Infrastructure Mobility: A What-if Analysis

Mahanth Gowda
University of Illinois (UIUC)

Nirupam Roy
University of Illinois (UIUC)

Romit Roy Choudhury
University of Illinois (UIUC)

ABSTRACT

Mobile computing has traditionally implied mobile clients connected to a static infrastructure. This paper breaks away from this point of view and envisions the possibility of injecting mobility into infrastructure. We envision a WiFi access point on wheels, that moves to optimize desired performance metrics. Movements need not necessarily be all around the floor of a home or office, neither do they have to operate on batteries, or connect wirelessly to the Internet. At homes, they could remain tethered to power and Ethernet outlets while moving in small areas (perhaps under the study table). In offices of the future, perhaps APs could move on tracks installed on top of false ceilings.

This paper explores the viability of this vision and presents early measurements from various home/office environments. We find that complex multipath characteristics of indoor environments cause large fluctuations in link quality even when the antenna moves in the scale of few centimeters. Mobile APs can leverage this spatial variation by relocating to a location that is strong for its own clients and yet weak from its interferers. Experiment results show that such micro-mobility itself can offer up to $2x$ throughput gains. Additional opportunities may emerge, such as in energy savings, security, QoS, and even in applications such as indoor localization. While this paper explores a small fraction of the landscape of opportunities, the results have been far more promising than what we had anticipated originally.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Wireless communication

Keywords

Wireless; Robotics; Mobility; Channel Diversity

1. INTRODUCTION

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
HotNets-XIII October 27 - 28 2014, Los Angeles, CA, USA
Copyright 2014 ACM 978-1-4503-3256-9/14/10...$15.00
http://dx.doi.org/10.1145/2670518.2673862.

The wireless research community, including hardware engineers, protocol designers, and information theorists, are actively exploring opportunities to improve the capacity of wireless networks. While significant advances have occurred in the last 30 years, there is growing agreement that gains from the lower layers (MAC and PHY) are reaching saturation. Many believe that the next "jump" in network capacity will emerge from fundamentally new ways of organizing networks. While pondering on ideas for new network architectures, we discovered substantial past work in this topic [1–5]. However, one assumption that all these propositions seemed to make is that infrastructure – WiFi APs, enterprise WLANs, cell towers – is static. As we considered the feasibility of relaxing this assumption, we began surveying the current state of robotics and the potential implications of physically moving wireless infrastructure (e.g., APs on wheels). We make a few observations

(1) Mobility is expected to bring a new degree of freedom to network design, but perhaps its more important to observe that this degree of freedom compliments existing dimensions of wireless innovation. Techniques for power control, channel allocation, beamforming, MIMO, localization, can all benefit if APs have the ability to move, even in the scale of inches.

Figure 1: Regimes of infrastructure mobility, ranging from centimeter scale micro-motions, to feet scale mini-motion under couches, to building scale macro-motion on tracks laid on ceilings. In future, flying quadcopters could serve as cell-tower extenders to meet client needs.

(2) Infrastructure mobility may not be viewed as a one-size-fit-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.

(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 fly out from the cell towers and position themselves strategically to meet users' demands. As networks become loaded and human tolerance goes down, infrastructure mobility might become an inexpensive alternative to over-provisioning.

Of course, some basic questions arise.
(1) Do we really need mobile infrastructure? What is the "killer app"? We admit that there may not be a killer app today to immediately embrace infrastructure mobility. Wireless capacity in homes and enterprises seem to be adequate; the infrequent problems of poor connectivity in corners of houses and airports are below tolerable thresholds. Perhaps enterprises can over-provision to solve most of the wireless networking problems they face today. Nonetheless, we believe that mobility is a powerful addition to the toolkit of networking techniques that is becoming feasible with personal robotics coming to the mainstream. The ability to understand the opportunities from this may trigger new ideas and applications in the future. Perhaps the cloud will control "swarms" of APs in an emergency relief operation; perhaps new forms of software defined mobility schemes will emerge. The benefits may not only be in terms of capacity, but also in energy reduction, security, QoS, fairness, network diagnostics, etc. This paper is motivated by the hope that bottom-up research is valuable in certain contexts, and believes this is one of them. However, to still ground the work in one instance of reality, we study infrastructure mobility in the context of homes and enterprises.

(2) Is moving infrastructure really practical? Concerns on feasibility are certainly valid, however, instead of trying to argue in favor of feasibility, we ask: *why should it not be feasible*? Advances in personal robotics, beginning from the popular Roomba [7] to the more recent $50 flying quadcopters [8], are already in the mainstream market. Hardware is rapidly becoming cheap and small; sensing and navigation algorithms are efficient [9]; robotic interfaces are rapidly maturing [10–12]. Based on where robotics technology stands [13], it is certainly not the fundamental barrier to infrastructure mobility.

Perhaps the viability concerns arise from architectural aspects, such as maintaining power/Internet connectivity to a mobile AP, tangling wires, awkward moving objects on the floor, etc. However, as mentioned earlier, we do not envision a one-size-fit-all solution. In certain cases, such as in homes, a mobile AP might just remain tethered to power and Ethernet, and only move under the table. In scenarios such as enterprises, airports, hotels, etc., AP motions may be instrumented on top of false ceilings, by placing tracks on which the robot-APs move. These tracks could be embedded with electric and network cables, such that the AP is always powered and connected to the backbone Internet. Movement need not be continuous – the time scales could evolve as the system matures. Building administration and other logistical/policy issues may arise, but we believe they can be overcome if benefits are compelling.

(3) How compelling are the gains? How much of the gains are achievable in practice? While the answer obviously depends on numerous factors, the high level message is that the upper bound can reach 2 to $4x$ compared to the average case. For example, in home environments, median throughput from one feet of mobility is around $1.6x$; with several neighboring interferers, it can increase up to $1.9x$.

Observe that gains are not always achieved by moving the AP close to one client – with many scattered clients, moving close to one client can mean moving away from others. Instead the gains come from the AP finding a nearby location from which the SNRs to all clients are strong, and the interference from other APs is weak (or inaudible). The former increases the transmission data rates while the latter enables spatial reuse. Fortunately, indoor environments cause wireless signals to be unexpectedly strong or weak at many scattered locations – investing a little bit of mobility to find these locations is profitable.

The above is the core intuition that makes appreciable gain realizable in practice. We envisage near-future systems focusing on small scale micro-mobility as a starting point – such systems can be plug and play, requiring no changes to the established WiFi eco-system. Over time, as demands grow and robotic systems get more accepted, perhaps macro-mobility will be of greater interest.

(4) Why not over provision the network (i.e., deploying many APs and picking the best ones any time)? For realistic densities, the gains from density and mobility are complimentary. Thus, the ideal strategy might be to increase density to the extent possible (i.e., under the constraints of hardware cost, re-wiring, and protocol changes [1]), and then inject mobility on APs. Sufficient AP density would obviate the need to macro-move.

The rest of this paper is designed to add specificity to these high level discussions. Our main contributions are: (1) Envisioning the landscape of infrastructure mobility. (2) Measuring the room for improvement, if unpredictable signal propagation in indoor environments are leveraged

through small scale AP mobility. (3) Preliminary validation of micromobility gains on Atheros WiFi cards.

2. EXPERIMENT SETUP

The next two sections focus on the upper bounds of gains as viewed from an oracle's perspective. Whether the bounds could be achieved is a separate question, and depends on how well new research problems could be solved. We will identify these problems as we progress through the paper. We introduce some terminology first.

2.1 Terminology

We define 3 regimes of AP mobility, shown in Figure 2(a).

• **Micro mobility** corresponds to moving in the granularity of centimeters, with the AP always remaining within a square region of 1 feet (called a "spot"). We also refer to this as *antenna mobility* and is envisioned to happen when the AP is, say, on top of a small table. Each location to which the AP (or the antenna) moves to is called a *pixel* and is a square of few centimeters. Benefits from micro-mobility arise mainly from wireless multipath effects, where multipath components constructively amplify the signal, or destructively nullify interference.

• **Mini mobility** refers to moving within three or four spots from the AP's installed location. We also call this *tethered mobility* since, in this mode, an AP can remain tethered to its power outlet and Internet cable modem. Moving under a table or couch are examples of mini-mobility. Benefits of mini-mobility arise mainly from alleviating wireless shadow effects, sidestepping interference, and from leveraging the vagaries of the wireless channel.

• **Macro mobility** corresponds to moving over the longest distances, in the granularity of 5 feet or more. Benefits of macro mobility arise mainly from pathloss (i.e., proximity to client) and from avoiding interference from other APs.

Measurement Set-up and Methodology

We perform measurements in four different settings – a 1600 square feet single-family **Home**, a graduate student **Apartment**, a student **Office**, and corridors and atriums of a **Lab**. Figure 2(b) shows a crude robotic AP designed for our experiments using the iRobot Create 2.1, laptops, webcams, and USRPN210 (a software radio) [14]. The laptop connects to the iRobot on the serial interface, to a USRP-N210 over Ethernet, and also to a webcam attached to the front of the robot. The laptop acts as the controller for the whole system, sending motion commands to the robot, while also controlling the transmissions from the USRP. 8 clients were uniformly scattered over the area and programmed to communicate to the robotic AP.

To guide this webcam-enabled robotic AP, colored tapes are pasted on the ground (Figure 2(c)). Every half meter the robot moves along the red line, it encounters a blue marker – this red-blue intersection is the center of a spot and the robot moves 40 pixels uniformly around this point. Once done, the robot continues on the red line to the next spot. The bandwidth for measurements was chosen as 1 MHz to decouple the effects of frequency diversity and evaluate pure mobility gains. The 5 GHz band was used to decouple interference in the crowded 2.4 GHz spectrum. We later present results from off-the-shelf Atheros WiFi cards showing gains comparable to USRP experiments.

2.2 Metrics

Our baseline for comparison will be a scheme in which the AP is placed randomly and remains static thereafter. To emulate this, we will characterize the performance at all pixels within a spot and pick the median pixel. Thus, throughput gain due to AP mobility will be:

$$Gain = \frac{\max_{spot}\ (\max_{pixel}\ (tput_{spot,pixel}))}{median_{spot}\ (median_{pixel}\ (tput_{spot,pixel}))}$$

Recall that we pick the *max* in the numerator to characterize the upper bounds on gain.

3. MEASUREMENTS

3.1 Upper Bounds with Micro-Mobility

Recall that upon arriving at a spot, the robot AP moves through the pixels within the spot transmitting a few packets from each pixel. Eight randomly scattered clients record the SNRs from every pixel in a spot – for each client this results in a SNR heatmap. If the AP moves through N spots, each client records N heatmaps. Figure 3 shows 4 heatmaps from 4 arbitrarily picked clients when the robot moved within a randomly picked spot. Darker shades in the heatmap indicate stronger SNR and the vice versa. We make two crucial observations:

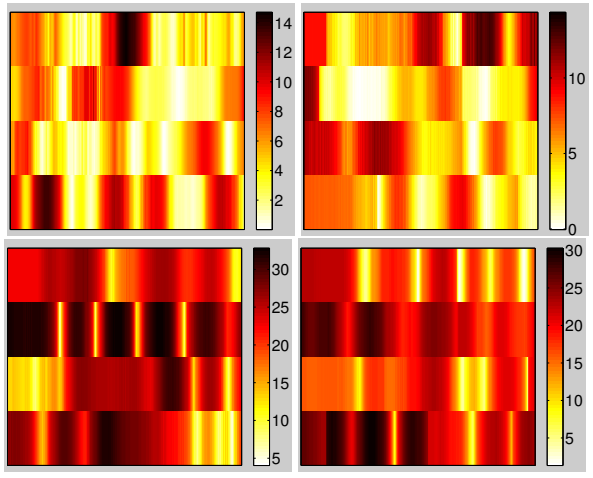

Figure 3: SNR heatmaps (dB): the top two are from clients far away from the AP; bottom two from nearby clients.

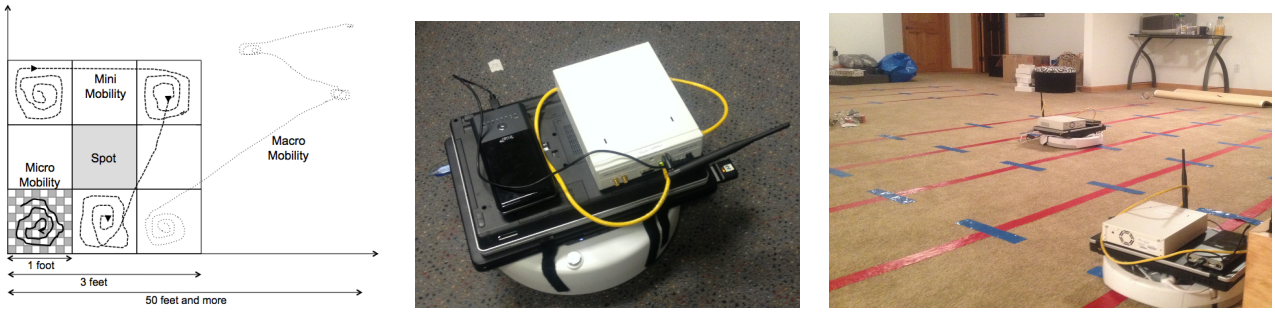

Figure 2: (a) Three regimes of AP mobility. (b) A laptop, USRP, and webcam mounted on a Roomba to emulate a line following AP robot. (c) Measurement in the home – red tapes laid out with blue periodic marks indicating spot locations.

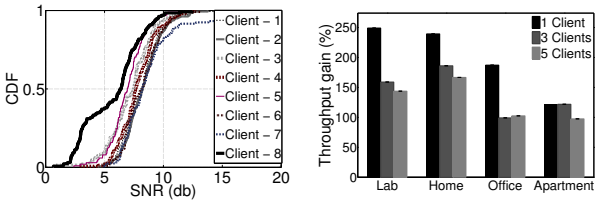

Figure 4: SNR variation due to micro-mobility within a spot, averaged across many spots: (a) CDF of (max - median) SNR. (b) Throughput gain with multiple clients and interferers.

(1) Dominantly light colored spots, indicating that client is far away from the AP, have several pixels that are dark. This suggests that it is *possible for an AP to significantly improve SNR to its client with centimeter scale mobility.*

(2) Spots that are dominantly dark, indicating that the client is close to the robot, has several pixels that are light. This suggests that it is *possible for a robot AP to move a little and avoid being interfered by other nearby APs, enabling parallel transmissions.*

Figures 4(a) reports the statistics from all spots for each client in all settings (Lab, Office, Apartment, Home). Figure 4(a) shows the CDF of the difference between maximum and median SNR from each spot. On average, this difference is at around 8 dB, implying that on any given spot, antenna mobility should offer appreciable gains to a client. The CDF of the difference between median and minimum SNR from each spot is also similar. For around 20% of the cases, the interfering signal can be suppressed by around 10 dB, just by moving the robot antenna to the pixel with minimum SNR.

To reason about throughput gains, we convert SNR to throughput using Shannon's equation (of course, this produces the upper bound and the protocol overheads will certainly diminish gains). We discuss them next.

Satisfying Multiple Clients

Figure 4(b) plots the throughput gains for different client densities and different settings – in a large home, with interferences from 3 neighboring APs, the throughput gain can be up to 150% for 5 active clients. Observe that 5 simultaneous clients is reasonably high density, since in reality, not all clients are active at the same time. If they are, we could optimize for the throughput hungry clients and still achieve substantial spectrum savings [15, 16].

3.2 Upper Bounds with Mini-Mobility

While micro-mobility is within one spot, recall that with mini-mobility APs have a longer leash (i.e., it can move to adjacent spots). Figures 5(a) and 5(b) plot the throughput gains without and with interference for varying number of clients under Mini-Mobility. With 5 active clients and co-channel interference (from 3 surrounding interferers), the throughput can increase to around 2.5x.

Channel Non-Monotonicity

Intuitively, it might seem that the client needs to be brought very close to AP for a 3x gain. Mini-mobility seems to be suggesting that this is not necessarily true, instead, *carefully searching for a good nearby pixel may be comparable to blindly moving close to the client.* We believe this could be a valuable intuition.The core opportunity arises from the fact that the indoor wireless channel has non-monotonicity, that is, *some* far away locations can be strong and *some* nearby ones can be weak.

To quantify this, we perform the following experiment on our measurement data. We position the AP at a random pixel P near the client – let's say the SNR at the client from this AP is S_P dB. We now scan all spots in the entire building and pick the maximum SNR pixel from them, say X_i, and check whether this maximum SNR is within 90% of S_P. If so, we draw a line joining P and X_i. Longer the line, stronger is the evidence of this opportunity. Figure 5(c) visualizes the scenario, corroborating the intuition that carefully searching local pixels can be as effective as blindly moving close to the client.

Dense AP Deployment

AP mobility is complimentary to AP density. We believe robot APs can be very cheap, and hence, all installed APs can be mobile. Thus, we compare between two schemes:

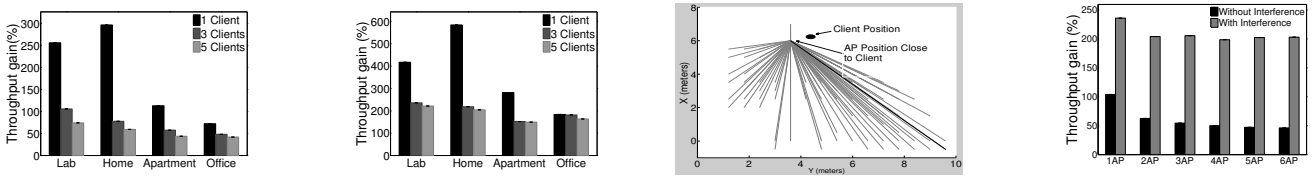

Figure 5: Mini Mobility: (a) Throughput gain w/o interference (b) With interference (c) Pixels of comparable SNRs connected with a line. Carefully chosen far away pixels offer strong SNRs. (d) Gains by adding mini-mobility to dense APs.

(1) K static APs installed at realistic locations and scattered clients associating to the strongest AP at any given time; (2) the same K APs but each AP capable of mini-mobility. Figure 5(d) shows that throughput gains from min-mobility can be up to $2x$ with interference (and less without). Injecting mobility to a high density system can still be useful, so long as the density is not extremely high.

3.3 Upper Bounds with Macro-Mobility

The benefits of macro-mobility arise not only from SNR improvement but also from avoiding large interference zones (recall that APs can move as much needed in this case). Interference footprints from surrounding APs may overlap to form many zones, each zone defined as a region in which a unique set of interfering APs are audible (Figure 6). With Macro-Mobility, APs are not only able to transition from stronger to weaker zones, but also control the footprint. The flexibility (to adapt to time varying interference zones) of mobility outperforms density.

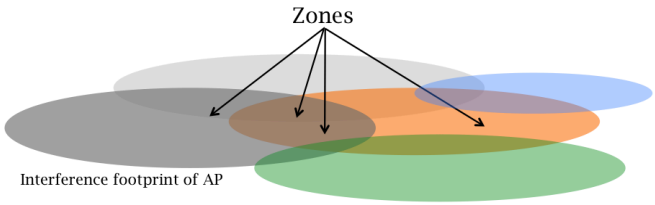

Figure 6: Zones created by overlapping interference footprints from surrounding APs.

Figure 7 quantifies this benefit by comparing a static AP (placed at the median pixel of the median spot) and a macro-mobile AP that starts from the same pixel, but moves to the best spot and pixel for the current network environment (i.e., macro-mobility includes mini and micro mobility). We note that fairness is an important criteria in macro-mobility, since it is possible to move very close to one client and super-optimize throughput. For this, we allow the mobile AP to optimize for throughput so long as the fairness (i.e., Jain's Fairness Index) is at least as much as the static AP. Still, the median improvements are $3x$ or more, as evident from Figure 7.

3.4 Validation on Atheros WiFi Cards

While USRP results help characterize upper bounds without protocol overheads, we were curious to investigate upper bounds with real cards too. We perform similar micro-mobility experiments with real laptops. Figure 8(a) shows

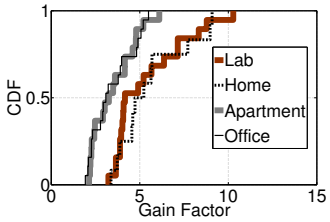

Figure 7: Throughput gains from Macro-mobility

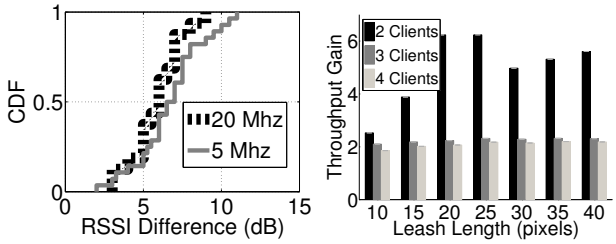

Figure 8: (a) RSSI variation over pixels due to micro-mobility in a spot (b) Throughput Gains over clients and leash size

the CDF of difference between maximum and median SNR in a spot. Up to 7dB improvement (median) is possible across both 5MHz and 20MHz bandwidths. Figure 8(b) shows the translation of SNR improvements into throughput. Gains extend over multiple clients and over various leash sizes of mobility (Each pixel roughly contributes $\frac{1}{40}^{th}$ to the spot of 1 sq feet size)

3.5 Coping with Temporal Fluctuations

The AP can adapt to changing locations of clients and traffic. However, it is complex to adapt if the environmental dynamism like human mobility induces fast channel changes. Our data shows that channel perturbations due to such dynamism is temporary. Figure 9(a) shows a 5 minute snapshot of a SNR trace. The signal is mostly constrained within 3db band around the median. Also, the channel reverts back to previous conditions (also observed in [17]) when the temporary environmental change disappears (e.g., human walking past the receiver at 1.5 minute in the graph). Figure 9(b) plots the fraction of time, the deviation was less than 3dB from recent (5 minute) median. Evidently, fluctuations are bursty, implying that AP can adapt in long time-scales. Permanent changes (Ex. big metal cabinet moved), happened at longer timescales.

4. MANY OPEN ISSUES, QUESTIONS

Needless to say, this paper is a small step towards the

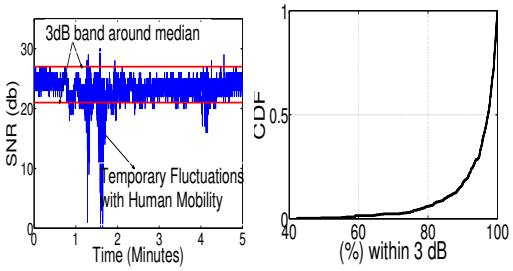

Figure 9: (a) A 5 Minute SNR trace. (b) Channel Resilience and Persistence with environmental dynamism

broader vision and much work remains as discussed here.

Searching and relocating to a new pixel: To minimize overhead, the AP would have to continue its normal operation (Tx and Rx) while it is moving and searching. This is sub-optimal since some transmissions will be from weak pixels. To minimize this, we should be able to predict the channel model and quickly move the AP in the direction of the best pixel. Further work is needed here. Overlapping search with idle time of the AP will also be explored.

Serving many clients with higher throughput may not be feasible (i.e., no single pixel may improve every client in a classroom). Our vision is to serve a subset of clients that are in greater need (e.g., high-bandwidth applications). Based on $80-20$ rule [15, 16], we would target these 20% users to generate huge spectrum savings, without degrading the quality of the other 80% users.

MIMO, beamforming, and other channel pre-coding techniques may achieve gains comparable to mobility. We do not believe infrastructure mobility (IM) is an alternative to these techniques, rather should be viewed as complimentary to them. In fact, we can design utility functions to optimize MIMO channels by the mobile AP.

Classical Problems: Motion adds a new dimension to problems of power control, channel allocation and cloud based network management. Accuracy of signal strength based localization systems can be improved by combining channel profiles from various AP locations and mitigating the multipath. Channel based security protocols can be made robust. With static APs, the channel changes over long time scales, hence secrecy rate is smaller. However, with mobile APs, channel changes can be accelerated by motion and hence increasing the rate of secrecy.

Cellular Networks: Perhaps infrastructure mobility will play a greater role in cellular networks, where high quality network coverage is a bigger problem and over-provisioning is prohibitively expensive. Our ongoing work is exploring a futuristic vision of quadcopters flying out from cell towers to address specific network needs in certain locations. The quadcopters may temporarily park and relay traffic to cell towers. Our initial measurements show that high altitude channels are far stronger than ground channels, enabling long distance communication between the quadcopter and the tower.

5. RELATED WORK

Robotics and wireless networks are rich, mature fields, however, the intersection of them is relatively less explored. Inherent spatial diversity has been opportunistically exploited in [18, 19]. In contrast, we create the opportunity by mobility and quantify the diversity with microbenchmarks. The work closest to this proposal is *MoMiMo* [6], where the receiver adjusts its antenna in centimeter scales to actively perform interference alignment. While *MoMiMo* is a specific scheme optimizing for a given user, this paper is a generalization to a broader architecture. In some sense, MoMiMo may be viewed as nano-mobility, while we extend the spectrum to various mobility regimes and shed light on their behaviors. Further, our schemes are complimentary since *MoMiMo* can be fully used along with mobile APs. Loon [20] provides Internet access to remote areas via an ad hoc network–style balloons drifting above the stratosphere. In robotics, [21, 22] cooperative robots achieve a common wireless communication goal. In one instance, robots plan their motion paths in a distributed manner to form constructive beamforming towards a specified receiver. Authors in [23] have envisioned on-the-fly robots forming a "chain" while first responders (e.g., fire fighters) move into a catastrophe stricken building. Our proposal, in contrast, brings a sense of "control" on persistent infrastructure, with the goal of better serving existing clients. Finally, beyond wireless networking, [24] is an example where the camera rotates and focuses on the user as she moves during a video conference. This is a creative form of infrastructure mobility that we believe is extendable to networking infrastructure as well.

6. CONCLUSION

We explore robotic wireless networks, as a new bridge between robotics and wireless networking. We present this proposal in the context of WiFi alone, but the core ideas to certainly be generalized to other networks and infrastructure. 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 committing serious research attention to it, we believe, is worthwhile.

7. ACKNOWLEDGMENT

We sincerely thank the anonymous reviewers for their valuable feedback on this paper. We are also grateful to NSF and Huawei. The research was partially funded by NSF (through grants – CNS-1441638 and CNS-1430033) and Huawei.

8. REFERENCES

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

[2] D. Raychaudhuri *et al.*, "Mobilityfirst: a robust and trustworthy mobility-centric architecture for the future internet," *ACM SIGMOBILE MCCR*, 2012.

[3] M. Jain *et al.*, "Practical, real-time, full duplex wireless," in *ACM Mobicom*, 2011.

[4] R. Ramanathan *et al.*, "Topology control of multihop wireless networks using transmit power adjustment," in *Infocom*, 2000.

[5] H. V. Balan *et al.*, "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] "Robots are becoming ready to work among us." https://openslam.org/.

[10] 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.

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

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

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

[14] "USRPN210 - Software Radio." `https://www.ettus.com/product/details/UN210-KIT`.

[15] "Top 1 percent of mobile users consume half of worlds bandwidth, and gap is growing." http://www.nytimes.com/2012/01/06/technology/top-1-of-mobile-users-use-half-of-worlds-wireless-bandwidth.html.

[16] "The law of vital few - 80-20 rule." http://en.wikipedia.org/wiki/Pareto_principle.

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

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

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

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

[21] M. Zavlanos *et al.*, "Network integrity in mobile robotic networks," in *IEEE Transactions on Automatic Control*, 2013.

[22] N. Chatzipanagiotis *et al.*, "Controlling groups of mobile beamformers," in *IEEE CDC*, 2012.

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

[24] "Robotic camera on balls." `http://www.gosphero.com/`.