

Interactive WiFi Connectivity For Moving Vehicles

Aruna Balasubramanian
University of Massachusetts, Amherst

Ratul Mahajan
Microsoft Research

Arun Venkataramani Brian Neil Levine
University of Massachusetts, Amherst

John Zahorjan
University of Washington

ABSTRACT

We ask if the ubiquity of WiFi can be leveraged to provide cheap connectivity from moving vehicles for common applications such as Web browsing and VoIP. Driven by this question, we conduct a study of connection quality available to vehicular WiFi clients based on measurements from testbeds in two different cities. We find that current WiFi handoff methods, in which clients communicate with one basestation at a time, lead to frequent disruptions in connectivity. We also find that clients can overcome many disruptions by communicating with multiple basestations simultaneously. These findings lead us to develop ViFi, a protocol that opportunistically exploits basestation diversity to minimize disruptions and support interactive applications for mobile clients. ViFi uses a decentralized and lightweight probabilistic algorithm for coordination between participating basestations. Our evaluation using a two-month long deployment and trace-driven simulations shows that its link-layer performance comes close to an ideal diversity-based protocol. Using two applications, VoIP and short TCP transfers, we show that the link layer performance improvement translates to better application performance. In our deployment, ViFi doubles the number of successful short TCP transfers and doubles the length of disruption-free VoIP sessions compared to an existing WiFi-style handoff protocol.

Categories and Subject Descriptors

C.2 [Computer Communication Network] Routing protocols

General Terms

Measurement, design, performance

Keywords

Vehicular networks, applications, WiFi, handoff, diversity

1. INTRODUCTION

Our work is driven by two observations: one a growing need and another an opportunity. Many users want cheap and high-quality Internet access from moving vehicles to stay connected while traveling. Cellular networks can provide such connectivity today, but they tend to be expensive. At the same time, there is an increasingly ubiquitous deployment of inexpensive WiFi (802.11) networks, and in many cases, entire cities are being covered [38, 39].

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'08, August 17–22, 2008, Seattle, Washington, USA.
Copyright 2008 ACM 978-1-60558-175-0/08/08 ...\$5.00.

The ubiquity of WiFi provokes an intriguing question: can WiFi deployments support common applications such as Web browsing, instant messaging, and voice over IP (VoIP), from moving vehicles? We are, of course, not the first to suggest allowing WiFi access from moving vehicles. Several recent works study connectivity from vehicles to open-access basestations. They propose techniques to improve connectivity to an individual basestation [8, 19]. Some also propose application-specific techniques or new applications that work well in such environments [4, 14]. The question we pose, however, pushes the envelope beyond this type of special-case usage to supporting common applications.

Our primary contribution is the design of ViFi, a protocol that minimizes disruptions in WiFi connectivity in order to support interactive applications from moving vehicles. ViFi's design is motivated by a rigorous measurement study of two vehicular testbeds in different cities. The goals of our study are to understand the fundamental challenges in supporting interactive applications and to explore opportunities that can be leveraged in this environment.

We find that with current WiFi handoff methods clients experience frequent disruptions in connectivity even when they may be close to WiFi basestations. Handoffs in WiFi today are *hard*, i.e., at any given time, clients communicate with only one basestation that is expected to offer the best connectivity. Hard handoffs are limited by gray periods in which connectivity drops sharply and unpredictably, the difficulty of estimating the continuously changing channel quality to near-by basestations, and the short-term burstiness of losses. Interestingly, we find that even though the impact on the performance of delay or disruption-tolerant applications is small, the user-perceived quality for interactive applications that need consistent connectivity deteriorates significantly.

We also find that *macrodiversity*, i.e., using multiple basestations simultaneously, can help reduce disruptions for vehicular clients. Its use has been successful in cellular networks [37]. In our context, it overcomes the limitations of hard handoff because of independence of packet losses across basestations and even outperforms an ideal hard handoff strategy with future knowledge of loss rates.

ViFi exploits macrodiversity and opportunistic receptions by near-by basestations to minimize disruptions for mobile clients. The challenge in designing ViFi is in coordinating among basestations that opportunistically receive packets. This coordination must be nimble enough to allow per-packet processing and must use the communication channel efficiently. ViFi addresses this challenge using a simple yet effective probabilistic algorithm. Basestations that opportunistically overhear a packet but not its acknowledgment, probabilistically relay the packet to the intended next hop, such that wasted transmissions are minimized. Unlike opportunistic routing protocols for wireless mesh networks [5, 9], the per-packet overhead of ViFi is low enough to not require batching. Batching tends to delay packets and is thus unsuitable for many interactive applications. And unlike diversity-based handoff protocols for enterprise WLANs [23, 25, 26], ViFi places little additional demand




Figure 1: The layout of BSes in VanLAN.

on the inter-basestation communication plane that is bandwidth limited in our setting.

We implemented and successfully deployed ViFi on our testbed in Redmond, Washington, for over two months. We evaluate ViFi using our deployed prototype and simulations driven by traces from our testbed in Amherst, Massachusetts. Our evaluation shows that the link-layer performance of ViFi is comparable to an ideal diversity-based handoff protocol and outperforms an ideal hard handoff protocol. At the same time, it uses the wireless medium as efficiently as a single-basestation handoff protocol.

We study the performance of ViFi for two commonly used interactive applications: VoIP and short TCP transfers that are typical in Web browsing. Based on our deployment results, we find that ViFi reduces disruptions for these applications and improves their user-perceived quality. Trace-driven simulations based on the second testbed corroborate our findings.

2. EXPERIMENTAL PLATFORMS

Our designs and evaluations are based on experiments performed on two vehicular mobile network testbeds: VanLAN and DieselNet. VanLAN is a vehicular testbed that we deployed in the Microsoft campus in Redmond, Washington, and we use it for measurements and to conduct live experiments. DieselNet [7] is also a vehicular testbed that we deployed in Amherst, Massachusetts. Unlike VanLAN, however, in DieselNet we cannot modify the BSes, and can only control the equipment on the vehicles. We profile the environment of DieselNet and use the collected traces for trace-driven studies of the environment and to evaluate performance. Our intent in using two distinct environments in two different cities is to gain confidence that our results apply in different settings.

2.1 VanLAN

VanLAN consists of eleven basestations (BSes) and two vehicles. The BSes are deployed across five buildings in the Microsoft campus, as illustrated in Figure 1. The box in the figure, covering a 828×559 m² area, bounds the region in which at least one packet is received by vehicles from any BS. Not all pairs of BSes are within wireless range of one another. The vehicles provide a shuttle service around the town, moving within a speed limit of about 40 Km/h. Each vehicle visits the region of the BSes about ten times a day.

The BSes and vehicles have small desktops with Atheros 5213 chipset radios. The antennae are omnidirectional and are mounted on the roofs of the respective buildings and vehicles. The vehicles are equipped with a GPS unit that outputs location information once every second. All nodes were set to the same 802.11 channel for the experiments in this paper.

The radios operate in the *ad hoc* (IBSS) mode using a locally modified Windows driver. The modifications let us use a fixed BSSID, which prevents network partitions and reconciliation delays. They also let us log received frames along with PHY layer information (e.g., RSSI) without place the radio in *monitor* mode.

2.2 DieselNet

The DieselNet testbed is in Amherst, a college town in Massachusetts. DieselNet consists of vehicles equipped with a computer, an 802.11 radio, and a GPS unit. To enable trace-driven studies, we set one vehicle to log all beacons heard from nearby BSes. Later sections describe how this information is used for analysis.

In this paper, we use traces from Channels 1 and 6. Each channel was profiled for 3 days in December 2007, during which the vehicle logged more than 100,000 beacons. The profiling channel was fixed so that beacons are not lost while scanning. We limit our analysis to BSes in the core of the town and to BSes that are visible on all three days. There are 10 such BSes on Channel 1 and 14 on Channel 6. About half of the BSes on each channel belong to the town's mesh network and the rest belong to nearby shops. Our traces are available at traces.cs.umass.edu.

3. A CASE FOR DIVERSITY

This section presents a case for using BS diversity to support vehicular WiFi clients, based on measurements from our testbeds. Client performance in this environment depends on the handoff strategy. We show that handoff strategies that use one BS at a time, as is common today, cause disruptions in connectivity that make supporting interactive applications challenging. We also show that using multiple BSes can mask these disruptions and improve application performance. These results set up our proposal for a new diversity-based handoff protocol, ViFi, which we present in the next section. To our knowledge, ours is the first study that evaluates various handoff strategies in vehicular environments.

3.1 Methodology

There are many possible handoff strategies; for a comprehensive evaluation, we study six different policies. Four of these are practical and are based on existing literature; the other two are idealized methods and lend insight into the inherent limitations of the practical policies. Our goal is to first understand the fundamental differences among the policies, so we ignore the time taken to switch between BSes and the time to scan for BSes. Research shows that careful design can tackle these issues effectively [31, 6, 20, 2].

We analyze the performance of these handoff policies using a trace-driven evaluation on VanLAN. During the evaluation, each BS and vehicle broadcasts a 500-byte packet at 1 Mbps every 100 ms. We verified that self-interference of this traffic is minimal by comparing its packet reception ratio with the case where only one node sends at a time. Even though unicast packets, followed by acknowledgments, are more common today, broadcast packets suffice to probe the underlying connectivity and let us measure connectivity from a sender to all receivers simultaneously [1]. Nodes log all correctly decoded packets and beacons. Our results are based on analysis of traces collected over a two-week period.

We evaluate a handoff policy using these traces as follows. The policy determines which BS a client associates with at a given time.

The client can communicate with only the associated BS when using a hard handoff policy. We assume that clients have a workload that mirrors our trace traffic; i.e., they wish to send and receive packets every 100 ms. The traces of broadcast packets and the current association determine which packets are successfully received.

We use two measures to provide insight into the performance of different kinds of applications: (i) aggregate performance; (ii) periods of uninterrupted connectivity. An aggregate performance measure considers the total number of packets delivered and the total time or distance over which the vehicle is connected. These are relevant to delay or disruption-tolerant applications that care most about total throughput, e.g., synchronizing mail folders in the background. The period of uninterrupted connectivity measures contiguous time intervals when the performance of an application is above a threshold, for some definition of performance and threshold. Measuring periods of uninterrupted connectivity will, for example, tell us the length of time a VoIP caller can talk before the call quality drops. Applications such as instant messaging lie between these extremes; interpolating our results can provide insight into their performance.

The six handoff policies that we study are the following.

1. **RSSI**, where the client associates to BSes with higher signal strength, measured as the exponential average of the RSSIs of received beacons. This policy is similar to what many clients, including the NICs in our testbed, use currently in infrastructure WiFi networks.

2. **BRR**, where the client associates to the BS with the highest exponentially averaged beacon reception ratio. This policy is inspired by wireless routing protocols that are based on the reception ratio of probes [12].

We use an exponential averaging factor of half for both methods above and find the results robust to the exact choice.

3. **Sticky**, where the client does not disassociate from the current BS until connectivity is absent for a pre-defined time period, set to three seconds in our evaluation. Once disassociated, the client picks the BS with the highest signal strength. This policy was used in the CarTel study [8].

4. **History**, where the client associates to the BS that has historically provided the best average performance at that location. Performance is measured as the sum of reception ratios in the two directions, and the average is computed across traversals of the location in the previous day. MobiSteer shows the value of history in vehicular environments [28].

5. **BestBS**, where at the beginning of each one-second period, the client associates to the BS that provides the best performance in the *future* one second. Performance is measured as the sum of reception ratios in the two directions. This method is not practical because clients cannot reliably predict future performance.

In cellular terminology, all of the policies above use *hard handoff* because the client associates with only one BS at a time. Using future knowledge, *BestBS* represents an upper bound on the performance of hard handoff methods. Comparing it with macrodiversity methods (below) reveals the inherent limitations of hard handoff.

6. **AllBSes**, where the client opportunistically uses all BSes in the vicinity. A transmission by the client is considered successful if at least one BS receives the packet. In the downstream direction, if the client hears a packet from at least one BS in an 100-ms interval, the packet is considered as delivered.

AllBSes is an ideal method that represents an upper bound on the performance of any handoff protocol. It exploits path diversity between the client and the set of nearby BSes. Because of differences in CDMA and CSMA, it is not identical to, but is inspired by, macrodiversity methods in cellular networks [37]. Path diversity is




Figure 2: Average number of packets delivered per day in VanLAN by various methods. Error bars represent 95% confidence intervals.

known to improve performance in WiFi mesh and indoor infrastructure networks as well [5, 25]. We study if such benefits materialize in vehicular WiFi settings.

3.2 Aggregate Performance Results

Figure 2 shows the packets delivered by the six handoff policies. To study the impact of BS-density, the independent variable in the graph is the number of BSes in the system. There are eleven BSes in VanLAN, and each point in the figure represents the average of ten trials using randomly selected subset of BSes of a given size.

The graph shows that more packets are delivered as the density of BSes increases but the relative performance of various methods is similar.¹ *AllBSes* performs best, followed by *BestBS*, and then by *History*, *RSSI*, and *BRR*; the performance of *Sticky* is the worst. Ignoring *Sticky*, all methods are within 25% of *AllBSes*. This result suggests that for non-interactive applications, the choice of the exact method is not critical — however, results below demonstrate that interactive applications manifest great differences among the policies. Because of space limitations, we omit similar results for other aggregate performance metrics, such as the total time or distance for which the methods provide some minimal connectivity.

History, *RSSI* and *BRR* perform similarly for all measures that we study. The competitive performance of *History* confirms recent results [28] about the feasibility of using past experience to predict future performance. For visual clarity, we present results for only *BRR* as representative of all three in the remainder of this paper.

3.3 Uninterrupted Connectivity Results

To compare the ability of different handoff methods in providing uninterrupted connectivity, we start with a qualitative example. Figure 3 shows the behavior of *BRR*, *BestBS*, and *AllBSes* during one example trip of the vehicle. In this example, we define adequate connectivity to mean at least 50% of the packets are received in a one-second interval. Consistent with our aggregate performance measurement, each method provides adequate connectivity for roughly the same total path length. However, *BRR* contains several regions of inadequate connectivity. *BestBS* has fewer interruptions because it uses the optimal BS. *AllBSes* performs best as it uses multiple BSes to further reduce the number of interruptions.

Frequent interruptions in *BRR* can be explained through a detailed analysis of the connectivity between a vehicular WiFi client and a BS. Contrary to earlier studies of controlled environments [29, 17], we find that in realistic environments this connectivity is of-

¹We find it surprising that the curves for number of packet delivered did not flatten. We expected the density of VanLAN to be sufficiently high for performance to be the same as a sparser deployment. This points to the challenge of deploying WiFi in outdoor areas such that performance is maximized and not just coverage.

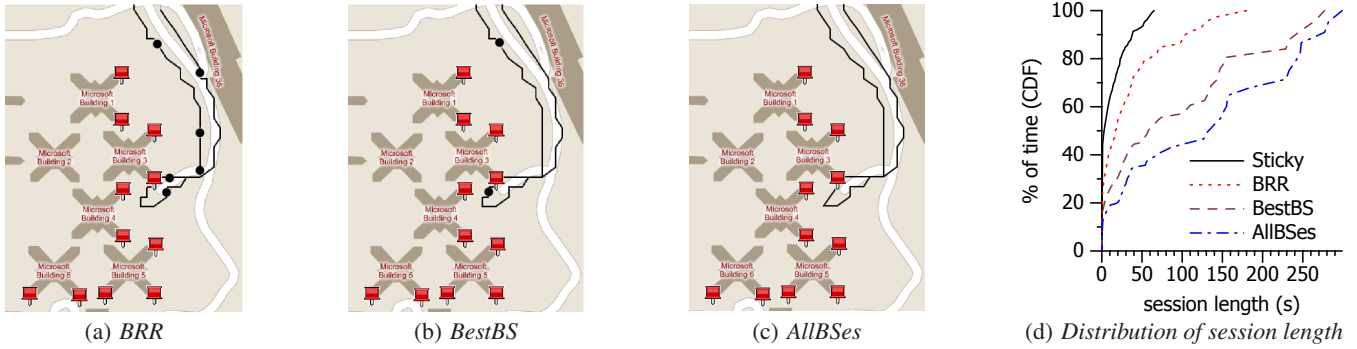


Figure 3: (a)-(c): The behavior of three handoff methods for an example path segment in VanLAN. Black lines represent regions of adequate connectivity, i.e., more than 50% reception ratio in a one-second interval. Dark circles represent interruptions in connectivity. (d): The CDF of the time the client spends in a session of a given length.

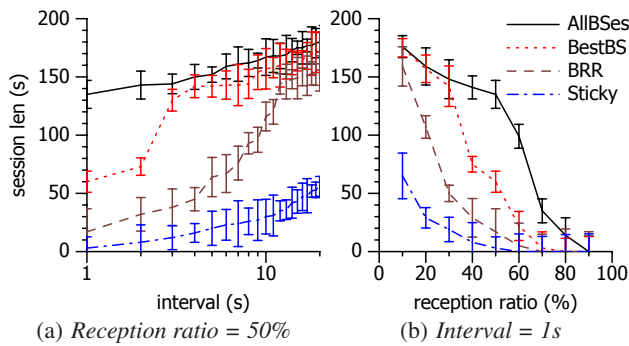


Figure 4: The median session length in VanLAN as a function of the time interval (a) and the minimum reception ratio (b) used to define adequate connectivity. The left graph has logarithmic x -axis. Error bars represent 95% confidence intervals.

ten marred by *gray* periods where connection quality drops sharply. Gray periods are unpredictable and occur even close to BSes. With BRR, the clients often find themselves experiencing a gray period with respect to the associated BS, which causes frequent disruptions in connectivity. But because they tend to be short-lived, gray periods do not severely impact aggregate performance. We have analyzed gray periods in our testbeds in more detail [3, 24] but omit that analysis from this paper.

Figure 3(d) quantitatively compares the handoff policies with respect to the cumulative time clients spend in an uninterrupted session of a given length. We see that the median session length of *AllBSes* is more than twice that of *BestBS* and more than seven times that of the more practical *BRR*. This suggests that a practical, multi-BS handoff policy can achieve significant gains over hard handoff.

To investigate how applications with different requirements can be supported, we now explore other definitions of adequate connectivity. Figure 4(a) varies the averaging interval while keeping the minimum reception ratio requirement fixed at 50%; Figure 4(b) varies the minimum reception ratio while keeping the averaging interval fixed at one second. A longer averaging interval represents less stringent requirements because the session is said to be interrupted only if there is no activity for a longer period. Similarly, a shorter reception ratio represents a weaker requirement. The results suggest that when the requirements are less stringent all methods other than *Sticky* perform similarly. But as the requirements




Figure 5: The CDF of the number of BSes from which a vehicle hears beacons in a 1-second period.

become more demanding the relative advantage of using multiple BSes increases. The right end of Figure 4(b) does not represent convergence but a degenerate point where the requirements are so strict that all methods have short sessions.

3.4 Why is using multiple BSes effective?

We now explain why *AllBSes* is significantly more effective than using only one BS even when that BS is judiciously chosen. We show that its effectiveness stems from two factors: (i) the vehicle is often in range of multiple BSes; (ii) packet losses are bursty and roughly independent across senders and receivers. In the next section, we leverage these findings in the design of ViFi.

3.4.1 Extent of diversity

To exploit BS diversity, a vehicle must be in range of multiple BSes on the same channel. As shown in Figure 5, this is true not only in VanLAN, which we have deployed, but also in DieselNet. The graphs plot the CDF of the number of BSes from which the vehicles hear beacons in one-second intervals. Our results are consistent with measurements in other cities [8]. While future deployments may be engineered for diversity (§6), we find sufficient diversity even in existing deployments.

In separate experiments (not shown here), to understand the extent of diversity actually needed, we find that using as few as two BSes brings most of the gain and there is no additional benefit to using more than three. A similar observation holds for cellular networks [37].




Figure 6: (a) Probability of losing packet $i+k$ from a BS to vehicle in VanLAN given that packet i was lost. (b) Unconditional and conditional packet reception probability from two BSes to the vehicle.

3.4.2 The nature of losses

Diversity can effectively tackle losses in the vehicular environment. In the upstream direction, diversity is effective because losses are roughly independent across BSes and a packet sent by the vehicle is received by at least one BS with a high probability. In other words, the fact that a packet from the vehicle is lost at, for instance, the closest BS, has little bearing on whether it is lost at another BS. This independence of losses at receivers has been shown previously for outdoor WiFi meshes [5]. We find that it holds in our setting as well but omit detailed results.

Diversity is effective in the downstream direction because it can tackle bursty losses better than single-BS systems [25]. Figure 6(a) shows evidence that losses are bursty in the vehicular setting. The figure plots the probability of losing the packet ($i+k$) from a BS to vehicle in VanLAN given that packet i was lost. In this experiment, a single BS sends packets every 10 ms; we pick a different sending BS for each trip by the vehicle. The probability of losing a packet immediately after a loss is much higher than the overall loss probability. Thus, even when a vehicle is associated to a BS with a low average loss rate, it can lose many packets in a small time period, hurting interactive applications.

Diversity helps overcome burst losses because when the vehicle is in a burst-loss phase with one BS a second BS can deliver packets to it. That is, most burst losses are path dependent (e.g., due to multipath fading) rather than receiver dependent. Figure 6(b) shows evidence that this holds for the vehicular environment and quantifies the effect for one pair of chosen BSes in VanLAN. Each BS sends a packet every 20 ms. $P(A)$ and $P(B)$ are the unconditional downstream packet reception probabilities from BSes A and B . $P(A_{i+1}|\neg A_i)$ is the conditional reception probability of receiving ($i+1$)-th packet from A given that the i -th packet from A was lost. Other probabilities can be similarly interpreted. We see that after a loss from a BS, the reception probability of the next packet from it is very low. But the second BS's probability of delivering the next packet is only slightly lower than its unconditional loss probability.

4. ViFi DESIGN AND IMPLEMENTATION

In this section, we present ViFi, a protocol designed to minimize disruptions in connectivity between moving vehicles and a network of BSes. We focus on improving the underlying link-layer connectivity. In some environments, providing continuous connectivity to applications may also require higher-layer techniques, for instance, to handle IP address changes; for these, we rely on existing solutions [35, 30].

4.1 Target environment

The design of ViFi assumes that its operating environment has the following characteristics.

- *Diversity*: A packet sent by a moving vehicle can often be heard by multiple BSes, and multiple BSes can often deliver packets to a moving vehicle. This assumption is fundamental for leveraging BS diversity. It is not necessary that the reception rate between the vehicle and each BS be high.
- *Bandwidth-limited inter-BS communication*: The BSes can communicate with each other. However, in WiFi deployments today, inter-BS communication tends to be based on relatively thin broadband links or a multi-hop wireless mesh. Accordingly, we assume that inter-BS communication is bandwidth constrained.

We also assume that some of the nearby BSes can overhear each other over the vehicle-BS channel. This assumption is not strictly necessary. For example, the functionality can be fulfilled using the inter-BS backplane.

4.2 Goals and approach

Motivated by the effectiveness of *AllBSes*, which itself is impractical, we seek to develop a protocol that leverages BS diversity to reduce disruptions. The key challenge is in coordinating among the BSes such that the coordination scheme: (i) imposes minimal additional load on the inter-BS and vehicle-BS communication media; (ii) does not increase per packet latency, as that hurts interactive applications; (iii) can handle rapidly changing sets of BSes.

Other works that leverage diversity in WiFi networks either assume a high-speed inter-BS backplane [26, 25] or batch packets to amortize overhead [5, 9]. In our setting, however, high-capacity backplane is often not available. We also cannot use batching because that increases latency for packets.

Our approach is to leverage opportunistic receptions by nearby BSes, followed by probabilistic relaying. Opportunistic receptions provide a low-overhead but unreliable means for disseminating information. With probabilistic relaying, each BS relays based on an independently computed relaying probability, which avoids the need for explicit coordination messages between BSes. The resulting protocol is lightweight and decentralized.

Another aspect of ViFi that differs from traditional WiFi hand-off is *salvaging*, in which BSes attempt to save packets that are stranded at old BSes when the vehicle moves away.

4.3 Protocol overview

In ViFi, the vehicle designates one of the nearby BSes as the *anchor*. The anchor can be selected using any of the association methods that clients use today to pick a BS. Our implementation uses *BRR*. The anchor is responsible for the vehicle's connection to the Internet—packets from the vehicle are forwarded through the anchor and packets from the Internet destined for the vehicle first arrive at the anchor. A client acquires its IP address from the anchor BS, if needed.

The vehicle designates other nearby BSes as *auxiliary*. We currently pick all BSes that the vehicle hears as auxiliaries. In certain environments, such as those that are highly dense, the list of auxiliaries may need to be more carefully selected.

The vehicle embeds the identity of the current anchor and auxiliary BSes in the beacons that it broadcasts periodically. Beacons enable all nearby BSes to learn the current anchor and the set of auxiliary BSes. Thus, changes to the identity of the anchor or the set of auxiliary BSes are communicated to the BSes by the vehicle

at the beaconing frequency. The vehicle also embeds the identity of the previous anchor for salvaging (§4.5).

The operation of ViFi is symmetric in both directions and is described below in terms of the source, src , and destination, dst , of the transfer. In the upstream direction, the vehicle is the source and the anchor is the destination. The roles are reversed in the downstream direction.

1. src transmits the packet P .
2. If dst receives P , it broadcasts an ACK.
3. If an auxiliary overhears P , but within a small window has not heard an ACK, it probabilistically relays P .
4. If dst receives relayed P and has not already sent an ACK, it broadcasts an ACK.
5. If src does not receive an ACK within a retransmission interval, it retransmits P .

Upstream packets are relayed on the inter-BS backplane and downstream packets on the vehicle-BS channel. A packet is considered for relaying only once, and packets overheard from other auxiliary BSes are not relayed. In Step 3, the overheard ACK that suppresses relaying by an auxiliary BS could be in response to either the source's transmission or a relayed transmission by another auxiliary BS.

It is instructive to understand why relaying by an auxiliary BS is better than a retransmission by the source itself. The first reason is that losses are bursty—if the original was lost, there is a high chance that an immediate retransmission will be lost as well. After a loss, other nodes are better positioned to deliver the packet to the destination (§3.4.2). An additional reason in the upstream direction is that relaying uses the inter-BS communication plane, which in many cases will be more reliable than the vehicle-BS channel.

4.4 Computing relaying probability

The key challenge in computing relaying probability for auxiliary BSes is to balance the trade-off between too few and too many relayed transmissions. With the former, the performance will degrade to that of no diversity; the latter will lead to excessive load on the vehicle-BS and inter-BS communication mediums.

The relay probability computation in ViFi is based on the following guidelines.

- G1: Account for relaying decisions made by other potentially relaying auxiliaries.
- G2: Prefer auxiliaries with better connectivity to the destination.
- G3: Limit the expected number of relayed transmissions.

The first two guidelines are easily motivated, but the third one is not immediately obvious. Should the number of relayed transmissions be low or be such that at least one of them reaches the destination? We use the former in ViFi, but we also considered a formulation based on the latter. We outline this formulation in §5.5.1, and show that it leads to too many relayed transmissions. Similarly, we study formulations that do not adhere to the other two guidelines and show that they do not perform well either.

Let B_1, \dots, B_K be the current set of auxiliary BSes. Let node s be the source of a packet and node d be its destination, where a *node* is a vehicle or anchor BS depending on the packet's direction. Let p_{ab} represents the probability that b correctly receives a transmission from a , for $a, b \in \{s, d, B_1, \dots, B_K\}$. ViFi estimates and disseminates the p_{ab} using periodic beacons (§4.6).

When some auxiliary B_x hears a packet but not an acknowledgment, it must use a locally computed probability to decide whether to relay. The overall strategy is to compute relaying probabilities so that the expected number of packets relayed across all auxiliary

BSes is equal to 1. Within this constraint, auxiliary BSes that are better connected to the destination are preferred.

We reflect the constraint on the expected number of packets relayed using

$$\sum_{i=1}^K c_i r_i = 1 \quad (1)$$

Here c_i is the probability that auxiliary B_i is *contending* on this packet, that is, that B_i has heard the packet but not an acknowledgment, and r_i is B_i 's relay probability. Strictly speaking, however, r_i is the number of times B_i should relay the packet. Except in pathological cases, r_i evaluates to less than one. We do not allow an auxiliary BS to relay a packet more than once.

We compute the c_i using an approach described below. We then pick r_i satisfying Eq. 1 in a way that favors auxiliaries that are better connected to the destination node d . Specifically, we choose r_i such that

$$\frac{r_i}{r_j} = \frac{p_{B_i d}}{p_{B_j d}} \quad (2)$$

implying that $r_i = r \cdot p_{B_i d}$ for some r . Each contending auxiliary B_x solves Eq. 1 uniquely for r , and then relays the packet with probability $\min(r \cdot p_{B_x d}, 1)$.

A contending relay B_x computes c_i for each B_i , including itself, as the unconditional probability:

$$c_i = p_{s B_i} (1 - p_{s d} p_{d B_i}) \quad (3)$$

Here the first term, $p_{s B_i}$, is the probability that B_i receives the original packet, the second is the probability that B_i does not hear an acknowledgment. We have assumed that the two are independent.

The probability computation method described above is but one of the possibly many that adhere to the guidelines above. We use it because it is simple and works well in our experiments. It strikes a balance between false positives, i.e., relaying packets that are already at the destination, and false negatives, i.e., no relaying for packets that are lost at the destination.

In practice, false positives are also reduced because relay attempts of auxiliary BSes are not synchronized. Each auxiliary BS has a timer that fires periodically. When that happens, the auxiliary BS uses the equations above to decide whether it needs to relay any unacknowledged packet. In some cases, an acknowledgment arrives at the auxiliary BS even before its timer fires. Such an event prevents the BS from relaying unnecessarily, even if the equations indicate that the packet should be relayed.

The combination of ViFi's relaying probability computation method, asynchronous relaying timers, and suppression based on overheard acknowledgments means that we do not need to explicitly, temporally order the relaying BSes based on their proximity to the destination; such ordering has been used in the past in both wired and wireless settings [5, 15].

4.5 Salvaging

Sometimes a vehicle moves out of range before the anchor BS can deliver packets from the Internet. Application performance, especially that of TCP, can suffer if such groups of back-to-back packets are lost frequently.

To avoid this problem in ViFi, newly designated anchors *salvage* packets by contacting the previous anchor over the backplane. The new anchor learns the identity of the previous anchor from the beacons. Upon contact, the old anchor transfers any unacknowledged packets that were received from the Internet within a certain time threshold. We set the threshold to one second in our experiments, based on the minimum TCP retransmission timeout. The

new anchor treats these packets as if they arrived directly from the Internet. Our salvaging mechanism is inspired by DTN routing and DSR [21], but it is based on pulling data rather than pushing.

4.6 Estimating packet reception probabilities using beacons

As WiFi BSes do today, ViFi nodes send periodic beacons. The beacons are used to disseminate information about the packet reception probabilities needed by auxiliary BSes, which include those between the other auxiliary BSes and the anchor and between the other auxiliary BSes and the vehicle.

A ViFi node estimates the reception probability from another node to itself using the number of beacons received in a given time interval divided by the number that must have been sent. Incoming reception probabilities are maintained as exponential averages ($\alpha=0.5$) over per-second beacon reception ratio. In their beacons, nodes embed the current incoming reception probability from all nodes that they heard from in the last interval. They also embed the packet reception probability from them to other nodes, which they learn from the beacons of those other nodes. This embedded information suffices for an auxiliary BS to learn all the packet reception probabilities that it needs.

4.7 Retransmission timers

In the current 802.11 standard, acknowledgments are sent immediately after packet transmission, so the source knows when to retransmit an unacknowledged packet. But acknowledgments in ViFi may be delayed if they are generated in response to a relayed packet. The delay depends on the time for relayed packets to reach the destination, and thus retransmission timers must be set based on current network conditions.

The ViFi source sets the retransmit timer adaptively based on the observed delays in receiving acknowledgments. The source keeps track of the delays in receiving acknowledgments for its transmissions. Each packet carries a unique identifier so that acknowledgments are not confused with an earlier transmission. The source then picks as the minimum retransmission time the 99th percentile of measured delays. Picking this high percentile means that sources err towards waiting longer when conditions change rather than retransmitting spuriously.

Transmission opportunities can arise for the source before the retransmission time for the earliest packet in the queue elapses. In such an event, instead of leaving the medium idle, the source sends the earliest queued packet that is ready for transmission. This can cause some amount of reordering when a later packet reaches the destination first. In our experiments, we find that the amount of reordering is small and does not hurt TCP performance. Hence, our current implementation does not attempt to order packets. If need be, it is straightforward to order packets using a sequencing buffer at anchor BSes and vehicles.

4.8 System Implementation

We have implemented ViFi on the Windows operating system. Almost all of our implementation sits in user space. A special in-kernel network driver receives outgoing packets from the OS and hands it to our process. This process then sends it back down to the wireless interface after adding appropriate headers. Upon receiving incoming packets from the wireless interface, this process strips the headers and hands the packet to the special driver which then passes it on to the OS. We embed our own sequence numbers as identifiers in transmitted packets, though it should be possible to use 802.11 sequence numbers with a tighter integration with the device driver.

Our current implementation uses broadcast transmissions at the MAC layer because this lets us disable the automatic retransmission behavior of the NIC. Instead, a ViFi node retransmits unacknowledged packets as described in §4.3. The ViFi node also send acknowledgments for received packets, since broadcast transmissions in 802.11 are not acknowledged. However, broadcast transmissions disable exponential backoff in response to losses which is intended to reduce collisions. Given that many losses in the vehicular environment will not be due to collisions but due to poor radio links, it is unclear if the standard 802.11 exponential backoff behavior is appropriate. To reduce collisions, our implementation relies on carrier sense. The implementation also ensures that there is no more than one packet pending at the interface, to prevent a node from sending multiple back-to-back broadcast packets.

As an optimization, ViFi packets carry a 1-byte bitmap that signals which of the last eight packets before the current packet were not received by the sender. This helps save some spurious retransmissions of data packets that are otherwise made due to loss of acknowledgment packets.

Our implementation of ViFi has been deployed on VanLAN, where it ran successfully for more than two months.

5. EVALUATION

In this section, we evaluate the performance of ViFi according to several criteria. We show the following:

- The link-layer performance of ViFi is close to ideal (§5.2).
- ViFi improves application performance two-fold compared to current handoff methods (§5.3).
- It does that without placing little additional load on the vehicle-BS wireless medium (§5.4).
- Its coordination mechanism has low false positive and false negative rates (§5.5).

In addition to the experiments presented in this paper, we have conducted a broader study of the performance of ViFi across a range of environmental factors. These factors include the density of BSes and the speed of the vehicle, which we could not control for either of our testbeds. Our results, which are presented in a separate technical report [3], show that ViFi performs well across these factors.

5.1 Methodology

Our evaluations use the deployment of ViFi on VanLAN and a trace-driven simulation based on measurements from DieselNet. The first approach provides results in the context of complete real-world complexities. The second approach allows us to verify that the results are not due to characteristics specific to our deployment on VanLAN. Below, results that are marked with VanLAN, are deployment-based and those marked with DieselNet, are based on trace-driven simulations.

The trace-driven simulations are based on beacons logged by the buses in DieselNet. The beacon loss ratio from a BS to the vehicle in each one-second interval is used as the packet loss rate from that BS to the vehicle and from the vehicle to the BS. This assumption ignores any asymmetry or finer-timescale behavior of packet loss. For inter-BS loss rates, we assume that BS pairs that are never simultaneously within the range of a bus cannot reach one another. For other pairs, we assign loss ratios between 0 and 1 uniformly at random. Our results are based on multiple trials and random seeds.

We use a QualNet-based implementation of ViFi to analyze performance. The loss rates are instantiated in the QualNet simulator by mapping them to the corresponding path loss values. This method allows us to program loss rates found in a real vehicular environment and therefore includes losses due to mobility and multipath fading, while still losing packets to events such as collisions.

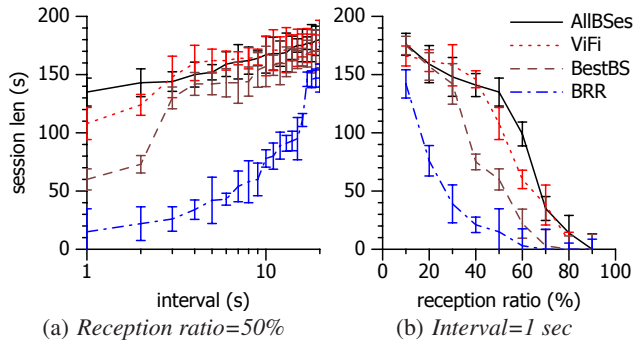


Figure 7: The median session length in VanLAN as a function of the reception ratio threshold and time interval used to define adequate connectivity. The curves for *AllBSes* and *BestBS* are identical to those in Figure 4.

We validate our trace-driven simulation method by collecting the same measurements from VanLAN and comparing its results to the deployment, i.e., we configure the loss rate for each one-second interval to be the beacon loss ratio between the vehicle and the BS in that one second. Because we have inter-BS beacon loss ratios in VanLAN, unlike DieselNet, we configure the inter-BS loss rates also as the inter-BS beacon loss ratio at each one-second interval. We find that the simulation results match the deployment results. For instance, the VoIP session lengths in the simulations are within five seconds of the session lengths observed for the deployed prototype. We omit details of this validation from this paper but include them in a technical report [3].

We compare the performance of ViFi against *BRR*, the practical, hard handoff protocol that we studied previously. To ensure a fair comparison, we implement *BRR* within the same framework as ViFi but with the auxiliary BS functionality switched off. Like ViFi, *BRR* uses broadcast transmissions without exponential backoff restrictions and uses bitmap acknowledgments. We omit experiments that show that *BRR* performs worse with unicast transmissions. The poor performance is because of backoffs in response to losses. In VoIP experiments, for instance, the length of disruption-free calls were 25% shorter.

Our experiments are based on a fixed 802.11b transmission rate of 1 Mbps to maximize range. Rate adaptation in vehicular networks is an open problem as current algorithms assume an environment that is less dynamic [19, 14].

Unless otherwise specified, results for VanLAN are based on at least three days of data for each protocol and workload configuration. Details for the data we use for the DieselNet simulations are provided in §2.2. All errors bars in the graphs below represent 95% confidence intervals.

5.2 Link-layer performance

We start by evaluating the basic link-layer connectivity provided by ViFi. This analysis is based on the VanLAN deployment and uses a methodology similar to §3.

Figure 7 quantifies the performance of ViFi in comparison to the *BRR*, *BestBS*, and *AllBSes* handoff policies. In this experiment, the van and a remote computer attached to the wired network send a 500-byte packet to each other every 100 ms. Since we focus on basic link-layer quality provided by each protocol, link-layer retransmissions are disabled. The figure plots the median uninterrupted session length for various definitions of interruptions, as in Figure 4. The performance of ViFi is even better than *BestBS* and




Figure 8: The behavior of *BRR* and ViFi along a path segment in VanLAN. Black lines represent regions where the reception ratio was more than 50% in 1-second intervals. Dark circles represent interruptions.




Figure 9: TCP performance in VanLAN.

closely approximates *AllBSes*. It is notable that our simple and practical opportunistic protocol is able to beat the performance of the ideal single-BS protocol and approximate the ideal multi-BS protocol. In §5.4, we show that this high performance does not come at the cost of efficiency.

Figure 8 illustrates the behavior of *BRR* and ViFi, in a format similar to Figure 3. These are average case examples for the performance of these two protocols; individual runs differ. The paths are similar but not identical as they represent different days. We see that with *BRR* the path has several interruptions. ViFi performs significantly better, with only one interruption.

5.3 Application performance

Our experiments also show that the resilient link-layer connectivity of ViFi translates into better performance for interactive applications. In particular, we evaluate short TCP transfers, which are motivated by typical Web workloads, and Voice over IP. In these experiments, unacknowledged packets are retransmitted by the source at most three times. We find that higher limits yield similar or slightly worse results.

5.3.1 Performance of TCP transfers

Our TCP experiments evaluate two performance measures: (i) the time to complete a transfer; (ii) the number of completed transfers in a session, where a session is a period of time in which no transfer attempt was terminated due to a lack of progress. The vehicle repeatedly fetches a 10 KB file from a machine connected to the wired network and the machine does the same in the other direction. Transfers that make no progress for ten seconds are terminated and started afresh; we impose this limit because some transfers either hang or take a very long to complete due to packet losses at inopportune times in the TCP exchange.




Figure 10: TCP performance in DieselNet.

Figure 9(a) shows the median time to complete a transfer. To isolate the benefits of diversity and salvaging in ViFi, the middle bar shows the median time for a configuration in which diversity was enabled but salvaging was disabled. The results show that ViFi’s median TCP transfer time is about 0.6 seconds, which represents a 50% improvement over *BRR*. This improvement is higher than what would be predicted based on the number of additional packets delivered by the link layer (Figure 2). This brings out the difference between improvement in aggregate performance versus performance of interactive applications when using diversity. The figure also shows that most of ViFi’s gain is a result from diversity, although salvaging does provide a noticeable gain of about 10%. Given that we find that only 1.2% of the packets are salvaged, this benefit of salvaging is disproportionate. It confirms our intuition that the few packets that get stuck at older basestations when the vehicle moves away can disproportionately hurt the TCP performance.

We find that ViFi’s TCP performs comparably to modern cellular technologies. We added an EVDO Rev. A based cellular modem to one of our vehicles and generated a similar TCP workload. The median connection time in the downlink was 0.75 sec and in the uplink was 1.2 sec. (Cellular data rates are asymmetric.)

Figure 9(b) shows the average number of completed transfers per session. The average for ViFi is more than twice of *BRR*. Combined with its lower transfer times, this implies that users of ViFi will experience fewer disrupted transfers as well as better performance for individual transfers.

5.3.2 Performance of VoIP traffic

We evaluated the performance of VoIP sessions over ViFi by measuring the length of uninterrupted sessions. Supporting VoIP is more challenging than TCP because quality is sensitive to both loss and delay. The industry-standard for evaluating a voice call is the *Mean Opinion Score* (MoS), which ranges from 1–5, with labels of perfect (5), fair (4), annoying (3), very annoying (2), and impossible to communicate (1). MoS is a perceptual measure, but it is commonly estimated from an R-factor score [11] as: 1, if $R < 0$; 4.5, if $R > 100$; and $1 + 0.035R + 7 \times 10^{-6}R(R - 60)(100 - R)$, otherwise. R-factor is sum of four terms $R = 100 - I_s - I_d - I_{ef} + A$, where I_s is the signal-to-noise impairments, I_d and I_{ef} are impairments due to delay and loss, and A is expectation factor, which is higher when users expect lower quality. The impairments are functions of the codec.

We use the G.729 codec, which is implemented on most VoIP devices. For simplicity, we set A to zero (though it may be higher given the challenging environment). Then, the R-factor reduces to [11]: $94.2 - 0.024d - 0.11(d - 177.3)H(d - 177.3) - 11 - 40\log(1 + 10e)$, where d is the mouth-to-ear delay which includes the coding delay, network delay, and the delay introduced by the jitter buffer, e is the total loss rate which includes losses in the network and losses due to late arrivals, and H is the Heaviside step function: $H(x) = 1$ if $x > 0$; 0 otherwise.




Figure 11: Median length of uninterrupted VoIP sessions.

Per the codec, we generate 20-byte packets every 20 ms. Following convention, we assume that the coding delay is 25 ms, the jitter buffer is 60 ms, and that the wired segment of the end-to-end path adds 40 ms (corresponding to cross-country paths in the USA) to each VoIP packet. Aiming for a mouth-to-ear delay of 177 ms (because the impairment due to delay increases significantly beyond that) means that packets that take more than 52 ms in the wireless part should be considered lost. We measure one-way delays by applying piecewise linear regression [27] to remove clock skew and assuming that the minimum one-way delay is identical in the two directions.

We quantify the VoIP performance using the median length of time between interruptions. We deem an interruption to have occurred when the MoS value drops below 2 for a three-second period. Three seconds is roughly the time it takes to enunciate a short English sentence and a MoS value lower than 2 represents a severe disruption in call quality. We are not aware of an existing method to evaluate voice calls when packet delay and loss varies with time, and this definition seemed reasonable to us. We also studied different MoS thresholds and time periods within this framework. Results for those fit the qualitative behavior of Figure 7: the relative advantage of ViFi over *BRR* increases as the definition of an interruption becomes more stringent.

Figure 11 shows the results for our deployment on VanLAN and trace-driven simulations based on DieselNet. Because our results indicate that salvaging brings little benefit for VoIP, we do not isolate diversity and salvaging components of ViFi in the figure. The results show that the average session lengths are much longer with ViFi: the gain is over 100% in VanLAN, over 50% in Channel 1 of DieselNet, and over 65% in Channel 6 of DieselNet. We find that the overall call quality with ViFi is better as well. In VanLAN, the average of three-second MoS scores is 3.4 with ViFi and 3.0 with *BRR*. Thus, our results show that users of ViFi experience better call quality and significantly fewer disruptions in their voice calls.

5.4 Efficiency of medium usage

The higher application performance of ViFi does not stem from it using the medium more aggressively; in fact, its overall efficiency is comparable to that of *BRR*. We measure efficiency as the number of application packets delivered per transmission, in the channel between the vehicle and the BSes.

We compare ViFi with *PerfectRelay* and *BRR*. In the *PerfectRelay* protocol, exactly one basestation relays only if the intended destination did not hear the packet. We estimate its efficiency using packet-level logs of ViFi. In the upstream direction, a packet is considered delivered by *PerfectRelay* if at least one BS hears it. In the downstream direction, a complication is that even if a BS relays the packet, the vehicle may not hear it. We get around this by: (i) assuming that the outcome of the relaying will be identical to that of ViFi if at least one of the BSes relayed the packet; and (ii) the relaying is successful if no BS relayed it in ViFi.

		Upstream	Downstream
A1	Median number of auxiliary BSes	5	5
A2	Average number of auxiliary BSes that hear a source transmission	1.7	3.6
A3	Average number of auxiliary BSes that hear a source transmission but not the acknowledgment	0.6	2.5
B1	Source transmissions that reach the destination	67%	74%
B2	Relayed transmissions corresponding to successful source transmissions (i.e., false positives)	25%	33%
B3	Average number of auxiliary BSes that relay when a false positive relay occurs	1.5	1.5
C1	Source transmissions that do not reach the destination	33%	26%
C2	Cases where at least one auxiliary BS overhears a failed source transmission	66%	98%
C3	Cases where zero auxiliary BSes relay a failed source transmission (i.e., false negatives)	10%	34%
C4	Relayed packets that reach the destination	100%	50%

Table 1: Detailed statistics on the behavior of ViFi in VanLAN.

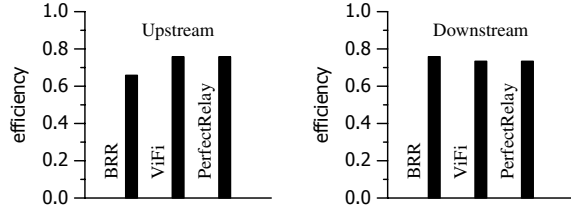


Figure 12: Efficiency of medium usage in VanLAN.

Figure 12 shows the results for the TCP experiments in VanLAN (§5.3.1). For upstream, we see that the efficiency of ViFi is better than *BRR* and nearly as high as *PerfectRelay*. For downstream, all three protocols have similar efficiency. *BRR* has slightly better efficiency because in ViFi the BS chosen to relay a packet may be distant. Considering both directions together, we find that ViFi is slightly more efficient.

5.5 Effectiveness of coordination

In this section, we present detailed statistics on the behavior of ViFi to provide insight into the effectiveness of its coordination mechanism. Table 1 shows data from the TCP experiments in VanLAN. Row B2 shows that ViFi has few false positives, that is, relayed packets that are already present at the destination divided by the number of successful source transmissions. Comparison with the average number of auxiliary BSes that receive the source transmission (Row A2), we can infer that the coordination mechanism of ViFi is effective at curtailing unnecessary relaying. Without that, the false positive rate would have been 170% and 360%. Row A3 reveals that overhearing acknowledgments sent in response to either the source or a relayed transmission is not sufficient to curtail false positives; probabilistic relaying is needed as well. If auxiliary BSes deterministically relayed whenever they hear source transmission but not an acknowledgment, the false positive rate would have been 60% and 250%.

Row C2 confirms that auxiliary BSes are often in a position to relay packets that do not reach the source. Row C3 shows that in such cases, the false negative rate is low. We define false negative rate as the number of times no auxiliary relays a failed transmission divided by the number of failed source transmissions. Combining the two rows, we can infer that roughly 65% of the lost source transmissions are relayed in each direction.

5.5.1 Comparison with other formulations

We compare ViFi’s coordination mechanism with three other formulations. Each formulation violates one of the three guidelines outlined in §4.4.

	ViFi	$\neg G1$	$\neg G2$	$\neg G3$
False positives	19%	50%	40%	157%
False negatives	14%	14%	12%	10%

Table 2: Comparison of different downstream coordination mechanisms for DieselNet Ch. 1.

- $\neg G1$: Auxiliary BSes ignore the presence of other potential auxiliary BSes. Each relays with a probability equal to its delivery ratio to the destination.
- $\neg G2$: Auxiliary BSes ignore loss rate to the destination. Each relays with a probability equal to $\frac{1}{\sum_i c_i}$, where c_i is that the auxiliary BS i is contending (Eq. 3)
- $\neg G3$: Auxiliary BSes relay such that the expected number of packets received by the destination is 1. (Recall that in ViFi, the expected number of packets relayed is 1.) Within this constraint, the objective is to minimize the number of relays. This formulation is an optimization problem: $\min \sum_i r_i \cdot c_i$ subject to $\sum_i r_i \cdot p_{B_i d} \cdot c_i \geq 1$.

An optimal solution to this optimization problem, is $r_i = 0$ if $s_i > 1$; $r_i = 1$ if $s_i + p_{B_i d} \cdot c_i < 1$; and $r_i = \frac{1-s_i}{p_{B_i d} \cdot c_i}$ otherwise; where $s_i = \sum_{j: p_{B_j d} \geq p_{B_i d}} p_{B_j d} \cdot c_j$. In simpler terms, this solution first picks the auxiliary BS_x with the highest $p_{B_x d}$ and sets $r_x=1$. It stops if that satisfies the constraint above. Otherwise, it picks BS_y with the next highest $p_{B_y d}$. If the constraint is satisfied by $r_y=1$, then y relays with a probability $\frac{(1-p_{B_x d} \cdot c_x)}{p_{B_y d} \cdot c_y}$. Otherwise, r_y is set to 1, and the BS with the third highest $p_{B_i d}$ is picked, and so on.

We find that compared to these other schemes ViFi strikes a good balance between false positives and false negatives. Table 2 shows the results for simulations over DieselNet’s Channel 1 environment. We see that while the false negatives for all schemes are roughly similar, ViFi has substantially lower false positives. Further, we observe in our experiments that the number of packets saved by $\neg G2$ is a lot lower than ViFi and that the false positive rate of $\neg G1$ increases rapidly with the number of auxiliary BSes. As shown in our technical report, application performance for all three schemes is worse than that for ViFi [3].

5.5.2 Limitations

Finally, we tested the relaying mechanism of ViFi in a range of simulated conditions to understand where it might perform poorly. We find two such conditions. First, when the number of auxiliary BSes is high (e.g., greater than 15). Second, all auxiliary BSes are equi-distant from both the source and the destination. In both conditions, while the average number of relays per packet is one

(Eq. 1), the variance in the number of relays per packet increases, resulting in higher false positives and negatives. Neither of these situations arise in our testbed environments. To make ViFi robust in environments where they might, it can be extended such that the number of auxiliary BSes is limited or the symmetry between them is broken. These extensions are subject of future work.

6. DEPLOYMENT ASPECTS OF ViFi

In this section, we comment briefly on the deployment related aspects of ViFi. ViFi requires changes to BSes and clients that may create an initial barrier to adoption but we believe that these barriers are surmountable. In the case of city-wide mesh networks [38, 39] operated by a single administrative entity, operators can unilaterally choose to deploy ViFi. In the case of organic deployments in individual residences and offices, service models pioneered by Fon [16], where a service provider supplies BSes for shared access, can pave an effective deployment path.

A different issue is whether WiFi deployments would be broad enough to enable more than a few city blocks of contiguous coverage; the lack of coverage between WiFi islands can render interactive applications unusable. However, a mix of ViFi and cellular modes can be used to maintain connectivity in such areas. Client devices can use ViFi — the cheaper option — where available and use cellular elsewhere. Some cellular providers already let users switch between WiFi and cellular to save the more expensive cellular minutes [32].

Finally, ViFi is beneficial only if clients often hear multiple BSes on the same channel. While already true of organic deployments (Section 3.4.1; CarTel [8]), this may not hold by default for city-wide meshes if they are engineered in a cellular pattern with neighboring BSes on different channels. In this setting, BSes can be equipped with an auxiliary radio such that neighbors of a BS are tuned to the same channel as the BS. These auxiliary neighbors interfere only minimally because they do not transmit often on the BS-client channel. They transmit data on that channel only when a downstream packet is overheard but the acknowledgment is not heard. Upstream packets are not relayed on the BS-client channel.

7. RELATED WORK

Our work benefits from and builds upon a large body of work in wireless handoffs and routing. What sets it apart is its goals and the unique constraints of its target environment: enabling common interactive applications from moving vehicles using WiFi. We divide prior work into four categories and contrast our work with examples of each.

Using multiple BSes ViFi is inspired by the successful use of macrodiversity in cellular networks [37], where multiple BSes act in concert to improve client performance.² The cellular methods, however, require tight integration with the physical layer and strict timing across BSes. These abilities need expensive BS hardware that is not suitable for commodity wireless deployments. ViFi is a macrodiversity method built on top of off-the-shelf WiFi radios.

In the WiFi context, Distributed Radio Bridges [23], Divert [25], and MRD [26] also use multiple BSes to improve client performance in enterprise WLAN deployments. The BS coordination mechanism in these systems assumes that a high-capacity LAN is available. For instance, in MRD, BSes coordinate by sending all received frames to a central controller that is responsible for forwarding only one copy to the Internet. Thus, if clients typically

²In contrast, microdiversity (e.g., MIMO) improves direct communication between two nodes. It brings complementary gains [13] and can be used in our setting as well.

reach three BSes, the required LAN capacity is at least three times the cumulative sending rate of all clients. Because a high-speed backplane is typically not available in our setting, the coordination mechanism of ViFi imposes little additional load on the backplane.

MultiNet [10], FatVAP [22], and PERM [36] enable clients to associate with more than one nearby BS, to increase throughput if the wireless capacity is greater than the capacity of wired links behind the BSes. The focus of this work is improving connectivity of the client-BS communication.

Opportunistic routing in static mesh networks Protocols such as ExOR [5] and MORE [9] share our goal and challenge in leveraging opportunistic receipt of packets with low coordination overhead. Their approach is to batch packets to amortize overhead across the batch; the authors recommend using a batch size of at least around ten. Batching, however, is unsuitable for most interactive uses. For instance, VoIP cannot afford the delay associated with waiting for ten packets. For short TCP transfers, the sender's congestion window will frequently be smaller than the batch size. Even for bigger transfers, batching may interact poorly with TCP's rate control, as mentioned by the authors of ExOR. ViFi, in contrast, uses a novel probabilistic coordination mechanism that operates on individual packets. In the future, we plan to study its performance in static mesh scenarios as well.

Network access from moving vehicles Early works on WiFi performance for vehicular access are based on controlled settings, with near line-of-sight connectivity and little interference [29, 17]. They find a relatively benign environment. Our study of more realistic settings, with WiFi and non-WiFi interferers and obstacles such as trees and buildings, reveals a challenging radio environment with frequent disruptions.

Several works consider the problem of transferring data using TCP through individual BSes as the vehicle drives by them, without maintaining connections across BSes [8, 19, 14, 18]. They find that performance in this setting is severely hindered by overheads at several layers, such as DHCP and aggressive TCP backoffs due to losses, and propose methods to lower these overheads. We investigate the possibility of continuous connectivity across BSes. We find that even if some of the overheads they observe (e.g., DHCP) are removed completely, the basic link layer connectivity remains problematic, especially for interactive applications. An interesting avenue for future work is to investigate the extent to which some of the methods proposed by these works (e.g., for aggressive TCP backoffs) are needed when the underlying link-layer connectivity is improved using ViFi.

MobiSteer shows that equipping vehicles with directional antennae can significantly improve performance [28]. Our work is based on omnidirectional antennae because, given the high cost and large form factor of directional antennae, typical clients (e.g., laptops, PDAs) are likely to have omnidirectional antennae. Further, while directional antennae extend reach, they do not prevent connectivity disruptions which we show can occur even close to BSes. ViFi can complement the gains from directionality when multiple BSes are visible in the current sector of the antenna.

Rodriguez *et al.* study the performance of vehicular clients while transferring data using cellular networks [33]. WiFi, the focus of our work, merits an independent examination. It differs from cellular in many ways, has a much shorter range, and operates in unlicensed spectrum. We find it interesting that even though the cellular technology is expensive and the networks carefully planned, like us, these authors find a challenging radio environment with unpredictable and sharp drops in connection quality.

Fast Handoffs There is a large body of work on minimizing the delay associated with handoffs in wireless networks [31, 2, 34,

6, 20]. This delay can be a major source of disruption in networks that otherwise have good wireless connectivity. Our work instead is focused on improving the basic connectivity itself which is quite challenging even if the handoff delays are minimal.

8. CONCLUSIONS

Our work improves WiFi performance for interactive applications. Using measurements from testbeds in two different cities, we showed that hard handoff methods that are used by WiFi clients today are poorly suited for the vehicular environment. These methods lead to frequent disruptions in connectivity. We also showed that methods that leverage basestation diversity are effective because they mask many disruptions.

We then designed, ViFi, a practical and efficient handoff protocol that exploits opportunistic receptions by nearby BSes to minimize disruptions for clients. The key to its effectiveness is a decentralized probabilistic algorithm that obviates per-packet coordination. Based on a two-month long deployment and trace-driven simulations, we showed that ViFi has close to ideal link-layer performance and significantly improves interactive experience. Our deployed prototype doubled the number of successful TCP transfers and doubled the length of disruption-free VoIP calls compared to a hard handoff protocol.

Acknowledgments We thank Victoria Poncini and Brian Zill for help with deploying the VanLAN testbed and thank Brian Lynn for help with collecting data on DieselNet. This work was supported in part by the NSF Grant CNS-0519881 and ARO award W911NF-07-1-0281.

9. REFERENCES

- [1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *SIGCOMM*, Aug. 2004.
- [2] Y. Amir, C. Danilov, M. Hilsdale, R. Musaloiu-Elefteri, and N. Rivera. Fast handoff for seamless wireless mesh networks. In *MobiSys*, June 2006.
- [3] A. Balasubramanian, R. Mahajan, A. Venkataramani, B. N. Levine, and J. Zahorjan. Interactive wifi connectivity for moving vehicles. Technical Report TR-2008-18, Dept. of Computer Science, University of Massachusetts, 2008.
- [4] A. Balasubramanian, Y. Zhou, W. B. Croft, B. Levine, and A. Venkataramani. Web search from a bus. In *CHANTS*, Sept. 2007.
- [5] S. Biswas and R. Morris. ExOR: opportunistic multi-hop routing for wireless networks. In *SIGCOMM*, Aug. 2005.
- [6] V. Brik, A. Mishra, and S. Banerjee. Eliminating handoff latencies in 802.11 WLANs using multiple radios: Applications, experience, and evaluation. In *IMC*, Oct. 2005.
- [7] J. Burgess, B. Gallagher, D. Jensen, and B. N. Levine. MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks. In *Proc. IEEE INFOCOM*, April 2006.
- [8] V. Bychkovsky, B. Hull, A. Miu, H. Balakrishnan, and S. Madden. A measurement study of vehicular Internet access using unplanned 802.11 networks. In *MobiCom*, Sept. 2006.
- [9] S. Chachulski, M. Jennings, S. Katti, and D. Katabi. Trading structure for randomness in wireless opportunistic routing. In *SIGCOMM*, Aug. 2007.
- [10] R. Chandra, P. Bahl, and P. Bahl. MultiNet: Connecting to multiple IEEE 802.11 networks using a single wireless card. In *INFOCOM*, Mar. 2004.
- [11] R. G. Cole and J. H. Rosenbluth. Voice over IP performance monitoring. *CCR*, 31(2), 2001.
- [12] D. D. Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *MobiCom*, Sept. 2003.
- [13] L. Dai, S. Zhou, and Y. Yao. A comparative study of the effects of microdiversity and macrodiversity on CDMA forward-link capacity. *IEEE Int'l Conference on Communications*, 3, 2002.
- [14] J. Eriksson, H. Balakrishnan, and S. Madden. Cabernet: A content-delivery network for moving vehicles. Technical Report TR-2008-003, MIT, CSAIL, 2008.
- [15] S. Floyd, V. Jacobson, C. Liu, S. McCanne, , and L. Zhang. A reliable multicast framework for light-weight sessions and application level framing. *IEEE/ACM ToN*, 5(6), 1997.
- [16] Fon. <http://www.fon.com/>.
- [17] R. Gass, J. Scott, and C. Diot. Measurements of in-motion 802.11 networking. In *Workshop on Mobile Computing Systems and Applications (WMSCA)*, Apr. 2006.
- [18] S. Giordano, D. Lenzarini, A. Puiatti, and S. Vanini. Enhanced DHCP client. Demo at CHANTS, Sept. 2007.
- [19] D. Hadaller, S. Keshav, T. Brecht, and S. Agarwal. Vehicular opportunistic communication under the microscope. In *MobiSys*, June 2007.
- [20] M. ho Shin, A. Mishra, and W. Arbaugh. Improving the latency of 802.11 hand-offs using neighbor graphs. In *MobiSys*, June 2004.
- [21] D. B. Johnson, D. A. Maltz, and J. Broch. DSR: The dynamic source routing protocol for multihop wireless ad hoc networks. In *Ad Hoc Networking*, 2001.
- [22] S. Kandula, K. Lin, T. Badirkhanli, and D. Katabi. FatVAP: Aggregating AP backhaul bandwidth. In *NSDI*, Apr. 2008.
- [23] V. Leung and A. Au. A wireless local area network employing distributed radio bridges. *ACM/Baltzer Wireless Network Journal*, 2, 1995.
- [24] R. Mahajan, B. Zill, and J. Zahorjan. Understanding WiFi-based connectivity from moving vehicles. In *IMC*, Nov. 2007.
- [25] A. Miu, G. Tan, H. Balakrishnan, and J. Apostolopoulos. Divert: fine-grained path selection for wireless LANs. In *MobiSys*, June 2004.
- [26] A. K. Miu, H. Balakrishnan, and C. E. Koksal. Improving Loss Resilience with Multi-Radio Diversity in Wireless Networks. In *MobiCom*, Sept. 2005.
- [27] S. B. Moon, P. Skelly, and D. Towsley. Estimation and removal of clock skew from network delay measurements. In *INFOCOM*, Mar. 1999.
- [28] V. Navda, A. P. Subramanian, K. Dhanasekaran, A. Timm-Giel, and S. Das. MobiSteer: Using directional antenna beam steering to improve performance of vehicular Internet access. In *MobiSys*, June 2007.
- [29] J. Ott and D. Kutscher. Drive-thru Internet: IEEE 802.11b for automobile users. In *INFOCOM*, Mar. 2004.
- [30] C. Perkins. Mobile IP. *IEEE Communications Magazine*, 35(5), 1997.
- [31] I. Ramani and S. Savage. Syncscan: Practical fast handoff for 802.11 infrastructure networks. In *INFOCOM*, 2005.
- [32] M. Reardon. Switching from cell to Wi-Fi, seamlessly. http://www.news.com/Switching-from-cell-to-Wi-Fi,-seamlessly/2100-1039_3-6113223.html, Sept. 2006.
- [33] P. Rodriguez, R. Chakravorty, J. Chesterfield, I. Pratt, and S. Banerjee. MARS: A commuter router infrastructure for the mobile Internet. In *MobiSys*, June 2004.
- [34] S. Seshan, H. Balakrishnan, and R. H. Katz. Handoffs in cellular wireless networks: The Daedalus implementation and experience. *Wireless Personal Communications (Kluwer)*, 4(2), 1997.
- [35] A. C. Snoeren, D. G. Andersen, and H. Balakrishnan. Fine-grained failover using connection migration. In *USITS*, Mar. 2001.
- [36] N. Thompson, G. He, and H. Luo. Flow scheduling for end-host multihoming. In *INFOCOM*, Apr. 2006.
- [37] A. J. Viterbi, A. M. Viterbi, K. S. Gilhousen, and E. Zehavi. Soft handoff extends CDMA cell coverage and increases reverse cell capacity. *IEEE JSAC*, 12(8), Oct. 1994.
- [38] City-wide Wi-Fi rolls out in UK. <http://news.bbc.co.uk/2/hi/technology/4578114.stm>.
- [39] Cities unleash free Wi-Fi. <http://www.wired.com/gadgets/wireless/news/2005/10/68999>.