

A Measurement Study of Inter-Vehicular Communication Using Steerable Beam Directional Antenna

Anand Prabhu Subramanian¹, Vishnu Navda², Pralhad Deshpande¹, Samir R. Das¹

¹ Computer Science Department, Stony Brook University, Stony Brook, NY 11794, U.S.A.

² Microsoft Research, India.

{anandps, pralhad, samir}@cs.sunysb.edu, navda@microsoft.com

ABSTRACT

We provide a measurement study of a single vehicle-to-vehicle (V2V) link using 802.11b as the link layer technology. Our goal is to investigate practical usage of steerable beam directional antennas to improve V2V communications. We conduct extensive experiments using commercially available phased-array antennas mounted on cars in two different environments – suburban roads and highways, with various drive patterns. It is observed that directional beamforming improves the link SNR significantly, that translates to significant range improvements. However, to achieve this performance gain both antenna beams must be steered appropriately in the right direction. We observe that often the best beams indeed point directly to each other (called ‘LOS beams’), in spite of various sources of reflections that could be present in the environment. We develop and evaluate a simple beam steering approach that uses LOS beams for communication. We present experimental data, demonstrating the performance gains (in terms of SNR and PHY-layer data rates) achieved by this approach. While we have studied a single V2V link, this method can be extended to a multihop V2V network.

Categories and Subject Descriptors

C.4 [Performance of Systems]: Measurement techniques;
C.2.1 [Network Architecture and Design]: Wireless Communications—*Vehicular Communications*.

General Terms

Measurement, Design, Experimentation, Performance.

Keywords

Inter-Vehicular Communication, Steerable Directional Antenna, Beam Steering.

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.

VANET’08, September 15, 2008, San Francisco, California, USA.
Copyright 2008 ACM 978-1-60558-191-0/08/09 ...\$5.00.




Figure 1: Beam steering to improve quality of V2V links between two moving vehicles in opposite direction.

1. INTRODUCTION

Vehicular ad hoc networking (VANET) has been an emerging area of interest in recent times. The idea here is to form a mobile ad hoc network where a group of vehicles form the network nodes, possibly with connectivity to static infrastructure. Such networks can enable applications of tremendous value, ranging from safety (e.g., collision warning) to infotainment (e.g., traffic information, just-in-time navigation, multimedia download, etc.). VANET research has progressed mainly in two areas (not necessarily independent) – vehicle-to-infrastructure (V2I), connecting moving vehicles to road-side access points (APs) in a one-hop fashion, and vehicle-to-vehicle (V2V), where a network is to be maintained among a set of moving vehicles. Our focus in this paper is on V2V network.

In this work, we study the use of a physical layer enhancement – directional communication – to improve the connectivity of a vehicle-to-vehicle (V2V) network. By focusing energy in one direction, a directional antenna can get a better transmit or receive gain in a target direction compared to their omni-directional counterparts [27]. Directional antenna also provide better immunity from co-channel interference [27] and multi-path fading [11]. This directly improves SINR, giving better range as well as better PHY-layer data

rates. However, when a directional antenna is used in a mobile V2V context, the direction must be steered appropriately to maintain the best link quality. Also, steering must be done in a continuous basis keeping with the motion of the end-points of the V2V link. See Figure 1.

In this context, our goal is twofold – (i) to understand the potential benefit of using *steerable beam directional antennas* in the context of V2V communication, and (ii) to develop practical beam steering techniques. We use the *MobiSteer* system we developed on our past work [23] as the experimental platform. *MobiSteer* is a 802.11-based mobile network node that uses a steerable-beam directional antenna with an appropriate beam steering protocol. We use 802.11b/g as the link layer because of its wide availability, though much of the techniques developed in this work are not link layer specific and would apply as well other link layers (e.g., 802.11p based DSRC [4]).

To put certain amount of focus on our work and also to make the logistics of experiments manageable, our work investigates the performance of a *single V2V link in isolation* and not an entire multihop network. The expectation is that if the performance of a single link can be improved, multiple instantiation of the same underlying technique with an appropriate directional MAC protocol like [16,29] would improve the performance of all network links. We will revisit extensions of our techniques for a network-wide use in the Section 4. Our technique resides in the link layer, right above the MAC, and is completely transparent to routing protocols.

Our work makes the following contributions.

1. This is the first work to systematically study the use of steerable beam directional antenna in the context of V2V communication with 802.11-based link layer.
2. We report extensive measurement study using two mobile nodes (cars) in two different environments. We have chosen a suburban area and a highway for the experiments, noting from prior works [23,30] that the multi-path behavior and hence the performance of directional antennas are dependent on the surrounding environments. The two environments also provide different driving patterns (speed, turns, etc.) that impact the performance of our techniques.
3. Based on this study, we show that the best practical strategy for beam steering is to point the directional beams at the two link end points directly towards each other. We call this the *line-of-sight or LOS* strategy.¹ While this may not be always the optimal in the sense of SNR, we show that the cost of finding the optimal beam directions is significant in practical platforms.

The rest of the paper is organized as follows. In Section 2 describes the experimental set up and scenarios used for the experiments. In Section 3, we present the measurement results. Section 4 describes our beam steering strategies and present related experimental results. Related work is presented in Section 5. Conclusions and future work are presented in Section 6.

¹Note that this is a slight abuse of terminology as the straight-line between the two nodes may not always provide a line-of-sight when there is an obstruction in between.




Figure 2: Different beam patterns for Phocus Array antenna: (a) Omnidirectional Pattern (b) Directional beam pattern. Taken from [26].

2. EXPERIMENTAL SETUP AND SCENARIOS

In this section, we describe our hardware setup, measurement tools and give some details about the measurement scenarios used for the experiments.

2.1 Hardware

The hardware setup is similar to *MobiSteer* [23]. However, for the sake of completeness, we provide a brief description here. More details are available in [23].

We used two cars each mounted with a *Phocus Array* antenna unit from Fidelity Comtech [1]. The Phocus antenna unit is a phased-array antenna with electronically steerable beams for the 2.4 GHz band used in IEEE 802.11b/g. Different beam patterns are produced by software control via serial-line commands from a single-board computer (SBC). The SBC has a 802.11 a/b/g card based on Atheros 5212 [3] chipset with an external antenna interface that connects to the directional antenna elements. The SBC runs *Pebble Linux* [9] with the Linux 2.4.26 kernel and the widely used *madwifi* [8] (version 0.9.14.9) device driver (with some patches to be described in the next section) for the 802.11 interface. This SBC acts as a network node.

While many beam patterns are possible using the phased array, we have used 9 beams² – one omni-directional beam and 8 directional beams, each with an approximately 45° half-power beam-width on the main lobe and small side-lobes. See Figure 2. The beams are equispaced and cover the 360° circle with the 8 beam patterns. The directional gain is about 15dBi and the omni-directional pattern has about 9 dBi gain. We refer to the omni-directional beam with beam index 0 and the 8 directional beams we use with beam indices 1 to 8. Adjacent beams are numbered successively.

Beam switching latency is an important issue for continuously switching beams. The switching latency is optimized to 250 μ s as described in [23]. As will be apparent from our experiments later, this latency is significantly smaller than switching intervals in our experiments, and thus does not play any role in the evaluations.

A separate laptop computer is used to run a GPS daemon *gpsd* [6]. It communicates the location information to the SBC via Ethernet. We have used a USB-based Garmin GPS 18 [5] GPS receiver. This setup allows us to get the

²Many more beams are possible. We are describing only what we have used.

location of the vehicle along the route it travels. The laptop also makes the experiments manageable by providing console access for the SBC and downloading data at the end of an experimental run for later analysis.

2.2 Measurement Tools

Our measurement tool primarily consists of a *sender* and a *receiver* program that operate on the two MobiSteer nodes. They serve as the end points of the link under evaluation. The *sender* program is a UDP application that transmits unicast packets of a given size at a constant rate to the receiver application. In addition to transmitting packets, it also steers the antenna beam appropriately depending on the operating mode. There are three operating modes – *fixed*, *switching* and *LOS*. In the fixed mode, there is no beam steering. The antenna uses a specific fixed beam out of the 9 possible (one omni and 8 directional beams). In the switching mode, the antenna beam is switched periodically, cycling through all the 9 beams, staying in each beam for a specified *hold* time. In the LOS mode, the beam chosen for communication is always the one such that the centerline of the beam is the closest to the straight line joining the two nodes. The sender and the receiver programs compute the LOS beams independently as described in Section 4. Each packet sent is annotated with (i) a sequence number, (ii) the sender nodes’s current GPS coordinates, and (iii) the index of the antenna beam used to transmit the packet. We will discuss more about accurate beam index annotation in Section 2.3.

The *receiver* program is a modified version of Kismet [7], a popular wireless packet sniffing software. We run Kismet to capture packets on a raw monitoring interface created using the `madwifi` [8] driver so that each packet captured has a *prism monitoring header*, which contains information about RSSI (received signal strength indicator), noise, PHY-layer data rate and channel. Similar to the sender program, the modified Kismet on the receive side also operates in three modes – *fixed*, *switching*, and *LOS*. For every received packet, Kismet also annotates it with the current time (time in the prism monitoring header provided by the driver), the receiver node’s GPS coordinates and the receive beam index. For each received packet, all relevant information is logged for later analysis. It includes sequence number, sender and receiver beam index, sender and receiver GPS coordinates, receive timestamp, RSSI, noise, PHY-layer data rate and channel.

2.3 Beam Index Annotation

When beams are switching fast, the beam index annotation must be carefully done. As observed in [13], doing this annotation in the application program can often lead to errors. Then the difference between the time at which this annotation is done to the time at which the packet actually reaches the device driver for transmission can vary significantly and non-deterministically. During this time the beam could switch to a different pattern. To avoid this, we modified the `madwifi` driver to annotate the packet with the current beam index right at the point the driver handles the packet to the Hardware Abstraction Layer (HAL). This is the final point we have any control over the packet, as HAL is implemented in binary.³ While this is the best we could

³We indeed experimented with an open source HAL, OpenHAL [10]; but found it to be unstable in our system.

do, annotation errors could not be completely eliminated without some more work. This is because packet retransmissions are possible and this could take time in order of tens of milliseconds at low data rates for large packet. This time is of the same order of the hold times that we have considered. Thus, beams can still switch at this time scale. To prevent this problem, we set the retransmission count of 802.11 to be 0 so that lost packets are never retransmitted. Note that we can safely ignore backoff times, as they are in microseconds.

Also, our UDP application sends data only at a very moderate rate (one packet every 10 ms) so packet queueing delay within HAL would be very small, if any. Packet queueing in the card can lead to a packet to be sent in a different beam than the beam intended for it. We also took care to maintain a constant inter-packet interval so that we can get the same number of packet samples on all beam combinations for the switching mode experiments.

Similarly on the receiver side, the time at which the packet is received by the kernel and the time it reaches the receiver program can be quite different. We again annotate the packet in the driver right after the HAL delivers the packet.

2.4 Measurements Scenarios

In our previous works using *MobiSteer* [23, 30], we observed that in congested urban environments (with buildings and foliage in close proximity), the directionality of the antenna is poor and reflections often dominate. Reflections have the ability to make the best directional beam on one end of a link point away from the other end. This makes the determination of the best beam difficult. While in these prior studies, we used only one directional antenna (the other end point used a regular omni-directional antenna), we hypothesized that similar situation would be likely even with two directional antennas. Thus, we carefully selected the measurement scenarios. Our first scenario is a *highway* (2-3 lanes on each direction) that is fairly straight with only few turns along our driving route. The cars here have a less obstructed, straightline path between them. The second scenario is a *suburban* scenario – a dense residential neighborhood with narrow streets (1 lane) with houses and trees on both sides fairly close to the street. There are also quite a few turns along our driving route often obstructing the straightline path between the cars.

3. EXPERIMENTAL RESULTS

One car runs the sender program and the other runs the receiver program. The *switching* mode experiment is done to collect data in all $9 \times 9 = 81$ beam combinations between the sender and receiver. A fixed transmit power of 18 dBm is used on both the nodes.⁴ The PHY-layer data rates of both the nodes is ‘fixed’ at 1 Mbps so as to have a longer link. It is important to disable auto-rate control, as then the packet delivery performance on one beam combination will impact the rate and hence the performance on another beam. This will make performance across beams very difficult to compare. Evaluation with auto-rate control enabled is reported in a more appropriate context in Section 4. As

⁴The equivalent isotropically radiated power (EIRP) when using this transmit power is 33 dBm, well below the FCC limit of 45 dBm [26].

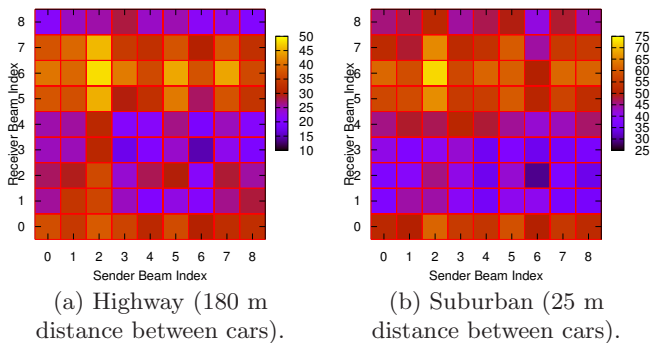


Figure 3: Average SNR (in dB) in each of 9×9 beam combinations for two representative samples.


explained before, the packet send rate is kept constant – 1 packet every 10 ms or 100 packet/s. Each packet has a 512 bytes payload.

After careful consideration, the hold time at the sender side is fixed at 30 ms and on the receiver side it is set 9 times this value, that is 270 ms. This is to ensure that once the receiver switches to beam i , it has opportunity to receive packets on all 9 beams from the sender before it switches to the next beam. A complete scan of 81 beam combinations thus take $81 \times 30 \text{ ms} = 2.43 \text{ s}$. Since the packets are sent at 10 ms interval, we need to hold on each beam for a duration longer than this so that in each beam combination, some packets are received to make meaningful comparisons. The cars can travel about 60 meters within 2.43 s even at 55 miles/hour (The fastest speed we used in highway scenario) and about 30 meters at 30 mile/hour (The fastest speed we used in suburban scenario). During our experiments, we found the SNR in the environment did not change much within these distances in the two scenarios considered.


The experiments are repeated 6 times in each scenario for (i) the two cars following each other and (ii) the cars driving in opposite directions. The cars drive normally appropriate for the road and traffic conditions. When the cars drive in the same direction, care is taken to produce as much variability as possible in the intervening distance between cars, as we sometimes categorize the performance results in terms of this distance. The car speed varied between 15–30 miles per hour in the suburban scenario and between 35–55 miles per hour in the highway scenario so as to take measurements with a wide range of distances between the two cars.

For each experimental run, information about all the received packets are logged (see Section 2.2). The log is post-processed to group all received packets for one complete scan of both send and receive beams, i.e., 81 beam combinations spanning over a continuous time of $81T_h$, where T_h is the hold time in each beam combination. Each group provides us with one ‘sample.’ In our analysis, we consider only those samples that have atleast one received packet in all beam combinations so that we can do meaningful comparisons. The ‘best beam combination’ in a sample is the one that provides the best ‘average’ SNR (Signal to Noise Ratio) for this sample. The averaging is done over all received packets within the hold period.

The first three subsections below analyzes results for the



(a) Highway scenario.



(b) Suburban scenario.

Figure 4: SNR improvement (in dB) for the best beam combination over the omni-omni combination on a scale of link distance.

cars following each other. The last subsection analyzes the results for the cars driving in the opposite direction.

3.1 Evidence of Directionality

To understand the effect of directionality we present two representative samples from the log in Figure 3. The figure shows the average SNR on all 81 beam combinations in a color-coded fashion. The first is a long link with the cars about 180 m apart and the second one is a short link with the cars about 25 m apart. Note that the difference in SNR between the best and worst beam combination is quite large — about 40 dB and 50 dB in the two samples respectively. Note also SNRs for all beam combinations around the best beam combination (sender 2, receiver 6) are relatively strong as well, going down about 5-10 dB for each beam difference. Note also the worst beam combination is approximately on the opposite direction (sender 6, receiver 2) as expected. The omni-omni combination (0,0) provides only a mediocre SNR, about halfway between the best and the worst. The combination with omni on one side and directional on the other side follow a predictable pattern. The line for beam index 0 on sender peaks at index 6 on receiver, and the line for beam index 0 on receiver peaks at index 2 on receive.

While qualitatively both scenarios produce similar results, there is a significant quantitative difference. The highway scenario gives similar SNR values even at much longer distance. This is due to the unobstructed view and openness of the surrounding space. Note that while we show here only two samples, similar behavior is observed throughout the drives.

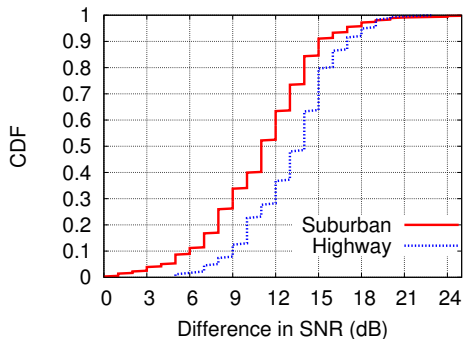


Figure 5: Cumulative distribution function of SNR improvement (in dB) for the best beam combination compared to the omni-omni combination (cars driving in the same direction).

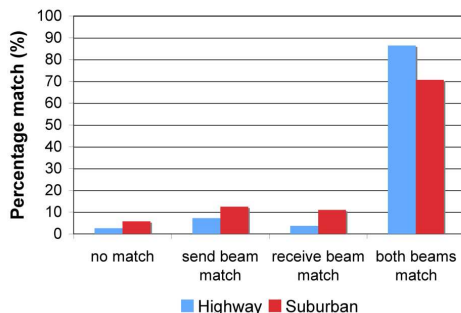


Figure 6: How often do LOS beams match with the best beams? (Cars driving in the same direction.)

3.2 Best Beam Combination Vs. Omni-Omni

Now we study how much improvement one can expect when using the best beam combination vs the omni-omni combination. The latter serves as the baseline as it represents the performance of using regular antennas without beamforming or beam steering abilities. Figure 4 shows the average SNR improvement (in dB) when using the best beam combination compared to the omni-omni case for each sample categorized by link distance (in 20 m divisions). We plot the improvement of the average SNR along with 95% confidence intervals for each distance category. Note that over a wide range of link distances the best beam combination provides excellent improvement when averaged, often between 10-14 dB. No specific behavioral difference is observed along the distance scale. However, there is some scenario-specific differences. See Figure 5 for the cumulative distribution function of the SNR improvements in the two scenarios for all samples. Note that the median SNR improvements are about 11 dB and 14 dB in the suburban and highway scenarios, respectively. Also, note that there is at least 6 to 10 dB SNR improvement 90% of the times. Slight better performance in the highway scenario (about 2-3 dB) is due to less scattering and longer segments of straight roads.




Figure 7: SNR differential (in dB) between the best beam combination and LOS beams (cars driving in the same direction).

3.3 Best Beam Combination Vs. LOS

Now, we will address a natural question: how often the best beam combination corresponds to the beams along LOS. We define the LOS beam as the beam whose main lobe is closest to the straight line direction to the other node in terms of angles. A detailed discussion on computing the LOS beams is described in Section 4.1. To determine the relation between LOS beams and the best beam combination, we look at the entire set of samples and present the fraction of times they match in Figure 6. We see a significant fraction of times (86% and 71% in highway and suburban scenarios, respectively) the LOS beams match with the best beams. Again, the higher fraction of matches in highway is due to less multipath and scattering. Also, quite often (about 96% of times in either scenario) there is a match of at least one beam. *This means that very often choosing the LOS beams actually provides the best beams.* With the availability of the GPS coordinates, LOS beams are easier to determine relative to the actual best beams. This is because the latter requires certain amount of probing and coordination among communicating nodes.

To see the dB difference between the two choices (best beams vs LOS), look at Figure 7. This figure shows the SNR differential between the best beams and LOS beams. The 90-percentile difference is less than 2 dB in both scenarios. Thus, very little improvement is expected by using the actual best beams relative to using the LOS beams.

3.4 Cars Driving In Opposite Directions

Now, we repeat similar experiments except that the cars drive in opposite directions. Again the experiments are repeated 6 times in each of the scenarios as before. These results give us insights as to whether the direction of travel influences our conclusions. Also, it is easier to study the range improvements when cars are driven in opposite directions. The cars start out of range and drive in opposite direction so that they come in range, stay in range for some time and then go out of range.

We next present similar data sets as in the previous experiments for studying (i) SNR improvements for the best beam combination over omni-omni combination and (ii) relationship between the best beams and the LOS beams. The results are presented in the same fashion in Figures 9(a),(b) and (c). We note that a high degree of similarity (both qual-




Figure 9: (a) Cumulative distribution function of SNR improvement (in dB) for the best beam combination compared to the omni-omni combination; (b) Frequency of match between the best beam combination and LOS beams; (c) SNR differential (in dB) between the best beam combination and LOS beams (cars driving in opposite direction).




Figure 8: Average communication ranges when using the best beam combination and the omni-omni combination (cars driving in opposite direction).

itative and quantitative) with the previous results (i.e., Figures 5, 6 and 7). Thus, we conclude that while surrounding environment plays a key role (e.g., highway vs. suburban), relative direction of travel does not play any significant role in use of directional antennas for V2V communications.

The range improvements are shown in Figure 8 for the best beam combination versus omni-omni combination. Note approximately 50% and 80% range improvements from use of directional antennas in highway and suburban scenarios respectively. Also, note in accordance with our previous observation with SNR, the communication range is almost twice in highways compared to suburban scenarios, with an average range of about 1.6 km.

4. PRACTICAL BEAM STEERING

So far, we have studied the performance of V2V links when using directional beams on both ends. While the performance improvement over using omni-directional beams is impressive, the performance potential has been demonstrated by scanning and probing on all beam combinations. Online scanning and probing on all beam combinations to select the best is evidently a high overhead operation. As discussed before in Section 2, the packet transmit time in the lowest bit rate (1 Mbps) of 802.11b is in the order of

several ms. Thus, even if the hold time is in order of a single packet transmit time, it will take an order of 100ms for probing on all 81 beam combinations. When cars drive in opposite direction, the best beam directions could change rapidly generating considerable overhead. Also, in a VANET setting a node may have more than one links. Each of them must be evaluated separately. This increases the overheads further. However, we have seen that the performance of the LOS beams very closely match the performance of best beam combinations in both scenarios and various driving patterns and speeds. LOS beams are relatively straightforward to compute. They only need the GPS coordinates of the both end points. Thus, each node just need to know only the coordinates of its neighbors. This is likely not an extra overhead, as some neighbor discovery protocol (e.g., hello messages [2]) must be present in the routing protocol. The GPS coordinates could simply be piggybacked onto this protocol messages. Also, neighbor location updates does not need to be very frequent. For example, predictions could be used based on speed history and driving route (from the navigation system, e.g.). Since the predictive model is known to all nodes involved, location updates can be sent only when the prediction deviates sufficiently.⁵

While orthogonal to our work, we mention a point here about neighbor discovery. This generates a subtle issue when directional communication is used. This is because many neighbor discovery protocols use broadcast messages. While omni-directional beam would be a natural choice for the broadcast, reduced range of the omni-directional beam presents a problem. There could be many solutions for this. For example, (i) such broadcasts could be sent with a higher transmit power on the omni beam; (ii) multiple broadcasts could be made on directional beams on all directions; or

⁵This can work as follows. Consider, two end points of a link U and V . U receives a location update from V , and from this point, predicts V 's location using a given model. The location update consist of V 's current location and mobility model. The model can use a range of information about V , e.g., speed history and route. Since the model is also known to V , V uses the same model to determine whether U 's prediction is sufficiently accurate. This can be done without any communication. When the deviation becomes significant, V sends another location update to U . See [20] for further details of this technique in a general setting.




Figure 10: LOS beam computation. Beam indices increase in the anti-clockwise direction.

(iii) reduced range is deemed acceptable, and only nodes reachable via the omni beam at regular power are deemed as neighbors. In the last case, the extra reach of the directional beams cannot be utilized, but higher SNR can certainly be used for better link quality and thus higher PHY-layer data rate. An appropriate choice here should take into account upper layer aspects, including the routing protocol and the needs of the application.

4.1 Computation of LOS Beam

For our experimental work here we use choice (iii) above. We use periodic broadcast beacons (at 1 sec intervals) on the omni beam to disseminate GPS coordinates. The same transmit power is used. No prediction is used. We consider using mobility prediction models like in [20] as part of our future work. Now, there are two technical issues to be addressed: (a) determining the LOS beam indices on either sides, and (b) consideration of the GPS error.

Since the antenna is fixed on the car, the orientation of the car needs to be determined to compute the LOS beam indices. We compute the car’s orientation by computing its ‘heading’. Heading can be computed based on past GPS locations. Note that we assume that GPS samples are taken every 200 ms intervals as we have done in the previous set of experiments. This sampling interval is deemed sufficient, as within this time the car can move a distance roughly equivalent to GPS error ranges even at highway speeds. A car can compute the LOS beam towards the other car using its own GPS coordinates, its own heading and the GPS coordinates of the other car. The beam with its center line closest to this angle is chosen to be the LOS beam. See Figure 10. Since we are considering only a single link in our experiments, LOS beams are computed whenever the car(s) move sufficiently and the antenna is always kept steered to this beam.

To determine when a new LOS beam computation should be done, we use estimates of GPS error. The goal is not to compute the beam too frequently that can be influenced heavily by GPS errors. We measured the error in the GPS unit we used (Garmin GPS 18 USB [5]) and found the median error to be about 5.5 m and 90-percentile error 7 m. Considering these values and several trials, we determined

a 25 m linear difference between two GPS readings of the same car is sufficient to trigger a new beam computation. This threshold is used to compute a heading or a new LOS direction.

4.2 Experimental Results

Our interest here is to do a realistic throughput experiment where the LOS beam is computed online. Note that we have computed the LOS beams only in the post processing step in the experiments described in the previous section using the logged GPS coordinates. Since all beams are no longer scanned as before, comparison between different beam choices (e.g., omni-omni and LOS beams) now presents a problem. This is because multiple drives are now needed with different beam choices for comparison. Since the position of two moving cars can vary a lot during different drives, it is hard to compare across experiments. To address this, we keep one car fixed and drive the other car such that the cars come in range, remain in range for some time and then go out of range. This does require fairly rapid change in the LOS direction when the cars cross each other. Thus, the LOS computation technique is exercised quite well. Also, now different experiments use the same car location and drive paths, making them directly comparable. The default auto-rate control algorithm [12] in the `madwifi` driver was enabled instead of using a fixed rate as we are not scanning in all beams. This is to emulate a realistic operational environment, and also to study improvement in the PHY-layer data rate using this technique. MAC level re-transmissions are also turned on since we are not switching between beams frequently. Otherwise, the software setup for these experiments is similar to that used in the previous section. Experiments are done again using the same two scenarios.

Two different experiments are performed - (i) the sender program and Kismet operate in the *fixed* mode with the beam index set to 0; (ii) the sender program and Kismet operate in the *LOS* mode where they use the LOS beams computed using the GPS coordinates of the two nodes as described before. In both the experiments the sender node is static and the receiver node moves at a constant speed. It comes in range, crosses the sender node and then goes out of range of the sender node.⁶ The speed of the car for both the experiments was kept constant (50 mph in highway scenario and 30 mph in the suburban scenario) via cruise control. The sender node transmitted 512 byte packets at a rate of 100 pkts/second. Each experiment was repeated 4 time in both scenarios.

Figure 11 shows the average SNR for each 20 meter segment along the path travelled by the mobile car. The x-axis shows the distance between the two nodes as the mobile node (receiver) comes in range, crosses the sender and goes out of range. Packets recieved in the 4 runs of each experiment were grouped based on the distance between the sender and receiver node and we plot the average SNR for each 20 meter segment. In the highway scenario, note a peak SNR of 56 dB and 36 dB in LOS beams and omni-omni combination respectively with a median SNR improvement of 18 dB. Note a peak SNR of 58 dB and 47 dB in LOS beams and omni-omni combination respectively in the suburban scenario with a median SNR improvement of 16 dB.

⁶The results were similar when the roles of the sender and receiver are changed.

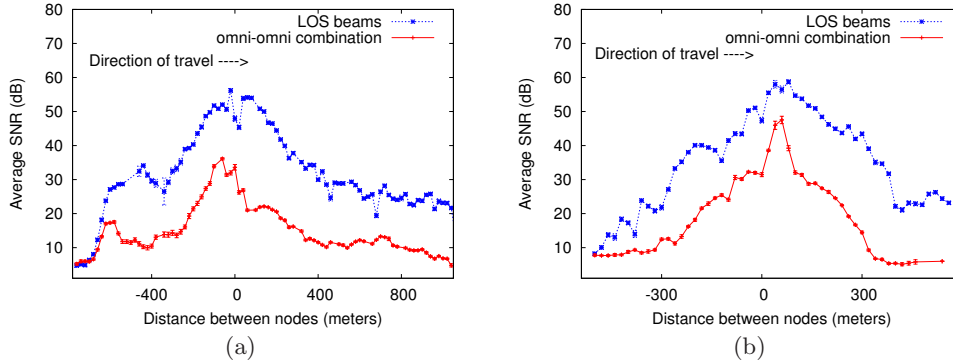


Figure 11: SNR improvement (in dB) when using LOS beams vs omni-omni combination in (a) Highway (b) Suburban scenario.

In Figure 12, we show the performance of LOS beams and omni-omni combination in terms of the PHY-layer data rate improvement. For each 20 meter segment, we plot the percentage of received packets in different data rates. The darker color in the plot denotes higher data rates. The PHY-layer data rate of the link is chosen by the auto-rate control algorithm based on the quality of the link. As seen from Figure 11, better SNR is observed even at a longer distance between the communicating nodes using the LOS beams and this is reflected on the data rate plot. While using the omni-omni combination, the auto-rate control algorithm switches the data rate to the highest value (11Mbps in 802.11b) when the distance between the communicating nodes is about 200 m or less. However, when using the LOS beams, the highest data rate is used as far as 800 m and 450 m distance between the sender and receiver nodes in the highway and suburban scenarios respectively. At farther distances, lower data rates are used due to poor link quality. As V2V links are short-lived due to high mobility of nodes, it is important to have a good quality link so that high throughput is possible for the duration of connectivity.

5. RELATED WORK

Much of the current research in vehicular networking is focused on analytical or simulation modeling for evaluations. Experimental studies are less common. Among the experimental studies reported in literature, work on V2I communication is more mature. Several works have demonstrated the feasibility of IEEE 802.11 based communication between moving vehicles and roadside APs and considered various performance improvement strategies at the link and transport layers. These works include [14, 18, 21, 23–25].

Fewer experimental studies have been reported in the V2V domain. In [28], the authors measure the performance of 802.11b-based V2V communication in different environments. They show that the performance of V2V links are greatly affected by the nature of the environment. Our experience in this regard has been similar. In [31], the authors demonstrate single hop and multihop V2V link with a single forwarder node. They present preliminary results on throughput vs. distance of V2V communication in a highway environment. Authors in [22] report experiences with static 1 and 3-hop scenarios and a mobile 3-hop scenario. TCP and UDP performance results are presented in a 2-hop vehicular network in [19]. In the DieselNet project [32], a bus-based

disruption-tolerant networking (DTN) testbed using 802.11 nodes are used for several studies – related to routing, mobility modeling, security and upper-layer protocol design. In comparison to these papers, our work adds a new dimension of using steerable beam directional antennas to improve the quality of V2V links.

Measurement studies using steerable beam directional antennas for 802.11-based networks are also quite limited. The *MobiSteer* project [23] studied V2I communication using the same steerable beam directional antenna as we have used here. The same antennas have been also used in [13] to develop a measurement tool and in [30] for a localization study. In [15], the authors also used beamforming antennas (a different variety) to improve network capacity in the context of static mesh networks and wireless backbones.

6. CONCLUSIONS AND FUTURE WORK

In this work, we have investigated the use of directional antenna beam steering to improve performance of 802.11 links in the context V2V communication. We have considered a single V2V link, and used extensive experiments in two different multipath environments to demonstrate the performance advantage over using regular omni-directional antenna. *In our experiments, we have seen median SNR improvements of about 11 dB and 14 dB in suburban and highway environments, respectively, relative to omni-directional communication. This also translates to significant range improvements (50% to 80% depending on the environment).* As expected, in highways with lower possibility of multipath reflections the improvements are considerably higher – relative to congested suburban environments.

We have also demonstrated that determining the best beams for communication can be simplified by a simple heuristic, where the beam pointing directly to the other node (LOS beam) is used. While in theory the LOS beams are not the best always, our results show that even in the suburban environment they are indeed the best for a significant fraction of times. Thus, instead of using an expensive scanning and probing method, we have developed a simple protocol, by exchanging the GPS coordinates, to determine the LOS beams in a continuous fashion. Our experience of using this protocol directly shows that the expected gain from directional communications is achievable in a practical setting. We show significant improvement in SNR and PHY-layer data rates. *The range in which higher data rates are used is*

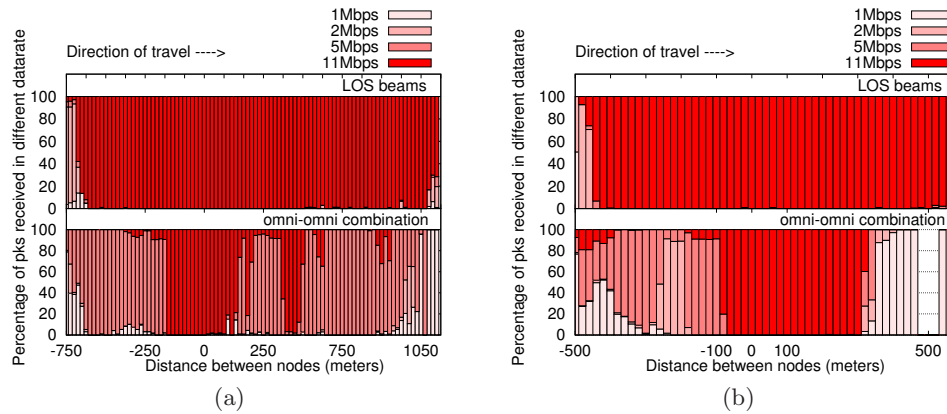


Figure 12: Physical layer data rate improvement when using LOS beams vs omni-omni combination in (a) Highway (b) Suburban scenario.

increased to about 2 to 4 times compared to omni-directional communication.

Note that our work has focused solely on improving link quality, when the link is noise-limited rather than interference-limited. In a high load situation, when the link is interference-limited, the same steering technique should still be useful. However, appropriate MAC protocols [17] must also be used for the best throughput performance.

While we have used a single V2V link, the technique can be applied easily in a more general ad hoc network environment. Use of LOS beams only needs dissemination of GPS coordinates of the nodes in the neighborhood. Such dissemination can be piggybacked in routing control messages, for example. Also, as we have discussed, predictive approaches can be used using historical information and knowledge of driving route to predict GPS coordinates within a given accuracy. A complete evaluation using an ad hoc network of multiple cars in the context of a real application will be a part of our future work.

7. ACKNOWLEDGMENT

This work was supported in part by NSF grants OISE-0423460, CNS-0519734, CNS-0721455 and from support from CEWIT (Center of Wireless and Information Technology). Vishnu Navda's work was done when he was a Ph.D. student in Stony Brook University.

8. REFERENCES

- [1] 802.11 Phocus Array Antenna System by Fidelity Comtech. <http://www.fidelity-comtech.com/>.
- [2] Ad hoc On Demand Routing Protocol (Internet RFC 3561). <http://www.ietf.org/rfc/rfc3561.txt>.
- [3] Atheros Communications. <http://www.atheros.com>.
- [4] Dedicated Short Range Communications. <http://grouper.ieee.org/groups/scc32/dsrc/index.html>.
- [5] Garmin GPS System. <http://www.garmin.com/>.
- [6] gpsd - a GPS service daemon. <http://gpsd.berlios.de/>.
- [7] Kismet Wireless. <http://www.kismetwireless.net/>.
- [8] MADWIFI Project. <http://sourceforge.net/projects/madwifi/>.
- [9] NYCwireless Pebble Linux. <http://www.nycwireless.net/pebble>.
- [10] OpenHAL. <http://madwifi.org/wiki/OpenHAL>.
- [11] T. Aubrey and P. White. A Comparison of Switched Pattern Diversity Antennas. In *IEEE VTC*, 1993.
- [12] J. Bicket. Bit-rate selection in wireless networks. *MIT Master's Thesis*, 2005.
- [13] 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 *MobiEval : System Evaluation for Mobile Platform*, San Juan, Puerto Rico, June 2007.
- [14] V. Bychkovsky, B. Hull, A. K. Miu, H. Balakrishnan, and S. Madden. A Measurement Study of Vehicular Internet Access Using In Situ Wi-Fi Networks. In *12th Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2006.
- [15] R. R. Choudhury, T. Ueda, J. Bordim, and N. Vaidya. BeamNet: Ad Hoc Networking Testbed Using Beamforming Antennas. In *IEEE VTC*, 2005.
- [16] R. R. Choudhury and N. Vaidya. Deafness: A MAC Problem in Ad Hoc Networks when using Directional Antennas. In *10th IEEE International Conference on Network Protocols (ICNP)*, 2004.
- [17] R. R. Choudhury, X. Yang, R. Ramanathan, and N. Vaidya. Using Directional Antennas for Medium Access Control in Ad Hoc networks. In *8th Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2002.
- [18] D. Hadaller, S. Keshav, T. Brecht, and S. Agarwal. Vehicular Opportunistic Communication Under the Microscope. In *ACM MobiSys*, San Juan, Puerto Rico, June 2007.
- [19] F. Hui and P. Mohapatra. Experimental characterization of multi-hop communications in vehicular ad hoc network. In *ACM VANET*, Cologne, Germany, 2005.
- [20] V. Kumar and S. R. Das. Performance of Dead Reckoning-Based Location Service for Mobile Ad Hoc Networks. *Wireless Communications and Mobile Computing Journal*, 4(2), 2004.
- [21] R. Mahajan, J. Zahorjan, and B. Zill. Understanding wifi-based connectivity from moving vehicles. In *ACM Internet Measurement Conference (IMC)*, San Diego, California, USA, 2007.

- [22] M. Moske, H. Fubler, H. Hartenstein, and W. Franz. Performance Measurements of a Vehicular Ad Hoc Network. In *VTC*, 2004.
- [23] 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 *ACM MobiSys*, San Juan, Puerto Rico, June 2007.
- [24] J. Ott and D. Kutscher. Drive-thru Internet: IEEE 802.11b for Automobile Users. In *23rd Annual IEEE Conference on Computer Communications (INFOCOM)*, 2004.
- [25] J. Ott and D. Kutscher. A Disconnection-Tolerant Transport for Drive-thru Internet Environments. In *24th Annual IEEE Conference on Computer Communications (INFOCOM)*, 2005.
- [26] Phocus Array Antenna System: Manual. Fidelity Comtech Inc. 2006.
- [27] R. Ramanathan. On the Performance of Ad Hoc Networks with Beamforming Antennas. In *2nd ACM International Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC)*, 2001.
- [28] J. P. Singh, N. Bambos, B. Srinivasan, and D. Clavin. Wireless LAN Performance Under Varied Stress Conditions in Vehicular Traffic Scenarios. In *IEEE VTC*, 2002.
- [29] A. P. Subramanian and S. R. Das. Addressing Deafness and Hidden Terminal Problem in Directional Antenna based Wireless Multi-Hop Networks. In *2nd International IEEE COMSWARE Conference*, Bangalore, India, January 2007.
- [30] A. P. Subramanian, P. Deshpande, J. Gao, and S. R. Das. Drive-by Localization of Roadside WiFi Networks. In *27th Annual IEEE Conference on Computer Communications (INFOCOM 2008)*, Phoenix, Arizona, April 2008.
- [31] H. Wu, M. Palekar, R. Fujimoto, R. Guensler, M. Hunter, J. Lee, and J. Ko. An empirical study of short range communications for vehicles. In *ACM VANET*, Cologne, Germany, 2005.
- [32] X. Zhang, J. K. Kurose, B. N. Levine, D. Towsley, and H. Zhang. Study of a bus-based disruption-tolerant network: mobility modeling and impact on routing. In *ACM MobiCom*, Montreal, Quebec, Canada, 2007.