Acoustic Communications for Deep-ocean Observatories: Results of Initial Testing at the MBARI MARS Node

Lee Freitag, Keenan Ball, Peter Koski, Sandipa Singh, and Eric Gallimore
Woods Hole Oceanographic Institution
Woods Hole, MA 02543

***Abstract*— A prototype system for connecting remote instruments to a cabled observatory is described. The system consists of base-station modems and a data acquisition system attached to the MBARI MARS observatory in 900 m water off Monterey, California, plus remote modems used for link testing. The objective of the system is to provide a drop-in capability to connect sensors on the sea-floor or on sub-sea moorings to a cabled node that would otherwise require an ROV to make a hard-wired connection. While ROV installations of sensor-to-node cables are practical for many instruments, the acoustic connection is a way of lowering the barrier to general access of cabled nodes.**

The system designed for demonstration purposes is a combination of commercial components and custom hardware and software. While the system is a prototype, considerable effort was expended to make the sub-sea portion as simple and reliable as possible so as to minimize the potential for failure and the reduce the associated costs of recovering and re-deploying the system. While the system uses real-time modems (the WHOI Micro-Modem) in the sub-sea portion, a data acquisition system and streaming feed of the acoustic array allows the use of one or more software modems in the future, which can be vendor independent or conform to an established standard.

The system was deployed in February 2010 and has been undergoing testing since. This paper reports on initial results obtained to-date. Several types of tests have been performed to establish performance over both short (24 hours) and medium (two week) time scales. At 2 km burst data rates of 5000 bps have been achieved with near-perfect reliability using a carrier frequency of 25 kHz and 5 kHz of bandwidth. Temporal variability is low, and the largest impacts on performance are the delayed surface bounce from the surface and acoustic interference observed during occasional ROV operations in the area.

I. Introduction

The installation of cabled, deep-water research nodes at sites around the world creates new opportunities for access to areas of the ocean that typically cannot be reached on a continuous basis. Cabled observatories in North America include the VENUS system now installed in Canada [1], and the US NSF Regional Scale Nodes, called NEPTUNE during its design phase [2] and now moving forward to a construction phase. The capability of the cabled observatories depends upon their size and construction, but for these two large systems, many gigabits per second of data throughput is available and hundreds of Watts of power is allowed for sensors on the network. The long duration (25 years) of the observations will provide data sets important for many ocean science disciplines, and it is likely that the network locations will attract extra attention because of the wealth of data that will be available.

While the principal means of connecting sensors to sea-floor observatories is by use of an ROV to install a cable between the sea-floor node and sensor, methods that avoid the expense of the ROV include optical and acoustic modems. While emerging technology for optical modems can provide links in clear-water at ranges of 100-200 m [3], acoustics is power-efficient for several kilometers or more, depending on the geometry and height of sources and receivers, and it is also cost-effective for a number of applications.

Acoustics are appropriate for moderate bandwidth, moderate power applications spanning periods of 1-2 years or less. Many modern instruments use very little power, and most sensor computers now have sophisticated control options that allow real-time data to be tailored to specific applications while storing all data to non-volatile memory for safe-keeping. Thus while new high-power, high-throughput sensors have been proposed (e.g. HD video), most standard oceanographic sensors do not require high-speed Ethernet or very much power, and they motivate the proposed approach.

Previous work in acoustically-connected observatories has focused on sea-floor to buoy links. These links are very useful when a surface buoy is available, but they differ from the system described here because they are potentially affected by weather and buoy watch circle, both of which can impact the instantaneous throughput. For the sea-floor sensor network the platform does not move, and in this results in a more stable link environment that can be optimized over the long term. An objective of this project is to explore the stability of the link over multiple time scales.

Examples of buoy-based, acoustically-linked observatories include [4], and it should also be noted that the NOAA tsunami-warning systems are examples of deep-sea acoustically-linked observatories, but with low throughput requirements.

The design concept described here was presented in [5],

978-1-4244-5222-4/10/$26.00 ©2010 IEEE

where the rationale and basic approach were outlined. That paper included examples of system requirements and lists of potential sensors and estimated daily throughput. Other work included the ASSEM project, a European effort in geological and hazard monitoring that used both cables and acoustic modems to telemeter real-time data from the sea-floor to shore [6].

II. System Design

The system is a hybrid data collection and acoustic modem system as shown in Figs. 1 and 2. Components of the system designed and built for this project include three sub-systems, a shore-based computer, the dry section of the ACOMMS node, and the wet section as described in detail below.

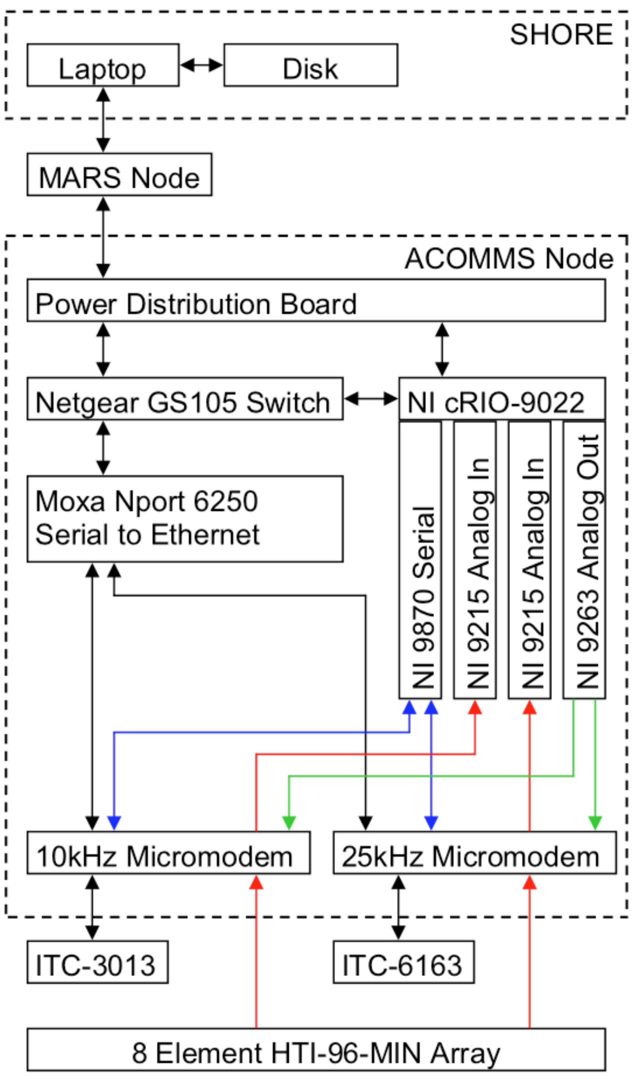

Fig. 1. Block diagram of the acoustic communications base station.

A. ACOMMS Node, Dry Section

The dry section of the node consists of two main sections, the two Micro-Modems and a National Instruments data acquisition system. The two modems are standard WHOI units, in this case they are configured to share the receiver array, but each drives their own source. The 10 kHz modem uses a hemispherical beam pattern transducer, while the 25 kHz modem has a toroidal beam pattern transducer with a 30 degree vertical beam. Both are oil-filled to allow operation in the deep-ocean. The modem serial ports are interfaced to the MARS Ethernet via a converter, and the modem console connections are accessed from shore.

Also in the dry section is the data acquisition system. It is a compact cRIO package that fits inside the 24 by 6.75 inch ID pressure vessel. The cRIO has the capability to act as a completely stand-alone system, or it can be run using National Instruments LabView from a host controller. The system can stream 8 channels at 80 kHz per channel in parallel with the acoustic modems operating over their serial ports on the network. While the data acquisition can run faster, at present the networks as configured cannot support higher sample rates, and this is being investigated.

B. ACOMMS Node, Shore Side

The shore system includes a laptop which is within the MBARI-MARS firewall. Security for the subsea network is handled carefully by MARS operators at MBARI, and the computer at MBARI lies within a secure area with physical access to the MARS network. Control of the system is done directly from the laptop, which is in turn accessed using remote-control software allowed from specific computers at WHOI. The serial ports on the modems are mapped to UDP packets, and custom testing software allows commands to be generated for the modems to initiate test transmissions.

C. ACOMMS Node, Wet Section

The wet section of the system includes the receiver array and cable plus a cage with the two sources. The array has one element right at the lower assembly, and the elements are spaced 1 meter apart. The total height of the assembly with the cage is about 9.5 meters. While from an acoustic propagation perspective a higher source and receiver system would be preferred, the main MARS node is only about 70 m away, and safe operation of the ROV is affected by the presence of structures in the area. Thus the ROV pilots requested that the height be kept to approximately 10 percent of the distance between the two. The size and shape of the system was dictated in part by the need to be able to deploy and recover the unit with the MBARI ROV *Ventana*. The two tubes on either side are made to allow the forks of the ROV to pick up the system.

D. Base Station at MARS

The MARS node is located on Smooth Ridge, in the center of Monterey Bay (Fig. 3). The 52 km power and data cable terminates at MBARI and the system is operated from a dedicated control room where an engineer can turn off and on power to individual ports and monitor current draw and other system parameters. Multiple experiments are on-going at the MARS site, with the acoustic modem base station being one of two that use active acoustics, the other being an upward looking sonar for fisheries work.

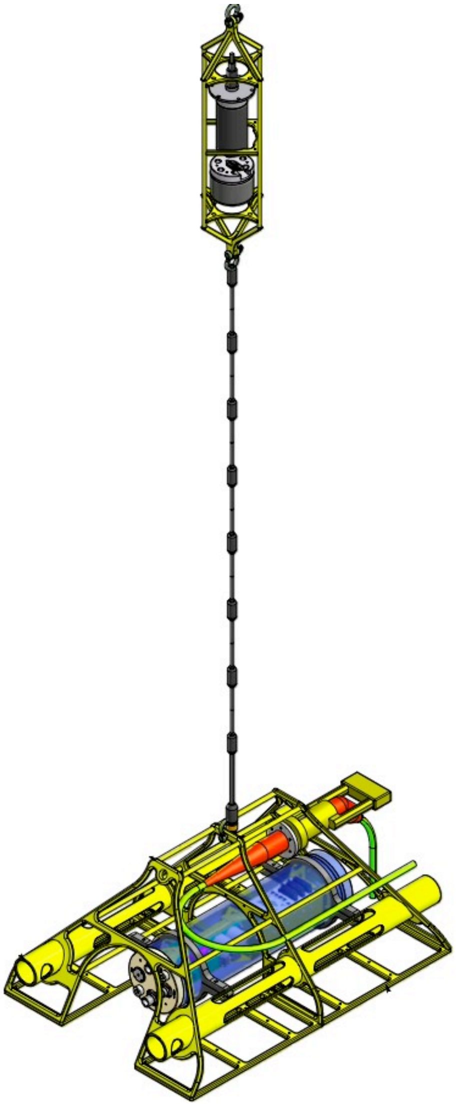

Fig. 2. Modem node assembly drawing showing the ROV-deployable cage with transducer cage and array above. A glass ball is used above the transducer cage for flotation.

E. Installation

The MARS node, shown in Fig. 4, is a large, trawl-proof unit that is the termination point for the cable from shore and it hosts the sub-sea sensors that are connected to it. There are a number of options for interfacing to the node, 100 Mbps Ethernet is necessary for the modem system due to the bandwidth of the acoustic array.

The installation by ROV *Ventana* was successfully accomplished February 17, 2010. Fig. 5 shows the ROV manipulator holding the wet-matable power-Ethernet connector and preparing to plug it into the node. Installation was accomplished by carrying the entire unit down to the node with the ROV, then plugging it in and then finally deploying the cage (with float) so that the array was suspended.

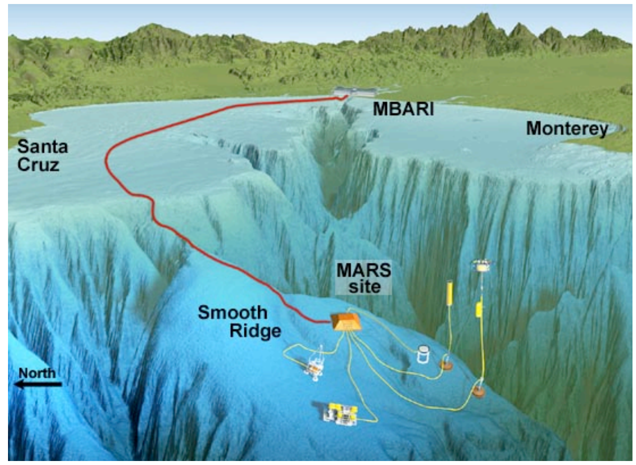

Fig. 3. MARS observatory in Monterey Bay. David Fierstein (c) 2005 MBARI.

Fig. 4. MARS node trawl-proof cage undergoing tank testing prior to deployment.

III. MODEM HARDWARE AND SIGNALING

The modems at the base station and those deployed at the two remote locations are standard WHOI Micro-Modems [7]. The remote 25 kHz systems transmit at 186 dB using 24 W of power, while the 10 kHz systems transmit at 184 dB using 10 W. The link efficiency is thus a function of the achievable data rate, which is described in more detail below.

A. Modem Coding, Modulation and Protocols

The modems configured for this experiment use 5 kHz of bandwidth at either 10 or 25 kHz carrier. The modulation is quadrature phase-shift keying (QPSK), which has a bandwidth efficiency of 2 bits per Hz, resulting in 10,000 bps raw burst data rate. To improve reliability several different types of spreading and coding are used, with low rates (300 and 600 bps) employing spreading codes, while for the highest two data rates a concatenated approach is used. The inner code for the latter is a BCH block code, while a Reed-Solomon (RS) code is used on the outer layer to correct occasional bad code words. The RS code adds 10% overhead, and the combination

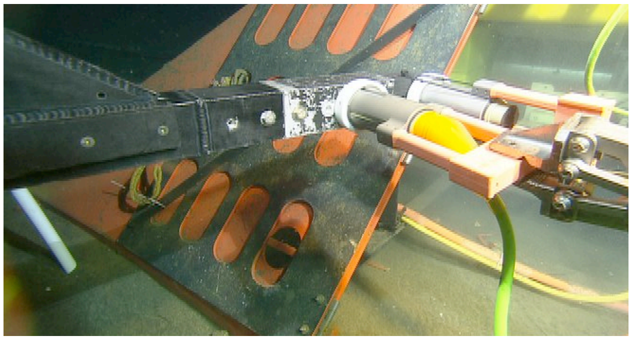

Fig. 5. MBARI ROV *Ventana* preparing to plug in the power and Ethernet cable from the modem node.

of the two codes plus training data results in the 10,000 bps raw rate being reduced to approximately 5300 bps.

Each packet is preceded by an FM sweep for initial acquisition, followed by a null time, training sequence and modulation selection header (transmitted at a lower rate). The modulation header contains bits that indicate the error-correction type that is applied, and the number of frames in the packet. The highest data rate packet using a (9,14) block code and contains 8 frames, each 256 bytes, for a maximum data payload of 2048 bytes. The packet length is approximately 3.5 s.

B. Receiver

A multi-channel adaptive equalizer is used as the receiver for the signals [8]. The receiver on the real-time system implements the adaptive equalizer using a Texas Instruments TMS320C6713 processor. The Decision Feedback Equalizer (DFE) is very similar to that described in previous papers, and is based on the original work of Stojanovic [8] [9]. The differences, some of which are related to the real-time nature of the implementation, include the use of the adaptive step-size LMS for the filter update algorithm [10], and a single phase-locked loop applied after the feed-forward channels are combined rather than n PLLs, one for each hydrophone channel. The adaptive step-size LMS is computationally simpler than the RLS algorithm, and more amenable than RLS to single-precision math (RLS is typically implemented in its square-root variant for single-precision). Adaptive LMS performance relative to RLS for the underwater channel has been described in [11].

C. Protocols and Throughput

Near real-time data (a few hours delay at most) is easiest to support with a time-multiplexed approach, and large data transfers (in this context up to a few hundred kilobytes) are more efficient than low-rate continuous operation. The preferred protocol for initial work is simple polling, where the central node operates as a master controller, sending a command to each unit requesting a packet of data at a time, then acknowledging each one if received without errors. More complex protocols are certainly possible, the most obvious being to request a much larger number of packets to be sent in a row, then acknowledging them all at once.

For the testing of the prototype system the overall link efficiency can be calculated from the packet success rate of the physical layer, and thus the simple, one-packet-at-a-time approach was used.

IV. Results to Date

Performance of a phase-coherent communications system is largely dependent upon: (a) the SNR at the receiver, (b) the amount of multipath and its length, and (c) the time rate of change of the multipath. For this system the SNR is determined by the source level and the receive noise, the former being constrained by the available power and the response of the transducer, and the latter set by the frequency band and any nearby interference.

A. Impulse Response Estimate (IRE)

The channel impulse response is measured by matched filtering for the FM sweep that precedes each packet. In this deep-water channel the FM sweep and packet are closer in time (0.2 s) than the delay between initial arrival and surface bounce (0.5 s), thus adding noise to the estimate. However, the gain of the FM sweep is sufficient to clearly show the surface bounce at 0.5 seconds, almost exactly what is calculated based on the source-receiver depths and range. The delay is much too long to be incorporated into the taps of the equalizer adaptive filters, thus the receiver depends on the minimum MSE solution that combines all four channels to help reduce the impact of the later arrival. It should also be noted that the surface bounce is not a single, clean arrival, but multiple arrivals spanning approximately 10 msec., and it is the sum of that delayed energy that affects performance.

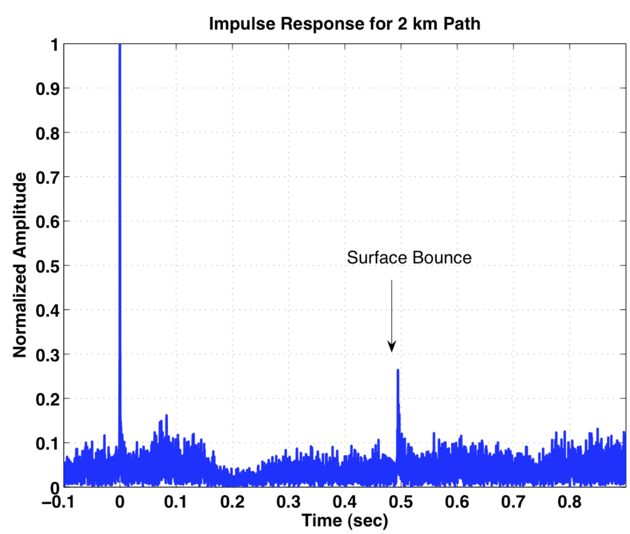

Fig. 6. The surface bounce for the 2 km path arrives almost exactly 0.5 s later than the direct path with a relative intensity of just less than 0.3 of the main arrival.

B. Equalizer Mean-Square Error (MSE)

The MSE at the output of the adaptive equalizer governs the achievable data rate given the specific coding gain of the concatenated approach. While the equalizer is good at tracking channel changes as long as the arrivals fall within the filter window, when the multipath is highly delayed, it falls outside the filter interval and the effectiveness of the DFE is reduced. However, use of multiple hydrophone channels helps to make up for this. In Fig. 7 the MSE is shown for the 2 km 25 kHz link. The 4-hydrophone receiver trains to a very good MSE, about -20 dB, before the surface bounce of the initial FM sweep arrives, driving the MSE up to almost -5 dB. The 40 msec FM sweep stops, but soon after the burst used to set the AGC arrives, with a similar effect. After the burst ends the MSE again goes down to -20 dB, but when the delayed version of the packet finally arrives the MSE goes to -10 dB and then stays there for the duration of the packet.

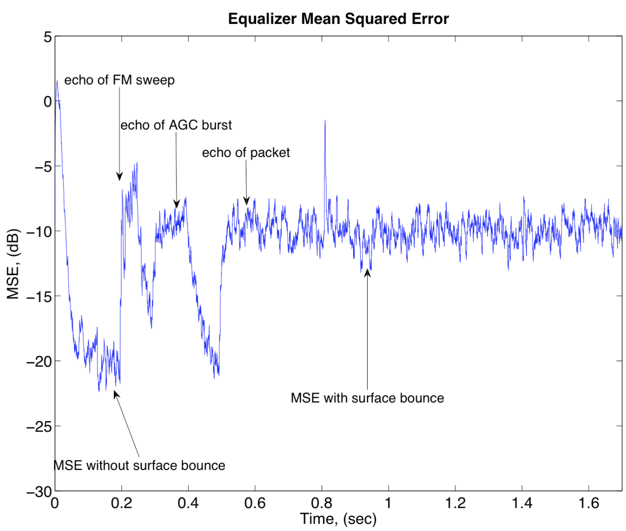

Fig. 7. Equalizer mean-square error for the 2 km 25 kHz case showing the impact of the delayed surface bounce on the receiver. The best case MSE is -20 dB, which would support at least 10-15 kbps throughput if there was no surface bounce.

C. 24-hr Test - 2 km, 25 kHz

A test to examine the performance of the 2 km link at the maximum coded burst rate of 5300 bps was performed on April 1, 2010. The base station was configured to poll the 2 km 25 kHz unit every five minutes. The performance metrics of the real-time modem were captured and analyzed. In Fig. 8 the output SNR from the equalizer is shown. For the coding method used in this test (block and RS code as described earlier), it is expected that if the output SNR is above 7 dB, then the the probability of packet success is nearly 100%.

Indeed, this is the case, and the resulting packet success rate over the 24-hour period is better than 97%.

D. 15-day Test - 2 km, 25 kHz

A longer test was performed with polls of several remote units performed approximately twice per hour (Fig. 9). The success rate for this sequence was approximately 92% after the periods with ROV operations at the node are taken out. Additional analysis of the data is needed to identify the conditions that resulted in this success rate. It was found that the only time where communications was poor for long periods was when the ROV was operating close to the node.

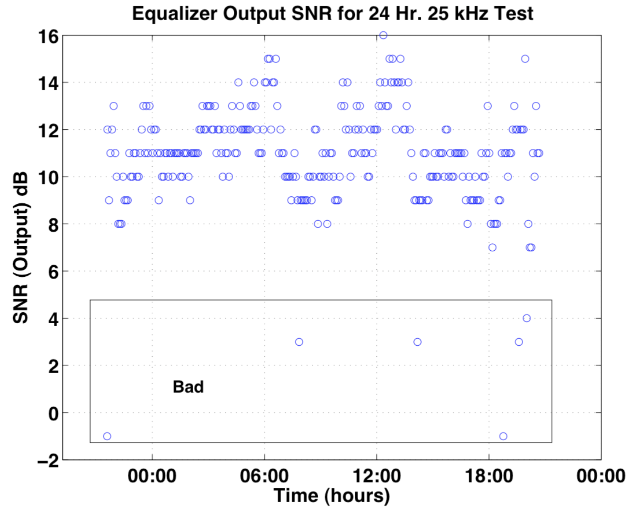

Fig. 8. Output SNR from the adaptive equalizer from a 24 hour test with packets being requested from a remote unit every five minutes. The success rate is 97%.

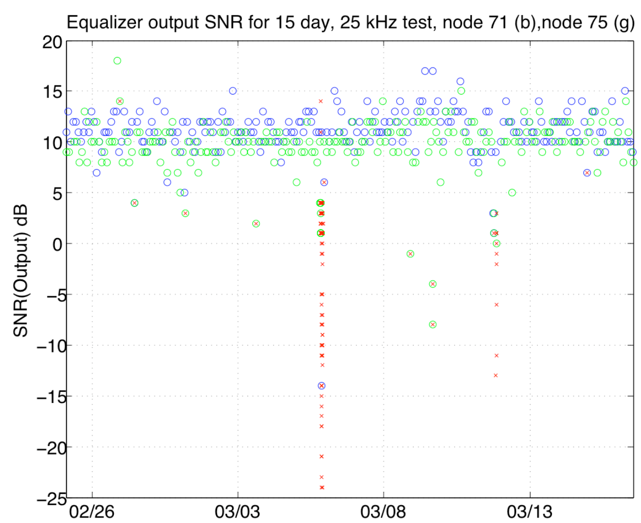

Fig. 9. Output SNR from the adaptive equalizer during the 15 day test for two of the 25 kHz nodes co-located at 2 km range. After the two days with ROV operations are removed the overall success rate is 92%.

V. Conclusions and Future Work

Interim results of testing a bottom-to-bottom acoustic communications link in 900 m water in Monterey Bay shows that reliable and power-efficient acoustic communications is both practical and feasible. Both remote and base station transducers are quite close to the bottom, 50 m for the remote nodes, and less than 10 m for the base station. Despite the proximity to the bottom, the 2 km link has more than enough SNR for reliable communication at 5300 bps, even with long-delay (0.5 sec) multipath from the ocean surface. The use of four elements selected from an eight element array (resulting in 2 m spacing) provides 3-5 dB improvement over a single channel, making a significant difference in possible data rate.

Short (1-day) and medium length (2 week) tests have been performed, with additional tests scheduled. The most important initial result is a very high success rate during a finely sampled 1 day test, and a very good success rate over a 30 minute sampled 2 week test. The surface bounce reduces the possible data rate significantly in comparison with a source-receiver link without any boundary interaction. Thus an important initial conclusion is that the source beam-pattern is critical to maximizing data rate and that narrow beam-patterns will help. A narrower pattern also will have the collateral benefit of improving source level for a given power level, further increasing the power efficiency of the link. While the beam pattern of the node transducer is narrow, those of the remote nodes are essentially omni, and thus they insonify the surface of the ocean quite well, which creates the multipath as observed in Fig. 6.

Future analysis will include comparison with local weather obtained from nearby ocean buoys and examination of the sound-speed and resulting ray paths. A goal is to identify the maximum typical data rate so that data transfer protocols can be designed for this type of centralized network.

Acknowledgments

This work was supported by the National Science Foundation under grant OCE-0532223. The authors would also like to acknowledge all of the great help and support from Steve Etchemendy, Craig Dawe and the ROV and ship crews at the Monterey Bay Aquarium Research Institute who made the deployment possible. At National Instruments Matthew Pollock and Jonathan Cotton provided continual support during integration and testing of the data acquisition system for this specific and non-standard application. The work was conducted in accordance with the MBARI Monterey Bay National Marine Sanctuary Research Permit.

References

[1] C. R. Barnes, V. Tunnicliffe, "Building the World's First Multi-node Cabled Ocean Observatories (NEPTUNE Canada and VENUS, Canada): Science, Realities, Challenges and Opportunities," *OCEANS 2008 - MTS/IEEE Kobe Techno-Ocean*, vol., no., pp.1-8, April 2008.

[2] J. Delaney, P. Beauchamp, M. McNutt, C. Barnes, A. Chave, J. Madden, "Project NEPTUNE: an interactive, regional cabled ocean observatory in the Northeast Pacific," *Proc. OCEANS 2003*, vol.2, no., pp. 1038-1042, Vol.2, Sept. 2003.

[3] C. Pontbriand, N. Farr, J. Ware, J. Preisig, H. Popenoe, "Diffuse high-bandwidth optical communications," *Proc. OCEANS 2008*, vol. 1, no., pp.1-4, 15-18 Sept. 2008.

[4] D. Frye, J. Ware, M. Grund, J. Partan, P. Koski, S. Singh, L. Freitag, J. Collins, R. Detrick, An Acoustically-Linked Deep-Ocean Observatory, *Proc. Oceans 2005 Europe*, Brest, France. June 2005.

[5] L. Freitag, M. Grund, S. Singh, M. Stojanovic, "Acoustic Communications for Seafloor Observatories," *Proc. 4th Scientific Submarine Cable Workshop (SSC06)*, Dublin, Ireland, Feb. 2006.

[6] J-F. Rolin, J. Blandin, V. Lykousis, J. M. Strout, G. Etiope, P. Favali, P. Briole, V. Ballu, G. Papatheodorou, G. Ferentinos, D. Cathy, M. Masson, "Common issues between cabled and non cabled observatories in ASSEM project," *Proc. Oceans 2005 - Europe* , vol.2, no., pp. 872- 877 Vol. 2, June 2005.

[7] L. Freitag., M. Grund, J. Partan , S. Singh , P. Koski , K. Ball , The WHOI Micro-Modem: An Acoustic Communications and Navigation system for Multiple Platforms, *Proc. Oceans 2005*, Washington, D.C. Sept. 2005.

[8] M. Stojanovic, J. Catipovic and J. Proakis, "Phase coherent digital communications for underwater acoustic channels", *IEEE J. Oceanic Eng.*, Vol. OE-16, pp. 100-111, Jan. 1994.

[9] M. Stojanovic, J. Catipovic and J. Proakis, "Adaptive multi-channel combining and equalization for underwater acoustic communications", *J. Acoust. Soc. Am.*, **94**, (3), part 1, pp. 1621-1631.

[10] S. Haykin, Adaptive Filter Theory (4th Edition). Prentice Hall, September 2001.

[11] L. Freitag, M. Johnson and M. Stojanovic, "Efficient equalizer update algorithms for acoustic communication channels of varying complexity," *Proc. Oceans '97*, pp. 580-585, Oct. 1997.