

MRMV: Design and Evaluation of a Multi-Radio Multi-Vehicle System for Metro-WiFi Access

Pralhad Deshpande
IBM Research India
Manyata Embassy Business Park
Bangalore - 560045, India

Zafar Ayyub Qazi, Samir R. Das
Computer Science Department
Stony Brook University
Stony Brook, NY 11794, USA

ABSTRACT

In this work, we design and evaluate the MRMV (Multi-Radio Multi-Vehicle) system for vehicular WiFi access in the 2.4 GHz band. The design essentially centers around the vehicular client that accesses typical metro-WiFi networks (V2I) under vehicular mobility. MRMV has two features that makes it unique – i) it uses multiple WiFi interfaces that intelligently associates to different APs to mask hand-off latencies, (ii) it is able to use other MRMV clients as relays (V2V) and is thus able to avoid coverage holes. The V2V link uses 900 MHz interfaces for interference avoidance and better ranges. We provide extensive performance evaluation of the MRMV system using a large scale metro-WiFi deployment. We show that both connectivity (periods of non-zero throughput) and median throughputs improve substantially over default cases. Overall performance numbers indicate that the MRMV system can be an excellent platform for offloading data from cellular data networks to unlicensed bands for ubiquitous and high-throughput vehicular connectivity.

Categories and Subject Descriptors

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

General Terms

Performance; Measurement; Experimentation.

Keywords

Vehicular Internet Access; WiFi.

1. INTRODUCTION

The ‘Connected Car’ concept has been gaining ground for the past several years. Examples include General Motors’ OnStart [5], Toyota’s Entune [3] and Ford’s SYNC [7]. The

general idea is to provide remote vehicle monitoring, navigational assistance, entertainment and general Internet access to the vehicle’s occupants. Dashboard and backseat access to maps, Internet radio, Internet video streaming sites like Youtube and Netflix are expected to be commonplace in the near future. The current prototypes, however, assume a cellular broadband connection using licensed spectrum as the backhaul to provide these services. This indeed provides ubiquitous connectivity. However, it is expensive and the demand on mobile broadband networks is increasing exponentially [1], motivating the exploration of alternative ways to have high speed data connectivity on the move.

In recent years, several metro-scale WiFi deployments have cropped up in US and other developed countries [28]. They have indeed succeeded in making WiFi ubiquitous providing very good coverage in urban spaces. Thus, offloading mobile data to urban WiFi networks has been considered by several research groups [17, 22, 18, 10]. This also provides an opportunity for cost savings for both cellular broadband providers and consumers [2].

One of the most vexing problems for vehicular WiFi access is intermittent connectivity. The loss of connectivity faced by a vehicular client has two reasons. First, the vehicular client experiences frequent and time consuming hand-offs. Prior measurements in [16] using a similar metro-WiFi deployment indicate that a vehicular client may experience handoffs every 45 seconds. Also, the lossy nature of outdoor environments results in loss of association related frames, making handoffs in vehicular WiFi access environments very expensive. Second, even in dense metro-scale WiFi deployments, coverage holes do exist given that such network operators do not make any concerted effort to cover all the roadways.

This paper is a natural followup of our prior work on vehicular access of metro-WiFi networks [16] where we have demonstrated competitive or better median throughput performance for WiFi access in realistic driving conditions relative to a 3G network. While the median performance was impressive, there were significant periods of zero throughput due to large handoff latency and/or coverage holes. Thus, in [22] the authors proposed a hybrid access solution so that WiFi and 3G networks are both used in a seamless fashion to mask the disconnection issue. In this paper, we take this investigation to its natural extreme, where we address the problems of handoff latencies and coverage holes by using a Multi-Radio-Multi-Vehicle (MRMV) system.

In the MRMV system, each vehicular client can associate with multiple road-side APs simultaneously thus masking

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

VANET’13, June 25, 2013, Taipei, Taiwan

Copyright 2013 ACM 978-1-4503-2073-3/13/06 ...\$15.00.

handoff latencies. Also, when a client finds itself in a coverage hole, it can treat other vehicular clients in the vicinity as relays for its data and communicate with them over a relay link. The general goal is to demonstrate a system that provides a competitive or superior experience both in terms of average throughput and periods of connectivity relative to cellular data networks, while using a unlicensed band. With the cost and size of WiFi radios reducing by the day [4], using multiple radios does appear to be a viable option.




Figure 1: Vehicular WiFi access model demonstrating the use of multiple radios and relays (see text for explanation).

A schematic of our MRMV system in operation is shown in Figure 1. In the schematic, the green vehicle is exiting the space covered by the green AP and entering the space covered by the blue AP. If a single radio interface is used, the client would only try to associate with the blue AP once its connectivity with the green AP is lost. This is not a quick process. The client will experience loss of connectivity even when the space is well covered by the WiFi network. However, in a multi-radio system, a second radio interface can associate with the blue AP even before the client loses connectivity with the green AP. This way, a multi-radio client can mask handoff latencies almost completely. Also in Figure 1, the red vehicle is in a coverage hole but it can use the green vehicle as a relay reducing the impact of the coverage hole.

In this paper we make the following contributions:

1. Design an architecture of the MRMV system with multiple V2I (vehicle-to-infrastructure) interfaces and a single V2V (vehicle-to-vehicle) interface. The V2V interface operates in the 900 MHz band to reduce interference and obtain better range. See Section 3.
2. Develop an AP filtering technique to exploit multiple radio interfaces on vehicular clients to mask handoff latencies. This reduces the perceived disconnection periods (periods of zero throughput) from about 40% to about 5% when full WiFi coverage is available. See Section 4.
3. Demonstrate that the relaying opportunity and relay performance both are excellent. Using a relay vehicle can potentially reduce the disconnection periods by factor of 2 when coverage holes are indeed present. See Section 5.
4. All experiments are done in realistic driving conditions over hundreds of miles of actual driving often using

two cars. With the developed protocol features a sustained median throughput (UDP) of over 3 Mbps are obtained.

2. RELATED WORK

2.1 Exploiting AP Diversity from Moving Vehicles

Several prior studies have explored using AP diversity from moving vehicles [11, 32, 37]. ViFi [11] exploits macro diversity i.e., using simultaneously multiple APs and opportunistic receptions by near-by APs to reduce disruptions for mobile clients. R2D2 [32] uses multi-lobe beam pattern switching on a smart antenna to use multiple access points while using directional antennas. However, in real metro-scale AP deployments, near-by APs may not be on the same channel to minimize interference. In such real settings, ViFi and R2D2 might not be able to exploit performance benefits from near-by APs which operate on different channels. Secondly, both ViFi and R2D2 require APs to coordinate with each other, which requires modifications to existing AP deployments. Our work on the other hand tries to exploit AP diversity in real metro-scale deployments without assuming any modifications to existing AP deployment. In a recent work, Spider [37], a node associates with multiple APs using a single radio to avoid association and dhcp delays. Spider performs scheduling to decide when to switch to an AP and how much time to spend on each AP. However, their protocol performs poorly in scenarios where nearby APs are in different channels because of inherent channel switching delays in commodity hardwares.

2.2 Vehicular WiFi Handoff

Several techniques have been proposed to improve the handoff mechanism of a mobile client when switching associations in a wireless network. The solutions address two issues related to regular handoff mechanism: reducing the handoff latency and improving handoff decisions. Handoff latency is improved by reducing the number of channels to probe [34], reducing the number of control messages exchanged [18] and using a second radio to pre-associate to an alternate AP [12]. Handoff decisions are improved either by using historical information [29, 17] or by scanning actively [20, 26, 33].

However only [29, 17, 20] propose handoff mechanisms that are designed for vehicular environment. The authors in [17] propose a scripted handoff using historical information of signal strengths to decide which AP to associate to and also eliminate scanning or probing delays. Mobisteer [29] uses steerable beam directional antenna and past historical information to improve AP selection process. The protocol in [20] actively scans even when a client is associated to select the best AP. However, all these schemes experience handoff latency. In this work, one of the problems that we address is how we can mask handoff latencies using multiple radios in vehicular settings.

2.3 Vehicle-to-Vehicle Communication

There have been several experimental studies in the context of V2V communication. In [35], the authors measure the performance of 802.11b-based V2V communication in different environments. They show that the performance of

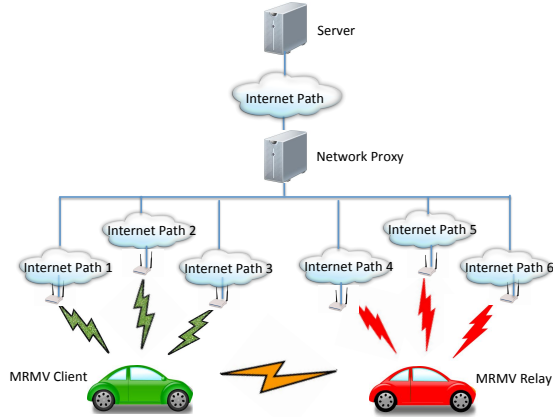


Figure 2: MRMV system architecture. Each vehicular client has multiple (3 in the experiments) WiFi interfaces to connect to multiple APs in the metro-WiFi network

V2V links are greatly affected by the nature of the environment. Authors in [27] report experiences with static 1 and 3-hop scenarios and a mobile 3-hop scenario. TCP and UDP performance results are presented in a 2-hop vehicular network in [24]. The work in [38] uses directional antennas to improve the transmission range in the context of V2V communication.

The work in [39] proposes a vehicle-vehicle relay scheme to improve coverage but only performs simulations and some controlled experiments using a single AP. In this work, one of the problems that we address is that given real metro-scale WiFi deployments, whether we can use V2V relay to reduce physical coverage holes.

3. THE MRMV SYSTEM

An architecture that supports switching between multiple communication paths is an integral part of the MRMV system. The architecture may be addressed at different layers of the protocol stack. Although link layer architectures [36, 14] exist, they are infeasible in networks that span different domains. Application and transport layer solutions have also been proposed [23, 25]. They however involve cumbersome application or server side changes. Network layer solutions, on the other hand are transparent to the applications and servers and are easy to deploy. While most limitations of other approaches are overcome using a network layer approach, efficiency concerns do exist because it operates further down the stack. Research efforts such as [13] are aimed at minimizing these inefficiencies. In light of these observations we use a network layer architecture to be part of the MRMV system. In this section we will describe our network layer architecture, the data transfer procedure using the architecture and our implementation of the MRMV design.

3.1 Network Layer Architecture

Figure 2 shows a high level view of the network layer architecture. Each MRMV node is potentially connected to




Figure 3: Data transfer between the vehicular client and a server on the Internet using the network proxy. (SIP (DIP) : Source (Destination) IP address.)

the Internet via multiple network interfaces each having a care-of IP address obtained from the metro-WiFi network. Apart from having multiple paths to infrastructure nodes, the nodes also form an ad-hoc network over V2V links. Each node can also use other nodes in the vicinity as relays for its traffic. The V2V interfaces operate in the 900MHz frequency range for interference avoidance with the V2I communications and also because 900MHz offers better propagation characteristics resulting in better ranges. The V2V interfaces are statically assigned IP addresses from the ad-hoc network's domain. The network layer architecture also comprises of a network proxy. MRMV nodes communicate with remote servers via the proxy. In addition to the V2I and V2V interfaces, each MRMV node also possesses a virtual interface. Each node acquires a fixed IP address for the virtual interface from the proxy and uses it to establish connections with the remote server over the best route from those presently available. Each node registers its multiple care-off IP addresses with the proxy. The proxy maintains a mapping table to map the various care-off addresses of a client with its fixed IP address. When application traffic originating from a node passes through the proxy, the care-of address used is marked as the active address. When application traffic originating from a remote server passes through the proxy's domain, the proxy intercepts the packets and tunnels them using IP-in-IP encapsulation to the client's active address. This mechanism is similar to that used in Mobile IP [30] but has been extended to handle multiple interfaces and relay nodes.

3.2 Data Transfer in the MRMV System

As in Mobile IP, each node obtains a fixed IP address from the network proxy. It also obtains care-of addresses from the metro-WiFi network for the V2I interfaces. The V2V interface is statically assigned an IP address belonging to the vehicular ad hoc network's domain.

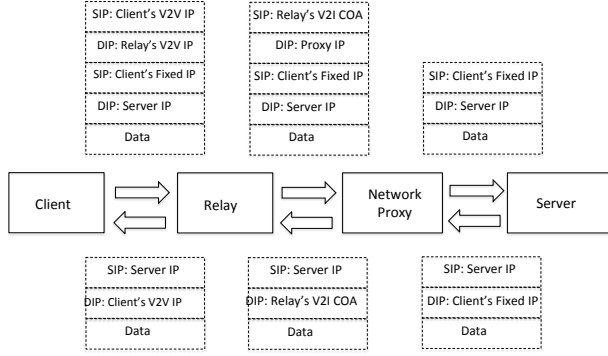


Figure 4: Data transfer between the vehicular client and a server on the Internet using the network proxy and an intermediate V2V relay. (SIP (DIP) : Source (Destination) IP address.)

Direct V2I data transfer (without V2V relay): An MRMV client uses its fixed IP address (assigned by the proxy) to communicate with servers on the Internet. Referring to Figure 3, the client first encapsulates the IP packet with another IP header and routes it to the network proxy via a V2I interface. The network proxy strips off the outer IP header and forwards the packet to the server. When the server wants to communicate with an MRMV client, the packet gets routed to the proxy's network. The packet is intercepted by the network proxy and the destination IP address is re-written as the client's V2I care-of address. Finally, all packets arriving on any of the client's V2I interfaces are forwarded to the client's virtual interface with the fixed IP address.

Data transfer using a V2V relay: When routing packets using a relay (Figure 4), the MRMV client first encapsulates the IP packet so as to forward it to the MRMV relay. The IP header is decapsulated at the relay node and re-encapsulated using the proxy's IP address as the destination. The packet is then forwarded to the proxy over the V2I interface. The network proxy strips off the outer IP header and forwards the packet to the remote server. On the return route, a packet from the server is routed to the proxy's network where it is intercepted by the proxy. The proxy replaces the destination IP address by the relay node's V2I care-of address and forwards the packet to the MRMV relay. Like the proxy, the relay node also maintains a mapping table for the client's fixed and V2V IP addresses. The MRMV relay then replaces the destination IP address with the client's V2V IP address and forwards it to the MRMV client over the V2V link. Any packet arriving on the client's V2V interface is forwarded to the client's virtual interface with the fixed IP address.

3.3 Relay Selection

An important part of the MRMV system is the relay selection procedure. As previously discussed, each node maintains multiple routes to the Internet. If any of the available routes are not being used for data transfer, then MRMV

nodes can offer to relay other nodes' data on those routes. In the following we describe the relay selection procedure.

1. If an available V2I interface is not being used for any data transfer, then the node may decide to operate it in relay mode and offer its use to the neighbors in the ad-hoc network. It does this by periodically broadcasting a message of the format,

$$RELAY_OFFER(IP_i, \langle RSS_1 \rangle, \dots, \langle RSS_n \rangle),$$

where, n is the number of unused available interfaces that the node (with fixed IP address i) has to offer and RSS_j is the recent average RSS observed for the V2I link on the j^{th} interface. Knowledge of RSS enables the MRMV system to estimate the rates on the V2I interfaces assuming a rate control protocol such as BRAVE [15]. BRAVE considers recent values of RSS experienced on a link for rate selection such as throughput is maximized by considering loss rates for all physical rate data rates available on the interface. Thus, RSS essentially translates to an estimated throughput.

2. When a client node looking for a relay receives the broadcast $RELAY_OFFER$ messages from its neighbors, it logs the RSS of the received packets along with senders' V2V IP addresses and the message contents. The choice of the actual relay depends on the estimated throughput through the V2V and V2I links. Once again a rate control protocol like BRAVE [15] is assumed. This helps estimate the throughputs.

The client node simply selects the relay that provides the best throughput, i.e., $\max(\min(\text{V2V throughput}, \max(\text{V2I throughput over each interface})))$ over all such relays. Then it sends a message to the chosen relay in the form:

$$RELAY_REQUEST(j, IP_{fixed}).$$

This message requests for the relay's j^{th} interface to be used by the client's fixed IP.

3. When a $RELAY_REQUEST(j, IP_{fixed})$ message is received the receiving node sets its j^{th} interface in relay mode and sends a $LEASE(IP_j, IP_{fixed}, t)$ message to the proxy. This message tells the proxy to tunnel the packets for IP_{fixed} using IP_j for t seconds. The relay node creates an entry in its mapping table in order to forward packets from the proxy to the V2V interface of the client. It then sends a $RELAY_ACCEPT(t)$ message to the client node indicating that the relay route is available for t seconds.
4. The client node can use the route for t seconds. If the route is no longer needed, then a $RELAY_RELEASE(j, IP_{fixed})$ message is sent to the relay node. The relay node then relinquishes the lease by sending an $UNLEASE(IP_j, IP_{fixed})$ message to the proxy to undo the IP_{fixed} to IP_j mapping for tunneling and stops forwarding packets for the client node. In case the client wants to use the route for more than t seconds, then it sends another $RELAY_REQUEST$ message before t seconds are up. In case no such message is received, then the relay relinquishes the lease by sending an $UNLEASE(IP_j, IP_{fixed})$ message to the




Figure 5: Components of an MRMV node.

proxy and also removes the entry for the client from its own mapping table.

5. If a relay node while forwarding packets to the client node on its j^{th} interface loses its V2I connection on that interface, it sends the client a *RELAY_UNAVAILABLE* message. The relay node then relinquishes the lease by sending an *UNLEASE*(IP_j, IP_{fixed}) message to the proxy to undo the IP_{fixed} to IP_j mapping for tunneling and stops forwarding packets for the client node.

We acknowledge that in the relay protocol presented here there may exist potential security and privacy issues. In this preliminary study we focus on performance issues and do not address security and privacy related issues. However, we do believe that most of these issues can be addressed by borrowing from a vast array of literature on cryptography. For instance, for controlling how a *LEASE* is distributed, we can use public-key cryptography and digital signatures whereby the proxy gives a *LEASE* to a relay node only if it receives a request signed by the client nodes' private key.

3.4 Prototype Implementation

Our implementation of the MRMV system consists of two MRMV nodes. The current implementation is targeted towards a comprehensive performance evaluation. One node operates as a dedicated MRMV client node and the other as a dedicated MRMV relay node. Each node has three V2I interfaces and one V2V interface. Each node consists of four Avila GW2348-2 single board computers (SBC) [9] as shown in Figure 5. Each board is fitted with a single



Figure 6: Vehicular antenna set up used in the driving experiments.

carrier-grade radio interface. Three boards are fitted with Ubiquiti XR2 [8] interfaces (acting as the V2I interfaces) and one is fitted with Ubiquiti XR9 [8] interface (acting as the V2V interface). The XR2 and XR9 uses 802.11b/g protocol in the 2.4 GHz band and 900 MHz bands respectively. The transmit powers of all interfaces are set to 25 dBm. The interfaces use Atheros chipset supported by madwifi-based custom driver provided by Ubiquiti. A Garmin 18x USB GPS device [19] is used to log the GPS coordinates.

The radio interfaces are connected to our vehicular antenna setup shown mounted on a car in Figure 6 using low-loss RF cables. The antenna setup comprises of three 2.4 GHz antennas and one 900 MHz antenna mounted on top of the car using a mounting rack. The antenna gains for the 2.4 GHz and the 900 MHz antennas are 8 dBi and 6 dBi respectively. Each vehicle also has a laptop which is connected to the four SBCs using a switch. The laptop acts as the actual client. The laptop decides which of the three SBCs to use for the V2I connectivity if at all available. It does this by monitoring the connectivity status and the link qualities (RSS) experienced by the three SBCs and changing its own routing table to route data through the SBC that is connected and has best link quality. The V2I connectivity status and link qualities are determined by running a per-second ping from each of the SBCs to the DNS server and keeping track of the lost ping messages and the RSSes.

If none of the SBCs has connectivity, the 900 MHz relay link can be used to forward the data to a neighboring MRMV relay node that is able to send the data on the client's behalf using one of its own V2I interfaces.

We use a metro-scale WiFi deployment in the Long Island area in New York. This service is called Optimum WiFi [6], and is provided by Cablevision, a local cable TV provider and ISP. The WiFi network covers most of the populous areas of Long Island where our study is conducted. The network also has coverage in parts of New York City, Pennsylvania, Connecticut and New Jersey, where Cablevision has service. The entire network has roughly 18,000 APs. The deployment consists of both indoor and outdoor WiFi hotspots. Note that there are several hundreds of such metro-scale WiFi deployments in USA alone [28]. While we expect that our general observations will be repeatable in

another metro area deployment, we do note that specific quantitative observations are likely to be strongly tied to deployment density, radio characteristics and any handoff control on the APs.

Our driving experiments are done along two routes; a short drive of 9 miles and a long drive of nearly 100 miles. The short drive is very urban and offers excellent Optimum WiFi coverage. The coverage along the long drive is variable.

4. MASKING HANDOFF LATENCY

In the first set of experiments, we study the throughput performance of a MRMV client node *under real driving conditions* along the short drive. The drive has been repeated several times until the statistics converged and the averages are reported here. The drives go with the prevailing traffic that is variable – roughly 30-45 miles/hour. There are several stops at the intersections. These drives are of similar nature reported in prior work related to vehicular access of metro-WiFi networks [15, 16, 22].

We first configure the MRMV node so that each V2I interface uses the default handoff strategy and functions independently i.e., without any coordination with other interfaces. This serves as the base case to understand performance limitations and issues. Based on the insights for these experiments, we then develop our own handoff mechanism for MRMV clients using a novel AP filtering technique, described in Section 4.1, and report results from experimental evaluation in Section 4.2 and Section 4.3. In all cases, each interface is configured to use the 11 Mb/s bit-rate, which uses the DSSS modulation scheme. Prior studies [18] and our own experience show that this rate works well under vehicular mobility. Also, a static bit-rate choice helps in interpreting the results better without being influenced by bit-rate selection protocols. The stock Madwifi driver is used (version 9.0.4). Each node runs three saturated UDP streams (upload) to a server with a public IP address located at our lab in the university. For ease of implementation, we restrict our experiments to upload only. The analysis of the base case reveals two observations about handoffs in dense AP deployments.

Handoffs are very frequent: Our results show that the average time per association for the three interfaces is about 61 sec and the average distance travelled per association is about 466 meters (less than half the average span of an AP). Also, the average link quality (RSS) when an interface associates with a new AP is about 31 dB that is very high. This is because an interface using the stock Madwifi driver in roaming mode performs handoff from one AP to another when its link quality with the current AP drops below a threshold (10 dB above noise floor). At this time the interface associates with the AP that has the strongest signal from among those in its cache that is populated by background scans. It is easy to imagine in a dense metro-WiFi deployment, a client following such a protocol remains associated with an AP over much less than its span – from very near to the AP where the link quality is excellent to when the link is barely usable (10 dB above noise floor). This results in frequent handoffs. With careful design it should be possible to reduce the number of handoffs substantially by increasing the length of association with APs.

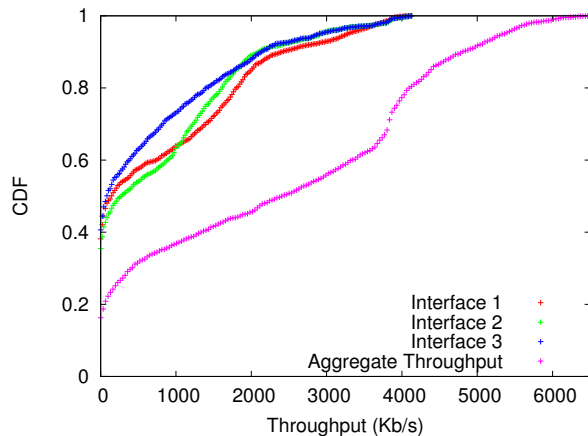


Figure 7: CDF of individual and aggregate per-second throughputs when each interface uses the default handoff strategy. Bit-rate is fixed at 11 Mbps.

Handoffs are very expensive: With reference to the above experiment, we plot the CDF of the per-second throughputs seen at the server in Figure 7. Each interface is unable to communicate with the server (zero throughput) about 40% of the time. Overall, looking at the aggregate throughput plot the MRMV node does not have connectivity to the server about 20% of the time. It demonstrates a significant handoff latency so that the vehicular client is unable to take advantage of the excellent deployment. This motivates a design where the V2I interfaces coordinate with each other strategically to successfully mask handoffs.

4.1 AP Filtering

To address the issue of loss of connectivity due to frequent and expensive handoffs, we propose a novel two-part AP filtering technique. The filtering technique assumes that radio fingerprinting data (RSS of APs with location) is available along the routes to be driven and an MRMV node can perform such filtering. Alternately, a central authority that is aware of the metro-WiFi network deployment may supply the filtering information directly. This strategy has some similarity with the ‘scripted handoff’ technique proposed in our prior work [17] except that handoffs are not directly controlled. This way the technique is straightforward to implement, has a low overhead and is agnostic to the driver software.

Outdoor vs indoor filtering: Outdoor APs have much larger coverage spans than indoor APs. To avoid frequent handoffs a MRMV node should only associate with outdoor APs. This is akin to associating with only macrocells for fast moving mobiles and not with microcells – a strategy followed often in a multilayer cellular network to reduce handoff overhead [31]. We filter out a large number of APs with small coverages (supposedly indoor APs). Such a filter is easily implemented by modifying the Madwifi driver to accept a whitelist of APs. So, the driver only processes beacons of APs belonging to the whitelist. The rest of the beacons are simply dropped.

Filtering for mutually exclusive AP assignment: The second level of filtering is based on the idea that to best ex-




Figure 8: CDF of aggregate per-second throughput with various filtering levels. Bit-rate is fixed at 11 Mbps.

exploit AP diversity different interfaces on the MRMV client should associate with different APs. To facilitate such mutually exclusive assignment without any added protocol overhead, we form a serial order of the APs (order in which they will be encountered along the drive) and assign them to the interfaces in round-robin manner. This assignment again is implemented using a whitelist.

Further, not all outdoor APs are used for the round-robin assignment. The average inter-AP distance in the AP whitelist may also impact handoff performance; e.g. if the whitelist APs assigned to a particular interface are too close to each other then that might induce frequent handoffs. Hence, we also experiment with different average inter-AP distances in the AP whitelist. By selectively filtering out outdoor APs one can arrive at a round-robin assignment that reflects a particular average inter-AP distance. For example, an ‘RR-300’ filtering is a round-robin assignment of outdoor APs that have an average inter-AP distance of approximately 300 m. At any given time we use a single radio for data communication. Intelligently using multiple radios concurrently will likely improve cumulative throughput performance but is unlikely to reduce handoff latency any further.

4.2 Performance of AP Filtering

Figure 8 shows the CDF of per-second UDP throughputs (aggregated over all three interfaces) as observed by the server.¹ Various filtering levels are used. It is clear that increasing the average inter-AP distance improves median aggregate throughputs and connectivity up-to RR-600 level of filtering. For RR-600 filtering, the median aggregate throughput at the server is approximately 3800 Kb/s and the MRMV node has connectivity to the server 96% of the times (versus 80% of the times in the default case). Beyond this the median aggregate throughput starts drop-

¹Note that aggregated performances are shown here in order to make comparisons between various scenarios. Multiple interfaces are solely used as an instrument to mask handoff latency.




Figure 9: CDF of RSS of received ACKs at different filtering levels.

ping and the connectivity reduces (see plots for RR-750 and RR-1000).

It is interesting to observe that this increase in median throughput and coverage happens despite the fact that the overall link qualities become progressively worse with increased AP filtering. Figure 9 shows the CDF of RSS values of the link layer ACKs received at the interfaces. Notice that for RR-600 filtering (best case) the median RSS is about half that of the default handoff case.

Table 1 captures some key statistics when using different average inter-AP distances for filtering. As can be seen from the table, the average time spent while being associated with a single AP increases with increasing average inter-AP distance. This translates to an increase in distance in meters over which a single association lasts. Note that the maximum range of association is around 800 m and RR-600 filtering achieves this range. Further increasing average inter-AP distance also impacts the distance travelled between associations. Also, note that staying associated with a single AP for longer time entails associating with it much earlier than when the client is associated with it for shorter times. This is seen in the last row of table 1. The RSS in dB when associating with a new AP progressively decreases with increased average inter-AP distance. In the default case (no filtering), a new association happens when the link quality is quite high (RSS over 30 dB). In contrast, this number drops down to nearly 10 dB for RR-1000. We conclude that the stock Madwifi driver misses out on a significant part of the production phase (RSS greater than 15 dB [21]) which can be better exploited if the vehicular client associates with an AP earlier like in the case of RR-600 filtering. As a result of better exploitation of the production phase, the resulting median throughput is higher for RR-600 filtering. Even though for RR-750 and RR-1000 filtering, the production phase of an AP is fully exploited the average distance travelled between associations is large enough to adversely impact the overall coverage and throughput.

	Default	RR-200	RR-300	RR-375	RR-600	RR-750	RR-1000
Avg. time per association (sec)	61	59	62	82	93	110	112
Avg. distance travelled per association (m)	466	497	524	736	802	806	814
Avg. time between associations (sec)	17	17	25	27	56	77	87
Avg. distance travelled between associations (m)	164	158	226	285	537	732	743
Avg. RSS when associating (dB)	32	27	19	20	13	13	10
Coverage (fraction of time with non-zero throughput)	84%	88%	93%	97%	96%	94%	92%

Table 1: Key observations with various levels of filtering.

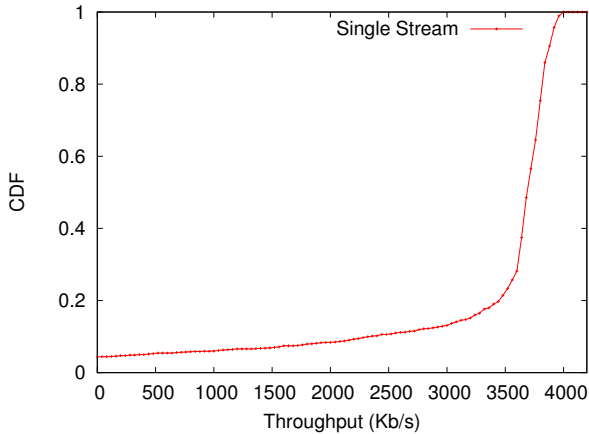


Figure 10: CDF of per-second single stream throughput using the best possible interface (RR-600 filtering is used).

4.3 Throughput Performance of MRMV Client

Here, we are interested in evaluating a single stream’s performance on an MRMV client node via the best possible interface. We implement the interface selection scheme as described in Section 3.4 on top of the AP filtering mechanism (RR-600 is used for its superlative performance). Recall that the selection process simply runs a per-second ping on all (connected) interfaces and uses the RSS of the ping responses to determine the interface to use. The UDP traffic is routed through the selected interface. The short drive is again used for evaluation. Figure 10 shows the CDF of the per-second throughputs as seen at the server. Note that the connectivity (fraction of time with non-zero throughput) is 95% and the median throughput is around 3500 Kbps. Also, 80% of the times the throughput is between 3300 Kbps and 4200 Kbps. This shows a promise of reduced variance in per-second throughputs. Overall, the throughput is excellent, significantly higher than 3G throughputs in similar driving conditions in the same region [16] and likely comparable with 4G throughputs.

5. OVERCOMING COVERAGE HOLES

So far, we focused on improving handoff performance so that each MRMV client node sees an improved connectiv-

ity and throughput. However, this assumes perfect WiFi coverage. When the coverage holes are present the MRMV system uses nearby cars with V2I connectivity as relays. To establish the performance potential of our design, we present several performance results. In particular, we present relay performance in realistic driving conditions along the long drive.

5.1 Evaluating Relay Performance

Recall that the V2V link operates in the 900 MHz band for its better propagation performance and for interference freedom from 2.4 GHz. We first evaluate the performance of the V2V link separately by using two MRMV vehicular nodes driving together with one car following the other at different distances. This experiments are done on the long drive to get a greater amount of variability in the road and driving conditions. Saturated UDP traffic is sent over the V2V link from the sender car to the receiver car. Figure 11, the average throughput drops gracefully with inter-car distance up-to a distance of 1750 m. After this, we barely see any connectivity. In general, the throughput performance under real driving conditions is very good – over 3000 Kbps even beyond a distance of 1000 m. This is given that often the cars do not have line-of-sight due to road curvatures and intervening vehicles.

We now evaluate a complete system performance with two MRMV nodes under the same condition with one vehicle following the other on the long route under real driving conditions. No specific attempt was done to maintain a specific distance and in fact it was not even possible due to traffic conditions. However, the drivers tried best not to follow very closely (this ensures that the V2I links of both cars behave in a correlated fashion making any performance gain from the relay unlikely) and also not to exceed 1750 m for the inter-car distance so that the V2V link remains connected. For the performance measurement, one MRMV node acts as a dedicated client and the other as a dedicated relay. The client node sends saturated UDP traffic to a server located in the lab and opportunistically uses the relay node to relay traffic as described in Section 3.3 when the client is in a coverage hole (zero V2I throughput). Since the relay is dedicated no relay selection protocol is used. Thus such protocol overhead does not creep in the measurements. For the V2I interfaces on both the client and relay nodes an AP filtering mechanism as presented in the previous section is used for optimized V2I performance.

Figure 12 shows the CDF of throughputs observed when client node uses the relay and when it does not. Note that




Figure 11: Average per-second throughputs at various distances for the 900 MHz V2V link under real driving conditions for a two vehicle experiment on the long drive.

using the relay improves the connectivity to 84% from 70% and median throughput from about 2000 Kbps to about 3100 Kbps. The V2V link for this experiment has been up for about 85% of the time with an average throughput of 4000 Kbps. Overall, when the client node is in a coverage hole, 48% of the packets received by the relay is successfully delivered to the server. We suspect a better fraction could be delivered when a multiple potential relays are available due to diversity. As mentioned before, when the vehicles are close to each other, it is likely that when the client node is in a coverage hole, the relay node is there as well.

6. CONCLUSIONS

To make the ‘Connected Car’ concept a reality one must develop inexpensive, ubiquitous wireless access technologies with abundant bandwidth. While WiFi provides this opportunity, prior work on vehicular WiFi access was either limited to small deployments or showed significant periods of disconnection when metro-scale deployments were indeed used. In this work, we have addressed these disconnection issues at its core. We developed the Multi-Radio-Multi-Vehicle (MRMV) system that (i) uses multiple radio interfaces to mask the handoff latencies, (ii) uses other MRMV nodes as relays when a client loses V2I connectivity on all its interfaces. We demonstrated a mechanism to exploit multiple interfaces effectively, it uses an AP filtering technique to exploit the knowledge of the deployment. We show that this technique is able to reduce the periods of disconnections substantially over the default single interface setting. The relaying part of the work shows excellent relaying opportunity and significant reduction of coverage holes. Overall, we have demonstrated that when WiFi coverage is excellent, it is possible to bring down the periods of disconnections to a mere 5% (from 40% in the default case) with median throughputs over 3.5 Mbps. When coverage holes are indeed present, the relay protocol is able to ‘plug’ about half of the coverage holes continuing to provide a median throughput of over 3 Mbps.




Figure 12: CDF of average per-second throughputs for the two vehicle relaying experiment on the long drive.

The current design has several limitations. It does not consider any form of interference or load balancing issues and focuses primarily on connectivity. The relay protocol does not address potential security and privacy issues. Actual driving experiments with a larger number of vehicles are also needed to understand performance potentials. Our future work will address these limitations. In addition, we plan to do extensive experiments to evaluate the proposed relay protocol and quantify the protocol overhead in vehicular settings.

Acknowledgements

This work was partially supported by NSF grants CNS-0831791 and CNS-1117719 and an equipment gift from Laufer Wind Group. The authors also acknowledge help from the students in the Wings Lab during experiments.

7. REFERENCES

- [1] 3G and 4G Wireless Speed Showdown: Which Networks Are Fastest? http://www.pcworld.com/article/253808/3g_and_4g_wireless_speed_showdown_which_networks_are_fastest_.html.
- [2] The business value of mobile data offload. <http://tinyurl.com/chy34zg>.
- [3] Entune. <http://www.toyota.com/entune/>.
- [4] Intel’s tiny wi-fi chip could have a big impact. <http://www.technologyreview.com/news/429299/intels-tiny-wi-fi-chip-could-have-a-big-impact/>.
- [5] Onstar service. <https://www.onstar.com/web/portal/landing>.
- [6] Optimum WiFi. <http://www.optimum.net/MyServices/WiFi/>.
- [7] Sync service. <http://www.ford.com/technology/sync/>.
- [8] Ubiquity Networks, Inc. <http://www.ubnt.com>.
- [9] Avila GW2348-2. <http://www.gateworks.com/products/avila/gw2348-2.php>.

- [10] A. Balasubramanian, R. Mahajan, and A. Venkataramani. Augmenting Mobile 3G Using WiFi. In *MobiSys*, 2010.
- [11] A. Balasubramanian, R. Mahajan, A. Venkataramani, B. N. Levine, and J. Zahorjan. Interactive WiFi connectivity for moving vehicles. *SIGCOMM Comput. Commun. Rev.*, 2008.
- [12] V. Brik, A. Mishra, and S. Banerjee. Eliminating handoff latencies in 802.11 WLANs using multiple radios: applications, experience, and evaluation. In *IMC*, 2005.
- [13] K. Chebrolu, B. Raman, and R. Rao. A network layer approach to enable tcp over multiple interfaces. *Wireless Networks*, 2005.
- [14] A. F. T. Committee et al. Inverse multiplexing for atm (ima) specification version 1.0. *Physical Layer Working Group*, 1997.
- [15] P. Deshpande and S. Das. Brave: bit-rate adaptation in vehicular environments. In *VANET*, 2012.
- [16] P. Deshpande, X. Hou, and S. Das. Performance comparison of 3G and metro-scale WiFi for vehicular network access. In *IMC*, 2010.
- [17] P. Deshpande, A. Kashyap, C. Sung, and S. R. Das. Predictive methods for improved vehicular WiFi access. In *MobiSys*, 2009.
- [18] J. Eriksson, H. Balakrishnan, and S. Madden. Cabernet: A WiFi-Based Vehicular Content Delivery Network. In *MobiCom*, 2008.
- [19] Garmin. <http://www.garmin.com/>.
- [20] A. Giannoulis, M. Fiore, and E. W. Knightly. Supporting vehicular mobility in urban multi-hop wireless networks. In *MobiSys*, 2008.
- [21] D. Hadaller, S. Keshav, T. Brecht, and S. Agarwal. Vehicular opportunistic communication under the microscope. In *MobiSys*, 2007.
- [22] X. Hou, P. Deshpande, and S. R. Das. Moving Bits from 3G to Metro-Scale WiFi for Vehicular Network Access: An Integrated Transport Layer Solution. In *ICNP*, 2011.
- [23] H. Hsieh and R. Sivakumar. A transport layer approach for achieving aggregate bandwidths on multi-homed mobile hosts. *Wireless Networks*, 2005.
- [24] F. Hui and P. Mohapatra. Experimental characterization of multi-hop communications in vehicular ad hoc network. In *VANET*, 2005.
- [25] L. Magalhaes and R. Kravets. Transport level mechanisms for bandwidth aggregation on mobile hosts. In *ICNP*, 2001.
- [26] V. Mhatre and K. Papagiannaki. Using smart triggers for improved user performance in 802.11 wireless networks. In *MobiSys*, 2006.
- [27] M. Moske, H. Fussler, H. Hartenstein, and W. Franz. Performance measurements of a vehicular ad hoc network. In *VTC*, 2004.
- [28] Muniwireless.com. List of cities and counties with large WiFi networks. <http://www.muniwireless.com/reports/Mar-28-2009-list-of-cities.pdf>.
- [29] V. Navda, A. P. Subramanian, K. Dhanasekaran, A. Timm-Giel, and S. R. Das. MobiSteer: using steerable beam directional antenna for vehicular network access. In *MobiSys*, 2007.
- [30] C. Perkins. IP mobility support, RFC 2002, October 1996.
- [31] G. Pollini. Trends in handover design. *IEEE Communications Magazine*, 1996.
- [32] K. Ramachandran, R. Kokku, K. Sundaresan, M. Gruteser, and S. Rangarajan. R2d2: Regulating beam shape and rate as directionality meets diversity. In *MobiSys*, 2009.
- [33] I. Ramani and S. Savage. SyncScan: practical fast handoff for 802.11 infrastructure networks. In *INFOCOM*, 2005.
- [34] M. Shin, A. Mishra, and W. A. Arbaugh. Improving the latency of 802.11 hand-offs using neighbor graphs. In *MobiSys*, 2004.
- [35] J. Singh, N. Bambos, B. Srinivasan, and D. Clawin. Wireless lan performance under varied stress conditions in vehicular traffic scenarios. In *VTC*, 2002.
- [36] K. Sklower, B. Lloyd, G. McGregor, C. D., and T. Coradetti. The PPP multilink protocol (MP), RFC 1990, August 1996.
- [37] H. Soroush, P. Gilbert, N. Banerjee, B. N. Levine, M. Corner, and L. Cox. Concurrent Wi-Fi for mobile users: analysis and measurements. In *CoNEXT*, 2011.
- [38] A. P. Subramanian, V. Navda, P. Deshpande, and S. R. Das. A measurement study of inter-vehicular communication using steerable beam directional antenna. In *VANET*, 2008.
- [39] J. Zhao, T. Arnold, Y. Zhang, and G. Cao. Extending drive-thru data access by vehicle-to-vehicle relay. In *VANET*, 2008.