Terrain Modelling for Immersive Visualization for the Mars Exploration Rovers

J. Wright, F. Hartman, B. Cooper, S. Maxwell, J. Yen, J. Morrison
Jet Propulsion Laboratory(JPL)
California Institute of Technology
M/S 168-514 4800 Oak Grove Drive
Pasadena, CA USA 91109-8099

ABSTRACT

Immersive environments are being used to support mission operations at the Jet Propulsion Laboratory. This technology contributed to the Mars Pathfinder Mission in planning sorties for the Sojourner rover and is being used for the Mars Exploration Rover (MER) missions. The stereo imagery captured by the rovers is used to create 3D terrain models, which can be viewed from any angle, to provide a powerful and information rich immersive visualization experience. These technologies contributed heavily to both the mission success and the phenomenal level of public outreach achieved by Mars Pathfinder and MER. This paper will review the utilization of terrain modelling for immersive environments in support of MER.

Keywords: Immersive environments, MER, terrain modelling, visualization, Mars, JPL.

Current developments in immersive environments for mission planning include several tools utilizing terrain visualization, which make up a system for performing and rehearsing missions. This system, known as the Rover Sequencing and Visualization Program (RSVP)[1][2], includes tools for planning long range sorties for highly autonomous rovers, tools for planning operations with robotic arms, and advanced tools for visualizing telemetry from remote spacecraft and landers. In addition, a Web-based tool, known as the Science Activity Planner (SAP), allows for collaboration by remote scientists in designating features of interest in a similarly immersive environment, using similar terrain models.

As rovers can range over greater distances with more autonomy when they have accurate self-locating systems, the operator paradigm shifts from a hands-on micromanagement level to a hands-off level of mission specification. This calls for a more immersive interaction with the environment with tools for designating waypoints, samples to be collected, regions of hazard and interest, and other types of features. This type of environment is applicable to both rover navigation and operations with robotic arms and sensors. This type of immersive environment is critical for maximizing operator understanding of the environment for reducing risk and optimizing science return[3][4]. A variety of immersive environments are possible and each carries a particular part of the information that needs to be conveyed to the operator [5][6].

The primary goal of the immersive mission planning tools in RSVP is to provide the mission planners with the best possible understanding of the region. This requires building three-dimensional terrain models from the stereo imagery captured by the three primary camera systems on the rovers, the PanCam, the NavCam, and the front and rear HazCams. Processing of the data from these different camera systems produces three-dimensional models of different resolution and spectral characteristics. Some models may be used independently but it is necessary to combine the models into a single, coregistered dataset to support free-roaming visualization and exploration of the mission area. A further requirement is to support multiple resolutions in the final model products for increased rendering performance, as well as producing models in a variety of formats to be utilized by different applications. Thus, there are requirements to produce individual terrain "wedges" from stereo pairs, register the wedges to each other, merge the wedges into a coherent model, and then produce terrain models tailored for different immersive visualization applications.

The process descriptions in this paper outline the production of terrain models to support immersing the operator in the environment of another planet, body, or space to make the mission planning function more intuitive and effective.

The System for Unifying Multiresolution Models and Integrating Three-dimensional Terrains (SUMMITT) task has the goal of developing the underlying modelling technology for supporting missions involving rovers. Three-dimensional models of terrain areas are an invaluable asset in planning operations and in reviewing the predicted and telemetered operations of a robot arm. The SUMMITT task had the initial goal of supporting the Mars Volatiles and Climate Surveyor (MVACS) team during the Mars '98 mission, which unfortunately failed. The next Mars surface operations missions were the Mars Exploration Rover missions that landed two rovers. These missions expected to have orbital imagery from Mars Global Surveyor (MGS) and Mars '01, descent imagery from the lander, and lander imagery from four sets of stereo imagers. These sources of imagery were to be combined to create a multiresolution terrain model with very high resolution detail available within the immediate area of operations of the rover. Data constraints, including the limited amount of descent imagery and the wide range of exploration of the rovers, has limited the use of the modelling technology to the data collected by the rovers themselves. This has had spectacular success.

The three primary sources of data are the NavCams, the HazCams, and the PanCams. Each type of imagery is partially processed independently, then combined with data derived from the other imagery to create the multiresolution terrain models.

A fundamental problem is the registration of the different terrain model pieces generated independently from the different data sources. The method employed at JPL uses volumetric primitives (voxels) to represent the terrain to be matched. Voxels have some advantages over the polygonal surface matching methods in that it is easier to represent unknown volumes, such as regions occluded by rocks and hills, and easy to use, multiresolution data structures are available in which to combine the models once they are matched. The registration of the voxel datasets uses an iterative closest points method based on the work of [7]. The entire voxel model is stored in an octree structure which supports multiresolution data and rapid access while utilizing significantly less memory than a three-dimensional grid. For the MER missions, only lander imagery was available which greatly simplifies the registration process. While not trivial, each of the terrain model pieces was created from calibrated and relatively colocated instruments making registration straightforward.

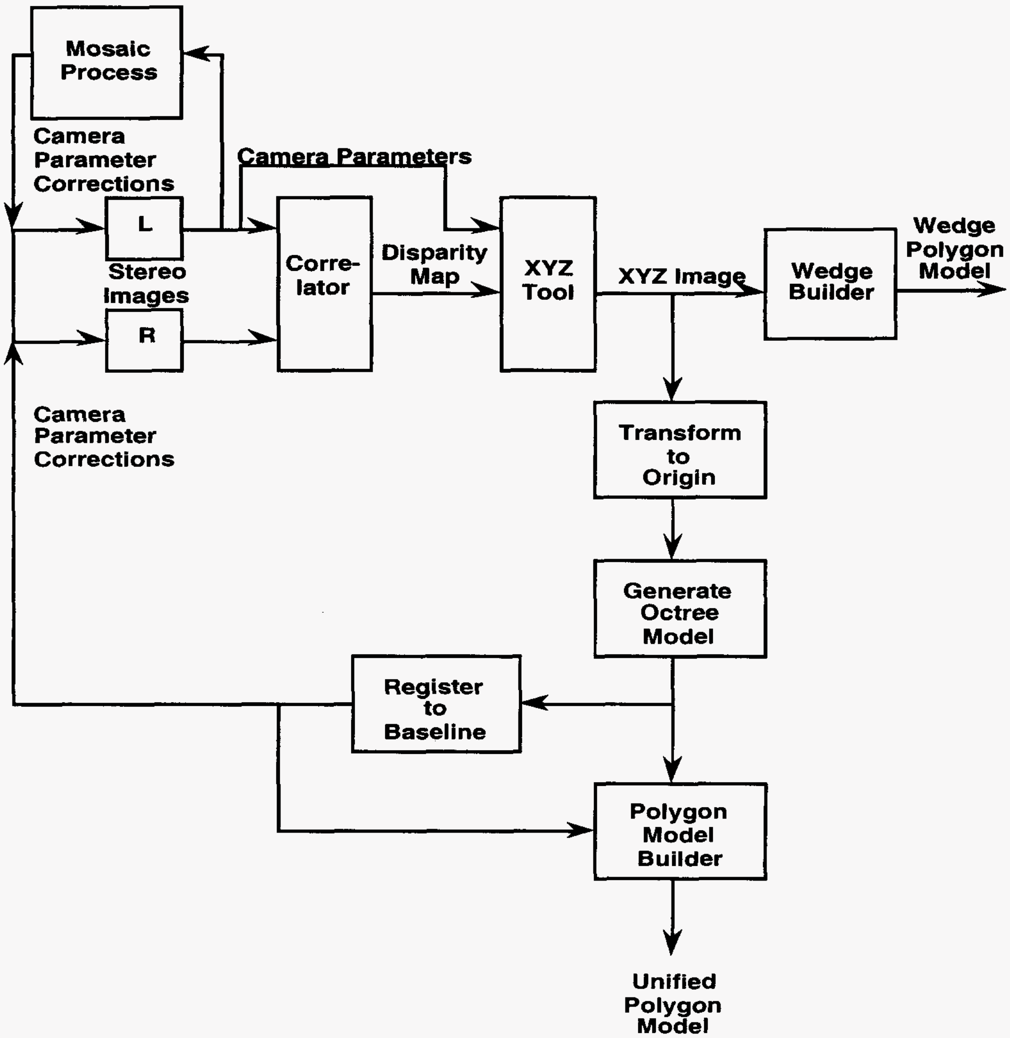

Figure 1 – Terrain Modelling Process Flow

Figure 1 illustrates the processing performed on the stereo pairs from the PanCams, NavCams, and HazCams. The PanCams and NavCams are mounted on the mast, about 1.5m above the terrain. The NavCams are mounted below the rover deck, about 0.5m above the terrain. Each has a left and right imager generating stereo pairs for evaluating range information and producing three-dimensional models. An example of a stereo pair is in Figure 2. Each stereo pair is first processed with a correlator that produces a disparity map that identifies matching features in each image. The camera model is then used to compute the range to each pixel in one image using the disparity to its matching pixel in the other image. Then the camera pointing information is used to project the pixels

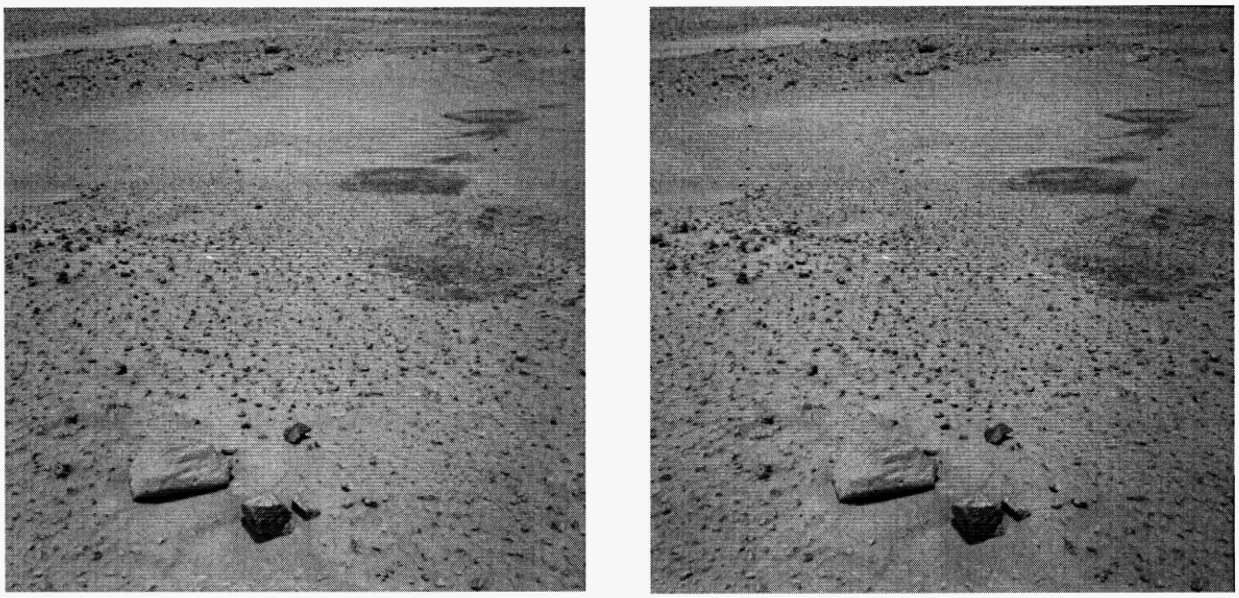

Figure 2 – Stereo Image Pair (view cross-eyed for 3D)

to an (x,y,z) location in three-dimensional space. These XYZ images contain three bands with the x coordinates of each point sample stored in one band, the y coordinates in a second band, and the z coordinates in a third band.

The XYZ images can be thought of as a cloud of point samples generated from a single image pair. These point samples can then be converted to a simple polygon model by merely connecting adjacent points as determined by adjacency within the XYZ image. Such a process creates a wedge of terrain model with the narrow end of the wedge pointing toward the rover. The SAP tool utilizes a set of wedges as a model and manages and renders the individual wedges appropriately for visualization by the scientists. However, for higher performance rendering, a unified terrain model, with multiresolution characteristics, is desired. The terrain models are registered and merged into an octree. The registration process is very simplistic in that the camera pointing information is assumed to be correct and thus the (x,y,z) coordinates of the samples is accurate. During Mars Pathfinder, such was not the case. However, for MER the pointing is much more accurate and the modelling process is helped greatly.

The point sample nature of the original data is modified to fit the octree concept which is volumetric in nature. Each point sample is given a volume determined by the pixel field of view of the associated instrument and the range to the sample. This volume is associated with the sample and used to locate the appropriate level of the octree. As data from multiple XYZ images is added, the octree absorbs them and results in a single model of the terrain in the region.

Due to the ubiquity and performance of polygon rendering systems, both software and hardware, a polygon model format was desired. The primary requirements of the polygon model are high fidelity in the terrain imagery with high rendering performance. Because the operator will be making planning decisions that require detailed local knowledge combined with general understanding of more distant terrain, the polygon models must be multiresolution also, or at least reflect the multiresolution nature of the underlying data samples. The use of rendering algorithms and tools in the original point sample space has been considered using such tools as the Volumizer from Silicon Graphics or Mitsubishi Electric's VolumePro real-time hardware. Unfortunately, these systems are optimized for medical data, the primary volumetric data source today. Thus, the datasets are three-dimensional grids of limited size (typically 512x512x512). The datasets expected on the Mars missions will be equivalent to 512^3 for each dataset generated from a single stereo pair. The baseline terrain may be 1024x1024 with tens or hundreds of higher resolution inserts. To avoid the memory usage required by a full 3D grid, an octree was utilized for intermediate storage of the point samples. Current tools do not work with specialized data structures such as an octree. Future work in exploring the use of such techniques as splatting [8] may prove effective.

The multiresolution nature of the point samples precludes the use of a simple algorithm such as Marching Cubes [9] for converting to a polygon surface. Other, more sophisticated methods such as Marching Triangles [10] also do not work well in this environment. To extract the polygon model from the octree, a unique combination of image space sampling and octree space sampling is used. It is desired to produce a tiled, multi-resolution polygon model, yet it is difficult to extract a valid surface from the merged point cloud within the octree without dramatically reducing the resolution and removing features that could be critical to safe operations. To accomplish the polygon model extraction, the connectivity of the samples is extracted from the original XYZ images. Each XYZ image produces one or more independent polygon mesh sections. The images are downsampled to produce lower resolution versions of each section. To achieve tiling, however, the octree is used to divide the points within each XYZ image into separate square regions in X and Y and, structurally, each section of an XYZ image, at all resolutions, is associated with a tile for that region. To avoid artifacts along the edges of the tiles, polygon edges that cross tile boundaries have both endpoints in both tiles, thus duplicating the polygons along the edges. While not optimal, the increased polygon count is minimal and does not affect rendering performance.

Each type of imager on the rovers is used to produce terrain models for a specific purpose. The HazCams are used to produce models for planning and rehearsing operations of the robot arm, the Instrument Deployment Device or IDD. Because the HazCams are on a fixed mount below the deck, they offer the best view of the IDD work volume and are not subject to any pointing errors that might result in the misplacement of an IDD target. Figures 3 and 4 are views of a HazCam model being used to rehearse IDD operations. The NavCam imagery is used primarily for producing terrain models for planning traverses. The NavCams have about a 45° field of view so several of these are

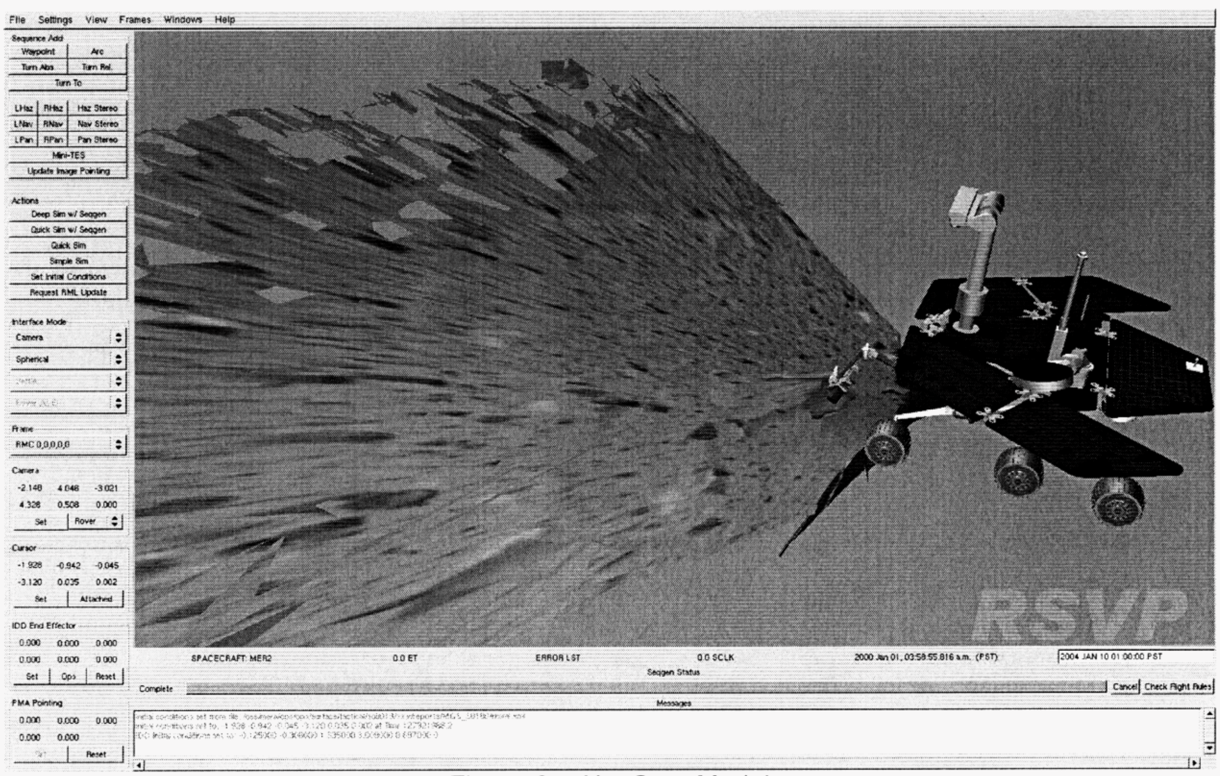

Figure 3 – HazCam Model

combined to model the terrain around the rover. An entire panorama of NavCam models constitutes a site within which several days worth of investigations may be performed. Figure 5 shows a portion of the complete panorama model from the site produced right after driving off the lander. The PanCams are typically used for science observations but are also used to produce terrain models at a distance along the expected direction of the next traverse. These higher-resolution images with a

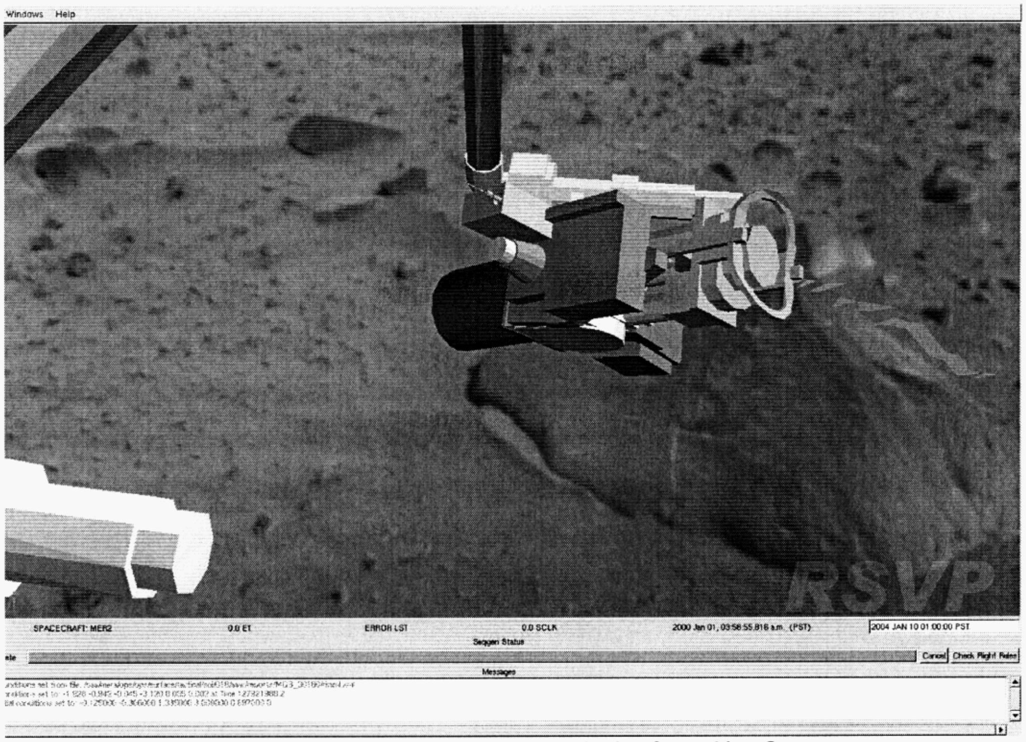

Figure 4 – IDD Work Volume Model from HazCam

field of view of about 16° produce good quality meshes out to a greater distance providing for longer traverses being planned safely. Figure 6 shows a merged dataset with both PanCam and NavCam terrain data. Note that where the NavCam data begins to lose quality, the PanCam data is still very good and useful for planning traverses.

The standard paradigm for MER operations is to drive some distance, declare a new site, capture a panorama of imagery, build terrain models, and then spend some number of days exploring the site. All coordinate information is specified in site coordinates, the origin of which is declared to be the rover's location when the site was declared. As the rover wanders, new sites are being declared and panoramas captured. By Sol 68, 18 sites had been declared for the Spirit rover. Figure 7 shows how the sites are declared. Figure 8 shows a view of the traverses performed up through site 17 for the Spirit rover.

The Jet Propulsion Lab is utilizing enhanced immersive technologies for supporting MER mission operations. Immersive technologies and systems are aiding the operations teams in making mission critical decisions. Creating models of the operational environment and providing visualization tools to explore and interact with that environment are the key aspects of the MER mission tool suite. Complete systems which take all the data available and use it to create an immersive environment, enable scientists to interact with the environment and establish science mission goals, collate science requests into mission operations sequences for uplink, rehearse expected activities within the

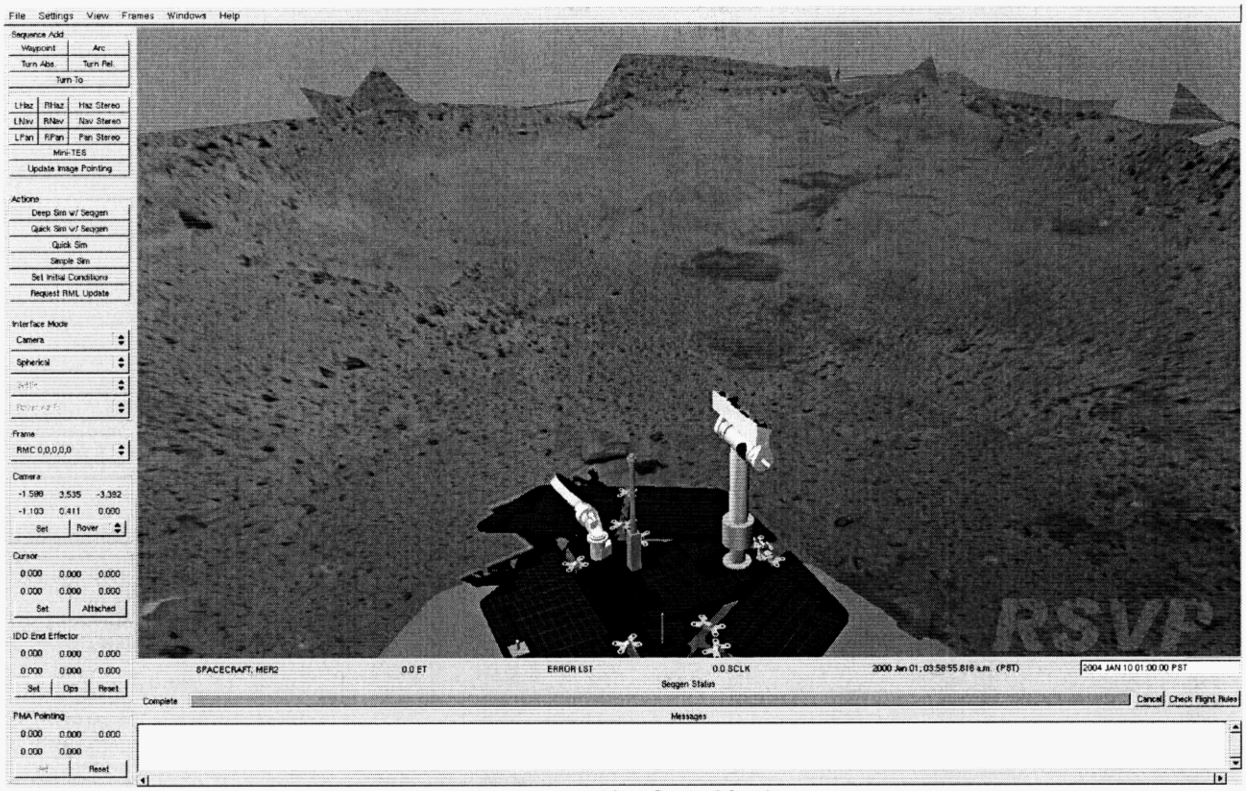

Figure 5 – NavCam Mesh

environment, replay actual mission activities based on telemetry, and compare actual operations to predicted operations are what have made MER so successful.

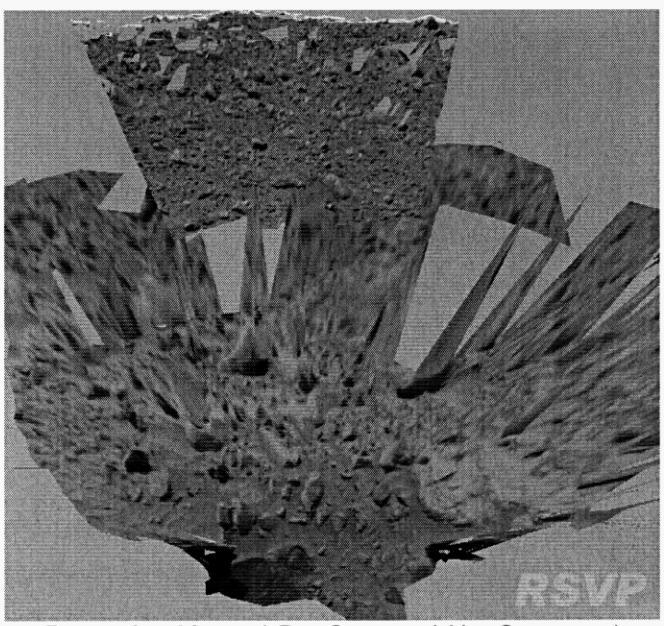

Figure 6 – Merged PanCam and NavCam mesh

Immersive technologies will continue to offer more capability to mission operations teams in the foreseeable future. This will be especially important as exploratory craft gain in autonomy and intelligence and begin to explore farther afield from their initial landing site. Operations will change from precisely specifying number of revolutions of each wheel to broadly defining goals and strategies. Such paradigms require a broader understanding of the environment to facilitate rapid decision making. While the level of immersion can vary from simple two-dimensional displays through stereo displays and virtual workbenches and on up to full immersion with head-mounted displays, haptic and tactile feedback [11], and intuitive manipulators [12], the main goal is to provide the operators with a better understanding of the operational environment.

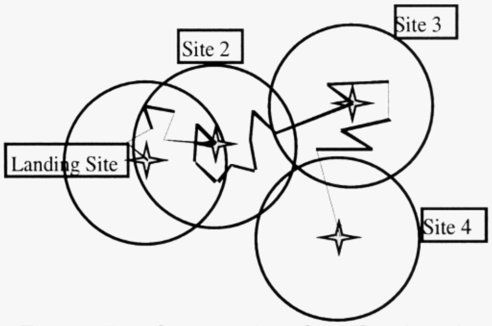

Figure 7 – Successive Site Declarations

Figure 8 – Traverse History

The work described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

[1] Wright, J., Hartman, F., and Cooper, B. (1998) Immersive Environments for Mission Operations: Beyond Mars Pathfinder, Proceedings of SpaceOps '98, Tokyo, Japan, 1998.

[2] Cooper, B. (1998) Driving on the Surface of Mars Using the Rover Control Workstation, Proceedings of SpaceOps '98, Tokyo, Japan, June, 1998.

[3] Ellis S.R. (1994). What are Virtual Environments? IEEE Computer Graphics and Applications. 17-22.

[4] McKenna, M. & Zeltzer, D. (1992). Three Dimensional Visual Display Systems For Virtual Environments. Presence: Teleoperators and Virtual Environments. 1(4). Cambridge, MA: MIT Press. 421-458.

[5] Stuart, R. (1996). The Design of Virtual Environments. New York McGraw Hill.

[6] Volbracht, S. (1996). An Experimental Comparison of 3-D Display Modes. IEEE Visualization '96. Hot Topics. Piscataway, NJ: IEEE. 8-11.

[7] Champleboux, G., Lavallee, S., Szeliski, R., and Brunie, L., (1992) From Accurate Range Imaging Sensor Calibration to Accurate Model-Based 3-D Object Localization, Proceedings of IEEE CVPR '92, pp. 83-89, 1992.

[8] Swan, J.E II, Mueller, K., Moller, T., Shareef, N, Crawfis, R., and Yagel, R. (1997). An Anti-aliasing Technique for Splatting. IEEE Visualization '97. Phoenix, AZ: IEEE. 197-204.

[9] Lorensen, W. and Cline, H. (1987) Marching Cubes: A High Resolution 3D Surface Construction Algorithm. Proceedings of SIGGRAPH '87, pp. 163-169, July 1987.

[10] Hilton, A., Stoddart, A.J., Illingworth, J., and Windeatt, T. (1996) Marching Triangles: Range Image Fusion for Complex Object Modelling. Proceedings of International Conference on Image Processing, 1996.

[11] Hughes, R.G. & Forrest, A.R. (1996). Perceptualisation Using a Tactile Mouse. IEEE Visualization '96 Proceedings Piscataway, NJ: IEEE. 181-188.

[12] Shimoga, K. B. (1993). A Survey of Perceptual Feedback Issues in Dexterous Telemanipulation. VRAIS '93: 1993 IEEE Virtual Reality Annual International Symposium. Piscataway, NJ: IEEE. 263-279.

John Wright is one of the developers of RSVP at the Jet Propulsion Lab and is currently one of the drivers for the MER rovers. He has been at JPL for ten years working on a variety of visualization and computer graphics related tasks. He holds a B.S. in Chemical Engineering from Purdue University and an M.S. in Computer Science from Wright State University.

Frank Hartman is one of the developers of RSVP at the Jet Propulsion Lab and is currently one of the drivers for the MER rovers. He has been at JPL for ten years working on a variety of visualization and computer graphics related tasks. Mr. Hartman developed tools for creating the three-dimensional models of Mars used during the Pathfinder mission for operation of the Sojourner rover. He holds a B.F.A in Sculpture from the Philadelphia College of Art and an M.S. in Aeronautical and Astronautical Engineering from Stanford University.

Brian Cooper is the lead developer of RSVP at the Jet Propulsion Lab and is currently one of the drivers for the MER rovers. He has been at JPL for 19 years and has worked on a variety of robotic vehicles tasks. He was the principal developer of the RCW used to control the Sojourner rover on Mars during the Pathfinder mission as well as the principal operator of the rover. Mr. Cooper holds a B.S. in Electrical Engineering from the University of California, Irvine, and an M.S. in Computer Engineering from the University of Southern California.

Scott Maxwell is one of the developers of RSVP at the Jet Propulsion Lab and is currently one of the drivers for the MER rovers. He is the author of the book Linux Core Kernel Commentary, whose second edition was published by Coriolis Press in 2001. Mr. Maxwell has worked for JPL since graduating from the University of Illinois at Urbana-Champaign in 1994.

Jeng Yen is one of the developers of RSVP at the Jet Propulsion Lab and is currently one of the downlink telemetry analysts for the MER rovers. Dr. Yen received his PhD in Applied Mathematics with minor in Mechanical Engineering from The University of Iowa at 1990.

Jack Morrison, a contractor for Northrop Grumman working at the Jet Propulsion Laboratory, designed the software on-board the Sojourner Mars rover, and was one of the primary "rover drivers" during Pathfinder surface operations. He is the key software developer for the SUMMITT tools. Mr. Morrison holds a BS degree in Math from the University of California Los Angeles.