Autonomy and Sensor Webs: The Evolution of Mission Operations

Rob Sherwood[1]
Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena, CA 91109

The world is slowly evolving into a web of interconnected sensors. Innovations such as camera phones that upload directly to the internet, networked devices with built-in GPS chips, traffic sensors, and the wireless networks that connect these devices are transforming our society. Similar advances are occurring in science sensors at NASA. NASA-developed autonomy software has demonstrated the potential for space missions to use onboard decision-making to detect, analyze, and respond to science events. This software has also enabled NASA satellites to coordinate with other satellites and ground sensors to form an autonomous sensor web. A vision for NASA sensor webs for Earth science is to enable "on-demand sensing of a broad array of environmental and ecological phenomena across a wide range of spatial and temporal scales, from a heterogeneous suite of sensors both in-situ and in orbit."

Several technologies for improved autonomous science and sensor webs are being developed both inside and outside of NASA. Each of these technologies advances the state-of-the-art in sensor webs in different areas including enabling model interactions with sensor webs, smart autonomous sensors, and sensor web communications.

Demonstration of these sensor web capabilities will enable fast responding science campaigns that combine spaceborne, airborne, and ground assets. Sensor webs will also require new operations paradigms. These sensor webs will be operated directly by scientists using science goals to control their instruments. We will explore these new operations architectures through a study of existing sensor web prototypes.

I. Introduction

Technology today is transforming the world into an interconnected web of sensors. As an example, the modern cell phone has up to five different ways that it can communicate with other devices including Wi-Fi, Bluetooth, a GPS receiver, and voice and data communications. The range of other devices that can communicate with your cell phone is virtually unlimited and growing daily. All US passports issued today contain Radio-Frequency Identity (RFID) chips. It's just a matter of time before we start seeing personally targeted advertising such as in the 2002 movie "Minority Report." NASA has also been investigating the use of interconnected sensors to solve complex interdisciplinary science problems such as global climate change, earthquake forecasting, or long-term weather predictions. Many of these questions cannot be answered with the suite of sensors from a single satellite. They require data from multiple satellites, ground sensors, airborne sensors and models. The operation and coordination of a sensor web can be very different than operating a single spacecraft. This paper explores some of those operations implications as they relate to current research and trends in NASA sensor webs.

The operation of spacecraft has evolved considerably over the past decade with the inclusion of autonomy both on the ground and on the spacecraft. Starting with the Remote Agent Experiment in 1999, followed by the EO-1 Autonomous Sciencecraft (ASE) in 2004, autonomy has changed the way we perform mission planning, sequencing, and science discovery. These changes were enabled by higher performance computers that allow complex planning and reasoning software to migrate onboard the spacecraft. This onboard autonomy allows a spacecraft to be directed by goals rather than detailed commands and sequences where every step has to be defined in advance. This same autonomy is being used in ground sensor networks to increase coordination between sensor assets. These autonomous sensor networks, or sensor webs, rely on a new generation of operations tools and procedures that link the

[1] Program Manager, Earth Science Information Systems, 4800 Oak Grove Dr., Pasadena, CA 91109/MS 180-401, Senior Member AIAA.

scientist directly with the instruments/sensors. The current focus for NASA funded sensor web technologies is enabling model interactions with sensor webs, smart autonomous sensors, and sensor web communications.

Enabling model interactions in sensor webs is focused on the creation and management of new sensor web enabled information products. Specifically, the format of these data products and the sensor webs that use them must be standardized so that sensor web components can more easily communicate with each other. This standardization will allow new components such as models and simulations to be included within sensor webs.

Smart sensing implies sophistication in the sensors themselves. The goal of smart sensing is to enable autonomous event detection and reconfiguration. This may include onboard processing, self-healing sensors, and self-identifying sensors.

The goal of communication enhancements, especially session layer management, is to support dialog control for autonomous operations involving sensors and data processing and/or modeling entities. These technologies may include antenna for tracking dynamic sensors, autonomous networks and protocols that can distribute data communication tasks among the sensors and control the flow of data, transmission schemes that optimize bandwidth use, and distributed data storage devices.

II. What is a Sensor Web?

Depending on who you ask, you may get a different definition of the term sensor web. Most of those definitions use related concepts that are evolving into a broadly accepted lexicon. Other terms such as sensor network and system-of-systems have been used interchangeably with sensor web. Some examples of sensor webs include the seismic GPS network, the A-Train constellation of satellites, and the tsunami early warning system. Each of these examples have varying levels of integration, coordination, and autonomy. Recent publications have defined sensor webs as follows:

- A coherent set of distributed “nodes,” interconnected by a communications fabric, which collectively behave as a single, dynamically adaptive, observing system (Talabac, GSFC).
- An interconnected “web of sensors” that coordinates observations by spacecraft, airborne instruments and ground-based data-collecting stations. Instead of operating independently, these sensors collect data as a collaborative group, sharing information about an event as it unfolds over time (NASA press release).

Both of these are good definitions, but they don’t really capture the idea of feedback between sensors to capture further measurements. As defined at the February 2007 NASA ESTO/AIST workshop on sensor web technologies[1], a sensor web is:

> A coordinated observation infrastructure composed of a distributed collection of resources that can collectively behave as a single, autonomous, task-able, dynamically adaptive and reconfigurable observing system that provides raw and processed data, along with associated metadata, via a set of standards-based service-oriented interfaces.

Some key sensor web features include the ability to obtain targeted observations through dynamic tasking requests, the ability to incorporate feedback (e.g., forecasts) to adapt via autonomous operations and dynamic reconfiguration, and improved ease of access to data and information. Some key sensor web benefits include improved resource usage where selected sensors are reconfigured to support new science questions; improved ability to respond to rapidly evolving, transient phenomena via autonomous rapid reconfiguration, contributing to improved tracking accuracy; cost effectiveness which derives from the ability to assemble separate but collaborating sensors and data forecasting systems to meet a broad range of research and application needs; and improved data accuracy, e.g., through the ability to calibrate and compare distinct sensor results when viewing the same event.

III. Research Areas for Sensor Webs

Current research in sensor webs is focusing on multiple aspects of the coordinated sensing problem. One area of research is enabling model interactions in sensor webs. This area is focused on the creation and management of new sensor web enabled information products. Specifically, the format of these data products and the sensor webs that use them must be standardized so that sensor web components can more easily communicate with each other. This standardization will allow new components such as models and simulations to be included within sensor webs. Some of the research topics being addressed are:

- Interoperable data ingest as well as easy plug-and-play structure for scientific algorithms;
- Data from emerging grid and web common languages input such as the Open Geospatial Consortium (OGC) SensorML;
- Flexible hardware interfaces that can adapt to rapidly-changing data ingest protocols as well as ever-evolving algorithms;

- Connections to major spacecraft schedulers and task managers; and
- Semantic metadata to enable the transformation and exchange of data as well as data fusion.

Another research area in sensor webs is smart sensing. Smart sensing implies sophistication in the sensors themselves. The goal of smart sensing is to enable autonomous event detection and reconfiguration. Research areas include:

- Communication of the sensor with the system, including interfacing with certain system protocols and sensor addressability, in which sensors can identify themselves and interpret selective signals from the system, providing output only on demand.
- Diagnostics to inform the system of an impending failure or to signal that a failure has occurred, as well as self-healing sensors.
- On-board processing (up to and including science data products, as appropriate), self-describing sensor languages and actuation logic.

Another important area of sensor web research is communications technology. The goal of communication enhancements, especially session layer management, is to support dialog control for autonomous operations involving sensors and data processing and/or modeling entities. Specifically, research is being performed in the following areas:

- Adaptive and directive beam-forming antennas that can track the dynamic movement of sensor platforms;
- Autonomous networks and protocols that can distribute data communication tasks among the sensors and control the flow of data;
- Transmission schemes that maximize data throughput and provide optimum use of assigned bandwidth; and
- Distributed network of storage devices that can be accessed by any node in the sensor web with minimum latency.

The following sections highlight ongoing work in the NASA funded sensor web program to address the goals of earth observation sensor webs and the supporting technology development.

IV. Adaptive Sky Cloud Monitoring Sensor Web

The Adaptive Sky Cloud Monitoring Sensor Web Project[2] (ASky) developed software that enables multiple sensing assets to be dynamically combined into sensor webs. In this project, lead by Dr. Michael Burl of NASA's Jet Propulsion Laboratory, a feature correspondence toolbox was developed that includes methods for automatically relating the observations of one instrument at time t to the observations of another instrument at time t' (see Fig. 1). The toolbox contains both purely data-driven methods (e.g., feature-based image registration), as well as hybrid approaches that utilize metadata and ancillary information (e.g., sensor footprint collisions, open-loop georeferencing based on sensing geometry and imaging parameters).

Components from the Adaptive Sky toolbox were combined to enable a demonstration sensor web scenario in which multiple ash clouds generated by the eruption of the volcano Bezymianny on the Kamchatka Peninsula on October 14, 2007 were tracked using GOES (Geostationary Operational Environmental Satellite) brightness tem-

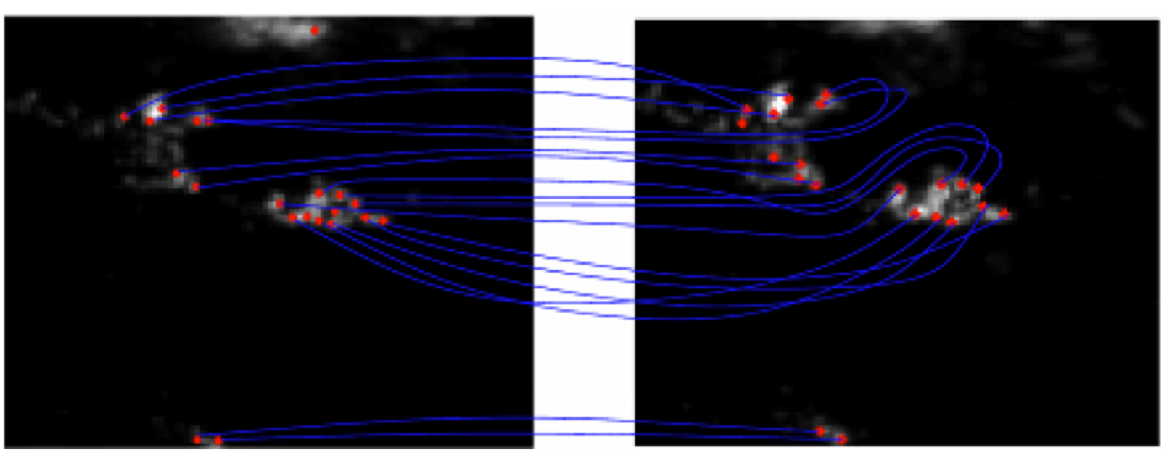

Figure 1. Two images taken approximately one minute apart by different cameras from the Multiangle Imaging SpectroRadiometer (MISR) instrument on the Terra satellite. Blue lines indicate corresponding features detected.

perature difference (BTD) images and reacquired using specialized instruments on both the Terra satellite and the A-train group. This sensor web scenario resulted in the first-ever definitive observation of a volcanic ash cloud from the CALIOP (Cloud-Aerosol Lidar with Orthogonal Polarization) lidar, as well as the first joint observations of a volcanic ash cloud by both the Multi-angle Imaging Spectroradiometer (MISR) and CALIOP. Without the feature correspondence capabilities provided by Adaptive Sky, these returns likely would have been attributed to cirrus clouds.

The operations implications of an ASky-like sensor web cut across multiple flight projects. The key operational capabilities that are required for this sensor web are data availability, data timeliness, data format standards, and data usage policies. In particular, the georeferenced data from multiple missions has to be made available relatively quickly after acquisition to scientists. The data must be in an easily usable and published format and with permission to combine it with other data products to build new science data products. ASky would also be more valuable if each mission had a published projection of its ephemeris data and pointing data (some number of days) into the future. For spacecraft that can be pointed, it would be helpful to have a published pointing range of motion for instrument observations, as well as any operational constraints on pointing. And lastly, it would be helpful for spacecraft to have a standardized and automated method for requesting observations rather than the proprietary and often manual methods used today. Of course this assumes that observation requests outside the mission/science team can be accommodated from time to time.

V. Semantically-Enabled Science Data Integration (SESDI)

The goal of the SESDI[3] project, which is lead by Dr. Peter Fox of NCAR in Boulder, CO, is to bring together diverse sets of scientific data from three very different scientific disciplines such that researchers can access the data quickly and transparently, without having to understand the details of the sources of the data and the peculiarities of each individual data source. This type of measurement-based, as opposed to instrument-based, approach has the potential to greatly speed up interdisciplinary research that would normally require collecting many kinds of data from many different sources, then figuring out how to meld the various datasets into something consistent with respect to location, time, units, definitions, etc. Rather than forcing scientists to be data set curators, SESDI provides a significant step away from that paradigm. SESDI can be used in areas such as climate research that involve data from ground-based, airborne, and spaceborne instruments studying the Earth's atmosphere, oceans, land surface and ice cover, as well as the Sun and the response of the Earth system to solar activity. SESDI will allow scientists to focus on creative new investigations of climate change, without spending undue resources on collecting and integrating disparate data sources.

Data discovery and access has traditionally been done using the specific terminology of those who have a deep understanding of the data sources (instruments or model output), those who have developed the databases, and those who are expert in a specific discipline. Moving from this basic syntactic (instrument-based) approach to a higher level (measurement-based) semantic approach will also make data and higher-level information more accessible to a wider group of people. At the same time, the semantic metadata will allow for a transformation and exchange of measurement data from diverse sensors within a sensor web. SESDI will also enable data fusion from diverse data sources within a sensor web, which will be important for understanding complex scientific phenomena. One example is large volcanic eruptions. Geologic databases provide the information about the seismic, geochemical and geophysical consequences of the eruption, and its impact on atmospheric chemistry and reflectance associated with particulate matter, requires integration of concepts that bridge terrestrial and atmospheric data sources (see Fig. 2).

The semantic web technologies in SESDI offer great promises in the ability to integrate multiple data sets. Operators may have to include the metadata within the processed and archived sensor data to enable this data integration but inference can also be used to fill in gaps in incomplete information thus reducing the need for additional ground processing unless the missing metadata cannot be inferred.

VI. Model-based Volcano Sensor Web

This project, lead by Dr. Ashley Davies at NASA's Jet Propulsion Laboratory, has created a fully-autonomous system that integrates multiple sensors (both on the ground and in space), advanced software on the Earth Observing 1 (EO-1) spacecraft that processes hyperspectral data to detect thermal emission from active volcanoes, other remote-sensing data processing systems, alerts from volcano monitoring organizations around the world, and models of volcanic processes. Alerts of volcanic activity trigger observation requests by EO-1. Data are rapidly processed to determine location of activity, thermal emission and lava effusion rate, all valuable information to volcanologists in the field. These products are distributed automatically.

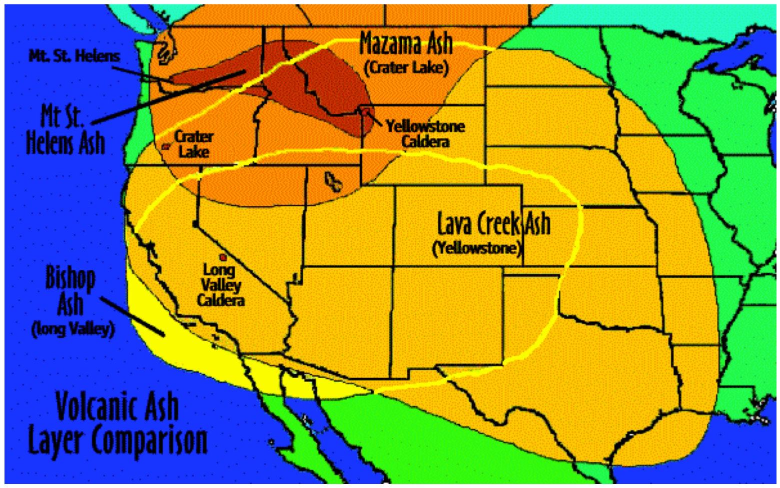

Figure 2. Compilation of distribution of volcanic ash associated with large eruptions. Note the continental scale ash fall associated with Yellowstone eruption ~600,000 years ago.

Assets, data products and other system components are described using Open Geospatial Consortium SensorML, and are linked with Web Services which allow additional assets to be incorporated into the system and for other users to access data and products. The Volcano Sensor Web proved its worth in late 2006 when it successfully delivered high-value products in a timely fashion to help mitigate effects of a volcanic eruption at Nyamulagira (Democratic Republic of the Congo).[4]

The operations for the volcano sensor web include collaboration between scientists and the EO-1 satellite operators. The scientists specify a campaign of volcano monitoring in the form of a goal set given to the EO-1 operations staff. This is an unprecedented method of operations where scientists issue goals that control the science instruments onboard EO-1. This requires a high level of trust, which was built up during early EO-1 autonomy testing. The volcano sensor web also requires close coordination with ground-based volcano sensor networks that act as triggers for satellite observations.

VII. Prototype Interoperable Sensor Web Architecture

The "Interoperable Sensor Architecture for Sensor Webs in Pursuit of GEOSS" task is prototyping a platform-independent message-based architecture to facilitate interoperability between a diverse set of sensors and science data processing algorithms. The vision is to provide users the capability to create "mash ups" (a web application that combines data from more than one source into an integrated experience), similar to that used by Google Earth users to create a composite map with overlays of information from other data sources such as weather, traffic, urban construction etc. (see Fig. 3). The team, lead by PI Dan Mandl (GSFC), created, integrated and is running a demonstration with four satellites, one Unmanned Aerial System, multiple ground sensors, and data algorithms and models. Operators can rapidly assemble customized workflows to produce science products to help manage wildfires. Whereas the present mash ups integrate a variety of data sources, this project's mash ups trigger workflows that actually task sensors via a common interface. Previously these sensors had unique interfaces. The basic components and features of the architecture are:

- Open Geospatial Consortium (OGC) compliant, platform-independent web service interfaces.
- Workflow engines.
- Decision support system.
- Self-describing sensor nodes, self-describing data processing nodes, and self-describing workflows that assemble customized data products.
- Discovery capability by wrapping the sensor nodes, the data processing nodes and the workflows in Internet news feeds which are then aggregated by Internet news aggregators such as Google burner and thus allow users to discover these capabilities using common terms.
- Use of open source software for all aspects of this architecture.

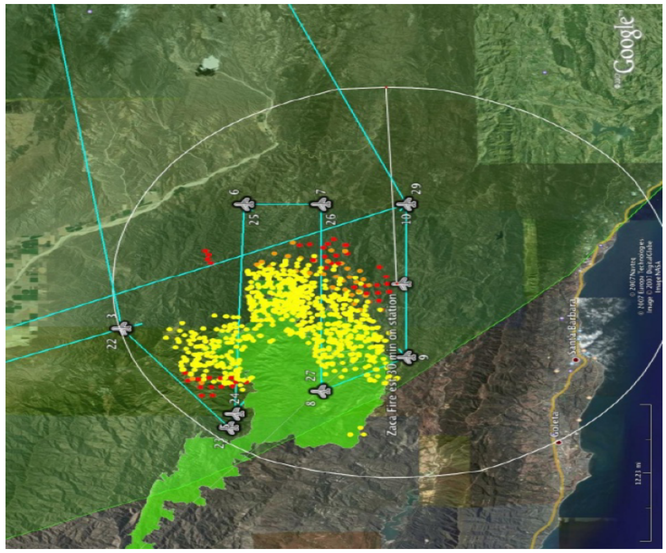

Figure 3. Data mashup using Google Earth with satellite data, UAV flight plans, and ground fire data.

The ultimate result will be to reduce the cost of obtaining customized science data products by an order of magnitude and decreasing the time to discover and receive customized science data products by as much as an order of magnitude.

This sensor web project is another example of shifting control of multiple sensors from operations staff to scientists. In this case, operators must make their sensors available as assignable assets through a standards-based interface. This level of science control is a radical departure from the current stovepipe operations paradigm, but is also necessary to answer many of the complex multi-variable problems in Earth science such as global climate change. Necessity may drive operations to this paradigm sooner rather than later.

VIII. Secure Autonomous Integrated Controller for Sensor Webs

This project, lead by William Ivancic of NASA Glenn Research Center, is developing and demonstrating architectures and protocols to enable secure, time-critical interaction between space and ground systems owned and controlled by various entities. The project has developed and ground demonstrated large file transfers over multiple terminals utilizing delay tolerant network (DTN) code and established a large consortium-based flight team with multiple-agency, industry, DoD, academic, and international involvement. The project will demonstrate an integrated sensor web response to an external trigger utilizing a spacecraft-based sensor provided by Surrey Satellite Technology Limited, ground stations provided by the Naval Research Laboratory, and a "virtual" mission operations center provided by General Dynamics.

The operations implications for a wide scale deployment of this architecture are the adoption of new network protocols for both space and ground sensor systems. In addition, these sensor systems operations team would have to make their sensors available for tasking from the larger sensor web.

IX. Change Detection On-Board Processor (CDOP)

One existing smart sensing project is the On-Board Processor for Direct Distribution of Change Detection Data Products.[5] CDOP, lead by Yunling Lou of NASA's Jet Propulsion Laboratory, is developing an autonomous disturbance detection and monitoring system for imaging radar that combines the unique capabilities of imaging radar with high throughput onboard processing technology and onboard automated response capability based on specific science algorithms.

Figure 4 contains a block diagram of CDOP. Raw data from the radar observation are routed to the onboard processor via a high-speed serial interface. The onboard processor will perform SAR image formation in real time on two raw data streams, which could be data of two different polarization combinations or data from two different interferometric channels. The onboard processor will generate real-time high-resolution imagery for both channels. The onboard processor will also execute calibration routines and science algorithms appropriate for the specific

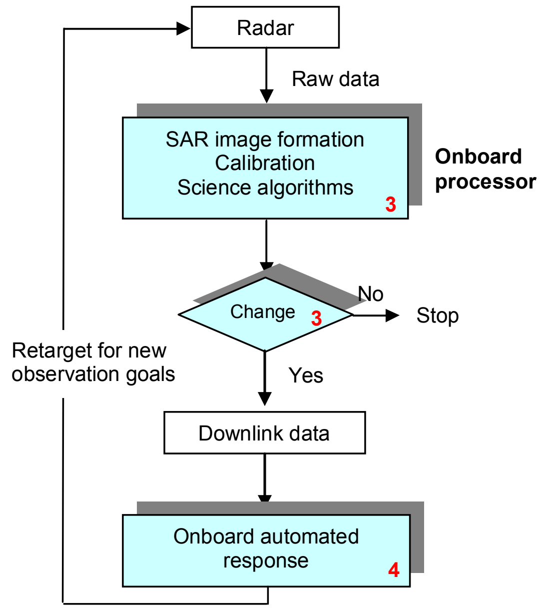

Figure 4. CDOP Block Diagram

radar application. Autonomous detection is performed by an intelligent software routine designed to detect specific disturbances based on the results of science processing. If no change is detected, the process stops and the results are logged. If "change" due to specific disturbances is detected, the onboard automated response software will plan new observations to continue monitoring the progression of the disturbance. The new observation plan is routed to the spacecraft or aircraft computer to retarget the platform for new radar observations.

The CDOP team is also developing interfaces to existing sensor webs to conduct autonomous observation of specific science events based on external triggers from other sensors in the sensor web. This smart sensor technology has the potential to provide key information for disaster and hazards management in the event of an earthquake, volcanic eruption, landslide, flood, and wildfire.

The operational implications of CDOP include a dramatically lower bandwidth for telemetering the processed radar data versus the raw radar data. Another result will be less ground-based science data processing. The ability to run the change detection algorithm requires a past radar image of the same area. This could require the uplink of radar images for change detection processing if no catalog exists on the satellite.

X. Sensor-Analysis-Model Interoperability Project

The Sensor-Analysis-Model Interoperability Technology Suite (SAMITS) project, lead by Stefan Falke of Northrop Grumman, is defining a framework for interoperability among sensors and models using service oriented architecture principles and Open Geospatial Consortium (OGC) standards. The project goal is to make it easier to use sensor data as model input and, conversely, allow model output to be used to direct where, when and what the sensors measure.

The team is partnering with the Forest Service in testing interoperability within wildfire smoke analysis and forecasting applications. Presently, locations of fires and smoke patterns detected by satellites and Unmanned Aerial Systems (UAS) are being served through standard interfaces and automatically pulled into analysis web services for comparison and validation with surface observations and smoke forecasts. Next steps include a more comprehensive system of connected and interacting sensors, models, and analysis services, including the use of forecasts in planning fire related measurements by satellites and UAS sensors.

For the recent Southern California wildfires, the team compiled smoke related resources, including the EO-1 and Ames UAS data. They are assessing the data, their accessibility and tasking, and addressing gaps to make the process more interoperable for future responses. This analysis and the integration of forecasting will create a system that will help the air quality planners (smoke forecasters and public health agencies) in the future.

The operations implications of the SAMITS project include coordination of spaceborne, airborne, and ground assets with multiple government agencies (USFS, NASA, FAA). The SAMITS project allows a manually operated group of sensors to function as a model driven sensor web. Without SAMITS, the EO-1 operators might receive a call or email to image a particular fire zone. SAMITS enables a more automated response with a direct standards-based connection between the ground sensors and the satellite.

XI. QuakeSim Project

The QuakeSim[6] project, lead by Dr. Andrea Donnellan at JPL, is applying a sensor web system to understand and study active tectonic and earthquake processes. Earthquake studies could be considered a classic case of a sensor web, with distributed seismic sensors that are coordinated in the study of a scientific process. But studying earthquakes is considerably more complex than just a network of seismic sensors. QuakeSim integrates both real-time and archival sensor data with high-performance computing applications for data mining and assimilation. The computing applications include finite element models of stress and strain, earthquake fault models, visualization, pattern recognizers, and Monte-Carlo earthquake simulations (see Fig. 5). The data sources include paleoseismic data, seismic sensors, GPS sensors for surface deformation, and spaceborne sensors such as interferometric synthetic aperture radar (InSAR).

The QuakeSim team is developing simulation and analysis tools to study the physics of earthquakes using state-of-the-art modeling, data manipulation, and pattern recognition technologies. This includes developing clearly defined accessible data formats and code protocols as inputs to the simulations. These codes must be adapted to high-performance computers because the solid Earth system is extremely complex and nonlinear, resulting in computationally intensive problems with millions of unknowns. Without these tools it will be impossible to construct the more complex models and simulations necessary to develop hazard assessment systems critical for reducing future losses from major earthquakes.

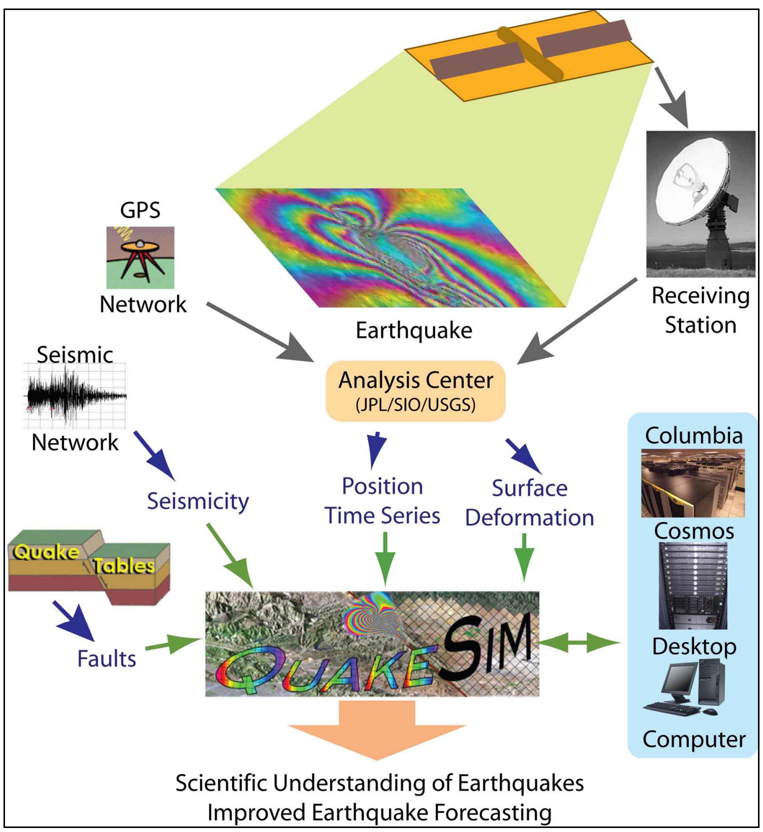

Figure 5. QuakeSim Architecture

The operations implications for using the QuakeSim architecture are still evolving. Operators for the upcoming Deformation, Ecosystem Structure and Dynamics of Ice (DESDynI) mission will have to interface with the QuakeSim architecture for post event earthquake analysis. This will require adhering to the data formats and protocols established within the architecture. Operators will have to work closely with earthquake scientists when using

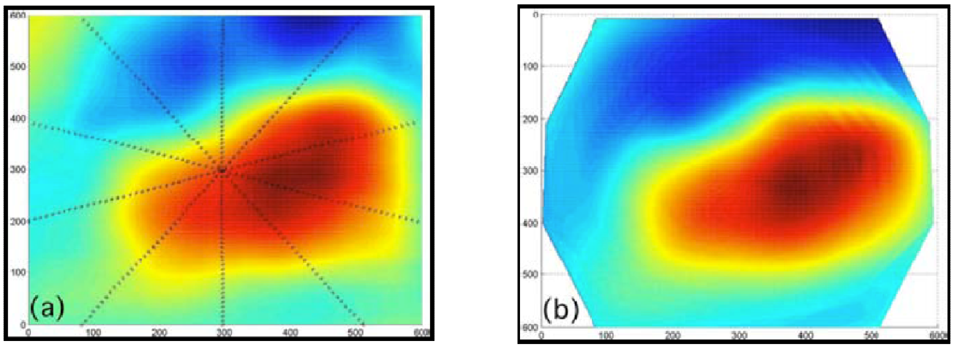

Figure 6. Example of a 2D field measured by a sensor web: (a) true field, and reconstructed field using (b) distributed wavelets.

QuakeSim to coordinate mission-planning activities. This coordination will also involve ground and possibly airborne assets (such as UAVSAR) that could be deployed after an event.

XII. ESCOMS Project

Another research project in sensor web communications is the Efficient Sensor Web Communication Strategies Based on Jointly Optimized Distributed Wavelet Transform and Routing (ESCOMS) project. The ESCOMS project, lead by Dr. Antonio Ortega of the University of Southern California, is developing algorithms for configuring a sensor network topology and for efficiently compressing the correlated measurements as data is shipped toward a central node, so as to minimize energy consumption while reproducing the underlying field as accurately as possible (see Fig. 6). This system enables the nodes to reconfigure the network automatically, taking into account variations in the node characteristics (node mobility, power consumption, addition of new sensors, and deletion of other sensors).

ESCOMS will implement advances in compression including entropy coding, filter optimization, path merging, joint compression and routing, and temporal coding. Advances in networking and routing will include techniques in node selection, network initialization, routing optimization, link quality robustness, inclusion of broadcast nodes, and automatic reconfigurability. These new capabilities are being tested in a lab and in a sensor web of about 100 nodes. Eventually they will be tested in an ocean monitoring application at the Great Barrier Reef in Australia for an extended period of time.

One ESCOMS scenario is the use of an in situ sensor web to monitor conditions and changes in the Antarctic ice shelf. Data collected from such a sensor web would be used in conjunction with a larger sensor web including airborne and spaceborne instruments. A second scenario that would benefit from ESCOMS is an in situ sensor web to monitor ecological conditions in a remote region such as a forest or a desert. ESCOMS is a great example of the use of autonomous networks and protocols, as well as optimized transmission schemes, which can be used in remote power-constrained Earth monitoring sensor webs.

ESCOMS is focused on reducing communications power and increasing bandwidth for ground-based sensor webs. The project is investigating multiple techniques to achieve these goals. The effectiveness of these techniques is dependent upon the format and quantity of sensor data and the network topology. Operators need to understand which techniques are most applicable to their sensor web in order to derive benefit from these algorithms. That choice could change as the sensor network evolves over time.

XIII. CARDS Project

The Autonomous In-situ Control and Resource Management in Distributed Heterogeneous Sensor Webs (CARDS) project is developing a general technique for resource management in heterogeneous sensor webs for environmental monitoring and other applications has been successfully completed. In this project, lead by Dr. Ashit Talukder of NASA's Jet Propulsion Laboratory, a novel Model Predictive Control (MPC) framework to adaptively control real-time resources of a variety of static and mobile sensor web assets with limited resources was developed and shown to improve the performance of the New York Harbor Prediction and Observation Sensor Web System (NYHOPS).

NYHOPS has been in operation for several years and comprises of a network of static and mobile sensors to monitor coastal and ocean parameters (salinity, temperature, water velocity, elevation) along the New Jersey Atlantic Ocean shoreline. The current NYHOPS sensor web operates in a fixed mode of operation, which is severely limiting for observing and modeling many kinds of spatiotemporal events such as plumes and storm surges.

CARDS uses an MPC control "model-based" solution for sensor web adaptation, where an optimal control policy is derived taking into accounts the system model and the physical constraints of the full system. The predictive capability of the MPC allows consideration of future states/events to make the current optimal control decision. The MPC adaptively controlled the full NYHOPS system, including the sampling rates of sensors (see Fig. 7), communication bandwidth, dynamic re-assignment of sensors to relay nodes, and optimal movement of mobile underwater unmanned vehicles in response to observed events (plumes and storm surges). Initial tests of the adaptive control technique have shown to improve the accuracy of NYHOPS models by at least 50%.

Further evaluation of CARDS sensor and UUVs control, and sensor web visualization tools in Google Earth is being conducted. CARDS is directly applicable to a variety of sensor webs and paves the way for coordinating multiple ground and space assets owned by NASA for faster and better detection, tracking and characterization of earth and extra-terrestrial environments.

The CARDS project has significant operations implications. The CARDS project transforms a largely manual sensor network into an automated model-driven sensor network. Operators will have to clearly understand the model inputs and the operation of the model predictive controller. Operators will be required to deploy autonomous mobile sensors when requested by the CARDS system.

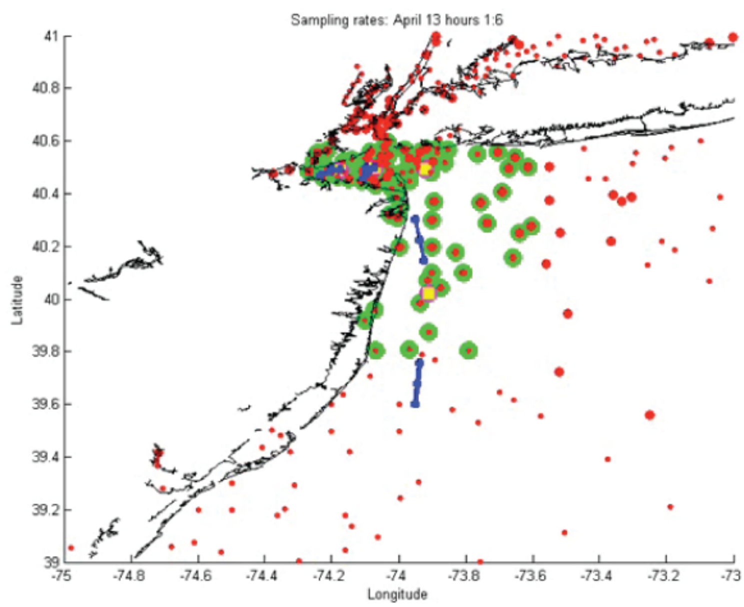

Figure 7. Based on model-predicted storm surge/plume event, the sampling rate is increased for some sensors. Size of green dots indicate additional sampling rate.

XIV. The EO-1 Sensor Web

One example of how onboard autonomy has enabled a sensor web is the Autonomous Sciencecraft (ASE) running on the EO-1 spacecraft. The ASE has demonstrated several integrated autonomy technologies to enable autonomous science. Several science algorithms including: onboard event detection, feature detection, and change detection are being used to analyze science data. These algorithms are used to downlink science data only on change, and detect features of scientific interest.

The use of automated planning onboard EO-1 has enabled a sensor web capability. The EO-1 satellite has been networked with other satellites and ground sensors to form an autonomous satellite observation response capability.[7] The EO-1 sensor web has been used to implement a global surveillance program of science phenomena including: volcanoes, flooding, cryosphere events, and atmospheric phenomena. Science agents for each of the science disciplines automatically acquire and process satellite and ground network data to track science phenomena of interest. These science agents publish their data automatically to the internet each in their own format. When the science agents discover a significant science event, the EO-1 satellite is automatically retasked to study the area of interest using its higher resolution instruments. Scientists can update the science agents based on specific scientific goals. In this case, the EO-1 satellite is being operated as a *goal-directed spacecraft (GS).* In a GS, the operator will specify what to accomplish rather than how to accomplish it.[8] Using a GS allows a complex system to select among alternatives for achieving a goal. The GS includes a spacecraft model of resources, flight and operations constraints, and goals to be achieved.

Advantages of GS include reducing the operations complexity of the mission. Modern missions have far too many states and transitions between states to test, verify, and operate. Missions can be more resilient to anomalies by trying alternatives rather than safing during anomalies. Goals are much more intuitive to operators than long command sequences. GS are more adept at dealing with unknown or uncertain environments. Lastly, GS are more efficient because they close the loop on spacecraft control.

When autonomy is added to a goal-directed spacecraft, the science feedback loop can be moved from the ground to onboard the spacecraft. Typically, spacecraft operations involve acquiring instrument data, downlinking that data, analyzing the data, and then issuing new commands to the spacecraft to acquire more data based on the analysis. With autonomy onboard, the decision about what data to acquire next can be made by science agents. These science agents are directed with goals rather than detailed commands. The goals are expanded using planning software that has a model of the spacecraft resources and operating constraints. For non-autonomous spacecraft, operators would typically run detailed safety checks of spacecraft sequences and commands each week before upload. Using autonomous spacecraft, these checks only have to be performed once on the onboard spacecraft model. This greatly simplifies the ground planning process. It also allows the scientists to be more directly involved in the planning process. Scientists will select and submit observation goals directly to the flight control team for upload to the spacecraft. These goals will be integrated with spacecraft-produced science goals and engineering goals from the spacecraft operations team. The onboard planning software will determine which goals are achievable.

Another way autonomy software is changing operations is anomaly resolution. Typically when a spacecraft experiences an anomaly, a safe hold mode will be entered to allow the spacecraft to remain safe until the ground operators can resolve the problem. Using an autonomous spacecraft, the planning software can still function with knowledge of the failed component. Since the components are modeled as resources, and the resources are used in achieving goals, the planner can just plan using the degraded state of the resource. In some instances, the planner can select alternative methods of achieving a goal that do not rely on the failed component.

XV. Operational Challenges for More Effective Sensor Webs

Building more effective sensor webs involves many different challenges in the areas of information standardization and autonomy. These challenges are driven by the increasing complexity of sensor webs and the sensors within them.

The challenges in information standardization have evolved from the difficulty in the collection and analysis of information from many different types of sensors. For future sensor webs to operate more effectively, we need to develop standards on how to operate sensor webs, as well as a standard representation of sensor data. The Open Geospatial Consortium (OGC) is contemplating adoption of a technology called Web Processing Service (WPS), which is one step towards standards-based exposure of sensor data on the web.[9] This is a great step forward for accessing sensor data, but we need to create data standards so that the different sensor data and the models that use them can be fused together to answer complex scientific questions. Much of these data are in different spatial and temporal resolutions, but by combining them, we achieve greater spatial coverage and resolution than by analyzing any one sensor.

Another issue with data standardization is that end users of sensor data have insufficient technical expertise and time to extract information from the sensor data. This problem is being addressed with the evolving Earth science ontologies being created that will infuse metadata into the sensor data. This will allow data that can be filtered, summarized, and transformed, and will also allow features to be extracted into higher-level features. In addition, the same data can be reused for different applications. Different users will be able to have different views of the sensor data depending on their particular needs.

Many of today’s sensor webs employ little autonomy. The deployment and usage of sensors is usually tightly coupled with the specific location, application, and the type of sensors being used. Future science applications for sensor webs will require data from many different types of sensors and even integrating multiple sensor webs (sometimes referred to as a *system-of-systems*). To be effective, this will require a capability for publishing and discovering sensor resources. Once this infrastructure is in place, it will be much easier to pull additional sensors into a particular sensor web application.

Operation of complex sensor web systems will require cooperation between different missions and agencies. For example, the EO-1 sensor web involved combining data from EO-1, Aqua, Terra, GOES, QuickScat, and several ground sensor networks. Operations organizations will require more flexibility because they will often be working with systems they don't control. One very complicated sensor web example requiring cooperation is the Global Earth Observing System of Systems (GEOSS). GEOSS is an agreement among 69 nations to share Earth science data and models to achieve comprehensive, coordinated and sustained observations of the Earth system, in order to improve monitoring of the state of the Earth, increase understanding of Earth processes, and enhance prediction of the behavior of the Earth system. The GEOSS architecture has a focus on open systems, standard interface, interoperable formats, service oriented architectures, and semantic data. The GEOSS vision will never be realized without significant cooperation at all levels (from operators and spacecraft developers to governments).

XVI. Summary

The increase in onboard computing performance has enabled new levels of autonomy on spacecraft. This autonomy in turn has allowed spacecraft to be linked with Earth-based sensor networks. These ground networks have also been incorporating increased autonomy to form sensor webs. The paradigm shift toward highly autonomous spacecraft will enable future NASA missions to achieve significantly greater science returns with reduced risk and reduced operations cost. Demonstration of these sensor web capabilities will enable fast responding science campaigns and increase the science return of spaceborne assets. Future research in sensor webs will allow model and simulation driven sensors. Earth system models can be treated as virtual sensors within sensor web, influencing the operation of remote sensors and sensor networks as well as reacting to near real time observations supplied by the sensor web. Autonomy capabilities are being developed for sensors to allow them to interact with other sensors. The increased operation complexity in integrating multiple sensors into systems-of-systems will be a challenge for future spacecraft operators. No longer will they only have to operate a single platform with a limited number of sensors. Complex science questions demand the use of multiple sensors and science models. Governance and security within system-of-system architectures will need to be addressed. New operations techniques, as well as data standards and coordination tools will have to be implemented to ensure successful sensor web operations.

Acknowledgments

Portions of this work were performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. We would like to acknowledge the important contributions of Ashley Davies, Michael Burl, Andrea Donnellan, Steve Chien, Yunling Lou, Sharon Kedar, Sam Dolinar, Dan Dvorak, Daniel Tran, Gregg Rabideau, and Aaron Kiely of JPL, Stefan Falke of Northrop Grumman, Will Ivancic of NASA GRC, Peter Fox of UCAR, Dan Mandl and Karen Moe of NASA GSFC, and Antonio Ortega of USC.

References

[1]NASA ESTO/AIST Sensor Web PI Workshop, URL: http://esto.nasa.gov/sensorwebmeeting/, cited [12 June 2007], February 2007.

[2]Burl, M. C., Garay, M. J., Wang, Y., Ng, J., "Adaptive Sky: A Feature Correspondence Toolbox for a Distributed Cloud Monitoring Sensor Web," *IEEE Aerospace Conference*, Big Sky, MT, (2008).

[3]Fox, P., McGuinness, D. L., Raskin, R., Sinha, A. K., "Semantically-Enabled Scientific Data Integration," *U.S. Geological Survey Scientific Investigations Report 2006-5201*, (Geoinformatics 2006).

[4]Davies, A. G., Castano, R., Chien, S., Tran, D. Mandrake, L., Wright, R., Kyle, P., Komorowski, J-C., Mandl, D., and Frye, S., "(2008) Rapid Response to Volcanic Eruptions with an Autonomous Sensor Web: The Nyamulagira Eruption of 2006," *Proc. IEEE Aerospace Conference Paper 1180*, March 2008, Big Sky, MT.

[5]Lou, Y., Hensley, S., Le, C., Moller, D., "On-Board Processor for Direct Distribution of Change Detection," *IEEE Radar Conference and Proceedings*, Philadelphia, April 2004.

[6]Donnellan, A., Parker, J., Norton, C., Lyzenga, G., Glasscoe, M., Fox, G., Pierce, M., Rundle, J., McLeod, D., Grant, L., Brooks, W., Tullis, T., "QuakeSim: Enabling Model Interactions in Solid Earth Science Sensor Webs," *Proceedings of 2007 IEEE Aerospace Conference*, Big Sky, MT, March 2007.

[7]Sherwood, R., Chien, S., Tran, D., Davies, A., Castano, R., Rabideau, G., Mandl, D., Frye, S., Shulman, S., Szwaczkowski, J., "Enhancing Science and Automating Operations Using Onboard Autonomy," *The 9th International Conference on Space Operations*, Rome, Italy, June 2006.

[8]Dvorak, D., Ingham, M., Morris, R., and Gersh, J., "Goal-Based Operations: An Overview," *AIAA Infotech 2007*, Rohnert Park, CA, May 2007.

[9]Open Geospatial Consortium (OGC) page, URL: http://www.opengeospatial.org/standards/requests/28, cited [5 June 2007].