An Ocean Observatory Sensor Network Application

Robert Herlien [*1], Tom O'Reilly [*2], Kent Headley [*3], Duane R. Edgington [*4], Sameer Tilak [#5], Tony Fountain [#6], Peter Shin [#7]

[*] *Monterey Bay Aquarium Research Institute (MBARI), 7700 Sandholdt Road, Moss Landing, CA 95039, USA*

[1]bobh@mbari.org, [2]oreilly@mbari.org, [3]headley@mbari.org, [4]duane@mbari.org,

[#] *California Institute of Telecommunications and Information Technology, UCSD, La Jolla, CA, 92093, USA*

[5]stilak@ucsd.edu, [6]tfountain@ucsd.edu, [7]pshinn@ucsd.edu

Abstract — **We describe our implementation of a novel deep ocean sensor network, the MBARI Free Ocean CO_2 Enrichment (FOCE). FOCE is a system designed for installation in the deep ocean to enable manipulative experiments that explore the impact of deep ocean increase in CO_2 and resulting pH change on ocean biogeochemistry and ecology. This system uses control feedback and pH sensors to inject CO_2 into a small volume of seawater, thus creating a controlled environment per science requirements. To implement this system, we utilized the MBARI-developed network middleware known as "SIAM", which provides a standardized interface to instruments on a sensor network. For the FOCE application we integrated Open Source DataTurbine (OSDT) into SIAM. OSDT provides asynchronous communication links between distributed components, and is particularly well-suited to streaming instrument data. Combined with the existing synchronous SIAM framework, these features enabled a straightforward and efficient architecture for our application. We describe how we achieved our goals of software reuse of infrastructure and instrument services, instrument-in-the-loop control, and rapid assembly of a scalable end-to-end sensor network system.**

I. Introduction

The burning of fossil fuels for energy production has produced cumulative emissions on the order of 1 trillion tons since the beginning of the industrial revolution. While approximately half of the CO_2 has remained in the atmosphere, the ocean is the predominant repository for the remainder of these emissions. This has resulted in a lowering of the preindustrial surface ocean pH by about 0.1 units. If society is able to stabilize atmospheric CO_2 levels at twice the preindustrial concentration, this will result in a lowering of surface ocean pH by another 0.2 units [1]

MBARI scientists and engineers have designed FOCE (Free Ocean Carbon Enrichment) technology to enable the study of the impact of this pH change on ocean biogeochemistry and ecology. The FOCE technology concept enables small-scale, *in situ* CO_2 enrichment experiments to be carried out, in a manner analogous to the terrestrial Free Air CO_2 Enrichment (FACE) experiments [2]. FOCE is used to control the pH within a small volume of seawater that exchanges freely with the surrounding region in the ocean. The technology uses feedback from pH sensors and other instruments to inject CO_2, creating ocean acidification conditions (reduced pH) in a small area, corresponding to future levels of greenhouse gases while maintaining other environmental parameters. In this way, FOCE is like an ocean acidification time machine, allowing researchers to peer into the ocean's future to see the effect on natural ecosystems such as coral reefs, cold water corals, and other sensitive benthic habitats.

FOCE is a general technology concept that is also being adapted outside of MBARI to study coral reefs and other shallow water environments. The MBARI deep ocean FOCE system is currently connected to the Monterey Accelerated Research System (MARS) (http://www.mbari.org/mars), a cabled observatory in Monterey Bay. The MARS cable extends 50 km from shore and reaches a depth of 890 m. The cable provides a total of 10 kW of available power and Gigabit Ethernet to experiments at the MARS site (fig. 1). Researchers remotely control a variety of parameters in real time while monitoring in situ ocean acidification experiments using FOCE and software components known as "SIAM" and "OSDT".

SIAM middleware was originally designed for the Monterey Ocean Observing System (MOOS), a portable, moored ocean observing platform that provides power and a fiber optic network to experiments (networks nodes) on the seafloor through its riser cable. The mooring also provides a wireless two-way communication channel to shore via satellite or line-of-site radio. Unlike cabled observatories such as MARS, moored observatories can be moved from one location to another. But available power and bandwidth to shore are limited on a mooring: 100 W continuous average power and 7200-38400 bits/sec. MBARI has deployed its Monterey Ocean Observing System at several locations within Monterey Bay [9].

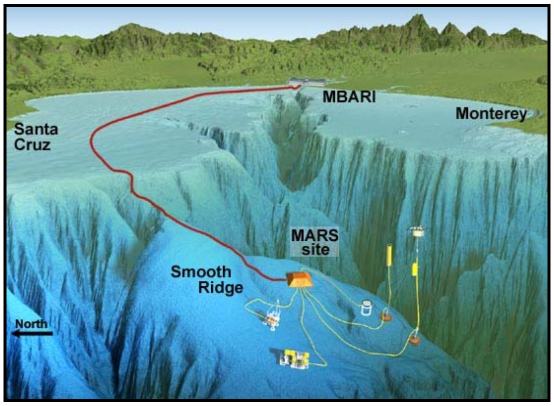

Figure 1: FOCE deployment at MARS node

Both cabled and moored observatories are designed to support a wide variety of science investigations and instruments, including sensors that measure chemical, physical, and biological properties of seawater, ocean floor, and sea-air interface.

MARS and other cabled observatories provide power and data connections to experiments on the seafloor, but do not, in general, provide software infrastructure to collect, archive and process data. Applications to perform these roles may be run either on dedicated hardware attached to the MARS node, or on shore, anywhere on the Internet. In contrast, a stand-alone system like MOOS must operate autonomously, but shares many of the same requirements for processing and archiving with cabled observatories.

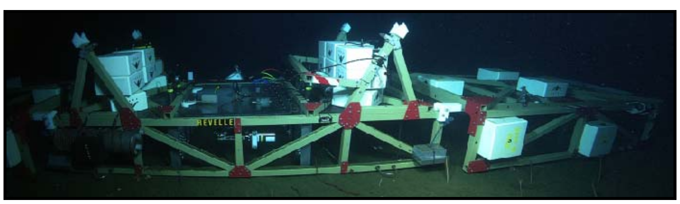

Figure 2: FOCE equipment Panoramic View

On observing systems like MOOS and MARS, most of these instruments communicate with a node's host computer through an RS-232, RS-485, or Ethernet interface. Apart from these standards, most oceanographic instrument manufacturers use proprietary software protocols and data formats for instrument control and data transport. An ocean observatory can include many kinds of instruments supplied by many manufacturers, and so these very diverse protocols and formats must be accommodated by the observatory's software infrastructure.

The primary purpose of SIAM middleware is to implement robust distributed data acquisition on sensor networks like MOOS and MARS. It does this by presenting a common network interface that obscures vendor specific instrument protocols. SIAM's distributed architecture enables it to collect and archive data locally and autonomously, across a network (as on a cabled observing system), or a combination of these. SIAM includes an instrument service framework that is used to write instrument drivers. The framework implements the generic network instrument interface, which controls the instrument through a set of callback methods for operations like initialization, data acquisition and configuration. Support for new sensors requires that callbacks be written to implement the proprietary instrument control protocol of an instrument. Figure 3 illustrates how this works.

II. FOCE Experiment Cyberinfrastructure

1. *System Requirements*

FOCE is intended to serve a broad range of science applications and support both experimental and on-going cyberinfrastructure activities. Key system requirements affecting software design include:

- ***Heterogeneous sensors/instruments***: The system must support integration of numerous instruments from a variety of vendors. These include both commercial and experimental instrument packages.
- ***Real-time processing***: The system must support acquisition and processing of observations in real-time across a range of operating conditions and computational loads.
- ***Security and Data Sharing***: Ensuring security of instruments and collected information is critical. Authentication must be implemented for privileged operations (e.g., command and control of deployed sensors). Data should be managed in a secure fashion, and the system should be flexible in supporting data sharing policies.
- ***Scalability and Flexibility***: The cyberinfrastructure must be extensible and easily configurable to adapt to evolving changes in sensors, networks, and processing platforms.
- ***Interoperability with other Environmental Observing Systems***: The system must support the sharing of resources, including data and analysis tools, with other research observatories, for example the Coral Reef Environmental Observatory Network (CREON) [7]. Exposing FOCE data by using accepted standards will increase its user base and ensure its broad applicability.

2. *Technical Risks and Challenges*

The FOCE equipment is designed for nominal deployment durations of up to two years. Deep ocean FOCE must function remotely in harsh environmental conditions that include extreme pressures, low temperature and corrosive fluids. Field operations require the use of ships and remotely operated vehicles (ROVs) for periodic (monthly) replenishment of CO_2. This is time-consuming and costly. In shallow water sites like coral reefs, high temperatures and biofouling also pose challenges. The remote management of software services and hardware components is essential to minimize expensive field operations. The ability to implement closed loop control across a network over long periods requires a robust software architecture and the ability to reliably stream data between software components comprising the control loop.

3. *Design Principles*

The real-world requirements and challenges outlined above directly inform our design principles and architectural decisions. It was important to design for real-time performance from the beginning. Real-time data acquisition, event detection, and instrument control are fundamental to the operational requirements. Therefore we start from the premise that streaming-data support should be built into the infrastructure from the beginning.

We also needed to employ scalable infrastructure components. All system components should be capable of handling high demands, including data acquisition, transmission, processing, and application servers. Even though the first deployment of FOCE experiment has moderate demands, i.e., few sensors, few processing workflows, and few clients, the system components were chosen to scale gracefully to orders-of-magnitude increases along each of these dimensions.

It was also necessary to architect for a phased deployment. The FOCE experiment is a complex and long-term project. However, there is an immediate science agenda that requires basic data services. So, rather than a multi-year design process followed by construction, we chose to deploy the foundational services of FOCE while additional components are still on the drawing board or under development in the laboratory.

4. *SIAM Middleware*

MBARI's Software Infrastructure and Applications for MOOS (SIAM) is software that manages observatory infrastructure, instruments, and their data [4].

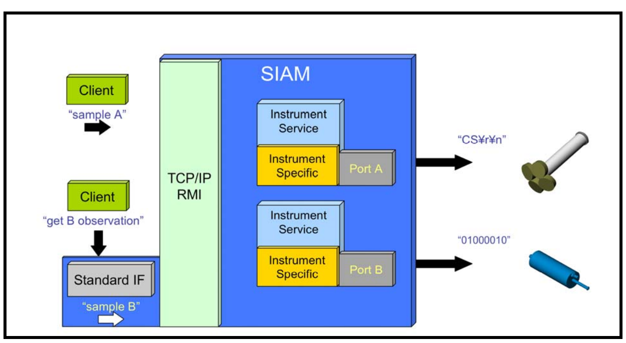

Fig. 3: SIAM provides distributed instrument services and hides proprietary instrument protocols behind a common network interface.

SIAM is implemented in Java, making it portable to a variety of computers and observatory architectures, including cable-to-shore and moored systems. To enable distributed services and clients across networked observatory instrument nodes, SIAM uses Java Remote Method Invocation (RMI). Using RMI, SIAM instrument services present a generic network interface to clients seeking to access the diverse instruments deployed on observatory nodes. This generic interface includes methods that expose instrument operations such as configuration, control, metadata retrieval, and data acquisition. An instrument service framework maps these operations to the proprietary protocol recognized by each instrument.

Every instrument in the observatory has a corresponding SIAM instrument service that presents a generic interface to the observatory network, but communicates with the instrument itself using the appropriate proprietary protocol, thus "hiding" the diverse instrument protocol details from network clients. This design approach results in a system that is scalable to many types of instruments, since clients can access all instruments through the same network interface. For example, a client can trigger sample acquisition from many diverse instruments through the same interface operation. The architecture is scalable to many nodes and different types of instruments and enables both flexibility and interoperability in partitioning observing systems.

The SIAM instrument service interface also defines methods for managing metadata describing the instrument in a standard way, e.g. using SensorML encoding. This metadata may include a standard description of the instrument's data record structure, thus enabling automated data processing and visualization from many kinds of instruments.

5. *SIAM and Open Source DataTurbine*

Combining data from multiple instrument services can be used to create advanced applications like closed loop control (as on FOCE), or automated event detection and response. Applications like these would benefit from access to asynchronous (streaming) data, but SIAM does not natively provide this capability.

The RMI mechanism used by SIAM supports synchronous data transfer between services: a client or service may use it to poll other instrument services for data. An asynchronous publish/subscribe mechanism, e.g. for streaming instrument data or propagating event messages across the observatory network would be more effective for advanced applications like those described above.

DataTurbine is such a real-time streaming data engine. It is an open-source middleware product supported by NSF, NASA, and private industry. It is managed by the NSF-sponsored Open Source DataTurbine Initiative at CalIT2 (www.dataturbine.org) [5] DataTurbine provides (1) a programming abstraction over heterogeneous devices, and (2) integrated network service for managing streaming data. DataTurbine makes disparate devices look similar; e.g., data streams from instruments as diverse as accelerometers and video cameras are easily integrated and managed by a common Application Programming Interface (API). DataTurbine is also portable and scalable. Our tests demonstrated that DataTurbine runs efficiently on platforms from cell phones to supercomputers[6].

DataTurbine can simultaneously support both pull and push modes over intermittent networks with guaranteed data delivery and persistence. This feature is at the core of DataTurbine and was motivated by real-world applications where data needs to be pulled from various devices and pushed to other devices during windows of network availability. From the perspective of distributed systems, the DataTurbine middleware is a "black box" to which applications and devices send and receive data. DataTurbine handles all the data management operations between data sources and sinks, including reliable transport, routing, scheduling, and security. DataTurbine accomplishes this through the innovative use of flexible network bus objects combined with memory and file-based ring buffers. Network bus objects perform data stream multiplexing and routing.

With OSDT, data sources (e.g. SIAM instrument services) parse raw data into discrete channels, and write them to an OSDT ring buffer. Clients may access data by subscribing to data streams at Ring Buffer locations on the network; Ring Buffers may also be mirrored across network segments and automatically synchronize if the network connecting them is intermittent.

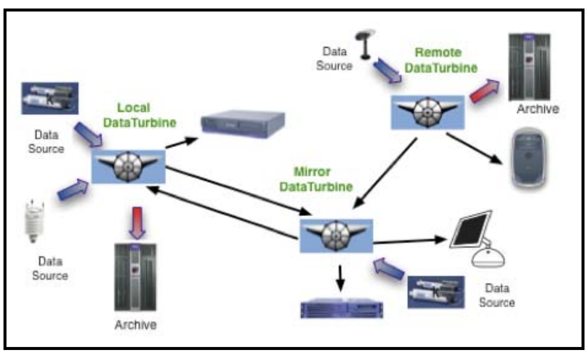

Fig. 4: OSDT is middleware for distributed data streaming.

Integrating OSDT into SIAM was very straightforward. A lightweight adapter class, Turbinator, was added to SIAM instrument services, and an OSDT Ring Buffer instance is run somewhere on the observatory sensor network. Turbinator streams instrument data to the OSDT Ring Buffer, which is then made available to clients (which may include other instrument services) on the network.

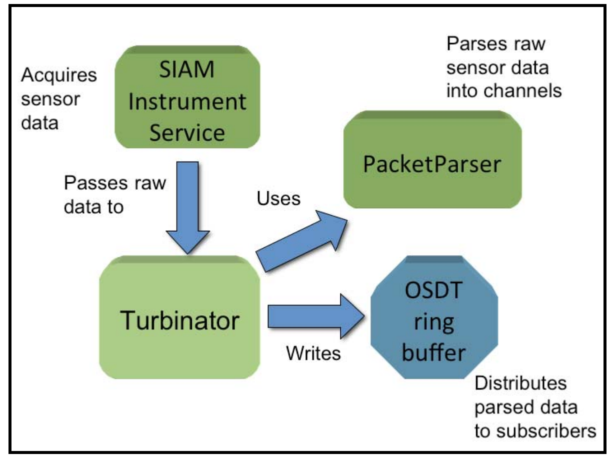

Fig. 5: Integrating SIAM and OSDT is straightforward.

Integrating OSDT with SIAM took only a week or two to complete. In addition to providing this needed data streaming capability, OSDT includes the Remote Data Viewer (RDV) graphical user interface for visualizing data. RDV allows users to playback real time or archived (ring buffer) data, using a powerful and simple drag and drop display interface.

III. FOCE OPERATION AND INSTRUMENTS

Though it is a sophisticated apparatus, the operating principle of FOCE is fairly straightforward. The FOCE flume, open at both ends and on the bottom, rests on the seafloor. For the most part, natural tidal currents maintain the flow of water through the flume. To maintain a pH offset inside the flume that is lower than the surrounding environment, CO_2 is introduced into the water at the appropriate (up-stream) end of the flume, where it reacts with the seawater as it makes its way towards the chamber in the center of the flume. The pH of the resulting CO_2-enriched seawater is decreased, while preserving the temperature and other environmental properties. At the low temperatures of the deep ocean, this reaction can take a few minutes to reach equilibrium [10].

A series of baffles provide a longer current path, to allow time for the reaction between the CO_2 and the seawater to complete before it reaches the chamber. Fans may be used to adjust the flow of water when ambient currents are small. Louvers on the CO_2 delivery subsystem adjust the mixture of CO_2 and seawater.

A block diagram showing the core sensors and actuators FOCE uses to maintain a pH offset from the surrounding environment appears below (fig. 6), followed by a brief description of each.

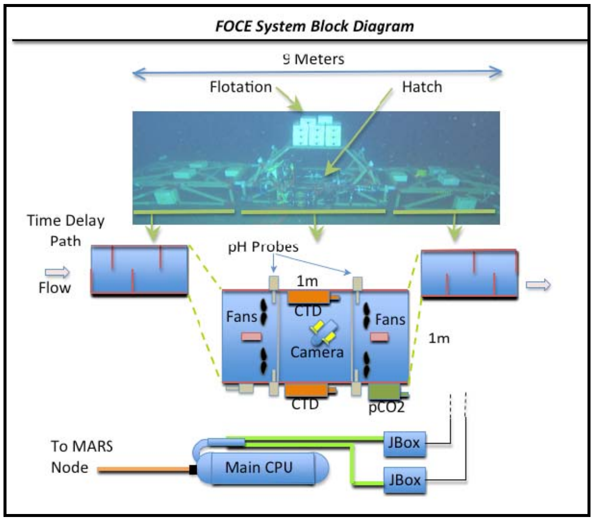

Fig. 6: FOCE system diagram, as deployed on MARS observatory.

pH – Four pH sensors in the chamber account for possible differences due to water transit time and eddies. Two external pH sensors monitor the ambient pH of the region outside FOCE.

CTD – Measures conductivity (proxy for salinity), temperature, and depth. This standard oceanographic instrument provides a very accurate temperature measurement that is essential to both the calculation of pH and the kinetics of the chemistry. Salinity is a necessary component to calculate corrections for pH sensor drift (see pCO_2). .

Velocimeter – Acoustically measures X, Y, and Z components of water velocity at a single point.

ADCP – Acoustic Doppler Current Profiler. Uses Doppler shift of acoustic pings to measure ambient (external) current at various distances from the instrument. The ambient currents are needed to direct the injection of pCO2 and control the fans, which augment current when ambient currents are small.

pCO_2 – Measures partial pressure of CO_2. Used in conjunction with the CTD to recalibrate the pH sensors, which are subject to drift.

Motors and Fans – Two Elmo Solo motor controllers, each attached to a brushless DC motor and two fans, to control the water velocity through the experimental chamber.

Cameras – Lit by LED light panels to visually observe conditions in the chamber and inspect the instruments. Currently there are two VGA Ethernet video cameras and a higher resolution still camera is used to photograph specimens.

CO_2 Injection Subsystem -- This subsystem is currently being tested. It consists of a DigiOne Ethernet-to-serial converter connected to an Elmo Solo motor controller. The motor controls a set of louvers in the CO_2 injection system, which sets the CO_2 injection rate. The CO_2 software control system will use this protocol to command and monitor the subsystem.

Monitors, Alarms, and Engineering Sensors -- Building a complex electrical system inside housings to withstand ocean pressure is a very difficult and error-prone business. We monitor numerous parameters inside the pressure can, including humidity, temperatures of various components, and a water intrusion detector.

IV. Results

A. Current Status

We have developed a prototype FOCE system that uses SIAM with OSDT. The prototype has been deployed in our working FOCE system at a depth of 890 meters. We have completed several engineering test deployments to verify the mechanical, electrical, and software design of the system. We have used the described software architecture to characterize the system response to CO_2 input *in situ*, and have maintained a manually controlled pH environment over several days (fig. 7).

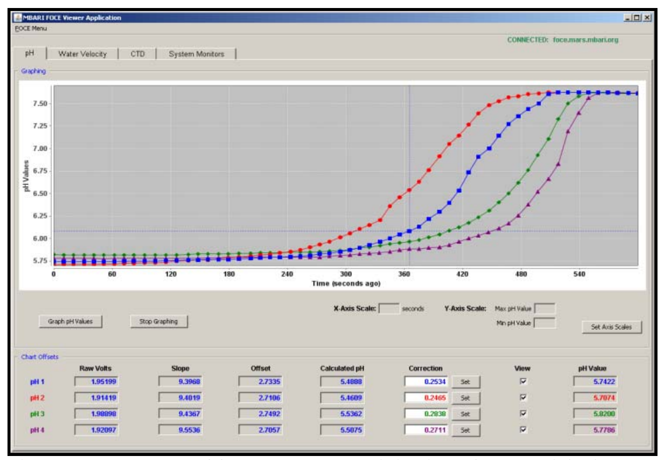

Fig. 7: Step response of four pH sensors to injection of CO_2 enriched seawater. Time proceeds from right to left in this chart (like a strip chart recorder). Note the time lag between the four sensors, representing their physical placement in the chamber and the movement of water past each sensor.

At this writing, the closed loop control architecture and algorithms are being finalized. We will write a control service that uses OSDT to fuse data from the FOCE sensors to control the various actuators (fig. 8). The integration and testing of SIAM and OSDT has been very successful. We routinely use RDV to monitor FOCE data streams, and are in the process of integrating OSDT with the FOCE control interface as well. We have used these to perform compelling live FOCE technology demonstrations to internal and external audiences.

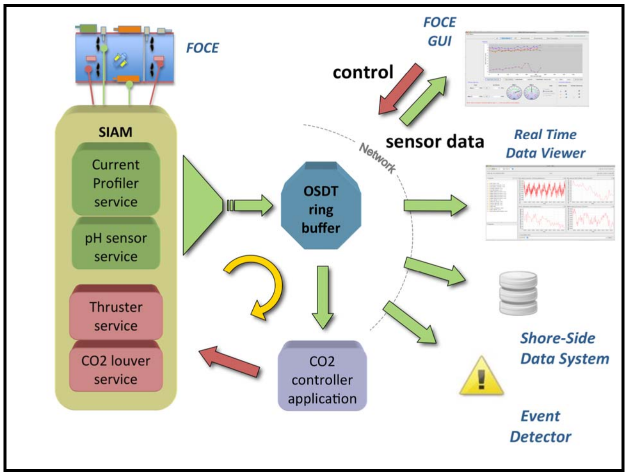

Fig. 8: OSDT enables closed loop control and other advanced applications.

The closed loop control implementation is expected to begin testing *in situ* by the end of 2010, followed by algorithm tuning and development and science experiments in 2011.

V. CONCLUSION

We used the SIAM middleware developed by MBARI [4] to operate instruments deployed on TCP-IP network-based ocean observatories. We integrated Open Source DataTurbine[5][6] (OSDT) middleware with SIAM to provide data streaming functionality. OSDT provides asynchronous communication links between distributed components, and is particularly well suited to streaming instrument data. Combined with the existing synchronous SIAM framework, these features enable a straightforward and efficient architecture for our application. With this approach we achieved our goals of software reuse of infrastructure and instrument services, instrument-in-the-loop control, and rapid assembly of a scalable end-to-end sensor network system.

ACKNOWLEDGMENTS

This work is supported at UCSD by a grant from the Gordon and Betty Moore Foundation, and NSF OCI Award 0722067. Work at MBARI is supported by generous support from the David and Lucile Packard foundation, and NSF OCE 03-30428.

REFERENCES

[1] Caldeira K, Wickett ME 2003 Anthropogenic carbon and ocean pH, NATURE 425 (6956): 365-365 Sep 25, 2003

[2] Long, Stephen P. et al, “Food for Thought: Lower Than Expected Crop Yield Stimulation with Rising CO_2 Concentrations”, Science 312, 1918 (2006); DOI: 10.1126/science.1114722

[3] P.M. Walz, W. J. Kirkwood, K.C. Hester, P.G. Brewer, “Creating controlled CO2 perturbation experiments on the seafloor: development of FOCE techniques,” in Proceedings of the Marine Technology Society and Institute of Electrical and Electronics Engineers Oceans '08 Conference, Kobe, Japan, April 8-11, 2008.

[4] O'Reilly, T. C. et al., “Instrument interface standards for interoperable ocean sensor networks,” in Proceedings of the Marine Technology Society / Institute of Electrical and Electronics Engineers Oceans'09 Conference, Bremen, Germany, May 11-14, 2009.

[5] Tony Fountain, S. Tilak, P. Shin, P. Hubbard, L. Freudinger, "The Open Source DataTurbine initiative: streaming data middleware for environmental observing systems," presented at 33rd International Symposium on Remote Sensing of Environment, Stresa, Italy, 4-8 May, 2009. (http://www.dataturbine.org/biblio)

[6] Sameer Tilak, P. Hubbard, M. Miller, T. Fountain, "The Ring Buffer Network Bus (RBNB) DataTurbine Streaming Data Middleware for Environmental Observing Systems," in Proceedings of the third IEEE International Conference on E-Science and Grid Computing (e-Science), 2007.

[7] CREON: *The Coral Reef Environmental Observatory Network*, http://www.coralreefeon.org/

[8] T. O’Reilly et al., “Software infrastructure and applications for the Monterey Ocean Observing System: design and implementation”, in Proceedings of MTS/IEEE Oceans ‘04 Conference, Kobe Japan, 2004.

[9] M. Chaffey, et al., “MBARI’s buoy-based seafloor observatory design”, in Proceedings of MTS/IEEE Oceans ‘04 Conference, Kobe Japan, 2004.

[10] E.T. Peltzer, P. Walz, P.G. Brewer, “Cabled Observatory Technology for Ocean Acidification Research,” Proceedings of IEEE Oceans ’09 Conference, Bremen Germany, 2009 (090107-010)