

Wireless Map-Based Handoffs for Mobile Robots

Richard Wang¹, Matthew K. Mukerjee¹, Manuela Veloso¹, and Srinivasan Seshan¹

Abstract—Most wireless solutions today are centered around people-centric devices like laptops and cell phones that are insufficient for mobile robots. The key difference is that people-centric devices use wireless connectivity in bursts under primarily stationary settings while mobile robots continuously transmit data even while moving. When mobile robots use existing wireless solutions, it results in intolerable and seemingly random interruptions in wireless connectivity when moving [1]. These wireless issues stem from suboptimal switching across wireless infrastructure access points (APs), also called AP handoffs. These poor handoff decisions are due to stateless handoff algorithms that make wireless decisions solely from immediate and noisy scans of surrounding wireless conditions. In this paper, we propose to overcome these motion-based wireless connectivity issues for autonomous robots using highly informed handoff algorithms that combine fine-grain wireless maps with accurate robot localization. Our results show significant wireless performance improvements for continuously moving robots in real environments without any modifications to the wireless infrastructure.

I. INTRODUCTION

Mobile robots experience wireless connectivity issues that are foreign to most of us due to their more strenuous wireless needs. For example, people browsing the web usually only require bursty and small transmissions of wireless data. Video conferencing is one of the more strenuous people-centric applications that requires sustained high throughput and low latency wireless connectivity but is typically used in primarily stationary settings. Mobile robots supporting features like telepresence have the same wireless requirements while also moving across the environment. The addition of motion presents significant wireless challenges since wireless devices need to continuously switch between wireless infrastructure access points (APs). Unfortunately, poor AP handoff decisions are likely to produce lengthy interruptions in wireless connectivity when moving.

Many valuable features of mobile robots depend on uninterrupted wireless connectivity. For example, telepresence robots rely on remote human operators for navigation [1]. Intermittent wireless connectivity not only results in a poor telepresence experience but also an unresponsive robot. Semi-autonomous robots that can navigate on their own use wireless transmissions to dynamically schedule and modify tasks to perform. Wireless connectivity also helps to overcome on-board limitations by using the network to overcome

perceptual, cognitive, and actuation limitations [2]. As a result, it is very important for mobile robots to have reliable network connectivity.

Existing AP handoff algorithms are designed to operate ubiquitously in any environment for any device so they employ general techniques that depend on reactive scans of the wireless environment whenever connectivity becomes poor. Not only are scans noisy, they take several seconds by which time the snapshots may no longer apply for a moving robot. This can lead to a sequence of suboptimal handoff decisions that cause significant intermittent wireless connectivity issues.

In this paper, we aim to show that motion-based wireless challenges are not some inherent failure of wireless technologies or poor wireless network management and that they can be addressed with highly informed AP handoff algorithms. Our approach is targeted specifically for devices like autonomous robots that are capable of continuously localizing with high accuracy. We provide our handoff algorithm with fine-grain wireless maps that are continuously combined with the robot's location in order to make intelligent handoff decisions. Our results show substantial improvements in wireless performance over existing scan-based algorithms while the robot is in continuous motion at several speeds.

II. CHALLENGES OF WIRELESS HANDOFFS

Enterprise wireless networks are typically composed of wireless infrastructure access points (APs) carefully spread throughout the environment to ensure full coverage. Scan-based handoff algorithms today gain access to wireless Internet connectivity by first scanning for nearby APs that are in range. During a scan, wireless devices measure the received signal strength indicator (RSSI), which loosely reflects signal reception quality. The handoff algorithm identifies the highest RSSI AP and then associates with it in order to communicate with other devices in the network. Once RSSI of the current AP falls below some threshold, the wireless device disassociates from it and then repeats the scanning process to select the next AP to associate with. This process is referred to as an AP handoff.

When moving across a building, a wireless client may need to perform many AP handoffs. Our building has over a dozen APs per floor. A wireless client moving from one end to the other typically associates with around five APs to avoid interruptions in wireless connectivity. Scan-based handoffs cannot handle motion well because wireless devices are unable to simultaneously associate with an AP while also scanning for alternative nearby APs on different channels.

¹R.Wang, M. Mukerjee, M. Veloso, and S. Seshan are with the Computer Science Department, Carnegie Mellon University 5000 Forbes Ave., Pittsburgh, PA. 15213, United States {rpw, mukerjee, mmv, srini}@cs.cmu.edu

This research was partially supported by the Office of Naval Research under grant number N00014-09-1-1031 and the National Science Foundation award number NSF IIS-1012733.

Aggressive Disassociations - disassociate when RSSI below threshold and then scan for APs

Require: Current state S with RSSI of associated AP S_{RSSI} . Given a minimum threshold RSSI with AP, T_{RSSI} .

```

1: function AGGRESSIVE( $S, T$ )
2:   if  $S_{RSSI} \leq T_{RSSI}$  then
3:      $APs \leftarrow APScan()$ 
4:      $AP_{curr} \leftarrow max_{RSSI}(APs)$ 
5:      $Connect(AP_{curr})$ 

```

Highest RSSI AP Policy - associate with the highest RSSI AP for every location using known wireless map

Require: Current state S with device's location S_{LOC} and current AP S_{AP} . Previously collected wireless map M_{WiFi} .

```

1: function RADIO( $S, M$ )
2:    $APs \leftarrow APsAtLocation(S_{LOC}, M_{WiFi})$ 
3:    $AP_{curr} \leftarrow max_{RSSI}(APs)$ 
4:   if  $S_{AP} \neq AP_{curr}$  then
5:      $Connect(AP_{curr})$ 

```

Location-Based AP Selection - aggressively disassociate and then select AP by location

Require: Current state S with RSSI of associated AP S_{RSSI} and device's location S_{LOC} . Given a minimum threshold RSSI with AP, T_{RSSI} . Previously collected wireless map M_{WiFi} .

```

1: function INFORMED( $S, T, M$ )
2:   if  $S_{RSSI} \leq T_{RSSI}$  then
3:      $APs \leftarrow APsAtLocation(S_{LOC}, M_{WiFi})$ 
4:      $AP_{curr} \leftarrow max_{RSSI}(APs)$ 
5:      $Connect(AP_{curr})$ 

```

Look-Ahead Plan - Follow a precomputed a sequence of AP transitions along a known path.

Require: Current state S with device's location S_{LOC} . Given a precomputed queue $Q_{LOC,AP}$ of AP assignments for each location along device's given planned path from Algorithm 1.

```

1: function LOOKAHEAD( $S, Q$ )
2:   if  $isAtLocation(Q_{LOC}, S_{LOC})$  then
3:      $AP_{curr} \leftarrow Q_{AP}$ 
4:      $Connect(AP_{curr})$ 
5:      $Q \leftarrow Q.next()$ 

```

Fig. 1: Pseudocode of the four proposed handoff algorithms

In the interest of maintaining wireless connectivity, scan-based handoffs often choose to remain associated whenever possible. Despite these limitations, scan-based efforts have been sufficient for stationary wireless usage that do not move.

Unfortunately, significant wireless challenges arise when these scan-based handoffs must contend with devices in continuous motion:

- 1) Scans incur a 3 to 5 second time cost that become significant with the increased frequency of AP handoffs.
- 2) Scan-based efforts often disassociate when the current AP becomes weak with no idea if there had been other better alternative APs available for quite some time.
- 3) Scans are noisy and may no longer reflect current wireless conditions.

As a result, scan-based efforts often result in suboptimal handoff decisions that significantly degrade wireless performance for moving devices.

Better AP handoffs require improved timing of disassociations and more intelligent selection of APs to associate with. Late disassociations when the device is already out of range results in periods of no connectivity. Excessively frequent disassociations cause many interruptions in connectivity due to the overheads of switching APs [3]. Intelligent AP selection is equally important because associating with APs that are no longer in range will require additional handoffs.

The purpose of this paper is to investigate the impact of



Fig. 2: Emerging autonomous robots that perform tasks in the environment.

handoff algorithms that are fully aware of actual wireless conditions and device location. This has only recently become possible due to the emergence of autonomous robots like the one shown in Figure 2 that continuously localize with high accuracy [4], [5]. Their ability to continuously localize also makes them ideal tools for collecting fine-grain wireless

Algorithm 1 Pre-compute an AP plan - Use known path and WiFi map to precompute a sequence of AP transitions to minimize AP handoffs and maximize RSSI.

Require: Device's path as sequence of locations to visit P_{LOC} . Previously collected wireless map M_{WiFi} . Given a minimum threshold RSSI with AP, T_{RSSI}

```

1: function ASSIGNAPS( $P, M$ )
2:   for  $i=0; i < P_{LOC}.length; i++$  do
3:      $P_{APs}[i] \leftarrow getAPsAtLoc(P_{LOC}[i], M_{WiFi}, T_{RSSI})$ 
4:
5:    $N_{open} \leftarrow \{\}$ ;
6:   for  $AP$  in  $P_{APs}[0]$  do
7:      $N_{open}.push(\{0, AP, [(0, AP)]\})$ 
8:
9:    $N_{done} \leftarrow \{\}$ ;
10:  while  $!N_{done}.isEmpty()$  do
11:     $N_{next} \leftarrow \{\}$ 
12:    while  $!N_{open}.isEmpty()$  do
13:       $curr \leftarrow N_{open}.pop()$ 
14:      for  $i=curr_{LOC}; i < P.length; i++$  do
15:        for  $AP$  in  $P_{APs}[i]$  do
16:          if  $AP == curr_{AP}$  then
17:            continue;
18:           $next \leftarrow$ 
19:             $\{i, AP, curr_{PLAN} + (i, AP)\}$ 
20:          if  $i == P.length - 1$  then
21:             $N_{done}.push(next)$ 
22:          else
23:             $N_{next}.push(next)$ 
24:          if  $!(curr_{AP} \text{ in } P_{APs}[i])$  then
25:            break;
26:     $N_{curr} \leftarrow N_{next}$ 
27:   $APPlan \leftarrow GetMaxRSSIPlan(N_{done})$ 
28:  return  $APPlan$ 

```

maps of the environment that will be used by our approach.

III. WIRELESS MAP-BASED AP HANDOFFS

Our proposed informed AP handoff algorithms take advantage of highly accurate wireless maps and continuous robot localization. We consider a set of wireless handoff algorithms that are iterations of one another. This will allow us to see what pieces are most effective at addressing motion-based wireless connectivity challenges. The four algorithms in order of increasing complexity are: **aggressive disassociation**, **location-based AP selection**, **highest RSSI AP policy**, and **look-ahead planning**. A comparison of their pseudocode is shown in Figure 1 and we explain each in further detail below.

a) Aggressive Disassociation: As shown in Figure 1, this mimics the default, scan-based handoff algorithm that disassociates when RSSI of the current AP falls below some threshold. Given strong AP coverage across our environment, we use an aggressive RSSI threshold T_{RSSI} of -70 dBm.

b) Location-Based AP Selection: Instead of scanning for surrounding APs, this algorithm queries the given wireless map for the best available AP at the device's current location. By using a wireless map, this approach removes the uncertainty of scans by providing the device with the actual highest RSSI AP at its current location. We will be able to see how much of an impact optimal AP selection has on wireless performance.

c) Highest RSSI AP Policy: A device that is continuously aware of its location and has access to accurate wireless maps does not need to rely on any wireless measurements. Instead, it could simply always associate with the highest RSSI AP at all times by using its location as reflected in the pseudocode. This avoids the need to specify a threshold for disassociation and ensure the device is always associated with the highest RSSI AP. An example wireless map showing the highest RSSI AP across our environment is shown in Figure 3a. This approach will highlight the importance of high RSSI on wireless performance.

d) Look-Ahead Planning: Finally, we consider a hand-off planning algorithm with full awareness about the device's location, wireless map, and planned future movements. This approach is unique for autonomous robots because few other devices plan their movement in advance. This provides an opportunity to pre-compute a plan that minimizes the total number of handoffs while also timing AP switches to occur at opportune times. As an example, the pre-computed AP assignment plan used in our evaluation is shown in Figure 3b. The pseudocode shows that executing this pre-computed plan simply requires the device to switch APs as it reaches waypoints where AP handoffs should occur. The challenge is pre-computing this AP handoff plan.

Details of pre-computing this AP handoff plan are reflected in Algorithm 1. The function *AssignAPs* is given the robot's future path and the wireless map. The path is represented as a sequence locations. Using the wireless map, we first compute the set of APs that exceed a minimum threshold RSSI, T_{RSSI} , for each location along the specified path. Next, we perform an iterative deepening depth-first search to find the set of nodes requiring a minimum number of AP handoffs that ensures connectivity along the entire path. Each node in our search tree is a tuple $\{LOC, AP, PLAN\}$. $PLAN$ is a sequence of waypoint and AP assignment pairs. Included for clarity are LOC and AP that indicate the last waypoint and AP assignment pair considered by the node. Each iteration i of the search generates the set of all possible i AP transitions from the initial waypoint. Nodes that assign the same AP to a longer sequence of waypoints will naturally be closer to finding a $PLAN$ that assigns APs to the entire path.

The search begins by initializing a set of nodes that consider the device associating with all possible APs at the starting position $\{0, AP, [(0, AP)]\}$. For each iteration of the search, each node in the set N_{curr} will either be placed in N_{done} or N_{next} . Nodes that assign APs along the entire path will be placed in N_{done} . All other nodes are placed in N_{next} and represent incomplete AP assignments that include all

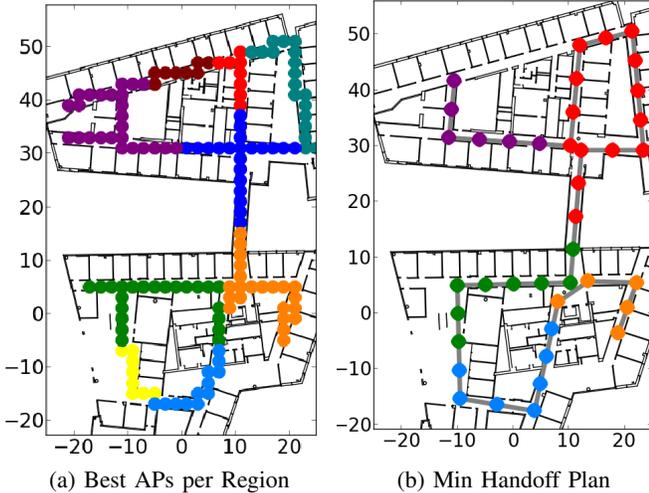


Fig. 3: Colors are used to uniquely identify each access point. The highest median RSSI APs per grid in our environment (left). Sequence of APs to connect to as computed by the pre-computed handoff plan that minimizes handoffs while maintaining a minimum RSSI of -70 dBm across the entire path (right).

	Aggr Diss	Loc AP Select	RSSI Policy	Look Ahead
Mean RSSI	-61.8	-59.4	-53.1	-55.3
Med RSSI	-61.0	-59.0	-53.0	-55
Mean Tput	29.13	41.47	43.60	45.44
Med Tput	33.00	46.00	46.53	48.51
# APs	6	4	8	5
Gap Time	35.50	10.62	18.30	10.26

TABLE I: Summary of aggregate measurements across four handoff algorithms. RSSI (dBms), throughput (Mbps), # AP switches, and gap time (s) without connectivity are shown.

possible ways in which the current node $curr$ can result in an additional AP transition $next$. Since an AP may be available for a lengthy sequence of path indexes, some nodes will be further in assigning APs for the entire path. This is why each node keeps track of its own last waypoint reached LOC and last associated AP .

Since our search considers switching APs at every waypoint location, we know that all possible combinations of AP switches that minimize the number of handoffs will end up in the set N_{done} . By selecting the plan that maximizes total RSSI across the entire plan, we will have a plan that minimizes the number of handoffs while optimally timing the handoffs so RSSI for both APs are strong at transition points.

IV. EVALUATION

We show how wireless performance is affected by these different informed handoff algorithms. We first look at fine-grain variations in wireless performance to reveal differences across the different handoff algorithms. We then perform

a thorough comparison of scan-based versus location-based handoffs at three different speeds over a total of 3.6 kilometers.

For our experiments, we use the omni-directional wheeled autonomous robot shown in Figure 2 that is equipped with a 802.11n WiFi dongle with a RT3575 chipset. During normal operation, the robot moves at speeds of up to $.75$ m/s. Prior to these experiments, the robot was driven around the environment while the dongle was set to monitor mode in order to capture RSSI of wireless signals transmitted by surrounding APs. The wireless measurements were bucketed by location into $1\text{m} \times 1\text{m}$ grids. The wireless map used in our work reflects the median RSSI of each AP observed in each grid.

A. Fine-Grain Wireless Performance

To show wireless performance differences across these handoff algorithms, we subject the wireless device to identical traversals of a complex path that requires the device to make many challenging AP handoffs. The path is shown in Figure 4 with starting location (S) and end location (E) marked. The numbers indicate the order that intermediate waypoints are visited. The device traverses each location at most once to make it easier to inspect performance variations with wireless maps.

Overall performance for one iteration is summarized in Table I. We see that **location-based AP selection** actually minimizes the number of AP handoffs, **highest RSSI policy** ensures highest median RSSI, and **look-ahead planning** achieves highest median throughput. Unsurprisingly, all three of our location-based handoffs clearly dominate the scan-based efforts.

Figure 4 shows more fine-grain details about variations in wireless performance. On top, we see a map showing RSSI while associated and successfully transmitting data to show where loss of connectivity occurs. The **highest RSSI AP policy** and **look-ahead planning** are most successful at ensuring consistently high RSSI across the entire path.

We also show variations over time for both RSSI (middle) and throughput (bottom). The labeled numbers correspond to the location waypoints on top. Notice that RSSI tends to dip suddenly during a handoff. The best performing algorithms immediately switch to another AP with higher RSSI. In contrast, notice that **aggressive disassociations** experiences two lengthy periods without connectivity that we see corresponds to RSSI lingering at levels of low RSSI. This is a result of poor AP selection from noisy scans.

These detailed measurements show why location-based handoffs are able to overcome systemic problems with scan-based efforts. We also see further opportunities for optimizing handoffs. For example, throughput in some cases is higher despite lower RSSI, suggesting other factors like congestion or interference may play a role. In addition, we can see for **location-based AP selection** how less frequent handoffs sometimes results in less jittery throughput so the duration of handoffs may be an important consideration in the future.

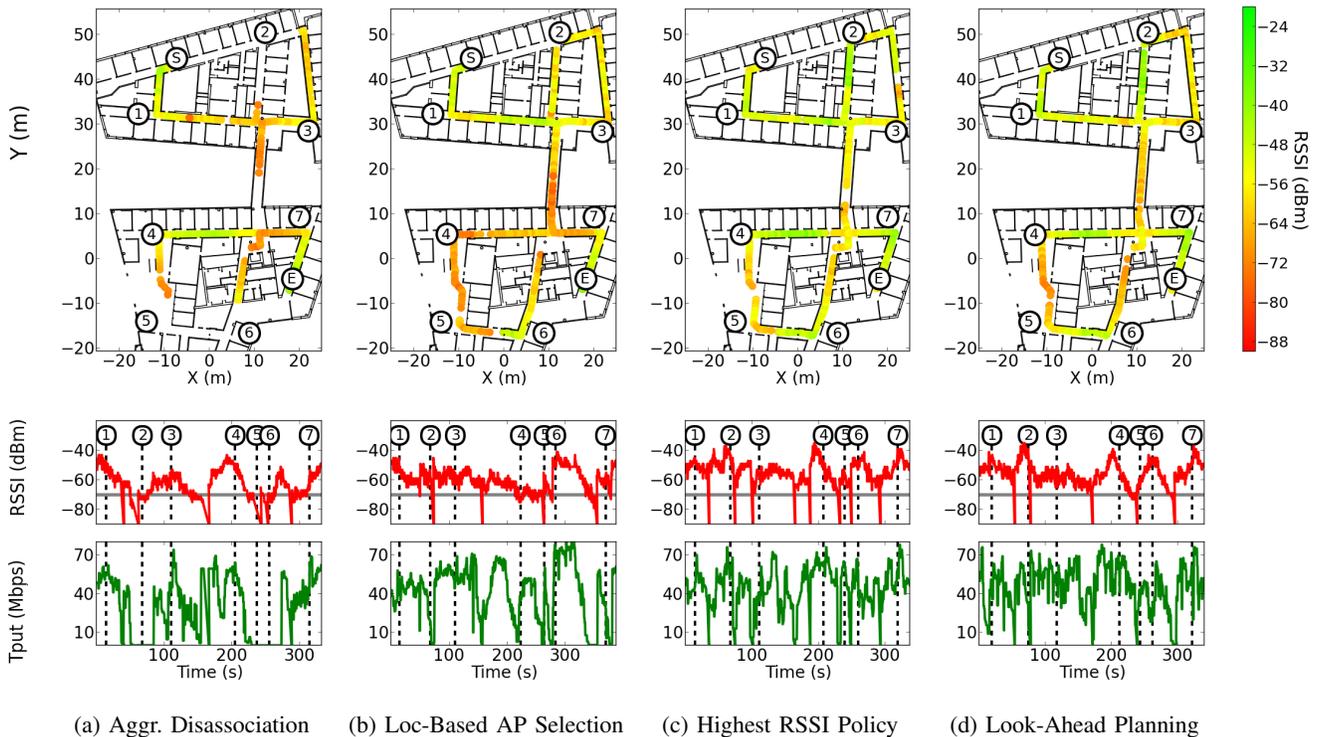


Fig. 4: Evaluation of AP handoff algorithms when subjected to the exact same motion path. RSSI of the associated AP at each location is shown in the spatial maps (top). The corresponding time-varying RSSI (middle) and throughput (bottom) with time marked when each intermediate path point was crossed. RSSI of -70 dBm is marked with a gray line as a reference (middle).

B. Aggregate Wireless Performance

We perform a robust comparison of scan-based and location-based handoff algorithms under continuous motion at several different speeds. In particular, we compare **aggressive disassociations** against **highest RSSI policy**. We select a 150 meter path and perform four iterations of each speed and handoff algorithm combination for a total of 3.6 kilometers.

Figure 5 shows a CDF reflecting the proportion of points along the path above some throughput. The path was divided every 2 meters and then we computed the average throughput for each 2 meter segment. The CDF shows the cumulative throughput for each of these segments. Notice the scan-based efforts do not have wireless connectivity for significant proportions of the path due to a combination of 3-5 second overheads for scans and increasingly frequent handoffs due to poor AP selection. Even when moving at the slowest speed, scan-based efforts are without connectivity for 10% of the path, which is intolerable for applications like telepresence. Wireless performance improves as the device moves more slowly because scans are more reflective of actual wireless conditions.

Notice how location-based handoffs at any speed almost always have wireless connectivity along the entire path. This

shows that many of the interruptions in connectivity for scan-based handoffs are not from switching APs but more from scans. As we can see, eliminating scans entirely and more intelligently selecting APs to switch to can significantly improve wireless performance. As a result, location-based handoffs are well-suited for ensuring reliable wireless connectivity while moving.

V. RELATED WORK

Many works have explored the collection of wireless maps. Initial efforts required exhaustive human effort to manually mark their locations while capturing wireless signals [6]. Location accuracy and timeliness can be improved by deploying grids of dedicated sensor hardware [7], [8], [9]. Attempts to decrease the human effort automate location estimates by using a combination of odometry, magnetometer, and WiFi found in cell phones with an accuracy of 1.69 m [10], [11]. Some efforts have used robots to predictably collect wireless signals [12], [13] but still require human assistance due to lack of autonomous navigation. Our work uses autonomous robots to collect fine-grain wireless maps without any modification to the environment or tedious human effort.

Research to understand connectivity issues has typically been from the *infrastructure point-of-view* as it is much easier

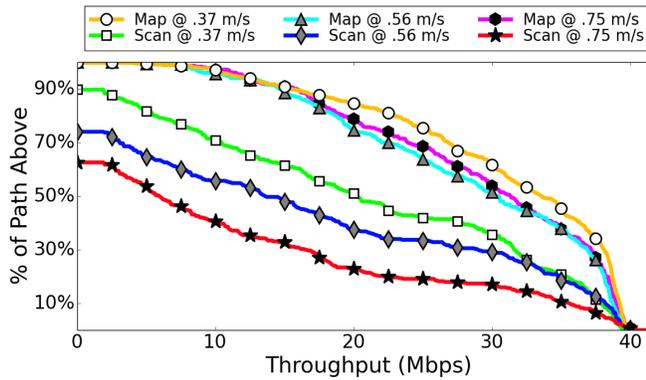


Fig. 5: Comparison of wireless performance for scan-based and location-based handoffs while continuously moving at three different speeds. CDF showing the % of measured segments along the evaluation path above some throughput. Notice that scan-based handoffs when moving at .37 m/s had no connectivity for 10% of the path.

to collect and aggregate measurements at APs rather than wireless clients. A global view of wireless signals observed across infrastructure APs can be used to infer aggregate performance metrics like number of active wireless clients, interference, loss rates, and utilization [7], [8] and even infer missing packets [9]. These approaches are only able to account for received wireless signals. This does not include failed transmissions that are not overhead or events that lead to particular wireless *client* handoff decisions, and thus are orthogonal to our work.

Prior work has proposed techniques to better predict and reduce the impact of handoffs as well as giving application the opportunity to prefetch data [14]. Nearby access points that happen to overhear wireless data despite not being the target AP have been used to opportunistically mitigate the effects of WiFi handoffs for vehicles moving across multiple buildings [15]. Efforts for faster handoffs have looked at reducing retry timeouts for vehicles [16] and synchronizing broadcast of beacon frames by modifying APs [3]. There are also orthogonal efforts to reduce the interruption associated with handoffs when they do occur [17]. More intelligent AP selection has been considered by using less accurate GPS location estimates [18]. Our work focuses on indoor environments and does not require other wireless devices or APs to adopt new wireless protocols.

VI. CONCLUSION

Mobile robots in continuous motion face significant wireless connectivity challenges. Current scan-based handoff algorithms are insufficient for meeting the strenuous connectivity demands for supporting features like telepresence for mobile robots. We have shown how wireless map-based AP handoffs can substantially improve the reliability of wireless connectivity for autonomous robots. In fact, these highly informed AP handoffs eliminate many of the interruptions in connectivity due to motion. While we did not observe significant changes in the wireless map across a few days,

future work is needed to investigate how one can maintain up-to-date wireless maps in dynamic environments. Future work is also needed to apply these same highly informed handoff algorithms to other mobile devices like cell phones that have much less accurate localization.

REFERENCES

- [1] "Life telepresent: working vicariously through the beam robot." [Online]. Available: "http://www.theverge.com/2012/11/7/3611510/suitable-beam-robot-aims-for-bulletproof-telepresence"
- [2] M. Samadi, T. Kollar, and M. M. Veloso, "Using the web to interactively learn to find objects." in *AAAI*, 2012.
- [3] I. Ramani and S. Savage, "Syncscan: practical fast handoff for 802.11 infrastructure networks," in *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, vol. 1. IEEE, 2005, pp. 675–684.
- [4] M. Veloso, J. Biswas, B. Coltin, S. Rosenthal, T. Kollar, C. Mericli, M. Samadi, S. Brandao, and R. Ventura, "Cobots: Collaborative robots servicing multi-floor buildings," in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*. IEEE, 2012, pp. 5446–5447.
- [5] J. Biswas and M. Veloso, "Multi-sensor mobile robot localization for diverse environments," *RoboCup 2013: Robot Soccer World Cup XVII*, 2013.
- [6] 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.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] S. Sen, B. Radunovic, R. R. Choudhury, and T. Minka, "Spot localization using phy layer information," in *Proceedings of ACM MOBISYS*, 2012.
- [13] 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.
- [14] 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.
- [15] 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.
- [16] J. Eriksson, H. Balakrishnan, and S. Madden, "Cabernet: vehicular content delivery using wifi," in *Proceedings of the 14th ACM international conference on Mobile computing and networking*. ACM, 2008, pp. 199–210.
- [17] S. Pack, J. Choi, T. Kwon, and Y. Choi, "Fast-handoff support in IEEE 802.11 wireless networks," *IEEE Comms Surveys and Tutorials*, vol. 9, no. 1-4, pp. 2–12, 2007.
- [18] A. J. Nicholson and B. D. Noble, "Breadcrumbs: forecasting mobile connectivity," in *Proceedings of the 14th ACM international conference on Mobile computing and networking*. ACM, 2008, pp. 46–57.