Reflections on modelling a sonar range sensor

Gregory Dudek, Michael Jenkin *,
Evangelos Milios*, David Wilkes †

CIM-92-9 December 21,1993

McGill Research Centre for Intelligent Machines
McGill University
Montréal, Québec, Canada

* Department of Computer Science, York University Downsview, Canada
M3J 1P3
† Department of Computer Science, University of Toronto Toronto, Canada
M5S 1A4 and Ontario Hydro Research Division.

Postal Address: 3480 University Street, Montréal, Québec, Canada H3A 2A7
Telephone: (514) 398-6319 Telex: 05 268510 FAX: (514) 283-7897
Network Address: mcrcim@larry.mcrcim.mcgill.edu

Reflections on modelling a sonar range sensor

Gregory Dudek, Michael Jenkin ,
Evangelos Milios*, David Wilkes

Abstract

With the advent of inexpensive off-the-shelf sonar ranging units sonar has become a common sensor on mobile robotic platforms. Many different algorithms have been proposed to obtain environmental layout and other structural information by integrating sonar measurements from multiple positions and/or sensors. Given perfect distance measurements and telemetry, this integration can be fairly simple. Although sonar has become a ubiquitous sensor in mobile robotic systems, surprisingly few results are available data interpretation strategies are often very simple and surprisingly few results are available that accurately model the typical behaviour of the sensor under such conditions. This paper considers some of the potential causes of sonar errors, their effect on surface or object recovery, and then develops a simulation-based model of sonar range sensing for robot navigation that accounts for multiple reflections of the sonar signal between transmission and reception. This gives more realistic results than previous models. The approach is based on simulation of the reflection and diffraction of sonar rays from reflecting surfaces until they are attenuated beyond detectability or return to the receiver. Parameters of the model include frequency, minimum and maximum range, and signal detection threshold (relative to emitted signal strength, after linear gain compensation). Finally, the usefulness of the model to the development of more effective algorithms for the interpretation of sonar data is discussed.

1. Introduction

Sonar range sensing refers to technologies for emitting sound signals and measuring the characteristics of the echo that returns. The most common form of sonar range sensing – the one we will address here – is based on emitting a sonar signal (a "chirp") and measuring the time delay until an echo is detected. This provides a very simple and economical method for distance measurement.

A single transducer is often used as both transmitter and receiver. From knowledge of the time taken for the sound to return and of the speed of sound, the distance to the object reflecting the pulse can be inferred. A typical and widely used ultrasonic ranging device is that produced by the Polaroid Corporation. Its use is described in detail elsewhere[1; 12]. It operates using sound signals with frequencies of roughly 50 kHz. Due to its low cost, wide availability, low weight and low power consumption, Polaroid sonar units can be found on many research and commercial robotic vehicles. Many researchers have proposed algorithms for object avoidance[2], map making[4; 13], and other complex robotic operations based on sonar measurements. In addition, many "simulation only" results have been reported showing potential uses of sonar rangers for even more complex processing.

Despite the prevalance of sonar in mobile robotic systems there is a paucity of models describing to behavious of sonar in realistic environments. Existing software simulations are often very simple and fail to reflect significant phenomena. Sonar simulation has many advantages in the early stages of robotic algorithm development. As is the case with simulations in general, simulated environments can be easily constructed, re-arranged and modified. Experi-

ments can be repeated, and specifically in the case of simulating sonar-based mobile robotics, it is possible to move a simulated robot much more easily than a real one. That being said, there are a number of "simulation only" sonar-based robotic projects which model the sonar response as the distance to the closest reflecting structure in the environment in the direction the sensor is pointed. Some simulations assume simple errors, such as modeling the error as an uncorrelated Gaussian process. While simulated sonar data is highly desirable for many reasons, it is important that the simulation accurately model the true characteristics of the sensor.

Through a few simple demonstrations of real sonar measurements we show that simple simulations are incapable of accurately describing the sonar errors (or effects) that occur in real environments. More complex simulations (such as [14]) can account for some, but not all, of the complex effects seen in real sonar data. A minimal sonar sensor model must correctly model both the physical nature of the sonar pulse and receiver, and also the way in which the sonar pulse interacts with structures in the environment. We will also examine some of the complex interactions that we have found when sonar responses are obtained in real environments. Sonar responses from real environments can be very complex, and some common features, such as corners and parallel straight walls can give rise to complex interactions between adjacent structures and the body of the sonar unit itself.

Sonar devices are well known for the apparent unreliability of their readings. In a typical sonar scan taken around the compass from a single position, very few of the distance readings returned are accurate direct ranges to the nearest object in the aiming direction. Several authors [3; 11] have produced clever interpretation methods that attempt to cope with this unreliable data. The methods essentially throw away all measurements that do not

have sustained support from multiple viewpoints. They do not incorporate sophisticated models of the measurement process. The model discussed herein accounts for much of the structure in the data that is typically thrown away. An encouraging conclusion that has been reinforced by this work is that many classes of ranging errors are in fact *repeatable* and *predictable,* features of the sonar device. The hope is that it may be possible to develop better sonar interpretation algorithms, given a deeper understanding of these features.

Sonar units are typically installed in a robotic platform in one of two ways: One approach is to mount a single sonar unit on top of rotating platform and point it in different directions. A second approach is to mount a collection of sonar units around the exterior of the robot. These two systems have subtle implications with respect to the types of sonar errors that can occur.

Ignoring the possibility that multiple sonar units may interact with each other (crosstalk), there are two major potential sources of sonar errors. The first arises from the physics of the sonar process itself: the sonar chirp is not an infinitely narrow beam of infinite power but has a finite angular extent with a complex cross-sectional energy distribution. A detailed study of the effects of off-axis energy in the sonar signal, and simulations of these effects can be found in the work of Kuc and Siegel[5], McKerrow[9; 10] and Wilkes *et al*[14]. Experimental studies of the side lobes have also been reported[6]. In addition, a distance threshold on measurements is imposed by the fact that when the sonar pulse is emitted by the transducer the sound begins to travel away from the sensor, the sensor goes through a transition phase during which it is unable to listen for the returning pulse, and then waits for a signal to return. The second critical source of error arises from the conventional assumption that (usable) measurements provide the straight-line distance to the nearest obstacle in the aiming direction. For typical sonar frequencies, most objects

in a man-made environment appear to be specular acoustic reflectors. Hence, the accidental alignment of surfaces in the environment may give rise to a path back to the transducer involving multiple reflections.

If only the interaction between a single sonar unit and a single obstacle in the environment is considered, the distance error associated with a given sonar response can be well modeled by a normal distribution[6] with a slightly longer far distance tail in the distribution than the near tail. A simulation for this type of sonar error model is quite simple to design. It also suggests that least squares or Kalman filtering type approaches would give rise to successful algorithms for integrating sonar measurements, as the noise associated with each measurement can be treated independently and each can be modeled by some Gaussian process. (See [8; 7] for examples of sonar measurement integration algorithms using this type of assumption.) Unfortunately very few environments consist of only a single surface. For more complex environments, sonar errors are not well modeled by independent Gaussian distributions, but are rather highly correlated and very complex.

The sonar simulation presented here can accurately model many of the complexities of sonar ranging, and so it is still possible to perform simulation-based sonar experiments. Although not capable of modelling all of the subtleties involved in real sonar ranging, the simulation presented here provides a more realistic model for testing simulations of proposed sonar processing algorithms.

2. Raw sonar data

Sonar data, as obtained via a sonar subsystem such as those available on mant robotic systems presents a simple digital model of the underlying analogue data

actually being returned by the sonar unit. Although the resulting data appears almost ideally suited for higher-level robotic tasks, the hidden preprocessing masks much of the complexity inherent in raw sonar data and also much potentially usable information. If we examine the raw analog return echo detected by a sonar transducer, deviations from the ideal case can be readily observed.

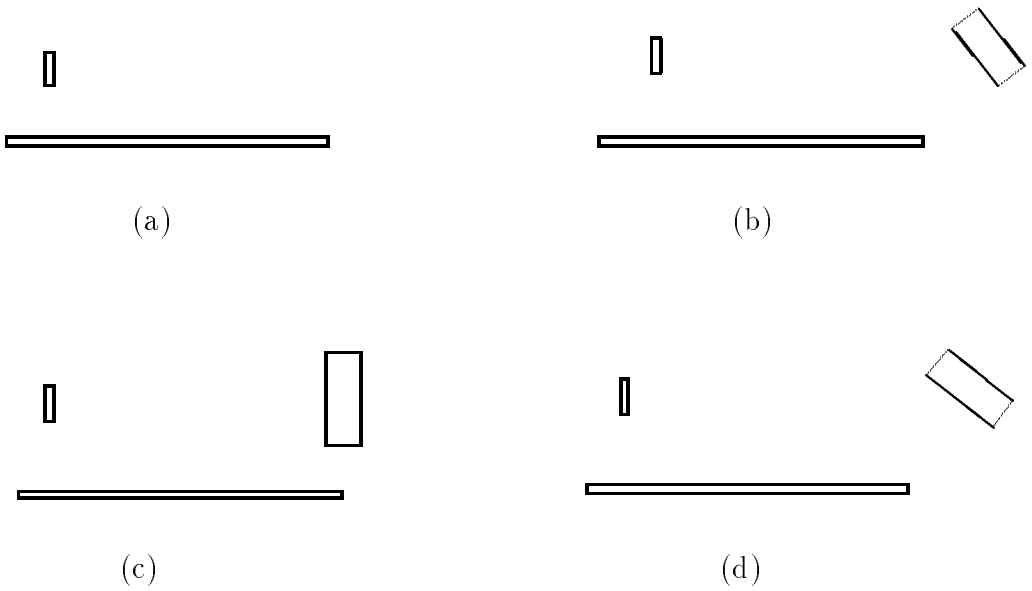

Figure 1: Experiment configuration
Sketch of the experimental setup used to obtain the traces given in Figure 2. Case (a) is the situation in which the transducer has no target within the timeout distance. Case (b) is the situation in which a rectangular target is inclined at a sharp angle towards the transducer. Case (c) is the situation in which a rectangular target is parallel to the transducer. Case (d) is the situation in which a rectangular target is inclinded at 45^o *towards the transducer.*

Figure 1 shows four simple "ideal" cases for a sonar transducer. A Polaroid

transducer was aligned with a movable rectangular object, and the raw sonar signal returned by the transducer was recorded (see Figure 2).

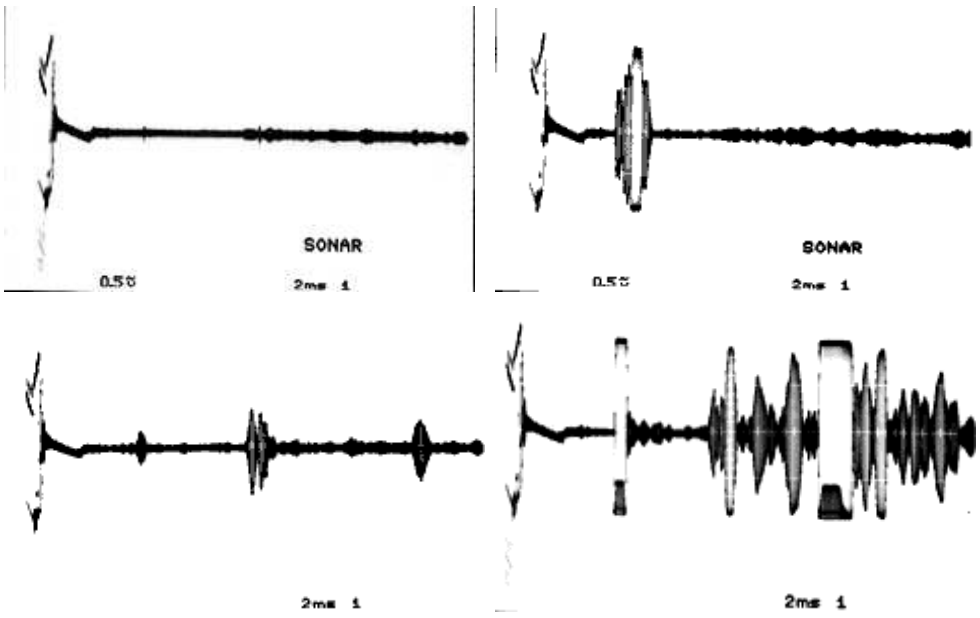

Figure 2: Low level sonar traces
The sonar traces shown have identical horizontal (time) and vertical (voltage) scales. The situations that gave rise to these traces are showin in Figure 1. These figures have the same horizontal (response strength) and vertical (time) axes. The traces were obtained between the amplified transducer return signal and ground from the Polaroid transducer module. For left to right, top to bottom: (a) Empty environment, (b)"Ideal" isolated obstacle, (c) "Realistic" environment, (d) Reflected obstacle.

Figure 2a shows the very clean sonar response for an empty environment (illustrated in Figure 1a). The tail end of the sonar "chirp" as it leaves the transducer is visible at the left hand side of the display, and the signal is not reflected by any close structure in the room and thus does not return to the transducer. Background noise in the room, and weak scattering of the signal due to irregularities in nearby oblique surfaces, provide some slight variations

from a zero response.

Figure 2b shows what one would expect would be the "ideal" sonar response for a simple isolated target. A single isolated chirp is returned to the sensor from which the time of flight can be computed. Unfortunately this response corresponds to the setup shown in Figure 1b, rather than Figure 1c. The actual angle in which the target was positioned was selected experimentally so as to produce an isolated response. When the target is placed parallel to the transducer (Figure 1c and Figure 2c), a large number of multiple echoes are recovered. The sonar signal takes multiple paths, all of which return to the sensor. Finally Figure 1d and 2d show the ideal reflected case in which the strongest returned echo is a path from the transducer, bouncing off of the target, reflecting off some far object and then returning via the target to the sensor. A small earlier response shows that some signal did return directly from the target. By modifying the angle of the target it was possible to completely remove this earlier echo.

The analog and digital preprocessing that is applied to this returned signal amplifies it so as to reduce the effect of signal attenuation due to distance, and then typically selects the onset of peaks which exceed some threshold as the "true" responses. The difference between the "expected" response, and the "true" response can be significant. It is a function of not only the relative positions of the transducer and the target, but it is also a function of other structure in the environment.

3. A real room example

Figure 3 shows a floor plan of a laboratory at McGill University. A single sonar unit, mounted on top of a stepper motor assembly, was located on a table top

to produce the scans. Figure 4 shows the positions of sonar responses obtained by taking one scan roughly every five degrees, and then joining a line between successive scans. A number of scans are overlaid to show the repeatability of the measurements.

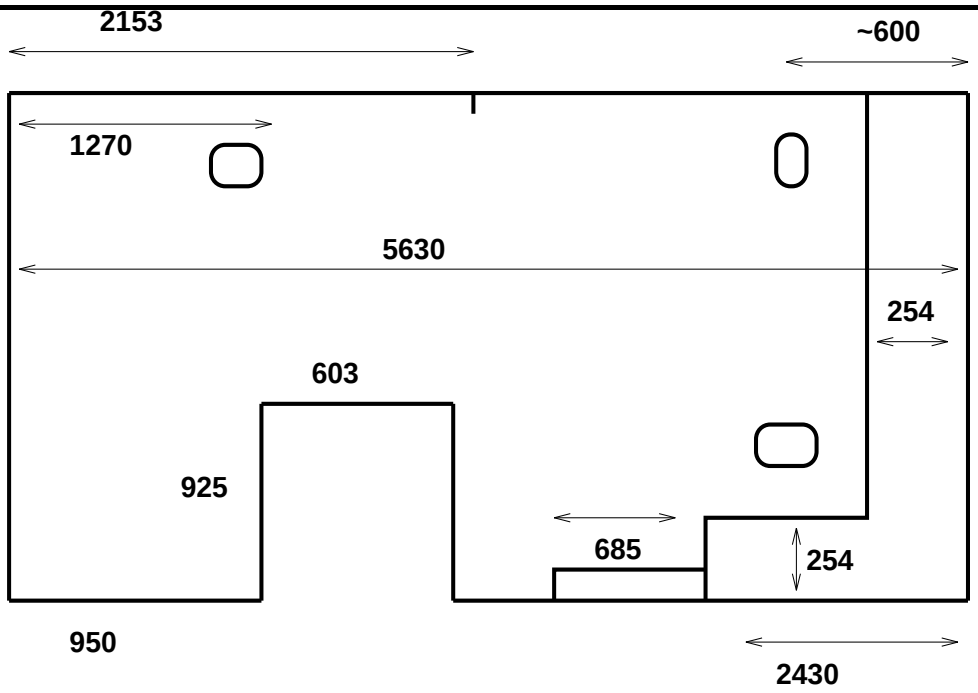

Figure 3: Measurements of a room (centimeters)

Perhaps the most obvious result is that the sonar measurements do not well approximate the position of objects in the room. It is illustrative to describe which features in the room seem to have been detected by the sonar unit and which features have not.

Directly up from the sonar unit, the nearer wall was recovered for a small distance, and then very long distances are reported. This artifact arises from any point in the central lobe of the sonar signal being reflected back to the transceiver from a surface perpendicular to the receiver, but then being re-

flected away for larger angles. If this reflected signal never returns to the sensor, the display software drops the data point. Responses displayed indicate that the reflected signal managed to follow a multiple-reflection path that returned to the sonar unit. Following the ring of responses counter-clockwise, the signal returns to a plane roughly at the wall depth, but slightly closer. This effect was caused by the edge of a blackboard protruding very slightly (about 20 mm) from the wall.

The protruding corner makes for a very good sonar response. One of the authors (Wilkes) was seated in the chair at the upper left of Figure 3, and authors also make for good sonar responses, because of the large variety of surface orientations that they present. Continuing around the bottom portion of the response, the pillar protruding into the room is partially recovered, but part of the signal is reflected away.

The lower portion of the figure shows the correct recovery of portions of this wall. The lower right hand part of the figure shows complex responses due to that portion of the room being cluttered with computer equipment and another author (Dudek). Continuing the signal around to the top of the figure, the reflection effect is once again seen, in which responses above the critical angle are reflected away from the true surface.

This example illustrates some significant points. Sonar responses are the result of complex interactions between the sonar beam and objects in the environment. A robot travelling down a straight corridor may have its sonar signal bounce and reflect away due to collisions beyond the critical angle. This signal may not be lost, but may rather be subsequently reflected back from complex structure in the environment, giving rise to misleading long distances to targets.

Consider what would happen if the sonar unit was translated parallel to

Note that these responses are not to the same scale as the floorplan shown in Figure 3.

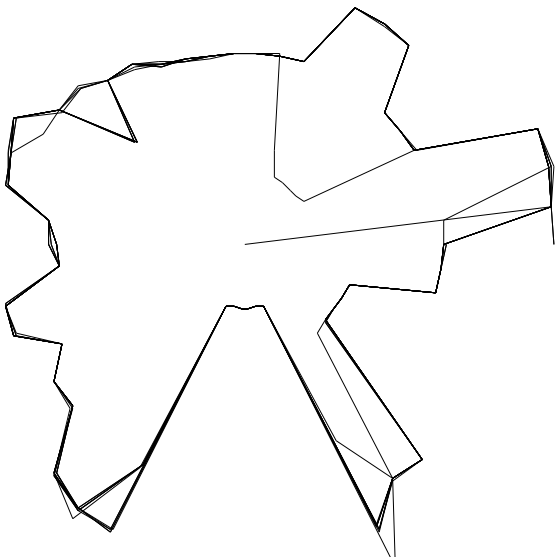

Figure 4: Sonar responses for a room at McGill University

a single wall in a long corridor. Some measurements associated with parts of the wall at right angles to the robot (the top and bottom of Figures 3 and 4), will be correctly recovered, while walls ahead and behind the robot may give rise to reflections, and long distance effects (like the effects seen in the above figures). These errors are not random, but rather are highly correlated from one transducer azimuth to the next, or from one robot position to the next.

Any simulation or model of sonar sensors cannot simply assume that the response from a particular surface can be modelled independently of other structure in the environment. This effect is similar to problems encountered in computer graphics with reflective surfaces. Rendering techniques for fully opaque matte surfaces can be much simpler than techniques for scenes containing reflective ones. A common computer graphics approach to rendering reflective surfaces is ray tracing. A similar approach can be used to more accurately model sonar sensors.

4. Model and Implementation

Earlier models of sonar sensors have accurately modelled the way in which a single sonar pulse is created, interacts with a single surface, and is detected [5]. The model presented here utilizes these earlier results in order to model the interaction between the sonar pulse and individual surfaces. In order to accurately model the complex interactions between the sonar pulse and multiple surfaces in the environment, a ray-tracing-like algorithm will follow the path taken by the pulse as it interacts with different structures in the environment. In a computer graphics ray tracing algorithm light rays are traced from the image plane, off (and through) structure in the environment, until the ray is either lost, attenuated beyond visibility, or intersects a light

source. In this sonar model, a sonar chirp will be traced from its emission, through interactions with environmental structure, until the signal is either lost, attenuated beyond detectability, or is reflected back to the transducer. In order to model the sonar signal, a number of assumptions must be made. These are stated here explicitly.

The model is premised on several assumptions about the behaviour of sonar signals in normal circumstances.

1. It is assumed that predominantly specular reflection will occur of the sonar signal from surfaces, and that floors and ceilings can safely be ignored. The degree to which the former assumption is valid depends on the surface texture and the wavelength of the sonar. For the Polaroid unit's typical frequencies, reflection from most wall surfaces is highly specular. Reflection from some surfaces, such as carpets, is not. The latter assumption regarding floors and ceilings may not always be valid. In this case, the two-dimensional simulation may be run on different cross-sections of the environment to recover reliable range readings. For example, the simulation could be run on one horizontal and one vertical cross-section intersecting the transducer, or (with slight modification) on several parallel cross-sections of the environment. Although this increases the running time of the simulation, running time on a single cross-section is sufficiently small that this would not pose significant difficulties. Given this assumption, the model can be considered as a two dimensional process, rather than as a more complex and time-consuming three dimensional one.

2. As is commonly the case, it is assumed that the sonar system responds only to the first occurrence of the reception of a signal over the threshold

strength, in a given trial.

3. It is assumed that the rangefinder incorporates an amplifier whose gain increases with time to compensate for the dispersion of the signal into space (this is also consistent with commonly-used technology [12]).

The results from Kuc and Siegel [5] that are used to set parameters of the model are related to the transducer impulse response and the ratio of the strengths of a specularly reflected signal and a signal diffracted from a straight edge. Under the assumption that the objects reflecting the sonar pulse of wavelength λ are at distances z much larger than $\frac{a^2}{\lambda}$ from a circular transducer with radius a, the impulse response of a transducer at angle α to the sonar wavefront is given by

$$h_R(t, z, a, \alpha) = \frac{2c * \cos\alpha}{\pi a * \sin\alpha} \sqrt{1 - \frac{c^2\,(t - 2z/c)^2}{a^2 \sin^2\alpha}} \tag{1}$$

for c the speed of sound in the environment, $t\epsilon \left[\frac{2z - a*\sin\alpha}{c}, \frac{2z + a*\sin\alpha}{c}\right]$, and $0 < |\alpha| \leq \frac{\pi}{2}$. Note that this expression is replaced with the delta function $\delta(t - 2z/c)$ when $\alpha = 0$ (the wavefront leaves/hits all parts of the transducer simultaneously).

When h_R is convolved with the sonar pulse waveform (e.g. a Gaussian-modulated sine wave), the degree to which the sonar signal is attenuated is determined as a function of the angle α at which the resulting signal leaves (or arrives at) the transducer. Figure 5 shows the result of the convolution, and a plot of signal attenuation as a function of transducer angle.

The second result of Kuc and Siegel that is used to parameterize the model concerns the relative strength of the signal due to diffraction from corners. Modelling corners as line sources of echoes with cylindrical wavefronts, the

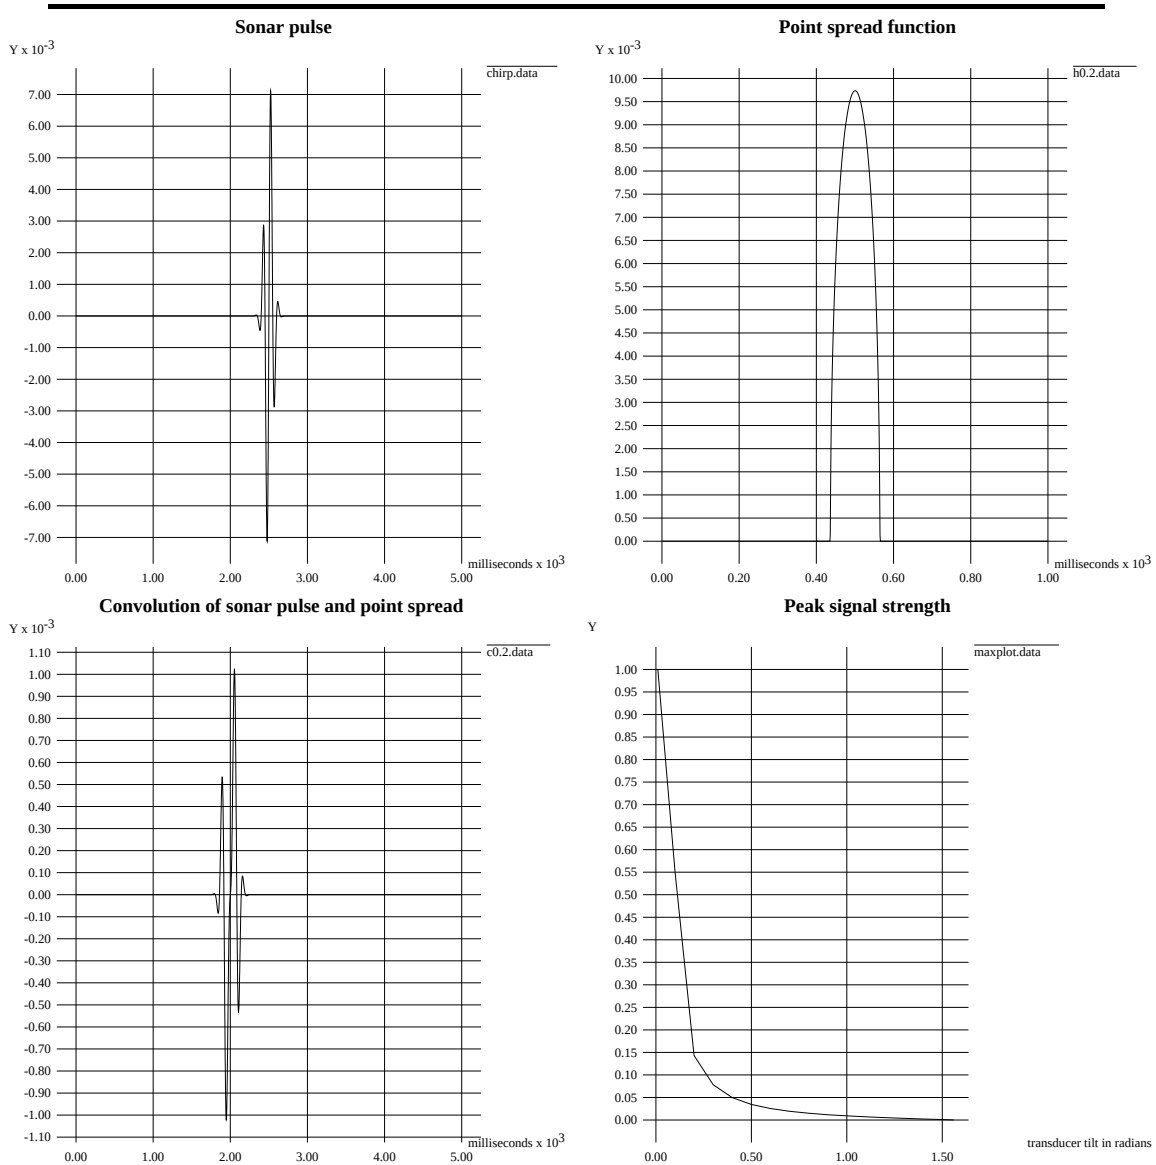

Figure 5: Convolution of transducer impulse response with sonar pulse.

relative strength is $\left(2\pi\sqrt{\frac{z}{\lambda}}\right)^{-1}$.

The sonar model is parameterized by the sonar frequency, the minimum and maximum measurable ranges, and the signal detection threshold relative to the strength of the outgoing signal (after gain amplification). From the 2D assumption, the environment can be modelled as a set of line segments with specified reflectance coefficients from which incident sonar signals are reflected specularly or diffracted. The recovery time of the transducer from the transmission transient is incorporated using the minimum range parameter. This time also gives a minimum distance consistent with the assumption that $z >> \frac{a^2}{\lambda}$. Maximum range accounts for limits on amplifier gain, and transducer sensitivity.

The simulation proceeds by generating a finely spaced fan of rays leaving the front of the transducer, and following their progress in the environment until they again hit the front face of the transducer or until they are sufficiently attenuated. Figure 14 illustrates this. For a single range measurement, the following steps occur. First, the fan of rays is generated from the transducer with an signal strength distribution that is a function of emittance angle. The intersections of this initial fan of rays with walls are computed. These intersections are added to a heap with the closest at the top. Subsequently, the closest (and hence earliest) intersection is examined. If it is within the ray-spacing of a corner, a diffracted ray is generated in the direction of the receiver (since it propagates in all directions). Otherwise, the direction of the reflected ray is computed. In both cases the current intersection is replaced in the heap with the new ray's next intersection with a model wall, and the signal strength is attenuated appropriately. The new closest intersection is then examined as the process repeats. This terminates when an intersection with a small line

segment representing the receiver is encountered as the earliest intersection in the heap. If the arriving signal has a sufficiently small angle of incidence (so that the further attenuation of the signal does not render it indetectible) then half the cumulative distance is returned as the range measurement. If no reception occurs within the maximum distance, then the maximum distance is returned as the range value. In computing the receiver line length and distance from model corners at which to generate a diffracted ray, allowance is made for the spacing of the model rays, so that features are not missed as a result of being between model rays. A high level description of the algorithm is given in the appendix.

5. Experimental Results and Model Validation

In order to validate the model against real sonar data, sonar responses obtained with the simulation have been compared with real sonar measurements obtained with two different sonar scanners, and also with measurements reported in the literature.

Figure 6 shows the results of running the simulation on a model of the room used by Kuc and Siegel to validate their predictions[5]. The data in Kuc and Siegel includes range measurements as a function of transducer orientation for a small sample room. Kuc and Siegel accurately account for a portion of this data, but their model does not account for those readings which are due to interactions of the signal with multiple surfaces. The experiments reported herein include the single-bounce reflectance measurements obtained by Kuc and Siegel (which are consistent with the actual sonar data), but also reproduce the multiple-reflection readings from the actual sonar data in Kuc and Siegel.

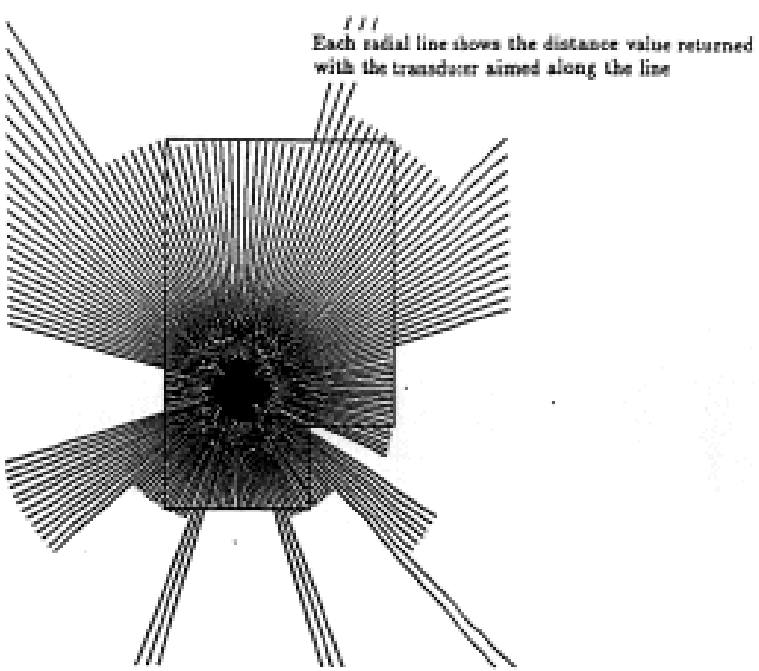

Figure 6: Validation on Kuc and Siegel's data.
The horizontal and vertical lines indicate the positions of the simulated walls of a room. The radial lines indicate the range measurements predicted as a function of angle.

The use of sonar data for robotic exploration is further complicated by the variable reflectances of various objects in the environment. Different materials attenuate sonar signals in different ways. The simulation reported here for sonar range measurement readily allows for surfaces with various isotropic sonar reflectance properties. By modelling this aspect of real surfaces in addition to multi-bounce effects, it is possible to replicate much of the complex nature of sonar sensing. More detailed experiments with an orientable sonar device have also been performed. Many of the peculiarities of the data demonstrated in the simulation show up with surprising consistency.

A rigid room, from within which we could make more controlled measurements, was constructed. The controlled room was constructed from shelving material and consisted of an enclosed rectangular region 928mm x742mm. Figure 7 shows the response for the sonar unit displaced 397mm from the top edge of the room and 94mm from the right edge. Note first that a number of sonar responses appear in a circle around the sonar unit. This circle is a result of the minimum time that the sonar unit can take to listen for a reflected sound. For responses on this circle, the sonar pulse repeatedly reflects off of structure in the environment and is finally sensed by the transducer just as the transducer switches to receiving mode. Note that the far wall is obtained as a wavy line. This is an artifact of the coincident nature of the two corners and A sonar simulation of this (shown in Figure 8) the far wall. A sonar simulation of this (shown in Figure 8) illustrates how the corner reflections (which appear as walls tangential to the corner) and the straight wall interact to form a wiggle on the far wall. In the simulated room sonar responses are drawn as vectors, originating from the sonar unit and terminating at the timeout distance for the unit, or at a distance representing the distance reported by the sonar unit. Note that the right hand sides of the two figures do differ. The simulated room

Figure 7: Actual sonar responses for the simple room.
The endpoints from successive scans are connected to indicate the region that would be inferred. Several scans are overlayed to indicate the degree of repeatability.

returns infinite responses while the real sonar unit returned "very close" responses. The sonar simulation used here [14] while accurately modeling many of the effects found in sonar measurements, does not accurately model all of the complexities of true sonar measurements.

Note that these responses are not to the same scale as the real sonar data.

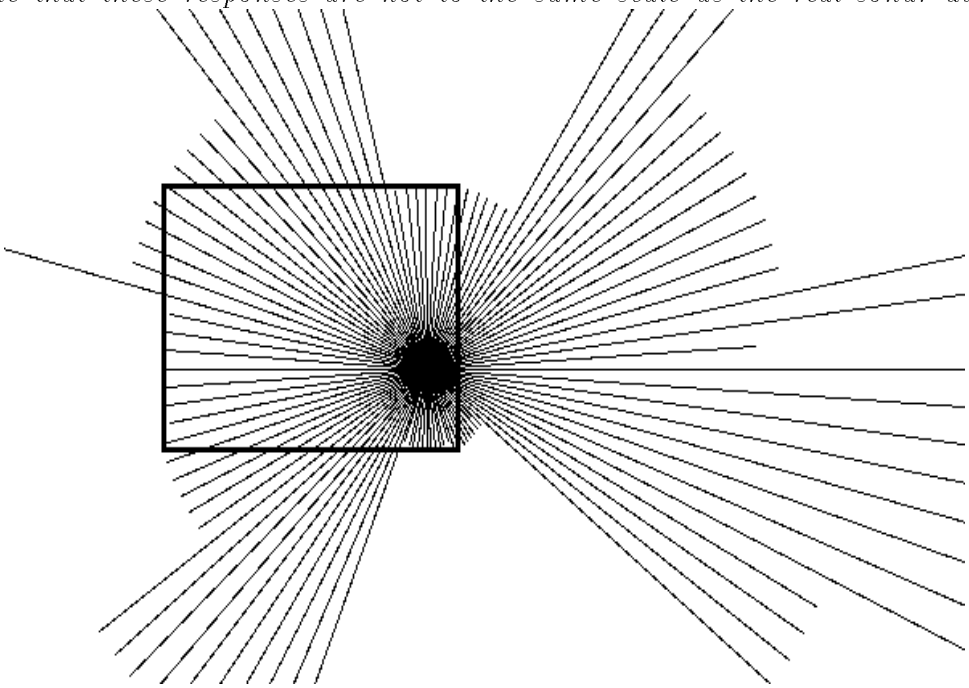

Figure 8: Simulated sonar responses for the simple room
Distance measurements are shown by lines. Non-returning echos by very long lines.

The sonar unit was then moved close to the bottom edge of the simple room (657mm from the top edge and 154mm from the right edge), and repeated the sonar scans. The resulting scan is shown in Figure 9. Once again, corners generated interesting effects and the far wall appears wavey. In addition to these two effects, surprising responses were obtained from the closer wall. The

sonar unit was now far enough away from the upper wall to obtain a number of responses along that wall, but a false lower wall has also appeared. This false lower wall is reported as the sonar signal bounces off the closer lower wall, past the sonar unit, off the far wall, past the rear of the sonar unit, off the near wall, and then back to the transducer. This is a terrible situation for a robot using the sonar sensor. Rather than "seeing" the close wall, the robot sees itself roughly in the centre of a corridor twice the true width. Worse still, the robot sees itself being slightly closer to the real wall than it is to the phantom wall. A robot which was designed to translate down the centre of the hallway using sonar data would correct itself directly into the closer wall.

For a practical comparison, this problematic effect was also investigated with a RWI (Real World Interface) B12 Robot equipped with a sonar ring. In this case, we found that the bulk of the robot is sufficient to block the echo from the near wall and the chirp simply reflects between the body of the robot and the sensor until the signal can be detected at the minimum sensor distance. In addition, the body of the robot may block the returning sonar pulse. (Our simulator[14] does not model this particular artifact as it simulates a robot with a smaller "body").

Corners provide many interesting effects for the interaction of the sonar signal and the two wall surfaces making up the corner. A number of effects are possible, and two of these are shown in Figures 10 and 12. In these displays, the robot is shown as a small circle, and the position of sonar responses are indicated by small squares. Both of these figures show sonar responses from the RWI mobile platform taken every 3 degrees. In both cases, a corner was constructed from rectangular obstacles placed in the lower left hand corner of the display. The two obstacles meet corming a corner. Simulations of these two cases are shown in Figures 11 and 13. Note that in the first case, the sonar

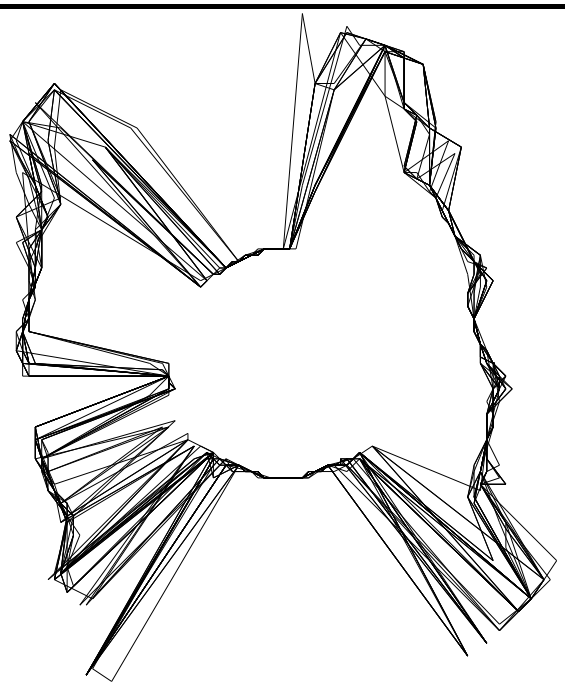

Figure 9: Sonar responses for the simple room near the bottom wall

data appears to show two separate walls which do not meet, and in fact the lower of the two walls is actually seen behind the true position of the closer (upper) wall. This is clearly seen in the simulated measurements. The second case shows the classical corner reflection effect in which the corner itself is lost and is replaced by a wall tangential to the corner, and two holes appear surrounding this "corner wall". In both of these cases, the local wall structure appears to have openings which do not exist at all. These effects remain almost completely unchanged as the robot translates towards the corner. Once again, the sonar errors are not uncorrelated, but rather indicate a highly structured and non-linear behavior.

6. Discussion

The modelling of sonar measurement as straight line distance measurements with simple Gaussian noise is excessively simplistic to capture most real phenomena that occur with these sensors. If sonar sensor simulations are to be used, and there are many potential advantages to using them, then the simulations must agree with not only the physics of the sensor being used, but interactions between the sound chirp and structure in the environment must be correctly modelled.

Failure to model either or both of these effects will result in sonar simulations which are much cleaner than the data that can be recovered. Algorithms which are based on simulated data which lacks this required realism may fail when presented with real data. Holes may appear in corners (when they are not really there), and walls in hallways may be further away than they seem. In addition, simulations exist (for example, [14]), which are capable of simulating many of the effects found in sonar measurements.

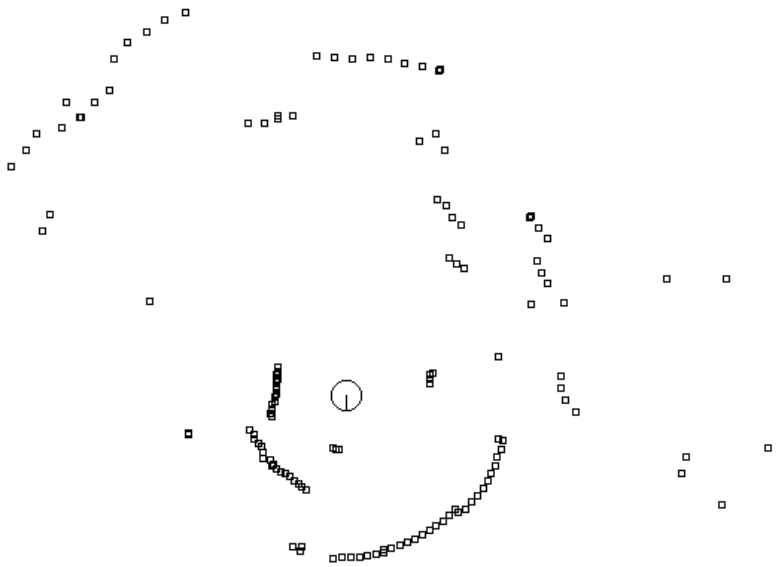

Figure 10: Actual responses for an oblique corner.
The distance measurements returns from the sonar sensors are shown by rectangles. The responses from the corner are the almost-linear clusters of readings just below and to the left of the robot (see following figure). The rest of the responses (above and to the right of the robot) are from other objects in the environment and are included only for completeness.

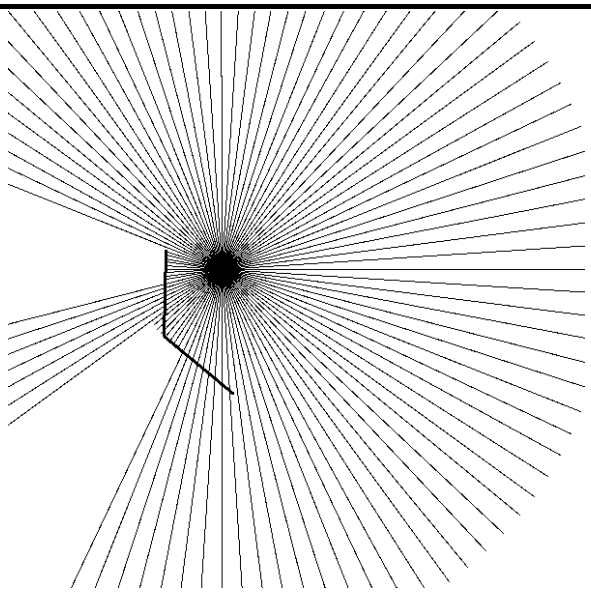

Figure 11: Simulated responses for an oblique corner
Distance measurements are shown by lines. Non-returning echos by very long lines.

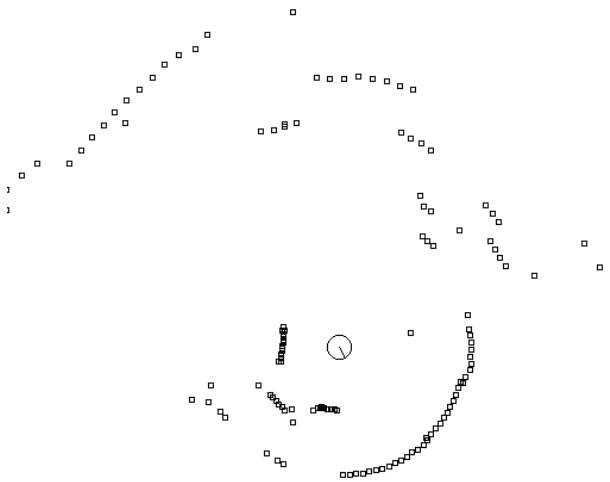

Figure 12: Actual sonar responses for a right-angle corner
The robot is indicated by a circle. The distance measurements returns from the sonar sensors are shown by rectangles. The responses from the corner are the almost-linear clusters of readings just below and to the left of the robot (see following figure). The rest of the responses (above and to the right of the robot) are from other objects in the environment and are included only for completeness.

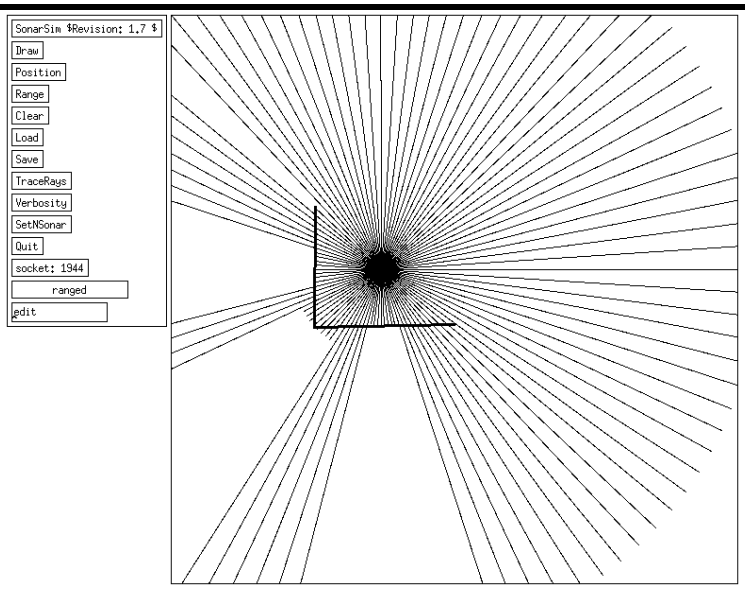

Figure 13: Actual and Simulated responses for a right-angle corner
Distance measurements are shown by lines. Non-returning echos by very long lines.

The model described here allows for the realistic simulation of sonar range sensing in real environments. As shown, it is consistent with real data collected from such devices. The major remaining limitation of this model is that it uses only a planar model of the world. As such, it fails to capture bounces of the sonar signal that might come from interactions with objects outside this plane. Experimentation with the sonar device and the simulation suggests that in many environments sufficient accuracy may be obtained by projecting all objects near the horizontal plane scanned by the device onto this plane.

The availability of realistic, effective and efficient simulation permits rapid and repeatable trials of experimental scenaria. Experimental work is essential to the enumeration of realistic problems in sensing. On the other hand, the ability to examine rapidly particular situations, and more importantly to repeat exactly a series of (simulated) measurements is critical to the rigorous scientific development of sensing strategies. In addition, this allows a large suite of test cases to be examined rapidly. It is expected that this will facilitate the development and evaluation of robust sensing methodologies.

Methods for efficiently inferring the reflecting-surface geometry from sonar range measurements are currently under investigation. It is hoped that the deeper understanding of these measurements acquired through efforts to simulate them will allow conversion of aspects of the ranging data that have long been considered "errors" into useful features that can be part of the interpretation process.

Appendix A. High level description of the sonar simulation

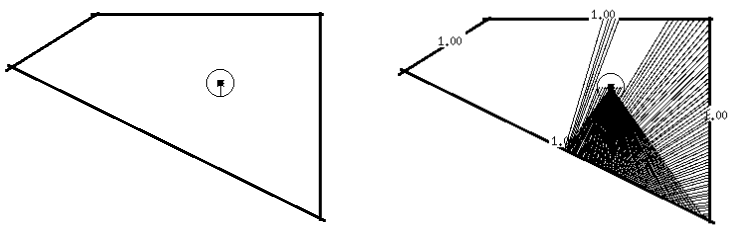

Figure 14:A single range measurement. Left: the room with robot. Right: some rays generated by a single range measurement. The numbers are relative wall reflectances.

The following gives a high level description of the algorithm

```
range(θ, x, y):  determine the range returned by a sonar measurement
    from position (x, y) of the sonar sensor facing direction θ

    initialize heap of intersections to be empty
    for each angle φ from θ − π/2
        to θ + π/2 in steps of size δ
        determine:  the first intersection I of a ray leaving (x, y)
        at angle φ with a model wall segment, the degree of
        attenuation of the signal as a function of |φ − θ|,
        the total distance travelled from the transducer, and angle
        of incidence at the intersection.
        add I to the heap (ordered by distance travelled
            from transducer)
    while the heap is non-empty
        if the intersection I at the top of the heap is at a
```

```
            distance d > maxdistance from the transducer
            return(maxdistance)
        if the intersection is with the model segment
            representing the transducer
            compute the further attenuation of the
            incoming signal as a function of the angle
            of incidence
            if the remaining signal strength is above
                detection threshold
                return(total distance travelled / 2)
        generate the reflected ray r leaving the intersection,
            based on the angle of incidence of the incoming ray
            and reduce its amplitude as a function of the
            characteristics of the segment
        if the intersection is within a single ray-spacing
            of the end of the model segment involved
            generate a diffracted-ray intersection with the
            receiver also, with appropriate signal strength and
            angle of incidence
        add the closest intersection between r and a
        model segment to the heap
    end while
```

Acknowledgements

Funding for this work was provided by the National Research Council of Canada and by the ARK Project which receives its funding from PRECARN Associates Inc., the Department of Industry, Science and Technology, the National Research Council of Canada, the Ontario Technology Fund, Ontario Hydro, Atomic Energy of Canada Limited, and the Government of Ontario inter-provincial collaboration program

Bibliography

[1] C. Biber, S. Ellin, E. Shenk, and J. Stempeck. The polaroid ultrasonic ranging system. In *The 67th Convention of the audio engineering society*, New York, 1980.

[2] J. Borenstein and Y. Koren. Histographic in-motion mapping for mobile robot obstacle avoidance. *IEEE Transactions on Robotics and Automation*, 7(4):535–539, 1991.

[3] J. Crowley. Navigation for an intelligent mobile robot. *IEEE Journal of Robotics and Automation*, 1(1):31–41, March 1985.

[4] M. Drumheller. Mobile robot localization using sonar. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(2):325–332, 1987.

[5] R. Kuc and M. W. Siegel. Physically based simulation model for acoustic sensor robot navigation. *IEEE Transactions on Pattern Analysis and Machine Intelligence,*, 9(6):766–768, 1987.

[6] S. Lange, L. Korba, and A. Wong. Characterizing and modelling a sonar ring. In *Proc. SPIE Mobile Robotics IV*, pages 291–304, Philadelphia, 1989.

[7] J. J. Leonard and H. F. Durrant-Whyte. Mobile robot localization by tracking geometric beacons. *IEEE Transactions on Robotics and Automation*, 7(3):376–382, 1991.

[8] R. Malik and E. T. Polkowski. Robot location densities. In *Proc. SPIE Mobile Robotics V*, pages 280–290, Boston, 1990.

[9] P. McKerrow. Simulation of sonar-echolocation. In *IARP 1st Workshop on domestic robots*, pages 127–136, Newcastle-upon-Tyne, 1989.

[10] P. McKerrow and J. Hallam. An introduction to the physics of echolocation. In *Third national conference on robotics*, pages 198–209, Melbourne, 1990.

[11] H. Moravec. Sensor fusion in certainty grids for mobile robots. In A. Casals, editor, *Sensor Devices and Systems for Robotics, NATO ASI Series, Vol. F52*, pages 253–276. Springer-Verlag, Berlin, 1989.

[12] Polaroid Corporation. *a: Untitled product spec sheet on Ultrasonic Ranging System. b: "Ultrasonic Ranging System" Users Manual.*

[13] Y. Roth-Tabak and T. Weymouth. Environment model for mobile robotics indoor navigation. Technical Report CSE-TR-53-90, The University of Michigan, 1990.

[14] D. Wilkes, G. Dudek, M. Jenkin, and E. Milios. A ray following model of sonar range sensing. In *Proc. Mobile Robots V*, pages 536–542, Boston, 1990.