The Case for Robotic Wireless Networks

Mahanth Gowda
University of Illinois (UIUC)
gowda2@illinois.edu

Ashutosh Dhekne
University of Illinois (UIUC)
dhekne2@illinois.edu

Romit Roy Choudhury
University of Illinois (UIUC)
croy@illinois.edu

ABSTRACT

This paper explores the possibility of injecting mobility into wireless network infrastructure. We envision WiFi access points on wheels that move to optimize user performance. Movements need not be all around the oor, neither do they have to operate on batteries. As a rst step, WiFi APs at home could remain tethered to power and Ethernet outlets while moving in small areas (perhaps under the couch). If such systems prove successful, perhaps future buildings and cities could offer explicit support for network infrastructure mobility.

This paper begins with a higher level discussion of robotic wireless networks – the opportunities and the hurdles – and then pivots by developing a smaller slice of the vision through a system called iMob. With iMob, a WiFi AP is mounted on a Roomba robot and made to periodically move within a 2 x2 sqft region. The core research questions pertain to nding the best location to move to, such that the SNRs from its clients are strong, and the interferences from other APs are weak. Our measurements show that the richness of wireless multipath offers signi cant opportunities – even within a 2 x2 sqft region, locations exist that are 1.7 x better than the average location in terms of throughput. When multiple APs in a neighborhood coordinate, the gains can be even higher. In sum, although infrastructure mobility has been discussed in the context of Google Balloons, ad hoc networks, and delay tolerant networks, we believe that the possibility of moving our personal devices in homes and of ces is relatively unexplored, and could open doors to new kinds of innovation.

Categories and Subject Descriptors

•Networks ! Network architectures; Wireless access points, base stations and infrastructure; Cyber-physical networks;

Keywords

Wireless; Robotic Networks; Infrastructure; Measurement

1. MOTIVATION AND VISION

The last 30 years have witnessed signi cant advancements in wireless networking, ranging from hardware improvements to breakthroughs in theory, algorithms, and protocols. In the recent years, however, there is growing agreement in the research community that gains from the lower layers (MAC and PHY) are reaching saturation. Many are beginning to believe that the next "jump" in network performance will emerge from new ways of organizing networks [1–5]. In considering new network organizations, we explore the possibility of merging wireless networking with robotics. Speci cally, we ask: what if network infrastructure of the future – WiFi APs, enterprise WLANs, cell towers – are empowered with the ability to move physically? In pursuit of this thought, we began surveying the current state of robotics, as well as the pros and cons of physically moving infrastructure (e.g., WiFi APs on wheels, or cell towers on drones). We make a few initial observations below.

(1) Infrastructure mobility may not be viewed as a one-size- t-all solution, rather as a spectrum of opportunities illustrated in Figure 1. The opportunities range from centimeter scale antenna mobility to exploit multipath opportunities [6], to feet scale tethered mobility to evade wireless shadows and interferences, to full scale macro-mobility that minimize distance to clients. Network designers can choose to operate at different points on this spectrum, depending on user's requirements, budget, applications, and psychological comfort.

(2) Mobility is expected to bring a new degree of freedom (DoF) to network design, but more importantly, this DoF compliments existing dimensions of wireless innovation. Techniques for power control, channel allocation, localization, topology control, can all bene t if APs are able to move, even in the scale of inches.

(3) The time scale of mobility can be regulated as necessary. Small scale mobility can be used to compensate for small changes in network conditions, while full scale mobility can be triggered occasionally when the system moves to a skewed state, or a strict QoS requirement is ordered. In cellular networks, for instance, quad-copters could occasionally y out from cell towers and position themselves strategically to meet users' demands – like a network cloudlet [2, 3]. Infrastructure mobility could evolve as an on-demand service, a cost-effective and scalable alternative to over-provisioning.

Of course, some basic questions arise.

(1) Is moving infrastructure really practical? Concerns on feasibility are valid, but could perhaps be alleviated by building the vision in small systematic steps. Advances in personal robotics, beginning from the popular Roomba [7] to the more recent quad-copters [8–11] are already mainstream. Hardware is rapidly becoming cheap and reliable – an Arduino based robot car chassis adequate for cradling WiFi APs is $16 today [12]. Based on where robotic technology stands today [13], it is certainly not the fundamental barrier to infrastructure mobility.

Questions on the architectural aspects are certainly more relevant, such as maintaining power/Internet connectivity to a mobile AP, tangling wires, awkward moving objects on the oor, etc. However, we do not envision an all-at-once technology deployment, rather we intend to activate functionalities incrementally.
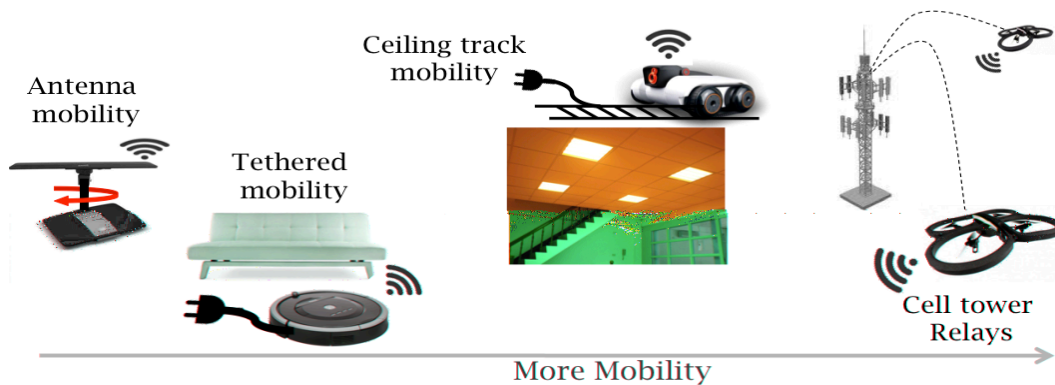
Figure 1: Regimes of infrastructure mobility, ranging from centimeter scale micro-motions, to feet scale mini-motion under couches, to building scale macro-motion perhaps on tracks laid on ceilings. Further into the future, perhaps ying quadcopters can serve as cell tower extenders, parking at strategic locations to meet client needs.

As a rst step in home settings, a mobile WiFi AP might just remain tethered to power and Ethernet, and only move in small spatial scales (say, under the couch or study table). In enterprises, airports, and hotels, the APs may also be tethered, but they could move in a coordinated manner (like a joint topology control problem) orchestrated by the cloud. Moreover, the AP movements need not be continuous; the time scales could slowly become more frequent as the system matures and gains social acceptance. Of course, facilities management and other logistical/policy questions will arise, but we believe they can be mitigated if the core performance gains are compelling.

(2) How compelling are the gains? While the answer obviously depends on numerous factors, the high level message is that the upper bound can reach 3x and more, compared to the static case. For example, in home environments, median throughput from 2 feet of mobility is 2 x for single clients, with the possibility of reaching 4x in 20% of the cases. With multiple homes, if APs coordinate to avoid mutual interference and optimize client SNR, median gain in overall network throughput can be 1.77 x or more.

It is crucial to recognize that the performance gains are not obtained by moving the AP close to one client – with multiple clients associated to an AP, moving close to one client will adversely affect others. The gains we observe actually arise from nding appropriate AP locations from which the SNRs to all its clients are strong. This is feasible due to rich spatial diversity in indoor environments, i.e., there exists certain nearby locations from which many clients experience strong channel conditions. In fact, the best AP locations could also experience lower interference from other APs and clients, enabling greater spatial reuse. On the other hand, blindly chosen AP locations can will fail to leverage these bene ts, resulting in far inferior performance.

iMob demonstrates the ability to improve throughput to 5 Å clients simultaneously. If too many more clients are active simultaneously, iMob can choose the top-K demanding clients and optimize their performance without affecting the others. If no solution is feasible, i.e., no AP location is able to satisfy the requirements, iMob could reduce the value of K. In the worst case, iMob will degenerate to a "static" AP and behave exactly as today's WiFi technology.

(3) Why move? Why not use MIMO, beamforming, or other software techniques? While these PHY layer techniques also leverage spatial diversity, mobility is still complimentary. Micro-shadowing scenarios are highly common in indoor environ-

ments [14, 15] – moving slightly can appreciably increase the rank of the channel matrix, resulting in higher MIMO gains. Our measurements con rm 3 x3 MIMO gains with today's 802.11 WiFi cards. Further, interference at the MAC layer is a function of energy, implying that AP1 would need to move out of AP2's carrier sensing range to enable spatial reuse. With beamforming/MIMO, AP1 will still sense AP2 and will defer communication. However, if AP1 could physically move out of AP2's range, or if AP1 and AP2 could jointly move to become "independent", system performance can improve further. Lastly, mobility and beamforming can be performed jointly to harness the best of both worlds.

The above is a high level vision (and qualitative arguments) aimed at motivating the overall research direction. We published a part of this vision in a workshop paper [16], along with toy measurements on USRPs using 1 MHz frequencies. This paper focuses on systematically characterizing the research landscape in real environments, and then builds a completely functional robotic AP system – iMob – using off the shelf 802.11n hardware. The key technical modules we develop are described next.

2. iMob: ROBOTIC WIFI ACCESS POINTS

As a rst step of the broad vision, we focus on small scale mobility in homes, in a way that is minimally disruptive to the established notions of a WiFi network. The iMob system we develop will allow WiFi APs to move on wheels while being tethered to the same power and Ethernet cable, as is currently used in most homes. Ideally, the APs could be placed away from human movement, such as underneath a couch or a side-table, or at the corner of a room [1]. In this setting, the iMob system will be tasked to offer performance gains to client devices. The main technical components we develop are as follows:

[2] We begin by measuring the upper bound on performance gain achievable through feet-length mobility of WiFi APs. These gains are measured using a testbed of 8 laptops mounted on Roomba robots – the laptops run 3x3 MIMO using Intel 5300 802.11n cards. Using one of the devices as a mobile AP and others as scattered clients, we nd the optimal AP location from which system performance is maximized. Besides serving as an Oracle, these measurements also offer insights into the nature of the gains, ultimately guiding the design of a real-time robotic networking system.

---

[1] This is anyway the case in many homes, given that network devices and wires are typically hidden from eyesight.

² We cross-check the Intel card results with USRPs and Atheros cards and verify that the gains scale across heterogeneous hardware (and not a function of our hardware idiosyncrasies).

² We then develop a practical iMob system in which the AP observes channel conditions and moves in real-time to the best estimated location. The motion planning algorithm uses insights from channel measurements, properties of the robot, and results from optimal stopping theory , to balance the tradeoff between exploration and exploitation (i.e., whether the AP should continue to explore more locations or should stop and perform remaining transmissions from its current location). This tradeoff naturally arises because the channel changes over space and time, and the AP does not possess the Oracle's view.

² We also build a coordinated iMob system in which the cloud moves multiple interfering APs (e.g., in neighboring apartments or houses) to optimize performance. This is essentially a topology control problem, with physical mobility as a degree of freedom. Both signals and the interferences can now be controlled to optimize desired performance metrics.

² We evaluate single AP iMob in faculty homes, student apartments, and in our lab. Coordinated iMob is evaluated with   4 APs deployed across 2 oors in our engineering building. Experiments are designed to evaluate a range of parameters and scenarios, including throughput and fairness, MIMO gains, impact of "leash length", impact of increasing number of clients, client mobility, etc. The overall gains are promising, and achievable without accurate prediction of wireless multipath and spatiotemporal channel variations. The inherent statistical nature of the environment offers viable opportunities.

3. MEASUREMENTS

To characterize performance upper bounds with mobility, we will exhaustively move APs in small spatial granularities and pick the best location that optimizes a given metric – we call this the   Oracle. We will then focus on understanding the nature of the gains, and utilize the insights to guide the design of a practical, real-time robotic WiFi system.

3.1 Experiment Platform and Methodology

Figure 2(a) shows a iMob AP assembled using a Roomba iRobot 2.1, a webcam, and a laptop equipped with Intel 5300 802.11n cards. The laptop is mounted on the iRobot and connected to it over the serial interface; it is also connected to a Microsoft live cam (attached in front of the iRobot) to guide its motion. The laptop acts as the controller for the whole system, sending motion commands to the robot (via the OSI interface), while also controlling the network interface for transmission/reception. 8 laptop clients were uniformly scattered at various locations and programmed to communicate back to the iMob AP.
The robot's mobility is con ned within a 2   x2 feet square region, demarcated by colored duct tapes pasted on the oor. If the robot drifts out of the square box, the camera detects the color of the duct tapes and triggers a change in heading direction. These square regions are selected from realistic areas in homes and apartments, i.e., near cable connection outlets. The AP performs "raster scans" within the square box (Figure 2(b)) at a speed of 10 cm/sec – during the scan, the AP continuously sends around 200 packets/second, equivalent to 60 packets per 3cms . Transmissions are performed on regular OFDM with 3x3 MIMO at both 2.4GHz and 5GHz bands. Clients record the per-packet channel state information   (CSI) for of ine analysis [17, 18].


Figure 2: (a) A laptop and a webcam mounted on a Roomba to emulate a iMob AP. (b) Raster scan in a box while communicating to scattered client(s).

The experiments were conducted in 4 different settings: (1) Student-of ce referred to as   Of ce. (2) Various corridors opening into the atrium of the engineering building, called   Lab . (3) Single bedroom graduate student apartment, called   Apartment. (4) Large single family home with APs placed in different rooms, called Home . In all cases, people moved naturally during the experiments, and clients scattered at realistic locations. Total measurements exceed 100 hours, generating 5TB of data.

Metrics: We evaluate performance in terms of data rates, throughput, and fairness. While the   Oracle selects the location with best data rate, our baseline scheme re ects today's static systems where the AP is placed at an arbitrary location near cable connection outlets. In light of this, the   median performance among all locations inside the 2 x2 feet square is treated as the baseline. Thus, the upper bound gain, for throughput say, is de ned as:

$$\text{Gain} \,\text{Æ}\, \frac{\max_{8i} \; throughput}{median_{8i} \; throughput}$$

where i denotes location i to which the AP can move to. Of course, when we design the real-time iMob system (later in Section 4), the median gain is not known to the AP since continuous raster scans are impractical. Still, the iMob AP should park itself at "good" locations from which the performance exceeds the median. We will discuss these later; for now, we focus on characterizing the system's upper bounds.

3.2 Characterizing Upper Bounds

The experiments are designed around 8 questions – the rst 4 focussed on the amount of performance gain, and the next 4 on understanding the nature of the gains.

(1) How much Data Rate Gain at Single Client?

Consider a case where the iMob AP moves within a box while continuously transmitting packets, and 8 scattered clients record the channel state information   (CSI) for every location of the AP. The CSI at each client can be accurately translated to the achievable data rate for communication between this client and the AP. For each tuple Ç Box_i , $Client_j$ È , we compute the max , median , and min data rates (to avoid outliers, we always use the 99^{th} percentile as max and the 1 percentile as min). Figure 3(a) plots the CDF of max minus median data rates due to the mobile AP, as well as the static AP, across all tuples. The key observation is that AP mobility induces large variations in data rates, far greater compared to the variations from temporal channel uctuations. Figure 3(b) plots the CDF of   median minus min data rates for both mobile and static APs, and shows that the reduction in data rates are also equally stronger due to mobility. Figure 3(c) further compares the range of data rates experienced in the same box by a mobile and static AP – the error
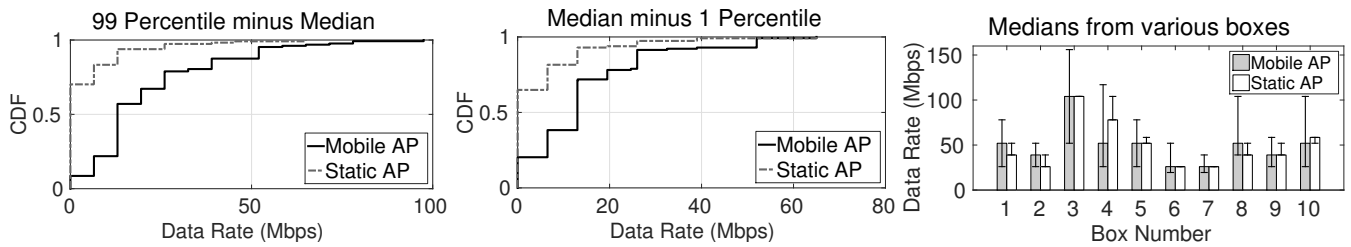
Figure 3: (a) Max minus Median data rates for mobile and static AP, verifying the greater diversity caused due to mobility. (b) Median minus Min data rates confirms that mobility also induces low data rates. (c) Comparison of the range of data rates for mobile and static APs in 10 randomly selected boxes (each bar representing Median and error bars representing [Max, Min]).

bars represent the max and min (Static's 1 percentile is sometimes the same as median due to low CSI variations). Clearly, mobility induces diversity.

While these results validate the known intuition that the wireless multipath signals interfere constructively or destructively in small spatial scales (causing diversity), it opens 2 specific opportunities for robotic WiFi applications.

(1) With centimeter scale mobility, an AP might appreciably improve data rate to a given client.

(2) With centimeter scale mobility, an AP can relocate to minimize interference from nearby APs/clients (potentially improving spatial reuse).

Assuming that the iMob AP is able to magically relocate to the best position, what is the gain possible compared to a static AP? Figures 4(a) plots the CDF of "rate gain" from 8 clients across 21 different boxes where the AP moved. We compute the rate gain as the ratio of $\frac{max}{median}$ data rate from each box. Evidently, an Oracle can easily double the data rate on average, and up to 4 x in ¼20% cases. Figures 4(b) now plots the CDF of "SNR reduction" to reflect how the mobile AP can move to avoid interference from nearby interferers. SNR reduction is computed as the difference between median and minimum SNR (note that interference is a function of energy and not the interferer's data rate, and hence plotted in terms of SNR). The achieved SNR reduction is around 4.5 dB on average, contributing to a modest improvement in spatial reuse and throughput. In summary, the potential gains seem substantial given that the AP moved within a box of side 2 feet.


Figure 4: (a) CDF of (max/median) data rates in a box indicates gain at a client. (b) CDF of (median − min) SNR in a box indicates gain from avoiding interference.

(2) Does Gain Scale to Multiple Clients?

In most realistic settings, the AP must serve multiple clients. So the natural question is: is there any AP location from which the data rates can be simultaneously improved for all clients? For this, we sum the data rates of all clients for each AP location within a given box – let S_i denote this sum for location i. Then we compute the average per-client data rate gain, β, defined as $\frac{max_{8i} (S_i)}{median_{8i} (S_i)}$. As before, the median represents the performance to be expected when the AP is placed statically at a

random location.

Figure 5 plots the CDF of β for increasing number of clients. The gains are obviously expected to diminish since the AP must satisfy a stricter condition, nonetheless, the gains are still upwards of 1.35x on average even with 7 clients, and up to 1.45 x for 3 clients. Homes mostly fall within this regime, where greater than 3 simultaneously backlogged connections are rare. In enterprises and hotspots (e.g., coffee shops), perhaps iMob can serve the 7 most data-hungry clients or the 7 weakest clients, improving the overall performance of the entire network. This result confirms the richness in indoor multipath diversity, offering support for robotic AP mobility even for the case of multiple clients.
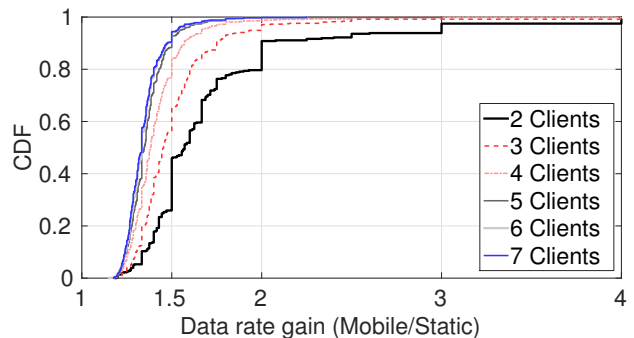

Figure 5: CDF of sum(data rate) gain over a static AP, where data rates are summed over multiple clients.

(3) How much Gain in Throughput?

Figure 6(a) plots the CDF of throughput experienced by each client due to AP mobility. If an Oracle were to pick the best AP location, the throughput gain (compared to a random location) is shown in Figure 6(b). Aligned with expectations, the throughput gains are proportional to the data rate gains, although slightly less due to wastage from backoff and DIFS/SIFS slots.


Figure 6: (a) CDF of throughput for each client during AP's mobility. (b) CDF of (max/median), i.e., the Oracle's gain over a randomly placed static AP.

(4) Does the Gain Scale across Environments?

Figure 7 reports the Oracle's median data rate gains from each of 4 environments, namely Office, Lab, Apartment, and Home. The

Figure 8: (a) CDF of number of high gain locations observed in one box (high gain defined as 0.95 of the max SNR in that box). (b) High gain locations within one box (marked with white dashes) shown as an example. (c) Distance traveled to reach high gain locations (distance defined as a percentage of raster scanning the entire box).


Figure 7: Oracle's median data rate gain in Office, Lab, Apartment, and Home in entirely uncontrolled settings.

reported gains are computed using the same metrics as above (i.e., max/median), and the experiments executed at 4 to 8 different places/rooms in each environment. The environment was entirely uncontrolled with natural human and object/furniture movements. Improvements are consistent, especially in the larger Office where the the clients are relatively further away from the AP (i.e., lower SNR). This is because modest improvement in SNR here can translate to greater rate improvement due to their logarithmic relationship.

To verify portability across hardware platforms, we performed similar measurements on USRPs and Intel cards. Figure 9 summarizes the results – this is loose in the sense that experiment conditions differed and some parameters were not identical (e.g., packet aggregation, MIMO, etc.) The key message is that the gains are consistent over static (single client), precluding any misgivings on our hardware.

	USRP (Software Radio)	Atheros Cards (AR 5418)	Intel Cards (5300agn)
Interference Suppression	8 dB	6 dB	4.5 dB
Throughput	1.7x	2.15x	2x

Figure 9: Comparison across platforms.

3.3   Understanding the Nature of Gains

While the upper bounds on performance are valuable, the extent to which the bounds can be achieved is also important. The next 4 questions are focussed on achievability.

(5) How Many High Gain Locations?

The existence of high gain locations is a necessary but not sufficient condition – if such locations are rare, the AP would have to spend a large time searching for it, affecting performance. Now, instead of targeting only the `maxDataRate` locations, we define high gain locations as those that achieve greater than 0.95 times the maximum data rate in that box. Figure 8(a) plots the CDF of the fraction of these high gain locations, computed across 64

boxes from all experiments (we define "locations" as a 3×3 cm^2 area as will be clear soon). Evidently, ≈ 40 high gain locations are available on average in a box, with some boxes offering far more. This is a favorable indication.

(6) How Scattered are High Gain Locations?

It is important to also characterize the scattering of the high gain locations within the box – if all the high gain locations are clustered in a small region, searching one of them can still be time consuming. Figure 8(b) shows one example of the scattering in one box – the white marks denote high gain locations and visually illustrate that they are "well scattered". However, to quantify this, we compute the distance, δ, that an AP must travel to encounter at a high gain location. Figure 8(c) plots the CDF of δ with randomly chosen starting positions, and with mobility similar to a 2D raster scan within the box. Evidently, δ is quite small for a large fraction of the cases, suggesting that high gain locations can be encountered without searching for too long. This brings hope that the potential gains might actually be achievable.

Of course, the above graph also suggests that in some cases, the AP needs to move a large distance to encounter a high gain location. However, this does not mean that for these cases, the performance will be poor. To capture this, we attempt to answer the following question: if the AP moves a pre-specified distance δ, what is the best performance that can be achieved? Specifically, for increasing values of δ, we record the best data rate encountered, and compare this data rate against a static AP (i.e., median data rate in the box) and the Oracle (i.e., the max data rate in the box). Figure 10(a) and (b) plot the two comparisons, respectively – δ is defined as a fraction of a full raster scan in the box. Figure 10(a) suggests that even when the AP travels a small distance ($\delta \approx 5\%$ of the raster scan), the data rate gain over static AP is still 1.5x. Figure 10(b) suggests that this gain reaches close to the Oracle. Thus, the overall message is that strong locations are not elusive – even if the best location is unavailable, "good" ones can be found quickly.

(7) How Predictable are High Gain Locations?

In designing a practical system, it would be useful if the existence of a nearby high gain location can be predicted. Such predictions may be possible if the locations surrounding the high gain location form a gradient, like a "hill". On the other hand, if the surrounding locations exhibit significantly less correlation to the high gain locations, then predictions are difficult. To this end, we compute the CSI at a given location and measure how the correlation degrades as we move gradually away from it. If the correlation degrades gradually, it would indicate the "hill" we desire. Figure 11 shows the results of this experiment. Unfortunately, we observe that CSI correlations are strong until separations of 2.5cms, but plummets drastically at separations of 3 cms
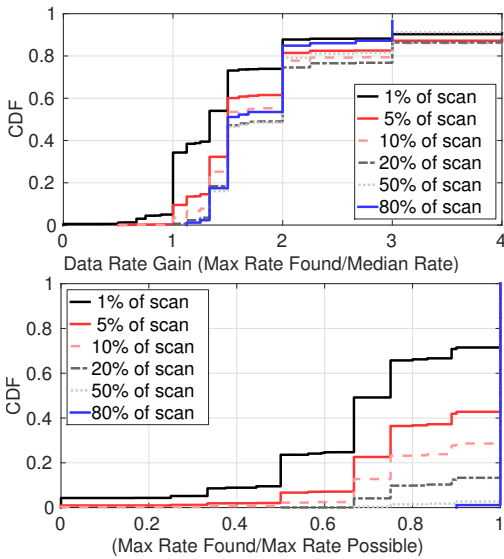
Figure 10: (a) The best location encountered after moving a small distance (local max) can still offer good gains over a static AP. (b) Local max is not too inferior compared to the *Oracle*.

and more. This implies that the coherence region of a signal is around 3 cms, and locations outside that region is a poor indicator of its neighborhood. We term this $3 \times 3\, cm^2$ coherence region as a pixel – which now defines a "location" – and recognize that neighboring pixels will vary drastically in SNR or data rate. Thus, the data rate landscape is like a "jagged mountain range" in the granularity of 3cms , making predictions difficult. These results and conclusions are consistent with multipath theory and independent measurements in literature [19–21].
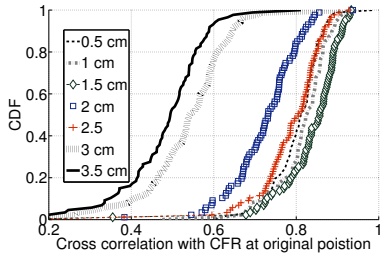

Figure 11: Data rates within a 5 cm shift of the mobile AP

(8) How Persistent are High Gain Locations?

If small changes in environmental factors cause the channel to change drastically, then iMob may not be worthwhile, since the AP will need to move very frequently. We classify environmental factors in 3 categories, namely human mobility, object mobility (e.g., doors, furniture), and device mobility (e.g., a smartphone moving in the user's hand). We then extensively investigate temporal stability by perturbing each of these factors – a human user typing on the keyboard, many people walking around, furniture moving, client laptops moving, etc. In the interest of space, we distill our key findings: (1) Client device mobility at the centimeter scale induces drastic change in the CSI, causing the channel to heavily fluctuate. iMob may not be beneficial to such devices (tablets, smartphones) when they are being held/carried in the hand. (2) For a static device (e.g., laptop, TV), human and object mobility impact the channel only when they block dominant signal components between the AP and the client. However, as shown in Figure 12(a) and (b), the channel revives once the human/objects have moved past. (3) Only when the human

or object moves to a new position, and also blocks the dominant signals, the CSI (and data rate) changes persist. However, such changes occur in the time scale of minutes [19] and can be detected by tracking changes in the CSI (detailed later). Thus, the take away message is that iMob could be effective even under dynamic environments, so long as the clients are static.


Figure 12: Data rate fluctuates when (a) humans, (b) objects go close to client, dwell for 10s, and walk past; the rate revives.

4. SYSTEM DESIGN

We take away 3 important messages from the measurements above: (1) The achievable performance improvement due to robotic AP mobility is substantial, available under realistic conditions (multiple clients and different indoor environments), and hence worth pursuing. (2) The high gain locations are challenging to model because they are randomly located, spatially small, and often juxtaposed next to poor SNR locations (making predictions difficult). (3) Although challenging, some opportunities offer hope – the high SNR locations are many, well scattered in a box, and stable for reasonable time scales even in real environments. This section is aimed at designing a practical AP motion planning algorithm that will suitably cope/leverage the above challenges and opportunities.

Some Design Guidelines

The core task of the algorithm is to search through different pixels (called exploration) and stop at a pixel that is expected to offer maximum performance gains (called exploitation). In the interest of space, we omit the various trials and deliberations that led to our final design; instead, we briefly discuss the key design guidelines that emerged from them. We will then assemble these guidelines into a practical iMob AP.

(1) Since AP mobility is at far slower time scales than packet transmissions, the *exploration* process must be speedy. Otherwise, an AP would spend unnecessary time at suboptimal pixels.

(2) Robotic motion is not accurate due to skidding of wheels, noisy compass values, mechanical turns – thus a robot cannot go back on the exact path on which it has traveled. This implies stopping decisions need to be made on-the-spot based on the SNR at that pixel. Performing a search and then retracing back to the *max* pixel on that path is not an option.

(3) The need to stop immediately at a high SNR pixel limits the maximum speed of the AP. Specifically, the inertial displacement after applying the brakes should be no more than a pixel width – to allow the AP to be within the same pixel once it decides to stop.
(4) Stochastic hill climbing or simulated annealing algorithms are not an option. Simulated annealing either incurs excessive time, or the starting point of the algorithm must jump to different random locations, which is impractical for the physically moving AP. Also, as mentioned earlier, these algorithms assume that backward motion is possible, which in our case is difficult.

(5) When clients move, or the environment changes too much, the CSI at the AP exhibits substantial change. This can be a trig-

ger for the AP to re-explore the best pixel, since the current one may have become sub-optimal. This is particularly necessary when this client is data hungry and optimizing its performance will boost the overall network performance.

Finally, and perhaps needless to say, the mobility heuristic must be lightweight to run on a simple robot in real time.

Optimal Stopping Theory

The crux of our heuristic is designed around a result from optimal stopping theory (OST) in applied statistics [22, 23]. The problem definition of OST is as follows. An employer intends to hire 1 individual out of n applicants (all of whom can be ranked based on quality). The applicants are interviewed one by one in a random order. However, unlike typical situations, in this case the interviewer must make a decision immediately after the interview; once rejected, an applicant cannot be recalled. Of course, during the interview, the interviewer can rank all candidates seen thus far, but is unaware of the quality of yet unseen candidates. OST asks: which candidate should be selected to maximize the probability of recruiting the best candidate. Selecting too early can leave many good candidates unseen; picking too late might mean that the best candidate is already rejected. The OST result dictates that the first $\frac{n}{e}$ candidates should be rejected, and among the subsequent candidates, the first on that ranks better than all $\frac{n}{e}$ candidates should be recruited.

OST bears a strong resemblance to our problem of selecting the best pixel, primarily because the pixels are scattered in an entirely random manner, with little spatial correlation (3 cms) (Figure 11). Hence, there is hardly a notion of "gradient" to leverage. Moreover, channel modeling or ray tracing seemed impractical since the iMob AP does not have details of the environment (floorplan, furniture, etc.) that would influence the multipath signal components. A statistical approach seems inevitable. In fact, since high SNR pixels are not rare and quite well scattered (recall Figure 8(b) and (c)), a statistical approach may be able to find such a pixel within a short time. The time to search can be reduced by moving the AP fast during the exploration phase, and slowing it down during exploitation (i.e., when its time to stop). With this background, we now describe the heuristic precisely.

Mobility Planning Heuristic

Figure 13 shows the flow-chart for iMob's mobility planning heuristic. The AP is placed at a random location by the user. Once it observes a stream of packets from a client, it begins an exploration phase. In this phase, it performs a raster scan at its maximum permissible speed, V_{max}, recording the channel state information (CSI) from each packet transmitted by client(s). Of course, the AP continues to communicate during exploration, moving through pixels of varying quality. The exploration continues until the AP has moved through $\frac{N}{e}$ pixels, where N is the total number of pixels in the box. At this point, the AP computes the best pixel among these $\frac{N}{e}$ pixels, where "best" is defined as an utility function of CSI:

$$U_{max} \approx \max_{p \in [1, \frac{n}{e}]} \left(\frac{\sum_i \log(SNR_i)}{I_p} \right)$$

where p denotes a pixel covered by the AP, i denotes the index of its own clients. I_p denotes the number of interfering APs and clients sensed at p. The AP now enters the exploitation phase. During exploitation, the AP computes every pixel's utility, and stops whenever a pixel's utility is $\geq U_{max}$. However, to brake and


Figure 13: Core flow diagram of iMob's heuristic

stop in the same pixel, the velocity of the AP must be reduced during exploitation. Otherwise, inertia and skidding of wheels will propel the AP forward, and returning back to this exact pixel will be time consuming. The reduced speed, V_{min}, is designed such that inertial displacement (after the application of brakes) is less than a pixel length (3cms) (discussed earlier). Once stopped, the AP continues communication with the client(s), expectedly at a near optimal data rate. The AP remains in this location until a new data hungry client joins, or until it observes a substantial change in the CSI of a client. Substantial CSI changes suggest mobility of the client or appreciable changes in the environment. Under both these conditions, the AP triggers the exploration phase again, and relocates to a new pixel. Additionally, the AP proactively relocates if it has been static for a very long time.

A common perception is that the exploration phase incurs a performance penalty because the AP is moving during this time and communicating from sub-optimal pixels. We observe that this sub-optimality is true with respect to the Oracle but not with respect to the static AP. Note that a mobile AP should statistically achieve the same performance as a static AP during exploration because the mobile AP will move through equal number of strong and weak pixels. Evaluation results confirm this (Figure 14(c)).

A natural question might be: what if the channel quality at other locations improve over time – an iMob AP will not be able to proactively exploit this opportunity. We observe that this is unlikely when CSI is used as the indicator function. If some other pixel has to improve substantially, then either the client must move to a new location, or the environment must change appreciably. Unlike SNR, both the effects will manifest in CSI variations.

Improvements to the Heuristic

We discuss a few optimizations to the core heuristic above.
(1) In some cases, the exploitation phase may not end quickly – the AP may not encounter a pixel offering U_{max} for a long distance. In such cases, the AP could be made to lower its expectations in proportion to the time spent in the exploitation phase. In other words, the AP starts with the hope to achieve U_{max}, but progressively lowers the bar to some fraction of this value. The rational is stop soon at a pixel that offers reasonable utility, as opposed to paying the cost for finding the perfect pixel.

(2) Data hungry clients, such as those that perform video streaming, are likely to be the highest beneficiaries of iMob. However, most video streaming clients buffer data, leaving bursts of time in which packet downloads are much less. The AP could exploit these gaps to explore – if new pixels are discovered with greater utility, it could relocate. Recall that the pixel at which the AP stopped moving is not guaranteed to be optimal – its only a statistical estimate using OST. Exploring more can still be beneficial.

Multi-AP Coordinated Motion Planning

We extend the above heuristic to multiple APs (e.g., in residential neighborhood) by engaging the cloud as a mobility coordinator.
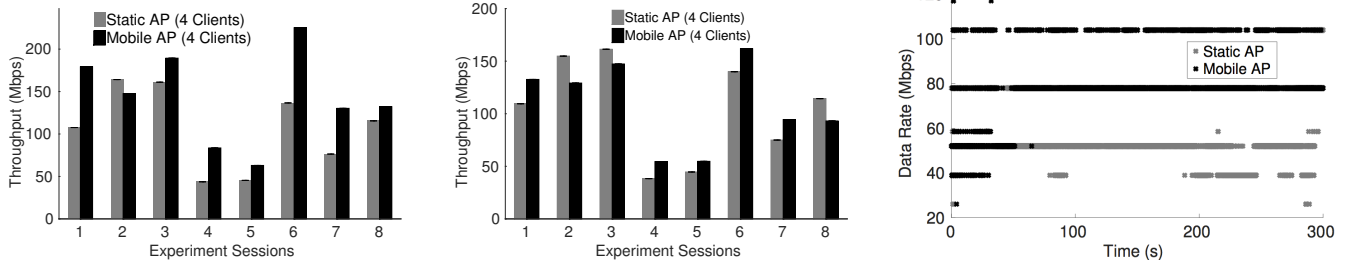
Figure 14: Throughput from real-time iMob with 4 clients: (a) Overall average throughput. (b) Average throughput when the AP is mobile, showing that AP mobility does not impose a performance penalty. (c) Data rate variation before and after stopping – the mobile AP's rates are comparable to the Static until it stops, and higher thereafter.

For ease of explanation, let us assume K APs numbered from 1 to K. The cloud instructs the APs to enter the exploration phase together, and each AP computes the maximum utility among the first $\frac{N}{e}$ visited pixels. Of course, these pixels not only experience different SNRs but also different interferences caused by the other (simultaneously moving) APs. Each AP reports at the end of the exploration phase, and once all APs have completed exploration, the cloud again instructs them to begin exploitation. However, the exploitation phase is executed in series, meaning that AP1 performs exploitation first, followed by AP2, and so on. This partially ensures that AP_i has at least accounted for interferences from all $(i - 1)$ APs. Each AP searches for a pixel that matches or exceeds its target utility, and stops upon finding one. Of course, this does not position APs in the optimal manner, but settles down in one reasonable configuration quickly. This also ensures convergence of AP movement.

5. EVALUATION

We evaluate a completely functional single and multi-AP iMob system and focus on (1) the throughput and fairness comparison with today's static APs, (2) gap from the Oracle, and (3) impact of various parameters, such as client density, traffic sessions, mobility area, etc. We begin with a description of the methodology.

5.1 Implementation and Methodology

The evaluation platform is similar to the measurement platform, with the following key differences. The iMob exploration/exploitation heuristic has been implemented in the Linux kernel (Ubuntu 10.04) to completely operate in real time (e.g., pixel search, utility computation, Roomba speed control, braking). Performance is measured on the wireless link only – the wired Internet connections at residences are the bottleneck, so connecting to the Internet would not reflect the actual wireless gains. We perform both single AP and multi-AP experiments. In the multi-AP case, a central server controls 4 APs – deployed across 2 floors of our university building – to extract holistic SINR and topological gains. Clients associate to our AP and upload/download packets over UDP/TCP while the AP moves to optimize performance. To compare against the Oracle, we performed experiments with continuous mobility and used the CSI data to precisely infer data rates [18] and throughput of each scheme. For realistic backlogged traffic, we record and use packet traces from YouTube, Google Hangout, and casual browsing sessions, captured from Wireshark. Across all experiments, the AP and clients were placed at realistic locations (to the extent possible). The environment was completely uncontrolled with people naturally moving, working, etc.

As a o21(atu-2188p)1(aco-1(n)1(g,)-213(o)F0(i)-1(r)g(u)-1(r)19(e)-208(a)11(r)52(

Figure 16: (a) AP detects when client moves and trigger relocation. (b) Detection accuracy for increasing client displacement. (c) Variation of data rates when human typing on a laptop versus the absence of humans.

the detection accuracy across all experiment sessions, as a function of the distance the client moved from its prior position. In some additional cases, the AP also triggered mobility because of CSI changed (even though the client did not move), but we are unable to verify if it was a valid trigger. This is because we do not know the ground truth on whether the environment truly changed or not, hence false positives cannot be computed in such cases. To shed more light, Figure 16(c) shows the CDF of throughput variation between two cases: (1) a human is typing and working with the client laptop, and (2) the client laptop without the human user. The similarity in deviation suggests that the channel does not vary due to the human working, obviating the need for iMob APs to move in such realistic cases.

5.2.2 Fairness and Leash Length

Figure 17(a) shows that throughput improvements with iMob is not obtained at the cost of fairness. Using Jain's Fairness Index, we nd comparable fairness performance as Static. Moreover, if desired, iMob can explicitly optimize for fairness, or even a combination of throughput and fairness. Figure 17(b) plots the variation of throughput with decreasing coverage area of the mobile AP. The performance does not degrade too much, indicating that the diversity is truly rich. This bodes well for iMob – even where the AP has less than a feet to move around, the single AP throughput gains can still be 40%.


Figure 17: (a) Fairness does not suffer with iMob. (b) Throughput loss for decreasing mobility area.

5.2.3 Comparison with Oracle

Figure 18 compares iMob's performance against Oracle and Static AP, for single client scenarios. The experiment sessions are derived from wireshark traces of YouTube, Hangout, and a casual browsing session. For example, for YouTube, active time windows were concatenated, while intermediate gaps (typical for buffered playback) were not considered. Evident from the graphs, increasing session lengths improve throughput because the sub-optimality during the exploration phase gets amortized over longer session lengths, and the performance at the best pixel begins to play a more dominant role. Figure 18(a) shows that iMob remains reasonably close to the optimal, around 0.9. Against Static AP, iMob continues to achieve around 40% gain on average, but exceeds 80% in few cases of longer traf c sessions.
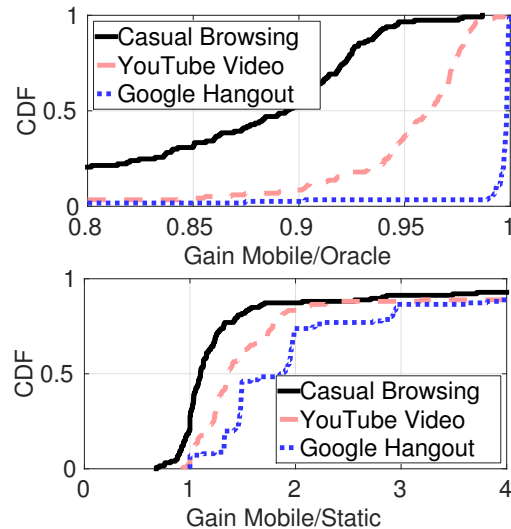

Figure 18: CDF of throughput gain for increasing traf c burst. (a) iMob over Oracle, (b) iMob over Static.

Figure 19 shows the variation of iMob's throughput against the Oracle and Static for increasing number of clients. iMob outperforms Static consistently and stays close to the upper bound. This suggests the ef cacy of the optimal stopping heuristic to nd a high quality pixel, even within 2 feet mobility.


Figure 19: Median throughput for increasing clients.

5.3 Real-time Multiple AP Experiments

Figure 20(a) shows the topology setup in our engineering building. The testbed is spread over two oors (2 nd and 4 th) and consists of 4 APs with a total of 6 clients (each AP associated to 1-2 clients). All APs were placed in the 2.4GHz channel such that the neighboring APs are at the edge of each other's interference range. Transmit powers were assigned at 8dBm to all the nodes; clients remain static for all the sessions. The topology mimics an EWLAN network of access points where the APs in the same channel are placed far from each other. A central server connects to each AP over WiFi and coordinates their movements to con g-

Figure 20: (a) iMob testbed deployed in 2nd and 4th oors. (b) Downlink throughput comparison. (c) Uplink throughput comparison. (d) Gain due to spatial reuse only, caused by sidestepping mutual interference from the other APs.

ure an effective topology that offers strong SNR to the AP's clients but avoids interference (to the extent possible) from other APs.

Figures 20(b) and (c) report the downlink and uplink UDP throughput comparison between the Mobile and Static AP. Gains are higher – 65% for downlink and 90% for uplink on average – implying that interference avoidance and better client SNR together contribute to net bene ts. Fairness remains greater than the static case (not shown here). Figures 20(d) zooms into this break-up and shows the improvements due to spatial reuse. The "Gain %" on the Y axis shows how much extra opportunity was created by evading interferers in comparison with the static AP case. The average gain was about 12%, considerably less than client throughput gains. This is because of the binary nature of the carrier sensing threshold (APs need to nd positions where the interferer is outside the sensing range). Nevertheless, the gains are still worthwhile because it combines multiplicatively with data rate gains resulting in net ampli cation in throughput.

6.   LIMITATIONS AND OPPORTUNITIES

This is an early attempt to characterize and exploit the landscape of robotic wireless networks. Much remains to be done.

² Moving client devices.   The key limitation with iMob is that constantly moving clients will not bene t from AP mobility, since the channel will change constantly. For such devices, however, the performance will still match the static AP. On the other hand, in favorable common scenarios where devices are static – video conferencing on laptops, streaming on smart TVs, even watching movies on a tablet on the table – gains are consistent.
² Joint Mobility and Power Control: Adding mobility to APs warrants revisiting classical problems in wireless networking. Power control and channel allocation can now be performed jointly with mobility, and adapted to changing traf c conditions.
² Localization and Security: Micro-moving APs may be able to mitigate the impacts of multipath, converging to a reasonably accurate pathloss index for their observed channel. Moreover, they could move macro distances to "look" at clients from different vantage points, ultimately improving the various techniques in triangulation and trilateration. Security bene ts also emerge from mobility, thereby changing the channel properties that are used as the "secret key" between the transmitter and receiver.

7.   RELATED WORK

Closest to this proposal is probably MoMiMo [6], where the receiver adjusts its antenna in centimeter scales to perform interference alignment. While MoMiMo is a speci c optimization for interference, this paper attempts to create a broader theme of robotic wireless networks, and presents a case for the regime of feet scale full-device mobility. Perhaps a further step in this

direction is "software de ned mobility" where the cloud controls the mobility of network infrastructure. Finally,   MoMiMo is complimentary to iMob – a WiFi AP can implement both. Google Loon [24] provides Internet access to remote areas via ad hoc network–style balloons drifting above the stratosphere. DARPA envisioned the use of self-autonomous network of LANdroid robots [25] to provide connectivity in warfare areas. Our broad proposal certainly bears similarities, but focuses on injecting controlled mobility to today's established infrastructure.

Spatial diversity has been exploited in MIMO, beamforming [26, 27], and through other opportunistic ideas [28, 29]. Infrastructure mobility is by no means an alternative to these. Our results show that moving within a 2 feet box can yield higher data rates even with a 3x3 MIMO interface – we believe that feet-scale mobility can offer higher ranked channel matrices. From the robotics side, authors in [30, 31] have researched how robots cooperate to achieve a common wireless communication goal. In one instance, robots plan their motion paths to constructively beamform towards a speci ed receiver. Authors in [32] have envisioned robots forming a "chain route" to maintain connectivity to rst responders (e.g., re ghters) moving into a catastrophe stricken building. Delay tolerant networks have also considered node mobility [33], even in under water [34] and mobile sensor networks [35]. We believe this paper is still different in the sense that it brings feet-scale controlled mobility to existing network infrastructure that are conventionally viewed as static.

In [16], authors envisioned robotic wireless networks and presented initial upper bounds on USRPs. This paper builds significantly on top of [16], including (1) full scale testbed with Intel 5300 802.11n 3x3 MIMO turned on, (2) insights into the nature of mobility gains, (3) a practical heuristic using optimal stopping theory, (4) a real time system running with single and multiple coordinated APs constructed out of Roomba robots.

8.   CONCLUSION

This paper envisions WiFi APs-on-wheels that move in controlled ways to optimize desired performance metrics. Early results are promising, although a deeper treatment is needed to fully characterize the interplay of many parameters underlying the success of such technology. Nonetheless, mobility is a valuable degree of freedom missing in today's network infrastructure, and extending research attention to it, we believe, is entirely worthwhile.

9.   ACKNOWLEDGMENTS

10. REFERENCES

[1] R. Murty, J. Padhye, R. Chandra, A. Wolman, and B. Zill, "Designing high performance enterprise wi- networks.," in NSDI, 2008.

[2] M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies, "The case for vm-based cloudlets in mobile computing," Pervasive Computing, IEEE, vol. 8, no. 4, pp. 14–23, 2009.

[3] G. A. Lewis, S. Echeverría, S. Simanta, B. Bradshaw, and J. Root, "Cloudlet-based cyber-foraging for mobile systems in resource-constrained edge environments," in Companion Proceedings of the 36th International Conference on Software Engineering, pp. 412–415, ACM, 2014.

[4] M. Jain, J. I. Choi, T. Kim, D. Bharadia, S. Seth, K. Srinivasan, P. Levis, S. Katti, and P. Sinha, "Practical, real-time, full duplex wireless," in ACM Mobicom, 2011.

[5] H. V. Balan, R. Rogalin, A. Michaloliakos, K. Psounis, and G. Caire, "Achieving high data rates in a distributed mimo system," in ACM Mobicom, 2012.

[6] F. Adib, S. Kumar, O. Aryan, S. Gollakota, and D. Katabi, "Interference alignment by motion," in ACM MOBICOM, 2013.

[7] "irobot roomba." http://www.irobot.com/us/learn/home/roomba.aspx.

[8] "Quadcopters." http://www.ted.com/talks/raffaello_d_andrea_the_astounding_athletic_power_of_quadcopters.html.

[9] R. R. Rajkumar, I. Lee, L. Sha, and J. Stankovic, "Cyber-physical systems: The next computing revolution," in Proceedings of the 47th Design Automation Conference, DAC '10, ACM, 2010.

[10] "Amazon prime air." http://www.amazon.com/b?node=8037720011.

[11] "Robots are becoming ready to work among us." http://www.technologyreview.com/view/522646/robots-are-becoming-ready-to-work-among-us/.

[12] "Smart robot car chassis kit for arduino (works with of cial arduino boards)." http://www.dx.com/p/smart-robot-car-chassis-kit-for-arduino-transparent-yellow-152984.

[13] "Advance in robotic technology." http://www.everything-robotic.com/.

[14] J. B. Andersen, T. S. Rappaport, and S. Yoshida, "Propagation measurements and models for wireless communications channels," Communications Magazine, IEEE, vol. 33, no. 1, pp. 42–49, 1995.

[15] P. Shankar, "Outage probabilities of a mimo scheme in shadowed fading channels with micro-and macrodiversity reception," Wireless Communications, IEEE Transactions on, vol. 7, no. 6, pp. 2015–2019, 2008.

[16] M. Gowda, N. Roy, and R. Roy Choudhury, "Infrastructure mobility: A what-if analysis," Hotnets, 2014.

[17] D. Halperin, W. Hu, A. Sheth, and D. Wetherall, "Tool release: Gathering 802.11n traces with channel state information," SIGCOMM Comput. Commun. Rev., vol. 41, pp. 53–53, Jan. 2011.

[18] D. Halperin, W. Hu, A. Sheth, and D. Wetherall, "Predictable 802.11 packet delivery from wireless channel measurements," ACM SIGCOMM Computer Communication Review , 2011.

[19] S. Sen, R. R. Choudhury, B. RadunoviÄG, and T. Minka, "You are facing the mona lisa: spot localization using phy layer information," in ACM Mobisys, 2012.

[20] S. Mathur, W. Trappe, N. Mandayam, C. Ye, and A. Reznik, "Radio-telepathy: extracting a secret key from an unauthenticated wireless channel," in Proceedings of the 14th ACM international conference on Mobile computing and networking, pp. 128–139, ACM, 2008.

[21] J. Zhang, M. H. Firooz, N. Patwari, and S. K. Kasera, "Advancing wireless link signatures for location distinction," in Proceedings of the 14th ACM International Conference on Mobile Computing and Networking , MobiCom '08, (New York, NY, USA), pp. 26–37, ACM, 2008.

[22] G. Peskir and A. Shiryaev, Optimal stopping and free-boundary problems. Springer, 2006.

[23] T. S. Ferguson, "Who solved the secretary problem?," Statistical science, pp. 282–289, 1989.

[24] "Google project loon." http://www.google.com/loon/.

[25] "Landroid robots to support communications in urban combat." http://defense-update.com/products/l/landroids.htm.

[26] D. Gesbert, M. Sha , D.-s. Shiu, P. J. Smith, and A. Naguib, "From theory to practice: an overview of mimo space-time coded wireless systems," JSAC, vol. 21, no. 3, pp. 281–302, 2003.

[27] J. Litva and T. K. Lo, Digital beamforming in wireless communications. Artech House, Inc., 1996.

[28] A. Miu, G. Tan, H. Balakrishnan, and J. Apostolopoulos, "Divert: ne-grained path selection for wireless lans," in ACM Mobisys, 2004.

[29] A. Balasubramanian, R. Mahajan, A. Venkataramani, B. N. Levine, and J. Zahorjan, "Interactive wi connectivity for moving vehicles," in ACM SIGCOMM, 2008.

[30] M. Zavlanos, A. Ribeiro, and G. Pappas, "Network integrity in mobile robotic networks," in IEEE Transactions on Automatic Control, 2013.

[31] N. Chatzipanagiotis, Y. Liu, A. Petropulu, and M. M. Zavlanos, "Controlling groups of mobile beamformers," in IEEE CDC, 2012.

[32] Y. Huang, W. He, K. Nahrstedt, and W. C. Lee, "Requirements and system architecture design consideration for rst responder systems," in Technologies for Homeland Security, 2007 IEEE Conference on, IEEE, 2007.

[33] P. N. Pathirana, N. Bulusu, A. V. Savkin, and S. Jha, "Node localization using mobile robots in delay-tolerant sensor networks," Mobile Computing, IEEE Transactions on , 2005.

[34] P. Corke, C. Detweiler, M. Dunbabin, M. Hamilton, D. Rus, and I. Vasilescu, "Experiments with underwater robot localization and tracking," in Robotics and Automation, 2007 IEEE International Conference on, 2007.

[35] L. Hu and D. Evans, "Localization for mobile sensor networks," in ACM MOBICOM, 2004.