

A Phased Array Antenna Testbed for Evaluating Directionality in Wireless Networks

Michael Buettner, Eric Anderson, Gary Yee, Dola Saha, Anmol Sheth, Douglas Sicker, Dirk Grunwald

ABSTRACT

One of the most important components of any mobile system is the antenna; antenna design can overcome or cause a number of problems that then must be addressed at other technology layers. Modern mobile platforms are beginning to include novel antenna technology such as MIMO and beam steering; these technologies increase the complexity of evaluating the effectiveness of topology formation algorithms, routing and overall performance due to the large number of configuration states the system can contain. Directional antennas allow for significant improvements in link quality and spatial reuse in wireless communication. Traditional antennas with fixed direction are effective but unable to respond to station mobility or a dynamic environment including such factors as wind and foliage growth. There is a growing body of work on using steerable and sectored antenna systems to harness the efficiency of directional antennas while retaining the flexibility of ad-hoc networks; however, there has been very little work on implementation and measurement of such networks.

We examined the physical-layer properties of directional links in two real RF environments, and have evaluated higher-layer strategies for utilizing these antennas. Our results indicate the topology formation process must be a *network operation*, and that simple link-by-link topology optimization is likely to lead to poor overall performance. These observations drive the formation of the testing and evaluation tools we have developed. This paper describes the tools, methodology and metrics we are using in the evaluation of topology formation algorithms using a dynamically steerable phase array system.

Categories and Subject Descriptors: C.2.1 [Computer-Communication Networks] Network Architecture and Design – *Wireless communication*; B.4.1 [Input/Output and Data Communications] Data Communications Devices – *Transmitters*

General Terms: Design, Experimentation, Measurement, Performance

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.

MobiEval'07, June 11, 2007, San Juan, Puerto Rico, USA.
Copyright 2007 ACM 978-1-59593-762-9/07/0006 ...\$5.00.

Keywords: 802.11, RSSI, SINR, antenna, directional, phased array, steerable, testbed

1. INTRODUCTION

Existing mobile network devices typically have few “configuration states” – the radios used in mobile data applications tend to be fairly simple. However, recent challenges and research agendas proposed by DARPA [1] suggest that many problems in existing radio networks can be traced back to those simple physical interfaces. In response to those problems, the DARPA WANN program is developing mobile radios that will combine MIMO (multi-input multi-output) antenna systems in conjunction with other radio technologies. Those technologies will be combined with a radio networking layer that then seeks to exploit the collection of capabilities to increase link and system reliability and throughput.

Systems based on MIMO or “smart antennas” can be used in a number of ways to improve the wireless network topology or to improve overall throughput. Most MIMO systems actually use *multiple radios* that each control a single antenna. Those radios and antennas can be used to establish multiple concurrent links (using distinct frequencies for each radio). They can also be used for *beam forming* and *null steering*; in beam forming, the individual radio signals are transmitted on the same frequency, but are phase and amplitude shifted to provide a “direction” to the beam. Null steering involves using phase shifting to “null out” specific signals that interfere with desired communications. Lastly, MIMO systems can be used for *space time coding*, where multiple streams can be sent and received in the same frequency band. The DARPA WANN effort envisions a multi-radio system composed of four radios and antennas along with associated signal processing.

In this paper, we demonstrate one step in the process of evaluating such multi-antenna systems using a less mobile platform than that envisioned by the DARPA WANN program. We use dynamically steerable phased array antennas with a single radio each in our evaluation; the primary difference between these antennas and the MIMO systems envisioned in the DARPA WANN program is the use of a single radio interface, limiting the opportunity for multi-channel topology formation and space-time coding. However, the system is capable of beam forming and null steering.

We have been using these systems to evaluate *topology formation and control* algorithms. In the course of that work, we have developed an extensive infrastructure to evaluate and test such systems and have learned many lessons about

the efficacy of such systems. In this paper, we describe the evaluation platform and methods used, as well as the problems that arise as we move from “uncluttered” environments (where multipath is limited) to more realistic environments (with extensive multipath).

2. CHALLENGES IN EVALUATING MULTI-ANTENNA SYSTEMS

The capabilities of phased array antennas open up exciting avenues of research. The added flexibility also introduces a number of challenges.

- *Understanding directionality:* Existing static directional antenna systems are difficult to reorient. This makes it impractical to study the impact of direction in detail: Typically the transmitter and receiver antennas are pointed towards each other and fixed in place once a good enough – or locally optimal – state is found. Actual antennas have complicated beam patterns, and typical environments have complex transmission characteristics, meaning that there may be many reasonable directions with different advantages and weaknesses. Electronically-steerable antenna systems make it possible to explore these options in ways that would be effectively impossible with fixed directional equipment.
- *Large state space:* The phased array antenna system consists of 16 directional states and one omni state. Thus, in order to systematically explore all possible antenna state permutations for just a single, one-way link we would need to evaluate a space of 289 states. As more nodes and links are introduced the amount of time required to test the network can grow exponentially to impractical levels. Mechanisms are required to quickly search through the state space as well as to prune down the search space to make the search feasible.
- *Time synchronization:* To search such a state-space quickly, tight time synchronization is required to ensure the correctness of the experiment. Several mechanisms exist, but each has limitations. GPS devices can provide synchronization on the order of nanoseconds, but can only work with clear view of the sky and require additional cost. We chose to use the standard 802.11 mechanism for synchronizing the interface clock via periodic beacons. We discuss this later in greater detail.
- *Dynamic environment:* We have used two test sites: One is campus-wide testbed with link distances of roughly 300-400 meters. The other is situated in an open field with distances between 1 and 2.5 kilometers. Initial experiments have shown that the same experiments can yield significantly different results from day to day, as well as from site to site. Hence, experiments should be carefully planned and repeated across a large time scale and multiple environments to improve the applicability of the results. An especially variable aspect of the environment is interfering transmissions by other equipment. Experimenters working with unlicensed spectrum must quantify and account for its effects.




Figure 1: Unidirectional Pattern




Figure 2: Omnidirectional Pattern

3. APPARATUS FOR EVALUATING MULTI-ANTENNA SYSTEMS

Our test apparatus consists of Fidelity Comtech Phocus 3000 access points running our experimental software, and commodity x86 computers used for control and off-line data processing. The Phocus systems, described in detail in the next section, contains a single-board computer, a stock 802.11 interface card, and the phase-array antenna itself. The system uses a Linux kernel driver to control the wireless interface and steer the antenna.

3.1 Hardware

The phased array antennas used in our study were designed and constructed by Fidelity Comtech. The antenna operates in the 2.4GHz ISM band and uses an 8-dipole circular array that supports a minimum 42° primary lobe when configured for a tight unidirectional pattern, as shown in Figure 1. Each dipole is controlled by a vector modulator controlled by a distinct embedded processor. Intrinsic antenna reconfiguration time is $\approx 10\mu\text{seconds}$, although the interface with the transceiver boards limits the effective reconfiguration time to $\approx 100\mu\text{seconds}$. The transceiver boards are controlled using a series of phase-amplitude settings stored in flash memory, which allows fast reconfiguration between set patterns. For example, the antenna can quickly change the direction of the pattern shown in Figure 1, or switch to the omnidirectional pattern in Figure 2, by indicating the pre-computed configuration to be used; several thousand pre-computed configurations can be stored. The unidirectional pattern has a gain of 18dBi, which allows long-range connections. Additionally, the ratio of the lowest null to the highest peak is $\approx 40\text{dB}$, which allows for selectively “nulling out” interfering signals.

The embedded computer is an SBC based on the Intel XS-scale IXP425 processor. The wireless interfaces card used is a Senao 5345MP MiniPCI adapter. The combined antenna and embedded computer is physically compact and can be mounted on vehicles, light poles and buildings.

3.2 Software

Our experiments used stock packet generation and capture tools (`pktgen` and `tcpdump`) along with custom drivers to control and monitor the wireless link itself. The drivers were based on the Multi-band Atheros Driver (Mad-WiFi) open-source project [2] and antenna-interface code from Fidelity Comtech. Our modifications added instrumentation and facilities to schedule antenna state changes, synchronize such changes across many nodes, and sched-

ule packet transmissions to match the appropriate antenna state.

3.2.1 Configuration

Since our experiments were designed to focus on link characteristics, we disabled many features of the 802.11 protocol including RTS-CTS and link layer acknowledgements. We would have preferred to disable CCA in our experiments, but were unable to do so, resulting in a coupling effect between transmitting nodes.

For all experiments the data rate was fixed at 11Mbps. Pktgen was configured to inject 60 byte packets at a rate of one every of 2 ms. The small packet size was intended to give greater resolution with regards to packet loss, while the high packet injection rate was meant to saturate the channel for more accurate channel sounding.

On each node, two VAPs (Virtual Access Points) were configured in the driver: one in ad-hoc mode and one in monitor mode. The monitor mode interface was used exclusively for collecting trace data for analysis. The monitor mode interface passed all packets received or transmitted, with modified “Prism” metadata headers, up to the tcp-dump capture program.

3.2.2 Antenna and Packet Scheduling




Figure 3: Transmission chain architecture




Figure 4: Reception chain architecture

The state-space of system configurations is potentially vast: For k interacting nodes with n possible antenna states each, there are n^k distinct system configurations. To consider a large number of states in a practical period of time requires the ability to cycle through them quickly. Further, the risk that the external environment has changed increases with the elapsed time between when two cases are considered.

To address these concerns, we have developed infrastructure for quickly switching states in a coordinated manner across the entire system. There are two main challenges to doing so: (1) To interpret the results, it must be possible to match each packet sent or received to the antenna configuration in effect at the time. (2) To conduct experiments involving multiple nodes, it must be possible to ensure that

they change states together so that the system state remains consistent.

We address both of these challenges by clocking our system off of the high-resolution clock included in the 802.11 adapter’s chipset. Most of the difficulty in connecting packets to antenna states comes from non-deterministic timing: On the sending side, the host can know when a packet is passed to the hardware (diamond 1 in Figure 3), but it cannot know exactly when the packet will hit the antenna, especially if the card performs CCA and CSMA/CA backoff. Similarly, there is a variable delay between when the packet passes through the receiving antenna and when the host’s interrupt handler is called to service the packet (diamond 4 in Figure 4).

While there is a large margin of error as to what the *system time* was when the packet was actually sent, the *MAC time* can be known much more precisely. The MAC time, used for calculating retransmission timeouts, back-offs and the like, is maintained by a high-precision clock on the interface card. Packets are stamped by the hardware with the MAC time upon arrival (diamond 3), so there is almost no non-deterministic delay between the actual reception and the time-stamp. Because the AR5213 chipset also makes this time available *to the host*, antenna transitions are scheduled relative to the MAC time.

Using the on-chip timer helps with the second challenge: MAC time synchronization is already required by the 802.11 protocol and is done in the interface hardware. In both BSS and IBSS (ad-hoc) modes, stations include their MAC time in beacon packets. Listening stations then set their own clocks off the beacons. Because this is done in the chipset (diamonds 2 and 3), the variability in delay is much lower – and thus the synchronization much tighter – than what can be achieved using software on the end hosts.

4. EXPERIMENTAL SETUP AND DESIGN

4.1 Sites

To evaluate our experimental tools and to gain insight into the behavior of phased array antennas, we performed experiments in two different environments. The first environment consisted of four car mounted arrays on the perimeter of an open field where the topology exhibited relatively long links with clear line of sight. The second consisted of 3 car mounted arrays and one fixed building-mounted array. The car-mounted arrays were in parking areas around the University of Colorado at Boulder campus, one of which was an elevated ramp. The fourth array was mounted on the seventh floor of the Engineering Center. In this environment, the links are shorter but the paths are highly occluded.

4.1.1 NxN Scanning

The first set of experiments was designed to characterize the link quality when using different combinations of transmitter and receiver antenna patterns. Given the 17 patterns, omni plus 16 directional, we explored the complete state space comprised of 289 transmitter/receiver tuples for each of 4 transmitters. In total, 3468 tuples were assessed.

The NxN scan consisted of four sub-experiments, each of which has one of the nodes transmitting for 180 seconds while rotating its antenna pattern. The other three nodes did not transmit, and instead received while rotating their antenna patterns. The TDM schedules of all nodes were

identical and consisted of 17 slots, each slot using one of the 17 antenna patterns. When a node was a receiver the slot length was set to 40 ms, resulting in one complete cycle of the receiver antenna patterns every 680 ms. When transmitting, the slot length of the node was set to 680 ms so as to allow each receiver to complete one cycle for every slot of the transmitter. The length of each sub-experiment was 180 seconds, resulting in each receiver gathering data for approximately 16 iterations of all 289 transmitter/receiver tuples.

4.1.2 Independent Link Pairs

To determine the effect of interferers, we considered the two links A-B and C-D. The experiment consisted of two scenarios performed in sequence. The first had node A transmitting to node B and node C transmitting to node D, while the second differed only in the fact that node D transmitted to node C instead of vice versa.

To avoid an exhaustive N^4 search of the possible antenna configurations, we considered only a subset consisting of the top 100 configurations for each of the two links. The ranking was determined by the number of received packets per slot for each configuration in the $N \times N$ scan. For both of the scenarios, nodes A and B had slot times of 40 ms while nodes C and D had slot times of 4 seconds. With this, A and B will step through all 100 slots for each slot of C and D. Hence, one complete cycle of system wide antenna configurations takes 400 seconds; we allowed for three complete cycles.

5. EVALUATION

In this section, we evaluate how well the experimental system addresses the challenges identified in Section 2.

5.1 Fidelity of Beam Schedule Timing

The efficacy of the TDM mechanism depends on two things. First, the synchronization of the schedules across nodes and second, the precision with which the antenna states change at slot boundaries. To verify that the schedules are synchronized and that the antenna pattern is being set appropriately, we append the current antenna pattern to outgoing packets. When a packet is received, this information along with the current receiver pattern is passed to the monitoring software. Using these packet annotations, we can determine what the actual patterns were as close as possible to when the packet was transmitted and received.

Since the transmitter antenna state is appended to the packet before it is sent to the hardware queue, it is possible that a packet may not actually be transmitted before a slot change takes place. To mitigate this, we calculate the expected time of transmission and, taking into consideration packets previously sent in this slot, we limit the number of packets sent to the hardware. However, as we are unable to turn off CCA, it is possible for a packet transmission to miss its slot. In addition, at the receiver the packets are not processed immediately upon arrival. Instead, this processing is done by a tasklet which is scheduled with a 10 ms resolution. Consequently, it is possible for packets to be received with a certain antenna pattern, but before they are processed and stamped with the receiver pattern a slot change takes place.

Upon reception, a packet is time stamped by the hardware, and this timestamp is reflected in the trace and can be used to determine which slot the packet actually arrived in. Assuming the propagation delay is negligible, if the sched-

ules were not synchronized or the antennas were not being changed precisely upon slot boundaries, one would expect to see a discrepancy between the packet annotations and the patterns specified in the schedule.

We first looked at the cases where the receiver annotation did not match the schedule. We found that between 11.4% and 13.2% of the receiver pattern annotations disagreed with the schedule. As stated previously, the receiver pattern annotation is performed by a tasklet which, in our system, can be scheduled at approximately 10 ms intervals. One of these intervals will span every slot boundary and, on average, the receive tasklet will be scheduled approximately 5 ms after the start of a slot. As the slot length was 40 ms and we were sending packets continuously, on average 12.5% of the packets will be received near the end of one slot but will not be annotated until the next slot.

Next, we considered transmitter misses and found similarly encouraging results. For the experiments, we saw an average of 0.2% of the packets having an incorrect transmitter pattern annotation, with the max being 1.28%.

5.2 Metrics

The experiments in Section 4 were conducted as part of an effort to characterize the impact of antenna direction on link-layer network performance. This brings up two related evaluation problems: (1) How does one measure the goodness of the network as a whole? (2) How does one measure the goodness of its components? We don't have any conclusive answers, but we can shed some light on the strengths and weaknesses of several possible metrics.

A good component metric would be something locally-measurable that's also a good predictor of overall goodness. In the context of radio links, two widely-used metrics are signal strength and signal to interference and noise ratio (SINR). Signal strength is relatively easy to measure, and is widely used in practice. For instance, most 802.11 implementations use signal strength to select a base station to associate with.

Figure 5 shows the relationship between RSSI and throughput over an aggregate of several experiments in the presence of interference. Several limitations are apparent to the naked eye: First, the RSSI bounds the maximum throughput fairly clearly, but leaves a great deal of variation unaccounted for. Second, its predictive value diminishes significantly above a (modulation-dependent) threshold.

If the variation in RSSI with relation to throughput is due to interferers, SINR would be expected to account for it. However, SINR is difficult to define – let alone measure – for intermittent interference like packet-oriented communication. The hardware platform we used does not provide meaningful noise floor measurements. In practice, the only measure of interference using this type of hardware is that which can be gleaned from 802.11 packets that are successfully received. As such, any SINR measurement is a poor approximation at best. For our analysis, we calculated SINR by taking the ratio of the average RSSI of the transmitter and the average RSSI of the strongest interferer. Fig 6 shows the correlation between this SINR measurement and throughput. As can be seen, the case is similar to that seen for RSSI. A correlation can be seen when the signals are separated by 20 dBm or more, but there is a large amount of variation unaccounted for.

While it is clear that signal strength and interference (at




Figure 5: RSSI and throughput aggregated across campus and field




Figure 6: SINR and throughput aggregated across campus and field

the receiver) affect the quality of any radio link, it is unclear what may be the best practically-observable predictive metric. When protocol-specific behavior (like CSMA/CA) is considered, it becomes more complicated, as conditions at the sender affect throughput as well.

We have considered two system-level metrics in these experiments: Throughput and packet error rate. For a fixed transmission rate, they are interchangeable. However, they give fairly different views of the system in practice: Because of the “clear channel assessment” (CCA) aspect of the 802.11 protocol, heavy traffic near the transmitter inhibits it from sending. This reduces the throughput but does not necessarily effect packet error rate, particularly when using directional antennas as the environment at the transmitter can be quite different than that at the receiver.

Neither metric gives a *truer* view of the situation than the other: Packet error rate is a purer representation of the radio link as such, but throughput is a better indicator of how useful the link is to users with this type of hardware.

5.3 Directionality Under a Dynamic Environment

Wireless systems in the real-world are difficult to model because they can change over time. In order to gain an understanding of how this variability affects our experimental platform, we conducted our NxN scan over several short periods of time (minutes), separated by a long period of time (days), as well as in two different locations (a flat farm field and a cluttered campus environment).

Figures 7, 8, and 9 show a selected sample of the experimental results of these scans. State “0” indicates the




Figure 7: RSSI values for campus C→D link on day 1 for each pair of antenna states.

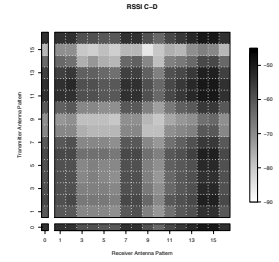


Figure 8: RSSI values for campus C→D link on day 2 for each pair of antenna states.

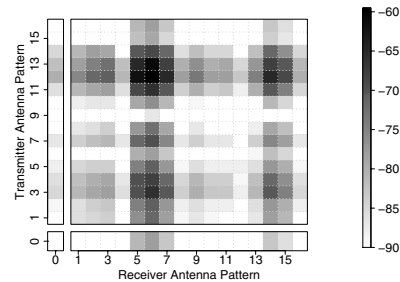


Figure 9: RSSI values for field link B→D for each pair of antenna states.

omnidirectional antenna pattern. In each figure we see clear indicators of each node’s main antenna lobe as well as side lobes and various “nulls”.

Looking at the variation of RSSI between two days, we see that the same campus link has changed significantly. Figures 7 and 8 depict the average RSSI on the campus link C → D for each antenna state combination. Figure 7 depicts the results on day 1 and Figure 8 shows day 2.

The two are both qualitatively and quantitatively different: Day 2 has higher overall values, and also a clearer band structure showing the antennas’ nulls and secondary lobes. These two differences do not automatically go together: Figure 9 shows a measurement taken at the field site. The signal strengths are lower than in either campus measurement, but the structure is at least as regular as in Figure 8.

We then evaluated the consistency of our RSSI measurements within and between 3 minute trials by calculating the coefficients of variation ($\frac{\sigma}{\mu}$), shown in Figures 10 and 11. We observed that the short-term variability was much lower than that seen either between days or between sites. It is also interesting to note that the *level of variability* is similar across different experiments, even when the results themselves are very different.

6. RELATED

Various researchers have examined the performance of real-world 802.11 based mesh networks [3]. A general finding is that these networks do not perform nearly as well as simulation might suggest. Aguayo et al. [3] provide an




Figure 10: Intra-experiment coefficient of variation: Link C→D RSSI, field exp. 1 (intra-experiment results for exp. 2 are comparable)

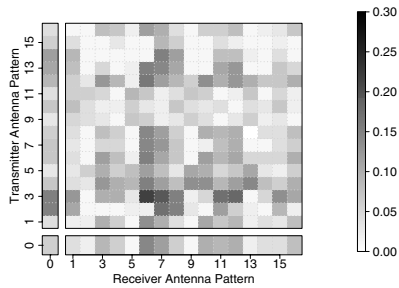


Figure 11: Inter-experiment coefficient of variation: Link C→D RSSI, field exps 1, 2 combined

assessment of a real-world 802.11b mesh network, Roofnet. They found that the demarcation between nodes that are in-range and out-of-range is indistinct and that multipath fading strongly influences loss rates for many of the links. Roofnet is based on a mesh of omnidirectional antennas. Ramanathan examines the performance of ad hoc networks built with beamforming antennas [4]. This work, based on simulation, shows under what conditions directional antennas might be beneficial. In later work, Ramanathan et al., describe an implementation of directional antennas and propose a new MAC protocol to manage backoff periods in a manner appropriate to such networks. Other researchers have examined aspects such as range control and routing in directional antenna MANETs [5, 6], scheduling algorithms [7] and MAC protocol variations [8, 9, 10, 11] for directional ad hoc networks.

7. CONCLUSIONS

In this paper we presented a platform and an experimental methodology by which the use of directional antennas can be evaluated systematically. Time-coordinated antenna steering makes it possible study the effects of directionality – for fixed and steerable systems – in ways that could not be done with only fixed equipment.

We discuss important metrics for characterizing these directional links. There are inter-node and environmental interactions that make measures like RSSI insufficient, and there are system-level effects that make it difficult to choose a metric which isolates the properties of interest.

Finally, we examined the consistency of measurements

across multiple environments and time scales. These results lead to two observations: First, the environment heavily influences the structure of how antenna patterns interact. Techniques which depend on predictable null and lobe effects may work in some environments, but are likely to fail in others. Second, environmental variability in the long term is much greater than in the short term. Experimental trials should be short, so that the different cases can be examined under comparatively consistent circumstances. Experiments should also be repeated over a longer term (and in different locations) to verify that results hold over a range of conditions.

8. ACKNOWLEDGEMENTS

This work was supported by NSF Projects #0435297, #0454404, #0435452.

9. REFERENCES

- [1] Preston Marshall. Wireless adaptive network node program. <http://www.darpa.mil/sto/solicitations/WANN/sn.htm>, Mar 2006.
- [2] Multiband atheros driver for wifi. <http://madwifi.org/>.
- [3] Daniel Aguayo, John Bicket, Sanjit Biswas, Glenn Judd, and Robert Morris. Link-level measurements from an 802.11b mesh network. In *Proc. Sigcomm 2004*. ACM, August 2004.
- [4] Ram Ramanathan. On the performance of ad hoc networks with beamforming antennas. In *MobiHoc '01: Proceedings of the 2nd ACM international symposium on Mobile ad hoc networking & computing*, pages 95–105, New York, NY, USA, 2001. ACM Press.
- [5] Mineo Takai, Junlan Zhou, and Rajive Bagrodia. Adaptive range control using directional antennas in mobile ad hoc networks. In *MSWIM '03: Proceedings of the 6th ACM international workshop on Modeling analysis and simulation of wireless and mobile systems*, pages 92–99, New York, NY, USA, 2003. ACM Press.
- [6] A Nasipuri, J Mandava, H Manchala, and R.E Hiromoto. On-demand routing using directional antennas in mobile ad hoc networks. *Computer Communications and Networks, 2000. Proceedings. Ninth International Conference on*, pages no.pp.535–541, 2000.
- [7] Lichun Bao and J.J. Garcia-Luna-Aceves. Transmission scheduling in ad hoc networks with directional antennas. In *MobiCom '02: Proceedings of the 8th annual international conference on Mobile computing and networking*, pages 48–58, New York, NY, USA, 2002. ACM Press.
- [8] Thanasis Korakis, Gentian Jakllari, and Leandros Tassioulas. A mac protocol for full exploitation of directional antennas in ad-hoc wireless networks. In *MobiHoc '03: Proceedings of the 4th ACM international symposium on Mobile ad hoc networking & computing*, pages 98–107, New York, NY, USA, 2003. ACM Press.
- [9] Yihu Li and Ahmed M. Safwat. Efficient deafness avoidance in wireless ad hoc and sensor networks with directional antennas. In *PE-WASUN '05: Proceedings of the 2nd ACM international workshop on Performance evaluation of wireless ad hoc, sensor, and ubiquitous networks*, pages 175–180, New York, NY, USA, 2005. ACM Press.
- [10] Siuli Roy, Dola Saha, S. Bandyopadhyay, Tetsuro Ueda, and Shinsuke Tanaka. A network-aware mac and routing protocol for effective load balancing in ad hoc wireless networks with directional antenna. In *MobiHoc '03: Proceedings of the 4th ACM international symposium on Mobile ad hoc networking & computing*, pages 88–97, New York, NY, USA, 2003. ACM Press.
- [11] Mineo Takai, Jay Martin, Rajive Bagrodia, and Aifeng Ren. Directional virtual carrier sensing for directional antennas in mobile ad hoc networks. In *MobiHoc '02: Proceedings of the 3rd ACM international symposium on Mobile ad hoc networking & computing*, pages 183–193, New York, NY, USA, 2002. ACM Press.