A Multi-Surface Model of Sonar Range Sensing

David Wilkes[1], Gregory Dudek[2], Michael Jenkin[3], and Evangelos Milios[1,3]

[1] Department of Computer Science, University of Toronto
Toronto, Canada M5S 1A4.

[2] McGill Research Center for Intelligent Machines
McGill University, 3480 University St.
Montreal, Canada H3A 2A7

[3] Department of Computer Science, York University
Downsview, Canada

ABSTRACT

We present 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 model also accounts for environments with surfaces having different reflectances. This gives more realistic results than previous models. The approach is based on simulation of the reflection and diffraction of sonar rays off reflecting surfaces until they are attenuated 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) and environmental characteristics.

Keywords:
Sensing, autonomous navigation, active vision, sonar, sonar simulation.

1. Introduction

Sonar range sensing refers to technologies for emitting sound signals and measuring the characteristics of the echo that returns. The commonest 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.

In work on algorithms for robotic exploration of unknown environments[2,3] a crucial problem is that of recovering the shape of the open areas of an environment. Performing this task accurately using sonar sensors is somewhat problematic due to the complex reflection, diffraction and echo phenomena which can occur. In order to develop algorithms for reliable shape recovery, we need both to understand and to be able to generate synthetic as well as real sonar data. The generation of realistic synthetic data facilitates the analysis of the algorithms in specialized scenarios. The performance of sonar-based methods can be studied in a controlled, repeatable framework. Furthermore, alternative environments can be simulated quickly and easily.

The class of sonar sensor we deal with here operates by measuring the time required for a pulse of ultrasound to travel from a transmitter, bounce off an object, and return to approximately the same point. The transmitter and receiver are often the same transducer. 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 in documents from Polaroid[4] and Biber et al.[5]. It operates using sound signals with frequencies of roughly 50 kHz.

There are several reasons why false range values may be inferred by such systems. The return signal detected by the system may be the result of multiple reflections of the sound pulse, or diffraction from edges of objects, or reflection of part of the wavefront off the line along which the transducer is aimed.

Kuc and Siegel[1] have done a careful analysis of the detailed physics of signal transmission, reflection, diffraction and reception. They give results on the behavior of sonar signals at right angled corners and walls. Their simulation accurately matches actual sonar measurements from such features, provided that the signal returns to the receiver without multiple intervening reflections. In real environments, however, phenomena such as multiple reflections, amplifier characteristics, and the room shape play a significant role. In order to generate a realistic simulation of sonar sensing, we have incorporated these phenomena. The result is a model that uses an approach somewhat akin to the "ray-tracing" models used in the computer graphics community. Sample signal paths are followed from transmitter, through bounces and diffraction, to the receiver.

Details of the approach follow in the next section followed by a discussion of our particular implementation. The fourth section summarizes the validation of the model. Finally, we discuss possible extensions to the model.

2. Model

The model is premised on several assumptions about the behavior of sonar signals in normal circumstances. First we assume predominantly specular reflection 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. The latter assumption regarding floors and ceilings may not always be valid. In this case, we may run our two-dimensional simulation 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. As is commonly the case, we assume that the sonar system responds only to the first occurrence of the reception of a signal over the threshold strength, in a given trial. Also, we assume that the rangefinder incorporates an amplifier whose gain increases linearly with time to compensate for the dispersion of the signal into space (this is also consistent with commonly-used technology[4]).

The results from Kuc and Siegel[1] that are used to set parameters of our 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

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

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 h_R is just the delta function $\delta(t - 2z/c)$ when $\alpha = 0$ (the wavefront leaves/hits all parts of the transducer simultaneously).

When we convolve h_R with the sonar pulse waveform (e.g. a Gaussian-modulated sine wave), we can determine the degree to which the sonar signal is attenuated as a function of the angle α at which the resulting signal leaves (or arrives at) the transducer. Figure 1 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 our model concerns the relative strength of the signal due to diffraction from corners. Modelling corners as line sources of echoes with cylindrical wavefronts, the relative strength is $\left(2\pi\sqrt{\frac{z}{\lambda}}\right)^{-1}$.

Our model parameters are the sonar frequency, minimum and maximum measurable ranges, and the signal detection threshold relative to the strength of the outgoing signal (after linear-gain amplification). The environment is modelled as a set of line segments with specified reflectance coefficients from which incident sonar signals are reflected specularly or diffracted. Minimum range accounts for the recovery time of the transducer from the transmission transient. Maximum range accounts for limits on amplifier gain.

3. Implementation

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 2 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 line, 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 undetectable) 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.

Figure 3 is a very high-level statement of the algorithm.

4. Experimental Results

Figure 4 shows the results of running our simulation on a model of the room used by Kuc and Siegel to validate their predictions. Their published data includes range measurements as a function of transducer orientation for a small sample room. They accurately account for a portion of this data, but their model does not account for

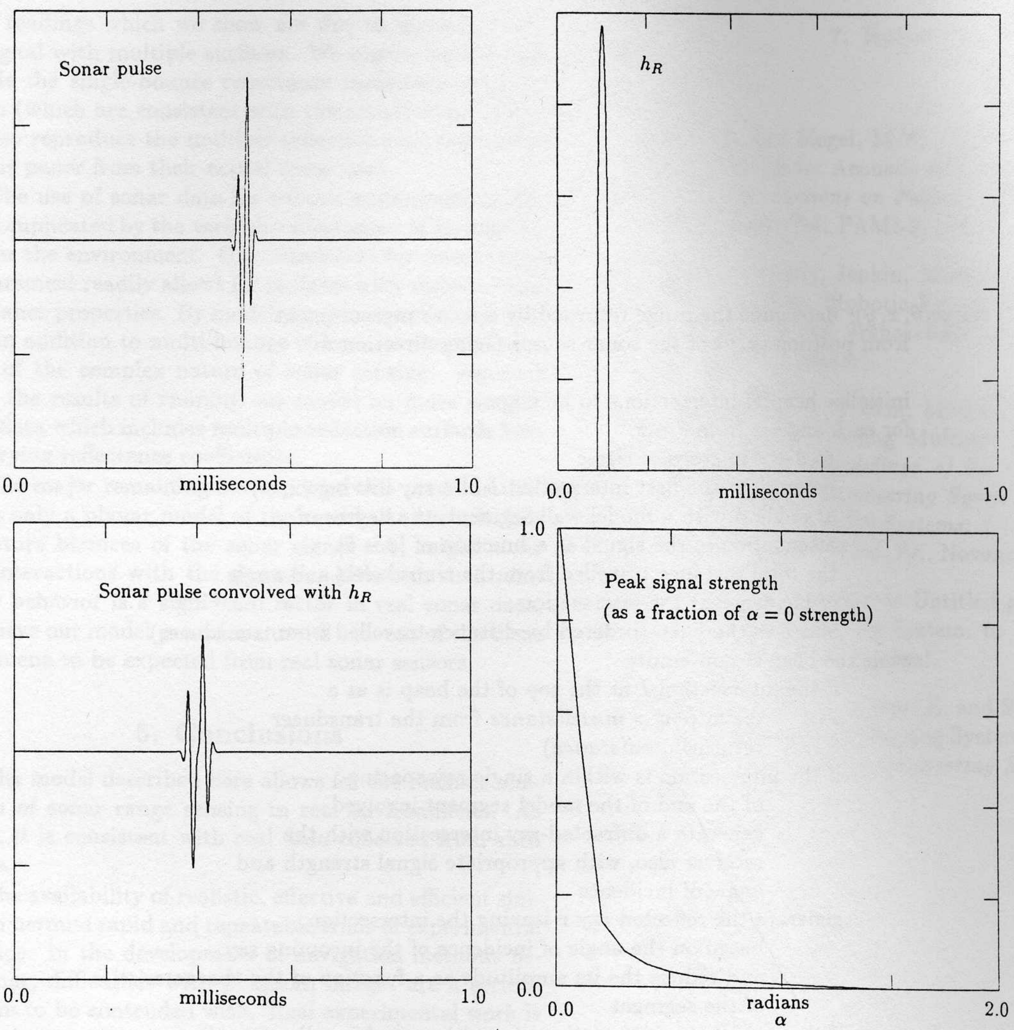

Figure 1: Convolution of transducer impulse response with sonar pulse

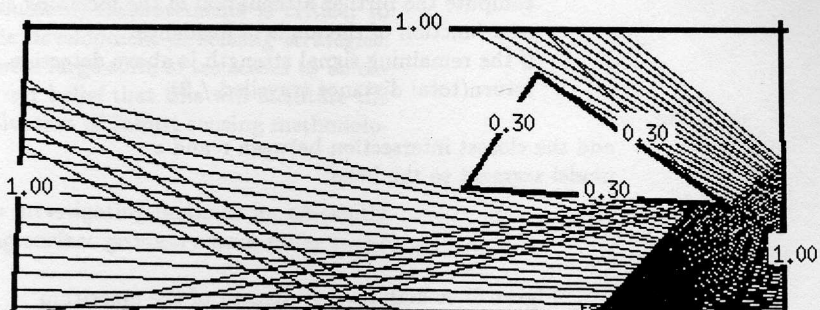

Figure 2: A single range measurement

```
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
                to θ + π 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 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
            generate the reflected ray r leaving the intersection,
                based on the angle of incidence of the incoming ray
                and reduce the its amplitude as a function of the characteristics
                of the segment
            if there exists an intersection of r with a model wall segment
                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)
            else
                add the closest intersection between r and a
                model segment to the heap
        end while
```

Figure 3: A high-level statement of the algorithm

those readings which we show are due to interactions of the signal with multiple surfaces. We obtain results that include the single-bounce reflectance measurements they obtain (which are consistent with the actual sonar data). We also reproduce the multiple-reflection readings shown in their paper from their actual sonar data.

The use of sonar data for robotic exploration is further complicated by the variable reflectances of various objects in the environment. Our simulation for sonar range measurement readily allows for surfaces with various sonar reflectance properties. By modelling this aspect of real surfaces in addition to multi-bounce effects, we can replicate much of the complex nature of sonar sensing. Figure 5 shows the results of running our model on more complicated data which includes multiple-reflection surfaces having varying reflectance coefficients.

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 the ground. Although such non-planar behavior is a significant factor in real sonar data, we believe our model provides a useful account of the basic phenomena to be expected from real sonar sensors.

5. Conclusions

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 availability of realistic, effective and efficient simulation permits rapid and repeatable trials of experimental scenarios. In the development of navigation methods using sonar, difficulties with phantom images are a serious problem to be contended with. Real experimental work is essential to the enumeration of realistic problems in sensing. On the other hand, the ability to rapidly examine 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 our belief that this will facilitate the development and evaluation of robust sensing methodologies.

We are currently investigating methods for efficiently inferring the reflecting-surface geometry from sonar range measurements.

6. Acknowledgements

The support of the Natural Sciences and Engineering Research Council of Canada and of the Government of Ontario through the Information Technology Research Centre is gratefully acknowledged. We also gratefully acknowledge the thoughtful comments of the anonymous reviewers.

7. References

1. Kuc, R. and Siegel, M.W., "Physically Based Simulation Model for Acoustic Sensor Robot Navigation," *IEEE Transactions on Pattern Analysis and Machine Intelligence,* Vol. PAMI-9, No. 6, pp.766-78. 1987.

2. Dudek, Gregory, Jenkin, Michael, Milios, Evangelos, Wilkes, David, "Robotic Exploration as Graph Construction," *IEEE Transactions on Robotics and Automation* (to appear).

3. Dudek, Gregory, Jenkin, Michael Milios, Evangelos, Wilkes, David, "Using Multiple Markers in Graph Exploration," *Proceedings of the International Society for Optical Engineering Symposium on Advances in Intelligent Robotics Systems: Conference on Mobile Robotics*, Philadelphia, PA, November 1989.

4. Polaroid Corporation, a: Untitled product spec sheet on Ultrasonic Ranging System. b: "Ultrasonic Ranging System," Users manual.

5. Biber, C., Ellin, S., Shenk, E. and Stempeck, J., "The Polaroid Ultrasonic Ranging System," *The 67th Convention of the Audio Engineering Society*, New York. 1980.

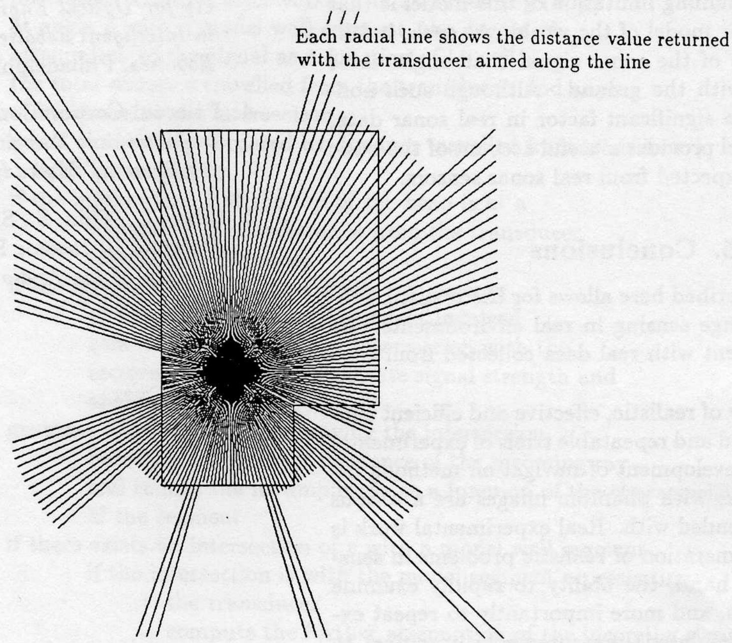

Figure 4: Validation on Kuc and Siegel's data

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.

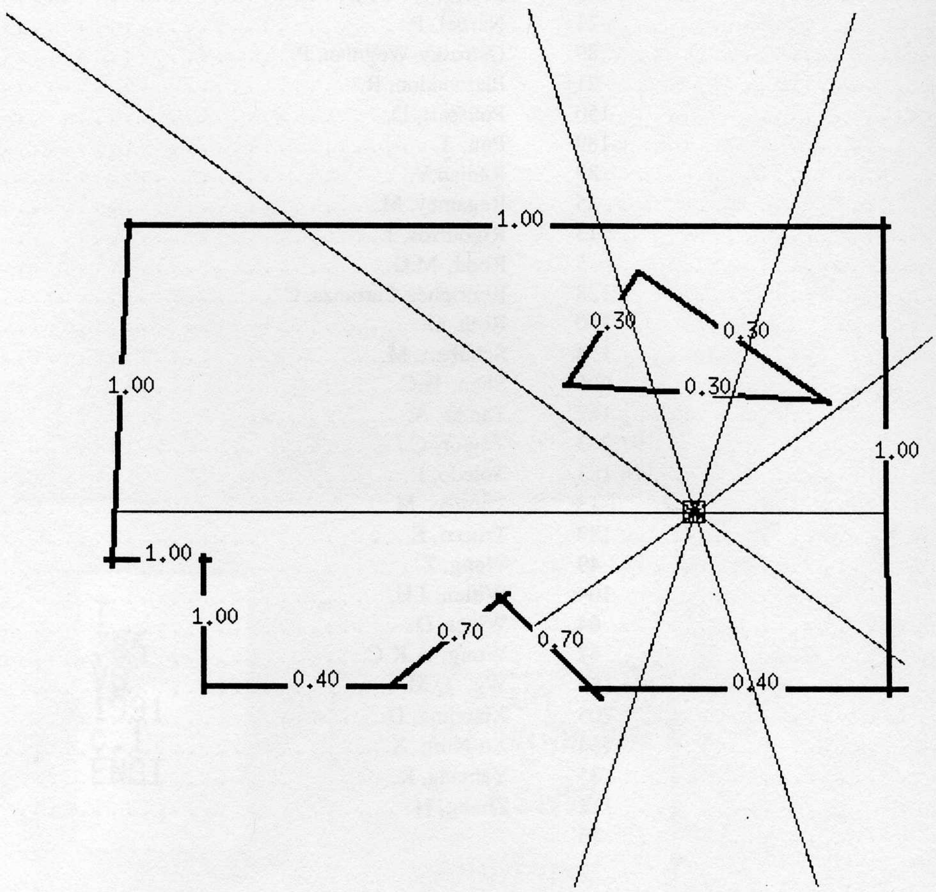

Figure 5: A more complicated example
The relative reflectance of each modelled wall is shown by a number appearing over it.