2012 IEEE International Conference on Robotics and Automation
RiverCentre, Saint Paul, Minnesota, USA
May 14-18, 2012

Motion Planning for Robust Wireless Networking

Jonathan Fink, Alejandro Ribeiro and Vijay Kumar

Abstract— **We propose an architecture and algorithms for maintaining end-to-end network connectivity for autonomous teams of robots. By adopting stochastic models of point-to-point wireless communication and computing robust solutions to the network routing problem, we ensure reliable connectivity during robot movement in complex environments. We fully integrate the solution to network routing with the choice of node positions through the use of randomized motion planning techniques. Experiments demonstrate that our method succeeds in navigating a complex environment while ensuring that end-to-end communication rates meet or exceed prescribed values within a target failure tolerance.**

I. INTRODUCTION

The use of mobile robotic systems to achieve tasks that are otherwise dangerous for humans is becoming more and more ubiquitous. However, in a real-world deployed system, the issue of communication between a human operator and the robotic system severely limits the system's usefulness. While commercial technologies exist to provide general purpose wireless data connectivity, they rely on infrastructure that is not available in the hazardous environments for which robots are most useful. Instead, we propose that agents within a team of mobile robots must act as communication relays to reliably support communication requirements.

To restrict the scope of our work, we focus on a motivating telepresence application where a team of robots move and wirelessly relay messages in order to maintain a particular rate of communication between a human operator at a base station and the "lead" robot he or she is controlling. There are two key challenges to this work. First, point-to-point wireless channel capability is notoriously uncertain and difficult to predict. Second, the communication maintenance problem inherently involves a high-dimensional search that must jointly consider network routing and robot placement or mobility.

Much of the communication-aware mobility control literature relies on the assumption that point-to-point wireless links can be predicted primarily based on proximity or visibility and that the existence of these links makes them feasible for communication. With this simple model of point-to-point connectivity, maintenance of network integrity reduces to finding robot motions that guarantee the connectivity graph retains a single connected component [1], [2], [3], [4]. These methods are able to leverage network-wide connectivity indicators such as the second largest eigenvalue of the graph Laplacian or the k-connectivity of the network's graph. These indicators are tractable and often attainable with distributed methods making these approaches appealing.

However, it is well accepted in the wireless networking literature that binary channel models are not accurate representations of wireless links and several methods have been proposed that consider the reliability of point-to-point links in order to choose routes that optimize end-to-end reliability metrics [5], [6]. Additionally, point-to-point reliability metrics have been successfully incorporated into mobility control algorithms [7], [8].

The existence of graph connectivity only implies that a multihop path from any source to any destination exists but does not determine whether the network formed by the robots is able to support desired communication rates. A more accurate metric of network integrity relies on the achievability of target end-to-end communication rates. Since end-to-end communication rates do not depend solely on the spatial configuration of the system but also on the manner in which packets are routed through the network, the maintenance of network integrity requires methodologies that concurrently plan robot positions and solve the network routing problem [9].

Due to shadowing and small scale fading, even small variations in robot positions lead to significant changes in channel capability [10], [11]. Furthermore, the application of motion planning algorithms to the team state necessitates predictions of channel quality in future configurations that have not yet been visited and measured. Thus, we consider point-to-point communication rates as random variables and robust methods must be employed in order to consider the maintenance of end-to-end rates [12].

We propose an architecture for robust communication maintenance that addresses a wide range of multirobot spatial applications as depicted in Fig. 1. We then propose a solution to the concurrent routing and mobility control problem. That is, jointly finding robot configurations $\mathbf{x}$ with wireless network routing solutions $\boldsymbol{\alpha}$ (see Fig. 1) that address the task specification in terms of visiting required goal locations while reliably maintaining adequate network integrity.

In contrast to most existing methods for communication maintenance on a team of mobile robots which rely on reactive control algorithms, we develop a deliberative planning technique. In Section III we present a randomized motion planning approach that finds sequences of robot configurations that satisfy spatial objectives while robustly maintaining the end-to-end communication requirements. We

J. Fink is with the US Army Research Laboratory, Adelphi, MD, USA. jonathan.r.fink3.civ@mail.mil
A. Ribeiro is with the Department of Electrical & Systems Engineering, University of Pennsylvania, Philadelphia, PA, USA. aribeiro@seas.upenn.edu
V. Kumar is with the Department of Mechanical Engineering& Applied Mechanics, University of Pennsylvania, Philadelphia, PA, USA. kumar@seas.upenn.edu

978-1-4673-1405-3/12/$31.00 ©2012 IEEE

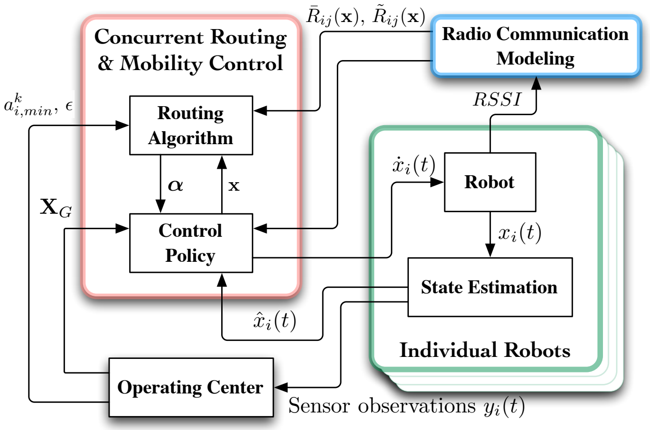

Fig. 1. System architecture. A human operator provides goal locations $\mathbf{X}_G$ for a lead robot as well as requirements on the rate of sensor updates that are relayed back. Individual robot components consist of the low-level controllers, estimators, and communication. The focus of this work is on developing concurrent methods for routing and mobility control.

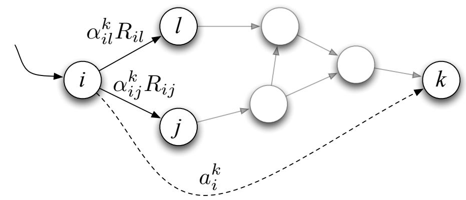

Fig. 2. Communication network. The nodes are deployed to support end-to-end rates from node i to fixed destination k. Routing variables α_{ij} determine the fraction of time node i sends packets to node j. R_{ij} is defined as the supported rate of the wireless channel from node i to node j.

finish with experimental results in Section IV and concluding remarks in Section V.

II. ARCHITECTURE FOR COMMUNICATION MAINTENANCE

The architecture depicted in Fig 1 is designed to support a fixed operating center that requires a lead agent to visit locations $\mathbf{X}_g$ while relaying back observations and maintaining network integrity metrics $(a_{i,min}, \epsilon)$. Individual robots are capable of state estimation, navigation, and pairwise measurements of point-to-point communication capability via received signal strength indicators (RSSI). For the purposes of this work, we make limited assumptions about the structure of radio communication modeling – only that point to point rate is a random variable and represented by mean $\bar{R}_{ij}$ and variance $\tilde{R}_{ij}$. The remainder of this section mathematically formulates our problem statement, details our assumptions on point-to-point communication rate modeling, and introduces the concept of robust network routing.

A. Problem Statement

Consider a team of N robots and denote their positions as $x_i \in \mathbb{R}^2$, for $i = 1, \ldots, N$. The robots are kinematic and fully controllable which allows us to consider simple mobility models of the form $\dot{x}_i(t) = u_i(t)$, where $u_i(t)$ is the control input to robot i. A human operator is located at the fixed operation center that we index as $i = 0$ at position x_0. Further define the vector $\mathbf{x} := (x_0, \ldots, x_N) \in \mathbb{R}^{2(N+1)}$ to group all positions.

The spatial task assigned to the team is specified through a generic scalar convex task potential function $\Psi : \mathbb{R}^{2(N+1)} \to \mathbb{R}$. E.g., when a designated leader agent ℓ must visit a target location $x_{\ell,g} \in \mathbb{R}^2$, we define $\Psi(\mathbf{x}) = \|x_\ell - x_{\ell,g}\|^2$.

While controlling to a physical configuration, the system must also maintain a certain network integrity that is specified by desired end-to-end communication rates which represent the flow of data from members of the team to the fixed operating center. For the purposes of this paper, we will focus on a single flow of information to a fixed operating center but our problem formulation and solution is generic to support several information flows [13]. The variable $a_{i,min}$ represents the required end-to-end communication rate between node i and the operating center.

We model point-to-point connectivity through a rate function $R_{ij}(\mathbf{x}) = R_{ij}(x_i, x_j)$ that determines the amount of information that an agent at x_i can send to an agent at x_j. Since direct communication between the source and the destination of an information flow is not always possible, terminals self-organize into a multihop network to relay packets for each other. Packet relaying is determined by routing variables α_{ij} which describe the fraction of time node i spends transmitting data to node j as depicted in Fig 2.

Thus, the product $\alpha_{ij} R_{ij}(\mathbf{x})$ determines the rate of point-to-point information flow from i to j. The information rate $a_i(\boldsymbol{\alpha}, \mathbf{x})$ available for source i is the difference between outgoing and incoming rates at node i

$$a_i(\boldsymbol{\alpha}, \mathbf{x}) = \underbrace{\sum_{j=0}^{N} \alpha_{ij} R_{ij}(\mathbf{x})}_{\text{Outgoing packets}} - \underbrace{\sum_{j=1}^{N} \alpha_{ji} R_{ji}(\mathbf{x})}_{\text{Incoming packets}}, \tag{1}$$

where we define the vector $\boldsymbol{\alpha} \in \mathbb{R}^{N^2}$ grouping all routing variables α_{ij}.

Routing variables $\boldsymbol{\alpha}$ and configuration-dependent rates $R_{ij}(\mathbf{x})$ determine the set $a_i(\boldsymbol{\alpha}, \mathbf{x})$ of end-to-end communication rates from each node i as per (1). The task specification requires that end-to-end rates exceed the minimum threshold $a_{i,\min}$. Therefore, integrity of the communication network necessitates that $a_i(\boldsymbol{\alpha}, \mathbf{x}) \geq a_{i,\min}$ for all i. Notice that $a_i(\boldsymbol{\alpha}, \mathbf{x})$ is a function of positions $\mathbf{x}$ and routing variables $\boldsymbol{\alpha}$. Thus, end-to-end connectivity is maintained by control of positions $\mathbf{x}$ *and* network routes $\boldsymbol{\alpha}$.

The primary mobility control problem is to find a sequence of robot configurations $\mathbf{x}(t)$ such that at time t_f the team configuration $\mathbf{x}(t_f)$ satisfies task completion in that $\Psi(\mathbf{x}(t_f))$ is minimized. The maintenance of network integrity dictates that $a_i(\boldsymbol{\alpha}(t), \mathbf{x}(t)) \geq a_{i,min}$ for all $t \in [t_0, t_f]$. In practice, it is difficult to ensure that the network integrity constraint is satisfied. As seen in (1), rates $a_i(\boldsymbol{\alpha}, \mathbf{x})$ depend on point-to-point link reliabilities $R_{ij}(\mathbf{x})$ that are difficult to estimate. We address the specific form of R_{ij} in the following sections but for now it suffices to assume it is a random variable. The important observation is then that if point-to-point rates $R_{ij}(\mathbf{x})$ are random, so are the end-to-end rates $a_i(\boldsymbol{\alpha}, \mathbf{x})$ [cf. (1)]. Thus, we introduce a reliability tolerance ϵ and require that

$$\mathrm{P}\left[a_i(\boldsymbol{\alpha}, \mathbf{x}) \geq a_{i,\min}\right] \geq \epsilon. \tag{2}$$

We can write the full joint mobility and robust routing problem as

$$\begin{aligned} \min_{\boldsymbol{\alpha}(t), \mathbf{x}(t), t\in[0,t_f]} \quad & \Psi(\mathbf{x}(t_f)) \\ \text{subject to} \quad & \mathrm{P}\left[a_i(\boldsymbol{\alpha}(t), \mathbf{x}(t)) \geq a_{i,\min}\right] \geq \epsilon \\ & \mathbf{x}(t) = \mathbf{x}(0) + \int_0^t \dot{\mathbf{x}}(\tau)\ d\tau\ . \end{aligned} \tag{3}$$

The mobility portion of this problem inherits the same issues that are studied in the robotic motion planning literature. That is, the search for a sequence of obstacle-free configurations from the initial configuration to goal. However, the network integrity constraint leads to a coupling of the mobility and network routing problems and forces us to consider planning in a very high-dimensional joint space of positions and network routes. The general problem has $N^2 + 2N$ variables and $(\boldsymbol{\alpha}, \mathbf{x}) \in \mathbb{R}^{N^2} \times \mathbb{R}^{2(N+1)}$. We decompose this space by fixing $\mathbf{x}$ and choosing $\boldsymbol{\alpha}$ in a manner that optimizes the reliability $\mathrm{P}\left[a_i(\boldsymbol{\alpha}(t), \mathbf{x}(t)) \geq a_{i,\min}\right]$ in Section II-C. First, however, we provide more detail on our assumptions of point-to-point rate modeling.

B. Point-to-Point Rate Modeling

In general, we seek to develop a probabilistic model for the communication rate $R_{ij}(\mathbf{x}) = R_{ij}(x_i, x_j)$ between robots at position x_i and x_j. However, we also stress that the approaches described in this paper are agnostic to the particular form of this underlying rate model. In fact, we only rely on the ability to predict the mean $\bar{R}_{ij}(x_i, x_j)$ and variance $\tilde{R}_{ij}(x_i, x_j)$ of the supported communication rate. It is well known that the supported rate of communication is a function of the signal-to-noise ratio (SNR) [14]. We therefore focus on models of the received signal strength $P_R(x_i, x_j)$ that we cascade into models of the packet error rate and the supported communication rate $R_{ij}(x_i, x_j)$.

While many methods exist for the modeling of received signal strength [12], in this work we rely on a relatively coarse model based on the distance between source and receiver as well as the existence of a line-of-sight between the nodes. This model is most easily described on a logarithmic scale of received power $P_{R,\text{dBm}}(x_i, x_j) = 10\log(P_R(x_i, x_j))$ measured in dBm

$$P_{R,\text{dBm}}(x_i, x_j) = \underbrace{L_0 - 10n \cdot \log(\|x_i - x_j\|)}_{\text{Path loss}} - \underbrace{W(x_i, x_j)}_{\text{Shadowing}} - \underbrace{\mathcal{F}}_{\text{Fading}}, \tag{4}$$

where the term $\mathcal{F}$ is a zero-mean Gaussian random variable with variance $\sigma_{\mathcal{F}}^2$ modeling fading effects. The term L_0 is the measured power at a reference distance 1 m from the source, n is a path loss exponent, and $W(x_i, x_j)$ is a non-smooth function to model shadowing as a function of the number of obstacles between source and destination. Experimental data collected in an indoor office environment at the University of Pennsylvania leads to parameters $L_0 = -51$ dBm, $n = 2.1$, $W(x_i, x_j) = 0$ for line of sight links and $W(x_i, x_j) = 7.6$ dB for non-line-of-sight links, and $\sigma_{\mathcal{F}}^2 = 32$ dB^2. For brevity, the application of an approximation that relates received signal strength to bit-error-rate and, in turn, supported communication rate is omitted [14].

The form of received signal strength in (4) yields a suitable model for the supported communication rate between nodes at x_i and x_j that not only takes into account path loss and shadowing due to environmental obstructions but also models random fading that is the result of the multi-path phenomenon in indoor environments.

C. Robust Routing

A key component to our approach is the ability to decompose the search of the space defined by network routes and robot configurations $(\boldsymbol{\alpha}, \mathbf{x})$. We do this by determining an $\boldsymbol{\alpha}(\mathbf{x})$ that optimally satisfies the reliability constraint (2) for the current configuration despite the uncertainty of point-to-point communication rates $R_{ij}(\mathbf{x})$.

A simple way to account for the uncertainty in $R_{ij}(\mathbf{x})$ is to discount individual $R_{ij}(\mathbf{x})$ in order to reduce the likelihood of having actual rates smaller than the estimated value. While this is a feasible solution, it results in underutilization of the communication network. Instead, we recall that the end-to-end rather than point-to-point failures are relevant. By splitting traffic and exploiting spatial redundancy, we can devise routes that guarantee small changes in end-to-end rates despite the large variability of point-to-point rates $R_{ij}(\mathbf{x})$ [13].

For a given configuration $\mathbf{x}$, we would like to find routes $\boldsymbol{\alpha} = \boldsymbol{\alpha}(\mathbf{x})$ that provide the maximum possible reliability. We do this by maximizing the minimum achievable rate with probability ϵ. Maximizing this minimum implies that the constraints in (3) are satisfied with significant slack and that there is liberty to change the physical configuration without violating communication constraints. To solve for such an optimal route $\boldsymbol{\alpha}(\mathbf{x})$, we can introduce a slack variable a_Δ and write the optimization problem

$$\begin{aligned} \boldsymbol{\alpha}(\mathbf{x}) = & \operatorname*{argmax}_{\boldsymbol{\alpha},\, a_\Delta} \quad a_\Delta \\ & \text{subject to} \quad \mathrm{P}\left[a_i(\boldsymbol{\alpha}, \mathbf{x}) \geq a_{i,min} + a_\Delta\right] \geq \epsilon. \end{aligned} \tag{5}$$

The form of the probability constraint is important when we write (5). By representing this probability based on the mean and variance statistics of $R_{ij}(\mathbf{x})$, it can be written as a second order cone constraint allowing us to solve this optimization as a second order cone program (SOCP) [12]. SOCPs are a particular class of convex optimization problem that can be solved by efficient polynomial time algorithms [15]. For the problems considered here, the computational complexity of these algorithms is represented as a polynomial function of the number of agents N as $\mathrm{O}\big((N^2)^{3.5}\big)$. In practical implementations, the N^2 term can be reduced by eliminating links for which rate estimates $\bar{R}_{ij}(\mathbf{x})$ are below a certain threshold.

III. PLANNING FOR MOBILITY

The original problem statement in (3) illustrates the full joint state space of network routes $\boldsymbol{\alpha}$ and positions $\mathbf{x}$ that must be determined by communication maintenance algorithms. We have demonstrated that for a fixed $\mathbf{x}$, we can efficiently compute the optimal $\boldsymbol{\alpha}(\mathbf{x})$. One approach to the optimization of (3) would then be to apply gradient descent algorithms and incrementally drive the system towards a goal configuration [12]. However, gradient-based approaches have major drawbacks with respect to local minima. These are due to obstacles in the environment and the complications of jointly optimizing $\boldsymbol{\alpha}$ and $\mathbf{x}$ under complex network topologies. Therefore, it is necessary to pursue deliberative motion planning that takes into account global constraints. In order to consider this type of approach, it is necessary to address some additional notation.

Let $\mathbf{X}$ be a bounded, connected open subset of $\mathbb{R}^{2N}$ that represents the joint state space for the team of robots where $\mathbf{x}_{init}$ is the initial configuration of the team. We then define the goal set for our physical task to be $\mathbf{X}_g = \{\mathbf{x} : \Psi(\mathbf{x}) < \Psi_{min} + \delta\}$ where Ψ_{min} is the minimum value of the potential function, δ is a parameter used to represent configurations *close* to this absolute minimum, and $\mathbf{X}_g \subset \mathbf{X}$. The obstacle region $\mathbf{X}_{obs}$ contains any configuration that places an individual robot on a physical obstacle and the infeasible region $\mathbf{X}_{inf}$ represents configurations where it is infeasible to satisfy the network constraint (2),

$$\mathbf{X}_{inf} = \left\{ \mathbf{x} : \mathrm{P}\left[a_i(\boldsymbol{\alpha}, \mathbf{x}) \geq a_{i,min}\right] \leq \epsilon \ \forall \boldsymbol{\alpha} \right\}. \quad (6)$$

The free space $\mathbf{X}_{free}$ is then $\mathbf{X} \setminus (\mathbf{X}_{obs} \cup \mathbf{X}_{inf})$. Finally, a path in $\mathbf{X}$ is parameterized by a scalar $s \geq 0$ and given by $\boldsymbol{\sigma} : [0, s] \to \mathbf{X}$. Concatenation of paths is defined by $\boldsymbol{\sigma} = \boldsymbol{\sigma}_1 | \boldsymbol{\sigma}_2$. A feasible path, and solution to our global-planning problem, is then $\boldsymbol{\sigma} : [0, s] \to \mathbf{X}_{free}$ such that $\boldsymbol{\sigma}(0) = \mathbf{x}_{init}$ and $\boldsymbol{\sigma}(s) \in \mathbf{X}_g$.

The dimensionality of our problem and the high computational cost of verifying a state is in X_{free} makes deterministic search algorithms impractical. Instead, we turn to probabilistic search methods that offer good space-filling properties and efficient exploration of an unknown space – e.g., rapidly–exploring random tree (RRT) algorithms [16]. The basic structure of an RRT, as detailed in Algorithm 1, is to start with an initial state $\mathbf{x}_{init}$ and expand to fully explore the workspace, adding states in a tree-like structure until a feasible point is added such that $\mathbf{x} \in X_g$. The tree is expanded by picking a random state $\hat{\mathbf{x}} = \textsc{RandomState}(X, \mathcal{T})$, finding the closest point on the existing tree $\mathbf{x}_{\min} = \textsc{Nearest}(\mathcal{T}, \hat{\mathbf{x}})$, and attempting to add a new point by extending from $\mathbf{x}_{\min}$, $\mathbf{x} = \textsc{Extend}(\mathbf{x}_{\min}, \hat{\mathbf{x}})$.

There are two difficulties that arise when applying standard RRT algorithms to solve the specific high-dimensional network connectivity problem in (3): (i) the verification of feasible states as Extend is used to expand the tree towards $\hat{\mathbf{x}}$ and (ii) the prohibitive cost of uniformly exploring X_{free} for our high-dimensional problem with slow-to-compute constraints.

Algorithm 1 Structure of the rapidly exploring random tree algorithm

Require: Initial state $\mathbf{x}_{init}$, goal region X_g, representation of the bounded configuration space X.

$\mathcal{T}.init(\mathbf{x}_{init})$
while $i < N$ **do**
  $\hat{\mathbf{x}} \leftarrow \textsc{RandomState}(X, \mathcal{T})$
  $\mathbf{x}_{\min} \leftarrow \textsc{Nearest}(\mathcal{T}, \hat{\mathbf{x}})$
  if $\mathbf{x}_{new} \leftarrow \textsc{Extend}(\mathbf{x}_{\min}, \hat{\mathbf{x}})$ **then**
    $\mathcal{T}$.add_vertex($\mathbf{x}_{new}$)
    $\mathcal{T}$.add_edge($\mathbf{x}_{\min}, \mathbf{x}_{new}$)
    if $\mathbf{x}_{new} \in X_g$ **then**
      return $\mathcal{T}$
    end if
  end if
end while
return $\mathcal{T}$

A. *Efficient verification of feasible states*

The $\textsc{Extend}(\mathbf{x}_{from}, \mathbf{x}_{to})$ algorithm attempts to virtually drive the system from $\mathbf{x}_{from}$ towards $\mathbf{x}_{to}$ by successively verifying that points along the line connecting $\mathbf{x}_{from}$ and $\mathbf{x}_{to}$ are in X_{free}. It returns the state $\mathbf{x}_{new}$ as the closest state to $\mathbf{x}_{to}$ such that all states sampled with precision $\Delta\mathbf{x}$ between $\mathbf{x}_{from}$ and $\mathbf{x}_{new}$ are in X_{free}. In traditional motion planning applications, verification that $\mathbf{x} \in X_{free}$ is based on an algebraic constraint or collision query via efficient computational methods, e.g., [18]. While the necessary computation to determine $\mathbf{x} \notin X_{obs}$ is typically small, computation of $\mathbf{x} \notin X_{inf}$ requires a solution of the SOCP (5) and can be costly.

To limit re-computation of optimal $\boldsymbol{\alpha}$, we store the full state $(\boldsymbol{\alpha}, \mathbf{x})$ for every node in $\mathcal{T}$ and rely on the fact that an optimal robust routing solution $\boldsymbol{\alpha}(\mathbf{x})$ will be feasible for neighboring states and will often be a feasible solution for the entire trajectory to $\mathbf{x}_{new}$. Thus, in Algorithm 2, we recompute $\boldsymbol{\alpha}$ only when reliability drops below the desired threshold to warrant the additional computation.

B. *Biased space sampling*

Random states $\hat{\mathbf{x}}$ are chosen to sample the space $\mathbf{X} \subset \mathbb{R}^{2N}$ according to a probability distribution $p_{\mathbf{x}}(\mathbf{x})$ representing the belief about configuration $\mathbf{x}$ being part of a feasible path $\boldsymbol{\sigma}(s)$. If nothing is known about $\boldsymbol{\sigma}(s)$, we choose $p_{\mathbf{x}}(\mathbf{x})$ uniform in the space X. In general, the final configuration is known in that $\boldsymbol{\sigma}(s) \in \mathbf{X}_g$. We can then bias the distribution towards $\mathbf{X}_g$ by making

$$p_{\mathbf{x}}(\mathbf{x}) = \frac{p_g}{v(\mathbf{X}_g)} \mathbb{I}\{\mathbf{x} \in \mathbf{X}_g\} + \frac{1 - p_g}{v(\mathbf{X} \setminus \mathbf{X}_g)} \mathbb{I}\{\mathbf{x} \notin \mathbf{X}_g\}. \quad (7)$$

where larger values of p_g make $\hat{\mathbf{x}}$ more likely to hit $\mathbf{X}_g$. Goal biasing as in (7) improves efficiency of RRT algorithms by

Algorithm 2 EXTEND($\mathbf{x}_{from}, \mathbf{x}_{to}$)

Require: Initial state $\mathbf{x}_{from}$, desired final state $\mathbf{x}_{to}$, verify segment with resolution $\Delta\mathbf{x}$

$\mathbf{x}_{new} \leftarrow \mathbf{x}_{from}$
$\boldsymbol{\alpha} \leftarrow$ argmax (5) for rates R_{ij} in configuration $\mathbf{x}_{new}$.
while $\mathbf{x}_{new} \neq \mathbf{x}_{to}$ **and** $\boldsymbol{\alpha} \neq \emptyset$ **do**
    $\mathbf{x}_{new} \leftarrow \mathbf{x}_{new} + \Delta\mathbf{x}$
    if $\min \{\mathrm{P}[a_i(\boldsymbol{\alpha}, \mathbf{x}_{new}) \geq a_{i,min}]\} \leq \epsilon$ **then**
        {Recompute $\boldsymbol{\alpha}$ if reliability low}
        $\boldsymbol{\alpha} \leftarrow$ argmax (5) for rates $R_{ij}(\mathbf{x}_{new})$.
    end if
end while
if $\mathbf{x}_{new} = \mathbf{x}_{from}$ **then**
    return $\emptyset$
else
    return $\mathbf{x}_{new}$
end if

reducing the number of samples necessary to find a feasible path $\boldsymbol{\sigma}(s)$ in the high dimensional space $\mathbf{X} \subset \mathbb{R}^{2N}$.

In our telepresence application where $\Psi(\mathbf{x}) = \|x_\ell - x_{\ell,g}\|^2$, the goal position of the leader $x_{\ell,g}$ is known, but the positions of the remaining robots are free. This means that volume of $\mathbf{X}_g$ is comparable to the volume of $\mathbf{X}$. In this case goal biasing offers little improvement over uniform sampling. That is, goal biasing would reduce the exploration cost along the components associated with x_ℓ but keep the cost of exploring the remaining $2(N-1)$ dimensions fixed. To further reduce exploration cost in this case, we construct a prediction $\tilde{\mathbf{X}}_g \subset \mathbf{X}_g$ of the final configuration and bias sampling towards this configuration prediction by recomputing the sampling distribution $p_{\mathbf{x}}(\mathbf{x})$ in (7) with $\tilde{\mathbf{X}}_g$ instead of $\mathbf{X}_g$.

Constructing a final configuration prediction $\tilde{\mathbf{X}}_g$ is task-specific. We describe here a method applicable to the telepresence–type application where the final position of a lead node is specified. To determine the prediction $\tilde{\mathbf{X}}_g$, we determine predictions $\tilde{\mathbf{X}}_{i,g}$ for each robot and compute $\tilde{\mathbf{X}}_g$ as the Cartesian product of these individual sets. Notice that for the lead robot we can make $\tilde{\mathbf{X}}_{\ell,g} = \mathbf{X}_{\ell,g} = \{x_\ell \in \mathbb{R}^2 : \|x_\ell - x_{\ell,g}\| < \delta\}$.

Observe now that $\mathbf{X} \subset \mathbb{R}^{2N}$ is the Cartesian product $\mathbf{X} = \prod_{i=1}^N \mathbf{X}_i$ of the N decoupled spaces $\mathbf{X}_i \in \mathbb{R}^2$ corresponding to each individual robot. If we further assume a homogeneous team of robots, then all robots operate in the same space $\mathbf{X}_i = \mathbf{Y}$, with a common set of physical obstacles $\mathbf{Y}_{obs}$, and consequently a common free space $\mathbf{Y}_{free} = \mathbf{Y} \setminus \mathbf{Y}_{obs}$. It follows that the joint free space $\mathbf{X}_{free}$ is also a Cartesian product of N identical sets $\mathbf{Y}_{free}$ minus those configurations for which a network cannot be established with sufficient reliability,

$$\mathbf{X}_{free} = (\mathbf{Y}_{free})^N \setminus \mathbf{X}_{inf}. \quad (8)$$

While infeasible network configurations are captured by $\mathbf{X}_{inf}$ as given in (6), $\mathbf{X}_{free}$ can otherwise be described by the free space of individual robots.

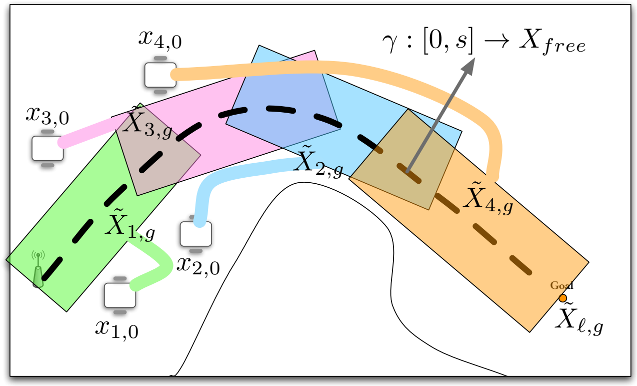

Fig. 3. Illustration of the biased space sampling. Since we only know one component of the goal state $x_{g,\ell}$ and it is expensive to expand our search space in the high-dimensional state of the entire system, it is beneficial to bias our search towards configurations that are deemed likely to succeed.

To exploit this observation, we first determine an obstacle free path $\gamma : [0, s] \rightarrow \mathbb{R}^2$ such that $\gamma(0) = x_0$ is the position of the operating center and $\gamma(s) \in \mathbf{X}_{\ell,g}$. Since the dimensionality of the space and the goal set $\mathbf{X}_{\ell,g}$ are small, it is possible to find this path with small computational cost using modern discrete planning algorithms [19]. The obstacle-free path $\gamma : [0, s] \rightarrow \mathbf{Y}_{free}$ is split into $N-1$ equal length segments γ_k such that $\gamma_k : [0, s_k] \rightarrow \gamma : [\frac{ks}{N-1}, \frac{(k+1)s}{N-1}]$. The ith robot is then assigned to a segment by the function $k(i)$ based on euclidian distance to its midpoint such that $\sum_{i \neq 0,\ell} \|\gamma_{k(i)}(s/2) - x_{i,0}\|$ is minimized; see Fig. 3. Segments are then enlarged to define the region $\tilde{\mathbf{X}}_{i,g}$ for $i \neq 0, \ell$. Since this is a heuristic for the goal configuration, the only requirement on $\tilde{\mathbf{X}}_{i,g}$ is that $\gamma_{k(i)} : [0, s] \rightarrow \tilde{\mathbf{X}}_{i,g}$. A typical choice is

$$\tilde{\mathbf{X}}_{i,g} = \{x_i : \min_s \|x_i - \gamma_{k(i)}(s)\| < \tilde{d}_g\}$$

where $\tilde{d}_g$ is a parameter controlling the enlarged size of $\tilde{\mathbf{X}}_{i,g}$. The predicted final configuration is then computed as the Cartesian product $\tilde{\mathbf{X}}_g = \prod_{i=1}^N \tilde{\mathbf{X}}_{i,g}$.

This procedure is summarized in Algorithm 3. In lines 1–3, the predicted goal configuration $\tilde{\mathbf{X}}_g$ is constructed. A random sample $\hat{\mathbf{x}}$ is then drawn uniformly from $\tilde{\mathbf{X}}_g$ with probability p_g or from $X \setminus \tilde{\mathbf{X}}_g$ otherwise. It should be noted that the construction of $\tilde{\mathbf{X}}_g$ described above is based on the heuristic that a feasible goal configuration in an environment with obstacles will *resemble* a line-of-sight communication chain. Increasing the size of $\tilde{\mathbf{X}}_g$ with large values of $\tilde{d}_g$ limits the implication of this assumption.

IV. RESULTS

We rely on a centralized implementation of the algorithms for concurrent solutions to the robust routing and mobility control problem that are presented in this paper. Since the algorithms implicitly maintain a connected network of agents, coordinated control commands can be robustly routed through the wireless network. Furthermore, a centralized implementation is not a shortcoming for the problem sizes

Algorithm 3 RANDOMSTATE(X)

Require: Configuration space description X, obstacle-free path $\gamma(s) \to \mathbb{R}^2$ such that $\gamma(0) = x_0$ and $\gamma(s) = x_{\ell,g}$, probability p_g.

$\tilde{\mathbf{X}}_{\ell,g} = \{x_\ell \in \mathbb{R}^2 : \|x_\ell - x_{\ell,g}\| < \delta\}$
$\tilde{\mathbf{X}}_{i,g} \leftarrow \text{Enlarge}\left(\gamma_{k(i)}\right)$
$\tilde{\mathbf{X}}_g \leftarrow \prod_{i=1}^{N} \tilde{\mathbf{X}}_{i,g}$
$p \leftarrow \text{Uniform}[0,1]$
if $p > p_g$ **then**
  $\hat{\mathbf{x}} \leftarrow \text{Uniform}(X \setminus \tilde{\mathbf{X}}_g)$
else
  $\hat{\mathbf{x}} \leftarrow \text{Uniform}(\tilde{\mathbf{X}}_g)$
end if
return $\hat{\mathbf{x}}$

we consider, e.g. 10 or fewer agents, since it is tractable to exchange and maintain global state.

All experiments are conducted on the *Scarab* ground platform at the University of Pennsylvania; see Fig. 4 inset. The robots are capable of accurate self-localization throughout the indoor environment and are equipped with off-the-shelf *Zigbee* radios that provide basic point-to-point communication capabilities in the 2.4 GHz spectrum.

Each experiment consists of a telepresence-type task, e.g. (3), that requires a single lead robot, indexed by ℓ, to visit one or more locations in the environment while maintaining a desired end-to-end communication rate $a_{\ell,min}$ with reliability ϵ to a fixed operating center. The algorithms introduced in this paper yield feasible configurations for the team – $\boldsymbol{\alpha}(t)$ and $\mathbf{x}(t)$ which represent the network and physical configurations respectively. During an experiment, each robot probes the communication channels with its neighbors to determine actual instantaneous measurements of the point-to-point received signal strength at a rate of 5 Hz. This data is logged locally and aggregated after each experiment to compute the supported communication rate $\hat{R}_{ij}(t)$ between node i and j at time t. Using these measurements in conjunction with the network routing solution $\boldsymbol{\alpha}(t)$, we can compute lower bounds on the actual achievable end-to-end rate at time t for each node i, $\hat{a}_i(\boldsymbol{\alpha}(t), \mathbf{x}(t))$.

Recall that the problem statement in (3) requires that $\mathrm{P}\left[a_i(\boldsymbol{\alpha}(t), \mathbf{x}(t)) \geq a_{i,\min}\right] \geq \epsilon$ for all nodes i. Thus, in our experimental verification each node must be able to maintain

$$\hat{a}_i(\boldsymbol{\alpha}(t), \mathbf{x}(t)) > a_{i,min} \tag{9}$$

with probability ϵ. To achieve the desired end-to-end rates, *all* nodes in the team must satisfy this constraint. Thus, in experimental analysis we will evaluate (9) across the duration of the experiment to determine the percent of time $\hat{a}_i(\boldsymbol{\alpha}(t), \mathbf{x}(t)) > a_{i,min}$ and use this as a metric for the success of that trial.

A. Experiments

Figure 4 depicts the series of waypoints that the lead node, x_5 must visit. Four additional mobile nodes, x_1, x_2, x_3, x_4 are available to relay data back to the fixed access point indicated in the lower left of Fig. 4 with end-to-end rate of $a_{5,min} = 0.25$ with probability $\epsilon = 0.75$. Each relay node must maintain end-to-end rates greater than zero. The predicted and measured end-to-end rates of each node are depicted in Fig. 7.

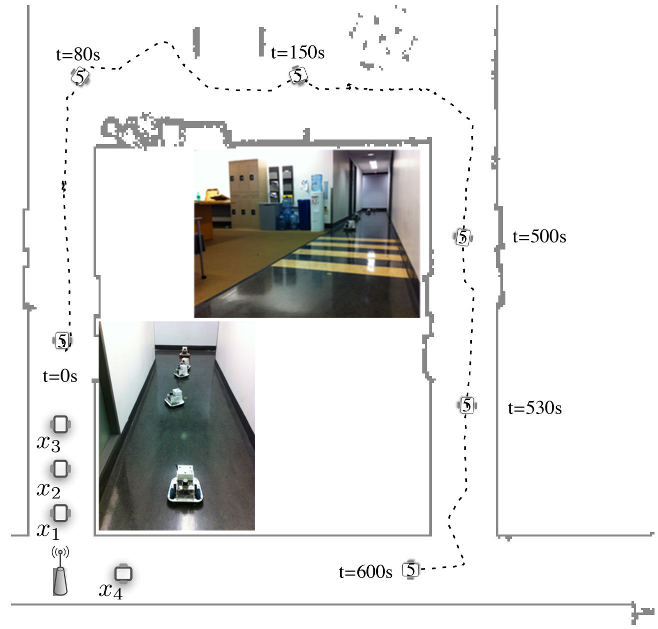

Fig. 4. The task specification for the Levine building experiment. It requires the lead node, x_5 to follow a sequence of waypoints. The initial configuration of the team is depicted above.

First, notice that the instantaneous rate $\hat{a}_5(\boldsymbol{\alpha}(t), \mathbf{x}(t))$ is almost always above its minimum threshold of $a_{5,min} = 0.25$. In fact, it drops below the minimum threshold only 2.9% of the time, well within the allowable reliability $\epsilon = 25\%$ for this problem specification. However, for that rate to be maintained in an end-to-end sense across the network, each node must be able to support the necessary rate margin $a_{i,min}$. The corresponding fraction of time spent below the minimum threshold for each of the instantaneous node rates $\hat{a}_1, \hat{a}_2, \hat{a}_3, \hat{a}_4$ is 9.2%, 0.8%, 0.3%, and 0.6%. This means the desired rates are not supported during, at most, 13.8% of the time.

Representative network configurations are depicted in Fig. 6. In Fig. 6a, at $t = 100$ s, the predicted goal state $\tilde{\mathbf{X}}_g$ assumes the shortest line of sight path which is the left hallway, i.e. a result similar to what one would expect from reactive methods. As the system transitions to Fig. 6b, where the lead node x_5 has been tasked to a waypoint in the right hallway, the prediction for $\tilde{\mathbf{X}}_g$ shifts to a chain of relays going through the right hallway. This shift in the basic topology of $\tilde{\mathbf{X}}_g$ re-focuses exploration of the joint state space so that x_4 moves towards a configuration that will *lower* the performance of the network over the short term. As node x_5 completes the desired loop, it utilizes x_4 as a relay channel and is able to maintain the desired end-to-end rate. This dramatic shift in network topology would not be possible with a purely reactive method and illustrates one of the advantages of our deliberative approach.

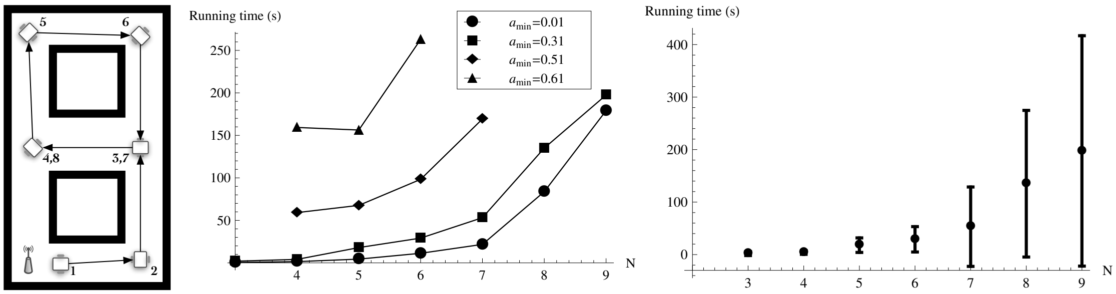

(a) (b) (c)

Fig. 5. Running time for benchmark environment with global planner solving the task depicted in (a). The average running time for tasks with different a_{min} and number of robots N are depicted in (b). The variance of the running time for a particular task is depicted in (c).

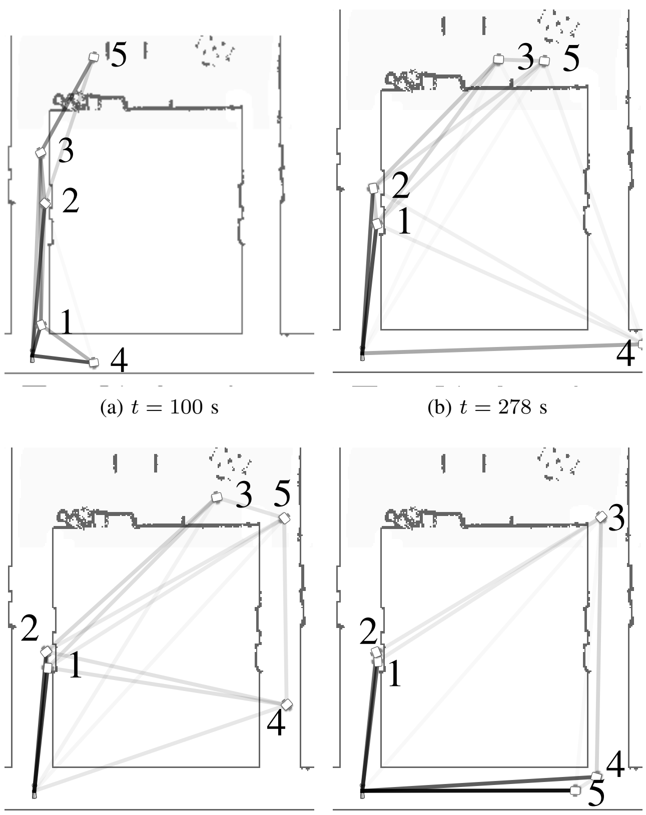

(a) $t = 100$ s (b) $t = 278$ s

(c) $t = 482$ s (d) $t = 622$ s

Fig. 6. Snapshots from the sequence of feasible network configurations that satisfy the task depicted in Fig. 4. Line weight indicates the expected amount of information to be transmitted over that point-to-point link.

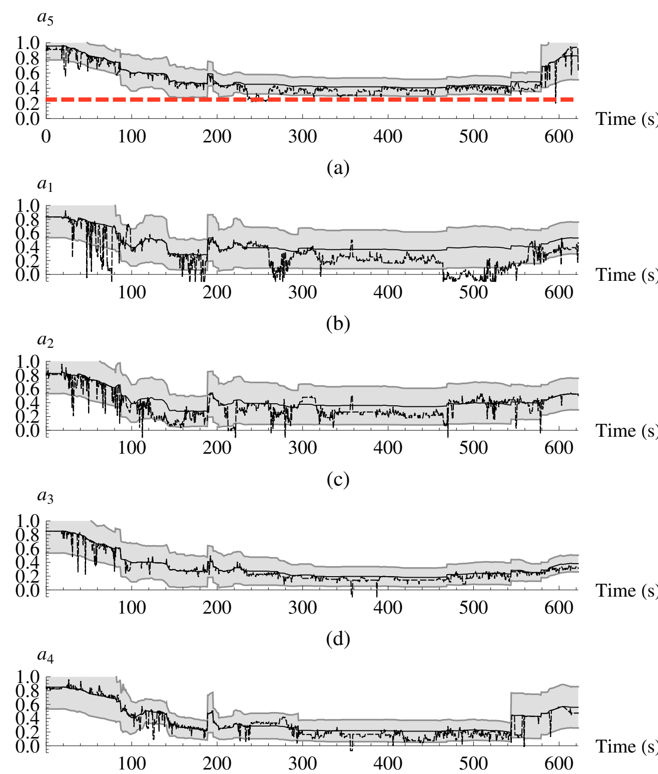

(a) (b) (c) (d) (e)

Fig. 7. The end-to-end rates of the nodes during the Levine building experiment depicted in Fig. 4. (a) depicts the prediction, $\bar{a}_5, \tilde{a}_5$, and instantaneous, $\hat{a}_5$, end-to-end rate for the leader and (b) – (e) depict the instantaneous rates of the relay nodes. In each plot, the solid line with shaded envelope depicts $\bar{a}_i$ and variations that occur with probability $\epsilon = 0.75$ based on $\tilde{a}_i$. The dashed black line represents the instantaneous end-to-end rate $\hat{a}_i$. The dashed red link in (a) depicts the threshold $a_{5,min} = 0.25$.

B. Benchmarking

We also study the computational complexity of the proposed approach through a benchmark task that can be solved many times with different problem parameterizations. The task, depicted in Fig. 5a, requires the lead robot to visit a series of positions in the environment (labeled 1–8) while communicating data at a specified rate, a_{min}, to the operating center located near waypoint 1. We parameterize the task by the number of robots N and the end-to-end rate of communication that must be maintained a_{min} while fixing the desired reliability $\epsilon = 0.8$.

The performance is measured by the running time to compute the series of network configurations necessary for the lead robot to visit its sequence of waypoints. The average performance is depicted in Fig. 5b based on 10 trials per task parameterization. As expected, increasing the number of robots adds to the complexity of both the individual SOCP solutions as well as the randomized search algorithm. Increasing the minimum end-to-end rate, a_{min}, has a similar effect on the complexity. Intuitively, increased a_{min} increases complexity because it reduces the number of feasible configurations that can support the required rate a_{min}. This increases the planning effort necessary to explore the workspace and find a feasible path. Additionally, as we increase the end-to-end rate requirement, more agents are necessary for task completion.

Randomized planning algorithms can only offer the guarantee of probabilistic completeness. Since there is no precise way to determine when a task cannot be solved with the current configuration, we test for task infeasibility by stopping the planning process after a specified timeout period. For the purposes of this benchmarking, that timeout is 300 s for each subtask. An artifact of this timeout is that tasks in extremely complex spaces (e.g. $a_{min} = 0.51, N = 8$) are not solved though we know a solution exists (e.g. the solution for $a_{min} = 0.51, N = 7$ is a subset of the possible solutions with $N = 8$). As the complexity of the task increases, so does the variance of the running time as depicted in Fig. 5c for a particular task, $a_{min} = 0.61$.

V. CONCLUSION

We propose an architecture and algorithms to address the full communication maintenance problem that includes solving for relay node mobility and network routing to maintain end-to-end connectivity. Because the performance of point-to-point wireless links is difficult to predict, we adopt a stochastic model for supported rates. By developing an optimal robust solution to the network routing problem for a given physical configuration, we are able to reduce the dimensionality of our search problem to only be the joint state-space of the robot positions. This allows us to pursue the connectivity maintenance problem in the framework of a high-dimensional motion plan where feasible states rely on not only the free space of the environment but also the feasibility of a robust wireless networking solution.

We present, to the best of our knowledge, the first example of an experimental verification for a communication maintenance system that relies on an end-to-end rate metric for network integrity. Furthermore, we are able to do this with limited assumptions about the model for point-to-point achievable rates. In fact, our experiments succeed with a coarse rate model that can be applied to a wide range of environments. Finally, our experimental results illustrate the value in pursuing global search methods rather than reactive gradient-based methods in that we are able to find a sequence of network configurations that would not emerge from local optimization of network integrity.

The use of randomized planning techniques implicitly casts our formulation as a feasibility problem. Unfortunately, recent developments found in [20] for optimal RRT–based planning are not directly applicable due to the high-cost of verifying inter-state connectivity in our problem. Future work will focus on the incorporation of techniques from gradient-based methods to decrease the time-to-plan and increase the optimality of our solutions. We will also pursue decentralized approaches that can be applied to larger team sizes.

ACKNOWLEDGMENTS

We would like to thank Brian Sadler and Nathan Michael for numerous discussions and their help in the development of the ideas and experiments described in this paper. Additionally, this research was supported by ARL Grant W911NF-08-2-0004, and ONR Grants N00014-08-1-0696 and N00014-07-1-0829.

REFERENCES

[1] M. Ji and M. Egerstedt, "Distributed coordination control of multiagent systems while preserving connectedness," *IEEE Transactions on Robotics*, vol. 23, no. 4, pp. 693–703, August 2007.

[2] E. Stump, A. Jadbabaie, and V. Kumar, "Connectivity management in mobile robot teams," in *Proc. IEEE International Conference on Robotics and Automation*, Pasadena, CA, May 2008, pp. 1525–1530.

[3] M. Schuresko and J. Cortes, "Distributed motion constraints for algebraic connectivity of robotic networks," *Journal of Intelligent and Robotic Systems*, vol. 56, no. 1-2, pp. 99–126, September 2009.

[4] O. Tekdas *et al.*, "Robotic routers: Algorithms and implementation," *The International Journal of Robotics Research*, vol. 29, no. 1, p. 110, 2010.

[5] A. Ribeiro, Z.-Q. Luo, N. D. Sidiropoulos, and G. B. Giannakis, "Modelling and optimization of stochastic routing for wireless multihop networks," in *Proc. 26th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, Anchorage, Alaska, May 2007, pp. 1748–1756.

[6] D. DeCouto, D. Aguayo, J. Bicket, and R. Morris, "A high-throughput path metric for multihop wireless routing," in *International ACM Conference on Mobile Computing and Networking*, San Diego, CA, September 2006, pp. 134–146.

[7] A. Hsieh, A. Cowley, V. Kumar, and C. J. Taylor, "Towards the deployment of a mobile robot network with end-to-end performance guarantees," in *IEEE Int'l Conf. on Robotics and Automation*, Orlando, Florida, May 16-18 2006.

[8] Y. Mostofi, "Communication-aware motion planning in fading environments," in *IEEE International Conference on Robotics and Automation*, 2008, pp. 3169–3174.

[9] M. Zavlanos, A. Ribeiro, and G. Pappas, "Mobility and routing control in networks of robots," in *Conference on Decision Control*, Atlanta, GA, December 15–17 2010.

[10] Y. Mostofi, A. Gonzalez-Ruiz, A. Gaffarkhah, and L. Ding, "Characterization and modeling of wireless channels for networked robotic and control systems - a comprehensive overview," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009, pp. 4849–4854.

[11] J. Fink, N. Michael, A. Kushleyev, and V. Kumar, "Experimental characterization of radio signal propagation in indoor environments with application to estimation and control," in *International Conference on Intelligent Robots and Systems (IROS)*, St. Louis, MO, Oct. 2009.

[12] J. Fink, "Communication for teams of networked robots," Ph.D. dissertation, University of Pennsylvania, 2011.

[13] Y. Wu, A. Ribeiro, and G. B. Giannakis, "Robust routing in wireless multi-hop networks," in *Conference on Information Sciences and Systems*. Johns Hopkins Univ., Mar. 14-16, 2007, pp. 637–642.

[14] R. Shorey, A. Ananda, M. C. Chan, and W. T. Ooi, *Mobile, Wireless and Sensor Networks: Technology, Applications and Future Directions*. Wiley, John & Sons, Inc., 2006.

[15] M. Lobo, L. Vandenberghe, S. Boyd, and H. Lebret, "Applications of second-order cone programming," *Linear Algebra and its Applications*, vol. 284, no. 1-3, pp. 193–228, 1998.

[16] J. Kuffner Jr and S. LaValle, "RRT-connect: An efficient approach to single-query path planning," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2000.

[17] A. Atramentov and S. M. LaValle, "Efficient nearest neighbor searching for motion planning," in *Proceedings IEEE International Conference on Robotics and Automation*, 2002, pp. 632–637.

[18] M. C. Lin, D. Manocha, J. Cohen, and S. Gottschalk, "Collision detection: Algorithms and applications," in *Algorithms for Robotic Motion and Manipulation*, J.-P. Laumond and M. H. Overmars, Eds. Wellesley, MA: A.K. Peters, 1997, pp. 129–142.

[19] M. Likhachev, G. Gordon, and S. Thrun, "Ara*: Anytime a* with provable bounds on sub-optimality," *Advances in Neural Information Processing Systems*, vol. 16, 2004.

[20] S. Karaman and E. Frazzoli, "Incremental Sampling-based Algorithms for Optimal Motion Planning," in *Robotics: Science and Systems (RSS)*, 2010.