

Experimental Characterization of Sectorized Antennas in Dense 802.11 Wireless Mesh Networks

Anand Prabhu
Subramanian
Stony Brook University
New York, USA
anandps@cs.sunysb.edu

Henrik Lundgren
Thomson
Paris, France
Henrik.Lundgren@thomson.net

Theodoros Salonidis
Thomson
Paris, France
Theodoros.Salonidis@thomson.net

ABSTRACT

Sectorized antennas can increase wireless network capacity through greater spatial reuse. Despite their increasing popularity, their real-world performance characteristics in dense wireless mesh networks are not well understood. This paper conducts a systematic experimental study on a mesh network testbed using commodity 802.11 hardware and multi-sector antennas.

Our study results in the following main observations. (i) Sector selection should be based on explicit measurement in all sectors, though the measurement overhead can be significantly reduced by exploiting spatio-temporal characteristics of the best sector. (ii) Multi-sector activation typically reduces the signal strength of a link compared to single sector activations due to antenna design constraints. (iii) Spatial reuse is constrained by characteristics of antenna radiation pattern in different sectors (iv) Physical layer capture reduces the effect of directional hidden terminal problem. Finally, we discuss the implications of these observations on the design of practical directional MAC and topology control protocols.

Categories and Subject Descriptors

C.4 [Performance of Systems]: Measurement Techniques;
C.2.1 [Network Architecture and Design]: Wireless Communication—*Dense Wireless Mesh Networks*

General Terms

Experimentation, Measurement, Performance

Keywords

Sectorized Antenna, Dense Wireless Mesh Networks, Sector Selection, Directional Hidden Terminal Problem

1. INTRODUCTION

Today, numerous urban areas around the world are covered by dense 802.11 wireless networks that enable a wide

variety of legacy and emerging applications, ranging from WLAN Internet access to content sharing in community wireless mesh networks [1]. These dense deployments are typically unplanned and operate in unlicensed bands shared by other wireless technologies such as Bluetooth devices, cordless phones, and microwave ovens. As a result, interference becomes a major problem for 802.11 wireless networks operating in this environment.

Sectorized antennas aim to reduce interference and increase network capacity and spatial reuse through directional transmissions. In practice, they have been successfully used in cellular network planning [12]. In the enterprise 802.11 WLAN space, several companies ([2], among others) are now offering sectorized or directional antenna systems in their products, though their focus is on WLAN base station to client communication. In dense wireless mesh network scenarios, there is little evidence on the real world performance of sectorized antennas. This environment poses two main challenges. First, sector selection is challenging in environments characterized by non-line-of-sight, multipath reflections and high angular spread. Second, sectorization can limit the performance of the 802.11 CSMA MAC protocol due to directional hidden terminals. This problem has been identified in ad hoc networks and addressed by directional MAC protocols [5, 10]. These protocols have been evaluated only by simulations and, as we show later, some of their fundamental design assumptions do not hold in dense reflection rich 802.11 wireless environments.

The primary goals of this paper are to characterize the performance of sectorized antennas in real-world dense 802.11 wireless mesh networks and to derive practical design guidelines for directional MAC protocols and topology control mechanisms that operate in this environment. Directional MAC protocols [5, 10] achieve higher performance with single sector switching at packet-level time scale. On the other hand, topology control mechanisms [7, 11] activate multiple sectors at slower time scales at the potential expense of performance, but do not require complex modifications to the standard 802.11 MAC protocol. We conduct an experimental study using an 802.11 wireless mesh testbed equipped with sectorized antennas in an environment rich in reflections and multipath. The measurements cover a wide range of network topology and traffic configurations, including single sector and multi-sector activations, interaction of 802.11 MAC protocol with sectorization, directional hidden terminals and interference. These measurements investigate three main issues related to usage of sectorized antennas: sector selection, spatial reuse and directional hidden terminals.

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.

MobiHoc '09, May 18–21, 2009, New Orleans, Louisiana, USA.
Copyright 2009 ACM 978-1-60558-531-4/09/05 ...\$5.00.

Our measurement study has the following main results. First, received signal strength value is a robust metric to compare the performance of a link across different sectors as it is directly influenced by the radiation pattern of the sectors and has good correlation with higher layer metrics. Second, in a dense reflection rich environment, selecting the sector that geographically points to a node to communicate with it, can lead to poor performance. This contrasts the fundamental assumption in directional MAC protocols proposed for ad hoc networks. Although all sectors need to be probed to identify the best sector, we find that the best sector exhibits several spatio-temporal properties that can be exploited to reduce the measurement overhead. Third, contrary to the assumption made by current topology control algorithms, performance of a link can be radically different under single-sector activation and multi-sector activation due to antenna design constraints such as the antenna array factor [9]. Fourth, the increase in spatial reuse due to sectorization is primarily constrained by the antenna radiation pattern characteristics, such as the difference in maximum and minimum gain, and the particularities of the shapes of the main and side lobes rather than environmental effects, such as multi-path reflection. Fifth, physical layer capture reduces the effect of the directional hidden terminal problem. We identify conditions for sector selection to reduce directional hidden terminals while maintaining spatial reuse. These conditions are specific to wireless hardware parameters rather than environmental factors. Qualitatively, all our results are generic and apply to different multi-path environments (both indoors and outdoors), sectorized antenna parameters (number of sectors, sector width, peak gain in main lobe and side lobes) and wireless hardware parameters, though the actual numerical values presented in this paper may vary depending on the above factors.

The rest of the paper is organized as follows: Section 2 describes our experimental setup and Section 3 characterizes our performance metric for comparing different sectors. In Section 4, we present results on impact of sectorization on the performance of links in isolation. In Section 5, we present results on spatial reuse opportunities achievable when using sectorized antennas. We discuss related work in Section 6 and conclude the paper in Section 7.

2. EXPERIMENTAL SETUP

In this section, we describe the antenna hardware and testbed setup we use in this study.

2.1 Sectorized Antenna

We use multi-sector antennas [13] of the type shown in Figure 1(a) that consists of four Vivaldi sectors, each covering one quarter of the azimuth plane. These sectors can be simultaneously activated in any combination, and when all four sectors are active, it represents omni-mode. Sector activation is controlled by a switch integrated on the antenna system, which we control from the parallel port of the host computer. Figure 1(b) and Figure 1(c) show the radiation patterns in the azimuth plane for a single sector and omni-mode, respectively. The radiation patterns of different sector activations are obtained by performing measurements in an anechoic chamber [8]. The Antenna under Test (AUT) is attached on a mast at far-field distance from a source antenna and is rotated around a vertical axis. A RF continuous wave signal is emitted by the source antenna and the

received signal level by the rotating AUT is recorded for each angular position. For the antenna gain measurements, the conventional gain-comparison test method [8] is used with a standard gain antenna. The measured radiation patterns are used to obtain data shown in Section 4.4 and 5.1.

The maximum gain for sectorized as well as omni-mode are comparable. This property eliminates the negative effects caused by directional hidden terminal problem due to asymmetry in gain [5]. Current solutions to eliminate this problem uses multi-hop RTS/CTS packets which are hard to implement in commodity 802.11 hardware. The sectorization of this antenna thus primarily targets interference reduction along the side lobes rather than link budget improvement. These antennas operate in the 5GHz band and do not require any additional signal processing such as needed in MIMO operations. They can thus be used with commodity IEEE 802.11a wireless cards.

2.2 Testbed

We deploy a 6 node testbed in an indoor office environment as shown in Figure 2. We purposely select a dense environment which is rich in multi-path reflection, since this has resemblance with the typical conditions in dense wireless mesh deployments. Each node is a Dell D610 laptop equipped with an Atheros IEEE 802.11a/b/g mini-PCI wireless card, to which we connect the multi-sector antenna. We set this wireless card to operate in Ad hoc demo mode to facilitate ad hoc communication between nodes. In order to accurately identify the geographically pointing sectors between different nodes, we use the same orientation for all sectorized antennas. Each node has a second wireless card with a regular omni-mode antenna that is set to monitor mode. This card is used to accurately measure the number of packets actually sent out in air by the first card. We perform all our experiment in channel 52 (in the 5GHz band of 802.11a) without any external interference. All experiments are conducted during nights, except when we investigate the channel variation in the presence of people mobility (Section 3.2 and 4.3). All nodes run Linux (Fedora Core 6) and use Madwifi driver (version 0.9.4) to control the wireless cards. We use a UDP traffic generator to generate broadcast packets and collect measurement data using Tcpcdump. We detail our experimental settings for each respective experiment in conjunction with presenting the results.




Figure 2: A map of the testbed with the dark squares showing the location of the 6 nodes and the orientation of their antennas.

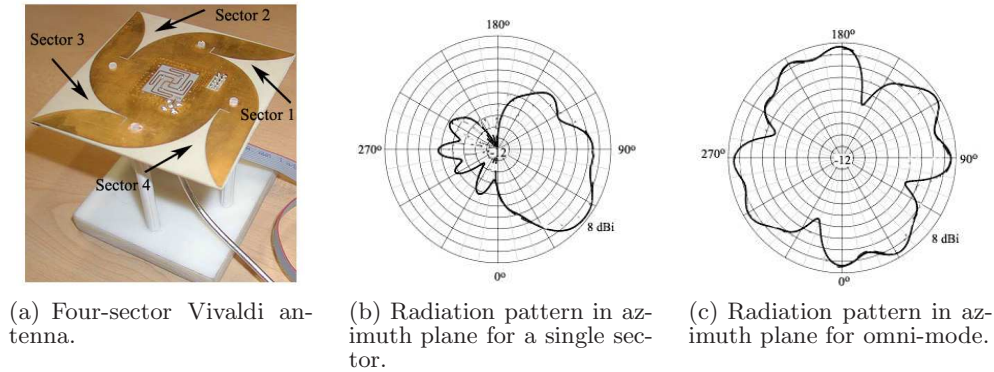


Figure 1: Sectorized antenna used in our study. The concentric circles in the radiation patterns are 2dB per division.

3. RSS CHARACTERISTICS

We use the *received signal strength* (RSS) value reported by the vast majority of commodity IEEE 802.11 wireless cards as our primary metric to compare the performance of different sector combinations. RSS is an estimate of the signal energy level at the receiver during packet reception and is measured when receiving the PLCP header of a packet. It is the fundamental physical layer quantity directly influenced by sector changes, because the antenna gain changes, and this in turn affects the energy level at the receiver.

We first study how well RSS of a link is related to the link’s delivery ratio and data rate. This is important in understanding how a physical layer phenomenon (e.g., a sector change) affects higher layer performance. Next, we quantify the variation of RSS in our testbed under different channel conditions. This is necessary to accurately compare different sector combinations, as it is impossible to measure the RSS of a link in all combinations at the same time instant.

3.1 Relation between RSS and Delivery Ratio

We set all nodes to omni-mode and instruct each node, one after another, to broadcast 1000 UDP packets of size 1024 bytes at a rate of 100 pkts/second. We measure the delivery ratio and the average RSS value of packets in each of the 30 links in our testbed. We repeat the same experiment by varying the transmit power in all the nodes from 18dBm to 1dBm, in steps of 1dBm, to obtain the average RSS value and corresponding delivery ratio for 540 different links. The above procedure is repeated for all data rates.

Figure 3 shows the relationship between average RSS and delivery ratio at different data rates. For each data rate, we obtain a scatter plot between RSS and delivery ratio and compute an interpolation for the same. For better readability, we show the interpolations for all data rates. The goodness of the interpolations is shown in the plot by the R^2 (coefficient of determination) values. The high values of R^2 show that the interpolations are good approximations of the relation between delivery ratio and average RSS. We see that for lower data rates the delivery ratio is 100% above a certain RSS value and drops sharply to 0% below this value. For higher data rates, there is a larger range of RSS values with intermediate delivery ratio. In these regions, the correlation between RSS and delivery ratio is weaker. In general, however, we observe that for each data rate the decrease in

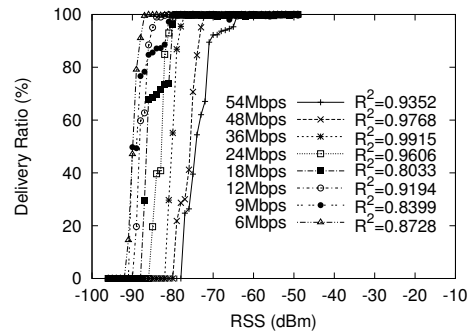


Figure 3: Relation between RSS and delivery ratio at different data rates.

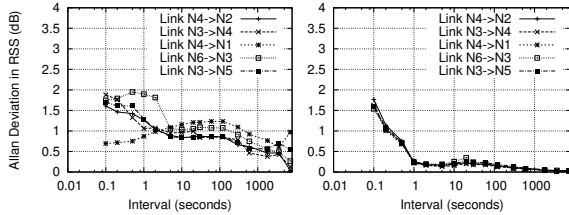
delivery ratio closely follows the decreasing RSS. *This relation is useful to map RSS of different sectors to higher layer performance.*

3.2 RSS Variation

We study 5 different links in our testbed, each in isolation and for a continuous duration of 5 hours. For each link, we set the antennas to omni-mode, send UDP packets of size 1024 bytes at a rate of 100 pkts/second, and record the RSS value of each packet. We conduct this experiment under two different channel conditions – during office hours when there are people movement, and at night when there are no people and thus a relatively more stable environment. We quantify the RSS variation of a link at different time intervals by computing the Allan deviation¹.

Figure 4 shows the Allan deviation of RSS for five links under the two different channel conditions. During office hours (Figure 4(a)), the Allan deviation is about 1-2 dB for all time intervals. It is important to note that under these conditions, sector combinations can only be compared with an accuracy of 2 dB. By comparing these results with the ones from the night time experiments in Figure 4(b), we conclude that most of the RSS variation during the office hours stems from people mobility. During night time, the Allan deviation is very similar across all five links and

¹Allan deviation measures variation burstiness (of any quantity), and can be used to understand RSS and loss rate variation of IEEE 802.11 links.



(a) During office hours. (b) During night time.

Figure 4: Allan deviation of RSS at different time intervals for five different links.

is consistently less than 0.5 dB for intervals greater than 1 second. In this channel condition, sector combinations can be accurately compared (within 0.5dB) even when the measurements are taken at different time instants. Moreover, RSS measured in a short interval (in the order of seconds or minutes) is representative of long term measurements. In all our experiments, we measure RSS of a link in different sector combinations for a duration of 10 seconds to quickly iterate over all combinations. In the next section, we study the characteristics of different sector combinations when links in our testbed operate in isolation.

4. SINGLE LINK BEHAVIOR

In this section, we study the behavior of links in isolation across different sectors. We define the following terms used in this section. ‘Sector combination’ denotes the sectors used by the sender and receiver node of a link. ‘Best sector combination’ refers to the sector combination with the highest RSS for a link. ‘Geographical sector combination’ refers to the sector combinations at the end-nodes that directly point towards each other. First, we study the RSS distribution across different single sector combinations. We then investigate how well the best sector combination is correlated with geographical direction, and next the best sector combinations’ temporal behavior. Finally, we quantify the impact of multi-sector activation on links’ RSS. Throughout this section, we use the following experimental setup. All nodes use 6 Mbps data rate. Each node takes turns and broadcasts 1000 UDP packets of size 1024 bytes at a rate of 100 pkt/second, and all other nodes record the number of received packets and their RSS values. This experiment is repeated for all single sector combinations and omni-omni combination, one after another.

4.1 Link Quality in Different Single Sector Combinations

We study the impact of using different single sector combinations on the RSS of a link. Figure 5 shows the average RSS value and MAC level delivery ratio for all 17 combinations (omni-mode and single sectors) for two sample links. In the first link, we observe that the delivery ratio is 100% for all sector combinations, although there is up to 16 dB difference in RSS among different combinations. The omni-omni combination has the highest RSS (recall from Section 2.1 that sectorization does not provide link budget gain) and is closely followed by the (4, 3) and (4, 4) sector combinations, while the other combinations have significantly lower RSS. This link can reduce interference to other nodes by switching from omni-mode to (4, 3) or (4, 4) while maintaining

the same link quality. In the second sample link, we can see that the delivery ratio drops to zero for several sector combinations whose RSS falls below the receive threshold (-92 dBm). We also see intermediate delivery ratios for combinations with RSS close to this threshold. This is representative of links that work only in certain sector combinations.

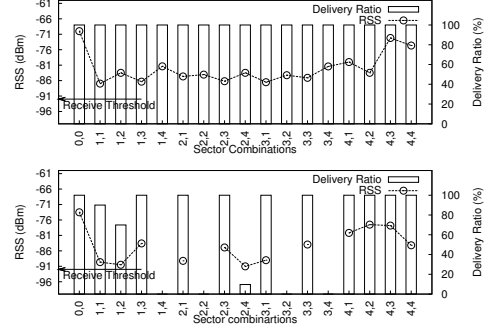


Figure 5: Average RSS and delivery ratio in two sample links.

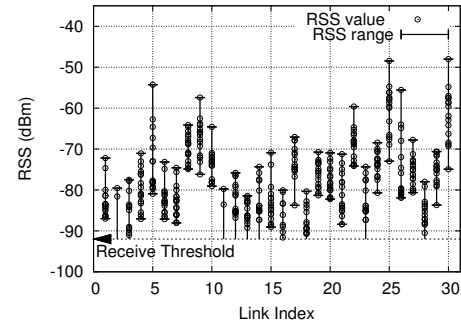


Figure 6: Range of RSS values in each link with different sector combinations.

Figure 6 shows the range of RSS values for the 16 single sector combinations of each link in our testbed. The circles on the vertical bars are the individual RSS values for each sector combination. The distribution of RSS values for each link is not uniform and differs for each link. We observe a maximum RSS range of 26 dB and the median RSS range is 14 dB. There are also several links for which the RSS range crosses the receive threshold. These links do not work in certain sector combinations. This result demonstrates the importance of identifying the best sector combination to achieve the highest possible data rate over a link (cf. Figure 3). Techniques to identify a link’s best sector combination is an important part of any MAC protocol that uses sectorized antennas. If the geographical sector combination of a link has the highest RSS, then the best sector combination can be identified without any measurement overhead by only knowing the location of the nodes and the orientation of the antennas. We therefore proceed to investigate whether we can identify the best sector combination using only knowledge about a link’s geographical direction.

4.2 Correlation of Best Sector Combination with Geographical Direction

A common assumption in directional MAC protocol design [5, 6, 10] is for any node to select the geographical sector for the node that it wants to communicate with. In this section, we study how often the best sector combination matches the geographical sector combination. We determine the best sector combination as described in the previous section, and the geographical sector combination by using the node locations and each node’s antenna orientation, as given in Figure 2. Figure 7(a) shows that the fraction of links in our testbed for which the best sector combination and geographical sector combination match (both sender and receiver side) is less than 40%. Our results show that between a majority of the nodes, there is poor correlation between the best sector combination and the geographical sector combination. The actual fraction of links in which the geographical sector combination has the highest RSS could vary depending on the amount of multi-path reflection that occur in the measurement environment. Qualitatively, our observations hold for all dense reflection rich wireless environments. Similar results have also been shown in both indoor [3] and outdoor scenarios [18]. The novelty in our work in comparison with existing works, is that, in addition to showing the poor correlation, we also quantify how bad the geographical sector combination performs in practice and propose practical sector selection mechanisms in dense wireless deployments (as discussed in Section 4.3).

Even though the geographical sector combination is not the best sector for a majority of the links, it is interesting to investigate how much it differs from the best sector combination in terms of RSS. This helps us understanding the trade-off between selecting geographical sector combination versus best sector combination.

Figure 7(b) shows the distribution of the RSS difference between the best sector combination and the geographical sector combination whenever these do not match. We can see that the median difference is about 8 dB and the maximum difference is as large as 24 dB, which is close to the maximum difference between the best and worst sector combination as seen in Figure 6. The large difference in RSS between the best and the geographical sector combination shows that relying on geographic location for sector selection can lead to a significant reduction of link performance. This result emphasizes the need to perform measurements in all sector combinations to identify the best sector combination. Measuring in all sector combinations can incur a high overhead. If there are k sectors in each node, the number of measurements needed is $O(k^2)$. If it is possible to optimize the number of combinations to measure, the best sector combination can be identified quickly. For this we investigate how much the geographical sector combination differs from the best sector combination in terms of angle. Figure 7(c) shows the percentage of time that a sector part of the best sector combination is the same, adjacent, or opposite compared to the sectors in geographical sector combination. We see that only in less than 10% of cases a geographical sector is directly opposite to the best sector. This gives us an opportunity to start with the geographical sector combination and measure only in the adjacent sectors to identify the best sector combination with high probability. This observation is also useful when using antennas with a higher number of narrow sectors as the measurements in

sectors in the opposite geographical direction need not be measured. Next, we investigate the temporal characteristics of the best sector combination in order to understand how often the measurements need to be repeated.

4.3 Temporal Characteristics of Best Sector Combinations

In this section, we study the temporal characteristics of the best sector combination for all the links in our testbed. We continuously repeat the previously described experiment for a duration of 30 hours. This duration includes both night and office hours. Recall that the Allan deviation during regular office hours is approximately 1.5 dB for intervals larger than 1 second (see Figure 4). We therefore use this as a threshold to distinguish between sectors when we identify the best sector combination. At each iteration, we select a sector combination to be the best if it has the highest RSS during the current iteration, and this value is more than 1.5 dB higher than the current RSS of the best sector combination in the previous iteration. Otherwise, the previous best sector combination remains the best for the current iteration. From our measurement data, we observed that for each link, a specific sector combination is selected as the best sector combination in a majority of iterations. We use the term ‘dominant sector combination’ to refer to this sector combination for each link. Figure 8 shows different characteristics of the dominant sector combination. Figure 8(a) shows that the probability of occurrence of the dominant sector combination is higher than 0.6 for all links and higher than 0.9 for 50% of the links. These high probabilities reveal that even under varying channel conditions there is good stability in the best sector combination. Figure 8(b) shows that whenever the best sector combination for a specific iteration differs from the dominant sector combination, it is within 90° (adjacent sector) with high probability. Another interesting characteristic is that this variation occurs only in one end of the link in most cases, as shown by Figure 8(c). Finally, Figure 8(d) shows the distribution of the difference in RSS value between the best sector combination and the dominant sector combination when they differ. Interestingly, in only 30% of cases is the dominant sector combination poorer by more than 1.5 dB and the difference does not exceed 3 dB in 90% of cases. This means that if we select the dominant sector combination even when it is not currently the best choice, the resulting difference in RSS would be typically low.

These spatio-temporal properties of the dominant sector combination can be exploited to reduce the measurement overhead involved in identifying the best sector combination for a link. *We conclude that practical MAC protocols using sectorized antennas cannot rely on geographic locations to identify sectors to communicate with other nodes. Instead, they should have a sector selection scheme that performs measurements in different sector combinations which can be optimized using our observations.*

4.4 Impact of Activating Multiple Sectors

In this section, we study the impact on the RSS of a link when activating multiple sectors at the communicating nodes. Activating multiple sectors is necessary when a node is communicating with multiple neighbors using different sectors and cannot switch between sectors on a per-packet basis. This idea has been used in various topology

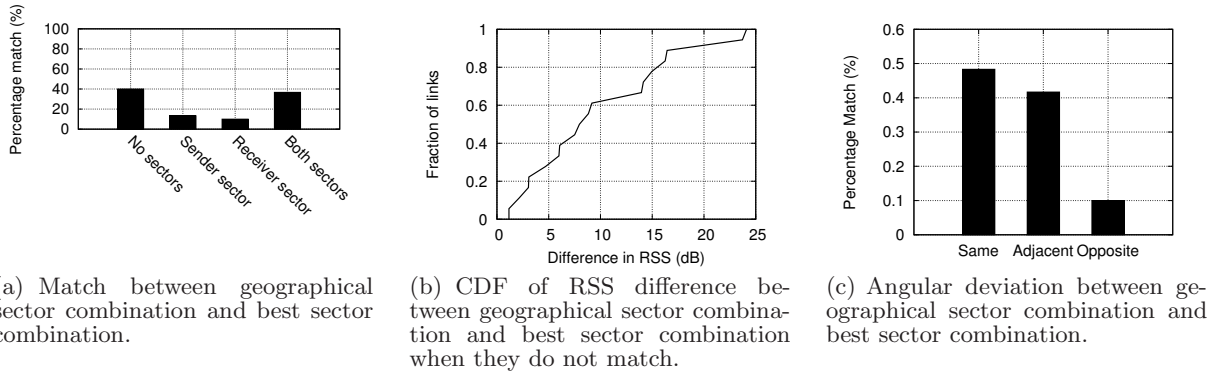


Figure 7: Comparing best sector combination and geographical sector combination.

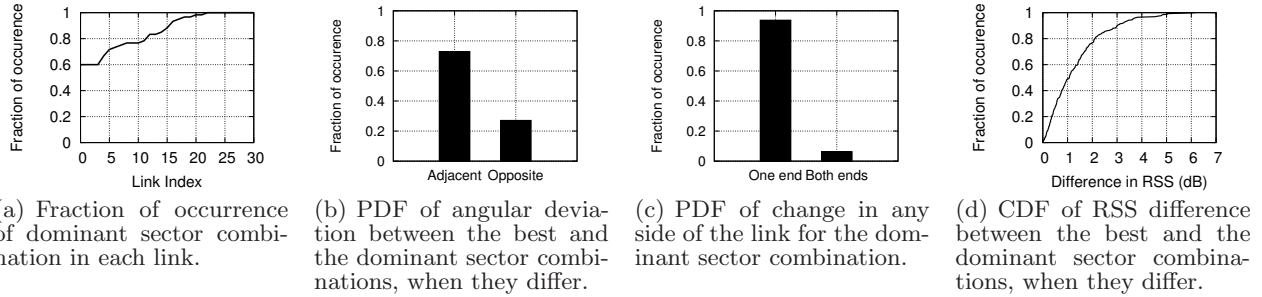


Figure 8: Characteristics of dominant sector combination.

control approaches using directional antenna [7, 11, 20]. A typical assumption in such topology control approaches is that activating an additional sector does not impact the already activated sectors. We refer to this as the ‘sector independence’ assumption. As an example, refer to Figure 2 and consider following scenario. The nodes $N4$ and $N2$ communicates using their best single sector combination. Node $N4$ then wants to communicate also with node $N3$ using a different sector, and must thus activate two sectors simultaneously. We investigate for each link in our testbed if such simultaneous activation of multiple sectors affects the performance of the original links (which previously used the best single sector). We conduct experiments similar to those previously described in this section. For each link in the testbed, we measure the average RSS in all sector combinations (including multi-sectors combinations). Recall that our antennas have four sectors, out of which any number of sectors can be simultaneously activated, leading to 225 sector combinations for each link.

Figure 9 shows the distribution of the difference in RSS value at the receiver node when the sender activates multiple sectors (two, three and all four sectors, respectively) compared to the best single sector combination of the link. We present two sets of data – one derived from sector radiation patterns measured in an anechoic chamber, and another based on measurements of the links in our testbed. According to the sector independence assumption, there should be no difference in RSS of the link when simultaneously activating multiple sectors compared to using single sectors, which would result in a curve that would be a vertical line at 0 dB in the figure. However, the results from the testbed as well as the anechoic chamber show that the RSS of the original

link is affected when activating two or more sectors. We note that our testbed results follow the same trend as the results from the anechoic chamber, but the latter typically shows a higher RSS decrease. We expect some differences due to signal propagation effects in real environments, such as different multi-path scattering for different radiation patterns.

Focusing on the testbed data, we see that for two and three sector combinations, the RSS decreases for about 40% of links and increases for the remaining links. For four sectors (omni-mode), we observe an increase of RSS for about 70% of the links. The reason for decrease and increase of RSS is due to the antenna radiation pattern irregularities when activating multiple sectors. *Radiation patterns of multiple sectors are not simple superimpositions of their component single sectors.* This stems from antenna design difficulties, such as the *antenna array factor* [9]. This factor quantifies the effect of combining radiating elements in an array without the element specific radiation pattern taken into account. This factor is not an artifact of the antenna used in our experiment and exists in all types of sectorized antennas based on antenna arrays, which is a popular method of constructing smart antennas (e.g. [2]). The antennas used in our experiments have indeed been optimized to reduce the antenna array factor [13]. We expect this effect to be more prominent in antennas that are not carefully designed.

The differences in RSS we observed have two implications. First, if the RSS of a link decreases by more than around 3 dB, there is a high probability that the link’s maximum achievable data rate decreases to a lower rate (refer Figure 3) or, in the worst case, a link operating at a maximum rate of 6 Mbps may break. Second, if the RSS of the

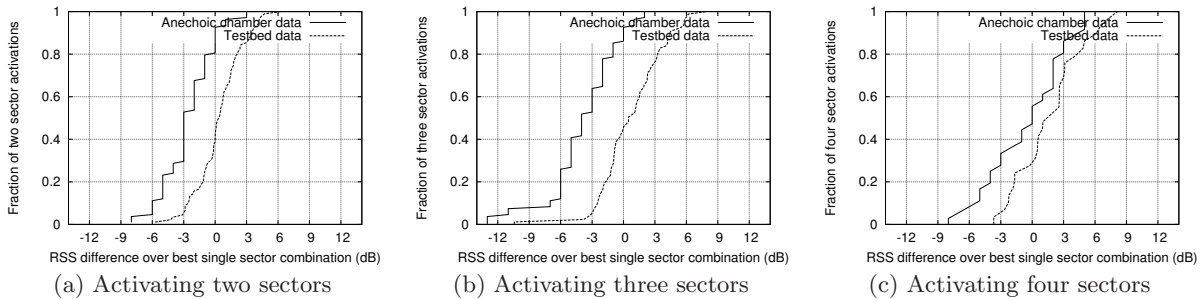


Figure 9: Effect of Activating multiple sectors on a link.

link increases when activating multiple sectors, it can cause interference (and thus collisions) at nodes which were not previously affected by this link. Our results show that the sector independence assumption does not hold in practice. It is insufficient to measure RSS only in the single sector combinations and assume that it will remain the same when activating multiple sectors. However, measuring RSS in all combinations is difficult in practice as the number of multi-sector combinations are exponential in number of single sectors. If there are k sectors in each node, there are $2^k - 1$ multi-sector combinations and the number of measurements for each link becomes $O((2^k)^2)$. Interestingly, as previously noted, the distribution of our testbed data closely follows the anechoic chamber data with a few dB offset (see Figure 9). This means that a topology control protocol could limit its RSS measurements to single sector combinations, and use the data from those measurements together with anechoic chamber data to extrapolate multi-sector RSS.

5. SPATIAL REUSE OPPORTUNITIES

We now investigate the spatial reuse opportunities offered by antenna sectorization in dense 802.11 wireless environments. We perform single-interferer and aggregate throughput experiments. In single-interferer experiments, we study the impact of a single interferer I on a link with sender S and receiver R. We first quantify the reduction in sender side interference when I uses single sector activations instead of omni-mode (Section 5.1). We then study to what extent increased spatial reuse at S is reduced by directional hidden terminal problems at R (Section 5.2). In aggregate throughput experiments (Section 5.3), we study the combined effect at both S and R by measuring the throughput achieved by multiple competing nodes using sectorization as opposed to omni-mode in our testbed.

Experimental Methodology: The setup in single interferer experiments is as follows. We select 23 links with 100% delivery ratio. For each link, we study the effect of four interferers, one at a time, which gives 92 (S,R,I) triplets. S and R always use their best single sector combination; I uses its four single sectors and omni-mode, one at a time. This results in 460 triplet configurations. For each configuration, we first instruct the nodes to broadcast back-to-back 1024-byte UDP packets in sequential rounds, 10s each, to collect (i) RSS for all 6 links among the three nodes (ii) number of packets (S_s) sent by S and the number of packets delivered at R (R_s) when S was broadcasting alone. Then, in a similar broadcast round, S and I broadcast simultaneously to col-

lect the number of packets (S_{si}) sent by S and the number of packets delivered at R (R_{si}). We describe the multiple interferer experiment in Section 5.3.

5.1 Medium Access Probability Improvement

In this section, we present results that quantify the reduction in sender side interference when the interferer I uses one of the sectors instead of omni-mode. We use the quantity S_{si}/S_s , referred to as *medium access probability*, as the amount of sender side interference reduction. Ideally, we would expect the medium access probability to be 0.5 for RSS values above a carrier sensing (CS) threshold value, (i.e., S and I equally share the channel). Similarly, ideally the medium access probability should be 1 for RSS values below the CS threshold, (i.e., I is completely hidden from S).

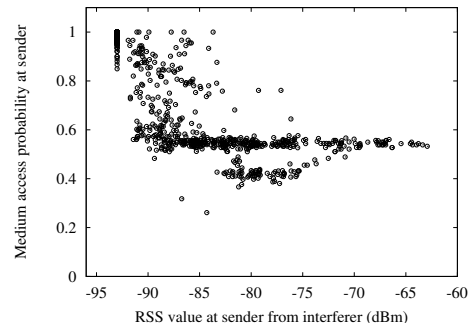


Figure 10: Relation between RSS value at the sender from the interferer and the medium access probability at the sender.

Figure 10 plots the medium access probability at S versus the RSS at S due to I for all 460 triplet configurations. We observe that for RSS values above -85 dBm, the medium access probability of S is between 0.40 and 0.55, indicating that S shares the medium with I. However, below -85 dBm, the medium access probability of S varies between 0.5 and 1. We conclude that in practice there is no sharp threshold below which S obtains full medium access. However, its medium access probability (hence spatial reuse) will still increase above 0.5 with high probability when sectorization at I causes the RSS to fall below -85 dBm.

Figure 11 depicts the CDF of RSS change at S, when I switches from omni-mode to each single sector. The first CDF corresponds to the 460 triplet configurations studied in

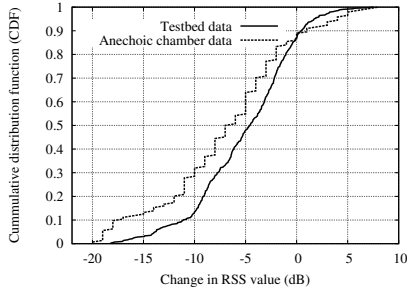


Figure 11: The distribution of change in RSS value at the sender when the interferer switches from omni-mode to one of the sectors.

our testbed and the second from antenna measurement data in the anechoic chamber. An important observation is that the two distributions have a similar shape and differ by only a few dB. This means that spatial reuse opportunities are primarily controlled by the antenna design characteristics rather than the testbed environment. The distribution from testbed data is shifted by a few dB to the right. Thus, there is typically less RSS difference between omni- and sectorized mode in the testbed environment. One possible cause is that sector combinations with large width (in this case omni-mode) suffer more from multi-path fading than single sectors.

In both environments there is about 10% probability that the RSS value at S increases when interferer I switches from omni-mode to a single sector. This case occurs when I uses its best sector towards S, while the omni-mode suffers low antenna gain in this particular direction (e.g., compare the respective gains at 60° angle in Figure 1(b) and 1(c)). The median RSS reduction is about 7 dB and 5 dB, and the maximum RSS reduction is about 20 dB and 18 dB, for the anechoic chamber and testbed, respectively. The maximum RSS reduction has an important implication. It limits the possible interference reduction, and thus also the maximum achievable improvement in spatial reuse. For example, if we refer to Figure 10, we can infer that with a maximum RSS reduction of 20 dB, we can in best case increase medium access probability (i.e., decrease an interferer's RSS below -85 dBm) only for sender nodes where the RSS from an omni-mode interferer is lower than -65 dBm. As a reference, we recall from Figure 3 that -65 dBm roughly corresponds to a loss-less 54 Mbps link.

Dense wireless environments typically suffer from interference with high RSS level, leading to increased need, but at the same increased challenge, to achieve spatial reuse. A high RSS reduction is therefore key to successfully improve spatial reuse in such environments. Note that the shape of the RSS reduction distribution directly depends on the antenna radiation pattern characteristics, such as the difference in maximum and minimum gain, and the particularities of the shapes of the main and side lobes. The achievable RSS reduction from sectorization² is therefore limited and subject to, often challenging, antenna optimization. As spatial reuse from sectorization increases, the directional hidden

²Combining sectorization with techniques such as transmit power control and carrier sense threshold adjustment to achieve some further spatial reuse is an interesting avenue for future work.

terminal problem increases. We therefore proceed to study to what extent hidden terminals exist and reduce the gain from the improved spatial reuse.

5.2 Directional Hidden Terminal Problem

The directional hidden terminal problem occurs when two nodes adjust their sectors such that they can transmit packets simultaneously (i.e., they no longer carrier sense each other), but their packets collide at a receiver. In this section, we investigate how the directional hidden terminal problem manifests in our testbed as a result of using sectorization.

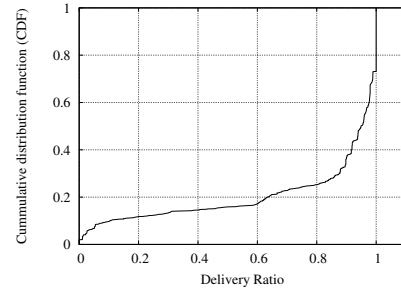


Figure 12: Distribution of delivery ratio of the links studied when the interferer is hidden from the sender node.

In Figure 12, we show the distribution of delivery ratio for the links in our testbed where sender S has medium access probability equal to 1 (i.e., interferer I is completely hidden from S). We see that approximately 70% of the links have more than 0.9 delivery ratio. This shows that even when there is an increased risk of directional hidden terminal problem, the delivery ratio does not significantly decrease, leading to increased effective throughput. Furthermore, about 82% of links have higher than 0.5 delivery ratio. In the case of broadcast traffic, this results in higher link throughput (as the medium access probability is 1) compared to the case when S and I perfectly share the medium with probability 0.5.

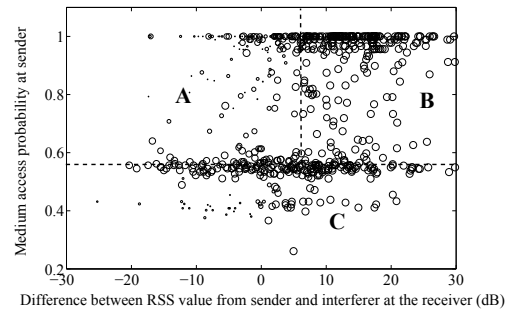


Figure 13: Scatter plot between difference in RSS value at the receiver from sender and interferer, and the medium access probability at the sender node. The size of the circle denotes the corresponding link's delivery ratio.

We explore the reasons for the sustained high delivery ratio in the presence of directional hidden problem using Figure 13. This figure shows a scatter plot between the difference in RSS value at the receiver from the sender and the

interferer, and the corresponding medium access probability. Positive RSS difference means that the RSS value at the receiver from the sender is higher than that of the interferer. The size of the circles is proportional to the corresponding link's delivery ratio. We note that there are three characteristic regions in the plot. We mark these regions using dotted lines and denote them A , B , and C .

In both regions A and B , we see that the sender obtains high medium access, which corresponds to the cases when the interferer with high probability is hidden from the sender. However, these two regions exhibit two main differences. First, links in region A typically have poor delivery ratio (small circles), while a majority of the links in region B have a delivery ratio above 90% (large circles). Second, the positive RSS difference from the sender and the interferer (at the receiver) is larger in region A than in region B . We see that when the RSS difference is higher than 6 dB (in Region B), the delivery ratio is high. This is due to the physical layer capture effect, that is, the packets from sender S are captured at the receiver R . The capture effect thus reduces, or in some cases completely eliminates, the directional hidden terminal problem. On the contrary, RSS differences below 6 dB result in collision at the receiver, which manifests the directional hidden terminal problem. Region C shows the cases when S carrier senses I , and thus shares the channel with it. However, despite S 's proper carrier sensing, we observe a few cases with poor delivery ratio in this region. These cases occur when I does not carrier sense S , which in turn causes collisions at the receiver (unless the sender's packets are captured due to the capture effect).

We conclude that the directional hidden terminal problem occurs in our testbed, but under certain conditions its effect can be eliminated by physical layer capture of the sender's signal at the receiver. *This implies that sector selection schemes for topology control algorithms should select sectors in different nodes that result in capture in the links that carry traffic.* This observation is generic and is not specific to the wireless hardware and antenna used in our experiments. A wide variety of current commodity wireless cards based on popular chipsets, such as Atheros and Prism, implement capture effect, though the actual thresholds can differ for different cards. The capture threshold also depends on the data rate used [17]. However, appropriate MAC modifications [5] are needed to address the remaining cases of the directional hidden terminal problem.

5.3 Impact of Sectorization on Aggregate Throughput

In this section, we quantify the impact of using sectorized antennas instead of omni-directional antennas on aggregate network throughput. We create 30 different topologies, each consisting of 3 sender-receiver pairs. We fix the data rate of all nodes at 6 Mbps. For each of these topologies, in turn, we perform experiments where all three links are simultaneously loaded with back-to-back UDP traffic (1024 bytes packet size). For each topology, we first conduct experiments with all nodes using omni-mode, and then repeat with all node pairs using their best sector combination. Figure 14 shows the relative improvement in aggregate sending rate and aggregate throughput when using the best sector combination for each link compared to using omni-mode. We order the topologies (x-axis) with increasing aggregate sending rate. We see from the graph with circles, that sectorization leads

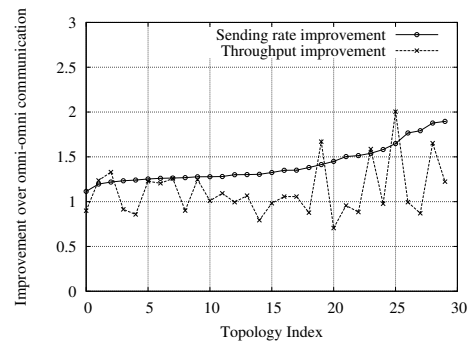


Figure 14: Improvement in aggregate sending rate and throughput when using best sector combination compared to using omni-omni combination.

to higher aggregate sending rates for all topologies. The graph with crosses shows that sectorization leads to throughput improvement in a majority of the topologies although we do not address hidden terminals. We observe that the relative throughput improvement compared to omni-mode ranges from 25% reduction to 200% increase. We notice that in some cases the improvement in aggregate throughput is higher than the improvement in aggregate sending rate. This occurs when using best sector combination results in less hidden terminal problems compared to using omni-mode.

6. RELATED WORK

The majority of past research on the use of sectorized antennas in 802.11 wireless networks has focused on analytical studies, protocol design and evaluation through simulations. Several directional MAC protocols have been proposed in the literature [5, 10] (see also references therein). These proposals focus on, either partially or completely, solving MAC related issues that arise from per-packet sector switching (e.g., deafness and directional hidden terminals). Topology control is another approach where nodes activate multiple sectors and form topologies that reduce interference in the network, thereby improving spatial reuse and capacity [7, 11, 20]. A few efforts have taken an experimental approach to study various aspects of using directional or sectorized antennas. Previous work studied V2I [14] and V2V communication [19] using steerable beam directional antennas and proposed practical beam steering strategies. In [4], a steerable beam directional antenna is used to study link quality in outdoor environments. The same antenna has been used in an indoor WLAN scenario [3] to study the extent of directionality and its impact on node localization and spatial reuse. Node localization has also been investigated in other indoor [15] and outdoor [18] settings. These works also observe that there is poor correlation between geographical direction and best directional sector due to multipath reflection. In [16], Ramanathan et al. studied the use of directional antennas in mobile ad hoc networks and developed a CSMA/CA based directional MAC protocol that operates on a custom-made radio board and showed performance gains in outdoor ad hoc network scenario with little reflections. Our work relates to these studies as follows. We conduct an experimental measurement campaign on a real testbed rather than simulations. In contrast to existing

experimental studies, we consider dense network environments with rich reflection and scattering, which pose great challenges in using sectorized antennas. Furthermore, rather than vehicular networks and steerable beam antennas, our work targets dense mesh networks using multi-sector antennas. While we confirm some results from the node localization studies, we extend beyond these studies and investigate the spatio-temporal characteristics of the best sectors. Moreover, we quantify the achievable interference reduction using commodity hardware, and provide implications that can guide design of practical protocols using sectorized antennas in these scenarios.

7. CONCLUSIONS

In this paper, we undertook an experimental approach to explore the performance characteristics of sectorized antennas in dense wireless mesh environments. We conducted a systematic experimental study using an 802.11 wireless mesh testbed equipped with sectorized antennas exploring the impact of single sector and multi-sector activations on link quality, effects of sectorization on spatial reuse opportunities, and the extent of directional hidden terminal problem in dense wireless environments. Based on our measurement results we made several observations that have implications on, and should guide the design of, future directional MAC and topology control protocols.

We draw the following main conclusions. Sector selection should be based on explicit measurements of sector performance. However, measuring in all sector combinations has high complexity. Our results suggest that we can significantly reduce this measurement overhead while maintaining good performance, by exploiting spatio-temporal characteristics of both the geographical sector combination and the dominant sector combination. Furthermore, antenna design, and in particular the radiation pattern characteristics, strongly influence the possible spatial reuse improvement. Anechoic chamber measurement data provides good indication for antenna performance in terms of resulting spatial reuse. In addition, this data could be exploited to further reduce measurements of multi-sector combinations in topology control protocols. Finally, physical layer capture reduces the effect of the directional hidden terminal problem and can help leverage the benefits of greater spatial reuse from sectorized antennas even without complex modification to the standard 802.11 MAC protocol.

8. ACKNOWLEDGMENTS

We would like to thank Christophe Diot for initiating this project, Ali Louzir, Philippe Minard and Jean-Luc Robert for designing and providing the Vivaldi multi-sector antennas, and the anonymous reviewers that helped improve the quality and presentation of the paper.

9. REFERENCES

- [1] FON WiFi Community. <http://www.fon.com/en/>.
- [2] Ruckus Wireless. <http://www.ruckuswireless.com/>.
- [3] M. Blanco, R. Kokku, K. Ramachandran, S. Rangarajan, and K. Sundaresan. On the Effectiveness of Switched Beam Antennas in Indoor Environments. In *Proc. PAM*, April 2008.
- [4] M. Buettner, E. Anderson, G. Yee, D. Saha, A. Sheth, D. Sicker, and D. Grunwald. A Phased Array Antenna Testbed for Evaluating Directionality in Wireless Networks. In *Proc. MobiEval : System Evaluation for Mobile Platforms*, June 2007.
- [5] R. Choudhury, X. Yang, R. Ramanathan, and N. Vaidya. Using Directional Antennas for Medium Access Control in Ad Hoc Networks. In *Proc. ACM MobiCom*, September 2002.
- [6] T. Elbatt, T. Anderson, and B. Ryu. Performance Evaluation of Multiple Access Protocols for Ad Hoc Networks Using Directional Antennas. In *Proc. IEEE WCNC*, 2003.
- [7] Z. Huang, C. Shen, C. Srisathapornphat, and C. Jaikao. Topology Control for Ad hoc Networks with Directional Antennas. In *Proc. IEEE ICCCN*, 2002.
- [8] IEEE Standard Test Procedure for Antennas. ANSI/IEEE Std 149-1979, 1979.
- [9] R. C. Johnson and H. Jasik. *Antenna Engineering Handbook*. Mc Graw Hill, Second edition, 1984.
- [10] Y. Ko, V. Shankarkumar, and N. Vaidya. Medium Access Control Protocols Using Directional Antennas in Ad Hoc Networks. In *Proc. IEEE INFOCOM*, March 2000.
- [11] U. Kumar, H. Gupta, and S. R. Das. A Topology Control Approach to using Directional Antennas in Wireless Mesh Networks. In *Proc. IEEE ICC*, 2006.
- [12] J. Liberti and T. S. Rappaport. Smart Antennas for Wireless Communications: IS-95 and Third Generation CDMA Applications. *Prentice Hall*, 1999.
- [13] P. Minard, J. Robert, and A. Louzir. Multi-beam Antenna System for Mesh Network with Omnidirectional Mode. In *Proc. IEEE Antennas and Propagation Society International Symposium (AP-S)*, 2008.
- [14] V. Navda, A. P. Subramanian, K. Dhanasekaran, A. Timm-Giel, and S. R. Das. Mobisteer: Using Steerable Beam Directional Antenna for Vehicular Network Access. In *Proc. ACM MobiSys*, June 2007.
- [15] D. Niculescu and B. Nath. VOR Base Stations for Indoor 802.11 Positioning. In *Proc. ACM MobiCom*, September 2004.
- [16] R. Ramanathan, J. Redi, C. Santivanez, D. Wiggins, and S. Polit. Ad Hoc Networking with Directional Antennas: A Complete System Solution. *IEEE Journal on Selected Areas in Communications*, 2005.
- [17] L. Ryu, J. Lee, S. Lee, and T. Kwon. Revamping the ieee 802.11a phy simulation models. In *Proc. MSWiM*, 2008.
- [18] A. P. Subramanian, P. Deshpande, J. Gao, and S. R. Das. Drive-by Localization of Roadside wifi Networks. In *Proc. IEEE INFOCOM*, April 2008.
- [19] A. P. Subramanian, V. Navda, P. Deshpande, and S. R. Das. A Measurement Study of Inter-vehicular Communication Using Steerable Beam Directional Antenna. In *Proc. ACM VANET*, September 2008.
- [20] K. Sundaresan, W. Wang, and S. Eidenbenz. Algorithmic Aspects of Communication in Ad-hoc Networks with Smart Antennas. In *Proc. ACM MobiHoc*, May 2006.