Autonomous Rover Technologies

Technical and Administrative Contact:

Name
emailname@cmu.edu
412-2 8-xxx (Fax: 412-2 8-xxx)

The Robotics Institute
arnegie Mellon Universit
5000 Forbes Avenue
Pittsburgh PA 15213

For the period Oct 1, 1999 through September, 2000

Amount requested: $ 00,000

____
Red Whittaker, Senior Research Scientist, Date
Computer Science Dept. & The Robotics Institute

____
Dimi Apostolopoulos, Systems Scientist Date
The Robotics Institute

____
Martial Hebert, Senior Research Scientist Date
School of Computer Science

_____________________________ _____
Reid Simmons, Research Scientist, Date
The Robotics Institute

__________________________ _____
Sanjiv Singh, Systems Scientist Date
The Robotics Institute

____________________ _____
Takeo Kanade, Director Date
The Robotics Institute

_______________ _____
Jim Morris, Dean Date
School of Computer Science

____________________________ _____
Susan Burkett, Associate Provost Date
Carnegie Mellon University

Autonomous Rover Technologies Task

e propose a research effort to advance the robotic capabilities of planetary rovers through the creation and validation of mechatronic, perception, and cognition technologies. This research will pursue breakthroughs and insights into aspects of robotic navigation, planning, and learning for autonomous exploration, precision positioning from tracking natural features, robust radar perception for mapping and safeguarding, and synthesis of reconfigurable robotic mobility for challenging planetary terrains.

1. Objectives

The primary goal of this task is to generate robotic technologies to enable autonomous regional exploration by planetary rovers. The specific research areas for FY00, are:

- Navigation, planning, and learning for autonomous Exploration
- Precise positioning from visual tracking of natural features
- Robust radar perception for mapping and safeguarding
- Synthesis of reconfigurable robotic mobility for all-terrain exploration

2. Justification

Ambitious planetary exploration agenda can be undertaken with capable and highly autonomous robots. Upcoming missions to our solar system and in particular Mars exploration call for intra-site traverses in unknown environments, and execution of complex science missions including sampling and geological classification. This program develops the fundamental technologies to enable such challenging missions.

Robot technologies developed by this effort could benefit strategic planetary rover concepts, such as the “rover geologist” and “autonomous planetary robotic systems.” The program will develop insights that are crucial to the realization of NASA’s long term objectives and will produce technologies to make robots more robust and cost effective, both in space and on Earth. This program, as past and current NASA programs at Carnegie Mellon, encourages an exchange of technology between research centers and universities. Publications of articles, conference papers, and technical reports spread our ideas and technologies to NASA and beyond. In addition, we are now engaged in partnerships with projects at NASA centers that will field technology developed in the course of the proposed work in demonstrations leading to near term space missions. Specifically, we are planning to port our technology to demonstration programs at NASA Ames and JPL.

3. Technical Approach

The main focus of the Autonomous Rover Technologies is to innovative and advance the state-of-the-art in planetary rover autonomous mobility, perception, position estimation, navigation, and integrated exploratory science from a robot, to the point where complex planetary missions can be achieved. Our approach is to identify customer needs from NASA enterprise requirements and to introduce novel concepts with merit to fundamental problems in robotic surface systems. This task will produce the techniques, models, prototypes, and validation that will bring the proposed technologies to higher readiness levels.

3.1 Navigation, Planning, and Learning for Autonomous Exploration

Exploration investigates an area to discover and record new information. It is a dynamic process where an explorer continually makes decisions based on its gathered information for the purpose of gathering new information. In the near future, robots will be called upon to perform difficult exploration tasks in remote locations such as searching for water ice on the Moon, life on Mars and Europa, and meteorites in Antarctica. These complex tasks require exploring robots to not only discover and map new terrain but also to gain and utilize information from many different sources. For example, a robot prospecting for water on the Moon needs to discover and use information about which regions are permanently dark, rock types to determine drilling conditions, crater ages and rover as it explores such that it stays best localized. It is presumed that a rover will continually localize itself using features that it can detect visually. Some of these features are known ahead of time while others are discovered as the rover discovers its environment. Such a method will trade-off the requirement to stay localized versus other requirements such as the need to explore and reach a goal location. The second method uses a bound on the estimate of position uncertainty to plan paths that succeed in the face of position drift. In essence, this kind of a method is able to cancel some of the effects of the drift. Currently, these methods are formulated separately. It is desirable that if possible, the two goals be merged into a single formulation.

3.1.3 Exploration

Robot explorers cannot limit themselves to simple, benign terrains but must be capable of making discoveries in rugged terrain with many obstacles. Some areas are near the limits of mobility, making them expensive to drive over, others may be smooth and inexpensive to drive through while still others are impossible to drive over. Obstacles can affect sensor coverage, creating holes in the viewed area. Conversely, high ground can offer better views of large areas. To efficiently explore, robotic explorers must take the effect of terrain, both on exploration cost as well as information gain, into account.

Exploration which considers multiple criteria of information and terrain effects leads to a marked improvement in complex exploration tasks over competing approaches. By investigating exploration planning techniques that maximize information gain from multiple sources while minimizing driving and computation costs, this research will show the benefits of information gain exploration for complex tasks. To investigate the exploration problem, this research will use a model shown below. Depending on the exploration task the robot can have different numbers and types of sensors. Each of these sensors feed information into the Knowledge Database which is the repository of all the information gathered by the explorer.

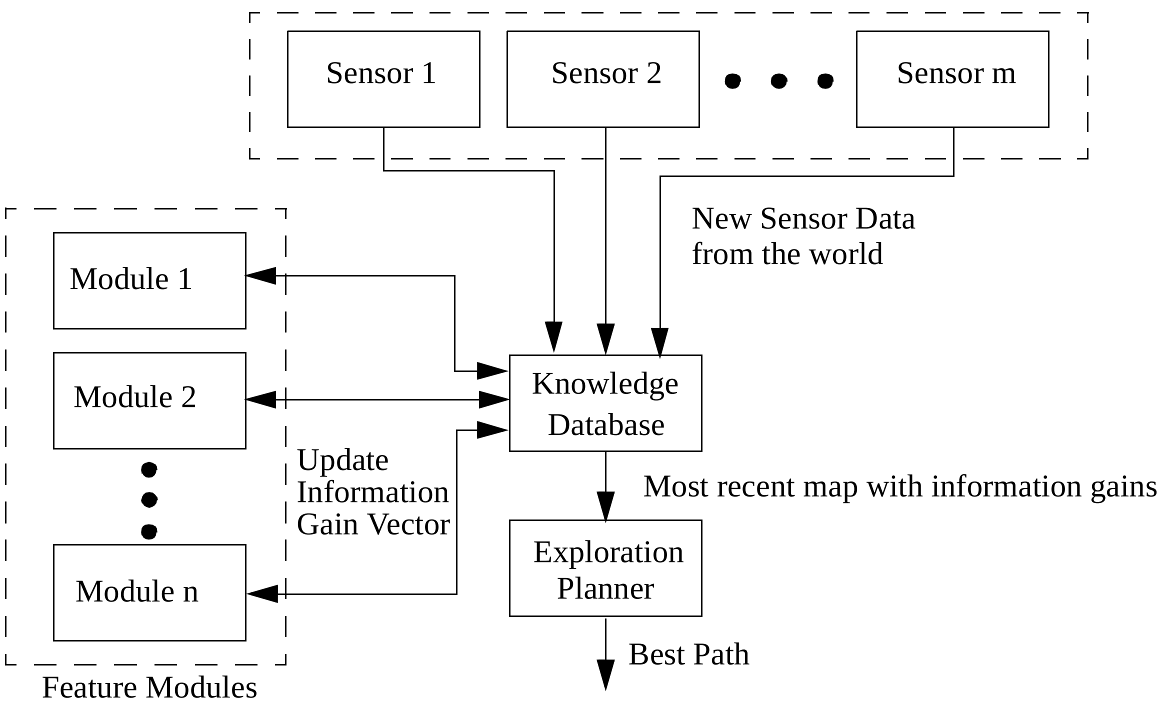

Figure 1: Architecture for autonomous exploration

The number and form of the feature modules also depends on the exploration task at hand. Each of these modules takes data from the knowledge base and uses this to assign values of information gain to specific areas of the region being explored. Each module uses a different criterion to assign an information gain. For example a robot exploring for lunar ice might have four feature modules. One module would look at the latitude and local terrain to determine which areas are likely to be permanently dark. It would then assign high values of information gain to those regions which are permanently dark and low values to regions with long exposure to sunlight. The other three feature modules would look at rock types which determine drilling conditions, crater ages and location of potential sites inside a crater respectively and assign an appropriate information gain value. The information gains determined by the feature modules are then stored in the knowledge database.

The exploration planner box tells the robot where to go. Using the terrain and information gains found in the knowledge database, the planner must create exploration paths which maximize information gain, causing the robot to explore, while minimizing the cost of travelling. Since the regional terrain feature being explored is unknown before exploration, the exploration planner must continually revise the exploration route as new information becomes available.

Essentially the path planning step requires the robot to choose a sequence of areas to visit. At each area it gets a prize, information gain, but there is also a cost to travel from one area to another. If the information gain in each area is independent then this is the Prize Collecting Travelling Salesman (PCTS) problem which is NP-complete so the best that can be hoped for is a sub-optimal solution. However, the exploration planning problem is more difficult than the PCTS because the information gains of areas are not independent and the amount of information to be gained at a location is subject to some level of uncertainty. The development of efficient solutions to the exploration planning problem will be a central part of this research.

The task will investigate the following issues:

- A general framework for the solution of exploration planning problems using the premise of choosing paths of greatest information gain.
- The investigation of efficient solutions to the exploration path planning problem.
- A simulation of an outdoor mobile robot exploring an unknown area using multiple sources of information gain.

3.2 Precision Positioning from Visual Tracking of Natural Features

Many of the autonomous missions currently envisioned feature limited communications opportunities and bandwidth. The ability to automate precision navigation within the confines of a science site and from science site to science site will thus dramatically increase the science return of these missions; more of the available communication bandwidth and rover life can be utilized for science telemetry rather than rover control telemetry.

The proposed research addresses a mode of precision navigation and positioning based on natural landmark, in which the position of natural features in images is tracked over time and integrated to produce an accurate estimate of the rover's position. The proposed work builds upon prior joint work between NASA Ames and CMU in the area of tracking and position estimation. The objective of the proposed work is to extend and integrate the existing techniques into a system leading to a fieldable approach to position estimation.

The present proposal builds upon the results of the FY99 Grant NAG2-1233. Those results include the development of an integrated approach to relative positioning and mapping, its implementation, partial delivery of the software to NASA Ames. The work proposed here involves the continuation of the development of the positioning and mapping techniques, as well as the integration into the NASA Ames system.

In natural landmark navigation, features are located and tracked from the rover's onboard cameras, and the rover's position is determined through triangulation and filtering against dead-reckoned estimates of vehicle motion. Precise relative navigation about a science site is achieved by tracking distinct image features, such as prominent rocks or ledges, and computing vehicle position in a local coordinate frame. Given uncertainty representation on the detection of the features, their position and the position of the rover. Local navigation estimates from matching and tracking will be combined by a pose filter to yield an overall estimate of vehicle position and orientation.

The main components of the landmark navigation system are:

- Landmark detection
- Landmark tracking
- Relative position estimation

e will address those three aspects with an emphasis on the tracking and position estimation aspects, since they are most critical in achieving the desired precision. In addition to detecting and tracking features in color CCD cameras, we will also demonstrate those algorithms using omnidirectional cameras and panoramic cameras as an alternative. Experimentation along those lines will be conducted using the results obtained by CMU in this area under prior NASA grants.

e have developed techniques for detecting landmarks in natural images. Because the landmarks are detected in static images, they may not be suitable for tracking; this is the problem of "feature selection". More precisely, the problem is to use only those features that can be reliably tracked between images and that can be easily discriminated again the other parts of the image. As part of

The first problem faced in natural landmark navigation is to detect landmarks in images. In NAG2-1233, we had developed an approach in which attributes such as discriminability, trackability, and saliency, which characterize the suitability of a feature, are computed over the entire image. The attribute images are then combined to form a feature preference image. Features with the largest confidence are selected as the optimal features to select for tracking.

Preliminary results from NAG2-1233 show that this algorithm does select the preferred features in typical natural environment. For example, Figure 1 shows two of the feature selection maps that are computed from an input image, the trackability and discriminability maps. The two maps are combined into a single confidence map which is used for the final feature selection.

The *proposed work* in this area includes the refinement of the strategy for combining feature confidence maps, the addition of other cues such as color discriminability, and the integration of the feature selection with the mapping system. This part of the work will be conducted at CMU by a graduate student, with input from a student at LAAS, France, who was supported as a visiting scientist during NAG2-1233.

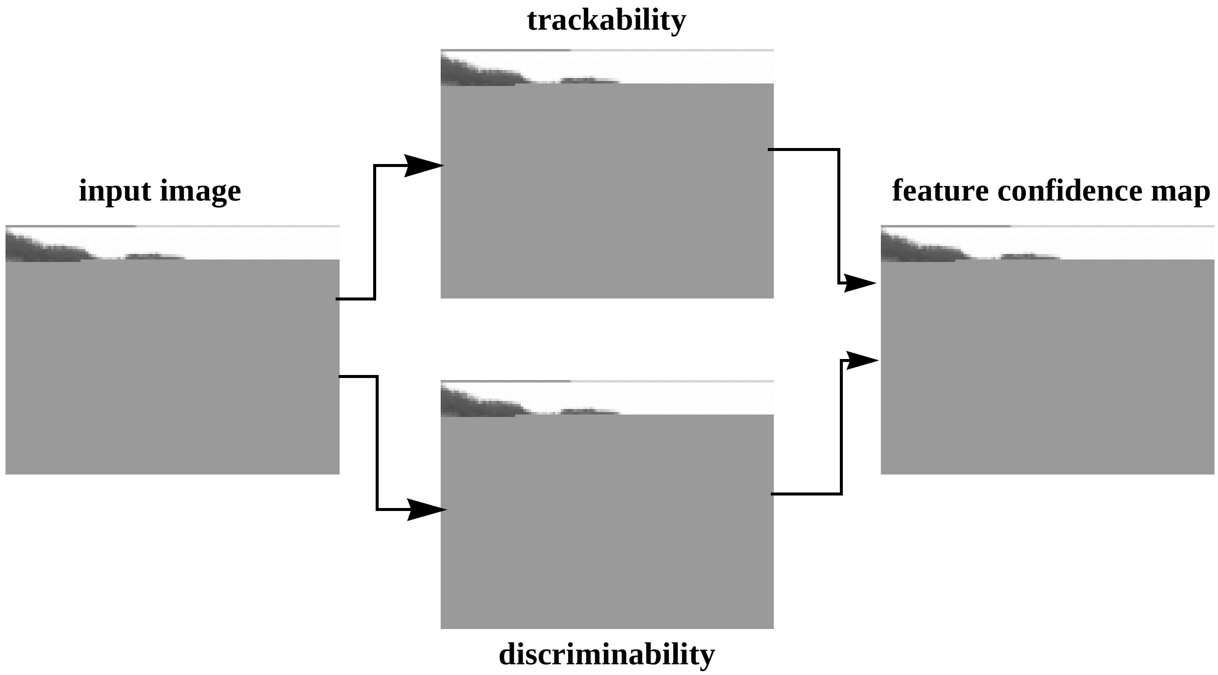

Figure 2: Example of feature confidence map generated by combining multiple criteria for feature selection.

Local vehicle position will be estimated based on two sources of information: triangulation on the currently tracked features, and filtering of previous position estimates. The challenge in this problem is that the errors in robot pose and in landmark positions are strongly coupled. This is known as the SLAM (Simultaneous Localization and Mapping) problem. In our case, the issue is complicated by the fact that we have only bearing measurements to the landmarks, and therefore the actual positions of the landmark cannot be measured directly.

A solution to the SLAM problem was developed during FY99 and will be extended and integrated as part of this FY00 proposal. The basic approach is to use virtual measurements that are independent of the robot position, such as mutual distances and angles between landmarks. Such virtual measurements can be used to build a relative map of the landmarks that is independent of the robot pose. The relative map may be converted to an absolute metric map using a standard iterative process.

One additional issue is that the measurement errors may not be normally distributed, as is typically assumed in traditional estimation techniques such as Kalman filters. The invariant SLAM approach described above is augmented by using a particle based approach for representing and estimating the robot pose. Such an approach can represent arbitrary distributions of uncertainty. The particle approach has been demonstrated in FY99 and will be integrated with the invariant SLAM algorithm as part of the proposed work for FY00.

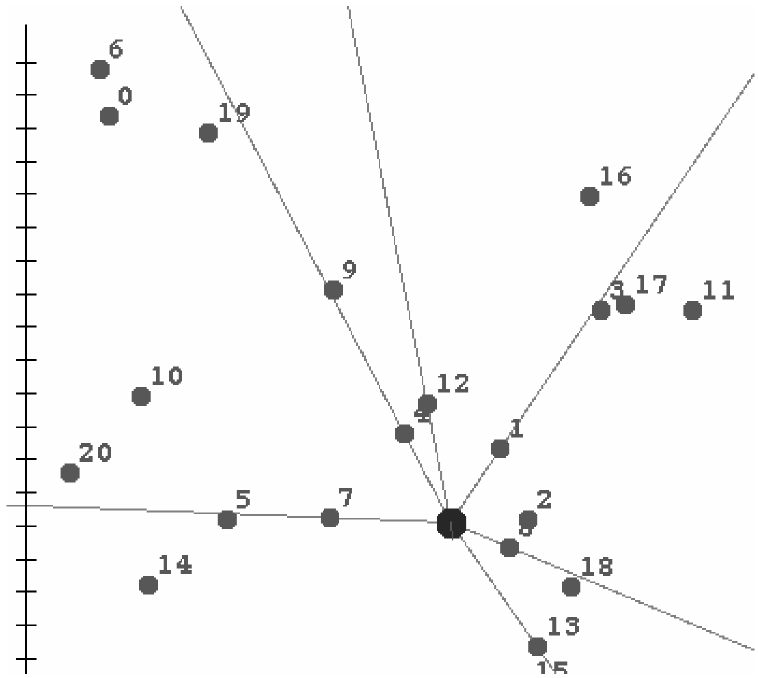

Figure 3: Typical example of bearing measurements from robot to features.

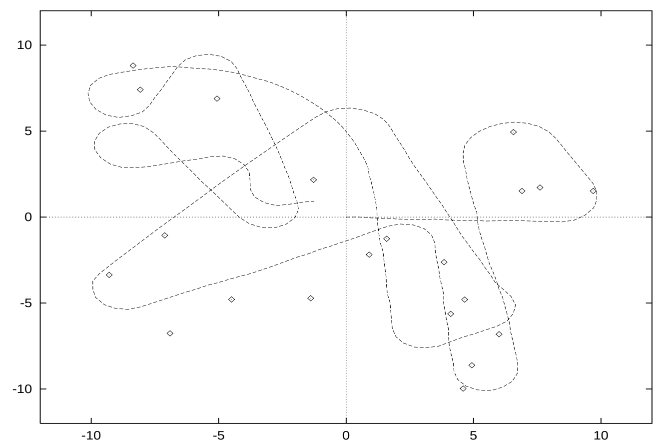

Figure 4: Example of robot trajectory and feature map recovered from bearing measurements.

The current version of the SLAM software has been delivered to ARC. The main goals of the *proposed work* in this area include further development of the invariant localization and mapping algorithms to handle large maps and long trajectories, the integration of particle-based methods and invariant filtering methods, the optimal selection of constraints for estimation, and the integration with visual tracking, and the NSAS Ames Mars system.

A critical aspect of estimating position from feature tracking is the determination of the uncertainty in extracting and tracking features in images. Specifically, the detection and tracking uncertainty may vary greatly depending on the feature. As a result, the accuracy of the position estimate may vary depending on the uncertainty model used for each feature. e will develop an uncertainty model that takes into account the error in feature tracking based on the matching error between the current image and the feature template from the previous image. The approaches based on particle filters developed in NAG2-1233 will be used for representing arbitrary uncertainty distributions. In particular, distributions that are non-Gaussian or multi-modal will be represented accurately.

3.3 Robust Radar Perception for Mapping and Safeguarding

Radar technology offers unique features for planetary navigation and safeguarding purposes for robotics applications. It is not affected by light levels and therefore can be used at any time. Dust clinging to the surface of the antenna does not affect radar performance. This is most relevant in planetary exploration where dust accumulation and darkness are the bane of long term missions. Millimeter wave (MM) radar has a number of unique features that make it an ideal robotic sensor for obstacle detection, mapping, and subsurface object detection, even in dusty environments. A radar unit could mount on the front of a rover and by phased array could electronically steer its fine beam in two dimensions. Immediate radar returns could be used to perform obstacle detection. Successive scans over the same area could provide a detailed terrain map that could be used for planning and safeguarding. Finally, if the radar is pointed normal to the planet's surface, some surface penetration can be achieved.

The generation of short-range perception models from radar is beyond the state of art. High angular resolution can only be obtained with inconveniently large antenna apertures. Mechanical-scanning equipment is heavy and unreliable. Similarly, robot position estimation and navigation based on radar succeed only when relying on artificial landmarks; natural landmarks are hard to distinguish from noise and false targets. Furthermore, the incorporation of motion-free scanning antennae complicates the applicability of radar because the antennas project different beam properties as they scan. Early attempts to build models using short-range radar compare poorly with the perception models built from sonar, laser and stereo sensors. Current short-range radar perception lacks adequate data-interpretation capabilities. Radar interpretation should account for properties such as sidelobes, reflection, antenna-radiation pattern variation and noise distribution. This presents a research opportunity that could benefit robotics with reliable perception in compromised visibility conditions.

This task will investigate radar sensing to create a new interpretation method to address radar-sensor shortcomings. Radar-data heuristics will provide early feature identification. Evidence grids, suggested by wide-beam sensing developments, will provide a representation tool appropriate for radar-sensing research. A sensor profile will capture radar multilobe sensing, and varying-radiation patterns characteristic of a motion-free scanning antenna. Estimation of surface properties will allow for accurate interpretation of radar data. The interpreter will integrate these techniques, and experimental testing will demonstrate enhancement of the perception models.

This research will proceed as follows:

- Interpreter design: Design an interpreter that accepts a raw radar signal, normalizes it and applies heuristic rules. The interpreter will also incorporate multilobe data into an evidence grid and will recognize surfaces. The result will be a template to serve as the context for component research and as the model for integrating components into a culminating capability of radar modeling.
- Radar data heuristics: Develop rules for extracting features such as points and planes from radar signals. The extraction of features considers increased resolution techniques to identify single targets and groups of close targets.

- Detailed radar profile modeling: Devise a methodology for incorporating the multiple lobes of radar-sensor data into an evidence grid. The development will consider variation of the antenna-radiation pattern and reflection extensions in the sensor profile. The task will result in a dynamic multilobe-sensor profile for radar-data incorporation into the evidence grid.
- Surface-extraction: Develop a technique for extracting surface properties from the evidence grid, that are relevant for modeling the specular reflection typical of radar sensing. The outcome will be a component that enhances an environmental representation estimating surface geometry and reflection properties.
- Integration of interpreter components: This task will integrate previous developments radar-data heuristics, radar-sensor profiles, and surface-extraction components into a functional interpretation technique that receives radar data and builds a perception model.

3.4 Synthesis of Reconfigurable Mobility for All-Terrain Exploration

This task will create a novel reconfigurable robotic surface system (R2S2) enabling access and in situ science at high value, high risk planetary locations such as escarpments, fissures, breakout channels. The technology concept is a modular system, as configured from a reusable, fixed parts inventory. The system can autonomously change shape, mechanical functions, and control actions in response to sensed internal state or external observations. Criteria for success are system ability to self-adapt to changing terrain characteristics, achieve high terrainability relative to mass-scale, and at the extreme, discretely reconfigure its functions to escape/recover from entrapment or other major operational anomalies. R2S2 has a further powerful property of being physically separable into two (or more) cooperating agents during use--allowing extended baseline observations, networked instrument activities, or even coordinated robotic work (e.g., one robotic agent anchors; the other rappels). Many different initial system configurations are possible, based on the target task. This “reusable”combinatorial design and fixed inventory greatly benefits mission cost, launch access, and survivability, and has broad application to robotic outposts. Technical products include a virtual prototyping environment, sensor-based autonomous control software, and task-level mobility synthesis and inventory optimization tools.

CMU will collaborate with JPL (task lead), MIT, and University of Nebraska to solve several fundamental problems related to the design of reconfigurable and modular R2S2 mobility. First is the creation of a multifunction system design specifying the kinematics, dynamical, structural, actuation and sensing configuration to meet given task performance constraints and design variable limitations. Automated synthesis and validation for R2S2 is essential because the design space is enormous and poorly understood. A second problem is terrain-adaptive motion control. The primary issue here is controlled reaction and reconfiguration to internally sensed vehicular state and externally estimated terrain parameters, both visual and dynamical (environmental contact forces, soil/tire interactions). As a generalization, this task will investigate cases where the robot, based on sensed/observed constraints, is at pathological limits and must discretely reconfigure its mechanical degrees of freedom and actuation to escape entrapments, re-establish a stable base, even provide a new robotic function (effector/mast/anchor) etc.

Task Technology Readiness Levels (TRL)

Robot Technology	FY00	FY01	FY02
Navigation, Planning, Learning	TRL1	TRL2	TRL3
Precise Localization	TRL2	TRL3	-
Robust Radar Perception	TRL2	TRL3	-
Advanced rover autonomy	TRL1	TRL2	TRL3

. Milestones and Deliverables

4.1 Navigation, Planning, and Learning for Autonomous Exploration

4.2 Precise Positioning from Visual Tracking of Natural Feature

Product Milestones	Product Performance Achievement	Significance & Impact of Product to Missions
FY 00 **LEVEL 1**	Demonstrate software to estimate accurately (i.e., within tens of cm) the robot's position and the position of features in its environment, permitting accurate simultaneous mapping and positioning. Use feature detection in images of unstructured terrain to build a landmark map and to do position estimation.	Product enables a level of accurate positioning and mapping for autonomous explorer not possible with the sensor suite envisioned for those missions. Remaining work: Integrate in navigation and mapping systems. Test with long distance travel and long image sequences.
FY 00 **LEVEL 2**	Produce software for simultaneous mapping and localization from natural landmarks based on the relative filter developed during FY99. This approach is the only that permits mapping and localization in the case of: - Bearing only measurements to landmarks (e.g., feature locations in images) - Noisy pose estimate - Non-Gaussian and multi-modal noise models Produce algorithms for automatic selection of natural features based on the saliency algorithms developed in FY99. Those algorithms select landmarks by combining several criteria - Discriminability of the landmark in images - Trackability of the landmark over sequences of images Test and integrate in the ARC navigation system. Partial delivery of the software took place in FY99.	Permits the use of cameras for localization without requiring additional instrumentation, e.g., range sensors, over the sensor suite envisioned for the autonomous exploration robots. General noise model applicable to typical platforms envisioned for autonomous exploration. Automatic system for mapping and localization as required by autonomous exploration missions.

4.3 Robust Radar Perception for Mapping and Safeguarding

Product Milestones	Product Performance Achievement	Significance & Impact of Product to Missions
FY 00 LEVEL 1 Milestone	Produce accurate terrain-perception models from millimeter-wave radar. Achieve 20 cm resolution and 5 cm accuracy at 20 m.	Radar perception will enable rovers to map and operate through planetary environments cluttered with dust and with low visibility. Missions such as Robotic Outposts for In-Situ Measurement and Communication, as well as Robotic Outposts for Heavy Duty Science Operations will benefit from reliable and accurate radar perception, as an alternative to stereo vision and laser perception.
FY 00 LEVEL 2 Milestone	Demonstrate robust radar perception with a novel data interpretation method that combines sensor and surface models. The interpreter consists of: * Heuristics for early feature identification * Radar evidence grids * A sensor profile that captures radar multi-lobe sensing and varying-radiation patterns *Estimation of surface properties for refinement of radar data The interpreter will integrate these techniques and experimental testing will demonstrate enhancement of perception models.	The proposed interpretation methods will improve the accuracy of radar data, produce rich terrain perception models, and enable reduction of radar size and weight.

4.4 Synthesis of Reconfigurable Mobility for All-Terrain Exploration

Product Milestones	Product Performance Achievement	Significance & Impact of Product to Missions
FY 00 LEVEL 1 Milestone	Produce a virtual prototyping environment to derive R2S2 designs and perform comparative analysis of candidate configurations.	R2S2 will enable access and in-situ science at high value, high risk planetary locations such as escarpments, fissures, channels, etc. This capability is essential for Space Science and robotic precursor to HEDS surface missions.
FY 00 LEVEL 2 Milestone	Design R2S2 control architecture, mechanization concepts and models, and perform task motion synthesis studies.	The design of robotic component technologies for R2S2 will facilitate the production of the virtual prototyping environment of modular rovers for all-terrain exploration.
FY 01 LEVEL 1 Milestone	Develop, test, and validate component technologies of R2S2, including actuation modules and control software.	Introduce advanced R2S2 mobility concepts to the planetary science community. Augment expected return missions through the use of more capable robotic mobility.
FY 02 LEVEL 1 Milestone	Produce an R2S2 prototype and evaluate its performance in real environments. Test R2S2 adaptation and reconfiguration across broadly representative, repeated examples to establish system performance benchmarks.	Establish a new class of planetary rovers for challenging surface missions, such as exploration of Martian canyons and volcanoes.

5. Intended Users

The proposed tasks target future applications in Space Science Enterprise's Solar System Exploration program (S3-Requirements), with important applications to autonomous and human-assisting robotics for Human Exploration and Development of Space (HEDS) Exploration (H3-Requirements) and Human Life Science (H4-Requirements).

6. Participating Organizations

- Funding Agency: NASA Cross Enterprise Technology Development Program.
- Funded Organization: The Robotics Institute, Carnegie Mellon University. Expertise includes perception and navigation for planetary rovers, sensors for modeling/ navigation/science, robots for field applications.
- Collaborating Organizations:
  - NASA Ames Research Center: Precision Localization, Radar Perception, Autonomous Exploration
  - NASA Johnson Research Center: Autonomous Exploration
  - Jet Propulsion Laboratory: Reconfigurable Mobility

7. FY00 Budget & Personnel

Budgets attached.