SAE 25th International Conference on Environmental Systems, San Diego, CA, July 1995 (pre-print)

Operator Interfaces and Network-Based Participation for Dante II

Terrence Fong, Henning Pangels, David Wettergreen
The Robotics Institute, Carnegie Mellon University

Erik Nygren
Massachusetts Institute of Technology

Butler Hine, Phil Hontalas
Intelligent Mechanisms Group, NASA Ames Research Center

Christopher Fedor
K^{2}T, Inc.

ABSTRACT

Dante II, an eight-legged walking robot developed by the Dante project, explored the active volcanic crater of Mount Spurr in July 1994. In this paper, we describe the operator interfaces and the network-based participation methods used during the Dante II mission. Both virtual environment and multi-modal operator interfaces provided mission support for supervised control of Dante II. Network-based participation methods including message communications, satellite transmission, and a World-WideWeb server enabled remote science and public interaction. We believe that these human-machine interfaces represent a significant advance in robotic technologies for exploration.

INTRODUCTION

The Dante Project develops semi-autonomous walking robots to advance robotic technologies for planetary exploration[1]. Our objective is to demonstrate the capability of this technology by deploying terrestrial analogs. Robotic technology is critical for future space missions that will require remote science in distant, hazardous and unstructured environments. In particular, we are seeking to unify humans and robots to maximize the advantages and minimize the limitations of each[2].

In July 1994, the Dante Project deployed Dante II into the active crater of Mount Spurr (Aleutian Range, Alaska). Shown in Figure 1, Dante II is an eight-legged, frame-walking robot that uses an anchored tether cable to descend in a rappelling-like manner. It was created to allow volcanologists to explore active volcanic craters from safe, remote locations[3]. We used Dante II to explore Mount Spurr by placing the robot and support equipment at the crater rim. Mission operators and volcanologists remotely supervised the descent from a control station in Anchorage, Alaska. An overview of the Dante II / Mount Spurr expedition is shown in Figure 2.

For Dante II, we needed to create a competent robot capable of continuous, reliable and self-reliant exploration. To do so, we had to satisfy operator needs for controlling a robot which exhibits localized autonomy as well as observer needs for remotely conducting field science. Thus, we endeavored to create appropriate operator interfaces and to use network-based participation methods which would result in synergistic human-machine interaction.

We designed Dante II to be operated in control modes ranging from teleoperation of individual actuators to autonomous

Figure 1: Dante II entering Mount Spurr (Aleutian Range, Alaska). The tether can be seen exiting at lower right. Cameras and a scanning laser ranger finder are mounted on the arched mast. (photo: NASA)

walking. This range of human-machine interaction is characteristic of *supervisory control*[4]. Our reasons for using supervisory control are motivated by factors relevant to planetary exploration robots: to make operations faster and unconstrained by the human sensorimotor system, to reduce operator workload and improve performance, and to compensate for limited communication bandwidth and transport delays between the operator and the remote system.

In the strictest sense, supervisory control means once the operator relinquishes control, the robot is allowed to function

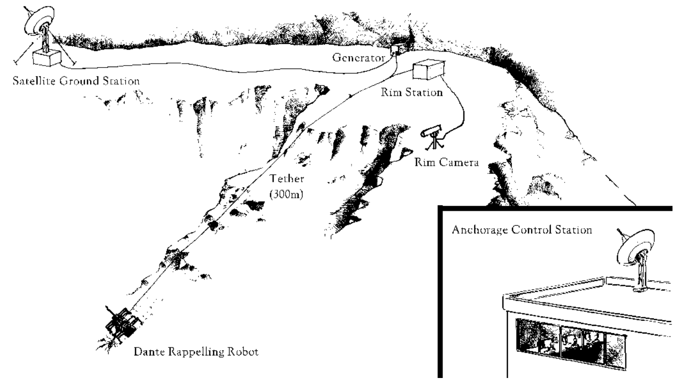

Figure 2: Overview of the Dante II / Mount Spurr expedition. Dante II and support equipment were deployed on the active crater of Mount Spurr. Mission operators and volcanologists supervised the descent from a control station located in Anchorage, Alaska.

completely autonomously. Our intent, however, was to be less restrictive and to let operators intermittently assume direct operation (*traded control*) or to supervise certain variables while directly handling others (*shared control*)[4]. We believe this approach affords greater human-machine synergy, benefiting and improving overall system performance.

To support traded and shared control of Dante II, we recognized that operators need interfaces to visualize and understand system state, as well as to specify and use appropriate control modes. Due to the complexity of Dante II's configuration and the risks associated with exploration, we saw a need for concise and efficient human-machine interaction.

The Dante II / Mount Spurr expedition was intended not only to advance robotic exploration technology, but to also support numerous remote mission observers and participants. We were interested in conducting collaborative research with remotely located volcanologists and planetary scientists. Thus, we needed to provide observer interfaces which would enable field science and which would allow distant researchers to directly participate in the mission. Additionally, we wanted to provide public access prior to, during, and after the Mount Spurr expedition. By offering compelling observer interfaces, we hoped to generate awareness, to broadly educate and to inspire others.

In this paper, we describe the operator interfaces and network-based participation methods used during the Dante II / Mount Spurr expedition. The development of these interfaces and methods was guided by two fundamental design principles: *ease of understanding* and *ease of use*. In the following sections, we show how these principles enabled us to create human-machine interfaces that represent a significant advance in robotic technologies for planetary exploration.

OPERATOR INTERFACES

A. OVERVIEW

This section presents the operator interfaces developed for the Dante II / Mount Spurr expedition. We first describe the *Virtual Environment Vehicle Interface* (VEVI) which provided real-time system and sensor visualization. We then present the *User Interface 2D* (UI2D) multi-modal control interface which supported traded and shared control operations. Finally, we describe the on-board and external video systems used by mission operators and observers.

B. VEVI

1. Background

The NASA Ames Intelligent Mechanisms Group (IMG) has been developing virtual environment interfaces since 1991. Virtual environments enable the efficient presentation, manipulation and visualization of complex data through immersive and spatially-oriented displays. They can both enhance an operator's situational awareness and help compensate for sub-optimal communication constraints, such as low bandwidth and transport delays, between a control station and remotely operated system [5].

VEVI is a modular operator interface for robotic vehicles [6][7][8]. It is a tool for planning and previewing high-level task sequences, monitoring system state and analyzing anomalous events. Through virtual environment technology and techniques, we can provide operators with spatial orientation and perspective superior to conventional control stations. Additionally, with network communication we are able to support distributed collaboration through shared virtual environments.

Our motivation for developing VEVI was to provide a consistent operator interface for supervisory control of a variety of mobile robot vehicles. To date, we have used VEVI to control and monitor numerous systems including wheeled, air-bearing, legged, and free-flying underwater vehicles [5][7][8][9][10].

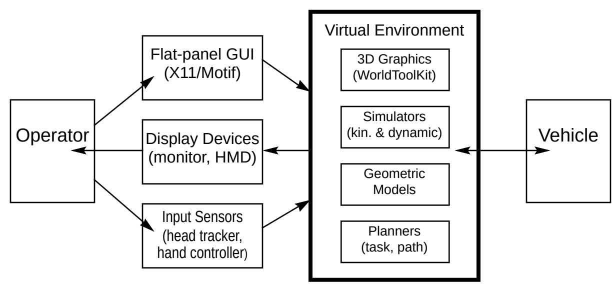

Figure 3: Virtual Environment Vehicle Interface (VEVI) structure. An operator interacts with a virtual environment through input sensors and display devices. The virtual environment may be simulated, a mirror of the real-world, or a combination of the two.

The overall structure of VEVI is shown in Figure 3. VEVI utilizes *real-time*, interactive 3D graphics and position-orientation sensors to support a range of interface modes from flat panel screen (windowed or stereoscopic) to head-mounted and head-tracked displays. VEVI executes as a loosely-synchronous process on Silicon Graphics workstations, rendering a scene containing graphical models of the controlled vehicle and its environment. Feedback from on-board sensors is used to update the simulated vehicle and environment. Operators interacting with VEVI have control over interface and display parameters including field-of-view and viewpoint. High-level task sequences can be created within VEVI using simulators and planners, then transmitted to the remote vehicle for execution[6]. We did not, however, use such task sequences for Dante II because control of walking has not reached the necessary maturity or simplicity.

2. Mission Needs

Dante II is a rappelling frame-walking robot with eight pantograph legs arranged in two groups of four on inner and outer frames. Each leg is independently actuated and can be positioned vertically to accommodate various step heights. Stroke and turn actuators allow its two frames to be translated and oriented relative to each other. To operate on steep slopes, a tensioned tether system provides reaction to gravity and minimizes shearing forces at the feet[3][11].

Although the motion of each actuator is well defined, the complexity of Dante II's configuration makes it difficult for people to visualize; even those who are familiar with the mechanism can have trouble. When Dante II is operating in unstructured and unknown environments, terrain features including steep embankments and obstacle strewn fields further compound the difficulty of visualizing system state.

Consequently, we chose to apply VEVI in a manner that enables operators and observers to readily comprehend Dante II's configuration and its surroundings. The challenge, however, is not merely to create an accurate kinematic representation, but to synthesize a unified interface in which visualization of Dante II and its sensory data appear concisely, clearly and *intuitively comprehensible*.

To achieve this, we defined requirements in three areas as guidelines for development. These were: operational metrics, configuration display, and terrain display. Operational metrics are those interface characteristics which are directly visible or directly impact the user such as frame rate, responsiveness, and interactivity. Configuration display includes the system configuration representation including model update rate, graphical detail, and color usage. Terrain display contains the features governing the depiction of perception sensor derived data including terrain representation, resolution, and update method. The specific requirements are summarized in Table 1.

Table 1: Requirements governing VEVI for Dante II

Operational Metrics	• 10 Hz minimum frame rate • high interactivity • support for multiple display stations • nominal performance on mid-range Silicon Graphics workstation (e.g., Indigo^2 Extreme) • unified kinematics & terrain display
Configuration Display	• visualize system configuration • 10-20 Hz update rate from robot's internal sensor data • low level-of-detail polygonal model • accurate kinematic representation • display sensed leg and tether forces
Terrain Display	• visualize external terrain environment • shaded polygonal grid representation of height field with coded (colored) features • 0.1 to 1.0 m grid resolution • up to 20m square grid • support patch based terrain mosaics • reflect haptic terrain perception

3. Design & Implementation

In modifying VEVI for Dante II, we learned the importance of accurately eliciting data flow parameters (data types, bandwidth) and graphics resource limits (rendering rate, polygon size). These constraints are significant since they directly constrain the design and impact the performance of virtual environment interfaces.

Our data flow analysis, for example, revealed that a significant portion of the Dante II telemetry is periodic and repetitive. This data type (e.g., state information) requires that the latest, most recently transmitted data be processed as rapidly as possible and that older, unprocessed data be discarded. Similarly, we determined that other data types (e.g., elevation maps) need to be handled in order and cannot be discarded. Having this knowledge was essential for partitioning processing functions and tasks.

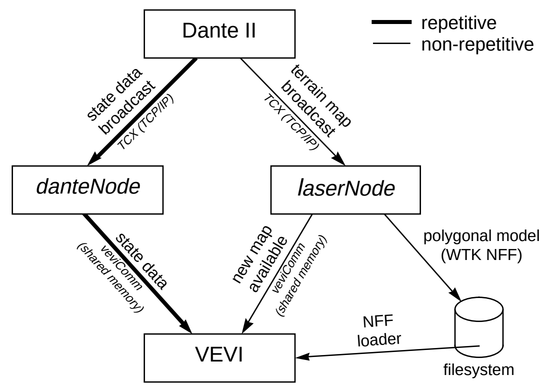

Figure 4: Dante II to VEVI data flow. Independent modules provide explicit processing of repetitive and non-repetitive types.

Figure 4 shows the data flow between Dante II and VEVI. The portion of the Dante II telemetry containing state data is received by the *danteNode* module. This module performs calculation of forward kinematics, matching of state variables to graphical objects, coloring of select model components based on sensed forces, and creation of quadrilateral polygons to reflect haptic terrain perception (sensing of ground location via correlation of leg position and force). The *danteNode* module relays the processed data to VEVI via a shared-memory data channel.

The remaining telemetry (elevation map data generated by the on-board scanning laser-rangefinder system) is received by the *laserNode* module. This module processes the regularly gridded, height field data into a form suitable for rendering: the elevation map is converted from height fields into a shaded tri-mesh representation and coded features (e.g., unknown regions) are colorized. At this point, the *laserNode* module stores the processed output then signals VEVI to retrieve and render the new data.

A representative display of Dante II in VEVI is shown in Figure 5. In this view, the graphical model of Dante II is oriented similarly to the actual vehicle shown in Figure 1. Several features of note are: (1) low level of detail Dante II model (approximately 1,000 tri-polygons), (2) coloring of vertical leg links and tether to indicate sensed force, (3) quadrilateral polygons indicating ground contact position sensed by haptic perception, and (4) gravity normal horizontal reference grid.

Figure 6 depicts VEVI with terrain maps added to the display. The left figure shows the visualization of elevation map data obtained from a single scanning laser range-finder elevation

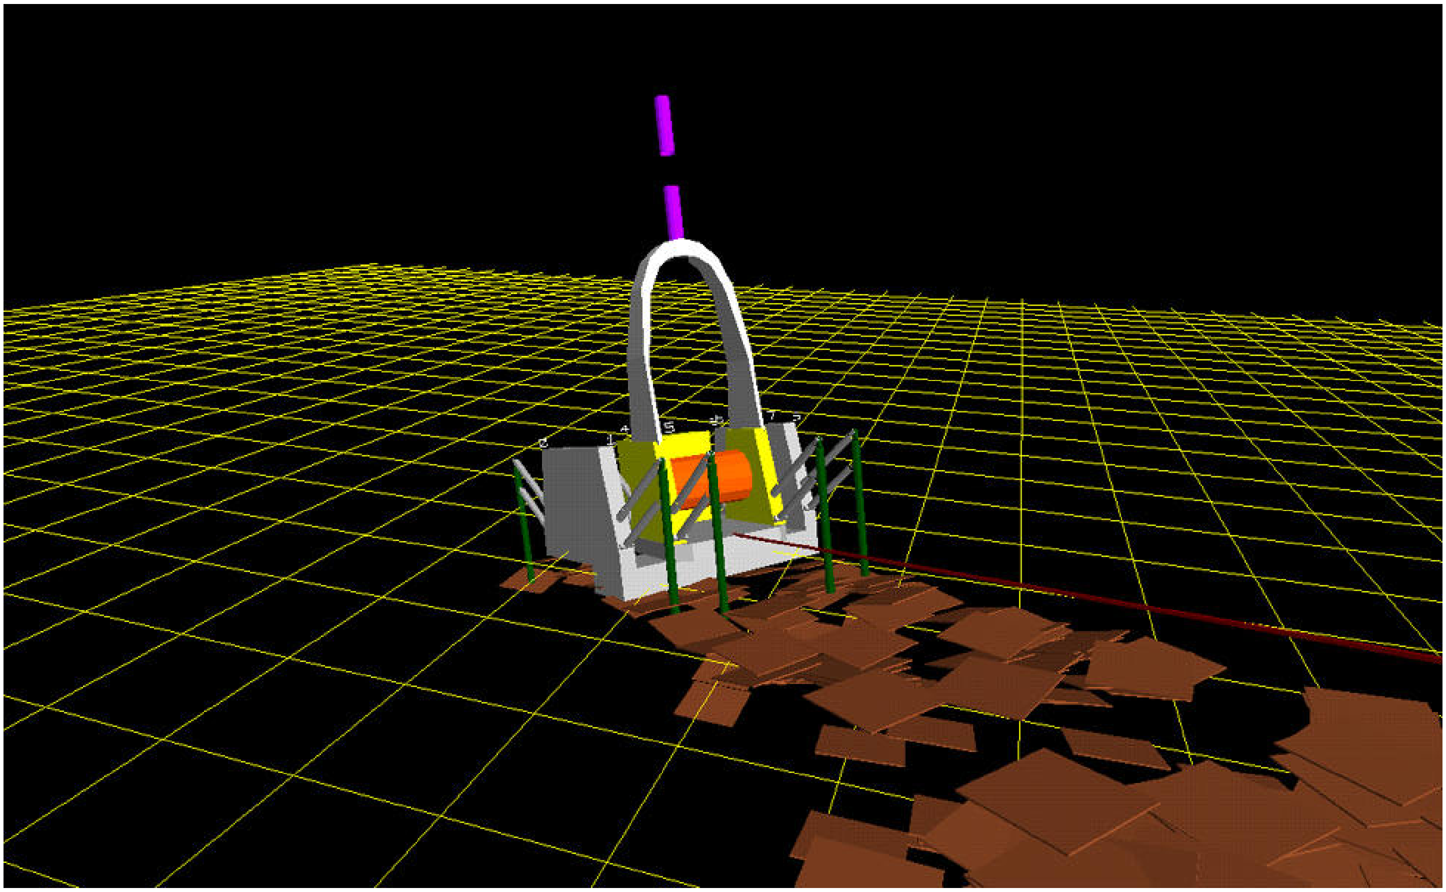

Figure 5: Dante II in the Virtual Environment Vehicle Interface (VEVI). The graphical model contains approximately 1,000 triangular polygons. The colored vertical leg links and tether represent sensed forces: shades of green (shown above as light gray) indicate normal loading range and shades of red (shown above as dark gray) indicate high or excessive loading. The quadrilateral polygon "footpads" mark the contact point for each leg step and are automatically created when contact pressure exceeds a threshold. The rectilinear grid provides operators with a gravity normal horizontal reference.

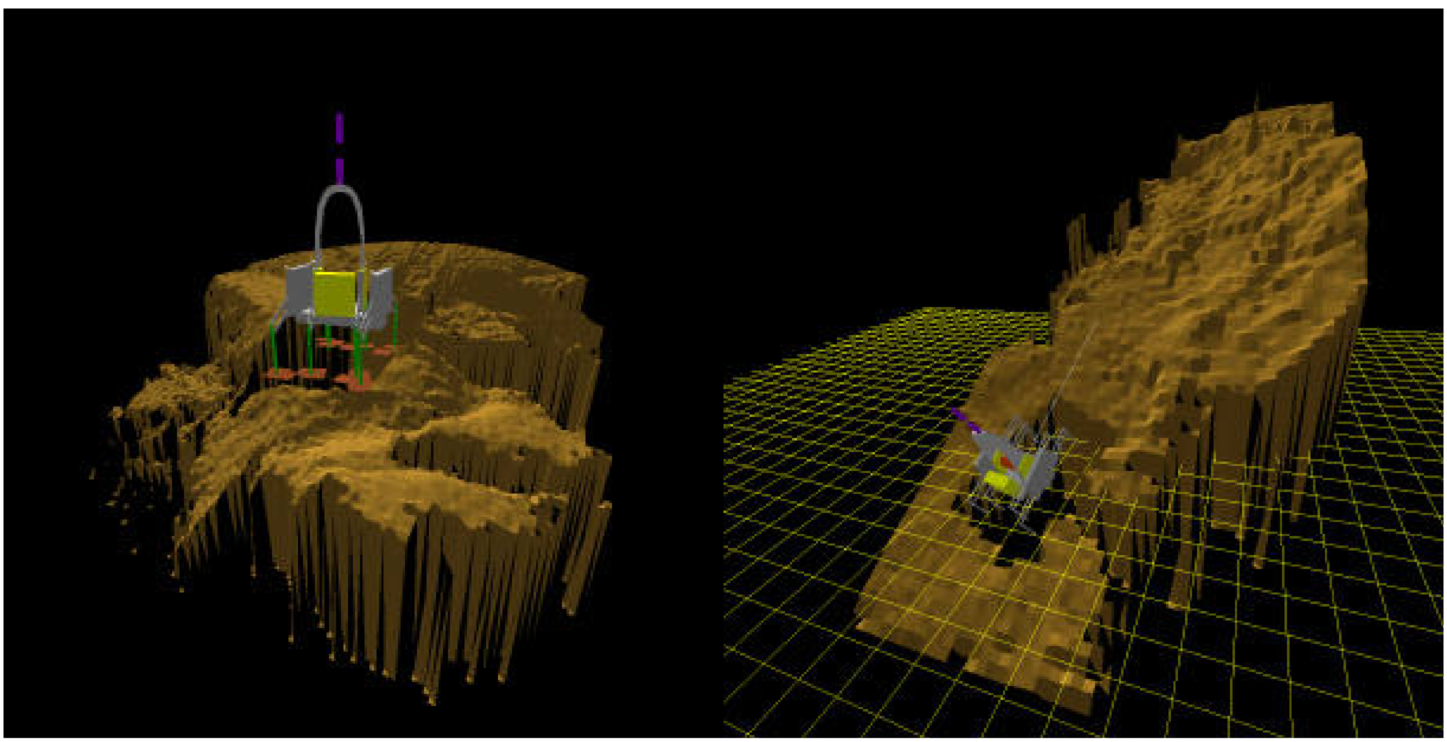

Figure 6: a(left)**, b**(right)
Visualization of scanning laser range-finder elevation maps acquired during Mount Spurr descent.
(note: the rectangular region below Dante II has been removed for clarity)
(a) Single elevation map (10m by 10m, 10cm grid resolution)
(b) Terrain mosaic created by merging successive elevation maps

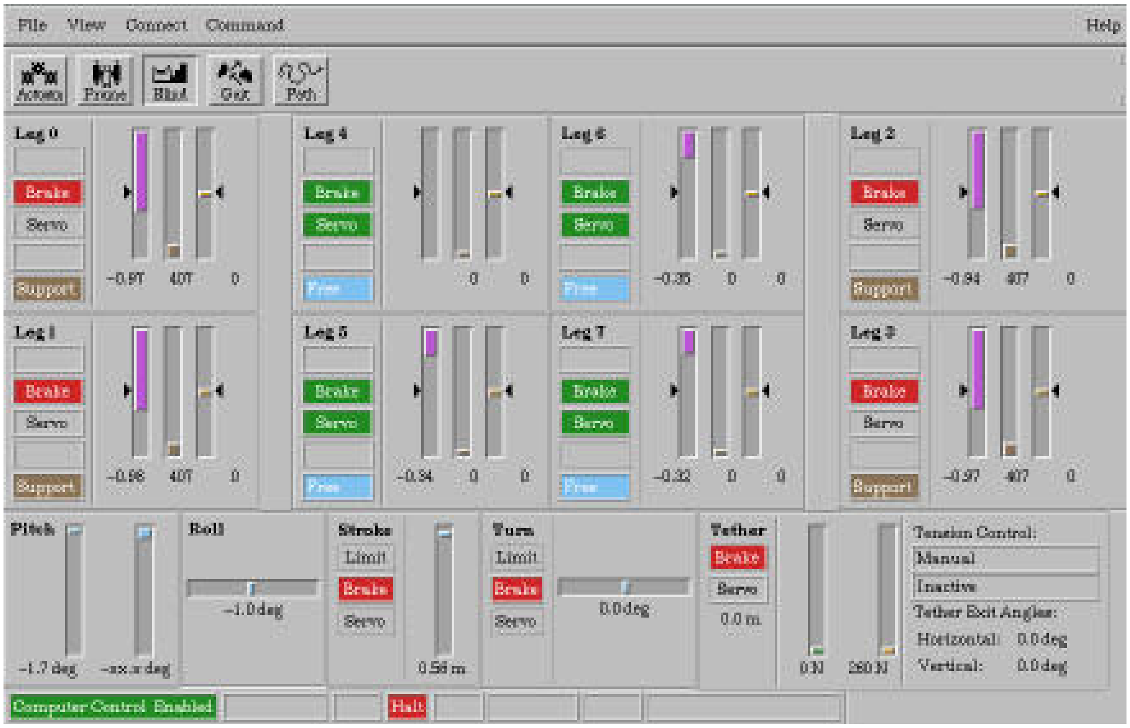

Figure 7: UI2D main window displaying positions and forces on Dante II's eight legs, its body, and its tether

sample. The terrain model represents data in a 10m by 10m square, sampled at 10 cm resolution and rendered as two tri-polygons per grid cell. The right figure provides a view of a terrain mosaic. This mosaic was created by combining successive elevation maps with dead-reckoned navigation data. We observed good correlation between data sets due to the stable and predictable manner in which a frame-walking robot traverses the terrain.

4. Experiences

We made extensive use of VEVI during Dante II's descent into Mount Spurr. Mission operators at the Anchorage base station and numerous remote observers (at NASA Ames and the National Air & Space Museum) were able to interact with the interface and view the graphical representations of Dante II's configuration and surroundings.

Upon integration of VEVI into the Dante II control station, we saw a substantial improvement in the operator's situational awareness and ability to assess the remote environment. Specifically, we learned that the interface provided a clear and easy to understand representation of both robot configuration and terrain features. We also discovered that the display of haptic terrain perception was extremely useful for gaining an understanding of explored terrain shape.

Operators who had experience with the display tools used during the Dante I / Mount Erebus expedition (precursor to Dante II / Mount Spurr), felt that VEVI was a revolutionary improvement over the earlier tools. The high frame rate, level of interactivity, and ease of use were all noted as contributing towards immersiveness and sense of presence. We agree, however, with Sheridan who states that *presence* is an extremely subjective sensation, and not amenable to objective physiological definition and measurement[12].

During the Mount Spurr descent, we observed several deficiencies of the integrated Dante II / VEVI system. A lack of redundant sensing and sensor health monitors was reflected in distorted graphical models. A failure of the on-board tether exit angle encoder, for example, resulted in erroneous tether configurations on the display. With additional information such as sensor health telemetry, this problem could have been alleviated. Additionally, when the scanning laser ranger-finder failed during ascent, the lack of current terrain elevation data severely limited VEVI's usefulness as a situational assessment aid.

Perhaps the most significant shortcomings we found in our virtual environment interface design were a failure to provide sufficient visual reference aids including grids, orientation display, world and local axes, and landmarks as well as a lack of high accuracy correlation between graphical models and physical devices. Insufficient visual aids lead to human judgement errors in displayed direction and scale, similar to those described by Ellis[13]. Insufficient viewpoint constraints exacerbated the problem by causing users to become disoriented. The lack of graphical-to-physical correlation decreased the level of confidence exhibited by mission operators when faced with conflicting information from camera systems. Higher trust would have been given to the interface if we had been able to reduce registration errors between the virtual environment and physical world to less than the five percent that was achieved.

C. UI2D

1. Overview

The UI2D is the multi-modal control interface that enabled shared and traded control of Dante II. We used UI2D to monitor sensor and state information, and to generate commands to Dante II. It is a graphical user interface, displaying buttons, knobs and slider-bars in a windowed, computer display. The main window of UI2D is shown in Figure 7.

2. Objectives

Our primary objective in creating a command interface was to make it as easy as possible to operate Dante II. We were motivated both to minimize mission operator workload and to make it possible for novices to quickly learn to control Dante II. Thus, we defined the following guidelines for UI2D development:

- *Consistent appearance and interaction*
  Consistent look and function allows operators to focus on robot actions rather than the mechanics of using the interface. Inconsistent and confused expectations make operation more difficult.
- *Functional organization*
  Grouping of similar functions and information display allows operators to apply different types of commands for different situations.
- *Uncluttered layout*
  Clean graphical design with qualitative representations of sensor and state information allows quick assessment of current conditions. Quantitative data provides precision but should support graphical features unobtrusively.
- *Simple command generation*
  Clear, easy-to-use controls allow efficient, rapid command sequences. Easily modified values and reusable commands are important for reducing operator workload during teleoperation.
- *Visual indication of safeguards*
  Different command safeguards are appropriate depending upon the situation and the types of commands being applied. Indicators that clearly reflect active safeguards reduce operator misconceptions and error.

3. Design & Implementation

To control Dante II, the operator needs to apply different types of commands for different situations. We wanted, therefore, to organize UI2D so that commands appropriate for a particular type of function or operation would be grouped together. To do this, we identified a collection of *operational control contexts*, which explicitly defined the set of commands and the information provided to the operator. In the case of Dante II, these contexts are *Individual Actuator*, *Frame*, *Blind Walk*, *Gait*, and *Path*. Each context other than *Individual Actuator* adds *elements* of autonomy, as shown in Table 2. We found that these operational control contexts allowed us to unify and simplify the human-machine interaction and to establish a range of functionality from direct teleoperation to full autonomy[3].

In the *Individual Actuator* context, the operator can teleoperate each actuator with no command safeguarding other than checks on the limits of motion. It is useful for debugging and system checkout. Since the tether is not automatically servoed and because manually tensioning it during body motion is impractical, this context cannot be used to make the robot walk. *Individual Actuator* is used occasionally during teleoperation. For example, when a single leg hits a rock during translation, the operator may wish to raise just that leg.

In the *Frame* context, actuators can be commanded in groups to allow walking by raising, lowering, stroking, and turning frames of legs. Figure 8 shows the *Frame* context pop-up window, which is activated from the UI2D main window *Frame* button. The *Frame* context enables teleoperation of Dante II as a frame-walking system. The key automatic feature is the coordination of tether payout with body motion. On level terrain the tether pays out in direct proportion to the amount the body moves. As the terrain steepens, inclinometers measure robot pitch and compute the component of the robot weight that the tether must support. The tether tension is automatically controlled to the proper value. The *Frame* context typifies shared supervisory control.

Table 2: Operational control contexts and elements of autonomy

Control Function	Individual Actuator	Frame Context	Blind Walk Context	Gait Context	Path Context
Servo tether					
Maintain body height					
Maintain posture					
Adjust leg step					
Surmount obstacles					
Perceive terrain					
Determine step height					
Determine body height		User Control		Automatic	
Determine posture					
Determine stride					
Correct leg placement					
Generate path					
Determine heading					
Avoid obstacles					

In the *Blind Walk* context, the operator specifies the set of parameters that describe the desired gait. This gait is commanded to the robot, which then executes it autonomously until told to stop or until an insurmountable obstacle is encountered. Perceptive sensors are disabled; only proprioceptive sensors are used to detected contacts and maintain posture, hence the name *Blind Walk*. We considered this context as traded supervisory control.

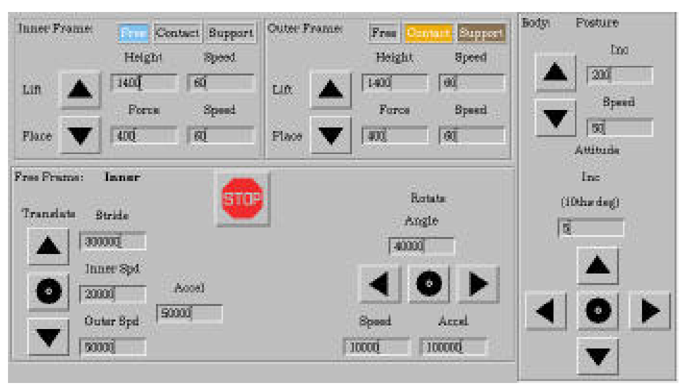

Figure 8: UI2D *Frame* context popup window enables operation of Dante II as a frame-walking system. Buttons are used to command groups of legs and the frames. Numerical entry fields allow quick but precise specification of desired motions.

In the *Gait* context, the operator commands only a desired trajectory for the robot to follow. Perceived terrain information from the scanning laser rangefinder is built into an elevation map. This map and other motion constraints that arise from robot kinematics, stance stability, and body flexibility are analyzed to identify a set of gait parameters for efficiently crossing the terrain. In contrast to *Blind Walk*, in which values are manually determined, this context uses a gait planner for automatic parameter generation. When the *Gait* context is selected, Dante II's planning capabilities are activated and it can plan its own actions autonomously.

The *Path* context is envisioned as the context in which the operator need only specify locations that the robot should visit. A global map of robot-scale resolution is needed to identify a path for Dante II to follow. The path avoids obstacles that are insurmountable and can be broken into shorter trajectories that the gait planner then transforms into executable gaits. In the *Path* context Dante II is fully autonomous.

4. Experiences

We have had a number of unsatisfying experiences with text-based operator interfaces for both Dante I and Dante II. They invariably provide a daunting array of numerical information and an indecipherable collection of command codes. Yet we were surprised at what an improvement this graphical interface was over previous interfaces. As soon as UI2D was available all other operator interfaces fell into total disuse.

Anecdotal evidence suggests that operators were able to run Dante II longer, faster, and safer with UI2D. We found that visitors watching operations were able to quickly grasp how the interface worked and to begin to suggest what to do. Although UI2D did not contain a graphical visualization of robot and sensor state, it provided sufficient information so that in several instances, operators stopped watching video monitors and the VEVI display and ran the robot based solely on the information that it provided.

D. VIDEO SYSTEMS

1. Overview

To further enhance the operator's sense of presence in teleoperating the robot, we equipped Dante II with seven color video cameras mounted in strategic on-board locations. Four fixed focal-length cameras, mounted low at the four corners of the robot's body, provided close-up views of the ground contact point for each leg. Atop the arched mast, a camera enclosure (containing a stereo camera pair and a variable-zoom camera) was mounted on a pan/tilt platform, allowing the operator to view the surroundings and a top-down view of the robot itself. We placed a variable-zoom camera on a pan/tilt platform at the volcano rim during the mission to obtain fixed-point, exocentric views of the robot and crater environment.

Due to bandwidth limitations of Dante II's tether conductor and satellite telemetry, all eight video signals could not be simultaneously transmitted to the control station at full video rates. Instead, we routed the signals to a pair of remotely-controllable four-channel switch and video compression units, then transmitted the compressed output signals via two 384 Kbps satellite channels.

At the control station, we used two color monitors to independently display either full-sized images from any camera or a combined quad-view image. This latter mode of display enabled simultaneous viewing of four cameras, but with a significant loss of picture resolution and quality. We also used a field-sequential stereo monitor that facilitated viewing of the stereo image pair through eyewear incorporating active liquid crystal shutters. Selection of images to be displayed, control of the quad splitter, and adjusting pan & tilt angles of the on-board camera pod and rim camera were all integrated through a graphical user interface.

2. Experiences

A significant portion of the Mt. Spurr descent required direct teleoperation with the UI2D *Frame* context due to the extreme ruggedness of the terrain. Operating Dante II in this mode required the operator to understand the remote environment, a task that was greatly aided by having direct visual feedback from live video images. These images also allowed us to make qualitative estimates of soil characteristics such as firmness and moisture content.

By studying images which contained parts of the robot as well as the natural environment, we were able to coarsely estimate the size of objects (e.g., boulders, ditches). This task was easier with the four corner cameras than the mast-mounted cameras, which frequently contained few or no parts of the robot body and were typically aimed at the path ahead. Changes in ambient lighting conditions, shadows and disturbances such as fumarole steam or fog occasionally impeded our assessment efforts. Slow video update rates did not have a significant impact due to the low speed of typical Dante II's motions. However, image resolution (especially when using the combined quad view) limited the observable detail.

During certain critical phases of operation, such as walking over boulders near the limits of Dante II's capabilities, we found a need for additional camera views to assess under-body clearance. Similarly, the fixed location and orientation of the leg cameras proved to be difficult to optimize a priori; at certain times during the mission the capability to realign them would have been advantageous.

In operating the pan and tilt motions of the mast-mounted camera pod, the importance of closed-loop control became apparent: by commanding a target pan or tilt *angle*, rather than explicitly starting and stopping the corresponding motors, the satellite transmission delay did not have a significant adverse impact.

We observed that due to concentration required by the operator, even small tasks such as selecting and switching between the desired camera views became a potential source of distraction. A preferable scenario would have employed more monitors with all available views present simultaneously, so that the operator would only have to look at the screen of choice. In general, we observed a trade-off between the amount of information available and the mental effort required to make effective use of it. For example, the stereo imaging added utility in only a few instances due to the effort required to don the eyewear, to focus on the monitor and to allow eye accommodation/convergence to the stereo display.

NETWORK-BASED PARTICIPATION

A. OVERVIEW

This section presents the network-based participation methods developed for the Dante II / Mount Spurr expedition. We used network methods including message communications, satellite transmission, and a WorldWideWeb (WWW) server to reach remote mission observers and participants at numerous locations. Specifically, these methods enabled us to provide high levels of public interaction, to disseminate mission status and information, and to conduct collaborative science throughout the mission.

We first present an overview of the communications infrastructure which supported network-based participation. We then present the WorldWideWeb services which were used for the real-time distribution of information and which provided significant access to the mission. Finally, we describe the impact of network-based participation methods at two remote observation sites.

B. DANTE II COMMUNICATIONS

To support the network communication required by the Dante II / Mount Spurr expedition, we utilized many technologies and techniques. Data transmissions were carried via point-to-point satellite and terrestrial communications links as well as over existing wide area networks. Audio and video transmissions were handled separately from data over additional point-to-point links. During a portion of the mission, these links included Ka-band communications through the Advanced Communications Technology Satellite (ACTS), an experimental multi-spot satellite system developed by NASA[14].

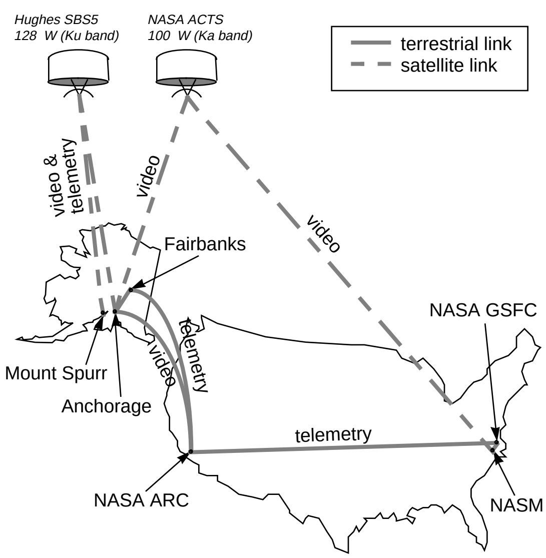

Figure 9: Dante II network and communication links. Video and telemetry from Dante II were relayed via commercial satellite from Mount Spurr to the mission control station in Anchorage. Terrestrial links carried these signals from Anchorage to the NASA Ames Research Center (ARC) in Moffett Field, CA. Satellite communications provided connectivity to the National Air and Space Museum (NASM) in Washington, D.C.

Shown in Figure 9 are the primary network and communication links used by Dante II. Video and telemetry originating on-board Dante II was relayed via commercial satellite from Mount Spurr to the mission control station in Anchorage. Terrestrial fiber and wire lines carried these signals to the NASA Ames Research Center for distribution to existing wide area networks: video on the NASA Select television system and telemetry on the NASA Science Internet. Satellite communications through ACTS provided video support to an exhibit at the National Air and Space Museum. We used ethernet network

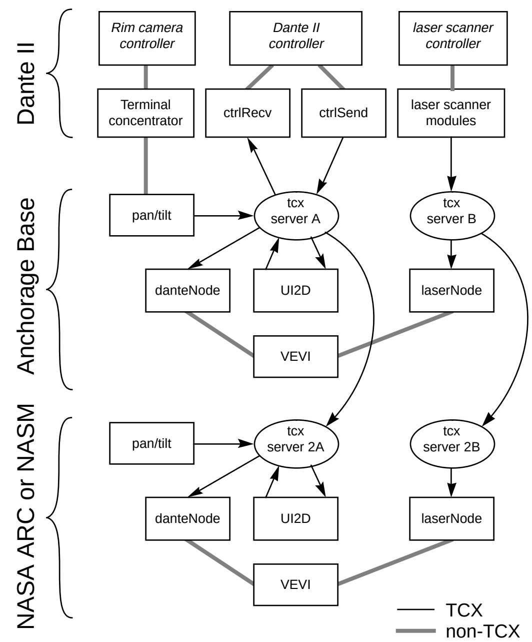

Figure 10: Dante II TCX communications structure. Process modules pass messages via centralized server modules. The server modules provide control of message routing to optimize performance across point to point links.

devices and the TCX messaging system for the majority of Dante II data communications. TCX is an interprocess communications package designed to support mobile robotic applications across multiple hardware platforms, operating systems and programming languages. TCX provides the communication building blocks to coordinate flow of control and flow of data between planning, perception, and real-time control in mobile robots[15]. Figure 10 depicts the data flow through TCX.

C. WORLD WIDE WEB

1. Overview

To provide real-time public access to the Dante II / Mount Spurr expedition, we operated a WorldWideWeb (WWW) server from the NASA Ames Intelligent Mechanisms Group laboratory[1]. This server provided current and archived mission information via the Internet to the general public, remote scientists, and others working on the mission. One of our objectives was to use the WWW server as a testbed of public interaction techniques for future robotic planetary missions.

2. Background

The WorldWideWeb, a networked hypertext system, was initially developed at CERN in 1991[16]. It has since grown to include thousands of servers world-wide. Uses of the Web currently range from distributing commercial product information to entertainment to disseminating educational and scientific materials. A dramatic example of real-time information distribution via the Web occurred during July 15-17, 1994 when over two million accesses were logged on NASA WWW servers distributing images and information of comet Shoemaker-Levy 9's collision with Jupiter[17]. Most of the documents on the Web are written in the Hypertext Markup Language (HTML)[18]

[1]The Dante II WWW server may be accessed via the Uniform Resource Locator <http://maas-neotek.arc.nasa.gov/dante>

and are distributed using the Hypertext Transfer Protocol (HTTP)[19]. These documents may contain links to images, sounds, and animations as well as documents residing on servers. During the Dante II mission, we operated a server on a workstation running the National Center for Supercomputing Applications' *httpd*[20].

The WWW is an extremely flexible media and does not suffer from many of the constraints of traditional public interaction methods, such as video broadcasts and fixed observation sites, which are used to provide public mission access. Video broadcasts provide a wide audience with live access to events, but offer limited opportunity for viewer participation. Viewers are typically constrained to a single audio/video stream and are unable to replay prior events or to focus the broadcast to portions of specific interest to them. Fixed observation sites can similarly provide live access, but only to a limited number of observers who must travel to the site. One advantage of these sites to video broadcast is that observers have much greater freedom to focus and explore areas of interest. With the WWW, however, participants are able to view compressed video and still images, read and search mission status reports, and explore technical details of their choosing. As an additional benefit, the WWW images and documents collected during the mission, including those generated via public interaction, provide a permanent mission record.

3. Server Design

We found during our WWW server development that creating hypertext documents in HTML is straightforward, but that designing a cohesive, efficient document structure can be challenging. When creating pages for the WWW, trade-offs must be made between aesthetic appearance, human interface issues and server limitations. We strove to design pages that are engaging and visually attractive yet present information which is both concise and complete. At the same time, we made an effort to minimize data transfer requirements. We found this second guideline to be necessary to reduce the load on the server, decrease the transfer time for users with slow connections, and enable users to rapidly find information of interest. Finally, we attempted to support as many WWW browsers as possible. Although NCSA Mosaic is the predominant WWW client, many other browsers including Lynx and NetScape™ exist. To ensure portability across all browsers, it is necessary to adhere closely to the HTML specification, provide text alternatives to image-based menus, and give text descriptions along with inlined images.

We set up the Dante II / Mount Spurr WWW server with a shallow hierarchical structure so that users could quickly locate areas of interest. Shown in Figure 11 is the top level Dante II document which contains a main menu followed by a brief description of the Mount Spurr expedition. The main menu has links to three child pages (*Status*, *Technical Info*, and *Images*) and was implemented both as an *imagemap* and as a text-only menu. Imagemaps are inlined images with *hot regions*, portions of the image which are linked to other documents. These image-based menus have more visual appeal than text-only menus and usually transfer faster then multiple graphical buttons. However, since *imagemaps* do not provide visual feedback (the individual selections do not highlight with pointer movement), the

Figure 11: Dante II top-level WWW page. The main menu contains links to the "Status", "Technical Info", and "Images" pages. Note the use of both *imagemap* image-based menu and an alternative text-only menu (located below).

boundaries of the menu options must be easy to determine from the image content.

The *Status* page contained two main sections. At the top of the page, current mission information was accessible via links to press releases and daily status reports. In addition to a chronological listing, a *forms* interface provided keyword search capability. Due to the length and delays associated with retrieving lengthy reports, we also provided a brief summary of the immediate mission status. This summary, which we updated as frequently as every ten minutes, proved useful to both the public and to those directly involved with the mission. At the bottom of the *Status* page, we included an interface for users to obtain the latest images captured from each of Dante II's cameras (see Figure 12). These images were captured by the vehicle operator in real-time and were automatically inserted into the WWW server. Since HTTP is a stateless protocol, there is no way for the server to asynchronously inform client browsers of changes. Thus, we observed that users found it difficult to detect when images had been updated. One solution to this problem is to display creation-time information along with each image, although we did not implement it for the Mount Spurr expedition. The *Technical Info* page of the WWW server was set up to provide educational and technical information about Dante II. Included were sounds, system descriptions, and animations of Dante II walking. We also provided descriptions of the mission's goals were under this hierarchy.

The *Images* page was created initially as a list of selected, interesting images and MPEG movies. Each entry in the listing consisted of a thumbnail image, a description, and image statistics. As the mission progressed, we discovered that large numbers of thumbnails severely impacted system responsiveness due to data transfer delays. As a result, we redesigned the page with a hierarchical structure using image categories. We also

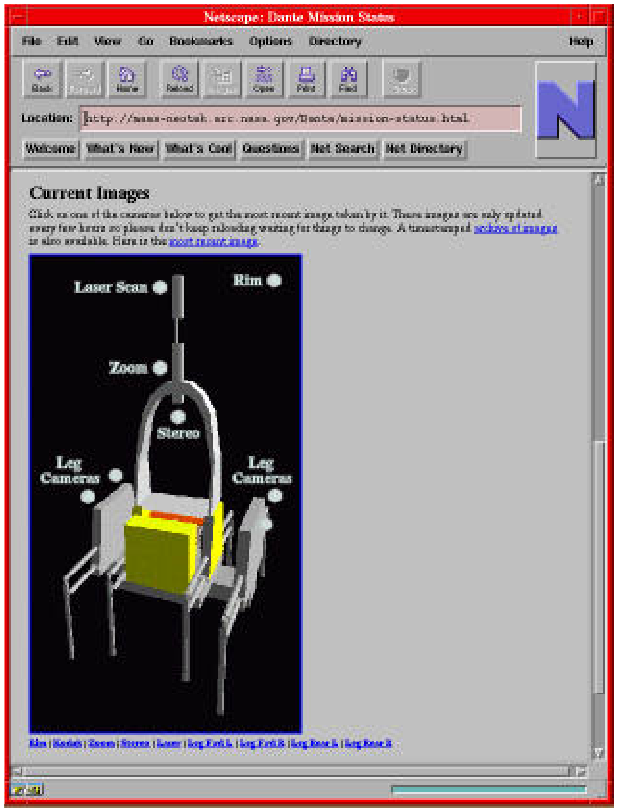

Figure 12: WWW current image interface. This interface enabled users to obtain the latest images captured from each of Dante II's eight cameras and laser scanner. Users found the graphical presentation to be much more intuitive and comprehensible than other formats.

included a canonical index of all the images in the archive, with timestamps incorporated into the filename. On request, an index was dynamically generated by converting the directory listing to HTML. However, as the mission progressed (and the index length increased) the time required to generate it also grew. By the end of the Mount Spurr expedition it took a considerable amount of time to generate listings on demand. A future refinement would be to periodically generate and store the listing and rather than attempting to generate it at each request[21].

4. Experiences

At the peak of the mission, our server was logging over 20,000 connections each day. With this load we started to notice an impact on our WWW server machine and on the fileserver which stored all documents and images. If future missions gain attention similar to the Shoemaker-Levy 9 collision, steps will need to be taken to protect the server and the network from overloading[17]. Also, if the WWW server is performing other functions in the mission, an excessively high load may interfere with mission operations. One solution to decrease the load on servers involves designing the document layout to discourage people from browsing, thus decreasing the number of accesses. Other possible solutions include using a local disk on a dedicated WWW server for storing documents and images, statically creating automatically generated documents (such as image lists), and distributing or replicating parts of the server to other servers. The latter solution is also useful for providing information to sites with slow network connections. During the Dante II / Mount Spurr expedition, we wanted participants at the National Air and Space Museum remote observation site to have access to the WWW server. Due to slow network communication links, we replicated most of the server at the site.

Overall, we received a great deal of positive feedback about the WWW server from researchers, educators, and enthusiasts. Whenever status reports were delayed for a few days, we usually received numerous email messages from people who were closely following the mission over the Web. We have also had requests from educational institutions for copies of the WWW server documents and images on a CD-ROM or tape. Hopefully this technology will be useful for providing real-time public access to future planetary missions.

D. REMOTE OBSERVATION SITES

1. Overview

As part of our effort to develop robotic planetary exploration methods, we are evaluating the ability of our network-based methods to enable public interaction and collaborative science at remote observation sites. For Dante II, we chose to assess these methods at two locations: (1) a public exhibit at the National Air and Space Museum (NASM) in Washington, D.C., and (2) a laboratory at the NASA Ames Research Center in Mountain View, California.

2. Public Interaction

As a precursor to Dante II's deployment on Mount Spurr, we performed tests at an Anchorage gravel pit to verify telemetry and system software. The tests took place over a two day period during which Dante II descended and ascended a path of 92m[3]. During these tests, we used a public exhibit at the NASM to provide continuous, real-time access to the Dante II mission. Ka-band satellite communications via ACTS provided live video transmissions (NTSC color monocular and stereo) from Dante II's cameras to display monitors in the NASM. Terrestrial wide area networks and TCX messaging relayed telemetry from Dante II, enabling real-time monitoring and remote camera control via operator interfaces (VEVI, UI2D) on workstations in the NASM. Additionally, hourly audio status updates were provided via telephone lines.

Our objective for the NASM exhibit was to provide a publicly accessible and interactive demonstration of robotics, telepresence and virtual environment technology in exploration. We were particularly interested in engaging the public, creating awareness of the Dante project and conveying the potential of these technologies. network-based methods provided the means for to achieve these goals.

We received substantial feedback concerning the exhibit from the public, the NASM staff, and NASA representatives. To a large extent, this feedback reflects the high-level of enthusiasm and interest the exhibit garnered. We found that this type of public demonstration is significant for three factors: (1) it provides dramatic, dynamic and stimulating mission coverage unparalleled by traditional media, (2) it encourages the public to actively participate and to share in the excitement of scientific exploration, and (3) it makes NASA missions *accessible* and *comprehensible* at a personal level.

3. Remote Science Collaboration

Throughout Dante II's descent into Mount Spurr, we used a laboratory at the NASA Ames Research Center to examine the use of network-based methods for supporting collaborative science. Terrestrial networks relayed video and telemetry from the Anchorage base station to NASA Ames. Planetary scientists and other observers located at NASA Ames were able to view and periodically control Dante II's on-board cameras. Operator interfaces provided continuous, real-time visualization of Dante II's configuration and sensor data.

After descending over 200m through snow and rocky slopes, Dante II reached the crater floor and remained in the area for several hours to permit study of the active fumarole region. During this period, a planetary scientist and a geologist located at NASA Ames remotely conducted an informal, real-time, video field survey and geologic interpretation of the crater floor.

We were able to gain several insights from this experience. Deficiencies of the remote camera systems (e.g., resolution, viewpoint) and lack of correlating sensors limited the scientists ability to accurately identify surface features. Additionally, the absence of a manipulator prevented us from collecting geologic samples. However, two single gas sensors did allow us to characterize the effluent gases emitted from the active fumaroles. Most importantly, though, we observed that remote field science can be performed given appropriate operator interfaces and network-based methods.

CONCLUSIONS

A. LESSONS LEARNED

We found that traded and shared supervisory control can improve system performance through synergistic human-machine interaction. To realize these benefits, however, requires the design and application of appropriate human-machine interfaces. For the Dante II / Mount Spurr expedition, we developed and used virtual environment and multi-modal operator interfaces.

Virtual environment interfaces can be used to improve an operator's situational awareness and to efficiently visualize complex information. VEVI provided a clear and easy to understand representation of Dante II's configuration and sensor data. We found that high frame rate, level of interactivity, and ease of use were all contributing factors for achieving immersiveness and sense of presence. However, insufficient visual reference aids and correlation level between graphical and physical models can degrade operator performance.

Multi-modal operator interfaces support a full range of control modes ranging from direct teleoperation of individual actuators to full autonomy. UI2D helped to minimize mission operator workload and to make it possible for novices to quickly learn to control Dante II. The use of *operational control contexts* provides a unifying and simplifying perspective on human-machine interaction. This approach enabled us to concisely organize UI2D so that commands appropriate for a particular type of function or operation could be grouped together.

The use of network-based methods including satellite and terrestrial communications allowed us to provide high levels of public interaction. A WWW server was used to disseminate mission status and information throughout the mission to remote scientists and the general public. The NASM exhibit demonstrated the significance and benefit of public demonstration. It showed that live, public access to mission is both educational and exciting. Remote field science conducted from NASA Ames demonstrated the capability of network-based methods to support collaborative science.

Overall, we believe that the operator interfaces and network-based participation methods developed and utilized during the Dante II / Mount Spurr expedition represent a significant advance. These techniques have the potential to greatly improve the capability, capacity and quality of robotics technology for supporting planetary exploration.

B. FUTURE WORK

The operator interfaces we developed for Dante II raised a number of human-machine interaction questions. Though we found that our separate interfaces were complimentary, operators often found it difficult to decide where to direct their attention. One solution would be to merge the interfaces into a single unified application. But, we do not yet know how to achieve this goal. Additionally, we are interested in quantitatively evaluating the immersiveness and sense of presence in our interfaces. The lack of suitable metrics, however, leaves this issue open and unresolved.

The use of shared and traded supervisory control for operating planetary exploration robots creates a number of concerns. These include detection of anomalies and faults, negative effects of intermittent operator involvement, improper training methods (i.e., how does training reflect the cognitive nature of supervising), and errors due to distributed, shared decision making[4].

Though there are currently no plans for a successor to Dante II, the technologies developed during the Mount Spurr expedition continue to be used for other applications. In February 1995, NASA Ames and others conducted lunar and planetary mission simulations with a Marsokhod rover deployed on the Kilauea volcano. These simulations utilized technology and techniques derived from Dante II including VEVI, the World-WideWeb services, and much of the communications infrastructure. Finally, NASA and the NASM have begun working to establish a permanent exhibit similar to the Dante II / Mount Spurr demonstration. The exhibit will showcase virtual environment and stereo video based telepresence as enabling technologies for robotic planetary exploration.

ACKNOWLEDGMENTS

The development of Dante II and the expedition to Mount Spurr were supported by the National Aeronautics and Space Administration under grant NAGW-1175 and NASA RTOP 233-02-04 (D. Lavery, program manager).

The authors wish to thank all the members of the Dante Project for their tireless efforts and dedication. We also acknowledge the invaluable contributions and support of numerous organizations including NASA Headquarters, the NASA Ames Research Center (Codes IC, ID, and SS), the NASA Science Internet, the NASA Lewis Research Center (ACTS program office), the StereoGraphics Corporation and the Sense8 Corporation.

REFERENCES

[1] Wettergreen, D., Thorpe, C., and Whittaker, W., "Exploring Mount Erebus by Walking Robot", Robotics and Autonomous Systems 11, 1993, pp. 171-185. also in

Proceedings of the Third International Conference on Intelligent Autonomous Systems, February 1993, p. 72-81.

[2] Lavery, D., and Weisbin, C., “Telerobotics Program Plan”, NASA Office of Advanced Concepts and Technology, January 1993.

[3] Wettergreen, D., Pangels, H., and Bares, J., “Gait Execution for the Dante II Walking Robot”, IEEE Conference on Intelligent Robots and Systems, 1995.

[4] Sheridan, T., “Telerobotics, Automation, and Human Supervisory Control”, MIT Press, Cambridge, MA, 1992.

[5] Hine, B., et. al., “The Application of Telepresence and Virtual Reality to Subsea Exploration”, The 2nd Workshop on Mobile Robots for Subsea Environments, Proceedings ROV ‘94, Monterey, CA, May 1994.

[6] Fong, T. “A Computational Architecture for Semi-autonomous Robotic Vehicles”, AIAA 93-4508, AIAA Computing in Aerospace 9, San Diego, CA, October 1993.

[7] Piguet, L., Fong, T., Hine, B., Hontalas, P., and Nygren, E., “VEVI: A Virtual Reality Tool for Robotic Planetary Explorations”, Virtual Reality World 95, Stuttgart, Germany, February 1995.

[8] Hine, B., Hontalas, P., Piguet, L., Fong, T., and Nygren, E., “VEVI: A Virtual Environment Teleoperations Interface for Planetary Exploration”, SAE 25th International Conference on Environmental Systems, San Diego, CA, July 1995.

[9] Stevens, H.D., et. al. “Object-Based Task-Level Control: A Hierarchical Control Architecture for Remote Operation of Space Robots”, AIAA/NASA Conference on Intelligent Robots in Field, Factory, Service, and Space, Houston, TX, April 1994.

[10] Garvey, J., “A Russian-American Planetary Rover Initiative”, AIAA 93-4088, AIAA Space Programs, and Technologies Conference and Exhibit, Huntsville, AL, September 1993.

[11] Apostolopoulos, D., and Bares, J., “Configuration of a Robust Rappelling Robot”, IEEE Conference on Intelligent Robots and Systems, 1995.

[12] Sheridan, T., “Musings on Telepresence and Virtual Presence”, Presence, 1(1), 1992.

[13] Ellis, S.R, Smith, S., Grunwald, A., and McGreevy, M., “Directional Judgement Error in Computer Generated Displays and Actual Scenes”. In Ellis et. al, “Pictorial Communication in Virtual and Real Environments”, Taylor and Francis, London, 1991.

[14] Wright, D., and Balombin, J., “ACTS System Capability and Performance”, AIAA-92-1961-CP, AIAA 14th International Communication Satellite Systems Conference, 1992.

[15] Fedor, C., “TCX: An Interprocess Communications Systems for Building Robotic Architectures”, internal document, Robotics Institute, Carnegie Mellon University, January 1994.

[16] Berners-Lee, T., “The World Wide Web Initiative”, CERN European Particle Physics Laboratory, 1995. <http://info.cern.ch/hypertext/WWW/TheProject.html>

[17] Towheed, S., “NASA’s Use of the World Wide Web to Deliver Shoemaker-Levy 9 Collision Data in Near-Real Time”, Electronic Proceedings of the Second World Wide Web Conference, Chicago, IL, October 1994. <http://www.ncsa.uiuc.edu/SDG/IT94/Proceedings/Astronomy/towheed/towheed.html>

[18] Berners-Lee, T., “HyperText Markup Language (HTML): Working and Background Materials”, CERN - European Particle Physics Laboratory, 1995. <http://www.w3.org/hypertext/WWW/MarkUp/MarkUp.html>

[19] “HTTP: A Protocol for Networked Information”, Internet Engineering Task Force, draft, 1994. <http://www.w3.org/hypertext/WWW/Protocols/HTTP/HTTP2.html>

[20] McCool, R., “NCSA httpd Overview”, National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign, 1994. <http://hoohoo.ncsa.uiuc.edu/docs/Overview.html>

[21] Jennings, D., Damon, P., Good, M., and Pisarski, R., “How to Present Lots of Volatile Information on the World Wide Web”, Electronic Proceedings of the Second World Wide Web Conference, Chicago, IL, October 1994. <http://www.ncsa.uiuc.edu/SDG/IT94/Proceedings/Astronomy/jennings/jennings.html>