Onboard Processing for Low-latency Science for the HyspIRI Mission

Steve Chien[*], David Mclaren[*], Gregg Rabideau[*] , Dorothy Silverman[*],
Daniel Mandl[**], Jerry Hengemihle[***]

[*] Jet Propulsion Laboratory, California Institute of Technology, USA
e-mail: firstname.lastname@jpl.nasa.gov

[**]Goddard Space Flight Center, USA
e-mail: daniel.mandl@gsfc.nasa.gov

[***] Microtel, LLC
e-mail: jerry.hengemihle@microtel-llc.com

Abstract

Hyspiri is evaluating a X-band Direct Broadcast (DB) capability that would enable data to be delivered to ground stations virtually as it is acquired. However the HyspIRI VSWIR and TIR instruments will produce 1 Gbps data while the DB capability is 15 M bps for a ~60x oversubscription. In order to address this data volume mismatch a DB concept has been developed that determines which data to downlink based on both: 1. The type of surface the spacecraft is overflying and 2. Onboard processing of the data to detect events. For example when the spacecraft is overflying polar regions it might downlink a snow/ice product. Additionally the onboard software will search for thermal signatures indicative of a volcanic event or wild fire and downlink summary information (extent, spectra) when detected.·

1 Introduction

Future space missions will produce immense amounts of data. A single image from the HiRise camera on the Mars Reconaissance Orbitter (MRO) spacecraft is 16.4 Gigabits (uncompressed). The future HyspIRI mission under study is proposed to have two instruments - the HyspIRI thermal infrared imager (TIR) instrument producing 1.2 million pixels per second with 8 spectral bands at 4 and 7.5-12 microns per pixel and the HyspIRI visible shortwave infrared (VSWIR) producing 300 thousand pixels per second with 220 spectral bands per pixel in the 0.4-2.5 micron range. Keeping up with these data rates requires efficient algorithms, streamlined data flows and careful systems engineering. HyspIRI is also considering using Direct Broadcast technology to rapidly deliver this data to application users on the ground. However, in order to leverage the existing DB network, this downlink path is limited to approximately 15 million bits per second. The question is – which data to downlink when, in order to maximize the utility of the DB system?

We are studying the desired products and spectral bands required by volcanic, wildfire, flood and ocean/coastal, snow/ice, dust, and vegetation/ecosystem applications to assess onboard processing and band selection strategies for the mission. Three baselines for study are being investigated:

1. downlink of the MODIS bands over all target areas;
2. downlink of specially selected subsets of the bands based on overflight targets; and
3. onboard development of custom products based on overflight masks.

Volcanic applications include thermal detection and signature analysis as well as plume tracking applications. These volcanic techniques enable spatial subsampling to the areas of interest for dramatic downlink reduction. Onboard (EO-1) detection and ground-based (MODIS, AVHRR) detection algorithms are well understood. Wildfire applications include active fire mapping based on thermal signature (onboard EO-1, ground-based MODIS) as well as development of burned area products (significant heritage with Landsat ETM+, prior work with EO-1/ALI). A significant range of other applications with strong heritage in MODIS, AVHRR, GOES, and other rapid data delivery sensors exist in a range of disciplines. Ocean/coastal applications include products such as such as sea surface temperature and sea color applications such as harmful algal bloom tracking and Chl indices. Snow/Ice applications include trafficability and commerce route safety products as well as science cryosphere uses. Dust applications include aviation hazard assessment and environmental applications. Vegetation applications include plant stress, fire hazard, and disease vector applications based on measures of plan health and species identification.

These applications were derived from existing DB applications, discussions with the HyspIRI working groups, and others working in the relevant areas. Processing algorithm under consideration were assessed for adaptability and heritage from relevant prior sensors including MODIS, AVHRR, ASTER, Hyperion and others. These products are being refined with science and applications inputs and tested on current datasets and missions (such as EO-1) as well as being tested on relevant flight processing hardware and software configurations.

In the remainder of this paper we describe the operations concept being developed for the direct broadcast option for HyspIRI, outline applications for this concept, and ongoing benchmarking on potential flight hardware.

.

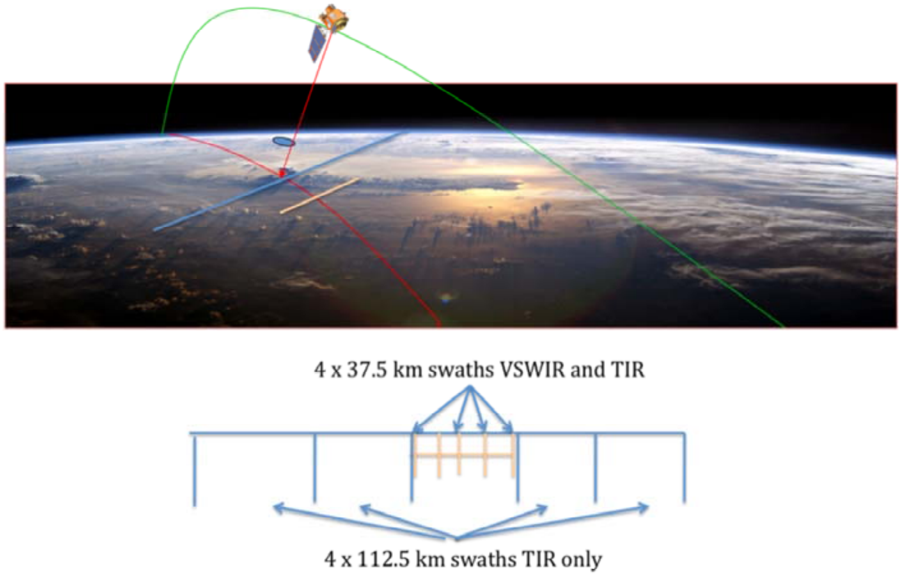

Figure 1: Projected Instrument Swaths

2 A Direct Broadcast Operations Concept for HyspIRI

The HyspIRI DB operations concept key drivers are:

1. Low or no sustaining operations costs
2. Low or no system development costs
3. Maximize utility of returned data
4. Graceful degradation/ high reliability of operations concept
5. Low risk, high heritage

With these drivers in mind, we have developed a highly automated operations flow for the DB component on HyspIRI consisting of the following steps. In order to reduce cost and risk we have used mature software systems.

1. Specification of geographical regions of interest (ROIs) by the DB applications team. In this step, the applications team has a set of geographical regions (in essence polygons on a map of the Earth). For each polygon there is an algorithm and a priority. These polygons may also be seasonal (e.g. January to March of each year) and may be derived based on external information (e.g. reports of flooding, or rainfall, or the National Interagency Fire Center (NIFC) fire reports.
2. The spacecraft operations team provides the current best projection of the orbit of the spacecraft (e.g. a ground track file).
3. The spacecraft orbit is combined with the ROIs using knowledge of the spacecraft instrument swaths (150km wide for VSWIR and 600km wide for TIR) using the CLASP coverage planner [Knight 2009]. This produces a timeline of overflights for each of the 8 instrument swaths.
4. The ASPEN [Chien et al. 2000] is then used to determine the top priority products/spectral bands to process for each timestep, respecting the product priorities as specified by the applications team. ASPEN produces an activity plan/sequence for the onboard processing module.

Thus the spacecraft orbit determines the type of terrain that will be overflown (e.g. land, ice, coastal, ocean, etc.). The TIR instrument has a 600km swath under the spacecraft and the VSWIR a 150km wide swath. In order to satisfy the high data rate from the instruments there are four interfaces from the instrument to the onboard processing. The VSWIR data is divided into four across track swaths of 37.5 km each. The TIR data includes four swaths matching the VSWIR swaths with the remaining 450km with of TIR only data divided into another four data paths (see Figure 1). Therefore each of the interface paths receives one 37.5km swath of VSWIR and TIR data and one 112.5k swath of TIR only data.

Each of the terrain masks implies a set of requested modes and priorities. And is evaluated based on the eight swaths from the instruments. For example, when overflying polar or mountainous regions, producing snow and ice coverage maps can provide valuable science data. Additionally, the science team can adjust

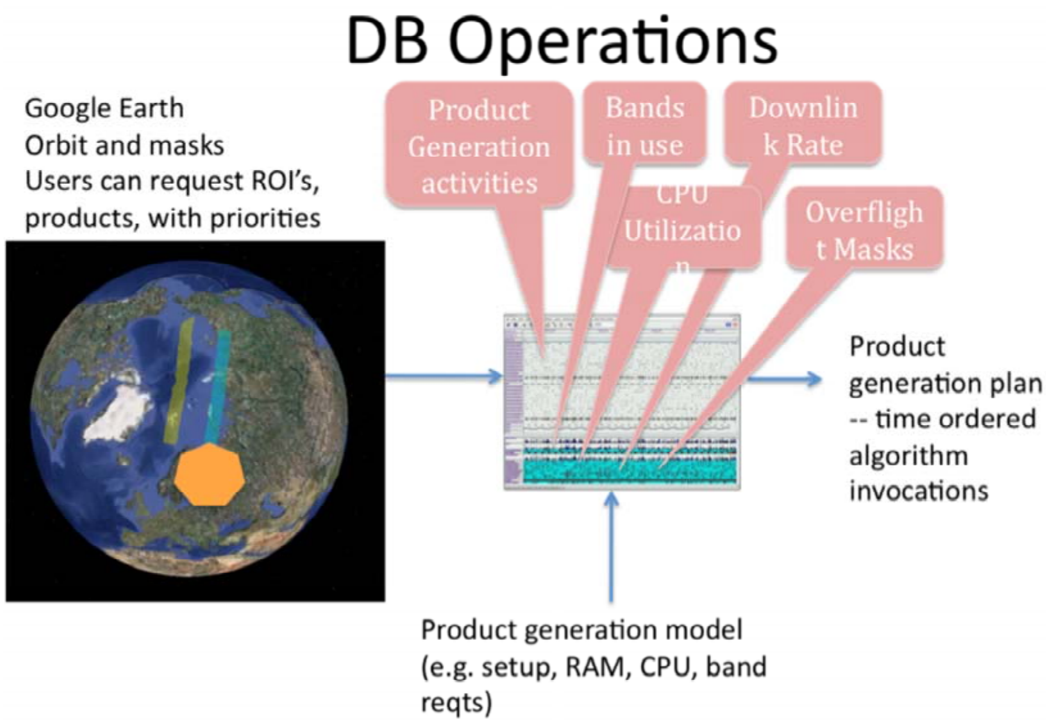

Figure 2: Direct Broadcast Operations Flow

these priorities based on additional information (e.g. external information that a volcano is active, knowledge of a flooded area, an active wildfire, or a harmful algal bloom). The mission planning tool accepts all of these requests and priorities, and determines which onboard processing algorithms will be active by selecting the highest priority requests that will fit within the onboard processing CPU resources, band processing limitations,

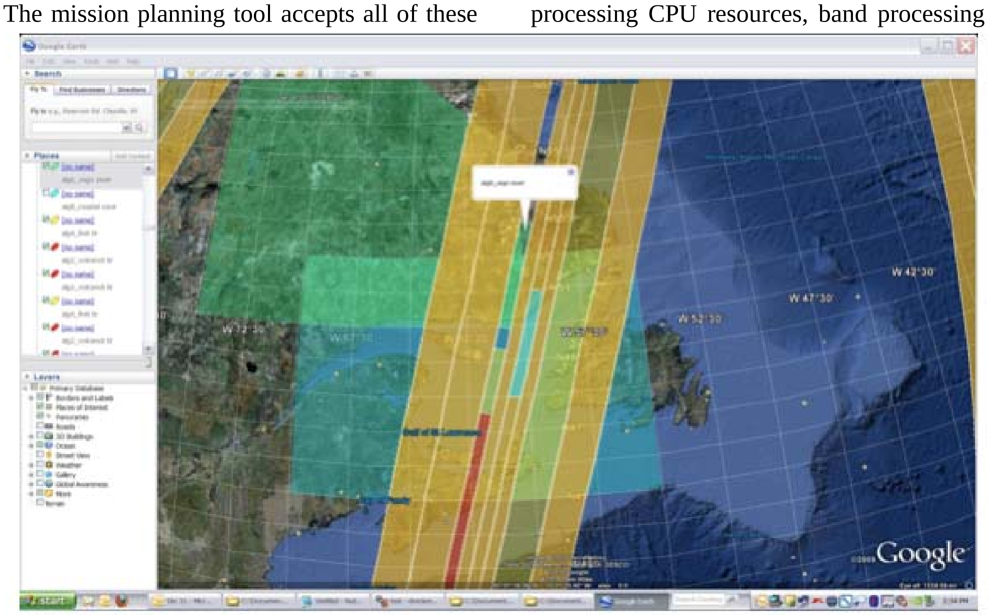

Figure 3: Instrument swaths overlayed against science regions of interest.

and downlink bandwidth.

The automated planning model tracks the limited spacecraft resources that in this case include: # of bands processed, onboard CPU (each algorithms places a different load on the CPU), and downlink bandwidth. These operations constraints represent the onboard restrictions that: 1) only a limited number of bands of the instrument data can be processed onboard (for example, on EO-1 we can only process 12 of the bands per image), 2) that we have limited CPU processing capability onboard and this may limit the products we can generate at any one time, and 3) that the downlink transmission rate is limited to 15 Megabits per second. Accounting for these operations constraints, the mission planning system chooses the highest priority products that can be produced. Figure 3 shows the instrument processing swaths and Figure 4 shows a sample mission plan generated based on CPU and downlink resources.

Onboard processing algorithms can use a wide range of techniques. Past algorithms have consisted of: expert derived decision tree classifiers, machine learned classifiers such as Support Vector Machines (SVM) classifiers and regressions, classification and regression trees (CART), Bayesian maximum likelihood classifiers, spectral angle mappers, and direct implementations of spectral band indices and science products.

For example, SVM's have been applied to learn to classify EO-1 Hyperion images into Snow, Water, Ice, Land, and Cloud pixels [Castano et al. 2005].

CART techniques [Breiman et al. 1984] have been applied to a wide range of classification problems including remote sensing [Castano et al. 2006]. CART techniques recursively split the decision classification or estimation problem until a stopping criterion of goodness of fit is met. Maximum likelihood classifiers have also been applied to classification of remote sensing imagery (e.g., [Goodenough et al. 2003]). Given a presumed parametric probability distribution, these techniques find the parameters that maximize the likelihood of the observed training set. Spectral angle mapping (SAM) is an instance-based classification technique. If one considers each pixel as an n dimensional vector if the remote sensing imagery has n spectral measurements, SAM selects as the matched class the one whose prototype spectral vector is closest in angular distance to the vector of the pixel in question.

In other cases the science disciplines have already developed science products (e.g. measures) to track a physical phenomenon. In oceanography, Flourescent Line Height and Maximum Chlorophyll Index are indicative of biological activity such as algal bloom. In vegetation and ecosystem monitoring Normalized Difference Vegetation Index and Photochemical Reflectance Index (PRI) are indicative of plant health.

3 Rapid Response Science and Applications

In this section we describe in greater detail a number of applications for rapid delivery data.

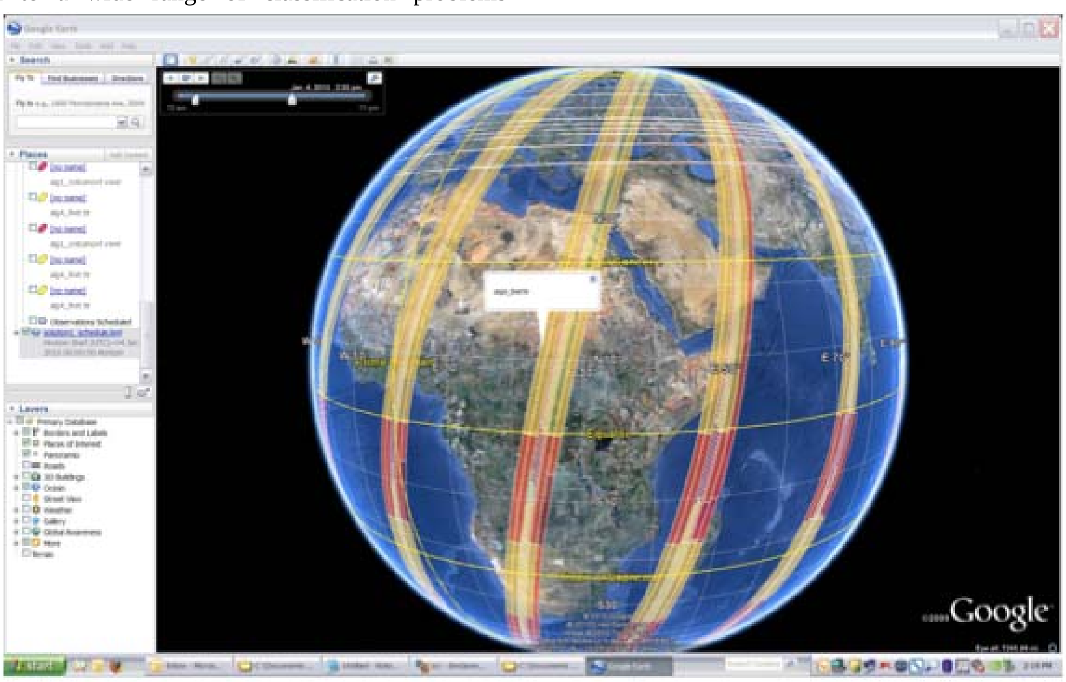

Figure 4: Sample processing plan.

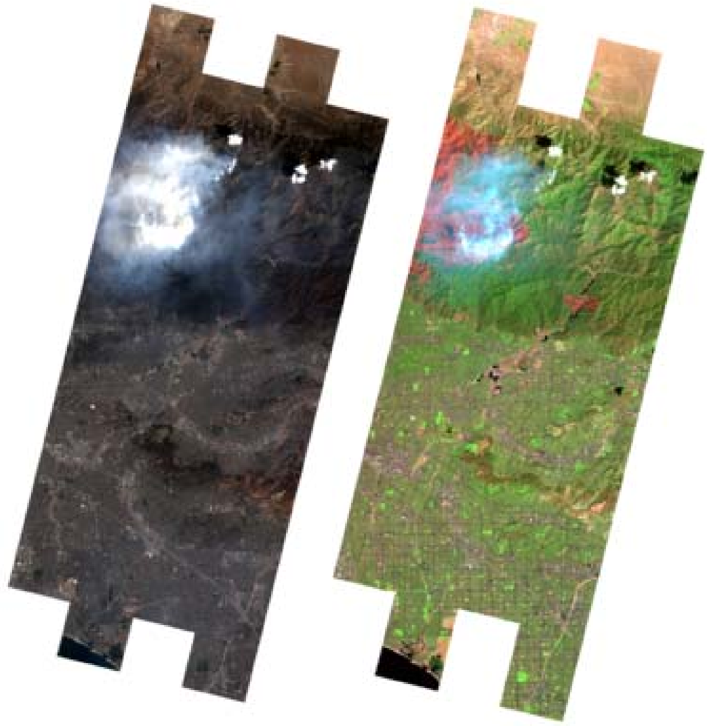

Figure 5: Visible and burn scar enhanced images from ALI instrument on EO-1 of Station Fire near Los Angeles
03 September 2009

Images courtesy EO-1 Mission NASA GSFC

Fire Products

The HyspIRI instruments are useful in producing a variety of fire products to support as active fire mapping and burn severity assessment for burned area reclamation. MODIS rapid response imagery has demonstrated this utility in providing fire location information that can be used by relevant agencies to allocate scarce assets. For example, the RapidFire active fire mapping project of the MODIS land rapid response team [Justice et al. 2002] provides data to a number of organizations including the National Interagency Fire Center (NIFC). Other products and data can be used provide broad coverage burn assessments such as difference normalized burn ratios (NBR) and Composite Burn Index (DCI) that assess burn damage to vegetation and the environment for reclamation planning. Such measures allow agencies to quickly assess the damage in an area and implement reclamation plans in a matter of days.

Operationally, active fire mapping algorithms would operate over large amounts of the Earth's surface (likely all land masses). These algorithms would use a small number of TIR channels (likely 3) to estimate the surface temperature and signal pixels that are likely to be hot. Performing active fire mapping onboard represents a tremendous data reduction because only an extremely small proportion of the landmass is burning at any one time, therefore the data volume required to downlink active fire data is tiny.

Burn scar data is likely to be requested based on external fire information. This could occur manually (e.g. an application user requests the data for a burned area explicitly) or automatically by requesting the data for all burned areas as reported by other agencies and electronically compiles (e.g., such as the NIFC site). Again, the total area that has recently burnt (e.g. the past fire season) while a large area is not large relative to the total landmass. Therefore request-based burn scar information represents a huge downlink savings compared to bringing down all of the data.

Volcano Applications

The HyspIRI TIR and VSWIR instruments have great applicability to volcano monitoring. The HyspIRI instruments can be used to measure the thermal signature of volcanic sources. While significant prior work has tracked thermal signatures of volcanic activity from space [Harris et al. 2000, Wright et al. 2003, 2004], HyspIRI offers more sensitive detection capability. The EO-1 Hyperion instrument has been used operationally with onboard software to track volcanic thermal activity [Davies et al. 2005] and automatically downlink full spectra of thermally active pixels. HyspIRI could use a similar capability to detect and summarize volcanic activity enabling the spacecraft to pick out the few key pixels of data out of literally billions of non-relevant pixels. Such an algorithm would use several TIR bands to estimate surface temperature and flag pixels likely to be hot due to volcanic activity. When such pixels are detected a notification with TIR and selected VSWIR data would be downlinked for each flagged pixel.

Onboard processing of volcanic data represents a huge win from both a data volume and timeliness perspective. Because volcanic eruptions represent a small number of pixels (even a major eruption might only cover a few hundred pixels), localization of the volcanic activity means analysis can focus on a relatively small fraction of pixels). Volcanic activity also represents an extreme example of timeliness. Given the large number of people living close to volcanos, accurate and timely information is critical to informed assessment of current and future risk to both lives and property.

Other Applications

Snow and Ice Products

The Hyspiri instruments (both VSWIR and TIR) will be useful for studying and monitoring Snow, Water, Ice, and Land (SWIL) phenomena. SWIL classification is important for monitoring climate change, assessing environmental sustainability, and regulating both land-based and sea traffic. Ice and snow products have been developed on the ground using a range of instruments including MODIS [MODIS Snow Ice] and onboard spacecraft using Hyperion [Doggett et al. 2006]. Because of the impact of snow and ice on commercial activities rapid delivery of this remote sensed data is important.

Flooding

HyspIRI instruments are useful for tracking surface water extent with applications to flooding and disaster response. Because flooding is the greatest natural hazard (both in terms of lives lost and property damage), any real-time capability to assist in humanitarian efforts is of tremendous importance. Both the Dartmouth Flood Observatory [Brakenridge and Anderson 2005] and the University of Maryland [Carroll et al. 2009] have used MODIS to provide near real-time flood mapping. EO-1 Hyperion has also been employed to detect floods using onboard software that enables rapid alert generation and retasking [Ip et al.

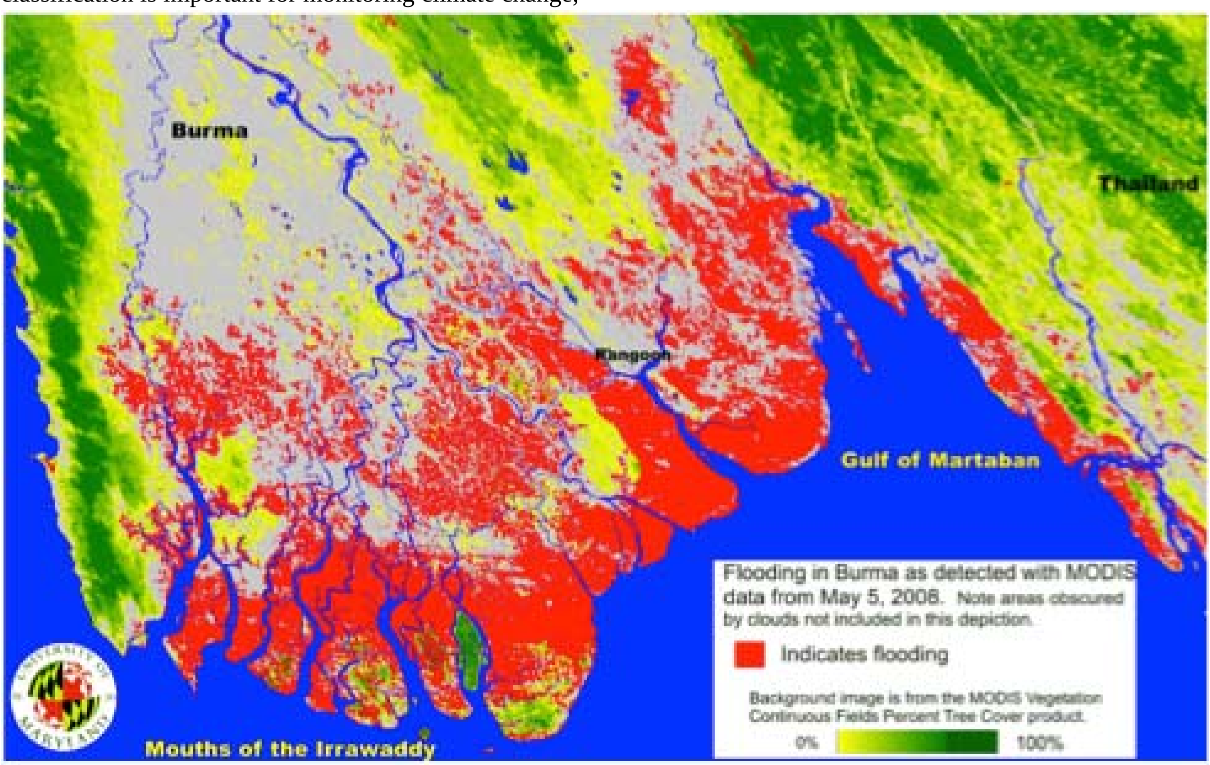

Figure 6: UMD Flood tracking of Myanmar using MODIS bands 1,2,5,7 (620-2155 nm) Courtesy M. Carroll et al. U. Maryland

2006]. Because of the tremendous human and economic impact of flooding, rapid delivery of flooding data is critical.

Dust Products

HyspIRI instruments can be used to track large-scale dust storms using both color (VSWIR) and thermal (TIR) information. These dust storms are hundreds of kilometers in extent and threaten human health and aircraft safety as well as having significant environmental impacts. Because of these major impacts and the dynamic nature of dust storms rapid delivery of relevant satellite data is key.

Vegetation

The HyspIRI instruments can be used to measure plant stress [Perrry and Roberts 2008] as well as identify plant species. Monitoring vegetation pigment levels with VSWIR in the 500-1200 nm range can identify plant stress to assist in predicting crop failures – a key rapid response application. VSWIR can also identify plant species with rapid response applications for disease risk estimation. Thermal plant stress as measured by the TIR instrument can also be used to estimate evapotranspiration (ET) [Anderson & Kustas 2008], a key indicator in predicting crop failure. Timely products from MODIS are currently being used by the USDA to assess crop health and yield (e.g., as affected by drought, fires, volcanic eruptions, or storms). For example, the U.S. Department of Agriculture's Foreign Agricultural Service (FAS) is using MODIS data to estimate predicted crop yields. This data enables FAS to make accurate crop-yield estimations, which ultimately affect decisions impacting U.S. agriculture, trade policy and food aid. All of the above applications require timely delivery of data for decision making based on crop and disease models.

Ocean/Coastal

The HyspIRI VSWIR and TIR instruments are useful for studying a wide range of oceanographic applications, many of which significantly benefit from rapid data delivery. Ocean color measurements using VSWIR can be used to detect and track harmful algal blooms (HAB's) and other biological events, and sediments that pose a threat to both exposed people and wildlife. Because of these hazards, rapid data delivery is of great value. Ocean/coastal applications present unique challenges because of the subtlety of the ocean signals. Image corrections (e.g., atmospheric correction) are critical and present tremendous challenges for onboard processing.

4 Benchmarking Onboard Algorithms in Flight Testbeds

As the HyspIRI mission is still in concept development, several hardware options are being evaluated for onboard procesing for the Direct Broadcast component. These platforms in clued the Space Cube [Flatley 2008], Isaac and follow on packages [He 2008], both based on the Vertex chipset, and the Opera [Opera 2008] package based on the Tilera chipset. All of these options provide significantly more computing power than conventional CPU's which is needed due to the very high data rates from the HyspIRI TIR and VSWIR instruments.

5 Discussion, Related Work, Conclusions

A number of prior missions have performed some aspect of onboard data processing to manage downlink. The Autonomous Sciencecraft on EO-1 [Chien et al. 2005] produced summary products for Volcano, Cryosphere, and Flooding science events. The Mars Exploration Rovers WATCH system processed images to summarize and detect dust devils and clouds in rover imagery [Castano et al. 200X]. The Mars Odyssey mission averages down THEMIS data to manage downlink product size [Odyssey 2007].

We have discussed a concept under development for Direct Broadcast for the HyspIRI mission. Because the HyspIRI TIR and VSWIR produce 1 Gbits per second of data and the heritage X band direct broadcast link can only downlink 15 M bits per second, onboard data reduction is required. We have presented a hybrid approach that uses scientist specified regions of interest, onboard processing and event detection, and product generation all as methods to reduce the amount of data for downlink. This approach is currently under refinement and evaluation using the EO-1 mission as a testbed and simulations. Prototype algorithms are also being benchmarked on flight option hardware.

6 Acknowledgements

The authors gratefully acknowledge inputs from many science and applications specialists including members of the HyspIRI Science working group. Special thanks are due to: M. Anderson, R. Brakenridge, P. Campbell, I. Csiszar, A. Davies, R. Furfaro, L. Giglio, G. Glass, J. Kargel, F. Kogan, R. Kudela, J. Mars, E. Middleton, S. Miller, F. Muller-Karger, A. Prakash, V. Realmuto, D. Roberts, J. Ryan, R. Sohlberg, and R. Wright. However, full responsibility for any errors or omissions lies with the authors.

Portions of this work were performed by the Jet Propulsion Laboratory, California Institute of

Technology, under contract with the National Aeronautics and Space Administration.

References

M. Anderson and W. Kustas, "Thermal Remote Sensing of Drought and Evapotranspiration," Eos, Vol. 89, No. 26, 24 June 2008, pp. 233-234.

L. Breiman, J. Friedman, C. Stone, R. Olshen, Classification and Regression Trees, 1984, Chapman & Hall CRC.

Carroll, M., Townshend, J., DiMiceli, C., Noojipady, P., Sohlberg, R. (2009), A New Global Raster Water Mask at 250 Meter Resolution. International Journal of Digital Earth, December 2009.

R. Castano, D. Mazzoni, N. Tang, T. Doggett, S. Chien, R. Greeley, B. Cichy, A. Davies, "Onboard Classifiers for Science Event Detection on a Remote Sensing Spacecraft," *Proceedings of the Twelfth Annual SIGKDD International Conference on Knowledge Discovery and Data Mining*, August 2006. Philadelphia, USA, ACM Press New York, NY pp. 845-851.

R. Castano, K. Wagstaff, S. Chien, T. Stough, B. Tang, Onboard analysis of uncalibrated data for a spacecraft at Mars, *Proceedings of the 13th International Conference on Knowledge Discovery and Data Mining*, August 2007, San Jose, CA.

S. Chien et al., "Using Autonomy Flight Software to Improve Science Return on EarthObserving One," J. Aerospace Computing, Information, and Communication, Apr. 2005, pp. 196–216.

S. Chien et al., "An Autonomous Earth-Observing Sensorweb," IEEE Intelligent Systems, vol.20, no. 3, 2005, pp. 16–24.

A.G. Davies et al., "Monitoring Active Volcanism with the Autonomous Sciencecraft Experiment (ASE) on EO-1," Remote Sensing of Environment, vol. 101, no. 4, 2006, pp. 427–446.

Ewert J. W. and C. J. Harpel (2004), In harm's way: population and volcanic risk, *Geotimes*, 49, no. 4, 14-17.

T. Flatley, SpaceCube Processor, http://hyspiri.jpl.nasa.gov/downloads/public/2.9%20Flatley%20-%20OnBoard%20Processing%20Technology.pdf

D. Goodenough, A. Dyk, K. O. Neimann, J. Perlman, H. Chien, T. Han, M. Murdoch, C. West, "Processing Hyperion and ALI for Forest Classification," IEEE Transactions on Geoscience and Remote Sensing, V. 41, No. 6, June 2003.

Harris, A. J. L., Flynn, L. P., Dean, K., Wooster, E. M., Okubo, C., Mouginis-Mark, P., et al. (2000). Real-time satellite monitoring of volcanic hot spots. Geophysical Monograph, vol. 116 (pp. 139–159) pub. AGU.

Yutao He Le, C. Zheng, J. Nguyen, K. Bekker, D. , ISAAC - a case of highly-reusable, highly-capable computing and control platform for radar applications, IEEE Radar, 2009.

Justice, C.O., Giglio, L., Korontzi, S., Owens, J., Morisette, J.T., Roy, D.P., Descloitres, J., Alleaume, S., Petitcolin, F., Kaufman, Y. 2002. The MODIS fire products. Remote Sensing of Environment 83:244-262.

R. Knight and S. Hu. "Compressed Large-scale Activity Scheduling and Planning (CLASP)," Proceedings of the Sixth International Workshop in Planning and Scheduling for Space (IWPSS-2009), July 19th-21st, Pasadena, California, 2009.

MODIS Snow Ice Web site, http://modis-snow-ice.gsfc.nasa.gov/

OPERA http://www.zettaflops.org/spc08/Malone-Govt-OPERA-FT-Spaceborne-Computing-Workshop-rev3.pdf, 2008

E. Perry and D. Roberts, "Sensitivity of Narrow-Band and Broad-Band Indices for Assessing Nitrogen Availability and Water Stress in an Annual Crop, Agronomy Journal, vol. 100 Issue 4, 2008, pp. 1211-1219.

Wright, R., Flynn, L. P., Garbeil, H., Harris, A., & Pilger, E. (2003). Automated volcanic eruption detection using MODIS. Remote Sensing of Environment, 82, 135– 155.

Wright, R., Flynn, L. P., Garbeil, H., Harris, A. J. L., & Pilger, E. (2004). MODVOLC: Near-real-time thermal monitoring of global volcanism. Journal of Volcanology and Geothermal Research, 135, 29–49.

F. Ip, J. Dohm, V. Baker, T. Doggett, A. Davies, R. Castano, S. Chien, B. Cichy, R. Greeley, R. Sherwood, D. Tran, G. Rabideau, Flood detection and monitoring with the Autonomous Sciencecraft Experiment onboard EO-1, Remote Sensing of Environment, Volume 101, Issue 4, 30 April 2006, Pages 463-481

G.R. Brakenridge and E. Anderson, MODIS-based flood detection, mapping, and measurement: the potential for operational hydrological applications. In: Transboundary Floods, Proc. of NATO Advanced Research Workshop, Baile Felix – Oradea, Romania, May 4-8, 2005

HyspIRI Mission, http://hyspiri.jpl.nasa.gov