MV-MAX: Improving Wireless Infrastructure Access for Multi-Vehicular Communication

David Hadaller University of Waterloo David R. Cheriton School of Computer Science

dthadaller@uwaterloo.ca

Srinivasan Keshav University of Waterloo David R. Cheriton School of Computer Science

keshav@uwaterloo.ca

Tim Brecht
University of Waterloo
David R. Cheriton School of
Computer Science

brecht@uwaterloo.ca

ABSTRACT

When a roadside 802.11-based wireless access point is shared by more than one vehicle, the vehicle with the lowest transmission rate reduces the effective transmission rate of all other vehicles. This performance anomaly [9] degrades both individual and overall throughput in such multi-vehicular environments. Observing that every vehicle eventually receives good performance when it is near the access point, we propose MV-MAX (Multi-Vehicular Maximum), a medium access protocol that opportunistically grants wireless access to vehicles with the maximum transmission rate. Mathematical analysis and trace-driven simulations based on real data show that MV-MAX not only improves overall system throughput, compared to 802.11, by a factor of almost 4, but also improves on the previously proposed time-fairness scheme [20, 22, 15] by a factor of more than 2. Moreover, despite being less fair than 802.11, almost every vehicle benefits by using MV-MAX over the more equitable 802.11 access mechanism. Finally, we show that our results are consistent across different data sets.

Categories and Subject Descriptors

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

Keywords

802.11p, Infostations, Vehicular Communication, Opportunistic Connectivity, Delay Tolerant Networking

1. INTRODUCTION

Users desire high bandwidth Internet connectivity on the road, much as in their office or home. Currently, 3G technologies such as EvDO and HSDPA address this need. However, because these operate on a licensed spectrum and aim for complete coverage, they are typically prohibitively expensive for transferring bulk data such as multimedia con-

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.

SIGCOMM'06 Workshops September 11-15, 2006, Pisa, Italy. Copyright 2006 ACM 1-59593-417-0/06/0009 ...\$5.00.

tent. Fortunately, even higher throughput, albeit with disconnection episodes, can be provided by 802.11-based road-side wireless access points or *infostations* [16, 5]. The latter solution is attractive due its high capacity, low capital cost, and incremental deployability.

We envision that vehicles of the future will be routinely installed with 802.11 NICs allowing occupants to opportunistically download travel information, shopping coupons, news summaries, and perhaps even music and videos, to be viewed by passengers or listened to by the driver or viewed at the destination. This vision is also shared by the IEEE 802.11p working group and the Dedicated Short Range Communications (DSRC) [4] branch of the US Intelligent Transport System (ITS), discussed further in Section 6. Figure 1 depicts multiple vehicles passing a roadside access point.

Because vehicles in motion have short connection durations with roadside access points, efficient use of this duration is important. With a single vehicle, this is relatively straightforward and has been studied in past work [16]. However, when multiple vehicles are within range of an access point, the wireless medium is shared. This creates the following problem: standard contention-based media access control used in 802.11 suffers from a performance anomaly [9] when multiple vehicles with different MAC-layer data rates are within range of the same access point. Essentially, the vehicle with the lowest data rate (and likely the poorest signal quality) slows all other vehicles down to its rate, resulting in poor use of the wireless medium and reduced performance for all vehicles.

MV-MAX solves this problem. Typically, as a vehicle passes a roadside access point, it experiences poor signal quality when entering wireless range, followed by a stronger signal as the vehicle nears the access point, and a weakening signal after the vehicle passes the access point. This typically translates into a low transmission rate at the fringes of the access point's coverage area and a high transmission rate near the center of the coverage area. When multiple vehicles are in range of an access point, the vehicles on the fringe of coverage degrade the performance of all other vehicles. Consequently, if we temporarily deny access to a vehicle when it has poor signal quality, permitting it access only when it



Figure 1: Vehicles with short range communication devices drive past an 802.11 access point

has good signal quality, not only does every vehicle eventually get an allocation of transmission capacity, making the resource allocation roughly equal over the long term, but the overall system throughput also increases.

The contributions of our work are as follows. To our knowledge, we are the first to identify the impact of the 802.11 performance anomaly on the performance of multivehicular wireless access. Second, we show how to remove this anomaly using MV-MAX, a simple opportunistic scheduling algorithm. Third, we present an accurate mathematical model of a multi-vehicular wireless access system and use it to analyze the performance of 802.11, MV-MAX, and the previously proposed time-fairness scheme [20, 22, 15]. Finally, we use trace-based simulations, using real data collected from two different environments, to study the performance benefits of MV-MAX and compare it with both 802.11 and the time-fair resource allocation scheme.

The paper is laid out as follows. We describe the performance anomaly problem and an overview of MV-MAX in Section 2. We then present a mathematical model of multivehicle access and use it to analyze MV-MAX in Section 3. In addition to mathematical modeling, we show, using trace-based simulations in Section 4, that this technique not only improves overall system throughput by 3.9 times that of 802.11, but also improves the performance received by nearly every vehicle. We also compare MV-MAX with an alternative scheme that shares time equally among vehicles (time fairness), and show that our scheme provides 2.2 times more throughput. Our results are consistent across data sets collected both by us and by researchers at Intel Research at Cambridge. We present a discussion of our work in Section 5 and related work in Section 6.

2. PROBLEM DESCRIPTION AND SOLUTION OVERVIEW

Similar to Drive-Thru Internet [16], we consider a single wireless access point (AP) placed at the roadside that is used for Internet access by computers embedded in vehicles as these vehicles drive by (see Figure 1). At a given moment, more than one vehicle may use the access point. The goal of each vehicle is to maximize the amount of data transferred to and from the access point as it drives by, while the goal of the access point provider is to maximize the sum of individual vehicular throughputs.

The (a) hardware configuration of the access point and vehicle-based client, such as its antenna and NIC, (b) modulation scheme and channel coding chosen by each NIC, and (c) RF environment jointly determine the transmission rate at every point in the coverage range of an access point. Each NIC chooses a modulation and encoding to maximize the transmission rate given the current RF environment and its own hardware capabilities. Although this decision process can be complex, the 802.11 standard decrees that the resultant channel transmission rate must be one of a pre-defined set of rates. For instance, 802.11g has 12 different combinations, resulting in transmission rates between 1 Mbps and 54 Mbps; 802.11b has four combinations yielding rates between 1 Mbps and 11 Mbps. Due to the high variability of RF signals, each point in the range potentially has a different wireless transmission rate, as experimentally verified in Figure 2. Moreover, at the same point in the roadway, different vehicles may choose different transmission rates depending on

the rate selection algorithms [19, 12] implemented at their NIC.

When there is a single vehicle in range of the AP, variability is not a problem. However, if more than one vehicle is in range of the AP, then a well-known performance anomaly arises [9]. Because transmitting one maxmium-sized 1500 byte IP packet at 1 Mbps takes approximately 50 times longer, depending on 802.11 preamble and header configuration [7], than transmitting one packet at 54 Mbps, a "slow" user, transmitting at 1 Mbps, can greatly slow down "faster" users. As we show in Section 4, this results in poor overall performance.

In contrast, MV-MAX allocates the wireless medium to the vehicle(s) experiencing the best SNR, which roughly corresponds to sharing the medium among vehicles with the best transmission rates. Although this is unfair over the short term to vehicles in locations with poor SNR, every vehicle is likely to be allocated a high transmission rate at some point in its trajectory. Hence, over a longer time frame, the scheme is more or less fair. We also considered a distance-based approach, where the vehicle closest to the AP would be assigned use of the medium; however, due to the multitude of environmental factors that affect the RF signal, distance alone does not accurately predict signal quality, as we show experimentally in Section 4.1.

MV-MAX is unfair because of two reasons. First, it shares the medium among all vehicles with the best SNR. Therefore, a vehicle that does not share the medium, because no one else is simultaneously in a good coverage area, will get better throughput than vehicles that share access. Second, in a naive implementation of MV-MAX, vehicles that happen to be shadowed and therefore never get a good SNR may never be allowed to transmit. Nevertheless, we show in Section 4 that under some assumptions, nearly every vehicle's performance improves using MV-MAX compared to 802.11. Therefore, it is incentive-compatible for every vehicle to switch to MV-MAX.

2.1 Implementation

One potential argument against MV-MAX is that it requires a change to the 802.11 MAC scheduling algorithm. To address this, we present three realistic implementation alternatives.

First, the 802.11k extensions to the base 802.11 protocol provide a dedicated signaling channel for network control. Therefore, if APs and clients support 802.11k, which is plausible in future systems, the AP can use this channel to ask each client to report its current SNR, then allocate the data channel to the client with the best SNR.

Second, if such a signaling channel is unavailable, but clients support PCF, then a PCF grant message can be used to allocate the channel to the client with the best SNR.

Third, if neither approach is feasible, then a third alternative is a "MAC ACK Hack", which works as follows. Here, the AP is modified to persistently refrain from sending an 802.11 ACK to clients who are to be denied service. This forces them to back-off, clearing the air for the chosen client(s). Using 802.11g, which has a maximum back-off value of 1024 slots and a slot time of 20 μ sec, clients would be forced to retransmit a packet at least every 20 ms [7]. Assuming symmetric RF channels, this also acts as a probe to provide the client's SNR to the AP. Therefore, it is simple for the AP to choose the client with the best SNR and send

Table 1: Rate coupling functions (k vehicles)

rateCouple	User Rate	System Rate
(r_1, r_2, \ldots, r_k)		
802.11		$k \times User\ Rate$
002.22	$\sum_{k=1}^{k} 1$	
	$\sum_{i=1}^{n} \frac{\overline{r_i}}{r_j}$	
	j=1 -	7
Time Fairness	$\frac{r_u}{k}$	$\frac{1}{k} \sum_{j=1}^{k} r_j$
	$\int r_u \text{ if } r_u =$	J
MV-MAX	$\left \begin{array}{ccc} & & & & \\ & & max(r_1,,r_k) \end{array} \right $	$max(r_1,,r_k)$
	0 otherwise	·

it an ACK, denying access to all other clients.

The "MAC ACK Hack" is not without problems. A client who is forced to retransmit many times may think that the AP is faulty or that it is out of range of the AP, and may disassociate from the AP and therefore may not be able to use the AP at all. Additionally, retransmissions can force TCP timeouts, which can reduce client throughput to nearly zero. Worse, if the client were to mark the network as unreachable, TCP sessions might simply close. Despite these challenges, the solution is at least backward compatible with all legacy clients. We plan to identify and eliminate these and other drawbacks in future work.

3. ANALYSIS

We now mathematically analyze MV-MAX and compare it with unmodified 802.11 and time-fair scheduling.

3.1 Modeling Assumptions

We make the following assumptions to simplify the mathematical model. The impact of these assumptions on the correctness of our results is discussed in Section 5.

- Vehicles are assumed to have a single radio and all vehicular radios use the same channel.
- Vehicles are assumed to move in the same direction.
- Vehicles always have data to either send or receive.
- The rate at which a vehicle can send or receive user-level data over the wireless medium depends on many factors. We assume that, at every point in time, there is a deterministic decision process that maps from the received signal strength at a particular point in the roadway to a transport-layer data transmission rate, taking into account the decision process embodied in IEEE 802.11 auto-rate fallback [19, 12]. We justify this assumption based on real traces in Section 4.
- We assume that all vehicles experience the same signal profile as they pass the access point. This implies that they travel at the same speed, and use identical equipment and software.
- We assume that the RF links are symmetric.
- All communication is between a vehicle and the AP and there is no vehicle-to-vehicle communication.

3.2 Model

Time Slots

We divide time into discrete time intervals (or *slots*) of length *slot length* (e.g. 100 ms). With fixed vehicle speed

and fixed slot length, the number of slots m in coverage range of the ${\bf AP}^1$ is shown in Equation 1.

$$m = \frac{coverage \, range}{slot \, length \times vehicle \, speed} \tag{1}$$

Because all vehicles move with the same speed, a time slot also corresponds to a constant-sized section of the roadway. A vehicle in the u^{th} time slot has an intrinsic user-level data transmission rate r_u . This represents the average rate this vehicle would achieve if it were the only vehicle in range of the AP. The choice of slot length therefore functions as a smoothing parameter, mitigating the effects of sporadic rate fluctuation. This is done without loss of generality as the slot length can be chosen to be the shortest possible data unit (i.e. the smallest 802.11 frame transmitted at the highest rate).

Vehicle Arrival Rate

We assume that the vehicle arrival process is Bernoulli with parameter p. That is, with probability p, a vehicle will enter range of the AP during one time interval (one slot) and with probability (1-p) no vehicle will enter range during that slot.

Effective Transmission Rate

We now derive the expected throughput using a particular medium access control (MAC) protocol in a multi-vehicular scenario. Given k vehicles in range of the AP and their intrinsic rates $r_1, r_2, ..., r_k$, we compute the actual rate (or effective rate) achieved by each vehicle due to the presence of other vehicles (which we call rate coupling). The rate coupling function depends on the MAC scheduling scheme used.

Using standard 802.11, each vehicle has an equal opportunity to transmit at its intrinsic rate r_u , resulting in perpacket fairness. Therefore, each vehicle is allocated the harmonic mean of the intrinsic rates of all k vehicles in range, as shown by the equation in the first row of Table 1. The effective system throughput is therefore k times the individual throughputs.

With the time fairness scheme [20, 22, 15], each vehicle is given an equal *time share* to use the wireless medium, which can be thought of as a partitioning of a single time slot. This leads to the vehicle in the u^{th} slot receiving an individual throughput of r_u/k as shown in Table 1.

With MV-MAX, the vehicle with the largest intrinsic rate gets the entire capacity. If more than one vehicle has the highest rate, the capacity is shared equally among them.

Effective System Throughput

The *expected* throughput of a system with Bernoulli vehicle arrivals to a section of roadway with m slots is:

$$\sum_{k=0}^{m} p^{k} (1-p)^{m-k} \frac{1}{k!} \sum_{j_{1}=1}^{m} \sum_{\substack{j_{2}=1\\j_{2}\neq j_{1}}}^{m} \cdots \sum_{\substack{j_{k}=1\\j_{k}\neq j_{1}\\\dots\\j_{k}\neq j_{k-1}}}^{m} rateCouple(r_{j_{1}},\dots,r_{j_{k}})$$
(2)

The outer sum considers all possible number of vehicles in the system, k, which is at most equal to the number of slots, m, as a slot can only contain up to one vehicle. The

¹We use the term *coverage range* to indicate the total length of the roadway in which a vehicle is in range of the AP.

probability that there are k vehicles in range is given by the binomial distribution B(m,p), where p is the vehicle arrival probability. These vehicles each occupy one of the m slots on the roadway; the k vehicles can be placed in m slots in m!/(m-k)! permutations. Only the simplified form of this combined with B(m,p) is shown. For each permutation, the k vehicles are in positions $j_1...j_k$ where they receive rates $r_{j_1}...r_{j_k}$ which are coupled using the rateCouple function, which depends on the MAC scheduling scheme used (see Table 1). The nested summations are chosen so that no two subscripts are the same because a chosen slot cannot contain more than one vehicle.

We show in Section 4.3 that the results from this equation closely match those from our simulations.

4. SIMULATIONS

To study the expected behavior of a system that implements MV-MAX, we developed a simulator that combines Monte Carlo simulation with real traffic traces. Simulated vehicles arrive at a stretch of road following a Bernoulli distribution and as they pass the AP, each vehicle encounters an identical set of transmission rates derived from experimental data. By tuning the Bernoulli arrival probability p, we can control the mean vehicle arrival rate, which is proportional to the average number of vehicles in range of the AP, $p \times m$, in a system with m slots. We used a slot length of one second for our simulations.

Each simulation run corresponds to at least three hours of real time. For each data point, we ran 100 simulations and 95% confidence intervals (error bars) are shown. On most graphs, however, the error bars are too small to be seen. We use two different data sets for comparison, each using IEEE 802.11b, as discussed next. We also compare the performance of three separate MAC scheduling algorithms: standard 802.11, Time Fairness, and MV-MAX.

4.1 Experimental Data

Our Data

We collected the first data set in December 2004 [8]. A roadside wireless access point (a Linksys WAP11, running 802.11b) was attached to a laptop running an FTP server. A vehicle with a wireless client mounted in the sunroof (a Soekris 4801 with an Aries 5354 MP 802.11b miniPCI card) drove past the access point and was configured to send FTP data as fast as possible once it detected the presence of a connection. Throughput over time (see Figure 2(a)) was then calculated using tcpdump logs of each run. Four vehicle speeds were tested and five runs at each vehicle speed were done in each direction (40 runs in total).

Due to space limitations only data from 80 km/h in one direction is shown in this paper. Using this data, with a slot length of one second, there are m=84 slots and therefore the AP has a range of 1867 m. As an example, using an arrival probability p=0.067 yields an arrival rate of four vehicles per minute, which implies that on average there are 5.6 vehicles in range of the AP with an average of 333 m between each vehicle.

Figure 2(a) shows a significant drop in throughput in the middle of the run. This drop was consistent across all runs (and in both directions), and we attribute this to a small dip in the road which caused line-of-sight to be briefly lost between the vehicle and the access point.

Intel Data

The second data set was collected by Intel Research [6], and is publicly available at the Dartmouth CRAWDAD archive [3]. The test procedure was similar to ours, except that the range is diminished without an external antenna. Only two runs per configuration were done. Throughput over time is shown in Figure 2(b).

We used data collected at 40 km/h in our simulations. There are m = 59 slots and therefore the AP has a range of 656 m. Here, four vehicles arriving per minute (p = 0.067) implies an average of 3.9 vehicles in range of the AP.

Maximum TCP Throughput over 802.11b

In an ideal environment, the maximum achievable throughput of TCP over an IEEE 802.11b link is influenced by many factors, including: TCP header and flow-control overhead, MAC header and ACK overhead, PHY preamble overhead and legislated 802.11 idle times such as SIFS and DIFS. The theoretical maximum throughput of a 1500 byte datagram sent over an 11 Mbit/s 802.11b link was calculated by Jun et al. [11] to be 6 Mbit/s. The maximum TCP throughput over the same link has been experimentally measured to be 5.1 Mbit/s [9], and as much as 5.5 Mbit/s in our lab experiments. We use a generous upper limit of 5.9 Mbit/s to compare our results in the following sections.

4.2 Consistency of Throughput Profile

For MV-MAX to be fair, the signal pattern observed by every vehicle must be relatively consistent. Otherwise, a vehicle that never gets a good signal is in danger of being shut out. Nevertheless, because MV-MAX focuses only on the best SNR, variability in SNR in the fringe areas does not affect its performance.


Figure 2(a) and 2(b) each show four runs of throughput vs. distance overlaid overtop of each other. Although not perfectly identical, the general shape of the signal profile is consistent across all four runs, including the drop in throughput in the middle of the run in Figure 2(a). Note that all runs, corresponding to different vehicles, achieve the highest throughput for a substantial fraction of the time.

4.3 Simulator Validation

Bugs in simulation can lead to unexpected and invalid results. Therefore, we first validated the correctness of our simulator against our mathematical model.

The mathematical model of expected system throughput (Equation 2 presented in Section 3.2) requires us to enumerate all k! permutations of vehicular placements on the road, which grows exponentially with k. Consequently, it can be evaluated only for small values of k. Our simulations, therefore, can be thought of as Monte Carlo simulations that approximate this equation. We now compare simulation results with those predicted by the equation for small values of k. To simplify matters, we created a smaller test data set; see Figure 2(c) for the signal profile, which is representative of real data.

Using this test data with several different arrival probabilities, we both solved our equation and performed simulations. The lines in Figure 3(c) show the analytical results, while the points represent the simulated results. As shown, the simulation results (with 95% confidence intervals too small to see) agree exactly with our analytical results. This validates the correctness of our simulator.



4.4 Effect of 802.11 Performance Anomaly

Standard 802.11 performs poorly in multi-vehicle scenarios due to the performance anomaly [9]. Our simulations confirm this, as can be seen in the 802.11 curves in Figure 3(a)-(c), which show the average system throughput versus vehicle arrival rate for standard 802.11. Here, system throughput is the average TCP throughput achieved from the access point over time. Note that as the vehicle arrival rate increases, the overall performance of 802.11 decreases because there are more vehicles in fringe areas which drags down overall performance.

4.5 MV-MAX Improves Overall and Individual Performance

Figure 3(a)-(c) also shows the increase in average system throughput achieved using MV-MAX as vehicle traffic density increases. MV-MAX improves overall system throughput by up to 3.9x over that of 802.11 using our data (3.6x using Intel data). Moreover, MV-MAX approaches the *maximum* achievable TCP throughput over an IEEE 802.11b link (derived in Section 4.1 to be 5.9 Mbit/s). At high vehicle arrival rates, MV-MAX attains TCP throughput of 4.81 Mbit/s using our data and 5.3 Mbit/s using Intel data.

Because the improvements reported above are averaged across many vehicles, it was not clear whether the performance improvement of some vehicles came at the expense of reduced performance of other vehicles. To analyze this, we simulated 100,000 vehicles passing by a roadside access point and recorded each vehicle's improvement ratio for each scheme compared to standard 802.11. An improvement ratio greater than one indicates that a particular vehicle transferred more data using an alternative scheme than using 802.11. In dense vehicle traffic this was always the case, as shown in Figures 4(a) and 5(a) which show each vehicle's

improvement ratio sorted in descending order. Therefore, in dense vehicle traffic, all vehicles benefit from using either of the two alternative schemes, making them incentive compatible. We next explore different traffic densities and compare the relative gain from MV-MAX vs. Time Fairness.

4.6 MV-MAX vs. Time Fairness

In the multi-vehicle scenario, using Time Fairness results in a significant performance gain compared to standard 802.11. In our simulations, we observe up to 1.8x greater system throughput (Figure 3), using our data (2.3x using Intel data). However, MV-MAX outperforms even Time Fairness. In particular, using our data, MV-MAX results in up to 2.2x more throughput than Time Fairness (Figure 3) (1.6x using Intel data). Overall, MV-MAX outperforms 802.11 by 3.9x using our data and 3.6x using Intel data.

Nevertheless, some vehicles may do better with Time Fairness than with MV-MAX. This is clear from Figures 4(a) and 5(a) where we observe some overlap between the improvement ratios of MV-MAX and Time Fairness. To explore this further, as well as to explore the effect of vehicle arrival rate, we performed the same simulation as above for different vehicle arrival rates, except we report the percentage of vehicles that transferred the same or more data using MV-MAX than with either Time Fairness or 802.11. Figures 4(b) and 5(b) show that, at worst, with an arrival rate of one vehicle/min., 97.2% (our data) and 95.8% (Intel data) of vehicles were able to transfer the same or more data using MV-MAX than they would have if they 802.11 had been used. The results are slightly lower for Time Fairness, where, in the worst case, with a vehicle arrival rate of 24 vehicles/min. and 34 vehicles/min. (our data and Intel data, respectively), only 80.8% (our data) and 81.6% (Intel data) of vehicles transferred the same or more data using MV-MAX than with Time Fairness.




Figure 5: Intel Data (40 km/h)

In general, using MV-MAX results in greater average individual improvement than with Time Fairness. However, due to its unfair nature, some vehicles experience poorer performance than they would have received using another scheme. In the worst case, this affects only approximately 4% of vehicles versus 802.11 and only approximately 19% versus Time Fairness.

4.7 MV-MAX Unfairness

As discussed above, MV-MAX trades off fairness for performance. Here we explore the degree to which MV-MAX is unfair.

To quantify this, Figures 4(c) and 5(c) present a Lorenz curve for vehicular throughput. A Lorenz curve is often used to represent income distribution in society, showing what fraction of households have what fraction of the cumulative income. For instance, a point (x,y) on the curve would indicate that x% of the housholds earn y% of total income in society. Clearly, in an equitable society, the Lorenz curve is a straight diagonal line, and a deviation from linearity is a visual indicator of unfair income distribution. In our case, we show the fraction of vehicles that transferred each cumulative fraction of the total data. Using the same 100,000 simulated vehicles used in Figures 4(a) and 5(a), we constructed a Lorenz curve representing the data, shown in Figures 4(c) and 5(c).

Both 802.11 and Time Fairness achieve near perfect fairness, as shown by the overlapping lines near the diagonal perfect fairness line. MV-MAX however, is somewhat unfair; using our data (Figure 4(c)), we see that the bottom 25% of vehicles transferred only 10.2% of the data, and the bottom 50% transferred only 30.1% of the data (14.5% and 35.5%, respectively, using Intel data as in Figure 5(c)). We next discuss the tradeoff between fairness and optimality.

5. DISCUSSION AND FUTURE WORK

5.1 Effect of Modeling Assumptions

We have made some simplifying assumptions in our analysis and simulations. Here, we discuss the degree to which these simplifying assumptions are likely to hold in practice.

- Single radio per vehicle: Having multiple radios is feasible, but then MV-MAX is applicable to each channel, separately.
- All vehicular radios on the same channel: This is a reasonable assumption because every vehicle has to be on the same channel as the access point anyway.
- All vehicles are assumed to move in the same direction: This assumption will be met as long as the APs on opposite sides of a roadway either are physically staggered or allocated different channels so that there is no interference between them. This is consistent with current channel allocation best practices.
- Vehicles always have data to send or receive: We make this assumption to simplify analysis. Of course, granting access to a vehicle with the best SNR, but with no data to send or receive, is not useful. This can be taken into account in MV-MAX by granting access to the vehicle with the best SNR, among all vehicles that have data to send or receive. Unfortunately, analysis of the resultant system is complex because it depends on the exact workload. We plan to study intermittent vehicular transmissions in future work.
- All vehicles encounter the same signal profile and thus achieve the same transport-layer throughput as they drive past the access point: This assumption is necessary to ensure long-term fairness among vehicles. It

is the most problematic of our assumptions, as the transport-layer throughput of a vehicle depends on many factors, such as the type of equipment it has, its antenna configuration, and MAC-level settings (such as the RTS threshold). Suppose one vehicle had a large antenna and therefore received the best signal while other vehicles with normal or no external antennas consistently had weaker signals. MV-MAX would then allocate all the bandwidth to the vehicle with the biggest antenna, making it unfair. Nevertheless, this assumption allows us to derive a simple expression for the gain from our approach. We intend to study the effect of rate heterogeneity, and to enforce fairness using some form of proportional fairness [13] in future work.

- RF links are symmetric: Measurements [1] indicate that although radio links are not completely symmetric, this assumption is valid to first order. Investigating the effect of asymmetric links is future work.
- All communication is between a vehicle and the access point: This assumption allows us to ignore highly variable effects introduced by the back-haul connection to the Internet. In practice, this assumption amounts to the access point functioning as a cache between vehicular users and the Internet.

5.2 Fairness vs. Optimality

MV-MAX makes resource allocations less fair, yet nearly every vehicle improves its performance. Therefore, it is incentive compatible for each vehicle to participate, even though the resulting resource allocation increases the degree of envy in they system. Designing a resource allocation that is both Pareto optimal and fair is a challenging problem, that, in general, is unsolved. We believe that devising such a scheme for the multi-vehicular environment is a fruitful direction for future work.

5.3 Dealing with Urgent Data

By denying access to vehicles with poor SNR, MV-MAX increases the delay experienced by their packets. This may be inappropriate for "urgent" data. If the fraction of urgent packets in the system is small (and there is a way to prevent a vehicle from marking all packets as urgent), then a simple two-priority scheme suffices to reduce the delays faced by urgent traffic. In this scheme, preference is given to urgent packets, and among urgent packets, to packets from vehicles with the best SNR. We intend to study such a scheme in future work.

5.4 Prototype Implementation

Besides simulating the effect of parameters such as intermittent workloads, dissimilar signal profiles, and asymmetric links on performance, we are in the process of implementing MV-MAX on a realistic testbed and will be running real-world experiments to verify and build on our simulation results. This will also allow us to study the effectiveness of the "MAC ACK Hack".

6. RELATED WORK

Our work assumes that mobile devices in vehicles will use high-bandwidth, short-lived opportunistic connections that arise as the vehicle is in motion. This vision is shared by many academic as well as industry initiatives.

Ott and Kutscher are extensively studying the Drive-thru Internet system [16]. Their work focuses on high level issues such as maintaining persistent connections through periods of disconnection [17] and automating hotspot authentication [18]. Our work is complementary as it provides a lower-layer optimization for scheduling media access when multiple vehicles are in range of a single access point.

Work at Rutgers WINLAB pioneered Infostations [5], which shares the same vision as Drive-Thru Internet. Here, vehicles engage in "infofueling" as they opportunistically drive past roadside access points. However, Infostation research focuses primarily on motivating the idea [5] and integration with cellular networks [2]. Work at the MAC layer includes improving packet delivery probability through the use of an adaptive retransmission scheme [10].

The driving force behind vehicular communication is vehicular safety. The Dedicated Short Range Communications group (DSRC) [4] along with IEEE 802.11p working group are heading this initiative. Their solution, backed by many major automobile manufacturers, uses spectrum in the 5.9GHz band allocated by the FCC for use by the US Intelligent Transport System (ITS) project. Work by the DSRC group addresses many important problems, such as location-based broadcast [21] and vehicle-to-vehicle multihop communication [14]. Our work complements this initiative, providing a scheduling scheme that improves performance in the vehicle-to-infrastructure mode of operation.

Our work is part of the broader field of scheduling access to a shared resource (in our case, the wireless medium). Existing scheduling disciplines are either centralized or decentralized. In a decentralized approach, such as CSMA/CD used in 802.3 Ethernet or CSMA/CA used in 802.11, users randomly contend for access to the medium, and perform some back-off procedure upon collision [7]. On the other hand, centralized media access schemes, such as those used in cellular technology, rely on a scheduler to grant media access to each user. The goal of a scheduler is to maximize some objective function (e.g. high throughput).

The 802.11 MAC protocol uses a decentralized CSMA/CA approach for sharing media access. This contention-based design aims for per-packet fairness giving no consideration to the amount of time a client occupies the medium (ignoring the client's transmission rate and the length of the packet). This results in a *performance anomaly*, identified by Heusse et al. [9], where one slow user with poor signal quality effectively slows down all other users to its rate.

This performance anomaly can be resolved using a time fairness approach [20, 22, 15]. Rather than users contending on a per-packet basis, users contend on a per-time basis. Therefore, each user obtains a fair share of time to access the medium and as a result, slow users don't impact the performance of fast users. While this solution provides an alternative way of achieving fairness, it isn't able to achieve optimal use of the medium in a multi-vehicular environment because it doesn't exploit the consistent signal profile experienced by vehicles as done by MV-MAX.

MV-MAX builds on the ideas of opportunistic scheduling [13]. Opportunistic scheduling, currently used by EvDO systems, exploits the constant small-scale fluctuation of user signal quality, known as multi-user diversity. The idea is that if two clients are communicating with a base station,

and one has poor signal quality during a particular time slot and the other has good signal quality, the channel should be assigned to the user with the good signal quality because on average, due to the constant fluctuation of signal quality, all users will experience good signal quality at some point in time.

MV-MAX differs from traditional opportunistic scheduling in several important ways. First, opportunistic scheduling ideas have never been applied to 802.11. Second, opportunistic scheduling exploits small-scale fluctuations (on the order of milliseconds), whereas MV-MAX considers much larger time scales (minutes). This allows MV-MAX to exploit the repeatability of the signal profile experienced by a vehicle passing a roadside access point; something opportunistic scheduling work has yet to explore.

MV-MAX deviates from traditional opportunistic scheduling because scheduling media access for cellular users is different from scheduling media access for vehicles that occasionally pass roadside access points. In particular, cellular users expect to be connected all the time, thus the scheduler does not have the flexibility to temporarily restrict access to a user, even if the user is currently experiencing poor signal quality and servicing the user would significantly degrade the performance of the entire system. Because users of opportunistic roadside connections inherently expect to experience disconnected episodes between access points, the vehicular user wouldn't even notice if his/her access was being restricted for a few extra seconds as the scheduler waits for the vehicle's signal quality to increase as it nears the access point.

7. CONCLUSIONS

We study the performance of an 802.11-based wireless access point that is shared by more than one passing vehicle. In this situation, due to the 802.11 performance anomaly [9], vehicles on the fringe of coverage with low signal-tonoise ratios, and therefore slow transmission rates, reduce the effective transmission rate of all other vehicles. This degrades both individual and overall performance. We propose a simple and intuitive opportunistic medium access mechanism, called MV-MAX, that grants wireless access only to vehicles with good SNR. Using both analysis and simulation, we show that MV-MAX not only improves overall system throughput, compared to 802.11, by a factor of 3.9, but also outperforms the previously proposed time-fairness scheme [20, 22, 15] by a factor of 2.2. MV-MAX is incentivecompatible for all vehicular users, even though its resource allocation is less fair than that of 802.11. Our results are consistent across different data sets.

8. ACKNOWLEDGEMENTS

We would like to thank Nabeel Ahmed for his insights into this work. We also gratefully acknowledge the National Sciences and Engineering Research Council of Canada for financial support for portions of this project.

9. REFERENCES

- D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level Measurements from an 802.11b Mesh Network. In ACM SIGCOMM, page 121, 2004.
- [2] J. Borras and R. Yates. Infostation Overlays in Cellular Systems. In $I\!E\!E\!E$ WCNC, page 495, 1999.

- [3] Dartmouth CRAWDAD archive. http://crawdad.cs.dartmouth.edu/.
- [4] Dedicated Short Range Communication group. http://grouper.ieee.org/groups/scc32/dsrc/index.html.
- [5] R. Frenkiel, B. Badrinath, J. Borres, and R. Yates. The Infostations Challenge: Balancing Cost and Ubiquity in Delivering Wireless Data. *IEEE Personal Communications*, 7:66–71, Apr. 2000.
- [6] R. Gass, J. Scott, and C. Diot. Measurements of In-Motion 802.11 Networking. In *IEEE Workshop on Mobile Computing System and Applications* (HOTMOBILE 2006), April 2006.
- [7] M. S. Gast. 802.11 Wireless Networks: The Definitive Guide. O'Reilly & Associates, Inc., 2002.
- [8] D. Hadaller, H. Li, and L. G. A. Sung. Drive By Downloads: Studying Characteristics of Opportunistic Connections. In *USENIX NSDI Poster Session*, 2005.
- [9] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda. Performance Anomaly of 802.11b. In *IEEE INFOCOM*, 2003.
- [10] A. Iacono and C. Rose. Bounds on File Delivery Delay in an Infostations System. In *IEEE VTC*, Spring 2000.
- [11] J. Jun, P. Peddabachagari, and M. L. Sichitiu. Theoretical Maximum Throughput of IEEE 802.11 and its Applications. In *IEEE NCA*, 2003.
- [12] M. Lacage, M. H. Manshaei, and T. Turletti. IEEE 802.11 Rate Adaptation: A Practical Approach. In ACM MSWiM, 2004.
- [13] X. Liu, E. Chong, and N. Shroff. A Framework for Opportunistic Scheduling in Wireless Networks. Computer Networks Journal, 41(4):451–474, 2003.
- [14] L. Michael and M. Nakagawa. Multi-hopping Data Considerations for Inter-vehicle Communication over Multiple Lanes. In *IEEE VTC*, 1997.
- [15] A. Munaretto, M. Fonseca, K. Agha, and G. Pujolle. Fair Time Sharing Protocol: A Solution for IEEE 802.11b Hot Spots. Lecture Notes in Computer Science, 3124:1261 – 1266, 2004.
- [16] J. Ott and D. Kutscher. Drive-thru Internet: IEEE 802.11b for Automobile Users. In *IEEE Infocom*, 2004.
- [17] J. Ott and D. Kutscher. A Disconnection-Tolerant Transport for Drive-thru Internet Environments. In IEEE INFOCOM, 2005.
- [18] J. Ott and D. Kutscher. Towards Automated Authentication for Mobile Users in WLAN Hot-Spots. In *IEEE VTC*, Fall 2005.
- [19] S. Pal, S. R. Kundu, K. Basu, and S. K. Das. IEEE 802.11 Rate Control Algorithms: Experimentation and Performance Evaluation in Infrastructure Mode. In PAM '06, 2006.
- [20] G. Tan and J. Guttag. Time-based Fairness Improves Performance in Multi-rate Wireless LANs. In *USENIX* Annual Technical Conference, 2004.
- [21] Q. Xu, R. Sengupta, and D. Jiang. Design and Analysis of Highway Safety Communication Protocol in 5.9GHz Dedicated Short Range Communication Spectrum. In *IEEE VTC*, Spring 2003.
- [22] S.-H. Yoo, J.-H. Choi, J.-H. Hwang, and C. Yoo. Eliminating the Performance Anomaly of 802.11b. Lecture Notes in Computer Science, 3421:1055–1062, 2005.