs are called *active TGWs*). The set of TCs that share the same active TGWs

constitute a tendril. Thus, as long as the tendril information is up-to-date in both the infrastructure and the tendril nodes, then the region of the network that must propagate the page message is greatly limited.

An important challenge is to keep the tendril information up-to-date without requiring excessive overhead due to registrations. The next section describes a method that reduces the overhead due to registrations by allowing the registration of a single TC to be gratuitously applied to all the TCs within the tendril. One complication arises when a client moves into a new tendril. In this case, the TC's registrations may be gratuitously applied to all the TCs within last tendril it was in, and the registrations by TCs within the new tendril will not be gratuitously applied to TC that just entered the tendril. This complication can be eliminated by a client rapidly detecting that it has changed tendrils. The algorithm for this detection is developed in Section 7.3.3.3.

7.3.3.2 Similarity and Gratuitous Registration

While restricting the page to remain within a specific tendril greatly reduces the overhead involved in a page, it requires that tendril clients keep the WFINs informed about the TGW that they are reachable through. This requirement can lead to excessive registration overhead when TGWs move or go off/online. For example, when a DC moves close to a building and is able to communicate with nodes within the building, it becomes a TGW for the tendril within the building. Similarly, when a node inside the building moves toward the edge of the building, it may be able to communicate with the infrastructure and become a TGW. When either of these events occurs, there is a risk that every node within the tendril will try to register with the information that it is reachable through the new TGW.

To reduce the number of registrations when the TGW change, the similarity of nearby nodes is exploited, where two nodes are *similar* if they share the same set of active TGWs (active TGW will be defined shortly). The idea is that when

a registration from a tendril client is received, the WFIN applies any changes in the registration to all nodes that are similar to the registering node. Thus, one registration will suffice for the entire set of similar tendril clients. To implement this idea, the tendril nodes must be able to determine whether a node has registered and whether the WFIN will apply this registration to it. The algorithm that implements this scheme is discussed after the section 7.3.3.2.1 where active TGWs are defined.

7.3.3.2.1 Active TGW

A TC's *active TGWs* are those TGWs that offer good connectivity between the tendril and infrastructure. Specifically, an active TGW must have offered connectivity from itself to the tendril client for a sufficiently long time and with high enough quality and offer connectivity between itself and the infrastructure for a sufficiently long time and with high enough quality. Here we specify that communication quality is the fraction of successful transmissions. While not investigated here, quality could also include other metrics such as SNR and interference.

To determine which TGWs are active TGWs, the TGW beacon contains a beacon number, the beaconing period, and a measure of the quality of communication from the TGW's "best" WFIN to the TGW. The "best" WFIN is the one that has high enough quality of communication to the TGW for the longest duration. By including a beacon number, the TC is able to determine if any beacons have been missed. The beaconing period allows the tendril client to determine whether and how many beacons have been missed since the last one was received. Together, these allow the TC to determine the probability of packet delivery from the TGW to the TC. The WFIN beacons also include beacon numbers and beaconing period and allow the TGW to compute the communication quality from the WFINs. It is possible to include other measures within the TGW beacon. For example, as the beacon is being flooded through the tendril, each node could append the received signal strength. On the other hand, the nodes could append whether the node is

a fully mobile node (e.g., mobile VoIP hand-set), or semi-permanent (e.g., a desktop or consumer owned device that is specifically designed to act as a relay). Such techniques are not investigated here.

As TGW beacons are received, the TC is able to determine whether the TGW is stable (i.e. the number of beacon received exceeds a threshold) and whether the communication quality from the TGW is sufficient (e.g., the probability of a beacon being successfully received exceeds a threshold). If both of these requirements are met, then the TGW becomes an active TGW. On the other hand, a TGW is removed from the list of active TGWs when the communication quality falls below a threshold or when a sufficient number of consecutive TGW beacons have not been received. It is important to note that it takes time for a TGW to become an active TGW. Thus, a mobile node that happens to momentarily be near a building will not become an active TGW. Furthermore, hysteresis is included so that there is limited "flap" in the list of active TGWs.

7.3.3.2.2 Gratuitous Registration

In order to support similarity-based implicit registration, for each TC, three lists of active TGWs are maintained. First, the TC maintains a list of currently active TGWs. Second, the WFIN that the TC registers with maintains an estimate of the TC's active TGWs. Third, the TC maintains an estimate of the WFIN's estimate of its list of active TGWs. In general, when the TC detects that the WFIN's estimate of active TGW differs from the TC actual list of active TGWs, the TC sends a registration. After the WFIN receives a registration, the WFIN's estimate of the TC's list of active TGWs matches the TGWs list. And when the acknowledgement of the registration is received by the TC, all three lists are identical.

The following scheme is used to reduce registration overhead. This scheme is also depicted in Figures 7.8, 7.9 and 7.10. When a TC receives a TGW beacon from a TGW, it determines its list of active TGWs (perhaps adding the TGW whose

beacon just arrived and perhaps deleting any TGW from which beacons have not recently arrived). The TC then compares this updated list of active TGWs to the TC's estimate of the WFIN's list of active TGW for this node. If these lists are different, then the node will perform two tasks; first it will prepare to register with the WFIN, and second, it appends its active TGW information on the TGW beacon and propagates the TGW beacon. As is the case with all flooding, the propagation of the TGW beacon follows an efficient flooding algorithm.

Now consider the propagation of the TGW beacon that has been appended with the list of active TGWs. We will refer to the TC that appended the TGW information as the *registering TC* and any TC that receives the appended TGW beacon as a *receiving TC*. The appended active TGW information includes both the registering TC's updated list of active TGWs and its estimate of the WFIN's list of active TGW for the registering client. Upon receiving this appended TGW beacon, the receiving TC first confirms that this beacon has not travel too many hops as compared to other TGW beacons that arrive from other active TGWs. If so, the TGW beacon is dropped. If not, the receiving TC compares its estimate of the WFIN's list of active TGW to the estimate that is included in the beacon. If these two lists match, then the registering TC and the receiving TC are similar. In this case, when the WFIN receives the registration from the registering TC, it will apply the updated active TGW list to all similar clients, which includes the receiving TC. Thus, if the registering TC and the receiving TC are similar, then the receiving TC updates its estimate of the WFIN's list of active TGW for it to match the list of active TGWs included in the beacon. The receiving TC then updates its actual list of active TGW according to the received TGW beacon (i.e., to reflect the new active TGW represented by the reception of the TGW beacon and/or to reflect the TGW deleted due to missed beacons). In the typical case, the actual list and the newly updated estimate of the WFIN's list match, and hence the receiving node does not

prepare to send a registration packet. Rather, it updates the number of hops that the TGW has traveled and propagates the TGW beacon (following the rules of the efficient flooding algorithm). On the other hand, if the receiving TC is not similar to the registering TC or if the receiving TC's estimate of the WFIN's list of active TGWs for it does not match the receiving TC actual list of active TGWs, then the receiving TC prepares to register, and append its estimate of the WFIN's list of active TGW for the it along with its actual list of active TGWs. This extended TGW beacons is then propagated. In this way, a single client registration will serve to register a large number of clients greatly reducing the number of registrations. But, if this registration does not suffice for some nodes, then these nodes will register as well.

The number of registrations is further reduced as follows. When a registering node prepares to register, it first listens for a random duration for other registrations. If it overhears another node transmit a registration and these two nodes are similar (i.e., they have their estimate of the WFIN's estimate of active TGWs is the same), then the node that has not yet transmitted its registration will terminate the registration process.

While this scheme reduces the number of registration messages, it increases the impact of a dropped or lost registration packet. Hence, a node that transmits a registration packet will repeat the transmission periodically until an acknowledgement from the WFIN is received.

7.3.3.3 Detecting a Change of Tendril and Appropriation of Neighbors' Active TGW Lists

7.3.3.3.1 Detecting a Change of Tendril

One problem with the scheme above arises when a client node changes tendrils. For example, a TC may be similar to a number of TCs within a tendril. Thus, any registration sent by this TC will be applied to these other TC. However, when

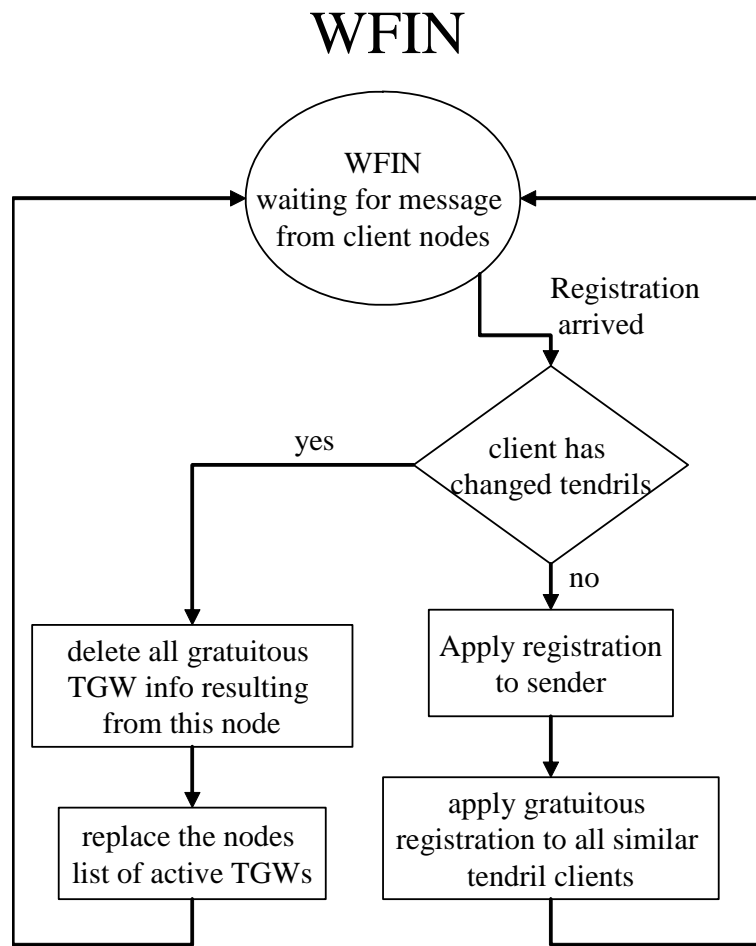


Figure 7.8: Flowchart describing the functionality of a WFIN on reception of a registration from a client.

this TC moves to another tendril, it will detect new TGWs and hence will eventually register with this new TGW information. Upon receiving such registrations, the WFIN will update the list of active TGWs for all the TCs that are similar to the registering node, which includes TC in the old tendril. Similarly, and more importantly, when a node changes tendrils, it will eventually remove TGWs from its active TGW list. When this TC sends registrations, the registrations will result in the incorrect deletion of the TGW from similar nodes. While the spurious assignment of additional TGWs to the WFIN's estimate of a TC's list of active TGWs will result in page messages being propagated through additional tendrils, the spurious deletion of TGWs could result in page messages not being propagated through the tendril where the TC resides.

In order to reduce such spurious updating of active TGW list, the WFIN will mark all TGW information it has applied to a client as gratuitous and it will record which node or nodes are responsible for the gratuitous registrations. Furthermore, each TC will attempt to detect when it changes tendrils. When a change of tendril is detected, the node will send a message to the WFIN which will delete any gratuitous TGW information that resulted from this node.

Furthermore, when a client node detects that it has changed tendrils, it will also appropriate the list of active TGW from its new neighbors. This list of TGWs is sent in a registration message to the WFIN. The WFIN then not only deletes any gratuitous TGW information due to this node, but also replaces the node's list of active TGWs with this new list. In this way, the client becomes instantly similar to the other nodes within the new tendril. These steps used to detect a change of tendril are described next. This algorithm is depicted in Figures 7.8, 7.9 and 7.10. The figures 7.9 and 7.10 shows the functionality of a client node as a flowchart. Furthermore, the fine grain details of this scheme are discussed in the analysis of the algorithm in Section 7.3.4.

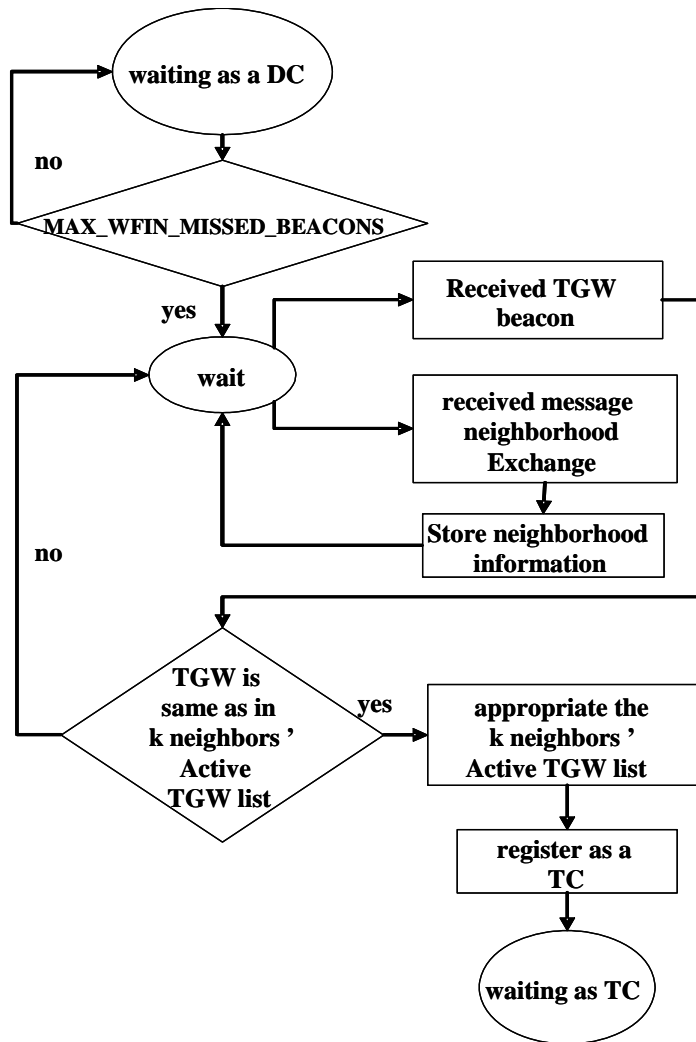


Figure 7.9: Flow chart depicting the functionality of a Tendril Gateway or Direct Client.

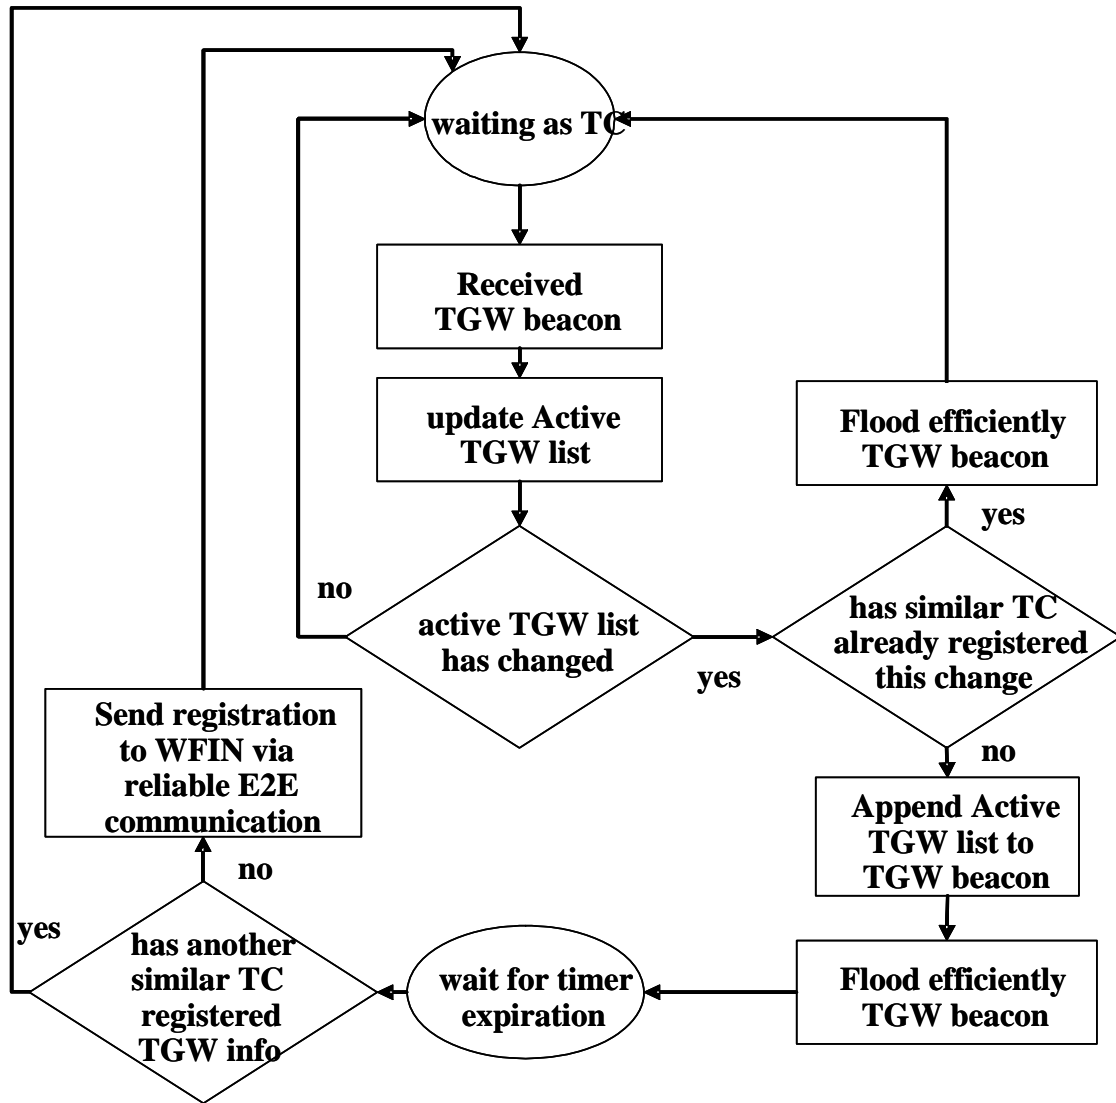


Figure 7.10: Flow chart depicting the functionality of a Tendril Client.

As mentioned above, in support of efficient flooding, TCs will periodically broadcast neighborhood information. For example, it is typical to exchange 1-hop neighbor information so that each client is able to construct its 2-hop neighborhood. Piggy-backed on the neighborhood broadcasts, TCs include their list of active TGWs. Thus, a node can easily determine whether it is similar to its neighbors. If a client is not similar to its neighbors, then it is likely that one or more nodes moved. But a difference in neighboring TCs list of active TGWs only indicates that some client has moved, it does not indicate which client moved. However, once a TGW beacon is flooded through the tendril, the TGW beacon will clarify that some active TGWs are correct and others are incorrect. When a sufficient number of TGW beacons confirm that which of the list of active TGWs is correct, the TC or TCs with an incorrect list of active TGWs will appropriate the list of active TGW from the neighbors with correct a list. This TC (or TCs) will send a special registration message indicating that the TC has changed tendrils and the new list of active TGWs will be added.

Note that the appropriating of neighbor's list of active TGWs greatly speeds up the generation of the list of active TGWs. More specifically, as mentioned above, to reduce the number of registrations and changes in active TGWs, it is important that the active TGW lists not include transitory TGWs (i.e., client nodes that are TGW for only a short time). This can easily be accomplished by only making a TGW active once a TGW when a large number of TGW beacons have arrived. A drawback of this approach is that it increases the amount of time that a TC must remain in a new tendril before it makes the TGWs that serve the tendril into active TGWs. However, by appropriating the list of active TGWs from its neighbors, a TC can determine its active TGW quickly and register with the WFIN, but still reduce the impact of transitory TGWs.

7.3.3.3.2 Detecting Entry into a Tendril

When a client moves within a large building it may move to a location where there are different active TGWs (i.e., two tendrils are overlapped). The scheme discussed above will detect such changes. A second way in which a client changes tendrils is when it moves from being an DC to a TC. It is important to distinguish the change of tendril problem and the change from a client becoming a TC. In the first scenario typically occurs when tendrils overlap. Hence, while the TC is in the area of overlap, it is reachable via both tendrils (i.e., both sets of TGWs). However, when a client becomes a TC, it means that it is no longer a DC and hence is no longer reachable by the infrastructure. Reachability will only be restored when the TC registers. The duration of the period until the TC become reachable is an important part of the performance of the paging scheme. As will be shown in Section 7.3.4, there is a trade-off between the average length of this duration and the overhead generated by the scheme.

The algorithm to detect which tendril a client has entered is similar to the scheme to detect that a TC has changed tendrils. We briefly describe the algorithm here. Further details are included in Section 7.3.4 where the trade-off between the duration of unreachability and overhead is examined.

Briefly, the a client becomes aware that it is no longer a DC when a sufficient number of WFIN beacons have been missed. The client then loses its status as a DC and becomes a TC. This event triggers the broadcast of its neighborhood information to support efficient flooding. This broadcast induces the client's new neighbors to broadcast their neighborhood information. This appended on to the neighborhood information are the TCs' list of active TGWs. The new TC records the received TGWs. When a TGW beacon is flooded, the new TC is able to confirm which of the neighbors' active TGWs are valid. The new TC then appropriates the active TGW from these TCs and registers with the WFIN.

7.3.3.4 Other issues

While there are many minor aspects to the paging scheme, some of the more important ones are as follows. While the registration process discussed above is sufficient for a node to always be reachable, the tendril nodes will periodically send registrations to the WFIN. These registration are useful to confirm gratuitous registrations and provide an upper bound on the worst-case duration that a client is not able to be paged due to incorrect or corrupted registrations (e.g., infrastructure node failure/rebooting. However, rebooting can be accommodated with specialized beacons that request registrations at a particular hop distance. Hence, the WFIN could request registrations in an expanding ring fashion.

In spite of the mechanisms described above, in many cases, when a new TGW becomes an active TGW, several nodes in the tendril will register. For example, if two nodes are within range of a newly active TGW but not within range of each other, they will both register that they are reachable via the newly active TGW. While this slightly increases the overhead due to registration, it also provides robustness in that it reduces the impact of any single TC's registration

Registrations only occur shortly after a TGW beacon is received from an active TGW. The registration follows the reverse path that the beacon took. Thus, the path is likely to exist. Furthermore, when a TC or DC forwards the registration, it appends its own registration information. Moreover, if a TC desires to make a connection with the WFIN or some node not in the mesh network, the TC can initiate a connection to the base station by using the reverse path followed by the most recently received TGW beacon.

The TGWs do not send TGW beacons into the tendril. Rather, they periodically register with the WFIN. TCs overhear the registration and flood it into the tendril as a TGW beacon.

There are several aspects of this scheme are not discussed, as they can be

achieved using techniques commonly used in mobility management and wireless networking. For example, the infrastructure must maintain and distribute information about which WFIN each client is reachable from. Also, deciding when a client is a DC, and allowing a TGW to become active more quickly when the TC and its neighbors do not have any active TGWs can be attacked using standard techniques that employ thresholds etc.

7.3.4 Analysis of Performance

One of the main performance aspects of the scalable paging algorithm described above is the overhead generated. This overhead is incurred when paging a node. Further, overhead is generated to maintain the tendrils and direct client information. This second type of overhead includes the registrations by tendrils and direct clients, the propagation of tendrils gateway beacons, the periodic broadcast of neighborhood information, and the periodic beaconing by WFINs.

The main objective of this work is to reduce the overhead involved in paging a node. Since this scalable paging algorithm restricts the paging to the tendrils where the destination is located, the overhead incurred to page a node is greatly reduced. However, besides reducing the overhead when paging a node, it is also important that clients are able to be paged, i.e., the infrastructure is aware which tendrils a client resides in. In general, as long as the network remains connected, nodes are reachable, except for a short time when the node moves from being a DC to become TC. As will be shown, there is a trade-off between the overhead generated to maintain the client information, and the duration when the node is unreachable. In this section we study this trade-off analytically. Section 7.3.5 studies this trade-off via simulations.

The TC information is maintained in three ways. First, the TCs will periodically register. However, due to the efficiency of the similarity technique, the

periodic registration can be made to occur infrequently, and hence, has little impact on the overhead. The second way in which client information is maintained is by TCs registering when a change in the direct clients for the tendril is observed. Recall (see flowchart of Figure 7.9) that when a TC detects a change in the active TGW, it propagates the TGW beacon that resulted in the detection of the change with information so that other TCs in the tendril do not register with the same information. In theory, only a single TC is required to register with the new tendril gateway information. In practice, a few TCs will register. In any case, only a small fraction of the nodes within the tendril register when the list of active TGW changes. Furthermore, this registration only occurs when the tendril gateways have changed. A node will not become a tendril gateway unless it is reasonably stable. For example, in our simulations, a direct client had to provide a stable connection for about five minutes before becoming a TGW. Thus, nodes that are currently in motion, will not become TGW. On the other hand, from studies of mobility [72], it is known that office workers are at their desk for an average of 18 minutes between trips. Thus, once a node becomes a TGW, it will stay as a TGW for an average of 13 minutes. As a result, the registration due to a change in the list of active TGWs for a tendril is rare, and does not result in significant overhead.

A third way in which the client information is maintained is by a mobile node detecting that it has entered a new tendril and then registering with the infrastructure. The registration itself results in little overhead. However, allowing the node to quickly detect that it has entered a new tendril results in significant overhead. Recall (refer flowchart of Figure 7.10) that a node detects that it has entered a new tendril with two mechanisms; the TGW beacons that the node receives are from TGWs that are not in its current list of active TGWs, and the node's own set of active TGWs differ from its neighboring nodes' set of active TGWs. The node

detects that its own set of active TGWs are different from its neighbors' by comparing its active TGWs to those within the neighborhood information packets that nodes periodically broadcast. This detection process takes some time and begins only after the client is no longer able to communicate with any WFIN. Thus, during the detection process, the node is not able to be paged. Thus, we closely examine this scenario by detailing the part of the algorithm that detects that a client is no longer a direct client and to determine which tendril it has entered. Once the node makes this determination, it registers with the WFIN (and includes its new tendril information) and is again reachable.

The first step in this process is when a DC detects that it is no longer within reach of the WFIN and occurs when a `MAX_WFIN_MISSED_BEACONS` beacons are missed. This takes $T_{WFIN} \times \text{MAX_WFIN_MISSED_BEACONS}$, where T_{WFIN} is the time between WFIN beacons. Once this occurs, the client node begins to search for TCs. Specifically, it transmits its neighborhood information. Typically, the neighborhood information includes the sender's list of active TGWs along with a list of the sender's neighboring clients. Of course, at this point, the client is not aware of these neighbors and has no active TGWs. Hence, two empty lists are included in the neighborhood information broadcast.

The neighboring TCs detect that a new node is in the vicinity and increase the rate at which they broadcast their neighborhood information and list of active TGWs. Specifically, these nodes will divide the time until they are scheduled to broadcast their neighborhood information by two. Thus, the time from when the client detects that the WFIN is no longer reachable until the time when a particular neighbor broadcasts its neighborhood information is uniformly distributed between 0 and $T_{NB}/2$, where T_{NB} is the average time between a node's neighborhood information broadcasts when a new neighbor has not been detected. Applying order statistics, the average time until k neighbors broadcast their neighborhood

information is given by

$$MNB(T_{NB}, \bar{D}, k) := \int_0^{T_{NB}/2} \frac{2\bar{D}}{T_{NB}} \binom{\bar{D}-1}{k-1} \times \left(\frac{2}{T_{NB}}x\right)^{k-1} \left(1 - \frac{2x}{T_{NB}}\right)^{\bar{D}-k} x dx, \quad (7.1)$$

where \bar{D} is the number of neighbors. It should be noted that the number of neighbors is not fixed, but a random variable. Setting \bar{D} to be the mean number of neighbors results in an approximation of the average time to receive k neighborhood information broadcasts. Figure 7.11 (a) shows this (approximate) average time until 3 such broadcasts are heard for different values of T_{NB} and \bar{D} .

The reception of neighborhood information broadcasts alone will not guarantee that a client has changed tendrils. For example, it is possible that a group of nodes has changed tendrils; hence, the neighbors of the node have moved, not the node itself. However, the entrance into a tendril is confirmed by the reception of a TGW beacon from a TGW that is included in k neighbors' lists of active TGWs (we insist that this TGW beacon is received with sufficient signal strength and reports a sufficiently long duration and high quality connection with the WFIN). The average time for a TGW beacon to arrive is

$$MTGWB(T_{TGW}, \bar{N}) := \int_0^{T_{TGW}} \bar{N} \frac{1}{T_{TGW}} \left(1 - \frac{x}{T_{TGW}}\right)^{\bar{N}-1} x dx, \quad (7.2)$$

where a TGW broadcasts a TGW beacon every T_{TGW} sec. and where there are \bar{N} TGWs for this tendril. Again, since the number of TGW is not deterministic, setting \bar{N} to be the average number of TGWs for each tendril results in an approximation of the average time until a TGW beacon is heard. Figure 7.11 (b) shows the value of $MTGWB$ for different scenarios.




Figure 7.11: (a). Mean time to receive three neighborhood broadcasts when a new node broadcasts its TGW history as a function of the total number of neighbors and T_{NB} , the time between broadcasting neighborhood information. (b) Mean time to receive a TGW beacon when there are different numbers of TGWs. The x-axis is the time between TGW beacon transmissions.

Finally, the mean time for a client node to detect that it has entered a new tendril is

$$T_{WFIN} \times \text{MAX_WFIN_MISSED_BEACONS} + \max(MNB(T_{NB}, D, k), MTGWB(T_{TGW}, N)) \quad (7.3)$$

This function is shown in Figure 7.12. The mean time to register after entering a tendril was also found from simulations; these results are shown in Figure 7.14. These figures are in reasonable agreement.

From (7.1), (7.2), and (7.3), it is clear that the duration that a node is unreachable after entering into a tendril depends on the density of nodes (i.e., \bar{D} is the average degree of a client that enters the tendril and \bar{N} is the average number of TGWs), on the rate that beacons are sent and the rate that neighborhood broadcasts are performed, i.e., T_{WFIN} , T_{TGW} , T_{NB} . It appears reasonable that these constants could be made adaptive. This tactic will be investigated in future efforts.




Figure 7.12: Mean time that a node is unreachable after it changes from a direct client to a tendril client. It is assumed that the $T_{INBeacon} = 0.5\text{sec.}$, $\text{MAX_IN_MISSED_BEACONS} = 6$, $T_{NB} = 15\text{ sec.}$, and the other parameters are as indicated.

7.3.5 Simulation results

From the analysis above, approximate parameter values can be found that satisfy the desired trade-off between the duration of unreachability and the overhead. Here we use simulations to perform further analysis. These simulations used the UDelModels [56] [91] and QualNet. The UDelModels [17] are a set of simulation tools that provide realistic mobility and propagation in urban areas. Several different scenarios were examined, namely, a 1-building/1-WFIN scenarios, a 2-building/1-WFIN scenario, a 4-building/2-WFIN scenario, and an 8-building/3-WFIN scenario. All buildings had 10 floors and the node population was approximately $100 \times$ the number of buildings. It is important to note that 100 nodes/building is considerably fewer than typically occupy a 10 story building. However, as the node density increases, efficient flooding plays an increasingly important role in the performance. While efficient flooding is well studied [102], there has been little work devoted to the specifics of flooding in an urban environment.

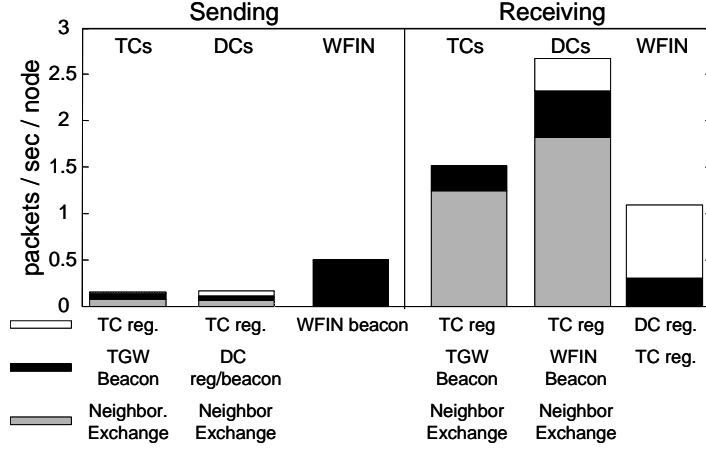


Figure 7.13: Rates and types of packets that are sent and received. These packet types are used to maintain the client information.

7.3.5.1 Client information maintenance

We begin by examining the types of overhead. Figure 7.13 shows the per node overhead to maintain the client information for $T_{WFINBeacon} = 0.5$ sec., $T_{NB} = 15$ sec., $T_{TGWBeacon} = 20$ sec. These results were generated from the two building scenario, but provide a good approximation into the other scenarios. As expected, little overhead results from the TC registration, but significant overhead is due to the TGW beacons and neighborhood exchanges.

Figure 7.14 shows the relationship between several parameters and the mean time for a node to become reachable (i.e., register) after entering into a tendril from being a direct client. The case where $T_{NB} = 15$ sec., $T_{WFINBeacon} = 0.5$ sec., should be compared to the analytical result shown in Figure 7.12.

The parameter values should be selected to achieve a trade-off between the overhead and the average duration until a node registers after entering a tendril. It is important to note that nodes move from being a direct client to being a TC relatively infrequently (indeed, once a node becomes a TC, it must remain a TC




Figure 7.14: Mean time for a client node to register after entering a tendrill. In these simulations, $MAX_IN_MISSED_BEACONS=6$ and $k = 3$. The values of the other parameters are as shown. The two building scenario was used in these simulations.

until either it has verified that it is no longer reachable from its TGW and is reachable from the WFIN, or the node has a strong connection with the WFIN for a sufficiently long time). Thus, during the course of the day, a node is only unreachable for a small fraction of time. Hence, in order to reduce overhead, it may be desirable to allow extended periods where the node is unreachable. Nonetheless, we select $T_{NB} = 15$ sec., $T_{WFINBeacon} = 0.5$ sec., $T_{TGWBeacon} = 20$ sec., and $MAX_IN_MISSED_BEACONS=6$.

7.3.5.2 Route Search/Paging Overhead

The main objective of the above algorithms is to find clients such that the overhead remains bounded as the size of the network increases. Figure 7.15 shows the overhead for performing a single page/route search using the algorithm presented above as compared to the route search technique used in AODV. In the case of




Figure 7.15: Comparison of overhead for performing a single page using the efficient paging algorithm and route search technique used in AODV. Each building has 10 floors and is populated with 100 nodes.

AODV, it is assumed that nodes have no prior information as to the location of the desired node. Hence, AODV will flood the network when performing route search.

As shown in Figure 7.15, the overhead for a single route search grows linearly with the number of buildings when AODV is used, while it remains constant when the scalable paging described here is used. Note that in the case of one building, AODV and the scalable paging perform similarly. This is expected since this paging algorithm aims to flood a single building. (As mentioned above, the impact of efficient flooding is minimal since the node density is low).

Chapter 8

CONCLUSIONS AND FUTURE WORK

It can be reiterated that simulation is one of the main methods for evaluating protocol performance for wireless networks. We have developed an elaborate framework for realistic propagation simulation of urban mesh networks. A general framework for overall simulation of the urban mesh networks, has also been explored. All the developed models have been extended to work with preferred simulators such as NS2 and QualNet. These models and multitude of trace files generated using these models for different urban environments are available for download at [17]. The thesis is aimed at improving the quality of simulations for wireless mesh and mobile ad hoc networks operating in urban environments.

8.1 Propagation simulator

Issues related to propagation simulation for urban wireless networking have been explored. Also, the design of a propagation simulator has been discussed. This propagation simulator is integrated with a mobility simulator discussed in [55]. This simulation package, which is available for download [17], demonstrates that it is possible to include realistic propagation into simulation of urban wireless networks such as the ones being deployed in several cities. As has been shown, traditional simulators that use random way-point mobility and the 2-ray propagation model are not suitable for urban environments. It is important to note that the simulation strategies discussed here and implemented in the simulation tools focus on realistic simulation, not accurate prediction. Thus, they are suitable for understanding the

performance of urban wireless networking protocols, but not for planning specific networks. However, as computational resources increase, we expect that more complicated and accurate techniques will be integrated into the simulation packages, and will, perhaps, allow for accurate prediction as well.

8.2 Graphical properties of urban mesh and mobile ad hoc networks

The urban environment brings many new problems and possibilities. In Chapter 3 heterogeneity of an urban mesh graph is explored. Various graphical properties such as degree distribution, duration in and out of largest connected component and centrality are examined for indoor and outdoor nodes operating in two contrasting environments. It is shown that the behavior of nodes inside and outside is quite different. Efficient protocols for urban mesh networks will have to cope with this disparity in the quality of communication depending on the position of the nodes. For both inside and outside nodes, the graph properties depend on the spatial position of the node. For example, nodes in some areas are well connected while nodes in other areas others are not. We see that these areas of well connected nodes are along major pedestrian thoroughfares. Thus, it seems likely that these thoroughfares will act as high bandwidth data communication paths. For example, they may be a natural backbone.

8.3 Variation in propagation characteristics

Chapter 4 is dedicated to examining the channel variability in a crowded environment. While the variability of the channel is often attributed to the mobility of the receiver or transmitter, it is also possible that objects in the environment are partly responsible for the channel variability. This work verifies that hypothesis. It was found that the standard deviation of the received signal strength is affected by the density of the mobile objects in the environment. A four-parameter, diffusion process was developed that approximates the observed variability of the signal

strength. To fully characterize the system, transition probabilities are investigated and modeled. It was demonstrated that a realistic variation can be synthetically generated using the developed diffusion model. Thus this model can be incorporated into hi fidelity simulators and test beds alike to introduce realistic variations in channel.

8.4 Effect of interference on performance of 802.11b/g

Chapter 5 explores the effect of signal to interference ratio (SIR) on the packet success probability. Specifically this work explored the behavior of 802.11b/g in the interference limited regime through a large number of laboratory experiments. Two observations were made. First, in the scenarios considered, synchronization error plays a critical role in the performance. As a result, in many cases, lowering bit-rate will not improve tolerance to interference. The implications that this observation has on bit-rate selection was investigated. A second observation is that during a collision, the probability of a bit error is independent of the frame size. This differs from the noise limited case where the probability of bit-error grows exponentially with the frame size.

8.5 Simulation framework

Chapter 6 presented simulation framework for urban mesh networks. Various aspects including number of nodes to be used in the simulation have been discussed. These techniques include methods to realistically simulate propagation and mobility. While realistic propagation modeling is computationally expensive, the propagation matrix needs to only be computed once for each urban map. Based on the findings from urban planning and traffic engineering research community, realistic mobility models can be developed. It is evident that these models are far more realistic than the random waypoint open-space propagation models that are widely used in current simulations. One challenge in realistic simulation is to keep the usage complexity

low. The methods, models, and model parameters developed reduce the complexity of use while still maintaining realistic simulation.

8.6 Insights gained from realistic simulations

The impact of using the realistic propagation simulator and the insights gained from these simulations has been discussed in Chapter 7. Specifically the issues relating to coverage and topology were investigated. While large-scale urban mesh networks (LUMNets) are being deployed around the world, there remain substantial challenges related to coverage and scalability. The coverage in terms of regions and achievable bit rates are investigated. The disparity in coverage of indoor and outdoor nodes is evident from this study. As mentioned earlier successful protocols will have to cope with this disparity in coverage and quality of communication. The study also investigates the coverage criteria when the mobile nodes are allowed/not allowed to act as relays. It was found that allowing mobile nodes to act as relays adds very less to the coverage when the infrastructure node densities are high. On the other hand it was seen that when the infrastructure density is low, coverage can be enhanced if mobile nodes are permitted to act as relays. In this case, traditional mobility management techniques are not applicable. While coverage is one aspect that dictates the performance of a network, quality of coverage is the one that is of more interest. By quality of coverage, we mean the achievable bit rates. A brief investigation on the achievable bit rates was conducted for nine different infrastructure node density scenarios. It was found that the average bit rate that can be achieved for an infrastructure node densities of 300 m. (scenario 9 in 7.1 which is the planned density for Philadelphia as of December 2006.) is around 4 Mbps. By increasing the density of the infrastructure nodes, bit rates of 10 Mbps can be achieved. It was found in [93] that these bit rates do not necessarily translate into application performance. However further work is required to ascertain

this fact whether good coverage and high achievable bit rates translate into network performance.

Topology of node connectivity has a very strong influence on the kind of protocols that can be developed. It was shown that using realistic simulations gives rise to realistic topologies. This topology information has been used to develop a new protocol that operates with low overhead in such scenarios. In particular, paging or searching for a mobile node in urban mesh networks has been addressed. Flooding, which is used in MANET routing protocols cannot be used since it does not scale to the size of the large scale urban mesh networks being developed. This work presents a scalable technique to page nodes in such networks. The scheme leverages the topology that arises in urban areas. Specifically, the fact that wireless propagation is good outdoors, but greatly restricted indoors produces a topology with highly connected core that is composed of outdoor nodes, and weakly connected tendrils made up of indoor nodes. The scheme presented restricts the flooding to the tendril where the mobile node resides and thus is more energy efficient and has very low overhead when compared to the other popular MANET routing protocols.

8.7 Future work

While the propagation simulation discussed above is significantly more realistic than the 2-ray and the open space models commonly used in mobile wireless network simulation, there are several areas that require further effort. For example, as discussed in Chapter 2 and Section 2.2.7, propagation is time-varying, and is linked to the mobility of objects, specifically people and vehicles. Since the integrated mobility and propagation simulator include the locations of people and vehicles, these locations can be included into the propagation calculations. However, more effort is required to model the impact of people and other mobile objects. One approach is to incorporate stochastic models. In this case, the simulation strategy would be site-specific and mixed deterministic and stochastic. See [30] for mixed

deterministic and stochastic propagation modeling.

Modeling propagation from indoors to outdoors and visa versa is an area of ongoing research [46]. One difficulty is that accurate propagation through non homogeneous walls is difficult to compute. As shown in Sections 2.2.3 and 2.2.4, propagation is greatly influenced by the material. However, these sections only considered propagation through walls made from a single material, whereas most walls are made of several types of material. While [44] develops a technique to model propagation through complicated walls, it is computationally complex. Furthermore, even if more computational power is available, there remains the difficulty of determining the materials and wall structure of realistic urban buildings which adds to complexity of simulation.

Throughout the thesis, it was assumed that transmitter and receiver antennas are vertically polarized. Furthermore, we also assumed that they are ideal. While it is straightforward to include the model of an ideal dipole antenna, in many cases the antenna is not an ideal dipole and is not vertically positioned. For example, see [47] for an example where the angle of inclination the antenna played a critical role in connectivity. Further work is required to get realistic models of antennas.

Capacity estimation is an other area which can utilize these models. While coverage and achievable bit rates have been addressed, further work is necessary to estimate the capacity of urban mesh networks.

In the work involving the investigation of time varying nature of propagation in crowded environments, it was found that the model parameters are closely related to one another and appear to follow a deterministic relationship. Thus, the only independent variable is the variance of the channel gain, which seems to be highly correlated to the pedestrian density. This prompts for further investigation to determine the relationship between them. A single parameter (pedestrian density) diffusion model can be developed once the relationship between variance in signal

strength and the pedestrian density is established.

Another important aspect of channel variation is the variation induced by small movements of transmitter or receiver. Currently due to computational constraints the resolution of a position in propagation simulation is unreasonable to be less than 0.5 m for large urban areas such as downtown Chicago (2km x 2Km). It is necessary to investigate the variations induced by movements that are less than the minimum resolution. These models can be incorporated into the packet simulator in conjunction with mobility simulator which can provide much smaller resolution in terms of node movement by correlating with the propagation characteristics at a coarse resolution provided by the propagation simulator.

Signal to interference plus noise ratio is one of the factors that affects the channel quality to a large extent. Due to the deployment of large and dense mesh networks, it has been receiving a lot of attention in recent times. Our work on the characterization of interference's impact on the packet error probability gave a lot insights into the way it differs from that of the impact of SNR. While this work examined the effect of a single interferer, more work needs to be done to model the effect of SIR on performance of wireless mesh networking protocols. Also further work needs to be done on making the rate fall back mechanisms to cope up with the interference.

While there can be no limit to the extent to which the quality of simulations can be extended, this concluding section addresses some of the issues that are of interest in the near future.

BIBLIOGRAPHY

- [1] http://www.cisco.com/en/US/products/hw/wireless_/ps4555/products-data-sheet09186a00801ebc29.html.
- [2] BVSystems. <http://www.bvsystems.com/Products/WLAN/Yellowjacket/yellowjacket.htm>.
- [3] Geographic information system (GIS). <http://www.gis.com>.
- [4] OPNET. <http://www.opnet.com/>.
- [5] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *Proceedings of the 2004 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*, pages 121–132, New York, 2004. ACM Press.
- [6] I.F. Akyildiz, J.S.M. Ho, and Yi-Bing Lin. Movement-based location update and selective paging for PCS networks. *IEEE/ACM Transactions on Networking*, 4:629–638, 1996.
- [7] J. B. Andersen. UTD multiple-edge transition zone diffraction. *IEEE Transactions on Antennas and Propagation*, 45:1093–1097, 1997.
- [8] A. Bar-Noy and I. Kessler. Tracking Mobile Users in Wireless Networks. In *IEEE Transactions on Information Theory*, volume 39, pages 1877–1886, 1993.
- [9] A. Bar-Noy, I. Kessler, and M. Naghshineh. Topology-based Tracking Strategies for Personal Communication Networks. In *Mobile Networks and Applications*, volume 1, pages 49–96, August 1996.
- [10] Henry L. Bertoni. *Radio Propagation for Modern Wireless Systems*. Prentice Hall, New Jersey, 1999.
- [11] C. Bettstetter. On the minimum node degree and connectivity of a wireless multihop network. *MOBIHOC*, 2002.
- [12] C. Bettstetter and C. Hartmann. Connectivity of wireless multihop networks in a shadow fading environment. *MSWiM'03*, 2003.

- [13] Giuseppe Bianchi, Fabrizio Formisano, and Domenico Giustiniano. 802.11b/g link level measurements for an outdoor wireless campus network. In *WOW-MOM '06*, pages 525–530. IEEE Computer Society, 2006.
- [14] S. Bohacek. Achievable performance improvements provided by cooperative diversity. In *WiOpt*, pages 73–82. IEEE Computer Society, 2006.
- [15] S. Bohacek, A. Ilic, and V. Sridhara. On the predictability of link lifetimes in urban manets. In *WIOPT '05: Proceedings of the Third International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks*, pages 209–218, Washington, DC, USA, 2005. IEEE Computer Society.
- [16] S. Bohacek and V. Sridhara. The graphical properties of MANETs in urban environments. In *The Forty-Second Annual Allerton Conference on Communication, Control, and Computing*, 2004.
- [17] S. Bohacek, V. Sridhara, and J. Kim. UDelModels - Models for Realistic Simulation of Urban Mesh Networks, 2004. <http://udelmodels.eecis.udel.edu>.
- [18] S. Bohacek and P. Wang. Toward tractable computation of the capacity of multihop wireless networks. In *To Appear in: INFOCOM'07*, 2007.
- [19] US Census Bureau. Topologically integrated geographic encoding and referencing (TIGER). Available at <http://www.census.gov/geo/www/tiger/>.
- [20] A. L. Cavilla, G. Baron, T. E. Hart, L. Litty, and E. de Lara. Simplified simulation models for indoor MANET evaluation are not robust. In *Proc. SECON*, 2004.
- [21] Cisco Systems. Aironet 1240AG Series 802.11A/B/G Access Point. Available at: <http://www.cisco.com/en/US/products/ps6521/products-data-sheet0900aecd8031c844.html>.
- [22] City and County of San Francisco. Request for proposals 2005-19 - TechConnect: Community wireless broadband network. Available at http://www.sfgov.org/site/tech_connect_index.asp, 2005.
- [23] CITY OF MINNEAPOLIS. Request for Proposal to Provide Broadband IP Data Access Services, April 13 2005. http://www.ci.minneapolis.mn.us/council/2006-meetings/20060224/docs/Mpls2005broadband_RFP.pdf.
- [24] City of San Francisco Department of Parking and Traffic. Daily traffic volume. Available at <http://www.sfgov.org/site/uploadedfiles/dpt/VolumesWeb.pdf>.

- [25] M. V. Clark, K. K. Leung, B. McNair, and Z. Kotic. Outdoor IEEE 802.11 cellular networks: Radio link performance. In *IEEE International Conference on Communications (ICC)*, pages 595–599, New York, NY, 2002. IEEE.
- [26] Atheros Communications. AR5212. Available at: <http://www.atheros.com/pt/AR5002GBulletin.htm>.
- [27] J.C. Cox, J.E. Ingersoll, and S.A. Ross. An intertemporal general equilibrium model of asset prices. *Econometrica*, 53:363–384, 1982.
- [28] S. K. Das and S. K. Sen. A New Location Update Strategy for Cellular Networks and Its Implementation Using a Genetic Algorithm. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM 1997)*, pages 185–194, September 1997.
- [29] L. Devroye. *A Course in Density Estimation*. Birkhauser, Boston, 1987.
- [30] A. Domazetovic, L. J. Greenstein, N. B. Mandayam, and I. Seskar. Propagation models for short range wireless channels with predictable path geometries. *IEEE Transactions on Communications*, 53(7):1123–1126, 2005.
- [31] O. Dousse, F. Baccelli, and P. Thiran. Impact of interferences on connectivity in ad hoc networks. *IEEE INFOCOM'03*, 2003.
- [32] O. Dousse, P. Thiran, and M. Hasler. Connectivity in ad-hoc and hybrid networks. *INFOCOM*, 2002.
- [33] Jean-Michel Dricot and P. De Doncker. High-accuracy physical layer model for wireless network simulations in NS-2. In *IWWAN*, 2004.
- [34] European Cooperation in the Field of Scientific and Technical Research. Urban transmission loss models for mobile radio in 900 and 1800 MHz bands. Technical Report Revision 2, The Hague, 1991.
- [35] S. Floyd and V. Paxson. Difficulties in simulating the internet. *IEEE/ACM ToN*, 9:392–403, 2001.
- [36] S. J. Fortune, D. M. Gay, B. W. Kernighan, O. Landron, R. A. Valenzuela, and M. H. Wright. WISE design of indoor wireless systems: Practical computation and optimization. *Computational Science and Engineering*, 2(1):58–68, 1995.
- [37] F. Graziosi and F. Santucci. A general correlation model for shadow fading in mobile radio systems. *IEEE Communication Letters*, 3:102–104, 2002.

- [38] L. J. Greenstein, V. Erceg, and Y. S. Yeh M. V. Clark. A new path-gain/delay-spread propagation model for digital cellular channels. *IEEE Transactions on Vehicular Technology*, 46(2):477–485, 1997.
- [39] B. Grigelionis. On generalized z-diffusions. Technical Report Preprint 2000-24, Institute of Mathematics and Informatics, 2000.
- [40] M. Gudmundson. Correlation model for shadow fading in mobile radio systems. *Electron. Lett.*, 27:2145–2146, 1991.
- [41] A. Hac and Z. Zhou. Locating Strategies for Personal Communication Networks: A Novel Tracking Strategy. In *IEEE Journal on Selected Areas in Communication*, volume 15, pages 1425–36, Oct. 1997.
- [42] M. Hata. Empirical formula for propagation loss in land mobile radio services. *IEEE transactions on vehicle technology*, 29(3):317–325, 1980.
- [43] O. Heckmann, M. Piringer, J. Schmitt, and R. Steinmetz. On realistic network topologies for simulation. In *Proceedings of the ACM SIGCOMM Workshop on Models, Methods and Tools for Reproducible Network Research*, 2003.
- [44] C. L. Holloway, P. L. Perini, R. R. DeLyser, and K. C. Allen. Analysis of composite walls and their effects on short-path propagation modeling. *IEEE transactions on Vehicle Technology*, 46:730–738, 1997.
- [45] IEEE Standards. IEEE standards wireless zone. <http://standards.ieee.org/wireless/overview.html>.
- [46] M. F. Iskander and Yun Zhengqing. Propagation prediction models for wireless communication systems. *IEEE Transactions on Microwave Theory and Techniques*, 50(3):662–673, 2002.
- [47] S. Jadhav, T. X. Brown, S. Doshi, D. Henkel, and R. G. Thekkekkunnel. Lessons learned constructing a wireless ad hoc network test bed. *1st International Workshop on Wireless Network Measurements*, 2005.
- [48] W. C. Jakes. *Microwave Mobile Communications*. Wiley, New York, 1974.
- [49] Amit Jardosh, Elizabeth M. Belding-Royer, Kevin C. Almeroth, and Subhash Suri. Towards realistic mobility models for mobile ad hoc networks. In *MobiCom*, 2003.
- [50] A. Jazwinski. *Stochastic Processes and Filtering Theory*. Academic Press, New York, 1970.

- [51] P. Johansson, T. Larsson, N. Hedman, B. Mielczarek, and M. Degermark. Scenario-based performance analysis of routing protocols for mobile ad-hoc networks. In *MobiCom*, 1999.
- [52] A. Kamerman and L. Monteban. *WaveLAN-II: A High-performance wireless LAN for the unlicensed band*, pages 118–133. Bell Lab Technical Journal, Summer 1997.
- [53] J. Kim. "realistic mobility modeling and simulation for mobile wireless network in urban environments". Master's thesis, University of Delaware, 2005.
- [54] J. Kim and S. Bohacek. Exploiting multihop diversity through efficient localized searching with cdma and route metric-based power control. In *MSWiM '06: Proceedings of the 9th ACM international symposium on Modeling analysis and simulation of wireless and mobile systems*, pages 155–164, New York, NY, USA, 2006. ACM Press.
- [55] J. Kim, A. Ilic, and S. Bohacek. A survey-based mobility model of people for simulation of urban mesh networks. In *MeshNets'05*, July 2005. <http://www.eecis.udel.edu/~bohacek>.
- [56] J. Kim, V. Sridhara, and S. Bohacek. Realistic mobility simulation of urban mesh networks. *Submitted ELSEVIER AD HOC NETWORKS*, page Submitted, 2007.
- [57] S. J. Kim and C. Y. Lee. Modeling and Analysis of the Dynamic Location Registration and Paging in Microcellular Systems. In *IEEE Transactions on Vehicular Technology*, volume 45, pages 82–90, Feb 1996.
- [58] Andrzej Kochut, Arunchandar Vasan, A. Udaya Shankar, and Ashok Agrawala. Sniffing out the correct physical layer capture model in 802.11b. In *ICNP '04: Proceedings of the Network Protocols, 12th IEEE International Conference on (ICNP'04)*, pages 252–261, Washington, DC, USA, 2004. IEEE Computer Society.
- [59] R. G. Kouyoumjian and P. H. Pathak. A uniform geometrical theory of diffraction for an edge in a perfectly conducting surface. *Proceedings of the IEEE*, 62(11):1448–1461, 1974.
- [60] T. Kurner and A. Meier. Prediction of outdoor and outdoor-to-indoor coverage in urban areas at 1.8 GHz. *IEEE Journal on Selected Areas in communication*, 20(3):496–506, 2002.
- [61] L. Li, D. Alderson, W. Willinger, and J. Doyle. A first-principles approach to understanding the internet's router-level topology. In *SIGCOMM*, 2004.

- [62] G. Liang and H. L. Bertoni. A new approach to 3-d ray tracing for propagation prediction in cities. *IEEE transaction on Antennas and propagation*, 46:853–863, 1998.
- [63] Y. Liu, F. L. Presti, V. Misra, D. Towsley, and Y. Gu. Fluid models and solutions for large-scale IP networks. In *SIGMETRICS*, 2003.
- [64] L. Luo, V. Sridhara, S. Bohacek, and A. Sethi. Paging in large-scale urban mesh networks. *Submitted*, 2006.
- [65] MADWiFi. Multiband atheros driver for wifi. Available at: <http://madwifi.org>.
- [66] Stefano Marinoni and Hannu H. Kari. Ad hoc routing protocol performance in a realistic environment. In *Proceedings of the Fifth IEEE International Conference on Networking (ICN 2006)*. Conference Publishing Services, to appear.
- [67] R. K. Mosvick and R. B. Nelson. *We’ve Got to Start Meetings Like This!* Scott, Foresman and company, Glenview, Illinois, 1987.
- [68] W. Navidi and T. Camp. Stationary distributions for the random waypoint mobility model. *IEEE transactions on Mobile Computing*, 3:99–108, 2004.
- [69] NS2. Network simulator - ns-2. <http://www.isi.edu/nsnam/ns>.
- [70] NS2. Wireless and mobility extensions to ns-2. <http://www.monarch.cs.rice.edu/cmu-ns.html>.
- [71] U.S. Department of Labor Bureau of Labor Statistics. American time use survey (ATUS), 2003. <http://www.bls.gov/tus/>.
- [72] R. R. Panko and S. T. Kinney. Meeting profiles: Size, duration, and location. In *Proceedings of The 28th Annual Hawaii International Conference on Systems Science*, 1995.
- [73] Y. I. H. Parish and P. Müller. Procedural modeling of cities. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, 2001.
- [74] C. E. Perkins and E. M. Royer. Ad hoc on-demand distance vector routing. *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications*, pages 90–100, 1999.
- [75] Philadelphia Mayor’s Office. Government leaders teleconference - wireless: The 21st Century technology. Online, 2004.

- [76] Philadelphia Mayor's Office. Government leaders teleconference - wireless: The 21st Century technology, 2004. [http://www.phila.gov/mois/press/multi-media/pdfs/Wireless Philadelphia092004.pdf](http://www.phila.gov/mois/press/multi-media/pdfs/Wireless%20Philadelphia092004.pdf).
- [77] F. Pozzi and C. Small. Vegetation and population density in urban and suburban areas in the u.s.a. *Presented at the Third International Symposium of Remote Sensing of Urban Areas*, June 2002.
- [78] Proxim Wireless. ORiNOCO 11b/g Card. Available at: <http://www.proxim.com/products/cp/pc.html>.
- [79] B. Pushkarev and J. M. Zupan. *Urban Space for Pedestrians*. MIT press, 1975.
- [80] T. Rappaport. *Wireless Communications - Principles and Practice*. Prentice Hall, 2002.
- [81] Charles Reis, Ratul Mahajan, Maya Rodrig, David Wetherall, and John Zahorjan. Measurement-based models of delivery and interference in static wireless networks. In *SIGCOMM '06*, pages 51–62. ACM Press, 2006.
- [82] K. Rizk, R. Valenzuela, S. Fortune, D. Chizhik, and F. Gardiol. Lateral, full-3d and vertical plane propagation in microcells and small cells. In *48th IEEE Vehicle Technology Conference*, pages 998–1002, 1998.
- [83] K. Rizk, J. Wagen, and F. Gardiol. Two-dimensional ray-tracing modeling for propagation in microcellular environments. *IEEE Transaction of Vehicle Technology*, 46:508–517, 1997.
- [84] K. Rizk, J. Wagen, J. Li, and F. Gardiol. Lamppost and panel scattering compared to building reflection and diffraction. In *COST 259 TD*, pages 158–167, Turin, Italy, 1996.
- [85] K. Rizk, J. F. Wagen, and F. Gardiol. Influence of database accuracy on two-dimensional ray-tracing-based predictions in urban microcells. *IEEE Transactions on Vehicle Technology*, 49:631–642, 2000.
- [86] N. C. Romano and J. F. Numamaker. Meeting analysis: Findings from research and practice. In *Proceedings of the 34th Hawaii International Conference on Systems Science*, 2001.
- [87] C. Rose and R. Yates. Minimizing the Average Cost of Paging Under Delay Constraints. In *Wireless Networks*, 1995.

- [88] S. V. Savov and M. H. A. J. Herben. Modal transmission-line modeling of propagation of plane radiowaves through multilayer periodic building structures. *IEEE Transactions on antennas and propagation*, 51(9):2244–2251, 2003.
- [89] Scalable Network Technologies. The QualNet simulator. <http://www.scalable-networks.com/>.
- [90] S. Y. Seidel and T. S. Rappaport. A ray tracing technique to predict path loss and delay spread inside buildings. In *Global Telecommunications Conference*, pages 649–653, 1992.
- [91] V. Sridhara and S. Bohacek. Realistic propagation simulation of urban mesh networks. *Accepted for publication: COMNETS*, 2007. Available at <http://udelmodels.eecis.udel.edu>.
- [92] V. Sridhara, J. Kim, and S. Bohacek. Models and methodologies for simulating mobile ad-hoc networks. In *The 3-rd IEEE International Workshop on Mobility Management and Wireless Access*. MobiWac’05, June 2005.
- [93] V. Sridhara, J. Kim, and S. Bohacek. Performance of urban mesh networks. In *MSWiM ’05: Proceedings of the 8th ACM international symposium on Modeling, analysis and simulation of wireless and mobile systems*, pages 269–277, New York, NY, USA, 2005. ACM Press.
- [94] V. Sridhara, H-C. Shin, and S. Bohacek. Observations and models of time-varying channel gain in crowded areas. In *The Third International Workshop on Wireless Network Measurement (WiNMee’06)*, Boston, Massachusetts, USA, April 2006. WinMee.
- [95] C. Steger, P. Radosavljevic, and J. P. Frantz. Performance of IEEE 802.11b wireless LAN in an emulated mobile channel. In *The 57th IEEE Semiannual Vehicular Technology Conference (VTC 2003-Spring)*, pages 22–25, 2003.
- [96] Taipei Times. Ma says taipei to be wireless in two days, December 19, 2005. <http://www.taipeitimes.com/News/biz/archives/2005/12/19/2003285096>.
- [97] Transportation Research Board. *2000 Highway Capacity Manual*. National Research Council, Washington, D.C., 2000.
- [98] R. A. Valenzuela and L. J. Greenstein. Performance evaluations for urban line-of-sight microcells at 900 MHz using a multi-ray propagation model. In *Global Telecommunications Conference GLOBECOM*, pages 1947–1952, 1991.
- [99] R. Vijayan and J. M. Holtzman. Foundations for level crossing analysis of handoff algorithms. In *ICC*, 1993.

- [100] D. M. Le Vinc and M. A. Karam. Dependence of attenuation in a vegetation canopy on frequency and plant water content. *IEEE Transactions on Geoscience and Remote Sensing*, 34(5):1090–1096, 1996.
- [101] WaveCall. WaveSight. <http://www.wavecall.com/>.
- [102] B. Williams and T. Camp. Comparison of broadcasting techniques for mobile ad hoc networks. In *Proceedings of the ACM International Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC)*, pages 194–205, 2002.
- [103] WirelessValley. LANPlanner. <http://www.wirelessvalley.com>.
- [104] M. Wist. Office space - how much is enough? Technical report, Gerald Eve, 2001.
- [105] M. Yarvis, K. Papagiannaki, and W. S. Conner. Characterization of 802.11 wireless networks in the home. In *1st Workshop on Wireless Network Measurements (WiNMee)*, 2005.
- [106] Y. Yi, M. Gerla, and T-J. Kwon. Efficient flooding in ad hoc networks using on-demand (passive) cluster formation. *Mobihoc*, 2003.
- [107] X. Zhang, J. G. Castellanos, and A. T. Campbell. P-MIP: Paging in mobile IP. In *Proceedings of the 4th ACM International Workshop on Wireless Mobile Multimedia*, 2001.
- [108] Y. Zhang, L. E. Owen, and J. E. Clark. A multi-regime approach for microscopic traffic simulation. *The 77th TRB Annual Meeting*, 1998.