

Maintaining Network Connectivity and Performance in Robot Teams

**M. Ani Hsieh, Anthony Cowley,
Vijay Kumar, and Camillo J. Taylor**
*GRASP Laboratory
University of Pennsylvania
Philadelphia, Pennsylvania 19104, USA
e-mail: mya@grasp.upenn.edu,
e-mail: acowley@grasp.upenn.edu,
e-mail: kumar@grasp.upenn.edu,
e-mail: taylor@grasp.upenn.edu*

Received 7 November 2006; accepted 15 August 2007

In this paper, we present an experimental study of strategies for maintaining end-to-end communication links for tasks such as surveillance, reconnaissance, and target search and identification, where team connectivity is required for situational awareness. Our main contributions are threefold: (a) We present the construction of a radio signal strength map that can be used to plan multi-robot tasks, and also serve as useful perceptual information. We show how a nominal model of an urban environment obtained by aerial surveillance, is used to generate strategies for exploration. (b) We present reactive controllers for communication link maintenance; and (c) we consider the differences between monitoring signal strength versus data throughput. Experimental results, obtained using our multi-robot testbed in three representative urban environments are presented with each of our main contributions. © 2008 Wiley Periodicals, Inc.

1. INTRODUCTION

There is growing interest in the convergence of the areas of multi-agent robotics and sensor networks. The goal of such work is to develop networks of sen-

sors and robots that can perceive their environment and respond to it. Such systems should anticipate the information needs of network users while repositioning and organizing themselves to best acquire and deliver that information. While there is significant literature on multi-robot control (Fierro, Song, Das & Kumar 2002), sensing (Cortes, Martinez, Karatas & Bullo, 2002), planning (Guo & Parker, 2002), and localization (Roumelilotis & Bekey, 2002), the focus has historically been on control and perception where it is

Contract grant sponsor: National Science Foundation. Contract grant number: NSF-IIS-0427313, NSF CCR02-05336
Contract grant sponsor: ARO
Contract grant number: ARO-DAAD19-02-01-0383
Contract grant sponsor: DARPA
Contract grant number: DARPA MARS NBCH1020012

assumed that robots can freely and successfully communicate with one another. Although control is necessary for successful mission execution, communication is essential in order to achieve coordination and cooperation within multi-robot teams. Much of the research in the mobile wireless network community has been devoted to the development of novel algorithms to handle packet routing, such as Xue & Nahrstedt (2004) and Agüero, de-las Heras-Quiros, Canas & Matallan (2003), resource allocation (Feistel & Stanczak, 2005), and bandwidth management (Nahrstedt, Shah & Chen, 2004) for mobile nodes. However, control of mobile robot teams provides us with the capability to shape the team's communication needs based on continuous evaluation of the demands on the network (Basu & Redi, 2004b).

In recent years, the communication network has evolved from being just a medium of information transmission to an actual sensor, where properties like connectivity and signal strength are used to maintain the quality of the medium (Sweeney, Brunette, Yang & Grupen, 2002, Powers & Balch, 2004, Basu & Redi, 2004b, Howard, Siddiqi & Sukhatme, 2003). Agents can use communication links to infer their individual locations with respect to those of their neighbors and other landmarks. Simultaneously, agents may also control their positions and orientations relative to other agents to sustain communication links. We are interested in developing robotic teams that can operate autonomously in urban and/or hazardous areas and perform tasks such as surveillance, target search and identification, and reconnaissance, all while maintaining team connectivity. These tasks are relevant in applications such as urban search and rescue, and environmental monitoring for homeland security, to name a few. We note that while the maintenance of network connectivity is required for useful situational awareness and system responsiveness, often the very environments we wish to operate in make it extremely challenging, especially when the mobile robots consist of small, lightweight ground vehicles that operate very close to the ground.

In general, it is difficult to predict radio transmission properties a priori due to their sensitivity to a variety of factors including transmission power, terrain characteristics, and interference from other sources. Most existing propagation models assume transmission distances of approximately 100–200 m with antennas placed high above the ground (Neskovic et al., 2000) and are not applicable to small, lightweight mo-

bile nodes operating with low transmission power. This is because at these small scales, the signal propagation mechanism is often dominated by the effects of reflection and scattering making modeling especially challenging in unexplored and unstructured environments. In this work, we present techniques for ground vehicles connected via a wireless network to collaboratively perform surveillance tasks while providing situational awareness to an operator. We first show how nominal models of an urban environment can be used to generate strategies for exploration and present the construction of a radio signal strength map that can be used to plan multi-robot tasks and also serve as useful perceptual information. Additionally, we present reactive controllers for communication link maintenance. These controllers can be used in conjunction with information gleaned from our radio signal strength maps to enable our robots to adapt to changes in actual signal strength or estimated available bandwidth. Inherent to this, we assume that the environments under consideration are traversable for our mobile ground vehicles. The focus of this paper is on maintaining group network connectivity while executing individual behavioral tasks. We describe techniques to aid in planning robotic missions subject to connectivity constraints, and a reactive technology layer that maintains those constraints that may be composed with other controllers. Both of these approaches are presented along with experimental results.

2. BACKGROUND

We consider the problem of a team of robots operating in an urban, potentially hazardous, environment for tasks such as reconnaissance and perimeter surveillance, where maintaining team connectivity is essential for situational awareness. In these tasks, robots must have the ability to align themselves along the boundaries of complex shapes in two dimensions while ensuring the successful transmission of critical data. Importantly, navigation based solely on the geometry of the environment will not always guarantee a connected communication network. In these situations, a rough model of the radio signal propagation encoded in a radio connectivity map, i.e., a map that gives the signal strength measurements from one position in the workspace to any other position, becomes extremely helpful in the planning phase (Hsieh, Kumar & Taylor, 2004). Furthermore, since




Figure 1. Signal strength measurements over various distances obtained in an environment similar to an urban park. Antennas were positioned 40 cm above the ground, and the signal strength (y axis) is normalized to a scale of 0–100.

real-world environments are often very complex and dynamic, it is important for robots to also have the ability to respond to real-time changes in link quality to ensure network connectivity (Hsieh, Cowley, Kumar & Taylor, 2006). In Figure 1, we show actual signal strength measurements obtained using two nodes for different separation distances obtained in an environment representative of an urban park. Although there is a strong correlation between signal strength and distance, there is also a lot of variability due to the various factors mentioned earlier (Neskovic et al., 2000). These kinds of information cannot always be accurately inferred from a radio connectivity map. Thus, successful mission execution will require both proper deliberative planning and suitably designed reactive behaviors to facilitate the operation of the team with little to no direct human supervision.

One of the earliest works studying the effects of communication on multi-agent systems is Dudek, Jenkin, Milios & Wilkes (1995) where the effects of two-way, one-way, and completely implicit communication and sensing in a leader follower task were considered. This along with other works Winfield (2000), Arkin & Diaz (2002) often assumed constant communication ranges and/or relied on line-of-sight maintenance for communication. Other examples in-

clude Pereira, Das, Kumar & Campos (2003) and Sweeney, Brunette, Yang & Grupen (2002), where decentralized controllers for concurrently moving toward goal destinations while maintaining relative distance and line-of-sight constraints were respectively presented; and Anderson, Simmons & Goldberg (2003) discussed the formation of communication relays between any pair of robots using line of sight. Although coordination strategies that rely on line-of-sight maintenance may significantly improve each agent's ability to communicate, it has been shown through simulation by Thibodeau, Fagg & Levine (2004) that line-of-sight maintenance strategies are often not necessary and may potentially be too restrictive. In their work, the authors were able to show through simulation that coordination strategies based on line-of-sight maintenance for cooperative mapping are overall less efficient than strategies based on inter-agent wireless signal strengths.

Recent works that consider coordination strategies based on inter-agent signal strength include Wagner & Arkin (2004), where the combination of planning and reactive behaviors for communication link maintenance in a multi-robot team conducting reconnaissance was used. In this work, robots were tasked to go to different goal positions while main-

taining communication links with the base station and/or a communication relay robot. In the event the robots sense a drop in the quality of their communication link(s), a contingency plan, i.e., a plan used to reestablish network connectivity, were triggered. In this case, the contingency plan re-tasked the robots to go to a location within the workspace selected a priori. Simulation results were presented for teams of two to four robots. In general, goal positions are determined and planned based on all available information including radio transmission properties. However, most reasonably ambitious missions run the risk of encountering situations that are not reflected in planning. In the case of radio signal propagation in urban environments, one could rely on simulation validation of a plan, however this would require one to be extremely conservative in mission planning due to the difficulty in accurately predicting radio transmission characteristics.

Navigation based on perceived wireless signal strength between robots for exploration was presented in Sweeney, Grupen & Shenoy (2004). Here a null-space projection approach was used to navigate each robot towards its goal while maintaining point-to-point communication links. This work included simulation results for a team of four planar robots. In Powers & Balch (2004), individual agents made control decisions based on their actual and predicted signal strength measurements while moving towards a goal. Simulation results for teams of one to four robots with and without the controller were presented. Although coordination strategies based on inter-agent signal strength can significantly improve overall performance, they do not account for the effects of team size on overall network performance. As team size increases, the issue of bandwidth becomes more important since an acceptable level of signal strength no longer guarantees a robot's ability to transmit critical data.

Figure 2(a) shows the number of transactions² per interval of time between four different robots, positioned at four distinct fixed locations, and a fifth stationary robot, which we will call the Base. Initially, one robot is transmitting at the maximum data rate supported by the network. As the second, third and fourth robots successively begin their transmissions to the Base, we see not only a drop in the bandwidth available to each robot, but also a drop in total network throughput as significant network resources are

spent coping with low-level packet collisions, retries and contention resolution. Situations such as this often occur in practice because a robot's sensing bandwidth typically exceeds network bandwidth. It is important to note that during this time, the wireless signal strength measurements between the individual robots and the Base are virtually constant, as shown in Figure 2(b), since inter-robot distances were kept constant.

Additional works considering the impact of communication include Pimentel & Campos (2003), where a distributed optimization approach for cooperative motion planning while maintaining network connectivity was proposed. Motion control algorithms for achieving biconnectivity in ad-hoc mobile networks were proposed in Basu & Redi (2004b), while deployment strategies for achieving k connectivity in sensor networks were considered in Bredin, Demaine, Hajiaghayi & Rus (2005). Basu & Redi (2004a) considered flocking strategies for placement of unmanned aerial vehicles to maintain connectivity of ground networks. The effects of time-varying communication links on control performance of a mobile sensor node over a wireless network and in distributed sensing and target tracking were considered in Mostofi & Murray (2004) and Mostofi, Chung, Murray & Burdick (2005), respectively. The use of wireless communication for localization was discussed in Howard et al. (2003) and for localization and navigation in Corke, Peterson & Rus (2003). Deployment strategies for a mobile sensor network to control sensor node density were considered in Zhang & Sukhatme (2005). An exploration methodology was presented in Hsieh et al. (2004) that enables the deployment of a multi-robot team to map the radio propagation characteristics of an urban environment.

Most prior works in the area of communication link maintenance leave the burden of performance specification to fixed metrics, typically based on the distance between nodes or on simulated signal strength. However, as mentioned earlier, radio signal propagation depends on a variety of factors that are often difficult to capture in simulation alone. Rather than rely on simulation, our approach entails the use of radio connectivity maps for planning as well as low level reactive controllers that respond to changes in actual signal strength or verified network bandwidth. The goal is to develop strategies that exploit information gathered during an initial exploration phase coupled with well-designed reactive behaviors that remain minimally disruptive to any high level

²This metric is defined more precisely later in Sec. 4.2.

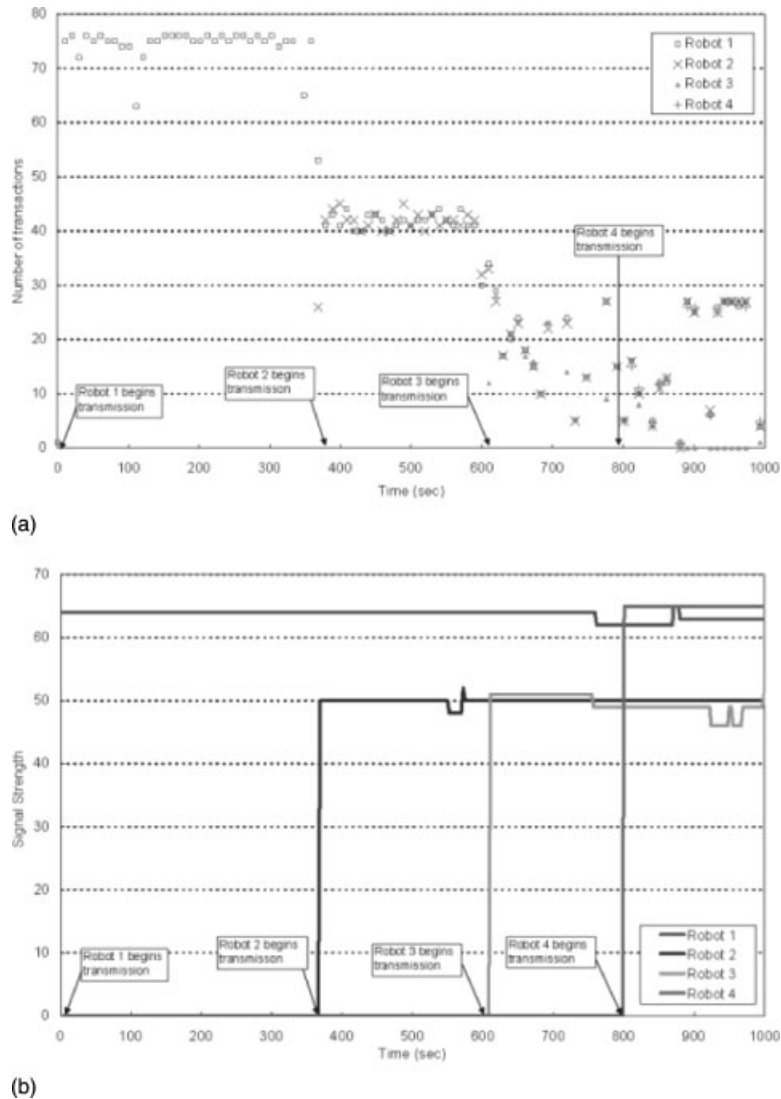


Figure 2. (a) Number of transactions per interval of time between four stationary robots, at four distinct locations, and the Base. The number of successful transactions between each robot and the Base drops as the number of transmitting robots increases over time. (b) Signal strength measurements from the robots to the Base for the same period of time.

deliberative plans in order to maximize the team’s ability to provide effective situational awareness to a base station. In essence, our strategies are based on metrics that do not rely on assumptions that may not be transferable or realistic in the physical workspace that the team is operating within.

In this work, we present an actual radio connectivity map for an urban environment acquired using the methodology described in Hsieh et al. (2004). Additionally, we present low-level reactive controllers

that can be used to constrain the motion of individual agents based on two link quality measures: signal strength and perceived network bandwidth. We present two sets of experimental results using these controllers in outdoor environments under different network interconnection topologies. In the first set of experiments, the radio connectivity map was used to determine a deployment strategy for a reconnaissance task. In the second set of experiments, we deployed our multi-robot team to execute a perimeter

surveillance task. The reactive controllers are designed to be minimally disruptive to the overall deliberative plan, and provide situational awareness to a base station including notification regarding potential failure points in the communication network.

3. MULTI-ROBOT RADIO MAPPING

In this section, we describe the methodology used to deploy our multi-robot team to obtain a radio connectivity map in an urban environment. We consider the special case of a team of three homogeneous planar robots and present our experimental results.

3.1. Modeling

For any given environment, denote the configuration space as $C \subset \mathbb{R}^2$ and the obstacle free portion of C as C_f , also referred to as the free space. Given any two positions $q_i, q_j \in C_f$, the *radio connectivity map* is a function $\varphi: (q_i, q_j) \rightarrow \mathbb{R}$ that returns the average radio signal strength between the two positions given by q_i and q_j . Since it is extremely difficult to obtain a connectivity map for all pairs of positions in C_f , our goal is to construct a map for pairs of locations in the set $Q = \{q_1, \dots, q_{n_1}\}$ such that Q is a subset of C_f .

Assume a convex cell decomposition can be performed on any given C_f such that each location in the set Q is located within a cell. This does not necessarily mean the signal strength will be the same for all pairs of positions in any two cells. However, since each cell is convex, it is possible to predict the signal strength for any two points given the line-of-sight property associated with points in a convex set and prior knowledge of the variation of radio signal transmission characteristics with distance. Our objective is to develop a methodology for the construction of the radio connectivity map, thus we will assume the decomposition is given instead of solving the problem of determining the appropriate cell decomposition.

Additionally, assume a connected roadmap can be constructed from the given cell decomposition of C_f and is represented as an undirected graph $G_1 = (V_1, E_1)$ where each cell is associated with a node in V_1 and every edge in the set E_1 represents the existence of a feasible path between neighboring cells. Given

$$V_1 = \{v_1^1, \dots, v_1^{n_1}\} \quad \text{and} \quad E_1 = \{e_1^1, \dots, e_1^{m_1}\},$$

the total number of nodes and edges in G_1 are denoted as n_1 and m_1 , respectively. Thus, for every $q_i \in Q$, there is a corresponding $v_1^i \in V_1$. The adjacency matrix for G_1 , denoted by A_1 , is defined as

$$A_1 = [a_{ij}] = \begin{cases} 1 & \text{if path exists between } v_1^i \text{ and } v_1^j \\ 0 & \text{otherwise} \end{cases}$$

The graph G_1 is called the *roadmap graph*. Since our team consists of homogeneous robots, the same G_1 applies to every member of the team.

Define the *radiomap graph*, $R = (V_1, L_1)$, which is used to encode signal strength information one would like to gather. The edge set L_1 represents signal strength measurements that must be obtained for pairs of nodes and is selected a priori based on the task objectives, the physical environment, prior knowledge of radio signal transmission characteristics, and may include all possible edges in G_1 . The adjacency matrix A_R for the radiomap graph, R , is given by

$$A_R = [a_{R,ij}] = \begin{cases} 1 & \text{if signal strength between } v_1^i \\ & \text{and } v_1^j \text{ is to be measured} \\ 0 & \text{otherwise} \end{cases}$$

Given the roadmap and radiomap graphs, G_1 and R , a *multi-robot exploration graph*, $G_k = (V_k, E_k)$, where k denotes the number of robots in the team, is constructed such that determining an optimal strategy to measure the edges in L_1 is equivalent to solving for the shortest path on the graph G_k . The methodology for the three robot case is presented in the following section.

3.2. Methodology

Given the roadmap, $G_1 = (V_1, E_1)$, and k robots, a *configuration* on G_1 is an assignment of the k robots to k nodes on the graph. Figure 4(b) shows some possible configurations of three robots on the roadmap graph G_1 , shown in Figure 4(a). Here solid vertices denote the locations of the robots. Since the graph G_1 is connected, a path always exists for k robots to move from one configuration to another. For certain configurations of k robots on G_1 , the complete graph, i.e., fully connected graph, consisting of the locations




Figure 3. (a) A typical surveillance picture from our fixed wing UAV taken at an altitude of 150 m. The area shown is approximately $90\text{ m} \times 120\text{ m}$. (b) A manual cell decomposition of the free configuration space for the site shown in Figure 3(a).

of the robots as vertices, contains some of the edges in L_1 . Figure 5(b) shows some three robot configurations on G_1 that can measure edges in L_1 – the edge set of the radiomap graph shown in Figure 5(a). Note the set of edges for the complete graph generated by the robots, shown as solid vertices and edges, consists of some edges in L_1 . And as such, the fully connected graph generated by the k robots may have more edges than the subgraph of R_1 consisting of the same vertices. Thus, a plan or exploration strategy to measure all edges in the set L_1 can be viewed as a sequence of robot configurations such that every edge in L_1 is measured by at least one of these configurations.

In general, given the graphs G_1 and R and k ro-

bots, the multi-robot exploration graph, G_k , is constructed such that every node in V_k denotes a k -robot configuration on G_1 that measures a subset of L_1 . An edge, $e_k^{ij} \in E_k$, exists between any two nodes $v_k^i, v_k^j \in V_k$ if the configuration associated with v_k^i is reachable from the configuration associated with v_k^j . Since G_1 is always connected, k robots can always move from one configuration to another, therefore, G_k is always a complete graph, i.e., a graph where every node is adjacent to every other node. Every edge in E_k is then assigned a minimum cost that represents the total number of moves required to move the robots from one configuration to another. For the configuration given by the nodes $\{2,3,4\}$ as shown in




Figure 4. (a) Roadmap graph, G_1 . The solid edges denote feasible paths between neighboring cells associated with each node. (b) Three different configurations three robots can take on the graph G_1 . The solid vertices denote the locations of the robots.




Figure 5. (a) Radiomap graph, R , for G_1 shown in Figure 4(a). The dashed edges denote links for which signal strength information must be obtained. (b) Three sample configurations of three robots on G_1 that can measure at least one of the edges in R . The solid vertices and solid edges denote the graph generated by the locations of the robots.

Figure 5, the minimum cost to move to the configuration given by nodes $\{1,2,3\}$ is 2. The optimal plan/exploration strategy simply consists of a sequence of configurations, such that moving through all configurations in the sequence results in covering all edges in L_1 while minimizing the number of total moves. And as such, since the multi-robot exploration graph, G_k , encodes all necessary information needed to determine an exploration strategy for the k -robot team to obtain, finding an optimal plan is equivalent to solving for a minimum cost path on G_k that covers all the edges of L_1 .

For example, consider the case where $k=3$ with the roadmap and radiomap graphs, G_1 and R shown in Figures 6(a) and 6(b). To determine the set of nodes in V_3 of G_3 , consider all three robot configurations on the graph G_1 that contain at least one edge in L_1 . Some of these configurations are shown in Fig-

ure 6(c) where the configurations given by nodes $\{1,5,6\}$, $\{2,3,6\}$, $\{3,4,5\}$, and $\{3,4,6\}$ would correspond to nodes $1'$, $2'$, $3'$, and $4'$ on G_3 respectively. Figure 6(d) is a subgraph of G_3 with the nodes associated with the configurations shown in Figure 6(c) as its vertices.

Shortest path computation between every node in G_1 is required to determine the weight of every edge in E_3 . For the three robot case, every edge in the set L_1 may potentially be associated with more than one node in V_3 . Thus, the optimal plan for the three robot case would result in a path that contains a subset of the nodes in V_3 . For this example, an optimal plan starting at the configuration given by node $1'$ is the path $\{1',2',4'\}$ with a total cost of 4 and does not contain node $3'$. In general, given a starting node on G_3 , a greedy algorithm is used to compute a path on G_3 such that traversal of each




Figure 6. (a) A roadmap graph, G_1 . (b) A radiomap graph, R . (c) Graph R overlaid with some G_3 nodes, denoted by \otimes . Nodes $1'$, $2'$, $3'$ and $4'$ refer to the configurations given by nodes $\{1,5,6\}$, $\{2,3,6\}$, $\{3,4,5\}$, $\{3,4,6\}$ respectively. (d) Subgraph of the radio exploration graph, G_3 , for the roadmap and radiomap graphs shown in 6(a) and 6(b).

node on the path increases the number of measured edges in L_1 . Thus, at any configuration, the next configuration is chosen as the one that increases the number of edges measured in L_1 and requires the least amount of moves to reach. During execution, robots are allowed to cross each other as they switch from one configuration to another since inter-robot collisions are prevented via each robot's local obstacle avoidance routines. We refer the interested reader to Hsieh et al. (2004) for details on the algorithms used to obtain the vertex set V_3 and the cost and adjacency matrices for G_3 .

3.3. Experimental Setup and Results

The objective of this experiment was to deploy a team of three robots to obtain a radio signal strength map for the Military Operations on Urban Terrain (MOUT) training site, located in Ft. Benning, Georgia, where radio signal strength data are important for operations such as surveillance, reconnaissance, and search and rescue. Figure 3(a) is an aerial view of the MOUT site. More information on the experiments conducted at the MOUT site can be found in Chaimowicz, Grocholsky, Keller, Kumar & Taylor (2004), Grocholsky, Bayraktar, Kumar, Taylor & Pappas (2004) and Grocholsky, Keller, Kumar, & Pappas (2006).

3.3.1. Hardware

Our multi-robot team consists of five unmanned ground vehicles (UGVs) modified from commercially available, radio controlled scale model trucks. Each vehicle's chassis is approximately 480 mm long and 350 mm high. Mounted in the center of the chassis is a Pentium III laptop computer. Each vehicle contains a specially designed Universal Serial Bus device, which controls drive motors, odometry, steering servos and a camera pan mount with input from the PC. A global positioning system (GPS) receiver is mounted on the top of an antenna tower, and an inertial measurement unit is mounted between the rear wheels. A forward-looking stereo camera pair is mounted on a pan mount which can pivot 180° to look left and right. The junction box (JBox) and an omnidirectional antenna handle wireless communication with the antenna placed approximately 40 cm off the ground. The JBox, jointly developed by the Space and Naval Warfare Systems Center and BBN Technologies, consists of a small



Figure 7. Our multi-robot testbed.

embedded computer with 802.11b wireless Ethernet and is used to handle multi-hop routing in an ad-hoc wireless network and provide signal strength measurements for all nodes on the network. Last, our robots are designed to travel at a fixed speed of 1 m/s. A picture of our multi-robot team is shown in Figure 7.

3.3.2. Results

Since the objective was to obtain the desired radio signal strength map from a given decomposition of the free space rather than to determine the appropriate decomposition, we assumed a cell decomposition of the free space shown in Figure 3(b), which was obtained by hand. The corresponding roadmap and radiomap graphs for this particular decomposition are shown in Figure 8. The edges for the radiomap graph were selected such that they cover the main North-South and East-West roadways on the MOUT site where other planar multi-robot experiments were often conducted (Chaimowicz et al., 2004, Grocholsky et al., 2004, Grocholsky et al., 2006).

Following the procedure outlined in the previous section, we obtained the three robot exploration graph, G_3 , which contained 188 nodes³. Rather than weight the edges of G_3 with the total moves to go from one configuration to another, we weighted each edge by the total Euclidean distance the team would

³The roadmap and radiomap graphs each contained 13 nodes.

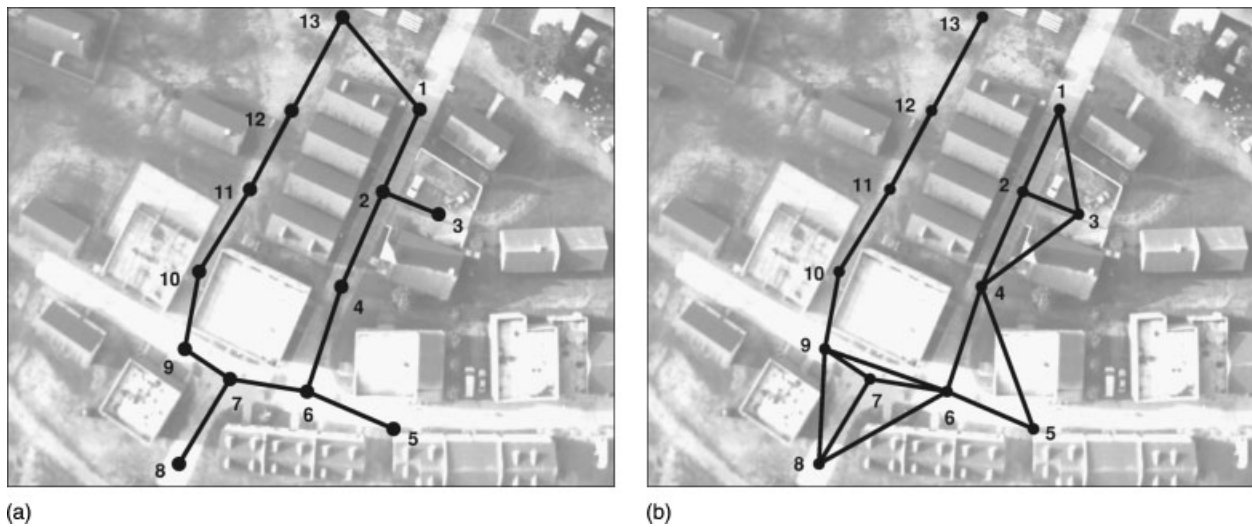


Figure 8. (a) Roadmap graph used for the site shown in Figure 3(a). (b) Radiomap graph for the site shown in Figure 3(a).

have to travel to get from one configuration to another to more accurately capture the cost associated with each configuration change. The starting configuration for the three-robot team was selected to be at nodes $\{1,2,3\}$ in Figure 8. The exploration strategy that was obtained would deploy the robot team in the following sequence: $\{1,2,3\}$, $\{1,2,4\}$, $\{2,3,4\}$, $\{4,5,6\}$, $\{6,7,9\}$, $\{6,8,9\}$, $\{7,8,9\}$, $\{9,10,11\}$, $\{10,11,12\}$, and $\{11,12,13\}$.

Once this strategy was obtained, a centrally located waypoint was selected for each cell and each robot was then assigned a set of waypoints to be traversed based on the exploration strategy. Using GPS, the robots navigated to each of their assigned locations. Upon arrival at each waypoint, the robots would synchronize and measure their signal strengths to other team members. If the synchronization failed, each robot would move on towards the next waypoint on its assigned list. To enable each robot to return to its starting position after completion, we assigned each robot its starting position as its last waypoint. The waypoints and cell decompositions were manually chosen to minimize the number of failed synchronizations during the execution based on prior knowledge of signal strength variation with distance (see Figure 1). In addition, each robot was continuously logging both signal strength and position data such that in the event of a failed

synchronization the information could be retrieved. There were no synchronization failures during the experiment.

Figure 9 shows the radio signal strength map constructed for the MOUT site. The numbers by each edge are the averaged normalized signal strength measurements obtained by the robots lo-




Figure 9. Radio signal strength map obtained for the MOUT site. The number on each edge is the average normalized signal strength for each position pair.

cated at each pair of positions. On average the GPS errors ranged from 2–3 m to as much as 5 m. However, the robots generally were able to stay within the boundaries of the convex cells.

4. REACTIVE CONTROLLERS FOR COMMUNICATION LINK MAINTENANCE

In this section, we consider the problem of guiding a group of N robots to a set of goals or simply a desired boundary (curve) while maintaining point-to-point communication links. We discuss the synthesis of reactive controllers that allow each robot to respond to changes in its perceived communication link quality with respect to other team members within its sensing range. We present experimental results with our multi-robot testbed in two separate outdoor environments. In the first experiment, reactive controllers were used in conjunction with the radio connectivity map shown in Figure 9 to determine a deployment strategy for a reconnaissance task on the MOUT site using a four-robot team. In the remaining experiments, reactive controllers capable of responding to changes in signal strength or data throughput were used to maintain point-to-point communication links in a perimeter surveillance task conducted in a separate outdoor environment.

4.1. Controllers

In general, for a team of N robots each with kinematics $\dot{q}_i = u_i$, where i denotes the i th robot, $q_i = (x_i, y_i)^T$ denotes the i th robot's position, and u_i denotes the i th robot's control input, consider the following controller

$$u_i = -k\nabla_i\phi_i(q_i) - \sum_{j \in \Gamma_i} \nabla_i g_{ij}(q_i, q_j) \quad (1)$$

where k is a positive constant scalar, ϕ is some artificial potential function, and Γ_i denotes the set of neighbors for agent i . The first term of the control law (1) guides each robot to its goal position and the second term maintains the constraints that need to be satisfied between robot i and a pre-specified set of neighbors. The functions, $g_{ij}: \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \mathbb{R}$, are artificial potential functions used to model inter-robot constraints. We are interested in maintaining radio connectivity, thus g_{ij} should model the radio propa-

gation characteristics among agents such that $-\nabla_i g_{ij}$ results in a policy that increases the quality of the communication link between robots i and j , where ∇_i denotes the gradient with respect to the coordinates of the i th robot.

As described in Sec. 3.3.1, our multi-robot team consists of five modified radio controlled scaled model trucks and therefore cannot be described by the simple kinematic model $\dot{q}_i = u_i$. However, taking inspiration from Eq. (1), our reactive controller is composed of two components: one for navigation to specific goal positions and one that modifies the navigation based on variations in a robot's link quality. These controller elements correspond directly to the first and second terms of Eq. (1). For each goal position, a reference heading, similar to the descent direction of a potential field controller, is computed by the navigation component. Based on this "descent direction," a look-ahead waypoint is generated based on the vehicle's speed and position. Then a simple proportional, integral, and derivative (PID) controller is used to steer the robot towards the look-ahead waypoint. The process is repeated until the goal position is reached.

To maintain constraints, each robot continuously monitors the quality of the communication link(s) to its specified set of neighbors. In our experiments, our robots have the capability to continuously monitor either the signal strength to its neighbors, as in Figure 2(b), or the number of successful transactions⁴ per unit time as in Figure 2(a). When the link quality drops below a minimum threshold, the constraint maintenance component will either stop the robot or move it closer to its neighbor until the quality returns to an acceptable level. When stopped, a robot can wait for a specified time interval before attempting to move towards its goal again. If the stopped robot perceives an increase in its link quality above the acceptable level, it can once again attempt to move towards its goal. Such recovery measures may be used to lessen the times a robot is caught in a spatio-temporal dip in link quality due to dynamic changes in the environment. In other words, local minima situations, in which a robot stopped before reaching its goal, caused by some temporary interference in the environment. Additionally, these measures also ensure that a robot is constantly minimizing its distance to the goal as long as all constraints are satisfied. We have pur-

⁴This metric is defined more precisely later in Sec. 4.2.

posefully incorporated these recovery measures in some of our experiments to highlight the reactive nature of our controller. The algorithm is summarized in Algorithm 1. In this algorithm, the “Recover” behavior drives the robot closer to the neighbor it is attempting to maintain its link quality with. While we chose a binary signal derived from the link quality to control the robots since they are designed to travel at a constant fixed speed, there is an effective deadband in the controller given by the “Minimum” and “Acceptable” quality thresholds in Algorithm 1.

Algorithm 1 Link Quality Constrained Navigation

```

if LinkQuality < Minimum then
  Recover;
  recover_flag=true;
end if
if Minimum < LinkQuality ≤ Acceptable then
  if recover_flag then
    Stop and wait;
    recover_flag=false;
    stopped=true;
    waitTime=current time;
  end if
  if (current time – waitTime) > MaxWaitTime then
    Retry going to goal;
    stopped=false;
  end if
end if
if LinkQuality > Acceptable then
  Go to goal;
end if

```

By design, our reactive controller favors constraint satisfaction above reaching the goal position. This is to ensure that the quality of the communication link is always maintained, thus providing the human operator at a remote base station the ability to monitor the various communication links in the network. At the base station, display panels with each robot’s imagery data and/or the signal strength measured by the various team members can be displayed. In the event the Base did not receive new data from a particular robot, or should some particu-

lar link exhibit low signal strength, an indication of the detected failure is relayed to the human operator. This capability, possible because the team always remains connected, enables the operator to decide whether or not to deploy additional robots, or to reorganize the team.

4.2. Link Quality Estimation

Signal strength between a sender and a receiver is a function of the transmission power, antenna gains, and signal attenuation. Our robots are equipped with JBoxes that, among other things, provide signal strength measurements to every node on the network. We refer the interested reader to BBN Technologies (2003) and Redi, Watson & Shurbanov (2002) for operational details on the JBox.

In contrast, it is difficult for an individual robot to estimate the available bandwidth at any given point in time since bandwidth is a function of the number of nodes, the amount of traffic on the network, as well as the signal strength. In multi-robot applications, it is often relevant to talk about bandwidth in terms of units of application level data that can be transmitted, therefore we define a *successful transaction* to be the transmission of one unit of application level data sent by a sender with an acknowledgment of receipt sent by the receiver. A robot’s conservative estimate of available bandwidth is determined based on the number of successful transactions it achieves over some interval of time.

For our experimental setup, we set one unit of application level data equal to a JPEG (Joint Photographics Expert Group) compressed image of approximately 10 KB in size. We define a successful transaction to be the transmission of such an image by the sender followed by the receipt of acknowledgment sent by the receiver. Then, based on the desired transaction rate, i.e., number of successful transactions per time interval, the robot, i.e., the sender, will periodically evaluate its connection with the receiver. The available bandwidth controller comprises two stages of response: network usage throttling, and robot repositioning corresponding to the “Recover” behavior in Algorithm 1. A robot that is streaming data over the network is capable of detecting when a network connection is not keeping up with the load being put on it, i.e., when some messages are dropped due to a full send buffer. When this type of application-level packet loss is detected, the system will automatically throttle com-

munication over this particular network link until a prescribed threshold is hit. This threshold is determined based on mission specifications; we typically specified video at a rate of three frames per second as a requirement because this was a realistic target for the 802.11b hardware in use, and provided sufficient video coverage to convey situational awareness. When the network throttling mechanism bumps into the lower threshold, the controller subsumes position control to move the robot closer to the peer it is attempting to communicate with. This action serves to increase signal strength, which reduces the number of low-level transmission retries caused by noise or attenuation.

All throughput estimation in this framework is conservative; we do not attempt to measure maximum data rates available on the network, but rather we verify that some prescribed minimum data throughput rate were available. This approach minimizes the amount of network traffic related solely to throughput measurement, and instead leverages throughput assessment on normal data traffic when such traffic satisfies the constraint. When normal traffic is not of sufficient volume to verify that the minimum available bandwidth constraint is met, a connection monitor will periodically verify available throughput. This latter mechanism, which simply sends data at a rate that verifies constraint satisfaction for a short period of time, allows us to deploy robots that do not maintain consistent data flow

back to the base station, e.g., robots that only send event data, yet still be confident that the available throughput would likely be available if it should be needed.

4.3. Experimental Results

In this section we present three sets of experimental results. The first experiment was conducted at the MOUT site, shown in Figure 3(a). The experiment was modeled after a reconnaissance application where the objective was to deploy a team of four robots to obtain surveillance imagery at a designated location out of line-of-sight and single-hop radio communication range with the base station. In this experiment we used information gleaned from our radio connectivity map, shown in Figure 9, to determine the deployment strategy, and coupled this with low level reactive controllers to enable the team to respond to unforeseen changes in signal strength. The second and third set of experiments were conducted at one of the University of Pennsylvania's soccer fields. A satellite image and a schematic of its surroundings are shown in Figure 10. These experiments were based on a perimeter surveillance application in which each robot was required to send imagery data back to a base station. Here we focused on the individual robot's capability to respond to changes in signal strength or perceived available bandwidth rather than the pairing of high level

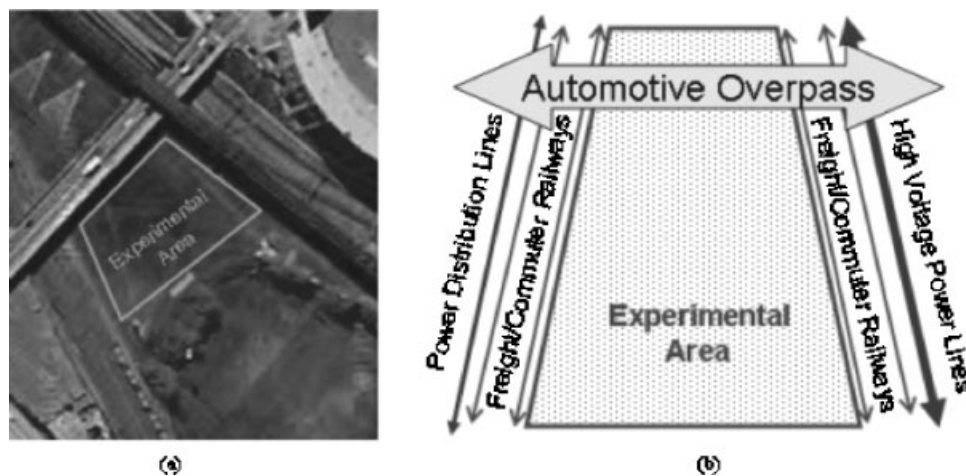


Figure 10. (a) A satellite image of the soccer field and its surroundings. (b) Detailed schematic of the experimental area and its surroundings.

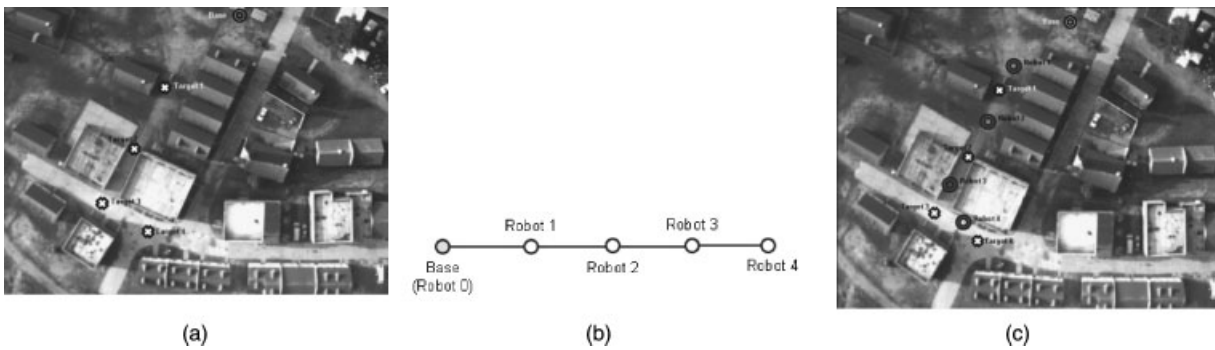


Figure 11. (a) An overhead view of the MOUT site taken from a fixed wing UAV at an altitude of 150 m. The area shown is approximately 90 m \times 120 m. The location of the Base is denoted by \circ and the target locations for the team are denoted by \times . (b) The underlying communication graph for the reconnaissance application. (c) The final positions attained by each robot and their designated target locations denoted by \circ and \times respectively.

planning with reactive controllers to ensure communication link maintenance. Once again, our multi-robot team for this set of experiments consisted of the five UGVs described in Sec. 3.3.1, one of which was chosen to be the base station (Base).

4.3.1. Reconnaissance at a MOUT Site

In this experiment, we deployed a team of four robots to obtain surveillance imagery at a designated location out of direct radio communication range. The objective was to deploy four robots to four separate goal positions such that the team formed a linear multi-hop network. Each robot's goal position was determined based on signal strength information given by our radio connectivity map, shown in Figure 9. The goal positions were chosen to be slightly beyond nodes 9, 10, 11, and 12 in the road-map graph shown in Figure 8(a) with respect to the Base location [see Figure 11(a)]. To account for unforeseen variations in signal strength during mission execution, we tasked the i th robot to monitor its signal strength to the $(i-1)$ th robot, and stop when the signal strength dropped below the acceptable threshold, i.e., $MaxWaitTime = \infty$ in Algorithm 1. The Base was considered the 0th robot. The radio connectivity map was then used to determine the minimum acceptable signal strength for each of the robots. The minimum acceptable signal strength was set to 55, 65, 65, and 60 for robots 1, 2, 3, and 4, respectively. Once the team stopped, the human operator at the Base requested images from the robot that was closest to the location of interest, i.e., the

4th robot. In this experiment, data were only transmitted between the 4th robot and the Base via the linear multi-hop network.

As shown in Figure 11(c), although the targeted locations were chosen to ensure team connectivity, these locations were not reached since each robot was also responding to changes in the real-time signal strength measurements to its designated neighbor, ensuring its signal strength was above the required threshold. Figure 12 shows the distance of each robot to its respective goal position over time. These measurements were obtained using each robot's raw GPS data with approximately 2–3 meters in accuracy and thus, the slight variations are due to GPS errors rather than robot movements. The final signal strength measurements for robots 1, 2, 3 and 4 were 52, 63, 64, and 57, respectively. While there exists some unpredictability in terms of each robot's ultimate destination, operations at the limits of hardware capabilities, such as demonstrated here, fall outside typical confidence intervals of reliable simulation. If dynamic responses are not allowed, then mission specification must be performed with such a level of conservatism as to severely limit system capabilities.

4.3.2. Perimeter Surveillance Application

These next experiments were based on a perimeter surveillance application where robots would navigate to positions on the desired perimeter and send imagery data back to a base station. In our experiments, we tasked a team of four robots to go to four




Figure 12. Robots' distances to their respective goals over time. The data are obtained using each robot's raw GPS data with approximately 2–3 m in accuracy.

separate goal positions where the positions were chosen to represent distinct locations on the perimeter of interest. In addition, each robot was also tasked to monitor its signal strength or estimated available bandwidth to the Base while continuously sending imagery data to it. Thus, every robot was sending approximately 10 KB JPEG compressed images back to the Base, as compared to the previous experiment, where only one robot was sending data to the Base.

At the Base, a display panel with each robot's imagery data was provided to the operator. In the event the display panel did not receive new data from a particular robot over a specified interval of time, the panel would highlight the display box for that particular robot. The objective of these experiments was to focus on an individual robot's ability to respond to changes in signal strength or perceived available bandwidth. Thus, we only considered single-hop network connections between each robot and the Base. Additionally, to avoid robots being caught in a "local minimum" due to dynamic changes in the environment that may affect signal strength measurements, and to better emphasize the reactive nature of our controllers, we set the `MaxWaitTime` in Algorithm 1 to a finite time.

The first such experiment conducted at this location demonstrates the reactive controller in the presence of dynamic network disturbances. In this experiment, the network disturbance was caused by the addition of a second robot to a network originally used by a single robot transmitting a video stream to the Base. As shown earlier in Figure 2, as new members are introduced into the team, the maximum bandwidth available to each robot drops.

Figure 14 shows how our controller responded

to the addition of the second robot. We first deployed a single robot, Robot 1, to a goal position and required that it continuously send imagery data to the Base while maintaining a minimum transaction rate of seven transactions/s. This can be seen in the top graph of Figure 14 where each marker denotes the number of transactions received in between the time of the current marker and the one before it, approximately 10 s. The `MaxWaitTime` variable in Algorithm 1 was set to 30 s. A schematic of the deployment strategy is shown in Figure 13. At around $t=60$ s, Robot 1 settled to a location about half way to the goal as shown in the bottom graph of Figure 14. Between $t=90$ s and $t=125$ s, the robot attempted to reach its goal a second time and settled to a similar location shown in the bottom two graphs in Figure 14.




Figure 13. Schematic of the experimental setup and underlying communication graph for the results shown in Figure 14. On the left, the dashed line denotes the communication link monitored by the robot. In this experiment, Robot 2 was used to cause a network disturbance by transmitting to the Base.




Figure 14. Top: Number of transactions received by the Base from Robot 1 and Robot 2 over time. Robot 2 began its transmission at around $t=130$ s. Center Top: 1 denotes Robot 1 achieved the target transaction rate and 0 otherwise. Center Bottom: Actual speed achieved by Robot 1. Positive speed denotes the robot is moving towards the goal and negative speed denotes the robot is moving towards the Base. Bottom: Robot 1's distance from the goal.

A second robot, Robot 2, transmitting to the Base was introduced to the network at approximately $t = 130$ s, as shown in the first graph of Figure 14. Immediately, Robot 1 was no longer able to maintain the required transaction rate and therefore began moving back towards the Base in an effort to boost its transaction rate. This can be seen in the last graph of Figure 14 where the robot's distance to its goal starts increasing.

Put simply, a robot was tasked with sending video back to a base station as it monitored a perimeter. In the early stages of the mission, the robot could transmit video at a high rate, but as the robot moved further away from the base station, the rate at which it could transmit video dropped. The robot continued moving until the transmission rate hit a pre-defined threshold. At this point, the robot stabilized its distance from the base station in order to maintain the minimum required transmission rate. Once another network user was added, the first robot had to move yet closer to the base station. This response is made because the transmission rate for a single robot is a function both of total network usage and signal quality. Once the robots are transmitting at the lowest acceptable rate, the other variable they

can individually have an effect on is radio signal strength, which is related to the distance between transmitters. Utilizing this last control, the first robot will settle on the maximum distance at which it can maintain the required transmission rate given the new networking situation. This repositioning is automatic, and does not require any changes in calibration or thresholds to reflect the new state of the network.

The behavior demonstrated by the two-robot experiment allowed us to successfully deploy a team of robots capable of maximizing network utilization while providing effective situational awareness. Subsequent experiments involved deploying a team of four robots to separate locations from a starting position by the Base. Each robot was tasked to continuously send imagery data from its camera to the Base at a rate above a pre-determined minimum transaction rate. Goal positions for each team member were chosen to provide a wide net of surveillance coverage. This type of goal specification is flexible in that it establishes a vector for the robots to move along, as opposed to specific way points to achieve. Thus




Figure 15. Schematic of the experimental setup and underlying communication graph for the results shown in Figures 16 and 17. Similarly, the dashed lines in the figure on the left denote communication links monitored by each robot.

success is a matter of degree, rather than a binary distinction: we wish to effectively cover as wide an area as possible.

Using the control algorithm described in Sec. 4.1, each robot would move towards its goal until the link quality dropped below the minimum threshold, at which point it would move back towards the Base and stop when the link quality rose back above the chosen minimum level. Once stopped, each robot would wait for a fixed time interval before attempting to go to its goal again. Two

sets of experiments were conducted in which each robot’s controller reacted based on changes in: (i) signal strength measurements and (ii) estimated transaction rate. A schematic of the deployment strategy is shown in Figure 15. Four trials for each experiment were conducted. Since the results are similar for all four robots in all four trials, we have selected one representative result for each set of experiments shown in Figures 16 and 17.

Figure 16 shows signal strength measured by the robot to the Base along with the corresponding com-




Figure 16. Top: Signal strength measured by Robot 1 to the Base. The solid black line denotes the minimum acceptable level. Center: Commanded speed based on the signal strength measurements. Positive speed denotes Robot 1 is moving towards the goal and negative speed denotes it is moving towards the Base. Bottom: Estimated speed achieved by Robot 1 based on the commanded speed. Data for Robots 2, 3, and 4 are similar and thus not shown.

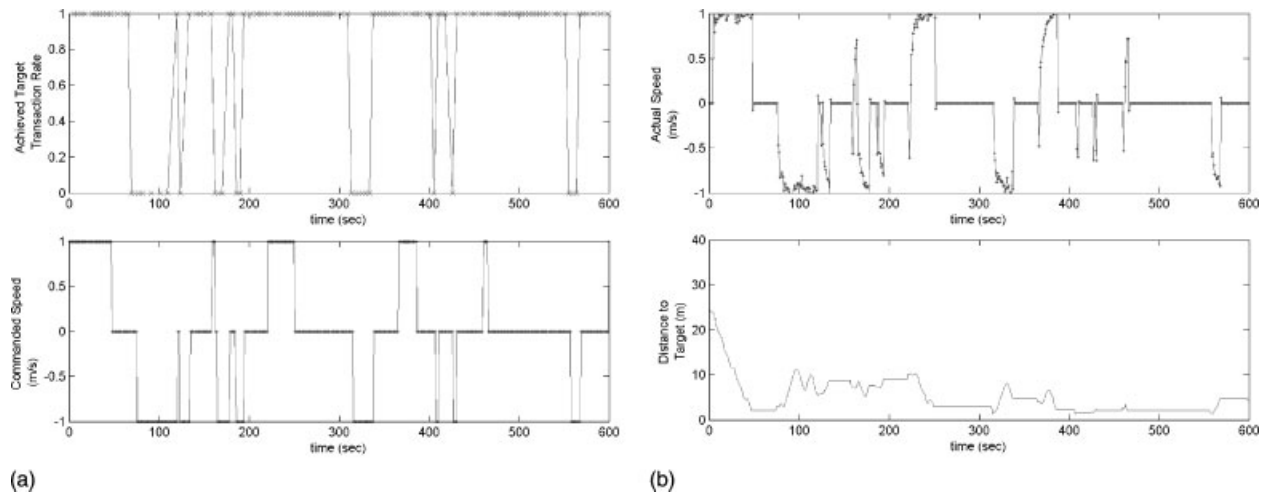


Figure 17. (a) Top: 1 denotes the target transaction rate was achieved and 0 otherwise. Bottom: Commanded speed based on whether the target transaction rate was achieved. Positive speed denotes Robot 1 is moving towards the goal and negative speed denotes it is moving towards the Base. (b) Top: Actual speed achieved by Robot 1 based on the commanded speed. Bottom: Robot 1's distance from the goal. Data for Robots 2, 3, and 4 are similar and thus not shown.

manded speed and actual speed. The MaxWaitTime in Algorithm 1 was set to 60 s. Initially, when the robot was close to the Base, the signal strength measurements were high. As the robot moved toward its goal, we see these measurements drop. The first time the signal strength dropped below the minimum threshold, around $t=45$ s, the robot attempted to move closer to the Base. Once the signal strength rose above the threshold, the robot stopped. Subsequently, the robot made additional attempts to move towards the goal but had to stop and move closer to the Base each time.

Similarly, Figure 17 shows the results for one of the four robots whose controller was reacting to changes in its estimated transaction rate. In these experiments, MaxWaitTime in Algorithm 1 was set to 120 s and the minimum rate was set to three transactions/s. Similar to the results shown in Figure 16, the robot's transaction rate dropped as it moved further away from the Base as shown in Figure 17(a) and the top graph in Figure 17(b). We note that it is possible for the robot to reach its goal location and achieve its target transaction rate. This can be seen in the bottom graph in Figure 17(b) where at approximately $t=50$ s the robot is within 2.5 m of the goal location. Around the same time, we see a change in the robot's speed from positive to zero as shown in the second and third graphs in the same

figure. When the transaction rate dropped, around $t=75$ s, the robot began to move back towards the Base, leaving its goal location.

4.4. Discussion

By design, our reactive controller favors constraint satisfaction above reaching the goal position. This was seen in both the reconnaissance and perimeter surveillance experiments. By designing our controller in this fashion, we ensure the human operator would always be able to get real-time status updates from the team. Should the operator notice certain robots not getting close enough to their goal positions, the operator could deploy additional robots to provide a multi-hop link to the Base or request a reconfiguration of the whole team should the original target connectivity prove unachievable. Similarly, should an intermediate robot fail in a configuration as shown in Figure 11(b), this information would be immediately reflected at the Base. Under these circumstances, the robot furthest away from the Base would surely lose connectivity, however this could be mitigated through the implementation of communication recovery measures, such as return to the last known location with good connectivity,

the dispatch of additional robots or the reconfiguration of the remaining robots still within communication range of the Base.

When considering multi-hop scenarios, it would be important to set the MaxWaitTime variable in Algorithm 1 to infinity. In our perimeter surveillance experiments, we were only considering single-hop communication links. One of the objectives in these experiments was to show the reactive nature of our controller along with minimizing the amount of time a robot was caught in a spatio-temporal “local minimum” due to dynamic interference in the environment. Therefore, in these experiments MaxWaitTime was set to a finite time. In contrast to our MOUT site reconnaissance experiment where the objective was to deploy a multi-hop network, we did not have these recovery responses because we did not want the constant back and forth motion to affect the robots’ ability to send data through the multi-hop network. Had the reactive response been present, it is very likely the constant back and forth motion would affect the team’s ability to reliably relay information to the Base.

In all our experiments, the minimum thresholds were chosen based on a combination of previously collected data and/or specific mission requirements. In the reconnaissance experiment, signal strength thresholds were determined based on information gleaned from a radio connectivity map. On the other hand, when considering perceived network bandwidth, the minimum acceptable threshold was determined based on hardware limitations in conjunction with acceptable transmission rates based on mission requirements specified by the human operator.

Last, our reactive controllers can be easily decentralized based on the methodology proposed in Hsieh & Kumar (2006), and thus scaled to a large number of robots. Rather than specifying specific goal positions for every robot in the team, Hsieh & Kumar (2006) specified a one-dimensional boundary curve for the team. This, however, does not necessarily mean the existing network would be able to handle the increase in traffic brought on by the increase in team size.

5. CONCLUSION

In this work, we have presented a paradigm and algorithms for deploying a mobile robot network with specifications on end-to-end performance. Our ap-

proach entails the automated construction of a radio map for a partially known urban environment which can then be used to deploy a team of robots, and control algorithms that drive the team to designated targets on some desired boundary (curve) while maintaining communication link quality.

There are two main contributions. First, we developed a method for obtaining radio signal strength maps that can be used to plan multi-robot tasks and also serve as useful perceptual information. Second, since a radio signal strength map only serves to create a nominal model, we have shown the importance for individual robots to have the ability to monitor communication links, in particular signal strength measurements as well as available data throughput. This method of link quality control provides scalability in the number of robots added to the network, and an abstraction of the underlying network architecture. Since the robots constantly strive to maximize network usage efficiency, robots may be added or removed from the network without changes to any thresholds or calibration numbers. This type of deployment characteristic is extremely important as robot team sizes scale, as we want teams to take advantage of bandwidth when it is available, and automatically scale back individual usage as available resources are stretched thin. Moreover, we have also shown that channel contention between multiple nodes can have a severe adverse effect on total network throughput. By monitoring successful transactions, we give our robots the ability to throttle their own network usage such that the transmission rates of each robot stabilize to levels that make efficient use of the network.

Additionally, as shown in our perimeter surveillance experiments, in dynamic environments where radio propagation characteristics may exhibit significant changes over time, it is good practice for agents to always attempt to move closer to the goal regardless of where they first come to a stop. The forward movement is the only way to confirm that positions closer to the goal violate communication constraints and to ensure the agents always minimize their distance to the goal location while remaining connected. Ideally, robotic agents should be deployed with the capability of monitoring inter-agent signal strengths as well as data throughput. In general, signal strength is a good indicator of potential connectivity while data throughput can efficiently be used to ensure minimum actual data throughput rates. Combining the two, good signal strength paired with unaccept-

ably low throughput may indicate a need for human attention to the network architecture and the demands being placed on it. As such, the communication medium becomes a useful sensor that can be used to monitor the effectiveness of any given multi-robot deployment.

Reactive navigation controllers such as the one presented provide a reliable foundation on which to build scalable, portable, high-level tasks. The reactive controller acts as a scenario-independent support that allows for the deployment of a robot team to any location, regardless of prior reconnaissance. As shown in our reconnaissance experiment, behaviors built on such a controller inherit respect for network constraints, thereby allowing both flexible goal specification and more deliberative trajectory planning done with environmental models that do not necessarily capture all static and dynamic aspects of an environment's radio propagation characteristics. Since the team always remains connected during mission execution, potential failure points in the communication network, as perceived by individual robots, can be relayed back to the base station to trigger contingency management routines, e.g., deployment of additional robots, or a re-allocation of resources. While the strategies presented in this work assign the highest priority to maintaining the network, applications that may benefit from a relaxation of this constraint provide a direction for future work.

Acknowledgments

The authors would like to thank Luiz Chaimowicz, Dan Gomez-Ibanez, Ben Grocholsky, Selcuk Bayraktar, Jim Keller from the University of Pennsylvania, the Penn Athletics Department, and Jason Redi and Keith Manning from BBN Technologies. We gratefully acknowledge the support of NSF Grant IIS-0427313, ARO Grants DAAD19-02-01-0383, DARPA MARS NBCH1020012, and NSF Grant CCR02-05336.

REFERENCES

Aguero, C., de-las Heras-Quiros, P., Canas, J.M., & Matalan, V. (2003). Pera: Ad-hoc routing protocol for mobile robots. In Proceedings of the 11th International Conference on Advanced Robotics, pp. 694–702, Coimbra, Portugal.

Anderson, S.O., Simmons, R., & Goldberg, D. (2003). Maintaining line of sight communications networks

between planetary rovers. In Proceedings of the 2003 IEEE/RSJ Intl. Conference on Intelligent Robots and Systems, pp. 2266–2272, Las Vegas, NV.

Arkin, R., & Diaz, J. (2002). Line-of-sight constrained exploration for reactive multiagent robotic teams. In AMC 7th International Workshop on Advanced Motion Control, Slovenia.

Basu, P., & Redi, J. (2004a). Coordinated flocking of uavs for improved connectivity of mobile ground nodes. In Proceedings of IEEE MILCOM 2004, Monterey, CA.

Basu, P., & Redi, J. (2004b). Movement control algorithms for realization of fault-tolerant ad-hoc robot networks. *IEEE Network* 18(4), pp. 36–44.

Bredin, J., Demaine, E., Hajiaghayi, M., & Rus, D. (2005). Deploying sensor nets with guaranteed capacity and fault tolerance. In Proceedings of Mobihoc 2005, Urbana-Champaign, Illinois.

Chaimowicz, L., Grocholsky, B., Keller, J., Kumar, V., & Taylor, C. (2004). Experiments in multirobot air-ground coordination. In Proc. IEEE International Conference on Robotics and Automation, New Orleans, LA.

Corke, P., Peterson, R., & Rus, D. (2003). Networked robots: Flying robot navigation using a sensor net. In Proceedings of the Eleventh International Symposium of Robotics Research (ISRR), Springer Tracts on Advanced Robotics (STAR). Springer-Verlag, Berlin.

Cortes, J., Martinez, S., Karatas, T., & Bullo, F. (2002). Coverage control for mobile sensing networks. In Proc. IEEE International Conference on Robotics and Automation (ICRA'02), pp. 1327–1332, Washington, DC.

Dudek, G., Jenkin, M., Milios, E., & Wilkes, D. (1995). Experiments in sensing and communication for robot convoy navigation. In Proceedings of the 2003 IEEE/RSJ Intl. Conference on Intelligent Robots and Systems, pp. 268–273, Pittsburg, PA.

Feistel, A., & Stanczak, S. (2005). Dynamic resource allocation in wireless ad-hoc networks based on QS-CDMA. In The 16th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), Berlin.

Fierro, R., Song, P., Das, A., & Kumar, V. (2002). Cooperative control of robot formations. In Murphey, R., and Parados, P. (Eds.), *Cooperative Control and Optimization: Series on Applied Optimization* (pp. 79–93), Dordrecht, The Netherlands. Kluwer Academic Press.

Grocholsky, B., Bayraktar, S., Kumar, V., Taylor, C., & Pappas, G. (2004). Synergies in feature localization by air-ground robot teams. In 9th International Symposium on Experimental Robotics (ISER'04), Singapore.

Grocholsky, B., Keller, J.F., Kumar, V., & Pappas, G. (2006). Cooperative air and ground surveillance: A scalable approach to the detection and localization of targets by a network of uavs and ugvs. *IEEE Rob. Autom. Mag.*, Vol. 13, pp. 16–25.

Guo, Y., & Parker, L.E. (2002). A distributed and optimal motion planning approach for multiple mobile robots. In Proc. IEEE International Conference on Robotics and Automation (ICRA'02), pp. 2612–2619, Washington, DC.

Howard, A., Siddiqi, S., & Sukhatme, G.S. (2003). An ex-

- perimental study of localization using wireless ethernet. In 4th International Conference on Field and Service Robotics, Lake Yamanaka, Japan.
- Hsieh, M.A., Cowley, A., Kumar, V., & Taylor, C.J. (2006). Towards the deployment of a mobile robot network with end-to-end performance guarantees. In International Conference on Robotics and Automation (ICRA) 2006, Orlando, FL.
- Hsieh, M.A., & Kumar, V. (2006). Pattern generation with multiple robots. In International Conference on Robotics and Automation (ICRA) 2006, Orlando, FL.
- Hsieh, M.A., Kumar, V., & Taylor, C.J. (2004). Constructing radio signal strength maps with multiple robots. In Proc. IEEE International Conference on Robotics and Automation, New Orleans, LA.
- Mostofi, Y., Chung, T., Murray, R., & Burdick, J. (2005). Communication and sensing trade offs in decentralized mobile sensor networks: A cross-layer design approach. In 4th International Conference on Information Processing in Sensor Networks (IPSN), Los Angeles, CA.
- Mostofi, Y., & Murray, R. (2004). Effect of time-varying fading channels on the control performance of a mobile sensor node. In Proc. of IEEE 1st International Conference on Sensor and Ad Hoc Communications and Networks (Secon), Santa Clara, CA.
- Nahrstedt, K., Shah, S.H., & Chen, K. (2004). Cross-layer architectures for bandwidth management in wireless networks. In Resource Management in Wireless Networking.
- Neskovic, A., Neskovic, N., & Paunovic, G. (2000). Modern approaches in modeling of mobile radio systems propagation environment. IEEE Communications Survey and Tutorials, Vol. 5, No. 2.
- Pereira, G.A.S., Das, A.K., Kumar, V., & Campos, M.F.M. (2003). Decentralized motion planning for multiple robots subject to sensing and communication constraints. In Proceedings of the Second Multi-Robot Systems Workshop, pp. 267–278, Washington, DC.
- Pimentel, B.S., & Campos, M.F. M. (2003). Cooperative communication in ad hoc networked mobile robots. In Proceedings of the 2003 IEEE/RSJ Intl. Conference on Intelligent Robots and Systems, pp. 2876–1881, Las Vegas, NV.
- Powers, M., & Balch, T. (2004). Value-based communication preservation for mobile robots. In 7th International Symposium on Distributed Autonomous Robotic Systems, Toulouse, France.
- Redi, J., Watson, W., & Shurbanov, V. (2002). Energy conserving reception protocols for ad hoc networks. In Proceedings of MILCOM 2002, Anaheim, CA.
- Roumelilotis, S.I., & Bekey, G.A. (2002). Distributed multi-robot localization. IEEE Trans. Rob. Autom. 18, 781–795.
- Sweeney, J., Brunette, T.J., Yang, Y., & Grupen, R.A. (2002). Coordinated teams of reactive mobile platforms. In Proceedings of the International Conference on Robotics and Automation (ICRA), pp. 99–304, Washington, DC.
- Sweeney, J.D., Grupen, R.A., & Shenoy, P. (2004). Active qos flow maintenance in controlled mobile networks. In Proceedings of the Fourth International Symposium on Robotics and Automation (ISRA), Queretaro, Mexico. IEEE, New York.
- BBN Technologies. (2003). BBN JBox Informal Manual. BBN Technologies.
- Thibodeau, B.J., Fagg, A.H., & Levine, B.N. (2004). Signal strength coordination for cooperative mapping. (Tech. Rep.) Amherst, MA: University of Massachusetts, Department of Computer Science.
- Wagner, A.R., & Arkin, R.C. (2004). Communication-sensitive multi-robot reconnaissance. In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), pp. 2480–2487, New Orleans, LA.
- Winfield, A.F.T. (2000). Distributed sensing and data collection via broken ad hoc wireless connected network of mobile robots. In L.E. Parker, G. Bekey, and J. Barhen (Eds.), Distributed Autonomous Robotic Systems 4 (pp. 273–282). Berlin: Springer-Verlag.
- Xue, Y., & Nahrstedt, K. (2004). Providing fault-tolerant ad-hoc routing service in adversarial environments. Wireless Personal Communications, Special Issue on Security for Next Generation Communications.
- Zhang, B., & Sukhatme, G.S. (2005). Controlling sensor density using mobility. In The Second IEEE Workshop on Embedded Networked Sensors, pp. 141–149, Sydney, Australia.