Towards the deployment of a mobile robot network with end-to-end performance guarantees

Mong-ying A. Hsieh, Anthony Cowley, Vijay Kumar and Camillo J. Taylor
GRASP Laboratory
University of Pennsylvania
Philadelphia, PA 19104
Email: {mya, acowley, kumar, cjtaylor}@grasp.cis.upenn.edu

***Abstract*— Communication is essential for coordination in most cooperative control and sensing paradigms. In this paper, we present an experimental study of strategies for maintaining end-to-end communication links for tasks such as surveillance and search and rescue where team connectivity is essential for providing situational awareness to a base station. We consider the differences between monitoring point-to-point signal strength versus data throughput and present experimental results with our multi-robot testbed in outdoor environments.**

I. Introduction

We are interested in developing robotic teams that can operate autonomously in urban areas and perform tasks such as surveillance and search and rescue where maintaining team connectivity is essential for situational awareness. In both these tasks, robots must have the ability to align themselves along the boundaries of complex shapes in two dimensions. However, navigation based solely on the geometry of the environment will not always guarantee a connected communication network. In general, successful mission execution not only requires proper deliberative planning but also suitably designed reactive behaviors since teams of robots will have to operate with little to no *direct* human supervision. Therefore, each robot must also have the ability to maintain constraints with respect to its neighbors to enable the transmission of critical information back to the base station.

In general, it is difficult to predict radio transmission properties *a priori*, especially in unknown and unstructured environments, due to their sensitivity to a variety of factors including transmission power, terrain characteristics, and interference from other sources [1]. Fig. 1 shows actual signal strength measurements between two nodes for different separation distances. Although there is a strong correlation between signal strength and distance, there is also a lot of variability. Even when it is possible to map the communication characteristics for a given site, [2], there is no guarantee the measurements will reflect actual performance during mission execution. Furthermore, as team size increases, the issue of bandwidth becomes important since an acceptable level of signal strength no longer guarantees a robot's ability to transmit critical data.

Fig. 2(a) shows the number of transactions[1] per interval of time between four different robots, positioned at four distinct 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 robot successively begins 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 due to over-saturation situations where significant network resources are spent coping with low-level packet collisions, retries and contention resolution. These situations often occur 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 is virtually constant as shown in Fig. 2(b) since inter-robot distances were kept constant.

[1]This metric is defined more precisely later in Section III.

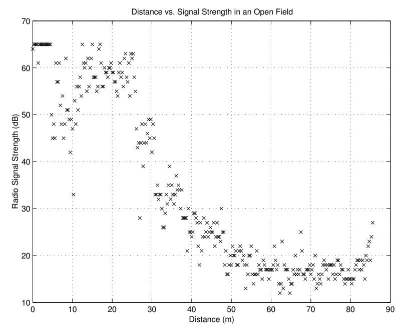

Fig. 1. Normalized signal to noise ratio measurements of the radio signal strength as a function of transmission distance in an open field. Transmitters and antennas were positioned 18.5 inches above the ground and the signal strength (y-axis) is normalized to a scale of 0 - 100 dB.

In this paper, we consider the problem of driving a team of robots to specific locations along a specified, parameterized, possibly closed curve while maintaining point-to-point communication links. We design low-level decentralized reactive controllers to constrain the motion of individual agents based on two link quality measures: signal strength or perceived network bandwidth. We present experimental results with our multi-robot testbed in outdoor environments using different interconnection topologies. Our strategies are designed to be minimally disruptive to the overall deliberative plan and provide situational awareness to the base station including potential failure points in the communication network. The strategies presented can be used in conjunction with information obtained from a radio connectivity map as shown in [2].

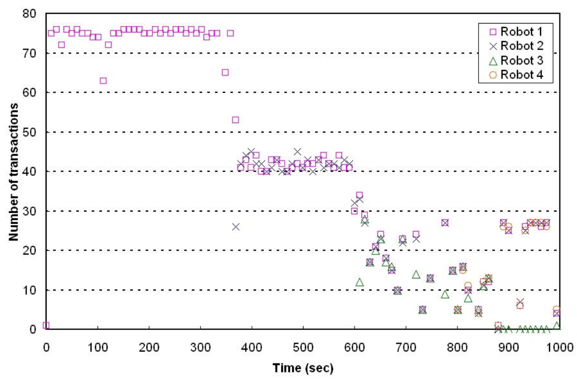

(a)

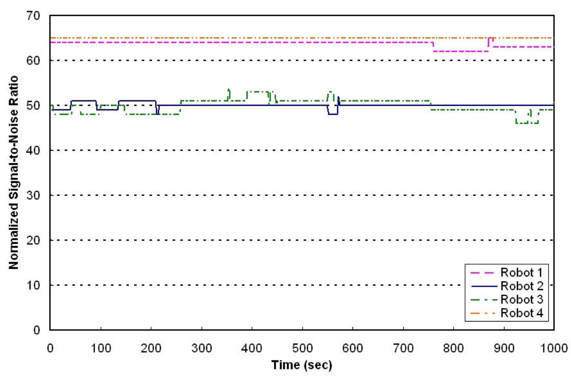

(b)

Fig. 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-to-noise ratio measurements between the Base Station and the transmitting robots for the same period of time.

II. Related Work

Communication is fundamental to most multi-agent coordinated and cooperative tasks. There is significant literature on multi-robot control [3], sensing [4], planning [5], and localization [6]. In the last couple of years, increased attention has been focused on the value of communication to multi-agent teams. Earlier works assumed static communication ranges, [7], and/or relied on coordination strategies that require direct line-of-sight, [8]. In [9] and [10] decentralized controllers were used for concurrently moving toward goal destinations while maintaining communication constraints by maintaining line-of-sight and assuming static communication/sensor ranges respectively. The formation of communication relays between any pair of robots using line-of-sight was discussed in [11]. Although line-of-sight may significantly improve the ability to communicate, often times it is not necessary and may potentially be too restrictive. Thibodeau et al., through simulation results, showed in [12] that coordination strategies based on line-of-sight maintenance for cooperative mapping are less efficient than strategies based on inter-agent wireless signal strengths.

Coordination strategies based on inter-agent signal strength include [13], where the combination of planning and reactive behaviors for communication link maintenance in a multi-robot team for reconnaissance was used. Here, reactive behaviors as well as contingency plans were designed by hand and simulation results were presented for teams of two to four robots. Navigation based on perceived wireless signal strength between robots for exploration was presented in [14]. Sweeney et al. used a null-space projection approach to navigate each robot towards its goal while maintaining point-to-point communication links and included simulation results for a team of four planar robots. In [15], 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. The use of wireless communication for localization is discussed in [16], and for localization and navigation in [17].

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. Our approach entails the design of low level decentralized reactive controllers that respond to changes in actual signal strength or verified network bandwidth. The goal is to develop strategies that remain minimally disruptive to any high level deliberative plans while maximizing the team's ability to provide effective situational awareness to a base station.

III. Methodology

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. Given a simple kinematics model, $\dot{q}_i = u_i$, where i denotes the i^{th} agent, $q_i = (x_i, y_i)^T$ and u_i denote the i^{th} agent's position and control input, we propose the following decentralized controller

$$u_i \quad = \quad -k\nabla_i \phi_i(q_i) - \sum_{j\in\mathcal{N}_i} \nabla_i g_{ij}(q_{ij}) \qquad (1)$$

where k is a positive constant scalar, ϕ is some artificial potential function, and $\mathcal{N}_i$ denotes the set of neighbors for agent i. The first term of the control law (1) guides each robot to its goal destination and the second term maintains connectivity with a pre-specified set of neighbors using the artificial potential functions, g_{ij}, to model the constraints between pairs of robots. ∇_i denotes the partial derivatives with respect to the coordinates of the i^{th} robot.

In our work, g_{ij} are measured quantities, either signal-to-noise ratio (SNR) measurements as in Fig. 2(b) or number of transactions per unit time as in 2(a). Further g_{ij} depends on an underlying communication graph that specifies the connections that need to be maintained.

A. 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 devices that, among other things, provide point-to-point signal strength measurements.

In general, 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, and the signal strength between sender and receiver. In multi-robot applications, it is often relevant to talk about bandwidth in terms of units of application level data that can be transmitted as opposed to low-level packet loss which is a network protocol dependent variable. Therefore, we define a *successful transaction* to be the transmission of one unit of application level data sent by a sender with an acknowledgement of receipt sent by the receiver. A robot's conservative estimate of available bandwidth is determined based on the number of successful transactions it can achieve over a specified interval of time.

For our experimental setup, we set one unit of application level data equal to a JPEG image of approximately 10 KB in size. Thus, a successful transaction is the transmission of a 10 KB JPEG image by the sender followed by the receipt of acknowledgement sent by the receiver. Then, based on the desired transaction rate, *i.e.* number of successful transactions per time interval, the robot will periodically evaluate its connection with the Base.

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 the prescribed minimum data throughput rate is 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 sufficient to verify that the constraint is met, a connection monitor will periodically verify available throughput. This latter mechanism allows us to deploy robots that do not maintain consistent data flow back to the Base, e.g. robots that only send event data, yet still ensure that the available throughput will likely be available if it should be needed.

B. Controllers

Our decentralized 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 equation (1). For each target waypoint, a reference heading, similar to the descent direction of a potential field controller is computed. Based on the descent direction, a look ahead waypoint is generated based on the vehicle's speed and position. Then a simple PID controller is used to steer the robot towards the look ahead waypoint. The process is repeated until the target location is reached.

To maintain constraints, each robot continuously monitors the quality of its communication link(s) to its specified set of neighbors. When the link quality drops below a minimum acceptable threshold, the controller can either stop the robot or move it closer to its neighbor(s) 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. Similarly, if the robot perceives an increase in its link quality above the acceptable level, it will once again attempt to move towards its target. Such recovery measures may be used to lessen the times a robot is caught in a spatial-temporal dip in the link quality due to dynamic changes in the environment. Furthermore, these measures also ensure that a robot is constantly minimizing its distance to the goal while respecting the constraints. The algorithm is summarized below.

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;
        stopped = true;
        waitTime = current time;
    end if
    if waitTime > MaxWaitTime then
        Retry going to goal;
        stopped = false;
        recover_flag = false;
    end if
end if
if LinkQuality > Acceptable then
    Go to goal;
end if
```

IV. Experimental Results

Our multi-robot team consists of four unmanned ground vehicles (UGVs) built from radio controlled scale model trucks each equipped with a laptop computer, odometry, stereo camera, GPS receiver, and a small embedded computer with 802.11b wireless Ethernet for network connectivity, called the Junction Box (JBox). The JBox provides point-to-point signal strength measurements for all pairs of nodes on the network and handles multi-hop routing in an ad-hoc wireless network. All programming and tasking of the vehicles are done using ROCI (Remote Object Control Interface) [18], [19], which is a high level operating system designed specifically for tasking and managing networks of robots and sensors.

Experiments were conducted in two separate outdoor environments. The first experiment was conducted at the Military Operations on Urban Terrain (MOUT) training site, shown in Fig. 3(a), located in Ft. Benning Georgia. The remaining experiments were conducted at a soccer field at the University of Pennsylvania. A satellite image of the soccer field and a detailed schematic of its surroundings are shown in Fig. 5.

(a)

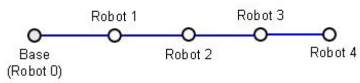

(b)

(c)

Fig. 3. (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 90m $\times$ 120 m. The location of the Base is denoted by $\bigcirc$ 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 which denoted by $\bigcirc$ and $\times$ respectively.

A. Reconnaissance at a MOUT Site

We considered a reconnaissance application where the objective was to deploy a team of four robots to obtain video surveillance imagery at a designated location out of single-hop radio communication range. The team of four robots were deployed to four separate target locations shown in Fig. 3(a). The i^{th} robot was tasked to monitor its signal strength with respect to the $(i-1)$ robot and stop when the signal strength dropped below the acceptable threshold. The Base was considered the 0^{th} robot. As shown in Fig. 3(c), the targeted locations were not reached since each robot was maintaining the signal strength to its designated neighbor above the required threshold.

B. Perimeter Surveillance Application

We based these experiments on a perimeter surveillance application in which each robot is required to send imagery data back to a base station. 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 is 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 Fig. 2, as new members are introduced into the team, the maximum bandwidth available to each robot drops.

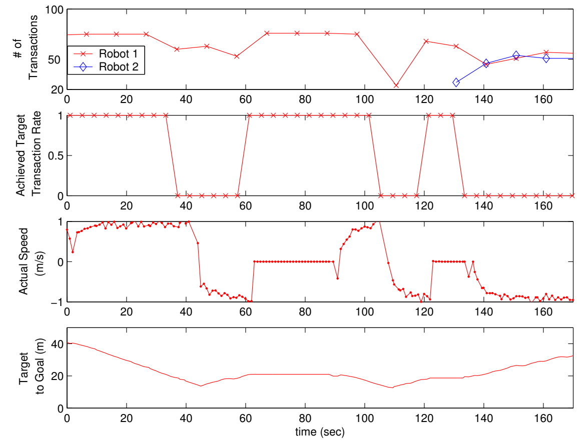

Fig. 4. Top: Number of transactions received by the Base from R1 and R2. R2 began its transmission at around $t = 130s$. Center Top: 1 denotes R1 achieved the target transaction rate and 0 otherwise. Center Bottom: Actual speed achieved by R1. Positive speed denotes the robot is moving towards the goal and negative speed denotes the robot is moving towards the Base. Bottom: The robot's distance from the goal.

In Fig. 4, we show how our controller responds to changes in the team size. We first deployed a single robot, R1, to a goal position and required that it continuously send imagery data to the Base while maintaining a minimum acceptable transaction rate. In the second graph of Fig. 6(a), 1 is used to denote the target transaction rate was achieved and 0 otherwise. A schematic of the deployment strategy is shown in Fig. 6(a). At around $t = 60s$, R1 settled to a location about halfway to the goal as shown in the bottom graph of Fig. 4. Between $t = 90s$ and $t = 130s$, the robot attempted to reach its goal a second time and settled to a similar location. A second robot, R2, transmitting to the Base was introduced to the network at approximately $t = 130s$, as shown in the first graph of Fig. 4. Immediately, R1 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 Fig. 4 where the robot's distance to the goal starts increasing.

The behavior demonstrated by the two-robot experiment allows 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 acceptable transaction rate. 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.

Goal positions for each team member were chosen to provide a wide net of surveillance coverage. This type of goal

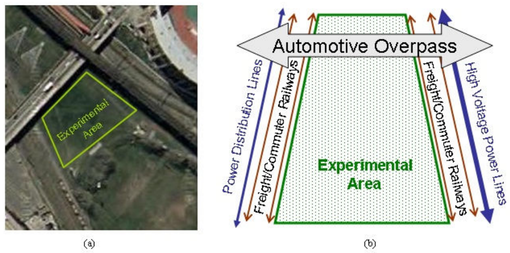

Fig. 5. (a) A satellite image of the soccer field and its surrounding. (b) Detailed schematic of the experimental area and its surroundings.

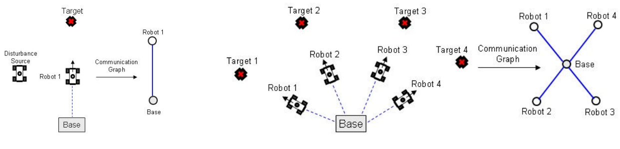

(a) (b)

Fig. 6. (a) Schematic of experimental setup and underlying communication graph for the results shown in Fig. 4. On the left, the dashed line denotes the communication link monitored by the robot. In this experiment, R2 is used to cause a network disturbance by transmitting to the Base. (b) Schematic of experimental setup and underlying communication graph for the results shown in Fig. 7 and 8. Similarly, the dashed lines in the figure on the left denote communication links monitored by each robot.

specification is flexible in that it establishes a vector for the robots to move along, as opposed to specific waypoints to achieve. Thus 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 Section III, each robot would move towards its goal until the link quality dropped below the minimum acceptable level, at which point it would move back towards the Base and stop when the link quality rose back above the chosen minimum acceptable 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 Fig. 6(b) and the results are shown in Fig. 7 and 8.

Fig. 7 shows the normalized SNR between one of the robots utilizing signal strength as a measure of link quality and the Base, along with the corresponding commanded speed and actual speed. The solid black line in the top graph in Fig. 7 denotes the minimum acceptable level. Initially, when the robot is close to the Base, the SNR measurements were high. As the robot moves toward its goal, we see these measurements drop. The first time the SNR dropped below the minimum acceptable level, around $t = 45s$, the robot attempted to backup and move closer to the Base. Once the SNR 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.

Similarly, Fig. 8 shows the results for a robot whose controller was reacting to changes in its estimated transaction rate.

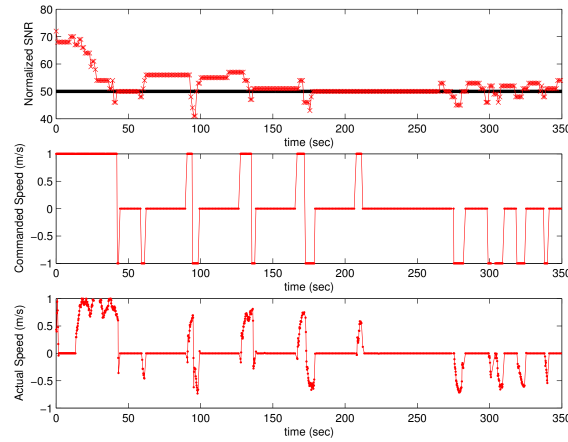

Fig. 7. Top: Normalized SNR between the robot and the Base. The solid black line denotes the minimum acceptable level. Center: Commanded speed based on the SNR measurements shown above. Positive speed denotes the robot is moving towards the goal and negative speed denotes the robot is moving towards the Base. Bottom: Estimated speed achieved by the robot based on the commanded speed.

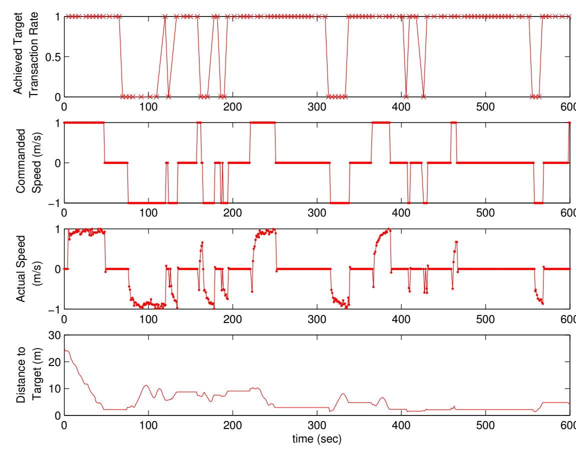

Fig. 8. Top: 1 denotes the target transaction rate was achieved and 0 otherwise. Top Center: Commanded speed based on whether the target transaction rate was achieved. Positive speed denotes the robot is moving towards the goal and negative speed denotes the robot is moving towards the Base. Bottom Center: Actual speed achieved by the robot based on the commanded speed. Bottom: The robot's distance from the goal.

Similar to the results shown in Fig. 7, the robot's transaction rate drops as it moves further away from the Base as shown in the first three graphs of Fig. 8. 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 last graph in Fig. 8 where at approximately $t = 50s$ the robot is within 2.5 meters 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 $t = 75s$, the robot began to move back towards the Base, leaving its goal location.

V. CONCLUSION

In this work, we have shown the importance for individual robots to have the ability to monitor actual signal strength measurements as well as available data throughput since signal strength alone may not guarantee successful transmission of critical data. 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 by avoiding excessive network clutter from low-level packet retries.

Additionally, 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 unacceptably low throughput may indicate a need for human attention to the network architecture and the demands being placed on it.

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. 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 deployment, 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 or a reallocation of resources.

Acknowledgments

The authors would like to thank the Penn Athletics Department, Jim Keller from the University of Pennsylvania 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

[1] A. Neskovic, N. Neskovic, and G. Paunovic, "Modern approaches in modeling of mobile radio systems propagation environment," *IEEE Communications Surveys*, vol. Third Quarter, 2000.

[2] M. A. Hsieh, V. Kumar, and C. J. Taylor, "Constructing radio signal strength maps with multiple robots," in *Proc. IEEE International Conference on Robotics and Automation*, New Orleans, LA, April 2004.

[3] R. Fierro, P. Song, A. Das, and V. Kumar, "Cooperative control of robot formations," in *Cooperative Control and Optimization: Series on Applied Optimization*, R. Murphey and P. Paradalos, Eds. Kluwer Academic Press, 2002, pp. 79–93.

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

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

[6] S. I. Roumelilotis and G. A. Bekey, "Distributed multi-robot localization," *IEEE Transactions on Robotics and Automation*, vol. 18, pp. 781–795, October 2002.

[7] A. F. T. Winfield, "Distributed sensing and data collection via broken ad hoc wireless connected network of mobile robots," in *Distributed Autonomous Robotic Systems 4*, G. B. L. E. Parker and J. Barhen, Eds. Springer-Verlag, 2000, pp. 273–282.

[8] R. Arkin and J. Diaz, "Line-of-sight constrained eploration for reactive multiagent robotic teams," in *AMC 7th International Workshop on Advanced Motion Control*, 2002.

[9] J. Sweeney, T. J. Brunette, Y. Yang, and R. A. Grupen, "Coordinated teams of reactive mobile platforms," in *Proceedings of the International Conference on Robotics and Automation (ICRA)*, Washington, DC, 2002, pp. 99–304.

[10] G. A. S. Pereira, A. K. Das, V. Kumar, and M. F. M. Campos, "Decentralized motion planning for multiple robots subject to sensing and communication constraints," in *Proceedings of the Second Multi-Robot Systems Workshop*, Washington, DC, March 2003, pp. 267–278.

[11] S. O. Anderson, R. Simmons, and D. Goldberg, "Maintaining line of sight communications networks between planetray rovers," in *Proceedings of the 2003 IEEE/RSJ Intl. Conference on Intelligent Robots and Systems*, Las Vegas, NV, October 2003, pp. 2266–2272.

[12] B. J. Thibodeau, A. H. Fagg, and B. N. Levine, "Signal strength coordination for cooperative mapping," Department of Computer Science, University of Massachusetts, Amherst, Tech. Rep., 2004.

[13] A. R. Wagner and R. C. Arkin, "Communication-sensitive multi-robot reconnaissance," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2004, pp. 2480–2487.

[14] J. D. Sweeney, R. A. Grupen, and P. Shenoy, "Active qos flow maintenance in controlled mobile networks," in *Proceedings of the Fourth International Symposium on Robotics and Automation (ISRA)*. Queretaro, Mexico: IEEE, August 2004.

[15] M. Powers and T. Balch, "Value-based communication preservation for mobile robots," in *7th International Symposium on Distributed Autonomous Robotic Systems*, 2004.

[16] A. Howard, S. Siddiqi, and G. S. Sukhatme, "An experimental study of localization using wireless ethernet," in *4th International Conference on Field and Service Robotics*, July 2003.

[17] P. Corke, R. Peterson, and D. Rus, "Networked robots: Flying robot navigation using a sensor net," in *Proceedings of the Eleventh International Symposium of Robotics Research (ISRR)*, ser. Springer Tracts on Advanced Robotics (STAR). Springer-Verlag, October 2003.

[18] L. Chaimowicz, A. Cowley, V. Sabella, and C. J. Taylor, "Roci: A distributed framework for multi-robot perception and control," in *Proceedings of the 2003 IEEE/RJS International Conference on Intelligent Robots and Systems*, Las Vegas, NV, 2003.

[19] A. Cowley, H. Hsu, and C. J. Taylor, "Distributed sensor databases for multi-robot teams," in *Proceedings of the 2004 IEEE International Conference on Robotics and Automation*, New Orleans, LA, 2004.