Hiding the Acoustic Signature of a Mobile Robot

E. Martinson

***Abstract*—A mobile robot can be a rather significant source of noise: noisy fans cool onboard computers, motors are spinning, rubber wheels are squeaking against the floor, and mechanical parts are grinding against each other. Despite these noise problems, robots are often suggested as ideal security and surveillance tools, since they can reliably patrol an area, and remove people from harms way. While the arguments for using a robot are persuasive, the use of a noisy robot for these tasks is questionable. Certainly, they might allow a human to be separated from immediate harm, but given the ease with which people and/or machines can detect changes in the ambient noise, a noisy robot might not be very effective at performing its duties. One solution is to make very small, very quiet robots. But such small robots do not usually do very much either, as their sensing and computational power is minimal. An alternative solution is to use a larger robot, but make the robot aware of its own acoustic signature. Combined with knowledge about sound sources and sound flow through the environment, an acoustically aware robot may hide its own acoustic signature in the ambient noise to reduce the risk of being detected.**

I. Introduction

The idea of a robotic security guard is not a new concept [1, 2]. By using a robot to patrol the hallways of an office building, or the perimeter of a military base, designers are hoping to reduce costs, improve surveillance efficiency, and remove some of the danger from human security guards. Outside the research community, there is also a large number of private firms working on similar concepts, so this field is very clearly of interest to the community at large. In most cases, however, the robotic platform being used is not a small, unobtrusive robot. Instead, the robot base is fairly large for the sake of robustness. As such, these robots also tend to be moderately noisy because of onboard cooling fans and motors designed to move heavier equipment. How can such a robot be used to quietly observe, or approach a target? We believe that the solution to this problem lies in making a robot aware of the surrounding auditory scene. By knowing something about the listener, the environment, the sound sources, and the physical principles that govern how they each affect sound flow, a robot can make predictions about how it will be perceived by a listener, and adjust its navigational strategies appropriately.

In this work, we propose and implement a navigational controller that incorporates some of these ideas of awareness into a stealthy approach scenario. Assuming that our listener is capable of recognizing either overall changes in volume or significant changes in volume from any given direction, a stealthy robot needs to recognize how its own movements will be perceived by each of these listener capabilities, and incorporate that into its own movement strategy. Specifically, to reduce the acoustic impact of the approach on the listener, the robot needs to first estimate the overall volume of ambient noise the listener is exposed to at their current location and the relative masking effects of each source in the environment. Then the robot predicts for all reachable locations in the environment how loud it will sound to the listener from that location and, combined with the previous information, identifies the path that will expose the listener to the perceptually least amount of robot-generated noise.

The remainder of this paper is organized as follows. In the next section, we discuss related work in auditory scene analysis and robot security. Then we will describe the scenario being tested and the algorithms we used to enhance robotic performance. Finally, we will discuss four scenarios under which a real robot stealthily approached a target with varying degrees of success.

II. Related Work

This work in hiding the acoustic signature of a mobile robot is strongly related to two fields: auditory scene analysis, and home security. The first area, auditory scene analysis allows a robot to grasp some aspects of the auditory scene around it with the goal of enhancing its recognition rate for acoustic phenomenon. By being able to recognize certain types of auditory scenes [3], a robot dropped into an arbitrary location could eliminate some unimportant sounds from consideration by knowing that they are common sounds in that area. Similarly, by separating sounds from each as much as possible using signal processing, the effects of masking noise on important events are reduced [4]. For a robot trying to be stealthy, knowing the location and innocence of certain sounds is important to disguising itself using the ambient noise. Another important capability is being able to predict the general noise levels in the environment [5], so that the robot can estimate how much noise it can generate while approaching the target.

Where auditory scene analysis has typically focused on separating events from the ambient scene, security robots have been more concerned with detecting acoustic phenomenon over a large area. The types of phenomenon that are interesting to the robot vary greatly between applications, often including non-human threats such as fires

Manuscript received March 27, 2007.
E. Martinson is with the Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: ebeowulf@cc.gatech.edu).

and leaking water [6], but also setting a goal of identifying intruders. Such work for detecting intruders employs cameras [2, 6], microphones, or combinations of both [1], and often is deployed in conjunction with sensor networks distributed through the environment [7]. What to do about the detected event, however, has remained nebulous, and the responsibility is usually passed off to a remote human controller who make the decisions about the nature of the appropriate action. However, a human controller may become habituated to the robot feedback, ignoring important events, if the robot requests human intervention for too many unimportant phenomenon. Under such circumstances, a robot may need to reduce the number of human intervention requests by first getting better information about the detected phenomenon. Approaching the target stealthily is a way to gather such necessary information.

III. How to Hide a Noisy Robot

The scenario proposed for testing the acoustic hiding abilities of a robot is the stealthy approach scenario. The target being approached is a 4-element microphone array capable of detecting changes in the overall volume, as well as identifying changes in the relative volume from each direction. This listening system is designed to mimic the perceptual capabilities of a human target, which can identify changes in overall volume, and separate sound sources from each other by angle. For now, our sensor system is not searching for differences in pitch.

Given this target listener, the robot's goal is to approach the target as quietly as possible, moving from some starting location to within a meter of the sound source. For this task, the robot is given knowledge a priori of significant sound source locations in the environment, their directivity, and a spatial evidence grid from which it can localize itself with respect to the environment. As demonstrated in Martinson and Schultz [5], these are all pieces of knowledge that could be acquired by the robot. Their acquisition, however, would require that the robot be deployed to that area at some time before being asked to approach the target.

The methodology used to hide the robot's acoustic signature is based on the capabilities of the target listening device. First, the robot estimates the volume of noise the observer is exposed to without the presence of the robot. Second, using the provided obstacle map, the robot identifies a set of discrete reachable locations in the environment. Then, for each location, the robot estimates a cost of visiting that location based on: (1) the absolute difference in volume at the receiver due to the robots presence at that location, and (2) the difference in the volume coming from the direction of the robot. Finally, these two cost estimates are combined together using weighted summation, and a path-planner identifies the path of minimal cost for the robot to travel.

A. Estimating Volume at Target

The first step in hiding a noisy robot is to estimate the overall volume detected by listener. This will be used to determine which areas of the environment are considered safe for the robot to enter undetected.

When making this estimate, there are two types of sound waves affecting the listener. The first type of waves are those that propagate directly from the source to the receiver without being reflected. Given a sound source (S_i) of volume (V_i), the angle (α_i) and distance (d_i) from that sound source to the listener, and the directivity function of that source ($Q_i(\alpha)$), we can use spherical spreading to estimate the direct sound as decaying with square of the distance:

$$S_i(d_i, \alpha_i) = V_i Q_i(\alpha_i) - 20\log_{10}(d_i) \qquad \text{Eq. 1}$$

Where S_i is in dB, and Q_i is determined on a dB-scale for sound-pressure level, instead of power.

The second type of waves affecting the listener is reverberant waves, which reflect off of one or more surfaces before reaching the listener. Although each wave is usually low volume due to the longer distances traveled, the large number of reflections in an average room mean a significant contribution to the overall volume of noise the listener hears. Modeling the reverberant field can be done using ray-tracing, but this needs a lot of information for an accurate model [8]. Such information includes detailed surface maps and the position of every known source in the room, which a robot may not have. A simpler assumption often used in architectural acoustics is that the reverberant field remains constant over the entire area [9]. By sampling the environment at some location relatively far away from any known sound sources, the robot can dynamically estimate the volume of the reverberant field (R) before planning an approach path.

The estimated combined volume of noise heard by the listener is then the logarithmic sum of the volume due to each source with the reverberant field:

$$T = 10\log_{10}\left(10^{R/10} + \sum_i 10^{S_i/10}\right) \qquad \text{Eq. 2}$$

B. Minimizing Changes in Volume

After estimating the volume of noise heard by the listener, the next step is to estimate how loudly the robot will be detected. Specifically, for each location in the environment that a robot can move through, how much additional noise will the listener hear due to the presence of the robot? This is accomplished by again using spherical propagation (Eq. 1) to estimate the volume of sound reaching the listener. Repeating this direct sound estimation for every unobstructed location in the environment, we can create a map of how loud the robot will appear to the listener for every location (Figure 1).

For now, this model does not include any reverberant obstacle or obstacle effects on sound propagation. Unlike the previous section, this section requires estimating the effects due to a single source, the robot. As such, the robots'

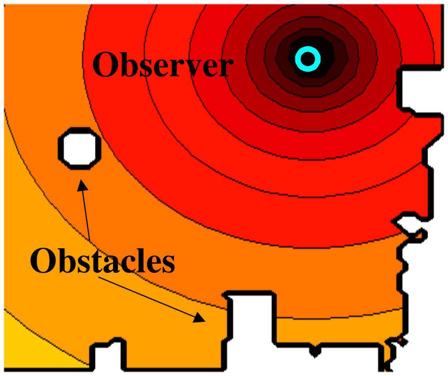

Figure 1. Contour map of the estimated noise at the observer due to the robot. Darker is louder.

own reverberant effects on the environment cannot be easily estimated without more knowledge of the environment, or measured in the presence of other noise sources. In the future, however, we intend to introduce a ray-tracing model to this estimate to include these effects from a single source.

C. Avoiding Directional Cues

Knowing just the volume of the robot at the target, however, is only part of the problem. Since the target is a microphone array, it is capable of estimating the angle to the detected sound source. So even if the overall volume of noise did not change significantly, it can still detect the robot if there is a significant deviation in angular energy from the baseline. Hiding the robot, therefore, requires choosing a path that also minimizes the change in angular energy. Such paths will be along the line from the source to the listener.

A source is only going to mask the robot's acoustic signature if the robot is between the source and the listener. How much the robot is masked by that source will depend on how loud the source is, and how far the robot is from the axis joining the source and the listener. For this purpose, we used a heuristic to estimate the directional occlusion of each source separately in dB, and summed the results together:

$$D_i(x,y) = W(l_i(x,y))[V_i Q_i(\alpha_i)]$$

$$D_{x,y} = 10\log_{10}\left(\sum_i 10^{D_i(x,y)/10}\right) \quad \text{Eq. 3}$$

Where V_i and Q_i are the source volume and directivity, $O_i(x,y)$ is the resulting directional occlusive effect at position (x,y) for source i, l_i is the distance from the robot to the line between source and listener, and W is a normalized bell curve with standard deviation of 1m.

D. Picking a Path

Now that we have finished estimating the volume at the listener, the volume of noise due to the robot, and an occlusive effect due to each source in the environment, we need to estimate the combined impact on the listener ($I_{x,y}$) for a robot being in each reachable location (x,y). This total impact will then be used with a path-planning algorithm to find the path with the smallest impact.

The first step in our heuristic for minimizing impact is to identify the environmental impact on the observer ($Env_{x,y}$). This is calculated as a log summation of the predicted total volume (T) at the observers location, plus directional occlusive effects (D) in viewing the robot position (x,y):

$$Env_{x,y} = 10\log_{10}\left(10^{T/10} + 10^{D_{x,y}/10}\right) \quad \text{Eq. 4}$$

Next, the impact of the robot traveling through that location ($I_{x,y}$) is the total impact on the listener (environmental impact plus the estimated sound heard by the listener due to the robot, $R_{x,y}$) minus the environmental impact:

$$I(x,y) = 10\log_{10}\left(10^{R_{x,y}/10} + 10^{Env_{x,y}/10}\right) - Env_{x,y} \quad \text{Eq. 5}$$

Finally, the robot picks a path a stealthy approach path by finding the shortest weighted path from the start to the goal using djikstra's single-source shortest path algorithm with impact being the weight of being in any given location. Figure 2 shows a contour map of the estimated impact for all unobstructed locations in the environment, using these equations with one 57 dB source.

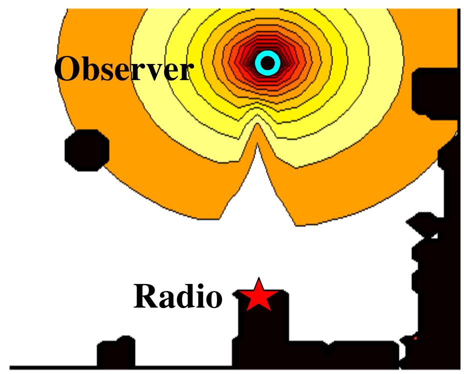

Figure 2. Contour map showing estimated impact on an observer due to a robot at some location in the environment. Darker means greater impact.

The reason for using an absolute difference in Eq. 5 is that we need the impact to vary with reverberant levels in the room. If the listener is overwhelmed by an 80dB noise in the area, then the impact of the environment should dominate the equation and reduce the impact of an approaching robot generating only 47dB. If, on the other had, the position of the target is a relatively quiet 40dB, then the approach of the robot should be a lot easier to detect. Using an absolute difference between total impact and environmental impact will reflect this difference, and allow the robot to adjust its path to the current level of reverberant sound in the environment.

IV. Experimental Setup

The robot hardware that was used for this task is a Pioneer-2dxe robot equipped with a SICK LMS200 for localization and obstacle avoidance (Figure 3). This robot platform emits roughly 47-dBA of noise in all directions (as measured by a Type II SPL-Meter) from its onboard cooling fans while standing still. Additional ego-noise in the form of

Figure 3. Pioneer 2-dxe used in the acoustic hiding scenario.

impulse sounds from the wheels rubbing on the tile floor is also occasionally observed during robotic movement.

The goal of the stealthy robot is to move from a specified start position to within 1-m of the observer's position as quietly as possible. Figure 4 shows the layout of one scenario setup in the Mobile Robot Laboratory, along with the two paths taken by the robot in the first scenario. The obstacles shown in the middle of the lab are all roughly 1-m in height.

Evaluation of the robot's performance involved analyzing the data collected from a 4-element microphone array located at the target's position. Sampling at 8192Hz, the array collects 1-s samples continuously over the duration of the run. This includes collecting 30-s of data with no robot in the room to set a baseline, and then, roughly 100 samples for longer paths, and 50 samples for shorter paths. Each sample was then analyzed to determine:

- Overall change in volume from baseline (dB)
- Change in volume from the direction of the robot.

This second metric required that each sample also include an estimate of where the robot was currently located in the room. For this purpose, we collected the believed location from the player/stage amcl driver whenever a sample was collected. Then, to estimate performance, we used a time-delay estimation algorithm, based on generalized cross correlation measurements, to estimate the energy at 1-m from the listener in the direction of the robot. The difference between this energy (in dB) and the mean energy at that angle from all noise samples (in dB) is the empirical measure of angular impact on the listener due to the robot.

V. Results

This work was tested with the Pioneer 2dxe robot in a total of four scenarios spread across two different environmental layouts. In each of these scenarios, the performance of the robot trying to approach the target stealthily is compared to a robot taking an alternative, usually shorter path.

A. *First Environmental Layout*

The layout of the first scenario is a relatively open 8x8m environment, with an observer located relatively far from any walls and a sound source pointed at the observer. In the first scenario, the robot uses its knowledge of the radio in the environment to hide itself better than an uninformed robot taking the shortest path to the target. In the second scenario, the performance effects of a significantly louder reverberant field are examined.

1) *Hiding in Front of a 67dB Source*

In this scenario, a 67dB source was placed 4-m to the left of the listening microphone array. That source was an fm radio with a typical cardioid directivity pattern generating static noise. Starting from a location below the listener in the map (Figure 4), the shortest path was to move upwards in a roughly straight-line while avoiding obstacles. The robot that was trying to hide its acoustic signature, however, would move upwards to get in line with the source before approaching the target. This scenario was repeated 30 times for each robot path.

Given our open environment, and the listener's positions being all relatively far from the wall, the first metric did not produce significantly different results for the two paths except in one region. For most of either path, the 47dB robot added little overall volume (<1dB) to the total energy in the room. This is not surprising, as the reverberant field averaged 54dB for this environment, while the reverberant field due to the robot (measured with the sound source turned off) added a significantly smaller 43dB. The exception to this rule, however, was part of the path taken by the acoustically hiding robot where the robot turned relatively sharply to get in line with the radio. This region is marked "turning region" in Figure 4. While turning, the robot generated a noticeably louder amount of noise, mostly tire squeaking and equipment rattling, which violated the original assumption of the robot as a constant 47dB source.

With the exception of the turning region, the first metric had very similar results for either path. The second metric, however, demonstrates a significant difference between the paths. In Figure 5 (Top), the average energy at a given distance from the observer is estimated by using a passing a Gaussian smoothing function over the sampled data (std=0.1m). Looking at the shortest path energy from 3.5m to the stopping point 1-m from the observer, a relatively steady volume can be detected until ~1.5m where the presence of the robot becomes more noticeable. In contrast,

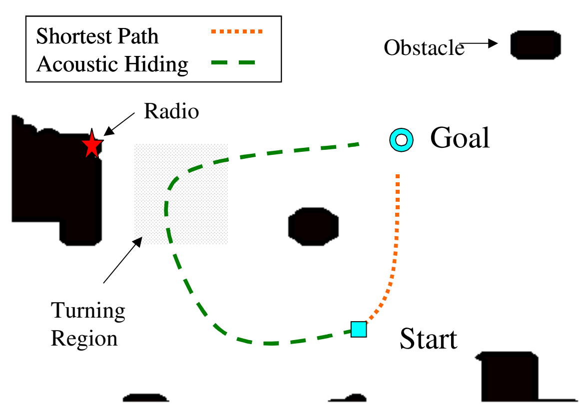

Figure 4. Layout of the acoustic hiding scenario. The robot that does not try to hide approaches the observer along the shortest distance path. The robot that tries to hide its acoustic signature moves in line with the radio source, before approaching the observer at the goal.

reverberant field scenario. The overall increase in volume detected by the observer was minimal (<1dB) for all parts of either path. Measuring angular detection energy saw similar results. Taking the shortest path meant a less than 1 dB increase in volume over the maximum reverberant field noise in 99% of the samples, while the robot on the longer path remained unobserved in 96% of the samples. With the longer path, 90% of the detected samples occurred in the "turning region" where the robot is aligning itself with the radio.

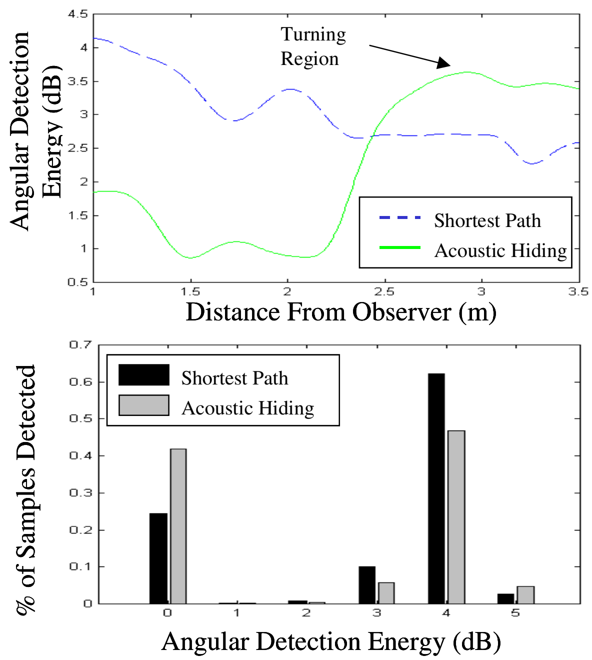

Figure 5. (Top) The robot taking the shortest path has a higher average angular detection energy over the last 1.5m than the robot trying to hide its acoustic signature. (Bottom) The robot hiding its acoustic signature remained undetected in 43% of the samples measured by the observer, as opposed to 24% for the robot taking the shortest path.

the robot trying to hide from the observer first demonstrates higher energy while it is getting in line with the source, but then quickly drops down into the unnoticeable range as the robot hides in the radio noise. Figure 5 (Bottom) shows the difference between the collected samples. While the shortest path was unnoticeable 24% of the time, the robot that hid in the radio static noise was unnoticeable 43% of the time on a similar sized sample set (~600 samples/path).

2) Loud Room Scenario

The second test using this room layout was a repeat of the 57dB source scenario, except that the reverberant field in the room was raised to over 60dB using a loud floor fan placed in a far corner of the room (away from the testing area). The hypothesis behind this test was that a loud enough room should eliminate the advantage of any particular path, because the addition of the robot will be too small.

The effect of this change to reverberant sound levels on the robot's path-planning algorithm was to logarithmically reduce the cost (or weight) of visiting any gridcell in the map. This applied nonlinear decrease in all weights means that the shortest-path becomes less costly than the longer path, because the robot travels across fewer grid cells to reach the goal. Therefore, after detecting the change in reverberant noise, the robot does not try to get in line with the source, but simply approaches the source from the shortest distance path. To determine whether or not this path was chosen correctly, we also tested the path chosen for the quieter room with just the 57dB source.

The two paths were each tested 15 times in this loud

B. Second Room Layout

The second room layout was designed to add a larger reverberant field component to the detection of the robot. Nearby walls would amplify the noise of the robot, making it easier to detect. Since this effect is not modeled in the path-planning algorithm, there should be a performance decrease from the previous layout.

In this second environmental layout, the first scenario shifts the radio source to a different location in the room and tries to duplicate the success of the first scenario. In the second scenario, a quieter fan source is substituted for the radio source.

1) Hiding in Front of a 67dB Source

In this scenario, the same radio source used in the previous room layouts was moved to a location 4-m below the listening microphone array. Starting from a location to the left of the listener in the map (Figure 6), the shortest path was to move to the right in a roughly straight-line while avoiding obstacles. The robot which was trying to hide its acoustic signature, however, would move down, along the wall, before moving upwards to get in line with the source to approach the target.

As expected, this scenario saw a significant decrease in performance, both overall, and relative to the other run. The total volume due to the robot remained small over the entire path, with no region exceeding the average noise level by more than 1dB. Looking at the angular energy, however, sees

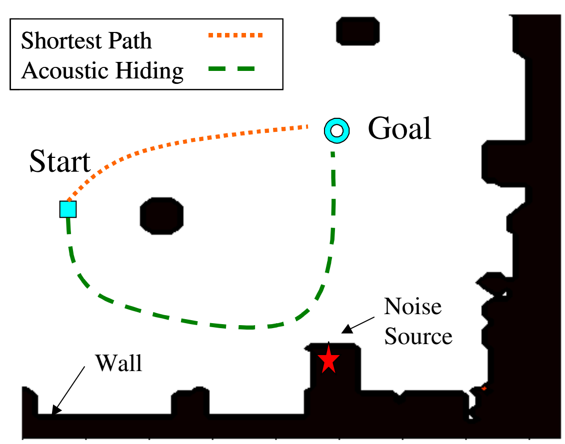

Figure 6. The second environmental layout used to test acoustic hiding performance. In this scenario, nearby walls make the robot more easily detected due to reverberant effects.

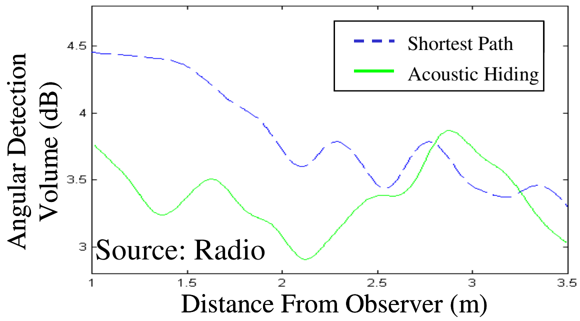

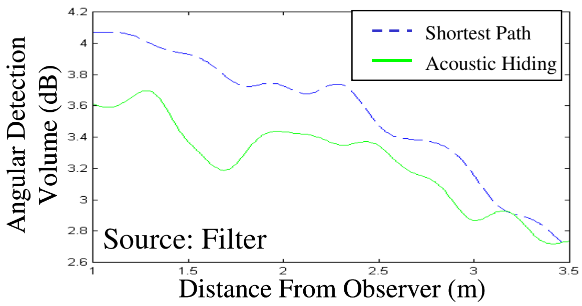

Figure 7. (Top) The robot hiding from the observer using the radio in the second environmental layout is detected more often than before, but still does better than the robot taking the shortest path. (Bottom) The quieter fan source is only slightly worse than the radio at hiding the robot. Notice that where the energy dips in the (Top) map as the robot moves in line with the radio, moving in line with the fan source does not produce this result.

that the robot trying to hide in the radio's noise was undetected (energy less than 1dB) in only 17% of the samples. While this was still better than taking the shortest path, where the robot remained unobserved in less than 9% of the samples, the difference between the two runs was not as significant as with the same source in the previous layout.

Figure 7 (Top) plots the angular energy vs distance from the observer for this scenario. The energy detected from the acoustically hiding robot is noticeably less than the robot taking the shortest path, but not by as much a margin as with the first room layout.

2) Hiding in Front of a 54dB Source

In this scenario, a 54dB source was placed 4-m below the listening microphone array. That source was an air filter with a bipolar directivity pattern generating wind noise. With the 3dB difference between this source, and the radio source, it was expected that this configuration would produce another drop in performance.

After 15 trials for each path, the robot's performance is as expected. Seen in Figure 7 (Bottom), the angular detection energy vs distance to the observer plot shows almost a constant offset for the two paths. Where the radio sources produce a sharp dip in the energy once the robot moves in front of the source, the fan did not produce such a dip. Instead, the presence of the fan appears to merely lower the overall detection of the robot by some small amount.

VI. CONCLUSION

The goal of this initial work in acoustic hiding was to demonstrate that a robot could hide its own acoustic signature in the ambient noise. Given some knowledge of the auditory scene, a robot can position itself between known sources and a target to reduce the chances of being detected by a listener at some arbitrary location in the environment. In this paper, we have demonstrated the effectiveness of this simple technique for two environmental layouts and four source/auditory scene configurations. Together, these scenarios explored the effects of different source volumes, reverberant field levels, and the general shape of the robot's path on performance in a stealthy approach scenario.

What has also been demonstrated in this work is the complexity of the acoustic hiding problem. If the volume of the source disguising the robot's approach decreases, then the robot will be detected easier. If the reverberant field increases substantially, then the robot may not need a stealthy approach to remain undetected. The presence of nearby walls in the environment may also make the robot more detectable, as will certain types of robotic movement that cause the robot to generate more noise. The combination of all of these factors is a complex task, which we have only touched upon in this paper, and hope to spend more time refining and improving upon in future work.

REFERENCES

[1] D. Carrol, G. A. Gilbreath, and H. R. Everett, "Extending Mobile Security Robots to Force Protection Missions," presented at AUVSI Unmanned Systems, Lake Buena Vista, FL, 2002.

[2] A. Treptow, G. Cielniak, and T. Duckett, "Active People Recognition Using Thermal and Grey Images on a Mobile Security Robot," presented at IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, Alberta, Canada, 2005.

[3] C. Couvreur, V. Fontaine, and C.G. Mubikangiey, "Automatic Classification of Environmental Noise Events by Hidden Markov Models," *Applied Acoustics*, vol. 54, pp. 187, 1998.

[4] H. Okuno, T. Ogata, K. Komatani, and K. Nakadai, "Compuational Auditory Scene Analysis and Its Application to Robot Audition," presented at Int. Conf on Informatics Research for Development of Knowledge Socieity Infrastructure, 2004.

[5] E. Martinson and A. Schultz, "Robotic Discovery of the Auditory Scene," presented at IEEE Int. Conf. on Robotics and Automation, Rome, Italy, 2007.

[6] J. Liu, M. Wang, and B. Feng, "iBotGuard: An Internet-Based Intelligent Robot Security System Using Invariant Face Recognition Against Intruder," *IEEE Trans. on Systems, Man, and Cybernetics - Part C: Applications and Reviews*, vol. 35, pp. 97-105, 2005.

[7] R. Luo, C. Lai, and C. Liao, "Rapid Environment Identification for Intelligent Mobile Robot," presented at IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, Beijing, China, 2006.

[8] L. Savioja, Ph.D., "Modeling Techniques for Virtual Acoustics", Helsinki University of Technology, 1999

[9] W. Cavanaugh, "Introduction to Architectural Acoustics and Basic Principles," in *Architectural Acoustics: Principles and Practice*, W. Cavanaugh and J. Wilkes, Eds. New York: John Wiley & Sons, 1999, pp. 1-54.