

Adaptive Communication in Multi-Robot Systems Using Directionality of Signal Strength

Stephanie Gil, Swarun Kumar, Dina Katabi, Daniela Rus
 Massachusetts Institute of Technology, MA.
 {sgil,swarun,dk,rus}@mit.edu

Abstract— We consider the problem of satisfying communication demands in a multi-agent system where several robots cooperate on a task and a fixed subset of the agents act as mobile routers. Our goal is to position the team of robotic routers to provide communication coverage to the remaining client robots. We allow for dynamic environments and variable client demands, thus necessitating an adaptive solution. We present an innovative method that calculates a mapping between a robot’s current position and the signal strength that it receives along *each spatial direction*, for its wireless links to every other robot. We show that this information can be used to design a simple positional controller that retains a quadratic structure, while adapting to wireless signals in real-world environments. Notably, our approach does not necessitate stochastic sampling along directions that are counter-productive to the overall coordination goal, nor does it require exact client positions, or a known map of the environment.

I. INTRODUCTION

There are many projects on today’s frontier that are pushing the capabilities of multi-agent systems. Swarm robotic systems perform many complex tasks through coordination, such as cooperative search of an environment, consensus, rendezvous, and formation control [Cortes et al., 2004; Jadbabaie et al., 2003; Olfati-Saber et al., 2007]. Google’s *Project Loon*, Facebook’s *Connectivity Lab*, and similar projects envision using a network of controllable routers to provide wireless communication infrastructure in remote areas of the world. At their core, these systems rely on coordination between agents [Cortes et al., 2004; Moreau, 2004; Spanos and Murray, 2005; Tahbaz-Salehi and Jadbabaie, 2007], making reliable communication of primary importance. Beyond simply maintaining connectivity, reliable communication may mean supporting heterogeneous and possibly time-varying communication rates amongst different pairs of agents. For example, some agents may need to use the network for transmitting video while others may simply wish to transmit status information.

We focus on problems where an auxiliary team of robot routers can be deployed to establish reliable wireless communication to a team of client agents who are performing an independent task. As depicted in Figure 1, we wish to control the positions of robot routers to establish communication links that are capable of supporting variable demanded rates to the client agents. The problem of providing wireless communication coverage amongst multi-robot systems requires tight feedback between spatial positioning of the robots and sensing of the communication quality. The richer the information on signal quality, the more effective the control. A key realization makes this problem very challenging: Robotic tasks leverage mobility in Euclidean space and thus require knowledge of how position effects communication. However, the relationship of signal quality with spatial position is notoriously hard to

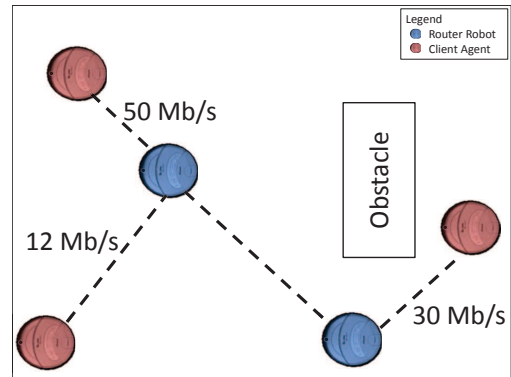


Fig. 1. Picture of a network of two robot routers satisfying the demands of three clients in an environment with an occluding obstacle whose position is unknown.

predict due to complex interactions with the environment such as multipath, where the signal is reflected and/or attenuated by multiple objects in the environment before arriving at a receiver [Goldsmith, 2005; Lindhe et al., 2007; Malmirchegini and Mostofi, 2012]. Past literature employs two broad strategies to address this challenge. On the one hand, there is the Euclidean disk model which assumes that the signal quality of a link is a function of distance between the communicating vehicles. This model is deterministic and simple, and hence when incorporated in a robotic controller, yields simple positional optimizations for a wide range of collaborative tasks [Cortes et al., 2004; Jadbabaie et al., 2003; Olfati-Saber et al., 2007]. Unfortunately, the Euclidean model is too simplistic and fails to represent wireless signals in realistic environments [Malmirchegini and Mostofi, 2012]. On the other hand, there are stochastic sampling methods [Fink et al., 2012; Malmirchegini and Mostofi, 2012; Yan and Mostofi, 2013a] that measure the wireless signal strength in a robot’s vicinity to fit parameters for intricate probabilistic communication models. While such methods are not oblivious to wireless channels, they require exploratory sampling [Lindhé and Johansson, 2010] along directions that may be counter-productive to the overall coordination goal. Further, they often assume the knowledge of parameters based on the structure and material composition of the environment.

Our objective is to i) present a novel method to capture the spatial variation of wireless signals in the local environment *without* sampling along counter-productive directions, or requiring information about the environment or the channel’s distributions and ii) derive a control formulation that maintains the structural (quadratic) simplicity allowed by the Euclidean disk model while accounting for wireless channel feedback.

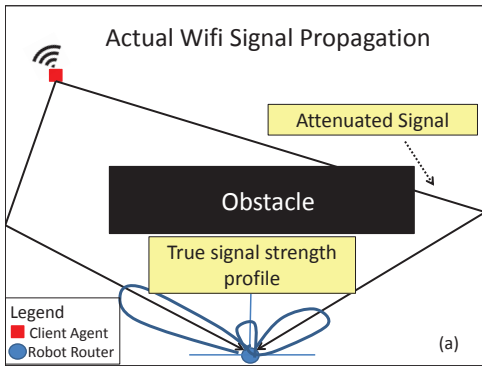


Fig. 2. Schematic drawing of a true signal strength profile in the local environment of a robotic router. Large lobes indicate directions of high signal strength.

First, we introduce an innovative approach for mapping communication quality to robot placement. We calculate a mapping between a robot’s current position and the signal strength that it receives along *each spatial direction*, for every wireless link with other robots (see Figure 2). This is in contrast to existing methods [Fink et al., 2012; Yan and Mostofi, 2013a], which compute an aggregate signal power at each position but cannot distinguish the amount of signal power received from each spatial direction. Our approach combines the best attributes of both the Euclidean disk model and the stochastic sampling methods: Like the disk model, we can compute our mapping without knowledge of the environment and its obstacles, or a model of the channel’s distribution. Like the stochastic methods, our approach uses feedback from the actual wireless signals and hence can help multi-robot systems satisfy their desired communication demands in a real-world implementation. A naive approach to achieve this would be to mount directional antennas atop the routers; but these antennas are bulky and prohibitive for small agile platforms [Networks, 2014]. Instead, we present a novel algorithm based on Synthetic Aperture Radar (SAR) [Fitch, 1988], where a single omnidirectional antenna emulates a high-resolution directional antenna. This paper presents the first such algorithm for implementing SAR using off-the-shelf wireless cards in a non-radar setting, a challenging task since these devices are not intended for this purpose.

Second, we construct an optimization for positioning a team of robot routers to provide communication coverage over client vehicles using the directional information provided by our mapping. Being able to measure the profile of signal strength across spatial directions in real-time yields a much more capable controller. For example, the direction that improves signal strength the most is immediately attainable from these profiles (see Figure 2 for a schematic interpretation). Therefore as a direct consequence, the controller has access to the gradient of communication quality for each of its “links”, or neighbors to which it communicates. While in the favorable scenario, there is a single recommended direction of movement, in real-world implementations it is possible for there to be multiple such directions due to multipath or even noise that may be affecting the wireless link. This information is important for gauging the confidence with which the controller can improve signal quality by navigating the robot along any of the recommended directions. To this end, we present a method for computing a confidence metric

from the data and show that this metric can accurately and automatically identify the three scenarios of strong single-peak, multi-peak, or noisy peak in actual experimental data. Our control algorithm leverages the gradient directions and their associated confidences to automatically tune the speed of the robot, improving both stability and convergence time. Finally, our controller optimizes communication with multiple robots by choosing a direction of movement corresponding to a strong signal that strikes trade-offs between competing demands.

The result of a tight integration between our wireless signal quality mapping and positional controller yields algorithms for router placements that do not rely on environment-dependent parameters, obstacle maps, or even client positions. The overall solution presented is adaptive to variable communication quality demands by the clients, as well as changes in the wireless channels due to natural fluctuations or a dynamic environment.

We implement our method in a multi-robot testbed that has two robotic routers serving three robotic clients. We conduct our experiments in different indoor environments without providing the robotic controller the environment map or the clients’ positions. We observe the following: 1) Our system consistently positions the robotic routers to satisfy the robotic client demands, while adapting to changes in the environment and fluctuations in the wireless channels; 2) Compared to the disk model [Cortes et al., 2004; Jadbabaie et al., 2003] and the stochastic approach [Le Ny et al., 2012; Spall, 2000] under identical settings, our system converges to accurately satisfy the communication demands, unlike the disk model, while significantly out-performing the stochastic method in terms of empirical convergence rate (see Fig. 15 in Sec. VI-E).

A. Contributions

We present a method to enable a robotic receiver to find the profile of signal strength across spatial directions for each sender of interest. To this end, we perform synthetic aperture radar (SAR) techniques using standard Wi-Fi packets exchanged between two independent nodes with single omnidirectional antennas. We derive a quantitative metric, the confidence, that can accurately and automatically identify the presence of multipath or noise for each communication link. This provides valuable information to the controller in gauging the effectiveness of each recommended direction of movement in improving communication quality. We develop an optimization that leverages the directional signal profiles and their confidences, to position robotic routers to satisfy heterogeneous (and possibly variable) communication demands of a network of robotic clients, while adapting to real-time environmental changes. Finally, we provide aggregate empirical data to show that our method outperforms existing Euclidean disk or Stochastic sampling methods both in convergence time ($3.4\times$ faster) and variability of performance ($4\times$ smaller variance).

II. RELATED WORK

Related work falls under two broad categories.

A. Multi-Robot Coordination

Our work is related to past papers on multi-robot coordination to achieve a collaborative task while supporting specific

communication demands [Fink et al., 2012; Le Ny et al., 2012; Malmirchegini and Mostofi, 2012; Yan and Mostofi, 2013a]. Past work on this topic fall under two classes of approaches.

Euclidean Disk Model: The first class employs Euclidean disk assumptions where signal quality is assumed to be deterministic and mapped perfectly to the Euclidean distance between the communication nodes. A Euclidean metric allows for quadratic cost for the edges of the network and enables a geometric treatment of an otherwise complex problem. In reality, signal strength suffers from large variations over small displacements [Goldsmith, 2005; Lindhe et al., 2007] that these models simply do not capture. Yet, the simplicity afforded by these models has led to significant contributions including i) multi-agent coordination for coverage and flocking [Martinez et al., 2007; Schuresko and Cortes, 2009], ii) assignment of routers to clients for attaining a prescribed level of connectivity [Feldman et al., 2013; Gil et al., 2012] or throughput [Craparo et al., 2011], and iii) connectivity maintenance based on graph theoretic approaches [De Gennaro and Jadbabaie, 2006; Michael et al., 2009].

Stochastic Sampling Methods: Recently, efforts have focused on giving the communication quality over each link in the network a more realistic treatment by sampling the signal strength and building closed-loop controllers using this feedback. Such stochastic sampling methods either supplement theoretical models for signal strength with a stochastic component based on the collected samples [Lindhe et al., 2007; Malmirchegini and Mostofi, 2012], or, use the collected samples to design stochastic gradient controllers [Le Ny et al., 2012; Twigg et al., 2013]. These papers have studied stochastic sampling patterns for i) acquiring sufficient signal strength (RSS) samples [Lindhe et al., 2007; Lindhé and Johansson, 2010], ii) co-optimizing communication quality and other higher level tasks like motion planning or message routing [Fink et al., 2013; Yan and Mostofi, 2013b], iii) used router mobility to escape “deep fades” or null points where connectivity may be lost [Vieira et al., 2013] or to map out the signal strength and resulting connectivity regions of the environment [Twigg et al., 2013]. Unfortunately these works necessitate at least one of the following prohibitive requirements: i) motion of the routers along counter-productive paths to collect sufficient RSS samples, ii) assumptions of a known environment map, static surroundings, and known positions of communicating agents, or iii) previously acquired signal strength maps.

In comparison to these papers, we introduce a system that captures the magnitude of the signal arriving from different *directions*, as opposed to only its total magnitude at a particular position. This allows us to combine the best of both the disk model and stochastic sampling methods: Like the disk model, we do not require prior knowledge of the environment and its obstacles, or a model of channel’s distribution. Like the stochastic methods, our approach accurately captures actual signal characteristics and hence can help multi-robot systems satisfy their desired communication demands in real-world environments. Figure 3 provides an illustrative example of how our system out-performs the Euclidean disk model and stochastic sampling, particularly in the presence of obstacles.

B. Angle of Arrival Systems

Our method builds upon a rich body of literature in wireless networking that estimates actual angle-of-arrival of each of the reflected paths of a signal at a receiving device. Past work has employed two classes of hardware to estimate angle-of-arrival:

Antenna Arrays: Past literature has leveraged arrays of antennas to estimate angle-of-arrival for localization [Joshi et al., 2013; Wang and Katabi, 2013; Xiong and Jamieson, 2013] and tracking [Pham and Sadler, 1997]. These use stationary multi-antenna receivers to locate the transmitter with sub-meter accuracy. Unfortunate for the robotics community, many of these techniques require bulky, specialized hardware such as customized software radios, and are thus difficult to place on small, agile, mobile platforms that are ubiquitous for robotics applications.

Synthetic Aperture Radar: Understanding how to attain this directional information using a moving platform is the subject of Synthetic Aperture Radar (SAR) [Fitch, 1988]. SAR allows even a single-antenna mounted on a flying aircraft or satellite to emulate a multi-antenna array. Unfortunately, most SAR applications [Fitch, 1988; Wang et al., 2013; Wang and Katabi, 2013] are geared towards radar-type problems (eg. imaging, RFID applications) where signals are transmitted and processed by the same node. Therefore, they cannot be used to analyze the direction of arrival of the signal from a distinct transmitter (e.g. Wi-Fi devices).

For an adaptive communication network of small router robots, we need a light-weight, single-antenna system that can perform SAR using two-way transmissions (unlike radar) on off-the-shelf Wi-Fi devices. In this regard, we develop a system that builds upon synthetic aperture radar meant for robotic routers and clients equipped with standard Wi-Fi cards.

C. Organization

Section III presents a formulation of the router placement problem for achieving communication coverage for clients with heterogeneous demands. The following sections describe each component of our solution to the problem:

- Section IV derives a new method for measuring rich directional information from wireless channel feedback.
- Section V presents an algorithm for finding a configuration of routers that *balances* the network, i.e. maximizes the signal quality of the weakest link for a fair network.
- In Section V-B we derive a confidence metric using channel feedback, that captures the effects of multipath and noise.

Finally, Section VI experimentally evaluates our approach against the disk model and stochastic sampling methods.

III. PROBLEM STATEMENT

We consider a mobile network with two classes of members, n robotic clients (or *clients*) whose positions are not controlled, and a team of k robotic routers whose mobility we control. Our goal is to position the robotic routers to provide adaptive wireless communication coverage to the clients, while allowing variable communication quality demands for all clients, and where exact client positions are unknown. For each client $j \in [n] = \{1, \dots, n\}$, we define demanded communication quality $q_j > 0$, and achieved communication quality p_{ij} to each router i (where $i \in [k]$), both expressed in terms of *Effective*

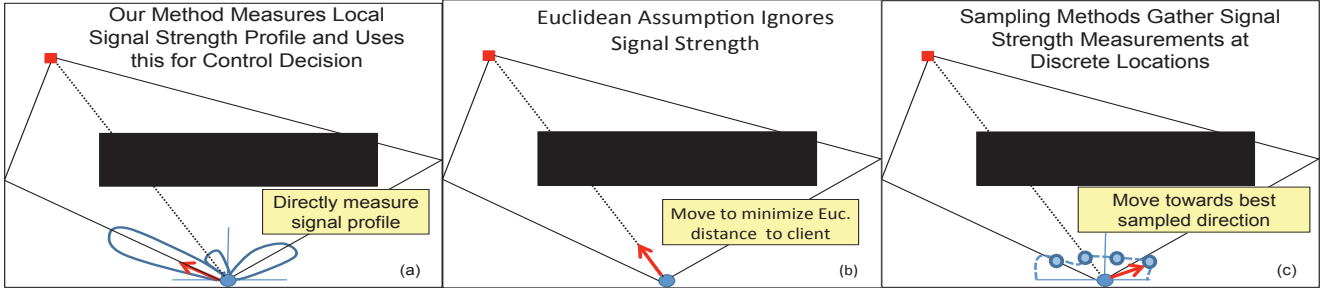


Fig. 3. Compares our method against the Euclidean disk model and a stochastic sampling. The figures depict the actual router (blue) and client (red) separated by an obstacle (black). The black lines indicate the different paths of the signal. (a) Our method estimates the actual signal power arriving from different angular directions, much like a high-resolution directional antenna would. This provides a sharpened peak in the direction of maximum signal strength. (b) The Euclidean Disk model guides the router along the shortest Euclidean path, which is greatly attenuated by the obstacle. (c) Stochastic methods measure the signal strength by moving the router and sampling at various positions (blue circles). The signal strength does not vary significantly between locations due to the lack of spatial resolution, i.e. at each sample location the signal strength is a combination of signals arriving from all angular directions. This leads to much less discernible peaks when contrasted with (a) (Note that the polar plot of signal power, shown here as dotted lines, have peaks in the same angular directions as our method though less sharp). This method guides the router towards the direction of the best sample, which often may not be the actual direction of maximum signal strength, as shown.

Signal to Noise Ratio (ESNR) that has a direct mapping to rate in Mb/s [Halperin et al., 2010].¹ Additionally, let every client j be given an importance $\alpha_j > 0$. We allow all quantities in this section (i.e. q_j, ρ_{ij}, α_j) to be time dependent though we omit this dependency henceforth for simplicity.

We define the notion of *service discrepancy* for each pair of robots (i, j) to be the difference between the demanded and achieved communication quality scaled by the importance of the client.

$$w_{ij} = \max(\alpha_j(q_j - \rho_{ij})/q_j, 0) \quad (1)$$

Physically, this is the fraction of the client's communication demand that remains to be satisfied, scaled by α_j . Denote by $c_i \in \mathbb{R}^d$ the position of the i th robot router and by $p_j \in \mathbb{R}^d$ the position of the j th client and $C_t = \{c_{1,t}, \dots, c_{k,t}\}$ is the set of all router positions at time t . In this paper we give explicit treatment to the case for $d = 2$ although all concepts are extensible to $d = 3$.

A. Problem Formulation

Given a cost g in terms of signal quality, communication demands, and agent positions, we wish to position each robotic router to minimize the largest discrepancy of service between routers and clients. However, the true form of this function g has an intricate dependence on the position of the client, router, and the environment. Thus an inherent challenge to solving this problem is approximating the influence of spatial positioning on communication quality that generalizes across environments. Our goal is to 1) find $f_{ij} : [-\frac{\pi}{2}, \frac{\pi}{2}] \rightarrow \mathbb{R}$ (a relation capturing directional information about the signal quality between i and j), and an approximation \tilde{g} of g , which is a cost, characterizing the anticipated communication quality for the router-client pair (i, j) at a proposed router position c_i , and 2) use this cost to optimize router positions to minimize the service discrepancy to each client. Formally,

Problem 1: Find i) a mapping $f_{ij} : [-\frac{\pi}{2}, \frac{\pi}{2}] \rightarrow \mathbb{R}$ that maps spatial direction to wireless signal strength directly from

channel measurements, and ii) a cost

$$\tilde{g}(c_i, C_t, w_{ij}, f_{ij}) > 0 \quad (2)$$

that is independent of the environment and satisfies the following properties:

Property 1: All link costs \tilde{g} are quadratic

Property 2: Minimization of a link cost \tilde{g} over c_i directly relates to increasing signal quality for client j and optimization over all link costs \tilde{g} allows trade-offs between clients with competing demands

Property 3: The link costs \tilde{g} are independent of client positions p_j

Given a known number k of routers, client demands q_d , and the mapping f_{ij} for all links in the network, position routers to minimize the worst-case link. Specifically, we aim to find a position for the routers that minimizes the maximum service discrepancy by solving for C in the following problem:

$$C_{t+1} = \arg \min_{c_i \in C} \{ \max_j \min_i \tilde{g}(c_i, C_t, w_{ij}, f_{ij}) \} \quad (3)$$

Intuitively, the solution to this optimization problem favors a “fair” network. Specifically, the solution aims to minimize the “worst service discrepancy” among clients in the network, at any point in time. The worst service discrepancy is given mathematically by the bracketed expression in Equation (3) and can be understood intuitively as follows: 1) The service discrepancy of a router-client link captures the difference between the measured quality of the link and the client's demanded communication quality (see Equation (1)). 2) Each client is served by a router that offers the minimum service discrepancy to it, at any given time (the innermost min over i in Equation (3) above). 3) The worst service discrepancy, is the maximum service discrepancy among all clients to their chosen routers in the network (the inner max over j in Equation (3)). Notice that the client with the worst service discrepancy may change at any point in time, depending on the configuration of the routers. The optimal configuration of the routers (given by the outer argmin term), is therefore the configuration that best satisfies communication demands across the entire network.

¹ESNR is a continuous signal quality measure that has a one-to-one mapping to the maximum data rate supported by a link [Halperin et al., 2010]. We work with ESNR values rather than rates since the latter are discretized (non-continuous).

B. Problem Scope

We specify that our aim in this paper is to position mobile routers to establish a communication network whose links have high enough ESNR to support given client communication demands, q_d . In other words, we are interested in providing the *infrastructure* to support the requested quality of communication. This is in contrast to solving for routing protocols that would optimize the communication traffic over the infrastructure to ensure successful message passing from a sending node to a receiving node. While this is another common metric of connectivity, it is often times treated as a layer on top of an existing communication infrastructure and is an out-of-scope problem with a vast body of dedicated literature (See [Fink et al., 2010] for an example of routing in robotic networks). Finally, we assume throughout that router-router links are high capacity and that router-client links are the limiting factor that must be optimized.

We dedicate the next sections of this paper to 1) Developing a method that computes f_{ij} as the profile of signal qualities along each direction θ for each link (i, j) found directly from channel measurements; and 2) Developing an optimization framework that utilizes this directional information to handle trade-offs between competing client demands, and position all routers to jointly minimize the maximum service discrepancy across links in the network.

IV. DIRECTIONAL POWER PROFILE OF A WIRELESS LINK

In this section, we develop the first component of the solution of Problem 1; namely, we derive a method to calculate $f(\theta)$, the mapping to capture the signal strength from a robotic client to its router along each direction θ , where this mapping can be updated often, roughly once every 6cm of motion.²

Before we explain how we compute $f(\theta)$, we describe this function to help understand what it captures. Assume we have a robotic client and router, where the router moves along some trajectory. We will define the direction θ relative to the tangent to the router's trajectory at each point. Consider the scenario in Fig.4(a), where the robotic client is in line-of-sight at -50° relative to the robotic router, which is moving along the horizontal axis. In this case, one would expect $f(\theta)$ to have a single dominant peak at -50° , as shown in Fig.4(b). Now consider the more complex scenario in Fig.4(c), where the environment has some obstacles and one of these obstacles obstructs the line-of-sight path between the router and its client. In this case, $f(\theta)$ would show two dominant peaks at 20° and -30° that correspond to the two reflected paths from surrounding obstacles, as shown in Fig.4(d).

Advantage over Sampling Methods: One may estimate $f(\theta)$ by sampling the signal power similar to stochastic techniques [Le Ny et al., 2012; Spall, 2000; Yan and Mostofi, 2013a]. In this case, one has to move the router along each direction, compute the power in all these new positions relative to the first, and draw the profile $f(\theta)$. Unfortunately, this approach leads to much wasted exploration. This is because the signal power does not change reliably when the robot moves. For example, if the robot moves for 5 or 10 centimeters, it is very likely that the resulting change in the signal power is below the variability in noise. Hence, measurements of

power over short distances are likely to be marred by noise or phenomena that affect the signal strength locally such as deep fades [Tse and Viswanath, 2005] (due to reflections of the signal at the receiver interfering constructively or destructively). To obtain reliable measurements of changes in the signal power, the robot has to move significantly along potentially counter-productive paths.

To address this limitation, our approach relies on the channel phase as opposed to the power. Specifically, at any position the wireless channel can be expressed as a complex number $h(t)$ [Rahul et al., 2012]. The magnitude of this complex channel captures the signal power (more accurately, its square-root). The phase of the channel has traditionally been ignored by robotic systems. However, the phase changes rapidly with motion. For Wi-Fi signals at a frequency of 5 GHz, the phase of the channel rotates by π every 3 cm. This far exceeds any rotation due to noise variability. Thus, by measuring channels as complex numbers and tracking changes in its phase as the robot moves, we reliably estimate signal variation without much exploration. In the next section, we explain how to use a technique called synthetic aperture radar (SAR) to extract the received signal strength along each direction from changes in channel phase. Note that SAR does not need exploration in all directions; the robot can move along its path without extra exploration or sampling. SAR uses the resulting variations in channel phase over distances of a few centimeters to find $f(\theta)$.

A. Synthetic Aperture Radar (SAR)

Synthetic Aperture Radar (SAR) enables a single antenna mounted on a mobile device to estimate the strength of the signal received along every spatial direction. As explained in Section II-B, SAR employs a single moving antenna to emulate a multi-antenna array and compute the directional profile of signal strength $f(\theta)$ (See Figure 5). Therefore, we can leverage the natural motion of a robotic router to implement SAR and measure $f(\theta)$ for each of its robotic clients using a single omni-directional antenna. To do so, the robotic router measures the channel $h(t)$ from its client as it moves along any straight line. The straight line path over which the router acquires data is on the order of half a wavelength (centimeters); assuming the source is stationary and the router either moves at a known constant velocity or its position is known for the traversal time window, then a sufficient amount of usable channel data can be collected. This means every few centimeters the router can have an updated measurement of $f(\theta)$, for all values of θ .

Specifically, Let $h(t)$ for $t \in \{t_0, \dots, t_m\}$ be the $m + 1$ most recent channel measurements, corresponding to the robot whose displacement from its initial position is $d(t_0), d(t_1), \dots, d(t_m)$. SAR computes the received signal strength across spatial directions $f(\theta)$ as:

$$f(\theta) = \left| \sum_t h(t) e^{-j \frac{2\pi}{\lambda} d(t) \cos \theta} \right|^2, \quad (4)$$

where λ is the wavelength of the Wi-Fi signal. The analysis of this standard SAR equation may be found in [Stoica and Moses, 2005]. At a high level, the terms $e^{-j \frac{2\pi}{\lambda} d(t) \cos \theta}$ in Eqn. 4 project the channels $h(t)$ along the direction of interest θ by compensating for incremental phase rotations introduced by the robot's movement to any path of the signal arriving along θ .

²For simplicity, we denote $f_{ij}(\theta)$ as $f(\theta)$ as we consider only the single link between robotic router i and client j for the rest of this section.

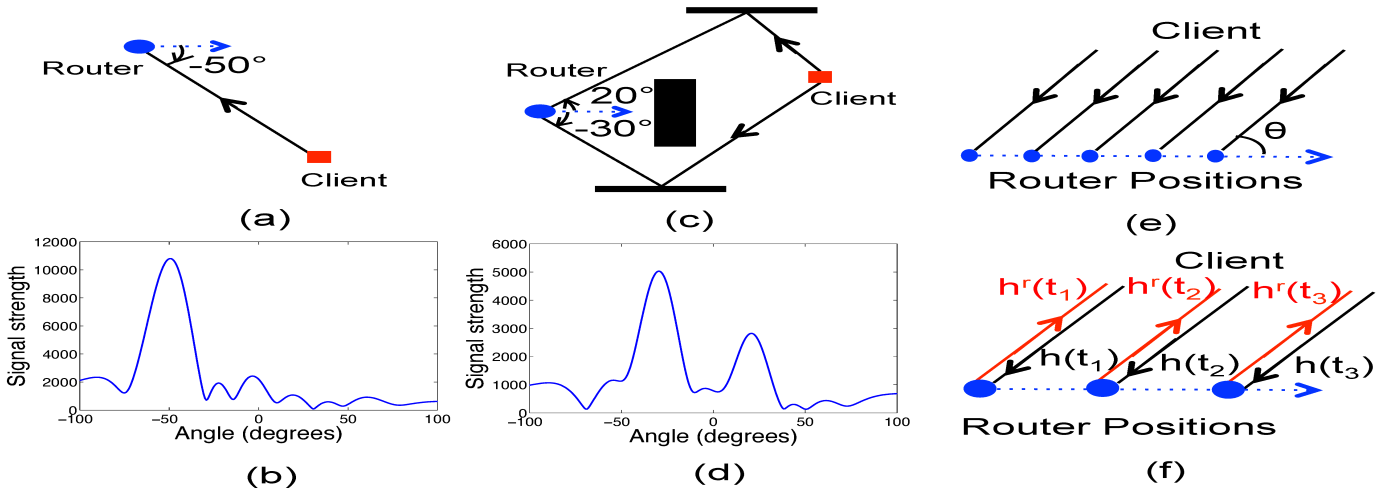


Fig. 4. (a)/(c) LOS and NLOS topologies annotated with signal paths. (b)/(d) $f(\theta)$ of the signal in LOS and NLOS. (e) Shows how θ is defined in SAR. (f) Shows $h(t_i)$, the forward channel from transmitter to receiver and $h^r(t_i)$, the reverse channel from receiver to transmitter at time t_i .




Fig. 5. Schematic representation of our method for emulating a directional antenna array with a single omni-directional antenna attached to a mobile off-the-shelf platform.

Note that SAR finds the signal power from every angle θ simply by measuring the channels, without any prior tuning to the given direction. Of course, the resolution at which θ is available depends on the number of channel measurements. In fact, moving by around a wavelength (about 6 cm) is sufficient to measure the full profile of $f(\theta)$.

Therefore, SAR is a natural choice for autonomous robotic networks since it exploits the mobility of the robots to compute $f(\theta)$. Further, it only requires the robot to move along a small straight line along any arbitrary direction, and does not require it to explore directions counter-productive to the overall coordination goal. Note that SAR requires only the relative position of the robotic router $d(t)$ and both the magnitude and phase of the channel $h(t)$. It does not require the topology of the environment nor the exact location of the transmitter.

B. Algorithm for Performing SAR on Independent Wireless Devices

A key challenge in adapting SAR to multi-robot systems is that all past SAR-based solutions [Adib and Katabi, 2013;

Fitch, 1988; Wang and Katabi, 2013] are for radar-like applications, where a single device transmits a radar signal and receives its reflections off an imaged object, e.g., an airplane. However, in our scenario the transmitter and receiver are completely independent wireless devices (i.e., the robotic client and router, respectively). This means that the transmitter robot and the receiver robot have different frequency oscillators. In practice, there is always a small difference between the frequency of two independent oscillators. Unfortunately, even a small offset Δ_f in the frequency of the oscillators introduces a time varying phase to the wireless channel.

For instance, let $h(t_0), h(t_1), \dots, h(t_m)$ be the actual wireless channel from the robotic client to the robotic router at times t_0, t_1, \dots, t_m . The channel observed by the router from its client $\hat{h}(t_0), \hat{h}(t_1), \dots, \hat{h}(t_m)$ are given by:

$$\hat{h}(t_0) = h(t_0), \quad \hat{h}(t_1) = h(t_1)e^{-2\pi\Delta_f(t_1-t_0)}, \dots, \\ \hat{h}(t_m) = h(t_m)e^{-2\pi\Delta_f(t_m-t_0)}. \quad (5)$$

Hence, the phase of the channels are corrupted by time-varying values due to the frequency offset between the transmitter and the receiver. Fortunately, we can correct for this offset using the well-known concept of channel reciprocity [Rahul et al., 2012]. Specifically, let $h^r(t)$ denote the reverse channel from the robotic router to its client, as shown in Fig. 4(f). Reciprocity states that the ratio of the forward and reverse channels stays constant over time, subject to frequency offset, i.e. $h^r(t) = \gamma h(t)$, where γ is constant. Further, the frequency offset in the reverse direction Δ_f^r is negative of the offset in the forward direction, i.e. $\Delta_f^r = -\Delta_f$. Thus, the observed reverse channels $\hat{h}^r(t_0), \hat{h}^r(t_1), \dots, \hat{h}^r(t_m)$ are given by:

$$\hat{h}^r(t_0) = h^r(t_0), \quad \hat{h}^r(t_1) = h^r(t_1)e^{2\pi\Delta_f(t_1-t_0)}, \dots, \\ \hat{h}^r(t_m) = h^r(t_m)e^{2\pi\Delta_f(t_m-t_0)}. \quad (6)$$

Multiplying Eqn. 5 and 6 and using $h^r(t) = \gamma h(t)$, we have $\hat{h}(t)\hat{h}^r(t) = h(t)h^r(t) = \gamma h(t)^2 \Rightarrow h(t) = \sqrt{\hat{h}(t)\hat{h}^r(t)}/\gamma$. Hence we re-write Eqn. 4 as:

$$f(\theta) = \left| \sum_t \sqrt{\hat{h}(t)\hat{h}^r(t)} e^{-j\frac{2\pi}{\lambda}d(t)\cos\theta} \right|^2, \quad (7)$$

where the constant scaling γ is dropped for simplicity. Hence, to measure $f(\theta)$, the router and client simply need to measure their channels at both ends. In practice, the router and client transmit back-to-back packets with a small gap $\delta \approx 200\mu\text{s}$ to obtain $\hat{h}^r(t+\delta)$ and $\hat{h}(t)$, respectively. The router collects these values and approximates $\hat{h}(t)\hat{h}^r(t)$ as $\hat{h}(t)\hat{h}^r(t+\delta)e^{-j2\Delta_f\delta}$. The router computes this 10 times per second (an overhead of just 0.1%) and obtains θ with a resolution of 1° . Algorithm 1 summarizes our above approach to compute the signal strength profile $f_{i,j}(\theta)$ for a general wireless link (i, j) .

We note the following important points about Algorithm 1: 1) It requires as input the relative displacement of the robot router $d(t)$ from its initial position at $t = 0$. In particular, if the robot moves at a known constant velocity v for the duration of SAR (i.e., corresponding to a total displacement of few cm), the algorithm only requires this velocity v , since it can readily compute the relative displacements as: $d(t) = vt$. 2) While the algorithm requires the client to be static, this requirement is only necessary for the duration that the router performs SAR (i.e., corresponding to a total router displacement of few cm). We note that i) the assumption of static channels is also necessary for stochastic sampling based methods since the channels and (thus sampled signal strengths) change otherwise and must be re-sampled and ii) the time scales are largely different between our proposed method and existing sampling methods; specifically, because our method allows for the attainment of rich channel data after a comparatively short measurement period, changes in the environment can be quickly adapted to.

In the following section, we explain how we leverage the signal strength profiles $f_{i,j}(\theta)$ on each link (i, j) output by Algorithm 1 to control the position of multiple robotic routers to meet the clients' communication demands.

Algorithm 1: Algorithm for finding directional signal strength profile for a wireless link (i, j)

input : Wireless Channels on the forward link $h_{ij}(t)$, and reverse link $h_{ji}^r(t)$ and robotic router's displacement from its initial position $d_i(t)$ at times $t = t_0, \dots, t_m$ on link (i, j)

output: A vector of directional signal strength values $f_{ij} \in \mathbb{R}^l$ for l discrete directionality angles in $[0, \pi]$

- 1 **for** $t \in \{t_0, \dots, t_m\}$ **do**
- 2 $\tilde{h}_{ij}(t) \leftarrow \sqrt{h_{ij}(t)h_{ji}^r(t)}$;
- 3 **end**
- 4 **for** $\theta \in \{0, \frac{\pi}{l-1}, \frac{2\pi}{l-1}, \dots, \pi\}$ **do**
- 5 $f_{ij}(\theta) \leftarrow \left| \sum_t \tilde{h}_{ij}(t) e^{-j\frac{2\pi}{\lambda} d_i(t) \cos\theta} \right|^2$
- 6 **end**

V. COMMUNICATION COVERAGE CONTROLLER

In this section, we target the problem of placing a team of mobile router vehicles at locations such that they provide wireless coverage to client vehicles, each with different communication demands. Specifically, using as input the channel feedback $f_{i,j}(\theta)$ derived in the previous section, we aim to find a function \tilde{g} that can be optimized over router positions such that:

$$C_{t+1} = \arg \min_C \{ \max_j \min_{c_i \in C} \tilde{g}(c_i, C_t, w_{ij}, f_{ij}) \}. \quad (8)$$

Where C_t are current router positions and w_{ij} are the current service discrepancies.

Our focus in this section is to find communication link costs \tilde{g} that have the three desirable properties 1, 2, 3 from Section III.

We show how to capitalize the rich spatial information provided by $f_{i,j}(\theta)$, to derive a cost \tilde{g} for each link possessing these three desired qualities. The resulting cost can then be optimized to complete our objective of robot router placement that best satisfies the communication demands of the clients.

A. A Generalized Distance Metric

We turn attention to the derivation of a quadratic cost whose minimization will improve signal strength. We derive a generalized distance that encodes the direction of steepest descent and the confidence around this direction. We begin with the case where all positions are known and extend to the position independent case in Section V-E.

Consider a single router-client pair (i, j) located at positions (c_i, p_j) . A Euclidean disk model approach similarly assigns distance, in the Euclidean sense, to be the cost of each communication link in the network. However, this disk model approach does not use $f_{i,j}(\theta)$ at all. Instead, it relates improving communication quality between the router and client to reducing the Euclidean distance between them, i.e. edges in the network take the cost $\tilde{g} := \text{dist}(p_j, c_i)$. The appeal of such a cost is in its simple quadratic form that can be easily optimized. Unfortunately, the cost is oblivious to the actual wireless channel at the client and fails to capture the current service discrepancy which can be large even at small distances (say, due to obstacles).

Our system avoids this pitfall, while retaining simplicity, by incorporating real-time channel feedback into a generalized distance metric. Intuitively, we employ a distance metric that effectively "warps" space so that the shortest distance for enabling better communication between two robots is not the straight line path between them, but rather the path along the θ_{\max} , the direction of maximum signal strength from the mapping $f_{i,j}(\theta)$. The advantage of using this distance metric as compared to a Euclidean distance metric becomes clear when an attenuating obstacle blocks the straight line communication path as shown in Figure 6.

Importantly, the recommended heading direction $\vec{v}_{\theta_{\max}}$ may exhibit variation due to noise or multipath on the wireless link. To account for these effects, while not over-fitting to noise, we leverage the entire $f_{i,j}$ signal profile to design a *confidence* metric σ_{ij} in the recommended heading direction. The exact form of the confidence metric is derived in the following section. The purpose of this confidence metric is to incorporate second-order information from $f_{i,j}$ that captures the presence of noise, or multipath, and can be used to alter the behavior of the controller accordingly (see Section V-B). By using a *Mahalanobis* distance metric for assigning costs to each communication edge in the network, we can encode both the recommended heading direction and its confidence. The mathematical definition of the Mahalanobis distance is:

Definition 1 (Mahalanobis Distance): Given a positive definite matrix $M \in \mathbb{R}^{d \times d}$, a vector $x \in \mathbb{R}^d$, and a vector $y \in \mathbb{R}^d$,

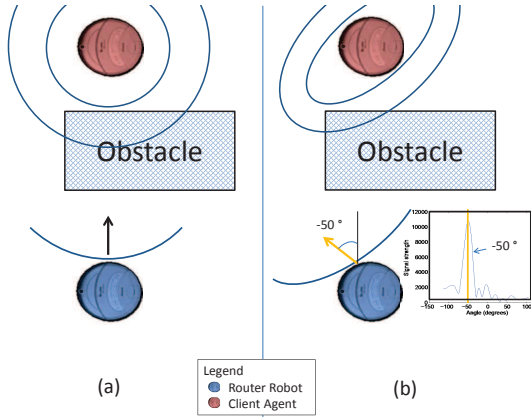


Fig. 6. Schematic depiction of the use of channel feedback for assigning cost to communication links in the network where edge cost is shown as circular contour lines. On the left, a Euclidean distance metric assigned lowest cost to the straight-line direction, whereas using the Mahalanobis distance (right) skews the distance contours to identify the direction about $\vec{v}_{\theta_{\max}}$ as the lowest cost. The amount of skew in the contour lines is determined by the confidence metric derived in Section V-B.

the Mahalanobis Distance between x and y is:

$$\text{dist}_M(x, y) = \sqrt{(x - y)^T M (x - y)} \quad (9)$$

Euclidean distance is a special case of the Mahalanobis distance (see Fig. 9(a)) with $M = I$ where I is the identity matrix of appropriate dimension.

Here, $M = Q\Lambda Q^T$ is a positive-definite matrix, where Q consists of orthogonal eigen-vectors and Λ contains their corresponding eigen-values. By a careful construction of the matrix M , we can encode channel feedback as a quadratic Mahalanobis distance cost for each communication link in the network. This construction requires both the recommended descent direction $\vec{v}_{\theta_{\max}}$ from $f_{ij}(\theta)$, and the confidence metric σ_{ij} that is also computed from $f_{ij}(\theta)$ in the following section.

B. Confidence Metric from Channel Feedback

We design a parameter σ_{ij} that is derived from the mapping $f_{ij}(\theta)$ and that we refer to as a *confidence* in the recommended heading direction $\vec{v}_{\theta_{\max}}$. Intuitively, σ_{ij} captures the “variance” of $f_{ij}(\theta)$ around θ_{\max} . We define σ_{ij} mathematically as the ratio of two quantities, $\sigma_{f_{ij}}$ and $\sigma_{N_{ij}}$. We define

$$F = \sum_{\theta \in \{-\frac{\pi}{2}, \dots, \frac{\pi}{2}\}} f(\theta) \quad (10)$$

$$\sigma_{f_{ij}} = \sum_{\theta \in \{-\frac{\pi}{2}, \dots, \frac{\pi}{2}\}} (\theta - \theta_{\max})^2 \frac{f(\theta)}{F} \quad (11)$$

$$\sigma_{N_{ij}} = \sum_{\theta \in \{-\frac{\pi}{2}, \dots, \frac{\pi}{2}\}} (\theta - \theta_{\max})^2 \frac{F}{L} \quad (12)$$

$$\sigma_{ij} = \frac{\sigma_{f_{ij}}}{\sigma_{N_{ij}}} \quad (13)$$

where L is the total number of θ values that make up the plot $f_{ij}(\theta)$. The term $\sigma_{f_{ij}}$ is the variance of the plot f_{ij} around its maximum $\theta = \theta_{\max}$ and $\sigma_{N_{ij}}$ is a normalization factor (it is the variance around θ_{\max} in the case that the mass under the $f_{ij}(\theta)$ curve was distributed evenly over the θ values). The ratio of these two quantities, $\sigma_{f_{ij}}/\sigma_{N_{ij}}$, characterizes

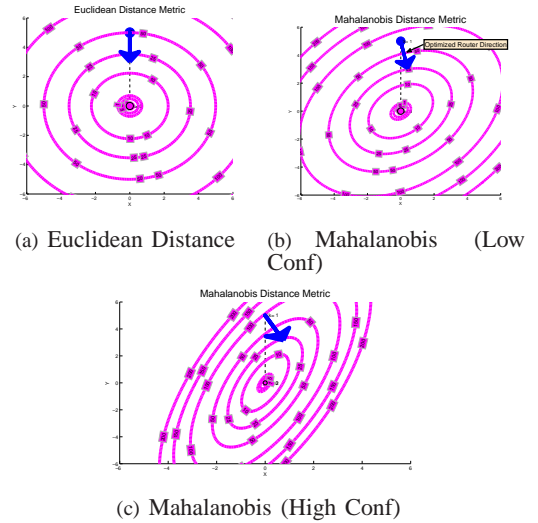


Fig. 9. These plots show the level sets of a Euclidean distance function and a Mahalanobis distance function.

the amount signal strength (mass under the $f_{ij}(\theta)$ curve) that is concentrated under the peak direction θ_{\max} versus the remaining parts of the curve. A ratio of $\sigma_{f_{ij}}/\sigma_{N_{ij}} = 1$ would mean that the $f_{ij}(\theta)$ plot does not provide evidence that the max direction θ_{\max} is of much significance and that indeed the plot is entirely noise. On the other hand a ratio $\sigma_{f_{ij}}/\sigma_{N_{ij}} < 1$ indicates that a significant portion of the signal strength curve in $f_{ij}(\theta)$ is concentrated around the max θ_{\max} and thus this peak is considered to have “high confidence.” Lastly, the case where $\sigma_{f_{ij}}/\sigma_{N_{ij}} > 1$ indicates the presence of high signal strength in other parts of the $f_{ij}(\theta)$ curve other than the θ_{\max} direction which suggests the presence of multipath. These three scenarios are demonstrated empirically in Figure 7 where three actual $f_{ij}(\theta)$ plots are automatically identified as being single peak “high confidence”, multiple peak “noise”, and multiple peak “multipath” scenarios respectively, by computing the ratio σ_{ij} for each plot. The figure demonstrates a graphic depiction of this ratio where areas of the $f_{ij}(\theta)$ plot above and below the uniform variance line determine the confidence value (compare with Equations in (10)).

We define these three cases below for reference:

Definition 2 (Confidence): Confidence in the direction of highest signal strength θ_{\max} . We define three cases captured by our confidence metric $\sigma_{ij} = \frac{\sigma_{f_{ij}}}{\sigma_{N_{ij}}}$:

- **High confidence peak:** $\sigma_{ij} < 1$
- **Noise:** $\sigma_{ij} \approx 1$
- **Multipath:** $\sigma_{ij} > 1$

See Figure 7 for examples of these regions identified automatically from actual experimental data.

Experimental results in the basement of the Stata Center building on the Massachusetts Institute of Technology campus show that the regions of high confidence, noise, and multipath defined above can be identified automatically from data using the confidence metric from Eq. (13) (see Figure 8a). As expected, areas of the environment with no significant occlusions to the client agent show strong evidence of high confidence profiles. Areas such as corridors with potential occlusions due to walls and corners show a much higher incidence of multipath, about 90% in the worst case.

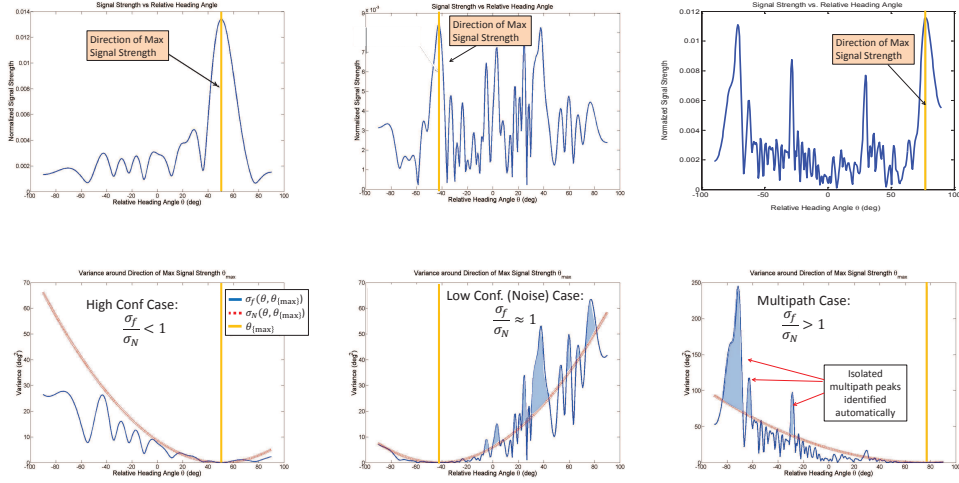


Fig. 7. These plots show directional signal strength profiles from actual experiments. They demonstrate how the confidence metric identifies cases of high confidence, low confidence, and multipath automatically from the f_{ij} signal strength profile. The dotted red line is the variance, σ_{Nij} , of a uniform signal strength profile $f_{Nij}(\theta) = 1/L$ centered around θ_{\max} . Comparing the variance (bottom row) σ_{fij} to σ_{Nij} indicates which of the three cases are occurring in the f_{ij} plot (top row): high confidence θ_{\max} ($\sigma_{fij} < \sigma_{Nij}$), low confidence (noise) θ_{\max} ($\sigma_{fij} \approx \sigma_{Nij}$), or multipath around θ_{\max} ($\sigma_{fij} > \sigma_{Nij}$).

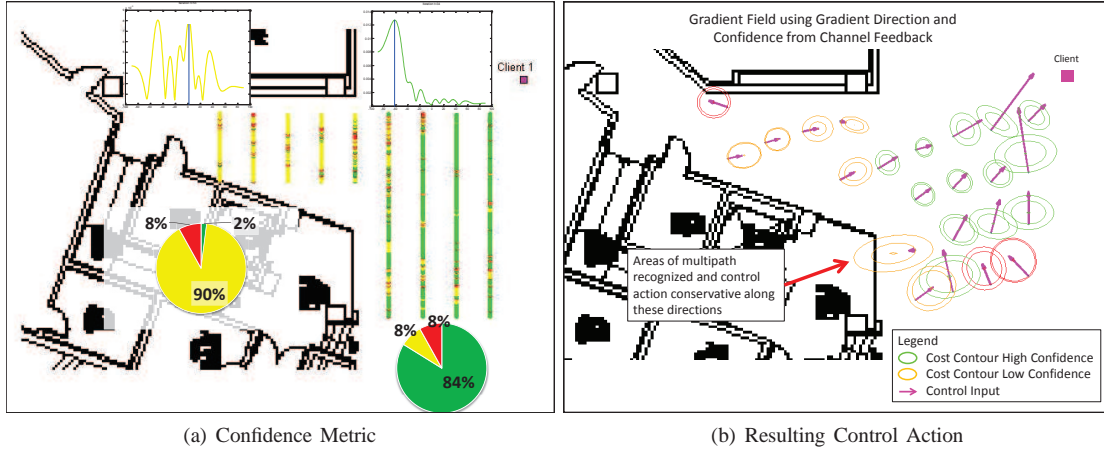


Fig. 8. Figure (a) shows data collected for a one-link system of one router and one client where the client is stationary at the top right corner of a basement environment and a mobile router is driven in a lawn mower pattern throughout the environment through line-of-sight and non-line-of-sight regions. Each colored data point represents an acquired directional signal profile (two example profiles are shown) and the color of the data point is the result of automatic mode detection from the data using the confidence metric from Eq. (13) where red=noise, yellow=multipath, and green=high confidence peak. In (b) the resulting edge cost contours (Equation (14)) and actual control command at each point in the environment is shown. Confidence values have a direct effect on velocity (as indicated by arrow length) where confident directions are pursued more aggressively.

An important observation from the data in Figure 8a is that even in line-of-sight regions of the environment (relative to the position of the client) there may be significant multipath present due to reflections from nearby concrete walls and this may cause the direction profile to have peaks in heading directions that are non-intuitive. Therefore this data suggests that metrics relying solely on the geometry of the environment, including visibility graphs, do not adequately capture the complexities of wireless signal quality in general environments.

C. Construction of Communication Link Costs

Our objective here is to construct the Mahalanobis distance matrix M_{ij} for each communication link (or edge) in the network using $f_{ij}(\theta)$. Specifically,

Problem 2 (Computation of M_{ij}): For each communication link (i, j) in the network where $i \in [k]$ and $j \in [n]$, find a heading direction $\vec{v}_{\theta_{\max}}$, a confidence metric σ , and a construction

of M_{ij} such that setting the edge costs \tilde{g} from Equation (8) to $\tilde{g} := \text{dist}_{M_{ij}}^2(p_i, c_j)$ satisfies Properties 1-3.

The direction along which the signal strength is maximum, θ_{\max} , is characterized by a peak in the $f_{ij}(\theta)$ plot and we define $\vec{v}_{\theta_{\max}}$ to be the unit vector along this recommended heading direction θ_{\max} . Using this direction alone does not provide enough information for effective position control of the routers however, due to the fact that this direction may experience corruption due to noise or multipath. In the previous section we showed that the presence of multipath or noise in the $f_{ij}(\theta)$ plot can be identified via the computation of a confidence metric σ_{ij} . Now, we encode the quantity σ_{ij} into our controller such that $\vec{v}_{\theta_{\max}}$ directions of high confidence are followed more aggressively (larger displacements along these directions), and the opposite is true of $\vec{v}_{\theta_{\max}}$ directions with low confidence. Figure 8b shows the effect of the confidence value on the commanded displacements made by the controller

in an actual implementation.

Specifically, for the three categories of σ_{ij} we desire the following behaviors for the routers: 1) $\sigma_{ij} < 1$: Indicates a high confidence in $\vec{v}_{\theta_{\max}}$ due to a sharp peak in f_{ij} . The robot is moved at higher speeds; 2) $\sigma_{ij} \approx 1$: Indicates that f_{ij} is noisy, so the robot moves slowly; 3) $\sigma_{ij} > 1$: Indicates that f_{ij} has multiple significant peaks owing to multi-path. We study this last case in Sec. V-D., and particularly the opportunity it presents for making trade-offs between clients.

We use the heading direction and confidence to design a cost function \tilde{g} that locally captures the cost of communication in the spatial domain. We express this cost as a Mahalanobis distance. The square of the Mahalanobis distance is a cost function (paraboloid) with ellipsoidal level sets (Fig. 9). We design our cost by orienting these level sets so that the direction of steepest descent is along $\vec{v}_{\theta_{\max}}$. We then skew the ellipsoidal level sets using the confidence σ_{ij} , so that a higher confidence translates to a steeper descent which leads to larger router displacements (speed) in the descent directions with high confidence.

Algorithm 2 provides a calculation of the matrix M_{ij} from Problem (2). We simply set one of the eigen-vectors of Q to the heading direction $\vec{v}_{\theta_{\max}}$. To skew the ellipsoid, we set the ratio of the eigen values $\{\lambda_1, \lambda_2\}$ in Λ to the confidence σ_{ij}^2 , i.e. $\lambda_2/\lambda_1 = \sigma_{ij}^2$, where λ_1 is the eigen-value corresponding to $\vec{v}_{\theta_{\max}}$. For example, in Fig. 9(b), where $\sigma_{ij} \approx 1$ (i.e. poor confidence), the level sets are nearly circular, leading to a shallow descent in cost; while Fig. 9(c), where $\sigma_{ij} < 1$ (i.e. high confidence), the level sets are skewed, leading to a steep descent in cost along $\vec{v}_{\theta_{\max}}$. In other words, the cost function has an elegant geometric interpretation, akin to Euclidean distance, but is derived directly from channel measurements. Further, the cost function $\tilde{g} := \text{dist}_{M_{ij}}^2(p_i, c_j)$ from Eqn. 9 is quadratic, a desirable property for optimizations.

Algorithm 2: Algorithm for constructing M_{ij} from channel feedback.

input : Directional signal strength map f_{ij} for every link (i, j) from Algorithm 1

output: A matrix M_{ij} for defining communication edge costs in Equation (8) using Mahalanobis distance from Problem 2.

- 1 $Q_{ij} \leftarrow [\vec{v}_{\theta_{\max}}, \vec{v}_{\theta_{\max}^\perp}]$ // a set of orthonormal basis vectors defined using $\vec{v}_{\theta_{\max}}$
 - 2 $\sigma_{ij} \leftarrow \frac{\sigma_{f_{ij}}}{\sigma_{N_{ij}}}$ // confidence in the $\vec{v}_{\theta_{\max}}$ direction
 - 3 $\Lambda = \text{diag}([\frac{1}{\sigma_{ij}^2}, 1])$ // Construct a diagonal matrix using confidence
 - 4 $M_{ij} = Q_{ij}\Lambda Q_{ij}^T$
-

D. Network Trade-offs

In this section, we show how our optimization framework readily extends to a multi-agent scenario and study the different trade-offs. We show that via the setting of two parameters, both set automatically from wireless channel data, the resulting positional controller can be made to greedily optimize one client's needs or alternatively, strike trade-offs between multiple clients. First, we focus on managing service discrepancies specified by w_{ij} . The quantity w_{ij} aims to bias

the controller by assigning higher weight to users with larger service discrepancies. To do this, we scale the cost function $\tilde{g} = \text{dist}_{M_{ij}}^2(p_i, c_j)$ by the square of the discrepancy w_{ij}^2 to optimize yield the network cost:

$$r_M(P, C) = \max_{p_j \in P} \min_{c_i \in C} \{w_{ij}^2 \text{dist}_{M_{ij}}^2(p_i, c_j)\} \quad (14)$$

Second, we highlight the subtle role played by the confidence σ_{ij} in managing network trade-offs. For instance, consider a scenario with two clients: 1 and 2, where client-1 demands greater communication quality (as specified by w_{ij} 's). Suppose client-1 has a highly confident $\vec{v}_{\theta_{\max}}$ as shown in Fig. 10(a) (i.e. $\sigma_{ij} < 1$). As expected, the robotic router is directed towards client-1 as shown in Fig. 10(c). In the more interesting scenario in Fig. 10(b), client-1's confidence is poor due to multiple peaks in the signal profile f_{ij} (i.e. $\sigma_{ij} > 1$). Here, the router strikes a trade-off and services client-2 instead, as this may potentially benefit client-1 as well due to the multipath recognized in client-1's $f_{ij}(\theta)$ map. The intuition behind this is simple. Equation 14 above, scales the ellipsoidal cost function based on the discrepancies w_{ij} 's. However, recall that the ellipsoidal cost function is steep (or shallow) depending on whether the confidence is high (or low) and this is attained by setting the ratio of eigenvalues λ_2/λ_1 of M_{ij} (See Line 3 in Algorithm 2). In extremely low confidence scenarios such as Figure 10(b), the higher value of discrepancy of client-1 is masked by its low value of confidence. Hence, this balances the trade-off in favor of client-2, despite having a lower discrepancy.

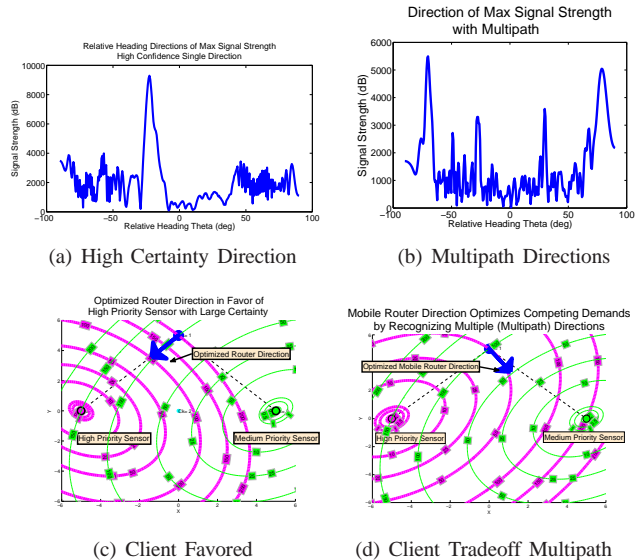


Fig. 10. Trade-offs between Clients: (a) – (b) show the $f_{ij}(\theta)$ map for the high demand client; (c) – (d) show the optimized router direction

Algorithm 3 demonstrates how the cost in Equation (14) can be used to find an updated set of router positions when both client and router positions are known at the current iteration.

The optimization in Equation (15) in Algorithm 3 is equivalent to a k -center optimization problem where the distance metric is a Mahalanobis distance. This is a generalized router placement problem similar to that studied in [Gil et al., 2012] for Euclidean distances. Thus the returned solution from this algorithm is the optimal placement of routers corresponding to

the optimal assignment of routers to clients, given the channel feedback at the current iteration t .

Algorithm 3: Algorithm for router placement with known client positions.

input : Directional signal strength map $f_{ij}(\theta)$ for every link (i, j) , demand q_j , relative importance $\alpha_j > 0$ for client j , current quality of each link ρ_{ij} , and current router and client positions

$$P = \{p_1, \dots, p_n\}, C = \{c_1, \dots, c_n\}$$

output: A configuration of optimal router positions C^* , $|C^*| = k$, given the current channel feedback for all links in the network.

1 **for** all links (i, j) in the network with $\rho_{ij} > 0$,

$i \in [k], j \in [n]$ **do**

2 $w_{ij} \leftarrow \max(\alpha_j \frac{q_j - \rho_{ij}}{q_j}, 0)$;

3 $M_{ij} \leftarrow$ result of Algorithm 2;

4 **end**

5 Compute:

$$C^* = \arg \min_C \{ \max_{p_j \in P} \min_{c_i \in C} w_{ij}^2 (p_i - c_j)^T M_{ij} (p_i - c_j) \} \quad (15)$$

return C^*

E. A Position-Independent Solution

A simple relaxation to the cost from the previous section frees the optimization of using client positions, while maintaining its simple structure and desirable properties developed above. Consider a user specified step-size $\gamma > 0$, that encodes the maximum permissible displacement for each router and denote $c_{i,t}$ to be the current router position. We replace client positions p_j in Equation (14) with “virtual” positions p'_{ij} :

$$p'_{ij} = c_{i,t} + \gamma w_{ij} \vec{v}_{\theta_{\max}}. \quad (16)$$

Intuitively, a client is no longer directly observed but rather estimated to be along the relative direction $\vec{v}_{\theta_{\max}}$ and at a distance of γw_{ij} with respect to the i th router. As before, $\vec{v}_{\theta_{\max}}$ is the heading direction associated with the maximum strength signal direction θ_{\max} . As a client’s demand is better satisfied by router i , the service discrepancy w_{ij} tends to 0 and the client is perceived as being closer to router i . The observation here is that routers better equipped to service a particular client as reflected by the w_{ij} term, will view the client as “closer” and those routers with a weaker signal to the same client will view this client as farther away. This results in a natural method of assigning client nodes to routers by effectively sensing over the wireless channels.

F. Controller for Router Positioning

We now present an algorithm for achieving router positions that minimize the edge costs \tilde{g} derived in the previous sections. Particularly we formulate \tilde{g} from Equation (3) to be

$$\tilde{g}(c_i, C_t, w_{ij}, f_{ij}) = (p'_{ij} - c_i)^T M_{ij} (p'_{ij} - c_i) \quad (17)$$

Where the dependence of \tilde{g} on C_t , w_{ij} and f_{ij} are captured indirectly by p'_{ij} and M_{ij} via Equation (16) and Algorithm 2

respectively. This choice of edge costs satisfy Properties 1-3. Namely, having a quadratic form, allowing optimization over the entire network with competing demands, and being independent of client positions. As described in Section III, minimization of these edge costs by Equation (3) results in the optimization of a network-wide metric, ie. minimizing the worst-case client service discrepancy.

The resulting optimization framework can be shown to exhibit other desirable properties relative to the instantaneous wireless channels over the network. An important remark is that we do not make assumptions on how the wireless channels may change over time, nor do we make assumptions on the underlying signal quality function in areas of the environment that are not currently being sensed by the routers. Unfortunately, this impairs our ability to prove certain desirable controller attributes such as convergence, that would require some additional assumptions on the signal quality such as a guarantee that this function is smooth, and can be strictly improved at every iteration. Such assumptions would be invalidated by small-scale fading alone [Goldsmith, 2005; Lindhe et al., 2007], in real wireless systems. However, by relying solely on instantaneous channel feedback, we retain the important ability to adapt quickly to changes in the wireless environment due to dynamic obstacles, for example. Based on current channel feedback, we highlight our controller’s network-wide properties. The following properties, and convergence to client demanded rates, are demonstrated extensively in actual implementations in the next section of the paper:

Property 4: The assignment of routers to clients is optimal based on the current feedback over wireless links in the network.

This can be seen from the observation that Line 8 from Algorithm 4 is the classic k -center solution [Feldman et al., 2013; Gil et al., 2012] under the Mahalanobis distance metric. A k -center solution will assign clients to their closest routers. In this case “closest” is defined in the signal quality sense where routers serve the clients to whom their signal strength is greater than the signal strength between any other router in the network to the same client. An example of this property in an actual hardware implementation can be seen in Section VI-C-VI-D where routers choose clients based on the strengths of their relative wireless links.

Property 5: Stability of router positions to solutions that satisfy client demands over the network.

Our final cost takes the form:

$$r_M(C) = \max_{j \in \{1, \dots, n\}} \min_{c_i \in C} \{ \text{dist}_{M_{ij}}^2(c_{i,t} + \gamma w_{ij} \vec{v}_{\theta_{\max}}, c_i) \} \quad (18)$$

By expanding the squared per-link cost $\text{dist}_{M_{ij}}^2(c_i + \gamma w_{ij} \vec{v}_{\theta_{\max}}, c_i)$ from Eqn. 14:

$$(c_i - c_{i,t})^T M_{ij} (c_i - c_{i,t}) - 2\gamma w_{ij} \lambda_{\theta_{ij}} \vec{v}_{\theta_{\max}}^T (c_i - c_{i,t}) + \gamma^2 w_{ij}^2 \lambda_{\theta_{ij}} \quad (19)$$

we note that as $w_{ij} \rightarrow 0$ the first term in Eqn. (19) favors stable solutions where $c_i = c_{i,t}$, ie. the router reaches a static solution when all of its assigned clients have zero service discrepancy. In the case where it is not possible to satisfy all client demands, for example if there are not enough routers k to provide communication coverage to the clients, Algorithm 4 returns the solution with the lowest service discrepancy that is within a user specified tolerance of optimal. Extensive empirical

validation of this property in actual hardware implementations is shown in Sections VI-C-VI-E.

Property 6: Adaptation of router positions to changes in the wireless channels and/or client demand.

The Equation (19) shows that for nonzero discrepancy, ie. $w_{ij} \neq 0$, the cost for edge (i, j) is also nonzero. Thus a change in the client demands q_j , or in the quality of the wireless link ρ_{ij} due to moving occlusions or changes in the environment, will equally change the weighting w_{ij} on the link (see Equation (1)). If the change in link quality ρ_{ij} is sufficient, ie. if $w_{ij} > 0$, the routers following Algorithm 4 will update their positions until a new solution is found (see Property 5). Empirical validation of this property in an actual implementation is demonstrated in Section VI-F.

An algorithm for finding router placements in the most general case of unknown client positions is presented as Algorithm 4 below.

Algorithm 4: Algorithm for router placement with unknown client positions.

input : Directional signal strength map $f_{ij}(\theta)$ for every link (i, j) , demand q_j , relative importance $\alpha_j > 0$ for client j , current quality of each link ρ_{ij} , step-size $\gamma > 0$, tolerance $tol > 0$ and current router positions $C = \{c_1, \dots, c_n\}$

output: A configuration of router positions C^* , $|C^*| = k$, and the achieved service discrepancy w^* .

```

1  $\delta \leftarrow \text{inf}$ ;
2 while  $\delta > tol$  do
3   for all links  $(i, j)$  in the network with  $\rho_{ij} > 0$ ,
    $i \in [k], j \in [n]$  do
4      $p'_{ij} \leftarrow c_{i,t} + \gamma w_{ij} \vec{\theta}_{\max}$ ; // compute virtual
     client  $j$  position as perceived by
     router  $i$ 
5      $w_{ij} \leftarrow \max(\alpha_j \frac{q_j - \rho_{ij}}{q_j}, 0)$ ;
6      $M_{ij} \leftarrow \text{result of Algorithm 2}$ ;
7   end
8   Compute:
   
$$C^* = \arg \min_C \{ \max_{p'_{ij} \in P} \min_{c_i \in C} (p'_{ij} - c_i)^T M_{ij} (p'_{ij} - c_i) \}$$
 (20)
   
$$(21)$$

    $C \leftarrow C^*$ ; // Update router positions
9   for all links  $(i, j)$ ,  $i \in [k], j \in [n]$  do
10     $w_{ij} \leftarrow \max(\alpha_j \frac{q_j - \rho_{ij}}{q_j}, 0)$ ; // Compute
    updated  $w_{ij}$ 
11  end
12  ;
13   $\delta \leftarrow \max_{(i,j)} (w_{ij})$ ; // Store max
14 end
15  $w^* = \min_{j \in [n]} \max_{i \in [k]} w_{ij}$ ;
16 return  $C^*, w^*$ 

```

VI. EXPERIMENTAL RESULTS WITH MOTION CAPTURE SUPPORT

We evaluated our system on a five-node testbed with two routers and three clients. Each node was an ASUS 1015PX netbook equipped with an Intel 5300 Wi-Fi card mounted on

an iRobot Create robot. We implemented SAR by modifying the iwlan driver on Ubuntu 10.04. We used the 802.11 CSI tool [Halperin et al., 2011] to obtain channel information ($\hat{h}(t)$ in Eqn. 7). The routers communicated with a central laptop emulating the base for control information and human input. Our first set of experiments were performed in a room with a Vicon motion capture system to measure the relative displacement of the robotic routers. Sec. VII describes results in complex indoor environments without motion capture support. Our testbed contains obstacles to simulate both line-of-sight and non-line-of-sight scenarios.

A. Measuring the Direction of Maximum Signal Strength

We first provide a microbenchmark to demonstrate that our system indeed provides the direction θ_{\max} that results in maximum improvement in client service quality. We consider two representative examples of a single robot router-client pair placed in i) line-of-sight configuration where the strongest signal path is also the shortest Euclidean distance path, and ii) non-line-of-sight configuration where the shortest Euclidean path between the router and client is obstructed by a cement column. These configurations are depicted in Figure 11(a). We measure the power profiles of signals from the robot router from different directions using the solution described in Sec. IV. We also compute the average service quality of the client (measured in terms of Effective Signal-to-Noise Ratio or ESNR) along various spatial directions by iteratively moving the robot router and exhaustively sampling the signal quality along each physical direction, at a total of 1800 samples (100 samples, about 1m, along each ten degree arc).

Results: Fig. 11(b) and (c) plots the power profile obtained by our system, as well as the service quality observed when moving along the different spatial directions. We note that the direction of maximum signal power measured by our system actually leads to maximum increase in service quality in both line-of-sight and non-line-of-sight settings. Notice that while the plots in both Fig. 11(b) and (c) capture similar trends, they are not identical. Specifically, the profiles output by our system isolate signal power arriving from individual spatial directions, and therefore have sharp peaks that are easy to discern and less prone to error. In contrast, the average service quality of the client varies much more gradually along different directions, and therefore needs to be sampled much more extensively to obtain accurate trends (for more details, see Sec. IV). This demonstrates that our system captures the direction of maximum signal strength with a higher accuracy, and without the need for exhaustive exploration, when compared to pure sampling-based approach.

B. Visualizing the Gradient Field of Signal Strength

In this experiment, we visualize the gradient field of the directions of maximum signal strength θ_{\max} , on a wireless link. We consider a single client, serviced by a robot router that is: 1) In direct line-of-sight (LOS) as shown in Fig. 12(a). 2) In possible non-line-of sight (NLOS) scenarios due to obstacles as shown in Fig. 12(b). We drive the robot router in a lawn-mover pattern and get θ_{\max} at regular intervals.

Results: Fig. 12(a) and 12(b) depict the gradient field with the arrows indicating θ_{\max} in LOS and NLOS, respectively. The gradient field in LOS accurately directs the robot router

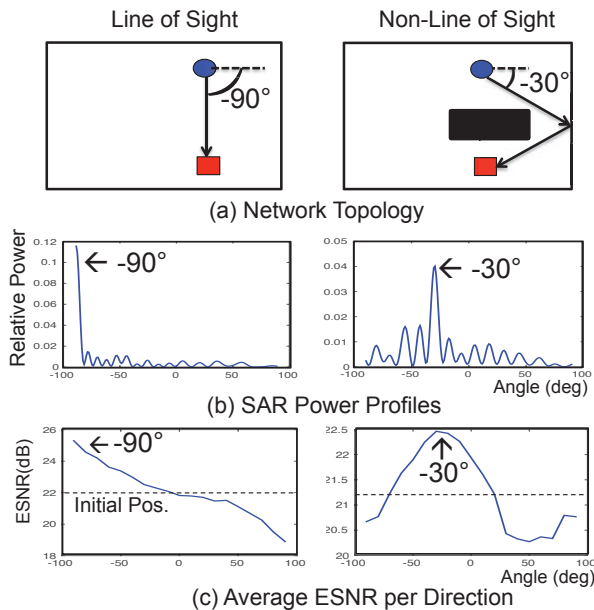


Fig. 11. We validate our computed direction of signal strength for two representative configurations of (a) line-of-sight and non-line-of-sight settings. (b) Power profiles indicating direction of maximum signal strength. (c) Service quality (average ESNR) measured along each spatial direction. The horizontal dotted line indicates the ESNR at the initial position.

towards the client regardless of its initial position. In NLOS, the robot is directed away from obstacles so that controller can route around obstacles to improve signal strength. We stress that θ_{max} is found locally at the router purely via wireless channels and its own position, *without* prior knowledge of the environment. Further, the plots are not static and naturally *change over time*, especially in dynamic settings. Thus our system obtains instantaneous θ_{max} values locally in real-time.

Fig. 12(c) and 12(d) plot $f_{ij}(\theta)$, the power profile of the signal along different directions, for a candidate location in line-of-sight and non-line-of-sight scenarios, respectively. Clearly, the power profile in line-of-sight is dominated by a single peak at θ_{max} , directed along the line-of-sight path to the client. In contrast, the power profile in non-line-of-sight close to an obstacle has two significant peaks, each corresponding to reflected paths along walls or other objects in the environment.

C. Controlling Router Trajectory to satisfy Client Demands

We evaluate how a single robotic router finds a trajectory to satisfy the demands of three clients (specified in terms of effective signal-to-noise ratio or ESNR) using θ_{max} on each link. We consider the candidate non-line-of-sight setting in Fig. 13(a). The router is unaware of exact client positions or the layout of the environment.

Results: Fig. 13(a) depicts the trajectory of the robotic router in blue. The colored arrows denote the recommended $\vec{v}_{\theta_{max}}$ directions for each client at every control point. The figure shows how the robot performs non-zero control actions until it eventually satisfies network demands. Fig. 13(b) tracks the ESNR of the clients across time (dotted lines). The plot shows that the ESNR demands of each client (solid lines) are satisfied upon convergence. Note that the whenever the robot decides to follow the $\vec{v}_{\theta_{max}}$ of a client at a control point (vertical line), the client's ESNR increases. This validates our claim that

following a heading direction based on $\vec{v}_{\theta_{max}}$ indeed improves the ESNR of the corresponding client.

D. Aggregate System Results

We evaluate our full system with two robot routers serving three clients with different ESNR demands. We perform the experiment in line-of-sight (LOS) and non-line-of-sight (NLOS) settings as shown in the inset maps of Fig. 14(b) and 14(d) respectively. We repeat the experiment five times in each setting and plot the results.

Results: Fig. 14(a) and 14(b) plot the mean and standard deviation of ESNR over time across experiments for each client (dotted colored lines) in LOS and NLOS. Clearly, each client's ESNR demand (solid lines) is satisfied at the converged position across experiments. Fig. 14(c) and 14(d) plot the corresponding aggregate link rate across time, which follows the same trend as the ESNR [Halperin et al., 2010].³ The inset plots in Fig. 14(c) and 14(d) depict the final converged position of the routers (blue dots) in LOS and NLOS. The results show that our system consistently satisfies client demands while adapting to real-time changes in wireless channels, even in the presence of obstacles.

E. Comparison with Existing Schemes

We test our method against two other popular approaches to the communication problem in robotics: 1) Euclidean Disk Model as used in [Cortes et al., 2004; Jadbabaie et al., 2003], where communication constraints are in terms of Euclidean distance; 2) Stochastic Gradient Approach, where we implement the Simultaneous Perturbation method (SPSA) [Spall, 2000] for estimating the gradient of signal power by sampling the ESNR (which provides greater granularity than RSSI), along randomized directions, similar to the approach utilized by [Le Ny et al., 2012]. For the generation of each direction in the SPSA method we use a Bernoulli random variable (as in [Spall, 2000]) and diminishing step sizes satisfying the conditions stated in [Spall, 2000] for convergence. Our largest step size was allowed to be the same maximum vehicle velocity of v_c for all experiments. We consider a robotic router and three clients, each with an ESNR demand of 20 dB. We repeat the experiment five times in the non-line-of-sight environment in Fig. 15(b)-(d). In each instance, we measure r_{max} , the maximum ratio of ESNR demand versus the ESNR achieved among all three clients. In particular, r_{max} is below one at the converged position (i.e. all client demands are satisfied), and above one otherwise.

Results: Fig. 15(a) plots the aggregate mean and standard deviation of r_{max} across time, for all the three approaches. Fig. 15(b)-(d) show a candidate trajectory adopted by the robotic router for the three schemes. The plots demonstrate while the disk model converges quickly to a solution, ignorance of the wireless channels leads to solutions not meeting client demands; especially in non-line-of-sight settings. In contrast, the stochastic gradient approach (in blue), which sample the instantaneous ESNR, eventually satisfies network demands. However, the convergence is often laborious as the router often traverses counter-productive directions (see Fig. 15 (c)). Indeed such techniques are noisy at low signal power, as even

³Note that the data-rate is capped by 60 Mb/s causing the plot to appear flat at times unlike ESNR.

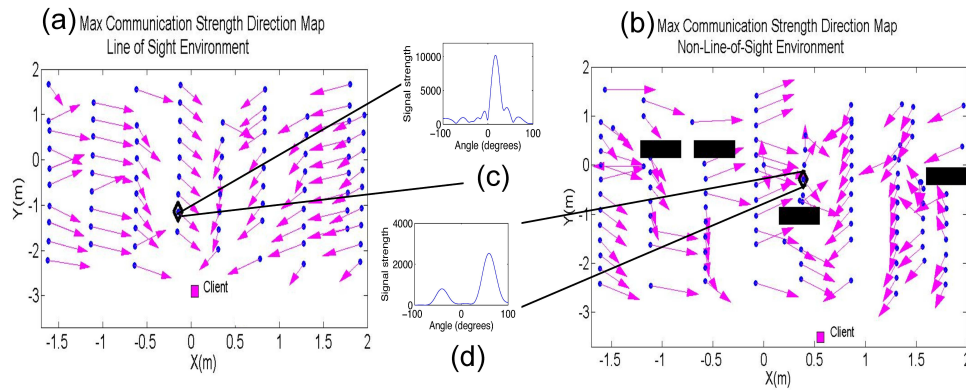


Fig. 12. Gradient field of θ_{max} and power profile for (a) Line-of-sight and (b) Non-Line-of-Sight.

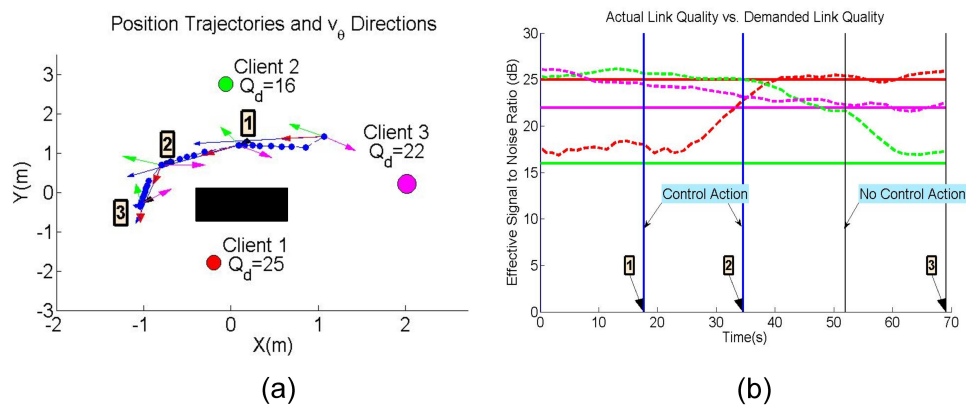


Fig. 13. (a) Depicts testbed with robot router servicing three clients in a candidate non-line-of-sight setting. The blue line depicts the trajectory, and colored arrows indicate instantaneous θ_{max} for the corresponding clients. (b) Plots the ESNR across time (as dotted lines) for each client through the experiment. Solid lines denote client demands.

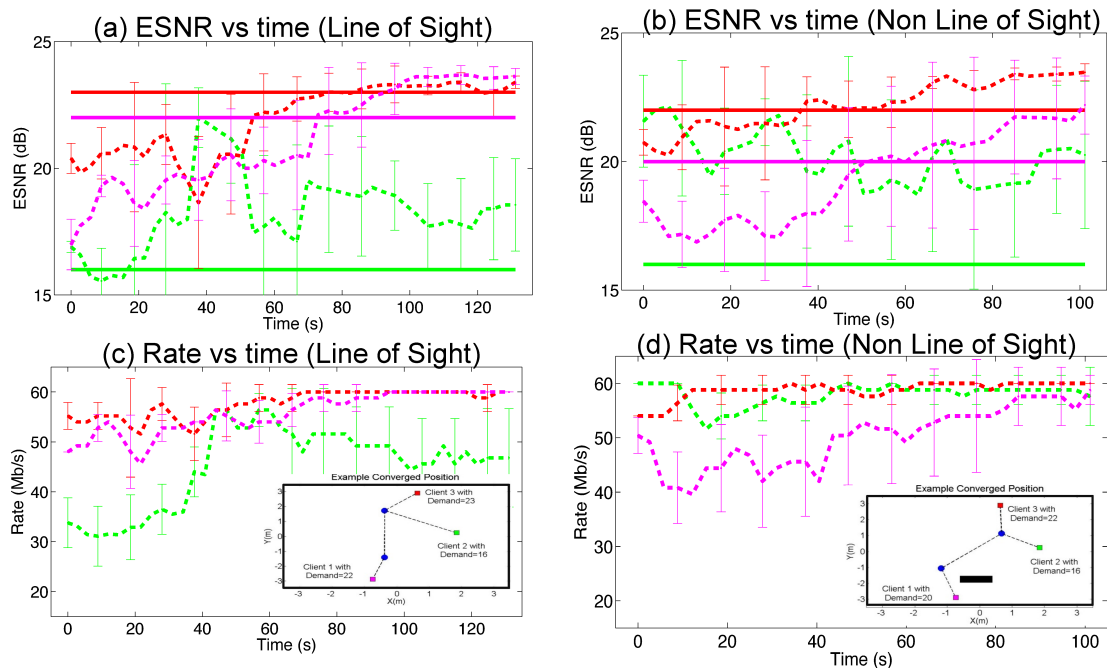


Fig. 14. Aggregate results obtained over 5 runs show demands are consistently met even in the presence of obstacles as demonstrated by the candidate converged solutions.

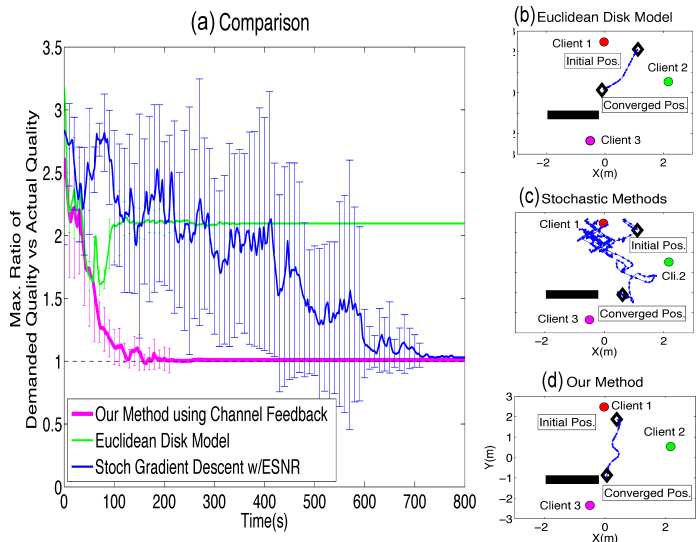


Fig. 15. Plots comparing our method against the Euclidean disk model and a stochastic gradient descent method based on ESNR. Our method both converges to a position that meets communication demands, and converges quickly along an efficient path.

a large change in distance translates to a small change in signal power (a well-studied problem, e.g. in [Chen et al., 2012; Joshi et al., 2013; Xiong and Jamieson, 2013]). Fig. 15(c) shows that this leads to areas at non-line-of-sight or far distances from the client, where the robot easily gets lost.

Our method leverages full channel information, including signal power and phase, to find the signal direction as opposed to just its magnitude. The result is an algorithm that converges to positions that satisfy network demands without the counter-productive exploration of a pure sampling approach.

F. Robustness to Dynamic Obstacle Positions

We evaluate how our system adapts to changes in the environment without an *a priori* known map. Consider two robotic routers and three clients in an environment with an obstacle located initially as shown in Fig. 16(a). We allow the robot routers to navigate to their converged positions. At $t = 120$ sec, we move the obstacle to a different location as in Fig. 16(c), and let the routers re-converge.

Results: Fig. 16(b) and Fig. 16(c) depict the converged position of the routers before and after the obstacle was moved. Fig. 16(d) plots the data-rate across time for each client. The plot shows that our system satisfies client demands at the initial position. It also recovers from the sharp fall in rate at one client and successfully re-converges after the obstacle is moved.

VII. EXPERIMENTAL RESULTS WITHOUT MOTION CAPTURE SUPPORT

In this section, we evaluate our system in a large complex indoor environment with concrete walls and columns without any motion capture support (see Fig. 17). Instead, we use a constant velocity assumption to infer the relative displacements, $d(t)$ (see Algorithm 1), of the Wi-Fi antenna on the router. The requirements for obtaining $d(t)$ as described in Section IV, are that the robot router moves at a constant known velocity over the time window required for computing SAR. Thus in our experiment we command the iRobot Create

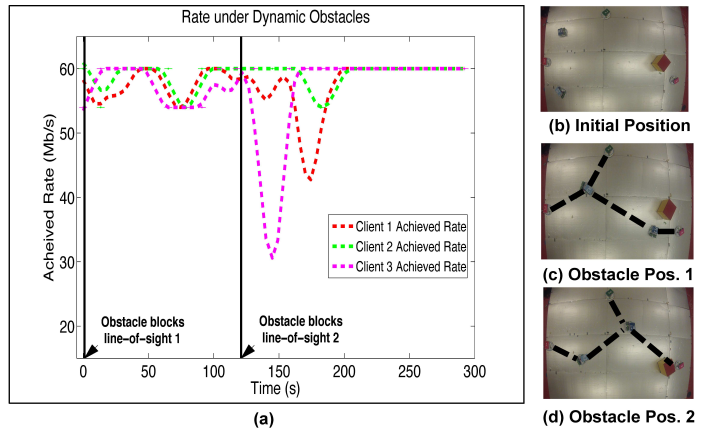


Fig. 16. These plots show the result of disturbing the wireless channels via movement of a line-of-sight obstructing obstacle. Actual testbed snapshots are shown on the right.

platform to move with a known constant velocity between control actions.

A. Gradient Field in Complex Environments

In this experiment, we measure the gradient field capturing the direction of maximum signal strength across spatial locations in the above testbed without Vicon support. We place a robotic router and client and line-of-sight (LOS) and non-line-of-sight (NLOS) as in Fig. 17. We trace the router's gradient field towards the client starting from multiple initial positions.

Results: Fig. 17 (a) and (b) plot of candidate trajectories (from gradient field) in LOS and NLOS across initial locations. The plots show that our system successfully navigates towards the client to satisfy its demands, without knowledge of the environment or client location.

B. Full-Scale Experiment in Complex Environments

We implement a full-scale experiment of two routers and five clients in the complex indoor environment described in Sec. VII-A above with no motion capture support. Clients in this case are static Asus 1015PX series netbooks and routers are AscTec Atom boards mounted on mobile iRobot Create platforms.

Clients are positioned in two clusters along orthogonal hallways, i.e. a non-convex environment. Routers are placed in the initial positions as shown in the floorplan in Figure 18(a). For these initial router positions, Client 1 and Client 2 are both out of direct line-of-sight as they are obstructed by a concrete wall.

The relative displacement of the Wi-Fi antenna, required by Algorithm 1 to obtain a directional signal strength profile, is measured by assuming the router moves at a constant velocity (see Sec. IV-B).

Before calculating the next waypoint, each router is commanded to move at constant velocity for a period of 24 seconds which is equivalent to two wavelengths in displacement. The commanded waypoint from the control Algorithm 4 is then provided as a heading/distance pair which is actuated by the router using dead reckoning.

Results: Figure 18(a) depicts the initial configuration of the network of routers and clients. The dotted lines indicate which




Fig. 17. Trajectories using measured $\vec{v}_{\theta_{\max}}$ directions satisfy a client's demand in line-of-sight and non-line-of-sight settings in complex indoor environments.

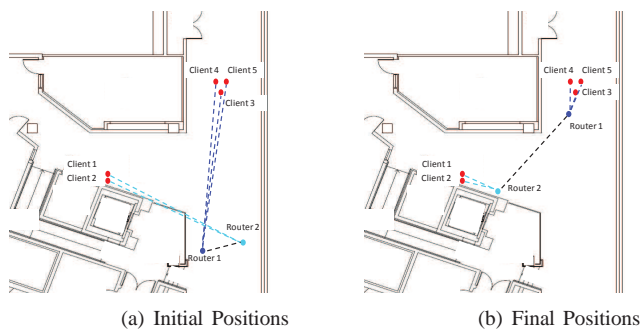


Fig. 18. Full-scale experimental setup with initial (a) and final (b) configurations for a two-router five-client network configuration.

cluster of clients each router is assigned to by the controller. These assignments are optimized by the controller based on the observed ESNR values as described in Algorithm 4. Figure 18(b) show the converged positions of the routers indicating that all client demands are satisfied at these positions. In particular, Figure 19 demonstrates the trajectories traversed by Router 1 and Router 2 (left top and left bottom respectively) and the corresponding ESNR curves for each router's assigned clients on the right column. The ESNR curves are averaged over a window of 24 seconds as the router moves along its trajectory, and the solid blue squares indicate the times where a control action was given. As shown by ESNR curves in Figure 19, all client demands are satisfied at the final router configuration.

VIII. DISCUSSION

Our primary focus in the body of this paper has been on developing a closed-loop controller that uses instantaneous feedback on wireless channels to position routers. This real-time feedback allows for routers to repair communication links on the fly as needed, for example in the case of dynamic obstacles that may occlude a link. However, here we point out that it is also possible to use the methods presented in this paper to obtain a static directional map of signal strength, or gradient field, throughout the environment. In fact, the richness of directional profiles derived here would allow mapping to a level of accuracy that was previously unattainable for small mobile platforms. Such a gradient field (as in Figure 12) can be used to plan router placements that are globally optimal, in contrast to the local solutions provided here. However, it is important to point out that in this case the ability to adapt to changes in the environment, for example if obstacles or clients

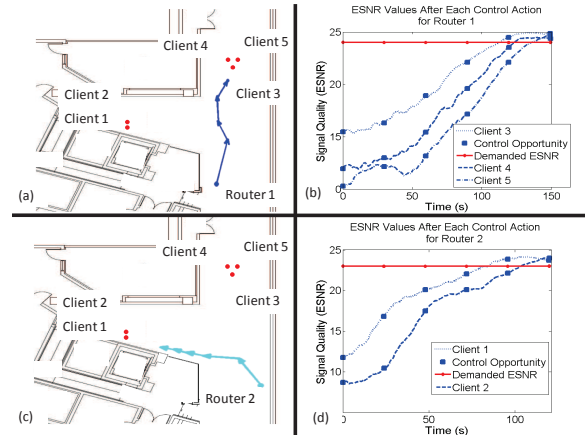


Fig. 19. Router trajectories resulting from execution of commanded waypoints from Algorithm 4. These paths were executed via dead reckoning for Router 1 (a) and Router 2 (c). Corresponding measured ESNR curves for Router 1's clients (b) and Router 2's clients (d) respectively.

move in the environment, is lost since this would invalidate a static map. Therefore this is a trade-off that would have to be evaluated carefully for each situation.

IX. CONCLUSION

In this paper, we present a framework to satisfy real-time variable communication demands for a changing network. We develop a solution enabling a robotic receiver to find the profile of signal strength across spatial directions for each sender of interest. While our technique retrieves these spatial signal profiles in real time, we note that it faces an important limitation: it assumes access to wireless channels from both the transmitter and the receiver. Developing a system that can work with unmodified transmitters remains an open challenge. Our system integrates the signal profiles with a controller that optimizes communication quality while maintaining quadratic edge costs, and thus has natural extensions to many communication-aware coordination problems such as coverage[Cortes et al., 2004], consensus[Olfati-Saber et al., 2007], formation control[Jadbabaie et al., 2003], etc. We believe our system provides the necessary robustness to bring the benefits of these important contributions to practical robotic systems.

Acknowledgments: The authors thank Dan Feldman for his comments on problem formulation and previous collaboration on k -center approaches that inspired much of this work, and

Brian Julian for his invaluable help in setting up the iRobot Create testbed. The authors acknowledge MIT Lincoln Laboratory and MAST project under ARL Grant W911NF-08-2-0004 for their support. We thank members of the MIT Center for Wireless Networks and Mobile Computing: Amazon, Cisco, Google, Intel, Mediatek, Microsoft, ST Microelectronics, and Telefonica for their interest and general support.

REFERENCES

- Adib, F. and Katabi, D. (2013). See through walls with wi-fi. In *SIGCOMM*.
- Chen, H.-C., Lin, T.-H., Kung, H., Lin, C.-K., and Gwon, Y. (2012). Determining rf angle of arrival using cots antenna arrays: A field evaluation. In *MILCOM*.
- Cortes, J., Martinez, S., Karatas, T., and Bullo, F. (2004). Coverage control for mobile sensing networks. In *IEEE Transactions of Robotics and Automation*.
- Cortes et al., J. (2004). Spatially-distributed coverage optimization and control with limited-range interactions. In *ESAIM*.
- Craparo, E., How, J., and Modiano, E. (2011). Throughput optimization in mobile backbone networks. *Mobile Computing, IEEE Transactions on*, 1.
- De Gennaro, M. and Jadbabaie, A. (2006). Decentralized control of connectivity for multi-agent systems. In *Decision and Control, 2006 45th IEEE Conference on*.
- Feldman, D., Gil, S., Knepper, R., Julian, B., and Rus, D. (2013). K-robots clustering of moving sensors using coresets. In *ICRA*.
- Fink, J., Ribeiro, A., and Kumar, V. (2012). Robust control for mobility and wireless communication in cyber-physical systems with application to robot teams. *Proceedings of the IEEE*, 1.
- Fink, J., Ribeiro, A., and Kumar, V. (2013). Robust control of mobility and communications in autonomous robot teams. *Access, IEEE*, 1.
- Fink, J., Ribeiro, A., Kumar, V., and Sadler, B. (2010). Optimal robust multihop routing for wireless networks of mobile micro autonomous systems. In *MILCOM*.
- Fitch, P. J. (1988). *Synthetic Aperture Radar*. Springer-Verlag.
- Gil, S., Feldman, D., and Rus, D. (2012). Communication coverage for independently moving robots. In *IROS*.
- Goldsmith, A. (2005). *Wireless Communications*. Cambridge University Press.
- Halperin, D., Hu, W., Sheth, A., and Wetherall, D. (2010). Predictable 802.11 packet delivery from wireless channel measurements. In *CCR*.
- Halperin, D., Hu, W., Sheth, A., and Wetherall, D. (2011). Tool release: Gathering 802.11 traces with channel state information. *ACM SIGCOMM CCR*, 1.
- Jadbabaie, A., Lin, J., and Morse, A. (2003). Coordination of groups of mobile autonomous agents using nearest neighbor rules. *Automatic Control, IEEE Transactions on*, 1.
- Joshi, K., Hong, S., and Katti, S. (2013). Pinpoint: localizing interfering radios. In *NSDI*.
- Le Ny, J., Ribeiro, A., and Pappas, G. (2012). Adaptive communication-constrained deployment of mobile robotic networks. In *ACC*.
- Lindhe, M., Johansson, K., and Bicchi, A. (2007). An experimental study of exploiting multipath fading for robot communications. In *RSS*.
- Lindhé, M. and Johansson, K. H. (2010). Adaptive exploitation of multipath fading for mobile sensors. In *ICRA*.
- Malmirchegini, M. and Mostofi, Y. (2012). On the spatial predictability of communication channels. *Wireless Communications, IEEE Transactions on*, 11(3):964–978.
- Martinez, S., Cortes, J., and Bullo, F. (2007). Motion coordination with distributed information. *Control Systems Magazine, IEEE*, 1.
- Michael, N., Zavlanos, M., Kumar, V., and Pappas, G. (2009). Maintaining connectivity in mobile robot networks. In *Experimental Robotics*, volume 54 of *Springer Tracts in Advanced Robotics*, pages 117–126. Springer Berlin/Heidelberg.
- Moreau, L. (2004). Stability of continuous-time distributed consensus algorithms. In *Decision and Control, International Conference on*.
- Networks, A. (2014). Outdoor antennas and rf coverage strategies.
- Olfati-Saber, R., Fax, J., and Murray, R. (2007). Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 1.
- Pham, T. and Sadler, B. (1997). Wideband array processing algorithms for acoustic tracking of ground vehicles. *US Army Research Laboratory Report*, 1.
- Rahul, H., Kumar, S. S., and Katabi, D. (2012). Megamimo: Scaling wireless capacity with user demand. In *SIGCOMM*.
- Schuresko, M. and Cortes, J. (2009). Distributed tree rearrangements for reachability and robust connectivity. *Intell. Robot. Syst.*, 1.
- Spall, J. (2000). Adaptive stochastic approximation by the simultaneous perturbation method. *Automatic Control, IEEE Transactions on*, 1.
- Spanos, D. and Murray, R. (2005). Motion planning with wireless network constraints. In *Proceedings of the ACC*.
- Stoica, P. and Moses, R. L. (2005). *Spectral Analysis of Signals*. Prentice Hall.
- Tahbaz-Salehi, A. and Jadbabaie, A. (2007). On recurrence of graph connectivity in vicsek’s model of motion coordination for mobile autonomous agents. In *ACC*, pages 699–704.
- Tse, D. and Viswanath, P. (2005). *Fundamentals of wireless communication*. Cambridge University Press.
- Twigg, J. N., Fink, J., Yu, P. L., and Sadler, B. M. (2013). Efficient base station connectivity area discovery. *Int. Journal of Robotics Research*, 32(12).
- Vieira, M. A. M., Taylor, M. E., Tandon, P., Jain, M., Govindan, R., Sukhatme, G. S., and Tambe, M. (2013). Mitigating multi-path fading in a mobile mesh network. *Ad Hoc Netw.*, 1.
- Wang, J., Adib, F., Knepper, R., Katabi, D., and Rus, D. (2013). RF-Compass: Robot object manipulation using rfids. In *MobiCom*.
- Wang, J. and Katabi, D. (2013). Dude, where’s my card? RFID positioning that works with multipath and non-line of sight. In *SIGCOMM*.
- Xiong, J. and Jamieson, K. (2013). Arraytrack: a fine-grained indoor location system. In *NSDI*.
- Yan, Y. and Mostofi, Y. (2013a). Co-optimization of communication and motion planning of a robotic operation under resource constraints and in fading environments. *Wireless Communications, IEEE Transactions on*, 1.
- Yan, Y. and Mostofi, Y. (2013b). Co-optimization of communication and motion planning of a robotic operation under

resource constraints and in fading environments. *IEEE Transactions on Wireless Communications*, 1.