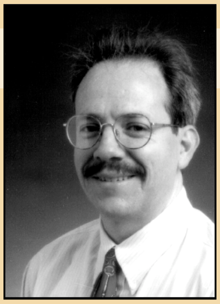

Editor: Richard Doyle
Jet Propulsion Lab
doyle@jpl.nasa.gov

An Autonomous Earth-Observing Sensorweb

Steve Chien, Benjamin Cichy, Ashley Davies, Daniel Tran, Gregg Rabideau, Rebecca Castaño, and Rob Sherwood, *Jet Propulsion Laboratory, California Institute of Technology*
Dan Mandl, Stuart Frye, Seth Shulman, Jeremy Jones, and Sandy Grosvenor, *Goddard Space Flight Center*

Near Antarctica, in a remote area of the South Atlantic Ocean, a volcano rumbles. Following a few minor tremors, fresh lava suddenly breaks to the surface, flowing out of an existing vent. In years past, such an episode might have passed unnoticed or have come to light only days or weeks after the event. Now, thanks to the Earth-observing sensorweb developed by the Jet Propulsion Laboratory and Goddard Space Flight Center, volcanologists around the world will have key science data about eruptions within hours.

The need

Although the South Sandwich Islands are uninhabited, NASA's Terra and Aqua satellites fly overhead four times per day, skimming past at 7.5 kilometers per second and an altitude of 705 kilometers. Each spacecraft carries a Moderate Resolution Imaging Spectroradiometer (MODIS) instrument, which acquires resolution data of 250 to 1,000 meters/pixel about the islands as part of a 2,700-kilometer-wide swath of imagery.

Streamed to Goddard Space Flight Center (GSFC), these data are processed at the Distributed Active Archive Center (DAAC) where MODVOLC (MODIS VOLCano Thermal Alert System) algorithms developed at the University of Hawaii (http://modis.higp.hawaii.edu) automatically detect the volcanic activity's hot-spot signature within hours of data acquisition. Software monitoring the MODVOLC Web site matches this new alert with a previously specified science team interest in volcanoes in this region, generating an observation request to the Earth Observing One (EO-1) ground system.

Based on the request's priority, the ground system uplinks the observation request to the EO-1 spacecraft. Onboard AI 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.

A wide range of operational satellites and space platforms make their data freely available, via either broadcast or the Internet, usually within from tens of minutes to several hours from acquisition. For example, data from the MODIS flying on the Terra and Aqua spacecraft are available via direct broadcast in near-real-time for regional coverage and from 3 to 6 hours from acquisition from Goddard's DAAC for global coverage. These data provide regional or global coverage with a wide range of sensing capabilities: MODIS covers the globe roughly four times daily (two day and two night overflights), while NASA's Quick Scatterometer (QuickSCAT) covers the majority of the globe daily.

Unfortunately, these global-coverage instruments don't provide the high-resolution data many science applications require. Their resolution ranges from 250 m to 1 km for the MODIS instruments to 1 km and above for the other instruments. Ideally, high-resolution data would be available continuously with global coverage. High-resolution assets typically can image only limited swathes of the Earth, making them highly constrained, high-demand assets.

In our sensorweb application, sensors, science event recognizers, and trackers are networked with an automated

Editor's Perspective

The Autonomous Sciencecraft Experiment on NASA's Earth Observing One mission has taken a significant and prominent step forward in validating onboard autonomous systems capability. ASE demonstrates onboard autonomy for both science and engineering functions of the mission by enabling autonomous science event detection and response. For the first time, the study of transient and dynamic scientific phenomena is made available, long assumed to be beyond the reach of spacecraft-based science investigations, particularly those at deep-space destinations. After operating for several months, ASE has passed that ultimate validation criterion: being promoted from a technology experiment to a baseline capability for the ongoing EO-1 mission.

The work described here reports on the application of ASE capability across a fleet of Earth-observing space platforms, further demonstrating the value and potential of this emerging autonomy-based science investigation paradigm.

—Richard Doyle

1541-1672/05/$20.00 © 2005 IEEE
Published by the IEEE Computer Society

Related Work in Sensorweb Research

Considerable effort has been devoted to closed-loop science for rovers at NASA's Ames Research Center, JPL, and Carnegie Mellon.[1–3] These efforts have some similarity in that they have science, execution, and, in some cases, mission-planning elements. However, because surface operations such as rovers are very different from orbital operations, they focus on integration with rover path planning and localization and reliable traverse, whereas our efforts focus on reliable registration of remotely sensed data, interaction with orbital mechanics, and multiple platforms. The Multi-Rover Integrated Science Understanding System also describes a closed-loop multirover autonomous science architecture.[4]

One closely related effort led by Keith Golden at NASA Ames seeks to enable real-time processing of Earth science data such as weather data.[5] However, this work focuses on the problem's information-gathering and data-processing aspect and thus is complementary to our sensorweb work, which focuses on operations. Indeed, we've discussed with Golden the possibility of a joint sensorweb information-gathering demonstration.

The Autonomous Sciencecraft Experiment on EO-1 demonstrates an integrated autonomous mission using onboard science analysis, replanning, and robust execution.[6] The ASE selects and autonomously retargets intelligent science data. ASE represents a single-spacecraft, onboard, autonomous capability. In contrast, the sensorweb uses multiple assets in concert and uses the ASE onboard capability to leverage ground-coordinated requests.

The Remote Agent eXperiment was the first flight of AI software to control a spacecraft.[7] RAX represented a major advance for spacecraft autonomy and operated the Deep Space One mission for several days in 1999. RAX operated DS1 during cruise and therefore performed primarily engineering operations. ASE has flown over a period of over 18 months and has been the primary science and engineering operations system for EO-1 since November 2003 and is expected to continue as such until the end of the EO-1 mission.

References

1. V.C. Gulick et al., "Autonomous Image Analysis during the 1999 Marsrokhod Rover Field Test," *J. Geophysical Research*, vol. 106, no. E4, 2001, pp. 7745–7764.
2. R. Castano et al., "Increased Mission Science Return through In-Situ Data Prioritization," *Proc. IEEE Aerospace Conf.*, IEEE Press, 2003.
3. T. Smith, "Science Autonomy in the Atacama," *Proc. Int'l Conf. Machine Learning Workshop on Machine Learning for Autonomous Space Applications*, 2003; www.lunabots.com/icml2003/docs/Smith_icml_wkshp_2003.pdf.
4. T. Estlin et al., "Coordinating Multiple Rovers with Interdependent Science Objectives," *Proc. 4th Int'l Joint Conf. Autonomous Agents and Multiagent Systems*, 2005, pp. 495–500.
5. K. Golden et al., "Automating the Processing of Earth Observation Data," *Proc. Int'l Symp. AI, Robotics, and Automation in Space,* 2003.
6. S. Chien et al., 'The EO-1 Autonomous Science Agent," *Proc. Autonomous Agents and Multiagent Systems Conf.* (AAMAS 2004), 2004, pp. 420–427.
7. N. Muscettola et al., "Remote Agent: To Boldly Go Where No AI System Has Gone Before," *Artificial Intelligence*, vol. 103, nos. 1–2, Aug. 1998, pp. 5–48.

response system to form a sensorweb. Our approach uses low-resolution, high-coverage sensors to trigger observations by high-resolution instruments (see figure 1). See the "Related Work in Sensorweb Research" sidebar for a discussion of other research activities in this realm.

Sensorweb scenario

As figure 2 shows, components in the EO-1 sensorweb architecture operate as follows:

1. A first asset *Asset1* (such as Modis) acquires data (usually global coverage at low resolution).
2. Data from *Asset1* is downlinked and sent to a processing center where it's automatically processed to detect science events.
3. Science event detections go to a retasking system (labeled "retasking" in the figure), which generates an observation request that's forwarded to an automated planning system. This automated planning system then generates a command sequence to acquire the new observation.
4. This new command sequence is uplinked to *Asset2* (for example, EO-1), which then acquires the high-resolution data.
5. *Asset2* then downlinks the new science data. On the ground this data is processed and forwarded to the interested science team.

To date, Asset2 has been EO-1, the first satellite in NASA's New Millennium Program Earth Observing series. EO-1's primary focus is to develop and test a set of advanced-technology land-imaging instruments.

EO-1 launched from Vandenberg Air Force Base on 21 November 2000. Its orbit allows for 16-day repeat tracks, with at least five overflights per cycle, with a change in viewing angle less than 10°. Because EO-1 is in a near-polar orbit, it can view polar targets more frequently.

EO-1 has two principal science instruments, the Advanced Land Imager and the Hyperion hyperspectral instrument. ALI is a multispectral imager with 10-m/pixel pan-

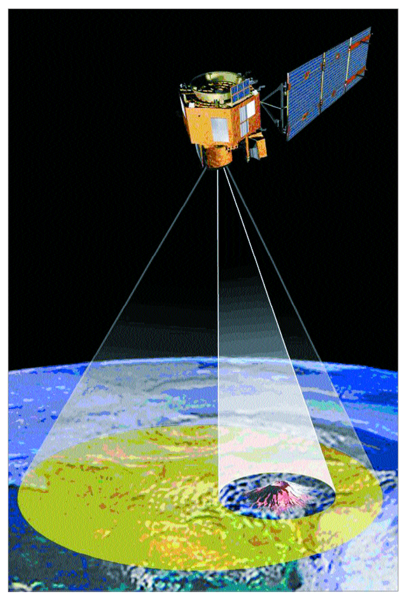

Figure 1. Sensorweb applications involve a networked set of instruments in which information from one or more sensors automatically serves to reconfigure the remainder of the sensors.

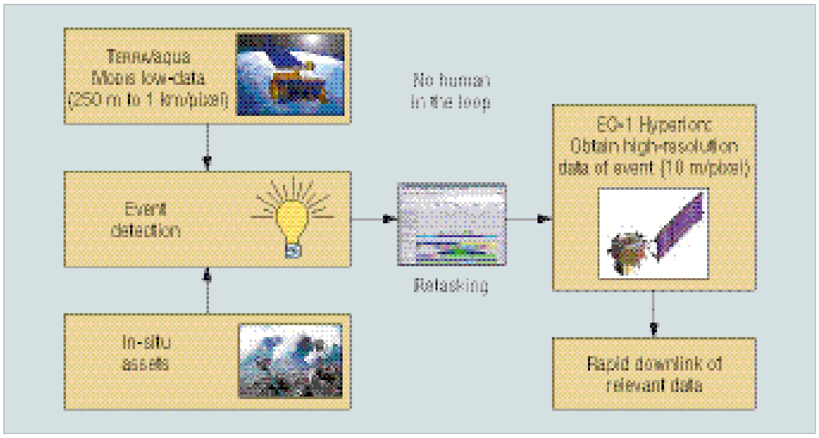

Figure 2. Sensorweb event detection and response architecture.

band resolution and nine spectral bands from 0.433 to 2.35 µm with 30-m/pixel resolution. ALI images a 37-km-wide swath. Hyperion is a high-resolution imager that can resolve 220 spectral bands (from 0.4 to 2.5 µm) with a 30-m/pixel spatial resolution. The instrument images a 7.5 by 42 km land area per image and provides detailed spectral mapping across all 220 channels with high radiometric accuracy.

EO-1 sensorweb architecture

As figure 3 illustrates, components in the sensorweb's automated retasking element work together as follows.

Science-tracking systems for each science discipline 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, each in their own format—in some cases through the HTTP or FTP protocol, in others via email subscription and alert protocols.

Science agents either poll these sites (HTTP or FTP) to pull science data or simply receive emails to be notified of ongoing science events. These science agents then produce science event notifications in a standard XML format, which sensorwebs log into a science event database.

The science event manager (SEM) processes these science event notifications and matches them with science campaigns, generating an observation request when a match occurs. The Aspen automated mission-planning system processes these requests, integrating them with already scheduled observations according to priorities and mission constraints. If a new request can fit within the existing schedule without removing a higher priority observation, an observation request is uplinked to the spacecraft. For an observation to fit within the schedule, the spacecraft must be able to acquire the observation without violating spacecraft constraints such as having adequate power for the observation, having a time of the observation that does not conflict with a higher-priority observation, or that there is adequate storage for the data onboard.

Onboard EO-1, the Autonomous Sciencecraft software will accommodate the observation request if feasible.[1] In some cases, onboard software might have additional knowledge of spacecraft resources or might have triggered additional observations, making some uplinked requests infeasible. Later, the spacecraft downlinks the science data to ground stations where it is processed and delivered to the requesting scientist.

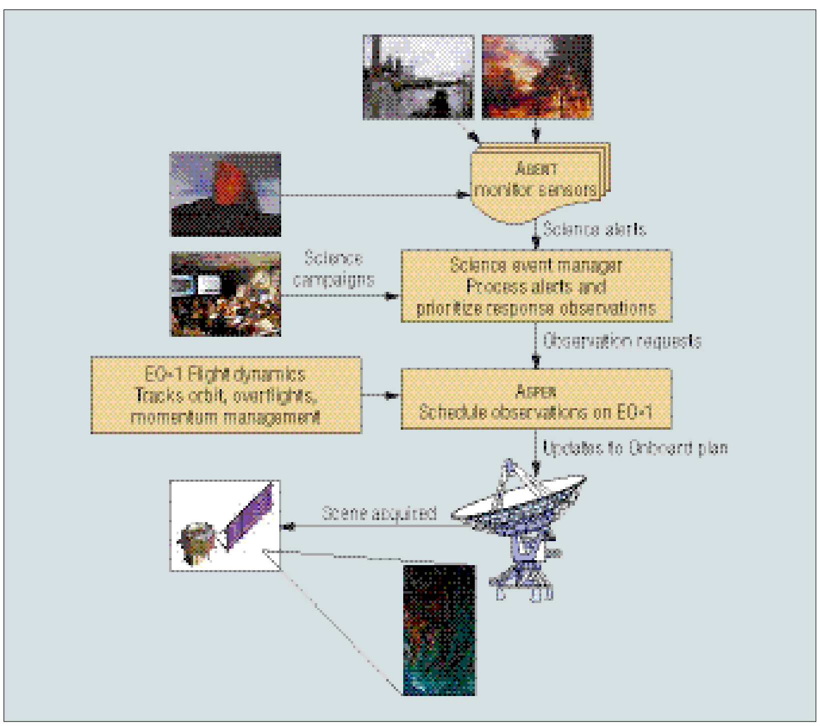

Figure 3. Sensorweb response, showing how automated retasking elements work together.

Event tracking and observation request generation

The science agents encapsulate sensor- and science-tracking-specific information by producing a generic XML alert for each science event tracked. The flexibility enabled by these modules lets users easily integrate with numerous science tracking systems even though each one has its own unique data and reporting format. These formats range from near-raw instrument data to alerts in text format, to periodic updates to a wide range of text formats. The posting methods have included HTTP, HTTPS, FTP, and email. Table 1 lists the science-tracking systems integrated into our system.

The science event manager lets scientists specify mappings from science events to observation requests, track recency and event counts, and perform logical processing—for example, triggering an observation if two Modvolc alerts and a Goesvolc alert occur in a 24-hour period. The SEM also permits tracking based on target names or locations and other event-specific parameters.

As an example, because the Kilauea volcano is often quite active, a volcanologist there might specify that several tracking systems would need to report activity with high confidence before an observation is requested. On the other hand, even a single low-confidence activity notification might trigger observation of Piton de la Fournaise or other less active sites.

Table 1. Science alert systems.

Discipline	Source	Detector
Volcanoes	Modis (Terra, Aqua)	Modvolc, Univ. of Hawaii
	GOES	GOESVolc
	POES	Avhrr—Volcano
	Air Force Weather Advisory	Volcanic ash alerts
	International FAA	Volcanic ash advisories
	Tunguratua, Eventador	In situ instruments, Harvard, UNH
	Hawaiian Volcano Observatory	Sensor alerts
	Rabaul Volcano Observatory	
Floods	QuikSCAT	Dartmouth Flood Observatory
	Modis	Dartmouth Flood Observatory
	AMSR	Dartmouth Flood Observatory
Cryosphere	QuikSCAT	Snow/ice, JPL
	Wisconsin Lake Buoys	UW Dept. Limnology
Forest fires	Modis (Terra, Aqua)	Rapidfire, UMD Modis, Rapid Response
Dust storms	Modis (Terra, Aqua)	Naval Research Laboratory, Monterey

Event response: Automated observation planning

To automate mission planning, we use the Aspen/Casper planning and scheduling system (Aspen is the ground-based batch planner and Casper is the embedded, flight-based planner; both share the same core planning engine).[2] Aspen represents mission constraints in a declarative format and searches possible mission plans for a plan that satisfies many observation requests (respecting priorities) and also obeys mission operations constraints. Aspen has served in a wide range of space mission applications, including spacecraft operations scheduling, rover planning, and ground communications station automation.

Aspen: Local, committed search for planning

Search in Aspen has focused on high-speed local search in a committed plan space, using a stochastic combination of a portfolio of heuristics for iterative repair and improvement algorithms.[3–5] In this approach, at each choice point in the iterative repair process, a stochastic choice is made by Aspen among a portfolio of heuristics (with probabilities the user can specify).[6] This approach has performed well in a wide range of applications.[2] The stochastic element combined with a portfolio of heuristics helps to avoid the typical pitfalls of local search. Using a committed plan representation enables fast search moves and propagation of effects (100 s of operations per CPU second on a workstation). To increase efficiency, we also use aggregates of activities.[7]

We've focused on an early-commitment, local, heuristic, iterative search approach to planning, scheduling, and optimization. This approach has several desirable properties for spacecraft operations planning.

First, using an iterative algorithm lets us use automated planning at any time and on any given initial plan. The initial plan might be as incomplete as a set of goals, or it might be a previously produced plan with only a few flaws. Repairing and optimizing an existing plan enables fast replanning when necessary from manual plan modifications or from unexpected differences detected during execution. Local search planning thus can have an anytime property, in which it always has a "current best" solution and improves it as time and other resources allow. Refinement search methods don't have this property.[8] Local search can also easily adapt for use in a "mixed initiative" mode for partial ground-based automation.

Also, it's easier to write powerful heuristics that evaluate ground plans. These strong heuristics let us prune the search, ruling out less promising planning choices.

Third, a local algorithm doesn't incur the overhead of maintaining intermediate plans or past attempts. This feature lets the planner quickly try many plan modifications for repairing conflicts or improving preferences. However, unlike systematic search algorithms, we cannot guarantee that our iterative algorithms will explore all possible combinations of plan modifications or that it will not retry unhelpful modifications. In our experience, these guarantees are not valuable because for large-scale problems complete search is intractable.

Finally, by committing to values for parameters, such as activity start times and resource usages, Aspen can efficiently compute effects of a resource usage and the corresponding resource profiles. Least-commitment techniques retain plan flexibility but can be computationally expensive for large applications.[9]

The sidebar "Unique Challenges of EO-1 Sensorweb Mission Planning" discusses some of the challenges we faced in adapting Aspen for the EO-1 sensorweb.

Science data access

A sensorweb project goal is to provide scientists easy access to multiple data sources on a single science event, such as a volcanic eruption or a forest fire (see the "Sensorweb Examples" sidebar). This data access portal for the sensorweb project is still under construction.

Another goal of the sensorweb effort was to enable easy tracking of spacecraft operations. This tracking would let scientists understand the images the spacecraft had acquired and view where science products are in request, acquisition, downlink, and processing phases. To accomplish this goal, we have an operational Web site for science team access. We will shortly make this site publicly available.

Ongoing extensions and deep space applications

Terrestrial dust storms are of significant science interest and can be detected using several sensors, including GOES, Avhrr, and Modis.[10] Growing to be as large as hundreds of kilometers long, these storms are important because of the amount of dust they can transport and their

Unique Challenges of EO-1 Sensorweb Mission Planning

The EO-1 sensorweb application presented a number of interesting challenges for automated planning, including prioritization, file system modeling, momentum management and maneuvers, and coordination of planners.

Priorities

In prioritization, the sensorweb application requires that the mission-planning element reason about relative priorities on observations as well as how their supporting activities relate to the goal observation priority. Within the EO-1 mission operations, we developed a strictly ordered set of priorities. In this scheme, each observation is assigned a value from 1 to 1,000 (with lower values denoting higher priority). Different user types can submit observation requests within an allotted range of priorities, with many of the users' ranges overlapping: a high-priority observation from user 1 might preempt a low-priority observation from user 2, but not a high priority observation from user 2. Aspen respects these priorities by the nature of the encoded search heuristics.

These search strategies first prefer plan repair operators that don't delete observations. However, if forced to delete observations, these heuristics prefer deleting lower-priority observations. In this scheme, priority levels are strictly dominating. For example, one observation of priority 500 will be preferred to two observations of priority 700.

File system

For file system modeling, one degree of scheduling flexibility involves separating science data processing from data acquisition. After a scene is acquired onboard, sensorweb can analyze it to detect science events (this applies only for certain types of science images: volcanoes, floods, and cryosphere). Sensorweb can then rapidly downlink this event summary to give scientists a snapshot of the activity; the complete image will take longer to downlink and process.

For example, if an image A and an image B are *X* minutes apart, there isn't enough time to process the science data from image A before the imaging of B. Playing back the image from the solid-state recorder (where it's streamed during data acquisition) into RAM to analyze it requires use of the SSR; acquiring image B also requires use of the SSR. However, image A must be analyzed prior to downlink and file deletion on the SSR.

Representing this properly within Aspen requires the ability to model a file system, which we've demonstrated in Aspen in the Generalized Timelines module. This capability isn't part of the core Aspen that has been used in many applications. To reduce risk, we require that images be analyzed before the next image is acquired and represent this as a simple protection in planning terminology. This representation decision slightly reduces the efficiency of Aspen-generated plans.

---

impact on aviation. A dust storm sensorweb would use low-resolution assets to track large-scale dust storms and autonomously direct high-resolution assets such as EO-1 to acquire more detailed data. Such data would improve scientific understanding of dust initiation and transport phenomena.

Figure 4 shows a large dust storm in the Persian Gulf as imaged by Modis in November 2003. Ground-based instrumentation—such as operated by the US Department of Agriculture in the American Southwest and the People's Republic of China's network of sites in the Gobi Desert—can also serve to detect these storms. Detection and tracking of dust storms is also of considerable interest on Mars where such storms can grow to cover the entire planet.

The sensorweb concept also applies directly to deep-space science applications, Sun-Earth connection science, and astrophysics applications.[11] On Mars, for example, surface instruments could detect or track active, transient atmospheric, and geologic processes such as dust storms. Already in place on Mars is a wide range of complementary assets. The Mars Global Surveyor (MGS) spacecraft is flying the Thermal Emission Spectrometer, which can observe dust storms at a global scale. We could use this instrument to detect dust storm initiation events, calling in higher-resolution imaging devices including Themis on Mars Odyssey, the MGS-MOC camera on MGS, and the HiRise camera on Mars Reconaissance Orbiter (scheduled to arrive in 2006).

Additionally, we could integrate surface assets such as the Mars Exploration Rovers (currently deployed) and the Phoenix Lander (2007) into a Mars sensor network. In the future, we expect even more assets to be on Mars, enabling even more integrated observation campaigns. An integrated network of Mars instruments is particularly critical to support extended Mars surface missions (particularly manned missions) as envisioned by the new NASA exploration initiative.

The automated sensorweb concept has broad application beyond planetary science. For example, in space weather, sun-pointed instruments could detect coronal mass ejections and alert Earth-orbiting magnetospheric instruments to reconfigure to maximize science data. This solar activity would also have ramifications on manned exploration on the Moon or Mars. Such an automated tracking system would be critical to ensuring the safety of such missions.

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 acknowledge the important contributions of Nghia Tang and Michael Burl of JPL, Thomas Brakke, Lawrence Ong, and Stephen Ungar of GSFC, Jerry Hengemihle and Bruce Trout of Microtel LLC, Jeff D'Agostino of the Hammers

Momentum management and maneuvers

Maneuver and momentum management presents particular challenges to EO-1 operations. EO-1 uses reaction wheels to point the spacecraft in imaging targets and downlinking data. Because EO-1 has only three reaction wheels, it must use magnetic torquers and the Earth's magnetic field to dump momentum from the reaction wheels. Without this process, the reaction wheels might either be unable to achieve the desired attitude (because they're already spinning as fast as they can, called *momentum saturation*) or a wheel might change spin direction (*zero crossing*) during an image, which causes jitter in the spacecraft and ruins the image.

However, desaturating the wheels (*zero biasing*) is time-consuming, so often there isn't enough time between images to perform this step. ASPEN attempts to acquire all images with zero biasing before and after each image. Still, if zero biasing for an image will not fit, ASPEN will still acquire the image: acquiring an image with only a small chance of being ruined is better than not acquiring the image at all. ASPEN implements this strategy in its search heuristics. If it has difficulty fitting an image into the schedule, it will first consider removing the momentum management activities for the lowest-priority image participating in the conflict. If it can still fit the image without these activities, it will retain the image.

This approach leads to another complication: the association of myriad activities the image requires. If after searching, ASPEN determines that an observation will not fit, it must remove all the associated activities. It does so by annotating them with the scene ID of the image requiring their presence in the plan and cleaning up the plan appropriately when the image is removed. As another complication, parameters of the momentum management activities depend on how ASPEN handled the immediately preceding scene, as that scene determines the reaction wheels' initial momentum. ASPEN handles this eventuality with an explicit dependency between the momentum management activities for an observation and the momentum state at the end of the prior observation. Again, this could be directly modeled in the Generalized Timelines module, but for expedience we modeled it with a state and parametric dependency.

Coordinating planners

Another challenge of sensorweb planning is the coordination of the onboard and ground planners. In the sensorweb, the onboard planner might have changed the plan since uplink because of execution variances (such as additional images being scheduled). To correctly handle this eventuality, the ground planner guesses, on the basis of the previously uploaded plan, whether a scene will be possible. If it is, the goal requesting the observation uploads. The onboard planner then receives this request and might add the observation (and any necessary supporting activities) to the onboard plan, deleting scenes as required (consistent with scene priorities). Because this replanning might require considerable search and hence onboard computing time—and the maneuvers commence 45 minutes prior to a scene—the ground ASPEN only uplinks observation requests that occur at least two hours after the uplink window.

Corp., Robert Bote of Honeywell Corp., Jim Van Gaasbeck and Darrell Boyer of ICS, Michael Griffin and Hsiao-Hua Burke of MIT Lincoln Labs, Ronald Greeley and Thomas Doggett of Arizona State University, and Victor Baker and James Dohm of the University of Arizona.

References

1. S. Chien et al., "The EO-1 Autonomous Science Agent," *Proc. Autonomous Agents and Multiagent Systems Conf. (*AAMAS 2004), ACM Press, 2004, pp. 420–427.

2. S. Chien et al., "Aspen: Automating Space Mission Operations Using Automated Planning and Scheduling," *Proc. 6th Int'l Conf. Space Operations,* 2000; http://ai.jpl.nasa.gov.

3. Huberman et al., "An Economics Approach to Hard Computational Problems," *Science*, vol. 275, 1997, pp. 51–54.

4. C. Gomes and B. Selmon "Algorithm Portfolios," *Artificial Intelligence*, vol. 126, 2001, pp. 43–62.

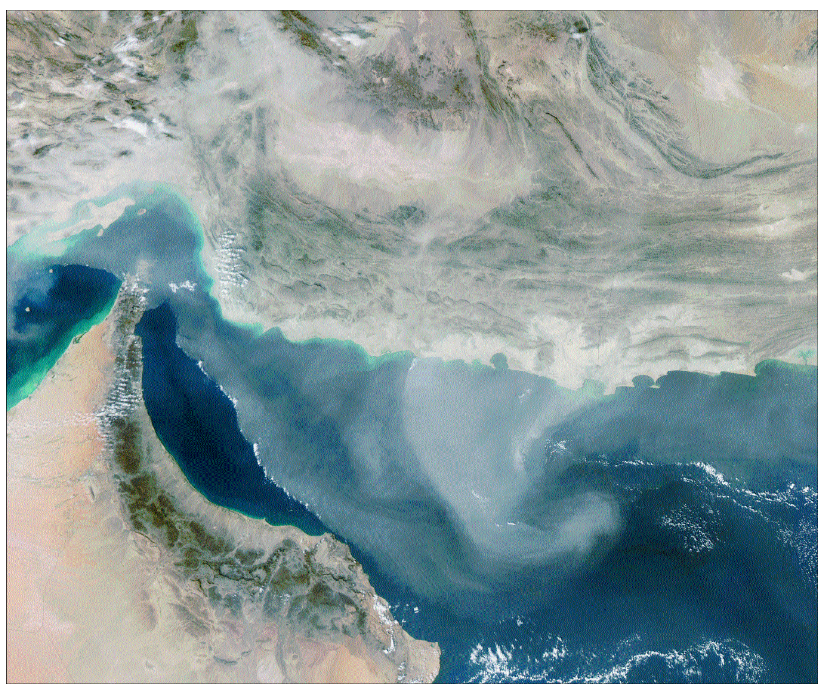

Figure 4. Dust storm in the Persian Gulf as captured by MODIS in November 2003.

Sensorweb Examples

Sensorweb has been used in such applications as wildfire and flood control, volcanology, and cryosphere monitoring.

The wildfire sensorweb

We have demonstrated the sensorweb concept using the MODIS active Fire Mapping System.[1] Both the Terra and Aqua spacecraft carry the MODIS instrument, providing morning, afternoon, and two night overflights of each location on the globe per day (coverage near the poles is even more frequent). The active fire mapping system uses data from the GSFC Distributed Active Archive Center (DAAC), specifically the data with the predicted orbital ephemeris, which is approximately three to six hours from acquisition.

The active fire mapping algorithm detects hot spots using MODIS thermal bands with absolute thresholds:

$T4 > 360K, 330K$ (night) or
$T4 > 330K, 315K$ (night) and
$T4 - T11 > 25K, 10K$ (night)

where T4 is the fourth band of the data and t11 is the eleventh band. This algorithm also uses a relative-threshold algorithm that requires six nearby cloud-, smoke-, water-, and fire-free pixels up to 21×21 square. This triggers if the thermal reading is three standard deviations above the surrounding area.

$T4 > \text{mean}(T4) + 3 \text{ std dev } (T4)$
and $T4 - T11 > \text{median } (T4 - T11) + 3$
$\text{std dev } (T4 - T11)$

Figure A shows the active fire map from October 2003 fires in Southern California. Figure B shows a context map of Southern California with the triggered sensorweb observation taken below it.

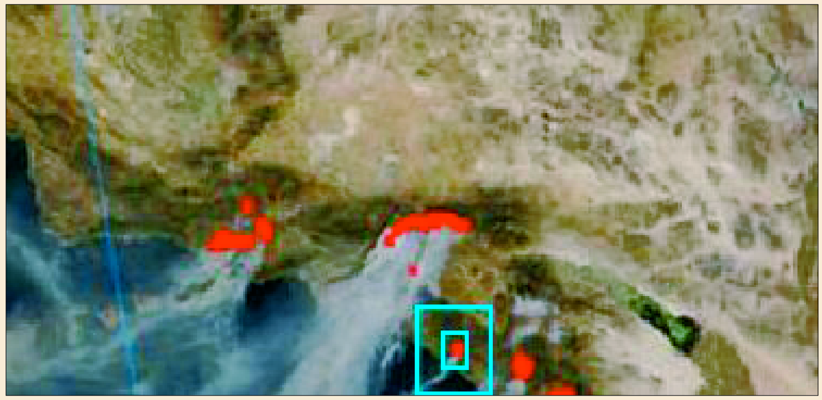

Figure A. Active fire alerts for the October 2003 Southern California fires. Red indicates active fires. The light blue box illustrates the background region used in the relative threshold detection.

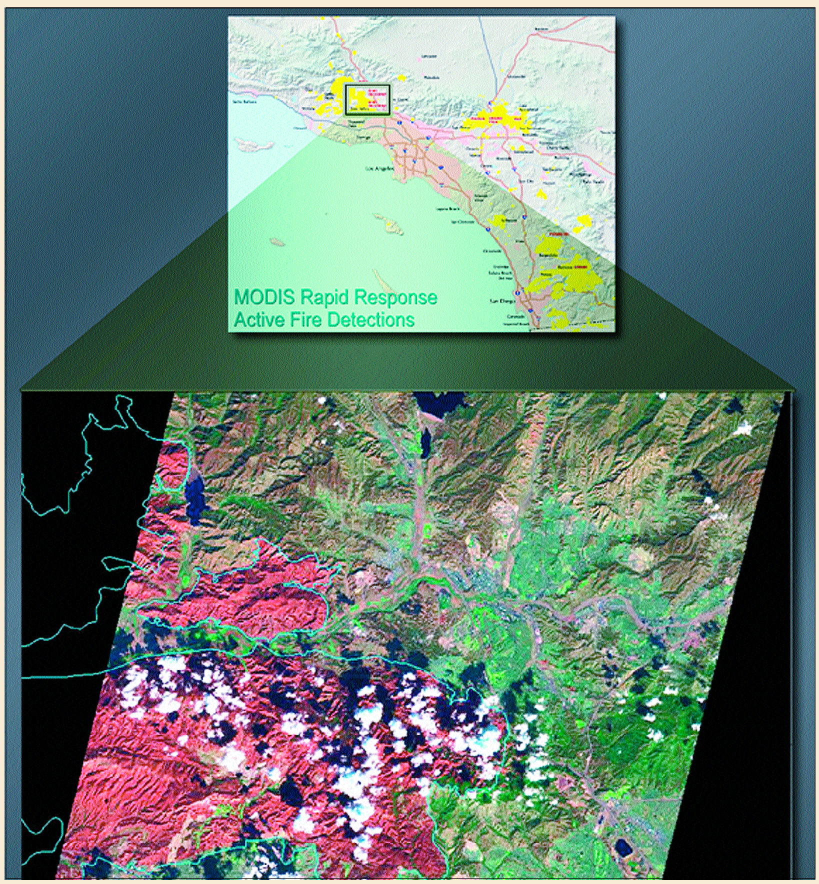

Figure B. Sensorweb trigger images for the October 2003 Southern California fires. Above is the MODIS Active Fire Map display. Below is the EO-1 Hyperion image acquired via sensorweb trigger of the Simi/Val Verde fire area used in Burned Area Emergency Reclamation.

The flood sensorweb

The flood sensorweb uses the Dartmouth Flood Observatory Global Active Flood Archive to identify floods in remote locations automatically, based on satellite data. The DFO flood archive generates flood alerts based on MODIS, QuikSCAT, and AMSR-E satellite data.[2] The DFO produces the DFO archive in collaboration with JPL. The flood sensorweb uses the DFO QuikSCAT atlas because it's not affected by cloud cover over flooded areas.

The DFO produces the DFO archive in collaboration with the JPL QuikSCAT team. In this process, the QuikSCAT scatterometer data help us assess surface water conditions.[3,4] Specifically, the VV/HH ratio serves to assess surface water properties of the areas in 0.25 latitude/longitude degree bins. The sensorweb uses the seven-day running mean to dampen effects of short-duration rainfall over urban areas. It then compares these data to the seasonal (90-day) average of the previous year season to screen out seasonal wetlands, publishing the screened alerts to a DFO Web site. Figure C shows an example of a global flood alert.

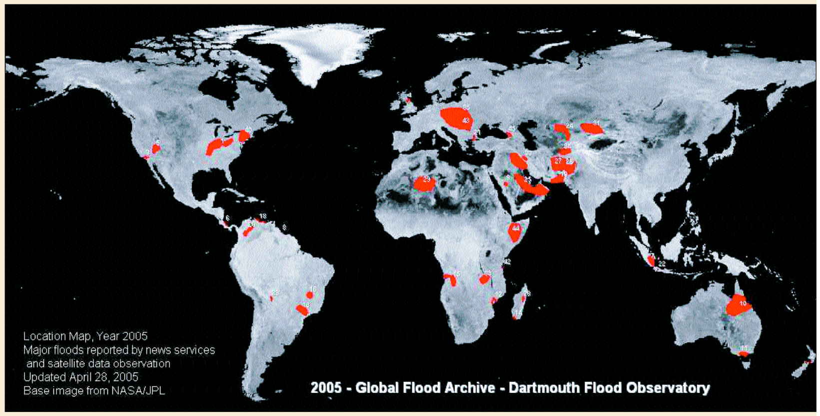

Figure C. Dartmouth Flood Observatory global flood alerts for October 2003.

In the flood sensorweb, indications of active flooding alerts trigger EO-1 observations at sites called *gauging reaches*. These are river locations whose topography is well understood. We can use flood discharge measurements at gauging reaches to measure the amount of water passing through a flooded region, comparing them with remotely sensed data. Ultimately, the flood sensorweb increases the amount of high-resolution remote sensing data available on flooding events in prime locations of interest (gauging reaches) and times of interest, such as when active flooding occurs. Figure D shows imagery from an August 2003 flood sensorweb demonstration capturing flooding in India's Brahmaputra River.

The volcano sensorweb

In the volcano sensorweb, MODIS, GOES, and AVHRR sensor platforms operate to detect volcanic activity. These alerts then trigger EO-1 observations. The EO-1 Hyperion instrument is ideal for studying volcanic processes because of its great sensitivity range in the infrared spectrum.

The GOES and AVHRR alert systems provide excellent temporal resolution and rapid triggering based on thermal alerts.[5] The GOES-based system looks for locations that are hot, have high contrast from the surrounding area, and are not visibly bright. Additionally, the system screens hits for motion (to eliminate cloud reflections) and persistence (to remove instrument noise). The GOES alert can provide a Web or email alert within one hour of data acquisition.

The MODIS alert system offers high instrument sensitivity but lower temporal resolution (MODIS generally has at least four overflights per day). MODVOLC derives the normalized thermal index (NTI) from MODIS raw radiance values by computing $(R_{22} - R_{32})/(R_{22} + R_{32})$, where R_i indicates use of the radiance value from MODIS band *i*. The system compares the NTI to a threshold to indicate alerts, generally making it available online within three to six hours of acquisition. We've also linked into in situ sensors to monitor volcanoes. We are working with a number of teams to integrate such sensors into our sensorweb. The Hawaiian Volcano Observatory has deployed numerous instruments in Hawaii's Kilauea region. These instruments include tiltmeters, gas sensors, and seismic instrumentation. These sensors can provide indications that collectively point to a high-probability near-term eruption, thereby triggering a request for high-resolution EO-1 imagery. The University of Hawaii has also deployed infrared cameras to a number of volcanic sites worldwide,

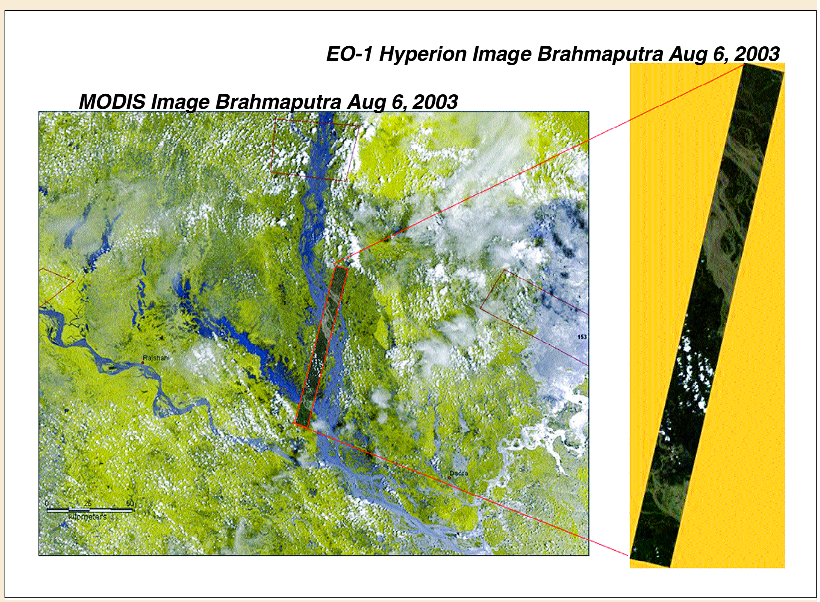

Figure D. Examples of low-resolution MODIS imagery (left) and high-resolution EO-1 imagery (right) from the Flood Sensorweb capturing Brahmaputra River flooding in India, August 2003.

including Kilauea, Hawaii; Erte Ale, Ethiopia; Sourfiere Hills, Montserrat; and Colima and Popocatepetl, Mexico.[7] These infrared cameras can provide a ground-based detection of lava flows based on thermal signatures, also alerting the sensorweb.

Cryosphere sensorweb

The Earth's cryosphere consists of freezing water in the form of snow, lake and sea Ice, and the corresponding thawing of these. Because it plays such a central role in creating the Earth's climate, a wide range of scientists are interested in studying the cryosphere. Planetary scientists also want to study cryosphere phenomena on other planets in the solar system; studying the Earth's cryosphere is an analogue for other planets' cryospheres and vice versa.

Using the EO-1 sensorweb, we can study numerous phenomena, including glacial ice breakup; sea ice breakup, melting, and freezing; lake ice freezing and thawing; and snowfall and snowmelt.

Using QuikSCAT data, we're tracking snow and ice formation and melting, automatically triggering higher-resolution imaging such as with EO-1. In collaboration with the University of Wisconsin–Madison's Center for Limnology, we have also linked into data streams from the Trout Lake stations so that temperature data can trigger imaging of the sites to capture transient freezing and thawing processes. These linkages let us request high-priority imaging from EO-1 to study short-lived—thaw, melt, breakup, snowfall, or ice formation—cryosphere phenomena.

References

1. C. Justice et al., 'the Modis Fire Products," *Remote Sensing of Environment*, vol. 83, 2002, pp. 244–262.
2. S.V. Nghiem, *Advanced Scatterometry for Geophysical Remote Sensing*, JPL Document D-23048, Jet Propulsion Lab, Calif. Inst. Tech., Mar. 2001.
3. B. Brakenridge et al., "Flood Warnings, Flood Disaster Assessments, and Flood Hazard Reduction: The Roles of Orbital Remote Sensing," *Proc 30th Int'l Symp. Remote Sensing of Environment,* 2003; http://isrse.pdc.org.
4. S.V. Nghiem, W.T. Liu, and X. Xie, "Polarization Reversal over Flooded Regions and Applications to Flood Mapping with Spaceborne Scatterometers," *Proc. Int"l Geoscience and Remote Sensing Symp.,* IEEE Computer Society Press, 1999; www.ewh.ieee.org/soc/grss/igarss.html.
5. A. Harris et al., "Web-Based Hot Spot Monitoring Using GOES: What It Is and How It Works," *Advances in Environmental Monitoring and Modelling*, vol. 1, no. 3, 2002, pp. 5–36.
6. R. Wright et al., "Automated Volcanic Eruption Detection Using MODIS," *Remote Sensing of Environment*, vol. 82, 2002, 135–155.
7. A. Harris et al., "Ground-Based Infrared Monitoring Provides New Tool for Remote Tracking of Volcanic Activity," *EOS*, vol. 84, no. 40 2003; www.agu.org/pubs/eos.html.

Steve Chien is technical group supervisor of the AI Group at JPL, where he leads efforts in research and development of automated planning and scheduling systems. He is the principal investigator of the Earth Observing Sensorweb Project and the Autonomous Sciencecraft Experiment. He's also an adjunct associate professor in the Department of Computer Science at the University of Southern California. Contact him at steve.chien@jpl.nasa.gov.

Benjamin Cichy is a member of the technical staff in JPL's Flight Software Development and Technology Group, where he develops, autonomy software for NASA's New Millenium Program. Contact him at benjamin.cichy@jpl.nasa.gov.

Ashley Davies is a research scientist in JPL's Earth and Space Sciences Division and lead scientist of the Earth Observing Sensorweb. Contact him at ashley.davies@jpl.nasa.gov.

Daniel Tran is a member of the technical staff in JPL's AI Group, where he's working on developing automated planning and scheduling systems for onboard spacecraft commanding. Contact him at daniel.tran@jpl.nasa.gov.

Gregg Rabideau is a member of the technical staff in JPL's AI Group, where he researches and develops planning and scheduling systems for automated spacecraft commanding. Contact him at gregg.rabideau@jpl.nasa.gov.

Rebecca Castaño supervises the Machine Learning Systems group at JPL and is also the team lead of the Onboard Autonomous Science Investigation System (OASIS). Contact her at rebecca.castano@jpl.nasa.gov.

Rob Sherwood is manager, Earth Science Information Systems. Contact him at rob.sherwood@jpl.nasa.gov.

Dan Mandl is the mission director for the Earth Observing One Mission At Goddard Space Flight Center. Contact him at dan.mandl@gsfc.nasa.gov.

Stuart Frye is the Lead Systems Engineer for the Earth Observing One Mission at Goddard Space Flight Center. Contact him at stuart.frye@gsfc.nasa.gov.

Seth Shulman is the Operations Lead for the Earth Observing One Mission at Goddard Space Flight Center. Contact him at seth.e.shulman.1@gsfc.nasa.gov.

Jeremy Jones was a software engineer at Goddard Space Flight Center.

Sandy Grosvenor is a software engineer in the Advanced Architectures and Automation Branch of Goddard Space Flight Center. Contact her at sandy.grosvenor@gsfc.nasa.gov.

5. M. Zweben et al., "Scheduling and Rescheduling with Iterative Repair," *Intelligent Scheduling,* Morgan Kaufmann, 1994.
6. G. Rabideau et al., "Iterative Repair Planning for Spacecraft Operations in the Aspen System," *Int'l Symp. AI Robotics and Automation in Space*, ESA Press, June 1999.
7. R. Knight et al., "Computing Valid Intervals for Collections of Activities with Shared States and Resources, *Int'l Conf. AI Planning Sys.,* AAAI Press, 2000.
8. A. Jonsson et al., "Planning in interplanetary Space: Theory and Practice," *Int'l Conf. AI Planning Sys.*, AAAI Press, 2000, pp. 177–186.
9. S. Chien, et al., "Automated Planning and Scheduling for Goal-Based Autonomous Spacecraft," *IEEE Intelligent Systems*, vol. 13, no. 5, 1998, pp. 50–55.
10. S. Miller, "A Consolidated Technique for Enhancing Desert Dust Storms with MODIS," *Geophysical Research Letters,* vol. 30, no. 20, 2003, pp. 2071–2074.
11. A.G. Davies et al., *ASC Science Report*, Jet Propulsion Lab., Aug. 2001; http://ase.jpl.nasa.gov.