THE MARS PATHFINDER SCIENCE DATA PROCESSING SYSTEM

Allan J. Runkle, Jean J. Lorre, Douglass A. Alexander, Susan K. LaVoie, Pamela J. Woncik, Elizabeth D. Duxbury, James M. McAuley, and William B. Green,
Jet Propulsion Laboratory,
California Institute of Technology
Pasadena, CA

Abstract

An integrated data processing system was used to support the science payload instruments on Mars Pathfinder, and was developed and operated at very low cost through introduction of new technology. The system provided level 0 processing of all payload data (from the Imager for Mars Pathfinder, the rover camera systems, the Alpha-Proton X-Ray Spectrometer, and the Meteorology experiment) in real time as the telemetry was received on the ground. The system also supported mission critical activities by performing rapid higher level processing of selected data. Computer generated mosaics, including stereo mosaics, were produced within minutes of data receipt on the ground to support rover deployment and navigation requirements. Color mosaics were produced rapidly to meet press release and science planning deadlines. An integrated data base management system provided science team access to data within seconds following completion of construction of science data records. The same data management system supported requirements of operations personnel, press access to released data in various formats, transfer of data rapidly to the world wide web for public access, query and retrieval requirements of science team members located at JPL or at their home facilities around the world, and preparation of final archival products for submittal to the Planetary Data System.

This paper will describe the system developed to support Mars Pathfinder, the technology that was used to provide sophisticated products at very low cost, and the variety of data products used to support Mars Pathfinder operations. This paper represents one phase of work performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

1. Introduction

JPL's Multimission Image Processing Laboratory (MIPL) has been a pioneer in Image Processing for over thirty years. During this time, MIPL has been given the responsibility for the design, implementation and operation of science data processing systems supporting NASA remote sensing spacecraft. Mars Pathfinder (MPF) is one of the latest missions in which MIPL has been a participant. The science and operational objectives designed for this mission required new automated capabilities for supporting rover deployment and navigation, and science planning activities. It is these capabilities, their design, adaptation and implementation that are described in this paper.

A general data flow for the science payload data from the Pathfinder lander and Sojourner rover is shown in Figure 1. Telemetry data from the lander was received by the Deep Space Network (DSN) then sent to the Mars Pathfinder Ground Data System (GDS) for validation and staging. The MIPL realtime subsystem queried the MPF GDS for the science telemetry and assembled the packet level data and ancillary ground data into science data records for each instrument's observation. The realtime subsystem populates the science database with the initial data records and associated index data, enabling user query and retrieval from a web-based system called the MPF Navigator. Remote science users were able to access the instrument data via a JPL/MIPL-developed capability called the "File Exchange Interface" or "FEI".

MIPL's operational science processing included automated mosaic production for mono, stereo or color map projections, generation of stereo data products including rover terrain visualizations and anaglyphs, specialized products for use as press releases, and video animation and rendering of data visualization products. As part of the data distribution role, data products were transferred to individual science teams, to the flight operations team, and to the Sojourner operations team, to a central web server used to support public access to released data via the Internet, and to a specific web site used by the press to access press release information. Data were released in a variety of standard digital and analog formats.

The remainder of this paper describes the processing performed in the areas shown in Figure 1. The processing was performed utilizing facilities within the MIPL and the Digital Image Animation Laboratory (DIAL). DIAL is a system that is interfaced electronically with the MIPL facility and digital mission databases maintained by MIPL. DIAL includes specialized equipment for production of science visualization and animation data products.

2. System Architecture

In 1993 the Mars Pathfinder project enlisted MIPL to supply the science data processing capabilities needed to perform surface operations and data archiving. By July 4th, 1997, MIPL's responsibilities were expanded to include the telemetry processing, distribution and archiving of data from the four primary science instruments, and supporting surface operations through production of automatically generated stereo and color mosaics in a variety of projection types for use by Lander, Sojourner and APXS science and operations teams. The realtime displays and automatically generated mosaics were used for time-critical decisions on spacecraft health assessment, rover navigation, evaluation of science objectives, and fulfilling NASA's goal of rapid public release of science data. The major objectives of the architecture of the system that supported these capabilities were:

Minimizing Cost

The MIPL science data processing system at JPL has continuously evolved over the past three decades by utilizing advances in computer technology. By 1992, the basic infrastructure supporting science data processing was operational on a variety of UNIX platforms. The conversion to a UNIX system also allowed for greater portability, lower hardware cost and the implementation of the client/server paradigm for a number of key services, database management, and data distribution. The high-level Mars Pathfinder system architecture was developed by adapting the existing MIPL capabilities. In cases where the Pathfinder capabilities required different processing, components supporting other missions were used to create the specialized capability.

Automated Product Generation

Mosaicking many small images to create a larger view is not new technology. Supporting critical flight operations activities (deployment of Sojourner rover, rover navigation and Lander science data acquisition strategy development support) required the development of an automated processing system capable of mosaicking hundreds of images within minutes of receipt of data on the ground. A database driven system based on MIPL's Multimission Data Management Subsystem (MDMS) was designed to support this capability. MDMS is a client/server package that generates SQL queries to a commercial database management system (Sybase), on behalf of the user or application, and isolates the lower level database from the rest of the environment.

Adapt Telemetry Processing.

The telemetry processing system for Mars Pathfinder was based on adaptation of the basic framework and support libraries used within MIPL's multimission system to support Galileo and other planetary missions. Changes had to be implemented to account for the different telemetry packets, decompression routines and metadata handling, although the design and external interfaces remained the same.

Data Archive Requirements.

NASA required that all final mission data be archived by the Planetary Data System (PDS). MIPL produced science products complying with PDS format and content standards. The PDS image format was defined as the standard transport format between the different Pathfinder science teams.

Data Distribution/Security

The baseline MIPL system includes the File Exchange Interface (FEI) which is a client/server implementation of a secure file management and distribution system. It uses a Kerberos authentication mechanism and database system to support its capabilities. The FEI capabilities allow for selected file transfers as well as automatic transmission of newly generated files.

User Interface

The wide diversity of science team members on Mars Pathfinder meant that their own image processing software use different image formats, different protocols, and different operating environments. Because the team members were consumed with their own tasks, it became apparent that there would not be adequate time to train the team members in the specifics of the MIPL image processing environment or any Pathfinder specific tools. For this reason, it was decided to implement a Web-based browser that would allow an easy way to query the database, browse images, generate FEI fetch command lists, and construct simple products in a variety of different formats. The existing web-based Planetary Image Atlas provided the core capabilities and was adapted for Pathfinder.

3. Realtime

The purpose of MIPL's Realtime subsystem is to obtain instrument telemetry packets from the project's Ground Data System (GDS) and build an Experimental Data Record (EDR) representing the sensor, engineering and ground data corresponding to an observation. Once the EDR has been created, it is quickly displayed for an initial assessment of the instrument's health and operation. As the software supporting these capabilities executes, anomalies, and tracking information are recorded into a running log and displayed at the Realtime Engineer's console to aid in error diagnostics and operational status.

The design of the realtime subsystem has its roots in the realtime system that supported the Voyager 2 fly-by of Uranus and Neptune. The subsystem consists of three mission independent programs and two types of project specific programs, (telemetry processing, and instrument display). For Mars Pathfinder, there were four project specific programs, the telemetry processor, MPFTELEMPROC, and the three instrument displays. These display programs were the first view of the data and is used to determine the spacecraft and instrument health, verify the proper operation of the realtime processing, and rapidly distribute the view of the images to the science team members and an interested public.

The critical component of this subsystem is the telemetry processor. It is responsible for requesting and processing the telemetry from the project's Telemetry Delivery System (TDS), the source of all the mission's telemetry data. Once the TDS starts sending telemetry packets to the telemetry processor, the packet's corresponding observation is identified and the science data is inserted into the appropriate location of the EDR. Normal processing also performed is an accounting of the telemetry associated with each EDR, an analysis of the basic statistics of the science data (e.g., data range, average value, standard deviation, etc.) and cataloging the EDR for easier identification and retrieval. Depending upon the instrument and operating mode, the realtime task may require additional capabilities to decompress the packet data, and integrate command parameters not available from the telemetry.

In reality, nothing works like it does in theory; as telemetry travels the millions of kilometers from Mars to Earth problems occur. The most common problem is that packets are lost. Since each observation is made of several packets, the telemetry processor must determine when it is no longer useful to wait for the missing packet(s) and create the EDR even though it will contain gaps. Occasionally just the opposite happened, the same packet was received more than once. This was a caused by the procedural errors in the GDS system and multiple transmissions of critical packets from the spacecraft. These were simply ignored because the error encoding performed on all of the telemetry by the spacecraft, guaranteeing the packets would be identical.

Mars Pathfinder was a very capable spacecraft, operational procedures also complicated the telemetry processing task. The spacecraft had the ability to stage hundreds of images into eleven different transmission queues. Depending upon the objectives for the day, any number of these queues would be allocated a percentage of the down-link bandwidth. This usually meant that 4 to 5 different images would be multiplexed in the telemetry at a given time. The telemetry subsystem required the additional capability to maintain up to eleven images simultaneously as it integrated and processed telemetry packets with the correct image.

4. Database Design

The need for timely turnaround of MPF instrument data to meet operational and science team objectives required only a modest retooling of the MIPL database systems used in previous Flight Projects. The end result was a highly efficient system that extended the role MIPL played in the trafficking of MPF data. It provided operational support to the telemetry processing and data file distribution subsystems, and was used by automated mosaicking and archive mastering software. The science teams accessed the database via a Web-based data browser, generic database applications for specialized tasks, and client applications of the file distribution subsystem.

The database design incorporated the operational characteristics of the instruments and the ground data system, and was based on the projected amounts of data to be inserted, updated, queried and extracted. The data stream began with the telemetry subsystem, which returned the observed science data and small amounts of related engineering data, and cataloged metadata needed to support scientific investigations and operational tasks.

To adequately support a given instrument, database designs typically relied on two types of metadata, commands and observation parameters associated with the EDRs. The command parameters needed to construct and annotate the data for scientific analysis are contained in a Command Table; most of these parameters are not available from the spacecraft telemetry. Observation parameters needed to meet operations and science requirements are stored in an EDR table. Only unique characteristics of the Pathfinder instruments that affected the database design will be discussed in this article.

The complexity of the IMP instrument, with its ability to generate up to four images per observation, imposed the most burden on the design. Another constraint was that much of the IMP operation could be reconfigured by uplinking flight parameters. These capabilities were conveniently handled by adding tables to the database for observation and flight parameters. The Observation Table contained information that was common to all four image data sets, such as pointing information and exposure time. The Flight Table contained information that affected every observation from a certain point in time until modified, such as the defined image area returned.

The contents of the tables were customized for each instrument; they contained all of the information to identify scientifically useful data sets, and perform a moderate level of automatic validation. The database design centralized the updating of the respective metadata type to one location in the database. While connectivity amongst internal interfaces increased in complexity, this was hidden from the user and applications programmer by a low level database interface.

5. File Exchange Interface (FEI)

To manage the traffic of data files flowing between the numerous Pathfinder client sites and the MIPL database system, a client/server application called the File Exchange Interface (FEI) was administered. FEI had been previously developed in-house by MIPL and first used by the Galileo Flight Project. Both local and remote sites run the FEI Client program (FEI) to transfer data to and from their home institutions and MIPL.

FEI is integrated with a user authentication scheme called Kerberos, which was developed through Project Athena at MIT. Kerberos provides a means of securely authenticating users over a network using encrypted passwords and time-stamped network packets.

In comparison with FTP, FEI provides several advantages in its ability to push files to client sites. First, since FEI works with Kerberos to maintain its own individualized user access control, FEI users do not require accounts on local data center machines to receive data. This is required for FTP users. Second, the FEI user requires no knowledge of the physical location of data products because FEI stores and controls access to data by FEI fileType. The incorporation of FEI into a newly developed Web-based browser named the "MPF Navigator" (described in the "Planetary Atlas" section), which accepts a variety of criteria such as azimuth/elevation coordinates, etc., precludes the user from having to know the actual fileType. Third, FEI users can optionally submit subscriptions for fileTypes so that the pertinent data files are automatically pushed to client sites as soon as they are generated. A fileType subscription is submitted only once and runs continuously to instruct the FEI Server to monitor the addition of any new files and push them to the user's local disk immediately.

6. Planetary Atlas

Organizing, locating and distributing data in support of Mission Operations and Science Data Analysis is a fundamental requirement for Flight Projects. An existing capability, The Planetary Image Atlas, provides a mechanism to quickly and intuitively query distributed data systems to locate and acquire publicly released planetary data without requiring an intricate knowledge of database organization or the physical location of the data products. By utilizing existing World Wide Web (WWW) technology, the Atlas supports any platform, in any location around the world provided the user has a Java-enabled Web browser. Although some data systems can be very complex, the Atlas strives to make science data access and analysis easy for researchers and is based on the assertion that users should not need an intimate understanding of the data organization to begin interacting with the collection.

Mars Pathfinder was the first planetary mission to adopt the Atlas to support active mission operations, calling their adaptation the "MPF Navigator". The Atlas provided the core capabilities of the MPF Navigator while allowing the implementation of additional project-specific routines, database information and security restrictions that may vary from project to project. These core capabilities include a visual map-based Java and a text-based forms user interface. Using either of these interfaces (Figure 2) the user can: formulate search criteria, browse thumbnails or full resolution images and PDS labels, retrieve data in a wide variety of data formats, generate ad hoc database reports, and order large data sets.

Using the map-based Java interface a user may scribe out a sub-area on the mosaic with the mouse. This sub-area will be used as one of many possible search criteria for database queries. This interface is designed to be a stand-alone query mechanism or be an aid to the text-based query mechanism so the user can augment the search parameters further. The Mars Pathfinder adaptation opted to augment the map interface with the text interface rather than provide just the stand-alone map-based interface. The text-based interface allows users to click on buttons and/or input values using common browser forms to build their database query. Both mechanisms can incorporate project defined search criteria, providing a means to tailor the Atlas their unique needs.

After formulating and submitting their search, the user can select from a number of processing options. The processing options can be extended to accept project defined options. The MPF Navigator adaptation took advantage of this flexibility to incorporate the capability to generate FEI commands and to provide a direct link to the Navigation and Ancillary Information Facility's (NAIF) Experimenters Notebook.

Utilizing the Planetary Image Atlas design made it easy to equip the MPF Navigator with an interface to quickly and easily query the database based on a variety of search criteria and incorporate secure data distribution using the FEI. The MPF Navigator coupled with FEI made it possible to provide automatic data distribution and database access in virtually realtime to the project in locations around the world. As data is validated and ready for public release, a few parameters of the Atlas are reconfigured to open access to any allowed subset of the data to the electronically connected public.

7. Autonomous Mosaic Generation

The mosaic software used a camera model [Yakimovsky and Cunningham, 1978] which is based upon a set of vectors referred to as CAHVOR. These vectors permit a point in space to be traced into the image plane. They also permit the camera to be translated and rotated in order to mimic the actual IMP at Mars.

Vector values come from a camera calibration at one fixed azimuth and elevation. Images are obtained of a calibration target with holes at known locations. From a set of these images the CAHVOR vectors can be derived which map target holes into image plane locations. This process avoids the costly operation of determining optical cardinal points individually and can be performed in a few minutes.

Mosaics were assembled autonomously by tracing a view ray from each mosaic location or pixel into the scene, determining it's intersection with a ground plane, and then querying each input image to determine if that point lay within it's field of view. In this fashion mosaics containing several hundred images were assembled for each spectral band in about 3 minutes each. It was necessary to refine the camera pointing in order to produce accurate mosaics. This required the determination of the actual azimuth and elevation of each image in order to correct for gear backlash. To do this tiepoints were acquired between all pairs of overlapping images and then camera pointing commands were estimated which caused the camera model to map the tiepoints to their correct locations. The simplex method [Press et al., 1986] was used. In some cases this could be accomplished automatically but in general it required human intervention to select tiepoints because of nonexistent overlap, changing lighting or compression/decompression artifacts.

There were many applications for the image mosaics. First was the assembly of small pieces into a larger field of view. This included tilting the camera model in the Mars coordinate system to model a tilted spacecraft which resulted in mosaics with a level horizon beginning and ending at Mars north. The science teams used these products to orient themselves. Another application was to provide each day to the rover planning team a small stereo mosaic which was registered to a fixed reference image. This permitted the triangulation of way points for the next day's maneuvering. The mosaics which were used most often are described below.

Perspective Stereo Mosaics

These are perspective projections with horizontal epipolar lines. The mosaics behave as though the "camera" which took them was an IMP with a much larger field of view. The images are in the Lander reference frame and are tilted to reflect the position of the lander relative to the horizon.

Cylindrical Mosaics

Figure 3 shows such a mosaic with azimuth and elevation grid lines superimposed. In this case each pixel represents a fixed angle in azimuth and elevation. Rows are of constant elevation in Mars coordinates. The horizon is level, and columns begin clockwise from Mars north.

Polar Mosaics

Concentric circles represent constant projected elevation. Mars nadir is at the convergent center and the horizon is corrected for lander tilt. North is up.

McAuley Projection

This is a perspective projection similar to a Perspective Mosaic except that the mosaic acts like a pinhole camera which follows the mosaic in azimuth.

Vertical Projection

It assumes that the field is a plane tangent to the Martian surface with up pointing north. This is not an orthonormal rendering but was found to be useful for rapid initial orientation.

8. Color Processing

The most important color processing task was to generate color images both on NTSC, High Definition Television, and hardcopy of true color mosaics. This allowed presenting to the viewer what they would see if they were transported to the landing site. This required solving three problems:

1. The imagery must be converted into radiance units;
2. Radiance must be converted into CIE chromaticity coordinates;
3. Chromaticities must be converted back into those rgb values which cause the eye to perceive the correct color when viewing the image on a particular device.

To accomplish this registered mosaics were generated in each filter in units of radiance. A spectrum was then constructed for each pixel by fitting a spline through the spectral radiance points, and the X, Y, and Z tristimulus values were computed by integrating the spectrum with the three color matching functions [Wyszecki and Stiles, 1982]. Colorimetric units represent the color of the scene. Chromaticity coordinates were computed from tristimulus values and from there to Lab coordinates in preparation for conversion into the gamut of a display device.

Color tables were then computed giving the chromaticity coordinates for a large set of combinations of rgb values. These tables, one for each display device, permitted an interpolation to discover the rgb values which gave the desired chromaticity values for each pixel. When viewed under D65 illuminant prints gave the sensation of true color.

The last step was to synthesize a left eye green image since this filter was missing. The right eye was used to establish the relationship between RGB and then applied this to the left eye RB images to create a green mosaic.

9. Conclusion

The system adapted for Mars Pathfinder incorporated several innovations based on the demands imposed by the time constraints of an in-situ operations scenario and the demands for rapid release of highly processed instrument data to the public. New systems for automated mosaicking and color reconstruction based on pre-flight calibration data were evolved, and the ability to electronically distribute data based on a variety of criteria for different users was show-cased in an intensive operations environment. The systems utilized for the first time to support Mars Pathfinder will become the baseline for future planetary exploration missions.

10. Acknowledgments

The authors wish to thank the Mars Pathfinder project management team for encouraging innovation and creativity within a flight project environment with strict operational deadlines and firm requirements. We especially thank Al Sacks, John Wellman, Richard Cook, Tony Spear, Brian Muirhead, Matthew Golombek, and Rob Manning for encouraging us to bring new capabilities to this mission. The support of Bill Tuk, Kris Capraro, Curt Hahn, Dave Hodges, Chris Cordell, Mike Mueller, Larry Goforth, Maher Hanna, Greg Earle, Ray Bambery in helping MIPL and DIAL function effectively in support of flight missions is gratefully acknowledged. The efforts of the Multimission Data Management System development team, including John Rector and Jeff Jacobson, and the long hours put in by Carol Stanley as the realtime engineer are hereby acknowledged. It was a pleasure working with the instrument PIs, including Peter Smith, Tim Schofield, and Tom Economou on this project. We thank the rover team, and especially Brian Cooper and Jake Matijevic, for their encouragement and assistance. We thank Tom Handley for his continuing support of this activity. The support of the Telemetry and Missions Operations Directorate and the Mars Exploration Directorate at JPL is gratefully acknowledged.

The Planetary Image Atlas and the Planetary Photojournal were designed and developed by the Planetary Data System (PDS) in collaboration with the Solar System Visualization Project (SSV). The URLs are: http://www-pdsimage.jpl.nasa.gov/PDS/public/Atlas/Atlas.html and http://photojournal.jpl.nasa.gov.

High Definition Television technology support was provided by a prototype stereo HDTV system which was recently added to JPL's Visualization and Analysis Testbed (VAT). The VAT is part of the Digital Image Animation Laboratory (DIAL) and Multimission Image Processing Laboratory (MIPL) and is an element of JPL's end-to-end testbed. The VAT was developed by the Solar System Visualization (SSV) project and New Millennium Program (NMP).

Work described in this paper was carried out by the Mars Pathfinder Project at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

11. References

Press,W. et.al., *Numerical Recipes*,p.292 (amoeba),Cambridge University Press, 1986.

Yakimovsky,Y. and R. A. Cunningham, System for Extracting Three-Dimensional Measurements from a Stereo Pair of TV Cameras, *Computer Graphics and Image Processing* 7, 195- 210, 1978.

Wyszecki,G. and W. Stiles,*Color Science: Concepts and Methods, Quantitative Data and Formulae*, 2nd Edition, 156,John Wiley & Sons, 1982.

12. Figures

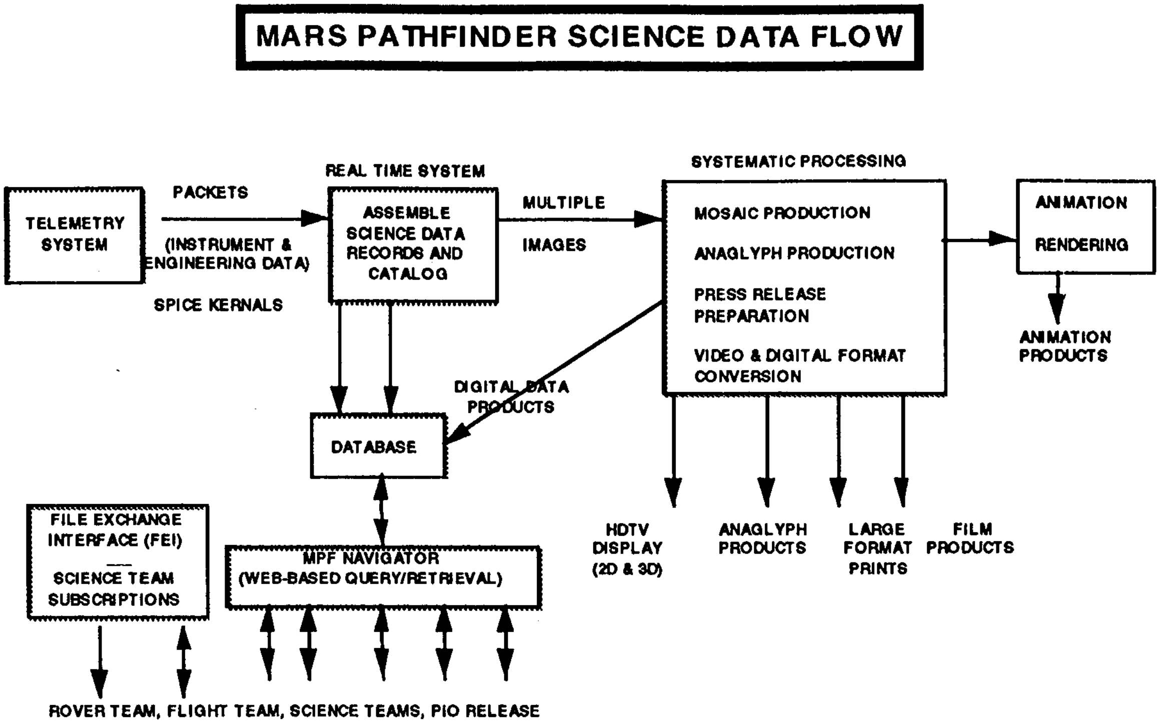

Figure 1: Mars Pathfinder Science Data Flow

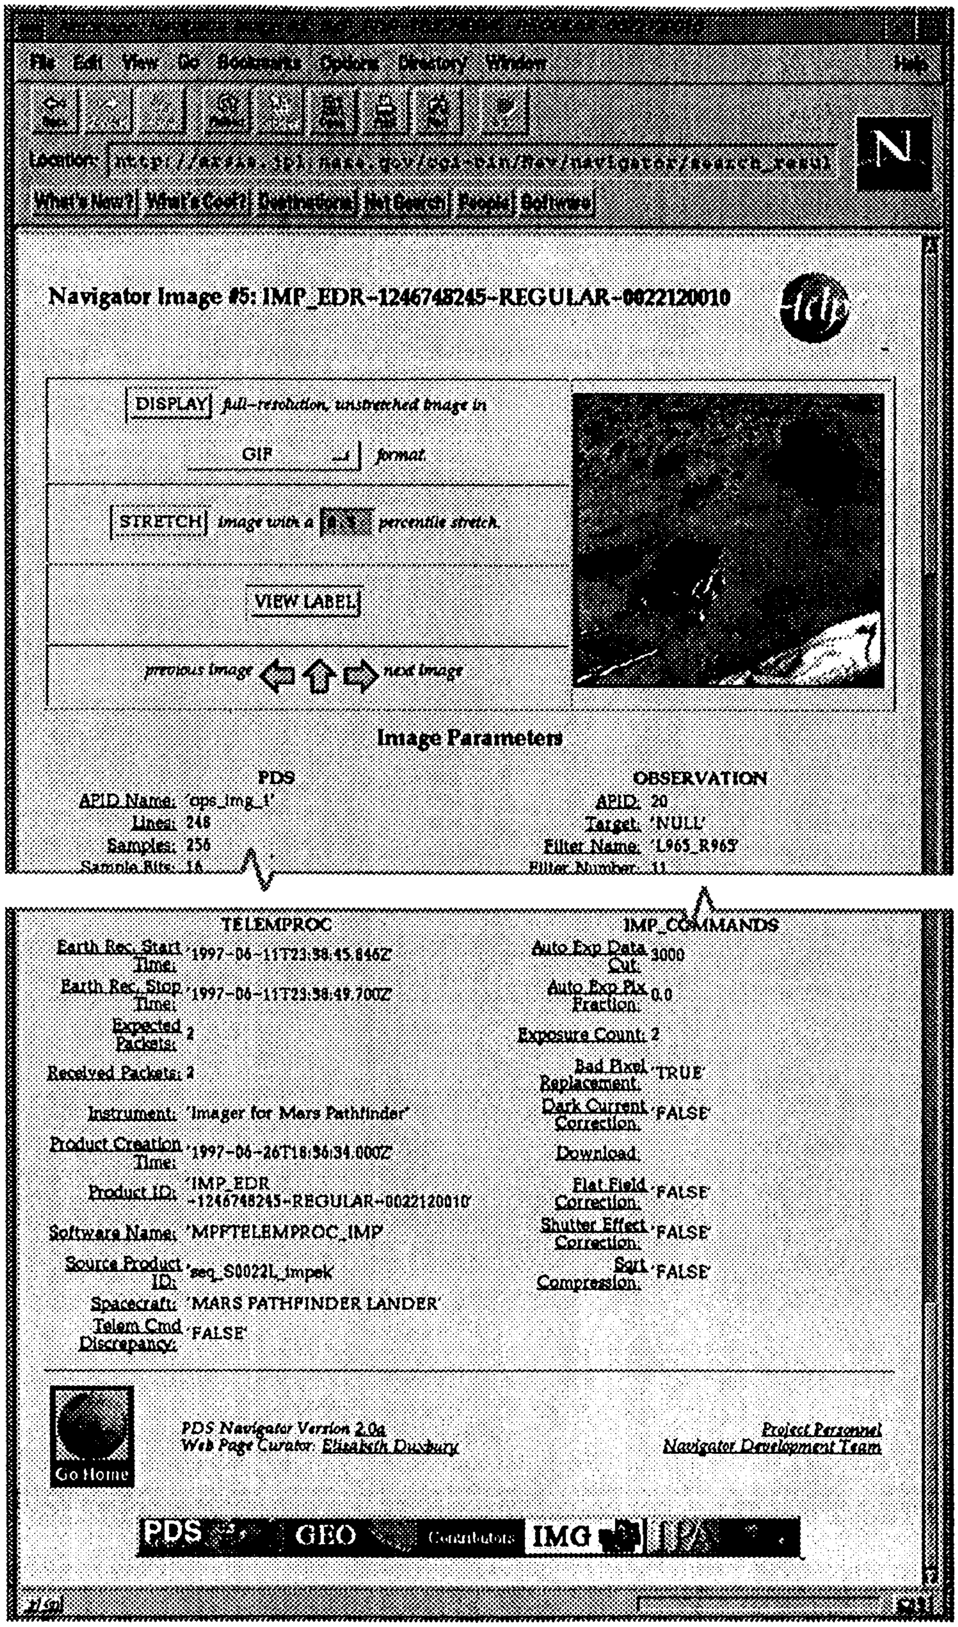

Figure 2. MPF Navigator - Image Browser page

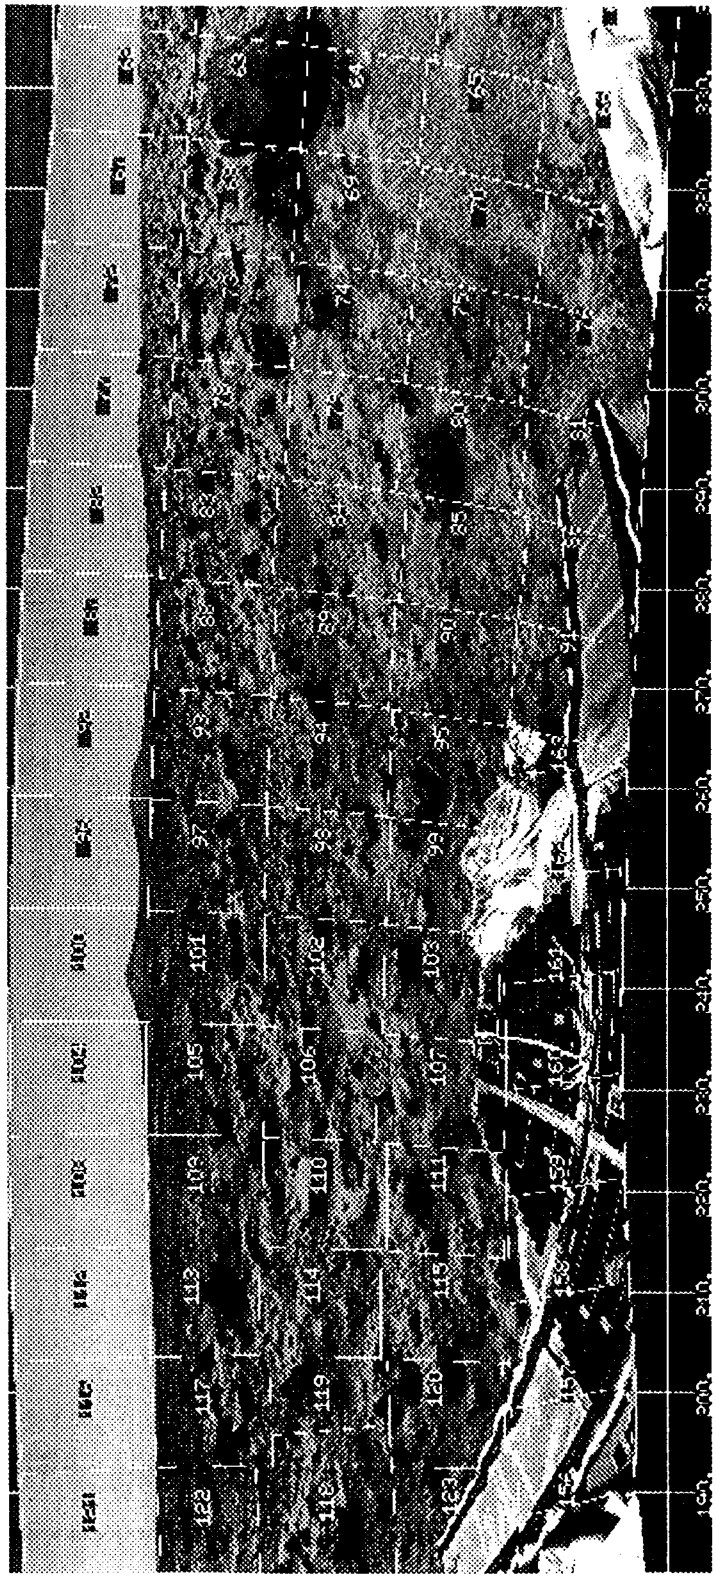

Figure 3. Cylindrical Mosaic