Proceedings of the 2001 IEEE/RSJ
International Conference on Intelligent Robots and Systems
Maui, Hawaii, USA, Oct. 29 – Nov. 03, 2001

Localization Techniques for a Team of Small Robots

Robert Grabowski, Pradeep Khosla
Department of Electrical and Computer Engineering
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213
rjg@andrew.cmu.edu, pkk@ece.cmu.edu

Abstract

Knowledge of position in the context of its surrounding is necessary for robots to build maps and develop path plans. Limitations in odometry and the lack of a priori knowledge reduce the effectiveness of a single robot to retain a sense of position for any extended duration. The problem is only compounded when the scale of the robot is reduced. However, by employing multiple robots we can exploit their distributed nature to provide an external context in which to evaluate sensor readings for mapping and localization. We have designed a team of centimeter-sized robots that coordinate sensing and action to establish and maintain position as they move throughout space. By utilizing low-cost ultrasonic sensors, the team is able to measure the range between each robot pair. We pose these measurements in terms of a position likelihood and combine them to find a global solution that best maximizes the position likelihood of each robot. We also address a unique multi-path interference mode that arises as a direct result of the reduced scale of the robot team. We present our experiences with localization and control of a small robot team.

1. Introduction

One of the most significant skills a robot can master is the ability to localize itself in the world. Knowledge of position and orientation in the context of its surrounding is necessary for avoiding obstacles and developing path plans. Moreover, without knowledge of position, it cannot exploit previous sensor readings or build maps. In an unknown environment, limitations in odometry and the lack of *a priori* knowledge reduce the effectiveness of a single robot to retain a sense of position for any extended duration. High-end solutions, such as video-based landmark detection, place significant burdens on the processing, size, cost and power requirements of a single robot. These restrictions become even more significant in terms of size and power when the scale of the robot is decreased. However, a team of robots can exploit distributed information about each other to regain effective position determination. Teams can share information and functionality as well as provide an external context in which to evaluate local and coordinated sensor readings. To this end, we have designed a team of small robots, called Millibots, which coordinate action and sensing to establish and maintain position as the team moves throughout space[1]. Once equipped with this skill, they are able to accomplish missions such as the mapping and exploration of unknown, hard-to-access spaces.

To achieve the ability to self-localize, we have equipped each robot in the team with a small, low-cost beacon transceiver that allows the robot to determine its range to other members in the group. Being able to measure range allows us to express a robot's position in terms of a likelihood function that correlates estimated robot position to measured range readings. Given multiple range readings, we express the position likelihood of the robot as the product of the multiple, individual likelihoods. We find the most likely position of the robot by finding the maximum

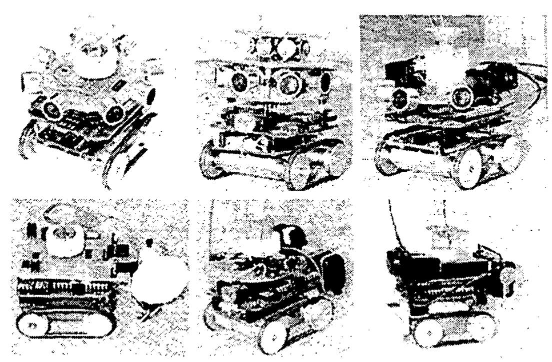

Figure 1. The Millibot Team

[1] see www.contrib.andrew.cmu.edu/~rjg/millibots/ millibot_project.html

of these products. Since we assume the beacon readings exhibit a Normal distribution, we can utilize optimization techniques to reduce the complexity of the computations. By taking the log of a series of probabilities, we reduce the equation from the product of a set of Gaussians to the sum of a series of square errors. In this way we can correlate the maximization of the products of likelihoods, to the minimization of the square error terms of those Gaussians. This reduction also allows us to exploit an optimization algorithm called *Broyden-Fletcher-Goldfarb-Shanno (BFGS)*. The BFGS algorithm finds a set of solutions that minimize the global error of a series of square errors while reducing the chances of becoming trapped in local minima. Tests in simulation have yielded very promising results on the team's ability to localize from an arbitrary geometry and maintain position as it moves throughout space.

Alas, experiments in the real world have exposed the susceptibility of the localization algorithm to outliers in the beacon measurement distribution. This problem is compounded by a unique multi-path interference mode that arises from the small scale of the robot team. Being so close to the ground opens a reflection path that results in distinct regions where interference causes significant error in the beacon measurements. Knowing this, we develop an interference model to adjust the beacon distribution utilized by the likelihood function. This correction alone reduces the sensitivity of the localization algorithm to small outliers. Given the same interference model, we then introduce methods for exploiting the kinematics of each robot to detect and filter for extreme outliers in the beacon measurements making the system more robust. Coupled together, we have been able to regain the effectiveness of the team to localize and maintain its sense of position as the team moves throughout space.

2. *Maximum Likelihood Estimator*

Measurements are obtained by coordinating transmission and reception of beacon signals between robots in the group. Each robot in turn generates and transmits a series of

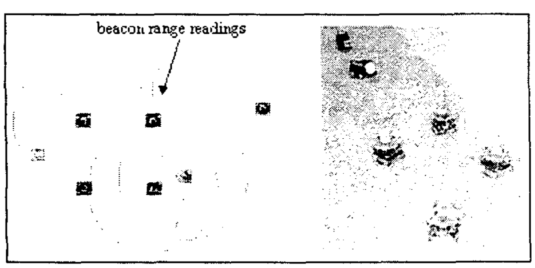

Figure 2. Localization of a Robot

beacon signals which are detected by the remaining robots in the group. By correlating the time between the signals, each receiving robot is able to derive its position with respect to the corresponding transmitting robot. Each robot, in turn, acts as the transmitter in order to collect distance pairs between all the robots in the group.

Visually, a single distance reading between two robots can be perceived as an annulus emanating from the corresponding receiving robot (readings around the end of a circle with radius equal to the measured distance). The rings of the annulus represents all the possible places the transmitting robot could have been in order to generate the measured distance. By adding additional measurements from other robots at different locations, we can constrain the possible positions and localize the robot.

Explicitly, we pose the likelihood that a robot at a given position could generate the measured distance as the probability function:

$$p(d|x_i,x_j,y_i,y_j) = \exp\frac{-[d_{i,j} - \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}]^2}{\sigma^2} \quad (1)$$

where : x_i, y_i is the position of the first robot
x_j, y_j is the position of the second robot
d_{ij} is the measured distance between the two

From here, we can employ a spatial search to find the most likely position of the robot. Since each measurement pair is independent from the others, the most likely position for a given robot is simply the product of all the corresponding individual likelihood probabilities. Maximizing the product of the combined set finds the most likely position of the robot. In practice, we simplify the calculations by taking the log of the total likelihood function for a given robot (the log of the product of a set of Gaussians). The solution reduces from a product of Gaussians to the sum of a set of square errors. Since the log function is monotonic and we are interested only in the global maximization of the probability, finding the minimum of the square errors is equivalent. Reducing the algorithm from a product to a sum also allows us to utilize a variable metric method of optimization called *Broyden-Fletcher-Goldfarb-Shanno (BFGS)*. The BFGS algorithm finds a set of solutions that minimize the global error while reducing the chances of becoming trapped in local minima. It achieves this property by iteratively developing a quadratic estimate of the both the objective function and its derivatives [3]. Iterative estimation has a desired property that we can ensure that even far away from the solution, the algorithm always moves us in the right direction. Closer to the solution, the estimation becomes better and enjoys quadratic convergence.

3. Anomalies in the Beacon Model

Unfortunately, distance measurements are obtained with real sensors that rarely exhibit ideal characteristics. Ultrasound-based sensors suffer from well-known failure modes such as specular reflection and multi-path from multiple objects [1]. For the most part, these failure modes are well understood and can be compensated with conventional filtering techniques such as mean and median averaging and rejection of extreme readings. However, the small size of the Millibots introduces a new systematic failure mode that is not so easy to detect and reduces the teams ability to localize. To understand the significance of the size issue, we must first look at the way range detection is accomplished for the Millibot team and how this mode of sensing is compounded by the scale of the team.

Millibots achieve robot-to-robot ranging by equipping each robot with a ultrasonic beacon transceiver (Figure 3a) and by coordinating sensing between multiple robots. A radio message is generated by the team leader that alerts the members about the beginning of a beacon sequence. One robot is identified as the transmitter and the remaining robots configure as receivers. Upon receiving the start command, the transmitter robot emits a single radio pulse immediately followed by a burst of ultrasonic pulses. Almost instantly, the receiving robots detect the radio pulse and immediately start their local timers. Shortly after, the receiver robots detect the ultrasonic burst and stop their local timers. The elapsed time is proportional to the time-of-flight of the ultrasonic burst.

Ideally, the recorded time measured by each of the receiving robots is proportional to the distance between it and the transmitting robot. A plot of actual distance-vs-measured time should be linear. However, the reduced size of the team places the beacon transmitter very close to the ground - on the order of a few centimeters. A secondary reflection path arises that travels from the transmitting robot, reflects off the ground and continues to the receiving robot (Figure 3a). Because the reflected wavefront must travel a slightly longer distance, it arrives at the receiving detector at a later time and different phase than the wavefront travelling the direct path. The result is the signal detected by the receiving sensor is the superposition of the signal from the direct and indirect path (Figure 3b). When the difference in path length is a multiple of the wavelength the two signals combine to produce a stronger response. However, when the difference in path length is on the order of a half wavelength, the two wavefronts are out of phase and will cancel. The result at these points is that no beacon signal is detected.

Figure 2c shows the distance vs. time plot for a given robot pair. Ten readings were collected at 10cm increments up to a distance of 1.2m. Given a simple interference model, we would expect to see the beacon plot exhibit linear behavior everywhere except where the two signals are sufficiently out-of-phase. At these points, we would expect no event to be detected at all. Indeed, anomalies occur at the predicted distances. However, instead of simply failing to produce a reading at these distinct locations, we see a spread in the measured times (Figure 3c). To understand why the sensor readings produce multiple readings at these nodes, we have to look closer at the detection circuitry.

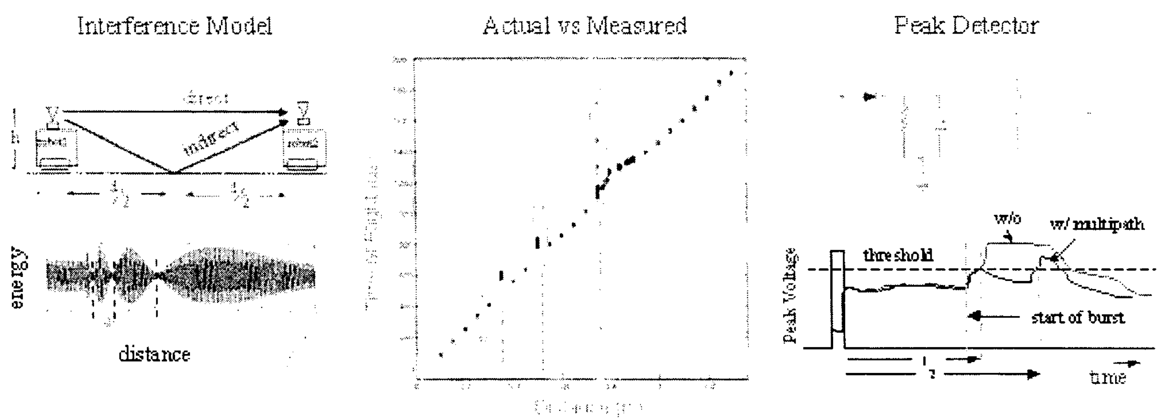

Figure 3. a) Test setup for a beacon measurement (upper left). b) Superposition of direct and indirect bursts as a function of distance (lower left). c) Plot of time-of-flight vs distance (middle). d) Schematic of peak detector (upper right). e) Output of peak detector (showing threshold level) with and without interference (lower right).

Detection of the ultrasonic burst is accomplished using a threshold circuit that triggers when the signal received by the detector exceeds a set amount. Incoming pulses are amplified and converted to a dc voltage proportional to peak amplitude of the ultrasonic pulses (Figure 3d). The sensor element is made from a piezo electric crystal that provides excellent rejection of unwanted sound so only the desired ultrasonic burst is heard. This feature eliminates the need for additional filtering circuitry and minimizes circuit size. However, good rejection comes at the cost of transient response. It takes several ultrasonic pulses to build the dc component above the threshold. Therefore, if the indirect pulses arrive before the direct pulses accumulate above the threshold, the direct and indirect pulses cancel and the dc component will not exceed the threshold. Therefore, a beacon event will not be detected.

Figure 3e shows the dc component of the burst at the output of the peak detector both with and without the multi-path interference. Without interference, the output of the peak detector builds quickly and triggers a detection (Figure 3e trace1). However, in the case of interference (when the indirect path is sufficiently out-of-phase with respect to the direct path), the build up from the direct pulse fails to exceed the threshold before interference commences (Figure 3e trace2). Since a beacon sequence is of finite duration, the direct burst eventually ends. However, the indirect burst, having arrived at a later time, still remains. Now without interference from the direct burst, the dc component of the peak detector begins to build again. The result is a longer detection time than the non-interference case. Experiments shown this effect results in an equivalent measured distances that ranges from 0 to 30 cm above the actual distance between the two. We verified this interference model by performing the same experiment with the robot raised on platforms. In this case, the indirect pulse was several wavelengths delayed from the direct pulse and did not produce interference. In the raised state, no anomalies were detected and the resultant plot generally exhibited linear behavior.

The implications of the interference model are profound. We estimate the position of the robots based on the assumption that we can express their position likelihood by modeling range measurements as Gaussian. We correlate the maximization of the product of these likelihoods to the minimization of their square errors and enjoy a guarantee of quadratic convergence in our search algorithm. However, the multi-path phenomenon breaks the Gaussian assumption and produces outliers in the distance measurement distribution. These outliers cause significant errors in the estimation of position. As a matter of fact, the errors are so significant that even a single outlier can generate such poor estimations that the team becomes lost and cannot recover without completely starting over. Unfortunately, we cannot distinguish between good and bad measurements based on the readings alone.

4. The Motion Model

If we exploit knowledge of the kinematics of the robot, we can develop an independent means for determining the expected position of the robot after any given command. This knowledge can be leveraged in two ways. First, we utilize the position from dead reckoning to test and filter for potential outliers in the distribution of robot-to-robot range readings. Second, we define a closed-form probability distribution for the position estimate itself and apply that as an independent set of constraints for the localization algorithm.

The motion model is based on the tread design of the Millibot team. Millibots employ two long rubber treads mounted on either side of the robot. Encoders provide feedback about both the velocity and number of turns of each of the motors. Many robot groups have developed precise motion models that integrate the incremental data reported from encoders [2]. However, limitations in local processing power and communication bandwidth make the implementation of an iterative motor model infeasible. However, we regain effectiveness of the motion-based approach by limiting robot motion to a series of polar steps. That is, the robot first rotates to the desired angle and then drives straight for the desired distance. By constraining robot motion to a single polar step between localizations, we can describe the uncertainty of the position estimation in a closed form as the product of two Gaussians - one representing rotation and one for translation (equation 2). The distribution of position after any given move is an ellipsoid.

$$\frac{1}{2\pi\sigma_d\sigma_\phi}\exp\left(-\left(\frac{(d-z)^2}{2\sigma_d{}^2}+\frac{(\phi-th)^2}{2\sigma_\phi{}^2}\right)\right) \quad (2)$$

where : d = commanded distance
z = actual distance of travel
θ = commanded rotation
th = actual rotation of robot

We exploit the motion-derived position estimation as an independent test against the individual distance pairs generated by the beacon sensors. We compare the measurement circle generated from a distance pair with the ellipsoid generated by the motor model. If the two curves do not intersect, we assume the distance pair represents an extreme outlier and reject it as an input for localization. Furthermore, since we can describe the ellipsoid in terms of a probability distribution of position, we can use the motion-based position estimation directly as a candidate in the localization algorithm. Simply in terms of simultaneous equations, we have added one more constraint to the derivation of position.

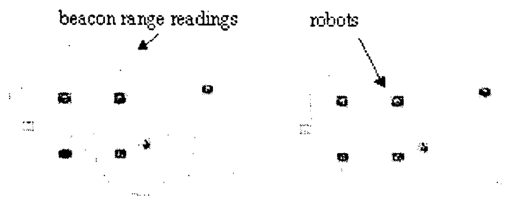

Figure 4: Localization Geometry

Motion-based position estimation has a potentially greater utility for effective localization in the nature of its distribution geometry. Consider the fashion in which range-based positions are derived. In essence the distribution of position is represented by the regions around the cumulative intersection of independent circles obtained via range measurements. If the range measurements are obtained from robots that are well spread out, the angles between the tangents of the intersecting circles are great (Figure 4 left). The corresponding geometry of the distribution is a small, confined region in space. However, when the team is not well spread out with respect to the robot being localized (Figure 4 right), the angles between the tangent measurements are small. The corresponding distribution geometry exhibits a largely distributed ellipsoid. Position estimation is then sensitive to small perturbations in the range measurements and limited precision in the search algorithm. On the other hand, the geometry of the motion-based derivation is independent of the geometry of the team. By adding the constraint of the motion command to the derivation, we can reduce the geometry of uncertainty even when the robot is not in the most idea position. This becomes especially important during exploration where the moving robots may need to move along the outside of the group. Motion-based position derivations allow the team to relax its dependence on team geometry while still being capable of localizing.

5. *Reducing Sensitivity to Outliers*

The addition of a motor model reduces the changes of an extreme outlier from completely disrupting the estimation of position. However, this correction only minimizes extreme outliers but does not eliminate them. Unfortunately, the BFGS in its current form is very sensitive to outliers. However, if we adjust the objective function to account for the anomaly, we can regain the effectiveness of the localization algorithm.

To achieve the new objective function, we adjust the model of probability to account for the small likelihood that a robot is actually closer than implied by its distance measurement (the returned distance is equal to or greater than the real distance). Now rather than relating than relating the Gaussian distribution to the sum of the square errors, we explicitly take the log of the likelihood function to produce the cumulative error term. Computation is more expensive but the system is less susceptible to outliers. Experiments show the addition of a uniform position likelihood reduces the sensitivity to outliers in the localization process. The team is able to localize even with the occasional outlier produced during distance measurements.

$$p(r|d) = \alpha\left[\frac{1}{\sqrt{2\pi\sigma}}\exp\left(\frac{(r-d)^2}{2\sigma^2}\right)-1\right] \qquad (3)$$

where : $\alpha = 0.9$ (10% chance of robot being anywhere)
r = range reading of robot
d = actual distance between robots

6. *Applications*

Providing techniques for establishing and maintaining position, gives the team a context in which to share range and sonar information. Figure 5 shows the composite map generated by a team of four robots exploring a small space both with and without global localization. The yellow lines represent the true objects and were added by hand only for verification. In the first image, we give the team information about the starting points of each robot. Beyond that, the team must rely on independent odometry to determine robot positions. By the time the robot team reaches the bottom corridor, odometry error has accumulated sufficiently enough that the team no longer has a viable context in which to share information. Furthermore, without knowing where it is, the team, as a unit, cannot agree on what it sees or coordinate to determine where to go. In the second image, the team has been placed in an arbitrary starting configuration. By applying the localization algorithm, it has been able to establish a common reference frame in which to share data. By coordinating sensing and action, they are able to compose a consistent map that allows them to continue operating and exploring the world.

Many researchers have proposed solution to establishing and maintaining robot position using various high-end

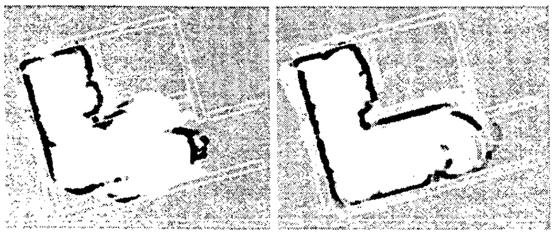

Figure 5. Team Maps

techniques such as employing high accuracy laser sensors [8], utilizing high-bandwidth, video processing [9] or matching sensor profiles to pre-defined maps [4]. While these systems enjoy potentially greater accuracy, speed and range, they are difficult to employ on smaller robots. We have developed a low-complexity alternative to these methods that is capable of localizing a team by measuring the distances between individual robots. Moreover, the size of the detector allows it to be employed by small robots with limited local processing.

7. Acknowledgements

The authors would like to thank the current members of the Millibot team for their contributions, particularly. Luis Navarro-Serment, Curt Bererton, Ben Brown and Mike Vandeweg. Specifically we would like to recognize Luis Navarro-Serment who designed and developed the beacon sensors and Dr. Christiaan Paredis who implemented the initial optimization routines. Finally I would like to thank Dr. Howie Choset for his many hours of discussion and guidance. This research is funded in part by the by the Distributed Robotics program of DARPA/ETO under contract DABT63-97-1-0003.

REFERENCES

[1] Borenstein, J., Everett, H. R., and Feng, L., 1996. *Navigating Mobile* Robots*: Sensors and Techniques*, Wellesley, MA: A. K. Peters, Ltd.

[2] Chong, K.S. and Kleeman, L. "Accurate Odometry and Error Modeling for a Mobile Robot", Proceedings 1997 IEEE International Conference on Robotics and Automation.

[3] Fletcher, R. 1987, Practical methods of optimization (second edition). New York, NJ: J. Wiley & Sons, Ltd

[4] Fox, D., Burgard, W., Kruppa, H., Thrun, S., "Collaborative Multi-Robot Localization", *Proceedings of the German Conference on Artificicial Intelligence*, Germany 2000.

[5] Grabowski, R., Navarro-Serment, L.E., Paredis, C.J.J. and Khosla, P.K., 2000. Heterogeneous Teams of Modular Robots for Mapping and Exploration. *Autonomous Robots. (special issue on heterogeneous and distributed robotics).* Volume 8, No. 3, June 2000. pp. 293–308.

[6] Navarro-Serment, L.E, Paredis, C.J.J., and Khosla, P.K. 1999. "A Beacon System for the Localization of Distributed Robotic Teams," In *Proceedings of the International Conference on Field and Service Robotics*, Pittsburgh, PA, August 29-31, 1999.

[7] Rekleitis, Y., Dudek, G., "Multi-Robot Collab-oration for Robust Exploration", *Proceedings of IEEE International Conference in Robotics and Automation*, San Francisco, CA, April 2000.

[8] Simmons, R., Apfelbaum, D., Burgard, W., Fox, D., Moors, M., Thrun, S., Younes, H., "Coordination for Multi-Robot Exploration and Mapping", *Proceedings of the International Conference on Artificicial Intelligence*, Austin Tx, August 2000.