R2D2: Regulating Beam Shape and Rate as Directionality meets Diversity

Kishore Ramachandran†, Ravi Kokku‡, Karthikeyan Sundaresan‡, Marco Gruteser†, and Sampath Rangarajan‡
†WINLAB, Rutgers University, North Brunswick, NJ, USA
‡NEC Laboratories America Inc., Princeton, NJ, USA
†{kishore, gruteser}@winlab.rutgers.edu, ‡{ravik, karthiks, sampath}@nec-labs.com

Abstract

We design, implement, and evaluate a vehicular communication system that improves uplink connectivity through multi-lobe beam pattern switching on a smart antenna. Directionality and base station-diversity are two well-known, independently developed mechanisms for improving the uplink connectivity of mobile clients. In this paper, we highlight that a system combining both mechanisms can achieve significant improvement in performance with multi-lobe beams that strike a tradeoff between directionality and diversity. This is in contrast to the mere steering of narrow beams used in conventional smart antenna systems. For tractability at vehicular speeds, our R2D2 *system searches through a limited set of beam patterns with different numbers of lobes, and includes a two-stage algorithm that uses both runtime adaptation and cached candidate patterns. We design and evaluate several variants of run-time adaptation that tune the number and angle of lobes in the beam, and the bit rate. The design of these algorithms is guided by both analysis and real-world measurements with a smart antenna system mounted on a vehicle. These measurements with our prototype implementation show that* R2D2 *can achieve an uplink throughput increase of up to 154% over pure beamsteering and 45% over pure basestation diversity.*

Categories and Subject Descriptors

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

General Terms

Algorithms, Design, Performance, Experimentation

Keywords

Outdoor Wireless Networks, Rate adaptation, Mobility, Beamforming, Uplink Capacity.

MobiSys'09, June 22–25, 2009, Kraków, Poland.

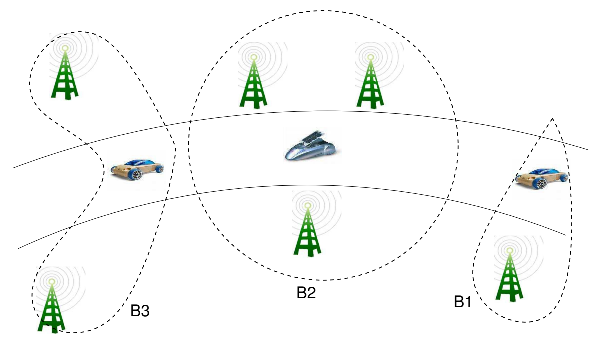

Figure 1: Different beam configurations for communication between clients and the receivers. B1 uses directionality with a single receiver, B2 uses diversity with all visible receivers, and B3 uses a combination of directionality and diversity with a subset of the visible receivers.

1. INTRODUCTION

The advent of truly mobile computing devices [7, 5] is leading to an unprecedented demand for affordable Internet access with both high capacity and coverage. In particular, users are increasingly expecting broadband Internet while actively *moving*, for example on public transit in trains and buses, and in cars. Meanwhile, novel applications such as web 2.0 [40], peer-2-peer media sharing [25], and video calls are shifting network workloads towards increased uplink usage, and reducing the applicability of existing downlink-intensive bandwidth allocation strategies.[1]

Directionality and Diversity. While several techniques are likely to be required to collectively meet these requirements, two fundamental mechanisms—directionality and diversity—have already individually found their way into several wireless standards, and we believe they form an integral part of future mobile networks. *Directionality* represents the idea of forming a beam towards a node (i.e., direct the transmitted energy in an intended direction) in order to increase the *average* signal-to-noise-ratio (SNR), while reducing interference to and from other surrounding nodes. This method is pictorially represented by beam B1 in Figure 1. Typically, it is implemented with an antenna array and is

[1]Existing 3G EVDO Rev A provides up to 3.1Mbps on the downlink and 300–400Kbps on the uplink. Next generation 3GPP Release 8 (LTE) supports a peak downlink bandwidth of 172.8Mbps and a peak uplink bandwidth of 86.4Mbps (2x2 antennas and 20MHz spectrum).

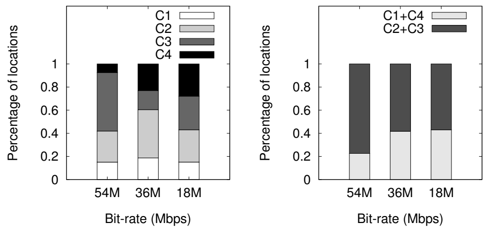

Figure 2: A combination of directionality and diversity (C2 or C3) achieves lower PER at a majority of locations relative to vanilla directionality (C1) and diversity (C4). C1, C2, C3, and C4 correspond to reception using 1, 2, 3 and 4 receivers respectively.

part of the WIMAX [4], LTE [9], and 802.15.3 [3] standards. While antenna arrays are usually first considered for base stations, their use has also been successfully demonstrated on vehicles [28, 37, 36] and could be particularly appealing for trunking solutions on mass transit vehicles [6, 23]. *Diversity*, in this paper, refers to base station macro diversity (a.k.a. receiver diversity) schemes, where multiple receivers at different locations overhear transmissions and combine received packets [39]. This scheme reduces packet losses due to SNR *variance* or deep fades at any individual receiver. It is already in use in CDMA cellular networks and is part of numerous additional standards and research proposals [13]. It may gain further prominence with the increasing availability of large numbers of small cells, in the form of WiFi access points (APs) [16] or customer-installed femto-cells [1]. Its use has also been successfully demonstrated for vehicular WiFi communications (ViFi [12]). This method is pictorially represented by beam B2 in Figure 1.

Need for beam shape adaptation. In this paper, we argue that a mobile system can *combine* directionality and diversity to significantly improve uplink throughput. This, however, raises the novel challenge of beam shape adaptation. Prior work, in contrast, has primarily focused on beam steering, to choose the direction of a single narrow beam (e.g., [28]). While a narrower beam increases SNR at one particular node, it tends to reduce diversity benefits because it covers fewer nodes (and vice versa). Intuitively, the optimal choice depends on node placement and the wireless channels to the APs. If only one AP is available, a narrow beam with no diversity performs best (e.g. Fig. 1-B1), while with multiple APs and highly variable channels, an omnidirectional (no beamforming) configuration with diversity (e.g. B2) may yield higher performance.

To verify this intuition, we measured SNRs and packet error rate (PER) for these different configurations from a moving vehicle to four road-side WiFi APs. More details on this experiment setup are provided in Section 4. Figure 2 shows the fraction of time (or locations of the moving vehicle) for which each of the configurations C1-C4 achieved the lowest PER (and hence highest throughput) among all configurations. Notably a majority of the time, a *mix* of directionality and diversity (C2 or C3) is more appropriate than using either pure directionality (C1) or pure diversity (C4). Further, the throughput-maximizing configuration varies over time. The graphs also show that the gain of using a combination of the mechanisms is higher at higher rates (54 Mbps); the lower rates (36 and 18 Mbps) are robust enough to reach maximum packet delivery ratios with either C1 or C4 alone. In other words, a beam that covers more than one but not all of the visible receivers (e.g. Fig. 1-B3) can enable the usage of higher rates, while also reducing packet loss. We confirm these results through analysis in Section 2.

R2D2's salient features. Our prototype system R2D2 provides beam shape adaptation, by switching among a set of selected single-lobe, double-lobe, and triple-lobe beam patterns that it preconfigures for a directional antenna. To rapidly converge on a suitable pattern on a highly mobile node, R2D2 uses a two-step approach. Nodes first report their GPS location to a centralized beam manager in the network infrastructure to obtain a set of candidate patterns with different numbers of lobes and PHY bit-rates for their upcoming locations. These candidate patterns are initially *learned*, and updated, through online learning algorithms based on client feedback. Next, to account for time-varying signal fading, we explore the effectiveness of two runtime adaptation algorithms, one performing only rate adaptation and the other performing both rate and beam shape adaptation. We implement the former in the prototype based on the complexity-to-performance tradeoff study among the two algorithms. The algorithm design is driven by both analysis and measurements. Note that R2D2's applicability is not limited to WiFi networks, although we build a prototype with a 802.11g Phocus Array smart antenna [8] mounted on a moving car and roadside WiFi APs.

In summary, this paper makes the following contributions.

- It identifies a fundamental tradeoff between directionality and diversity, and highlights the need for joint beam shape and rate adaptation.
- It provides a two-stage algorithm for joint beam shape and rate adaptation, with a location-based candidate set and runtime adaptation for fast convergence.
- It demonstrates feasibility and significant throughput gains over state-of-the-art techniques that use beam-steering and diversity in isolation — we observe gains of up to 154% and 45% respectively in our experiments.

The rest of the paper is organized as follows. In Section 2, we expose the tradeoff between directionality and diversity using an analytical model, which also serves to guide the design and implementation of R2D2 in Section 3. This is followed by an extensive empirical evaluation, using both trace-driven simulations and real-world experiments, in Section 4. Sections 5 and 6 discuss R2D2's limitations and related work. Finally, Section 7 concludes.

2. BASIC CONCEPTS

In this section, we provide the necessary background to understand the individual benefits of directionality and diversity. We then highlight a fundamental tradeoff between directionality and diversity that makes combining the two mechanisms for maximum benefit non-trivial. This motivates the need for a sophisticated beam adaptation system. The analysis also leads to several observations that will guide the design of R2D2.

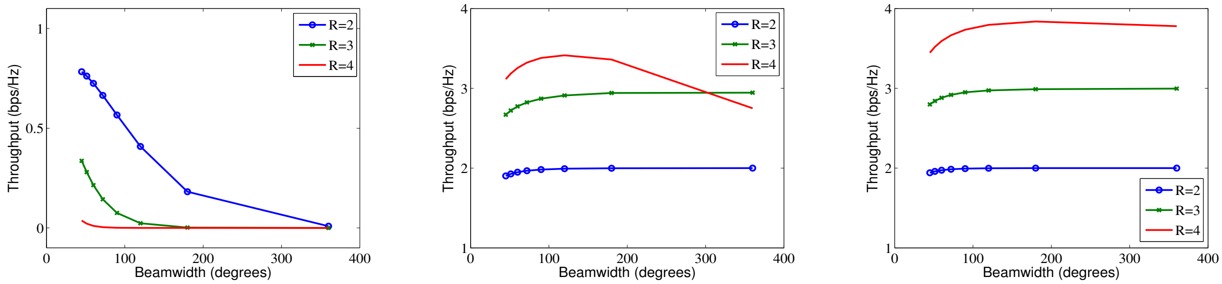

a) Low SNR (-4 dB) b) Medium SNR (8.75 dB) c) High SNR (11 dB)

Figure 3: Tradeoff Illustration. Throughput is maximized at a specific combination of rate and beamwidth, and individually controlling either the beamwidth or the rate leads to local maximas (making the case for joint adaptation).

2.1 Directionality

The notion of *directionality* corresponds to the ability of antennas to direct (beamform) energy in a desired direction, while suppressing the energy in all other unwanted directions. The footprint of the beam in the direction of maximum energy is often termed mainlobe. Increased directionality results in improved average link SNR in the desired direction, which is referred to as the *beamforming* gain. One way of achieving directionality is to use arrays of antenna elements placed in circular, linear, rectangular or other geometries. The signal sent to each of the elements is weighted in both magnitude and phase. The specific set of weights applied to the antenna elements is responsible for the antenna radiation pattern that is created. The antenna radiation pattern for an n-element uniform array is given by,

$$A(\theta) = a_1.e^{j(2\pi/\lambda)kd_1} + a_2.e^{j(2\pi/\lambda)kd_2} \ldots + a_n.e^{j(2\pi/\lambda)kd_n}$$

where a_n is a complex number representing the magnitude and phase applied to the n^{th} antenna element, k is the distance between two consecutive antenna elements, λ is the wavelength of transmission, and d_n is a geometry-specific function of θ, $0 \leq \theta < 2\pi$. Together, $(2\pi/\lambda)kd_n$ represents the shift in phase introduced by the displacement of each antenna element. An illustration of several beam patterns that we generated is provided later in Figure 5. If ρ is the average received SNR at a client due to an omni-directional transmission, then a beamformed transmission from an n-element array focused towards the client will result in a received SNR of at most $n\rho$, i.e. gain increases by a factor proportional to the number of elements n [19].

Observe that there exists a tradeoff between the beamforming gain and the mainlobe width. With increasing elements, the beamforming gain increases by a factor proportional to the number of elements, n. However, this is achieved by focusing energy in a thin lobe of width $\frac{2\pi}{n}$. In other words, increased beamforming gain is achieved at the cost of reduced spatial coverage. Also note that, since practical antenna arrays cannot completely eliminate the energy radiated in undesired directions, they do result in some spillover of energy in these directions, referred to as side-lobes. These side-lobes also increase with thinner main lobes.

2.2 Diversity

While directionality helps improve the average link SNR, it does not alleviate the pitfalls of variance in SNR, or deep fades, resulting in channel outages and consequent packet errors. Diversity is a mechanism that tackles such deep fades. Diversity constitutes the idea of leveraging the broadcast nature of the wireless medium to receive the transmitted signal at multiple receivers and exploit the statistical independence between the channel paths to the different receivers to successfully decode the packet. Essentially, with multiple observations of the same signal, the probability that all of the independent paths fail simultaneously reduces significantly, thereby alleviating channel outages and packet errors.

While diversity-combining can be considered at multiple layers—bit, symbol, packet, etc.—we focus on packet-level diversity that is amenable to implementation using off-the-shelf equipment, and is also shown to provide a large fraction of the benefits of diversity (relative to bit-level combining) on our chosen evaluation platform (with WiFi devices) [26]. If p_i is the packet error rate (PER) at a receiver i, then the PER after diversity combining reduces to $\prod_{i \in L} p_i$, where L is the set of receivers involved in diversity combining. Thus, it can be seen that the resulting diversity gain is dependent on the number of receivers involved, which in turn depends on the broadcast nature of the transmission.

2.3 Tradeoff and Observations

When the average SNR on a link is improved through beamforming, it reduces the broadcast nature of the transmission, thereby limiting its ability to leverage diversity combining to reduce PER, and vice versa. Consequently, there exists a fundamental tradeoff between using the available elements at a transmitter for directionality and diversity. However, in general, this tradeoff depends on a number of factors such as the number of antenna elements, the channel SNRs, the available set of discrete rates, the number of the receivers and their locations. Since an experimental approach to exhaustively capture the effect of all these factors is impractical, we take an analytical approach. Using a simple model (details in the Appendix), we capture this tradeoff. More importantly, we find that the link throughput ($T(R, n)$) is related to the number of antenna elements (n), the spectral efficiency in bps/Hz (R), and the number of receivers ($\ell(n)$) as follows:

$$T(R, n) = R \cdot \left(1 - \prod_{i \in \ell(n)} \left(1 - e^{-\left(\frac{S}{n\rho_i}\right)} \right) \right)$$

where $S = 2^R - 1$, and ρ_i is the average SNR at receiver i corresponding to an omni transmission.

The above equation captures the tradeoff between diversity and directionality. To illustrate this tradeoff, we quantitatively evaluate throughput as a function of n, ρ_i and R under simplifying assumptions. We assume $\rho_i = \rho$, $\forall i \in \ell(n)$, and $\ell(n) = \frac{8}{n}$ such that there are eight receivers reachable with omnidirectional transmission, and they decrease proportionally with increasing antenna elements (i.e. decreasing beamwidth). We allow fractional values of $\ell(n)$ for analytical simplicity. Figure 3 shows the results as an increasing function of average SNR (between -4 to 11 dB) across the graphs. Within each graph, we examine the behavior of different spectral efficiencies (R) for a given average SNR with increasing beamwidths on x-axis. The values for R (= 2, 3 and 4 bps/Hz) are chosen to be in the regime for future wireless broadband standards. Increasing n (the number of antenna elements) increases directionality (beamwidth of $\frac{360^o}{n}$), with $n = 1$ representing pure diversity from an omnidirectional transmission, and $n = 8$ representing pure directionality with an eight element array, with $1 < n < 8$ corresponding to combinations of diversity and directionality. The results reveal complex interactions between directionality, diversity, and rate. Specifically, we observe that:

Need for beam pattern adaptation: Comparing figures 3(a)-(c) shows that the throughput-optimal beamwidth depends on SNR and available rates. In case (a), throughput is maximum with a low beamwidth (i.e. increased directionality), whereas a moderate beamwidth is best in case (b), and full beamwidth (i.e. omni transmissions) maximizes throughput in case (c). In case (c), wide beams offer higher throughput because diversity is more effective in reducing packet errors due to deep fades than the average SNR increase possible through directionality. On the contrary, in case (a), when the average SNR is low, beamforming gain is required to increase the average SNR and maximize throughput. This motivates the need for beam pattern adaptation. It also means that the optimal beamwidth cannot statically be determined based on receiver locations alone, but must adapt to the actual observed SNRs.

Need for joint adaptation with rate: The throughput-optimal beamwidth depends on the rate choice and vice versa. Consider Fig. 3(b), where at $R = 3$, the highest throughput is achieved at a beamwidth of 360° while the throughput peaks at about 180° when $R = 4$. This behavior can be explained as follows. Any given wireless channel, at a point in time, supports a certain maximum R (that maps to a particular PHY rate) that can be achieved provided the set of rates are continuous. However, since the set of usable PHY rates are discrete in practice, there is a need for joint adaptation of beamwidth and rate so as to come as close as possible to the maximum achievable rate. Without joint adaptation, a system may get stuck in local maxima such as $R = 3/360^\circ$ in Fig. 3(b).

Summary: The analytical results in Figure 3 convey two guidelines for system designers: (a) throughput is maximized at a specific combination of rate and beamwidth, and (b) individually controlling either the beamwidth or the rate can lead to local maximas, thus making the case for joint adaptation of these parameters.

3. R2D2 DESIGN AND IMPLEMENTATION

Goals. Following these guiding observations, R2D2 is designed to improve uplink throughput through joint beam pattern and rate adaptation. Further, R2D2 aims for:

1. *Fast convergence suitable for vehicular speeds*: At high client speeds, the radio propagation environment changes rapidly. For example, at 50mph a car travels about 20m/s. It is not feasible to search through thousands of possible beam patterns to find the optimal configuration.

2. *Antenna technology independence*: As wireless technologies evolve, clients are bound to different beamforming capabilities. For example, the number of antenna elements or transmission powers may differ. The solution should be agnostic to these issues.

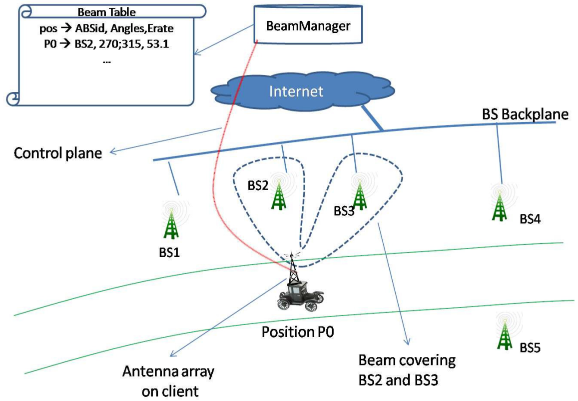

Figure 4: Solution Overview.

We also focus on *switched beam* directional antennas since they provide a good tradeoff between implementation complexity and beamforming gain [15]. We leverage the Phocus Array [8] platform in our implementation that can change or switch beam patterns, but cannot adaptively form new patterns, at runtime.

3.1 Approach

To achieve fast selection of beam pattern and rate, we (i) work with a limited set of the most useful beam patterns and (ii) use a two-step approach that includes long-term learning of parameters and short-term adaptation. Placement of receivers and their surrounding environment usually limits the set of useful beam patterns significantly. In some cases, even the optimal radio parameter choices can remain relatively stable for timescales of minutes to hours [20] at a given location. This allows learning of reasonable default parameters for each client location, particularly if clients can share the learned parameters. Thus, the long-term learning component is based on a centralized location-indexed *Beam-Manager* database through which clients obtain and share optimal beam patterns for each location. Clients determine their own position through GPS and access this database through a control channel to obtain a set of beam configurations for each of the upcoming locations.

Since the propagation environment can vary rapidly, we also use a short-term adaptation component on the client that selects the optimal beam pattern out of the received set and adapts bitrate from the default configuration. This short-term adaptation is based on actual observed SNRs and packet error rates. In essence, location-based learning provides reasonable starting parameters and a restricted search space to speed up convergence for the adaptation algorithm. Figure 4 shows a diagram of the overall architecture.

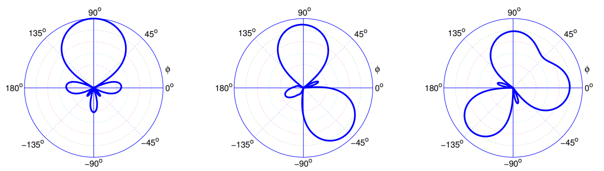

(a) 1 lobe at 90^o (b) 2 lobes at 90^o,-45^o (c) 3 lobes at 0^o, 90^o, -135^o

Figure 5: Example beams with different number of main lobes at different angles. The motion of the car is along zero degrees.

3.2 Selecting a Limited Set of Beam Patterns

To enable rapid switching between beam patterns, the Phocus Array [8] steerable antenna requires storing a predefined set of patterns in antenna memory. A pattern is defined through the phase and magnitude values applied to individual antenna elements, and when triggered by a command to the antenna the Phocus Array can switch among stored patterns with 150-200μs delay. A large number of possible patterns exist. For example, the 8 element Phocus Array antenna supports a minimum beamwidth of 45°, which can be rotated in 22.5° steps, yielding 16 possible patterns with the single narrowest beam. Considering all possible combinations of wider single- and multi-lobe beams, yields on the order of 2^{16} patterns, while antenna memory is limited to about 1000 patterns.

This raises the question of which patterns to store—e.g., is a 90° degree beam or a beam with two 45° lobes in opposite directions more useful? We select a subset of the possible patterns using the following heuristics:

1. *Select multi-lobe beams over wider beams.* It is well known that the channel between receiving antennas separated by distances greater than $\lambda/2$ is uncorrelated, and that the greater the inter-receiver distance, the lesser the correlation [30]. Accordingly, we expect that the channels to two receivers in different directions are more likely to be independent than the channels to two co-located receivers. We verify this using an experiment with three receivers, and a single mobile sender (see Figure 9 for the topological details). Two receivers (R2 and R3) are placed on the same side of the road (Independence Way), at a distance slightly greater than $\lambda/2$, whereas the third receiver (R1) is placed on the opposite side. The sender broadcasts 1500 byte packets at 54 Mbps, and we report the packet delivery rate using selection diversity across pairs of receivers in Figure 6. As expected, receive diversity is more beneficial (provides higher PDR and increased range) with widely dispersed receivers.

   From the perspective of directional antenna systems, multi-lobe patterns more efficiently cover widely dispersed receivers. To see this, consider an antenna communicating with minimum width of 30° with base station B1. To include another base station located at an angular separation of 180°by widening the beam leads to a loss of 9 dB power at B1 (3 dB loss for every doubling of beam width). A 2-lobe beam with 30° each covers both base stations but leads to a loss of only 3 dB at B1. Thus, we expect higher gains from multi-lobe, rather than wider single-lobe patterns.

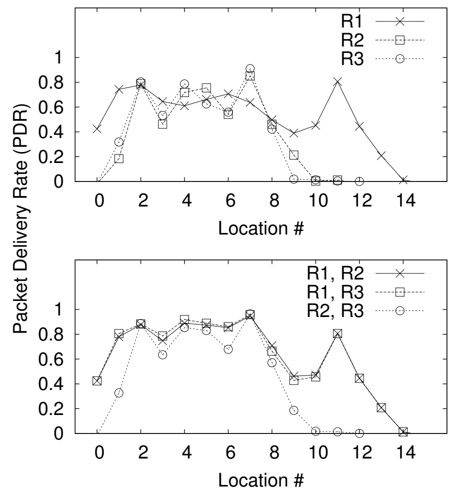

Figure 6: Receive diversity across receivers on opposite sides of the road (R1+R2/R3)is more beneficial than that from receivers on the same side (R2+R3).

2. *Select patterns with few main lobes.* Should we emphasize patterns with smaller or larger numbers of lobes? Prior work on wireless diversity [39, 12] has shown that diversity gains increase most significantly for a few additional receivers, and quickly level off after covering more than three receivers. We therefore include patterns with few main lobes.

Based on these heuristics we included 1-lobe, 2-lobe, and 3-lobe beams, as shown in the examples in Fig. 5, plus the standard omnidirectional pattern. To further reduce the number of beams, we only included patterns where each lobe has equal gain and omitted overlapping beams. In all, this leaves us with (a) 8 single-lobe beams—each beam shifted clockwise by the beam width of 45° from the previous beam such that they together cover the entire 360° space without overlap, (b) all $\binom{8}{2}$ possible two-lobe combinations of the single-lobe beams and (c) all $\binom{8}{3}$ three-lobe combinations of the single lobe beams. Together with the omnidirectional pattern, this yields a total of $1 + \sum_{i=1}^{3} \binom{8}{i} = 93$ patterns.

The single lobe beams are adopted from our previous work [15]. They possess the characteristic of having very low side lobes (a front-to-side lobe ratio of 18 dB), and about 8 dB extra gain over the omnidirectional pattern. To generate the multi-lobe beams, we super-impose the single lobe beams in MATLAB to derive the combined weights according to conventional antenna theory [30]. The algorithms to generate these patterns, and the exact element weights can be found in our extended technical report [32].

3.3 Long-term Learning of Parameters

Once a limited set of beam patterns are available, two key requirements to enable the use of parameter choices learned over a period of time are (a) an antenna-hardware-agnostic database, and (b) an online parameter update protocol.

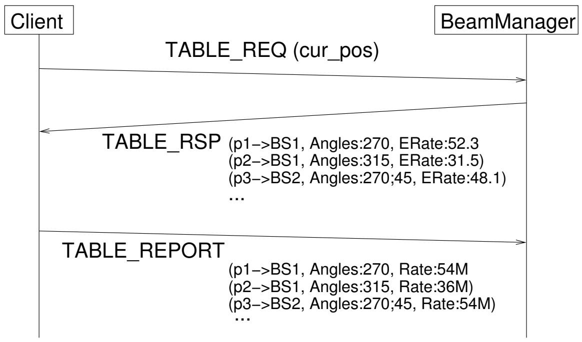

Figure 7: Protocol Overview. p1, p2, p3 represent GPS locations, and BS1, BS2 represent the anchor base stations. We define a GPS location (area) by considering only the first four digits after the decimal of the latitude and the longitude.

Antenna-agnostic database. The information stored in the BeamManager must be usable across clients with different hardware capabilities (e.g. number of antenna elements). To enable this feature, we leave to the client the choice of the exact beam and only identify the *direction* in which a beam should be formed. For each location, the BeamManager provides a ranked list of the single-lobe, two-lobe, three-lobe, and the omni beam that are likely to yield the highest throughput.

In addition, the direction must be stored relative to a common reference orientation. Clients can perform angular localization (similar to Mobisteer [28]) but this determines direction only relative to the current orientation of the antenna. We propose storing the direction relative to the global geographical north. By obtaining their heading from GPS and knowing the orientation of their antenna elements relative to their direction of movement (which depends on antenna mounting on the vehicle), the direction information in the database can be used in a client hardware independent manner. If antenna mounting information is not known, clients can use a one-time calibration procedure to determine the angular offset, whereby they rotate through all single-lobe beams to determine the throughput maximizing beam direction to a particular receiver. The offset is then the difference to the angle stored in the BeamManager for this receiver, corrected by GPS heading of the vehicle. Repeating this process for multiple receivers or locations improves the estimate.

BeamManager Online Update Protocol. Client probe configurations can be used by the BeamManager to continuously update its database. To this end, clients exchange and update the beam configuration table using a protocol shown in Figure 7. For every TABLE_REQ from a client, the BeamManager returns a TABLE_RSP containing the parameters corresponding to a set of nearby locations that the mobile client may traverse. The parameters include the anchor base station ID that must be added to the packets for correct diversity routing by the receiving BSes (we address this in Sec. 3.5 in more detail), the angles in which the beams should be formed, and the EWMA of bitrate observed by the clients at those locations. The clients choose the best supported bitrate below the ERate as starting point.

Algorithm 1 BeamManager pseudo-code

```
1: FUNCTION on_TABLE_REQ (pos)
2: for each posi close to pos do
3:     if exists unexplored combo then
4:         add posi, combo to ret_table;
5:     else if random() <= 0.9 then
6:         add combo with max. throughput;
7:     else
8:         randomly pick non-max. combo;
9:     end if
10: end for
11: send_TABLE_RSP (ret_table);
```

Algorithm 2 *on_position_change ()*

```
1: t = get_mapping_table(pos);
2: rcvr_angles = lookup_table(t, pos);
3: beam = form_beam (rcvr_angles);
4: configure_beam(beam);
5: tx_rx_packets();
```

The training-based approach used to build the BeamManager database is shown in Algorithm 1. Lines 2 and 3 in the on_TABLE_REQ pseudo code represent the initial phase when the BeamManager lets clients try out all possible parameter combinations. Lines 7-9 ensure that periodically the BeamManager table is updated with recent best configurations (using the TABLE_REPORT message (Figure 7)). We use a heuristic value of 0.1 to try out non-max combinations (and keep the overhead to 10%); more experiments are needed to determine the appropriate value that strikes the right tradeoff between convergence to changes and accuracy of the BeamManager.

The BeamManager applies a weighted update to resource parameters to ensure that a single observation due to momentary fluctuations or rare client capabilities (such as antennas with either too low or too high gain) will not change the settings significantly. Note that our idea is to use the BeamManager to capture settings that are generally applicable across several clients, and let the clients do further improvements through run-time adaptation. Also, since cycling through all the available parameter choices at a location requires a significant number of training samples to gain confidence in the parameter settings, an alternative approach involves measuring performance using a subset of beams (only the single-lobe beams), and estimating the composite beams (two-lobes and higher) that have the potential to improve performance. We leave the evaluation of this optimization as part of future work.

3.4 Client Adaptation

The basic client-side algorithm for location-based parameter changes is shown in Algorithm 2. The client looks up the angles at which a beam should be formed at the current location (Line 3), uses its antenna elements to form the appropriate transmit beam (Line 4 and 5), and transmits and receives packets till the location changes (Line 6). We call this basic algorithm R2D2-LOC. Further, we consider two variants that perform run-time adaptation to better match resource parameters to changing link conditions: R2D2-R that only adapts rate when the location changes, and R2D2-BR that adapts both beam and rate.

Algorithm 3 R2D2-R runtime rate adaptation.

```
FUNCTION on_pkt_loss_summary ()
if PER < low thresh then
    increase rate;
else if PER > high thresh then
    decrease rate;
end if
```

Table 1: Run-time adaptation.

SNR ↓ / PER→	High	Low
High	1. ↑ Diversity 2. ↓ Rate	↑ Rate
Low	1. ↓ Diversity 2. ↓ Rate	Maintain settings

R2D2-R: R2D2-R performs rate adaptation based on packet error rate (PER) much like RRAA [41], as shown concisely by Algorithm 3. The basic idea is that if the packet loss rate is too high, the rate is reduced, and if the loss rate is too low, the rate is increased. The rate is reset to that suggested by the BeamManager when the location changes.

R2D2-BR: The basic idea of R2D2-BR is captured in Table 1. First, it aims to increase diversity when packet errors are primarily due to high SNR variance (i.e., deep fades) and reduces diversity (increases directionality) if packet errors are primarily due to low average SNR for the given rate. It infers this from both SNR and PER measurements as follows. Deep fades are the likely cause if the SNR is sufficiently high (3dB above the SNR threshold for each rate in our implementation) to support communication at the current rate, but many packet errors are still observed. As indicated in the table, in this case the algorithm increases diversity. If no higher-diversity configuration is available, it reduces rate. Second, it aims to maintain the highest possible rate, by first exhausting the possible directionality and diversity configurations before lowering rate. For example, if SNR is low for the current rate and PER is high, the algorithm first reduces diversity (thereby increasing directionality). Once diversity reductions are exhausted, because it arrives at the narrowest possible beam, it reduces rate. An increase or decrease in diversity is obtained by switching to a different beam suggested by the BeamManager that has one more or one less lobe respectively.

Both R2D2-R and R2D2-BR depend on aggregate feedback from the infrastructure regarding the SNR and PER across multiple BSes. To calculate PER, we use packet selection diversity across the BSes, and for SNR, we use the highest SNR across packets from multiple BSes. While more sophisticated techniques can be used for aggregate PER and SNR calculation, we find that significant gains can be achieved even with these simple approaches.

3.5 Diversity Protocol

As shown in Fig. 4, all the BSes in this system are connected through the backplane to the Internet. For a set of consecutive locations, we designate an *anchor* BS to collect packets forwarded by all the receiving BSes [12]. The anchor BS also determines the performance (e.g. packet loss/reception and/or SNR) summary across all the receiving BSes (required for runtime adaptation).

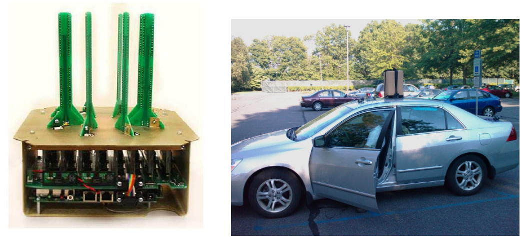

(a) (b)

Figure 8: (a) Transmitter with beamforming antenna, (b) Transmitter enclosed in a box and mounted on a car.

Essentially, we *split the data and control path functionalities* so as to effectively leverage the benefits of directionality and diversity: the receiving BSes are used as bridges to enable diversity in the data path (while not sending back individual acknowledgments), whereas the packet loss summary (i.e. control functionality) is sent back from the anchor BS, which is the actual point of association. In addition, since the anchor BS can change less frequently than the beams and receivers themselves (made possible by the assumption that the BSes are connected on a backplane), the effect of handoffs is also reduced [12]. Finally, to ensure that the mobile client is not *deaf* to downlink transmissions (including packet loss summary transmissions), the anchor BS at each location uses one of the currently receiving BSes (which is easily determined from the receiving packets).

MRD [27] and ViFi [12] discuss several components of the data transfer protocol required to leverage the benefits of receive diversity, which include (a) a block acknowledgment technique to ensure that the transmitter retransmits unrecovered packets, (b) the use of a reorder buffer on the anchor BS to reduce packet reordering and its consequent effects on higher level protocols, and (c) optimizations such as probabilistic relaying of packets to reduce the amount of forwarding traffic on the backhaul. However, both ViFi and MRD enable synchronous ACKs from one of the BSes primarily because their focus is on WiFi networks. We do not make the assumption of synchronous ACKs and build an ARQ protocol along with broadcast of packets with a reorder buffer on the anchor BS, both to reduce the complexity of determining which of the receiving base stations at that instant should send back the synchronous ACK, and to make the solution independent of WiFi. Further, we overprovision the backhaul and do not implement optimizations in the current prototype, to focus our concentration on the client-to-base station link and the tradeoff involved between directionality and diversity.

3.6 Prototype Implementation

We have implemented a prototype of R2D2 (Figure 8) by using the Phocus Array beamforming node [8] mounted on a car as the mobile client, and a set of stationary receivers that are small form-factor PCs with 6dBi gain external antennas. One of the receivers acts as the anchor BS to which the other receivers forward packets through additional wired and wireless interfaces (on orthogonal channels that the mobile client does not use).

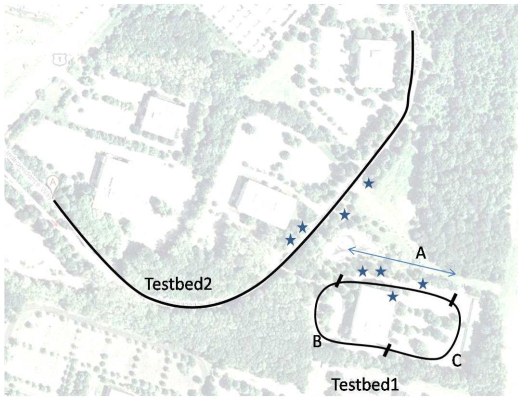

Figure 9: Testbeds: 1. Parking Lot (15-20mph), 2. Independence Way (35-40mph).

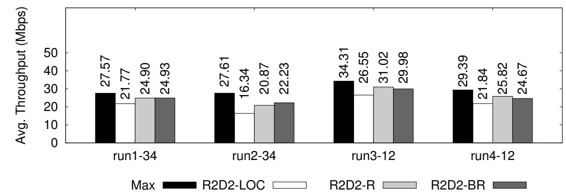

(a) Parking lot

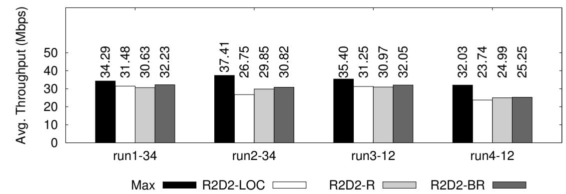

(b) Independence Way

Figure 10: Variants of R2D2. R2D2-R is a good compromise between performance and complexity.

On the Phocus Array, the directional antenna consists of an array of eight elements arranged in a regular octagon (see Figure 8). This directional antenna is electronically steerable, i.e., a specific beam pattern out of the several precomputed beams can be chosen on the fly via software control with less than 200 μs switching delay. The software control of the antenna is affected through an embedded computer running Linux [8, 28]. To disable the automatic retransmission behavior of the card (and the associated exponential backoff), we modified the device driver (MadWifi v0.9.2) on this node so that all *unicast* packets are sent only once. We implement R2D2's client adaptation algorithms (R2D2-LOC and R2D2-R) in user-space on this embedded computer. Our trace-driven simulations reveal that R2D2-R is effective in obtaining most of the benefits; hence we do not implement the prototype of the more complex R2D2-BR. We also implement prototypes of a beamsteering system and a receiver diversity system to match the behavior of Mobisteer [28] and ViFi [12] for comparison study.

On the stationary receiver nodes, we use a kernel-level click script [21] to forward packets between the *access* (mobile client to receivers) and *backhaul* (receivers to anchor) links. On the anchor BS, we use a multi-threaded C application in user-space to aggregate packets received from multiple BSes, and send periodic (every 100ms) block ACKs. We use libnet [22] and libpcap [2] to send and receive packets from user space. The block ACKs carry the PER information for packets received in the last 100ms interval. For ease of implementation, we (a) disable retransmissions on the backhaul links also to avoid issues related to delayed block ACKs (measured backhaul loss rate is <1% even without retries), and (b) implement the BeamManager on the client.

4. EVALUATION

We evaluate the performance of R2D2 in terms of the uplink throughput and average SNR compared to a state-of-the art receiver diversity system, a beamsteering system, and an oracle solution (that has the power to adapt on a per-packet basis to the best parameter combination for maximizing the throughput). In particular we use:

1. *AR-ViFi*: An enhanced version of ViFi [12], a system that exploits receiver diversity for vehicle-to-infrastructure (V2I) transmissions using regular omnidirectional antennas. The enhanced version also performs bitrate adaptation based on observed packet loss, so that comparison to R2D2 is fair. We also report selected results of the base version (i.e., transmitting at fixed rate) which is labeled ViFi in the graphs.

2. *Mobisteer [28]*: A system implementing beamsteering through location-based beam adaptation, for V2I communications. The beam is steered to a single receiver at any point in time, and bitrate adaptation works independently.

Our evaluation uses a combination of trace-driven simulations and prototype implementation. The trace-driven approach makes a rigorous evaluation tractable given the size of the search space (dictated by the number of beam patterns, bit-rates, receivers, GPS locations, receiver placement, etc.). In addition, it also enables us to approximate the performance of the oracle solution (referred to as MAX). As for the prototype implementation, it helps demonstrate the achievable throughput gains in real settings. We report results from our trace-driven simulations first.

4.1 Trace-Driven Methodology

For trace-driven analysis, we use two different realistic vehicle-to-roadside communication settings, as shown in Figure 9. The four receivers for each setting are placed as shown by the stars close to the paths. In each setting, we consider two receiver placements at different distances from the path, to emulate both high SNR conditions (when the receivers are close to the path), and low SNR conditions when the receivers are further away (corresponding to up to 20 dB drop in average SNR relative to the high SNR case).

We collect a set of packet traces with the transmitter in omni-directional mode and then emulate the effects of beamforming by adjusting the received SNRs at the access points. The transmitter broadcasts 200 ICMP packets (1350 bytes payload) per second using different 802.11g PHY rates (18, 36, 54Mbps) and transmission power levels. The use of broadcast mode suppresses several MAC-level features of 802.11 such as retransmissions, acknowledgments and RTS/CTS and enables us to measure the packet error encountered due to impairments suffered at the physical (PHY) layer. The receivers operate in monitor mode, in

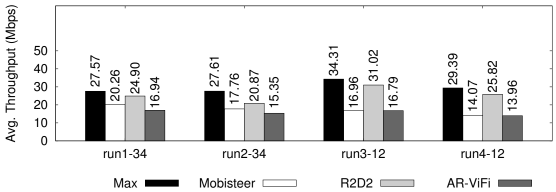

(a) Parking Lot

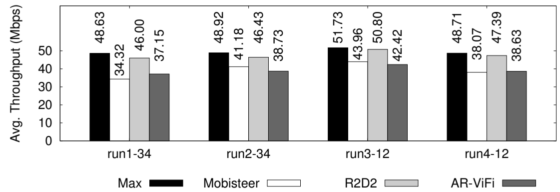

(b) Parking Lot, High SNR

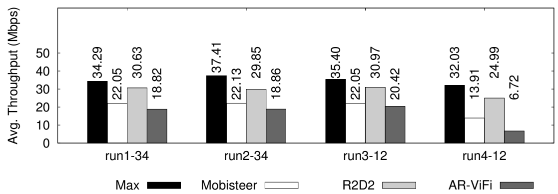

(c) Independence Way

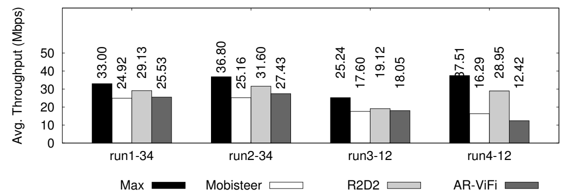

(d) Independence Way: High SNR

Figure 11: Average throughput obtained by R2D2-R, AR-ViFi, Mobisteer, and the oracle (MAX).

which the node can passively listen to all data on a particular channel without being associated with any AP. Receivers utilize tcpdump [2], which give them relevant information on a per-packet basis from both the PHY and MAC layers. In addition, all the nodes (clients and BSes) continuously log their location and speed information using a GPS device. The system time on each node is set to the GPS time so that the system clocks of all nodes are synchronized. The transmitter includes its timestamp in the ICMP packet's payload so that the receiver can correlate the location from which each packet was transmitted. Two sets of four runs each are taken on a circular path in a parking lot (Testbed1) around a building, and on a road named Independence Way (Testbed2) with a speed limit of 40 mph.

Using these traces, we can emulate the effect of beamforming by scaling received SNRs at all access points. We observe (a) with a phase array directional antenna [8] available to us, and (b) theoretically with MATLAB simulations, that a beam covering two well separated receivers (i.e. having two main lobes) will have approximately 3dB lower gain than a beam with a single main lobe pointing to a receiver, and a beam covering three receivers will have 3dB lower gain than the two receiver case. Based on this observation, we pick four transmit power levels: 17dBm, 14dBm, 11dBm and 8dBm. When transmitting at 17dBm, we interpret this as the beam being pointed to one of the receivers, at 14 dBm to any combination of two receivers, at 11dBm to any combination of three receivers, and at 8dBm to cover all four receivers. We assume that packets are dropped for receivers outside the emulated beam. We process these traces, and implement all the candidate algorithms using a simulator written in C (details in the extended technical report [32]).

4.2 Trace-Driven Results

We first present the performance of R2D2 compared to the oracle and other approaches, and then elaborate on the adaptation algorithm behavior.

4.2.1 Performance Improvement

Variants of R2D2 vs MAX: Figure 10 shows the average throughput with the three variants of R2D2 in four runs in each of the different settings when the average SNR is low, i.e. the receivers are farther away from the path. The location-based database is trained using two runs other than the one under consideration. For instance, run1-34 indicates the performance in run1 when using 3rd and 4th runs for training the database.

The difference between R2D2-LOC and R2D2-R or R2D2-BR makes the case for additional run-time adaptation. In both settings, the performance of R2D2 with run-time adaptation comes close to MAX. Between R2D2-R and R2D2-BR, the minor difference across all settings and runs shows that just run-time rate adaptation with location-based beam adaptation is quite effective in reaping most of the benefits. This is a significant observation given that the implementation complexity of R2D2-BR is substantially higher than that of R2D2-R; in particular, R2D2-BR involves tuning several parameters at run-time in a short period of time, and requires significant control messaging between the beam manager and a transmitter. Hence, in the rest of the paper, we only report results for R2D2-R. We also implement only R2D2-R in our prototype.

R2D2-R vs AR-ViFi vs MobiSteer: Figure 11 shows the average throughput obtained by R2D2-R relative to AR-ViFi, MobiSteer, and MAX. Note that we build the location-based database for MobiSteer in the same manner as that for R2D2 (using two runs other than the run under consideration). Figure 11(a) shows the throughput for the parking lot case when the average SNR is low. In this case, R2D2 performs significantly better than existing approaches and their extensions. The difference between AR-ViFi and MAX clearly shows that even after extending ViFi, there is still a significant scope for improvement. Finally, R2D2's performance gain over AR-ViFi and Mobisteer shows that carefully trading off directionality and diversity and jointly adapting rate can take the uplink throughput significantly closer to MAX. Similar observations can be made on the Independence Way path in Figure 11(c).

Figures 11(b) and 11(d) plot the throughput under high SNR conditions, i.e. when the receivers are very close to the paths. Even in this case, R2D2 performs better than other algorithms and is close to what MAX can achieve. Observe

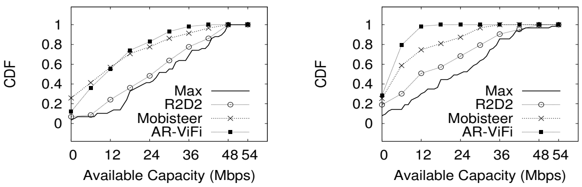

(a) Parking Lot (b) Independence way

Figure 12: CDF of throughput obtained by the candidate algorithms at several client locations.

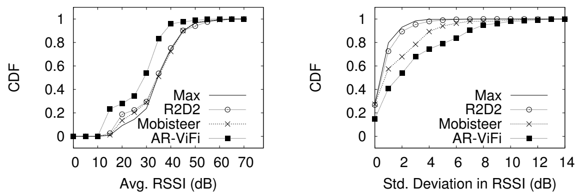

(a) Mean RSSI distribution (b) Std. deviation in RSSI

Figure 13: The distribution of average, and std. deviation of, RSSI for candidate algorithms relative to MAX. R2D2 **improves link robustness by increasing the average RSSI and reducing the variance.**

also that AR-ViFi performs close to or better than Mobisteer in most cases, unlike the low-SNR case. This is inline with the theoretical predictions in Figure 3.

Figure 12 shows the CDF of throughput for R2D2, Mobisteer and AR-ViFi for run 3, using runs 1 and 2 for training the database. The CDF clearly demonstrates that R2D2 has high throughput at higher percentage of locations. For instance, in the parking lot case, the median throughput with R2D2 is 25 Mbps, whereas that with AR-ViFi and Mobisteer are 9 and 8 Mbps respectively.

Improved mean SNR and reduced SNR variance: To highlight the effect on SNR better, Figure 13 shows the distribution of the mean and standard deviation in the received signal strength indicator (RSSI) achieved by each of the algorithms relative to MAX. RSSI is an estimate of the signal energy at the receiver during packet reception, measured during the PLCP headers of arriving packets and reported on proprietary (and different) scales. The Atheros cards we use, for example, report RSSI in dB [33]. For a single link, since SNR = RSSI / noise floor, and since the noise floor is relatively constant, here, we use RSSI to represent SNR. The CDFs shown in Figure 13 use the average and standard deviation in RSSI calculated over successive 100-millisecond intervals. As RSSI is available only if the corresponding packet was successfully decoded, and since we need the RSSI of lost packets to calculate the variance, we make the assumption that lost packets have an RSSI value just below the receive threshold for the bit-rate used to send that packet. Using this assumption, we report the "best-case" average and standard deviation of RSSI.

Figure 13(a) shows that both R2D2 and Mobisteer come close to approximating the maximum achievable average RSSI. The median for both algorithms is 5 dB higher than the

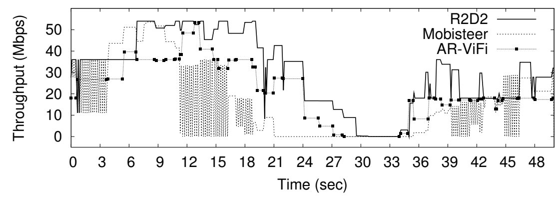

(a) Parking Lot

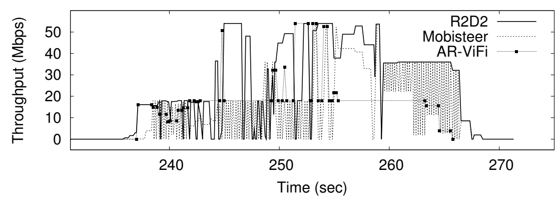

(b) Independence Way

Figure 14: Instantaneous throughput obtained by R2D2 **compared to AR-ViFi and Mobisteer.**

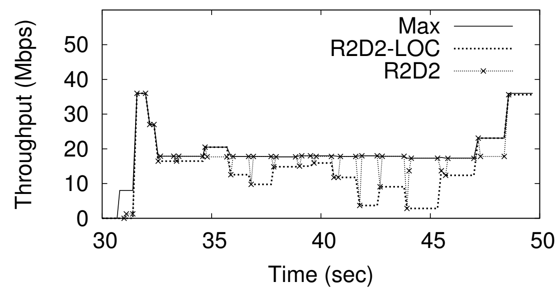

Figure 15: Parking Lot: zoomed-in.

pure diversity based scheme. However, Figure 13(b) shows that both pure directionality and diversity based schemes are unable to effectively deal with SNR variance (due to deep fades). The distribution of standard deviation for both algorithms has heavy tails with a maximum deviation of 14 dB. The graph also shows that even MAX observes standard deviations of up to 4dB. R2D2 is able to closely approximate the behavior of MAX (and eliminate a large part of the heavy tails of the distribution).

4.2.2 Algorithm Behavior

Efficient tracking of channel fluctuations: We compare the the behavior of the candidate algorithms in Figure 14. Figure 14(a) shows the throughput with time (averaged every 200 milliseconds) for the parking lot case with the low SNR setup for run3 using runs 1 and 2 for training the database. The graph clearly shows that in many regions R2D2 makes a better choice of parameters than AR-ViFi or Mobisteer. Figure 14(b) shows similar result for the Independence Way setting with the low SNR case.

In Figure 15, we zoom into the behavior of R2D2-R and compare its performance with R2D2-LOC and MAX, to show the efficacy of additional run-time adaptation over using location-based training for parameter selection. The graph clearly shows that location mispredicts the best choice in several instances making R2D2-LOC perform worse than

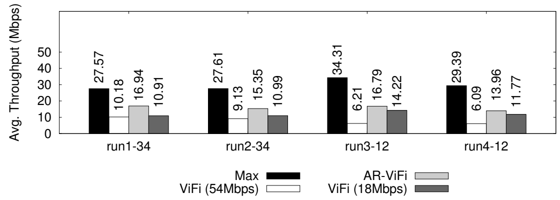

(a) Parking lot

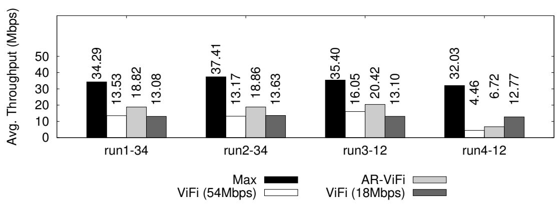

(b) Independence Way

Figure 16: Variants of ViFi. Runtime adaptation benefits pure diversity-based approaches as well.

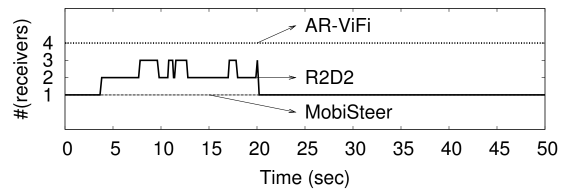

(a) Parking lot

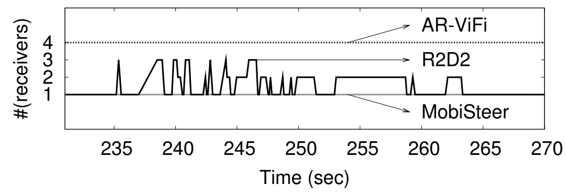

(b) Independence way

Figure 17: Receiver combination selections: #(receivers) = 1 corresponds to pure directionality while #(receivers) = 4 corresponds to pure diversity.

MAX; whereas R2D2-R detects these conditions and adapts well to match the performance of MAX.

Variants of ViFi: For the sake of completeness, we also demonstrate the benefits of run-time rate adaptation to diversity-based solutions, which have been acknowledged but not implemented in past solutions [27, 12]. In Figure 16, we compare AR-ViFi with two static rate settings for ViFi (54 Mbps and 18 Mbps) to show these expected gains.

Intelligent choice of BSes: Figure 17 shows the combination of receivers chosen with time by the transmitter when using R2D2, in comparison to the extremes of choosing all receivers (as in AR-ViFi) and only one receiver (as in Mobisteer). We use information from run3-12 for the parking lot case and run4-12 for the Independence Way case. The graph demonstrates that in several locations a middle-ground between the extremes, i.e. choosing a subset of visible receivers is the best choice. Note that while Mobisteer and R2D2 overlap in Figure 17(a) after 20 seconds, the choice of rate is different due to run-time adaptation, leading to R2D2's improved performance over Mobisteer (see Figure 14(a)).

We now describe our prototype implementation results.

4.3 End-to-end Evaluation

We perform our prototype evaluation in the parking lot setup, as shown in Figure 9. The BeamManager is trained on the same path and the best set of receivers at each location is determined. A location is uniquely determined by considering only the first four digits after the decimal point of the GPS coordinates (latitude and longitude). This method defines unique rectangles (of area 11x8m) that enables sharing beam configuration and rate selection information across clients. We perform angular localization on the same path to determine the beam angle that is best at each location for each receiver, similar in methodology to Mobisteer [28]. We also take the conservative approach of lowering the bit-rate in response to a block ACK loss since a majority of these losses occur on the access link (validated via experiments).

Figure 18 shows the total amount of data transferred by AR-ViFi, R2D2-R, and Mobisteer in three runs of the parking lot, for both the high and low SNR cases. R2D2 is able to transfer larger amounts of data in the same time duration, relative to both diversity-based (AR-ViFi), and directionality-based (Mobisteer) approaches. For instance, in the low SNR regime, R2D2 transfers 45% and 154% more data than AR-ViFi and Mobisteer respectively. In the high SNR regime, the gains are lower (23% over AR-ViFi, and 65% over Mobisteer) since the high SNR itself provides immunity from deep fades, while supporting higher bit-rates.

Figure 19 shows the CDF of the bitrates used by the mobile transmitter with AR-ViFi, R2D2-R, and Mobisteer in three runs of the parking lot, for both the high and low SNR cases. The graph demonstrates that R2D2's throughput gains are primarily a result of using higher bitrates (which in turn arise from the improved average and std. dev. in SNR). Note that Figures 18 and 19 are not directly comparable since the bitrate CDF plots are based on the bitrate used by the transmitter (and not the bitrates of the successfully decoded packets).

5. LIMITATIONS AND DISCUSSION

Compared to an omni-directional beam, the cost of choosing a "wrong beam", e.g. due to lack of enough information at a location or due to inaccurate position estimation due to GPS errors [20], can be very high. This can be easily appreciated by looking at Figure 20, where six beams out of the eight provide significantly lower packet delivery rate than the omni-directional beam. This observation indicates that the training phase is a very important component of the system, requiring a significant number of training samples to gain enough confidence in the parameter setting at each location, contrary to what we had initially anticipated. This problem is especially challenging due to (a) high client mobility, and (b) hardware differences between clients. For (a), we envision that, by distributing the training phase across a large client-base, resource parameter settings at a location can be used across different clients. As for (b), we acknowledge that our current approach may need extensions as it does not account for hardware differences between clients.

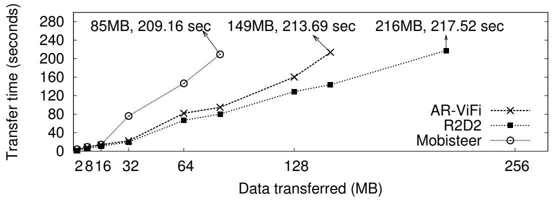

(a) Parking Lot

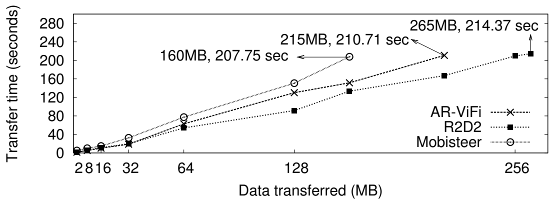

(b) Parking Lot, High SNR

Figure 18: Data transfer time (in seconds) using AR-ViFi, R2D2, and Mobisteer. R2D2 transfers up to 45%, and up to 154% more data than AR-ViFi and Mobisteer respectively.

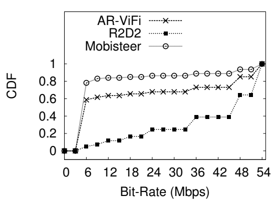

(a) Parking Lot

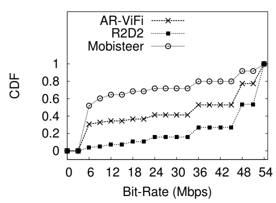

(b) Parking Lot, High SNR

Figure 19: CDF of bitrate used by the mobile transmitter with AR-ViFi, R2D2, and Mobisteer. The median bitrate used by R2D2 is ∼48Mbps in the low SNR regime while that used by AR-ViFi and Mobisteer is ∼6Mbps.

In Section 2, we discussed the effect of number of transmit antenna elements on the theoretical tradeoff between directionality and diversity. An additional issue that we do not address in this paper is that while a greater number of elements can form thinner beams, the duration of validity of a beam reduces because of client mobility. In general, the effect of the aging of training data, and the BeamManager's efficacy in keeping the parameter database up-to-date is an interesting avenue for future work.

A key focus of this paper was to introduce, and address the fundamental tradeoff between directionality and diversity, on the uplink, in a standards-independent manner. While we used WiFi devices (primarily due to their ready availability) to achieve this objective, we believe that our solution itself is equally applicable to other wireless technologies such as WIMAX [11] and LTE [9] assuming the necessary extensions are in place (such as the ability to connect to multiple base stations simultaneously). Note that current CDMA standards already allow data combining from multiple base stations on the uplink [10]. Also, while we focus on uplink connectivity, diversity and directionality can also be leveraged in the downlink direction. However, this requires additional functionality on the base stations e.g. synchronization. While we do not expand on this direction in this paper, this topic is already of significant interest to the LTE and WIMAX standards bodies, and forms an interesting problem for future work. Finally, this work does not consider the issues associated with multiple clients and base stations, which are beyond the scope of this paper.

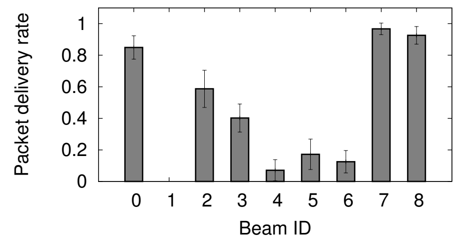

Figure 20: Angular Localization. Cost of choosing the "wrong" beam is high. Beam ID 0 = Omni.

From a deployment perspective, while it is unlikely that mobile devices such as cellphones and PDAs will incorporate multi-antenna systems with more than three elements due to their small form-factor, we envision the evolution of relay devices on moving vehicles that can afford to carry such large form-factor antenna systems. A similar proposal for a different purpose was made by Rodriguez et al. [34].

6. RELATED WORK

Smart antennas are an integral part of most future wireless standards (WiMAX [4], LTE [9], 802.15.3c [3], etc.). Beamforming (directionality) is one of the core features adopted by operators to meet the high spectral efficiency requirement of future mobile applications. As mentioned before, macro-diversity on the other hand, is already being used in CDMA cellular networks, following the work of Viterbi et al in [39] to improve performance for cell-edge users. While several of the future wireless standards advocate the concept of directionality and diversity, they deal predominantly with protocol issues and do not consider algorithms or systems that instantiate the core mechanisms. The algorithms are left open for implementation and innovation by individual vendors. Further, the bulk of mechanisms and sophistication in today's cellular networks is downlink-oriented. Note that the directionality-diversity tradeoff explored in this work is uplink-specific, where the receivers (BS or AP) can collaborate unlike the case of mobile clients on the downlink.

Several research works have looked at leveraging the benefits of directionality and diversity, albeit in isolation. Mobisteer [28] attempts to improve the uplink performance by forming a beam directed at a single receiver. Solutions in the WiFi domain [12, 14, 17, 26] also utilize *opportunistic reception* of packets due to omni-directional transmissions in order to mask-off packet loss at any individual receiver due to SNR fluctuations. However, none of the above works explore the issues involved in combining diversity and directionality. Tao et al. [38] address this issue to some extent in an ad hoc network setting. They conclude that diversity and directionality have conflicting parameter settings and that further exploration is required to incorporate both approaches. They also focus on the complementary issue of developing a MAC protocol to enable nodes in an ad hoc network to leverage directionality and diversity together. In contrast, this paper addresses the directionality-diversity tradeoff by devising the system support and run-time adaptation required for choosing the beams, receivers and rate jointly in a mobile environment.

Recently, a number of research efforts have focused on different aspects of improving connectivity to moving vehicles. Bychkovsky et al. [16] study the possibility of using organic WiFi deployments for providing improved uplink connectivity to moving vehicles. They investigate the effectiveness of a caching technique to reduce the overhead of IP address acquisition. Ott et al. [31] discuss an architecture and protocol to make applications disconnection-tolerant by maintaining application sessions despite connectivity interruptions. System support for fast association to APs and optimizations at the TCP level to improve throughput for moving vehicles is discussed by Eriksson et al in [18]. Rodriguez et al [34] introduce a wireless multi-homed device (MAR) for moving vehicles that dynamically aggregates channels (and hence bandwidth) across several technologies to meet the bandwidth requirements of moving users. Detailed studies on the factors affecting connectivity to moving vehicles is performed by Hadaller et al in [20], where they conclude that lack of environmental awareness is the fundamental underlying cause of several problems. Our exploration in this work is complementary to the above approaches and hence could be integrated. Further, our location-based beam and rate selection algorithm instantiates environmental (location) awareness into the adaptation process, the benefits of which have already been shown in several other works [28, 29, 35, 24].

7. CONCLUSIONS

In this paper, we highlight that directionality and diversity—two commonly used mechanisms for improved uplink connectivity from mobile clients—are only extremes of a solution, and a combination of these mechanisms, can significantly improve throughput. This, however, requires beam shape adaptation in addition to commonly used beamsteering. A narrow beam tends to cover fewer receivers, thereby reducing diversity. To obtain gains from both directionality and diversity, the system has to select a channel-dependent multi-lobe beam that covers multiple nearby receivers. To rapidly converge to suitable beam configurations on a fast moving node, we have proposed a two-stage algorithm that (a) obtains a candidate set of beam shapes from a centralized repository based on the current node location, and (b) selects the appropriate number of lobes within this candidate set through run-time adaptation. This run-time component increases diversity when packet loss is likely due to deep fades (i.e. high SNR variance), and reduces diversity (or increases directionality) if packet loss is likely due to low mean SNR. In addition, the two-stage algorithm also jointly adapts modulation rates.

Experiments with a smart antenna system mounted on a vehicle and roadside WiFi access points have demonstrated feasibility and significant throughput gains. We believe that this tradeoff, and our solution are not limited to WiFi networks, and that they may be particularly suited to providing high throughput communications to mobile users via relay devices on vehicles.

ACKNOWLEDGMENTS

We thank all the anonymous reviewers for their very useful comments. Many thanks to Rajesh Mahindra for all the help during experiments.

8. REFERENCES

[1] 3G and WiMax Femtocells: The Operator View. http://tinyurl.com/22w9uj.

[2] http://www.tcpdump.org/.

[3] IEEE 802.15 WPAN Millimeter Wave Alternative PHY Task Group 3c (TG3c). http://www.ieee802.org/15/pub/TG3c.html.

[4] IEEE-SA Get IEEE 802.16 LAN/MAN Broadband Wireless LANs. http://tinyurl.com/2nvbwn.

[5] Infonetics research. http://www.infonetics.com.

[6] Microsoft giving workers free ride — with its own bus service. http://tinyurl.com/djnf5b.

[7] Notebooks surpass pass desktops in U.S. retail sales. http://news.cnet.com/2100-1044_3-6033967.html.

[8] Phocus Array Antenna by Fidelity Comtech. http://www.fidelity-comtech.com/.

[9] Technical Specificati on Group Radio Access Network: Physical Layer Aspects for Evolved Universal Terrestrial Radio Access (UTRA). 3GPP TR 25.814 v.7.0.0 3GPP.

[10] UE Radio Transmission and Reception (FDD). 3GTS 25.1010 3GPP.

[11] WiMAX Forum. http://www.wimaxforum.org/home/.

[12] A. Balasubramanian, R. Mahajan, A. Venkataramani, B. Levine, and J. Zahorjan. Interactive WiFi Connectivity For Moving Vehicles. In *SIGCOMM*, '08.

[13] S. Barberis and E. Berruto. A CDMA-based radio interface for third generation mobile systems. *Mob. Netw. Appl.*, 2(1):19–29, 1997.

[14] S. Biswas and R. Morris. ExOR: opportunistic multi-hop routing for wireless networks. *SIGCOMM CCR*, 35(4):133–144, 2005.

[15] M. Blanco, R. Kokku, K. Ramachandran, S. Rangarajan, and K. Sundaresan. On the Effectiveness of Switched-beam Antennas in Indoor Environments. In *Passive and Active Measurement Conference (PAM)*, 2008.

[16] 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 *MobiCom*, 2006.

[17] S. Chachulski, M. Jennings, S. Katti, and D. Katabi. Trading structure for randomness in wireless opportunistic routing. *SIGCOMM CCR.*, 37(4):169–180, 2007.

[18] J. Eriksson, H. Balakrishnan, and S. Madden. Cabernet: vehicular content delivery using WiFi. In *MobiCom*, 2008.

[19] D. Gesbert, M. Shafi, D. shan Shiu, P. J. Smith, A. Naguib, and S. Member. From theory to practice: an overview of MIMO space-time coded wireless systems. *IEEE JSAC*, 21:281–302, 2003.

[20] D. Hadaller, S. Keshav, T. Brecht, and S. Agarwal. Vehicular opportunistic communication under the microscope. In *MobiSys*, 2007.

[21] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek. The Click modular router. *ACM Trans. Comput. Syst.*, 18(3):263–297, 2000.

[22] The Libnet Packet Construction Library. `http://www.packetfactory.net/libnet/`, Feb 2007.

[23] M. Corner and B. Levine. UMass DieselNet.

[24] R. Mahajan, J. Zahorjan, and B. Zill. Understanding WiFi-based connectivity from moving vehicles. In *IMC*, 2007.

[25] Mitch Wagner, Information Week. How iPhone 3.0 May Revolutionize The Smartphone Industry. http://tinyurl.com/d4lbsp, March 2009.

[26] A. Miu, H. Balakrishnan, and C. E. Koksal. Improving loss resilience with multi-radio diversity in wireless networks. In *MobiCom*, 2005.

[27] A. Miu, H. Balakrishnan, and C. E. Koksal. Improving loss resilience with multi-radio diversity in wireless networks. In *MobiCom*, 2005.

[28] V. Navda, A. P. Subramanian, K. Dhanasekaran, A. Timm-Giel, and S. Das. Mobisteer: using steerable beam directional antenna for vehicular network access. In *MobiSys*, 2007.

[29] A. J. Nicholson and B. D. Noble. Breadcrumbs: forecasting mobile connectivity. In *MobiCom*, 2008.

[30] S. J. Orfanidis. *Electromagnetic Waves and Antennas*. ONLINE BOOK: http://www.ece.rutgers.edu/ orfanidi/ewa/, 2004.

[31] J. Ott and D. Kutscher. A disconnection-tolerant transport for drive-thru internet environments. In *In IEEE Infocom*, March 2005.

[32] K. Ramachandran, R. Kokku, K. Sundaresan, M. Gruteser, and S. Rangarajan. R2D2: Regulating Beam Shape and Rate as Directionality meets Diversity. Technical report, NEC Laboratories America, 2008.

[33] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan. Measurement-based models of delivery and interference in static wireless networks. *SIGCOMM CCR*, 36(4):51–62, 2006.

[34] P. Rodriguez, R. Chakravorty, J. Chesterfield, I. Pratt, and S. Banerjee. Mar: a commuter router infrastructure for the mobile internet. In *MobiSys*, 2004.

[35] P. Shankar, T. Nadeem, J. Rosca, and L. Iftode. CARS: Context Aware Rate Selection for Vehicular Networks. In *ICNP*, 2008.

[36] A. Subramanian, P. Deshpande, J. Gaojgao, and S. Das. Drive-by localization of roadside wifi networks. In *IEEE INFOCOM*, April 2008.

[37] A. P. Subramanian, V. Navda, P. Deshpande, and S. R. Das. A measurement study of inter-vehicular communication using steerable beam directional antenna. In *ACM VANET*, 2008.

[38] Z. Tao, T. Korakis, F. Liu, S. Panwar, J. Zhang, and L. Tassiulas. Cooperation and directionality: Friends or foes? In *IEEE ICC*, 2008.

[39] A. Viterbi, A. Viterbi, K. Gilhousen, and E. Zehavi. Soft handoff extends cdma cell coverage and increases reverse link capacity. *IEEE JSAC*, 12(8):1281–1288, Oct 1994.

[40] Wikipedia. Web 2.0. http://tinyurl.com/dqt86.

[41] S. H. Y. Wong, S. Lu, H. Yang, and V. Bharghavan. Robust rate adaptation for 802.11 wireless networks. In *MobiCom*, 2006.

Appendix

In this section, we use outage probability to model the resulting throughput on a link for a system that exploits both directionality and diversity. Outage probability refers to how often (probability) does the bit error rate (BER), or equivalently SINR, experienced falls below a certain threshold. It is both a popular and practical measure for robustness to fading, especially for block fading where it can directly be related to frame/packet error rate. It can be measured by determining the probability that the mutual information of communication (capacity) is less than the information rate. We consider an independent, quasi-static, frequency nonselective Rayleigh fading (complex channel coefficients being uncorrelated and circularly symmetric Gaussian random variables with zero mean and unit variance) along with free space path loss for our channel model. Let a channel realization between an n-element Tx and an omni-directional Rx be denoted by h_{tr} (captures fading). Now, the mutual information for a given channel realization in the case of beamforming is given by the asymptotic Shannon capacity formula,

$$C(h_{tr}) = \log_2(1 + n\rho|h_{tr}|^2))$$

where $\rho|h_{tr}|^2$ is the average receiver SNR (absolute value) at Rx for an omni-directional transmission (ρ is the average receiver SNR based on path loss alone). If R is the information (data) rate applied to the system (normalized to bandwidth in bits/sec/Hz), then the outage probability can be given as,

$$P_o(R, n) = Pr[C(h_{tr}) < R] = Pr\left[|h_{tr}|^2 < \frac{2^R - 1}{n\rho}\right]$$

By way of definition of $|h_{tr}|$, $|h_{tr}|^2$ follows an exponential distribution. Hence, on averaging over all possible channel realizations, we have,

$$P_o(R, n) = 1 - e^{-\left(\frac{2^R-1}{n\rho}\right)} = 1 - e^{-\left(\frac{S}{n\rho}\right)}$$

where $S = 2^R - 1$. When a beam of width $\frac{360^o}{n}$ is formed with n elements, let the number of receivers falling in the reception zone of the beam be given by the function $\ell(n)$. Note that, $\ell(n)$ is a monotonically decreasing function, with number of accessible receivers decreasing with finer beamwidths. Now, the resulting probability of successfully receiving the packet/frame can be given as,

$$P_s(R, n) = 1 - \prod_{i \in \ell(n)} \left(1 - e^{-\left(\frac{S}{n\rho_i}\right)}\right)$$

where ρ_i is the average SNR at receiver i corresponding to an omni-directional transmission. The resulting throughput can be given as,

$$T(R, n) = R \cdot \left(1 - \prod_{i \in \ell(n)} \left(1 - e^{-\left(\frac{S}{n\rho_i}\right)}\right)\right)$$