

Using Autonomous Robots to Diagnose Wireless Connectivity

Richard Wang, Manuela Veloso, and Srinivasan Seshan
Carnegie Mellon University

Due to the proliferation of wireless devices, many wireless users treat wireless connectivity as a black box. When wireless performance does not meet expectations, it can be a frustrating experience to try and resolve wireless issues. Wireless problems are more significant for mobile robots due to strenuous requirements for sustained wireless connectivity while moving [1]. Unfortunately, it can be difficult to understand the cause of wireless problems in real environments. First, wireless signals transmitted across the wireless medium are susceptible to attenuation, interference, and reflections from the surrounding environment and other wireless devices. Second, wireless connectivity depends on decentralized cooperation across heterogeneous devices. As autonomous robots are introduced in our environments, we believe they can be a perfect tool to capture detailed snapshots about our wireless environments to help diagnose wireless connectivity issues. In this paper, we show how these insights helped us to diagnose our robot's own motion-based wireless connectivity issues.

I. INTRODUCTION

Understanding wireless connectivity in real environments is hard. Much of the complexity stems from wireless transmissions occurring over an open, shared medium with a mixture of decentralized, heterogeneous devices [2]. Once devices begin to move, wireless problems become even more difficult to diagnose since wireless conditions around the device can change rapidly. The emergence of telepresence robots has shown that wireless devices in motion struggle to sustain uninterrupted wireless connectivity [1]. In this paper, we will show that autonomous robots can be a valuable tool for identifying the cause of poor wireless performance with direct observations of the wireless environment.

We focus on enterprise wireless networks composed of access points (APs) distributed throughout the environment to provide Internet access to devices at all locations. Today, motion-based wireless connectivity issues are difficult for users to resolve because:

- 1) wireless infrastructures are complex and vary over space and time
- 2) users have visibility and control over only their own device
- 3) wireless communication problems can require significant domain knowledge to deal with the range of hardware, drivers, and protocol layers

As a result, a natural reaction is to submit trouble tickets and wait some time for network administrators to come and resolve the problem. Even network administrators may

struggle to resolve the wireless issues because: 1. they have limited time due to the large number of users to administrators (25,000 to 6 in our case), 2. the problem must be easy to replicate, and 3. network administrators control the infrastructure APs but have limited visibility of the wireless medium.

Autonomous robots as a wireless tool can augment diagnosis of wireless problems by:

- 1) capturing fine-grain wireless maps reflecting actual propagation of wireless signals
- 2) serving as a vehicle to subject wireless devices to repeatable motions

This is made possible due to their ability continuously localize with high accuracy and autonomously and precisely navigate without human assistance. Detailed wireless maps help to reveal how the wireless medium is being used in order to eliminate unlikely causes of poor connectivity. They would also allow wireless users to diagnose simple dead zone coverage issues and perhaps also empower them to create more meaningful trouble tickets. Since wireless problems with motion are often short-lived, the ability to reliably repeat motions is essential for understanding more complex motion-based wireless connectivity issues.

In this paper, we will first show that autonomous robots can be used to collect detailed wireless measurements. Next, we show fine-grain insights allow us to better understand how our wireless infrastructure uses the wireless medium. Finally, we show how we were able to diagnose our device's own motion-based wireless connectivity issues.

II. INSIGHTS ABOUT SURROUNDING WIRELESS CONDITIONS

We now show the detailed insights that autonomous robots can capture without access to any sensitive wireless infrastructure APs. With these insights, we will be able to understand how the wireless medium is being utilized and see if possible infrastructure configuration issues may be causing our wireless connectivity issues.

A. AP Coverage

AP coverage ensures every location has at least one AP in range. Avoiding wireless dead zones is the responsibility of network administrators who manage the wireless infrastructure. They often try to place APs to provide a high minimum received signal strength indicator (RSSI) at every location. Our network administrators target a minimum RSSI of -60 dBm, which is much higher than -90 dBm that generally signifies no connectivity. The process of verifying

coverage simply requires sampling RSSI at all locations in the environment. Unfortunately, there are no practical solutions that require little human effort and achieve fine-grain sampling of the environment. As a result, there are situations where trouble tickets result in the discovery of wireless dead zones in practice.

We can automate this search for wireless dead zones by deploying autonomous robots to measure coverage across the environment. We were able to cover four floors of our enterprise environment. Figure 1 shows a histogram of median RSSI of the best available AP after dividing the environment into 1m x 1m grid regions. We see that AP coverage across two floors is very strong with few regions falling below the -60 dBm target. If there had been wireless dead zones, they would have been apparent in these histograms. As a result, wireless issues for these floors are unlikely to be due to wireless dead zones.

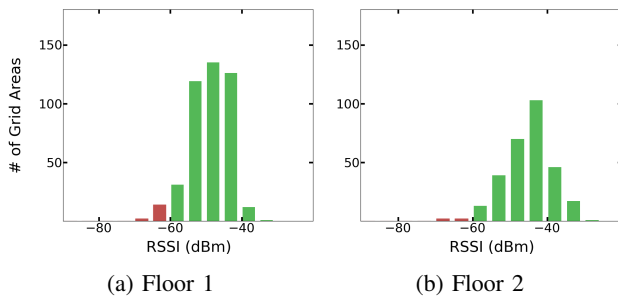


Fig. 1: Coverage summary showing histograms of the best median RSSI for each floor. Network administrators typically aim for a minimum of -60 dBm coverage.

B. Throughput Samples

Coverage is an important pre-requisite for wireless connectivity but not necessarily reflective of the actual rate of data transmission. Unlike RSSI that are instantaneous measurements, throughput samples depend on state and coordination with other wireless devices. Throughput tends to vary more than RSSI since congestion and dropped packets affect the rate of data transmission. As a result, throughput maps are unlikely to be a predictable as the coverage maps.

Figure 2 shows throughput maps collect by the robot as it moved across the environment. These measurements show how wireless performance varies over space. We can see that our robot’s own wireless connectivity problems are not isolated to small regions but spread across large regions of our building. This points to more systemic wireless issues that our robot is struggling with. If the robot was facing region-specific wireless issues due to excessive congestion, these types of throughput maps would have been helpful.

III. DIAGNOSING MOTION-BASED WIRELESS CONNECTIVITY

We have shown that our wireless infrastructure is well-configured and AP coverage is not an issue. Nevertheless, our throughput maps showed that motion-based wireless

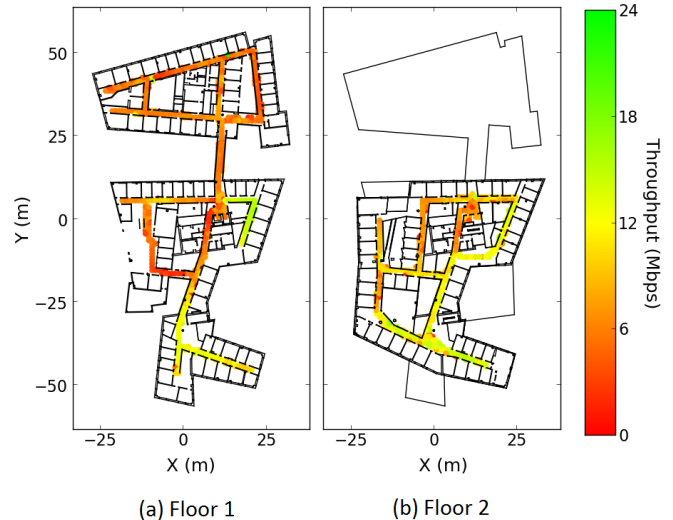


Fig. 2: Median throughput across two floors.

connectivity issues still persist. From our own empirical observations, these wireless issues appear intermittent and seemingly random while moving around. When we bring the robot back to revisit locations where it lost connectivity, the connectivity issues would not occur again so these problems must arise with motion. Autonomous robots will help to better understand these motion issues since they can continuously collect of wireless performance measurements while also reliably executing controlled motions. With the autonomous robots, we will methodically diagnose the root cause by enabling humans to search for similar patterns that lead to these poor connectivity situations.

A. Repeated Motions

Many factors including location and speed of motion can cause variations in wireless performance so we subject the wireless device to nearly identical situations. An autonomous robot itself is perfectly suited for subjecting the device to repeated traversals over the same path with the exact same speeds and device orientations. Deploying an autonomous is much preferred over fixed contraptions that are cumbersome and require modifications to the environment [3].

The robot is instructed to follow a simple three-quarter loop path around three hallways in the environment where connectivity issues occur frequently, as shown in Figure 3. We even instruct the robot to move in both clockwise (Figure 3b and 3d) and counterclockwise (Figure 3a and 3c) directions. We intentionally select a path where the robot traverses each location at most once. With no overlapping measurements at any location, it will be much easier to analyze the wireless performance variations using wireless maps.

B. Analyzing Variations in Wireless Performance

While being driven along the given path, the wireless device simultaneously captures RSSI, throughput, and current AP it is associated with. We present four noteworthy runs

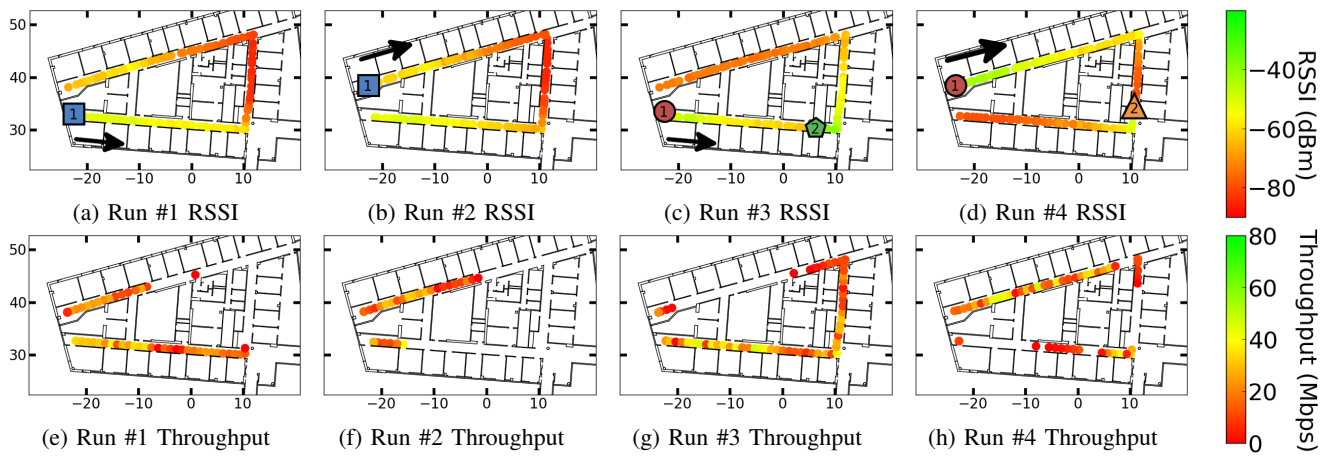


Fig. 3: Simultaneous associated RSSI and throughput for 4 runs over the same locations. Runs 1 (3a) and 3 (3a) began in the bottom left corner with the robot moving counterclockwise while Runs 2 (3b) and 4 (3d) started in the top left and moved clockwise. Numbered labels reflect the first point of association with each AP while the shape and color reflect a unique AP whose corresponding coverage is shown in Figure 4.

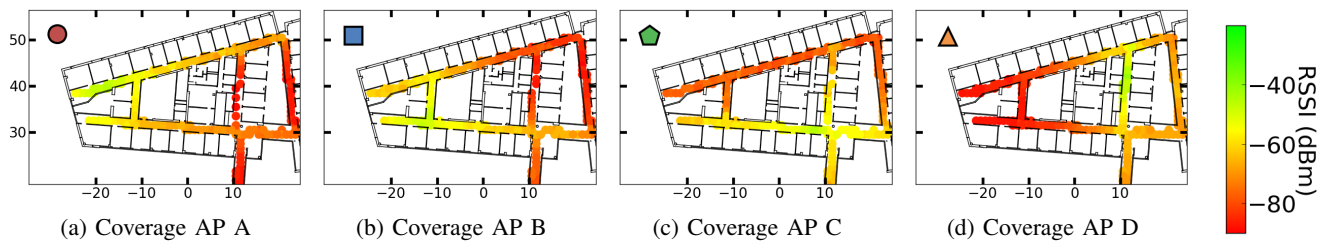


Fig. 4: Coverage for each AP corresponding to APs in Figure 3 identified with a unique shape and color.

in Figure 3 that shows RSSI (top), throughput (middle), and coverage for each AP (bottom). In the RSSI maps (top), the unique shapes reflect the location where the device first associated with the corresponding AP as identified by the color and shape. The numbered labels identify the order in which they were visited. We show AP coverage for each of these uniquely identified APs (bottom). Corresponding throughput while moving (middle) is also shown where large stretches of white space reflect absence of connectivity.

These four runs provide some noteworthy insights. First, RSSI changes gradually over several meters as a function of the device’s distance from the AP. Notice that run #1 and #2 remained associated with the same AP for the duration of the traversal. Irrespective of the direction of motion, RSSI for these runs nearly perfectly matches corresponding AP coverage. For these runs, throughput resulted in lengthy stretches of no connectivity since the AP was out of range.

In run #3 and #4, the device switches to another AP in the middle of the path. This AP switch particularly benefits run #3 but not as much for run #4. The difference is that run #3 switches APs just as it is about to enter the strongest AP coverage region for the selected AP. In contrast, run #4 switches to an AP that is almost out of range.

We can see in Figure 4 that there is at least one AP with high RSSI along the entire path so AP coverage is strong. The motion-based challenges must stem from poor AP handoffs. The key challenges appear to be centered

around timing disassociations before connectivity degrade significantly and then intelligently selecting the next AP to switch to. With the help of our autonomous robot, we are able to distinguish the effects of AP coverage, changing wireless conditions, and device motion to conclude that poor AP handoffs are the cause of our robot’s wireless issues.

IV. RELATED WORK

Past efforts to collect wireless measurements are unable to ensure fine-grain accuracy, densely cover spatially diverse areas, and provide timely updates. Unfortunately, it is difficult to predict the propagation of wireless signals in realistic, indoor environments so fine-grain wireless maps require measuring signals captured at each location. Measurement studies have been performed by having humans carefully traverse a building and mark their locations on a map [4], [5]. This is a tedious process that suffers from accuracy issues due to human errors that make it undesirable to repeat often so it will be difficult to ensure maps are up-to-date.

Dense deployments of static WiFi monitors can ensure timeliness but are limited by placement options for fixed location monitors and incur significant human effort and costs to deploy so typically they cannot achieve high spatial granularity. While some use dedicated sensor hardware [6], [7], [8], others reduce costs by adding WiFi dongles to available USB slots [9]. Distributed synchronization and hardware calibration enables creation of a single, unified

view from measurements collected across all WiFi monitors. A global view can be used to infer aggregate performance metrics like number of active wireless clients, interference, loss rates, and utilization [6], [7], [10] and even infer missing packets [8]. These approaches are limited to the perspective of the wireless infrastructure and have difficulty accounting for unreceived wireless client transmissions. In this paper, we view the wireless network from the perspective of the wireless client by accounting for the client's movement and considering the client's intent of transmitting wireless data.

Other efforts attempt to crowd-source collection of wireless maps. These approaches end up sacrificing accuracy in order to easily collect measurements. GPS can be used to provide location estimates [11], [12] but it operates primarily in outdoor environment and suffers from poor location estimates of around 3 meters. FM signals [13] similarly suffer from the effects of indoor environments and cannot achieve accurate location estimates. Recent efforts to take advantage of powerful sensors including odometry, magnetometer, and WiFi found in cell phones have been shown to have an accuracy of 1.69 m [14], [15]. Roomba robots have also been used to collect wireless coverage maps by spinning in small grid areas [16], [17] but they cannot autonomously navigate to reduce human time and effort costs or execute complex motions like our robot can. Our work takes advantage of much more powerful sensors that can localize within 10 cm using a wheeled platform that can reproduce complex movements.

Previous efforts have proposed techniques to use predictions to reduce the duration of handoffs or inform applications to allow for prefetching data and reduce the impact of handoffs [18], [19]. Nearby access points have also been used to opportunistically help mitigate WiFi handoffs for moving vehicles when moving across multiple buildings [20]. Our work that helps to expose and reproduce fine-grain failures in AP handoffs for moving devices is orthogonal to these efforts as it provides a mechanism for robustly evaluating handoff solutions.

V. CONCLUSION

Diagnosing wireless connectivity issues can be difficult due to the many factors that potentially impact wireless performance. We showed how autonomous robots can help to methodically drill down to the root cause by capturing detailed wireless measurements that eliminate unlikely factors. We were able to identify AP handoffs as the reason for our own robot's motion-based wireless connectivity issue by analyzing variations in wireless performance while subjected to repeatable motions. This was a challenging wireless problem that arose from poor decisions dependent on accurate timing and it is unclear that we could have uncovered them without the accuracy and control of autonomous robots. Opportunities for future work include using these detailed wireless maps for better management of enterprise wireless networks, ensuring timely maps for wireless localization solutions, and automated diagnosis of wireless problems.

REFERENCES

- [1] "Life telepresent: working vicariously through the beam robot," <http://www.theverge.com/2012/11/7/3611510/suitable-beam-robot-aims-for-bulletproof-telepresence>. [Online]. Available: <http://www.theverge.com/2012/11/7/3611510/suitable-beam-robot-aims-for-bulletproof-telepresence>
- [2] D. Eckhardt and P. Steenkiste, "Measurement and analysis of the error characteristics of an in-building wireless network," in *ACM SIGCOMM Computer communication review*, vol. 26, no. 4. ACM, 1996, pp. 243–254.
- [3] R. Fish, M. Flickinger, and J. Lepreau, "Mobile emulab: A robotic wireless and sensor network testbed," in *IEEE INFOCOM*, 2006.
- [4] P. Bahl and V. N. Padmanabhan, "Radar: An in-building rf-based user location and tracking system," in *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 2. Ieee, 2000, pp. 775–784.
- [5] M. Youssef and A. Agrawala, "The horus wlan location determination system," in *Proceedings of the 3rd international conference on Mobile systems, applications, and services*. ACM, 2005, pp. 205–218.
- [6] Y.-C. Cheng, J. Bellardo, P. Benkö, A. C. Snoeren, G. M. Voelker, and S. Savage, *Jigsaw: solving the puzzle of enterprise 802.11 analysis*. ACM, 2006, vol. 36, no. 4.
- [7] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, "Measurement-based models of delivery and interference in static wireless networks," in *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 4. ACM, 2006, pp. 51–62.
- [8] R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, "Analyzing the mac-level behavior of wireless networks in the wild," *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 4, pp. 75–86, 2006.
- [9] P. Bahl, J. Padhye, L. Ravindranath, M. Singh, A. Wolman, and B. Zill, "Dair: A framework for managing enterprise wireless networks using desktop infrastructure," in *HotNets IV*, 2005.
- [10] S. Rangwala, R. Gummadi, R. Govindan, and K. Psounis, "Interference-aware fair rate control in wireless sensor networks," in *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 4. ACM, 2006, pp. 63–74.
- [11] Y.-C. Cheng, Y. Chawathe, A. LaMarca, and J. Krumm, "Accuracy characterization for metropolitan-scale wi-fi localization," in *Proceedings of the 3rd international conference on Mobile systems, applications, and services*. ACM, 2005, pp. 233–245.
- [12] R. Mahajan, J. Zahorjan, and B. Zill, "Understanding wifi-based connectivity from moving vehicles," in *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*. ACM, 2007, pp. 321–326.
- [13] Y. Chen, D. Lymberopoulos, J. Liu, and B. Priyantha, "Fm-based indoor localization," in *Proceedings of the 10th international conference on Mobile systems, applications, and services*. ACM, 2012, pp. 169–182.
- [14] H. Wang, S. Sen, A. Elgohary, M. Farid, M. Youssef, and R. R. Choudhury, "No need to war-drive: Unsupervised indoor localization," in *Proceedings of the 10th international conference on Mobile systems, applications, and services*. ACM, 2012, pp. 197–210.
- [15] A. Rai, K. K. Chintalapudi, V. N. Padmanabhan, and R. Sen, "Zee: Zero-effort crowdsourcing for indoor localization," in *Proceedings of the 18th annual international conference on Mobile computing and networking*. ACM, 2012, pp. 293–304.
- [16] S. Sen, B. Radunovic, R. R. Choudhury, and T. Minka, "Spot localization using phy layer information," in *Proceedings of ACM MOBISYS*, 2012.
- [17] O. Rensfelt, F. Hermans, P. Gunningberg, L.-Å. Larzon, and E. Björnemo, "Repeatable experiments with mobile nodes in a relocatable wsn testbed," *The Computer Journal*, vol. 54, no. 12, pp. 1973–1986, 2011.
- [18] U. Javed, D. Han, R. Caceres, J. Pang, S. Seshan, and A. Varshavsky, "Predicting handoffs in 3g networks," in *Proceedings of the 3rd ACM SOSP Workshop on Networking, Systems, and Applications on Mobile Handhelds*. ACM, 2011, p. 8.
- [19] S. Seshan, H. Balakrishnan, and R. H. Katz, "Handoffs in cellular wireless networks: The daedalus implementation and experience," *Wireless Personal Communications*, vol. 4, no. 2, pp. 141–162, 1997.
- [20] A. Balasubramanian, R. Mahajan, A. Venkataramani, B. N. Levine, and J. Zahorjan, "Interactive wifi connectivity for moving vehicles," in *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 4. ACM, 2008, pp. 427–438.