Integrated Operational Control of Unattended Distributed Coastal Sensor Web Systems With Mobile Autonomous Robots

Ashit Talukder, Anand V. Panangadan, Nickitas Georgas, Thomas Herrington, and Alan F. Blumberg, *Member, IEEE*

***Abstract*—Unattended autonomous systems of the future will involve groups of static and mobile sensors functioning in coordination to achieve overall task objectives. Such systems can be viewed as wirelessly networked unmanned heterogeneous sensor networks. We discuss a distributed heterogeneous sensing system with static sensors and mobile robots with novel adaptive control optimization algorithms for dynamic adaptation, coordinated control and end to end resource management of all sensors in response to detected events to achieve overall system goals and objectives. While our system design is applicable to a host of domains, it has been applied to and tested offline on an existing, functional maritime sensor web system, the New York Harbor Observation and Prediction System (NYHOPS) comprised of a host of maritime ocean and land sensors. Our goal is to enable adaptive control technologies to make the NYHOPS sensor web react faster and more effectively to threats or changing conditions, and to further maritime homeland security for the New York Harbor. Our contribution allows static sensors to work seamlessly with unmanned vehicles that can be deployed autonomously in response to detected events, and dynamically adjust operational parameters of static and mobile assets in the sensor web. Results for large area coastal monitoring are presented. Offline results using actual modeled data from *in situ* sensory measurements from the NYHOPS sensor web demonstrate how the sensor parameters can be adapted to maximize observability of a freshwater plume while ensuring that individual system components operate within their physical limitations.**

***Index Terms*—Adaptive control, autonomy, coastal sensor web, resource management, unmanned vehicles.**

I. Introduction

WHILE many fields in automated sensing and sensors have progressed based on rapid advances in sensing technologies and computing and information systems, the field of distributed and cooperative sensing over wide areas has been catalyzed in particular by new network and communications technologies. In the context of distributed sensing over large areas, the linkage of information systems by wired, wireless, or optical channels and protocols combined with advanced control and resource management solutions facilitates interactions that scale from a few (two to five) coordinating sensors to a large set (tens to hundreds) of distributed sensors combined with "swarms" of mobile robots demonstrating cooperative behaviors. Such systems, referred to as wireless sensor networks, have potential in a number of applications, ranging from defense systems for battlefield monitoring, to homeland security to science and commercial uses in environmental and habitat monitoring. We use the terms sensor networks and sensor webs interchangeably in this paper.

Sensor networks have been deployed and operational in a few applications over the past few years. Examples include the SEA-MONSTER system [1], the NIMS sensor network on the San Joaquin river system [2], and the New York Harbor Observation and Prediction system (NYHOPS) coastal sensor web [3]. The currently operational NYHOPS comprises of a network of sensors to monitor coastal and ocean parameters in the densely populated regions of the Hudson-Raritan Estuary and the New Jersey Atlantic Ocean shoreline [3]. The readings from the sensors are provided to a predictive model of the environment, the ECOMSED/POM model [4]. The modeled area covered by NYHOPS is shown in Fig. 1. The ECOMSED model describes the physical properties of the entire water mass in the NY/NJ harbor area using a set of differential equations. This model uses the sensor readings as input and outputs predictions of environmental parameters for the entire area. The NYHOPS coastal predictions are generated and released every 24 hours for a 48-hour forecast period and are used by local and federal agencies including NOAA, emergency first responder units along the east coast of the United States, Department of Homeland Security, local Departments of Health, harbor boat captains, and individuals that engage in recreational boating.

The NYHOPS model outputs are used to detect and predict events and regions of interest such as freshwater plumes, coastal storm surges, or man-made events. One limitation of the initial NYHOPS sensor web system was that all the sensor operational parameters were pre-determined and fixed; the sensor web sampled data at fixed, pre-set time intervals, and data communication occurred at pre-determined times. This limits the ability of the sensor web to reliably detect fast evolving events, and dynamically react to rapidly changing conditions. Ideally, as these

Manuscript received October 20, 2009; revised January 21, 2010; accepted March 22, 2010. Date of publication July 08, 2010; date of current version December 15, 2010. This work has been sponsored by the NASA Applied Information Systems Technology Program under AIST-QRS Award AIST-QRS-06-0017, and supported in part by the Office of Naval Research under ONR Grant Number N00014-08-1-0782. The research described in this work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

A. Talukder and A. V. Panangadan are with the Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109 USA (e-mail: ashit.talukder@jpl.nasa.gov).

N. Georgas, T. Herrington, and A. Blumberg are with the Department of Civil, Environmental, and Ocean Engineering, Stevens Institute of Technology, Hoboken, NJ 07030 USA.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSTARS.2010.2050463

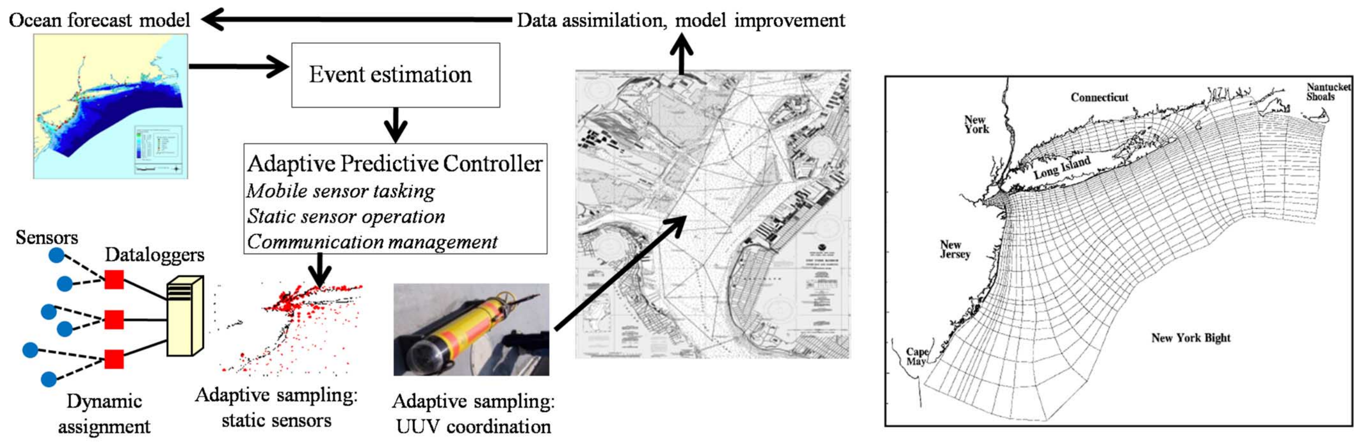

(a) (b)

Fig. 1. (Left): MPC-based control framework within the NYHOPS system. (Right) Output area of the ECOMSED hydrodynamic model.

events occur, the sensing system should be able to study the affected regions in greater detail by increasing the sampling resolution of sensors that are in that area. Thus, the goal of an *adaptive* sensor network control protocol is to dynamically regulate the operation of the sensors so as to maximize the measurement accuracy at all points in areas of interest while ensuring that the physical limitations of the sensor network are not exceeded.

We introduce a general mathematical control framework which will enable the resources of a sensor network to be regulated to guarantee a certain system lifetime while continuing to capture and measure information according to the needs of the application. As on-board computation in general uses significantly lower *energy* resources compared to wireless communication, it is beneficial to implement a sophisticated controller if it brings about significant savings in the amount of energy expended. A model-based controller is the most appropriate mechanism for this situation since the internal resources of the sensor system can be parameterized in advance and the expected changes in the environment are usually known for a given application. Model Predictive Control (MPC) is an established technique for controlling complex continuous systems [5]. MPC assumes that a model that describes how the system state responds to control inputs is available. At each control iteration, the values of the controlled inputs are obtained by solving an optimization problem that utilizes this state model. Limits on the range of the control, and other domain-specific requirements can be specified naturally as equality and inequality constraints in the optimization step. The relationship between the optimal MPC-based control framework and the modules in the NYHOPS system is shown in Fig. 1. This flexibility in problem specification and the ability to derive optimal control are some of the chief advantages of this control technique. Successful applications of MPC include chemical process control and resource management in the battlefield [6]–[9]. We believe that this approach is sufficiently flexible that it has the potential of assimilating both high-resolution *in situ* measurements and large scale remote-sensed data from controllable sensors into a predictive model of the physical environment. We have demonstrated in prior work [10], [11] that the MPC improves the lifetime and efficiency of a mobile sensor network system for health monitoring, where the data is temporal in nature; in contrast the coastal monitoring sensor network is spatiotemporal and involves larger number of nodes.

In addition, the sensors that are deployed along the coast or in the NY/NJ harbor have to transmit their measurements to a central data acquisition/control computer. We simulate the use of *dataloggers* as intermediate relay stations situated between the sensors and the central computer. The bandwidth across the network can be optimized by intelligently assigning sensors to dataloggers. If two high throughput sensors need to transmit simultaneously, they should be assigned to different dataloggers. In other instances, the assignment should be such that total energy cost of transmission is reduced. We demonstrate how this assignment can be computed within the MPC framework.

We first describe related work on adaptive control of sensor webs. We then introduce the NYHOPS coastal monitoring network. We next describe the model predictive controller framework and how this framework can be applied to the NYHOPS coastal monitoring network. We then present results from simulated modeled data obtained from actual observations. Though the adaptive MPC control method has not been implemented on the NYHOPS sensor network, these results from model outputs for actual historical dynamic coastal events indicate that the method is likely to improve the accuracy of system predictions and enable faster responses to critical transient events as compared to assimilating sensor data at a fixed rate.

II. Related Work

Adaptive sampling has been applied to minimize energy usage in a sensor network [12]–[15]. These approaches define a stochastic model of the environment and minimize the resources used to sample from this event model. These works model the environment but not the dynamics of the sensing system itself. Thus, though methods exist to explicitly control resource utilization in specific sensor networks, there is no general framework for optimal resource management that integrates both the application-specific information demands from the environment and the physical resource constraints imposed by the network. In our approach, we assume that sensors can be directly controlled by a central model predictive controller.

However, distributed versions of MPC [16]–[18] have been studied if our approach has to be extended to a completely distributed scenario.

Optimization of sensing resources is studied by the wireless sensor network community as this is critical to ensure that wireless embedded systems remain in operation long enough to be useful in real-world applications. For example, Krause *et al.* [19] describe a method of selecting sensor placement locations in a water distribution network for optimal detection of contaminants. In our target applications, all the sensors are expected to be operational and it is their relative sampling rates that are adapted.

Data from a network of sensors is increasingly being directly assimilated into a computational model of the environment to improve forecast accuracy. Haupt *et al.* [20] use a genetic algorithm to assimilate sensor data into a dispersion model of air pollutants. Gadgil *et al.* [21] describe a network of sensors to track the release of airborne pollutants. Dolan *et al.* [22] and Howe *et al.* [23] describe systems that coordinates multiple mobile sensors to study the ocean surface and sub-surface. Researchers have also assimilated data from remote-sensing satellites into model outputs [24].

III. The New York Harbor Observing and Prediction Ocean Sensor Network System (NYHOPS)

NYHOPS is comprised of a network of sensors and a model of the ocean environment to monitor and predict coastal and ocean conditions in the densely populated regions of the Hudson-Raritan Estuary and the extended New York City metropolitan area [3]. The readings from the sensors are provided to the model of the environment, the ECOMSED/POM model [6]. ECOMSED is a hydrodynamic model that describes the physical properties of the entire water mass in the NY/NJ harbor area using a set of differential equations (representing conservation of mass and momentum, and heat and salt transfer). The inputs to the model are ocean elevation (which depends on tides, offshore weather, cross-shelf elevation change), salinity and temperature at the open and coastal boundary of the model, and weather (air temperature, humidity, pressure, wind speed, solar radiation, cloud cover obtained from NOAA weather stations and forecasts). The model outputs include elevation, salinity, temperature, and water velocity. The model predictions are calculated over a high resolution orthogonal but curvilinear three-dimensional grid. The resolution is highest in the inland water bodies and decreases toward the open ocean. The model is run daily and the predictions (along with hindcasts) are displayed as images on a webpage.

Boundary conditions for the model are available from ground-based sensors and weather from NOAA. Further, unmanned underwater vehicles (UUVs) may be deployed to obtain additional information. Sensors may also be attached to passing ships to collect oceanographic parameters along the paths of these vessels referred to as mobile sensors.

In the NYHOPS system, sensors that are deployed along the coast or in the NY/NJ Harbor have to transmit their measurements to a central data acquisition/control computer. The distant location of the coastal sensors with respect to this central computer requires that *remote dataloggers* act as an intermediate relay station. A datalogger (which is a PC) compresses data files, establishes a connection to the Internet via a local ISP, and pushes the data to the data acquisition server. Data transmission to local sensors located in the harbor is through a line-of-sight serial radio modem system. A sensor can establish a 1200 baud, two-way simplex communication link with any of the remote dataloggers. Mobile sensors utilize serial cellular modems for data transmission to the remote logger. Currently, the data collection schedule is adjusted based on the power source and sampling requirements of the platform. Sites that are on the power grid can collect and transmit data at a high frequency. Typical sampling schedules consist of the measurement of up to 20 parameters that are retrieved every 5 minutes. Sensors that do not have access to the power grid measure an average of 10 parameters every 15 minutes and data is retrieved every hour.

IV. Model Predictive Control for Resource Management

We have designed a MPC controller that takes into account the utility of operating the sensors and the limitations of the NYHOPS network. Optimal controls are generated for (outputs of the MPC controller):

1) sampling rates of fixed sensors;
2) positions of mobile sensors such as UUVs;
3) assignment of sensors to dataloggers.

The constraints in the optimization problem represent:

1) maximum sensor sampling rates;
2) physical limitations of the UUVs (such as speed, and energy capacity);
3) bandwidth of the wireless communication network.

A. Mathematical Formulation of the Objective Function and System Constraints

The function to be optimized represents the sensing error and the energy required for data transfer. The constraints represent the physical limitations of the system such as bandwidth of the wireless networks. Let S represent the set of sensors and let p_s, $s \in S$, represent the position of a sensor s. We assume that the errors in a sensor s's measurements, are random, zero-mean Gaussian with variance σ_s^2. Let the number of independent measurements that are averaged to generate one sample (the sensor sampling rate) be u_s. Then, the variance of error of the averaged samples is σ_s^2/u_s.

We assume that the sensor's measurements at location p_s are correlated with values that would be obtained at location p and that this correlation decreases linearly with the distance $d(p, p_s)$ between p and p_s (this is the case when the region around a sensor is wide-sense stationary for short durations). This is illustrated in Fig. 2(a). The variance of error at location p when sensor s is used for measurements is $(\sigma_s^2/u_s) + d(p, p_s)\sigma_D^2$ where σ_D is a constant of proportionality that quantifies the rate at which the correlation decreases with distance. We selected a value for σ_D such that all points in the modeled region would have more than one sensor that is significantly correlated to the true value. Ideally, this value should be set based on the true response of the sensor. We use the theory of Kalman filters to optimally fuse multiple correlated sensor measurements. Then

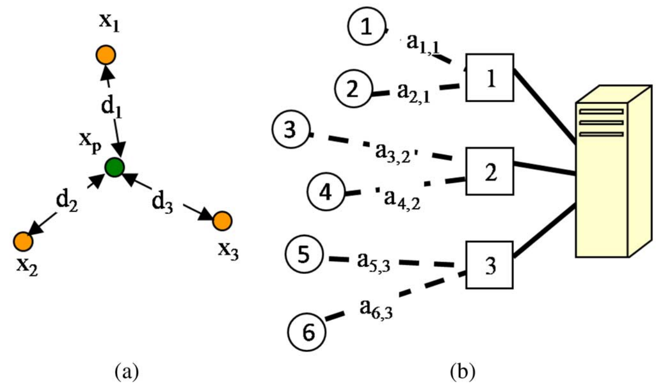

Fig. 2. (a) Three sensors correlated to point x_p. (b) Assignment of $\mathrm{S} = 6$ sensors to $\mathrm{M} = 3$ dataloggers.

the expected variance of error, σ_p^2, error at a location p, is given by

$$\frac{1}{\sigma_p^2(\mathbf{u})} = \sum_{s \in S} \frac{1}{\frac{\sigma_s^2}{u_s} + d(p_s, p)\sigma_D^2} \tag{1}$$

where $\mathbf{u} = (u_s)$, $s \in S$ is the vector of sensor sampling rates. Note that fusion using a Kalman filter can also include the case when sensors have non-zero cross-correlation (in (1), the cross-correlations were assumed to be zero). Equation (1) gives the expected variance at one point. We next extend the definition to include the case when the region of interest has a spatial extent. Let R denote the entire expanse of the region being modeled. As a measure of the overall accuracy of the sensor network for a particular set of sensing rates, we will use the average of the inverse variances of the expected error over all points in R:

$$f(\mathbf{u}, R) = \frac{1}{|R|} \sum_{p \in R} \frac{1}{\sigma_p^2(\mathbf{u})}. \tag{2}$$

The extent of the region can change over time. For instance, the region may represent a freshwater plume that changes shape and location, assumed predictable by the ECOMSED/POM model. The MPC framework can use a predictive model of the system to optimize control not just for a single control step, but also for a finite number of control steps in the future. Denote by $R_M(t)$ the extent of the region at time t as predicted by the spatial event detection model, M. A new objective function can then be defined that sums the objective functions for every control step.

$$f^T\left(u(t+1), u(t+2), \ldots, u(t+T)\right) = \sum_{t'=t+1}^{t+T} f\left(\mathbf{u}(t'), R_M(t')\right). \tag{3}$$

Here, t denotes the current time and T the number of future control steps to be optimized. Note that in this formulation, no distinction is made between the sampling rates of static and mobile sensors. When the paths of mobile sensors have to be optimized, the sequence of positions of the mobile sensor is treated as a time-varying independent variable in the optimization step.

We now present the formulation that models the limitations of the data links between dataloggers and sensors. When this is incorporated into an optimization framework, then the data transfer schedule that maximizes throughput can be determined. Let b_s denote the total data generated by sensor s (b_s is an integer multiple of its sampling rate u_s). Let there be M dataloggers. Each sensor can establish a connection with exactly one datalogger at a time. Denote by binary variable $a_{s,m}$ if sensor s is connected to the mth datalogger, i.e., $a_{s,m} = 1$ if they are connected and $a_{s,m} = 0$ if they are not. This is illustrated in Fig. 2(b). Let D_m denote the bandwidth of the data link to the mth datalogger (for example, 1200 baud). Then, the network limitations can be expressed by the following constraints:

$$\begin{aligned} &\sum_{s \in S} b_s a_{s,m} \le D_m, \quad \forall m = 1, 2, \ldots, M \\ &\sum_{m=1}^{M} a_{s,m} = 1, \qquad \forall s \in S \\ &\qquad a_{s,m} \in \{0, 1\},\ \forall m = 1, 2, \ldots, M,\ \forall s \in S. \end{aligned} \tag{4}$$

In practice, some of the $a_{s,m}$ will be preset to 0 indicating that it is physically impossible to establish a communication link between sensor s and datalogger m. The above constraints lead to an integer programming formulation as the $a_{s,m}$ are integers. A given assignment of sensors to dataloggers determines the energy expended while transmitting data over wireless links from the sensors to the dataloggers. We assume that the energy expenditure is proportional to the volume of data (the sampling rate) and the square of the distance between the transmitter and receiver (an assumption that is valid for radio transmission). Then the total energy expenditure is proportional to the function f_W:

$$f_W(\mathbf{u}, \mathbf{a}) = \sum_{m=1}^{M} \sum_{s \in S} u_s a_{s,m} d^2(s, m) \tag{5}$$

with $d(s, m)$ here being the distance between sensor s and datalogger m.

A single objective function cannot model all the system components that are to be optimized. Hence, the optimal control is obtained as the solution to a series of objective functions (*multi-objective optimization*):

1) maximize the utility of measurements from fixed sensors;
2) maximize the utility of measurements from mobile sensors;
3) minimize the energy expended in moving mobile sensors;
4) minimize the energy expended in wireless data transmission.

We integrate the components related to dynamic regions and constraints imposed by remote dataloggers. The objectives to be minimized are the sampling variance ($f_M = f^{-1}$) and the energy expended during wireless data transfer (f_W). We use the lexicographic approach to multi-objective optimization [25]. The lexicographic approach optimizes the functions in decreasing order of importance. For the NYHOPS system, we presume it is more important to minimize the sensor sampling error than to minimize the energy expenditure. In this case, we can solve the optimization problem in the following two steps.

$$(\mathbf{u}_1^*, \mathbf{a}_1^*) = \underset{\mathbf{u}\in U^{|S|}}{\arg\min}\,[f_M(\mathbf{u}, R_M)]$$
$$u_s > 1, \quad \forall s \in S; \sum_{m=1}^{M} a_{s,m} = 1, \quad \forall s \in S$$
$$\sum_{s\in S} b_s a_{s,m} \le D_m, \quad \forall m = 1, 2, \ldots, M. \tag{6}$$

The second step in the multi-objective optimization is

$$\mathbf{a}_2^* = \underset{\mathbf{u}\in U^{|S|}}{\arg\min}\,[f_W(\mathbf{u}_1^*, \mathbf{a})]$$
$$\sum_{m=1}^{M} a_{s,m} = 1, \quad \forall s \in S;$$
$$\sum_{s\in S} b_s a_{s,m} \le D_m, \quad \forall m = 1, 2, \ldots, M. \tag{7}$$

The optimal control parameters are then obtained as $\mathbf{u}_1^*$ (sampling rates) and $\mathbf{a}_2^*$ (assignment of sensors to dataloggers). The actual transmission schedule [such as the transmission slots in a Time Division Multiple Access (TDMA) scheme] can be calculated analytically after the datalogger assignment is completed.

This optimization problem is computationally difficult to solve as it is nonlinear and non-convex, with the presence of local minima. We therefore implemented the following approximations to enable real-time operation. We used the gradient descent algorithm as a fast solver for the optimization problems. However, this does not guarantee a globally optimal solution. The parameters of 50 static sensors could be optimized in less than 10 seconds using a PC with a 2.8GHz CPU. In order to solve for larger number of sensors, we clustered sensors into groups to reduce the problem size (every sensor in a group would then be assigned the same operational parameters). We used the k-means clustering algorithm for this step. This algorithm requires the positions of the initial clusters to be specified. We chose the initial clusters from sensor locations so as to maximize inter-cluster distance. Fig. 3 shows the results of the clustering step. Clustering to reduce the problem size of the assignment problem has been suggested for multiple target tracking [26] and stochastic matching [27] and found to improve performance in large problem sizes by reducing computational time without compromising quality of results.

In addition, the datalogger assignment involves mixed integer optimization. The assignment problem ("bin packing") is known to be NP-complete and hence algorithms to determine optimal assignment scale exponentially. Hence, we use a suboptimal "greedy" assignment of sensors to dataloggers—assign a sensor to nearest datalogger if bandwidth is available; else repeat with next farthest datalogger. This process is suboptimal since it does not consider the total energy expenditure (a global measure).

V. Results

True ground truth sensor observations are not available at the required resolutions throughout the modeled area. Hence, we studied the end-to-end performance of the proposed approach to improving model prediction accuracy using model output as ground-truth. Sensor data was simulated by subsampling from

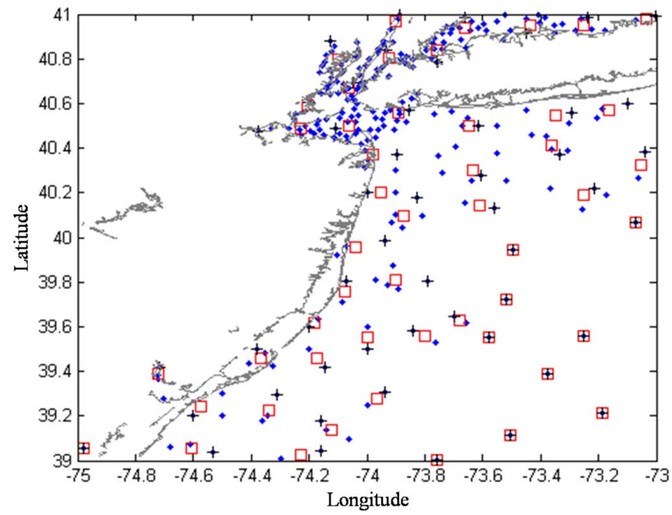

Fig. 3. Reducing the problem size by clustering sensors into groups using the k-means clustering algorithm. Blue dots represent sensor locations, + indicates initial clusters, and red squares are final clusters after convergence.

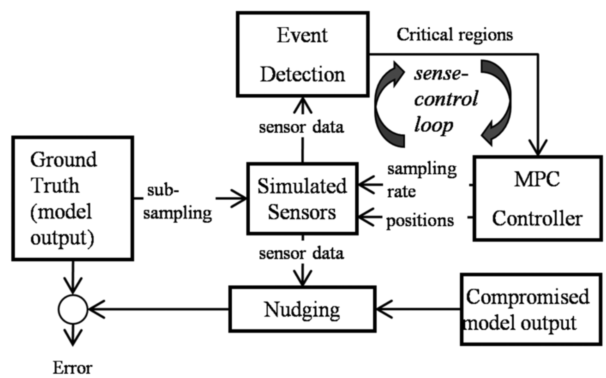

Fig. 4. Framework for evaluating the adaptive controller using simulated model data.

high resolution ECOMSED model output (NYHOPS forecasts and hindcasts). The base model predictions were obtained by intentionally "compromising" the ECOMSED model, i.e., certain model parameters were set to historic values instead of real-time observations in order to increase the divergence from true observations.

We evaluate our system by computing the root mean square error (RMSE) between the base (compromised) model outputs to which the sensor measurements have been assimilated and the uncompromised model (ground-truth) outputs. The squared error of a model output variable is calculated at every point of the model grid. The mean of the squared errors is computed over all grid points. The evaluation framework is shown in Fig. 4.

The locations of local universities, and a few coastal spots were chosen as positions of dataloggers for this simulation study. A few of the dataloggers were also situated in the open ocean in order to keep the maximum data transmission distance between sensors and dataloggers within practical radio range. We assumed that UUVs could navigate freely in the Lower New York Bay and along the New Jersey coast. The paths of ships that regularly ply in the area was available and we simulated sensor measurements along these paths.

During the week of April 15, 2007 unusually heavy rainfall welled the rivers flowing into the Hudson-Raritan Estuary ("Tax day flood"). A freshwater plume formed in the NY/NJ Harbor

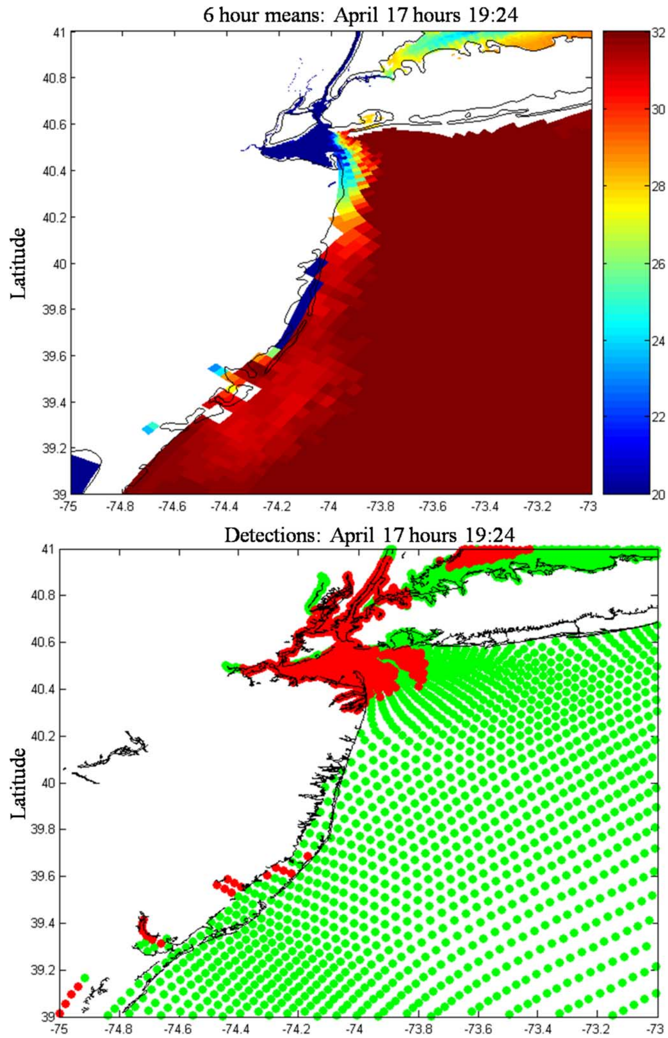

Fig. 5. Top: Sea surface salinity (ps units) observations (at high spatial resolution) at one time instant. Bottom: Locations of detected events.

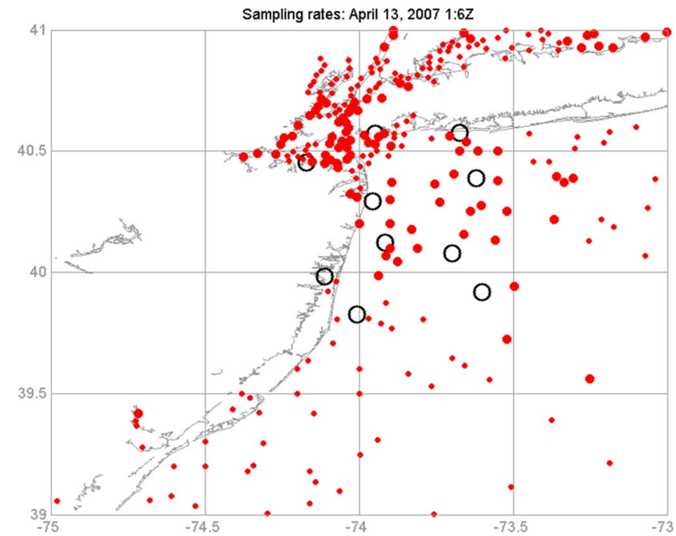

Fig. 6. Location of sensors (filled red/dark circles) and the detected critical region (unfilled circles). The size of the filled circles is proportional to the sampling rates.

and spread out along New Jersey's coastal waters. Such events may carry pathogens from overwhelmed metropolitan treatment plants into the coastal ocean, Atlantic Seashore beaches, and back bays. For this study, we set the inflow rates of freshwater in the ECOMSED model to their historic median rates. This compromised "base" ECOMSED model did not predict the freshwater plume but the simulated sensors observed the plume (as these were obtained by subsampling the uncompromised "ground-truth" ECOMSED model output from the NYHOPS forecasts and hindcasts).

We implemented an event detection algorithm to determine those (time varying) regions in the modeled area that would benefit the most from assimilation of sensor observations in indentifying the freshwater plume. To demarcate the extent of the plume we used the difference in surface salinity from historic expected levels. Sensors observing salinity values that were beyond one standard deviation from the historic mean were designated to be in the critical region (Fig. 5).

The number of simulated sensors in the sensor network is 250. The locations of the sensors were computed by clustering all the points in the model grid (using the k-means clustering algorithm). Note that the sensor locations are calculated independently of the event detection step. The locations of the simulated sensor nodes are shown in Fig. 6. ECOMSED model outputs were available at one hour intervals. Each control step was set to cover 6 hours, so that the MPC controller could determine an optimal (sub-) sampling rate from simulated ground-truth samples per sensor (later assimilated into the compromised base model). At the maximum sampling rate, the simulated sensor net data included all 6 hours of available data; at the minimum sampling rate only data from one hour of the available six hours was used. In order to directly compare the adaptive sampling method with uniform sampling, the total amount of data generated was held fixed: individual sensors sampled at optimized rates constrained by the total intra-sensor sum. The uniform sampling rate was this fixed number divided by the number of sensors. In our adaptive sampling system, this was accomplished by setting this total data limit in the datalogger constraint of the optimization problem. In Fig. 6, the size of each dot is proportional to that sensor's relative sampling rate as computed by the constrained multi-objective optimization step focusing on the event.

The decrease in RMSE of the sea surface salinity output variable using our adaptive sampling approach is compared with that achieved with uniform sampling in Fig. 7. The RMSE shown is the mean over one day (four control steps each assimilating six hours of data). The uncompromised ECOMSED predictions which correctly modeled the plume were used as ground truth. Fig. 7 compares the errors after assimilating surface salinity simulated sensor measurements into the ECOMSED model output. The sensor node sampling rates were computed using our MPC technique (adaptive sampling) and are compared to the equivalent fixed rate (uniform sampling). The surface salinity after assimilation matches the ground truth closely since the MPC dynamically adapted the sampling rates of the sensors that are close to the flooding event while reducing sampling rates of sensors at other locations to ensure that the power consumption or overall bandwidth used by the system does not overwhelm the sensor network and violate physical constraints of the system.

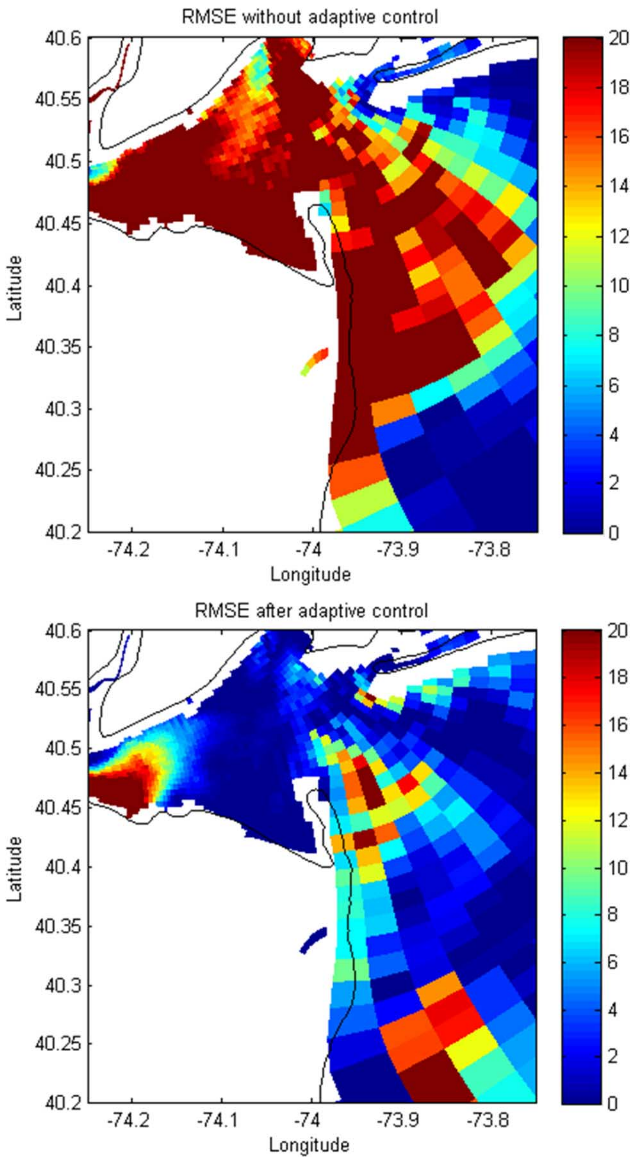

Fig. 7. Root mean square error (psu) between ground-truth and compromised model output of sea surface salinity with assimilations from sensor data on a moving coastal plume event on 4/26/2007. Top: Sensors without adaptive control. Bottom: Sampling rates are adapted using MPC. Bright red indicates high error rates as evidenced in fixed sampling, and blue indicates low error as obtained with adaptive sampling.

The effect of assimilating sensor data from four available UUVs (whose paths are optimized by our MPC controller) and three passing ships is shown in Fig. 8. The "model output" case shows the RMSE of the compromised ECOMSED prediction over the entire model domain. The other cases include assimilation of the different available sensor observations. As expected, increasing the number of sensor observations that are assimilated decreases the RMSE. However, the relative benefit of using a mobile asset depends on its location. For instance, one of the ships equipped with sensors (mobile sensor web node) moves in the Upper New York Bay, an area where there are few static sensors and no UUVs. Hence, assimilating data from this ship (mobile node) leads to a large decrease in RMSE.

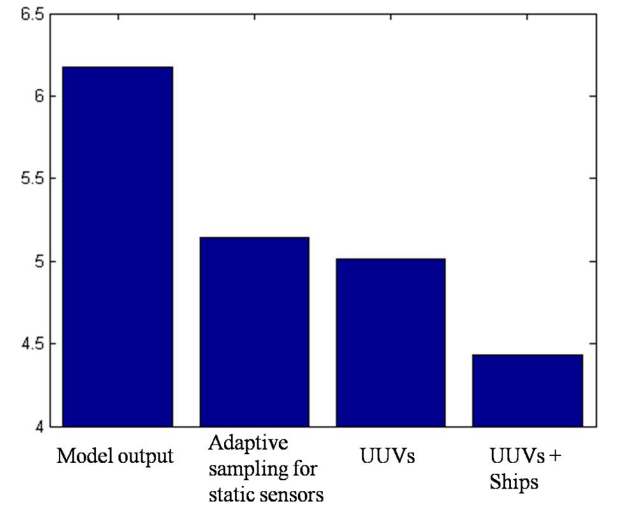

Fig. 8. Decrease in RMS error (psu) with assimilation of data from static sensors, four UUVs, and cruise ships.

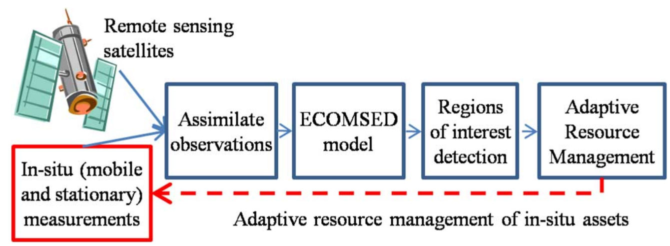

Fig. 9. Potential integration of remote-sensed and *in situ* measurements into an ocean prediction model. Solid lines indicate flow of data. Dashed lines indicate flow of control.

VI. Conclusion and Future Work

We demonstrated a framework for resource management in ocean observation sensor networks. The technique is general enough that it can produce control for static sensors, optimal paths of mobile sensors, and wireless transfer of sensor data. The control technique adapts these sensor and communication resources to changes in the ocean environment. The changes in the environment can be calculated from a variety of data sources. We demonstrated how static sensor measurements, and data from UUVs and passing ships can be assimilated into the outputs of a hydrodynamic ocean model to increase the prediction accuracy of the system forecasts.

This approach can also be generalized to accept data from remote sensing satellites and to dynamically generate observation tasks for satellites that could be re-tasked on demand. The concept is illustrated in Fig. 9. In particular, in the future, we intend to assimilate remote-sensed sea surface temperature (from ongoing AMSR, GOES missions and the Decadal Survey PATH mission) and ocean surface vector winds measurements (from the existing QuikSCAT mission and planned Decadal Survey XOVWM mission) into the ECOMSED ocean model along with *in situ* sensor measurements in order to improve prediction accuracy and provide more accurate estimates of coastal and maritime events. Satellites such as EO-1 could be re-tasked to provide better measurements and models of critical tracked events.

ACKNOWLEDGMENT

The authors thank Wei Jiang and Mitchell Kerman for their discussions and assistance with the event detection framework.

REFERENCES

[1] W. R. Otte, J. S. Kinnebrew, D. C. Schmidt, G. Biswas, and D. Suri, "Application of middleware and agent technologies to a representative sensor network," in *Earth Science Technology Conf.*, 2008, University of Maryland.

[2] W. J. Kaiser, G. J. Pottie, M. Srivastava, G. Sukhatme, J. Villasenor, and D. Estrin, "Networked Infomechanical Systems (NIMS) for ambient intelligence," in *Ambient Intelligence*, W. Weber, J. Rabaey, and E. Aarts, Eds. New York: Springer, 2005, pp. 83–113.

[3] T. O. Herrington, M. S. Bruno, and J. W. Rankin, "The New Jersey coastal monitoring network: A real-time coastal observation system," *J. Marine Environ. Eng.*, vol. 6, pp. 69–82, 2000.

[4] A. F. Blumberg, L. A. Khan, and J. P. St John, "Three dimensional hydrodynamic model of New York harbor region," *ASCE J. Hydraulic Eng.*, vol. 125, pp. 799–816, Aug. 1999.

[5] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, pp. 789–814, Jun. 2000.

[6] D. J. Sandoz, M. J. Desforges, B. Lennox, and P. R. Goulding, "Algorithms for industrial model predictive control," *Comput. Control Eng. J.*, vol. 11, pp. 125–134, Jun. 2000.

[7] V. Gopal and L. T. Biegler, "Large scale inequality constrained optimization and control," *IEEE Control Syst. Mag.*, vol. 18, pp. 59–68, 1998.

[8] J. Jelinek and D. Godbole, "Model predictive control of military operations," in *Proc. 39th IEEE Conf. Decision and Control*, Piscataway, NJ, 2000, p. 2562.

[9] J. E. Tierno, "Distributed autonomous control of concurrent combat tasks," in *Proc. American Control Conf.*, 2001, vol. 1, pp. 37–42.

[10] A. Talukder, A. Panangadan, L. Chandramouli, and S. Ali, "Optimal sensor scheduling and power management in sensor networks," presented at the SPIE Defense and Security Symp., Optical Pattern Recognition XVI, Invited Talk, Orlando, FL, 2005.

[11] M. Venugopal, K. E. Feuvrel, D. Mongin, S. Bambot, M. Faupel, A. Panangadan, A. Talukder, and R. Pidva, "Clinical evaluation of a novel interstitial fluid sensor system for remote continuous alcohol monitoring," *IEEE Sensors J.*, vol. 8, pp. 71–80, 2008.

[12] W.-P. Chen and L. Sha, "An energy-aware data-centric generic utility based approach in wireless sensor networks," presented at the 3rd ACM/IEEE Int. Symp. Information Processing in Sensor Networks (IPSN), Berkeley, CA, 2004.

[13] N. Sadagopan and B. Krishnamachari, "Maximizing data extraction in energy-limited sensor networks," presented at the IEEE INFOCOM, Hong Kong, 2004.

[14] R. Willett, A. Martin, and R. Nowak, "Backcasting: Adaptive sampling for sensor networks," presented at the 3rd ACM/IEEE Symp. Information Processing in Sensor Networks (IPSN), Berkeley, CA, 2004.

[15] B. Zhang, G. Sukhatme, and A. Requicha, "Adaptive sampling for marine microorganism monitoring," presented at the IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS), Sendai, Japan, 2004.

[16] E. Camponogara, D. Jia, B. H. Krogh, and S. Talukdar, "Distributed model predictive control," *IEEE Control Syst. Mag.*, pp. 44–52, 2002.

[17] W. B. Dunbar and R. M. Murray, "Model predictive control of coordinated multi-vehicle formations," presented at the 41st IEEE Conf. Decision and Control, Las Vegas, NV, 2002.

[18] A. N. Venkat, J. B. Rawlings, and S. J. Wright, "Stability and optimality of distributed model predictive control," presented at the CDC-ECC Joint Conf., Seville, Spain, 2005.

[19] A. Krause, J. Leskovec, C. Guestrin, J. VanBriesen, and C. Faloutsos, "Efficient sensor placement optimization for securing large water distribution networks," *J. Water Resources Planning Manag.*, vol. 134, pp. 516–526, 2008.

[20] S. E. Haupt, C. T. Allen, and G. S. Young, "A genetic algorithm method for sensor data assimilation and source characterization," in *Proc. 2006 Int. Joint Conf. Neural Networks*, Vancouver, BC, Canada, 2006, pp. 5096–5103.

[21] A. Gadgil, M. Sohn, and P. Sreedharan, "Rapid data assimilation in the indoor environment: theory and examples from real-time interpretation of indoor plumes of airborne chemicals," in *Air Pollution Modeling and Its Application XIX*. New York, NY: Springer, 2008, pp. 263–277.

[22] J. Dolan, G. Podnar, S. B. Stancliff, K. H. Low, A. Elfes, J. Higinbotham, J. Hosler, T. Moisan, and J. Moisan, "Cooperative aquatic sensing using the telesupervised adaptive ocean sensor fleet," in *SPIE Conf. Remote Sensing of the Ocean, Sea Ice, and Large Water Regions*, 2009, vol. 7473.

[23] B. M. Howe, N. Parrish, L. Tracy, A. Gray, Y. Chao, T. McGinnis, P. Arabshahi, and S. Roy, "A smart sensor web for ocean observation: Integrated acoustics, satellite networking, and predictive modeling," presented at the NASA Earth Science Technology Conf. (ESTC), Univ. Maryland, College Park, MD, 2008.

[24] K. Regner, M. Goodman, H. Conover, B. Zavodsky, M. Maskey, G. Jedlovec, X. Li, J. Lu, M. Botts, and G. Berthiau, "Using sensor web processes and protocols to assimilate satellite data into a forecast model," in *Proc. IEEE Int. Geoscience and Remote Sensing Symp. (IGARSS)*, Boston, MA, 2008.

[25] J. Ignizio and T. Cavalier, *Linear Programming*. Englewood Cliffs, NJ: Prentice-Hall, 1994.

[26] M. R. Chummun, T. Kirubarajan, K. R. Pattipati, and Y. Bar-Shalom, "Fast data association using multidimensional assignment with clustering," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 37, pp. 898–913, 2001.

[27] M. Hauskrecht and E. Upfal, "A clustering approach to solving large stochastic matching problems," in *Proc. 17th Conf. Uncertainty in Artificial Intelligence (UAI)*, J. S. Breese and D. Koller, Eds., Seattle, WA, 2001, pp. 219–226.

Ashit Talukder received the Ph.D. degree in electrical and computer engineering from Carnegie Mellon University, Pittsburgh, PA, in 1999.

He is currently a Senior Researcher (Level A) at the Jet Propulsion Laboratory/NASA, California Institute of Technology, Pasadena, California, and also holds a joint appointment as a research adjunct faculty member at the University of Southern California, and Senior Researcher at the Children's Hospital Los Angeles in Los Angeles, California. He is a principal investigator and technical lead on several NASA, DHS, NIH, NSF, DoD and JPL-funded projects. He has over 55 refereed journal and conference publications and has authored a book chapter. His research interests and expertise include pattern recognition, image and signal processing, machine learning, AI, controls, computer vision, optimization, and mathematics, applied to data mining, robotics, autonomous systems and sensors, sensor networks, biometrics, ATR, and human-machine interaction.

Dr. Talukder is on the organizing committee of the Annual SPIE Defense and Security Symposium. He was a Technical Program Committee member of the International Conference on Digital Telecommunications ICDT 2006 and chair of the First IEEE/ACM Workshop on Sensor Networks for Earth and Space Science (ESSA-2009) held at IEEE/ACM IPSN-2009. He has a patent for a biomedical data analysis and visualization system which was nominated for the NASA software of the year, and holds a provisional patent for a remote biometric identification system. He is a recipient of the NASA Space Award and a PACE award at Iowa State University. He has chaired several SPIE conference sessions and given five invited conference and keynote talks. He is a reviewer on several technical journals.

Anand V. Panangadan received the B.Tech. degree in computer science and engineering from the Indian Institute of Technology, Bombay, in 1996. He received the Ph.D. degree in computer science from the University of California, Los Angeles, in 2002. His Ph.D. dissertation describes a framework for building autonomous agents that can collectively construct arbitrary structures in a simulated environment.

He is currently a Postdoctoral Affiliate at the NASA Jet Propulsion Laboratory and a Research Specialist at the Childrens Hospital Los Angeles. He is developing algorithms for energy conservation in wireless sensor networks. These algorithms are to be deployed in remote human health monitoring systems and environmental sensor networks. He has also worked at the Robotics Research Lab of the University of Southern California. While there, he developed probabilistic models for describing human movement and region detection algorithms for sensor networks.

Nickitas Georgas received the B.S. degree in physics from the University of Crete, Greece, in 1996. He received the M.S. degree in marine environmental science from Stony Brook University, New York, in 2001. He is currently a Doctoral Student and Research Faculty at the Stevens Institute of Technology Center for Maritime Systems.

At CMS, he is responsible for the continuous development and daily forecasts of the New York Harbor Observation and Prediction System (NYHOPS: www.stevens.edu/maritimeforecast). His main interests are numerical modeling and forecasting, hydrodynamics, and water quality of inland and coastal waters. As a researcher and consultant he has participated in a variety of coastal engineering studies around the eastern seaboard of the United States, involving feasibility, impact evaluation, use-attainability, TMDL/WLA/LA, BMP, and dredging assessments, by designing and coupling hydrodynamic, sediment transport, and water quality models.

Thomas Herrington received the Ph.D. degree in ocean engineering from Stevens Institute of Technology in 1996.

He is currently an Associate Professor of ocean engineering and the Assistant Director for Ocean Observations and Forecasting at the Center for Maritime Systems at Stevens Institute of Technology. His research focuses on the improved understanding and prediction of coastal waves, tides and currents and their impact on coastal processes and wave attack on structures. To better understand these physical processes, Dr. Herrington and his colleagues have developed a real-time coastal observing and prediction system for the NY Bight, which is now an integral part of the U.S. Integrated Ocean Observing System.

Alan F. Blumberg is George Meade Bond Professor of Ocean Engineering and Director of The Center for Maritime Systems at Stevens Institute of Technology. Dr. Blumberg is recognized as one of the pioneer and leading experts in modern estuarine and coastal ocean prediction and has contributed significantly to the creation of integrated modeling and observing systems. Dr. Blumberg is the recipient of the 2001 American Society of Civil Engineers Karl Emil Hilgard Hydraulic Prize and the 2007 Denny Medal from the Institute of Marine Engineering, Science and Technology. He received a doctorate in ocean physics from The Johns Hopkins University and did postdoctoral work with Princeton University in their Geophysical Fluid Dynamics Program.