# A NOVEL STRATEGY FOR HIGH THROUGHPUT IN WIRELESS Ad- Hoc NETWORKS

*A Thesis submitted to*

**SHOBHIT UNIVERSITY, MEERUT**

*For the award of the degree of*

**DOCTOR OF PHILOSOPHY**

*In*

**COMPUTER ENGINEERING**

**By**

**Arun Kumar Singh**

*Under the Supervision of*

**Prof. (Dr.) R. P. Agarwal**
**&**
**Prof. (Dr.) Neelam Srivastava**


**Faculty of**
**Computer Engineering and Information technology**
**Shobhit University, Meerut-250110**

# Declaration by the Candidate

I, hereby, declare that the work presented in this thesis, entitled **A NOVEL STRATEGY FOR HIGH THROUGHPUT IN WIRELESS Ad- Hoc NETWORKS** in fulfillment of the requirements for the award of Degree of Doctor of Philosophy, submitted in the School of/Centre for **Computer Engineering and Information technology** at Shobhit University, Modipuram, Meerut is an authentic record of my own research work carried out under the supervision of Prof. (Dr.) Neelam Srivastava.

I also declare that the work embodied in the present thesis

(i) is my original work and has not been copied from any Journal/thesis/book, and

(ii) has not been submitted by me for any other Degree or Diploma of any university/institution.

**Signature of the candidate**

# Certificate of the Internal Supervisor(s)

This is to certify that the thesis entitled "**A NOVEL STRATEGY FOR HIGH THROUGHPUT IN WIRELESS Ad- Hoc NETWORKS**" submitted by **Arun Kumar Singh** for the award of Degree of Doctor Philosophy in the School of / Center for **Computer Engineering and Information technology** of Shobhit University, Meerut is a record of authentic work carried out by him under my supervision.

To the best of my knowledge, the matter embodied in this thesis is the original work of the candidate and has not been submitted for the award of any other degree or diploma.

It is further certified that he has worked with me for a period of Dec-2008 to Feb-2013 in the School/ Center/ Department of **Computer Engineering and Information technology** of Shobhit University, Meerut.

**Prof. (Dr.) R. P. Agarwal**
**(Administrative Supervisor)**

# Certificate of the External Supervisor(s)

This is to certify that the thesis entitled "**A NOVEL STRATEGY FOR HIGH THROUGHPUT IN WIRELESS Ad- Hoc NETWORKS**" submitted by **Arun Kumar Singh** for the award of Degree of Doctor Philosophy in the School of / Center for **Computer Engineering and Information technology** of Shobhit University, Meerut is a record of authentic work carried out by him/her under my supervision.

To the best of my knowledge, the matter embodied in this thesis is the original work of the candidate and has not been submitted for the award of any other degree or diploma.

It is further certified that he has worked with me for a period of Dec-2008 to Feb-2013 in the School/ Center/ Department of Electronics and Communication Engineering, IET (A Constituent College of GBTU, Lucknow) Lucknow.

**Prof. (Dr.) Neelam Srivastava**
**(External Supervisor)**

# Acknowledgments

Some graduates write terse acknowledgments thanking exactly the select few who sped them along their way. Others ramble on, trying to give thanks to everyone who helped throughout the years of toil, only to glance at their dissertation two weeks later and realize they left out key individuals.

You've probably already guessed this is more like the second than the first. This work took me a long time, and I had a lot of help.

I am most grateful to my guide Prof. (Dr.) Neelam Srivastava and Prof. (Dr.) R.P. Agarwal who shepherded this work through the lengthy stages of implementation, debugging, experimentation and writing. She/he held me to the highest standards of quality and accuracy, and any value in my work is thanks to those high standards. I am also grateful for her/ his frank opinions and advice, and for listening when I was confused.

Dr. Neelam Srivastava made significant contributions to this work from the beginning. She helped deploy the first version of the test network, designed and ran many of the earlier experiments which motivate this work, and wrote the DSR-DSDV routing protocol implementation that I used.

Amar Deep Gupta non-stop hacking made possible the second version of the test-bed network, which connected every node to the wired Ethernet. His work made debugging and running experiments at least an order of magnitude easier.

Ankur Garg created the Click modular router, which allowed elegant implementations of PTC and the routing protocols. Click was especially useful for debugging, testing, and simulating my protocol implementations. Click is an extremely powerful tool. Mr. Bhoopendra Kumar spent many hours fixing Click so that it could do everything I needed. Mr. Varun Kumar Singh also wrote the writs typing break program, without which my wrists and back would hurt very much.

Mr. Vipin Tyagi's work on the capacity of multi-hop routes was an important result that led to the idea of the PTC metric.

I want to thank everyone in the Parallel and Distributed Operating Systems research group and on the floor at Shobhit University for making it such a stimulating place to work and learn. Mr. Kapil Agarwal, Sandeep Rana, and Vivek Sen Saxena were especially helpful, and

I've learned a lot from them. My office mates at DJCET-Modinagar, Gzb and DIT school of Engineering, greater Noida; deserve special thanks for putting up with me over the years, and for being patient with my questions: Mr. Raghav, Praveen Sharma, Mr. Ajay Verma and Mr. Ankit.

I also owe a very special thanks to all the lab members who graciously agreed to host a test node in their office. They put up with me tramping in and out of their offices for over two years, crawling under their desks, and making the machines beep and grind their disks when I rebooted the network several times a day. The poem in the front of this dissertation is dedicated to them.

Life, however, is not all about the lab, and I was lucky enough to have many good friends to remind me of that.

Finally, my Wife Ranjana Singh supported me over the years with my lovely daughter Arunjana Singh (BEBO), even though I didn't always know how to explain what I was trying to do.

My graduate program was partially supported by the Prof. (Dr.) G.S. Sandhu (Director, DIT School of Engineering, Greater Noida) and Prof. (Dr.) S.C.Gupta (Director, DJCET, Modinagar, Ghaziabad).

Last but not least my personally heart full thanks to Sir Kunwar Shekhar Vijendra (Pro-Chancellor of Shobhit University, Meerut) and Prof. (Dr.) R. P. Agarwal (VC-Shobhit University, Meerut), without his support it's not possible to submit this.

Arun Kumar Singh

# Table of Contents

# List of Figures

# Chapter 1: Introduction

## 1.1 Introduction

The present research deals with the ways to find high-throughput routes in multi-hop wireless packet networks. With the use of the *Potential Transmission Count (PTC)* metric presented here, routing protocols can find multi-hop routes that have up to twice of the throughput that has been found using the minimum hop-count metric. Most routing protocols reduce the hop-count metric, the number of wireless links in a route, regardless of the performance of each link. Routes preferred by the hop-count metric often contain lossy links, as a multi-hop wireless networks are expected to contain many lossy links, which reduce throughput. PTC prefers shorter routes with better links because it always selects high-throughput routes therefore; the PTC metric is based on the loss ratio of each link in a route, as well as, the number of links in a route.

Throughput is not the only property that requires concern of the network users. For example, voice and interactive users prefer low delay, while video users want to minimize jitter, which is the variability in delay and throughput. The focus of this work is to find out these applications and many other benefits from increased throughput.

## 1.2 Multi-hop Wireless Networks

A multi-hop wireless network is a network of computers and devices (*nodes*) connected by wireless communication *links*. The links are most often connected with digital packet radios. Because each radio link has a limited communication range, many pairs of nodes cannot communicate directly, and must forward data to each other via one or more cooperating intermediate nodes. We will often use ‘hop-count metric’ to mean the minimum hop-count metric.

A source node transmits a packet to a neighboring node with which it can communicate directly. The neighboring node in turn transmits the packet to one of its neighbors, and so on until the packet is transmitted to its ultimate destination. Each link that a packet is sent over is referred to as a *hop*; the set of links that a packet travels over from the source to the destination is called a *route* or *path*. Routes are discovered by running a distributed *routing protocol* on the

network. Figure 1-1 shows an example of a multi-hop wireless network. These networks are often called ‘mesh’ networks, in reference to the topology formed by the links and nodes. Typically a mesh network does not operate in isolation, and often has one or more gateways that connect it to a larger internet.

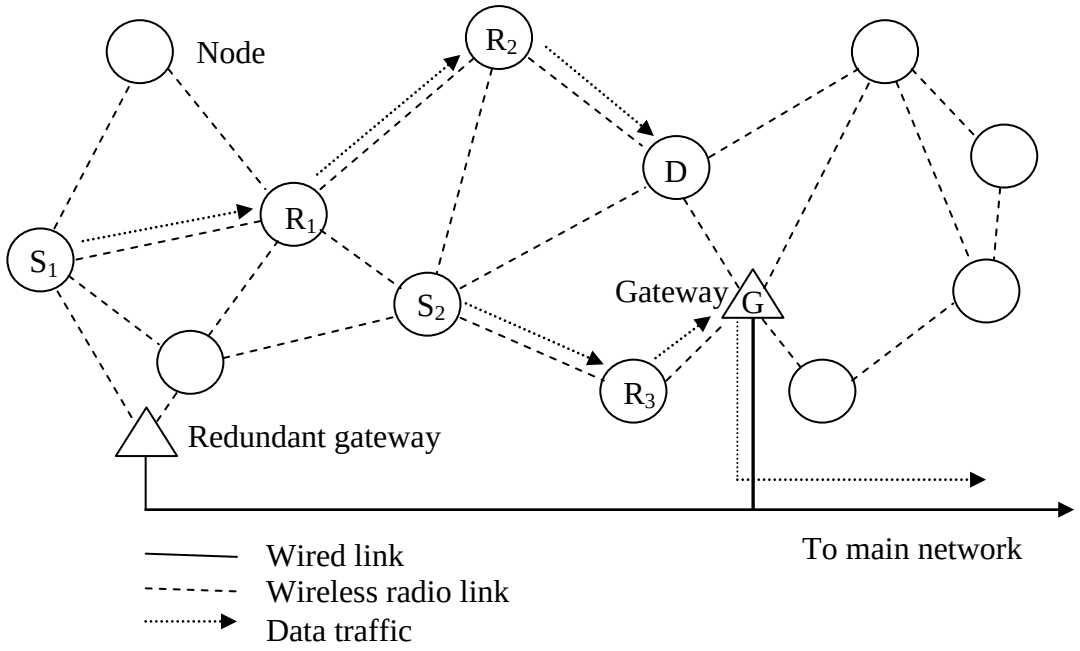


Figure 1-1: A multi-hop wireless network.

Node $S_1$ sends data to node $D$ via cooperating nodes $R_1$ and $R_2$, while node $S_2$ sends data out of the network via node $R_3$ and the gateway $G$.

### 1.2.1 Antennas

Wireless networks can be built using *omnidirectional* antennas, *directional* antennas, or some combination of the two. An omnidirectional antenna transmits and receives radio signals equally in all directions, forming links with other nodes in all directions. A directional antenna transmits and receives radio signals in a single direction, only forming links with nodes in that direction. Figure 1-2 illustrates the difference.

Links built using directional antennas can be approximated as wired links, and traditional wired network routing techniques will work well over these links [76]. Because each directional antenna is one end of a single point-to-point link, network designers can individually engineer each point-to-point link to have a very low loss ratio [69]. Also, each link can be considered

independently by the routing protocol, because the narrow coverage area of the directional antennas greatly reduces interference between links.

The disadvantage of point-to-point directional links is that they are difficult to install and engineer. Antennas must be aimed, link budgets must be calculated, and the network topology must be determined in advance, as each link requires its own antenna at each end. To add a new node to the network the place for adding new links must be explicitly decided. The network designer must add multiple point-to-point links for each new node in order to control the redundancy and fault tolerance.

On the other hand, a node with a single omnidirectional antenna cans form multiple links with many other nodes in any direction. The network designer can easily add a new node by placing it within range of any other node. Because the antenna is omnidirectional, it does not need to be aimed, and forms multiple links simultaneously with any nearby neighbors, providing redundant links with little extra effort.

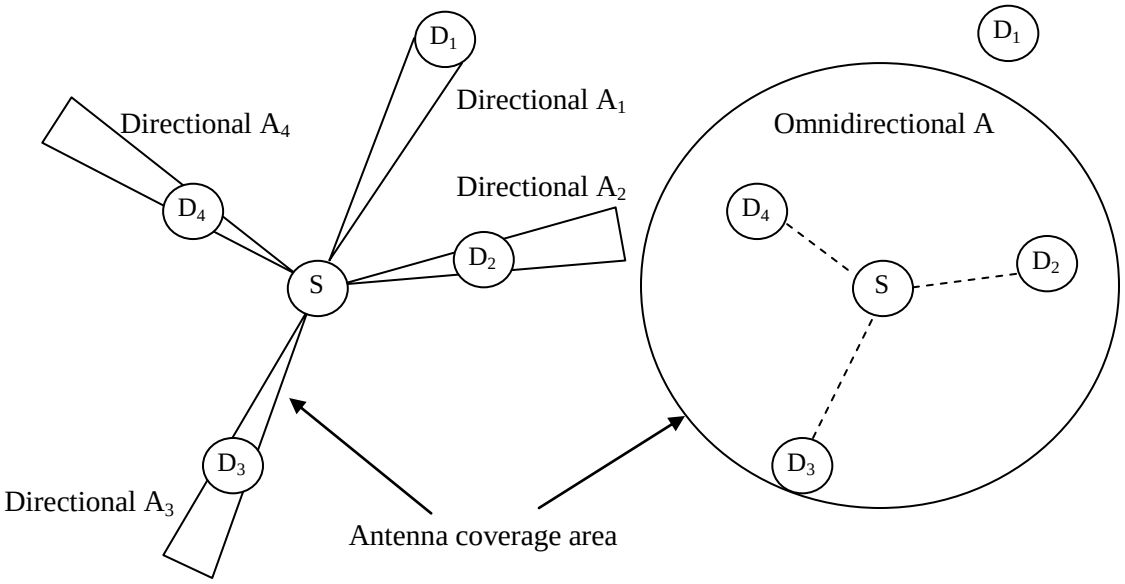


Figure 1-2: Building wireless networks with directional versus omnidirectional antennas.

The left side shows links using directional antennas, while the right sideshows links using an omnidirectional antenna. In both cases, the network operator would like to build links between node *S* and each of its neighbors *D*1 through*D*4. On the left, each directional antenna has a very narrow, long-range coverage area, and Node *S* has a good link to its neighbors *D*2 through *D*4. *S* has a marginal link to *D*1, since it as the edge of the antenna range. Each link

requires a separate antenna, *A*1 through *A*4. On the right, node *S* uses a single omnidirectional antenna *A*, with a very broad but relatively short-range coverage area. Node *S* has good links to nodes *D*2 and *D*4, a marginal link to *D*3, and no link to *D*1, which is out of range.

Although omnidirectional antennas make it easy to deploy new nodes, they have their own drawbacks. Each antenna is the end-point of multiple links, it is not feasible to independently engineer most of the links in the network, and many links will be lossy. Furthermore, the overlapping antenna coverage patterns of nearby nodes will cause them to interfere with each other, reducing the throughput of each link.

The rest of this work is about how to find high-throughput routes in multi-hop wireless networks built with omnidirectional antennas. Antennas are a significant part of the cost of a multi-hop wireless network, and unlike digital radios, their cost and functionality do not scale according to Moore's law. However, digital radios follow the steeply increasing performance curve of computer processing power, and will continue to become cheaper, with increasingly sophisticated signal processing, coding, and routing capabilities. As a result, systems built with a single omnidirectional antenna at each node will likely remain much cheaper than those built with multiple directional antennas at each node, even as their performance gap narrows.

### 1.2.2 Why Not Cellular?

A multi-hop wireless network can be incrementally expanded by adding nodes to the network, typically at the edges as its physical area grows. In this sense it is self-expanding: since the network nodes using the network cooperate to provide connectivity to each other, the network exists wherever there are nodes. This is in contrast to a cellular network, where data travels directly from wireless nodes to fixed base stations. Data typically travels from a base station to its destination over a wired network, as shown in Figure 1-3. Since each base station provides a fixed amount of network coverage to a fixed geographical area (the 'cell'), there is only network connectivity where base stations have been redeployed. Cellular base station locations and radio configurations are carefully chosen not to interfere with adjacent cells while avoiding coverage gaps between cells. Although cellular networks can be incrementally deployed and expanded, the overhead and planning required to setup a base station is much larger than to deploy a few extra new nodes in a multi-hop wireless network. In addition, unlike multi-hop wireless networks,

installing a cellular base station requires some preexisting or additional network infrastructure, such as land lines or long distance radio links to obtain network connectivity for the base station.

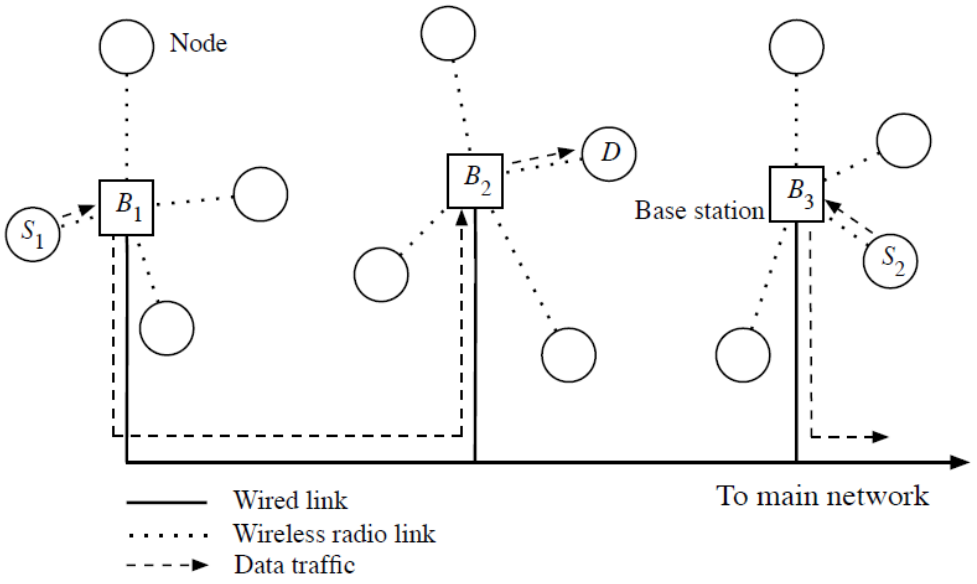


Figure 1-3: A cellular wireless network.

Node $S_1$ sends data to node $D$ via base stations $B_1$ and $B_2$, which communicate over the wired network. Nodes $S_1$ and $D$ might also share the same base station. Node $S_2$ sends data out of the network via base station $B_3$.

### 1.2.3 The Problem

The same flexibility that makes it easy to deploy multi-hop wireless networks with omnidirectional antennas also makes it difficult to find good links and routes. Unlike wired networks or wireless networks with point-to-point wireless links, it is difficult to engineer the communications links. When a new node is deployed, it will form communication links with all the nodes within the range, including those that are on the edge of communication range. As discussed in Chapter 4, links at the edge of communications range will have very poor signal strength, and packets sent over these links will often be lost completely or will be corrupted and discarded at the receiver end. Since lost or corrupted packets do not transfer any useful data over the link, the effective bandwidth of a lossy link is less than that of a good link. The percentage of transmitted packets that are lost or discarded is termed as the *loss ratio*; its complement, *delivery*

*ratio*, is the percentage of transmitted packets that are successfully received. Some of the links formed by adding a new node to the network will have low delivery ratios and low throughput, some will have high delivery ratios and high throughput, and many will have intermediate delivery ratios and throughput.

In general, there will be many potential routes between each pair of nodes in the network; because each route uses a different set of links, these routes will have different throughputs. The routing protocol selects the route with the highest throughput. Routing protocols use a *route metric* to decide which route to use between a pair of nodes. A route metric is a number assigned to each route; the routing protocol then selects the route with the best metric. The route metric is based on some underlying property of the route. For example, the commonly used hop-count metric is the number of links in a route. Protocols choose a route with the minimum hop-count; there may be many minimum hop-count routes, in which case protocols often choose arbitrarily between them.

## 1.3 Design Constraints

Multi-hop wireless networks have a very rich design space, and designers must make choices in many dimensions when building these networks. We make several assumptions about the underlying network that constrain the design space.

As discussed above, we assume that the network uses omnidirectional antennas, as they are cheaper and more convenient.

We implicitly assume that the network is a store-and-forward network which decodes and retransmits packets at each hop, according to predetermined routes that are decided by a routing protocol. This is one traditional way of operating data networks, and fits in well with current practice. However, it precludes techniques like network coding [5, 43, 48], which make more efficient use of the underlying network capacity.

We also assume that all network nodes have a single radio and antenna, operate on the same shared channel, and use the same fixed bit-rate and transmission power. But in reality, many radios can reduce link bit-rates for increasing reliability and variable transmission power can be used to trade off transmission range for total network capacity. Some radios can switch between multiple channels, or transmit on multiple frequencies simultaneously, as in orthogonal

frequency division multiplexing (OFDM) [16, 19, and 21]. Finally, placing multiple radios and antennas into each node can reduce or eliminate interference between links in the same route.

## 1.4 Contributions of this Work

The work emphasizes on the design, implementation, and evaluation of the potential transmission count (PTC) metric, through which high-throughput routes can be searched by routing protocols. The PTC consists of the total number of transmissions and retransmissions of packets for the purpose of sending a packet across the route, assuming that each link in the route retransmits the packet until it is successfully received across the link. PTC is designed for links with link-layer acknowledgments (ACKs) and retransmissions, as provided by IEEE 802.11 radios [18]. For a route the PTC metric can be calculated by measuring the lossless of each link in the route. Routing protocols select routes with the minimum PTC. For short routes (up to and including 3-hop routes), the minimum- PTC route is the maximum-throughput route; for longer routes, the minimum- PTC route is still a high-throughput route. The design of the PTC metric does not depend on a particular routing protocol; PTC improvement shows the throughput of both Dynamic Source Routing (DSR) [37], an on-demand source routing protocol, and Destination-Sequenced Distance-Vector (DSDV) [61] routing, a proactive table-driven distance-vector routing protocol. We also present a set of design changes and implementation techniques that allow DSR and DSDV to work well with PTC.

Additional contributions are a detailed exploration of the performance of minimum hop-count routing on a wireless test using 802.11b radios and a simple model of how link loss ratios vary with packet size. Chapter 3 explains why minimum hop-count often finds routes with significantly less throughput than the best available throughput, and quantifies the throughput difference between the typical minimum hop-count route and the highest throughput route. Chapter 4 shows how to use a few link loss ratio measurements to predict loss ratios at different packet sizes; these predictions can be used to decrease the protocol overhead of the PTC metric, by allowing PTC to measure links with small packets. PTC is also likely improving network capacity.

In order to demonstrate that PTC is effective, Chapter 7 presents measurements taken from the test network. These measurements show that PTC improves the throughput of multi-hop routes by a factor of two over the minimum hop-count metric. PTC provides the highest paths

with two or more hops, suggesting that PTC offers increased benefit as networks grow larger and paths become longer.

## 1.5 How to Read This Dissertation

Chapter 2 surveys various works in wireless ad-hoc routing. Chapter 3 reviews the 802.11 radios used in this work, and can be skipped by readers familiar with 802.11. Chapter 4 describes the test-bed network and its throughput problems. Chapter 5 explains how digital packet radios work, and gives a simple model of how packet size affects loss ratios; it can also be skipped by readers familiar with digital communications, although they might wish to read Section 5.4 to learn about the packet size model. Chapters 6 and 7 present the design and implementation of PTC and explain how it is used by the routing protocols. Chapter 8 evaluates how well PTC works on a real wireless network, while Chapter 9 shows how PTC might be improved in future, and how it is useful when the wireless network design space is expanded. Finally, Chapter 10 concludes the strategies and findings of the research.

# Chapter 2: Literature Survey

Much of the recent work in ad hoc routing protocols for wireless networks [61, 37 and 62] has focused on coping with mobile nodes, rapidly changing topologies, and scalability. In general these works have been carried out using analysis and simulation that focuses on how mobility affects link connectivity and routing protocol behavior. To this end, researchers have used link connectivity models based on radio ranges: in-range links work well, while out-of-range links are broken. This model works well when the dominating factor in determining link connectivity is node location and motion. However, this on-off link model is not useful for fixed networks, where the individual link performances are no longer decreased by the effects of motion. In particular, protocols that seem promising for simulated mobile networks often don't provide the best performance for fixed networks. Many of the ways that simulation and analysis could be altered to model wireless links more accurately are detailed by Kotz et al. [46]. Recent protocol design and evaluation work has started to focus on the detailed behavior of real wireless links.

The behavior of routing protocols over lossy links has been addressed and evaluated by real implementations in several recent papers. Lundgren et al. [50] relate their experiences with an 802.11-based multi-hop network with four nodes, and they coin the term 'gray zones' to refer to links that deliver routing protocol packets but not data packets. They propose using link handshaking and counting route broadcasts to filter out gray zone links. Link handshaking requires both ends of a link to acknowledge the other end before using the link. Chin et al. [15] also describe a four-node multi-hop network based on 802.11 radios, and independently propose link handshaking to filter out asymmetric links.

The CMU/Rice Monarch project has had several years of experience with their DSR implementation [25], which has provided important lessons about wireless network emulation [41], and the performance of real implementations [51]. Hu and Johnson describe an eight-node mobile test using DSR that reliably transmits video and audio streams [34]. Because mobility causes new links to form and older links to break, they modify DSR to preemptively issue DSR route requests when a link's signal-to-noise ratio (SNR) drops below a given threshold. The assumption is that links with low SNRs are likely to break soon because of motion. Preemptively issuing a route request allows the source node to have a fresh route ready to use before the link

actually does break. However, even with this modification, DSR does not discriminate between links that are deemed functional: it treats all working links as equivalent, and finds minimum hop-count routes.

The ideas presented in this work are motivated by experiments on a relatively large-scale test, with 5 or 6 times more nodes and an order of magnitude more links than the tests used in most previous work. Chapter 3 showed that these links have many different qualities, evenly spread from best to worst, and that there is no easy division of links into 'good' and 'bad' categories. However, experiments on smaller tests are not as likely to reveal such a wide distribution of link quality, and previous work has been more focused on categorizing links as 'good' or 'bad'. In contrast, the mechanisms and protocols introduced in this work are specifically designed to accommodate and take advantage of links of all qualities.

Some of the earliest and most important work in multi-hop wireless networking was the DARPA Packet Radio Network (PRNet) research [39], which was carried out in the 1970s and 1980s.The PRNet was specially designed for multi-hop wireless networking at all layers of the system, including a new spread-spectrum packet radio design [28]. The system was also designed to handle lossy links, and used link-level packet loss ratio measurements to quantify link quality. However, the PRNet used loss ratio thresholds to distinguish good links from bad, which has a few drawbacks, as discussed below. PRNets were deployed and experimented with by many research groups, at a scale similar to the test network used in this work [29].

One unanswered question is whether or not the PRNet deployments had similar link loss rate distributions to those shown in Chapter 3. It is conceivable that the different radio design and deployment scenarios produced a different distribution of link loss ratios, and that the threshold technique used by the PRNet was indeed the most suitable technique.

The PRNet research also produced a new routing metric, called Least Interference Routing (LIR) [74]. This was shown in simulation to increase the total network throughput of random packet radio networks. Routing protocols using LIR select routes that interfere with the fewest number of other nodes, by choosing routes to minimize the sum of the number of one-hop neighbors of each node in the route. LIR differs from the PTC metric presented here in that it does not directly evaluate actual link performance, in terms of packet loss ratios; instead, it uses the interference metric to predict performance.

There is also much related work in the field of sensor networks, even though on the face of it sensor networks are very different from the multi-hop wireless networks we are concerned with. In general, the nodes in sensor networks have orders of magnitude fewer resources than a typical PC, in terms of processing power, memory, and available energy. In addition, because of size and power constraints, sensor radios are much less sophisticated than 802.11 radios, and the application requirements are much less: sensor networks typically carry very low-bandwidth data streams. However, recent work has shown that despite these differences, sensor networks must deal with the same sort of link lossiness and variation that multi-hop wireless data networks face. Indeed, the PTC metric presented here is useful in both sorts of networks.

Zhao and Govindan [80] provide a detailed analysis of the performance of wireless links in three different network deployment scenarios, using BerkeleyMotes, a low-power sensor network platform. They find that many links can be lossy to varying degrees, and, as in Chapter 4 and 5, find that signal strength is not an adequate predictor of packet loss ratios, because of effects such as multipath interference.

Woo et al. [78] also provide measurements showing the variability of links in aMote-based network. They use the measurements to drive an analysis of link estimator sand routing protocol techniques. They propose the *Minimum Transmission* (MT) metric, and show that it greatly improves the end-to-end packet delivery ratio in their experiments compared to minimum hop-count. MT is the same metric as PTC, except that the link delivery ratios are estimated by passively snooping on data packets over a fixed time window, rather than sending special probes. Furthermore, the link measurements are smoothed with an exponentially-weighted moving average, and a hysteresis is applied when deciding whether to switch routes based on a metric change.

Yarvis et al. [79] also observe that hop-count performs poorly as a routing metric for a Mote-based sensor network. They present a path metric which approximates the product of the per-link delivery ratios. As argued in Chapter 5, this metric is likely to use low-loss paths with many hops in situations where a path with a smaller number of higher loss links would perform better.

A number of existing ad hoc wireless routing algorithms collect per-link signal strength information and apply a threshold to avoid links with high loss ratios [15, 22, 26, 30, 34, 39, 50 and 67]. One problem with these approaches is that given the complexities of radio links such as

multipath interference and fading, it is unlikely that SNR measurements can be used to accurately predict packet loss ratios; this is shown experimentally in Chapter 5 for 802.11 radios and by Zhao and Govindan for Motes [80]. It is true that SNR measurements provide some useful information about the underlying quality of a link: links with extremely poor SNR values are likely to have high packet loss ratios. But, it is not clear what the appropriate SNR threshold should be. Setting the SNR threshold too high may eliminate links that are necessary for connectivity, while setting the threshold too low will allow the routing protocol to use poor links with low throughput. In general, any approach that uses thresholds will suffer a similar problem. For example, Woo et al. [78] show that a routing protocol which simply ignores links with delivery ratios below a given threshold is either unable to keep the network connected, or unable to find routes with good end-to-end delivery ratios, depending on the threshold used. PTC avoids these problems by allowing any link to be directly compared to another link, and by assigning metrics that allow all routes to be compared and ranked. This avoids leaving out any links that are required to connect the network, but still enables the routing protocol to choose better links and routes over worse.

A complementary solution to high link loss ratios is to improve the effective loss ratio with some form of redundancy. Forward error correction, MAC level acknowledgment and retransmission, and solutions such as Snoop-TCP [7] and Tulip [60] all take this approach. However, even with these techniques it is preferable to use links with low loss ratios rather than links with high loss-ratios: retransmissions (or other redundancy) reduce useful link capacity and increase interference.

Like most of the work discussed above, this work treats wireless networks from the bottom up, trying to find good techniques and abstractions for characterizing and distinguishing between links and routes in wireless networks. However, there is a large body of work that is concerned with Quality of Service (QoS) in networks, especially wireless multimedia networks. Wireless QoS algorithms approach route selection from the top down. Some techniques explicitly schedule transmission slots in time or frequency division-MAC layers to provide bandwidth guarantees [14, 33, 49, 53 and 81], while others treat the MAC as opaque, and rely upon it for bandwidth and delay information and constraints [13, 70 and 72]. Unfortunately, this link information is hard to determine in the sort of distributed multi-hop networks we are concerned with, because of unknown data traffic patterns or network topology, and inter-node

interference. Indeed, although the PTC metric enables routing protocols to find high-throughput routes, it doesn't provide enough information about links to QoS algorithms to allow them to make bandwidth or latency guarantees, especially when multiple nodes are sending data.

We assume that the loss ratio of a given link cannot be controlled by the system. More sophisticated hardware might allow transmit power levels to be changed to make links better behaved. Existing systems exploit this idea, often with a focus on minimizing the energy consumption required to successfully deliver data [32, 40 and 68]. Energy consumption is primarily a concern for sensor networks, where radio transmissions consume the majority of each node's energy budget. However, fixed networks can benefit by using power control to reduce transmission ranges and increase network capacity [31, 71].

Since the PTC metric assumes that all links run at the same bit-rate, it does not properly find high-throughput routes when links run at multiple bit-rates. A werbuchet al. [6] presents the *Medium Time Metric* to help find high-throughput paths when links can run at different bit-rates. Since their metric does not account for losses, it is complementary to PTC.

# Chapter 3: Overview of 802.11

## 3.1 Introduction

This chapter provides an overview of the IEEE 802.11b used in this work.802.11 is an IEEE standard for the physical and medium access control (MAC) layers of wireless LANs [18]. The standard specifies several layers of a packet radio system, including radio modulation and coding, packet formats, and the MAC protocol for managing contention between multiple senders. The original 802.11standard specifies radios that can operate at one and two megabits per second; the follow-on 802.11a [19], 802.11b [20] and 802.11g [21] standards specify additional bit-rates and packet formats.

The main focus of the 802.11 standard is on the networks with a star topology, and almost all 802.11 radios are used this way. In these networks wireless clients exchange packets with specially-designated wireless access points. The access points then relay client packets between the wireless clients and a wired LAN. In this scenario, wireless clients do not exchange packets directly with each other; all packets pass through an access point. This mode of operation is often referred to as *infrastructure* mode.

However, the protocols and experiments described in this work do *not* use the radios in infrastructure mode. Instead, the radios are used in a *peer-to-peer* mode where they can directly send and receive packets from any radio which might be in range. This mode is also sometimes called *ad hoc* mode. The 802.11 standard refers to radios operating in this mode as an Independent Basic Service Set (IBSS).

+1, −1, +1, +1, −1, +1, +1, +1, −1, −1, −1

Figure 3-1: Barker spreading sequence used by 802.11.

Table3.1: IEEE 802.11b bit-rates and their associated modulation.

| Bit-Rate | Modulation | Bits/Symbol | Chips/Symbol |
|---|---|---|---|
| 1 Mbps | DBPSK | 1 | 11 |
| 2 Mbps | QPSK | 2 | 11 |
| 5.5 Mbps | CCK | 4 | 8 |
| 11 Mbps | CCK | 16 | 8 |

## 3.2 Physical Layer

The IEEE 802.11 standard describes three physical layers: an infrared layer, afrequency-hopping spread-spectrum layer, and a direct-sequence spread-spectrum (DSSS) layer. Almost all 802.11 radios use the DSSS physical layer, as do the802.11b radios we used.

The DSSS physical layer specifies how bits and packets are transmitted overthe radio air interface. IEEE 802.11b specifies four bit-rates, with associated modulationtechniques, as summarized in Table 3.1.

After modulation, the data symbols are encoded by an 11-chip Barker spreadingsequence, at 11 megaships per second. The Barker sequence used is shown in Figure 3-1. Table 3.1 shows how many chips encode each symbol for each bit-rate.

In the United States, 802.11 and 802.11b specify 11 channel center frequencies,starting at 2,412 MHz, and spaced 5 MHz apart. Since after spreading with the Barker code, the main lobe of the transmitted signal has a frequency width of22 MHz, these channels actually overlap significantly with each other. However, it is possible to choose three channels without significant overlap.

Each 802.11 packet transmission consists of a 148-bit preamble and a 48-bitphysical layer header, followed by the 802.11 payload. The preamble and physical layer header bits are sent at 1 Mbps, and the 802.11 payload bits can be sent at any of the 802.11 bit-rates. The contents of the preamble are specified by the802.11 standard. The physical layer header specifies the total length of the packet and the bit-rate used for the 802.11 payload. The 802.11b standard specifies additional optimizations that decrease the time required for the preamble and physical layer header when higher bit-rates are used for the payload; this reduces packet overhead at 802.11b's 5.5 and 11 Mbps data rates.

## 3.3 MAC Layer

The 802.11 MAC protocol is a carrier-sense multiple-access scheme with collision avoidance (CSMA/CA). The standard refers to this scheme as the distributed coordination function (DCF). The goal of the MAC protocol is to allow multiple competing senders to share the radio medium without interfering with each other.

### 3.3.1 Medium Access

Before sending a packet, a potential sender listens to see if any other transmission is in progress. If there is no such transmission, or once such a transmission is over, the sender waits for a mandatory of time called the DCF inter frame space (DIFS). After the DIFS time has passed, the sender chooses a random back-off time *b* from its *contention window*. The contention window has a minimum length of 620 microseconds, and a maximum length of 2,460 microseconds. After each successful transmission, the contention window is set to its minimum value; after each failed transmission, the contention window is doubled, up to the maximum value. The sender waits for the back-off time *b* to pass before attempting to send its packet. If some other radio transmits while the sender is waiting for *b* to elapse, the sender does not count that time as waiting, and resumes waiting at the end of the transmission. That is, the sender waits for *b* amount of *idle* medium time before attempting to send.

The 802.11 standard also specifies an optional request-to-send/clear-to-send(RTS/CTS) protocol which can further reduce radio contention in some scenarios. As RTS/CTS is not used in this work.

### 3.3.2 Retransmissions and Packet Timing

The 802.11 MAC supports two kinds of data packets: broadcast and unicast. Broadcast packets are intended to be received by any radio which hears them, and are delivered to the networking layer on that radio's node. Unicast packets are directed to a specific destination node. When a radio receives a unicast data packet directed to it, it immediately sends back an acknowledgment (ACK) packet after a short interframe space (SIFS), and delivers the incoming data packet to the networking layer. Other radios may receive the same unicast packet, but they discard it and do not send an ACK response. Each packet includes a destination address so that a radio can decide if the packet was intended for it. Figure 3-2shows the formats of data and ACK packets.

**(a) 802.11 data frame, 59 +*n* bytes over the air.**

| **Preamble (18 bytes)** | **Physical layer header and CRC (6 bytes)** | **802.11 and Ethernet headers (31 bytes)** | **Ethernet Payload (*n* bytes)** | **Data CRC (4 bytes)** |
|---|---|---|---|---|

**(b) 802.11 ACK frame, 38 bytes over the air.**

| Preamble (18 bytes) | Physical layer header and CRC (6 bytes) | ACK frame (10 bytes) | Data CRC (4 bytes) |
|---|---|---|---|

Figure 3-2: Packet formats for 802.11 data and acknowledgment packets.

**(a) Packet timing for 802.11 broadcasts.**

| 802.11 Data (*n* data bytes) (8× [*n* + 59] µs) | ← DIFS → (50 µs) | ← Backoff → (≈310 µs ) | 802.11 Data |
|---|---|---|---|

**(b) Packet timing for 802.11 unicasts.**

| 802.11 Data (*n* data bytes) (8× [*n* + 59] µs) | ← SIFS → (10 µs) | 802.11 ACK (304 µs) | ←DIFS→ (50 µs) | ←Back off→ (≈310 µs ) | 802.11 Data |
|---|---|---|---|---|---|

Figure 3-3: Packet timing diagram for 802.11 data traffic, assuming no contention for the radio channel.

The total time required to send an 802.11 data broadcast at 1 Mbps with an *n*-byte data payload is 8 × [*n* + 59] + 50 + 310 = 832 + 8*n*microseconds. The total time for a unicast is increased by 10+304 microseconds because of the ACK packet, and is 1,146 + 8*n* microseconds.

If a unicast sender does not receive an ACK packet after a specified period of time (SIFS + DIFS time after sending the data packet), it marks the transmission as failed. The sender then increases its back-off window, enters back-off, and tries to resend the packet. A sender will repeatedly try to retransmit a packet up to a specified maximum number of tries1 giving up and discarding the packet.

Figure 3-3 shows the packet exchanges and timings for broadcast and unicast packets at 1 Mbps, assuming that every packet transmission is successful and that there is no contention. The figure shows the average expected back-off time of310 microseconds. In the absence of contention, the back-off window should beat its minimum size of 620 microseconds, and the average expected random back off is one-half of that. The maximum broadcast and unicast

throughputs of a given packet size can be calculated in packets per second by inverting the time required to send a single packet. For example, for the 134-byte payload used through out this work, the unicast throughput $B$ can be calculated as

$$\frac{\qquad\qquad}{}$$

packets per second. For unicast packets with a 1,386-byte payload, the throughput is 82 packets per second.

# Chapter 4: The Throughput Problem

## 4.1 Introduction

The minimum hop-count route metric is used by most of the existing wireless routing protocols. They select routes with the fewest links. The minimum hop-count metric implicitly preset that links either work well, or do not work at all and that all working links are equivalent. Furthermore, most protocols assume links that deliver routing control packets such as DSDV route updates or DSR route queries will also successfully deliver data packets.

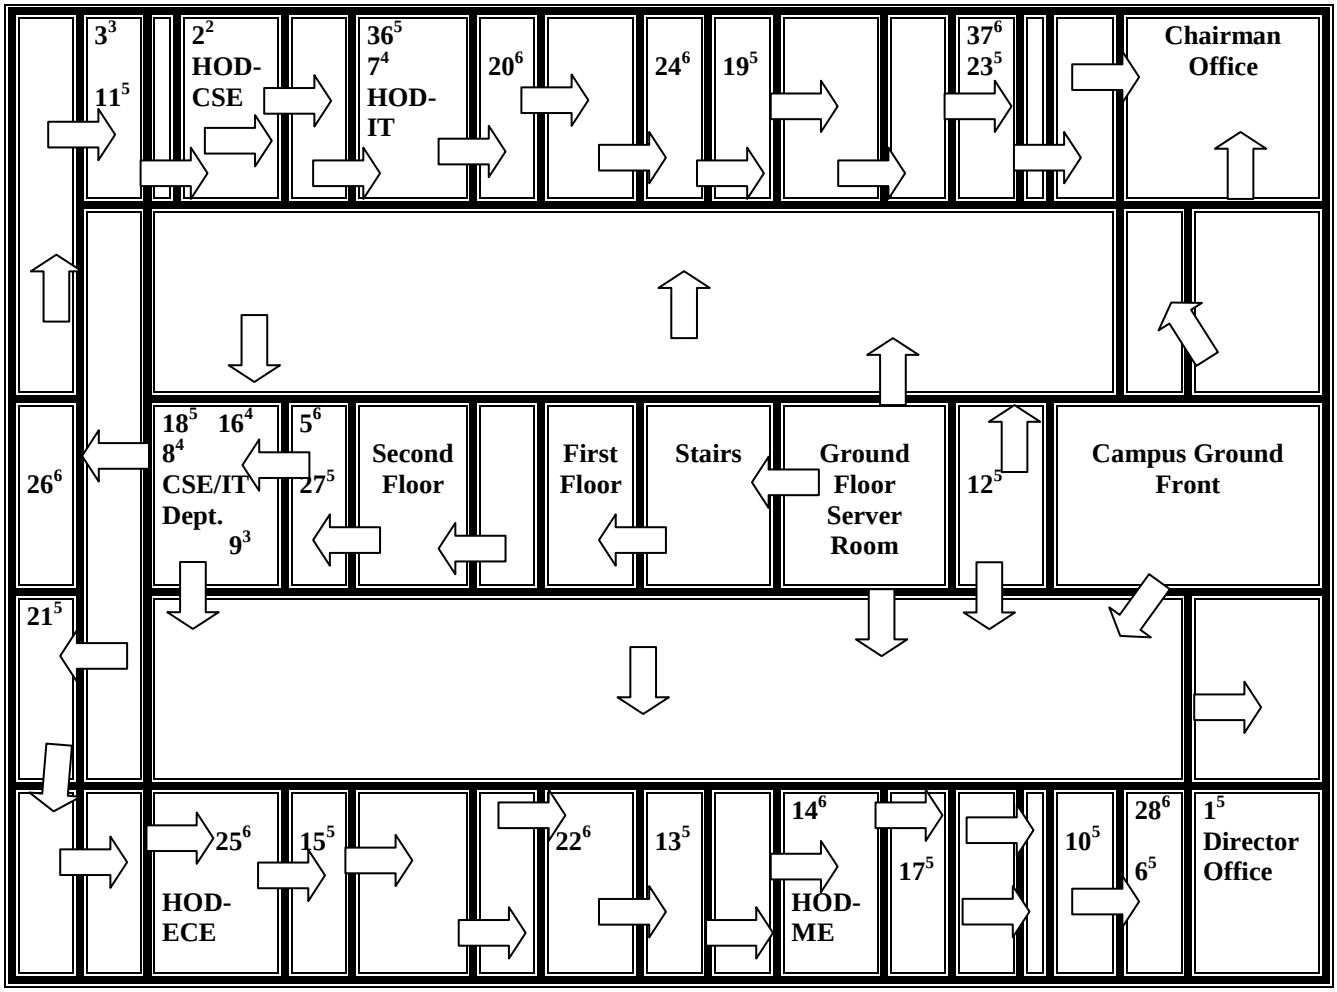


Figure 4-1: A map of the test.

Each circle is a node; the large number is the node identifier, and the superscript indicates which floor of the building the node is on.

However, these assumptions are incorrect for multi-hop wireless networks with omnidirectional antennas. Unlike wired and wireless networks with point-to-point links, where the performance of each link can be tightly controlled and engineered, networks with omnidirectional antennas have many wireless links with a wide range of intermediate loss ratios. These lossy links are not useful for data, but deliver enough routing control packets so that the routing protocol uses the link. Measurements in Section 4.4 illustrate the even distribution of link loss ratios for an indoor 802.11 test; others have measured similarly even distributions for an outdoor 802.11 network [4], and for indoor and outdoor sensor networks [12, 78 and 80].

Given a broad variation in link loss ratios, hop-count will choose links poorly.This is because minimizing the hop-count of a route maximizes the distance traveled by each hop, which reduces the received signal strength and increases the loss ratio. Even if the best route is a minimum hop-count route, there may be many routes with the same minimum hop-count, but with widely varying qualities. The arbitrary choice made by minimum hop-count is not guaranteed to be the highest-throughput route. This chapter shows that minimum hop-count routing typically finds routes with significantly lower throughput than the best available, using measurements of the DSDV routing protocol on a test network. We explain why minimum hop-count does poorly by looking at the distribution of route throughputs and link loss ratios.

## 4.2 Experimental Test

All the data in this paper are the result of measurements taken on a 29-node wireless test. Each node consists of a stationary PC with an Intel processor 2.4 GHz (Core 2 Duo) PCI/PCI-X 802.11b card, Server with a Geon Processor (Core 2 Duo-4 Processor) Intel Pro/1000 MT PCI/PCI-X 802.11b card and an omnidirectional 2.2 dBi dipole antenna, also called a ‘rubber duck’ antenna. Each PC runs the Linux operating system. The nodes are placed in offices and cabins on four consecutive floors of an office building. Their positions are shown in Figure 4-1.

The 802.11b cards are set to transmit at one megabit per second (Mbps) with one millwatt (mW) of transmit power. RTS/CTS is turned off, and the cards are set to ‘ad hoc’ (IBSS, DCF) mode. Each data packet in the following measurements consists of 24 bytes of 802.11b preamble, 31 bytes of 802.11b and Ethernet encapsulation header, 134 bytes of data payload, and

4 bytes of frame check sequence: 193 bytes in total. An 802.11b ACK packet takes 304 microseconds to transmit, the inter-frame gap is 60 microseconds, and the minimum expected mandatory back-off time is 310 microseconds, resulting in a total time of 2,218 microseconds per data packet. This gives a maximum throughput of 451 unicast packets per second over a loss-free link.

While the test itself carried only the data and control traffic involved in each experiment, interference of various kinds was inevitably present. In particular, each floor of the building has four 802.11b access points, on various channels.

The test was designed to experiment with wireless routing protocol implementations, and is one of the larger 802.11-based multi-hop wireless tests currently described in the research literature [2, 3, 76, 34, 15, 50 and 51]. There are also many commercial multi-hop wireless networks, which are not publicly documented; some of these are smaller than the test described here, but many are much larger, both in the number of nodes and in the area covered by the network. We believe that although radio link performance is known to be quite different indoors than outdoors [69], the conclusions drawn from measurements of this test are still valid for many other networks. This is for two reasons: first, initial measurements of a larger outdoor rooftop network corroborate many of the findings described in this chapter; and second, the main effects we observe stem from the underlying network design, not the specific performance of any particular link.

## 4.3 Path Throughputs

Figure 4-2 compares the throughput of routes found with a minimum hop-count metric to the throughput of the best static routes that could be found. Each curve shows the throughput cumulative distribution function (CDF) for 100 node pairs; the pairs are randomly selected from the 29 × 28 = 812 total ordered pairs in the test. A point's *x* value indicates the throughput between the pair, in packets per second; the *y* value indicates the fraction of pairs with fewer throughputs. The left curve is the throughput CDF achieved by routing data using DSDV with the minimum hop-count metric. The right curve is the throughput CDF for the best known path between each pair of nodes. Packets were only sent between one pair at a time. That is, there was no data cross traffic. For each pair, the DSDV and best-path tests were run immediately after one another, to limit variation in link conditions over time.

The ‘best’ static route between each pair of nodes was found by sending dataalong ten potential best paths, one at a time, and selecting the path with the highest throughput for each pair. Potential best paths were identified by running an off-line routing algorithm, whose inputs were measurements of per-link loss ratios similar to those in Section 4.4.

The algorithm also incorporated a penalty to reflect the education in throughput caused by interference between successive hops of multihop paths. New link measurements were collected roughly every hour during the experiment; the best paths for each pair were generated using the most recently available loss data. Each node ran a user-level program which forwarded packets according to source routes in the packet headers.

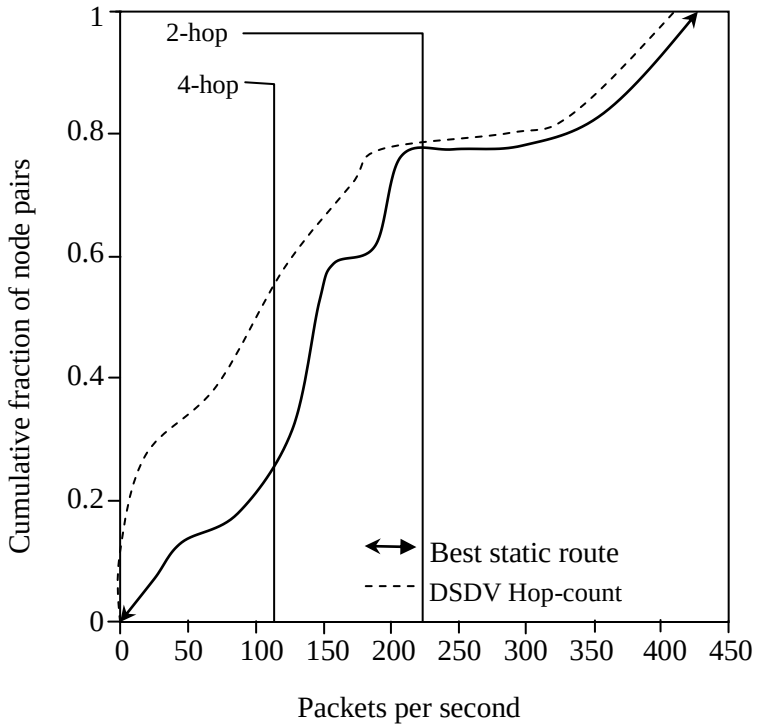


Figure 4-2: When using the minimum hop-count metric, DSDV chooses paths with far less throughput than the best available routes.

Each line is a throughput CDF for the same 100 randomly selected node pairs. The left curve is the throughput CDF of DSDV with minimum hop-count. The right curve is the CDF of the best throughput between each pair, found by trying a number of promising paths. The average throughput difference is 42 packets per second ($\sigma = 61$). The vertical lines mark the theoretical maximum throughput for routes of each hop-count.

The throughputs in Figure 4-2 are split into two main ranges, above and below 225 packets per second. Pairs with throughputs above 225 sent data along single-hop paths; pairs with throughputs at or below 225 sent data over multi-hop paths or very poor single-hop paths. Multi-hop paths have fewer throughputs because transmissions on the successive hops interfere with each other. In a two-hop path, the middle node cannot receive a packet from the first node at the same time it is sending a packet to the last node, limiting throughput to one-half the link throughput. Similar effects cause the fastest three-hop route to have a capacity of about 450/3 = 150 packets per second [47].

Minimum hop-count performs well whenever the shortest route is also the fastest route, especially when there is a one-hop link with a low loss ratio. A one hop link with a loss ratio of less than 50% will outperform any other route. This is the case for all the points in the right half of Figure 4-2. Note that the overhead of DSDV route advertisements reduces the maximum link capacity by about 15 to 25 packets per second, which is clearly visible in this part of the graph.

The left half of the graph shows what happens when minimum hop-count has a choice among a number of multi-hop routes. In these cases, the hop-count metric usually picks a route significantly slower than the best known. The most extreme cases are the points at the far left, in which minimum hop-count is getting a throughput close to zero, and the best known route has a throughput of about 100 packets per second. The minimum hop-count routes are slow because they include links with high loss ratios, which cause bandwidth to be consumed by retransmissions. The zero-throughput points on the left are due to asymmetry: DSDV with hop-count chose asymmetric links which delivered routing packets in the reverse direction, but no data packets in the forward direction.

## 4.4 Distribution of Path Throughputs

Figure 4.3 illustrates a typical case in which minimum hop-count routing wouldnot favor the highest-throughput route. The figure shows the throughputs of several static routes from node 23 to node 36. The routes are the eight highest-throughput routes between 23 and 36 which were found in the ‘best’ static route experiments described above.

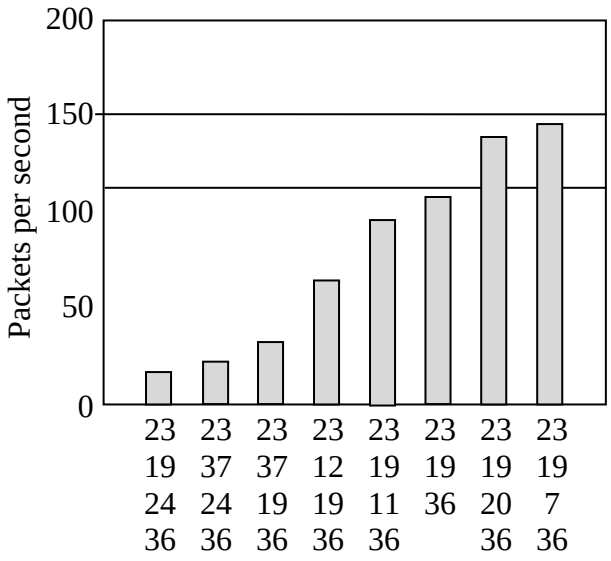


Figure 4-3: the measured throughputs from node 23 to node 36, along the eight highest-throughput routes found in the ‘best’ static route tests.

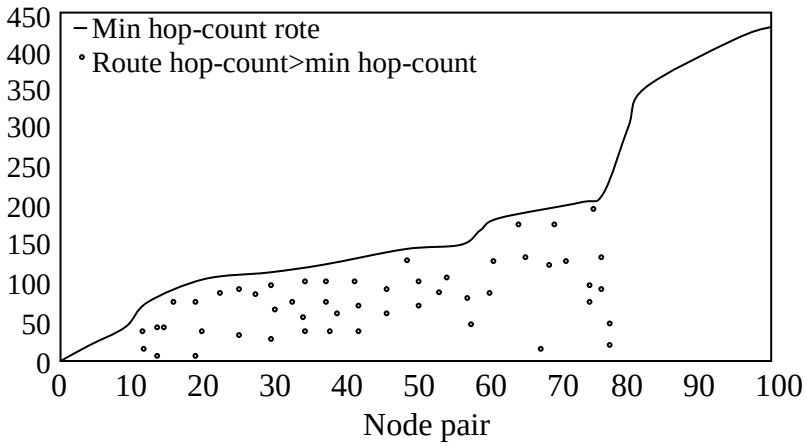


Figure 4-4: Measured throughput of all static routes.

Circles mark the throughput of minimum hop-count routes; longer routes have their throughput marked with triangles. 99 pairs are shown here; a minimum hop-count route had the highest throughput on 73 of those pairs. Multi-hop routes were not tested for pairs with a one-hop throughput of greater than 225 packets per second, as that is faster than any multi-hop route can deliver packets.

The graph shows that the shortest path, a two-hop route through node 19, does not yield the highest throughput. The best route is three hops long, but there are a number of available three-hop routes which provide widely varying performance. Figure 4-4 shows the ‘best’ static route results for all the node pairs tested. Although the fastest route many pairs was a minimum hop-

count route, 35 pairs have multiple minimum hop-count routes, typically with very different throughputs. Further more, the minimum hop-count route was not the fastest route for a quarter of the pairs. A routing protocol that selects randomly from the shortest hop-count routes is unlikely to make the best choice, particularly as the network grows and the number of possible paths between a given pair increases.

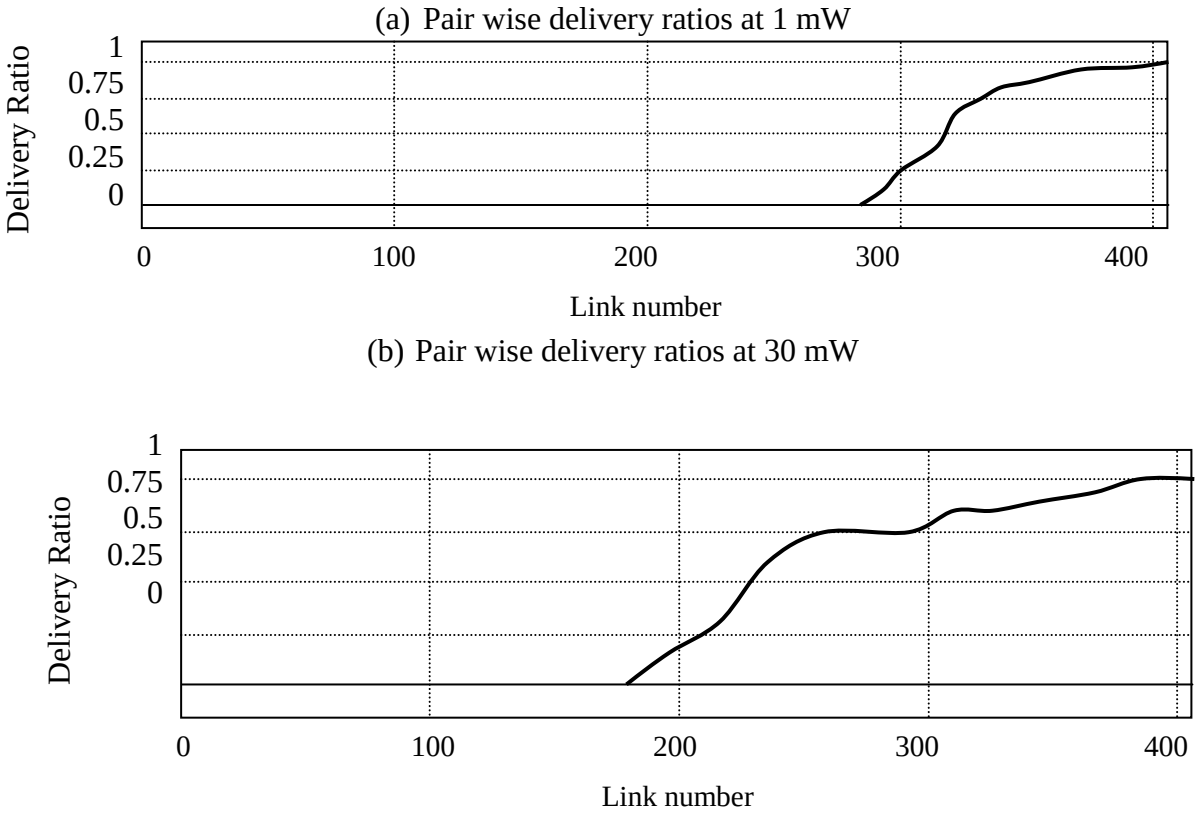


Figure 4-5: One-hop packet delivery ratios between each pair of nodes at 1 mW (above) and 30mW(below).

The top and bottom ends of each vertical line indicate the delivery ratios in the two directions. The bars in each graph are sorted by the minimum of the two directions, so the link numbers do not necessarily match between the two graphs. The packet size is 134 bytes of 802.11b data payload. Data for all 406 pairs of hosts are shown. Many links are asymmetric, and there isa wide range of loss ratios.

## 4.5 Distribution of Link Loss Ratios

Figure 4-5 shows the underlying delivery ratios of each link in the network, which helps explain why high-throughput paths are difficult to find. Each vertical bar corresponds to the direct radio link between a pair of nodes; the two ends of the bar mark the broadcast packet delivery ratio in the two directions between the nodes.

To measure delivery ratios, each node took a turn sending a series of broadcast packets for two seconds, and counted the number of packets that the radio reported as transmitted. Packets contained 134 bytes of 802.11b data payload, and were sent at a rate of 40 packets per second. Every other node recorded the number of packets received. The delivery ratio from node *X* to each node *Y* is calculated by 38 dividing the number of packets received at *Y* by the number sent by *X*. The loss ratio of a link is one minus its delivery ratio. We use the term 'ratio' instead of 'rate' to avoid confusion with throughput delivery rates, which are expressed in packets per second.

Note that 802.11 broadcasts don't involve acknowledgments or retransmissions. Because 802.11 retransmits lost unicast packets, a link's unicast packet loss ratio at higher layers is potentially far lower than the underlying broadcast loss ratio, depending on the maximum number of retransmissions allowed. Since only one node was broadcasting at a time in the network, any packet losses are due to interference from the environment or from transmitters outside the network.

Figure 4-5 has three important features. First, large fractions of the links have an intermediate delivery ratio in at least one direction. That is, they are likely to deliver some routing protocol packets, but would lose many packets if used for data. Second, there is a full spectrum of link delivery ratios, so some advantage can be expected from making fine-grained choices between links when choosing paths. Third, many links have asymmetric delivery ratios.

As discussed in Chapter 1, using omnidirectional antennas makes it easy to deploy a wireless network, but hard to engineer any particular link to have a very low loss ratio. As a result, many of the links in the network are operating in situations they were not designed for, and therefore have non-negligible loss ratios. These links are operating with low SNRs, high noise, and excessive multipath due to the wide variety of obstacles indoor, such as doors, walls, furniture, and people.

The network has a wide range of link loss ratios because the links are operating in a wide range of conditions, despite being in the same network. For example, there is a wide distribution of link distances, and therefore a wide range of received signal levels. This is further compounded by the different levels of receiver noise for each link and the various obstacles blocking and reflecting each link's signal; these produce a wide range of SNR levels throughout the network. Also, the different obstacles to each link produce different multipath effects, further affecting loss ratios in unpredictable ways. Multipath effects are discussed further in Chapter 5.

Even though the effects of attenuation and multipath should be symmetric for each link, many of the links are asymmetric. There are a few explanations for this asymmetry. As just described, receiver noise levels affect the SNR and therefore loss ratios; since each receiver is in a different environment it is likely to have a different noise level and SNR. Some receivers will be in high-noise environments, producing very asymmetric links. Second, although the radios used in the test are identical models, they may come from different manufacturing batches, with slightly different components. Even though the radios were set to use the same power, their actual power outputs may vary, producing differences in the received signal level at each end of the link. Finally, because the test radios are half-duplex the measurements in each direction of a link occur at different times. It is possible that link conditions changed between the measurements in each direction. For example, a door may have been closed, or a person may have moved their chair; these effects have been informally observed to affect link measurements in the test. Although these time variations may seem to make the link measurements less reliable, they are in fact an accurate reflection of the sorts of link behavior that a wireless routing protocol will encounter. In a real network, perceived asymmetry will occur as a result of link changes over time. The routing protocol chooses routes before sending data over them, by using protocol packets sent in the reverse direction; links may change between the time protocol packets are sent in the reverse direction and data packets are sent in the forward direction.

Of the 406 node pairs in Figure 4-5a (1mW), there are 124 with links which delivered packets in at least one direction. Of those links, 28 are asymmetric, with forward and reverse delivery ratios that differ by at least 25%. The 28 asymmetric links involve 22 different nodes, indicating that asymmetry is prevalent throughout the whole network, and not isolated to only a few nodes and links. Because 802.11b uses link-level acknowledgments (ACKs) to confirm delivery, both directions of a link must work well in order to avoid retransmissions. Since most

nodes in the network are involved in at least one asymmetric link, routing protocols must cope with asymmetry to be effective.

Figure 4-5b shows similar data, but with the radios set to the 30mW transmit power, which is about a 15 dB increase in transmit power. As a result, 229 links deliver packets, almost twice as many as in the 1mW experiment. Also, many more links have very high delivery ratios: at 1mW there are 69 links (17% of all links) that deliver at least 95% of their packets; at 30mW there are 121 such links (30% of all links). However, the fraction of working links with high delivery ratios is about the same in both experiments, at just over one half. There are still a large number of asymmetric links at the higher power: 76 links are asymmetric, and 28 nodes are end-points for at least one asymmetric link. This is about 33% of the non-zero links, while only 23% of the non-zero links were asymmetric in the 1mW experiment. These measurements illustrate that turning up the transmit power does not eliminate the variations in link delivery ratios across the network. Although increased transmit power will increase the delivery ratio of any particular link, it will also add new non-zero links to the network; these new links will be marginal, with intermediate delivery ratios, and the overall shape of the network's delivery ratio distribution will probably stay the same.

# Chapter 5: Wireless Model

## 5.1 Introduction

This chapter gives a simplified description of how digital radios transmit and receive data packets, along with a description of the sorts of problems radios face when transmitting packets. The purpose of this chapter is two-fold: first, to give a rough sense of why the packet losses described in Chapter 4 occur, and second, to explain the experimentally observed fact that packet loss probabilities vary with the size of the packet. As we will see in later chapters, the accuracy of the PTC metric proposed in Chapters 6 and 7 can be improved by properly accounting for packet sizes. This chapter describes a model that accurately predicts loss ratios at different packet sizes based on the measured loss ratios at two other sizes. Since the model is based on the operation of digital packet radios, we start with a description of how radios work.

## 5.2 Digital Packet Radios

This section provides a brief outline of how a data packet is transmitted as a radio frequency (RF) signal, and how that signal is converted back to bits at the receiver. For a thorough description, see a standard text such as Sklar [73], Proakis [65], or Rappaport [69].

There are essentially three main steps in transmitting the bits in a packet: *coding* and *modulation*, together with packet *framing*. Coding converts the stream of bits in the packet into a stream of *symbols*; modulation converts each symbol into a RF waveform which is then transmitted. Framing is the process of grouping bits into packets and transmitting them with extra information, which is used by the receiver to know when to start demodulation. Demodulation converts the stream of RF signals into symbols, which are then decoded into bits. Although coding, modulation, and framing are logically separate steps, radio designs often interleave parts of each step.

There are many different types of coding schemes; they are typically designed to make the resulting signals more robust to problems in the RF channel. One relevant effect of many codes is that multiple adjacent bits in a packet may be grouped together into one symbol. That is, one symbol represents multiple bits, such as two or four bits. Because some coding schemes

effectively scramble the bits in a packet, bits that are coded into the same symbol may not be near to each other in the original packet.

Just as there are many sorts of coding schemes, there are many modulation schemes. The modulation scheme describes what sort of RF signal is sent for each symbol. Some schemes indicate which symbol is sent by changing the amplitude of the signal, some by changing the frequency or phase of the signal, and some by a combination of all three techniques. No matter what modulation scheme is used, however, the demodulation scheme needs to know where to look for each symbol in the incoming RF signal in order to correctly demodulate it. That is, the receiver must know when in time each symbol starts and ends. Because packet radio systems are typically asynchronous, a radio may receive a packet at any time, and symbol timing information must be re-established for each packet. This is done by adding extra framing information to each packet, such as a *preamble*. A preamble is a predetermined sequence of symbols transmitted at the beginning of each packet. Since the receiver knows what preamble to look for, it can adjust its symbol timing until it finds the expected preamble; at this point the receiver knows it is receiving a packet, as well as where the symbol boundaries lie.

## 5.3 Channel Model

The RF signal travels from the transmitter to the receiver over the RF *channel*. The channel could be a cable, free space, obstacles, or some combination of the three. The *channel model* describes how the RF signal is affected by the channel. In general, a channel has two main characteristics: path loss and delay. In addition to these two characteristics, the receiver's version of the signal is affected by noise, which is received in addition to the transmitted signal. Although noise is not strictly part of the channel model, we consider it here as it also affects wireless link behavior.

### 5.3.1 Path Loss

The transmitter's RF output does not reach the receiver in its original form. The amplitude of an RF signal decreases with distance, as the signal spreads out in space. This attenuation is typically on the order of $d^{-2}$ to $d^{-4}$ for a distance $d$, depending on the environment (e.g. free space or in-building) and the radio frequencies being used [69]. Receivers will receive weaker signals on

longer links. In addition, there may be obstacles blocking parts of the transmitter's signal, such as walls or foliage, which will further attenuate the signal seen at the receiver. Finally, loss in radio hardware such as cables and connectors can also decrease the power of the received signal. The total attenuation is referred to as path loss, and is typically constant over time for a given radio link, assuming that neither end is moving, and that the environment is also static.

### 5.3.2 Multipath

In addition to path loss, a transmitter's signal may be subject to *multipath* effects. When an RF signal is reflected by obstacles, copies of the signal travel to the receiver over multiple paths simultaneously. In general, each of these paths will have a different path loss, and since each path will be a different length, each path will have a different delay. The net result is that the receiver will see several copies of the transmitter's signal, each with a different magnitude and delay. These shifted copies of the transmitted signal will combine together, either reinforcing or degrading each other.

Because the behavior of multipath effects depends greatly on the exact details of the environment, small (or large) changes in the environment or in the locations of the receiver or transmitter can cause the received signal to vary suddenly over time. This variation generally occurs in mobile radio systems, but can also occur in static networks. For example, obstacles such as people, vehicles, doors, or leaves may move in and out of the way of signal paths.

### 5.3.3 Noise

The receiver will see RF signals from sources besides the link's transmitter. Since these signals are not carrying information from the transmitter, they are referred to as noise. A common assumption is that the noise is additive white Gaussian noise (AWGN). AWGN has three main features. First, because it is *white noise*, its power is uniform across the whole radio spectrum; that is, the noise has the same amount of energy, on average, in all frequency bands. Second, white noise is uncorrelated in time; the noise during one time period cannot be predicted from the noise during a previous time period. Finally, because the noise is *additive*, it is simply summed with the transmitter's signal (as modified by the channel) at the receiver. Multiple Gaussian noise sources can be added together to form a single Gaussian noise source.

In real systems there are often many sources of noise that is not AWGN. For example, other transmitters may be using the same radio spectrum. Noise from these transmitters would not be white: it would be focused in one part of the spectrum, and correlated in time. Machinery such as cooling fans or microwaves may also produce predictable time-dependent noise. However, in this chapter, we will only consider AWGN.

Finally, transmissions from adjacent radios in the same network can add noise to an RF link. In many multi-hop wireless networks, all the radios use the same coding and modulation, and transmissions from adjacent radios are likely to have RF signal strength on the same order as the local link. This sort of interference can be particularly damaging because network traffic patterns make the interference highly correlated. To avoid intra-network interference, most wireless networks use a *medium access control* (MAC) protocol to coordinate adjacent transmissions in the same network. MAC protocols, like that used by 802.11 [18], include mechanisms to prevent the *hidden terminal* problem illustrated in Figure 5-1. They often precede each data transmission with a Ready-to-send/Clear-to-send (RTS/CTS) packet exchange between the sender and receiver. The sender transmits a very short RTS packet to the receiver, which replies by transmitting a similarly short CTS packet. The CTS tells nodes around the receiver about the following data packet, so that they can avoid interfering. Unfortunately, RTS/CTS cannot be used for broadcast packets, because there is no unique receiver specified to send back CTS.

### 5.3.4 Asymmetry

If we build a radio link with two identical transmitters and receivers, we might think that the link would be symmetric. That is, the performance of the link, measured as throughput or percentage of packets received correctly, should be identical in each direction. Indeed, path loss and multipath effects are symmetric. However, receiver noise might be different at each end of the link. In addition, in a practical system, especially in a low-cost system, it is unlikely that the radios at each end of the link are precisely the same. For example, although both radios might be set to use the same transmit power, manufacturing and calibration differences and power supply differences (e.g. differences in battery level) can cause the transmit powers to be different.

Because nodes *A* and *C* are out of range of each other, they are hidden terminals to each other, and neither can tell if the other is transmitting. As a result, they may transmit simultaneously, interfering at node *B*.

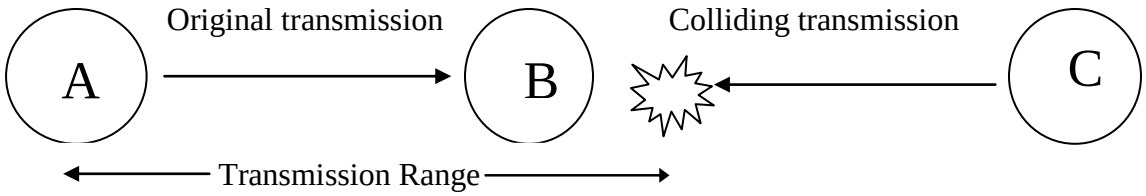


Figure 5-1: The hidden terminal problem.

## 5.4 Effect of Spread-Spectrum

Many modern packet radio systems use *spread-spectrum* techniques [63], which have numerous applications for radio and timing systems. This section briefly describes how spread-spectrum techniques can improve the performance of digital radios in the face of the narrow-band interference and multipath effects discussed in Section 5.2.

The basic idea behind spread-spectrum is that the signal transmitted by a radio is spread out over a much larger range of frequencies than necessary to convey the signal's information. For example, a signal that originally occupied 10 MHz is spread by a factor of ten to occupy 100 MHz The receiver dispreads the signal to its original frequency width before demodulating. Because the power of the signal is spread over a wider frequency range, the signal is less susceptible to interference that occurs in a narrow band of frequencies, such as that from other narrow-band transmitters. This interference only affects a fraction of the spread signal. The advantage that spreading gives over narrow-band interference is termed *processing gain*; as an example, 802.11 radios have about 10 dB of processing gain [18] at 1 Mbps. Spread-spectrum can also help mitigate frequency-selective noise and path loss, by limiting their effects to a small fraction of the original signal. However, spread-spectrum does *not* provide any processing gain over white noise. Because white noise has the same power at all points in the spectrum, it affects a transmitted signal the same amount regardless of how widely the signal is spread.

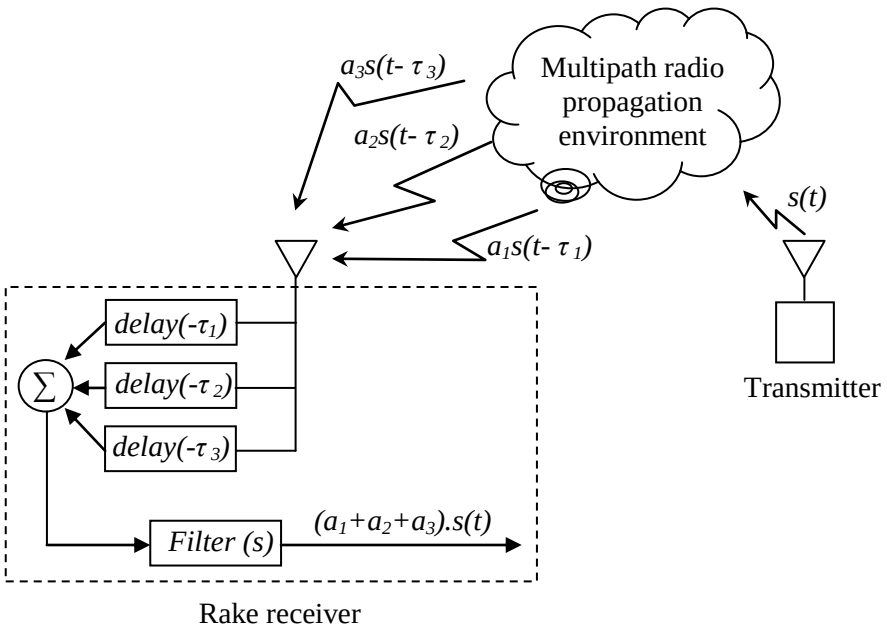


Figure 5-2: A Rake receiver.

The radio propagation environment causes the sum of three delayed and attenuated versions of the original signal $s(t)$ to arrive at the receiver. The Rake receiver splits the received signal into three copies, which it delays appropriately before adding them back together and filtering to recover an attenuated version of the original signal: $(a_1 + a_2 + a_3) * s(t)$. The name of the Rake receiver comes from the structure of its multiple delay lines, which resembles a garden rake.

A second advantage of spread-spectrum is that it can be used to combat multipath effects, using a special sort of receiver design known as a *rake receiver* [64, 65]. The key feature of a rake receiver is that it is able to identify and compensate for the effects of multipath signals, either by filtering out delayed copies of the original signal, or by shifting them in time and recombining them into the original signal, as shown in Figure 5-2. Two important design parameters for a rake receiver are the number of delayed copies of the signal it can identify, and the range of delay for which it can compensate. These parameters can be chosen to match the receiver to the environment in which it operates. For example, the Intersil Prism 802.11 receiver chip, designed for indoor office wireless LAN applications, can handle up to 250 nanoseconds of delay spread at 5.5 Mbps, and 125 ns at 11 Mbps [35]. The measured delay spread of 2 GHz signals in an office building environment ranges from 50 to 150 nanoseconds [57, 58 and 23].

## 5.5 Error Model

The characteristics of the RF channel discussed in Section 5.2 cause most wireless links to have some degree of error. Each link can be characterized in terms of its *loss ratio*, which describes what fraction of packets sent over the link will be incorrectly received. We use the term loss ratio because all damaged packets are treated as lost: any errors are detected using a checksum and discarded by the radio. This section describes a model for predicting the loss ratio at different packet sizes based on the measured loss ratios at a few known sizes. The experiments in Chapter 8 use spread-spectrum 802.11 radios, in the sort of indoor office environment for which they were designed. Therefore, we assume that the radios are robust to most narrow-band and multipath interference, and that errors are a result of a poor signal-to-noise-ratio (SNR) in an AWGN channel.

Before a receiver can correctly demodulate and decode a packet, the receiver must notice that the packet is being transmitted. The details of this depend on the receiver design, but typically the receiver will first notice a higher amount of RF power being transmitted on the frequency used by the radio link. The receiver will then try to synchronize with the signal by looking for framing information such as the preamble. Given that a packet has been transmitted over a particular link, the probability that the receiver successfully detects and synchronizes to that packet frame is $P_f(\text{SNR})$, which is a function solely of the wireless link's SNR (since we are assuming that all errors are due to poor SNR values). The exact form of $P_f$ can be determined from the details of the radio design and implementation.

Once the receiver has detected that a packet is being transmitted, and the receiver is synchronized with the transmitter, the receiver can demodulate and decode each data symbol in the packet. The receiver may incorrectly demodulate a symbol; we again make the common assumption that any error is due to a poor SNR, and that the noise is AWGN. Under this assumption, symbol errors are independent, since the noise which causes any error is uncorrelated over time. Given that a receiver successfully detects and synchronizes to a packet over a particular link, we will write the per-symbol probability of each symbol in that packet being correctly demodulated as $P_s(\text{SNR})$. Like $P_f$, $P_s$ is a function solely of the link's SNR, and the form of $P_s$ depends on the details of the coding and modulation scheme being used.

For a packet with $n$ data symbols, we can write the probability that all symbols are correctly received as $P^n_s$, since the probabilities of correctly receiving each symbol are independent. Therefore the probability of correctly receiving an entire packet is

$$P_p\,(\mathrm{SNR}, \mathrm{n}) = P_f\,(\mathrm{SNR}) \times P_s(\mathrm{SNR})^{\mathrm{n}} \qquad (5.1)$$

That is, $P_p$ is the probability that the receiver detects and synchronizes to the packet, and successfully demodulates every data symbol in the packet. Since *Pf* and *Ps* are solely functions of the link's SNR, $P_p$ is a function of the link's SNR and the packet size *n*. As described above, if we know all the relevant details of the receiver design and the radio's modulation scheme, we ought to be able to write out the function $P_p$. This is actually the link's *delivery ratio*, which is the complement of the link's loss ratio. Once $P_p$ is determined, we can predict the delivery ratio of a link for a given packet size given that link's SNR. That is, if a link's SNR is measured to be *s* (perhaps using some statistics from the radio itself), we can calculate the delivery ratio for a packet of size *n* as $P_p(s, n)$.

However, determining the loss ratio of a link using $P_p$ and the SNR is impractical, because the SNR information *s* can be hard to determine and $P_f$ may not be known. Some radios do not report accurate SNR information, or only report it for successfully received packets, biasing the SNR statistics. $P_f$ may be unknown for several reasons. First, the design of the radio may be too complicated to model accurately; for some radio designs $P_f$ is determined using Monte Carlo simulations. Second, the detailed design of the radio may not be available for analysis to produce $P_f$. This is especially true if commodity radios are being used, as manufacturers are loath to give out the details of their hardware. Finally, although the per-symbol probability functions $P_s$ can be looked up from a textbook for many modulation schemes (including those used by 802.11b), the radio's actual performance may differ from the theoretical performance by some margin. For example, the Intersil Prism 802.11 chipset has a measured symbol error performance of about 3 dB less than the theoretical performance at the 1 Mbps bit-rates [35]. The magnitude of this performance margin may not be known for a particular radio.

We sidestep these problems by measuring each link to determine its $P_f$ and $P_s$. We assume that each link has some fixed, but unknown SNR, at least over a period of time long enough to take measurements. Then $P_f$ and $P_s$ are fixed but unknown quantities for each link. By

measuring $P_p$ at two known packet sizes $n1$ and $n2$, and using equation 4.1, we end up with two equations which can be solved for the two unknowns, $P_f$ and $P_s$:

(5.2)

Let $R = P_p(n1)/P_p(n2)$, $\Delta = n1 - n2$, and assume that both packet sizes had non-zero probabilities of being successfully received. Then

(5.3)

$$\Delta \ln P_s = \ln R \tag{5.4}$$

$$P_s = e \tag{5.5}$$

Substituting equation 4.5 into equation 4.2 gives

$$P_f = \tag{5.6}$$

$$= \tag{5.7}$$

### 5.5.1 Model Inaccuracies

The loss model presented above is extremely simple. For example, it assumes that each link's SNR does not change over time, or if it does, that it changes slowly enough that we can get accurate and consistent measurements for $P_f$ and $Ps$. Also, like Modiano [54], this model assumes that each symbol error is independent. In general, channel noise and interfering transmissions are not AWGN, and will be correlated in time: a symbol is more likely to be received in error if the previous symbol also encountered an error. The model doesn't account for coding techniques such as forward error correction (FEC), which allow a receiver to correctly reconstruct the data in a packet despite some number of errors occurring. The details of how many symbol errors can be tolerated depend on how many errors there are, where the errors are

in relation to one another, and how many bits are affected by each symbol error. Accounting for all of these details, even assuming AWGN, requires relatively complex analysis that is beyond the scope of this chapter.

Despite the model's simplicity, it is consistent with some previous network measurement results. For example, Nguyen et al. [56] report indoor loss and error measurements of the AT&T Wave-LAN, a 900 MHz spread-spectrum radio. Their results show that packet delivery ratios decrease exponentially with increasing packet size. Duchamp and Reynolds [27] also report the results of indoor experiments with the Wave-LAN radio, concluding that the average number of errors per bit in received packets is independent of packet size. Although this result does not imply that bit or symbol errors are independent, it is consistent with that assumption. Willig et al. [77] present error measurements of a radio implementing the IEEE 802.11 physical layer at 2.4 GHz. Their results show that although bit errors are highly correlated, the number of errors per bit does not seem to depend on the packet size, as in the 900 MHz Wave-LAN measurements. The next section shows that although the symbol independence assumptions used to derive the loss model may be too strong, the model still provides accurate delivery ratio predictions.

### 5.5.2 Model Evaluation

To determine the accuracy of the model's loss ratio predictions, seven sets of broadcast experiments were run over two days. During each experiment, for each one-hop link, the source node sent broadcast packets of eighteen different Ethernet sizes, from 50 to 1,500 bytes. Packets were sent at the 1 Mbps bit-rate. The destination node of each link recorded how many packets it received of each size from each sender. The delivery ratio of each packet size over each link is calculated as the number of packets of that size received over the link divided by the number of packets of that size that were actually sent over the link. To smooth out variability over time, the results from each of the seven experiments are averaged, resulting in a single delivery ratio for each one-hop link and packet size.

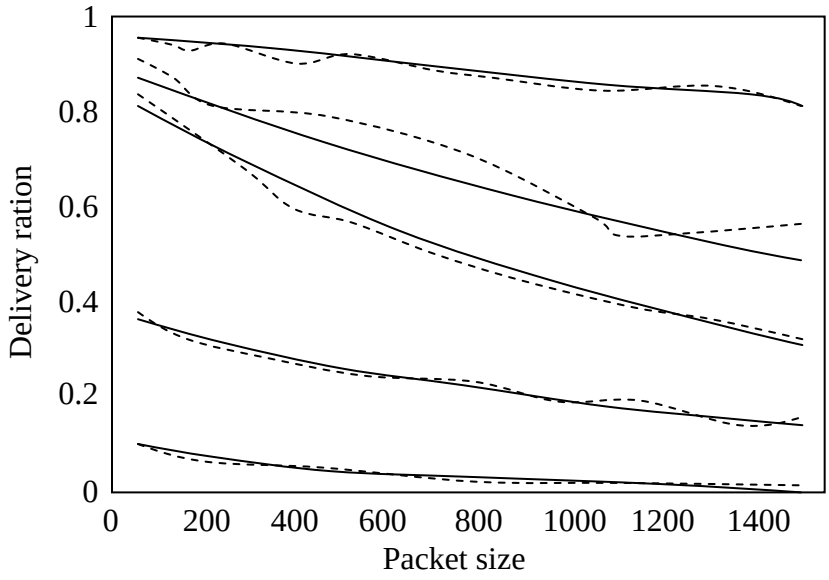


Figure 5-3: Predicted loss ratios for a few links, using the model from Equation 5.1.

The parameters $P_f$ and $P_s$ were calculated using the measured loss ratios of 200- and 1,200-byte packets, which give the best overall results, as shown in Figure 5-5. The smooth gray lines show the predicted delivery ratios for each link, and overlay the dashed black lines with points which show the measured delivery ratios for the same link at various packet sizes. Data are shown for five links: 12 →1, 11 →13, 23 →36, 21 →19, and 17 →18. The measured packet sizes are 50, 100, 150, 200, 250, 300, 400, 500, 600, 700, 800, 900, 1,000, 1,100, 1,200, 1,300, 1,400, and 1,500 bytes.

Figure 5-3 shows the measured loss ratios at each size for a few example links. The figure illustrates how packet delivery ratios decrease with increasing packet size. The figure also shows the predicted delivery ratios, calculated using Equation 5.1, and superimposed over the measured delivery ratios for each link. Although a few links do not have a perfect exponential delivery ratio structure, in general the model closely matches the measured delivery ratios.

Figure 5-4 shows the loss ratio prediction accuracy of Equation 5.1 for each packet size shown in Figure 5-3, and all links, using the measured loss ratios of 200- and 1,200-byte packets. For each size and link, a tiny circle shows the predicted delivery ratio of that size over the given link, versus the measured delivery ratio of that link. Most points lie very close to the line $y = x$, indicating that overall, the model is very accurate. The median prediction error is 0.006.

Figure 5-5 shows how the model performs when we vary the size of the packets whose loss ratios are measured. Each line shows the delivery ratio prediction performance using measurements at a given pair of packet sizes. Each point on a line shows the average prediction error at that size across all links. Clearly, the model more accurately predicts delivery ratios for sizes closest to the sizes whose delivery ratios were actually measured. Using two small packet sizes will lead to inaccurate predictions for large packets; similarly, using two large packet sizes will give poor predictions for small packets. Using one small and one large packet size gives the smallest average error, as this spreads the errors out across all packet sizes. However, the best choice of which sizes to measure at will ultimately depend on the packet sizes we are interested in.

The distribution of prediction errors by size is particularly relevant for the route metric calculation, discussed in Chapter 6, which depends on the delivery ratio of 802.11 ACK packets. We would like to know how accurately we can predict the delivery ratio of ACK packets over a link. Unfortunately, we cannot directly evaluate the prediction accuracy for ACK packets, primarily because with the 802.11 protocol there is no way to measure the delivery ratio of ACK packets in isolation. It is impossible generate ACK-sized packets without first generating a much larger data packet in the reverse direction. Unlike data packets, whose loss ratio we can measure directly by sending broadcasts, the transmission of an ACK packet is conditional on the successful reception of the data packet.

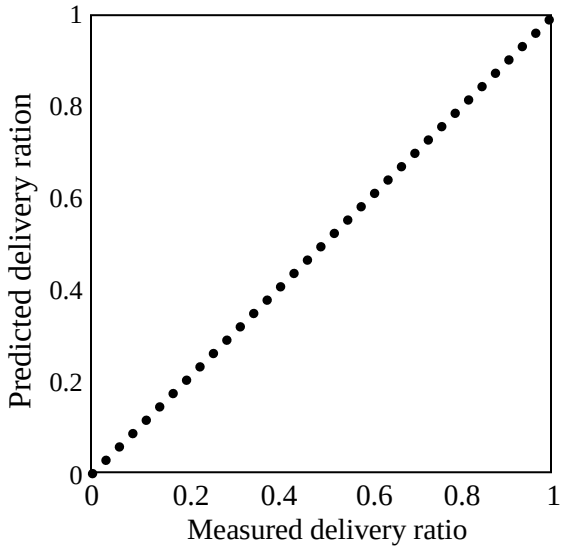


Figure 5-4: Scatter plot of predicted delivery ratio versus measured delivery ratios.

Each point shows the predicted versus measured delivery ratios for one link andone packet size. The predicted delivery ratio for each point is calculated using Equation 5.1 and the measured delivery ratios of 200- and 1,200-byte packets, as in Figure 5-3. The packets sizes used are the same as in Figure 5-3, except for 200-and 1,200-byte packets, which were used for the curve; fit (200 and 1,200 bytes).22 of 339 links were omitted because 1,200-byte packets had a higher delivery ratio than 200-byte packets. The graph has 5,062 points.

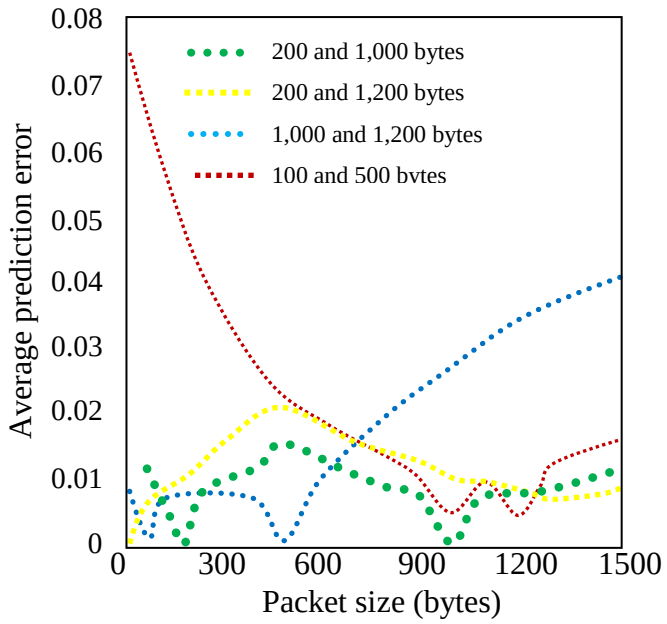


Figure 5-5: Distribution of delivery ratio prediction errors by packet size, using the two-size exponential model in Equation 5.1.

Each line shows the prediction errors at various packet sizes for a given pair of sizes used to find *Pf* and *Ps*. Each line touches the *x* axis at the packet sizes used to make predictions; the predictions become worse as the sizes move further away from those used for the prediction. The best average error across all sizes and links is obtained using the measured delivery ratios of 200- and 1,200-byte packets.

# Chapter 6: Design of PTC

## 6.1 Introduction

This chapter describes the design of the Potential Transmission Count (PTC) metric for finding high-throughput routes. The PTC design is motivated by the causes of low throughput described in Chapter 4: lossy and asymmetric links, and contention between links in a route. This chapter also includes a discussion of why PTC is better than some other proposed metrics.

## 6.2 PTC Intuition

The goal of PTC is to find high-throughput routes. The main intuition behind the PTC design is that because links in a route share the wireless spectrum, protocols can increase throughput in packets per second by decreasing the amount of time each packet uses that spectrum. One way to do this is for protocols to choose routes with fewer links, that is, find minimum hop-count routes.

However, as Chapter 4 showed, a minimum hop-count route may not be the highest-throughput route, if it uses lossy links. Since the 802.11 protocol uses link-level retransmissions, it takes more time to send a packet over a lossy link. This time reduces route throughput in the same way that adding links to a router reduces route throughput: while the sender is retransmitting packets over a lossy link, other links in the route are unable to send. So, in addition to using shorter routes, protocols should also try to use less lossy links.

The second intuition behind PTC is that these two criteria can be combined in to one: the extra transmissions due to adding links can be lumped with the retransmissions on lossy links, producing a total number of transmissions for a path. Protocols should find routes that reduce that total number of transmissions per packet. Routes with fewer total transmissions per packet have higher throughput, because they take less time to send a packet.

## 6.3 Design Criteria

The goal of PTC is to find high-throughput routes by choosing routes with the fewest transmissions per packet. Chapter 4 described several aspects of link behavior that affect route throughput:

**Broad Distribution of Link Loss Ratios** The distribution of link loss ratios is relatively evenly spread from very lossy links to very good links, as shown in Figure 3-5. As a result, the metric should avoid discarding links based on loss ratios. This is primarily because even a lossy link may provide higher throughout over a one-hop route than any available multi-hop routes. It would also be hard to select which threshold should be used classify links. For any reasonable threshold, there are likely to be many links which could be useful, but whose loss ratios are slightly greater than the threshold.

**Asymmetric Loss Ratios** The loss ratios in both directions of a link are often different. For example, a link may deliver all of its data packets in on direction, but drop most of the 802.11 acknowledgment packets in the reverse direction. The metric should consider loss ratios in both directions, and should not draw conclusions about one direction of a link based on its performance in the other direction.

**Multi-hop Interference** As described in Chapter 4 and Li et al. [47], successive hops of a route interfere with each other, reducing throughput even when all links successfully deliver every packet. The metric should account for this intra-route interference as well as the effects of lossy links.

## 6.4 The PTC Metric

The PTC metric for a link is the expected number of data transmissions required to send a packet over the link, including retransmissions. The PTC metric of a route is the sum of the PTC metrics for each link in the route. For example, the PTC of a three-hop route with perfect links is three; the PTC of a one-hop route with a 50% delivery ratio is two.

The PTC of a link is calculated using the forward and reverse delivery ratios of the link. The forward delivery ratio, $d_f$, is the measured probability that a data packet successfully arrives at the recipient; the reverse delivery ratio, $d_r$, is the probability that the ACK packet is successfully received by the data sender, given that the data packet was received successfully. The probability that a data transmission is successfully received and acknowledged is $d_f \times d_r$. A sender will retransmit any data packet that is not successfully acknowledged. Because each attempt to transmit a packet can be considered a Bernoulli trial, the expected number of transmissions for a link is approximated as:

(6.1)

This equation assumes that the probabilities $d_f$ and $d_r$ re constant for a given link, or are at least constant for the duration of link measurements.

PTC has several important characteristics:

- PTC is based on packet delivery ratios, which directly affect throughput.
- PTC detects and appropriately handles asymmetry by incorporating loss ratios in each direction.
- PTC can use precise link loss ratio measurements to make fine-grained decisions between routes.
- PTC penalizes routes with more hops, which have lower throughput due to interference between different hops of the same path [47].
- By minimizing transmission counts, PTC tends to minimize spectrum use, which should maximize overall system capacity.

In addition, PTC may decrease the energy consumed per packet, as each transmission or retransmission may increase a node's energy consumption.

The delivery ratios $d_f$ and $d_r$ are measured using dedicated link probe packets. Each node broadcasts link probes of a fixed size, at an average period $\tau$ (one second in the implementation). To avoid accidental synchronization, $\tau$ is jittered by up to ±10% per probe. Because the probes are broadcast, they are not acknowledges or retransmitted. Every node remembers the probes it receives during the last *w* seconds (ten seconds in our implementation), allowing it to calculate the delivery ratio from the sender at any time *t* as:

Count $(t - w, t)$ is the number of probes received during the window *w*, and $w/\tau$ is the number of probes that should have been received. In the case of the link $X \rightarrow Y$, this technique allows *X* to measure *dr*, and *Y* to measure *df*. Because *Y knows* it should receive a probe from *X* every $\tau$ seconds, *Y* can correctly calculate the current loss ratio even if no probes arrive from *X*.

Calculating a link's PTC requires both *d f* and *dr*. Each probe sent by a node *X* contains the number of probe packets received by *X* from each of its neighbors during the last *w* seconds. This allows each neighbor to calculate its *d f* to *X* whenever it receives a probe from *X*.

The PTC of a route *R* is the sum of the link metrics for each link *l*:

(6.2)

——— (6.3)

where and are the forward and reverse delivery ratios for each link *l* inthe route *R*. DSDV accumulates this metric sum as it forwards route updates. DSR can accumulate the metric sum as it forwards queries, or at the querying host once all route replies have been received.

If the highest-throughput path has three or fewer hops, PTC is likely to choose it: the throughput of these paths is determined by the total number of transmissions, since all of the hops interfere with each other [47]. If the best path has four or more hops, PTC may choose a slower path with fewer hops, since the extra transmissions required by extra hops do not slow down throughput beyond three hops. Route throughput is also affected by the amount of back-off at each node. Two links with the same average PTC but different patterns of packet loss overtime can have different throughputs, which PTC does not account for.

### 6.4.1 PTC Assumptions

PTC makes several assumptions about the link layer. First, PTC is designed for networks with per-link retransmissions, like 802.11 provides. Networks with end- to- end retransmissions will have a different expression for the number of transmissions per packet.

Second, PTC assumes that radios have a fixed transmit power level. With variable power radios, it might be preferable to maximize hop-count, thereby decreasing interference and minimizing the energy used by each packet [68, 32 and 40].

Third, PTC assumes that each node has a single half-duplex radio, and that no two links can use the same radio spectrum simultaneously. If successive links transmit on different logical or physical radio channels, or if nodes have multiple radios, it may be possible for every link in a route to send packets at the same time.

Finally, PTC assumes that all links operate at the same bit-rate. However, when links can run at different bit-rates, it might be faster to use a lossy high bit-rate link than a perfect low bit-rate link.

PTC does not attempt to route around congested links, and in theory should not suffer from the oscillations that sometimes plague load-adaptive routing metrics such as end-to-end delay [9, 42]. To a first approximation, the loss measurements used by PTC do not reflect how busy a link is; a busy link may cause a probe broadcast to be deferred, but won't ordinarily cause it to be lost. This is not true, however, when the network is subject to heavy load like the UDP streaming used in Chapter 8. Because RTS/CTS cannot be used for broadcasts, the 802.11 probe broadcasts are vulnerable to collisions from hidden terminals. Also, 802.11 MAC unfairness can prevent probes from being sent on time. Because the most recent successful sender always resets its back-off window to the minimum size, it is most likely to succeed in the next contention window [8, 10 and 55]. To address this unfairness in the 802.11 MAC, senders that do not get a chance to transmit will continuously decrement their back-off timers during any idle time, including during other senders' back-off. This ensures that eventually one of the nodes waiting to send will win back-off, send a packet, and continue to win back-off for a while. However, even though a node receiving heavy data traffic will eventually get to win back-off and send a probe, that probe (and following probes) could be substantially delayed, reducing the number of probes sent during a given time window. As a result, the node's neighbors will believe that the reverse delivery ratios are very small, and calculate a large metric, causing the routing protocol to avoid using the link. PTC does not specifically account for mobility. PTC may choose good paths despite mobility if the underlying routing protocol can propagate route metrics quickly enough, and if accurate link measurements are available. One way to quickly obtain good PTC estimates might be to use the number of retransmissions per packet reported by the 802.11 interface, but these metrics would still need to be propagated around the network. In general, there is a trade-off between the accuracy of link measurements and a routing protocol's responsiveness to mobility.

## 6.5 Alternative Metric Designs

There are many other techniques and routing metrics proposed for finding highthroughput paths in multi-hop wireless networks. This section discusses some of those techniques with an eye to the design criteria in Section 6.2.

**Masking Errors** A simple approach to handling loss links is to mask transmission errors, either with retransmissions or error correcting codes. For example, the 802.11 ACK mechanisms resend lost packets, making all but the worst 802.11 links appear loss-free. However, retransmission and coding do not make lossy links more desirable for use in routes, as they simply convert lossy links into slow links. The routing protocol should instead find links with lower loss ratios.

**Thresholds** Minimum hop-count routing could be augmented by ignoring links with loss ratios above a specified threshold. However, a below-threshold link may be the best way to reach some node. Also, there may be significant loss ratio differences even among the above-threshold links.

**End-to-End Delivery Ratio** the end-to-end delivery ratio of a route is the product of the per-link delivery ratios along that route; protocols choose routes with the highest product. This metric fails to account for inter-hop interference; it would view a perfect two-hop route as better than a one-hop route with a 10% loss ratio, when in fact the one-hop route would have almost twice the throughput.

**Bottleneck Bandwidth** Protocols choose routes with the highest bottleneck link throughputs. Although this approach may work well for networks where links are relatively isolated from each other, such as wired networks or wireless networks with directional links, it doesn't account for the inter-hop contention in wireless networks with omnidirectional antennas. For example, the bottleneck metric would consider a one-hop route with a 10% loss ratio to be equivalent to a two-hop route consisting of two links, each with a 10% loss ratio. However, the one-hop route would actually have about twice the throughput.

**End-to-End Delay** End-to-end delay is influenced by several factors, including the number of transmissions along a route, queuing time, and MAC protocol back off time. All else being equal, it is probably desirable to choose a route that decreases the delay due to any of these factors. However, end-to-end delay changes with network load as interface queue lengths vary, while the goal of this work is to design a metric that is independent of network load. Load balancing and traffic engineering to decrease queuing and back-off times can be performed with separate algorithms.

**Signal Strength** Many radios can provide measurements of the received signal strength for each link. In theory, link signal strength should predict the probability of packet errors on that

link. Figure 6-1 shows the measured relationship between short-term packet delivery ratios and the received signal strength reported by the radios for a few links in the test network. In practice, using signal strength does not seem to be a practical approach for commodity 802.11 hardware, as there is no good relationship between the signal strength and delivery ratios. Since signal strength measurements are only reported by the radio for successfully received packets, the data is biased. Also, in addition to low signal-to-noise ratios, packet errors can be caused by multipath effects, which are not captured by the receiver's signal measurements.

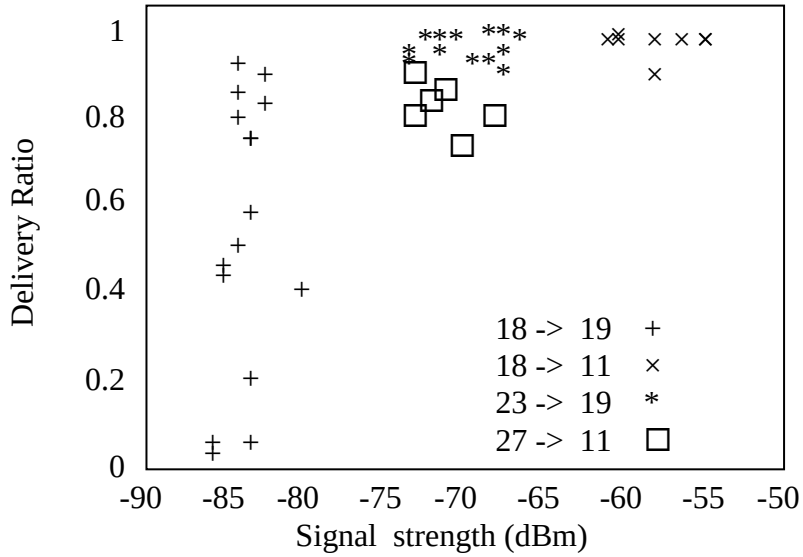


Figure 6-1: Delivery ratios measured over 1 second versus the average received signal strength during that time, for four sample links.

# Chapter 7: Protocol Implementation

## 7.1 Introduction

The routing system in which PTC is implemented has four main parts: the Click toolkit [44], Click-based implementations of the DSDV [61] and DSR [37] routing protocols, and the PTC link measurement algorithm. This chapter describes the implementation details of DSDV, DSR, and PTC, as well as the route metric abstraction that enables these DSR and DSDV implementations to work with many different route metrics.

Because the Click toolkit can run in both user-space and in the kernel, so can the protocol implementations described here. However, running in the kernel provides a few important implementation advantages, such as priority queuing, and easy access to transmission failure notification from the 802.11 MAC layer.

The DSDV protocol is implemented following the description by Perkins and Bhagwat [61], with ambiguities resolved by consulting Broch et al. [11] and the Rice/CMU implementation in the *ns* simulator [59, 66]. The DSR implementation follows the IETF Internet-Draft, version 9 [38].

## 7.2 Operation of DSDV

DSDV is a distance-vector protocol, which uses sequence numbers to ensure freshness, and a *settling time* mechanism to avoid unnecessarily propagating any routes with inferior metrics. We made four changes to the original DSDV design in order to ensure that it uses the path with the best known metric. Before describing those changes, we present an overview of how the published version of the protocol selects routes. Every node has a routing table entry for each destination *D*, which contains four fields: *D*'s identifier (IP address), the next hop on the route to *D*, the latest sequence number heard for *D*, and the route metric. A node forwards packets tothe next hop specified by the current contents of its routing table.

Every node periodically broadcasts a route advertisement packet containing its complete routing table. This advertisement is known as a *full dump*, and occurs at the *full dump period*. Each node maintains a sequence number for itself, which it increments and includes in its own

entry in every full dump it originates. The node copies the sequence numbers for the other entries in the full dump from its routing table. The effect is that the sequence number field in a routing table entry or advertisement entry reflects the age of that entry's routing information.

When a node receives another node's route advertisement broadcast, it updates its own route entries as follows. Suppose node *X* receives an advertisement from *Y* for destination *D* with metric *m* and sequence number *n*. If *n* is newer than the sequence number in *X*'s current entry for *D*, *X* replaces its current entry with the new route through *Y*. *X* also accepts the new route if the sequence number is the same, but *m* is better than the metric of the current route. If *X* has no route to *D*, it accepts the new route. Otherwise *X* ignores the advertised route.

Each route entry has an associated *weighted settling time* (WST). The settling time of a route entry with a given sequence number is the amount of time between when a route with the sequence number was first received, and the time when the best route with the same sequence number was received. The WST is the weighted average of the settling times for recent sequence numbers, and is updated whenever a route with a new sequence number is received.

The WST is used together with *triggered updates* to quickly propagate good routes through the network, while avoiding an explosion of broadcasts. Whenever a node replaces a route entry with a newly received entry, it propagates the new route to its neighbors by sending a triggered update which contains only the changed information. However, triggered updates are not sent until at least 2×WST has passed since first hearing the current sequence number. This prevents nodes from advertising a new route which will likely be replaced later by a better route. In addition, regardless of each route entry's WST, triggered updates are sent at no more than a maximum specified rate. Triggered updates that are delayed are batched together and sent at the next available time.

Finally, DSDV specifies that triggered updates can become full dumps if a large enough fraction of the routes need a triggered update. In this case, all routes with an elapsed WST are included in the full dump, and the node's sequence number is incremented.

## 7.3 Changes to DSDV

The DSDV algorithm we implemented differs from the CMU *ns* DSDV implementation in four ways that improve its performance in the test network. The first two changes were made based

on observations from the literature, while the third and fourth changes were motivated by pathological DSDV behavior observed on the indoor network using detailed packet traces.

The first change affects how the WST is used. The *ns* DSDV implementation does not advertise a route entry until 2×WST has passed since that *particular* route entry to the destination was heard. However, according to our interpretation of the original DSDV description [61], the waiting time before advertising a route should start when the *first* route of each sequence number is heard. Because each node's WST is an estimate of the time between when the node first hears a given sequence number for a destination and when the node hears the best metric with the same sequence number for that destination, the node assumes that it has the best route for a given sequence number after 2×WST has passed. Then it is likely that no better route will be heard for that sequence number, and the best route heard so far should be propagated.

The second change is that our implementation does not use link-level feedback (i.e. 802.11 transmission failure notices) to detect broken links and produce broken-route advertisements. Broch et al. [11] report that broken-route advertisements due to link-level feedback typically because all routes to the particular destination to be broken throughout the whole network, not just those that use the broken link. This makes the destination effectively unreachable from anywhere until its next route advertisement. Our implementation still generates broken-route advertisements when routing table entries time out, but this rarely occurs during the experiments.

The third change is that full dumps are never sent on a triggered update, even if many routes have changed. Triggered updates contain only the changed routes, and full dumps are only sent at the full dump period. This change significantly decreases the routing protocol overhead on our network. Because the indoor test-bed is dense, each node exchanges advertisements with a large fraction of the network's nodes. If a full dump were sent on a triggered update, the sender's new sequence number would in turn trigger a cascade of triggered updates from its neighbors, increasing the amount of protocol overhead.

The fourth and final change (called *delay-use*) is that a route is not used until it is allowed to be advertised. That is, a new route is not used until 2×WST has expired since its sequence number was first heard. With this change, the best route heard for the previous sequence number

is used until the current sequence number's WST has expired. Unmodified DSDV always uses the latest route accepted for a given destination, even if it cannot yet advertise that route.

Delay-use prevents DSDV from prematurely using routes with bad metrics. For example, if there is an asymmetric one-hop route, a node will always hear new sequence numbers along the one-hop link first. Without delay-use, DSDV is forced to immediately use the new one-hop route for routing, even if the PTC metric is poor. In general, shorter routes deliver new sequence numbers first, causing the original DSDV to use shortest paths for some fraction of the time between successive sequence numbers, regardless of the metric in use. With delay-use, DSDV will use the best route with the previous sequence number until the WST has expired and the best route with the new sequence number has likely been heard. Section 8.1.2 shows that delay-use improves the performance of DSDV with PTC.

Figure 7-1 shows pseudo-code for the DSDV routing table update and packet forwarding algorithms, including our changes. The full dump period was 15 seconds, and routing table entries were timed out after 60 seconds. Triggered updates were issued at a maximum rate of one per second. All the DSDV experiments used the four protocol changes described above, unless otherwise noted.

The PTC implementation measures link loss ratios with small probe packets, as described in Chapter 6. Probes contain 134 bytes of 802.11 payload. A PTC node broadcasts one probe per second, and remembers probes received from neighbors over the last ten seconds. Using relatively small probes saves bandwidth; Chapter 8 shows that predictions based on small packets are still useful even when the data traffic consists of large packets.

## 7.4 DSR Implementation

Our DSR implementation follows revision 9 of the IETF Internet-Draft specification[38], following the requirements for networks which require bidirectional links to send unicast data. The implementation is derived from Click-based DSR implementations originally developed at the University of Colorado at Boulder [24, 75]. This section reviews DSR's basic operation as described in the draft, and describes our modifications to support PTC and other metrics.DSR is a reactive routing protocol, in which a node issues a *route request* only when it has data to send. Route requests are flooded through the network, each node appending its own address to each

request it receives, and then rebroadcasting it. Each new request includes a unique ID, which forwarders use to ensure they only forward each request once.

The DSR implementation.

```
Handle_route_ad (Packet p) {
Foreach Route r in p do
handle update(r);

}

Handle_update (Route r) {
// current[]: best route for current seq
// old[]: best route for previous seq

// add link-metric to r.metric
update metric(r);

if (r.seq == curr[r.dest].seq && r.metric < curr[r.dest].metric) {
curr[r.dest] = r;
curr[r.dest].best time = now;
schedule triggered update(r);

}
else if (r.seq > curr[r.dest].seq) {
// save best route of last seq no
old[r.dest] = curr[r.dest];

curr[r.dest] = r;
curr[r.dest].first time = now;
curr[r.dest].best time = now;

// update settling time
old_wst = old[r.dest].wst;
best_t = old[r.dest].best_time;
first t = old[r.dest].first_time;
curr[r.dest].wst = 0.88×old wst + 0.12×(best t − first t);

schedule_triggered_update(r);

}
        // ignore old seqnos and bad metrics
}

        // returns next hop ip address for dst
Lookup_route(IPAddressdst) {
        // use old route if we haven't yet advertised current route
```

```
if (curr[dst].first time + 2×curr[dst].wst> now)
        return old[dst].next hop;
else
returncurr[dst].next hop;
}
```

Figure 7-1: DSDV pseudo-code, including the modifications described in Section 7.2. The WST parameters 0.12 and 0.88 are chosen to produce a reasonable average.

The request originator issues new requests for the same destination after an exponentially increasing back-off time. Route requests are issued with increasing time-to-live (TTL) values, to minimize the range and cost of flooding.

The destination issues a *route reply* in response to every forwarded request it receives. Each reply, which includes the route which was accumulated as the request was forwarded through the network, is source-routed back to the originator along the reverse route. The source node chooses a route using information from the route replies it receives, and source-routes data along this route.

Our implementation stores the results of route replies in a *link cache*, which stores information about each link separately. A node runs Dijkstra's shortest-path algorithm on its link cache to find the best route to a destination.

DSR uses feedback from the link layer to react to link failures. When the 802.11 card signals that no acknowledgment was received after the maximum number of retries, the forwarding node issues a *route error* back to the source, which removes the link from its link cache and then computes a new route. If the source cannot find a route using its link cache, it issues a new route request.

To deal with asymmetric links, each node maintains a *blacklist*, which lists immediate neighbors with unidirectional links to the node. These are links over which the node might receive broadcast requests, but which are unsuitable for unicast traffic. If a transmission failure occurs when forwarding a route reply, the neighbor to which the node was trying to forward the reply is added to the black list, with an entry of *unidirectionality probable*. From that point, the node will not forward route requests received over that link. If the asymmetry of the link is not positively determined for some time, its entry is downgraded to *unidirectionality questionable*. If a route request is received over such a link, the node delays forwarding it while it issues a direct,

one-hop unicast route request back to the questionable neighbor. If a reply is received, the node forwards the original route request and removes the blacklist entry. Otherwise, the node drops the request. Entries are removed from the blacklist when the link is determined to be bidirectional, e.g. by a successful unicast transmission. The DSR specification describes optimizations in which nodes update their link caches using data from packets they forward or overhear on the air. We did not implement any of the optimizations which require the wireless interface to operate in promiscuous receive mode. We also did not implement 'reply from cache,' in which forwarding nodes can respond to route requests with information from their own link caches. All link caches were flushed between experiments, so these decisions should not affect the results in this work.

The nodes do not perform packet salvage, where forwarding nodes try to find alternate routes for queued packets when the head-of-queue packet has a transmission failure, or a route error is received. Instead, queued packets with invalid route information are simply removed from the queue and dropped. Because this implementation has only a five-packet queue, a maximum of five packets could be salvaged at each error, which would not increase throughput appreciably.

To use PTC and other metrics besides hop count, the implementation was modified in a few simple ways. Link probes are used to measure delivery ratios, as in the DSDV implementation. When a node forwards a request, it appends not only its own address, but also the metric for the link over which it received the request. These metrics are included in the route replies sent back to the sender. When anode receives a request which it has already forwarded, it forwards the request again if the accumulated route metric is better than the best which it has already forwarded with the same ID. This increases the chances that the originator will hear about the route with the best metric.

Entries in the link cache are weighted by the metrics which were included inthe route replies. The Dijkstra algorithm finds the route to the destination which has the minimum metric.

## 7.5 Router Configuration Details

If a node is sending large volumes of data, there is a danger that probe packets or routing protocol packets may be dropped or delayed due to a full queue. To mitigate this problem, the implementation maintains separate Click queues for data packets, protocol packets, and link

probes. Each of these queues can hold five packets. These queues all drain into a single queue in the wireless adapter's memory, managed by the driver, which has a capacity of three packets. Loss-ratio probes enter the adapter's queue first, followed by protocol packets, then data packets.

The DSDV implementation looks up a packet's destination in the routing tableafter dequeuing the packet from the data queue, and just before handing the packetto the 802.11b card. This avoids committing to the next hop before queuing, and makes forwarding more responsive to changes in the routing table. This technique depends on the fact that the nodes have only one wireless interface. Figure 7-2shows the DSDV queuing configuration.

The DSR implementation, on the other hand, adds the source-route header to data packets before inserting them into the queue. On a transmission failure or a received route error, a node removes and drops all enqueued packets which include the broken link in their source route. This ensures that the node experiencing the transmission failure does not spend additional time and spectrum retransmitting more packets over the broken hop.

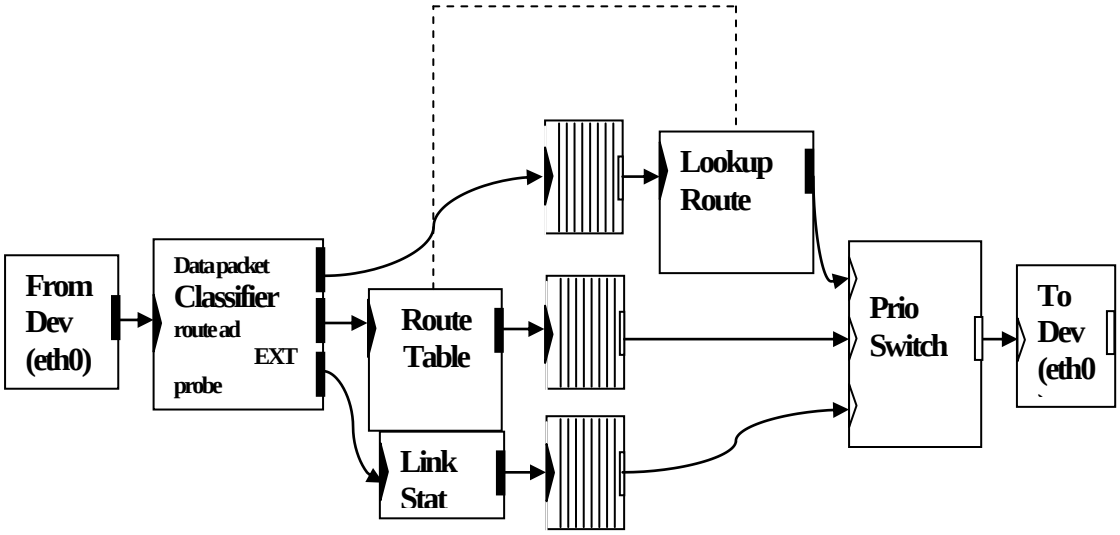


Figure 7-2: DSDV queuing configuration.

Incoming packets are classified as data, route broadcasts, or PTC probes. Data packets are immediately queued; *LookupRoute* determines each data packet's next hop by looking in the *RouteTable*when the packet is dequeued and sent to the interface. Route broadcasts are sent to the *RouteTable*, while PTC broadcast probes are sent to the *LinkStat* module, which maintains link delivery ratio estimates. *RouteTable* periodically produces new route advertisements for the local node, and *LinkStat* periodically produces new PTC probes; these are enqueued separately.

*PrioSwitch* pulls packets from the queues in priority order, sending them to interface. It gives priority to PTC probes, then route broadcasts, then data packets.

## 7.6 Modular Route Metrics

The protocol implementations take advantage of Click's modularity to use modular route metric implementations. That is the metric implementations are not part of the protocol implementations, but are instead modules that the protocol implementations are configured with at run time. This increases the maintainability of the protocol implementations by allowing them to work with new route metrics without modifying the protocol code, and allows different protocols to share the same metric implementations, reducing the work to test a new metric with different protocols.

The drawback of modular route metrics is that the metrics and protocols must use a fixed generic metric abstraction which may not map well to a particular metric or protocol. In practice the benefits outweigh the drawbacks.

```
// Abstract metric datatype
struct metric_t {
bool valid;             // Is this metric valid?
int metric_val;         // Actual metric data; opaque to protocol code
}

// Is the abstract value of M1 'better' than M2? (M1 <M2)
boolmetric_val_lt(metric_t M1, metric_t M2)

// Return the link metric from this node to neighbor N. DATA_SENDER?
// is true if this node is sending data to N over the link, false if N is
// sending data to this node over the link.
metric t get link metric(LinkAddress N, bool DATA_SENDER?)

// Return the metric for the route formed by appending the link with
// metric L to the end of the route with metric R.
metric t append metric(metric_t R, metric_t L)
// Return the metric for the route formed by prepending the link with
// metric L to the front of the route with metric R.
metric t prepend metric(metric_t R, metric_t L)
```

Figure 7-3: Generic metric abstraction. This interface is compatible with the structure of both the DSDV and DSR protocols.

```
// Return measured delivery ratio for sending packets to neighbor N,
// as 0-100 percent
int get_forward_ratio(LinkAddress N)

// Return measured delivery ratio for receiving packets from neighbor N,
// as 0-100 percent
int get_reverse_ratio(LinkAddress N)
```

Figure 7-4: Link measurement interface. All metric implementations that require link delivery ratio measurements use this interface.

The generic metric abstraction is shown in Figure 7-3. A distributed routing protocol can calculate metrics incrementally at each node as it build routes one link at time. When each node adds a new link to a route it is building, the protocol first calls get link metric (), then append metric () (or prepend metric () as appropriate) to combine the new link's metric with the route metric so far. In a centralized protocol, each node can call get link metric () for each of its links, then send those link metrics to a central node for route computation; that central node will repeatedly call append metric () to calculate the metric for each route. This route metric abstraction only works for route metrics than can be calculated incrementally.

The value of the DATA SENDER? flag passed to get link metric() specifies whether or not the node calculating the link metric will be sending or receiving data over the link. This distinction is important because not all metrics are symmetrical, and different routing protocols calculate the link metric at different ends of a link. For example, DSDV calculates link metrics at the data sender end of each link, as the route advertisements flood through the network away from the destination. DSR calculates link metrics at the data recipient end of each link, as route requests flood through the network away from the data sender.

Another modular feature of the protocol implementations is that the link delivery ratio measurement code is contained in its own module that is also configured at runtime. Since many of the metric modules need to know link delivery ratios (such as PTC, bottleneck loss ratio, and end-to-end loss ratio), they can share a single link measurement implementation. The interface to this module is shown in Figure 7-4.

# Chapter 8: PTC Evaluation

## 8.1 Introduction

This chapter presents experimental results that show that PTC often finds higher throughput paths than minimum hop-count, particularly between distant nodes. It also explores the effects of a few individual design decisions in the PTC algorithm, and explains why there is a performance gap between the throughput of the routes with the lowest PTC, and the 'best' routes found by searching the network.

We evaluated PTC by running three kinds of experiments. *Routing protocol tests* evaluate how well PTC improved the performance of the DSR and DSDV protocols. *Static throughput tests* show how the underlying throughput of a particular route can change quickly over time, which is a fundamental limitation on how well PTC can predict which route to use. *Single link tests* characterize the accuracy of PTC predictions over a single link at a time, as well as illustrate howdelivery ratios of a single link can vary quickly, which also affect how accurately PTC can choose good routes. The following sections describe each set of experiments in more detail.

## 8.2 Routing Protocol Tests

The routing protocol tests show well PTC improves the throughput of a complete routing system. As a result, they include protocol-specific behavior and overheads. We tested PTC with both the DSDV and DSR routing protocols, which are described further in Chapter 7.

### 8.2.1 Experimental Setup

Unless otherwise stated, the experimental setup is as follows. The test, radio configuration, and packet size are as described in Section 4.1: 134-byte UDP payloads, 1mW transmit power, RTS/CTS disabled. The DSDV implementation includes the improvements described in Section 7.2 for both PTC and the hop count metric. The DSR implementation is as described in Section 7.3.

The protocol performance data presented below were collected during a few separate 'runs'. An entire run takes anywhere from 18 to 72 hours, depending on the experiment

parameters. A run considers each pair of nodes in turn. For each pair, one experiment is performed for each routing protocol variant. At the start of each experiment, the routing software is reset (all tables and protocol state are cleared), then the routing protocol and PTC probe algorithm, if PTC is used, are allowed to run long enough to stabilize and setup forwarding and routing state(typically 90 seconds for DSDV). Next, the sending node of the pair sends UDP data packets as fast as the radio allows through the routing system to the destination. The destination counts how many packets arrive over 30 seconds to calculate the average throughput.

After the protocol tests run for each pair, the 'best' static route is identified for that pair by testing the throughput of 10 candidate routes, as described in Section 4.2. Like the routing protocol tests for each pair, the static routes are also tested by sending UDP packets as fast as possible and counting how many are received for 30 seconds. However, packets are forwarded along a static route according to source routes in each packet header, rather than running a dynamic routing protocol. The per-pair protocol interleaving ensures that the results from different routing protocols and the static routing are comparable for the same pair of nodes, since the experiments for each protocol are run within a few minutes of each other for a given pair.

Each graph below is labeled with the run from which it came. Graphs with the same run number are comparable. Graphs with different run numbers should not be compared, since the network's behavior changes substantially with time. The graphs below do not include error bars, but are representative of the many runs performed.

In DSDV experiments using PTC or minimum hop-count, the routing protocol runs for 90 seconds, immediately after which the source sends data packets as fast as possible for 30 seconds. As described in Chapter 6, the heavy load causes the MAC protocol to become extremely unfair, distorting the PTC measurements. To minimize the effects of MAC unfairness, every node routes packets using a snap-shot of its route table taken at the end of the 90-second warm-up period, before any data is sent. The snapshot also makes the DSDV results more comparable to the 'best' static route results, since the static route tests are not allowed to switch routes in the middle of testing a particular route.

All experiments run with the appropriate routing overhead. That is, while measuring the throughput of routing with the PTC metric, nodes send periodic PTC broadcast probes. While measuring the throughput of DSDV (with either metric),nodes sends DSDV routing advertisements, just as a production routing system would.

### 8.2.2 DSDV Performance

Figure 8-1 compares the throughput CDFs of paths found by DSDV using PTC and minimum hop-count, between 100 randomly chosen node pairs. This data is taken from the same run as in Figure 4-2, and shows that DSDV using the PTC metric often finds much faster routes than the minimum hop-count metric.

There are two main regions in Figure 8-1. The right half shows node pairs that could communicate directly, with loss ratios less than about 50% (i.e. with throughput greater than the maximum possible two-hop throughput of 225 packets per second). In these cases the minimum hop-count metric finds the one-hop route, which is the best route, and there is no opportunity for PTC to perform better. The left half corresponds to node pairs with a high direct loss ratio, for which the best route has more than one hop. In this region, the sensitivity of PTC to differences among the many different paths of the same length allows it often to find better paths than hop-count.

Figure 8-2 shows the same data as Figure 8-1, but organized as a scatter plot toallow a direct comparison between the performances of each metric for individual pairs. Each pair is represented by one point; the point's *y* value is the throughput obtained by DSDV using PTC, and the *x* value is the throughput obtained by DSDV using minimum hop-count. The upper-right quadrant shows pairs where PTC and minimum hop-count both used the one-hop path.

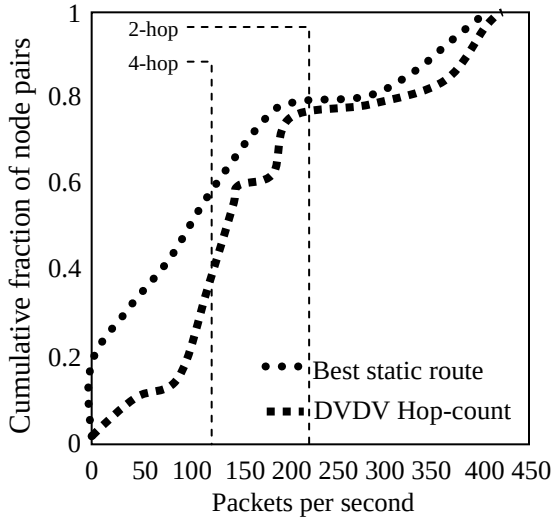


Figure 8-1: PTC finds higher throughput routes than minimum hop-count. This data is taken from the same experimental run as Figure 4-2. Each point represents one of 100 node pairs.

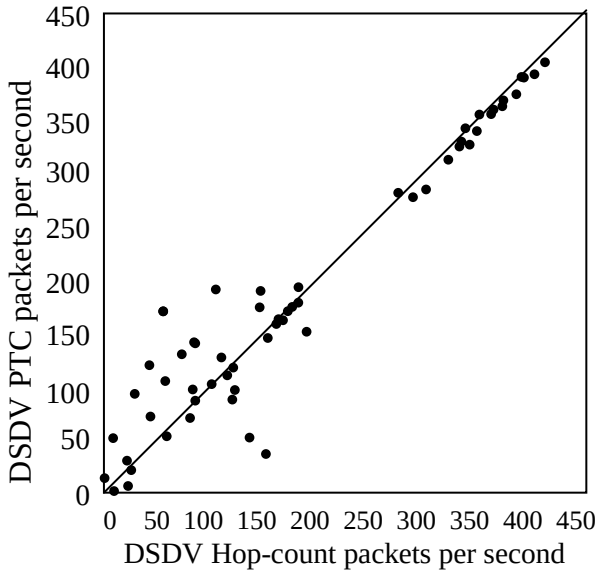


Figure 8-2: The PTC and hop-count data from Figure 8-1, plotted on a per-pair basis. The *x* value of each point shows that pair's throughput for DSDV with minimum hop-count; the *y* value shows the throughput for DSDV with PTC. Points above the line $y = x$ are pairs where PTC outperformed hop-count.

PTC outperforms minimum hop-count by the greatest margin when the hop count metric uses links with very asymmetric loss ratios. This is illustrated by the points with *x* near zero and with *y* relatively large. In these cases, minimum hop count chooses links that deliver routing updates in one direction but deliver few or no data packets in the other, while PTC correctly avoids those links.

The points for two pairs in Figure 8-2 lie well below the $y = x$ line; this is because of variations in link quality between the PTC and minimum hop-count tests for those pairs. For the first pair, both PTC and hop-count used the same route, so the difference is due to an underlying change in the route's throughput. For the second pair, PTC used a slower 3-hop path while hop-count used a two hop path; PTC avoided using one of the links in the two-hop path because the measured delivery ratios were very poor. It is likely that the link's quality was different for the PTC and hop-count tests.

Figure 8-3 shows the throughput CDF for TCP traffic routed using DSDV with PTC and minimum hop-count. The figure also shows the 'best' static route TCP throughput found for each pair. TCP sent data for 30 seconds between each pair. All experimental parameters were the

same as for the UDP tests, except that the packet size was varied by TCP according to its congestion control algorithm. Hop-count does particularly poorly for TCP. Although we have not examined the results in as much detail as the UDP results, we conjecture that the hop-count performance suffers for two main reasons. First, since TCP traffic requires good routes in both directions in order to send back end-to-end TCP acknowledgments, there are twice as many chances for hop-count to select a bad route: once in each direction. Second, the TCP back-off algorithm amplifies the effects of any errors in the underlying route.

Figure 8-4 shows the UDP throughput for packets with a 1,386-byte payload. Although PTC still offers an improvement over minimum hop-count, the gain is not as large as for small packets. This is because PTC is still using small probes to estimate the link metrics. Since small packets are more likely to be delivered, PTC is incorrectly over-estimating the quality of each link and causing DSDV to pick sub-optimal routes. For example, if the single-hop direct route between two nodes has a PTC probe delivery rate of 51%, PTC will use it; however, the delivery rate of 1,386-byte packets on such a link is likely to be much smaller, so a route with a larger number of higher-quality links would have been preferable.

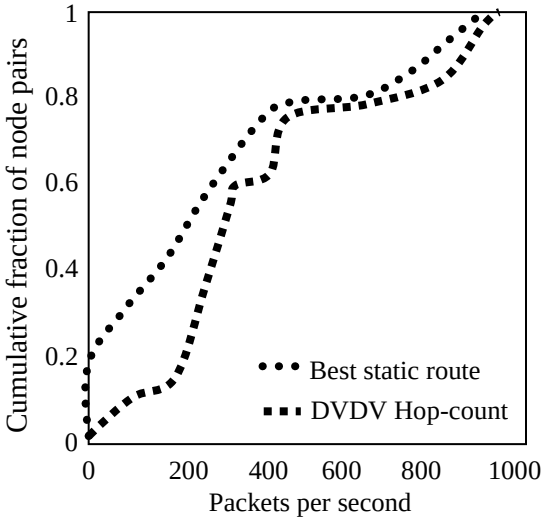


Figure 8-3: PTC finds higher throughput routes than minimum hop-count for TCP traffic. This data is from an experiment similar to Figure 8-1, except that data was sent using the TCP protocol rather than fixed-size UDP packets. The same 100node pairs are tested as in Figure 8-1.

Figure 8-5 shows the results of PTC versus minimum hop-count from a third run with the radios transmitting at 30mW instead of 1mW. The packet size is134 bytes. When nodes send at

the higher transmit power they have more links, as shown in Figure 4-5. This makes the network much more connected, decreasing the average hop-count required for nodes to communicate. As a result, PTC has fewer routes to choose from, and minimum hop-count has a lower chance of choosing a bad route. Figure 8-5 shows that PTC still provides some advantage in the more highly connected network.

**Impact of Asymmetry**

Some fraction of PTC's gains comes from avoiding extremely asymmetric links. The problem of routing when there are asymmetric links has been addressed in previous work by Lundgren et al. [50] and by Chin et al. [15]. These authors propose a link handshaking scheme to detect and avoid asymmetric links. In this scheme, a node *X* only accepts route updates from a neighboring node *Y* if *Y* is advertising a direct route to *X*. A node bootstraps the handshake by advertising provisional route entries, which indicate that the node has 'seen' another node, but not yet accepted routes from it.

We implemented the handshaking scheme for DSDV with the minimum hop count metric. Figure 8-6 compares link handshaking to the PTC and minimum hop-count metrics. Although link handshaking often improves throughput over minimum hop-count alone, PTC finds faster routes. PTC's link measurements allow PTC to discriminate between links with varying degrees of asymmetry and quality.

**Effects of DSDV Modifications**

Section 7.2 described modifications to DSDV designed to increase its responsiveness to metrics. The *delay-use* modification causes DSDV to delay using a newly received route until it is permitted to advertise the route (i.e. 2×WST has passed). Figure 8-7 shows that the delay-use modification improves the performance of DSDV with PTC.

The small packets used to measure link loss ratios incorrectly predict the actual transmission counts for large packets. This graph shows 40 pairs randomly chosen from the 100 pairs used in the previous figures. The maximum 1-hop throughput of 1,386-byte data packets at1 Mbps is 82 packets per second.

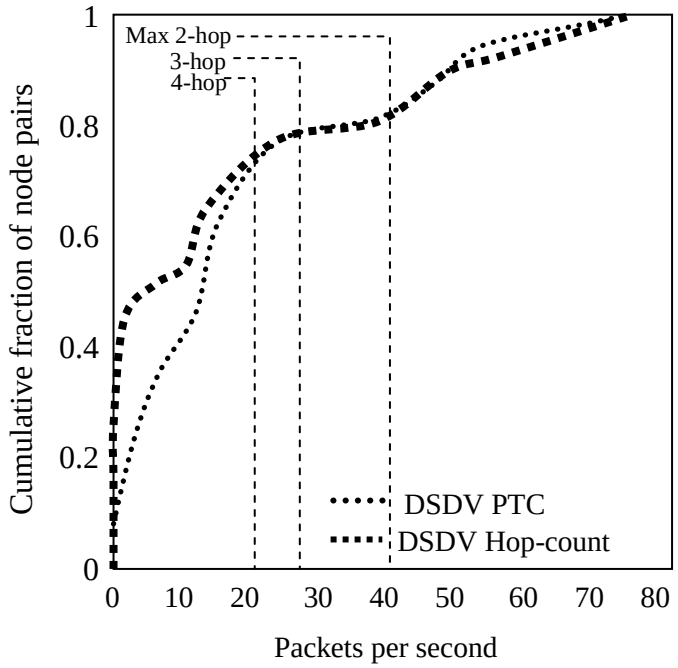


Figure 8-4: PTC provides less of a throughput advantage over minimum hop-countwhen using large (1,386-byte) packets.

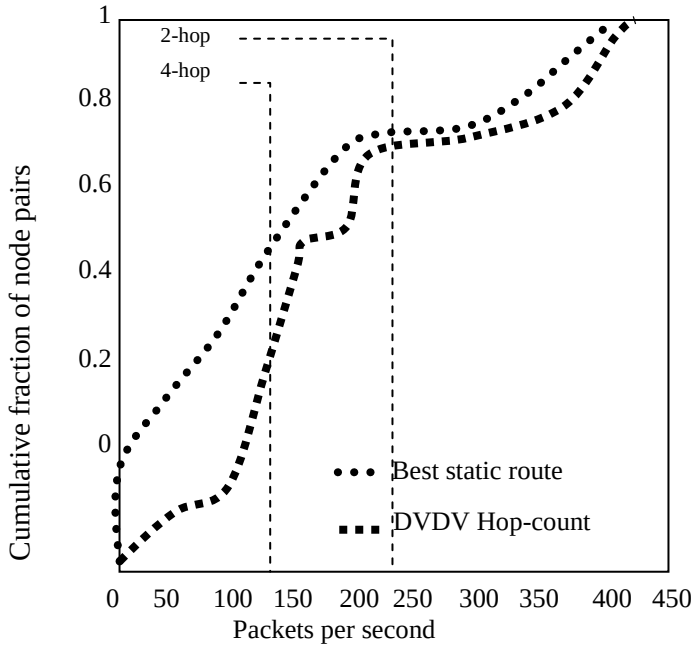


Figure 8-5: PTC versus minimum hop-count.

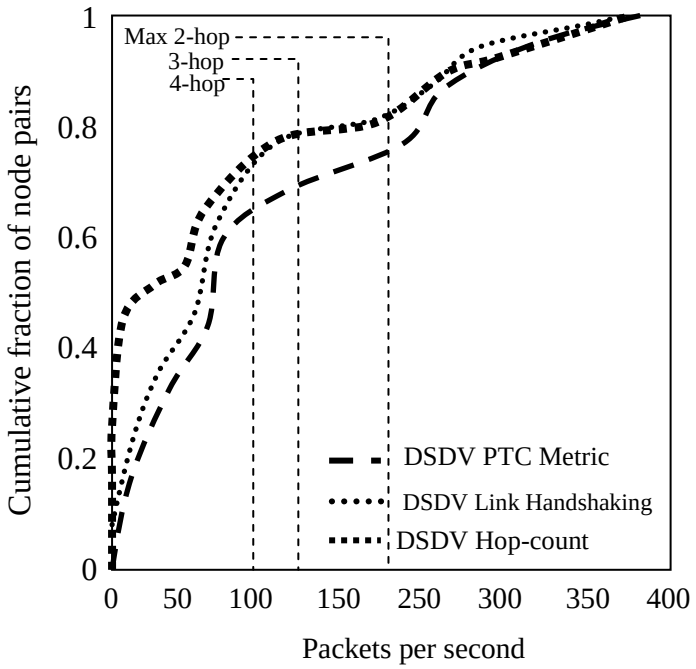


Figure 8-6: PTC provides a significant throughput advantage over a simple handshaking scheme which avoids very asymmetric routes. This is because PTC can make fine-grained decisions between links with varying degrees of asymmetry and quality.

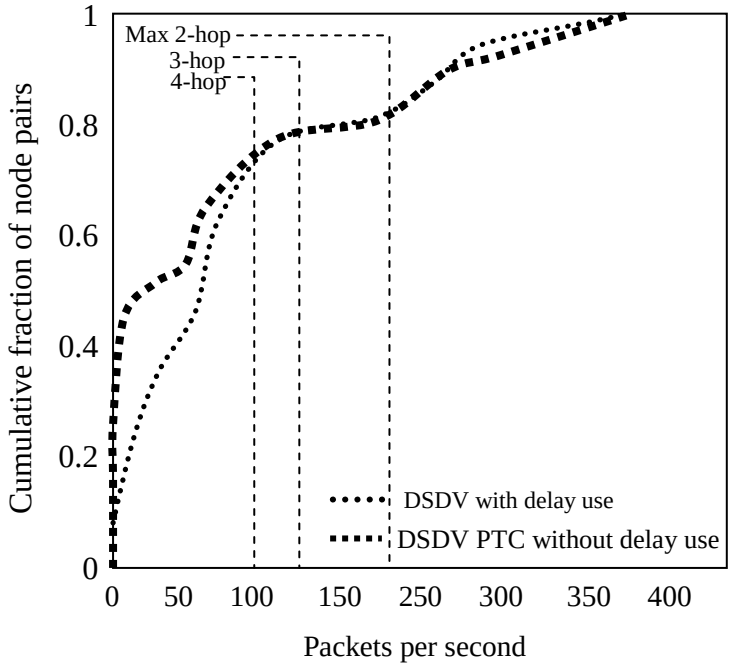


Figure 8-7: DSDV PTC with and without the delay-use modification to DSDV.

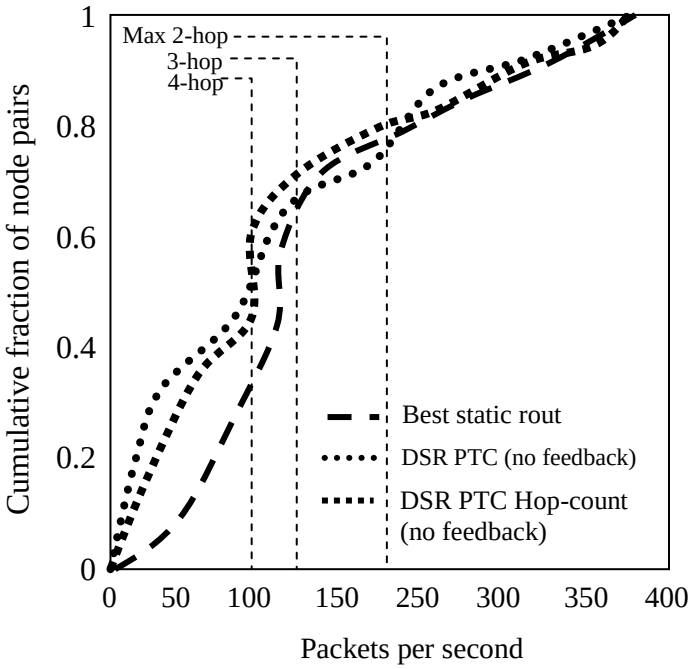


Figure 8-8: Throughput CDFs for DSR PTC compared with DSR hop-count.

### 8.2.3 DSR Performance

This section evaluates the performance of the DSR routing protocol with the PTC metric. As described in Section 7.3, DSR uses link-layer transmission failure feedback to avoid bad routes. To isolate the effects of using PTC with DSR, we evaluated DSR performance both with and without link-layer feedback enabled.

Figure 8-8 shows the effect of using the PTC metric with DSR without link layer feedback, for the same 100 pairs as in Figure 8-1. Because DSR never learns about transmission failures, no forwarding node ever issues any route errors. Thus DSR uses only the best route found by the initial route request, as determined by the metric.

The figure shows that PTC greatly improves initial route selection in DSR compared to minimum hop-count. This is consistent with the DSDV results in Section 8.1.2. Minimum hop-count essentially chooses randomly from all the shortest routes the source obtains from the initial route request; as illustrated in Figure 4.3, this is often not the best route. PTC helps the source pick an initial route with high throughput. Figure 7-9 illustrates the performance of PTC with DSR's link-layer feedback enabled. PTC provides a small benefit to some pairs in the

intermediate and low throughput ranges (the middle and bottom of the CDF). However, failure feedback alone allows DSR to perform almost as well as DSR with PTC.

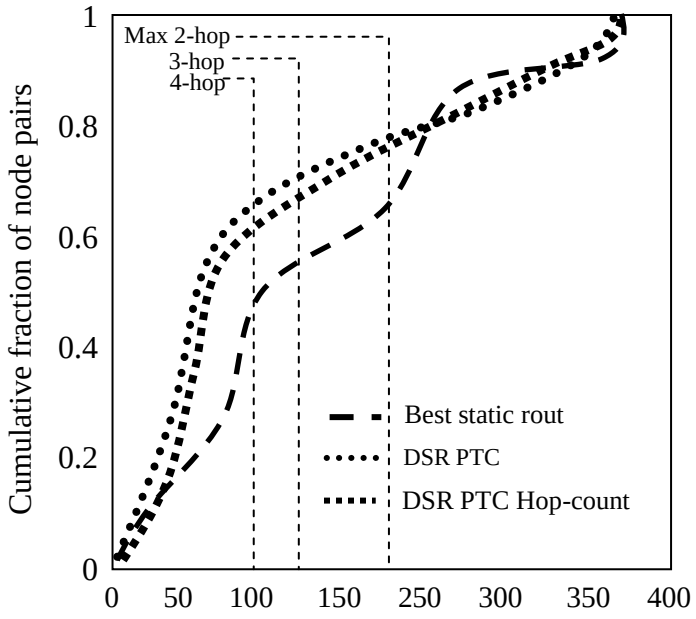


Figure 8-9: DSR PTC compared with DSR hop-count, with link-layer transmission feedback enabled.

### 8.2.4 PTC versus 'Best'

One main question is why there is a gap between the throughput distribution of the routes found using PTC, and 'the' best static routes found by searching the network. Although a route's underlying transmission count does a good job of explaining the route's throughput, routing protocols only use PTC estimates of those transmission counts. PTC mispredicts actual transmission counts due to packet size effects and time variation between when link measurements are made, and when route throughput is tested. Furthermore, the underlying throughput of many routes can change significantly over very short periods of time.

Figure 8-10 shows how the PTC estimates available to the routing protocol may wrongly predict the actual transmission counts for each route. The PTC estimates are calculated using the broadcast delivery ratios for each link measured with the broadcast probes. The actual transmission counts for each route are measured using transmission counters supplied by the 802.11 radio interface. Because the PTC estimates are often incorrect, the metric can disorders routes, causing the routing protocol to use a route with a lower throughput than the best route available at a given time.

Figure 8-11 shows that if a routing protocol could get a more accurate estimate of the transmission count for each route, the protocol could more accurately choose the highest throughput route. The graph shows the predicted throughput for each route versus the route's actual measured throughput. The predicted throughput is calculated as *B*/TXC, where *B* is the maximum link throughput for that packet size (451 packets per second), and TXC is the route's average measured transmission count, obtained from the 802.11 radio interface.

The predicted throughput has a roughly linear relationship to the measured throughput, with some errors, and a systematic offset. The causes of the error and offset are unknown, although there are a few likely suspects about which we can speculate. For example, a route's throughput is affected by both the underlying throughput of its links (determined by their lossiness, and reflected in the transmission count), as well as other 802.11 traffic in the vicinity of those links which contends for the same piece of radio spectrum. That is, a one-hop route over a perfect link might only have half the link throughput if another nearby 802.11 radioswas operating. Since our experimental environment is filled with many other802.11 radios, this case cannot be ruled out. As for the systematic offset, it might be explained by an inaccurate estimate of the maximum ideal throughput for each route. This might happen if the 802.11 interfaces are performing back-off in as lightly different manner than described in the 802.11 specification. Furthermore, the throughput predicted using transmission counts does not account for the exponential back-off performed when a radio has to retransmit a packet multiple times.

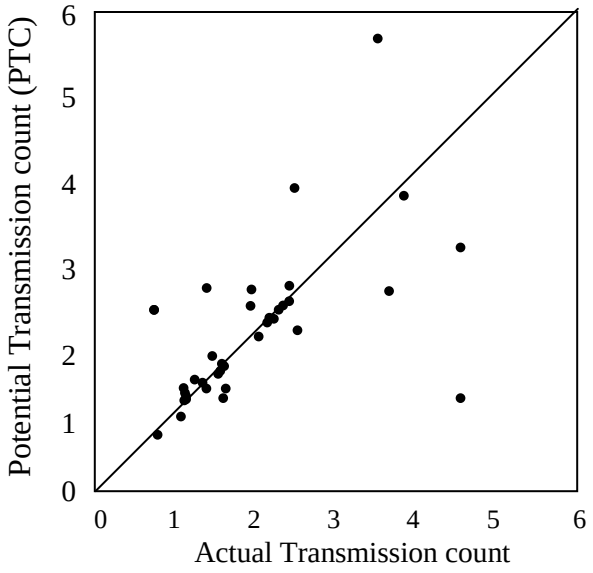

Figure 8-10: The PTC estimates used by DSDV mispredict the actual transmission counts incurred by a route.

The *x* value of each point shows the average measured transmission count for a route found by DSDV, while the *y* value shows the PTC metric for that route. This data is from the same run as Figure 8-1.

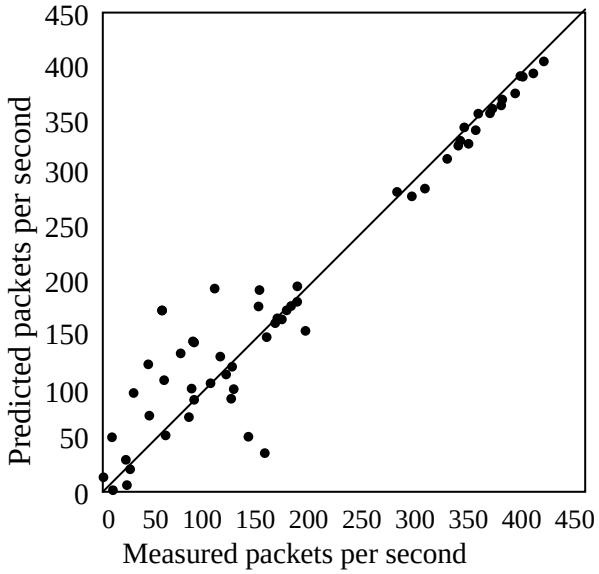


Figure 8-11: A route's underlying transmission counts.

Do a good job of predicting that route's throughput. The *x* value shows the measured throughput of a route found by DSDV using PTC, while the *y* value shows the throughput predicted using the average measured transmission count. This data is from the same run as Figure 8-10.

## 8.3 Static Throughput Tests

In addition to the routing protocol tests, we ran simple throughput experiments using static routing to explore how the throughput of a particular route changes over time. Each experiment tested the throughput of several routes. For each route, UDP packets were sent as fast as possible for 30 seconds, then after a wait of 30seconds, packet were again sent as fast as possible for 30 seconds. Figure 8-12shows the results of these tests for various packet sizes and times of day. During the daytime route throughputs can change substantially, although the throughputs seem

very stable at night. One explanation for the daytime variation is the increased level of activity in the lab. People are moving around in lab, opening and closing doors, and running equipment. Another daytime source of variation is the lab's 802.11 access point infrastructure. The 802.11 access points are essentially idle at night, but under heavy use during the day as many lab members and visitors use their laptops to access the lab network wirelessly. This increased and variable802.11 traffic load can reduce the throughput available to the test.

These experiments show that part of the difference between 'best' and DSDV with PTC in the routing experiments can be explained by underlying variation in the route throughputs. That is, the throughput of the 'best' route may not have been available in the network when the DSDV with PTC experiment was performed. Or, the route selected by DSDV and PTC was no longer the highest throughput route when the data packets were sent. Furthermore, since the 'best' throughput is the maximum throughput of several different routes is not sensitive to changes and reordering of those route throughputs: it just selects the highest throughput, and will likely always do better than DSDV. Although the throughput variation does not occur at night, since the route experiments spanned several days, many of them were performed during busy lab hours and would be affected by the route throughput variations.

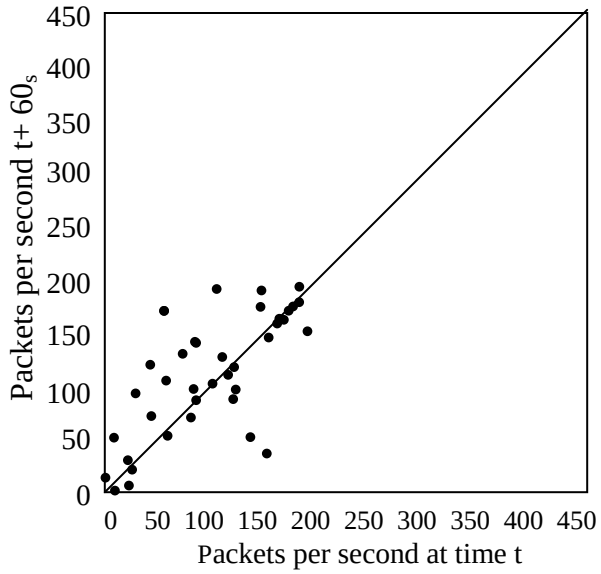


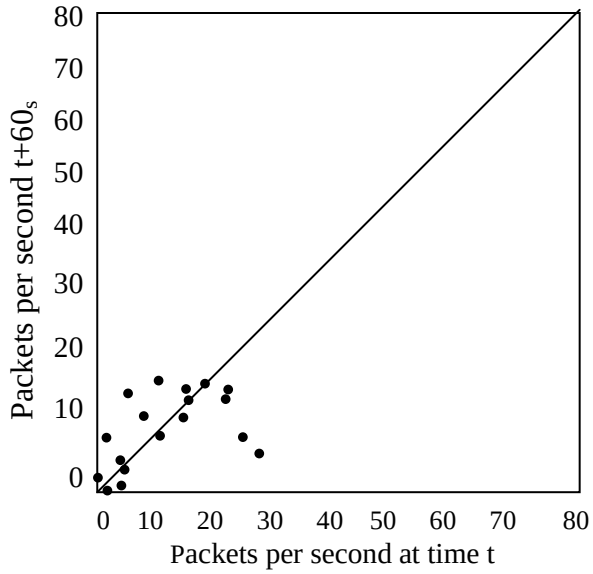

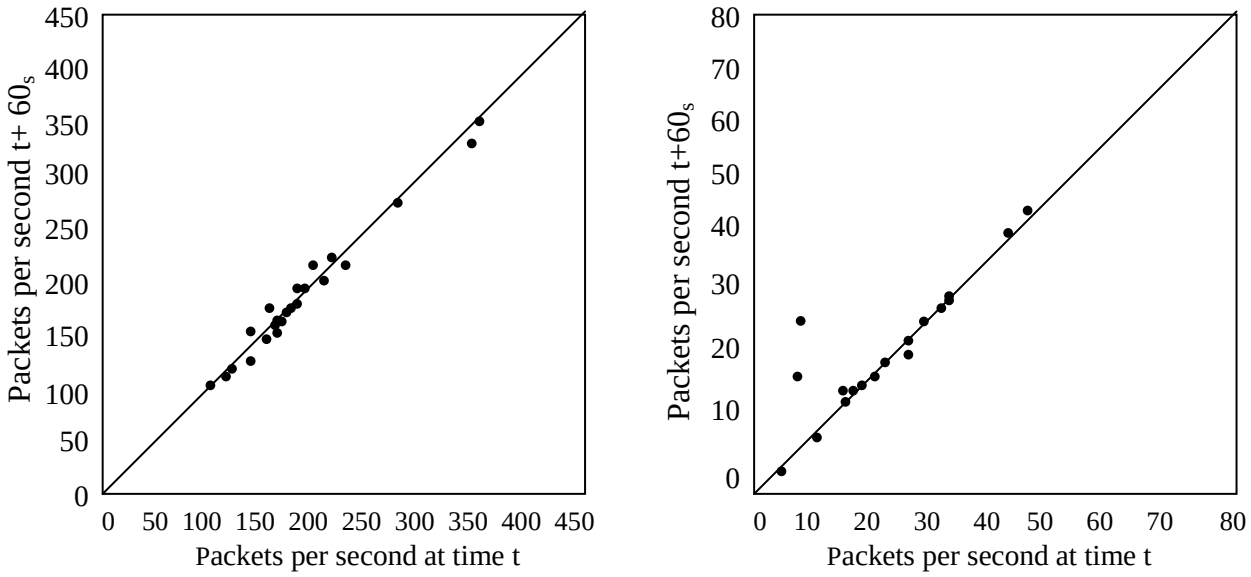


Figure 8-12: The throughput of a route can sometimes vary significantly over a very short period of time.

The *x* value for each point shows the throughput of a route measured over 30 seconds. The *y* value shows the throughput of the same route measured 60 seconds later, also over 30 seconds. The left graphs show measurements of 134-byte packets; the right graphs show 1,386-byte packets. The top row shows measurements from the daytime, the bottom row shows measurements from the nighttime.

## 8.4 Single Link Tests

To better understand the accuracy of link PTC measurements, we performed several experiments with individual links. In each experiment, broadcast delivery ratios are measured to estimate the PTC for a link, and then unicast packets are sent to measure the actual transmission count of the link. Broadcasts are typically sent before and after the unicasts to measure how much the delivery ratio of a link varies during the unicast part of each experiment.

Figure 8-13 shows how well the PTC estimates actually predict the link transmission counts. The PTC estimates are calculated as described in Equation 5.1,where the values $d_f$ and $d_r$ are measured using the forward and reverse link broadcast tests respectively. The graphs illustrate two main points. First, measuring the reverse ACK delivery ratio of each link (from destination to source) using the minimum-size Ethernet packets provides more accurate PTC estimates, which can be seen by comparing the graphs on the left (data-size reverse

measurements)with the graphs on the right (minimum-size reverse measurements). Second, as with the static route throughput tests, there is considerable day-night variation. At night, almost all links work very well, while during the day there is a wider distribution of link performance, along with more short-term time variation of link performance.

Figure 8-14 shows how well PTC estimates actual link transmission counts for larger 1,386-byte packets. The general behavior of the graphs is the same as in Figure 7-13, except that there are still considerable prediction errors even at night. This is likely because the much larger 1,386-byte packets do an even worse job of predicting the ACK delivery ratios in the reverse direction than the 134-bytepackets.

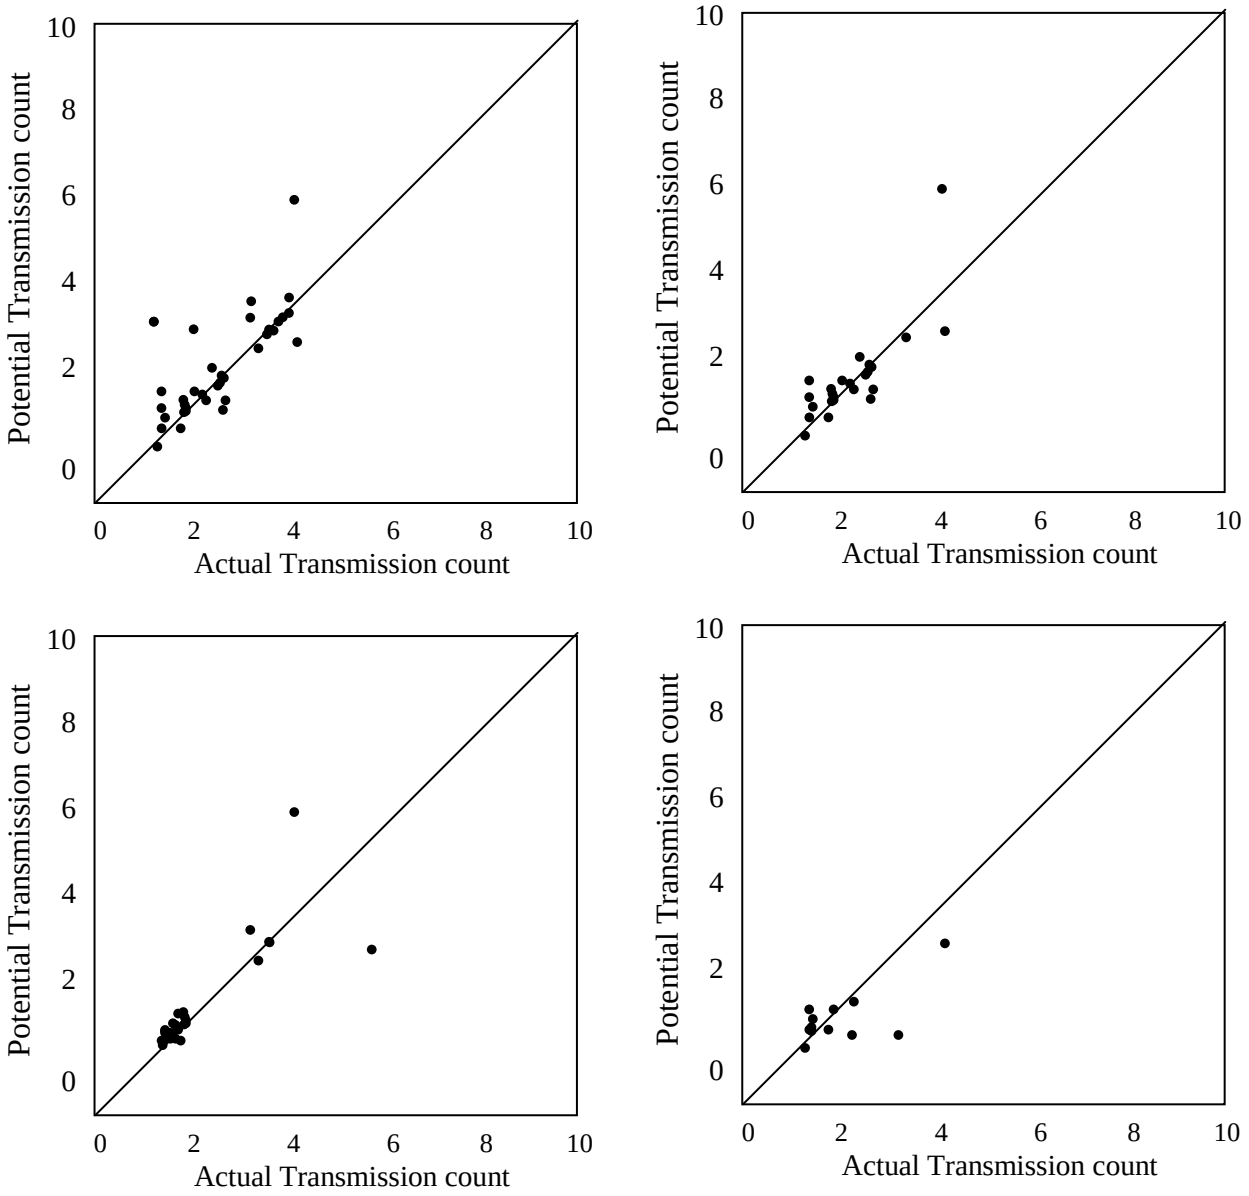

Figure 8-13: PTC versus measured transmission count for 134-byte data packets over a single link.

PTC is calculated using the link's broadcast delivery ratio in the forward and reverse directions, measured immediately before sending the unicast data. The left graphs measure the reverse delivery ratio using broadcasts that are the same size as data packets; the graphs on the right use minimum-size Ethernet packets in the reverse direction (14 bytes). The top row shows measurements from the daytime, the bottom row shows nighttime. The nighttime graphs have many points overlapping at (1, 1). Some points are not plotted because they are out of range of the graph.

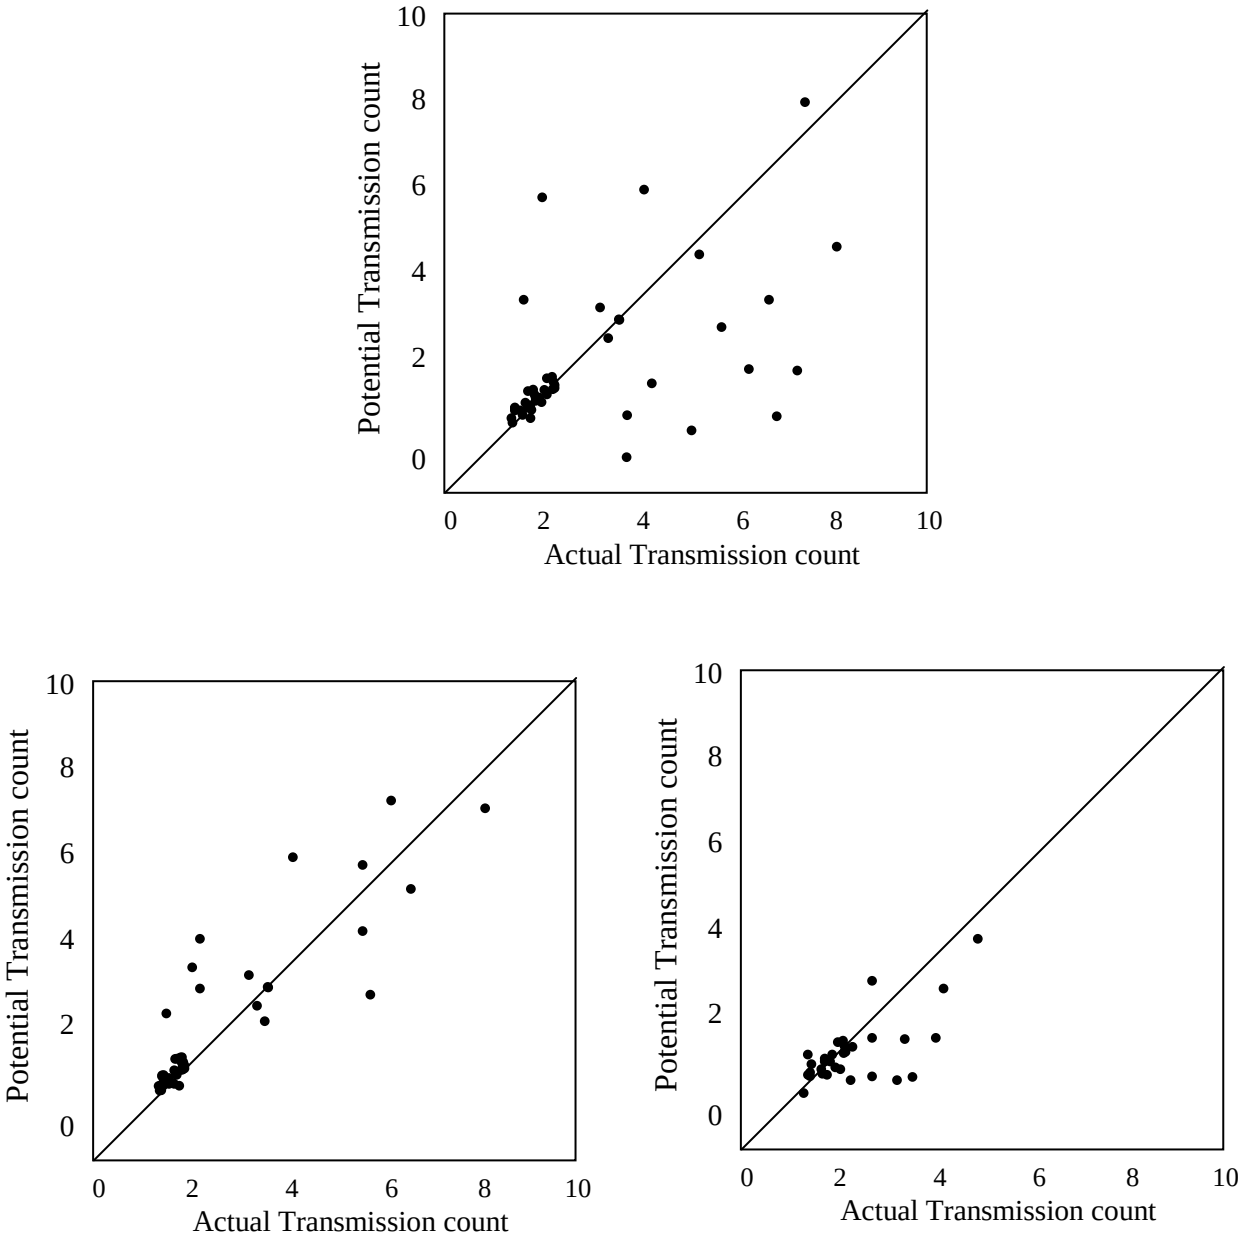

Figure 8-14: PTC versus measured transmission count, as in Figure 7-13, but for1, 386-byte data packets.

Figure 8-15 shows how much link delivery ratios can change over time. The graphs show the measured forward delivery ratio of each link at the beginning and end of each single link test. The graphs show results for both 134- and 1,386-bytepackets, during day and night. In all cases there is significant variation between the two measurements. This is less so for the 134-byte packets, which have relatively high delivery ratios across the board. Both the 134- and 1,386-byte sets of experiments show diurnal variation, as in the static throughput tests. The implication of these graphs is that the transmission counts of the links which PTC is trying to estimate can change by a significant amount in a short time. Even if PTC chooses a maximum-throughput route, that route may not be the fastest route 30seconds later.

## 8.5 Evaluation Summary

This chapter showed how PTC increases the throughput performance of the DSR and DSDV routing protocols. It also used more focused static throughput and single link experiments to understand the gaps between the throughputs of routes found using PTC and the 'best' routes found using static routes. We identified two main causes of the discrepancy. First, PTC mispredicts the transmission count of links because it measures the reverse ACK delivery ratios using the wrong packet size. Second, underlying time variations in link delivery ratios and throughputs make it hard for PTC to make accurate predictions.

For each link, the graph shows the link's broadcast delivery ratio at sometime versus its broadcast delivery ratio 50 seconds earlier. The left graphs show measurements of 134-byte packets; the right graphs show 1,386-byte packets. The top row shows measurements from the daytime, the bottom row shows measurements from the nighttime.

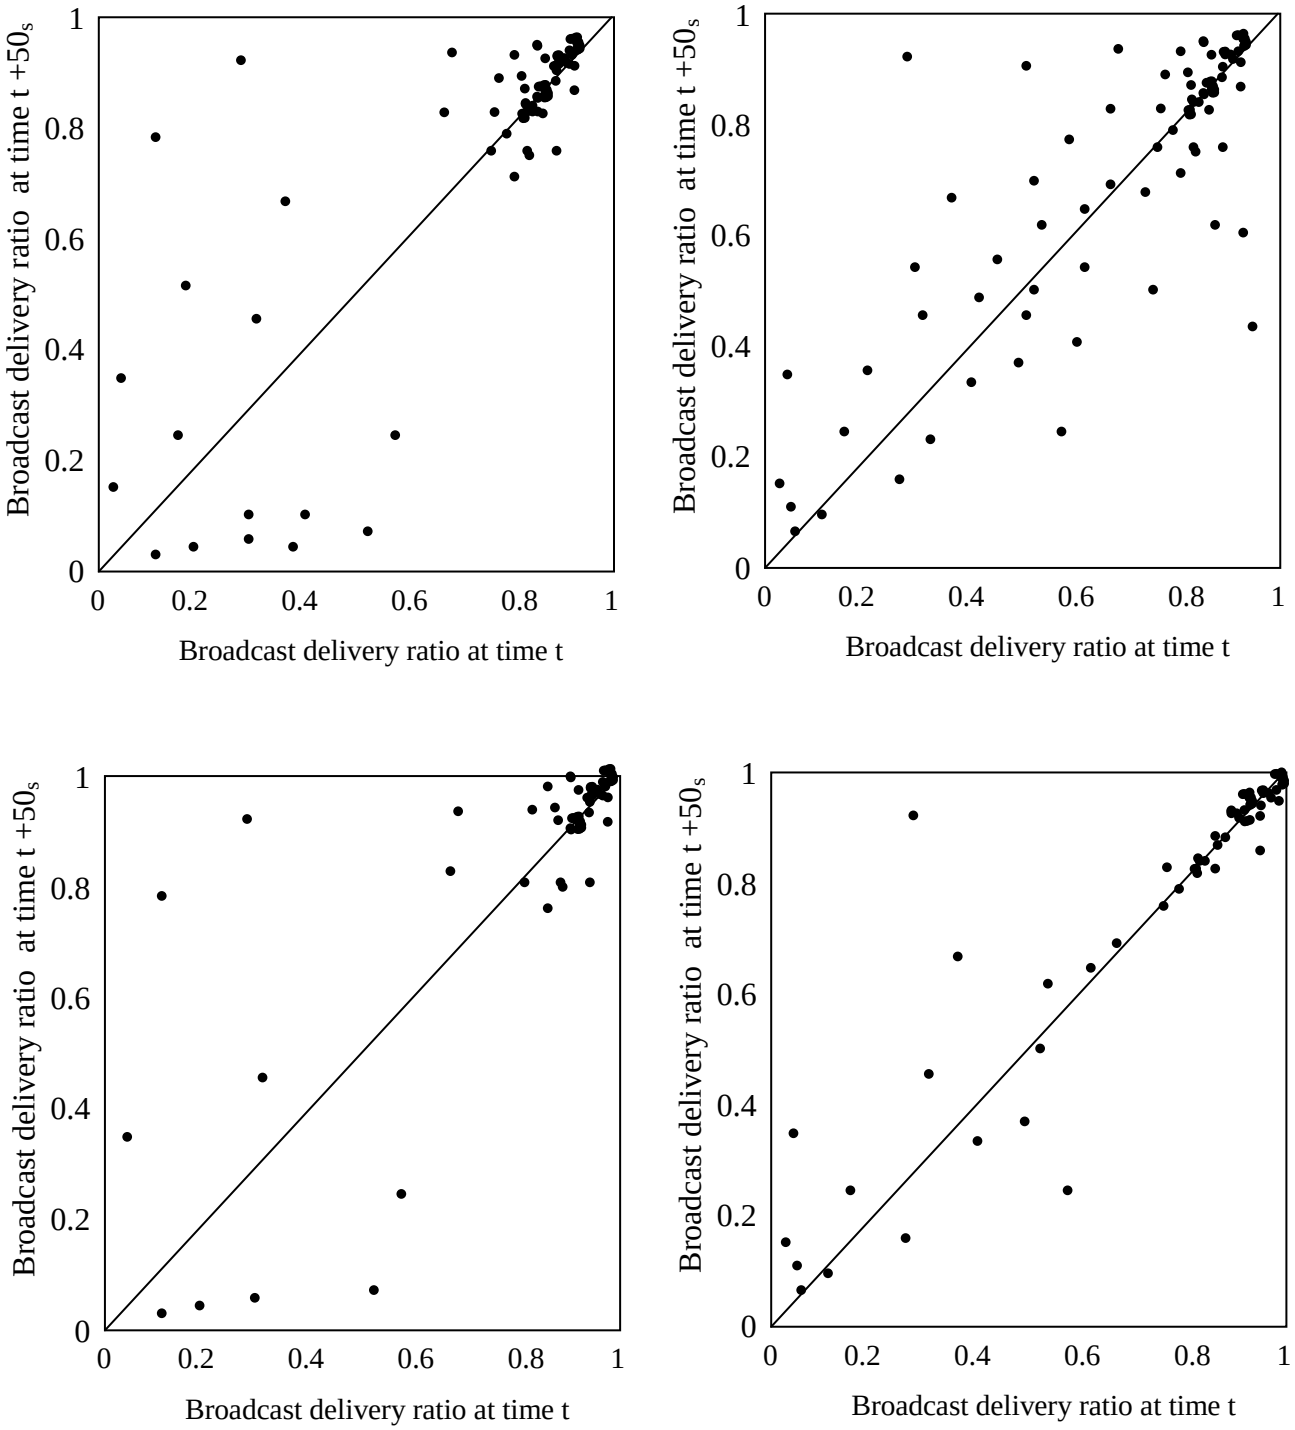


Figure 8-15: The broadcast delivery ratio of a link can change significantly overtime.

# Chapter 9: Future Directions

## 9.1 Introduction and PTC Improvements

This chapter outlines possible future work that could expand upon, extend the results in previous chapters and also describes how PTC fits into the larger design space of multi-hop wireless networks.

This section describes how to improve the performance of PTC and increase its applicability. Many PTC design constraints were chosen to simplify the PTC implementation and design, but can be relaxed with a corresponding increase in complexity.

PTC assumes that all packets are the same size, but the PTC predictions can be adjusted for packet size. One approach would be to use the packet size model in Chapter 4 to correct for ACK and data packet sizes using measurements at only two sizes. Another approach would be to directly measure delivery ratios at all relevant packet sizes, estimating the ACK delivery ratio as the delivery ratio of a minimum-size data packet. In practice this would only involve measuring at two or three packet sizes, since the distribution of data packet sizes has only a few distinct peaks [17]. A related problem is how to modify the routing protocol to handle multiple packet sizes. Because the PTC of each link can vary differently with packet size, each packet size should have its own minimum-PTC route. This can be done simply by including the intended data packet size in each route entry or route control packet (e.g. DSR route requests and replies), at the expense of multiplying the routing state by the number of data packet sizes.

PTC would also benefit from taking into account multiple bit-rates. PTC should be able to trade off a lossy link with high bit-rate for a low-loss, low bit-rate link. This can be done by combining the Medium Time Metric [6] with PTC to produce the *potential transmission time* (PTT) metric [3]. PTC calculations are simplified by using transmission counts as a proxy for the time each data packet keeps the radio medium busy. PTT makes time explicit, which allows it to combine the PTTs from links with different bit-rates into a single route PTT. PTT can also consider links with different amounts of per-packet overhead.

Many radios can change their bit-rates, including 802.11 radios. This allows links to be made more reliable by decreasing the bit-rate. Nodes can locally decide on the optimal bit-rate for each link by measuring the link at every bit-rate, and choosing the bit-rate which results in

the lowest PTT for that link [3]. Another approach would be to treat each physical link as multiple virtual links, one for each bit-rate. The PTT for each virtual link is determined by measuring the link at the specified bit-rate, and the routing protocol finds the best route using the virtual links.

One flaw in the current PTC implementation is that it is highly sensitive to load because of the effects of unfairness and interference on the probe broadcasts. Although a load-sensitive metric might be useful in some applications, PTC is intended to reflect the underlying quality of a route, independent of network traffic. Using a MAC protocol that supports priority traffic might isolate the PTC probe measurements from heavy data traffic, since the data traffic wouldn't prevent PTC probes from being transmitted. Another approach to maintaining accurate transmission count measurements in a busy network would be to use per-packet transmission feedback from the 802.11 interface, which provides a direct measurement of the number of times each data packet is transmitted over each link. Tracking these per-packet data transmissions would incur no extra overhead for routes that are already carrying data traffic.

Because the current implementation of PTC is load sensitive, routing protocols using PTC may oscillate between routes. This can be reduced with appropriate damping [9]. Another approach might be to use multipath routing, choosing paths with the best PTC metrics. Each path will be less loaded than in the single path case, and will have similar load-sensitive PTC effects, reducing the metric discrepancies which cause the routing protocol to switch routes.

A final problem is how routing protocols should handle route metrics that change over time. The DSDV protocol continuously searches for new routes with better metrics, but has no sense of hysteresis or uncertainty: slight changes in the PTC metric can cause DSDV to switch routes, which can negatively impact TCP traffic due to reordering. On the other hand, DSR doesn't switch routes unless there are too many retransmissions for a single packet, allowing DSR to keep using a poor route for too long. Routing protocols should notice and react appropriately to changing metrics, without excessive flapping and overhead.

## 9.2 Wireless Routing and PTC

This section describes larger problems in the area of wireless routing which are related to PTC.

This work did not address mobility, which is important for many wireless networks. One unanswered question is whether or not PTC is effective in a mobile or even partly mobile

network. In these networks it is a challenge to distribute accurate topology information and find routes; adding the problem of distributing accurate and consistent route quality metrics like PTC is a further challenge. PTC is likely to be helpful in mobile networks if the routing protocol can obtain accurate metrics and propagate them in a timely manner.

Another problem is how routing metrics and protocols should model the time variation of links and routes. This is a fundamental problem: route metrics must predict the future performance of a route based on past measurements. The PTC design assumes that the loss probability of each link is constant, but this is not realistic. As Chapter 8 showed the underlying performance of routes and link scan change between when a metric is calculated and when a route is used, causing protocols to choose the wrong routes. Protocols might be able to use models of how links change over time to make better predictions. These models could help reduce PTC overhead, for example, by identifying links that don't change much over time, and which require fewer probes to characterize. Also, modeling the uncertainty bounds or time variation of each metric may be useful for choosing routes when the variation of a route's performance is also important [45], as with streaming multimedia traffic, which prefers routes with bounded delay.

This work looked at how PTC improved the performance of sending a single flow at a time. However, most networks have multiple flows of traffic. It may be the case that the minimum-PTC route will not provide the highest throughput to a flow when there is cross-traffic in the network. However, the individual link PTC metrics may be useful as input to higher-level traffic engineering algorithms, such as the work by Jain et al. [36].

An important assumption made by PTC is that each network node has a single radio and antenna, and that the radio uses link-level retries, as in the 802.11specification. There is a great opportunity to improve network performance by relaxing these assumptions. For example, equipping each node with multiple radios on different bands and allowing them to transmit and receive simultaneously would significantly reduce or eliminate inter-hop interference, if neighboring link sare assigned to non-interfering channels [2]. With enough independent channels, the throughput of a route can be made independent of the number of hops in the route, and throughput is likely determined by the bottleneck bandwidth of each link, rather than it's PTC. When there is cross-traffic in a network where each link can operate on multiple independent channels, channel assignments for each link should consider both the ordering of

link channels within a route, and the interactions between links from different routes which share one or more nodes.

Even though PTC does not predict throughput in networks with multiple flows or multiple channels and radios, it is still useful for increasing total system capacity in busy networks. The capacity of a wireless network with fixed transmission power is determined by its area: the number of transmissions per time per unit area is constant. By choosing minimum-PTC routes for each flow, routing protocols are maximizing the number of packets each flow can send per second. It is not clear; however, what are the fairness properties of PTC in this situation. Also, in networks with many idle regions, total network may be improved by routing some flows out of the way through these idle regions, rather than having them share the throughput of a busy region.

Some wireless networks use end-to-end retransmissions instead of link-level retransmissions. The PTC of a route can be calculated for these networks, albeit in a much more complicated form than Equation 6.3. The exact form of the equation depends on the ordering of the links within each route. PTC will not predict throughput for these networks, because unlike networks with link-layer retransmissions, significant amounts of time can be spent waiting for end-to-end acknowledgments to timeout. However, by incorporating timeout information, the end-to-end PTC can be converted into an end-to-end PTT metric which does predict throughput. And, like networks with multiple radios, PTC is still useful for increasing the total network throughput.

PTC is also useful in networks with variable transmission power. The transmission power at each node is generally decreased as much as possible to reduce energy consumption and increase network capacity, while still keeping the network connected [31, 71]. For any given allocation of transmit powers to nodes, there will likely to be many lossy links, and PTC will be helpful for identifying good routes. And since most radios will have a fixed maximum transmit power, there will always be some links that have marginal performance.

Using more sophisticated radio technology may reduce some of the underlying link loss problems that routing protocols can detect using PTC. Improvements in radio coding and modulation, and multi-user access can improve packet delivery ratios over a link. For example, the orthogonal frequency-division multiplexing (OFDM) [16] technique adaptively chooses sets of frequencies to avoid noise and interference, thereby decreasing packet error rates. However,

this incurs a throughput cost because OFDM is avoiding some fraction of the available radio spectrum. Each link has an intrinsic throughput related to its individual RF characteristics and interference, which cannot be exceeded. Increasingly advanced modulation or coding techniques will only allow radio systems to obtain a larger fraction of the fundamental throughput of each link. So, there will always be an opportunity to use a metric like PTC to find high-throughput routes using the best links in the network.

# Chapter 10: Conclusion

The main contribution of this work is a simple way for multi-hop wireless routing protocols to choose high-throughput paths in networks with link-layer retransmissions. By measuring the delivery ratios of each link in a route using fixed size broadcast packets, protocols can estimate the throughput of the route as the inverse of that route's expected transmission count, which is called PTC. The PTC of a route *R* is calculated as

(10.1)

Where and are the measured delivery ratios in the forward and reversedirections of each link *i* in *R*.

The inverse of PTC predicts throughput for routes with small to medium hop-counts. Since at most one node in those routes can transmit at any time, the throughput of the route is limited by the number of transmissions, or equivalently, time, required to transmit each packet over the route. Measurements on a real test bed network show that PTC helps the DSR and DSDV routing protocol find routes with significantly higher throughput than the default minimum hop-count metric. The overhead of the PTC delivery ratio probes depends on the spatial density of the network, and is relatively small compared to the amount of data traffic that can be sent over each link.

This work also characterized the delivery ratios and asymmetry of the test bed network, and showed how lossy and asymmetric links affect route throughput. Lossy links require more retransmissions, and therefore have lower effective throughput. However, a route with few lossy links can be preferable to a route with many higher-quality links, since contention between links also reduces route throughput.

Finally, this work proposed a simple model for how link delivery ratios vary with packet size. The packet delivery $P_p$ for a packet with $n$ data symbols is

$$P_p(n) = P_f \times \qquad (10.2)$$

Where $P_f$ is the per-packet probability that a receiver successfully acquires and synchronizes to a packet frame, and $P_s$ is the per-symbol probability that the receiver successfully decodes that symbol. Measurements on the test show that this model can accurately predict the delivery ratios at many packet sizes using measurements at two packet sizes over each link.

# Bibliography


**[1]** C.-H. Lee and M. Haenggi, "Delay Analysis of Spatio-Temporal Channel Access for Cognitive Networks," in *IEEE International Conference on Communications (ICC'11)*, (Kyoto, Japan), June 2011.

**[2]** X. Zhang and M. Haenggi, "A Location-Based MAC Scheme for Random Networks," in *IEEE International Conference on Communications (ICC'11)*, (Kyoto, Japan), June 2012.

**[3]** Z. Gong and M. Haenggi, "Temporal Correlation of the Interference in Mobile Random Networks," in *IEEE International Conference on Communications (ICC'11)*, (Kyoto, Japan), June 2011.

**[4]** A. Sarkar and M. Haenggi, "Percolation in the Secrecy Graph," in *2012 Information Theory and Applications Workshop (ITA'11)*, (San Diego, CA), Feb. 2012.

**[5]** P. Vizi, S. Vanka, S. Srinivasa, M. Haenggi, and Z. Gong, "Scheduling using Superposition Coding: Design and Software Radio Implementation," in *IEEE Radio Wireless Week*, (Phoenix, AZ), Jan. 2011.

**[6]** Z. Gong and M. Haenggi, "Mobility and Fading: Two Sides of the Same Coin," in *IEEE Global Communications Conference (GLOBECOM'10)*, (Miami, FL), Dec. 2010.

**[7]** S. Vanka and M. Haenggi, "Coordinated Packet Transmission in Random Wireless Networks," in *IEEE Global Communications Conference (GLOBECOM'10)*, (Miami, FL), Dec. 2010.

**[8]** K. Stamatiou and M. Haenggi, "Optimal Spatial Reuse in Poisson Multi-hop Networks," in *IEEE Global Communications Conference (GLOBECOM'10)*, (Miami, FL), Dec. 2010.

**[9]** A. Sarkar and M. Haenggi, "Secrecy Coverage," in *44th Asilomar Conference on Signals, Systems, and Computers (Asilomar'10)*, (Pacific Grove, CA), Nov. 2010.

**[10]** S. Srinivasa and M. Haenggi, "Throughput-Delay-Reliability Tradeoffs in Multihop Networks with Random Access," in *2010 Allerton Conference on Communication, Control and Computing*, (Monticello, IL), Sept. 2010.

**[11]** M. Haenggi, "Local Delay in Poisson Networks with and without Interference," in *2010 Allerton Conference on Communication, Control and Computing*, (Monticello, IL), Sept. 2010.

**[12]** M. Haenggi, "Delay-Based Connectivity of Wireless Networks," in *International Symposium on Mathematical Theory of Networks and Systems (MTNS'10)*, (Budapest, Hungary), Aug. 2010.

**[13]** S. Srinivasa and M. Haenggi, "The TASEP: A Statistical Mechanics Tool to Study the Performance of Wireless Line Networks," in *2010 International Conference on Computer Communication Networks (ICCCN'10)*, (Zurich, Switzerland), Aug. 2010.

**[14]** C.-H. Lee and M. Haenggi, "Interference and Outage in Doubly Poisson Cognitive Networks," in *2010 International Conference on Computer Communication Networks (ICCCN'10)*, (Zurich, Switzerland), Aug. 2010.

**[15]** K. Stamatiou and M. Haenggi, "The Delay-optimal Number of Hops in Poisson Multi-hop Networks," in *IEEE Symposium on Information Theory (ISIT'10)*, (Austin, TX), June 2010.

**[16]** S. Vanka and M. Haenggi, "Analysis of the Benefits of Superposition Coding in Random Wireless Networks," in *2010 IEEE International Symposium on Information Theory (ISIT'10)*, (Austin, TX), June 2010.

**[17]** R. K. Ganti, Z. Gong, M. Haenggi, C.-H. Lee, S. Srinivasa, D. Tisza, S. Vanka, and P. Vizi, "Implementation and Experimental Results of Superposition Coding on Software Radio," in *2010 IEEE International Conference on Communications (ICC'10)*, (Cape Town, South Africa), May 2010.

**[18]** Brian H. Davies and T. R. Davies. The application of packet switching techniques to combat net radio. In *Proceedings of the IEEE* [1]. Special Issue on Packet Radio Networks.

**[19]** IEEE Computer Society LAN/MAN Standards Committee. *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: High-speed Physical Layer in the 5 GHz Band*. The Institute of Electrical and Electronic Engineers, New York, New York, 1999. IEEE Std. 802.11a–1999.

**[20]** IEEE Computer Society LAN/MAN Standards Committee. *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Higher-Speed Physical Layer Extension in the 2.4 GHz Band*. The Institute of Electrical and Electronic Engineers, New York, New York, 1999. IEEE Std. 802.11b–1999.

**[21]** IEEE Computer Society LAN/MAN Standards Committee. *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Amendment 4: Further Higher Data Rate Extension in the 2.4 GHz Band*. The Institute of Electrical and Electronic Engineers, New York, New York, 2003. IEEE Std. 802.11g–2003.

**[22]** D. M. J. Devasirvatham, M. J. Krain, and D. A. Rappaport. Radio propagation measurements at 850 MHz, 1.7 GHz and 4 GHz inside two dissimilar office buildings. *Electronics Letters*, 26(7):445–447, March 1990.

**[23]** SheetalkumarDoshi, ShwetaBhandare, and Timothy X. Brown. An ondemand minimum energy routing protocol for a wireless ad hoc network. *Mobile Computing and Communications Review*, 6(2), July 2002.

**[24]** CMU/Rice Monarch DSR Source Code Distribution. http://monarch.cs.rice.edu/dsrimpl.html.

**[25]** RohitDube, Cynthia D. Rais, Kuang-YehWang, and Satish K. Tripathi. Signal stability-based adaptive routing (SSA) for ad hoc mobile networks. *IEEE Personal Communications*, February 1997.

**[26]** William C. Fifer and Frederick J. Bruno. The low-cost packet radio. In *Proceedings of the IEEE* [1]. Special Issue on Packet Radio Networks.

**[27]** Dan Duchamp and Neil F. Reynolds. Measured performance of a wireless LAN. In *Proceedings of the 17th Annual Conference on Local Computer Networks*, September 1992.

**[28]** M. S. Frankel. Advanced technology testbeds for distributed, survivable command, control, and communications. In *Proceedings of the IEEE MILCOM Conference*, volume 3, October 1982.

**[29]** Tom Goff, Nael B. Abu-Ghazaleh, Dhananjay S. Phatak, and RidvanKahvecioglu. Preemptive routing in ad hoc networks. In *Proc. ACM/IEEE Mobi- Com*, July 2001.

**[30]** Piyush Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions of Information Theory*, 46(2):388–404, March 2000.

**[31]** Wendi RabinerHeinzelman, AnanthaChandrakasan, and HariBalakrishnan. Energy-efficient Communication Protocols for Wireless Microsensor Networks. In *Proceedings of the Hawaiian International Conference on Systems Science*, January 2000.

**[32]** Yu-Ching Hsu, Tzu-Chieh Tsai, Ying-Dar Lin, and Mario Gerla. Bandwidth routing in multi-hop packet radio environment. In *Proceedings of the 3rd International Mobile Computing Workshop*, 1997.

**[33]** Eun-Sun Jung and NitinVaidya. A power control MAC protocol for ad hoc networks. In *Proc. ACM/IEEE MobiCom*, September 2002.

**[34]** Yih-Chun Hu and David B. Johnson. Design and demonstration of live audio and video over multihop wireless ad hoc networks. In *Proceedings of the MILCOM*, 2002.

**[35]** *HFA3863 Direct Sequence Spread Spectrum Baseband Processor with Rake Receiver and Equalizer Data Sheet*. Intersil Americas Inc., December 2001.

**[36]** Kamal Jain, JitendraPadhye, VenkatPadmanabhan, and LiliQiu. Impact of interference on multi-hop wireless network performance. In *Proc. ACM/IEEE MobiCom*, September 2003.

**[37]** David B. Johnson. Routing in ad hoc networks of mobile hosts. In *Proc. Of the IEEE Workshop on Mobile Computing Systems and Applications*, pages 158–163, December 1994.

**[38]** David B. Johnson, David A. Maltz, and Yih-Chun Hu. The Dynamic Source Routing protocol for mobile ad hoc networks (DSR). Internet draft (work in progress), IETF, April 2003. http://www.ietf.org/internet-drafts/ draft-ietf-manet-dsr-09.txt.

**[39]** John Jubin and Janet D. Tornow. The DARPA packet radio network protocols. In *Proceedings of the IEEE* [1]. Special Issue on Packet Radio Networks.

**[40]** Chenxi Zhu and M. Scott Corson. QoS routing for mobile ad hoc networks. In *Proc. IEEE Infocom*, June 2001.

**[41]** QifaKe, David A. Maltz, and David B. Johnson. Emulation of multi-hop wireless ad hoc networks. In *Proceedings of the Seventh InternationalWorkshop on Mobile Multimedia Communications (MOMUC 2000), IEEE Communications Society*, October 2000.

**[42]** AtulKhanna and John Zinky. The revised ARPANET routing metric. *ACM SIGCOMM Computer Communication Review*, 19(4), September 1989.

**[43]** Ralf Koetter and Muriel Medard. An algebraic approach to network coding. *IEEE/ACM Transactions on Networking*, 11(5):782–795, October 2003.

**[44]** Eddie Kohler, Robert Morris, Benjie Chen, John Jannotti, and M. FransKaashoek. The Click modular router. *ACM Transactions on Computer Systems*, 18(4), November 2000.

**[45]** C. EmreKoksal and HariBalakrishnan. Quality-aware routing in timevarying wireless networks. In preparation, 2004.

**[46]** David Kotz, Calvin Newport, and Chip Elliott. The mistaken axioms of wireless-network research. Technical Report TR2003-467, Dept. of Computer Science, Dartmouth College, July 2003.

**[47]** Jinyang Li, Charles Blake, Douglas S. J. De Couto, Hu Imm Lee, and Robert Morris. Capacity of ad hoc wireless networks. In *Proceedings of the 7th ACM International Conference on Mobile Computing and Networking*, pages 61–69, Rome, Italy, July 2001.

**[48]** Shuo-Yen Robert Li, Raymond W. Yeung, and NingCai. Linear network coding. *IEEE Transactions on Information Theory*, 49(2):371–381, February 2003.

**[49]** Chunhung Richard Lin. On-demand QoS routing in multihop mobile networks. In *Proc. IEEE Infocom*, April 2001.

**[50]** Henrik Lundgren, Erik Nordstr¨om, and Christian Tschudin. Coping with comunication gray zones in IEEE 802.11b based ad hoc networks. In *5th ACM international Workshop on Wireless Mobile Multimedia (WoWMoM 2002)*, September 2002.

**[51]** David A. Maltz, Josh Broch, and David B. Johnson. Quantitative lessons from a full-scale multi-hop wireless ad hoc network testbed. In *Proceedings of the IEEE Wireless Communications and Networking Conference*, September 2000.

**[52]** James L. Massey. Optimum frame synchronzation. *IEEE Transactions on Communications*, COM-20(2), April 1972.

**[53]** AnastassiosMichail and Anthony Ephremides. Algorithms for routing session traffic in wireless ad-hoc networks with energy and bandwidth limitations. In *Proceedings of 12th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, 2001.

**[54]** EytanModiano. An adaptive algorithm for optimizing the packet size used in wireless ARQ protocols. *Wireless Networks*, 5(4):279–286, July 1999.

**[55]** ThyagarajanNandagopal, Tae-Eun Kim, Xia Gao, and VaduvurBharghavan. Achieving MAC layer fairness in wireless packet networks. In *Proc. ACM/IEEE MobiCom*, pages 87–98, August 2000.

**[56]** Giao T. Nguyen, Randy H. Katz, Brian Noble, and MahadevSatyanarayanan. A trace-based approach for modeling wireless channel behavior. In *Proceedings Winter Simulation Conference '96*, December 1996.

**[57]** P. Nobles, D. Ashworth, and F. Halsall. Indoor radiowave propagation measurements at frequencies up to 20 GHz. In *Proceedings 44th IEEE Vehicular Technology Conference*, June 1994.

**[58]** P. Nobles and F. Halsall. Delay spread and received power measurements within a building at 2 GHz, 5 Hz and 17 GHz. In *Proceedings 10th International Conference on Antennas and Propagation*, April 1997.

**[59]** The Network Simulator—ns-2, 2003. http://www.isi.edu/nsnam/ns.

**[60]** Christina Parsa and J. J. Garcia-Luna-Aceves. TULIP: A link-level protocol for improving TCP over wireless links. In *Proc. IEEE Wireless Communications and Networking Conference 1999 (WCNC 99)*, September 1999.

**[61]** Charles E. Perkins and PravinBhagwat. Highly dynamic Destination- Sequenced Distance-Vector routing (DSDV) for mobile computers. In *Proc. ACM SIGCOMM Conference (SIGCOMM '94)*, pages 234–244, August 1993.

**[62]** Charles E. Perkins and Elizabeth M. Royer. Ad hoc On-Demand Distance Vector Routing. In *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications*, 1999.

**[63]** Raymond L. Pickholtz, Donald L. Schilling, and Laurence B. Milstein. Theory of spread spectrum communications—A tutorial. *IEEE Transactions on Communications*, 20(5):855–884, May 1982.

**[64]** R. Price and P. E. Green. A communication technique for multipath channels. *Proceedings of the IRE*, 46:555–570, March 1958.

**[65]** John G. Proakis. *Digital Communications*. McGraw-Hill, 2001.

**[66]** Rice Monarch Project. Wireless and mobility extensions to ns-2. http:// www.monarch.cs.rice.edu/cmu-ns.html.

**[67]** Ratish J. Punnoose, Pavel V. Nitkin, Josh Broch, and Daniel D. Stancil. Optimizing wireless network protocols using real-time predictive propagation modeling. In *Radio and Wireless Conference (RAWCON)*, August 1999.

**[68]** Ram Ramanathan and Regina Rosales-Hain. Topology control of multihop wireless networks using transmit power adjustment. In *Proc. IEEE Infocom*, March 2000.

**[69]** Theodore S. Rappaport. *Wireless Communications: Principles and Practice*. Prentice Hall, Upper Saddle River, New Jersey, 1996.

**[70]** Samarth H. Shah and KlaraNahrstedt. Predictive location-based QoS routing in mobile ad hoc networks. In *Proceedings of IEEE International Conference on Communications*, 2002.

**[71]** Timothy Shepard. A channel access scheme for large dense packet radio networks. In *Proc. ACM SIGCOMM Conference (SIGCOMM '96)*, pages 219–230, August 1996.

**[72]** PrasumSinha, RaghupathySivakumar, and VaduvurBharghavan. CEDAR: A core-extraction distributed ad hoc routing algorithm. In *Proc. IEEE Infocom*, March 1999.

**[73]** Bernard Sklar. *Digital Communications: Fundamentals and Applications*. Prentice Hall, Upper Saddle River, New Jersey, second edition, 2000.

**[74]** James Almon Stevens. *Spatial Reuse Through Dynamic Power and Routing Control in Common-Channel Random-Access Packet Radio Networks*. PhD thesis, The University of Texas at Dallas, 1988.

**[75]** AudunTornquist. Modular and adaptive ad hoc routing in Click. Master's thesis, University of Colorado, 2001.

**[76]** Rudi van Drunen, Jasper Koolhaas, HuubSchuurmans, and Marten Vijn. Building a wireless community network in the Netherlands. In *USENIX*/*Freenix Conference*, June 2003.

**[77]** Andreas Willig, Martin Kubisch, Christian Hoene, and Adam Wolisz. Measurements of a wireless link in an industrial environment using an IEEE 802.11-compliant physical layer. *IEEE Transactions on Industrial Electronics*, 49(6), December 2003.

**[78]** Alec Woo, Terence Tong, and David Culler. Taming the underlying challenges of reliable multihop routing in sensor networks. In *Proc. ACMSensys*, November 2003.

**[79]** Mark Yarvis, W. Steven Conner, Lakshman Krishnamurthy, JasmeetChhabra, Brent Elliott, and Alan Mainwaring. Real-world experiences with an interactive ad hoc sensor network. In *Proceedings of the International Workshop on Ad Hoc Networking*, August 2002.

**[80]** Jerry Zhao and Ramesh Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *Proc. ACM Sensys*, November 2003.

## Author Biography

Arun Kumar Singh was born in Azamgarh, UP, India in 1978. He attended the Government Inter College in Allahabad for the high school and intermediate certification.

He then studied at the Dr. B. R. Ambedkar. University, Agra-UP (Formerly Agra University, Agra) where in 2002 he received his B.E. in Electronics and Communication Engineering, and M.Tech.in Information Technology (WCC-Wireless Communication and Computing) from Indian Institute of Information Technology (IIIT-Allahabad), Allahabad in 2005.

While at IIIT–Allahabad he worked on a variety of computer/Information/Artificial/Architecture systems projects, in particular implementing and studying ad hoc and multi-hop wireless networks. He also worked on dynamic MANET scheduling for mobile computing, and applications of the Global Positioning System (GPS) for enhancing throughput. He also worked as a coordinator in CSE/IT Dept. in Shobhit University, Meerut for 4 Year. Currently he is working as a head of the information Technology in DIT School of Engineering, Greater Noida, UP (Affiliated to Gautam Budha Technical University, Lucknow, UP.)

# List of Publications/Conferences (2008-2013)

| S. No. | Title of research Paper | Name of Journals/Publications | Volume No./ Page No. | ISBN/ISSN | Paper Status |
|---|---|---|---|---|---|
| **1** | A Novel Strategy for High Throughput in Ad Hoc Networks using Potential transmission Count (PTC) Metric | WSEAS (The World Scientific and Engineering Academy and Society, Canada) | Issue 8, Volume 10/Page No.223-232 | ISSN: 1109-2742 | Published |
| **2** | Experimentally Modified Protocols with Transmission Count Metric for Efficient Throughput in Ad-Hoc Networks | UBICC (Ubiquitous Computing and Communication Journal), USA | Volume 6 No.2/ Page No.807-813 | ISSN: 1992-8424 | Published |
| **3** | Throughput Increments Phenomena in Ad Hoc Network using Potential transmission Count (PTC) Metric with Protocols | IEEE Explore ICCCT-2011 MNNIT Allahabad | 2nd IEEE International/Page No. 562-567 | 978-1-4577-1386-611$26.00©2011 IEEE | Published |
| **4** | A WiMAX Segmentation for 3rd Generation N/W Simulator | IEEE Explore ICCCT-2010 MNNIT Allahabad | Paper 257 Session 2/Page No. 135-140 | 978-1-4244-9032-5/10/$26.00© 2010 IEEE | Published |
| **5** | A Count Metric Strategy with Protocols to Increase Throughput in Ad Hoc Networks | IEEE Microwave Theory and Technique Society India chapter/ NECE-2010 | MITS-Gwalior Page No.-131 | NVIS Technologies Pvt. Ltd. IEEE India council | Published |
| **6** | An Effective Technique of Routing for MANETs | IEEE Microwave Theory and Technique Society India chapter/ NECE-2010 | MITS-Gwalior Page No.-106 | NVIS Technologies Pvt. Ltd. IEEE India council | Published |
| **7** | A NOVEL COMPARISION AND SIMULATION OF ROUTING PROTOCOLS IN Ad-hoc NETWORKS | UBICC, IJARCS, INNS International Conference on Issues and Challenges in Networking, Intelligence and Computing Technologies (ICNICT' 2011) | Page No.-683-686 | ISSN: 978-93-81126-27-1 | Published |