Lights Out Operations of a Space, Ground, Sensorweb

Steve Chien[1], Daniel Tran[1], Mark Johnston[1], Ashley Gerard Davies[1], Rebecca Castano[1], Gregg Rabideau[1], Benjamin Cichy[1], Joshua Doubleday[1], David Pieri[1], Lucas Scharenbroich[1], Sharon Kedar[1] ,Yi Chao[1], Dan Mandl[2], Stuart Frye[3], WenZhan Song[4], Philip Kyle[5], Rick LaHusen[6], Patrice Cappelaere[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

We have been operating an autonomous, integrated sensorweb linking numerous space and ground sensors in 24/7 operations since 2004. This sensorweb includes elements of space data acquisition (MODIS, GOES, and EO-1), space asset retasking (EO-1), integration of data acquired from ground sensor networks with on-demand ground processing of data into science products. These assets are being integrated using web service standards from the Open Geospatial Consortium. Future plans include extension to fixed and mobile surface and subsurface sea assets as part of the NSF's ORION Program.

I. Introduction

Recent developments have seen a dramatic increase in the deployment of sensors to observe dynamic terrestrial science events. These assets include spaceborne remote sensors, air, land and water (including subsurface) sensor networks.

Because of the explosion in sensor technology and its applications, controlling, allocating, and using these sensor resources is an increasing challenge. As the number and complexity of sensor networks has increased, managing these suites of sensors to extract the answers to key science questions has emerged as a central issue.

In this paper we describe an agent and service oriented approach to automated management of sensor networks designed to alleviate the burdens of managing sensor networks and their commensurate data. In this approach, sensors, alerts, and data processing are presented as services and science campaigns can be achieved using a range of automation from software agents that automatically compose and utilize these services to scientists that manually access the services to conduct their investigations.

These concepts have been embodied in an autonomous, integrated sensorweb linking scores of space and ground sensors in 24/7 operations since 2004. This sensorweb includes space-based elements (MODIS, GOES, and EO-1), space asset retasking (EO-1), ground networks (volcano observatories), and air assets (e.g., UAVSAR testbed). Future plans include extension to marine surface and subsurface sea assets as part of the NSF's ORION Program.

The central theme of the sensorweb is as follows:

"An intelligent sensorweb is a networked set of sensors that modifies its own observation patterns based on the sensor data observed"

The basic rationale of an intelligent sensorweb is that by using the current sensor data to infer information about the science phenomena being studied; one can subsequently observe more effectively by utilizing this understanding of the science phenomena. A key point is that complete understanding of the underlying science phenomena is not required to enhance the sensor network performance – often partial information is sufficient to improve data acquisition.

This general principle applies to an incredibly wide range of potential sensor networks. For example, it applies to various concepts including the following.

1. Directing a pointable satellite to track volcanic activity as measured by ground sensors.
2. Directing a subsurface submersible to investigate a possible eddy event as detected by surface radar.
3. Noting the signature of earthquakes from seismographic and GPS sensors and allocating a higher data rate of acquisition from those best placed to measure the areas likely to have future activity.

Thus, the sensorweb is a tool that enables scientists greater control over observation, measurement, and monitoring of complex natural phenomena. For example, in the science flow described below in Figure 1, science measurements are used to develop models of science phenomena, enabling characterization and understanding of the phenomena. These models and understanding subsequently enable monitoring and forecasting of said phenomena. The intelligent sensorweb is a tool that enables automation to leverage the scientists' model, understanding, and characterization of the phenomena to direct measurement (sensing) to more effectively monitor and forecast science events.

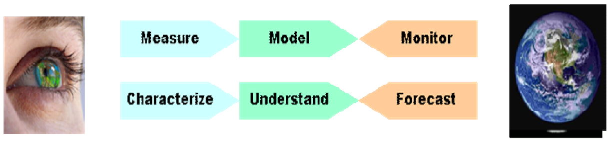

Figure 1: Science Understanding Flow

In this paper we first describe the extensible sensorweb architecture under development. We then describe project results from operational experiences along with our ongoing efforts to develop web service interfaces to task, retrieve, and process science data. We then describe our future plans for extension of these concepts into new applications and uses.

The Open Geospatial Consortium (OGC) Sensor Web Enablement (SWE) initiative [Botts et al] has been developing general service oriented interfaces to enable networked sources to utilize the sensorweb as generic information and processing resources. Several of these protocols are described below.

1. The Sensor Planning Service (SPS): used to determine if a sensor is available to acquire requested data. For example, using the SPS, an observation request to a space asset can be issued to acquire science data, determine the status of an existing request and cancel a previous request.
2. The Sensor Observation Service (SOS): used to retrieve 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): used to perform a calculation on the acquired remote sensing data. This includes processing the raw data into derivative products such as vegetation indices, soil moisture, burn areas, lava flows and effusions rates, etc.
4. The Sensor Alert Service (SAS): used to publish and subscribe to alerts from space, ground, and air assets. Users register with this service and provide conditions for alerts. When these conditions are met by the acquired data, alerts containing the data along with time and location of the event are automatically issued to the user.
5. A description of the space, air, and ground instruments 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 to discover instruments on the web along with services to task and acquire sensor data (such as the SPS, SOS, SAS, and WPS).

II. A Space Directed Sensorweb

Thus far our efforts have focused on control of space based assets. Our efforts to date have centered on automated tasking of the Earth Observing-1 Satellite [Chien et al. 2005]. This control flow is highlighted in Figure 2 and consists of the following steps.

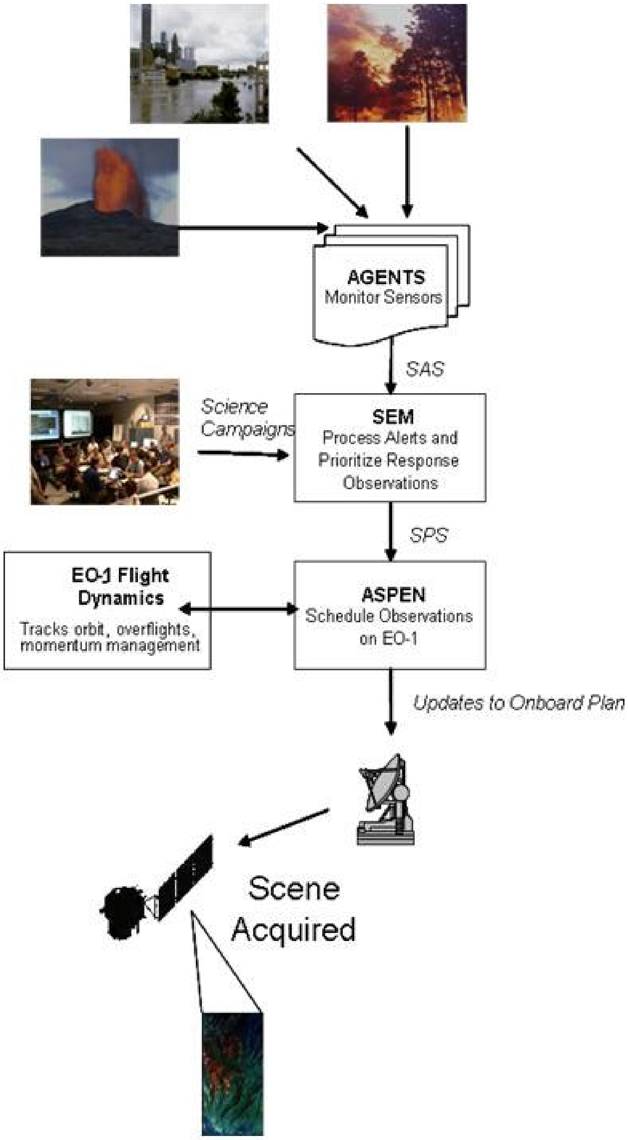

Figure 2: Space Sensorweb

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 publish alerts conforming to the Sensor Alert Service (SAS) interface, 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 SAS to retrieve the science alerts, 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 Sensor Planning Service (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 [Chien et al. 2005] 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 downlink, processed, and delivered to the requesting scientist.

This 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
* under development		

Table 1. Science Alert Systems

This sensorweb has achieved numerous impacts including:

1. Routine re-tasking of assets and sensor reconfiguration based on sensor alerts and scientist-defined campaigns with no human in the loop.
2. Rapid response within hours to changing observation requests based on weather, science phenomena, and operational concerns.
3. Network reconfiguration and observation via triggers from space-borne assets, ground-based in situ instrumentation, and derived sensors (e.g. automatic interpretation of multiple sources and/or automatically generated updates such as aviation advisories).
4. Dramatic reduction in operations costs (over $1M US per year).

The remainder of the paper describes a generalization of this space-driven sensorweb in several respects:

- to include control of a wider range of assets including ground networks, autonomous aerial, land, or marine vehicles
- to enable control of passive sensor networks (such as fixed sensors)
- to enable closer integration of science modeling and processing in the control of the network
- to enable hierarchical construction of sensor networks with multiple levels of external linkages.

We first describe the general sensorweb architecture. Then we describe a number of applications under development.

III. Generalized Sensorweb Architecture

We are working to generalize our sensorweb architecture as illustrated below in Figure 3. Key generalizations are:

1. Separation of the general processes or event detection and response. Event detection involves analysis of the sensor data to synthesize it into higher level conclusions such as science events. Response includes the science

requests caused by the new data (e.g. requests for new data, changes to the observation strategy, policy, or allocations) as well as prioritization and resolving scheduling conflicts required by limited sensing resources.

2. Explicit consideration of science modeling and its direct influence on both event detection and response. In event detection, the current estimation of the overall physical state of the environment often drives the interpretation of current measurements – hence the centrality of the science model. In response, often the purpose of new observations is to fill gaps in the model or to reduce some uncertainty in the model – again calling for an explicit, central role for the model.

3. Making the sensorweb hierarchical. At one level a local sensor network or installation may have its own event detection, modeling, and closed loop control. This sensorweb may also provide external notifications which may be treated as data to another sensor web. Likewise, this sensorweb may also accept external requests for observation.

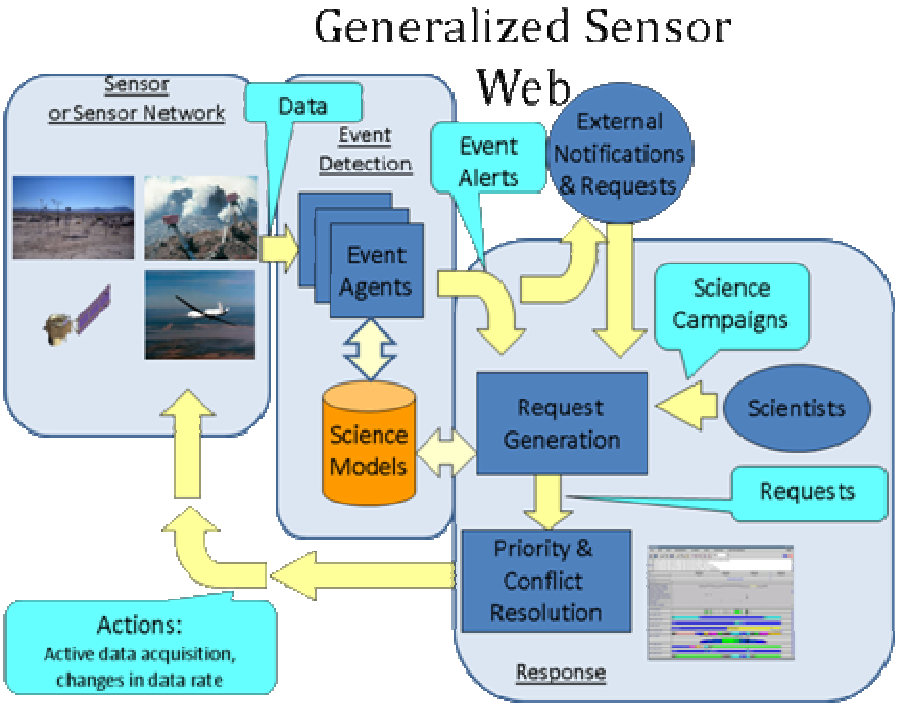

Figure 3: Generalized Sensor Web Architecture

IV. Example Sensorweb Applications

Our Earth observing sensorweb has been operational since late 2003, responding to many science disciplines. 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 generalized sensor web architecture and new web services infrastructure.

Cascade Volcano Observatory

In late 2004, Mt. St. Helens began a process of building a new lava 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 a platform and location to develop an integrated system which is applicable to 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 two target scenarios, one highlighting the local control of the in-situ sensor network and one highlighting the interaction between the local sensor network and the space sensor network.

Case 1: In-situ Sensor Network Reconfiguration

1. In-situ sensors acquire data and transmit from Mt. St. Helens back to CVO in Vancouver, WA for storage and analysis.
2. At CVO, sensor data are automatically analyzed to assess response conditions.
3. If response conditions are met, the in-situ network may automatically be reconfigured. For example, increased seismographic activity may drive a re-allocation of bandwidth to a geographic area of the volcano of greatest interest. Or degradation of a sensors signal may lead to re-allocation of bandwidth to better performing sensors.
4. Throughout this process the control of the in-situ network proceeds in the context of external requests for data (in this case both raw and derivative data products.

Case 2: In-situ network triggers space segment through web services.

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 services are 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. Prototype ground sensors developed by USGS

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 [Davies *et al.,* 2007].

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, a remote camera enables both visible and infra-red data to be acquired on demand (this camera is not currently operational due to hardware and environmental issues).

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 the integration of science modeling software into the sensorweb. 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) which in turn is used to estimate the lava flow. 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 can increase 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.

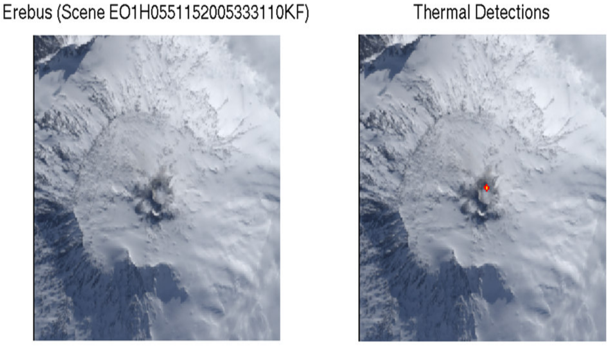
Erebus (Scene EO1H0551152005333110KF)

Thermal Detections

Figure 6. Thermal hotspots detected on Mount Erebus

Volcano Processing Workflows

JPL is also developing automated workflow processing for the Volcano sensorweb. In this effort, the following steps occur in response to a volcano alert.

1. Sensors perform ongoing monitoring of volcano sources (e.g. VAAC, MODVOLC, in situ,…)
2. When trigger conditions are met, alerts are triggered (SAS).
3. This alert is sent to the ASPEN planning system which constructs a workflow to satisfy the pre-defined science requirements. Currently this workflow includes the following steps.
    1. Gathering past acquired data from MODIS (WCS), EO-1 (SOS), in-situ (SOS)
    2. Tasking assets (EO-1, in-situ) to acquire future data (SPS)
    3. Actually getting tasked future data (SOS)
    4. Processing past and current data to compute volcano context products including thermal output, effusion rate, correlated with inflation/deflation events, seismic events, acoustic events (WPS).
    5. Delivery of above data products to relevant parties. Notification of parties as to the availability of said data (SAS).
4. Newly acquired data may trigger subsequent workflows, alerts (goto step 2. above).

Unpiloted 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 on a Gulfstream Jet. 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, and 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.

V. Related Work and Conclusion

There has been considerable effort devoted towards closed loop science for rovers at NASA Ames [Gulick et al. 2001], JPL [Castano et al. 2003], and Carnegie Mellon University [Smith 2003]. 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 [Estlin et al. 1999] also describes a closed-loop multi-rover autonomous science architecture.

One closely related effort is led by Keith Golden [Golden et al. 2003] 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 [Chien et al. 2005] demonstrates an integrated autonomous mission using onboard science analysis, replanning, and robust execution. The ASE performs intelligent science data selection and autonomous retargeting [Davies et al., 2006; Doggett et al., 2006; Ip et al., 2006]. 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 and we have described several new deployments as well as the updates to the sensorweb software to support the OGC web services interface.

Acknowledgments

The research described in the paper was performed at the Jet Propulsion Laboratory, California Institute of Technology under a contract with the National Aeronautics and Space Administration.

References

M. Botts, G. Percivall, C. Reed, J. Davidson, “OGC Sensor Web Enablement: Overview and High Level Architecture.”, http://www.opengeospatial.org.

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.

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.

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.

S. Chien et al., Automated Scheduling and Planning Environment (ASPEN), project home page, aspen.jpl.nasa.gov, 2003.

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

A. G. Davies, J. Calkins, L. Scharenbroich, R. G. Vaughan, R. Wright, P. Kyle, R. Castańo, S. Chien, and D. Tran (2007) Multi-Instrument Remote and In Situ Observations of the Erebus Volcano (Antarctica) Lava Lake in 2005: a Comparison with the Pele Lava Lake on the Jovian Moon Io, submitted to *J. Volc. Geotherm. Res,* 2007

A. G. Davies, A. G., S. Chien, V. Baker, T. Doggett, J. Dohm, R. Greeley, F. Ip, R. Castaňo, B. Cichy, G. Rabideau, D. Tran and R. Sherwood, Monitoring Active Volcanism with the Autonomous Sciencecraft Experiment on EO-1, *Rem. Sens. Environ*., 101, no. 4, 427-446, 2006.

T. Doggett, R. Greeley, A. G. Davies, S. Chien, B. Cichy, R. Castano, K. Williams, V. Baker, J. Dohm and F. Ip, Autonomous detection of cryospheric change with Hyperion onboard Earth-Observing 1, *Rem. Sens. Environ*., 101, no. 4, 447-462, 2006.

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.

Goddard Space Flight Center, EO-1 Mission page: http://eo1.gsfc.nasa.gov

Goddard Space Flight Center, MODIS instrument page: http://modis.gsfc.nasa.gov

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.

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.

Ip, F., J. M. Dohm, V. R. Baker, T. Doggett, A. G. Davies, R. Castaňo, S. Chien, B. Cichy, R. Greeley, R. Sherwood, Flood detection and monitoring with the Autonomous Spacecraft Experiment (ASE) onboard EO-1, *Rem. Sens. Environ*., 101, no. 4, 463-481, 2006

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.

P. Schut, A. Whiteside, “OpenGIS Web Processing Service”, OGC Public Discussion Paper, OGC 05-007r4.

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.