

Characterizing Multi-Way Interference In Wireless Mesh Networks

Saumitra M. Das, Dimitrios Koutsonikolas, Y. Charlie Hu, and Dimitrios Peroulis
School of Electrical and Computer Engineering, Purdue University
Center for Wireless Systems and Applications
West Lafayette, IN, USA
{smdas, dkoutson, ychu, dperouli}@purdue.edu

ABSTRACT

Wireless mesh networks (WMNs) have been proposed as a solution for ubiquitous last-mile broadband access. A critical limiting factor for many WMN protocols in realizing their throughput potential is the interference between nodes in the WMN. Understanding and characterizing such interference is important for a variety of purposes such as channel assignment, route selection, and fair scheduling. Instead of using ad hoc heuristics, a recent study proposed characterizing interference in a WMN by *measuring* two-way interference, i.e., interference between each pair of communicating links.

In this paper, we study the extent of *multi-way interference*, i.e., the interference caused by multiple transmitters to a communicating link. We find through simulations and through measurements of a 32-node wireless testbed that even if these transmitters individually do not interfere significantly with a given communicating link, simultaneous transmissions of them have the potential to significantly affect the throughput of the communicating link. This implies that pairwise interference measurements may be optimistic when used to drive protocols in wireless mesh networks. Encouragingly, we find that this phenomenon, although significant when it occurs, is not widespread. In particular, multi-way interference caused significant additional throughput degradation compared to pairwise interference to a small fraction of the links in the testbed over our measurement period. In addition, we find that there is a strong correlation between the impact of multi-way interference and the quality of the link under consideration. We conclude with recommendations on how protocols should take multi-way interference into account.

Categories and Subject Descriptors

C.2.1 [Computer Communication Networks]: Network Architecture and Design—*Wireless Communication*

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.

WiNTECH'06, September 29, 2006, Los Angeles, California, USA.
Copyright 2006 ACM 1-59593-538-X/06/0009 ...\$5.00.

General Terms

Measurement, Performance

Keywords

Wireless mesh networks, interference, testbed measurements

1. INTRODUCTION

Wireless mesh networks are characterized by static mesh routers connected by wireless links to each other and to a few gateway nodes. The WMN routers effectively form a multi-hop wireless access backbone. Recently, the deployment and use of WMNs have increased significantly and several cities have planned and/or deployed WMNs ([17, 21, 20, 15, 18, 19]). Thus, improving WMN performance will have a direct impact on a growing population of users. Since the most significant application of such networks is to provide broadband Internet access to static or mobile hosts in areas where wired infrastructure is difficult or economically infeasible to deploy, it is important to optimize the network throughput of WMNs.

A well-known fundamental technique to improve throughput is to exploit “parallelism”. In wireless networks, parallelism is achieved through spatial reuse, i.e., enabling simultaneous transmissions of packets at multiple sender-receiver pairs. Since WMNs operate over a shared broadcast medium, such parallelism is fundamentally limited by signal interference, i.e., the nature and amount of interference caused by simultaneously operating transmitters to other receivers determines the amount of parallelism that can be exploited and consequently the network throughput achievable. Thus, it is important to study the nature and extent of interference in WMNs and its impact on the network throughput.

Improved characterization of interference can aid several other techniques proposed for WMNs. For example, multi-radio multi-channel protocols [10, 9] perform channel to interface assignment based on channel reuse possibilities which in turn depend on interference. As another example, packet scheduling based techniques [5, 14] also rely on the knowledge of interference in order to schedule transmissions in the network to achieve certain fairness objectives.

Characterizing interference through closed-form expressions is infeasible due to the computational complexity involved, difficulty in modeling complex multi-path fading, and other signal propagation phenomena such as the waveguide effect in hallways. Due to these difficulties, protocols proposed for

WMNs have relied on simplistic representation of interference and drive their algorithms with heuristic rules such as “everyone interferes with everyone else” and “nodes in twice the transmission range interfere”. Although simple to calculate, such heuristics can be far from accurate in modeling the actual interference. An important result showing the inaccuracy of such heuristics was presented in [7]. This work was also the first to provide a realistic and useful way to characterize interference through measurement. It proposed the BIR metric that can range from 0 to 1 denoting different levels of interference. To make the measurement process feasible in real deployments, it proposes to perform pair-wise measurements, i.e., it measures whether each pair of links in the network interfere with each other.

In this paper, we further investigate and characterize interference in WMNs. Specifically, we study whether the above model of considering two transmitters at a time is accurate enough to characterize interference in WMNs and if not, whether there are significant benefits from more thorough measurement and interference characterization. Fundamentally, interference can be caused to a link $S \rightarrow R$ under a vanilla 802.11 MAC if (1) S senses the carrier of another transmitter and backs off, or (2) the noise due to signals and multi-paths from other transmitters is enough to decrease the SINR below the packet reception threshold¹ at R . Thus, interference actually depends on all the transmitters and their effect on the link $S \rightarrow R$. For example, under the previous method, even if an interferer I_1 does not affect link $S \rightarrow R$ and interferer I_2 does not affect link $S \rightarrow R$, it is possible that the combined transmissions from I_1 and I_2 cause interference. This is critical since, under the previous model, a scheduler could for example make a bad decision by allowing these transmissions in parallel for spatial reuse. Thus, there is need to understand whether we need to consider this 3-way, ..., k -way interference² possible from simultaneously operating transmitters. We acknowledge that measuring k -way interference is time-consuming as the number of measurements grows as $O(n^k)$. The goal of our study is to understand the impact of such k -way interference so that it can be taken into account as appropriate in WMN protocols.

We perform our study using a detailed wireless simulator as well as measurements from a 32-node wireless mesh network testbed [6]. We find that even if a set of transmitters individually do not interfere significantly with a given communicating link, multiple such transmitters transmitting simultaneously have the potential to significantly affect the throughput of the communicating link. This is important since it implies that pairwise interference measurements may be optimistic when used to drive protocols in wireless mesh networks. Encouragingly, we find that this phenomenon, although significant when it occurs, is not widespread. In particular, multi-way interference caused significant additional throughput degradation compared to pairwise interference only to a small fraction of links in the testbed over our measurement period. In addition, we find that the impact of multi-way interference depends on the quality of the link under consideration.

¹This is determined by hardware receiver specifics and the rate (modulation) used.

²We define k -way interference for link $S \rightarrow R$ as the case when $k-1$ interferers transmit simultaneously with the sender S .

The rest of the paper is structured as follows. Section 2 details how interference occurs and motivates the possibility of multi-way interference and consequently the need to study and characterize it. Section 3 studies multi-way interference first through controlled-distance-based experiments in a detailed wireless simulator that models multi-path fading, capture and other signal propagation effects. We then present the results of measurements of multi-way interference in our 32-node wireless mesh network testbed. Finally Section 4 outlines related work and Section 5 concludes the paper with a summary and recommendations.

2. CAUSES OF INTERFERENCE IN WIRELESS MESH NETWORKS

We first review the primary causes of interference in wireless mesh networks and their characteristics. A variety of factors may cause interference in wireless networks. The most common ones are listed below.

Intentional interferers that transmit in the same band and in the same area. These usually include other 802.11 nodes which makes the interfering signal to have very similar structure with the desired signal. The IEEE 802.11b medium access protocol handles this problem by allowing nodes to transmit packets only when there is no other transmitting node. If traffic is sensed in the medium, nodes wait a predetermined amount of time before they attempt to listen and transmit their packets. As a result, this source of interference causes a direct reduction to the network throughput.

Non-intentional interferers that transmit in the same band and in the same area. Bluetooth nodes, microwave ovens, cordless phones and similar equipment are examples of these interferers. These sources typically emit signals whose structure is very different from the desired signal. For example, unlike 802.11b nodes that occupy a relatively wide bandwidth of nearly 30MHz, the spectral mask of a bluetooth signal is limited to 1MHz. In addition, bluetooth employs the frequency hop spread spectrum (FHSS) technique that causes nodes to hop over 79 frequencies of 1MHz bandwidth [16]. Such non-intentional interferers may cause two effects: (1) they may occupy the medium not allowing the desired nodes to transmit; or (2) they may transmit their signal while a desired transmission is in progress leading to damaged packets that need to be re-transmitted. In both cases, the network throughput will be impaired.

Multipath fading that leads to inter-symbol interference. Multipath fading occurs when the desired signal arrives at the intended node through several different paths (at least two). Multipath is caused by objects (and/or humans) that happen to exist in the vicinity of two communicating nodes (Figure 1). The physics that cause multipath signals are typically quite complex and described statistically by appropriate models such as in [13]. Nevertheless, they can be qualitatively described by three fundamental phenomena: reflection/transmission, diffraction and scattering. Reflection and transmission take place when the desired propagating wave impinges on an object whose dimensions are large compared to its wavelength (e.g. building walls, large desks, etc.). Diffraction occurs when the desired signal impinges on sharp edges such as wall edges. The wave bends around these edges and therefore can reach locations that are not optically visible from its source. Scattering hap-




Figure 1: Multipath caused by reflected, scattered and diffracted signals.

pens when the desired wave meets objects with dimensions significantly smaller than its wavelength (e.g. foliage). Scattering causes the wave to disperse in many different directions.

In typical indoor and outdoor environments, all these mechanisms occur several times as the desired wave propagates from its source to its destination. As a result, several copies of the desired signal arrive at the intended node. This effect is called delay spreading which is described by an average time delay, which represents the time window that delayed copies of the signal reach the receiver. Typical values range from $0.1\mu\text{s}$ in weak multipath environments to over $1\mu\text{s}$ in dense urban environments [8]. The delayed signal copies are typically weaker than the line-of-sight (LOS) signal (if it exists) but exhibit various phases depending on the followed wireless path and the objects they encountered. This may lead to intersymbol interference in the receiver that degrades the signal-to-noise ratio (S/N), thus leading to reduced throughput for the overall network. This is typically seen by an automatic reduction of the bit rate in the 802.11b protocol. In strong multipath environments, the protocol falls back to simpler coding schemes with higher redundancy in the error correction code and lower throughput.

In this paper, we specifically study *multi-way interference* (from multiple transmitters to a desired signal). When multiple interferers exist in a multipath environment, the previously described phenomena (intentional interferers and multipath fading) become coupled and possibly exacerbated. Depending on the constellation of the interfering nodes their signals and their delayed copies may add constructively or destructively at the receiving node of the desired signal leading to a complicated case. Even for static nodes in mesh networks, this can be dynamically changing if the multipath environment varies versus time (e.g. due to moving objects). In addition to the receiver, the combined signal power from multiple transmitters could affect the sender of the desired signal through carrier sense. We now proceed to studying multi-way interference using simulations and measurements.

3. MEASUREMENTS OF MULTI-WAY INTERFERENCE

To understand interference behavior in a controlled manner (e.g. with controlled node distances), we first study the impact of multiple interferers on a transmitting link using the Qualnet simulator.




Figure 2: Simulation topology (S:source, R:receiver, I1, I2, I3: interferers.)

3.1 Controlled simulation experiments

3.1.1 Methodology

The simulated topology is shown in Figure 2. In this figure S is the source node, R is the receiver node, and $I1, I2, I3$ are the interferers. The interferers are placed symmetrically around the receiver and the distance between the receiver and each interferer is the same. The distance between S and R (denoted as $S - R$ distance) is always smaller than the transmission range (e.g., R can always receive packets from S), while the distance between R and each of the interferers $I1, I2, I3$ (denoted as $I - R$ distance) is larger than the transmission range (e.g., R cannot receive any packet from $I1, I2$ or $I3$), but possibly smaller than the interference range (e.g., simultaneous transmission from at least one of $I1, I2, I3$ along with S may affect the throughput at R from S).

We ran three different sets of simulations; in each of them the $S - R$ distance is kept constant, equal to 100m, 200m, and 250m, respectively. For each set of simulations, we first measured the throughput at R from S when only S transmits, i.e., there is no interference. Then we added one, two, and three interferers, having them transmit simultaneously with S in distances varying from 1800m (out of interference range) to 500m, and we measured again the throughput from S to R . Each node broadcasted 1500-byte probe packets for 30 seconds. The nominal bit-rate was 2Mbps.

Qualnet can simulate a very realistic physical model. In our experiments we used the two-ray propagation model, along with thermal noise and Rayleigh fading. The noise factor was set to 7dB. The Rayleigh fading model is appropriate for modeling environments with many reflectors, e.g., trees and buildings, where the sender and the receiver are not in Line-of-Sight of each other. Such environments are common in WMNs. Since the effect of fading and thermal noise is random, we ran each set of simulations with 10 different seeds and we averaged the results. This is equivalent to repeating an experiment 10 times in a real testbed. The radio propagation range according to Qualnet is 255m. However, due to fading this range is not fixed, and we found that R can receive packets from S even at a distance of 400m.

3.1.2 Results

The results are shown in Tables 1, 2 and 3 for three different $S - R$ distances, namely 100m, 200m, and 250m. In these tables, column “None” contains the throughput at R when only the sender S transmits, and columns “ $I1$ ”, “ $I1 + I2$ ”, and “ $I1 + I2 + I3$ ” contain the throughput at R when one, two or

three interferers transmit simultaneously with S . In addition to the throughput result, we also provide the *receiver throughput ratio* defined as the throughput at R when only S transmits divided by the throughput at R when S and at least one interferer transmit simultaneously. The receiver throughput ratio varies between 0 and 1 and answers the question whether the other nodes interfere with S or not. In a binary model, there is no interference when the ratio is equal to 1 and there is interference when it is less than 1. However, such a binary model can be very strict by finding that most nodes in a network interfere with each other, preventing, for example, any spatial reuse in a scheduling algorithm. Hence in practice we consider that two nodes do not interfere when the ratio is kept above a threshold (set to 0.9 in this paper). In addition to the three tables, Figures 3(a), 3(b) and 3(c) show the percentage decrease in throughput in case of one, two, or three interferers compared to the case when there is no interferer.

As a general observation, in all the three tables the throughput remains almost unchanged, i.e., the receiver throughput ratio remains close to 1, for $I - R$ distances larger than 1200m, and it is significantly reduced only for distances smaller than 800m. Although these thresholds are not exactly the same for all cases, as they depend on the quality of the $S - R$ link, it becomes obvious that in a realistic environment the interference range can be much larger than twice the transmission range, a commonly used heuristic. Another general observation is that adding more interferers usually causes further throughput degradation, but in most cases, this degradation is not proportional to the number of interferers.

Table 1 and Figure 3(a) show the results for $S - R$ distance equal to 100m, which is much smaller than the theoretical transmission range (255m). In this case the link $S \rightarrow R$ is very good, and the throughput at R is high, equal to 1.299 Mbps, in absence of interferers. We also observe that such high quality links are in general unaffected from interferers, unless they are very close to R . Even in case of three interferers in an 800m distance from R , the receiver throughput ratio remains higher than 0.9 and the reduction in throughput is less than 10%. Only in distances smaller than 800m, the interference becomes critical. However, in these smaller distances, we observe cases where the pairwise interference model is not enough to describe interference. For example, for $I - R$ distances between 600m and 800m, the presence of one interferer causes a tolerated throughput reduction of less or equal to 10%, but the presence of two or three interferers causes a throughput reduction of more than 10%, lowering the throughput ratio to unacceptable levels. In such a case, the pairwise interference model would allow, for example, nodes S , I_1 , and I_2 to transmit simultaneously, resulting in lower throughput at R than the expected value. Only for very small $I - R$ distances, the pairwise interference model becomes again enough to correctly predict interference. For example, when interferers are 500m away from R , the presence of only one interferer reduces the throughput by 26%, preventing nodes S and I_1 from transmitting simultaneously.

Table 2 and Figure 3(b) show the results for $S - R$ distance equal to 200m, close to the theoretical transmission range. The longer distance between S and R worsens the link quality (throughput at R is now 983Kbps in absence of interferers), and makes the link more sensitive to interference. Reduction in throughput here starts at a larger distance (1100m) and for

Table 1: S-R distance: 100m

| I-R dist.(m) | Throughput(Kbps)/Throughput ratio | | | |
|--------------|-----------------------------------|------------|-------------|-------------------|
| | None | I_1 | $I_1 + I_2$ | $I_1 + I_2 + I_3$ |
| 1800 | 1299 | 1299/1 | 1299/1 | 1299/1 |
| 1700 | 1299 | 1302/1.002 | 1300/1 | 1300/1 |
| 1500 | 1299 | 1300/1 | 1300/1 | 1300/1 |
| 1200 | 1299 | 1296/0.998 | 1292/0.995 | 1292/0.995 |
| 1100 | 1299 | 1294/0.996 | 1288/0.991 | 1283/0.987 |
| 1000 | 1299 | 1290/0.993 | 1277/0.98 | 1268/0.97 |
| 800 | 1299 | 1261/0.97 | 1231/0.95 | 1199/0.92 |
| 600 | 1299 | 1181/0.90 | 1079/0.83 | 986/0.76 |
| 500 | 1299 | 960/0.74 | 801/0.62 | 650/0.5 |

Table 2: S-R distance: 200m

| I-R dist.(m) | Throughput(Kbps)/Throughput ratio | | | |
|--------------|-----------------------------------|-----------|-------------|-------------------|
| | None | I_1 | $I_1 + I_2$ | $I_1 + I_2 + I_3$ |
| 1800 | 983 | 983/1 | 983/1 | 983/1 |
| 1700 | 983 | 982/0.998 | 982/0.998 | 982/0.998 |
| 1500 | 983 | 979/0.995 | 979/0.995 | 979/0.995 |
| 1200 | 983 | 974/0.99 | 970/99 | 955/0.98 |
| 1100 | 983 | 965/0.98 | 950/0.97 | 944/0.96 |
| 1000 | 983 | 952/0.97 | 922/0.94 | 892/0.91 |
| 800 | 983 | 888/0.90 | 798/0.81 | 723/0.74 |
| 600 | 983 | 702/0.71 | 510/0.52 | 378/0.38 |
| 500 | 983 | 485/0.49 | 270/0.27 | 154/0.16 |

distances smaller than 800m the presence of only one interferer is enough to make the link quality unacceptably bad. Hence the pairwise interference model works in most cases. However there is still an area between 800m and 1000m, where the presence of one interferer keeps the ratio above 90%, but the addition of one or two more interferers reduces it to possibly unacceptable levels.

Finally, Table 3 and Figure 3(c) show the results for $S - R$ distance equal to 250m, almost equal to the theoretical transmission range. In such a long distance, fading causes significant losses and the receiver throughput is quite low (705Kbps) even in absence of interference. Throughput reduction because of interference now starts at 1200m and is much more intense than in the previous two cases. For small distances (500m) the throughput is reduced by about 65% with one interferer and it almost drops to 0 with two or three interferers. Again, we observe a zone of distances larger than 800m and smaller than 1100m where the pairwise interference model would fail to correctly predict interference. However, the low value of the receiver's throughput in such a link would probably cause a good routing metric (e.g., ETX [2]) to avoid such a link, even in absence of interference.

In summary, these controlled simulation experiments showed that in most cases the pairwise interference model gives satisfactory results in predicting interference. When an interferer is close to the receiver, it affects the sender-receiver link heavily, reducing the throughput beyond acceptable thresholds. On the other hand, when interferers are far away from the receiver, even the presence of many of them does not cause significant further throughput reduction compared to that caused by only one. However, in all three sets of simulations, we observed some distance ranges where the presence of more than one interferers can result in throughput reduction larger than the ac-

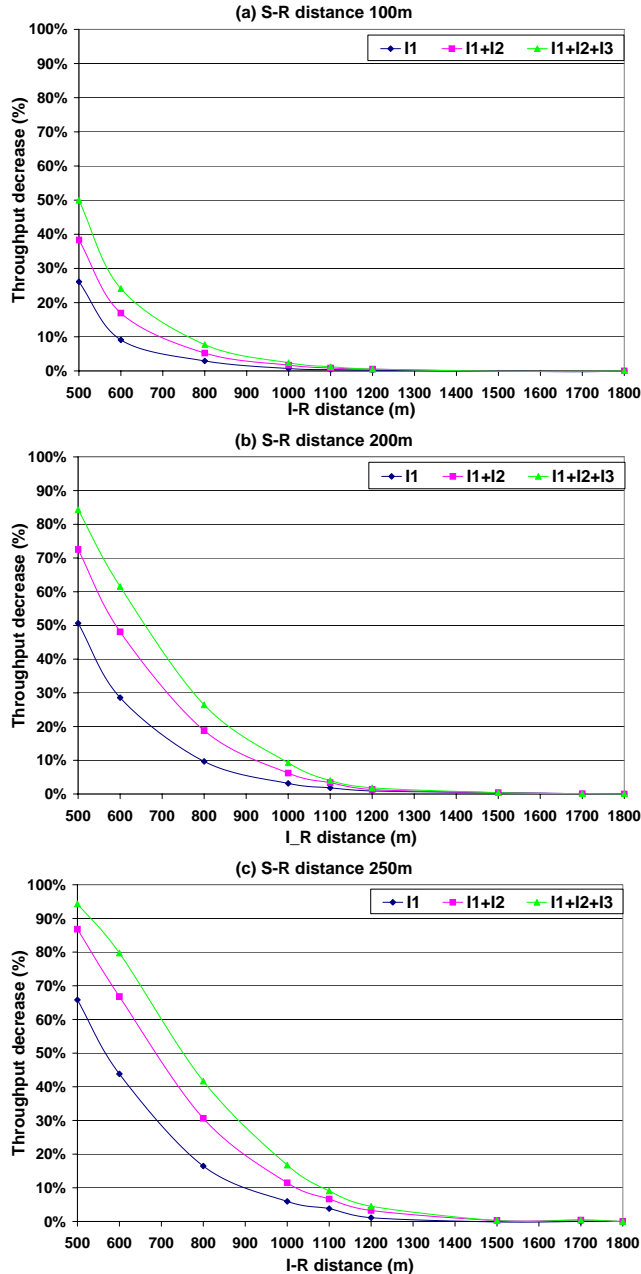


Figure 3: Throughput decrease vs. S-R distance for three different I-R distances.

Table 3: S-R distance: 250m

| I-R dist.(m) | Throughput(Kbps)/Throughput ratio | | | |
|--------------|-----------------------------------|----------|-------------|-------------------|
| | None | I_1 | $I_1 + I_2$ | $I_1 + I_2 + I_3$ |
| 1800 | 705 | 705/1 | 705/1 | 705/1 |
| 1700 | 705 | 705/1 | 702/0.995 | 702/0.995 |
| 1500 | 705 | 706/1 | 703/0.997 | 703/0.997 |
| 1200 | 705 | 697/0.99 | 682/0.97 | 673/0.95 |
| 1100 | 705 | 678/0.96 | 658/0.93 | 641/0.91 |
| 1000 | 705 | 663/0.94 | 624/0.89 | 587/0.83 |
| 800 | 705 | 589/0.84 | 489/0.69 | 411/0.58 |
| 600 | 705 | 396/0.56 | 234/0.33 | 143/0.2 |
| 500 | 705 | 241/0.34 | 93/0.13 | 40/0.06 |

ceptable levels, and in such cases considering 3-way, or higher class interference would be necessary. In the following section, we study if such cases appear and at what frequency in a real-world WMN testbed where the placement of senders and interferers is not controlled.

3.2 Testbed results

3.2.1 Testbed description

Our testbed, MAP [6], shown in Figure 4, currently consists of 32 mesh routers (small form factor desktops) spread out across four academic buildings on the Purdue campus (EE, MSEE, PHYSICS and ME). Each router has two radios. In this paper, we use one of them: the Atheros 5212 based 802.11a/b/g wireless card. Each radio is attached to a 2dBi rubber duck omnidirectional antenna with a low loss pigtail to provide flexibility in antenna placement. Each mesh router runs Mandrake Linux 10.1 and the open-source *madwifi* drivers are used to enable the wireless cards. IP addresses are statically assigned. 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. Apart from structural impediments, interference exists in our deployment from other 802.11b networks (the Purdue Airlink network). We used channel 11 of 802.11b to operate our network since it was the band furthest away from those being already used in the deployment environment. We performed our measurements only during the night to minimize interference from the other 802.11b networks and other sources such as microwaves.

3.2.2 Methodology

The 32 nodes of the testbed can form up to 992 (directional) links. However many of these links may not exist, for example, if two nodes are far away from each other or if they are separated by many obstacles such as walls and metallic objects. To find out the exact number of links in our testbed and to get a qualitative estimate of their quality, we first transferred a series of ping messages between each pair of nodes. If all ping messages between a pair of nodes are lost, we classified this link as non-existing. This experiment made sure that our network is not partitioned. As expected, the link quality and the node distances do not directly correspond. For example, there is a perfect link (0% loss) between nodes 1 and 13, which are placed in different buildings, because these nodes are close to windows and signal is propagated in the low-loss outdoor environment. On the other hand, there is no link between nodes 19 and 20, although their distance is half compared to the dis-



Figure 4: Top view topology of the MAP testbed

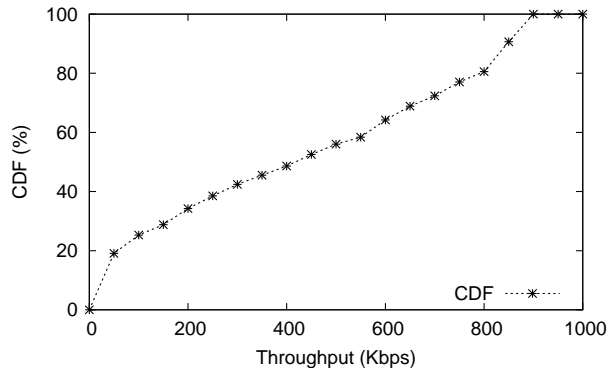


Figure 5: CDF of individual throughputs for the 257 links of the testbed.

tance of link 1 – 13. In total, we found 257 links with loss rate less than 100%.

To measure interference, we executed the following series of experiments with the rate of the radios set to 2Mbps. In the first set of experiments, each node in turn broadcasted 1500-byte packets as fast as 802.11 allowed it³ for 30 seconds. All other nodes that could receive packets measured the throughput in these 30 seconds. This experiment gave us the throughput of all existing links in our testbed. Figure 5 shows the cumulative distribution function (CDF) of the throughput for the 257 existing links.

In the second series of experiments, we measured the pairwise interference among all nodes in our testbed, following a methodology similar to [7]. For this experiment, we had each pair of nodes broadcast 1500-byte packets as fast as they can together for 30 seconds. In each 30 second interval, all other nodes except for the two senders measure the throughput from

³This method of sending packets has been previously used in studies that calculate link-level characteristics such as [2, 1].

the two senders. In the end of this experiment, for each pair of nodes, one viewed as the sender S and the other the interferer I (and vice versa), we calculated the receiver throughput ratio (RTR) at each one of the rest 30 nodes (R) as follows: $RTR_R^{S,I} = Throughput_R^{S,I} / Throughput_R^S$, where $RTR_R^{S,I}$ is the receiver throughput ratio at receiver R when S is the sender and I the interferer node, $Throughput_R^S$ is the throughput at R from node S when only S transmits and $Throughput_R^{S,I}$ is the throughput at R from node S when both S and I transmit simultaneously. If $RTR_R^{S,I} < 0.9$, we consider that node I is an interferer for link $S \rightarrow R$, else it is not. In this way we found out all nodes that are not interferers for each particular link $S \rightarrow R$ according to the pairwise interferer model. We found 45 links for which all the rest 30 nodes were acting as interferers. We removed these links, since the pairwise interference model is enough for them. We were left with 212 links.

In the third set of experiments, we considered each of the remaining 212 links in turn, along with their non-interferers (say n). For each of these links we had again the sender broadcast packets for 30 seconds along with 1, 2, 3, ..., n interferers simultaneously and every time we measured the sender's throughput at the receiver and calculated the receiver's throughput ratio. Comparing the ratio in case of 2, 3, ..., n interferers at a time with the ratio in case of 1 interferer will tell us for which links the pairwise model cannot accurately predict interference.

3.2.3 Does multi-way interference occur?

Figure 6 shows three examples in our testbed where the pairwise interference model is not enough. In these figures S is the sender, R is the receiver, and I_1, I_2, I_3, I_4 are nodes that were found not to interfere according to the pairwise model, but they might interfere if two, three, or four of them transmit together. We call these nodes interferers. Table 4 shows the receiver throughput ratio for different number of interferers in the three examples of Figure 6. The second column shows the minimum ratio at the receiver when interferers are considered one at a time, the third one the minimum ratio when interferers are considered in pairs, and the fourth one the minimum ratio when interferers are considered three at a time. Finally, the last column shows the ratio when all four interferers transmit simultaneously. As we observe, the ratio for the second column is above 0.9 in all the three examples (no interference according to the pairwise model).

In the first example, node 11 is the sender, node 2 is the receiver, and nodes 3, 10 the interferers. As Table 4 shows, if we allow both nodes 3 and 10 to transmit simultaneously with node 11, the receiver throughput ratio is reduced to 0.74, quite below the threshold of 0.9 and the quality of link 11 \rightarrow 2 worsens significantly. Note again that the distance of the nodes is not directly mapped to the quality of the links. For example, nodes 11 and 10 are almost equal distance away from node 2, however only link 11 \rightarrow 2 exists.

In the second example, node 16 is the sender, node 7 is the receiver, and nodes 8, 9, 10 the interferers. In this example, all the three interferers are very far from link 16 \rightarrow 7 and they cannot affect it (throughput ratio is 1 when each of these three nodes transmits simultaneously with the sender). Even when two of these nodes transmit together, the ratio is not reduced a lot – it only drops to 0.88, very close to the threshold. But

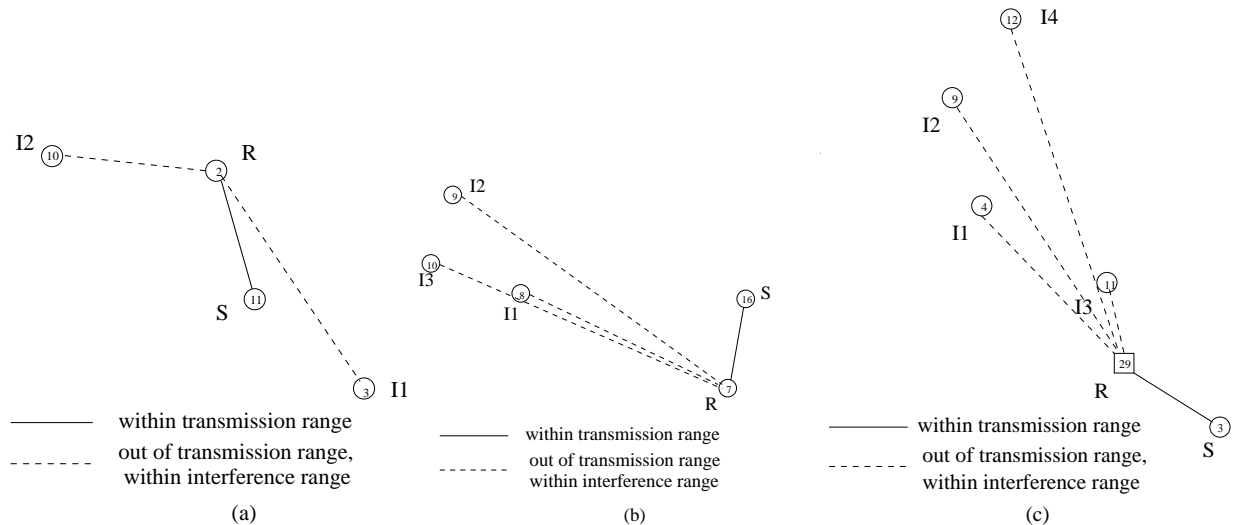


Figure 6: Three examples in our testbed showing that pairwise interference is not always enough.

when all the three nodes 8, 9, 10 transmit simultaneously, the ratio becomes 0.74 which makes the link quality unacceptable.

Finally, the third example shows that in some cases even 5-way interference has to be considered. In this example node 3 is the sender, node 29 is the receiver and nodes 4, 9, 11, and 12 the interferers. As Table 4 shows, the throughput ratio remains above the threshold when two or three of the interferers transmit simultaneously (with minimum values equal to 0.96 and 0.91 respectively) but it drops to 0.86 when all four interferers transmit simultaneously.

Table 4: Three examples in our testbed showing that pairwise interference is not always enough.

| Example | Throughput ratio | | | |
|---------|------------------|-------------|-------------|---------|
| | min 1 intf. | min 2 intf. | min 3 intf. | 4 intf. |
| (a) | 0.92 | 0.74 | - | - |
| (b) | 1 | 0.88 | 0.74 | - |
| (c) | 0.93 | 0.96 | 0.91 | 0.86 |

3.2.4 How widespread is multi-way interference?

Out of 212 links in our testbed, we found 16 for which pairwise interference model could not accurately predict interference. Out of these 16 links, without any interferers, 4 had throughput between 200 and 400 Kbps, 9 had throughput between 600 and 800 Kbps, and the rest three had throughput higher than 800 Kbps. Some general observations are as follows. As we observe in Figure 5, about 20% of the links have lower than 50 Kbps throughput. For those links, no observation can be made about interference, since the quality is so bad that adding more interferers cannot make it worse. Actually, when we repeated our experiments, many of them got zero throughput in some cases. Such links will probably not be selected by a routing protocol that uses a link-quality based routing metric (e.g., ETX [2] or SPP [12]). Hence, in the rest of the paper, we ignore these links. Similar methodology is

followed in [7] where low quality links are rejected using an ETX-based threshold.

We did not find any case of 3-way interference for the 32 links with throughput between 50 and 200 Kbps (about 15% of the total 212 links). For these links, throughput is still very low, although non-zero. For many of them we observed large variations in throughput when the number of interferers changed. For example, in many cases throughput was increased when we added interferers, compared to the case where only the sender transmitted. For the rest of them, one interferer was enough to reduce the throughput ratio to very low levels, hence the pairwise model was enough. Note that a reasonable routing algorithm should already avoid most of these links.

For the 59 links of medium quality with throughput between 200 and 600 Kbps (28% of the total 212), the pairwise interference model was successful in predicting interference in almost all cases. For those links, we did not observe the strange variations described in the previous paragraph, but in most cases one interferer was enough to change the link quality from medium to low and reduce the throughput ratio below acceptable levels. Only 4 links remained unaffected by single interferers, but were affected when two or three interferers transmitted simultaneously.

The majority of cases (9 out of the 16 links) where pairwise interference model was not enough were observed for the 36 links (17% of the total 212) of medium to high quality with throughput varying between 600 and 800 Kbps. Since throughput is high enough for these links, there is margin for gradual decrease by adding more interferers. Hence we had cases where one interferer reduced the ratio only slightly but without crossing the 0.9 threshold, the second interferer sent the ratio close to the threshold, and three or more interferers resulted in large throughput reduction.

Finally, for the 42 high quality links (20% of the total 212) with throughput higher than 800 Kbps, the common case is that if such a link is not affected by other nodes, when they are

considered one at a time, it is also not affected when the other nodes are considered more than one at a time. Hence, again the pairwise interference model gives the correct answer in most cases. But we still found three cases where 3-way interference should be considered.

Thus, while multi-way interference does occur and when it does it significantly affects throughput, we found that the phenomenon is not widespread and depends on the original link quality.

4. RELATED WORK

The most closely related work to this paper is [7] which proposed the first practical approach to estimating interference in wireless multi-hop networks. We extend this important work to include the study of multi-way interference. Various papers have proposed different heuristics for estimating interference [2, 3, 4]. However [7] showed that a measurement based approach is more accurate than using heuristics in a real network. Finally, [11] showed that interference can also occur between radios on the same node in multi-radio networks. In this case, multi-way interference could occur even at the node itself if all the radios do not have enough frequency separation and physical separation.

5. CONCLUSIONS AND FUTURE WORK

We conclude that multi-way interference can indeed significantly affect throughput performance in wireless mesh networks. However, due to the impracticality of actually measuring multi-way interference continuously, it is difficult to take this into account in real deployed protocols. Our recommendation is that protocols should perform pairwise experiments (that are practically feasible) and then be conservative in scheduling/allowing multiple interfering transmitters for spatial reuse. For example, one way of being conservative is to limit the number of simultaneous transmitters to some K and to select these K nodes in the sorted order such that the first one scheduled is the one with the lowest pairwise interference impact as measured by the BIR. Such a technique allows interference-aware protocols to take multi-way interference into account using a practically feasible approach.

There are several avenues for further investigation on multi-way interference such as characterizing multi-radio multi-way interference. In this case, radios on a single node could particularly interfere even if operated on orthogonal channels due to hardware imperfections. Other avenues of interest are to study how environmental attributes and the antenna type and placement affect multi-way interference. While we have shown that multi-way interference can sometimes significantly affect performance, an interesting related problem is to find the largest subset of links from a set of pairwise non-interfering links that are least affected by multi-way interference. It is hard to derive this from pairwise interference information since the joint behavior of multiple links depends on complex phenomena ranging from MAC layer behavior to signal propagation characteristics.

Acknowledgment

This work was supported in part by NSF grant ANI-0338856 and the Purdue Research Foundation.

6. 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 SIGCOMM*, 2004.
- [2] 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.
- [3] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, 2000.
- [4] M. Kodialam and T. Nandagopal. Characterizing achievable rates in multihop wireless networks: The joint routing and scheduling problem. In *Proceedings of ACM Mobicom*, 2003.
- [5] H. Luo, S. Lu, and V. Bharghavan. A new model for packet scheduling in multihop wireless networks. In *Proc. of ACM MobiCom*, 2000.
- [6] Purdue Mesh Testbed. <http://www.engineering.purdue.edu/MESH..>
- [7] J. Padhye, S. Agarwal, V. Padmanabhan, L. Qiu, A. Rao, and B. Zill. Estimation of Link Interference in Static Multi-hop Wireless Networks. In *Proceedings of IMC*, 2005.
- [8] J. Proakis. *Digital Communications*. McGraw Hill, 4th edition, 2000.
- [9] K. Ramachandran, E. Belding, K. Almeroth, and M. Buddhikot. Interference-aware channel assignment in multi-radio wireless mesh networks. In *Proc. of INFOCOM*, 2006.
- [10] A. Raniwala and T. Chiueh. Architectures and algorithms for an IEEE 802.11-based multi-channel wireless mesh network. In *Proc. of IEEE INFOCOM*, March 2005.
- [11] J. Robinson, K. Papagiannaki, C. Diot, X. Guo, and L. Krishnamurthy. Experimenting with a multi-radio mesh networking testbed. In *Proc. of WiNMe*, 2005.
- [12] S. Roy, D. Koutsonikolas, S. M. Das, and Y. C. Hu. High-Throughput Multicast Routing Metrics in Wireless Mesh Networks. In *Proc. of ICDCS*, 2006.
- [13] A. Saleh and R. Valenzuela. A statistical model for indoor multipath propagation. *IEEE Journal on Selected Areas in Communication (JSAC)*, 5(8), Feb 1987.
- [14] N. B. Salem and J.-P. Hubaux. A fair scheduling for wireless mesh networks. In *Proc. of WiMesh*, 2005.
- [15] Champaign-Urbana community wireless network. <http://www.cuwireless.net>.
- [16] M.B. Shoemake, Wi-Fi (IEEE 802.11b) and Bluetooth Coexistence Issues and Solutions for the 2.4 GHz ISM Band, Texas Instruments White Paper. <http://focus.ti.com/pdfs/vf/bband/coexistence.pdf>, Feb. 2001.
- [17] MIT Roofnet. <http://www.pdos.lcs.mit.edu/roofnet>.
- [18] Southampton wireless network. <http://www.sown.org.uk>.
- [19] Wireless leiden. <http://www.wirelessleiden.nl>.
- [20] Seattle wireless. <http://www.seattlewireless.net>.
- [21] Bay area wireless users group. <http://www.bawug.org>.