John E. Bares
The Robotics Institute
Carnegie Mellon University
bares@cs.cmu.edu

David S. Wettergreen*
Department of Systems Engineering
The Australia National University
Canberra, ACT 0200, Australia
dsw@syseng.anu.edu.au

Dante II: Technical Description, Results, and Lessons Learned

Abstract

Dante II is a unique walking robot that provides important insight into high-mobility robotic locomotion and remote robotic exploration. Dante II's uniqueness stems from its combined legged and rappelling mobility system, its scanning-laser rangefinder, and its multilevel control scheme. In 1994 Dante II was deployed and successfully tested in a remote Alaskan volcano, as a demonstration of the fieldworthiness of these technologies. For more than five days the robot explored alone in the volcano crater using a combination of supervised autonomous control and teleoperated control. Human operators were located 120 km distant during the mission. This article first describes in detail the robot, support systems, control techniques, and user interfaces. We then describe results from the battery of field tests leading up to and including the volcanic mission. Finally, we put forth important lessons which comprise the legacy of this project. We show that framewalkers are appropriate for rappelling in severe terrain, though tether systems have limitations. We also discuss the importance of future "autonomous" systems to realize when they require human support rather than relying on humans for constant oversight.

KEY WORDS—rappelling, walking robot, supervised autonomy, behavior-based control, terrain visualization, volcano exploration

1. Introduction

Relying on support from an actively tensioned tether cable and vision from a scanning-laser rangefinder, Dante II rappelled into the crater of an Alaskan volcano and operated alone for nearly a week, exploring previously uncharted terrain and sampling volcanic gases (see Fig. 1). The legged robot traveled one-quarter of its 165-m descent into Mount Spurr autonomously, relying only on onboard sensors and computers to plan and execute its motion. Challenges included crossing 1-m boulders on ash-covered slopes, as well as navigating areas of deep snow, ditches, and rubble. When terrain difficulty exceeded the capability of the automated planning system, supervisory teleoperation was used to guide the robot. Throughout its exploration mission, Dante II streamed a variety of scientific data to operators and volcanologists located at remote sites. Several times the surrogate robot scientist was struck by boulders cascading down the steep crater walls. After arriving at the heavily bouldered crater floor, the robot measured the gas composition of several large fumarole vents. While climbing out of the crater, Dante II lost stability on steep, muddy terrain, and fell on its side, ending its mission.

Dante II's expedition to Mount Spurr during 1994 was a significant achievement in remote robotic exploration, and has provided important insight into high-mobility robotic locomotion. Several terrestrial robots have operated with more autonomy than Dante II, but in much less-demanding terrains and usually with close human oversight. As an example, the NavLab, with safety observer onboard, autonomously traveled several kilometers of moderate terrain in search of an obstacle-free path to a goal (Stentz and Hebert 1995). Other robots have achieved long field missions, but required daily or more frequent human support (Thomas, Hine, and Garvey 1995; Wettergreen et al. 1999).

The principal objective of the Dante project was to develop and demonstrate technologies which could lead to solutions for robotic exploration of the most rugged lunar and planetary terrain. In at least the latter case, a certain amount of robot autonomy is essential, due to the latency and bandwidth of communication to Earth: planetary explorers must be self-capable, and succeed with only occasional human guidance.

*Author was a PhD candidate at Carnegie Mellon's Robotics Institute at the time of the Dante project.

The International Journal of Robotics Research
Vol. 18, No. 7, July 1999, pp. xxx-xxx,

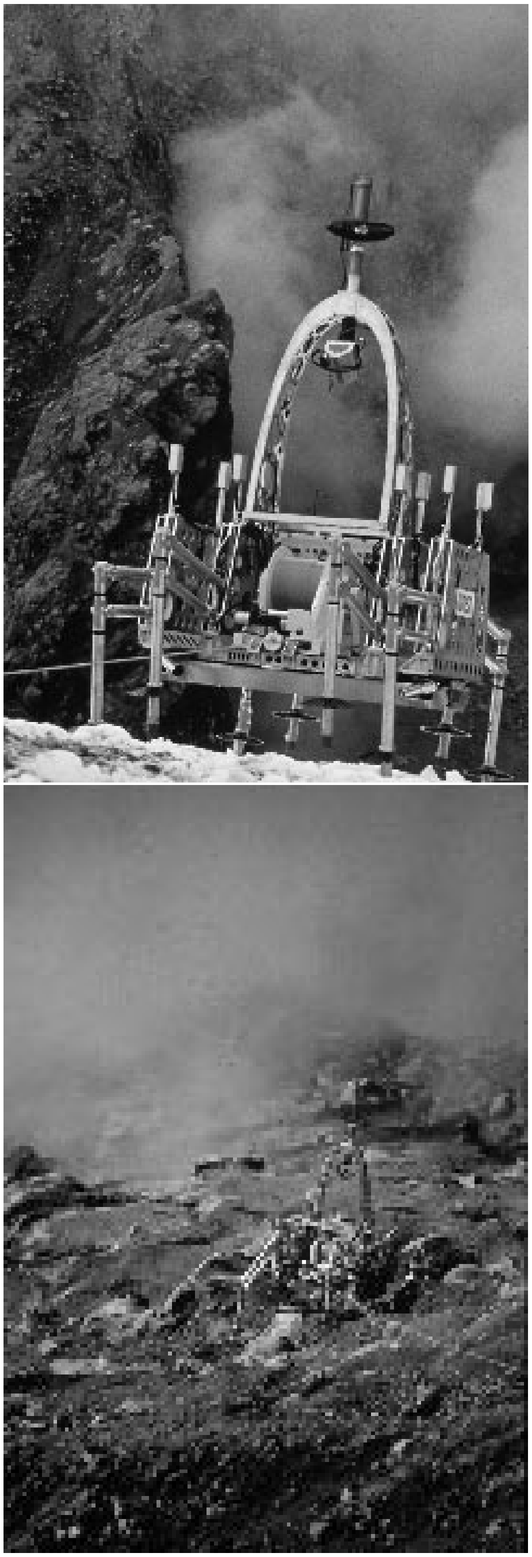

Fig. 1. Dante II enters the Mount Spurr crater by crossing a snow field on the rim of the volcano (top). Deep in the crater, Dante negotiates several large boulders and a ridge as it begins its ascent (bottom). There is low tension on the tether (note the slight sag), because the robot is on the relatively horizontal terrain of the crater floor.

Furthermore, human assistance for repair or rescue from entrapment is virtually impossible. In addition to offering lunar and planetary terrain analogs, Mount Spurr demanded some autonomous capabilities, because human operators were located far from the volcano for safety reasons and communication bandwidth to the robot was limited due to site conditions. Given the active state of the volcano, human entry to the crater floor was prohibited, which forced a certain amount of self-reliance on the robot. Mount Spurr also offered a volcanically active crater that was unexplored by humans, and therefore of some scientific interest. This formed the secondary objective of the project: to procure scientific data from the volcano and thereby broaden experience in the use of robots as surrogate field scientists.

The key accomplishments of the Dante project were:

- The development and demonstration of rappelling locomotion in natural terrain. We believe Dante II traversed the most severe terrain attempted by a mobile robot.
- The development and demonstration of a contextual control architecture with enhanced teleoperation and autonomous behavior-based operational modes.
- The demonstration of a self-reliant robotic system with a 5-day unattended mission in a harsh environment.
- The demonstration of remote volcanology by providing video and gas-sensor data to scientists who could view, interact, and interpret exploratory results safely and instantly.
- The generation of widespread public and scientific interest in the mission, robotic exploration, and robotics in general.

This article fully describes the Dante II robot system, explains the important results, and then draws out the technical and programmatic lessons learned from the project. This article is organized into three major parts: Section 2 offers a technical description; Section 3 describes our technical results; and Section 4 relates the lessons learned. Section 2 begins with an overall description of the Dante II system, followed by a summary of the approach that guided our design and development process. A complete technical description of the Dante II robot system, including associated remote and support equipment, is then presented. Design rationale are given throughout the technical description. In Section 3, technical results from the extensive testing program and mission are described. Finally, Section 4 enumerates technical and programmatic lessons that the project can teach or demonstrate.

2. Technical Description

2.1. System Summary

Dante II, shown in Figure 2, is a framewalker; its eight pantographic legs are arranged in two groups of four, on inner and

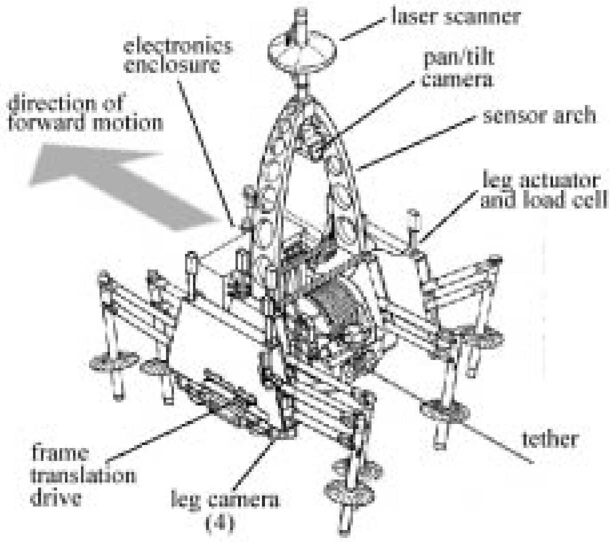

Fig. 2. The Dante II mechanism and sensors.

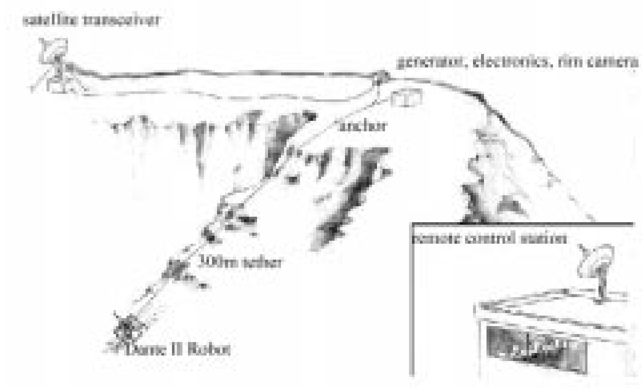

Fig. 3: Schematic of the equipment deployed to Mount Spurr. The remote-control station was 120 km distant.

outer frames. Each leg can individually adjust its position vertically to avoid obstacles and adapt to rough terrain. Body translation is actuated by a single drive train that moves the frames with respect to each other. The frames can turn relative to one another, to change heading. The maximum turn per step is 7.5°, so it is best to avoid obstacles in advance to minimize repeated turning.

Unlike most walking robots, Dante II pays out tether cable from an onboard winch as it advances. On slopes, tether tension is continuously adjusted to maintain stability and minimize adverse structural leg loading. Dante II is statically stable—it has no dynamic (balancing) phase in its gait cycle—but dynamic events, like bumps and slips, do occur and can destabilize the robot. To ascend and descend steep slopes, the tensioned tether, mounted on the inner frame, provides the reactive force to gravity. To walk, the legs on one frame raise up, while legs on the other frame support. The free (unsupporting) legs recover to new locations as the frame translates, propelling against the supporting legs. When the inner frame is in motion, the tether winch spools in relation to robot inclination to counteract the down-slope component of gravity, and to minimize shearing forces at the feet. The legs can lift and lower the body as well as adjust the body pitch to remain parallel with the nominal terrain slope. With this locomotion scheme, the robot can surmount large obstacles and traverse terrain in the full range from horizontal to vertical.

In addition to strength fibers, the rappelling cable also includes conductors for power, data, and video telemetry. Power is supplied from an offboard generator through the tether (see Fig. 3). A central electronics enclosure houses all equipment for real-time computing, actuation control, and communication. Navigation and control sensors include force transducers on each of the legs and tether winch, seven video cameras including a steerable stereo pair, and a two-axis laser scanner. In addition to the video cameras, which are used both for navigation and science, scientific data sensors include gas concentration sensors and temperature thermocouples.

Dante II is controlled through a variety of methods, ranging from direct teleoperation to supervisory control that enacts autonomous walking on rough terrain. In its autonomous mode, the robot uses control behaviors and data from force sensors on the legs and inclinometers in the body to generate motion sequences that advance the robot and safely adapt to the terrain.

Power-generating equipment, video and data telemetry equipment, as well as a crater overview camera, are all located offboard the robot near the tether-anchor point on the volcano rim. A satellite-band communications antenna at the anchor site sends video and data between the robot and the remote-control station. Operators in the remote-control station use graphical user interfaces to monitor and command the robot. Operators are also provided with a variety of video images and terrain maps built with the scanning laser. Key system specifications are given in Table 1.

2.2. The System-Design Approach

The Dante project required a new level of robotic competence to be developed, tested, and deployed in a 14-month period. Technologies previously demonstrated only in laboratory settings were to come together for a multiday mission in the harsh environment of an active volcano. From the outset, we were clear that success would hinge on adopting an approach that was tailored to the time constraints and technical challenges. Early in the project we detailed the measures of mission success:

- rappel to the floor of Mount Spurr's crater,
- procure scientific data from inside the volcano, and
- conduct the entire exploration without human presence.

These essential metrics of success became the standards against which we tested numerous technical, logistic, and

Table 1. Dante II System Specifications

Robot size (length × width × height)	3.7 m × 2.3 m × 3.7 m
Robot weight; payload	770 kg; 130 kg
Typical rappelling speed	1 cm/sec
Total tether length	300 m
Typical total system power draw during rappelling (including all telemetry equipment)	2,000 W
Typical bandwidth: control station to robot; robot to control station	56 kb/sec; 1 Mb/sec
Maximum rappelling slope angle; maximum cross-slope angle	90°; 30°
Maximum boulder or step on level terrain; on 30° slope	1.3 m; 1.0 m

programmatic questions and proposals throughout the project and actual mission. In general, anything that did not contribute directly to these essential metrics was set at a low priority or struck from the project. This approach was quite practical in countering the inevitable "mission creep" that increases technical risk, schedules, and budgets.

Another major aspect of our approach was to develop and test incrementally. In an effort to avoid late surprises in the fast-track program, we forced as much testing as early as possible. To ensure the appropriateness of tests, we simulated expected conditions whenever possible and attempted to push components to failure in order to understand failure modes and estimate safety factors.

Although the robot would be designed for autonomous operation, we decided to prioritize the development, testing, and hardening of the teleoperation system. Given the difficulty of the Mount Spurr terrain, we believed that even a rather capable autonomous control system would have to be assisted occasionally by human supervisors using teleoperation. While the various levels of autonomous controls were desirable from research and programmatic standpoints, they were not essential for mission success. Our baseline mission-readiness target was a solid teleoperation system. This approach would guarantee a robot that could traverse a variety of terrain challenges.

As the teleoperation capability was tested and proven, we began to layer autonomous modes of control on it. The "teleoperation-first" strategy led the program from easy to difficult: direct teleoperation uses simpler algorithms than more autonomous control schemes. Furthermore, autonomous robots tend to use the motion-control primitives that are the essence of a teleoperated device. It has been our experience that robot-development projects frequently prioritize development of autonomous control modes only to realize too late that robust teleoperation is needed and nontrivial to implement. With a solid teleoperation capability at the core of our system, we felt that it would be acceptable to experiment with less-proven autonomous modes up to and throughout the mission. We could always revert to teleoperation if necessary.

Another element to our approach was to overdesign components and subsystems, as this enables rapid integration and margin for unexpected conditions. We have found that robot-system integration commonly stresses components more than actual operating conditions. Rather than integrating electromechanics slowly in an attempt to avoid control "mistakes," we believe that it is more expedient to overdesign components and thus permit integration to proceed rapidly and with less concern for equipment damage. Of course, overdesign is no excuse to skip proper control design and testing. By nature, an exploration mission will encounter unexpected conditions. We believe that a systemic development approach that applies judicious overdesign will also result in a system more suited for an exploration mission.

A common result of overdesign is increased weight. In the case of a walking robot, increased body weight causes increased leg loads, leading to larger structures and actuators. We prioritized reliability and durability over weight minimization, although we did make systematic efforts to reduce weight.

2.3. The Rappelling Locomotion Mechanism

Dante II's locomotion mechanism is specialized for rappelling on steep slopes composed of variable materials, boulders, ditches, and ridges. The legs and tensioned tether are coordinated to advance the body and adjust elevation and attitude, while minimizing undesirable forces on the legs and maintaining a safe margin of stability (Apostolopoulos and Bares 1995). Because of the nature of rappelling, the vehicle is optimized to move in straight lines, typically down the slope's fall line (path of steepest descent). While on steep slopes, turning is used only to make minor heading changes, as the tether exit must remain nominally aligned with the anchor point to utilize the stabilizing effect of the cable.

The locomotion mechanism consists of legs, two frames, and the winch. Four legs are mounted on each frame. The legs individually lift and lower to compensate for terrain irregularities and adjust the body's height and pose. With four legs of a frame lifted off of the terrain, the frames move relative to each other to advance and turn the robot (see Fig. 4). While four legs of one frame support the robot, the other frame translates forward and places its legs on the ground. The winch is mounted on the inner frame and winds the tether in and out as needed to maintain a tension force on the robot. This section describes the legs, the body frames, and the rappelling winch system.

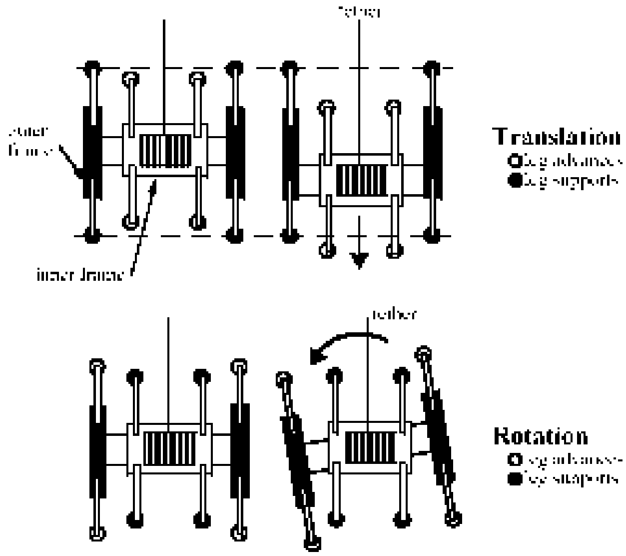

Fig. 4. Dante II's framewalking motions: translation and rotation.

2.3.1. *The Legs*

The design philosophy for the leg was to make a tough appendage that could withstand repeated falls and collisions with obstacles, and support its load on any type of terrain material ranging from mud to hard rock. Each leg is a pantograph mechanism that amplifies the motion of a linear drive actuator to generate a straight-line foot motion (see Fig. 5). The actuator moves the upper control point, while the lower control point is fixed to the robot's body. The linkage geometry amplifies actuator motion by four times at the foot, and amplifies axial foot forces by the same factor at the actuator. The actuator is compact and located far away from the foot and leg shank area, which greatly reduces the chance of damage during walking and rappelling. The pantograph also eliminates so-called geometric work energy losses during body lifting and lowering, the former of which is typically the peak energy demand for a slow-moving walking robot (Hirose 1983). The geometry of the moving and fixed control points results in minimal angular excursion of the shank; such excursions can cause leg entrapment, especially when lifting and lowering the body (Bares and Whittaker 1993).

The fundamental design criterion was to enable the legs to carry very high in-plane transverse (or shearing) foot forces, which could occur during rappelling, especially if the winch cable-tension control faltered or failed. The geometry of the leg is such that in-plane transverse foot forces result in amplified reactions at the two control points. The control points are very stiff and transfer all loads to the robot frame except the axial load borne by the actuator. A linear rail and sliding bearing connects the upper control point to the body. A wobble nut above the sliding bearing ensures that only axial loads

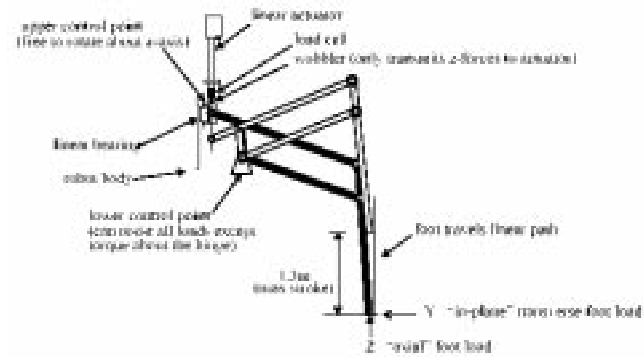

Fig. 5. The pantograph-style leg mechanism and actuator.

are transmitted to the actuator. The result is a compact control point and drive region without actuator side loads, making the actuator design simpler and more weight efficient.

The leg's linear actuator is designed to maximize the strength-to-weight ratio, and is based on electric drives used in the aircraft industry (see Table 2). A brushed 28-V DC motor drives a ball screw through a three-stage 60:1 gear reduction. Additional gearing beyond the load path drives a 10-turn absolute potentiometer and a small ball screw whose nut is used to trigger travel-limit switches. A power to open (fail-safe) brake on the motor locks the actuator whenever power is removed. This feature is used to save power and reduce actuator wear when the leg is supporting the body. Additionally, the robot can be locked in position merely by disconnecting the power supply. The actuator is sealed against water intrusion. One shortcoming with such a weight-optimized actuator is that the motor is only sized for intermittent high load use, and will overheat if used constantly for load holding. Exacerbating the situation is the lack of a simple and reliable method for measuring armature temperature on a small brushed motor. The result is that the actuator amplifier and control software system must protect the motor from overheating conditions that can occur during long periods of high or stall loads. Although the control software is designed so that the motor is never energized for long periods of high load, a custom current integration and cutout circuit on the amplifier provides an additional layer of motor protection.

The leg must withstand all cases of foot loading (axial to the foot, transverse to the pantograph plane, and transverse normal to the pantograph plane) that might occur throughout the spectrum from flat terrain walking to vertical slope rappelling. There was specific concern for minor falls during rappelling which might result in the entire robot's dynamic load acting on only a few feet, primarily in a shearing (transverse in-plane) direction. The pantograph leg structure is designed for maximum strength-to-weight performance; all links are thin-wall tubes with tapered wall thickness (to follow moment loading) where possible. The four principal hinge joints (see Fig. 6) experience a variety of forces and moments resulting from axial and transverse foot loads. These joints employ a

Table 2. Leg Actuator and Linkage Specifications

Actuator (integrated motor, brake, gearing, potentiometer, limit switch)	Sussex GB-106L
Continuous actuator-force output	11,564 N @ 1.4 cm/sec
Resulting force and speed at foot	2,891 N @ 5.6 cm/sec
Maximum speed at foot (no load)	7.9 cm/sec
Maximum body-lifting speed (assuming 8 legs supporting and lifting)	7.1 cm/sec
Actuator weight	5.4 kg
Actuator stroke	0.33 m
Leg weight (excluding actuator, guide bearing, pivot block, load cell)	16.1 kg
Leg-link material	Aluminum 6061-T6
Vertical leg stroke	1.31 m
Lower-shank pivot throughout vertical stroke	$< 5°$

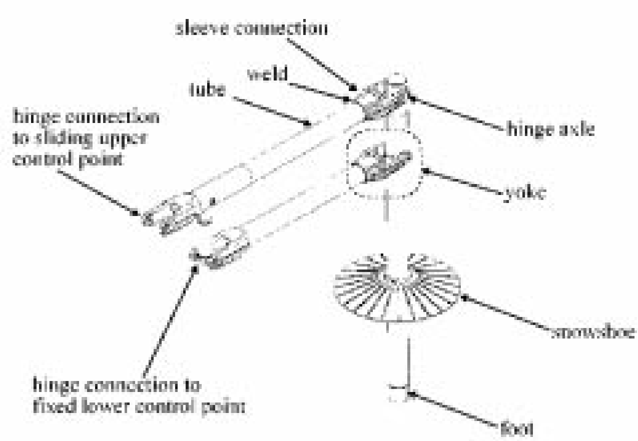

Fig. 6. The leg's structural details. All major connections are yoke and tube style, to minimize reliance on welds.

torsionally stiff yoke with a sleeve that transmits the loads directly to the tubes with little reliance on the sleeve/tube welds. Each yoke is machined from solid aluminum to minimize the need for welding and improve strength. The leg structure and specifically the yoke design was developed through a lengthy iterative cycle of design, finite-element analysis, fabrication, and destructive testing. Leg-load and stiffness test results are shown in Table 3.

A snowshoe rings the lower shank, and a hard rubber foot is appended to the bottom of the shank. Much like a ski pole, the hard foot is used when walking on solid surfaces while the shoe (like a ski-pole basket) supports the leg load on soft surfaces. The shoe is a multisegment design that can only support the leg load when most radial segments are sharing the load, as in the case of snow and mud. For instance, when walking among rocks, the segments readily deflect so that moment loads cannot be applied to the leg shank. The snowshoe also freely rotates on the leg to avoid jamming with terrain obstacles or other legs and snowshoes during forward frame translation and turning.

2.3.2. *The Frames*

The eight pantograph legs are arranged on the robot in two groups of four, on body sections termed the inner frame and outer frame. The inner frame supports the winch, electronics enclosure, and arched sensor mast. The frames can translate and rotate with respect to one another in a fully decoupled manner. For instance, to advance forward, the legs on one frame are raised, the frame is translated, and the legs are lowered to the ground. To continue advancing, the process repeats with the legs on the other frame lifting, followed by a frame translation (in which the single frame actuator drives back in the reverse direction). Turning is similar except that the frame with raised legs rotates with respect to the supporting frame instead of translating. Though the legs are attached to the frames in fixed rectangular constellations, each leg can individually move its foot up and down with its pantograph actuator. All of the legs can thus adjust lengths to best suit the terrain and objectives. For instance, if walking on a side slope, the legs adjust their lengths differentially so the body remains level. In another common situation, where the terrain slope changes, legs individually change length to adjust the robot's pitch as it crosses the region of slope change.

Decoupled translation and rotation of the inner and outer frames is achieved with the mechanism shown in Figure 7. The outer frame is composed of a transverse middle frame section, as well as left and right leg towers, each of which provide attachment for a pair of legs. The inner frame rotates on a bearing interface to the mid-frame. The turning actuator is attached to the inner frame, and drives an output pinon gear along a spur gear sector fixed to the mid-frame for a total turning motion of $\pm 11°$. Yaw position of the inner frame is measured by a potentiometer mounted within the central rotary bearing. The leg towers (on the outer frame) translate with respect to the mid-frame on bearing rails attached to the underside of the towers. A single actuator drives both towers simultaneously via a driveshaft that spans the mid-frame. The position of the towers with respect to the mid-frame is measured with a multiturn potentiometer attached to the final gear train. For the turning drive system, output

Table 3. Leg-Testing Results

Leg Position	x-Direction Loading F (N)	δ (cm)	k (N/cm)	y-Direction Loading F (N)	δ (cm)	k (N/cm)	z-Direction Loading F (N)	δ (cm)	k (N/cm)
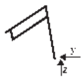	1,886	10.6	178	3,159	5.3	596	No test	No test	No test
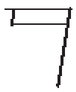	1,935	13.6	142	3,186	5.6	568	4,445	3.7	1,201
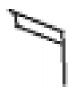	1,917	16.7	115	3,204	5.6	572	5,050	3.5	1,451

Loads shown are peak successful loading; in some cases, loads were increased beyond the values shown, but resulted in link or joint failure.

torque is quite high (1,618 Nm) to provide overcapacity in the event that lateral forces from the tether place a counter-rotation moment on the vehicle during a turning motion in the opposite direction. In the case of the translation drive system, a worst case of propulsion up a 30° slope without any assistance from the tether system was assumed. This could occur, for instance, if the robot were climbing a small incline at the crater floor and could not rely on propulsive assistance from the tethering system. Specifications for the turning and translation actuators are given in Table 4.

2.3.3. The Tethering System

Dante II's tethering system enables it to traverse steep and nearly vertical terrain by reducing tractive forces at the feet and providing stability. As shown in Figure 8, tension from the tether cable is used to assist the robot in climbing by lifting much of the weight directly with the cable, instead of through transverse shearing forces at the feet (Krishna, Bares, and Mutschler 1997). Without assistance from the tether, the maximum transverse load possible at any foot is limited by either the tractive characteristics of the terrain material, or the transverse in-plane leg strength. On loose material such as that found on many volcanic slopes, a tetherless walking robot would thus be limited to mild inclines. Tether tension also provides stability on steep slopes (also shown in Fig. 8) where the robot would tip forward without the tether force. Note that the tether force also reduces the transverse in-plane foot forces (H_1 and H_2) which lowers the need for leg structure and thus lowers weight. Force components of the tether tension can also have destabilizing effects, which commonly occur when the robot turns its longitudinal axis away from the departing path of the cable, thus creating a lateral "restoring" force at the tether exit point. (A lateral force on the

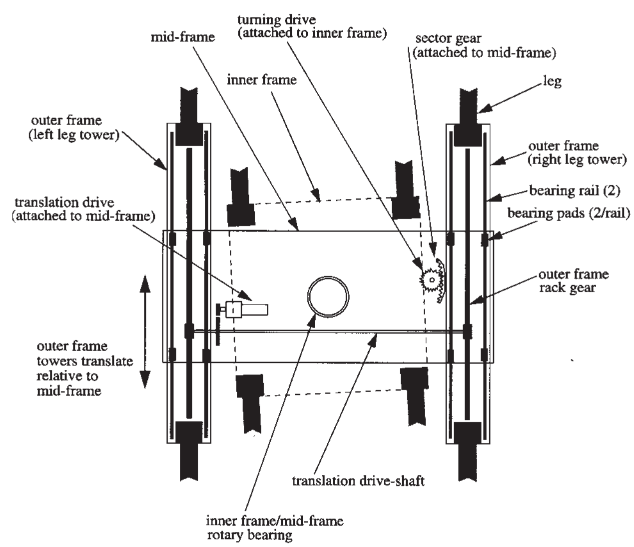

Fig. 7. Robot frames, translation, and turning mechanisms (plan view). The robot body consists of three "frames": an outer frame comprised of two leg towers, a mid-frame that houses the translation drive, and an inner frame that houses the turning drive as well as the winch and electronics enclosure.

Table 4. Translation and Turning Actuator Specifications

	Turning Actuator	Translation Actuator
Motor	Pacific Scientific R35	Pacific Scientific R35
Maximum continuous output torque	2.85 Nm @ 2,000 rpm	2.85 Nm @ 2,000 rpm
Brake	Inertia Dynamics 1905-2521 fail safe (i.e., power to hold open)	Inertia Dynamics 1905-2521 fail-safe (i.e., power to hold open)
Gearhead at motor	Micron HPI-75-50 planetary, 50:1	Micron HPI-60-82 planetary, 8.2:1
Final gearing	Pinon and spur gear, 32.8:1	Pinon and multiple spur gear, 8:1
Total gear reduction	1640:1	65.6:1
Maximum continuous output torque/force at frame	1,618 Nm (assuming efficiency of 85% for planetary and 95% for single spur)	3,765 N (assuming efficiency of 90% for planetary and 86% for multiple spur)
Absolute position	ETI Systems SOP30B 1-turn	ETI Systems MH20S 10-turn

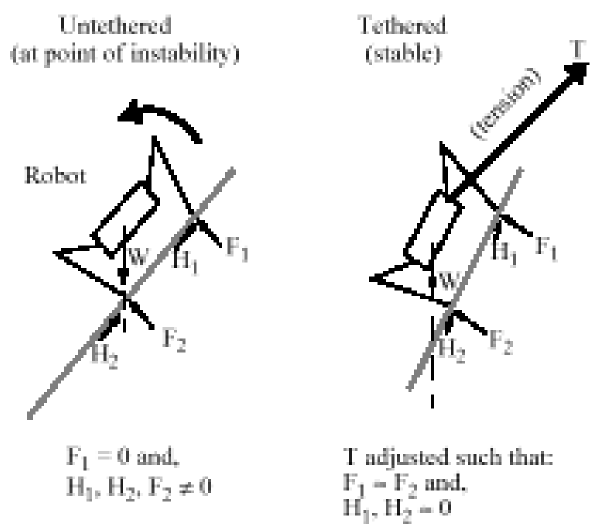

Fig. 8. The tether stabilizes the robot and equalizes leg loads on steep terrain.

robot from the tether contributed to Dante II tipping over as it climbed out of Mount Spurr.) A less common destabilizing effect occurs when the terrain geometry is such that an upward cable-tension component reduces the normal force on the uphill legs.

Vehicle pitch relative to gravity, as well as tether exit angle relative to the robot, are measured with sensors, and a desired tether tension is computed to eliminate the transverse in-plane shearing forces on the feet and maintain robot stability. The winch controller uses a hybrid position/force feedback algorithm to maintain the target tension on the cable at all times. This approach simplifies both teleoperated and autonomous operation, because no operator-generated commands are required. As the robot translates or changes pose, the tether is automatically wound or unwound as needed to maintain the desired tension.

In addition to Kevlar fibers that are used for strength, the tether cable includes conductors for power, data, and video (see Fig. 9 and Table 5). The custom-made lightweight tether was developed and tested iteratively using a battery of temperature, abrasion, and flexibility tests (Krishna, Bares, and Mutschler 1997). Most notably, the cable was designed and tested to withstand high tension while crossing a sharp corner such as a rock edge (kinking).

The tether is wound on a winch that is located on the inner frame and is driven by a brushless motor via a spur gearbox and chain drive. A multiturn potentiometer geared to the output of the spur gear head measures cable pay-out. From the winch drum the tether cable passes over a sheeve on a level-wind mechanism then under the winch drum, and finally exits the vehicle through an exit cone (see Fig. 10). The level-wind mechanism is a reversing ball screw that is timed via a chain to lay the cable on the drum, one layer at a time. The level-wind is placed on the opposite side of the drum as the exit cone to improve compactness of the overall winch system. At the center of the drum the tether cable passes to the inside of the drum, and into a rotary slip ring.

The winch and exit cone are mounted to a rigid frame that is suspended from four small cantilever beams, or *flexures*. The flexures are stiff in all directions except along the robot's lon-

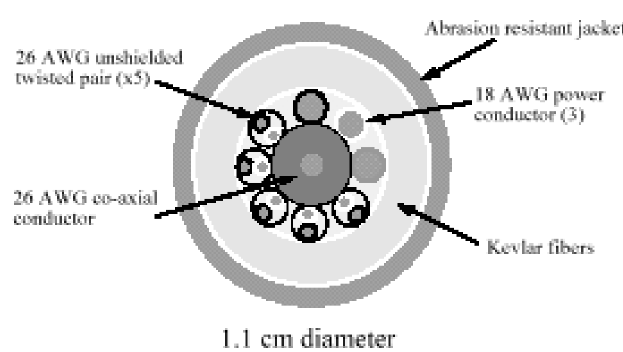

Fig. 9. Composite tether cross-section.

Table 5. Tether System Specifications

Winch Manufacturer	InterOcean Systems, Inc.
Motor	Pacific Scientific R35
Maximum continuous output torque	2.85 Nm @ 2,000 rpm
Brake (fail-safe, i.e., power to open)	Inertia Dynamics 1905-2521
Gear head at motor	Sussex Gear four-stage spur; 178:1
Final gearing	Chain drive; 7:1
Total gear reduction	1246:1
Maximum line-pull test with one layer of cable on drum	17,800 N
Maximum no-load line speed	3.5 rpm
Drum diameter	32.3 cm
Overall winch dimensions	1.0 m × 0.75 m × 0.75 m
Slip-ring unit (24 rings, 17 used)	IEL-BX-24-MOD
Tether-cable tension sensing	Two T-Hydronics TH-DL5M (SPL) in series mounted to suspended winch platform
Winch weight (including actuator, excluding cable)	114 kg
Winch weight (including actuator and cable)	169 kg
Cable manufacture	Gore Manufacturing
Length	300 m
Maximum static load; dynamic load; ultimate load	7,700 N; 13,600 N; ∼44,000 N
Weight	0.186 kg/m

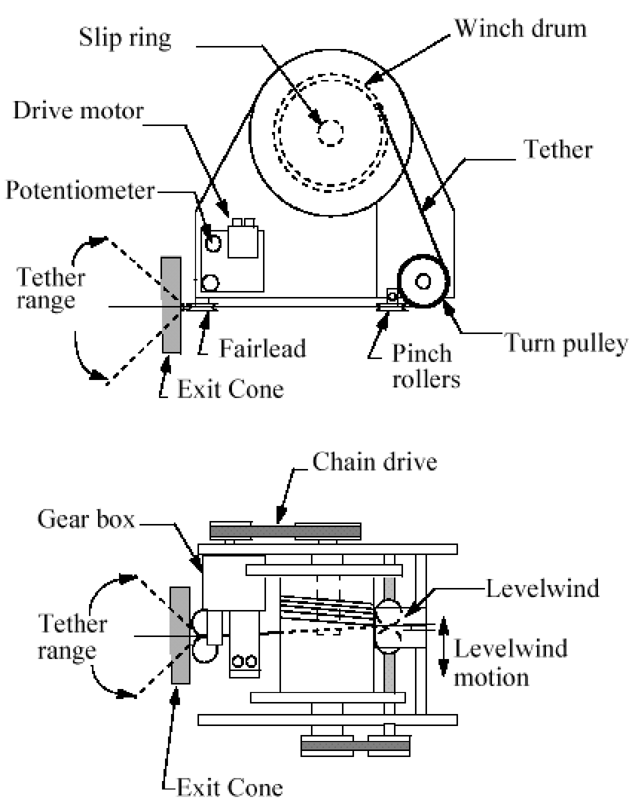

Fig. 10. Tether winch with level-wind mechanism (side and top views).

gitudinal axis. An axial load cell restrains the winch frame and measures the tether force in the longitudinal direction, regardless of the tether angle relative to the exit cone. The flexure/load-cell system can accurately measure the longitudinal tether-tension component to within 22 N. The flexures proved to be an excellent means of suspension for force measurement, as they have no hysteresis, deadband, or stiction, and are generally unaffected by weather conditions. Two hoops with potentiometers affixed to the exit cone measure the lateral and vertical exit angle of the tether relative to the winch frame. Given the two exit angles and the longitudinal tension component, the tension can be calculated.

2.4. Sensors

Dante II has a suite of navigation and scientific sensors, including inclinometers in the body, force sensors on the legs, video cameras, a scanning laser, and ambient-gas-concentration sensors. Terrain topography—crucial to planning and executing rappelling—is sensed with perceptive-imaging devices as well as the leg position and force sensors. Total distance traveled by the robot is measured by tether payout (though some error occurs when the tether assumes a more direct path than the vehicle has actually traveled).

2.4.1. Force and Position Sensors

Each leg is instrumented with an axial load cell mounted between the vertical actuator and the pantograph mechanism to measure vertical foot force, and a pair of differential strain gauges on the upper section of the leg shank to measure in-plane transverse forces on the shank (including those transmitted via the foot). The strain gauges detect continuous high loads or transient bumps of small magnitude (approximately 1 N). Potentiometers encode absolute joint positions, and inclinometers measure gravity-relative posture. These sensors characterize the positions of and forces on the 11 actuated motions and the posture of the body.

The load cell at the pantograph-drive actuator of each leg measures axial foot loads (the actual foot load is one-quarter the load-cell reading). The load cells are crucial for smooth locomotion: as recovering feet are placed on the terrain, the load-cell signals are used to attain a preload onto the legs (and soil) before locking each actuator brake. In this manner, when the supporting feet are lifted, a minimal amount of body "sagging" occurs as the full load is transferred onto the other legs. The vehicle controller continues to monitor the foot loads, and readjusts the load apportioning as needed to distribute foot loads.

The strain gauges affixed to the leg shank measure bending in the plane of the pantograph. Although located high on the leg shank, the strain gauges are susceptible to damage from scraping and bumping large boulders. The strain gauges are primarily used to detect collisions of the leg shank with obstacles during translation, which produces bending moments. When the leg is supporting, shank bending can also arise from improper tether control, which places transverse forces on the foot. Excessive bending in any number of legs may indicate failure in tether-tension control. In the event of failure of the tether-tension sensor, the shank-bending sensors could be used as a coarse means of inferring tether tension.

2.4.2. Video Cameras

Dante II is equipped with a total of seven video cameras for navigation, teleoperation, and scientific observation. One camera is positioned at each corner of the outer frame, viewing the pair of legs on the opposite corner of the robot. The remaining three are in an enclosure on a pan/tilt mechanism mounted to the sensor arch. Two of these are used to provide stereo images, and are separated by a distance equal to the human ocular baseline. The third is equipped with a remote-controlled zoom/focus lens for detailed viewing of areas of interest. A windshield wiper keeps the front window of the enclosure free of dirt and water. For the Mount Spurr mission, imagery from the video cameras provided the principal source of scientific information used for volcanology.

2.4.3. Scanning-Laser Rangefinder

Geometric terrain information is measured by a spherical scanning-laser rangefinder mounted atop the sensor arch. A custom two-axis scanner sweeps a distance-measuring laser beam around the vehicle in a concentric pattern (see Fig. 11) with a 350° × 45° (vertical) field of view. Elevation maps generated from the range data are used for teleoperation interfaces and planning. Each elevation map covers a circular area with diameter of 15 m, with an average grid resolution of 10 cm. A single scan requires about 24 sec to gather 20,000 range readings (constrained by mirror speed, not laser-ranging rate).

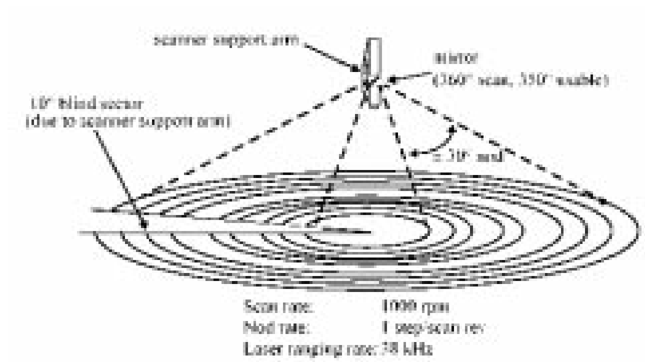

Fig. 11. The two-axis range scanner: scanning pattern and field of view.

The scanner consists of a fast scan (or slew) axis driven by a DC motor and a slow nodding axis driven by a stepper motor (see Fig. 12). The nodding stepper increments the mirror-nod angle once for each scan-axis mirror rotation. The scan motor and coupled encoder are hollow core to permit passage of the driveshaft from the nodding-motor speed reducer to the nodding assembly. The slip ring permits unrestrained rotation of the nodding-motor and speed-reducer assembly. The rangefinder (Acuity 300) has a 10-mW, 780-nm laser with a sample acquisition rate of 38 kHz. The transmitter and receiver are coaxial, with the receiver having an aperture diameter of 7.6 cm. The laser beam is initially 3 mm in diameter, and has very low divergence. Maximum range on poorly reflective targets such as wet mud is limited to about 7 m. A further limitation of the laser is that it cannot penetrate fog or steam, a common condition in volcanos. Although a window enclosing the spinning mirror assembly would offer protection, we could not find a method to keep a window clean enough so as not to cause internal laser reflections. As a result, cleanliness of the spinning mirror is a major weakness of the system. (Late in the Mount Spurr mission, a fine film of volcanic ash settled on the mirror and rendered the laser scanner inoperable because of transmitted laser reflection back to the receiver.)

2.4.4. Science Sensors

Science sensors provide temperature and gas composition information: a thermocouple near the foot on one of the downhill legs can be positioned over any area of interest. Underbody sulfur dioxide and hydrogen sulfide sensors provided real-time gas-concentration measurements. Science sensor data is sampled continuously, and correlated with time and location.

2.5. Actuation Electronics and Computing

A layout of Dante II's actuation electronics and power distribution is shown in Figure 13. A 240-V AC generator stationed

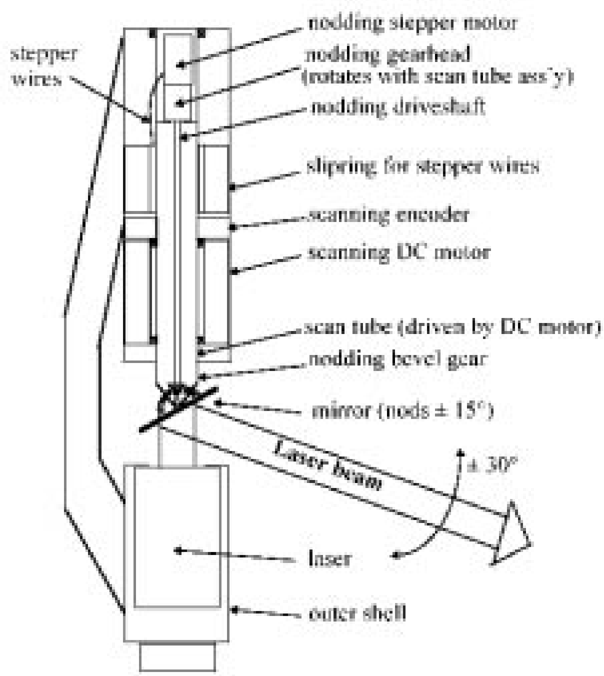

Fig. 12. The two-axis range scanner's internal mechanism.

at the rim of the volcano powers the robot. Power is stepped up to 1,000 V AC to reduce conductor size in the tether. A toroidal step-down transformer onboard the robot produces the various voltages required by the electronic system, and tends to smooth any voltage ripple caused by the generator.

The onboard computing hardware consists of three Motorola 68030 boards, one Sparc2 board (for perception processing), two Creonics motion-control boards, a DAADIO board, and several custom PCBs (for power, brakes, etc.), all mounted in a 6U VME backplane, shown in Figure 14 (Boehmke and Bares 1995). Although the custom boards do not use the backplane bus, mounting them in the cage provides security and easy access. The interface boards condition all signal input and output, provide safety monitoring of onboard systems, and provide fail-safe brake actuation. The robot controller, running on three 68030 processor boards, digitally filters all sensor data before supplying it to other processes via shared memory.

Rain, snow, steam, acidic gases, extreme sunlight, and ash are common in the volcanic craters. To combat these hazards, all onboard electronics are centrally located in a single, shock-mounted, removable NEMA4X-rated enclosure. All interconnections to the enclosure use waterproof connectors and Teflon insulated wire for sustained flexibility in cold temperatures. All connections are made on the rear door of the enclosure, which can be hinged open for access to wiring on the rear of internal components. The front door is easily removable for access to the front of components including the VME cage. Since most low-cost, high-performance electronic and computing components are typically specified to function from 0–50°C, a cold-start sequence warms the enclosure before powering computers and associated equipment. Two heat-pipe-based passive-cooling devices installed on top of the electronics enclosure dissipate up to 900 W of heat when the internal enclosure temperature exceeds 25°C (as occurred during summer testing in Pittsburgh). The components are positioned so that the air circulation forces the coolest air over the most sensitive computing components first and the least sensitive components (like transformers) last, before reaching the heat exchanger again.

2.6. Communication Electronics and Telemetry

The Dante II telemetry system transmits data and video between the robot, rim station, and remote-control station. The major components of the telemetry system are diagrammed in Figures 15, 16, and 17. The satellite uplink is 192 kB, with a round-trip delay of about 4 sec. This is sufficient for monitoring robot state, although transmission of large data packets and network anomalies can cause delays of 30 sec or more. As with most remotely controlled systems, telemetry constraints encourage minimizing communication and maximizing onboard self-reliance.

The primary data path between the robot and the rim station is an Ethernet link through a pair of transceivers running over a single unshielded twisted pair in the tether. Serial line drivers provide full-duplex RS-232 communications to each CPU and to video quad splitters. A 9,600-baud spread-spectrum radio modem serves as a backup data path between the robot and the rim station in the event of Ethernet telemetry failure.

Video telemetry is fundamentally limited to two compressed video feeds by the cost of satellite bandwidth between the rim station and the remote-control station. Onboard the robot (Fig. 15), the quad splitters each receive all seven video streams, and can select any signal for a full screen view, or a split screen of any four camera views. Output from each quad splitter is transmitted to the rim station via wireless video links at 2.0 GHz and 915 Mhz. The intended primary video path was originally a mini coax in the tether instead of the wireless links, which were to serve only as backups. However, complications with the extremely small coax in the tether yielded it unusable, forcing full reliance on the wireless video links.

An intermediate telemetry station, diagrammed in Figure 16, enables communication between the robot and the distant control station. Two 384 kb/sec video-teleconferencing CODECs digitize and compress the video streams. One receives a direct stream from one wireless video link, while the second is switched via computer control between the signal from the other wireless video link, a rim-observation camera, and the stereo cameras. When desired, a stereo view is formed by transmitting the two images over the wireless links from the robot to the rim station, where they are merged into a single interleaved image by a stereo encoder box. A terminal server connected to the Ethernet enables operators to directly control attached devices. The Ethernet and video-data streams

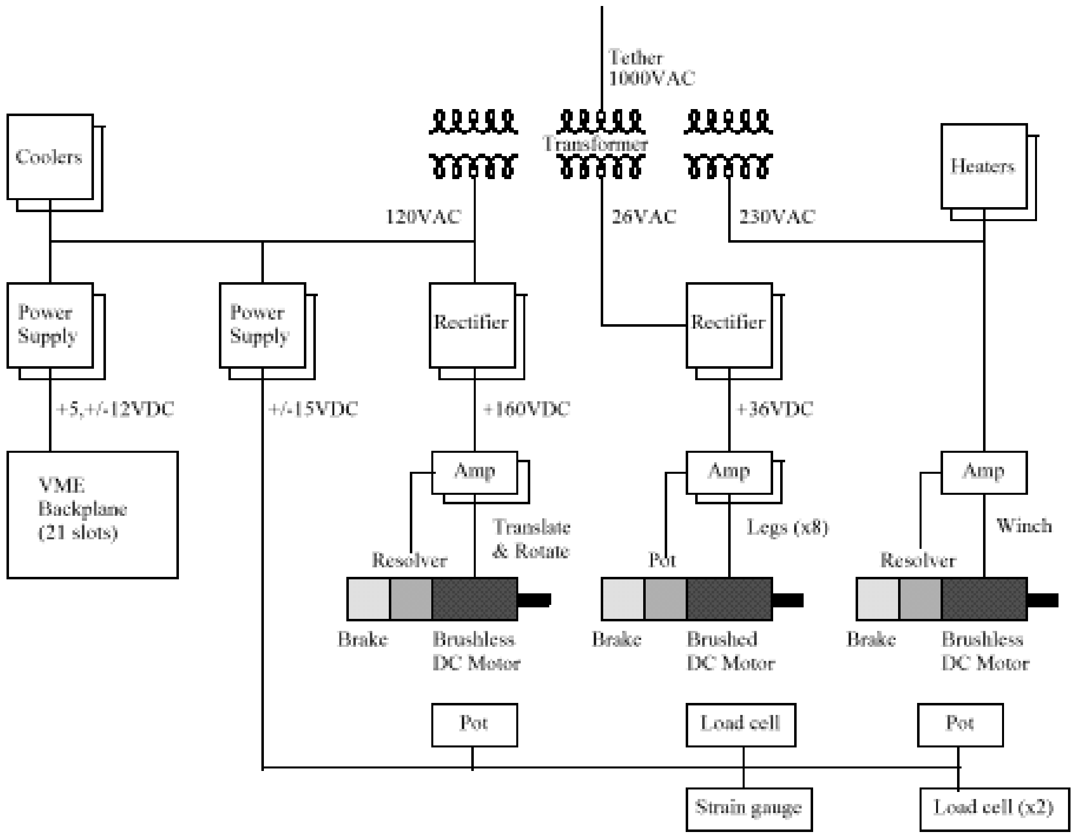

Fig. 13. Actuation and power-distribution electronics.

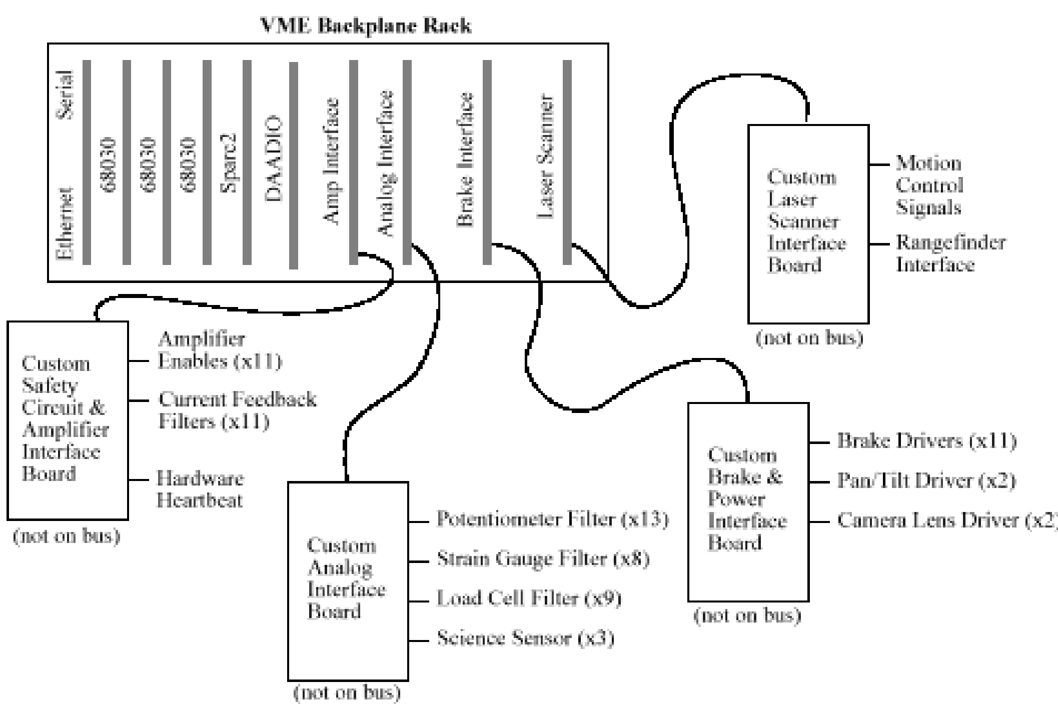

Fig. 14. Computing and electronics backplane cards.

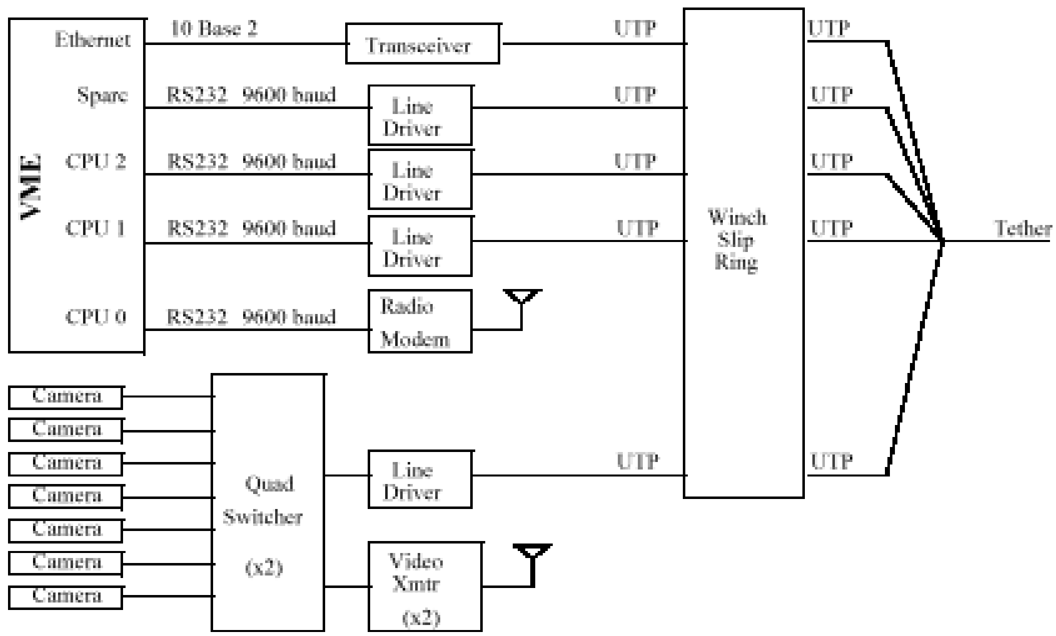

Fig. 15. Robot onboard telemetry components.

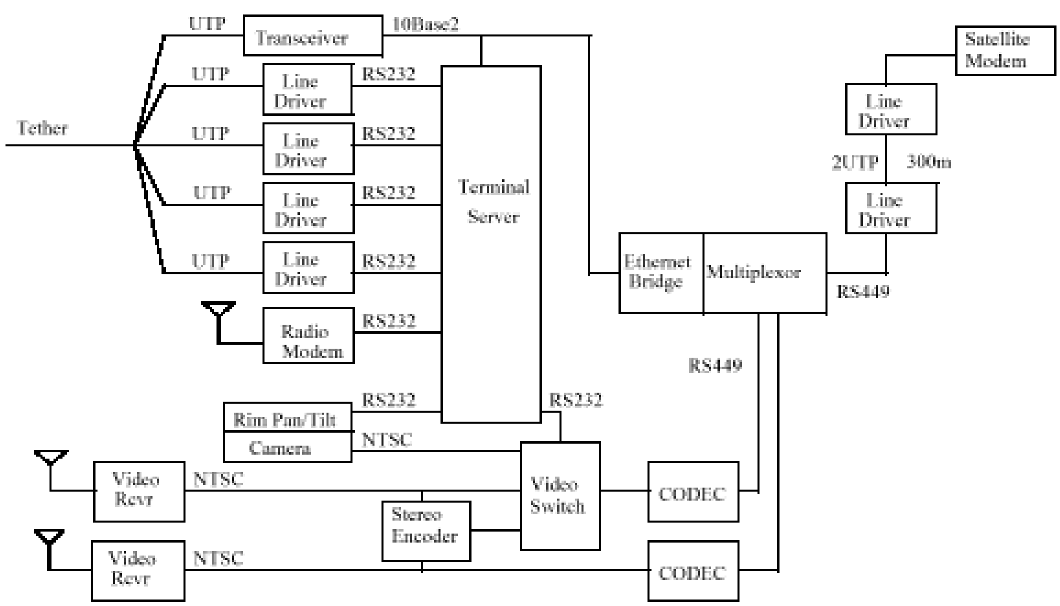

Fig. 16. Rim station telemetry components.

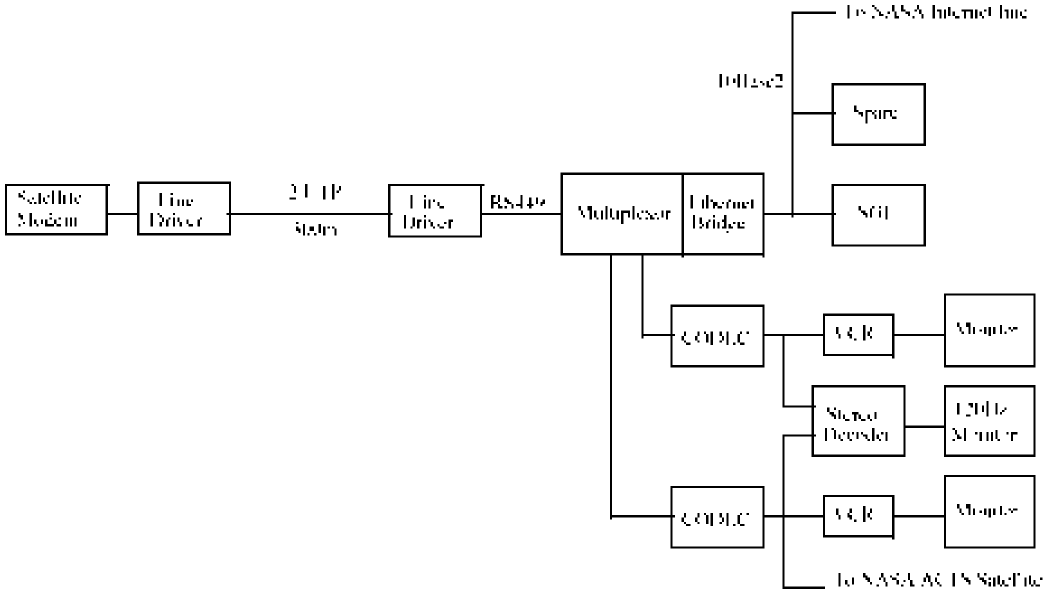

Fig. 17. Remote-control station telemetry components.

are fed into a digital multiplexor equipped with an Ethernet bridge and a high-speed serial card. The 1-Mb/sec composite RS449 telemetry signal is sent 300 m to the satellite modem over a pair of line drivers. (The modem, up/down converters, 8-W amplifier, and 2.4-m satellite dish were leased for the duration of the Mount Spurr mission.) All telemetry and power equipment located at the rim station is enclosed in waterproof cases with internally shock-mounted racks. Power is supplied from the same generator that powers the robot. An uninterruptable power supply provides power backup for all rim equipment as a means to enable limited remote troubleshooting in the event of a generator failure.

The remote-control station telemetry is shown in Figure 17. An existing 8.1-m satellite dish mounted on top of an office building in Anchorage served as the downlink for the telemetry stream. Inside the control station, two color monitors and video recorders displayed and recorded the video streams which were first decompressed by two decoding CODECs. One video stream was also routed through a stereo decoder to a 120-Hz stereo display monitor, which gave operators wearing LCD-shuttered glasses a sense of depth perception when viewing images from the stereo camera pair.

2.7. The Operator Interface

We designed Dante II to be operated in control modes ranging from direct teleoperation of individual actuators to fully autonomous control of walking. The typical human-machine interaction with Dante II is characterized by supervisory control. Our reasons for base-lining supervisory control were 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 delays between the operator and the remote system. In the strictest sense, supervisory control means that when the operator relinquishes control, the robot is allowed to function autonomously. Our intent, however, is to be more flexible and to let operators intermittently assume direct operation (traded control) or to supervise certain variables while directly handling others (shared control). This approach affords greater human-machine synergy, benefiting and improving overall system performance.

To support the traded and shared control of Dante II, we recognized that operators needed interfaces to visualize and understand system state, as well as to specify and use appropriate control modes. Because of 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 Mount Spurr expedition was intended not only to advance robotic exploration technology, but to also involve numerous remote mission observers and participants. We were interested in conducting collaborative research with remotely located volcanologists and planetary scientists. Thus, we provide observer interfaces which enable field science and allow distant researchers to directly participate in the mission. Additionally, we provided 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.

2.7.1. The Graphical Interface

Our aim in creating the graphical user interface was to make it as easy as possible to operate Dante II. For a human operator, trying to coordinate all of Dante II's degrees of freedom to make it walk requires concentration and stamina. We were motivated both to minimize operator workload and to make it possible for novices to quickly learn to control the robot. Dante II's graphical user interface (UI2D) enables shared and traded control to monitor sensor and robot-state information, and to generate commands to the robot. The interface presents graphical buttons, dials, slider bars, and numerical displays in a collection of display windows. The UI2D's main window is shown in Figure 18.

During the course of a robotic exploration, different situations are encountered and a variety of actions must be performed. To control the robot differently in these various situations, operators need to be able to conveniently access related commands. We organized UI2D so that commands appropriate for a particular type of function or operation are 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. These contexts are *individual actuator*, *frame*, *behavior*, *gait*, and *path*. Each context, other than individual actuator, adds elements of autonomy, as shown in Table 6. 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.

2.7.1.1. The Individual Actuator Context

In the individual actuator context, the operator can teleoperate each actuator with no command safeguarding other than checks on the limits of motion. This 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 context may be applied occasionally during exploration; for example, when a single leg hits a rock during translation, the operator may wish to raise just that leg and then switch contexts to resume walking.

2.7.1.2. The Frame Context

In the frame context, actuators can be commanded in groups to allow walking by raising, lowering, stroking, and turning frames of legs. The frame context enables teleoperation of Dante II as a frame-walking system. The key automatic feature is the coordination of tether pay-out 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, and the tether tension is automatically controlled to the proper value.

Figure 19 shows the frame-context window, which is activated from the UI2D main window. Frame lift-and-place buttons command groups of legs, and numerical fields allow precise specification of desired motions. The translate and rotate arrow buttons move to specified numerical values. Similarly, the body posture and elevation can be adjusted with a mouse click.

2.7.1.3. The Behavior Context

In the behavior context, the operator specifies the set of parameters that describe the desired gait for the robot. The interface provides a window for entering these parameters and then downloading them to the robot, along with a "go" command. The gait is commanded to the robot, which executes it autonomously until told to stop or until an insurmountable obstacle is encountered. Perceptive sensors are disabled; only proprioceptive sensors are used to detect contacts and adjust posture. We considered this to be supervisory control, as the robot responds automatically to the terrain as it senses it, but is guided by the parameters that the operator provides.

2.7.1.4. The Gait Context

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 the behavior context, 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. Only preliminary implementation of the gait context for Dante II was completed and tested with the laser rangefinder prior to departing for the Mount Spurr mission. Final implementation and testing occurred in simulation after the exploration of Mount Spurr was completed.

2.7.1.5. The Path Context

The path context is envisioned as the mode 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. The path context was designed, but not implemented, as the last step to full autonomy for Dante II.

2.7.2. Terrain Modeling and Visualization

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

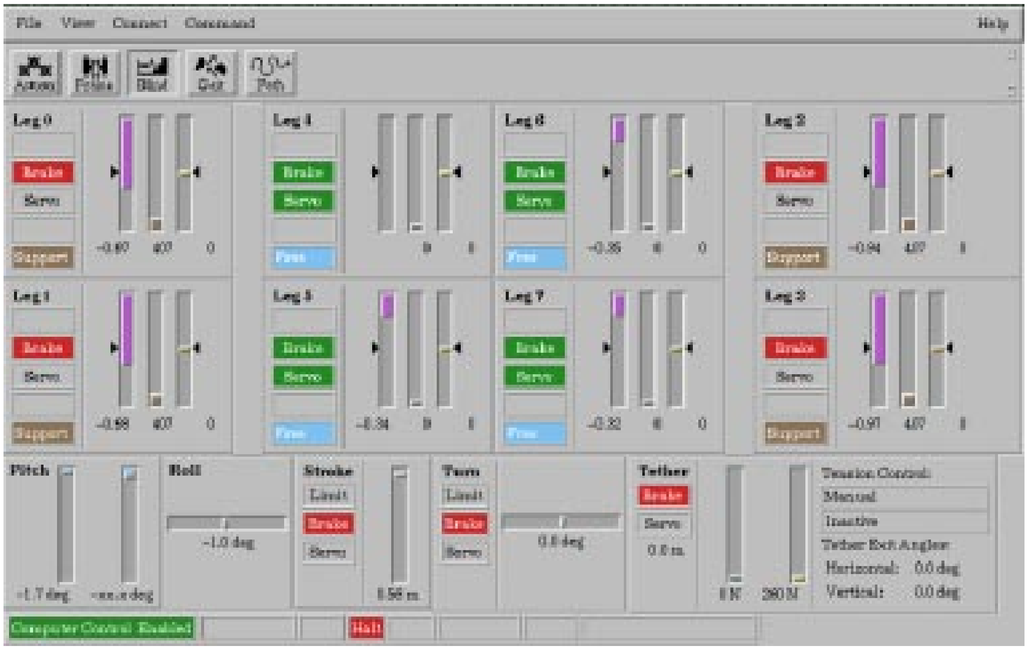

Fig. 18. The main display window of the graphical user interface, UI2D, showing current robot-state information.

Table 6. Operational Control Contexts and Elements of Autonomy

Control Function	Indiv. Actuator Context	Frame Context	Behavior Context	Gait Context	Path Context
Servo tether					
Coordinate leg motions					
Maintain body height					
Maintain posture					
Adjust leg step					
Surmount obstacles					
Perceive terrain					
Determine step height					
Determine body height		Operator Control		Automatic Control	
Determine posture					
Determine stride					
Correct leg placement					
Generate path					
Determine heading					
Avoid obstacles					

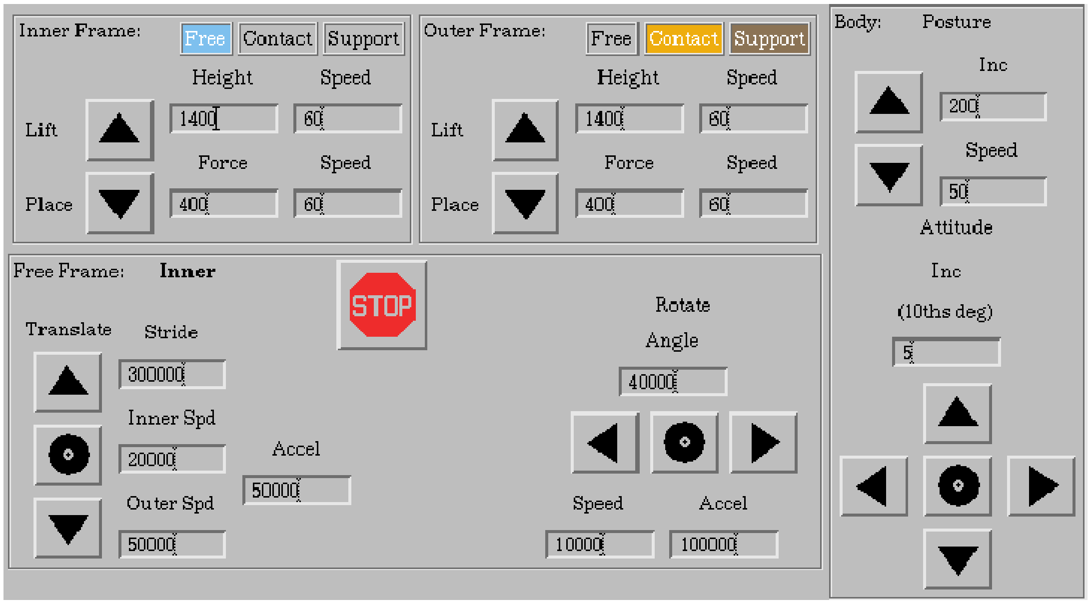

Fig. 19. The graphical user interface frame-context window with controls for commanding groups of legs to allow direct teleoperation of walking.

suboptimal constraints such as low bandwidth and delays to the remotely operated system.

To improve our ability to control Dante II and make correct decisions about where to send the robot, we developed UI3D, a three-dimensional visualization of Dante II and its surrounding terrain (Fong et al. 1995). The scanning-laser rangefinder measures the distance to points in the terrain, and this data together with robot-position information is used to generate local elevation maps. The UI3D incorporates these capabilities for terrain modeling and robot localization with a virtual environment package called VEVI (Virtual Environment Vehicle Interface). The VEVI is a modular operator interface for robotic vehicles developed by the Intelligent Mechanisms group at NASA Ames Research Center (Hine et al. 1995). It is a tool for planning and previewing high-level task sequences, monitoring system state, and analyzing anomalous events. Through its virtual environment, the VEVI provides operators with spatial orientation and perspective superior to conventional control stations. The VEVI utilizes real-time, interactive, 3-D graphics and feedback from onboard sensors to update the simulated vehicle and its environment. 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 understanding its state. When Dante II is operating in unstructured and unknown environments, terrain features such as steep embankments and obstacle-strewn fields further compound the difficulty of comprehending the system state.

The UI3D enables operators and observers to easily envision Dante II's configuration and its surroundings, as shown in Figure 20, where Dante II is descending a steep slope.

2.8. Motion Planning, Behavior-Based Execution, and Real-Time Control

Making a robot walk poses the implicit problem of generating a coordinated sequence of leg and body motions—a gait. We have been convinced, by our experiences with previous walking robots, that gait planning alone, without adaptation during execution, is insufficient for guiding a walking robot in natural terrain. Too much is unpredictable; bumping into obstacles or slipping off precarious footholds cannot always be foreseen while planning. Behavior-based architectures address this need to adapt to actual conditions by continuously relating sensation to action (Brooks 1986). Sense-act mappings establish planned reactions to expected, but unpredictable, events. Biological systems also provide evidence for simple sense-act reflexes, as well as decentralized control in walking (Grillner 1985). Both systems exhibit properties necessary for gait execution: reaction to unexpected events, concurrency of reflexes, and coordination of actions. For Dante II, we took a behavior-based approach to gait execution to benefit from relating sensation directly to action in simple reflexes, and have developed a network of coordinated behaviors to stand, posture, step, and walk (Wettergreen 1995).

To guide Dante II's actions, each behavior is parameterized (for example, the height to raise a leg or the stroke to propel the body) in direct relation to physical actions. The sensitivity (such as to contact forces, and to pitch and roll errors) can also be quantified and adjusted. Motion planning for Dante II is thus simplified to determining the most effective (fastest, most stable) parameterization of the gait behaviors. Because many parameters are independent (for example, leg-lift height is

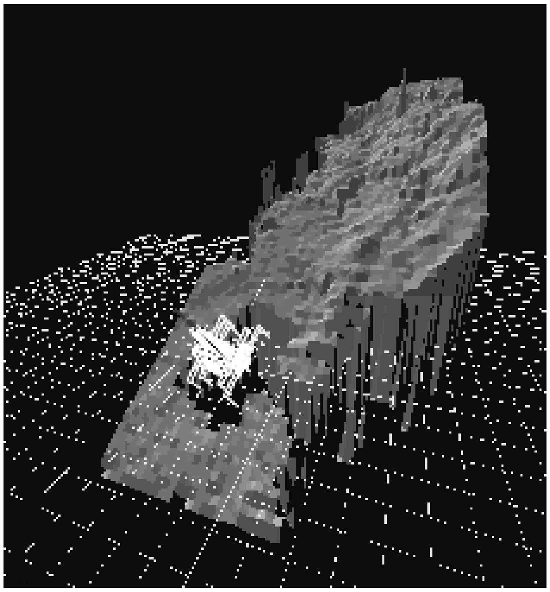

Fig. 20. A virtual-environment visualization of Dante II in the terrain (the grid denotes the horizontal reference plane). As the robot progresses down the slope, multiple laser scans are merged to construct an elevation map of the terrain's "history." Note that there is no terrain data under or immediately around the robot: this is because the scanner is blind in these areas. However, as the robot continues to move forward, these unknown areas will be rescanned and filled into the terrain model. Operators use this model to "fly around" the robot, and thereby build a complete understanding of the surrounding terrain conditions. In addition, the robot's legs as well as the tether change colors as a function of force loading, giving the operator a quick sense of the vehicle's force state. (Note: the robot in the image was retouched for clarity.)

related to terrain elevation and little else), independent planning modules can be applied to set specific parameters. A collection of planners can guide the robot quickly and safely. By manually setting the parameters, a human can guide the robot, directing its overall performance: this forms the basis for supervisory control of Dante II.

2.8.1. Behavior-Based Gait Execution

Asynchronous control processes, gait behaviors, act independently to achieve or maintain desired states, and interact to walk. These behaviors are parameterized to allow external modification and direction. A contact foot behavior (for each of the eight legs) causes the foot to maintain contact with the terrain, and acts to lower the foot to the ground whenever a vertical force is not sensed. When exhibiting a lowering of a foot, the contact-foot process inhibits the movement of the robot body. The emergent behavior is a reflex that returns the foot to the ground if it loses contact, and a coordination with body motions to halt motion until contact (and stability) is reinstated.

Conversely, a free-foot behavior causes a foot to stay free, out of contact from the terrain. When exhibiting and detecting either vertical or horizontal terrain contact, the free-foot behavior raises the leg. This initiates the stepping motion and also creates a reflex, when a leg bump occurs, that causes legs to raise up in coordination with a momentary pause in body motion. Figure 21 shows the extensions of all eight legs (inner-frame legs are dotted) during a portion of the Mount Spurr descent. At minute 118.5, a raised leg bumped the terrain and continued to raise up. The reaction time was less than 0.5 sec.

Five behaviors—raise legs, move frame, turn frame, lower legs, and sit still—when sequenced together, enable walking. The raise-legs behavior coordinates the lift of a group of legs. It sends a signal to a set of four free-foot processes that it maintains until all four have raised. It then sends simultaneous exhibition signals to the move-frame and turn-frame behaviors. They then signal the lower-legs behavior, which in turn signals the sit-still behavior (if image and data capture is needed), and completing the cycle, the raise-legs behavior again.

To correct for rolling terrain, the roll behavior adjusts robot posture about the longitudinal axis. Typically, roll is minimized to maximize stability, although in some situations it is reasonable to lean to one side. A coordinated motion of all legs—some raising, some lowering—rolls the robot to the correct value.

On level terrain, the pitch behavior can function identically to the roll behavior: monitoring an inclinometer, measuring the pitch about the lateral axis, and coordinating corrective leg motions. However, Dante II climbs slopes and must follow the pitch of the terrain. By fitting a plane to the position of

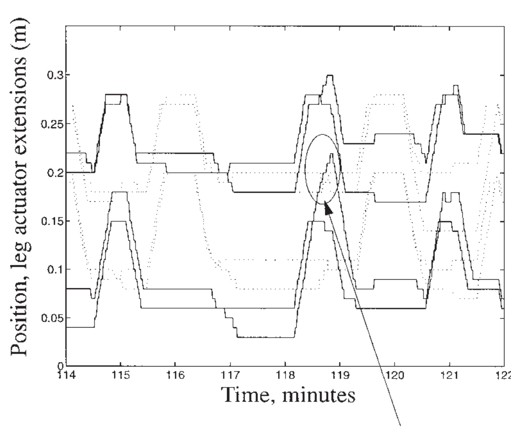

Fig. 21. Leg-reflexive motion on contact with the terrain. Note that "leg extension" is actually the actuator extension; so for instance all legs at an actuator extension of 0 would correspond to the walker standing at full height. The point at which a raised leg contacts the terrain during body motion and continues to raise higher is circled.

all the supporting legs, a coarse-estimate terrain-relative pitch can be made proprioceptively. A large object under one foot can bias the pitch estimate (and the clearance estimate), but adjusting to surmount the obstacle is not harmful. With the pitch behavior estimating relative body-terrain pitch, Dante II can negotiate transitions between differing slopes.

The gait behaviors are implemented as 24 asynchronous processes: eight contact-foot behaviors to stand, eight free-foot behaviors to step, one each of raise-legs, move-frame, turn-frame, lower-legs, and sit-still behaviors to walk, and roll, pitch, and clearance behaviors to posture. These behaviors are networked by binary links that carry inhibit- and exhibit-control signals. Each process has the same structure: it executes a nonterminating loop waiting for an incoming exhibit or inhibit signal. The inhibition/exhibition logic is simply, "exhibit when receiving one or more exhibit signals and no inhibit signals." When the process exhibits its behavior, it watches for signaled events and sensed conditions, and produces signals and actions.

The gait-control processes determine how the robot will respond to the environment, and do so by changing the position of each of the 11 degrees of freedom. To cause motions to occur, the behavior processes adjust the set points of servo-control loops associated with each degree of freedom, running on Dante II's onboard real-time computers. This approach effectively produced terrain-adaptive walking, and because of its distributed structure, enabled incremental development and layered complexity.

2.8.2. Real-Time Control

Real-time motion control is distributed among the three onboard 68030 processors, which all run the multitasking VxWorks operating system. The processor boards, shown in Figure 14, reside in a VME card cage along with a Sparc processor dedicated to the laser rangefinder, and digital-to-analog as well as analog-to-digital input/output (DAADIO) boards.

The first 68030-processor board collects sensor information and writes the state into shared memory at 120 Hz. Most sensors are wired to the DAADIO board so that their analog signals are converted to digital values. These signals are sampled and filtered to make their values usable to the control processes.

The second processor drives eight leg PD servo loops, and services dedicated motion-control boards for the translation, turn, and tether actuators. The servo loops generate trapezoidal velocity profiles from leg-encoder values at 150 Hz to produce smooth motion. An integral term is not necessary, because short motions up and down tend not to accumulate systematic error. The tether is controlled in a hybrid position-force control mode: the desired tether force is determined by the robot's inclination on the slope, and the servo loop makes small position adjustment to the amount of tether paid out to achieve the desired force. The tether force is servoed at 60 Hz.

The third board runs the gait-control processes, which can access sensor values and servo loops. It has cycles available for other functions, including external communication. One external communication process is dedicated to broadcasting status packets, typically at 10 Hz, to the remote-control station for display on the graphical user interfaces.

3. Technical Results

3.1. The Development Program

We followed a systematic and incremental approach to developing and testing the robot's components and systems. This section describes the activities and results of the development effort.

3.1.1. Leg and Tether

We devoted much of our structural engineering effort toward design of a leg and rappelling tether that would survive the Mount Spurr mission. These two components were expected to see the highest and most variable loading. They also are in contact with the terrain, and are the direct recipients of shock loading due to control errors or terrain failure. The robot has many legs and carries much tether, so a high strength-to-weight ratio is important for both. At the core of the leg and tether developments was an iterative program of design, manufacture, and testing to failure.

In the case of the tether, the principal challenge was to find a lightweight method of protecting the multiple conductors that would reside inside the cable. Without proper protection and load transfer, inner conductors can, for instance, fail when the tether is kinked under high load. We conducted various abuse tests on tether samples, including shock loading, cold-flex testing, high-tension kinking, and repeated abrasion (Krishna, Bares, and Mutschler 1997). Though costly, this process resulted in an extremely tough, lightweight tether. Strength-to-weight concerns in the leg design led to a highly optimized design that could withstand a variety of extreme loading conditions. The leg-development program addressed the slender leg links, the joints, and the welding details and procedures used to attach the links and joints. A complete finite-element leg model was constructed and optimized. Test joints and entire legs were constructed and loaded to failure. The process was repeated several times to trim additional weight as well as to improve high-stress areas. In parallel, we were able to tune and test the leg-motion controller and force-sensing scheme on separate test legs.

3.1.2. Winch and Tension Control

The basic winch was fully load tested by the vendor. We relied on the vendor's extensive experience with offshore winches and decided to forgo environmental related tests. Though the winch used industry-proven details, the method of cable stacking and the tension control were somewhat unique, and thus required special testing. The cable-routing method was complicated and we were especially concerned about the cable winding so that it would not jam on the drum (Krishna, Bares, and Mutschler 1997). A key factor in proper winding is maintaining a minimum tension throughout, so that an upper wrap of cable cannot knife down through the lower coils. After testing and verifying the accuracy of the tension-measurement system separately, we tested tension-controlled winding and unwinding. These final tests included the actual control electronics, control software, and user interface for winch operation. At the successful conclusion of these tests, the entire winch system had reached the level of being verified independently of the walking mechanism.

3.1.3. Indoor Walking

After the major components of Dante II had been assembled (legs, body, electronics, winch), the next major systemic testing phase was flat-floor walking while maintaining the required minimal tension on the tether cable. While on a flat floor, we exercised most functions, including body translation, turning, maximum body lifts, and coordinated winch operation. Since in most cases the components had been previously tested, this phase proceeded efficiently and without major problem. As we began to operate the complete robot mechanism and winch, we were able to iteratively improve the teleoperation interface. A critical portion of the testing program began as we started to control the robot to walk on a 30°, 7-m-long ramp while maintaining a gravity-balancing tension on the tether cable. At first, a safety cable was used that would stabilize the robot in the event of loss of tether tension. As part of the ramp testing, we began to teleoperate the robot to walk on and off the ramp. Doing so required changes in the robot's pitch, which caused large internal leg stresses, as well as changes in the target tension value for the tether cable. After several practice attempts, we identified a method of adjusting pitch that minimized leg-stress buildup. A feature of the teleoperation system which was very important to this test was that the tether cable would automatically control its tension, independent of robot motion.

Once the teleoperation interface and controller were operational and proven on floors, ramps, and transitions, we began testing the behavior-based gait controller. Again we tested flat-floor walking, turning, and obstacle-crossing capabilities. Specifically, the controller's ability to raise a foot after bumping an obstacle was tested and refined. As the gait controller evolved, we continued the tests with the ramp and transition crossings. We quickly found that behavior-based control tended to advance the robot much faster than teleoperation. However, human oversight was still required to halt the machine when it encountered a danger that the force sensors that guided behavior-based walking could not sense.

3.1.4. Outdoor Mixed-Terrain Testing

After testing Dante II and the basic control modes on a variety of indoor features, we began a series of outdoor tests on a hillside adjacent to the Robotics Institute at Carnegie Mellon. The path included a 10-m steep slope (50°) followed by a 5-m flat area, which transitioned into a 30-m variable-slope region (20–50°) to the top. The path also included some minor (about 10°) cross-slopes. The terrain was hard soil covered with light vegetation. During this series of tests, we set up a remote-control station so that the operators could not view or hear the robot directly. The full suite of onboard cameras and scanning-laser rangefinder were used to provide terrain information to the operators. We also limited communication between the robot and the control station to the bandwidth that would be available for the actual mission. In addition to the most challenging terrain to date, these tests brought extremely hot weather and rain, which tested the environmental readiness of the robot and support equipment, including the generator and the communication gear that would be placed at the rim of the volcano.

To conclude the mixed-terrain testing, we covered the 5-m flat section of the path with large boulders (0.5–1-m tall) in an effort to emulate the worst-case conditions expected at the crater floor of Mount Spurr (see Fig. 22). The dense boulder field proved to be the most challenging terrain yet, slowing teleoperation to as slow as 1 m/hr and resulting in a few jarring falls as feet slipped off of boulders. (Sections of the Mount Spurr ascent proved to be more challenging than even this boulder field!)

3.1.5. Endurance Testing

3.1.5.1. Early Development of Autonomous Walking

Before travel to Alaska, we tested Dante II on full-scale volcanolike terrain in Pittsburgh. The Pittsburgh slag heaps are expansive slopes of hardened slag, a byproduct of steelmaking. In early 1994, we conducted tests along a 170-m path. The upper portion of the path is level for 40 m, and then slopes into a smooth escarpment of 30–40° for 70 m and 40–50° for 5 m, and then follows a moderate but trenched uphill grade for 60 m (see Fig. 23). The major escarpment is piled at the angle of repose, and comprises a very unstable surface. The operators were again separated from the robot, this time in a control trailer at the top of the slag heap. The operators had no audible or direct line-of-sight information.

The longest autonomous run was 182 steps over 111 m in 219 min (3:39) for an average speed of 0.51 m/min. The slope varied from 30–40°, and the cross-slope (lateral to the direction of travel) was ±5°. Roll and pitch were maintained to within ±2°.

Dozens of leg bumps occurred during these tests, and the (free-leg behavior) reflex was so effective that accurate specification of leg-lift height was unnecessary; feet could skim the ground, providing protection against tipping, and raise up if they bumped. The reaction occurred so quickly that body advance slowed almost imperceptibly. In these tests we discovered how difficult and exhausting it is to

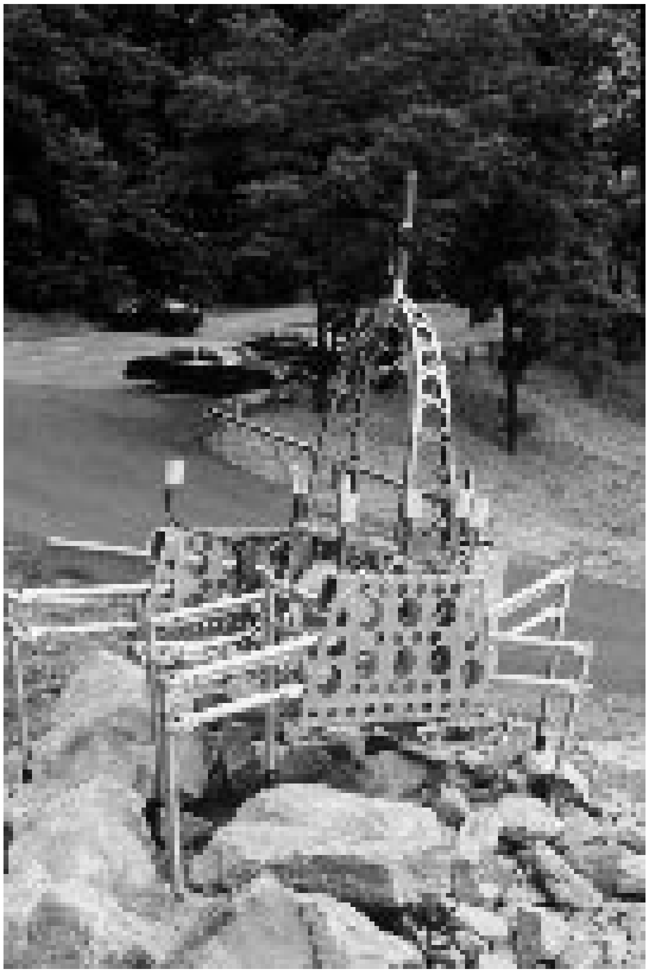

Fig. 22. Dante II in a boulder field at the Carnegie Mellon test site.

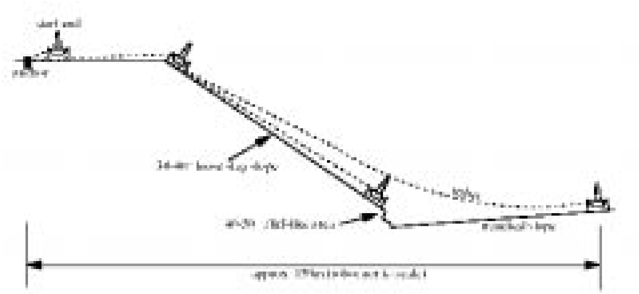

Fig. 23. The endurance test course (at the Pittsburgh slag heap). Autonomous walking was effectively used on the long sloped and level areas (see Fig. 24). Perhaps the most challenging portion of this course was the clifflike dropoff at the base of the main slope. A detail photo of Dante II crossing this terrain feature (while under teleoperation) is shown in Figure 25.

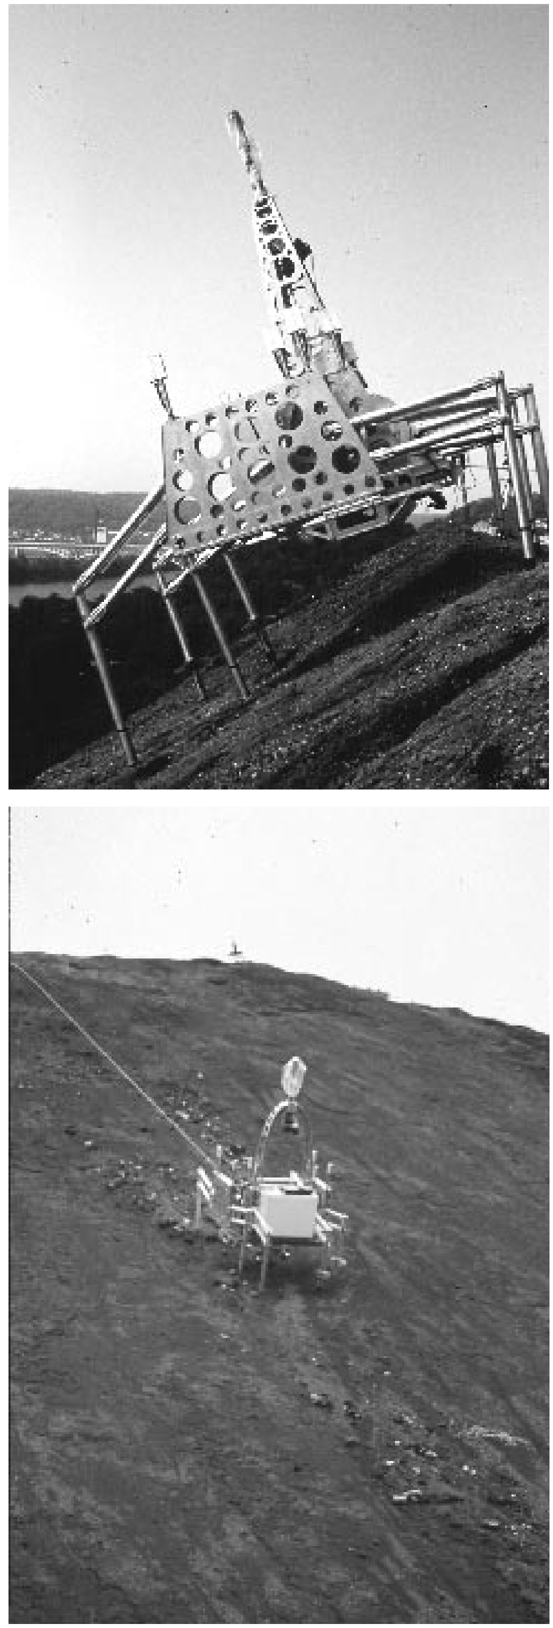

Fig. 24. Dante II crossing the transition at the top of the slag-heap main slope (top); and Dante II proceeding down the slope under autonomous control (bottom).

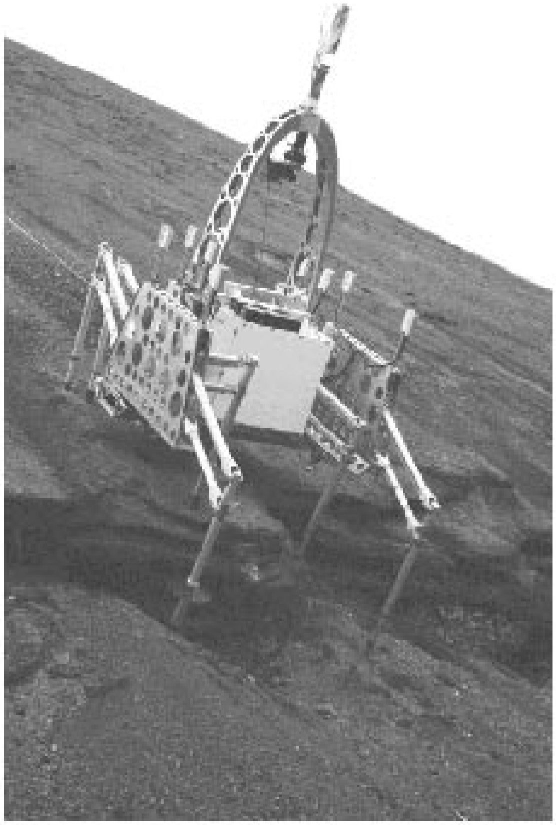

Fig. 25. Dante II crossing a 1.5-m clifflike dropoff on 40° slope.

teleoperate a walking robot. Autonomous walking is faster, and with fatigued operators, more reliable. Most of the 100 m of slag-heap operation was done using the behavior-based gait control (see Fig. 24). The operators did, however, control the robot in the several areas of major slope transition. In one particular transition (see Fig. 25), the robot was nearly at its kinematic limit as operators guided it over a 1.5-m clifflike step and onto the 40° escarpment.

3.1.5.2. Premission Shakedown

A premission test evaluated the robot, communication, and team readiness for the actual mission. The robot and all equipment were packed for the mission, then using only the four-person deployment team, we installed the robot and rim equipment including satellite telemetry on the edge of a large gravel pit in Anchorage. This was the first time that the rim station included a remotely controllable camera that could be used to track the robot's progress. The remote camera was found to be more useful for causal observers in the control trailer rather than operators due to poor perspectives, especially as the robot advanced far down the slope. Onboard cameras and the laser scanner were preferred sources of terrain information. Over a 30-hr period, we conducted tests and demonstrated complete mission readiness.

The gravel pit provided a moderate slope, on average 30°, with sandy soil eroded with 10–50 cm ruts. The path of descent and ascent was 92 m. From the top, the slope descends approximately 50 m before it abruptly transitions to a level bench. From this bench, it again drops (to 30°) for 5 m, and then slowly transitions to level over 30 m.

The descent into the gravel pit was primarily teleoperated, although brief portions were autonomous. The descent required 321 min (5:21) for an average speed of 0.28 m/min. This is less than half the maximum speed dictated by the motor/amplifier configuration, and was due primarily to human delays in interpreting sensors, considering information, and making plans. Ascending the same terrain autonomously with the behavior context required only 179 min (2:59). In the autonomous mode, the gait controller averaged 0.51 m/min, and in some areas averaged 0.67 m/min, more than twice the human-controlled speed (shown after the pause in Fig. 26).

The ascent did require one instance of operator intervention: the first step above the level bench onto the 30° slope was placed in a depression, masking the imminent uphill transition. The pitch correction required in the next step would have exceeded 20° (which is physically stressful to the mechanism), and was apportioned by the operator to take place over two steps (instead of one). This instance of shortsightedness on the part of the autonomous system points out the need to foresee some situations and prepare in advance.

During the premission test, we also used communication links to NASA Ames Research Center in Moffett Field, California, and the National Air and Space Museum in Washington, DC to view and interact with the robot.

3.2. The Mount Spurr Expedition

Dante II's final destination, Mount Spurr, erupted three times in 1992, spreading 200,000,000 m^3 of ash over portions of

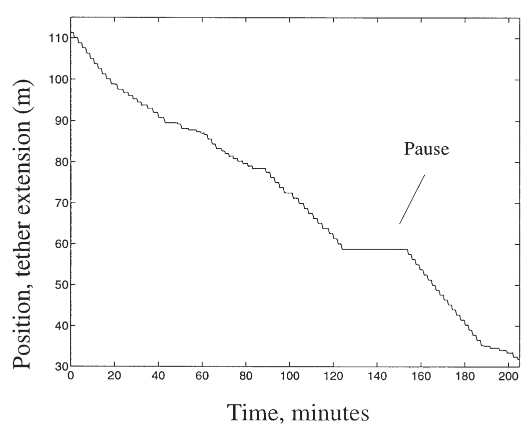

Fig. 26. The Dante II tether extension when ascending from an Anchorage gravel pit (the actual robot travel is approximately equal to the tether extension).

western Alaska. It is of interest because of its proximity to Anchorage and its potential for further eruptions. Because of debris instability from previous eruptions, volcanologists were unable to enter the crater to assess the likelihood of future volcanic activity.

The active crater is on a secondary peak, Crater Peak, at elevation 2,300 m. One side of the crater is composed of a 350-m vertical wall with talus slopes at the base. The other side is blown out, with a broad flat rim and a 20–45° slope down to the crater floor. The slope is covered with snow, wet ash, and mud, which are deepest in the middle of long chutes that run downhill. Ridges divide the chutes and, like the crater floor, are littered with meter-size blocks. Fumarole vents of interest are located on the crater floor. Dante II would enter the crater from a launch point located on the broad flat rim. In July 1994, Dante II and support equipment were transported by barge, truck, airplane, and finally helicopter to Mount Spurr. A short setup period was required to install and arrange the various components on the volcano (see Fig. 3). Dante II's tether was anchored to a large beam buried deep in hard-pack snow, which covered much of the upper extents of the crater at the time of deployment. A diesel generator flown to the volcano's rim was sized to provide 2 kW to power the robot and communication equipment for two weeks without additional fuel or service.

Problems with a high-speed data link between the robot and the rim telemetry station, as well as a damaged load cell delayed the start of the mission for several days. Because of the inherent danger on the volcano, the setup crew could only make brief visits to the deployment site. Some attempts to helicopter to the volcano rim were thwarted by bad weather. Finally these problems were resolved, and Dante II began its descent on July 29, 1994. Dante II was truly embarking on a descent to explore unknown terrain: humans hadn't entered Mount Spurr's crater since before the last eruption in 1992. Dante II would operate without problem for the next five days with no humans on or near the volcano.

The mission consisted of three segments: descent to the crater floor, floor exploration, and ascent. As a result of many ash and rock ridges that extended down into the crater, during most of the descent Dante II contended with complex slope conditions. The upper section of the crater was covered with hard-pack snow, across which Dante could move at about 1 cm/sec autonomously using leg-force sensors, pitch and roll sensors, and a behavior-based control scheme. As the receding snow-pack boundary was crossed, the terrain became much more rugged and forced frequent use of an enhanced teleoperation mode (frame context) that placed all critical control decisions on the human operators. Meter-scale boulders were negotiated on a frequent basis. Footing conditions worsened as Dante II progressed lower into the crater and left the hard snow pack for deep mud and steep areas of ash. The snowshoes mounted to each leg were invaluable in limiting sinkage in the soft mud.

On the upper slopes of the crater, Dante II walked autonomously twice: 9.8 m in 23.3 min (0.42 m/min) and 9.6 m in 19.78 min (0.49 m/min). Shown in Figure 27, this is about twice the typical speed of teleoperation.

Because it was necessary to navigate across the chutes and ridges during the descent, Dante II experienced cross-slopes up to 30°. This severity was unexpected. We hoped to travel directly downslope as much as possible, but the robot reached dead ends that had taken hours or days to discover—apparently this is the nature of exploration. The most direct exits were to turn across the slope and climb over a ridge into the next chute.

Near the bottom of the descent, Dante II made two autonomous descents down a snow chute, 8.3 m in 35.2 min (0.24 m/min) and 6 m in 12.3 min (0.49 m/min). The snow was hard packed, and rocks lying on its surface were easily cleared when bumped.

As Dante II advanced farther into the crater, progress slowed not only due to the challenging terrain conditions but also because sensors could not see far enough ahead to avoid dead ends and nontraversable areas. At least half of the time was spent retreating up the slope to try a new path around an especially large boulder or ridge. Another complicating factor was an ongoing concern that the tether would become trapped between boulders far up the mountain. As Dante II moved across the slope relative to the anchor point, the tether would at times straighten to attain the shortest path and in doing so move laterally across the rough terrain. Dante II's lateral moves were therefore limited to minimize the chance that the tether would become trapped, and thus preclude future ascent. Though overall terrain slope lessened in the crater-floor region, boulder size and density increased—with some boulders larger than cars—making continued travel treacherous.

The laser-built 3-D elevation maps proved crucial to efficiently guide the robot and minimize path reversals. In a common mode of operation, humans would set the

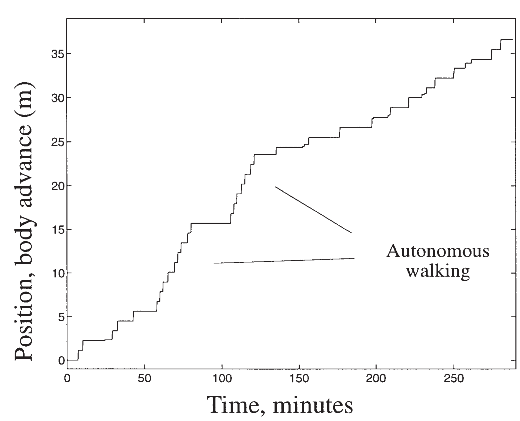

Fig. 27. The Dante II body advance when descending into the Mount Spurr crater.

direction of travel and then turn over control to the automated behavior-based controller, which would then enact efficient safe walking in the intended direction. When forced to revert to enhanced teleoperation, progress was halved as human operators attempted to contend with a huge amount of sensory data including video and laser data as well as vehicle-state information.

Throughout most of the mission, and especially during the crater-floor exploration, imagery was transmitted beyond Anchorage to NASA Ames Research Center, where geologists studied the crater and controlled Dante II's cameras while the robot operators were resting.

After two days of exploring the crater floor and detecting no water and no sulfur compounds in the fumarole gasses, Dante II began its ascent out of the crater. Early in the ascent, all power and communication was lost to the robot. We suspected an onboard problem, which would have ended the mission, as the robot was still far below the region in the crater to which human access was deemed to be safe. However, the culprit was found to be a moisture-related short circuit in the power cabling at the rim, which was quickly remedied, and the mission continued. The remedy, however, required an engineer to fly to the crater rim, thus ending Dante II's period of true unattended operation. Up to this point, the robot had operated five days without any human presence on the volcano, and successfully achieved all of the mission objectives.

As Dante II progressed midway out of the crater, the most difficult terrain of the mission was encountered. This was due in part to the fact that much snow had melted since the descent, revealing harsh underlying slopes and boulders. Also, by the fifth day, the laser-scanner mirror had become obscured by airborne volcanic ash and thus video images became the only means of viewing the terrain. Determination of path became more difficult, and much time was spent retreating back into the crater to try alternative exit routes. Ultimately, Dante II fell on its side while under teleoperated control. The accident was due to a combination of factors including steep slope and cross-slope conditions, soft unstable slope material, a destabilizing tether-exit angle, and a control algorithm that had never been tested in such perilous stability conditions. Dante II was fully operational, but unable to right itself.

In an initial attempt to salvage the robot, we removed the end of the tether connected to the anchor and attached it to a helicopter. The helicopter then tried to lift the robot using the 150-m length of tether. This attempt failed when the tether broke near the robot and caused the robot to tumble several times down the steep, boulder-covered slope. Ultimately we decided that humans could attach a sling with minimal danger, since the robot was about halfway out of the crater. One team member and an experienced salvage expert hiked down the slope and attached a sling to the robot so that it could be lifted out by helicopter. All other rim equipment was retrieved, and the mission was complete.

4. Lessons Learned

4.1. Technical Lessons

The lessons learned in the development and deployment of Dante II are described in the following subsections.

4.1.1. Framewalking Is Suitable for Rappelling

Walking mechanisms are excellent as the underlying locomotion means for a rappelling robot: legs can enable the robot to cross very large terrain features without deviating from the fall line and jeopardizing lateral stability. In very rough terrain, a walking robot can avoid undesirable footholds, optimize stability, and move its body independent of terrain details (Bares and Whittaker 1993). Using a variety of control modes, walking robots have demonstrated rough terrain capabilities undersea and on a variety of terrestrial sites (Ishino 1983; Onodera 1989; Pugh et al. 1990; PlusTech Oy Corporation 1995). At issue in the configuration of a walking robot is whether the mechanism should be a mechanically simple "framewalker," which advances groups of legs together, or a more complex "terrain-adaptive" robot, which moves legs individually and is closer to animal analogs. Terrain-adaptive walking robots require more cognition and control to move the legs, but offer more flexibility in leg placement and gait. Our experience with Ambler, however, showed that terrain adaptability may offer more flexibility than needed, and that stability margins must not be underemphasized in rugged terrains (Bares and Whittaker 1993).

Framewalking places fixed "constellations" of feet, which offers a basic level of stability without relying on planning and control. However, if the foothold for any single foot is poor (e.g., on the edge of a boulder), the entire frame must be repositioned until satisfactory footholds are found for all feet on the frame. Once, crossing a 10-m-long boulder field required more than 4 hr while Dante II operators labored to avoid placing feet on the edge of boulders. A major improvement to the foothold-selection problem would be to automate the search using a terrain-elevation map.

In addition to inherent stability, framewalkers offer several key advantages for rappelling. First, since a rappeller is basically constrained to move down the fall line, a mechanism specialized for straight-line motion, such as a framewalker, is preferred. Second, only two motions need to be synchronized to unwind the rappelling cable: winch rotation and frame translation. Even when crossing large obstacles and slope transitions, Dante II rappelled smoothly with this simple control scheme. Alternately, a rappelling terrain-adaptive robot would have to coordinate all supporting legs and the winch pay out to achieve smooth motion.

4.1.2. Rappelling Extends Capability, but Limits Scope

A tensioned support cable can enable exploration of terrain otherwise impossible for a terrestrial vehicle to safely traverse.

However, the obvious shortcomings of rappelling locomotion are that a reliable anchor is needed and the exploration range is limited to the cable length. Dante II's rappelling line is 300 m; ultimately, a line several times longer may be feasible for mobile robots, but the size and weight of the rappelling system are real concerns. However, if all conductors could be eliminated from the tether through the use of onboard power and wireless communication, a dramatic decrease in diameter and thus increase in cable length could be realized.

While a tensioned cable can greatly improve the a robot's stability and ability to traverse various terrains along the fall line of the slope, travel away from the fall line creates restoring forces on the robot that try to move it back toward the fall line. Restoring forces can cause the robot to slip sideways or even tip over (a side pull from the tether aggravated the conditions that led to Dante II's tip-over.) On extremely steep slopes, if the restoring forces overcome the frictional forces on the feet, the robot will swing much like a pendulum—a very dangerous condition for a mobile robot, especially if external sensors and appendages might be damaged. Maximizing a robot's obstacle-crossing capability can reduce the need for travel away from the fall-line path. For this reason, Dante II was designed with the ability to traverse 1.3-m obstacles.

4.1.3. Self-Righting is Unrealistic in Most Scenarios

After Dante II tipped onto its side, a frequent question was whether some means of self-righting had been considered. There are several areas of technical concern with self-righting scenarios: first, during a tip-over and possible subsequent tumble, it is likely that some devices such as terrain sensors, communication equipment, and actuators will bear the brunt of energy absorption and be damaged or rendered inoperable. Even if not damaged, communication and solar-array equipment may only function when aligned upright.

The second fundamental challenge to self-righting is that when exploration-robots tip over, the reason is more often due to static instability conditions, rather than dynamic instability. As a result, the robot cannot simply be "stood up" or righted again at the same location, as it will immediately repeat the fall. Rather, some intelligence is needed to attempt self-righting that will not result in a repeated fall. A tip-over and a tumble into a steep ravine could require a complex series of recovery steps to self-right to a safe stable stance.

Finally, even if the major issues of preventing damage and avoiding repeated tip-over can be overcome, we are faced with the final question of developing devices to enable self-righting. Such devices need to be of high strength, and should be capable of self-righting the robot from a variety of tip-over conditions. Rocket thrusters, airbags, helium balloons, and highly geared linkages have been suggested. Since such a device would be used only rarely, it would need to be small and of low mass, as it represents lost payload capacity.

Several years ago Waldron proposed a multibody wheeled robot that could self-right through a series of body twists (Waldron, Kumar, and Burkat 1987). Several clever robot configurations address the issue with mechanisms that can either walk upside down (Angle and Brooks 1990), or have no definition of "upright" (Pai 1995). Even so, given the concerns for component damage during a tip-over or tumble and the need to avoid a repeated tip-over, it may be most appropriate to focus future research efforts on hardware and software approaches to anticipate and prevent tip-over.

4.1.4. A Single Electronics Enclosure Lowers Risk

All Dante II electronics were placed in a single sealed and shock-isolated enclosure. The oversized enclosure was provided with heating and cooling for all expected development and mission conditions. From a reliability viewpoint, wiring connections are minimized, and most are located within a protected enclosure. Since the enclosure environment was temperature maintained and shock protected, conventional "office-grade" telemetry and computing equipment could be used for which hardened versions were cost prohibitive or nonexistent. As a result, we could quickly and economically explore experimental designs and options with computing, telemetry, and video. During development, the single box could easily be removed for troubleshooting and repair.

4.1.5. Reflexes and Coordinated Behaviors Produce Tactical Autonomy

Dante II has shown us that gait adaptation during execution is necessary for walking in natural terrain. Too much is unpredictable, like bumping obstacles or slipping off footholds, and cannot be foreseen while planning.

Biological systems provide evidence for simple sense-act reflexes and decentralized control in walking. In their neurological basis, these systems are constructed of inhibitory and excitatory links between neurons to create reflexes and behaviors to sequence fixed patterns of action. They possess properties necessary to gait execution: reaction to unexpected events, concurrency of reflexes, and coordination of actions. For Dante II, we took an approach that benefits from these properties, and developed a network of behaviors to use leg force and contact information as well as body-attitude measurements to advance the framewalker across variable terrain. Together they establish a nominal gait to achieve forward progress, and can individually react quickly to unexpected conditions. The ability to adjust behaviors and guide their collective performance established Dante II's supervised-autonomy control mode. During testing and later during the mission, the behavior-based control operated the robot at its maximum speed, which was significantly faster than the most skilled human teleoperators.

4.1.6. Operational Control Contexts Structure the Degree of Autonomy

A robot performs actions as the result of the coordinated control of its actuators. In one situation, it may be necessary to control a single motion in a direct manner, while at other times it may be appropriate to broadly guide only the coordinated result of many individual motions. A robot may be autonomous in some respects, but eventually it must take external guidance. We found that it was effective to construct a software-control system based on the idea that for an exploration robot, where the unforeseeable is commonplace, it may be necessary for a human operator to intervene at any level to control the actions of the robot. We designed control software in layers so that control signals could be generated by a human operator or by other software modules. We then developed the operator interface with specific control contexts that described which functions are performed by the robot and which are performed by the human operator.

As an architecture, an organization into control contexts aids development by partitioning the system and clarifying interrelations. For instance, force-based leg control could be refined and tested even before the robot's body was completed. Efforts proceeded in parallel to build each control context. Though ultimately targeting autonomy, our development strategy was to build a solid teleoperation capability after which more autonomous capabilities could be layered as developed. Then, if schedules slipped or technical problems arose with the implementation of autonomous modes, the mission could proceed using only teleoperation.

As an approach to robot teleoperation, the contextual architecture simplified the challenging task of controlling a walking robot. The ability to access low-level functionality was crucial to extracting Dante II from several exceptional conditions and tight spots. High-level functionality built on robust low-level capability enabled more efficient operation in easier terrain.

4.1.7. Constant Oversight Limits Autonomy Benefits

We found that the ideal control mode for a remote explorer is a supervisory mode in which a human occasionally monitors the robot's progress but focuses on scientific observation and mission strategy. The robot is given high-level directives such as "move ahead 10 m," from which it must sense the terrain, plan its actions, and execute the actions in a safe manner. Dante II completed 25% of its descent into Mount Spurr under this type of supervisory control. Since its behavior-based control only uses proprioceptive data, it is "blind," and cannot reason about complex terrain situations. As a result, continuous human oversight is required to monitor progress and reinstate teleoperation when automated performance worsens or the robot's safety becomes a concern.

Continuous human oversight and safety response is impractical for terrestrial robots, and impossible for planetary systems for which large communication latencies exist. An important research focus must be to embed robots with sufficient intelligence to know when they should cease autonomous operation and request human input. Most "autonomous" mobile systems developed to date relegate this key decision to humans. Field robots will become more useful and practical when they are able to conduct safe operations without continuous oversight, and request human assistance when needed. As robots take on increasingly challenging autonomous tasks, this requirement escalates in importance.

4.1.8. Graphical Presentation of Telemetry Improves Operator Performance

A teleoperated robot that is instrumented with many sensors, each of which is sampled at a high rate, can produce an overwhelming amount of information for an operator to comprehend. Much of this information can be distilled to a dozen or so critical values. With the addition of video imagery, which itself represents a substantial volume of information, monitoring robot telemetry is difficult to sustain reliably, and is ultimately fatiguing. A number of studies confirm this observation (Sheridan 1992). We developed a multiwindow graphical user interface for this purpose.

We have had a number of unsatisfying experiences with text-based operator interfaces: they invariably provide a daunting array of numerical information and an indecipherable collection of command codes. The graphical interface provided a refreshing change that was quickly adopted. We have found that a graphical presentation of robot telemetry simplifies monitoring of the critical information and improves the operator's performance, increasing his/her abilities to comprehend the robot's state and teleoperate and interact with the robot over long periods of time.

Our primary objective in creating a graphical operator interface was to make it as easy as possible to operate Dante II. We were motivated to minimize operator workload and to make it possible for novices to quickly learn to control the robot. The following guidelines emerged:

- Consistent appearance and interaction. Similar or identical design throughout the interface allowed operators to focus on robot actions rather than the mechanics of using the interface.
- Functional organization. It is appropriate to embed the functional layout within the interface to avoid operator confusion. The use of operational control contexts provides a unifying and simplifying perspective on human-machine interaction. This approach enabled us to concisely organize the interface so that commands appropriate for a particular type of function were grouped together.
- Uncluttered layout. Clean graphical design with qualitative or simple quantitative representations of sensor and state information allows for quick assessment of current conditions. Numeric data provides precision, and 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 on the situation and the types of commands being applied. Indicators that clearly reflect active safeguards reduce operator misconceptions and error.

Anecdotal evidence suggests that operators were able to run Dante II longer, faster, and safer with the graphical interface as compared with numeric interfaces on previous robots. We also found that visitors watching operations were able to quickly grasp how the interface worked.

4.1.9. Terrain Visualization Is Essential to Exploration

A frequently updated model and visualization of the robot and surrounding terrain is essential if efficient and thorough exploration is to take place. Visualization of the robot and terrain together aides situational awareness. A full 3-D "virtual environment" interface was developed for Dante II to provide operators an easy means to visualize the robot's stance, forces, and surrounding terrain topology. Using this 3-D interface, operators could rather easily plan safe moves after studying the virtual robot and terrain from various viewpoints. In comparison, teleoperation from 2-D video imagery was difficult because of a lack of depth information. Furthermore, the 2-D images offered no visual cues that would help an operator reconcile object size and distance from the robot.

Virtual environment interfaces can be used to improve an operator's situational awareness and to efficiently visualize complex terrain and vehicle-state information. We found that high frame rate, level of interactivity, and ease of use were all contributing factors for achieving immersiveness and a sense of presence. However, insufficient visual-reference aids and accuracy of correlation between graphical models and physical objects can degrade operator performance and confidence.

While climbing out of the volcano, the laser rangefinder became dysfunctional after airborne volcanic ash coated its spinning-mirror surface. Teleoperation became extremely tedious as operators were left to infer terrain geometry based only on video and leg-position information. Progress slowed and safety margins declined until the tip-over occurred.

4.1.10. Self-Sufficient Mobile Robots Are Suited for Remote Science

A ground-based mobile robot is an excellent platform from which to conduct scientific and data-logging activities, in part due to its ability to precisely position using internal and external sensors. Temporal and environmental data can be added to the positional information to create a complete sampling history. Using ranging sensors, a mobile robot can quickly map surface topology to a resolution and accuracy that would require a human team many days to complete. A robotic vehicle can dwell for days, weeks, or even months while collecting, analyzing, and streaming data to scientists at remote stations. In the event of a dangerous event such as a rock fall, only the robot is jeopardized.

Once a mobile robot is placed in the field to stream data back to humans, highly interactive robotic science missions involving many people are possible using live video and data links. As demonstrated by Dante II and more recently by the Nomad robot (Wettergreen et al. 1998), the nature of exploration has evolved to the point where scientists in several locations can simultaneously be in "mission control" and view progress, analyze data, and advise.

Critical to this new era of exploratory robots is a system design and method of deployment that eliminates human danger and minimizes costs. While deploying Dante II at Mount Spurr, significant human effort and expense was required to emplace the anchor and configure power, communications, and logistical equipment at the rim of the volcano. Including the large robot, nearly 3 tons of equipment were transported to the volcano. Self-deploying compact systems (e.g., air-drop "packages") are more practical on Earth and are essential for space missions.

4.1.11. Overdesign Can Enable Rapid Development

A key challenge of the Dante II project was to design and increase the ruggedness of the robot and support equipment to enable repeated testing, transport, and deployment at a variety of sites, culminating with the harsh volcanic environment.

Our approach was to overdesign components and subsystems to enable rapid integration and unexpected conditions during deployment and exploration. It has been our experience that robot-system integration commonly stresses components more than actual operating conditions. One approach is to slow integration in an attempt to avoid control "mistakes" that could damage the electromechanical systems. We believe, however, that it is much more expedient to overdesign components where possible, and thus permit integration to proceed rapidly and with less concern for equipment damage. For instance, motors were sized to withstand long durations of stall current, even though this condition was designed to be protected by software. Of course, overdesign is no excuse to avoid proper control-code design, simulation, and off-line testing.

Overdesigned hardware also helps the robot survive the rigors of deployment. (It may be that deployment and transport activities—rather than exploring in the volcano crater—were the harshest conditions that Dante II had to face.) By nature, an exploration mission will encounter unexpected conditions. We believe that a systemic development approach that stresses overdesign will also result in a system more suited for an exploration mission. A common result of overdesign, however, is increased weight. Our approach was to prioritize reliability and durability, although we also made continued efforts to reduce weight.

4.2. Programmatic Lessons

The Dante expedition to Mount Spurr made for an atypical project, because the intention was to develop a capable robotic system and then demonstrate it in the field, rather than under controlled laboratory conditions. The is the ultimate objective of field robotics, to put machines to work in the uncontrolled natural environment. During the course of the project we learned several programmatic lessons that aided in our getting Dante II to the field.

4.2.1. Set Specific Objectives and Success Criteria

For a robotic field experiment, it is beneficial to identify specific objectives that are clearly defined and measurable. Defined objectives serve two important purposes. First, they help to reveal what will be necessary to perform the field experiment. It is inefficient to conduct applied research without goals and a rational plan. And worse, going to the field without a clear idea of what you are trying to accomplish seldom results in a successful experiment. Second, clearly defined objectives help to eliminate tangential efforts that creep into the scope of work. It is essential to restrain any ballooning effect and keep the project focused on only activities that move toward the objectives. (This may not be true of basic or more speculative research.)

Similar to objectives, which lay out the areas in which research will be conducted, success criteria for a field demonstration tell you what the robot must be able to do and when it is done. What were the success criteria for Dante II's expedition to Mount Spurr? Get to the crater floor and transmit images and scientific data. That was it. Even though Dante II tipped over while we were trying to bring it up from the crater floor, NASA already considered the experiment a success. Both during development and, more critically, during a field experiment, activities that do not contribute directly to success should be avoided.

4.2.2. Harsh Field Experiments Drive Program

It is important to know the complexities of the proposed field experiment. In addition to knowing the robotics issues involved, it may be necessary to visit the site to aid in the design process. If the experiment is to involve science, it is important to find expert partners who can advise regarding appropriate sensors, methods, and analysis. It takes time and thorough examination to realize all the subtle issues and various aspects of an intended field deployment.

One of the major accomplishments of the Dante II project was that it was structured not only to develop new technology and approaches to remote robotic exploration, but that it was to demonstrate the technology in real, unforgiving environments. While the field deployment generated much publicity and interest, the key technical reason to take on such a committing experiment was that it focused the entire team onto a clear, quantifiable goal, it and forced technology that was capable of performance under rugged conditions. As a result, we believe that the Dante II project contributed to robotic science as one of the first unattended field expeditions.

4.2.3. Software Necessarily Lags Hardware

Software development necessarily lags hardware development, initially because hardware specifications are not complete and finally because integration and testing cannot occur without working hardware. In a fast project, where delays begin to compound early, it is beneficial to build hardware components in parallel and, whenever feasible, overdesign to compensate for other, less reliable components. By defining intermodule communication methods first and implementing all the communication messages with each of the software modules "stubbed-out," we were able to begin implementing software modules in parallel, and we laid substantial groundwork before the physical hardware was available.

4.2.4. Operational Failure Modes Are Significant

Overall, reliability and failure modes were stressed in the design and preparation phases of the project, because the working assumption was that if the robot failed while in the volcanic crater, human entry and repair would be forbidden. Component-related failure modes and subsequent implications nearly became an obsession during the development process. Interestingly enough, however, throughout the development program we did not seem to spend sufficient time contemplating operationally related failures and possible implications. Indeed, even as the actual mission wore on, we became increasingly fixated on the likelihood of component failures and their effect. Operational failures such as teleoperating the robot into a tip-over condition or tether entanglement were not considered to the same extent. In retrospect, additional focus on possible operational failures would have been prudent.

4.2.5. Unforeseeable Events Sometimes Occur

When exploring the unknown, expect to find it. This is the nature of exploration, and operation with a field robot is no exception. The robot will eventually, either technically or physically, meet a circumstance for which no contingency has been devised.

4.3. Conclusions

The purpose of the Dante project was to develop technology relevant for use in space exploration, and demonstrate the concepts in a harsh terrestrial environment. To this end, the Dante II robot system incorporated a variety of new technologies and put them to ultimate test in a demanding, planetary-like field setting. Conclusions we can draw from this effort span the realm of technical and programmatic:

- Rappelling with a legged robot is an excellent means of locomotion in severe terrain that includes a combination of complex slopes, slope transitions, obstacles, and soft materials. Furthermore, with a framewalker, coordinating rappelling and walking control is readily achievable.
- Future autonomous mobile robots operating in unknown hazardous environments must be capable of realizing when they need human assistance, instead of relying on humans to monitor and redirect operation.
- Immersive and readily comprehensible terrain and robot-state information is crucial for effective teleoperation in severe terrain.
- The possibility of catastrophic failure is very real in severe terrains. Careful system design as well as operational training can help to reduce—but not eliminate—the possibility of a mission-ending event. The focus should be on avoiding destabilizing conditions, rather than engineering recovery methods.
- The requirement for a "real" exploration mission can be a galvanizing factor in driving technology and focusing a development program.

Is the Dante II system appropriate for space exploration or to serve as a terrestrial volcanologist? Perhaps not in its present form. However, we believe that some of the lessons learned can be used to form the basis for next-generation terrestrial explorers, and future generations of their planetary counterparts.

Acknowledgments

The development of Dante II and the expedition to Mount Spurr were supported by the National Aeronautics and Space Administration (NASA) under grant NAGW-1175.

The authors would like to thank Dave Lavery of NASA Code SM for his unflagging support of the Dante missions and telerobotic research in general. We would also like to thank Red Whittaker for his mentorship and guidance throughout this challenging endeavor. Scott Boehmke helped on this paper, as well as being a tremendous leader in the Dante II project. To the rest of the Dante team: thanks for your incredible efforts and teamwork that made it all possible!

References

Angle, C., and Brooks, R. 1990 (Tsuchiura, Japan). Small planetary rovers. *Proc. of the Intl. Conf. on Intell. Robots and Sys.*

Apostolopoulos, D., and Bares, J. 1995 (Pittsburgh, PA). Configuration of a robust rappelling robot. *Proc. of the Intl. Conf. on Intell. Robots and Sys.*

Bares, J., and Whittaker, W. 1993. Configuration of autonomous walkers for extreme terrain. *Intl. J. Robot. Res.* 12(6):535–559.

Boehmke, S., and Bares, J. 1995 (Boulder, CO). Electronic and telemetry systems of Dante II. *Proc. of the 41st Intl. Instrumentation Symp. of the Instrumentation Soc. of Am.*

Brooks, R. 1986. A robust layered control system for a mobile robot. *IEEE J. Robot. Automat.* 2(1):14–23.

Fong, T., Pangels, H., Wettergreen, D., Nygren, E., Hine, B., Hontales, P., and Fedor. 1995 (San Diego, CA). Operator interfaces and network-based participation for Dante II. *Proc. of the SAE 25th Intl. Conf. on Environmental Sys.*

Grillner, S. 1985. Neurobiological bases of rhythmic motor acts in vertebrates. *Science* 228:143–149.

Hine, B., Hontales, P., Piguet, L., Fong, T., Nygren, E. 1995 (San Diego, CA). VEVI: A virtual environment tleoperations interface for planetary exploration. *Proc. of the SAE 25th Intl. Conf. on Environmental Sys.*

Hirose, S. 1983. Adaptive gait control of a quadruped walking vehicle. *Proc. of the First Intl. Symp. of Robot. Res.*

Ishino, Y. 1983. Walking robot for underwater construction. *Proc.s of the 1983 Intl. Conf. on Adv. Robot.*

Krishna, M., Bares, J., and Mutschler, E. 1997 (Albuquerque, NM). Tethering system design for Dante II. *Proc. of the IEEE Intl. Conf. on Robot. and Automat.* Washington, DC: IEEE.

Onodera, S. 1989 (Seattle, WA). Several examples of robot and automation on the underwater construction work. *Proc. of the 16th Meeting of US/Japan Marine Facilities Panel*.

Pai, D. 1995 (Pittsburgh, PA). Platonic walking beasts. *Proc. of the Intl. Conf. on Intell. Robots and Sys.*

PlusTech Oy Corp. 1995. Tampere, Finland. Technical Brochure.

Pugh, D., Riggle, E., Vohnout, V., Bihari, T., Walliser, T., Patterson, M., Waldron, K. 1990. Technical description of the adaptive suspension vehicle. *Intl. J. Robot. Res.* 9(2):24–42.

Sheridan, T. 1992. *Telerobotics, Automation, and Human Supervisory Control.* Cambridge, MA: MIT Press.

Stentz, A., and Hebert, M., 1995. A complete navigation system for goal acquisition in unknown environments. *Autonomous Robots* 2(2).

Thomas, H., Hine, B., and Garvey, J. 1995 (Princeton, NJ). Kilauea: A terrestrial analogue for studying the use of rovers in planetary science. *Proc. of the 1995 Conf. of the Space Studies Institute*.

Waldron, K., Kumar, V., and Burkat, A., 1987 (Versailles, France). An actively coordinated mobility system for a planetary rover. *Proc. of the 3rd Intl. Conf. on Adv. Robot.*

Wettergreen, W. 1995. Robotic walking in natural terrain, PhD thesis, Carnegie Mellon University.

Wettergreen, D., Bapna, D., Maimone, M., Thomas, G., 1999. Developing Nomad for robotic exploration of the Atacama Desert Robotics and Autonomous Systems. 26(2-3) 127–148.