ADAPTATION & USE OF OPENGIS© WEB TECHNOLOGIES FOR MULTI-DISCIPLINARY ACCESS TO PLANETARY DATA

Elaine R. Dobinson, David W. Curkendall, and Lucian Plesea

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, California 91109-8099
USA

and

Trent M. Hare

U.S. Geological Survey
2255 N. Gemini Dr.
Flagstaff, AZ, 86001

Email: {Elaine.Dobinson,David.Curkendall,Lucian.Plesea}@jpl.nasa.gov;thare@usgs.gov

ABSTRACT

.
We are adapting the fast developing and well-supported OpenGIS standards and technologies for the access, processing, and display of geo-spatial data to the planetary domain. In particular, we are implementing WMS and WCS servers to present high level products derived from NASA's Mars and Lunar datasets to produce scientific composites of our knowledge of the planet. The specific information technologies used for this work are the existing OpenGIS methodologies for Earth data; we are extending and generalizing these standards for the planetary case. The basic motivation is that investigators need a unified methodology for accessing higher-level products that serve as substrate, background, and current known detail pertinent to their ongoing work. The goal and promise is that, when ultimately placed in operation at the relevant data archival sites, the data can continue to be archived at these distributed sites, while affording the user a single point of access, and, that specialized access points (*clients *in OGC parlance) can be built by anyone. The nature of our work will lead to the broad publication of global multi-phenomenological maps of lunar and planetary bodies. These maps will be hierarchical and afford both synoptic and high-resolution views of each body published. The multiple views will make clear what differences exist, and, because they present the different phenomena as a composite, this capability should provide a basis to better understand planetary processes.*

INTRODUCTION

The OpenGIS© Consortium (OGC) has developed a broad complex of specification and implementation for the interoperable publication and sharing of Earth related geospatial data [1]. Developed originally in support of commercial applications, the Earth science community is now capitalizing on these technologies to support the *scientific* publication and analysis of geospatial data [2].

The planetary science community is now widely using commercially available GIS tools for the scientific mapping, categorization, and analysis of planetary data [3]. However, there is not the critical mass of support that delivers planetary data to these tools using the OGC developed means of interoperable publication and retrieval. The analysis techniques of GIS are being embraced, but the infrastructure has been left undeveloped.

The availability and accessibility of planetary science data products to all members of the science community has greatly improved over the past decade. In particular, the incorporation of World Wide Web access to the data archived in NASA's Planetary Data System (PDS) has revolutionized the data discovery and retrieval process. However, the PDS-available products are primarily the lower level data archives (Experiment Data Records and Reduced Data Records) and not the higher-level aggregated and/or gridded products needed for comprehensive studies [4]. Moreover, the explosion of these low level archives from 3- to nearly 300- TB in just the last few years now acts as a formidable *barrier*, rather than *enabler* of planetary science, requiring a huge effort on the part of the individual scientist. The day when meaningful science could be performed simply by browsing these vast archives has all but vanished. Now every researcher has to produce higher-level products and base the ongoing science on them, rather than on the appeal to individual records alone. Yet, there is no generally accepted methodology for constructing and publishing these higher level data products so that the community can interact and build collectively upon the work as it proceeds. Rather, multiple high-level products are generally produced ad hoc, as they are needed, and shared on informal bases.

What is needed is the development of an overall meta-architecture for the PDS permitting publication of higher-level data products in ways that the data are both easily combinable with each other, and continuously, dynamically accessible to the community. As the planetary community moves into the era of data abundance, it is urgent that an architecture permitting the transformation of this abundance into published 'characterization', and ultimately a 'model', be properly crafted and put into place now.

Fortunately, the planetary community – and in particular the planetary geo-science community – can benefit greatly from work already done by their Earth science counterparts, the Earth geospatial community. The OpenGIS© Consortium (OGC) has developed a broad complex of specification and implementation that plays an increasingly important role in linking the similarly huge distributed reservoirs of Earth related geospatial data. With provenance dating from the early 80's, the OGC is now comprised of over 200 international government, university, and commercial organizations. In particular, NASA's Earth Science Enterprise – via now the Geospatial Interoperability Office – has supported and been active in OGC development since 1994.

Thus we are pursuing a body of work that will provide the individual investigator with comprehensive access to high-level products and, at the same time, afford a publication venue for communicating new results. Investigators need a unified methodology for *accessing* higher-level data products that serve as substrate, background, and current known detail pertinent to their ongoing work. As stated earlier, the archive for these data products does not currently exist. While the Geoscience and Imaging nodes of the PDS produce and archive *some* higher-level products, many other products are produced elsewhere by the community and do not make it into any official archive. The goal here is that our WMS and WCS servers will ultimately be placed in operation at these PDS discipline nodes. By so doing, we assure that these valuable data can be managed and curated by the PDS. In addition, once the basic data products are in place, investigators can not only use them directly in creating their derived products, *they can publish their own results within this same framework, adding to the depth of knowledge available*. The result is that the community is afforded a single, unified point of access for the entire collection.

BACKGROUND

The JPL WMS Server

In early 2000, an early Landsat mosaic of the Continental U.S. was adopted for use and hosted in accordance with what was then called the Web Mapping Testbed (WMT) Server Specification. The WMT Specification, once adopted, became WMS. Since our server was one of the first open, on-line WMT providers, we were requested to participate in a series of NASA-led demonstrations of it. Once the WMS standard was ratified, our WMT server was updated to be in compliance with it, becoming a cornerstone of this technology. Building on the success of this first server, we have been improving the technology behind the WMS server, and at the same time, providing a global 15m-resolution coverage using NASA purchased Landsat7 2000 data set [5]. Many aspects of the evolving WMS standard are perfectly suited to showcase a multispectral sensor such as Landsat; accordingly a second implementation of a WMS server was completed in 2003.

The result of this implementation is a fully featured raster data WMS server, many of its features being unique in the field. This server has already been deployed on a public access Internet server, initially hosting more than one terabyte of data, and providing more than one hundred thousand custom maps per month. With the completion of the WMS Landsat Mosaic in March 2004, this server, currently hosting over 5 terabytes of data, now provides full world coverage, high resolution map image data, to any WMS client application [6].

WMS Functional Description

The JPL OnEarth WMS server, which will be the basis for the new OnMars and OnMoon servers for planetary data, implements the latest version of the WMS protocol, including Style Layer Description (SLD) support for image (raster) data [7].

The most notable features are:

- **Multispectral data support** The server implements the band selection scheme provided by the Style Layers Descriptor (SLD) standard, and in addition, it can perform the specific Landsat operation known as *pan-sharpening*, in which a high resolution panchromatic band is used to sharpen the lower resolution, narrow band multi-spectral images.
- **Multiple projection support** While the data are being stored in one projection only, the server is capable of producing map images in an assortment of projections such as Orthographic, Gnomonic and UTM.

  Other projections can be simply added by providing the code to implement the coordinate transformation
- **Multiple layer support** Multiple layers (datasets) can be specified in one request, each one of them having a specified transparency model, the resulting image being a single composite view. This approach is much more powerful than reliance on the client for compositing support. This is not a metaserver feature, all the layers need to exist on the same server.
- **Multiple styles** Any layer can have many predefined presentation styles, in which the server administrator can provide access to certain features of the data. Any supported feature can be included in such a style. For example, a grayscale image set can be colorized in some predetermined way, and this color representation be made available as a predefined style. Additionally, a specific contrast, brightness and scaling can be specified, matching a certain type of display or print media. This feature is implemented as a software function, any number of styles can be generated
- **SLD support** A companion specification to the WMS standard, an SLD provides a standardized description of a style, controlling features such as band selection, image normalization, and color transformations. Since the location of an SLD is specified using a URL, a separate SLD file can be retrieved and interpreted for each request. This gives a potential application full control over any data transformation that the server is capable of providing.
- **Multiple streaming output format** An important design feature of the server is the generation of output maps as they are being sent to the requestor, without using any server side storage. This contributes to the ease of administration and improves the response times, since the map image starts to be sent to the requestor before the whole map is completed. Both lossless and lossy formats are supported, as implemented by the JPEG and the PNG formats, providing various levels of quality/size ratio.
- **Performance** The server implementation makes full use of available parallel processing capabilities of the host machine. The response times depend on the nature and the complexity of the requested image as well as the server machine capabilities, usually being less than one second under the loads we have been experiencing, 100,000 request/mo.

Our efforts with respect to the global mosaic can be explored at http://onearth.jpl.nasa.gov/. An explicit discussion with dynamic examples of the server's features listed above can be found at: http://onearth.jpl.nasa.gov/testing.html.

EXTENSION TO PLANETARY DATA

The main strength of the OGC approach resides in the existence of well-supported collaborative standards that allow individual development of components. Since these standards are well supported by the GIS industry, once enough planetary data are present on the Internet in this form, many of the very powerful and complex commercial applications developed for Earth science can be directly utilized.

Our overall approach, then, is to first develop baseline servers of lunar and planetary data, host multiple Level 3 datasets on them, and deploy them for use in conjunction with our PIGWAD and JMARS use-case applications. In parallel, we will work with WorkGroup IV/9 -Extraterrestrial mapping of the International Society for Photogrammetry and Remote

Sensing [ISPRS] and with the OGC committees to support a NASA led effort broadening the OGC standard to embrace the lunar and planetary case. With the servers in place, the use cases demonstrated, and a coherent set of standards adopted, we are confident that the critical mass will have formed, and additional servers and applications will begin to appear as they have in the Earth science field.

WMS/WCS Servers and VCG

The centerpieces of this project are the Mars and Lunar Topography & Image Servers (M/LTIS). These servers will host the various datasets as shown in Table 1, and will be capable of performing the specific functions listed in the left-most column of Table 2. It is important to realize that, while many commercial GIS applications now contain interfaces for access to WMS data and can directly utilize them, the WMS specification supports server custom processing of the fundamental datasets hosted for specific customer delivery requests. Thus, while the commercial GIS applications themselves can perform sophisticated client-side processing, emplacement of custom, on demand server side processing can offer processing options and analysis capabilities independent of any specific connecting application. It can be used to add new functions, or processing, specific to a certain dataset in a consistent, recommended way, without any changes to the client side application.

Figure 1 below represents an overall description of the envisioned system of prototype elements and how they will work together.

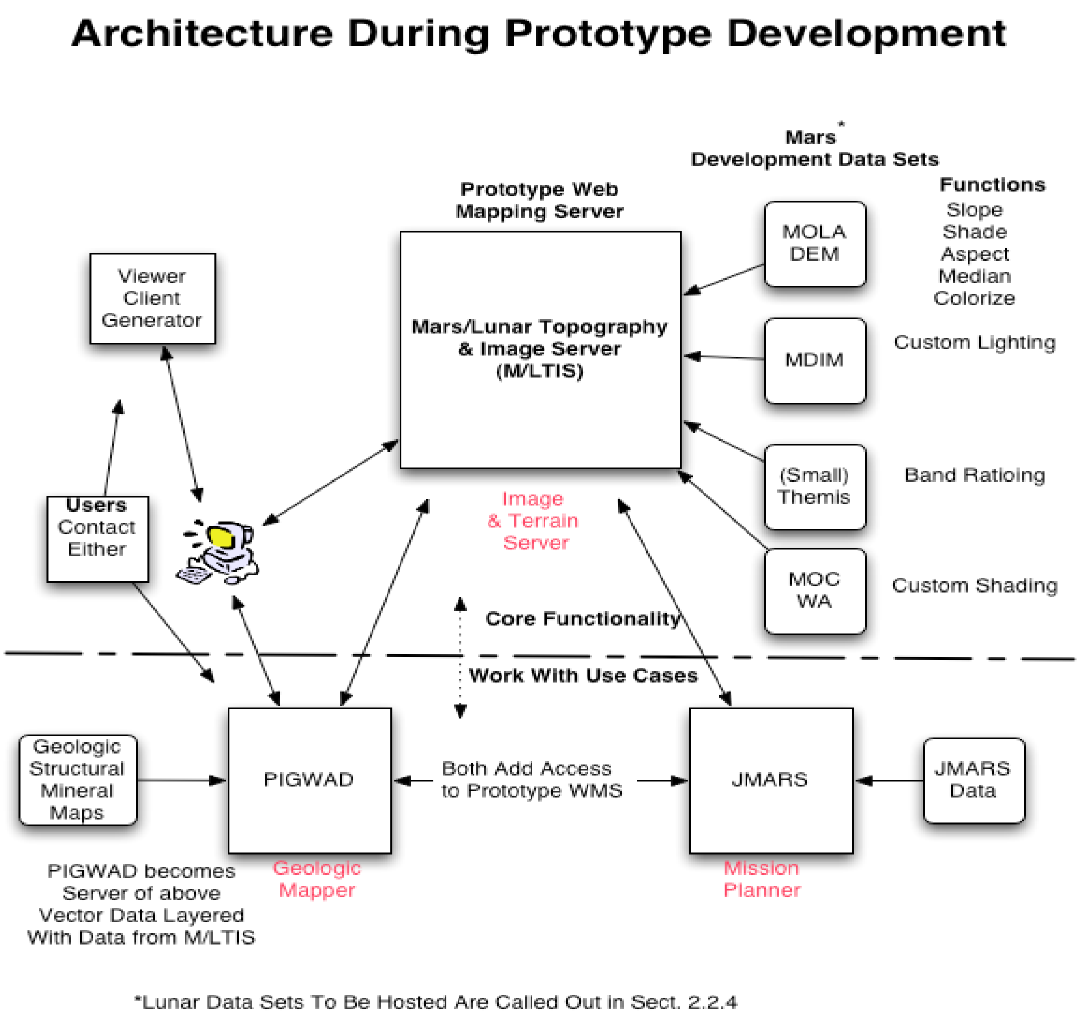

Fig. 1 Conceptual Block Diagram of the Project Elements, Their Functionality, & Interactions

In order to make the WMS server universally available even over a simpler Web browser, a 'helper application' – in OGC parlance, a Viewer Client Generator, VCG – must be configured. The VCG appears to the web based client as a website. Exercising this website in effect fashions URL commands which are returned to the browser for forwarding to the WMS server, invoking its functionality. See Fig. 2 (a). The requested custom image is then delivered directly to the initiating user and presented in his browser. The VCG application is not involved in any direct map data manipulation, being just a navigation engine. *It is this separation and independence of function that gives the OGC protocols much of their power*. Once a server is in place, anyone can develop a VCG that can invoke the server's functionality in specific ways dictated by the VCG's implementer and make both it and the server available to anyone. Additional servers hosting other datasets can be added independently as well; a knowledgeable VCG can fashion coordinated commands to each and a composite image product will be delivered to the user. In this manner, very large and sophisticated distributed data stores can be grown and exploited in independent, entrepreneurial ways. In our local case, we will be concentrating on the server side but build enough VCG functionality to test and demonstrate the server's capabilities. This prototype VCG will also serve as a user manual for the server's capabilities, providing guidance for anybody else that needs to implement a VCG.

GIS applications, when they are *OGC-enabled* can dispense with the VCG and bring its functionality directly into the application in the form of a GIS/OGC interface. See Fig. 2 (b). As shown, the data flows have the same character and meaning but are somewhat simpler. It is important to recognize that the server sees the same commands in either case and need not be knowledgeable as to which is configured.

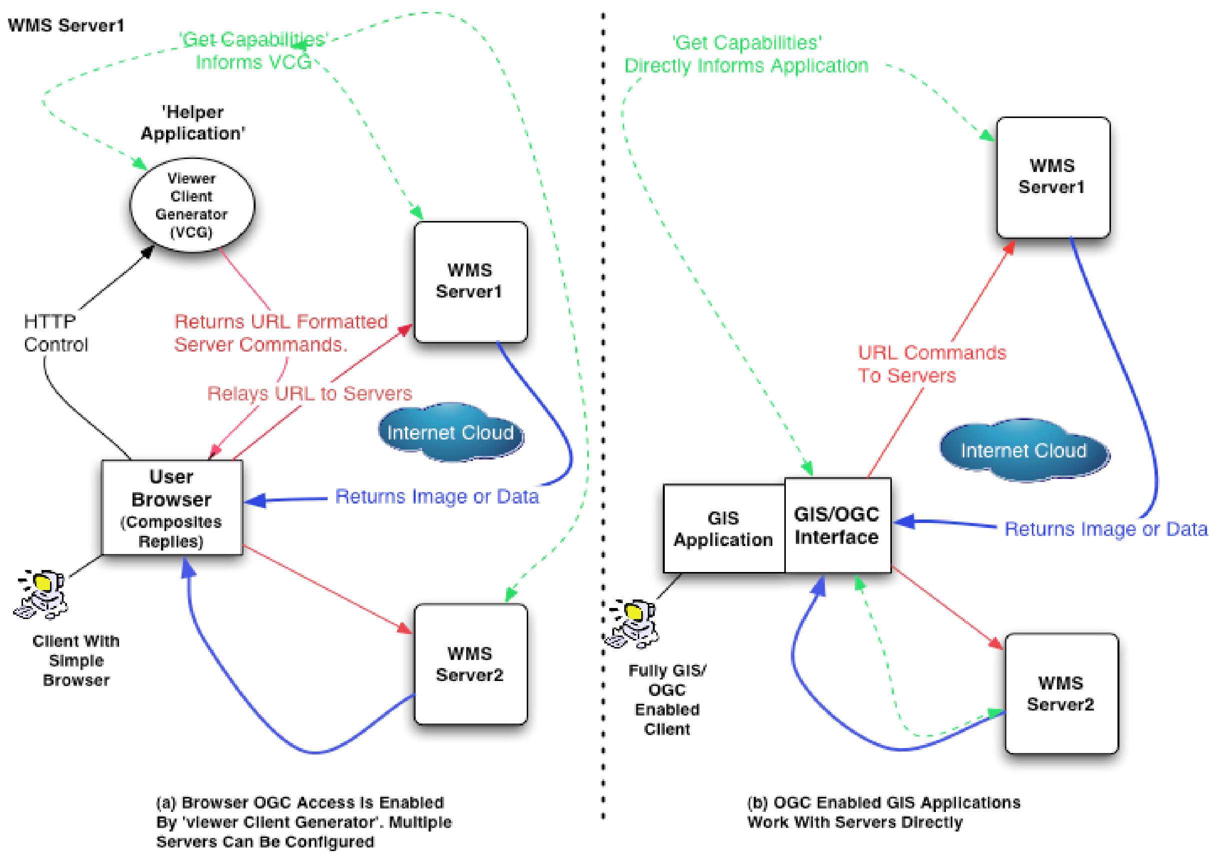

Fig. 2 Basic Information Flow For (a) Web Browser and (b) GIS Application Clients

In advanced configurations, multiple servers may not be individually capable of performing all the data manipulation and server side compositing that is required. If so, the OGC specification permits configuring a meta- or *cascading server* to communicate with the WMS servers directly and operate upon and composite their returns (not shown here but to be demonstrated in the PIGWAD use case as discussed below).

In contrast with a WMS server, which is mostly concerned with generating an image of a map, a WCS server contains a much more refined mechanism for selection and retrieval of data. While it is still using HTTP as the transport mechanism, it has no direct connection with a simple web browser, the target applications being complex scientific or GIS applications. Several WCS server implementations exist, usually built around a database engine. As a critical volume of planetary datasets becomes available using the WMS protocol, and as planetary GIS applications become more complex, a WCS server will be required. In the third year of this effort, we will assess and adopt an existing WCS implementation and deploy and adapt it for planetary use. A suitable client, most probably a commercial GIS application, will be demonstrated in conjunction with this server.

PIGWAD and JMARS Use Cases

Another class of WMS client is represented by special purpose applications, not necessarily related to the web browser. These applications contain a fully functional WMS interface, making maps generated by a WMS server appear as a direct data connection. They usually offer much more functionality than is possible with a generic web browser, yet the same WMS protocols are used to interrogate the server about it's data holdings and capabilities, and retrieve the user-selected maps.

In the field of planetary science, there exist several web-enabled tools from simple image maps, more advanced java tools or full-fledged applications. However, none of these applications will meet the needs of all researchers. *We stress that our proposed OGC system is not an attempt to create a single tool that will fit everyone, but rather a backbone to support several diverse applications.* The availability of these WMS enabled client applications is growing each year. Many powerful commercial and Open Source GIS, remote sensing, and other geospatial applications can now use WMS streamed datasets as just another layer. Thus, the investigator can still choose the application that will match his research needs.

We are currently working with selected planetary facilities to help them incorporate the WMS served data within their applications. The OGC standards are well documented and the methods to request data from the servers are straightforward. Arizona State University has written a Mars mission-planning tool called JMARS [8]. This group has expressed interest in supporting WMS protocols so that they can utilize the data in other servers. Of course, the PIGWAD [Planetary Interactive G.I.S.-on-the-Web Analyzable Database] mapping site, supported by members on this team who specialize in GIS applications for Mars, will immediately use this server within its supported interfaces. As these two facilities (and others) leverage the datasets served from the planetary servers, they will be freed to concentrate on their local and unique expertise.

These principles are illustrated by the proposed PIGWAD use case. PIGWAD's current implementation is built around a commercial web-based version of ESRI GIS applications ArcIMS and ArcView IMS. Currently PIGWAD hosts its own planetary datasets locally and adds to these their locally generated store of geological and mineral maps, etc. The most basic transformation we will effect in the PIGWAD case is to substitute web-based appeal to the Mars and Lunar Topographic Image server, replacing the local store and curation of those data products. This mirrors the current situation in many planetary science centers – personnel must locate and transfer to local storage the products they need for their investigations. Not only are these products likely to be uncoordinated – different projections, casual co-registration, etc.– but also as they evolve and are improved, the local curation is burdened with the additional tasks of staying current. All of this will be circumvented when these datasets are available on the web and are hosted by the institution of their original creation and/or current curation.

The PIGWAD use case will also demonstrate the ability to 'cascade' the servers. PIGWAD will receive data from the Mars server and will continue to overlay that data with its own image and vector data. Data from both sites can then be delivered to a user employing the PIGWAD as a VCG and operating in the web browser client mode. In the development of the PIGWAD use case, ESRI has offered their assistance to provide any ESRI software, including needed customization assistance without fee [9]. Using commercial GIS applications and JMARS, as OGC–enabled clients, we will demonstrate the ability to receive data from the Mars server and also overlay data from their own systems.

Technical Support to Broaden the OGC Standards

We will provide technical support for this essentially programmatic goal of broadening the OGC standards to encompass the planetary case. Officials from NASA will lead the process, which involves petitioning the OGC to consider this broadening of their current scope and goals. We have discussed this with the Geospatial Interoperability Office, NASA's current liaison to the OGC. They have created a Lunar and Planetary Ad Hoc working Group as the first step towards official incorporation of planetary reference systems and other standards into the body of recognized OGC definitions and procedures. This development will be fully coordinated with the WorkGroup IV/9 - Extraterrestrial mapping of the International Society for Photogrammetry and Remote Sensing [ISPRS], to make sure that the community requirements are reflected in the resulting standards. These standards extensions fall mainly into two categories: projections/coordinate systems and standards for processing data.

A significant obstacle to use of GIS tools for planetary science purposes is represented by the lack of coordinate system encodings. We will generate and publish via OGC channels an encoding scheme for the most common coordinate systems used for Mars and Lunar datasets. Since OGC members encompass the research, commercial and freeware community, the presence of this encoding scheme will encourage compatible development of applications.

The current mechanism for controlling the aspect of mapping products generated via WMS is controlled by the SLD. A common set of controls is defined, but discussions within OGC are already underway to refine and enhance this mechanism. While the SLD specification provides a standard way to control the appearance of data, it is not intended to standardize controls for server side processing. Such processing is essential for enabling scientific use of the data, and will usually constitute an SLD extension particular to each server and application. Several envisioned use cases of planetary data require this type of processing.

Since the JPL WMS server is one of the most refined SLD enabled applications, members of development team already participate on the SLD working group that is tasked with further refinement and development of the SLD standard. Because of this, we will be able to provide a planetary science perspective to this working group, conveying its requirements as part of the standardization process, and ensuring further development of the SLD and server side processing standards are aligned with the planetary science requirements.

Datasets	**Description**	**Year**
Mars		
• MOLA Topography	Global topographic DEM 500 m/p	1
• Color MDIM 1	Global color Mars digital image mosaic (Viking) 1km/p	1
• MDIM 2.1	Global b/w image mosaic (Viking) 250m/p	1
• MSSS Atlas	Global b/w image mosaic (MGS) 250 m/p	1
• THEMIS (thermal imager)	Example area mosaics 3 to 11 bands (MO) 100 m/p	2
The Moon		
• Clementine UVVIS 750nm	Global b/w mosaic 500 m/p	1
• LIDAR Topography	Lunar topographic DEM ~2 km/p	1
• Clementine UVVIS (other)	Add in the four remaining UVVIS bands besides 750nm	2
• Clementine NIR	Global 5 band near-infrared mosaic 500 m/p	2
• Lunar Orbiter	Global b/w mosaic ~65 m/p (100 microns)	3

Table 1. Data Sets to be Hosted On the Servers and Their Schedules

Data Sets

Data sets to be hosted are listed above in Table 1. The image bases we plan to implement will demonstrate the server capability to handle very different data types. For example, the MDIM 2.1 and MSSS Atlas demonstrate the ability to serve single band, several gigabyte large datasets. The MOLA and LIDAR datasets, global topographic map of Mars and the Moon, will exhibit the ability to handle 16 bit data types and the option to apply stretched color maps. The THEMIS and the Clementine datasets demonstrate working with multiple band datasets.

Server Functions

Along with supporting the standard OGC WMS protocols, we will also enable many image processing functions that can also be used by the client. This list includes functions that can be applied to topographic layers including: colorized, slopes, hillshading, aspect, median and potentially other filters. For satellite images we will allow the user to set one of several stretch types; min/max, standard deviation, histogram equalization, guassian, etc. Multiple band raster sets can be mix into RGB (red, green, and blue) channels to create color images or to help emphasize features. We also will support band ratioing and image mixing using transparency routines.

Most of these processes are standard means to help the researcher extract information from the datasets. For example, shading topographic datasets from several different angles can highlight structural features, which may not have shown from only angle. Large median filters on a topography, also known as de-trending, will help remove the regional slope leaving shape localized topographic signatures. Band ratioing is critical to geologic mappers and can be applied to Lunar Clementine and Mars THEMIS images.

To elaborate on why this is important, the images from Clementine comprise the first color global digital data set of the Moon captured in 11 spectral bands at ~200 meters per pixel using a UV/Visible and a near-infrared (NIR) imaging system. The coloration of the Moon in the visible to near-infrared spectrum is sensitive to differences in both the composition of the surface material and the amount of time material has been exposed. The different bands characterize parts of the spectrum that are known to contain absorption bands diagnostic of iron-bearing minerals and plagioclase feldspar. Combining and/or ratioing these spectral bands can be essential to help map the distribution of rock and soil types on the Moon.

The precise representation of data in WMS is controlled by the Styled Layer Descriptor specification. The current version of this specification does cover some of the features we identified in the previous paragraphs. These features will be supported by the WMS servers in the first year of the project. Additional features, such as band ratioing, will be implemented in the second year of this project, and the implementation specification will be documented with OGC.

Processing Routines	**Description**	**Year**
Topography		
• Colorize	Assigning color maps	**1**
• Hillshade	Simulated shading from user defined angles	**1**
• Slope	Angle of surface pitch in degrees	**2**
• Aspect	Direction of slope in degrees	**2**
• Median filter	Help to removes regional slope of features (de-trending)	**2**
Imagery		
• Band mixing	Build a RGB image from user selected bands	**1**
• Transparency	Blending of datasets	**1**
• Shading	Shading imagery with topographic layer	**1**
• Standard stretches	Min/max, contrast, linear, standard deviation	**1**
• More advanced stretches	Histogram equalization, Gaussian (normal curve)	**2**
• Band ratioing	Band division	**2**

Table 2. Server Functionality With Schedule

EXPECTED IMPACT/BENEFITS

Because GIS is an expansion of cartography and mapping, the use of the technology fits naturally into planetary research. It has been used in planetary research for a number of years but recently, we have seen its use rapidly grow. This growth can be related to increased opportunities for GIS training, increased availability of GIS applications, a higher-level of maturity of GIS tools, and a greater volume of digital planetary datasets, requiring a solution that can handle multiple and diverse datasets. This capability to overlay data has been significantly enhanced by the OGC's specifications to use streamed geospatial data from WMS servers.

Planetary researchers must now juggle the dozens of datasets available. The ability for Mars geologic mappers to compare these layers within one interface during the mapping process is crucial to help express their geologic story. To gather and prepare these datasets into a usable form can take weeks to months before the researcher can even start. Once our WMS servers are in place, the datasets can be accessed by the researcher *within a matter of seconds.*

A prime application for our servers is the creation of geologic maps for Mars and the Moon. Because of the wealth of new data now available for Mars and the Moon, the current geologic maps are in desperate need of updates. There is now a great deal of interest in the planetary community to begin the systematic geologic and structural remapping of Mars and the Moon. Depending on the mapping scale, there would probably be at least thirty quads per body requiring numerous researchers. Our project will directly contribute to the efficiency of this research and research projects like it. Investigators would be assured they would be working with the best data available already co-aligned and projected for ease of use. Once their derived maps are completed and reviewed, they could also be added to the WMS catalog of available data sets whether they are raster or vector.

RELEVANCE TO PLANETARY MISSIONS

In general, our effort directly relates to *all* NASA missions to terrestrial planets, and, specifically in this prototype, to the Moon and to Mars. Planning these missions will require data of the utmost accuracy and detail. As a specific example, our work is directly relevant to NASA's Office of Space Science ambitious program of Mars explorations and associated technology development. The current missions of the Mars Exploration Program, the Mars Global Surveyor (MGS), Mars Odyssey, and the Mars Exploration Rovers (MERs) have together returned an unprecedented volume of data. Future missions, such as the Mars Reconnaissance Orbiter, will move from the realm of Gigabits to Terabits (26-Gb to 34-Tb) of raw returned data. As part of the exploration strategy, each mission feeds forward into future activities. The wide variety of data sets collected to date and those to be obtained from soon to be launched missions, will provide a basis for continued exploration. Specifically, in addition to supporting overall scientific analyses, orbital data has a key goal in providing a basis for the selection of landing sites for the upcoming Phoenix and Mars Science Laboratory (MSL) missions. As demonstrated by landing site selection for the Mars Exploration Rovers [10], this process involves detailed analysis of multiple data sets. Although many of these data sets are available via the PDS, they often are only available as individual data records, and not as compiled or higher level products. In addition, there is a need to deal with data in different projections (planetocentric vs. planetographic). As such, the MER landing site selection process required data manipulation by nonstandard means. It is clear that the ability to access and manipulate multiple co-registered data sets is important for both mission planning purposes and scientific analyses.

This proposed effort becomes even more important in the light of the new U.S. Presidential initiative of returning to the Moon and sending humans to Mars. Not only do we need good quality maps to plan these missions (assess resources, landing sites, etc.), but also, once humans are involved in landing and working on these bodies, the existence of detailed, topographic, and geologic maps is a critical need. We can create them by hand as we do now, we can use stand-alone GIS applications as we do now, or we can create a system of standardized products that interoperate and serve many needs. It is this latter choice that we have begun to implement.

The nature of our work will lead to the broad publication of global multi-phenomenological maps of lunar and planetary bodies. These maps will be hierarchical and afford both synoptic and high-resolution views of each body published. The multiple views will make clear what differences exist, and because they present the different phenomena as a composite, this capability should provide a basis to better understand planetary processes.

Also, we fully anticipate that these browseable, synoptic, as well as detailed planetary views will become a source for *educational and literate public education* about the Moon and planets. It is a unique characteristic of the OGC standards, once the servers are in place, that they permit individual VCG's to be crafted for different constituencies. The OGC compliant publication we enable will be equally accessible to scientist and student alike.

Finally, at the end of the day, when the missions have ended, and the instruments have ceased collecting data, what is left to us, as legacies, are *the data sets*. They are the reason the missions fly. Enabling the entire planetary community ***to increase the utility of these data, to provide higher-level derivative data sets, and to transform the data into physical knowledge of the planet*** are the main objectives of this work.

ACKNOWLEDGEMENT

This work is being performed at the Jet Propulsion Laboratory, California Institute of Technology, and the US Geological Survey, Flagstaff, AZ, under a contract with the National Aeronautics and Space Administration.

REFERENCES

[1]. OpenGIS Consortium, Vision & Mission, http://www.opengis.org/about?page=vision

[2]. Nittel, Sylvia "Interoperable Data Services for Earth Science Data" available at: http://www.esipfed.org/knowledge_center/interoperable_data_access.pdf

[3]. Hare, T. M., Tanaka, K.L., Skinner Jr., J.A., GIS 101 For Planetary Research, "Advances in Planetary Mapping 2003", International Society for Photogrammetry and Remote Sensing, Lunar Planetary Institute, Houston, Texas, March 2003.

[4]. Dobinson, E., Planetary Data System: Challenges and Solutions, CNES Data Archiving Workshop, Toulouse, France, November 2002

[5]. GeoCover 2000, Landsat 7, NASA Scientific Data Purchase, http://www.esad.ssc.nasa.gov/datapurchase/

[6]. Plesea, L., WMS Global Mosaic, OGC Technical Committee and Planning Committee, United Nations Headquarters, Jan 2004

[7]. OpenGIS Project Documents OGC 02-070 Styled Layer Descriptor, OGC 03-065 Web Coverage Service, OGC 01-109 Web Map Service, http://www.opengis.org/specs?page=baseline

[8]. Gorelick, N. S., Weiss-Malik, M., Steinberg, S., B., JMARS: A Multimission Data Fusion Application. Lunar Planet. Sci., XXX, Abstract #1849, Lunar and Planetary Institute, Houston (CD-ROM).

[9] Dangermond, J. , ESRI President, Unpublished letter to Dobinson, Mar. 1, 20004.

[10] Golombek, M. P., Grant, J. A., Parker, T. J., Kass, D. M., Crisp, J. A., Squyres, S. W., Haldemann, A. F. C., Adler, M., Lee, W. J., Bridges, N. T., Arvidson, R. E., Carr, M. H., Kirk, R. L., Knocke, P. C., Roncoli, R. B., Weitz, C. M., Schofield, J. T., Zurek, R. W., Christensen, P. R., Fergason, R. L., Anderson, F. S., and Rice, J. W., Selection of the Mars Exploration Rover landing sites, *Journal of Geophysical Research*, 108, 13-1 – 13-48, 2003.