

Indoor Target Intercept Using an Acoustic Sensor Network and Dual Wavefront Path Planning

Lynne E. Parker, Ben Birch, and Chris Reardon

Department of Computer Science, The University of Tennessee, Knoxville, TN 37996–3450

Email: {parker, birch, reardon}@cs.utk.edu

Abstract—This paper presents an approach that enables a mobile “Interceptor” robot to intercept targets in an indoor environment using information from a distributed acoustic sensor network. The approach assumes the indoor environment has been previously mapped and that the sensor nodes know their position in the map. The targets are localized in the sensor network based upon local maxima of the acoustic volume. The current target localization information is reported to an Interceptor robot, which utilizes a dual wavefront path planner to move from its current location to a location that is within visibility range of a target. Results of the complete implementation of this approach using 70 sensor net robots in the Player/Stage multi-robot simulator are reported, as well as implementation results to date on a team of physical robots. To our knowledge, this is the first implementation of a multi-robot system that combines the use of an acoustic sensor net for target detection with an Interceptor robot that can efficiently reach the moving position of the detected target in indoor environments.

I. INTRODUCTION

In many security, surveillance, and reconnaissance applications, an indoor area must be monitored for unusual activity. For example, in warehouse applications, a sensor network could be utilized to detect unexpected motion, which could then trigger a call to a more capable mobile response robot, or *Interceptor*. The Interceptor robot would typically have more sophisticated sensors that allow it to move to the location of the target and then transmit images or other relevant data back to the human monitor.

In this paper, we describe an approach enabling an Interceptor robot to use information from a distributed acoustic sensor network to intercept a target in an area of interest. The distributed acoustic sensor network we use in this research is very simple, involving only a single low-fidelity microphone at each sensor node. In our current work, we assume that the only sounds in the environment are those from the target that moves through the environment. The sensor nodes are able to communicate with each other and with the Interceptor robot. We assume that the environment has been previously mapped by other robot team members using a technique such as [1], and that the sensor nodes know where they are in the map. We also assume that the Interceptor robot has a vision system onboard that enables it to visually detect and track the target once it is within visual range. This paper describes (1) how the distributed acoustic sensor network determines an estimation of the target location, and (2) how an Interceptor robot uses this information to dynamically replan and execute a path to move within visibility range of the target.

In Section II, we describe research related to this project. Section III then discusses the technique used by the sensor net to localize the acoustic target. In Section IV, we describe the dual wavefront path planning technique we use to enable the Interceptor to move within visibility range of the target. Section V describes our results from implementation of this approach in simulation and on a team of physical robots. We conclude with some summary remarks in Section VI.

II. RELATED WORK

In recent years, an extensive amount of research has investigated the topic of sensor networks for acoustic target localization. The types of acoustic sensors vary – some of the sensors are able to detect the distance to an acoustic target (e.g., [2]), while others can detect the direction to an acoustic target (e.g., [3], [4]), and still others assume only acoustic volume is detectable (e.g., [5]). Some techniques (e.g., [2]) involve the use of a known strength of signal at the source in order to estimate the distances to the target and combine them in real time. The approach in [6] assumes a time synchronized microphone array and uses time difference of arrival between microphones and microphone arrays to localize targets. In all of these cases, the sensor net integrates the detected information to localize the target.

In our application, each sensor node has only one low-fidelity microphone. Thus, individual acoustic sensors are only able to detect volume; the distance and direction to the target cannot be detected. Our real world constraints also prevent us from using a known energy output level from the target, both because of practical issues of varying sound levels at the source and declining microphone sensitivity as battery power declines. Sophisticated algorithms for target localization under these constraints have been developed in [5], which uses mathematical analysis to estimate the target position based upon the strengths of the signals heard by neighboring nodes. However, these analytical models were designed for open-air environments where acoustic signal propagation models are known. Unfortunately, there exist no known models for acoustic propagation in indoor environments; thus we are not currently able to take advantage of these precise analytical methods. As [5] notes, the complexity of the more sophisticated algorithms is justified only if their accuracy is higher than the node spacing itself. Since this will likely not be the case in our application, we make the same conclusion as [5] that a reasonable approximation to target localization is the

location of the node hearing the highest volume. In the case of multiple targets, the local maxima of volumes heard will be the locations of multiple targets, within the resolution of the sensor node placements.

Likewise, a significant amount of research has addressed the issue of moving target interception. Techniques for enabling a mobile robot to intercept a moving target have been developed in a wide range of application domains. These strategies typically depend on an ability to predict the path of the target over time. However, unlike bouncing balls or ballistic missile motions, it is generally quite difficult to predict the motions of a human target moving through an indoor environment unless prior models have been developed over time that track the typical behavior patterns of humans in the environment. Since this modeling is not relevant to our application, it is not meaningful to try to predict a target's future path. Our approach, therefore, is to plan a path to the current estimation of the target position, and then dynamically replan the path as the target position changes.

While an extensive number of path planning techniques are possible, we use a dual wavefront propagation algorithm to plan a path to the target position. Several variations of the wavefront propagation algorithm have been implemented, such as Murphy's Trulla system [7]. Our approach is similar to the work of Behring, *et al.* [8]. However, we make modifications to the Behring algorithm, which are described in section IV.

We are not aware of any other previous work that has combined the use of a distributed acoustic sensor network for target localization with an Interceptor robot that is able to efficiently move to the estimated target location in an indoor environment.

III. DISTRIBUTED ACOUSTIC SENSOR NET TARGET DETECTION

In our research, the acoustic sensor nodes are simple AmigoBot robots that use an iPAQ computer for computations, along with wireless communications capabilities. The iPAQ computer has an installed simple microphone that is used for acoustic sensing. In our larger research project, we address the issue of sensor net deployment of these simple robots using heterogeneous herding techniques (see our companion paper [9] for details). For this paper, we assume that a map of the environment has been previously built by robot team members (e.g., using a technique such as [1]), and that the acoustic sensor nodes have been deployed to known positions within the map.

Once the acoustic sensor robots have been deployed, they activate their distributed acoustic sensor net to detect targets that are moving through the environment. Since our robots are currently equipped with only a very simple microphone, we assume that the target moving through the environment is making some detectable noise, and that the target is the only source of sound. As illustrated in Figure 1, after initially filtering the volume information, our approach has each robot reporting its volume heard to its local neighbors (indicated by the edges in the figure). Because of noisy data generated

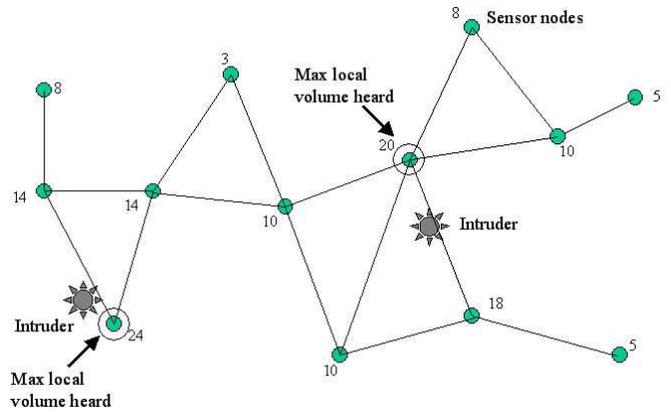


Fig. 1. Illustration of how the distributed acoustic sensor net operates. The small circles represent sensor node positions. The edges represent nearby neighbor communications paths. The numbers represent the volume heard at that node. The circled sensor nodes are the local maxima, determined via communications with neighboring nodes.

by our simple microphones, merely reporting the location of the highest raw volume heard is not sufficient to define the target's current position. We therefore define a *qualifying ratio*, which gives the fraction of time a sensor node must hear the highest volume (relative to its immediate neighbors) during the most recent predefined (short) time interval. If the robot's volume heard over the time interval exceeds the qualifying ratio, it reports its position as the current target position estimate. If there is more than a single sensor robot that exceeds the qualifying ratio (as is sometimes the case when the ratio is less than 0.5), then all such robots will report their positions. If multiple targets are in the environment, all the target locations in the sensor network are reported. If there is a tie in the detected maximum volumes of two neighbors, then both report their positions. If the target is moving, these ties will be broken very quickly. In Figure 1, the numbers at each sensor node indicate the volume heard at that node. In this figure, two nodes detect local maxima of the sensor network, corresponding to two targets. These two nodes are the only nodes whose volume heard is higher than all of their neighbors. The pseudocode for this distributed acoustic detection algorithm is shown in Table I.

IV. DUAL WAVEFRONT PATH PLANNING FOR TARGET INTERCEPT

Like the integration of the Trulla path planner into a reactive architecture described in [5], our target intercept algorithm can also be thought of as being composed of a deliberative component (path planning) and a reactive component (navigation along the path to the target while avoiding obstacles). Since the essence of the intercept function is being able to quickly react to changes in the target position, we wanted a path planner with execution times that would not hamper the operation of the reactive component of the algorithm. Our empirical results show that the dual wavefront path planner approach described here is capable of meeting this requirement.

TABLE I
THE DISTRIBUTED ACOUSTIC SENSING ALGORITHM.

Distributed Acoustic Sensing Algorithm

While (forever)

- Filter sound – For each sound instant:
 - Subtract out all noise below a predefined volume;
 - Use only sound data within specified frequency range;
 - Average sounds over a short period (approx. 1/3 second)
- Communicate my filtered sound (h) to my nearest neighbors;
- Receive $V[1..n]$ volumes from my n nearest neighbors;
- Update fraction of highest volume heard for most recent time interval. (i.e., if $h > V[i]$ for all i , then fraction increases; else, it decreases.)
- If my fraction of highest volume heard $>$ qualifying ratio,
 - Broadcast to the Interceptor robot my position, g , as nearest to the detected target.

Our approach to path planning takes as input a map generated by mapping robots and uses dual-wavefront propagation in an 8-connected occupancy grid to generate an optimal path from the starting position to the goal position. This path consists of a series of waypoints that the robot should travel through. Beginning with the initial occupancy grid, obstacles are expanded to enable the generation of paths that keep robots away from potential collisions. This algorithm then propagates wavefronts from both the starting position and the goal position until the wavefronts meet. A path is then extracted from these wavefronts that minimizes the Euclidean distance between the starting and goal positions. This path, which is a series of waypoints, is then smoothed to eliminate unneeded waypoints and jagged path edges. The result is a series of waypoints that take the robot from the starting to the goal position.

Our approach is similar to the work of Behring, *et al.* [8]. However, we make two modifications to the Behring algorithm. First, instead of propagating a single wavefront from the starting position, we implement dual wavefront propagation from the starting and goal positions, which is much more computationally efficient.

Second, instead of associating the cells on the wavefront boundary with the current number of iterations, we associate these cells with the Euclidian distance that has accumulated as the wave has propagated outward through neighboring cells. In other words, after n iterations, cells on the wavefront boundary can take on values ranging from n (no diagonal travel) to $n\sqrt{2}$ (all diagonal travel). When the wavefronts meet, the two neighboring cells on opposite sides of the boundary with the smallest combined distance from the initial and goal positions are placed on the path. From these points, a path is constructed first back to the initial position, then onward to the goal position, by adding to the path in iterative steps the neighboring cell having the smallest accumulated distance.

TABLE II
THE DUAL WAVEFRONT PATH PLANNER.

Dual Wavefront Path Planner (map : M , position : s , position : g)

- 1) Convert map M to occupancy grid.
- 2) Expand obstacles.
- 3) Simultaneously propagate wavefronts from Start (s) and Goal (g) points as follows:
 - While (not finished propagating)
 - Propagate wave from s , marking adjacent grid cells with the minimum Euclidian distance from s to the centers of these cells;
 - Propagate wave from g marking adjacent grid cells with the minimum Euclidian distance from g to the centers of these cells;
 - If two wavefronts meet, we have finished propagating.
- 4) Extract path, p , in two phases: a) from wavefront meeting point back to s , and b) from wavefront meeting point onwards to g , as follows (execute twice, using temporary variable z , where $z = s$ then $z = g$):
 - Select the two neighboring cells on opposite sides of the wavefronts' boundary which have the lowest combined marked Euclidian distance;
 - Place the center points of these two cells on the path;
 - While (not extracted back to z)
 - Select adjacent grid cell with the minimum marked Euclidian distance;
 - Add the center of this grid cell as a waypoint in the path;
 - If this grid cell contains z , the extraction is complete.
 - Return planned path p , which is a series of waypoints.
- 5) Smooth path, p , as follows:
 - Remove all waypoints from p that are colinear with their adjacent waypoints.
 - Remove any waypoint i from p such that the line between waypoints $i - 1$ and $i + 1$ does not intersect with an obstacle. Iterate until no remaining waypoints can be removed without causing path to intersect with an obstacle.
 - Return smoothed path p , which is a reduced series of waypoints.

Our second modification has thus enabled us to ensure that the extracted path will always be a shortest path. The pseudocode for this algorithm is shown in Table II.

Once the distributed acoustic sensor net has detected approximate target positions, this position information is sent to the Interceptor robots. The appropriate algorithm for target interception depends upon the resolution of the current target position calculation. In the case we have outlined in this paper, the resolution is very coarse, in which the target position is determined only to the nearest sensor net robot position. With this coarse information, it is not possible to estimate the current trajectory of the target; instead, the information has to be combined with map reachability constraints to estimate a possible target trajectory. Future research may explore this

TABLE III
THE TARGET INTERCEPT ALGORITHM.

<i>The Target Intercept Algorithm</i> (<i>map</i> : <i>M</i>)
<ul style="list-style-type: none"> • Repeat until done: <ul style="list-style-type: none"> - g = Target position from Acoustic Sensor Net. - s = Interceptor position from Localizer. - p = Path from Dual_Wavefront_Path_Planner(M, s, g) - distance_to_target = L = Path_Length (p) - If (distance_to_target < visual_range), exit loop. - distance_traveled = 0 - While (distance_traveled < $L/2$) <ul style="list-style-type: none"> * Move along path p * s = Interceptor position from Localizer. * Update distance_traveled • Transfer to visual target tracking.

possibility.

The algorithm for target intercept, outlined in Table III, involves continually replanning and executing a path to the current detected location of the target using the dual-wavefront propagation algorithm. We assume the Interceptor robot has a Localizer module (e.g., [1]) that enables it to localize itself in the environment. Once the Interceptor has moved halfway along its planned path, it replans to take advantage of better solutions that may be possible due to target motion.

Once the Interceptor robot is within a pre-defined distance from the target, the Interceptor robot activates a visual search behavior to search for the target, which we call “visual tracking” mode. Once the target has been detected visually, the Interceptor robot uses a vision-based tracking algorithm to keep the target within view. If the target is lost, the Interceptor robot again obtains information from the distributed acoustic sensor network on the approximate detected position of the target, repeating the process. In this case, the dual wavefront path planner would again be used to move to the vicinity of the target and repeat the attempt for visual detection and tracking.

V. RESULTS

We have implemented this approach both in simulation and on a team of physical robots. We report here the results of each.

A. Simulation Results

For our simulation studies, we used the Player/Stage simulator [10] using 70 pre-positioned acoustic sensor robots in an indoor environment. The indoor environment for these studies was the floor plan of the decommissioned hospital at Fort Sam Houston, in San Antonio, Texas, which as overall dimensions of 140 meters by 53 meters. The cell size of the occupancy grid was approximately 20 centimeters. For the example environment reported in this paper, obstacles are expanded in the grid representation by approximately 20 centimeters to enable generation of robot paths that keep the robot team away from obstacles. Figure 2 shows a snapshot of the distributed acoustic sensor net from these simulation studies. In this

TABLE IV
TIMING RESULTS FROM PATH PLANNER.

Path Length (meters)	Planning Time (sec)
100	22
50	10
25	4
12	2
6	1

implementation, we model volumes as falling off with the square of the distance from the sound source. However, we have not modeled the physics of sound propagation through and around walls in the simulation.

We have fully implemented the dual wavefront propagation algorithm for target intercept in Player/Stage, and it operates in a variety of environmental maps. Figure 3 shows an example of the flooding process of the dual wavefront path planner, as well as the resulting smoothed path that is extracted from the wavefront propagation process.

Figure 4 shows a simulation result of the dual wavefront path planner applied for target intercept. In this figure, the gray path from the left wing of the hospital, into and out of several rooms to the central region of the building, is the actual path taken by the target. The black path from the right wing of the hospital to the central region is the actual path dynamically planned and executed by the Interceptor robot in response to the target motion. The Interceptor robot converts to a visual tracking behavior once it is within sight of the target, in one of the rooms near the center of the building.

To evaluate the run time capabilities of our algorithm, we collected data on the planning time required for different path lengths. Table IV shows the execution times of our path planner as it plans paths of varying lengths in the Fort Sam Houston hospital map. These times include wavefront propagation, path extraction, and path smoothing (steps 3-5 in Table II). Not included in these times are steps 1 and 2 in Table II, which are performed once, in preparation for assuming the Interceptor role, and have a combined execution time of 4 seconds for this example.

As can be seen from the table, planning time decreases roughly linearly with decreasing path length. It is this property that allows the path planner to replan paths in a nearly transparent manner as the robot moves closer to the target, and to complete the interception in what appears to be a purely reactive way. A video of a sample run in Player/Stage of a target moving through the distributed acoustic sensor net and being intercepted by the Interceptor robot is available.

B. Physical Robot Results

The distributed acoustic sensor net algorithm has been implemented on a team of AmigoBot robots, as shown in Figure 5. These robots have an iPAQ computer onboard running Linux. The iPAQs have a low-fidelity microphone that we use for detecting sound. The robots communicate using an ad hoc wireless network. The behavior software interfaces with the robot and microphone hardware using Player. In our

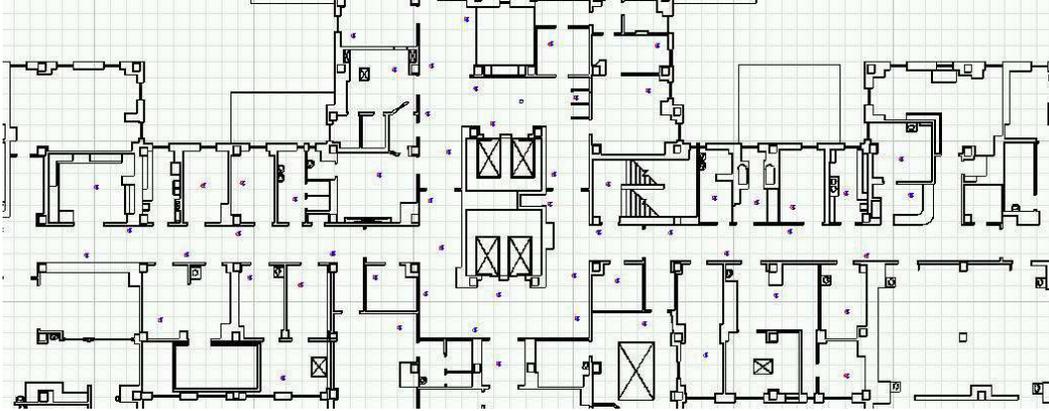


Fig. 2. Deployed sensor net in Player/Stage simulator. Each small robot icon (or dot, at this small image resolution) shows the position of a mobile sensor node.



Fig. 3. Dual wavefront path planner in operation (read from left to right, from top to bottom). The robot starting position is in the right wing of the building. The robot goal position is near the center of the building. The final figure shows the smoothed path (thin black line) extracted from the dual wavefront propagation.

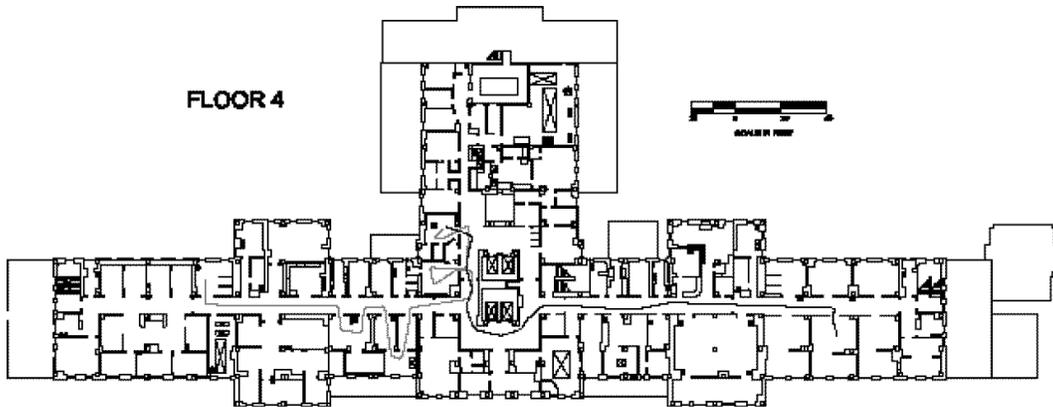


Fig. 4. Simulation results showing our approach for intercepting targets. The gray path from the left is the actual path taken by the target. The black path from the right is the actual path dynamically planned and executed by the robot interceptor.



Fig. 5. Physical robot implementation of the acoustic sensor network on a team of AmigoBot robots.

implementation, the robot that hears the loudest volume rotates slightly in each direction effecting a “wiggle,” which provides visual feedback to the experimenter as to which robot is closest to the detected target. Using our approach, these robots are able to determine which robot is hearing the loudest volume and which robot is most closely tracking a target and then report the position of that robot to a control station.

In our experiments, we have used a number of sounds for target localization, including buzzers, walking with loud shoes, moving while talking, and so forth. All of these sounds are detected fairly reliably by the sensor network. Of course, different sounds propagate differently in the indoor environment, and thus the responsiveness of the robots is somewhat different depending on the sound characteristics. Our best results were generated while using a buzzer operating at a high, audible frequency and inter-robot distances around 3-5 meters worked well for accurate localization. We have had success with physical robot implementations of 10-15 robots and are working to expand that number. A video of our physical robot implementation is available.

Implementation of the Interceptor robot functions on physical robots is underway, involving the incorporation of localization functions on an ATRV-mini robot. Once this implementation is complete, the Intercept robot function will be integrated with the acoustic distributed sensor network.

VI. CONCLUSIONS

This paper has presented algorithms for indoor target intercept using an acoustic sensor network and a dual wavefront path planning process. Our approach involves simple sensor

nodes reporting the local maxima of detected sound as the approximate position of a target. This information is dynamically fed to an Interceptor robot, which uses a dual wavefront path planning process to continually replan a path to intercept the moving target. We have successfully implemented our approach in the Player/Stage simulator and presented results of this implementation. We have also implemented the distributed acoustic sensor net on a team of physical robots, and have demonstrated success in detecting the location of targets to the resolution of the distances between sensor nodes. Work is underway to complete the implementation of the robot Interceptor capabilities on a physical robot, and to integrate these capabilities with the distributed acoustic sensor network. To our knowledge, this is the first implementation of a multi-robot system that combines the use of an acoustic sensor network for target detection with an Interceptor robot that can efficiently reach the moving position of the detected target in indoor environments.

ACKNOWLEDGMENTS

The authors thank Yifan Tang, Kannan Balajee, and Xiaoquan Fu for their valuable discussions regarding this research. This research was sponsored in part by DARPA/IPTO’s Software for Intelligent Robotics program, through Science Applications International Corporation, and in part by The University of Tennessee’s Center for Information Technology Research. This paper does not reflect the position or policy of the U. S. Government and no official endorsement should be inferred.

REFERENCES

- [1] S. Thrun, D. Fox, and W. Burgard, “A probabilistic approach to concurrent mapping and localization for mobile robots,” *Autonomous Robots*, vol. 5, pp. 253–271, 1998.
- [2] Z. Deng and W. Zhang, “Localization and dynamic tracking using wireless-networked sensors and multi-agent technology: first steps,” in *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, 2002, pp. 2386–2395.
- [3] D. Estrin, R. Govidan, J. Heidemann, and S. Kumar, “Next century challenges: Scalable coordination in sensor networks,” in *Proceedings of MOBICOMM 1999*, 1999, pp. 263–270.
- [4] L. Kaplan, Q. Le, and P. Molnar, “Maximum likelihood methods for bearings-only target localization,” in *Proceedings of IEEE International Conference on Acoustic, Speech, and Signal Processing*, Salt Lake City, Utah, 2001.
- [5] D. Li, K. Wong, Y. Hu, and A. Sayeed, “Detection, classification and tracking of targets in distributed sensor networks,” *IEEE Signal Processing Magazine*, pp. 17–29, March 2002.
- [6] M. S. Brandstein, J. E. Adcock, and H. F. Silverman, “A closed-form location estimator for use with room environment microphone arrays,” *IEEE Transactions on Speech and Audio Processing*, vol. 5, no. 1, pp. 45–50, 1997.
- [7] R. Murphy, K. Hughes, and E. Noll, “An explicit path planner to facilitate reactive control and terrain preferences,” in *Proceedings of IEEE International Conference on Robotics and Automation*, 1996, pp. 2067–2072.
- [8] C. Behring, M. Bracho, M. Castro, and J. A. Moreno, “An algorithm for robot path planning with cellular automata,” in *Proceedings of the Fourth International Conference on Cellular Automata for Research and Industry*, Karlsruhe, Germany, 2000.
- [9] L. E. Parker, K. Balajee, X. Fu, and Y. Tang, “Heterogeneous mobile sensor net deployment using robot herding and line-of-sight formations,” in *Submitted to IEEE International Symposium on Intelligent Robots and Systems (IROS '03)*, 2003.
- [10] B. Gerkey, R. Vaughan, K. Stoy, and A. Howard, “Most valuable player: A robot device server for distributed control,” in *Proc. of 2001 IEEE/RSJ International Conference on Intelligent Robotics and Systems*, 2001, pp. 1226–1231.