

Estimation of Link Interference in Static Multi-hop Wireless Networks

Jitendra Padhye[†], Sharad Agarwal[†], Venkata N. Padmanabhan[†], Lili Qiu[‡], Ananth Rao[§], Brian Zill[†]

[†]Microsoft Research

[‡]University of Texas, Austin

[§]University of California, Berkeley

We present a measurement-based study of interference among links in a static, IEEE 802.11, multi-hop wireless network. Interference is a key cause of performance degradation in such networks. To improve, or to even estimate the performance of these networks, one must have some knowledge of which links in the network interfere with one another, and to what extent. However, the problem of estimating the interference among links of a multi-hop wireless network is a challenging one. Accurate modeling of radio signal propagation is difficult since many environment and hardware-specific factors must be considered. Empirically testing every group of links is not practical: a network with n nodes can have $O(n^2)$ links, and even if we consider only pairwise interference, we may have to potentially test $O(n^4)$ pairs. Given these difficulties, much of the previous work on wireless networks has assumed that information about interference in the network is either known, or that it can be approximated using simple heuristics. We test these heuristics in our testbed and find them to be inaccurate. We then propose a simple, empirical estimation methodology that can predict pairwise interference using only $O(n^2)$ measurements. Our methodology is applicable to any wireless network that uses omni-directional antennas. The predictions made by our methodology match well with the observed pairwise interference among links in our 22 node, 802.11-based testbed.

1 Introduction

Multi-hop wireless networks have been a subject of much study. Most of the original work in this area was motivated by scenarios in which the nodes were highly mobile. Recently, interesting commercial applications of *static* multi-hop wireless networks have emerged. One example of such applications is “community wireless networks” [2, 11]. Several companies [9, 14] are field-testing wireless “mesh” networks to provide broadband Internet access.

Interference among wireless links significantly impacts the performance of static multi-hop wireless networks. Several researchers have studied this issue. The impact of interference on the capacity of wireless networks is studied in [8, 10, 12], while the impact on the performance of transport-level protocols is considered in [6, 7, 13]. The need for routing protocols to take link interference into account has been underscored in [4, 5]. Information about link interference is also needed for optimal channel assignment [16].

Many of these studies build upon the knowledge of which links in the network interfere with each other. Yet, the problem of *estimating* the interference among links in a multi-hop wireless network has not been adequately addressed.

The problem of estimating link interference can be described informally as follows: given a set of wireless links, estimate whether (and by how much) their aggregate throughput will decrease when the links are all active simultaneously, compared to when they are active individually.

This is a challenging problem for several reasons. Accu-

rate modeling of radio signal propagation is difficult since many environment and hardware-specific factors must be considered. Empirically testing every group of links for interference is not practical: a network with n nodes can have $O(n^2)$ links, and even if we consider only pairwise interference, we may potentially have to test $O(n^4)$ pairs. The interference pattern could change due to environmental factors, so interference estimation is not a one-time task. Hence, it is important to do it efficiently.

Given these difficulties, some researchers simply assume that the information about which links in the network interfere with each other is known a priori [10]. Others assume that it can be approximated by using simple heuristics. One common heuristic states that the interference range equals a small multiple (typically, a factor of 2) of the communication distance [8, 18].

In this paper, we study the phenomenon of interference among links in a 22-node IEEE 802.11 a/b/g based indoor wireless testbed. The paper makes two contributions. First, we show that the simple heuristics described in the previous literature fail to accurately predict the interference among links in our testbed. Second, we propose a simple, empirical estimation methodology to predict pairwise interference that requires only $O(n^2)$ measurement experiments. We show that the predictions made by our methodology match well with the observed pairwise interference among links in our testbed under a variety of conditions. Our methodology is useful for any wireless network where nodes use omni-directional antennas. We focus on omni-directional antennas since these are cheap and easy to deploy, and hence popular. Network architectures based on omni-directional antennas are quite common [17].

Paper outline: First, we formalize the notion of pairwise interference among wireless links. Next, we present a brief description of our testbed. We show that the simple heuristics used in previous work do not accurately model the interference in our testbed network. We next present our empirical methodology, and show that it accurately predicts pairwise link interference in our testbed. Finally, we summarize related work and present our conclusions.

2 Interference among wireless links

In this section, we define a metric to measure interference between a pair of wireless links. We assume that nodes communicate using the IEEE 802.11 protocol; parameters such as transmit power, data rate etc. are all set to fixed values; and the background noise level is constant. We also assume that RTS/CTS handshake is disabled for all nodes, which is the default behavior for most wireless cards.

We start by defining what constitutes a wireless link. Unlike in a wired network, the links in a wireless network are

not well-defined. For the purposes of this paper, we define wireless links using packet loss rate. We say that a link from node A to node B , denoted by L_{AB} , exists if the packet loss rate in either direction does not exceed some threshold. We defer a detailed discussion of the definition until Section 4.2.

We now define a metric to measure interference between a pair of links. Consider links L_{AB} and L_{CD} . For some fixed packet size, let U_{AB} denote the unicast throughput of the link L_{AB} , when no other links are active in the network. Similarly define U_{CD} for link L_{CD} . Now assume that both L_{AB} and L_{CD} are active simultaneously. Let their respective unicast throughput be denoted by $U_{AB}^{AB,CD}$ and $U_{CD}^{AB,CD}$. Define the *link interference ratio* as:

$$LIR_{AB,CD} = \frac{U_{AB}^{AB,CD} + U_{CD}^{AB,CD}}{U_{AB} + U_{CD}} \quad (1)$$

Thus, LIR is the ratio of aggregate throughput of the links when they are active simultaneously, to their aggregate throughput when they active individually.

LIR takes values between 1 and 0. The maximum value of LIR is 1, which means that the aggregate throughput does not decrease when the links are active simultaneously. Thus, $LIR = 1$ implies that the links do not interfere. A value of LIR less than 1 means that the aggregate throughput of the links decreases when they operate simultaneously. Thus, $LIR < 1$ implies that the links interfere with each other. The links can interfere with each other due to several reasons, listed below. Consider two links, L_{AB} and L_{CD} :

Carrier Sense: The 802.11 protocol requires the sender to monitor the radio channel for signs of activity, prior to transmitting a packet. If any activity is detected, transmission is deferred until a later time¹. This is known as *carrier sensing*. If the two senders, A and C are within the carrier sense range of each other, then only one of them will transmit at a time. Otherwise, they may both transmit, and one of the following may occur.

Data-Data Collision: The transmission by C may generate sufficient noise at B to interfere with reception of the packet being sent by A . A similar “collision” may occur at D . This is known as the hidden terminal problem.

Data-ACK Collision: For unicast communication, the 802.11 protocol requires the receiver of a packet to transmit an acknowledgment to the sender. If node D successfully receives the data packet sent by C , it will transmit an ACK. This transmission may interfere with ongoing reception of data packet at B . A similar collision may occur at D .

ACK-Data Collision: The data packet sent by C may interfere with ongoing reception of ACK sent by B at A . A similar collision may occur at C .

ACK-ACK Collision: The ACK sent by D may interfere with the reception of ACK sent by B at A . A similar collision may occur at C .

¹This is a simplified description of the actual protocol.




Figure 1: Layout of our testbed

A typical value of LIR is 0.5, which means that the aggregate throughput of the links is halved when they are active together. This usually (but not always) happens when the senders are within carrier sense range of each other. The minimum value of LIR is 0. This means that the links get zero throughput when they operate together. This can happen if the senders are not within the carrier sense range of each other, and collisions at the receiver are frequent.

In practice, we see a range of LIR values instead of just the three described above. They can result from packet losses, variable nature of background noise etc. Many of the simple heuristics used to estimate link interference only predict whether a pair of links interfere with each other or not. In other words, they only predict LIR is less than 1. We term this as the “binary” notion of interference.

3 Testbed

The experimental data reported in this paper was collected using a 22-node wireless testbed, located on one floor of a typical office building. The nodes are placed in offices, conference rooms and labs (Figure 1). All rooms have floor-to-ceiling walls and wooden doors. The nodes were not moved during testing. Each node is equipped with two 802.11 wireless cards: a Proxim ORiNOCO a/b/g Combo-Card Gold, and a NetGear a/b/g WAG 511. For experiments described in this paper, the two cards were never active simultaneously. RTS/CTS handshake was disabled (by default). All cards operated in the 802.11 ad-hoc mode. We used the built-in antennas of these cards, which are roughly omni-directional.

4 Performance of Simple Heuristics

In this section, we consider the performance of three simple heuristics from previous literature. The experiments described in this section use the following settings. All Orinoco cards were turned off. All Netgear cards were set to operate in 802.11a mode, on channel 36, at full transmit power. The transmit rate of each card was fixed at 6Mbps. We use 802.11a mode for these tests since our building has an operational 802.11b network.

4.1 The Heuristics

The first heuristic that we consider is used in [4, 5]. It assumes that all links on a multi-hop wireless path interfere with each other. In a connected network, we can always construct a path that includes a given pair of links. In short, this heuristic assumes that any two links in the network interfere with each other. This is clearly a pessimistic model.

The second heuristic [12] assumes that two links in the network interfere only if they share an endpoint. This is an optimistic heuristic. It is commonly known as the *point interference* model.

We do not expect either of these two heuristics to work well in our testbed. We include these in our study because they represent the two extreme ends of the approximations that have been used in the literature. Our measurements indeed show that these heuristics perform poorly.

The third heuristic [8, 18] is more sophisticated. Consider two links L_{AB} and L_{CD} . Let d_{AB} be the distance between nodes A and B . Similarly define d_{CD}, d_{BC} and d_{AD} . The model says that L_{AB} and L_{CD} will interfere with each other if either $d_{BC} \leq K * d_{AB}$ or $d_{AD} \leq K * d_{CD}$. A commonly used value for K is 2. Intuitively, the model says that if a node is receiving a transmission, a second transmitter can interfere with that reception only if it is sufficiently close. This model is generally paraphrased as *interference range is twice the communication distance*.

We term these three heuristics as $M1, M2$ and $M3$, respectively. These heuristics are “binary” models, since they only predict whether a given pair links interfere with each other or not. They do not predict the actual *LIR*.

To see how these three models perform in our testbed, we compare their predictions for several pairs of links against experimentally measured *LIR*. The first step in this process is to select a set of links in the testbed to experiment with.

4.2 Which links to use?

We define wireless links using packet loss rate as in [4]. Packet loss rates are easy to measure and reflect the link quality experienced by higher layers. For a pair of nodes that communicates using the 802.11 protocol, packet loss rate in both directions matters, since a unicast packet transmission is considered successful only if the sender successfully receives the ACK sent by the receiver. To discover links that have reasonably low packet loss rate in both directions, we carried out the following experiment.

We had each node in our testbed broadcast 1000 byte packets for 30 seconds. Only one node was active at a time. We measured the packet reception rate at all other nodes in the network. The entire test was repeated 50 times. This data gives us the average packet loss rate between every ordered pair of nodes in our testbed. For two nodes A and B , let P_{AB} be the packet loss rate from A to B , and let P_{BA} be the loss rate from B to A . We say that links L_{AB} and L_{BA} exist if: $1/((1 - P_{AB}) * (1 - P_{BA})) \leq \alpha$, where $\alpha \geq 1$ is some threshold value. This definition was proposed in [4], where the ratio is called the ETX (expected transmissions) value of the link. For the purposes of our paper it is sufficient to note that a high ETX value implies that the link is lossy in either one or both directions.

We use the threshold of $\alpha = 3$ in the rest of the paper, to weed out highly lossy links. Any reasonable routing protocol will avoid such poor quality links. Of the $22 * 21 = 462$ possible links in our testbed, 152 links have $\alpha \leq 3$. The average loss rate of these 152 links is 2.9%. We have ex-




Figure 2: Median LIR of 75 link pairs

perimented with lower values of α ; lower values reduce the number of links, but our interference results remain similar.

Note that we measure packet loss rate in both directions with 1000 byte packets, even though ACK packets are much smaller than data packets. It is not our aim to accurately characterize the loss rate of a link - we only want to eliminate highly lossy links that a routing protocol will avoid. Similar approach has been used in [4, 5].

4.3 Estimation of interference

Given the 152 links as defined above, there are a total of 11476 possible link pairs in our network. We ignore pairs in which the links share at least one endpoint, since such links will always interfere with one another. After removing such pairs, we are left with 9168 link pairs, which are still too many to be tested exhaustively. In this paper, we present results for 75 of these pairs, selected at random. We have also done some experiments with larger groups of link pairs, and have seen similar results.

For each selected link pair, we measured *LIR* as follows. For each link in the pair, we measured the unicast throughput using 1000 byte UDP packets for 30 seconds. Immediately afterwards, we measured the aggregate throughput of the two links operating together, again using unicast UDP packets for 30 seconds. Using the definition in Equation (1) we calculated the *LIR* for this pair. Testing links in a pair in quick succession helps mitigate the impact of environmental variations. We repeated the experiment 5 times for each of our 75 link pairs. Thus the total duration of the experiment was just under 10 hours. The median *LIR* value for these 75 link pairs are shown in Figure 2. Note that testing all 9168 pairs would have required more than 1100 hours.

First, note that we have several link pairs with intermediate *LIR* values between 0.5 and 1. In other words, interference is not a binary phenomenon. To compare this data with the binary predictions of $M1, M2$ and $M3$ models, we must pick a threshold, β . If $LIR < \beta$, we deem the links to have interfered in our experiment. If $LIR \geq \beta$, we deem that the links did not interfere.

Of the 75 node pairs, 24 pairs have *LIR* of 1. Thus, in each of these 24 pairs, the two links do not interfere with each other. Five other link pairs have *LIR* values between 1 and 0.9. Given the minimal interference, we classify these link pairs as non-interfering as well. Thus, we set $\beta = 0.9$. With this threshold, we have 29 pairs in which links do not interfere, and 46 pairs in which the links do interfere.

We see that the $M1$ model is too pessimistic for our net-

	M3 Prediction	
	Interference	No interference
Observed Interference	46	0
Observed No Interference	10	19

Table 1: Performance of M3 model

work, since we do have 29 non-interfering link pairs. On the other hand, the *M2* model is too optimistic. The two links in each pair do not share an endpoint, so according to the *M2* mode, none of the link pairs should show any interference. Yet, we have 46 link pairs in which the links *do* interfere.

The *M3* model is harder to verify. It is defined in terms of distance between nodes. We found that the predictions made using distance are quite inaccurate in our testbed. In an indoor testbed like ours, the radio signal propagation is also affected by office walls and other obstacles. There is no easy way to incorporate this information in the model. Therefore, we define a variant of the *M3* model that does not rely on physical distance between nodes. We will say that a pair of links L_{AB} and L_{CD} interfere if there is a 2 hop (or shorter) path from C to B , or from A to D . In other words, the modified model says that a pair of links will interfere if the sender of one link is within two hops of the other link’s receiver. Note that “hop” is just another term for a wireless link. This variant of the *M3* model predicted that 56 of the 75 link pairs will show interference. In our experiments, we observed interference in only 46 of these 56 pairs. The other 10 pairs did not show interference in our experiments. On the other hand, the model predicted no interference for 19 pairs. We indeed did not observe interference in any of these 19 pairs. These numbers are summarized in Table 1. The conclusion is that the model is pessimistic: it errs on the side of predicting interference even when there is none.

It may appear that the model seems pessimistic because we used $\beta = 0.9$ to classify experimental observations, and it is too low a threshold. However, even if we use $\beta = 1$ to classify experimental observations (and hence classify more pairs as interfering), the model still incorrectly predicts interference in 7 pairs that do not see any interference.

The pessimistic nature of the model is probably due to the indoor setting of our testbed. In such an environment, the radio signal degrades much faster than it would in free space, thus limiting the overall interference. We also evaluated a 1-hop variant of the model, which turned out to be optimistic. We believe that it may be possible to modify the 1-hop variant further to provide better predictions. However, there is no guarantee that the predictions of the 1-hop model will be accurate in other environments. Furthermore, even the improved model will provide only binary predictions. In the following section, we present a measurement-based approach which automatically takes into account the impact of environmental factors, and is capable of predicting intermediate values of *LIR*.

5 Proposed empirical methodology

In the previous section, we showed that the simple models proposed in the literature do not accurately predict the inter-




Figure 3: Median LIR and BIR of 75 pairs.

ference in our testbed. We now present a simple empirical methodology to estimate *LIR*.

Recall that in Section 2, we listed several reasons why two links may impact each other’s throughput. If we ignore the ACKs (given their relatively small size), we can then use the following simple methodology to estimate the impact of carrier sensing and collision of data packets as follows.

First, have one node, say A , broadcast packets as fast as it can. Only one node is active at a time. Denote the send rate by S_A . Keep track of the delivery rate of packets at all other nodes in the network. For example, the delivery rate at node B will be denoted by R_{AB} . Have each node broadcast in turn. Then, select a pair of nodes, say A and C , have them broadcast packets together. Denote their send rates by S_A^{AC} and S_C^{AC} . At all remaining nodes measure the delivery rate of packets they receive from each of the two broadcasting nodes. For example, at node B , the delivery rate of packets from A is denoted by R_{AB}^{AC} . Similarly, at node D , the delivery rate of packets from C is denoted by R_{CD}^{AC} . Have each pair broadcast in turn. Thus, we have carried out a total of $O(n^2)$ experiments.

Consider links L_{AB} and L_{CD} . Using the data gathered from the above methodology, we can define the “broadcast interference ratio” (*BIR*) as follows.

$$BIR = (R_{AB}^{AC} + R_{CD}^{AC}) / (R_{AB} + R_{CD}) \quad (2)$$

Our hypothesis is that the *BIR* is a good approximation of *LIR*. If the hypothesis is true, we can estimate interference for every pair of links using only $O(n^2)$ experiments, while testing each link pair will require far more (potentially $O(n^4)$) experiments. This is a substantial improvement: if we use 30 second transfers, and repeat each experiment 5 times, calculating *BIR* for every pair of links requires just over 28 hours. The key idea is that we can estimate unicast interference using broadcast packets, if we ignore impact of ACKs.

There are several reasons to believe that our hypothesis is correct. It is easy to see that *BIR* captures impact of carrier sensing on the two senders. It also captures the impact of data packet collisions at the receivers. The ACK packets are quite small (only 14 bytes) and the chance of them colliding with each other is also small. There are also some reasons to believe that the hypothesis is not correct. First, if a broadcast packet is lost (say due to collision), it is not retransmitted. On the other hand, a lost unicast packet is retransmitted multiple times, and the mean wait time (802.1 backoff) before sending the next retransmission is doubled.

Thus, a lost unicast packet has a higher impact on throughput measured at the user level. Second, while ACK packets may not collide with one another, data and ACK packets can still collide. We now test our hypothesis experimentally.

5.1 Evaluation: Baseline scenario

To test the hypothesis, we performed the following experiment. We use the same settings (802.11a, full transmit power, transmission rate fixed at 6Mbps) that we used in Section 4 and consider the same 75 link pairs as shown in Figure 2. To minimize the impact of environmental factors, the broadcast experiments designed to measure *BIR* were performed just before the unicast experiments designed to measure *LIR*. The median values of *BIR* and *LIR* for each link pair are shown in Figure 3. We see that *BIR* matches *LIR* well in most cases. The CDF of the absolute error ($|LIR - BIR|$) is shown in Figure 4. The median of absolute error is zero, and the mean is 0.026. Given that $|LIR - BIR|$ can range from 0 to 1, the mean and the median are quite low. Thus, our methodology works quite well in this scenario.

These results bring up several interesting questions. First, does the methodology work for other scenarios? Second, note that we carried out the broadcast and the unicast experiments back-to-back. In reality, we must do all the broadcast experiments together, and then use the results to predict link interference. The question then becomes: if we do broadcast experiments separately, will *BIR* obtained at some point in time still match *LIR* observed at some later point? Third, is the model capable of telling us *why* two links interfere? We discuss these questions next.

5.2 Other scenarios

We considered three other scenarios to evaluate our methodology. In the first scenario, we turned on the autorate feature for each card. When the autorate algorithm is on, the transmission rate for unicast packets may vary over time, in response to changing noise levels etc. The rate selection algorithm is not standardized. The broadcast packets, however, are always sent at the lowest data rate (6Mbps for 802.11a). Note that in the baseline scenario, the unicast transmission rate was also fixed at 6Mbps. With autorate on, we would expect more mismatch between *BIR* and *LIR*.

In the second scenario, we reduced the transmit power on each card to 50% of the full power. We fixed the transmission rate at 6Mbps. At 50% transmit power, the network has fewer links: only 128, instead of 152. Thus, the link pairs used in this scenario are different from the link pairs used in the previous, full-power scenarios. The average loss rate of these 128 links is 4.6%, while the average loss rate at full power was 2.9%. Since, the links are more lossy in this scenario, we would expect slightly higher mismatch between *BIR* and *LIR* in this setting.

All the experiments so far were done in 802.11a mode, using the NetGear cards. In the third scenario, we turned off the Netgear cards, and used Orinoco cards, set to operate in 802.11g mode (i.e. in 2.4GHz spectrum), at full power, with rate fixed at 1Mbps (i.e. the lowest data rate

for 802.11g). We have an infrastructure mode 802.11b network in our building, which operates in the same frequency band. We tested this scenario at night, to minimize the impact of interference from the WLAN, however, we would still expect to see higher error in this scenario.

For each of these scenarios, we measured *BIR* and *LIR* of 75 link pairs, using back-to-back experiments as before. The CDF of absolute error in each of the three cases is shown in Figure 5. The results show that our methodology performs generally well in each scenario. As expected, the mismatch is somewhat high for the autorate scenario. In the other two cases, the median error is only 0.01. Even in autorate case the median error is only 0.03, while the mean is 0.065. These three experiments increase our confidence in the general applicability of our method.

5.3 *BIR* and *LIR* measured 5 days apart

We used the same settings as the baseline scenario, but did only the broadcast experiments. We compare *BIR* calculated from these experiments with the *LIR* measured in the baseline experiment. The two experiments were done 5 days apart. The CDF of absolute error is shown in Figure 6. The graph also shows the baseline error CDF (labeled “Back-to-back”) for comparison purposes. We see that *BIR* is still generally a good predictor of *LIR*, but as expected, the error is somewhat higher compared to the baseline (back-to-back) case. The median error is only 0.01, and mean is 0.049. The results show that even in a static environment like ours, the interference patterns are slightly different at different times. The need to repeat interference measurements underscores the need for an inexpensive experimental methodology to measure interference.

5.4 Why do links interfere?

Our methodology also helps determine *why* two links interfere with one another. Consider links L_{AB} and L_{CD} that interfere with one another. During the broadcast experiments done to determine *BIR*, we had nodes *A* and *C* broadcast alone, as well as together. Consider the ratio of their *send rates*, when they were broadcasting together to when they were broadcasting alone. Define *carrier sense ratio*:

$$CSR = (S_A^{AC} + S_C^{AC}) / (S_A + S_C). \quad (3)$$

Note that we are using broadcast packets, so both senders send at the same data rate. If two senders are within the carrier sense range of each other, then only one of them would be able to send at a time, resulting in a *CSR* value of 0.5. If the senders are not within each other’s carrier sense range, *CSR* will be 1. Intermediate values can result from noise, differences in sensitivity of antennas, signal strength fluctuations due to environmental factors etc.

In the baseline scenario shown in Figure 2, 46 link pairs have $LIR < 0.9$, indicating some degree of interference. Of these, 34 link pairs had a *CSR* of 0.5. Thus, carrier sensing seems to be the major cause of interference in our testbed. We see similar results for the other three scenarios considered in Section 5.2. We believe that this is one of the reasons why *BIR* and *LIR* show a good match under all scenarios. We are currently investigating this issue further.




Figure 4: Baseline Scenario




Figure 5: Three other scenarios.




Figure 6: Measured 5 days apart

6 Related Work

The importance of studying wireless interference has long been recognized. For example, in [15] the impact of interference on fairness is considered. In [19] it is shown that paths with high degree of interference suffer disproportionately. Several researchers [8, 10, 12, 13] have considered impact of interference on the overall capacity of a multi-hop wireless network. However, each of these papers assumes that the information about link interference is available, but they do not describe how to estimate it. Thus our work on a practical method for estimating link interference helps generate the information taken for granted in previous work.

As we discussed earlier, various heuristics for estimating link interference have been proposed [4, 8, 12, 18]. We have shown that our empirical methodology can provide a more accurate estimation of pairwise link interference.

The knowledge of which links interfere with one another can benefit a number of network operations. For example, it can improve routing algorithms [4, 5], help in engineering and managing multi-hop wireless networks [3], and aid in channel assignment [16].

There is a large body of work on measuring various properties of wireless networks. Here we list some of the recent work. Aguayo et al. [1] analyze the causes of packet loss in an outdoor multihop 802.11b network. Yarvis et al [20] use testbeds in three different houses to study the properties of home wireless networks. Gupta et al [7] experimentally study the performance of TCP in a multi-hop wireless network. Our work contributes a practical technique for measuring another key property, viz., wireless interference.

7 Conclusion and Ongoing Work

In this paper, we considered the problem of estimating pairwise interference among links in a multi-hop wireless testbed. Using experiments done in a 22-node, 802.11-based testbed, we showed that some of the previously-proposed heuristics for predicting pairwise interference are inaccurate. We then proposed a simple, empirical methodology to estimate pairwise interference using only $O(n^2)$ measurements. We showed that our methodology accurately predicts pairwise interference among links in our testbed in a variety of settings. Our methodology is applicable to any 802.11-based wireless network where nodes use omni-directional antennas.

There are several avenues for future work. We hope to increase the accuracy of our methodology by accounting for the impact of four factors that we ignored in this paper. These four factors are: (i) retransmissions of lost unicast

packets, (ii) RTS/CTS handshake (iii) collisions between data and ACK packets (iv) autorate algorithms.

We would like to extend our approach to estimate interference among larger groups of links, instead of just pairwise interference.

Finally, we note that our methodology requires nodes to generate broadcast traffic, and existing traffic on the network can significantly reduce the accuracy of our approach. We are currently exploring the possibility of determining interference patterns by simply observing correlation between existing traffic flows on the network.

References

- [1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *SIGCOMM*, 2004.
- [2] Bay area wireless users group. <http://www.bawug.org/>.
- [3] R. Chandra, L. Qiu, K. Jain, and M. Mahdian. On the placement of internet taps in wireless neighborhood networks. In *ICNP*, 2004.
- [4] D. De Couto, D. Aguayo, J. Bicket, and R. Morris. High-throughput path metric for multi-hop wireless routing. In *MOBICOM*, 2003.
- [5] R. Draves, J. Padhye, and B. Zill. Routing in multi-radio, multi-hop wireless mesh network. In *MOBICOM*, 2004.
- [6] Z. Fu, P. Zerfos, H. Luo, S. Lu, L. Zhang, and M. Gerla. The Impact of Multihop Wireless Channel on TCP Throughput and Loss. In *INFOCOM*, 2003.
- [7] A. Gupta, I. Wormsbecker, and C. Williamson. Experimental evaluation of TCP performance in multi-hop wireless ad hoc networks. In *MASCOTS*, 2004.
- [8] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Trans on Info Theory*, Mar 2000.
- [9] Invisible Networks. <http://www.invisible.uk.net/how/>.
- [10] K. Jain, J. Padhye, V. Padmanabhan, and L. Qiu. The impact of interference on multi-hop wireless network performance. In *MOBICOM*, 2003.
- [11] R. Karrer, A. Sabharwal, and E. Knightly. Enabling Large-scale Wireless Broadband: The Case for TAPs. In *HotNets*, 2003.
- [12] M. Kodialam and T. Nandagopal. Characterizing achievable rates in multi-hop wireless networks: The joint routing and scheduling problem. In *MOBICOM*, Sep. 2003.
- [13] J. Li, C. Blake, D. S. J. De Couto, H. I. Lee, and R. Morris. Capacity of ad hoc wireless networks. In *MOBICOM*, 2001.
- [14] Mesh Networks Inc. <http://www.meshnetworks.com/>.
- [15] T. Nandagopal, T. Kim, X. Gao, and V. Bharghavan. Achieving MAC Layer Fairness in wireless packet networks. In *MOBICOM*, 2000.
- [16] A. Rainwala and T. Chiueh. Architecture and algorithms for an IEEE 802.11-based multi-channel wireless mesh network. In *INFOCOM*, 2005.
- [17] MIT roofnet. <http://www.pdos.lcs.mit.edu/roofnet/>.
- [18] K. Xu, M. Gerla, and S. Bae. How effective is the IEEE 802.11 RTS/CTS handshake in ad hoc networks? In *GLOBECOM*, 2002.
- [19] X. Yang and N. H. Vaidya. Priority scheduling in wireless ad hoc networks. In *MOBIHOC*, June 2002.
- [20] M. Yarvis, K. Papagiannaki, and W. S. Conner. Characterization of 802.11 wireless networks in the home. In *WinMeE*, 2005.