The 2009 IEEE/RSJ International Conference on
Intelligent Robots and Systems
October 11-15, 2009 St. Louis, USA

An Advanced Sonar Ring Design with 48 Channels of Continuous Echo Processing using Matched Filters

Damien Browne and Lindsay Kleeman
Intelligent Robotics Research Centre
Department of Electrical and
Computer Systems Engineering
Monash University 3800, Australia

Abstract— **Advanced sonar systems produce both accurate range *and* bearing measurements rather than just range alone as in conventional systems. Previous advanced sonar rings do not process incoming echo data as it arrives, but only after a completed measurement cycle and only on limited data samples that are above a noise floor threshold. The system described in this paper can process all echo data, as it is produced, with full matched filtering tuned to each receiver on all echo samples directly implemented in hardware. This minimises measurement latency, important for real time robotics applications and also provides optimal arrival time estimates due to the matched filters. Previous systems have used sequential firing of many transmitters around the ring to prevent interference between transmitters or used different pulses fired simultaneously. This paper presents a single transmitter solution where a reflector disperses a single ultrasonic pulse evenly around the sonar ring. Processing is performed with a dedicated hardware data processing architecture implemented with a Field Programmable Gate Array to achieve the desired real time performance. The system reports range and bearing results at a rate of 30 measurement cycles a second in a full 360 degree coverage to ranges up to 4 metres with the prototype processing over 4.9 Giga-arithmetic operations per second. Experimental results are presented that show the performance of the system is suitable for high speed mapping and localisation applications.**

I. Introduction

Sonar sensor systems for robots are a low cost and complementary sensor modality to laser time of flight systems. Sonar can provide very accurate range and bearing to acoustic reflectors that are not detectable by laser based systems [6]. An obvious example is a glass window, but more subtle features like doorjambs are good acoustic reflectors but below the range resolution of laser systems. Examples exist where advanced sonar data is directly beneficial to a robot employing laser scan matching along a corridor where the sonar can produce sufficient measurement of features along the corridor to maintain good localisation in a SLAM implementation. [2].

Sonar signal processing aims to obtain arrival times of echoes from discrete time digital samples. In its crudest form, a simple threshold can be used often combined with a varying gain as in the popular Polaroid ranging module employed in many early sonar ring designs [5], [10]–[12]. Simple threshold based approaches are susceptible to noise and cycle hopping where the threshold is surpassed on a second or subsequent cycle of the received echo pulse. The matched filter (often called template matching) has been shown to offer the highest accuracy in sonar arrival time estimation [6] at the expense of the high computational load of correlating the entire echo signal with expected possible pulse shapes.

Sonar rings based on matched filtering have suffered due to processing demands of frequently updated data on many different receiver channels. Given the amount of data returned from a high frequency sonar ring, processing the data at a real time rate ($\approx$30 Hz for a 5 metre range) is not a simple task. For example in the current system, data at a rate of 36 Mbytes per second is processed (48 channels at 0.5 Msample/sec for 12 bit samples) and 2.47×10^9 multiplies per second and the same number of additions. In previous attempts at sonar ring design with full matched filtering, DSP boards have been used to process the data [1], [4]. While a successful sonar ring has been demonstrated [3] with a repetition rate of around 15 Hz, the processing is limited to small blocks of data that exceed a threshold based on the noise floor and matched filtering is performed as a batch on these blocks after all the echo data has been logged at the end of a measurement cycle. Interference effects between neighbouring transmitters around the ring is a difficulty in multiple transmitter systems [4], [12]. This paper uses a single transmitter.

We propose and demonstrate the use of Field Programmable Gate Array(FPGA) technology to provide the processing required to identify Time of Flight of echoes from 48 sonar transducers using all data samples, even those below the noise floor threshold used in [4]. This potentially allows the identification of very weak pulses. The system is also designed to cope with the dispersive properties of air and/or the angle of reflection by dynamically changing the template used in matched filtering such that it more closely matches the expected signal as the pulse travels further.

The processing system was designed on a commercially available FPGA development board using a XILINX Virtex 2 pro FPGA. This board has around 30,000 logic blocks each containing a 4 input look up table and a flip flop, 2.4 Mbit of RAM and 136 18 bit$\times$18 bit multipliers. The electronic interface between the FPGA and the transducers is via a set of custom designed Printed Circuit Boards

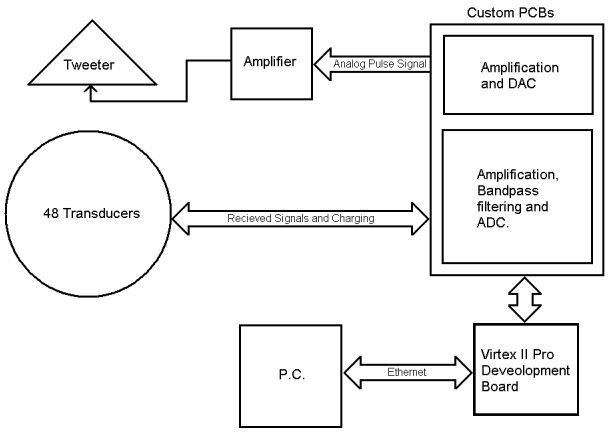

Fig. 1. A block diagram of the sonar ring system

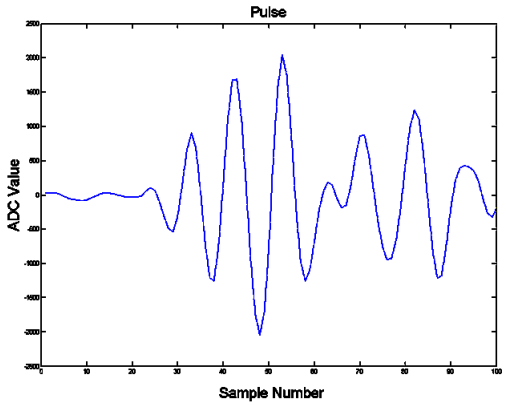

Fig. 2. A template used to match to incoming signal

(PCB). The sonar ring consists of a single transmitter, a Murata ESTD01 "super tweeter" [9] above a CNC machined parabolic reflector [8] allowing relatively uniformly shaped pulses in all directions from an effective point source. This simplifies template selection and resolves crosstalk as all pulses originate from the same point as discussed later in the paper. The ring of receivers is made up of 48 polaroid 7000 series transducers. Transducers are placed around the ring in pairs of two facing the same angle. Each pair is rotated around the ring at a progression of $15°$. To reduce the diameter of the ring transducers are placed at two different heights as shown in Figure 6. A block diagram of the entire system can be seen in Figure 1.

The paper is organised as follows. Section II briefly reviews the matched filter technique that optimally estimates the arrival time of a pulse when Gaussian noise is present. The sonar ring hardware, including the transmitter design, the receiver structure and gain stages are presented in section III. Section IV covers the hardware design of the data path that processes the sonar echo data, with section V describing the post processing of data to produce sub-sample arrival time estimation and bearing angle estimate. Results of experiments are presented in section VI and concluding remarks in section VII.

II. Template Matching

The processing method used to estimate the arrival time of echoes is called template matching or equivalently matched filtering. The receiver data samples are compared with a template based on an expected pulse shape. The comparison is implemented with a cross correlation. The template we used is 100 samples in length and an example is shown in Figure 2.

A. Cross Correlation

A cross correlation is performed on the incoming signal with the template as shown in equation (1).

$$\sum_{n=1}^{100} S_n \times T_n \tag{1}$$

The cross correlation in Equation (1) can be normalised so that the results do not depend on the power of S and T, but only the agreement in shape between the two. Equation (2) shows the correlation coefficient that takes on a maximum of 1 when the incoming signal S is a scalar times that of the template T. We consider any signal with correlation coefficient greater than a threshold, say .85, to be a valid returned pulse. The arrival time is estimated from the relative position of the template to the received pulse that gives rise to the best correlation.

$$-1 \leq \frac{Crosscorrelationvalue}{\sqrt{Power_{signal} \times Power_{template}}} \leq 1 \tag{2}$$

B. Template Selection

A problem with comparing a template with a received signal is that secondary matches occur. That is, parts of the signal that returns before and after the actual pulse arrival time can match the template almost as well as the desired peak position of (2). The template start and finishing positions is selected to maximise the difference between the peak correlation and its neighbouring side peaks, since this only needs to be done once, a simple search method is used based on collected echoes from the sonar ring itself.

A signal is recorded from the sonar ring from a manually configured target so that it contains very little noise. A template is extracted from this data such that the correlation at the arrival time of the pulse was most unique. That is, template is selected such that the second highest correlation squared value across the recorded data is minimized. By doing this we ensure that there is a greater boundary between the threshold of the correlation of the pulse and its self correlation. Figure 3 shows the variation in the square of the correlation as the pulse arrives. Provided the training target behaves acoustically similar to targets seen in a testing environment the pulse shapes will correlate well.

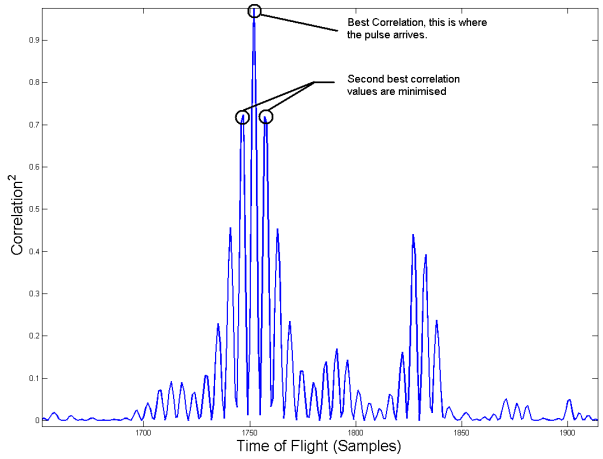

Fig. 3. Correlation as a pulse arrives

III. Sonar Ring

The sonar ring presented in this paper has a number of innovations that improve on a previous DSP based design [3]. These improvement are summarised as: a reduction in physical size and power consumption since a single FPGA chip replaces 7 DSP devices in [3]; a higher receiver sample rate of 500 kHz; potentially more accurate feature localization due to the higher sample rate and uniformity of transmitted pulse; and the ability to implement many different sonar ranging techniques not previously possible on a large scale ring by rapidly reconfiguring the FPGA data path.

A. Point Source Transmitter

Initial work on the design of a virtual point source transmitter and reflector arrangement is reported in [8], written in part by one author of this paper. This paper employs the same transmitter design and reports new work on the receiver array and the full implementation of the processing data path. Figure 6 shows a photo of the ring prototype. A CNC machined parabolic conical reflector sits just below the Murata ESTD01 "Super Tweeter Drive" [9] that is at the focal point of the parabolic shape.

The operation of the Murata ESDT01 is such that it closely approximates a point source. This was verified experimentally in [8]. A diagram of the "Super Tweeter" is shown in Figure 4. The whole transmitter structure is designed to produce cylindrical wavefronts that are reasonably uniform in all bearing angles around the ring.

By having a single transmitter we remove the problem of dealing with crosstalk between transmitters, see Figure 5. Though crosstalk can be detected using pulse coding techniques, the resulting overlapping echoes must be discarded and hence their time of flight information is lost. The current approach has a single pulse transmitted that is used by all receivers and hence their is no issue with separate transmitted pulses overlapping.

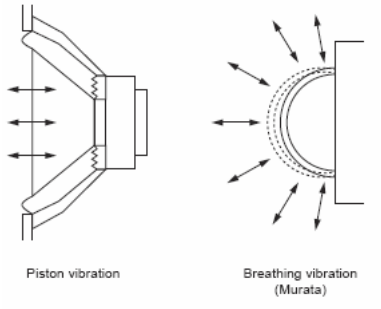

Fig. 4. Cross sectional comparison of vibration of a piston versus the ESTD breathing vibration taken from [9]

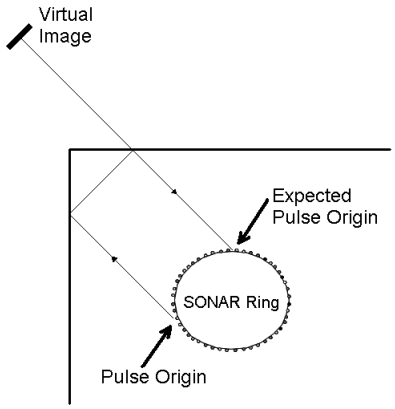

Fig. 5. An Example of Crosstalk on a traditional sonar ring

Interestingly this transmitter configuration also lends itself to other ranging techniques besides the pulse echo approach used here, such as Continuous Transmission Frequency Modulation (CTFM) where the problems of crosstalk that arise with separate discrete transmitters in a traditional ring cannot be resolved easily with pulse coding techniques [7].

B. Transducer Pairs

The ring of receivers shown in Figure 6 is made up of 24 pairs of Polaroid 7000 series transducers. Each transducer in a pair is closely positioned next to its partner to improve the data association between receivers for robust bearing estimation as established in [6]. The time of flight noise along two transmission paths in air with closely spaced receivers is highly correlated as can be seen in Figure 10. The angle estimate which relies on the difference of these times of flight is thus not affected by the common air column noise.

Each transducer is biased to 300 V. This charge is periodically refreshed by switching on relays via a signal from the controlling FPGA. The objective is to isolate electrical noise generated by the 300 V source from the receiver high gain electronics. Currently the charge is conservatively refreshed every 5 seconds, with charging taking about 1 second, during which no pulse can be recorded. As the bias discharges, the template shape modifies slightly and the sensitivity of the

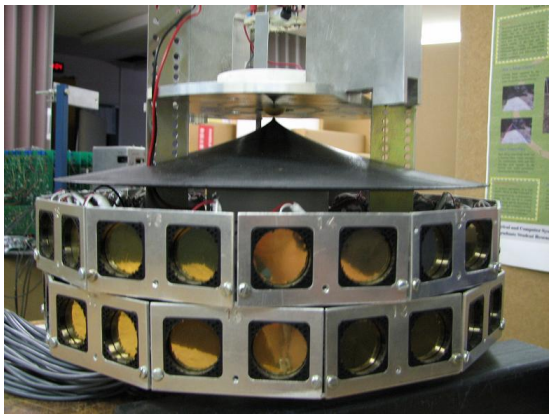

Fig. 6. The assembled sonar ring

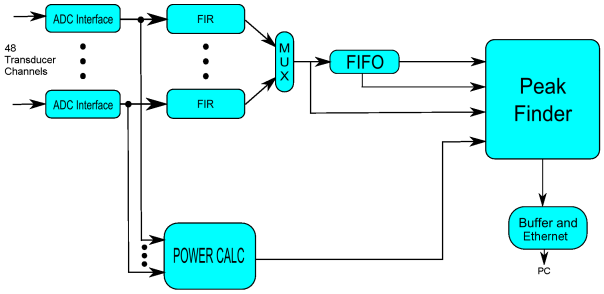

Fig. 7. The FPGA Processing Data Path

receivers decreases.

Each transducer in a pair points in parallel directions. In general the pair will detect very similar reflections but with a difference in arrival time that depends on the relative angle to the reflector.

C. Sample Rate and Resolution

The signal from each transducer is bandpass filtered, amplified and sampled by a 12 bit Analogue to Digital Converter (ADC). The sample rate is currently set at 500 kHz with the ability of the hardware to increase to 1 Mhz. Previous sonar rings have shown that a high sample rate, relative to the frequency of the transmitted pulse, helps reduce problems such as cycle hopping [3], [11].

D. Adjustable Amplification

The signals received by the transducers are amplified before being sampled. The amplifier gains are controlled from a digital word sent from the FPGA. Using this we can progressively increase the gain to compensate for energy losses and geometric spreading as the pulse travels in air.

IV. FPGA Processing

By designing a custom hardware pipeline for the sonar ring, a large data rate can be processed. The whole pipeline can process up to the equivalent of 7.0 Giga Instructions Per Second (GIPS) with less than 75% of clock cycles used for processing. The processing functions consist of template to receiver signal cross correlation, received signal power computaton, and time of flight extraction and validation. The information consisting of the extracted ToF, power and correlation coefficient is packetised and buffered for 100 Mbit ethernet communication with a host computer. A diagram of the sonar processing core on the FPGA data path can be seen in Figure 7. Each part of this data path is explained in greater detail in the following sub sections. Verilog Hardware Description Language (HDL) was used to specify the design of the datapath.

The most significant advantage this HDL data path has over previous sonar ring processing approaches is that the correlation between the signal and the template is calculated for every value of the signal. Due to DSP processing constraints, previous sonar ring correlation calculations were only performed once the power of the signal reached a certain threshold [4]. What this now means is that a pulses can be identified when the signal to noise ratio is less than 1. This provides the capability of lifting pulses out of the noise floor and increasing the sensitivity of the sonar system. Despite the increase computation load, all calculations are performed on a single FPGA chip, significantly reducing the size of hardware required to interface with the sonar ring.

A. ADC Interface

All ADCs use a common 10 MHz clock and common chip select signals. Data is transmitted serially at 10 Mbits per second from the ADC to the FPGA. The data is parallelized and synchronized with the faster 50 MHz clock speed of the rest of the data path.

B. Finite Impulse Response (FIR) Filtering

FIR filters are used to correlate the signal with the expected pulse shape. These filters are generated using XILINX FIR compiler 3.2. Though the data sheet is no-longer publicly available most features are very similar to those in [13]. These FIRs exploit the difference between the central clock speed and the speed of the incoming samples to reduce the number of fixed point multipliers and logic required.

The FIRs selected for this system are designed to have multiple sets of filter templates which correspond to the filter coefficients of the FIR. These multiple filters can be used to compensate for the changes in the expected reflected pulse shape as the speed of sound varies with frequency. That is, as the pulse travels further in air the pulse changes shape. By dynamically changing the filter template as the time of flight increase we can have a template that more closely represents the returned pulse shape. Though this is implemented in the logic design, only one template has been used to generate prototype results seen in this paper. Despite this the good correlations at long distances are still acceptable due to the ability of the system to process the entire incoming signal.

C. Power Calculation

The equation for calculating power of a digital signal, within a certain window, is a similar but different to an

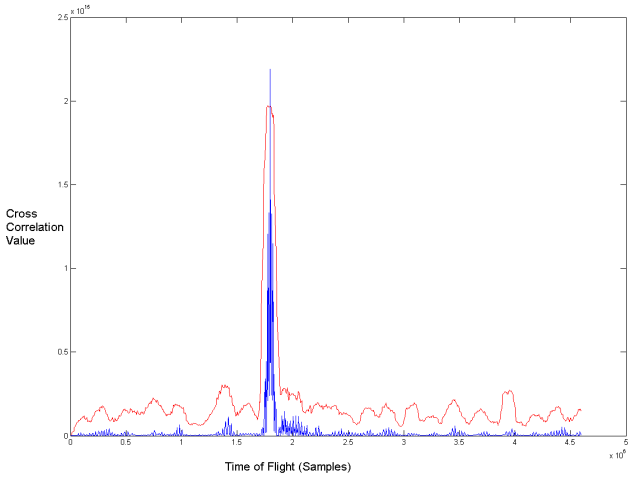

Fig. 8. Results showing the movement of the dynamic threshold relative to the square of cross-correlation value.

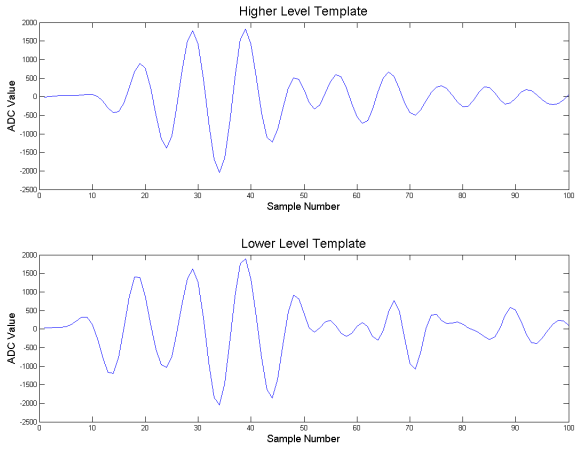

Fig. 9. Examples of two templates gathered by receivers at similar flight times but differing heights.

FIR filter. A custom designed logic block has been designed for this purpose and we are able to exploit the difference between clock speed and data rate and also the continuity of the calculation to significantly reduce the required logic. A recursive formulation is shown below for the calculation of the windowed power in Equation (3).

$$\sum_{n=i}^{i+k} a_n^2 \equiv \left[\sum_{n=i-1}^{i+k-1} a_n^2\right] + a_{i+k}^2 - a_{i-1}^2 \tag{3}$$

This shows that the entire power calculation sum does not have to be recalculated as each new sample comes in, but rather the previous power is adjusted with the square of the incoming signal minus the square of the last sample that is no longer included in the calculation. This results in the reduced logic path for the power calculation unit shown in Figure 7.

D. Pulse Identification

Equation (4) can be changed to create a condition that is easily implemented in logic as shown in Equation (5). The equation gives a dynamic threshold where the correlation value produced by the FIRs is relative to the power of the incoming signal.

$$Threshold < \frac{Crosscorrelationvalue}{\sqrt{P_{signal} \times P_{template}}} \tag{4}$$

$$P_{signal} \times P_{template} \times Threshold^2 < Crosscorrelationvalue^2 \tag{5}$$

A graph demonstrating the dynamic threshold can be seen in Figure 8. The red line represents the threshold, while the blue line represents the correlation. As the power of the signal increases, the correlation value is expected to increase. When the correlation value peaks above this line a pulse is considered to have arrived since its correlation coefficient is above threshold.

E. Pulse and Template Selection

Through experimentation a short two period pulse was transmitted via the speaker. The pulse used in this paper has 2 periods of a 48 kHz sine wave.

Experimentation confirms that the two different heights of transducers receive significantly different pulses due to the angle of reflection. This is most apparent with flight distances of less than a metre. Therefore two separate templates have been created, one for each height of receivers. Examples of the two different templates can be seen in figure 9

F. Network Interface

A simple custom protocol is used to transmit the pulse data from the FPGA using 100 Mbits per second ethernet to a host computer. Each packet of data contains the channel number the pulse arrived on, the correlation data for the pulse as well has one sample before and one sample after, the power at the arrival time and the time of flight. Each packet is 60 bytes long with room available to transmit additional data if it required later.

V. Post Processing

Once a pulse has been identified from the signal data, the information is sent via an ethernet to a computer. At this point there is a far lower amount of data bandwidth required so the information can easily be processed in real time to create more usable data. This is done in three main ways: Firstly the sub sample pulse arrival time is calculated; Secondly the bearing of signals to each pair is computed; and thirdly additional validation constraints are applied to remove redundant or unreliable pulses.

A. Sub-sample Arrival Time

Given that the correlation of signal to the template will be at a peak of the continuous correlation, we can use this to our advantage to estimate a sub-sample arrival time of pulses. By taking the three closest correlation points to the pulse arrival time we fit a parabola to these points and calculate

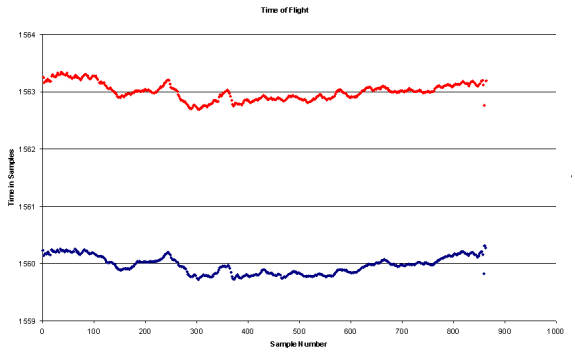

Fig. 10. Time of Flight Results for two paired receivers

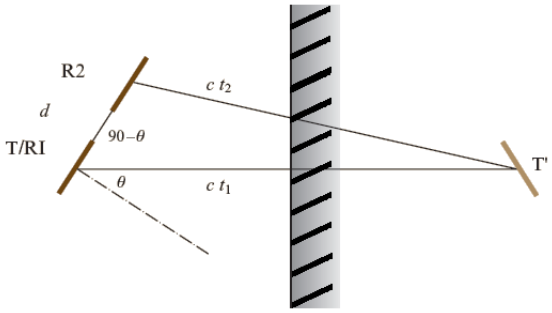

Fig. 11. A diagram showing the bearing estimation geometry. T is the transmitter, RI and R2 are receivers and T' the virtual transmitter location.

the peak and hence estimate the sub-sample arrival time. Figure 10 shows pulse arrival times for two channels of the same pair. The sub-sample arrival times vary as is expected due to thermal and turbulent effects on the speed of sound. Both channels vary mostly in the same way, consistent with the common air column that they experience.

B. Bearing Calculation

Since all receivers are organized in pairs we can calculate bearing estimations using multiple transducers as seen in [6]. Figure 11 shows the basic geometry of the estimates. The post processing software checks the arrival time of pulses on transducer pairs. If the time between these pulses is less than can be produced by the maximum possible arrival angle then the pulses are assumed to have come from the same reflector. The bearing estimation is calculated for these two pulses. However, due to occasional variation in side peaks in the template correlation combined with noise, the sonar system may have multiple arrival times for the same reflection. To resolve this, only the best correlating pulse within an arrival time equal to the time of the template is taken, all other large correlations within this time frame are assumed to be invalid. Future improvements mentioned above to the templates will reduce the occurrence of these multiple arrival times per pulse.

Fig. 12. Photo of part of the corridor used for testing.

VI. Results

Results from a proof of concept prototype system are presented in this section. The sonar ring was positioned at known locations along a corridor that was manually measured with a tape measure to obtain a ground truth. Careful note was taken of the position of moldings and doorjambs present in the corridor that are known to be good sonar reflectors.

A. Processed Sonar Data

Figure 13 shows as coloured circles the sonar ring reported location of reflections where pulses have been detected on both transducers of a pair and a bearing calculation can be performed. The ground truth is represented with blue lines and the robot locations as large circles. Figure 12 shows a photo of part of the corridor environment.

The vast majority of the results align with the ground truth map - this may not be immediately apparent since measurements at the same correct locations overlap on the diagram. There is however the occasional erroneous bearing measurement, where the reported positions do not correspond to a corridor feature. These errors can be caused by a mismatch between the template used in matched filtering and the arrival pulse shape. A mismatched template causes less discrimination between the desired correlation peak and the secondary peaks either side. Therefore an arrival time error of around one cycle can result.

As mentioned above the prototype employed common templates for each of the two layers of receivers in the ring. In practice there is variation in the pulse shape between different receivers around the ring on the same layer due to subtle changes in the reflector shape, imperfections in the tweeter and variations in receiver transducer properties.

VII. Conclusion

The paper has presented a new FPGA based sonar ring incorporating central a single point transmitter. The system can process all echo data, as it is produced, with full matched filtering tuned to each receiver on all echo samples directly

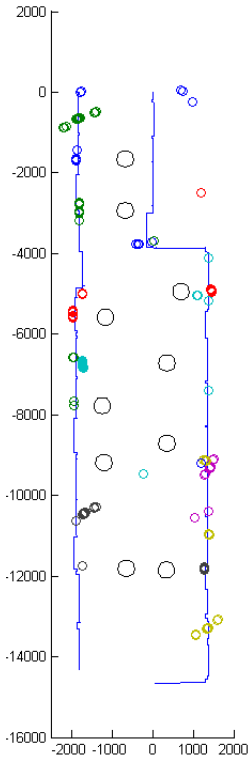

Fig. 13. Sonar Data Vs Map (Large circles are robot locations. Small are sonar reflections with bearing)

implemented in hardware. This minimises measurement latency, important for real time robotics applications and also provides optimal arrival time estimates due to the matched filters. Previous systems have used sequential firing of many transmitters around the ring to prevent interference between transmitters or used different pulses fired simultaneously. This paper has presented a single transmitter solution where a reflector disperses a single ultrasonic pulse evenly around the sonar ring. Processing is performed with a dedicated hardware data processing architecture implemented with a Field Programmable Gate Array to achieve the desired real time performance. The system reports range and bearing results at a rate of 30 measurement cycles of multiple range/bearing targets a second in a full 360 degree coverage to ranges up to 4 metres with the prototype. Results have been presented that show the potential application of the system to mapping and localisation problems.

VIII. Future Work

Currently the system is a first prototype and improvements are expected before the final version of this paper is due. Raw data transmitted from the FPGA based processing system is as would be expected and can be improved through better selection of templates, better calibration of the reflector and higher gain amplification. We also hope to use different ranging techniques using a similar hardware architecture as this sonar system, such as CTFM [7] and Multiple Pulse Coding. It is hoped that these two methods will further increase the refresh frequency of ranging data beyond 30 Hz.

IX. Acknowledgements

We gratefully acknowledge the technical input of Steve Armstrong.

References

[1] D. Bank and T. Kampke, *High-resolution ultrasonic environment imaging*, Robotics, IEEE Transactions on **23** (2007), no. 2, 370–381.

[2] A. Diosi and L. Kleeman, *Advanced sonar and laser range finder fusion for simultaneous localization and mapping*, Proceedings 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems IROS2004 (2004), 1854–1859.

[3] S. Fazli and L. Kleeman, *A low sample rate real time advanced sonar ring*, Australasian Conference on Robotics and Automation 2004 (2004).

[4] _____, *A real time advanced sonar ring with simultaneous firing*, Intelligent Robots and Systems, 2004. (IROS 2004). Proceedings. 2004 IEEE/RSJ International Conference on **2** (2004), 1872–1877 vol.2.

[5] Sewan Kim and Younggie Kim, *Robot localization using ultrasonic sensors*, Intelligent Robots and Systems, 2004. (IROS 2004). Proceedings. 2004 IEEE/RSJ International Conference on **4** (2004), 3762–3766 vol.4.

[6] L. Kleeman and R. Kuc, *Mobile robot sonar for target localization and classification*, The International Journal of Robotics Research **14** (1995), 295–318.

[7] _____, *Springer handbook of robotics*, ch. Chapter 21: Sonar Sensing, pp. 491–519, Springer-Verlag Berlin Heidelberg, 2008.

[8] L. Kleeman and A. Ohya, *The design of a transmitter with a parabolic conical reflector for a sonar ring*, Australasian Conference on Robotics and Automation 2006 (2006).

[9] MURATA, *Estd01 data sheet*.

[10] S. Shoval and J. Borenstein, *Using coded signals to benefit from ultrasonic sensor crosstalk in mobile robot obstacle avoidance*, Robotics and Automation, 2001. Proceedings 2001 ICRA. IEEE International Conference on **3** (2001), 2879–2884 vol.3.

[11] A.Ohya T.Yata and S.Yuta, *A fast and accurate sonar-ring sensor for a mobile robot*, Proceedings of the 1999 IEEE International Conference on Robotics and Automation (1999), 630–636.

[12] S. Walter, *The sonar ring: Obstacle detection for a mobile robot*, Robotics and Automation. Proceedings. 1987 IEEE International Conference on **4** (1987), 1574–1579.

[13] Xilinx, Inc., *Fir compiler v4.0 product specification*, 4.0 ed., June 2008.