

Supporting Vehicular Mobility in Urban Multi-hop Wireless Networks

Anastasios Giannoulis
Dept. of Electrical and
Computer Engineering
Rice University
Houston, TX 77005
agiannou@rice.edu

Marco Fiore
Dipartimento di Elettronica
Politecnico di Torino
Corso Duca degli Abruzzi 24
10129 Torino, Italy
marco.fiore@polito.it

Edward W. Knightly
Dept. of Electrical and
Computer Engineering
Rice University
Houston, TX 77005
knightly@rice.edu

ABSTRACT

Deployments of city-wide multi-hop 802.11 networks introduce challenges for maintaining client performance at vehicular speeds. We experimentally demonstrate that current network interfaces employ policies that result in long outage durations, even when clients are always in range of at least one access point. Consequently, we design and evaluate a family of client-driven handoff techniques that target vehicular mobility in multi-tier multi-hop wireless mesh networks. Our key technique is for clients to invoke an association change based on (i) joint use of channel quality measurements and AP quality scores that reflect long-term differences in AP performance and (ii) controlled measurement and hand-off time scales to balance the need for the instantaneously best association against performance penalties incurred from spurious handoffs due to channel fluctuations and marginally improved associations. We utilize a 4,000 user urban deployment to evaluate the performance of a broad class of hand-off policies.

Categories and Subject Descriptors: C.2.1 [Computer-Communication Networks]: Wireless Communication; C.4 [Performance of Systems]: Measurement Techniques

General Terms: Algorithms, Design, Experimentation, Measurement, Performance

Keywords: Mesh, Vehicular, Wireless, Urban, Roaming, Differential Capacity Handoff, AP Quality Scoring, Throughput Outage, Handoff Timescale

1. INTRODUCTION

Over 1,000 cities worldwide, from Taipei, Taiwan to Mountain View, USA, have deployed large-scale 802.11 networks. Such networks employ a three tier architecture including an access tier for mobile and residential client access, a backhaul “mesh” tier for wirelessly interconnecting access points, and a capacity injection tier in which directional or high capacity

This research was supported by NSF grants CNS-0325971 and CNS-0721894.

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.

MobiSys’08, June 17–20, 2008, Breckenridge, Colorado, USA.
Copyright 2008 ACM 978-1-60558-139-2/08/06 ...\$5.00.

wireless links inject capacity in order to increase available resources to nodes farther away from fiber gateways.¹

Unfortunately, despite providing signal coverage over large contiguous areas, our experiments indicate that today’s mesh networks cannot support vehicular mobility, as the clients and network infrastructure inherited design choices targeted towards 802.11 WLAN architectures. In fact, no standard handoff scheme is specified for 802.11 architectures, including 802.11s, the mesh standard, so that in practice, mobility is addressed according to poorly performing and proprietary manufacturer dependent mechanisms and policies. For example, a commonly implemented policy is for a client to maintain the current AP association until it fails, and then to initiate a new association to the new AP having the strongest signal strength: We will show that application of this policy to a vehicular client yields an average client handoff outage of 18 seconds, despite the client remaining within coverage of at least one access point at all times.

Prior research and standards have addressed vehicular handoff in cellular networks (see for example [1] and the references therein). However, such networks rely critically on network signaling to control handoff decisions, utilizing a high speed wireline backhaul (or a dedicated point-to-point wireless link) at each base station in order to provide a dedicated control channel. In contrast, large scale mesh networks employ multiple tiers and multiple hops of random access transmission, thus far precluding deployment of such finely coordinated control mechanisms. Likewise, WLANs that have each AP connected to Ethernet can also employ network coordinated handoff, e.g., [2]. Finally, existing work on client-driven policies also targets WLANs and therefore targets pedestrian client velocities, e.g., [3, 4]. Thus, no prior work addresses vehicular mobility in mesh networks.

In this paper, we present the following two contributions. First, we develop a client-side family of handoff policies, representing the first such design able to support vehicular mobility in large-scale mesh networks. Our key technique has two components. (i) We jointly incorporate AP-to-client channel quality with a long-time-scale per-AP quality score in the handoff decision. The quality score reflects a critical component of mesh networks: due to the multi-tier multi-hop architecture, different backhaul AP associations can yield significant long-time-scale differences in client

¹Despite only the middle tier having a mesh topology, all three tiers are collectively commonly referred to as a “mesh network.”

throughput. For example, an association to an AP that is connected directly to the capacity injection tier can typically provide superior service compared to APs that require multiple omni-directional hops to reach a gateway. Likewise, different backhaul links have differing inter-AP distances and propagation characteristics that can yield long-term differences in individual backhaul link capacities. Consequently, under our proposed policies, clients favor associations with such high quality-score APs, provided that their channel quality is sufficiently high to allow it. *(ii)* We employ mechanisms to control handoff frequency, as we will show that spurious handoffs as well as excessively maintaining a degrading association both yield significant performance penalties in mesh networks. In particular, the timescales of channel measurement and association decisions must be selected to respond to inherent changes in channel conditions of vehicular clients, while avoiding outages and throughput degradations that will occur if clients excessively search for a marginally improved association. We refer to this family as Differential Capacity Handoff policies, as they incorporate the observation that in mesh networks, capacity is non-uniformly distributed across spatial locations.

Second, we present the first experiments of vehicular mobility in an urban mesh network. Our research platform is the Technology For All network, a 4,000 user urban mesh network that covers an area of over 3 km² in an under-resourced community of Houston, TX [5]. Moreover, to enable client-side realization of the above policy family, we developed a HostAP-based wireless interface driver that supports all of the required functionality. Our experiments consist of 55 drives around a 2.5 km reference loop representing approximately 140 km and 825 performed handoffs. All points within the reference loop are covered by at least one AP so that our experiments contain no outages due to being “out of range.” Moreover, numerous additional experiments are presented for comparison purposes including forced handoffs in non-mobile situations and handoffs in WLANs. Our methodology characterizes the isolated and joint impact of each component of the Differential Capacity Handoff policy. Our key findings are as follows.

- **Baseline policies.** To guide our experimental study of the factors controlling handoff performance and provide a baseline for comparison, we study three baseline policies, the latter two being special cases of the Differential Capacity family: Maintain Until Broken, Always Strongest Signal, and Averaged with Hysteresis. The Maintain Until Broken policy selects the strongest-signal AP and does not handoff until that connection fails. This policy represents the default configuration of the SMC wireless interface and of many commercial clients [6]. We show that despite our client always being with range of at least one AP, this policy yields outage times that average 18 seconds per handoff. While a policy that proactively always selects the strongest signal yields vastly reduced outages, it unfortunately comes with a high penalty in client throughput due to an excessive number of spurious handoffs induced by channel fluctuations. Finally, we show that reducing handoff frequency by signal strength averaging at appropriately chosen time-scales coupled with a hysteresis mechanism used to control handoff frequency can dramatically improve client performance.

- **Evaluation of Differential Capacity Handoff.**

First, we show that the average client throughput obtainable while associated with different APs varies by a factor of up to three. Second, we show that the attained client throughputs for each AP have minimal variation over long time scales (days), indicating that differences among AP performance are not due to transient load or channel conditions but rather to architectural aspects of the network as described above. Third, we evaluate mechanisms for AP quality scoring in order to reflect these long-term throughput disparities. We show that weighting APs according to the averages of the (minimally varying) client throughput measurements yields a 50% percent throughput gain compared to decisions using only channel quality measurements and not AP quality scoring.² Moreover, we show that Differential Capacity Handoff outperforms the current Maintain Until Broken policy by 300%. Finally, we approximate the maximum throughput achievable in the reference loop via a hypothetical mobile client that always connects to the highest throughput AP and performs a minimum number of handoff. We find that the Differential Capacity policy obtains a throughput of 81% of this value.

- **Origins of throughput outage.** We explore the origins of throughput outages by designing a set of experiments to isolate the components of the outage duration. We use a non-mobile client and forced handoffs as a comparison baseline to show that *(i)* multi-hop wireless backhaul induces an order-of-magnitude increase in successful association time as compared to wireless LANs, *(ii)* failed association attempts dominate outage time as compared to the association delay, *(iii)* connection outages in the Differential Capacity family are close to the minimum average achieved in a non-mobile forced-handoff setting, and *(iv)* the outage duration measured by the receiver is sometimes masked by network effects; namely, packets in transit at an old association can continue to arrive to the receiver while the client is delayed in obtaining a new higher performing association.

The remainder of this paper is structured as follows: Sec. 2 describes the experimental testbeds and methodology. Sec. 3 provides a description of the proposed handoff policies. Sec. 4 experimentally quantifies the performance of signal-based policies. Sec. 5 investigates the Differential Capacity Handoff policy and Sec. 6 explores the origins of throughput outages. Finally, Sec. 7 discusses the related work and Sec. 8 concludes the paper.

2. EXPERIMENTAL PLATFORMS AND METHODOLOGY

In this section we provide a description of our client and network experimental infrastructure and our experimental methodology.

2.1 Client Platform

All experiments were conducted using a laptop inside a car, with a Linux 2.6.17 Operating System and an SMC

²Considering performance at locations covered by more than one AP, such that clients had a choice in association.

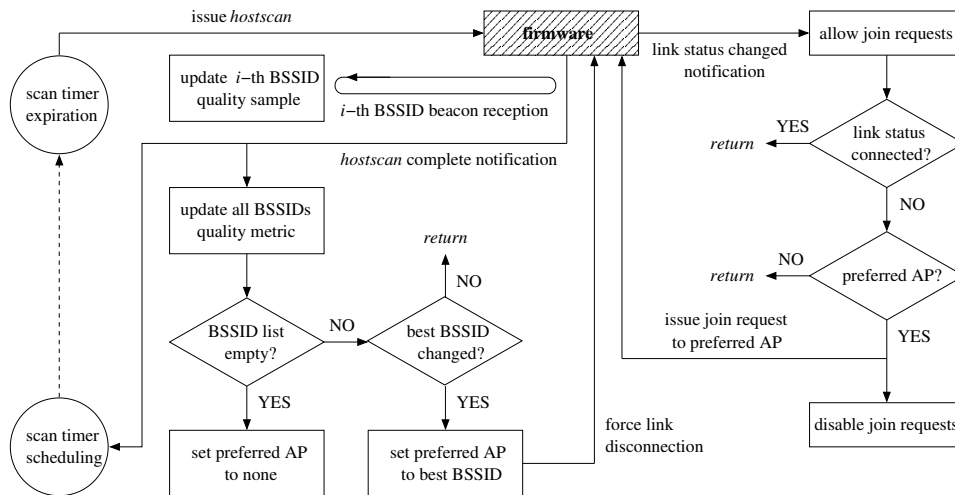


Figure 1: Diagrammatic description of the driver operation

2532-B card as an 802.11b wireless interface. The interface uses HostAP drivers, Prism 1.1.1 primary firmware and 1.8.0 secondary firmware. This card was connected to a 7 dBi external antenna, approximately 20 cm tall, which was mounted on top of the car. A GPS device was also connected to a USB port of the laptop to record location information.

By default, the handoff process of the SMC wireless interfaces is performed by the *firmware*. This only initiates handoffs after the recession of beacon reception from the associated AP. At that time, a new association is performed to the AP having the highest channel quality. In order to realize experiments for a broad family of handoff policies, we implemented an alternative device driver for the SMC wireless interface. The driver disables the above firmware operation. It is HostAP-based and provides a clean-state, tunable and flexible way for implementing a variety of handoff policies. The software is publicly available to allow further research on related topics.³

The additional functionality that this driver provides primarily consists of: (i) a periodic scheduling of AP scanning, employing the firmware’s *hosts can* function, (ii) Access Point quality evaluation following the reception of AP-beacons based on each handoff algorithm’s functions, and (iii) a mechanism to force a handoff via disconnection and subsequent re-association when a policy or experiment requires it. A diagrammatic description of the above driver operation is provided in Fig. 1.

2.2 Network Infrastructure Platforms

Our experiments were performed in an urban mesh network with the exception of a set of baseline measurements performed in an indoor WLAN.

TFA Network: The TFA Network is an urban mesh network deployed in southeast Houston by Rice University in cooperation with non-profit organization Technology For All (TFA).⁴ As of November 2007, the network consists of 20 Access Points, 4,000 users, and is under expansion. Fig. 2 depicts the topology and connectivity of the network.

³<http://networks.rice.edu/software.html>

⁴See [5], <http://www.techforall.org>, and <http://tfa.rice.edu>.

First, the network consists of an *access tier* in which clients access mesh nodes via a client to AP link. Clients, including our vehicular clients, typically use antennas with lower gain and from near ground locations with propagation obstacles intervening. The TFA network uses 20 nodes to provide coverage to an area of approximately 3 km². Within that area, nearly all locations are within range of at least one AP [5]. All client connections are rate-limited to 500 kbps, except for a single AP that employs 250 kbps rate limits for network management purposes.

Second, the network employs a *backhaul tier* in which the APs wirelessly interconnect to forward traffic amongst each other. In most cases, the origin or destination of a flow is the wireline Internet. Hence, most traffic ingresses or egresses at a single *fiber gateway* which is currently rate limited to 100 Mbps.

Finally, to limit the path-length of flows traversing the backhaul tier, a *capacity injection tier* employs high-performance directional links (depicted by darker lines). These links yield virtual gateways as they provide an independent channel compared to other access and backhaul links.

Each mesh node uses a high gain 15 dBi omnidirectional antenna placed approximately 10 meters above the ground, higher than most of the houses and some of the trees in the neighborhood. Access points located on either end of directional links employ two radios whereas other access points employ a single radio. The three tier architecture resembles architectures employed by large-scale commercial networks.

Rice Indoor Network: As a baseline for comparison of handoff performance, we also consider a WLAN network in Duncan Hall at Rice University. Its Access Points are placed approximately 2 meters above the floor, with Access Points located in all 3 floors of the building. A detailed topology of the network can be found at [7, page 4], for all 3 floors of Duncan Hall. The Rice WLAN has all APs directly connected to Ethernet, hence the network is neither multihop wireless nor multi-tier.

2.3 Experimental Methodology

Here we describe our experimental methodology for the measurements reported in this paper.



Figure 2: Topology of the TFA network and Reference Loop

In all experiments, the vehicle drives around a loop that is 2.5 km long following the route shown in Fig. 2. The vehicle’s average speed is 30 mph and the range is from complete stops at stop signs to 40 mph. This path was selected to ensure continuous coverage based on our measurements of signal propagation in the TFA network. Moreover, this route passes through an area with high AP density, hence enabling multiple connection and handoff options. Furthermore, this loop is within the coverage of APs that have the greatest possible range in their spatial and hop distance from the wired gateway. We refer to this route as the “reference loop.”

Our default platform employs a different range of IP addresses for clients associated with each Access Point. Hence, the client changes IP address every time it hands off. To avoiding execution of DHCP, we instrumented the vehicular client to manually sets its own IP address to a prespecified and reserved value for the experiments. The incurred delay of such an action is approximately 100 ms.

Measures for Evaluation: While roaming, the laptop in the vehicle sends UDP traffic to an external server at Rice University using *iperf*. We consider the impact of the intermediate wired path between the TFA gateway and Rice University to be negligible, whereas the wireless mesh network is the bottleneck in terms of throughput and delay. The client transmits 500 kbps of UDP upload traffic to provide a characterization of handoff behavior decoupled from congestion control effects.

Our key performance metrics are (i) throughput as measured by the received packets of the wireline receiver and (ii) throughput outage due to handoff. We define throughput outage as the time interval during which the server receives traffic at zero rate from the client.

3. HANDOFF POLICY DESIGN FOR MESH NETWORKS

In this section, we propose a family of client-driven hand-

off policies for mesh networks. Our key techniques are (i) balancing client estimated channel quality of the available mesh-to-client links with AP quality scoring, a measure of long-time-scale differences of backhaul node performance and (ii) controlling the timescales of channel inference and handoffs to avoid spurious handoffs. With this family of policies, clients can favor APs with better performance in multi-hop backhaul, limit handoff frequency to the minimum required level to avoid performance penalties for excessive handoffs, and ensure that the AP-client link is of sufficient quality to provide a high-performance connection.

To guide our experimental study of the factors controlling handoff performance and provide a baseline for comparison, we define three alternate policies that include special cases of the above framework and solely rely on the channel quality metric. Moreover, one such policy, Maintain Until Broken, represents existing systems.

3.1 Differential Capacity Handoff

We propose a general family of policies that couple channel quality assessment with a broad class of Access Point quality scoring criteria, with the objective of maintaining quality connectivity at vehicular speeds. Moreover, we control the handoff frequency to avoid throughput degradations that occur from spurious handoffs.

Channel Quality. A critical input to a handoff decision is the quality of the mesh-to-client links for the available APs as the channel quality limits the modulation rate, affects the packet loss rate, etc. Ideally, channel quality is based on the received SINR; in practice, it can be a scaled estimate of the received signal strength, and potentially include other factors such as packet loss. Thus, to measure the channel quality and smooth the client’s inference of channel quality to aid in controlling handoff frequency, we utilize an Exponential Weighted Moving Average (EWMA) filter for determining the signal-based link quality q_i for each wireless link to the Access Points.

Let γ denote the period for client scanning of all Access

Points. Denote $\sigma_i(k)$ as the client’s signal indicator from Access Point i at the k -th scan. Then the EWMA quality metric $q_i(k)$ for this wireless link is:

$$q_i(k) = \alpha q_i(k-1) + (1-\alpha)\sigma_i(k) \quad 0 \leq \alpha \leq 1 \quad (1)$$

The parameter α determines the memory of the filter and the weight of older scans, and together with γ , controls the measurement timescale.

AP Quality Score. To weight each AP, we associate to each a *quality score* $0 \leq w_i \leq 1$, with higher values indicating a stronger client preference for the AP. The quality score can incorporate a range of long-time-scale performance properties such as the long-term client throughput disparities among APs that exist in large-scale urban 802.11 deployments. Likewise, it can reflect long-term differences in delay or load. In Section 5 we show that even with consistently high channel qualities, different APs support differing maximum throughput due to factors such as backhaul connectivity and topology. Thus, AP scoring enables clients to favor the highest-quality AP from the perspective of backhaul performance when channel conditions allow. Such AP quality scores vary slowly and can be advertised by the network to the client in a one-time exchange or can even be measured by the client directly.

Thus, to balance long-term AP quality with rapidly varying channel conditions, we transform the signal-based quality indicators $q_i(k)$ to joint indicators $\hat{q}_i(k, w_i)$ that encapsulate both channel information and AP scores. In particular, we propose a transformation given by the following *piecewise linear* function:

$$\hat{q}_i(k, w_i) = \begin{cases} \frac{\delta \times w_i + T_1}{T_1} q_i(k), & \text{if } q_i(k) < T_1 \quad (2) \\ q_i(k) + \delta \times w_i, & \text{if } q_i(k) \in [T_1, T_2] \quad (3) \\ \left[1 - \frac{\delta \times w_i}{Q_{max} - T_2}\right] (q_i(k) - Q_{max}) + Q_{max}, & \text{if } q_i(k) > T_2 \quad (4) \end{cases}$$

where δ is a weighting control parameter, and T_1 and T_2 are quality thresholds that define the ranges for the different constituting parts of the transformation.

For a certain quality range $[T_1, T_2]$, signal-based link quality estimates are weighted by $\delta \times w_i$ (3), i.e., according to the quality score of an Access Point w_i and the weight of AP-quality vs. channel quality δ . However for the low channel quality range ($q_i < T_1$), lesser weighting of AP quality indicators is performed. This operation (2) is critical: without this functionality in the low quality range, weighting of link quality indicators by $\delta \times w_i$ may induce associations to APs that are out of transmission range. The degree of weighting in (2) is chosen to yield a continuous transformation function. Values for the threshold T_1 are such that a negligible weighting is performed for channel qualities for which the wireless interface is unable to successfully transmit to an Access Point. The last component of the transformation (4) is of secondary importance, as its only functionality lies in the preservation of an identical quality range between the transformation and the signal-based region. Example transformation filters are presented in Sec. 5.

Controlling Handoff Frequency. Finally, as channel conditions can vary rapidly, mechanisms are required to mit-

igate spurious handoffs and control the frequency of handoff. For example, we will show that it is critical to prevent clients from continually searching for an incrementally better association, as doing so would lead to a significant throughput degradation.

To achieve this objective, we employ a hysteresis mechanism as follows: Let β denote a hysteresis threshold that controls the client’s tolerance to handoff events. A handoff occurs from AP i to AP j when the *transformed* time-averaged quality of an alternate AP j , $\hat{q}_j(k)$, exceeds that of the currently associated AP $\hat{q}_i(k)$ by at least β . That is, handoff is invoked if

$$\hat{q}_j(k) > \hat{q}_i(k) + \beta \quad (5)$$

Thus, the hysteresis mechanism and the time scale of the smoothed channel-quality measurements jointly control the handoff frequency and limit spurious handoffs.

3.2 Baseline Handoff Policies

Here we define three signal-based policies. By their evaluation, we study factors of the handoff process and characterize the need for AP quality scoring as employed in the Differential Capacity handoff policy.

Maintain Until Broken. This policy maintains a connection between the client and an Access Point until the client considers the link to be broken. This break occurs when no beacons are received from the Access Point for a client-configured timeout duration. Upon disconnection, the client initiates a new connection to the Access Point that yielded the largest SINR for the received beacon.

This is a handoff policy that does not employ averaging, hysteresis, nor weighting depending on AP preference. This policy is the default configuration of the SMC wireless interfaces that we use in our experiments. Moreover, many clients of commercial mesh deployments utilize this technique [6].

Always Strongest Signal. This policy is a pro-active one which targets to be connected to the AP yielding the strongest received signal strength at all times. In particular, the client continually monitors the SINR of received beacons from the available APs. If the AP with the strongest signal is different from the current association, the client will initiate a handoff to the new AP. This is achieved by forcibly disconnecting with the current AP and associating with the new AP having the greater SINR.

This policy can be viewed as a degenerate case of Differential Capacity Handoff having no filter memory ($\alpha = 0$), no hysteresis ($\beta = 0$) and no weighting ($\delta = 0$). The performance study of this policy brings out the limitations of a myopic handoff initiation.

Averaged with Hysteresis. The Always Strongest Signal policy is vulnerable to invoking excessively rapid handoffs due to channel fluctuations. Thus, some dampening of the handoff decision can mitigate this effect. The Averaged with Hysteresis policy limits the frequency of handoffs by time-averaging the signal strength estimates and employing hysteresis. It can be considered as a special case of Differential Capacity Handoff, where no AP-score weighting is performed ($\delta = 0$). That is, the smoothed link-quality estimator q_i is used jointly with hysteresis such that handoff is invoked if $q_j(k) > q_i(k) + \beta$.

Finally, for handoff decisions based on signal strength, utilizing a maximum value for a handoff-decision threshold can prevent clients from unnecessarily handing off when they al-

ready have a link with the highest possible modulation rate. Hence, in our experiments, we ensure that for the Averaged with Hysteresis policy, handoffs are not invoked whenever the signal strength exceeds an experimentally determined maximum level.

4. EVALUATION OF SIGNAL-DRIVEN POLICIES

This section evaluates the baseline policies described in Sec. 3.2 via a set of experiments and measurements performed on the platform described in Section 2. Our objective is to assess the durations of zero throughput that result from handoff policies that are driven by the channel’s signal-quality indicators, and to study their impact on the connection’s attained throughput. As a result, factors and policies that determine the operation of the signal-based component of Differential Capacity Handoff are studied in isolation.

4.1 Maintain Until Broken

Here, we drive the reference loop with the client employing the Maintain Until Broken policy described above. We first consider a 7 dBi client antenna, a typical value for an end-user device, and next consider a 15 dBi client antenna, a value that matches the mesh node’s antenna.

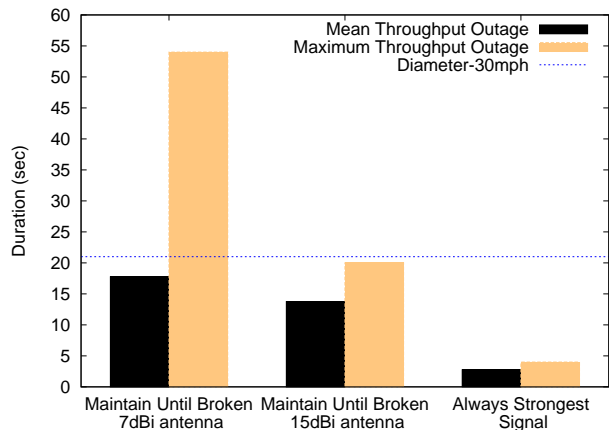


Figure 3: Throughput Outages for Maintain Until Broken and Always Strongest Signal Policies.

Mean Throughput Outage is 18 Seconds. Fig. 3 depicts the mean and maximum outage time for both 7 and 15 dBi client antennas. It also depicts a horizontal dashed line at 21 seconds. 21 seconds is the time required at 30 mph to traverse 250 meters, the mean distance between cell boundaries on the reference loop.

The figure indicates that with 7 dBi antennas, the mean and maximum outage duration is 18 and 54 seconds respectively, durations that are unacceptably large for many applications. Note that these outages are not due to lack of coverage, as all points in the reference loop are covered by at least one access point. Thus, the outages are due solely to the maintain until broken policy itself. In particular, the long outage times arise due to a client successfully receiving beacons from the AP (and hence maintaining its association), yet being unable to successfully transmit data due to the channel quality being too poor.

High-Gain Client Antennas Marginally Improve Average Performance. Comparing the 7 dBi bars with the 15 dBi bars in Fig. 3, further explores this effect. The results show that when clients employ high gain 15 dBi antennas matching the gain of the AP antenna, they obtain a significant decrease in the maximum and mean throughput outage. Furthermore, twice as many throughput outages occurred with a 7 dBi antenna, resulting to longer total outage duration. This is due to the 15 dBi antenna yielding a reduced spatial region in which the client can receive a beacon but not transmit data. Unfortunately, while improving worst-case performance significantly, high-gain client antennas yield average outages that are still quite poor.

4.2 Always Strongest Signal

We next repeat the 30 mph drive around the reference loop, yet now consider the Always Strongest Signal policy, as a prospective for the signal-based component of Differential Capacity Handoff. In this case, a handoff is invoked by the client any time an AP with superior channel conditions comes within range. Henceforth, all measurements have client antennas with 7 dBi gain.

Throughput Rapidly Decays with Switching Frequencies > Once per 10 Seconds. Fig. 3 indicates that this policy significantly reduces the outage times as compared to the Maintain Until Broken policy, even with high-gain client antennas. In particular, the Always Strongest Signal Policy yields a mean and maximum outage of 2.75 and 4 seconds respectively.

On the other hand, while yielding a significant benefit in outage time, this policy yields frequent handoffs. In particular, the above experiment resulted in an average of 5.7 seconds between handoffs, significantly less than the mean time within a cell of 21 seconds.

To isolate the effect of the handoff frequency on throughput we design the following experiment. A non-mobile client is placed within range of two APs such that the links from either AP to the client are high quality. The client chooses a handoff interval and switches between the two APs periodically according to the interval by initiating a handoff. During this continued handoff process, the client transmits a single long-lived UDP flow and we measure the received throughput of this flow.

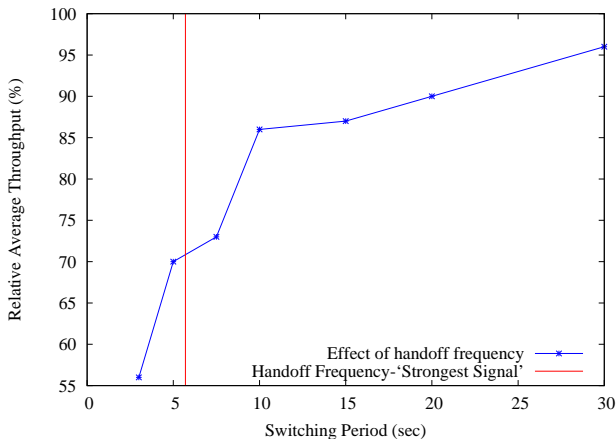


Figure 4: Impact of switching frequency

Fig. 4 depicts the results of the experiment. The x-axis indicates the period of the client’s handoff between the two APs and the y-axis indicates the throughput relative to the case of never handing off. The figure indicates that while the attained throughput monotonically decreases with increasing handoff frequency, the rate of the decrease is higher for low values of the handoff period. The handoff period of the Always Strongest Signal policy (6 sec) lies below a critical threshold (approximately 10 seconds), above which handoffs are performed at reduced throughput cost.

4.3 Time-Averaged with Hysteresis

Because excessive handoffs yield a throughput degradation, here we evaluate the signal-based components of Differential Capacity Handoff that limit handoff frequency. In particular, time averaging of signal-quality together with decision hysteresis jointly limit handoff frequency.

To evaluate the Time-Averaged with Hysteresis policy, we first perform a set of experiments to obtain the best filter parameters via exhaustive search, i.e., considering combinations of α , the filter memory parameter, β , the hysteresis parameter, and γ the scanning period.

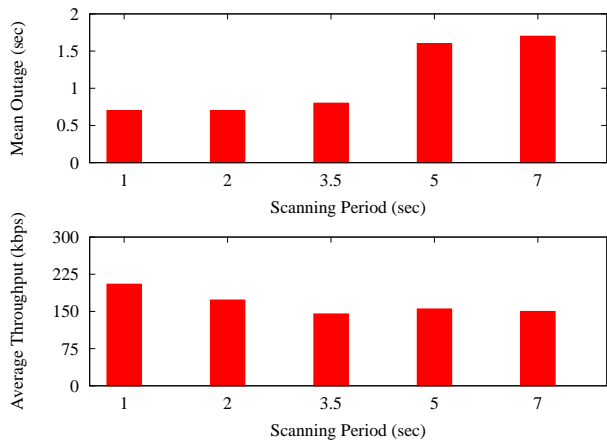


Figure 5: Impact of scanning frequency

Fast Scanning More than Halves the Outage Duration. We first evaluate the impact of the scanning period γ . Fig. 5 depicts the average duration of throughput outages and the average attained throughput for different values of the scanning period γ . While fast decision making yielded a throughput penalty, fast scanning has a positive effect provided that the values are adequately averaged. In particular, the figure shows that the duration of the outages monotonically decreases with scanning frequency. Thus, at vehicular speeds, channel variations are significant thereby requiring a short scanning period. Note that the Averaged with Hysteresis policy can outperform the two memoryless policies in terms of throughput outage duration: for low scanning periods ($\gamma = 1$ sec), the mean outage lasts for 0.7 seconds (Fig. 5). Recall that mean outage durations for Maintain Until Broken and Always Strongest Signal are 18 and 2.75 seconds respectively, while both of those policies also performed scanning with a period of 1 sec.

Time averaging and hysteresis yield shorter outages than Always Strongest Signal, as in the latter policy, clients initiate consecutive handoffs without any intermediate data

transmission. Contrary to the impact of frequent handoffs on attained throughput (Fig. 4), the average throughput in Fig. 5 indicates that in single channel deployments where the scanning duration is 20 ms, the cost of frequent scanning does not overwhelm the gain from an accurate estimation of the link conditions.

Finally, we find that throughput outages were minimized for $0.5 \leq \alpha \leq 0.8$.

The Average Throughput of the Averaged with Hysteresis Policy is 250% Greater than the Maintain Until Broken Policy and 40% Greater than the Always Strongest Signal Policy. To compare the throughput obtained by the Averaged with Hysteresis policy with the two memoryless policies, we consider three respective experiments for the reference loop, where the scanning period is 1 second for each policy under comparison.

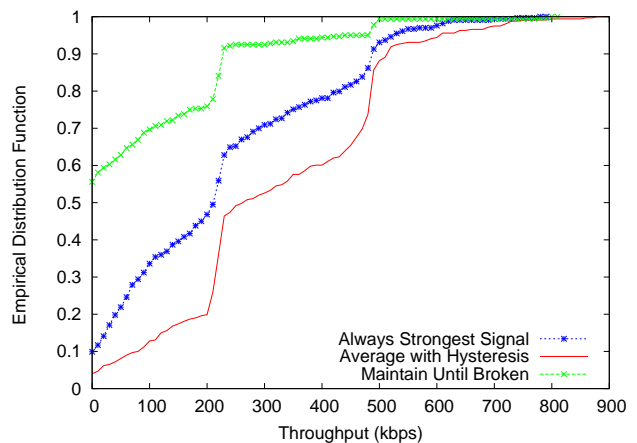


Figure 6: Throughput Comparison of Signal-Driven Policies

Fig. 6 depicts the Empirical Cumulative Distribution Function of the connection’s instantaneous reception rate. This can be shaped by rate-limiting,⁵ and due to asynchronous packet delivery can be also higher than transmission rate. As Fig. 6 demonstrates, Maintain Until Broken attains significantly lower throughput than the other two policies, a result that is due to its excessively long throughput outages that constitute more than half of the vehicle’s roaming time (55%). The Averaged with Hysteresis policy attains higher throughput than Always Strongest Signal, a result that is due to the throughput reduction that the excessive tendency for handoff is shown to induce (Fig. 4). We note that the Always Strongest Signal policy issued 3 times more association requests than Averaged with Hysteresis. As a result, this policy attains only 70% of the throughput attained in our experiments (Fig. 6) by the Averaged with Hysteresis policy.

5. EVALUATION OF DIFFERENTIAL CAPACITY HANDOFF

In this section, we evaluate the Differential Capacity Handoff with a set of experiments that (i) demonstrate the need for AP Quality Scoring by analyzing long-term differences

⁵Clients are rate-limited to 500 kbps for all APs, except for one at 250 kbps.

in client throughputs for different APs, (ii) characterize the weighting of channel conditions with AP quality scores for handoff decisions, and (iii) present and evaluate two AP quality-score metrics. Moreover, to provide perspective of the policy’s performance, we design experiments to approximate the maximum achievable throughput under a hypothetical idealized policy that always connects to the highest throughput AP and incurs no handoff delay.

5.1 Differential Client-AP Throughput

The use of per-AP quality scores in the Differential Capacity Handoff family addresses long-time-scale differences in each AP’s performance. To isolate this effect, we design an experiment in which AP weighting is turned off. In particular, we employ the Averaged with Hysteresis policy corresponding to $\delta = 0$. While driving the reference loop, the client attempts to transmit upload UDP traffic at 500 kbps, the same rate that the mesh network limits all clients to.

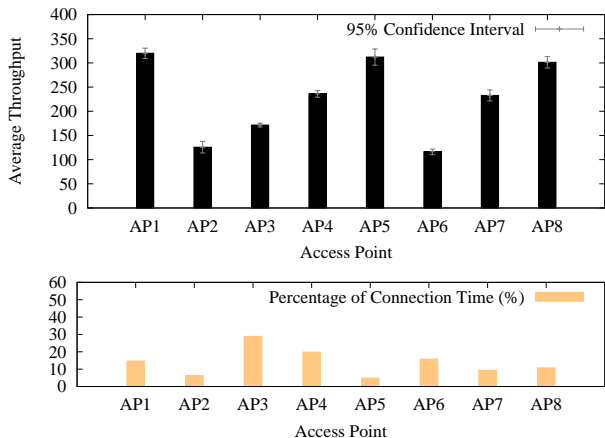


Figure 7: Average per AP Throughput

Long-term Disparity Exists among APs in Client Throughput. Fig. 7 depicts the average throughput while the client is associated with each Access Point chosen while driving the reference loop, along with the fraction of time associated with each AP. The figure indicates that certain APs yield throughput up to three times as high as others. The figure also depicts the 95% confidence intervals which yield a range of less than 50 kbps, despite the fact that the experiments consist of measurements performed at different hours over a 7 day period. Thus, for a particular AP, the variation in client throughput is minimal. However, among APs, significant differences in average throughput exist. Consequently, we conclude that APs have *long-term* disparities in the throughput that they can provide to clients. We verified this finding with additional experiments consisting of non-mobile measurements at locations that are as close as possible to each Access Point antenna. Despite the average reception rates being higher due to better link quality (eventually reaching the 500 kbps rate limit), the above conclusions still hold.

These throughput disparities arise due to the multi-hop multi-tier architecture: some APs are connected to directional antenna links, others to fiber gateways, and others must use multiple omni-directional 802.11 links to reach ei-

ther a fiber gateway or directional link. However, hop count to the gateway or the existence of directional links alone do not account for the disparity. For example, APs that are within the same hop distance from the gateway are observed to attain different throughput due to dissimilar inter-AP channel conditions (e.g., differing path loss exponents due to differing intervening foliage), differences in client rate limiting policies, etc.

Finally, we note that because the per-AP quality scores have little variation over time, these values can either be measured by the client and used repeatedly, or they can be provided by the network, e.g., via a one-time announcement at login or via other mechanisms such as infrequent beacons.

5.2 AP Quality-Scoring Policies

Here, we define and study two policies to assign AP quality scores to capture the disparity among AP performance capabilities.

The first policy is *gateway proximity weighting*. In this case, the quality score (weighting) of AP i is given by $w_i = 1 - \frac{h_i}{h_{max}}$, where h_i is the minimum distance in hop count between AP i and the gateway. We consider directional links that operate in a separate channels to yield a partial contribution to the hop-count metric h_i .⁶ The weighting is normalized to the maximum diameter of the network, h_{max} . In practice, the client can obtain this value either by direct advertisement from the network or via estimation from the network’s physical topology. For example, the Mountain View, CA topology has GPS coordinates of all nodes and gateways publicly available, from which clients can estimate the minimum hop distance from the gateway.⁷

The second AP quality score that we consider attempts to maximize client throughput by using per-AP measurements as described in Figure 7. In particular, we consider quality scoring given by $w_i = \frac{c_i - c_{min}}{c_{max} - c_{min}}$, where c_i is the average attained throughput through sampling from Access Point i and c_{max} and c_{min} are network wide client minimum and maximum rates. As with the previous policy, these weights can either be estimated by the client or advertised by the network. For the former case, clients can refine weights over time after repeated visiting of APs. For the latter case, the network can measure the required values and convey them to clients in occasional advertisements or at network login time. For example, for 10,000 Access Points that can correspond to a coverage area of 100’s of square miles, only 300 msec is required to advertise all quality scores for 1 byte scores transmitted at 250 kbps.


For gateway proximity weighting, Fig. 8(a) depicts the AP scoring transformations that are induced on the signal indicators for each Access Point. As described in Section 3.1, the transformation yields a three-segment piece-wise linear function. Fig. 8(b) provides the assigned quality scores, for both scoring policies.

5.3 Weighting AP Quality Scores and Channel Quality

A goal of coupling AP quality scores with signal-based indicators is to favor APs with high-quality backhaul capacity whenever the client has multiple association choices.

⁶In our experiments, we considered a scoring that accounted for directional links as $\frac{1}{5}$ of a hop.

⁷<http://wifi.google.com/city/mv/nodes.xml>



(a) Hop-count Weighting Transformation

AP	Gateway Proximity Weighting	Throughput-Sample Weighting ^a
AP1	1	1
AP2	0.5	0.2
AP3	0.9	0.4
AP4	0.4	0.6
AP5	0.4	0.9
AP6	0	0
AP7	0.5	0.6
AP8	0.5	0.9

^aA backhaul-tier modification resulting to a lower throughput for AP6, motivated a different sampling than Fig. 7.

(b) AP-Quality Scoring

Figure 8: Integration of AP Quality Scores and Channel Quality

However, excessively weighting the AP quality score could yield a policy that attempts to maintain associations to favored APs that have poor quality links or that are even out of transmission range. Here, we study the impact of the weighting of AP quality score vs. link quality. Thus, unlike the evaluation of purely signal-based decisions with $\delta = 0$ as considered in Section 4, here we consider $\delta > 0$.

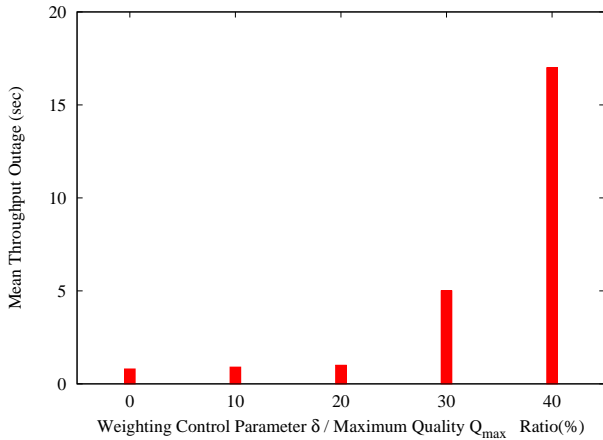


Figure 9: Throughput Outages for varying $\frac{\delta}{Q_{max}}$ ratio

AP Quality Scores Can be Weighted up to 20% without an Outage Penalty. Our experiments consists of multiple measurements taken on the reference loop, with each experiment employing a different weighting δ , and all experiments having a fixed set of AP quality scores. Fig. 9 depicts the average throughput outage duration as a function of the ratio of δ to the maximum link quality Q_{max} , with the measurement for $\delta = 0$ representing the Averaged with Hysteresis policy. The figure shows that for weights below 20%, the outage duration is near that of the Average with

Hysteresis policy. For weights larger than this range, the outage duration increases, eventually reaching levels similar to the Maintain Until Broken policy.

Thus, we conclude that a weight of less than 20% is needed by the Differential Capacity Handoff policy in order to maintain the low outages associated with signal-based policies. Advantages in client *throughput* for a weight of 20% as opposed to a weight of 0 (representing Averaged with Hysteresis) are evaluated below. Finally, we note that because this value is not known *a priori* by clients, in practice, clients can gradually increase their weighting as long as their throughput outages do not exceed the average duration incurred for $\delta = 0$ (Averaged with Hysteresis). For the remainder of our experiments, we use $\delta = 20\% \times Q_{max}$ as this value can weight the handoff decision based on the long-time-scale quality scores without increasing outages from those obtained with the Average with Hysteresis policy.

5.4 Throughput and AP Quality Scores

To evaluate the performance of these two AP quality-scoring techniques, we performed multiple experiments on the reference loop and depict the resulting average throughput in Fig. 10, along with Averaged with Hysteresis ($\delta = 0$) as a baseline.

The figure indicates that gateway proximity weighting yields a 20% increase in throughput as compared to Averaged with Hysteresis, improving throughput from 192 kbps to 229 kbps. Moreover, throughput sample weighting improves the throughput further to 240 kbps, yielding a 25% increase compared to Averaged with Hysteresis. Fig. 10 also provides a comparison of Differential Capacity Handoff against the signal-based policies with uncontrolled timescale of handoff initiation and indicates a throughput gain of 240%, as compared to the Maintain Until Broken policy.⁸

⁸As there was a change in network topology over the course of the project, we repeated some experiments under both topologies to allow comparison of all experimental results.




Figure 10: Throughput Gain

Thus, these experiments illustrate the throughput gains of each component of the Differential Capacity Handoff policy family. In particular, they characterize the relative importance of incorporating and weighting client information from signal strength to hop count to a gateway and finally to differences in AP throughput.

We also note that for such a gain in throughput to be established, associations to Access Points with highest quality score ($w_i = 1$) are increased by up to 5% of total association time and area, as compared to signal-based ($\delta = 0$) handoff. This is approximately equivalent to 125 m. and 18 sec. of extended association area and time respectively, values comparable to typical associations to Access Points. For the least weighted Access Points ($w_i = 0$), an analogous decrease (5%) of association time and area is observed. The difference with associations from signal-based handoff ($\delta = 0$) are regulated within the $\pm 5\%$ range, according to the value of the quality score w_i .

5.5 Maximum Achievable Throughput

To further evaluate the effectiveness of the family of Differential Capacity Handoff policies, we consider the following experiment to obtain an approximation for the maximum throughput that a user can achieve while driving the reference loop. This estimate will represent the throughput that a hypothetical client would obtain if always connecting to the maximum throughput AP and initiating a minimum number of handoffs, while roaming identically to the previous experiments.

Step 1. We first obtain the maximum throughput for each AP obtainable at the highest modulation rate. This is achieved by non-mobile measurements performed at each AP used within the reference loop. We select a location close to the AP to minimize degradation of the client-AP channel.

Step 2. Second, we estimate the maximum client throughput at each point of the reference loop with the following sub-steps: (i) we obtain a channel estimate for each point in the reference loop by obtaining the area’s “spectral footprint.” In particular, we measure channel conditions and modulation rates at a dense set of locations throughout the reference loop. Using GPS coordinates of the reference loop, we have the average signal quality of each point in the loop for each AP. (ii) Because the per-AP rates in Step 1 are

obtained via the maximum modulation rate (11 Mbps for 802.11b), we *scale* the throughput estimate at each location and to each AP according to an estimate of the modulation rate r , within $\{1,2,5,5,11\}$ Mbps. The rate is estimated according to the channel estimate of part (i) of Step 2. (iii) We scale the throughput for each location and each AP obtained in Step 1 by $r/11$ to reflect the effect of the reduced modulation rate. (iv) For each location, we consider that the client associates with the AP that provides the best throughput according to part (iii) of Step 2.

Step 3. Finally, we consider each point in the reference loop and client roaming similar to the experiments performed above as described by GPS traces. We compute a time average of the throughput, while also accounting for the minimum possible number of handoffs (8, one for each AP), and the average outage time during handoffs (0.8 sec, Sec. 6).

Under this methodology, we obtain a value of 295 kbps. Thus, 295 kbps represents an approximation of the *maximum* throughput that any policy can obtain, as it considers a hypothetical client that always associates with the highest throughput AP but incurs minimum handoff cost. This value compares favorably with those presented in Fig. 10 for the Differential Capacity Handoff family of policies.

5.6 Gain from Available Decisions

Finally, we design an experiment to isolate the differences among policies due solely to the handoff decisions. Hence, we isolate a subset of the performance data that includes only client spatial locations that are covered by a single Access Point by employing the signal data from the spectral footprint. All points with multiple APs in range are excluded from this baseline by considering the hypothetical client to obtain zero throughput at all of these locations. Thus, the baseline spatial locations do not involve an AP choice, thereby providing a minimum throughput for any handoff policy.

Using GPS coordinates, we correlate each of those locations with the respective time instants of the experiments and throughputs of Sec. 5.4. The average throughput for this hypothetical roaming to single-AP locations only is 98 kbps.

Subtracting this default value from the attained throughput in Sec. 5.4, quantifies the throughput gain due to *handoff decisions* to be 40% and 50% for gateway proximity weighting and throughput-sample weighting respectively, as compared to the Averaged with Hysteresis policy. Considering the Maintain Until Broken policy, subtracting the throughput attained at those locations increases the relative gain of our framework to 300%.

6. ORIGINS OF THROUGHPUT OUTAGE

In this section, we study the duration of throughput outages perceived by the receiver, when employing Differential Capacity Handoff. We perform measurements that quantify outages induced by *every* handoff policy and study the two integral parts of handoff: (i) the time interval during which the client defers from transmitting, dominantly comprised by the association delay and (ii) resumption of transmission to a *different* Access Point. We show that reception outages that are of higher duration than the association delay even occur in non-mobile conditions, under a scenario of forced association switching to different in-range Access

Points. Moreover, we perform experiments with Differential Capacity Handoff and show that it incurs only the minimal outages induced by association switching in *non-mobile* cases.

6.1 Delay of Successful Associations

Each handoff incurs an association delay, i.e., the delay between when a client initiates an association request to its selected access point and the time that the AP replies so that the client can transmit data.

To evaluate this factor, we measure the association delay of successful attempts. We also measure the percentage of failed association attempts, i.e., instances in which an association request was not followed by a response. These experiments consist of traveling the reference loop with handoff initiated according to the Differential Capacity policy and $\delta = 20\%$ and the signal-based policy with $\delta = 0$. As a baseline for comparison, we also measure association delays for an indoor wireless LAN.

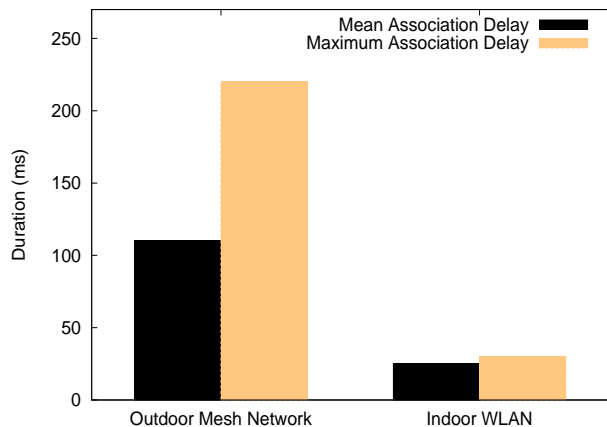


Figure 11: Association Statistics in Indoor and Outdoor Environment

Fig. 11 depicts the average and maximum association delay for the reference loop and the wireless LAN. The association delay in the outdoor environment has a mean of 110 ms, with values ranging up to 220 ms. However, we note that in the experiments, the percentage of failed association requests is 40% on average. Each failure also incurs an inter-association-request delay of one second so that the total delay incurred can be much larger than 220 ms. We refer to this total delay that incorporates both failed association attempts (if any) and the final successful association as an *association hole* and study it further below.

Multi-Hop Wireless Backhaul Induces an Order-of-Magnitude Increase in Successful Association Time vs. WLAN. For comparison, in a Rice indoor WLAN, the association delay is always observed to lie within a small range around 25 ms. Furthermore, association attempts are always successful. This disparity between WLAN and mesh arises because transmissions occurring on the wireless backhaul constitute traffic that is absent in WLAN deployments; such traffic can incur delays, collisions, etc. Furthermore, most locations in the indoor WLAN are under better signal coverage, as the deployment is denser with 27 Access Points

within a building whose area is significantly smaller than the footprints in the TFA network.

6.2 Outage and Association Holes

The experiments above presented association delays only for successful associations and considered client-side measurements. Here, we consider the total outage time consisting of both the successful association delay and the duration of failed attempts, if any. Moreover, we present receiver-side measurements which we will show can mask or exacerbate the outage duration, depending on the quality of the old and new paths.

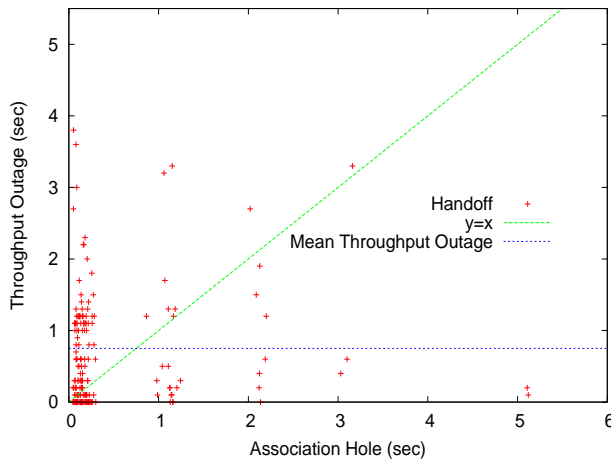


Figure 12: Association Holes / Throughput Outage Relation

We return to the reference loop for 200 handoffs while using the Differential Capacity Handoff policy with $\delta = 20\% \times Q_{max}$. Fig. 12 depicts the result of the experiments. Each handoff yields an association hole (total duration required for the client to successfully associate, including the time required for failed associations if any) and a throughput outage (total duration of zero receiver throughput). Each point on the graph represents such an x-y pair of association hole and outage. Moreover, a horizontal line depicts the mean duration of throughput outage for the 200 handoffs, 750 ms. A 45 degree line of reference is also depicted and labeled “ $y = x$ ” as discussed below.

Outage Durations Can Be Less Than Association Delays. A striking aspect of the figure is that the data points do not lie above the 45 degree line. In particular, recall that the total outage time as seen by a receiver consists of three components: the time required for a successful association, the time required for association failures (possibly zero), and an additional network component. The figure indicates that the network component can in some cases *reduce* the throughput outage as characterized by points below the 45 degree line. For these points, throughput outages are *less* than the association delay due to the network infrastructure: When the client hands off from the first to second AP, traffic is already in flight from the first AP at (potentially) multiple nodes along the backhaul path. Even while there is a client-side outage due to (for example) failed association attempts, packets may still be arriving from such packets queued during the old association. If the new as-

sociation yields a higher performing path (e.g., due to low queueing delay or low path length), packets may quickly arrive on the new path, thereby partially *masking* association outages from the receiver. On the other hand, many points also lie above the 45 degree line. In these cases, the converse occurs and network conditions increase the total duration of the throughput outage.

Outage Duration under the Differential Capacity Policy is Similar to the Non-Mobile Case. As a baseline for comparison, we next study outage durations incurred by non-mobile clients having forced handoffs, thereby eliminating mobility effects and eliminating policy-dependent effects.

To achieve this objective, we design an experiment in which a *non-mobile* client is placed within range of two Access Points for 10 minutes. Every 10 seconds, the client initiates a handoff between the two Access Points. A total number of 4 such experiments were performed for 4 different locations and pairs of Access Points. The experiment yielded an average duration of all throughput outages of 820 ms. As this is quite close to 750 ms obtained in the mobile case, we conclude that the Differential Capacity policy yields close to ideal outage durations.

The Majority of Long-Duration Handoffs are Due to Association Failure and Asynchronous Delivery among APs. Finally, we note that different applications will have differing tolerance levels for outage durations. For example, voice requires 200 ms for high quality voice and 400 ms for low voice quality voice.⁹ For mesh handoffs, additional delays in this range can only be achieved if association attempts are successful so that throughput outages are not magnified, and network delays are not excessive. Unfortunately, in an urban mesh environment, association attempts fail with a percentage that can reach 40%; with the addition of network factors as discussed above, exceeding such thresholds is inevitable. In our experiments, handoffs that yielded an outage exceeding 400 msec can be classified as follows: 29% exceeded due to incurring at least one association failure and 26% exceeded due to the asynchronous delivery of different streams of traffic from the old to new association.

7. RELATED WORK

7.1 Cellular Networks

Voice and 3G cellular networks can support vehicular speeds of client devices. However, the architecture of cellular networks differs from mesh networks in two critical ways: (i) mesh networks employ random access at up to three tiers whereas 3G networks employ scheduled access and typically employ only a single wireless tier and (ii) cellular networks employ signaling protocols that convey uplink and downlink quality information that is exploited by the *network infrastructure* for handoff initiation, whereas our approach employs *client-driven* handoff. See [1] for a survey of handoff techniques in cellular systems.

⁹Source: International Telecommunication Union, <http://www.itu.int>

7.2 Non-Mesh 802.11 Networks

7.2.1 Vehicular Mobility in Outdoor Networks

Traffic exchange under vehicular speeds in 802.11 networks was performed in [8, 9, 10, 11]. In those papers, data was transmitted between cars, or between a car and static Access Points. “Gray periods” in connectivity are recognized in [12] as a combined result of the variability in the urban radio environment and the vehicle’s traversal from regions under poor coverage. In [13], vehicular clients connect to open-access residential wireless routers of Boston, MA. Reference [14] makes use of directional antennas for maintaining high throughput physical-layer connections as the vehicle moves.

In contrast to our work, none of the papers above considered a multi-hop or multi-tier architecture, as all transmissions were destined to Access Points with wired access. Moreover with the exception of [14], they don’t propose novel handoff techniques.

7.2.2 Handoff in Indoor WLANs under Pedestrian Mobility

As no standard handoff protocol is specified by the 802.11 protocol, in practice wireless interfaces follow manufacturer dependent policies. In [2, 3, 4, 15], these policies are shown to be simplistic and yield substantial degradation of the link quality until handoff initiation. Empirical studies employing wireless sniffers reverse engineered the handoff behavior of network interfaces yielding a sequence of scanning, flushing, authentication and association phases [16, 17]. Velayos [15] targets reducing the AP discovery duration, with the introduction of a signal quality threshold. Mathre [3] introduced filtering-based techniques for handoff initiation. Wu [4] reduces the AP-discovery duration as scanning actions are performed proactively, interleaving with standard traffic. Hence, prior research focused on reduction of handoff duration, primarily by addressing the problem of handoff initiation and Access Point discovery.

Our work contrasts in the network architecture (mesh vs. WLAN), our consideration of vehicular mobility, and our use of long-time-scale AP quality scores.

7.3 Network Assisted Handoff

Ramani proposes SyncScan, an algorithm that exploits synchronization of the Access Points in a WLAN and transmits beacon frames at known instants on different channels [2]. A synchronization-aware mobile station monitors AP signals with minimal resource consumption, hence reducing the scanning overhead. *SMesh* defines a network architecture and a set of protocols for a *mesh* network [18]. The protocols target fast handoff of mobile stations as Access Points are responsible for monitoring the link quality of the client. Association decisions for clients are derived with inter-AP message exchange.

Our work differs in that we consider vehicular vs. pedestrian mobility and client- vs. network-driven handoff. In our work, clients can estimate the AP weight parameters directly; however, if they decline to do so, network assistance can be employed in a one-time access to the set of weights, information that can be exchanged at client log-in time only.

In contrast to all of the aforementioned work, our paper is the first to consider vehicular mobility in a mesh network having a scalable multi-tier architecture. Such an archi-

ecture targets serving a large user population over a large coverage area.

8. CONCLUSIONS

We addressed vehicular mobility in large-scale multi-tier multi-hop mesh networks. We designed a family of techniques that employ smoothed AP-client signal quality coupled with per-AP quality scores. The AP quality scores characterize the inherent inability of the mesh architecture to provide uniform bandwidth to all spatial locations. We performed an extensive set of experiments on an operational testbed covering over 3 km^2 . We designed numerous experiments to isolate the performance factors that control a handoff policy's performance. We find that for vehicular clients, the Differential Capacity family of handoff policies provides performance close to experimentally obtained ideal values in terms of both hand-off outage durations and average throughput.

9. REFERENCES

- [1] D. Wong and T. Lim, "Soft handoffs in CDMA mobile systems," *IEEE Personal Communications Magazine*, vol. 4, no. 6, pp. 6–17, Dec 1997.
- [2] I. Ramani and S. Savage, "Syncscan: practical fast handoff for 802.11 infrastructure networks." in *Proc. IEEE INFOCOM*, Miami, FL, Mar. 2005.
- [3] V. Mhatre and K. Papagiannaki, "Using smart triggers for improved user performance in 802.11 wireless networks," in *Proc. ACM MobiSys*, Uppsala, Sweden, June 2006.
- [4] H. Wu, K. Tan, Y. Zhang, and Q. Zhang, "Proactive scan: Fast handoff with smart triggers for 802.11 wireless LAN," in *Proc. IEEE INFOCOM*, Anchorage, AK, May 2007.
- [5] J. Camp, J. Robinson, C. Steger, and E. Knightly, "Measurement driven deployment of a two-tier urban mesh access network," in *Proc. ACM MobiSys*, Uppsala, Sweden, June 2006.
- [6] Earthlink, personal communication with network operators, June 2007.
- [7] A. Haeberlen, E. Flannery, A. M. Ladd, A. Rudys, D. S. Wallach, and L. E. Kavraki, "Practical robust localization over large-scale 802.11 wireless networks," in *Proc. ACM MobiCom*, Philadelphia, PA, USA, Sept. 2004.
- [8] P. Bergamo, M. Cesana, D. Maniezzo, G. Pau, K. Yao, D. Whiteman, and M. Gerla, "IEEE 802.11 wireless network under aggressive mobility scenario," in *Proc. International Teletraffic Congress (ITC)*, Las Vegas, NV, Oct. 2003.
- [9] R. Gass, J. Scott, and C. Diot, "Measurements of in-motion 802.11 networking," in *Proc. of 7th IEEE Workshop on Mobile Computing System and Applications (HOTMOBILE)*, Washington, DC, Apr. 2006.
- [10] D. Hadaller, S. Keshav, T. Brecht, and S. Agarwal, "Vehicular opportunistic communication under the microscope," in *Proc. ACM MobiSys*, San Juan, Puerto Rico, June 2007.
- [11] J. Ott and D. Kutscher, "Drive-thru Internet: IEEE 802.11b for automobile users," in *Proc. IEEE INFOCOM*, Hong Kong, China, 2004.
- [12] R. Mahajan, J. Zahorjan, and B. Zill, "Understanding WiFi-based connectivity from moving vehicles," in *Proc. Internet Measurement Conference (IMC)*, San Diego, CA, Oct. 2007.
- [13] V. Bychkovsky, B. Hull, A. Miu, H. Balakrishnan, and S. Madden, "A measurement study of vehicular internet access using in situ Wi-Fi networks," in *Proc. ACM MobiCom*, Los Angeles, CA, Sept. 2006.
- [14] V. Navda, A. P. Subramanian, K. Dhanasekaran, A. Timm-Giel, and S. Das, "Mobiteer: using steerable beam directional antenna for vehicular network access," in *Proc. ACM MobiSys*, San Juan, Puerto Rico, June 2007.
- [15] H. Velayos and G. Karlsson, "Techniques to reduce IEEE 802.11b handoff time," in *Proc. International Conference on Computer Communications (ICC)*, Paris, France, June 2004.
- [16] A. Mishra, M. Shin, and W. Arbaugh, "An empirical analysis of the IEEE 802.11 MAC layer handoff process," *SIGCOMM Computer Communication Review*, vol. 33, no. 2, pp. 93–102, 2003.
- [17] J.-O. Vatn, "An experimental study of IEEE 802.11b handover performance and its effect on voice traffic," <http://www.it.kth.se/~vatn/research/handover-perf.pdf>.
- [18] Y. Amir, C. Danilov, M. Hilsdale, R. Musăloiu-Elefteri, and N. Rivera, "Fast handoff for seamless wireless mesh networks," in *Proc. ACM MobiSys*, Uppsala, Sweden, June 2006.