

Mark-and-Sweep: Getting the “Inside” Scoop on Neighborhood Networks

Dongsu Han
Carnegie Mellon University
dongsuh@cs.cmu.edu

Michael Kaminsky
Intel Research Pittsburgh
michael.e.kaminsky@intel.com

Aditiya Agarwala
Carnegie Mellon University
adityaag@andrew.cmu.edu

Konstantina Papagiannaki
Intel Research Pittsburgh
dina.papagiannaki@intel.com

David G. Andersen
Carnegie Mellon University
dga@cs.cmu.edu

Srinivasan Seshan
Carnegie Mellon University
srini@cs.cmu.edu

ABSTRACT

Residential Internet connectivity is growing at a phenomenal rate. A number of recent studies have attempted to characterize this connectivity—measuring coverage and performance of last-mile broadband links—from a various vantage points on the Internet, via wireless APs, and even with user cooperation. These studies, however, sacrifice accuracy or require substantial human time. In this work, we present a novel two-pass method to characterize neighborhood networks. We demonstrate that the two pass method dramatically reduces the time spent in active measurement while retaining accuracy. A case study on two neighborhoods in Pittsburgh provide new and accurate insights into broadband connectivity, including throughput, broadband coverage (DSL vs. cable vs. fiber), NAT configurations, DHCP, DNS usage. The results further characterize 802.11 connectivity in the neighborhood.

Categories and Subject Descriptors

C.2.5 [Computer-Communication Networks]: Local and Wide-Area Networks—*Internet*

General Terms

Measurement

Keywords

Broadband connection, access point, access network, measurement tool

1. INTRODUCTION

Over the past few years, residential network connectivity has undergone a number of dramatic changes. Residences have moved from using dialup to using a broad range of wireless technology within the home and a mix of broadband DSL,

cable modem and fiber-based technologies to connect homes to the Internet. The result of the increase in broadband connectivity is that residential Internet use has far more impact on the Internet than it did in the past.

Unfortunately, there is relatively little work that characterizes typical network connectivity to the home user. Most existing measurement studies, datasets and tools target the core of the Internet. A few recent studies have attempted to characterize residential broadband connectivity. These studies have used three types of measurements: Internet-based [6], home-user driven [9], and wireless access point (AP) based measurements [8]. Internet-based measurement studies, which probe residential links from remote Internet locations, have been the largest in scale. However, these studies suffer from measurement noise and inaccuracy since the measurements are performed far from the target links. While studies that rely on users to perform tests at their homes greatly improve accuracy, they require significant user participation to collect results. Finally, measuring connectivity using wireless APs also provides accurate characterization of connectivity but requires significant time from the individuals performing the measurements. This tends to limit the scale of such studies.

In this paper, we adopt the wireless AP-based approach to measuring residential connectivity. Unlike the other two approaches, AP-based measurement allows us to focus on neighborhood-level connectivity and answer questions about how residential connectivity varies within a small geographic area. In addition to the measurement study, one of the key contributions of this paper is a set of measurement tools and methods, called *Mark-and-Sweep*, that ensures the accurate and time efficient collection of wireless-AP based measurements in a neighborhood (Section 2). We achieve these improvements by dividing the measurement task into a two stage process. In the first stage, we drive around the neighborhood and passively record all transmissions we observe along with GPS coordinates. This gives us a map of the APs within a neighborhood and key properties (e.g. SSID, security settings, channel, signal strength at different locations) of the APs. Using this data, we prune out APs with low signal strength and identify ideal locations to perform active measurements for the rest of the APs. In the second stage, we drive to the chosen locations and perform detailed active probing of the APs to measure properties of their wireless network and last-mile connectivity.

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.

IMC'08, October 20–22, 2008, Vouliagmeni, Greece.

Copyright 2008 ACM 978-1-60558-334-1/08/10 ...\$5.00.

We have performed this two-pass measurement of residential connectivity in a densely populated neighborhood near Carnegie Mellon University and a more sparsely populated suburban neighborhood in Pittsburgh. In this paper, we analyze these measurements to evaluate the benefits of *Mark-and-Sweep* (Section 3) and to answer key questions about residential connectivity (Section 4). We compare the measurement performance of *Mark-and-Sweep* with the traditional approach of stopping every few homes (approximately every 75 ft) and performing a set of active measurements. Our results show that *Mark-and-Sweep* identifies just as many APs and provides similar or better accuracy for its active probes as the traditional approach; more importantly, *Mark-and-Sweep* obtained these measurements almost six times faster.

We use the above measurements to consider how effectively an ISP could provide roaming wireless connectivity in a neighborhood and how wireless connectivity could be used to allow users to access wired last-mile links in other homes. While some of the observations are obviously specific to the neighborhoods measured, we believe that our results highlight useful measurement and analysis methods for answering such questions. Our tools and data can also be used to evaluate a wide range of questions regarding the design of access networks and the feasibility of neighborhood-wide applications or networks. *Mark-and-Sweep* and the data presented in this paper are publicly available.¹

2. MARK-AND-SWEEP

Our method for collecting network information focuses on achieving two goals: (1) ensuring an accurate set of collected data; (2) maximizing the amount of data collected for the amount of human time invested in taking measurements.

Our experience, and that of others [5, 8], demonstrates that the largest percentage of time associated with obtaining detailed measurements collected through wireless networks is in the scanning, per-AP association, DHCP, and subsequent broadband tests. These stages of the measurement process are particularly slow when measurements are taken in an area of poor signal strength. In addition, poor signal strength can degrade measurements of broadband performance (the wireless medium may be too slow or lossy to permit accurate measurement).

To address this potential inefficiency, we use a two-pass measurement scheme. In the first pass, we drive through an area without stopping, collecting extensive passive measurements of all APs. We do not associate with APs during this pass. We use these measurements to create a plan for the second pass, in which we measure each access point and its associated network in-depth exactly once at the approximate location where its signal strength was the strongest.

Pass 1: Access point identification and location. The goal of the first pass is to determine which APs are in an area, approximately where those APs are located, and where their signal strength is the strongest. This information has been the focus of a number of war-driving efforts (see Wiggle, WiFiMaps, PlaceLab), and data from previously war-driven areas could be used as input to our tool.

Movement pattern: During these runs, we scan all streets in the targeted area at least twice (usually once per side)

using *kismet*. Since we are only passively recording data, we can cover a given geographic area rapidly and drive at approximately 20 mph during this pass. Multiple passes through the same location ensure that we identify all APs that have a reasonably strong signal and are transmitting packets including data and beacons.

Measurements collected: For each AP detected, we record its bssid, essid, channel, encryption methods, and supported rates. In addition, we record the signal strength of every packet heard, and the GPS coordinates of the observation.

Between passes: Access point preprocessing. Based on the collected information, we determine which access points to measure in depth and the optimal locations for such measurements. The preprocessing has two parts:

AP Pruning: The list of APs is pruned to include only unencrypted APs with a Signal-to-Noise Ratio (SNR) of 20 dB or larger. In Section 3, we offer justification for this threshold. This pruning eliminates weak APs where the measurements either 1) are likely to fail; or 2) are likely to be wireless-bottlenecked to the extent that they do not provide insight on the actual residential broadband link.

Path Planning: Given the list of APs to measure and the map of all AP observations for a neighborhood, we identify the location where each AP’s signal was the strongest. We then determine an order in which to visit these locations that minimizes driving time. Currently, path-planning is done manually; a future version of the tool will use path planning software to suggest an efficient route to visit the APs.

Pass 2: In-depth network characterization. The second pass of war driving collects detailed wireless and wired network measurements from targeted access points. We use a combination of *kismet* (for discovery) and *wicrawl*² (as an infrastructure to run individual tests on each AP).

Movement pattern: When taking measurements during this pass, we use “war-parking” instead of war-driving. The driver stops at the locations specified during the preprocessing phase and searches for access points in range for testing. We use the *Navit* open source mapping software to plot the testing locations and the current location of the testing vehicle. A custom application integrates the mapping, location checking, and filtering of the specific access points to be tested. This application informs the driver when to stop to conduct the tests, starts a filtered version of *wicrawl* specific to access points to be tested at that location, and finally updates the map and removes access points that have been tested successfully.

Measurements collected for each AP: The tool first attempts to connect to each access point. Next, it sends five DHCP requests in a row to quickly obtain an IP address, using the first response it receives. The tool records the entire DHCP response for later analysis. If the association or DHCP attempts fail, the tool prompts the user to reposition the vehicle and attempt the test once more.

Once it has an IP address on the target wireless LAN, the tool performs five tests:

1. *Ping.* Ping a test server in the CMCL lab at Carnegie Mellon University.
2. *Port Availability.* Attempt to open a TCP connection to the test server on five ports: 25 (SMTP), 80 (HTTP),

¹The tool and data are available at <http://www.cs.cmu.edu/~dongshuh/Mark-and-Sweep/>

²*wicrawl* is an open source tool available at <http://midnightresearch.com/projects/wicrawl>

443 (HTTPS), 587 (authenticated SMTP) and 56123 (a high-numbered port).

3. *Traceroute*. Run traceroute to the test server.
4. *NAT*. Use the open source STUN client/server to determine the NAT type.
5. *Throughput*. Send UDP packets at 15 Mbps for 4 seconds to measure upstream and downstream bandwidth to the test server (if the success rate exceeds 90% we increase the transmission rate). We use a version of *nuttcp* that we modified to work through NATs.

While collecting these measurements, we simultaneously use *tcpdump* to record a packet dump of all sent/received traffic on the laptop’s second wireless interface (running in monitor mode). We use this packet trace to help determine which measurements are affected by poor quality in the wireless medium. At the end of this process, we parse the collected data files (*kismet* output and GPS log from pass 1; association, GPS and experiment log from pass 2; and packet dump from the monitor NIC in pass 2) and insert the results in a database for further analysis.

3. METHOD EVALUATION

Our two-pass method has two potential benefits: 1) reducing the total amount of time required to characterize a neighborhood network, and 2) improving the accuracy of the collected measurements, compared to alternate approaches. Time savings come from careful planning that allows us to stop and collect detailed measurements only once for each AP. Accuracy improvements (particularly for performance metrics such as latency and throughput) come from choosing a single measurement location with a high-quality wireless link, thereby removing, as much as possible, any wireless medium bottlenecks.

To validate the two benefits of our method, we compared *Mark-and-Sweep* to two alternate measurement schemes in a smaller geographical area with 40 unencrypted APs. The alternate schemes work as follows: We drive through the given area, stopping to collect measurements every 75 ft. At each stop, we discover all APs, associate with each one and collect all measurements of *Mark-and-Sweep*’s pass 2. We then define two measurement schemes: one that performs detailed active measurements only once per AP at the time that DHCP first succeeds (*Measure-First*) and one that performs detailed active measurements for all APs at each stop (*Measure-All*). *Measure-First* aims to reduce the time taken by performing active measurements as soon as possible and only once; *Measure-All* aims to improve accuracy by measuring the same AP multiple times and reporting the best value.

We compare the total amount of time spent in active measurement, number of APs that were successfully measured and UDP throughput for each scheme in Table 1. We report the average UDP throughput for the APs whose maximum SNR is greater than 20dB to provide a fair comparison, and show the UDP throughput distribution for all APs from *Measure-All* in Figure 1. The measured throughput values for *Mark-and-Sweep* are as accurate as *Measure-All*’s but are collected almost 6 times faster. The faster *Measure-First* is 2.7 times slower than *Mark-and-Sweep* and reports throughput measurements that rarely exceed 2 Mbps. These results attest to the importance of measuring at the location of the best wireless signal. The APs that were pruned by *Mark-and-Sweep* for pass 2 provide very low throughput (Figure 1)

Methods	Time (s)	# APs	Avg xput
Measure-All	3885	15	3.3 Mbps
Measure-First	1814	15	1.3 Mbps
Measure-All(Thresh)	1099	10	3.6 Mbps
Mark-and-Sweep	656	11	3.4 Mbps

Table 1: Comparison with alternative methods

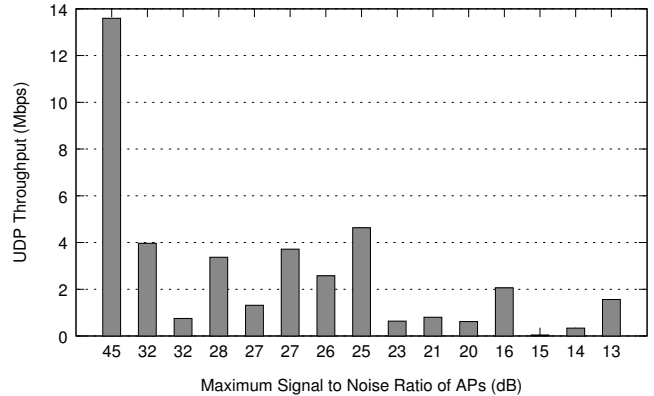


Figure 1: UDP throughput versus Maximum SNR of AP

and will always be wireless-bottlenecked even at the best location. We prune 4 APs, a relatively high fraction of APs, in this test area because the test area is small and we observe many weak signals of APs that are located outside the test area. Later in this section, we report the actual number of DHCP-able APs we might have missed because of the threshold.

Mark-and-Sweep uses a 20 dB threshold to decide whether to collect detailed measurements from an AP. A similar criterion could be used by *Measure-All*. We report on the results of such a scheme, *Measure-All(Thresh)*, in Table 1.

To further validate the choice of 20 dB as a cutoff threshold, we quantify the number of plausible measurement points we lose due to such thresholding in a neighborhood where we actually performed the measurement (SQ in Section 4). Given the poor quality of the wireless link at less than 20 dB SNR values, DHCP leases rarely succeed. The total DHCP successes from APs that had maximum SNR value less than 20dB was only 9 out of 181 such APs.

4. RESULTS

Mark-and-Sweep enables us to gain an accurate view of neighborhood networks today. Below, we provide a sample of the findings one can derive using the collected measurements (*Mark-and-Sweep* can be easily extended to include more measurements). We measured two neighborhoods in Pittsburgh: Squirrel Hill (denoted SQ) and part of Ross and McCandless Township (RMT) in suburban Pittsburgh.³ While both neighborhoods are mainly residential, the Squirrel Hill area is more densely packed with detached homes, townhouses and small apartments. The areas we wardrove span 1.3 sq. km for Squirrel Hill and 3 sq. km for Ross and McCandless Township.

³These results cover 802.11b/g only.

	SQ	RMT
Total APs	1200	965
Unencrypted APs	354 (29.5%)	302 (31.3%)
2nd pass APs	173	184
Association succeeded	156	178
DHCP succeeded	89	126
Internet available	80	115

Table 2: Summary statistics of our trace data.

Vendor	# of APs seen	% encrypted
Linksys	977	64.48
Actiontec Electronics	383	97.91
Netgear	264	76.14
AboCom Systems	249	77.51
D-Link	232	55.17
Apple	161	71.43
Belkin	112	67.86
Cisco	81	54.32
Agere Systems	38	5.26

Table 3: AP encryption ratio by vendor (Top 9 vendors shown).

4.1 A View into Wireless Neighborhoods

Neighborhoods feature a large number of unencrypted APs. The first pass of *Mark-and-Sweep* revealed 1200 APs for SQ and 965 APs for RMT (Table 2). 30% were unencrypted. A similar study based on pass 1 measurements was carried out in Pittsburgh in 2005 [3]; over two years, the percentage of encrypted APs increased from 50% to 70%. We believe this is partly due to increased awareness in security and privacy as well as vendors shipping APs with security settings. However, we have yet to see more sophisticated usage of APs that allows users to have both encrypted and unencrypted APs.⁴ Such usage pattern may stabilize or even reverse the trend of increasing percentage of encrypted APs. The APs for a part of RMT are shown in Figure 2(a). The right-hand figure (b) shows ISPs associated to open and Internet available APs.

Vendor/ISP partnerships influence neighborhood security. Using the MAC address of the discovered APs we computed the fraction of vendor APs that were encrypted. Our results are shown in Table 3. With two exceptions, the fraction of encrypted APs does not correlate strongly with AP vendor; approximately 60%-70% of the APs were encrypted regardless of the vendor. The two exceptions were Agere systems APs, which were rarely encrypted, and Actiontec Electronics, whose APs were almost always encrypted. This finding was intriguing, since other APs also ship with encryption by default but do not achieve the same fraction once deployed (e.g., only 71% of Apple APs were encrypted). One key difference is that Verizon ships pre-configured Actiontec APs to their customers, which appears to reduce the chance that the customer might disable encryption during configuration or debugging performance or connectivity problems.

⁴APs from Meraki and British Telecom allow to have encrypted and non-encrypted APs at the same time.

ISP	SQ	RMT
Comcast	49	87
Verizon DSL	24	10
Verizon FiOS	-	18
aspStation	2	-
Covad	2	-
Nauticom	1	-
Speakeasy	1	-
Full Service Computing	1	-
Total	80	105

Table 4: ISP distribution of open Access Points

Open APs cover up to 98% of the urban area. Using the measurements collected in pass 1 we estimate each AP’s coverage range by computing the distance between the location of the AP’s strongest signal measurement and the farthest location where the AP was heard at a SNR over 20 dB. The maximum ranges of APs for SQ and RMT were 327 m and 507 m and median values were 152 m and 80 m respectively. Using the individual coverage areas computed for each AP⁵, we then compute the fraction of the wardriven area with 802.11 connectivity to Internet available open APs. Surprisingly, 98% of SQ had access to an open AP, as did 48% of RMT. As can be seen in Figure 2(a), the discovered (but not necessarily open) APs provide 100% geographic coverage.

802.11n has started penetrating neighborhoods. A small percentage of the APs supported 802.11n (which has a claimed effective throughput up to 130 Mbps). While IEEE 802.11n is still a draft standard, up to 6% of the APs in SQ, and 2.2% in RMT, support it.

4.2 Last-mile Internet Connection

Table 2 shows that while both neighborhoods had almost 300 unencrypted APs, the number of APs tested in pass 2 was much smaller. The difference arises from selecting only APs whose SNR is over 20 dB. AP association and DHCP failures further reduce the number of locations where we can collect detailed broadband measurements to 195.

Residential network speeds are increasing. Table 4 lists the ISPs accessible via open APs in the two neighborhoods. We use the domain name of the external IP address to identify the ISP. We then classify APs by their ISP and by the type of technology they use (e.g., cable, DSL, or fiber).

We notice that neighborhood networks are not constrained to slow speed cable and DSL connections, but may feature high capacity fiber lines. In addition, RMT appears to feature a much higher number of cable links than DSL links, which is probably an artifact of the range limitations of DSL.

Cable throughput is higher and more variable than DSL throughput. Using *nuttcp* we compute the maximum instantaneous UDP throughput sustained through the broadband link for one second during a 4 second measurement duration. (We chose 4 seconds to balance measurement quality with the amount of time and intrusiveness of the measurement.) To reliably estimate the *broadband* capacity, however, requires that the measurements are not artificially

⁵We trim very large range numbers to the median observed range to avoid over-estimating coverage.

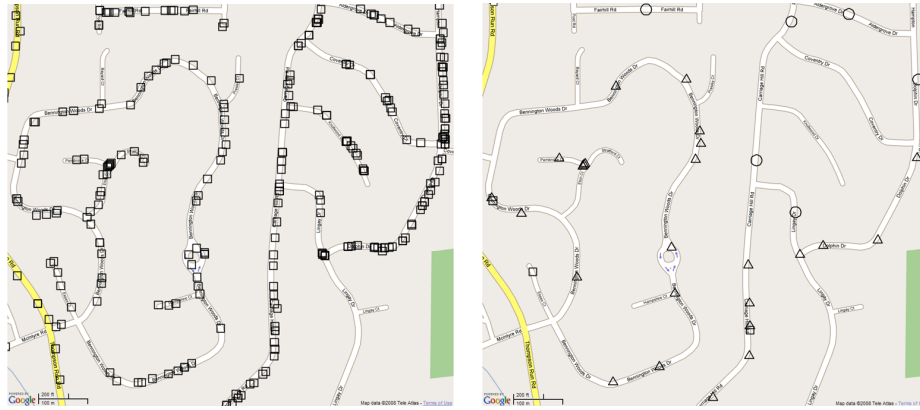


Figure 2: (a) Locations of all discovered APs from pass 1 in RMT. (b) Locations of open APs and associated ISPs. Comcast:Triangle, Verizon DSL:Square, Verizon FiOS:Circle

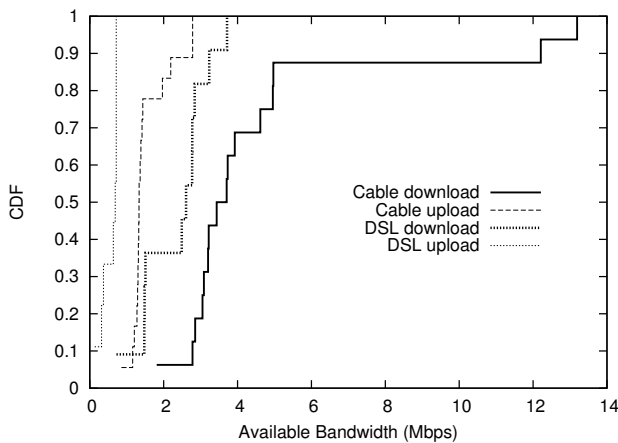


Figure 3: Broadband link bandwidth estimation

limited by the wireless link. We ensure this by using the packet trace collected on the monitor interface to compute the effective wireless transmission rate. We compute the amount of time required for the transmission of each packet at the reported MAC layer tx-rate and remove retransmissions. If the computed wireless rate is over 10 Mbps, we assume the measurements are wired-bottlenecked. If the number is between 5 and 10 Mbps, we compare it with the measured UDP throughput. If the throughput is more than half the average transmission rate, then we assume that wireless is not the bottleneck and accept the estimated bandwidth number. For all other cases, we assume the collected measurements are biased by the wireless quality. Figure 3 shows upload and download bandwidth for DSL and cable for SQ.

As expected, cable throughput speeds appear higher than those of DSL, and upload speeds are much smaller than download. The incumbent cable provider, Comcast, offers 6 and 12 Mbps connectivity, which is significantly higher than the 3 Mbps offered by the incumbent DSL provider. Identifying the available rates is harder for cable, which we believe may be due to cable modems sharing more throughput with neighbors and, as a result, actually offering lower throughput than advertised [6]. On the other hand, Comcast’s “12 Mbps for the first ten megabytes” *Powerboost* service shows up

clearly at the right of the figure. More detailed and longer term measurements will be required to discern those two effects. Having said that, *Mark-and-Sweep* could use any throughput estimation technique in order to obtain a more accurate view.

4.3 Home Network

Wireless APs often also serve as the hub of the residential network, in many cases also acting as NATs that gateway to the Internet. We expect that wireless access points will play a greater role in in-home networks as an increasing number of devices, such as printers, PDAs, storage devices, and even HDTVs, become 802.11 enabled. In this section, we take a deeper look into home networks by looking at two specific configuration parameters, regarding DNS and NAT.

Most home users do not change their ISP provided DNS configuration. During pass 2, our tool records DHCP lease information. 47% of the 215 DHCP-able access points advertised DNS resolvers only in the local private address space. 48% of them were using only remote/public DNS, and 5% used both local and remote DNS. 99% of the remote DNS servers were provided by the direct upstream ISP. We further examined whether these DNS servers were located in close proximity to the wireless network using ip2geo [1] which maps IPs to geographical area. Surprisingly, only 2.5% of the remote DNS servers were located far away (e.g., Georgia, Washington). 97.5% of the servers were located in Pittsburgh, Virginia, or New Jersey. Previous studies have shown that an appreciable fraction of computers use quite remote DNS servers [7]. Many content distribution networks (e.g., Akamai) use the source address of a DNS query to direct users to a nearby content replica. Our results suggest that these techniques work particularly well for residential users. Further analysis also showed that there is a strong correlation between vendors and DNS settings.

NAT implementations violate the RFC. To characterize the behavior of NATs, we used the open source STUN client/server [2]. Out of 189 APs that we tested, 4 could not be tested since they reported that the server was not reachable or blocked. Table 5 shows a classification of the remaining 185, all of which used NAT. RFC 4787 [4] describes the Best Current Practices for NATs with respect to these properties. Interestingly, 16 out of 185 NAT-enabled APs (9%) did not use endpoint-independent port/address map-

Property	# of APs
<i>Address Mapping</i>	
Endpoint-dependent	16 (9%)
Endpoint-independent	169 (91%)
<i>Hairpinning</i>	
Yes	90 (49%)
No	95 (51%)
<i>Filtering Behavior (Endpoint Independent)</i>	
Independent	63 (37%)
Address Dependent	30 (18%)
Address & Port Dependent	76 (45%)

Table 5: Breakdown of NAT behavior.

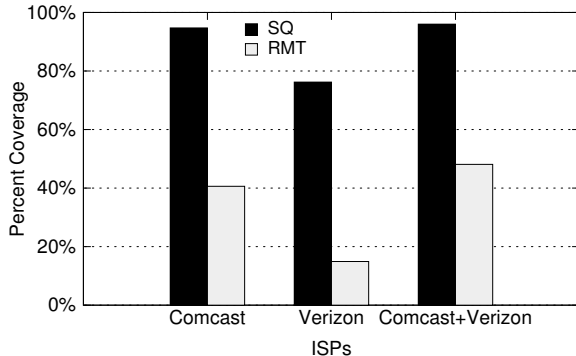


Figure 4: Coverage provided by ISPs.

ping which the RFC requires (REQ-1). Similarly, over half (51%) of the APs did not support hairpinning, also required by the RFC (REQ-9). The overwhelming majority of these RFC violations are specific to a small number of vendors.

4.4 Neighborhood Network

The strength of our 2 pass measurement is that it allows us to characterize wireless and Internet connectivity of a neighborhood as a whole, which can provide useful insight for neighbor-aware or neighbor-cooperative system design and deployment strategies.

Open APs offer diverse access to the dominant ISPs in the region. Recent efforts, such as community WiFi deployments and collaborative wireless access schemes raise interesting questions about the use of neighborhood networks for collaborative access to one or multiple ISPs for roaming (e.g., FON), for bandwidth sharing, or for improved reliability. We use *Mark-and-Sweep* measurements to help shed insight on the practicality of such schemes by identifying the area covered by each AP, and through that, the area covered by wireless APs serviced by each ISP. Figure 4 shows that open APs attached to Verizon and Comcast offer 96% coverage in SQ and 48% coverage in RMT, which has a much lower population density. (Note that these results mean that any area covered by *any* open wireless point is covered by either Verizon or Comcast.)

Home networks can be made robust to individual ISP outages. A second collaborative neighborhood application is that of making residential Internet access robust to access link and neighborhood-wide failures (e.g., a failed DSLAM or cable head-end). The performance of such a

system depends on the number of wireless hops required to reach an AP served by a different ISP. As a first approximation, we assume that the number of hops is roughly proportional to distance. We then measure the distance from every point on the map (actually, a uniform grid of points) to a “secondary” ISP, using the AP nearest the point to determine the “primary” ISP. The median distance to the “secondary” ISP was 93 and 150m for SQ and RMT. We also observe that the median difference between the distance to the “primary” ISP and the distance to the “secondary” ISP was low in both neighborhoods: 26 and 68 meters for SQ and the suburban RMT areas, respectively. This very preliminary measurement suggests that such alternate-ISP schemes could have traction, assuming the other technical and legal barriers to their adoption can be overcome.

5. CONCLUSION

Mark-and-Sweep is a new tool for measuring residential wireless and broadband network properties. Its two-pass method—quickly finding all access points in an area followed by detailed measurements from targeted locations—provides equivalent accuracy to previous methods in a fraction of the time. Our initial experience with *Mark-and-Sweep* produced several interesting insights, such as vendor influence on wireless security, NAT RFC non-compliance, 802.11n penetration, and coverage provided by open APs.

6. ACKNOWLEDGMENTS

We would like to thank Swapnil Patil for sharing his earlier experience in wardriving, and Vijay Reddy and Dexter Rietman for their help in measurement and data parsing. Finally, we thank our anonymous reviewers for their valuable feedback.

7. REFERENCES

- [1] IP Address to Geographic Location. <http://www.ip2geo.net/>.
- [2] STUN server. <http://www.stunserver.org>.
- [3] A. Akella, G. Judd, S. Seshan, and P. Steenkiste. Self Management in Chaotic Wireless Deployments. *Wireless Networks Journal (WINET), Special Issue on Selected Papers from MobiCom 2005*, 13(6):737–755, Dec. 2007.
- [4] F. Audet and C. Jennings. *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP*. Internet Engineering Task Force, Jan. 2007. RFC 4787.
- [5] V. Bychovsky, B. Hull, A. K. 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.
- [6] M. Dischinger, A. Haeberlen, K. P. Gummadi, and S. Saroiu. Characterizing Residential Broadband Networks. In *Proc. ACM SIGCOMM Internet Measurement Conference*, San Diego, CA, USA, Oct. 2007.
- [7] Z. M. Mao, C. D. Cranor, F. Douglass, and M. Rabinovich. A Precise and Efficient Evaluation of the Proximity between Web Clients and their Local DNS Servers. In *Proc. USENIX Annual Technical Conference*, Berkeley, CA, June 2002.
- [8] A. J. Nicholson, Y. Chawathe, M. Y. Chen, B. D. Noble, and D. Wetherall. Improved Access Point Selection. In *Proc. ACM MOBISYS*, Uppsala, Sweden, June 2006.
- [9] C. R. Simpson, Jr. and G. F. Riley. NETI@home: A Distributed Approach to Collecting End-to-End Network Performance Measurements. In *Passive & Active Measurement (PAM)*, Antibes Juan-les-Pins, France, Apr. 2004.