

Peer-Reviewed Technical Communication

Development of a Second-Generation Underwater Acoustic Ambient Noise Imaging Camera

Venugopalan Pallayil, *Senior Member, IEEE*, Mandar Chitre, *Senior Member, IEEE*, Subash Kuselan, Amogh Raichur, Manu Ignatius, *Member, IEEE*, and John R. Potter, *Senior Member, IEEE*

Abstract—A nominally circular 2-D broadband acoustic array of 1.3-m diameter, comprising 508 sensors and associated electronics, was designed, built, and tested for ambient noise imaging (ANI) potential in Singapore waters. The system, named Remotely Operated Mobile Ambient Noise Imaging System (ROMANIS), operates over 25–85 kHz, streaming real-time data at 1.6 Gb/s over a fiber optic link. By using sensors that are much larger than half-wavelength at the highest frequency of interest, so with some directionality, good beamforming performance is obtained with a small number of sensors compared to a conventional half-wavelength-spaced array. A data acquisition system consisting of eight single-board computers enables synchronous data collection from all 508 sensors. A dry-coupled neoprene cover is used to encapsulate the ceramic elements as an alternative to potting or oil filling, for easier maintenance. Beamforming is performed in real-time using parallel computing on a graphics processing unit (GPU). Experiments conducted in Singapore waters yielded images of underwater objects at much larger ranges and with better resolution than any previous ANI system. Although ROMANIS was designed for ANI, the array may be valuable in many other applications requiring a broadband underwater acoustic receiving array.

Index Terms—Ambient noise imaging (ANI), broadband array design, underwater acoustics, data acquisition.

I. INTRODUCTION

THE idea of using ambient noise for underwater imaging applications has been explored by several researchers [1]–[6]. The first ambient noise imaging (ANI) camera, the

Manuscript received July 19, 2014; accepted December 11, 2014. Date of publication January 16, 2015; date of current version January 11, 2016.

Associate Editor: A. Morozov.

V. Pallayil and M. Chitre are with the Acoustic Research Laboratory (ARL), Tropical Marine Science Institute (TMSI), National University of Singapore (NUS), Singapore 119227, Singapore (e-mail: venu@arl.nus.edu.sg; mandar@arl.nus.edu.sg).

S. Kuselan was with the Acoustic Research Laboratory (ARL), Tropical Marine Science Institute (TMSI), National University of Singapore (NUS), Singapore 119227, Singapore.

A. Raichur was with the Acoustic Research Laboratory (ARL), Tropical Marine Science Institute (TMSI), National University of Singapore (NUS), Singapore 119227, Singapore, he is now with Finisar Inc., 569059, Singapore.

M. Ignatius was with the Acoustic Research Laboratory (ARL), Tropical Marine Science Institute (TMSI), National University of Singapore (NUS), Singapore 119227, Singapore, he is now with Subnero, Singapore.

J. Potter was with Acoustic Research Laboratory (ARL), Tropical Marine Science Institute (TMSI), National University of Singapore (NUS), Singapore 119227, Singapore. He is now with the Centre for Maritime Research and Experimentation (CMRE), NATO Science & Technology Organization (STO), La Spezia 19126, Italy.

Digital Object Identifier 10.1109/JOE.2014.2386783

Acoustic Daylight Ocean Noise Imaging System (ADONIS), was successfully built and tested in 1994 at the Scripps Institute of Oceanography (La Jolla, CA, USA) [7]. Since then, two other ANI systems have been built: the Remotely Operated Mobile Ambient Noise Imaging System (ROMANIS), at the Acoustic Research Laboratory (ARL), National University of Singapore (NUS, Singapore), and the imaging array built at the Defence Science and Technology Organization (DSTO, Maritime Operations Division, Sydney, Australia) [8]. A fourth ANI system is currently being developed at the National Defense Academy (Yokosuka, Japan) [9]. Table I shows a comparison of various attributes of the four imaging systems.

The ADONIS system was able to produce images of underwater objects, both static and moving, at ranges of about 40 m, using ambient noise as the main source of illumination [7]. It was also able to discriminate, to some extent, between various materials of the object through “acoustic color” processing. Nevertheless, due to the specific system design decisions, it had several limitations. The frequency spectrum for each beam was estimated using an analog filter, which was switched to each of the 16 frequencies used for imaging in turn. Allowing for settling time and an extra period at the end of each frame cycle, a lot of data were effectively discarded. The frame rate was too low to track rapid temporal features of ambient noise. Finally, as energy estimates in frequency bins were recorded, phase information was effectively discarded. Therefore, only incoherent imaging algorithms using first- and second-order statistics of the energy could be applied [10].

With regard to the Australian and Japanese ANI systems, performance evaluations are not available in the open literature at present.

In 1998, we initiated the development of ROMANIS—a second-generation broadband ANI camera that could effectively address many of the limitations of ADONIS. After four years of development work, the first prototype ROMANIS array was completed in 2002. Data collected during a deployment in Singapore waters in 2003 produced an image of an underwater object at about 70-m range [11]. Although this preliminary result was encouraging, ROMANIS was plagued by electronic stability problems due to the large power consumption and related thermal issues and the extremely high data acquisition rate. Further, the computing technologies available at that time did not permit real-time analysis of the data. Post-processing took hours of computing time for each second of collected data.

As new and faster technologies became available, solutions to these problems became feasible. The electronics and software of the ROMANIS system were completely refurbished in 2009 to deliver efficient and reliable operation [12], [13]. Newer and faster processing platforms coupled with optimized algorithms allowed near-real-time imaging in the field. In March 2010, the upgraded ROMANIS system was deployed for testing in Singapore waters. Data collected during this experiment allowed us to apply novel algorithms to produce high-quality acoustic images and videos of various underwater objects, and in some cases determine the range to those objects passively [14].

In a decade of research and development of ROMANIS toward a reliable, maintainable, and near-real-time ANI system, several novel ideas in terms of sparse array design, mechanical encapsulation using dry coupling, distributed data acquisition, and highly parallelized beamforming on a graphical processing unit (GPU) were developed. In this paper, we provide an overview of the key design features of ROMANIS as well as some of the results of ANI from a local field trial conducted in 2010. Many of the design ideas presented here may also be applied to other systems that require high-speed (gigabits per second) data acquisition and real-time processing capability.

Section II focuses on the array design for optimal beamforming performance. This is followed by details of the electronics needed to acquire the data from the array in Section III. Section IV provides details of the mechanical design and encapsulation of the sensors. In Section V, the design of the distributed data acquisition and GPU-based beamforming software architecture is presented. Section VI presents some results from the 2010 field experiments. Finally, Section VII draws conclusions and presents some future directions for the research.

II. ARRAY DESIGN

The sensor array or “acoustic eye” of ROMANIS was a key part of the overall design. The main considerations were to obtain a wide frequency band of operation with the best possible angular resolution within an overall package that was as compact as possible. The original idea was that the array should be a conformal array on a mobile underwater vehicle. It quickly became clear that a 3-D conformal array, coupled with the considerable other challenges associated with the very high data rate, would be too great an engineering challenge to take on in one step. Therefore, ROMANIS was designed as a flat, static array as an intermediate step to this goal.

It is well known that snapping shrimp are the major contributors to high-frequency ambient noise in warm shallow waters and that their acoustic spectrum spans a wide frequency band [15], [16]. For a given aperture size, higher frequencies provide better resolution at the cost of lower effective range as a result of increased attenuation. The use of high frequencies also incurs the need for a higher sampling rate and hence computational power for processing. The anticipated effective imaging range using frequencies above about 85 kHz was expected to drop below 50 m due to absorption. Below 25 kHz, the resolution of a 1.3-m array drops to about 5 m at 100-m range, which was considered borderline useful, most objects of interest having characteristic dimensions of this size or smaller. Since considerable

energy from snapping shrimp is available in the 25–85-kHz frequency band [17], [18], we selected this band for ROMANIS operation.

Our objective was to be able to resolve a 1-m² target at 80-m range, corresponding to an angular resolution of roughly $0.7^\circ \times 0.7^\circ$. This requires a circular planar array aperture of about 1.3-m diameter at the highest frequency of operation. Taking the traditional approach of building a fully populated array with omnidirectional sensors at half-wavelength spacing would require more than 17 000 sensors, each sampled at a minimum of 170 kSa/s. This was obviously prohibitive in cost and required high computing power. Therefore, we sought ways to reduce the number of sensors. ADONIS had achieved this by performing analog beamforming using a reflector. We wanted to retain the phase information for more sophisticated processing, so were obliged to accept the necessity of a sparse array, since neither the aperture nor the frequency could be reduced without compromising resolution. We estimated about 500 sensors to be the limit of what was achievable within the technology available at an acceptable cost, and set out to design a sparse array of 1.3-m diameter that would provide good beamforming performance.

It is well known that sparse arrays exhibit grating lobes due to spatial aliasing. In the case of digitally beamforming a set of identical directional elements, the final beampattern obtained is the product of the directivity of individual elements and the beamformer performance for point-like sensors [19]. If element directionality is derived from the geometric shape and size of the receiving ceramics, then in the case where the ceramic elements can be tessellated, without any gaps in the sensing surface, it is easy to show that for broadside beamforming the first null of the element directivity falls on the first grating lobe of the digital beamforming output, as illustrated in Fig. 1 for our rectangular elements placed side by side. In this way, providing the entire 2-D sensor surface fully populated by sensors (of any size and shape that will tessellate), there is no performance loss when beamforming broadside, perpendicular to the element surface, compared to a fully populated array or single monolithic piston sensor. Once the digital beamformer is steered off-broadside, however, the grating lobes move out of the element nulls (unless these are also physically steered to match) and the degradation of main to grating lobe sensitivity eventually becomes a problem [20].

To reduce the grating lobe problem, we investigated the possible advantage of pseudorandomly placing the individual elements, rather than arranging them in a rectilinear grid. It is well known that in the case of point-like sensors (with dimensions much less than the interelement spacing and hence of negligible directionality) a pseudorandom placement yields advantages. This is because, for a regular periodically spaced array, the spatial aliasing due to the spacing of each pair of sensors occurs at the same angle for a given frequency. By using aperiodic spacing, aliasing occurs at different angles for different sensor pairs and thus the grating lobes tend to become “smeared out” over angle.

In our case, we wanted only a limited field of view (FoV) centered on broadside. In this case, the directionality of elements that were of comparable size to the interelement spacing gave significant benefits, outperforming the pseudorandom spacing

TABLE I
COMPARISON OF ANI SYSTEMS

	ROMANIS	ADONIS	Australia	Japan
Aperture shape	circular	circular	square	circular
Aperture size	1.44 m	3 m	2 m × 2 m	1 m
Beamforming	phased array	spherical reflector	phased array	refracting lens
Bandwidth	25–85 kHz	8–80 kHz	10–150 kHz	10–200 kHz
No. of sensors	508	130	256	128*
Sensor type	directional	omnidirectional	omnidirectional	omnidirectional
Array design	compact, with offsets between rows	elliptical, at the focus	four sub-arrays, randomly populated	unknown shape, at the focus
Approximate weight	500 kg	1000 kg	1000 kg	250 kg
Construction	highly modular	single system	modular	single system
Realtime Imaging	yes	yes	no	yes

* Only 15 sensors were used to evaluate the 1-dimensional performance of the system. The final system is expected to have about 128 sensors arranged in a 2-dimensional pattern.

of point-like sensors. But it seemed that we might yet benefit from a hybrid approach, with physically large sensors that were also perturbed about their nominal regular grid-like placement to mitigate grating lobe artefacts. We investigated this possibility by running a simulated annealing optimization process, combined with elements of a genetic optimization algorithm and principal component analysis, to search the performance space (in terms of central beam sensitivity, beamwidth, and main beam to first sidelobe signal-to-noise ratio) as a function of placement of the 500 sensors [20]. While these numerical experiments (which took several weeks, running in parallel on several computers) were largely unsupervised, they tended to converge on solutions that pointed to two primary factors to optimize performance:

- 1) the ideal offset from a regular grid-like arrangement converged to an integer multiple of the half-wavelength at the highest frequency of operation (8.5 mm);
- 2) if “gaps” opened up in the sensing surface, these caused a loss of performance that outweighed the benefits of the aperiodic element offsets that caused the gap(s).

As a result, we saw time and again that the best results from the optimization processes converged on a contiguous placement of elements, as if they had been “pushed together” to leave as few gaps as possible, with elements forming rows or columns that were offset, like geological strike-slip faultlines, by integer multiples of half-wavelength at 85 kHz (8.5 mm). The half-wavelength offset made intuitive sense, since this would be the optimal spacing if the array were not obliged to be sparse. If gaps were not to be tolerated, however, offsets could only be applied either in rows, or, alternatively, columns. Since we wanted a “letterbox” FoV (i.e., a FoV that is wider than it is high) the main grating lobe problem appeared in azimuth, rather than in elevation. Therefore, it seemed better to create contiguous rows of elements, and stagger these rows by multiples of 8.5

mm, attempting to populate all possible offsets (8.5, 17, 25.5, 34.0, 42.5 mm) for the 50-mm² square sensors (separated one from another by approximately 1 mm of Corprene insulation) as evenly as possible. When such an element map was generated, as illustrated in Fig. 2, it indeed tested as performing better than any of the simulated annealing results. Even though we were not able to state and solve a closed analytical form for this task, the numerical simulation and optimization process guided us to an intuitive understanding of the optimization that enabled a better solution to be found than could be obtained from numerical work alone. Fig. 3 illustrates the simulated beamforming performance of the array shown in Fig. 2 for two specific frequencies and along the broadside and edges of FoV. We were able to achieve an angular resolution of about $0.7^\circ \times 0.7^\circ$ across a FoV of about $17^\circ \times 8.5^\circ$, at the highest frequency of operation. For the shallow waters of Singapore where depths are typically less than 40 m and anticipated ranges are of the order 100 m, this “letterbox” format FoV provides excellent coverage.

The sensors selected for the ROMANIS array were EC-97 ceramic tiles from EDO Ceramics Corporation (now part of ITT Corporation, White Plains, NY, USA). Each sensor is a 49.53 mm × 49.53 mm tile with a thickness of 12.7 mm and weighs about 150 g. These sensors respond to pressure changes only in one plane. This allowed the ceramics to be encapsulated without decoupling pressure release material. Each sensor has a sensitivity of about -190 dB (± 1.5 dB) re 1 V/ μ Pa across the band of interest. The sensor exhibits a sharp resonance-related phase nonlinearity at around 49.5 kHz, and hence a small band around this frequency is not used in data processing. Each sensor module is fitted with a dual-channel preamplifier directly mounted onto the rear of the ceramic, separated only by a decoupling plate. The sensors are mounted on a 1 mm thick steel plate with a Corprene sheet for acoustic decoupling. The sides of the sensors in the modules (excluding the face) are also lined

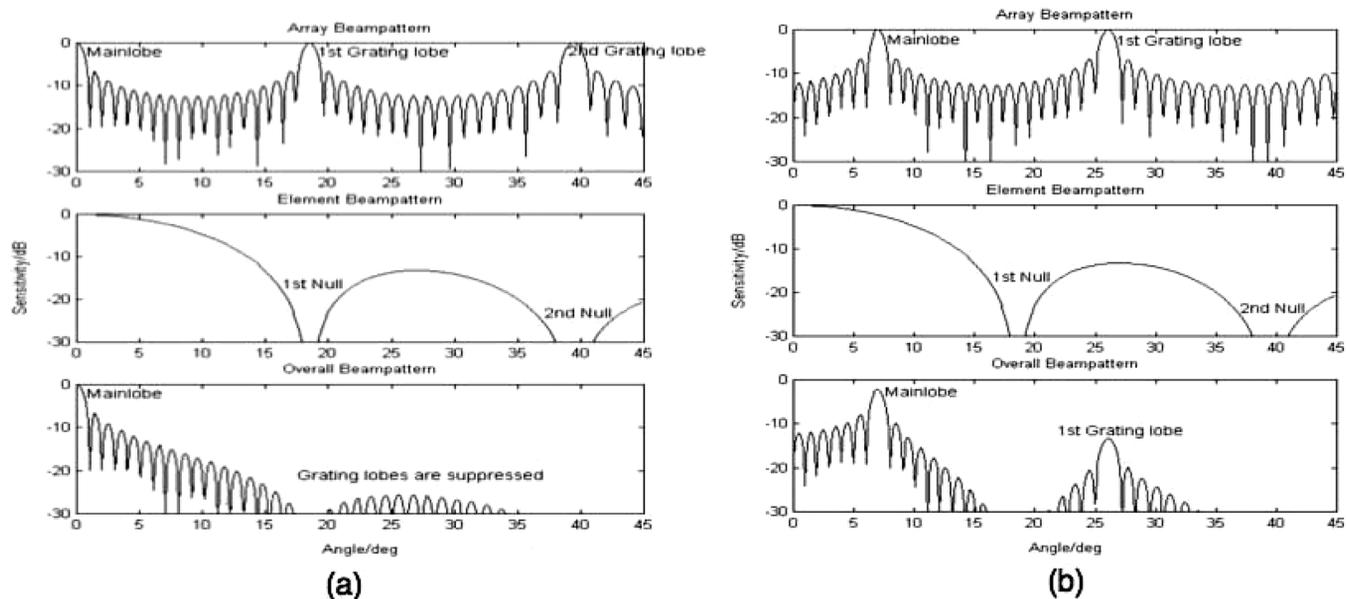


Fig. 1. Directional array beampattern is formed as a product of the omnidirectional array beampattern and the sensor directionality [20]. When the array is steered toward broadside, as shown in (a), the grating lobes of the omnidirectional array beampattern fall into the nulls of the sensor beampattern and are therefore canceled. When the array is steered away from broadside, as shown in (b), the grating lobes move but the sensor directionality does not change. This leads to poorer performance as the grating lobes are not completely canceled by the nulls in the sensor directionality.

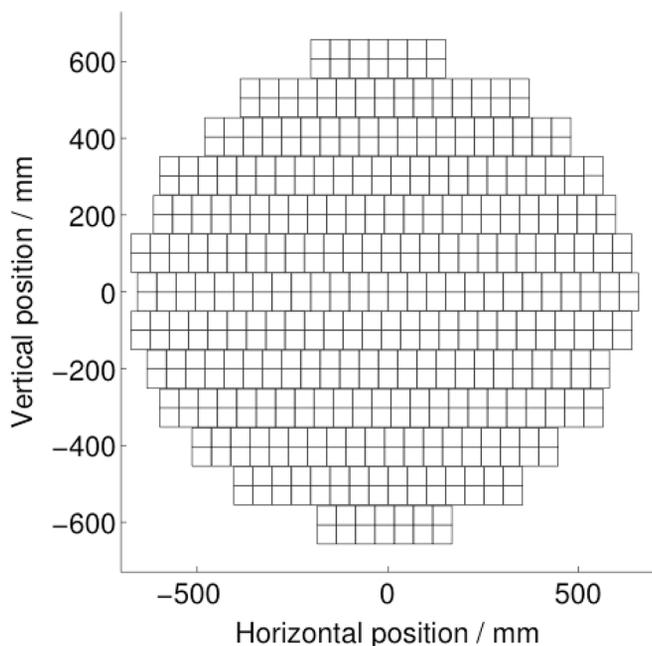


Fig. 2. Sensor placement of ROMANIS array. Rows of sensor modules (pairs of adjacent sensors placed vertically) are offset from each other to enhance beamforming performance

with a 1 mm thick layer of Corprene for acoustic baffling and electrical isolation. Photographs of the sensor module with its preamplifier and the sensor array are shown in Fig. 4. Further assembly details of the array are covered in Section IV.

III. ELECTRONIC DESIGN

A. Data Acquisition

The data acquisition system was required to sample and record data from all the 508 sensors simultaneously to preserve the phase information. This is important for beamforming and

other coherent processing. To avoid temporal aliasing at the highest frequency of operation (85 kHz), a sampling frequency of at least 170 kHz per channel is necessary. We chose a slightly higher sampling rate of 196 kHz with sigma-delta anti-aliasing filters. For 16-bit samples, the overall data acquisition and streaming requires about 1.6 Gb/s. There were no off-the-shelf solutions available to meet these data acquisition and transfer needs when ROMANIS was designed (1997–1999). Therefore, a customized solution was developed.

In the first version of ROMANIS, the data acquisition and management was implemented using fiber channel arbitrated loop (FC-AL) technology, the only high-bandwidth data acquisition and streaming solution that was available at that time. Though the technology was in its infancy, we successfully implemented the FC-AL solution and demonstrated ROMANIS in the field [21]. However, the system suffered from reliability problems due to both hardware and software limitations. Large power dissipation needs further limited the endurance of the system for field operations. Therefore, ROMANIS was rebuilt in 2009 and the FC-AL technology was replaced by the more recent yet robust and widely used gigabit ethernet technology. This approach, along with the availability of high-speed and low-power embedded processors, resulted in a reduction in power consumption of about 70%.

The signal conditioning electronics are built around a very low noise preamplifier, LT1169, from Linear Technology (Bensalem, PA, USA) with a fixed gain of 20 dB. It has low voltage noise as well as very low current noise, both of which are important when working with high impedance sensors. The device provides dual preamplifiers within a very small package. Another low voltage noise amplifier, AD797 from Analog Devices (Norwood, MA, USA), with a gain of 26 dB followed the low noise preamplifier to provide an overall gain of 46 dB for the signal conditioning stage. A 3-bit digitally programmable gain amplifier (PGA), LTC6910-2, provides further gain adjustments

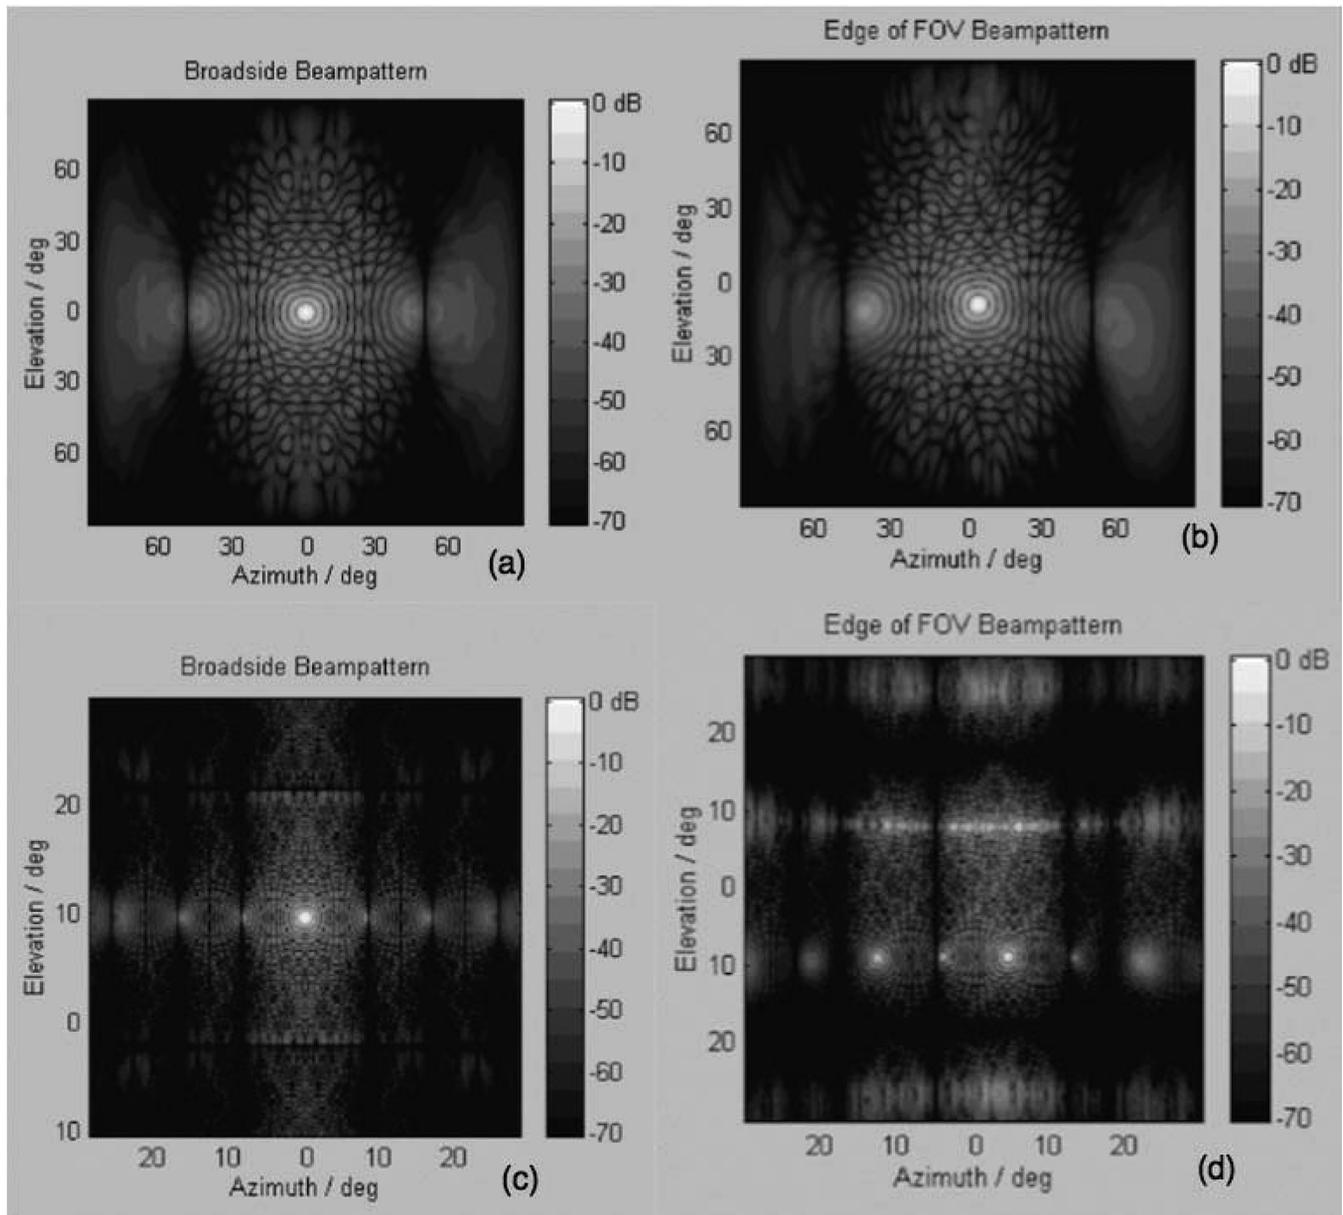


Fig. 3. Simulated beamforming performance of ROMANIS array at the (a)–(b) lower and (c)–(d) higher ends of its frequency bands [20]. The patterns are shown for the best (broadside) and worst (edge of FoV) FoV scenarios. The presence of grating lobes at the edges of FoV can be observed with its sensitivity being higher at higher frequency.

between 0 and 36 dB (in steps of 6 dB). The gain of the PGA can be set in the field, depending on the ambient noise level, to make best use of the full dynamic range of the system. The preamplifier frequency response is shaped using a first-order highpass filter with a cutoff frequency of 25 kHz. The PGA is followed by fourth-order Butterworth lowpass filter, implemented using LTC1563-2 from Linear Technology, with unity gain and 3-dB cutoff frequency of 85 kHz. The low voltage and current requirements of the devices ensure that the power consumption by the signal conditioning card is kept to a minimum. This is a key system power driver, since we require 508 signal conditioning channels in the system. The present signal conditioning electronics consumes approximately 180 W, which is about 40% of the overall power consumption of ROMANIS.

A functional block diagram of the data acquisition module (DAM) is shown in Fig. 5. Each of the eight DAMs consists

of a high-performance single board computer (SBC) in PC104 plus form factor and a 64-channel data acquisition card (DAQ). The SBC used is an Intel Atom 1.6-GHz processor running Tiny Core Linux OS, a lightweight distribution of the Linux OS. The DAQ card, PMC66-16AI64SSC from General Standards (Huntsville, AZ, USA), is capable of simultaneously sampling 64 channels at a rate of 200 kHz per channel with 16-bit resolution. It has provision for an external trigger to initiate the sampling of all the channels synchronously. The DAQ board also features self-test and autocalibration functionality for each channel. The self-test feature helps the host to verify the integrity of the card, while the autocalibration procedure allows offset and gain corrections to be applied to respective channels. The input voltage range is software selectable [22]. The data acquisition process is triggered by the application of external clock to the DAQ card by the system controller (SYSCON).

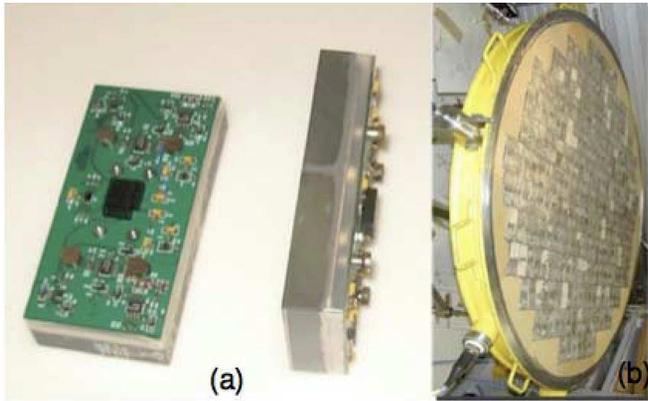


Fig. 4. ROMANIS (a) sensor module and (b) sensor array assembly.

Each channel contains a dedicated 16-bit ADC and the resulting sampled serial data are deserialized and multiplexed into a parallel stream. A 512 000 sample first-in–first-out (FIFO) buffer accumulates the samples for retrieval by the host SBC via direct memory access (DMA) over a peripheral component interface (PCI) bus. A sampling frequency of 196 kHz yields a data transfer rate of just over 200 Mb/s for 64-channel data acquisition. The SBC running Tiny Core Linux operating system (OS) is able to handle the data transfer in real-time. Efficient transfer of the high volume of data from all the sensors to the surface is achieved by the use of Gigabit Ethernet. The SBC has an onboard ethernet network interface controller (NIC) with a TCP offload engine (TOE) to reduce computational burden on the central processing unit (CPU). A custom TCP/IP-based protocol is used to transfer data to the surface receiver node without data loss. Since the combined data rate from all the DAMs is just over 1.6 Gb/s, two parallel Gigabit Ethernet interfaces are required for the data transfer. Data from four DAMs are multiplexed using a Gigabit Ethernet switch, and two such switches are used to multiplex data from the eight DAMs. Two commercial off-the-shelf RJ45 to SFP media converters (MCs) were used for converting digital data into an optical stream that is finally streamed to the surface over a pair of multimode fiber optic cables. The fiber optic cables as well as the cables for supplying power to ROMANIS are combined into a 160-m-long electro-mechanical cable. The data are received at the surface on an industrial computer system running on two Intel Nehalem processors with two onboard Gigabit Ethernet network controllers. The data then flow into a data storage system through another pair of Gigabit Ethernet NIC. The data storage system consists of four 1-TB hard disks configured as a RAID-0 array. This RAID-0 configuration increases the disk write speed by transferring data to all the four disks in parallel.

B. System Controller

The overall system operations are controlled by a system controller (SYSCON). The functions of the SYSCON include:

- power up the signal conditioning boards;
- power up the DAMs;
- supply synchronous clock for all the DAMs;
- supply gain control bits to the signal conditioning boards;
- monitor auxiliary sensors (temperature and leakage);

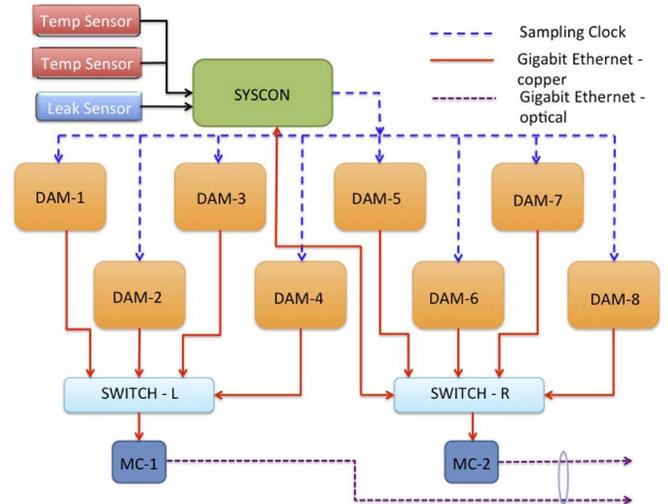


Fig. 5. Block schematic of the data acquisition system architecture [12].

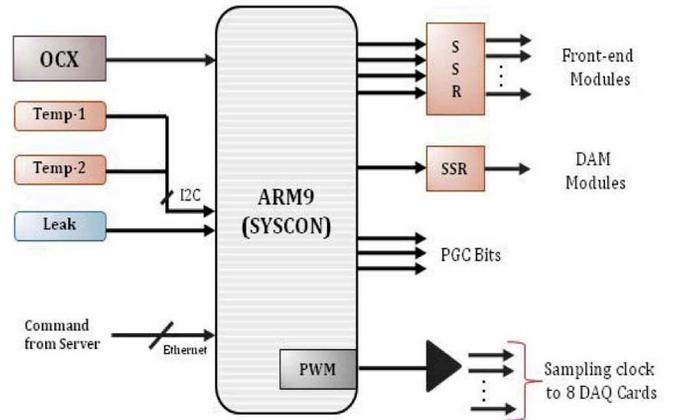


Fig. 6. Block schematic of the system controller for ROMANIS [12].

- warn and shut down the system automatically if any leakage is detected;
- warn the operator if the system temperature exceeds a threshold value.

The SYSCON is built around an ARM STR912, as shown in Fig. 6. The system receives commands from the surface over an ethernet optical link and generates required output to satisfy the various functionalities listed above. The status of the various environmental sensors can also be monitored from the surface by sending appropriate queries.

C. Power Supply Distribution System

ROMANIS is powered using a 148-V direct current (dc) 10 AH battery source from the surface over the 160-m umbilical that links ROMANIS with the surface node. The array electronics runs at 5-V dc; this voltage is generated using dc–dc converters housed inside the array. To keep the initial in-rush current at a minimum, it is necessary to power up the system in stages. Hence, a power supply distribution and control circuitry incorporating solid state relays (SSRs) is used. The electronics sections can be powered selectively by sending commands from the surface to the relevant SSR through SYSCON. The dc–dc converter units employ Densi–Lambda power supplies. Although they occupy substantial space inside the array,

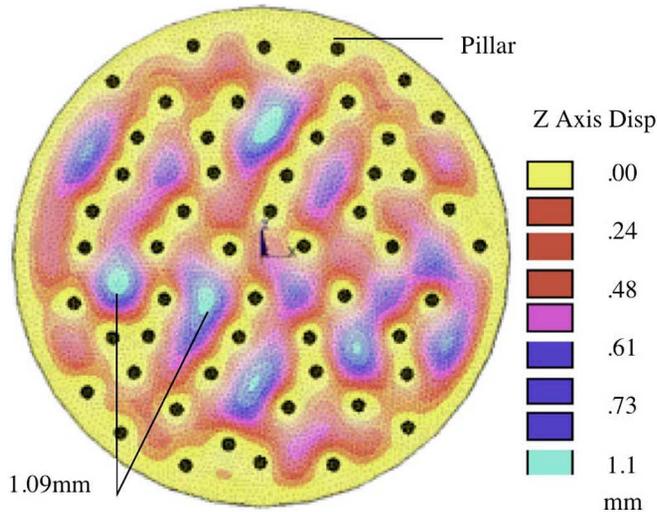


Fig. 7. The Von Mises strain distribution on the intermediate plate. The black dots indicate pillar positions.

they provide better noise performance as compared to the conventional small form-factor dc-dc converter modules. To conserve power, the SYSCON switches on power to sensor modules only during data acquisition.

IV. MECHANICAL DESIGN

ROMANIS consists of a cylindrical disc-shaped pressure casing of about 1.4 m in diameter and 100-mm thickness, constructed out of SS316L stainless steel. An intermediate plate inside the casing carries the electronics on its rear and sensor modules on the front. The interconnection between the data acquisition modules and the sensor modules is via a 32-pin connector [black rectangular connector on the sensor module in Fig. 4(a)]. Circular holes are provided on the intermediate plate to facilitate this connection for all the 254 modules. To take the mechanical load off the connectors, the sensor modules are provided with four magnetic “feet” which lock into four miniature magnets (with opposite polarity) embedded in the intermediate plate. The intermediate plate is a thin stainless steel plate of 3-mm thickness and rests on an O-ring embedded in a groove around a circular ridge in the casing. This provides slight cushioning and vibration isolation for the sensors mounted on the plate. The system was designed to operate at a water depth of 50 m. To ensure that the plate does not buckle under the pressure at depth, aluminum pillars are provided as support between the rear of the plate and inside back of the casing. The support pillars had to be distributed in such a way that they share the load evenly but do not obstruct the electronics. The positions of the pillars were optimized through finite element modeling and stress analysis. The final arrangement of the pillars and the resultant strain patterns are shown in Fig. 7. Any small perturbation at the back plate of the casing can result in an uneven distribution of stress along the pillars, leading to bending of the intermediate plate. Therefore, it is necessary to ensure that the back of the casing was machined to a flatness accuracy of about 1 mm. Because of the large diameter of the



Fig. 8. ROMANIS acoustic camera on its stand.

casing, special welding and machining processes were required to avoid warping.

ROMANIS uses a unique sensor-to-seawater acoustic interface. Conventional systems usually employ potted ceramic sensors or oil-filled neoprene boots to electrically isolate sensors from seawater while providing good acoustic transmission. To ease maintenance and reduce weight, ROMANIS uses a “dry-coupled” neoprene sheet to isolate the sensors from seawater. The neoprene sheet is seated over the sensors and held in place by pulling a vacuum. The vacuuming process removes any gas between the sensor surface and the neoprene, and hence provides good acoustic transmission. A partial backfilling of helium to 0.5 bar allows good heat dissipation. The neoprene sheet is sandwiched at the edges between the casing and a stainless steel retaining ring using clamps. The performance of dry coupling on sensitivity and beam pattern was measured using a smaller prototype array at the acoustic facilities of EDO Ceramics Corporation, New York, USA. Test results comparing the performance modules in the EDO test tank showed that sensors with dry-coupled neoprene perform as well (within 1 dB) as the ones with conventional polyurethane-based encapsulation. The beam pattern of the sensors was also confirmed to closely match theoretical predictions [23].

The array casing has three watertight connectors: two optical links for data transfer and one electrical link for power.

A leakage sensor is mounted inside the array and at its bottom. In case of any water ingress, the sensor sends a trigger signal to the SYSCON, which in turn sends a warning signal to the surface and initiates an automatic shutdown of the power to the electronics subsystems. A photograph of ROMANIS is shown in Fig. 8.

V. SOFTWARE DESIGN

A. Data Acquisition

A nested producer–consumer model is used for acquiring and transferring data to the surface. Eight producers (DAMs, also known as “clients”) and one consumer (surface node, also known as “server”) are involved in the transfer. Each producer generates data at the rate of about 200 Mb/s so that the consumer receives an overall data at the rate of 1.6 Gb/s. In addition to accepting sensor data from the embedded clients, the surface node also acts as a net-boot server for each client. This allows the clients to obtain a Tiny Core Linux OS image from the server at boot time, thus eliminating the need for an onboard flash or disk for booting. Each of the eight clients receives an IP address from a dynamic host configuration protocol (DHCP) server (running on the surface node) at bootup. The clients are controlled from the server using a secure shell (ssh). The server application software waits for and handles client connections.

1) *Server Architecture:* The server's main job is to read data from the clients. It checks for the integrity of data before writing them on to the storage disks. A TCP/IP-based protocol is used for reliable data transfer. Since the server communicates with eight clients concurrently, two quad-core processors with Intel hyperthreading technology are required to run the server. This allows up to 16 threads to run concurrently. The data acquisition process is initiated by the server once it is connected to eight DAM clients and the ARM-based system controller (SYSCON), as shown in Fig. 9.

Once the acquisitions begin, each process creates two threads: a read thread (producer) and a write thread (consumer). The producer then reads the data from the network interface controller (NIC) to a linked list of buffers in the main memory [random access memory (RAM)]. The consumer reads the data from these buffers and writes the data to files on the secondary disk storage.

2) *Client Architecture:* On startup, each client software establishes a connection to the server. On request from the server, each client starts data acquisition to transfer data from its 64-channel ADC card to the server. When the SYSCON enables the sampling clock, all the ADC cards start acquiring data and filling their respective FIFOs. Each client's job is to read data from its ADC card FIFO to its main memory (RAM) and write it to the NIC for transmission over Ethernet (using TCP/IP) to the server. The client software architecture and flow diagram are shown in Fig. 10.

B. Beamformer

Once the data are acquired, beamforming is necessary to identify acoustic energy from different directions. With a letterbox FoV of $17^\circ \times 8.5^\circ$, we require 24×12 beams to cover the FoV at the highest resolution of $0.7^\circ \times 0.7^\circ$. With 508 sensors producing

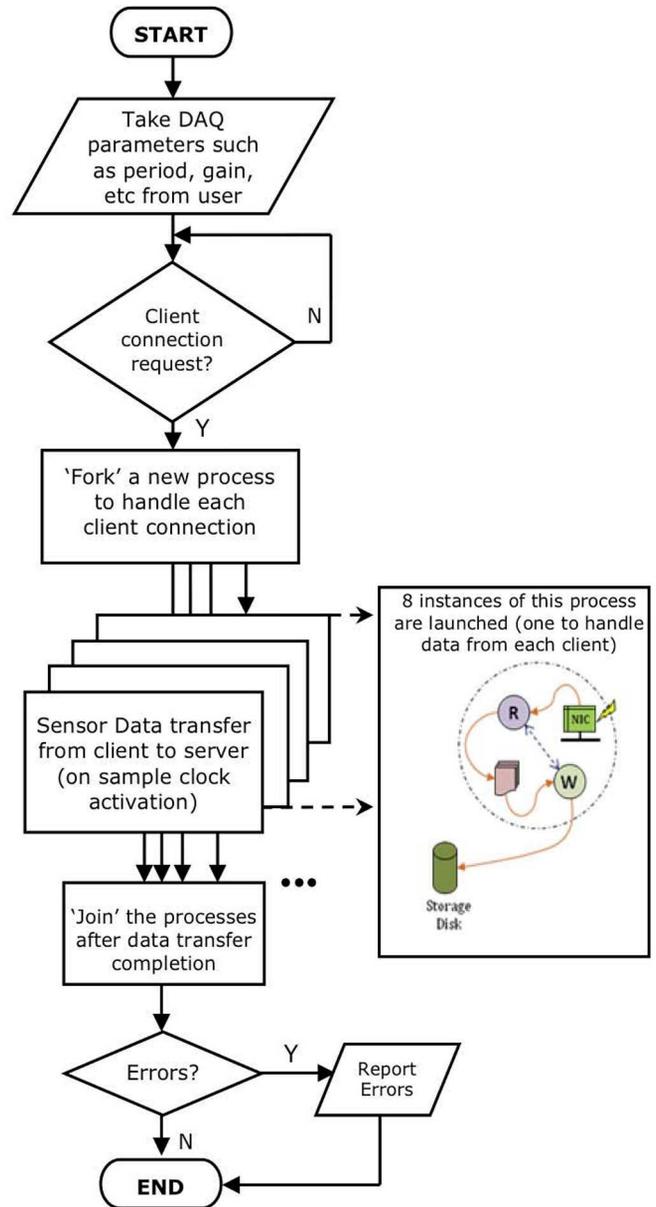


Fig. 9. Server software flow diagram and architecture (inset) [12].

200 kSa/s of data, the computational power required for beamforming is large. While it is possible to optimize the processing chain [20], the computational load is still $\mathcal{O}(10^2)$ Gflops.

The deployment and recovery of ROMANIS are resource-intensive operations, and hence *in situ* analysis of the data is invaluable for data quality assurance and control. This approach allows corrective measures, if needed, during the experiment rather than having to discover anomalies during post-processing. A 2-s data set from ROMANIS takes more than 5 min to process on a standalone server with dual quad-core Xeon (Nehalem) processors, while it takes about 20 s to process on a 36 CPU LS-20 IBM Blade cluster. Although there is a significant improvement in the processing time by using the PC cluster, it is not practical to carry the system to the field for *in situ* analysis of the data due to its large size and weight.

GPU processing has become popular in recent years for computationally demanding tasks. We implemented a frequency-domain beamformer using NVIDIA's compute unified

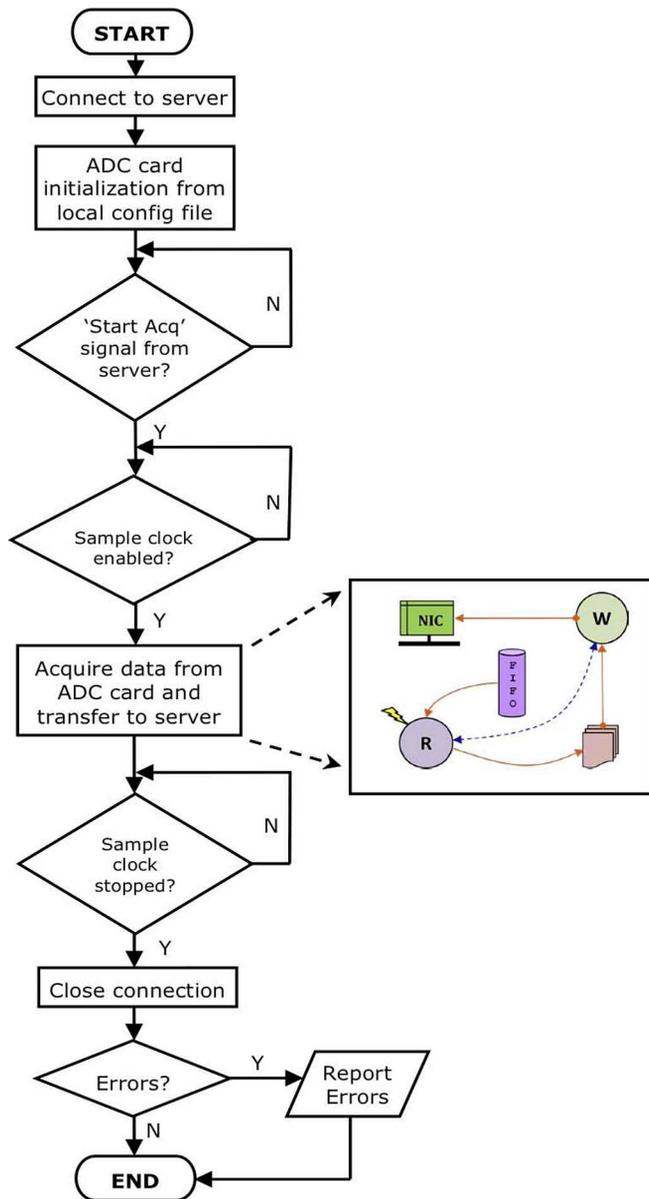


Fig. 10. Client software flow diagram and architecture (inset [12]).

device architecture (CUDA). CUDA is a heterogeneous computing model where the CPU and the GPU are used where they are the strongest. The cluster-based beamformer was implemented using C Language under Linux environment and used a message passing interface (MPI) algorithm for parallel processing. In the new GPU-based beamformer, the CPU handles the serial portions of the algorithm while the parallel computations are delegated to the GPU by executing a large number of threads in parallel. CUDA provides its own library, CUDA fast Fourier transform (CUFFT), for the computation of fast Fourier transform (FFT). The implementation of the beamformer on a Intel dual quad-core Xeon 5600 Super Server hosting a Tesla C1060 PCIe General Purpose GPU (GPGPU) provides real-time performance. A 2-s data set can now be beamformed in about 1 s, allowing substantial processing time for other image processing algorithms in a real-time streaming mode.

C. Imaging and Ranging

The beamformed data are processed using various algorithms to form images [10], [14] and to obtain range to imaged objects [14]. Imaging involves the computation of higher or fractional moment statistics, filtering using model-based filters (such as a Kalman filter) and combining information from various frequencies into acoustic-color images. The computational load from these algorithms is much lower than the beamforming, and hence these are currently implemented in MATLAB. The algorithms can be invoked in the field to perform *in situ* analysis of the recorded data.

VI. FIELD EXPERIMENTS

A. Location and Experimental Setup

The imaging performance of the ROMANIS system was evaluated through a series of field experiments. The experiments were conducted near Pulau Semakau, just south of Pulau Hantu in Singapore waters. The average water depth at the site was about 17 m with a reasonably flat bathymetry. The sea bottom was a mix of sand and mud. The location was surrounded by small islands and reef patches; these are good habitat for snapping shrimp. The closest island was about 400 m south of ROMANIS and directly behind the target frame. A cluster of islands were located to the north of ROMANIS at about 1-km distance with a shipping channel in between. There was also a mooring buoy in the FoV of ROMANIS; this may be a good habitat for snapping shrimp to colonize [24]. The experimental area is known to have significant population of snapping shrimp [25].

ROMANIS was deployed from a barge using an overhead crane and secured on to the sea bottom by divers. Both static and moving objects were used as targets. The main static target used was built using five sections of 1 m × 1 m × 6 mm closed cell neoprene sheets pasted on 2-mm thick aluminum plates and secured on to a 3-m × 3-m stainless steel frame, as shown in Fig. 12. The closed cell neoprene sheets served as good reflectors of acoustic energy due to the air bubbles trapped inside them. Another static target imaged was a 0.5-m diameter subsurface buoy. The mobile targets tested included scuba divers, who were helping with the target placement and alignment, and also a 2-m long, 200-mm diameter autonomous underwater (AUV) vehicle STARFISH [26]. Fig. 11 shows the experimental geometry.

B. Target Alignment

The main target frame was positioned at a range of about 65 m from ROMANIS with the help of a crane and divers. Once the setup was ready, ROMANIS was turned on for data acquisition and the ambient noise levels were measured. The gains of the signal conditioning sections were then adjusted to get the best dynamic range. The next step was to align ROMANIS to ensure that the object to be imaged is in its FoV. The initial placement of the object by the divers was, by necessity, approximate. To achieve a better alignment, an acoustic pinger was attached to the target frame and the pings were then recorded on ROMANIS. The pinger transmitted 100-ms acoustic pulses at a frequency of 37.5 kHz, once every second. Data from pairs of sensors along the horizontal and vertical axes of ROMANIS

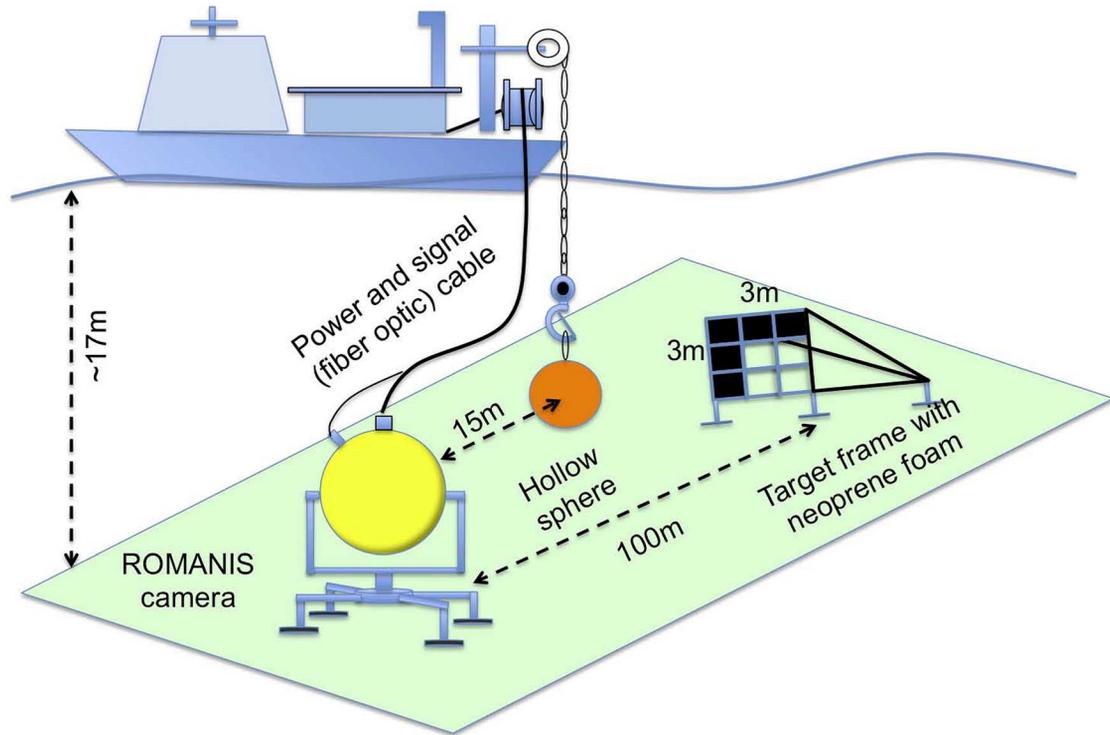


Fig. 11. Geometry of experimental setup.



Fig. 12. Static target formed by five 1-m \times 1-m neoprene panels on a 3-m \times 3-m steel frame. Lift bags used to deploy and retrieve the target can be seen hanging from the frame [14].

were used to compute the time difference of arrival, and azimuth/elevation angles to the pinger were estimated. Divers then rotated ROMANIS in the azimuthal direction to correct for the azimuthal bearing offset and locked it in position. Manual rotation of ROMANIS along the elevation was difficult to perform due to mechanical limitations, and hence the alignment in the elevation was achieved by adjusting the height of the target frame and ROMANIS before the deployment, and also using prior knowledge of the approximate bathymetry at the site. Once the target was aligned, the pinger was removed. These procedures were again adopted when the target was repositioned.

C. Data Processing

The primary sources of illumination in the area were snapping shrimp. Although the area also has significant shipping activity, the contribution from shipping noise was expected to be small since much of that noise would be at frequencies much lower than ROMANIS frequency band. Data recordings were obtained at various times of the day over several days, with some sessions lasting more than 30 min continuously. The main static target was stationed at ranges of 65 and 100 m from ROMANIS. The subsurface buoy was hung from the barge at a range of about 15 m, as illustrated in the experimental setup. The divers who were working on the positioning of target frame served as a moving targets for imaging purposes. The STARFISH AUV swimming through the ROMANIS FoV at a range of about 100 m served as another mobile target for imaging.

About 2 TB of data were collected over three weeks of experiments. The data processing chain included data quality checks followed by normalization, beamforming, and image processing. Though some of the data were processed onsite, most of the results presented here are from the postprocessing of data later in the laboratory. A rapid assessment of data quality at each sensor is made by generating a plot of energy received against element position, as shown in Fig. 13.

About 95% of the sensors provided good quality data during the experiment. A sample time series from one of the ROMANIS sensors is given in Fig. 14. These data confirmed that the experimental site was inhabited by snapping shrimp and there was sufficient energy in the processing band of ROMANIS for imaging purposes.

1) *Images of Static Targets:* After normalization of the data from the sensors, the beamformer yielded 288 (24 \times 12) beams. The energy received in each of the 288 beams was then mapped

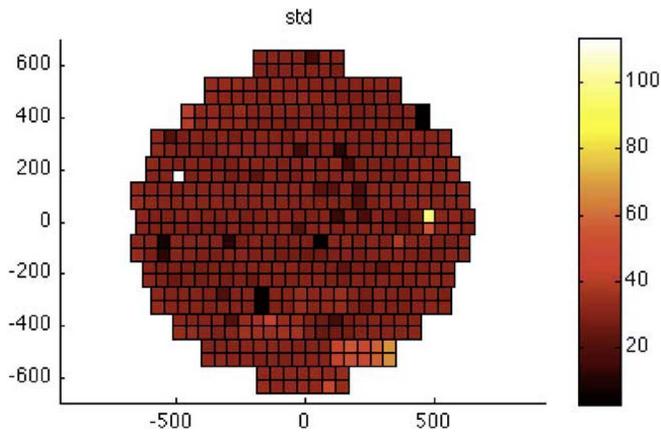


Fig. 13. Standard deviation plot showing the energy received by each sensor. The color bar indicates the levels on a linear scale as observed at the output of ADC and the white color indicates a failed sensor.

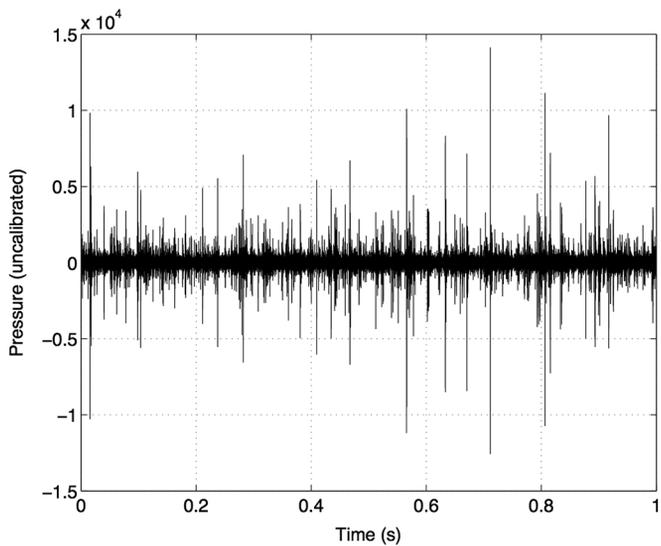


Fig. 14. Time series of the signal received on one of the sensors of ROMANIS. The impulses in the time series indicate snaps produced by snapping shrimp.

into pixel values, with the red color indicating a high intensity value compared to the blue ones over the band of frequencies mentioned. The color bars in these images was produced by mapping the highest pixel value to 1 and the lowest value to 0 on a linear scale. Other temporal and spatial statistics were used offline to produce better images. Fig. 15 shows the image of the pinger used to align ROMANIS and the target. It appeared as a bright spot near the center of the ROMANIS FoV, indicating that the target was well aligned. The two light patches at the top and bottom of the center image are believed to be caused by the grating lobes entering the edges of ROMANIS FoV.

The image of the submersible buoy, which was deployed using the crane at about 15-m range from ROMANIS, is shown in Fig. 16. The buoy was deployed using a weight at the bottom and then hung from a hook using a chain. The hook holding the buoy is also visible in the acoustic image.

The noise produced by snapping shrimp is impulsive, characteristic of cavitation bubble collapse. The resulting acoustic pressure field can be modeled by using symmetric α -stable ($S\alpha S$) distributions [27]. Imaging algorithms based on statistical measures such as fractional low order moments

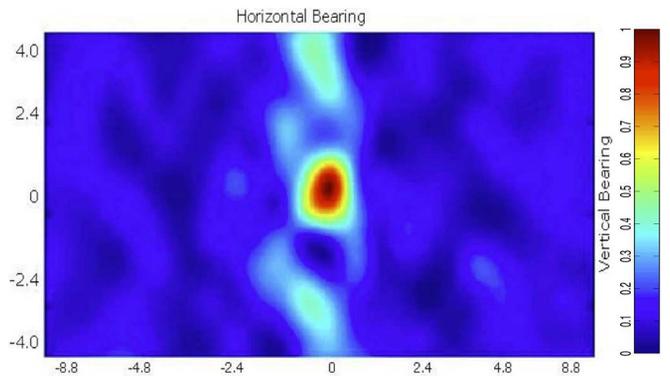


Fig. 15. Pinger (37.5 kHz) formed using data from ROMANIS. The two light patches on top and bottom are believed to be due to the effect of array grating lobes at the edges of FoV.

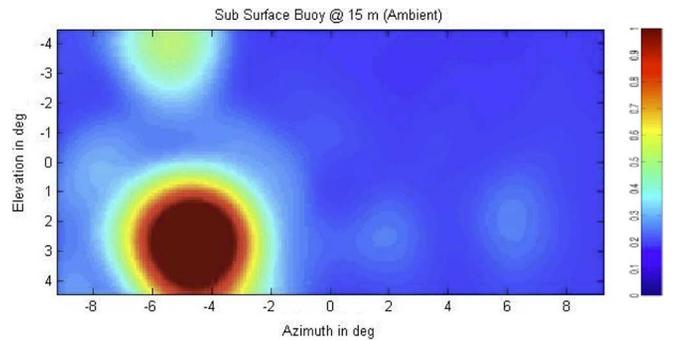


Fig. 16. Ambient noise image of the subsurface buoy at 15-m range from ROMANIS. The metal hook holding the buoy is also visible as a light colored patch near the top of the image. The frequency band covers the full range of ROMANIS, 25–85 kHz.

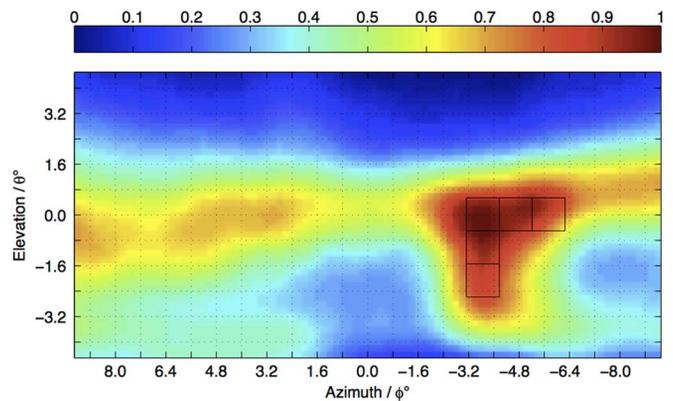


Fig. 17. Ambient noise image of the target formed using data (25–50 kHz) from ROMANIS using fractile imaging [14]. The estimated position and size of the target is superimposed on the image for reference.

(FLOM) and fractile estimators have been found to produce better images as compared to those based on second and higher order statistics [14]. Typically, fractile estimators produce more stable images as compared to FLOMs. One such image of the main target at about 65-m range is shown in Fig. 17. This image was produced from data in the 25–50-kHz band. At higher frequencies, due to rapid absorption, most of the ambient acoustic energy comes from nearby snaps. Since the number of snaps from the nearby area is low, we can expect higher statistical variability in the resulting pixel estimates and consequently poorer robustness in imaging. From the data, this was indeed found to be so. Conversely, lower absorption at low frequencies results in more energy contribution from snaps

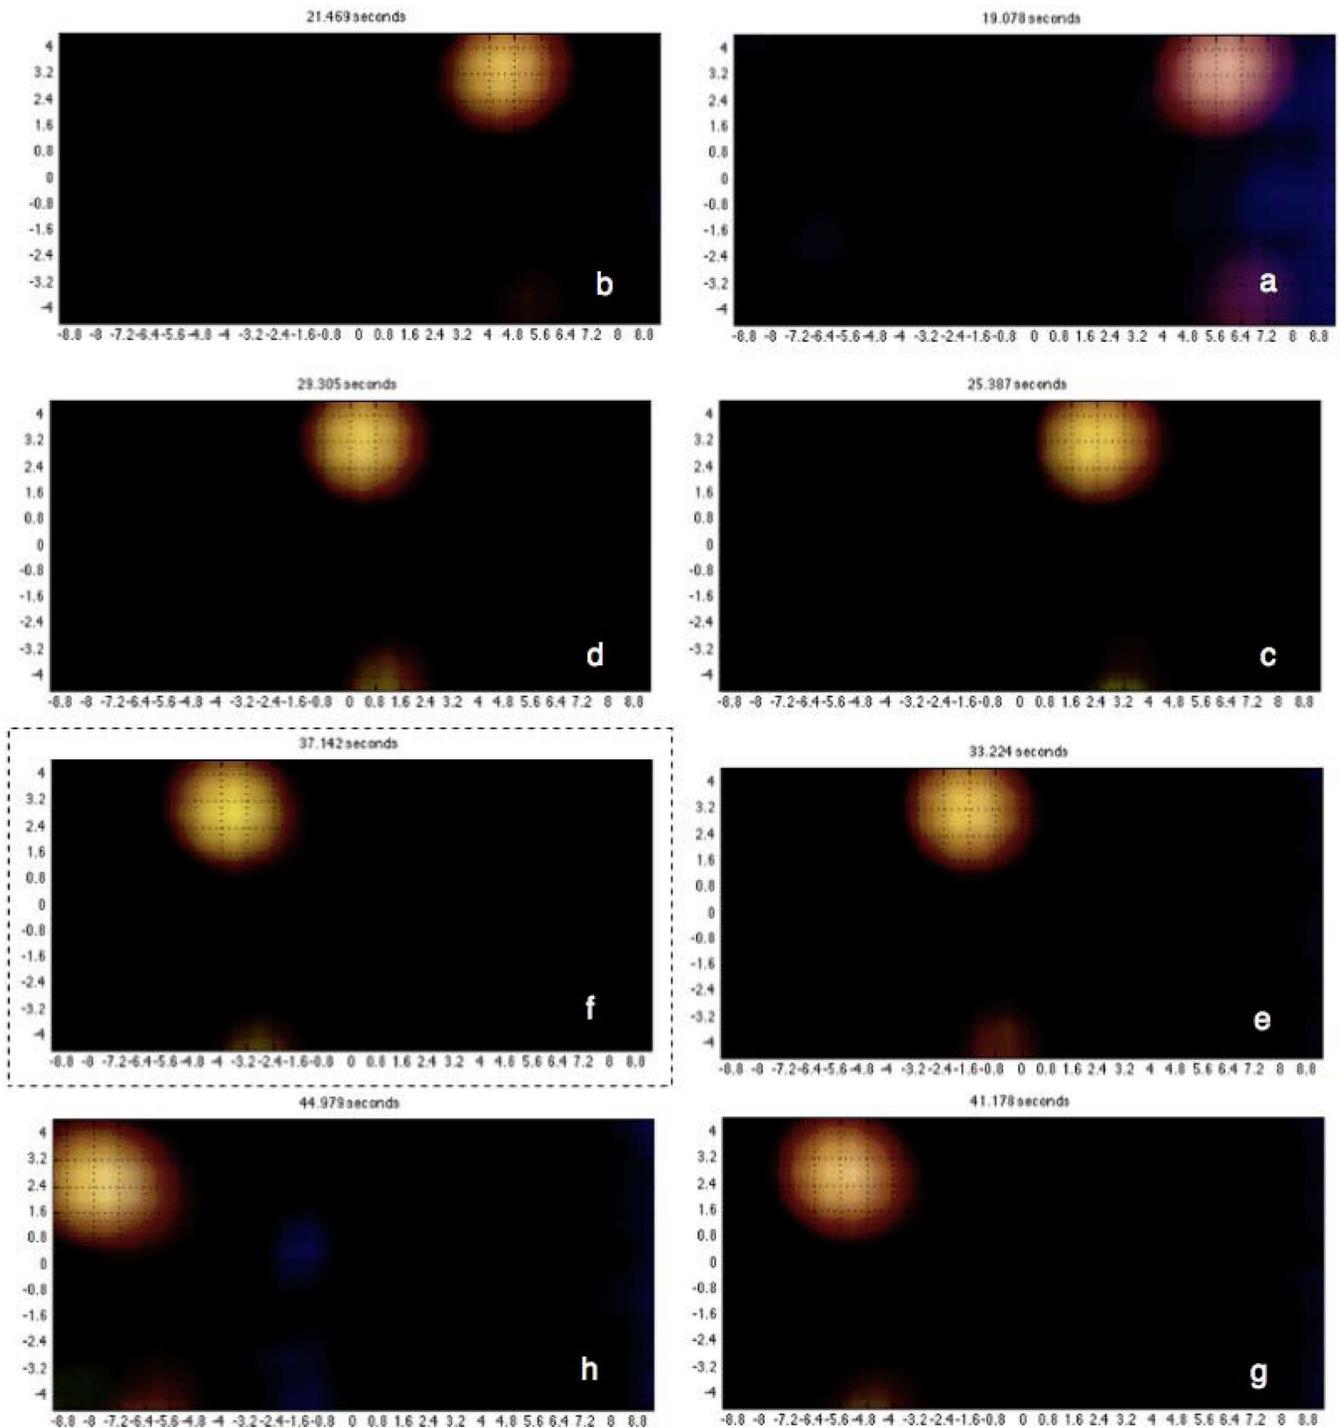


Fig. 18. Frames captured from an acoustic movie of generated using data from ROMANIS when the AUV was swimming across its FoV at about 100-m range. In addition to the ambient noise, the images seen here had contributions from a 37.5-kHz acoustic pinger on the AUV and also noise from its propulsion system. A weak reflection of the AUV in the seabed can be observed near the bottom of frames (b)–(g).

originating from farther ranges providing a stable illumination for ANI.

2) *Images of Moving Targets:* The STARFISH AUV was engaged in a separate field experiment and swam across the FoV of ROMANIS at a range of about 100 m. The data from corresponding recording were analyzed to create both static and moving images. A few frames from the movie showing the AUV motion are shown in Fig. 18. It may be noted that the AUV was also carrying a 37.5-kHz pinger for ground truth and hence the

energy from the pinger as well as the propulsion noise of the AUV would have contributed to these images. From the acoustic video generated it was also possible to estimate the speed of AUV as well as its altitude. The estimated speed of 2 kn matched expectations. The altitude of the AUV was computed from the reflection of the AUV image in the seabed and simple geometrical considerations. The estimated altitude of 6 m was verified using the depth information at the site (15 m) and the AUV dive depth (9 m). A static image of the AUV was also generated after

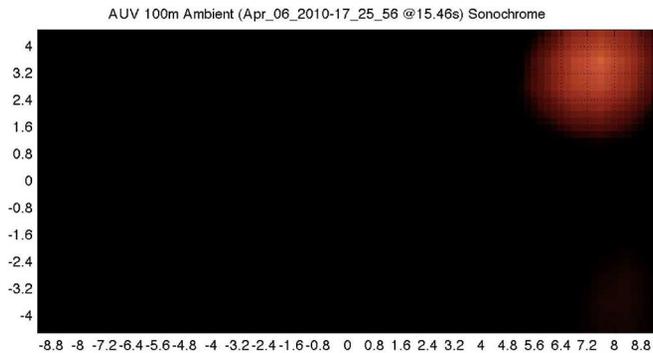


Fig. 19. Sonochrome acoustic image of the STARFISH AUV formed by ROMANIS at 100-m range. Low frequencies (20–40 kHz) are mapped to the red channel, medium frequencies (40–60 kHz) are mapped to the green channel, and high frequencies (60–80 kHz) are mapped to the blue channel of the image.

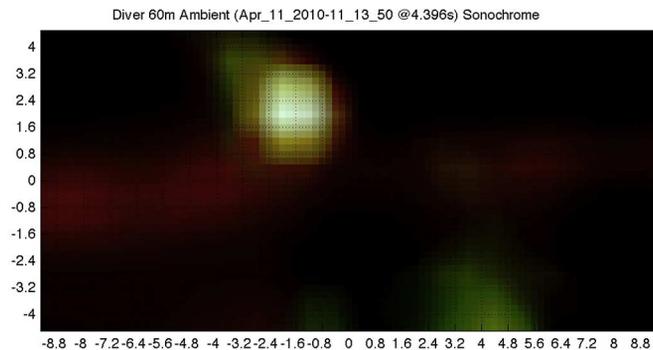


Fig. 20. Sonochrome acoustic image of a diver at 60-m range from ROMANIS. Low frequencies (20–40 kHz) are mapped to the red channel, medium frequencies (40–60 kHz) are mapped to the green channel, and high frequencies (60–80 kHz) are mapped to the blue channel of the image.

filtering out the 37.5-kHz pinger contribution, and this image is shown in Fig. 19. The sonochrome image was generated through RGB mapping and this type of image has both intensity and frequency information incorporated into it. In this case, the low frequencies (20–40 kHz) are mapped to red channel, medium frequencies (40–60 kHz) to the green channel, and high frequencies (60–80 kHz) to the blue channel.

When open-circuit divers were deploying the target frame, we seized the opportunity to use ROMANIS to image them. Bubbles released during exhalation by these divers act as strong sources and reflectors of sound. Fig. 20 shows an acoustic color image of a diver at a range of about 60 m from ROMANIS. The image may have contributions from ambient noise illumination as well as breathing noise from the scuba apparatus used by the diver. The data collected from this experiment were recently processed using generalized cross-correlation techniques, and the presence of divers was confirmed from the correlation spectrum associated with their breathing pattern [28].

3) *Passive Multistatic Ranging*: In addition to producing ambient noise images of underwater objects, we have also demonstrated the feasibility of estimating of their ranges passively using ROMANIS [14]. The method relies on first identifying the source (or snap) locations on the seabed and then applying bistatic sonar processing techniques with shrimp as the deterministic source and ROMANIS as the receiver. The challenges associated with this technique are in identifying and associating an echo with its snap in a multiple snaps scenario. It was initially

shown that this was achievable by carefully picking snaps and echoes [29]. Later, this process was automated [14]. Applying the algorithm to the ROMANIS data yielded a target range estimate of 67-m when the target was deployed at about 65-m according to GPS.

VII. CONCLUSION AND FUTURE WORK

The design and development of ROMANIS, a compact 2-D digital acoustic array, is presented. The system has been developed primarily for ANI applications in shallow waters and its performance was evaluated through a series of experiments in Singapore waters. ROMANIS, along with novel signal and image processing algorithms, was able to produce high-resolution images of both static and moving underwater objects using ambient noise as the source of illumination. The range and resolution performances of ROMANIS were found to be higher than any previously demonstrated ANI system. The implementation of a GPU beamformer and image processing algorithms has made it possible to generate images in real-time. Not only were we able to form images of targets, but also estimate the range of some of the targets from the camera using ANI. Although primarily built for ANI application, ROMANIS is a broadband 2-D array that may find use in other underwater applications. For example, ROMANIS may be able to generate acoustic videos of bubble evolution in ship wakes or during wave-breaking events. It could also be a great spatial diversity receiver for underwater communication applications. This may allow high-speed communication to AUVs at farther ranges than otherwise possible.

ACKNOWLEDGMENT

The authors would like to thank M. Panayamadam, S. P. Tan, and N. Chandavarkar for the wonderful support provided during the development and deployment of ROMANIS.

REFERENCES

- [1] M. Buckingham, B. Berkhout, and S. Glegg, "Imaging the ocean with ambient noise," *Nature*, vol. 356, pp. 323–329, 1992.
- [2] M. Buckingham, "Theory of acoustic imaging in the ocean with ambient noise," *J. Comput. Acoust.*, vol. 1, pp. 117–140, 1993.
- [3] J. Potter, "Acoustic imaging using ambient noise: Some theory and simulation results," *J. Acoust. Soc. Amer.*, vol. 95, no. 1, pp. 21–33, 1994.
- [4] J. Potter, M. Buckingham, C. Deane, and C. EpifanioN. Carbone, "Acoustic daylight: Preliminary results from an ambient noise imaging system," *J. Acoust. Soc. Amer.*, vol. 96, no. 5, pt. 2, p. 3235, 1994.
- [5] M. Buckingham, J. Potter, and C. Epifanio, "Seeing underwater with background noise," *Sci. Amer.*, vol. 274, no. 2, pp. 40–44, 1996.
- [6] M. Olivieri, S. Glegg, and R. K. Coulson, "Imaging of passive targets in shallow water using ambient noise acoustic imaging system," *J. Acoust. Soc. Amer.*, vol. 99, no. 4, pt. 2, p. 2453, 1996.
- [7] C. Epifanio, J. Potter, G. Deane, M. Readhead, and M. Buckingham, "Imaging in the ocean with ambient noise: The ORB experiments," *J. Acoust. Soc. Amer.*, vol. 106, no. 6, pp. 3211–3225, 1999.
- [8] M. L. Readhead, "Acoustic daylight—Using ambient noise to see underwater," *J. Acoust. Australia*, vol. 29, pp. 63–68, 2001.
- [9] K. Mori, H. Ogasawara, T. Nakamura, T. Tsuchiya, and N. Endoh, "Extraction of target scatterings from received transients on target detection trial of ambient noise imaging with acoustic lens," *Jpn. J. Appl. Phys.*, vol. 51, pp. 07GG10-1–07GG10-7, 2012.
- [10] J. Potter and M. Chitre, "Ambient noise imaging in warm shallow seas: Second order moment and model based imaging algorithms," *J. Acoust. Soc. Amer.*, vol. 106, no. 6, pp. 3201–3210, 1999.

- [11] V. Pallayil *et al.*, "First deployment of ROMANIS and preliminary data analysis," in *Proc. OCEANS Conf.*, 2003, pp. 882–888.
- [12] S. Kuselan, A. Raichur, and V. Pallayil, "Design and development of a giga-bit ethernet based high-speed broadband data acquisition system for an underwater imaging array," in *Proc. OCEANS Conf.*, 2010, pp. 1–7.
- [13] V. Pallayil *et al.*, "Ambient noise imaging: Experiments with ROMANIS," in *Proc. Int. Conf. Theor. Comput. Acoust.*, 2011, p. 40.
- [14] M. Chitre, S. Kuselan, and V. Pallayil, "Ambient noise imaging in warm shallow waters; robust statistical algorithm and range estimation," *J. Acoust. Soc. Amer.*, vol. 132, no. 2, pp. 838–847, 2012.
- [15] V. Albers, *Underwater Acoustics Handbook-II*. University Park, PA, USA: Pennsylvania State Univ. Press, 1965.
- [16] D. Cato, "The biological contribution to the ambient noise in waters near Australia," *J. Acoust. Australia*, vol. 20, pp. 76–80, 1993.
- [17] J. Potter, L. Wei, and M. Chitre, "High frequency ambient noise in warm shallow waters," in *Sea Surface Sound 1997*, T. Leighton, Ed. Norwell, MA, USA: Kluwer, 1997, pp. 45–54.
- [18] W. L. Au and R. Penner, "Target detection in noise by echolocating Atlantic bottlenose dolphins," *Aus. J. Phys.*, vol. 20, pp. 687–693, 1981.
- [19] B. Mills and A. Little, "A high resolution aerial system of a new type," *J. Acoust. Australia*, vol. 6, pp. 272–278, 1953.
- [20] M. Chitre and J. Potter, "Optimisation and beam forming of a two dimensional sparse array," *High Performance Comput.*, vol. 1, pp. 452–459, 1998.
- [21] H. Zhang, T. Koay, V. Pallayil, Y. Zhang, and J. Potter, "Fibre channel storage area network design for an acoustic camera with 1.6 gbits/s bandwidth," in *Proc. IEEE Region 10 Int. Conf. Electr. Electron. Technol.*, D. Tien and Y. C. Lian, Eds., 2001, vol. 1, pp. 143–148.
- [22] General Standards Corporation, "User manuals," [Online]. Available: http://www.generalstandards.com/user-manuals/pmc66_16ai64ssa_c_man_092807.pdf
- [23] V. Pallayil, P. Deshpande, S. Badiu, C. Constantin, and J. Potter, "A 1.6 gigabits/s, 25–85 kHz acoustic imaging array—Novel mechanical and electronic design aspects," in *Proc. OCEANS Conf.*, 1999, vol. 1, pp. 352–358.
- [24] B. Ferguson and J. Cleary, "In situ source level and source position estimates of biological transient signals produced by snapping shrimp in an underwater environment," *J. Acoust. Soc. Amer.*, vol. 109, no. 6, pp. 3031–3037, 2001.
- [25] T. Koay, E. T. Tan, M. Chitre, and J. Potter, "Estimating the spatial and temporal distribution of snapping shrimp using a portable broadband 3-dimensional acoustic array," in *Proc. MTS/IEEE OCEANS Conf.*, 2003, vol. 5, pp. 2706–2713.
- [26] T.-B. Koay *et al.*, "Starfish—A small team of autonomous robotic fish," *Indian J. Geo-Mar. Sci.*, vol. 40, pp. 157–167, 2011.
- [27] M. A. Chitre, J. R. Potter, and S. H. Ong, "Optimal and near-optimal signal detection in snapping shrimp dominated ambient noise," *IEEE J. Ocean. Eng.*, vol. 31, no. 2, pp. 497–503, Apr. 2006.
- [28] L. Fillinger, M. Chitre, and P. Venugopalan, "Passive diver detection in warm waters," in *Proc. Underwater Acoust. Int. Conf. Exhibit.*, J. S. Papadakis and L. Bjorno, Eds., 2013, vol. 5, pp. 1459–1464.
- [29] S. Kuselan, M. Chitre, and V. Pallayil, "Ambient noise imaging through joint source localisation," in *Proc. MTS/IEEE OCEANS Conf.*, 2011, pp. 1–6.



Venugopalan Pallayil (S'90–M'99–SM'04) received the M.Sc. degree in physics and the Ph.D. degree in microwave electronics from Cochin University of Science and Technology, Kerala, India, in 1981 and 1992, respectively.

From 1986 to 1998, he served as an R&D Scientist in the Defence Research and Development Organisation, India, contributing to the development of airborne ASW sensors. In late 1998, he joined the Acoustic Research Laboratory (ARL), Tropical Marine Science Institute (TMSI), National University of

Singapore, Singapore, as a Research Fellow. Currently, he is a Senior Research Fellow and Deputy Head at ARL. He has also served as Manager (Operations) for TMSI from 2007 to 2014 overseeing the finance and facility operations for the Institute. He has a strong interest in the development of underwater sensor systems and applications and has led many related projects in ARL. He was one of the leading members of the team which won the Defence Technology Prize 2004 for the project ROMANIS on ambient noise imaging.

Dr. Pallayil has served the local Oceanic Engineering Society (OES) Chapter in many capacities, including its Chairmanship. He was the Chair for Finance for the 2006 IEEE OCEANS Conference in Singapore and also the Co-Chair and Chair of the Singapore AUV Challenge, an AUV competition for tertiary students, for 2013 and 2014, respectively. He is a member of the Acoustical Society of America and the Society of Acoustics, Singapore. Currently, he is a member of the organizing committee for the Western Acoustic Pacific Conference (WESPAC) to be held in December 2015 in Singapore.



Mandar Chitre received the B.Eng. and M.Eng. degrees in electrical engineering from the National University of Singapore (NUS), Singapore, in 1997 and 2000, respectively. He also received the M.Sc. degree in bioinformatics from the Nanyang Technological University (NTU), Singapore, in 2004, and the Ph.D. degree from NUS in 2006.

From 1997 to 1998, he worked with the ARL, NUS in Singapore. From 1998 to 2002, he headed the Technology Division of a Regional Telecommunications Solutions Company. In 2003, he rejoined

ARL, initially as the Deputy Head (Research) and is now the Head of the Laboratory. He also holds a joint appointment with the Department of Electrical and Computer Engineering at NUS as an Assistant Professor. His current research interests include underwater communications, autonomous underwater vehicles, and acoustic signal processing.

Dr. Chitre has served on the Technical Program Committees of the IEEE OCEANS, the International Conference on Underwater Wireless Networks and Systems (WUWNet), the Defence Technology Asia (DTA) International Conference, and the Offshore Technology Conference (OTC). He has served as a Reviewer for numerous international journals. He was the chairman of the student poster committee for IEEE OCEANS'06 in Singapore, and the chairman for the IEEE Singapore AUV Challenge 2013. He is currently the IEEE Ocean Engineering Society Technology Committee Cochair of underwater communication, navigation and positioning, and an Associate Editor for the IEEE JOURNAL OF OCEANIC ENGINEERING.



Subash Kuselan received the B.Eng. degree in instrumentation and control engineering from the National Institute of Technology, Trichy, India, in 2003, and the M.Eng. degree from the National University of Singapore (NUS), Singapore, in 2012, for his research in the field of underwater ambient noise imaging in warm shallow waters.

He has worked as an R&D Scientist at the Naval Physical and Oceanographic Laboratory under the Defence Research & Development Organisation, India, from 2003 to 2007, making major contribu-

tions to the development of active towed array sonar system. From 2008 to 2012, he was a research staff member at the Acoustic Research Laboratory (ARL), NUS, and contributed to many projects by designing and implementing embedded system hardware and software solutions. His interests include underwater acoustics system development, signal processing, and ambient noise imaging.



Amogh Raichur received the B.E. degree in electronics engineering from the University of Mumbai, Mumbai, India, in 2007 and the M.Sc. degree in smart product design from the Nanyang Technical University, Singapore, in 2009, respectively.

He is currently a Senior Engineer working with test automation in the embedded systems division of Finisar Inc., an optical telecommunications company. He has previously worked as a Research Engineer with the Acoustic Research Laboratory (ARL), Tropical Marine Science Institute (TMSI),

National University of Singapore, Singapore, where the focus of his work was embedded system design for data acquisition and sensing applications, in the marine acoustics domain.



Manu Ignatius (M'14) received the B.S. degree in electronics and communication engineering from Mahatma Gandhi University College of Engineering, Kerala, India, and the M.S. degree in electrical and computer engineering from the National University of Singapore (NUS), Singapore, in 2004 and 2010, respectively.

Currently, he is a member of the core technology team at Subnero, a Singapore-based startup, which focuses on developing underwater communication, navigation, monitoring, and sensing technology and solutions. Before joining Subnero, he worked at several multinational companies in the field of embedded systems and computer networks.



John R. Potter (M'96–SM'98) received a joint Honors degree in mathematics and physics from the University of Bristol, Bristol, U.K., in 1979, and the Ph.D. degree in glaciology and oceanography from Cambridge University, Cambridge, U.K., in 1985, studying the Antarctic ice mass balance, where he spent four summers.

He worked for a while in Italy on acoustic propagation fluctuations, then sailed across the Atlantic to San Diego, CA, USA, to work at Scripps Institution of Oceanography on ambient noise imaging and

marine mammal acoustics. In 1995, he sailed to Singapore where he founded the Acoustic Research Laboratory (ARL) as part of the Tropical Marine Science Institute (TMSI). He was Head of the ARL and Associate Director of the TMSI for 12 years, working on passive imaging, marine mammals, distributed autonomous intelligent sensing, underwater communications networking, and cooperative behavior. In 2004–2005, he sailed with his family around the Indian Ocean on a sponsored voyage of research, public outreach, and education to draw attention to the urgent need to become better custodians of our fragile planet. Returning to Italy in 2007, he led a team to develop underwater communications and networking, leading to an underwater internet for robots, and is now working on strategic development for the NATO Science & Technology Organization (STO) Centre for Maritime Research & Experimentation (CMRE), La Spezia, Italy.

Dr. Potter is an Associate Editor of the IEEE JOURNAL OF OCEANIC ENGINEERING, PADI Master Scuba Diver Trainer, and an International Fellow of the Explorer's Club.