

Directional Communications to Improve Multicast Lifetime in Ad Hoc Networks

Kerry Wood

Final Examination submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in
Computer Engineering

Dr. Luiz A. DaSilva, Chair
Dr. Scott F. Midkiff
Dr. Y. Thomas Hou
Dr. Sanjay Raman
Dr. C. Patrick Koelling

July 26, 2006
Blacksburg, Virginia

Keywords: directional antennas, power control, maximum lifetime, multicast routing,
cross-layer optimization
Copyright 2006, Kerry Wood

Directional Communications to Improve Multicast Lifetime in Ad Hoc Networks

Kerry Wood

(ABSTRACT)

Wireless ad-hoc networks are easily deployed, untethered to infrastructure, and have virtually an unlimited number of applications. However, this flexibility comes at the cost of finite and often unreplenishable power supplies. Once a node has consumed all of its power, it can no longer receive, transmit, gather information, or otherwise participate in the network. Therefore, reducing the amount of energy necessary for node communication has been an area of intense research.

Previous work has investigated the use of directional antennas as a method to reduce inter-node power requirements. However, most proposed methods ignore inter-session interference, propose heuristic solution methods, and ignore the use of directional antennas for signal *reception*.

We develop a flexible mixed-integer linear program (MILP) designed to optimize max-min multicast path lifetime for directional antenna equipped networks in the presence of interference. The MILP is utilized to perform a comparison directional antenna use for signal *transmission* and *reception*. Results indicate that directional reception is slightly superior to transmission for the defined max-min lifetime metric, and vastly superior when considering cumulative power use.

We further analyze the performance of interference-ignorant link-based heuristics designed for both directional transmission and directional reception as they perform in our more realistic model. Our results show that interference-ignorant methods cannot find feasible solutions unless all nodes are equipped with high gain, high efficiency directional antennas.

Even in these cases, directional reception outperforms directional transmission.

Because of the superiority of directional reception, we focus our attention on this method. A heterogeneity study is performed, and two heuristic methods for approximating the MILP optima are developed. We find that even under heterogeneous conditions, directional reception can increase network lifetime.

Finally, a genetic algorithm (GA) and semi-distributed heuristic method are developed as alternatives to the MILP. Results show that the GA often can find solutions with lifetimes 85% as long as the optimal. Our semi-distributed heuristic, designed to be even more computationally simple than the GA, and to serve as a basis for a distributed protocol, is almost as effective as the GA as approximating optimal solutions.

We conclude that directional reception is the superior method of antenna use for extending max-min multicast tree lifetime, that it works well in heterogeneous conditions, and lends itself well to heuristic design.

Acknowledgments

I would like to first thank my wonderful wife Elizabeth. For all you went through, I'm not sure which one of us had a tougher time in graduate school. Thank you for your love, understanding, and patience.

I would also like to express my thanks to Dr. DaSilva for being an amazing advisor and mentor. Without your help, there is absolutely no way I would have made it through. There just aren't words to describe what I've learned, and I can only hope I was *half* as good a student as you are a teacher. You allowed me the freedom to think, experiment, fail, and try again. I know that I've grown personally and professionally during the last three years, and I'm sure that you were directly responsible for most of it. My most sincere thanks for everything.

My thanks to Dr. Midkiff for allowing me to be part of the wonderful IREAN program here at Virginia Tech, and for all of his help and understanding as I fumbled through the Ph.D. process. Additionally, my immeasurable gratitude to Dr. Koelling for being a great source of ideas, and especially for his assistance in providing me CPLEX access (without which I never could have accomplished half of the work I did). Finally, thanks to Dr. Hou and Dr. Raman for their continued assistance and guidance throughout my Ph.D. as members of my committee.

Two folks I can't thank enough are Michael Thompson and Josh Edmison. The friendship, the chats, and an occasional Tuesday bender are undeniable contributors to what's left of my sanity. In twenty years, when this document is woefully out of date, I'll still remember the good times with you guys. You are just ridiculously wonderful people, and Blacksburg will never be the same without you. Thanks for making this fun.

Many, many thanks to Sastry Kompella. I think that a cup of coffee at Squires and an hour of laughing might be the cure for everything.

And, of course, no list of acknowledgments would be complete without thanking Ryan Thomas and Juan Suris. I'm surely going to miss Durham 465, especially when I want to crack a rude joke, yell out a question, or just jot down some ideas on a whiteboard. It actually *was* fun coming to the office.

I have to thank my Uncle Pete, because without him I might be working at a gas station.

You make a great role model and I could write another hundred pages just listing out all the things you've done for me over the years. Don't think I don't remember, or that I'll ever forget.

Thanks to Aunt Laurie, Brian, Tim and Jenna for always making me feel like I was part of the family. It sure is nice to have a permanent job as "Bucky".

To Auntie Linda, my thanks for continually keeping me grounded and reminding me that I do, in fact, have a softer side. I certainly need that now and again, and without it, my head might be bigger than it is.

I have to thank my sister, who has grown up and turned into a fantastic mother. Don't think I don't notice how wonderful your kids are and how you take care of Grammie. I'm sure it makes your Mom as proud as it makes me.

To all the boys and gals back home, thanks for reminding me that "hermit" is not a compliment, and that blue-collar is. Believe it or not, I'm looking forward to the lumps I'm going to take for disappearing during this whole process.

Dedication

To Elizabeth, whose love and dedication are infinite and amazing:

In a sky bejeweled with the temptation of sparkling stars, I may spend eternity reaching for them. But I'll always follow the brightest one in the heavens, the one that leads me home to you.

To Grammie, my hero:

No matter what I might ever do, what I might ever accomplish, I can only hope that someday I can measure up to you. I've known you my whole life, and you're still the most amazing person I've ever met. Thanks for everything.

To Mom, who I miss very much:

As time passes, I realize more and more how much you sacrificed for us. I wish you were here to share a celebratory beer, and I sure hope that I make you proud.

Contents

1	Introduction	1
1.1	Energy-Efficient Multicast Delivery in Ad-Hoc Networks	2
1.2	Problem Statement	4
1.3	Methodology	6
1.4	Thesis Statement	7
2	Related Work	8
2.1	Wireless Networking and Energy Efficiency	11
2.1.1	Minimum-Energy Optimization	12
2.1.2	Maximum Lifetime Optimization	15
2.1.3	Directional Antennas	17
2.2	Optimization and Modeling	19
2.2.1	Disjunctions and Fixed-Cost Problems	19
2.2.2	Logic Programming	20
2.3	Summary and Comparison	21
3	Problem Statement and Approach	24
3.1	Introduction	24
3.2	Directional Antennas and Ad-Hoc Networks	25
3.3	Metrics	26
3.3.1	Justification of Metric	27
3.4	Communication and Antenna Model	28

3.4.1	MAC	28
3.4.2	Channel Model	29
3.4.3	Antenna Model	30
3.5	Solution Methods	32
3.5.1	Mixed Integer Linear Programs	32
	SINR Constraint and Interference Modeling	33
	Beam Enumeration	35
	D-TX MILP Model	39
	D-RX MILP Model	40
	MILP Complexity	40
	MILP Summary	43
3.6	Heuristic Methods	43
3.6.1	DRIP	44
3.6.2	Semi-distributed Algorithm	45
	Design Observations	46
	Power Control Algorithm	47
	Semi-Distributed D-RX Algorithm Overview	48
3.6.3	Genetic Algorithm / Metaheuristic	49
	Structure	51
	Objective and Feasibility	51
3.7	Summary	53
4	Results: Directional Transmission vs. Directional Reception	54
4.1	Experimental Setup	55
4.2	Metric Comparison	57
4.3	Intuition	62
4.4	Heuristic Analysis	64
	4.4.1 Mapping Heuristic Topology into Interference Model	64
	4.4.2 Experimental Results	66

4.4.3	Summary	69
5	Results: Directional Reception	71
5.1	Heterogeneity Study	72
5.1.1	Number of Directional Antenna Equipped Nodes	72
5.1.2	Network Size	74
5.1.3	Number of Receivers	75
5.1.4	Beamwidth	77
5.1.5	Sidelobe Intensity	78
5.2	Distributed Heuristic Results	78
5.2.1	MILP Comparison	79
5.2.2	DRIP Comparison	83
5.2.3	Summary	85
5.3	Genetic Algorithm Results	86
5.3.1	GA Performance vs. Optimal	88
5.3.2	GA Performance vs. Distributed Heuristic	91
5.3.3	Genetic Algorithm Results – Domain Reduction	95
6	Conclusion	98
6.1	Summary of Contributions	99
6.2	Future Work	101
A	Power and Semi-Distributed Algorithms	112

List of Figures

3.1	Beam Illustration	30
3.2	Bulb-and-Cone Antenna Model	31
3.3	D-RX Beam Configurations	37
3.4	D-TX: Formulation for Max-Min Network Lifetime with Inter-Node Interference	39
3.5	D-RX: Formulation for Max-Min Network Lifetime with Inter-Node Interference	41
3.6	D-TX Beam Limit Illustration	42
3.7	DRIP Link Cost	44
3.8	Overhearing	46
3.9	Iterative Power Algorithm	47
3.10	Genetic Algorithm Overview	50
3.11	Genome Structure	52
4.1	D-TX / D-RX Performance Ratio Histograms	57
4.2	D-TX / D-RX Lifetime Ratio vs. Network Parameters	59
4.3	D-TX / D-RX Cum. Power Ratio vs. Network Parameters	60
4.4	D-TX / D-RX Ratios vs. Single Network Parameter	61
4.5	Intuition Example	63
4.6	Heuristic Mapping Procedure	65
4.7	Heuristic Lifetime Ratio vs. $ \mathcal{R} $	66
4.8	Heuristic Lifetime Ratio vs. Network Size	67
4.9	Actual Lifetime Values (Heuristics and MILP)	70

5.1	Lifetime and Cumulative Power vs. $ \mathcal{D} $	73
5.2	Lifetime and Cumulative Power vs. Network Distance	75
5.3	Lifetime and Cumulative Power vs. $ \mathcal{R} $	76
5.4	Lifetime and Cumulative Power vs. Min. Beamwidth	77
5.5	Lifetime and Cumulative Power vs. Min. Beamwidth	79
5.6	Distributed Heuristic to MILP Lifetime Ratios	80
5.7	Distributed Heuristic to MILP Cumulative Power Ratios	81
5.8	Distributed Heuristic Lifetime and Cumulative Power Bar Graphs	82
5.9	GA and MILP Lifetime and Cumulative Power	89
5.10	GA / MILP Lifetime Ratios	90
5.11	GA / MILP Power Ratios	91
5.12	GA and Distributed Heuristic Metrics	92
5.13	GA / Distributed Heuristic Lifetime Ratios	93
5.14	GA / Distributed Heuristic Power Ratios	94

List of Tables

3.1	Notation.	33
3.2	Complexity of MILP Formulations	42
5.1	Percentage of Runs where Distributed Heuristic is Feasible and DRIP is Infeasible	84
5.2	Average DRIP / Dist. Heuristic Lifetime Ratios ($ \mathcal{D} = 15$) only	84
5.3	GA Parameters	86
5.4	Average GA Run Times (s)	87
5.5	Average GA Run Times (s)	87
5.6	Percent of Binary Variables Pruned in MILP	96
5.7	Average Number of Binary Variables Pruned in MILP	96

List of Algorithms

1	DRIP Multicast Routing Algorithm	45
2	Genetic Algorithm Objective Function – $\text{objective}(\text{genome})$	53
3	$\text{powerAlgorithm}(\mathcal{L})$	112
4	Semi-Distributed Algorithm	113

Chapter 1

Introduction

Ad-hoc networks have a virtually limitless number of possible applications. By their definition, these networks do not require any preconfigured or preinstalled infrastructure and are self-organizing. *Wireless* ad-hoc networks have the additional benefit of using wireless channels for inter-node communication. Nodes are untethered and independent.

Applications, potential uses, and deployment strategies for wireless ad-hoc networks cover a broad spectrum. For instance, nodes may be within radio range by chance, and be designed to deliver services to an individual customer. Imagine a busy subway station, where waiting passengers' PDA devices, laptops, and even cell-phones build an ad-hoc network to enable applications such as email, video streaming, web browsing, or instant messaging. On the other hand, wireless networks may be designed with a high degree of specificity for an intended purpose and location. Picture a military application, where a cloud of tiny sensor nodes are dispersed through a random, air-dropped deployment. This network may be patrolling and monitoring a potentially inhospitable region of a battlefield, filling a need that would otherwise be costly and inconvenient.

Clearly, the term “wireless ad-hoc network” is a super-class covering a wide variety of more specific instantiations. As the price and size of devices with substantial computing power

continues to steadily decrease, the number of feasible applications for ad-hoc networks correspondingly increases. Further, the *number* of nodes that will form a single network is likely to increase. The size of the network could grow accordingly, or nodes could be more densely packed in an area.

While wireless networks have a wide variety of uses, they have two qualities that bring up important research questions. Regardless of purpose, portable or easily deployed nodes have some common major characteristics: finite energy reserves and wireless communication. Either of these qualities on their own encourage network nodes to use energy as efficiently as possible, and the combination of the two makes energy conservation acutely important. Moreover, increasing the number of network nodes exacerbates one issue and potentially assists with the other. As tightly packed nodes compete for channel access, free spectrum becomes more scarce and nodes must cooperate to keep the network alive. However, a dense group of nodes can also cooperate and form a multi-hop topology, assisting sources by forwarding packets to more distant neighbors.

In this chapter, we introduce the reader to the research area we investigate, that of energy-efficient multicasting in wireless ad-hoc networks. We discuss the importance of this field and delve into the specifics of our initial assumptions. After motivating the problem, we briefly introduce our thesis statement and lay out the remainder of the document.

1.1 Energy-Efficient Multicast Delivery in Ad-Hoc Networks

Truly flexible and independent ad-hoc networks do not rely on any kind of infrastructure. In these networks, a node is equipped with a finite amount of initial energy, generally in the form of a battery. Some deployments, such as the devices in our subway station example, nodes could potentially be replenished. However, foolishly expending energy decreases user

utility (and increases frustration) as the battery must be more frequently recharged.

Applications such as sensor networks, on the other hand, have more drastic consequences when a node's energy reserves are depleted. These nodes are entirely independent and once energy stores are depleted, the node is essentially "dead" as it is unable to participate in the network. Note, though, that nodes with replenishable power supplies also have reason to limit the amount of energy expended in wireless communication. Unnecessary or excessive transmission power merely increases interference, meaning the wireless channel is used inefficiently.

The death of a network node due to expended power reserves can have consequences that affect the entire network. Nodes may have different functionality, data gathering capabilities, processing power, or other characteristics that make them unique. The loss of a node may be synonymous with the loss of a network function. If nodes are deployed randomly, the loss of a node may partition the network. Thereby, a network designed to sink information to a single location could suffer catastrophic failure, as an entire subset of nodes may actually be unreachable.

For this document, we consider the general situation where a network intends to deliver a source-initiated multicast to other network nodes. Multicasts are the most general form of communication, containing unicast and broadcast as subsets. As stated by Comer [1]:

The advantage of multicasting lies in the ability to limit broadcasts: every computer in a multicast group can be reached with a single packet transmission, but computers that choose not to participate in a particular multicast group do not receive packets sent to the group.

In other words, only a subset of nodes that desire to receive the multicast have packets delivered to them. Packets can be replicated at intermediate hops, and multiple receivers can hear the same packet. Utilizing multiple receivers for a single packet lends itself well to wireless communication. Since wireless communication can be "overheard," wireless networks get the benefit of multiple receivers at no additional cost.

However, communicating wirelessly can place a large drain on a node's energy reserves. It is well known that the strength of wireless signals falls off quickly as a function of distance. The power required to traverse even small distances can be quite high. This is especially true when a node utilizes omni-directional antennas. When transmitting or receiving omni-directionally, radiation is either sent or collected from all angular directions equally.

Directional antennas represent a way in which nodes can more efficiently use transmitted power. A directional antenna focuses radiation, which provides additional gain in the desired direction. The directional antennas we consider have a main beam "lobe" and smaller, less focused side-lobes. The main lobe represents the direction of optimal gain, and usually has the majority of the focused power. If a neighboring node lies within the main lobe, a specific amount of gain is achieved (in this document we assume that in-beam gain is inversely proportional to beam-width). We also consider the nodes *outside* the main beam that might lie in the antenna side-lobes. Side-lobes are generally thought to be a detriment to inter-node communication because signals transmitted or received through them may contribute to interference. Therefore, higher signal power might be necessary in the *intended* direction. Note that this focusing of radiation can also be used to *receive* directionally. In this case, the main lobe gathers transmitted radiation from a specific angular direction. Side-lobes are a concern in this case as well, since sources whom the node does *not* wish to receive from may be acting as interference.

1.2 Problem Statement

Our research spans two different, but related areas. First, energy conservation has been a major focus of ad-hoc and sensor networks. It is generally assumed that nodes are battery (or similarly) powered, and are equipped with a finite amount of non-replenishable energy. Therefore, by making nodes more energy-efficient, the lifetime of a deployed device, and thereby the utility of the entire network, can be enhanced.

Second, directional antennas have numerous benefits when communicating wirelessly. Omni-directional antennas are widely in use in simulations because of their simplicity and intuitiveness. In practice, real antennas always have some form of directionality. Unfortunately, this directionality is usually an artifact of device design, and the antenna cannot be pointed or configured for optimal communication. Since this directionality is often neglected, we consider this to be analogous to omni-directional communication. However, omni-directional communication wastes a significant portion of transmitted power, as it sends or receives in all directions rather than towards a specific node. Directional communication provides gain, and also allows for spatial multiplexing. These two properties result in lower interference, longer transmission distances, and lower overall transmission powers.

Intuitively, if power can be reduced by using directional antennas, a node's energy reserves can be made to last for longer periods of time. However, omni-directional communication has a simplifying property; it does not require knowledge of a neighbor's angular location. Instead, it can be assumed that all nodes within a specified distance may receive a transmission sent at a suitable power level. Though directional communication has the aforementioned advantages, it comes at the cost of additional complexity.

Specifically, passing messages becomes a cross-layer problem. The available links in the network are a function of transmission power, antenna direction, main lobe and side-lobe gain, and interference from other nodes. Interference is clearly the most interesting parameter. Although directional antennas help to reduce medium contention, they cannot eliminate it. Therefore, to use the medium most effectively, nodes must control not only the antenna parameters, but the transmission power settings. We assume that nodes communicating wirelessly share a common channel, and are therefore coupled because of interference.

1.3 Methodology

In this document, we are concerned with finding optimal multicast delivery schemes for networks equipped with steered-beam directional antennas in more realistic wireless models. We develop mixed-integer programs for the cases where nodes transmit directionally, and where nodes receive directionally, both incorporating interference and inter-node coupling. Accounting for medium competition allows us to determine how effective spatial multiplexing alone is at extending network lifetime. In contrast to heuristic techniques, MILP formulations guarantee that when a solution is found it is optimal for the given problem. This optimality guarantee removes the stochastic artifacts of simulations studies, and allows us to directly draw conclusions on the impact of antenna modeling and parameters. The operations research (O.R.) framework is then utilized as a black-box to determine which antenna use paradigm is most effective at prolonging network lifetime. Once this determination is made, we design heuristic methods to approximate the optima found from the MILP.

Since we are most interested in the actual effect that directional antennas have on network lifetime, we vary the network parameters to isolate complicating variables. The number of directional antennas in the network, the network size, the receiver set size, antenna parameters, and the method the antennas utilize (directional transmission, directional reception) are changed during the experiment. Results generated by the MILP are then compared to identify important characteristics and to isolate trends.

Based on our comparison results we develop two heuristic techniques for directional reception. One is a semi-distributed heuristic designed to approximate the network optima, with the goal of also being implementable in real deployments. Further, we utilize the fairly common meta-heuristic technique of genetic algorithms to provide an additional method of optima approximation.

1.4 Thesis Statement

The results presented in this document lead to the conclusion that directional reception is superior to directional transmission for optimizing network lifetime when delivering multicasts in wireless ad-hoc networks. Directional reception programs are smaller, more intuitive, and easier to solve. Further, the output metrics show that directional reception improves lifetime and cumulative power statistics over directional transmission. Finally, we show that directional reception lends itself well to heuristic design, and that our proposed heuristics closely approximate the optima found by the mixed-integer program.

Our results also indicate that output metrics are strongly influenced by interference and more realistic antenna parameters. Whereas previous work has assumed that directional antennas allow for perfect spatial multiplexing, or assume the operation of a perfect MAC scheme, we model inter-session competition for the channel. Our conclusion is that without *another* method of multiplexing (coding, frequency, time, etc.), directional antennas cannot completely mitigate the effects of inter-node interference. We finally conclude that not accounting for interference when it is present can lead to infeasible results, and at best, inefficient network operation.

Chapter 2

Related Work

The research presented in this document spans work in Electrical and Computer Engineering and a blend of Industrial Systems Engineering and Operations Research. Our main focus is the development of techniques for optimizing lifetime for multicast delivery in networks in which nodes are equipped with directional antennas. This collection of tools is then used to bound and characterize the effects directional antennas have on lifetime. From Electrical Engineering, the topic involves the combination of two major networking areas: that of energy awareness (and lifetime optimization), and the efficient use of directional antennas in ad-hoc networks. The development of the model and the application of solution methods is tied to well-known Operations Research and Industrial Engineering specialties. Our first contribution is the creation of a flexible, powerful model that uses well-known optimization techniques to find guaranteed optimal solutions. The model is then utilized to better understand how directional antennas impact network lifetime, and to illustrate the best method by which to incorporate these antennas.

The topic of energy efficiency and prolonging network lifetime in ad-hoc networks has been the central research area of numerous previous papers. In short, researchers have recognized that for true independence and ease of deployment, wireless networks nodes will be battery-powered. Obviously, unless replenishable, once a node has consumed all its energy

reserves, it can no longer participate in the network. Therefore, a significant amount of work has been dedicated to the problem of conserving energy and prolonging lifetime in ad-hoc networks. The specific research goal and the methods proposed have varied. For instance, a large portion of the initial research attempted to minimize energy consumption when delivering a broadcast in networks where nodes are equipped with omni-directional antennas [2][3][4][5][6][7]. Note the specificity in both metric and communication architecture. Minimum energy schemes are defined as those that minimize the *cumulative* power expenditure for a specific routing topology. Further, “broadcast” indicates that all network nodes must be included in the resulting routing tree.

While the previously mentioned papers investigate minimum energy broadcast delivery schemes, others have investigated maximizing network lifetime for multicasts [8][9][10]. Put simply, every communication architecture from unicast to broadcast represents a subset of multicast. Multicast, therefore, is clearly the more flexible and extensible architecture. The apparently slight difference in metric actually represents a significant change in overall performance. For instance, while optimizing for minimum energy intuitively seems to directly lead to longer lifetimes, Wieselthier et al. [11] showed that the definition of the term “lifetime” is important. Chang and Tassiulas [8] define the lifetime of a network as the time until death of the first node. Minimum energy schemes do not optimize this metric directly. Instead they focus on minimizing cumulative power. Since this metric may use network energy unevenly, it may potentially drain a single node’s battery quickly. The maximum lifetime definition, on the other hand, recognizes that every node in the network may have unique functionality, or may be a critical link in overall network connectivity. The death of even a single node could render the overall network inoperable, or represent a large decrease in utility. Furthermore, when considering a single routing topology, node death necessitates that a new routing tree be found.

We also find that the solution and modeling methodology presented in many papers is heuristic in nature and therefore difficult to analyze. With the exception of [2][12][13] which use classical optimization and mathematical programming techniques to find a solution, other

papers present heuristic methods. To validate the algorithm’s contribution, heuristics are often implemented in simulation packages, and compared against other proposals. Although this can show that one method has advantages over another for the given network setup, this comparison makes no claims about the “true” network optimal.

Directional antennas have characteristics that lead some to believe their inclusion could increase network lifetime. When transmitting or receiving directionally, a node can focus power towards/from a specific direction [14]. In addition to providing gain, this helps to mitigate interference by reducing the amount of power transmitted/received in non-useful directions. The benefits of directional communication are evident, and recent papers have included directional antennas while optimizing for multicasts [15][16][12]. This combination of directional communication and the flexibility of multicasts is the basis for our research area. Guo and Yang [12][13] develop a mathematical programming formulation for the problem. If solvable, a MILP is guaranteed to generate optimal solutions, removing the heuristic uncertainty regarding output values. However, their work does not include inter-node interference and medium contention, instead assuming the underlying operation of a MAC scheme or individual and non-overlapping inter-node channels will completely avoid such interference.

Maximizing lifetime for ad-hoc networks is clearly an interesting and difficult problem. The inclusion of directional antennas in the network adds the ability to spatially multiplex, but comes at the cost of additional dimensions of optimization, and additional complexity. However, the benefits of using directional antennas seem clear. Interference is reduced by transmitting or collecting radiation more efficiently, transmission range is increased with antenna gain, and overall power use is decreased. Intuitively, these characteristics should allow the network to utilize energy stores more efficiently and prolong overall lifetime.

In this document, we propose a mathematical program designed to optimize multicast path lifetime as defined in [8][17][18] for multicast delivery in a network where nodes are equipped with directional antennas. We recognize that multicast is the more flexible super-class of

other useful communication techniques. Further, the max-min lifetime metric better characterizes the spreading of energy usage among all network nodes, and represents the longest lived single topology. Though heuristic techniques have advantages such as computational simplicity, they cannot make quantitative claims about performance. Instead results are verified through comparisons to other techniques under a given set of simulation assumptions. We propose a MILP that can model inter-node interference and medium contention, that also maintains generality and configurability. Though MILPs are known to scale poorly, a solution to a MILP is guaranteed to be optimal for the defined program. Furthermore, there are a wealth of theoretical methods that help to reduce solve times, and numerous commercial and free software packages are available for simulations. Our model has the important characteristics that it returns the optimal for the given model, and can leverage widely available expertise and software. The program can then be utilized to generate data corresponding to networks with different parameters. These data are then analyzed to ascertain how directional antennas affect the network, and more importantly, how to best leverage these antennas. Patterns and simple “rules of thumb” can be inferred from the analysis to develop computationally cheap heuristic methods for the problem.

2.1 Wireless Networking and Energy Efficiency

Wireless networking has advantages over wired networking including lack of requirements for pre-built infrastructure, easy deployment, and configurability. A significant advantage of communicating wirelessly is the “wireless multicast advantage” as described in [19]. In short, this name reflects the idea that a node can communicate with multiple neighbors with a single wireless transmission. Any neighbors receiving the signal at a sufficient SINR can receive data. Directional antennas modify this behavior. When transmitting directionally, a node may change the direction, the width, or the power of a beam in order to cover a subset of neighboring nodes. While there is no penalty for changing beam direction, when a node chooses to widen a beam, this comes at the cost of reducing the main lobe gain.

Therefore, optimizing the configuration of the communication profiles within the network is a problem with numerous degrees of freedom. To find the best possible solution any network node should be able to change all of the aforementioned parameters. Logical connections are coupled, as our model considers contention for the wireless channel.

To find the global optimal solution we approach the problem as a cross-layer design problem. The goal of any proposed method is to develop a routing topology for the wireless network. Unlike wired networks where logical links are fairly static and pre-defined, wireless networks can have a combinatorially large number of possible network links. Depending upon the choice of transmission power, and if so equipped, directional antenna parameters, a node can choose how it would like to reach any possible subsets of neighboring network nodes. This leads us to the observation that each logical link does not have a pre-defined power requirement. Instead, since interference is considered, power settings and antenna choices at a node affect all other nodes. Therefore, we are faced with the combinatorial problem of beam choice and power variables subject to the feasibility requirements of SINR constraints and multicast forwarding topology. In this document, we remove the operation of a MAC scheme to better analyze *only* the effects of spatial multiplexing. This further allows us to primarily focus on the optimization of the physical and network layers.

2.1.1 Minimum-Energy Optimization

A large volume of research has been dedicated to the optimization of energy consumption when sending information through a multi-hop wireless network. In actuality, however, energy optimization is a super-class for a variety of problems that may only be weakly correlated.

It is intuitive that nodes with finite energy reserves should be cognizant of the power necessary to communicate in the network, and the network itself must optimize energy consumption to keep nodes alive [20] [21]. One problem that has been extensively studied is optimizing the delivery of a single message to *all* network nodes (the minimum-energy broadcast prob-

lem). For these selected papers, the chosen problem definition was specific: the metric chosen was minimum cumulative tree power for a broadcast in networks with omni-directional antennas. This problem has been investigated in [22][7][4] and [23]. Note that these papers do *not* consider any inter-node interference and only consider omni-directional antennas. Techniques for minimum-energy broadcasting are proposed in [4][23][19][24][25][3][6][5].

Further work investigates power-efficient routing and topology control for wireless ad-hoc networks [26][27]. Here, problem definition differs from the minimum-energy broadcast setup. Generally, communication in these papers is unicast in nature, allowing for information conservation at intermediate (trans-shipment) nodes. Of interest in these works is the use of power control to reduce inter-node interference and to maintain sufficient network connectivity. It is shown that interference, node connectivity, and logical topology are coupled, and cross-layer techniques are used to optimize power use for effective routing. As with other works, however, these come with limitations. Although these papers investigate the cross-layer optimization problem, they often fix the settings of a single layer to reduce the number of variable parameters. Therefore, while these works provide intuition about the global optimization problem, they often do not consider the entire optimization space.

A slight change in the approach to the research allowed proposed methods to handle a much wider range of problem types. Wieselthier et al.'s [19] seminal work recognized that optimizing for multicast provides a more flexible solution than that of specifically focusing on broadcast. Whereas broadcast requires delivery of a message to all nodes, and unicast to a single node, multicast is defined as the delivery of a message to *any subset* of neighbor nodes. Clearly then, unicast and broadcast can be viewed as special cases of the more flexible multicast class. Wieselthier et al. introduce the multicast incremental protocol (MIP), a centralized, online, tree-building heuristic that formed a base for much subsequent work, including [28][9][16][12][13][10][29][30]. This heuristic is designed to be computationally low-cost, and to utilize the nature of the wireless medium to build minimum energy delivery trees. In short, the tree is built by adding a node at each step according to a defined metric, and removing unneeded branches upon completion. The novel observation in this work is

that a transmitting node may only have to increase transmit power *slightly* to add another neighbor to the multicast tree. For instance, at step N , node i is added to the multicast tree. Node j is *slightly* further away than node i and in the same angular direction. The cost, then, to add node j to the tree is only the cost of the *increase* in transmission power. Therefore, at each step, at least one inter-node cost is changed as the tree grows from node to node. This phenomenon was given the moniker the “wireless multicast advantage.” Results indicate that protocols that exploit this characteristic of wireless networks can achieve substantial energy savings over protocols that do not. Further enhancements were made to the basic MIP algorithm and observations. Some of note include the G-REMiT (Group-shared: Refining Energy-Efficiency of Multicast Trees) as presented in [28], and associated work in [9] and S-REMiT (Source: Refining Energy-efficiency of Multicast Trees) [31]. As with previous research in energy-efficient protocols, these methods are heuristic in nature, and make no claim as to the network optimal.

The majority of the work documented thus far contains heuristics that are based on a variety of modeling assumptions. In general, models either use a simple graph-based approach to approximating network behavior, or utilize pre-existing packages such as NS-2 [32] to generate results. NS-2 and other Monte Carlo simulations provide useful insights into the performance of a proposed algorithm, but can make no claims as to the true optimal value. However, there is at least one work that looks at the same problem type, but proposes an entirely different solution method. Das et al. [2] introduce classical optimization techniques to the problem and model it as a mixed-integer linear program with the goal of constructing a minimum power broadcast tree for nodes equipped with omni-directional antennas. This program models the network as a set of logical links having pre-defined power costs. The power required to transmit to a group of neighboring nodes is simply the power required to reach the farthest node, which is consistent with the wireless multicast advantage. Although the model incorporates these strengths, it has one important weakness common to previous work in the area. Namely, transmissions are considered to be independent of one another and it is impossible to model contention for the medium (leading to static, pre-defined link

costs). Without a perfect MAC scheme or independent channels for each logical link, this assumption is rather unrealistic. However, if a solution is found, the MILP model guarantees that the solution is optimal. As mentioned earlier, the guarantee of optimality is important in that it generates values that can be used to assess the effectiveness of heuristic solutions.

2.1.2 Maximum Lifetime Optimization

Chang and Tassiulas defined network lifetime as the time until the *death of the first node* in the network [8]. Wieselthier et al. showed that minimizing power usage may not be directly related to enhancing network lifetime [15]. While minimum energy trees are effective at conserving overall network energy, the uneven drain rate can unfairly utilize a single node’s energy resources, leading to premature death. Note the important difference in optimization goal. Minimum energy schemes are intended to conserve *overall* network power, but ignore individual nodes. In more general terms, minimum energy heuristics have no knowledge of node’s energy stores, and can therefore burn out nodes quickly. Max-min lifetime schemes, on the other hand, maximize the lifetime directly, and optimize for the worst-case.

The slight change in objective function represented an important shift in the goal of subsequent research. Papers investigating maximum lifetime as defined in [8] (max-min network lifetime) for broadcast delivery appeared in [33][22][34][35]. There has been further interest in max-min lifetime *routing* as indicated in [17][18][36]. While Sankar and Liu [17] investigate the general problem of max-min lifetime routing, Hua and Yum [18] add the additional complexity of data aggregation at each node. Orda and Yassour compare omni-directional and directional techniques for routing in their work [36]. This work relies on the “single recipient” assumption, that directional antennas can isolate at most one neighboring network node. Though this paper introduces some interesting observations regarding polynomial time optimal algorithms, its dependence upon this assumption limits comparison with contributions in our work. Il et al. [35] and Wieselthier et al. [33] again present heuristic techniques and qualitative observations for the max-min lifetime problem. In all cases, however, the

definition of network lifetime is identical: “time until death of the *first* network node.”

Kang and Poovendran [22][34], introduce an additional consideration, that of using multiple trees to deliver a broadcast over time. They develop a graph-theoretic method for optimal determination of trees in what they define as static and dynamic networks. As defined in [22][34][35][10], a network is defined as static if it has only one multicast tree for delivery. Dynamic networks, on the other hand, may switch the tree structure one or more times to better spread energy consumption.

Regardless of metric, multicast provides the most flexible solution method, and therefore was at a considerable advantage compared to broadcast solutions. Floreen et al. contribute two distinct pieces of research in this direction in that they develop maximum lifetime heuristics for multicast with omni-directional antennas and also show the problem to be NP-hard [29][10]. Specifically, they find that when the network consists of nodes with set transmit powers, a polynomial time algorithm exists for optimal maximum lifetime routing. Unfortunately, the introduction of configurable transmit power makes the problem much more difficult, and it is proven to be NP-hard. In addition to the proof, Floreen et al. develop and describe yet another heuristic for approximating the optimal solution.

Maric and Yates [37] propose a novel approach using multiple multicast trees over time and “hitch-hiking” as presented in [38] to extend network lifetime. This scheme utilizes omni-directional antennas, but incorporates an advanced communication technique which uses cumulative received power to reconstruct a signal. It is assumed that any transmission above the noise floor has some utility to a receiving node. If enough cumulative energy is received over time, a transmitted message is assumed to have been received. Transmissions must be scheduled, accounted for, and the topology can be constructed. The “hitch-hiking” assumption represents an advanced signal processing technique that is beyond the scope of this document.

2.1.3 Directional Antennas

Much of the previous research mentioned to this point in the document utilizes omni-directional antennas for inter-node communication. However, it was soon realized that directional antennas have significant advantages when compared with their omni-directional counterparts. One of the initial works to analyze the use of directional antennas in multi-hop radio networks was [39]. However, it was Ramanathan’s seminal work in [40] that employed simulations to characterize the effects of directional antennas on a variety of chosen metrics. For simplicity, Ramanathan summarily ignored the need for a MAC layer in his simulations, and assumed global knowledge of node placement and routing paths. The metrics observed did not consider power use, but the benefits of spatial multiplexing on throughput were evident.

Unfortunately, it was quickly realized that protocols designed for use with omni-directional antennas such as the 802.11 Medium-Access Control protocol [41] were not equipped to handle directional communication. Solving these issues became an area of intense focus from the research community. Directional antennas suffer from conditions at the MAC layer such as deafness [42] and new varieties of the hidden node problem [43]. To this date, no proposed MAC scheme is known to solve all the potential problems brought about by including directional antennas. Instead, papers illustrating enhancements to specific network layers or behaviors include:

- New MAC protocols designed to use directional communication [44][42][45][46][47][48][49][50];
- Schemes that investigate directional routing, as exemplified by [51][52][53];
- Yi et al. perform a theoretical analysis of the increase in capacity resulting from the use of directional antennas [54] based on the paper by Gupta and Kumar [55].

In our work, and in light of the conclusions of many of this related work, we make the assumption that the MAC protocol is only able to probabilistically prevent medium contention.

Therefore, for communication to occur, a node must be able to overpower other transmitting sources and ambient noise at the intended receiver. The benefits of this approach are twofold. First, the final routing tree is not subject to any of the previously mentioned limitations. And second, the results from this study are independent of MAC layer functionality.

From our discussion of the failings of the MAC and routing layers when directional antennas are in use, it is clear that this is a cross-layer design problem. Wieselthier et al. [16] recognized this, and combined cross-layer design for multicast with directional antennas. The result was the D-MIP (Directional Multicast Incremental Protocol) protocol. Using this protocol, a multicast tree is built by adding a single node at each iteration based on a specified cost. Beam configuration, in this case, includes the boresight and the main lobe width, which is proportional to beam gain. Though D-MIP is designed to create minimum energy routing trees, the metric reported is network lifetime. Similar to previous work, it is defined as the time until death of the first node. As with other proposed methods, this algorithm is heuristic and cannot claim a proximity to optimality.

Later work attempted to find optimal values to better evaluate metrics and performance. Guo and Yang bring optimization techniques to bear on this problem in [12] [13]. In both, mixed-integer optimization models are presented that seek to create minimum energy multicast delivery schemes for networks utilizing directional transmission between nodes. The network is represented as a system of logical links where link cost is equivalent to the inter-node energy requirement as in [2] and mirroring the network model of [16]. Binary variables are used to indicate the nodes that are within the main beam lobe. The scheme by which nodes are determined to be “in-beam” is novel and scales better than simple enumeration. However, it also does not allow the model to account for inter-node coupling and interference. Instead, the MILP is specifically designed to generate minimum power multicast trees when it can be assumed that network transmissions are independent. In short, the method presented is similar to that of Das et al. [2], with extensions for directional antenna modeling.

More recently, Hou et al. have shown that the max-min lifetime problem with directional

transmission is also NP-hard [30]. In this work, they also present a heuristic algorithm for the determination of max-min lifetime optimal multicast trees. Their maximum lifetime routing for multicast with directional antennas (MLR-MD) iteratively improves the multicast tree by replacing bottleneck links with better alternatives to improve overall network lifetime. Results indicate that MLR-MD performs substantially better than D-MIP in a head to head comparison.

2.2 Optimization and Modeling

Mixed integer, and especially binary, programming has received considerable attention from the optimization community. Strictly linear programming problems containing thousands (and possibly tens or hundreds of thousands) of constraints and variables can usually be easily solved by commercial packages. However, the use of integer or binary variables quickly limits problem scalability. We utilize binary variables for our Signal to Noise Ratio sufficiency constraint (SINR constraint). Put simply, we assume that DSP hardware can decode received messages if the ratio of signal strength to all other received noise is above a set threshold. The mixed-integer program presented here uses a single binary variable to indicate that this SINR ratio is sufficient and communication can occur. Therefore, our proposed SINR sufficiency constraints require one binary variable per possible logical link, and therefore can grow quickly with the number of nodes. Fortunately, the model does bear a resemblance to binary programming models that are well-known to have useful internal structures. Some similarities and possible solution methods are discussed in detail below.

2.2.1 Disjunctions and Fixed-Cost Problems

Certain specific MILP problem types are known to have structures that can assist in quickly finding an optimal solution. While not identical to any specific case, our flow routing problem with SINR sufficiency constraints has some of these characteristics. Namely, it closely

resembles a “fixed-cost” network problem as described in [56][57]. In short, a fixed-cost problem has binary variables representing whether or not nodes are contained in a possible solution. If a node is included, it comes at an associated cost. In the MILP presented in this document, if a logical link is chosen, it comes at the cost of satisfying the associated SINR condition. All beams for all nodes are enumerated, and some combination of the constraints must be chosen for feasibility. The fixed-charge problem is a specific instantiation of another phenomena, that of programming disjunctions [58][59]. Disjunctive programming is described in [60] as:

A disjunctive program is a linear program complicated by disjunctions, that is, sets of constraints of which at least one must be true. The commonest example is the fixed cost problem in which either production of an item is zero or a fixed cost is incurred.

and in [56] as:

A typical disjunctive set of constraints states that a point must satisfy at least k of m sets of linear constraints.

Our mixed-integer program is structured such that SINR constraints are relaxed with binary variables. We utilize Big-M notation as described in [61] to enforce feasibility for all constraints. Big-M constants are introduced to “expand” the problem polytope by relaxing the SINR bounds when binary indicators are not positive. Integer flow variables indicate what subset of the set of SINR inequalities are binding at a given time. Numerous papers [60][62][63] show that disjunctions can be used to tighten relaxations and remove non-optimal sections of the polytope. Disjunctions are included in our MILP model as logical constraints via our “super-flow” variables, especially relating to logical multicast topology.

2.2.2 Logic Programming

In addition to disjunctions, additional information can be encoded in the model to eliminate “don’t be stupid” constraints as described in [64]. It has been shown that logical inference can be applied to binary models to promote solution efficiency as introduced in [65]. Logic programming is best described in [62]:

Finding a good design may require examining a large number of possible alternatives. This paper uses the methods of logical inference, along with more traditional mathematical programming methods

...

to derive logical rules that any design worth examining must satisfy.

This topic is covered extensively in the book by Hooker [57] and given a brief overview in his paper [64]. We use some of the methods of logical inference to tighten the number of possible combinations in the mixed-integer linear program. For future work, we plan to investigate leveraging many more of the methods described in [57] to solve even larger network models.

2.3 Summary and Comparison

As wireless networks become ubiquitous, the need to conserve limited resources such as energy will be imperative. Clearly, a large volume of research has been dedicated to conserving expended energy when communicating wirelessly. Although the research has been broadly based in terms of applications and approaches, three major areas of interest have emerged:

- Energy Efficiency and Maximum Lifetime Metrics;
- Multicast Delivery (as a superset of other routing requirements);
- Cross-Layer Optimization.

The “broadcast” nature of wireless communications results in potential advantages for messages that are intended for multiple destinations. Initial research sought to minimize cumulative power usage in a delivery tree for broadcast applications. However, it was shown that minimum power trees are not max-min lifetime optimal due to the unbalanced drain of resources among nodes. Directional antennas have also been considered for their potential advantages. Directional transmission or reception allows a node to optimize communication with only a subset of the neighboring nodes, and to limit interference with those nodes it wishes to ignore.

From our discussion, it should be clear that the optimization methods presented later in this document are most similar to those of Guo and Yang [12] [13]. Both utilize mixed-integer programs to optimize energy consumption for multicast delivery in wireless ad-hoc networks with directional antennas. However, the MILP introduced in this document has important differences in the modeling of the communication medium and the target objective.

Whereas both of Guo and Yang’s models assume that nodes are decoupled, and ignore interference between competing sessions, our model is capable of more realistic interference (specifically side-lobe interference) modeling and antenna assumptions. We recognize that each network node may have an important role in delivering functionality, and therefore our MILP is designed for the metric of max-min network lifetime as defined in [8]. Conversely, the programs in [12] [13] are designed to find minimum energy trees, and make no mention of extensibility to other optimization goals. On the other hand, our model is carefully formulated with independent power variables so that different optimization functions could be defined. For instance, by removing our auxiliary max-min lifetime variable, and replacing it with the summation of all power variables, we are able optimize for minimum energy trees as in [12] [13].

In the next chapter we present our proposed mixed-integer linear program. The flexibility of the model is key as we analyze and challenge how directional antennas are used to conserve energy under realistic assumptions about wireless operation. We first study a new paradigm for optimizing max-min lifetime in directional antenna equipped networks by investigating the case where nodes use their antennas to *receive* only. With the exception of a cursory investigation by Takai [50], we know of no papers that consider directional *reception*. We then utilize the model to perform a head-to-head analysis of multicasting with directional transmission and directional reception. Since mixed-integer linear programs are known to scale poorly, we also study simple heuristics designed for directional transmission and directional reception. The interference modeling capability and flexibility of our model is important for this analysis. While most heuristics are based on a non-contentious model, we use our MILP to map their logical routing hierarchy into an interference scenario and into our opti-

mization space. Because heuristics do not consider inter-node coupling, they do not consider the numerous beam configurations represented in our MILP model. We take the layer three hierarchy returned by the heuristic and optimize over the remaining variable space, which includes the available beam configurations. Subsequently, we can effectively compare the metrics of the heuristic with the true optimal.

Chapter 3

Problem Statement and Approach

This chapter defines the specific problem space and presents our solution methods. First, we describe the general use of directional antennas in ad-hoc networks. While directional antennas have numerous advantages and associated research challenges, a specific subset of the problem space has been chosen for investigation. After this general introduction, our chosen metric is defined, followed by modeling assumptions and setup.

3.1 Introduction

In wireless ad-hoc networks, it is likely that wireless nodes will be equipped with finite, and relatively small energy resources. Energy reserves will not be easily replenishable, and should be conserved as best possible. In our context, a node that consumes its energy stores is no longer able to participate in the network, and can be considered “dead.”

With limited battery capacity, and the non-linear degradation of signal over large distances, transmit power composes a significant part of any power budget. Nodes that foolishly waste power will remove themselves from the network quickly. From a network perspective, as a network loses constituent nodes due to expended batteries, the overall system utility can

decrease dramatically.

For instance, sensor networks may only operate when all nodes are operational and able to deliver sensed information to the sink. The death of a single node could remove reporting functionality available only at that node. Battlefield monitoring networks are a clear example of the catastrophic results of node death. The loss of a single node could leave entire sections of perimeter unmonitored, *drastically* reducing user utility from the network. Worse, the network could be partitioned, precluding the ability to relay messages to destinations or to the sink.

Recognizing this limitation, much previous work has been dedicated to reducing power usage when communicating. In this document, we explore the possibility of using directional antennas as a method of decreasing power consumption. To facilitate our discussion, we first describe a very general model for directional “smart-antennas.” We then describe the focus of our work, which can best be described as a system-wide optimization problem for the delivery of multicasts in directional-antenna equipped networks. Finally, we present and justify our modeling assumptions regarding the channel and MAC functionality.

3.2 Directional Antennas and Ad-Hoc Networks

Directional antennas are able to enhance network operation by focusing radiation either toward, or from, a single angular direction [14]. This behavior serves a variety of purposes. For example, inter-session interference is reduced, since radiated power from a directional antenna is sent predominantly in “useful” directions. Conversely, a directional antenna may use this gain to focus power *received* from a neighbor. Moreover, the power required to traverse inter-node distances is decreased because of the gain resulting from directional transmission / reception. Or, if we consider power to be constant, the maximum distance over which nodes can communicate is extended.

Using directional antennas comes at a cost of additional complexity. When communicating

omni-directionally, transmission power is the only configurable variable to adjust routing topology. Clearly, over a homogenous channel, the received power from a transmitting node is dependent only upon the distance from the source. From a modeling perspective, received power can be expressed in terms of only inter-node distance.

Incorporating directional antennas leads to an entirely new set of optimization variables. Depending upon the width of the antenna beam and the boresight (direction of the main lobe), the gain to neighboring nodes can vary. For this document, we not only assume that the main lobe can be pointed in any direction, but that the beam lobe can be tuned to variable width.

Therefore, any routing scheme must choose (at each equipped node):

- Transmission power (P_i^t);
- Antenna main lobe direction or “boresight” (σ^i);
- Antenna main lobe width (θ^i).

3.3 Metrics

The idea of using directional antennas to enhance the operation of wireless ad-hoc networks is not new. Spatial multiplexing and the associated benefits have been investigated in previous work. While these works often differ in model and assumptions, they most often differ in the “goal” of antenna use. Some possible objectives include:

- minimum cumulative transmission power trees;
- maximum end to end throughput;
- min-max network lifetime.

In this document, we are concerned with energy conservation as it relates to network lifetime. Unfortunately, the term “lifetime” also has numerous definitions. For instance, the lifetime of a network could be the time until a specified number of nodes have depleted their battery power. Or, it could represent the time until the network becomes disconnected. In the case of dynamic networks (multiple multicasts over time), lifetime is often defined as the time until the network is incapable of delivering a multicast. An excellent description of the three main classes of lifetime-maximization problems for ad-hoc networks is provided by Shi et al. in [30]. In our work, we consider the problem of maximizing the lifetime of a single static multicast tree.

In other words, we are interested in the case of a static deployment of nodes and the arrival of a *single* multicast. Lifetime, in this case, is defined as the time until the death of the *first* node in the multicast delivery tree. Then, as defined in [18], path lifetime can be expressed in terms of node power variables P_i^t , remaining energy stores E_i^r , and the set of nodes in the multicast tree \mathcal{M} as:

$$Life = \min_{v_i \in \mathcal{M}} \frac{E_i^r}{P_i^t} \quad (3.1)$$

The problem then, is to find the combination of tree topology, beam choice, and power settings that leads to the routing tree with the longest overall lifetime. This expression will avoid overly taxing any individual node, and also avoid nodes with limited energy reserves.

3.3.1 Justification of Metric

The goal of this work is to develop tools for modeling directional communication in ad-hoc networks and to perform an algorithmic analysis of their use and performance. It has been shown in previous work, and later in this document, that minimum power and maximum lifetime metrics do not correlate well [24] due to the fact that power minimization often leads to an uneven drain rate at individual nodes. Uneven power consumption results

in the death of some network nodes quickly. Further, the poor correlation results in the possibility that even when found, a multicast tree determined by min-power metrics may become disconnected in a short amount of time because of node death. Upon failure of a node in the tree, the tree is no longer feasible, and a new solution must be found. The assumption in this work is that multicast duration will be non-trivial, and that maximizing the worst-case node lifetime will be an important consideration to fairly utilize all available energy reserves.

However, we do recognize the importance of metrics such as cumulative power, and track them throughout our results section. Though we study only a single, static multicast, metrics such as cumulative power use allow us to infer performance for future multicasts. For example, our analysis shows that directional reception uses *much* less cumulative power than directional transmission for identical network layouts. We would expect that since directional reception leaves more energy at nodes in the network, it would be better able to deliver future multicasts.

3.4 Communication and Antenna Model

In this section our communication model is discussed in detail. Directional antenna research relating to energy conservation and routing, and especially research dedicated to multicast delivery has often ignored interference. A major strength of the model in this document is the ability to incorporate side-lobe interference and inter-node competition for the medium.

3.4.1 MAC

For the entirety of this document we assume an imperfect or probabilistic MAC scheme. That is, for an $i \rightarrow j$ session it is possible that any (or all) ongoing non- $i \rightarrow j$ sessions might interfere. For two nodes to be guaranteed a connection, the signal-to-noise ratio (SINR) at

the receiver from the transmitter must exceed a predetermined value. When both of these considerations are combined, the power setting required for $i \rightarrow j$ communication must be sufficient to overpower the interference on ongoing sessions *and* surpass the threshold requirement.

MAC schemes are intended to *efficiently* mediate access to a wireless medium. For example, current algorithms such as 802.11 use an exchange of RTS/CTS packets to “reserve” a neighborhood of the network for communication. Doing so prevents channel contention in local neighborhoods and avoids situations such as the “hidden-node” problem [43]. Allocating medium access (and preventing problems such as “hidden-node”) becomes more difficult when directional antennas are available in the network, and a MAC scheme must be specially designed to operate with them. Even if specially designed for directional antennas, the MAC is unlikely to take care of effects such as co-channel interference due to sidelobes.

Finally, the purpose of this document is to analyze the performance of directional antennas as it pertains to energy efficiency and spatial multiplexing. Therefore, we focus on the use of directional antennas to spatially multiplex the wireless channel. We assume that the MAC is allocating access to the medium as efficiently as possible, utilizing the ability of directional antennas to focus transmitted or received radiation and prevent interference.

Therefore, since nodes are required to “overpower” interference in the network to achieve communication, we avoid the disparity between communication distance and interference distance. Whereas previous work must cope with the fact that interference distance is larger than communication distance, our work recognizes that directional antennas may overcome this limitation by multiplexing spatially.

3.4.2 Channel Model

In order to keep the model manageable, we make some assumptions regarding wireless propagation. We do not consider multi-path, fading, or shadowing. Rather, we assume that

inter-node path losses are homogenous between nodes.

Also, it is assumed that inter-node propagation characteristics are identical. Path loss due to inter-node distance is directly proportional to the distance to some power, where $2 \leq \alpha \leq 4$. Although the model presented in this document can represent different path characteristics, the beam enumeration techniques illustrated here mandate homogeneity for the grouping of neighbor nodes.

3.4.3 Antenna Model




Figure 3.1: Beam Illustration

Antenna hardware and associated models range widely in the propagation pattern and performance characteristics. In this work, we consider two major antenna classes shown in Figure 3.1. Omni-directional transmission or reception is shown in Figure 3.1 (a). The signal from sender S is sent or received equally in all angular directions. Steered-beam antennas in Figure 3.1 (b) have a single main beam lobe. While multiple beams are possible, the DSP power costs and antenna complexity requirements are often prohibitive. We therefore assume only a single, tunable beam. Throughout this work, we assume that the DSP power cost associated with single-lobe directional communication is negligible when compared with transmission power. Therefore, we ignore the additional power requirement of using smart-antennas for transmission or reception. Our modeled antennas can vary the boresight of the beam (σ) and the beamwidth (θ) with infinite granularity ($\theta \in [\theta_{min}, 360^\circ]$),

$\sigma \in [0^\circ, 360^\circ]$). Beamwidth is lower bounded at a given value θ_{min} . Further, as in [16] and shown in Equation 3.2, we assume that beam gain (G) is inversely proportional to beamwidth θ .

$$G = \frac{360^\circ}{\theta^\circ} \quad (3.2)$$

Later in the document, we discuss a method of discretizing the beams to generate a finite set of feasible beams for the model. The granular antenna parameters lead to a large number of potentially useful beam configurations, all of which must be considered.




Figure 3.2: Bulb-and-Cone Antenna Model

Additionally, we consider sidelobe interference via the bulb-and-cone model shown in Figure 3.2. The main lobe is shown again with a beamwidth θ , and sidelobes are abstracted into a uniform circle surrounding the source. A static percentage of the radiated/received power is “in-beam” (P_{ct}^{inbeam}). The remaining power is radiated/received through the sidelobe bulb.

For directional transmission, we assume that a node is transmitting directionally and receiving omni-directionally. The converse applies for directional reception. Theoretically, it is possible for both the transmitter and receiver to utilize directional antennas simultaneously. However, as discussed in Chapter 2, timing and alignment become extremely difficult in these cases. Furthermore, as previously stated, smart-antenna directionality comes with an associated power cost. In this work, we assume this power requirement to be negligible when compared to transmission power, and therefore it is ignored. Our assumptions in this area

are consistent with previous research [11][66][16][12][30][2][13].

3.5 Solution Methods

In the following sections, we present solution methods for determining optimal max-min lifetime trees for multicast delivery in directional antenna equipped ad-hoc networks. The majority of the work presented is centered on a new mixed-integer linear program (MILP). Mixed-integer mathematical programs can be difficult to solve, but come with an important characteristic: if found, the solution to a mathematical program is *optimal*. We use this optimality guarantee to objectively compare different models, and to design and analyze computationally less expensive heuristics that approximate the MILP.

3.5.1 Mixed Integer Linear Programs

Mathematical programming provides a framework through which we can determine a route that is optimal for our chosen metric. Though the approach of using such techniques for determination of optima is not new [2][13][12], the novelty of our work lies in the structure of our proposed programs. The papers cited above are based on network models that ignore inter-node and side-lobe interference. Rather, they assume that interference has a negligible effect, or that a method exists to prevent it (MAC scheme, orthogonal channels/codes). As such, previous work models the network as a directed graph where inter-node distances alone determine power levels necessary for communication. Generally, binary variables are attached to all possible network links. A link that is chosen is assigned the value of 1, others 0. This approach constructs a more straightforward and seemingly intuitive mixed-integer program. However, it also completely neglects inter-node coupling and competition for the medium. In a deployment with a non-perfect MAC coordination scheme, nodes that are not immediate logical neighbors will still interact as transmitted radiation spreads throughout the network.

Symbol	Definition
\mathcal{N}	Set of network nodes.
\mathcal{D}	Set of nodes with directional antennas.
\mathcal{R}	Set of nodes that are receivers.
\mathcal{B}_i	Set of beams available at node i .
$\mathcal{B}_i(j)$	Set of beams at node i where j is within main lobe.
s	Source node.
P_i^t	Transmission power level at node i
$P_{i,b}^t$	Transmission power, node i using beam b .
\vec{G}_i	Gain / path loss vector at node i .
$F_{i,j,b}$	1 if flow possible from i to j when i using beam b .
$F_{i,j}$	total flow indicator from node i to node j (“super-flow”).
$M_{i,j}$	Message (information) flow from i to j .
$B_{i,b}(k)$	Beam gain from node i to node k , using beam b ($b \in \mathcal{B}_i$)
$U_{i,j,b}$	Bound of interference at j when receiving from i using beam b .
Q	Large integer (Big-M).
S^i	SINR ratio required at node i .
N_t	Thermal / ambient noise.
$E_{i,j}^r$	Signal power received at node j from node i .
R_i	Energy remaining (battery) available at node i .
P_{max}	Maximum transmit power setting for nodes.
Pct^{inbeam}	Fraction of power in main beam lobe.
θ_{min}	Minimum beamwidth for a node.
D-TX	Directional transmit MILP model.
D-RX	Directional receive MILP model.

Table 3.1: Notation.

We present a model based on a signal-to-noise and interference sufficiency requirement. Our directional transmit (D-TX) and receive (D-RX) mixed-integer programs optimize over possible beam configurations, while also accounting for inter-node coupling. For clarity, the model can be thought of as incorporating three distinct components: beam choice, power settings, and logical topology.

SINR Constraint and Interference Modeling

The principle behind the signal-to-noise and interference sufficiency requirement is identical for both the directional transmit and directional receive models. First, we discuss the

requirement in general to familiarize the reader with the assumptions and encoding of information. Notation directly utilized in each model is omitted in this general introduction. After discussing the methods we use for beam enumeration, the SINR requirements are translated into our MILP notation.

A simple signal-to-noise and interference requirement for feasibility of communication from node i to node j is expressed in Inequality 3.3.

$$\frac{E_{i,j}^r}{\sum_{\forall k \neq j, i} E_{k,j}^r + N_{thermal}} \geq S^j \quad (3.3)$$

Put simply, the power gathered at the receiver from the intended source as a ratio to all other received power must be above some minimal requirement. The denominator of Inequality 3.3 represents the received power from all non i nodes, along with thermal noise in the receiver. Because of non-negativity of the denominator, Inequality 3.3 can be manipulated into the single-line format in Inequality 3.4.

$$\underbrace{E_{i,j}^r}_A - S^j \underbrace{\left(\sum_{\forall k \neq j, i} E_{k,j}^r + N_{thermal} \right)}_B \geq 0 \quad (3.4)$$

Clearly, Inequality 3.4 is not feasible when interference and noise (B) exceeds received transmitted power (A). However, by placing the maximum amount that Inequality 3.4 can be violated (U) on the right hand side, feasibility is guaranteed. This represents the use of the “Big-M” technique [67]. Equation 3.5 and Inequality 3.6 show this relationship by replacing numerator and denominator with simple notation A and B respectively.

$$U = |\min_{\forall A, B} \{A - B\}| = \max\{B\} \quad (3.5)$$

$$A - B \geq -U \quad (3.6)$$

Although Inequality 3.6 is consistently feasible, the introduction of the maximum violation has modified the logic, and is not equivalent to Inequality 3.4. Therefore, the model cannot determine if the *original* constraint was met. We introduce a binary variable I to rectify this in Inequality 3.7.

$$A - B \geq U \cdot I - U \quad (3.7)$$

$$I \in \{0, 1\} \quad (3.8)$$

When $I = 1$, Inequality 3.7 reduces to the original form of Inequality 3.4. Thereby, only when $I = 1$ is the signal-to-noise ratio is sufficient for communication. In our model, inter-node path loss, antenna gain characteristics, and P_{max} are assumed known. It follows that we may compute the maximum possible non- i, j interference ($U_{i,j,b}$) for an i, j transmission. The structure of each SINR sufficiency constraint is therefore consistent with Big-M binary programming notation [61]. Note that a sufficiently large value may also be “guessed.” However, it is well known that tightening the constraints of a mixed-integer linear program reduces solve times.

Beam Enumeration

For both D-TX and D-RX, beam configurations are discretized for problem setup. Each node maintains a finite set of all possible beams, denoted (\mathcal{B}). Also, a matrix of actual gain values is recorded (B), where $B_{i,k}(j)$ represents the numeric value for inter-node gain $i \rightarrow j$ when node i chooses beam k . Note that this notation holds for both models if node roles are transposed. For D-TX, node i chooses beam k for transmit, and therefore the power received at j is equal to the gain value multiplied by node i 's transmit power. For D-RX, node i chooses beam k to *receive*, and the received power at i is equal to the gain value multiplied by node j 's transmit power. Thereby, the received power at j from i when i is

using beam k is equal to the multiplication of i 's transmit power by $B_{i,k}(j)$. Equation 3.9 shows the expression used to compute gain (including inter-node path loss) from node i using beam k to node j . Here, $r_{i,j}$ represents the distance between nodes i and j , and α represents the path loss exponent (typically, $2 \leq \alpha \leq 5$). As in Table 3.1, $\mathcal{B}_i(j)$ is the set of beams at node i that contain node j in their main lobe.

$$B_{i,k}(j) = \begin{cases} \frac{1}{r_{i,j}^\alpha} \cdot (1 - P_{ct}^{inbeam}) & k \notin \mathcal{B}_i(j) \\ \frac{360}{r_{i,j}^\alpha \cdot \theta_i} \cdot P_{ct}^{inbeam} & k \in \mathcal{B}_i(j) \end{cases} \quad (3.9)$$

• D-RX Gain Computation

A node with a steered-beam directional antenna may point the main lobe in any angular direction. The D-RX model utilizes the following observation: a node will seek to maximize **only** the SINR ratio of the intended source. Contrast this with directional transmission. While *transmitting* directionally, a node may wish to cover any number of angularly contiguous neighbors. However, while *receiving*, listening to more than one transmitter will only increase interference. A node may point the antenna lobe at whatever subset of neighbors leads to the best SINR for the logical link in question. Thus, there is no need for the antenna to have variable beamwidth.

To allow the solver to choose the global optimal topology, we must consider situations where node positions cause beam overlaps. In the event that two neighboring nodes cannot be isolated in a single beam, a beam for each combination is recorded.

Consider Figure 3.3, where node i can listen to neighboring nodes 1, 2, or 3. In configuration (a), node i has suitably narrow beamwidth that it can focus on any of the three sources individually. However, in the case that neighbors are too close to discern with a single beam, as shown in Figures 3.3 (b) and (c), the model simply has multiple entries for the node that cannot be isolated. Figure 3.3 (a) exemplifies the case of one beam per node, (b) and (c)

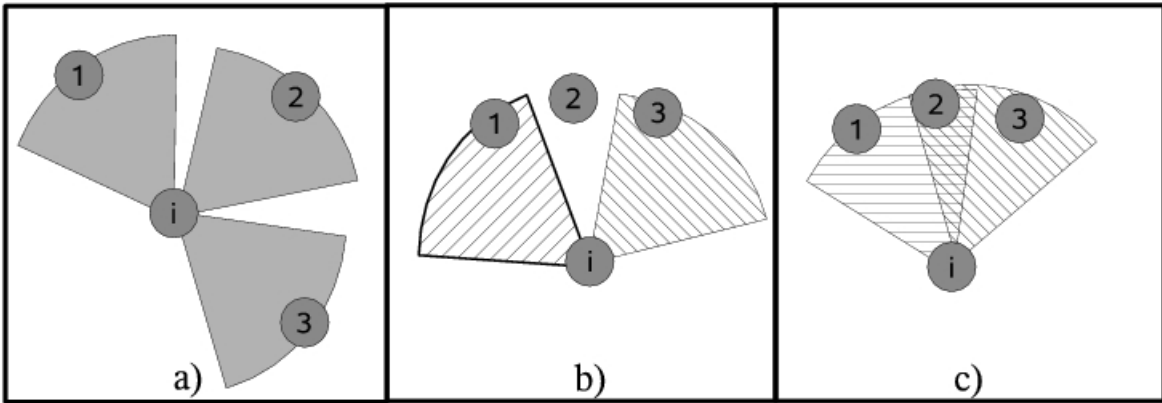


Figure 3.3: D-RX Beam Configurations

illustrate how beam overlaps are incorporated.

As shown in Figure 3.3 (b), node i has identical beam sets for nodes 1 and 3 as those in Figure 3.3 (a). That is, there is a single beam entry for both nodes 1 and 3. Figure 3.3 (c) shows that node i can cover node 2 with the following neighbor sets: $\{1, 2\}$ or $\{2, 3\}$. Since either node 1 or 3 may be transmitting on another logical link, both beam configurations must be included in the model so that node i can avoid interference.

In the worst case, it is possible that D-RX has $O(|\mathcal{N}|^3)$ beam configurations for an entire network model. For many cases, and especially those with narrow beamwidth, however, this number will be closer to $O(|\mathcal{N}|^2)$.

• **D-TX Gain Computation**

For the D-TX case, beams are enumerated based on the set of nodes included in the main lobe. A new beam is generated for each subset of neighbor nodes that are covered by the beam cone to exploit the “*wireless multicast advantage*” [16]. Below, we illustrate that the number of possible beams at a single node is bounded by $O(|\mathcal{N}|^2)$, leading to $O(|\mathcal{N}|^3)$ for the entire network model.

D-TX must also be constrained to choose only one beam at any transmitting node. In the reception model (D-RX), we are certain that a receiving node will cover as few transmitting nodes as possible to increase SINR as per the *directional reception observation*. However, since D-TX would like to have multiple receivers per beam, we must prevent the case where a single node attempts to choose multiple transmitting beams with high gain. Additional auxiliary variables and constraints are introduced into D-TX to prevent this condition. New individual continuous power variables $P_{i,b}^t$ are added for each beam at node i . An additional large integer Q is summed with the maximum violation term $U_{i,j,b}$. Penalty terms are then added to each SINR inequality such that only one $P_{i,b}^t$ may be positive, shown by example below.

Assume that only node i is transmitting ($P_k^t = 0 : k \in \mathcal{N} \setminus i$). Inequality 3.10 is then an $i \rightarrow j$ SINR inequality at i , complete with terms introduced above. Inequality 3.11 indicates that for any SINR inequality to hold at i , the sum of all non- b beam powers is upper bounded by an extremely small value. In most cases, other nodes will also be transmitting, limiting this value further. For all experiments in this dissertation we set $Q = 10,000$ and $P_{max} = 100$.

$$P_{i,b}^t \cdot B_{i,b}(j) - \frac{Q}{P_{max}} \cdot \sum_{\substack{\forall q \in \mathcal{B}_i \\ q \neq i,j}} P_{i,q}^t \geq 0 \quad (3.10)$$

$$\sum_{\substack{\forall q \in \mathcal{B}_i \\ q \neq i,j}} P_{i,q}^t \leq \frac{P_{max}}{Q} P_{i,b}^t \cdot B_{i,b}(j) \quad (3.11)$$

$$\begin{aligned}
& \mathbf{min} \quad \mathbf{T} \\
& \mathbf{s.t.} \\
& P_{i,b}^t \cdot B_{i,b}(j) - S^i \cdot \left[\sum_{\substack{k \in \mathcal{N} \\ k \neq i,j}} \sum_{l \in \mathcal{B}_k} P_{k,l}^t \cdot B_{k,l}(j) + N_t \right] - \frac{Q}{P_i^{max}} \cdot \sum_{\substack{k \in \mathcal{N} \\ k \neq i,j}} P_{i,m} \geq & : \quad \forall i, j \in \mathcal{N}, \forall b \in \mathcal{B}_i, \\
& & : \quad j \neq s, i \neq j \\
& F_{i,j,b} \cdot [U_{i,j,b} + Q] - [U_{i,j,b} + Q] & \\
& F_{i,j} = \sum_{\forall b \in \mathcal{B}_i} F_{i,j,b} & : \quad \forall i, j \in \mathcal{N}, i \neq j \\
& F_{i,j} + F_{j,i} \leq 1 & : \quad \forall i, j \in \mathcal{N}, i \neq j \\
& \sum_{\substack{\forall k \in \mathcal{N} \\ k \neq j}} F_{k,j} \geq 1 & : \quad \forall j \in \mathcal{R} \\
& M_{i,j} \leq |\mathcal{R}| \cdot F_{i,j} & : \quad \forall i, j \in \mathcal{N}, i \neq j \\
& M_{i,j} \leq \sum_{\substack{\forall k \in \mathcal{N} \\ k \neq i,j}} M_{k,i} - 1 \cdot (1 : i \in \mathcal{R}) & : \quad \forall i, j \in \mathcal{N}, i \neq j \\
& P_i^t \geq P_{i,b}^t & : \quad \forall i \in \mathcal{N}, \forall b \in \mathcal{B}_i \\
& T \geq \frac{P_i^t}{R_i} & : \quad \forall i \in \mathcal{N} \\
& P_i^t \leq P_{max} & : \quad \forall i \in \mathcal{N} \\
& F_{i,j,b} \in \{0, 1\} & : \quad \forall i, j \in \mathcal{N}, i \neq j, \\
& & : \quad j \neq s, b \in \mathcal{B}_i
\end{aligned}$$

Figure 3.4: D-TX: Formulation for Max-Min Network Lifetime with Inter-Node Interference

D-TX MILP Model

Figure 3.4 shows the entire model translated into the MILP notation in Table 3.1. The first inequality represents each inter-node’s SINR constraint. Interfering terms, beam penalty terms, and Big-M integers are shown. The second inequality sums all integral beam variables at a single node. These “super-flow” variables simplify throttling actual flow, preventing loops, and serve as effective cutting planes for the solution space. In effect, the “super-flow” variables are representative of the logical topology used throughout the network. The third and fourth expressions prevent one-hop loops and mandate that all receiving nodes have

one positive incoming link, respectively. Continuous variables M are introduced to track the flow of messages throughout the network in the fifth and sixth inequalities. The fifth ensures that $i \rightarrow j$ flow is allowed only when SINR is sufficient, and the sixth allocates a unit of flow to all receivers, while also conserving flow units. Finally, auxiliary power variables return the maximum power used over all beams. The remainder of the problem generates the reciprocal of the max-min lifetime value $\frac{1}{T}$, caps maximum power, and sets variable integrality. Note that this formulation reflects the goal of optimizing max-min network lifetime. The optimization of other metrics are easily formulated. For instance, to compute minimum cumulative power trees, remove the auxiliary variable T and replace it in the objective function with $\sum_{vi \in \mathcal{N}} P_i^t$.

D-RX MILP Model

Figure 3.5 shows the optimization model for directional reception translated into the MILP notation in Table 3.1. The format is similar to that of the D-TX model in structure, except that the single-beam constraints are omitted.

MILP Complexity

It is well known that the difficulty in finding a solution to a MILP problem is most often directly related to the number of binary variables in the model. An accurate D-TX model will be extremely complex, due to the combinatorics of beam enumeration discussed above. The upper limit on possible beams is illustrated by example below and referencing Figure 3.6.

Denote the transmitting source as node s , and the set of neighbors as \mathcal{N} , and the set of neighbors covered by a beam as \mathcal{B} . Assume that s 's neighbors are sufficiently spaced so that a beam of width θ_{min} may cover one node at a time (no overlap). Without loss of generality, order s 's neighbors by increasing σ_i . Denote δ_i as the beam setting covering angular neighbors i and $i + 1$. Because we assume s has infinitely granular beamwidth, s

min T

s.t.

$$\begin{aligned}
P_i^t \cdot G_i(j) \cdot B_{j,i}(b, i) - S^i \cdot \left[\sum_{\substack{k \in \mathcal{N} \\ k \neq i, j}} P_k^t \cdot G_{k,j} \cdot B_{j,i}(b, k) \right] &\geq F_{i,j,b} \cdot U_{i,j,b} - U_{i,j,b} &: i, j \in \mathcal{N}, b \in \mathcal{B}_{j,i}, \\
&&: j \neq s, j \neq i \\
\sum_{\forall b \in \mathcal{B}_{j,i}} \sum_{\substack{\forall i \\ i \neq j}} F_{i,j,b} &\geq 1 &: \forall j \in \mathcal{R} \\
\sum_{\forall b \in \mathcal{B}_{j,i}} F_{i,j,b} + \sum_{\forall c \in \mathcal{B}_{i,j}} F_{j,i,c} &\leq 1 &: \forall i, j \in \mathcal{N} \\
M_{i,j} &\leq (|\mathcal{N}| - 1) \cdot \sum_{\forall b \in \mathcal{B}_{j,i}} F_{i,j,b} &: \forall i, j \in \mathcal{N}, i \neq j, j \neq s \\
M_{i,j} &\leq \sum_{\substack{\forall k \\ k \neq i}} M_{k,i} - (1 \text{ if } i \in \mathcal{R}, 0 \text{ else}) &: \forall i \in \mathcal{N}, i \neq j, j \neq s, i \neq s \\
T &\geq \frac{P_i^t}{R_i} &: \forall i \in \mathcal{N} \\
F_{i,j} &\in \{0, 1\} &: \forall i, j \in \mathcal{N}, i \neq j \\
P_k^t &\leq P_{max} &: \forall i \in \mathcal{N} \\
F_{i,j,b} &\leq 1 &: \forall i, j \in \mathcal{N}, i \neq j, \\
&&: j \neq s, \forall b \in \mathcal{B}_{j,i} \\
T, M_{i,j}, P_k^t &\geq 0 &: \forall i, j \in \mathcal{N}, i \neq j, \forall k \in \mathcal{N}
\end{aligned}$$

Figure 3.5: D-RX: Formulation for Max-Min Network Lifetime with Inter-Node Interference

may cover any subset of contiguous neighbors.

1. When $|\mathcal{B}| = 1$, the beam is pointed directly at the receiving node, for $|\mathcal{N}|$ choices.
2. When $|\mathcal{B}| = 2$, s can choose any 1 of the δ_i arcs, for $|\mathcal{N}|$ choices.
3. When $|\mathcal{B}| = 3$, s chooses any 2 of the δ_i arcs, for $|\mathcal{N}|$ choices.
4. \vdots
5. When $|\mathcal{B}| = |\mathcal{N}|$, there is one optimal choice.

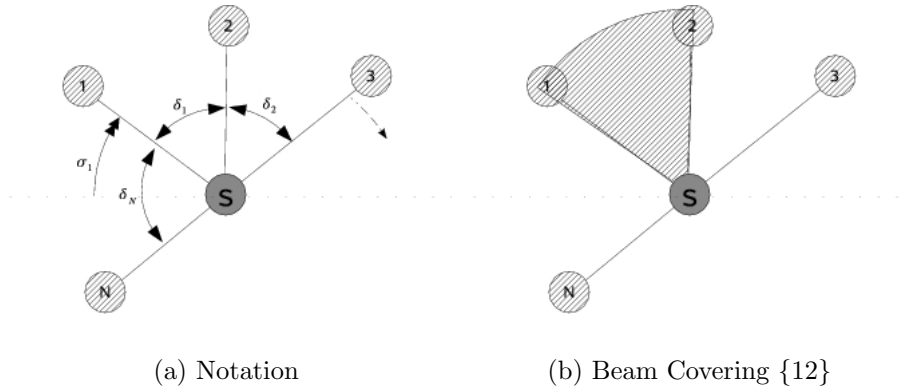


Figure 3.6: D-TX Beam Limit Illustration

Therefore, the number of beams is upper bounded by $|\mathcal{N}|(|\mathcal{N}| - 1) + 1$ or $O(|\mathcal{N}|^2)$ beams.

Please refer to Table 3.2 for a summary of the complexity bounds for both D-TX and D-RX. As shown above, D-TX nodes can maintain $O(|\mathcal{N}|^2)$ individual beams, for a total of $O(|\mathcal{N}|^3)$ for the network overall. While in the D-RX model we are sure that a receiver will seek to maximize gain to a single transmitter, any node in D-TX may wish to transmit to any subset of neighbor nodes. For completeness, we must even consider the situation where a neighbor node is not in the main beam lobe. Rather, it is possible (albeit unlikely) that communication is intentionally taking place via the antenna sidelobes. Therefore, $|\mathcal{N} \setminus \{s\}| - 1$ binary variables are required *per beam*, increasing the total number to $O(|\mathcal{N}|^4)$.

D-TX:

Type	Bound	Main Contributor
Constraints	$O(\mathcal{N} ^4)$	SNR
Binary Variables	$O(\mathcal{N} ^4)$	$F_{i,j,b}$
Continuous Variables	$O(\mathcal{N} ^3)$	$P_{i,j,b}$

D-RX:

Type	Bound	Main Contributor
Constraints	$O(\mathcal{N} ^3)$	SNR
Binary Variables	$O(\mathcal{N} ^3)$	$F_{i,j,b}$
Continuous Variables	$O(\mathcal{N} ^2)$	$M_{i,j}$

Table 3.2: Complexity of MILP Formulations

MILP Summary

In the preceding sections of this chapter, we introduce a mixed-integer linear program that is built upon a SINR sufficiency constraint. The inclusion of the SINR constraint provides modeling flexibility, and allows the program to effectively track interference and medium contention. Two sub-models are based on these initial assumptions, one handling the situation where nodes transmit directionally and receive omni-directionally (D-TX) and the converse case (D-RX). Both models are discussed in detail, and the complexity of each is illustrated. Noteworthy is the fact that the D-RX case has significantly fewer beam combinations because of the *“directional reception observation.”*

While mixed-integer linear programs are known to produce optimal solutions, they also do not scale well. This is especially true as the number of binary variables grows. Therefore, computationally inexpensive heuristics are used to generate solutions that are “close” to the optimal. There are many heuristics for the directional transmission multicast problem solved by D-TX. However, to our knowledge, there are none that consider the directional reception case as modeled by D-RX. In the remainder of this chapter, we introduce and discuss three heuristic methods for approximating the solution to the defined mixed-integer program.

3.6 Heuristic Methods

Mixed-integer programs are known to be increasingly difficult to solve as the number of integer variables grows. Therefore, computationally cheap heuristics are used to approximate the optimal solution of such a problem. While heuristic schemes for directional transmission have been introduced, we are unaware of any such schemes for directional reception. Firstly, we introduce our heuristic for multicast routing with directional reception, the Directional Reception Incremental Protocol (DRIP). DRIP is a centralized technique that takes into account a node’s energy reserves when attempting to build a lifetime optimal multicast tree. The tree is built incrementally, and a maximum weight link is chosen at each step. Later, we




Figure 3.7: DRIP Link Cost

introduce a semi-distributed heuristic technique with the same purpose. Instead of requiring global information about the network, this method takes advantage of the flexibility afforded by directional reception. Nodes “volunteer” to assist in the forwarding of a multicast, taking on the load of possible bottleneck links and extending network lifetime.

3.6.1 DRIP

As mentioned in the introduction, DRIP attempts to find a near-lifetime-optimal routing tree for the given network. To explain the algorithm, we introduce some notation and briefly describe the link weighting method. The network is modeled as a directed graph \mathcal{G} where an $i \rightarrow j$ link $l_{i,j}$ has a weight $c_{i,j}$ assigned as in Equation 3.12.

$$c_{i,j} = \frac{R_i \cdot 360^\circ \cdot P_{ct}^{inbeam}}{r_{i,j}^\alpha \cdot \theta_{min}^j \cdot S^j} \quad (3.12)$$

Link weight is therefore proportional to the amount of battery power remaining at the transmitting node, denoted by R_i , and inversely proportional to the power required for inter-node communication. All network links are represented by \mathcal{L} , where $\mathcal{L}(i)$ denotes all links with node i as a *destination*. A forwarding tree \mathcal{T} is built from source to all receivers incrementally, until either all receivers are covered or all network nodes are included in \mathcal{T} . At each step, a link with *maximum* lifetime to a non- \mathcal{T} node is chosen. Once the algorithm

terminates, post-processing eliminates any branches that do not contain receivers.

Algorithm 1 DRIP Multicast Routing Algorithm

```

1:  $\mathcal{T} \leftarrow s$ 
2:  $\mathcal{N} \leftarrow \mathcal{N} \setminus s$ 
3:  $\mathcal{L} \leftarrow \mathcal{L} \setminus \mathcal{L}(s)$ 
4: while  $! \mathcal{R} \subseteq \mathcal{T}$  do
5:    $i \leftarrow \text{maxLifetimeToNode}(\mathcal{L})$ 
6:    $\mathcal{T} += i$ 
7:    $\mathcal{L} \leftarrow \mathcal{L} \setminus \mathcal{L}(i)$ 
8: end while
9:  $\text{removeUnNeededBranches}(\mathcal{T})$ 

```

DRIP does not immediately take interference into account when selecting routes. It does, however, specifically tie link weight to the amount of remaining battery power at a node. Also, note the similarity in DRIP to similar incremental protocols such as D-MIP. The cost of adding an extra node is the use of extra power at a node *already included in the tree*. As opposed to D-MIP, which must increase power or beamwidth to cover other nodes, DRIP allows potential receivers to focus a narrow beam at upstream neighbors. Later in this document, we show that though directly ignorant of interference, DRIP performs fairly well even in interfering environments.

3.6.2 Semi-distributed Algorithm

In this section, we introduce our semi-distributed heuristic method for finding near-optimal multicast routing trees. Methods such as DRIP require global knowledge of the network layout, and are therefore potentially unrealizable. Our semi-distributed protocol takes advantage of the fact that nodes *receiving* directionally can more easily act as distribution centers for more distant nodes. Receivers can make local decisions as to which node represents the most desirable upstream neighbor.

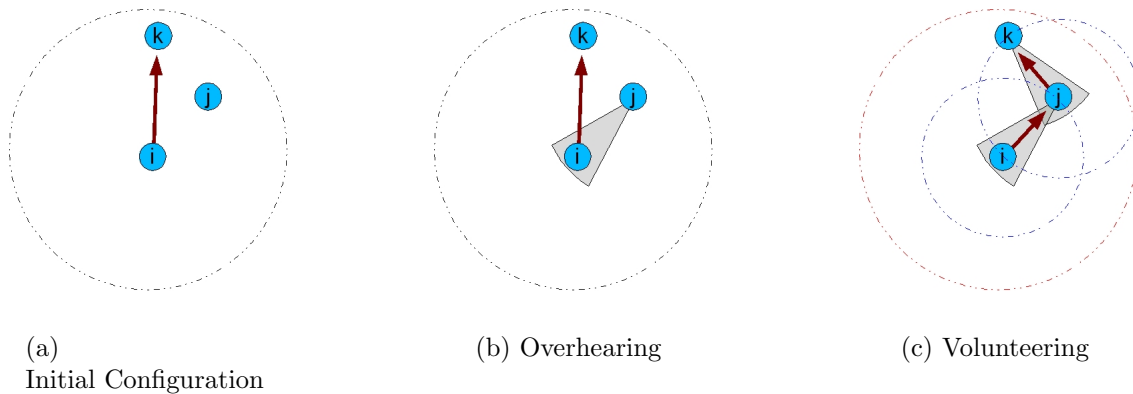


Figure 3.8: Overhearing

Design Observations

Before delving into the specifics of the algorithm, we discuss this observation as it is central to the operation of the heuristic. Refer to Figure 3.8(a). In this simple example, node i is transmitting at a power sufficient to reach node k . Assuming no interferer in the same angular direction, node j receives a signal from i at a level that is more than sufficient for communication, as shown in Figure 3.8(b). If it is so inclined, node j can “volunteer” to act as an intermediate hop between nodes i and k . Assume that node i now communicates with node j as a hop towards k . Node i can now reduce its transmitter power setting. The children of node i (along with others) can then consider using node j as their next upstream node.

A node who volunteers can limit “interested children” by only advertising at a power level such that it does not become the bottleneck node for the network. (Obviously, it makes little sense for node j to begin transmitting at a level that would decrease overall network lifetime.) If a child node chooses to use the new hop as its upstream neighbor, the power levels then decrease to meet the requirements of this new logical topology. Figure 3.8(c) illustrates that node i has lowered its transmitted power rate, utilizing node j .

In the implementation discussed in this paper, after a node volunteers, *all* other receivers are iteratively given the opportunity to use the new intermediary node. Although this may

inspire skepticism about implementation, in reality, receivers will only target the new node if received power makes it an appealing alternative. Therefore, by advertising at a power level that does not make the node the bottleneck, it will not gather any children that will result in a decrease of lifetime.

Power Control Algorithm

Proper operation of the semi-distributed algorithm requires that a child node be able to find an antenna setting for maximum gain to the intended parent, and that the parent find an efficient power setting for all children. Unlike directional transmission, this process is *entirely local at the receiver node* for directional reception. Assuming that power settings are consistent, a node will simply modify its beam settings to maximize gain.

We assume that every node uses this procedure to find the beam with maximum utility to the intended parent. Once each node completes this process, power settings throughout the network can be iteratively modified to meet SINR requirements with as little wasted power as is feasible. To determine the amount of power necessary, a source incrementally decreases its transmission power until a child is dropped. At this point, it reverts to the last feasible power setting. This process is shown diagrammatically in Figure 3.9.




Figure 3.9: Iterative Power Algorithm

The process of iteratively decreasing power settings requires that all nodes be connected at iteration 0. Moreover, the step-size can dramatically influence convergence times and

granularity of power settings. In our implementation, power is updated according to a difference equation as in [68].

At each iteration, children find the best beam to their parent, then the child with the worst SINR ratio is found (representing the farthest receiver). Assume node i is the transmitting source, and the worst child is node j , where node j is using beam b . Node i will then use Equation 3.13 to find appropriate power setting for a link to node j at time τ . $SINR(i, j)$ is computed from our beam gain matrix (in implementation, this would come from DSP hardware measurements). Equation 3.13 shows that the power at time τ is simply the difference between the required SINR value at j and the actual received power, divided by the inter-node gain. Finally, ω serves as a damping coefficient to prevent oscillations. The algorithm terminates when the SINR for the worst-case child is sufficiently close to the SINR threshold. The power algorithm is documented in pseudocode in Appendix A.

$$P_i^t[\tau] = P_i^t[\tau - 1] + \omega \frac{(SINR^{req} - SINR(i, j)[\tau - 1])}{B_{i,b}(j)} \quad (3.13)$$

It is also possible that there exists no feasible beam combination or power setting for the given configuration. In this case, the algorithm would continually increase the power settings, and infeasible configurations are detected by power settings that are above our allowed cap value.

Semi-Distributed D-RX Algorithm Overview

The algorithm is described in general terms in the following paragraph, and is shown in pseudocode in Appendix A. At iteration 0, we assume that the source is the only transmitting node in the network. Note that SINR for communication to the receivers is not expressly required, only that the source be the only node transmitting, and that it be transmitting at maximum power.

At each stage of the algorithm, all non transmitting nodes iteratively look to a transmitting node who is upstream from which they receive surplus transmission power. This node can then advertise its intention to act an intermediate hop for the upstream node. All receiving nodes for whom a link to the intermediate node does not form a loop are allowed to make a determination of whether or not to use the node as their new upstream neighbor. New logical configurations are tested when the aforementioned power algorithm is applied to the new topology, returning near-optimal power settings (or infeasibility) for the tree. A new node is added only if network lifetime can be increased through its use. This decision process is the cause of the “semi-distributed” moniker, since we require the use of a lifetime oracle for determination of the node’s contribution to overall network metrics. After node addition, all receivers are again given the opportunity to now use the new node as their upstream neighbor. This second reorganization allows other receivers to gain additional utility from the addition of the single node. If no new node is added, then the algorithm terminates.

3.6.3 Genetic Algorithm / Metaheuristic

Genetic algorithms (GA) are well known metaheuristic techniques designed to approximate solutions for hard problems. First introduced by John Holland [69], they draw their inspiration from biological evolutionary processes, where traits are passed from parents to offspring via genetic material.

In the parlance used in our implementation package [70], a single characteristic of a solution is encoded into a gene. A sequence of genes taken together is a genome, and represents a single problem solution. An objective function is defined that evaluates genomes based on the desirability of the solution they contain. Upon GA startup, an initial population of genomes is generated, and is modified between generations (iterations) through two major processes, crossover and mutation. Crossover takes two (or more) genomes, and combines their genes together to produce offspring. This process is intended to mimic sexual reproduction in that subsets of genes from each parent are represented in the child genome. Mutation, on the

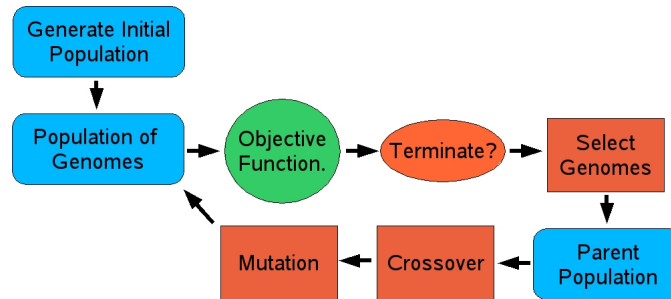


Figure 3.10: Genetic Algorithm Overview

other hand, probabilistically changes the value of a single gene within a genome to another random genome value.

In short, there are three main components of the operation of a Genetic Algorithm, as illustrated in Figure 3.10. First, individuals (genomes) are evaluated via an objective function, a new population is selected based upon this evaluation, and finally members of the population combine (mate) and generate a new population of offspring. Genomes with better objective values are preferred among generations. Based on this principle of “survival of the fittest,” genomes with more desirable attributes, and thereby representing stronger solutions, survive through generations and iterate toward optimal values.

The structure of a genetic algorithm lends itself well to our defined problem of maximum lifetime multicast for directional antenna equipped networks. Our problem could be described as the combination of topology requirements (logical tree) subject to resource allocation (power variables). While it is well known that GA’s do not perform well on combinatorial optimization problems, our problem can be split into two components. By encoding logical topology into the genome, feasibility and lifetime determination can be determined in the objective function. This natural separation of the problem into coupled components allows us to easily describe the problem as a GA. We can then leverage the wealth of genetic algorithm software and techniques available when finding a solution.

Structure

Encoding of a problem into a genetic algorithm is problem specific. Often, the method by which the problem is encoded and evaluated determines the algorithm’s performance. First, we describe the information encoded into each gene and genome and later describe the evaluation and objective function. For all GA instantiations described in this document, we use the GALib Library [70]. GALib is an object oriented open-source set of utilities for developing genetic algorithms written in C++. As such, it is fast, lightweight, and extremely configurable.

Though we are actually optimizing for lifetime, and therefore power variables are of direct interest, we choose to encode the logical topology into the GA code. The objective function later tests feasibility and sets power variables accordingly. As shown in Figure 3.11, the gene at position i represents the *parent* for node i . Therefore, each gene has $|\mathcal{N} - 1|$ values, including a value representing “null,” or no link. This encoding comes directly from the observation that any node in the network can be listening to *at most* one upstream node. Although this encoding does not capture the feasibility requirements of the logical topology, it does help to limit the search space of the problem by putting the additional complexity of beam and power choice in the objective function.

Objective and Feasibility

During evaluation, each genome is first translated in to an $|\mathcal{N}| \times |\mathcal{N}|$ boolean link matrix. For instance, suppose the gene at position 3 is set to value 0. This would indicate that node 3 is receiving from node 0, and entry $(0, 3)$ is set to *TRUE* in the link matrix.

Once all links are encoded in the link matrix, a depth-first search [71] is performed from each receiver towards the source. If a path from any receiver to the source does not exist, then this topology is infeasible. Therefore, the objective function immediately terminates and the genome is assigned a score of zero.

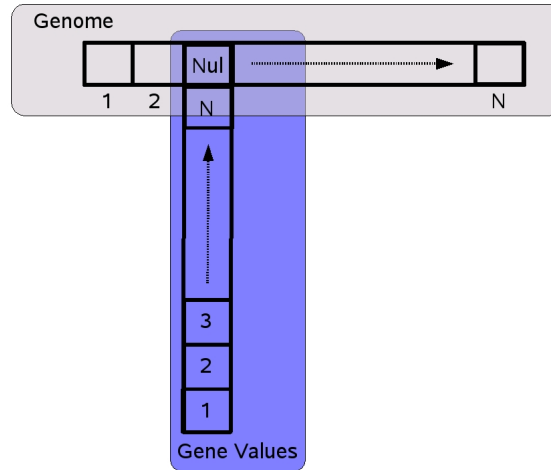


Figure 3.11: Genome Structure

If, on the other hand, all receivers have a feasible logical route to the source, the objective function continues its evaluation of the genome. At this point, the logical topology is handed off to the previously mentioned power algorithm. This set of difference equations either converges to a near-optimal set of power values, or continually increases the power settings beyond the maximum allowable power cap. Note that during the run of the algorithm the power variables may be slightly above the maximum value before reducing and converging to an allowable setting. Therefore, we decide to terminate the power-setting algorithm only when power settings are greater than twice their allowed cap. Non-allowable power values are also returned as genome values of zero. Otherwise, the lifetime is computed as the minimum lifetime of in-path nodes, and returned as the genome objective value.

Obviously, this process of testing feasibility, power convergence, and lifetime computation is computationally non-trivial. Since a GA will tend to reproduce individuals with high scores, the algorithm also caches feasible values. Once passing through the objective function, trees are assigned a hash value and cached. At the beginning of evaluation, the genome is first checked against the cache. Doing so prevents numerous calls to the GA objective function for identical inputs. The process of genome evaluation is documented in Algorithm 2.

Algorithm 2 Genetic Algorithm Objective Function – objective(*genome*)

```

1: if checkCache(genome) == true then
2:   return cacheValue(genome)
3: end if
4:  $M \leftarrow$  convertGenomeToLinkMatrix()
5: if ( $\forall i \in \mathcal{R}$ ) : pathToSource( $M, i$ ) then
6:   if powerAlgorithm( $M$ ) == feasible then
7:     return findLifetime( $M$ )
8:   else
9:     return 0.0
10:  end if
11: else
12:   return 0.0
13: end if

```

3.7 Summary

In this chapter, we introduce our mathematical program for both directional transmission and reception, our semi-distributed directional reception heuristic, and our directional reception genetic algorithm. All proposed methods are designed to find the maximum path lifetime of a directional antenna equipped ad-hoc network. The mixed-integer program (MILP) finds the optimal lifetime value, and the two heuristic methods find approximate values.

In the next chapter, we use the MILP to directly compare directional transmission and directional reception. We further use the MILP to analyze the performance of simple link-based heuristic methods when compared to the optima.

Chapter 4

Results: Directional Transmission vs. Directional Reception

In this section, we present the results of our experiments employing the two mixed-integer linear programs developed for directional communication in ad-hoc networks. The main goal of this chapter is to directly compare two different antenna use paradigms in a multi-hop wireless network: directional transmission, and directional reception. One might take note of the modeling similarities, and expect the methods to return qualitatively similar results. It will be shown in this chapter that the receiver-centric nature of directional reception makes it much more flexible and powerful for extending lifetime in wireless ad-hoc networks. First, we introduce the setup of the experiment. Then, we document the differences in optimal values between directional transmission and directional reception over the primary metric of multicast path lifetime, and over a secondary metric of cumulative power. After presentation of these data, we briefly discuss the intuition behind the results and conclusions. Results in this section are also documented in [72].

In the second section of the chapter, we perform a similar type of analysis. In this case, however, our directional receive heuristic (DRIP), and a previously defined transmission heuristic (D-MIP) [16] are the subjects of study. The heuristics are compared to their

respective MILP optima and head-to-head in an effort to characterize their performance and discover trends. Finally, we summarize the heuristic analysis and present our conclusions.

4.1 Experimental Setup

To perform a fair comparison, both the D-RX and the D-TX models were applied to identical networks. An experimental set of network configurations was generated, and for each network entry, the D-RX and D-TX models are individually applied. In this way, the models are directly compared with identical inputs.

Due to the complexity of the D-TX mixed integer program, we restrict network size to 10 nodes ($|\mathcal{N}| = 10$). Network sizes larger than this require extremely long solve times under D-TX and are impractical for study. In this section, we assume that inter-node and intra-session interference plays a part in each network model and simulation run. As part of this assumption, and as documented in Section 3.4.3, we assume that the directional antennas in use to be the more realistic model incorporating sidelobes. Referring to Figure 3.2, we assume that 70% of transmitted/received energy is in the main beam lobe with 30% leaking through sidelobes ($Pct^{inbeam} = 0.7$). Literature indicates that depending upon antenna design and implementation, Pct^{inbeam} can vary from 50% to approximately 90%, and we choose 70% as a conservative estimate.

Minimum beamwidth is fixed at 45° for both models, and D-TX is allowed to span up to 360° ($\theta_{min} = 45^\circ$). (Remember that in order for D-TX to reach neighboring nodes, it often must widen the main beam lobe.) Section 3.5.1 shows that the number of possible beams *increases* with smaller θ_{min} for the D-TX model and *decreases* with smaller θ_{min} for the D-RX model. The minimum beamwidth chosen is intended to be a suitable compromise between these competing methods, and also to reflect realistic hardware capabilities. When considering software steerable antennas (and especially phased-array configurations) extremely narrow beamwidths come at the cost of additional elements, power requirements, and complexity.

Therefore, we choose a suitably narrow beamwidth to illustrate the effects of gain and spatial multiplexing, but not at a level that would be unrealizable in implementation or make for intractable simulation times.

S^i (SINR ratio required) and $N_{thermal}$ (receiver thermal noise) are both set arbitrarily to one. Transmission power at any transmitter is capped at 100 units ($P_{max} = 100$), and the path loss exponent is set to two ($\alpha = 2$).

Nodes are randomly placed in 5x5, 10x10 and 15x15 unit areas to experiment with various node densities and transmission distances. At time zero, all nodes are assumed to be equipped with 300 units of energy ($R_i = 300$). Receiver set size was varied among five values $|\mathcal{R}| \in \{1, 3, 5, 7, 9\}$. For each combination of network dimension and receiver set size, 20 individual runs were performed, for a total of 300 runs. Network lifetime is calculated as the time until the first node in a multicast tree depletes its energy stores. We also track a secondary metric, cumulative multicast tree power. Cumulative power is defined as the sum of all node transmission powers for a given multicast tree.

Note that cumulative power is *not* directly optimized by the design of our MILP programs. It should be fairly clear that the max-min lifetime metric can lead to numerous trees having the same lifetime value. (The lifetime of the tree is influenced by the bottleneck link, leading to multiple possible combinations of non-bottleneck links.) When *ignoring* interference, there could be a large number of these metric-equivalent configurations. However, note the structure of the SINR requirement as shown in Equation 3.4. Globally, the MILP will also minimize cumulative power, as *any* extra power on an $i \rightarrow j$ link will increase the costs of all non- $i \rightarrow j$ links. Therefore, cumulative power is not directly minimized, but tends toward the minimum by the SINR constraints and structure of the program.

Mixed-integer linear programs are solved with either a freely available MILP package GLPK [73], or CPLEX [74], a professional solver suite from ILOG.

4.2 Metric Comparison

Figure 4.1 displays histograms representing the multicast path lifetime and cumulative network power for D-TX and D-RX. The histogram is constructed by applying D-TX and D-RX to identical networks, taking the ratio of their output results, and placing the ratio in the appropriate bucket. Therefore, the histogram represents a graphical view of the head-to-head per-run performance.

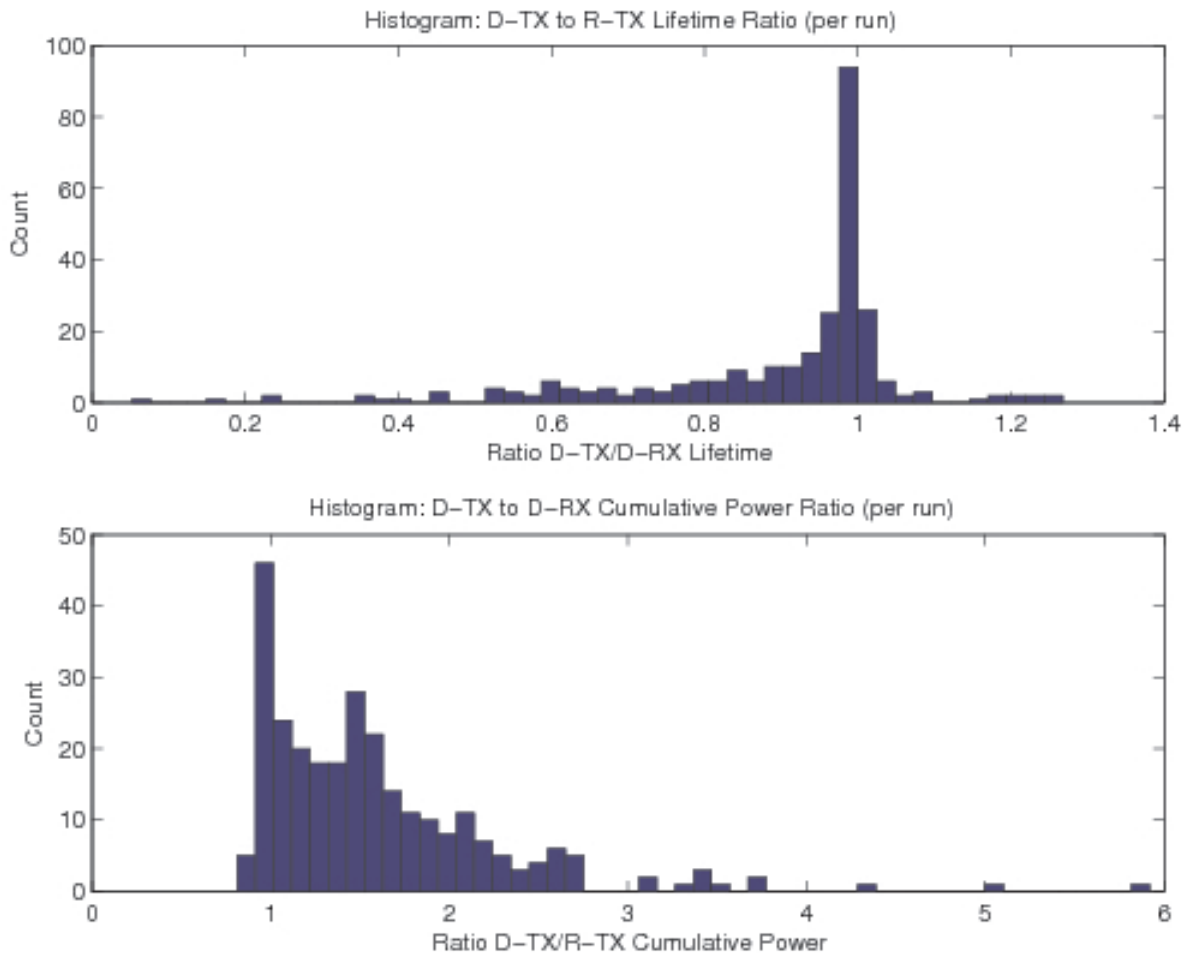


Figure 4.1: D-TX / D-RX Performance Ratio Histograms

The peak in Figure 4.1 (a) clearly indicates that for the vast majority of network layouts, D-TX and D-RX return approximately the same network lifetime. However, it is interesting

to note the bias to the left of the peak. There are a significant number of runs where D-TX’s lifetime is only 60% to 80% as long as D-RX. Moreover, the outliers indicate that there are runs where D-TX can only achieve 20-40% of the lifetime of D-RX. Compare this with the cases where D-RX is inferior. In only one case does D-TX beat D-RX by 40%, and only a handful where the difference is as high as 20%.

Figure 4.1 (b) shows the primary area of D-RX’s superiority, the cumulative power metric. D-TX almost never consumes less power overall than D-RX for any optimal tree configuration. In fact, cumulative power usage under D-TX is approximately 100 – 250% the power under D-RX for the same network layout. Only in *a handful* of cases does D-TX utilize less power than D-RX, and in those cases, the difference is negligible. It seems that while D-TX may achieve comparable lifetime values to that of D-RX, it does so at the expense of increased overall power utilization.

Figure 4.2 expands the results above and shows trends over all tested network parameters. The data point at the intersection of two indicated model parameters is the average of all runs for that parameter combination. For instance, at $|\mathcal{R}| = 7, Y = 5$, Z represents the average D-TX to D-RX lifetime ratio of all networks with seven receivers and a network area of 5x5 units.

When $|\mathcal{R}| = 1$, the two schemes have comparable lifetimes. This is somewhat expected, since with one receiver, the multicast case degrades into a unicast. In a unicast scenario, there is no need to replicate packets at any node, and therefore, no benefit of having multiple receivers connected to any transmitter. Nodes in both D-TX and D-RX will use the narrowest beamwidth available, and the multicast tree simplifies into a path. Therefore, any discrepancy in performance between the two programs is due only to network layout. Depending on source and receiver location for this small set of scenarios, directional transmission may outperform directional reception for a given configuration.

Clearly, though, as $|\mathcal{R}|$ increases, D-RX outperforms D-TX, at times by almost 20%. We would expect that the ratio would follow a clear and well-defined trend, but Figure 4.2 shows

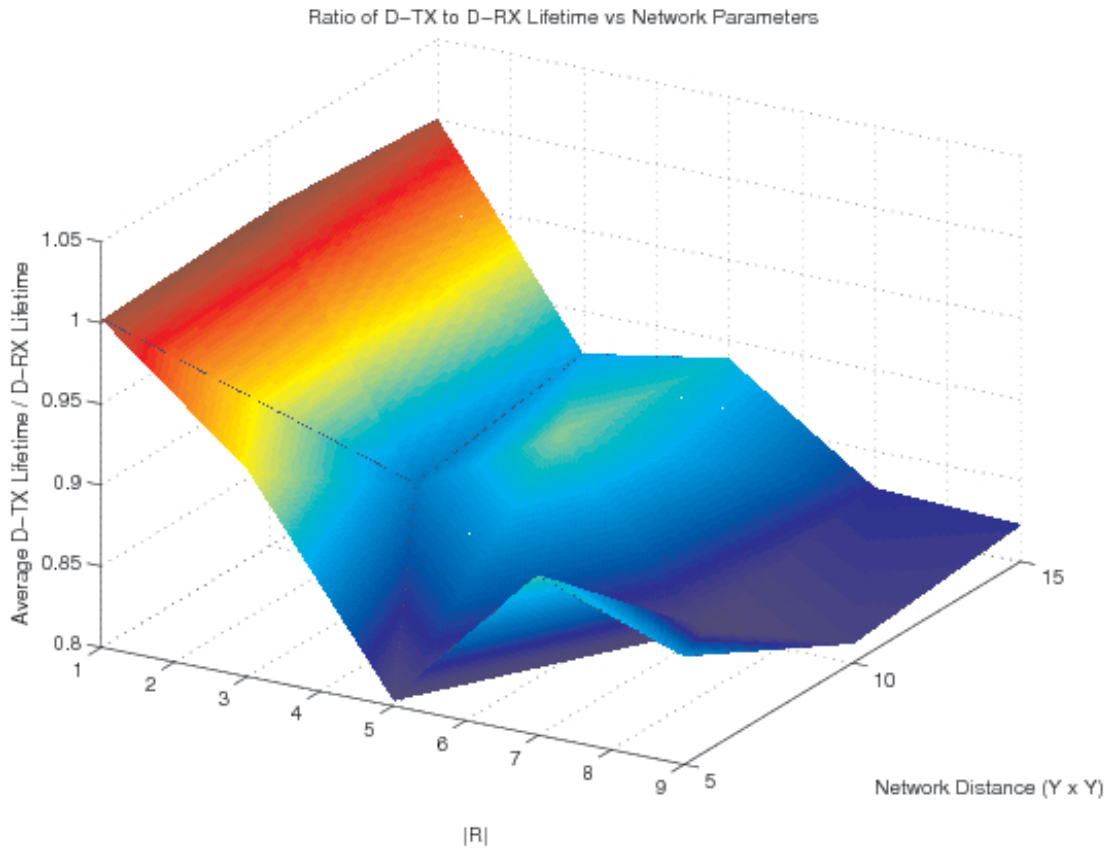


Figure 4.2: D-TX / D-RX Lifetime Ratio vs. Network Parameters

some non-smooth characteristics. Notice, however, that the plot actually displays a small number of points. Though the graph plots a three dimensional surface, there are actually only 15 data points representing simulated parameter combinations. Smoothness aside, we see that as the number of receivers increases, the average lifetime of D-RX can be as much as 20% higher than D-TX for the same network. Also note the counterintuitive result that network size does not affect the comparison. We would expect that large networks and low node density would allow D-TX to be more competitive with D-RX. However, Figure 4.2 shows that the number of receivers has a much more dramatic effect on the lifetime ratio.

The secondary metric of cumulative power presents a more powerful case that directional reception is the superior alternative to directional transmission.

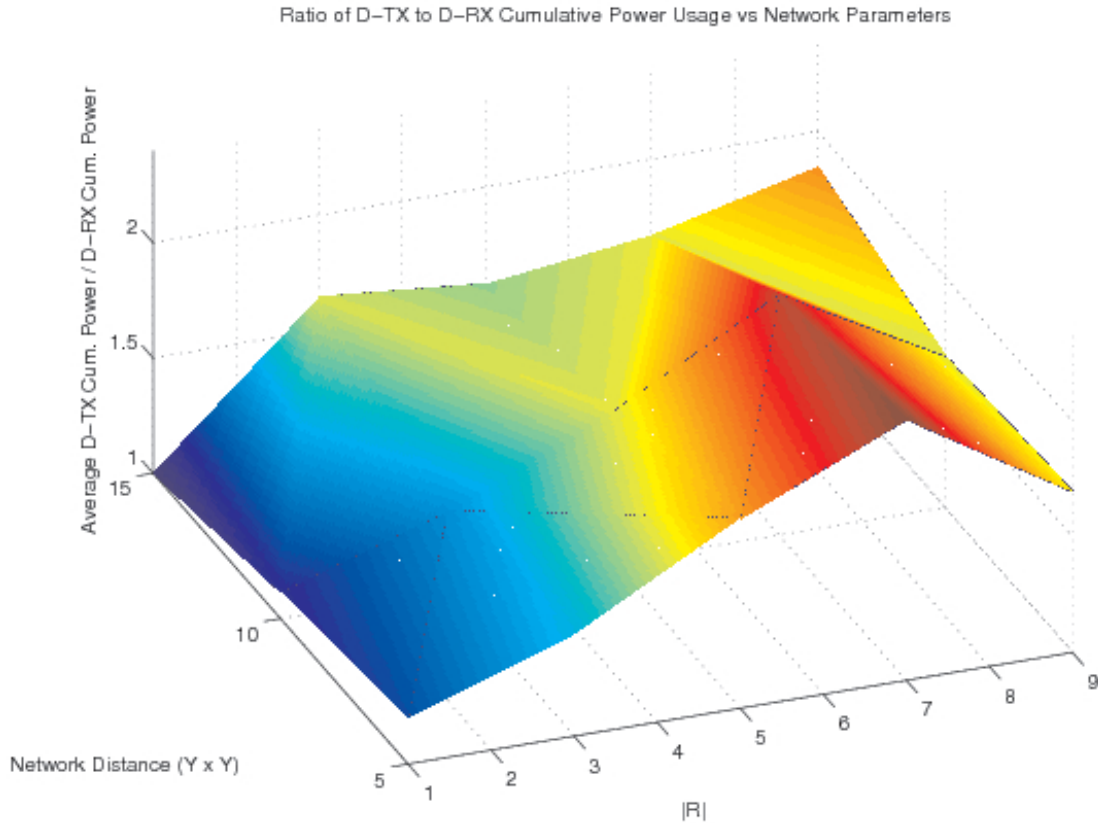


Figure 4.3: D-TX / D-RX Cum. Power Ratio vs. Network Parameters

Figure 4.3 displays the second metric of cumulative consumed power ratio for D-TX/D-RX (note that the graph's viewpoint has been changed for clarity). Again, for one receiver, lifetime values are similar. Note, that D-TX shows best performance when compared with D-RX over longer network distances and small receiver set sizes. Intuitively, we expect this to be the case, as node density decreases and D-TX is transmitting in a single angular direction. As $|\mathcal{R}|$ increases, and especially over smaller network distances, D-RX clearly outperforms D-TX. For $|\mathcal{R}| = 7$, in a 5x5 area, D-TX consumes approximately *twice* the cumulative power of D-RX. Referring back to Figure 4.2, we see that for this parameter combination, the lifetime of D-TX jumps slightly (compared to neighboring points) to 90% of D-RX. Quite clearly, this jump in lifetime comes at the cost of significant additional cumulative transmission power.

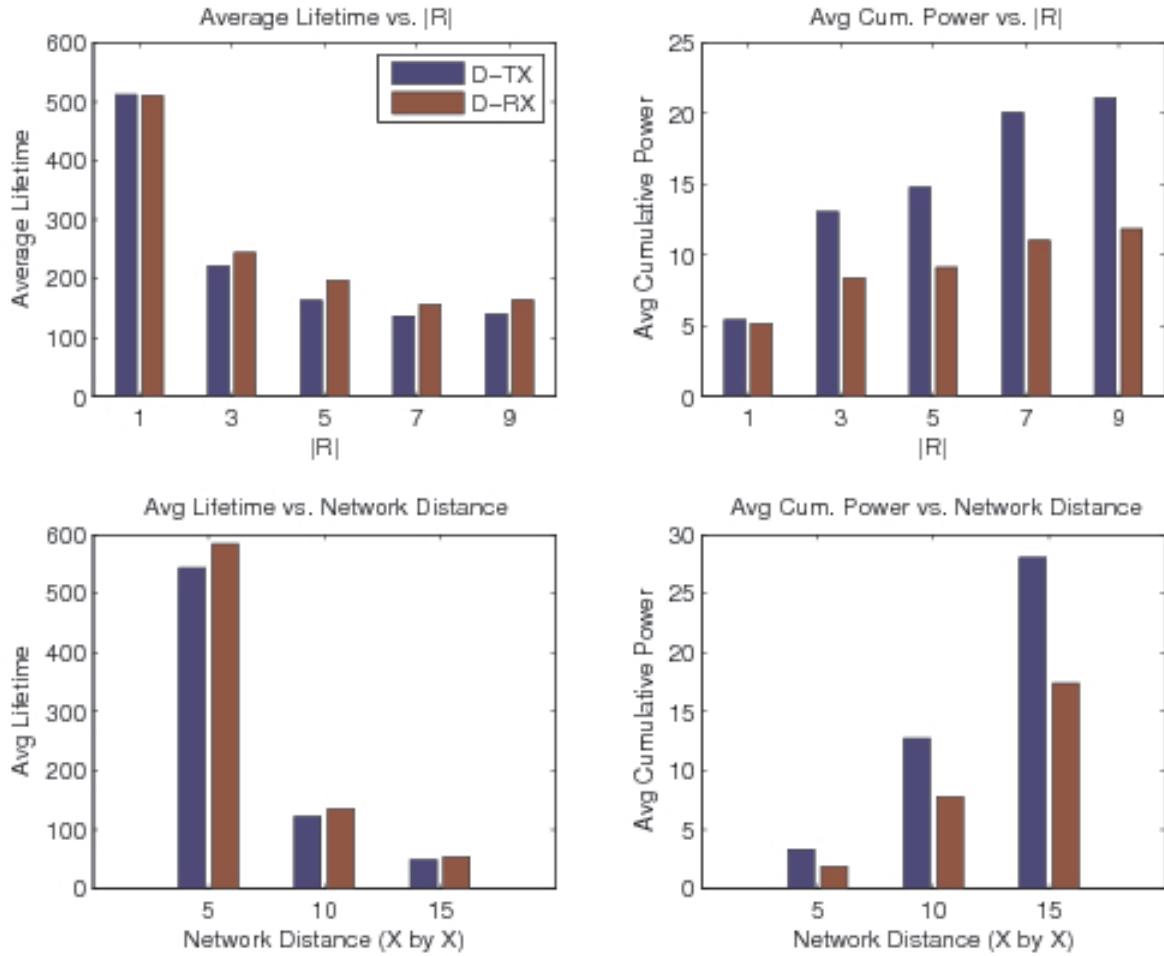


Figure 4.4: D-TX / D-RX Ratios vs. Single Network Parameter

The reader may question why the peak of the graph lies at this point. Recall that network layout is identical among *individual* runs, not among different parameter combinations. Therefore, the random layouts in the case where $|\mathcal{R}| = 7$ and in a 5x5 area show better performance under D-RX.

From this plot, it is obvious that D-RX has a distinct advantage over D-TX in power consumption. The plot clearly shows that the ratio is seldom, if ever, less than 1.

Figure 4.4 dispenses with the ratio and displays *actual* lifetime and cumulative power data over a single network parameter. The weak correlation between network lifetime and cumu-

relative power is evident. Regardless of $|\mathcal{R}|$, and network size, both schemes return comparable network lifetime values. On the other hand, the cumulative power use in D-TX is often much higher than that of D-RX over the same parameters. Surprisingly, this is the case even in sparse networks, where network size is 15x15. In these cases, we might expect that the longer transmission distances would be conducive to the narrow antenna patterns preferably employed by D-TX. However, these data show that this is not the case. The only exception are networks having small receiver set sizes, where D-TX can narrowly focus *all* antenna beams participating in the network. In the next section of the chapter, we attempt to provide the reader with some insight as to why this is the case.

4.3 Intuition

As shown in the previous figures, the D-TX and D-RX only differ slightly in performance of max-min lifetime statistics. But, over the same set of simulation results, cumulative power can be drastically different between schemes.

Most likely, this behavior is the result of the additional freedom afforded by D-RX. D-TX is best suited to point at receivers that are close in angular direction (σ), leading to a small beamwidth (θ) and higher gain. By the *directional reception observation*, D-RX will always use narrow beams with maximum gain. D-RX allows multiple receivers to point their main lobes at a source nearest in Cartesian distance, regardless of angular direction. This also assists in interference mitigation. Nodes can point at any potential transmitter, limiting the magnitude of interference from other active sources.

Figure 4.5 demonstrates a network layout and possible solution that could reflect the disparity between our primary and secondary metric. Suppose nodes are located as in Figure 4.5 (a), where node j is transmitting to the receivers k, l, m . In this layout, all nodes are separated by an equal distance d . Also, for our the purposes of this example, assume that all nodes have equal initial energy reserves.

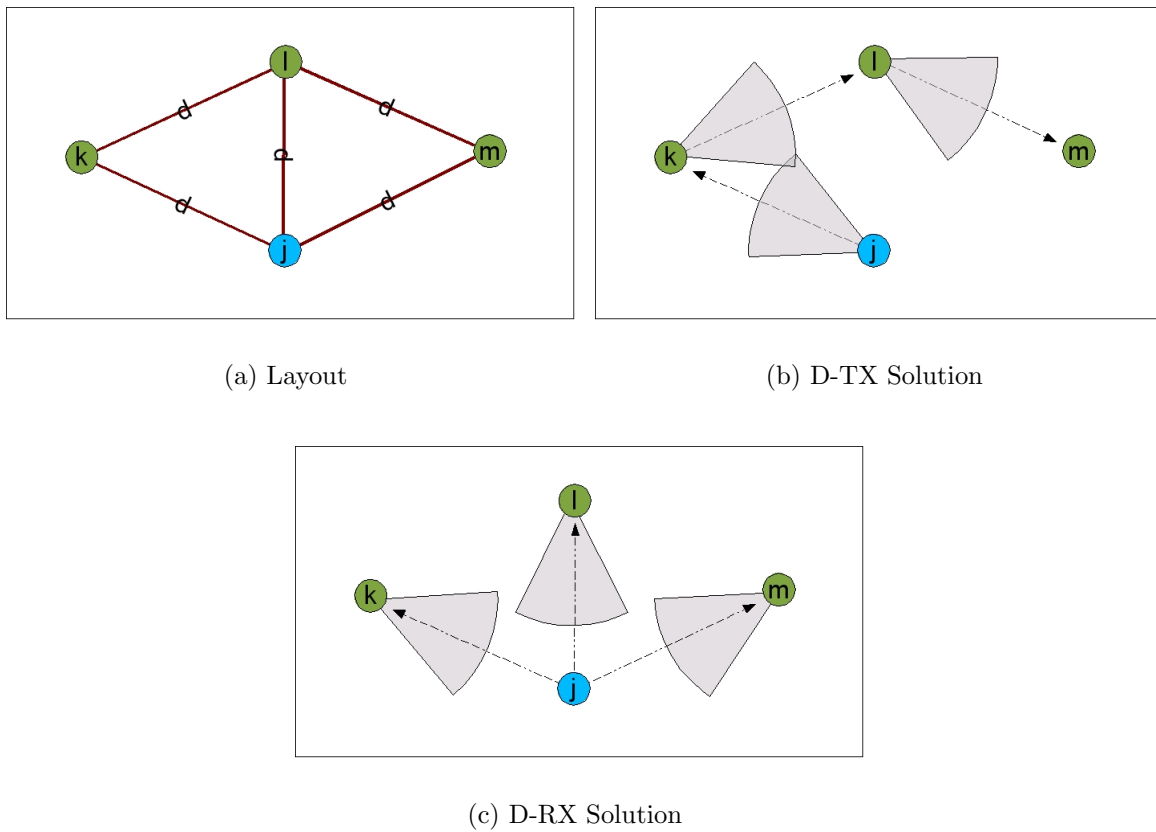


Figure 4.5: Intuition Example

Figure 4.5 (b) illustrates a possible D-TX solution. Nodes j , k , and l are all transmitting. Assuming negligible side-lobe interference, all three transmitting sources could be transmitting at identical power levels. Therefore, since we have identical initial energy stores, each node's *lifetime* is also equal.

Because the goal is to maximize *lifetime*, node j will not widen its beam significantly to cover a larger receiver set. Doing so would significantly reduce its own (and therefore the tree's) lifetime. Instead, numerous nodes use the highest gain possible by using minimum beamwidth, and covering only a single neighbor. (Obviously, j could reverse the initial transmission direction and use m as the first hop without changing the result.)

Contrast Figure 4.5 (b) with Figure 4.5 (c), which represents a possible D-RX solution. Here, the receiver nodes are free to utilize the same narrow beamwidth, but *all* receivers are free

to point it at a single transmitter. Node j 's power setting is sufficient to reach the entire receiver set using no additional hops.

If we assume minimal side-lobe power loss, then the D-TX solution and the D-RX solution should have comparable lifetimes. Since the node antenna gain is identical in either case, the power requirement *at each node* can be considered to be roughly equal. Therefore, the path lifetime in both cases is similar. However, as Figure 4.5 (b) indicates, the cumulative power for the D-TX solution is much higher.

In this document, we primarily focus on the performance regarding max-min path lifetime per *individual* multicast. The implications of higher cumulative power usage by D-TX are most likely to show their effects over multiple multicasts delivered by the network, and we consider this an area of future investigation.

4.4 Heuristic Analysis

This section builds upon the analysis of the D-TX and D-RX MILP models. Here, we characterize the performance of computationally inexpensive heuristic methods designed to approximate the optima found by our mixed-integer programs. D-MIP [15] is compared to the MILP optima for directional transmission, and our own directional reception heuristic DRIP is compared to the directional reception optima.

4.4.1 Mapping Heuristic Topology into Interference Model

For the directional transmit case, we implement D-MIP as specified in [16]. For directional reception, we model DRIP as specified in Section 3.6.1. Note that neither of these heuristics are specifically designed to operate in environments that consider interference, or with our more realistic antenna model. Both of these methods model the network as a directed graph with fixed link costs. In other words, the heuristics have no knowledge of beam enumeration,

side-lobes, or SINR constraints. Instead, they simply choose a series of $i \rightarrow j$ links that form a multicast delivery tree. Recall from our beam enumeration discussion that a single $i \rightarrow j$ link may have multiple SINR constraints in the MILP corresponding to multiple beams. For a fair comparison then, the logical topology of a DRIP or D-MIP solution must be mapped into our MILP model, which does account for beams, interference, and side-lobes. Below, we briefly explain the mechanism by which we map the heuristic solution into the interference model and find a solution over the remaining subspace.

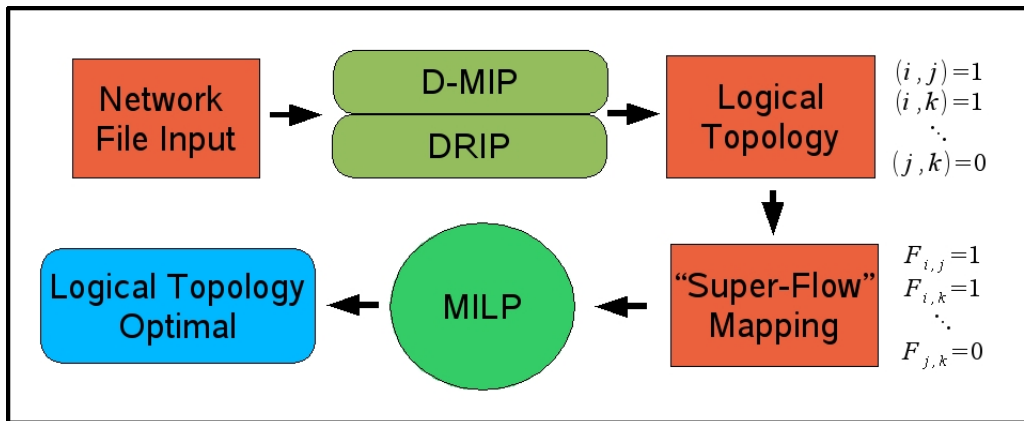


Figure 4.6: Heuristic Mapping Procedure

To obtain the resulting lifetime under the assumption that interference plays a role, the logical solution from each heuristic solution is mapped into a subspace of the D-TX or D-RX model as shown in Figure 4.6. Referencing Figure 3.4 and Figure 3.5, the logical routing tree specified by D-MIP or DRIP is input to the model via the $F_{i,j}$ “super-flow” variables. For instance, an active i to j link in D-MIP is input as $F_{i,j} = 1$ in the D-TX MILP model.

Recall that these continuous “super-flow” variables are only equal to one if at least one SINR constraint from i to j is sufficient. Once the logical topology is set, the MILP optimizes the remaining subspace of power variables and beam choice. Loop prevention and single input link constraints prevent the MILP from choosing alternative routing paths. As such, the optimal beam configuration (and therefore optimal max-min lifetime) for the given logical topology is returned.

The same set of networks is used as in input to the heuristic methods, and to the mixed-integer programs. In this way, the heuristics can be compared directly to one another, and to their respective optima. Note that as part of this comparison, the metrics for D-MIP and DRIP have been adjusted to reflect the percentage of power in the main lobe.

4.4.2 Experimental Results

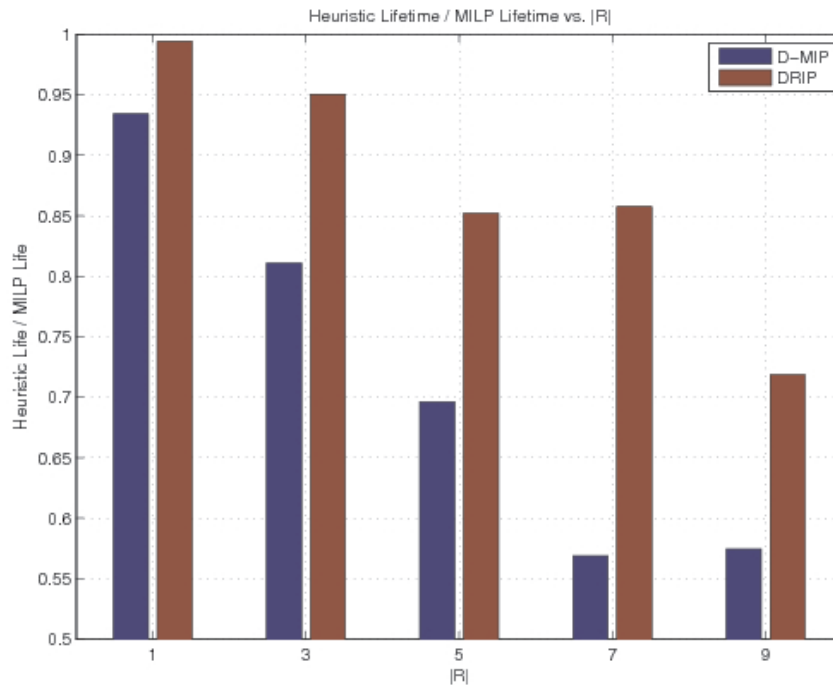


Figure 4.7: Heuristic Lifetime Ratio vs. $|\mathcal{R}|$

Figure 4.7 illustrates the ratio of heuristic lifetime to MILP lifetime versus receiver set size. Here, we see two interesting trends. First, and most obvious, is that as the receiver set size increases, both heuristic methods have difficulty approximating the MILP optima. The intuition behind this observation should be relatively clear. That is, as the number of receivers increases, the delivery tree becomes more complex and constrained. Whereas small receiver set sizes can be reached with a few branches, larger set sizes require additional link choices. Therefore, it becomes increasingly unlikely that the heuristic method will return

the optimal tree found by the MILP.

Our second observation is that DRIP performs better than D-MIP in all situations, and especially as receiver set size increases. For the unicast case, both heuristics are fairly comparable. This is in line with our intuition regarding directional transmission and directional reception. When there is a small set of receivers, directional transmissions can use narrow beams with high gains. However, as $|\mathcal{R}|$ increases, D-MIP's performance degrades. DRIP also exhibits a degradation of performance, but does so much more gracefully than D-MIP. Even when $|\mathcal{R}| = 9$ and DRIP is performing at its worst, it still outperforms D-MIP by almost 15%, and is within 30% of the optimal. Though not specifically designed for the interfering environment, we find that DRIP can approximate the MILP optima in most cases.

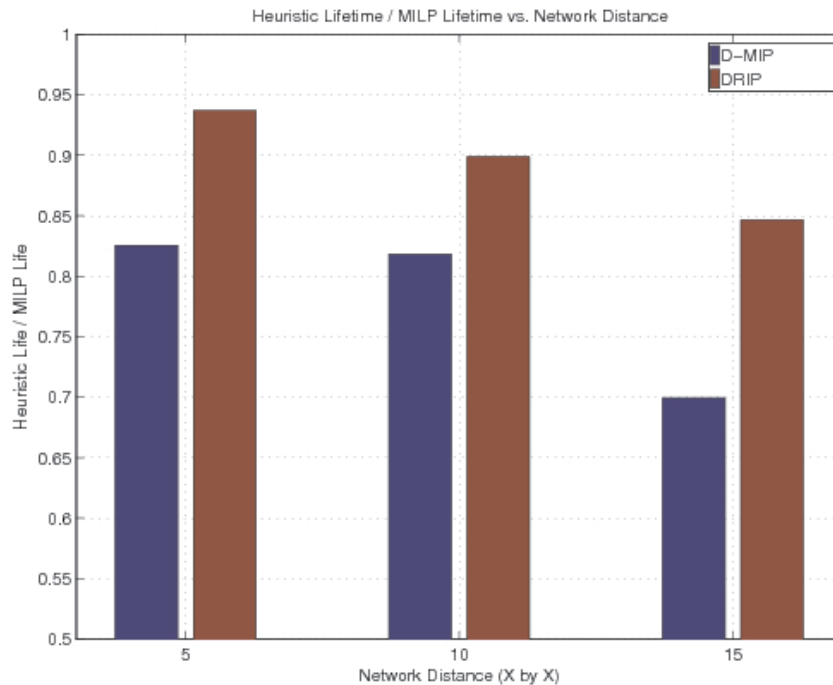


Figure 4.8: Heuristic Lifetime Ratio vs. Network Size

Figure 4.8 shows the lifetime ratio results as averaged by network distances. We see that, again, DRIP consistently outperforms D-MIP over all tested network parameters. While we

might expect that DRIP outperforms D-MIP on smaller networks, we see the counterintuitive result that DRIP outperforms D-MIP over larger, less dense networks. Also, we see that DRIP’s performance degrades only marginally over larger networks, while D-MIP suffers most when network distance is 15x15. From the previous section, we know that directional transmission is most comparable to directional reception when networks are large and node density is low. However, Figure 4.8 shows that the link-based heuristic chooses trees that do not exploit this situation, leading to D-MIP being 30% below the D-TX optima.

It is also noteworthy that Figure 4.7 and Figure 4.8 show the ratio of the heuristic to their respective MILP optima. The results may be slightly misleading. From the previous section, we know that the D-RX lifetime optima are often *larger* than the corresponding D-TX lifetime. Therefore, not only does DRIP better approximate the optimal values, but often, that optimal value is *larger*. For clarity, we dispense with the ratio and display actual lifetime values below.

Figure 4.9 shows the actual average lifetime values for the schemes in question. Readers should note that the bars for $|\mathcal{R}| = 1$ and network distance of 5x5 have been scaled down by constant values of 300 units, and 500 units respectively, for clarity of displayed data.

The difference between ratio and actual lifetime data are now evident. For instance, notice the difference between D-MIP and DRIP for a network size of 5x5 units (again, these results are scaled down by 500 units). Whereas Figure 4.8 shows that DRIP is approximately 10% closer to its optimal than D-MIP, the *actual* difference in lifetime is almost 150 units. The optima of D-RX are almost 50 units higher than that of D-TX for networks of size 5x5. When these two sets of results are combined, we see that D-RX optima outperform D-TX, and DRIP better approximates the optimal than D-MIP over all configurations. Therefore, DRIP is *clearly* the better alternative for extending network lifetime.

4.4.3 Summary

The investigation of simple link-based heuristics in this section furthers the conclusions obtained from our mixed-integer results. Directional transmission suffers in situations where there are numerous multicast receivers, or where node density is high. Most likely, this is due to the previously discovered characteristic of D-TX that it requires significantly more cumulative power to match the lifetime results of D-RX. When transmitting directionally with interference, nodes must choose to either widen their beam and lose gain, or to employ other neighboring nodes as intermediary hops. The first solution requires the transmitter to increase power, the second requires additional cumulative power throughout the network. In either case, D-RX can provide superior results because of its flexibility.

Based on our results, we conclude that directional reception is the superior alternative for extending multicast path lifetime. In the succeeding chapter, we study directional reception exclusively as it performs on heterogeneous networks. We also document the results of our genetic algorithm and semi-distributed heuristic as they compare to the network optima.

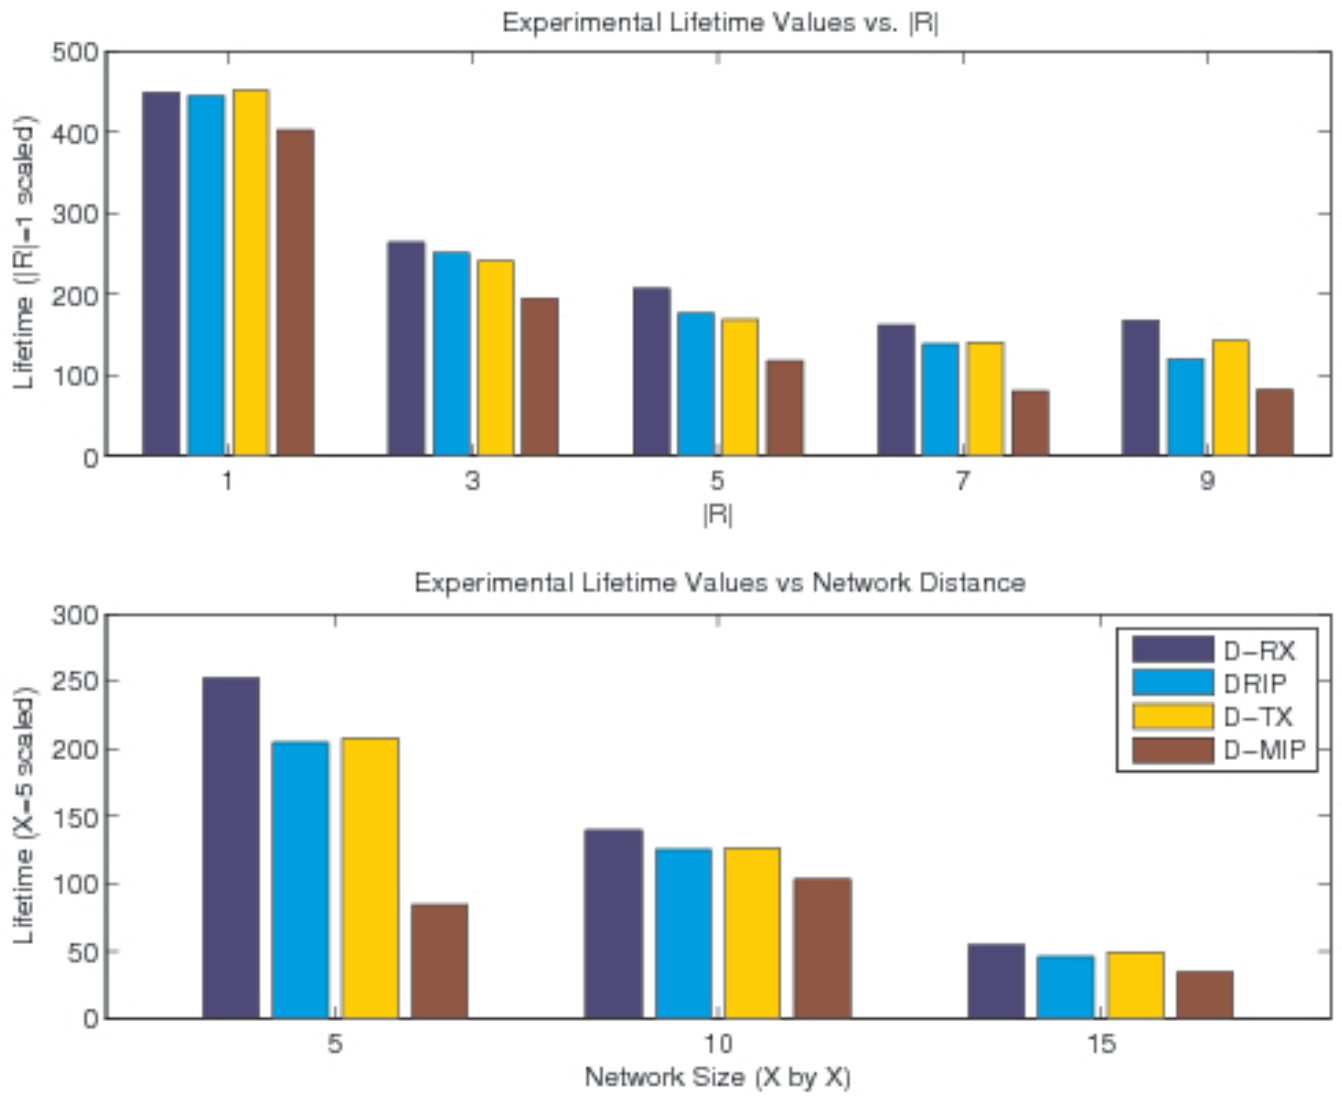


Figure 4.9: Actual Lifetime Values (Heuristics and MILP)

Chapter 5

Results: Directional Reception

In Chapter 4 we discussed and illustrated the differences between directional transmission and directional reception when applied to identical networks. From these data, it is quite apparent that directional reception is the superior alternative for extending lifetime and conserving cumulative power. In this chapter, we study directional reception exclusively. We seek to characterize the performance of directional reception under a variety of different inputs, including: antenna models, network distances, receiver set size, directional antenna equipped nodes, and minimum beamwidth. Investigating the model's performance under heterogeneous conditions and performing analysis on the results allows us to develop intuition as to what characteristics are most influential to the system. The first section of this chapter is devoted to the heterogeneity study. We present a series of figures and tables documenting the behavior of the optima under the influence of different input parameters.

Later in the chapter, we display the results of our semi-distributed heuristic as compared to the optimal values. The semi-distributed heuristic was designed using the results obtained from the comparison study, and this section allows us to critique the performance of the semi-distributed heuristic as a method of approximating the network optimal. Finally, we present our genetic algorithm results. The GA was designed to serve as a low-computational cost alternative to the MILP optimization, allowing us to study larger networks and more

models. Our results indicate that the semi-distributed heuristic and genetic algorithm have comparable performance and similar trends.

5.1 Heterogeneity Study

In this first section of the chapter, we use the mixed-integer program to find network optima of heterogeneous network layouts. Numerous network setups were simulated, each having a specific set of interesting characteristics, and optimized using our MILP. By varying the parameters of the the network the MILP operates on, we have the ability to treat the mixed-integer program as an input-output system.

Throughout this chapter, we present the results output by the mixed-integer program from an input network with different attributes. As we vary the inputs and study the results, we see that our modeling assumptions are justified, and that our more realistic network model has important implications. Moreover, we will show that directional reception operates well in varied, heterogeneous environments, and is a powerful alternative to directional transmission.

5.1.1 Number of Directional Antenna Equipped Nodes

For all sections of this chapter, all simulated network files have with 15 total nodes including source and all receivers. However, as part of our study, we vary the number of directional antennas in the network among four values, $|\mathcal{D}| = [0, 5, 10, 15]$. Further attributes are also modified throughout our study, and will be discussed with the appropriate set of results. Two levels of main-lobe efficiency are simulated, the first where 50% of the power is assumed to be in the main lobe (50% side-lobe), and the second where 90% of the power is in the main lobe (10% side-lobe).

Figure 5.1 illustrates the results for differing numbers of directional antennas, and also for the two different levels of side-lobe intensity. The average value for lifetime and cumulative

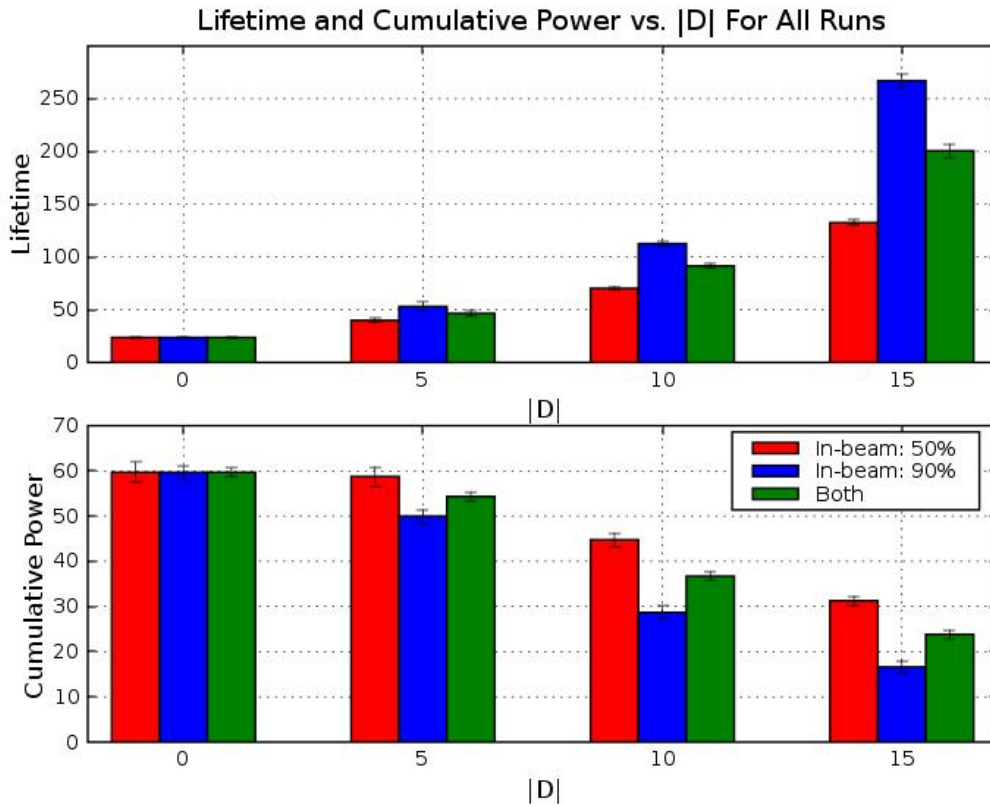


Figure 5.1: Lifetime and Cumulative Power vs. $|\mathcal{D}|$

power and confidence intervals are plotted against $|\mathcal{D}|$. Notice that for $|\mathcal{D}| = 0$, sidelobe intensity has no effect on the lifetime or cumulative power results. This is to be expected, since clearly, omni-directional transmission has no side-lobes. As the number of directional antennas increases, we see a clear correlation with higher lifetime and lower cumulative power. When all nodes are equipped with directional antennas ($|\mathcal{D}| = 15$), lifetime is at its highest, but is dramatically different depending upon side-lobe intensity. The effects of side-lobes are quite evident, especially in the cumulative power metric. This behavior can be attributed to two phenomena: the loss of efficiency in directional reception (losing an additional 50% of power through the side-lobes), and the higher amount of interference in the network. When 90% of the transmitted radiation is in the main beam lobe, less transmission

power is required to reach neighbors *and* less power acts as interference to other ongoing sessions. However, the overall trends relating to the number of directional antennas are still apparent. The more directional antenna-equipped nodes in the network, the higher the resulting lifetime, and lower the expended cumulative power.

5.1.2 Network Size

In addition to varying the number of directional antennas in the network, the network *size* is also varied among three values ($X = [5 \times 5, 10 \times 10, 15 \times 15]$). Our chosen path loss exponent is $\alpha = 2$, and thereby, the power required for transmission is proportional to the square of the distance. Since the number of network nodes remains constant, as the network size increases, network density consequently decreases as the 15 nodes are spread over a larger area.

Figure 5.2 shows the lifetime and cumulative power metrics associated with the three network distance settings. To better compare with the previous section, results are further broken down according to the number of directional antennas in the network. Remember from our setup parameters, node power is capped at 100 units. Therefore, in large networks, multi-hop network topologies are necessary for feasibility. (Note that networks that returned infeasible results have been removed from this analysis for clarity.) The trends from the previous section are more evident here. For small networks, the difference in lifetime when nodes are directional-antenna equipped as compared to omni-directional nodes is dramatic. It is interesting that at large network distances, and thereby sparse node layouts, directional antennas do not compensate enough to substantially increase lifetime when compared to dense networks. One of the potential benefits of directional communication is the ability to connect distant nodes at a much smaller energy cost. However, Figure 5.2 would indicate that node density is a much more important consideration, *even when directional antennas are included*.

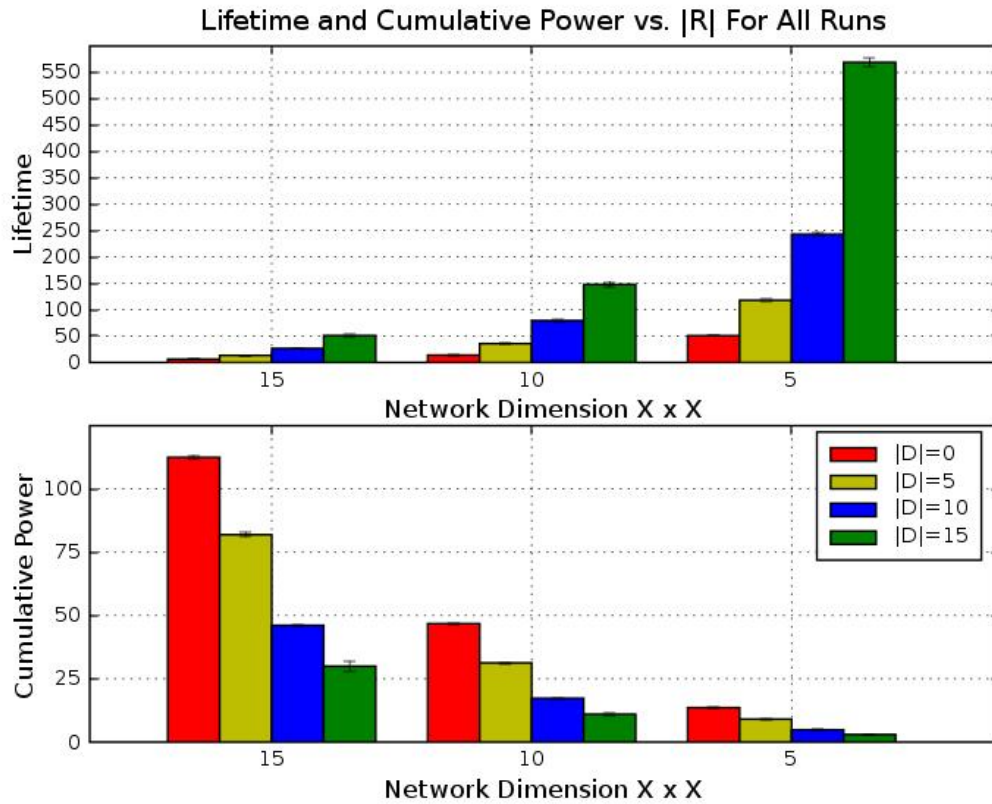


Figure 5.2: Lifetime and Cumulative Power vs. Network Distance

5.1.3 Number of Receivers

Our analysis would be incomplete without an investigation of the effects of receiver set size on performance. For these data, receiver set size was set at three individual values ($|\mathcal{R}| = [2, 7, 12]$). Receivers were randomly chosen among all non-source network nodes, independent of location and whether or not the node was equipped with a directional antenna. We would expect that because a larger receiver set requires more coverage in the network, it should require more power, and have smaller lifetimes.

Figure 5.3 shows this behavior exactly. Notice, however, the effect of having even *some* omni-directional nodes in the network. Especially in the case when $|\mathcal{R}| = 12$, lifetime is

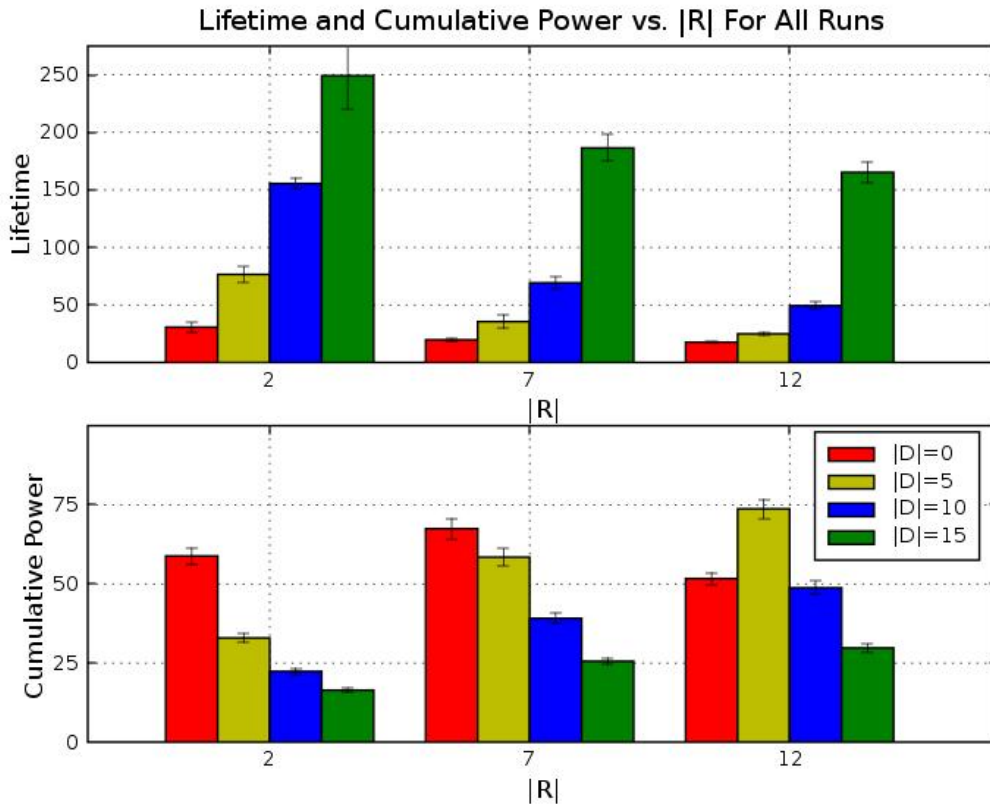


Figure 5.3: Lifetime and Cumulative Power vs. $|\mathcal{R}|$

much lower, and cumulative power is much higher for all cases where $|\mathcal{D}| < 15$. We can conclude that in this case, more receivers must be reached by the tree. Moreover, some participating nodes do not have directional antennas (including receivers), and therefore the power required is much higher.

Notice the results for $|\mathcal{R}| = 12, |\mathcal{D}| = 5$. Here, the MILP uses much more cumulative power compared to the scenario where all nodes are directional antenna-equipped, or even when all nodes are omni-directional. In this case, 25% of network nodes are directional antenna equipped and capable of spatial multiplexing. The MILP intends to maximize network lifetime, and likely uses these antennas to avoid interference. However, because most nodes *do not* have this capability, medium contention is still an issue and cumulative power is

slightly higher than the omni-directional case where *no* nodes have directional antennas.

5.1.4 Beamwidth

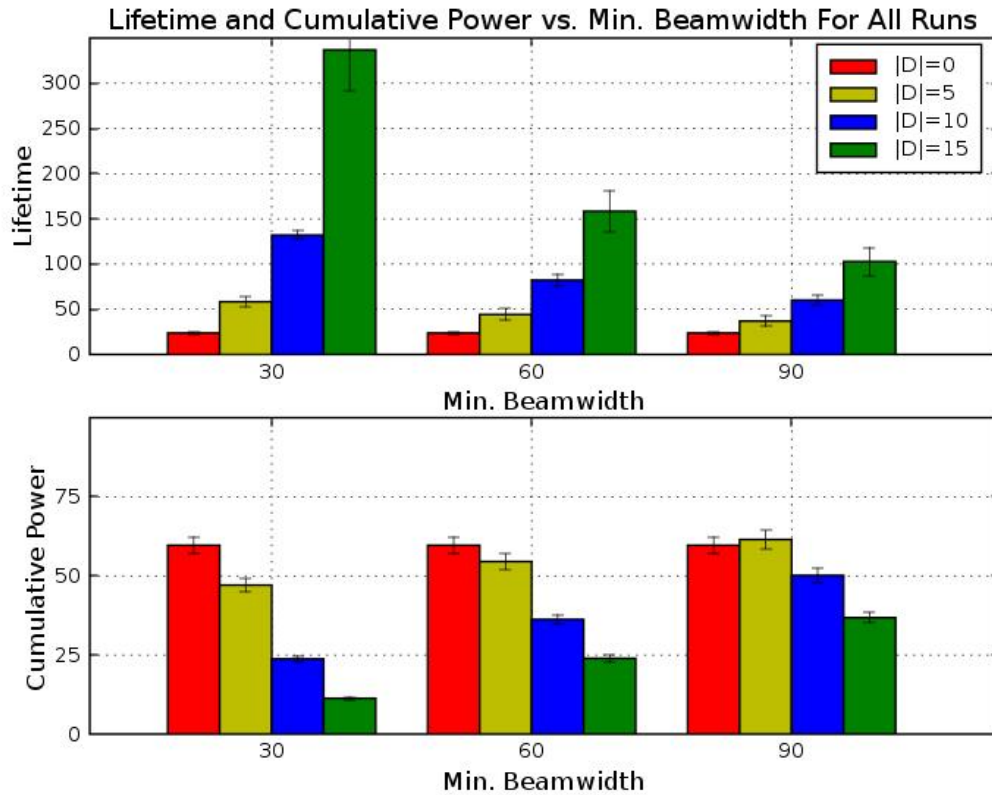


Figure 5.4: Lifetime and Cumulative Power vs. Min. Beamwidth

By the *directional reception observation*, when receiving directionally, nodes prefer beams with narrow beamwidth. Our antenna model assumes that narrow beams also have the highest gain.

Figure 5.4 shows the resulting lifetime and cumulative power over our three minimum beamwidth settings and number of directional antenna equipped nodes. The results illustrated support our assumptions regarding antenna beamwidth. Directional reception has

the greatest impact when numerous nodes are equipped with antennas having narrow beams and high gain. While lifetime exhibits expected trends, correlating well to $|\mathcal{D}|$ and minimum beamwidth, cumulative power is more interesting. As with our previous discussion, the inclusion of only a small number of antennas ($|\mathcal{D} = 5|$) means the program will maximize lifetime at the cost of using extra cumulative power. Moreover, for wide beamwidth with marginal gain, cumulative power is only slightly different among the chosen values for $|\mathcal{D}|$. This enforces our previous observation that cumulative power and lifetime are poorly correlated. In some cases, the MILP optimizes the primary metric at the expense of the secondary metric.

5.1.5 Sidelobe Intensity

Further evidence for this phenomenon is shown in Figure 5.5. Here, data are not grouped by $|\mathcal{D}|$, but rather by side-lobe intensity. As antenna efficiency decreases and more power is lost via the side-lobes, the system must work harder by expending additional cumulative power to maximize lifetime. However, because of interference, nodes cannot simply transmit at an extremely high power level. Instead, it becomes increasingly important to account for inter-session interference while maximizing path lifetime. Therefore, the mixed-integer program finds a balance between adding more transmission power to the system to maximize lifetime and accounting for detrimental interference.

5.2 Distributed Heuristic Results

In this section of the chapter, we analyze the performance of our proposed semi-distributed heuristic. To do so, the semi-distributed heuristic was applied to the same simulated data set as the previously documented D-RX results. As was performed in Chapter 4 for DRIP, we then compare the heuristic to the optimal on a run-by-run basis. Since the MILP represents the best possible bound on our metrics, we can compare best-case performance of

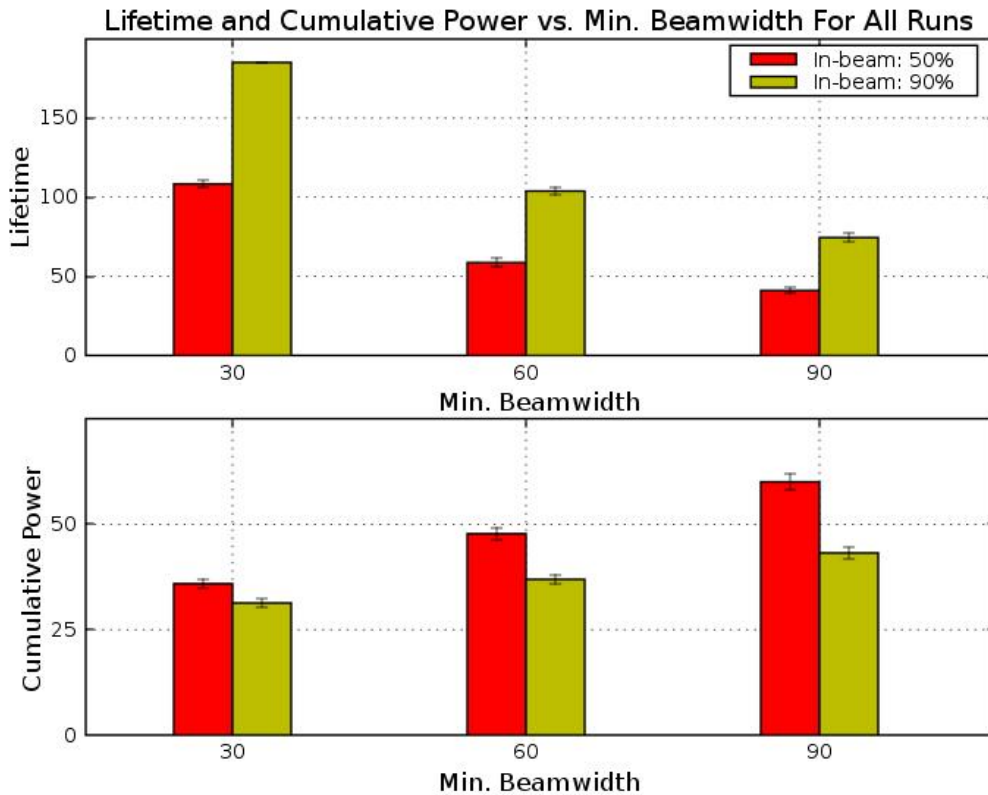


Figure 5.5: Lifetime and Cumulative Power vs. Min. Beamwidth

the heuristic method under many different input configurations to the upper bound of the system as returned by the MILP.

5.2.1 MILP Comparison

For this section, we use box plots to efficiently display to the reader some important statistics of our data. Median (as displayed by the red line), lower and upper quartile, and outliers are clearly depicted. Displaying average values alone can disguise the spread of the data over all runs. In this study, since we are comparing methods head-to-head, box plots provide an intuitive and informative representation of overall performance.

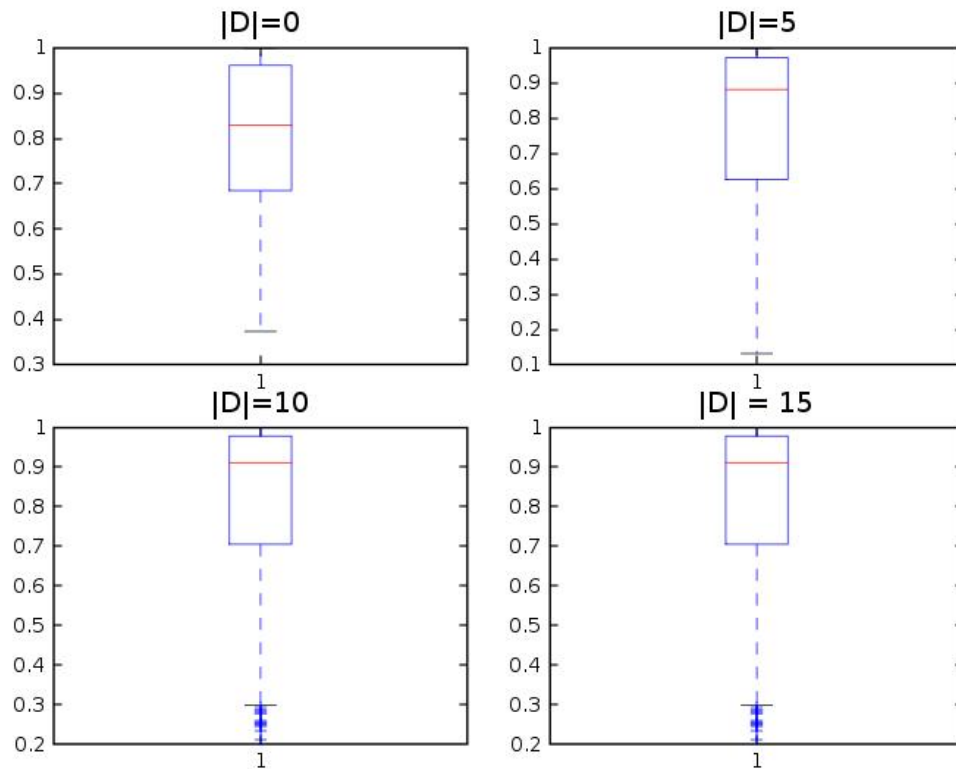


Figure 5.6: Distributed Heuristic to MILP Lifetime Ratios

Figure 5.6 shows box plots representing the ratio of the lifetime returned by our distributed heuristic versus the MILP optima. With the exception of the exclusive omni-directional antenna case, the heuristic produces, on average, over 80% of the mixed-integer optimal. Clearly, by the optimality criteria of the MILP solution, this ratio should never exceed one. Therefore, the outliers representing poor heuristic runs can drag the arithmetic mean to a much lower value than is actually representative of performance. With this in mind, one might consider the median to be a more descriptive statistic, and for most settings of $|\mathcal{D}|$, the median ratio is above 90%. While there are some extreme outliers in all cases, for the majority of runs, the ratio is greater than 70%.

Also, notice the plot for $|\mathcal{D}| = 5$. Here, we see that heuristic performance can suffer, as

the smallest non outlier ratio is approximately 15%. Recall from our previous section that this setting for the number of directional antennas had the most non-intuitive statistical trends. Cumulative power was higher than similar configurations, and only a slight increase in lifetime over the omni-directional case was observed. Figure 5.6 shows that the heuristic exhibits similar struggles with the simulated dataset. The heuristic was designed based on trends and intuition gained from our previous comparison study, and therefore is ill-equipped to handle these configurations.




Figure 5.7: Distributed Heuristic to MILP Cumulative Power Ratios

However, the cumulative power ratios as shown in Figure 5.7 better reflect the cost of the heuristic algorithm. Here, we display the ratio of cumulative power use in the distributed heuristic to the cumulative power of the MILP solution. The reader should notice the

immediate difference between this figure and Figure 5.6. Whereas the ratio of Figure 5.6 is capped at one (guaranteed by the optimality of the MILP solution), cumulative power may be less or greater than one. For most cases, the heuristic terminates with solutions requiring more power than the MILP topology. This might be expected, as the topology found by the heuristic does not represent the optimal layout for the network. On the other hand, the heuristic may stop at solutions that use *less* cumulative power. Recall that the distributed algorithm terminates when it can find no node to add that improves the overall optimal. Clearly, depending upon initial conditions, the heuristic can get stuck in local optima where cumulative power use may be less than the optimal values.




Figure 5.8: Distributed Heuristic Lifetime and Cumulative Power Bar Graphs

Figure 5.8 displays the average lifetime and cumulative power results for the MILP and

heuristic under the two sidelobe intensities studied. The performance of the distributed heuristic in all cases is encouraging. Regardless of intensity or number of directional antenna equipped nodes, the heuristic closely approximates the MILP lifetime. Recall from Section 4.4.2 that the centralized, interference-ignorant heuristic (DRIP) presented suffers from performance issues under some network configurations. Specifically, as the number of receivers increases, DRIP cannot account for the inter-node interference, and performance consequently suffers.

Also, the secondary metric of cumulative power is comparable to the MILP solution. In fact, when all nodes are directional antenna equipped ($|\mathcal{D}| = 15$), the heuristic uses no more cumulative power. Clearly, though, this comes at the cost of the lifetime metric, where the MILP is slightly superior.

5.2.2 DRIP Comparison

In Section 3.6.1, we introduced our centralized lifetime-centric heuristic method: the directional reception incremental protocol DRIP. Though we showed that it performs better than directional transmission based heuristics (even in the presence of interference), we did not show performance in heterogeneous environments, or compared to our distributed heuristic. Based on our previous section where we perform a direct comparison of our semi-distributed heuristic to the MILP, we use this section to perform a similar comparison of DRIP to the semi-distributed method.

For all data presented here, DRIP and our distributed heuristic were applied to identical network topologies. Recall that DRIP returns a logical routing tree, and does not specify beam choice or power settings. To find a comparable solution, DRIP must be mapped into the interference environment. DRIP's logical tree is input to the MILP as a logical routing tree. The MILP then returns the *best case* power settings and beam choice for the given tree topology. In short, we show the lifetime that the network could expect to achieve if DRIP were making routing decisions.

	$ \mathcal{D} $				
Pct^{inbeam}	0	5	10	15	\forall
50	100%	98.4%	78.3%	37.0%	76.5%
90	100%	94.8%	36.4%	3.4%	55.4%
\forall	100%	96.6%	56.7%	20.0%	54.7%
Beamwidth					
30	100%	93.3%	35.1%	5.1%	55.0%
60	100%	98.0%	61.0%	14.7%	65.6%
90	100%	98.5%	74.7%	41.0%	76.8%
\forall	100%	96.6%	56.7%	20.0%	54.7%

Table 5.1: Percentage of Runs where Distributed Heuristic is Feasible and DRIP is Infeasible

	Beamwidth			
Pct^{inbeam}	30	60	90	\forall
50 %	66.7%	46.7%	52.4%	61.5%
90 %	81.9%	74.7%	69.3%	77.9%
\forall	76.9%	68.0%	68.0%	73.4%

Table 5.2: Average DRIP / Dist. Heuristic Lifetime Ratios ($|\mathcal{D}| = 15$) **only**

Since DRIP has no knowledge of interference, the topology returned can be infeasible when subjected to an interference constrained medium. In fact, we find this to be an extremely common occurrence. Table 5.1 shows the percentage of topologies returned by DRIP that are *infeasible* when mapped into the interference model. We immediately see that omnidirectional communication is the worst case, as DRIP returns trees that are never feasible. This is not surprising, as DRIP was designed specifically to incorporate and utilize directional antennas. Moreover, no spatial multiplexing or perfect MAC means there is continuous contention for the channel, a situation DRIP is unable to predict.

On the other end of the spectrum, when all nodes are equipped with high gain directional antennas, DRIP finds *feasible* configurations for over 90% of the simulated runs. Under these assumptions, the network model more closely resembles the directed graph model that DRIP anticipates. Therefore, only this specific sub-case of heterogeneous parameters is chosen for study. The large number of infeasible runs for other configurations precludes fair comparison, and we show that our semi-distributed heuristic is superior for even this small set of results.

Table 5.2 itemizes the average performance ratio between our distributed heuristic and DRIP over feasible runs for different antenna characteristics. As in Table 5.1, we see that DRIP is at least somewhat competitive when antennas have narrow beamwidth and do not leak power excessively through side-lobes. Here, DRIP returns topologies capable of matching 81.9% of the distributed heuristic’s lifetime value. However, it is quite clear that DRIP’s ignorance of interference severely inhibits its ability to approximate the network optimal. Even when considering only this specific sub-case, on average, our distributed heuristic can generate approximately 30% higher path lifetimes than DRIP.

5.2.3 Summary

In this section, we discuss the performance of our proposed semi-distributed heuristic as compared to the MILP optima, and to our previously defined heuristic DRIP. Results show that while computationally much simpler than the mixed-integer program, the semi-distributed heuristic can approximate its results. For almost all simulated network layouts, the heuristic achieves over 70% of the MILP’s lifetime, and for nearly half of all simulations, lifetime is at least 90% of the optimal. This behavior was consistent independent of the antenna characteristics of the model. That is, even with broad beamwidth and leaky sidelobes, the heuristic achieves excellent results.

We contrasted this performance with that of our previously defined lifetime centric heuristic, DRIP. DRIP was intended for networks where inter-node communication is independent of other sessions (i.e., where interference is not a consideration). The data indicate that this places DRIP at a disadvantage in that it is unable to account for and avoid inter-session interference. For most cases, DRIP cannot even find a *feasible* routing topology under interference conditions.

Though Section 4.4 shows that DRIP at times can approximate network optima, we show here that this only occurs when all nodes are equipped with directional antennas, and when those antennas are suitably efficient. Our distributed heuristic, conversely, operates well

under heterogeneous conditions.

5.3 Genetic Algorithm Results

Finally, we will describe the results from our final heuristic method, our genetic algorithm. Again, the same network files input for both the MILP heterogeneity study and the semi-distributed heuristic study were used for comparison.

GA Parameter	Value
Population Size	30
Max Generations	800
Mutation Percent	5%
Crossover Percent	80%
Crossover Type	Single Point

Table 5.3: GA Parameters

Before entering our discussion of the GA's performance and metrics, we must first discuss setup and assumptions. Specifically, we focus on our setup for GA termination. GA termination can be defined in numerous ways, including: as a function of the change from generation to generation, as a specific objective value, or as a max number of generations. In our study, to facilitate an effective comparison, we terminate based on a maximum number of generations evaluated by the GA. That is, termination is deterministic and is based on the number of GA populations that have been investigated. This is inherently more fair than a time constraint, which might punish a GA for evaluating many possible solutions. We discuss this concept further below.

Table 5.3 shows the GA setup parameters used for all runs. A set number of generations does not directly influence the running time of the GA. By design, as specified in Chapter 3.6.3, the running time of our objective function may not be consistent among genomes. Each genome contains information to build a single *logical* routing tree. Some genomes will be determined to be infeasible immediately. For instance, a logical tree that does not reach all

	$ \mathcal{D} $				
$ \mathcal{R} $	0	5	10	15	\forall
2	12.18	8.40	4.70	2.42	6.92
7	22.70	20.18	10.75	6.44	15.02
12	42.10	40.81	26.79	14.66	31.09
\forall	25.67	23.13	14.00	7.84	17.68

Table 5.4: Average GA Run Times (s)

	Pct^{inbeam}		
Beamwidth	50%	90%	\forall
30	14.49	12.10	13.30
60	20.45	14.47	17.46
90	27.83	16.74	22.28
\forall	20.92	14.44	17.68

Table 5.5: Average GA Run Times (s)

receivers cannot possibly represent a feasible solution. Thus, it will be rejected immediately upon evaluation, and avoid the majority of the objective function.

For some network layouts, there will be a large number of feasible trees that must traverse the entire objective. In others, there will be only a few logically feasible trees, and running time will be small. Moreover, genomes are cached so as not to continually be evaluated by the objective function unnecessarily. So, as the GA converges to a population of similar individuals, the cached values dominate, and the GA quickly reaches the maximum generation count.

These two specifics of the GA instantiation can lead to drastically different running times, and we document this below.

Table 5.4 and Table 5.5 break down average running times for the GA in seconds, grouped by our familiar pairs of network parameters. There is a clear pattern of parameters where the GA converges in a short amount of time. For networks with many directional antennas, and antennas with high gain, the time to termination is rather short. This is likely due to the effects that the antennas have on overall lifetime. A near-optimal genome will have

a large lifetime value (and thereby evaluation function value), and will quickly dominate the population. Since the GA does not repeatedly check identical solutions, the evaluation function is essentially skipped, and the max number of generations is reached quickly.

The goal of the GA is to find a near optimal with less computational cost than the mixed-integer program. As we see, the average time to termination can vary between 2.5 to slightly over 40 seconds. Though there are identifiable trends with the network characteristics that are operated on, these data do not present immediate conclusions regarding performance. This section is intended to provide the reader with some data regarding computational cost, not as a comparison to the MILP or semi-distributed heuristic, but in terms of “user-utility.” Remember that the GA is written in a standard language and is therefore machine-independent. Contrast this with the commercial solver package CPLEX used to find the solution to the mixed-integer model. Because of the relatively short running times per simulated network, and the flexibility in using multiple machines, we find the GA to be a useful tool for our study.

In the next section, we focus on the performance of the GA in terms of our defined metrics.

5.3.1 GA Performance vs. Optimal

For the remainder of this chapter, we produce results that illustrate the performance of the GA relative to the MILP optima. As with the semi-distributed heuristic, we compare the GA to the MILP optima in both lifetime and cumulative power. However, unlike the previous chapters, we also compare the GA to the semi-distributed heuristic.

If placed on a scale of increasing computational complexity to find a solution, the methods presented in this document would fall, in order: semi-distributed heuristic, genetic algorithm, and mixed-integer program. Since the GA and the semi-distributed method are heuristic, they are both upper bounded by the MILP optima. Therefore, determining which of the two produces the closest approximation will assist us in deciding which has the best payoff

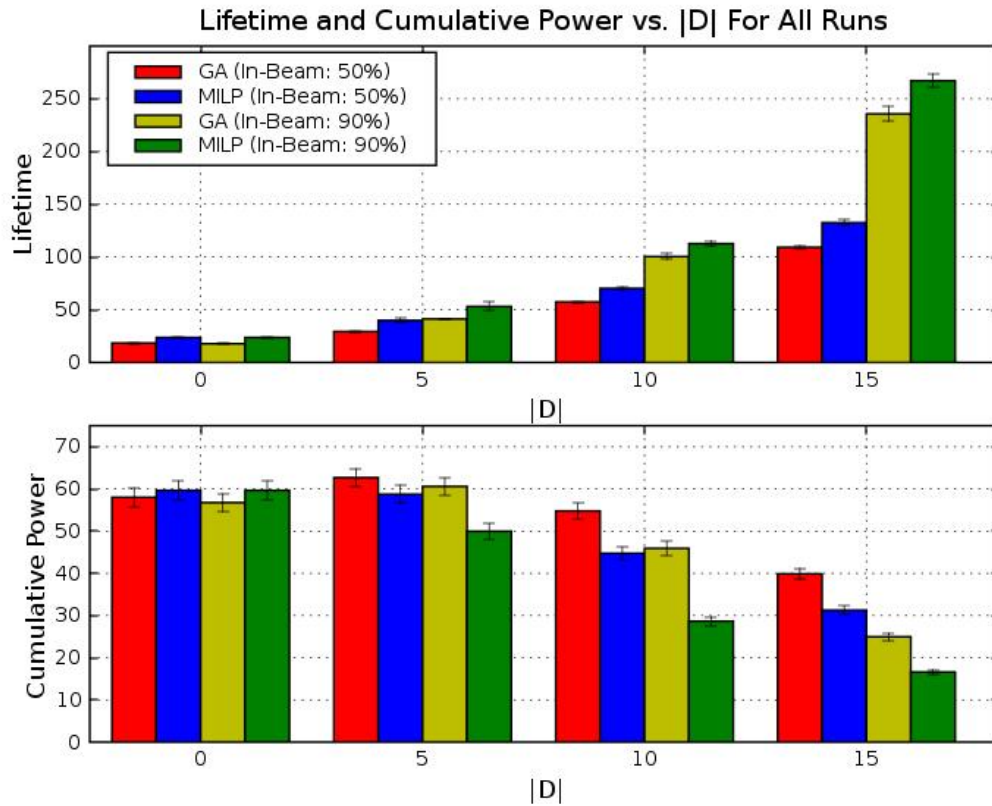


Figure 5.9: GA and MILP Lifetime and Cumulative Power

as an approximation method, and thereby, help the reader determine which method is most useful for the study.

Figure 5.9 shows the two metrics of concern for the GA and MILP, grouped and averaged by network characteristic. The trends first exhibited by our semi-distributed heuristic are also evident here, as the GA does an excellent job approximating lifetime values but at a significantly higher power cost. It is also quite apparent that the GA does not approximate the lifetime optimal decidedly better than the semi-distributed heuristic. Plots shown below further this analysis.

Figure 5.10 presents another view of the same data. Here, the GA solution lifetime is shown as a ratio on a per-run basis compared to the MILP lifetime. As with our semi-




Figure 5.10: GA / MILP Lifetime Ratios

distributed heuristic, we see that the GA is at its best when optimizing networks with numerous directional antennas. In many cases, more than half of the network runs are within 10% of the MILP optima. Notice that the omni-directional case shows rather poor performance for the GA as compared to the optimal. To understand this, we must reference Table 5.4 and Table 5.5. Termination times in these cases is longer than all others in the data set. This indicates that the objective function is traversed numerous times, most likely because there are many feasible trees that have approximately the same lifetime value. The GA struggles to find the tree that has *the optimal* solution to the problem. With a higher maximum generations count, the procedure has a higher chance of finding this configuration, but this comes at the cost of a much longer run time.




Figure 5.11: GA / MILP Power Ratios

Figure 5.11 illustrates that the GA also, at times, uses much more cumulative power than the MILP. However, note the vertical width of the box in each plot. On the average, the GA doesn't require an inordinate amount of extra cumulative power. However, the outliers do indicate that it can find topologies that are *much worse* in the secondary metric.

5.3.2 GA Performance vs. Distributed Heuristic

The GA is a heuristic method designed to approximate optimal values, similar to that of our semi-distributed heuristic. While GA's are well studied and generalized for a variety of problems, the semi-distributed heuristic is designed specifically for this problem with the

intention of having structure that could be exploited during implementation. In this section, we compare the two methods head-to-head to give the reader a better understanding of their respective performance.

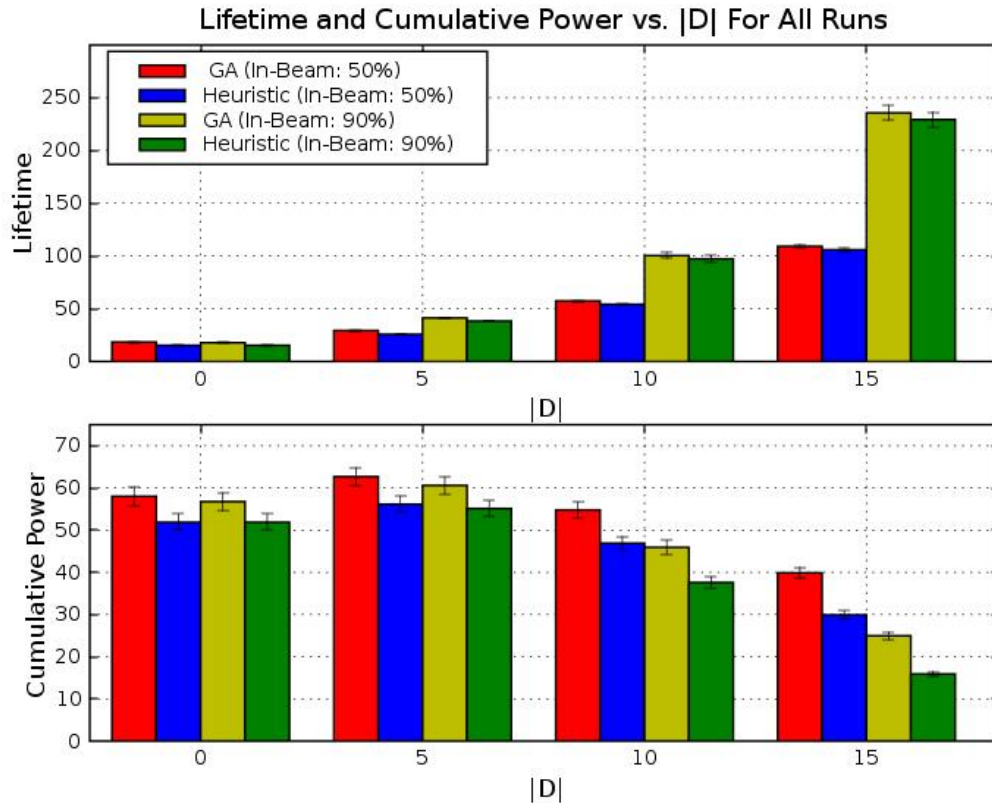


Figure 5.12: GA and Distributed Heuristic Metrics

Figure 5.12 shows the averaged performance of the GA when compared to our heuristic method. The GA does represent a slight improvement over the heuristic when finding optimal lifetime, but the overall difference is almost negligible. Worse, the cost of this marginal gain is the use of significantly more cumulative power.

Figure 5.13 is telling of the performance of the GA. The box is virtually invisible in the plot, since the GA and the heuristic often return virtually identical values. However, the outliers represent situations where the GA found a solution superior to the heuristic's. From




Figure 5.13: GA / Distributed Heuristic Lifetime Ratios

the semi-distributed heuristic study we know that, often, there exist topologies that have superior metrics. However, both our semi-distributed heuristic and our genetic algorithm have difficulty in finding these topologies.

Finally, Figure 5.14 validates our conclusion that the GA can use much more power than the heuristic or MILP. There are cases where this cumulative power value is *extremely* high, representing a power sub-optimal tree that also happens to have a comparable lifetime. The MILP and heuristic indirectly minimize power as they maximize lifetime. The GA, on the other hand, finds logical topologies that may or may not have small power settings. Instead, our defined objective function is based purely on lifetime, and therefore trees with *many* transmitting sources can be chosen over those with only a few.




Figure 5.14: GA / Distributed Heuristic Power Ratios

Because the GA is population-based, as time progresses, the entire population may have numerous genomes that are similar. Unfortunately, while these individuals have near-optimal lifetime values, they may also have little or no resemblance to the actual optimal configuration. From generation to generation, there is little chance that the GA will stumble upon the exact change that will lead to the optimal configuration.

Note that for most runs, as with our comparison to the MILP, the cumulative power ratio lies between 80% and 125% as shown by the small size of the box plot along the y-axis. The outliers represent a small number of cases, but skew the plot due to their *extremely* poor performance. While the GA certainly does use more cumulative power on average, the difference is actually fairly small.

5.3.3 Genetic Algorithm Results – Domain Reduction

One key use of a heuristic method is that it can be applied to the SINR constraints to significantly reduce the domain of the mixed-integer problem. From Equation 3.4 we know that a specific amount of received power (A) vs the interfering power (B) is required for communication. Recall that path loss and antenna parameters also determine how much power is received from a neighboring node. Because inter-node distances and beam configurations are assumed known on problem initialization, we can determine an absolute lower bound on power setting required for inter-node communication.

$$P_i^t \cdot G_i(\vec{j}) \cdot B_{j,i}(b, i) - S^i \cdot \left[\sum_{\substack{k \in \mathcal{N} \\ k \neq i, j}} P_k^t \cdot G_{k,j} \cdot B_{j,i}(b, k) + N_{thermal} \right] \geq \quad (5.1)$$

$$F_{i,j,b} \cdot U_{i,j,b} - U_{i,j,b}$$

$$P_i^t \cdot G_i(\vec{j}) \cdot B_{j,i}(b, i) - S^i \cdot N_{thermal} \geq 0 \quad (5.2)$$

In Equation 5.1 we repeat the SINR constraint for a single $i \rightarrow j$ link in the directional receive model. If we assume no other interfering sources (which clearly represents the min-power case for this SINR constraint) we are left with Equation 5.2. Best-case lifetime of the link can then be expressed as a function of energy reserves and the lower bound on the power requirement. This best-cast lifetime of the link represented in Equation 5.2 is shown in Equation 5.3.

$$Life_{i \rightarrow j} \leq \frac{R_i}{\{P_i^t : P_i^t \cdot G_i(\vec{j}) \cdot B_{j,i}(b, i) = S^i \cdot N_{thermal}\}} \quad (5.3)$$

We know that the MILP represents the network optimal, and that it upper bounds either heuristic method. Also, we know that by the max-min criterion of our lifetime metric, the network lifetime is determined by the *worst-case* link in the forwarding tree. Since the heuristic is a lower bound on MILP performance, and we have an expression for the upper

	$ \mathcal{D} $				
$ \mathcal{R} $	0	5	10	15	\forall
2	44.3%	52.1%	67.8%	77.7%	60.5%
7	27.5%	26.2%	46.1%	68.9%	42.4%
12	17.1%	16.1%	31.0%	65.8%	32.5%
\forall	29.7%	31.5%	48.3%	70.8%	45.1%

Table 5.6: Percent of Binary Variables Pruned in MILP

	$ \mathcal{D} $				
$ \mathcal{R} $	0	5	10	15	\forall
2	60.27	135.36	304.47	576.51	295.15
7	53.92	103.96	286.37	552.04	249.07
12	33.65	64.22	194.89	552.65	203.85
\forall	58.15	124.89	298.44	568.35	262.46

Table 5.7: Average Number of Binary Variables Pruned in MILP

bound on link lifetime, we can remove all links from the MILP having *best-case* lifetimes that are less than the heuristic solution. In other words, the MILP cannot possibly include these links, since doing so would lead to a lifetime *less than* the lower bound of the heuristic solution. Remember that each SINR constraint also requires one binary variable, and that the number of binary variables can have a profound effect on solve times.

Table 5.6 illustrates the average number of binary variables that can be pruned by using the GA lifetime as a lower bound when constructing the MILP. Remember that networks with large numbers of directional antennas and few receivers have the highest lifetime, which is reflected in the pruning results. The higher lifetime means we can prune a larger number of low lifetime SINR constraints, as many as 77% in the best case.

Table 5.7 shows the *actual* number of binary variables that are pruned on average. This table is intended to show that while a high percentage of pruned links is certainly a modeling improvement, that the actual *number* of SINR links remaining may still be quite large. For instance, the worst pruning percentage is when $|\mathcal{R}| = 12$ and $|\mathcal{D}| = 5$. Here, only 64 links are pruned from the network, as shown in Table 5.7, leaving 400 binary variables. When $|\mathcal{R}| = 2$ and $|\mathcal{D}| = 15$, almost 78% of the SINR links can be removed on average. However,

as per Table 5.7, 577 total links are removed, leaving 749 SINR variables in the program.

As you can see, a higher *percentage* of pruned links does not necessarily indicate a simpler program. In many cases, the initial program had more variables to begin with, and the pruning only helps to reduce the number of binary variables. However, any reduction in the number of binary variables and the overall problem size helps reduce solution times.

Chapter 6

Conclusion

New applications and users for wireless networks are appearing every day. Infrastructure-based wireless systems are nearly pervasive. Wi-fi is available at colleges, coffee shops, hotels, hospitals, even city parks. Cellphones are nearly ubiquitous and have become commoditized, with users ranging from grade-schoolers to grandmothers. People have found that the ability to communicate and gather information without being tethered to a phone line or network connection provides immense convenience.

While these products represent a complete paradigm shift from the standard “Internet” grid connectivity, they still require massive investment of capital and effort for the supporting infrastructure. Prognosticators have conjectured that the next shift in technological innovation for networks will be the introduction of truly infrastructure-less, self-organizing systems. Rather than requiring that infrastructure be in place a priori to provide connectivity to user devices, the user devices themselves will provide their own connectivity. Wherever the devices roam, they will carry the network with them. By working together, each user device or network node becomes a “service-provider” to neighbors.

The ability to communicate without wires has been a boon to mobile communication. Unfortunately, removing the necessity to be “wired” to communicate has not removed the need

for grid connectivity altogether. Instead, it has brought the biggest weakness of independent devices to the forefront.

Mobile devices still require some source of power.

For any kind of independence, nodes must carry their own power reserves, reserves that certainly have one characteristic in common. Regardless of the type of power stores, the amount of energy available is finite.

Transmitting information via RF has non-trivial power requirements. Moreover, the amount of power required to send a message is not a set quantity, it can be a function of the network layout and wireless characteristics. The rate at which nodes expend power determines how long the system can operate before nodes must be recharged. Alternatively, in the event that nodes *cannot* be recharged, nodes who have expended their energy reserves can be considered “dead” and can no longer participate in the network.

Finding the best possible configuration of the network topology can save a significant amount of network resources, and consequently, enhance overall utility. In wireless networks where nodes transmit their information on a shared channel, a single session must account for other transmitting sources when determining the power level necessary for communication. Power requirements are not static, rather they are dependent upon the state of all other network nodes at the time of communication. Clearly, this becomes a global optimization problem, where system performance depends on the behavior of all constituent network nodes.

6.1 Summary of Contributions

In this document, we present an algorithmic study of the effects of directional antennas on multicast lifetime in wireless ad-hoc networks.

Numerous previous papers have investigated the case where directional antennas are used to *transmit* a signal to a set of receivers listening omni-directionally. Unfortunately, most

of this analysis was heuristic based, making objective performance claims impossible. These papers have also assumed that some other form of multiplexing was active, allowing different sessions to operate without interference.

Here, we consider using directional antennas for the *reception* of signals in environments where interference is an issue. Because this method has not received attention from the research community, we first compare directional transmission and directional reception directly. Mixed-integer programs designed to optimize multicast lifetime are presented for both the case where antennas are used for directional transmission and where antennas are used for reception. Individual network files are input to each MILP, and upon completion, optimal path lifetime and a secondary metric of cumulative power are measured and compared among the two competing methods.

Our results overwhelmingly indicate that directional reception is the superior alternative to directional transmission. Directional reception has superior path lifetime when compared to directional transmission, and accomplishes this with *less* cumulative power. This superiority is consistent across all our tested network parameters.

We also show that directional reception heuristics perform much better than directional transmission heuristics. Both our specifically designed semi-distributed heuristic method and the genetic algorithm implementation closely approximate the network optima. As part of this study, we also analyzed directional transmission and reception heuristics designed to operate in non-interfering environments. It is shown that interference-ignorant heuristics produce routing topologies that are substandard (and in the worst case infeasible) when mapped into interfering environments.

Both the heterogeneity study and the heuristic study show strongly that regardless of how they are used, directional antennas cannot entirely overcome interference and coupling of wireless environments. Methods designed to operate in non-interfering environments cannot be mapped into environments where interference plays a role. Doing so leads to networks that are infeasible, or that have metric results *much* poorer than our suggested heuristics.

In short, this document illustrates two major conclusions. First, directional reception is superior to directional transmission for extending multicast path lifetime and conserving power. Second, we find that interference is a significant contributor to network performance even when directional antennas are available for spatial multiplexing. Previous research that ignores this fact is essentially considering an entirely different problem, and is *dependent* upon the operation of complementary multiplexing techniques.

6.2 Future Work

From the evidence presented in this document, it is clear that directional reception is superior to directional transmission. This conclusion opens up numerous new research areas for future investigation. Some are direct extensions of the work presented here, and others are more tangential, but derived from the intuition gained from this study. In this final section, we discuss some of these areas and possible approaches.

- **Cooperation**

The MILP and heuristic methods presented in this document (along with most previous research) assumes that the entire network is comprised of nodes willing to cooperate with the suggested scheme. Our semi-distributed heuristic, for instance, requires that nodes who “hear” a transmitting source at a power level surpassing the SINR constraint to “volunteer” to forward the transmitter’s packets. By doing so, they may act as an intermediate node for other neighbors. A branch of contemporary research is investigating the effects of cooperation on network performance. The number of nodes willing to participate, their probability of participation, and possible nefarious nodes could all affect the overall system optimum. Conversely, nodes may be magnanimous, willing to sacrifice themselves for the performance of the overall system. The possibilities for study in this area are varied and exciting, and we envision this problem to lend itself well to future research.

- **Theoretical Relationship: Max-Min Lifetime and Min Power**

As we have discussed and shown with our results, max-min lifetime and minimum power metrics are not necessarily strongly correlated. Unfortunately, as yet, a theoretical relationship between these and other optimization goals is unknown. From our MILP program construction, we can change the optimization target by simply modifying the objective function. The MILP might then provide a tool by which the two goals can be compared and contrasted to develop a better understanding of how the two metrics are related. Approximation algorithms and duality relationships may provide a framework by which the MILP can be leveraged to discover theoretical bounds on the relationship between different, but related metrics.

• Performance Over Multiple Multicasts

Based upon our research proposal, the goal of this document was to compare the performance of directional transmission and reception for a single multicast. Our results show that directional reception has higher lifetimes for a given network, but they also indicate that the cumulative power use is less than that of directional transmission. Therefore, we expect that over multiple multicasts, a network using directional reception should perform well *much* longer than a network using directional transmission. This is clearly the most direct descendant of the work presented in this document, and we expect it to further the case that directional reception is a superior alternative to directional transmission.

• Interference Pricing

From our SINR constraint and discussion of variable pruning, we know that all non i, j transmitting nodes contribute to interference for an $i \rightarrow j$ session. In a sense, we could consider this interference to be an increase in “price” for the i, j session. The final tree topology is dependent upon node positions, beam parameters, and prices for the entire system. From an O.R. perspective, the study of “prices” leads naturally into a study of our mathematical program’s dual. This duality study could yield more effective heuristics, and potentially loosely coupled distributed methods. A similar type of study was performed by Low and Lapsley [75] to optimize flow rates on networks with shared links. Instead of

optimizing flow, we are optimizing global network utility based on a shared channel model. However, the idea of pricing of shared resources is similar in either case.

Bibliography

- [1] D. E. Comer, *Internetworking with TCP/IP Principles Protocols and Architectures*. Prentice Hall, 2000.
- [2] A. Das, R. Marks, M. El-Sharkawi, P. Arabshahi, and A. Gray, “Minimum Power Broadcast Trees for Wireless Networks; Integer Programming Formulations,” in *Proc. of IEEE INFOCOM*, pp. 1001–1010, March 2003.
- [3] S. Singh, C. Raghavendra, and J. Stepanek, “Power-aware Broadcasting in Mobile Ad-hoc Networks,” in *Proc. of IEEE PIMRC’99*, vol. 1, pp. 153–157, September 1998.
- [4] M. Zagalj, J.-P. Hubaux, and C. Enz, “Minimum Energy Broadcast in All-Wireless Networks: NP-Completeness and Distribution Issues,” in *Proc. of ACM MobiCOM*, pp. 172–182, 2002.
- [5] O. Egecioglu and T. Gonzalez, “Minimum-energy Broadcast In Simple Graphs With Limited Node Power,” in *Proc. of IASTED International Conference on Parallel and Distributed Computing and Systems*, pp. 334–338, August 2001.
- [6] P. Wan, G. Calinescu, X. Li, and O. Frieder, “Minimum-Energy Broadcast Routing In Static Ad-Hoc Wireless Networks,” in *Proc. of IEEE INFOCOM*, vol. 2, pp. 1162–1171, 2001.
- [7] W. Liang, “Constructing Minimum-Energy Broadcast Trees in Wireless Ad-Hoc Networks,” in *Proc. of MobiHOC*, pp. 112–122, 2002.

- [8] J. Chang and L. Tassiulas, "Energy Conserving Routing in Wireless Ad-Hoc Networks," in *Proc. of IEEE INFOCOM*, vol. 1, pp. 22–31, 2000.
- [9] B. Wang and S. Gupta, "On Maximizing Lifetime of Multicast Trees in Wireless Ad-Hoc Networks," in *International Conference On Parallel Processing (ICPP-03)*, pp. 333–340, October 2003.
- [10] P. Floreen, J. Kohonen, P. Kaski, and P. Orponen, "Multicast Time Maximisation in Energy Constrained Wireless Networks," in *Proc. of the 2003 Joint Workshop on Foundations of Mobile Computing*, pp. 50–58, September 2003.
- [11] J. Wieselthier, G. D. Nguyen, and A. Ephremides, "Algorithms For Energy-Efficient Multicasting in Static Ad-Hoc Wireless Networks," in *ACM Mobile Networks and Applications (MONET)*, pp. 251–263, June 2001.
- [12] S. Guo and O. Yang, "Antenna Orientation Optimization for Minimum-Energy Multicast Tree Construction in Wireless Ad Hoc Networks with Directional Antennas," in *Proc. of MobiHoc*, pp. 234–243, May 2004.
- [13] S. Guo and O. Yang, "Minimum Energy Multicast Routing for Wireless Ad-hoc Networks with Adaptive Antennas," in *Proc. of International Conference on Network Protocols (ICNP)*, pp. 151–160, October 2004.
- [14] C. Balanis, *Antenna Theory: Analysis and Design*. John Wiley and Sons, 1997.
- [15] J. Wieselthier, D. Nguyen, and A. Ephremides, "Energy-aware wireless networking with directional antennas: The case of session-based broadcasting and multicasting," in *IEEE Transactions on Mobile Computing*, vol. 1, pp. 176–191, 2002.
- [16] J. Wieselthier, G. Nguyen, and A. Ephremides, "Energy-Limited Wireless Network with Directional Antennas: The Case of Session-Based Multicasting," in *Proc. of IEEE INFOCOM*, pp. 190–199, 2002.

- [17] A. Sankar and Z. Liu, “Maximum Lifetime Routing in Wireless Ad-Hoc Networks,” in *Proc of INFOCOM*, vol. 2, pp. 1089–1097, March 2004.
- [18] C. Hua and T.-S. P. Yun, “Maximum Lifetime Routing and Data Aggregation for Wireless Sensor Networks,” in *Proc. of 5th International IFIP-TC6 Networking Conference*, pp. 840–855, May 2006.
- [19] J. Wieselthier, G. Nguyen, and A. Ephremides, “On the Construction of Energy-Efficient Broadcast and Multicast Tree in Wireless Networks,” in *Proc. of IEEE INFOCOM*, pp. 585–594, March 2000.
- [20] W. Heinzelman, A. Chandrakasan, and H. Balakrishnan, “Energy-Efficient Communication Protocol for Wireless Microsensor Networks,” in *Proc. of the 33rd Hawaii International Conference on System Sciences (HICSS 2000)*, vol. 2, pp. 10–12, 2000.
- [21] V. Rodoplu and T. Meng, “Minimum Energy Mobile Wireless Networks,” in *IEEE Journal on Selected Areas in Communications*, vol. 17(8), pp. 1333–1344, August 1999.
- [22] I. Kang and R. Poovendran, “Maximizing Network Lifetime Of Broadcasting Over Wireless Stationary Ad-Hoc Networks,” in *ACM/Kluwer Mobile Networks and Applications (MONET)*, vol. 10, December 2005.
- [23] A. Ahluwalia, E. Modiano, and L. Shu, “On the Complexity and Distributed Construction of Energy-Efficient Broadcast Trees in Static Ad-Hoc Wireless Networks,” in *Proc. Conf. Inf. Sci. Syst.*, pp. 807–813, March 2002.
- [24] J. Wieselthier, G. Nguyen, and A. Ephremides, “Energy-Efficient Broadcast and Multicast Trees in Wireless Networks,” in *ACM/Kluwer Mobile Netw. Appl. (MONET)*, vol. 7, pp. 481–492, December 2002.
- [25] F. Li and I. Nikolaidis, “On Minimum Energy Broadcasting in All-Wireless Networks,” in *Proc. of IEEE Conference on Local Computer Networks (LCN)*, pp. 192–202, November 2001.

- [26] R. Wattenhofer, L. Li, P. Bahl, and Y. Wang, "Distributed Topology Control For Power Efficient Operation In Multihop Wireless Ad Hoc Networks," in *Proc. of IEEE INFOCOM*, pp. 1388–1397, April 2001.
- [27] Q. Li, J. Aslam, and D. Rus, "Online Power-Aware Routing In Wireless Ad-Hoc Networks," in *Proc. of MobiCOM*, pp. 97–107, 2001.
- [28] B. Wang and S. K. S. Gupta, "G-REMiT: An Algorithm for Building Energy Efficient of Multicast Trees in Wireless Ad Hoc Networks," in *Network Computing and Applications (NCA)*, pp. 265–272, April 2003.
- [29] P. Floreen, P. Kaski, J. Kohonen, and P. Orponen, "Lifetime Maximization for Multicasting in Energy-Constrained Wireless Networks," in *IEEE Journal on Selected Areas in Communications*, vol. 23, pp. 117–126, January 2005.
- [30] Y. T. Hou, Y. Shi, H. D. Sherali, and J. E. Weieselthier, "Online Lifetime-Centric Multicast Routing for Ad Hoc Networks with Directional Antennas," in *Proc. of IEEE INFOCOM*, pp. 761–772, March 2005.
- [31] B. Wang and S. K. S. Gupta, "S-REMiT: A Distributed Algorithm for Source-based Energy Efficient Multicasting in Wireless Ad Hoc Networks," in *IEEE GLOBECOM*, vol. 6, pp. 3519–3524, December 2003.
- [32] "The Network Simulator – ns-2," July 2006; <http://www.isi.edu/nsnam/ns/>.
- [33] J. Wieselthier, G. Nguyen, and A. Ephremides, "Resource Management in Energy-Limited, Bandwidth-Limited, Transceiver-Limited Wireless Networks For Session-Based Multicasting," in *Computer Networks*, vol. 39, pp. 113–131, June 2002.
- [34] I. Kang and R. Poovendran, "Maximizing Static Network Lifetime of Wireless Broadcast Ad-Hoc Networks," in *Proc. ICC*, vol. 3, pp. 2256–2261, May 2003.
- [35] R. M. II, A. K. Das, and M. El-Sharkawi, "Maximizing Lifetime In An Energy Constrained Wireless Sensor Array Using Team Optimization of Cooperating Systems," in

- Proc. International Joint Conference of Neural Networks, IEEE World Congress on Computational Intelligence*, pp. 371–376, May 2003.
- [36] A. Orda and B.-A. Yassour, “Maximum-Lifetime Routing Algorithms for Networks with Omnidirectional and Directional Antennas,” in *Proc. of MobiHOC*, pp. 426–437, May 2005.
- [37] I. Maric and R. D. Yates, “Cooperative Multicast for Maximum Network Lifetime,” in *IEEE Journal on Selected Areas in Communications*, vol. 23, pp. 127–135, January 2005.
- [38] M. Agarwal, J. Cho, L. Gao, and J. Wu, “Energy Efficient Broadcast in Wireless Ad-Hoc Networks with Hitch-Hiking,” in *Proc. of IEEE INFOCOM*, pp. 2096–2107, 2004.
- [39] J. Zander, “Slotted ALOHA Multihop Packet Radio Networks With Directional Antennas,” *Electronics Letters*, vol. 26, no. 25, pp. 2098 – 2100, 1990.
- [40] R. Ramanathan, “On the Performance of Ad Hoc Networks with Beamforming Antennas,” in *Proc. of ACM MobiHoc*, pp. 95–105, 2001.
- [41] IEEE 802.11, *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, Aug. 1999.
- [42] R. R. Choudhury and N. Vaidya, “Deafness: a MAC Problem in Ad Hoc Networks when Using Directional Antennas,” tech. rep., University of Illinois at Urbana-Champaign, 2003.
- [43] P. Karn, “MACA - A New Channel Access Method for Packet Radio,” in *Proc. of ARRL/CRRL Amateur Radio 9th Computer Networking Conference*, pp. 134–140, Sep. 1990.
- [44] Y.-B. Ko, V. Shankarkumar, and N. H. Vaidya, “Medium access control protocols using directional antennas in ad hoc networks,” in *Proc. of INFOCOM*, pp. 13–21, 2000.

- [45] A. Nasipuri, S. Ye, J. You, and R. Hiromoto, “A MAC Protocol for Mobile Ad Hoc Networks Using Directional Antennas,” in *Proc. of IEEE WCNC*, pp. 1214–1219, 2000.
- [46] T. Korakis, G. Jakllari, and L. Tassiulas, “A MAC protocol for full exploitation of Directional Antennas in Ad-Hoc Wireless Networks,” in *Proc. of ACM MobiHoc*, pp. 98–107, 2003.
- [47] R. R. Choudhury, X. Yang, R. Ramanathan, and N. Vaidya, “Using Directional Antennas for Medium Access Control in Ad Hoc Networks,” in *Proc. of ACM MobiCom*, pp. 59–70, September 2002.
- [48] T. El-Batt, T. Anderson, and B. Ryu, “Performance Evaluation of Multiple Access Protocols for Ad Hoc Networks Using Directional Antennas,” in *Proc. of IEEE WCNC*, pp. 982–987, March 2003.
- [49] M. Takai, J. Martin, R. Bagrodia, and A. Ren, “Directional Virtual Carrier Sensing for Directional Antennas In Mobile Ad Hoc Networks,” in *Proc. of ACM MobiHoc*, pp. 39–46, 2002.
- [50] M. Takai, J. Zhou, and R. Bagrodia, “Adaptive Range Control Using Directional Antennas In Mobile Ad Hoc Networks,” in *Proc. of ACM MSWiM*, pp. 92–99, September 2003.
- [51] R. Choudhury and N. Vaidya, “Ad hoc routing using directional antennas,” tech. rep., University of Illinois Urbana Champaign, 2002.
- [52] R. Choudhury and N. Vaidya, “MAC-Layer Anycasting in Ad Hoc Networks,” tech. rep., University of Illinois Urbana Champaign, 2003.
- [53] S. Roy, D. Saha, S. Bandyopadhyay, T. Ueda, and S. Tanaka, “A Network-Aware MAC and Routing Protocol for Effective Load Balancing in Ad Hoc Wireless Networks with Directional Antenna,” in *Proc. of ACM MobiHoc*, pp. 88–97, 2003.

- [54] S. Yi, Y. Pei, and S. Kalyanaraman, “On the Capacity Improvement of Ad Hoc Wireless Networks Using Directional Antennas,” in *Proc. of MobiHoc*, pp. 108–116, June 2003.
- [55] P. Gupta and P. Kumar, “Capacity of wireless networks,” tech. rep., University of Illinois, Urbana-Champaign, 1999.
- [56] G. Nemhauser and L. Wolsey, *Integer and Combinatorial Optimization*. John Wiley and Sons, 1988.
- [57] J. Hooker, *Logic-Based Methods for Optimization: Combining Optimization and Constraint Satisfaction*. Wiley-Interscience, 2000.
- [58] E. Balas, “Management Sciences Research Report No. 330,” tech. rep., Graduate School of Industrial Administration, Carnegie Mellon, 1974.
- [59] E. Balas, *Nonlinear Programming 2*. Academic Press, 1975.
- [60] N. Beaumont, “An Algorithm For Disjunctive Programs,” in *European Journal of Operational Research*, vol. 48, pp. 362–371, 1990.
- [61] H. Williams, *Model Building in Mathematical Programming*. John Wiley and Sons, 1999.
- [62] J. N. Hooker, H. Yan, I. E. Grossman, and R. Raman, “Logic Cuts for Processing Networks with Fixed Charges,” in *Computers and Operations Research*, vol. 21, pp. 265–279, 1994.
- [63] E. Balas, “Disjunctive Programming and a Hierarchy of Relaxations for Discrete Optimization Problems,” in *SIAM Journal on Algebraic and Discrete Methods*, vol. 6, pp. 466–485, July 1985.
- [64] J. Hooker, “Convex programming methods for global optimization,” in *COCOS*, 2003.
- [65] E. Balas and R. Jeroslow, “Canonical Cuts On the Unit Hypercube,” in *Journal on Applied Mathematics*, vol. 23, pp. 61–69, 1972.

- [66] K. Sundaresan and R. Sivakumar, “A Unified MAC Layer Framework for Ad-hoc Networks with Smart Antennas,” in *Proc. of ACM MobiHoc*, pp. 244–255, May 2004.
- [67] F. Hillier and G. Lieberman, *Introduction to Operations Research*. McGraw Hill, 2001.
- [68] R. D. Yates, “A Framework For Uplink Power Control In Cellular Radio Systems,” in *IEEE Journal on Selected Areas in Communications*, vol. 13, pp. 1341–1347, September 1995.
- [69] J. Holland, *Adaptation in Natural and Artificial Systems*. MIT Press, 1975.
- [70] ”GAlib: Matthew’s Genetic Algorithms Library,” July 2006; <http://lancet.mit.edu/ga/>.
- [71] T. E. Cormen, C. E. Leiserson, and R. L. Rivest, *Introduction to Algorithms*. MIT Press, 2001.
- [72] K. N. Wood and L. A. DaSilva, “Directional Reception vs. Directional Transmission for Maximum Lifetime Multicast Delivery in Ad-Hoc Networks,” in *Proc. of 5th International IFIP-TC6 Networking Conference*, pp. 1169–1175, May 2006.
- [73] ”GLPK – GNU Project – Free Software Foundation (FSF),” July 2006; <http://www.gnu.org/software/glpk/>.
- [74] ”ILOG CPLEX: High Performance software for mathematical programming and optimization,” July 2006; <http://www.ilog.com/products/cplex/>.
- [75] S. H. Low and D. E. Lapsley, “Optimization flow control – i: Basic algorithm and convergence,” in *IEEE/ACM Transactions on Networking*, vol. 7, pp. 861–874, December 1999.

Appendix A

Power and Semi-Distributed Algorithms

Algorithm 3 powerAlgorithm(\mathcal{L})

```
1: while  $\delta \notin [0, \Delta]$  do
2:   for all  $n \in \mathcal{N}$  do
3:     if children( $n$ )  $\neq \{\}$  then
4:        $\mathcal{C} \leftarrow$  children( $n$ )
5:       worst_SINR =  $\min_{c \in \mathcal{C}} \text{SINR}(n, c)$ 
6:       if  $|\text{worst\_SINR} - \text{SINR}^{req}| > \delta$  then
7:          $\delta = \text{worst\_SINR} - \text{SINR}^{req}$ 
8:       end if
9:       
$$P_n^t = P_n^t + \omega \cdot \left[ \frac{[\text{SINR}^{req} + \frac{\Delta}{2}] - \text{worst\_SINR}}{B_{n,b}(c)} \right]$$

10:    end if
11:    if  $P^t \geq 2 \cdot P^{max}$  then
12:      return INFEASIBLE;
13:    end if
14:  end for
15: end while
```

Algorithm 4 Semi-Distributed Algorithm

```

1:  $P_s^t = P^{max}$ 
2: //  $\mathcal{T}$  denotes transmitting nodes
3:  $\mathcal{T} \leftarrow s$ 
4: //  $\mathcal{I}$  denotes nodes included in tree
5:  $\mathcal{I} \leftarrow s$ 
6: //  $\mathcal{L}$  denotes active links
7: for  $\forall r \in \mathcal{R}$  do
8:    $\mathcal{I} \leftarrow r$ 
9:    $\mathcal{L} \leftarrow \{(s, r)\}$ 
10: end for
11: //  $\mathcal{S}$  denotes all silent nodes
12: for  $\forall n \in \mathcal{N}, n \notin \mathcal{T}$  do
13:    $\mathcal{S} \leftarrow n$ 
14: end for
15: powerAlgorithm( $\mathcal{L}$ )
16:  $\delta = 1$ 
17: while  $\delta > 0$  do
18:    $\delta = 0$ 
19:    $l = \text{getLifetime}(\mathcal{L})$ 
20:   for  $\forall s \in \mathcal{S}$  do
21:     for  $\forall t \in \mathcal{T}$  do
22:       if  $\text{SINR}(t, s) \geq \text{SINR}^{req}$  then
23:         
$$P_s^t = \frac{R^s}{\text{upstreamLife}(t)}$$

24:         for  $\forall i \in \mathcal{I}, i \neq s, i \neq t$  do
25:            $\hat{\mathcal{L}} = \mathcal{L}$ 
26:           if  $\text{SINR}(s, i) < \text{SINR}^{req}$  then
27:              $\hat{\mathcal{L}} \setminus \text{getParent}(i)$ 
28:              $\hat{\mathcal{L}} \leftarrow \{(s, i)\}$ 
29:              $\hat{\mathcal{L}} \leftarrow \{(t, s)\}$ 
30:             powerAlgorithm( $\hat{\mathcal{L}}$ )
31:             if  $\text{lifeTime}(\hat{\mathcal{L}}) - l > \delta$  then
32:                $\delta = \text{lifeTime}(\hat{\mathcal{L}}) - l$ 
33:                $\text{toAdd} = \{(s, i), (t, s)\}$ 
34:                $\text{toRemove} = \text{getParent}(i)$ 
35:             end if
36:           end if
37:         end for
38:       end if
39:     end for
40:   end for
41:   if  $\delta > 0$  then
42:      $\mathcal{L} \setminus \text{toRemove}$ 
43:      $\mathcal{L} \leftarrow \text{toAdd}$ 
44:   end if
45: end while

```
