eRTIS: a fully embedded real time 3D imaging sonar sensor for robotic applications

eRTIS: A Fully Embedded Real Time 3D Imaging Sonar Sensor for Robotic Applications

Robin Kerstens[1], Dennis Laurijssen[2], Jan Steckel[3]

***Abstract*— Many popular advanced sonar systems provide accurate and reliable measurements containing crucial info needed by robotic applications such as range, bearing and reflection strength of the objects in the field of view. While these sensor systems provide these crucial pieces of information accurately, they are often limited by a lack of processing power and/or size which leads to them needing an external computing device to process all the information generated by the microphone array on the sensor. In this paper we present two versions of a novel fully embedded 3D sonar sensor which have different sensing architectures which enable 3D perception for robotic application in harsh conditions using ultrasound at low cost. Experimental results taken from an office environment will show the 3D localization capabilities and performance of the sensor, showing the sensor has a large field-of-view (FoV) with accurate 3D localization combined with real-time capabilities.**

I. INTRODUCTION

The last decade has seen a large increase in popularity of robotic systems. Both stationary and moving robots are being developed and manufactured for numerous applications. From simple machines to advanced robots, a large part of these systems rely on the feedback from an operator or the environment. With today's demand for increasing levels of autonomous behavior also comes an increased need for accurate and reliable sensing mechanisms, either for added safety or for completely autonomous control. This is especially true for challenging applications such as autonomous navigation, SLAM or object recognition. Among the most popular sensing modalities for these advanced applications are Time-of-Flight (ToF) sensors [1] [2], Radar sensors [3] and RGB-D or traditional camera's (e.g. MonoSLAM [4]) to provide exteroceptive sensor data. While these systems are known to perform well in their specific fields, they struggle in harsh environmental conditions (fog, dust, etc.), which impacts important key performance parameters such as measurement resolution or even hinder the capability of object detection completely. Furthermore, many optical systems exhibit a limited field of view (ToF cameras) or lack 3D sensing performance (many commercial LIDAR sensors). Despite of these issues these optical sensing systems still remain popular in the field of robotics. When looking at nature for solutions, various animals use echolocation to orient themselves in complex environments or for tasks like hunting prey [5]. A prime example of an animal genus that expertly makes use of echolocation for a wide variety of tasks are Microchiroptera [5]. These creatures have, with the help of evolution, developed various species using different techniques for using echolocation to handle these complex tasks [6][5]. In previous work, our group has developed several 3D sonar sensors which mimic the echolocation system of bats [7][8] which we further developed into an industrially-relevant 3D imaging sonar [9][10]. As stated before, while these systems achieve accurate 3D localization, they rely on an external computer running MATLAB software for the signal processing tasks, limiting the real-world applicability of the sensor solution. In this paper we present two architectures that build upon our previous developments in which we transfer the computational signal processing tasks from a dedicated computing device to an embedded system. Without the need for an external computing platform these are very compact yet powerfull systems capable of imaging the complete frontal hemisphere and performing 3D localization, suitable for applications like SLAM [10] or complex control tasks such as navigation in corridor-like environments [11]. The rest of this paper is structured as follows: In the next section (II) the sonar architecture will be discussed to give an overview of how the sensor systems acquire, process and evaluate environmental data. Section III will cover the embedded implementation and all the components of the embedded platform. In Section IV the signal processing is discussed. After this, the experimental results are shown and discussed in section V, where we will detail about the point spread function obtained and the 3D mapping results will be discussed. To conclude the paper, we discuss the results and briefly mention possible future work on this topic.

II. SENSOR ARCHITECTURE

The sensor architectures described in this paper serve as active acoustic 3D cameras. They operate by emitting a broadband pulse using a single transducer, which is reflected by the environment. These reflections are picked up by an array of microphones which are all sampled synchronously, which allows recording the acoustic wave field impinging on the array. Through array signal processing techniques, these raw microphone signals can be transformed into a 3D acoustic image of the environment. A block schematic of the two architectures' components can be seen in Fig. 1. In this paper

[1]Cosys-Lab, Faculty of Applied Engineering, University of Antwerp and the Flanders Make Strategic Research Centre robin.kerstens@uantwerpen.be

[2]Cosys-Lab, Faculty of Applied Engineering, University of Antwerp and the Flanders Make Strategic Research Centre dennis.laurijssen@uantwerpen.be

[3]Cosys-Lab, Faculty of Applied Engineering, University of Antwerp and the Flanders Make Strategic Research Centre jan.steckel@uantwerpen.be

we focus on two methods for implementing the digitizer and processing architecture using embedded systems engineering techniques. To capture the signals that are reflected by

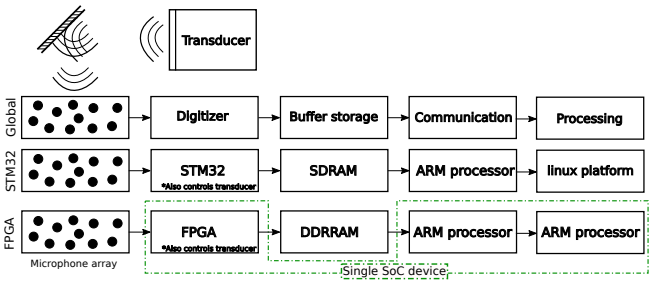

Fig. 1. Schematic global overview of both architectures.

objects in front of the sensor, a 32-element microphone array is used. This array consists out of 32 microphones with a built-in $\Sigma\Delta$-ADC that uses Pulse-Density-Modulation (PDM) [12]. The benefit of this is that no amplification of the microphone signal is necessary and through which analog noise which is added by the circuitry between the microphone, the amplifier and the digitizer is minimised. The microphones need to be connected to a digitizer that is capable of handling the data coming in at a high sampling rate of 4.5 MHz (because of the PDM). At this point there is the first difference between the two architectures, where one architecture is based on an STM32 microcontroller featuring an ARM Cortex-M4 processor and the other is based on the Cyclone V SoC which has an FPGA and an ARM Cortex-A9 dual core processor in the same package. In the case of the STM32 architecture the digitization will be handled by the Cortex-M4 processor itself, while in the FPGA-SoC architecture, the digitizing is handled by the FPGA fabric. This allows for the Cortex-A9 to handle all communication and signal processing required for the imaging algorithms. The digitized sensor data is stored on an on-board SDRAM memory where it can be accessed later for processing. In the case of the STM32 platform the signal processor is an of-the-shelf computing platform running Linux (such as the Odroid XU-4) which will do the final processing where the acquired microphone data will be transformed into acoustic 3D images. A controller on the robot can interpret these acoustic images for application specific tasks such as obstacle avoidance, corridor following or SLAM. In the other case where the Cyclone V SoC is used, there is no need for an additional computing platform as the FPGA is already accompanied by a dual-core ARM A9 processor running at 800 MHz [13]. For this FPGA-SoC architecture we used an of-the-shelf board from Terasic, the DE10-nano, that features this SoC. This greatly reduces the cost of the system while retaining the system performance. More details will be giving later on in the paper.

III. Signal Processing

Fig. 2 shows an overview of the signal processing steps that are taken by the system to transform the acquired microphone data into 3D acoustic images. A transducer (in this case the SensComp 7000 [14]) emits a broadband ultrasonic chirp between 20 kHz and 80 kHz. This signal will be reflected by objects in the environment. As the acoustic energy of the transducer is distributed quite broadly in the frontal hemisphere, and the microphones are near omnidirectional, the field-of-view of the sensor consists of the frontal hemisphere [9][12]. These reflections will then be received by each microphone in the 32-element array individually. The first signal processing step will be a matched filter that will optimally increase the signal to noise ratio while at the same time compressing the emitted pulse into it's auto-correlation function. This leads to the pulse compressed signal $s_i^{MF}[k]$, where i refers to the microphone channel:

$$s_i^{MF}[k] = \mathcal{F}^{-1}\left\{ S_i^M[j\omega] \cdot S_b^*[j\omega] \right\} \tag{1}$$

In equation (1) $\mathcal{F}^{-1}$ is the inverse Discrete Fourier Transform (DFT) applied on the discrete fourier transforms of signal coming from the i-th microphone channel ($S_i^M[j\omega]$) and it's complex conjugate of the fourier transform of the emitted signal ($S_b^*[j\omega]$). The next step involves the beamformer. The beamformer is used as an acoustic lens which is steered into a particular direction ($\boldsymbol{\psi}$) located in the frontal hemisphere. Here $\boldsymbol{\psi} = [\theta\ \varphi]^T$ with θ the azimuth angle and φ the elevation angle. Equation (2) shows the operation of the beamformer. Currently a delay and sum (or time-domain Bartlett) beamformer is used because of it's simplicity and robustness against calibration errors but other, more advanced, beamformers can also be implemented to achieve better peak-to-sidelobe ratios in the imaging point-spread function [15]. However, it should be noted that most data-dependent beamformers require multiple data snapshots to be acquired, which might not always be feasible in robotic applications where the sensor moves fast through the environment with respect to the sampling rate of 10-15Hz of the sonar sensor. To overcome the limitations of data-dependent beamformers, sparsity based segmentation techniques of the acoustic images might be used to increase the object localization accuracy, as demonstrated before in [16]. In the case of the time-domain Bartlett beamformer, time-delays $\tau_i(\boldsymbol{\psi})$ are added to each channel to compensate for the angle-dependent difference in time-of-arrival caused by the geometry of the array:

$$s_{\psi}^{BF}[k] = \sum_{i=1}^{32} w_i \cdot s_i^{MF}\Big[k + \tau_i(\boldsymbol{\psi})\Big] \tag{2}$$

The weigths (w_i) assigned to the channels are chosen from a gaussian window over the array aperture to minimize sidelobe levels in the point spread function (PSF). As a last step, an envelope detection is perfomed for every direction $\boldsymbol{\psi}$. When doing this a collection of range-energy profiles ($s_{(\theta,\varphi)}^{EN}[k]$) is obtained. The different range-energy profiles from the desired directions are all combined into a single image which we call the Energyscape (ES) of the environment [9].

$$E(k,\psi) = \begin{bmatrix} s^{EN}_{(\theta_1,\varphi_1)}[k] & s^{EN}_{(\theta_2,\varphi_1)}[k] & \dots & s^{EN}_{(\theta_n,\varphi_1)}[k] \\ s^{EN}_{(\theta_1,\varphi_2)}[k] & s^{EN}_{(\theta_2,\varphi_2)}[k] & \dots & s^{EN}_{(\theta_n,\varphi_2)}[k] \\ \dots & \dots & \dots & \dots \\ s^{EN}_{(\theta_1,\varphi_m)}[k] & s^{EN}_{(\theta_2,\varphi_m)}[k] & \dots & s^{EN}_{(\theta_n,\varphi_m)}[k] \end{bmatrix} \tag{3}$$

IV. EMBEDDED IMPLEMENTATION

An important feature of the proposed hardware implementations of the 3D imaging sonar sensor is that they are completely embedded, meaning that they have no need for external computational resources. This embedded character encapsulates the needed processing from a user's point of view, transforming it into a plug-and-play sensor from a robotics point of view. Here we will briefly discuss the key-components to these systems.

A. Microphone Array

Recent evolution in the microphone industry have led to small digital MEMS microphones such as the Knowles SPH0641LU4H [17]. The benefit of these microphones compared to other microphones like the Knowles FG-23329 used in [9] is that the SPH0641LU4H has a one-bit digital interface, making the need for additional amplification and ADC circuitry non-existant. Furthermore, the signal to noise ration of the digitized signals is high (approx; 60dB), which is in part due to the inclusion of the ADC into the microphone package. The highly oversampled Pulse Density Modulated (PDM) signal originating from the microphones can be converted to a Pulse Code Modulated (PCM) representation using an IIR filter in low-pass configuration in conjunction with decimation. As discussed and demonstrated in [12] these microphones can easily be integrated in ultrasound arrays used for 3D localization. The digital interface of these microphones allows us to connect them directly to the GPIO pins of a platform capable of acquiring the data at these speeds. In this case 32 microphones are placed in an array, one having a rectangular boundary similar to the array used in [8] [10] [11], while the second type of array makes use of a new method where the outline is an ellipsoidal shape and the microphones are irregularly placed using Poisson disc sampling [18]. Their corresponding Point Spread Function (PSF) is calculated using a time-domain model of the array sensor, which is shown in Fig.3.

B. STM32 platform

The ARM-based STM32 platform consists out of two custom PCB's (shown in Fig. 4), one featuring the ellipsoidal microphone array mentioned above and one contains the ARM daughterboard. The ARM processor (a STM32F429 from ST Microelectronics featuring a ARM Cortex-M4 processor) is responsible for obtaining all data from the microphones, formatting, and storing it into the 4 MB of SDRAM available. The STM32 will also initialize a DMA proces to transfer the data from the SDRAM onto the next processing block, which can be connected through USB. This allows for the flexibility of using the sensor as a feature of an existing system where there is no shortage of computing power. A dedicated embedded computing platform can also be used as an interface between the sensor and an existing platform. In the experiments for this paper an Odroid XU-4 single board computer is used. This is a small linux computer featuring a Cortex-A15 running at 2 GHz and an additional Cortex-A7 octacore CPU, 2 GB of LPDDR3 RAM along with 3 USB connections and ethernet for communication [19]. This embedded computational platform contains sufficient computational resources for the imaging algorithms in conjunction with basic robotic control processes.

C. FPGA-SoC platform

The use of FPGAs for large ultrasound arrays has already been explored in earlier literature [12][20][21]. Yet the newer System-on-Chip (SoC) architectures, combining FPGAs with processors capable of running a fully-featured OS like Linux, have still not found their way to this kind of research. This platform relies on the Altera Cyclone V SoC. This SoC has a large FPGA (Intel Cyclone V SE 5CSEBA6U23I7NDK) combined with a dual core ARM A9 processor running at 800 MHz [13] all in the same package, connected through AXI or Avalon bridges. This allows for a high performance all-in-one device. The FPGA will acquire the incoming data, do the initial processing such as the computational-heavy tasks like the IIR filtering on each channel, this gives the processor-side the possibility to run other tasks such as processing the signals and ethernet communication with other devices.

Sampling all 32 elements of the microphone array at 4,5 MHz generates 144 Mbps of microphone data. Taking in this amount of data and performing the necessary filtering requires a lot of processing power so using an FPGA to handle these tasks is beneficial for the rest of the system. Along with handling the incoming data, the FPGA also handles the initialization of the transmitting part of the system. The signal that is to be emitted sent to FPGA Block RAM at boottime and from there on the FPGA will continuously send the pulse to the corresponding circuitry on the system. The pulse that is used in this system is a 3 ms frequency sweep between 80 and 20 kHz (similar to what is used in previous works like [12][22]), inspired by the signals emitted by bats for localization [7].

D. High Voltage Amplifier

Both architectures mentioned use the popular Senscomp 7000 series transducer [14]. This transducer has a few specific operating conditions like a steady +150 V_{bias} and there is also a need for a linear frequency response in order for the frequency sweep to be emitted as accurately as possible with this transducer. To meet these requirements a custom High Voltage (HV) amplifier is designed that is tailored to this transducer and the signal that is to be emitted. A block schematic of this amplifier is shown in Fig. 5.

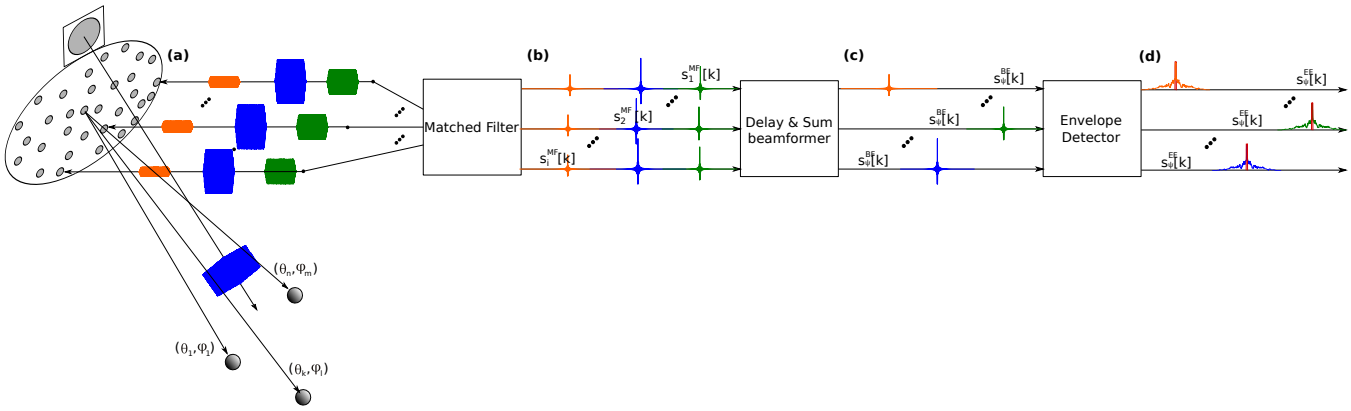

Fig. 2. The sonar system is made up out of (a) a single emitter sending a 3 ms ultrasound frequency sweep which will be reflected by the environment and then be recorded again using an array of 32 microphones placed pseudo-randomly in a ellipsoidal boundary. (b) is a matched filter used to increase the signal to noise ratio and compress the emitted pulse into it's auto-correlation function. (c) a delay-and-sum beamformer which is used to steer the array-data in certain directions and (d) an envelope extraction module to obtain the range-energy profiles.

V. EXPERIMENTAL RESULTS

A. Point Spread Function

When designing an imaging system, the key characteristic to defining the quality of a sensor is it's point spread function (PSF). This function can show the sensitivity of the sensor for point source impulses coming from different locations, the PSF's of both architectures presented in Fig. 3. Note that the two architectures use a different version of the sensor, one with a rectangular array and the other with an ellipsoidal array where the microphones are placed using poisson disk sampling. For each architecture three different PSF's are shown with different positions for a single point source. The first thing that one should notice is that, contrary to the PSF obtained from optical systems, the PSF of a planar array is anisotropic across the field of view of the sensor. The PSF is dependant on the angle of the source to the surface normal of the array, with the shape becoming wider as the angle strays further from the perpendicular location. This is caused by the visible sensor surface becoming smaller with increasing angles, making it less precise in these areas. What is interesting about this case is that the ellipsoidal array has an increased peak-to-sidelobe ratio when compared to it's rectangular counterpart, which is caused by the irregular placement through poisson disk sampling, as a poisson disk sampling distribution exhibits blue noise characteristics [23].

B. 3D mapping

In order to test the 3D mapping capabilities of our sensor, it was mounted on top of a mobile robot (Pioneer P3DX), aimed towards the forward direction and driven through the corridor an office environment. To make this environment more interesting three objects are placed in the trajectory of the robot along the length of the corridor (a small cabinet and two standing metallic cylinders). To get a ground truth of the robot position within the measurement environment at time t ($P(t)$), the robot is also equipped with a laser-range scanner and a small computer running a fastSLAM algorithm and AMCL [24]. This way it is possible to compare the 3D sonar data against the data of the laser range scanner and validate the measurements. For this experiment the robot moves at a speed of 0.3 m/s while the 3D sonar sensor mounted on-top measures the environment at a rate of 3 Hz. Note that this is not the maximal measurement-rate at which the sensor can operate. The maximum frames-per-second (FPS) of the sensor is 10 Hz, limited only by the USB2.0 interface which is upgraded in a version that is currently under development. Each sensor measurement results in an Energyscape $E(r, \psi)$, a 3D representation of the environment visualising the incident energy from all directions in the frontal hemisphere [9]. The Energyscape is created using a uniform sampling of the frontal hemisphere which is calculated by an equal area zone partitioning algorithm (EQSP [25]) and a peak detection algorithm analyses the image for object localization. Each sensor measurement at time-step t_s gives a local 3D point cloud $P_L(t)$:

$$P_L(t) = \begin{bmatrix} x_1 & x_2 & \dots & x_n \\ y_1 & y_2 & \dots & y_n \\ z_1 & z_2 & \dots & z_n \end{bmatrix} \tag{4}$$

Because the robot's path is not perfectly linear i.e. it also contains turns, the local 3D point cloud P_L needs to be mapped onto a higher-level 3D point cloud P_W in order for the measurements to be compatible to each other, which can be achieved using a standard coordinate transformation:

$$P_W(t) = T^z(\alpha_r) * P_L(t) + T\begin{pmatrix} x_r \\ y_r \end{pmatrix} \tag{5}$$

Here, α_r is the rotation component of the robot pose which is rotated around the Z-axis and $T\binom{x_r}{y_r}$ is the translation matrix, with x_r and y_r being the X and Y components of the robot pose. For each time-step we aggregate $P_W(t)$ into a global point cloud.

$$P_G = [P_W(t_1)\ P_W(t_2)\ \dots\ P_W(t_m)] \tag{6}$$

After the trajectory is completed we can combine the measurement data from both the 3D sonar sensor and the laser-

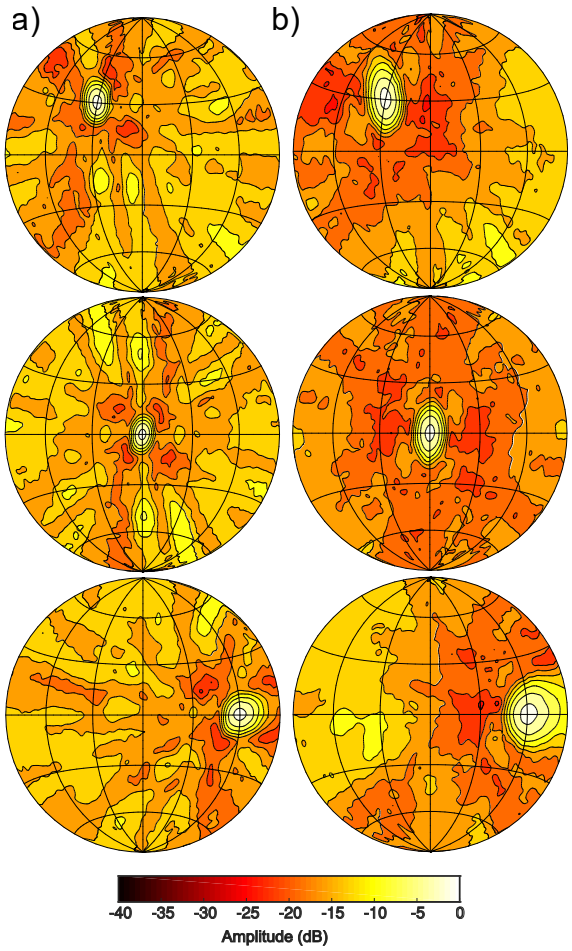

Fig. 3. The point spread function (PSF) of both microphone arrays is shown for three different source directions. The PSFs are plotted using a Lambert azimuthal equal area projection. Gridlines are spaced 30 apart and the contour lines indicate 3dB differences. First the source is located at -30°, 30° (azimuth, elevation), then at 0°, 0° (azimuth, elevation), and lastly 60°, 0° (azimuth, elevation). a) shows the PSF of the mirophone array with rectangular boundary and random microphone placement. b) shows the PSF of the microphone array with the ellipsoidal boundary and where the microphones are distributed over the area using poison-disk sampling. Notice that the Peak-to-Sidelobe (PSR) ratio is larger when using the ellipsoidal array and that the main lobe is wider, which is expected from random array literature [9].

range scanner to obtain a 3D image of the environment, this is shown in Fig. 6. This figure illustrates the 3D mapping capabilities of the system and a detailed description is provided in the figure caption.

VI. Conclusion & Future work

In this paper we have demonstrated two fully embedded real-time 3D imaging sonar (eRTIS) architectures. Both systems are capable of generating accurate 3D images of an office environment and detecting objects that are located within the field of view of the sensor, which spans the entire frontal hemisphere. Both systems do this in a fully embedded way so that no external computational power is needed, allowing for a portable system that is well suited for robotic applications. The two sensor versions that were presented have a different hardware architecture with identical performance where one

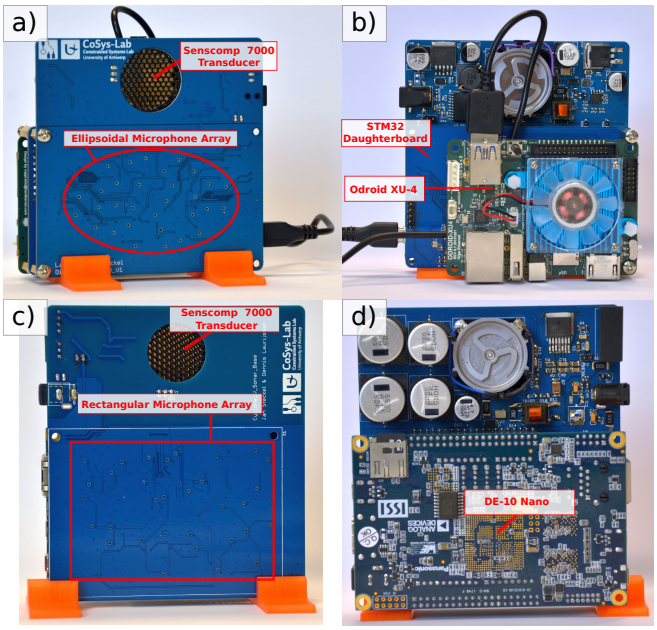

Fig. 4. a) shows the front of STM32 platform with ellipsoidal microphone array featuring the 32-element ellipsoidal microphone array and the Senscomp 7000 transducer. b) shows the back of the STM32 platform where the daughterboard and Odroid XU-4 are visible. c) Shows the front of the Cyclone V SoC platform featuring the 32-element rectangular array and the Senscomp 7000 transducer. d) shows the back of the Cyclone V SoC platform where the circuitry needed to drive the Senscomp transducer and the DE10-Nano are shown. Note that the two underlying hardware architectures behind the microphone arrays, despite their differences, offer identical performance and any differences in accuracy (as shown in Fig. 3) are caused only by the shape of the microphone array.

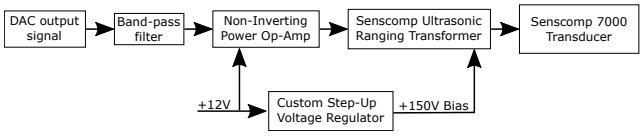

Fig. 5. The schematic of the custom designed amplifier, giving a steady $+150V_{bias}$ to operate the Senscomp transducer. Along with an linear frequency response to limit loss of information on the data that is being transmitted

is based on an STM32 processor featuring an ARM Cortex-M4 which handles the incoming microphone data and sends it to a small of-the-shelf Linux computer (the Odroid XU-4). The signal processing is handled by the Odroid which will send object locations to the robot. The second sensor architecture is based on the Cyclone V SoC featuring an FPGA and a dual-core ARM Cortex-A9 in the same package. Here the FPGA will handle all the incoming microphone data and filtering while the Cortex-A9 will take care of the remaining signal processing and communication with the robot. Also two different microphone arrays were tested in this system, where the array with ellipsoidal boundary and poisson-disk sampled microphone placement yielded a larger PSR. We demonstrated the use of the sensor in an office environment where the system succeeds in detecting the complete environment as well as the obstacles placed in front of the moving robot. Future work on this topic will include a system where the microphone array can be

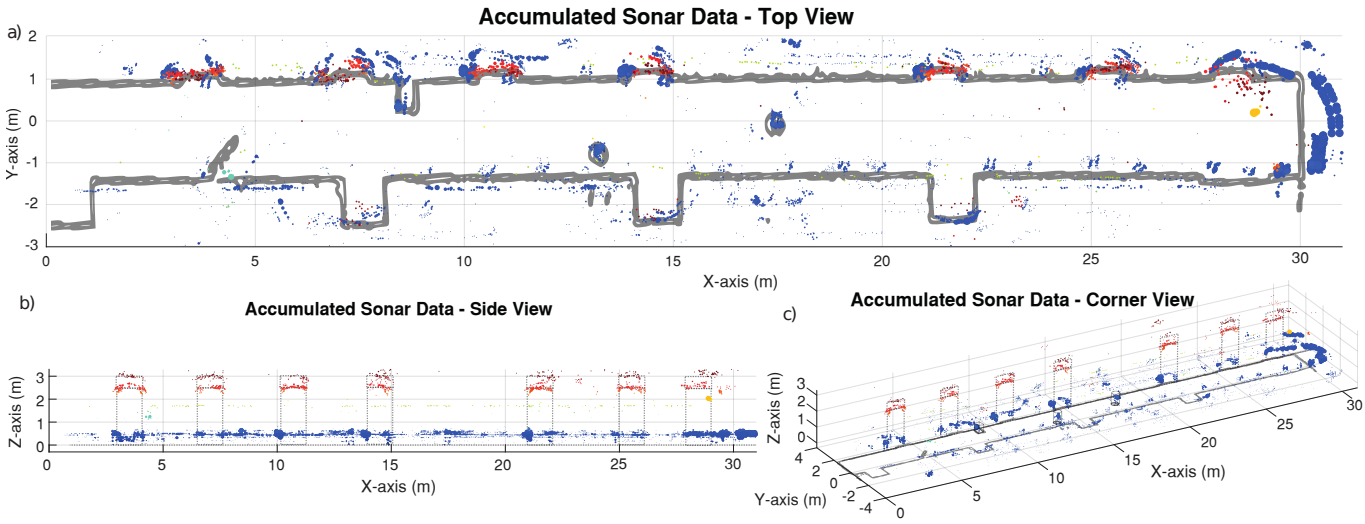

Fig. 6. The plots shown here display the data collected by the 3D sonar sensor (colored data), the data obtained by the laser-range finder (grey bottom layer), and some added visual cues to help the reader in understanding what has been measured (dotted grey lines). The color of the 3D sonar data is related to the height of the reflector found (making the 3D plots easier to interpret), while the size of the dot represents the strength with which the object's reflection was detected. Three versions of the same plot are shown. a) shows the top-view of the generated 3D map, here we can see that the 3D sonar sensor is capable of locating the boundaries of the environment along with the objects that were put in its path. b) shows the side-view where the data from the X-Z plane is visible. One can see that the walls are detected at the height of the sensor itself, and that objects such as doors also give strong reflections that are received by the sensors (the doors have windows above them that give similar reflections), the green dots found at approx 1 m height are handles for closets that are built into the walls. c) gives a corner view that shows off the 3D view. Notice the upright metallic cylinders and cabinet that are accurately detected.

changed depending on the application, making the system more versatile and reducing the cost for experimenting with multiple array geometries. This paper also illustrates the 3D mapping capabilities of the 3D sonar sensor, encouraging extended testing in more challenging environments. Another topic to be investigated are methods to increase the framerate of the imaging sensor using pulse encoding techniques, as demonstrated before in [21].

REFERENCES

[1] Gerald Rauscher, Daniel Dube, and Andreas Zell. A Comparison of 3D Sensors for Wheeled Mobile Robots. *2014 International Conference on Intelligent Autonomous Systems (IAS-13)*, pages 1–12, 2014.

[2] Stefan May, David Droeschel, Dirk Holz, Stefan Fuchs, Ezio Malis, Andreas Nüchter, and Joachim Hertzberg. Three-dimensional mapping with time-of-flight cameras. *Journal of Field Robotics*, 26(11-12):934–965, 2009.

[3] Peter Hoppen, Thomas Knieriemen, and Ewald Von Puttkamer. Laser-radar based mapping and navigation for an autonomous mobile robot. *IEEE International Conference on Robotics and Automation*, 2:948–953, 1990.

[4] Andrew J. Davison, Ian D. Reid, Nicholas D. Molton, and Olivier Stasse. MonoSLAM: real-time single camera SLAM. *Pattern Analysis and Machine Intelligence (PAMI), IEEE Transactions on*, 29(6):1052–67, 2007.

[5] Hans-Ulrich Schnitzler, Cynthia F. Moss, and Annette Denzinger. From spatial orientation to food acquisition in echolocating bats. *Trends in Ecology & Evolution*, 18(8):386–394, aug 2003.

[6] Donald R. Griffin. *Listening in the dark: The acoustic orientation of bats and men.* 1958.

[7] Jan Steckel and Herbert Peremans. A novel biomimetic sonarhead using beamforming technology to mimic bat echolocation. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 59(7):1369–1377, 2012.

[8] Jan Steckel and Herbert Peremans. BatSLAM: Simultaneous Localization and Mapping Using Biomimetic Sonar. *PLoS ONE*, 8(1), 2013.

[9] Jan Steckel, Andre Boen, and Herbert Peremans. Broadband 3-D sonar system using a sparse array for indoor navigation. *IEEE Transactions on Robotics*, 29(1):161–171, 2013.

[10] Jan Steckel and Herbert Peremans. Spatial sampling strategy for a 3D sonar sensor supporting BatSLAM. *IEEE International Conference on Intelligent Robots and Systems*, 2015-Decem:723–728, 2015.

[11] Jan Steckel and Herbert Peremans. Acoustic Flow-Based Control of a Mobile Platform Using a 3D Sonar Sensor. *IEEE Sensors Journal*, 17(10):3131–3141, 2017.

[12] Robin Kerstens, Dennis Laurijssen, and Jan Steckel. Low-cost One-bit MEMS Microphone Arrays for In-air Acoustic Imaging Using FPGA ' s. In *IEEE SENSORS*, pages 1–3, Glasgow, 2017.

[13] Altera. Cyclone V Devices Datasheet. 2012.

[14] Senscomp. Series 7000 Ultrasonic Sensor Datasheet. (734):1–3, 2014.

[15] H L Van Trees. *Optimum Array Processing: Part IV of Detection, Estimation, and Modulation Theory*. Detection, Estimation, and Modulation Theory. Wiley, 2004.

[16] Jan Steckel and Herbert Peremans. Sparse decomposition of in-air sonar images for object localization. *IEEE SENSORS 2014 Proceedings*, (1):1356–1359, 2014.

[17] Knowles. SPH0641LM4H Datasheet, 2014.

[18] Daniel Dunbar and Greg Humphreys. A spatial data structure for fast Poisson-disk sample generation. *ACM Transactions on Graphics*, 25(3):503, 2006.

[19] Rob Roy. ODROID-XU4 Manual.

[20] Damien Browne and Lindsay Kleeman. An advanced sonar ring design with 48 channels of continuous echo processing using matched filters. *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2009*, pages 4040–4046, 2009.

[21] Damien C. Browne and Lindsay Kleeman. A double refresh rate sonar ring with FPGA-based continuous matched filtering. *Robotica*, 30(7):1051–1062, 2012.

[22] Jan Steckel. Sonar system combining an emitter array with a sparse receiver array for air-coupled applications. *IEEE Sensors Journal*, 15(6):3446–3452, 2015.

[23] Robert Bridson. Fast Poisson disk sampling in arbitrary dimensions. *ACM SIGGRAPH 2007 sketches on - SIGGRAPH '07*, pages 22–es, 2007.

[24] Patrick Pfaff, Wolfram Burgard, and Dieter Fox. Robust {M}onte-{C}arlo Localization using Adaptive Likelihood Models. 22(Mcl):181–194.

[25] Paul Leopardi. A partition of the unit sphere into regions of equal measure and small diameter. *Electronic Transactions on Numerical Analysis*, 25(309-327), 2006.