Complex real-time environmental monitoring of the Hudson River and estuary system

H. R. Kolar
J. Cronin
P. Hartswick
A. C. Sanderson
J. S. Bonner
L. Hotaling
R. F. Ambrosio
Z. Liu
M. L. Passow
M. L. Reath

Multiparameter and multiscale real-time environmental monitoring of a river and estuary system will be realized through the River and Estuary Observatory Network (REON) for the Hudson River in New York. In this paper, we describe a system under development that provides a holistic view of this complex and dynamic natural environment for scientific research, education, management, and environmental policy-related applications. The system incorporates a complex array of sensor technologies encompassing the physical, chemical, and biological measurement domains. REON supports Lagrangian, Eulerian, and autonomous robot sensor deployments, as well as flexible telemetry options through an open and consistent middleware architecture with advanced device management capabilities. Multimodal data streams are ingested and analyzed by an intelligent distributed streaming data analysis system known as ***System S****. The challenges of managing high volumes of complex heterogeneous data are addressed via a distributed network of intelligent computational nodes that incorporate both autonomic algorithms and active knowledge management including a temporal component. REON provides an adaptive computing environment that provides isotropy in terms of data access and collaborative computation in contrast to traditional hierarchical control systems for sensor environments. Also presented is the underlying information infrastructure that supports a robust and integrated data modeling, simulation, and visualization manifold.*

Introduction

A number of technological advances exist that are applicable to environmental monitoring and that make it possible to improve our knowledge of river and estuarine systems. Rivers and estuaries are dynamic ecosystems, and natural perturbations, such as precipitation or snowmelt, can dramatically alter their physical and chemical characteristics. Anthropogenic perturbations, such as those from chemical spills, pollutant discharges, water withdrawals, and physical alterations to land and water, can also dramatically affect the river and estuarine system. These changes can occur in a matter of hours, with attendant changes in the ecology of the system and the quality of human use of the resource. An example of a deleterious effect is the harmful impact on commercial and recreational fishing of storm water runoff, which transports nutrients (e.g., phosphorus, nitrate, and nitrite) into the ecosystem from fertilized ground sources. These nutrients can stimulate phytoplankton (microscopic algae) "blooms," which in the case of some phytoplankton species (e.g., *Karenia brevis*) produce toxins that can poison fish, birds, marine mammals, and other organisms. Large algal blooms may also induce hypoxic conditions in which the dissolved oxygen in the water is less than 2 mg/L, as oxygen is consumed during the decomposition of phytoplankton.

Improved understanding of these complex and dynamic systems is of interest to the global environmental community given that human population densities in the geographic areas around rivers and estuaries are typically high. These systems are naturally complex, and the addition of factors related to coupled human and natural systems adds further complexity across the environmental, hydrological, and water management

0018-8646/09/$5.00 © 2009 IBM

domains. Often our knowledge of riverine and estuarine systems is based on limited measurements, both spatially and temporally, that result from focused studies driven by a particular environmental concern or issue. Such studies are generally limited to snapshots in time that cannot reflect the dynamic nature of the ecosystems being measured. Since these systems can undergo significant perturbations over relatively short periods with long-lasting effects, field studies that are not underway during these short-lived but far-reaching phenomena are hindered from contributing data that create valid baseline information or establish predictive models. This can have significant ramifications for public policy, such as that enunciated in the 1972 Clean Water Act, which calls for the restoration of the biological, physical, and chemical integrity of the waters of the United States yet provides no models or references against which to measure the restoration of aquatic ecosystems or the efficacy of the law itself [1].

A number of observatory systems have been deployed in estuarine environments, including the Shoreline Environmental Research Facility (SERF) in Corpus Christi Bay [2], MYSound in the Long Island Sound [3], and the Columbia River Environment (CORIE) system [4]. Monitoring programs are underway in the Chesapeake Bay [5], between Maryland and Virginia, the Albemarle-Pamlico Sound in North Carolina [6], the Parker Rowley and Ipswich Rivers at the National Science Foundation (NSF) Long-Term Ecosystem Research (LTER) Plum Island site in Massachusetts [7], and RiverNet on the Neuse River in North Carolina [8]. While these systems are making significant contributions to the science of river and estuarine environments, additional benefits would be realized with the addition of real-time data analytics and an integrated data environment across a multitude of sensor and modeling domains. As observatories realize increasing numbers of sensor types and numbers, real-time analytics are useful not only for monitoring, but also for helping to trigger or guide modeling of key events and to ensure that model input data is available in the proper format. Autonomic analyses and data preprocessing can reduce or eliminate manual preparation of data and can assist in model selection, expedite model execution, and aid in model tuning. For example, rapid prediction of the movements of contaminants may help improve emergency response or mitigation actions for a vulnerable public water supply. More complex models will naturally evolve as data access and availability improve in an integrated system and as data sharing and collaboration improve across observatories.

To provide next-generation advanced monitoring for rivers and estuaries, Beacon Institute for Rivers and Estuaries and IBM are in the process of developing REON, the River and Estuary Observatory Network for the Hudson River in New York. This collaborative project combines the development and integration of a number of technologies and approaches to enable the extraction of new levels of knowledge using a multitude of physical, chemical, and biological sensors, providing data in real time for the Hudson River and estuary system. The architecture for REON is unique, as it provides flexibility and adaptability that enable researchers to collect large amounts of streaming data from a mixture of sensor feeds across a large area (along 315 miles of the Hudson River). REON also facilitates intelligent analytics near the "edge" of the network (e.g., a physical edge of a network to which devices are connected) and supports an integrated information infrastructure through which scientists, engineers, policy makers, and educators are able to navigate and access all data. The architecture also establishes an extensible framework for the operational aspects of advanced water management through device management capabilities critical to ensure cost-effective deployments and ongoing operations. The latter is an important area that impacts the total cost of ownership (TCO) and that often has not been addressed at the time of system design.

Knowledge domain and data requirements

A wide range of research interests in river and estuary systems exist that require various physical, chemical, and biological sensor inputs. These may include water quality study and evaluation, habitat and ecosystem monitoring (e.g., invasive species monitoring), environmental impact analyses, and water management. The sensors required range from those monitoring traditional variables (such as temperature, salinity, dissolved oxygen, turbidity, chlorophyll, pH, nitrates, and pressure) to more complicated biological sensors that may sometimes be as sophisticated as self-contained miniature laboratories (such as the Autonomous Microbial Genosensor designed by University of South Florida researchers to provide field detection and quantification of harmful algae [9]). Special applications are also under consideration and could include the radio or acoustic tagging and tracking of fish as in the Rutgers University striped bass (*Morone saxatilis*) tracking project [10], and acoustic sensors could also be used to detect the proximity of newly hatched fish to power plant cooling intakes: A short delay in the onset of pump activity could provide protection for at-risk populations or species.

For REON, the initial sensor deployment consists of a sensor array fitted on a moveable robotic arm for depth profiling and is located in the mid-Hudson region near the Beacon Institute headquarters, approximately 50 miles north of New York City [11, 12]. Additional deployments of similar robotic arrays are planned for

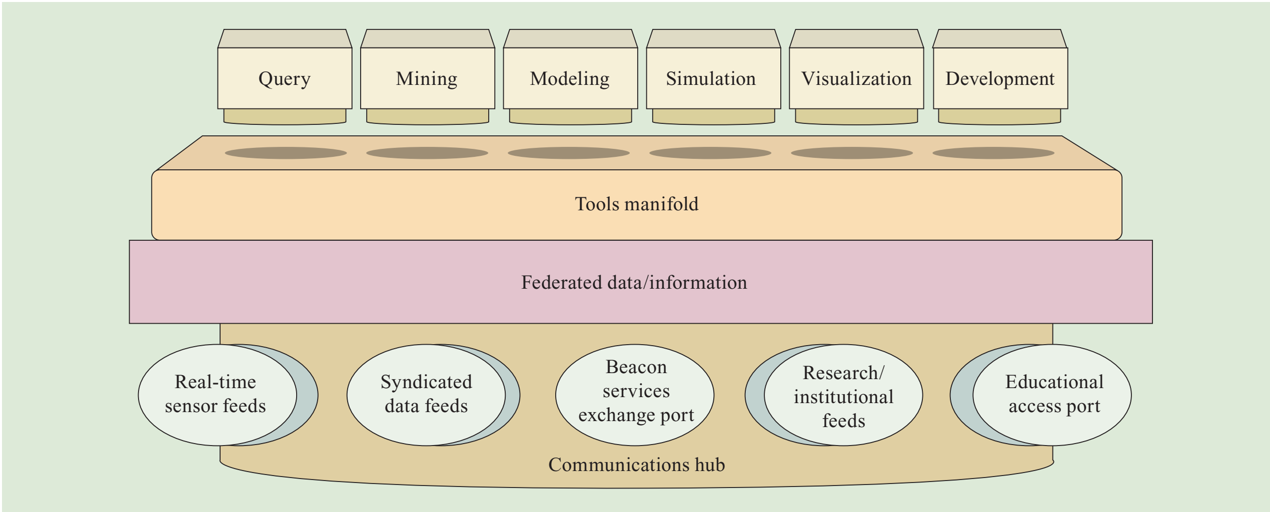

Figure 1

The information architecture for REON supports heterogeneous data types and bidirectional multiscale data feeds in order to serve scientific, engineering, policy-making, and educational users. A comprehensive query and analytical toolset is supported for the diverse user community via an open-access tools manifold.

multiple locations along the river. The sensors in this array provide measurements of salinity, turbidity, temperature, conductivity, dissolved oxygen, chlorophyll a, particle size (using laser *in situ* scattering and transmissometry, or LISST), colored dissolved organic matter (CDOM), and total petroleum hydrocarbons (TPHs) for multiple depths along the water column. The platform also hosts an acoustic Doppler current profiler (ADCP) to measure water current structure and a meteorological sensor package to monitor meteorological conditions such as wind speed and direction and atmospheric pressure.

A comprehensive observatory system must be able to support and manage a diverse set of sensors, in large numbers, spatially distributed, and with varied telemetry requirements for fixed (Eulerian) and moving (Lagrangian) deployments (including robotics-based autonomous data-acquisition and data-collection platforms) [13]. We do not discuss here the complexities of the physical sensor placement and management, which include positioning, mooring, power options, cabling, corrosion, and biofouling. These are all important aspects in the deployment, efficacy, and maintenance of observatory systems.

Bidirectional dataflow is necessary for effective sharing of logistical information and sensor setup and control commands for agility in data-collection strategies. To cope with changing conditions, adaptive sensing techniques provide modifications to sampling plans that may be generated by device-based autonomous computing algorithms or sent to the sensor platform by researchers through telemetry. Collection of real-time data, with ensured quality and minimum latency, is an imperative in order to allow researchers, policy makers, and water management experts to advance the knowledge frontier and subsequently protect and conserve natural resources. Data quality may be validated at or near sensors in a distributed architecture to prevent erroneous measurements. The implications of poor data quality are broad and include unnecessary expense associated with the power needed for data transmission (in the case of remote self-powered sensors) and actual transmission costs (e.g., for satellite or commercial communication networks). By making this data available in a standardized manner, supporting "fusion" of data with external data and information sources (such as sources of geospatial and meteorological data), and providing it to a range of users with appropriate analytical tools, REON differentiates itself from existing observatory systems. An information architecture supporting multiple data sources and users is a crucial aspect of REON that ensures that the massive amounts of data may be consistently exploited to extract new information and knowledge about the river and estuary system (**Figure 1**).

The information architecture of REON relies on data federation to ensure straightforward access to all data sources. Federation provides a single virtual data source or data service conforming to a services-oriented architecture (SOA) design philosophy [14]. Metadata is a critical aspect of federation. For example, metadata may be used to describe the type of data provided by a particular sensor type. Multiple data stores of various types, platforms, and physical locations, involving multiscale measurements that typically span mesoscopic and macroscopic dimensions, will be available to users. A federated data infrastructure supports interoperability and helps eliminate the complexities of data warehouse approaches that require data replication and synchronization via extract, transform, and load (ETL) processes.

Data and information are exchanged through the communications that support sensor inputs, syndicated data from external data sources (such as data that is enhanced or modified by external service providers), access for other research institutions, and bidirectional exchange with educational institutions. One notable example is a pilot program whereby local secondary schools will be deploying sensors in the Hudson River, collecting data, and transferring the data to the REON system for analysis and visualization. Through a services exchange port, REON will also be capable of supplying various data and information (including modeling output) to private-sector collaborators (who are performing technology development), municipalities, or organizations with unique data requirements.

The REON tools manifold provides a flexible and consistent set of interfaces and services to support a wide range of tools to facilitate both hypothesis-based and data-driven studies, as well as an assortment of modeling and advanced visualization technologies from varied sources (e.g., open-source, research, and commercial sources).

Cyberphysical systems: Integration of sensor networks and information technology

Sensors and their immediate connectivity may be characterized as a control-system layer in the overall system hierarchy. These layers are designed with very specific characteristics including asynchronous signaling, real-time and high-performance functioning, and high availability. They regularly involve specifically designed hardware and software optimized for this purpose, and the associated performance metrics differ from those of traditional IT (information technology) systems. These technologies may span deterministic closed-loop proprietary fieldbus-based systems for control I/O (input/ouput) used in automation applications with mixed analog and digital devices to more esoteric scientific and engineering sensors that may be referred to as *instrumentation*. (Here, the term *fieldbus* refers to a family of industrial computer network protocols used for real-time distributed control.) These scientific and engineering sensors may be characterized as more complex, research-grade (i.e., non-mass-produced), and nonstandardized sensors exemplified by many of the leading-edge developments in chemical and biological sensing. Although recent trends have supported the migration of Internet Protocol (IP)-based technologies into control systems, there still remains a significant technology difference between these domains, as they are designed and managed by very different user groups using different technologies, including software programming models, tools, and frameworks.

Closing the gap between control systems and traditional IT is important in order to provide a new level of operational effectiveness for applications such as environmental observatories. A closer integration is necessary through a consistent architecture that provides enhanced interoperability and optimization across a number of application domains and that allows maximized exploitation of data within these systems. An example of this is a holistic and federated information infrastructure that supports varied types of users. The user types range from researchers using sophisticated and computationally demanding modeling and simulation algorithms and advanced visualization on supercomputers to K–12 (kindergarten through 12th grade) educational users examining simple datasets for school projects. An effective architecture permits accessibility to data at all locations from any point in the system in a consistent manner using a complement of query and analytical tools. Exercising an SOA approach and incorporating appropriate data standards are key design elements important not only to REON as a local observatory system, but also for the integration and data sharing with other environmental observatories.

The architectural foundation for REON relies on the use of advanced middleware developed by IBM Research to create a distributed system of heterogeneous automation objects. These software objects perform various services and interact in a consistent manner. We are adapting a foundational application framework for REON using Internet-scale Control Systems (iCS) [15] technology, and we use Harmony [16] as the network management layer. These technologies have been designed for large-scale implementations where we expect to support hundreds to thousands of sensors and control devices (i.e., actuators) in an inherently heterogeneous environment.

The role of iCS is to provide an efficient abstraction and event-based programming model that establishes an integration environment for data collection and processing of automation objects spanning the sensor and IT domains. When extended to an operational model, iCS supports the implementation of functionality encompassing information collection, decision making, and command execution. iCS uses a sense-and-respond approach that spans the physical control loop domain

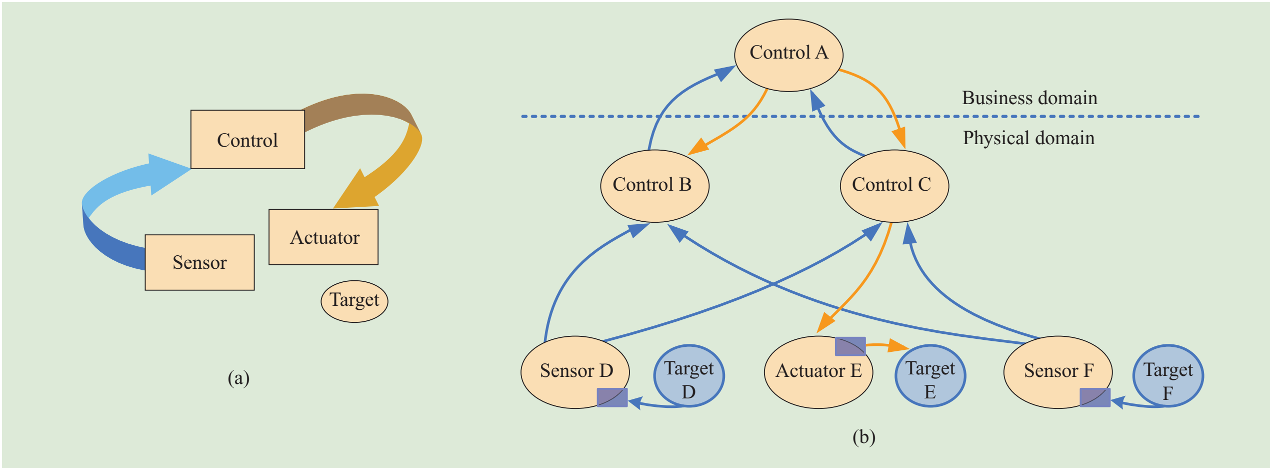

Figure 2

The iCS object model for the cyberphysical environment: (a) the abstraction of control system elements required to support sense-and-respond functions that interact with a target in the physical or business worlds; (b) a schematic representation of an application that has been composed across the sensor and IT (physical and business) domains through event-based programming. In the figure, blue arrows represent sensor outputs (even outputs from a control object appear as sensor outputs to other objects), and the gold arrows represent actuator inputs.

and higher-level decision processes characteristically found in IT systems. Here, the phrase *physical control loop domain* refers to the operations of sensors and actuators. In REON, iCS (combined with Harmony) provides the sensor and control network integration into System S (described below), in which the primary information analysis and management will take place.

iCS provides standardization through common information taxonomies and lexicons and open-standards design points such as the use of a common XML (Extensible Markup Language) schema-based application abstraction model to describe objects for both the control systems and the IT environments (the physical world and business world domains, respectively). Applications are created from sensors, actuators, and control objects [**Figure 2(a)**] by defining the connections between output and input variables of the various objects, as depicted in **Figure 2(b)**. To support heterogeneous environments, the iCS framework enables the encapsulation and abstraction of existing objects so that they may be used in the new integrated cyberphysical environment. This is accomplished by instantiating XML description objects as runtime code objects that supply management functions for the asynchronous event communication links between the objects. Object functions may be implemented either directly as new application logic encapsulated by the iCS object or through logic that provides an interface between existing code, such as low-level device drivers, and sensors. The asynchronous event communication links between objects define an abstract event bus that is used as the uniform interprocess communication model in iCS. This event bus is implemented in memory for local (same-node) communication between iCS objects or is utilized across any combination of underlying transport protocol layers for distributed communication between iCS objects. As mentioned, in REON, iCS uses Harmony as the underlying network communication and management layer. This aspect is described more fully below.

Within the construct of an overlay network, these iCS abstraction and encapsulation objects preclude the creation of traditional protocol-level gateways, as each object essentially performs application-level bridging between a particular sensor or control network and the iCS abstract event-programming environment. This provides a high degree of isolation from the underlying protocols used in sensor and control networks and thereby supports extremely heterogeneous and changing environments without a significant impact on the actual application.

For data interoperability, iCS defines a set of XML-schema physical measurement data-type primitives on the basis of seven fundamental types derived from basic physics such as length, mass, and time, as well as binary (digital) types and the provision to extend the data-type schema using logical (or abstract) types from the business domain, including such types as currency, price to earnings (P/E) ratio, and key performance indicators (KPIs). An electrical energy example is shown in **Figure 3**. The physical data-type primitives can be combined as required to create any complex physical data type.

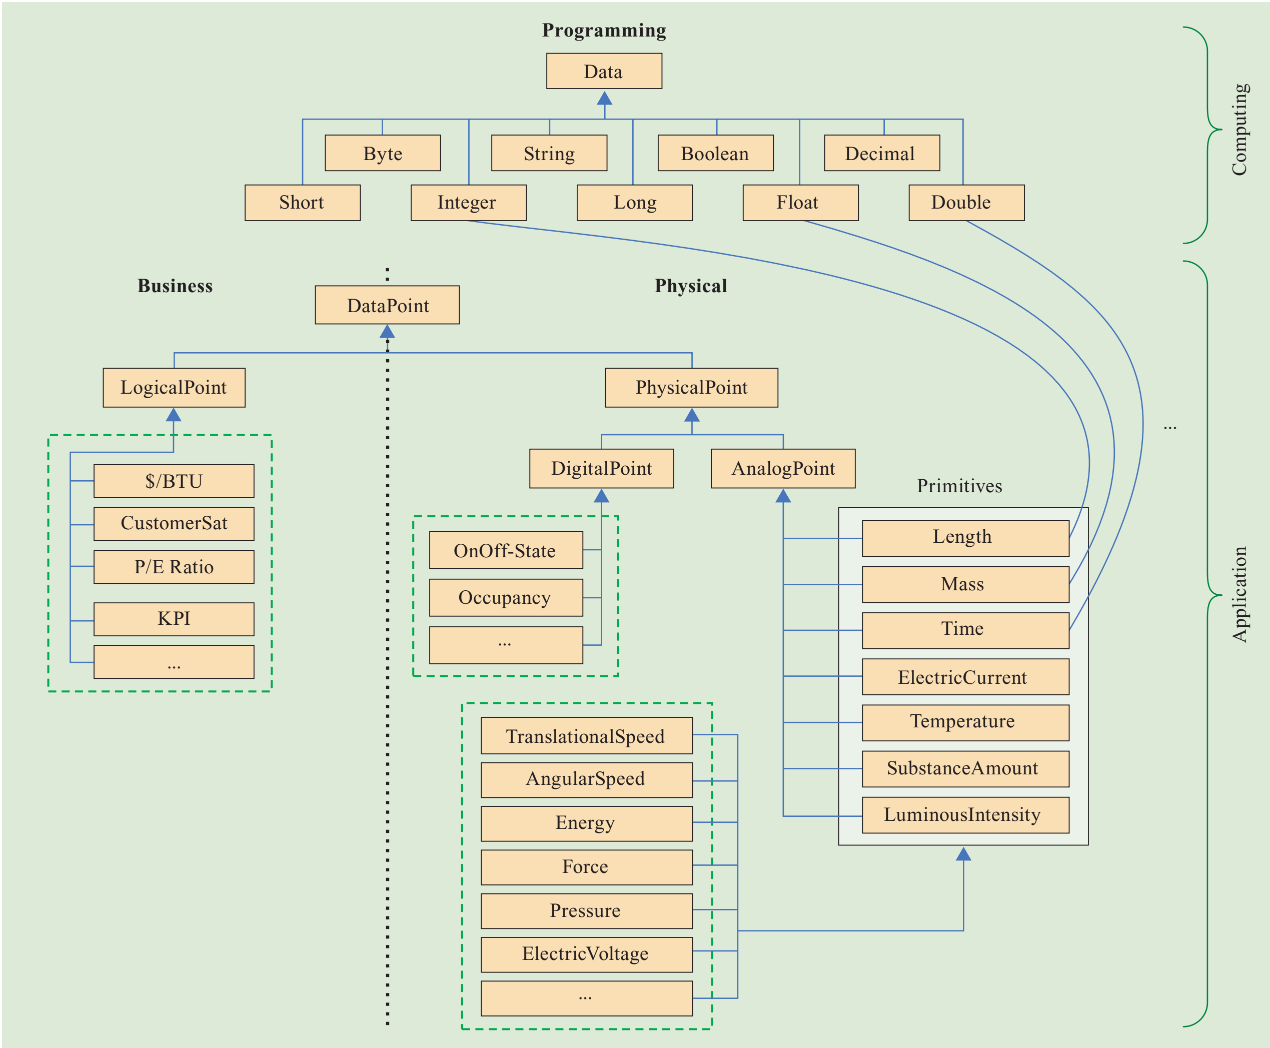

Figure 3

The iCS XML schema, a portion of which is schematically represented here, provides a consistent data-type framework for mapping the computing space for both the IT (business) and operations (physical) domains. Data types are extensible to different applications and provide a data-type platform for model object inputs and outputs.

A key to providing heterogeneous interoperability through iCS is that there is an unambiguous, one-to-one mapping from each of the seven data-type primitives to a single computer language data type in the underlying language used in a particular implementation of iCS (Figure 3). From that, all complex physical data types can be mapped to a combination of underlying computer language representations, moved as needed through the event bus, and reconstituted on another node even if the implementation language of iCS is different on that node.

Specific properties of devices such as sensors are described via the object schemas and can include I/O details (including data-type and quality of service [QoS] requirements that can be passed to the underlying messaging layer if needed). The schemas can also include general object properties and any associated object logic. The same object construct applies to the IT domain where a virtual sensor may be designed to monitor particular state changes with a pointer to logic that could, for example, invoke a Web services interface to access real-time weather data. This is one of the methods iCS uses to bridge the event programming and SOA domains. Software encapsulation is used for conformity to a common sensor-control-actuator object abstraction that allows for the creation of a uniform distributed application model across a highly heterogeneous set of underlying systems.

As mentioned, the underlying network management layer is provided by Harmony, innovative messaging middleware designed for distributed, event-driven applications. Harmony transports and delivers event information to applications while simultaneously satisfying requirements related to QoS guarantees, resiliency and availability, and security and privacy. The deployment of Harmony in the REON setting is particularly appropriate because of the QoS, availability, and security requirements that Harmony provides for REON.

Because of the heterogeneity of sensor types and sensor data, different QoS requirements on the message delivery middleware should be supported. Some sensors, such as those mounted on autonomous vehicles, need to transmit real-time data (including location information) to enable control in real time. Other sensors, such as video cameras, may require network communication throughput guarantees. A more challenging issue in sensor networks is resiliency and availability because of the nature of unreliable environments where sensors are deployed. The messaging system must have the capability to dynamically reconfigure itself to ensure message delivery in response to node and link failures, in particular those caused by correlated failures within a particular geographic area. Another critical capability that is needed is the support of security and privacy requirements. In particular, some of the information collected in water sensor networks needs to be protected.

Harmony middleware provides these essential capabilities required by event-driven applications in a holistic way [16], and none of the commercial products or research tools currently available can achieve this goal. The key technologies developed in Harmony are a comprehensive quality model, overlay routing support, and trusted virtual domain management. The Harmony communication infrastructure supports a scalable sensor environment with its underlying network management layer and is designed for distributed, event-driven applications. In the REON deployment, sensor data feeds into Harmony infrastructure through an Internet access point. The data are then forwarded to the processing center with the overlay network managed by Harmony.

Data collected in the sensor network are then processed in real time in a stream computing platform known as *System S*. Stream computing is well suited for the processing of sensor data from environmental monitoring, where data collection is continuous and occurs at diverse rates because of the heterogeneity of sensors and applications. The stream computing paradigm supports real-time analysis of such data. Various tools have been developed for the streaming of data into processing platforms. Of particular interest is DataTurbine, an open-source streaming data system that can ingest a variety of data sources [17]. It is important to note that in our architecture, System S is used for the deployment of analytics functionality such as filtering, aggregation, correlation, and modeling. System S analyzes incoming (streaming) data continuously and in real time. Such stream processing technology is complementary to streaming solutions such as DataTurbine, which can be used at the data-ingestion points of System S. DataTurbine can also be used as an output adapter for outward streaming of analytical results from System S.

System S, an IBM Research prototype [18–20], has been designed with the goal of functioning well under high loads and dynamic input, and the system continually adjusts its resource allocations to support the highest-priority activities. System S allows different types of analytic functionality (from simple filtering to sophisticated analysis and modeling) to be run cohesively on the same platform. Jobs can be deployed in a System S runtime environment and interact with Web services and back-end databases. Jobs can also interact with the physical world and to control sensors and control devices such as actuators. System S allows REON to easily increase and decrease the number of sensor sources and the amount of data to be processed. This is important for observatory systems since implementations are often phased and driven by focused research requirements and policy issues.

Summary and concluding remarks

The potential value of REON to fundamental research, environmental policy, and water management cannot be overstated. River and estuarine ecosystems, as well as understanding of the human activities that affect the water and watershed, will benefit greatly from data collection that is as dynamic as the ecosystem itself. Resource management programs that can respond dynamically to chemical, physical, and biological alterations will revolutionize protection of the ecosystem and help to manage and control ecosystem-dependent human activities. Although REON is currently under development, this integrated system for a real-time observatory network will provide new insights and approaches that will be globally applicable to rivers and estuaries. While standardization is a key goal, REON is far from mature. However, the design philosophy of REON provides a reference architecture applicable to similar environmental systems and a flexible testbed for further development of advanced environmental sensor technologies and supporting distributed intelligent infrastructure. Interest in the integration of observatory systems continues to be high, as evidenced by the

National Science Foundation Ocean Observatories Initiative (OOI) and Water and Environmental Research Systems (WATERS) Network initiative [21]. The open, scalable, and extensible architecture of REON facilitates the integration of multiple observatories with their respective heterogeneous technologies, and planning is already underway to extend the architecture of REON to the St. Lawrence River and to the Corpus Christi Bay and estuary system, which includes the Nueces River, Nueces Bay, Oso Creek, and Oso Bay.

References

1. Federal Water Pollution Control Act Amendments of 1972, Public Law 92–500 (October 18, 1972); see *http://www.epa.gov/history/topics/fwpca/*.
2. Shoreline Environmental Research Facility (SERF); see *http://www.serf.tamus.edu*.
3. University of Connecticut, Department of Marines Sciences, MYSound; see *http://www.mysound.uconn.edu*.
4. Center for Coastal and Land-Margin Research, CORIE; see *http://www.ccalmr.ogi.edu/CORIE*.
5. Chesapeake Bay Monitoring; see *http://www.dnr.state.md.us/Bay/monitoring*.
6. Albemarle-Pamlico National Estuary Program; see *http://www.apnep.org*.
7. Plum Island Ecosystems Long Term Ecological Research (PIE-LTER); see *http://ecosystems.mbl.edu/PIE/*.
8. RiverNet; see *http://rivernet.ncsu.edu*.
9. E. T. Casper, S. S. Patterson, P. Bhanushali, A. Farmer, M. Smith, D. P. Fries, and J. H. Paul, "A Handheld NASBA Analyzer for the Field Detection and Quantification of *Karenia brevis*," *Harmful Algae* **6**, 112–118 (2007).
10. StriperTracker.org; see *http://www.stripertracker.org/*.
11. M. S. Islam, J. Bonner, T. Ojo, and C. Page, "Integrating Observation with Numerical Modeling for the Exploration of Hypoxia," *Environmental Science and Technology Conference*, Vol. 2, S. Starrett, J. Hong, and W. Lyon, Eds, American Science Press, 2006, pp. 436–440.
12. M. S. Islam, J. Bonner, T. Ojo, and C. Page, "Using Numerical Modeling and Direct Observation to Investigate Hypoxia in a Shallow Wind-Driven Bay," *Oceans '06 MTS/IEEE Boston Technical Program*, September 18–24, 2006, pp. 1–5.
13. D. O. Popa, A. C. Sanderson, R. J. Komerska, S. S. Mupparapu, D. R. Blidberg, and S. G. Chappell, "Adaptive Sampling Algorithms for Multiple Autonomous Underwater Vehicles," *Proceedings of IEEE/OES AUV2004*, 2004, pp. 108–118.
14. "Special Issue on Service-Oriented Architecture," *IBM Syst. J.* **44**, No. 4 (2005, entire issue).
15. R. Ambrosio, A. Morrow, and N. Noecker, "e-Business Control Systems," *Proceedings of the International Conference on Computing, Communications and Control Technologies*, Vol. 4, 2004, pp. 91–96.
16. P. Dube, N. Halim, K. Karenos, M. Kim, Z. Liu, S. Parthasarathy, D. Pendarakis, and H. Yang, "Harmony: Holistic Messaging Middleware for Event-Driven Systems," *IBM Syst. J.* **47**, No. 2, 281–287 (2008).
17. Open Source Data Turbine Initiative; see *http://dataturbine.org/*.
18. L. Amini, N. Jain, A. Sehgal, J. Silber, and O. Verscheure, "Adaptive Control of Extreme-Scale Stream Processing Systems," *IEEE ICDCS Conference*, Lisbon, Portugal, July 2006, pp. 71–77.
19. M. Branson, F. Douglis, B. Fawcett, Z. Liu, A. Riabov, and F. Ye, "CLASP: Collaborating, Autonomous Stream Processing Systems," *Eighth International Middleware Conference*, November 2007, pp. 348–367.
20. B. Gedik, H. Andrade, K.-L. Wu, P. S. Yu, and M. C. Doo, "SPADE: The System S Declarative Stream Processing Engine," *Proceedings of the ACM International Conference on Management of Data (ACM SIGMOD)*, June 2008, pp. 1123–1134.
21. J. L. Montgomery, T. C. Harmon, W. J. Kaiser, A. Sanderson, C. N. Haas, R. Hooper, B. Minsker, et al., "The WATERS Network: An Integrated Environmental Observatory Network for Water Research," *Environ. Sci. Technol.* **41**, No. 19, 6642–6647 (2007).

Received June 21, 2008; accepted for publication July 10, 2008

Harry R. Kolar *IBM Systems and Technology Group, 294 Route 100, Somers, New York 10589 (kolar@us.ibm.com).* Dr. Kolar received B.S. and M.S. degrees in physics and an interdisciplinary Ph.D. degree in the science and engineering of materials from Arizona State University. He joined IBM in 1983 from Arizona State University where he coordinated external user research efforts for the National Science Foundation John M. Cowley Center for High Resolution Electron Microscopy. He was also an independent scientific consultant for the Atlantic Richfield Company Advanced Materials Development Laboratory in Chatsworth, California. Dr. Kolar has worked in magnetoresistive and magneto-optical thin-film development at the IBM General Products Division laboratory in Tucson Arizona and has worked across several IBM divisions in technical, management, and executive roles to advance cross-industry application of new technologies. These include advanced analytical methods (involving signal processing, data mining, massively parallel processors, and very large databases), information and knowledge management, pervasive and embedded real-time intelligent systems, and sensor-based and cyberphysical systems. Dr. Kolar is currently the Chief Architect for Sensor-Based Solutions in the IBM Microelectronics Division and concentrates on environmental monitoring and management (Advanced Water Management Solutions) through the IBM Big Green Innovations initiative. He is an Adjunct Professor of Physics at Arizona State University and a District Advocate for the American Physical Society at the U.S. congressional level.

John Cronin *Beacon Institute for Rivers and Estuaries, 199 Main Street, Beacon, New York 12508 (jcronin@thebeaconinstitute.org).* Mr. Cronin is the Director and CEO of Beacon Institute for Rivers and Estuaries and has a long and distinguished career in environmental policy, advocacy, and education. He began his work on the Hudson River in 1973 and since then has worked as an advocate, a state and congressional legislative aide, and as director of the Pace Academy for the Environment at Pace University School of Law. For 17 years, he served as Hudson Riverkeeper, where his work in environmental policy and enforcement gained widespread prominence, and served as the inspiration for 170 similar programs around the world. He has been the subject of print, broadcast, and documentary profiles that have brought international attention to the story of the rebirth of the Hudson River. *The Wall Street Journal* called him "a unique presence on America's major waterways" and *Time* magazine honored him as a "Hero for the Planet."

Perry Hartswick *IBM Systems and Technology Group, 2070 Route 52, Hopewell Junction, New York 12533 (hartsp@us.ibm.com).* Mr. Hartswick is an IBM Distinguished Engineer and the Chief Engineer of the 300-mm semiconductor facility for the IBM Systems and Technology Group, where he was responsible for the strategic direction and architecture of the facility, which includes IT infrastructure, application platforms, and vendor direction, as well as commercialization of the products and improvement of schedules. He holds a B.Sc. degree in computer mathematics. He has recently become the delivery technology executive in the IBM Big Green Innovations effort. As the lead technologist, he is responsible for driving technology into areas of water management, solar power, and environmental monitoring. He provides technical input and direction to a number of external companies as a consulting architect, and he regularly provides industry direction to standards groups. He is a member of the IBM Technical Leadership Council and Technical Vitality Council, and he serves as the chief consulting engineer in several large semiconductor companies. Previously, he served as lead architect for the Advanced Semiconductor Technology Center (ASTC), where he was responsible for designing and developing the architecture required to create the ASTC. Mr. Hartswick joined IBM Poughkeepsie in 1984 as an electrical technician. During his career, he has held a number of technical and management positions in both Poughkeepsie and East Fishkill and has numerous patents.

Arthur C. Sanderson *Rensselaer Polytechnic Institute, Troy, New York 12180 (sandea@rpi.edu).* Dr. Sanderson is Professor in the Department of Electrical, Computer, and Systems Engineering at Rensselaer Polytechnic Institute. He received his B.S. degree from Brown University and his M.S. and Ph.D. degrees from Carnegie Mellon University. He held faculty positions at Carnegie Mellon University from 1973 to 1987, where he was co-director of the Robotics Institute. In 1987, he joined Rensselaer Polytechnic Institute as professor and department head of the Electrical, Computer, and Systems Engineering Department. He was instrumental in establishing the IEEE Robotics and Automation Society in 1989 and served as its president. He is a Fellow of the IEEE and AAAS. From 1998 to 2000, he was division director of electrical and communications systems research at the National Science Foundation in Arlington, Virginia. From 2000 to 2004, he served as vice president of research at Rensselaer Polytechnic Institute. Dr. Sanderson is author of more than 250 publications and proceedings and four books in the areas of biomedical signal processing, robotics and automation systems, sensor-based control, and computer vision. In 2005, he was a visiting distinguished professor in the College of Engineering, University of South Florida, and a visiting research fellow of the Autonomous Undersea Systems Institute, conducting research on underwater robotics and distributed sensor networks and their application to environmental observation and monitoring.

James S. Bonner *Clarkson University, Potsdam, New York 13699 (jbonner@clarkson.edu).* Dr. Bonner is the Director of the Center for the Environment at Clarkson University, where he is a member of the civil and environmental engineering faculty. He is also the REON Director of Research at Beacon Institute for Rivers and Estuaries. Dr. Bonner received his B.S. degree from the State University of New York–Plattsburgh, his M.S. degree from Clarkson College of Technology, and his Ph.D. degree from Clarkson University. Previously, he held a faculty position at Texas A&M University from 1986 to 2007 in the Zachry Department of Civil Engineering. While part of the Texas A&M University system, he was regional director for the Texas Engineering Experiment Station–Corpus Christi (1997–2005) and was the director of the Conrad Blucher Institute for Surveying and Science associated with the Texas A&M University–Corpus Christi (1997–2005). Dr. Bonner established the Shoreline Environmental Research Facility in Corpus Christi, Texas, in 1997. His research interests include particle contaminant transport, fate and effect modeling, environmental sensor development and deployment, and environmental hazard assessment and site remediation.

Liesl Hotaling *Beacon Institute for Rivers and Estuaries, 199 Main Street, Beacon, New York 12508 (lhotaling@thebeaconinstitute.org).* Ms. Hotaling is the Chief Education Officer for Beacon Institute for Rivers and Estuaries. She has a B.S. degree in marine science, an M.A.T. degree in science education, and an M.S. degree in maritime systems. She has worked as a field biologist, informal educator, and formal educator (K–14). Prior to joining the Beacon staff, she worked with the Center for Innovation in Engineering and Science Education at Stevens Institute of Technology, working to bring real-time data resources into classrooms throughout the world. She is a partner in Centers for Ocean Sciences Education Excellence–Networked Ocean World (COSEE-NOW) and specializes in educational projects supporting observing networks.

Ron F. Ambrosio *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (rfa@us.ibm.com).* Mr. Ambrosio received a B.S. degree in computer science with a mathematics emphasis from Hofstra University on Long Island, New York, and performed additional graduate study at Columbia University and New York Institute of Technology. He joined the IBM T. J. Watson Research Center in 1981 before graduating and has spent most of his career there. From 1995 through 1997, he was Director of Research and Development for Microware Systems, a real-time operating system developer in Des Moines, Iowa, and then returned to IBM Research. Mr. Ambrosio is the IBM Global Research Leader for Energy and Utilities and is a member of the U.S. Department of Energy GridWise Architecture Council working on future smart electric grids. He is also directly involved in IBM Energy and Environment initiatives, including frequent support of the IBM Big Green Innovations program. Mr. Ambrosio is the editor and workgroup leader for the ISO/IEC JTC 1/SC 25/WG 1 international standard 18012 on application interoperability in residential and building automation systems. His research interests include event-based programming environments, sensor and control system integration, real-time systems, and real-world/physical system interaction.

Zhen Liu *Nokia Research Center Beijing, 5 Dong Huan Zhong Road, BDA, Beijing 100176, China (zhen.38.liu@nokia.com).* Dr. Liu received a Ph.D. degree in computer science from the University of Orsay in Paris, France. He worked with France Telecom Research and Development as a research associate from 1986 to 1988. He joined INRIA (the French national research center on information and automation) in 1988, first as a researcher and then as a research director. He joined the IBM Thomas J. Watson Research Center in 2000 and was a senior manager before leaving for Nokia in 2008. Dr. Liu is currently the Head of Nokia Research Center Beijing. He is a fellow of the IEEE and has served on National Science Foundation panels and a number of conference program committees. He was the program cochair of the Joint Conference of ACM Sigmetrics and IFIP (International Federation for Information Processing) Performance 2004, and the general chair of ACM Sigmetrics conference in 2008. His current research interests are in distributed and networked systems, stream-processing systems, sensor networks, performance modeling, distributed optimization, and control.

Michael L. Passow *IBM Systems and Technology Group, 2070 Route 52, Hopewell Junction, New York 12533 (passow@us.ibm.com).* Dr. Passow is a Senior Technical Staff Member and the Manager of Advanced Process Control, sensor integration, and e-Diagnostics at the IBM 300-mm semiconductor facility in East Fishkill, New York. He received B.S. and M.S. degrees in nuclear engineering from the University of Oklahoma in 1984 and 1986 and a Ph.D. degree in nuclear engineering and plasma physics from the University of Michigan in 1991. Dr. Passow has worked at IBM since 1991 where he has been focused on semiconductor factory architecture, automation, sensors, and control. Currently, he leads a division-wide, cross-functional team responsible for the development, integration, and deployment of advanced semiconductor control solutions. Lessons learned and solutions developed in multiple facilities have been applied by the team to create a state-of-the-art control environment for the fully automated 300-mm manufacturing and development facility in East Fishkill. Dr. Passow is a technical engagement leader in similar automation and control projects in semiconductor and other industries, including pharmaceutical, water, and other "green" initiative projects worldwide.

Mark L. Reath *IBM Systems and Technology Group, 2070 Route 52, Hopewell Junction, New York 12533 (mreath@us.ibm.com)*. Mr. Reath received his B.S. and M.S. degrees in chemical engineering from Michigan State University in 1986 and 1989, respectively. He joined the IBM General Technology Division in 1989 where he worked on microcontamination detection and control at the IBM semiconductor manufacturing facility in Essex Junction, Vermont. Mr. Reath is the lead engineer responsible for remote sensor deployments within the IBM semiconductor manufacturing facilities in Essex Junction, Vermont, and Hopewell Junction, New York. His interests include real-time data collection and processing from residual gas analyzers, optical spectrometers, high-speed plasma characterization systems, and endpoint detection systems. He is a member of the American Institute of Chemical Engineers.