Interference Map for 802.11 Networks

Dragoş Niculescu
NEC Laboratories America
4 Independence Way
Princeton, NJ 08540, USA
dragos@nec-labs.com

ABSTRACT

The interference map of an 802.11 network is a collection of data structures that can help heuristics for routing, channel assignment and call admission in dense wireless networks. The map can be obtained from detailed measurements, which are time consuming and require network down time. We explore methods and models to produce the interference map with a reduced number of measurements, by identifying interference properties that help to extrapolate complex measurements from simple measurements. Actual interference in an 802.11a testbed is shown to follow certain regularities – it is linear with respect to packet rate of the source, packet rate of the interferer, and shows independence among interferers. When multiple cards are available, they behave differently, and even different channels of the same card have different performance. We find that while current methods of gathering the interference map may be appropriate for characterizing interference in one card networks, they are unscalable for multiple card networks when considering: 802.11 characteristics (card and channel asymmetries, time variation), required downtime, and complexity of the measurement procedure.

Categories and Subject Descriptors

C.2.1 [**Network architecture and design**]: Wireless communication; C.2.3 [**Network operations**]: Network monitoring

General Terms

Measurement, Experimentation

Keywords

802.11, interference, measurement, model

1. INTRODUCTION

For large populations of wireless devices, such as sensor networks with single RF nodes, or meshes with multiple RF nodes with small number of independent channels, interference is the major performance factor that is part of a circular feedback of interdependence. Interference determines how channels are assigned, how flows are routed, and how calls are admitted, but all these may in turn cause changes in the interference. For these dense networks, it is therefore important to get an understanding of interference that can help heuristics for all the mentioned problems, and provide prediction for the quality of service obtainable.

When enough orthogonal channels are available to cover with channel reuse an entire 802.11 network, interference between base stations is not of concern. However, for the current crop of available channels in the unlicensed spectrum 802.11b/g the number of channels is limited such that channels need to be reused in an indoor environment that requires higher deployment density. The availability of 802.11a, with its shorter range, higher bit rates, and more orthogonal channels provides one way of scaling up the wireless service: by increasing the bandwidth per area ratio. But this is done by increasing the density of access points and the number of wireless cards per access point. In such a dense wireless network, be it multiple hop (ad-hoc, mesh), or single hop, several base stations and clients operating on the same channel are bound to interfere each other. It is an accepted fact that indoors the nature of interference is generally unpredictable for these carriers (2.4GHz and 5GHz) due to variability in building construction, people movement, and other uncontrolled sources, such as microwave ovens. Adding that channel usage is unregulated by most institutions, it is generally hard to predict what the quality of service can be achieved even in a one hop setup. For multiple hops, the problem becomes harder because backhaul traffic interferes with itself when carried on the same channel. The nature of the traffic also makes the amount of interference hard to predict - TCP traffic depends on the congestion, which means that the interference it produces depends on conditions on other orthogonal channels. This tends to link together the problems of interference, routing, and channel allocation.

The problem that we address in this paper is to predict the effects of interference under some restrictive conditions:

1. we assume that traffic is completely controlled for the channels of interest. This can be achieved through administrative ways, and through call admission policies, and has the role of making all traffic at all nodes on all channels accounted for.

2. the traffic on all nodes is constant bit rate (CBR), which avoids problems caused by time variable behavior of congestion aware protocols, like TCP. These two requirements are actually met by VoIP networks, both in WLAN and mesh networks.

3. measurements of interference can be made in the absence of real traffic. They can be performed periodically at low traffic times, for example midnight/weekend. The purpose of this last requirement is the building of the interference map –

IMC'07, October 24-26, 2007, San Diego, California, USA.

a collection of measurements in which the relation between source, destination, and the interferers is unaffected by exogenous unpredictable factors such as live traffic.

One hypothesis that we validate in this work is that interference measurements for simple configurations can be used to predict interference effects for more complex scenarios. The simple configurations have the advantage of taking a reduced time to test, and of requiring collaboration of fewer nodes (in our case three). Scenarios with many interfering nodes are harder to test mainly because of complexity - the number of tests can grow exponentially with the size of the group. If interactions of triplets of nodes can be used to predict interaction for larger groups, the short measurements (that require network downtime) can be performed more often, improving the accuracy of the prediction.

While there is research that needs interference information for the mentioned problems (channel allocation, routing, call admission), interest in the actual measurement and modeling of interference is quite recent. In [1], Padhye et al. proposed a pairwise interference measurement method, which we also use in this work. The broadcast based interference estimation is shown to be an adequate estimation for interference produced between unicast flows. They also found that carrier sense is the major cause for interference. Das et. al [2] show that remote nodes which have no interference effects in isolation may combine to produce interference when acting together, but the occurrence is rare. Our proposed model closely predicts their results for close interferers. Another study of carrier sense [3] shows that it is not always a good predictor of transmission success, and also suffers of the exposed terminal problem (close by senders cannot send simultaneously even if their destinations can actually receive packets), and is overly conservative with respect to the capture effect. The question of interference is acknowledged to be central for the problems of channel assignment, bandwidth allocation and routing in [4]. The authors find that there is a circular dependency between these problems and interference, and propose a centralized algorithm. Other researchers have identified interference as being a cause for unfairness [5]. The most studied problem is that of capacity being affected by interference [6, 7, 8, 9], but complex interference scenarios are considered an input for the optimization process, without addressing the problem of obtaining them. More recent works [10, 11] try to reduce the complexity of measuring the interference by only considering pairwise communication, but this ignores the remote interferers (outside carrier sense range), which as we will show are the main source of uncontrolled loss.

One contribution of this work is a model of interference that allows complex traffic scenarios be predictable using simple, low complexity measurements. In an 802.11a testbed, the the model is shown to be very accurate (correlation of 0.97 between analysis and measurement). This is a positive result in that it greatly reduces the complexity of obtaining the interference map for one channel, and one card per node. The second contribution is showing that the complexity of the complete interference map is much higher than previously thought: each card and each channel have in fact to be measured independently. This latter result is rather pessimistic, as it implies that in a real world setup when nodes have multiple cards, and are supposed to use orthogonal channels, the complexity of producing the complete map is prohibitive for dense networks of even moderate scale.

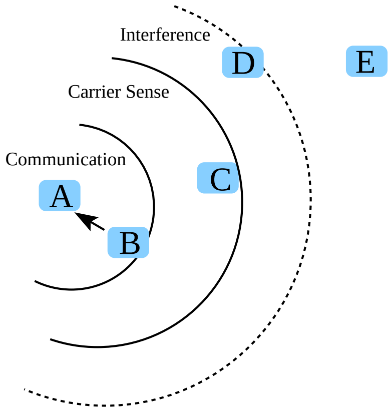

Figure 1: Possible interference relationships.

2. INTERFERENCE MAP

If we consider two nodes, there are several relative positions of interest with respect to each other (Figure 1). If they are close enough, like A and B, they are in communication range, meaning that packets sent by B can be received at A. The actual distance depends on conditions, carrier frequency, and output power. For example a 5004 MP Atheros a/b/g card claims 300m outside and 30m inside when operating in 802.11b at 11Mbps. 'Distance or range' in this description are just convenient terms because in reality delivery ratio decreases from 1 to 0 in a progression that describes a donut around A, rather than a circle. The circular shape is of course true only in void, whereas indoors the shape of coverage is highly irregular. The next range of interest is that of carrier sensing (CS). The carrier sense range includes the communication range, regardless of their actual shape. When A and C are this far apart, even if C cannot send packets to A, its carrier can be sensed at A which backs off when C has a transmission of its own. A will then defer transmission so that it doesn't destroy packets originated by C (if it is the case, see next paragraph). Both communication and CS can be asymmetrical: B can send to A, but not vice-versa, or C can sense A, but not vice-versa.

The next range of interest is the interference range. This range is not in a fixed relationship with the other two ranges as indicated in the figure, but rather depends on the source as well: the interference range is defined for the ordered tuple (B, A), in which B is the source and A the destination. Assuming that B and D are outside CS range of each other and therefore send packets at the same time, the question is if A is able to receive B's packets. The interference range is then defined by all positions of D that destroy some of B's packets at A. If B's power at A does not exceed D's power at A by some capture threshold, then B's packet is lost. In literature, when D is outside CS range of B, but in communication range of A, it is called a hidden terminal. In this paper however, we refer to all D nodes that destroy packets at A as hidden terminals or interferers: close hidden terminals are ones inside CS range of the destination, whereas remote hidden terminals are outside.

Finally, if far enough apart, station E is completely out of interference range of A, meaning none of its packets can interfere with packets arrived at A, regardless of B's position. As shown recently [2], such nodes may together generate enough power to have non negligible effect over B→A, but the occurrence of this situation is fairly rare. Interference for the link B→A is defined as the cumulative effect other nodes in the network have on throughput achievable in that link. Based on the previous definitions, some nodes may reduce B's sending capacity by being in its CS range, or may destroy packets after they arrive at A. We therefore want to differentiate between sending and receiving interference because they are qualitatively different: while sending interference caused by CS is nondestructive, the receiving interference destroys packets that require retransmission. When we say that sending interference is nondestructive, we mean not that it is beneficial, but that it grabs the resource (the carrier) to support some useful transfers, whereas receiving interference is much more wasteful in that it corrupts packets at the destination, after the air resource has been used already.

Having a large population of wireless devices in an area operated on a small number of channels brings the question of how devices on the same channel interact. Any two devices are in one of the four situations described above, but the amount of interference they create depends on the amount of traffic they carry, and is therefore linked to problems of routing, load, call admission and channel allocation. The interference map is a data structure that characterizes this interaction, so that one can answer questions like: given a channel allocation and routing, is a particular traffic matrix supported? Can an additional call be accommodated? When a link goes down, can the current service be maintained? What is a channel coloring that favors particular patterns of traffic (tree, mesh)?

The information in the interference map has three disjoint, but dependent parts: delivery ratio matrix, the carrier sense matrix, and the hidden terminal relationships. The delivery ratio matrix describes the capacity for each pair of nodes in the network. This includes effects of SNR degradation because of local geography, fading, multipath, as well as external external interference sources. The carrier sense aspect governs **what can be sent** into the air, which is the first step in getting the data across in the wireless network. The hidden terminal aspect is the actual interference information and describes **what can be received** at a destination under interference from other nodes.

The interference map of n nodes on the same channel consists of:

- d_i^k delivery ratio from node i to node k without any interference, $1 \leq i, k \leq n$
- cs_{ij} what fraction of the maximum capacity node i can put on air, when j is active at maximum capacity as well, $1 \leq i, k \leq n$
- d_{ij}^k delivery ratio from i to k with interferer j sending at maximum capacity

In addition, we introduce the following notations:

- s_i^k traffic sent from i to k
- $s_i = \Sigma s_i^k$ all traffic flowing out of node i,$\forall k$ neighbor of i

All these values are normalized to the interval [0,1], and they can be easily obtained from throughput measurements as we detail in the next section. Traffic sent out is divided by the nominal capacity, while delivery ratios are directly measured as throughput of broadcast traffic, divided by nominal capacity. For example, $s_i^k = 0.3$ would mean that node i sends at 30% of maximum capacity. $d_{ij}^k = 0.7$ means that in the presence of interferer j the throughput $i \longrightarrow j$ is 70% of the maximum supported by the channel.

2.1 Carrier sense (sending) interference

When two nodes sense each other they share the medium completely, and the sum of their maximum output rates on to the air is 1. When they are out of CS range of each other, each of them can send at full throttle, yielding a total output rate of 2. Any value between 1 and 2 is possible, because CS is not a symmetric or discrete phenomenon - one node may sense the carrier from a source only a fraction of the time, or the CS may behave asymmetrically. cs_{ij} is an $n \times n$ mostly symmetric matrix describing the capacity that two nodes can put on air when sending at the same time. $cs_{ij} + cs_{ji}$ represents the total capacity placed on the air ranging from 1 to 2. To gather cs_{ij}, we send broadcast traffic at full throttle from nodes i and j. We then use the packet rate reported as sent by each node - producing cs_{ij} and cs_{ji}. The complexity is $O(n^2)$, where n is the number of nodes involved.

Matrix cs_{ij} can also be seen as a directional CS graph: $cs_{ij} = 1$ indicates no link from i to j in the the CS graph, because i is sending at full throttle even when $j$is active. When $cs_{ij} + cs_{ji} = 1$, there are two bidirectional links - each node pointing to the other. $cs_{ij} + cs_{ji} = 1.5$ usually indicates a one direction CS link: one node has an output of 1, and the second of 0.5 because the first node doesn't hear the second. The direction of the link in the CS graph indicates the direction of sensing. To compute how much of that traffic can actually be received is the role of receiving interference mapping, which is described in the next section.

2.2 Hidden terminal (receiving) interference

This component is sometimes described in literature by a model called *conflict graph* [7]. The conflict graph indicates which groups of links mutually interfere and hence cannot be active simultaneously. In this paper, we quantify receiving interference by emphasizing the interferer in isolation, without considering him as part of a link. This model is appropriate when considering multiple interferers, regardless of which their destinations are.

The purpose of the receiving interference map d_{ij}^k is to have an estimate of the effect a remote source j has over traffic sent from i to k. When i and j are in CS range, they share the medium, and j does not destroy i's packets at k. However, when j does not sense i's transmission, packets received at k may be garbled - this is known as the hidden terminal problem. Example: when j is silent and i sends, throughput $i \longrightarrow k$ $d_i^k = 0.8$. When i and j send, throughput $i \longrightarrow k$ becomes $d_{i,j}^k = 0.4$. Conclusion: traffic leaving j produces a degradation of 50% for traffic $i \longrightarrow k$. We collect the measurements d_i^k for all pairs (i, k) and $d_{i,j}^k$ for all triplets (i, j, k) in the network. The complexity of collecting the entire data set is $O(n^2)$.The measurement process has these steps (this procedure was first proposed in [1]):

1. node i alone broadcasts and all other nodes k record d_i^k = throughput of $i \longrightarrow k$. All nodes take their turn in broadcasting - complexity $O(n)$
2. nodes i and j broadcast simultaneously and all other nodes k record: $d_{i,j}^k$ = throughput of $i \longrightarrow k$ jammed by j and $d_{j,i}^k$ = throughput of $j \longrightarrow k$ jammed by i. All possible pairs (unordered) take their turn - time complexity is $O(n^2)$.

The storage requirement for the second step is $O(n^3)$ as there is one interference measurement stored for each ordered (i, j, k) triplet in the network. For the first step, the required storage is $O(n^2)$ - throughput measurement for each directed link in the network. Note that these are in fact upper bounds for when the communication ranges (for step 1) and interference ranges (for step 2) of all nodes extend over a constant fraction of the entire network. In reality, the number of nodes that can produce interference at a destination are limited to a donut shaped region around the destination, the radius of the region being dependent on the hardware, bitrate, antenna, etc. In this case, the storage complexity could be reduced to $O(dn^2)$ where d is the degree of the node, or some other spatial density measure. Another factor in reduction of the complexity of measurements is the fact that not all links of the network are interesting, as will be seen in the experiments section.

This measurement procedure can produce a reasonably accurate image of what happens when a triplet (i, j, k) is involved in a sending/jamming process. But in reality, there are several nodes sending at the same time on the same channel, and one node's traffic is another node's interference. Having analyzed the complexities of measurement and storage for one interferer scenario only, it becomes clear that a measurement based approach is not scalable to the entire network: arbitrarily large groups of nodes can send at the same time, effectively jamming each other on the same channel. What is needed is a method that can predict the effect of several interferers acting simultaneously from single interferer measurements.

2.3 Analytical model

The model we propose comprises of the following two relations, the first expressing the limit for sending, and the second one the limit for receiving:

$$s_i + \sum_{j \in CS(i)} s_j \quad < \quad 1 \tag{1}$$

$$d^k_{i,all} \quad = \quad d^k_i \prod_{j \in I(i \rightarrow k)} [1 - (1 - d^k_{i,j}) s_j] \tag{2}$$

In the first condition (1), sending capacity of node i is limited by contention with all nodes it has to defer to. These are the nodes towards which i has a directed link in the CS graph.The second equation (2) models the delivery ratio of link $i \rightarrow k$ when all its interferers are active. Each interferer j contributes with an amount of interference that is measured separately as $d^k_{i,j}$. This property of independence between different interferers makes the procedure scalable as $d^k_{i,j}$ can be measured in $O(n^2)$ time for the entire network. If this independence wouldn't hold, a complete interference map would have to measure each possible group of interferers which at run time might affect the capacity of link $i \rightarrow k$. In the worst case, this is the power set of all nodes, of exponential size. The s_i factors in these equations represent the sending rates at the current node i and its neighbors j. These rates are considered known for the entire network, as stated by conditions 1 and 2 in the introduction. To understand the rationale of this second equation assume that $s_j = 1$ meaning that all the interferers send at full capacity. In this case $d^k_{i,all}$ becomes $d^k_i \prod_{j \in I(i \rightarrow k)} d^k_{i,j}$, showing that the final delivery ratio is merely a product of delivery ratios achieved when each interferer acts in isolation.

These relations can be used in any heuristics for solving problems that implicitly depend on interference: routing, call admission and channel assignment. For example, a call admission decision should first use the first relation to assess whether the proposed new traffic can be sent onto the air. There is no point sending voice traffic that is dropped even before it leaves the access point, so a call should not be admitted unless sending capacity for all nodes remains valid under condition (1). The second relation would then estimate the delivery ratio achievable across various links, using the measured values of $d^k_{i,j}$, and the s_i values accepted by the first step.

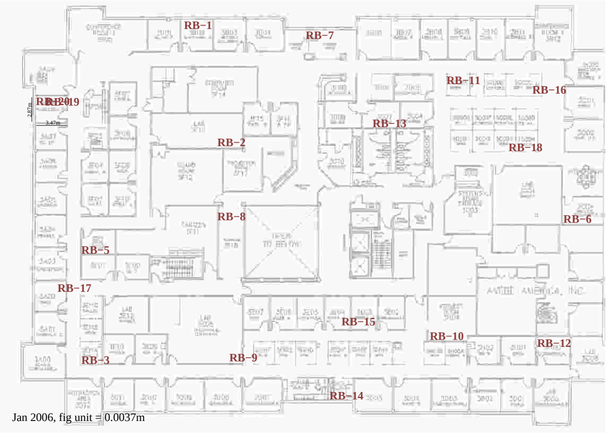

Figure 2: 802.11a testbed: 20 nodes in a 45m x 60m building.

Alternately, for a route or flow optimization procedure, these relations would participate as constraints in the optimization process (albeit nonlinear).

The rest of the paper is devoted to validate equation (2), namely confirming the fact that effect of different interferers can be measured independently and used to predict complex scenarios with several interferers. In the next section, we present several experimental results that explore the dependence of packet delivery ratio on several variables: sending rate, jamming rate, and number of interferers, distance, actual card used, actual channel used.

3. EXPERIMENTAL RESULTS

3.1 Testbed setup

We use a 20 node testbed deployed in a 45m x 60m building (Figure 2). Each node is equipped with two 802.11a/b/g cards tuned to 802.11a, running Linux with `madwifi-old` driver for Atheros chipsets. In order to cover the entire building, we use the lowest bit rate setting (6Mbps) which allows the longest range indoors. We employ Click modular router [12] to generate broadcast traffic for all the measurements: delivery ratio, carrier sense and interference. One particular feature that is needed for the measurement of carrier sense is the tx feedback: the driver gives a report for each packet submitted to the card - whether it was ACK-ed successfully, retried to the maximum and dropped. For broadcast packets the feedback only says that the packet made it on the air, as there is no ACK or retry. This allows for each node to directly measure its access to the medium, without the need of other receivers. For all measurements, we collect rates in packets per second and divide them by the nominal capacity of the channel, so that all the values handled are between 0 and 1: delivery ratio, traffic sent out, damage produced by an interferer, etc. Broadcast at all bitrates is possible with `madwifi`, so this measurement method is not limited to the basic 6Mbps rate.

We run the procedure outlined in section 2.2, accumulating the structures cs_{ij}, d^k_i and $d^k_{i,j}$ for all ordered triplets (i, j, k). For n nodes, bw^k_i is an $n \times n$ matrix, usually asymmetric, containing the throughput from node i to node k. Knowing the maximum capacity $lcap$ of a link, and assuming that bidirectional communication is usually necessary, only links with good delivery ratio are considered for the rest of the experiments.

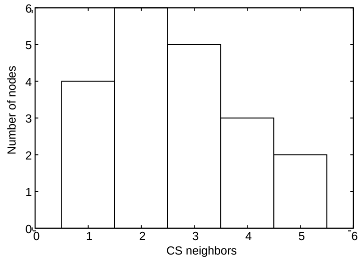

Figure 3: Histogram of carrier sense (CS) degree of nodes - on average, 2.6 nodes are within CS range in a mesh of 20 nodes.

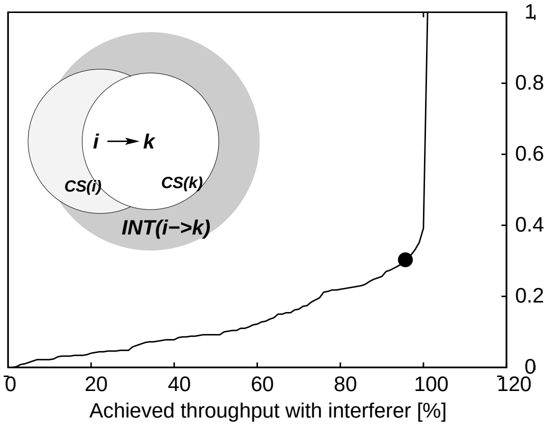

Figure 4: Inset: relative positions of CS and interference areas. Graph: CDF of throughput achieved in the presence of all possible remote interferers. These are the ones outside the CS range of both the source i and the destination k, therefore cannot be detected by ether the source or the destination. 70% of the potential interferers allow the link to function at 95% or more of its capacity, but this includes nodes outside the interference range of the link $i \rightarrow j$. The other 30% of jamming situation produce sizable damage on the capacity of the link.

- $d_i^k = \frac{bw_i^k}{lcap}$ is the delivery ratio for each directional link, normalized to the interval [0,1].
- $ETX_{ik} = \frac{1}{d_i^k d_k^i}$ metric describing quality of a bidirectional link - the expected number of transmissions required to send a packet from i to k or from k to i.

Assuming an ETX [13] value of 4 as a cutoff point, we are left with only 19 bidirectional 'good' links (a particular example of a bidirectional link with ETX=4 is one for which delivery ratio in both directions is 0.5). $d_{i,j}^k$ is then a matrix $p \times n$, where p is the number of interesting unidirectional pairs - 38 in our setup. Each value in this matrix represents the capacity of the pair in the presence of the interferer.

The cs matrix contains a directional CS graph as described in section 2.2. In Figure 3, we see how the number of CS neighbors is distributed among the 20 nodes - the average CS degree is 2.6 (compared to the average node degree with the 'good' links of 1.9). These sources of interference are not the most dangerous, since nodes in the carrier sense range take turns in sending packets, as opposed to nodes in interference range (hidden terminals).

The most critical question for any link is how many interferers are out there, and how bad are they? In Figure 4, we look at potential interferers for all 'good' links. The inset picture shows a source i and a destination k with their respective CS ranges. The large gray circle around the destination i labeled $INT(i \rightarrow k)$ represents the area of potential interferers that can affect the transmission $i \rightarrow k$. We plot the CDF of all $d_{i,j}^k$ for the selected 'good' links, and all their potential interferers . The CDF does not include interferers which are in CS range with the sender or the receiver $CS(i) \cup CS(k)$, but does include nodes which are far away from both sender and receiver, outside of $INT(i \rightarrow k)$. This last category is the largest, as we see that more than 70% of the potential interferers allow links to operate at more than 95% capacity. The rest are real remote hidden terminals producing real damage on the links - packets which are garbled at the destination.

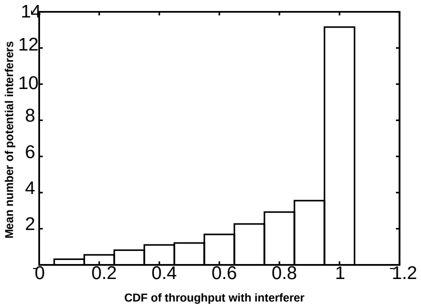

Figure 5: Cumulative distribution of the number of possible interferers and the amount of damage they produce. Nodes outside interference range are not included. On average, there are 2 interferers which reduce the capacity of the link to 60% or less.

Figure 5 shows the number of potential interferers and what effect they have on the link. Again, a large number of nodes (13 out of the total of 18 possible interferers) leave the capacity almost intact, meaning that the interference range covers about a third of our 20 node network. The real hidden terminals however are quite present as well: there are on average 2 hidden terminals which reduce the link capacity to 60% or less. From the statistics we eliminated non interferers outside $INT(i \rightarrow k)$. From the remaining, we also eliminated nodes which are in CS range with the sender $CS(i)$, but allowed the ones which are in $CS(k)$: these are all the effective hidden terminals (close and remote), contained in the region $INT(i \rightarrow k) - CS(i)$. In Figure 6 we also eliminated interferers which produce less than 5% damage, to have a closer look at the worst offenders. The sum of the first two bins in this histogram shows that there is on average one interferer which reduces the throughput to 20% or less.

These statistics confirm that remote hidden terminals are a significant presence even in a sparse wireless network like ours, with an average degree of 1.9. These numbers are likely to be much worse in a better connected network, as all the regions surveyed here would be more populated. These remote hidden terminals will adversely affect routing, channel allocation, and call admission, as all these issues directly influence the amount of traffic on each link. Previous work [1] found that most interference is in the form of carrier sense, and we attribute this to hardware / software differences: examining interference map (cs_{ij} and $d_{i,j}^k$ structures) we found that the carrier sense range is almost perfectly overlapped over the communication range for our hardware/software configuration. This means that there are almost no cases when nodes are in CS range, but no packet can fly across (node C in Figure 1). The more important aspect however, is that interference range starts im-

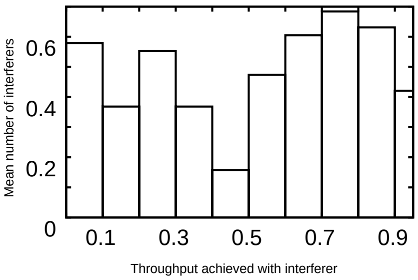

Figure 6: Histogram of the number of interferers for each interval of achieved throughput. interferers allowing more than 95% of the throughput are omitted. The sum of the first two bins shows that there is on average one interferer reducing the capacity to 20% or less.

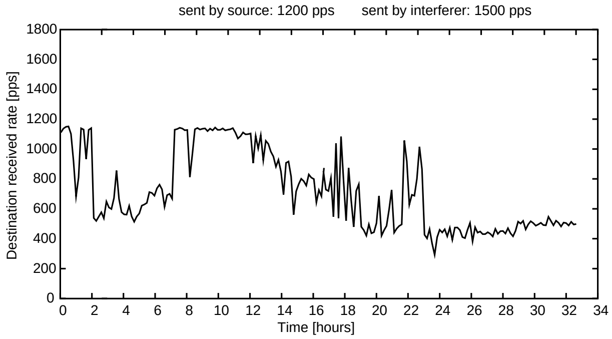

Figure 7: Time line: due to external factors, interference measurements taken at different times cannot be compared. We use a round robin scheme to alternate between measurements and identify stable periods, which allow for meaningful comparisons.

mediately beyond communication range, which makes this testbed appropriate for the study of remote interference - area designated by node D in Figure 1 is quite large. Two recent contributions [10, 11] proposed modeling of the interference only on the basis of packets which are successfully received between stations, easily achieving $O(n)$ complexity for measuring the interference for the entire network. But this clearly ignores remote hidden terminals which are out of the CS range of both sender and receiver, which have a considerable effect even in sparse networks as ours.

3.2 Interference properties

After having established the extensive presence of both close and remote hidden terminals, we set to explore in more depth equation (2). Some of its more useful features are the linearity with respect to interferer rate, shown by the presence of s_j inside the product, and the linearity with respect to source rate, which is implied by the absence of s_i. For these measurements we want to compare the achieved throughput for different source/interferer rate, but how relevant is this comparison if the measurements are not taken at the same time? In most situations we want to compare configurations that cannot possibly be ran at the same time because they inherently affect each other - and this is always the case with interference measurements. In addition, the wireless medium is highly variable indoors, depending on the level of human activity. Both these factors make the measurement of the interference difficult to setup, reproduce, and interpret.

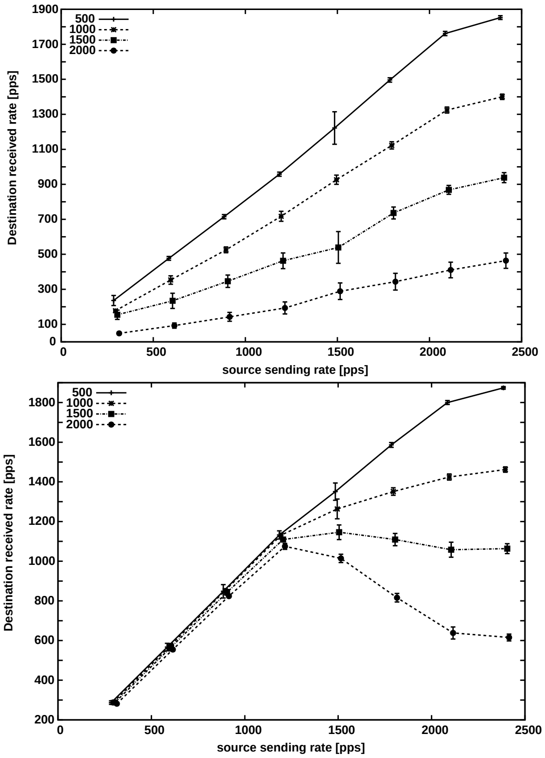

Figure 8: Delivery rate with one interferer. Top: in most cases, the achieved throughput in packets per second is linear with the sending rate for the entire range of sending rates. Bottom: low rate flows from the source are not affected by the interferer, but for higher rates, the behavior is still linear.

3.2.1 Measurement methodology

The methodology we use is to have a round robin of short runs for each experiment (E_1, E_2, E_3, E_1, E_2, ...) over longer periods of time in order to identify stable periods during which external conditions do not vary much and measured values show some stability. During those periods, we may compare the results of experiment E_1 with the results of experiment E_2, even if they are not run exactly at the same moment, assuming that conditions were comparable since each measured value shows stability. This method of comparison is important especially for the experiments which would conflict over resources. One example is using the same channel by two cards over the same period of time, as in section 3.2.5. Another one is testing the same source-destination pair of cards over different channels as in section 3.2.6. In all the experiments, there is a time sharing between the experiments so that each experiment has exclusive use of the resource, and yet its result can be compared with a virtually parallel experiment using the same resource.

For example, we setup three nodes - source, destination, and interferer to operate on the same channel in 802.11a, using the 6Mbps rate to send 200 byte packets in broadcast mode. The channel capacity for this setup (packet size, bit rate, SNR) is about 2300 pps (packets per second) for broadcast packets - no ACK, and no retry. In addition to other measurements mentioned below, we record the packet rate received by the destination under the conditions that the source sends 1200 pps, and the interferer sends 1500 pps. In Figure 7, we follow the delivery ratio at the destination over a period of 33 hours and attribute the high variation to external factors (mov-

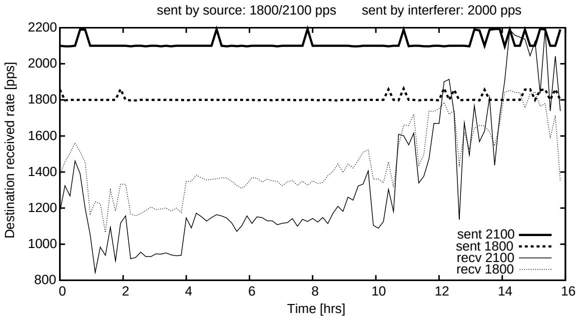

Figure 9: Anomaly: packets sent at higher rate use a lower signal strength, yielding in a lower delivery ratio. This behavior is persistent for many hours.

ing people, doors). There is no institutional use of 802.11a in our building, and to the best of our knowledge there is no unaccounted traffic on the channels used. While there are large variations in the delivery ratio, we used times 8-12 and 24-32 as relatively stable periods to investigate for our purposes. In fact, all the following measurements were performed during the above time line, virtually time-sharing with the experiment in Figure 7.

3.2.2 *Linearity with source and interferer rates*

In the Figure 8(top) we vary on the horizontal axis the sending rate of the source from 300pps to 2400pps and measure the delivery rate for four packet rates of the interferer: 500,1000,1500, and 2000, represented by separate lines in the graph. Each point also shows the standard deviation over all the samples used. Because curves are mostly straight, we infer that delivery ratio is stable for different sending rates of the source. For example a delivery ratio of about 83% is maintained for the top curve when the source sends between 300 and 2100pps, and the interferer sends at 500 pps. The source and the interferer are confirmed to be outside each other CS range by periodical verification of sustained simultaneous output of 2300pps. Figure 8 bottom corresponds to period 24-32, and Figure 8 top to period 8-12, which we deemed as stable for the purposes of comparing results. For period 8-12, we can see that the linearity with respect to sending rate is not followed anymore, especially for low packet rates. Specifically, when the source sending rate is below 1200pps, the transmission is not affected by the interferer. The anomaly we believe is caused by a behavior of the Atheros chipset which sends weaker packets when the packet rate is high. A separate 16 hours experiment measuring delivery ratio between another source and a destination (Figure 9) shows that for longer than 10 hours, the higher packet sending rate (2100pps) consistently gets a lower delivery ratio than the lower sending rate (1800pps). We found that there were 5dB more for the signal strength of the slower rate, thus justifying the results shown at the bottom of Figure 8. This and other nonstandard features of Atheros based chipsets are confirmed by other researchers [14], and mostly explained by aggressive power saving implementations.

Fortunately, this behavior is sporadic, and in most cases we can observe the linearity of the interference with respect to the sending rate of the sender. Although somewhat visible in Figure 8, the linearity with respect to interferer rate is plotted in Figure 10 for two different interferers. The source sent 2300pps for all experiments, whereas the cumulative interferer rate is shown horizontal axis. For the middle curve both interferers sent simultaneously each with half the rate, showing that independent measurements for each interferer can be used to derive the effect of several interferers sending concurrently. This linear combination of effects of different interferers is also maintained even in the case of anomalous delivery cases mentioned in the previous paragraph, as shown in Figure 11. One of the interferers produces almost negligible damage, while the second one is worse for high rates at the source, but their combined effect is piecewise linear.

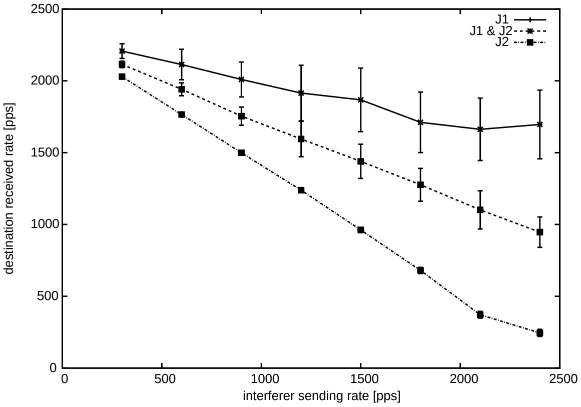

Figure 10: The amount of interference is linear with respect to interferer rate. Separate interferers acting simultaneously also create interference that is linear with their combined rate.

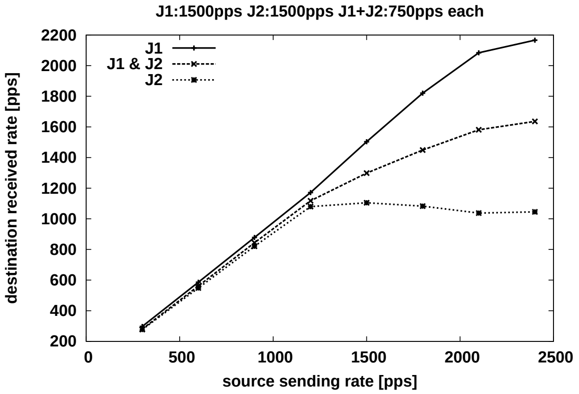

Figure 11: Even with anomalies in power of emitted packets, effect of interferer is linear with respect to sending rate for single interferers, and for combinations of two.

3.2.3 *Independence of different interferers*

The other crucial aspect of the model we propose (equations (1) and (2)) is the fact that interferers have effects that are independent of each other. Basically, the probability of a packet being delivered in presence of several interferers is the product of delivery probabilities when the interferers act in isolation. This is the major reason why the network wide interference can be characterized with a small number of measurements: the ability to combine simple measurements for isolated interferers to predict the effect of possibly every node sending, as it happens in a network under normal use.

In a separate 34 hour experiment we verified that the delivery ratios with two different interferers can be treated as independent variables across various delivery ratios in time. We send data from a source S and two interferers J1 and J2 at the maximum capacity for several situations:

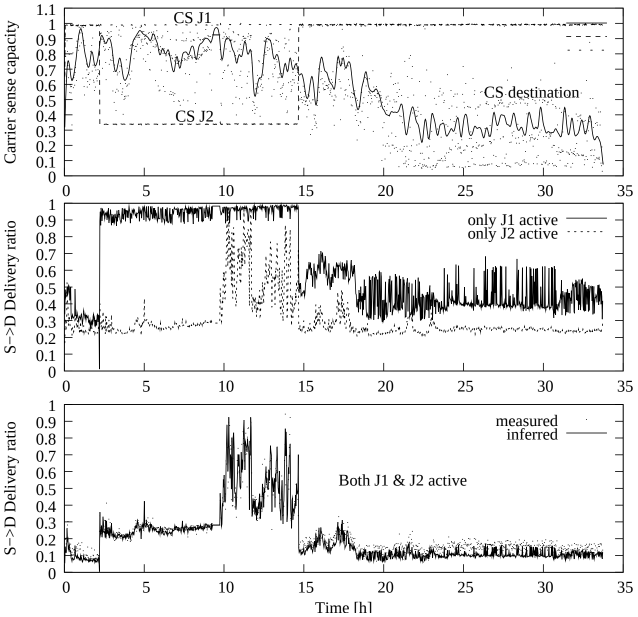

Figure 12: Top: History of access to the medium for the interferers and the destination. Middle: delivery ratios sampled independently for each interferer. Bottom: When both interferers are active, measured delivery ratio confirms the independence of the two interferers .

- S→D: verifies the nominal capacity of the link
- S, J1 → D: measures the capacity with J1 alone
- S, J2 → D: measures the capacity with J2 alone
- S, J1, J2 → D: measures the capacity with both interferers, monitors the CS between S,J1, and J2
- J1, J2, D: monitors the CS relation between the J1, J2, and D

The placement of the source, destination and the two interferers is set such that the source cannot sense the carrier of any of the two interferers, so it is always sending at full throttle. On the destination side, interferer J1 is far enough not to defer to any node (Figure 12 top). The second interferer has a clear deferral period, and an independent period. The destination (shown with dots and also with a smoothed graph, because of high variation) experiences a more ambiguous situation with respect to the interferers in which it senses either none, one, or both of the interferers throughout the course of the experiment. The middle of Figure 12 shows the delivery ratio achieved with each interferer independently, ranging from 20% to 100% depending on the time, and position of the interferer. We chose this scenario from a larger set of experiments with similar placement of nodes because its diversity makes it appropriate for verifying the independence assumption between the two interferers.

In the bottom figure, we plot the measured delivery ratio with both interferers active, together with the product of the delivery ratios measured separately. We see that the modeled value is highly correlated (0.96) with the measurement, indicating real independence. More than actual correlation, the value inferred as the product of the separate delivery probabilities closely tracks the measured delivery probability with J1 and J2 active.

Finally, we verify the accuracy of all aspects of equation (2) in a 3.5 hour long experiment that is similar to the one in the previous paragraph, except that both the rate of the source and the rates of the interferers are selected randomly and independently in the interval $[0..1]$, from a uniform distribution. The delivery ratio for the non interfered link is around 95% for the entire length of the experiment, but with the presence of both interferers, it can drop below 5% as seen in Figure 13. The analytical model using delivery ratios measured for independent interferers $d^k_{i,j}$ closely follows the values measured for two simultaneous interferers. Although it tends to overestimate the delivery ratio for higher values, it has a good correlation (0.97) with the measured data. For the entire duration of the experiment the interferers are out of the carrier sense range of each other and of the receiver. After each interference measurement, a separate broadcast of both interferers and the receiver was run to confirm that the interferers are able to send at maximum throughput, therefore do not defer to the receiver or each other. The receiver however is carrier sensing one of the interferers for the duration of the experiment. This shows that a combination of interferers can be modeled with around 5% error in delivery ratio when only measurements for individual interferers are available.

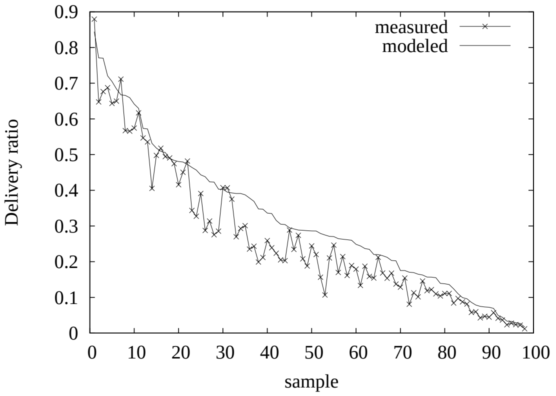

Figure 13: A source and two interferers send data at random rates - we sort the experiments based on delivery ratio achieved. Using equation (2), we predict the delivery ratio based on separate measurements for each interferer. The correlation is quite high (0.97) but the model overvalues the delivery ratio by about 4.5%.

After the validation of equation (2), which is the main way of reducing the measurement complexity of the interference map for one channel, we turn to other aspects of interference: first, we show that an approximation of the cs_{ij} structure can be inferred if positions of the nodes are known. Second, we look at the complexity of the interference map for the multiple card case.

3.2.4 Correlation with distance

Since producing the interference map has high cost ($O(n^2)$ network down time), it would be desirable to produce at least a good approximation of it by some cheaper method. Knowing that in theory interference range is linked to the communication range, we want to see how predictable the interference is for our particular indoor setup with respect to distance. In most static networks when access points are deployed in a building, a map association is usually available so that distances can be estimated with reasonable precision from a map drawn at scale. We used a map of our building as the one shown in Figure 2, and assigned coordinates to each node based on its relative position, and using the known dimensions of the building. We then associated each measurement d^k_i with the distance between i and k. In Figure 14, we see that delivery ratio

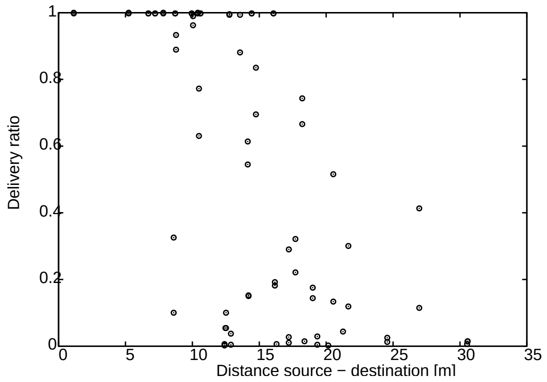

Figure 14: Delivery ratio is weakly correlated with distance.

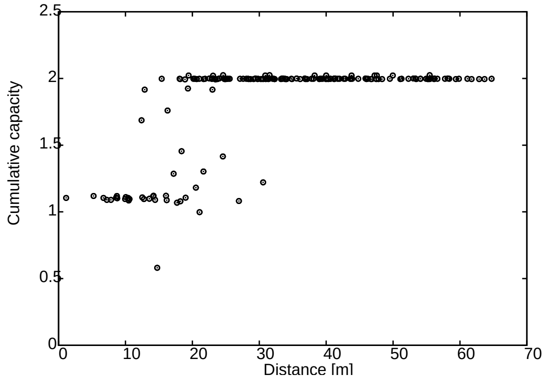

Figure 15: Carrier sense depends strongly on distance.

is very weakly correlated with distance, and this corroborates well with other findings in literature showing that delivery ratio (and also signal strength) and distance do not correlate well [15]. For carrier sense however, distance is a much better indicator - Figure 15 plots $cs_{ij} + cs_{ji}$, so that a value of 1 indicates carrier sense, while a value of 2 indicates independence. CS on 802.11a 6Mbps shows a well defined threshold at 18m for indoors 802.11a. This means that at least part of the $O(n^2)$ complexity can be avoided by getting an estimate of the CS graph based on the distances between access points or mesh nodes.

In Figure 16 we look at the relation between the damage produced by hidden terminals and distance. The damage produced to communication $i \rightarrow k$ by a hidden terminal j is computed as $1 - \frac{d^k_{i,j}}{d^k_i}$. Interferers in CS range of the source or destination are excluded. Although the amount of interference and distance are correlated (cor=-0.61), the correlation is not strong enough to produce a prediction based on distance. For example, for a hidden terminal at 30m, we cannot really say what amount of damage it will cause. We also examined correlation of the damage with ratio of distances $(\frac{jk}{ik})^2$ as classical communication theory would indicate, but the obtained correlation is weaker (cor=-0.29).

As mentioned in section 2.2, the $O(n^2)$ complexity is driven by both measurements of carrier sense and of hidden terminals. The conclusion of these distance based statistics is that while we can infer a non-directional CS graph just by using coarse node positions, the hidden terminals effect is not sufficiently correlated with distance. Therefore, pairwise measurements would still be necessary to measure receiving interference, so in the process, they might collect the CS directional graph as well.

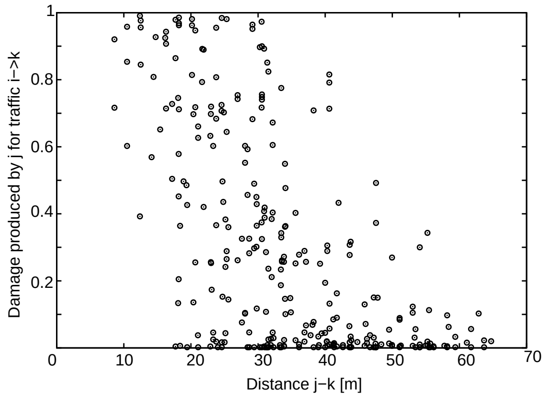

Figure 16: Effect of the hidden terminals is correlated with distance, but not strong enough for a prediction.

3.2.5 Consistency across interfaces

In all the experiments so far we considered that all nodes have one available wireless card, tuned to the same one channel for the entire mesh. For scalability reasons however, it is desirable that each node use several cards tuned to different channels. The allocation of channels and channel reuse will clearly affect the amount of interference, and this is one of the main targets of producing an interference map. A solution to channel allocation assigns a channel to each card so that connectivity is preserved and higher throughput becomes available. Several solutions have been proposed in the literature [4, 16, 17, 18, 19], but all implicitly assume that links may use any card in the same machine, implicitly assuming they are equivalent in terms of their interference patterns.

To verify this hypothesis, we measured delivery ratios for each pair of nodes in a group of 4 nodes for a total duration of 110 hours spread over a period of two weeks. Two sets of measurements were taken for two 'parallel' networks - one created over *ath0* interfaces, and the other over *ath1*. In more than half of the links measured, the delivery ratio across one interface is completely different from the other one in both quantity and quality, even if the channel used is the same. In figure 17 we follow one particular pair of parallel links with series of measurements spanning the entire period, comprising of a mix of busy weekday mornings and quiet weekend nights. We can see that while the link across *ath1* interfaces shows solid performance across the entire period, the link across *ath0* interfaces ranges from acceptable to less than 10%, and from steady to highly variable in the samples taken. The differences in the other 11 directional links (obtained among 4 nodes) ranged from high variation to steady delivery and from maximum capacity to no link - in fact half of these links showed differences like those in Figure 17 or worse. Given that the wavelength of 802.11a is about 6cm, it is very likely that shadowing and multipath would create variation between points that are that close. Our antennas are spaced 40cm apart, and as the experiments confirm, there is very little correlation between the performances of cards in the same node.

These measurements show that between two communicating nodes, we cannot consider a logical link that can use any two pair of interfaces. In fact, each of the four physical links between two nodes has to be treated as a different link, and we conclude that f**or the purpose of the interference map, each physical link has to be measured separately**.

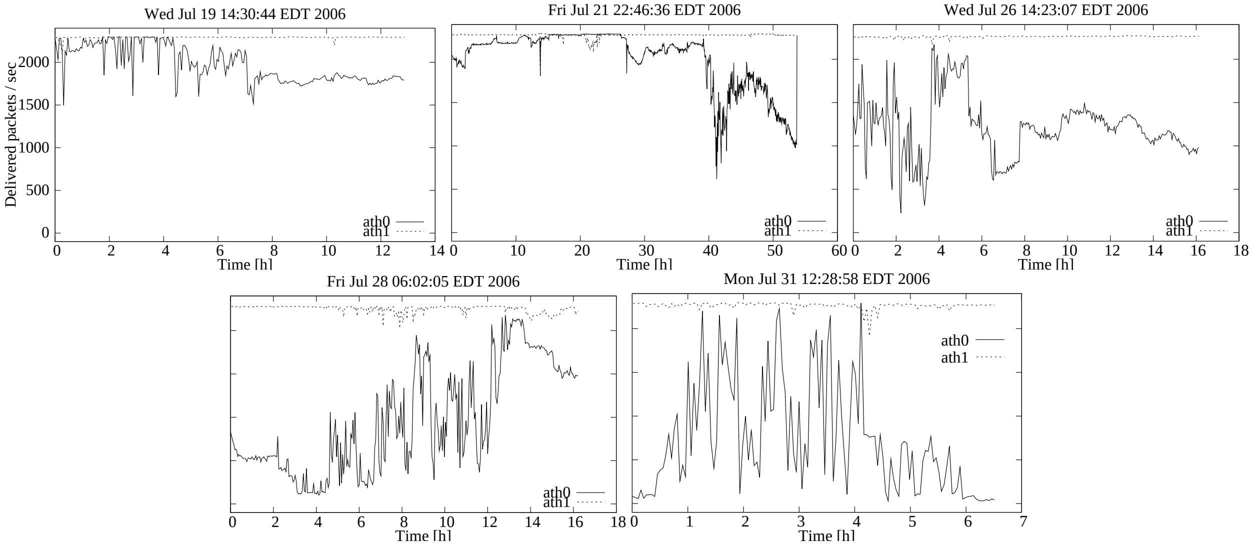

Figure 17: Delivery ratio sampled for a total of 110 hours spread over a period of two weeks. Delivery ratio differs widely across interfaces. Channel allocation algorithms may not assume equivalence of link performance based on sampling of links from a single interface.

3.2.6 Consistency across channels

Another assumption made by channel assignment solutions that employ variations of edge or vertex coloring schemes is that carrier channels are basically equivalent - so they can be assigned any color (channel). This assumption also turns out to be too optimistic, at least for the case of 802.11a. In Figure 18, we examine a 28 hour trace of delivery ratios across channels 36, 44, 52, 64, 149, 157 and 165 under the same conditions used in previous experiments. The last three channels (recommended for outdoors) performed worse, all having an almost negligible delivery ratio, barely visible next to 0; channel 64 had 100% performance and is not visible at the top. Channel 44 also shows a strong and consistent performance maintaining 80%-90% delivery ratio throughout the day. 36 and 52 showed periods of stability mixed with periods of high variation, but even during the stable periods, the performance across channels differs greatly. The high variation experienced by channels 36 and 52 is all happening when other channels show steady (high or low) performance. In order to validate these results, we ran additional measurements for different sets of nodes, different power settings, and with longer settling time after the channel switch, but are not including them for the sake of brevity. As in the card comparison, the consistency across channels is rather the exception than the rule, with very few cases in which a link performs the same across all frequencies. We conclude that **for the purposes of interference map, possible channels of any link have to be measured separately.**

4. DISCUSSION AND SUMMARY

The most relevant works in this field are [1], [2] and [10, 11]. We extend the work in [1] by clearly defining the interference map as a collection of: delivery ratios, carrier sense matrix, and hidden terminal matrix and by modeling the effect of multiple interferers. We characterize the complexity of gathering the map, and propose an analytical model to allow the use of pairwise measurements. The model reduces measurement complexity by using certain properties of interference: linearity with respect to source rate, interferer rate, and independence of multiple interferers. [2] looks at remote interferers (out of CS, that produce damage), and 'no impact' nodes which produce no damage. Their numerical results are properly captured by our model: remote interferers – can be combined in a linear fashion using relation (2); 'no impact' nodes – although they may become interferers when acting together, the occurrence is very rare. In addition, we show that the occurrence of remote interference is quite severe, phenomenon which is prone to rise with increase in deployment density. The remote interferers are ignored in [10, 11] by considering only the ones from which signal strength can be read. Signal strength however can only be used when packets are received properly, which means inside communication range. We model interferers outside communication range, which can be inside CS for sending interference, or inside interference range for remote hidden terminals (these are the causes of the $O(n^2)$complexity).

However, our model also has a few drawbacks:

- requires a global view of the network. This stems from the fact that interference has non local effects which we believe are best tackled in a global, centralized manner. However, many measurements can be performed in a distributed, asynchronous fashion. For example, delivery ratios d_i^k between nodes can be monitored passively on live traffic. Carrier sense and hidden terminal measurement require network down time to obtain d_{ij}^k, but they can be performed one triplet at a time, during periods of relative silence in an area between i and j.

- only models CBR traffic. This limitation is an attempt to eliminate the time factor from the model. Because each flow of traffic is a potential creator of interference somewhere else in the network, non constant flows such as TCP would create highly variable interference patterns for otherwise steady conditions. Voice networks handle only CBR traffic, and also have stringent loss requirements, so are good candidates for the controlled interference environment proposed here.

- is verified extensively only for 802.11a networks. If density of WiFi devices increases at the current pace, one way

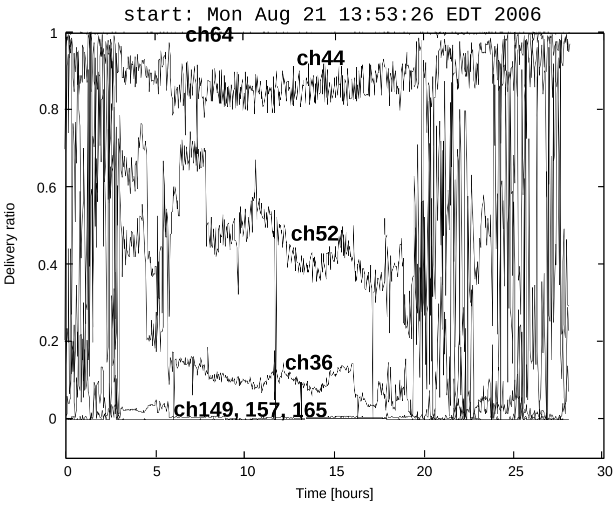

Figure 18: Delivery ratio differs widely across channels. Channel allocation algorithms may not assume that channels are interchangeable in terms of performance.

to increase bandwidth per area is by using more independent channels, and more cards. While 802.11b allows for only three orthogonal channels, most chipsets also suffer from electrical interference so that two cards must be at least 60cm apart, regardless of the channel. Depending on regulations, in 802.11a there are 12 orthogonal channels available and we found the electrical interference to be almost negligible.

- requires network down time. The very core of measuring interference is to quantify the effect other nodes have over a particular communication. Therefore any uncontrolled traffic has the potential of skewing the effects produced by an interferer: it can either increase the packet drop at the destination, or it can have the opposite effect, by contending for the medium with the interferer, and therefore produce a better delivery ratio at the destination.

The main factors driving the complexity of measurement of the interference map are the pairwise style measurements, and the asymmetries of the cards and channels. Pairwise measurements are the direct way of determining $d_{i,j}^{k}$ (the delivery ratio from i to k when j is interfering) when j is not in contact with either k or i. Since no packets from j can be received at k or i, no delivery ratio, or signal strength can be employed to determine the potential damage j can produce. In this case, generating traffic from j is a reliable, albeit costly way of gauging its effect over $i \rightarrow k$ communication. Obviously, for networks spreading over a wide area, only nodes in a circular region around the receiver are candidates for being remote interferers, so the complexity in these cases instead of $O(n^2)$, becomes just $O(n)$ - with a constant depending on the size of interference range. However, for setups that are small in area, but large in the population of wireless nodes, the potential interferers can be in a large fraction of the network. In our testbed, about one third of the nodes can interfere, so effectively, the complexity is still $O(n^2)$ – for the single card case.

The second cause of complexity is the asymmetry in card / channel behavior: we conclude that a more accurate estimation of the time complexity to obtain the interference map for all channels and across all possible links should be adjusted to $O(fcn^2)$ where c is the number of cards, and f the number of available orthogonal channels (the total number of experiments would be $O(fc^2n^2)$ but a node can run c at a time whem $f \geq c$). This is a considerable difference from the original $O(n^2)$, given the desirability of large number of cards to make use of the channel parallelism and decreased range/ increased density of 802.11a. To put this into perspective, a 20 node network, one card, one channel, 20 second measurements and associated overheads required for our experiments about 2.5 hours of network downtime - corresponding to $O(n^2)$. If we consider for example the case of a dual card 802.11a node ($f = 12$ and $c = 2$), the required downtime becomes unacceptable.

4.1 Summary

We proposed a model for interference in dense wireless networks that enables a low complexity procedure to collect the interference map in one card networks. We confirmed experimentally that interference measurements for isolated triplets of nodes (source/destination/interferer) can be used to predict the damage from several simultaneous interferers. The interference from distant interferers behaves linearly with respect to rate of the source and rate of the interferer, and shows independence between interferers. The most important result is that behavior of complex interference scenarios can be estimated based on measurements that have relatively low complexity $O(n^2)$, which could otherwise depend exponentially on the group size. On the negative side, measurement of the interference map faces asymmetries in the card and channel behavior, which make the complexity still prohibitive for dense multiple card networks.

5. REFERENCES

[1] J. Padhye, S. Agarwal, V. N. Padmanabhan, and L. Qiu, “Estimation of link interference in static multi-hop wireless networks,” in *Internet Measurement Conference*, 2005.

[2] S. M. Das, D. Koutsonikolas, Y. C. Hu, and D. Peroulis, “Characterizing multi-way interference in wireless mesh networks,” in *WiNTECH ’06: Proceedings of the 1st international workshop on Wireless network testbeds, experimental evaluation & characterization*, pp. 57–64, 2006.

[3] K. Jamieson, B. Hull, A. Miu, and H. Balakrishnan, “Understanding the real-world performance of carrier sense,” in *ACM SIGCOMM E-WIND Workshop*, 2005.

[4] A. Raniwala, K. Gopalan, and T. Chiueh, “Centralized channel assignment and routing algorithms for multi-channel wireless mesh networks,” in *ACM Mobile Computing and Communications Review(MC2R)*, vol. 8, April 2004.

[5] T. Nandagopal, T.-E. Kim, X. Gao, and V. Bharghavan, “Achieving mac layer fairness in wireless packet networks,” in *MobiCom ’00: Proceedings of the 6th annual international conference on Mobile computing and networking*, (New York, NY, USA), pp. 87–98, ACM Press, 2000.

[6] J. Li, C. Blake, D. S. J. De Couto, H. I. Lee, and R. Morris, “Capacity of ad hoc wireless networks,” in *ACM MobiCom*, (Rome, Italy), pp. 61–69, July 2001.

[7] K. Jain, J. Padhye, V. Padmanabhan, and L. Qiu, “Impact of interference on multi-hop wireless network performance,” in *ACM MobiCom*, (San Diego, CA), September 2003.

[8] M. Kodialam and T. Nandagopal, “Characterizing achievable rates in multi-hop wireless networks: the joint routing and scheduling problem,” in *MobiCom ’03: Proceedings of the 9th annual international conference on Mobile computing*

and networking, (New York, NY, USA), pp. 42–54, ACM Press, 2003.

[9] M. Kodialam and T. Nandagopal, “Characterizing the capacity region in multi-radio multi-channel wireless mesh networks,” in *MobiCom ’05: Proceedings of the 11th annual international conference on Mobile computing and networking*, (New York, NY, USA), pp. 73–87, ACM Press, 2005.

[10] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, “Measurement-based models of delivery and interference in static wireless networks,” in *ACM SIGCOMM*, (Pisa, Italy), September 2006.

[11] A. Kashyap, S. Ganguly, and S. Das, “A measurement-based approach to modeling link capacity in 802.11-based wireless networks,” in *ACM MOBICOM*, 2007.

[12] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek, “The click modular router,” *ACM Transactions on Computer Systems*, vol. 18, pp. 263–297, August 2000.

[13] D. De Couto, D. Aguayo, J. Bicket, and R. Morris, “A high-throughput path metric for multi-hop wireless routing,” in *9th ACM MobiCom*, (San Diego, CA), September 2003.

[14] G. Bianchi, “Surprises in experimental assessment of wireless LAN networks,” in *EXPONWIRELESS*, (Helsinki, Finland), June 2007. keynote speech.

[15] J. Bicket, D. Aguayo, S. Biswas, and R. Morris, “Architecture and evaluation of an unplanned 802.11b mesh network,” in *ACM MobiCom*, (Cologne, Germany), pp. 31–42, 2005.

[16] A. P. Subramaniam, H. Gupta, and S. Das, “Minimum-interference channel assignment in multi-radio wireless mesh networks,” tech. rep., Computer Science Dept, SUNY Stony Brook, 2006.

[17] B. Raman, “Channel allocation in 802.11-based mesh networks,” in *IEEE Infocom*, 2006.

[18] A. Sen, S. Murthy, S. Bhatnagar, and S. Ganguly, “Topology formation in multi-channel mult-radio mesh networks,” tech. rep., NEC Laboratories America, 2006.

[19] M. Alicherry, R. Bhatia, and L. E. Li, “Joint channel assignment and routing for throughput optimization in multi-radio wireless mesh networks,” in *MobiCom ’05: Proceedings of the 11th annual international conference on Mobile computing and networking*, (New York, NY, USA), pp. 58–72, ACM Press, 2005.