The Arc-Transversal Median Algorithm: an Approach to Increasing Ultrasonic Sensor Accuracy

Keiji Nagatani — Mechanical Engineering
Howie Choset — Mechanical Engineering
Nicole Lazar — Statistics and CALD
Carnegie Mellon University

Abstract

This paper describes a new method for determining range information about a robot's surroundings using low resolution ultrasonic sensors. These sensors emit ultra-sound which bounces off of nearby objects and returns to the sensor. The time-of-flight for the sound to return to the sensor is the distance between the sensor and the object. A sonar arc represents the possible locations of the object. We model these locations with a simple uniform probability distribution on the sonar arc. We then introduce a new method to fuse sonar data to determine the actual obstacle location. This new method is termed the Arc-Transversal Median method because the robot determines the location of an object by intersecting one arc with other arcs whose angle-of-intersection exceeds a threshold and then taking the median of the intersection. The median is a robust estimator that is insensitive to noise because a few stray readings will not affect the median. We show via some simple geometric relationships, that this method can improve the accuracy of the sonar sensor by a specified amount, when certain assumptions were in place. Finally, experimental results on a real mobile robot verify this approach.

1 Introduction

Our task is exploration of an unknown environment using a robot equipped with ultrasonic sensors. In our work, the robot enters an unknown environment and, relying solely on sensor information, it builds up a map of that environment which can be used for future excursions. We use a standard mobile robot with sixteen sonar sensors. In the course of implementing our mapping method, we noticed that the low azimuth resolution of sonar sensors causes problems, especially at narrow openings. Effectively, the wide beam patterns and low resolution information create the illusion that narrow passageways do not exist.

Using a better sensor would automatically solve this problem. One such sensor is a laser range system which provides high resolution information, both in distance and in azimuth. Laser range systems providing full 360^o coverage are quite costly, sometimes more than the mobile base itself. Moreover, the laser sensor modality does not work in all environments, such as those with glass obstacles.

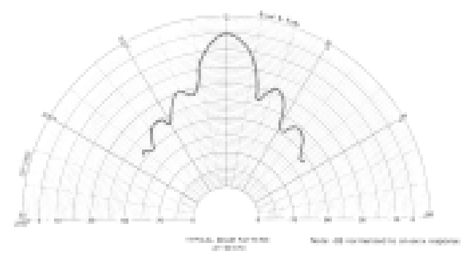

Fig. 1. Beam pattern for the Polaroid transducer installed on many mobile robots.

In this work, we increase the effective resolution of the sonar sensors by appropriately moving the robot and fusing previous sonar data over time. The underlying math of this procedure specifies the improvement in resolution. First, we consider a naive sensor model and show why it does not work. Next, starting with the naive model, we introduce our procedure, termed the *Arc Transversal Median* (ATM) method. We then give some experimental results demonstrating the utility of the ATM method, including an example of a robot exploring an unknown environment.

2 Prior Work

Mobile robots use sonar sensors to navigate and explore in their environments. Conventional sonar sensors measure distance using time of flight. Under nominal conditions, the speed of sound in air is constant, so the time sound requires to leave and return to the transducer is proportional to the distance to the object. In actuality, it is proportion to the distance to the point of reflection on the object [3]. This object, however, can be located anywhere along the perimeter of the sonar sensor's beam pattern (Figure 1). Therefore, the distance information that sonars provide is fairly accurate in depth, but not in azimuth. Nominally, prior work has considered three methods to process sonar data: storing echos in certainty grids, tracking sonar arcs, and processing the entire echo signal. See [6] for an excellent overview of sonar sensor use.

Alberto Elfes is among the early researchers to use the occupancy grid method in [4]. An occupancy grid

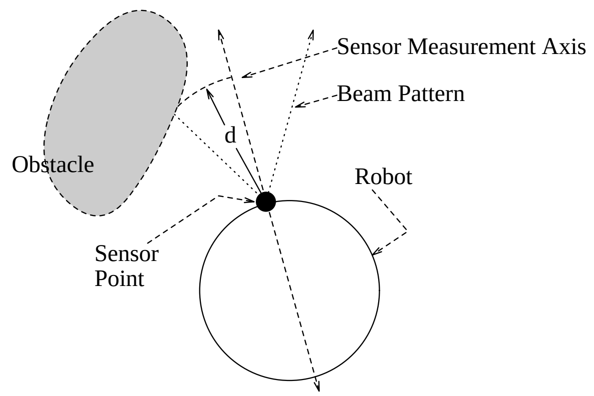

Fig. 2. Simplified Distance Measurement Sensor Model

is a planar discrete representation of the robot's environment where the value of the cell or pixel represents its occupancy status along with a certainty value: 0, unknown; $[-1, 0)$, empty; and occupied, $(0, 1]$.

Elfes assumes a sonar sensor model reflecting the probability that all points in the range of the sensor are occupied by an object. The model is essentially a planar cone with an arc-base, as opposed to a straight line segment. The height of the cone is the distance at which the sonar sensor detects an obstacle.

This model essentially places a Gaussian distribution centered at the arc's midpoint reflecting the likelihood that an obstacle is located along the arc. The Gaussian assigns a higher likelihood that the object is located in the center of the arc. Since the likelihood of an obstacle being located in the interior of the arc is quite low, the model assigns a -1 for almost all points in the interior. Points on the "edge" of the cone, between the interior and arc, are assigned intermediate values.

Finally, this model is then discretized into cells, also, to match the grid representation of the world. Initially, all cells in the world map have a zero value, corresponding to unknown. As the robot acquires new sensor data, the cells of the world map are updated using Bayes Rule. Other work such as [9] use a similar grid-based and updating approach, but instead using a uniform distribution along the arc, as opposed to a Gaussian.

Leonard stepped away from the conventional pixel based approach by identifying corners and edges in a polygonal environment using ultrasonic sensors with the range of constant depth method (RCD) [7]. Essentially, RCD assumes the point of reflection is the center of the beam pattern (See Fig. 2). The RCD assumption is reasonable in [7] because it uses transducers whose associated beam patterns have a main lobe of 1.65^o.

The RCD approach fuses several sonar readings as the robot moves through its environment. As the robot moves through space, these center points are tracked. If they continuously move in a straight line, the RCD hypothesizes that the robot is tracking a wall. If the center points remain relatively fixed as the robot moves, then the RCD hypothesizes that the robot is tracking a corner. Upon first inspection, the ATM approach seems similar to RCD, but in actuality, they can compliment each other quite well.

Professor Kuc, on the other hand, developed a biomimetic sonar system that relies on the quality of echo signal, in addition to the time-of-flight [5]. His work processes the echo signal to perform object recognition. Initially, three sonar sensors are used: a center one to transmit a signal and two later (left and right) ones to receive the signal. Based on the time-of-flight of the return echos, the lateral sensors rotate inward toward an obstacle so that the echo reflecting from the object is normal to the lateral sensors. This effective focuses the lateral sensors on the object. Once focussed, this system processes the signal and then discretized the result into 16-tuple vector. This vector represents the geometry of the object and thus can identify it. Professor Kuc's results were successful at distinguishing among machine washers, ball bearings, O-rings, and paper clips. This work does not address the issue of mapping, but the authors feel his work can augment localization work currently underway [8].

3 Sensor Model

Many mobile robots have sonar sensors that are rigidly attached to their perimeter, pointing radially outward from the robot. Sound is emitted from the sonar sensors, bounces off an obstacle, and returns to the robot. The time of flight is proportional to the distance between the object and the sensor. This object, nominally , can be located anywhere along the sensor arc. The size of this arc varies with different sonar sensors, but is 22.5^o for the Polaroid ultrasonic sensor that is found on most robots.

3.1 Centerline Model Breakdown

Initially, we used the centerline model which produces a reasonable estimation for the actual location of objects if the arc of the cone is small, as was the case in [7]. Likewise, the centerline model also produces a reasonable estimation for the actual location of objects that are close to the sonar sensor.

Using the standard Polaroid transducer, the centerline model breaks down quickly in approximating the location of obstacles far from the robot. In Figure 3, the robot passes by an opening which is wider than the diameter of the robot. In other words, the robot can pass through this opening. The sonar cone from the downward pointing sensor is drawn at the two robot locations, each depicted by a circle. The solid dot in the center of the arc represents the robot's perception of

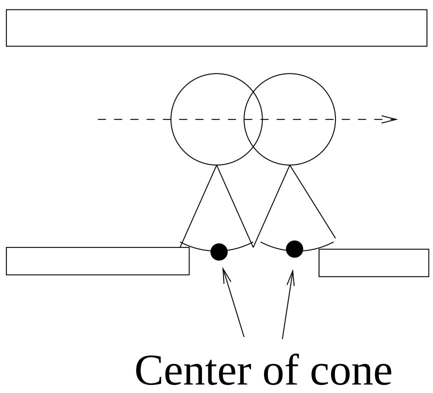

Fig. 3. Center of cone model makes the passageway seem narrower

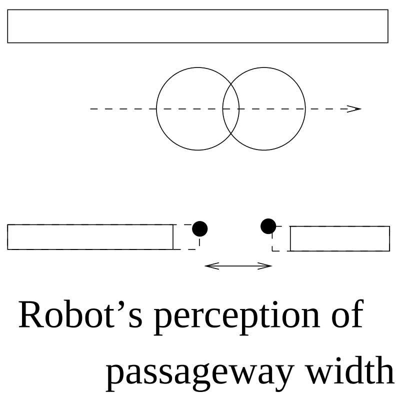

Fig. 4. Center of cone makes the passageway seem narrower.

where obstacles are located. Based on this perception, the robot believes the passageway is too narrow for it to pass through (Figure 4).

Note that, using the same reasoning as above, simply rotating the robot by half a sonar cone width and interleaving the additional center-cone measurements will *not* double the resolution of the sonar sensors and will also give a similar false impression that corridors are more narrow than they are in actuality. See Figure 5. If the sonar cones had also halved in width, then the rotate-in-place method would have worked, in theory.

A common method for detecting walls is to fit a line through cones of constant depth. This method also does not work in the presence of narrow openings using wide angled sonar detectors. See Figure 6.

Figure 7 contains an environment with a narrow opening, similar to the examples above. In this experiment, the Nomad 200 Mobile Robot used the naive centerline model to sample the location of objects. The dots represent the centers of the sonar arcs. This method gives the robot a false impression of the world because the robot perceives there *not* to be an opening.

Others have studied the problem of negotiating narrow door-ways, but Schultz first pointed out this problem to the authors [11], [10]. Also, the architecture pre-

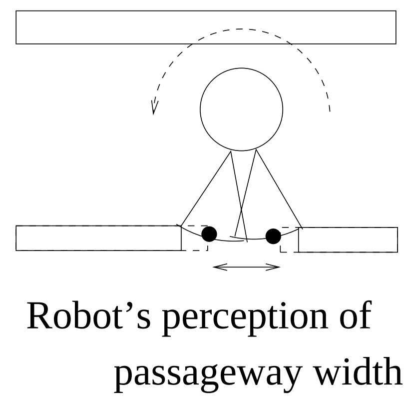

Fig. 5. Center of cone makes the passageway seem narrower

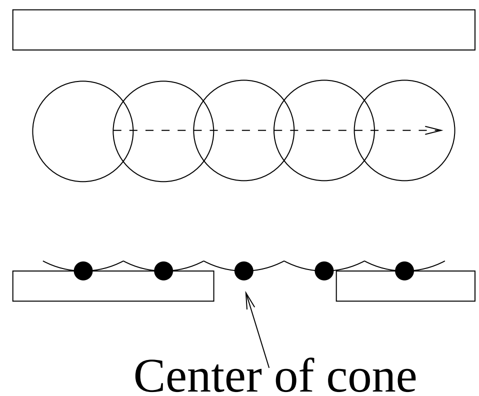

Fig. 6. Passing a line through the centers of the cones gives the robot the impression there is one wall.

sented in [1] handles a robot negotiating through a narrow doorway.

3.2 Uniform Distribution Model

Simply assuming that the echo comes from one fixed location on a sonar arc is not sufficient. Instead, the entire arc must be considered. There has been considerable success in using a Gaussian distribution to model the location of an object along a sonar arc; the center of the distribution is taken to be the centerline of the arc [4].

According to the documentation [3] of the widely-used Polaroid ultrasonic sensors, there should be a uniform probability that an echo, and hence the reflection point of an object, will be anywhere along the arc of the sonar cone. This is model we use in the ATM approach and we have experimentally verified this beam pattern. Other researchers such as [9] also use a uniform distribution model as well. Certainly, this means that we cannot assume that the object lies along the centerline, unless the sensor reading is quite small. We have to consider the arcs themselves. We assume that each sonar arc corresponds to only one point of reflection, and hence one obstacle.

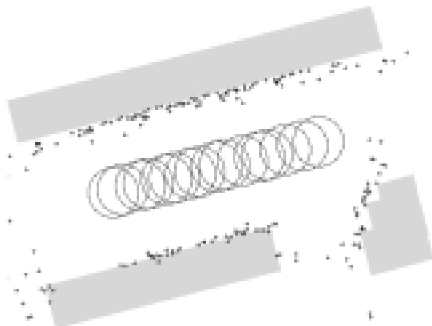

Fig. 7. Centerline method was used. The points correspond to the center of arcs from real sonar data and the light grey obstacles were drawn for the sake of display. The actual obstacles were walls in the lab.

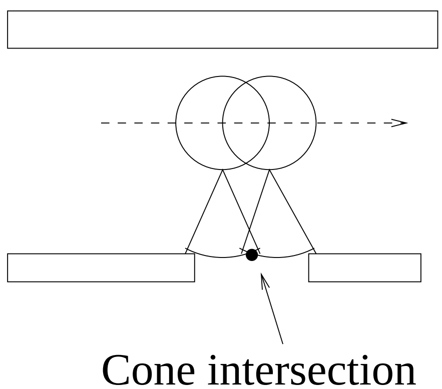

Fig. 8. Intersection of cones implies that there is an object in the middle of the opening.

4 Arc Transversal Median Method

Using the uniform distribution model from the previous section, we explain one more naive method to infer obstacle location, and build up our approach from there. Once we explain our procedure, we will demonstrate some experimental results, and then in the following section explicitly derive the ATM method itself.

4.1 Arc Intersections

Initially, we considered arc intersections of two arcs, labeled arc1 and arc2. The point of reflection can lie anywhere on arc1; likewise, this point can lie anywhere on arc2. If these two arcs intersect, then possibly the point of reflection is more likely to lie at the intersection. However, considering single intersections gives the robot a false impression about the locations of objects. In Figure 8, the robot receives individual echoes from the two lower objects. The two corresponding sonar cones intersect in the middle of the opening, and thus using all intersections does not work.

The ATM method considers several intersections on an arc. If many sonar sensor arcs all intersect at one point, the probability of there *not* being a point of reflection from an object is low, as described in Section 5.1. In actuality, due to sensor resolution in distance and slight error in dead-reckoning, many arcs may not intersect exactly at one point even if the source of the echo is constant. Instead, their intersections will form a cluster on one arc corresponding to the same point on an object. Any element of the cluster serves as an excellent approximation to the exact location of the object along the arc.

4.2 Median of Intersections

As a candidate element of a cluster, we use the median of all intersections on one arc. Using the median bypasses the need to employ an explicit clustering routine. Furthermore, the median is robust with respect to noise; it automatically eliminates bad sonar echoes and spurious intersections that correspond to other objects. These would manifest themselves in our data as outlying readings which we want to ignore.

If an arc has no intersections, however, then the robot uses the center of the arc as the location of the echo. When the robot has three or more intersections, the median operation can then applied. In the scenario where there is an even number of intersection, the robot uses the mean of the two medial values. That is, the median of $\{1, 4, 6, 10\}$ is 5.

4.3 Transversal Intersections

Finally, we do not consider all intersections, just those that "stably" intersect. Two cones stably or *transversally* intersect, if their intersection does not significantly change after one of the sets is slightly perturbed. The two cones in the upper portion of Figure 9 transversally intersect because if one cone were slightly perturbed, then the intersection would not significantly change. On the other hand, the two lower cones do not transversally intersect because if one cone were slightly perturbed, then the location of the intersection would change. We only consider transversal intersections when computing the median. (This is a similar criterion to stereo vision or structure from motion.)

Since we consider the median of only the transversal intersections, we termed our approach the *arc-transversal median method* (ATM method). This median gives a more accurate location of the sonar echo from the object.

4.4 Experimental Verification

We now apply the ATM method to the environment that was described in Figure 7. Recall that the centerline approach effectively inflated objects giving the robot the false impression that there was no opening present. First consider Figure 10 which displays all of the sonar arcs and all their respective transversal inter-

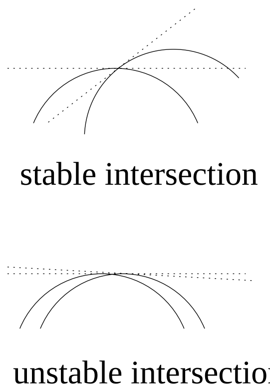

Fig. 9. Intersection of cones implies that there is an object in the middle of the opening

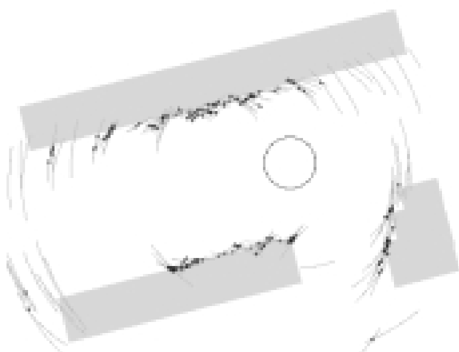

Fig. 10. All arcs corresponding to sonar readings when the robot drives down the corridor passing by the opening. The arcs correspond to real sonar sensor data and the light grey obstacles were drawn for the sake of display. The actual obstacles were walls in the lab.

sections denoted by dots. For this experiment, we used intersections whose tangents are thirty-degrees or more. The choice of thirty-degrees is derived in Section 5.2. The dots in Figure 11 represent the medians of each arc's intersection points. These points represent a more accurate view of the environment, most notably at the entrance way, which was not apparent in Figure 7.

5 Derivation of the ATM

Why does this method work so well? We first give a probabilistic argument based on the properties of the uniform distribution, to confirm our intuitive sense that areas of many intersections should correspond to locations of objects. The choice of thirty degrees for the transversal angle is then derived mathematically.

5.1 Probabilistic Reasoning: The Role of the Uniform Distribution

Given a uniform distribution on an interval of length Θ, the probability of a point landing at random within a subinterval $(d, d+\Delta)$ of length Δ is $\frac{\Delta}{\Theta}$. Clearly $\frac{\Delta}{\Theta} < 1$. This probability only depends on the length of the subinterval, not its location. If we now consider the probability of n points falling into this same small interval, when no object is present, this probability can be written as

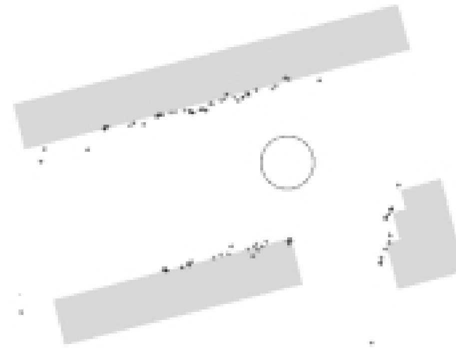

Fig. 11. Median method was used. Gray area denotes the location of objects and the plus marks represent the medians of transversal intersections.

$$\Pi_{i=1}^{n} P(X_i \in (d, d+\Delta)) = \Pi_{i=1}^{n} \frac{\Delta}{\Theta} \qquad (1)$$

which is equal to $\frac{\Delta}{\Theta}^n$. As n increases (or Δ decreases), this probability approaches zero. In other words, as more readings fall into the same small interval, the more we believe that there is an object that is being detected by the sonar.

In terms of statistical hypothesis testing, we have tested the null hypothesis that there is no object present against the alternative that there is an object. If the probability of the event is low, i.e., there are many intersections, then we reject the null hypothesis in favor of the alternative.

5.2 Transversal Intersections: Why we used 30^o in our experiments

In this section, we discuss our choice of using thirty degrees as the threshold for transversal intersections. For the following calculation, we assume that each sonar arc only measures distance to one obstacle. This assumption is quite reasonable and we have not encountered a configuration of objects in our experiments where this assumption was not upheld.

The sonar sensors on the Nomad robot produce range readings with a one inch resolution, and thus we use inches, instead of centimeters. In Figure 12, a sonar sensor detects an object d inches away from the robot. Since we assume a uniform distribution for possible obstacle locations along the arc, this object is equally likely to lie anywhere along an arc of 22.5^o degrees (i.e., $\phi = 22.5^o$). Furthermore, our range resolution is one inch, and thus the object can lie in a one inch band-arc of 22.5^o. The length of this arc is $d\frac{\pi}{180}22.5$ inches.

Since we are only considering objects whose distance

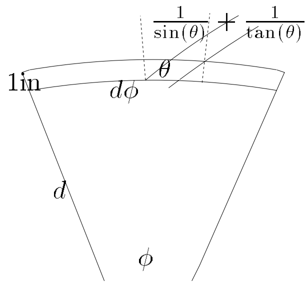

Fig. 12. The intersection of two cones is approximated by a trapazoid.

is significantly greater than one inch away from the robot, the sonar band can be viewed as a 1 inch by $d\frac{\pi}{180}22.5$ inches rectangular strip. Along these lines, the intersection of two sonar cones can then be approximated by a rhombus.

The length of the major axis of the rhombus is an upper bound to the distance between two echoes that lie in the intersection of the two bands. To determine the azimuth range of object location for a particular sensor, we project the rhombus onto the sonar cone's arc d inches away from the sensor. The length of the projected rhombus along the arc is $\frac{1}{\sin(\theta)} + \frac{1}{\tan(\theta)}$.

Therefore, when considering intersections of arcs whose angle is θ, the value $\frac{1}{\sin(\theta)} + \frac{1}{\tan(\theta)}$ is the length of an interval along the arc where an object could be located. Note that this length does not vary with distance d. For example, if $\theta = 30$ degrees, then the sonar cone has an accuracy of 3.73 inches in azimuth, regardless of obstacle distance. In other words, an object can be located in an interval on the cone centered at the intersection of length 3.73 inches.

At a distance d inches for a given θ, the new resolution can be computed by

$$d\phi = n\left(\frac{1}{\sin(\theta)} + \frac{1}{\tan(\theta)}\right). \qquad (2)$$

For $d = 100$ inches, $\phi = 22.5$ degrees, and $\theta = 30$ degrees, $n = 10.5$. This means that at a distance of 100 inches from the robot, we have at least a 10.5-fold improvement in resolution.

If d is small, then this range is probably longer than the arc itself, and in such situations, we use the standard center of the cone approach since that suitably approximates the location of the object.

The above analysis implies that as d increases, the resolution along the arc of the sonar sensor seemingly improves. However, as d increases, the accuracy of the sonar range does not necessarily remain fixed at $\pm.5$ inch. Since this accuracy changes over large distances, the 1 inch accuracy in range can also become a parameter of the model.

5.3 Median

In actuality, we are projecting a number of trapezoids onto the arc under consideration. The trapezoids project onto intervals, not point intersections. As stated above, the maximum length of these intervals correspond to the minimum resolution provided by the ATM method. So, in essence, we have to take the median of intervals, which is not clearly defined. If all intervals were the same length, then we could easily take the median of the mid-points of all intervals to approximate the location of the echo. However, the desired resolution (i.e., the threshold for transversal intersections) limits the variation in the size in the intervals, so taking the median of the mid-points is still a meaningful calculation.

5.4 Discussion

Detecting Flat Walls. This method does not segment the environment into corners and walls; it simply looks for echos that comprise the boundary of the robot's free space. However, it is worth pointing out how well it detects points on corners and flat walls. The ATM method detects corners quite well because corners have many opportunities to receive echos from a variety of directions.

A specular flat wall, on the other hand, can only receive an echo when the sonar beam is normal to the wall. In such cases, the ATM method reverts back to the centerline approach when passing along a wall. In other words, the robot will use the center of the cone to approximate the location of a point on the wall.

Points on non-specular walls can receive echos from several angles, allowing for multiple transversal intersections to form as the robot drives along the wall. In a sense, the wall comprises little bumps or corners that the sonar sensors can detect. See Figure 13. In our experiments, both inside the lab and in the halls outside of our lab, we found the walls to be sufficiently non-specular to allow for multiple arcs to intersect at several points along the wall (see Figure 10).

The ATM method may give the robot the impression that non-specular walls are closer to the robot, but only by a small amount, on the order of the depth of the roughness of the wall surface. Certainly, this amount is less than one inch, the depth resolution of the sonar sensors used in our experiments.

Limitations on accuracy improvement. According to Equation 2, if we increase θ, then the ATM method provides higher resolution information. Practically, there is a limit on θ. We cannot simply increase θ to achieve any desired resolution since a higher θ threshold would result in too few intersections. With a smaller number of intersections, we have less information avail-

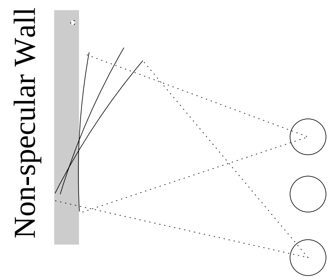

Fig. 13. As the robot passes by one point on the wall, it can collect enough arcs that intersect near the same point on the wall.

able about the environment. Since a 10-fold improvement in resolution for our experiments provided excellent results, we did not need to increase θ beyond 30 degrees for improved resolution.

Length of sonar history. When computing the intersections of the sonar cones, we only use the "recent" sonar sensor readings because of localization error. Keep in mind that the arc locations are with respect to the robot's encoder coordinates. Sonar data that was acquired over a short period of time is self-consistent because the robot accrues minimal localization error, and thus sonar readings "near" each other in time are comparable. However, sonar readings that are acquired at significantly different times are not comparable because one arc's position has significant accumulated localization error with respect to the other.

Although the authors' main interests are exploration and mapping of unknown static environments, using only the "recent" sensor readings makes sense in dynamic environments where readings made in the past may not reflect the true current location of obstacles.

In our experiments, we used the ten most recent sonar sensor readings which is derived from the desired resolution. The desired improvement in resolution when detecting objects as far away as 100 inches is 10-fold. This leads to the 30 degree transversal cut-off. Therefore, for the robot to obtain an intersection at the same point 100 away from the robot, it must move roughly $len = 50$ inches, i.e.

$$\frac{\text{len}}{2} = \sin(\frac{30}{2}) \times 100 \text{ inches} = 25.8 \text{ inches}.$$

Since we use a 1 second update rate for the sonar sensors and drove the robot at $5in/sec$, the robot requires 10 updates to move 50 inches, thereby guaranteeing it will get at least one intersection.

Robot dynamics. The location and the transversality of the intersections also depend upon the speed of the robot and the bandwidth of the sensor data. Currently, the robot moves at the rate of four inches per second and reads in all sixteen sonar sensor readings approximately once a second. In our implementation, the time to acquire sonar readings currently limits the robot's speed. The number of prior readings, robot speed, and sensor throughput are all parameters that future work will integrate into the ATM method.

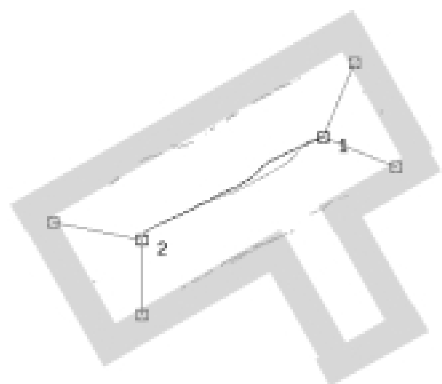

Fig. 14. Incrementally constructing the GVG using naive centerline approach.

6 ATM Application Example: Mapping Unknown Environments

Virtually any mapping procedure, such as [7], can use the ATM method. In this work, because of the authors' prior experience, we applied the ATM method to a generalized Voronoi graph (GVG) mapping approach [2], [8]. Recall that the GVG in the plane is the set of points equidistant to two objects. In other words, it is the set of points $\{x : d_i(x) = d_j(x)\}$, where $d_i(x)$ is the distance between a point x and an object C_i, i.e.,

$$d_i(x) = \min_{c_o \in C_i} \|c_o - x\|.$$

Since the robot can use the GVG to plan paths between any two points in the environment, exploring an unknown environment is reduced to incrementally constructing the GVG using the robot's sensors.

We conducted two experiments in the same T-shaped environment with a narrow corridor extending from the base of the T. In the first experiment, the robot used the center of the cone model when incrementally constructing the GVG and missed the opening. The gray regions in Figure 14 correspond to the object locations, which are not a priori known to the robot. The dots correspond to the location of the center of the arcs, and thus the robot's understanding of obstacle locations. Notice how there are dots in the mouth of the opening, giving the robot the impression the opening is too narrow for it pass. Therefore, the robot did not exhaustively explore the environment and produced an incorrect GVG, represented by solid lines.

Using the ATM method (Figure 15), the robot quite easily detected the opening and explored it. The dots,

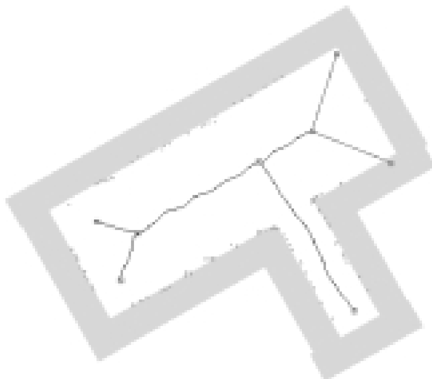

Fig. 15. Incrementally constructing the GVG using the ATM approach.

here, represent the medians of the transversal intersections of all of the arcs, and thus the robot's understanding of the obstacle locations. The robot was successful because the ATM method gave the robot more accurate information about its surroundings.

7 Conclusion

Sensor based exploration with mobile robots of unknown environments has motivated the work described in this paper. We addressed the issue of uncertainty of sonar sensors in azimuth. Using simple sonar models, this limitation in accuracy prevents the robot from fully exploring an environment because the robot does not detect narrow openings. In this paper, we address this problem by use a uniform distribution model and develop a method of processing sonar data to improve sonar sensor accuracy. This method is called the Arc-Transversal Median (ATM) method.

Experiments with a mobile robot mapping an unknown environment with a narrow opening verify the strength of this method. However, the ATM method can still be further refined. We need to determine the tradeoff between information gain and accuracy. Furthermore, the ATM method assumes that *most* readings are true. The median filter, by iteself, takes care of spurious intersections on an arc, but there is a problem when several false readings intersect to form a "ghost" obstacle. Specular reflections and multi-pass echos are both forms of false readings that can give rise to these ghost obstacles. In our experiments, the ghost obstacles formed far away from the robot, but on several occassions, the updating process of our local map automatically deleted them as the robot approached the vicinity of the false obstacles. Future work will consider this problem more carefully.

Acknowledgments. The authors would like to thank Dr. Teresa McMullen at ONR, Grant 97PR06977 and Dr. Howard Moraff, Dr. Jing Xiao, Dr. Larry Reeker, and Dr. Ephraim P. Glinert at NSF, Grant IRI-9702768 for supporting this work, as well as the NSF CISE Post-doctoral Program Grant 98-08271 for supporting Dr. Nagatani. We would also like to thank Teruko "Tel-Tel" Yata in Professor Yuta's lab at the University of Tsukuba for her insight into this work. We would like to thank Dr. Steve Seitz in the Robotics Institute and Jong Chul Park in Mechanical Engineering at Carnegie Mellon. Finally, we would like to send many thanks the anonymous reviewers for their constructive suggestions to our paper.

References

[1] J. Budenske and M. Gini. Why is it so difficult for a robot to pass through a doorway using ultrasonic sensors? In *IEEE International Conference on Robotics and Automation*, San Diego, CA, May 1994.

[2] H. Choset and J.W. Burdick. Sensor Based Planning, Part I: The Generalized Voronoi Graph. In *Proc. IEEE Int. Conf. on Robotics and Automation*, Nagoya, Japan, 1995.

[3] S. Ciarcia. An Ultrasonic Ranging System. *Byte Magazine*, pages 113–123, October 1984.

[4] A. Elfes. Sonar-based real-world mapping and navidation. *IEEE Journal of Robotics and Automation*, RA-3:249–265, June 1987.

[5] R. Kuc. Biomimetic soanr locates and recognizes objects. *IEEE Journal of Oceanic Engineering*, 22:616–624, October 1997.

[6] J. J. Leonard and H.F. Durrant-Whyte. *Directed Sonar Sensing for Mobile Robot Navigation*. Kluwer Academic Publishers, Norwell, MA, 1992.

[7] J. J. Leonard, B. A. Moran, I.J. Cox, and M. L. Miller. Underwater Sonar Data Fusion Using an Efficient Multiple Hypothesis Algorithm. In *IEEE Int. Conf. on Robotics and Automation*, Nagoya, Japan, May 1995.

[8] K. Nagatani, H. Choset, and S. Thrun. Towards Exact Localization without Explicit Localization with the Generalized Voronoi Graph. In *IEEE Int. Conf. on Robotics and Automation*, Lueven, Belgium, May 1998.

[9] D. Pagac, E.M. Nebot, and H. Durrant-Whyte. An Evidential Approach to Map-Building for Autonamous Vehicles. *IEEE Transactions on Robotics and Automation*, 14:623–629, August 1998.

[10] A. Schultz, Yamauchi B., and Adams W. Integrating map learning, localization and planning in a mobile robot. In *Proc. of the 1998 Conference on Computational Intelligence in Robotics and Automation*, pages 331–336, Gaithersburg, MD, Sept. 1998.

[11] Alan C. Schultz and John J. Grefenstette. Using a genetic algorithm to learn behaviors for autonomous vehicles. In *Proceedings of the American Institute of Aeronautics and Astronautics Guidance, Navigation and Control Conference (AIAA)*, pages 739–749, August 1992.