Using Automated Planning for Sensorweb Response

Steve Chien, Ashley Davies, Daniel Tran, Benjamin Cichy,

Gregg Rabideau, Rebecca Castano, Rob Sherwood
Jet Propulsion Laboratory, California Institute of Technology
Jeremy Jones, Sandy Grosvenor, Dan Mandl, Stuart Frye[1], Seth Shulman[4],
Stephen Ungar, Thomas Brakke, Jacques Descloitres
Goddard Space Flight Center

Chris Justice, Rob Sohlberg, University of Maryland

Rob Wright, Luke Flynn, Andy Harris, University of Hawaii

Robert Brakenridge, Sebastien Cacquard, Dartmouth Flood Obervatory   Son Nghiem JPL/Quikscat

Ronald Greeley, Thomas Doggett, Arizona State University

Victor Baker, James Dohm, Felipe Ip, University of Arizona

Contact: steve.chien@jpl.nasa.gov

Abstract— A recent explosion in availability of timely satellite data has resulted in operational systems that track wildfires, floods, and volcanic activity worldwide. This paper describes efforts to link these science event detection systems with an automated response system to retarget remote sensing assets to observe these important but transient science events. Of course, automated mission planning is a key element of the overall tracking and response system. We describe the current prototype system which utilizes the Earth Observing One spacecraft, MODIS flying on Terra and Aqua, QuickSCAT, GOES, and AVHRR platforms as well as future plans for expansion.

1. INTRODUCTION

In a remote area of southern Africa, lightning strikes grassland, sparking a local fire that rapidly spreads across the dry, grassy region. Overhead, the Terra & Aqua satellites, each using a MODIS instrument, acquire moderate resolution (250m) data of every point on the Earth twice per day. These data are streamed to the Rapidfire center at the University of Maryland and Goddard Space Flight Center where raw imagery is automatically classified into fire alerts within hours of acquisition. Software monitoring the Rapidfire web site matches this new alert with a previously specified science team interest in fires in this region and generates an observation request to the Earth Observing One (EO-1) Ground System. This observation request is processed within a ground system automation system for EO-1 to develop a command sequence to acquire the observation. After the commands are uplinked, at the appropriate time, the spacecraft slews and acquires the high resolution (pan-band up to 10m) image with hyperspectral (220+ bands) data for science analysis. This autonomous response has enabled detailed follow-up data on the science event of interest within 48 hours of the initial detection. Additionally, the science return of EO-1 is optimized by targeting known science events.

        --The EO-1 sensorweb has demonstrated this and similar scenarios since its first operations in August 2003.

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 the Moderate Resolution Imaging Spectrometer (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) (for global coverage). These 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). QuickSCAT covers the majority of the globe daily.

Satellite/ Platform	Instrument	Overflight Frequency	Data timeliness
Terra, Aqua	MODIS	Every 12 hours day, night	3-6 hours global, near real-time regional
Quikscat	Scatterometer radar	~1 day	daily
NOAA-POES	AVHRR	Variable, frequent	< 1 hour
GOES	Visible, Infra-red	continuous	10s of minutes

Table 1: Timely satellite data streams

Unfortunately, these global coverage instruments do not provide the high resolution data desirable for many science applications. The above instruments range in resolution

---

1 - MitreTek, 2 - Microtel LLC, 3 – the Hammers Company, 4- Honeywell.

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 are constrained in one of two ways:

- They are nadir pointing which means infrequent overflights due to limited swath. For example, Landsat-7 provides high resolution (15m/pixel) imagery, but its overflights are 16 days apart.
- They are point and shoot which means that they can only observe relatively small, pre-designated areas; thus competing targets in the same portion of an orbit often cannot both be taken.

In this paper we describe initial efforts to network sensors and science event recognizers/trackers with an automated response system to form a sensorweb, defined as follows.

Sensorweb A networked set of instruments in which information from one or more sensors is *automatically* used to reconfigure the remainder of the sensors

Specifically, in our application, In our application we use low resolution, high coverage sensors to trigger observations by high resolution instruments[2].

In the remainder of the paper we first describe our preliminary sensorweb efforts to track

- Wildfires
- Floods, and
- Volcanoes.

We also describe ongoing efforts to expand the sensorweb to other remote sensing and in-situ assets.

2. THE EO-1 SENSORWEB ARCHITECTURE

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)
2. Data from Asset1 is downlinked
3. This data is automatically processed to detect science events

---

[2] Note that there are many other rationales to network sensors into a sensorweb. For example automated response might enable observation using complementary instruments such as imaging radar, infra-red, visible, etc. Or automated response might be used to apply more assets to increase the frequency of observation to improve the temporal resolution of available data.

4. Science event detections are forwarded to a re-tasking system. This system generates an observation request which is forwarded to an automated planning system.
5. This automated planning system then generates a command sequence to acquire the new observation.
6. This new command sequence is uplinked to Asset2 which then acquires the high resolution data.
7. This data is then downlinked, processed, and forwarded to the interested science team.
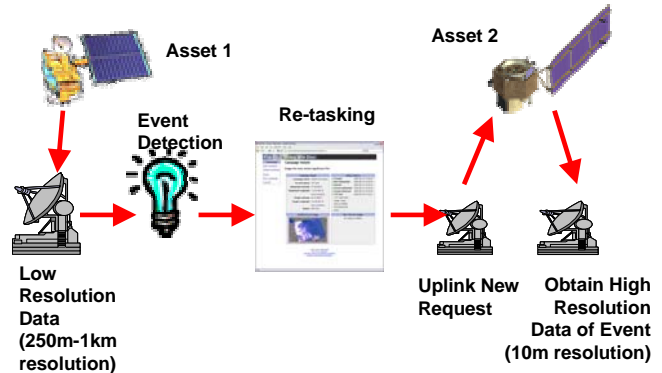

Figure 1. Sensorweb Detection and Response Architecture

In our implemented demonstrations thus far Asset2 has been the Earth Observing One 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 on a Delta 7320 from Vandenberg Air Force Base on November 21, 2000. It was inserted into a 705 km circular, sun-synchronous orbit at a 98.7 degrees inclination. This orbit allows for 16-day repeat tracks, with 3 over flights per 16-day cycle at a less than 10-degree change in viewing angle. Because EO-1 is in a near polar orbit, polar targets can be viewed more frequently.

For each scene, over 20-Gbits of data from the Advanced Land Imager (ALI), Hyperion, and Atmospheric Corrector (AC) are collected and stored on the onboard solid-state data recorder at high rates.

EO-1 is currently in extended mission, having more than achieved its original technology validation goals. As an example, over 5,000 data collection events have been successfully completed, against original success criteria of 1,000 data collection events.

EO-1 has 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. ALI images a 37km wide swath. 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.

In the following sections of the paper we describe:

1. the automated response;
2. the three principal science detectors that have been implemented; and
3. ongoing efforts to extend the sensorweb to additional orbital and ground-based assets.

3. AUTOMATING EO-1 TASKING

The automated retasking element of the sensorweb consists of several components working together as follows.
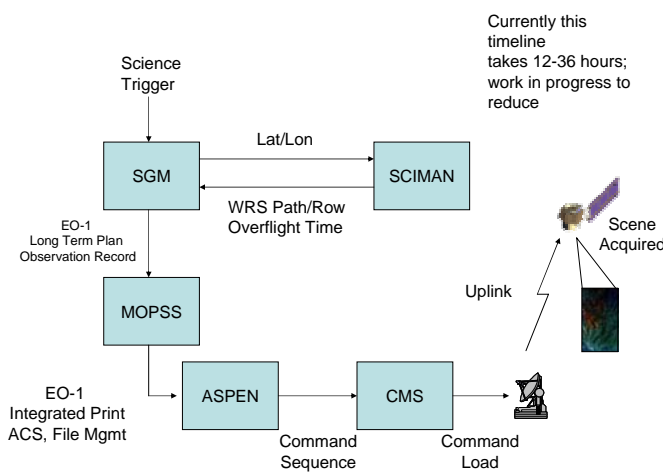

Figure 2. Sensorweb Response

1. The science alert (trigger) is noted by the Science Goal Monitor (SGM) [Jones et al]. The alert is specified by the latitude and longitude of the requested observation.
2. SGM then queries SCIMAN (the EO-1 observation planning system). SCIMAN uses information on the EO-1 orbit to determine the EO-1 overflight time and the worldwide reference system (WRS) path and row which also specifies the overflight time.
3. SGM uses this information to construct an EO-1 Long Term Observation plan record.
4. This observation record is ingested by MOPSS, the principal component of the EO-1 ground operations system. MOPSS is used in our automated architecture to perform file management and maneuver planning functions.
5. An integrated print of the observation specifying file management, maneuver, and observation parameters is passed to the Automated Scheduling and Planning Environment (ASPEN) system [Chien et al. 2001]. ASPEN uses this information

to generate a low level command sequence for the observations.
6. This command sequence is assembled for uplink by the command management system (CMS). This sequence is then uplinked.
7. Onboard EO-1 the command sequence is executed to acquire the science scene. Later, the science data is downlinked, processed, and delivered to the requesting scientist.

As an alternate architecture we are investigating uplinking the observation goal to the Autonomous Sciencecraft Experiment software [Chien et al. 2003] onboard EO-1 and enabling the response to be planned onboard.

The SGM system enables the science user to select triggers for areas of specific science interest. For example a user might specify a trigger such as "Image all fire alerts with confidence > 0.8 and > 10 km square area in this specified region of northern Montana" or "Image the Colima volcano every time the GOES volcano alert system confirms a MODVOLC volcano alert with > 0.8 confidence". The basic idea here is capture a transient science process through an automated trigger based on *content* or events, as opposed to the traditional specification of requests based on *location* and *time,* which often misses short-lived events such as floods, fires, or volcanic activity. SGM also tracks the request and data production for the scientist.
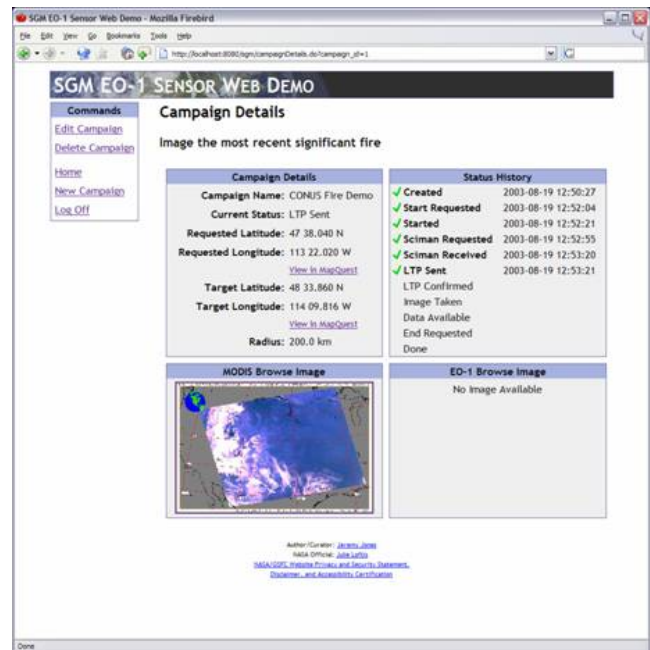

Figure 3: Science Goal Monitor Request Tracking

Our science response work also uses the ASPEN/CASPER planning & scheduling system [Chien et al. 2000][3]. Search in ASPEN has focused on high speed, local search, in a committed plan space, using stochastic combination of a portfolio [Huberman et al. 1997, Gomes and Selman 1997] of heuristics for iterative repair and improvement algorithms [Zweben et al., 1994. In this approach, at each choice point in the iterative repair process [Rabideau, et al. 1999], a stochastic choice is made among a portfolio of heuristics (with probabilities specifiable by the user). This approach has performed well in a wide range of space mission applications [Chien et al 2000b] including spacecraft operations scheduling, rover planning, ground communications station automation and autonomy for unpiloted aerial vehicles. 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 (100s of operations per CPU second on a workstation). To increase efficiency, we also make use of aggregates of activities [Knight et al 2000].

We have focused on an early-commitment, local, heuristic, iterative search approach to planning, scheduling and optimization. This approach has a number of desirable properties for spacecraft operations planning.

One of the first benefits is that using an iterative algorithm allows automated planning to be utilized at any time and on any given initial plan. The initial plan may be as incomplete as a set of goals, or it may 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. This enables local search planning to 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 [Jonsson et al 2000] do not have this property. Local search can also be easily adapted for use in a "mixed initiative" mode for partial ground-based automation.

A second benefit is that it is easier to write powerful heuristics that evaluate ground plans. These strong heuristics allow the search to be pruned, ruling out less promising planning choices.

Third, a local algorithm does not incur the overhead of maintaining intermediate plans or past attempts. This allows the planner to quickly try many plan modifications for repairing the conflicts or improving the preferences. However, unlike systematic search algorithms, it cannot be guaranteed that our iterative algorithms will explore all possible combinations of plan modifications or that it will not retry unhelpful modifications. In our experience, these

---

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

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, the effects of a resource usage and the corresponding resource profiles can be efficiently computed. Least-commitment techniques retain plan flexibility, but can be computationally expensive for large applications. Further discussions on this topic can be found in [Chien, Muscettola, et al., 1998].

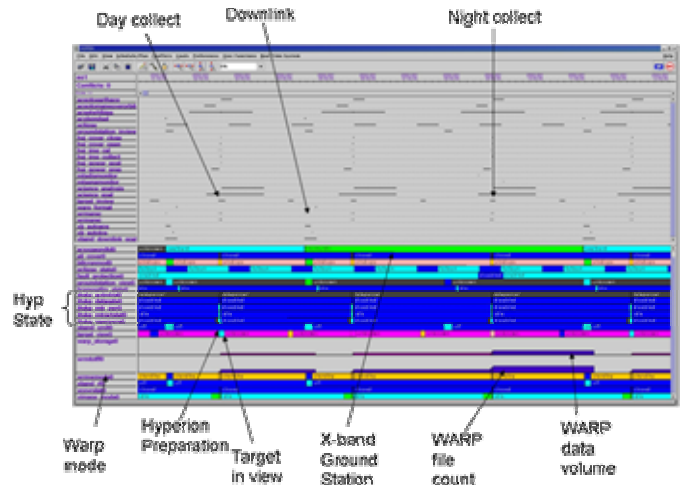Figure 4 shows the graphical user interface of an EO-1 operations plan.

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

ASPEN then generates an EO-1 command sequence from this activity plan (see below).

```
2003:233:16:49:57 CMD ACSETWHLBIAS(INERTIAL,X=0.341589,Y=1.1749,Z=-0.118046);
2003:233:17:56:57CMD
ACGOTOMANEUVER(ORBITAL,TIME=900,XLIMDEG=0.02,YLIMDEG=0.062699,…);
2003:233:18:07:06 CMD I_SETFPEPOWER(POWER_MASK=5);
2003:233:18:07:06 CMD YHEASTBY;
2003:233:18:07:16 CMD YHEASETSWIR(GAINA=1,GAINB=1,GAINC=1,GAIND=1,…);
2003:233:18:07:26 CMD YHEASETVNIR(VNIRALV8,VNIRBLV8,VNIRCLV8,VNIRDLV8);
2003:233:18:11:06 CMD I_CONFIGFPE(CONFIG_COMMAND=16908); …
2003:233:18:17:06 CMD BCMMODESCRS422;
2003:233:18:17:16 CMD WRMSREC(IDWS=65535,IDWV=65535,…);
2003:233:18:17:54 CMD I_SET_FPE_DG(DURATION=-1);
        …
```

4. THE WILDFIRE SENSORWEB

We have demonstrated the sensorweb concept using the MODIS active fire mapping system. 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 (cover 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 3-6 hours from acquisition.

The active fire mapping algorithm [Justice et al. 2002] detects hotspots using MODIS bands Ti using absolute thresholds:

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

It also uses a relative threshold algorithm which requires 6 nearby pixels in an up to 21x21 square that are cloud, smoke, water, and fire free. This triggers if the thermal reading is 3 standard deviations above the surrounding area.

- T4 > mean(T4)+ 3stddev(T4)
  and T4 - T11 > median(T4-T11)+ 3stddev(T4-T11)

Figure 5 shows the active fire map from October 2004 fires in Southern California. Figure 6 shows the context active fire map and a sensorweb trigger observation taken during this demonstration.


Figure 5: Active fire alerts for the recent October 2004 Southern California Fires. Red indicates active fires. The light blue box illustrates the background region used in the relative threshold detection.
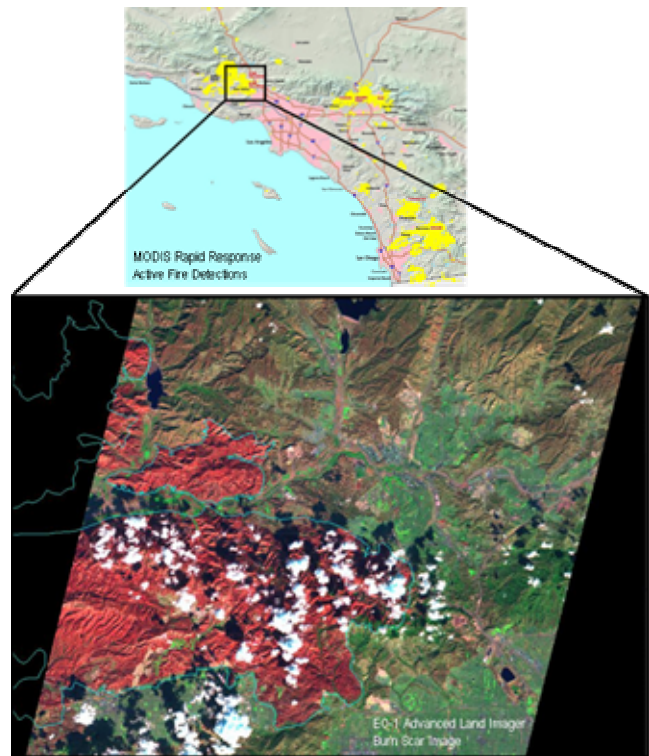

Figure 6: Sensorweb trigger images for October 2004 Southern California Fires. Above is the MODIS Active Fire Map display. Below is the EO-1 Hyperion image acquired via sensorweb trigger of the Simi/Val Verde fire area used in Burned Area Emergency Reclamation (BAER).

5. THE FLOOD SENSORWEB

The flood sensorweb uses the Dartmouth Flood Observatory [DFO] Global Active Flood Archive to identify floods in remote locations automatically based on satellite data. The DFO flood archive generates flood alerts based on both MODIS and QuikSCAT [Nghiem 2001] satellite data. The flood sensorweb utilizes the DFO QuikSCAT atlas because it is not affected by cloud cover over flooded areas.

The DFO archive is produced by the DFO in collaboration with JPL/QuikSCAT team. In this process the QuikSCAT Scatterometer data is used to assess surface water conditions [Brakenridge et al. 2003, Nghiem et al. 1999]. Specifically the VV/HH ratio is used to assess surface water properties of the areas in 0.25 lat/lon degree bins. The 7 day running mean is used to dampen effects of short-duration rainfall over urban areas. These data are then compared to the seasonal (90 day) average of the previous year season to screen out seasonal wetlands. The screened alerts are then published to a DFO website.
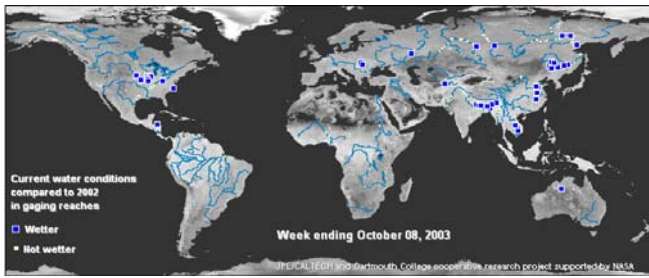
Figure 7: Dartmouth Flood Observatory Global Flood Alerts for October 2003.

In the flood sensorweb, active flooding alerts prime locations of known scientific interest trigger EO-1 observations at gauging reaches. Gauging reaches are river locations whose topography is well understood. Flood discharge measurements at gauging reaches can be used to measure the amount of water passing through a flooded region and can be compared with remotely sensed data. The end effect of the flood sensorweb is to increase the amount of high resolution remote sensing data available on flooding events in prime locations of interest (e.g., gauging reaches) and times of interest (e.g. when active flooding occurs). Imagery from an August 2003 flood sensorweb demonstration capturing flooding in the Brahmaputra River, India, is shown below.


Figure 8: Examples of low-resolution MODIS imagery (left) and EO-1 imagery (right) from the Flood Sensorweb capturing Brahmaputra River flooding in India, August 2003.

6. THE VOLCANO SENSORWEB

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

The GOES [Harris et al. 2002] and AVHRR alert systems provide excellent temporal resolution and rapid triggering based on thermal alerts. The GOES-based system looks for locations that are: hot, are high contrast from the surrounding area, and not visibly bright. Additionally, hits are screened for motion (to eliminate cloud reflections) and persistence (to remove instrument noise). The GOES alert can provide a web or email alert within 1 hour of data acquisition.

The MODIS alert system [Wright et al. 2002] has the advantage of high instrument sensitivity but has lower temporal resolution (MODIS generally has at least 4 overflights per day). MODVOLC deerives the normalized thermal index (NTI) from MODIS raw radiance values by computing $(R22 – R32)/(R22+R32)$ where Ri indicates the use of the radiance value from MODIS band i. The NTI is compared to a threshold to indicate alerts and is generally available online within 3-6 hours of acquisition.

7. ONGOING EXPANSION

Terrestrial dust storms are of significant science interest and can be detected using several sensors including GOES, AVHRR, and MODIS [Miller 2003]. These storms can become quite large (100s of kms long) and are of interest because of dust transport and aviation impact. A dust storm sensorweb would utilize 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 9 shows a large dust storm in the Persian Gulf as imaged by MODIS in November 2003. Dust storms can also be detected by ground-based instrumentation, such as operated by the U.S. Department of Agriculture in the U.S. Southwest and the Peoples Republic of China network of sites in the Gobi Desert. Detection and tracking of dust storms is also of considerable interest on Mars where such storms can grow to cover the entire planet.

Many freeze/thaw applications are also of interest. This includes the phenomena of glacial ice breakup, sea ice breakup, melting, and freezing, lake ice freezing and thawing, and snowfall and snowmelt. All of these phenomena can be detected and tracked using both Quikscat and MODIS and used as triggers for higher resolution imaging such as with EO-1.

6

Figure 9: Dust Storm in Perian Gulf as captured by MODIS November 2003.

A wide range of in-situ instrumentation also exists for detecting science events of interest. We are working with a number of teams to integrate such sensors into our sensorweb. The Hawaiian Volcano Observatory [HVO] has deployed numerous instruments on the Kilauea region in Hawaii. 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 infra-red cameras [Harris et al. 2003] to a number of volcanic sites worldwide (e.g., Kilauea, Hawaii; Erte Ale, Ethiopia; Sourfiere Hills, Montserrat; Colima and Popocatepetl, Mexico). These infra-red cameras can provide a ground-based detection of lava flows based on thermal signatures, thereby alerting the sensorweb.

In a collaboration with the Center for Limnology of the University of Wisconsin at Madison, we have linked into data streams from the Trout Lake stations to use temperature data to trigger imaging of the sites to capture transient freezing and thawing processes.

8. EXTRA-TERRESTRIAL APPLICATIONS

The sensorweb concept is directly applicable to deep space science applications, sun-earth connection science, and astrophysics applications. For example, on Mars surface instruments could detect and/or track active, transient atmospheric and geologic processes such as dust storms. Alternatively, sun-pointed instruments could detect Coronal Mass Ejections (CMEs) and alert Earth orbiting magnetospheric instruments (e.g. IMAGE, MMC, MMS, …) to reconfigure to maximize science data.

9. RELATED WORK & SUMMARY

There has been considerable effort devoted towards closed loop science for rovers at ARC [Gulick et al. 2001], JPL [Castano et al. 2003], and CMU [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 in the FY04 timeframe.

The Three Corner Sat (3CS) mission [Chien et al. 2001] was scheduled to fly onboard data validation, execution, and replanning in August 2002. However, 3CS was delayed and was a shuttle launch and thus its launch future is uncertain. However, the basic hardware and technology was flown in a balloon sat to an altitude of 100,000 feet in August 2003. 3CS is now scheduled for launch in July 2004.

The Autonomous Sciencecraft Experiment on EO-1 [Chien et al. 2003] 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 automated science event tracking system with an autonomous response capability based on automated planning technology. Demonstration of these sensorweb capabilities will enable fast responding science campaigns and increase the science return of spaceborne assets.

REFERENCES

B. Brakenridge et al., Flood Warnings, Flood Disaster Assessments, and Flood Hazard Reduction: The Roles of Orbital Remote Sensing, Proc 30th Intl Symp on Remote Sensing of the Environment, Honolulu, HI, Nov 2003.

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, R. Knight, A. Stechert, R. Sherwood, and G. Rabideau, "Using Iterative Repair to Improve Responsiveness of Planning and Scheduling," Proceedings of the Fifth International Conference on Artificial Intelligence Planning and Scheduling, Breckenridge, CO, April 2000. (see also casper.jpl.nasa.gov)

Chien, S. et al. (2003) Autonomous Science on EO-1, I-SAIRAS, 7th Symposium on Artificial Intelligence, Robotics and Automation in Space, Nara, Japan, May 2003.

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

S. Chien, N. Muscettola, K. Rajan, B. Smith, G. Rabideau, "Automated Planning and Scheduling for Goal-based Autonomous Spacecraft," *IEEE Intelligent Systems*, September/October 1998, pp. 50-55.

Davies, A. G. et al. (2001) ASC Science Study Report, available from http://ase.jpl.nasa.gov

Davies A. G. et al. (2002) LPSC-33 abstract 1130, on CD-ROM.

A.G. Davies, R. Greeley, K. Williams, V. Baker, J. Dohm, M. Burl, E. Mjolsness, R. Castano, T. Stough, J. Roden, S. Chien, R. Sherwood, "ASC Science Report," August 2001. (downloadable from ase.jpl.nasa.gov)

Davies, A. G., E.D. Mjolsness, A.G. Gray, T.F. Mann, R. Castano, T.A. Estlin and R.S. Saunders (1999) Hypothesis-driven active data analysis of geological phenomena using semi-autonomous rovers: exploring simulations of Martian hydrothermal deposits. EOS, Trans. Amer. Geophys. Union, 80, no. 17, S210.

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://EO-1.gsfc.nasa.gov

K. Golden, W. Pang, R. Nemani, and P. Votava, "Automating the Processing of Earth Observation Data," *Proceedings of the International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS)*, Nara, Japan, May 2003.

C. Gomes, B. Selman, "Algorithm Portfolio Design: Theory vs. Practice," *Proceedings of the Conference on Uncertainty in Artificial Intelligence,* Providence, RI. 1997.

Greeley et al., (2003) Lunar and Planetary Sciences Conference-34 abstract 1779, on CD-ROM.

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.

Harris, A. et al., (2002) Web-Based Hot Spot Monitoring using GOES: What it is and How it Works, Advances in Environmental Monitoring and Modelling http://www.kcl.ac.uk/advances Vol. 1 No. 3 (2002) pp.5-36.

Harris et al., Ground-based Infrared Monitoring Provides New Tool for Remote Tracking of Volcanic Activity, EOS,Vol. 84, No. 40, 7 October 2003

Hawaiian Volcano Observatory (HVO), http://hvo.wr.usgs.gov/

B. Huberman, R. Lukose, T. Hogg, "An Economic Approach to Hard Computational Problems, *Science* v. 275, 3 January 1997, pp. 51-54.

J. Jones et al., (2003) Science Goal Monitor Project, http://aaa.gsfc.nasa.gov/SGM.

Justice, C. et al., (2002) The MODIS fire products, Remote Sensing of Environment 83 (2002) 244–262

R. Knight, G. Rabideau, S. Chien, "Computing Valid Intervals for Collections of Activities with Shared States and Resources," *Proc 5th Intl Conf Artificial Intelligence Planning and Scheduling*, Breckenridge, CO. April, 2000.

Kratz, T., et al. "University of Wisconsin at Madison, Center for Limnology, Trout Lake Station," http://limnology.wisc.edu/tls/troutlake.html

M. C. Malin and K. S. Edgett, "Evidence for recent groundwater seepage and surface runoff on Mars," Science, 288, 2330-2335, 2000.

S. Miller, "A consolidated technique for enhancing desert dust storms with MODIS," Geophysical Research Letters, Vol. 30, No. 20, 2071.

S. V. Nghiem, W. T. Liu, and X. Xie, "Polarization Reversal over Flooded Regions and Applications to Flood Mapping with Spaceborne Scatterometers," International Geoscience and Remote Sensing Symposium, Hamburg, Germany, June 28 - July 2, 1999.

S. V. Nghiem, Advanced Scatterometry for Geophysical Remote Sensing, JPL Document D-23048, 40 pages, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, March 2001.

G. Rabideau, R. Knight, S. Chien, A. Fukunaga, A. Govindjee, "Iterative Repair Planning for Spacecraft Operations in the ASPEN System," International Symposium on Artificial Intelligence Robotics and Automation in Space, Noordwijk, The Netherlands, June 1999.

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.

R. Wright et al. (2002). Automated volcanic eruption detection using MODIS. Remote Sensing of Environment, 82, 135-155.

M. Zweben, B. Daun, E. Davis, and M. Deale, "Scheduling and Rescheduling with Iterative Repair," *in Intelligent Scheduling*, Morgan Kaufman, San Francisco, 1994.

ACKNOWLEDGEMENT