1991 Year End Report
Autonomous Planetary Rover

Eric Krotkov Reid Simmons William Whittaker
Principal Investigators

CMU-RI-TR-92-02

The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213

February 1992

This research was sponsored by NASA under Grant NAGW-1175. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of NASA or the United States Government.

Contents

List of Figures

Abstract

This report describes progress in research on an autonomous robot for planetary exploration performed during 1991 at the Robotics Institute, Carnegie Mellon University. The report summarizes the achievements during calendar year 1991, and lists personnel and publications. In addition, it includes several papers resulting from the research.

Research in 1991 focused on understanding the unique capabilities of the Ambler mechanism and on autonomous walking in rough, natural terrain. We also designed a sample acquisition system, and began to configure a successor to the Ambler.

Understanding Mechanism Capablities — In 1991 we concluded an investigation of the interaction between the Ambler and the terrain. This work led to new stability measures, novel force redistribution models, and reactive control schemes. In order to quantify the performance of the Ambler, we measured the power consumption of the Ambler while walking, and while raising and lowering the body. We also calibrated more exactly the kinematics of the mechanism.

Autonomous Walking — The experimental program in walking on rough terrain continued and expanded. Highlights include the following:

- Long-term autonomous walking over challenging terrain, including a live demonstration to sponsors.
- Autonomous walking outdoors, including night-time navigation.
- Installation of all computing and electronics on-board the Ambler.

Sample Acquisition — In 1990 we demonstrated sampling capabilities on a testbed separate from the Ambler. In 1991 we designed a sampling system for the Ambler. It consists of a commercial manipulator, a short-range light-stripe sensor, and a storage receptacle. The arm and sensor are mounted on the horizontal link of a leg, rather than underneath a body. This design was not fabricated during 1991 for financial reasons.

Configuration of a Successor to the Ambler — We began to configure an integrated lander/walker for planetary exploration. This preliminary work is incomplete, and we expect this topic to be one of the central concerns for the research program in 1992.

1 Introduction

This report reviews progress during 1991 at the Robotics Institute, Carnegie Mellon University, on research sponsored by NASA titled "Autonomous Planetary Rover." This program to develop an Earth-based prototype of an autonomous planetary rover is organized around three teams that are developing the locomotion, perception, and planning subsystems. A joint task is to integrate the three subsystems into an experimental robot system. We will use this system for evaluating, demonstrating, and validating the concepts and technologies developed in the program.

The technical objectives of the research include the following:

- To develop and demonstrate an autonomous Earth-based mobile robot that can survive, explore, and sample in rugged, natural terrains analogous to those of Mars.
- To provide detailed, local representations and broad, 3-D descriptions of rugged, unknown terrain by exploiting diverse sensors and data sources.
- To demonstrate robot autonomy through a planning and task control architecture that incorporates robot goals, intentions, actions, exceptions, and safeguards.

This report is organized as follows. The next four sections describe key accomplishments of the project research from January 1991 to December 1991. These accomplishments span four research areas: 1) understanding the capabilities of the Ambler mechanism, 2) autonomous walking, 3) sample acquisition, and 4) preliminary configuration of a successor to the Ambler. The report lists the members of the research group in Section 6, and their publications in Section 7. Finally, the report includes detailed papers representative of specific areas of research.

2 Understanding Mechanism Capabilities

2.1 Walker/Terrain Interaction

Walker/terrain interaction underlies all of the stability and reliability issues associated with walking on unstructured, natural terrain. In his Ph.D. thesis [66], Peter Nagy employed modeling, simulation, analysis, and experiments to investigate and characterize this interaction in the context of reliable, autonomous walking on natural terrain, where important effects include 1) ground compliance, and 2) supports that might fail due to slope failures, slipping off the edges of rocks, or the like.

Walker/terrain issues that were investigated in the thesis are combined to form a viable walking prescription, where the state of the walker is continuously monitored, and used to characterize the nature of the walker/terrain interaction. One characterization of the

interaction is by walker stability; for this, Nagy developed new stability measures that take into account the effects of compliant terrain. If the interaction is favorable, planned machine motions may be executed by using nominal control, which is normally used to control the walker on natural terrain. If the interaction is unfavorable, feet are repositioned to place the walker in a more favorable stance.

It may be necessary to accept some poor footholds in order for a walker to progress; in these instances force redistribution models—that describe how vertical forces redistribute under a set of compliant feet due to body motion—may be used to ensure the stability of subsequent motions. This conservative method of carrying out walker motions is detailed in Appendix A.

If anomalous conditions arise that may affect the stability of the walker—such as support failures, unexpected foot forces, and low stability—reactive control is employed. The structure of the reactive controller is shown in Figure 1.

Portions of the both reactive and nominal control have been implemented, and the ongoing debugging and progression of the walking prescription is leading to more reliable autonomous locomotion on unstructured terrain.

2.2 Power Consumption

One of the central considerations in developing the Ambler has been energy efficiency (see Appendix D). In 1991, we acquired power consumption data for the current implementation of the Ambler design.

To measure power consumption, we installed a digital power meter on the line between the 208 V supply and the Ambler. We digitized the analog output of the meter at 10 Hz, and synchronized the readings with the real-time robot controller. We then commanded the Ambler to perform various motions, and recorded the sum of the three phases of effective power.

Figure 2 illustrates the power consumed while walking 2 m in four steps on sandy terrain. The figure reveals that circulating a single leg consumes 150 W, and that propelling the body horizontally at 7.5 cm/sec (roughly one-half the maximum velocity) requires 600 W. To our knowledge, this level of propulsive power is unprecedented for a 2500 kg vehicle traversing rough terrain. The figure shows that the steady-state power consumption of the motors, amplifiers, and associated electronics is about 1400 W. Using more efficient components, particularly multiplexor power supplies and servo amplifiers, could substantially reduce this steady-state power draw.

Figure 3 illustrates the power consumed while raising and lowering the body at different rates. The figure shows that lifting the body at 7 cm/s (maximum velocity) consumes approximately 1800 W. Lifting the body, power consumption increases approximately linearly with velocity. Lowering the body, power consumption is constant and equal to the steady-state level, because lost potential energy is dissipated as heat by the amplifiers. As shown

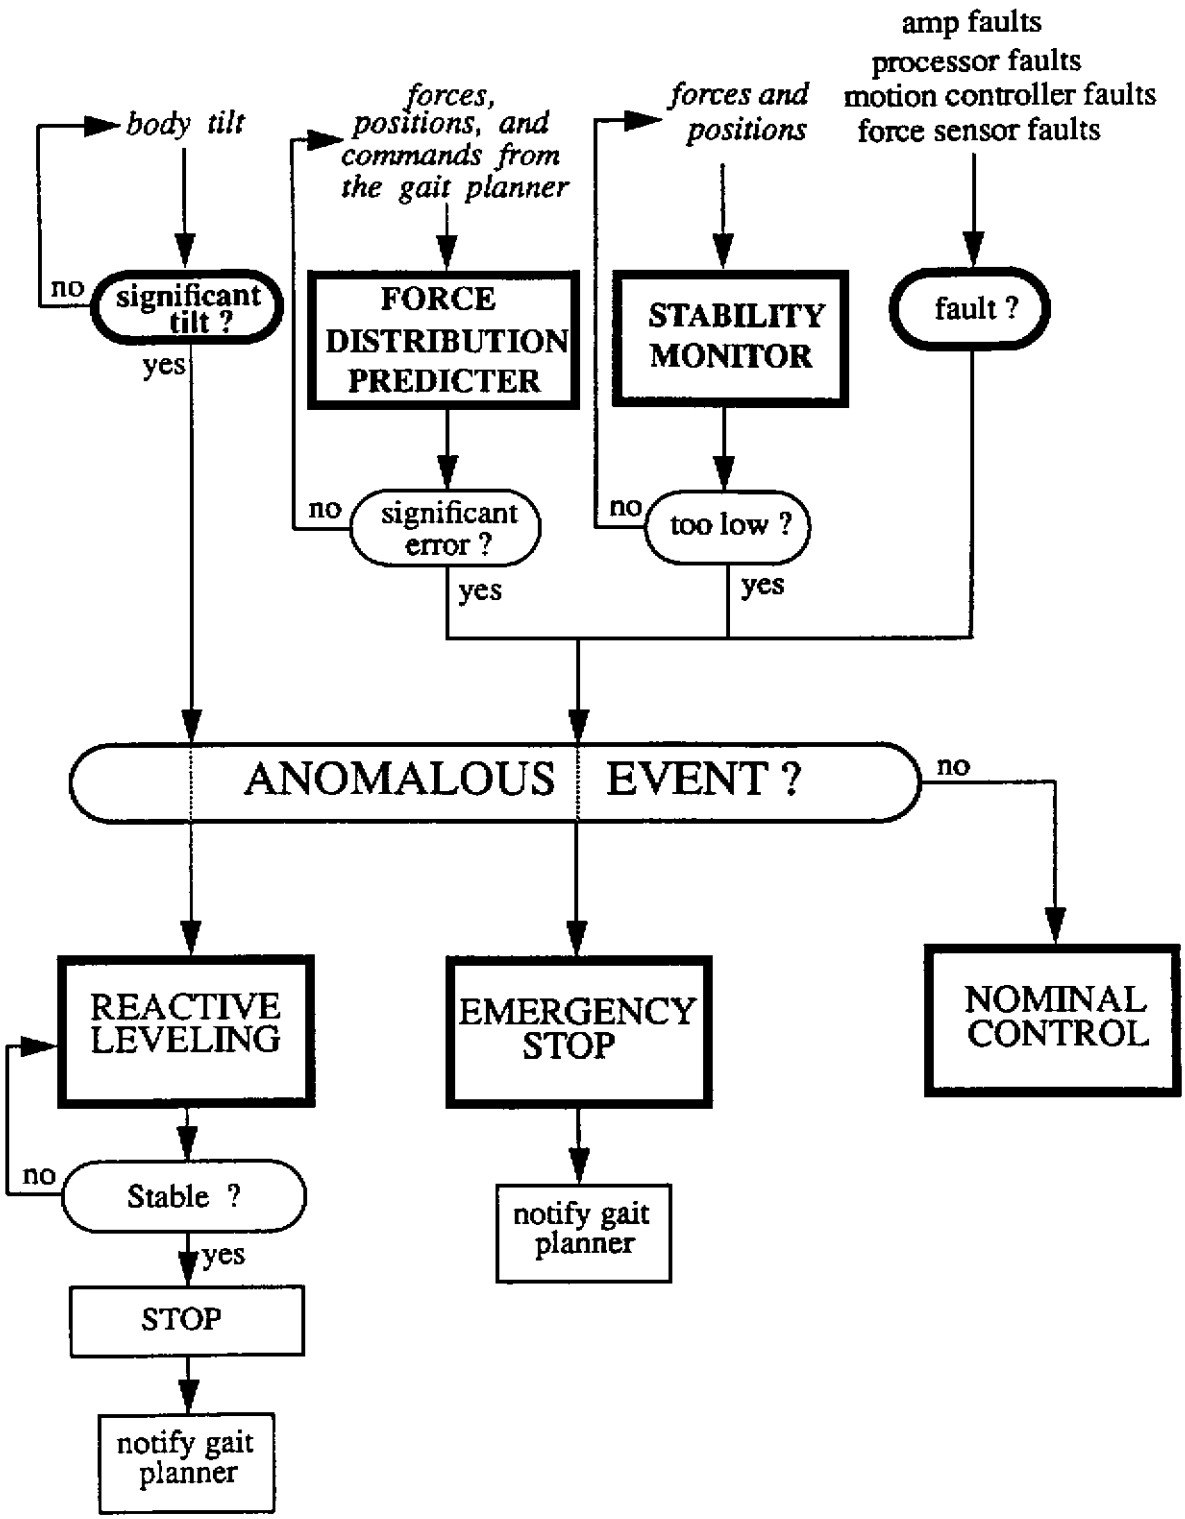

Figure 1: Reactive walking control

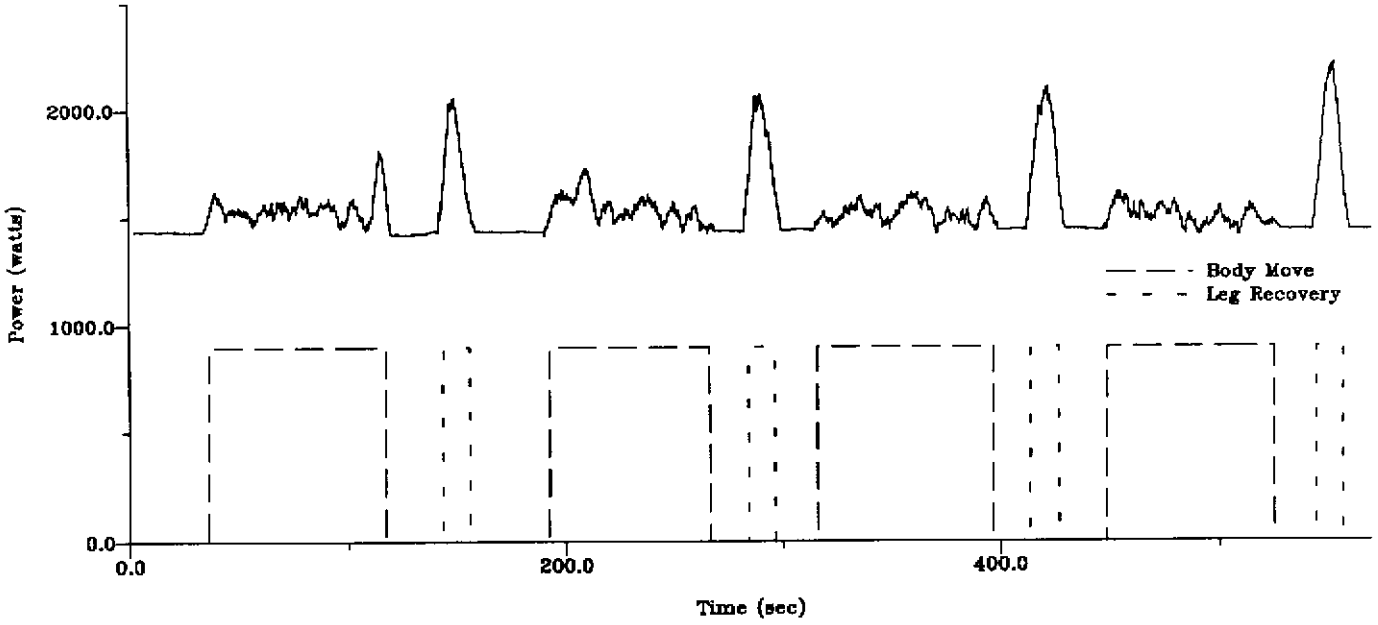

Figure 2: Power consumption during walking cycle

Power consumed while walking 2 m in four steps on sandy terrain. The distance that the body moves is 50 cm. The peak body velocity is 7.5 cm/sec (4.5 m/min).

in Figure 3d–f, the power profile does not depend significantly on acceleration. From the experimental data, we computed the efficiency of the mechanism as the ratio of mechanical output power to measured input electrical power. While lifting the body at 7.5 cm/sec, the efficiency is 70% [3].

2.3 Kinematic Calibration

Other advances in 1991 have refined our models of the Ambler mechanism, and quantified them more exactly. We calibrated separately each of the six Ambler legs, identifying offsets and gear ratios related to the "sag" of the body, and "spread" of the two central shafts under load. The calibration improved the accuracy of dead reckoning from 3.5 cm/step to less than 1 cm/step.

3 Autonomous Walking

Major accomplishments in 1991 include autonomous walking outdoors, autonomous walking at night, and autonomous long-duration walking.

In September of 1991, the Ambler ventured outdoors for the first time, walking in the parking lot adjoining the Planetary Robotics Building (Figure 4). In one afternoon, it took

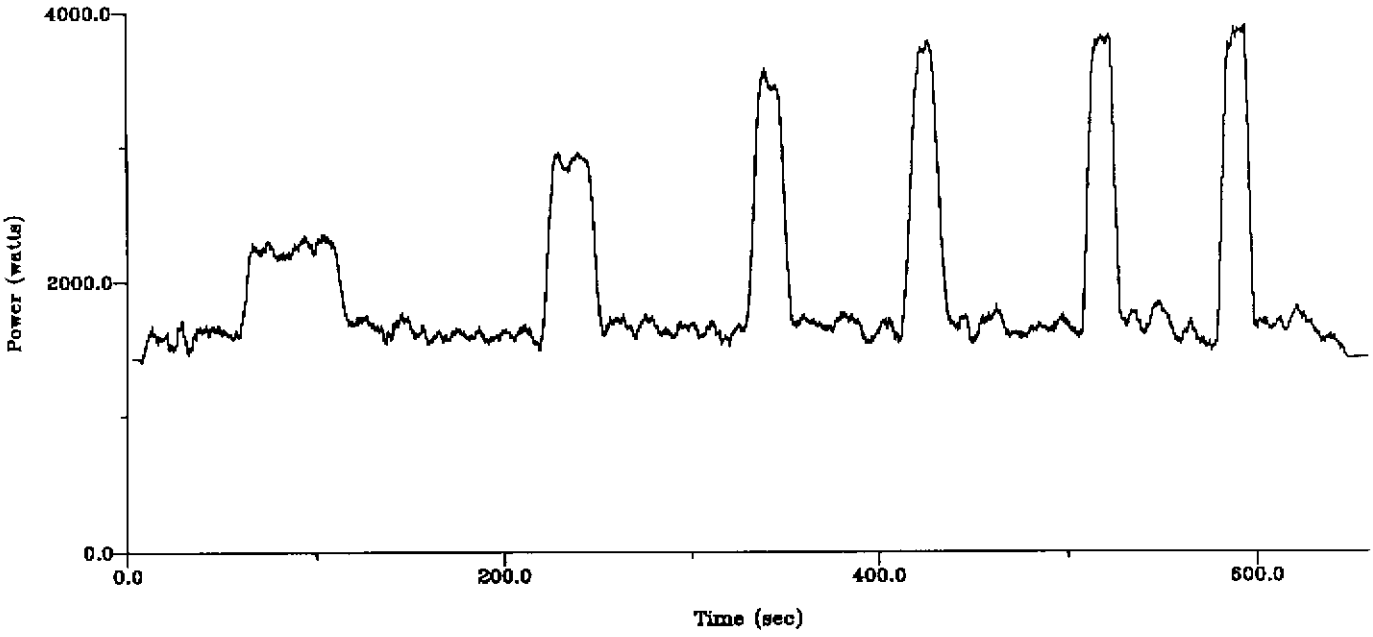

Figure 3: Power consumption during vertical body motion

Power consumed while raising and lowering the body 1 m at different velocities and accelerations: (a) 2 cm/s, 1 cm/s^2, (b) 4 cm/s, 1 cm/s^2, (c) 6 cm/s, 1 cm/s^2, (d) 7 cm/s, 1 cm/s^2, (e) 7 cm/s, 10 cm/s^2, (f) 7 cm/s, 100 cm/s^2.

100 steps along a gently curving arc, traveling about 25 meters over a variety of obstacles, including wooden boxes and ramps. The software system operated for about six consecutive hours.

Another first in 1991 was for the Ambler to walk at night, without lights. The scanning laser rangefinder does not require ambient illumination, unlike the human eye and ordinary cameras. In fact, the laser rangefinder images are sharper at night than during the day, because the signal-to-noise ratio is higher without ambient illumination. With this sensor, the Ambler can operate 24 hours a day, independently of lighting conditions.

During the Fall of 1991 the Ambler set new endurance records walking over an obstacle course—rolling, sandy terrain studded with numerous large boulders and a long wooden ramp—inside the Planetary Robotics Building. It walked a number of "figure-eight" circuits, each covering about 35 meters and 550 degrees of turn. In one three-day trial without being reset, it took 300 steps, traveling about 100 meters.

In addition, the Ambler autonomously crossed over a 1.5 meter tall, 4 meter long boulder. To achieve this, the software system automatically raised the height of the Ambler to near full extension.

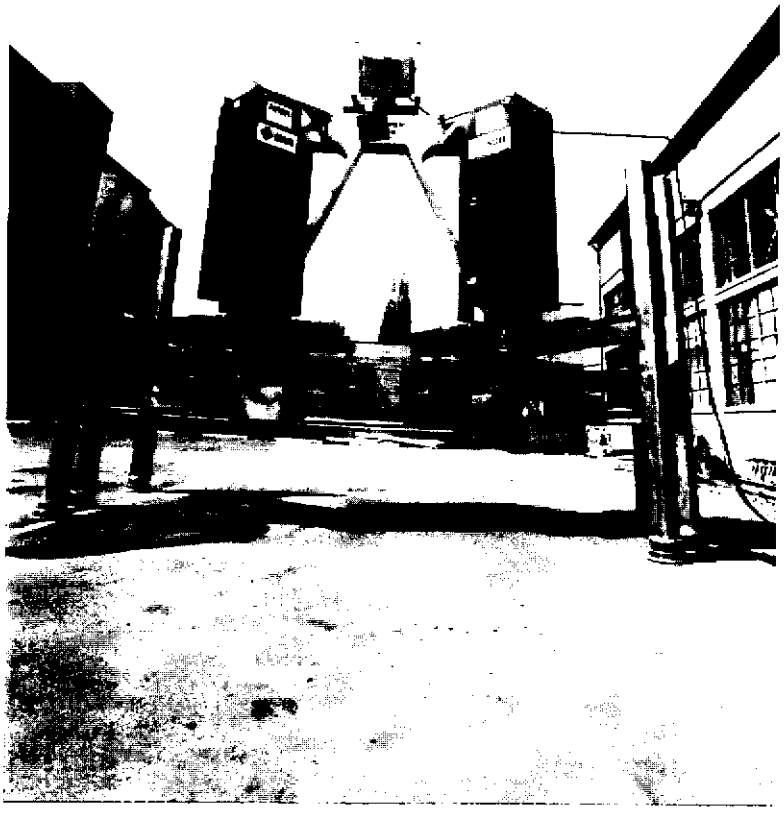

Figure 4: Ambler walking outdoors

3.1 Mechanism and Real-Time Control

In 1991 all computing for perception, planning, and control was installed on-board. Three Sun workstations provide the computing for perception and planning. Two processors and nine motion control boards provide real-time control. As a consequence of this migration, the tether was reduced to a power line and thin ethernet.

Other mechanism changes included installing two new shoulder gears, replacing the signal multiplexing system, changing from analog to digital force sensing, and installing a camera for visual position estimation.

Significant extensions of the real-time controller occurred in 1991. One extension implemented horizontal leg-terrain collision handling, so that significant horizontal forces cause all motions to terminate. Another extension implemented automatic leveling, so that significant tilt (measured by on-board clinometers) cause corrective leveling movements by the Z-axes. A third extension implemented velocity profiling for leg moves, producing smoother and faster motions.

3.2 Perception

In 1991, we directed the manufacturer of the Perceptron scanner to modify the avalanche photo-diode circuits in the device. The changes substantially reduced the range drift due to thermal excitation that caused many problems in 1990 (see Appendix B). In concert with these hardware improvements, we implemented a new terrain mapping system that constructs maps in a reference frame affixed to the Ambler rather than in a fixed reference frame affixed to an arbitrary point. Other developments include supporting video input, managing memory usage, and achieving concurrency by interleaving map computations with image access.

Work on position estimation evolved in three directions during 1991. First, we completed the implementation and evaluation of dead reckoning. Second, we developed a procedure to identify the position and orientation of the Ambler from a single image acquired by a black-and-white CCTV camera mounted on the robot. The approach requires knowing in advance the positions of "landmarks," which were structural features of the Planetary Robotics Building, such as windows and doors. Third, we implemented a position history mechanism that records dead reckoned and visually reckoned position estimates.

In a more theoretical vein, in 1991 we developed an approach to model rough terrain with fractal functions. We implemented algorithms to compute the fractal dimension of terrain viewed with scanning laser rangefinders, and tested them successfully on real Perceptron images.

3.3 Planning

In 1991 we continued to develop the hierarchy of planning algorithms for trajectories, gaits, leg recoveries, and footfalls. For trajectory planning, we implemented a simple technique for choosing arcs to traverse, and incorporated it into the graphical user interface (cf. next section).

For gait planning, we re-modularized the existing planner to streamline the control flow of the walking system and to facilitate incorporation of different planning algorithms (at least three have been developed). In addition, we developed a new constraint analysis method, and used it to implement a simple "crab" (side-to-side) gait plus an algorithm to shuffle the legs into standard configurations while maintaining conservative stability. Finally, we investigated the use of the energy stability margin to plan the most stable body move.

We completed testing and evaluating the neural-net footfall learning algorithm, and incorporated its results in the footfall planning module. See Appendix C for details.

3.4 Task-Level Control

Substantial progress in execution monitoring, error detection, and error recovery was achieved in 1991, thus expanding the Ambler's repertoire of software safety features. Error detection has been implemented for the following conditions:

- Kinematic limits. This required developing an accurate as-built model of the Ambler and using the model to to predict leg/leg and leg/body collisions in real-time.
- Stability problems. This required first developing an accurate model of the Ambler's weight distribution, in order to compute the center of gravity and to evaluate instantaneous stability, and then implementing a stability monitor that watches commands going to the Ambler and rejects those that would cause tipover.
- Unstable footholds. We developed an algorithm that uses a neural net, trained on actual data from the Ambler, to analyze the forces experienced when legs contact the ground.
- Unanticipated terrain collisions.
- Intermittent hardware faults.

Error recovery strategies have been implemented that enable the Ambler to continue after failures. The strategies include the following:

- Retrying leg moves. When an unstable foothold is detected, the Ambler moves the leg in the vicinity until a stable foothold is reached.

- Replanning moves. In concurrent operation, the Ambler plans ahead several leg and body moves. When the time comes to execute a move that was planned some time ago, there may be a difference between the actual and projected body geometry. In this case, the Ambler kills off the planned moves, and restarts planning from the current configuration.

- Lifting the body. When planned motion is impeded by tall obstacles, the Ambler generates and executes plans to elevate itself enough to clear the obstacles.

- Shuffling the legs. When planning cannot proceed (typically because of awkward stances), the Ambler shuffles its feet into a standard pose. The algorithm chooses a sequence of leg moves that maintains stability while shuffling the feet.

In 1991 we continued to address issues of memory usage by the Task Control Architecture (TCA) and by the component software (controller, perception, planning). These efforts to develop sophisticated memory management techniques eliminated memory leaks to the extent that thousands of steps can be taken.

We implemented an initial version of a graphical user interface that provides the ability to specify and edit routes and to view the progress of the Ambler.

We ported the TCA software to the NeXT machine for use by NASA Ames, and provided Ames users with information and debugging help. We also patched various bugs in both the C and LISP versions, enhanced the functionality of the "wiretap" mechanism, and developed direct (point-to-point) communications schemes.

4 Sample Acquisition

We have developed and demonstrated sampling capabilities on a testbed separate from the Ambler. The system demonstrated in the laboratory is now mature and can be used as part of a full-fledged sampling system. Consequently, we have investigated possible configurations for a sampling system on the Ambler. This section summarizes the possible designs that were investigated and reports on the selected design. This design was not fabricated during 1991 for financial reasons.

4.1 Design Considerations

The overall scenario calls for the Ambler navigating through unknown terrain, stopping at specific locations to collect samples. The sampling system that we envision comprises the following components:

- A manipulator: The number of degrees of freedom and the exact configuration are the main parameters. One driving consideration in the selection and placement of

the manipulator is to take advantage of the existing degrees of freedom of the Ambler itself. For example using the Ambler for approximate positioning in the vicinity of the samples, and the manipulator for fine positioning. Two important consequences are that a small reachable area would be acceptable, and that a full six degree-of-freedom design may not be necessary.

- A short-range 3-D sensor: For efficiency reasons, we considered at first using the same sensor as used on the current sampling testbed. It is a short-range (up to two meters) imaging range finder that uses visible light patterns. Using the same sensor allows for using the same perception software as on the testbed. However, the sensor does have some limitations such as relatively large size, small field of view, and small operational range due to the use of conventional cameras and visible light. More efficient laser rangefinders could and should be used; this would not affect the physical configuration significantly.

- Computing: The computing used on the sampling testbed consists of a Sun4-equivalent workstation. In the current scenario, the Ambler would not be performing any other task while sampling, so existing on-board computing could be used to perform sampling operations.

It became apparent very early in the design process that one issue would override any other considerations of power, weight, or speed: the available space on the vehicle. Available space is severely limited for several reasons. First, the vertical clearance of the Ambler, that is the smallest height between terrain and bottom of vehicle body is extremely small, about 6 inches, which limits the size of the equipment that may use space under the body. This limit may be increased by placing the limit switches higher on the leg. This would guarantee the safety of additional equipment while providing more space. However, this solution would reduce the available leg stroke, and therefore would restrict the terrainability of the vehicle. We decided to keep the current vertical clearance. Second, there is little room on the legs themselves to add equipment, because the additional equipment may cause leg collisions during leg recovery.

4.2 Designs Investigated

Several designs were investigated. In this section we give a brief overview of two possible designs and our reasons for not selecting them.

The first, and simplest design, is to mount a gripper on a vertical actuator at the center of one of the bodies (Figure 5). This is a simple design that addresses directly the concerns with vertical clearance. This concept was rejected because the Ambler would have to be used to control the position of the gripper in the horizontal plane.

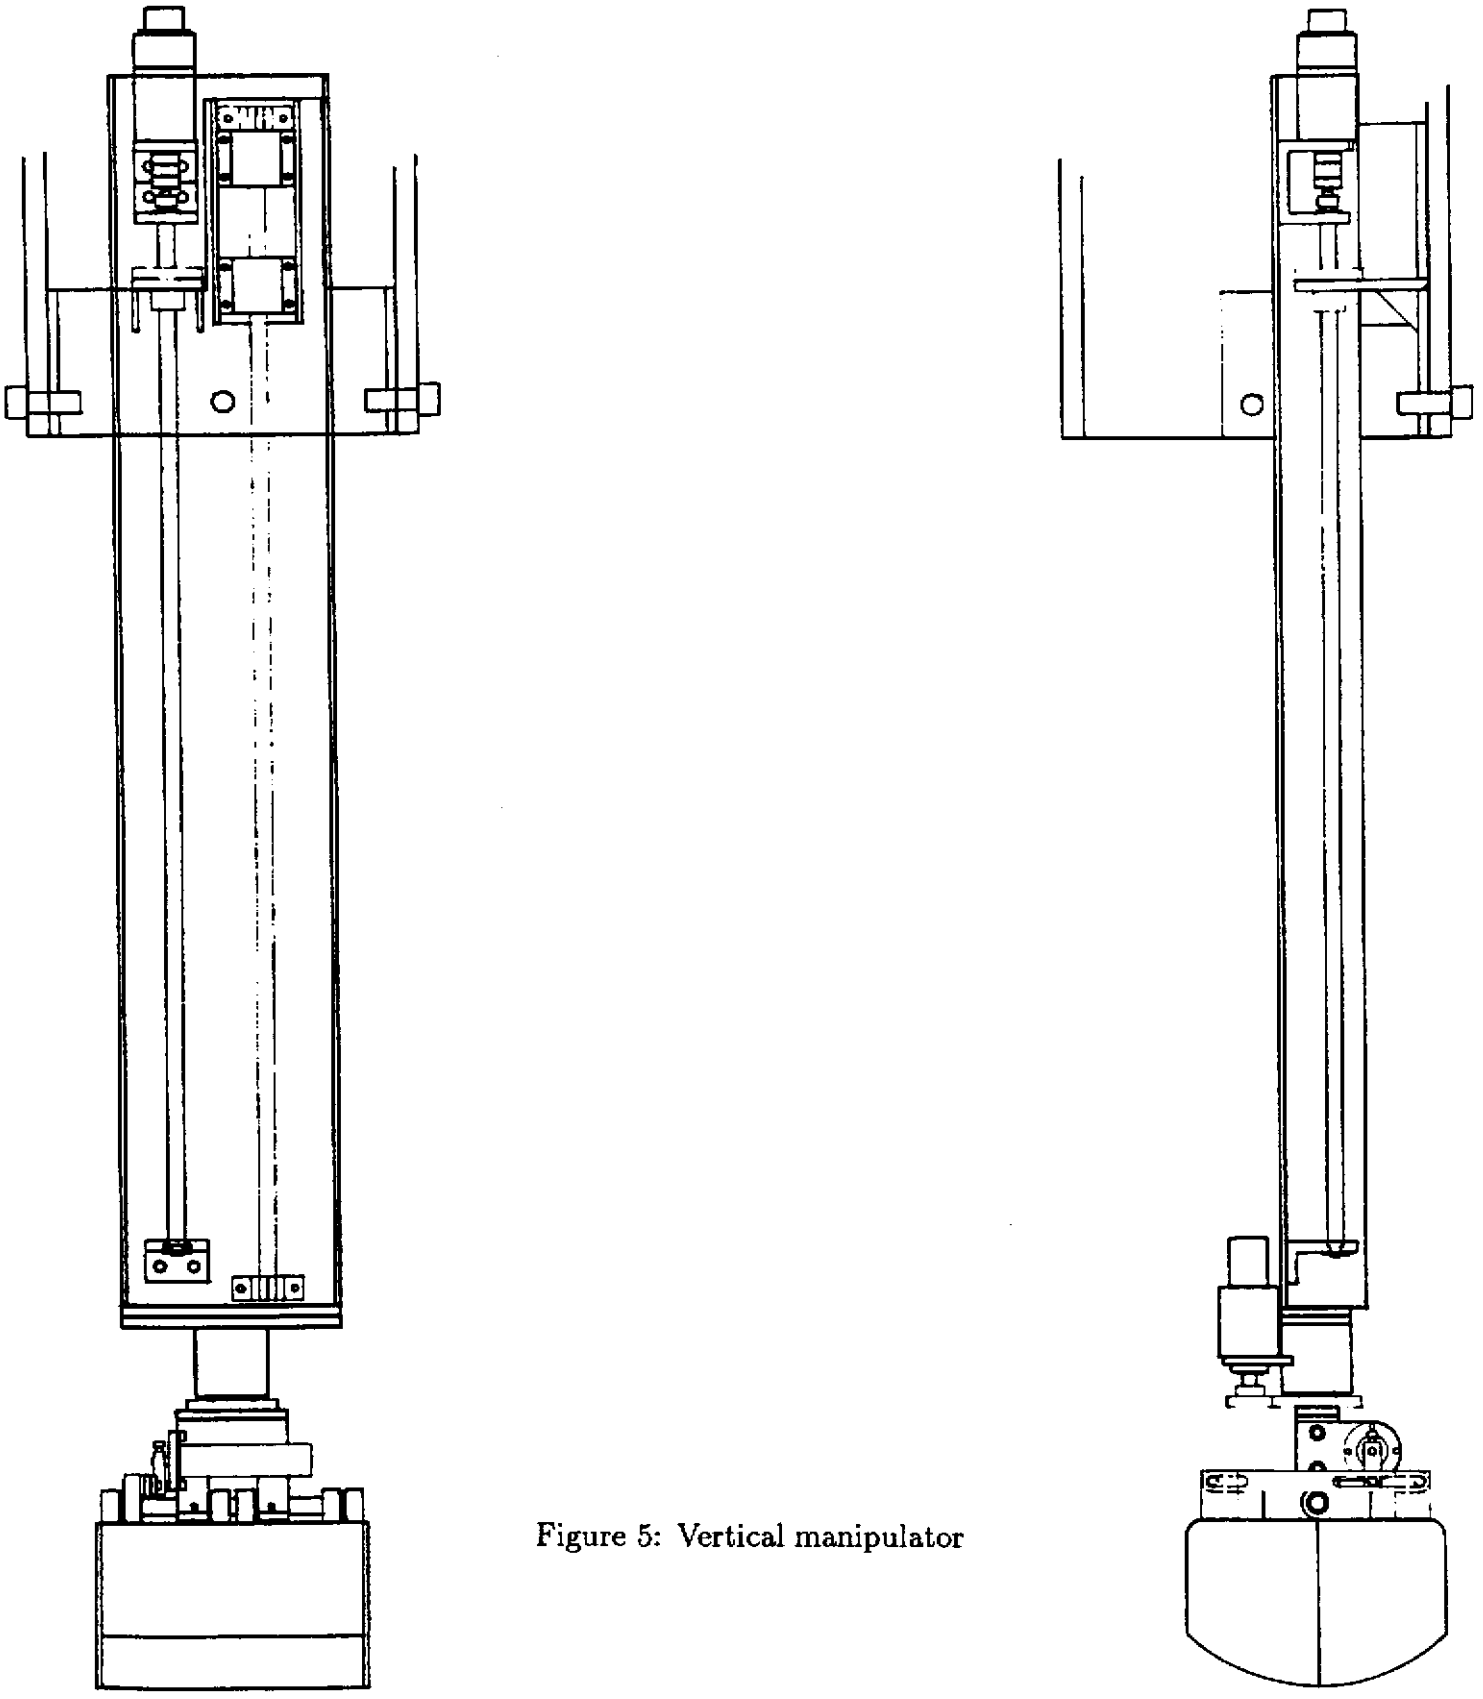

Figure 5: Vertical manipulator

In a second design, a 3-dof arm is mounted in a cylindrical sampling pod under one of the bodies. The arm can be retracted to fit inside the cylinder, depositing the sample in receptacles placed at the circumference of the cylinder. A first problem is the design of a manipulator that can at the same time expand to reach samples, retract to fit within the vertical clearance, and to deposit samples in the receptacles. No off-the-shelf manipulator is appropriate and new components would have to be designed. A second, more serious problem, is the placement of the sensor. Because of vertical clearance, the sensor cannot be placed directly under the body in an orientation such that the intersection of its field of view with the reachable space of the manipulator is large enough. One possible solution is to mount the sensor on one of the legs, while keeping the manipulator under the body. Although attractive, this solution introduced another problem related to the calibration of the sensor with respect to the manipulator. Specifically, since the sensor may move with respect to the arm, the transformation between arm and sensor reference frames has to be recomputed every time the vehicle is in a different configuration.

4.3 Selected Configuration

The configuration that was selected is to attach both sensor and manipulator to the side of the leg. Figure 6 shows top and side views of the configuration. To guarantee that no collisions occur between sampling equipment and other legs, the placement shown in those figures is the only possible one. In the configuration, sensor and arm are fixed with respect to each other, thus avoiding the calibration issue. The other advantage is that the sampling system may be easily placed such that the vertical clearance constraint is satisfied. A receptacle, currently a rectangular bin, is also placed on the leg to receive the collected samples. Based on a trade study, a commercial arm, the CRS-Plus was identified as having the characteristics needed for this configuration, and is the one shown in the figures. The configuration can be designed to optimize the overlap between sensor field of view and arm reachable region while satisfying the other constraints. The optimal configuration is shown in Figure 7 in which the intersection region is shaded.

One drawback of this configuration is that signals from the sensor and the actuators have to be routed to the manipulator controller and the computing equipment. Since the system is mounted on a leg and the controller and computing are in the body, an additional slip ring is required to carry the signals to the racks inside the body. Although adding some complexity to the overall machine, it was determined that commercial slip rings would suffice for the number and type of signals required.

In conclusion, a complete design of a sampling system on the Ambler is now available. The design uses mostly off-the-shelf components and is intended to make maximum use of the techniques developed on the sampling testbed in the laboratory.

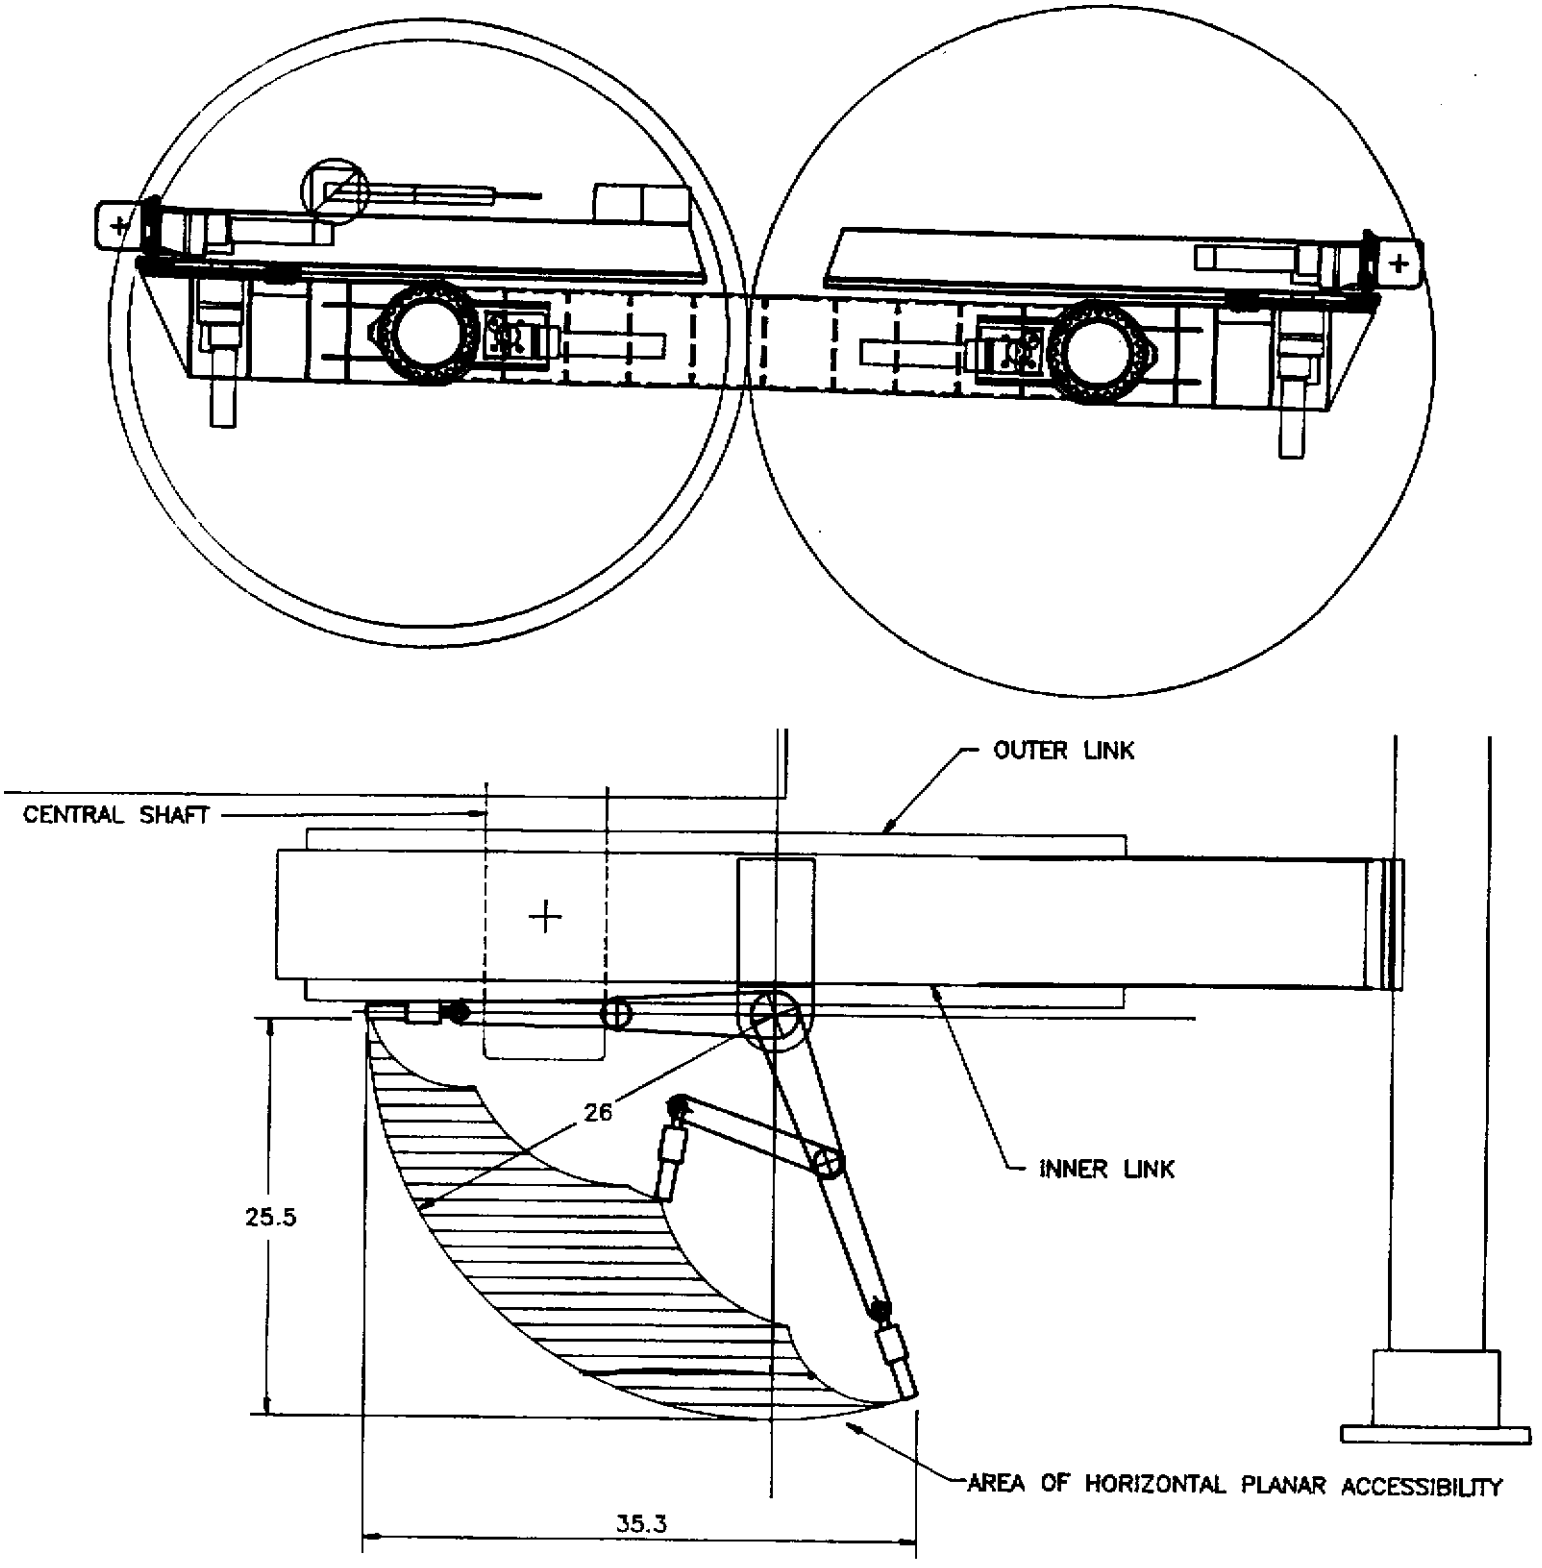

Figure 6: Selected design (top and side views)

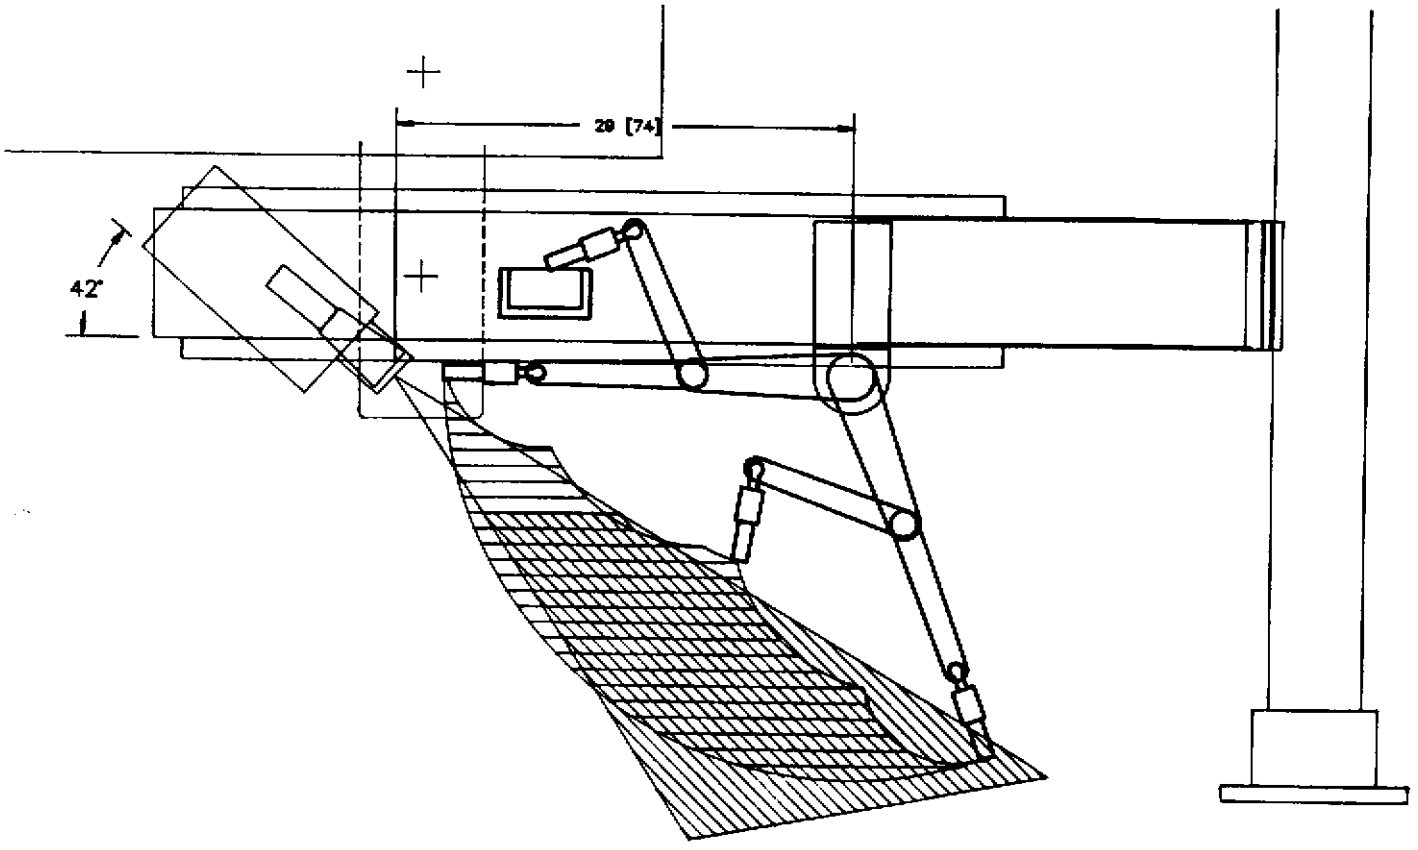

Figure 7: Intersection of field of view and reachable area

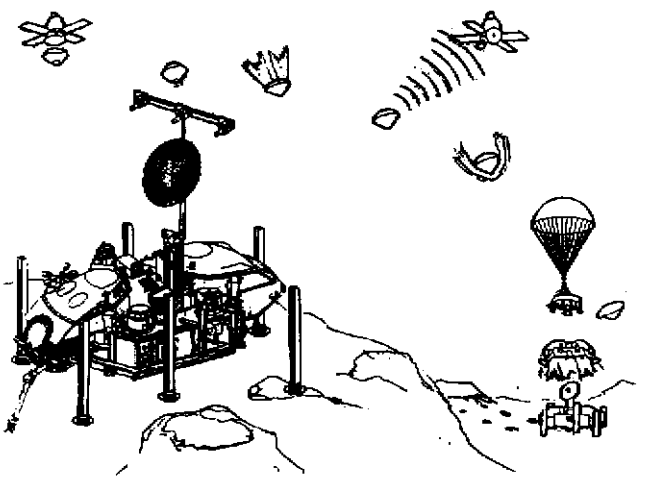

Figure 8: Integrated lander/walker concept

5 Configuration of a Successor to the Ambler

To actualize the unfulfilled potential of articulated machines for roles in space, there is a need to configure a successor to the Ambler. In 1991, we began to configure an integrated lander/walker (Figure 8) for planetary exploration.

We evaluated a number of alternatives, including the Ambler, the Soviet three-cab, the Martin Marietta beam walker, the Viking lander, and the Lunokhod. We identified *minimalism* and *spaceworthiness* as key issues that must play decisive roles in the configuration. We investigated mission specifications and constraints, stowage, scale, power, telemetry, and thermal control. We revisited the model developed in the Ambler research program of perception, planning, and control.

Our preliminary work is incomplete. We expect this topic to be the central focus for the research program in 1992.

6 Personnel

The following personnel were directly supported by the project, or performed related and contributing research in 1991:

Faculty: Martial Hebert, Katsushi Ikeuchi, Takeo Kanade, Eric Krotkov, Tom Mitchell, Reid Simmons, Chuck Thorpe, William Whittaker.

Staff: Brian Albrecht, Purushothaman Balakumar, Gary Baun, Mike Blackwell, Kevin Dowling, Christopher Fedor, Kerien Fitzpatrick, Regis Hoffman, Jim Martin, Clark McDonald, Jim Moody, Dave Pahnos, Henning Pangels, Bill Ross.

Visiting Scientists: Kenichi Arakawa, Herve Delingette, Pablo Gonzalez de Santos, Fabrice Noreils.

Graduate Students: John Bares, Lonnie Chrisman, Richard Goodwin, Goang Tay Hsu, Peter Nagy, Gerry Roston, David Wettergreen.

Undergraduate Students: Steve Baier, Jonathan Burroughs, Doug DeCarlo, John Greer, Nathan Harding, Terry Lim, Hans Thomas.

7 Publications

References

[1] K. Arakawa and E. Krotkov. Estimating Fractal Dimension from Range Images of Natural Terrain. Technical Report CMU-CS-TR-91-156, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania, July 1991.

[2] J. Bares. Orthogonal Legged Walkers for Autonomous Navigation of Rugged Terrain. Ph.D. Thesis Proposal, Department of Civil Engineering, Carnegie Mellon University, December 1988.

[3] J. Bares. *Configuration of Autonomous Walkers for Extreme Terrain.* PhD thesis, Department of Civil Engineering, Carnegie Mellon University, May 1991.

[4] J. Bares, M. Hebert, T. Kanade, E. Krotkov, T. Mitchell, R. Simmons, and W. Whittaker. Ambler: An Autonomous Rover for Planetary Exploration. *IEEE Computer*, 22(6):18–26, June 1989.

[5] J. Bares and W. Whittaker. Configuration of an Autonomous Robot for Mars Exploration. In *Proc. World Robotics Conf. on Robotics Research: The Next Five Years and Beyond*, pages 37–52, Gaithersburg, Maryland, May 1989.

[6] J. Bares and W. Whittaker. Ambler: A Walking Robot for Autonomous Planetary Exploration. In *Proc. International Symposium on Artificial Intelligence, Robotics, and Automation in Space*, Kobe, Japan, November 1990.

[7] J. Bares and W. Whittaker. Walking Robot with a Circulating Gait. In *Proc. IEEE Intl. Workshop on Intelligent Robots and Systems*, pages 809–818, Tsuchiura, Japan, July 1990.

[8] J. Bares and W. Whittaker. Orthogonal Legged Walking Robot. United States Patent. Allowed 1990, Awarded 1991.

[9] C. Caillas. Imaging Sensing to Identify Footfall Positions for a Legged Robot. In *Second Workshop on Military Robotic Applications*, Kingston, Ontario, August 1989. Royal Military College of Canada and Civil Institute of Environmental Medicine.

[10] C. Caillas. Autonomous Robot Using Infrared Thermal Camera to Discriminate Objects in Outdoor Scene. In *Proc. SPIE Conf. Applications of Artificial Intelligence VIII*, Orlando, Florida, April 1990.

[11] C. Caillas. Thermal Imaging for Autonomous Vehicle in Outdoor Scenes. In *Proc. IEEE Intl. Workshop on Intelligent Robots and Systems*, pages 651–658, Tsuchiura, Japan, July 1990.

[12] C. Caillas. Thermal Imaging for Robotic Applications in Outdoor Scenes. Technical Report CMU-RI-TR-90-8, Robotics Institute, Carnegie Mellon University, April 1990.

[13] C. Caillas, M. Hebert, E. Krotkov, I. S. Kweon, and T. Kanade. Methods for Identifying Footfall Positions for a Legged Robot. In *Proc. IEEE Intl. Workshop on Intelligent Robots and Systems*, pages 244–250, Tsukuba, Japan, September 1989.

[14] T. Choi, H. Delingette, M. DeLuise, Y. Hsin, M. Hebert, and K. Ikeuchi. A Perception and Manipulation System for Collecting Rock Samples. In *Proc. NASA Symposium on Space Operations, Applications, and Research*, Albuquerque, New Mexico, June 1990.

[15] L. Chrisman and R. Simmons. Sensible Planning: Focusing Perceptual Attention. In *Proc. AAAI*, Los Angeles, California, July 1991.

[16] H. Delingette, M. Hebert, and K. Ikeuchi. Deformable Surfaces: A Free-Form Shape Representation. In *Proc. SPIE Workshop on Geometric Methods in Computer Vision*, San Diego, California, July 1991.

[17] H. Delingette, M. Hebert, and K. Ikeuchi. Shape Representation and Image Segmentation Using Deformable Surfaces. In *Proc. Conf. Computer Vision and Pattern Recognition*, Hawaii, June 1991.

[18] H. Delingette, M. Hebert, and K. Ikeuchi. Shape Representation and Image Segmentation Using Deformable Surfaces. *Intl. Journal of Image and Vision Computing*, To appear 1992.

[19] C. Fedor, R. Hoffman, G. Roston, and D. Wettergreen. Software Standards and Guidelines. Technical Report PRWP-89-2, Robotics Institute, Carnegie Mellon University, 1989.

[20] C. Fedor and R. Simmons. Task Control Architecture User's Manual. Technical Report PRWP-89-1, Robotics Institute, Carnegie Mellon University, 1989.

[21] C. Francois, M. Hebert, and K. Ikeuchi. A Three-Finger Gripper for Manipulation in Unstructured Environments. In *Proc. IEEE Intl. Conf. Robotics and Automation*, Sacramento, California, April 1991.

[22] P. Gonzalez de Santos, P. Nagy, and W. Whittaker. Leveling of the Ambler Walking Machine: A Comparison of Methods. Technical Report CMU-RI-TR-91-13, Robotics Institute, Carnegie Mellon University, July 1991.

[23] Planetary Rover Group. Experiments in Perception, Planning, and Control for the Carnegie Mellon Mars Rover. Robotics Institute, Carnegie Mellon University, unpublished working document, December 1988.

[24] Planetary Rover Group. Integration of Perception, Planning, and Control in the Carnegie Mellon Mars Rover. Robotics Institute, Carnegie Mellon University, unpublished working document, December 1988.

[25] M. Hebert, T. Kanade, E. Krotkov, and I. S. Kweon. Terrain Mapping for a Roving Planetary Explorer. In *Proc. IEEE Intl. Conf. on Robotics and Automation*, pages 997–1002, Scottsdale, Arizona, May 1989.

[26] M. Hebert, T. Kanade, and I. Kweon. 3-D Vision Techniques for Autonomous Vehicles. Technical Report CMU-RI-TR-88-12, The Robotics Institute, Carnegie Mellon University, 1988.

[27] M. Hebert and E. Krotkov. 3-D Measurements for Imaging Laser Radars: How Good Are They? In *Proc. IEEE Intl. Workshop on Intelligent Robots and Systems*, pages 359–364, Osaka, Japan, November 1991.

[28] M. Hebert and E. Krotkov. 3-D Measurements From Imaging Laser Radars: How Good Are They? *Intl. Journal of Image and Vision Computing*, To appear April 1992.

[29] M. Hebert, E. Krotkov, and T. Kanade. A Perception System for a Planetary Explorer. In *Proc. IEEE Conf. on Decision and Control*, pages 1151–1156, Tampa, Florida, December 1989.

[30] R. Hoffman. Autonomous Legged Robots in Agriculture. In *Proc. ASAE Conf. Automated Agriculture for the 21st Century*, pages 218–225, Chicago, Illinois, December 1991.

[31] R. Hoffman and E. Krotkov. Terrain Roughness Measurement from Elevation Maps. In *Proc. SPIE Conf. on Advances in Intelligent Robotics Systems, Mobile Robots IV*, pages 104–114, Philadelphia, Pennsylvania, November 1989.

[32] R. Hoffman and E. Krotkov. Perception of Rugged Terrain for a Walking Robot: True Confessions and New Directions. In *Proc. IEEE Intl. Workshop on Intelligent Robots and Systems*, pages 1505–1511, Osaka, Japan, November 1991.

[33] G. T. Hsu and R. Simmons. Learning Footfall Evaluation for a Walking Robot. In *Proc. Machine Learning Workshop*, Chicago, Illinois, June 1991.

[34] T. Kanade, T. Mitchell, and W. Whittaker. 1988 Year End Report: Autonomous Planetary Rover at Carnegie Mellon. Technical Report CMU-RI-TR-89-3, Robotics Institute, Carnegie Mellon University, January 1989.

[35] E. Krotkov. Active Perception for Legged Locomotion: Every Step is an Experiment. In *Proc. IEEE Intl. Symposium on Intelligent Control*, pages 227–232, Philadelphia, Pennsylvania, September 1990.

[36] E. Krotkov. Laser Rangefinder Calibration for a Walking Robot. Technical Report CMU-RI-TR-90-30, Robotics Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania, December 1990.

[37] E. Krotkov. Qualitative Perception of Shape and Material. In *Proc. AAAI Workshop on Qualitative Vision*, pages 209–213, Boston, Massachusetts, July 1990.

[38] E. Krotkov. Active Perception of Material and Shape by a Walking Robot. In *Proc. Intl. Conf. Advanced Robotics*, pages 37–42, Pisa, Italy, June 1991.

[39] E. Krotkov. Laser Rangefinder Calibration for a Walking Robot. In *Proc. IEEE Intl. Conf. Robotics and Automation*, pages 2568–2573, Sacramento, California, April 1991.

[40] E. Krotkov. Mapping Rugged Terrain for a Walking Robot. In *Preprints of Proc. Intl. Symp. Experimental Robotics*, Toulouse, France, June 1991.

[41] E. Krotkov. Performance of a Six-Legged Planetary Rover: Power, Positioning, and Autonomous Walking. In *Proc. IEEE Intl. Conf. Robotics and Automation*, Nice, France, To appear, May 1992.

[42] E. Krotkov, J. Bares, M. Hebert, T. Kanade, T. Mitchell, R. Simmons, and W. Whittaker. Design of a Planetary Rover. *1988 Annual Research Review, The Robotics Institute, Carnegie Mellon University*, pages 9–24, 1989.

[43] E. Krotkov, J. Bares, M. Hebert, T. Kanade, T. Mitchell, R. Simmons, and W. Whittaker. Ambler: A Legged Planetary Rover. *1990 Annual Research Review, The Robotics Institute, Carnegie Mellon University*, pages 11–23, 1991.

[44] E. Krotkov, J. Bares, and R. Simmons. Ambler: A Six-Legged Planetary Rover. In *Proc. Intl. Conf. Advanced Robotics*, pages 717–722, Pisa, Italy, June 1991.

[45] E. Krotkov, C. Caillas, M. Hebert, I. S. Kweon, and T. Kanade. First Results in Terrain Mapping for a Roving Planetary Explorer. In *Proc. NASA Conf. on Space Telerobotics*, pages 247–256, Pasadena, California, January 1989. Jet Propulsion Laboratory Publication 89-7, Vol. 2.

[46] E. Krotkov, G. Roston, and R. Simmons. Integrated System for Single Leg Walking. Technical Report PRWP-89-3, Robotics Institute, Carnegie Mellon University, 1989.

[47] E. Krotkov, R. Simmons, and C. Thorpe. Single Leg Walking with Integrated Perception, Planning, and Control. In *Proc. IEEE Intl. Workshop on Intelligent Robots and Systems*, pages 97–102, Tsuchiura, Japan, July 1990.

[48] I. Kweon, R. Hoffman, and E. Krotkov. Experimental Characterization of the Perceptron Laser Rangefinder. Technical Report CMU-RI-TR-91-1, Robotics Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania, January 1991.

[49] I. S. Kweon. Modeling Rugged 3-D Terrain from Multiple Range Images for Outdoor Mobile Robots. Ph. D. Thesis Proposal, School of Computer Science, Carnegie Mellon University, July 1989.

[50] I. S. Kweon, M. Hebert, and T. Kanade. Perception for Rugged Terrain. In *Proc. SPIE Mobile Robots III Conf.*, Cambridge, Massachusetts, November 1988. Society of Photo-Optical Instrumentation Engineers.

[51] I. S. Kweon, M. Hebert, and T. Kanade. Sensor Fusion of Range and Reflectance Data for Outdoor Scene Analysis. In *Proc. NASA Symposium on Space Operations, Applications, and Research, NASA Conf. Publication 3019*, pages 373–382, Dayton, Ohio, July 1988.

[52] I. S. Kweon and T. Kanade. High Resolution Terrain Map from Multiple Sensor Data. In *Proc. IEEE Intl. Workshop on Intelligent Robots and Systems*, pages 127–134, Tsuchiura, Japan, July 1990.

[53] I. S. Kweon and T. Kanade. Extracting Topographic Features for Outdoor Mobile Robots. In *Proc. IEEE Intl. Conf. Robotics and Automation*, pages 1992–1998, Sacramento, California, April 1991.

[54] I.S. Kweon. *Modeling Rugged Terrain by Mobile Robots with Multiple Sensors.* PhD thesis, Robotics Institute, Carnegie Mellon University, January 1991.

[55] L.-J. Lin, T. Mitchell, A. Phillips, and R. Simmons. A Case Study in Robot Exploration. Technical Report CMU-RI-TR-89-1, Robotics Institute, Carnegie Mellon University, January 1989.

[56] L.-J. Lin, R. Simmons, and C. Fedor. Experience with a Task Control Architecture for Mobile Robots. Technical Report CMU-RI-TR-89-29, Robotics Institute, Carnegie Mellon University, 1989.

[57] S. Mahalingam. Terrain Adaptive Gaits for the Ambler. Master's thesis, Department of Mechanical Engineering, University of North Carolina, 1988.

[58] S. Mahalingam and W. Whittaker. Terrain Adaptive Gaits for Walkers with Completely Overlapping Leg Workspaces. In *Proc. Robots 13*, pages 1–14, Gaithersburg, Maryland, May 1989.

[59] D. Manko. Models of Legged Locomotion on Natural Terrain. Ph.D. Thesis Proposal, Department of Civil Engineering, Carnegie Mellon University, June 1988.

[60] D. Manko. *A General Model of Legged Locomotion on Natural Terrain.* PhD thesis, Dept. of Civil Engineering, Carnegie Mellon University, April 1990.

[61] D. Manko. *A General Model of Legged Locomotion on Natural Terrain.* Kluwer, Assinippi Park, Massachusetts, To appear, 1992.

[62] D. Manko and W. Whittaker. Inverse Dynamic Models used for Force Control of Compliant Closed-Chain Mechanisms. In *Proc. ASME Design and Automation Conf.*, pages 61–66, Montreal, August 1989.

[63] D. Manko and W. Whittaker. Planar Abstraction of a Prototype Walking Machine. In *Proc. 20th Pittsburgh Conf. on Modeling and Simulation*, pages 1817–1823, Pittsburgh, Pennsylvania, May 1989.

[64] P. Nagy. Attitude and Altitude Control for a Novel 6-Legged Robot. Research Prospectus, Mechanical Engineering Department, Carnegie Mellon University, January 1989.

[65] P. Nagy. Motion Control for Multi-Legged Walking Vehicles on Rugged Terrain. Ph.D. Thesis Proposal, Dept. of Mechanical Engineering, Carnegie Mellon University, June 1990.

[66] P. Nagy. *An Investigation of Walker/Terrain Interaction.* PhD thesis, Dept. of Mechanical Engineering, Carnegie Mellon University, August 1991.

[67] P. Nagy. Use of Attitude Sensors for Platform Leveling. In *Proc. ASME Winter Annual Meeting, Session on Instrumentation and Components for Mechanical Systems*, pages 39–42, Atlanta, Georgia, December 1991.

[68] P. Nagy, P. Gonzalez de Santos, and W. Whittaker. Reactive Leveling on the Ambler Walking Machine. In *Porc. IFAC Symp. Intelligent Components and Instruments for Control Applications*, Malaga, Spain, To appear, May 1992.

[69] P. Nagy, D. Manko, S. Desa, and W. Whittaker. Simulation of Postural Control for a Walking Robot. In *Proc. IEEE Intl. Conf. Systems Engineering*, pages 324–329, Dayton, Ohio, August 1991.

[70] P. Nagy and W. Whittaker. Experimental Program for the CMU Mars Rover Single Leg Testbed. In *Proc. 20th Pittsburgh Conf. on Modeling and Simulation*, pages 1825–1829, Pittsburgh, Pennsylvania, May 1989.

[71] P. Nagy and W. Whittaker. Motion Control for a Novel Legged Robot. In *Proc. IEEE Symp. Intelligent Control*, pages 2–7, Albany, New York, September 1989.

[72] P. Nagy, W. Whittaker, and S. Desa. A Walking Prescription for Statically-Stable Walkers Based on Walker/Terrain Interaction. In *Proc. Intl. Conf. Robotics and Automation*, Nice, France, To appear, May 1992.

[73] P. Nagy, B. Wu, and K. Dowling. A Testbed for Attitude Control of Walking Robots. In *Proc. ISMM Intl. Symp. Computer Applications in Design, Simulation and Analysis*, pages 120–123, Las Vegas, Nevada, March 1991.

[74] G. Roston and E. Krotkov. Dead Reckoning Navigation for a Six Legged Walking Robot. Technical Report CMU-RI-TR-91-27, Robotics Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania, November 1991.

[75] R. Simmons. An Architecture for Building and Controlling Mobile Robots. In *Proc. Workshop on Domain-Specific Software Architectures*, Hidden Valley, Pennsylvania, July 1990.

[76] R. Simmons. An Architecture for Coordinating Planning, Sensing, and Action. In *Proc. DARPA Workshop on Planning*, San Diego, California, November 1990.

[77] R. Simmons. An Architecture for Coordinating Planning, Sensing, and Action. In *Proc. of DARPA Planning Workshop*, San Diego, California, November 1990.

[78] R. Simmons. Concurrent Planning and Execution for a Walking Robot. Technical Report CMU-RI-TR-90-16, Robotics Institute, Carnegie Mellon University, July 1990.

[79] R. Simmons. Robust Behavior with Limited Resources. In *Proc. AAAI Spring Symposium*, Stanford, California, March 1990.

[80] R. Simmons. Concurrent Planning and Execution for a Walking Robot. In *Proc. IEEE Int. Conf. on Robotics and Automation*, pages 300–305, Sacramento, California, April 1991.

[81] R. Simmons. Coordinating Planning, Perception, and Action for Mobile Robots. *Sigart Bulletin*, 2(4), August 1991.

[82] R. Simmons. Coordinating Planning, Perception, and Action for Mobile Robots. In *Proc. AAAI Spring Symposium on Integrated Architectures*, Stanford, California, March 1991.

[83] R. Simmons and E. Krotkov. Perception, Planning and Control for Walking on Rugged Terrain. In *Proc. NASA Symposium on Space Operations, Applications, and Research*, Albuquerque, New Mexico, June 1990.

[84] R. Simmons and E. Krotkov. An Integrated Walking System for the Ambler Planetary Rover. In *Proc. IEEE Intl. Conf. Robotics and Automation*, pages 2086–2091, Sacramento, California, April 1991.

[85] R. Simmons, E. Krotkov, and J. Bares. A Six-Legged Rover for Planetary Exploration. In *Proc. AIAA Computing in Aerospace 8*, Baltimore, Maryland, October 1991.

[86] R. Simmons, E. Krotkov, and G. Roston. Integrated System for Single Leg Walking. Technical Report CMU-RI-TR-90-15, Robotics Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania, July 1990.

[87] R. Simmons, L.-J. Lin, and C. Fedor. Autonomous Task Control for Mobile Robots. In *Proc. IEEE Intl. Symposium on Intelligent Control*, pages 663–668, Philadelphia, Pennsylvania, September 1990.

[88] R. Simmons and T. Mitchell. A Task Control Architecture for Autonomous Robots. In *Proc. NASA Symposium on Space Operations, Applications, and Research*, Houston, Texas, July 1989.

[89] R. Simmons and T. Mitchell. A Task Control Architecture for Mobile Robots. In *Proc. AAAI Spring Symposium*, Stanford, California, March 1989.

[90] H. Thomas, D. Wettergreen, C. Thorpe, and R. Hoffman. Simulation of the Ambler Environment. In *Proc. 23rd Pittsburgh Conf. Modeling and Simulation*, May 1990.

[91] C. Thorpe. Outdoor Visual Navigation for Autonomous Robots. In *Proc. Intl. Conf. on Intelligent Autonomous Systems*, pages 530–544, Amsterdam, Netherlands, December 1989.

[92] D. Wettergreen, H. Thomas, and C. Thorpe. Planning Strategies for the Ambler Walking Robot. In *Proc. IEEE Intl. Conf. Systems Engineering*, August 1990.

[93] W. Whittaker, T. Kanade, and T. Mitchell. 1989 Year End Report: Autonomous Planetary Rover at Carnegie Mellon. Technical Report CMU-RI-TR-90-4, Robotics Institute, Carnegie Mellon University, February 1990.

[94] W. Whittaker, T. Kanade, and T. Mitchell. 1990 Year End Report: Autonomous Planetary Rover at Carnegie Mellon. Technical Report CMU-RI-TR-91-19, Robotics Institute, Carnegie Mellon University, August 1991.

To appear in Proc. IEEE Intl. Conf. Robotics and Automation, May 1992

A Walking Prescription for Statically-Stable Walkers Based on Walker/Terrain Interaction

Peter V. Nagy[1], William L. Whittaker[1], and Subhas Desa[2]

[1]Field Robotics Center
The Robotics Institute
[2]Department of Mechanical Engineering
Carnegie Mellon University
Pittsburgh, PA 15213, USA

ABSTRACT

Statically-stable walking robots offer advantages over their wheeled and tracked counterparts for autonomous planetary exploration; viz.: enhanced agility, power efficiency, and smooth platform motion. A prescription for how to walk on unknown terrain with an unprecedented degree of reliability is needed for this, and other applications. To fulfill this objective, what is required is an appropriate way to deal with walker/terrain interaction. This interaction underlies all of the stability and reliability issues associated with walking. This work addresses walker/terrain interaction in the context of walking on general terrain in order to achieve stable performance of autonomous walking robots in the face of the multitude of terrain conditions that exist in unstructured, natural terrain.

Walker/terrain interaction is described for nominal motions of the machine. The nature of the interaction is expanded to consider what occurs during support failure, and what may be done to counter its effects. To evaluate the stability of the walker, a new measure has been developed which takes into account the effects of terrain compliance on walker stability. These research elements are combined into a viable prescription for how to walk statically on unstructured terrain.

1. INTRODUCTION

The ability of some walkers to select where their feet contact the ground greatly increases their ability to traverse rugged terrain [1]. Furthermore, selective terrain contact may be used to enhance the walking robot's stability. Discrete terrain contacts ideally lead to lower locomotive power consumption, as no energy is required for rolling friction or "plowing" through terrain. Unless feet slip, legged vehicles only put energy into the terrain when making foot contacts. These advantages are only possible through appropriate control of such walkers, which is the subject of this work. What makes this work difficult is contending with unstructured, natural terrain, while providing absolute reliability.

There is a growing recognition of the importance of accounting for the effects of natural terrain in walking control. For example, many recent works on servo level control of walkers focus on the effects of terrain compliance. Orin's force redistribution control method to minimize power consumption [2] has been supplanted by more efficient methods that emphasize foot/terrain contact to avoid foot slippage [3],[4]. Considering foot/terrain interaction phenomena has increased the state-of-the-art in walking control research. However, we believe that further advances will be made by approaching walking modeling and control from a systemic viewpoint. Each individual foot contact is just part of the entire mechanical system. Therefore, we wish to address interaction phenomena of the entire walker with terrain.

The walker/terrain interaction phenomena and methods of dealing with it addressed in this work are broadly applicable to statically-stable walkers. However, these will primarily be discussed in the context of AMBLER, a hexapod walking machine developed at Carnegie Mellon University [5],[6]. This walking robot was used extensively in the experimental portion of this work.

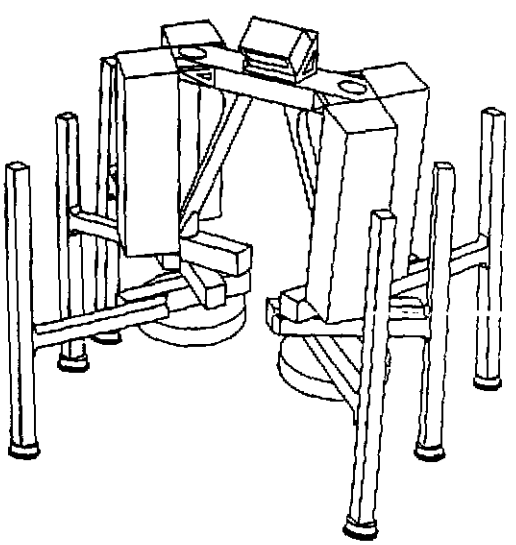

Figure 1. The AMBLER Walking Machine.

The walker/terrain interaction phenomena for nominally controlling a statically-stable walking machine are described in the next section. This is followed by a section describing the interaction and methods developed to avoid tipover due to support failures. A new stability measure that takes into account the effect of compliant natural terrain on stability is then presented. This measure is useful for evaluating planned motions, and as a real-time monitor to preserve the safety of the robot. The algorithms, measures, and knowledge of walker/terrain interaction phenomena are then combined to form a prescription for how to walk on general terrain.

2. WALKER/TERRAIN INTERACTION

Interaction with the terrain occurs during the different elemental walking motions. The elemental walking motions are: leg positioning, propulsion, attitude, altitude, and force redistribution control. Altitude control of the robot does not lead to any significant walker terrain interaction. However, swinging a leg or propelling the body shifts the body c.g., and consequently causes some interaction to occur as detailed later in this section. The method by which attitude control is achieved also will affect the nature of the interaction.

There was little coupling of the elemental walking motions. Forces passively redistribute during motion of the c.g., and body tilt is affected by any redistribution of the vertical foot forces, whether caused by this interaction mechanism, or vertical foot force redistribution control. For the relatively stiff terrain of the AMBLER testbed, this coupling was not too important.

In the control simulations it was observed that simple position and velocity motion control sufficed in achieving stable and accurate motion [8]. The control simulations used a sophisticated dynamic model which included non-linear, non-conservative foot/terrain interaction models [7]. Consequently, motion control boards were procured for the AMBLER, and over several thousands of hours of operation, they have performed well.

For a gravity-decoupled robot, it is possible to decouple most of the elemental walking motions into the vertical and horizontal directions. It was found that peak power consumption could be significantly reduced by carrying out motions that require horizontal and vertical actuations sequentially.

2.1 Attitude Control

Attitude (leveling) control is conventionally achieved by extending and retracting vertical actuators by an amount determined by using a simple calculation that uses small angle approximations:

$$\Delta Z_i = y_i \sin(\alpha) - x_i \sin(\beta) \qquad (1)$$

where: y_i is the horizontal distance of leg_i to the pitch axis

x_i is the horizontal distance of leg_i to the roll axis

α and β are the pitch and roll angles respectively

ΔZ_i is the change in the vertical length of leg_i.

Other leveling methods that use only the vertical actuators but utilize more kinematic information have been derived and evaluated [9]. There is not a large difference between those methods. All of these methods cause the top of the walker to translate, mainly in a horizontal direction. They also lead to foot slippage and/or build-up of internal link forces. Another method was derived that uses all axes to level the body. By using this method, an arbitrary point on the body may be kept from rotating. More importantly, this method avoids foot slippage and build up of internal linkage forces due to flexure. These phenomena occur with z-axes only leveling, and their effects are amplified when footfalls are at disparate elevations. Therefore, the method that uses all axes to level the body should be used when footfalls are at dissimilar elevations. If the peak power consumption needs to be minimized, one of the methods that use only the vertical axes to level the body should be used for gravity-decoupled robots when footfalls are at similar elevations.

2.2 Body Propulsion

During body propulsion there is an active redistribution of forces underneath the feet of a walker, even if the brakes are applied to the vertical actuators. When the machine moves, so does its c.g. location. The foot forces must therefore redistribute to obey the laws of statics. If more than three legs are in ground contact, then the vertical force distribution is indeterminate. There are an infinite number of possible distributions for a given c.g. location, not all of which will be desirable. With foot force sensing, the actual distribution may be found. It is desirable to be able to predict how these forces will redistribute due to a planned machine motion, in order to determine if the planned motion is safe, and if the motion itself is being achieved with predictable walker/ terrain interaction.

Simulations of body propulsions with the dynamic model showed that foot forces changed approximately linearly with a linear change in body position (which roughly corresponds to a linear change in c.g. location). This indicated that a relatively simple model for how foot forces redistribute should be attainable.

Two methods were developed to predict how vertical forces redistributed due to c.g. motion. The *least-variance model* predicts the force distribution after a c.g. move that is as close as possible to the original distribution in a second-norm sense. This method does not utilize terrain/leg compliances. A second method, the *compliance model*, utilizes leg/terrain compliance and kinematic constraints to determine the new force distribution. It can be shown that these two models give exactly the same predictions for when the compliance under each leg is equal [10]. In this paper, only the compliance model will be presented.

For a walker with n ground-contacting legs, n > 3, there are three statics equations that apply. The sum of the vertical forces equals the weight of the machine, and there are no net rolling or pitching moments. These may be expressed mathematically as:

$$F_1 + F_2 + F_3 + \ldots + F_n = W \qquad (2)$$

$$X_1F_1 + X_2F_2 + X_3F_3 + \ldots + X_nF_n = X_cW \qquad (3)$$

$$Y_1F_1 + Y_2F_2 + Y_3F_3 + \ldots + Y_nF_n = Y_cW \qquad (4)$$

where: F_i is the vertical force on leg_i

X_i , Y_i is the location of leg_i

X_c , Y_c is the location of the center of mass

W is the weight of the walker.

Equations (2)-(4) are three equations in n unknowns. A further n-3 equations are derived through kinematic constraints to yield n linearly independent equations in n unknowns, so the unknown foot forces may then be solved. To derive these additional equations, a basis of three legs is chosen. The footpad elevations of these define a plane that will be used as a reference. As foot forces redistribute due to c.g. motion, this plane will move as a function of the change of force on, and the compliance of, each foot. If we consider other ground-contacting feet, they will rise or drop (analogous to unloading or loading a spring under the foot) such that the footpads of these feet also touch the reference plane. Thus a constraint equation is derived for each of the feet outside of the reference tripod (three legs) by separately combining them with the reference tripod, and constraining all four legs to lie on the reference plane. Each of these equations has the form:

$$Y_{ix}F_i + m_{jx}F_j + m_{kx}F_k + m_xF_x = a_x \qquad (5)$$

where: legs i, j, and k form the reference tripod
leg x is the fourth leg under consideration
all m_i are a function of leg coordinates and spring compliances
a_x is a function of leg coordinates, spring compliances, and previous foot forces.

To test the validity of these models, experiments were carried out on the AMBLER. The body was propelled various distances, and the actual foot forces experienced by the AMBLER were compared to the values predicted by the models. For example, the AMBLER was propelled forward by 1 meter on sand, as shown in Figure 2. During this traverse the c.g. moves forward in the Y-coordinate direction by a smaller amount, 0.702 meters, as the legs have significant mass and they don't travel as far as the body during propulsion.

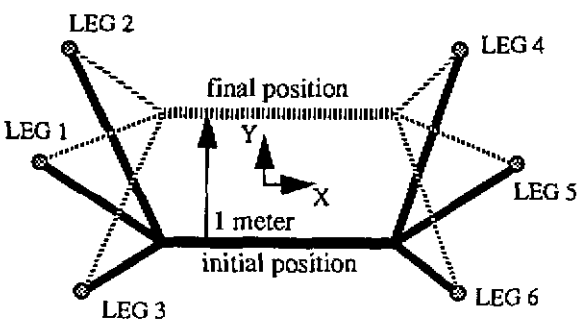

Figure 2. Propulsion Example.

The vertical foot forces logged from the AMBLER are shown in Figure 3(a). The predictions from the compliance model are shown in Figure 3(b). The predictions are close enough that the model is suitable for usage.

The vertical foot force redistribution models give a reasonably accurate portrayal of how the vertical forces redistribute during machine motions. The horizontal forces and bending moments are not incorporated in this model. These forces and moments were monitored during the experiments with the AMBLER. In these experiments they varied, but not in any predictable manner. Fortunately, the effect of these forces and moments was on the order of one to two magnitudes less importance than the vertical force variations. Therefore it was sufficient to consider only the vertical forces in this work.

3. REACTIVE WALKING CONTROL

When anomalous conditions arise, it is desirable to take reflexive measures in order to ensure the safety of the walker. Reactive control is the method by which this safety is assured. For slowly moving, statically-stable walkers, such as the AMBLER, it is sufficient to simply halt the robot should an unexpected event arise. However, if the unexpected event is threatening to tip over the machine, then a more elaborate response is required. In these instances, the legs of the robot should be moved to counteract the tipping motion. The *reactive leveling* algorithm, described below, has been created for this purpose.

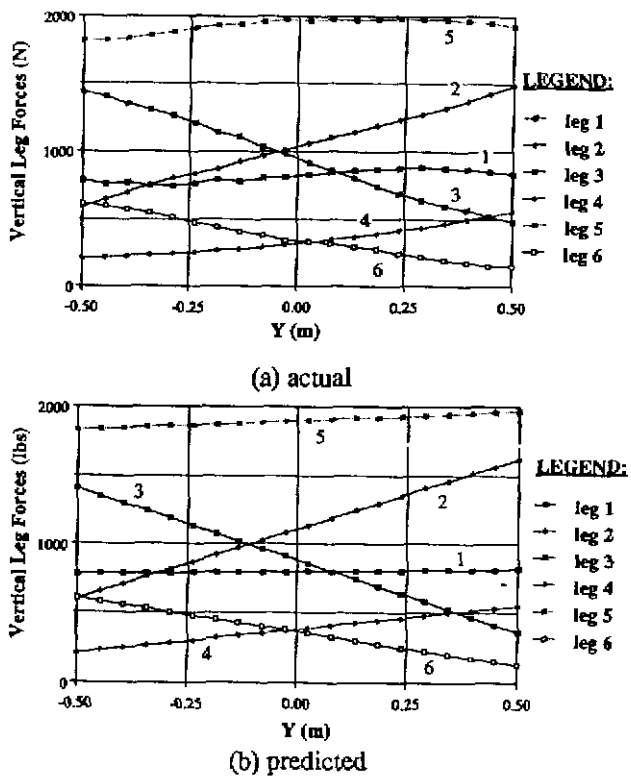

Figure 3. Vertical Force Redistribution Due to Body Motion.

3.1 Reactive Leveling

Slowly-moving statically-stabilized walking machines are in danger of tipping over when support failure(s) occur. The reactive leveling algorithm is able to respond appropriately to these support failures without knowledge of which supports failed, nor by how much. The reactive leveling algorithm is depicted in Figure 4.

Reactive leveling incorporates tilt sensing within the control loop. Given the current roll and pitch of the machine, the vertical leg extensions required to level the machine are calculated. The ratio of the required extensions between legs serves as the ratio of velocity commands that are sent to each leg. The speed that the machine levels is determined by the specified velocity of the leg that moves with the highest speed (the leg with the highest calculated extension). For large values of tilt, this is set to the maximum speed of the leg. For small values of tilt, it is set to zero, so that the machine does not "hunt" about the level position.

The ground-contacting legs will bring the machine up to level with this algorithm. Unless the machine is free-falling, there will always be a sufficient number of ground contacts to form a stable support polygon, which allows the algorithm to succeed. However, support failures that cause tilts usually cause one or two legs to end up in the air. These legs are also on the high side of the machine. As a result, these legs further retract if the reactive leveling algorithm is applied to them, and they will not contribute to body support. This significantly lessens the stability of the machine.

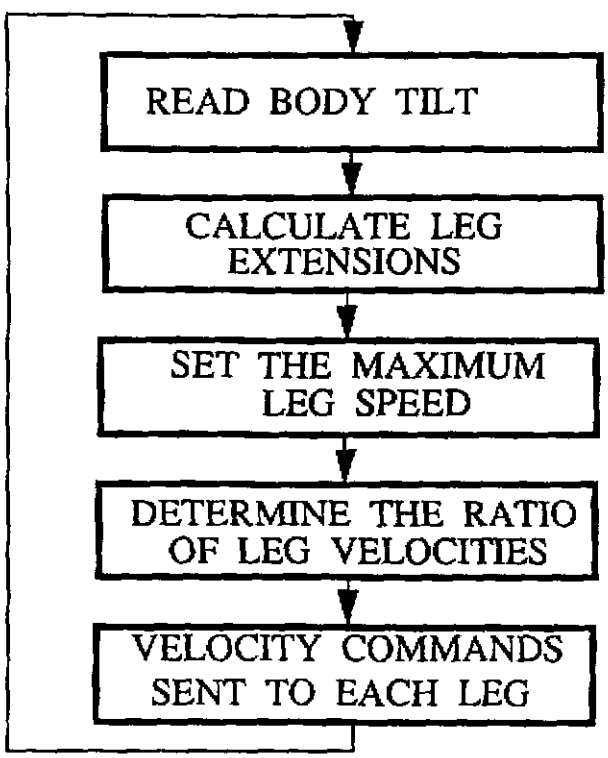

Figure 4. The Reactive Leveling Algorithm.

To maximize stability while still bringing the machine to level with this maneuver, the active leveling algorithm is applied only to the legs that are in ground contact for that control cycle. Legs that are in the air extend slowly, until they contact the ground. After having made contact, they join the subset of legs that are participating in the reactive leveling algorithm.

4. WALKER/TERRAIN INTERACTION AND STABILITY

In the previous sections of this paper there have been a number of references to machine stability. To ensure the safety of the robot, its stability should be quantified. Stability measures may be used in the planning of walker motions, such that the planned motions do not unduly jeopardize the safety of the robot. These measures may also be incorporated as a safety measure incorporated in the real-time controller.

A walker is said to be stable if the vertical projection of its c.g. onto a plane lies inside the polygon formed by the vertical projections of the feet on the same plane. An example of the support polygon for five-legged ground contact is shown in Figure 5. For the walker to remain stable, the projection of its c.g. must lie inside this polygon. For more conservative (safer) walking, the c.g. may be constrained to move above a smaller polygon, the *Conservative Support Polygon* (CSP), which is a subset of the support polygon [11]. If the motion of the c.g. is confined to this smaller area, the machine remains stable even if the support of any one of the legs fails. Our planning algorithms constrain the body c.g. to lie within the CSP.

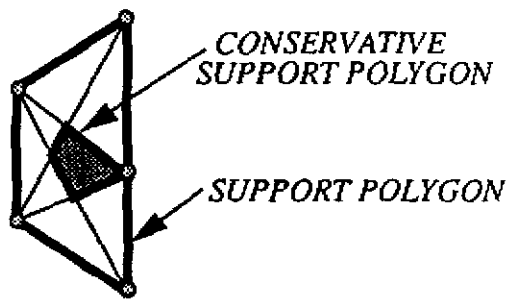

Figure 5. Plan View of the Support Polygons for Five Ground-Contacting Feet.

To quantify the stability, the distance of the c.g. projection to the boundary of either support polygon may be used. However, a better measure is the *Energy Stability Margin* (ESM) developed by Messuri and Klein [12]. This measure calculates the minimum energy required to tip over the walker. This is found by calculating the energy to tip over each pair of adjacent legs of the support polygon. The way to calculate it graphically is shown in Figure 6. To tip the walker over these two legs, the c.g. has to raise by the height h, requiring energy mgh, where m is the mass of the walker, and g is the force of gravity. The analytical determination of the ESM has been derived [10].

This type of measure is better since it quantifies the energy of a disturbance (such as support failure) that is required to topple the robot. However, it does not take into account the compliant effects of natural terrain. We have augmented the ESM by taking this into account by developing the *Compliant Energy Stability Margin* (CESM).

To calculate the Compliant Energy Stability Margin, the ESM for each edge is first calculated. However, compliant footfalls will further compress, as the two legs that the c.g. is swinging over now bear the full weight of the walker. Therefore these feet sink, requiring a smaller height that the c.g. needs to rise for incipient tipover. The geometry of incipient tipover is

shown in Figure 7. The ESM calculates the energy required for the c.g. to reach point j. With compliant footfalls, the c.g. rises by a smaller amount, to point k, in this example. Therefore the CESM predicts smaller stability.

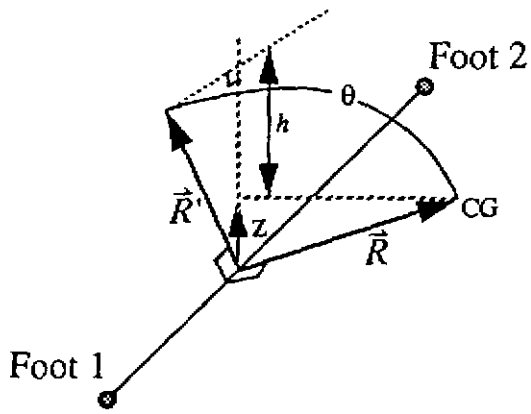

Figure 6. Graphical Calculation of the Energy Stability Margin.

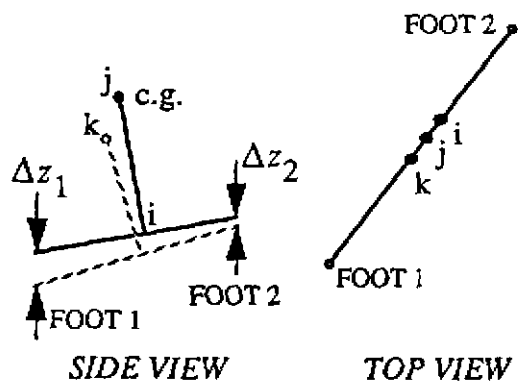

Figure 7. Geometry of Incipient Tipover.

The ESM and CESM measures for the same 1 meter propulsion example discussed in Section 2 are shown in Figure 8. There is no significant difference between these two measures, though the terrain was very stiff. The propulsion example was done on flat terrain. If we superimpose a slope of

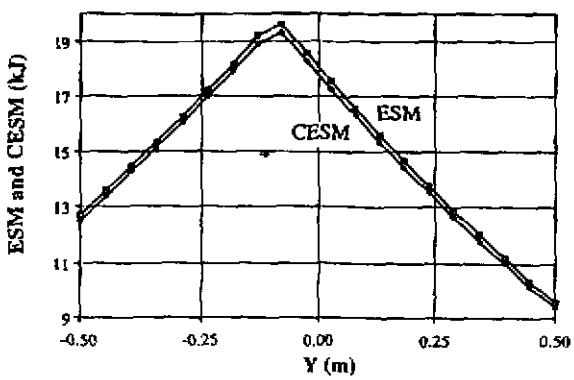

Figure 8. The ESM and CESM for the 1 meter Propulsion Example.

30° on the same propulsion example, the two measures differ by about 13-15%, as shown in Figure 9. This occurs since the compliance of sand on a slope increases significantly. For the flat ground example, the maximum stability is near the middle of the trajectory. With the machine propelling forward horizontally as part of ascending the slope, the point of maximum stability moves forward due to the machine geometry.

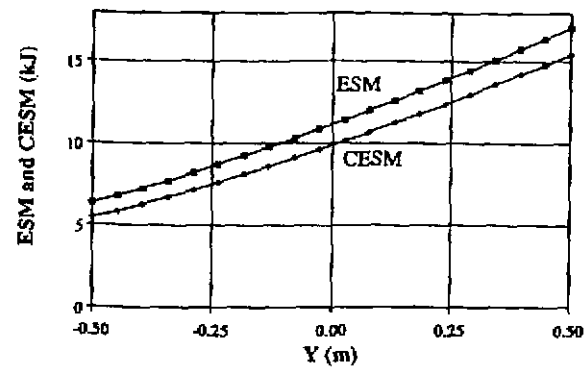

Figure 9. The ESM and CESM for 1 meter Propulsion on a 30° Slope.

5. A PRESCRIPTION FOR WALKING

The various components of important walker/terrain interaction phenomena have been described thus far in this paper. They will now be combined to form a prescription for how to walk on general terrain. To walk on such terrain (e.g. unstructured planetary terrain), the method chosen should be conservative in order to ensure the safety of the robot.

To walk on general terrain while taking interaction phenomena into account, the state of the walker is continuously monitored, and the nature of the walker/terrain interaction is characterized. If the interaction is favorable, planned machine motions may be executed by using nominal control. If the interaction is unfavorable, feet are repositioned to place the walker in a more favorable stance. It may be necessary to accept some poor footholds in order for a walker to progress; in these instances the force redistribution models may be used to ensure the stability of subsequent motions. If anomalous conditions arise that may possibly affect the stability of the walker, reactive control is employed to respond to such events.

For maximum stability, only one leg is repositioned at a time, leaving five legs in ground contact. The walker moves by alternately taking steps and propelling the body. The action of propelling the body and picking up the rear-most leg and placing it in front of the body is called a gait cycle.

Before carrying out a gait cycle, the stability provided by the set of terrain contacts should be evaluated before moving to ensure the safety of the robot. The nominal prescription for carrying out a gait cycle is shown in Figure 10. Before advancing the walker, the controller first checks to see if there are toe-holds. If there are any, the vertical force redistribution that will occur during the gait cycle is predicted. If the machine may carry out the motions of the gait cycle with a stable set of leg forces, the gait cycle is executed. If not, the gait planner is informed of which feet have toe-holds, and it will move these

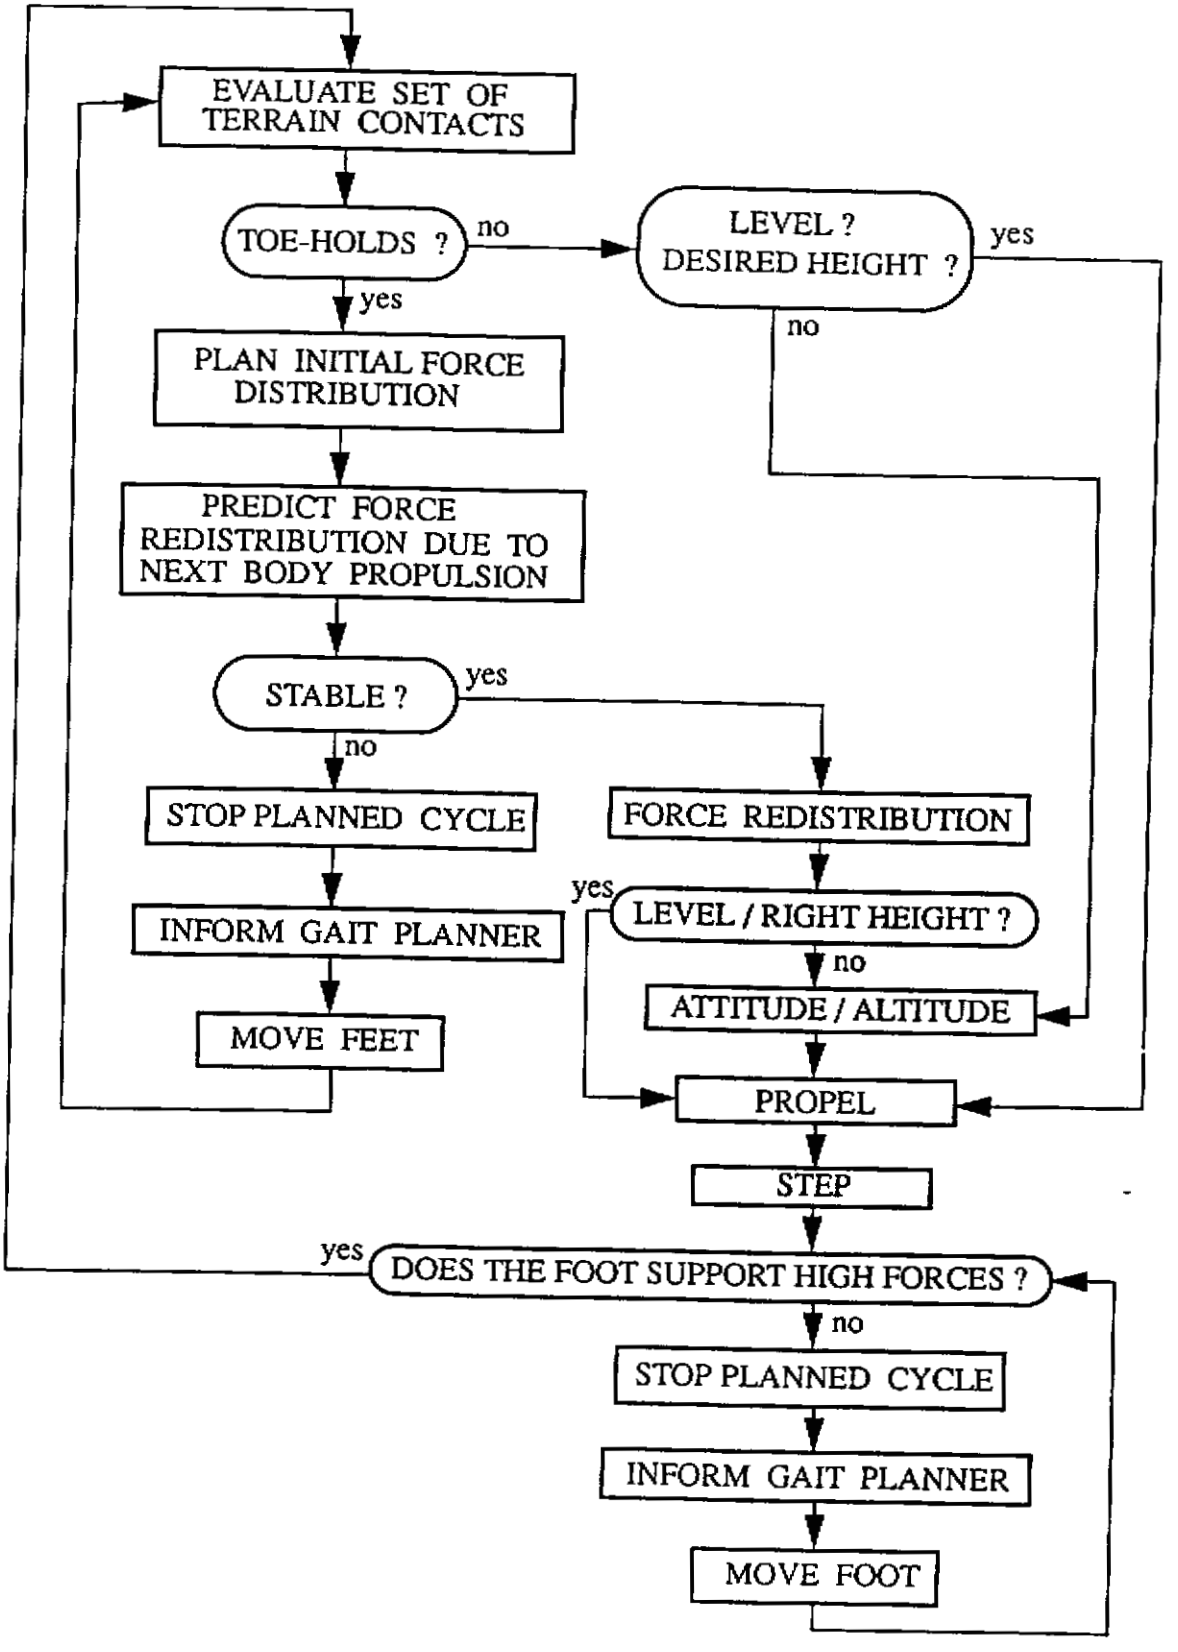

Figure 10. Prescription for Nominal Walking on General Terrain.

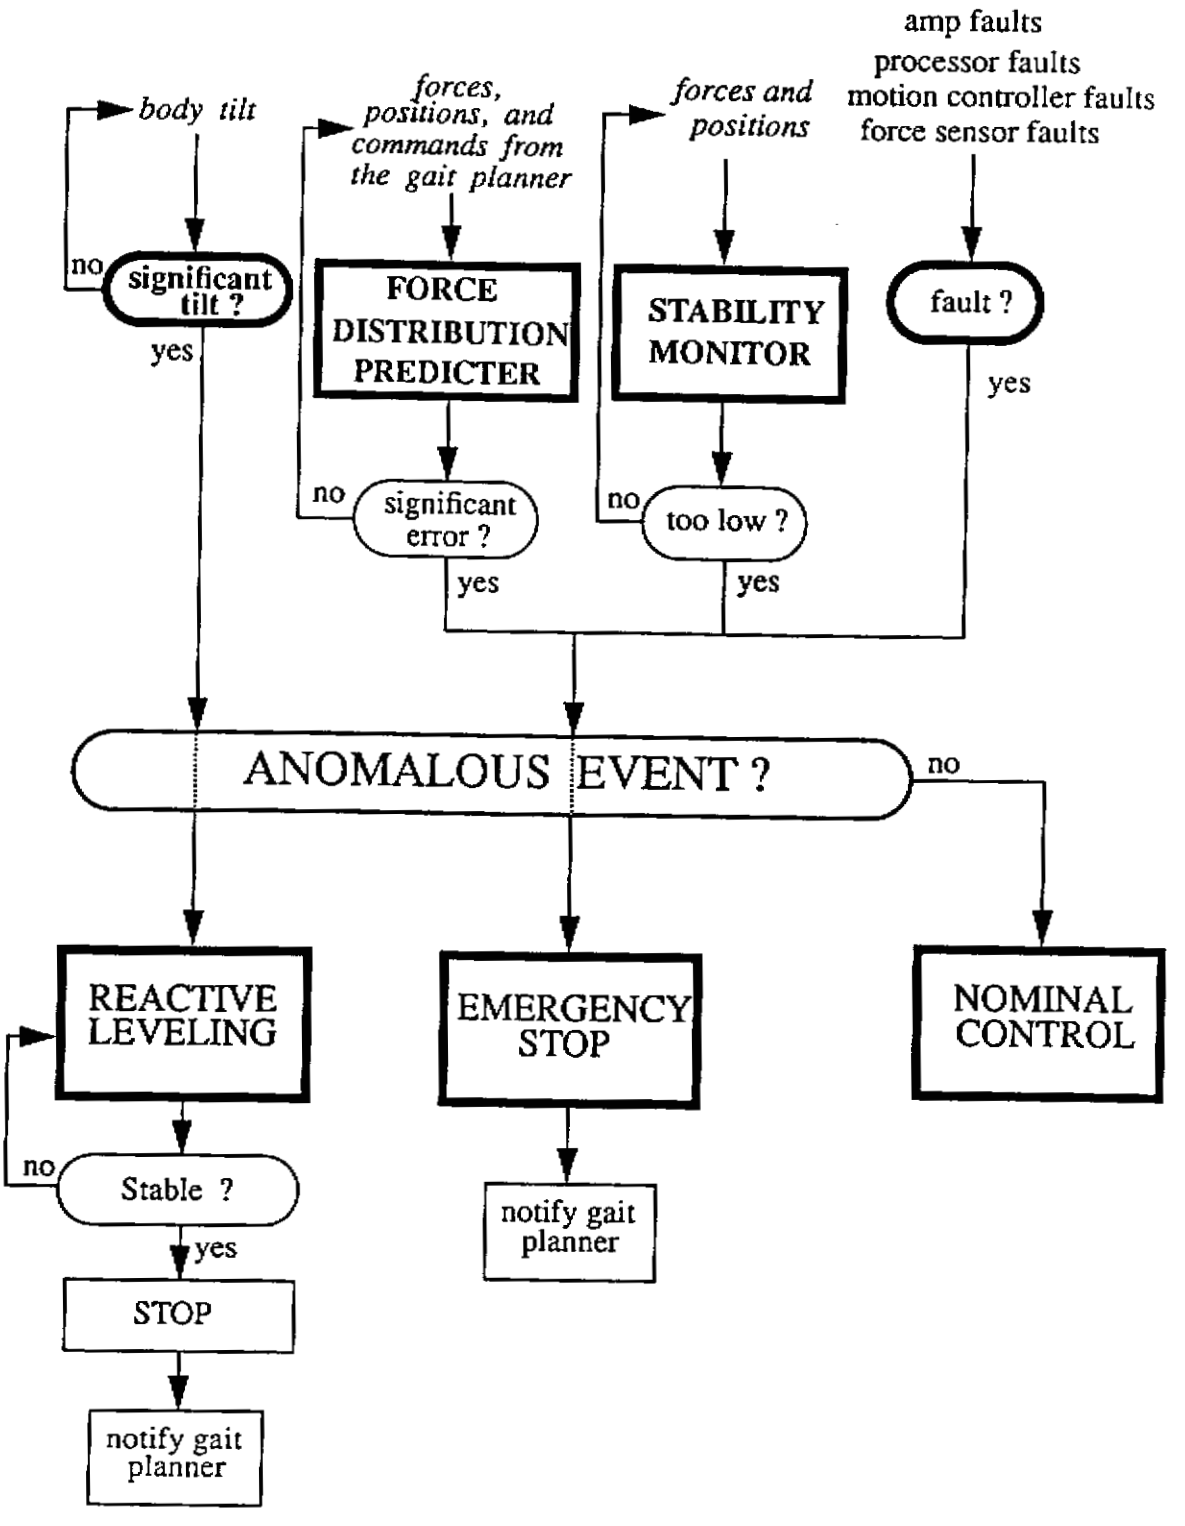

Figure 11. Prescription for a Walking Controller for General Terrain.

feet to new locations. If the feet have to be re-located, the overall stability of the new set of foot contacts is assessed to see if the intended motion may then be carried out safely. When a new footfall is made, the bearing capacity of that foothold is tested by putting a large force on that leg.

The load-bearing capacity test does not preclude the possibility of foot/terrain contact from being unfavorable if it is a toe-hold. By itself, a toe-hold does not endanger the stability of the robot. Some toe-holds form stable ground contacts. Stability of the walker is threatened when a number of toe-holds co-exist in the constellation of ground-contacting feet. If only one footfall is a toe-hold, then vertical forces may be biased away from this foothold, and stability is not a problem. If a new toe-hold is formed that will endanger the robot, the footfall should be changed to a better position if possible. If only toe-holds exist in the footfall selection area, other feet which have toe-holds may then be repositioned.

This nominal method of walking should suffice to allow walking on unstructured terrain, unless anomalous conditions arise. Therefore the walker/terrain interaction and the state of the machine are continuously monitored as shown in Figure 11. Reactive control should be used to preserve the safety of the robot in these instances. Since the machine is walking statically, in most cases it is sufficient to simply halt the machine. If the failure is dynamic, thereby threatening the survival of the robot, then more involved response is necessary.

Static failure events include unexpected foot forces, crossing an allowable stability threshold, and machine failures, such as amp faults, processor faults, motion controller faults and sensor faults. All of these failure modes should be monitored to achieve reliable walking. The foot forces may be monitored in real time, and compared to their predicted values (using the model described in Section 2.2) in order to determine if the walker/terrain interaction is unusual, in which case the machine can be brought to a halt, and a new, stable stance established. The stability of the machine should be monitored using the method described in Section 4 in real-time in order to stop motion if some anomalous walker/terrain interaction lessens the machine's stability.

The dynamic failure that is the greatest concern is when the walker starts to tip over. This may occur due to support failures. The reactive leveling algorithm (see Section 3) was developed for the walker to respond to such an event. This algorithm is designed to bring the body close to level, and when the walker and terrain has stabilized, the walker is halted. A new stance that is favorable is then established, and nominal walking may then continue.

6. CONCLUSION

There is a need for determining the effects of walker/terrain interaction so that autonomous walking robots may traverse unknown terrain in a reliable manner. Walker/terrain interaction phenomena have been characterized through simulation, experimentation, analysis, and use of previous research in the field. Knowledge of these phenomena has led to the formulation of a prescription of how to walk. While this prescription has been posed in the context of AMBLER, it is largely applicable to other statically-stable walking machines. Portions of this prescription have been implemented on the AMBLER. Further implementation, testing, and refining of the prescription will lead to safe autonomous walking on unstructured, natural terrain.

ACKNOWLEGMENTS

This research was sponsored by NASA under Grant NAGW 1175. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of NASA or the US Government.

REFERENCES

[1] J. Bares, et. al., "AMBLER: An Autonomous Rover for Planetary Exploration," IEEE Computer Magazine, June 1989, pp. 18-26.

[2] D.E. Orin, Interactive Control of a Six-Legged Vehicle With Optimization of Both Stability and Energy, Ph.D. Thesis, Ohio State University, 1985.

[3] C.A. Klein and Kittivatcharapong, "Optimal Force Distribution for the Legs of a Walking Machine with Friction Cone Constraints," IEEE Journal of Robotics and Automation, Vol. RA-6, No. 1, Feb. 1990, pp. 73-85.

[4] V. Kumar and K.J. Waldron, "Force Distribution in Walking Vehicles," ASME Journal of Mechanical Design, Vol. 112, March 1990, pp. 90-99.

[5] J. Bares and W. L. Whittaker, "Configuration of an Autonomous Robot for Mars Exploration," Proc. 1989 World Conference on Robotics Research: The next five years and Beyond, Gaithersburg, MD, May 7-11, 1989, pp. 1:37-1:52.

[6] J. Bares and W. L. Whittaker, "Walking Robot with a Circulating Gait," Proc. IEEE International Workshop on Intelligent Robots and Systems, Tsuchiura, Japan, July 1990, pp. 809-818.

[7] D.J. Manko, A Model of Legged Locomotion on Natural Terrain, Ph.D. Thesis, Dept. of Civil Engineering, Carnegie Mellon University, Pittsburgh, PA, April 1990.

[8] P.V. Nagy, et. al., "Simulation of Postural Control for a Walking Robot," Proc. of IEEE Int. Conf. Systems Engineering, Dayton, OH, Aug 1-3, 1991, pp. 324-329.

[9] Gonzalez de Santos, P.V. Nagy, and W.L. Whittaker, "Leveling of the AMBLER Walking Machine: A Comparison of Methods," CMU Robotics Institute Technical Report CMU-RI-TR-91-13, July 1991.

[10] P.V. Nagy, An Investigation of Walker/ Terrain Interaction, Ph.D. Thesis, Dept. of Mechanical Engineering, Carnegie Mellon University, Pittsburgh, PA, August 1991.

[11] S. Mahalingham and W.L. Whittaker, "Terrain Adaptive Gaits for Walkers With Completely Overlapping Leg Workspaces," Proc. of Robots 13, Gaithersburg, MD, May 7-11, 1989, pp. 6:1-6:14, reprinted as SME Technical Paper MS89-298.

[12] D.A. Messuri and C.A. Klein, "Automatic Body Regulation for Maintaining Stability of a Legged Vehicle During Rough Terrain Locomotion," IEEE Journal of Robotics and Automation, Vol. RA-1, No. 3, Sept. 1985, pp. 132-141.

IEEE/RSJ International Workshop on Intelligent
Robots and Systems IROS '91, Nov. 3-5, 1991,
Osaka, Japan. IEEE Cat. No. 91TH0375-6

3-D Measurements From Imaging Laser Radars: How Good Are They?

Martial Hebert and Eric Krotkov

The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

In this paper, we analyze a class of imaging range finders—amplitude-modulated continuous-wave laser radars— in the context of computer vision and robotics. The analysis develops measurement models from the fundamental principles of laser radar operation, and identifies the nature and cause of key problems that plague measurements from this class of sensors. We classify the problems as fundamental (e.g., related to the signal-to-noise ratio), as architectural (e.g., limited by encoding distance by angles in $[0, 2\pi)$), and as artifacts of particular hardware implementations (e.g., insufficient temperature compensation). Experimental results from two different devices—scanning laser rangefinders designed for autonomous navigation—illustrate and support the analysis.

1 Introduction

Range sensing is a crucial component of any autonomous system. It is the only way to provide the system with three-dimensional representations of its environment. The classical computer vision approach to range sensing is to use passive techniques such as stereovision, or shape from X. However, those techniques are not yet sufficiently reliable or fast to be used in many applications, most notably real-time robotic systems. Active sensors, which generate the illumination instead of using only the ambient illumination, have received increasing attention as a viable alternative to passive sensors. Many such sensors were developed [1] and many have been used in real computer vision and robotics applications such as obstacle avoidance [2, 16] and autonomous navigation [9] of mobile robots. Surveys of rangefinding sensors and their applications in robotics can be found in [4, 7, 11]. A review of their use in autonomous navigation of mobile robots can be found in [5].

Although significant theoretical results have been derived in the area of laser radar characterization [15], experimental work is needed to evaluate performance in robotics applications. In this paper, we characterize and evaluate a class of sensors, the imaging laser radars[1]. Those sensors have been proposed as a good compromise between accuracy, resolution, and speed requirements, especially in the context of mobile robotics. Our intent is to present measurement and noise models for those sensors, to identify problems that are specific to this class, and to provide experimental data to support our conclusions. Our emphasis is on identifying limitations and capabilities that have an impact on the use of standard image analysis algorithms for those sensors.

We concentrate on laser radars because of our experience with two such sensors, Erim and Perceptron, in our work in mobile robotics. The theoretical and experimental results presented in this paper use in part results from earlier analyses from [3, 8, 17, 19] for the Erim sensor, and from [10, 14] for the Perceptron sensor. However, we suggest that more analyses of this type are needed in order to better grasp the state of sensor technology from the point of view of computer vision and robotics.

[1] An optical-wavelength radar is also called Lidar, which is an acronym for Light Detection And Ranging.

The paper is organized in three parts. Section 2 describes the principle of the sensors, the theoretical models of noise and measurement geometry, and the two sensors that we use in our experiments. Section 3 describes some problems that are specific to this class of sensors. Those problems can significantly impact the quality of the data and limit the use of those sensors. We distinguish between problems that are inherent to the physics of the laser radars, and problems specific to the hardware currently available for robotics applications. Section 4 describes experimental results obtained from actual sensors.

2 Sensors

In this section we address imaging laser radars, covering both their principles and practical characteristics. First, we describe the general principle of operation, define a sensor reference frame, and present measurement models for the range and intensity data. Then, we describe two particular sensors that we used for experimentation.

2.1 Principle of Operation

The basic principle of a laser radar is to measure the time between transmitting a laser beam and receiving its reflection from a target surface. Three different techniques can be employed to measure the time of flight, which is proportional to the range: pulse detection, which measures the time of flight of discrete pulses; coherent detection, which measures the time of flight indirectly by measuring the beat frequency of a frequency-modulated continuous-wave (fm-cw) emitted beam and its reflection; direct detection, which measures the time of flight indirectly by measuring the shift in phase between an amplitude-modulated continuous-wave (am-cw) emitted beam and its reflection.

Experimental devices have been developed using both pulse and coherent detection technologies (for a survey, see Besl [1]). They are not yet widely in use in computer vision and robotics applications. In this paper, we concentrate on am-cw laser radars.

For am-cw laser radars, the range to a target is proportional to the difference of phase; if $\Delta\varphi$ is the difference of phase, then the range is $r = \frac{\lambda}{4\pi}\Delta\varphi$, where λ is the wavelength of the modulation. Since the phase is defined modulo 2π, the range is defined modulo r_a, where $r_a = \lambda/2$ is the distance (or *ambiguity interval*) corresponding to the maximum phase difference of 2π. Therefore, an inherent limitation of this principle of operation is that it cannot measure range uniquely, i.e., it measures range only within an ambiguity interval.

For many applications, *imaging* laser radars are essential. Imaging sensors generate a dense set of points structured as an image. Typically, image generation is achieved by two mechanically controlled mirrors that raster-scan the beam across a scene, measuring the range at regularly sampled points.

In addition to range, am-cw scanners measure the "strength" of the reflected beam, thus generating a second image that some call the *reflectance* image. To avoid confusion with surface reflectance, we will refer to it as the *intensity* image.

2.2 Sensor Reference Frame

It is useful to convert the range pixels to points in space expressed with respect to a reference frame. In this section we define a

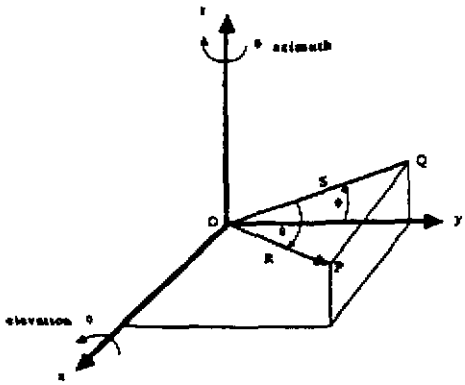

Figure 1: Reference frame for scanning laser range finder In the figure, the symbol "R" denotes range. To be consistent with the majority of the text, we should denote range by "r."

standard reference frame attached to the scanner, and the relation between a pixel (row, column, range) and the coordinates of the corresponding point.

Figure 1 illustrates the reference frame. As shown, the y-axis coincides with the direction of travel of the laser beam projected through the central point of the scanner (i.e., the principal ray). The angle θ (azimuth) corresponds to a rotation about the z-axis. The angle φ (elevation) corresponds to a rotation about the x-axis.

Given the sensor measurement (u, v, d) (i.e., row, column, range), the transformation to spherical-polar coordinates is

$$\varphi = u\Delta_\varphi + \varphi_0 \ , \quad \theta = v\Delta_\theta + \theta_0 \ , \quad r = \Delta_R d \ . \tag{1}$$

where

- Δ_φ is the angular increment, in degrees/row, of the nodding mirror,
- Δ_θ is the angular increment, in degrees/column, of the panning mirror,
- φ_0 is the initial orientation, in degrees, of the nodding mirror,
- θ_0 is the initial orientation, in degrees, of the panning mirror, and
- Δ_R is the scanner range resolution in meters/grey-level.

Given the spherical polar coordinates φ, θ, r, the transformation to Cartesian coordinates is given by

$$x = r\sin\theta \ , \quad y = r\cos\theta\cos\varphi \ , \quad z = r\cos\theta\sin\varphi \ . \tag{2}$$

2.3 Measurement Models

An am-cw range sensor would approach perfection if it emitted a zero-width laser beam and observed the returned signal through an infinitely small receiver. In reality, the beam subtends a non-zero angle, and the receiver detects signals subtended by a solid angle we will call the instantaneous field of view (IFOV). Assuming a circular field stop, the beam projects to an ellipse on the target surface, the *footprint* of the beam. Every point within the intersection of the footprint and the IFOV contributes a range value and an intensity value to the final range and intensity measurements.

We may model a (range, intensity) pair as a complex number z, or equivalently as a vector (Figure 2a). The phase of z represents the sensed range (or phase shift), and the magnitude of z is the sensed intensity. According to this model, the range measured at a pixel is the integral of $k_0 z$ over the IFOV of the receiver, where k_0 is scalar and depends on parameters of the instrument (transmitted power, and the electro-optics of the receiver) and parameters of the environment (the angle α between the surface normal and the direction of measurement, the reflectance ρ of the target surface, and the range r as shown in Figure 2b).

Ultimately, the fidelity of the range measurement depends on the power of the signal reaching the photodetector, which in turn depends on k_0. More precisely, the time-average radiant flux $\overline{F}_P$ (in watts) reaching the photodetector can be shown [15] to be

Figure 2: Model of (range, intensity) pair as complex number

$$\overline{F}_P = k_1 \frac{\rho\cos\alpha}{r^2} \ . \tag{3}$$

where k_1 is a function of the transmitted radiant flux, the capture area of the receiver, and filter bandwidths. Assuming that the output power of the photodetector is proportional to $\overline{F}_P$, the output signal is proportional to $\rho\cos\alpha/r^2$.

When the received power is small, the output signal is small, and noise is significant. There are many sources of noise, including photon noise, laser noise, ambient noise, dark-current noise, secondary emission noise, and subsequent amplifier noise.

Nitzan et al. [12] identify photon noise as the dominant source, and assume that photoemission is a Poisson process to obtain the signal-to-noise ratio

$$\mathrm{SNR} = k_2 \left(\frac{\rho\cos\alpha}{r^2}\right)^{\frac{1}{2}} \ . \tag{4}$$

Assuming a modulation index of 100 percent, they go on to derive the approximate effect of photon noise on the measured range value. Their analysis indicates that the standard deviation of the range can be estimated by

$$\sigma_r \approx \frac{r}{\sqrt{2}\pi\ \mathrm{SNR}} \ . \tag{5}$$

Combining (4) and (5) yields

$$\sigma_r \approx \left(\frac{r^2}{\sqrt{2}\pi k_2}\right)\left(\frac{r^2}{\rho\cos\alpha}\right)^{\frac{1}{2}} \ . \tag{6}$$

where the first term depends only on the physical, optical, and electronic characteristics of the sensor, and the second term depends only on the observed scene.

This equation states that σ_r is approximately linear in r, or equivalently, that the variance σ_r^2 is approximately quadratic in r. Our experimental results follow this model. In particular, Figure 12 shows that σ_r^2 is quadratic in r, consistent with (6).

2.4 Examples

We used two sensors for experiments: Erim and Perceptron, whose geometric parameters are listed below. Other sensors based on the same principle exist [1, 13]. The two sensors that we use are typical of the operation and performance of this type of sensor.

The Erim scanning laser rangefinder is designed for applications in outdoor autonomous navigation. Several versions of the scanner exist. We refer to the version used for research on autonomous land navigation [2, 5], which is the successor of a sensor used for legged locomotion [18].

The Erim scanner uses a 100 mW laser diode operating at 820 nm. The sensor volume is roughly 90 × 50 × 90 cm and the mass is about 45 kg. The scanning mechanism consists of a vertical rotating mirror for horizontal scanning, and a horizontal nodding mirror for vertical scanning. The acquisition rate is 0.5 s for a 64 × 256 × 8 bit image. Table 1 summarizes the characteristics of the sensor. Figure 3 shows a range image (top) and a reflectance image (bottom) from the Erim scanner. The images show an outdoor scene of a road surrounded by a few trees.

The Perceptron scanning laser range finder is more recent than the Erim, and has higher resolution and bandwidth. It is currently used for terrain mapping in support of legged locomotion [6].

The sensor volume is roughly 50×45×35 cm and the mass is about 30 kg. The sensor uses a 180 mW laser diode operating at 810 nm. The image acquisition rate is 0.5 s for a 256 × 256 × 12 bit image. The scanning mechanism is similar to the Erim design, except that the vertical field of view and the vertical image size are

Param.	Units	Erim	Perceptron	Description
γ_{FOV}	deg	30	60	Vertical FOV
θ_{FOV}	deg	80	60	Horizontal FOV
N_{rows}	pixel	64	256	Rows
N_{cols}	pixel	256	256	Columns
Δ_γ	deg	0.47	0.24	Vert. step (Note A)
Δ_θ	deg	0.31	0.24	Hor. step (Note B)
r_a	—	64 ft	40 m	Ambiguity interval
N	bit	8	12	Number of bits/pixel
Δ_R	—	3.0 in	0.98 cm	Range unit (Note C)

Table 1: Nominal values of sensor parameters

Note A: $\Delta_\gamma = \frac{\gamma_{FOV}}{N_{rows}-1}$. B: $\Delta_\theta = \frac{\theta_{FOV}}{N_{cols}-1}$. C: $\Delta_R = \frac{r_a}{2^N-1}$.

Figure 3: Erim intensity (top) and range images
The scene contains a tree (visible to the left), and a person (visible in the upper center) on a path.

programmable. Table 1 summarizes the operating characterisucs of the sensor as specified by the manufacturer [14].

Figure 4 shows the eight most significant bits of a typical range image (right) and intensity image (left) from the Perceptron. The distance between the scanner and the floor is approximately 4 m.

Figure 4: Perceptron intensity (left) and range images
One of the authors is visible in the intensity image standing in front of box-like targets used for calibrating the sensor. The images have been enhanced for printing.

3 Problems

In this section, we describe four effects that may lead to corrupted or degraded data. Two effects, mixed pixels and range/intensity crosstalk, are due to fundamental limitations of am-cw laser radars, although they are sometimes compounded with problems in the design of the actual sensors that we used. The two other effects, distortion due to scanning and range drift, are more specific to the particular sensors that we use. However, it is important to be aware of those effects since they do affect the quality of the data. Furthermore, simple cost-effective remedies do not seem to exist at the moment even though those problems are theoretically avoidable.

3.1 Mixed Pixels

Significant problems occur at pixels that receive reflected energy from two surfaces separated by a large distance. From an image analysis point of view, we would like the range at such points to be measured on either of the surfaces, or at least to fall in between in some predictable way. The fact that the range is measured by integrating over the entire projected spot leads to a phenomenon known as *mixed pixels* in which the measured range can be anywhere along the line of sight. In practical terms, this means that occluding edges of scene objects are unreliable, and that phantom objects may appear due to mixed measurements that are far from the real surfaces. This is a problem inherent to direct detection am-cw laser radars and it cannot be completely eliminated.

Figure 5 shows the geometry of the measurement at an occluding edge: two objects at distances D_1 and D_2 from the sensor ($D_1 < D_2$) are separated by a distance D and a point is measured at the edge between the two surfaces. Due to the angular width of the beam, the range is formed by integration over a spot that contains reflections from both surfaces.The relevant parameter is the ratio p between the spot surface due to the near surface and the total area. The combination of the two measurements is better explained using a model in the complex plane like the one presented in Section 2.3. Each portion i of the spot generates a measurement that can be represented by a complex number z_i. The phase of z_i, φ_i is proportional to the range D_i and its magnitude depends on the reflectivity ρ_i of the material. The resulting measurement is given by the sum $z = z_1 + z_2$. The final measured range is proportional to the phase of z.

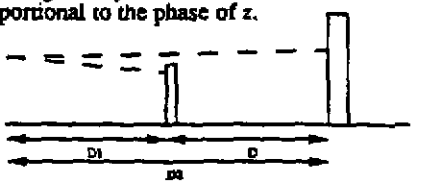

Figure 5: Surface 1 occludes surface 2 (side view)

The phenomenon becomes clear with this formulation: the phase of z can be anywhere between φ_1 and φ_2 depending on the ratio of the lengths of z_1 and z_2 which depends on p, ρ_1, and ρ_2. In practice, this phenomenon is observed as soon as there are discrete objects in the scene. Figure 6 shows the effect of mixed pixels in a real image. The top panel shows an Erim image with a small window (shown as a white rectangle) containing an object (a tree) and its left and right edges. The bottom panel shows an overhead view of the 3-D points in this region as calculated by Eq. (2), and using the range from the image pixels. The points at the center of the distribution correspond to the smooth surface of the object. Away from the center are two lines of mixed pixels that appear at the object's edges. In this example, all mixed pixels are located between the two targets, tree and background surface. This result is typical of the mixing effect in laser radars. The mixed pixel effect complicates the processing and interpretation of the range images. Its main consequence is that strong range edges are unreliable. One approach to the problem is to apply a median filter to the image. Another approach is to transform all the points to a 3-D coordinate frame. There, mixed pixels appear as isolated points, which makes them easier to remove than in image space.

3.2 Scanning Pattern

With currently available commercial technology, imaging can be achieved only by scanning the beam using rotating and nodding mirrors. The scanning mechanism introduces additional errors into the sensor. They are probably the hardest to quantify and to correct. They result in a correct range measurement being stored at the wrong pixel in the image. Limitations due to the scanning mechanism are not inherent to the am-cw technology, but are due to the lack of alternatives to electro-mechanical scanning devices.

The main problem, *synchronization*, is due to the fact that three systems (horizontal mirrors, vertical mirrors, and range measuring system) must be synchronized exactly. In particular,

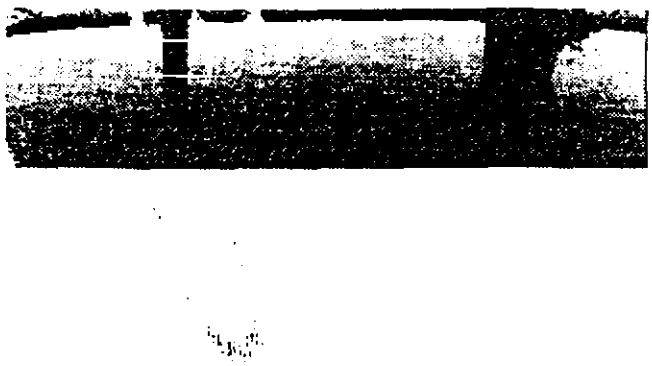

Figure 6: Mixed pixels in real image
Top: selected region in Erim image. Bottom: overhead view of the corresponding points

the motion of the two mirrors must be exactly synchronized with the sampling of the range measurements. A small error in synchronization results in an error in either φ or θ depending on which mirror is affected. Even though the range measurements themselves are not affected, the angular errors will propagate to the coordinates computed from Eq. (2). For example, poor synchronization may occur at the top of the image because the nodding mirror takes a finite amount of time to go from zero speed at its starting position to its normal scanning speed. During this interval of time, a few scanlines that are not correctly sampled are collected. Figure 7 shows a Perceptron image in which a rectangle in the scene projects to a skewed shape instead of a rectangle. In general, there is a discrepancy between the nominal values of the scanning angles and the actual values. This error is difficult to quantify. We describe an experimental setup for estimating the angular error distribution in Section 4.3.

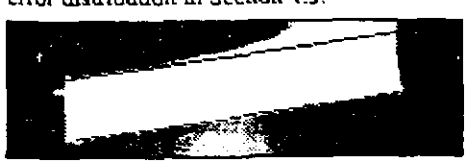

Figure 7: Synchronization error in Perceptron image
The black wire-frame rectangle marks the correct position.

3.3 Range/Intensity Crosstalk

Ideally, we would like the range measurements to be completely independent of the reflective properties of the observed object. Unfortunately, they do influence the range measurements and can even render range useless in some cases. This *crosstalk* effect between range and intensity is due to a number of causes.

The first cause for the crosstalk effect is a fundamental property of direct am-cw range measurement. The standard deviation given by Eq. (6) depends on the reflectance of the observed material. Roughly speaking, the lower the intensity, the higher the range noise. This affects only the variance of the measurement, not its mean value.

Another source of crosstalk is in the implementation of the detection electronics. Typically, the receiver electronics operate optimally only in a narrow range of intensities compared to the large dynamic range of intensities that can be observed. As a result, surfaces that reflect intensities outside of the optimal operating range will produce noisy or even erroneous range measurements. This effect can be reduced by dynamically adjusting the operating range according to the intensity. The Perceptron scanner implements such a solution. However, the low intensities (dropout) and the high intensities (saturation) still produce spurious range readings. There are ways to increase the dynamic range of the receiver but they are not implemented in most scanners available to date.

The crosstalk effect becomes more noticeable at edges or on textured surfaces. In those cases, the beam illuminates a region that contains points of different surface reflectance. Since those points may generate slightly different range measurements, the situation is similar to the one discussed in Section 3.1 except that the mixing is due to the variation of reflectance within the spot instead of the variation of range. A consequence of this effect is that the behavior of the range measurement at the edge between two surfaces with different reflectance may be unpredictable.

The crosstalk problem cannot be completely eliminated. However, some additions to the basic sensor design can diminish its effects. For example, the Perceptron scanner adjusts dynamically the operating range of the receiver, and uses a lookup table built using an off-line calibration procedure to correct the range as a function of intensity.

Figures 8 to 10 illustrate the crosstalk effect. In order to quantify the crosstalk effect, we designed an experiment in which a target with low reflectance is observed against a background of higher reflectance (Figure 8). Considering one scanline in the range image, the dark target is located between columns 121 and 135. We computed the mean and variance of the range and intensity distributions at each pixel in the scanline by taking 100 images of the scene. Figure 9 shows the mean values as a function of the column number. The mean intensity drops sharply at the edges of the black target and remains at a low level between them, as expected. The mean range remains roughly constant except for a sharp discontinuity at each edge. The reason is that the intensity from both materials is mixed at the edges, therefore the range is not properly corrected. Figure 10 shows the variance of the range and intensity distributions. This clearly shows a sharp increase in m^{-2} between the high intensity background and the black target, as expected from the theoretical expression of range noise.

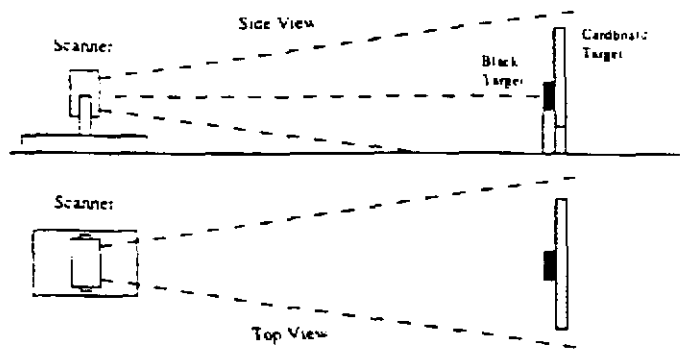

Figure 8: Experimental setup to study range/intensity crosstalk

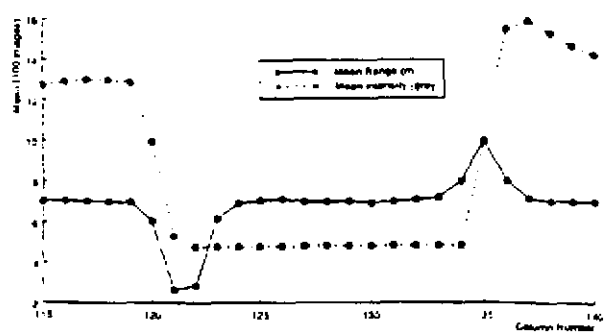

Figure 9: Mean range and intensity for black and white targets

3.4 Range Drift

We observed a significant drift of range measurements over time. To illustrate this effect, we placed a target 6 m from the origin of the Perceptron scanner and acquired one image per minute over 24 hours, during which the scene was static.

Figure 11 plots the sensed range at one target pixel. The

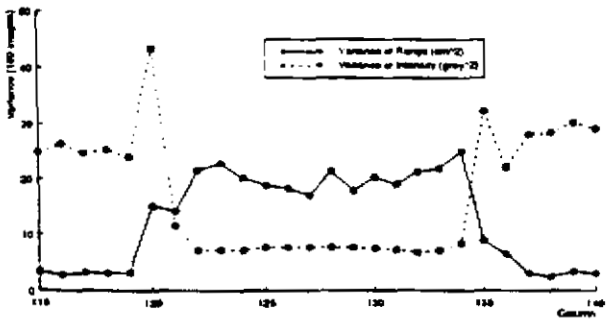

Figure 10: Variance of range and intensity for black and white targets

figure shows a dramatic variation over time. Between hours 0 and 3, the ranges climb approximately one meter, as if the sensor were translating away from the target. After this four hour "warm-up" period, the sensed ranges reach a plateau where they remain, with apparently random variations, for the rest of the day.

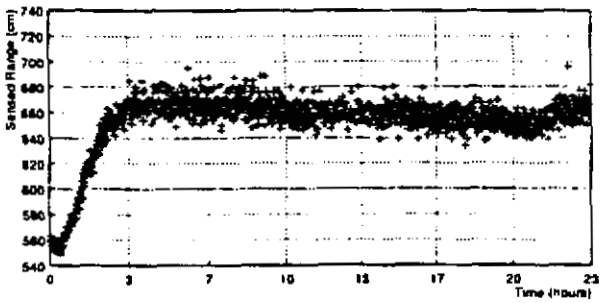

Figure 11: Range measurements vary over time
The target is a sheet of cardboard 6m in front of the scanner.

We hypothesized that some of the variations might be due to temperature changes. To test this hypothesis, we placed an electric heater directly behind the scanner, and repeated the above trial, acquiring images at 2 Hz over two hours (14,000 images). Without the heater, the temperature was $21°C$, and we observed approximately constant range measurements. We turned on the heater and after 30 minutes the temperature climbed to $45°C$. During this time, the sensed ranges fell about 40 cm, a substantial decline. When we turned off the heater, the sensed ranges gradually increased until they regained their original level. This demonstrates conclusively that temperature changes cause the distribution of range values to translate by significant amounts.

It is clear that heat cannot directly affect the phase shift which carries the range information. Therefore, the drift observed in those experiments must be due to poor temperature compensation in the electronics used in the scanner. An important lesson is that such effects may dominate the errors due to the physics of the measurements.

4 Accuracy and Precision

We have introduced a theoretical framework in Section 2.3 that leads to a characterization of expected sensor accuracy for am-cw laser radars. However, it is important to verify that the sensors do indeed follow the theoretical model. In particular, real sensors include effects such as those described in Section 3 that are not part of the theoretical framework. Actual sensor accuracy is important in determining what algorithms and what applications are appropriate for a given sensor. In this section, we describe a series of experiments designed to measure range accuracy for the Erim and Perceptron sensors under different conditions and to compare it with the predicted theoretical values. Following [1], we distinguish between *accuracy*, the difference between measured range and actual range, and *precision*, the variation of measured range to a given target. To separate the errors due to scanning and the errors due to actual range measurements, we distinguish between range precision and angular precision.

	a (deg)	r_0 (m)	Error (m)
Black (Sunny, lights)	45.59	-1.32	0.28
Black (Cloudy, lights)	44.34	-0.73	0.20
Black (Cloudy, no lights)	44.11	-0.81	0.29
Cardboard (Cloudy, lights)	44.92	-1.11	0.12
Wood (Cloudy, lights)	45.57	-1.26	0.09

Table 2: Accuracy results for different targets and lighting conditions
The table shows accuracy results for various combinations of targets (one untreated cardboard slab, one cardboard slab painted black, and a planar piece of wood) and lighting conditions (sunny, cloudy, with and without room lights).

4.1 Accuracy

To determine the accuracy of the range measurements is to identify the distance between them and ground truth ranges. For a target point lying in the direction (φ, θ), let $r_{\varphi,\theta}$ be the range measurement reported by the scanner and let $d_{\varphi,\theta}$ be the true distance from the geometric origin, measured with a tape measure. Under ideal conditions, we expect to observe a linear relationship:

$$r_{\varphi,\theta} = a\, d_{\varphi,\theta} + r_0 \quad . \tag{7}$$

where r_0 is the *offset distance* from the origin to the (conceptual) surface corresponding to a range measurement of zero and a is the slope.

To determine the parameters a and r_0, we acquire range measurements of targets at six known distances between 6 and 16 m, and fit a line to the data. We illustrate the results for the Perceptron scanner in Table 2, which shows the extracted parameters from five trials under different conditions. Because of range drift (cf. Section 3.4), we do not assign high confidence to the particular slope, intercept, and rms error entries in the table.

Nevertheless, the variation with surface material and lighting conditions is obvious. This suggests that the accuracy of the scanner depends significantly on variables in Eq. (7), including surface material, ambient illumination, and temperature. It also suggest that the effect of those variables is amplified by the particular hardware used in those sensors, as described in Section 3. We conclude by remarking that preliminary analysis of this and other data suggests that the accuracy does not depend significantly on target distance.

4.2 Range Precision

To determine the precision of the range measurements is to identify by how much repeated range measurements vary. Here, we quantify the precision as the standard deviation of a distribution of measurements.

We have conducted a number of experiments in which we take 100 images at each target position, and compute the standard deviation of the depth measurements. In the experiments, we have examined how precision changes as a function of ambient illumination conditions, surface material of the target, distance from the scanner to the target, and beam incidence angle at the target. In this section, we report on the effect of ambient illumination for the Perceptron scanner, and refer readers interested in the other properties to Appendix A of [9].

To study the effect of ambient light on sensed range, we place a target (in this experiment, a cardboard slab painted black) at a known distance, take 100 images, and compute the variance in range at particular pixels. We repeat this procedure for six target distances between 6 and 16 m under different indoor lighting conditions.

Figure 12 plots the results, which show that the precision decreases with intensity of illumination. The results strongly support the conclusion that the brighter is the ambient light, the larger are the temporal variations in the range measurements. As in the case of accuracy, above, the effect of the variables in Eq. (6) is clearly visible in those experiments, and it is amplified by the particular hardware implementation used. The results also illustrate the dependence of precision of the square of the target distance (cf. Eq. (6)).

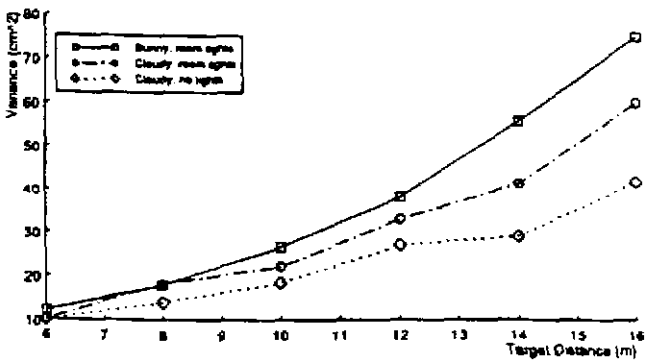

Figure 12: Range variance under different lighting conditions

4.3 Angular Precision

To measure the angular precision of a scanner, or its repeatability in pixel position, we fix the scanner configuration and scene, and take a series of images. We position the scanner in front of a vertical wall on which we have drawn white circles of radius 12 cm, surrounded by black squares. We acquire a sequence of 10^4 images. From each intensity image we extract the white circle by thresholding, and then compute its centroid. We compute the standard deviations of these centroids for several row and column positions.

We find that for both scanners, the standard deviation of the centroid is an order of magnitude smaller than the nominal horizontal and vertical angle increments, varies little over time, and varies little over different pixels in the image. The Perceptron has significantly better angular precision than the Erim.

These findings suggest that the angular precision of the scanners is not a limiting factor. However, the experimental setting—stationary sensor and scene, gathering information over a region—represents a best case, for which the angular precision should be zero. Thus, the findings do not reveal the limitations introduced by relative sensor motion, and do not justify neglecting angular variations as a source of random disturbances in the measurement process.

5 Discussion

In this paper, we have examined in detail the 3-D measurements supplied by amplitude-modulated laser radars. We presented measurement models for this class of sensors, identified problems unique to the technology, and problems specific to the implementation of sensors currently available to the robotics community, presented experimental performance results, and related our practical experience with the sensors.

How good are the three-dimensional measurements? In terms of speed and reliability for medium-range operations, we are not aware of any sensors with superior performance. So the short answer is that they are very good.

The special problems—mixed pixels, crosstalk, deviations from the scanning pattern, and range drift—make it necessary to pre-process the images. For some problems, e.g. mixed pixels, no solutions exist. Thus, range data interpretation algorithms, like so many others in machine perception, must tolerate spurious data.

The quantitative performance, in terms of accuracy and precision, is highly variable. It depends on geometric factors such as incidence angle and target distance. We understand these reasonably well. But it also depends significantly on non-geometric factors such as temperature, ambient illumination, and surface material type. Even though detailed measurement model which include some of those variables have been developed, there is still a significant discrepancy between the theoretical sensor characteristics and the observed performance.

Acknowledgements The authors wish to thank M. Blackwell, R. Hoffman, and I. Kweon who performed many of the experiments described in this paper and kept the scanners in working condition. This work was supported in part by NASA under Grant NAGW-1175 and by DARPA through ARPA No. 4976 monitored by the Air Force Avionics Laboratory under contract F33615-87-C-1499. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the policies of NASA, DARPA, or the US government.

References

[1]P. Besl. Active, Optical Range Imaging Sensors. *Machine Vision and Applications*, 1:127–152, 1988.

[2]R. Dunlay and D. Morgenthaler. Obstacle Detection and Avoidance from Range Data. In *Proc. SPIE Mobile Robots Conference*, Cambridge, Massachusetts, 1986.

[3]Environmental Research Institute of Michigan. *Proc. Range and Reflectance Processing Workshop*, Ann Arbor, Michigan, 1987. Limited distribution.

[4]H. Everett. Survey of Collision Avoidance and Ranging Sensors for Mobile Robots. *Robotics and Autonomous Systems*, 5:5–67, 1989.

[5]M. Hebert, T. Kanade, and I. Kweon. 3-D Vision Techniques for Autonomous Vehicles. Technical Report CMU-RI-TR-88-12, Robotics Institute, Carnegie Mellon University, August 1988.

[6]M. Hebert, E. Krotkov, and T. Kanade. A Perception System for a Planetary Explorer. In *Proc. IEEE Conf. on Decision and Control*, pages 1151–1156, Tampa, December 1989.

[7]R. Jain and A. Jain. *Analysis and Interpretation of Range Images*. Springer-Verlag, New York, 1990.

[8]E. Krotkov. Laser Rangefinder Calibration for a Walking Robot. Technical Report CMU-RI-TR-90-30, Robotics Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania, December 1990.

[9]I. Kweon. *Modeling Rugged Terrain with Multiple Sensors*. PhD thesis, School of Computer Science, Carnegie Mellon University, January 1991.

[10]I. Kweon, R. Hoffman, and E. Krotkov. Experimental Characterization of the Perceptron Laser Rangefinder. Technical Report CMU-RI-TR-91-1, Robotics Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania, January 1991.

[11]D. Nitzan. Assessment of Robotic Sensors. In *Proc. Intl. Conf. Robot Vision and Sensory Controls*, pages 1–11, London, England, April 1981.

[12]D. Nitzan, A. Brain, and R. Duda. The Measurement and Use of Registered Reflectance and Range Data in Scene Analysis. *Proc. IEEE*, 65(2):206–220, February 1977.

[13]Odetics, Inc., Anaheim, California. *3D Laser Imaging System*, 1989.

[14]Perceptron, Inc., Farmington Hills, Michigan. *LIDAR Scanning System*. U.S. Patent Appl. No. 4226-00015.

[15]W. Pratt. *Laser Communication Systems*. Wiley, 1969.

[16]P. Veatch and L. Davis. Efficient Algorithms for Obstacle Detection Using Range Data. *Computer Vision, Graphics, and Image Processing*, 50:50–74, 1990.

[17]R. Watts, F. Pont, and D. Zuk. Characterization of the ERIM/ALV Sensor—Range and Reflectance. Technical Report , ERIM, Ann Arbor, Michigan, 1987.

[18]D. Zuk and M. Delleva. Three-Dimensional Vision System for the ASV. Final report no. 170400-3-f, DARPA, Defense Supply Service-Washington, January 1983.

[19]D. Zuk, F. Pont, R. Franklin, and V. Larrowe. A System for Autonomous Land Navigation. Technical Report IR-85-540, ERIM, Ann Arbor, Michigan, 1985.

Proceedings of the Eighth International Workshop on Machine Learning,
Evanston, IL, 6/27-29, 1991.

Learning Footfall Evaluation for a Walking Robot

Goang-Tay Hsu Reid Simmons
The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

Abstract

The Ambler is a six-legged robot designed to walk in irregular and rugged terrain. Footfall evaluation is the problem of predicting the goodness of footfall locations given the current status of the Ambler and properties of the terrain. We use an inductive learning technique that implicitly correlates terrain features to footfall stability and traction. The learning method also possesses the desirable characteristics of adaptability to different environments, noise tolerance, ability to be trained using relative measures, efficiency and extensibility. This paper describes how a feature-based neural net and a user-specified cost classification scheme can be used to do the evaluation and training. The approach sketched here is applicable to cases where a real-valued function is to be learned from relative measures.

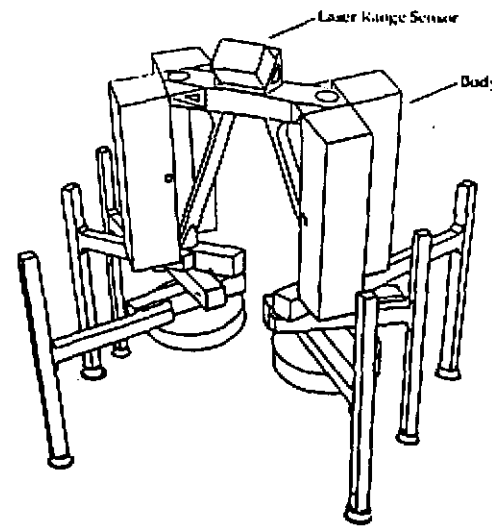

Figure 1: The Ambler

1 INTRODUCTION

The Ambler is a six-legged robot designed for planetary exploration [Bares,1989] (Figure 1). In order to explore new regions and collect samples, the Ambler must traverse unknown and geographically diverse terrain.

The Ambler autonomously navigates terrain using perception to guide its movements [Simmons,1991]. The perception system of the Ambler takes a sequence of laser range images and constructs a geometric terrain representation — the elevation map [Hebert,1989]. The elevation map is a 2D grid in which each pixel value is the elevation of the corresponding location (Figure 2).

The Ambler walks by lifting one leg vertically, swinging it horizontally, extending it down until terrain contact, and then sliding the body forward. The sequence of leg moves is determined by a *gait planning* algorithm which uses geometric constraints (such as the leg reachability and body support) to limit the feasible range of footfall locations. Once a region of geometrically feasible footfalls is selected, the *footfall planner* evaluates and chooses the best possible footfall within that region, given the current status of the Ambler and properties of the terrain [Wettergreen,1990]. Criteria for goodness may include maximal stability, maximal forward travel for the body, minimal power dissipation, best maneuverability, and/or enough traction for the movement. Different criteria may be important depending on the task of the Ambler or its current situation. For simplicity and integrity, the same evaluation function should have the capability of addressing individual criterion by changing the parameters of the function. To achieve real-time performance, efficiency is also a consideration.

In [Caillas,1989], several measures for the goodness of footfall locations from an elevation map are proposed. As indicated in that paper, a single feature is not sufficient to ensure that a footfall position is completely safe. Some form of fusion to combine multiple features is thus required.

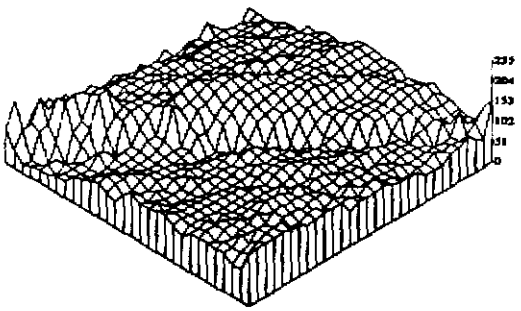

Figure 2: An example of a terrain elevation map

In order to combine features to reliably choose footfall locations, we decided to train a neural net to evaluate the goodness of the terrain. This choice was based mainly on the fact that no explicit model exists relating terrain features to footfall stability and traction. As will be discussed in Section 3, off-line training of the neural net is desirable. For off-line training, people can visually pick out good and bad footfall locations fairly reliably. Typically, however, the user can specify only the relative goodness between two footfall locations, not exact numerical values. An absolute value for the goodness of a footfall location is only obtainable for extreme cases, such as totally flat regions.

Due to sensor noise, the learning method should be tolerant of noise and exhibit graceful degradation. The ability to use feedback from real runs as training data is necessary for adapting to different environments, or for fine tuning to distinguish subtle differences not covered by the off-line training.

Also, when new properties of the terrain or new status of the Ambler are available, the method used should have the ability to incorporate them easily.

Section 2 of this paper describes the method of using a feature-based neural net as the evaluation function. Section 3 describes how users specify training cases and Section 4 describes the way training is done using that specification. Then, in Section 5, experimental results are discussed. Finally, the paper concludes with a plan for future extensions.

2 THE EVALUATION FUNCTION

We can intuitively identify those terrain features that are likely indicators of good footfalls — roughness, and the first and second order change of the terrain. However, there is no direct mapping between terrain features and the desired footfall characteristics of stability and traction, and no single feature suffices to determine the characteristics. It is also very difficult for a human to come up with a reliable combination function.

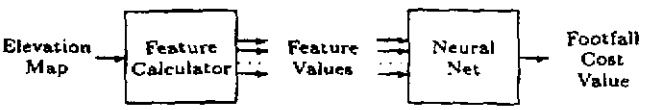

Figure 3: The footfall evaluation function

A feature-based neural net approach is thus proposed as the way of computing the evaluation function (Figure 3). We use a three level feed-forward net with one output node, 8 input nodes, and 8 hidden level nodes. The output of the net is a real value ranging from 0 to 1.

Each feature at a given location is calculated by examining the 5x5 surrounding region of the elevation map. The features currently used include:

- *slope and residual in plane fit* : Fit a plane to the region by least square error approximation. The slope of the plane is then calculated as the tangent of the angle between the fit plane and a referencing horizontal plane. The residual is the average of least square error of the fit.
- *max-min* : The difference between the maximum and the minimum elevation in the region.
- *free volume* : The unoccupied volume located between the foot plane and the surface of the terrain.
- *mean, maximum, and Gaussian curvature* : Fit the region into a quadratic plane $z = f(x, y)$, where z is the elevation. Then, calculate the mean curvature, Gaussian curvature and the maximum curvature directly from the parametric function of the quadratic plane.
- *normal change* : Fit planes to the x, y, and z component of the surface normals respectively. Sum the slopes of the fit planes to get the normal change.

These features are chosen because they approximately model the roughness of the terrain (the residual in plane fit), the first order change (the slope of the fitting plane), the second order change (the various curvatures and normal change), and combinations of these (max-min and free volume).

The computation of these features is fast because all the derivations are analytical and some of the intermediate results (like the surface normal) are shared among features. The complexity of the calculation is basically $O(nm)$, where n is the size of the surrounding region based on which the computation is done (in this case, 25), and m is the number of candidate footfall positions to be evaluated.

3 USER SPECIFIED COST CLASSES

Although we will ultimately train the net using feedback from actual Ambler steps, it is preferable to first train the system off-line on synthetic data. The reason is as follows: First, as the cost of failure for the Ambler is too high, it is more desirable to start with a reasonable footfall evaluation function. Second, training on synthetic data allows us to easily create pathological cases. Third, we can easily train the evaluation function from the conceptually easiest cases to the hardest ones, which may help the learning algorithm converge faster.

One problem with this approach is that the user is not able to give the precise values required to train a neural net. A scheme for the user to specify relative relationships between footfall locations is thus required.

The user classifies the sample terrain as regions of different cost classes (say, "very low", "low", "high", and "very high") using a graphics-based interface. Overlap among regions is allowed to account for inexact classification or uncertainty. Unsure locations are left unclassified.

In addition, to emphasize different aspects of the terrain, multiple classifications can be specified for the same elevation map. This enables the user to do the classification in a hierarchical manner, in which a coarse classification is done, followed by finer classifications to address subtle differences. Thus, Figure 4 shows a region classified as "high" in one case, and then further subdivided on a subsequent training run. This approach not only facilitates the classification but also speeds up the convergence of the neural net.

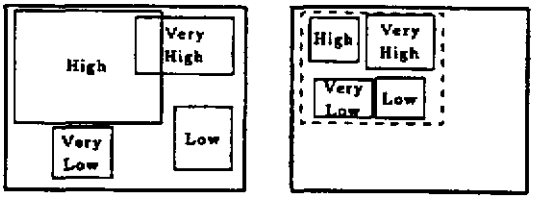

Figure 4: Hierarchical user classification (left: a coarse classification, right: a fine one on High region)

4 THE TRAINING METHOD

As the training data supplied by the user classifications consist of relative measures, two problems must be handled. One is that the output of the net may cluster in a small range of values. The other is that target values have to be estimated in order to make the learning algorithm work.

The first problem can be solved by feeding extreme cases (those with output of exactly 0 or 1) into the net periodically (e.g., every 10 training cases). The conjecture is that the distribution of the output will be anchored by the extreme cases. Experimental results evidence our supposition.

The second problem can be solved by estimating the target value from current output of the net. Pairs of locations from different cost classes are compared to see if the neural net evaluations of each pair are in the relative order specified by the user. Out-of-order pairs are selected as the training cases. The target value for a training case is then formed as the mean value of the current goodness of the pair plus (or minus) some offset (which affects the learning rate). The offset will increase the difference between the target value and current outcome of the net. The use of either a mean or the value of the other location in the pair will not make much difference. The use of an offset, however, is significant to the success of learning. As the neural net learning is incremental, the values of the out-of-order pair tend to approach each other gradually, rather than just be switched at once. When the values of the pair are too close to each other, the offset can improve not only the speed of learning but also the ability to switch the ordering. The value of the offset should not be too large or too small (in our case, it is set to 0.07). In the former case, overshooting may happen and the learning curve may become damped. In the latter case, the learning will be too slow.

The training procedure is thus as follows:

1. Initialize the weights of the net by random values or load them from some previously trained results.
2. Get a sample terrain and a user specified cost classification as described in Section 3.
3. Calculate the feature values of every classified location.
4. For each pair of locations from different cost classes, if the current neural net evaluation for the pair indicates a different ordering from the user specified cost classes, then construct a target value as described above and modify the net weights using the back propagation algorithm [Rumelhart,1988].
5. Train the net on extreme cases periodically to help stretch out the distribution of output values.
6. Repeat steps 4 and 5 until no out-of-order location pairs can be found.

5 EXPERIMENTS

Three synthetic terrains were created and classified based on user's determination of stability and traction. One terrain consists of ramps with different slopes (Figure 5, the elevation is scaled by a factor of 4 for better viewing). Another consists of planes with different roughness. The third terrain contains hills of various shapes. In addition, for each terrain,

finer classifications were provided to emphasize subtle differences. For example, in the terrain of various slopes (Figure 5), one classification distinguishes high slope planes from low slope ones; another distinguishes locations along the edges from those inside a plane; a third distinguishes higher locations along an edge from lower ones.

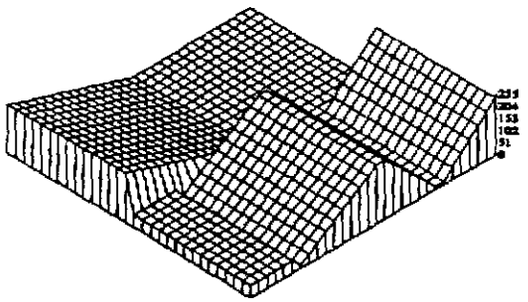

Figure 5: A synthetic training terrain

The net is initialized with random values, then the synthetic terrains are provided for training one after the other. The training converges very fast. For example, for the training on the terrain with different slopes, initially 648 pairs are out of order. After the net is provided with 30 training cases (where a training case is either from an out-of-order pair or an extreme value), no additional out-of-order pairs remain. For the two finer classifications, which are conceptually harder, it totally takes about 80 training cases to converge. This high speed of convergence probably comes from the small size of the net, the relevance of the chosen terrain features, and the correct sequencing of the coarse-to-fine approach. When the training is done by presenting the finer classifications first, the convergence takes about 180 training cases.

The footfall cost map of the synthetic slopes after the training is shown in Figure 6. As can be seen from the graph, the edge locations are very bad footfall locations (with high values) and the difference among ramps of various slopes is not very significant (note that the elevation is exaggerated). This is intuitively what one could expect.

An actual terrain map (Figure 2) with a user cost classification was obtained as a reference terrain to see how well the function learned from synthetic data can work on real, noisy terrain data. The results are quite encouraging: the trained net produces only 16 out-of-order pairs, out of around 50,000 possible pairs. The problem of noise seems to be handled well. This ability to tolerate noise derives partially from the inherent interpolating ability of a neural net and partially from the abstracting and filtering ability of our feature-based approach.

The footfall cost map of the reference terrain is shown

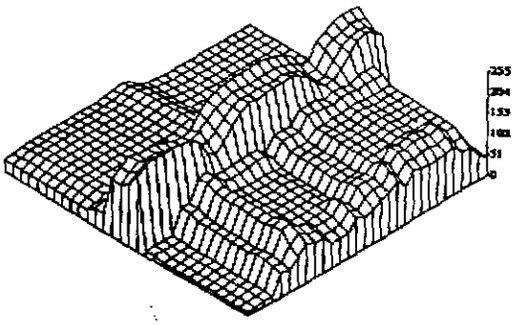

Figure 6: Footfall cost map of the training terrain

in Figure 7. As can be seen from the graph, the result of training seems promising. Those bad footfall locations (along the step-like edge and around small protrusions) stand out very well, while mainly flat areas have very low values. Figure 8 shows the learning curve in terms of the number of out-of-order position pairs in the reference terrain. The asymptotic convergence of the curve also demonstrates the feasibility of the learning approach.

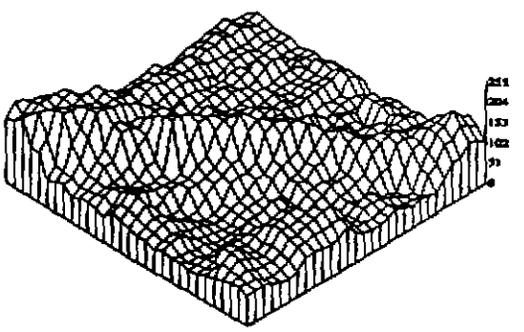

Figure 7: Footfall cost map of the reference terrain

The computation time for evaluating a 30x30 elevation map (i.e. 900 points) is 10 seconds on a Sun Sparc workstation. However, the number of feasible footfall locations received from a gait planner is typically around 100. A computation time of 1 to 2 seconds is thus expected for the real run. The off-line training for the curve shown in Figure 8 takes about 20 minutes.

6 CONCLUSION AND FUTURE WORK

We have developed a method to learn footfall evaluation that relates terrain features to the desired criteria of stability and traction. The approach sketched here is applicable to cases where a real-valued function is to be learned from relative measures.

In future work, we plan to test the generality of the ap-

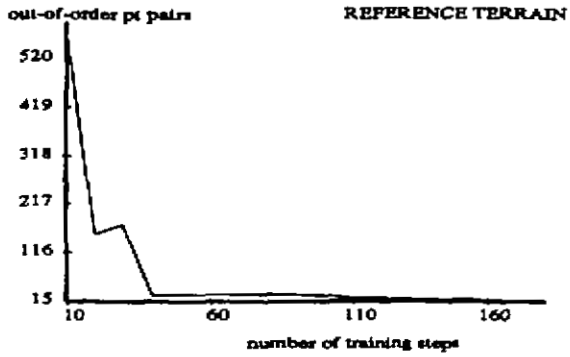

Figure 8: Learning curve for the reference terrain

proach by doing additional experiments on more real terrain data. We further plan some extensions to the method presented here. One is to fine-tune the evaluation function as the Ambler is actually walking, by using force/torque sensor and tilt angle feedback as a measure of stability and traction. Other extensions include incorporating more features from other sensors, such as camera images, and exploring the footfall evaluation function for other criterion, such as minimal power consumption.

Acknowledgments

This research was sponsored by NASA under Contract NAGW-1175.

Thanks to Long-Ji Lin for his provision of back propagation training program and to Tom Mitchell for many helpful discussions.

References

[Bares,1989] J. Bares, M. Hebert, T. Kanade, E.Krotkov, T. Mitchell, R. Simmons, and W. Whittaker. Ambler: An Autonomous Rover for Planetary Explorations. *IEEE Computer*, 18-26, June 1989.

[Caillas,1989] C. Caillas, M. Hebert, E. Krotkov, I.S. Kweon, and T. Kanade. Methods for Identifying Footfall Positions for a Legged Robot. In *Proc. IEEE International Workshop on Intelligent Robots and Systems*, 244-250, Tsukuba, Japan, September 1989.

[Hebert,1989] M. Hebert, C. Caillas, E. Krotkov, I.S. Kweon, and T. Kanade. Terrain Mapping for a Roving Planetary Explorer. In *Proc. IEEE Robotics and Automation*, 997-1002, Scottsdale, Arizona, May 1989.

[Rumelhart,1988] D.E. Rumelhart, G.E. Hinton, and R.J. Williams. Learning Internal Representations by Error Propagation. In D.E. Rumelhart and J.L. McClelland, eds., *Parallel Distributed Processing, Volume 1: Foundations*, 318-362, MIT Press, Cambridge, Massachusetts, 1988.

[Simmons,1991] R. Simmons and E. Krotkov. An Integrated Walking System for the Ambler Planetary Rover. In *Proc. IEEE International Conference on Robotics and Automation*, 2086-2091, Sacramento, California, April 1991.

[Wettergreen,1990] D. Wettergreen, H. Thomas, and C. Thorpe. Planning Strategies for the Ambler Walking Robot. In *Proc. IEEE International Conference on System Engineering*, Pittsburgh, Pennsylvania, August 1990.

Computing in Aerospace 8.
Baltimore, MD, October 1991.

A SIX-LEGGED ROVER FOR PLANETARY EXPLORATION

Reid Simmons, Eric Krotkov and John Bares
School of Computer Science and Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

To survive the rigors and isolation of planetary exploration, an autonomous rover must be competent, reliable, and efficient. This paper presents the Ambler, a six-legged robot featuring orthogonal legs and a novel circulating gait, which has been designed for traversal of rugged, unknown environments. An autonomous software system that integrates perception, planning, and real-time control has been developed to walk the Ambler through obstacle strewn terrain. The paper describes the information and control flow of the walking system, and how the design of the mechanism and software combine to achieve competent walking, reliable behavior in the face of unexpected failures, and efficient utilization of time and power.

1 Introduction

To explore the surface of another planet, a mobile robot must be highly competent, reliable, and efficient. A robot competent enough for planetary exploration must be highly mobile, perceive rugged terrain accurately, and successfully plan and execute paths through extreme environments. To ensure reliable behavior, the robot must never try to exceed its capabilities, should monitor its own health and safety, and be capable of reacting to failures in a timely and appropriate manner. Finally, it is highly desirable for the robot to be efficient in terms of its resource utilization, especially power and computation, since it must carry these on board.

We have built and are currently testing the Ambler, a six-legged robot designed for competent, reliable, and efficient planetary exploration [1, 2, 21]. The Ambler has a novel configuration consisting of a stacked arrangement of orthogonal legs (Figure 1). It features a unique *circulating gait*, where trailing legs recover past all others to become leading legs, that significantly decreases the average number of steps needed for travel. The Ambler has been designed to stably traverse a 30° slope while crossing meter sized surface features (e.g., ditches, boulders, and steps).

A legged configuration offers several advantages over wheeled locomotion [1]. Walking machines isolate the robot's body from the underlying terrain and can propel the body (along with terrain sensors and scientific equipment) on a trajectory that is relatively independent of terrain details. Due to the ability to choose footfalls that conform to the terrain, the Ambler provides a high degree of stability and safety while traversing very uneven terrain. Walkers are theoretically power efficient in part because the body can be maintained at a fairly constant orientation and elevation, and because power losses due to terrain interactions are minimized by discrete foot placements.

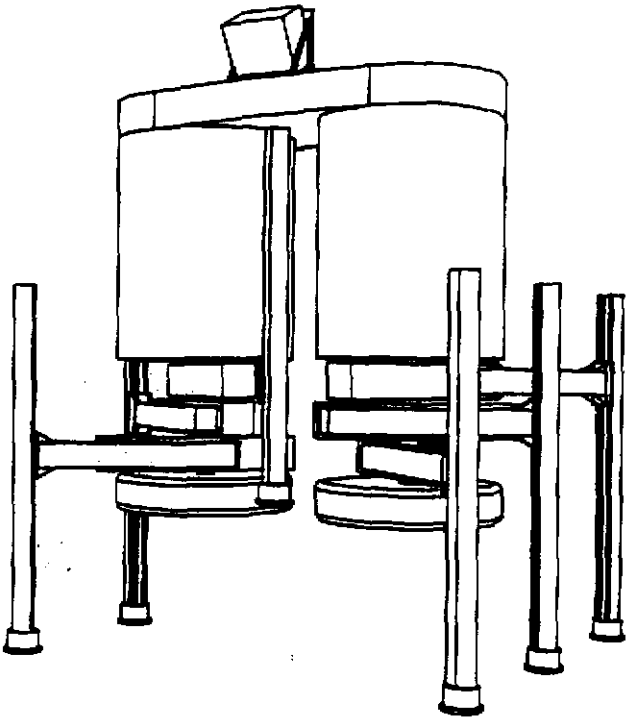

Figure 1: The Ambler Robot

While other walking machines have demonstrated advanced capabilities (e.g., [6, 8, 19, 23]), the majority are either teleoperated or operate under supervisory control. In contrast, we have developed an autonomous software system for navigating the Ambler in rugged, unknown terrain. This is especially important in planetary exploration: due to long signal delay times from Earth, teleoperation would not be very productive or reliable.

The walking system integrates perception, planning, real-time control, and task-level control. The per-

ception subsystem uses data from a scanning laser rangefinder to autonomously calibrate the rangefinder and to build 3D maps of the terrain. The planning subsystem combines kinematic, terrain and pragmatic navigational constraints to find leg and body moves that provide good forward progress and stability. The real-time control coordinates the Ambler's joints to perform accurate leg and body moves, maintains the dead-reckoned position, and monitors the status of the robot. The task-level control facilitates concurrent operation of the subsystems, reliable execution monitoring and error recovery, and management of the Ambler's computational and physical resources.

In numerous trials, the Ambler has autonomously walked through rolling, sandy terrain and across meter tall boulders, ditches, and ramps. The system incorporates many safety features to prevent the Ambler from harming itself, such as ensuring that the rover remains stable and monitoring for unexpected terrain collisions. The safety features are hierarchically layered, with critical low-level checks implemented directly in hardware, mid-level safety checks performed by the real-time controller, and additional checks performed by the planning and task-control subsystems.

The next section presents the mechanical configuration of the Ambler and Section 3 describes the walking system and our experiments to date. Sections 4, 5 and 6 describe the design rationale for the system in terms of the desired characteristics of competence, reliability and efficiency, respectively. Finally, Section 7 concludes with plans for future work.

2 The Ambler

The Ambler is configured with six legs, arranged in two stacks on central shafts (Figure 1). Its height ranges from 4.1 to 6.0 meters, and its width can vary between 4.5 and 7.1 meters. The shafts are connected to an arched body that supports four enclosures housing electronics and computing. This includes two CPU boards for real-time control, nine Creonics motion control cards, and three Sun workstations. Communication with the outside world is via a tethered Ethernet.

The Ambler's legs are orthogonal mechanisms that decouple horizontal and vertical motions (Figure 2). Motion of the vertical links is accomplished using a rack and pinion drive. The Ambler's feet can passively rotate about the vertical axis. The two degrees of freedom of horizontal motion are provided by a rotary link around the central shaft and an offset extensional link. The rotary joint is continuous, with the leg electronics connected to the body via multi-ring sliprings. To reduce the number of conductors that must pass through the sliprings, non-time-critical signals are multiplexed.

The stacked legs and body cavity between the stacks enable the Ambler to exhibit a novel *circulating gait*,

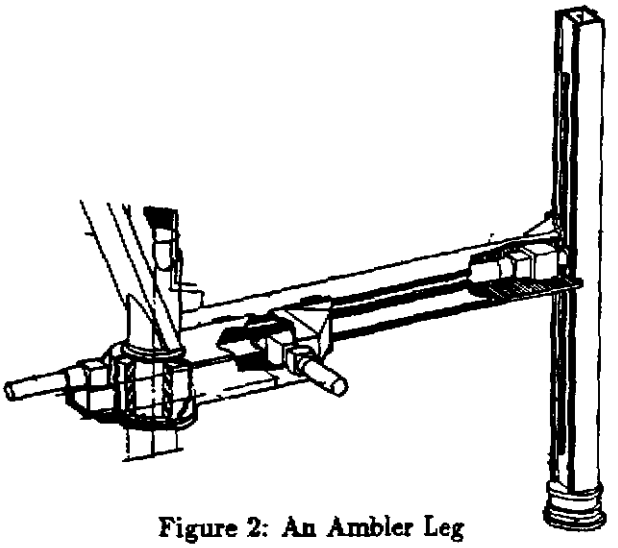

Figure 2: An Ambler Leg

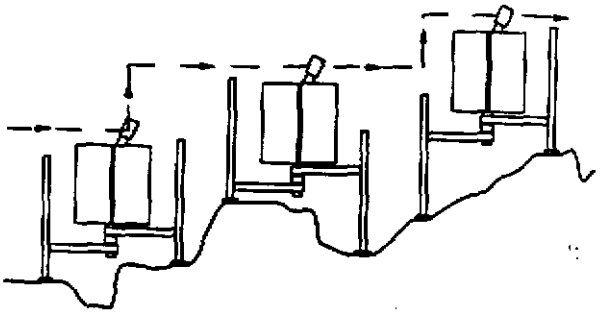

Figure 3: Level Body Motion

where a trailing leg recovers through the body cavity and past the other two legs on the stack to become the new leading leg. The circulating gait is depicted in Figure 4, where a sequence of leg and body moves is shown from left to right across the page. As the Ambler moves forward, the bold leg (leg 1) completes a full counter-clockwise revolution about the left body shaft. During the same period, all other legs also circulate to their original positions. The use of a circulating gait enables the Ambler to use approximately one-half to one-third the footsteps of a follow-the-leader type walker with similar legs.

The Ambler is designed to walk with a level body. This enables body moves to be decoupled into motion in the horizontal plane (translation and rotation) and vertical raising and lowering of the body (Figure 3). Legs adjust individually to terrain roughness and maintain the body in a level orientation over the terrain. Equal displacements on all legs are used to lift the body in climbing slopes, steps, etc. Level body motion simplifies the control problem by reducing the number of joints that must be simultaneously coordinated and provides the Ambler's terrain sensors with a predictable field of view.

The Ambler has a number of sensors to monitor its progress and safety. Each motor has a fail-safe load

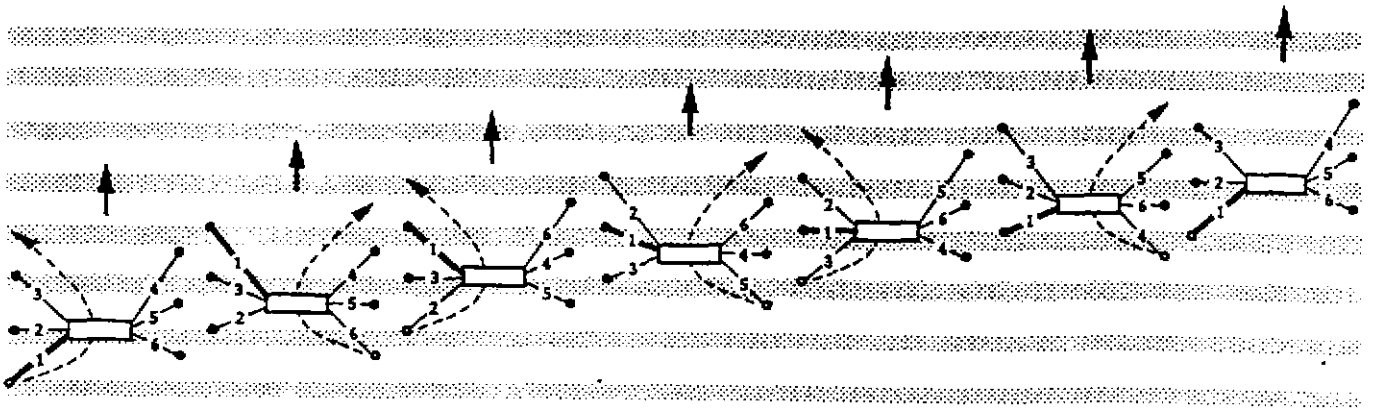

Figure 4: A Circulating Gait

holding brake that activates automatically when power to the rover is lost. All joints have both absolute and incremental (optical) encoders for servo control and dead-reckoning. The prismatic joints each have limit switches to detect excessive travel of the links, and six axis force/torque sensors mounted on each foot are used to detect terrain contact. Two inclinometers on the body indicate tilt and roll from the horizontal plane. Perception of the terrain is provided by a forward-pointing scanning laser rangefinder and a black-and-white CCD camera.

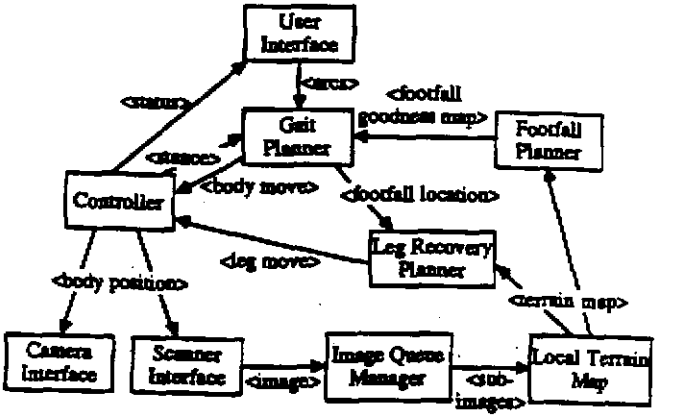

Figure 5: Information Flow Between Modules

3 Autonomous Walking

The Ambler walking system consists of a number of distributed modules (processes), each with a specific functionality. Roughly, the modules divide into planning, perception, and real-time control functions. Figure 5 illustrates the information flow between modules. In this section, we briefly describe the functionality of the modules, detail the information and control flow that combine to produce autonomous walking, and present some of our experimental results to date.

3.1 System Modules

The Controller module is the interface to the Ambler mechanism. It controls the execution of leg and body moves, and monitors the electronics and sensors to detect faults. The Controller runs on two CPU boards in a real-time, multi-tasking environment. One CPU is dedicated to trajectory generation and communication with other modules. The other executes horizontal and vertical trajectories, sending out updates to the motion control boards every 10msecs.

The Controller also maintains the dead-reckoned stance (body and feet locations) of the Ambler. Given an initial stance, the positions of all feet in the global frame are computed. Thereafter, following each body move the Ambler's location is estimated by finding the rotation matrix and translation vector that, with the least squared error, transform the positions of the feet in the vehicle frame to their stored positions in the global frame. This method of dead-reckoning is fairly accurate (typically within 1-2cm), and can often detect when individual feet have slipped.

The Scanner Interface and Camera Interface modules acquire images from a 2D scanning laser rangefinder and a black-and-white CCD camera, respectively. The modules tag the images with the current dead-reckoned position. The Image Queue Manager receives range and camera images, and stores them for use by other modules. This module also pre-processes images, performing such tasks as edge detection and flagging bad data pixels.

The Local Terrain Map module utilizes the range images to construct elevation maps of the terrain (Figure 6). It uses the *Locus Method*, which is an efficient algorithm for transforming and interpolating range data from the scanner frame into Cartesian coordinates [5]. To construct maps useful for planning, occluding legs are masked out of the images, the maps are filtered to reduce noise, and occluded and unknown regions are flagged [12].

Three planning modules are used to calculate sequences of leg and body moves that enable the Ambler to follow arcs of different radii. The Gait Planner uti-

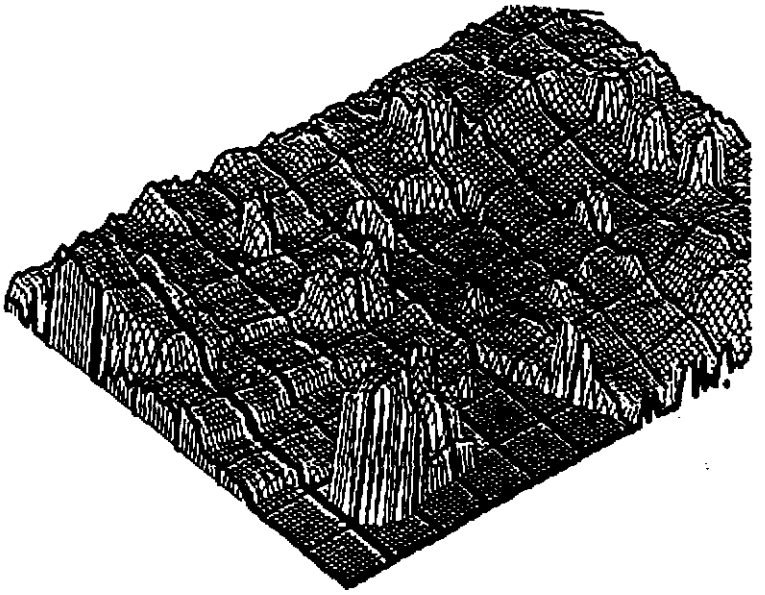

Figure 6: Local Terrain Map

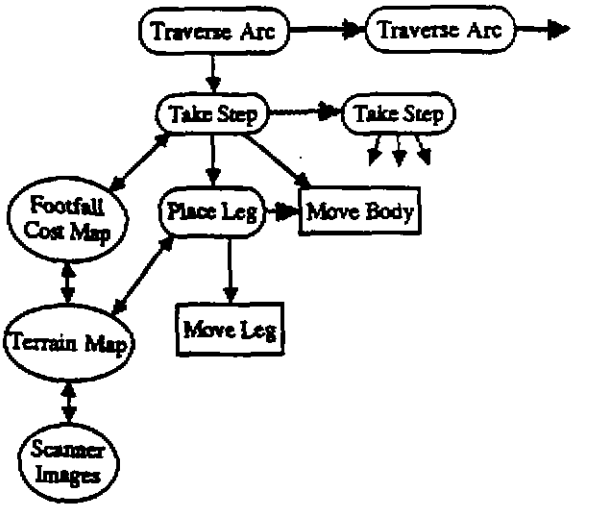

Figure 7: Task Tree for Ambler Walking

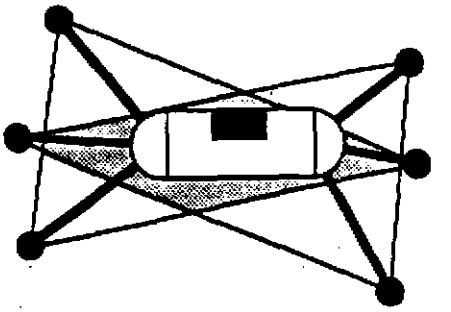

Figure 8: Conservative Support Polygon

lizes constraints on the Ambler's kinematic motion to determine maximal body moves along the arc. Kinematic and pragmatic constraints (e.g., "do not place a leg in front of the body") are combined to determine feasible moves for the legs. The Footfall Planner analyzes elevation maps to determine desirable locations to place the feet. The Leg Recovery Planner produces efficient trajectories to move the legs to their desired locations.

We are currently developing a graphical User Interface module for interacting with the system. The module will provide users with an overhead view of the area around the Ambler, enable users to interactively specify a series of arcs to be followed, and provide feedback regarding the progress and status of the Ambler and the software system.

The modules are integrated into a complete system using the Task Control Architecture (TCA) [14, 20, 22]. While conceptually the modules communicate directly with one another, in reality they send messages to a centralized *task control module*, which logs the messages and routes them to the appropriate modules to be handled. TCA also maintains hierarchical *task trees* (Figure 7) that are used to coordinate the planning, task execution, and monitoring and error recovery needed by the walking system.

3.2. Control Flow

This section describes a typical autonomous walking cycle. The system first calibrates the laser scanner. We have developed an automatic calibration procedure that moves a leg to various positions within the scanner's field of view, uses image processing techniques to locate targets on the leg, and calculates the transformation that minimizes the error between the target locations in the images and their projected positions based on the kinematics of the leg [11]. The procedure is quite reliable and accurate, on average yielding a r.m.s error of 5cm.

To begin walking, a series of arcs is input. TCA forwards each arc in turn to the Gait Planner, which issues a "takeStep" message to begin the planning process (Figure 7 illustrates the control flow: double-headed arrows denote queries to other modules, single-headed arrows denote task decomposition, and gray arrows indicate sequentiality). The Gait Planner first finds a body move along the arc that does not violate the limits on leg motion and keeps the Ambler's center of gravity within the *conservative support polygon (CSP)* [16]. The CSP, which is the intersection of all five-legged support polygons, is the area within which the Ambler will remain stable even if any single leg fails (Figure 8). The chosen body move is then sent via TCA to the Controller module.

The Controller executes the planned body move by computing, for each leg, the trajectory of horizontal joint motions that will simultaneously translate and rotate the body to achieve the commanded position. The Controller precisely synchronizes the twelve horizontal joints to achieve smooth acceleration and deceleration of the body. When the body move completes, TCA notifies the Scanner Interface and Camera In-

terface modules to acquire new images. If any errors occur, the Controller halts all motion and engages the brakes. A separate Error Recovery module is then invoked to handle the situation (see Section 5).

While the body move is occurring, the Gait Planner plans to move the leg that maximizes subsequent body motion. It finds a region of geometrically feasible footfalls based on the limits of the leg's motion, the need to avoid colliding with the current leading leg, the desire to avoid placing the leg in the path of the body, and the need to maintain conservative support in subsequent moves [26]. The Footfall Planner provides estimates, within the geometrically feasible region, of the goodness of the footfalls with respect to the underlying terrain, and the Gait Planner chooses the best one.

While the stability and traction of a footfall cannot be determined directly from the elevation maps, they can be estimated from features such as the slope, roughness, and curvature of the terrain. The Footfall Planner combines these features using an evaluation function whose coefficients are learned by a neural net [7]. The net is trained off-line in advance using human-supplied preferences on footfall locations (e.g., "good", "very bad"). For each pair of locations where the evaluation function indicates a contrary preference, back-propagation is used to update the weights in the net based on the difference between the two values.

The Gait Planner sends off the chosen footfall to the Leg Recovery Planner and, if the end of the arc has not been reached, sends a "takeStep" message to itself to plan out the next move (Figure 7). In this way, the system simultaneously executes one step while planning the next [20].

The Leg Recovery Planner (LRP) requests an elevation map that encompasses the possible paths between the current leg location and the chosen footfall. The Local Terrain Map module, in turn, requests from the Image Queue Manager a series of the most recent scanner subimages that contain relevant range data. The Locus Method is then used to compute the desired elevation map.

The LRP first determines a trajectory for the horizontal joints that avoids collisions with the terrain and the mechanism itself. It then optimizes the vertical trajectory to minimize travel time. This is accomplished by projecting an envelope of vertical locations created by assuming that the leg raises/lowers at full speed while traveling horizontally. Purely vertical moves are added only when the envelope would intersect the terrain. The LRP then issues a leg move command, which TCA queues until the previous body move is completed.

The Controller generates leg move trajectories that are linear in joint space. As with the body move, the three leg joints are coordinated so that all motions start and end simultaneously. The Controller checks the leg move for feasibility, refusing to perform the move if the leg would exceed safe limits, or collide with other legs or the Ambler's body. During execution, the force sensor is set to trigger an interrupt if a force threshold is exceeded, signifying terrain contact. This is treated either as the successful completion of a leg move or as an unexpected collision, depending on where in the trajectory it occurred. Typically, by the time the body and leg moves have been executed, the planner is ready with a new pair of commands.

3.3 Experiments

We are in the midst of an extensive experimental program to test the limits of the Ambler and the walking system. The Ambler currently operates within a large indoor area that can be sculpted to provide a variety of terrains (Figure 9). We have also implemented 2D and 3D graphical simulators to facilitate development of planning and task control software [24], and developed a full dynamical simulation program to help us better understand the real-time control problems [17].

The Ambler has autonomously walked along a variety of arcs. The system has little trouble traversing arcs of different radii, although transition between arcs is not always smooth and sometimes feet must be shuffled. The Ambler has walked over boulders up to one meter tall, crossed a 30° wooden ramp, and negotiated sandy, rolling terrain. Observed slippage of the feet is minimal, and the change in tilt varies by only a fraction of a degree per move. To compensate for accumulation of tilt, however, the Controller periodically adjusts leg heights to level the Ambler.

Obstacle avoidance has sometimes been a problem, due to unreliable readings from the scanner caused by mechanical and electrical problems. This is compensated for, to some degree, by maintaining a large safety zone around the foot as it moves through space. In any event, the Controller can stop the leg in less than a second when the force sensors detect unexpected collisions. We anticipate that on-going repairs to the scanner will alleviate much of this problem.

Average walking speed, including all computation, is 35cm/min (each body move is about 50cm). Moving the mechanism is the main limitation to the speed. During operation, the Controller is active about 80% of the time, while the planners and perception subsystems are each active about 50% of the time, and the centralized Task Control module is active only about 3% of the time (the total is greater than 100% because operations occur concurrently). To date, the Ambler has walked autonomously well over a kilometer.

We have also used preplanned motions to test the limits of the Ambler's mobility. This has included stepping down the sheer face of a meter high rock into a meter deep trench, propelling with several feet poised on the edges of boulders, walking with one foot in the

Figure 9: The Ambler Testbed

air (to demonstrate conservative support), and raising the body to its full height of six meters.

Current experiments involve long-term autonomous operation: the goal is to have the Ambler perform many hours of figure-eight circuits without human intervention. One difficulty is that with each step the Ambler's dead-reckoned position drifts from its actual position, causing it to veer off course. To compensate for this, we are developing a Position Estimation module that visually determines position with respect to known landmarks (in this case, window frames). The method identifies landmarks by extracting strong vertical edges from black-and-white images, and searches an interpretation tree to determine the position from which the identified landmarks would be visible [9]. This position estimate is then combined with the dead-reckoned position based on their relative uncertainties. The new estimate of the Ambler's position can then be used for more reliable navigation.

4 Competence

The following three sections discuss how the configuration of the Ambler and the components of the walking system combine to produce competent, reliable, and efficient behavior.

By competence, we mean the ability to traverse extreme terrain conditions. The Ambler is designed to step over obstacles nearly two meters tall, and to clear one meter sized features while climbing a 30° slope. In addition, it is omni-directional: able to walk backwards and sideways, and turn in place.

In very rugged terrain, the choice of good footfalls is limited. The Ambler's design increases the chances of attaining a good footfall because, unlike traditional pantograph and knuckle-jointed legs, the vertical links of its orthogonal legs sweep out no additional volume as the body moves. Thus, legs can be placed in tight spaces, such as along rock faces. The Ambler's wide foot (30cm) increases its ability to catch toeholds and avoid sinking into soft terrain, and its rigid ankle facilitates stepping on and climbing sloped surfaces. The long leg stride enabled by the circulating gait reduces the total number of foot placements needed and avoids the need for feet to place adjacent to one another. This advantage is diminished somewhat, however, by the stacking of the legs, which reduces the area of available footfalls as compared with designs that space orthogonal legs along the sides of the body [1].

To take advantage of the ability to step in tight spaces, motion control must be extremely accurate. The Ambler is built to be very rigid: the measured sag of the body while standing is only 2cm. The joints are similarly accurate: the prismatic joints can be controlled within millimeters of the commanded motion, and the backlash of the rotary joint is less than a hundredth of a radian. In any envent, deflections in leg links and actuators have insignificant effect on the planar linkage geometry of orthogonal legs, enabling accurate body motion.

The ability to step in tight spaces also depends on the accuracy of the terrain maps. The laser scanner we use digitizes to 12 bits over a range of approximately 40m, which provides a range resolution of 1cm. Taking into account sensor noise and the uncertainty inherent in the calibration procedure, the effective resolution of the elevations maps is around 10cm. This means that we can confidently place the 30cm foot within an area about 50cm wide, which is quite adequate for all but the most extreme terrains.

The planning algorithms provide a fairly general capability for walking in cluttered terrain. In particular, the Gait Planner combines kinematic, terrain and pragmatic constraints to plan moves both forwards and backwards along arcs of any radius, including straight-line paths and point turns. In addition, we are investigating other modes of walking, such as "crippled" gaits (walking with fewer than six legs) and "crab" walking (moving from side to side), which might be expedient in certain situations.

Besides competence in walking, the Ambler provides a competent platform for scientific investigation. Its level body motion provides a stable, predictable platform for cameras and scientific equipment. Instruments placed on the bottom of the stacks can be lowered to contact the ground, which is important for tasks such as seismography and coring. Finally, the legs themselves can be used to perform experiments, such as utilizing the force sensors to investigate material properties of the underlying terrain [10].

5 Reliability

Given its distance from Earth, both in time and space, it is clear that a planetary rover must be extremely self-reliant. The rover must detect unexpected conditions, including hardware and software problems, and react to them in a timely and intelligent manner. This may include requesting assistance from humans when situations arise that the rover itself cannot handle.

Our approach is to provide a hierarchy of safety features, extending from hardware, through real-time control, to the planning and task-level control software. The lower level safety features react on a short time scale and reflexively act to stabilize the Ambler (typically by halting its motion) when unexpected events are detected. Once stabilized, higher level procedures are activated that analyze the situation and take corrective action.

The Ambler has a number of hardware features to prevent damage to the mechanism (see Section 2). Prominent is a safety circuit that continually monitors the motor amplifiers, motion control cards, and limit switches, disabling all motion commands and enabling the brakes if a fault is detected. The safety circuit also monitors for a "heartbeat", a periodic signal from the Controller module that indicates a functioning software system.

Before executing moves, the Controller uses kinematic models of the Ambler to verify that the commands generated by the planners are valid. In particular, it ensures that joint limits will not be exceeded, that leg/leg and leg/body collisions will not occur, and that the Ambler's center of gravity will remain within the support polygon of its legs. During execution of leg moves, the force sensors are used to detect both liftoff and terrain contact (either planned or unplanned).

A major concern with traversing unknown terrain is that the rover could tip over if the terrain collapses under its weight. While the Ambler will remain stable up to about 17° of tilt, for reasons of both safety and efficiency it is preferable to keep it nearly level. We are investigating several methods for dealing with this problem. One is to use the leg to preload the soil before committing to a step. Another method involves active leveling, where feedback from the inclinometers is used to servo the joints to maintain a level posture. Research continues in this area, utilizing both the Ambler and dynamical simulator to determine optimal strategies for maintaining both body attitude and altitude [18].

Once the Controller stabilizes the Ambler in response to failures, a high-level Error Recovery module is invoked to analyze the problem and determine appropriate recovery steps. The simplest procedure, used mainly in cases of unexpected terrain collisions, is to cancel steps that have already been planned and replan from the current stance. Other recovery strategies include shuffling legs when the planners cannot find feasible moves, lifting the body to pass over unexpectedly tall obstacles, and informing human operators to handle unrecoverable hardware faults. If all else fails, the Ambler can be teleoperated through the Controller module interface.

The conservative, deliberative approach of the Ambler walking system contrasts sharply with the more reactive approach [3] favored in micro-rovers such as Ghengis [15]. While both Ghengis and Ambler can react to terrain contact, Ambler uses higher-level perceptual and cognitive processes to plan its way out of difficulties. We feel that in planetary missions, where failure cannot be tolerated and response time is not critical, the deliberative approach is potentially more reliable because it can compare alternatives based on global information.

Sensor reliability is extremely important in autonomous systems. Planning decisions and error detection depend on the quality of the sensed data. Since no sensor is perfect, we utilize several methods to improve reliability. One is to use redundant sensors and sensor modalities. For example, if a force sensor fails to detect ground contact (which has actually occurred) the leg will drive into the ground, jacking up the rover. This can be detected both by the inclinometers and by unexpected force redistributions in the other feet. Another method for improving sensor reliability is to utilize a model of the sensor to characterize and/or reduce its uncertainty. Such a model has been developed for the laser scanner [13] and has proven useful in doing calibration, filtering images, and producing high quality elevation maps.

6 Efficiency

Power is at a premium in space missions, so planetary rovers must be extremely efficient. In addition, they should minimize the time spent in performing tasks, to accomplish as much as possible during their relatively limited missions.

The Ambler's orthogonal legs and level body motion produce an efficient walking machine. Because horizontal and vertical motions are decoupled, orthogonal legs eliminate energy losses due to geometric work, a principal cause of inefficiency for many walkers [25]. With level body motion, the vertical links carry most of the rover's weight, making body propulsions more efficient. While vertical body lifts require significant power, they are only performed when walking up slopes or over very large obstacles.

Currently, using only off-the-shelf components, the steady-state power consumption of the motors, amplifiers, and associated electronics for the 2500 kg Ambler is about 1800 watts. Lifting and moving a single leg takes an additional 150 watts, propelling the body hor-

izontally takes 450 watts, and lifting the body uses an additional 1800 watts above steady state.

One consequence of reducing power consumption is slow speed. While speed could be increased, without affecting the power budget, by using more efficient motors, amplifiers, and lightweight materials, the mechanism itself only partly determines overall walking speed. Perception and planning are also significant. We reduce this time both in algorithm design and with concurrency.

While laser scanners use more power than passive sensors such as cameras, there is a large savings in the computation needed to produce terrain maps. This is primarily because laser scanners determine distance directly, greatly simplifying the transformation between image data and elevation maps. In addition, the Locus Method used to do the transformation is quite efficient [5]. To reduce the number and sizes of images that must be examined, knowledge of the scanner's field of view is utilized to examine only those subimages that may contain data relevant for producing a map of a desired region.

The Gait Planner reduces the need for look-ahead search by combining constraints on the Ambler's kinematic limits and with heuristic constraints on future feasible moves. These constraints also reduce the area that the Footfall Planner must consider, which in turn reduces the demand on perception. The Leg Recovery Planner takes advantage of the Ambler's orthogonal leg design to find near optimal trajectories efficiently. In particular, it decomposes the problem into horizontal and vertical subproblems, producing ramping trajectories that carry legs over obstacles with minimal vertical motion.

The Task Control Architecture is used to achieve concurrent perception, planning and control. TCA notifies the perception subsystem to asynchronously acquire (and preprocess) images following each body move. For planning and control, the TCA's task trees and temporal constraint mechanisms are used to coordinate the concurrent execution of one leg/body move cycle while planning the next [20]. All in all, the use of concurrency enables the Ambler to achieve nearly continuous walking motion.

7 Conclusions

To survive the rigors and isolation of planetary exploration, an autonomous rover must be competent, reliable, and efficient. To investigate such issues, we have designed and built the Ambler, a unique six-legged robot featuring orthogonal legs and a circulating gait. A comprehensive software system has been developed that combines perception, planning, real-time and task-level control to walk the Ambler through rugged terrain. To date, the Ambler has autonomously walked over a kilometer in rolling terrain, negotiating ramps, boulders, and trenches.

This paper has focused on the design decisions that produce competent, reliable, and efficient behavior. The Ambler is highly mobile, capable of walking in extreme terrain, and the associated perception, planning and control algorithms enable the system to take advantage of this high degree of mobility. A hierarchy of hardware and software monitoring and error recovery strategies produce reliable behavior over a wide range of conditions. The Ambler's orthogonal legs and level body motion combine to eliminate geometric work, making it a power efficient walker. Finally, careful algorithm design and the use of concurrency reduce the time spent in computation, which yields a more time efficient rover.

The Ambler is part of a comprehensive project to develop a complete robotic explorer. To this end, we intend to add on-board power and wireless telemetry in order to test the Ambler outdoors in more difficult terrain and varied conditions. Experiments in long-term, autonomous operation will continue, both indoors and out. We also plan to extend previous work in sample acquisition [4], which will include adding a manipulator arm to the Ambler. We anticipate that these efforts will lead to new insights into the technologies needed for autonomous planetary exploration by robots.

Acknowledgements

Many members of the Planetary Rover project have participated in developing the Ambler and the walking system. In particular, we acknowledge the efforts of Brian Albrecht, Lonnie Chrisman, Chris Fedor, Regis Hoffman, Goang-Tay Hsu, Inso Kweon, Terry Lim, Pete Nagy, Henning Pangels, Gerry Roston, Hans Thomas, Jay West, and David Wettergreen. Red Whittaker, Takeo Kanade, Tom Mitchell, Chuck Thorpe, and Martial Hebert have provided valuable insight and guidance. This research is sponsored by NASA under grant NAGW-1175.

References

[1] J. Bares. *Configuration of Autonomous Walkers for Extreme Terrain*. PhD thesis, Civil Engineering, Carnegie Mellon University, June 1991.

[2] J. Bares, M. Hebert, T. Kanade, E. Krotkov, T. Mitchell, R. Simmons, and W. Whittaker. Ambler: An autonomous rover for planetary exploration. *IEEE Computer*, 22(6):18–29, 1989.

[3] R. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robots and Automation*, RA-2(1), 1986.

[4] T. Choi, H. Delingette, M. DeLuise, Y. Hsin, M. Hebert, and K. Ikeuchi. A perception and manipulation system for collecting rock samples. In *Proc. Workshop on Space Operations Applications and Research*, Albuquerque, NM, July 1990.

[5] M. Hebert, T. Kanade, E. Krotkov, and I. Kweon. Terrain mapping for a roving planetary explorer. In *Proc. IEEE International Conference on Robotics and Automation*, pages 997–1002, Scottsdale, Arizona, May 1989.

[6] S. Hirose. A study of design and control of a quadruped walking vehicle. *International Journal of Robotics Research*, Summer 1984.

[7] G. T. Hsu and R. Simmons. Learning footfall evaluation for a walking robot. In *Proceedings of Machine Learning Workshop*, Chicago, IL, June 1991.

[8] Y. Ishino, T. Naruse, T. Sawano, and N.Honma. Walking robot for underwater construction. In *Proc. International Conference of Advanced Robotics*, pages 107–114, 1983.

[9] E. Krotkov. Mobile Robot Localization Using a Single Image. In *Proc. IEEE International Conference on Robotics and Automation*, pages 978–983, Scottsdale, Arizona, May 1989.

[10] E. Krotkov. Active perception for legged locomotion: Every step is an experiment. In *Proc. IEEE Symposium on Intelligent Control*, pages 227–232, Philadelphia, PA, September 1990.

[11] E. Krotkov. Laser Rangefinder Calibration for a Walking Robot. In *Proc. IEEE International Conference on Robotics and Automation*, Sacramento, CA, April 1991.

[12] E. Krotkov, C. Caillas, M. Hebert, I. Kweon, and T. Kanade. First results in terrain mapping for a roving planetary explorer. In *Proc. NASA Conf. on Space Telerobotics*, Jet Propulsion Laboratory, Pasadena, CA, January 1989.

[13] I. Kweon, R. Hoffman, and E. Krotkov. Experimental Characterization of the Perceptron Laser Rangefinder. Technical Report CMU-RI-TR-91-1, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, January 1991.

[14] L. J. Lin, R. Simmons, and C. Fedor. Experience with a task control architecture for mobile robots. Technical Report CMU-RI-89-29, Robotics Institute, Carnegie Mellon University, 1989.

[15] P. Maes and R. Brooks. Learning to coordinate behaviors. In *Proc. AAAI-90*, pages 786–802, Boston, MA, 1990.

[16] S. Mahalingam and W. Whittaker. Terrain adaptive gaits for walkers with completely overlapping leg workspaces. In *Proc. Robots 13*, pages 1–14, Gaithersburg, MD, May 1989.

[17] D. Manko and W. Whittaker. Inverse dynamic models used for force control of compliant closed-chain mechanisms. In *Proc. ASME Design and Automation Conf.*, pages 61–66, Montreal, Canada, August 1989.

[18] P. Nagy and W. Whittaker. Motion control for a novel legged robot. In *Proc. IEEE Symposium on Intelligent Control*, pages 2–7, Albany, NY, September 1989.

[19] M. Raibert. *Legged Robots That Balance*. MIT Press, 1986.

[20] R. Simmons. Concurrent planning and execution for a walking robot. In *Proc. IEEE International Conference on Robotics and Automation*, pages 300–305, Sacramento, CA, April 1991.

[21] R. Simmons and E. Krotkov. An integrated walking system for the Ambler planetary rover. In *Proc. IEEE International Conference on Robotics and Automation*, pages 2086–2091, Sacramento, CA, April 1991.

[22] R. Simmons, L. J. Lin, and C. Fedor. Autonomous task control for mobile robots. In *Proc. IEEE Symposium on Intelligent Control*, Philadelphia, PA, September 1990.

[23] S. Song and K. Waldron. *Machines that Walk: The Adaptive Suspension Vehicle*. MIT Press, 1989.

[24] H. Thomas, D. Wettergreen, and C. Thorpe. Simulation of the Ambler environment. In *Proc. Modeling and Simulation Conference*, Pittsburgh, PA, May 1990.

[25] K. Waldron and G. Kinzel. The relationship between actuator geometry and mechanical efficiency in robots. In *4th CISM-IFToMM Symposium on Theory and Practice of Robots and Manipulators*, pages 305–316, Warsaw, Poland, 1983.

[26] D. Wettergreen, H. Thomas, and C. Thorpe. Planning strategies for the Ambler walking robot. In *Proc. IEEE International Conference on Systems Engineering*, August 1990.