LIGHTS OUT AUTONOMOUS OPERATION OF AN EARTH OBSERVING SENSORWEB

Steve Chien(1), Daniel Tran(1), Ashley Davies(1), Mark Johnston(1), Joshua Doubleday(1), Rebecca Castano(1), Lucas Scharenbroich(1), Gregg Rabideau(1), Benjamin Cichy(1), Sharon Kedar(1), Dan Mandl(2), Stuart Frye(3), WenZhan Song(4), Philip Kyle(5), Rick LaHusen(6), Patrice Cappaelare(7)

(1) Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA, USA, 91109, {firstname.lastname}@jpl.nasa.gov
(2) Goddard Space Flight Center, 8800 Greenbelt Road, Greenbelt, MD, USA, 20771, dan.mandl@gsfc.nasa.gov
(3) Noblis, 3150 Fairview Park Dr, Falls Church, VA, USA, 22042, stuart.frye@gsfc.nasa.gov
(4) Washington State University, 14204 NE Salmon Creek Ave, Vancouver, WA, USA, 98686, songwz@wsu.edu
(5) New Mexico Tech, 801 Leroy Place, Socorro, NM, USA, 87801, kyle@nmt.edu
(6) Cascade Volcano Observatory, 1300 SE Cardinal Court #100, Vancouver, WA, USA 98642, rlahusen@usgs.gov
(7) Vightel Corporation, 8303 Winter Wind Ct, Ellicott City, MD, 21043, pat@vightel.com

ABSTRACT

The Earth Observing-1 Mission has been operating an autonomous, integrated sensorweb linking dozens of sensor nodes in 24/7 operations since 2004. Key to the successful lights out automation has been flexible web-based tasking with back-end mission operations automation.

This sensorweb has achieved a number of impacts including:

1. Routine re-tasking based on sensor alerts and scientist-defined campaigns with no additional operational effort.
2. Rapid response within hours to changing requests based on weather, science phenomena, and operational concerns.
3. Integration of triggers from the EO-1 mission, other space-borne assets, ground-based in situ instrumentation, and composite synthetic sensors (automatically generated updates such as aviation advisories).
4. Dramatic reduction in operations costs (over $1M US per year).

We describe our extensible sensorweb operations architecture, engineering and project results from operational experiences, and our ongoing effort to develop a web service interface to task, retrieve, and process science data.

1. INTRODUCTION

On Ross Island in remote Antarctica, the volcano Mount Erebus (Fig. 1) has a permanent lava lake at its summit. The volcano is forever rumbling, but occasionally, enhanced activity indicates an increase in effusion rate. Seismic tremors and large Strombolian explosions in the lake, result in the eruption of lava. In years past, such an event might have passed unnoticed, or might only have come to light days or weeks after the event. Now, thanks to the Earth Observing Sensorweb developed at the Jet Propulsion Laboratory (JPL) and Goddard Space Flight Center (GSFC), and critical in-situ instrumentation deployed by the Mount Erebus Volcano Observatory (MEVO), operated by New Mexico Tech., volcanologists around the world will have key science data about this eruption within hours.

Scientists from New Mexico Tech. have deployed an integrated suite of seismographic, tilt, acoustic, and camera sensors on Mount Erebus. Additionally, NASA's Terra and Aqua satellites fly overhead four times per day, skimming past at 7.5 kilometers per second. Each spacecraft carries a Moderate Resolution Imaging Spectroradiometer (MODIS) instrument, which acquires observations of the volcano at visible and infrared wavelengths, at resolutions of 250 to 1000 meters/pixel. A MODIS image is typically some 2700-km wide [1].

Figure 1. Mount Erebus, Ross Island, Antarctica

Streamed to GSFC, MODIS data are downlinked to ground stations and accessed by scientists worldwide. Additionally, the MEVO in-situ data are streamed via the internet to New Mexico Tech. in Socorro, New Mexico. At the Jet Propulsion Laboratory in Pasadena, CA, these data are monitored and will be assimilated into physical models of lava effusion. This processing and modeling indicates a dramatic increase in lava and results in a request to acquire more data.

This request drives a search for assets to acquire the needed data. For example, in response, in-situ cameras at the volcano are requested to acquire additional images of the increased activity. In addition, the software makes an observation request to the Earth Observing-1 (EO-1) mission system via the internet. When evaluating the observation request, based on the request's priority, the EO-1 ground system uplinks the observation request to the spacecraft. Onboard software evaluates the request, orients the spacecraft, and operates the science instruments to acquire high-resolution (up to 10 m/pixel) images with hyperspectral (220 or more bands) data for science analysis. Onboard, EO-1 processes this data to extract the volcanic eruption's signature, downlinking this vital information within hours. The EO-1 sensorweb has demonstrated this and similar scenarios since its first operations in August 2003 [2].

A wide range of operations satellite/platforms make their data freely available (e.g. broadcast or internet) in a rapid fashion (tens of minutes to several hours from acquisition). For example, data from MODIS flying on Terra and Aqua are available via Direct Broadcast in near real-time for regional coverage and 3-6 hours from acquisition from the GSFC Distributed Active Archive Center (DAAC). This data provides regional or global coverage with a wide range of sensing capabilities. For example, MODIS covers the globe roughly 4 times daily (two day and two night overflights). The Quikscat instrument flying on the SeaWinds spacecraft covers the majority of the globe daily.

Unfortunately, these global coverage instruments do not provide the high spatial and spectral resolution data desirable for many science applications. The above instruments range in resolution from MODIS with 250m-1km resolution to 1km and above for the other instruments. While ideally high resolution data would be available continuously with global coverage, typically, high resolution assets (e.g., ASTER, ALI, Hyperion) can image only limited swathes of the Earth – thus making them highly constrained and in high demand.

In this paper we describe our efforts to network sensors and science event recognizers/trackers with an automated response system to form a sensorweb. Our application uses data from a range of low resolution, high coverage sensors and constant operation sensors to trigger observations by controllable instruments. We also trigger observations from ground-based in situ instrumentation data. In turn, the data from the high resolution instruments may trigger observations or tasking of other sensors. Note that there are many other rationales to network sensors into a sensorweb. For example, an automated response might enable subsequent observations using complementary instruments such as imaging radar, infra-red, visible, etc. Or an automated response might be used to apply more assets to increase the frequency of observation to improve the temporal resolution of available data.

Our Earth observing sensorweb has been successfully operational since late 2003, responding to five different science disciplines and acquiring data from over 10 different sources. Table 1 displays a list of the science tracking system integrated into our system.

Discipline	Source	Detector
Volcanoes	MODIS (Terra, Aqua)	MODVOLC. U Hawaii
	GOES	GOESVolc
	POES	AVHRR - Volcano
	Air Force Weather Advisory	Volcanic Ash Alerts
	International Aviation Authorities	Volcanic Ash Advisories
	Tungurahura, Reventador	In-situ instruments, Harvard, UNH*
	HVO	Sensor alerts
	CVO	In-situ instruments*
	MEVO	In-situ instruments*
Floods	Quikscat	Dartmouth Flood Observatory
	MODIS	Dartmouth Flood Observatory
	AMSR	Dartmouth Flood Observatory
Cryosphere	Quikscat	Snow-ice, JPL/Nghiem
	Wisconsin Lake Buoys	UW Dept. Limnology
	SSM/I (DMSP F-13)	NSIDC*
Forest Fires	MODIS (Terra, Aqua)	RAPIDFIRE, UMD, MODIS Rapid Response
Clouds	EPOS	DoD
* in development		

Table 1. Science Alert Systems,

This remainder of the paper describes a common sensorweb scenario, as well as our efforts to expand the software and architecture to make use of the Open Geospatial Consortium (OGC) web services standards to seamlessly connect data from various sources and task other sensor sources. We also describe several of our new sensorweb examples as well as the targeted deployments making use of the new web services architecture.

2. SENSORWEB SCENARIO

The EO-1 sensorweb architecture consists of a number of components which operate in the following sequence of steps.

1. Asset1 acquires data (usually global coverage at low resolution or localized in-situ data).
2. An alert is published to the web indicating Asset1 data are available.
3. Data from Asset1 is downlinked via Assets1's web service.
4. These data are automatically sent to a web processing service to process the data to detect science events.
5. Science event detections are forwarded to a tasking system. This system generates a task request which may forward a request for more data to several other assets which may result in a request issued to an asset's automated planning system to schedule a collect for data.

6. The automated planning system then generates a command sequence to acquire a new observation.
7. This new command sequence is uplinked to Asset2 which then acquires the high resolution data.
8. An alert indicating Asset2 data are available is posted to the web.
9. These data are then available to be downlinked by connecting to Asset2's web service.

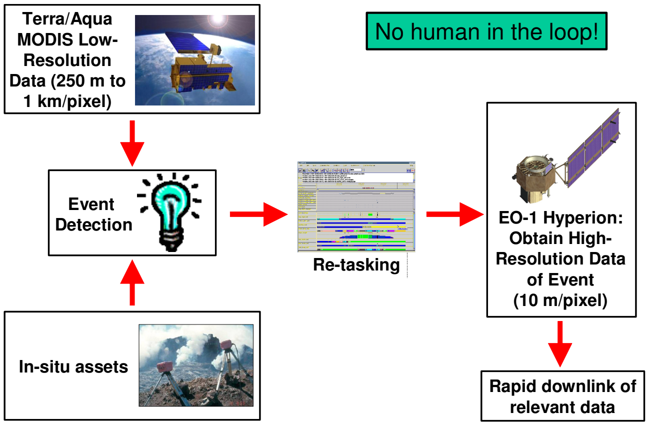

Figure 2. Sensorweb Detection and Response Architecture

In our operational system thus far, Asset2 has been the Earth Observing-1 spacecraft (EO-1). EO-1 is the first satellite in NASA's New Millennium Program Earth Observing series. The primary focus of EO-1 is to develop and test a set of advanced technology land imaging instruments.

EO-1 was launched from Vandenberg Air Force Base in 2000 with two principal science instruments, the Advanced Land Imager (ALI) and the Hyperion hyper spectral instrument. The ALI is a multi-spectral imager with 10m/pixel pan-band resolution and 9 spectral bands from 0.433 to 2.35 μm with 30m/pixel resolution. The Hyperion is a high-resolution imager capable of resolving 220 spectral bands (from 0.4 to 2.5 μm) with a 30m/pixel spatial resolution. The instrument images a 7.5 km by 42 km land area per image and provides detailed spectral mapping across all 220 channels with high radiometric accuracy [3].

3. EO-1 SENSORWEB ARCHITECTURE

The EO-1 sensorweb architecture is currently being updated to support the interface standard in development by the Open Geospatial Consortium (OGC) Sensor Web Enablement (SWE) group. The role of the SWE group is to develop common standards to [4]:

- Discover sensor systems observations and observation processes.
- Determine a sensor's capabilities and quality of measurements.
- Access sensor parameters, allowing software to process and geolocate observations.
- Retrieve real-time or time-series observations and coverage in standard encodings.
- Task sensors to acquire observations of interest
- Subscribe and publish alerts issued by sensors or sensor devices based on certain criteria.

This common interface is achieved by developing a standard XML-based web service that can be invoked in either a REST or SOAP/WSDL methodology. The goal of providing a web service interface is to create a web enable sensorweb, and to begin standardizing the discovery and communication protocol between various heterogeneous sensors. It begins the process of allowing sensors to seamlessly integrate other sensors together into a global network.

Several of the services and the interfaces we are implementing to discover, task, and process EO-1 science data are:

1. The Sensor Planning Service (SPS): used to determine if an EO-1 observation request can be achieved, re-task the satellite to acquire science data, determine the status of an existing request, cancel a previous request, and obtain information about other OGC web services.
2. The Sensor Observation Service (SOS): used to retrieve EO-1 engineering or science data. This includes access to historical data as well as data requested and acquired from the SPS.
3. The Web Processing Service (WPS) [5]: used to perform a calculation on the EO-1 data. We intend to make available the opportunity to run image classifiers on EO-1 data. These classifiers will range from the set available onboard EO-1 (thermal, cryospheric, and fluvial algorithms) to custom defined classifiers defined by scientists for determining areas of interest.
4. The Sensor Alert Service (SAS): used to publish and subscribe to alerts from EO-1. Alerts will be issued when the engineering and science data products are available.
5. A description of the Hyperion and ALI instrument and their associated products and services using the Sensor Model Language (SensorML). SensorML provides a high level description of sensors and observation processes using an XML schema methodology. It also provides the functionality for users[1] to discover the EO-1 instruments on the web along with services to task and acquire sensor data (such as the SPS, SOS, and WPS).

---

[1] Users include humans as well as software agents acting as proxies for humans

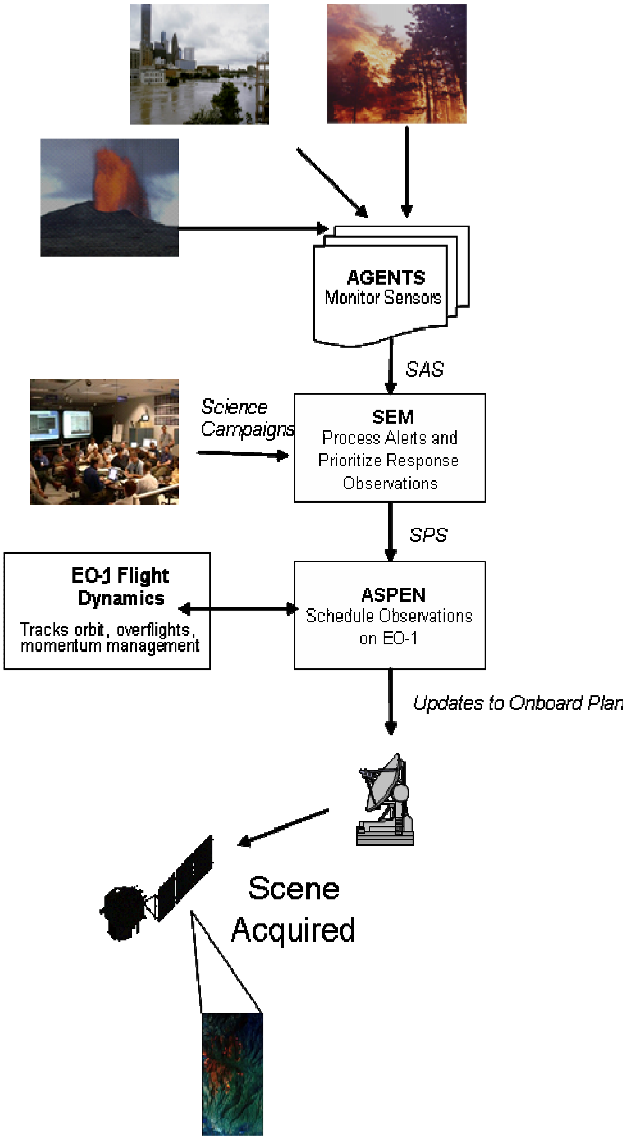

Figure 3. Sensorweb Response

With a flexible standardized web services interface in place, we have several options in developing the client software to invoke the service. The direct method is to invoke the service by developing custom client software which will access these services, parsing and handling the results returned. However, we are also working to develop a process model of the user intent, such as a request to acquire image data for a region of the world, and automatically discover the available sensors and processes from a catalogue service. This would be achieved by encoding the description sensors and their associated processing services in the OGC SensorML. SensorML is an XML-based description of the sensors and processes which allow them to be discovered through a catalogue or web search.

Development has begun to invoke the web services by discovering the sensors available, but at this time, the Earth Observing Sensorweb invoke these services directly. We describe several of the major components developed to automate the re-tasking of the EO-1.

1. Tracking systems for each of the science disciplines automatically acquire and process satellite and ground network data to track science phenomena of interest. These science tracking systems publish their data automatically to the internet in their own format. In some cases this is via the http or ftp protocol; in others, via email subscription and alert protocols.
2. Unless these tracking systems develop a Sensor Alert Service (SAS) to clients to register for published alerts, science agents act as the front end web service interface. These science agents either poll these sites (http or ftp) to pull science data or simply receive email notifications of ongoing science events. These science agents then publish these alerts via the SAS to any consumers registered to receive them. Agents also implement a Sensor Observation Service (SOS) to allow clients to retrieve the tracking system's science data.
3. A science event manager, registered to receive alerts, connects to the SOS to retrieve the science data, processes the notifications and matches them up with a science campaign. When a match occurs, as specified in the science campaign, a task request is generated and processed. A task request is a list of objectives to be achieved, where the user has the flexibility to specify a wide range of objectives to respond to the alert. These include submitting an observation request to EO-1 to requesting data processing of science data.
4. EO-1 observation requests are processed by the EO-1 SPS, using the ASPEN automated mission planning system. ASPEN integrates these requests and schedules observations according to priorities and mission constraints. For observations that are feasible, the science event manager issues a request to EO-1, and the uplinks the request to the spacecraft.
5. Onboard EO-1, the Autonomous Sciencecraft software [6] accommodates the observation request if feasible. In some cases onboard software may have additional knowledge of spacecraft resources or may have triggered additional observations so several uplinked requests may not be feasible.
6. Later, the data are downlinked, processed, and delivered to the requesting scientist.

3.1. Science Agents

The SAS is the primary method of publishing and subscribing to the alerts by various tracking systems. However, most tracking systems currently publish data in their own format. These formats have ranged from near raw instrument data, to alerts in text format, to periodic updates, to a wide range of text formats. To resolve this inconsistency and provide a common interface to users, science agents are front end web services that encapsulate sensor, science, tracking specific information. Science agents implement the SAS to publish alerts to subscribers and also implement the SOS to provide users access to the tracking system data.

3.2. Science Event Manager and Science Campaigns

The science event manager enables scientists to specify mappings from science events to task requests. It enables them to track recency and count of events and perform logical processing. It also enables them to track based on target names or locations, and other event specific parameters (for example, some tracking systems produce a confidence measure). As an example, a volcanologist might specify, for the Kilauea site, that several tracking systems would need to report activity with high confidence before a spacecraft observation is requested. This is because Kilauea is quite often active. On the other hand, even a single low confidence activity notification might trigger a task request of imaging of Piton de la Fournaise or another less active site.

The science event manager receives published alerts from the SAS, implemented by the science agent front end, or the original tracking system. These alerts are compared against a science campaign, and if a match occurs, the specified task request is executed. Task requests range from acquiring a single EO-1 acquisition of the target to acquiring as many collects of a target within 10 day. They also specify any science processing to be ran either onboard EO-1, or on the ground. To process a task request for an EO-1 data collect, the science event manager invokes the EO-1 SPS, requesting the feasibility of an observation request, and submitting the observation request, re-tasking the spacecraft.

3.3. Automated Observation Planning

To determine the feasibility of an observation request to EO-1, we automate the mission planning processing using the ASPEN/CASPER planning & scheduling system [7][2]. ASPEN represents mission constraints in a declarative format and searches possible mission plans for a plan that satisfies many observation requests while respecting priorities, and also obeys mission operations constraints. ASPEN has been used in a wide range of space mission applications [8] including spacecraft operations scheduling, rover planning, and ground communications station automation. Fig. 4 shows the graphical user interface of an EO-1 operations plan.

ASPEN/CASPER maintains a baseline schedule of activities to be performed on EO-1. As requests for observations are received by the SPS, ASPEN/CASPER is invoked to determine the feasibility of the schedule. If the observation request is feasible and uplinked to the spacecraft, the baseline schedule is updated to reflect this new goal.

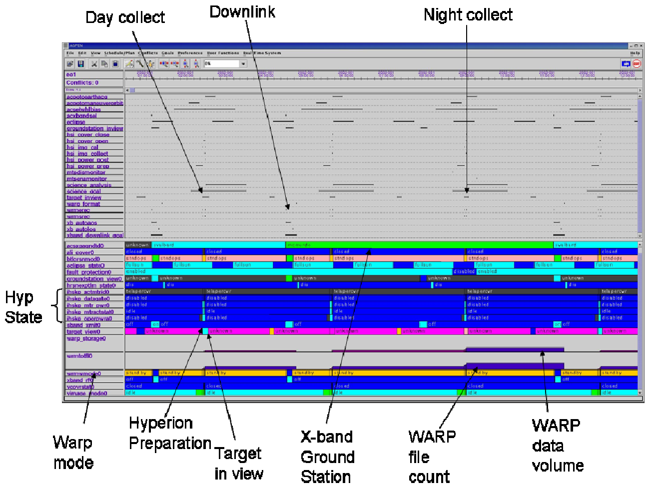

Figure 4. EO-1 Observation Plan as displayed by the ASPEN GUI

3.4. Science Data Access

One of the design goals of the sensorweb project has been to provide scientists easy access to multiple data sources on a single science event (such as a volcanic eruption, forest fire, etc). This is achieved by providing the Sensor Observation Service (SOS). EO-1 provides a SOS and the science agents also implement this service to allow clients to retrieve and access the data. Scientists are notified by the SAS when data are available for downlink and viewing.

4. SENSORWEB EXAMPLES AND APPLICATIONS

Our Earth observing sensorweb is operational since late 2003, responding to many science disciplines [2]. In this section, we describe several of our new science tracking systems as well as our targeted deployments of the sensorweb, making use of the new web services infrastructure.

4.1 National Snow and Ice Data Center

Many cryosphere applications are of interest to scientists worldwide. Cryosphere includes phenomena such as glacial ice breakup, sea ice breakup, melting, and freezing, lake ice freezing and thawing, and snowfall and snowmelt. Using data acquired from the National Snow and Ice Data Center (NSIDC), data we track ice formation and melting and use this information to automatically triggering higher resolution imaging with EO-1.

The NSIDC make available a sea ice index containing the near real-time sea ice concentrations of the Earth's polar region. These data are generated from the Defense Meteorological Satellite Program (DMSP) F13 Special Sensor Microwave/Imager (SSM/I) and created using brightness temperatures from NASA's Global Hydrology

---

[2] ASPEN is the ground, batch planner; CASPER is the embedded, flight planner. Both share the same core planning engine.

Resource Center (GHRC) at Marshall Space Flight Center (MSFC). The index provides sea ice concentration values at 25km resolution with a temporal coverage of three to six months [9].

We acquire the sea ice concentration values on a daily basis, and perform a simple day-to-day comparison of the values to determine regions of the poles that are freezing or thawing with a preference for regions that have the largest change area. Our algorithm for observation selection is as follows:

1. Acquire the sea ice concentration values from NSIDC.
2. Compare the values with the previous day's concentrations for the same locations.
3. Determine the polar locations where the difference in concentration values is greater than 40 units. We consider these to be the locations of interest.
4. Order the locations of interest with a preference for those that are nearest other locations with a large discrepancy in concentration values.
5. Issues an EO-1 observation request for the top 10 locations

The integration of the NSIDC sensorweb has been successfully operational since early 2007, processing over 300 events from the data acquired.

4.2 Cascade Volcano Observatory

In late 2004, Mt. St. Helens began a process of building a new lave dome within its crater. Hundreds of small tremors were measured as well as the release of ash and steam into the air. Our goal is to acquire high resolution infrared data of Mt. St. Helens when these events occur and provide them to geologists and volcanologists as quickly as possible, with the hope that it also serves as platform and location to develop an integrated system which is applicable other less-accessible and less-studied volcanoes.

JPL is collaborating with Washington State University, Vancouver (WSU), and the Cascade Volcano Observatory (CVO) of the United States Geological Survey (USGS) to install a new network of in-situ sensors, ranging from seismometers to acoustic flow monitors on Mt. St. Helens and integrate them into the Earth observing sensorweb.

We describe our targeted scenario and detail its interaction with the provided EO-1 web services. Figure 5 shows the interaction between a satellite and a network of ground sensors.

1. In-situ sensors readings are transmitted from Mt. St. Helens back to CVO in Vancouver, WA for storage and analysis.
2. Sensor data are automatically analyzed and if a triggering condition is detected, as setup by the volcanologist and geologists, the EO-1 web service is accessed.
3. Through the SPS, EO-1 is tasked to acquire high resolution data of the target.
4. Through the SOS, EO-1 science data just acquired is transmitted and stored at CVO.
5. These science data are then sent to an available Web Processing Service (WPS), configured to run a thermal detection algorithm to determine the hot-spot regions. These results are also sent back to CVO for storage.
6. Results of the thermal detection algorithm translate to another triggering condition, causing the volcano sensors to be re-configured and reprioritized for transmission.

This scenario demonstrates re-tasking of the EO-1 spacecraft, and based on the science data collected, a re-tasking of the ground sensor network.

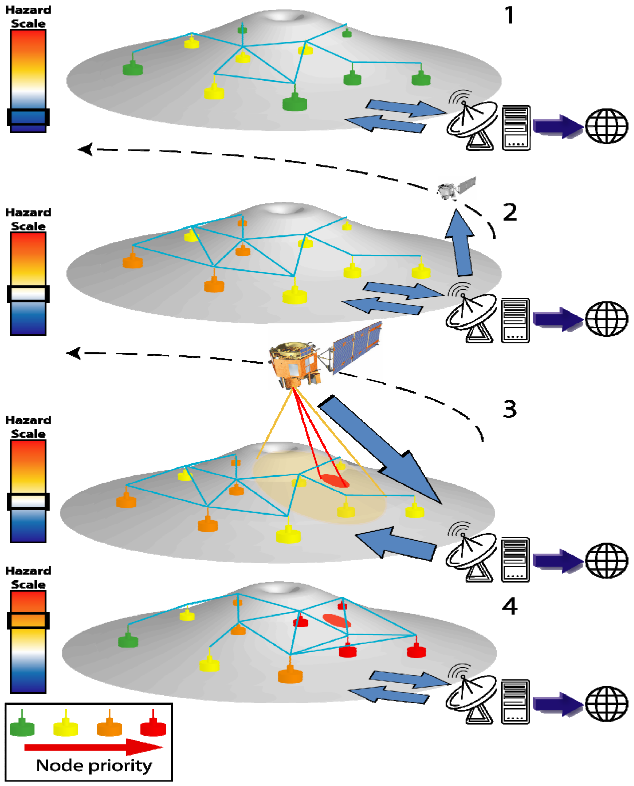

Figure 5. Volcano Sensorweb Deployment

4.3 Mount Erebus Volcano Observatory

JPL and the EO-1 mission are collaborating with New Mexico Tech. (NMT), which operates the Mount Erebus Volcano Observatory (MEVO) in developing an integrated space ground sensorweb for monitoring Mount Erebus.

NMT has deployed a wide range of sensors to the Mount Erebus summit which provide seismographic, acoustic, tilt, and image data on volcanic activity. These sensors can be maintained and upgraded during the Antarctic summer, but during the remainder of the year, they must be operated remotely from NMT. While the majority of the sensors follow a regular unalterable policy for acquiring data, the remote camera enables both visible and infra-red data to be acquired on demand.

As with the in-situ integration of the CVO assets, our collaboration with MEVO/NMT utilizes both ground and space assets, with each segment potentially causing a change in the operations of the other. For MEVO, acoustic, seismographic, and infrared data can cause an alert of increased activity which then triggers spaceborne observations. Correspondingly, space-based observations can detect activity which then triggers imaging with the remote MEVO camera to provide ground-based imagery of the phenomena.

Another aspect of the NMT/MEVO sensorweb is that JPL has developed preliminary models to track the evolution of the Erebus lava activity. In this effort infra-red imaging capability (mostly from space) is used to estimate the thermal output of the Erebus volcano, see Fig. 6. This can be compared to historical data to determine if there is a significant increase in activity (e.g. activity that is above and beyond normal fluctuations). Introducing a physical model increases the accuracy and reliability of sensorweb triggers. A promising area of work is to continue to enhance the physical models that are used to drive sensorweb operations to realize this potential for improved performance.

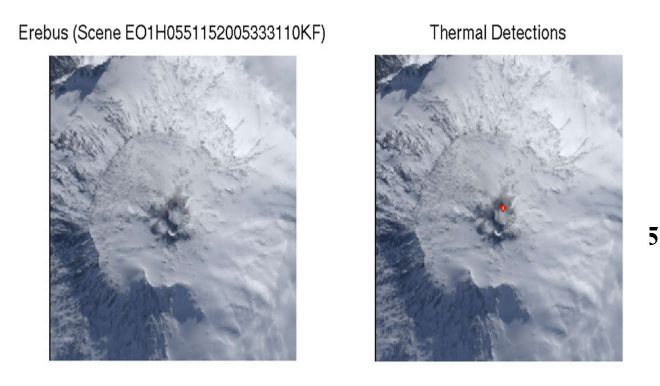

Figure 6. Thermal hotspots detected on Mount Erebus

4.4 Uninhabited Aerial Vehicle - Synthetic Aperture Radar

JPL is developing an onboard autonomy software package to integrate a Synthetic Aperture Radar (SAR) payload on an Uninhabited Aerial Vehicle (UAV) into an earth observing sensorweb. In the near term tests are being conducted with a SAR instrument flying in a Gulfstream Jet, (Fig. 7). The end goal is to develop and demonstrate autonomy software that would enable a UAVSAR to (a) acquire data as directed by other nodes of a sensor network (e.g be tasked as a node in the sensorweb) and (b) based on its own data acquisition, the UAVSAR would make requests of other sensorweb assets, and (c) the UAVSAR also represents an autonomous sensorweb node that may autonomously respond to changing events, goals, and conditions.

Specifically, we are working towards demonstration scenarios where:

1. The UAVSAR acquires SAR imagery of an ongoing forest fire.
2. This imagery is processed to develop an updated fuel map and integrated with wind and fuel estimation to derive new areas for observation (e.g. locations to which the fire is likely to spread).
3. The UAVSAR and other assets (space, ground, air) are automatically tasked to gather additional data of these areas.

In such scenarios the UAVSAR is an autonomous node interacting with other nodes in the sensorweb to achieve the overall sensorweb goals of tracking the forest fire.

Figure 7. SAR instrument on a Gulfstream Jet, courtesy http://uavsar.jpl.nasa.gov

5. RELATED WORK AND SUMMARY

There has been considerable effort devoted towards closed loop science for rovers at NASA Ames [10], JPL [11], and Carnegie Mellon University [12]. These efforts have some similarity in that they have science, execution, and in some cases mission planning elements. However, because surface operations (e.g. rover) are very different from orbital operations, their focus is on integration with rover path planning and localization, reliable traverse, etc., whereas our efforts focus on reliable registration of remote sensed data, interaction with orbital mechanics, and multiple platforms. The MISUS system [13] also describes a closed-loop multi-rover autonomous science architecture.

One closely related effort is led by Keith Golden [14] at NASA Ames to enable real-time processing of Earth Science data such as weather data. However, this work focuses on the data processing and information gathering aspect of the problem, and thus is complementary to our sensorweb work which focuses on the operations aspect of the problem. Indeed, we have discussed with Golden the possibility of a joint sensorweb information gathering demonstration.

The Autonomous Sciencecraft Experiment on EO-1 [6] demonstrates an integrated autonomous mission using onboard science analysis, replanning, and robust execution. The ASE performs intelligent science data selection and autonomous retargeting. ASE represents a single spacecraft onboard autonomous capability. In contrast the sensorweb uses multiple assets in concert.

This paper has described ongoing work to link together an automated science event tracking system with an autonomous response capability based on automated planning technology. The Earth Observing Sensorweb enables fast response science campaigns and increases the science return of spaceborne assets. These capabilities have been demonstrated since August 2003 [2] and we've described several new deployments as well as the updates to the sensorweb software to support the OGC web services interface.

REFERENCES

1. Goddard Space Flight Center, MODIS instrument page: http://modis.gsfc.nasa.gov
2. S. Chien, B. Cichy, A. Davies, D. Tran, G. Rabideau, R. Castano, R. Sherwood, D. Mandl, S. Frye, S. Shulman, J. Jones, S. Grosvenor, "An Autonomous Earth-Observing Sensorweb", *IEEE Intelligent Systems*. May/June 2005.
3. Goddard Space Flight Center, EO-1 Mission page: http://eo1.gsfc.nasa.gov
4. M. Botts, G. Percivall, C. Reed, J. Davidson, "OGCSensor Web Enablement: Overview And High Level Architecture.", http://www.opengeospatial.org.
5. P. Schut, A. Whiteside, "OpenGIS Web Processing Service", OGC Public Discussion Paper, OGC 05-007r4.
6. S. Chien, R. Sherwood, D. Tran, B. Cichy, G. Rabideau, R. Castano, A. Davies, D. Mandl, S. Frye, B. Trout, S. Shulman, D. Boyer, "Using Autonomy Flight Software to Improve Science Return on Earth Observing One," *Journal of Aerospace Computing, Information, and Communication*. April 2005.
7. S. Chien et al., Automated Scheduling and Planning Environment (ASPEN), project home page, aspen.jpl.nasa.gov, 2003.
8. S. Chien, G. Rabideau, R. Knight, R. Sherwood, B. Engelhardt, D. Mutz, T. Estlin, B. Smith, F. Fisher, T. Barrett, G. Stebbins, D. Tran, "ASPEN - Automating Space Mission Operations using Automated Planning and Scheduling", *International Conference on Space Operations (SpaceOps 2000)*, Toulouse, France. June 2000
9. J. Maslanik, J. Stroeve. 1999, updated daily, "Near real-time DMSP SSM/I daily polar gridded sea ice concentrations", Acquired Daily, Boulder, Colorado USA: *National Snow and Ice Data Center*. Digital Media.
10. V. C. Gulick, R. L. Morris, M. A. Ruzon, and T. L. Roush, "Autonomous image analysis during the 1999 Marsrokhod rover field test," *J. Geophysical Research*, Vol.106, No. E4, pp. 7745-7764, 2001.
11. R. Castano, R. C. Anderson, T. Estlin, D. DeCoste, F. Fisher, D. Gaines, D. Mazzoni, and M. Judd, "Increased Mission Science Return Through In-Situ Data Prioritization", *IEEE Aerospace Conference*, Big Sky, Montana, March 2003.
12. T. Smith, "Science Autonomy in the Atacama," Proceedings of the International Conference on Machine Learning *Workshop on Machine Learning for Autonomous Space Applications*, Washington, DC, August 2003.
13. T. Estlin, T. Mann, A. Gray, G. Rabideau, R. Castano, S. Chien, and E. Mjolsness, An Integrated System for Multi-Rover Scientific Exploration, *National Conference on Artificial Intelligence*, Orlando, Florida, August 1999.
14. K. Golden, W. Pang, R. Nemani, and P. Votava, "Automating the Processing of Earth Observation Data," *Intl Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS)*, Nara, Japan, May 2003.

ACKNOWLEDGEMENT

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.