A Sonar Sensor with Random Double Pulse Coding

Andrew Heale and Lindsay Kleeman

Intelligent Robotics Research Centre
Department of Electrical and Computer Systems Engineering
Monash University, Victoria 3800, Australia
Andrew.Heale@eng.monash.edu.au Lindsay.Kleeman@eng.monash.edu.au

Abstract

Double pulse coding is used to give sonar sensors a distinctive 'voice', allowing several to operate concurrently in the same acoustic space. This paper describes the use of random spacing between the two pulses to overcome interference problems arising when the same spacing occurs between targets or is used by another sensor. Results demonstrate the robustness of this method in comparison with previous approaches. The echoes are digitised by a 1 MHz, 12 bit ADC and processed on a custom DSP system.

1 Introduction

As use of sonar sensing methods increases [1-6, 8-11], situations arise where multiple sensors must share the same airspace. When several sensors, whether on one robot or many, must work in the same acoustic environment, crosstalk occurs [1, 2, 9]. With simple sonar pulses, a sensor cannot determine whether a signal it receives originated from itself or some other source, resulting in many incorrect measurements. Two robots attempting to pass each other in the corridor could become hopelessly confused unless, by some means, each is able to distinguish its own sonar signals.

Other researchers have attempted to solve the interference problem using random inter-sensing times. For example, Borenstein [1] has employed alternating firing patterns to statistically eliminate crosstalk problems. These techniques require more than one measurement cycle to reject interference. Pseudo random pulse sequences have been used in sonar [2], but incur a significant processing overhead and require long pulse sequences.

To reject sonar interference, each sensor may be given a distinctive 'voice' by transmitting two pulses with a specific interval separating them *within one sensing cycle*. This is known as double pulse coding [9]. Only if two pulses are received with the correct separation and matching amplitude are they accepted. Pulses without a partner of the same amplitude and specific separation are discarded. This method can reject most signals from other sonar sources.

In our implementation in this paper, random double pulse separations down to about 100 microseconds (approaching the duration of the pulse itself) and up to several milliseconds can be used. A difference as small as 1 microsecond can be distinguished, although differences, proportional to the pulse separation, of several microseconds can be introduced by movement of targets or the sensor.

The paper is organised as follows. Section 2 introduces the double pulse coding approach and the signal processing required for implementation. Section 3 describes a pulse overlap problem when the environment contains two targets at a critical separation. Assigning pulse spacings is discussed in section 4. The sensor hardware and signal processing are described in sections 5 and 6, whilst experiment results that validate the random pulse approach are presented in section 7. Results are presented from multiple targets in the environment and also from two interfering sonar systems, where one system is building a map of a wall. Conclusions are presented in section 8.

2 Double Pulse Coding

Previous work [9] validates echoes by finding any two returned pulses with approximately the required separation and the same shape. This is achieved by a simple processing approach whereby the maximum difference between pulse waveforms, taken sample by sample, must fall below a threshold to be validated. Doppler-like effects slightly change the separation of returns from a moving target, and since the robot itself moves, many returns exhibit this effect. To allow for this variation, the two returns may be shifted slightly relative to each other and tested again.

This work uses a similar algorithm. The chief difference is the use of correlation rather than difference, and the result is rather like using the first return to define a matched filter that will identify the second.

Pairs of returns with approximately the required separation are selected, each one's correlation coefficient is tested against a threshold, and the waveforms may be shifted to compensate for Doppler effects. If shifting is used, the maximum correlation coefficient is estimated by interpolation. Additionally, the amount of shift required corresponds to the radial velocity of the target relative to the sensor; although we presently make no use of this, it could be valuable for path planning in busy environments, or to exclude moving objects from map building.

3 Pulse Overlap due to Target Spacing

Sonar is poor at distinguishing two targets that are very close to each other in range - even with matched filtering the interference between the two echoes makes accurate arrival time estimation difficult or impossible due to overlapping pulses. Of course, the more airborne sound the greater potential for pulse overlap, and with double pulse coding an echo of the first pulse might overlap an echo of the second. For example, if the separation between two transmitted pulses is 1 ms, and there are two targets separated by 17 cm range, interference will occur because the time of flight, two ways, over the 17 cm is about 1 ms.

4 Choosing a ‘voice’

There are a number of ways each sensor could be assigned a particular pulse separation, ranging from preset to random. However, they all have the same two flaws: a permanent ‘blindness’ to targets with a corresponding spacial separation, and the possibility that at some time, two sensors operating with the same voice will encounter one another. A dynamic solution is required.

We choose a random pulse separation for each each sensing cycle - because it is simple. Both the ‘blindness’ and ‘same voice’ problems remain, but in most cases they will not persist over repeated firings of the sensor. The occasional spurious results when two sensors use the same separation are left for higher level processing, which must handle false returns and noise problems anyway. In theory, two sensors could follow the same pseudo-random number sequence – however this is extremely unlikely due to timing differences resulting in lack of synchronisation and also since the random number seed can be generated using noise from, for example, the A/D converter.

5 Hardware

The sensor is self-contained in a box about 15 cm x 10 cm x 7 cm, requiring a 5 V supply and a serial link to a host computer. Polaroid 7000 series transducers are used for two upper transmitters (only one is used in this work) and two lower receivers in a square with 40 mm between centres. See figure 1.

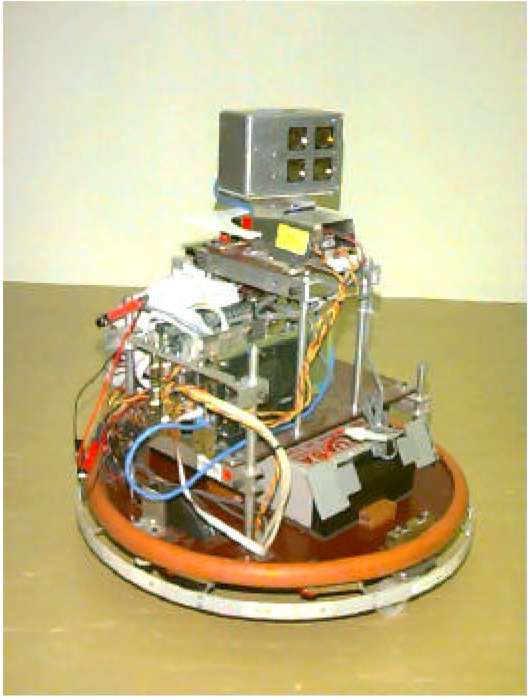

Figure 1 - DSP Sonar mounted on the robot Werrimbi

Two sonar receiver channels are digitised at 1 MHz and 12 bit precision after low pass filtering and amplification. An Analog Devices 33 MHz ADSP2181 is the central controller and workhorse of the system, featuring an on-chip 80 kbyte RAM, allowing echo pulses to be stored on the chip. It is interrupted at 1 MHz for pulse forming and pulse capture operations. A high speed buffered UART interfaces to the serial port.

6 Processing

The processing algorithm is summarised here. The two receiver channels are treated separately, though concurrently.

6.1. Pulse forming and Pulse capture

The DSP is interrupted every microsecond. During pulse forming, it writes directly to the transmit hardware thus forming the pulse. The pulse shape (e.g. length, double pulse separation) can be changed every transmit cycle. During pulse capture, the DSP reads a sample from each receiver.

6.2. Pulse Extraction on the Fly

Any sample exceeding a predetermined threshold is deemed part of an echo, along with a window of samples before and after it. Windows that overlap are merged on the assumption that they are adjacent cycles of a pulse, and stored.

6.3. Pulse splitting

Frequently, two or more returns will overlap, or be merged during step 6.2. If the distortion is not too severe, they are separated based on the relative amplitudes of waveform peaks.

6.4. Double pulse recognition

Pairs of returns with approximately the required separation are compared and if found to match, the first is retained. All other pulses are discarded. Waveforms may be shifted to compensate for Doppler effects, and the maximum correlation coefficient estimated by interpolation.

6.5. Template matching

Time-of-flight is determined by matched filtering [5]. Because pulse shape depends on range and the angle of arrival, several filters are trialled, each being defined by a template, or ideal pulse shape for a particular range and angle. Parabolic interpolation gives a time-of-flight estimate to a fraction of a microsecond accuracy. [5]

6.6. Correspondence

From each receiver channel we obtain a list of arrival times, pulse amplitudes, pulse durations, and correlation coefficients with the template set. Lists from two receivers are associated to obtain bearings and ranges to targets. The process of association solves the correspondence problem. Because the receiver physical spacing is just 40 mm, the correspondence problem is usually easy to solve and is simply based on matching the arrival times and pulse amplitudes within predefined tolerances.

6.7. Triangulation

Bearing estimation to targets is performed using geometric modelled discussed in [5]. This processing is performed by a host computer separate from the DSP sensor.

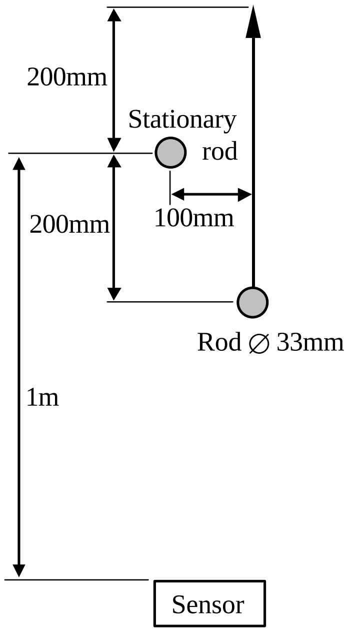

Figure 2a– Configuration of two target experiment.

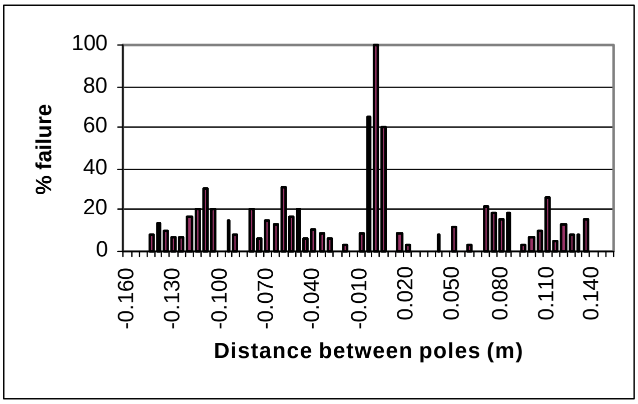

Figure 2b– Interference rate, constant separation

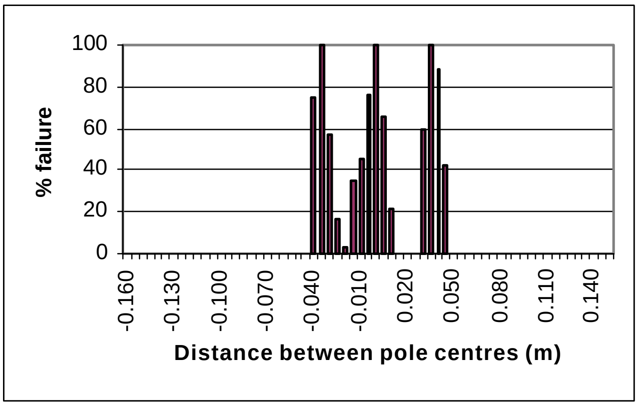

Figure 2c – Interference rate, random separation

7 Experimental Results

7.1 Multiple target interference

The purpose of this first experiment is to compare double pulses having constant separation with those having random separation, in terms of their susceptibility to interference caused by multiple targets.

As in Figure 2a, the sensor is stationary, measuring the positions of two cylinders (wooden rods of 33 mm diameter) upright in front of the sensor. One rod is 1 m from the sensor at about +3 degrees bearing. The second rod starts 20 cm closer and is moved away until it is 20 cm further than the first cylinder; keeping about 100 mm to its right (roughly -3 degrees bearing). As the rod is moved, the sensor makes the measurements, alternating between preset (200 µs) and random (200 µs to 799 µs) separation modes.

The typical length of a received pulse is 60 µs, so two pulses arriving within 60 µs of each other are likely to overlap and be rejected. This corresponds to features with a physical separation of ±10 mm, so it is expected that results from both methods will be poor when the rods are separated by less than 10 mm in range. Similarly, the 200 µs constant separation corresponds to 34.6 mm, so poor results are also expected for this range and method. In contrast, the random separation method should have a lower error rate spread over a wider range of target spacings: from 34 mm to 137 mm the error rate should be "range of separations causing interference ÷ range of separations available" = 120/(800-200) = 20%.

The analysis divided the 40 cm total travel into regions of 0.5 cm, and calculated the failure rate in each region, where 'failure' means that a target was overlooked. No false returns occurred. Figure 2b and 2c show that these results correspond well with expected results. Note that the 'poor results' expected actually manifest as 100% failure rate, indicating total blindness to those targets. The constant separation method has three such regions of blindness, whereas the random separation method has only one. Also the failure rate for the random separation method is somewhat less than the 20% predicted because 60 µs is a generous estimate of pulse length, and because often two pulses with only a small degree of overlap can be separated, enabling them to be correctly received.

7.2 Interference from another sensor

To demonstrate the severity of cross-talk between sensors, two sensors were mounted on the mobile robot Werrimbi, moving parallel to a wall at 1 m range – see figure 3. One sensor trialled three sensing methods to map the wall in the presence of four types of interference generated by the other sensor. The maps produced are presented in a grid in figure 4. Each run comprises about 230 readings. The sensing and interference methods used are: single pulse, double pulse with constant separation (200 µs), and double pulse with random separation (200 to 499 µs, in steps of 1 µs). Additionally, no interference is used to give a baseline for comparison. The maps show the results of each combination of sensor and interference. Good targets are marked with tiny black circles. Targets that are deemed 'unreliable' (for example, if they only barely passed the matched filter test) are grey. Targets that were only detected by one of the two receivers are yellow/pale grey, and although shown at the correct range the bearing may be somewhat wrong. The path taken by the robot is marked by a series of crosses. Scales are in metres.

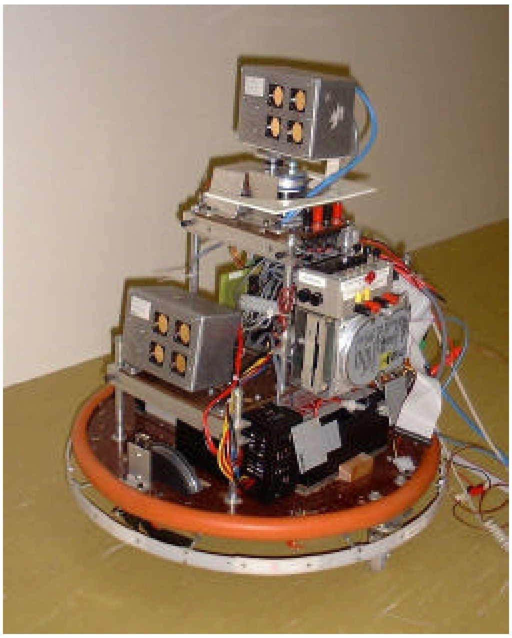

Figure 3 – Two sensors aligned to cause interference.

All the maps show the wall at 1 m, and the corner of the wall with the floor in pale grey immediately beyond it. Most also show several *double echoes* of the wall. The double echoes arise when sound returned from the wall reflects off various parts of the robot, then back via the wall to the receiver. The total path length depends on which part of the robot reflects it, hence several may appear. They are all phantom targets.

The results from the single pulse sensor are all in the left column in figure 4. Clearly, it is highly susceptible to interference. Double pulse coding with constant separation, in the second column, has much better noise rejection [9]. The effect of interference with the same 'voice' is seen in figure 4i (the graph labelled IEX-I.LOG, for 'interference experiment I'). It is this result which motivated development of the random double pulse method.

Double pulses with random separations in the presence of double pulse interference – whether random or not – could be expected to produce false targets about 1 time in 300 in this experiment. This is the probability that the random separation chosen for a given reading will be the same as that currently in use by the interfering sensor. Since each part of this experiment has only about 230 readings, we might not see a false target at all.

The results show many more false targets than this theory predicts – see figures 4m, 4n and perhaps 4k. Even figure 4g, where the interference is only single pulses that should never be accepted, displays two obvious false targets.

This increase in interference must be expected for complex objects, such as a bookcase. When times of flight from two targets differ by the same separation as is in use by the sensor, they effectively forge its signature and are incorrectly accepted as a valid target. In this experiment the complexity arose from the double echoes off Werrimbi itself. When the experiment was repeated with the sensors 2 m from the wall, the double echoes are weakened and the error rate drops to the expected level of about 1 in 300.

Results from the sensor are limited to the five nearest returns; consequently only the five nearest targets can be triangulated, and mapped. This limitation can be seen, for example, in figure 4d. The false targets, apparently sprayed across the map, cease abruptly at around 2 m range. At this time, double echoes arrive from the wall, filling up the last available slots, and any further returns are discarded. For a similar reason, most of the false targets in figures 4h and 4l appear closer than the wall – those apparently further away were lost.

Early observations made while setting up this experiment indicate that the source of interference should be close to the sensor to cause many false targets. Beamwidth effects and the discriminating ability of the matched filter eliminated most false returns. However, interference that causes jamming is less well rejected, and leads to missed targets.

8 Conclusions

This paper has presented a new approach to crosstalk rejection between multiple sonar sensors. The approach is to use randomly spaced double pulses on each sensing cycle. By using random spacing the method can tolerate multiple identical sensors operating in the same environment with error rates better than 1 per 300 measurements. Previous work [9] requires sensors to be configured individually to have different double pulse spacing. Also, two reflectors at a critical spacing can cause ‘blindness’ due to pulse overlap problems. By using random pulse separation, as in this paper, these problems will not occur consistently between measurements. Extensive experiments have been performed to show the effects of different interference rejection strategies. The random pulse separation approach was found to be a practical and reliable technique.

Acknowledgements

Financial assistance of a Large ARC grant and the Engineering technical assistance of Steve Armstrong are gratefully acknowledged.

References

[1] J Borenstein and Y Koren, “Error eliminating rapid ultrasonic firing for mobile robot obstacle avoidance”, IEEE Trans. Robotics and Automation, Vol 11, No. 1 pp 132-138, 1995.

[2] K. Jorg and M. Berg, “Mobile robot sonar sensing with pseudo-random codes”, Proceedings 1998 IEEE Conference on Robotics & Automation, Leuven, BelgiumMay 1998, pp 2807-2812.

[3] T. Tsubouchi, “Nowadays trends in map generation for mobile robotics”, Proceedings IROS’96, Osaka Japan 1996, pp 828-833.

[4] J. Budenske and M. Gini, “Why is it so difficult for a robot to pass through a doorway using ultrasonic sensors?,” IEEE Conference on Robotics and Automation, San Diego CA, 1994 pp. 3124-3129.

[5] L. Kleeman and R. Kuc, "Mobile robot sonar for target localization and classification", International Journal of Robotics Research, Volume 14, Number 4, August 1995, pp 295-318.

[6] H. Peremans, K. Audenaert, and J. M. V. Camperhout, “A high-resolution sensor based on tri-aural perception,” IEEE Transactions on Robotics and Automation, vol. 9, pp. 36-48, 1993.

[7] P M Woodward, Probability and Information Theory with Applications to Radar. 2nd ed. 1964, Oxford: Pergamon Press.

[8] K S Chong and L. Kleeman, ”Feature-based mapping in real, large scale environments using an ultrasonic array”, International Journal Robotics Research, Vol 18, No. 1, Jan 1999, pp. 3-19.

[9] L. Kleeman, “Fast and accurate sonar trackers using double pulse coding”, IEEE/RSJ International Conference on Intelligent Robots and Systems, Kyongju, Korea, October 1999, pp.1185-1190.

[10] H. Akbarally and L. Kleeman, "3D robot sensing from sonar and vision", IEEE International Conference on Robotics and Automation 1996, Minneapolis, Minnesota, April 1996 pp. 686-691.

[11] Bozma and R Kuc. "Characterizing the environment using echo energy, duration, and range: the ENDURA method", in IEEE/RSJ International Conference on Intelligent Robots and Systems. 1992. Raleigh, NC: pp. 813-820.

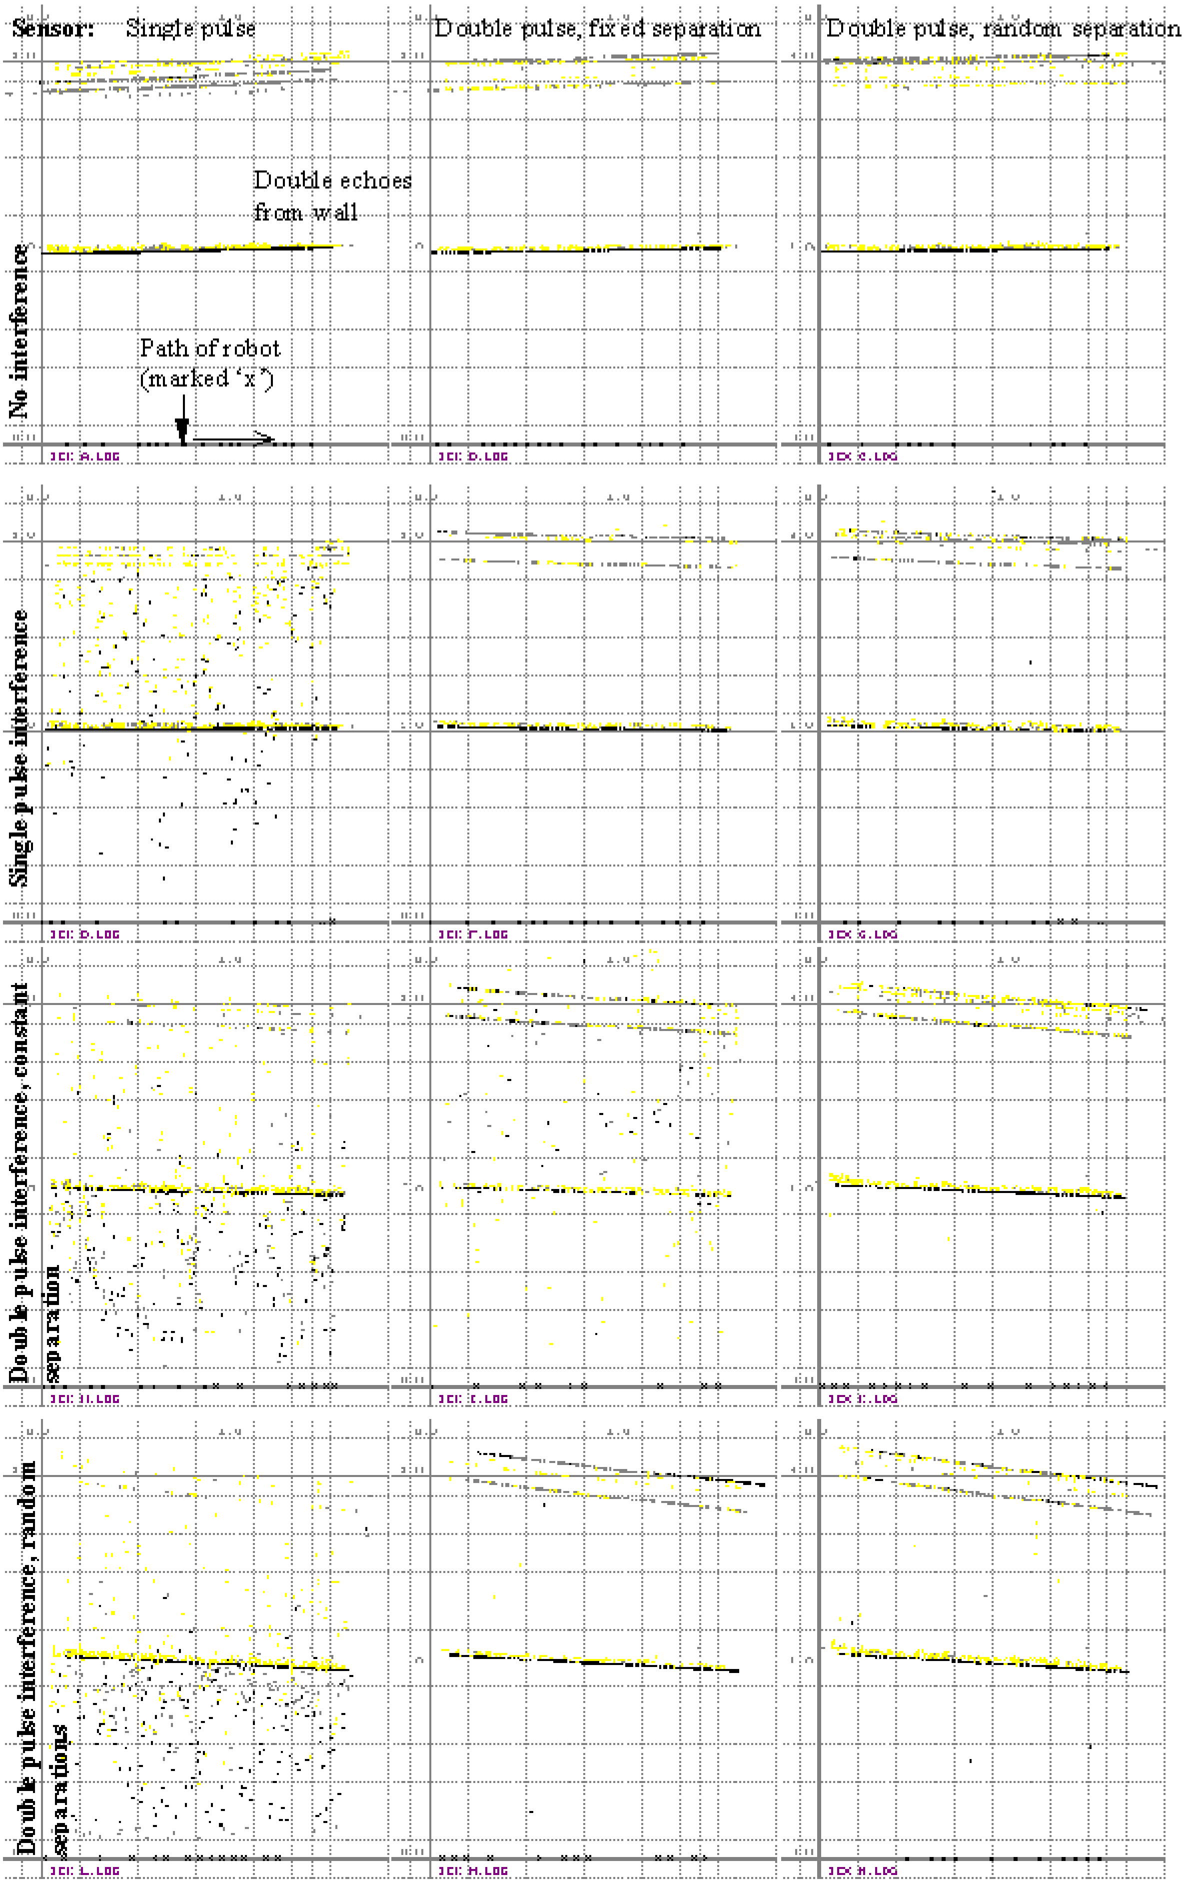

Figure 4 – Performance of sensing methods in the presence of interference. (Scales in metres)