

Studying Wireless Routing Link Metric Dynamics

Saumitra M. Das¹, Himabindu Pucha¹, Konstantina Papagiannaki², Y. Charlie Hu¹

¹School of ECE, Purdue University and ²Intel Research, Pittsburgh

ABSTRACT

Multi-hop wireless mesh networks are increasingly being deployed for last-mile Internet access. Typically, network algorithms such as routing, channel assignment and topology control for such networks rely heavily on metrics that intend to capture link "quality" across the network. However, the underlying dynamics of the proposed link metrics themselves have not yet been studied in detail. In this paper, we study the dynamics of the most popular link metrics in real network deployments. Using two wireless mesh testbeds, we measure a number of link metrics across different hardware platforms and network environments. The collected measurements allow us to study the stability and sensitivity of the different metrics to various conditions. Our study provides several insights and future research directions on how network algorithms need to adapt to link dynamics as well as how popular and widely used link metrics can be improved.

Categories and Subject Descriptors : C.2.3 [Computer Communication Networks]: Network Operations

General Terms : Measurement, Performance

Keywords : Wireless networks, link metrics, link variation

1. INTRODUCTION

Multi-hop wireless mesh networks are composed of static nodes equipped with one or more radios that use each other to obtain network connectivity (through multi-hopping). Such networks have a variety of envisioned applications; the most important being the last-mile extension of the Internet in rural, underprivileged or under-provisioned neighborhoods.

Challenges that routing and management protocols face in such networks stem from the vagaries of the wireless channel. This leads to the need to adapt and account for link dynamics due to fading, interference, obstacles etc. For example, distributed channel assignment algorithms typically periodically re-evaluate their assignment to deal with link dynamics. Routing protocols periodically probe the links to determine appropriate routes across the network. Networks with electronically steerable directional antennas may re-

quire this information to reconfigure the network topology.

Due to the complexity and inaccuracy of propagation and interference models, a link's performance is in practice typically characterized via some periodically measured link metric such as ETT (expected transmission time), ETX (expected transmission count), WCETT (weighted cumulative ETT), RTT (routing trip time), loss rate etc. While a lot of effort has been put into the design of link metrics; there has been little focus on the dynamics of these metrics, i.e. their stability and sensitivity. As a consequence, it is still not known how upper layer protocols should adapt to the link dynamics. In addition, it is not known how sensitive these metrics are to various perturbation triggers such as interference, obstacles and traffic.

In this paper, we study link metric dynamics by measuring the stability and sensitivity of link metrics in two wireless mesh deployments, at Purdue University and Microsoft Research. We measure various link metrics across different hardware, operating frequencies and physical spaces over several days. We also study the variability of different link metrics in real deployments such as how much and how often variation can occur, the timescale of such variation, correlation with time-of-day, whether the variation is similar across different hardware devices and physical layer standards. We also study the sensitivity of different link metrics to variation, i.e. whether some link metrics are more stable than others as well as whether and by how much the link metrics vary due to existing traffic. Our study has practical significance since it uses two widely used and mature protocol implementations.

The main findings of this study are: (1) Link quality can vary considerably over time and while bad quality links vary more, a small number of excellent quality links also experience significant variation. (2) Perceived link quality depends on how protocols implement link metric probing and thus there is a need to correlate signal level changes to observed link metric behavior to design better probing mechanisms. (3) Typically links do not change in quality every few seconds, so it may be advantageous to perform higher overhead, more accurate link quality probing at longer time scales. (4) Some link metrics exhibit higher variation than others and thus the choice of link metric will affect the upper layer protocol behavior. (5) Some channels exhibit higher variation than others on a given link and this varies on a per link basis. This behavior creates restrictions on what channel can be used on which link and motivates the need for more intelligent upper layer protocols to optimize performance. (6) Most importantly, we find that link metrics as defined and used in today's mesh networks are negatively impacted by background traffic. While some metrics are caused to inflate to meaningless values, other metrics do not respond at all. Thus, route selection *in the presence of multiple flows* using the state-of-the-art protocols operates on essentially random values. Our results indicate that there is a definite need for a redesign of routing metrics to

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IMC'07, October 24-26, 2007, San Diego, California, USA.
Copyright 2007 ACM 978-1-59593-908-1/07/0010 ...\$5.00.

make them either agnostic of other flows or able to correctly incorporate other flows' effect on the desirable property. The conclusion of the paper lists various future research directions based on the reported findings. Our study motivates that research needs to focus on link metric dynamics rather than just the design of the metrics themselves.

2. BACKGROUND

We begin with a quick review of link metrics used in wireless mesh networks. Link metrics typically involve the measurement of a particular quantity for the link between a pair of nodes. Traditional metrics like hop-distance result in the selection of high-loss paths by traveling long distances in each hop or cannot differentiate similar hop-count paths with drastically different packet delivery performance. Other traditional metrics such as RTT have been shown to be troubling [4] due to distortion from queuing delays. In this study, we look at some state-of-the-art metrics used in mesh networking deployments. Draves et. al. [4] provide a comparison study of these well known metrics with other traditional ones such as hop-count and RTT.

ETX (Expected Transmission Count) : The Expected Transmission Count (ETX) metric was proposed in [3] to model the expected number of transmissions required to send a unicast packet over a link, including retransmissions. To calculate ETX, each node measures the probability that a packet successfully reaches the receiver, denoted as d_f , and the probability that an ACK is successfully received by the sender, denoted as d_r . The ETX value of the link is given by $ETX = \frac{1}{d_f \times d_r}$. The routing algorithm then selects the path with the least sum of ETX values of all its constituent links. To measure d_f and d_r , each node broadcasts a probe packet every second. Each such probe contains the number of probes the node received from each of its neighbors in the previous 10 seconds. Since the 802.11 MAC layer protocol does not retransmit broadcast packets, nodes use this information to estimate the forward and reverse delivery probabilities. One of the proposed qualities [3] of this metric is that it is intuitively load-independent (hence should oscillate less).

BW (Link Bandwidth) : Measuring link bandwidth using the Packet Pair method was first proposed by Keshav in [6]. The use of a pair of successive probe packets eliminates the effect of queuing delays. To measure BW in [4], each node periodically unicasts two back-to-back probe packets, one small and one large, to each of its neighbors. Each neighbor computes the delay between the arrival of the two probes and sends this delay back to the sender. An Exponentially Weighted Moving Average (EWMA) of the delays is maintained at each node for each of its neighbors. This delay value is used to calculate the link bandwidth. This metric is intuitively load-dependent and hence should vary with offered traffic load. An advantage is that it can differentiate between low and high bandwidth links which occur frequently due to the use of heterogeneous radios or variable link quality and rate control algorithms.

ETT (Expected Transmission Time) : ETT is estimated as $ETT = ETX \times \frac{S}{B}$, where S is the size of the probe packet and B the bandwidth of the link. To calculate ETT for each link, one needs to measure both ETX and bandwidth as described above. This metric is also intuitively load-dependent. Its advantage, compared to ETX and BW, is that ETT takes into account both loss as well as the raw bandwidth of the link and hence gives a more complete picture of the link performance to upper layer protocols.

3. METHODOLOGY

We use two wireless networks in our study: (1) a 29 node two-radio 802.11b network at Purdue (Network A) and (2) a 14 node



Figure 1: Network A schematic.

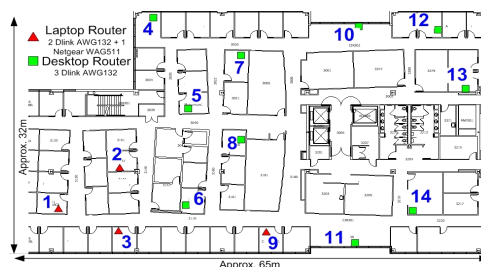


Figure 2: Network B schematic.

802.11a/b/g three-radio network at Microsoft Research, Redmond (Network B).

3.1 Network A

Network A is located on an academic campus and is illustrated in Figure 1. It consists of 29 mesh routers in 4 buildings on the Purdue campus. Each router has two radios: an Atheros 5212 based 802.11a/b/g wireless card and a Senao Prism2 based 802.11b card. Each radio is attached to a 2dBi omnidirectional antenna with a low loss pigtail to provide antenna isolation. Each mesh router runs Mandrake Linux 10.1 and the open-source *madwifi* and *hostap* drivers are used to enable the radios. The testbed deployment environment is not wireless friendly, having floor-to-ceiling office walls instead of cubicles as well as some laboratories with structures that limit the propagation of wireless signals. Potential interference exists in our deployment from other 802.11b networks (the Purdue Airlink network). The radios on each node operate on 2412 and 2462 Mhz. The testbed runs the OLSR [9] routing protocol which is a very popular open-source implementation used by several community wireless networks. OLSR implements the ETX metric. The d_f and d_r in ETX are determined by counting the reception of the periodic (every 2 seconds) HELLO messages (that form part of the OLSR protocol) over a sliding time window of 20 seconds. We collected link metrics on both radios of each node.

3.2 Network B

Network B is located on the 3rd floor of an office building and is illustrated in Figure 2. Ten of the nodes are small form factor HP desktops each with 3 DLink AWG132 802.11a/b/g USB cards while four nodes are Toshiba tablet PCs with one Netgear WAG511 802.11a/b/g PCMCIA card and 2 DLink AWG132 802.11a/b/g cards. The driver configurations are modified to allow multiple cards from the same vendor to co-exist in a machine. The radios on each node

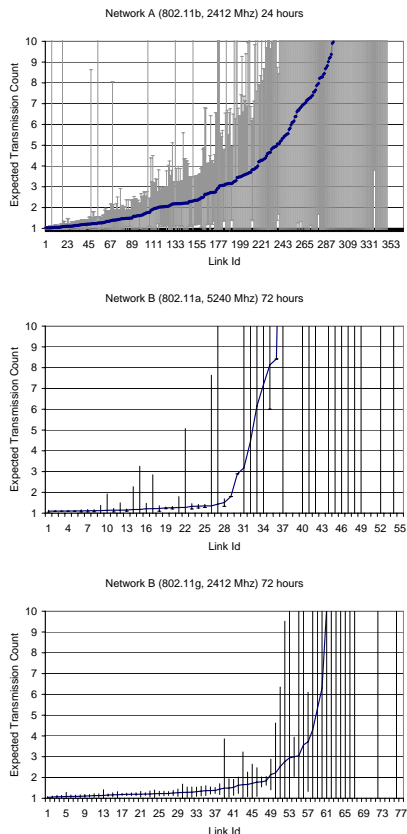


Figure 3: Average and standard deviation of the ETX metric for different links. Note that the links are ordered in decreasing order of their ETX value.

operate on 5180, 5240 and 5300 Mhz. Each radio works in ad-hoc mode and performs autorate control. The machines all run Windows XP and each machine implements TCP-SACK. The results are expected to be unaffected by external interference since no other 802.11a network is in the area. We also operate this network in the 802.11g band on channels 2412, 2437 and 2462 Mhz where an existing Microsoft corporate network operates and can potentially interfere. The testbed runs the LQSR [4] protocol with the multi radio extensions proposed in [5]. This protocol is also widely used in many studies and academic testbeds with 263 university departments having obtained the software. The protocol calculates the following link metrics during its operation: (1) ETX is calculated by sending broadcast packets every second and measuring their reception. (2) Bandwidth is measured using the packet pair technique every one minute. This is done via unicast packet exchanges with each neighbor (hence is less frequent). (3) ETT is calculated by dividing the ETX with the link bandwidth. These measurements are collected individually for each radio on each node.

4. MEASUREMENT STUDY

We now detail the results of our measurement study of link metric dynamics. In the first part of the study we study the stability of link metrics.

4.1 Link Metric Stability

4.1.1 Are there significant link metric variations?

We compute the average and standard deviation of all link metrics measured for each link in network B in 802.11g for 72 hrs, network B in 802.11a for 72 hrs and network A in 802.11b for 24 hrs. In network A, the link metric was sampled every 1 second, while in network B, the link metric was sampled every 5 seconds.

The overall link dynamics as exposed through ETX are captured in Figure 3. Note that the number of links in each configuration is different since they depend on the number of nodes in the testbed as well as the physical layer (a, b or g). The results show that ETX varies significantly in network A over a 24 hour period. Similarly, the same frequency in 802.11g in network B shows variations in ETX. Note that both networks have interfering networks operating on that frequency. Interestingly, network B in 802.11a (which has no interfering networks) also shows variability. Thus, link quality changes not related to interference can also be significant. Another interesting finding is that while there is correlation between link quality and link metric variability, this may not always be true and a small fraction of good quality links can also exhibit significant link variation. Additionally, the protocol implementation has a lot to do with the perceived link variability. The OLSR implementation used in network A by default uses 10 probes in 20 seconds while the LQSR implementation in network B uses 30 probes in 30 seconds to calculate ETX. In addition, LQSR also averages the past and new ETX observations for smoothing. Thus, one can argue that network A appears very dynamic because the samples in OLSR are wider apart than in LQSR and OLSR does not smooth the observations made. It is an avenue of our future work to see which is more in line with actual signal level variation so as to guide the design of hysteresis in protocol implementations.

Another interesting observation from Figure 3 is the difference between 802.11b and 802.11g/a networks which use an OFDM physical layer. 802.11 variants with OFDM show a clearer distinction between good links and bad links (the knee of the curve is stronger), while 802.11b tends to have a large number of intermediate quality links as has been observed earlier. Thus, the dynamics can depend on the physical layer technology in place which in turn affects upper layer protocols. For example, schemes that exploit opportunistic receptions may work significantly better in 802.11b networks due to the variety of lossy links than 802.11g/a networks.

4.1.2 What does the time variation look like?

We now look at the time variation of one sample link from each network configuration (Figure 4). Interestingly, the results show that perturbations in the link metric typically occur at important transition periods in the day such as when people leave the premises, around 5 to 6 p.m. Random signal fades may lead to link metric variations, such as when doors open/close etc. Network A in a university does not show a marked perturbation in the morning unlike network B which is an industrial lab. Perhaps this is because a campus population tends to come in at evenly distributed times based on class timings. This suggests that network performance could be improved by learning these patterns and adapting features such as probing frequency or channel assignment accordingly.

4.1.3 What are the time scales of link dynamics?

For each link, we identify time intervals between significant metric changes (i.e. a change of +/- 20%). We then calculate the average and standard deviation of the duration between metric changes for each link. The time scales of link dynamics for ETX are shown in Figure 5.

The results show that link metrics change very frequently in network A with over 80% of the links exhibiting variation in less than 10 seconds. However, the same frequency used in network B re-




Figure 4: Time variation of an example link from each network under the ETX metric.




Figure 5: Time scales of link dynamics under the ETX metric.

sults in 80% of the links exhibiting variation on average after more than 100 seconds. This again suggests that the dynamics cannot be attributed to the frequency of operation but rather to the link probing technique that can significantly affect the *perceived* link quality and cause frequent reconfiguration and overhead in upper layer protocols. Network B is relatively stable with more than half the links only changing after more than 1000 seconds. These results also suggest the possibility that link metric probing need not be done at very small timescales and more accurate/higher overhead probing (e.g. using unicast instead of broadcast) done less frequently will better match the real link quality.

4.1.4 Are different channels different in performance?

Figure 6 shows the average ETX on each channel for every link id. We can see that in network B it is very common for the link metric to be vastly different for different channels. While we do not plot the standard deviation (for readability), this also varies widely. Additionally, even in network A, it is very common for one channel to lead to node disconnection while another leads to nearly perfect communication. This problem possibly occurs less frequently in network B since most radios are from the same manufacturer whereas network A uses two different cards on each machine.

Thus, certain frequencies may have different propagation or interference characteristics from others depending on the environment. Thus, in multi-radio mesh networks certain frequencies should be preferred over others on a given link. In addition, this shows that certain radios on a node may perform badly for a given link because of features such as antenna placement. Overall, this means that using Dijkstra's algorithm with WCETT [5] for finding routes (which is a state-of-the-art practice) is likely to be suboptimal. Since different channels are bottle-necked on different links, selecting a path needs to somehow avoid these bottlenecks. For example, if channel 5180 is bad on the third hop of a route, it should be utilized before and some other channel should be used on the third hop. Similarly, channel assignment algorithms need to consider which channels are best used on which links in order to avoid such bottlenecks.

4.1.5 Are different metrics different in terms of dynamics?

Here, we measure the time between significant link metric changes observed under each metric. We only study this in network B since the network A protocol implementation uses only the ETX metric. The results for the ETX, BW and ETT metric are shown in Figure 7. The results show that in both networks, the BW metric is far more stable, only changing after more than a minimum of 1000 seconds. While the probing rate of the BW metric is lower than the other two, it is significantly shorter than 1000 seconds. The ETX metric overall causes more variation. Finally since ETT is composed of both loss rate and bandwidth it tends to follow the more dynamic ETX metric. Thus, because of the instability of the ETX metric, the widely used ETT metric also exhibits significant variation. At first this says that bandwidth on a link varies slower compared to loss rate. However, since loss rates should affect bandwidth, this seems counter-intuitive. The reason for this behavior becomes apparent in the next section which studies link sensitivity.

4.2 Link Metric Dynamics: Sensitivity

In this section, we study the sensitivity of link metrics. This provides insight on how link metrics are affected by the time-of-day and background traffic.

4.2.1 Sensitivity to time-of-day

For this experiment we look at network B in 802.11a over a 24 hour period and measure the standard deviation of the ETX link metric observed during four different time periods: at night, morning, lunchtime and at the end of the work day. While the time variation results in the previous section showed that time-of-day matters, that was just for a single link. In this section, we look at whether time-of-day matters network wide. Figure 8 clearly shows that link metrics are significantly more stable at night because there are fewer movements and changes in the environment at that time. In addition we can see that the variability in the links is maximum around lunchtime (with a lot of movement to and from the cafe-




Figure 6: Behavior of different channels under the ETX metric.




Figure 7: Comparison of dynamics between BW, ETT and ETX.

teria) and when people leave work. Additionally, note that many links are not affected by time-of-day changes and the links affected at different times are also different. This suggests that it may be worthwhile to learn these patterns in a network on a per-link basis to adapt the probing mechanisms to be more effective.

4.2.2 Sensitivity to Small Flow

We now measure the sensitivity of link metrics to background traffic. We first consider a small flow, i.e., we perform 100 pings (one per second) between two nodes in network B on channel 5240 Mhz and look at the link metric (ETT) observed before and after




Figure 8: Impact of time-of-day on ETX in Network B using 802.11a.

the perturbation across the entire network. The impact of these 100 pings on the link metrics is shown in Figure 9(a). The results show that around 10% of the links see a 200% or more increase in ETT! It's quite perplexing why the transmission count of 10% of all the links in the network would triple in number. We verified that the small flow traffic being synchronized with ETT measurement traffic was not the reason by repeating this experiment multiple times. The most likely cause of this behavior is the presence of hidden terminals in the network which causes losses in the ETT measurement packets when a small flow is initiated.

4.2.3 Sensitivity to Large Flow

We next look at what happens to link metrics if just a single mesh user downloads a file. We perform a one minute TCP transfer between two nodes in the network and look at the link metrics observed before and after the perturbation. The TCP transfer was performed between nodes 6 and 8 on the radio at 5240 Mhz. The impact of this TCP transfer on the links on all the different channels is shown in Figure 9(b).

The results show that almost 70% of the links in the network are severely perturbed by a single TCP transfer with 55% of the links showing an ETT increase greater than 300%. Even links operating on the orthogonal channel of 5300 Mhz are affected with 20% of those links (which have no traffic) seeing an ETT increase greater than 200%. The links on 5180 Mhz however remain stable. Note that the measurements on all channels were taken at exactly the same time. The full impact of this TCP transfer on the links on the affected channel (5240 Mhz) is shown to full scale in Figure 9(c). Amazingly, the link metrics show a totally disproportionate increase with 30% of the links having more than a 10000% increase in ETT! This is an important finding which illustrates that with even the most mature implementations, link metrics calculated after a single transfer are almost meaningless and subsequent route selection operates on essentially random values. We were able to recreate this behavior in network A as well as using two different implementations (OLSR and AODV-ST) which shows that this is a fundamental problem. Such variations directly impact routing. After the single hop transfer from node 6-8 was complete, we initiated another TCP transfer from node 7-6 (see Figure 2). Because the metrics on both radios on 5240 and 5300 were huge, the path chosen was 7-8-6 on channel 5180 Mhz, which gave a throughput of 6 Mbps. We monitored the path that this flow would originally have taken before we initiated the perturbation (TCP flow) and measured its throughput to be 14 Mbps. Even on single-radio networks, such large changes in metrics cause a subsequent flow to choose longer routes in order to avoid links with high metrics, thereby resulting in low throughput.

On investigating why so much variation occurs, counter to our intuition, we find that the TCP transfer does not even change the




Figure 9: Impact of traffic on ETT in different networks.

BW metric. On the other hand, the loss rates measured go up to 70% from 0% resulting in the high ETT values. This suggests that broadcast based metrics are not a good choice in a wireless mesh network due to their high sensitivity to background traffic. The BW metric is not affected since it is possible that the two unicast packets (which don't suffer from collisions due to virtual carrier sense) still go through quick enough because of 802.11 binary exponential backoff. However, unicast based probing has a problem in that it gives the false impression that no other traffic exists in the network, unless the probing mechanism or the existing traffic fully saturates the medium. This motivates the need for research on new link metrics or probing mechanisms that can correctly account for existing traffic or always remain unaffected to traffic. Finally, we were unable to find a correlation between the amount of metric change with original link quality, or particular nodes. Thus, it is not easy to find if some links are more amenable to perturbation. A lower layer signal-level study may be needed to find such correlation.

5. RELATED WORK

A large body of work on wireless mesh networks focuses on finding better link metrics [3, 2]. However, no work has specifically focused on measuring the dynamics of these metrics and its ramifications for upper layer protocol design. Aguayo et. al. [1] studied link level measurements from the RoofNet [8] 802.11b wireless network deployment. However, the work primarily focuses on the characteristics of the links in the network in terms of delivery probability, loss rate and not on the dynamics over time. The link variations depicted in the paper are over only a 90 second period and measure changes in delivery ratio with real traffic being sent and not changes in metric values. Additionally, they do not consider different metrics to describe the link behavior. Ramachandran et. al. [7] studied the routing stability of wireless mesh networks. This is the only directly related work to ours. While their paper is an important step in understanding routing behavior in wireless mesh networks, our study looks at link dynamics which is of broader applicability than routing stability. Our study is also conducted on live networks instead of post-processed traces which allow us to study the impact of traffic on link dynamics. We also study both the 5Ghz and 2.4Ghz physical layers. We also report on variations across channels and study time of day effects. Another interesting way we differ from this work is that we are likely to capture more link dynamics since while routes may remain stable, it is not clear if that implies that links are stable. For example, there may not be a good alternate route which is causing the route stability despite links being dynamic. Finally, our measurement period is on the order of every 5 seconds compared to 60 seconds in this work which gives us a better observation window on link dynamics.

6. CONCLUSION AND OUTLOOK

This paper presents an initial study of link metrics and their dynamics in wireless mesh network deployments and forms the basis for future avenues of work. This work hints that performance optimizations may be possible but need to take into account the inherent link dynamics i.e. choosing more stable links in routing or dealing with high churn time-of-day based effects. Another major finding is that current link metrics completely fail when simultaneous flows occur and the numbers they provide are essentially meaningless. This is because the broadcast-based metrics are highly sensitive to background traffic while the unicast-based metrics are insensitive to the existing traffic. Thus, there is a need to design link metrics that are either unaffected by background traffic (which can be accounted for by other means) or are able to change correctly to reflect the expected performance, given the existence of background traffic. This is an important research problem that this paper highlights. Finally, our measurements can potentially be used to derive better link behavior models for improving the realism of protocol simulation in single radio and multi-radio wireless networks.

Acknowledgement

This work was supported in part by NSF grants CNS-0338842 and CNS-0626703. We are particularly grateful to Yunnan Wu and the CCS group at Microsoft Research Redmond, for allowing use of their testbed for trace collection. We also thank our shepherd Prof. Dina Katabi and the anonymous reviewers for their feedback.

7. REFERENCES

- [1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *Proc. of ACM SIGCOMM*, 2004.
- [2] S. Biswas and R. Morris. ExOR: Opportunistic multi-hop routing for wireless networks. In *Proc of ACM SIGCOMM*, 2005.
- [3] D. S. J. D. Couto, D. Aguayo, J. C. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *Proc. of ACM MobiCom*, 2003.
- [4] R. Draves, J. Padhye, and B. Zill. Comparison of routing metrics for static multi-hop wireless networks. In *Proc. of ACM SIGCOMM*, 2004.
- [5] R. Draves, J. Padhye, and B. Zill. Routing in multi-radio, multi-hop wireless mesh networks. In *Proc. of ACM MobiCom*, 2004.
- [6] S. Keshav. A control-theoretic approach to flow control. *Proc. of the conference on Communications architecture and protocols*, pages 3–15, 1993.
- [7] K. Ramachandran, I. Sheriff, E. Belding-Royer, and K. Almeroth. Routing stability in static wireless mesh networks. In *Proc. of PAM*, 2007.
- [8] MIT Roofnet. <http://www.pdos.lcs.mit.edu/roofnet>.
- [9] OLSR routing protocol. <http://www.olsr.org>.