

ADVANCED SYSTEM DESIGN OF IN-BUILDING WIRELESS COMMUNICATION NETWORKS USING VENTILATION DUCTS

A DISSERTATION IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY IN ELECTRICAL AND COMPUTER ENGINEERING

SUBMITTED BY

BENJAMIN E. HENTY

TO THE FACULTY OF THE DEPARTMENT OF ELECTRICAL AND COMPUTER ENGINEERING
IN THE CARNEGIE INSTITUTE OF TECHNOLOGY

CARNEGIE MELLON UNIVERSITY
PITTSBURGH, PENNSYLVANIA, UNITED STATES OF AMERICA



DECEMBER 11, 2006

Copyright 2006 Benjamin E. Henty.
All Rights Reserved.

*This work is dedicated to all those who have pursued the cause of Science
through the ages and their efforts to push back the darkness.*

*Not least among those is Ralph Edward Batchelder,
whose keen mind and endless fascination has so inspired me.*

Abstract

This thesis describes techniques to successfully design an IEEE 802.11g wireless system that uses ventilation ducts to distribute the signals throughout the building. To describe these techniques we present a wide array of measurements that have been performed on test duct networks, but setup in a laboratory and in buildings. We present trials of different duct modifications and antenna excitations we have tried, chronicling how to best install wireless local area network equipment in ventilation ducts. Using the installation techniques we developed, we present measurements of the pathloss through different duct networks in two different buildings in Stavanger and Oslo, Norway. Using our measurements we develop a model of propagation that predicts power levels inside a building based on only the duct locations. A rigorous calculation of propagation through the ducts would require detailed modeling of mode scattering at each component in the path. Since our goal was to develop reasonable prediction tools that do not require burdensome collection of detailed information about duct structure, we have chosen to develop an empirical model with parameters based on a physical understanding of the system. By describing this model and the accuracy of its use, we show that the model can be used reliably for designing a ventilation duct wireless distribution system. In fact, we present the design and performance of a building with a fully functional ventilation duct based IEEE 802.11g network. Our design covers a three story, 2000 square meter building. We fully characterize this design both with signal strength measurements and throughput performance measurements. As part of these measurements we present a simple technique for estimating the throughput present at particular locations in a building.

In addition to describing how to design a ventilation duct based wireless system, we explore several techniques for improving channel statistics and capacity in such a system. We present a novel antenna array that can be used to excite two different, single modes in a common type of ventilation duct. We also present some antenna array diversity measurements that show diversity can be an effective technique for improving the performance of IEEE 802.11b communications transmitted through a ventilation duct. Further, we consider an extensive array of antenna array techniques utilizing time-reversal focusing methods. We first demonstrate that time-reversal can be used at microwave frequencies, something not known prior to this work. We then show that time-reversal focusing and time-reversal nulling techniques can be effective in ventilation ducts, though

some techniques, antennas and duct configurations work better than others. We present the theoretical system capacity of some of these antenna and duct configurations. Lastly, as evidence that the time-reversal techniques are not just smoke and mirrors, we demonstrate both eigenvector diagonalization based MIMO and time-reversal based communications through ventilation duct channels. These demonstrations use actual frequency shift keying and IEEE 802.11g data transmissions to show that the techniques can substantively improve ventilation duct communications.

Acknowledgments

This research would not have been possible without the support of the ABB Corporation, the YIT Corporation, the National Science Foundation and the Defense Advanced Research Project Agency under the Army Research Office. I'd like to acknowledge and thank Anders Bredesen, Arne Malonæs, Dagfin Brodtkorb, Odd Drage, Oddvar Hetland, Astrid Nilsen, May Huth, Linda Keim, Geir Asbjornsen, Leif-Arne Amundsen, Stig Riise, and many others at YIT and ABB, in Norway and elsewhere in Scandinavia. Their hospitality, helpfulness and kindness is beyond belief and exceedingly appreciated. I am also amazed at their tolerance for an odd American student's quirks and shortcomings.

I can not possibly overstate the impact that Dan Stancil has had upon both this work and upon me. I've learned tremendously as a student working with him and benefited greatly from his friendship. I count myself very lucky to have met Professor Stancil so early in my academic career and to have worked with him for so long. I continue to be impressed by his dedication to scientific advancement and his considerate disposition. It is with an unquantifiably-large appreciation that I would like to thank him.

I would also like to thank my thesis committee for the time and consideration that each member has given to this work. The valuable advice and feedback provided by Carnegie Mellon Professors Rohit Negi and Elias Towe and Georgia Tech Professor Gregory Durgin is extremely appreciated and has significantly improved this work.

I would like to acknowledge the many helpful discussions with José Moura, Jian-Gang (Jimmy) Zhu, Ozan Tonguz, Yuanwei Jin, Yi Jiang, Kevin Borries, Ahmet Kaya, Ariton Khafa, Pavel Nikitin, Ratish Punnoose, Damon Becknel and Glenn Judd. The value of another thoughtful person to bounce ideas off of has been exceedingly helpful.

I'd also like to thank Paul Barter, Zenovy Mogilevsky, Rudy Ruelas, Robert Pulse and many others at Innovative Building Concepts and i5wireless. Their interest in this work has buffered my faith in the commercial viability of much of the material in this thesis and aided me in ensuring that the work has some practical application.

I'd like to thank my co-workers and friends who have contributed to my sanity and productivity, especially Ahmet Gurkan Cepni, JP Van't Hof, Jessica Hess and Nancy Dean. Many a long hour in the office was made enjoyable by their companionship. I would also like to thank Alan Grupe, Lou Anschuetz, Jack George, the Pennsylvania Department of Transportation, Mark Robertson, and Jim Martin for assistance with things I could not

have gotten to on my own.

I'd also like to thank the huge number of friends who have helped me over the long span of years that have led up to this work. It is impossible to successfully enumerate them all, but some of those foremost on my mind are Charles Lepple, Jonathan Betz, William Holtz, Megin Wardle, Jim Foraker, Chris Marshal, Hallie Parry, Samantha Lipson, Carla Geisser, Meg Richards, Aileen D'Esposito, Jen Smith, Audra Vest, Forest Godfrey, Chris Tuttle, David Gauthier, Colin Campbell, Feige Grundman, Gena Miller, Jason Aron, Jeff Lanza, Matt Silverstein, Krista Krahn, Zeb Krahn, Amanda Taratuski, Beth Meier, Shawn Maloney, Jeff Colletti, Mark Duchesneau, Matt Denton, Violet Blue, Adam Kramer, Maya Nigrosh, Adam Pennington, Alison Greenwald, Fred Merkle, Ros Neplokh, Lindsay Spriggs, Matt Sheby, Matt Theisz, Rob Siemborski, Ross Kinder, Susan Swithenbank, Will Ferry, Doctor Richard Kimbal, Rus Heywood, Sam Listopad, Hannah Listopad, Jeff Grafton, Jess Hand, Jim Hays-Wehle, Kevin Miller, Mike Kasick, Peter Nelson, Brian Gold, Katherine Crawford, Mike Sokolsky, Roger Skidmore, Ted Rappaport, Susie Bright, Sam Abtek, and many, many others who I have neglected or mentioned elsewhere.

I'd like to especially thank my mother-bff, father-friend, sister-brother, brother-brother, doctor-brother, and sister-cousins. Your tremendous support, love and familial compassion were and are vital support for all that I do. I'd like to thank my maternal grandmother, for her stalwart humor and confidence in me. I'd like to thank my paternal grandparents for their fascination and for their early inspiration, that still glows above me when I drift off to sleep each night (or morning). I'd like to thank my other family members who are remarkably kind and supportive, despite being stuck with me as a relative.

I'd like to thank you the reader for opening this work up in the first place. Without you, all this writing was for naught. I'd also like to thank you, you whose name I will not write here. I could not have done it without you.

Table of Contents

Title Page	i
Copyright Notice	ii
Dedication	iii
Abstract	iv
Table of Contents	viii
List of Tables	xii
List of Figures	xiii
Chapter 1 _____	
Introduction	1
Chapter 2 _____	
Physical Architecture	4
2.1 Introduction	4
2.2 Ventilation Duct System Characteristics	5
2.3 Investigations of Possible Duct Modifications	7
Testbed Duct Modifications	7
Screen and No Screen Comparison Measurements	12
Design Decisions	15
2.4 Radiating From Ventilation Ducts	16
Comparing Antenna Re-Radiators	16
Comparing AP Antennas	21
Design Decisions	22
2.5 Conclusion	24
Chapter 3 _____	
Propagation Measurements and Models	25
3.1 Introduction	25
3.2 RF Leakage from Ventilation Ducts	25
3.3 Building leakage	28
3.4 Signal Strength Measurements	29
Initial Measurements	29
Comparison of Duct System to Conventional System	31
3.5 Propagation Model	41
Description of the Model	41
Model Advantages	42

3.6	Computing Model Parameters	43
	Purely Empirical Model Fit	43
	Parameter Measurement Based Model	44
	Propagation Model Modifications for Range Limitations	53
3.7	Physical Interpretation of Propagation Model	53
3.8	Conclusions	58
Chapter 4		
	Throughput Measurements and Modeling	60
4.1	Introduction	60
4.2	Throughput versus Signal Strength Measurements	60
4.3	Throughput Improvements	65
	Antenna Diversity Improvements	65
	Bi-directional Amplifier Improvements	66
4.4	Full Building Implementation	68
4.5	Conclusions	72
Chapter 5		
	Channel Improvements via Antenna Arrays	78
5.1	Introduction	78
5.2	Single Mode Antenna Array	78
	Antenna Element Design	78
	Mode Attenuation Analysis	85
	Antenna Array Testing	86
5.3	Antenna Array Combination Diversity	89
5.4	Conclusion	94
Chapter 6		
	Channel Capacity Improvements via Time-Reversal and MIMO	97
6.1	Introduction	97
6.2	Demonstrating Time-Reversal Focusing at RF	98
	Time-Reversal Techniques	98
	Experimental Time-Reversal Setup	99
	Experimental Time-Reversal Results	101
6.3	Capacity Benefits from Time-Reversal Focusing	106
	Time-Reversal Focusing Simulation	106
	Duct Channel Measurements	109
	Calculated Time-Reversal in Measured Channels	112
6.4	Capacity Benefits from Interference Nulling	115
	Interference Nulling Techniques and Simulations	117
	Calculated Time-Reversal Nulling in Measured Channels	125
	Single Frequency Based Time-Reversal Nulling	128
	Comparing Nulling Techniques	132
6.5	Theoretical Capacity Available via MIMO Arrays	134
6.6	Conclusion	143

Chapter 7	
Demonstrating Time-Reversal and MIMO Communications	144
7.1 Introduction	144
7.2 Vector Modulator Based MIMO Communications	145
7.3 Time-Reversal MIMO Enabled IEEE 802.11g	148
Experimental Setup	148
MISO Time-Reversal Focusing	150
MIMO Time-Reversal Focusing	152
7.4 Conclusion	156
Chapter 8	
Conclusion	158
8.1 Summary of Contributions	158
8.2 Future Work	160
Appendix A	
System Layout Software	162
A.1 Basic Interface	162
A.2 Building Image Import Procedure	166
A.3 Duct System Structure Modeling	168
Locating Ventilation Ducts	168
Locating APs and Re-radiators	172
Locating Building Perimeter	174
Adding Notes	174
Completed Building Model	175
A.4 Measurements	175
Auto-numbered Points	178
Wireless NIC Measurements	178
AP Reported RSSI	178
Throughput Measurements	182
YellowJacket Measurements	182
A.5 Simulations	184
Setting Model Parameters	186
Displaying Simulations	188
Ideas for Simulating Transmit Rate and Delay Spread	193
Using Simulation Plots with Measurements	195
A.6 Data Analysis Plots	195
Throughput versus RSSI and Delay Spread Plots	195
RSSI versus Duct Distance Plots	197
RSSI versus RSSI Plots	197
Appendix B	
Load Balancing Software	200
B.1 Software Functions	200
B.2 Software Design	202
B.3 Speed Improvements	203

References

205

List of Tables

2.1	Performance metrics of unmodified testbed duct. Each table entry represents a single measurement.	12
2.2	Performance metrics of testbed duct with RF blocking screens installed. Each table entry represents a single measurement.	12
3.1	Computed parameters using empirical fit to propagation model. The fits were based on 98 different measurements in each case.	43
3.2	Error statistics for the five measurement cases using the modified prediction model. The statistics shown were calculated based on predictions at 98 locations for which power measurements were made in each of the 5 cases. ..	47
3.3	Error statistics for modified propagation model accounting for restricted propagation range. The statistics shown were calculated based on predictions at 98 locations for which power measurements were made in each of the 5 cases.	54
3.4	Final propagation model parameters. α , L_{in} and L_{out} were computed as a fit to 27 measurement points in the Oslo building. The remaining model parameters were chosen based on physical meaning and the quality of fit to the measurements made at 98 locations in each of the 5 cases in the Stavanger building.	59
4.1	Measured throughput at two measurement locations for various antenna diversity settings. Each table entry is based on a single measurement of a 1 Megabyte file transfer.	65
4.2	Throughput and RSSI measurements with and without a bi-directional amplifier installed. Each RSSI measurement is based on the average of about 20 measurements and every throughput entry in the table is based on a single 60 second measurement.	67
4.3	Reported maximum delay spreads tolerated at different transmit rates of a Cisco AP, and a Cisco Network Interface Card (NIC).	68
5.1	IEEE 802.11b signal strength and transmission rate measurements for antenna array phase shift measurements of U-shaped ventilation duct setup. Each entry in the table is based on a single measurement.	93
6.1	Measured channel isolation and resulting bit error rate in two phase-conjugate channels.	106

7.1	Measured gains and BERs of MIMO based channels. Each entry in the table is based on a single measurement.	147
7.2	Throughput measurements and channel statistics for 4 by 1 antenna array with equal amplitude and time-reversal coefficients. Each entry in the table represents a single measurement.	152
7.3	Throughput measurements of IEEE 802.11g transmission performance using MIMO time-reversal arrays. All data is presented in Mbps. Each table entry represents a single measurement.	156

List of Figures

2.1	A typical ventilation duct with a neck of the type used in installations.	6
2.2	An artistic rendering of the pressure valves used to control the airflow through ventilation duct outlets. Air enters the valve from the upper left and exits through a mesh of plastic slots to the right of the figure. This allows the air to leave the duct while maintaining the air pressure required inside the duct.	7
2.3	The most typical KlimaTak acoustical panels shown here in a hotel room. The ventilation duct is run behind the white sheet metal panels.	8
2.4	Another common KlimaTak acoustical panel. The ducts are hidden by the acoustical paneling and a drop ceiling all around.	9
2.5	This picture shows a larger KlimaTak acoustical panel setup. These allow more space for infrastructure above the ceiling but otherwise are similar to the panels shown in Figure 2.3.	10
2.6	These are a less common type of KlimaTak panel which is similar to a flat ceiling tile. The panels are made of the same perforated sheet metal but are incorporated into a standard drop ceiling.	11
2.7	Overhead diagram of duct layout of 315 mm diameter circular duct used to compare duct modifications with locations of pressure valve outlets. The end of the duct on the right side is a reduction in diameter that continues back to the air supply equipment.	11
2.8	Schematic of duct section used in experiment.	13
2.9	A ventilation duct with a neck of the type used in the measurements. A metal screen is shown as it was installed in the measurements.	13
2.10	Measured frequency response (a) and computed impulse responses (b) for channel measurements of laboratory ventilation duct setup.	14
2.11	Channel impulse response and computed propagation delays of propagating modes.	15
2.12	A simple half-wavelength re-radiator. It is normally inserted into a ventilation duct to allow RF signal to escape the duct.	16
2.13	A more complicated duct re-radiator system. The upper antenna is inserted into the ventilation duct while the lower antenna is mounted between two metal KlimaTak acoustical panels in the ceiling and a short run of coaxial cable is run between the two.	17
2.14	A picture of the panel re-radiator in which a monopole is wedged between two KlimaTak panels.	18

2.15	A Berkeley Varitronics' YellowJacket IEEE 802.11b scanning receiver of the type used for measuring received signal strength in the re-radiator comparison measurements.	19
2.16	This figure shows a comparison of conventional, duct based, and duct with re-radiator techniques.	20
2.17	Leakage comparison measurements of signal strength using a monopole antenna connected to an AP.	22
2.18	Leakage comparison measurements of signal strength using a dipole antenna connected to an AP.	23
3.1	Setup for duct leakage measurements.	26
3.2	Attenuation of signals leaked out of ventilation ducts due to screens placed over louver air outlets.	27
3.3	Loss in signal strength due to closed KlimaTak panels.	28
3.4	Measurement setup of power leakage outside of a building due to a conventional or duct installed AP.	29
3.5	Comparison of power leakage outside a building from conventional or duct based wireless AP.	30
3.6	First measurement set of ventilation duct based propagation. The AP is connected in the center of the building to the duct network highlighted in white. The signal strength is plotted in dBm.	32
3.7	Second measurement set of ventilation duct based propagation. The AP is in the center of the building and re-radiator was added as shown by the pink dot in the bottom right. The signal strength is plotted in dBm.	33
3.8	Third measurement set of ventilation duct based propagation. The AP is connected in the center of the building to the duct network highlighted in white. The signal strength is plotted in dBm.	34
3.9	Fourth measurement set of ventilation duct based propagation. The AP is connected in the right side of the building to the duct network highlighted in white. The signal strength is plotted in dBm.	35
3.10	Fifth measurement set of ventilation duct based propagation. The AP is connected in the right side of the building to the duct network highlighted in white. The signal strength is plotted in dBm.	36
3.11	Comparison of conventional and duct based coverage areas for an AP installed in the same location in the same building. The measurements are plotted by interpolating the colors between each measurement point.	37
3.12	A second comparison of conventional and duct based coverage areas for an AP installed in the same location in the same building. The measurements are displayed by interpolating the colors between each measurement point. ..	38
3.13	Plot of the difference in signal strength between conventional and duct based propagation using the data presented in Figure 3.11.	39
3.14	Plot of the difference in signal strength between conventional and duct based propagation using the data presented in Figure 3.12.	40
3.15	Blueprint of measurements in the Oslo building. The measured signal strength values are shown as points in the figure labeled in dBm.	45
3.16	Measured signal strength values from Figure 3.15 plotted versus computed distance through ventilation duct to AP.	46

3.17	Comparison of case 1 measurements to simulated signal strength. A visual comparison is shown in (a) by plotting the measurement points on top of a colored background based on simulated signal strength. The same points are plotted in figure (b). (a) allows us to see where points vary the most while (b) is more useful for seeing how much variation in prediction and measurement is present.	48
3.18	Comparison of case 2 measurements to simulated signal strength. A visual comparison is shown in (a) by plotting the measurement points on top of a colored background based on simulated signal strength. The same points are plotted in figure (b). (a) allows us to see where points vary the most while (b) is more useful for seeing how much variation in prediction and measurement is present.	49
3.19	Comparison of case 3 measurements to simulated signal strength. A visual comparison is shown in (a) by plotting the measurement points on top of a colored background based on simulated signal strength. The same points are plotted in figure (b). (a) allows us to see where points vary the most while (b) is more useful for seeing how much variation in prediction and measurement is present.	50
3.20	Comparison of case 4 measurements to simulated signal strength. A visual comparison is shown in (a) by plotting the measurement points on top of a colored background based on simulated signal strength. The same points are plotted in figure (b). (a) allows us to see where points vary the most while (b) is more useful for seeing how much variation in prediction and measurement is present.	51
3.21	Comparison of case 5 measurements to simulated signal strength. A visual comparison is shown in (a) by plotting the measurement points on top of a colored background based on simulated signal strength. The same points are plotted in figure (b). (a) allows us to see where points vary the most while (b) is more useful for seeing how much variation in prediction and measurement is present.	52
3.22	Schematic description of propagation model.	55
3.23	Power losses in a duct segment with a hole in it.	56
3.24	Power losses if we consider the duct broken up into many small duct segments each leaking a small amount proportional to the length of the segment.	57
4.1	Mist plot of measured throughput for Oslo measurements.	61
4.2	Plot of measured throughput versus AP reported (uplink) signal strength	63
4.3	Plot of measured throughput as points on top of a colored contour plot of the predicted throughput. The contours are drawn at 20, 16, 12, 8, and 4 Mbps levels.	64
4.4	Plot of best signal strength values measured on the first floor of the building. The outline of each square indicates the AP which had the best signal at each location.	70
4.5	Plot of best signal strength values measured on the second floor of the building. The outline of each square indicates the AP which had the best signal at each location.	71
4.6	Plot of best signal strength values measured on the third floor of the building. The outline of each square indicates the AP which had the best signal at each location.	72

4.7	Plot of throughput values measured on the first floor of the building. The outline of each square indicates the AP which had the best signal at each location.	73
4.8	Plot of throughput values measured on the second floor of the building. The outline of each square indicates the AP which had the best signal at each location.	74
4.9	Plot of throughput values measured on the third floor of the building. The outline of each square indicates the AP which had the best signal at each location.	75
4.10	Throughput versus RSSI as measured across all three floors in full building layout.	76
4.11	Wireless data transmission rates versus RSSI as measured across all three floors in full building layout. Each bar indicates the range of transmission rates seen.	77
5.1	Calculation of \mathcal{P} and \mathcal{T} parameters for a monopole antenna in a 315 millimeter diameter cylindrical ventilation duct.	80
5.2	Predicted mode content for various length monopole antennas. Each plot is labeled with the length of the monopole antenna. All plots are calculated using software written by Pavel Nikitin.	81
5.3	Mode content of adjusted monopole calculated using software written by Pavel Nikitin.	83
5.4	Cross-sectional field patterns of the major modes excited by a monopole 205 mm long in a 315 mm diameter cylindrical duct. The solid lines indicate electric fields and the dashed lines indicate magnetic fields.	84
5.5	Diagram of how to use a 180 degree hybrid splitter to excite different single modes depending on which AP antenna port is used.	84
5.6	Mode content excited by a 0.209 meter long monopole in a 0.305 meter diameter cylindrical duct as calculated using software by Pavel Nikitin.	86
5.7	Attenuation curves for the two modes excited by the long monopole antenna array.	87
5.8	Measured frequency responses comparing a quarter-wavelength monopole in the top plot and a long monopole array in the bottom plot. Note that while the data was measured from 2.4 to 2.5 GHz, the ISM band only extends from 2.4 to 2.4835 GHz.	88
5.9	Power splitter and vector modulator setup for exciting two antennas with arbitrary phase offsets.	89
5.10	Measurement duct setup of phase shifting antenna array.	90
5.11	Contour plot of the frequency responses.	91
5.12	Contour plot of frequency responses for different phase shifts between antennas in dBm.	92
5.13	Plot of the measured peak-to-peak variation in the frequency response across the 2.4 to 2.5 GHz frequency range using an antenna array with particular phase offsets and the resulting IEEE 802.11b transmit rate achieved.	94
5.14	Plot of the measured peak-to-peak variation in the frequency response across the 16 MHz range in the center of IEEE 802.11b channel 1. The peak-to-peak variation as well as the IEEE 802.11b transmit rate achieved are plotted versus the particular phase offset applied to the antenna array elements.	95

5.15	Plot of the measured mean channel gain of the frequency response across the 16 MHz range in the center of IEEE 802.11b channel 1. The channel gain as well as the IEEE 802.11b transmit rate achieved are plotted versus the particular phase offset applied to the antenna array elements.	96
6.1	Graphic representation of numerical aperture of an antenna array, (a), and the enhancement of numerical aperture of an antenna array via salutary use of multipath, (b).	100
6.2	Schematic of laboratory in which experiments were performed. Note the significant amount of clutter and the metal plate obstructing the line-of-sight path. All furniture and the metal plate were removed in the empty lab case. .	100
6.3	Setup of apparatus for experiment. All cables were matched lengths. The vector network analyzer can be switched internally to measure from the target antenna to the array or vice versa.	101
6.4	Normalized, received power as measured in the area around the target location. Measurements were made manually every 5 cm at 81 grid points.	102
6.5	Comparison of the cross-range focusing in the empty and cluttered lab measurements. In addition, the calculated focusing in an ideal multipath free setup is shown. The full widths at half maximum for the curves are simulated: 77cm, empty: 90 cm, cluttered: 6 cm.	103
6.6	Focusing effect in the cross-range direction by multipath focusing for band of frequencies about focus frequency used to compute conjugate-phase coefficients.	103
6.7	Measured channel responses for different target locations. The legend indicates the measured target location and the intended target location.	104
6.8	Apparatus setup for bit error rate measurements of two phase-conjugate channels to two target locations separated by approximately 5 centimeters. ..	105
6.9	Time-reversal simulation scenario used for investigating a basic time-reversal algorithm.	107
6.10	Simulated signal received at the user location for a (a) normal transmission and a (b) time-reversal transmitted signal.	109
6.11	Computed channels heard by eavesdropping AP in the normal case (a) and with the time-reversal coefficients applied (b).	110
6.12	Ventilation duct setup for time-reversal measurement cases.	110
6.13	Duct network used for large scale MIMO measurements.	112
6.14	Setup for measurement based time-reversal focusing simulations.	113
6.15	Average received power relative to the total input power across all transmit antennas for normal and time-reversal channels. The simulation bars represent randomly generated Rayleigh faded data while the large network bars represent large duct measurements.	114
6.16	Standard deviation of the calculated received power using normal or time-reversal signals. The simulation bars represent randomly generated Rayleigh faded data while the large network bars represent large duct measurements.	115
6.17	Average received power relative to the total input power across all transmit antennas for normal and time-reversal channels for 5 small duct measurements.	116

6.18	Standard deviation values of the calculated received power using normal or time-reversal signals for 5 small duct measurements.	117
6.19	Simulation scenario for comparing different nulling techniques.	117
6.20	Diagram showing how triangle nulling works in the frequency domain.	120
6.21	Cumulative distribution function of received signal levels for different nulling algorithms.	124
6.22	Cumulative distribution function of received signal levels at user location when different nulling algorithms are applied to AP_2 antennas.	125
6.23	Cumulative distribution functions for different sized, statistically generated, Rayleigh channels.	127
6.24	Comparison of the mean change in channel power compared to the normally received power in the calculated received signals for different nulling techniques. The first four items are statistically generated channels and the last two are measured duct channels.	128
6.25	Comparison of the mean change in channel power compared to the normally received power in the calculated received signals for different nulling techniques.	129
6.26	Calculated single frequency based nulling in the small duct measured channel with the axial array of short monopoles, absorber endcaps and no duct hole.	130
6.27	Calculated single frequency based nulling in the small duct measured channel with the axial array of short monopoles, metal endcaps and no duct hole.	131
6.28	Calculated single frequency based nulling in the small duct measured channel with the axial array of short monopoles, metal endcaps and a duct hole.	132
6.29	Calculated single frequency based nulling in the small duct measured channel with the circumferential array of short monopoles, metal endcaps and no duct hole.	133
6.30	Calculated single frequency based nulling in the small duct measured channel with the array of long monopoles, metal endcaps and no duct hole.	134
6.31	Calculated single frequency based nulling in the 3 by 3 large duct measurements.	135
6.32	Calculated single frequency based nulling in the 6 by 6 large duct measurements.	136
6.33	Capacity CCDF curves for Rayleigh generated random matrices and for two small duct measurement sets. The two measurement sets are the base case of an axial array of short monopoles, no duct hole and metal endcaps compared to a duct with absorber endcaps.	138
6.34	Capacity CCDF curves for Rayleigh generated random matrices and for two small duct measurement sets. The two measurement sets are the base case of an axial array of short monopoles, no duct hole and metal endcaps compared to a duct with a circumferential array of short monopoles.	139
6.35	Capacity CCDF curves for Rayleigh generated random matrices and for two small duct measurement sets. The two measurement sets are the base case of an axial array of short monopoles, no duct hole and metal endcaps compared to a duct with a hole in it.	140

6.36	Capacity CCDF curves for Rayleigh generated random matrices and for two small duct measurement sets. The two measurement sets are the base case of an axial array of short monopoles, no duct hole and metal endcaps compared to a duct with an array of long monopoles.	141
6.37	Capacity CCDF curves for Rayleigh generated random matrices, the two large duct measurements and the small duct base case of an axial array of short monopoles, no duct hole and metal endcaps.	142
7.1	Setup of channel and equipment used in experiments.	146
7.2	Measured channel responses forming the H matrix	147
7.3	Measured MIMO based channel responses created using an eigenvector basis, X of $H(f_p)$	148
7.4	Measurement setup of time-reversal arrays for IEEE 802.11 measurements. .	150
7.5	Uplink and Downlink Throughput versus Signal Strength as Computed from VNA Measured Frequency Responses.	153
7.6	Frequency Responses for Antenna Array Excited by Equal Amplitude Coefficients, Single Frequency Coefficients and for Ideal Time-Reversal Coefficients.	154
7.7	Signal and interference plots for the 2 by 2 MIMO created communications channels in the ventilation duct network.	155
A.1	A new window that opens when the Duct Mapper software is launched from MATLAB®.	163
A.2	The Duct Mapper <code>File</code> menu.	164
A.3	The Duct Mapper <code>Import</code> submenu.	165
A.4	First two steps in building image import process checking with user whether or not to mirror and/or rotate the image.	166
A.5	Further steps to importing a building image into Duct Mapper software.	167
A.6	Duct Mapper display after importing a building blueprint image.	169
A.7	Locate submenu for adding, moving and deleting ventilation ducts.	170
A.8	Instructions displayed in the MATLAB® command window that instruct the user how to draw ventilation ducts in the figure.	171
A.9	Picture of the user adding new ducts to a Duct Mapper layout.	172
A.10	Query as to whether or not the user is sure about replacing all of the ducts in the existing drawing with those from another figure file.	173
A.11	An example building modeled in Duct Mapper.	176
A.12	<code>Measurement</code> menu in Duct Mapper software.	177
A.13	Auto-numbered measurement points displayed in Duct Mapper.	179
A.14	Entering transmit rate data in Duct Mapper.	180
A.15	Configuration menu items for AP based RSSI measurement using Duct Mapper.	181
A.16	Configuration menu items for throughput measurement using Duct Mapper.	183
A.17	YellowJacket measurements recording marker numbers.	185
A.18	YellowJacket measurements being imported to Duct Mapper.	186
A.19	Duct Mapper <code>Edit</code> menu used for setting simulation model parameters for the Basic RSSI Model.	187
A.20	Duct Mapper <code>Edit</code> menu used for setting simulation model parameters for the Junction RSSI Model.	189

A.21	Example of In Duct RSSI simulation.	190
A.22	A colored mist simulation of the predicted signal strength in a building.	193
A.23	A contour simulation of the predicted throughput in a building.	194
A.24	Plot menu in Duct Mapper.	196
A.25	Example plot of throughput versus RSSI.	197
A.26	Example plot of measured RSSI versus distance through nearest duct to best AP.	198
A.27	Example plot of measured wireless NIC measured RSSI (cardrssi) versus AP measured RSSI (aprssi).	199
B.1	Load Balancer immediately after startup.	201
B.2	Load Balancer after clicking the Update button to get a list of connected users and RSSI values.	202
B.3	Load Balancer after forcing a user to move from AP 1 to AP 2.	203
B.4	Structure of balance2layout.m program and associated scripts and interfaces.	204

The focus of this thesis is on an alternative idea for providing inexpensive, reliable wireless coverage inside a building. The idea is to use a building's ventilation ducts as a means to distribute wireless local area network (WLAN) signals. The enclosed metal ducts act as waveguides with low loss to distribute the signal throughout a building. The idea to use ventilation ducts as a means to distribute signals has been under active research for a number of years. The original idea was encapsulated in a patent written in 1997 [1]. Since that time, the idea has been further detailed [2, 3, 4]. At the same time, significant work has been undertaken to demonstrate the practicality of the idea [5]. To begin to understand how to excite the duct channel, research into antenna behavior and design [6, 7, 8] was undertaken and channel models of the duct environment [9, 10, 11, 12] were developed. Some early work even developed techniques to analyze the mode content in ventilation ducts [13, 14]. Researchers were excited to realize that significant information capacity is available in ventilation ducts [15, 16, 17] and that multi-carrier transmissions can be used successfully in duct channels [18]. So, to begin to focus the duct idea towards a commercial application, a simple empirical pathloss model for duct propagation was formed [19, 20] and concepts for handover [21, 22] were developed. Work has even been undertaken to investigate the idea's applicability for cellular signals [23]. It is in this phase of focusing on the commercial viability of ventilation ducts that the work in this thesis was begun.

Chapter 2 of this thesis presents our work into how to best connect IEEE 802.11g access points (APs) to duct networks and ensure that the signal in the duct is sufficiently radiated to building users in such a way that the signal strength and signal quality are as high as possible. In this chapter we investigate the use of two different types of antennas combined with absorber and screens as possible duct modifications. We uncover some surprising results as to how using metal screens to prevent mode scattering inside a ventilation duct can actually cause the performance of wireless signals to degrade. Instead, allowing power to leak out of the ducts can actually work to the advantage of the system. In addition, we measure the leakage from ducts into the building environment using different re-radiators.

In Chapter 3, we delve further into the leakage of wireless signals from ventilation ducts without the use of re-radiating antennas into both the indoor and outdoor areas of

the building. In this chapter we also perform extensive measurements of signal strength throughout areas of two different buildings using an IEEE 802.11g AP connected to different ventilation ducts. We also use these measurements to develop an empirical pathloss model of signals transmitted through ventilation ducts. This model and fits of the model parameters are described fully in this chapter.

We extend our signal strength measurements in Chapter 3 to throughput measurements in Chapter 4. The chapter's main focus is actually a testbed system installation of an IEEE 802.11g network of APs. The testbed system was designed using the empirical pathloss model developed in Chapter 3. The pathloss model allows us to design a system for a 2000 square meter, three story building. After predicting the coverage throughout the building, we also measured the signal strengths and throughput. These measurements allowed us to validate our work in Chapter 3 and to develop a formula for estimating the saturation throughput available to a user at a particular location in a building based on signal strength.

At this point, we also mention that significant custom software was developed during much of this thesis. The software, implemented in MATLAB[®], is described in detail in Appendices A and B. The software in Appendix A is useful for measurements, simulation and analysis of ventilation duct based wireless distribution systems. The software allows the user to create simple models of ventilation duct systems including where APs are connected. The software allows the user to simulate the expected signal strength and throughput that the AP will provide in the building. The software will also allow the user to record various signal strength and throughput measurements at different locations in the building. These measurements can be displayed in various forms or used to empirically determine parameters of the pathloss and throughput models developed in Chapters 3 and 4. The software described in Appendix B is a smaller software utility that interfaces with the particular Linksys WRT54G APs that we used throughout this thesis. The software allowed us low level control of which wireless network interface cards (NICs) were connected to which APs. The ability to control which user was connected to which AP also demonstrates the viability of load-balancing in a ventilation duct based wireless distribution system.

Once a full scale installation was shown to be viable, we turned our attention towards improvements to such a system. One particular avenue of investigation explored was the use of antenna arrays. Some of our initial antenna array measurements are covered in Chapter 5. In this chapter, we introduce a novel two element antenna array that is capable of exciting one of two particular single modes of duct propagation in a commonly sized ventilation duct. This novel array can be used to select either mode by simply driving the elements in phase or out of phase, which can easily be implemented without complex phase shifters. However, we also present measurements using an ordinary two element array using a complex phase shift. We utilize this array with IEEE 802.11b devices to

show that an antenna array can have a significant effect on a duct channel by allowing or preventing communications depending on the relative phasing of the array elements.

With the success encountered in using simple antenna arrays in Chapter 5, we turn our attention to more sophisticated antenna array techniques in Chapter 6. In the hopes of improving the link and system capacities, we have analyzed the performance of time-reversal focusing, time-reversal nulling and MIMO techniques for different ventilation duct configurations and antenna excitations. The focus of these measurements is to achieve some insight into how best to excite a ventilation duct to maximize the effect of different antenna arrays on capacity. However, as part of this work we also showed that the concept of time-reversal focusing could be implemented at microwave frequencies with electromagnetic waves. Previous work in the area had been limited to acoustics [24] and ultrasonics [25]. After showing that time-reversal is a valid technique at microwave frequencies, we go on to compare the computed time-reversal focusing and nulling coefficients in 5 small, laboratory configurations and two larger duct network configurations. In this chapter we also present a novel time-reversal-like nulling algorithm which we term the “triangle-nulling” algorithm. We describe this nulling algorithm as well as the performance of several time-reversal focusing and nulling algorithms in various duct configurations. In addition, we use the measured duct channels to compute the theoretical channel capacity bounds for MIMO arrays used with different ventilation duct arrangements.

Finally, in Chapter 7 we wish to produce a practical demonstration of some of the advanced techniques explored in Chapter 6. Thus, in this chapter we setup two by two MIMO communications in a ventilation duct. We use eigenvector diagonalization to create two MIMO communications channels and show that individual bit streams can be transmitted simultaneously over these two channels. In this chapter we also provide a practical demonstration of time-reversal focusing implemented with actual IEEE 802.11g devices. Thus, we configure use a four element antenna array with time-reversal focusing coefficients to focus on one wireless user through a ventilation duct channel. We also use the time-reversal focusing coefficients for the same four element array that would focus on two different wireless users simultaneously with different data streams to demonstrate a time-reversal enabled MIMO communications technique. Throughput measurements of the actual IEEE 802.11g communications provide a concrete validation of the time-reversal techniques for use with actual IEEE 802.11g hardware.

Lastly, we discuss the results of our work in Chapter 8. In this chapter we discuss the implications and major results of our work. We also mention future areas of research that appear promising.

2.1 Introduction

In this chapter, we will discuss the measurements and system design options that we considered in order to develop a system installation concept. The goal was to consider different antenna and duct modification options in order to provide the maximum coverage with the easiest installation. To achieve these goals, we relied heavily on prior work on the ventilation duct concept. Nikitin et. al. has published significant work on the area. In addition to his Ph.D. thesis [26], he has published a number of papers relevant to the work in this chapter. Nikitin et. al. designed the monopole and dipole antennas that we work with in this chapter using theory developed in [6, 7, 8]. The antennas used in this work were based on these designs and the theoretical understanding developed.

In addition to antenna designs, Nikitin et. al. developed a model of the duct channel impulse response in [11, 12]. In these papers, the three different causes of dispersion in ventilation ducts are described. In order from largest to smallest cause of dispersion, they are multipath, intermodal distortion and intramodal distortion. Multipath is due to metal end-caps, discontinuities in the duct size or shape, and duct junctions. At these discontinuities, power can be reflected and multiple echoes can be setup causing multipath phenomenon. The more echoes, the longer a channel's delay spread and the more channel dispersion. Intermodal distortion is due to different modes having different propagation speeds in a waveguide. The different speeds result in power arriving at delayed intervals and increasing dispersion. Intramodal dispersion is due to different frequencies in the same mode traveling at different speeds in the waveguide. This again results in power arriving at different times and increased dispersion. Throughout this chapter we consider different techniques for combating these different types of dispersion and consider how best to design a system taking them into account.

Lastly, Nikitin et. al. have developed a technique for analyzing mode content in a ventilation duct in [13, 14]. The techniques presented in this chapter are similar, but are intended to provide insight into the approximate relative mode content rather than the precise mode amplitudes. We also make use of the mode content calculated using Nikitin's techniques to better understand our measurements. Some of the results presented in this section have been presented in [27] but are presented here for clarity of the

topic matter.

2.2 Ventilation Duct System Characteristics

To facilitate the development of an installable indoor system for wireless LAN use, we have concentrated our research on duct systems similar to the “KlimaTak” product used by Building Systems division of the YIT Corporation in Scandinavia. These duct systems have many characteristics that make them well-suited for wireless signal distribution. The testbed building we used in this chapter was owned by YIT and was located in Stavanger, Norway. The building is 3 stories tall and about 2000 square meters of it were available for our measurements. The unique aspects of the KlimaTak duct system used in the building are outlined here.

HVAC (Heating, Ventilation, and Air Conditioning) Ducts have a number of characteristics that make them well suited for distributing wireless signals in the indoor wireless environment. Indoor channels present significant challenges for wireless systems. Users rarely have line-of-sight to access points and so the losses and scattering caused by walls and various building materials have significant effects on signal propagation that are not always easy to predict. In contrast, the metal HVAC ducts act as waveguides to distribute RF (Radio Frequency) signals with low losses. Similar to much more expensive leaky feeder systems, the ducts act to evenly disperse the wireless signal everywhere ventilation ducts are run. This can dramatically improve the ability to predict the covered areas of a building and allows for a system that will not change with building renovations that leave the duct system in tact.

In addition to the advantages of using HVAC ducts, several characteristics of KlimaTak ventilation ducts make them ideal for use with wireless signals. These are primarily:

- Few changes in duct size or cross-section shape
- Almost no internal partitions
- Plastic or otherwise RF transparent louvers
- Ducts run directly through offices

Each office has a ventilation outlet cut in the side of the duct which allows air to enter the office. An example of this is shown in Figure 2.1. The neck piece shown is ordinarily attached directly to the duct with screws but is temporarily mounted for measurements in the figure. The circular hole in the neck piece is usually a 0.2 or 0.25 meter diameter hole. The outlets are covered either with a cloth filter bag or with a controllable pressure valve which, although containing a few metal parts, is made primarily of plastic. An image of the pressure valve is shown in Figure 2.2. Both of the outlet coverings are reasonably transparent to RF signals. Since the system uses the filter bags or pressure valves to control how much air is let out of each air outlet, the ducts are primarily unobstructed

runs of spiral circular duct with few changes in diameter. The most common diameter ducts used are 0.315, 0.4 and 0.5 meters. Additionally, air directing fins and other metal internal partitions are rarely needed in the ducts.




Figure 2.1: A typical ventilation duct with a neck of the type used in installations.

Another unique feature of the KlimaTak duct system is the metal acoustic panels which hide the ducts from view. There are several types of these panels used, but they are all made of sheet metal with small perforated holes (roughly 1 mm in diameter) across the surface. Several pictures of these panels are shown in Figures 2.3, Figures 2.4, Figures 2.5 and Figures 2.6. These panels attenuate RF energy leaving the ventilation ducts and entering the building environment.

These are the most important features of the KlimaTak system as it relates to their use for RF distribution. While many features of the system make KlimaTak ducts ideal for use with RF signals, it is hoped that the insight and design concepts developed for this system will carry over to other duct networks.




Figure 2.2: An artistic rendering of the pressure valves used to control the airflow through ventilation duct outlets. Air enters the valve from the upper left and exits through a mesh of plastic slots to the right of the figure. This allows the air to leave the duct while maintaining the air pressure required inside the duct.

2.3 Investigations of Possible Duct Modifications

2.3.1 Testbed Duct Modifications

Initially, we considered launching and maintaining a single mode of propagation inside the ventilation ducts. To explore this concept we selected a section of duct inside the YIT building that could be used as a testbed. This testbed allowed us to try different duct modifications in an actual building but on a reasonably small scale. A diagram of the duct setup is shown in Figure 2.7. Measurements were made through the ventilation duct using monopole or dipole antennas. The pressure valve air outlets were used as normally installed and also with wire screens installed between the neck piece and the duct itself. These screens were installed in such a way as to maintain the curved sidewall shape of the ducts while providing a minimal blockage to airflow. We also tested both a standard flat metal endcap and one with radio frequency (RF) absorber placed on the inside surface to prevent reflections. All together, these three variations (antennas, screens and endcap) made for eight different measurement cases. By comparing the channel statistics and throughput we were able to evaluate the effects of the duct modifications.

To characterize the testbed duct, an Agilent E8358A Vector Network Analyzer (VNA) was connected directly to the antennas in the ducts. After calibrating to remove the effects of the cables, the frequency response of the channel from 2.4 to 2.5 GHz was measured for all eight of the cases. The pathloss of the channel was computed by averaging the channel loss across the 22 MHz frequency band used by a particular IEEE 802.11b/g



Figure 2.3: The most typical KlimaTak acoustical panels shown here in a hotel room. The ventilation duct is run behind the white sheet metal panels.

signal. Additionally, the peak-to-peak variation in the frequency response of the channel was calculated using all of the data from 2.4 to 2.5 GHz. For reference, 1601 points were recorded to ensure the response was sufficiently sampled to detect deep nulls in the response. The pathloss and peak-to-peak values were recorded and are presented in Tables 2.1 and 2.2.

After characterizing the frequency response of the channels, an IEEE 802.11g Access Point (AP) was connected to the duct. The AP was a hardware version 1.1 Linksys WRT54G. This AP has Reverse Polarity Threaded Neill-Concelman Connectors (RP-TNC) that allowed us to connect the AP to the duct antennas using a short jumper cable. Note that even though the AP has two antenna ports, only one was connected to the ventilation duct. The other was terminated with a matched 50 Ohm load. A laptop was connected to the AP via a twisted pair, Category 5E (Cat-5E) cable. Additionally, a Buffalo Wireless Air Station, WLI-CB-G54A IEEE 802.11g Personal Computer Memory Card International Association (PCMCIA) card was used with a second laptop. The wireless card had an



Figure 2.4: Another common KlimaTak acoustical panel. The ducts are hidden by the acoustical paneling and a drop ceiling all around.

external antenna connector that was connected to the duct antenna via a short jumper cable. Note that using the external antenna connector does not prevent direct radiation and reception by the wireless card, so signal leakage from the AP and from unscreened ventilation duct air outlets could reach the wireless card without going through the duct. Care was taken during the measurements to ensure that this leakage signal was at least 20 dB weaker than the signal received directly through the duct.

To measure the throughput, version 1.1.1 of an open source tool named “Iperf” was used [28]. The tool has a client and server function and can send test packets using the User Datagram Protocol (UDP) or the Transmission Control Protocol (TCP) over a standard internet protocol (IP) link. The software was used in TCP mode with a 1 Megabyte TCP socket buffer window size and measurements were performed for 60 seconds each. The server was run on the laptop connected via ethernet to the AP. Additionally, the transmit rate used by the wireless card was noted and the maximum and minimum rates used were recorded. The transmit rates and measured throughput values are given in megabits per second (Mbps) for all 8 cases in Tables 2.1 and 2.2.

Table 2.1 and 2.2 show a comparison of all measurements made of the testbed duct. As can be seen from the tables, the pathloss in all eight cases is quite low. As a result, the AP should be able to operate at the maximum transmit rate in all cases. Further,



Figure 2.5: This picture shows a larger KlimaTak acoustical panel setup. These allow more space for infrastructure above the ceiling but otherwise are similar to the panels shown in Figure 2.3.

some of the lowest pathloss cases have the lowest throughputs. As a result, we believe that modal behavior causes the variation in throughput. As mentioned previously, the original idea behind the measurements was to try to launch and maintain a single mode of propagation in the ducts. The monopole and dipole antennas used were originally designed by Nikitin et. al. [26]. The monopoles simply provide a good impedance match with the duct and radiate primarily 4 or 5 modes [26]. However, the dipoles are designed to provide a good match while exciting primarily one mode, TE_{01} , in cylindrical ducts 12 inches in diameter.

The TE_{01} mode is thought to be advantageous for use because in waveguides made of metals with finite conductivity, the loss per meter experienced by the TE_{0n} modes decreases as frequency increases [29]. Further, TE_{01} has a low cutoff frequency and so will have low intramodal dispersion [11]. Since only one mode is excited, there is no intermodal dispersion. Thus, as long as the single mode propagation can be maintained and reflections are minimized, dispersion in the duct channel should be minimal. This is the reason we installed metal screens and RF absorber inside the duct. The metal screens prevent mode conversion at the air vents and the absorber prevents reflections at duct endpoints. We thus hypothesized that the case where we use dipoles with screens and absorber installed that there will be the least channel dispersion and the best throughput.




Figure 2.6: These are a less common type of KlimaTak panel which is similar to a flat ceiling tile. The panels are made of the same perforated sheet metal but are incorporated into a standard drop ceiling.




Figure 2.7: Overhead diagram of duct layout of 315 mm diameter circular duct used to compare duct modifications with locations of pressure valve outlets. The end of the duct on the right side is a reduction in diameter that continues back to the air supply equipment.

Launching and maintaining a single mode and preventing reflections does seem to allow the ventilation ducts to perform better as seen in Table 2.2. However, comparing

Table 2.1: Performance metrics of unmodified testbed duct. Each table entry represents a single measurement.

	Dipoles		Monopoles	
	Absorber	No Absorber	Absorber	No Absorber
Transmit Rate	54 Mbps	54 Mbps	54 Mbps	48-54 Mbps
Throughput	24.5 Mbps	25.3 Mbps	26.1 Mbps	22.5 Mbps
Pathloss	-36.6 dB	-34.1 dB	-38.1 dB	-37.1 dB
Peak-to-Peak	17.4 dB	12.0 dB	13.7 dB	10.7 dB

Table 2.2: Performance metrics of testbed duct with RF blocking screens installed. Each table entry represents a single measurement.

	Dipoles		Monopoles	
	Absorber	No Absorber	Absorber	No Absorber
Transmit Rate	48-54 Mbps	11-18 Mbps	24-48 Mbps	11-18 Mbps
Throughput	20.8 Mbps	4.1 Mbps	11.7 Mbps	8.6 Mbps
Pathloss	-41.1 dB	-34.2 dB	-32.7 dB	-32.4 dB
Peak-to-Peak	20.7 dB	25.7 dB	14.9 dB	31.4 dB

Table 2.1 with Table 2.2 there also appears to be an unexpected behavior taking place. First, note that in all cases, the measured throughput was higher for all of the cases without screens installed than for the cases with screens installed. However, the case which provides single mode propagation using screens, absorber and dipoles did perform markedly better than the other cases with screens installed. We thus undertook additional experiments to attempt to explain the apparent benefit of not installing metal screens.

2.3.2 Screen and No Screen Comparison Measurements

To identify the mode scattering phenomenon in a ventilation duct T-junction, a laboratory duct setup as shown in Figure 2.8 was used. The setup had seven duct segments of 305 mm diameter, connected together in a straight run 21.3 meters long with open ends. Quarter-wavelength monopole probe antennas, tuned to the 2.4 to 2.5 GHz frequency band, were installed in the duct 20.6 meters apart and a network analyzer was used to characterize the channel response in three different cases. The antennas remained fixed over each of the three cases; only the second duct segment was varied. In the first case the duct segment was unmodified. In the 2nd and 3rd cases a circular hole was cut in the duct segment. In the 2nd case a metal neck, as shown in Figure 2.9, with a circular, 190 mm diameter hole (matching the hole cut in the duct) was placed over the hole to form a T-junction representative of the sort used in actual installations. In the 3rd case a metal

screen was installed between the metal neck and the duct segment. The metal screen was 18 gauge metal wire in a 6mm square grid. The screen was installed such that it maintained the curvature of the duct and made good electrical contact with the duct and neck.




Figure 2.8: Schematic of duct section used in experiment.




Figure 2.9: A ventilation duct with a neck of the type used in the measurements. A metal screen is shown as it was installed in the measurements.

The measured frequency responses for the 3 cases are shown in Figure 2.10a. As can be seen from the frequency response, the screen placed over the air outlet has nearly an

identical response to the unmodified duct case. Further, the loss in power is an average of 7.5 dB from the unmodified to open air outlet cases, in reasonable agreement with the 8.9 dB expected loss, reported in [19]. Using the amplitude and phase of the measured frequency response, the impulse responses of the 3 cases were also computed and are shown in Figure 2.10b.

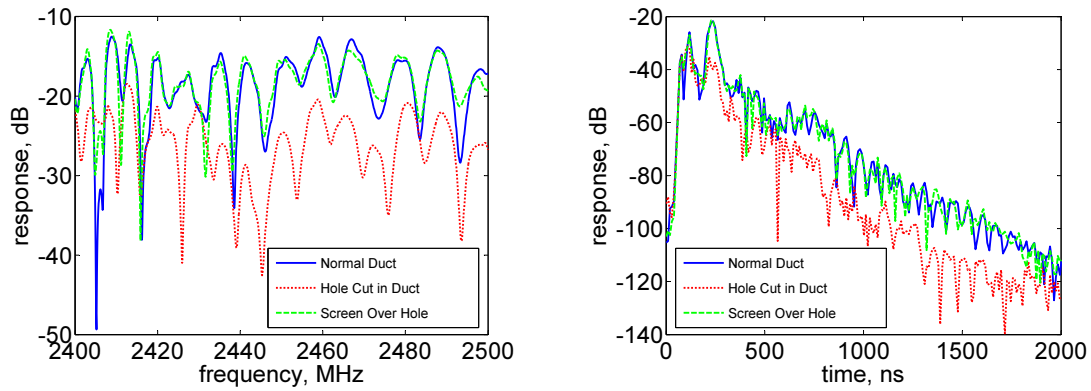


Figure 2.10: Measured frequency response (a) and computed impulse responses (b) for channel measurements of laboratory ventilation duct setup.

To analyze the mode content of the responses, the propagation delay of each of the 17 propagating modes was calculated using the theoretical group velocity for propagating modes in a cylindrical waveguide (see for example [29]) and then plotted in Figure 2.11 with an enlarged portion of the impulse response. The lower order modes appear to bunch up in the early part of the impulse response, while the higher order modes arrive at more spread out intervals. Our results are qualitatively consistent with mode excitation calculations reported in [29] where the most strongly excited modes were found to be TE_{61} , TE_{51} , and TE_{41} . We now wish to contrast the impulse response of the unmodified and screened air outlet cases with the open air outlet case and observe the mode filtering introduced by the hole cut in the duct. Interestingly, we note that there is little change in the early arriving power, but a significant decrease in the power of the late arriving modes is observed, particularly mode TE_{61} . This result is consistent with the hypothesis of [19] in which mode filtering results in an increased power loss at the first T-junction encountered due to mode filtering effects.

These measurements present the surprising result that allowing the natural mode filtering present in ventilation duct based wireless distribution systems due to ventilation outlets can be more beneficial than using metal screens to keep power within the ducts. These measurements thus explain the IEEE 802.11g throughput measurements.

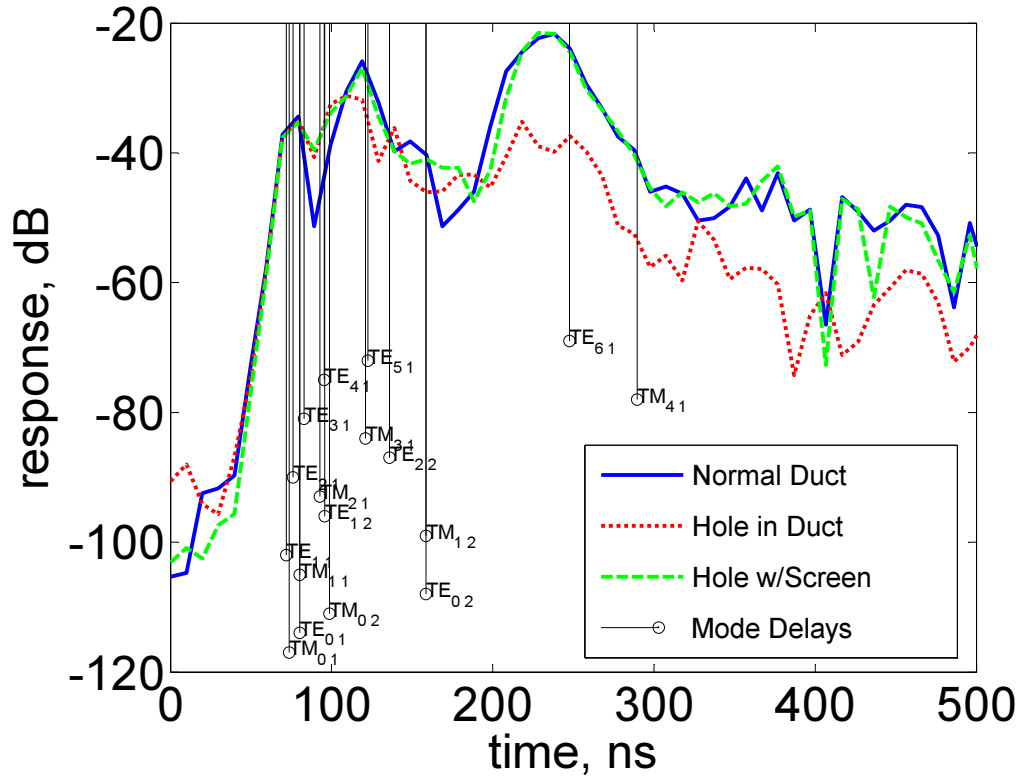


Figure 2.11: Channel impulse response and computed propagation delays of propagating modes.

2.3.3 Design Decisions

In light of the measurements presented in Sections 2.3.1 and 2.3.2, several recommendations for antenna and necessary duct modifications can be made. Firstly, we did not find any compelling reason to justify the extra effort and expense of installing screens and absorber in the ducts. In fact, the duct vents actually improve the duct propagation channel, making the use of screens undesirable. In regards to RF absorber, we note that monopoles without screens but with absorber did function somewhat better than without absorber. However, we are not confident that the difference in throughput between the two cases is significant due to the variable nature of throughput test measurements. For this reason, and because avoiding unnecessary complexity is preferred, we recommend that installations be made without RF absorber. If it seems that performance may be suffering from excessive reverberation, then absorber should perhaps be considered, but it is not anticipated to be needed. Finally, we saw in Table 2.1 that any combination of an-

tenna or absorber performed well in the unmodified duct. However, monopoles are easier to install because they can be installed in a duct simply by drilling a single small hole while dipoles require a large hole or slot to insert the antenna. For these reasons, it is recommended that systems use monopole antennas without metal screens or absorbers.

2.4 Radiating From Ventilation Ducts

2.4.1 Comparing Antenna Re-Radiators

After various duct modifications were explored as discussed in Section 2.3, it was decided that monopoles without metal screens would be used for further propagation measurements. Additionally, RF absorber would not be used initially, but would be considered in situations where performance was unsatisfactory. These decisions cover how Access Points would be connected to the ventilation ducts, but not how power would be radiated from within the ducts to users in the building.

Several techniques were considered for allowing wireless signals to propagate out of the ducts into the building environment. The simplest of these is to use the existing air outlet holes in the ducts that allow air to leave the ducts to also allow the wireless energy to escape. Two re-radiator antenna designs were also considered. The simplest of these is a single brass rod about half a wavelength long inserted into a plexiglass cap as shown in Figure 2.12. A hole is drilled in the duct and the smaller portion of the plexiglass disc is inserted into this hole. Often friction will hold the disc in place but it can be secured to the duct using duct tape or sealant. The length of the brass rod is chosen so that after it is inserted into the duct the lengths protruding on the inside and outside of the duct were both around one quarter wavelength.




Figure 2.12: A simple half-wavelength re-radiator. It is normally inserted into a ventilation duct to allow RF signal to escape the duct.

A second re-radiator design was also considered. Since the KlimaTak ventilation duct system has sheet metal acoustical panels that cover the ventilation ducts, there was a concern that the simple re-radiator would be shielded from the users. To combat this, two separate quarter-wavelength monopoles were used as shown in Figure 2.13. The monopole shown in the upper portion of the figure was inserted into a ventilation duct as with the AP antennas. The other antenna was wedged between two KlimaTak acoustical

panels in such a way that the panels acted as a ground plane for the monopole. A picture of this configuration is shown in Figure 2.14.



Figure 2.13: A more complicated duct re-radiator system. The upper antenna is inserted into the ventilation duct while the lower antenna is mounted between two metal KlimaTak acoustical panels in the ceiling and a short run of coaxial cable is run between the two.

To compare the potential improvement resulting from using the re-radiator antennas, the signal strength from an access point was measured at various locations for five scenarios. The first was a conventional installation of an AP not connected to a duct. The other four scenarios used an AP connected to a duct in a fixed location. The second of the five scenarios used a completely unmodified duct and no re-radiators. The third added RF absorber to the metal endcap terminating the illuminated duct segment, again with no




Figure 2.14: A picture of the panel re-radiator in which a monopole is wedged between two KlimaTak panels.

re-radiators. The fourth scenario was the same as scenario three, but with the addition of a simple re-radiator of the type shown in Figure 2.12. The fifth scenario was the same as scenario three, but with the addition of a dual monopole panel re-radiator as shown in Figure 2.13.

To minimize direct radiation and ensure the power measured was only from the ducts, the AP was placed inside an aluminum box filled with RF absorbing material. One AP antenna connection was terminated, the other was connected to a duct via RG-142 (double shielded) cable. A monopole antenna was used inside the duct. The AP used was a Linksys WRT-54G (hardware version 1.1) with unmodified firmware with a beacon interval of 1.024 seconds. To measure the signal strength, a Berkeley Varitronics YellowJacket IEEE 802.11b scanning receiver, shown in Figure 2.15 was used to measure the signal strength. The measurements were made by averaging the signal strength values recorded over a minimum of 20 seconds at each location. The YellowJacket was placed on a rolling cart and the position was recorded by estimating the location on a blueprint of the building.

Figure 2.16 shows the measurements made for the 5 scenarios considered. The measurements, AP and re-radiator locations are overlaid on a picture of the building blueprint. The AP is connected to the ventilation duct at the location indicated by the purple dot labeled "AP". When used, the re-radiator is at the locations labeled "Rrad" with another purple dot. The other colored dots indicate the received signal power in dBm at the marked location. The blueprint is somewhat confusing in that two separate duct



Figure 2.15: A Berkeley Varitronics' YellowJacket IEEE 802.11b scanning receiver of the type used for measuring received signal strength in the re-radiator comparison measurements. Picture excerpted from [30].

networks are present. The only duct that is being “lit up,” i.e., connected to the AP, is the run that proceeds straight up, jogs slightly left, proceeds to the top of the building, makes a right bend and stops in a flat endcap. This length of duct is marked in dark gray in the figure. All other ducts shown are not connected to this network and are shown as outlines. The duct used is about 30 meters long with filter bags over ventilation air outlets which are cut in the duct throughout its length.

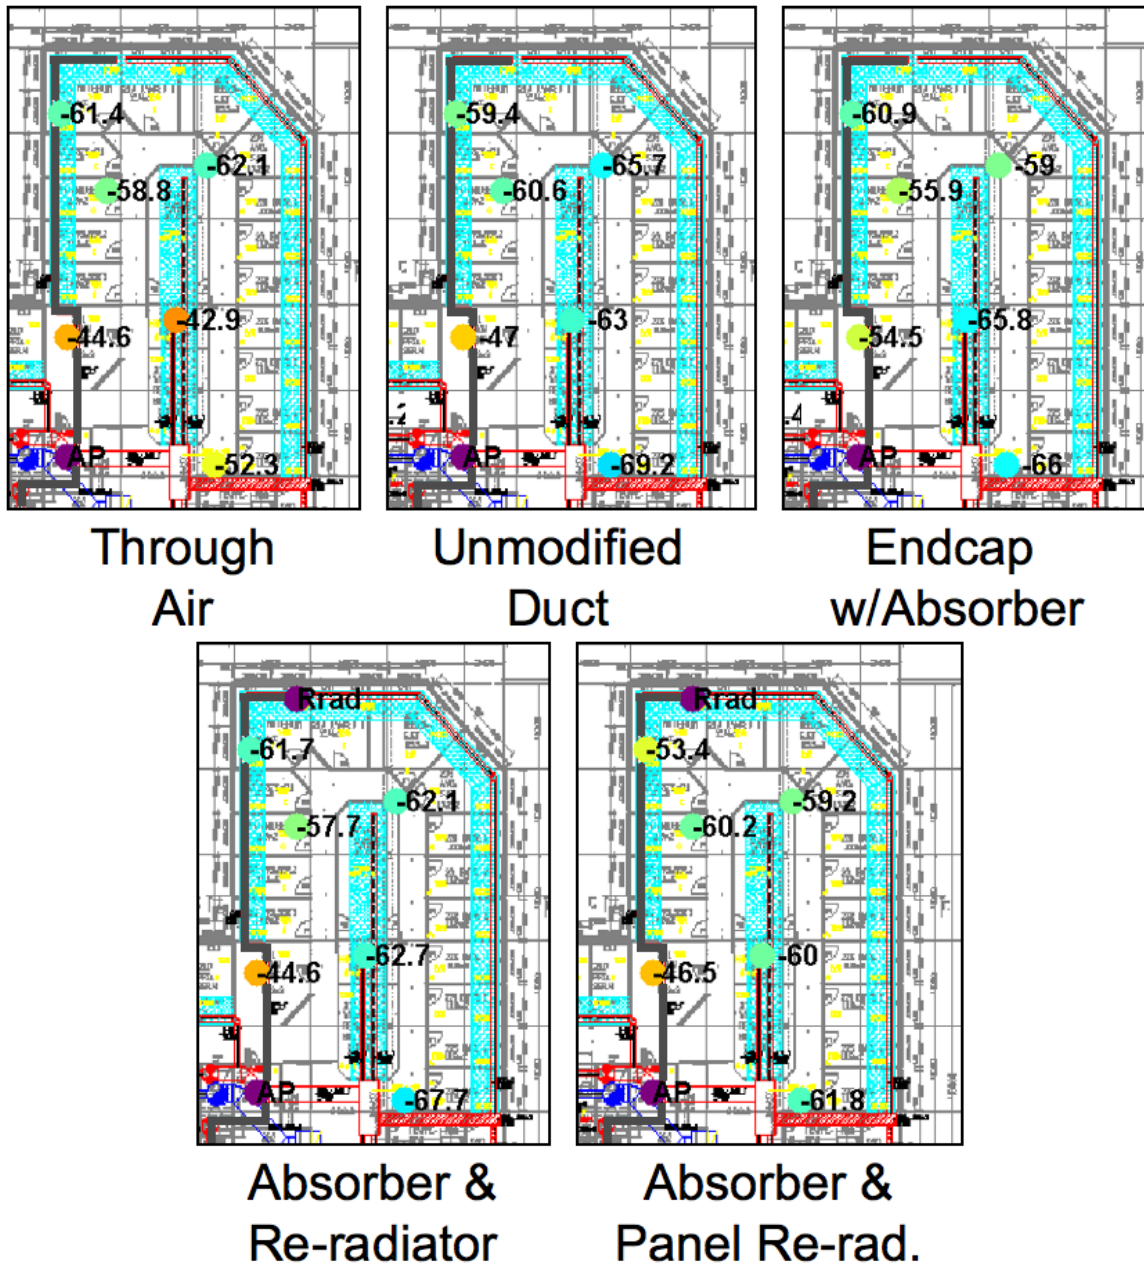


Figure 2.16: This figure shows a comparison of conventional, duct based, and duct with re-radiator techniques.

Of the measurements shown in Figure 2.16, the 4 duct scenarios are of most interest. First we note that, considering all of the measured signal strengths in each scenario, no case stands out as being clearly superior. There is a fair degree of variation in the signal strengths in each measurement case, but there is no systematic improvement in the signal strengths across each scenario. This is somewhat expected due to the difficulty in measuring in the exact same location across all of the different scenarios. It might be argued that the case with the panel re-radiator in the last case makes a difference to the signal strength in the room with the re-radiator, but it is hard to rely on a single point to make such a conclusion.

2.4.2 Comparing AP Antennas

To ensure we have chosen the best antennas for use in the ventilation ducts, we recorded two measurement sets of the power leaked from ventilation ducts at two different locations in the building. These measurements again used a Linksys WRT-54G AP placed inside a metal box filled with RF absorbing material. One AP antenna connection was terminated, the other was connected to a double shielded cable passed through the metal box and connected to a duct in a building ventilation duct. The antenna installed in the ventilation duct was different in each of the two measurement sets. The antennas used were a quarter-wavelength monopole, and a half-wavelength curved arm dipole. The monopole and dipole antennas were of the same type used in Section 2.3.1. For each AP antenna, we moved a laptop with a Buffalo Wireless Air Station, WLI-CB-G54A, IEEE 802.11g wireless LAN card to various locations in the building. Using custom MATLAB® software, described in detail in Appendix A, we recorded the uplink signal strength measured by the AP of the laptop wireless card transmissions. The signal strength recorded for each laptop location using the monopole AP antenna is shown in Figure 2.17 and using the dipole AP antenna is shown in Figure 2.18. In these figures, the measured signal strength in dBm is plotted overlaid on top of a picture of the building structure and is indicated by a colored dot with the color proportional to the signal strength. The measured value is also shown in text. While we took more measurements for the dipole case than the monopole case, 34 points were taken in both cases at as close to the same location as possible. Since the laptop was moved between measurement sets, some positioning error is inevitable. We attempted to minimize this error by always placing the laptop in the same place on the rolling cart and by aligning the cart with office furniture. As a result we estimate that repositioning error was maintained to within half a wavelength or better in the horizontal plane.

Firstly, the leakage measurements shown in Figures 2.17 and 2.18 indicate that an area of the building can be well covered without the need for re-radiating elements. Further, these measurements can be used to compare the two AP antennas. To do this we have computed the mean difference in signal strength over the 34 measured locations common to both measurements. These 34 points include of the measured points in the

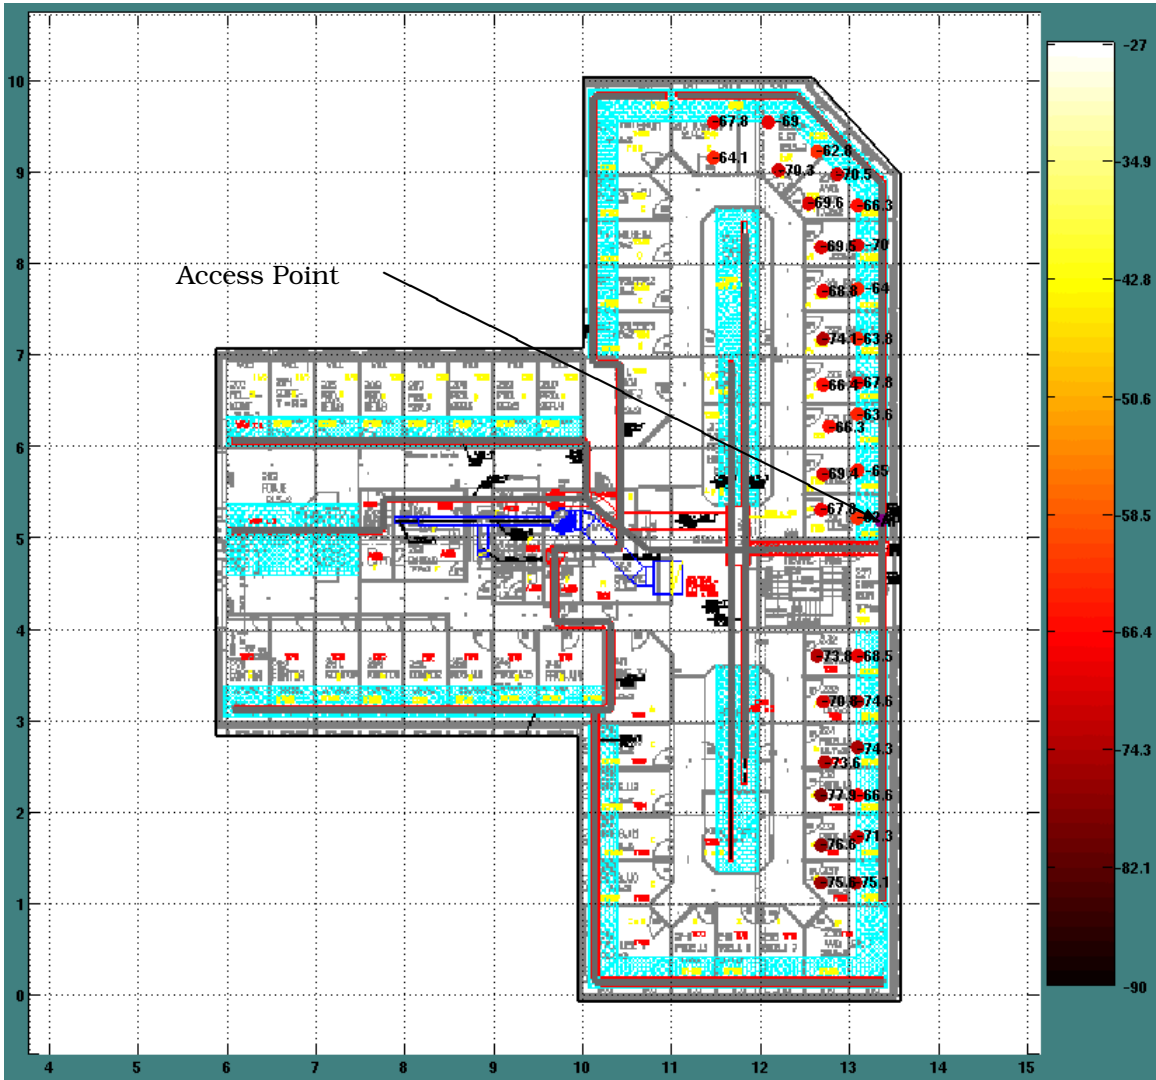


Figure 2.17: Leakage comparison measurements of signal strength using a monopole antenna connected to an AP.

monopole case and a little more than half of the points in the dipole case. The mean difference shows a 2.2 dB gain in signal strength using dipoles compared to monopoles.

2.4.3 Design Decisions

We further confirmed that re-radiators do not have a significant impact to the received signal strength except in a small area around the re-radiator. We did this by walking around the re-radiator while watching the received signal strength indication (RSSI) measured by the YellowJacket. It became clear that the only significant increase in power over the average was within about 1 meter of the re-radiator.

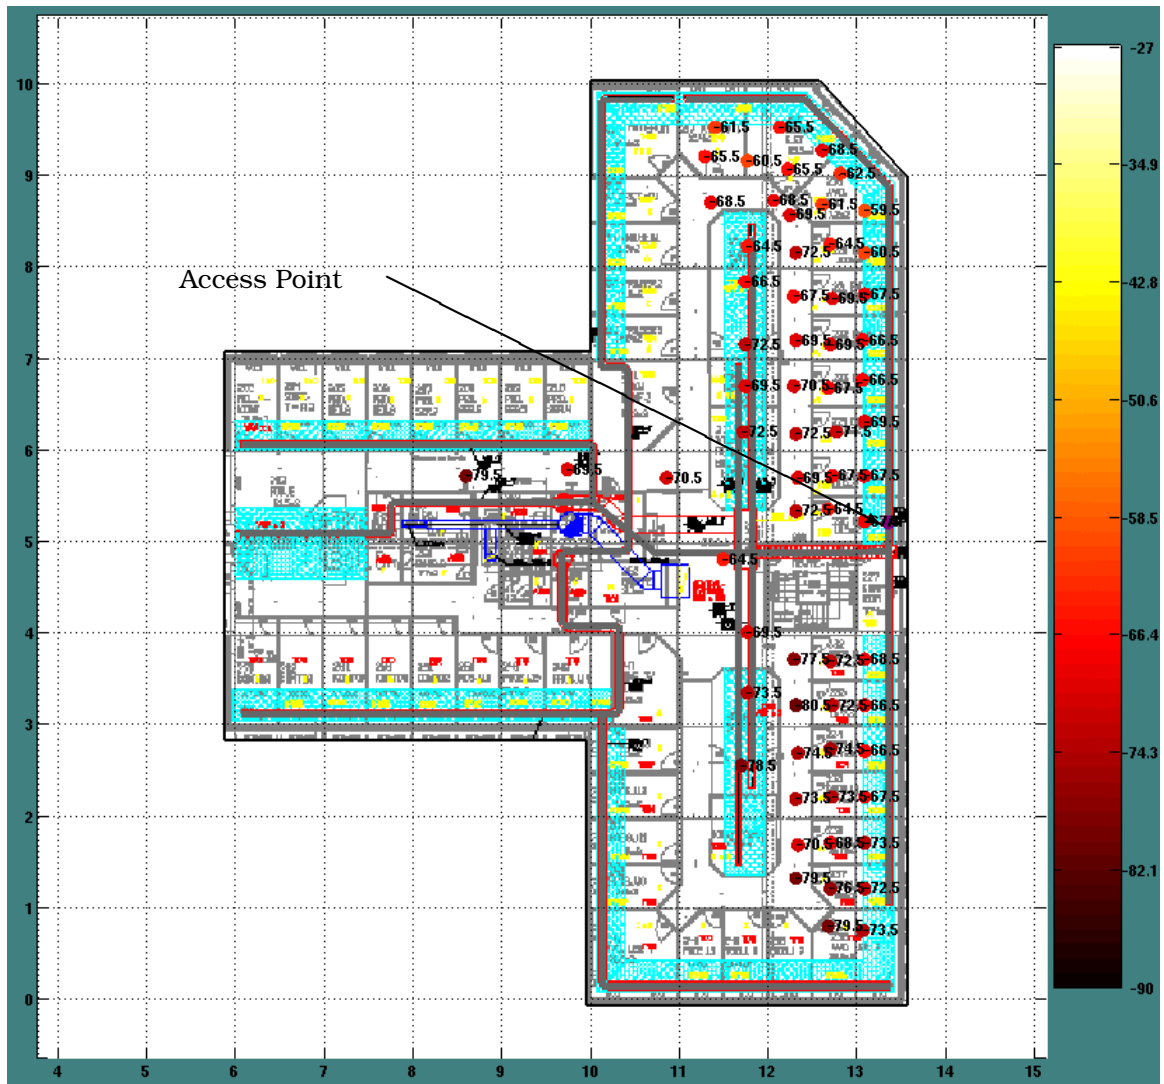


Figure 2.18: Leakage comparison measurements of signal strength using a dipole antenna connected to an AP.

Since the measurements in Section 2.4.1 indicated that there was no significant improvement in coverage using re-radiators, we believe installations should be made without the need for re-radiators. Apparently the ventilation duct seams, joints and ventilation air outlets leak a significant amount of power along the length of the ducts. As a result, the power from a monopole re-radiator is only higher than that leaked from the ducts close to the re-radiator. While this result is somewhat unexpected, it is fortunate since it implies that the extra cost and effort to install re-radiators is not necessary.

The measurements in Section 2.4.2 confirm that significant coverage can be provided without the use of re-radiating elements. Further, the comparison between using a

quarter-wavelength monopole instead of a half-wavelength curved dipole results in an average signal loss of 2.2 dB. We judge that this difference might be significant in some installations, but is probably not enough of a difference to warrant the increased difficulty in installing dipoles instead of monopoles.

2.5 Conclusion

Despite our initial expectations, and having tried several sophisticated concepts for designing a ventilation duct based RF distribution system, it is surprising and yet also satisfying that the some of the simplest approaches are also some of the most effective. After considering monopole and dipole antennas, RF absorber for metal endcaps, metal screens covering the ventilation duct outlets, and two different re-radiators, no combination was found to significantly out-perform the simplest configuration: monopoles in un-modified ducts. In fact, we found that omitting the metal screens helps filter dispersive modes and allows a significant amount of power to escape into the building environment. The duct outlets appear to be so effective at minimizing dispersion down that the simpler monopole antennas are just as effective as dipole antennas in maximizing the throughput of the duct channel. Further, the duct outlets leak sufficient power into the building environment that re-radiators are unnecessary. Combining these results, we conclude that the only modification needed for a KlimaTak style duct system is to couple the APs to the duct system using simple monopole antennas inserted into the ducts.

3.1 Introduction

The bulk of this chapter focuses on measurements made in 2 testbed buildings to develop a method for estimating the coverage area of an AP connected to a ventilation duct network. In this chapter we present measurements of the RF signal strength radiated from ducts excited by IEEE 802.11g APs. We also discuss how the measurements are used to develop empirical propagation models that can be used to design a duct based distribution system.

Previously, Nikitin et. al. presented methods for predicting the frequency response and impulse response of a ventilation duct system using detailed knowledge of the duct and antenna geometry [9, 10, 26]. In contrast, the work presented here relies on less detailed knowledge of the duct environment because we wish to use it for analyzing large, full building sized duct networks. It is similar in nature to empirically based models of the pathloss through different duct types as presented in [19, 20]. However, we deviate from the values presented in those works owing to the large number of duct outlets in the buildings considered. It should also be mentioned that this work is influenced by Carter's method of predicting the coverage of leaky feeder coaxial cables indoors presented in [31]. However, our approach is somewhat different in that we do not chop the duct up into segments but rather assume the power leaks uniformly from all duct locations. Further, our prediction methods are influenced by indoor propagation models such as [32, 33]. We extend the concepts of pathloss exponents and multi-breakpoint models by combining them with duct propagation concepts. These combinations are discussed in detail below.

3.2 RF Leakage from Ventilation Ducts

Early in the development of this project, it was expected that little power would leak from ventilation ducts into the environment except at duct louvers where plastic pressure valves or cloth filter bags provide easy egress for RF energy. It was felt that there may be occasions in which we would like to prevent RF energy from leaking out of the duct. In an attempt to prevent such leakage we measured the normal leakage from ventilation ducts with and without various metal screens over the air outlets. To quantify the effect of different screen materials, we placed a transmitting antenna inside a ventilation duct and then measured the signal strength just outside the ventilation duct air outlet. This

setup is illustrated in Figure 3.1.




Figure 3.1: Setup for duct leakage measurements.

The duct shown in Figure 3.1 is 315 mm in diameter and is part of an actual installed duct network found in the 3rd floor break room in the YIT building in Stavanger, Norway. The duct section in the room is shown in the figure but extends outside of the room to both the left and right. The duct is normally hidden behind KlimaTak paneling made of thin sheet metal. However, these measurements were made with all panels open. Also as shown in Figure 3.1, there are two 200 mm diameter neck pieces installed that normally have pressure values installed. These necks cover an oval hole cut in the duct that matches the size of the 200 mm diameter hole in the neck piece. Measurements were made with three different types of screens blocking the air outlets. Two were made of thin gauge wires woven in a mesh with 3 mm or 6 mm square holes. The third was made of sheet metal with 5 mm diameter circular holes punched out of the material. These screen materials were placed directly over the air outlet holes cut in the ventilation duct. A transmitting dipole tuned to 2.45 GHz and with curved arms was placed inside the ventilation duct, oriented to excite primarily a TE_{01} mode. A second dipole was placed outside of one of the pressure valves installed on a neck piece to measure the signal leakage. Measurements were made with the receive dipole oriented in both the vertical and horizontal polarizations and with and without each screen installed. To prevent any leakage effects from the second louver, aluminum foil was used to seal the duct opening.

The screens were installed to fit the curve of the duct and placed between the duct and the neck piece. As can be seen from Figure 3.2, the screens do have a shielding effect on the signal leakage measured right up against the duct air outlet. In addition, the screens with smaller holes tend to decrease the signal more than screens with larger holes.

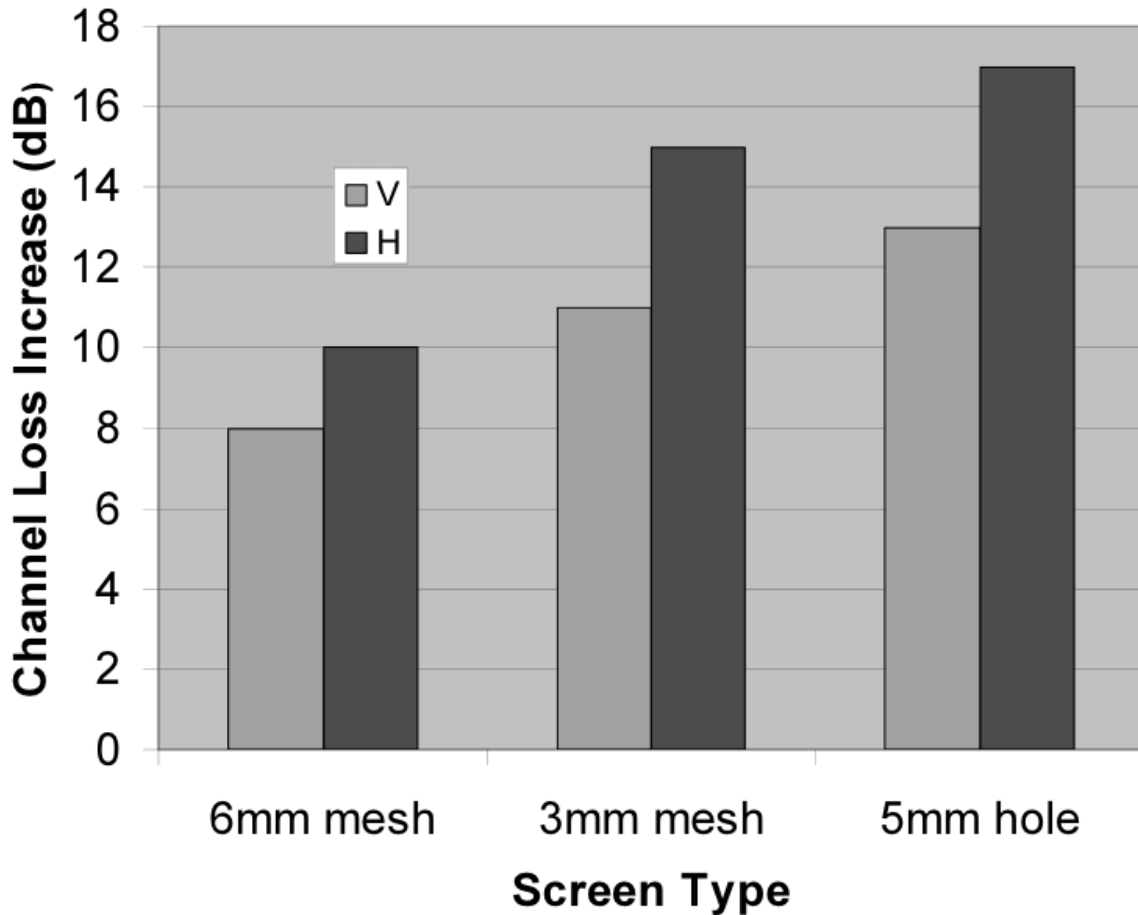


Figure 3.2: Attenuation of signals leaked out of ventilation ducts due to screens placed over louver air outlets.

We also wished to quantify the effect on signal strength due to the KlimaTak panels. To do this we again placed a transmitting antenna inside the ventilations ducts, this time a quarter-wave monopole antenna. We then measured just the leakage at four distances away from the ventilation duct that all had line-of-sight to the ventilation duct. The difference in the receive signal strength at each point with the KlimaTak panels up compared to down is plotted in Figure 3.3. As can be seen from the figure, the signal strength difference is approximately 5 dB lower with the KlimaTak panels in position.

Despite these measurements, in practice we have found only an 8.5 dB reduction in the average signal when all louvers are covered with screens in an area. This is believed to be due to leakage from the seams at duct junctions. Using the YellowJacket IEEE 802.11g scanning receiver described in Section 2.4.1, we were able to confirm this by searching around ducts for the highest signal strength. Areas near junctions and duct seams often

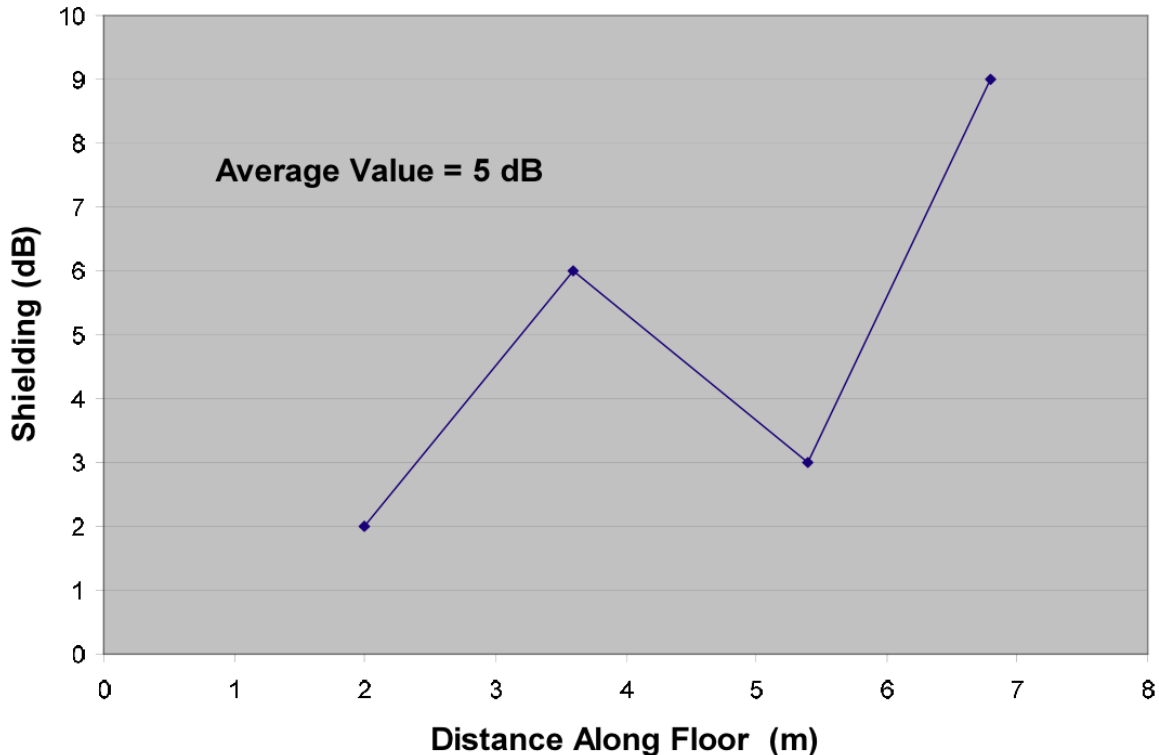


Figure 3.3: Loss in signal strength due to closed KlimaTak panels.

had significant power levels. It is possible that the power could be contained in the duct by sealing the joints better with aluminum tape or similar. However, it is likely that this procedure would be more expensive than it would be worth. This both illustrates the difficulty in preventing ducts from leaking power and the advantage to using the ducts for radiating power.

3.3 Building leakage

An important consideration of a ventilation duct system is how much the system affects the amount of power which leaks outside of the building. If a duct based system increases or decreases the amount of average power that leaks outside of the building, this could have important impact on the security of the system.

To investigate the differences between a conventional and duct based system, we compared the power leakage outside of a building from two identical access points. One AP was connected to a ventilation duct, the other was placed in the same location but using the standard AP antennas. The location of the APs is indicated by a star in Figure 3.4. Also shown are the 16 measurement locations outside the building. It should also be noted that the APs were near ceiling level on the second floor of the building and mea-

measurements outside were made at ground (i.e. first floor) level. The signal strength was measured using the YellowJacket so that both APs could be measured simultaneously. Both APs were set to the same channel, but had different SSIDs so that beacons could be distinguished and so that power measurements could be made simultaneously. Each signal strength measurement was averaged over 15 or more seconds of data. Since the beacon interval of the APs were set to 0.1024 seconds, this results in about 150 measurements at each measurement point.

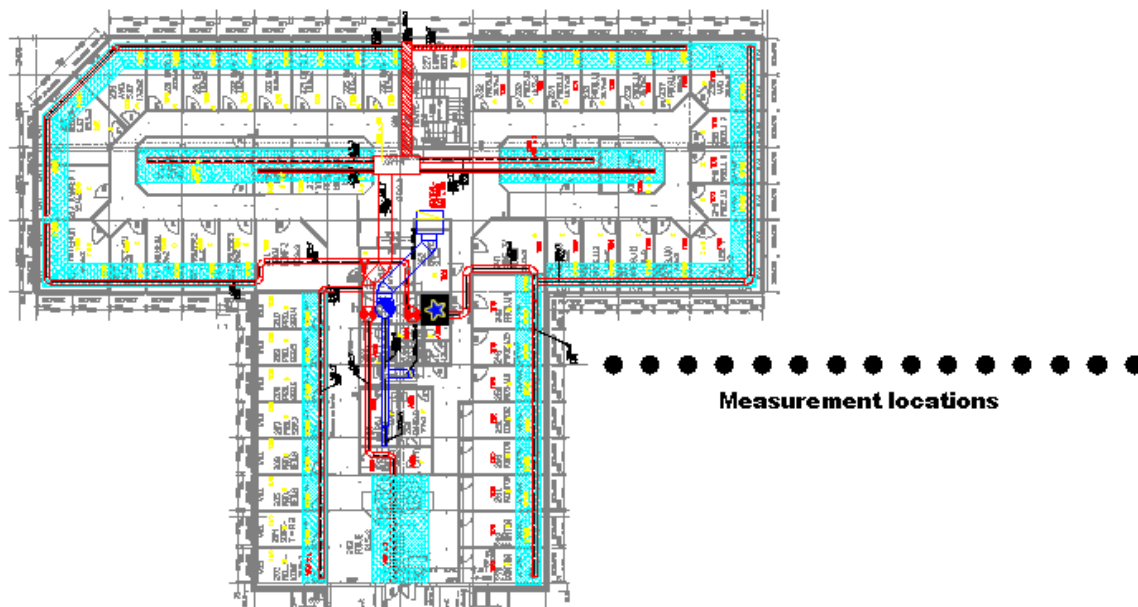


Figure 3.4: Measurement setup of power leakage outside of a building due to a conventional or duct installed AP.

The power level measured at each location is plotted in Figure 3.5. Both the conventional and the duct connected AP power levels are shown. As can be seen from, there is hardly any difference in the signal strength between the two systems. As a result, we conclude that there is no security advantage or disadvantage to a duct connected AP. We were somewhat surprised by this result, having expected a duct based system to leak less outside of the building than a conventional system. However, we feel that the duct conveys higher power levels to the far corner of the building, thereby reducing the rate at which power falls off with distance.

3.4 Signal Strength Measurements

3.4.1 Initial Measurements

To develop a model for predicting the signal strength at various points inside ventilation ducts and in the building environment, it was decided that we should take a wide range




Figure 3.5: Comparison of power leakage outside a building from conventional or duct based wireless AP.

of measurements in the testbed building we had available to us. By combining prior knowledge of the performance of signals in ducts and the prior propagation models for ventilation ducts we have created a propagation model based on physical propagation mechanisms with parameters that can be adjusted based on measurements.

Initially, 98 measurements were made in each of five cases on the second floor of the testbed building in Stavanger, Norway. In each case, the APs were in different locations, but the measurements were made at approximately the same location in the building environment. In addition, the second case used a panel re-radiator of the type mentioned in Section 2.4.1. In all measurements, the AP was again sealed in an aluminum box and a single antenna terminated and disabled, the other connected to a quarterwave monopole inserted in the duct. The receiver used for the measurements was the YellowJacket IEEE 802.11b receiver used in Section 2.4.1. The receiver was placed on a small, wooden rolling cart about waist height. The signal strength was measured at each location using

a horizontally polarized 2 dBi antenna. Signal strengths were recorded and averaged over at least a 20 second period at each location.

The measurement cases are shown in Figures 3.6, 3.7, 3.8, 3.9, and 3.10. The measurements are plotted as points on top of an image of the building blueprint. Each dot is colored as shown by the scale of the colorbar on the right side of each figure, with the units in dBm of received power. Each colored dot is labeled with the actual power level. The AP is labeled with a purple dot and the duct network actually excited by the AP is highlighted in white. Since the floor has two separate duct networks, the other duct is shown in gray. In addition, a semi-transparent colored mist is drawn behind the measurement dots on top of the blueprint. This colored mist is an interpolated mist of the actual measured values to better visualize the measurements. To get a reference of the scale, each grid square shown in the figures is 4.8 meters on a side.

The measurements presented were used for validating models and computing model parameters. The models developed will be discussed in Section 3.5.

3.4.2 Comparison of Duct System to Conventional System

In this section we present measurements that compare the coverage from a single AP versus the coverage provided by a conventional AP simply placed in the building environment. These measurements were performed by measuring the signal strength at the same 98 locations as used in Section 3.4.1. In fact, the two cases presented here correspond to cases 1 and 3 already presented in Figures 3.6 and 3.8. In these two cases, we also measured the signal strength at the locations from a conventional AP placed in the same location as when the AP was connected to the ventilation ducts. Figure 3.11 shows a comparison of the duct and conventional coverage maps generated from measurements for the first case. Figure 3.12 shows the coverage maps for the second case (i.e. the third case presented in Section 3.4.1).

Note that in both cases presented in Figures 3.11 and 3.12, the conventional propagation case has a brightly colored “hot spot” around the AP location. This hot spot is not present in the duct case because the power is conveyed away from the point of excitation by the duct. This does not imply that the conventional case is “better.” In fact, this hot spot represents effectively wasted energy. The power in the conventional case must be blasted away to maximize the coverage area. In contrast, the power is distributed more uniformly by the duct network. It is also interesting to note that the conventional coverage cases have possibly unexpected or hard to predict dark spots in the coverage areas. These dark spots represent areas of low received power that could represent a hole in coverage.

To better compare the coverage provided by the duct and conventional systems, we plot the difference in signal strength measurements. The differences are computed at each of the 98 measurement points and the differences plotted as before and shown in Figures 3.13 and 3.14. In these figures blue colors indicate better conventional signal

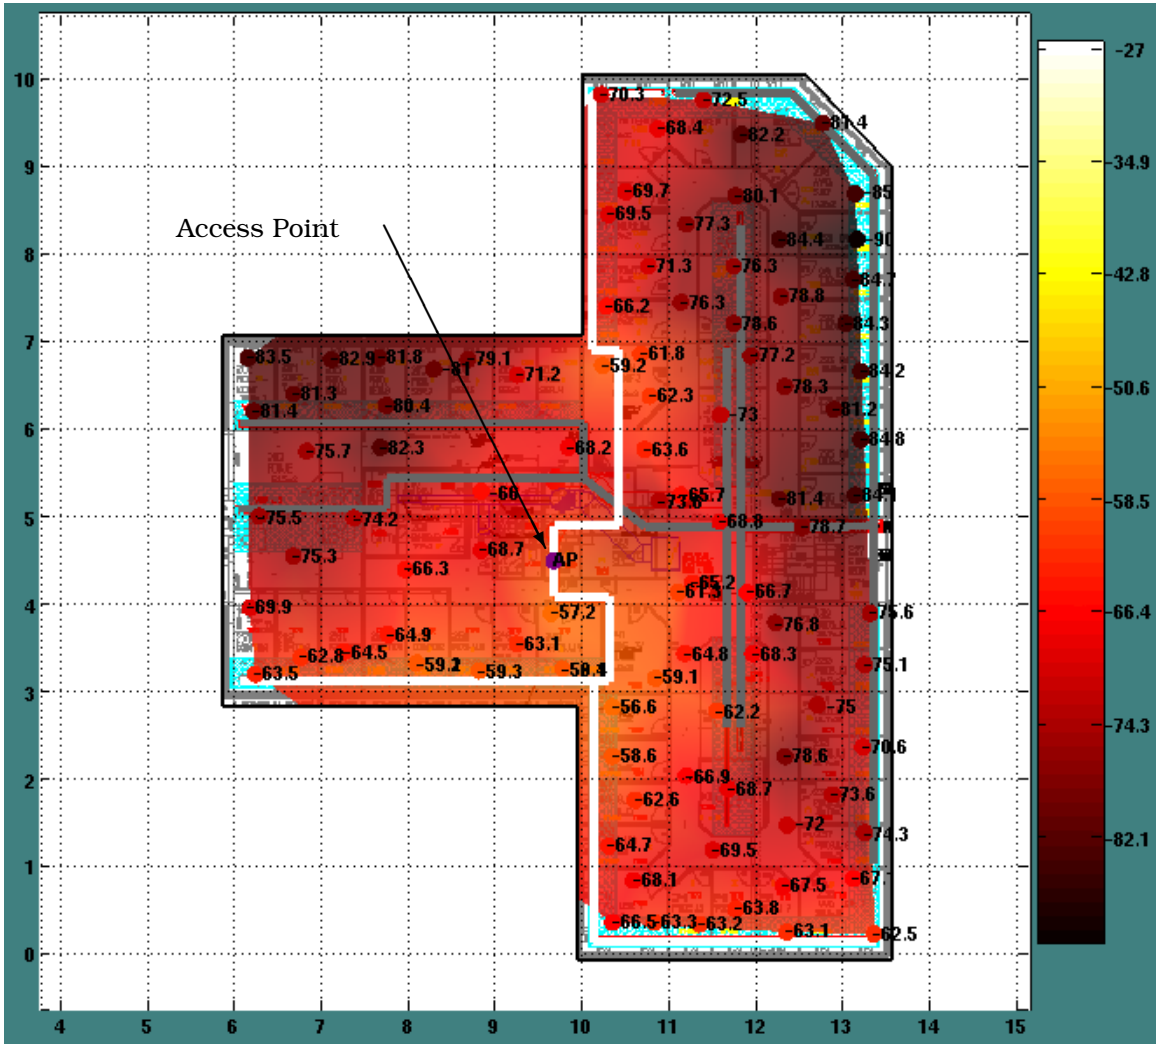


Figure 3.6: First measurement set of ventilation duct based propagation. The AP is connected in the center of the building to the duct network highlighted in white. The signal strength is plotted in dBm.

strength, red colors indicated better duct signal strength and greenish colors indicate approximately equal signal strength. Each plot has a color scale in dB on the right side of the figure. While the duct system only overcomes the conventional system in a few locations, it is in locations that matter most. The duct system reaches further than a conventional system and this is evident from the figures. Further, this is for a somewhat worst case building environment that is open with few walls. Those walls that do exist are mainly glass. As a result, the conventional system is rather competitive.

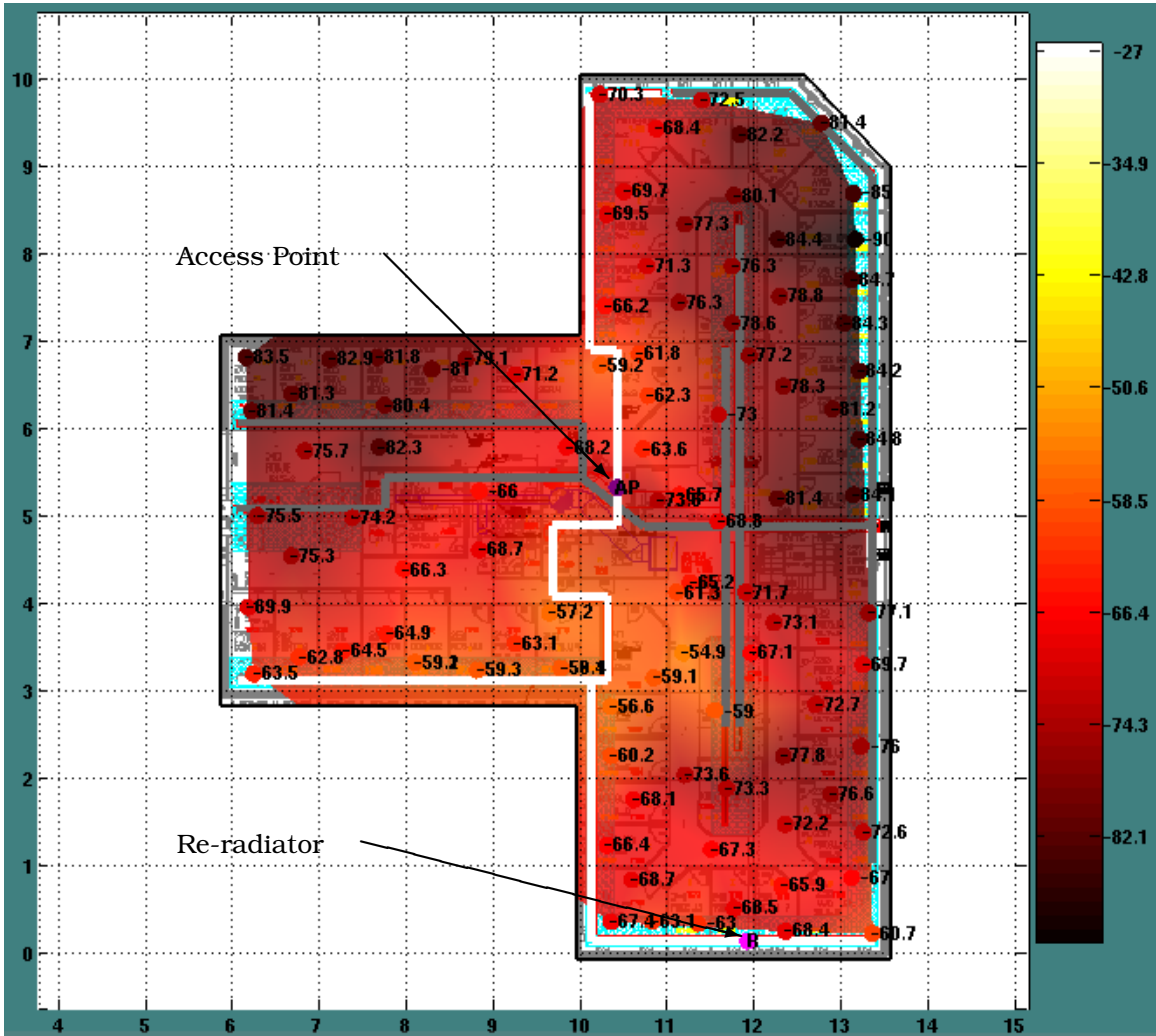


Figure 3.7: Second measurement set of ventilation duct based propagation. The AP is in the center of the building and re-radiator was added as shown by the pink dot in the bottom right. The signal strength is plotted in dBm.

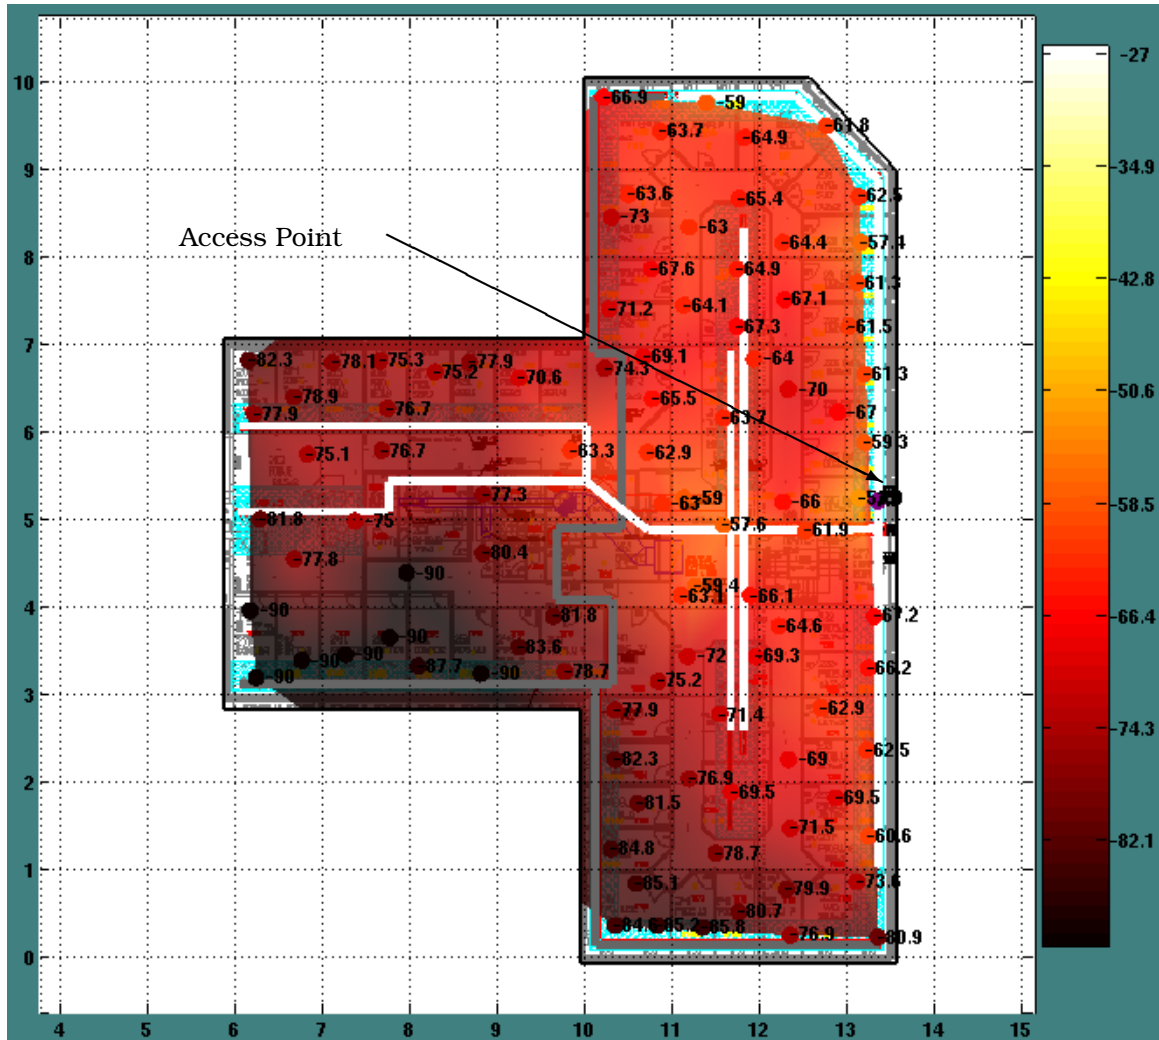


Figure 3.9: Fourth measurement set of ventilation duct based propagation. The AP is connected in the right side of the building to the duct network highlighted in white. The signal strength is plotted in dBm.

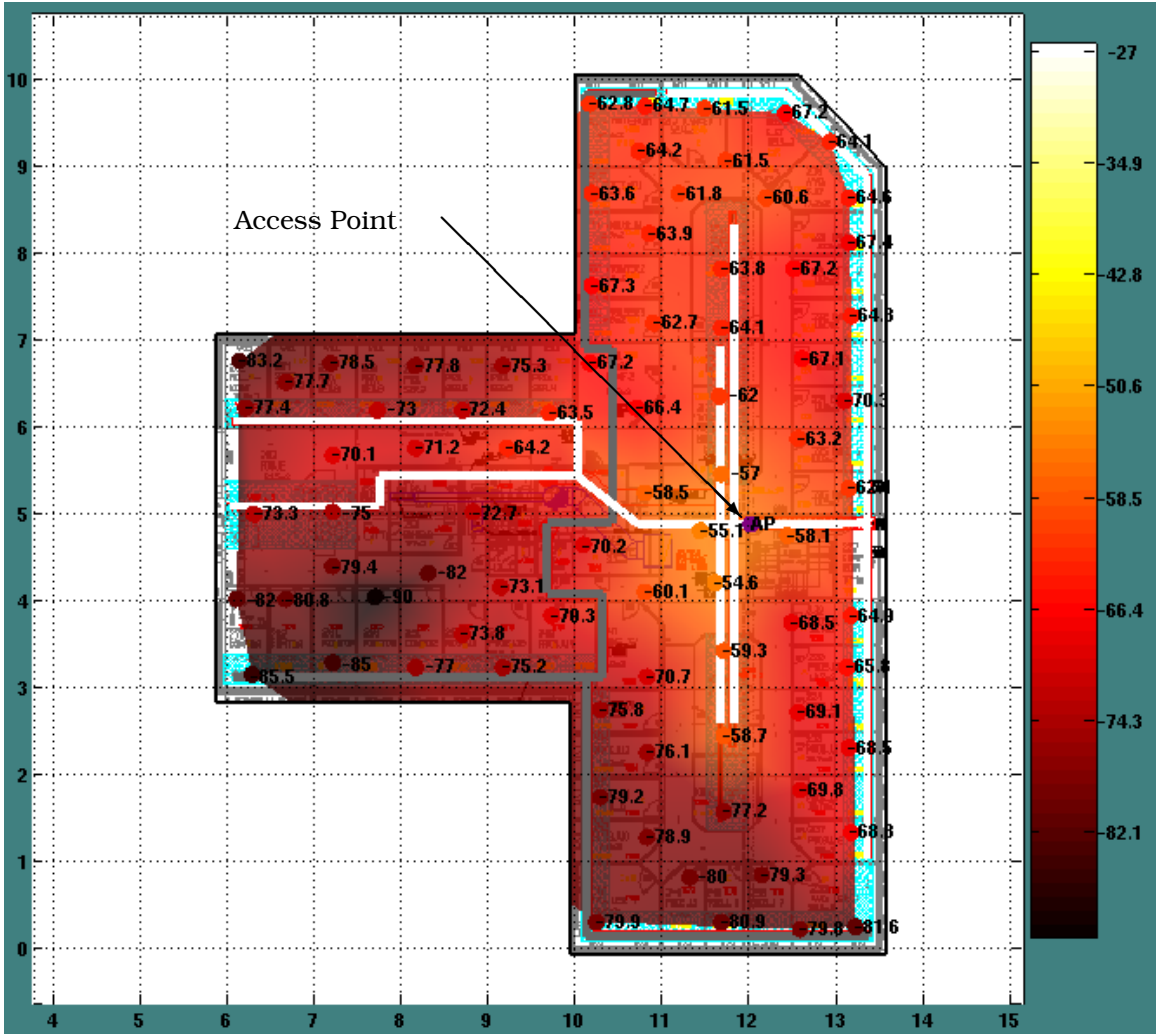
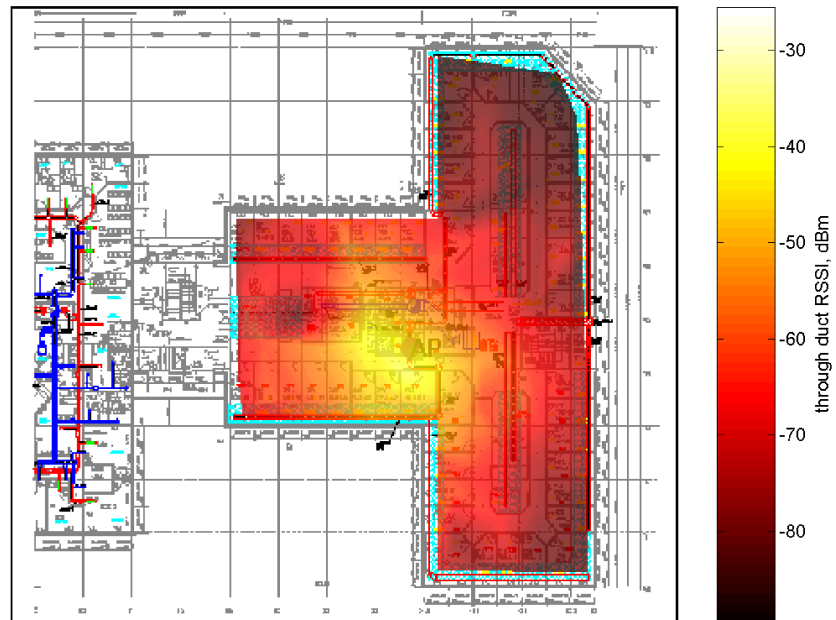
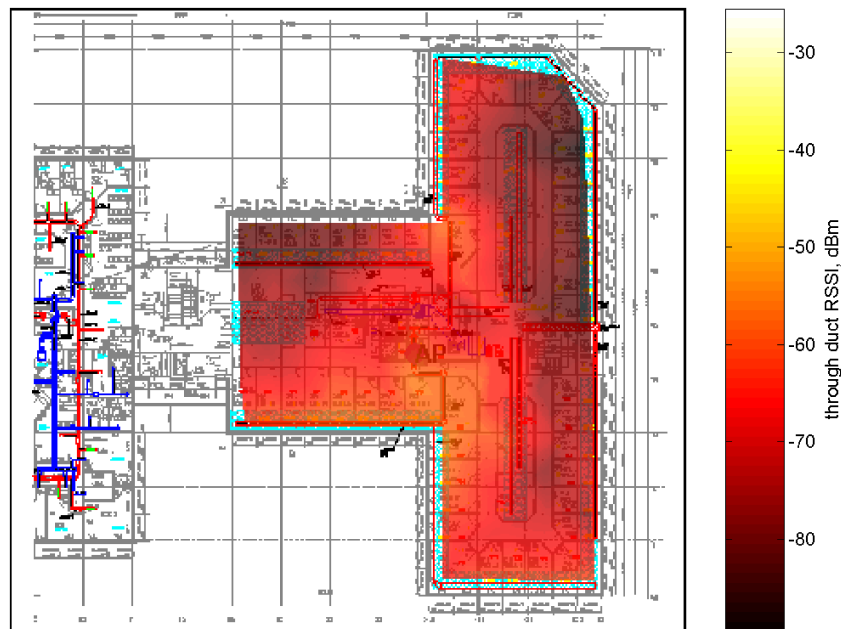


Figure 3.10: Fifth measurement set of ventilation duct based propagation. The AP is connected in the right side of the building to the duct network highlighted in white. The signal strength is plotted in dBm.

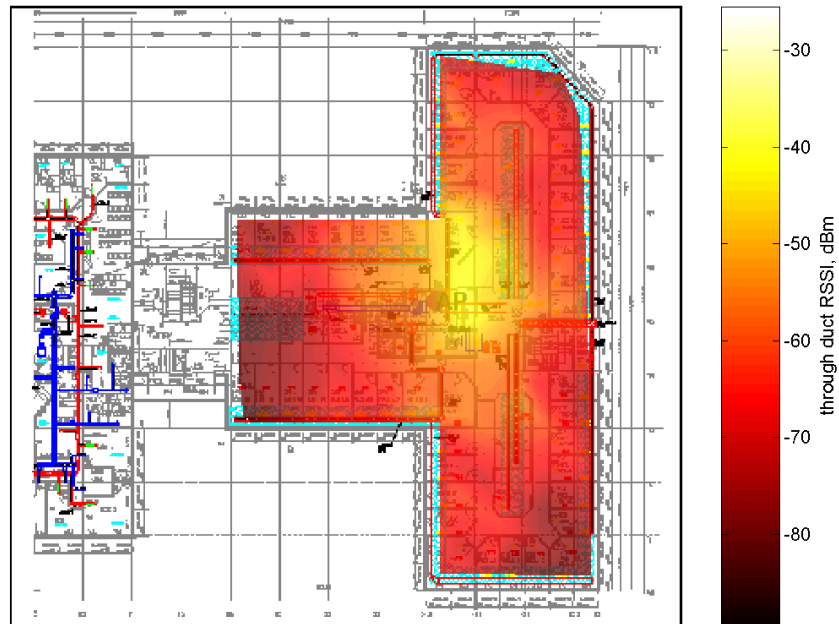


(a) Conventional Coverage

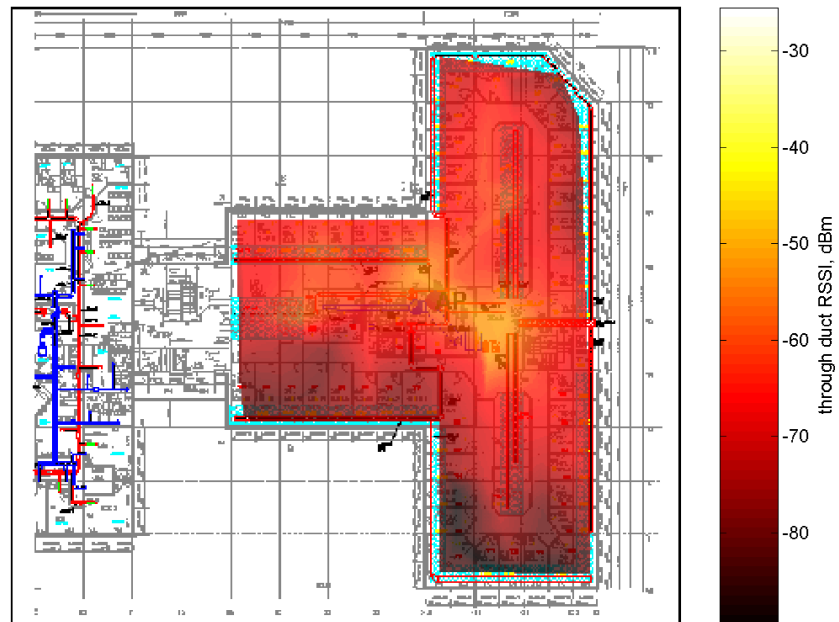


(b) Duct Based Coverage

Figure 3.11: Comparison of conventional and duct based coverage areas for an AP installed in the same location in the same building. The measurements are plotted by interpolating the colors between each measurement point.



(a) Conventional Coverage



(b) Duct Based Coverage

Figure 3.12: A second comparison of conventional and duct based coverage areas for an AP installed in the same location in the same building. The measurements are displayed by interpolating the colors between each measurement point.

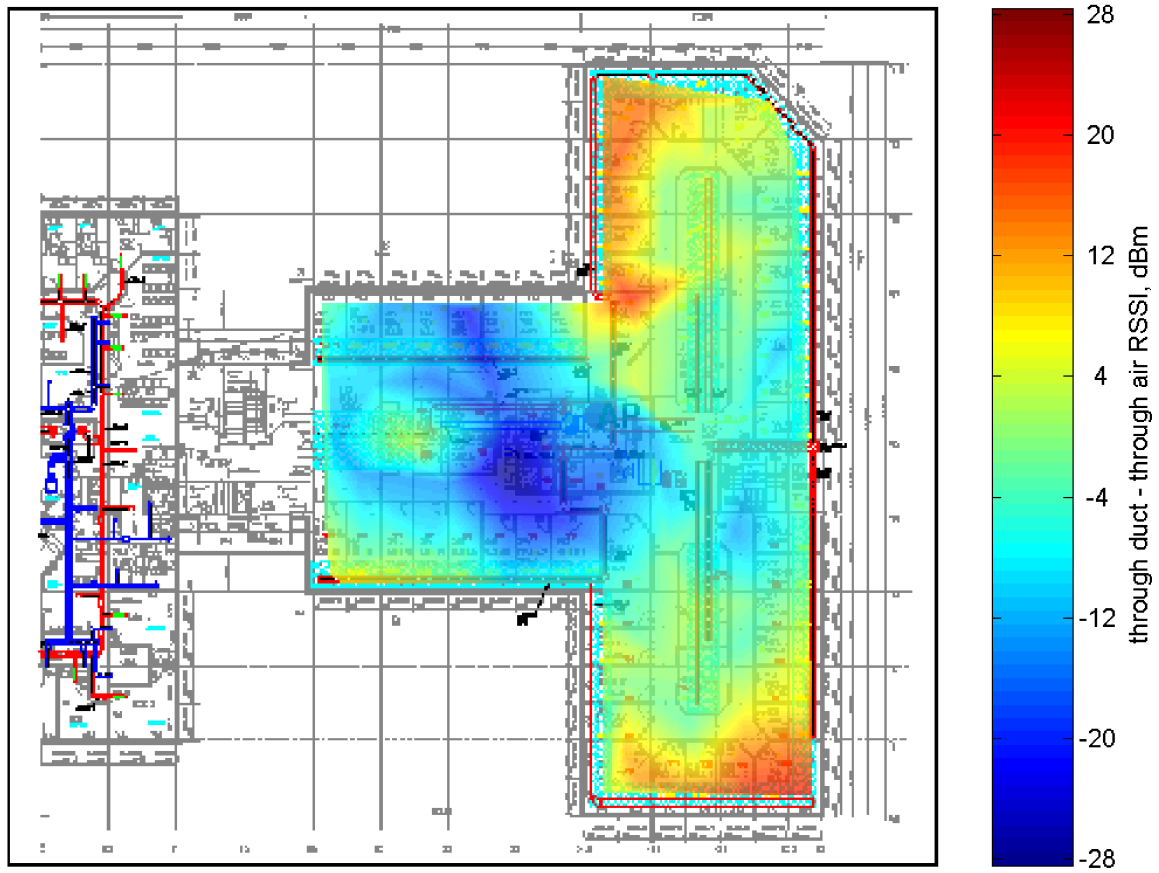


Figure 3.13: Plot of the difference in signal strength between conventional and duct based propagation using the data presented in Figure 3.11.

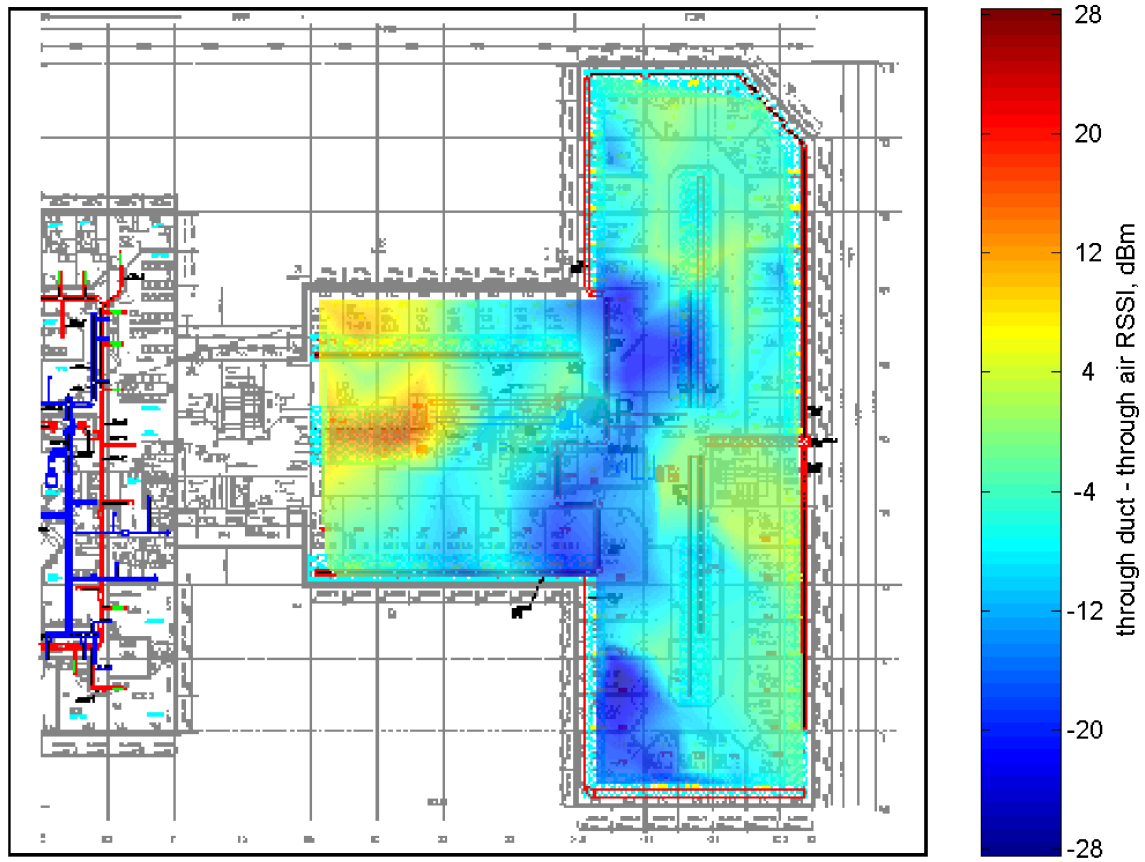


Figure 3.14: Plot of the difference in signal strength between conventional and duct based propagation using the data presented in Figure 3.12.

3.5 Propagation Model

3.5.1 Description of the Model

To aid in the design of a wireless duct system, a propagation model has been developed that can be used to predict the received signal strength, R , at locations inside a duct network and at locations inside a building near the ventilation duct. This model is simple to use but has still been found to be sufficiently accurate for coverage prediction. The model relies on 8 parameters:

r is the distance to the closest point on the nearest duct from a particular location in the building

d is the distance through the ventilation duct network from the nearest AP to the closest point on the nearest duct

J_m is the number of m -way splitting duct junctions between AP and closest point on nearest duct

P_0 is the transmit power used by the AP

L_{in} is the coupling loss into the duct from the AP

α is the loss per meter through the ventilation duct

L_{out} is the coupling loss from inside the ventilation duct to 1 meter away from the duct

n is the pathloss exponent used to compute the propagation from the duct to the receive location

Once these 8 parameters have been specified, the signal strength at a particular location can be computed using

$$R \text{ [dB]} = P_0 - L_{in} - \alpha d - 10 \sum_{m=2}^M J_m \log_{10}(m) - L_{out} - 10n \log_{10}(r). \quad (3.1)$$

All variables in the above equation have already been specified except for the value of M , which is the largest number of junctions found in any duct component. For example, if a 3-way splitter such as a cross junction is the duct junction type present with the largest number of splits in it, then M is set to 3. The large summation term then is just a calculation of the split loss at each junction. So if there are two 2-way junctions and one 3-way junction in the length of duct between an AP and the closest point on the nearest duct to a particular location, then the summation would be evaluated as

$$10 \sum_{m=2}^M J_m \log_{10}(m) = 10J_2 \log_{10}(2) + J_3 \log_{10}(3)$$

$$\begin{aligned}
 &= 2 \cdot 3.0 + 1 \cdot 4.8 \\
 &= 10.8.
 \end{aligned}$$

The other terms in the simulation equation are straightforward. As an example, suppose you wanted to know the received signal strength, R at a location that was 3 meters from the nearest duct, and the closest point on the nearest duct was 15 meters from the closest AP and had two 2-way junctions and one 3-way junction between it and the AP. Then, if the AP transmitted with a power of 20 dBm, the loss per distance of duct used was 0.34 dB/m, the pathloss exponent was 1, and the coupling losses were 8 and 60 dB for the AP to duct and duct to air coupling losses respectively, then the power received at the simulated location would be

$$\begin{aligned}
 R &= P_0 - L_{in} - \alpha d - 10 \sum_{m=2}^M J_m \log_{10}(m) - L_{out} - 10n \log_{10}(r) \\
 &= 20 - 8 - 15 \cdot 0.34 - 10.8 - 60 - 10 \log_{10}(3) \\
 &= -68.7 \text{ dBm}.
 \end{aligned}$$

3.5.2 Model Advantages

The most useful aspect of the above simulation model is that it requires minimal modeling of the building environment. By assuming that all points in the simulation are within a line-of-sight to a ventilation duct component, the losses introduced due to walls and building structure will not have a significant impact on the signal strength at a particular location in the building. It should be noted however that locations that do not have line-of-sight to a ventilation duct should be expected to have a lower than predicted signal strength.

Ducts do need to be modeled, however, and some care is needed to properly identify duct junctions, but the model does not distinguish between different types of ducts or different duct diameters. Instead, because KlimaTak installations tend to be similar in their installation, we can rely on ducts to be “typical” and simplify the modeling process. Thus, we can expect some additional error in the simulation model presented for duct networks that are significantly different than a “typical” KlimaTak system, either in duct dimensions or in the number and size of the pressure valve air outlets in the duct network.

The simulation model can be extremely quick for computer based calculation. It is true that finding the “nearest” duct to a particular location is not a trivial computation. However, since the locations of ducts are expected to be fixed, all of the nearest duct and distances to duct calculations can be precomputed at any number of desired points inside the building. As a result, as soon as an AP location and transmit power is specified, the only calculation required to find the power level at all selected points inside the building

Table 3.1: Computed parameters using empirical fit to propagation model. The fits were based on 98 different measurements in each case.

Parameter	Case 1	Case 2	Case 3	Case 4	Case 5
P_0 (dBm)	18				
$L_{in} + L_{out}$ (dB)	73.3518	73.1139	72.1801	71.5980	79.3876
α (dB/m)	0.07739	0.10658	0.039665	0.14106	0.30545
n	2.0124	1.9633	1.9934	1.3413	1.0237

is that of the power levels at all points in the duct network. However, because the loss through ducts is computed on a per meter basis, power only needs to be computed at duct junctions where split loss can cause a power step. Thus, the simulation time will be dominated by the number of duct junctions. As we do not expect more than a few hundred duct junctions in any one building, the simulation time will not be significant if the proper pre-calculations are performed.

3.6 Computing Model Parameters

In this section we describe the empirical techniques and additional measurements performed to compute parameters for a propagation model. The model used is described in Section 3.5. Also, many of the measurements used were presented in Section 3.4.1.

3.6.1 Purely Empirical Model Fit

The first attempt at computing parameters for the prediction model described in Section 3.5 was to simply attempt a best fit to the previously collected data. A fit was performed on the data collected in each of the 5 cases described in Section 3.4.1. The fitting procedure was performed using the MATLAB[®] command `nlinfit`. This command attempts to adjust model parameters using known inputs and outputs for the system. The command relies on initial estimates of the model parameters and uses a least squares estimation to minimize the error in the predicted output and actual output. For each of the 5 measurement cases, the fit model parameters computed are reported in Table 3.1.

Note that the nature of the prediction model prevents `nlinfit` from distinguishing between L_{in} and L_{out} . For this reason, the sum of the parameters are reported instead of individually. The initial guesses used for the parameters were $L_{in}=8$ dB, $L_{out}=60$ dB, $\alpha=0.2$ dB/m, and $n=2$. These parameters were picked based on prior measurements and assuming free space propagation between the duct and the end users. In all cases, `nlinfit` reported a mean error of 0 and standard deviations of between 5 and 5.7 dB. Although the values of each parameter in Table 3.1 could be averaged to obtain a “standard” model, the spread in the estimated values motivated a more refined attempt at determining the model parameters.

3.6.2 Parameter Measurement Based Model

In this section we discuss a more detailed attempt at measuring individual parameters for the model we developed in Section 3.5. To do this, we used a different building in Oslo, Norway. The building had a similar layout as the building in Stavanger, but was larger. In addition, the Oslo building had the same KlimaTak duct system installed. A blueprint of half the building in which we made measurements is shown in Figure 3.15. The measurements shown in the figure were made at a constant distance from a single long run of ventilation duct. The measurements were made about every 2 meters along a hallway that ran along the single duct run. The same Linksys IEEE 802.11g, WRT54G AP was connected at the end of a duct, again inside a metal box, with one antenna connector terminated and one connected to the duct with a monopole antenna. In Figure 3.15, the measured signal strength values are shown at the locations at which they were measured. The power levels are shown in dBm and the ventilation ducts in the building are shown in gray. The duct consisted entirely of circular spiral duct of 0.315 meter diameter. At least one ventilation outlet was present in each office or about 1 outlet every 2.4 meters. The signal strength was measured using a software utility designed to obtain the received signal strength directly from a Buffalo IEEE 802.11g wireless LAN card. The signal strength measurements from the LAN card were calibrated using the YellowJacket scanning receiver. The signal strength was recorded by visually watching the signal strength display for several seconds and recording the most commonly displayed value.

The intent of the measurements in Figure 3.15 was to compute the duct loss per meter for the propagation model. To do this, we extracted the duct length from the AP to each point on the duct closest to each measurement location. We then plotted the measured signal strength at each point versus the computed duct length distance. This plot is shown in Figure 3.16.

A linear best fit was computed for the data and is plotted as a blue line in Figure 3.16. The slope of this line is a good measurement of the duct loss per meter for the KlimaTak duct system. The value obtained of 0.34 dB per meter is somewhat higher than the 0.2 dB per meter value used previously. The added loss is accounted for in the decreased loss from the duct to the user location. The value of 0.34 dB per meter is also significantly higher than the value of 0.066 dB per meter reported in [20]. However, those measurements do not consider the loss due to leakage at the ventilation duct outlets. Extra loss from power leaking out the duct outlets could easily account for this added loss.

Once a number for the duct loss per meter parameter, α , is computed, the remaining parameters are the coupling losses, L_{in} and L_{out} and pathloss exponent, n . However, another physical insight into the propagation environment is that the ducts seem to leak power in many locations along its length. If this behavior were taken to the limit in which the duct were radiating power continuously along its length, then we would expect the propagation similar to a cylindrical loss with distance or a pathloss exponent of $n = 1$.

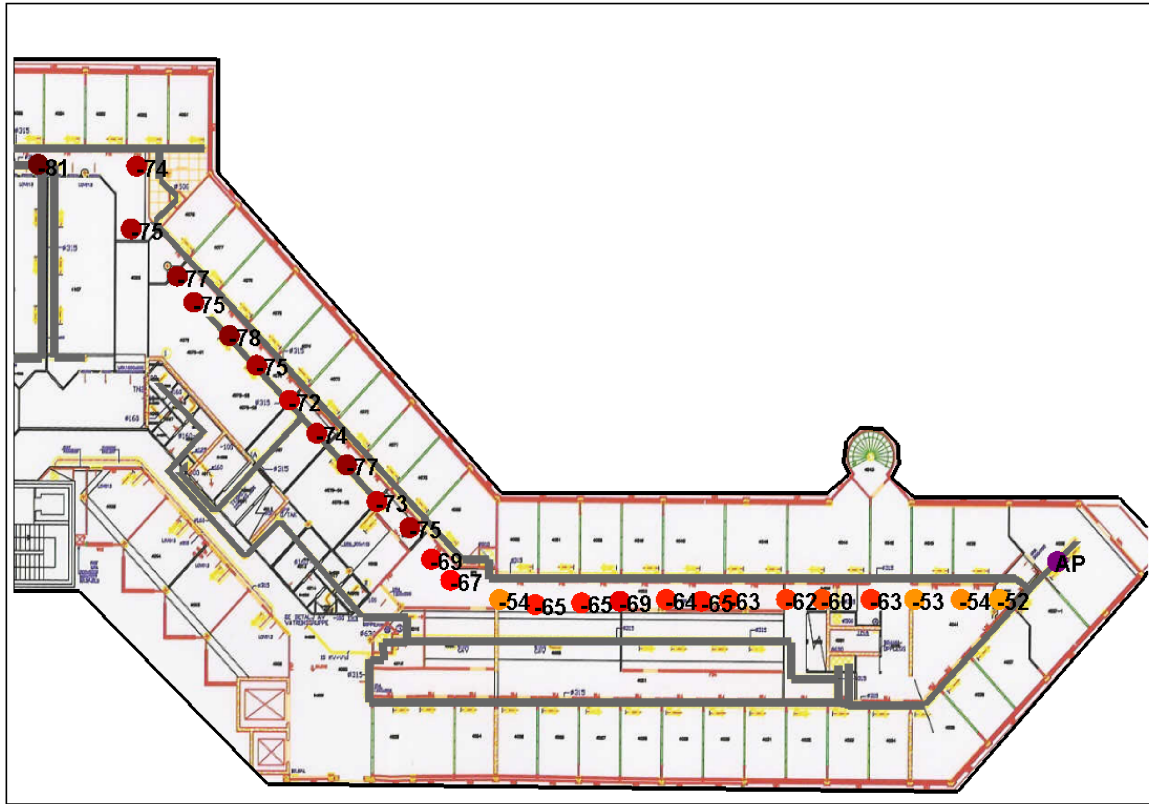


Figure 3.15: Blueprint of measurements in the Oslo building. The measured signal strength values are shown as points in the figure labeled in dBm.

When we account for the loss from the duct to the measured point and the duct junction loss, the value for $L_{in} + L_{out}$ works out to be 71.95 dB.

To check these measurement values, we used the values worked out from the Oslo measurements, to predict propagation in the Stavanager building, and compared the predictions to the 5 measurement sets described in Section 3.4.1. All 5 cases were simulated with the parameters calculated. A summary of the mean error and error standard deviation statistics for the 5 measurement cases are given in Table 3.2. Further, comparisons of simulated versus measured data for these cases are shown in Figures 3.17, 3.18, 3.19, 3.20, and 3.21. To visually compare the results, part (a) in each figure shows the measured points on top of a colored mist of the simulated signal strength overlaid on the building floor plan. In addition, we plotted the measured versus simulated values in part (b) of each figure.

Here we make a few comments to properly interpret Figures 3.17, 3.18, 3.19, 3.20, and 3.21. First, we note that in part (a) of the figures, when a measured point is the same color as the background color, the measured and simulated signal strength are about





Figure 3.16: Measured signal strength values from Figure 3.15 plotted versus computed distance through ventilation duct to AP.

the same. Thus, measurement points that are far from the simulated value appear to stick out while accurate simulations cause the measurement points to blend in. In this way we can see where in the building the simulations tend to be in error and where the measurements tend to be accurate predictors. Similarly, parts (b) of the figures allow us to see if there is any particular signal strength range that shows significant error. We observe that Figures 3.19(b), 3.20(b) and 3.21(b) show a linear scattering of error. However, Figures 3.17(b) and 3.18(b) show a bias for low measured signal strengths. That is, there tends to be higher prediction error at low measured signal strengths. Considering Figures 3.17(a) and 3.18(a), it appears the points with low measured signal strength but high predicted signal strength tend to be located more than 5 meters from an excited duct. Generally, the further from a duct you get the more likely there are to be walls and building partitions between the measurement point and the ducts. Therefore, a pathloss exponent of $n = 1$ is probably a poor model of the propagation environment as we consider

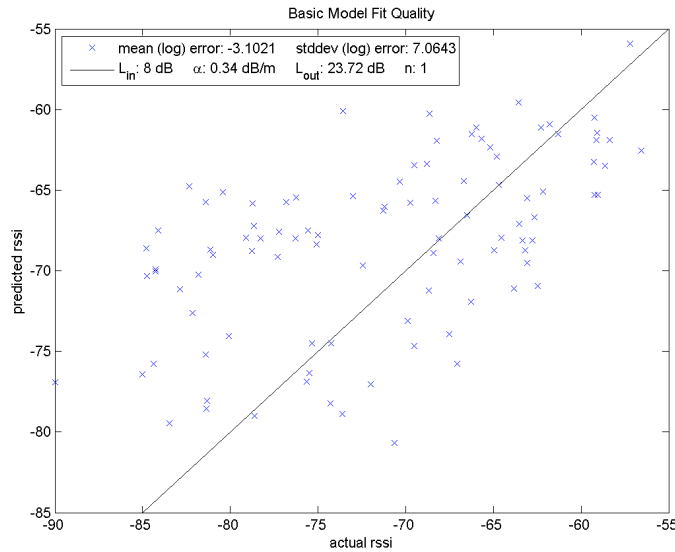
Table 3.2: Error statistics for the five measurement cases using the modified prediction model. The statistics shown were calculated based on predictions at 98 locations for which power measurements were made in each of the 5 cases.

Case	Mean Error (dB)	Error Standard Deviation
1	-3.10	7.06
2	-2.75	8.82
3	0.42	7.27
4	3.10	6.22
5	2.99	5.70

points further and further from a duct. It is for this reason that we further modify the propagation model as described next in Section 3.6.3.




(a) Visual Comparison




(b) Plot Comparison

Figure 3.17: Comparison of case 1 measurements to simulated signal strength. A visual comparison is shown in (a) by plotting the measurement points on top of a colored background based on simulated signal strength. The same points are plotted in figure (b). (a) allows us to see where points vary the most while (b) is more useful for seeing how much variation in prediction and measurement is present.




(a) Visual Comparison




(b) Plot Comparison

Figure 3.18: Comparison of case 2 measurements to simulated signal strength. A visual comparison is shown in (a) by plotting the measurement points on top of a colored background based on simulated signal strength. The same points are plotted in figure (b). (a) allows us to see where points vary the most while (b) is more useful for seeing how much variation in prediction and measurement is present.



(a) Visual Comparison



(b) Plot Comparison

Figure 3.20: Comparison of case 4 measurements to simulated signal strength. A visual comparison is shown in (a) by plotting the measurement points on top of a colored background based on simulated signal strength. The same points are plotted in figure (b). (a) allows us to see where points vary the most while (b) is more useful for seeing how much variation in prediction and measurement is present.

3.6.3 Propagation Model Modifications for Range Limitations

To improve the measurement based prediction model developed in Section 3.6.2, we modified Equation 3.1 to use a dual pathloss exponent. We maintain the same prediction model out to some critical distance r_c past which the pathloss exponent is increased to a higher value. The model then becomes

$$\begin{aligned} &\text{if } r < r_c: \\ R &= P_0 - L_{in} - \alpha d - 10 \sum_{m=2}^M J_m \log_{10}(m) - L_{out} - \\ &10n_1 \log_{10}(r) \end{aligned} \quad (3.2)$$

otherwise :

$$\begin{aligned} R &= P_0 - L_{in} - \alpha d - 10 \sum_{m=2}^M J_m \log_{10}(m) - L_{out} - \\ &10n_1 \log_{10}(r_c) - 10n_2 \log_{10}\left(\frac{r}{r_c}\right). \end{aligned} \quad (3.3)$$

As shown, we have renamed the original pathloss exponent n to be n_1 with the intention of keeping the same value, and introduced a second pathloss exponent n_2 .

To choose values for these parameters, we first choose r_c to be 5 meters in the horizontal direction. For the measurements we presented so far we have always used a difference in duct height and measurement height of 1 meter. For this reason, the actual value we use is $r_c = \sqrt{5^2 + 1^2} = 5.1$ meters. This number is based on the fact that the depth of almost all of the offices in the buildings we considered was about 5 meters.

For the second pathloss exponent we somewhat arbitrarily choose a value of $n_2 = 3$. Indoor propagation measurements often report a pathloss exponent of between 2 and 4 indoors, so a value of 3 is in this range. In addition, some trial and error of different values of n_2 for simulations with the original 5 Stavanger measurements show this to be a reasonable value. This is illustrated in Table 3.3. By comparing Table 3.3 to the prior error statistics in Table 3.2, it is clear that this final modification has improved our prediction accuracy, though both models should be sufficiently accurate for propagation modeling in other buildings.

3.7 Physical Interpretation of Propagation Model

The propagation models presented in Sections 3.5 and 3.6 are empirical in nature but are inspired by several different understandings of the underlying physical mechanisms. In this section we discuss the physical interpretation of the propagation models we developed and the mechanisms underlying its development.

Figure 3.22 shows a diagram view of the propagation model we have developed. As shown in the figure, the models we have developed start with the power transmitted by an AP and then we account for various loss mechanisms, ending with the power received

Table 3.3: Error statistics for modified propagation model accounting for restricted propagation range. The statistics shown were calculated based on predictions at 98 locations for which power measurements were made in each of the 5 cases.

Case	Mean Error (dB)	Error Standard Deviation
1	-0.38	5.94
2	-0.06	7.88
3	1.44	6.44
4	4.12	5.97
5	3.96	6.19

at a particular location in a building by a typical laptop card. The first loss mechanism is an antenna coupling loss parameter, L_{in} . This accounts for the efficiency of the antenna as a radiator into the ventilation ducts. We have simply used a value of 8 dB for this parameter as reported in [20].

The second loss mechanism we account for in the ducts is an exponential loss with distance through the duct, αd . We have chosen to model the loss through the duct as exponential for several reasons. The first reasons are pragmatic in that an exponential model seemed to be a good fit to the data collected and presented in Figure 3.15. Such a relationship is also a natural one from the point of view of an indoor wireless system designer who is used to looking at cable losses, which are presented as exponential losses with distance. As with coaxial cable transmission lines, waveguides with structures made of conductors with finite conductivity will also have an exponential loss with distance due to ohmic losses of currents in the waveguide structure. We might expect then that the waveguide propagation channel in ventilation ducts would also lead to an exponential loss with distance. However, ventilation ducts are overmoded waveguides. Thus, each waveguide mode will individually have an exponential loss with distance, but the falloff rate of the exponential loss with distance will vary significantly for different waveguide modes. As a result the net loss with distance behavior considering all of the waveguide modes is not likely to be exponential in nature.

Blaunstein has presented an explanation of where the exponential pathloss mechanism may come from in [34]. In this work, he considers the waveguide-like propagation that occurs in an urban street lined with buildings. He has considered such an environment as a waveguide with slits where buildings form the waveguide structure and open spaces separating the buildings are openings in the waveguide. These openings are considered to have randomly distributed lengths. Blaunstein demonstrates that the intermittent spacings along a waveguide formed by such distributions can be shown through analysis of Green's function to result in an average loss behavior that is exponential with distance. He also presents a comparison of his theoretical work with several




Figure 3.22: Schematic description of propagation model.

measurement data sets that also indicate an exponential loss with distance in similar urban environments.

Blaunstein's work is quite similar to our situation in the ventilation ducts. The ventilation ducts are overmoded waveguides with intermittent openings as air outlets. These duct openings are similar to the spaces between buildings in Blaunstein's work and thus the same mechanism should lead to an exponential loss with distance result. Our measurements confirm this as well. In [20], the authors found a loss of 0.066 dB per meter as a loss mechanism in straight ducts without any holes. However, we found a much more

significant loss with distance of 0.34 dB per meter in the ducts with holes. This indicates that the loss of power to the air outlets is a much more significant loss mechanism than attenuation due to ohmic losses in the duct walls and so the exponential loss due to the holes should dominate the ohmic losses.

Mathematically speaking, we could consider a duct segment with a single hole in it as shown in Figure 3.23. If we consider the power that leaves the duct through the hole as proportional to the power that enters the duct, P_{in} , then we can represent the power that exits the duct into the next duct segment, P_{out} as $P_{out} = P_{in} - aP_{in}$.




Figure 3.23: Power losses in a duct segment with a hole in it.

If we further subdivide the ventilation duct into smaller differential segments of length Δz and consider each segment to have a single hole, we can represent the power entering the segment at a distance z along the duct length as $P(z)$ as shown in Figure 3.24. The power leaving such a duct segment then would be $P(z + \Delta z)$. The power leaking out the hole in each duct segment then should be proportional to the segment length so we represent this as $\alpha \cdot \Delta z \cdot P(z)$ where $\alpha = aL$ and L is the average distance between duct holes. The resulting relationship between the power entering and exiting differential duct segment can be expressed as

$$P(z + \Delta z) = P(z) - \alpha \cdot \Delta z \cdot P(z), \quad (3.4)$$

which can be re-arranged to express the change in power versus distance along the duct length as

$$\begin{aligned} \Delta P &= P(z + \Delta z) - P(z) = -\alpha \cdot \Delta z \cdot P(z) \\ \frac{\Delta P}{\Delta z} &= -\alpha \cdot P(z). \end{aligned} \quad (3.5)$$

This is a discrete approximation to the differential equation

$$\frac{dP}{dz} = -\alpha P(z), \quad (3.6)$$

with the well known solution

$$P(z) = Ce^{-\alpha z}. \quad (3.7)$$

Obviously we have an exponential decay in the power with distance and thus the motivation for using an exponential decay term in our pathloss model to account for the loss with distance through the duct.




Figure 3.24: Power losses if we consider the duct broken up into many small duct segments each leaking a small amount proportional to the length of the segment.

The third loss mechanism in our pathloss models is the losses at duct junctions. We have simply modeled this as a split loss based on the number of branches at the junction as would be done for an ideal power divider. This behavior, assuming symmetric junctions, should be accurate on average, but will vary somewhat for specific duct junctions. It will also vary for non-symmetric junctions and for those with different duct cross-sectional sizes and areas in different branches. However, for the sake of model simplicity, we have not tried to model these effects.

The fourth loss mechanism is an output coupling parameter, L_{out} . This parameter is analogous to a reference pathloss parameter present in indoor, log-normal propagation

models [35]. Such models use a measurement of a reference pathloss at a distance close to the transmit antenna and model the pathloss as increasing logarithmically with distance from that reference distance. We have used such a model for our fifth loss mechanism, the loss based on the distance from a user to the nearest duct. Thus, the output coupling parameter is a coefficient that takes into account the losses coupling out of the duct to a distance 1 meter away from the duct and then a log-normal with distance loss is applied from that 1 meter reference distance to the users location. While this technique is slightly similar to the work done in [31] with leaky feeder propagation models, it is different in that it only considers pathloss to the nearest point on the nearest duct to the users location.

The log-normal propagation model we have used applies a significantly lower pathloss exponent, 1, than is typically used with indoor models. This value is used because it is what we would expect if the duct was an infinitely long, straight cylindrical radiator. In such a case, the loss with distance would go as $1/r$, corresponding to a pathloss exponent of 1 in the log-normal pathloss model. Since there are a large number of air outlets in the ventilation ducts and usually no building partitions between the ducts and users, such a loss with distance is not unreasonable.

3.8 Conclusions

In this chapter, we have outlined the measurements, modeling and statistical analysis of the measurements we made in two buildings in Stavanger and Oslo, Norway. The intent of the measurements was to develop a useful propagation model to allow the design of duct based distribution systems. For this reason, we have attempted to develop a model that can be used to design ventilation duct based systems. Further, we would like a model that can be implemented simply in software with the minimum amount of modeling or that can be reduced to a simple rule of thumb for paper design techniques. We believe we have managed to do this with the model described by Equations 3.2 and 3.3. The computed values for the model are summarized in Table 3.4.

To develop this model, we first investigated the degree to which ventilation ducts radiated power both inside and outside of buildings. Then, using test setups we took a number of measurements of the radiation from ducts with APs connected to develop a propagation model. The model described in this chapter is simple enough to be summarized as a few simple rules for paper design and can be implemented in software for simulation with a bare minimum of building modeling. For instance the model does not take into account the duct size and dimension. This eliminates the need to model different types of ducts. Unfortunately, this may also limit the accuracy of the prediction model. The measurements presented in this chapter indicate that the model can already provide significant accuracy to new buildings not yet encountered, but the degree to which the model can be generalized will need to be developed as experience is gained with additional buildings.

Table 3.4: Final propagation model parameters. α , L_{in} and L_{out} were computed as a fit to 27 measurement points in the Oslo building. The remaining model parameters were chosen based on physical meaning and the quality of fit to the measurements made at 98 locations in each of the 5 cases in the Stavanger building.

Parameter	Symbol	Value
AP Coupling Loss	L_{in}	8 dB
Duct Loss Per Meter	α	0.34 dB/m
Duct to Building Environment Coupling Loss	L_{out}	63.95 dB
First Pathloss Exponent	n_1	1
Pathloss Breakpoint Critical Distance	r_c	5 meters
Second Pathloss Exponent	n_2	3

Some potential ways to improve our propagation model would be to include losses for different duct junctions including bends, not just multi-way splits. Further, the model currently suffers from relying on the nearest duct to have the highest signal strength. This is not always the case and some simulation points may be closest to a duct that is far from an AP but still close (but not as close) to a duct that is near an AP. Thus, the model may at times underestimate the actual signal strength at a location. This situation could be avoided by simulating the signal strength from the closest points on all ducts in the environment and either adding them up or taking the largest value. This would be more difficult to implement in software, but could be added if the added complexity was deemed useful. However, the model as it is expressed herein seems to be a good tradeoff between complexity and accuracy.

Throughput Measurements and Modeling

4.1 Introduction

Up until this point, we have primarily discussed how to layout a ventilation duct based wireless LAN using signal strength as a figure of merit for measurements of a single AP. In this chapter, we extend our work to consider whether ventilation duct systems can provide adequate throughput to areas covered and if there are ways to improve the throughput provided through the use of antenna diversity or additional amplification. Most importantly, we put together everything learned in this chapter and previous chapters and implement a full ventilation duct based system in a 3 story, 2000 square meter building. We analyze the throughput performance of the entire system and find that the concept of a ventilation duct based system is quite promising.

4.2 Throughput versus Signal Strength Measurements

To ensure that the high delay spread or other concerns do not prevent an IEEE 802.11g system from achieving good throughput, we have performed several throughput measurements of a duct based system. The measurements we present here were measured in the building in Oslo, Norway, which was described in Section 3.6.2. We measured the signal strength as well as the throughput in these measurements to enable the correlation of these parameters to be examined. To measure the throughput, we used the same open source throughput measurement tool described in Section 2.3.1 called Iperf [28]. We measured the throughput at each location using tcp packets sent for 20 seconds using a 1 Megabyte transmission window. The measured throughput was displayed to the user every 3 seconds to allow measurement monitoring. Each of the measured throughput values is plotted on a picture of the building blueprint as shown in Figure 4.1. Again, each grid square in the figure is 4.8 meters on a side. An AP was connected to the ventilation duct as shown in the figure. Once again, the AP was sealed inside a metal box, one antenna terminated and the other connected to a monopole antenna.

To compare the measured throughput to signal strength, we used uplink rather than downlink measurements. To obtain this measurement we polled the Linksys WRT54G AP for the signal strength that it measured on the uplink from the wireless LAN card at the user location. We chose this AP because of its low cost and runs an embedded Linux kernel which has been extensively explored by Internet communities. This allowed us low level control of the AP and the ability to upload a custom Linux kernel if we desired. As a result, we were able to use techniques described by Seattle Wireless [36], to directly access the wireless driver on the AP. Our approach was to use a security hole in the AP configuration webpage to effectively run console commands on the Linux system. We wrote a custom perl script that would repeatedly open a HTTP session to the AP and query the AP for the signal strength of a client with a specific MAC address. This process takes about 3 seconds so it was repeated for 30 seconds and the signal strengths averaged together and recorded.

After taking measuring at the 55 points shown in Figure 4.1, a plot of the throughput versus signal strength for all locations was made as shown in Figure 4.2. A linear regression analysis resulted in the equation

$$T = 0.94 \cdot R + 76.54. \quad (4.1)$$

Since the slope of the linear regression is close to 1, we felt it would be useful to do a fit that forced the slope to be 1. Such a fit would be more useful for a quick rule of thumb because you could convert a signal strength to a throughput by just adding a constant. The result of this second fit is

$$T = R + 81.15. \quad (4.2)$$

Consequently, a quick rule of thumb is that 81 added to the RSSI in dB gives the expected maximum throughput in Mbps.

As can be seen from Figure 4.2, a linear fit performs reasonably well. However, owing to the spread in the experimental data, along with the variations in RSSI measured by different LAN cards and APs, Equations 4.1 and 4.2 should only be taken as rough estimates of actual performance and may need to be adjusted for particular NICs or APs. We should also note that we have previously presented similar throughput prediction techniques in [37].

An example of the accuracy of these estimates is shown in Figure 4.3. The figure shows a contour plot of the predicted throughput using Equation 4.1 and the same measured throughput data given previously. If the predictions are working well, the points in the area should be as bright or more brightly colored than the background. As can be seen from the figure, the results are reasonable, but a number of discrepancies can be seen. This is expected given the uncertainties discussed above along with the intrinsically variable nature of throughput measurements.




Figure 4.2: Plot of measured throughput versus AP reported (uplink) signal strength

Table 4.1: Measured throughput at two measurement locations for various antenna diversity settings. Each table entry is based on a single measurement of a 1 Megabyte file transfer.

Location	AP Antennas	Throughput	Transmit Rates
1	A	11.49 Mbps	24 Mbps
	B	12.71 Mbps	24 Mbps
	Both	12.32 Mbps	24 Mbps
2	A	0.76 Mbps	2-5.5 Mbps
	B	0.17 Mbps	1-2 Mbps
	Both	0.99 Mbps	2-5.5 Mbps

4.3 Throughput Improvements

To improve the throughput inside a building, we considered several incremental improvements to the system layout techniques described previously. The most obvious of these was using the antenna diversity built in to most APs. We also considered bi-directional amplifiers. We discuss these three possible improvements below.

4.3.1 Antenna Diversity Improvements

Previous measurements have all used a Linksys WRT54G AP inside a sealed metal box with one antenna terminated and one connected to a ventilation duct antenna. To determine improvements in performance from using both antennas, we connected an AP to the ventilation duct in the same location as measurement case 2, discussed in Section 3.4.1 and shown in Figure 3.7. We left the AP out of the metal box this time and connected both antenna ports to monopole duct antennas. This allowed us to enable the AP antennas one at a time or use full diversity, again using the wireless driver on the Linksys AP by way of the security hole in the firmware [36]. We then performed throughput measurements at two different locations in the building, one at a good signal strength location, one at the edge of the covered area. We measured the throughput back to a computer connected via 100 Mbps ethernet to the AP from a laptop with a IEEE 802.11g NIC. A program called QCheck was used to measure the throughput. This is a free program once available from NetIQ [38] but now distributed by IXIA [39]. The program was used in TCP mode to transfer a test file 1 Megabyte in size from the wireless client to the laptop connected via ethernet to the AP. Three measurements were performed and the results averaged. The data transmission rate on the wireless link was also noted and the range of values used over the measurement was recorded. These measurements were repeated for three cases in which each of the APs antennas was enabled individually and with both antennas enabled. To clarify, we disabled both transmit and receive diversity in the AP when only one antenna was used. The measurement results are given in Table 4.1.

As can be seen in the table, turning on both antennas resulted in a throughput better

than the worst throughput using a single antenna in both cases. We therefore conclude that even though we might achieve slightly worse throughput than with the best antenna in some cases, the added reliability from the antenna diversity implies that diversity should be used in duct based installations.

Incidentally, when using the AP with only a single antenna connected, it is crucial for accurate measurements to disable the diversity of the terminated antenna. The Linksys WRT54G AP uses a simple, yet common diversity algorithm. When the AP receives a packet that fails a cyclical redundancy check (CRC) or is otherwise known to be invalid, it switches to the other antenna (both for transmit and receive). If one antenna is terminated or if the two antennas are connected to different ventilation ducts, this can cause disruption of the link. Near the AP, there are rarely bit errors and so measured throughput will function normally. However, as you move towards the edge of the covered area, a bit error is more likely to occur. When it does, the AP will switch antennas. If it switches to the terminated antenna or to the antenna connected to a different duct, the AP will never again receive a packet from the user, valid or not. As a result, the AP will simply continue to listen for the remote user on the wrong antenna and the user is effectively disconnected from the AP.

4.3.2 Bi-directional Amplifier Improvements

In Section 3.4.2, we compared conventional wireless system design with duct based design. Unlike a conventional system, the ducts distribute the signal more evenly through a building so that there are no areas that have the high signal strengths normally found close to an AP. For this reason, it is conceivable that an amplifier could be installed at an AP to improve the coverage area and throughput provided by a duct system. A bi-directional amplifier (BDA) would also improve the uplink to the AP by increasing the receive sensitivity of the AP.

We tested two commercially available bi-directional amplifiers, but obtained similar results with both. We present measurements with one of the amplifiers here. The amplifier used was a SSB Electronic Model WL-2400, 1 watt bi-directional amplifier [40]. It has a specified transmit gain of 10 dB, a receive gain of 14 dB and a receive noise figure of 1.4 dB. We again used the same AP location as measurement case 2, discussed in Section 3.4.1 and shown in Figure 3.7. We then made measurements in three locations. The first was at a high signal strength location, the second at a moderate signal strength location and the last at a low signal strength location. To check the performance of the BDA, we compared throughput and signal strength values at each of these three points with and without the BDA installed. To perform this test, we put the AP in a metal box, terminated one antenna connector of the AP, and connected the other to either the BDA and the duct antenna or directly to the duct antenna. As before, we used a monopole antenna. Antenna diversity was again disabled for the AP. At each of location we recorded the signal strength measured by the AP, the throughput, and the wireless transmit rate. A laptop

Table 4.2: Throughput and RSSI measurements with and without a bi-directional amplifier installed. Each RSSI measurement is based on the average of about 20 measurements and every throughput entry in the table is based on a single 60 second measurement.

Without Amplifier		
RSSI	Throughput	Transmit Rate
-71.2 dBm	20.2 Mbps	48-54 Mbps
-83.9 dBm	9.2 Mbps	24 Mbps
-89.4 dBm	1.0 Mbps	2-5.5 Mbps
With Amplifier		
RSSI	Throughput	Transmit Rate
-54.9 dBm	15.7 Mbps	36-48 Mbps
-69.3 dBm	8.8 Mbps	24 Mbps
-78.8 dBm	3.3 Mbps	11 Mbps

was connected via 100 Mbps ethernet cable to the WRT54G AP for throughput testing. Another laptop with a wireless NIC was moved from location to location, but was left in the same exact location when switching in and out the amplifier. The signal strength was recorded by polling the AP for the signal strength it measured from the client card once every 3 seconds for 60 seconds. Throughput was measured using the Iperf software to transmit TCP packets to the remote laptop. The test was performed using a TCP window size of 1 Megabyte and was conducted for 60 seconds with user status updates every 3 seconds. During the throughput tests, the range of wireless transmission rates the card operated at were recorded. All of this data is presented in Table 4.2.

The data in Table 4.2 was unexpected. While the signal strength measurements indicate that the amplifier is functioning properly and providing gain to the system, there does not appear to be any consistent improvement in the throughput measured. Strangely, at the higher signal strength locations, the amplifier actually made the throughput worse, though it did seem to improve things at the low signal strength level. In an attempt to explain this behavior, we investigated the specifications of various wireless LAN APs and NICs. We discovered that not only do wireless LAN cards require a certain signal strength to operate at a particular transmission rate, but each transmission rate can only tolerate a particular maximum delay spread. Not all manufactures publish these specifications, but some example specifications for an AP and a NIC are shown in Table 4.3.

The values in Table 4.3 provide a possible explanation for the unexpected throughput results presented in Table 4.2. If the throughput measured without the external BDA is being limited by the delay spread of the channel, and, then the improved signal to noise ratio (SNR) provided by a BDA will not improve the throughput. However, if the throughput is limited by signal strength rather than delay spread, as it appears to be in the third measurement in Table 4.2, then the BDA can provide an improvement in the

Table 4.3: Reported maximum delay spreads tolerated at different transmit rates of a Cisco AP, [41], and a Cisco Network Interface Card (NIC), [42].

Rate Type	Transmit Rate	Cisco AP [41] Maximum Delay Spread	Cisco NIC [42] Maximum Delay Spread
g	54 Mbps	120 ns	70 ns
	48 Mbps	120 ns	90 ns
	36 Mbps	240 ns	100 ns
	24 Mbps	240 ns	160 ns
	18 Mbps	300 ns	220 ns
	12 Mbps	300 ns	250 ns
	9 Mbps	300 ns	250 ns
	6 Mbps	300 ns	400 ns
b	11 Mbps	140 ns	130 ns
	5.5 Mbps	300 ns	200 ns
	2 Mbps	400 ns	300 ns
	1 Mbps	500 ns	350 ns

throughput. We will discuss this explanation further in Section 4.4.

4.4 Full Building Implementation

To further explore the throughput achieved in a complete duct based RF distribution system, we conducted the next set of measurements in a 3 floor building in Stavanger, Norway. Coverage for the entire building was provided using 4 APs connected to different ducts in the building. The covered area was about 2000 square meters, though the building itself was slightly larger than this. Two APs were connected on the second floor since it had two separate duct networks. One AP was also placed on each of the two other floors. Unlike previous measurements, a slightly different AP was used. Four Linksys WAP54G APs were installed. To control the APs with our software, these APs were upgraded to a custom firmware called MustDie, version 2.07r1 [43]. In addition, both antennas were connected to monopoles inside the ducts. The antennas were placed 90 degrees apart along the circumference of circular ducts and in the middle of two adjacent sides of rectangular ducts.

Measurements of signal strength and throughput were made throughout the building. At each location, custom software discussed in Appendix B was used to force the client card to rescan for the best AP and reconnect to it. After doing this, the signal strength and throughput was measured at the location. The signal strength was measured using a utility program for the Buffalo Wireless Air Station, WLI-CB-G54A, IEEE 802.11g wireless LAN card used as the client NIC. The signal strength was estimated by watching the signal and estimating the mode of the values displayed while a 30 second throughput measurement was performed. The throughput measurement was exactly the same as

described in Section 4.2 with the same configuration parameters. The results of these measurements are shown in Figures 4.4, 4.5, 4.6, 4.7, 4.8, and 4.9. In all six figures, each AP has a different color. Each measurement point is outlined with one of the 4 AP colors to indicate which AP it was connected to when the measurement was taken.

Pictures of the best signal strength at all locations in the building are shown in Figures 4.4, 4.5, and 4.6. Further, the measured throughput at each location in the building is shown in Figures 4.7, 4.8, and 4.9.

Several interesting aspects of the above measurements stand out. Firstly, the overall coverage area is smooth and even as was expected, though there are some surprisingly high signal strengths in some locations. While in most cases the nearest duct with the nearest AP seems to provide the best signal strength, the AP on the first floor seems to be providing more signal strength to the floor above it than we would predict. Our best explanation as to the reason for this is that the ventilation outlets on this floor are cut towards the top of the ducts whereas they are cut in the center of the ducts on other floors. This may serve to aim power upward to higher floors.

In terms of throughput, the throughput measured is generally good at all locations and there is little variation from location to location in the building, especially on the first and third floors. However, there are some locations on the second floor, shown in yellow, that have higher throughput than most. It is interesting to note that these locations tend to be fed by ducts with few or no junctions and that are entirely circular. It seems likely that this sort of duct would have lower delay spread than ducts with changes in shape and many junctions. It is in these sorts of ducts that the delay spread may not be limiting the throughput. To explore this concept further, we have plotted the signal strength versus throughput for all of the measurement points shown in Figures 4.4, 4.5, 4.6, 4.7, 4.8 and 4.9. These measurements are shown in Figure 4.10.

From Figure 4.10, it is clear that every location in the building has a good throughput and a high signal strength. While it is tempting to apply the previously calculated throughput prediction model to the measured data, shown in the figure, that is not actually valid. One reason for this is that the throughput prediction model presented in Section 4.2 was for a different AP. The intercept points for this AP will most likely be different than for the other AP. The difference should be small, but different cards and APs have varied up to 10 dB from one another in our experience. Despite this issue, we can still make some observations about the measured throughput. First, we notice that the throughput often does not increase with the signal strength as expected. This is also apparent in Figure 4.11 where we plot the range of wireless transmission rates seen at each location versus the measured signal strength at that location.

Figure 4.11 illustrates that despite the high signal strengths present throughout the building, rarely does the wireless LAN operate at the highest data rates. Instead, 18 to 36 Mbps data rates are the most common. In [11] where Nikitin et. al. reported a maximum

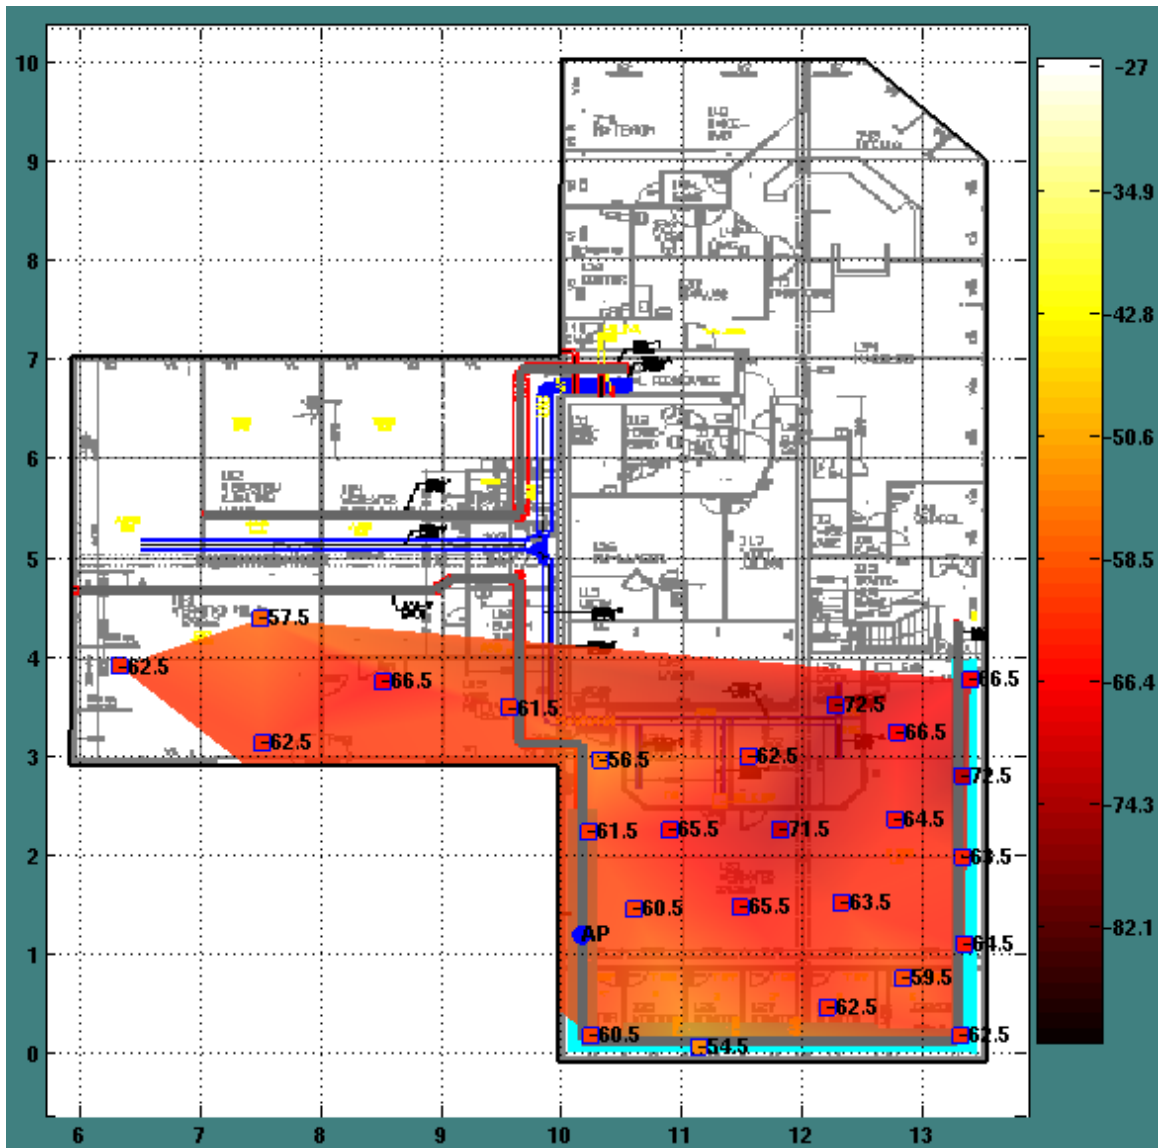


Figure 4.4: Plot of best signal strength values measured on the first floor of the building. The outline of each square indicates the AP which had the best signal at each location.

delay spread of 200 to 300 nanoseconds in duct based channels. Interestingly, these data rates have maximum delay spreads, as shown in Table 4.2, close to these values. Thus, we feel that we could probably improve this installation using an AP that has a greater delay-spread tolerance, but it is doubtful that BDAs would improve the throughput of the system.

One positive outcome of this result is that almost all users of the system will operate

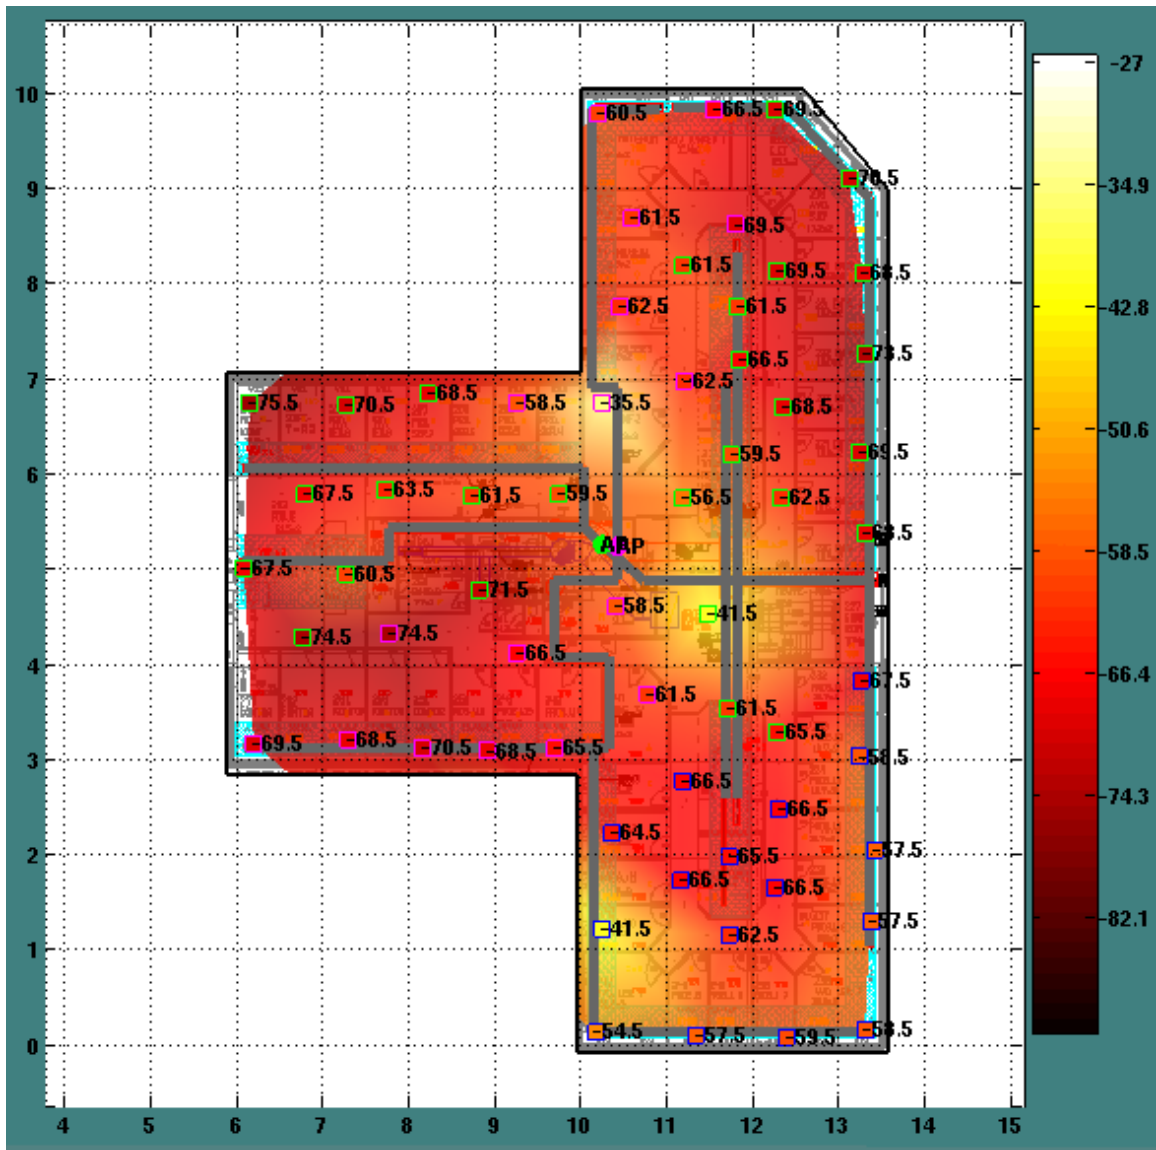


Figure 4.5: Plot of best signal strength values measured on the second floor of the building. The outline of each square indicates the AP which had the best signal at each location.

at about the same throughput. This means that all users should experience about the same throughput even at different locations in the building. Further, since there are not any locations in the building with low signal strength or low data transmission rates, one user cannot drag down everyone else's throughput, since it is a shared communications medium. As a result, the duct system provides an unexpected improvement to the




Figure 4.6: Plot of best signal strength values measured on the third floor of the building. The outline of each square indicates the AP which had the best signal at each location.

fairness of an IEEE 802.11g system due to its uniform coverage.

4.5 Conclusions

In this chapter, we have explored how a ventilation duct based system performs if designed and laid out using the techniques and models presented in Chapters 2 and 3. We have found that in the simple, long runs of cylindrical ducts with regularly spaced ventilation duct outlets, we can achieve maximum available throughputs, and better SNRs

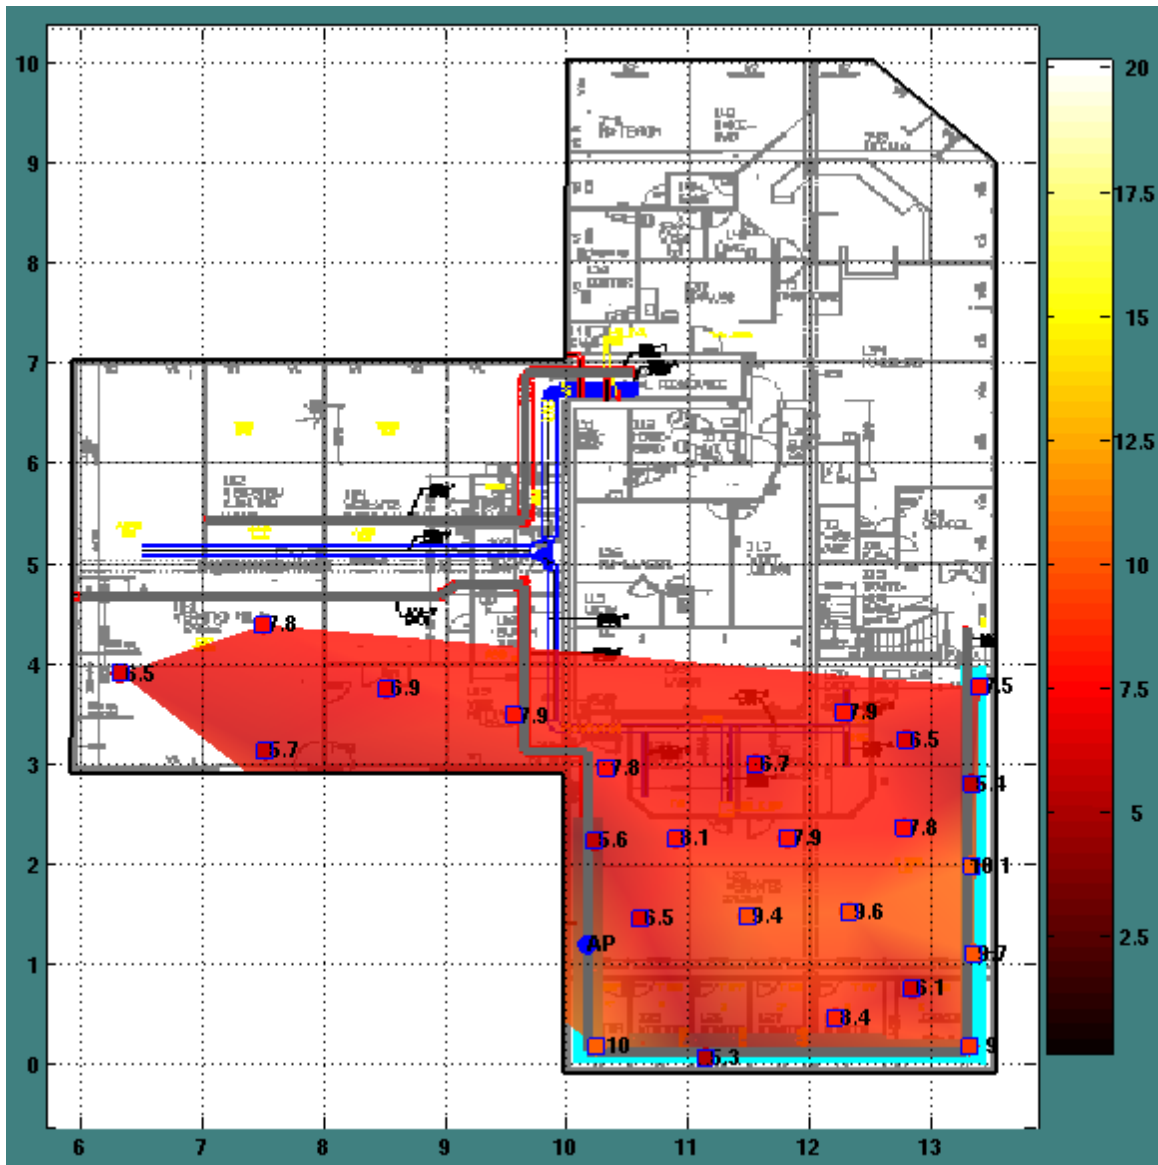


Figure 4.7: Plot of throughput values measured on the first floor of the building. The outline of each square indicates the AP which had the best signal at each location.

tend to lead to better throughput levels. In fact, we presented two simple equations that model this general behavior and allow simple rule of thumb calculations of the expected maximum throughput available at a particular location for IEEE 802.11g networks.

We also investigated two possible improvements to ventilation duct based networks. We discovered that using the antenna diversity available in some IEEE 802.11g APs will function as designed and will ensure throughput is at least as good or better than the




Figure 4.8: Plot of throughput values measured on the second floor of the building. The outline of each square indicates the AP which had the best signal at each location.

worst from either antenna alone. We also discovered, surprisingly, that commercially available BDAs designed for wireless LAN use improve measured signal strengths but do not seem to improve the available throughput.

To further explore the performance of the ventilation duct system, we installed 4 APs in a 3 floor, 2000 square meter building. The full system provided extremely even coverage throughout the building. In fact, most areas of the building operated in the 18 to 36 Mbps




Figure 4.9: Plot of throughput values measured on the third floor of the building. The outline of each square indicates the AP which had the best signal at each location.

transfer rates during throughput testing. As a result users throughout the building can expect consistent network performance and the system overall provides better fairness than conventional IEEE 802.11 network layouts.




Figure 4.10: Throughput versus RSSI as measured across all three floors in full building layout.




Figure 4.11: Wireless data transmission rates versus RSSI as measured across all three floors in full building layout. Each bar indicates the range of transmission rates seen.

5.1 Introduction

In Section 4.3.1, we measured the throughput improvement using the antenna diversity built into an IEEE 802.11g AP. Antenna arrays used for diversity is a natural idea for improving the performance of a wireless link. However, various authors (e.g. [44, 45]) have pointed out that antenna arrays can provide more than just a diversity benefit. In this chapter we investigate using antenna arrays in alternative ways. First, we describe a new antenna array that we have discovered that can be used to excite a single mode inside certain ventilation ducts. This unique array can actually be used to excite two different single modes depending on how the elements are driven. Secondly, we investigate how a simple two element phased array can be used to affect the connectivity and signal strength of an IEEE 802.11b link. These two experiments show the potential applications of antenna arrays for use in ventilation duct channels. We will further investigate the uses of antenna arrays in Chapter 6, but this chapter serves as an introduction to that material by presenting specific applications.

5.2 Single Mode Antenna Array

In Chapter 4, we discovered that there are cases in which the delay spread of a channel, rather than the signal to noise ratio, limit the throughput that can be achieved in that channel. For this reason, antenna designs capable of reducing the delay spread of channels are of interest. Towards this end, a new antenna array designed to excite primarily a single mode has been designed. The array is practical for easy installation in ventilation ducts and can be incorporated to complement existing 2 antenna diversity techniques in current hardware. These characteristics make it an attractive alternative to the single-mode dipole antennas designs presented in [26]. Below, we discuss the antenna array design we have invented and some advantages of it.

5.2.1 Antenna Element Design

The antennas are designed to operate in the 2.4 GHz, ISM (Industrial, Scientific and Medical) frequency band using 315 millimeter or 304.8 millimeter (12 inch) diameter ducts. The 2.4 GHz frequency band is allocated throughout the world, while we have encountered the 315 mm diameter frequently in our installations in Norway, and the 12

inch diameter is widely used in the United States. For this frequency band, 17 modes can propagate in the 304.8 mm diameter and 18 modes can propagate in the 315 mm diameter duct. Most unfortunately, the cutoff frequency of the highest order mode in the 315 mm diameter duct falls in the 2.4 GHz frequency band, potentially causing sharp nulls in the frequency range of interest.

The antenna array design uses two identical, properly driven array elements to excite a single mode. Thus, we begin by analyzing the mode excitation of a single array element and then analyze how the elements are driven to excite primarily a single mode. To describe how the antenna was designed, we will describe the process used to identify the final antenna element. This procedure was similar to that used by [26]. In fact, we rely heavily on simulation software developed by Pavel Nikitin, and much of the following analysis is derived from his work. As was done by Nikitin in Chapter 3 of [26], we have calculated two parameters, \mathcal{P} and \mathcal{T} , for monopole antennas in a 315 mm diameter cylindrical duct for all possible monopole lengths. \mathcal{P} is a mode power coefficient, which is the power carried by the most strongly excited mode normalized by the power carried by all modes. \mathcal{T} indicates the quality of the antenna's impedance match and is the power transmission coefficient of the antenna. The original design goal outlined by Nikitin was to maximize both the \mathcal{P} and \mathcal{T} parameters. Maximizing \mathcal{P} maximizes the amount of power in just one mode. Maximizing \mathcal{T} maximizes the impedance match of the antenna to the duct. The results of our calculation are shown in Figure 5.1.

At this point in the analysis, we diverge from Nikitin's analysis. We are most interested in the power transmission coefficient, \mathcal{T} and we will disregard the parameter \mathcal{P} . As can be seen in Figure 5.1, the parameter \mathcal{T} is maximum at five monopole lengths. Our intention is to simply consider the mode content of those five monopole lengths and see if we can use other techniques to reduce the mode content. With this in mind, we again used software and techniques developed by Nikitin in [26] to simulate the modes excited by a monopole of the lengths that maximize the parameter \mathcal{T} . These lengths are 33, 86, 156, 209 and 279 millimeters. The first length corresponds to the quarter-wavelength monopoles we have been using extensively throughout this work. The results of these simulations are shown in Figure 5.2.




Figure 5.1: Calculation of \mathcal{P} and \mathcal{T} parameters for a monopole antenna in a 315 millimeter diameter cylindrical ventilation duct.

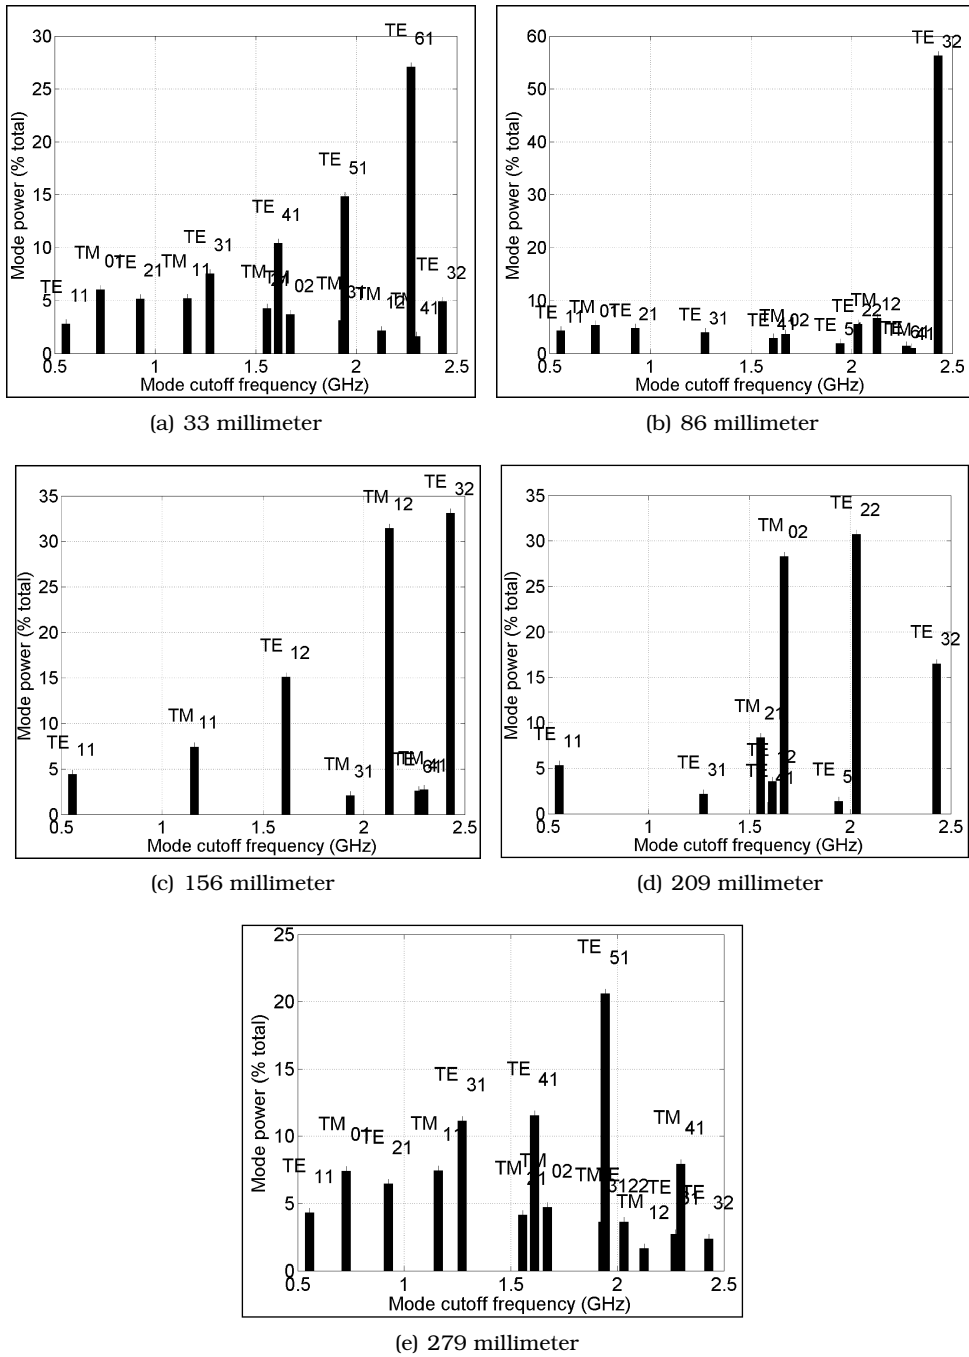


Figure 5.2: Predicted mode content for various length monopole antennas. Each plot is labeled with the length of the monopole antenna. All plots are calculated using software written by Pavel Nikitin [26].

In viewing the mode content, we were most interested in a mode mixture that would allow us to leverage multiple antennas to cancel or emphasize some of the most common modes. By expanding the number of antenna elements considered, we allow ourselves to use a new mechanism for finding ideal mode mixture. Now, rather than looking for a mode mixture dominated by a single propagating mode, we can look for a mode mixture in which a few modes are present at about the same power and all other modes are only slightly present. Then we could attempt to use the added parameter of multiple antenna elements to try and select one of the dominating modes while canceling out others. The most intuitive technique for doing this seemed to be to consider circumferential antenna arrays. That way we could look at the cross-sectional field patterns of the major modes of each of the cases in Figure 5.2 for patterns that are inconsistent with one another. We would also like to avoid the TE_{32} mode because it is in the frequency range of interest and could introduce significant dispersion issues. Thus, the overwhelming presence of TE_{32} in Figures 5.2(b) and 5.2(c) would seem to immediately eliminate those mode distributions. Also, the mode distributions in Figures 5.2(a) and 5.2(e) seem to be unsuitable for array use because of the predomination of differing amounts TE_{m1} modes, which have field patterns not useful to us. We say this because it is not clear how to create a circumferential array of elements with particular excitations that would result in a single dominant TE_{m1} mode. However, except for the presence of some TE_{32} , the mode mix in Figure 5.2(d) has a lot of potential. Thus, we tried slightly adjusting the length of the monopole around 209 mm to see if the mixture could be improved.

As seen in Figure 5.3, with a monopole length of 205 mm, the TE_{22} and TM_{02} modes are present equally and the TE_{32} mode is greatly reduced. With this advantage, we next consider the field patterns of the two major modes present in the setup, which are shown in Figure 5.4 as excerpted from [29].

The electric field (E-field) patterns of the two modes shown in Figure 5.4 present an intriguing opportunity. Since the E-field pattern of the TM_{02} mode is circularly symmetric, but the E-field pattern of the TE_{22} mode alternates in sign in each quadrant of the cross-section, a circumferential antenna array with elements spaced 90 degrees apart around the diameter is an antenna array nirvana, providing the desired mode selectivity. In this geometry, each array element can be thought of as exciting both the TM_{02} and TE_{22} modes with relative phases determined by the field patterns of the modes. For example, if two 205 mm long monopoles were driven in phase, both TM_{02} modes would be excited in phase resulting in a strong mode excitation. However, since the radial component of the TE_{22} mode reverses sign every 90 degrees around the circumference, the modes excited by the array elements will be 180 degrees out of phase and cancel. Various other modes will be enhanced or degraded, but the result should be a mode mixture dominated by the TM_{02} mode. By a similar argument, if the array elements are driven out of phase by 180 degrees, the mode mix will be dominated by TE_{22} .




Figure 5.3: Mode content of adjusted monopole calculated using software written by Pavel Nikitin [26].

This idea for an antenna array is quite compelling. Since both elements need only be driven either 0 or 180 degrees out of phase, the array can be implemented with minimal cost. In addition, the ease of installation of monopole antennas is maintained while only needing longer antennas. It should be noted that some non-conducting material might be needed on the array elements to prevent the antennas from shorting to one another since both are longer than the radius of the duct. The phase shifts of 0 and 180 degrees can of course be obtained by simply ensuring the antenna connections are made with equal or opposite polarities. However, one could also use a 180 degree hybrid power splitter for an even greater advantage as shown in Figure 5.5. If you connect the AP antenna ports to the sum and difference ports of the hybrid, then each AP antenna port will excite a single mode in the duct, but each port will excite a different single mode. This is because the signal on the sum port will excite the duct antennas in phase and excite the TM_{02} mode




Figure 5.4: Cross-sectional field patterns of the major modes excited by a monopole 205 mm long in a 315 mm diameter cylindrical duct. The solid lines indicate electric fields and the dashed lines indicate magnetic fields. Excerpted from [46].

and the difference port will excite the duct antennas 180 degrees out of phase and excite the TE_{22} mode.

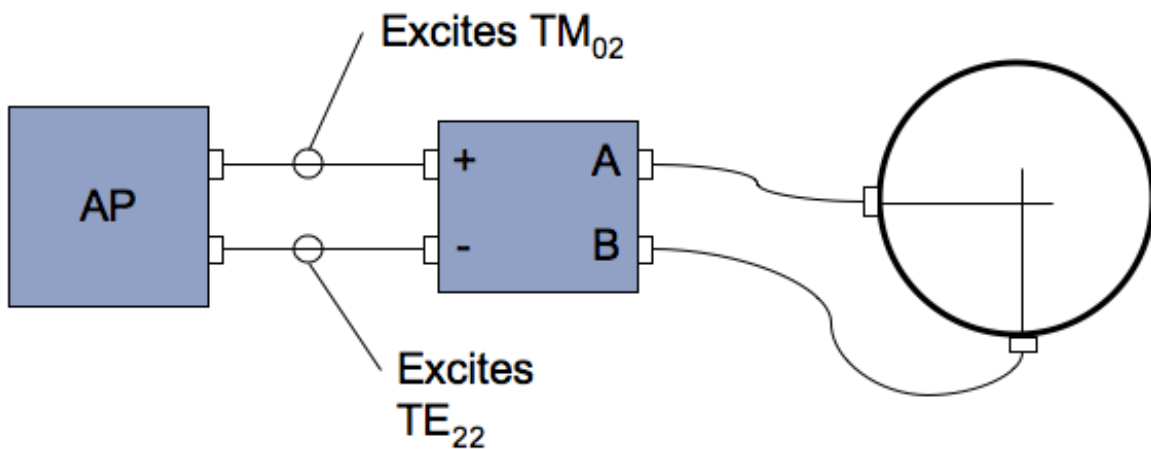


Figure 5.5: Diagram of how to use a 180 degree hybrid splitter to excite different single modes depending on which AP antenna port is used.

The major downside of this antenna is that shifting away, even slightly, from a maxima

in the T parameter implies that the antenna is no longer a good match to a 50 ohm impedance system. Thus, in order to use the antenna, the user must be willing to contend with a higher insertion loss in the system or impedance matching circuitry needs to be provided. There is a potential, though, to combine impedance matching circuitry with a 180 degree hybrid power splitter to maximize the efficiency and minimize the hardware installed with the antennas.

It should also be mentioned that the antenna will not work the same way in different duct sizes. For example, we also considered a 0.209 meter monopole in an duct 0.305 meters (12 inches) in diameter. It turns out that this diameter is close enough to the 0.305 meter size that the mode content is similar to the mode content of the 0.315 diameter duct. Figure 5.6 shows the mode content of a 0.209 meter monopole in a 0.315 diameter duct. The mode content excited by this length monopole is still dominated by the TM_{02} and TE_{22} modes. However, as shown in the figure, the modes are no longer present equally, but rather the TE_{22} is slightly more prevalent. Driving the antenna elements in or out of phase will still excite primarily TM_{02} or TE_{22} respectively, because the array arrangement ensures either TE_{22} or TM_{02} is canceled out. Despite this, the mode mix when TM_{02} is selected will be less dominated by TM_{02} because of the lower excitation of the mode by each array element individually. This should be taken into account when using the long monopole antenna arrays in 0.315 meter diameter ducts. As a result, it may be preferable to use TE_{22} to achieve a mode mix more dominated by a single mode.

5.2.2 Mode Attenuation Analysis

As we saw in Section 5.2.1, the proposed antenna array can be used to excite primarily modes TE_{22} or TM_{02} . To help choose the better mode, we wish to consider the attenuation with distance each mode experiences. The attenuation of a particular mode with distance in a vacuum filled waveguide is given by Equation 5.1 for TE modes and Equation 5.2 for TM modes [29] as

$$\alpha_{TE_{mn}} = \frac{R_{surf}}{r\eta_0} \frac{\left(\frac{f_{c_{mn}}}{f}\right)^2 + \frac{m^2}{(\chi'_{mn})^2 - m^2}}{\sqrt{1 - \left(\frac{f_{c_{mn}}}{f}\right)^2}} \quad (5.1)$$

and

$$\alpha_{TM_{mn}} = \frac{R_{surf}}{r\eta_0} \frac{1}{\sqrt{1 - \left(\frac{f_{c_{mn}}}{f}\right)^2}}. \quad (5.2)$$

As can be seen from the equations, the mode attenuation for modes with indices m and n is directly proportional to the waveguide surface resistance, R_{surf} , and inversely proportional to the waveguide radius, r . Additionally, the mode attenuation is related to the waveguide cutoff frequency, $f_{c_{mn}}$, the frequency of the propagating wave, f , and for TE modes, the χ'_{mn} value which is the location of the n^{th} zero in the derivative of the m^{th}




Figure 5.6: Mode content excited by a 0.209 meter long monopole in a 0.305 meter diameter cylindrical duct as calculated using software by Pavel Nikitin [26].

order Bessel function of the first kind, J'_{mn} . The mode attenuation constant for TM modes is instead related to the value of χ_{mn} , which is the location of the n^{th} zero of the m^{th} order Bessel function of the first kind, J_{mn} . In the equations, η_0 is the characteristic impedance of free space.

Figure 5.7 shows the computed attenuation curves for a waveguide diameter of 315 millimeters, surface resistance 10^{-5} ohms. It is interesting to note that the attenuation with distance for both curves are approximately equal. This is something of an advantage because it implies that there should be no major advantage to using either of the two single modes excited by the long monopole array.

5.2.3 Antenna Array Testing

To validate the concept of the antenna array for exciting single modes of propagation, we performed measurements on an experimental long monopole array. It should be noted that these measurements were conducted primarily by Jessica Hess and are presented

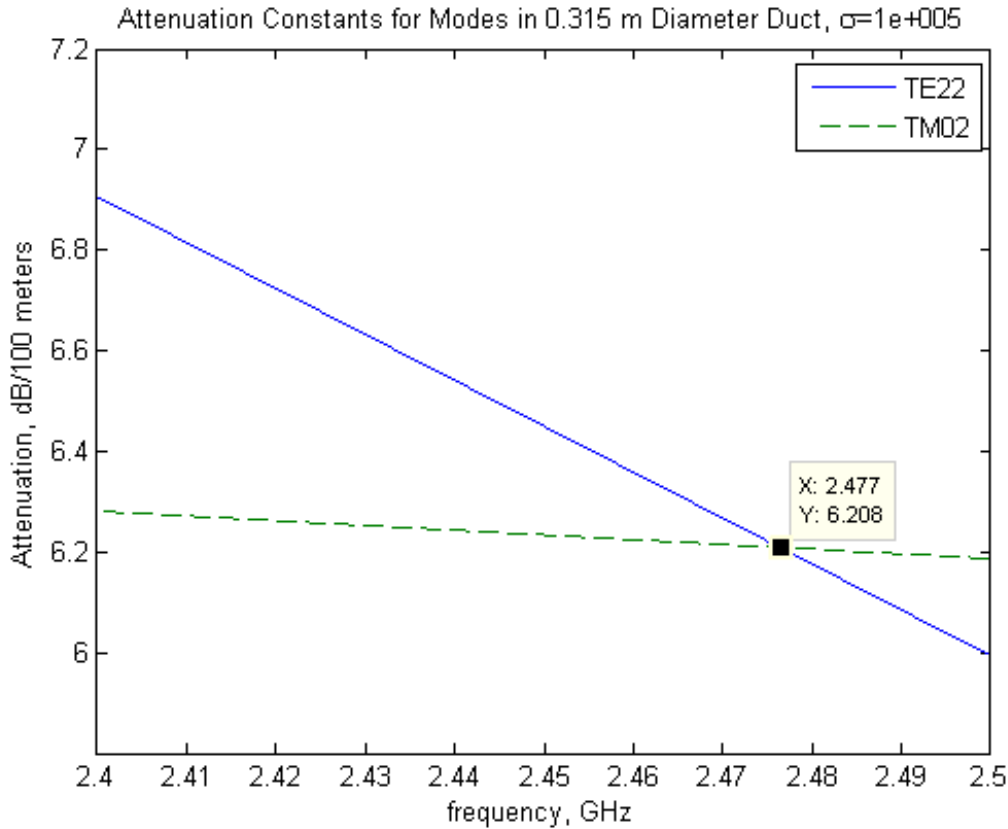


Figure 5.7: Attenuation curves for the two modes excited by the long monopole antenna array.

here as they validate the theoretical work performed for this thesis. The measurements were conducted in a straight run of cylindrical, 0.3048 meter diameter duct. The duct used was 3 meters long and the ends were left open. The antenna on one end was always a quarter-wavelength monopole. The other antenna was either another quarter-wavelength monopole or a 2 element, long monopole array with the elements driven in phase. The duct channel response for these two different antennas was recorded using a vector network analyzer at 1601 points between 2.4 and 2.5 GHz. The responses are plotted in Figure 5.8.

As seen in Figure 5.8, the peak-to-peak variation of the channel has been vastly reduced through the use of the long monopole array. In the ISM band, between 2.4 and 2.4835 GHz, the peak-to-peak variation in the frequency response is about 8 dB. While this does not directly tell us the mode content present, it does indicate that the response is dominated by a single propagating mode and that the array is functioning as expected.




Figure 5.8: Measured frequency responses comparing a quarter-wavelength monopole in the top plot and a long monopole array in the bottom plot. Note that while the data was measured from 2.4 to 2.5 GHz, the ISM band only extends from 2.4 to 2.4835 GHz.

5.3 Antenna Array Combination Diversity

In Section 5.2, we described a particular antenna array that leverages known mode behavior of a certain size ventilation duct to excite a single mode. In this section we consider more general antenna arrays used for diversity-like applications. Since the ventilation duct channel is static in nature, antenna arrays with fixed amplitude and/or phase shifts of elements could be installed reasonably inexpensively after being calibrated at installation time. However, we will consider the use of phased arrays for improving and worsening a particular communications channel.

To explore the feasibility of using antenna arrays for improving and worsening a channel, we used a digitally controlled Herley Farmingdale 7122 vector modulator which could provide magnitude and phase shifts between two antenna elements. The modulator was interfaced to a computer parallel port via a custom circuit board and custom MATLAB® software. The modulator is bi-directional and multiplies the RF signals passing through the modulator in either direction by a complex exponential, $Ae^{j\theta}$. The modulators are calibrated to provide amplitude coefficients, A , ranging from 0 to -70 dB and phase shifts, θ , of 0 to 360 degrees with an accuracy of ± 0.2 dB and ± 5 degrees over the magnitude and phase shifts and frequency ranges used in the experiment.

For the measurements presented below, we considered only a phase shift between two antennas. To use a single vector modulator to accomplish this, we created the setup shown in Figure 5.9. The setup used carefully chosen cables to ensure the electrical length of both paths are as close to identical as possible. Measurements with a vector network analyzer indicated the difference between the two paths was 28 picoseconds or about 1 degree across the 2.4 to 2.5 GHz bandwidth.




Figure 5.9: Power splitter and vector modulator setup for exciting two antennas with arbitrary phase offsets.

The measurement setup we considered for an antenna array is shown in Figure 5.10. The setup shows a U setup of 0.305 meter diameter cylindrical ducts. The setup has 5 lengths of 3.048 meter long straight ducts and 2 smoothly curving right angled ducts. The end of the ducts were covered with solid metal endcaps. A closed system such as this is representative of the most hostile environments found in duct networks. Quarter-

wavelength monopole antennas were used as shown. Lastly, the splitter setup from Figure 5.9 was used with a Vector Network Analyzer to characterize the ventilation duct setup.

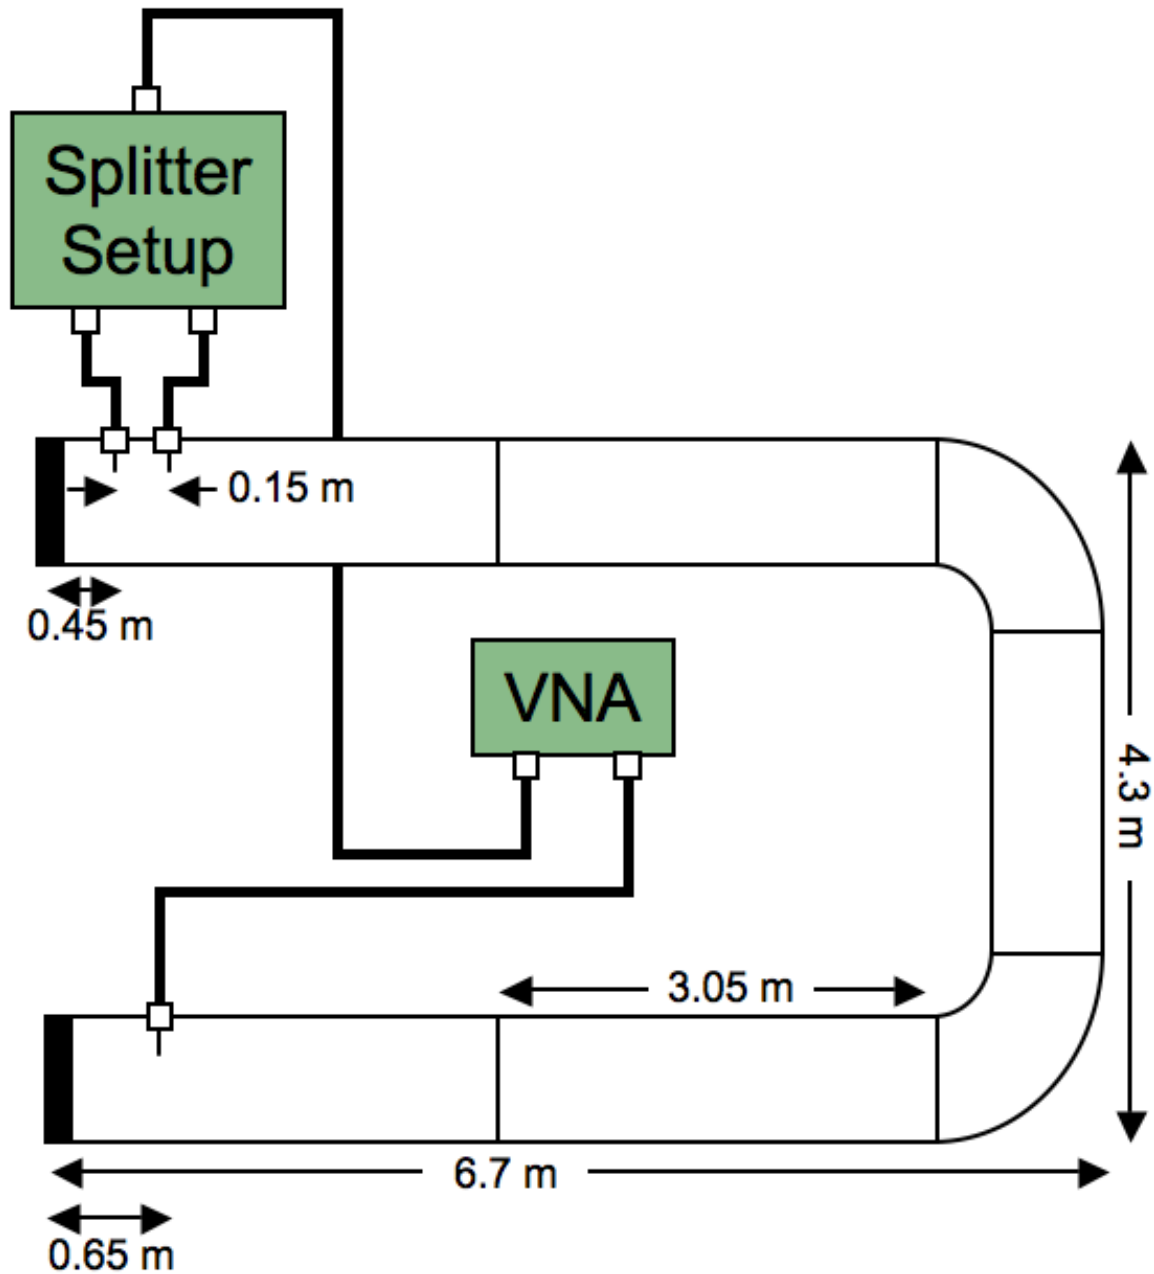


Figure 5.10: Measurement duct setup of phase shifting antenna array.

To characterize the setup in Figure 5.10, the splitter setup was used to create phase

shifts between -175 and 180 degrees in 5 degree increments between the two antennas. For all phase shifts, the attenuation selected for the vector modulator was selected to match the attenuation of the 10 dB attenuator shown in Figure 5.9. For each phase shift, the frequency response of the duct channel from 2.4 to 2.5 GHz was measured at 1601 evenly spaced points using the network analyzer. All 72 of these responses were plotted together as a contour plot as shown in Figure 5.11.




Figure 5.11: Contour plot of the frequency responses.

The peak-to-peak variation in signal strength across each of the frequency responses measured at each phase offset was calculated. For different phase offsets, the peak-to-peak variations ranged from 28 to 58 dB over the 2.4 to 2.5 GHz frequency range. This is a substantial variation in the quality of the channel response. It seems likely that by selecting the correct phase offset, we could improve wireless LAN performance through the duct channel. To check this hypothesis, we consider more closely the 22 MHz frequency range of IEEE 802.11b channel 1. These responses are shown in Figure 5.12. The peak-to-peak variation ranged from 18 to 41 dB over the center 16 MHz of bandwidth

in the center of IEEE 802.11b channel 1. This range was selected because it corresponds well to the occupied bandwidth of a IEEE 802.11b signal.




Figure 5.12: Contour plot of frequency responses for different phase shifts between antennas in dBm.

As can be seen in Figure 5.12, there are some phase shifts that have flat frequency responses and some that have sharp nulls. To see how a wireless LAN would perform in these channels, we connected an IEEE 802.11b AP to the splitter setup and the antenna array instead of the VNA shown in Figure 5.10. An Orinoco, IEEE 802.11b card was then connected via an adapter cable to the single antenna in the same setup. In addition, some attenuators were added between the AP and the splitter setup to allow the wireless NIC to operate in a typical range of received signal strengths. To ensure the wireless cards operated in a reasonable power range, 30 dB of attenuation was added at the AP and an additional 30 dB of attenuation was added at the wireless LAN card. The splitter setup was then used to provide equal magnitude phase shifts between the two antennas. The phases used were in 45 degree steps between -135 and 180 degrees. For each phase

Table 5.1: IEEE 802.11b signal strength and transmission rate measurements for antenna array phase shift measurements of U-shaped ventilation duct setup. Each entry in the table is based on a single measurement.

Phase Offset	Signal Strength	Transmit Rate	Peak-to-Peak Variation
180	-71 dBm	11 Mbps	20.1 dB
135	-72 dBm	11 Mbps	19.9 dB
90	-88 dBm	5.5 Mbps	23.2 dB
45	-78 dBm	5.5 Mbps	31.6 dB
0	No connection		27.8 dB
-45	-76 dBm	5.5 Mbps	20.8 dB
-90	-73 dBm	11 Mbps	18.6 dB
-135	-71 dBm	11 Mbps	23.5 dB

offset, a wireless LAN card utility was used to measure the received signal strength and operating transmission rate of the wireless LAN. These measurements are reported in Table 5.1. In addition, we also applied a maximum attenuation to the vector modulator to check the channel response of a single antenna setup. In this case, no communication between the AP and the wireless LAN card was possible.

Figure 5.12 seems to indicate that the phase shifts with the flattest responses and thus the lowest peak-to-peak variation, such as the response at 180 degrees of phase shift, have the best transmit rates. The peak-to-peak values shown Table 5.1 do not correlate with the transmit rate as well as might be expected. We have thus plotted the peak-to-peak variation of the 2.4 to 2.5 GHz frequency response along with the IEEE 802.11b transmit rate in Figure 5.13.

As shown in Figure 5.13, the peak-to-peak value is a decidedly poor predictor of the transmit rate. For this reason we have also plotted the peak-to-peak variation computed for just the center 16 MHz of the IEEE 802.11b channel 1 frequency range. While this plot indicates a somewhat better correlation between the peak-to-peak variation of the channel. However, we also have computed the mean gain of the frequency response over the same 16 MHz bandwidth for the different phase offsets as shown in Figure 5.15.

The correlation between channel gain and transmit rate is clear from Figure 5.15. This leads us to suspect that the ability to enable or disable communications between an AP and a user is more to do with the channel gain than the peak-to-peak variation in the channel. However, since a decrease in the peak-to-peak variation can lead to the increased channel gain, it is likely that the benefits of decreased gain variation and decrease channel loss both act to improve the channel for an IEEE 802.11b signal.

Lastly we present a hint of what's to come in the next chapter. When we check the phase response at the IEEE 802.11b channel 1 center frequency of the two measured frequency responses between each AP side antenna and the user side antenna, we find




Figure 5.13: Plot of the measured peak-to-peak variation in the frequency response across the 2.4 to 2.5 GHz frequency range using an antenna array with particular phase offsets and the resulting IEEE 802.11b transmit rate achieved.

a phase difference of 170.9 degrees. Interestingly, if we were to apply time-reversal coefficients, as we will discuss extensively in Chapter 6, we would use a phase offset equal to the conjugate phase of the measured response. This would result in applying a phase shift of -170.9 degrees. As shown in Figure 5.15, we would expect such a phase offset to perform well.

5.4 Conclusion

In this chapter we have presented some initial experimentation with antenna array applications for ventilation ducts. The results presented have illustrated two important uses of arrays. The first is that antenna arrays can improve the channel properties of a wireless link. This was demonstrated using the long monopole antenna array for selective mode excitation and by phasing the elements in a quarter-wavelength monopole array to improve the measured signal strength and throughput of an IEEE 802.11b link. Secondly,




Figure 5.14: Plot of the measured peak-to-peak variation in the frequency response across the 16 MHz range in the center of IEEE 802.11b channel 1. The peak-to-peak variation as well as the IEEE 802.11b transmit rate achieved are plotted versus the particular phase offset applied to the antenna array elements.

an antenna array can prevent connectivity on a particular link as we have also shown through the use of appropriate phasing of the quarter-wavelength monopole array. While the long monopole antenna array integrates nicely with the two element antenna diversity available in current IEEE 802.11b and 802.11g access points, the concepts presented in this chapter go beyond antenna diversity for simply improving link reliability. The usefulness of antenna arrays will be extended in Chapter 6 by more formally exploring the use of antenna arrays for use in improving the link and system capacity of ventilation duct based communications.




Figure 5.15: Plot of the measured mean channel gain of the frequency response across the 16 MHz range in the center of IEEE 802.11b channel 1. The channel gain as well as the IEEE 802.11b transmit rate achieved are plotted versus the particular phase offset applied to the antenna array elements.

6.1 Introduction

In Chapter 5 we demonstrated that antenna arrays can improve both the channel statistics and affect link connectivity in ventilation ducts. In this chapter we will extend the use of antenna arrays to improve the system capacity available in a ventilation duct based RF distribution system. While system capacity may be an area of difficulty for ventilation duct based systems, the highly reverberant nature of ducts means there is significant potential for capacity improvements [16]. The key reason for the concern over capacity, somewhat ironically, comes from the low loss nature of ventilation ducts. Since buildings are generally covered by a small number of duct networks, APs separated even significant distances from one another may only have a small pathloss. Since the loss with distance through ventilation ducts is so low, 0.34 dB per meter as measured in Section 3.6.2, the pathloss through the duct from AP to AP is relatively low. Further, since IEEE 802.11 transmissions use carrier sense multiple access (CSMA), if APs can detect other AP transmissions, they will defer their transmissions. However, since the loss to a user from a duct is significantly higher than the loss through the duct, both APs can likely transmit to different users simultaneously without issue.

Prior work with ventilation ducts has illustrated that multicarrier and OFDM (Orthogonal Frequency Division Multiplexing) transmissions could increase the capacity of ventilation duct channels [15, 18]. However, we already anticipate using IEEE 802.11g APs with ventilation ducts. Consequently, OFDM transmissions will be used in the ducts. We wish to investigate methods of further increasing the capacity of the duct networks. As such, the focus of this chapter is to overcome the possible loss in system capacity due to adjacent channel interference with different antenna array techniques. The first technique we will consider is time-reversal focusing for improving link capacities. Second we consider different antenna interference nulling techniques to allow APs to null out one another's signals. Lastly, we compare the theoretical capacity of duct based MIMO channels to ascertain the suitability of using MIMO communications devices such as IEEE 802.11n APs and to identify duct arrangements most suitable for increased capacity.

6.2 Demonstrating Time-Reversal Focusing at RF

Prior to the work that is presented in this section, no work had been published showing that time-reversal techniques could be applied to RF at microwave frequencies. It had been shown that if multiple transmit and receive antennas are used, multipath propagation can be exploited to realize multiple simultaneous information channels [47, 48, 49]. It had also been shown that multipath can be exploited for communications with acoustics and ultrasonics using time-reversal focusing [24, 25]. Using ultrasonic waves in a water tank, Derode, et al [25] focused short pulses simultaneously at five spatially distinct locations using multipath propagation introduced intentionally by a randomly-arranged forest of steel rods.

In this section we demonstrate that time-reversal focusing techniques can also be applied to electromagnetic waves at multi-gigahertz frequencies in indoor environments. In contrast to the acoustic measurements, our electromagnetic results involve signal bandwidths less than the coherence bandwidth and hence derive their focusing ability from multipath rather than bandwidth. This is an important demonstration to show before time-reversal focusing can be applied in ventilation ducts. As such, we have previously published some of the material in this section in [50, 51, 52].

6.2.1 Time-Reversal Techniques

A time-reversed signal is represented in the frequency domain by the phase-conjugate of the original signal. In addition to ultrasonics, phase-conjugate techniques have long been used in electromagnetics at both radio and optical frequencies. Phase-conjugate electromagnetic antenna arrays were developed in the early 1960's [53, 54, 55], and have been used extensively as retroreflective arrays. Van Atta [53] described a wideband technique in which symmetrically-placed antennas in an array are connected by transmission lines that all exhibit the same time delay. A technique using nonlinear mixing with a locally generated radio frequency signal was introduced by Pon [54]. Nonlinear refraction techniques were developed in the 1970's to realize phase-conjugate optical mirrors [56, 57].

The use of radio frequency retroreflecting arrays in multipath environments has also been explored prior to this work. Karode [58, 59] discussed the use of phase-conjugate arrays as multipath sensors, and Tuovinen, et al [60] showed that radio links using a phase-conjugate array resulted in significantly less fading caused by multipath propagation. Most relevantly, Lerosey et al [61] demonstrated a time-domain technique for performing time-reversal with electromagnetic waves. To the best of our knowledge, no one had demonstrated frequency domain time-reversal super-resolution focusing with electromagnetic waves in multipath environments, nor demonstrated the potential for parallel channels via this focusing effect prior to our publication of this work.

In this section we show that the focusing phenomenon can be realized with a bandwidth sufficient for significant RF communications to take place. This is an important

result because it allows two users that are extremely close in space to use the same frequency band to communicate to the same array with significant isolation. This work also clears the way for using time-reversal techniques in ventilation ducts.

The electromagnetic time-reversal principle states that replacing t with $-t$ in any valid solution to Maxwell's Equations results in a second valid solution, if the propagation environment is reciprocal. As an example, consider the fields from a source antenna that are completely characterized at a surface enclosing the source (it is sufficient to characterize the tangential components of the electric and magnetic fields). If sources are then placed at the surface and reproduce the characterized fields but with t replaced by $-t$, then a wave solution will be generated that propagates inward from the surface and focuses on the location of the original source antenna. This is the basic principle underlying time-reversal focusing.

In practice, a finite-size array antenna sparsely and incompletely measures the fields radiated by the source antenna. Consequently, the time-reversed signal retransmitted from the array will not precisely reproduce the time-reversed signal from the source antenna. However, the signals from the array elements that do arrive at the location of the original source antenna will be phased such that they interfere constructively, resulting in a localized region of increased field intensity. Since the fields from the array elements combine coherently at the target location and randomly elsewhere, the ratio of the amplitude at the focus to the surrounding values increases as \sqrt{N} where N is the number of antenna array elements.

In our work, the term “super-resolution” is used to refer to an improved spatial focusing of power from an antenna beyond what is predicted by the Rayleigh criterion (i.e. cross-range focusing) and the depth of focus (i.e. range focusing). The Rayleigh criterion states that the minimum resolution is limited by $\frac{0.6\lambda}{NA}$ where λ is the RF wavelength, $NA = \sin \theta$ is the numerical aperture, and θ is the half angle subtended by the array as viewed from the target antenna, as shown in Figure 6.1(a). Thus, it can be seen from Figure 6.1(b) that an increase in scatterers in the environment allows an increase in the effective numerical aperture of the antenna array. However, the minimum size of the focused spot is of order $\frac{\lambda}{2}$.

6.2.2 Experimental Time-Reversal Setup

Our measurements were conducted using a phase-conjugate, frequency-domain technique rather than a record and playback, time-domain technique. The experiments were conducted in a cluttered laboratory setting, as shown in Figure 6.2. A four element antenna array was used to focus power at a target antenna. The array consisted of four panel antennas with 9 dBi of gain across the 2.4 to 2.5 GHz frequency band. A broadband discone antenna with an approximate gain of 2 dBi was used as the target antenna. To ensure that no line-of-site path was present, a large metal plate was placed about 1 meter in front of the antenna array.




Figure 6.1: Graphic representation of numerical aperture of an antenna array, (a), and the enhancement of numerical aperture of an antenna array via salutary use of multipath, (b).




Figure 6.2: Schematic of laboratory in which experiments were performed. Note the significant amount of clutter and the metal plate obstructing the line-of-sight path. All furniture and the metal plate were removed in the empty lab case.

Power splitters and matched length cables were used to feed each array element with matched magnitude and phases, as shown in Figure 6.3. A Herley Farmingdale 7122 vector modulator was used with each element to provide magnitude and phase shifts. The modulator is bidirectional and multiplies the RF signals passing through the modulator in either direction by a complex exponential, $Ae^{j\theta}$. The modulators are calibrated to provide amplitude coefficients, A , ranging from 0 to -70 dB and phase shifts, θ , of 0 to 360 degrees with an accuracy of ± 0.2 dB and ± 5 degrees over the frequency and complex exponential ranges used in the experiment.

A vector network analyzer (VNA), shown in Figure 6.3, was used to measure the magnitude and phase of the frequency response between the antenna array and target antenna. To perform the time-reversal, the network analyzer was first configured to measure the reverse channel, from the target antenna to each of the array antenna elements, as follows. All modulators except one were set to the maximum attenuation value (>70 dB) and a single modulator was set with the minimum attenuation and zero degree phase shift. The VNA then measured the magnitude and phase response of the channel from the tar-

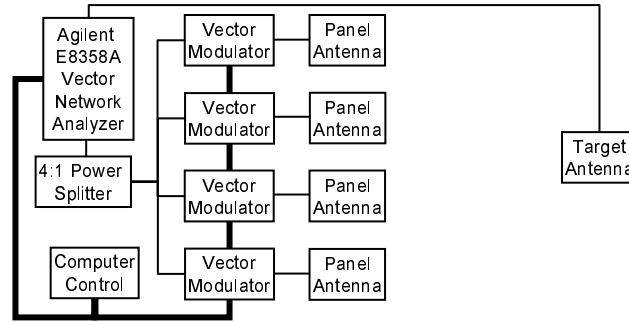


Figure 6.3: Setup of apparatus for experiment. All cables were matched lengths. The vector network analyzer can be switched internally to measure from the target antenna to the array or vice versa.

get antenna to the antenna array element at 2.45 GHz. This procedure was repeated for each array element. The measured magnitudes and phases were then normalized to the array element with the highest gain and the time-reversal (i.e. phase-conjugate) settings were computed. These settings were applied to each of the modulators to complete the time-reversal. All signals sent through the modulators with these settings were thus transmitted with a phase conjugate to that of the channel.

After the complex coefficients are applied to the vector modulators, the VNA was used to measure the channel response by recording the frequency response in the forward direction, from the antenna array to the target antenna (although, by reciprocity, the channel is independent of whether the target is the transmit or receive antenna). The amplitude received by the target antenna was measured at 81 points arranged in a square grid with 5 cm spacings and aligned manually via plumb bob with an estimated precision of ± 0.5 cm.

6.2.3 Experimental Time-Reversal Results

A contour plot generated from the measurements is shown in Figure 6.4. The figure clearly shows the spatial focusing that results from the time-reversal antenna array. From the figure, the cross-range null-null width of the focused spot is about 12 cm. For comparison, the Rayleigh criterion for one-dimensional focusing reduces to $\sim 1.5 \frac{L\lambda}{D}$. With $\lambda = 12.2$ cm, $L = 6.7$ m, and $D = 0.73$ m, the null-null spot size is 1.73 m, or over 10 times larger than the measured spot.

The result is even more striking in the range focusing direction. Figure 6.4 shows the spot size in the range direction to be about 25 cm, or roughly two wavelengths. The classical depth of focus for the array with line-of-sight and no multipath is $\sim 2\lambda(L/D)^2 = 20$ m, or roughly 80 times larger than the measured spot.

To better compare the measured focusing, the above measurements were repeated in

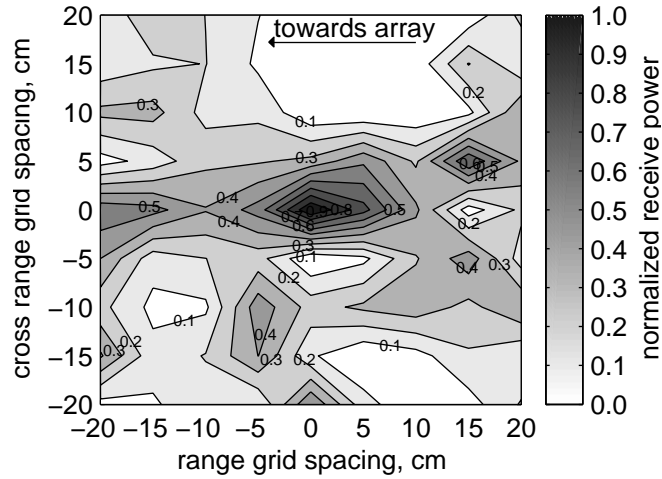


Figure 6.4: Normalized, received power as measured in the area around the target location. Measurements were made manually every 5 cm at 81 grid points.

the same room, but with all of the furniture and the metal plate removed. In addition, flat surfaces behind the target antenna were covered with RF absorbing material. This enabled a practical measurement of a setup in which the direct, free-space path significantly dominated the received signal, with comparatively minor multipath contributions. In this empty lab case, a cross-range measurement was made again in 5 cm steps, covering a range of 1 meter to either side of the target location. The results are shown in Figure 6.5 with the measurements from the cluttered case, as extracted from the data shown in Figure 6.4. Also shown in Figure 6.5 is the calculated focusing that is expected in a perfectly free space scenario. By comparing this to the empty lab scenario, we see reasonable agreement with the size of the observed focused spot, considering some multipath scattering will be present in the nominally empty room. The deviation from a smooth curve is due to the remaining multipath that could not be suppressed without significant additional RF absorbing materials. We note that the multipath has a small bias in one direction so the focusing appears slightly offset from the center point.

To quantify the potential bandwidth of spatial focusing, measurements of the received power were made in a range of frequencies about the focus frequency and at several locations in the cross-range direction about the target antenna location. The results of these measurements are shown in Figure 6.6. As can be seen from the figure, a definite focusing of power is present at the target location and frequency. Also, the spatial half power bandwidth is consistently about 5 cm for a band of frequencies about the focus frequency. This indicates the multipath focusing phenomenon is valid for a range of frequencies, not just the focus frequency, despite the constant phase shift with frequency provided by the vector modulators. It is also interesting to note that the best focusing




Figure 6.5: Comparison of the cross-range focusing in the empty and cluttered lab measurements. In addition, the calculated focusing in an ideal multipath free setup is shown. The full widths at half maximum for the curves are simulated: 77cm, empty: 90 cm, cluttered: 6 cm.

occurs at a frequency other than the focus frequency.




Figure 6.6: Focusing effect in the cross-range direction by multipath focusing for band of frequencies about focus frequency used to compute conjugate-phase coefficients.

To measure the isolation the focusing seen in Figure 6.6 could provide to two nearby antennas, phase-conjugate coefficients were computed using measurements for two tar-

get locations separated by 5 cm in the cross-range direction. Four frequency responses were then measured using the target 1 and target 2 phase-conjugate coefficients with the target antenna at locations 1 and 2. These four cases allow the comparison of the isolation that would be provided if 8 vector modulators were used to simultaneously create different communications channels to the two target locations. The channel response measurements are shown in Figure 6.7. From the figure, it is clear that 10 dB or more of isolation is present between the two target antenna locations over a bandwidth of about 15 MHz centered at the calibration frequency. This appears to be consistent with Figure 6.6, which showed a focusing bandwidth of similar size. This isolation could be improved by adding additional elements to the phase-conjugate antenna array.




Figure 6.7: Measured channel responses for different target locations. The legend indicates the measured target location and the intended target location.

Using the above techniques for focusing, two simultaneous communications channels were created. Four of the modulators in Figure 6.8 were configured to focus on one target antenna location, and the other four were configured to focus on a target antenna location 5 cm away in the cross-range direction. The signals from each set of modulators were combined pairwise to drive the 4 elements in the array. This configuration created two simultaneous communication channels at the same frequency. The isolation between the two channels was found by comparing the amplitudes of the channel 1 and 2 signals at each antenna location. Signal to interference ratios (SIR) at the center frequency of focusing are reported in Table 6.1.

To further characterize the viability for communications provided by the phase-conjugate enabled channels, a pseudo-noise (PN) sequence, encoded via 2 tone Fre-

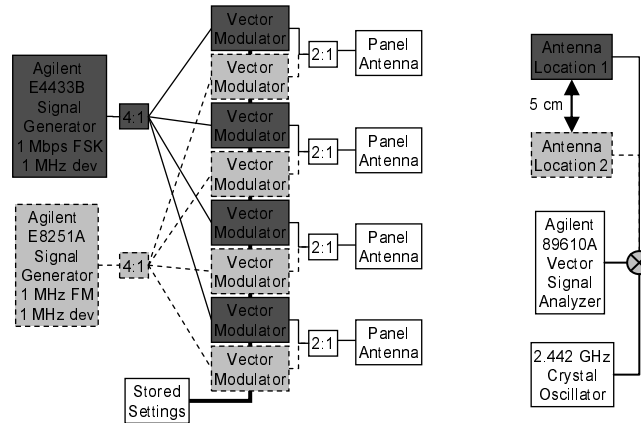


Figure 6.8: Apparatus setup for bit error rate measurements of two phase-conjugate channels to two target locations separated by approximately 5 centimeters.

quency Shift Keying (FSK-2) was repeatedly transmitted at a 1 Mbps rate. This test signal was transmitted using one phase-conjugate channel while an interfering 1 MHz, Frequency Modulated (FM) tone was transmitted at the same frequency on the second phase-conjugate channel. Both signals had a 1 MHz maximum frequency deviation and so occupied a nominal 2 MHz bandwidth about the center frequency used for the phase-conjugate focusing. The 50% correlation bandwidth in the cluttered environment was estimated to be approximately 14 MHz between 2.4 and 2.5 GHz.

A signal analyzer connected to the receive antenna at the focusing location of the intended channel demodulated and recorded the received bit stream. These bits were compared to the transmitted PN sequence to allow the bit error rate to be computed. This measurement was recorded for the first channel and then the FSK data signal was swapped with the FM signal to characterize the second phase-conjugate channel. This set of measurements was made in the empty lab and the cluttered lab scenario. The measured bit error rates (BER) are also shown in Table 6.1. The reported BERs are high compared to an expected BER in the presence of an additive white Gaussian noise (AWGN) signal of the same power due to the similarity of the interfering signal to the transmitted signal. It should be noted that in the baseline case of the interfering FM signal turned off, the bit error rate was less than 10^{-4} and in the case of both signals turned on but no phase-conjugate coefficients applied, the bit error rate was approximately 50%.

It is interesting to note that for the narrowband case discussed here, the time-reversal technique naturally yields the optimum coefficients for adaptive beamforming used with smart antennas [62]. Thus, in the presence of multipath, adaptive beamforming techniques should result in super-resolution focusing.

In conclusion, we have shown that time-reversal focusing techniques in a multipath

Table 6.1: Measured channel isolation and resulting bit error rate in two phase-conjugate channels.

Environment	Channel	SIR	BER
Empty	1	1.9 dB	4.5%
Empty	2	0.2 dB	~50%
Cluttered	1	15.5 dB	0.1%
Cluttered	2	11.2 dB	2.2%

rich environment can be used to achieve focused spots that are more than an order of magnitude smaller in the cross-range direction than would be possible in a line-of-sight configuration without multipath, and almost two orders of magnitude smaller in the range direction. Using the focusing technique, two communications channels were established using a 4 element phase-conjugate array. These channels had sufficient signal isolation in a multipath environment to allow minimal bit error rates for a 2 MHz wide signal. Only one channel was achievable in the minimum multipath environment. Thus, this technique is promising for RF communications with increased capacity. Further, it allows us to use the time-reversal techniques in ventilation duct channels.

6.3 Capacity Benefits from Time-Reversal Focusing

In Section 4.2 we measured a clear correlation between signal strength and the measured throughput available on an IEEE 802.11g link. This suggests that a basic way to improve the capacity of duct based wireless networks would be to simply decrease the pathloss to each user from an AP. Increasing the throughput of individual links thereby increases the system capacity. However, we also noted in Section 4.3 that some limit to the available link throughput may be present due to wireless cards being able to only tolerate some maximum delay spread. In Section 6.2 we demonstrated that time-reversal can be used to decrease the pathloss of a particular channel at RF, even in non-line-of-sight scenarios. In this section we will consider time-reversal for use in the ventilation duct channel and explore how it can be used to improve the capacity of individual wireless links.

6.3.1 Time-Reversal Focusing Simulation

To understand the impact of various time-reversal algorithms we first conduct a number of basic simulations using statistically generated random channels. We considered a situation shown in Figure 6.9. The situation is that an AP, AP_1 , is communicating with a user. The AP has N_1 antennas and the user has a single antenna. We assume that the AP is able to perfectly characterize the channel, H_u from AP_1 to the user. Then, to provide a better link to the user, AP_1 will apply time-reversal coefficients to its antennas to focus on the user location. We will also consider how the time-reversal coefficients affect the channel, H , from AP_1 to an eavesdropping AP_2 . Ideally, AP_1 should be able to improve

the channel to the user without adding any additional power to AP_2 .

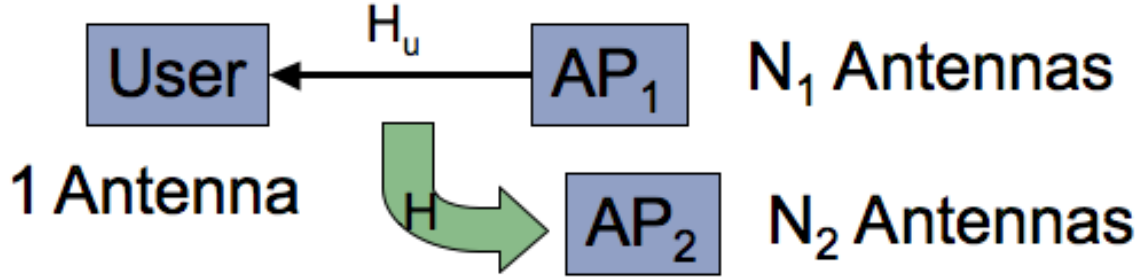


Figure 6.9: Time-reversal simulation scenario used for investigating a basic time-reversal algorithm.

For channels H_u and H we use the Rayleigh channel model. We generate independent, identically distributed (i.i.d.), complex, zero mean, unit variance entries. Both the real and imaginary parts of the entries are generated using Gaussian random variables. H_u is a 1 by N_1 matrix and H is a N_2 by N_1 matrix. We add a frequency dependence to H and H_u by generating 1601 instances so that we have $H_u(f)$ and $H(f)$. This is equivalent to a Rayleigh faded channel sampled with a sample frequency significantly smaller than the correlation bandwidth of the channel. While this is likely not the case for a 2.4 GHz to 2.5 GHz span in a duct channel, it does provide us with better statistical significance for averaging. For the following discussion we will leave off the frequency dependence of all matrices. This is simply for notational simplicity as the frequency dependence in no way affects the mathematics.

Throughout the remainder of this chapter we will use S_N to represent transmitted signals. In each case we use an equal magnitude, zero-phase transmission on all antennas as a comparison for time-reversal transmissions. Thus, S_N is an N by 1 column matrix of equal magnitude, zero-phase components. In this way, equal power will be transmitted on each antenna. We could easily loosen our restriction from zero-phase transmissions to equal phase transmission, but there is no significant change to the results and using zero-phase simplifies the notation. For the time-reversal transmissions we will compute a transmitted signal that reallocates an equal power column vector according to the particular calculation indicated by each time-reversal technique. For this simulation, after generating H_u and H , we compute the time-reversal coefficients using H_u . We first imagine a transmission from the user to the AP. The vector of received signals at the AP antenna is $R_{AP} = H_u^T S_1$. It is well known in signal processing that the time-reversal operator is the Fourier transform pair of phase-conjugation in the frequency domain [63]. Therefore, an ideal time-reversed signal that would be transmitted from the AP would be R_{AP}^* . If we normalize the transmission, the time-reversed signal received by the user from

AP_1 would be $R_{u,tr} = kH_u R_{AP}^*$, or

$$R_{u,tr} = kH_u H_u^h S_1, \quad (6.1)$$

where we have used the fact that S is purely real. $(\cdot)^h$ indicates the Hermitian operator, which is a conjugate transpose operation. We use k to ensure that the total AP transmitted power is the same for both the equal power and time-reversal cases. We also require that the total power in both cases is unity. In other words, we require that

$$\sum_{\kappa=1}^{N_1} |S_{N_1}|_{\kappa}^2 = \sum_{\kappa=1}^{N_1} k^2 |H_u^h S_1|_{\kappa}^2 = 1, \quad (6.2)$$

where $|\cdot|_{\kappa}^2$ indicates the magnitude squared of the κ^{th} element of the contained column vector. Thus, for the equal power, we have $S_{N_1} = [1, 1, \dots, 1]^T / \sqrt{N_1}$. We can also determine the normalization constant k as

$$k = \frac{1}{\sqrt{\sum_{\kappa=1}^{N_1} |H_u^h S_1|_{\kappa}^2}}. \quad (6.3)$$

We have computed R_u and $R_{u,tr}$ for a random instantiation of the matrix H_u using $N_1 = 4$. These two results are plotted in Figure 6.10. As can be seen in the Figure 6.10(a), R_u shows an expected average received power of unity over the frequency range. The received time-reversal signal in Figure 6.10(b) shows an improvement in the mean received power as indicated by the mean power gain of about 6 dB. Note that this gain comes only from time-reversal ensuring that the transmitted signals on all antennas add up in phase at the receive antenna location. Interestingly, time-reversal has also reduced the standard deviation and peak-to-peak variation in the received waveform.

As a side note, it is somewhat non-intuitive that the simulated time-reversal channels would show a gain above 0 dB, relative to the total channel input power. To understand this, we note that the N_1 AP antennas along with the single user antenna can be viewed as an $N_1 + 1$ port network. Such a network can be represented by a $(N_1 + 1)$ by $(N_1 + 1)$ scattering matrix in which the channel matrix H_u is contained in N_1 elements of one of the rows of the scattering matrix. Each element in H_u was generated using i.i.d., complex, zero mean, unit variance Gaussian random variables. As a result the mean value of the magnitude squared of an element of H_u is 1. Since the sum of the squares or the magnitudes of the elements in each row of a lossless network must sum to 1 [64], the network in the simulation, on average, could only be realized with active components. Clearly, wireless channels are passive in nature and have large attenuations. In such cases the 6 dB gain shown in Figure 6.10(b) would simply represent a reduced attenuation. Thus the active nature of the channel matrix in the simulation is merely an artifact of the unit gain statistical channel generated by the Gaussian statistics. The simulation still accurately

represents the gains realized by the proper phasing of the signals from the AP antennas via time-reversal.




Figure 6.10: Simulated signal received at the user location for a (a) normal transmission and a (b) time-reversal transmitted signal.

We have also computed the impact of using the time-reversal transmission at AP_1 on the received signal at an eavesdropping AP_2 . We first generate another Rayleigh faded channel matrix for the eavesdropping channel, H . We then calculate the normally received eavesdropped signal, $R_e = HS_{N_1}$, values and the eavesdropped signal received when time-reversal is used, $R_{e,tr} = HH_u^h S_1$. In this case, we set the number of antennas, N_2 , at AP_2 to be 4. Rather than plot each of the 4 received signals individually, we plot the sum of the signals in a single curve as shown in Figure 6.11. The normally received signal in Figure 6.11(a) shows an average value of about 6 dB over the total transmitted power, as we would expect. This gain over unity is due to the multiple receive antennas. We can conclude from Figure 6.11(b) that the time-reversal signal has not impacted the eavesdropped channel because there is no significant change to the mean power or standard deviation of $R_{e,tr}$ relative to R_e .

These simple simulations illustrate that time-reversal is a useful technique for improving the channel from an AP to a user without affecting channels to other users. For this reason we shall consider time-reversal further for use in ventilation ducts.

6.3.2 Duct Channel Measurements

To extend the simulation work presented above, we measured the frequency response of several test ventilation duct configurations. In this section we will describe the setup and characterization of these duct channels. We will then use the data for the comparison of different time-reversal focusing and nulling techniques in Sections 6.3.3 and 6.4.2




Figure 6.11: Computed channels heard by eavesdropping AP in the normal case (a) and with the time-reversal coefficients applied (b).

respectively.

The intention of the measurements is to allow us to identify how different duct layouts and antennas affect the efficacy of time-reversal algorithms. As a result, we started with a basic duct channel and then looked at several other configurations, each with a single aspect different from the base case. For our base case, we considered a 9.1 meter long, 0.3048 meter diameter, straight duct with metal endcaps, shown diagrammatically in Figure 6.12. At each end of the duct four holes, separated by a wavelength (12.24 cm), were drilled for quarter-wavelength monopoles. A vector network analyzer was used to characterize the 4 by 4 duct channel at 1601 points between 2.4 and 2.5 GHz. Since the network analyzer used was only capable of characterizing a two port system, the duct channel was measured by inserting monopole antennas in one hole at a time and recording each of the 16 frequency responses separately. This means that mutual coupling between antenna elements were not measured.




Figure 6.12: Ventilation duct setup for time-reversal measurement cases.

After characterizing the base case, we considered four modifications to the duct setup. The reasoning for this is to investigate how different parameters affect different time-

reversal algorithms. With this in mind, we kept everything the same, but changed one of four parameters of the measurement setup to investigate endcap reflections, mode scattering, antenna arrangement and mode excitation count. To explore the effect of endcap reflections we added RF absorber to the endcaps. Putting a duct with a 0.2 meter diameter hole (measured along the length of the duct) in the middle of the measurement case, as indicated by the dashed circle in Figure 6.12, indicates the effect of mode scattering. We also changed the arrangement of the quarter-wavelength monopoles from being aligned linearly along the length of the duct to being arranged circumferentially around the duct, each separated by 45 degrees. Lastly, we wished to consider the effect of reducing the number of modes excited by the antennas. To do this we have used the long monopole antenna element as a way to reduce the mode count. The long monopoles antennas were developed in Section 5.2.1 and excite primarily two modes as shown in Figure 5.6. The quarter-wavelength monopoles excite a significant number of modes as reported by [6] and is shown in Figure 5.2(a). These figures demonstrate that using the long monopole antennas will reduce the mode count excited in the duct channels. The holes in the duct for the long monopoles were drilled with two holes separated from each other by a full wavelength (12.24 cm) along the length of the duct and from the other two holes 90 degrees along the circumference of the duct. This orientation is how it is anticipated the long monopole antennas would be installed ordinarily.

For each of the four alternative configurations, the network analyzer was once again used to characterize the 4 by 4 MIMO (multiple-in multiple-out) channel by measuring in a pairwise manner the 16 different channels. The results of these measurements were five, 4 by 4, frequency dependent matrices, $\mathbf{H}(f)$.

In addition to the small scale duct measurements described, we wished to compare our results to a larger duct network. With this in mind, we used measurements presented in [15, 17] in a manner not previously published. We, however, re-describe the measurement details here for clarity.

A large duct network was used for two sets of measurements and is shown in Figure 6.13. The duct network was constructed of spiral cylindrical duct, 0.3048 meters in diameter. The duct structure was made up of 3.048 meter long straight sections, smooth right angle junctions and smooth t-junction pieces. All endcaps were left open and unterminated.

Two sets of measurements were taken using the large duct network. The first set, indicated by points A_1 and A_2 in Figure 6.13, was a 6 by 6 channel characterization. The measurements were conducted in the same manner as the 5 small scale duct measurements. A vector network analyzer was used to characterize all 36 channels in a pairwise manner using quarter-wavelength monopole antennas. In this case however, the holes used for the antennas were separated along the length of the duct by 5 centimeters instead of 12.24 centimeters. The second set of measurements is indicated by the points B_1




Figure 6.13: Duct network used for large scale MIMO measurements.

and B_2 in Figure 6.13. This measurement set was a 3 by 3 channel characterization, again using quarter-wave monopoles with hole locations spaced 5 centimeters apart along the length of the ducts.

6.3.3 Calculated Time-Reversal in Measured Channels

In this section we will focus on the channel impact available from time-reversal to various duct channels. As such, we will consider the scenario shown in Figure 6.14. As we did in Section 6.3.1, we consider the case where an AP with N_1 antennas applies time-reversal coefficients to focus on one antenna at a receiving user location but we will not consider the eavesdropping AP. We will use randomly generated Rayleigh faded channels for comparison but we will also select 1 by N_1 matrices from the measurement data taken in Section 6.3.2 for the H_u channels. We will also use the same time-reversal techniques

described by Equations 6.1 and 6.3. Rather than plot the frequency responses as was done in Section 6.3.1, we will instead compare the mean and standard deviation statistics of the normally received and time-reversal received signals.




Figure 6.14: Setup for measurement based time-reversal focusing simulations.

Figure 6.15 shows the mean statistics using the simulated channels and the measurements of the large duct networks. As can be seen, the gains of the normal channels are unitary by design, while the time-reversal technique provides increased channel gain as the number of antennas is increased. The large duct network results shown in Figure 6.15 also show a comparable gain from time-reversal for the 1 by 3 and 1 by 6 large duct networks. That is, the difference in the average strength of the received time-reversal signal relative to the normally received signal has increased. Interestingly, the gain due to time-reversal is slightly higher for the measured channel than for the simulated channels.

We also computed the standard deviation of the computed received channels. These are shown in Figure 6.16. As can be seen in the figure, the normal channels of the generated Rayleigh channels all have about the same standard deviation. However, the time-reversal signal is reducing the channel standard deviation and reducing it increasingly as the number of antennas increases. The large network measured data shows the same result that the increased number of antennas further reduces the channel standard deviation. Further, the reduction in standard deviation for the time-reversal channels is about the same for the measured cases as the simulated channels with the same number of transmit antennas. The results shown in Figures 6.15 and 6.16 indicate that time-reversal focusing is an effective technique for use in ventilation ducts and that Rayleigh faded channels can be used to predict how much impact time-reversal will have on a duct channel.

After calculating the normal and time-reversal received signals for simulated and large duct network channels, we performed similar calculations for the five small duct channels described in Section 6.3.2. We selected 1 by 4 matrices from the measurements taken for our H matrices. Once again, we computed the received channel matrices, R , using the techniques used in Section 6.3.1 and Equations 6.1 and 6.3. The resulting received power levels relative to the total power transmitted across all antennas are shown in Figure 6.17. As shown in the figure, time-reversal does provide some gain in all cases. Further, it appears that there is a correlation between the number of modes excited in the duct and the gain in received power of time-reversal signals over normal. This is evident




Figure 6.15: Average received power relative to the total input power across all transmit antennas for normal and time-reversal channels. The simulation bars represent randomly generated Rayleigh faded data while the large network bars represent large duct measurements.

because the first two and fourth case all have short monopole antenna arrays exciting about the same number of modes. The third case has a hole in the duct and hence filters out some of the modes and has a reduced gain for time-reversal. The fifth case uses long monopole antennas to excite the fewest number of modes and has the least gain for time-reversal. While this difference is apparent, there is overall little difference in the efficiency in time-reversal over any of the 5 cases.

We have plotted the standard deviations of the received signals for the 5 measured duct cases in Figure 6.18. This figure shows that again time-reversal has decreased the standard deviation of the received signal. It is not clear why some channels show more decrease in standard deviation than others. However, it does appear that short axial monopoles out perform both short circumferential monopoles and long monopoles.




Figure 6.16: Standard deviation of the calculated received power using normal or time-reversal signals. The simulation bars represent randomly generated Rayleigh faded data while the large network bars represent large duct measurements.

6.4 Capacity Benefits from Interference Nulling

While time-reversal focusing is a useful technique for improving the gain of a particular link, additional capacity gains are possible using nulling techniques. We have built our work primarily on time-reversal nulling techniques including recent work on a time-reversal based clutter nulling [52, 65, 66, 67]. Much work in interference nulling has focused on creating nulls in antenna array patterns [68]. Since the ventilation duct channel is inherently non-line-of-sight, conventional direction of arrival estimation techniques (e.g. [69]) are not applicable. However, some work on blind null-steering is relevant to these results [70], though to our knowledge no one has considered interference nulling in ventilation ducts.

Our reasoning for using interference nulling for improving the capacity of a duct based RF distribution system follows from the nature of the duct propagation environment. In




Figure 6.17: Average received power relative to the total input power across all transmit antennas for normal and time-reversal channels for 5 small duct measurements.

Chapter 2 we describe a system layout technique in which APs are connected to antennas inside ventilation ducts to provide coverage indoors. We have provided a propagation model in Chapter 3 which describes how signal strength varies inside ventilation ducts and throughout the building environment. This simulation model described a low, linear loss with distance model of the propagation inside ducts and an effectively large coupling loss from inside the duct to the building environment. For typical building dimensions these propagation mechanisms imply a far lower loss from AP to duct nearest a user location than the loss from that duct location to the user, even for users close to the duct (e.g. 1 meter away). As a result, two APs located to provide adjacent coverage areas that happen to be connected to the same duct will likely have a low pathloss between the two APs. Since IEEE 802.11 uses carrier sense multiple access (CSMA) for controlling access to the shared channel [71], APs that can detect packets sent by one another will defer transmission. This deferral reduces the capacity available to each AP. Ordinarily, system design would use the different frequency channels (e.g. channels 1, 6 and 11 in the United States of America) to reduce this co-channel reduction in capacity. However, the low loss with distance nature of ventilation ducts means that even APs far apart are likely to defer to one another. To combat this problem, we could use interference nulling antenna arrays to prevent APs from detecting one another's packets, thereby increasing




Figure 6.18: Standard deviation values of the calculated received power using normal or time-reversal signals for 5 small duct measurements.

the system capacity.

6.4.1 Interference Nulling Techniques and Simulations

To compare different nulling algorithms we first consider a simulation scenario shown diagrammatically in Figure 6.19. This scenario allows us to use various nulling techniques so that AP_2 does not detect packets from AP_1 . The techniques we will use all work on the receive side at AP_2 . We also want to ensure that the nulling algorithms do not affect the channel from AP_2 to a particular user, on average.




Figure 6.19: Simulation scenario for comparing different nulling techniques.

Here we describe the different nulling techniques we will consider. In Equation 6.4,

we describe the transmission from AP_1 to AP_2 :

$$R = \sum HS_{N_1}. \quad (6.4)$$

S_{N_1} is the transmitted signal which is always an N_1 sized column vector of equal amplitude, equal phase signals with a total power of unity. H is the matrix representing the N_2 by N_1 MIMO channel from AP_1 to AP_2 . R then represents the signal normally received by AP_2 summed over all of the receive antennas. We sum over the receive antennas because we are considering antenna array nulling techniques. Each nulling technique applies a complex coefficient for each antenna and then combines the resulting signals. We thus simply add the signals over the receive antennas in the normal case to have a basis for comparing to the nulling techniques.

Time-Reversal Nulling

Time-reversal nulling is a technique that was first described in [65]. This technique creates a time-reversal signal at each receive antenna and then inverts the sign on alternate antennas prior to summing the time-reversal received signals. This process is described in Equation 6.5:

$$R_n = \sum k_n PQHS_{N_1}. \quad (6.5)$$

H is once again the MIMO channel. The time-reversal signal, Q , is generated from the ordinarily received response, R . Q is a diagonal matrix with the conjugate of R as the elements of the diagonal of the matrix. That is, each element of Q is computed as

$$q_{\kappa,\kappa} = |r_{\kappa}| \exp(-j\angle r_{\kappa}) \quad (6.6)$$

where $q_{\kappa,\kappa}$ is the κ^{th} diagonal element of Q and r_{κ} is the κ^{th} element of R , which is the normally received signal given in Equation 6.4. Note that the negative sign inside the exponentiation is equivalent to conjugation. P is a diagonal matrix of alternating +1 and -1 values. This implements the alternating sign on alternating receive antennas. k_n is a power normalization factor. Since the nulling algorithms are all applied on the receive side, we normalize somewhat differently than was done in Section 6.3.1. We still want to ensure that the transmission power is fixed. So we normalize the antenna coefficients so that when power is transmitted from AP_2 using the same receive coefficients, the transmit power is still unity. For the time-reversal nulling technique in Equation 6.5, we compute the power normalization factor k_n as

$$k_n = \sqrt{\frac{N_1}{\left(\sum_{\kappa=1}^{N_2} \left| (PQ)^T S_{N_1} \right|_{\kappa}^2\right)}}. \quad (6.7)$$

Note that S_{N_1} is a column vector used for transmitting a signal from AP_2 with time-

reversal nulling coefficients. S_{N_1} also has a total power of unity but is a column vector of size N_1 . In Equation 6.7, the κ subscript indicates the κ^{th} element of the column vector enclosed. Note that we use $(\cdot)^T$, indicating transpose, because we are normalizing for transmissions from AP_2 .

Inverse Nulling

Time-reversal has the characteristic that it causes the received signals to be in-phase with one another. This is what allows the idea of inverting the sign of the signals on alternate receive antennas as implemented as time-reversal nulling to be effective. However, time-reversal does not cause signals to be of equal amplitude. Thus, while the time-reversal nulling technique described above should cause the signals to add destructively, perfect nulling will not be achieved, since the signals are of different amplitude. Ahmet Cepni has suggested a modification to the time-reversal nulling technique that is more effective. His idea is to both change the amplitude and phase of the time-reversal received signals. By making all of the received signals equal phase and by making alternating antennas opposite in sign, perfect nulling should be achieved. This is described in our simulations by first computing an equalization matrix E which is a pseudo-time-reversal matrix in that it uses the conjugate-phase, but modifies the magnitude of the signal. That is, E is a diagonal matrix with each element computed as

$$e_{\kappa,\kappa} = \frac{1}{|r_\kappa|} \exp(-j\angle r_\kappa) \quad (6.8)$$

where $e_{\kappa,\kappa}$ is the κ^{th} diagonal element of E and r_κ is the κ^{th} element of R , which is the normally received signal given in Equation 6.4. Note that the negative sign inside the exponentiation is equivalent to conjugation. The inverse nulling technique also uses a matrix of alternating signs to cause the nulling to occur. This is computed in our simulations as

$$R_i = \sum k_i P E H S_{N_1} \quad (6.9)$$

where P is once again a diagonal matrix of alternating +1 and -1 values. Each element of E on the diagonal is computed using Equation 6.8. Off diagonal elements are zero. We again introduce a power normalization constant k_i which is calculated as

$$k_i = \sqrt{\frac{N_1}{\left(\sum_{\kappa=1}^{N_2} |(PE)^T S_{N_1}|_\kappa^2\right)}} \quad (6.10)$$

Triangle Nulling

We now introduce a new nulling technique that we term “triangle nulling.” The technique is based on a geometric argument. First we consider a case where a receiver has three

antennas. In the normal case, we consider the received signals in the frequency domain as complex exponentials at each frequency. Each complex exponential can be visualized as a vector in the complex plane as shown in Figure 6.20. The figure shows three normally received signals r_1 , r_2 and r_3 . If we were simply to combine the three signals equally, they would sum to the indicated green vector, in the bottom right corner. However, if we introduce relative phasing angles as shown in the top right to create new received vectors r'_1 , r'_2 , and r'_3 , then we may be able to get the three vectors to sum to zero as shown in the bottom right of the figure.




Figure 6.20: Diagram showing how triangle nulling works in the frequency domain.

The idea behind triangle nulling is that given the magnitudes of the three received signals, there is a unique set of angles that can be applied to the signals that will result in the formation of a triangle. When the three signals form a triangle, then we have perfectly nulled the signal because the signals sum to zero. Since geometrically, the lengths of the sides of a triangle uniquely determine that triangle, there is only one set

of angles that will create the nulling triangle. We can calculate the angles of the triangle using the law of cosines:

$$\theta_1 = \cos^{-1} \left(\frac{|r_2|^2 + |r_3|^2 - |r_1|^2}{2|r_2||r_3|} \right). \quad (6.11)$$

In Equation 6.11 we use the magnitudes of the received signals, r , to compute the desired angles of the nulling triangle, θ . Once we have computed the angles of the triangle, we can compute the two phase shift angles ϕ_1 and ϕ_2 that rotate the r_2 and r_3 vectors to form the desired triangle.

It is important to note that a triangle cannot always be formed by three arbitrary length vectors. If the length of any one side is longer than the sum of the lengths of the other two, then no triangle is possible. However, if the arccosine function in Equation 6.11 is defined carefully for arguments with absolute values greater than 1, then the algorithm will still function well. That is, when the amplitude of one signal is longer than the some of the amplitudes of the other two, then the best way to phase the signals is so that the two smaller signals are 180 degrees offset from the largest amplitude signal. It turns out that this is exactly what the triangle nulling algorithm reduces to if we define the arccosine function as equal to 0 for input arguments greater than 1 and equal to π for input arguments less than -1. We could include a test on the inputs to the arccosine function shown in Equation 6.11. However, when the arccosine function is defined based on natural logarithms, then the arccosine function is defined for inputs with absolute values greater than 1 but simply has a complex result. It turns out then we can simply take the real part of the arccosine result as shown in Equation 6.12:

$$\theta_1 = \text{Real} \left\{ \cos^{-1} \left(\frac{|r_2|^2 + |r_3|^2 - |r_1|^2}{2|r_2||r_3|} \right) \right\} \quad (6.12)$$

So far we have only defined the triangle nulling algorithm for three receive antennas. We now extend the algorithm to more than three antennas. We would like to do this in a way that maximizes the chances that a nulling triangle can be formed. This is a simple process of combining the signals with the smallest amplitude signals into a single signal with an amplitude equal to the sum of the individual signal amplitudes which can be done by just phase shifting one of the signals to have the same phase as the other signal. To choose which received signals to phase align together, we use a recursive technique. To reduce N received signals, r , to three signals for nulling, we first pick the two smallest signals, r_a and r_b . We compute the phase difference between r_a and r_b and phase shift antenna a by that amount so that the signals for antennas a and b will sum in phase. The amplitudes of r_a and r_b will then be added together and considered a single vector with the phase of r_b and the amplitude of the sums of the amplitudes of r_a and r_b . We are then left with $N - 1$ vectors. By repeating this process until we are left with only 3 vectors, we ensure that we have combined the N signals into three signals that are closest in

magnitude to one another.

The triangle nulling algorithm described above is implemented in the context of our simulations as

$$R_{\Delta} = \sum k_{\Delta} T H S_{N_1} \quad (6.13)$$

where T is an N_2 by N_2 diagonal matrix with the triangle nulling angles as the arguments of complex exponentials as elements of the diagonal. The nulling angles are calculated from the normally received signal vector R from Equation 6.4. We again compute the power normalization factor k_{Δ} for the triangle nulling algorithm using the same unity power criterion used previously:

$$k_{\Delta} = \sqrt{\frac{N_1}{\left(\sum_{\kappa=1}^{N_2} |T^T S_{N_2}|_{\kappa}^2\right)}}. \quad (6.14)$$

Time-Reversal Triangle Nulling

Since triangle nulling does not provide perfect nulling when one of the magnitudes of the received signals is larger than the sum of the amplitudes of the other signals, we consider two other techniques with triangle nulling. The first is a time-reversal based triangle nulling. It is implemented by first calculating the time-reversal focusing received channel:

$$R_t = H^h H S_{N_1} \quad (6.15)$$

We then compute the triangle nulling angles based on R_t and create the time-reversal triangle nulling channel, $R_{\Delta t}$:

$$R_{\Delta t} = \sum k_{\Delta t} T_t H^h H S_{N_1} \quad (6.16)$$

where T_t is the diagonal matrix of time-reversal triangle nulling coefficients and $k_{\Delta t}$ is the power normalization factor computed as:

$$k_{\Delta t} = \sqrt{\frac{N_1}{\left(\sum_{\kappa=1}^{N_2} |(T_t H^h)^T S_{N_1}|_{\kappa}^2\right)}} \quad (6.17)$$

which is the same power normalization technique as previously used.

Eigen-channel Triangle Nulling

The second alternate means of performing triangle nulling is to first create a diagonal channel via eigenvectors used as transmit and receive antenna coefficients. This channel is given as:

$$R_m = V H U S_{N_1} \quad (6.18)$$

where V is a matrix formed from combining the eigenvectors of H as column vectors into a single matrix. U is the matrix inverse of V . This construction yields the result that VHU is a diagonal matrix D . The elements on the diagonal of the matrix are the eigenvalues of the channel H .

After creating the diagonal channels, the triangle nulling coefficients are computed based on the corresponding eigen-channel gains as created by the eigenvector transmit coefficients, which are the same as the eigenvalues that correspond to the eigenvector coefficients. The simulation channel result is:

$$R_{\Delta m} = \sum k_{\Delta m} T_m V H k_U U S_{N_1} \quad (6.19)$$

where T_m is a diagonal matrix of the triangle null angles of the eigenvalues of H . k_U is a power normalization constant that ensures the total power transmitted from AP_1 is unity. k_m is another power normalization constant that ensures that AP_2 transmits a total power of unity when it uses the computed receive antenna coefficients for transmission. k_U is computed as:

$$k_U = \sqrt{\frac{N_2}{\left(\sum_{\kappa=1}^{N_1} |U S_{N_1}|_{\kappa}^2\right)}} \quad (6.20)$$

while $k_{\Delta m}$ is calculated as:

$$k_{\Delta m} = \sqrt{\frac{N_1}{\left(\sum_{\kappa=1}^{N_2} |(T_m V)^T S_{N_2}|_{\kappa}^2\right)}} \quad (6.21)$$

Simulation Results

To compare the plethora of nulling techniques described above, we have generated random H channel matrices using the Rayleigh channel model. As was done in Section 6.3.1, we generated 1601 matrices representing the H channel from AP_1 to AP_2 . Each matrix entry was generated using complex Gaussian distributed, i.i.d., zero mean, unit variance random variables. For the size of each matrix, we chose $N_1 = N_2 = 4$. We also generated separate AP_2 -to-user channel matrices H_u that are uncorrelated to the H matrices but with identical statistics. We generated 1601 instances of the H_u matrices, which are N_2 by 1 in size.

For each instantiation of the channel matrices, we computed the normally received signal, R , described by Equation 6.4. We also computed the received signals for the 5 different nulling techniques described by Equations 6.5, 6.9, 6.13, 6.16, and 6.19. Since we are interested in how well each method nulls, we considered the received power levels for each of the 6 simulations. We computed the cumulative distribution function of the received power over the 1601 matrices for these computed receive power levels. Using the

AP_1 to AP_2 transmissions we have the results given in Figure 6.21.




Figure 6.21: Cumulative distribution function of received signal levels for different nulling algorithms.

The normally received signal levels shown in Figure 6.21 by a blue line with circle markers indicate the baseline with which to compare the nulling achieved. As one might expect, the inverse nulling always provides perfect nulling in all simulated cases. This is reasonable because the equalization matrix used ensures that the received signals always sum to a perfect null. In the simulation cases, all of the triangle nulling techniques work well. The triangle nulling technique does not always ensure perfect nulling, yet all three techniques provide perfect nulling in over 90% of the cases. The triangle nulling algorithm appears to be slightly improved by time-reversal triangle nulling and even more improved by eigen-channel triangle nulling. Somewhat unexpectedly, the time-reversal nulling technique appears to have much less effect on the channel. We believe this is because the time-reversal nulling technique does not guarantee ideal nulling, merely the the received antenna signals will add out of phase.

We have also computed and plotted the cumulative distribution function for the computed received signals at the user location. These values are plotted in Figure 6.22. These calculations assume that the various nulling algorithms apply their receive side coefficients to the antenna array used by AP_2 and then a total power of unity is transmitted equally across this antenna array to another user. The resulting received signals shown in Figure 6.22 show, as they should, no significant difference in distribution. This indicates that the power normalization constants, k are calculated correctly.




Figure 6.22: Cumulative distribution function of received signal levels at user location when different nulling algorithms are applied to AP_2 antennas.

6.4.2 Calculated Time-Reversal Nulling in Measured Channels

In this section we continue to compare the different channel nulling techniques described in Section 6.4.1. In this section, however, we use duct channel measurements to calculate the nulling responses. To do this we use the duct channel measurements described in Section 6.3.2. We have used the 5 small duct measurements and the 2 large duct measurements in combination with 4 statistically generated Rayleigh channel matrices.

We consider four square matrix sizes: 2, 3, 4 and 6. These random H matrices will be used for comparison to the large duct measurements. The cumulative distribution functions of the received signals computed using the 6 nulling channels are shown in Figure 6.23. First we note a few important features. First, we see that time-reversal nulling and time-reversal triangle nulling perform about the same in the 2 antenna case shown in Figure 6.23(a). This is as expected because the time-reversal triangle nulling algorithm reduces to the time-reversal nulling algorithm for 2 antennas. Next, we notice that the concept of using +1 and -1 on alternating antennas performs poorly in cases where we have an odd number of antennas as shown in Figure 6.23(b). Alternating +1 and -1 values on antennas is used for time-reversal nulling and inverse nulling. Both of these algorithms perform worse in the 3 antenna case than in the other cases.

The plots in Figure 6.23 also show the progression of the effectiveness of the nulling algorithms. Other than the 3 antenna case, time-reversal nulling appears to do provide about the same nulling or perhaps slightly less nulling compared to the baseline case as the number of antennas increases. This is reasonable because time-reversal nulling merely ensures that the received signals sum with opposite signs on alternating antennas. It is conceivable that, depending on the statistical distribution of the magnitude of the received signals, the time-reversal nulling technique may do worse as the number of antennas increase. The triangle nulling techniques, however, show consistent improvement as the number of antennas increases. Lastly, inverse nulling, except for an odd number of antennas, functions ideally in all cases.

We now compare the gain of the nulling techniques versus the normally received channel. For each of the 1601 channel matrices we computed the difference in the gain from the normally received signal to the nulling based received signal and then calculated the average. Figure 6.24 shows the four statistically generated channels and the two large duct network measurements.

Figure 6.24 shows many of the same previous conclusions. As the number of antennas increases, the time-reversal nulling technique provides slightly less average nulling in the three simulation cases with an even number of antennas. The triangle nulling techniques however show increasingly effective nulling. We have not plotted the inverse nulling technique since it is completely effective at nulling, there is no difference across any of the cases. We also note that the figure shows that the measured large duct network channels have similar changes in received power as the statistically generated channels with the same numbers of antennas. However, it is not clear which triangle nulling technique is the most effective from the results. In the statistically generated cases, the eigen-channel triangle nulling is generally the most effective. However, triangle nulling is the most effective in the 3 by 3, large duct measurement case. It is also hard to make a conclusion based on the 6 by 6 measurement case because the triangle nulling techniques are all equally effective.

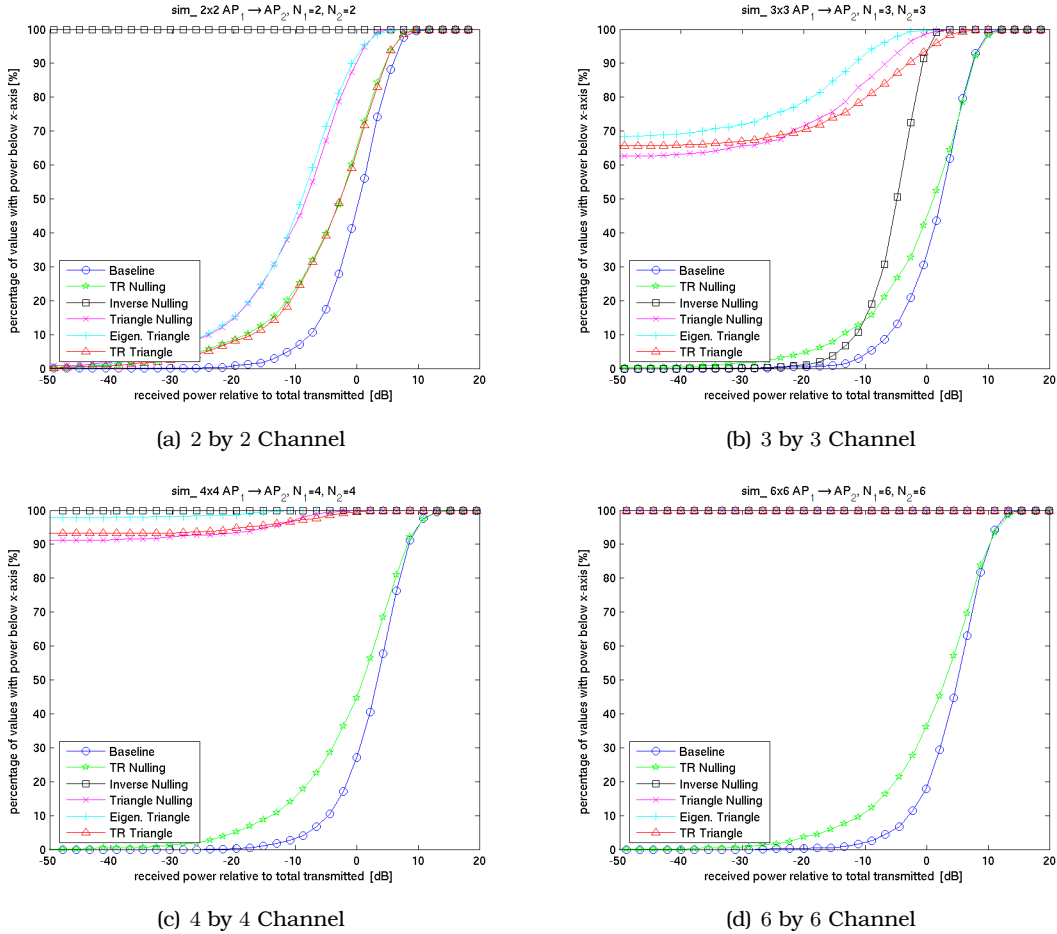


Figure 6.23: Cumulative distribution functions for different sized, statistically generated, Rayleigh channels.

The five small duct measurement cases are shown in Figure 6.25. Interestingly, there does not seem to be a stand-out nulling technique or measurement case that performs best. However, we do note that the long monopole antenna array scenario results in the almost the least nulling for all four of the nulling techniques plotted. Thus, long monopole antenna arrays are probably not a good choice for use with nulling techniques. We also notice that while the eigen-channel triangle nulling performs inconsistently across measurement cases, the triangle nulling and time-reversal triangle nulling techniques performed consistently well, while time-reversal nulling performed poorly or not at all. Triangle nulling and time-reversal triangle nulling would seem to be good choices as nulling algorithms. We might slightly recommend ordinary triangle nulling because it performed better in almost all of the simulation cases and in all of the large duct measurement cases




Figure 6.24: Comparison of the mean change in channel power compared to the normally received power in the calculated received signals for different nulling techniques. The first four items are statistically generated channels and the last two are measured duct channels.

shown in Figure 6.24. In either case, eigen-channel triangle nulling and time-reversal nulling are not the best techniques to select.

6.4.3 Single Frequency Based Time-Reversal Nulling

In the previous sections we have everywhere assumed that we can set the ideal coefficients at each frequency across the entire frequency range of interest. However, practical implementations of arbitrary magnitude and phase shifts over large frequency ranges are difficult to implement. A far simpler implementation is to simply select a single magnitude and phase shift based at a center frequency and apply the resulting complex coefficient to a range of frequencies. In this section we calculate how effective this technique can be in implementing the nulling algorithms we have explored previously. Since we are interested in IEEE 802.11b and 802.11g signals, we will consider a 20 MHz range of fre-




Figure 6.25: Comparison of the mean change in channel power compared to the normally received power in the calculated received signals for different nulling techniques.

quencies around a 2.45 GHz center frequency. We will once again analyze the small duct and large duct measurements described in Section 6.3.2 to investigate how well nulling at a single frequency can work in duct channels. We leave out the time-reversal nulling technique in these plots because we have found it to be ineffective as a nulling technique.

The first duct channel we will consider is the small duct measurement with RF absorbing end caps. The calculated nulling is shown in Figure 6.26. The normally received signal is plotted as a baseline of comparison for the nulling techniques. Also shown on the plot are the mean differences between each received nulling signal and the normally received signal, R , across the 20 MHz bandwidth plotted. For these numbers, positive values indicate better nulling. This first case shows some exciting potential for the single frequency nulling. The triangle nulling technique especially shows a difference in the average power gain of 17.5 dB across the 20 MHz span. Interestingly, the eigen-channel and time-reversal triangle nulling techniques as well as the inverse nulling technique do not appear to be effective over the 20 MHz range. Instead they seem to have a 2 to 4 MHz bandwidth of impact. The inverse nulling is moderately effective across the whole 20 MHz range. However, the sharp nulling effect in all four cases is mainly confined to the same 2 to 4 MHz wide bandwidth.

The case with metal endcaps instead of absorber is shown in Figure 6.27. The curves

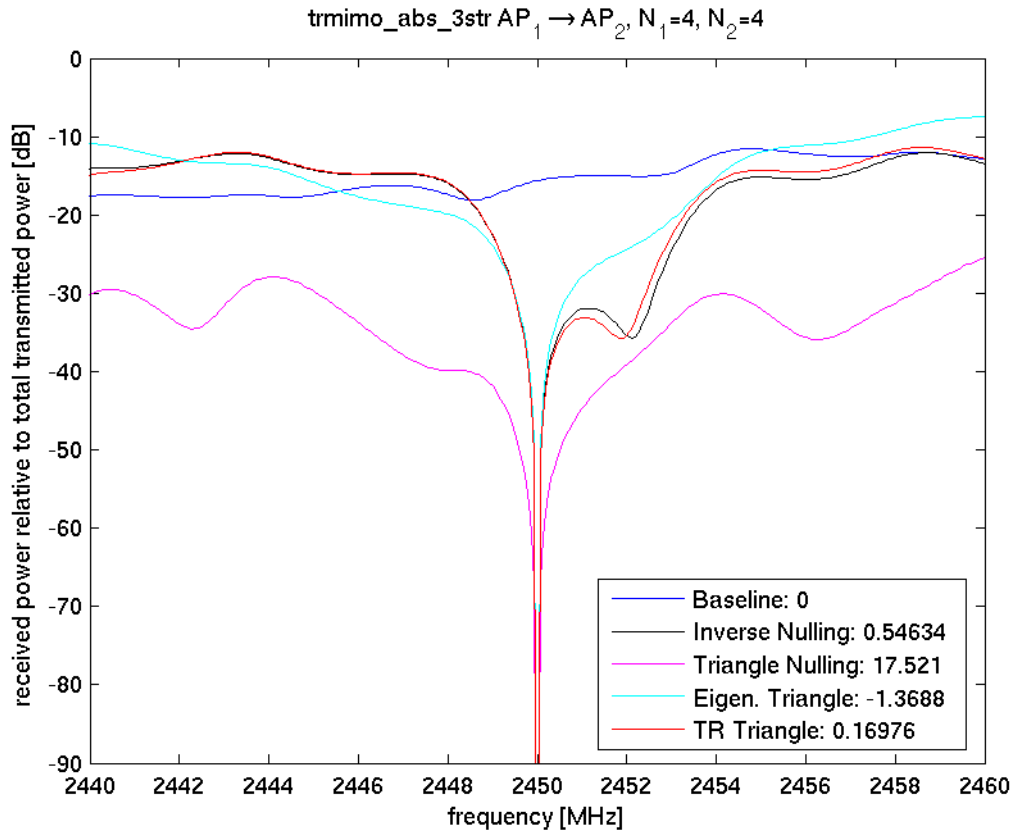


Figure 6.26: Calculated single frequency based nulling in the small duct measured channel with the axial array of short monopoles, absorber endcaps and no duct hole.

and differences in average gain figures seem to indicate that triangle nulling is again most effective while inverse nulling and time-reversal triangle nulling are slightly effective. However, the other techniques may not be as effective. In addition, the sharp nulls in the center of the range appear to be narrower in extent than they were for the absorber endcaps case. It is reasonable to expect that the bandwidth of nulling would decrease because the metal endcaps case has a narrower correlation bandwidth so we would expect the bandwidth of effective nulling to be smaller.

The remaining three small duct measurement cases show a more complicated picture. The calculated channels are shown in Figures 6.28, 6.29, and 6.30. As can be seen in these figures, there is generally a null achieved at the center frequency, but little impact on the overall 20 MHz channel. All of the nulling techniques appear to do fairly poorly at achieving significant nulling in these cases.

We have computed the effective nulling that could be expected in the large duct mea-




Figure 6.27: Calculated single frequency based nulling in the small duct measured channel with the axial array of short monopoles, metal endcaps and no duct hole.

surement cases. These results are shown in Figures 6.31 and 6.32. In these two figures we see a certain amount of nulling is achieved in all of the cases. Interestingly, the 3 by 3 case illustrates a potential difficulty with the triangle nulling algorithm. It appears that for the time-reversal triangle nulling algorithm, one of the three antennas is dominating the other two. As a result, perfect nulling cannot be achieved at the center frequency. This clearly impedes the time-reversal triangle nulling technique across the frequency range. This problem is far less likely to occur in the 6 by 6 case because antennas with small amplitudes are combined coherently to prevent one antenna from dominating the response. In fact, Figure 6.32 does show that the nulling techniques do achieve a decrease in the average received power. However, with a range of 3 to 9 dB improvement, the case for the single frequency nulling is not very compelling.




Figure 6.28: Calculated single frequency based nulling in the small duct measured channel with the axial array of short monopoles, metal endcaps and a duct hole.

6.4.4 Comparing Nulling Techniques

In the previous sections we have compared five different techniques for interference nulling. The intent of these techniques is to provide APs connected to the same ventilation ducts some degree of added isolation from one another without affecting communications to wireless users. We have found that the time-reversal nulling technique is the least effective of all the techniques because it is unable to null at multiple antenna locations simultaneously as explained above. Thus, while the technique is useful in clutter nulling as reported in [65, 66], it is not useful for the application we have in mind here.

The inverse nulling technique presented here is a modification of the time-reversal nulling algorithm. We found it to work well in all cases with an even number of antennas, though it was perhaps not the most effective in the single frequency nulling measurements. However, one issue with the inverse nulling that it needs to provide arbitrary gain to each antenna in order to provide perfect nulling. Actually implementing an RF




Figure 6.29: Calculated single frequency based nulling in the small duct measured channel with the circumferential array of short monopoles, metal endcaps and no duct hole.

system capable of this could be challenging. Further, the inverse nulling scheme is likely to increase the noise in the system because it is amplifying low SNR signals more than high SNR signals.

Triangle nulling provides something of a compromise compared to inverse nulling. Since basic triangle nulling only adjusts the phases of different received signals, it does not as significantly impact the noise of the system as inverse nulling does. The tradeoff is that triangle nulling does not always achieve perfect nulling. However, triangle nulling is significantly easier to implement in a ventilation duct system than inverse nulling. Since ventilation ducts are a static channel over time, it is conceivable that nulling coefficients could be applied to an antenna array once at installation and then left over the lifetime of the ventilation duct based RF distribution system. Installing the appropriate phase shifters in an antenna array would likely be less expensive than the equipment needed for an inverse nulling algorithm.




Figure 6.30: Calculated single frequency based nulling in the small duct measured channel with the array of long monopoles, metal endcaps and no duct hole.

The basic triangle nulling algorithm appears to be effective across many of the different measurement cases above. Despite considering two modifications to the triangle nulling algorithm, time-reversal and eigen-channels, neither appeared to have significantly better performance than basic triangle nulling. The measurement cases indicated that often triangle nulling was the best choice to use of the three. Further, both time-reversal and eigen-channel triangle nulling require added complexity. This added complexity is not justifiably beneficial based on the measurements and simulations presented above.

6.5 Theoretical Capacity Available via MIMO Arrays

One obvious potential route for improved system capacity is through space time coding of information on MIMO antenna arrays. At the time of writing, an exciting possibility for using these techniques is the IEEE 802.11n draft standard. The draft standard has been fraught with argument over how to implement MIMO techniques in the context of an IEEE 802.11 system. However, we have performed some basic research into how effective




Figure 6.31: Calculated single frequency based nulling in the 3 by 3 large duct measurements.

these systems can be inside ventilation ducts from an information theory perspective.

There has been significant research into extending information theory principles first developed by Shannon in [72] to MIMO communication channels. Of particular interest is the Shannon-Hartley theorem, [72, 73], given here as

$$C = \log_2(1 + \rho). \quad (6.22)$$

This theorem is used ubiquitously throughout communications. Here, C is the channel capacity in bits per second per hertz, and ρ is the signal to noise ratio. Several researchers have extended this work to multi-element antennas (MEA) and MIMO channels [47, 48, 49]. We will express the generalized capacity formula as

$$C = \log_2 \left(\det \left[I_{N_R} + \left(\frac{\rho}{N_T} H H^h \right) \right] \right) \quad (6.23)$$




Figure 6.32: Calculated single frequency based nulling in the 6 by 6 large duct measurements.

where N_T is the number of transmit antennas, N_R is the number of receive antennas, I_{N_R} is an identity matrix of size N_R , H is the N_R by N_T sized channel matrix and H^h is the complex conjugate transpose of the H matrix. This equation assumes that the transmitted signals are statistically independent, equal in power and with Gaussian distributions. We will use this equation to analyze the capacity available in the ventilation duct channels we measured in Section 6.3.2. This technique is somewhat similar to the technique used in [74, 75, 76] in which indoor channels are characterized for their MIMO capacity. It is the same technique we used in [17] in which we published the large duct data presented in this thesis. We have, however, extended our work to include the small duct channel measurements. We should also mention that our work is similar to [77] in which the authors present a brief analysis of the MIMO capacity of rectangular waveguides that support 1 or 2 modes of propagation.

To statistically characterize the MIMO capacity of each duct channel, we consider

each of the 1601 frequency bins as a separate channel matrix. At each frequency bin, we construct a matrix G using a measured value from each of the measured frequency responses. In this way we will be able to compute 1601 channel capacities. Below we plot the complementary cumulative distribution function (CCDF) of these 1601 channel capacities to get an idea of what percentage of the time we can expect the channel capacity to be above a certain value.

To analyze the capacity of the MIMO communications channels we need to normalize the channel matrices so that the signal to noise ratio is separated from the matrix channel characteristics. To do this we use the Frobenius normalization of each measured channel matrix. Thus, the H matrix used in Equation 6.23 is the Frobenius normalization of the constructed channel matrix G . We also use a signal to noise ratio, ρ of 10 dB in our calculations. Note that in the small duct measurements $N_T = N_R = 4$. In one of the large duct measurements $N_T = N_R = 3$ and in the other $N_T = N_R = 6$. For comparison purposes we have also generated the capacity CCDF curves for randomly generated, Rayleigh faded channels for 2 by 2, 3 by 3, 4 by 4 and 6 by 6 sized channel matrices.

Figures 6.33, 6.34, 6.35 and 6.36 show the capacity CCDFs of the measured small duct channels. Note that each figure has the Rayleigh faded data and the base small duct channel measurement plotted and one other small duct channel measurement. While the difference in the CCDF curves between measurements are small, the implications of the results make sense. Taken as a whole, a slight increase in capacity is observed when circumferential antenna arrays are used, and when a hole in the ventilation duct is present. A slight decrease in the capacity is observed when RF absorber is added and when long monopole antennas are used. These results agree with the work presented in [77] in which a reduction in the number of modes reduced the channel capacity and arranging transmit and received antennas at different locations along the cross-sectional perimeter of the waveguide from one another increased channel capacity. The long monopole antennas reduce the number of modes excited in the duct and show a reduction in capacity. The circumferential arrays displace the antennas from one another around the perimeter of the duct cross-section and shows an increase in the channel capacity.

Lastly, we plot the base case, small duct and large duct capacity CCDFs in Figure 6.37. As we can see the 3 by 3 duct measurements are close to achieving the full theoretical capacity bound and the 6 by 6 measurements are significantly reduced from the theoretical capacity bound.




Figure 6.33: Capacity CCDF curves for Rayleigh generated random matrices and for two small duct measurement sets. The two measurement sets are the base case of an axial array of short monopoles, no duct hole and metal endcaps compared to a duct with absorber endcaps.




Figure 6.34: Capacity CCDF curves for Rayleigh generated random matrices and for two small duct measurement sets. The two measurement sets are the base case of an axial array of short monopoles, no duct hole and metal endcaps compared to a duct with a circumferential array of short monopoles.




Figure 6.35: Capacity CCDF curves for Rayleigh generated random matrices and for two small duct measurement sets. The two measurement sets are the base case of an axial array of short monopoles, no duct hole and metal endcaps compared to a duct with a hole in it.

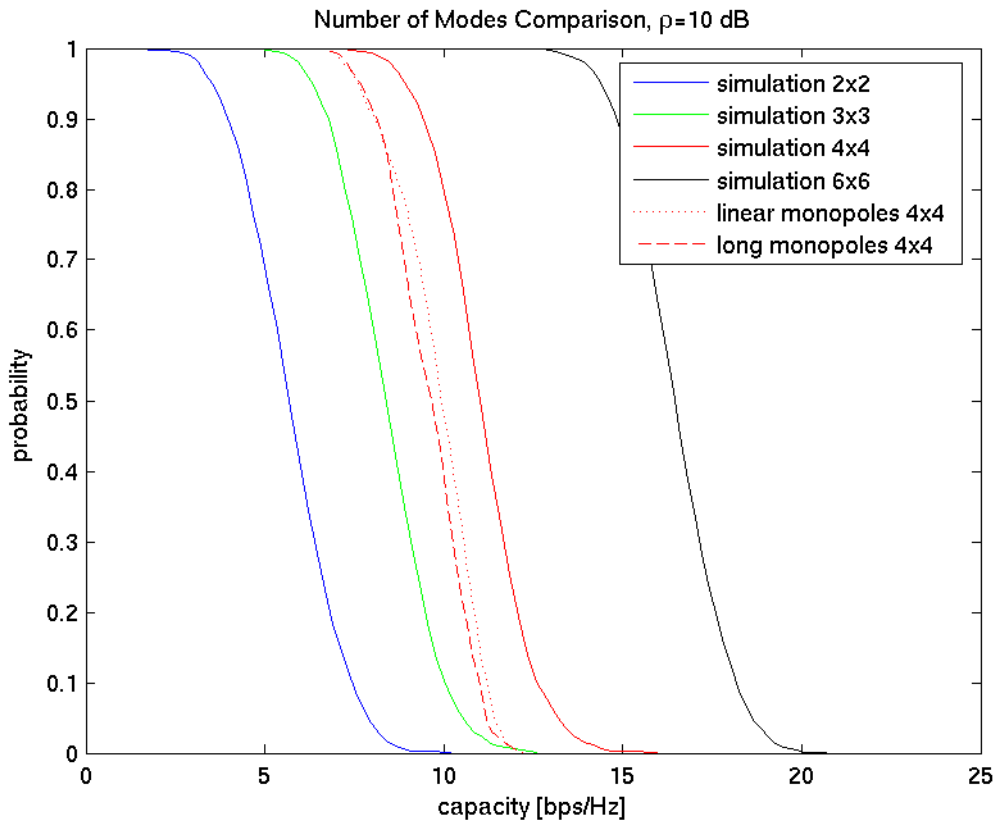


Figure 6.36: Capacity CCDF curves for Rayleigh generated random matrices and for two small duct measurement sets. The two measurement sets are the base case of an axial array of short monopoles, no duct hole and metal endcaps compared to a duct with an array of long monopoles.




Figure 6.37: Capacity CCDF curves for Rayleigh generated random matrices, the two large duct measurements and the small duct base case of an axial array of short monopoles, no duct hole and metal endcaps.

6.6 Conclusion

The major goal of this chapter was to investigate antenna array techniques for improving the system capacity of a ventilation duct based RF distribution system. To achieve this we first validated time-reversal as a technique for use at RF frequencies using indoor measurements. Building on this work, we have first investigated the impact that time-reversal focusing can have on duct based channels. We have measured several duct channels and used these measurements to show time-reversal focusing has a significant improvement in the gain and standard deviation in computed duct channel responses. We therefore believe that time-reversal focusing can provide a benefit in the link capacity for ventilation duct channels, thereby increasing the system capacity.

We have also investigated several interference nulling techniques. Since concepts of beam-forming and null steering do not directly apply in a ventilation duct environment which is inherently non-line-of-sight, we have considered several time-reversal techniques and a triangle nulling algorithm of our own design. These algorithms show some potential for use in ventilation ducts but are limited to at most about 10 dB of nulling. Whether this amount is of significant help depends on the installation.

Lastly, we have considered the potential impact of duct layout on the theoretical MIMO capacity of duct channels. Our results are in line with other published works showing significant capacity improvements and indicate a slight impact in channel capacity with duct layout. In general, it is advisable to install antennas at different locations around the perimeter of ventilation ducts and to use antennas that excite a rich mix of available propagation modes to maximize the available capacity.

7.1 Introduction

In Chapter 6, we have shown based on calculations with duct channel measurements that many time-reversal and MIMO communications techniques can be used to improve the system capacity available to a ventilation duct based wireless distribution system. In [17], the authors report a tremendous amount of potential bandwidth available if multi-element antenna arrays are used inside ventilation ducts. However, actual MIMO communications have not yet been attempted. The goal of the work presented here is to demonstrate MIMO based communications using antenna arrays inside ventilation ducts. Further, we show that even when using antenna array coefficients calculated at a single frequency but applied across a range of frequencies, significant RF bandwidths can still be used for communications. The first means of accomplishing this will be to use eigenvector decomposition of a MIMO duct channel matrix to form a diagonal transmission matrix and use FSK-2 (Frequency Shift Keying) data transmissions to validate the MIMO technique. Further, we show that even when using antenna array coefficients calculated at a single frequency but applied across a range of frequencies, significant RF bandwidths can still be used for communications.

The recent interest in the application of a time-reversal for use at electromagnetic, microwave frequencies will also be further explored in this chapter. Previously, several authors have presented demonstrations of communications at radio frequencies (RF) indicating that the technique is a valid one for use with electromagnetic waves [51, 60, 61]. Further, it has been noted that time-reversal can provide a reduction in the delay spread of a communications channel [78]. Additional work with time-reversal has shown it to be a technique with great potential for use in MISO (multiple-input single output) arrays [79, 80, 81]. In this section, we will explore the use of time-reversal for MISO array communications, thereby extending the theoretical work in Section 6.3. We will also investigate the use of time-reversal coefficients, rather than eigenvector diagonalization to create MIMO communications channels. To demonstrate the validity of time-reversal techniques, we use actual IEEE 802.11g equipment and quantify the system performance using network throughput measurements. While there has been significant interest in using time-reversal techniques with wireless LANs [82, 83], this will be the first demon-

stration of time-reversal used with actual wireless LAN transmissions of which we are aware.

7.2 Vector Modulator Based MIMO Communications

The relevant experiments were performed in the 2.4 to 2.5 GHz band using a hollow, cylindrical, metal ventilation duct, 9.1 meters long and 0.3 meters in diameter with flat metal endcaps as the physical transmission environment. As illustrated in Figure 7.1, the 2 by 2 MIMO setup was achieved using quarter-wavelength monopoles tuned to 2.45 GHz. The antennas were connected to the duct and aligned longitudinally along the length of the duct. The transmit and receive antennas were separated individually by 1 wavelength and by 8.0 meters from one another. Since we will be applying complex coefficients to the transmitted and received RF waveforms, we used 8 digitally controlled vector modulators that operate at RF to apply the desired magnitude and phase shift directly to the test waveforms. The vector modulators are capable of repeatably setting complex coefficients ranging from 0 to 40 dB of attenuation and any phase shift from 0 to 360 degrees across the frequency range of interest. The modulators can also provide a maximum attenuation with unknown phase shift of 70 dB. As seen in Figure 7.1, two vector modulators are paired with each antenna. In this manner, any arbitrary combinations of two input signals, T_1 and T_2 can be fed to both transmit antennas, and any arbitrary mix of signals at the two receive antennas can be fed to the two receive output ports, R_1 and R_2 .

The first step in the measurements was to characterize the channels between the transmit and receive antennas, $H_{mn}(f_q)$ using a Vector Network Analyzer (VNA). To do this, the transmit side of the VNA was connected through a power splitter to the two transmitter ports T_1 and T_2 . Likewise, the receiver side was connected via power splitter to the two receiver powers R_1 and R_2 . We then characterized the four H_{mn} channels between the transmit and receive antennas by selectively setting all but 1 transmit side and 1 receive side vector modulator to their maximum attenuation setting. In this way, $H(f_q)$ could be characterized without swapping cables. The measured channels are shown in Figure 7.2.

After the $H(f_q)$ matrix was characterized, we could compute the transmit and receive coefficients that would create two virtual channels between the transmit and receive ends. However, the vector modulators provide a single coefficient at all frequencies, while the channel responses, as seen in Figure 7.2, vary rapidly with frequency. We thus choose a single frequency bin, f_p of 2.45 GHz at which to create the MIMO channels.

It is a well know property of eigenvector bases that if we construct

$$D = X^{-1}HX \quad (7.1)$$

such that X is a matrix of the set of eigenvectors of H and X^{-1} is the matrix inverse of X , then D will be a diagonal matrix. In fact, the eigenvalues of H will make up the diagonal

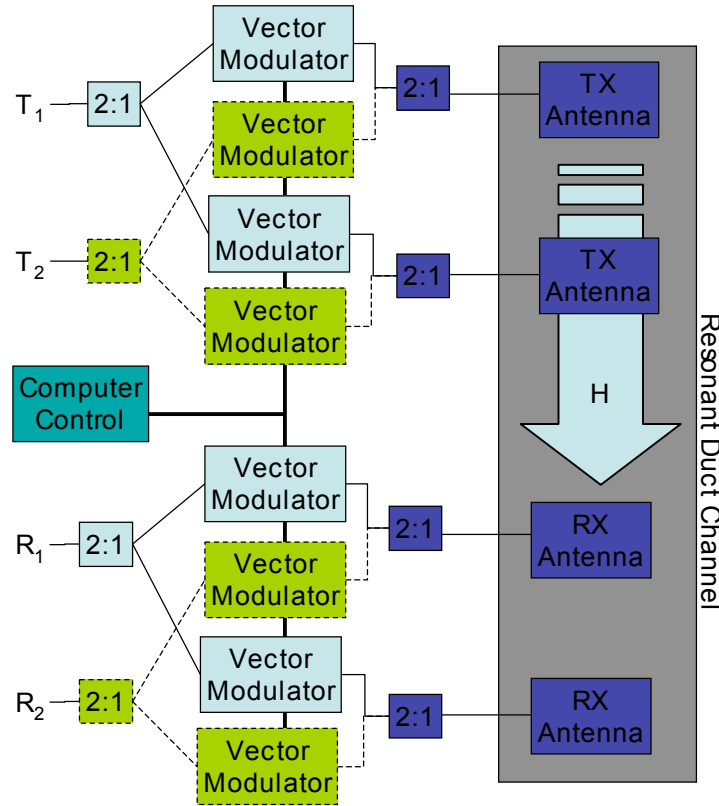


Figure 7.1: Setup of channel and equipment used in experiments.

elements of D . Equation 7.1 allows us to create two virtual channels between transmitter and receiver. This calculation was performed using the previously measured $H(f_p)$. The coefficients X and X^{-1} were applied to the appropriate vector modulators. We then utilized the same network analyzer measurement procedure to measure D across a range of frequencies around f_p . The resulting measured channels are shown in Figure 7.3.

Next we wished to confirm that two virtual channels had been formed and that they could be used over a reasonable bandwidth, despite being created using a single frequency point, f_p . To do this we used two signal generators and a vector signal analyzer to perform bit error rate (BER) measurements. Specifically, a pseudo-noise (PN) sequence, encoded via 2 tone Frequency Shift Keying (FSK-2) was repeatedly transmitted at a 1 megabit per second (Mbps) rate into one transmit side port (i.e., T_1 or T_2). A 1 MHz, frequency modulated, square wave tone was transmitted at the same frequency on the other transmit side port. Both signals had a 1 MHz maximum frequency deviation and so occupied about a 2 MHz bandwidth at 2.45 GHz. The resulting BERs and the measured channel gains are shown in table 7.1. The BERs clearly indicate that two MIMO channels


Figure 7.2: Measured channel responses forming the H matrix

Table 7.1: Measured gains and BERs of MIMO based channels. Each entry in the table is based on a single measurement.

Input	Output	MIMO Channel Gain, $ D(f_p) $	Bit Error Rate
T ₁	R ₁	-42.4 dB	$< 10^{-5}$
T ₁	R ₂	-81.6 dB	$\sim 50\%$
T ₂	R ₁	-67.6 dB	$\sim 50\%$
T ₂	R ₂	-53.0 dB	2.3%

were successfully created and that reasonable communications could be performed even using imperfectly applied MIMO channel settings.

In this section we have successfully shown that a 2x2 MIMO communications channel can be created in a reverberant environment. Even when channel variations are present across frequency, significant communication bandwidths can be achieved using MIMO coefficients applied based on measurements at a single frequency. The channel and BER measurements presented here confirm that MIMO techniques work well to create virtual communications channels even in the context of a reverberant ventilation duct environment.




Figure 7.3: Measured MIMO based channel responses created using an eigenvector basis, X of $H(f_p)$.

7.3 Time-Reversal MIMO Enabled IEEE 802.11g

In this section we present a practical demonstration of wireless local area network communications (WLAN) in the form of IEEE 802.11g signals, using a time-reversal focusing antenna array. We present throughput performance measurements of actual IEEE 802.11g transmissions using MISO and MIMO antenna arrays. Since the channel is static, it is possible to characterize the channel, perform calculations as necessary for time-reversal and apply the time-reversal to data transmissions, all without concern that the channel may have changed. Further, the highly reverberant nature of the duct channel should allow us to observe any improvement in performance due to a reduction in the delay spread of the channel from time-reversal coefficients.

7.3.1 Experimental Setup

Our measurement setup is shown Figure 7.4. Five cylindrical spiral duct segments, 0.3048 m in diameter and totaling 15.2 m in length were used for the measurements and had metal endcaps in place. Quarter-wavelength monopole antennas were inserted into the ventilation duct for all transmissions. A four element, axial antenna array with element spacings of one wavelength (12.24 cm) was installed halfway between two single element antennas. To apply time-reversal coefficients, we used vector modulators with

the antenna array. We used Herley Farmingdale 7122 vector modulators which use digital input signals and are interfaced to a computer parallel port via a custom circuit board and controlled via custom MATLAB[®] software. Each modulator is bi-directional and multiplies the RF signals passing through the modulator in either direction by a complex exponential, $Ae^{j\theta}$. The modulators are calibrated to provide amplitude coefficients, A , ranging from 0 to -70 dB and phase shifts, θ , of 0 to 360 degrees with an accuracy of ± 0.2 dB and ± 5 degrees over the magnitude and phase shifts and frequency ranges used in the experiment.

We wish to attempt two simultaneous connections via time-reversal, so we have placed two Linksys WRT54G, IEEE 802.11g Access Points (APs) inside metal project boxes. We filled the empty space inside these boxes with RF absorbing material and connected one antenna port to an SMA pass-through connector on the boxes. Power and ethernet cables were also passed through a hole in the boxes. The holes and all seams were sealed with aluminum and copper tape. When the SMA pass-through was terminated with a matched, 50Ω load, the APs could not be detected with a wireless LAN card, even right next to the boxes. Two laptops, 1 and 2, were connected via ethernet cables, one to each AP. These laptops were used for throughput measurements. Two other laptops, labeled 3 and 4, with Buffalo Wireless Air Station, WLI-CB-G54A, IEEE 802.11g NICs (Network Interface Cards) were used to connect to the duct endpoints and were also used for throughput measurements. Each AP was configured with a different SSID (Service Set Identifier) so that each NIC would connect only to the correct AP by setting the correct SSID. A final two laptops, labeled 5 and 6 in Figure 7.4, were used to control the settings of the vector modulators, which are labeled “IQ” in Figure 7.4 because they can also be conceptualized as IQ modulators. Lastly, 60 dB attenuators were added to the outputs of the APs. These attenuators were added to shift the total pathloss down to a level that allowed the IEEE 802.11 cards to use a moderate to low transmit rate. In this way any improvements due to time-reversal would be apparent. The attenuators also serve to reduce the power to a level that will not damage the vector modulators and also prevents the APs from detecting one another via leakage in the power divider modules. We also note that each of the cables used in parallel paths were of the same type (RG-142) and of matched length. Paired power dividers were also of the same type. These techniques were used to ensure each path through the RF circuitry had as close as possible the same phase and magnitude response.

To apply time-reversal coefficients, we needed to first characterize the full H matrix. To measure each component of the 4 by 2 matrix, we used a vector network analyzer (VNA). Rather than measure just the H matrix of the duct, we characterized each component including the vector modulators. This was done by connecting the network analyzer to the insertion points A , B , C and D indicated in Figure 7.4. Since our VNA can only characterize a two port network, we connected the VNA pairwise to the ports A - B , A - C ,

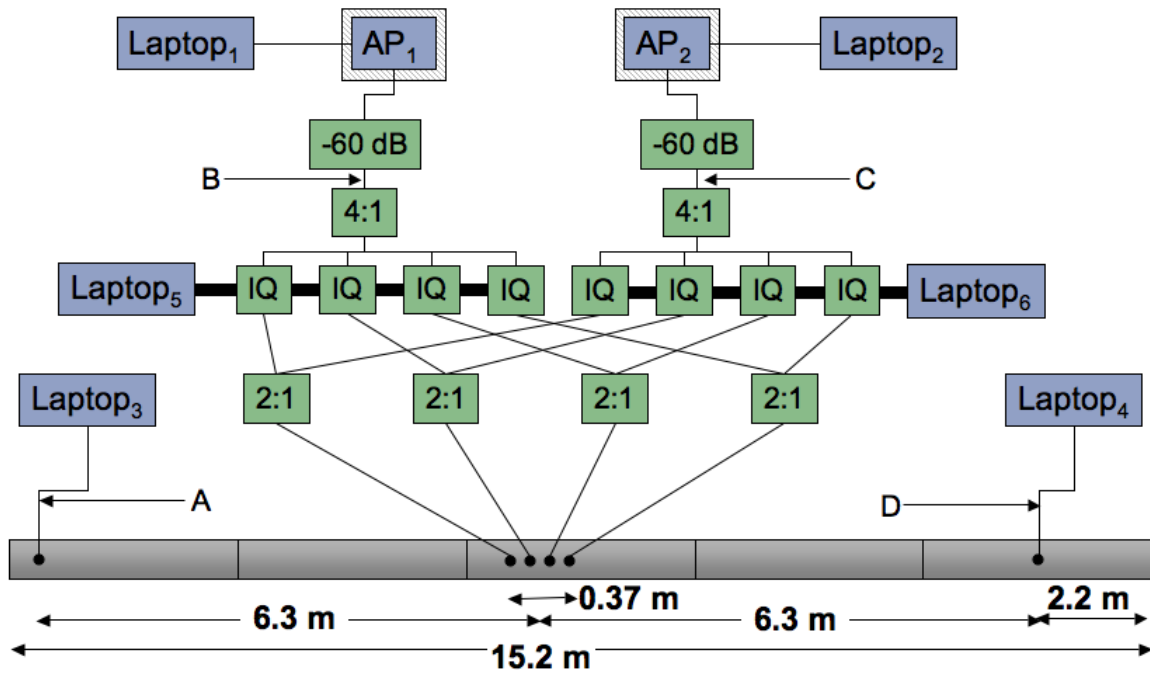


Figure 7.4: Measurement setup of time-reversal arrays for IEEE 802.11g measurements.

$D-B$ and $D-C$. For each connection, we used the vector modulators to characterize the $4 \text{ by } 2$ H matrix. This was accomplished by setting all of the vector modulators to a maximum attenuation and then setting a single modulator in one of the RF paths to a known, zero magnitude, zero degree reference point. In this manner each of the paths through the vector modulators and the full $4 \text{ by } 2$ H matrix through the ventilation duct could be characterized. As part of this process the magnitude and phase responses at the center frequencies of the first six IEEE 802.11g channels were recorded to use for computing time-reversal coefficients.

7.3.2 MISO Time-Reversal Focusing

After characterizing the ventilation duct H matrix with a VNA, we then connected the laptops shown in Figure 7.4. We first attempted time-reversal focusing using a single AP and single user location. This effectively reduces the $4 \text{ by } 2$ H matrix to a $4 \text{ by } 1$ matrix. These results are a practical demonstration of the theoretical work performed in Section 6.3. To measure the impact of time-reversal focusing on IEEE 802.11g communications, we have performed throughput measurements over the wireless LAN link using two different antenna array coefficients. The first set of coefficients was simply each vector modulator set to its 0 magnitude, 0 phase position. The second set of coefficients was the appropriate time-reversal coefficient setting on each vector modulator.

Due to limitations of the vector modulators, a minimum attenuation of 0.5 dB is required to ensure that any phase can be set. For this reason, when applying time-reversal coefficients, the computed coefficient with the lowest required attenuation was set to 0.5 dB. All other coefficients were scaled appropriately to maintain the correct time-reversal coefficient. This process ensures that the range of attenuations set by the vector modulators (18 dB in the measurements presented below) and the needed phase shifts are both well within the calibrated ranges of the devices. However, this procedure also prevents the time-reversal signal from being sent with the same transmit power as the equal phase case. While this prevents the comparison from being strictly fair, it is not necessarily unreasonable. While we do not anticipate an actual time-reversal system using the same technique as we use for these measurements, we do expect that the waveform shaping needed for time-reversal will reduce the maximum power output of a practical device somewhat, as is the case in our measurements. Therefore, we allow time-reversal to be at a disadvantage in terms of power, knowing that some limitations will be imposed in a practical deployment.

Throughput measurements in the two cases were performed using the Iperf measurement tool [28], which is an open source, multi-platform software utility. Throughput was measured in both the uplink and downlink directions by running a server and client on each computer as needed. TCP throughput was measured for 60 seconds for each test using a TCP window size of 1 megabyte. Laptops 1 and 3, shown in Figure 7.4, were used for the throughput measurements. Measurements were performed at each of the first six IEEE 802.11g channel center frequencies: 2.412, 2.417, 2.422, 2.427, 2.432 and 2.437 GHz. The resulting uplink and downlink throughput measurements are shown in Table 7.2. The table also has the calculated receive power and standard deviation of the frequency response for the total 4 by 1 channel. The received power shown is measured by first calculating the sum response of the four duct channels each multiplied by the appropriate equal amplitude or time-reversal coefficients. Then, taking into account the 18 dBm transmit power of the APs and the attenuators added, the expected receive power at the user laptop is computed by integrating over the 16 MHz span around the center frequency. The gain standard deviation shown in the table is also computed over the 16 MHz range.

The data in Table 7.2 indicates that the time-reversal settings are functioning as we would expect given the results found in Section 6.2. The standard deviation in the channel response generally decreased, though not as much as might be anticipated because of the single frequency nature of the time-reversal coefficients. Further the improvement in the channel gain is not the anticipated 6 dB, it does increase somewhat. This is reasonable given the limitations imposed by the vector modulators used. It is gratifying to see that the frequency response improvements do translate into improved throughput. Figure 7.5 is a plot of the receive power and throughput values from Table 7.2. The figure

Table 7.2: Throughput measurements and channel statistics for 4 by 1 antenna array with equal amplitude and time-reversal coefficients. Each entry in the table represents a single measurement.

Frequency	Receive Power	Gain Standard Deviation	Uplink Throughput	Downlink Throughput
Equal Amplitude Coefficients				
2.412 GHz	-74.6 dBm	5.1 dB	3.5 Mbps	2.9 Mbps
2.417 GHz	-75.1 dBm	5.1 dB	2.1 Mbps	2.1 Mbps
2.422 GHz	-74.4 dBm	5.6 dB	2.2 MBps	2.3 Mbps
2.427 GHz	-73.3 dBm	5.7 dB	3.5 Mbps	3.3 Mbps
2.432 GHz	-73.5 dBm	6.6 dB	3.5 Mbps	3.1 Mbps
2.437 GHz	-74.0 dBm	6.4 dB	2.1 Mbps	2.3 Mbps
Time-Reversal Coefficients				
2.412 GHz	-73.8 dBm	6.1 dB	3.7 Mbps	3.4 Mbps
2.417 GHz	-70.8 dBm	3.5 dB	4.2 Mbps	3.9 Mbps
2.422 GHz	-73.2 dBm	4.1 dB	3.9 Mbps	3.7 Mbps
2.427 GHz	-72.5 dBm	3.6 dB	5.9 Mbps	4.3 Mbps
2.432 GHz	-71.9 dBm	2.6 dB	4.7 Mbps	4.5 Mbps
2.437 GHz	-73.9 dBm	5.8 dB	4.4 Mbps	4.1 Mbps

suggests that the improvements in throughput observed could be explained solely by the improvement in signal strength due to the time-reversal focusing.

It is somewhat surprising that the single frequency time-reversal coefficient is effective over the range of frequencies that are used by the IEEE 802.11g communications. Figure 7.6 shows a comparison of the computed frequency responses for the two sets of antenna array coefficients applied for the first center frequency, 2.412 GHz. The figure also shows the computed response if ideal time-reversal coefficients were applied to the antenna array at each frequency, using the same attenuate-only power limitations imposed by the vector modulators.

It is surprising to note that the ideal time-reversal frequency response significantly differs from the single-frequency time-reversal frequency response except at the center frequency, where they are identical. However, there is some similarity in the two responses over the entire range. As seen in the figure the peaks and nulls of the single frequency time-reversal response seem to occur at many of the same frequencies as for the ideal time-reversal frequency response. Despite the fact that the single frequency time-reversal settings are only equivalent to ideal time-reversal over a narrow range, the time-reversal coefficient has some effect over a wide range.

7.3.3 MIMO Time-Reversal Focusing

We have also performed similar measurements for both APs operating on the same channel simultaneously. We have used the measurements of the ventilation duct frequency responses from each of the four AP side antennas to the two user antennas to create a




Figure 7.5: Uplink and Downlink Throughput versus Signal Strength as Computed from VNA Measured Frequency Responses.

2 by 2, time-reversal based, MIMO channel. This was done by applying single frequency time-reversal coefficients that would focus power onto one user location into one set of vector modulators and loading coefficients to focus on the other user location into the second set of modulators. The frequency responses formed in this way were calculated from the originally measured duct channel matrix using the equation

$$R = kHH^h$$

where H^h is the conjugate transpose of the duct channel matrix, H , and k is a power correction factor to accommodate the gain restrictions of the vector modulators. The above equation considers the antenna array to be the transmit side so H is a 4 by 2 matrix and R is the resulting MIMO channel matrix and is 2 by 2. The calculated R channels are shown in Figure 7.7. The figure shows frequency responses for the intended communications paths, R_{11} and R_{22} as well as the cross-talk channels R_{12} and R_{21} .




Figure 7.6: Frequency Responses for Antenna Array Excited by Equal Amplitude Coefficients, Single Frequency Coefficients and for Ideal Time-Reversal Coefficients.

The time-reversal coefficients applied to the vector modulators shown in Figure 7.4 allow transmissions from AP₁ to add coherently to the user location of Laptop₃, while the transmissions from AP₂ should add coherently to the user location of Laptop₄. To test this, we again measured throughput from Laptop₁ to Laptop₃. However, we also wished to transmit over both links simultaneously, so we also measured throughput from Laptop₂ to Laptop₄. Thus, we measured the throughput for each AP transmitting individually and for both APs transmitting simultaneously. As a baseline, we measured throughput for equal amplitude, zero phase coefficients and for time-reversal coefficients applied to the antenna array elements.

We again used the Iperf [28] throughput measurement software to measure uplink and downlink throughput. Throughput measurements of each link individually were performed using a 60 second measurement duration as before. However, it is difficult to perform throughput measurements exactly simultaneously. Thus, for the simultaneous




Figure 7.7: Signal and interference plots for the 2 by 2 MIMO created communications channels in the ventilation duct network.

transmission measurements, the throughput test for AP_2 was run for a duration of 2 minutes and started first. The throughput test for AP_1 was started about 10 seconds later and configured to last for 1 minute. Both throughput measurements were of the TCP throughput using a TCP window size of 1 Megabyte. Again we have measured the throughput in the case of equal amplitude and single frequency time-reversal coefficient excitations.

The throughput measurement results are presented in Table 7.3. The table presents the measurement results across several dimensions. First, we note that in all cases the time-reversal coefficients have resulted in an increase in throughput over the equal amplitude coefficients. We also notice that when transmissions are performed simultaneously, throughput decreases in all circumstances as would be expected. However, if we consider the total capacity of both APs, some interesting results emerge. In Table 7.3, we have

Table 7.3: Throughput measurements of IEEE 802.11g transmission performance using MIMO time-reversal arrays. All data is presented in Mbps. Each table entry represents a single measurement.

Individual Transmissions								
	Uplink				Downlink			
	Equal Amplitude		Time-Reversal		Equal Amplitude		Time-Reversal	
	AP 1	AP 2	AP 1	AP 2	AP 1	AP 2	AP 1	AP 2
	3.5	0.4	3.7	1.4	2.9	0.6	3.4	1.5
Mean	2.0		2.6		1.8		2.5	
Simultaneous Transmissions								
	Uplink				Downlink			
	Equal Amplitude		Time-Reversal		Equal Amplitude		Time-Reversal	
	AP 1	AP 2	AP 1	AP 2	AP 1	AP 2	AP 1	AP 2
	1.3	0.3	2.1	1.1	0	0.5	0.6	1.3
Total	1.6		3.2		0.5		1.9	

computed the mean throughput of the individual transmissions over the two APs and the sum of the throughput of both APs for the simultaneous transmissions. The direct comparison of these statistics is a way to compare the total throughput available to the AP. In this manner, it is apparent that for the uplink measurements, the interference has reduced the total capacity available from 1.95 Mbps to 1.6 Mbps when equal amplitude coefficients are used. On the other hand, the time-reversal coefficients actually enabled the AP to increase the available throughput from 2.6 Mbps to 3.2 Mbps. This result is not borne out in the downlink direction, however. In the downlink direction the total capacity for both the equal amplitude and the time-reversal coefficients has been reduced from 1.8 to 0.5 Mbps and from 2.5 to 1.9 Mbps respectively. However, in the equal amplitude case, transmissions were not possible at all for one of the two users whereas the time-reversal coefficients have allowed both users to transmit.

It is not clear why the time-reversal coefficients have enabled an increase in capacity in the uplink direction but not the downlink direction. It may be due to the imbalance in the measurement setup in that the APs cannot detect each others transmissions but the user NICs can detect and defer to transmissions from both APs and each other. Regardless, the above throughput measurements indicate that time-reversal may be a useful technique for increasing the capacity available to an AP in a single location and possibly as a method for reducing the effects of co-channel interference on IEEE 802.11g APs. Additionally, this technique may be able to reduce the network latency for wireless communications channels when multiple user transmissions take place by allowing users to transmit data simultaneously. Validation of this idea, however, is beyond the scope of this work.

7.4 Conclusion

In this chapter we have focused on demonstrating MIMO communications inside ventilation duct channels. Above we presented two practical demonstrations of actual MIMO

communications through ventilation duct channels. We have successfully shown that a 2x2 MIMO communications channel can be created in a resonant environment using an eigenvector diagonalization of duct channels. Even when channel variations are present across frequency, significant communication bandwidths can be achieved using MIMO coefficients applied based on measurements at a single frequency. The channel and BER measurements presented here confirm that MIMO techniques work well to create virtual communications channels even in the context of a reverberant ventilation duct environment.

The above work has also explored the impact of time-reversal based, multi-element arrays on the actual transmissions of IEEE 802.11g signals. We have demonstrated that in MISO configurations, an increase in throughput is readily apparent, even for practically constrained, single frequency time-reversal coefficients. We have demonstrated that this improvement in throughput is also apparent for MIMO array coefficients. We demonstrated that while the total capacity available to two APs operating simultaneously can be improved using time-reversal arrays, it is not necessarily always the case. However, the potential use of time-reversal antenna arrays for reducing co-channel, and perhaps even adjacent channel interference, has been identified. Lastly, we have shown that ventilation duct channels could benefit from the use of time-reversal antenna arrays.

8.1 Summary of Contributions

This work has focused on how to move from prior work investigating the feasibility and potential of ventilation duct based wireless signal distribution indoors to designable installations. As such, we have begun by focusing on how best to connect access points to ventilation ducts and how to ensure users in the building can reliably use the signal. Despite having considered several configurations, we found that one of the simplest configurations was also one of the most highly effective. Connecting an access point to a duct and letting power leak from inside the duct through air outlets was the technique that we decided was most effective. Our measurements have shown that exciting and preserving a single electromagnetic mode inside ventilation ducts is feasible. However, we have also found that frequent air outlets in the ventilation duct network serve to reduce the delay spread sufficiently to make such modifications unnecessary. We have also found that the large number of air outlets present in the ventilation duct networks we worked with allowed significant power to radiate from the ventilation ducts. As a result, re-radiating antennas connected to ducts were found to be unnecessary. Thus, we recommend in this work a system architecture of directly connecting access points to ventilation ducts and allowing power to radiate from air outlets in the ducts.

To facilitate design of ventilation duct based wireless distribution networks, we have gathered a significant amount of propagation data and developed a pathloss model from it. This model does not require any building structure information other than the location of the ventilation ducts, making it easy to implement in software simulation. It can also be reduced to simple design rules for pencil and paper design. We have shown with the development of custom MATLAB[®] software that the prediction model can be implemented in such a way as to allow quick simulation times. Despite the ease of use of the model, we have shown that it can be accurate for system design even when applied to a building structure previously unconsidered. The model and the parameter values in the model that we have computed from several measurements are fully tabulated in this work.

We have used our system design concepts and design model to fully install a 2000 square meter, 3 story building with an IEEE 802.11g, ventilation duct based system. We have fully characterized the signal strength and throughput provided by the system to

areas throughout the building. Further, our installation provides extremely even signal strength and transmission rates to all covered areas. As a result the system should provide a more constant and fairer experience for the user base. As part of characterizing the throughput performance of our demonstration system, we have developed a simple linear conversion that can be used to estimate the saturated user throughput available based on a signal strength measurement. In addition, we have explored the use of switching diversity mechanisms in current IEEE 802.11g access points and the use of additional bi-directional amplifiers for system improvements. We have found the switching diversity mechanism to be compatible and effective with ventilation duct based wireless distribution, but that the high delay spread found in ventilation ducts limits the usefulness of additional amplification.

As part of our efforts to further improve a ventilation duct based distribution system, we have presented some of our work using antenna arrays for improved signal quality. We presented a novel antenna array utilizing long monopole antennas that is capable of exciting one of two particular propagation modes in a commonly sized ventilation duct. The mode excited depends on whether the elements are driven in phase or out of phase. We have also shown that such an array could be used with existing access point switching diversity via a 180 degree hybrid splitter to good effect. Such a splitter would allow the access point to excite each of the modes individually depending on the selection diversity setting of the AP. In this way, the antenna array allows diversity to be maintained while still providing the benefits of single mode propagation. We should mention that the impedance of the antenna array is not 50 ohms, so the performance of the array would improve with the use of impedance matching circuitry.

In addition to the single mode antenna array, we present some measurements of combination diversity that we have implemented using vector modulators that function at 2.4 GHz. These modulators have allowed us to demonstrate the effectiveness of antenna arrays for use with IEEE 802.11b equipment. These measurements indicated that even simple antenna arrays can be used to significantly change the performance of a wireless local area network connection using only phase shifts of the antenna elements. Further, these measurements foreshadowed the potential for time-reversal focusing.

A significant contribution of this work was to demonstrate that time-reversal focusing was a valid procedure for use with microwave frequency range, electromagnetic waves. Previous effort had only shown the usefulness of time-reversal in acoustics [24] and ultrasonics [25]. As a result, we present our demonstration of time-reversal focusing in a cluttered laboratory environment. In this work, the presence of significant multipath enabled spatial focusing beyond the Rayleigh criterion and depth of focus that would be expected for the size of the antenna array and distance from the receive location. This focusing was enabled by multipath present in the environment.

To investigate how time-reversal can be used in ventilation ducts, we have performed

a significant computational study regarding the potential of time-reversal focusing and time-reversal nulling techniques. This study compared the impact of time-reversal focusing and nulling on different ventilation duct configurations. We showed that channel improvements using time-reversal focusing with measured ventilation duct frequency responses, were inline with the improvements seen in simulated Raleigh channel models. Time-reversal focusing enabled an increase in channel gain and a decrease in gain standard deviation in both simulated duct channels, small laboratory duct setups and larger duct networks.

Time-reversal nulling measurements with the same channels were also considered. As part of the work, we developed a new nulling algorithm termed “triangle nulling.” The triangle nulling algorithm utilizes a geometric argument to perform nulling even in the presence of multipath that would prevent conventional antenna array null steering techniques that use direction of arrival estimation from functioning. Three different versions of the triangle nulling algorithm, and two different time-reversal nulling techniques were compared for different duct channels. In these measurements, we found that the triangle nulling techniques all functioned well, especially for increasing numbers of antennas. We also realized that time-reversal nulling using antenna arrays is limited because the nulling algorithm cannot effectively focus at multiple antenna array elements simultaneously.

Since the finalization of the IEEE 802.11n standard is a highly anticipated addition of MIMO technology to IEEE 802.11g, we also analyzed the theoretical channel capacity of different ventilation duct arrangements. This work should give some insight in how best to connect IEEE 802.11n equipment to a ventilation duct when it becomes available. To do this we have considered several ventilation duct arrangements and considered the impact to the theoretical channel capacity.

Finally, we have implemented MISO and MIMO communications in ventilation duct channels using actual data transmissions. We successfully demonstrated two simultaneous FSK-2 transmissions through a ventilation duct over a two by two channel using eigenvector diagonalization to compute channel coefficients. We have also demonstrated tangible improvements to data throughput for IEEE 802.11g signals transmitted over a four by one MISO channel through a ventilation duct and shown that time-reversal focusing enabled MIMO communications has the potential to improve the capacity of an AP. Using simultaneous transmissions from two APs over the same antenna array, to two different NICs showed that time-reversal MIMO might be able to increase the capacity available to an AP or perhaps decrease co-channel or adjacent channel interference in a multiple access point installation.

8.2 Future Work

While this thesis has focused on creating a full solution for designing ventilation duct based wireless distribution networks, there are avenues of research still open. For in-

stance some implementation details may still be needed to be worked out. Extending the propagation model for duct based propagation to other duct dimensions and types may be needed as installations are performed. Also, details such as a means to provide impedance matching for the long monopole antenna arrays or a means of combining such circuitry with a 180 degree hybrid splitter will need to be worked out when an array is installed in a commercial system. Another implementation detail could be researched is how to implement time-reversal focusing techniques in a more commercially viable manner than the vector modulators used throughout this thesis. One idea is to implement the needed time-reversal waveform shaping at baseband and then mix the signal up to the microwave carrier frequency.

As part of conducting measurements and developing models of ventilation ducts, it became increasingly useful to develop custom software to model building duct networks. We will describe the interface and workings of this software in this appendix. We refer to the software we have developed as Duct Mapper.

A.1 Basic Interface

Figure A.1 shows a new window that opens when launching the Duct Mapper software. As can be seen from the window, there's a blank area for a blueprint of a building to be displayed. Menus at the top of the window are used to launch various commands. The software is launched from MATLAB®. The software uses MATLAB® figure files as the starting point. The blank area in the figure is a plotting axes where we can place images of buildings, plot ducts, measurements and predictions. The `File` menu is for opening, closing, printing and saving new Duct Mapper files. The `Edit` menu contains various prediction model prediction settings. The `Display` menu allows the user to turn on and off the display of many different parameters including building structures, user notes, measurements and simulations. The `Locate` menu provides commands for setting the location of building elements and user notes in the model of the building. The `Measure` menu is for measuring actual system performance parameters and bringing the results into the software. Lastly, the `Plot` menu has some analytical tools.

Basic file commands are kept in the `File` menu shown in Figure A.2. As can be seen, several standard commands are present. The `New` command opens a new Duct Mapper window exactly as shown in Figure A.1. The `Open...` command will allow the user to open a previously saved Duct Mapper figure. Duct Mapper files are saved as standard MATLAB® `.fig` files so they can be opened by any MATLAB® user without any of the associated Duct Mapper software. The `Close`, `Save` and `Save As...` commands close the Duct Mapper window or saves it to a file. It should be noted that there is no feature implemented to prevent the user from saving a modified building layout. As a result, a user can change something, close the window and loose any work done without warning. We will discuss the `Import` submenu further below. The `Export as Image...` command saves the axes area to an image file. Lastly, the `Page Setup...`, `Print Setup...`, `Print Preview...`, and `Print...` menu commands open standard MATLAB® printing




Figure A.1: A new window that opens when the Duct Mapper software is launched from MATLAB®.

windows for printing the layout.

The `Import` submenu is shown in Figure A.3. In this menu, we have various commands that are used to import images, building structures or measurements into the




Figure A.2: The Duct Mapper File menu.

current layout. We will discuss the building structures available in Section A.3 and measurement data in Section A.4. In the meantime we will mention that the import of a building image is the usual starting place for modeling a duct network in a building. We will therefore use this as a starting point to explain the normal process of importing an

image of the building to use as a background image for later modeling of the duct layout etc.

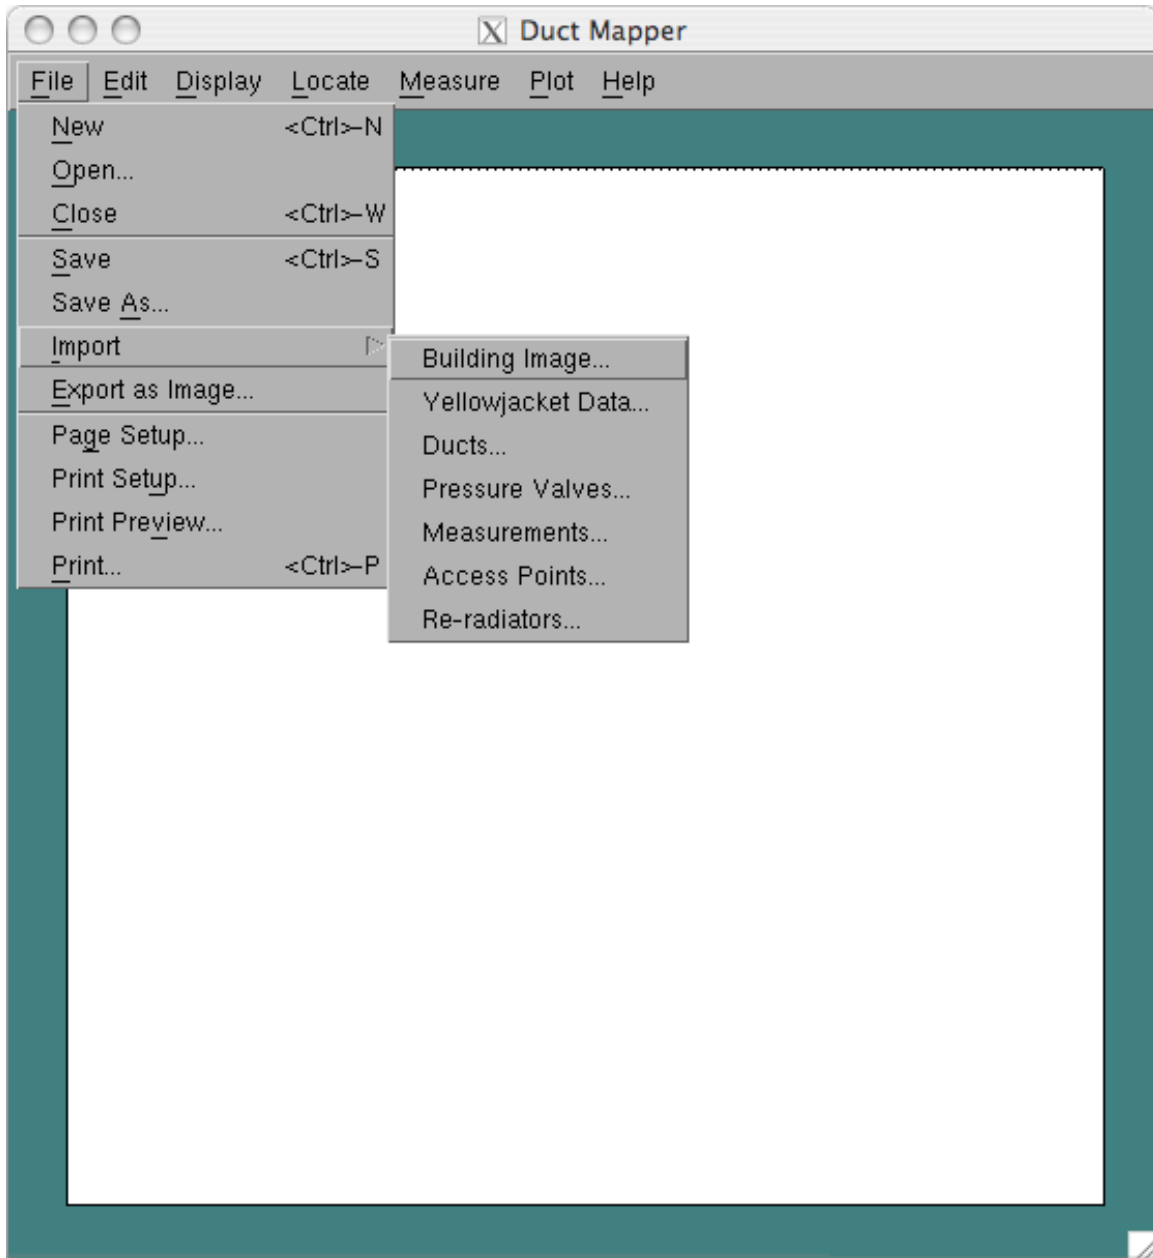


Figure A.3: The Duct Mapper `Import` submenu.

A.2 Building Image Import Procedure

To import a building blueprint image, we select the Import, Building Image... command, which starts a process of importing an image in a format useful for modeling work. The first step is to select an image file of any format MATLAB® is capable of reading. After selecting an image, the first step is to check if the image needs to be mirrored or rotated. The image is first plotted as MATLAB® imports it and asks if the image should be mirrored, if it should be rotated and if so, by how much. An example of this is shown in Figure A.4.

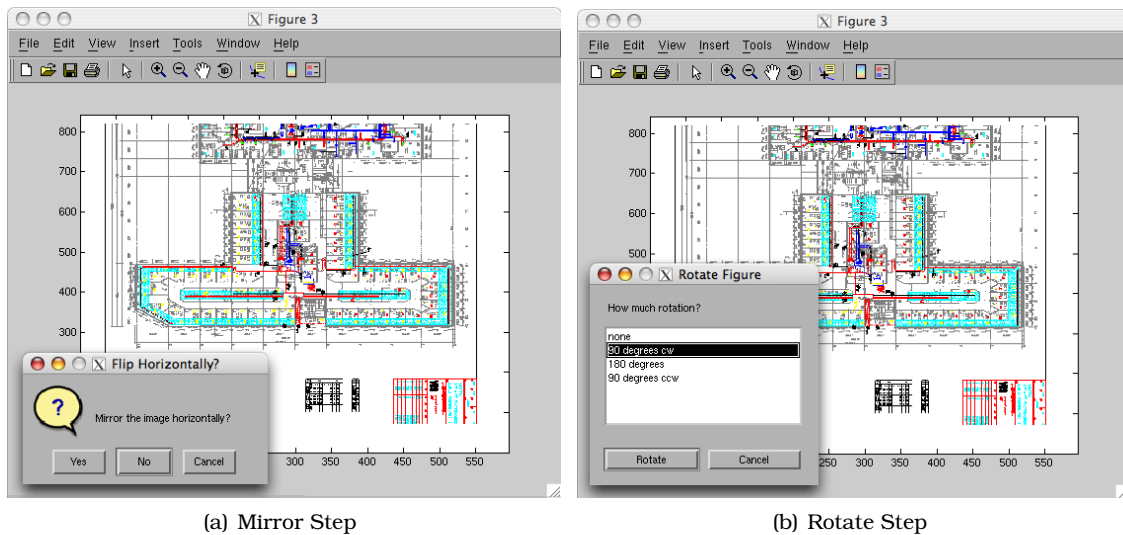
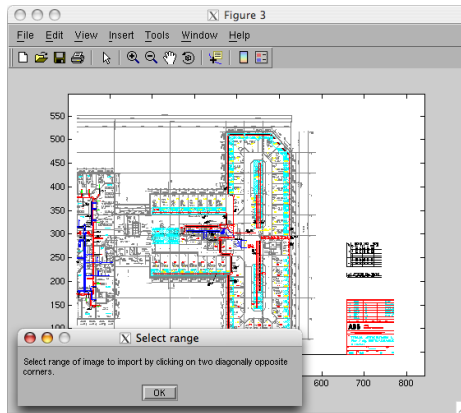
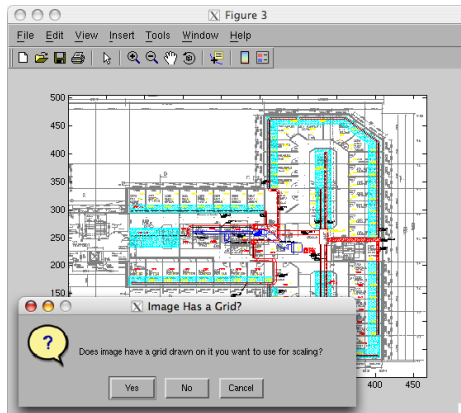


Figure A.4: First two steps in building image import process checking with user whether or not to mirror and/or rotate the image.

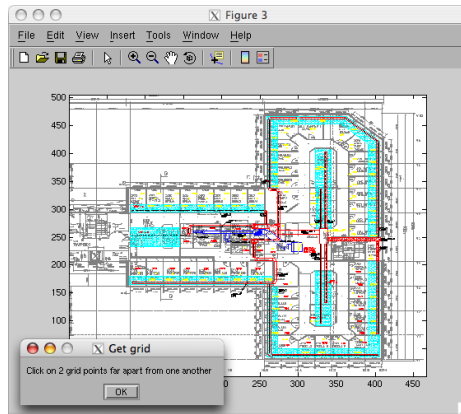
After each mirror and rotation step, the software redisplay the image based on the user's responses. The next step is to pick the range of the image the user wants to keep by clicking on two corners of a rectangular area to keep which is shown in Figure A.5(a).



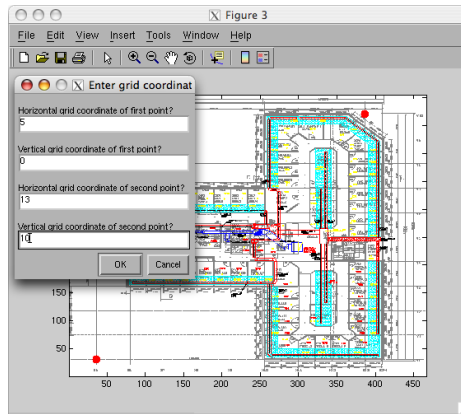
(a) Select Image Range



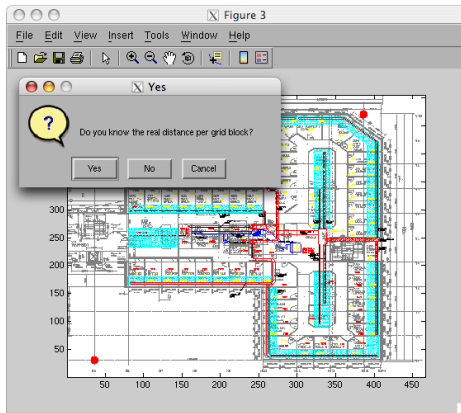
(b) How to Scale Image



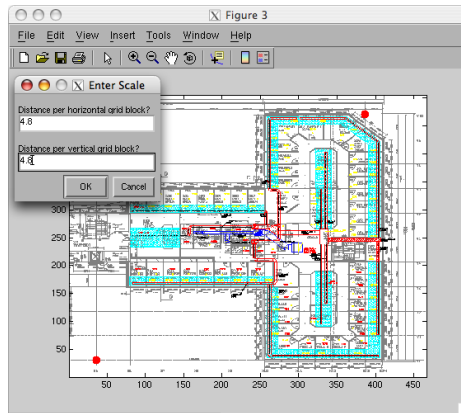
(c) Scaling Instructions



(d) Identify Grid



(e) Grid Size Question



(f) Grid Size Entry

Figure A.5: Further steps to importing a building image into Duct Mapper software.

The final step in importing a building image is to associate a real distance with the imported image. To do this we introduce a concept of an arbitrary grid on top of the image. This idea came about because many blueprints we had available had regularly spaced grids drawn on them with simple numbering schemes that are convenient for locating points in the building on paper. By drawing this grid over the image we can easily match a particular location on paper to a location on the image. This can drastically simplify entering measurement data. For this reason, as part of the import process, the next step is to check if the image imported has a grid drawn on it as shown in Figure A.5(b). The next step is to instruct the user to click on two grid points that are separated as much as possible from one another as shown in Figure A.5(c). The purpose of selecting points far apart is to minimize the error associated in slight imprecision in clicking on the image while it is shown on a screen. It is therefore recommended that the user maximize the image size while clicking on the grid points. After clicking on two grid points, the clicked points are highlighted in red and the user can enter the horizontal and vertical coordinates for these two points as shown in Figure A.5(d). After identifying the grid, we can provide scale to the drawing by entering the size of a grid block by entering its size in meters as shown in Figures A.5(e) and A.5(f). Once this process is complete, the image is displayed in the Duct Mapper window as shown in Figure A.6. As can be seen in the figure, dotted lines and axes labels indicate the software calculation of the grid location in the image. If the import process was performed correctly the dotted lines will align with the grid in the image.

A.3 Duct System Structure Modeling

A.3.1 Locating Ventilation Ducts

After an image has been imported into Duct Mapper, the `Locate` menu can be used for modeling the locations of ventilation ducts and various other items in the software. To add, move or delete ducts to the software, one uses the `Ducts` submenu under the `Locate` menu as shown in Figure A.7.

After clicking on the `Add New` submenu item, the software reminds the user of how to draw new ducts by displaying helpful text information in the MATLAB® command window as shown in Figure A.8. These instructions explain that ducts are drawn as line segments one after another. With each click (after the first one) the software adds a new line segment that represents a duct and draws a wide grey line to represent that section of duct. Line segments cannot be broken into subsegments so it is important for the user to draw segments that do not have junctions in the middle of the line segment. As the user draws ducts, a small blue plus sign is drawn to indicate the starting point of the line segment. To stop continuing a duct run and begin a new run, the user presses the escape key. This will erase the blue plus sign to indicate a new duct run is about to begin. At any time, the user can use a right-click or ctrl-click to force the newly added duct to be drawn

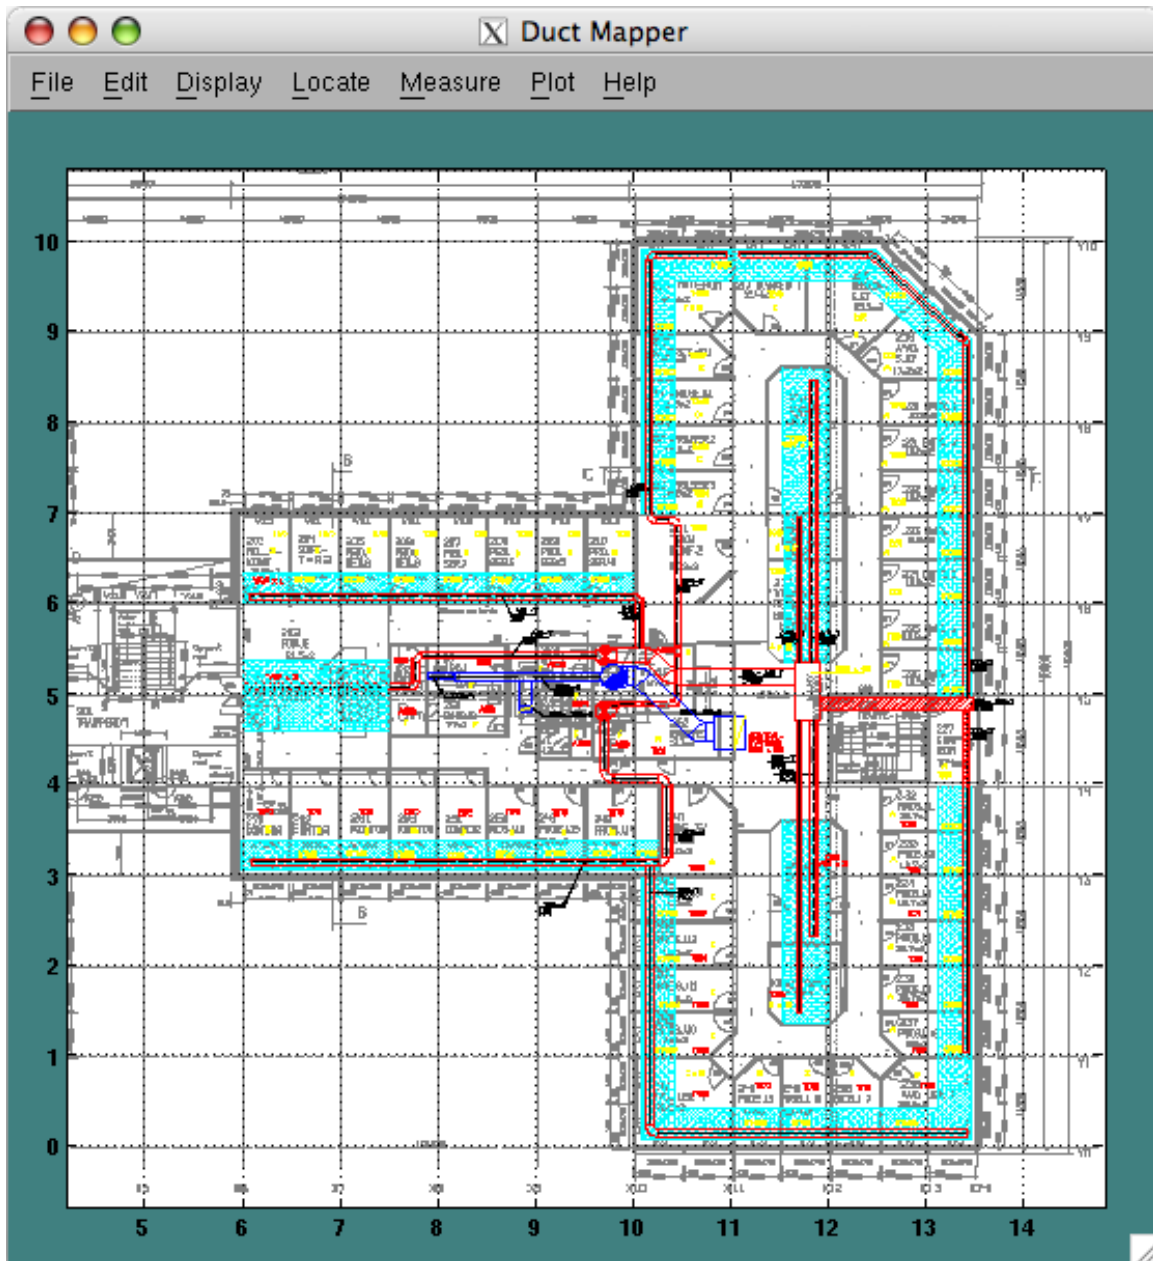


Figure A.6: Duct Mapper display after importing a building blueprint image.

at a right angle. It is highly recommended that the user always use this technique for drawing straight ducts. In this way the software can automatically distinguish between right angle and slant angle duct junctions. In addition, the user can always middle-click or shift-click to attach the next line segment to an existing duct segment endpoint. This

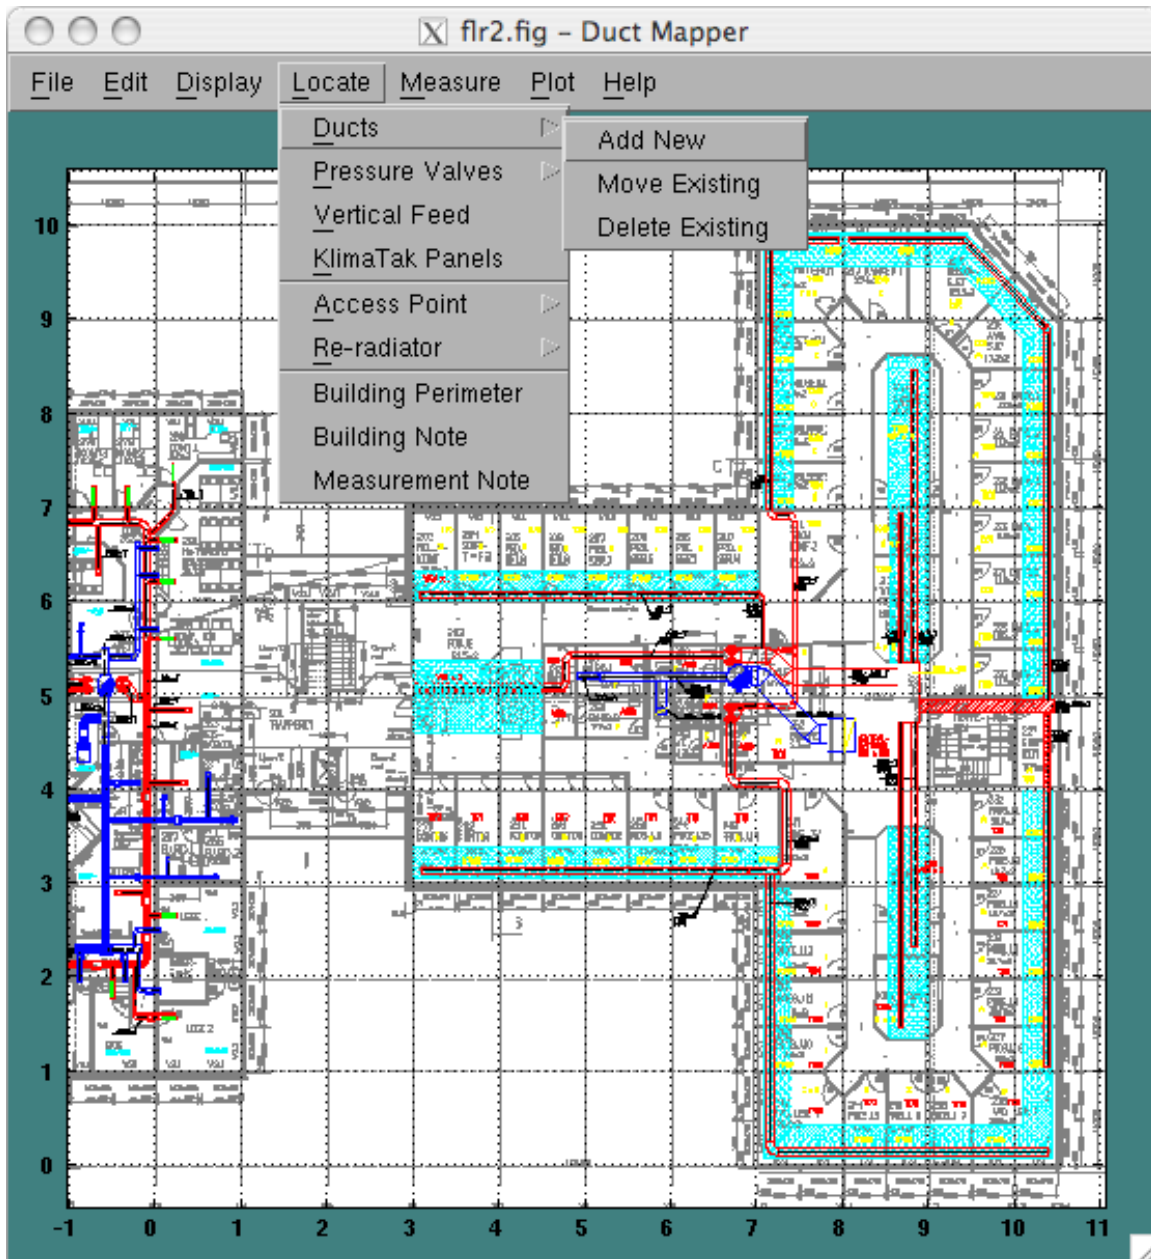


Figure A.7: Locate submenu for adding, moving and deleting ventilation ducts.

is also an import function for the user to always use. This function allows the software to distinguish between two duct segments that are connected together and duct segments that are merely close together. Attempting to manually click on another duct endpoint will not connect the two segments together. Lastly, the user can finish adding ducts by

hitting the escape key twice.

```
>>  
Left click at duct corners to add ducts to the drawing.  
Hit ESC to end a run of duct. Hit ESC again to stop adding ducts  
Right-click or Ctrl-click to draw ducts at right angles.  
Middle-click or Shift-click to attach to the nearest duct junction.  
>>
```

Figure A.8: Instructions displayed in the MATLAB® command window that instruct the user how to draw ventilation ducts in the figure.

A picture of the process of drawing ducts in the layout software is shown in Figure A.9. As can be seen in the figure the user is tracing over the red duct outlines with the grey line segments that model the duct location in Duct Mapper. As can be seen, the next duct segment will be drawn from the blue plus sign near grid location (7,0).

If a mistake is made while drawing ducts, the user must first stop the process of adding ducts by hitting the escape key twice. To delete a duct the user then selects the `Delete Existing` submenu item and then clicks on a grey duct segment to delete it. The nearest duct to the clicked location will be deleted. There is no undo function, so deleting ducts should be handled with care. In addition, the duct delete function is still enabled until the user hits the escape key. Thus, with every mouse click a duct is deleted unless there are no ducts left in the drawing. Moving an existing duct is a function that has not yet been implemented in the software.

After all ducts have been modeled in the software, the user can begin to model other items in the building. The `locate` menu allows the user to add and remove pressure valves with the `Pressure Valves` submenu item. To add a pressure valve, the user selects the `Add New` submenu item and then clicks on a duct segment to place a pressure valve. With each click a new valve is added as indicated by a black dot drawn on the nearest point on the duct to which the user clicked. When finished adding valves the user presses the escape key. Valves can be deleted by selecting the `Delete Existing` submenu and clicking on each pressure valve to remove.

Another way to add ducts to a figure is to import these locations from another Duct Mapper figure. By selecting the `File, Import, Ducts...` item. After choosing a MATLAB® figure file to import ducts from, the user will be asked if they wish to replace the ducts in the figure with those from the other file as shown in Figure A.10. It should be noted that this procedure will only work if the grid setup for the two drawings is the same. The software stores all items based on grid location, which is an advantage for manually adding items or referencing paper based drawings. However, it means that care should be taken to ensure the grids in the two figures are setup the same way before attempting to import ducts. Further, importing ducts should not be attempted if there are APs or simulations in the drawing as this could confuse the software.

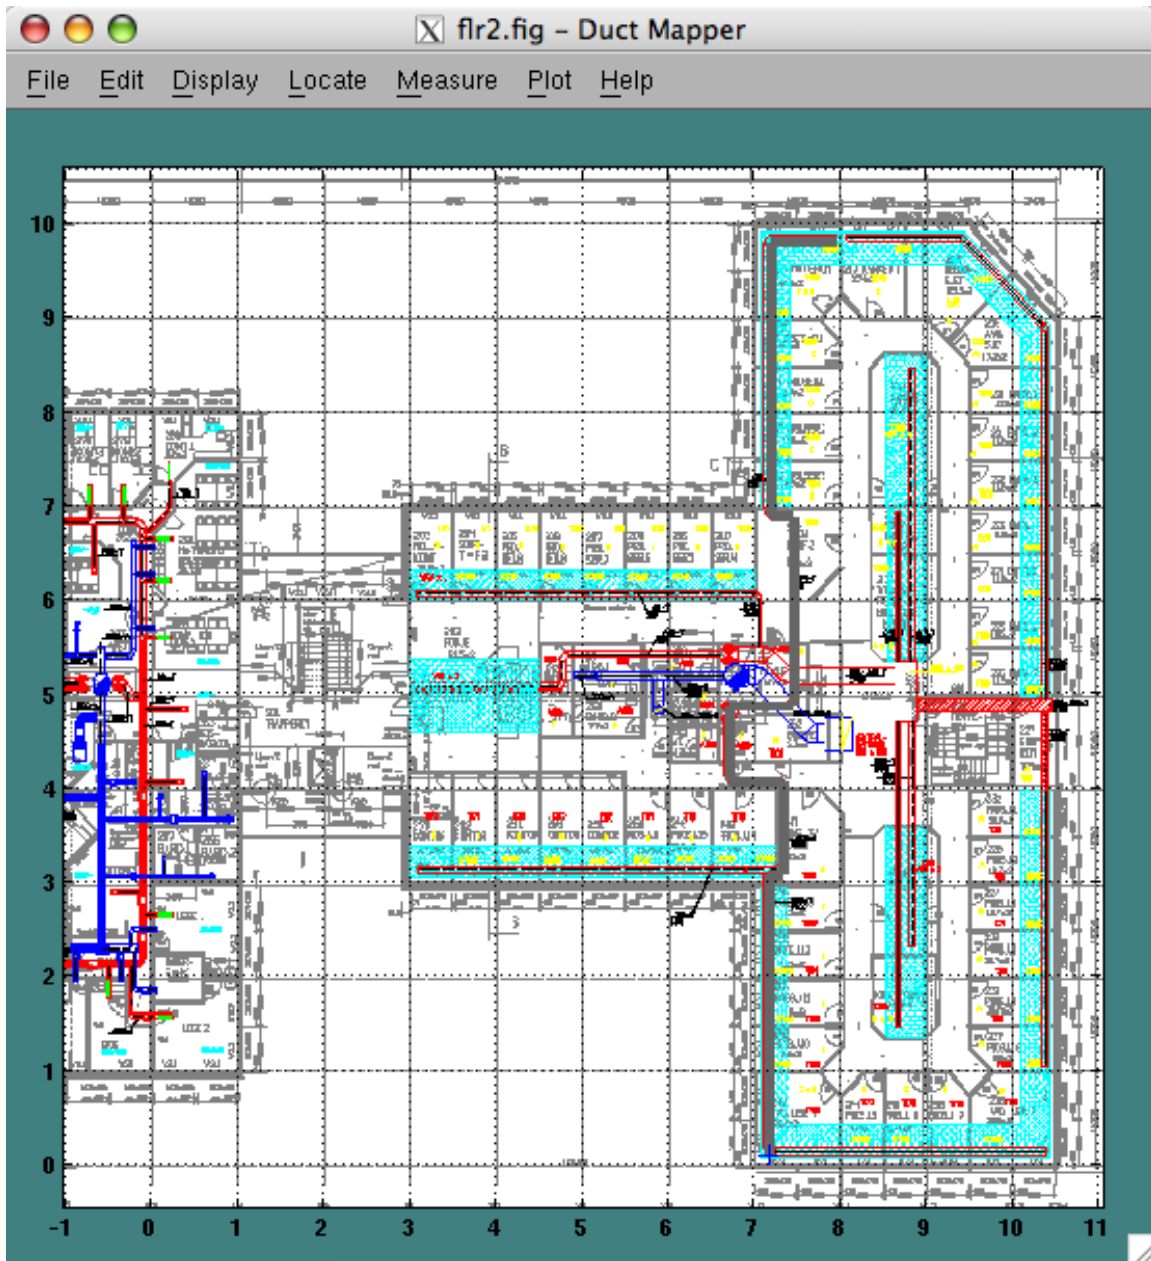


Figure A.9: Picture of the user adding new ducts to a Duct Mapper layout.

A.3.2 Locating APs and Re-radiators

As can be seen in Figure A.7, there are menu items for locating vertical feeds and KlimaTak panels. Unfortunately these menu items do not yet have implementations. However, a user can add, relocate and delete APs and re-radiators using the related `Locate`

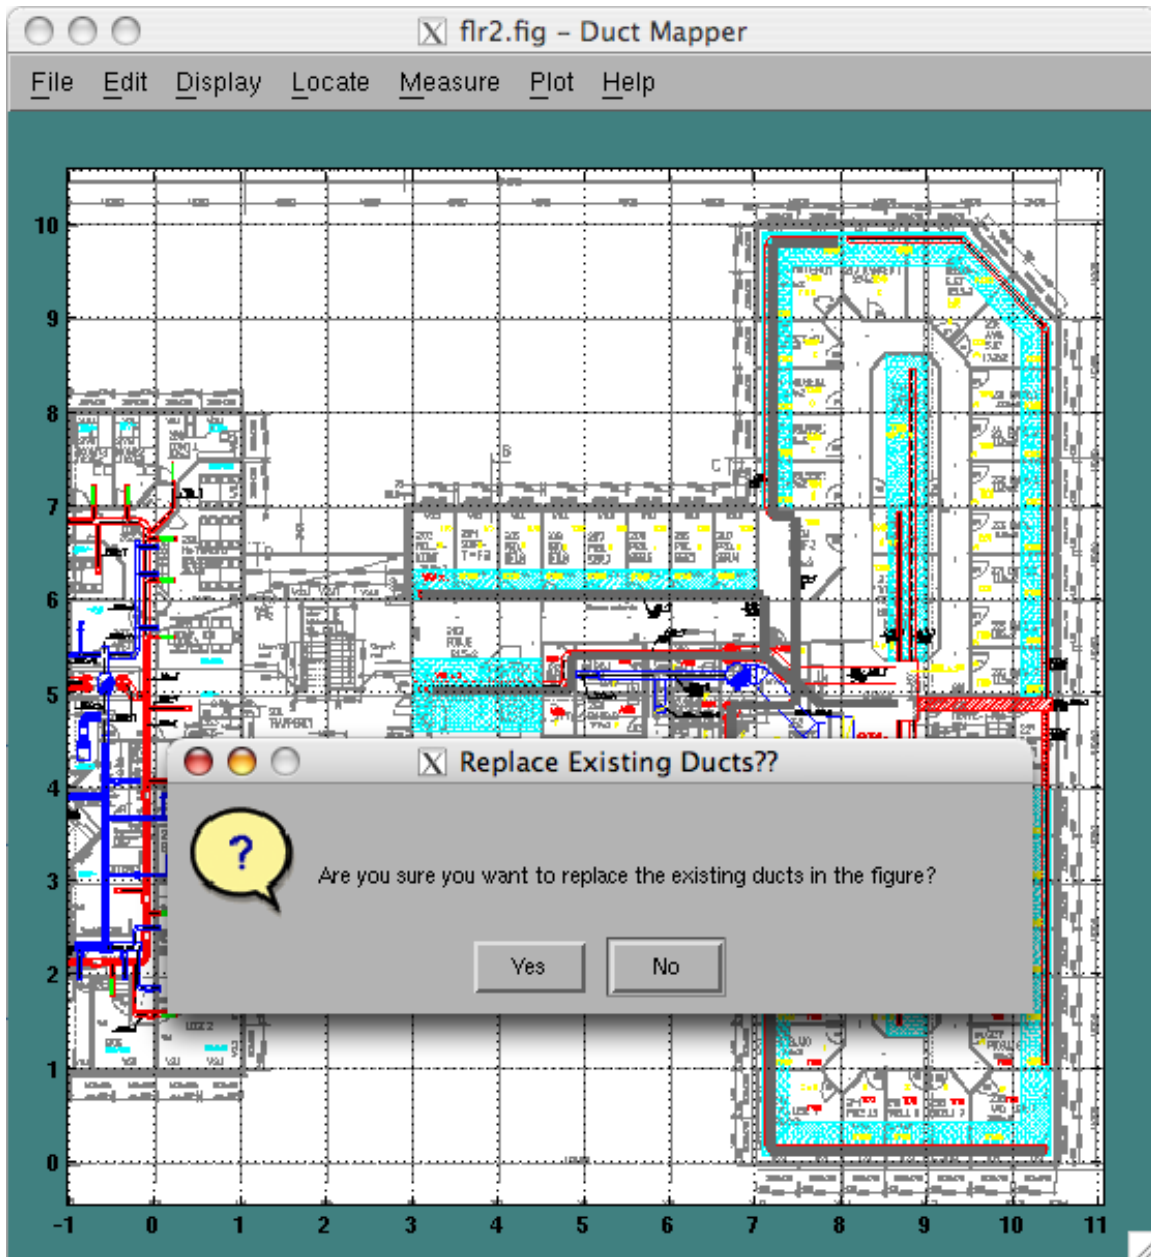


Figure A.10: Query as to whether or not the user is sure about replacing all of the ducts in the existing drawing with those from another figure file.

submenus. The processes for APs and re-radiators is the same, although APs and re-radiators are drawn by the software slightly differently. An AP is plotted as a purple dot labeled "AP" and a re-radiator is plotted as a pink dot labeled "R." Since the processes are

identical other than this, we simply describe the processes for APs.

To add an AP simply select the `Locate, Access Point, Add New` menu item. Then, each mouse click will add an AP to the nearest point on the nearest duct to the mouse click. To finish adding APs the user presses the escape key. To delete a duct simply use the `Delete Existing` submenu item. Again, the escape key finishes deleting APs. The `Move Existing` submenu item allows the user to first select an AP by clicking on it, so that it will be highlighted. The user then clicks a duct to move the AP to a new location.

It should be noted that the software automatically assigns a transmit power of 18 dBm to each AP. The software also has the ability to associate an IP address and MAC address with each AP. However, no user interface has been created for this. Instead use the MATLAB® commands `»getappdata` and `»setappdata` to access the “aps” parameter stored in the figure. For example: `»getappdata(gcf, 'aps')`.

A.3.3 Locating Building Perimeter

Since building blueprints can sometimes shown more than just the building structure the user is interested in, Duct Mapper allows the user to hide everything but the building by drawing the building perimeter. To model the building perimeter in the software the user draws line segments indicating the outside edge of the building. This process is activated using the `Locate, Building Perimeter` menu item. This process is similar to the procedure for adding ventilation ducts. A black line is drawn for a single continuous running line segment. A blue cross indicates the starting point of the next line segment. Again, right-click or ctrl-click draws line segments at right angles and hitting escape finishes the process. Only a single continuous line segment can be drawn around the building. After you have drawn the segments a final segment will be added from the first endpoint to the last endpoint forming a closed loop. Everything outside of the loop will be hidden from the user’s view. Segments cannot be added, moved or deleted as with duct segments. Instead, each time the user selects the `Building Perimeter` menu item the entire perimeter is replaced by the new one drawn. This feature will be of particular use for hiding parts of simulations that are outside the building and thus not necessarily valid or useful for display.

A.3.4 Adding Notes

The user can add two types of notes to the software as a way to record building structure peculiarities or measurement procedure details. The processes for adding these notes is the same so we will only describe the process of adding a building note. First the user selects the `Locate, Building Note` menu command. Then the user is instructed in the MATLAB® command window to select a point in the building that the note will refer to. After clicking on a point, the user is asked to select the top left corner to start the text from. After selecting these two points, the user is allowed to enter a text message which is then displayed.

A.3.5 Completed Building Model

Once the process of modeling the ducts, duct valves, APs, re-radiators, and building perimeter is complete, the display of the building will look something like Figure A.11. The figure shows a properly imported building image with ducts drawn on top, a building perimeter indicated and an AP and re-radiator located. The Duct Mapper software allows the user to turn on and off the display of these various items. The display of the background building image, ventilation ducts, APs, re-radiators, building and measurement notes can all be toggled on or off by simply selecting the appropriate `Display` menu item. Whether or not each item is displayed is indicated by marks of the appropriate menu item, which can also be seen in Figure A.11.

A.4 Measurements

One advantage of the Duct Mapper software is the ability to coordinate the measurement of various performance parameters like signal strength or throughput and to record the data in a map of the environment. In this way the locations of APs, ducts and the user can all be easily recorded. The measurements that can be configured and initiated from the `Measure` menu within the software as seen in Figure A.12. In later sections we will discuss each of the measurements available in this menu. First we discuss the basic measurement procedure.

To perform a measurement the user simply enables the desired measurements in the `Measure` menu and then starts measurements using the `Get Data` submenu item. The measurement proceeds in which the user moves to a particular location in the building and clicks on that location in the Duct Mapper software. The software then initiates each measurement type in turn and stores the data in the figure. When all types of measurements have been performed the user can move to the next location and repeat the process. When the measurements are completed, the user selects the escape key to end the measurements. During a measurement run the user can erase the data taken at the last measurement point by right clicking or ctrl-clicking anywhere in the figure. Also during measurements, the user can turn on and off the display of the measured data using the `Display` menu which can be seen in Figure A.11. The bottom half of the menu allows each measurement type to be plotted as a colored point with an associated number. There is also a colored mist and contour plot of the measured points available for some of the measurement data. These sorts of plots will be covered in greater detail in Section A.5. However, we should note that contour plots only work for data taken in regularly spaced grids so contour plots of measured data will only work if the measurement locations are fixed into a grid before hand. Clicking on the locations for these points will not generally provide a regularly spaced grid so measurement points may need to be adjusted manually using the `»setappdata` command on the `»measurements` parameter. For example: `»getappdata(gcf, 'measurements')`.

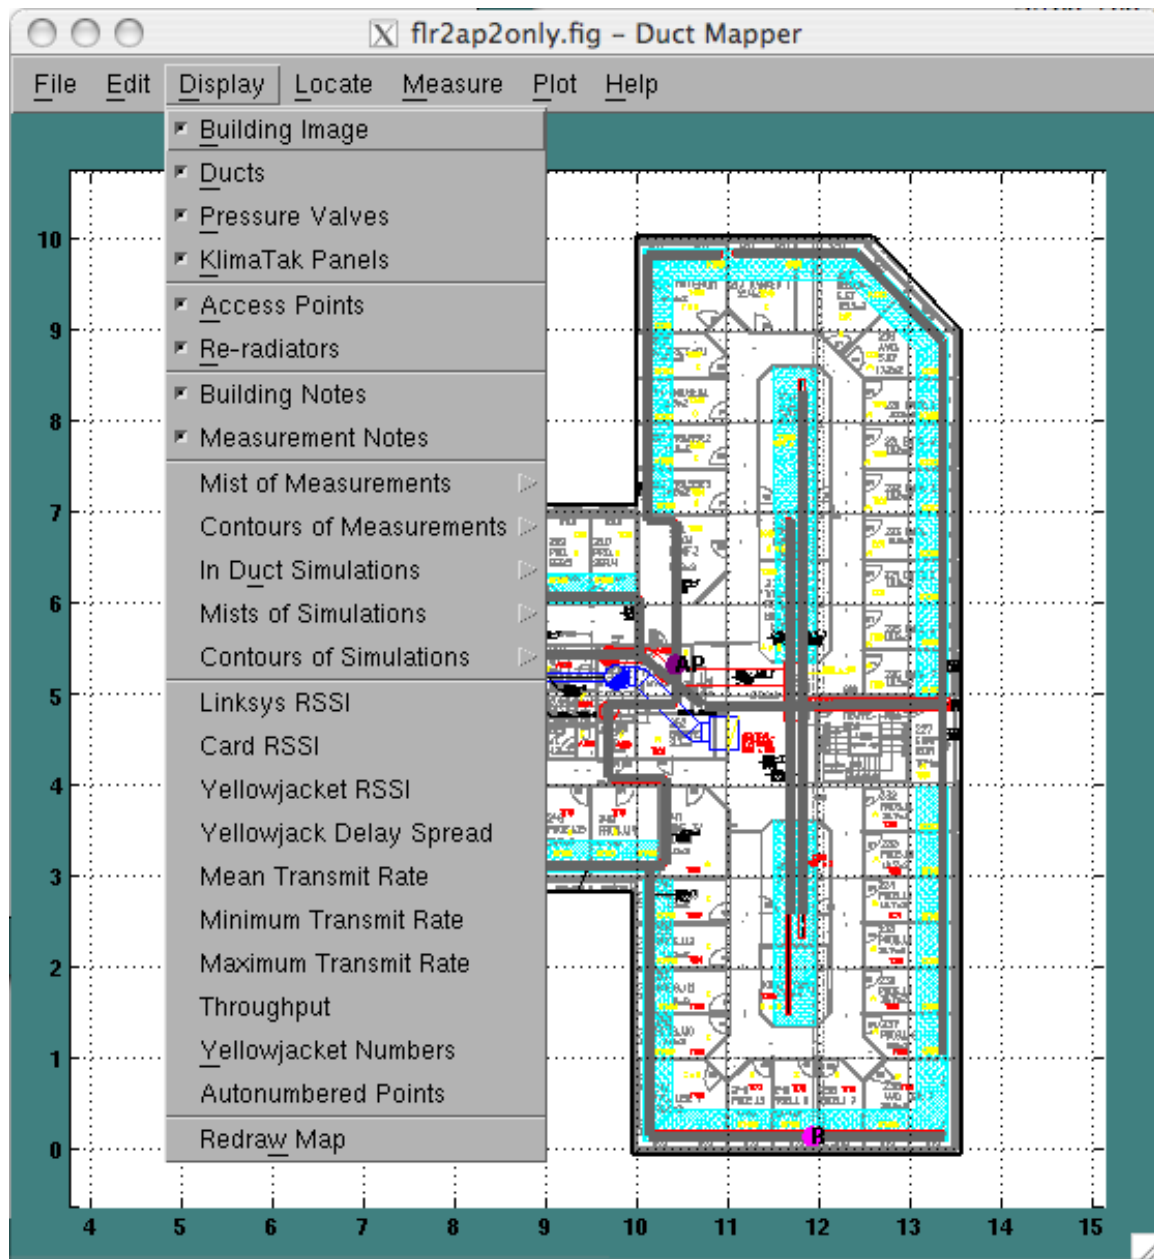


Figure A.11: An example building modeled in Duct Mapper.

If during measurements the user encounters a problem or the software crashes, the user may be able to recover some of the work using a backup file that is kept only during measurement runs. This file is the same as the name as the figure file with an added extension of “.meas.bkup” at the end. So a file “test-meas.fig” would have a backup file

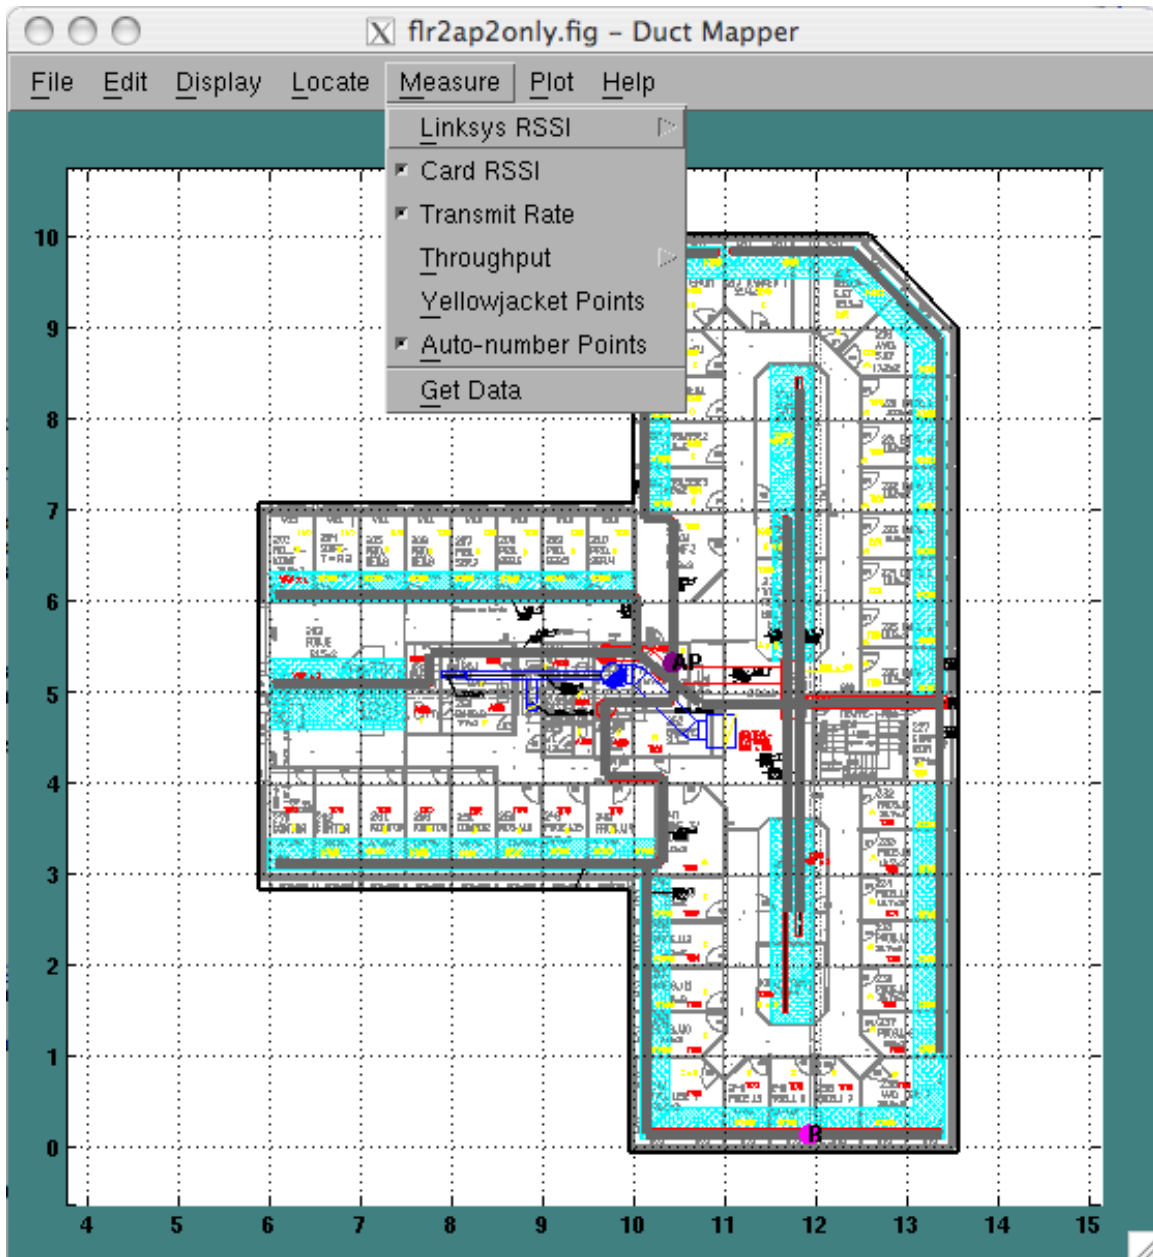


Figure A.12: Measurement menu in Duct Mapper software.

named "test-meas.fig.meas.bkup" in the same directory. To recover this file rename it to have a .fig extension and open it using Duct Mapper.

A.4.1 Auto-numbered Points

The most basic measurement type is auto-numbered points. With this measurement type the user clicks on a location in the building and a number is associated with that point. Each measurement point is of the form $x.y$. x is a counter of the current measurement set that has been initiated. y is a counter of the measurement location number. This mode is most useful for simply recording a location in the drawing that is of interest to the user or for saving locations that go with data that Duct Mapper can't directly record. To see the auto-numbered points the user should select the `Display, Autonumbered Points` menu item. An example display of this type of measurement point is shown in Figure A.13.

A.4.2 Wireless NIC Measurements

Wireless NIC measurements are somewhat unique in that the user must type in the measured value at each location because the Duct Mapper software cannot obtain it automatically. The Wireless NIC measurements include the card measured signal strength, and the minimum and maximum data transmission rates. These values are obtained through a software program provided with the particular wireless LAN card used for the measurements. The user reads the values off the screen and enters them manually into a dialog box as shown in Figure A.14. In addition to the data entry dialog box, the figure shows the point the user clicked on to identify the measurement location with a blue plus. Further, a previous measurement of the NIC measured signal strength is displayed as -83 (i.e. dBm) with a colored measurement point. It should be noted that the card RSSI entry is just a single value so the user must estimate an average value for the point. The transmit rate entry allows the user to watch the transmit rate change, preferably during a throughput measurement test and then record the maximum and minimum values in the dialog window. The software will then compute the mean of these two values and allow the minimum, maximum or mean value to be displayed using the `Display` menu.

A.4.3 AP Reported RSSI

A more automated method of measuring the pathloss from a location to an AP is to use the AP reported RSSI measurement technique that can be used with a Linksys WRT54G AP. The technique uses a security hole in the firmware of the AP to directly access the wireless driver software installed on the AP. This driver is capable of reporting the signal strength measured for the last received packet of a particular wireless client based on the MAC Address. This signal strength corresponds to an uplink signal strength.

The process of querying the AP for the signal strength takes about 3 seconds to complete. The measurement procedure then is to repeatedly query the AP for the signal strength for a specified duration. Care should be taken to ensure a new packet arrives from the specific wireless card to the AP between each measurement. If the user initiates the measurement using the wireless card to the AP this is taken care of automatically because the query process transmits packets that allow the RSSI measured by the AP to

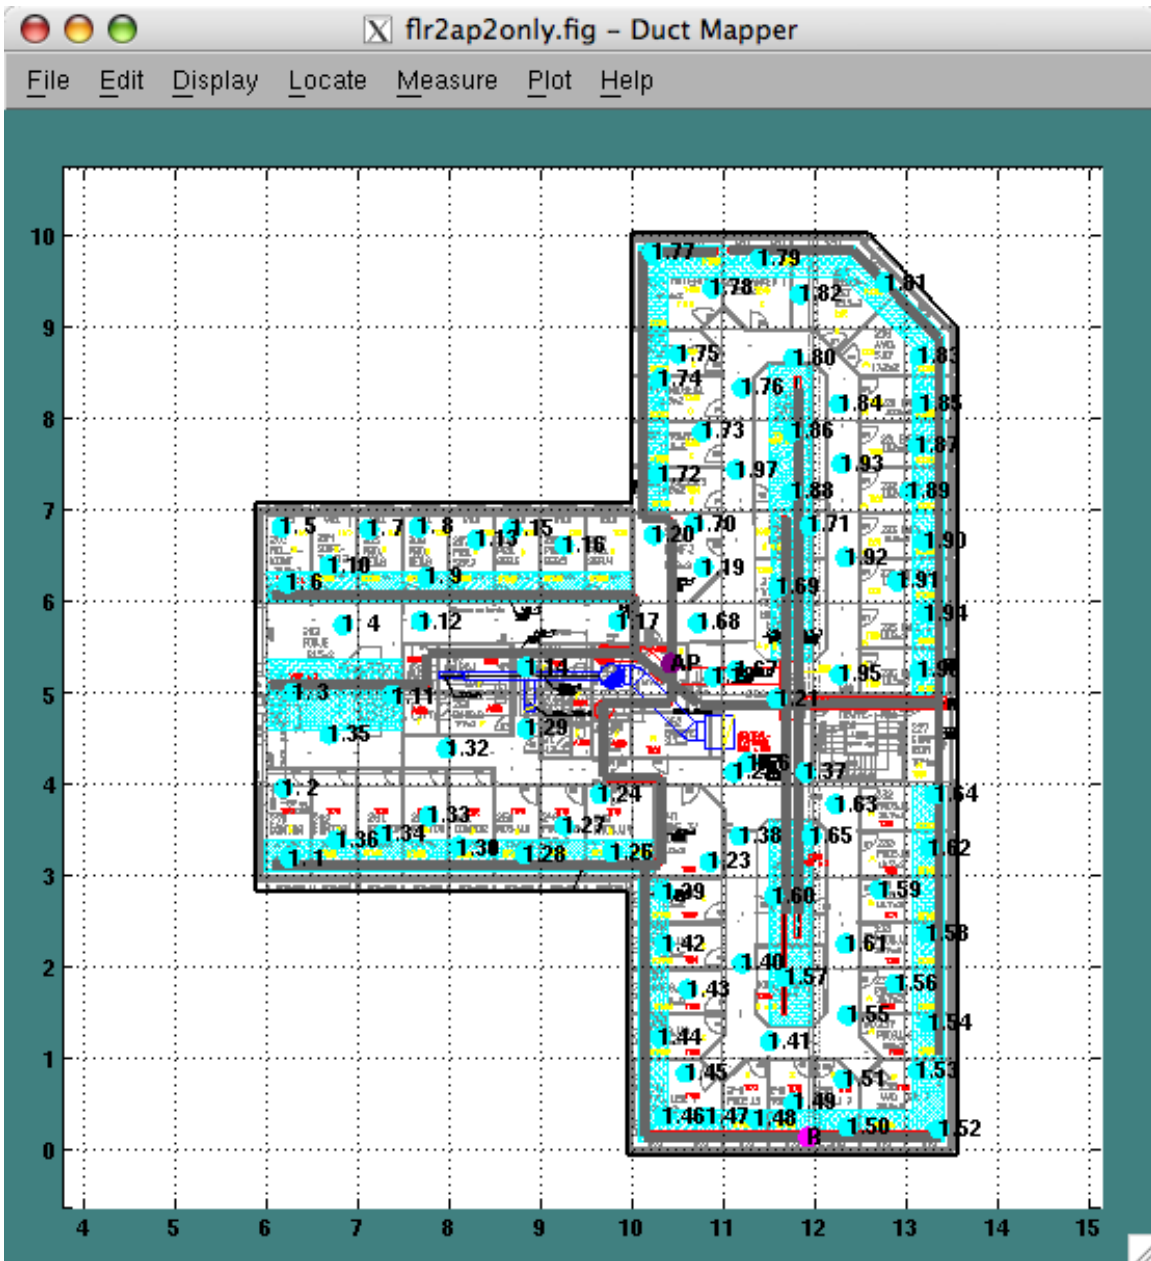


Figure A.13: Auto-numbered measurement points displayed in Duct Mapper.

be updated between each query to the AP.

The above measurement procedure is implemented as a perl script which is launched by the Duct Mapper program at each measurement point. The parameters that must be passed to this script are set using the Measure, Linksys RSSI submenu items shown in

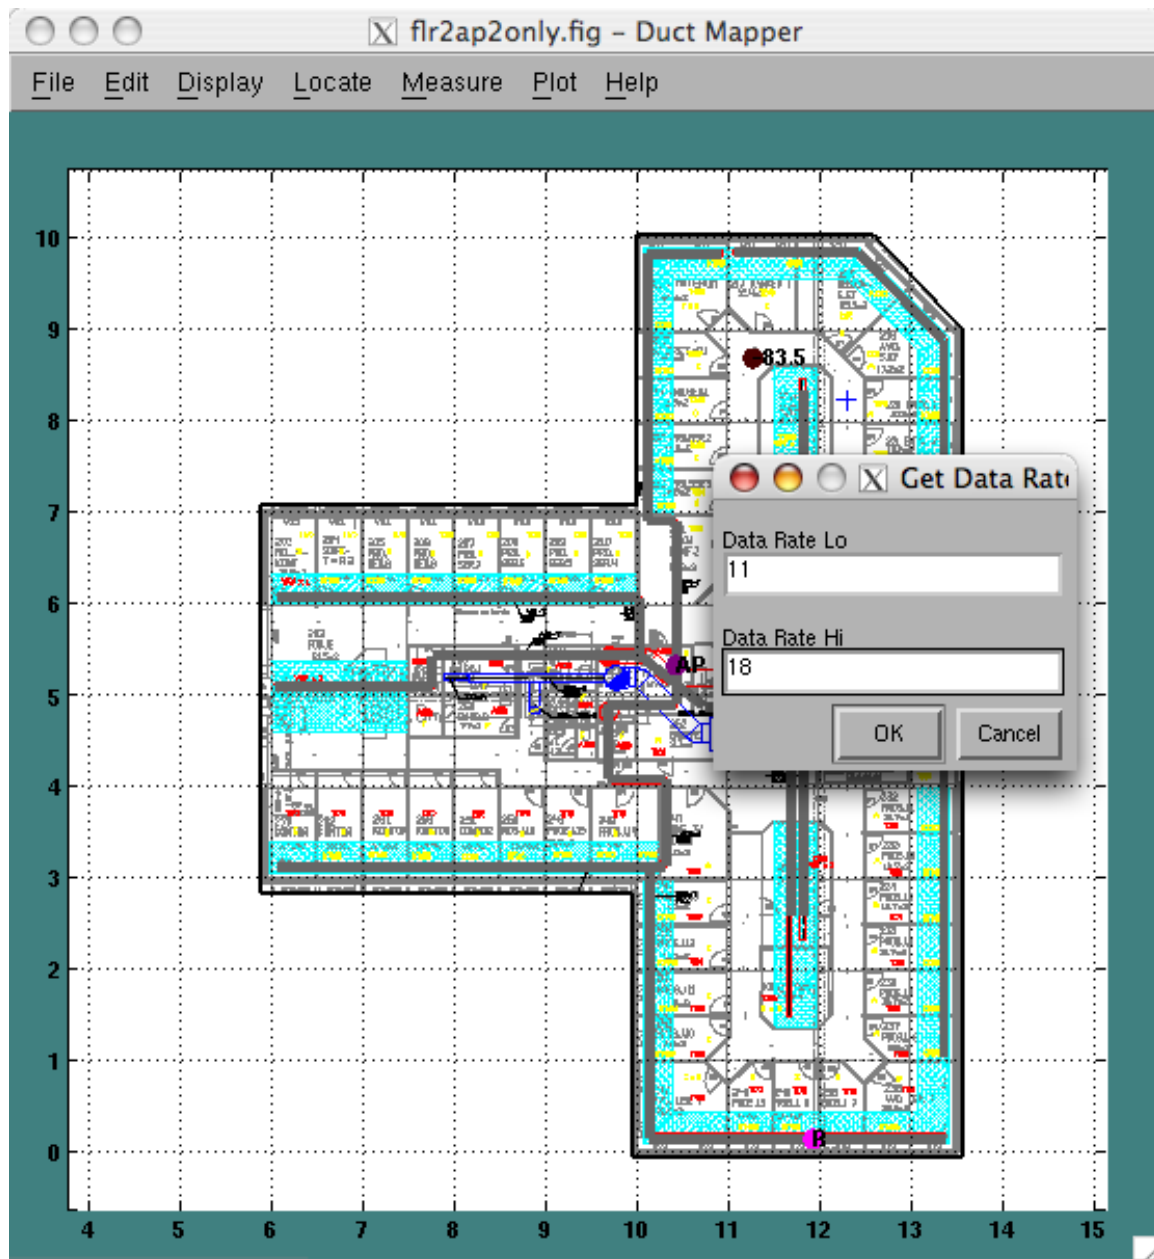


Figure A.14: Entering transmit rate data in Duct Mapper.

Figure A.15. The first submenu item allows the measurement technique to be enabled and a mark in the submenu item indicates whether or not the measurement type is enabled.

To access the correct AP, the user must enter the IP address of the AP using the Edit

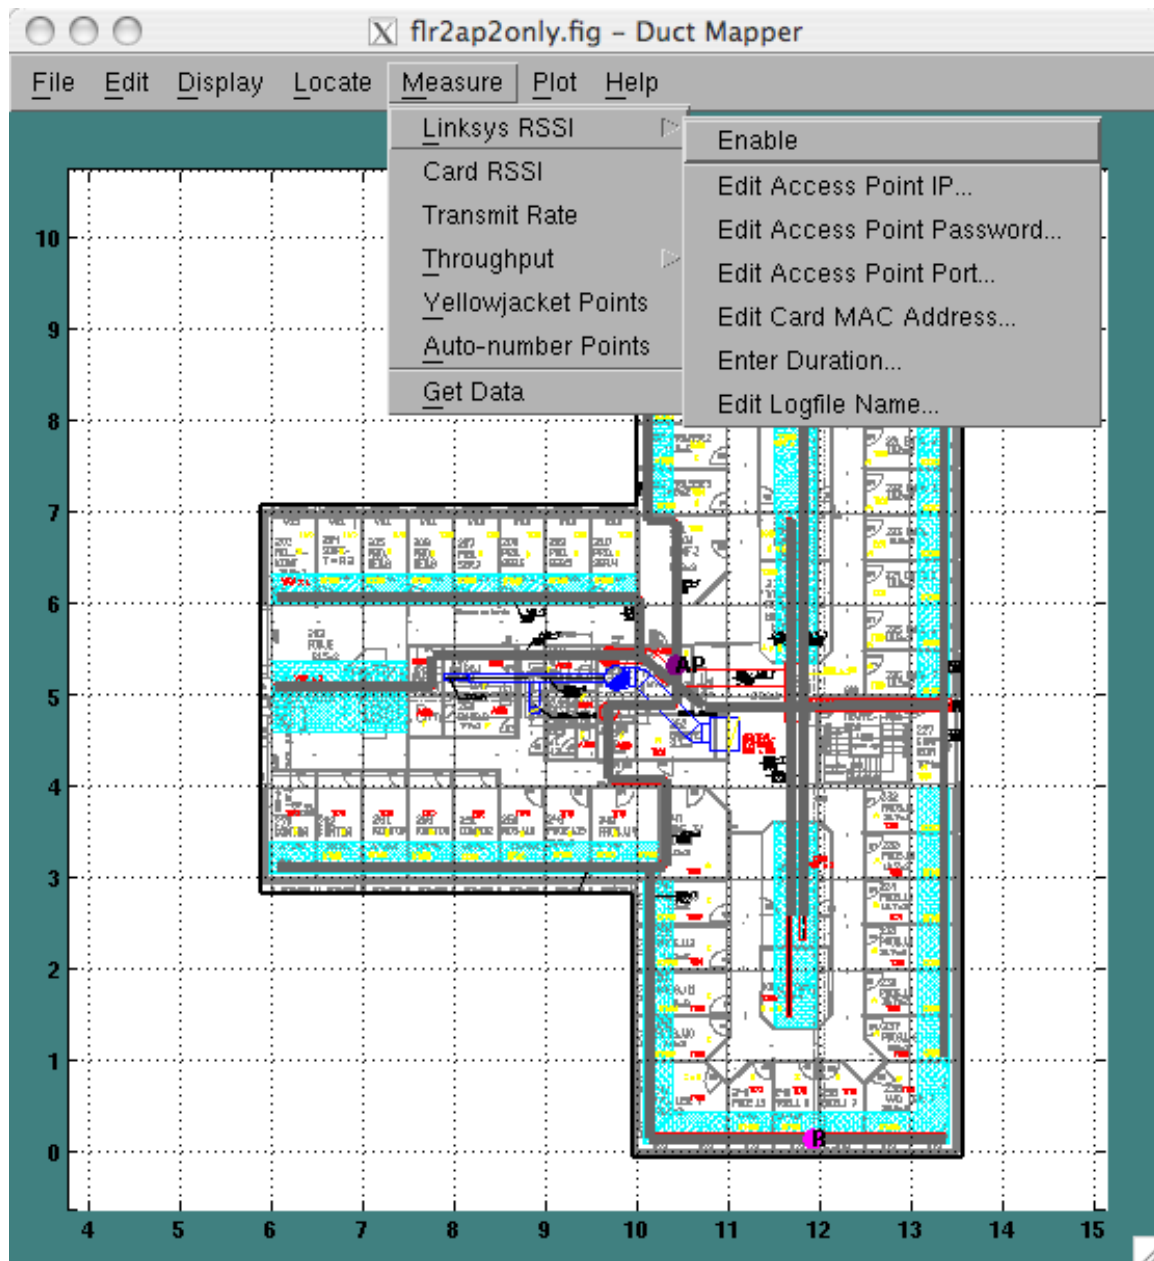


Figure A.15: Configuration menu items for AP based RSSI measurement using Duct Mapper.

Access Point IP... menu item. Additionally, the user must enter the Administrative Password and HTTP port for the Linksys APs management interface using the appropriate menu items. The default values for an WRT54G AP are a password of “admin” and

a port number of 80. It should be noted that the password is transmitted completely unencrypted during measurements and can be easily obtained via a packet sniffer. It is also necessary to enter the MAC Address of the NIC used to connect to the AP. The user should also enter a time in seconds over which to continuously probe the AP for the cards signal strength. Each query takes about 3 seconds so the number of measurements taken will be about the number of seconds entered divided by 3. The final RSSI value stored will be the linear average of the signal strength measured at each point. Lastly, a temporary file name to use for passing data to and from the measurement perl script should be entered using the `Edit Logfile Name...` menu item. This logfile can be used for data recovery or troubleshooting but is generally not needed.

A.4.4 Throughput Measurements

Throughput measurements are performed using the IPerf software program [28] that has been interfaced to the Duct Mapper software. This software uses a client and server setup to transfer data and measure the throughput of a link. Duct Mapper assumes that you will start the IPerf software in server mode manually on another computer. Then the user configures IPerf in client mode with various parameters set using the Duct Mapper menu commands shown in Figure A.16. The user first enters the IP address of the computer running the IPerf server using the `Edit Server IP...` menu item. The user also enters 4 parameters normally specified to the IPerf client software. These are whether to transmit TCP or UDP packets, the TCP window size or UDP packet size, the measurement duration and the update interval at which measured values are reported. These parameters are passed directly to IPerf. The last parameter is a file name to use for passing data to and from IPerf. This logfile is kept simply for backup and debug purposes. All of these parameters are set using the appropriate submenu item under the `Measure, Throughput` submenu.

To briefly explain the meaning of the throughput measurement parameters we describe a throughput measurement run. After clicking on the map of the building to indicate the measurement location, Duct Mapper will start a throughput measurement run. It will connect to the server at the specified location. If the data transmission technique selected was TCP, then TCP packets would be sent using the window size parameter as the socket buffer size. If the data transmission technique was UDP then the window size parameter is the receive buffer size. The measured throughput is displayed for the user in the MATLAB[®] command window area at a rate equal to the specified update interval. After the specified duration has passed, the sockets are closed and the throughput over the entire run is recorded and associated with the measurement location.

A.4.5 YellowJacket Measurements

YellowJacket measurements are handled by Duct Mapper differently from other measurements. Since the YellowJacket is a separate measurement device that cannot be directly

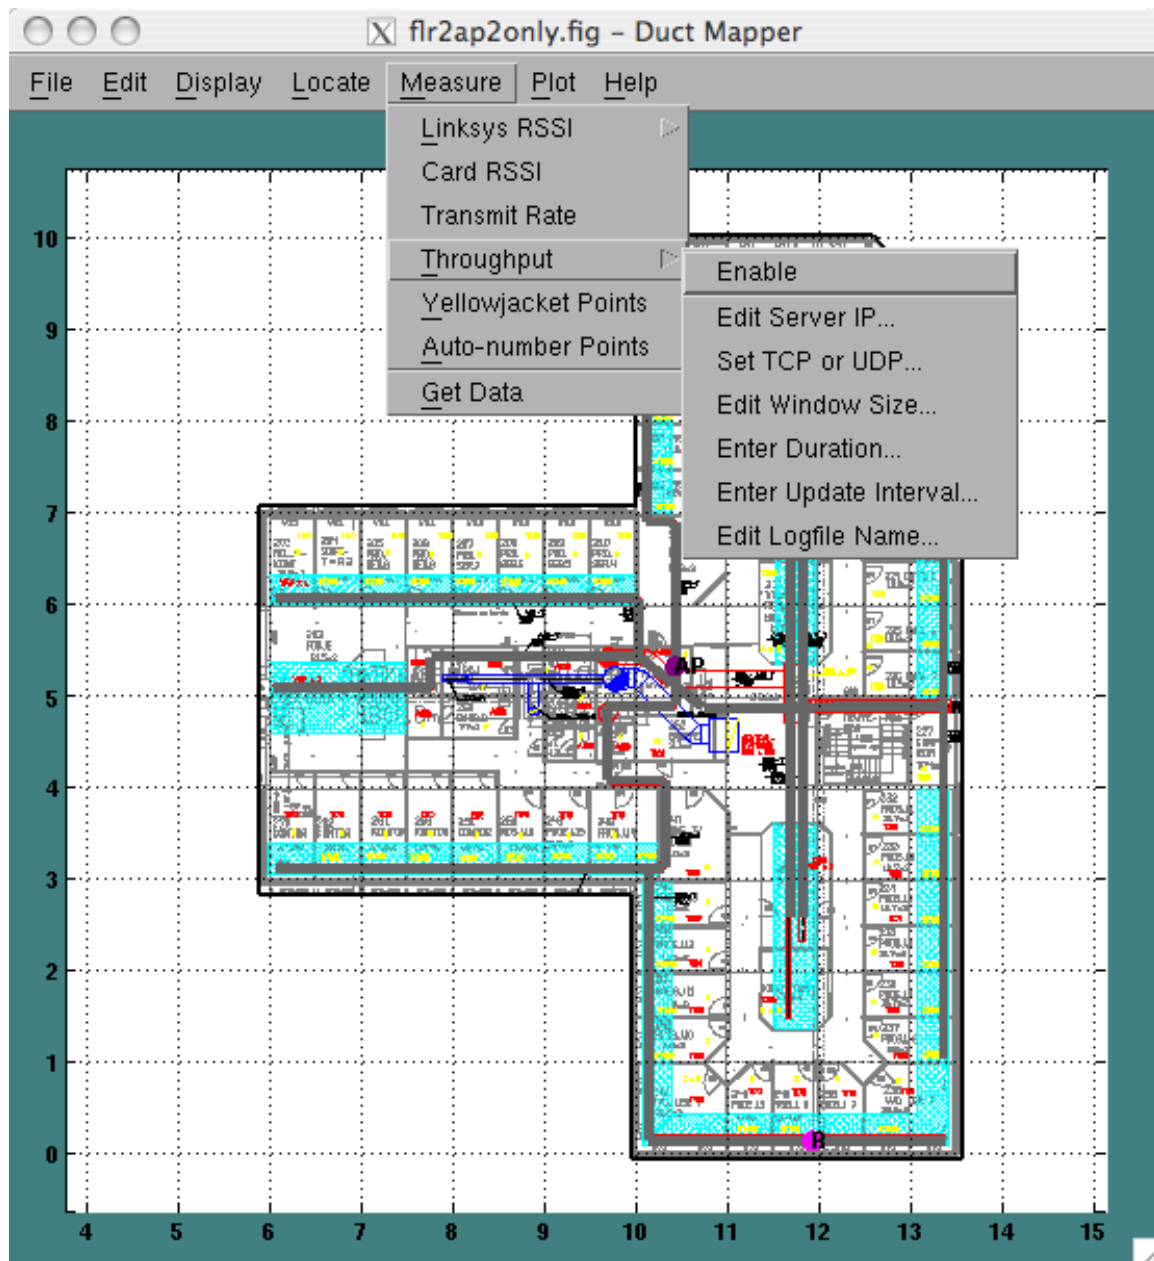


Figure A.16: Configuration menu items for throughput measurement using Duct Mapper.

interfaced to MATLAB®, Duct Mapper uses an import process to use the YellowJacket measurement data. This works by recording the YellowJacket markers at each location and later importing the YellowJacket logfiles. This process is a little strange because of the way YellowJacket markers work.

The basic procedure works as follows. After the user has enabled the recording of YellowJacket markers in the `Measure` menu, the Duct Mapper and YellowJacket measurements are started. The user moves to a particular location and then clicks on that location in the DuctMapper map. Then the user increments the marker number on the YellowJacket by clicking on the marker number. After incrementing, the displayed marker number will later be recorded. After waiting for the desired amount of time without changing location, the user again increments the marker number. The marker number from before is then entered into Duct Mapper in the provided dialog box as shown in Figure A.17. As can be seen, YellowJacket markers are plotted with yellow dots and are labeled with a number in the format $x.y$ where x is a measurement run counter number and y is the YellowJacket marker number entered by the user. In the figure, there is a measurement point, labeled 3.1, that is from a prior measurement run. It is important to make sure that you only ever associate 1 YellowJacket logfile with each Duct Mapper measurement run. However, 1 YellowJacket file may be associated with several Duct Mapper measurement runs. That is, if a new YellowJacket measurement logfile is created, a new Duct Mapper measurement run must be started. However, a new Duct Mapper measurement run does not need a new YellowJacket logfile. You may however wish to try and maintain a 1 to 1 correspondence between Duct Mapper measurements and YellowJacket logfiles to avoid confusion.

After completing the measurement run, the user will copy the logfile off the YellowJacket to the computer running Duct Mapper. The YellowJacket logfile should be converted from a `.yj3` binary file to a tab delimited text `.log` file using the Chameleon software provided with the YellowJacket. In Chameleon all fields should be exported, but order is not important. Then, the Duct Mapper is used to import the resulting `.log` file using the `File, Import, Yellowjacket Data...` menu item. This command will import the YellowJacket logfile. If the Duct Mapper file has more than one measurement data set with YellowJacket marker numbers, it will ask the user which measurement run to associate the YellowJacket data with as shown in Figure A.18.

After the user selects which measurement run to associate the YellowJacket logfile with, Duct Mapper will identify which markers in the logfile were recorded in the associated Duct Mapper measurement run. For each of these marker numbers, all data is averaged from the YellowJacket data and recorded as a single value for each measurement point. Then, after the data is imported the user can choose to display any of the recorded YellowJacket data, such as the RSSI or delay spread, as colored markers at the measured locations with the measurement value displayed as text.

A.5 Simulations

In Chapters 3 and 4 we develop our signal strength and throughput prediction models. As part of developing and calculating these models, we implemented two RSSI prediction models and one throughput prediction model in Duct Mapper. We found the simpler

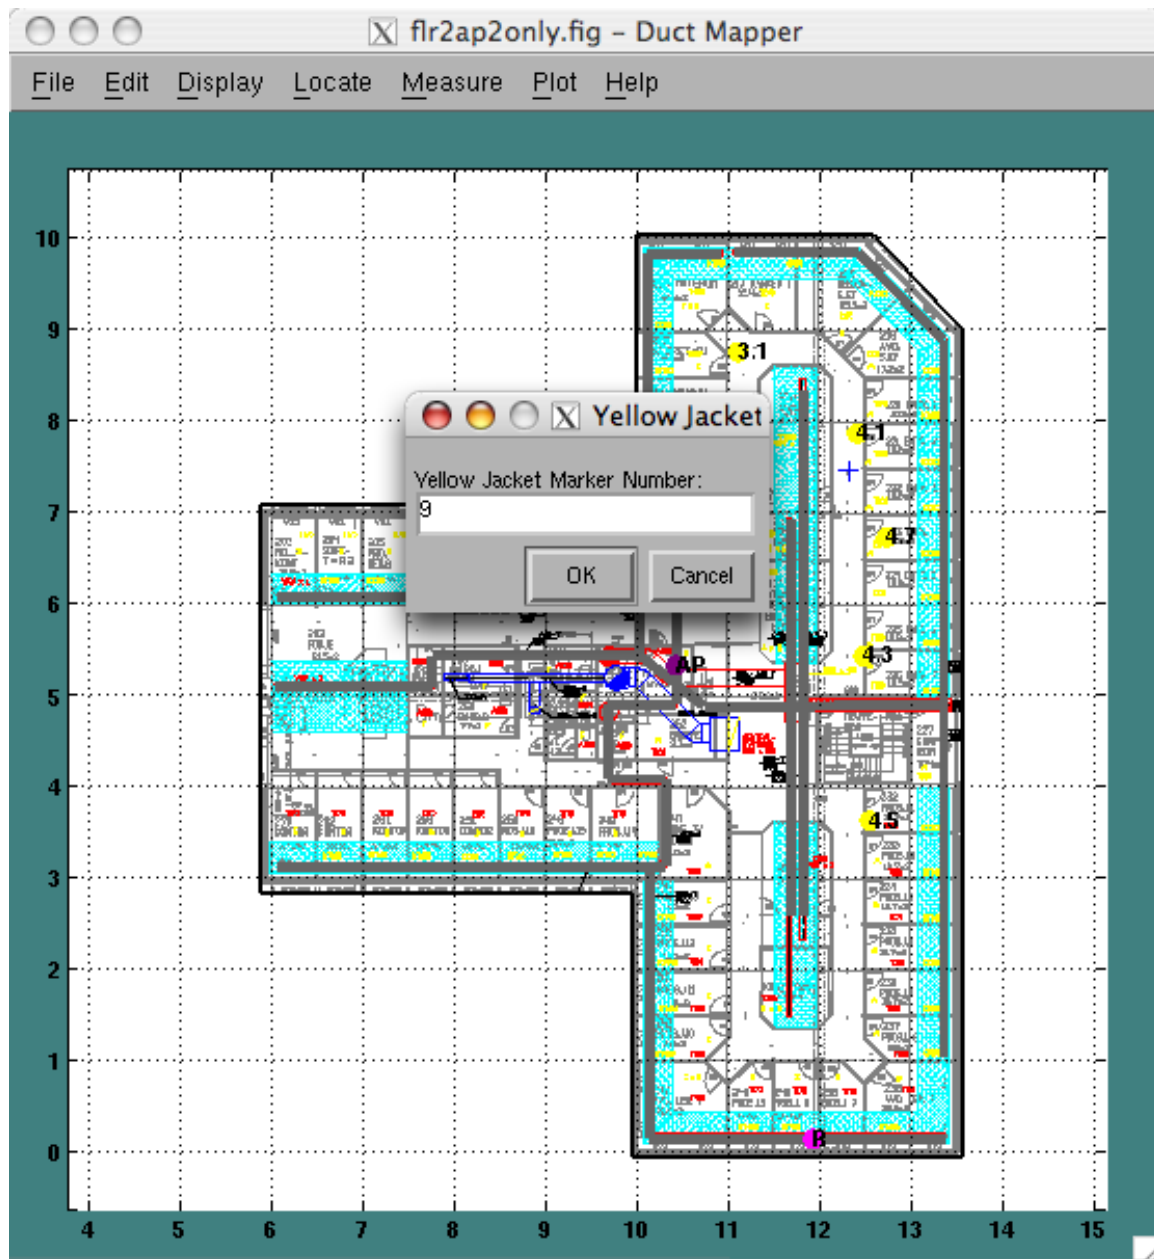


Figure A.17: YellowJacket measurements recording marker numbers.

of the RSSI models was more accurate for predicting propagation behavior so we only describe the first RSSI prediction model in Chapter 3. In this section we will explain the details of the implementations of the RSSI and throughput predictions and show some sample plots of the simulations.

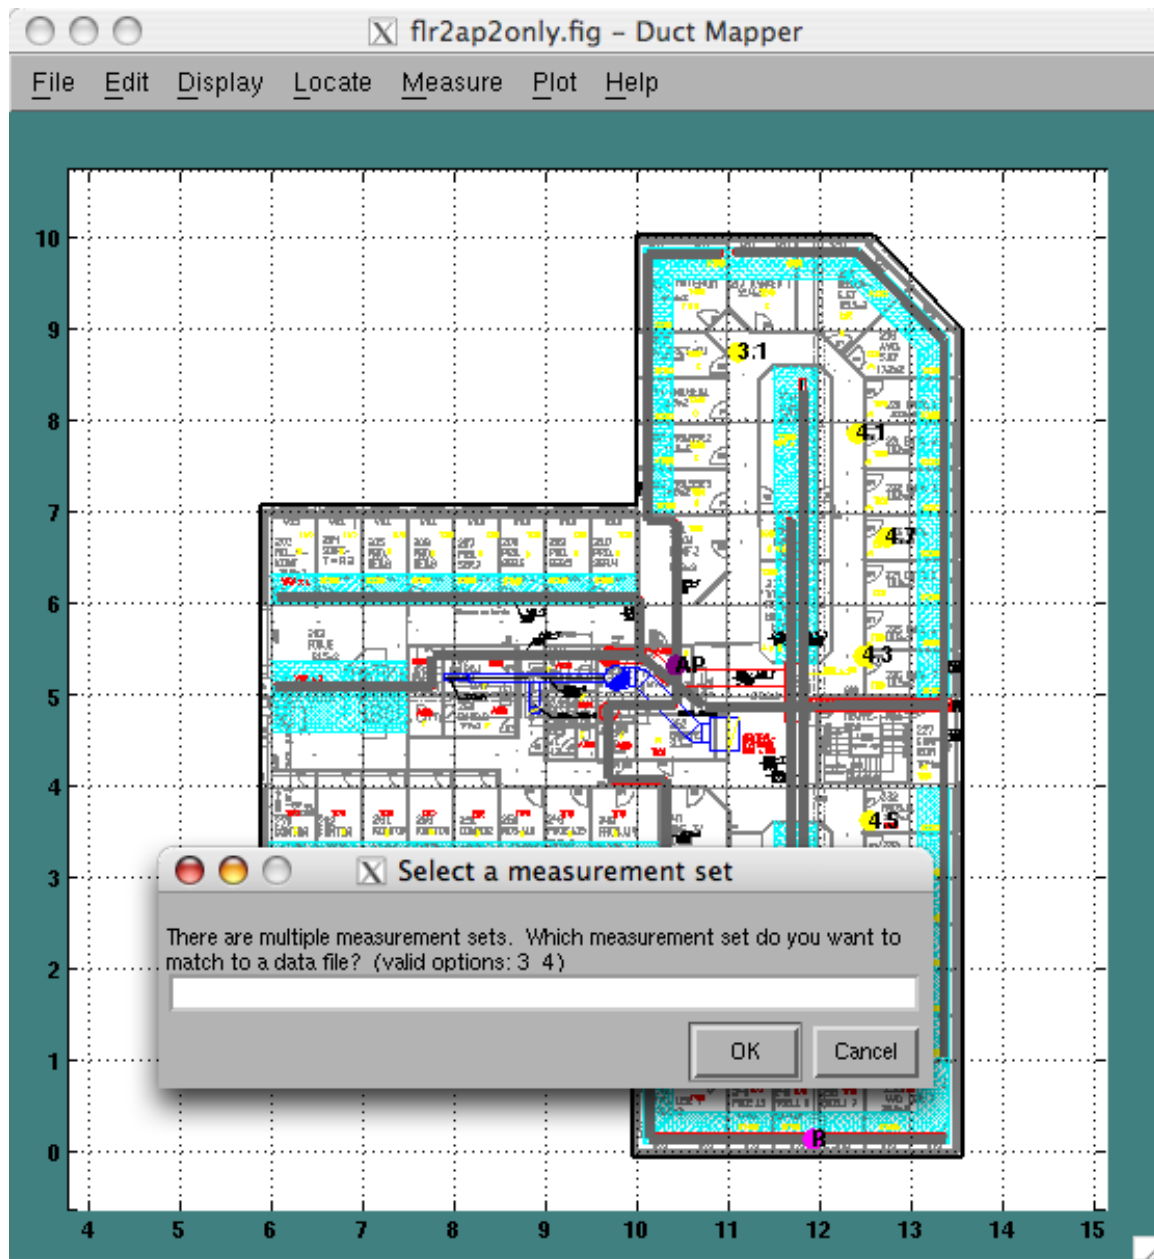


Figure A.18: YellowJacket measurements being imported to Duct Mapper.

A.5.1 Setting Model Parameters

The first step in performing a simulation is to input the simulation model parameters. This is done from the `Edit` menu shown in Figure A.19. As shown in the figure, there are two RSSI models listed as the Basic RSSI Model and the Junction RSSI Model

submenus. These submenus allow the user to select the simulation model to use and to set the parameters of the model.

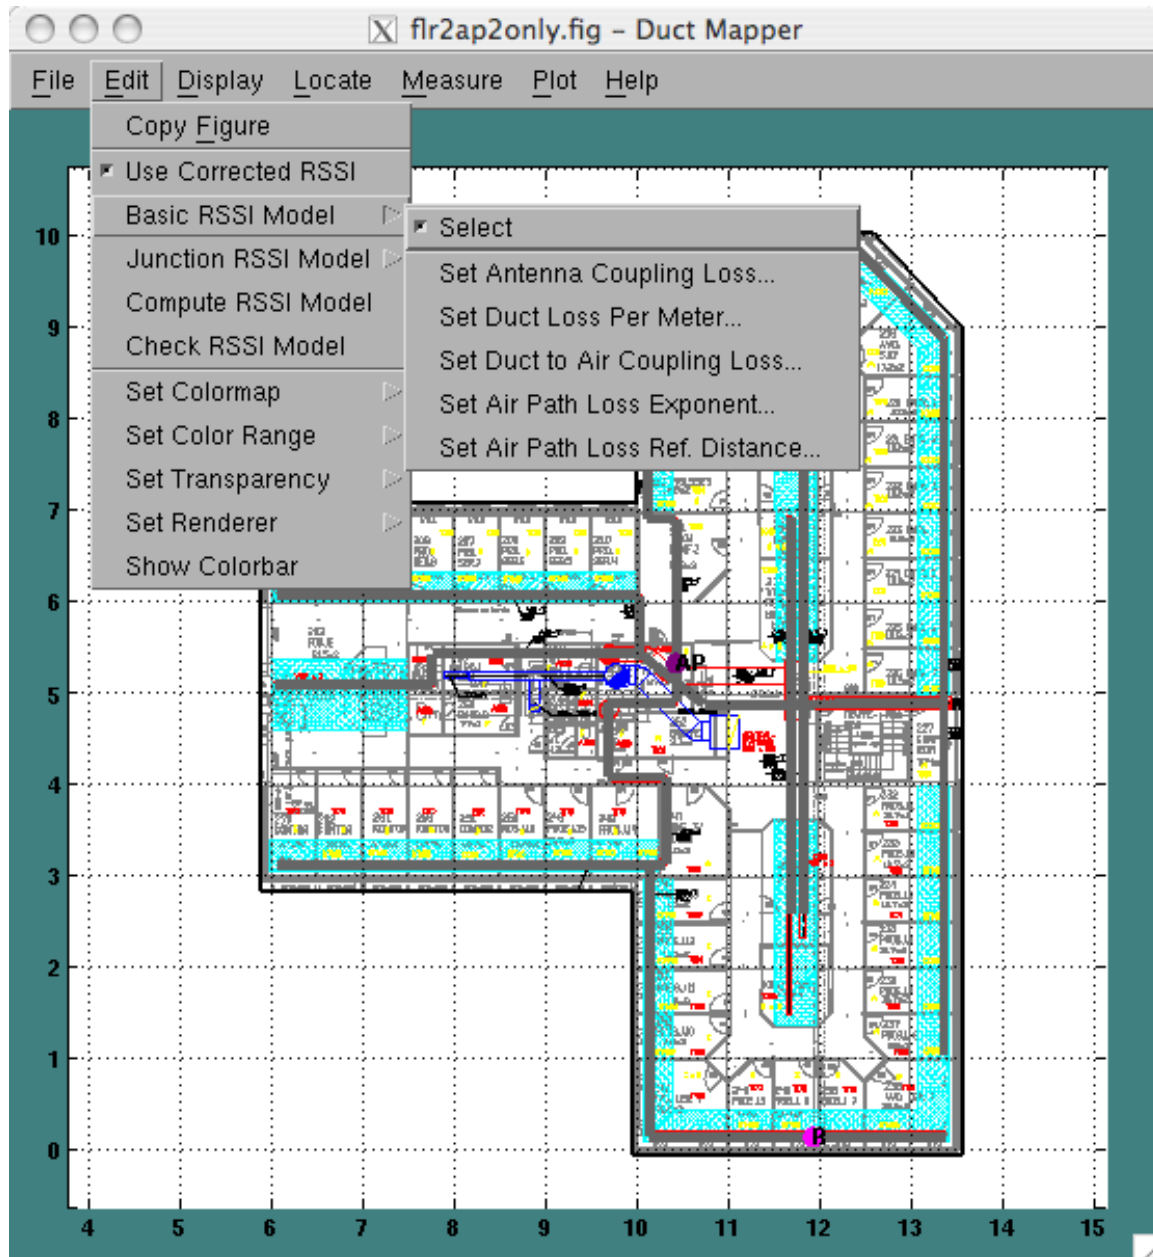


Figure A.19: Duct Mapper Edit menu used for setting simulation model parameters for the Basic RSSI Model.

The Basic RSSI Model is an implementation of the prediction model described in Chap-

ter 3. The Junction RSSI Model is nearly the same as the Basic RSSI Model except that it allows the user to specify the loss value at each junction type, rather than computing it based on a standard split loss. The `Edit, Junction RSSI Model` submenu is shown in Figure A.20. As shown, the model allows the user to set the same 5 parameters as the Basic RSSI Model and adds three parameters for overriding the Basic RSSI Models calculations of the junction loss values. It should be noted that both models assume that all ducts are at the same height and the receiver is 1 meter below the ventilation duct height.

Also shown in the `Edit` menu are two commands for comparing simulations to measurements. If the Duct Mapper detects that you have made measurements of signal strength values and stored them in the figure, then the simulation will calculate the signal strength values at all of the same locations as you have performed measurements. If there are no measurements, the simulation will use an automatically determined, regularly spaced grid of locations for simulations. Since the simulation calculates the signal strength at the same locations as the user has performed measurements, the simulation model can be directly compared to measurements. The `Check RSSI Model` menu command will plot the measured versus predicted signal strength values and display the mean and standard deviation of the error of the simulation model. Several examples of these plots were presented in Section 3.6, in figures such as Figure 3.17(b). In Section 3.6, we also note that we have performed a best fit of the model parameters to the measured data. Such a best fit using the MATLAB® command `»nlinfit` can be performed from Duct Mapper using the `Edit, Compute RSSI Model` command. This command will compute the best model parameters for the currently selected prediction model and store the parameters in the figure. After computing the best model parameters, Duct Mapper will also execute the `Check RSSI Model` command and plot the measured versus predicted values using the new model values.

A.5.2 Displaying Simulations

Simulation times in Duct Mapper are sufficiently fast that the user can simply turn on the display of a simulation and then let the software update the display as the user changes the locations of ducts or APs. Thus, the simulations are all enabled from the `Display` menu. In the upcoming sections we will describe the simulations that are available and show some examples. As we do so we will mention several of the commands in the `Edit` menu that are used for adjusting the colors and scaling of the plots.

In-Duct Simulations

The most basic simulation is to color the ventilation ducts based on the signal strength or throughput. For instance, if the user selects the `Display, In Duct Simulations, RSSI` menu command, then each time Duct Mapper redraws the display it will automatically recompute the signal strength at all locations inside the ventilation ducts if needed. Then

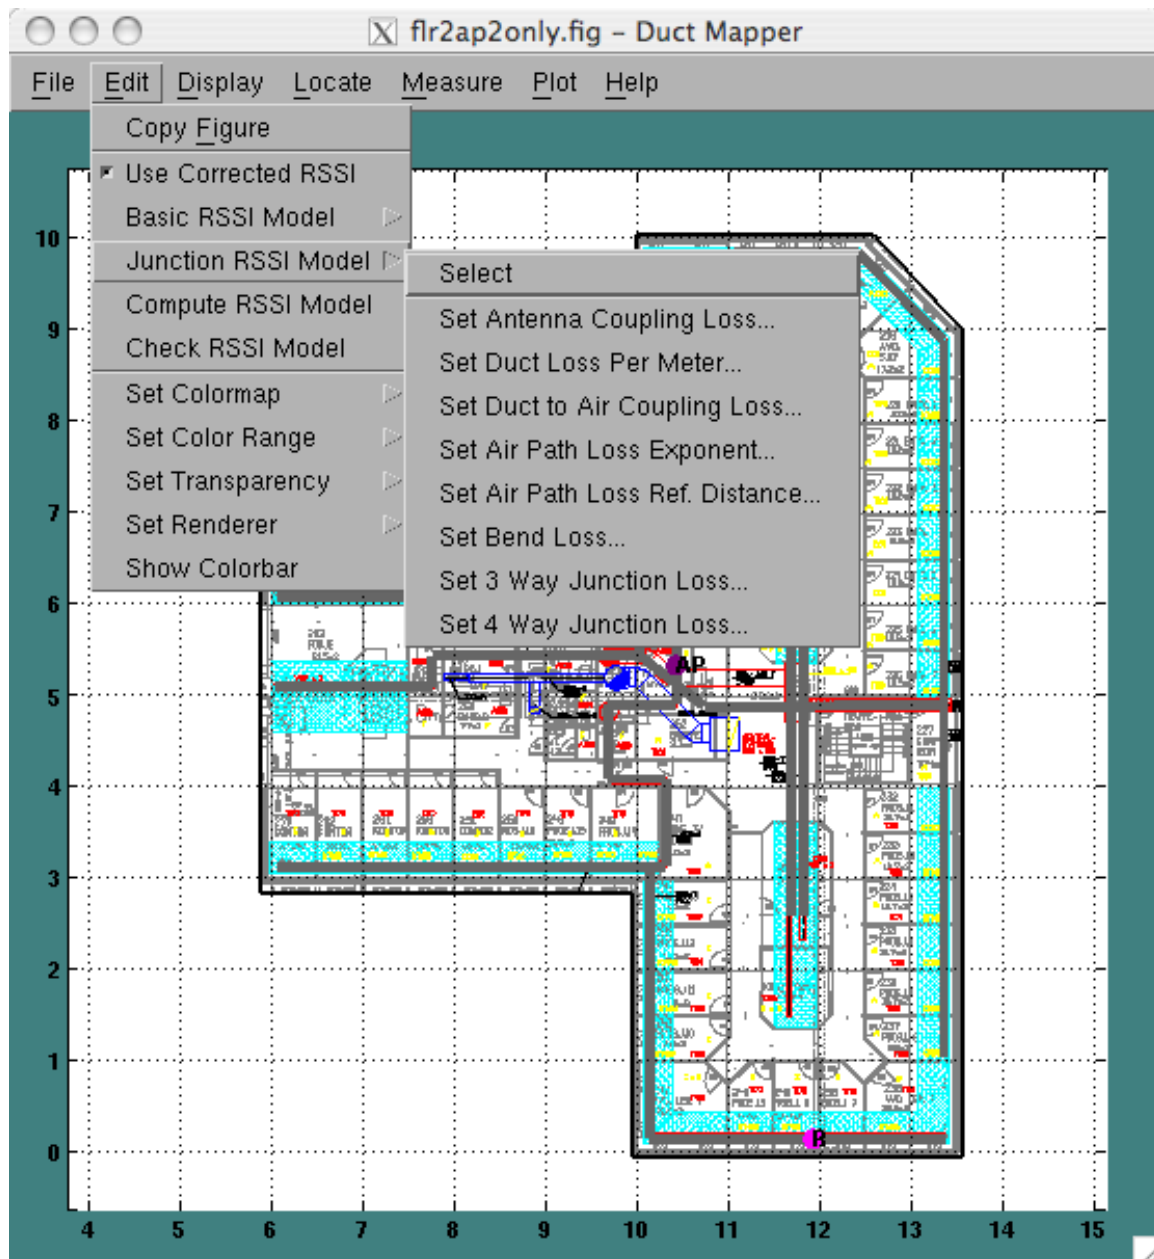


Figure A.20: Duct Mapper `Edit` menu used for setting simulation model parameters for the Junction RSSI Model.

Duct Mapper will draw colored lines scaled according to the signal strength level inside the duct. An example of this is shown in Figure A.21.

As can be seen in the figure, the duct network with an AP connected is colored while

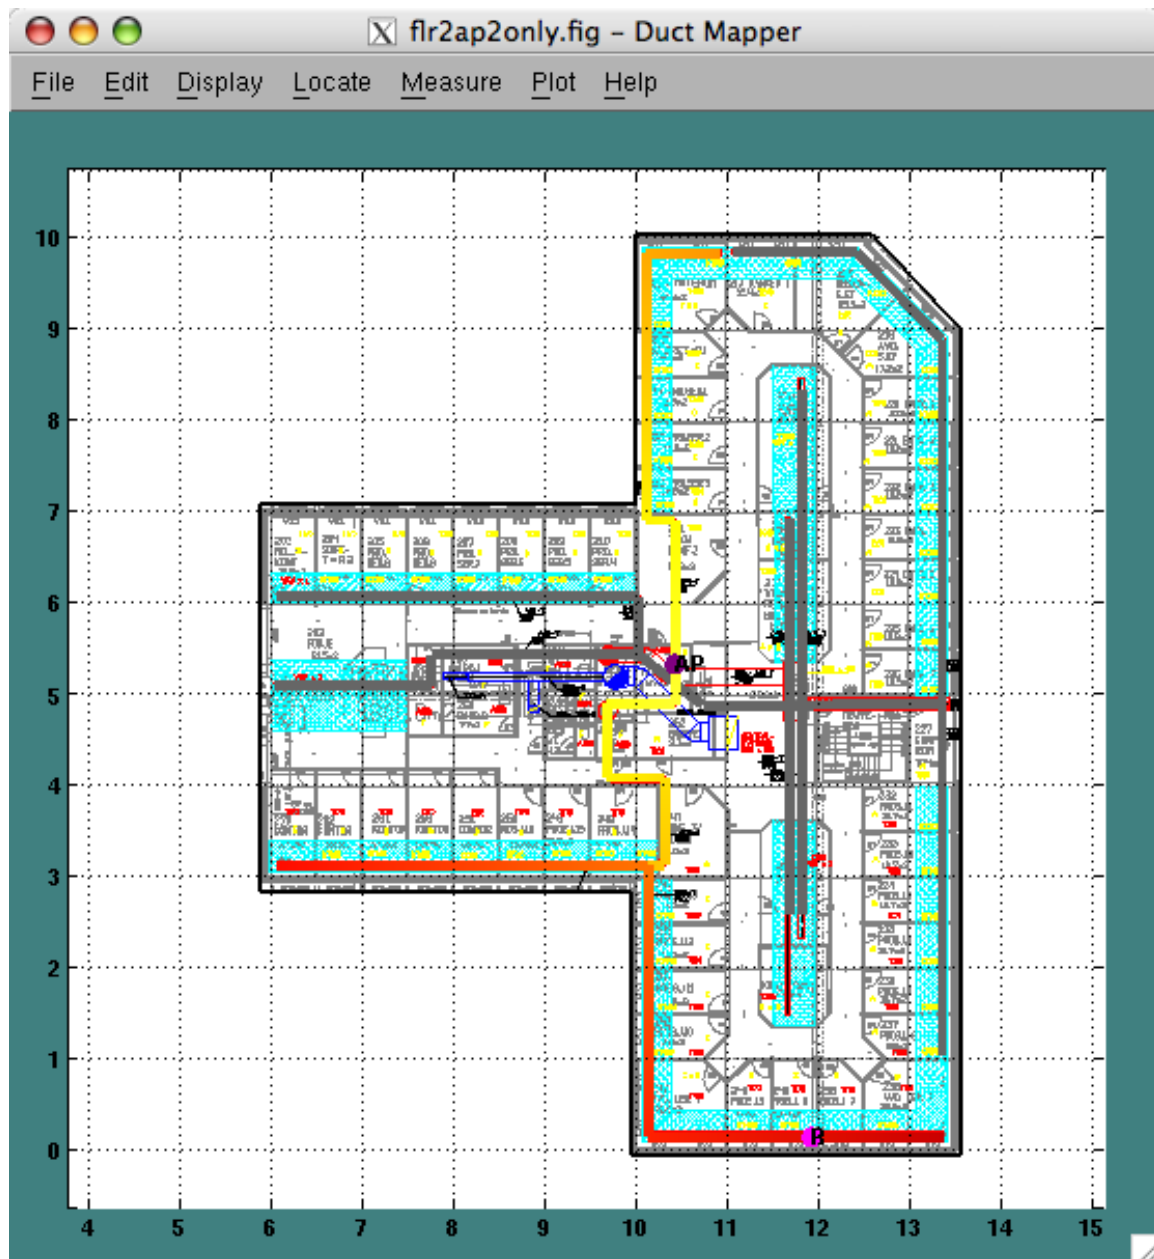


Figure A.21: Example of In Duct RSSI simulation.

the duct network without an AP is still the normal gray color. The coloring of the duct fades continuously from bright yellow to red as the signal strength decreases in value. The color scheme used can be set using the Edit, Set Colormap menu command or using the standard MATLAB[®] »colormap command. The colors used correspond to

particular signal strength values. Ordinarily, the signal strength in a duct changes slowly because of the low loss with distance associated with waveguides. However, in the figure we have set the colors to correspond to a narrow range of colors. To set the range the colors correspond to the user can use the `Edit, Set Color Range` submenu. This submenu has a few standard ranges of colors. There is no user menu for setting a custom range however, it can be done manually using the MATLAB® `»setappdata` command. For example, the color range in Figure A.21 was set using the command `»setappdata(gcf,'rssirng',[-10 15])`. After setting the range manually you will have to tell Duct Mapper to redraw the results using the new color range using the `Display, Redraw Map` command. To see how the colors map to a signal strength value you can turn on a MATLAB® colorbar using the `Edit, Show Colorbar` menu command. You should not use the MATLAB® `»colorbar` command as it will display an incorrect range.

The in-duct simulations are useful for checking that the user has drawn ventilation ducts correctly and that simulations are functioning as the user expects. The display can be a little confusing when more than one AP is present or if the ventilation ducts form closed loops. First, if more than one AP is present in the same duct network, the simulation tool calculates the signal strength at each location and then displays the maximum signal strength over all points.

Closed loops of ducts are more difficult for the software to accommodate. To calculate the signal strength in the ducts, Duct Mapper uses a recursive algorithm in which it starts at an AP, and computes the signal strength at each endpoint of the duct that AP is in. The algorithm then recurses to all of the ducts connected to those endpoints, calculating the power in the connected ducts. The recursion returns when there are no further ducts connected to the endpoint. Unfortunately, if a closed loop of ducts is present this algorithm will never complete. There is also no guarantee that the first path encountered by the algorithm to a duct is the one with the lowest signal strength. For this reason, we add the restriction that the algorithm will stop recursing to all connected ducts when it encounters a duct which has already been found to be closer in distance to the AP than would be calculated using the path encountered by the recursion. This means that the power in each duct should reflect the power from the least loss path back to the AP. It may, however, lead to some non-intuitive displays of signal strength in duct networks with closed loops.

Sometimes ducts can be laid out in some subtle way that leads to hard to understand or wrong simulations. The most common of these is when two ducts are drawn exactly on top of one another. To help identify these problems the user can access the array of duct endpoints using the MATLAB® command `»getappdata(gcf,'ducts')`. Also, another debugging tip is to have Duct Planner label each duct with its index number in the ducts array. To do this set the variable `DEBUG_NUMBER_DUCTS` to 1 in the `»drawmap.m` Duct

Mapper function.

Colored Mist Simulations

The colored mist simulation is one of the most useful simulations available in Duct Mapper. It can be enabled from the `Display, Mists of Simulations` submenu. Once enabled, Duct Mapper will display a semi-transparent colored overlay that represents the expected signal strength or throughput at locations in the building. The user can change how transparent the overlay is using the `Edit, Set Transparency` submenu. In this menu the user can use a predetermined value or enter a custom value. A value of 1 corresponds to completely solid and a value of 0 is completely transparent. An example of a colored mist plot of the predicted RSSI in a building is shown in Figure A.22. Note that the in-duct RSSI simulation is still enabled and that a colorbar is used to show the color range of signal strength values. A transparency value of 0.8 was used to generate the given plot.

Unfortunately, depending on the user's video card, MATLAB® often has difficulty rendering transparent overlays such as a colored mist simulation. The user may be able to overcome these issues by changing the display rendering method using the `Edit, Set Renderer` submenu.

Contour Plot Simulations

Contour plot simulations are the most useful for displaying expected throughput results. This is because users often think about throughput in terms of wanting at least a certain amount. Contour plots allow the user to visualize areas of the building that have at least a certain amount of throughput. This type of plot can be enabled using the `Display, Contours of Simulations` submenu. An example simulation of this type is shown in Figure A.23. As shown in the figure, the areas of the building with throughput above a 20 Mbps are highlighted with a yellow-orange color. This is indicated by the numerical labels on each contour and the colorbar in the right of the figure. When displaying such a plot it is recommended that the user turn off all transparency using the `Edit, Set Transparency, 1` menu command. This is because MATLAB® draws filled contour plots as overlapping contours so transparency will show underlying contours, not just the background image of the building blueprint.

Unfortunately, as of this writing, there is no direct user interface for specifying the parameters of the throughput prediction model. However, the model is implemented exactly as described in Chapter 4. You can set the parameters of this model manually from the MATLAB® command window using the `»setappdata` command. For example, you `»setappdata(gcf,'thrputmodelslope',1);` and `»setappdata(gcf,'thrputmodelintercept',80.3);` set the throughput simulation parameters for the current figure.




Figure A.22: A colored mist simulation of the predicted signal strength in a building.

A.5.3 Ideas for Simulating Transmit Rate and Delay Spread

It should be noted that Duct Mapper only implements signal strength and throughput simulations at this time. While there are menu items that indicate the software can predict the data transmission rate and delay spread, these functions are not implemented.




Figure A.23: A contour simulation of the predicted throughput in a building.

While these functions might be interesting, it is felt that the signal strength and throughput values are the most useful. However, it is likely that the data transmission rate could be predicted in much the same way that throughput is by establishing a scaling factor from throughput to the transmission rate.

A more difficult value to predict is the delay spread. However, the recursive algorithm for calculating the signal strength inside ventilation ducts described in Section A.5.2 could be modified to attempt this. The recursive nature of the algorithm is such that it can simultaneously compute the losses and delays through all paths through a reverberant duct environment. A slight modification would be needed to quantify the loss at reflections off of endcaps and possibly introduce new reflections at duct junctions, but this would be a relatively minor modification. The recursive algorithm would then need to be modified to continue to recurse down paths until a maximum allowable pathloss or path length is reached. This algorithm is similar to an in-building ray-tracing type algorithm except the number of paths needed to be considered is significantly reduced from an in-building simulation because of the limited scattering possible in a ventilation duct.

A.5.4 Using Simulation Plots with Measurements

It is important to note that Duct Mapper has the ability to display measurements using the same colored mist plot as used for simulations. This allows the user to view a colored mist interpolated across the measured values. To turn on such an overlay the user selects a command from the `Display, Mists of Measurements` submenu. Duct Mapper also has a submenu for displaying contour plots of measurements. However, contour plots require that the measurements all be on a regularly spaced grid of locations so rarely will a contour plot of measurements be feasible.

Another technique for displaying measurements and simulation plots is to display measurement points simultaneously with a colored mist plot. This allows the user to look for spatial variations in the expected versus measured values. It is recommended that any transparency be disabled in such a case to ensure colors match up between the measurements and the simulations.

A.6 Data Analysis Plots

Throughout the development of Duct Mapper, certain plots of data were used over and over to analyze measurements. Since they were used so often, a `Plot` menu was added to allow the user to initiate plots directly. Three basic plot types are available as shown in Figure A.24. These three plot types will be presented in the upcoming sections. However, we also note that the `Plot` menu also contains a command for computing the area inside the perimeter of the building which is initiated using the `Plot, Compute Building Area` menu command. This command will display a dialog box with the building area in square meters displayed.

A.6.1 Throughput versus RSSI and Delay Spread Plots

In Section 4.2, we discussed the correlation between signal strength and throughput. This was a topic of interest in many different measurement sets. Hence, the ability to plot the measured throughput versus measured signal strength or measured delay

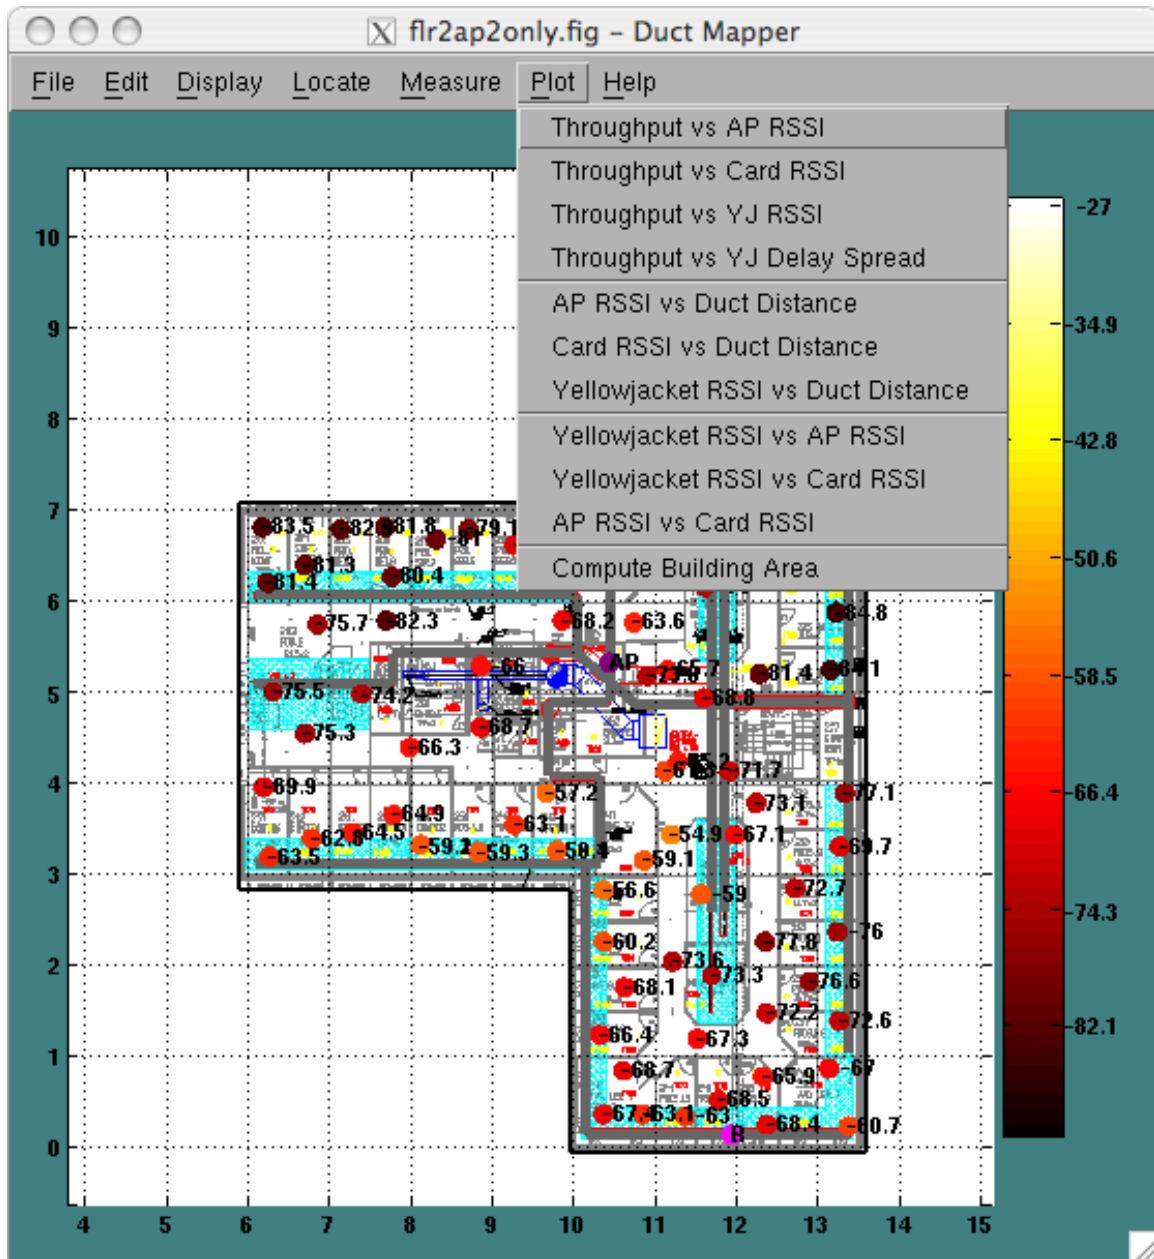


Figure A.24: Plot menu in Duct Mapper.

spread was useful. The plotting functions used for this purpose would simply combine all measurement sets that included throughput and a particular type of signal strength or delay spread measurements and plot the results. An example of such a plot is shown in Figure A.25.




Figure A.25: Example plot of throughput versus RSSI.

A.6.2 RSSI versus Duct Distance Plots

The RSSI versus duct distance plot was used in Section 3.6.2 to compute model parameters for a RSSI prediction model. This plot is generated using a set of RSSI measurements with a modeled duct system and AP that match the setup of the measurements. For each measured RSSI, Duct Mapper finds the nearest duct to the measurement point. It then computes the shortest distance from that point in the duct to the nearest AP. Duct Mapper then plots the measured signal strength versus the distance through the duct for all points. An example of this is shown in Figure A.26.

A.6.3 RSSI versus RSSI Plots

The third plotting type added to the Duct Mapper software is a plot that compares different measurements of signal strength at the same locations. This plotting technique




Figure A.26: Example plot of measured RSSI versus distance through nearest duct to best AP.

concatenates all of the points in a Duct Mapper figure that have the two desired types of signal strength measurements. It then plots each of these points as shown for example in Figure A.27.

In Figure A.27 are two computed best fits. The first is a simple linear fit. The second is a linear fit with the slope force to be 1. The reason for these fits is that it was discovered that different measurement devices tend to measure different RSSI values. By averaging several sets of measurements taken with the same devices over several different sets we could use these plots to compute correction factors for different measurement devices. These correction factors are stored in Duct Mapper and are applied automatically when the Edit, Use Corrected RSSI menu item is enabled.




Figure A.27: Example plot of measured wireless NIC measured RSSI (cardrssi) versus AP measured RSSI (aprssi).

Load Balancing Software

B.1 Software Functions

A screen capture of the software as it starts up is shown in Figure B.1. The software is designed to show the status of 2 access points and has the IP addresses 10.0.0.1 and 10.0.0.2 as well as the password “wirelessducts” hardcoded into it. When the software is started, it will disable the MAC Blocking functionality of the two Access Points but it will not update the lists of Connected Users, RSSIs or Blocked Users. To update these lists you must click the `Update` button.

To update the display with the currently connected clients and signal strengths as well as any MAC addresses in the block lists of the APs, click the `Update` button. This will initiate a series of commands to read the user MAC addresses and RSSI values from the two APs. The results after doing this in a test scenario are shown in Figure B.2.

As can be seen in Figure B.2, 2 clients are connected to Access Point number 1 with low signal strengths. Repeated clicks of the `Update` button will show any changes as users move around the environment. RSSI values will update (assuming the user is sending data to allow the AP to measure the RSSI value of the data packets) and if the user moves out of range of AP 1 and into range of AP 2, it will disappear from the AP 1 user list and appear in the AP 2 user list.

Below the connected users list are three buttons. The `Block` and `Unblock` buttons will add or remove a MAC address from the Blocked Users list for that AP. To choose which MAC address to block, select the MAC address from the Connected Users list and click on the `Block` button. To Unblock a user, select the MAC address from the Blocked Users list and click the `Unblock` button. Adding or removing users to the block list will do nothing unless the Blocking function is enabled by toggling the `Enable Blocking` button. This button will remain depressed when Blocking is enabled.

Adding a user to the Blocked User list won't prevent a user from continuing to use an AP that it is already connected to. The block list is only checked when a user first associates with that AP. If the user is already associated, it can continue to use the AP. To kick the user off of the AP you need to use the `Disassociate` button. This will disassociate the selected MAC address from the AP. The card will then scan all wireless channels and attempt to associate to the AP with the highest RSSI that isn't blocking the

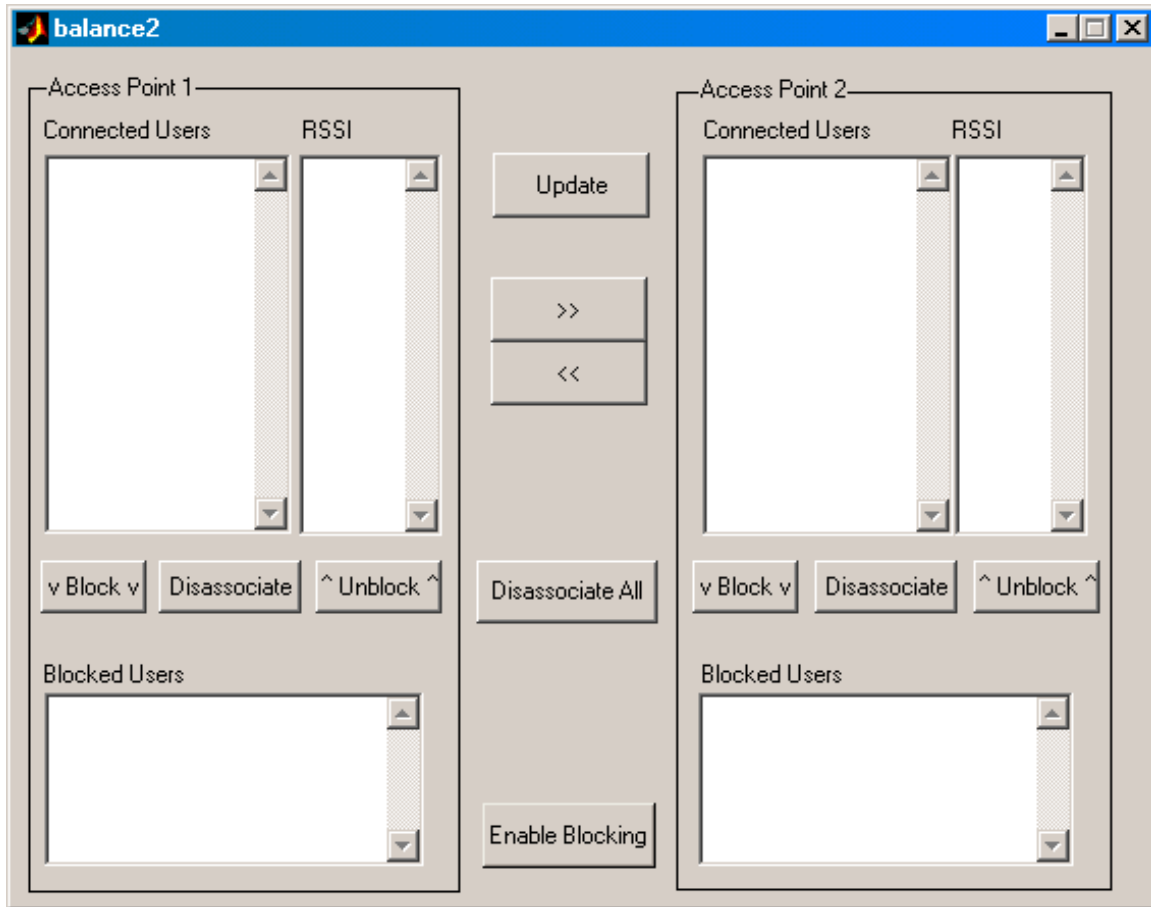


Figure B.1: Load Balancer immediately after startup.

user's MAC address.

The » and « buttons allow a user to be moved from one Access Point to another. If for example you clicked the » button, the Load Balancer software would move the user selected in the AP 1 Connected User list to AP 2. This is done by Unblocking the MAC address from AP 2, Blocking the user on AP 1 and then Disassociating the user from AP 1. Figure B.3 shows the software after exactly this was done.

It's important to note that there is no guarantee that a user, once forced off of AP 1 will associate with AP 2. In this example, the only other AP in range for the user was AP 2, so the kicked off user associated to AP 2. In actuality, the user will associate with the AP not blocking the user having the highest RSSI. This is actually the desired behavior because if AP 1 were overloaded, by forcing a user to go elsewhere, the card will then connect to the AP with the next best signal strength. Thus by letting the card choose what AP to choose the card will maximize its throughput. Of course, the user can be restricted from

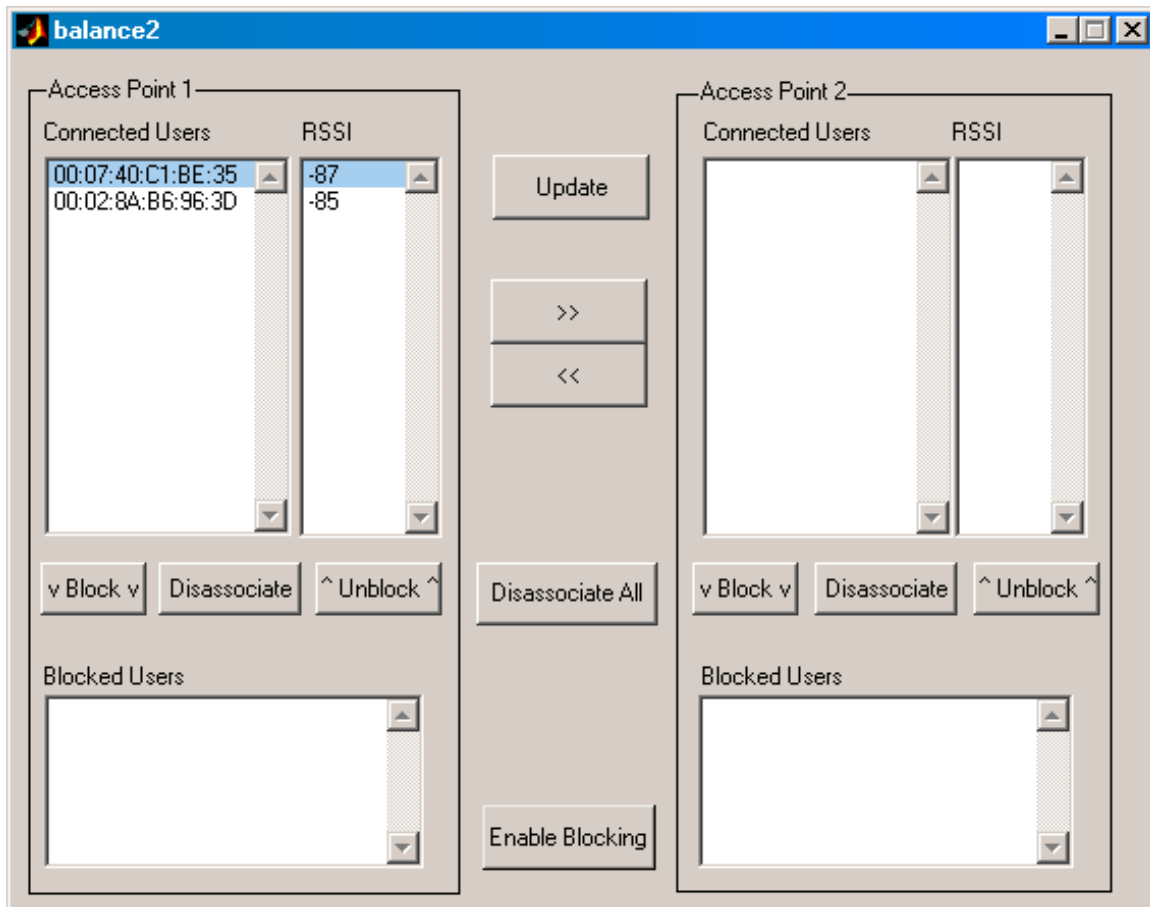


Figure B.2: Load Balancer after clicking the `Update` button to get a list of connected users and RSSI values.

connecting to other APs or forced to connect to a specific AP by adding its MAC address to the block list of other undesired APs prior to forcing the client to disassociate.

Lastly, note that there is also a `Disassociate All` button. This button will disassociate all clients from both APs. This function is useful for ensuring that all clients are connected to the highest signal strength AP available because by forcing a user to disassociate the user's card will scan for the best signal strength and associate to that AP.

B.2 Software Design

The load balancing software was created to illustrate issues with that wireless LAN systems experience under high load conditions. The software provides a basic interface to Linksys WRT54G access points with firmware versions 2.04.4 though it should work with firmware versions 1.30 - 2.04.4. A basic diagram of the software is shown in Figure B.4. As shown, the MATLAB[®] script `balance2.m` is simply a user interface to the perl script

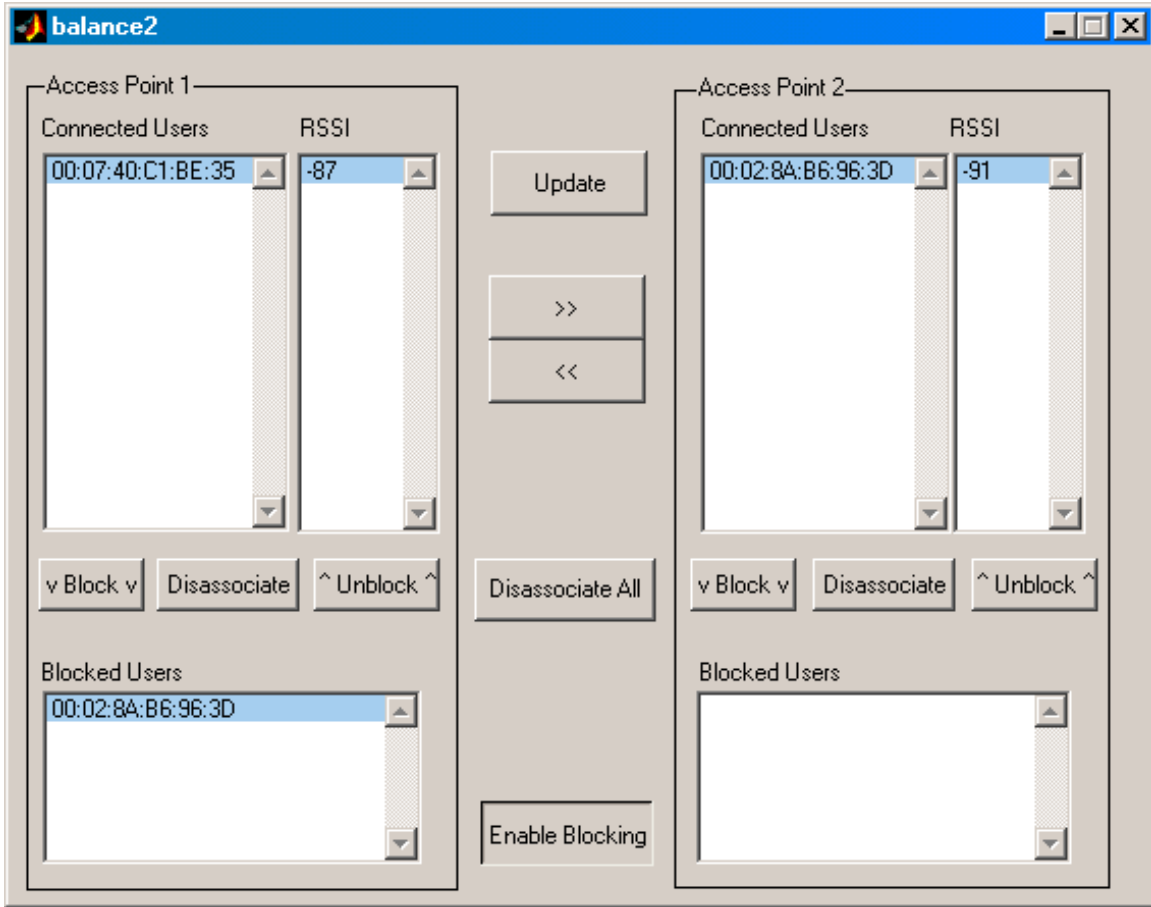


Figure B.3: Load Balancer after forcing a user to move from AP 1 to AP 2.

wrt54gcmd.pl The perl script in turn uses a validation exploit of the webpage interface of the WRT54G access point to access the wireless driver utility on the WRT54G, w1, which is located in the /usr/sbin directory on the access point. Most Access Points do not allow users to access the low level functions that can be accessed by the w1 script so care should be taken in accessing the script.

The Load Balancer software uses the w1 script to add and remove users from the wireless MAC filter lists, get lists of associated users and their signal strengths, and to force users to dissociate from an AP. These combination of these functions are used to demonstrate how load balancing could occur.

B.3 Speed Improvements

The Load Balancer needs to create an HTML connection to each AP many times for each click of a button on the user interface. Each HTML connection requires about 3 seconds

to complete. Thus, the responsiveness of the software is extremely slow. This limitation could be easily overcome by installing a custom firmware image on the AP that allows an SSH or other more responsive connection to the AP. Several such firmware images are available on the Internet, though they may be encumbered by licensing issues. An alternative would be to use the same validation exploit to upload a processing script directly to each AP. This could prove more difficult to implement, but is an alternative option. A more promising option would be to use different AP hardware.

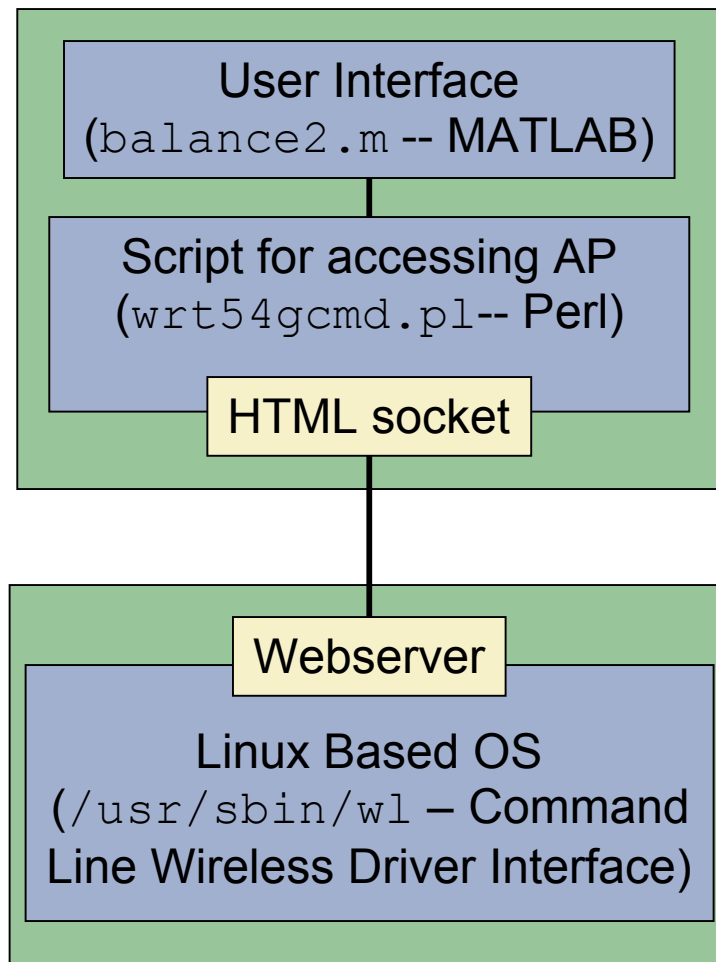


Figure B.4: Structure of `balance2layout.m` program and associated scripts and interfaces.

References

- [1] D. D. Stancil and C. P. Diehl, "Wireless signal distribution in a building hvac system," U.S. Patent 5 977 851, Nov., 1999.
- [2] D. D. Stancil, "Wireless signal distribution in a building hvac system," U.S. Patent 5 994 984, Nov., 1999.
- [3] —, "System and method for increasing the channel capacity of hvac ducts for wireless communications in buildings," U.S. Patent 6 781 477, Aug., 2004.
- [4] D. D. Stancil, O. K. Tonguz, A. Xhafa, A. Cepni, P. Nikitin, and D. Brodtkorb, "High-speed internet access via hvac ducts: a new approach," in *Globecom '01. IEEE Global Telecommunications Conference*, vol. 6, San Antonio, TX, USA, 2001, pp. 3604–3607.
- [5] C. P. Diehl, B. E. Henty, N. Kanodia, and D. D. Stancil, "Wireless rf distribution in buildings using heating and ventilation ducts," in *Wireless Personal Communications: Emerging Technologies for Enhanced Communications*, W. H. Tranter, T. S. Rappaport, B. D. Woerner, and J. H. Reed, Eds. Norwell, MA, USA: Kluwer Academic Publishers, 1999, pp. 61–70. [Online]. Available: http://www.eirp.org/~henty/refs/henty_mprg_1998.pdf
- [6] P. Nikitin, D. Stancil, A. Cepni, A. Xhafa, O. Tonguz, K. Areklett, and D. Brodtkorb, "Antennas in a waveguide propagation environment," in *IEEE Antennas and Propagation Society International Symposium*, vol. 2. IEEE, Jun. 2003, pp. 1181–1184.
- [7] P. Nikitin and D. Stancil, "Antenna radiation resistance in waveguide and in free-space," *IEEE Transactions on Antennas and Propagation*, vol. 53, no. 6, pp. 2126–2128, Jun. 2005.
- [8] —, "Connection between radiation resistances of antenna in rectangular waveguide and in free-space," in *IEEE Antennas and Propagation Society International Symposium*, vol. 2. IEEE, 2004, pp. 2035–2038.
- [9] P. Nikitin, D. Stancil, A. Cepni, O. Tonguz, A. Xhafa, and D. Brodtkorb, "Propagation model for the hvac duct as a communication channel," *IEEE Transactions on Antennas and Propagation*, vol. 51, no. 5, pp. 945–951, May 2003.
- [10] P. Nikitin, D. Stancil, A. Cepni, A. Xhafa, O. Tonguz, and D. Brodtkorb, "Propagation modelling of complex hvac networks using transfer matrix method," in *IEEE Antennas and Propagation Society International Symposium*, vol. 2. IEEE, Jun. 2003, pp. 126–129.

- [11] P. Nikitin, D. Stancil, O. Tonguz, A. Xhafa, A. Cepni, and D. Brodtkorb, "Impulse response of the hvac duct as a communication channel," *IEEE Transactions on Communications*, vol. 51, no. 10, pp. 1736–1742, Oct. 2003.
- [12] —, "Rf propagation in an hvac duct system: impulse response characteristics of the channel," in *IEEE Antennas and Propagation Society International Symposium*, vol. 2, 2002, pp. 726–729.
- [13] P. Nikitin, D. Stancil, A. Cepni, A. Xhafa, O. Tonguz, and D. Brodtkorb, "A novel mode content analysis technique for antennas in multimode waveguides," *IEEE Transactions on Microwave Theory and Techniques*, vol. 51, no. 12, pp. 2402–2408, Dec 2003.
- [14] —, "Novel mode content analysis technique for multimode waveguides," in *IEEE MTT-S International Microwave Symposium Digest*, vol. 3. IEEE, Jun. 2003, pp. 1827–1830.
- [15] A. Xhafa, O. Tonguz, A. Cepni, D. Stancil, P. Nikitin, and D. Brodtkorb, "On the capacity limits of hvac duct channel for high-speed internet access," *IEEE Transactions on Communications*, vol. 53, no. 2, pp. 335–342, Feb. 2005.
- [16] —, "Theoretical estimates of hvac duct channel capacity for high-speed internet access," in *IEEE International Conference on Communications*, vol. 2, 2002, pp. 936–939.
- [17] A. G. Cepni, D. D. Stancil, A. E. Xhafa, B. E. Henty, P. V. Nikitin, O. K. Tonguz, and D. Brodtkorb, "Capacity of multi-antenna array systems for hvac ducts," in *IEEE International Conference on Communications*, Piscataway, NJ, USA, 2004. [Online]. Available: http://www.eirp.org/~henty/refs/henty_icc_2004.pdf
- [18] A. Cepni, A. Xhafa, P. Nikitin, D. Stancil, and O. Tonguz, "Multi-carrier signal transmission through hvac ducts: experimental results for channel capacity," in *Proceedings IEEE 56th Vehicular Technology Conference*, vol. 1, 2002, pp. 331–335.
- [19] O. K. Tonguz, D. D. Stancil, A. E. Xhafa, A. G. Cepni, P. V. Nikitin, and D. Brodtkorb, "A simple path loss prediction model for hvac systems," *IEEE Transactions on Vehicular Technology*, vol. 53, no. 4, pp. 1203–1214, July 2004.
- [20] —, "An empirical path loss model for hvac duct systems," in *Globecom '02. IEEE Global Telecommunications Conference*, vol. 2. IEEE, Nov. 2002, pp. 1850–1854.
- [21] A. Xhafa, P. Sonthikorn, and O. Tonguz, "Handover performance of hvac duct based indoor wireless networks," in *Globecom '03. IEEE Global Telecommunications Conference*, vol. 2. Dallas, TX, USA: IEEE, 2004, pp. 903–907.
- [22] A. E. Xhafa, P. Sonthikorn, O. K. Tonguz, P. V. Nikitin, A. G. Cepni, D. D. Stancil, B. Henty, and D. Brodtkorb, "Seamless handover in buildings using hvac ducts: A new system architecture," in *Globecom '03. IEEE Global Telecommunications Conference*, 2003. [Online]. Available: http://www.eirp.org/~henty/refs/henty_globecom_2003.pdf
- [23] H. Andersson, P. Larsson, and P. Wikstrom, "The use of hvac ducts for wcdma indoor solutions," in *IEEE 59th Vehicular Technology Conference*, vol. 1. IEEE, May 2004, pp. 229–233.

- [24] M. Fink, "Time-reversed acoustics," *Scientific American*, vol. 281, no. 5, pp. 91–97, 1999.
- [25] A. Derode, A. Tourin, J. de Rosny, M. Tanter, S. Yon, and M. Fink, "Taking advantage of multiple scattering to communicate with time-reversal antennas," *Physical Review Letters*, vol. 90, p. 014301, Jan. 2003.
- [26] P. V. Nikitin, "Analysis of heating, ventilation and air conditioning ducts as a radio frequency communication channel," Ph.D. dissertation, Carnegie Mellon University, Pittsburgh, PA, Aug. 2002.
- [27] B. E. Henty and D. D. Stancil, "Improved wireless performance from mode scattering in ventilation ducts," in *IEEE Antennas and Propagation Society Symposium*, July 2005. [Online]. Available: http://www.eirp.org/~henty/refs/henty_aps_2005.pdf
- [28] Nlanr/dast : Iperf 1.7.0 - the tcp/udp bandwidth measurement tool. <http://dast.nlanr.net/Projects/Iperf/>. [Online]. Available: <http://dast.nlanr.net/Projects/Iperf/>
- [29] C. A. Balanis, *Advanced Engineering Electromagnetics*. John Wiley & Sons, 1989.
- [30] Yellowjacket 802.11b wi-fi receiver. <http://www.bvsystems.com/Products/WLAN/Yellowjacket/yellowjacket.htm>. [Online]. Available: <http://www.bvsystems.com/Products/WLAN/Yellowjacket/yellowjacket.htm>
- [31] K. Carter, "Predicting propagation loss from leaky coaxial cable terminated with an indoor antenna," in *Proceedings of Virginia Tech's Eighth Symposium on Wireless Personal Communications*, Blacksburg, VA, USA, June 10-12 1998.
- [32] A. A. M. Saleh and R. A. Valenzuela, "A statistical model for indoor multipath propagation," *IEEE Journal on Selected Areas in Communications*, vol. 5, no. 2, pp. 128–137, Feb. 1987.
- [33] H. Hashemi, "The indoor radio propagation channel," *Proceedings of the IEEE*, vol. 81, pp. 943–968, July 1993.
- [34] N. Blaunstein and M. Levin, "Vhf/uhf wave attenuation in a city with regularly spaced buildings," *Radio Science*, vol. 31, no. 2, pp. 313–324, 1996.
- [35] T. S. Rappaport, *Wireless Communications, Principles and Practice*. Prentice-Hall, Inc., 1996.
- [36] Linksyswrt54g – seattle wireless. <http://www.seattlewireless.net/index.cgi/LinksysWrt54g>. [Online]. Available: <http://www.seattlewireless.net/index.cgi/LinksysWrt54g>
- [37] B. E. Henty, "Throughput measurements and empirical prediction models for ieee 802.11b wireless lan (wlan) installations," Master's thesis, Virginia Tech University, Blacksburg, VA, USA, Aug. 2001. [Online]. Available: <http://scholar.lib.vt.edu/theses/available/etd-08142001-172523/>
- [38] Netiq corporation. <http://www.netiq.com/>. [Online]. Available: <http://www.netiq.com/>

- [39] Products :: Ixia - leader in ip performance testing. http://www.ixiacom.com/products/performance_applications/pa_display.php?skey=pa_q_check. [Online]. Available: http://www.ixiacom.com/products/performance_applications/pa_display.php?skey=pa_q_check
- [40] Wireless lan amplifiers. <http://www.ssbusa.com/wireless.html>. [Online]. Available: <http://www.ssbusa.com/wireless.html>
- [41] Cisco aironet 802.11g 1100 and 1200 series access point upgrade kit [cisco aironet 1200 series] - cisco systems. http://www.cisco.com/en/US/products/hw/wireless/ps430/products_data_sheet0900aecd801b9233.html. [Online]. Available: http://www.cisco.com/en/US/products/hw/wireless/ps430/products_data_sheet0900aecd801b9233.html
- [42] Cisco aironet 802.11a/b/g cardbus wireless lan client adapter [cisco aironet wireless lan client adapters] - cisco systems. http://www.cisco.com/en/US/products/hw/wireless/ps4555/products_data_sheet09186a00801ebc29.html. [Online]. Available: http://www.cisco.com/en/US/products/hw/wireless/ps4555/products_data_sheet09186a00801ebc29.html
- [43] Wap54g - seattle wireless. <http://seattlewireless.net/index.cgi/WAP54G#head-ff6d4d4d079fadd77e4592ed5155cdd24ee4cccd>. [Online]. Available: <http://seattlewireless.net/index.cgi/WAP54G#head-ff6d4d4d079fadd77e4592ed5155cdd24ee4cccd>
- [44] J. Winters, "On the capacity of radio communication systems with diversity in a rayleigh fading environment," *IEEE Journal on Selected Areas in Communications*, vol. 5, no. 5, pp. 871– 878, 1987.
- [45] J. Winters, J. Salz, and R. Gitlin, "The impact of antenna diversity on the capacity of wireless communication systems," *IEEE Transactions on Communications*, vol. 42, no. 3, pp. 1740–1751, 1994.
- [46] C. S. Lee, S. W. Lee, and S. L. Chuang, "Plot of modal field distribution in rectangular and circular waveguides," *IEEE Transactions on Microwave Theory and Techniques*, vol. 33, no. 3, pp. 271– 274, Mar. 1985.
- [47] G. J. Foschini and M. J. Gans, "On limits of wireless communications in a fading environment when using multiple antennas," *Wireless Personal Communications*, vol. 6, no. 3, pp. 311 – 335, March 1998.
- [48] I. E. Telatar, "Capacity of multi-antenna gaussian channels," *European Transactions on Telecommunications*, vol. 10, no. 6, pp. 585–595, Nov./Dec. 1999.
- [49] A. L. Moustakas, H. U. Baranger, L. Balents, A. M. Sengupta, and S. H. Simon, "Communication through a diffusive medium: Coherence and capacity," *Science*, vol. 287, no. 5451, pp. 287 – 290, 14 Jan. 2000.
- [50] B. E. Henty and D. D. Stancil, "Bandwidth limitations of phase-conjugate arrays used for multipath focusing," in *IEEE Antennas and Propagation Society Symposium*, Piscataway, NJ, 2004. [Online]. Available: http://www.eirp.org/~henty/refs/henty_aps_2004.pdf

- [51] —, “Multipath-enabled super-resolution for rf and microwave communication using phase-conjugate arrays,” *Physical Review Letters*, vol. 93, no. 24, p. 243904, Dec. 2004.
- [52] D. D. Stancil, A. G. Cepni, B. E. Henty, Y. Jiang, Y. Jin, J.-G. Zhu, and J. M. F. Moura, “Super-resolution focusing and nulling in rich multipath environments using time-reversal techniques,” in *International Conference on Electromagnetics in Advanced Applications (ICEAA)*, 2005. [Online]. Available: http://www.eirp.org/~henty/refs/henty_iceaa_2005.pdf
- [53] L. C. V. Atta, “Electromagnetic reflector,” U.S. Patent 2 908 002, Oct., 1959. [Online]. Available: [http://patft.uspto.gov/netacgi/nph-Parser?Sect1=PTO2&Sect2=HITOFF&p=1&u=%2Fmetahtml%2FPTO%2Fsearch-adv.htm&r=4&f=G&l=50&d=PALL&S1=\(%22electromagnetic+reflector%22.TI.\)&OS=TTL/%22electromagnetic+reflector%22&RS=TTL/%22electromagnetic+reflector%22](http://patft.uspto.gov/netacgi/nph-Parser?Sect1=PTO2&Sect2=HITOFF&p=1&u=%2Fmetahtml%2FPTO%2Fsearch-adv.htm&r=4&f=G&l=50&d=PALL&S1=(%22electromagnetic+reflector%22.TI.)&OS=TTL/%22electromagnetic+reflector%22&RS=TTL/%22electromagnetic+reflector%22)
- [54] C. Y. Pon, “Retrodirective array using the heterodyne technique,” *IEEE Transactions on Antennas and Propagation*, vol. 12, no. 2, pp. 176–180, Mar. 1964.
- [55] R. Miyamoto and T. Itoh, “Retrodirective arrays for wireless communications,” *IEEE Microwave Magazine*, vol. 3, no. 1, pp. 71–79, Mar. 2002.
- [56] A. Yariv and D. M. Pepper, “Amplified reflection, phase conjugation, and oscillation in degenerate four-wave mixing,” *Optics Letters*, vol. 1, no. 1, pp. 16–18, Jul. 1977.
- [57] D. M. Pepper, “Applications of optical phase conjugation,” *Scientific American*, vol. 254, pp. 74–83, Jan. 1986.
- [58] S. Karode and V. Fusco, “Use of an active retrodirective antenna array as a multipath sensor,” *IEEE Microwave and Guided Wave Letters*, vol. 7, no. 12, pp. 399–401, Dec. 1997.
- [59] —, “Multiple target tracking using retrodirective antenna arrays,” in *IEE National Conference on Antennas and Propagation*, Aug. 1999, pp. 178–181.
- [60] J. Tuovinen, G. Shiroma, W. Forsyth, and W. Shiroma, “Multipath communications using a phase-conjugate array,” in *IEEE MTT-S International Microwave Symposium Digest*, vol. 3, Jun. 2003, pp. 1681–1684.
- [61] G. Lerosey, J. de Rosny, A. Tourin, A. Derode, G. Montaldo, and M. Fink, “Time reversal of electromagnetic waves,” *Physical Review Letters*, vol. 92, p. 193904, 14 May 2004. [Online]. Available: <http://link.aps.org/abstract/PRL/v92/e193904>
- [62] J. Litva and T. K.-Y. Lo, *Digital Beamforming in Wireless Communications*. Artech House, Inc., 1996.
- [63] A. V. Oppenheim, A. S. Willsky, and S. H. Nawab, *Signals and Systems*, 2nd ed. Upper Saddle River, New Jersey, United States of America: Prentice Hall, 1997.
- [64] D. M. Pozar, *Microwave and RF Design of Wireless Systems*. John Wiley & Sons Inc., Jan. 2001.

- [65] J. M. F. Moura, Y. Jin, D. Stancil, J.-G. Zhu, A. Cepni, Y. Jiang, and B. Henty, "Single antenna time reversal adaptive interference cancellation," in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, vol. 4, Mar. 2005, pp. 1121–1124. [Online]. Available: http://www.eirp.org/~henty/refs/henty_icassp_2005.pdf
- [66] J.-G. Zhu, Y. Jiang, D. Stancil, and J. Moura, "A novel time reversal method for target detection in cluttered media," in *IEEE Antennas and Propagation Society International Symposium*, vol. 4B, 3-8 July 2005, pp. 135–138.
- [67] A. G. Cepni, "Experimental investigation of time-reversal techniques using electromagnetic waves," Ph.D. dissertation, Carnegie Mellon University, Pittsburgh, PA, USA, Dec 2006.
- [68] C. A. Baird and G. G. Rassweiler, "Adaptive sidelobe nulling using digitally controlled phase-shifters," *IEEE Transactions on Antennas and Propagation*, vol. 24, no. 5, pp. 638–649, Sep. 1976.
- [69] B. Friedlander and B. Porat, "Performance analysis of a null-steering algorithm based on direction-of-arrival estimation," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 37, no. 4, pp. 461–466, Apr. 1989.
- [70] S. Hara, S. Hane, and Y. Jia, "A simple nullsteering adaptive array antenna in ofdm," in *Proceedings*, 2003.
- [71] L. S. C. of the IEEE Computer Society, *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*. The Institute of Electrical and Electronics Engineers, Inc. Std. ANSI/IEEE Std 802.11, Aug. 1999.
- [72] C. E. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, pp. 379–423 and 623–656, Jul. and Oct. 1948. [Online]. Available: <http://cm.bell-labs.com/cm/ms/what/shannonday/paper.html>
- [73] R. V. L. Hartley, "Transmission of information," *Bell System Technical Journal*, Jul. 1928.
- [74] R. Stridh and B. Ottersten, "Mimo channel capacity on a measured indoor radio channel at 5.8 ghz," in *Conference Record of the Thirty-Fourth Signals, Systems and Computers*, vol. 1, Pacific Grove, CA, USA, 2000, pp. 733–737.
- [75] A. Swindlehurst, G. German, J. Wallace, and M. Jensen, "Experimental measurements of capacity for mimo indoor wireless channels," in *IEEE Third Workshop on Signal Processing Advances in Wireless Communications (SPAWC '01)*, 2001, pp. 30–33.
- [76] J. Wallace, M. Jensen, A. Swindlehurst, and B. Jeffs, "Experimental characterization of the mimo wireless channel: data acquisition and analysis," *IEEE Transactions on Wireless Communications*, vol. 2, no. 2, pp. 335–343, Mar. 2003.
- [77] P. Kyritsi and D. C. Cox, "Expression of mimo capacity in terms of waveguide modes," *Electronics Letters*, vol. 38, no. 18, pp. 1057–1058, Aug. 2002.
- [78] P. Kyritsi, G. Papanicolaou, P. Eggers, and A. Opera, "Miso time reversal and delay spread compression for fwa channels at 5ghz," *Antennas and Wireless Propagation Letters*, vol. 3, pp. 96–99, 2004.

-
- [79] P. Kyritsi, G. Papanicolaou, P. Eggers, and A. Oprea, "Time reversal techniques for wireless communications," in *Vehicular Technology Conference*, vol. 1, Sep. 2004, pp. 47–51.
- [80] H. T. Nguyen, J. B. Andersen, G. F. P. P. Kyritsi, and P. C. F. Eggers, "Time reversal in wireless communications: A measurement-based investigation," *IEEE Transactions on Wireless Communications*, vol. 5, no. 8, pp. 2242–2252, Aug. 2006.
- [81] H. Nguyen, J. Andersen, and G. Pedersen, "The potential use of time reversal techniques in multiple element antenna systems," *IEEE Communications Letters*, vol. 9, no. 1, pp. 40–42, Jan. 2005.
- [82] P. Kyritsi, P. Stoica, and G. Papanicolaou, "Time reversal and zero-forcing for wlan applications," in *Wireless Personal Multimedia Communications*, 2005.
- [83] P. Kyritsi and G. Papanicolaou, "One-bit time reversal for wlan applications," in *Proceedings of the Personal Indoor and Mobile Radio Communications Conference*, vol. 1, 2005, pp. 532 – 536.