

A Low-Frequency Sonar for Sensor-Adaptive, Multi-Static, Detection and Classification of Underwater Targets with AUVs

Donald P. Eickstedt and Henrik Schmidt
Massachusetts Institute of Technology
292 Main Street, Bldg. E38-300
Cambridge, MA 02139-4309
eicksted@mit.edu, henrik@keel.mit.edu

Abstract- The Generic Ocean Array Technology Sonar (GOATS) Joint Research Program is currently exploring the development of sensor-adaptive autonomous underwater vehicle (AUV) technology specifically directed toward rapid environmental assessment and mine countermeasures in coastal environments. As part of the effort, MIT is developing the GOATS multi-static sonar concept with the goal of being able to detect and classify targets on and within the seabed in very shallow water. To support this project, a high performance, multi-channel sonar was developed as a payload for the Odyssey-III class AUVs used by MIT and a number of other institutions. This sonar payload is capable of low/mid frequency target insonification, real-time processing of 16 array channels and communication with the main vehicle computer for real-time, adaptive mission coordination. The sonar payload also has the capability to time synchronize with the payloads on other AUVs with microsecond accuracy and the ability to maintain tight time synchronization while all vehicles are underwater, thereby enabling multi-static processing of sonar returns. The entire sonar payload was designed and implemented as a virtual sensor with respect to the main vehicle computer. Results from two major international experiments are presented in which both the time synchronization capabilities and the online detection capabilities of the sonar were tested and in which we believe the first ever online target detections were achieved with a system of this type.

I. INTRODUCTION

As was dramatically seen recently in the mine clearing operation in the shipping channel leading to the Iraqi port of Umm Qasr, the need for a robust, automated method to detect and classify underwater targets has never been greater. In that case, both human divers and dolphins were put at risk to search for and identify underwater mines buried in an attempt to deny access to the port. The fact that the clearing operation took several weeks and involved a large number of ships, helicopters, and personnel to open a single channel is a testament to the need for better methods.

One new paradigm, which may be instrumental in addressing problems of this type, is that of the Autonomous Oceanographic Sampling Network (AOSN) [1]. AOSN is comprised of a flexible network of surface vessels, fixed moorings, and AUVs linked together with robust acoustic communications. One of the features of the AOSN paradigm that is ideally suited for the target search and identification problem is its sensor-adaptive nature, i.e. the pre-planned movement of the sensors (the AUVs in this case) can change according to the nature of their sensor readings. This allows

for the possibility of adaptively cooperating sensors and the optimization of sensor movement.

For a number of years, in collaboration with the NATO SACLANT Undersea Research Center (SACLANTCEN) and supported by the Office of Naval Research, the MIT Ocean Engineering Department has been participating in the Generic Ocean Array Technology Sonar Joint Research Program (GOATS) [2]. The GOATS program, a subprogram of AOSN, has as its goal the detection and classification of both proud and buried targets in very shallow water. In line with the AOSN paradigm, this is to be done by enabling a fleet of adaptively cooperating AUVs (communicating via acoustic modems) equipped with acoustic receiving arrays to process multi-static sonar return data from targets insonified with a low-frequency acoustic source mounted on one of the vehicles. Because the scattered field is not spatially isotropic in general it is thought that, by analyzing the spatial and temporal nature of the multi-static returns, we can thereby simultaneously detect and classify the target in real-time.

We believe that our proposed method has significant advantages over traditional sonar methods for finding and identifying proud and buried targets. First and foremost, since the scattered field is non-isotropic and varies by the insonification angle, the mono-static sonar method can only, by its very nature, capture a small part of the structure of the scattered field. Second, by using a relatively low-frequency acoustic source we not only gain significant bottom penetration for detection of buried targets but we also operate in the mid-frequency band (with respect to the target size) where both geometric and resonant scattering are significant. [3]. We believe that capture of this scattering data is the key to real-time detection and classification. Third, by using the AOSN paradigm, we enable real-time transmission of target information to other nodes in the network.

The scope of the GOATS program includes the development of all the technology associated with the fleet of cooperating AUVs including the AUVs themselves, the navigation technology, acoustic communications technology, sonar payload and the associated algorithms for target detection and classification and adaptive vehicle control. Three GOATS experiments have tested various pieces of the system with the ultimate goal of a complete system test in 2004. During the GOATS 1998 experiment, acoustic scattering data was gathered by an AUV from targets insonified by an acoustic source mounted on a fixed tower. During the GOATS 2000 experiment, the time

synchronization for the sonar payload was tested, again using data collected by an AUV from targets insonified by a tower-mounted source. During the GOATS 2002 experiment, the full sonar payload was tested using a mono-static configuration with a single AUV. It is anticipated that the GOATS 2004 experiment will test the full multi-static target detection and classification scenario described above. The remainder of this paper will describe in detail the design and implementation of the sonar payload used to support the goals of the GOATS program as described above. Results gathered from the sonar during two international experiments will be presented. Finally, a short discussion of ongoing research related to the sonar will be presented.

II. SONAR DESIGN GOALS

Given our principal goal of being able to perform multi-static detection and classification of proud and buried targets in real-time using AUVs, the following design goals were adopted for the sonar payload:

A. Mechanical/Electrical

The sonar payload should integrate with the Odyssey-III AUV manufactured by the Bluefin Robotics Corporation and currently used by both MIT and the U.S. Navy as research vehicles (Fig. 1). The power consumption of the payload should allow at least three hours of vehicle run-time before recharging of the batteries is necessary. Because the application demands operation only in shallow or very shallow water, a depth rating of 100 msw is considered sufficient. The pressure vessel should be capable of dissipating up to 150 watts of internally generated heat while keeping the internal temperature at a maximum of 45 degrees Celsius while underwater.

B. Receiving Array Integration

The sonar payload should be capable of integrating with both single and dual 16-element line arrays.

C. Acoustic Source

The sonar payload should be able to drive an acoustic source with up to 200 watts of power in the 4-24 kHz range. The acoustic source should be side-looking with respect to the main vehicle axis and be adjustable in angle from fully horizontal to fully vertical downward. The sonar payload should have the capability to drive the acoustic source with a number of different waveforms that can be chosen on the fly.

D. Data Acquisition

Given the need for high fidelity processing of the sonar data, the sonar data acquisition should provide 16-channel simultaneous sampling with a precision of at least 16 bits and a maximum sampling rate of 100 kHz. In order to obviate the need for analog anti-aliasing filters, sigma-delta conversion will be used on all analog to digital (A/D) converters. All acquired data samples should be saved to hard disk for offline processing.



Fig. 1. Odyssey-III AUV with single line array.

E. Time Synchronization

The sonar payload should have the ability to synchronize itself to Coordinated Universal Time (UTC) with an accuracy of at least 1 microsecond while the AUV is on the surface. While submerged, the sonar time reference should drift no more than 1 microsecond per hour. The payload time reference system should be able to interface to the data acquisition system such that time tags can be generated by the time reference system in response to hardware triggers from the data acquisition system (for the purpose of time-tagging acquired data samples).

F. Vehicle Communications

The sonar payload should have the ability to communicate with the main vehicle computer via Ethernet.

G. Integration as a Virtual Sensor

One of the key robotic techniques that will enable advanced, high-level vehicle control in a cooperative, multi-vehicle framework is the concept of the virtual sensor. Under this concept, a sensor (the sonar payload in this case) will communicate with the main vehicle operating system to send meta-data and requests and to receive commands. Meta-data is high-level, processed data as opposed to raw data samples. For example, a piece of meta-data that might be sent from the sonar payload to the vehicle operating system is "target detected at coordinate (x,y,z)". As can be seen, this technique uses a level of data abstraction that is one level higher than is normally seen from a sensor, e.g. converting a voltage signal into a pressure. Another hallmark of the virtual sensor concept is that it allows the sensor to send requests to the vehicle operating system. For example, the sonar payload may send requests for specific vehicle movements as needed to optimize the target detection and classification based upon its processing of received data. The vehicle operating system determines how best to deal with requests from all virtual sensors in accordance with mission parameters. The integration of the sonar payload as a virtual

sensor requires it to have full-duplex communications with the main vehicle operating system as well as the capability to run whatever algorithms are necessary for the real-time detection and classification of targets.

H. *Ease of use as a research platform*

Because this system is intended to be used as a research platform maintained and operated by graduate students, ease of use and the ability to upgrade or change the system are of considerable importance. Every effort should be made to use commercial off-the-shelf hardware components and software that is well known and widely available to program the system.

III. SONAR IMPLEMENTATION

A functional block diagram of the sonar implementation is shown in Fig. 2.

A. *Mechanical/Electrical*

A standard Bluefin Robotics Corporation aluminum pressure vessel was used to house the sonar electronics. This pressure vessel is cylindrical with hemispherical end caps. It can be seen in Fig. 3 resting in a standard Odyssey-III payload section. This pressure vessel is capable of dissipating up to 150 watts of internally generated heat while maintaining a temperature below 45 degrees Celsius while underwater. The total power consumption of the payload is approximately 115 watts while in operation with the acoustic source transmitting three to five pings per second. The only connection between the sonar payload and the main vehicle consists of a 28-volt (nominal) power connection and a 10 Mbps Ethernet connection. Note the acoustic source mounted on a mechanism that allows it to be rotated between fully horizontal and fully vertical.




Fig. 2. Sonar functional block diagram.




Fig. 3. Sonar pressure vessel in Odyssey-III payload section.

B. *Receiving Array Integration*

The sonar payload was configured to operate with both a 16-element single line array and 16-element dual line array. Fig. 1 shows the Odyssey vehicle with the single line array while Fig. 4 shows it with the dual line array.

C. *Analog Processing Section*

A high-performance preamplifier with low noise characteristics is vital to obtaining quality sonar data. The preamplifier used in the MIT sonar system was designed and built at the Woods Hole Oceanographic Institution. It is needed to boost the signal levels from the hydrophones up to the 4-volt p-p needed by the digitizers in the signal processing section. It is capable of three levels of gain, 0 dB, +20 dB, and +40 dB, selectable on the fly from the signal processing section via a control line. The preamplifier was designed to have a relatively flat pass band in the 4-24 kHz region.



Fig. 4. Odyssey-III AUV with dual line array.

D. Host Computer Architecture

The major issue driving the architecture of the host computer platform is that of having enough bandwidth on the host computer's data bus to successfully stream raw data samples to the hard disk from the data acquisition system. Sampling 16 channels at 100 kHz using 16-bit samples produces a data rate of 3.2 megabytes per second. The two computer architectures that were considered were the PC-104 and the Peripheral Component Interface (PCI) bus architectures. The advantages of the PC-104 architecture are in its compact size and low power consumption. However, one major flaw is its 16-bit data bus as opposed to the 32-bit bus of the PCI architecture. It was felt that the PC-104 data bus would be only marginally capable of performing the continuous data transfer without even taking into account use of the bus by other peripheral cards. A lack of available off-the-shelf, high-performance, 16-channel signal processing peripheral cards was also a factor.

Taking those factors into consideration, the PCI architecture was chosen for the sonar host computer. In order to minimize space requirements in the pressure vessel, a passive backplane construction was chosen utilizing a single-board computer. A single board computer is a complete computer system on a PCI-compatible board. The single-board computer card is inserted into the passive backplane PCI bus thereby giving other peripheral cards access to the CPU. The board chosen for this implementation was based on a 266 MHz Pentium processor with 128 megabytes of main memory. Since this board did not have an integrated network card, a stand-alone network card was used in one of the free PCI slots.

Another significant decision for the host computer architecture is the choice of operating system. Both Linux and Microsoft Windows were considered. In this case, Linux was chosen for its cost (free), ease of programming (most graduate students are familiar with programming in Linux but not in Windows), and its speed. One drawback to the choice of Linux is the lack of availability of commercial device drivers for many peripheral cards of interest.

E. Data Acquisition Subsystem

The data acquisition subsystem is the most critical part of the sonar. High quality, low noise data is essential for detecting and classifying targets in the ocean environment, especially when the data analysis must be done in real-time without the benefit of powerful computers and hours of offline processing.

The acquisition system chosen for the sonar payload is based on the Heron modular digital signal processing (DSP) system manufactured by Hunt Engineering. This DSP system is specifically designed for demanding, real-time applications. It utilizes a modular, extensible hardware architecture that accommodates multiple TMS320C6000 DSP processors, multiple Virtex-II floating point gate arrays (FPGAs), and fully integrated ultra-fast A/D, D/A and Digital I/O interfaces.

For this system, the Heron HEPC8 Module Carrier Board was chosen. This PCI form factor module carrier board supports up to four Heron modules which can be multiple



Fig. 5. HEPC8 module carrier board.

combinations of DSP, FPGA, A/D, D/A, or I/O modules. The HEPC8 provides 32-bit first-in, first-out (FIFO) buffers between each module slot and the other modules slots on the board for data transfer between Heron modules. One FIFO on the board is also connected to the board PCI interface for data transfer between a Heron DSP module and the host computer. The HEPC8 is shown in Fig. 5.

The HEPC8 in the sonar payload contained four Heron modules; one Heron4-C6701 floating point DSP, one HEGD5 D/A converter, and two HEGD-12 A/D converters. A functional block diagram of the data acquisition subsystem utilizing the Heron modular system is shown in Fig. 6.

The DSP in this system has a number of roles. First, it waits for commands from the host computer (sent over the PCI bus) to begin sampling. Second, it outputs the transmit waveform to the D/A converter module and triggers the power amplifier via an output control line. Third, it sends a TTL pulse trigger to the timing subsystem to generate a time tag via another control line. Fourth, it collects incoming data samples from the two A/D converters and transfers them to the host computer over the PCI bus. The DSP in this system was not used for signal processing but may be in future revisions. The Heron4-C6701 DSP module contains a 167 MHz Texas Instruments TMS320C6201 floating point DSP capable of 1 gigaflop of performance. The DSP module also provides numerous digital input and outputs that can be used for control of or communication with other system components.




Fig. 6. Data acquisition subsystem functional diagram.

The HEGD5 is a four-channel, 16-bit D/A converter that is used to receive a digitized waveform from the DSP via a 32-bit FIFO and to output an analog signal to the power amplifier for the acoustic source. The sampling rate of the HEGD5 is 230 kHz.

The HEGD12 is an eight-channel, 16-bit A/D converter that is used to digitize the input waveform received from the preamplifier. The HEGD12 provides simultaneous sampling on all eight channels and utilizes sigma-delta conversion on all channels. The sigma-delta conversion technique involves over-sampling the input waveform by a factor of eight times, digitally filtering the over-sampled data, and then down-sampling the resulting data by a factor of eight. This technique eliminates the need for analog anti-aliasing filters. An onboard crystal oscillator provides the sample clock for the HEGD12. The sample clocks of the two HEGD12 modules are tied together in order to provide 16-channel simultaneous sampling. Digitized samples are pushed into a 32-bit FIFO where they are read by the DSP.

F. Time Reference Subsystem

The time reference subsystem is the key to being able to use multi-static sonar techniques for detection and classification due to the need for accurate travel times between waveform transmission by the acoustic source and reception of the scattered waveform by each of the receiving arrays on multiple vehicles. In order to accomplish this, the sonar payloads on each vehicle must be synchronized in time. This synchronization must be maintained for the duration of the mission. The solution to this problem is to synchronize each vehicle on the surface via GPS and to maintain this synchronization with an internal oscillator while the vehicles are underwater and unable to access the GPS satellites.

This was accomplished in the MIT sonar payload by using a digital clock card. This card, the Synclock-32 manufactured by JXI2 incorporated, has an onboard GPS receiver which can be used to synchronize the clock on the card to GPS time to within 500 nanoseconds of UTC. Once synchronized, however, all clocks will drift (either fast or slow) and the Synclock-32 will begin to drift once the GPS input is gone. The rate of drift will depend on the quality of the oscillator that is used to keep time on the clock. Unfortunately, standard crystal oscillators do not have the capability to meet the stringent drift requirement of 1 microsecond per hour that is needed for this application. The solution to this problem is to use a rubidium oscillator to keep time. A rubidium oscillator is actually a small atomic clock. The rubidium oscillator chosen for the sonar payload is the RMO rubidium oscillator manufactured by Temex in Switzerland. This oscillator has a drift specification of 700 nanoseconds per hour and consumes 10 watts of power. See Fig. 7 for a picture of the RMO.

Time tags can be generated by the Synclock-32 upon reception of a digital control input from the DSP. Once the time is latched, the Synclock-32 generates a PCI bus interrupt that is intercepted by a device driver that reads the time tag from the Synclock-32 and places it in a file.



Fig. 7. RMO rubidium oscillator.

Since the payload on each vehicle can be synchronized within 500 nanoseconds of UTC, the maximum initial time error between the transmitting vehicle and each and receiving vehicle is 1 microsecond.

G. Acoustic Source

The acoustic source used in the payload is the acoustic source from the SB-24 sub-bottom profiler tow sled manufactured by EdgeTech. It has a frequency range of 4-24 kHz and a beam width of 15-30 degrees depending on the transmit frequency. The power amplifier used to drive the source has a maximum output power of 200 watts.

H. Software Architecture

The software architecture of the sonar includes processes running on two different processors; the DSP in the data acquisition subsystem and the host computer.

The software on the host computer is responsible for bi-directional communications with the main vehicle operating system, booting the DSP with an executable image, bi-directional communications with the process running on the DSP, transferring sample data from the DSP and saving it to the hard disk, and for running any data analysis algorithms such as detection and classification algorithms. A single multi-threaded process on the host computer accomplishes the first four of these tasks while the detection and classification algorithms are run simultaneously as individual processes. The main communications process is compiled with a library provided by the main vehicle operating system which makes communication with the operating system as simple as making a function call. An Application Programming Interface (API) provided by Hunt Engineering allows booting of the DSP with an executable image. Communications with the DSP is accomplished through a set of function calls provided by the API. All code on the host computer is written in C or C++ using the Linux compiler.

The process running on the DSP is responsible for communicating with the host computer to send status or to read commands, reading samples from the A/D converters, transmitting the output waveform, triggering the time tag subsystem and for transferring sample data to the host computer. All code development for the DSP was done using the Texas Instruments Code Composer Studio. Communications between the DSP and host processor was accomplished by compiling the DSP code with an API provided by Hunt Engineering.

IV. EXPERIMENTAL RESULTS

The sonar payload was tested in two different international experiments. The GOATS 2000 experiment, conducted in conjunction with SACLANTCEN off of Elba Island Italy, tested only the time synchronization subsystem while the GOATS 2002 experiment, also held in conjunction with SACLANTCEN off the coast of Italy, tested the full sonar payload in a mono-static configuration with a single AUV. During the GOATS 2000 experiment, the time synchronization subsystem was successful in synchronizing the digital clock card to GPS time and in maintaining the time synchronization with the rubidium oscillator while the vehicle was underwater. The system was also successful in time-tagging sonar data collected by the AUV. During this experiment, the underwater targets were insonified by a tower-mounted acoustic source. The transmissions from the acoustic source were GPS time-tagged by the shore station that was controlling the transmissions. Fig. 8 and Fig. 9 show a set of results obtained during one AUV run to try and determine the accuracy of the time tagging concept. The vertical axis of each figure represents a sequential set of sonar reception windows (one for each ping of the acoustic source) as the AUV swims past the target being insonified. The horizontal axis represents sequential time in each window. Fig. 8 shows an attempt to align the sonar reception windows without the benefit of the time tag information as was done in

previous experiments. Fig. 9 shows the alignment using the time tag data from the shore station and the AUV. This allows each data reception window to be aligned (in time) exactly with the previous window. As can be seen by examining the two figures, the best manual alignment method was clearly inferior to the alignment achieved with the time-tagged data. Fig. 8. clearly shows that a lack of time synchronization between the transmitting station and the receiving station results in a large variance in the time-of-flight estimates for the sonar ping while Fig. 9 clearly shows the familiar hyperbolic pattern expected from an AUV swimming past a point target.

These results have important implications for our multi-static detection and classification concept. They show that our concept for synchronizing multiple sonar transmitter and receiver platforms is valid. Extremely accurate time-of-flight data is necessary for meeting our goal of sub-meter target localization accuracy.

A major goal of the GOATS 2002 experiment was to test the complete sonar payload in an operational environment. This involved programming the AUV to conduct a sonar search of a patch of sea bottom where several buried and proud targets had been placed. A simple target detection algorithm, developed by several graduate students in the department, was run on the sonar host computer.




Fig. 8. Sonar data without time tags.




Fig. 9. Sonar data with time tags.

This detection algorithm was fed a continuous stream of sonar data as the AUV conducted a search pattern in the area of interest. Detections were logged in a file for offline analysis. Fig 10 shows a plot of a representative survey mission. This type of search profile is called a “Zamboni” and is shaped this way due to the fact that the sonar’s acoustic source can only transmit on one side of the AUV.

The online detection algorithm was successful in detecting a number of targets in real time. Fig 11 shows the raw sonar data for one such online detection [4]. The sonar returns from the detected target can be seen inside the white box on the figure. Note the familiar hyperbolic shape of the target plot. This is exactly the pattern we expect to see from an AUV moving in a straight line past a point target. These extremely important results validate our sonar design and show that the data quality from the sonar is high enough to be used for real-time target detection.

Although more extensive online target tracking and classification algorithms were not available at the time of the GOATS 2002 experiment, extensive offline analysis of the sonar data continues to show its superb quality and suitability. One of the challenges to online target detection is in trying to sort out spurious returns from those returns from actual point targets. Fig. 12 and Fig. 13 show sonar data from one AUV run which has been processed using algorithms developed by a department graduate student [5]. Fig. 12 shows a plot of the raw sonar data. Fig. 13 shows a plot of the same data after it has been run through a constraint-based algorithm that is designed to pick out point targets in the data. As can be seen from the plots, almost all of the clutter, including bottom and surface bounces, has been removed leaving only the tracks of several point targets. This once again validates our sonar design and the quality of the data. It is anticipated that these algorithms will be incorporated into the sonar in time for the GOATS 2004 experiment.




Fig 11. Online target detection.




Fig 12. Raw sonar data.




Fig. 10. Target survey pattern.




Fig 13. Offline-processed sonar data.

V. CURRENT RESEARCH

Now that a proper sonar instrument has been built and validated, current research is focused on the development of a number of algorithms for online use on the sonar. One major focus is on the development of a multi-static detection algorithm. This algorithm will use sonar data from multiple AUVs to detect, track, and localize targets. A second major focus is on the development of an algorithm for online target classification. It is anticipated that this will be a model-based algorithm that will be able to distinguish between targets from a finite set of different target shapes (e.g. cylinders, spheres, etc). A third area of focus is on developing algorithms for determining optimal vehicle motion for target detection and classification.

REFERENCES

- [1]. T. Curtin, J.G. Bellingham, J. Catipovic, and D. Webb, "Autonomous Oceanographic Sampling Networks," *Oceanography*, vol. 6, issue 3, pp. 86-94, 1993.
- [2]. H. Schmidt, J.R. Edwards, and K.D. LePage, "Bistatic Synthetic Aperture Sonar Concept for MCM AUV Networks," *International Workshop on Sensors and Sensing Technology for Autonomous Ocean Systems*, Kona, Hawaii, 2000.
- [3]. H. Schmidt and J. Lee, "Physics of 3-D Scattering from Rippled Seabeds and Buried Targets in Shallow Water," *J. Acoust. Soc. Am.*, vol 105, pp. 1605-1617, 1999.
- [4]. M. Montanari, J.R. Edwards, and H. Schmidt, "Using Autonomous Underwater Vehicles for Concurrent Detection and Classification by Means of Higher-Order Spectral analysis," *unpublished*.
- [5]. R. Rikoski, "Dynamic Sonar Perception," *Ph.D. Thesis*, 2003.