Autonomous Robots 14, 103–126, 2003

Planetary Rover Developments Supporting Mars Exploration, Sample Return and Future Human-Robotic Colonization

PAUL S. SCHENKER, TERRY L. HUNTSBERGER, PAOLO PIRJANIAN, ERIC T. BAUMGARTNER AND EDDIE TUNSTEL

Jet Propulsion Laboratory, 4800 Oak Grove Drive, Pasadena, CA 91109, USA

paul.s.schenker@jpl.nasa.gov

Abstract. We overview our recent research on planetary mobility. Products of this effort include the Field Integrated Design & Operations rover (FIDO), Sample Return Rover (SRR), reconfigurable rover units that function as an All Terrain Explorer (ATE), and a multi-Robot Work Crew of closely cooperating rovers (RWC). FIDO rover is an advanced technology prototype; its design and field testing support NASA's development of long range, in situ Mars surface science missions. Complementing this, SRR implements autonomous visual recognition, navigation, rendezvous, and manipulation functions enabling small object pick-up, handling, and precision terminal docking to a Mars ascent vehicle for future Mars Sample Return. ATE implements on-board reconfiguration of rover geometry and control for adaptive response to adverse and changing terrain, e.g., traversal of steep, sandy slopes. RWC implements coordinated control of two rovers under closed loop kinematics and force constraints, e.g., transport of large payloads, as would occur in robotic colonies at future Mars outposts. RWC is based in a new extensible architecture for decentralized control of, and collective state estimation by multiple heterogeneous robotic platforms—CAMPOUT; we overview the key architectural features. We have conducted experiments with all these new rover system concepts over variable natural terrain. For each of the above developments, we summarize our approach, some of our key experimental results to date, and our future directions of planned development.

Keywords: mobile robots, cooperating robots, all terrain mobility, robotic colonies, robot architecture, reconfigurable robots

1. Introduction

There is growing international interest in wide-ranging exploration of the Martian surface. A better understanding of Mars' surface geology, morphology, geochemistry, and atmospheric science will yield important insights about comparative planetary origins, potential for past/present life, and capabilities of the Mars surface environment to sustain a permanent human-robotic colonized presence.

Thus, institutions worldwide are pursuing development of Mars mission platforms/payloads, both fixed and mobile, toward these science objectives. There are many options for such Mars surface exploration: stationary landers with affixed instruments/samplers, gravity-impact penetrators, shallow and deep drills, subsurface/tethered "moles", light airplanes, touch-and-go atmospheric balloons, and *semi-autonomous* surface mobility. The word "semi" connotes earth-based planning, command-sequencing and analysis of rover activity sequences and data products—as done by a science-engineering team through periodic data down-link and command up-links. We have done past related work on dexterous landed manipulators (Schenker et al., 1995, 1999a) resulting in a concept for NASA's Mars Polar Lander mission of 1998. More recently we have focused on developing *mobile science platforms*—science rovers, such as the FIDO technology prototype shown in Fig. 1.

Figure 1. FIDO Rover in desert field test on a cobbled lake bed, mast/science arm extended (inset, rover continuously traverses a sand wash, rear view, the mast/arm stowed (Weisbin et al., 1999)).

1.1. Evolution of Mars Robotic Surface Mobility

Near-term mission objectives include long-range mobility and highly instrumented in situ science operations. Such remotely-commanded, over-the-horizon/OTH, semi-autonomous mobile science platforms will enable *remote field geology*. As one specific example, NASA's upcoming Mars'03 (Mars Exploration Rovers) mission seeks to dramatically extend physical and observational scope of the 1997 Mars Pathfinder/Sojourner exploration (Shirley and Matijevic, 1997; Shirley et al., 1997)—from 10's of meters about a nearby lander, with the rover carrying a single rear mounted instrument (AXPS/Alpha X-ray Proton Spectrometer)—to 1000's of meters over open terrain with an on-board integrated science package (Mast instruments include a high resolution multi-spectral panoramic camera, bore-sighted NIR point spectrometer, and integrated thermal emission spectroscopy; rover arm instrumentation includes a color micro-imager, Mössbauer spectrometer, and rock abrasion tool).

Beyond this near-term vision there are major technical challenges and diverse opportunities confronting development of later Mars surface systems: Challenges include extending the spatial range and duration of autonomous science operations (including on-board science analysis); enabling Mars sample return to earth, providing mobile access to increasingly high risk, scientifically rich areas; and broadening robotic operations to teams of cooperating agents, e.g., robot work crews that support one another's objectives (coordinated assembly, inspection, maintenance of both science and habitat) and extended robotic presence (health maintenance & self-repair).

In the sections ahead we report on our approach to some of these problems. We begin with a summary of FIDO rover, Section 2, whose computing and electronics architecture is shared by a number of our other robotics research vehicles, and whose technical concept and terrestrial field experimentation support the NASA MER'03 payload. We progress from FIDO work to Section 3, next describing a smaller class of light, agile, highly autonomous research rovers that we concurrently developed. These vehicles, the Lightweight Survivable Rover (LSR) and Sample Return Rover (SRR) have novel mechanical design, materials structure and on-board sensory guidance. In particular, SRR provides a very rich technical infrastructure for visually-guided navigation, manipulation and the integration of these two functions in precision field rendezvous and payload transfers.

As an outgrowth of these efforts on rover science/ autonomy, we recently began work on *terrain-adaptive mobility*. The objective, detailed in Section 4, is to have a rover that adapts, in a physically optimal, behaviorally intuitive way, to variable terrain—reacting autonomously, quickly and definitively to perceived changes in terrain characteristics so as to improve stability and maneuverability. Factors of immediate concern include gravitational instability on steeper slopes (tip-over, high-siding); variable traction, floatation and drift (soil-tire mechanics), compensation of vehicular dynamics at high speeds, and "de-trapping" from surrounding obstacles and encumbrances.

Utilizing a version of SRR with an actively-articulated suspension and re-positionable center of gravity (c.g.), we have developed and experimentally demonstrated on-board autonomous rover reconfiguration in response to visually/inertially sensed changes in terrain slant-tilt. We explicitly model for the terrain interaction kinematics, idealized surface friction effects, and from these obtain a metric strategy for optimizing stability of traverse. With this technique, we have successfully extended the range of SRR operations from

12-to-15 degree slopes (given a fixed suspension geometry) to 40-to-50 degrees (variable shoulder strut angles and c.g. offset by repositioning of attached manipulator). Beyond this initial effort, our work progresses toward "reconfigurable mobility": a rover, or collective robotic modules at higher levels of granularity, that can autonomously restructure their overall electro-mechanical and sensor-based control organization in optimal response to environmental and internal system states.

Following this, in Section 5, we next discuss a new robotics architecture—CAMPOUT (Control Architecture for Multi-Robot Planetary Outposts)—for closely coordinated operations of two or more mobile robots operating under tight kinematics and force constraints. CAMPOUT utilizes hierarchically layered and collective *behaviors* to enable efficient distributed sensing, communications, and control among an extensible set of robots. We have implemented preliminary demonstrations of CAMPOUT on two SRR-derived platforms, with experiments on the coordinated manipulation and transport of large payloads and coordinated access to areas of terrain slope in excess of 60-to-70 degrees.

2. Rovers for in situ Science Exploration

There is a relatively well-established art for autonomous terrestrial mobile robots, to much lesser degree, robots for remote planetary surface science (Weisbin et al., 1999; Volpe et al., 2000). Various conceptual prototypes are under development and field testing at institutions worldwide (e.g., CMU and NASA/JPL in the US, CNES and CNRS in France, NASDA/MITI in Japan et al.). As noted above, NASA has planned a related Mars mission for 2003 wherein two such planetary rovers, operating independently, will carry out investigations of surface geology, mineralogy, atmospheric and biological features. Toward this end and NASA missions beyond, our group is developing the Field Integrated Design & Operations (FIDO) rover (Schenker et al., 2001).

FIDO is a technology integration and mission operations testbed for semi-autonomous in situ science exploration. The general operational paradigm for this class of rover is as follows: Based on down-link panoramic imagery, as obtained from a rover-mounted camera, scientists will designate nearby target(s) of interest to which the rover navigates via intermediate way-points. These are panorama coordinate locations referenced to the world frame, possibly situated on features recognizable by on-board sensing (which taken together constitute part of the sequence planning). The rover visually detects and avoids local obstacles, while also updating its absolute trajectory coordinates. Localization over longer distances is confirmed by ground analysis—comparing the actual latest rover imagery with views expected based on estimated position, as derived from onboard sensing.

In the particular case of FIDO, remote command and control is implemented with *WITS* (Web Interface for TeleScience), a JPL-developed toolset for cooperative, geographically distributed Mars robotic science operations (Backes et al., 1999). WITS (example shown in Fig. 2) provides resources for science planning, 3D pre- and post-visualization of sequences, uplink command-telemetry, science & engineering data product downlink display and more. See also http://wits.jpl.nasa.gov.

Table 1 lists FIDO's major design features. FIDO capabilities include wide-area panoramic imaging (a mast-mounted color stereo pair), 3D terrain mapping and hazard avoidance (B/W stereo navigation camera on mast; chassis-mounted front/rear stereo), visual self-localization (visual map registration/tracking), local

Table 1. FIDO rover system features; for more detailed information on the various rover subsystems, see the JPL FIDO public web site http://fido.jpl.nasa.gov.

Mobility and manipulation
- 6-wheel rocker-bogie, all wheels independently driven/steered
- Max speed 9 cm/sec, 20 cm wheels, ground clearance 23 cm
- Multiple mobility modes (turn-in place, "crab", passive/active wheel drive); max obstacle clearance ∼1.5 wheel diameters
- Rover dimensions, 1.0 m (L) × 0.8 m (W) × 0.5 m (H); 68 kg mass
- 4 d.o.f. articulated mast with integral science instrumentation
- 4 d.o.f. fully actuated and instrumented front science arm

Navigation and control
- PC104+, 266 MHz Intel Pentium, PCI/ISA bus, 64 MB RAM
- ANSI C software architecture under V×Works 5.3 real-time OS
- Front/rear hazard avoidance stereo camera pairs (115° H-FOV)
- Mast-mounted navigation stereo camera pair (43° H-FOV)
- Inertial measurement unit (IMU) and CCD-based sun sensor
- Differential GPS for ground-truth reference of traverse

Science instrumentation
- Mast-mounted multi-spectral stereo camera pair (650, 740, 855 nm, 10° FOV, .34 mrad IFOV); full extent is 1.94 m
- Mast-mounted near-infrared point spectrometer (1.3–2.5 microns, 9.3 mrad projective field of view)
- Arm-mounted color micro-imager (RGB color, 512 × 496 pixel, 1.5 × 1.5 cm^2 FOV at approx. 3 mm standoff), and Mössbauer spectrometer; arm reach is ∼50+ cm)
- Rover-mounted Mini-Corer with belly stereo camera

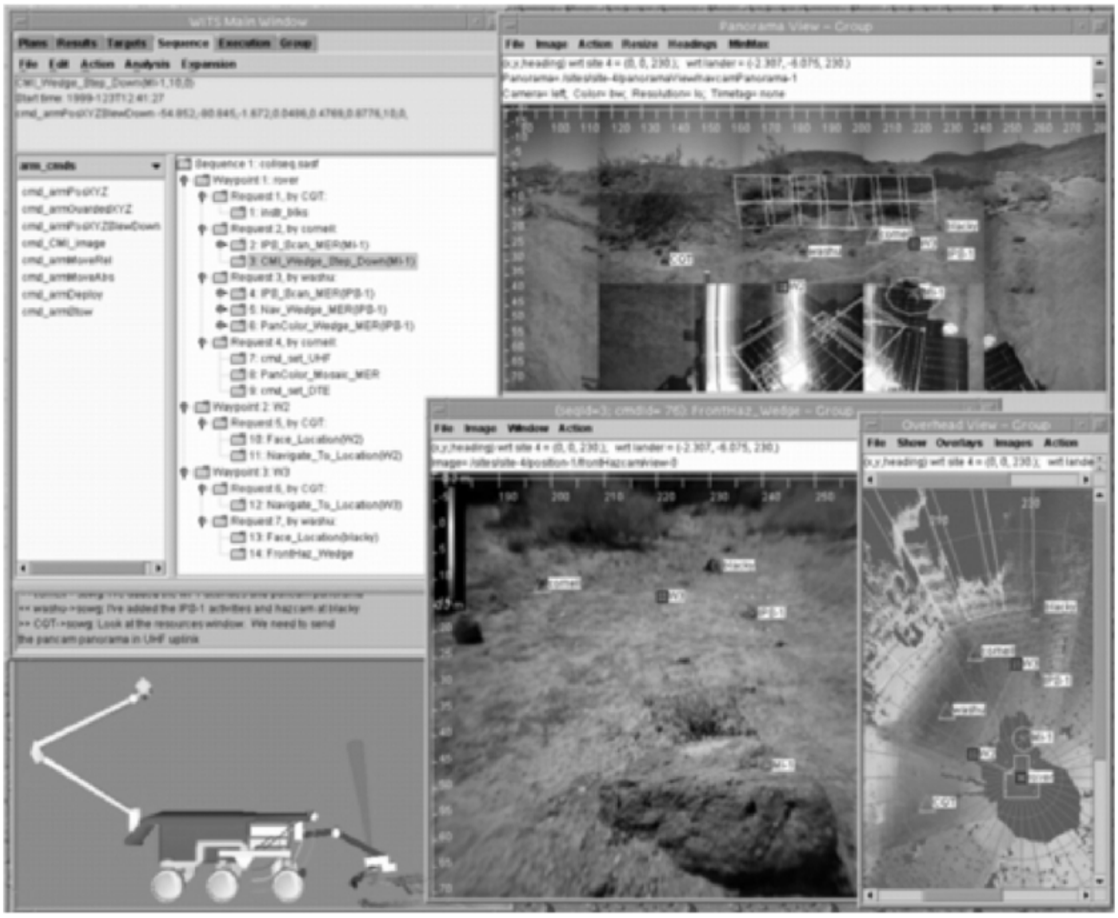

Figure 2. Web Interface for Telescience (WITS) display as seen by a single operator at a PC/Unix-based workstation.

path planning (via area stereo/navcam maps), inertial and celestial navigaitional reference (accelerometers, gyros, and sun sensor), and finally, fused state estimation supporting long range navigation (viz., statistical integration of odometry, visual, inertial, sun sensor and other data sources via Extended Kalman Filtering and related techniques (Hoffman et al., 1999; Baumgartner et al., 2000)).

We are characterizing FIDO—the rover design at large; underlying sensing, control, sampling technologies; and remote science operational strategies—in an increasingly challenging set of science field trials under direction of NASA's MER'03 flight science team (PI Steven Squyres, Cornell University, co-I Raymond Arvidson, Washington University). A first trial at Silver Lake, California, in the Mojave desert (1999) demonstrated a full "local sampling loop" about a putative lander site: panoramic imaging from the lander area, 3D navigational mapping to ground-designated targets of interest, open-loop traverses to selected targets, bore-sighted IPS imaging of targets in both stand-off scanning and proximity pointing modes, kinematics-referenced 3D visualization and placement of mast/arm instruments and tools, targeting and extraction of rock samples, and return to the immediate area of the lander.

A sequel field trial in spring 2000 at Black Rock Summit, Nevada, added significant elements of mission realism and complexity. In particular, operations were *blind* and fully remote. The science team controlled the rover by satellite communications from JPL; their prior site information was limited to large area thematic and descent imagery typical of real Mars orbital observations. This trial was followed by another in spring 2001 in the Soda Mountains of the Mojave Desert in California, in which the resource and time constraints of the MER mission were explicitly included in the command sequencing and data downlink activities. The MER science team with the MER mission operations personnel ran a 20-sol emulation that serves as direct training for the upcoming mission.

The first action of the *Science Operations Working Group* (SOWG) stationed at JPL was to acquire a full panorama looking out ∼50-to-100 meters and correlate this extensive visual data set with the multi-source overhead thematic visible and infrared imagery (incl. LANDSAT7+, TIMS, calibrated AVIRIS, typically at 10-to-30 meters^2 per pixel resolution; available data also included uncalibrated aerial photographs, perspective views, etc.). Once so "situated," the SOWG performed a prospective analysis of nearby targets of opportunity, ranking their science values against hypotheses about geological and mineralogical structure. Some targets were close enough to allow an immediate near-IR analysis via pointing of the mast-mounted IPS. This work done, the SOWG picked primary targets and commanded rover approaches.

The terrain, as shown in Fig. 3, was quite challenging and rich. This motivated a very "opportunistic" incremental exploration, in which the investigators frequently stopped the rover, deploying its arm-mounted micro-imager to examine ground soils and rocks en route to a primary target. A sense of the overall activity is depicted in Fig. 3, the rover having already acquired and down-linked a panorama, now beginning its local science in near field of the 1:1 scale lander mock-up.

In the aggregate, this simulated Mars in situ science using FIDO rover was akin to terrestrial field geology (Arvidson et al., 2000, 2001)—a nonlinear process of scientific discovery and discernment wherein multiple hypotheses were incrementally formed based on initial data and area history, then progressively updated, refuted or confirmed, with at times the overall investigation being redirected as a new observation of yet higher perceived priority was made (example of data products shown in Fig. 4). The SOWG, science investigators and engineering/operations staff learned a great deal in this multi-week experiment. Insights were gained about: (1) preferred science operational strategies and command-data sequencing protocols under realistic time and bandwidth constraints; (2) limitations and impacts of open loop localization of the rover and instrument arm placement during target acquisition (processes involving coordination of rover motion with inverse kinematics positioning of arm-mounted instruments, also, positioning of a rover-mounted mini-corer, wrt. terrain maps derived from hazcam-bellycam-navcam); and (3) continuing directions of development for 3D visualization (supporting rover activity planning and instrument operations), resource models for sequence planning (time, power, data volume, etc.), command-dictionary structure, downlink telemetry processing, and finally, automated report generation

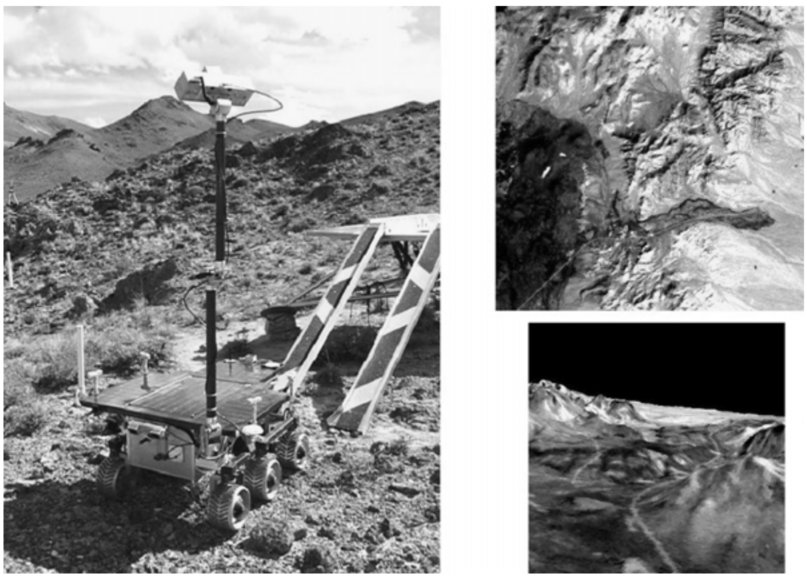

Figure 3. FIDO Rover at the Nevada blind field test, egress from lander complete, and beginning its science mission. Pictures at upper and lower right: composite LANDSAT data and LANDSAT overlay of 3D TIMS reconstruction. See also http://wufs.wustl.edu/fido and examples, Fig. 4.

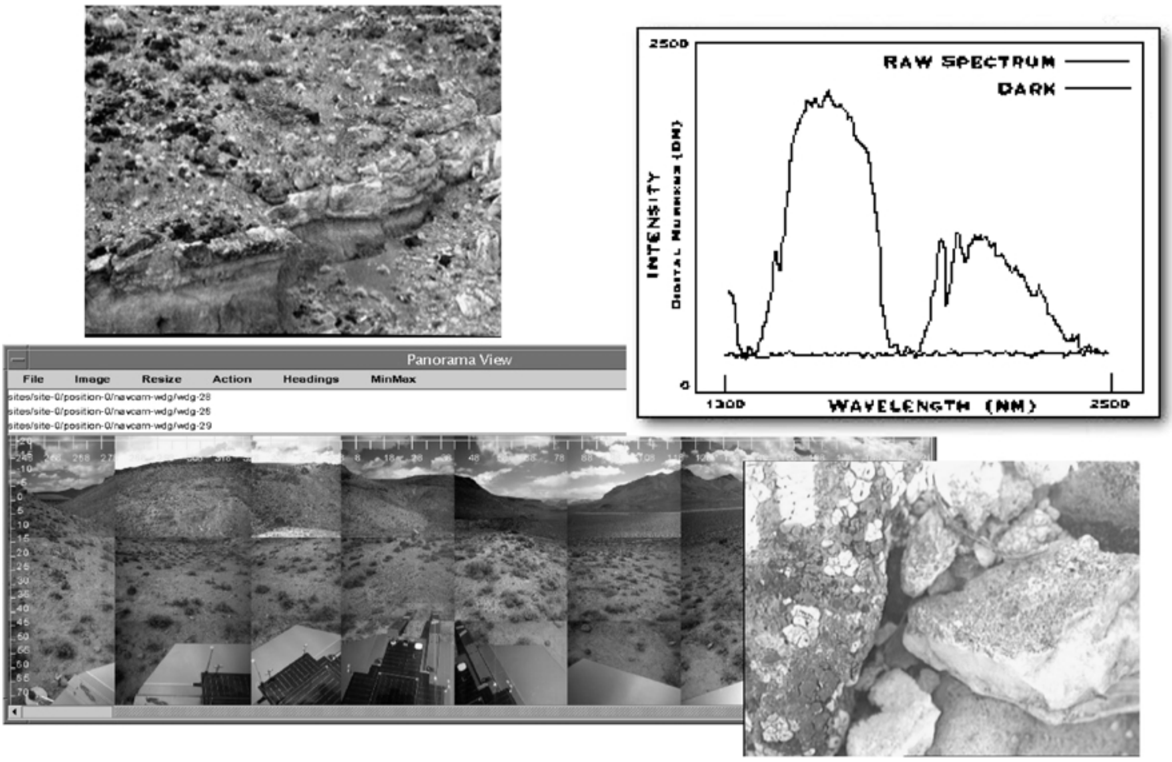

Figure 4. Representative data products from FIDO Rover field test at the Nevada. Upper left, near-field Panoramic Camera sector; lower left, Navigational Camera mosaic; upper right, Infrared Point Spectrometer analysis of target; lower right, close-up of ground rock structure taken with micro-imager. See also http://wufs.wustl.edu/fido.

& data archiving, and overall task simulation within WITS.

In summary, robotics experiments at this level of system integration and scale yield not only significant insights to component technology capabilities and limitations, but also serendipitous findings about system operations, e.g., the interactive staging of PanCam, IPS and micro-imager observations during driving; trends in resource utilization; and; the most useful roles and relative merits of visualization and simulation tools. At the conclusion of 2000 field experimentation, FIDO science investigators noted they had "... learned new ways to do rover science we wouldn't have thought of [sic]"... (in addition to evaluating scope and technical feasibility mission activities). See also Schenker et al. (2000c) [Mars'03 PI Steve Squyres' interview within on rover science, roles of field testing, etc.].

3. Rovers and Mars Sample Return

The current NASA Mars mission focus is mobile in situ science from larger platforms, of which the FIDO rover is representative; there are later plans for a Mars Sample Return (MSR). In such scenarios, a rover not only carries out sample access and in situ analyses, but also performs sample extraction, containment, and ultimately a sample transfer for orbital ascent. There are many approaches to MSR, each based in system trades against launch payload allocations, orbital transfer characteristics, rover/robot complexity, science diversity, mission cost and risk, etc. Various implementations might include: (1) sampling from a fixed lander and launch to orbital transfer (or direct-to-earth from accompanying ascent vehicle); (2) a "sample grab" using a simple, small, light rover and lander-based ascent vehicle (assumes an also simplified mechanism of sample transfer); (3) in-field retrieval of samples from an already present science rover (or field repository) by a small, faster "sample return rover", and from there to a lander with an integral Mars ascent vehicle (MAV); (4) direct and likely repeated rendezvous of the science rover with a lander-MAV complex for sample cache transfers; or (5) incorporation of a MAV and appropriate planetary protective containment directly in a large science rover. Several of these scenarios, particularly (3) and (4), pose significant robotics challenges in the area of precision autonomous visual terminal guidance/servoing.

In recent years we have developed a class of light, fast, agile, small, and volume-efficient planetary rover designs that have attractive properties for Mars surface missions such as those described above. These robots are research prototypes. They are less mature than the FIDO rover, indeed in some cases are the sources of algorithms, computing architecture, and mechanical designs that continue to be frequently migrated to FIDO, and from there to flight. These new research vehicles have been operated in fairly realistic terrestrial scenarios, e.g., variable outdoor terrain, and stand-alone continuous driving sequence simulations of conceptual mission models and functions. The rovers are highly autonomous, can traverse surprisingly rugged terrain (Viking Lander 1/VL1 densities or greater), and carry robot arms that can be autonomously sequenced and visually servoed for instrument deployment, sample cache pickup, and payload transfers. The potential range of applications for these 5-to-10 kg, 20+ cm/sec instrumented rovers, used individually, or in cooperative multi-robot activities (later reported) is broad: local area, lander–based operations; mid-range over-the-horizon (OTH) science, precursor scout missions, sample cache rendezvous and retrieval, networked science, human-robot interactive field activities/support (astronaut's "tag-along"), and tightly coordinated multi-rover Mars outpost operations.

Our earliest rover work was motivated by objectives of providing a vehicle that was very light, small in stowage, thermally robust, and would significantly advance near-NASA flight capabilities beyond Sojourner—including more richly instrumented science. The driving concern of was limited launch payload resources (mass/volume/etc.). The resulting development, *Lightweight Survivable Rover* (LSR) (Schenker et al., 1997c), shown in Fig. 5, fostered several novel component technology advances, particularly in mechanical design.

These developments included all-composite linkages, integrated structural-thermal chassis (WEB), a collapsible mobility structure (wheels that auto-deploy from 30% volume), and high terrain-ability. LSR, $\sim 45 \times 70 \times 100$ cm^3 in size, has over four times the deployed volume of Sojourner, weighs about 1/3 less at 7 kilograms, and has successfully traversed almost Viking Lander 2 terrain. LSR, in the spirit of a near-

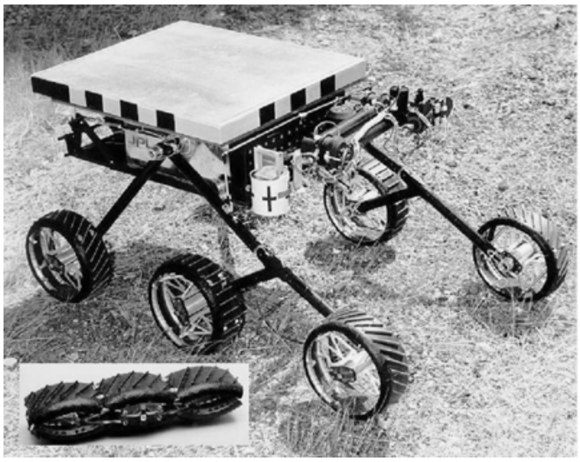

Figure 5. Lightweight Survivable Rover (LSR), with rear-mounted sampling arm, driven by rotary ultrasonic motors (Schenker et al., 1997). The arm carries an opposable gripper, rock abrader, and color micro-imager. The affixed "can" is a simulated sample container. The 20 cm dia. wheels are collapsible (*inset*) and self-deploying on compression release. LSR carries a front-mounted, multi-spectral science imager.

term mission application, was initially demonstrated with Sojourner flight avionics—a rad-hard 80c86—which was sufficient to support a hybrid on-board laser/CCD-stereo spot-pushbroom sensor (Schenker et al., 1997c) for coarse 3D terrain mapping and local obstacle detection- avoidance. LSR terrain traverses were 1.5-to-2 cm/sec, by comparison to Sojourner's .3-to-.4 cm/sec (in both cases requiring a "stop-and-look" mode of operation during obstacle avoidance (Compare: FIDO and other current R&D rover capabilities for continuous motion are 6-to-10 cm/sec or more). Later LSR experimentation utilized a 68040/VME computing environment.

The LSR was first demonstrated in local area, open-loop (odometry-based) navigation over sandy surfaces, en route to pre-designated targets. LSR autonomously navigated in-course obstacles, maintaining approximate heading, and with a final visual confirmation of target approach by the ground operator, successfully performed multi-spectral image acquisition with an on-board position-able front-mount camera.

Following on this and related work with LSR, we began development of a new small rover for different purposes. *Sample Return Rover* (*SRR*), Fig. 6 and Table 2, is a Pentium-class autonomous rover having fully collapsible and self-deploying mobility, and is based on a hybrid metal-and-composite design, stowing to $\sim 25\%$ field volume (Schenker et al., 1998; Huntsberger et al.,

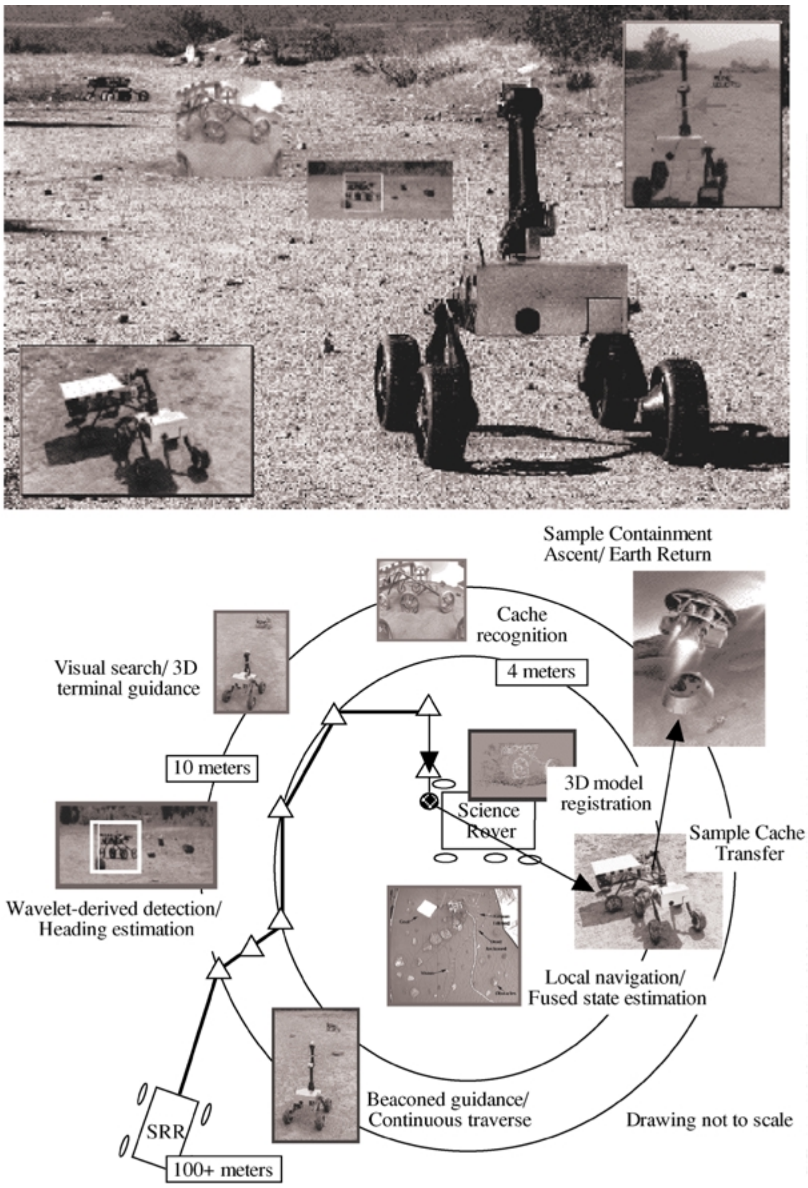

Figure 6. Rover-to-rover rendezvous, sample cache retrieval in early SRR field testing at Arroyo Seco near JPL.

1999a). SRR when deployed is ~35 × 55 × 85 cm^3 and 9 kg and also has a capability to independently set its two shoulder strut angles via an internal active differential, e.g. to alter undercarriage obstacle clearance and and/or vehicle stance and weight distribution. SRR carries a very novel all-composite 70 cm robot arm (Schenker et al., 1997a), the end-point positioning of which (for science instrument placement and object or cache retrieval) is automatically controlled by on-board visually-servo functions, including some capabilities for autonomous detection of objects of interest. In Fig. 6, SRR is seen performing the final phases of a simulated semi-autonomous "in-field" Mars rendezvous and sample cache retrieval. LSR is acting as the target "science rover", and is being approached from a distance of about 20 meters. SRR invokes a series of visual goal detection, navigation, terminal guidance, 3D object recognition-and-localization, and arm-servo functions for closed-loop control of both rover positioning and mobile manipulation (Huntsberger et al., 1999a). We have also implemented a similar longer range scenario from 100+ meters wherein a RF beacon

Table 2. Summary of SRR design features; see Schenker et al. (1998) and Huntsberger et al. (1999a) for detailed information on the various rover subsystems and experimentation.

Mobility and configuration control
- 4 wheel independent steering, fully instrumented, 3 N-m & 3 rad/sec; 20 cm dia. encoded wheels, 19 N-m & 21 cm/sec
- Passive, instrumented, rocker-type suspension with active spur-gear differential articulated shoulder joint
- Parallel linkage on suspension enables simultaneous operation of articulated shoulder/passive rocker/steering

Manipulation and effectors
- Fully instrumented 4 DOF (pitch, roll, yaw, lateral translation) gimbal with compliant gripper for extended payload
- Alternatively, 4 DOF instrument arm (MicroArms 1, 2) with end-arm opposable grippe for sample acquisition/cache transfer

Computing and electronics
- Pentium 266 MHz/32 MB, V × Works 5.4 RTOS, Solid State Disk; Ethernet (1.5 Mb/s) wireless modem; 24 v batt. pack/1.0–1.5 hr.
- 2 × 4-axis mot-ctrl., 2 × 640 × 480 color frame-grab, 12 bit × 16 ch D/A
- Front mounted stereo b/w pair, 120° FOV for hazard avoidance
- Arm-mounted color pair 45° FOV for "mast" observations
- Arm-mounted 20° b/w FOV elbow camera for goal tracking

initially guides SRR inward to the target rover. A number of important functions flowed out from this work, per illustration below, and were further integrated/developed in JPL rover activities to next be described. These functions included: fused state estimation for high accuracy rover position-and-pose determination, wavelet-based visual detection of man-made objects and features, 3D visual registration with an observed dense feature set (for relative localization of the two rovers), computationally efficient, means of robot arm visual end-point positioning with minimal pre-calibration (cache pick-up, instrument placements, etc.). The collage of Fig. 7 shows LSR and SRR in various related activities, focusing visually guided detection/guidance/rendezvous.

More recent development of the SRR concept focused on its extension to *rover-to-lander* terminal rendezvous over significant distances, at high levels of on-board autonomy. This work culminated in a demonstration of SRR that, like the last given example of rover-to-rover rendezvous, was conducted in the Arroyo Seco near JPL, Pasadena, CA. By comparison, the scope, complexity, and continuity of operations in Fig. 8 were somewhat greater than in Fig. 6. This work involved autonomously detecting a Mars'03 replica lander geometry from over 125 meters, tracking to a mid-range of 20+ meters, visually acquiring a more detailed multi-point map of lander locations of interest, approaching closely (several meters), then developing a very accurate and robust fused feature map of lander structure, moving into closure of a meter or less, with final registration of SRR to within 1-to-3 cm and 1-to-2 degrees accuracy at the lander ramp entry point. This was all done under *sequentially staged autonomy*, starting from fairly arbitrary approach directions.

This integrated operational capability is obviously important for MSR-type missions, as it implies great time efficiency in sample return legs for MAV transfer and launch (as opposed to nominal multi-day uplink-downlink cycles). Further, note that the two just described precision terminal rendezvous paradigms—including visual manipulation from mobile platforms—will be generic to rover-to-rover and rover-to-lander/habitat surface operations across many science, servicing, and human-robot outpost scenarios.

4. Rovers and High Risk Access Operations

The logical and desired evolution of science rovers would be to more all-terrain capabilities. There are numerous known and posited areas of the Mars surface that are not currently within safe reach of conventional rover designs, yet promise to be high in science content. For example, there have been recent orbital observations suggesting water out-flows and extended regions of rich mineralogy near cliff edges. Thus, development of mechanization and control architectures that enable roving into adverse, challenging terrain—areas changing markedly over short distances—is of considerable importance. Such developments of course also have terrestrial applications (military, rescue, etc.).

We have recently conducted related R&D based around SRR, with emphasis on having the rover autonomously adapt its real time control behaviors and configuration to observed/estimated terrain conditions and internal state. Figure 9 reflects one direction we have undertaken in our supporting research. The general philosophy is to have the rover image its forward-looking terrain, build a 3D local map, analyze traversability characteristics relative to kinematics/quasistatics based maneuverability-stability of progress, and enact an appropriate behavior to optimize a rover performance index. The behavior is implemented on SRR in terms of reposing its stance and c.g. This is done in two ways: by independent articulation of the rover shoulder strut angles, and repositioning of the rover top-mounted robot arm. Per Fig. 9,

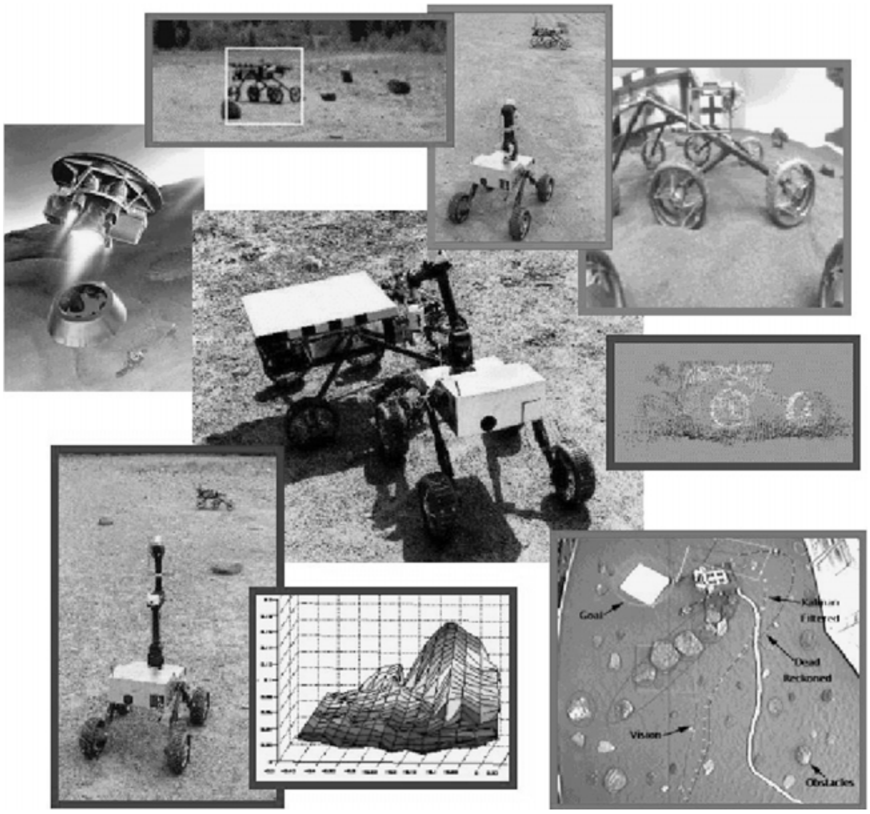

Figure 7. Starting from lower left, the SRR (cache retrieval rover) in near-field approach to LSR (science rover) and mid-field obstacle avoidance; Mars ascent vehicle (MAV) depiction; wavelet-based image localization of LSR from SRR goal camera; terminal goal-camera guidance and staging for normal vector approach; eigenvector-based recognition/localization of cache; 3D feature set of LSR used in final approach 3D registration/localization; rover experiment on fused visual tracking/odometry navigation; and bottom center, derived 3D map for hazard avoidance. In the middle, visually referenced sample cache pick-up.

the arm is treated as a "reconfigurable resource" to be used in both kinematically unconstrained and closed-loop fashions, e.g., in the latter case, as another drive actuator, pivot point, or other causal element in rover-ground interactions. No consideration is yet given to rover dynamics, as this is not a major contributory factor in the 5-to-10 cm/sec regime and low mass/volume envelope we are treating. We do account, however, in conjunction with both thin and super-quadric surface contact models, for static friction and slip effects. We discuss this at length in Schenker et al. (2000b) and references therein, including details of the terrain state estimation. Figure 10 sketches a related experimental infrastructure that we have developed.

In summary, our interest is in predicting the future state of the rover based upon look-ahead stereo range imaging, on-board IMU, and any other derived state information that can be sensed, e.g., stall conditions, inferred slip from accelerometry; etc. The general approach we have pursued in characterizing the rover-ground interaction and obtaining a measure of optimal (re)configuration is as follows (Schenker et al., 2000b):

1. Determine the surface shape of terrain ahead of the rover (model by appropriate spatial representation).
2. Solve the configuration kinematics to predict rover configuration on the modeled terrain, i.e. roll, pitch, yaw, internal angles, and wheel contact points.
3. Given a friction coefficient that characterizes wheel-ground interactions, determine if the span of nominal frictional and normal forces at the predicted contact are sufficient to resist the gravity wrench (and any other disturbance forces) in both the nominal and re-configured kinematics/c.g. (Reconfiguration consists of independent left-right shoulder angle changes and center-of-gravity shifts using the manipulator).

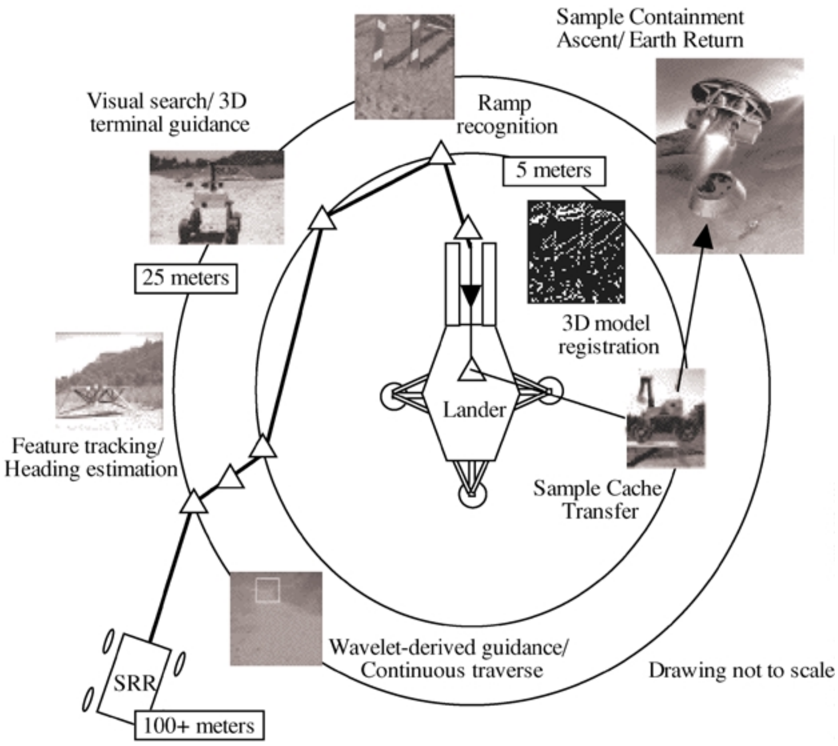

Figure 8. Rover-to-lander visual rendezvous for sample return cache transfer, performed at Arroyo Seco near JPL.

Posture & Mobility Modes ⇨ / Detectable/Predictable Conditions ⇩	Center of Gravity Rebalance	Shoulder Raise/Lower	Arm/Ground Contact	"Belly Down"	"Crabbing"	Join/Split & Tethering
Traction Loss	Visual Odometry & Wheel Current				Roll/Pitch Sensing & Wheel Current	
Steep Slope	Roll/Pitch Sensing & Range Map	Roll/Pitch Sensing & Range Map		Range Map	Roll/Pitch Sensing & Wheel Current	
Wheel Trap			Wheel Current			
Support Loss			Acceleration & Tilt Sensing			Acceleration & Tilt Sensing
Crevasse						Range Map
Tip-Over			Roll/Pitch Sensing			

Figure 9. Mobility reconfiguration in response to adverse terrain conditions (matrix entries are "trigger conditions").

4. Determine the minimum coefficient of friction in Step 3. This term is interpreted to be a *Locomotion Metric* indicative of the quality of the given configuration (or reconfiguration).

Step 1 is achieved by stereo imaging—that is correlating L/R images along epipolar lines to establish disparity, and consequently the range, via a camera model. Step 2 is achieved by means of an iterative Newton Solver. Step 3 involves setting up polyhedral inequality approximations to the friction cone at each contact point, and expressing as inequalities the unidirectional constraints on the wheel normal forces and the wheel torque constraints. These linear relationships are then transformed to the vehicle frame using the vehicle kinematics. An equality constraint characterizes

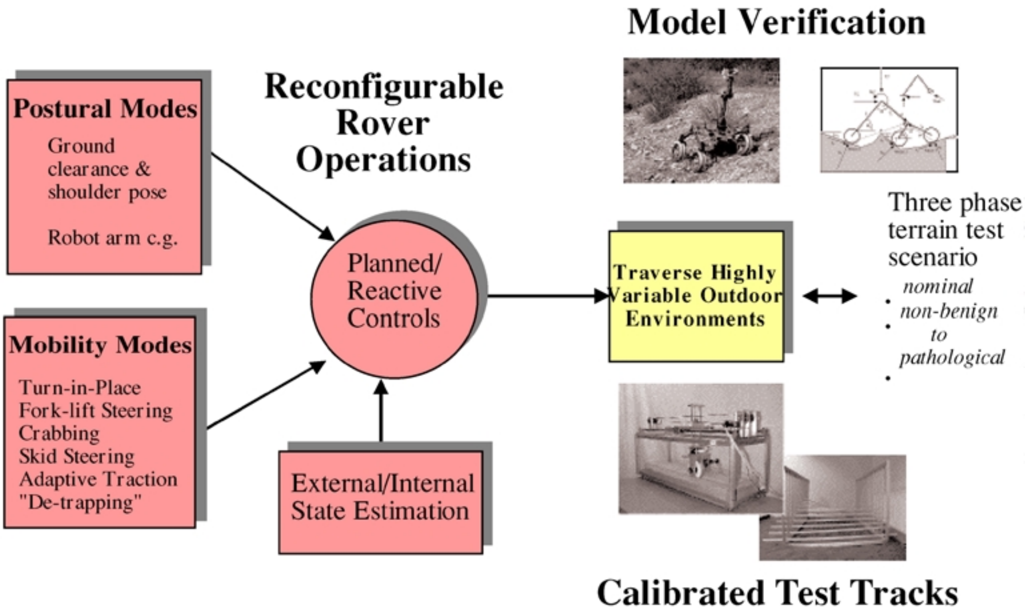

Figure 10. Experimental concept for development and test of SRR reconfigurable control over adverse terrain.

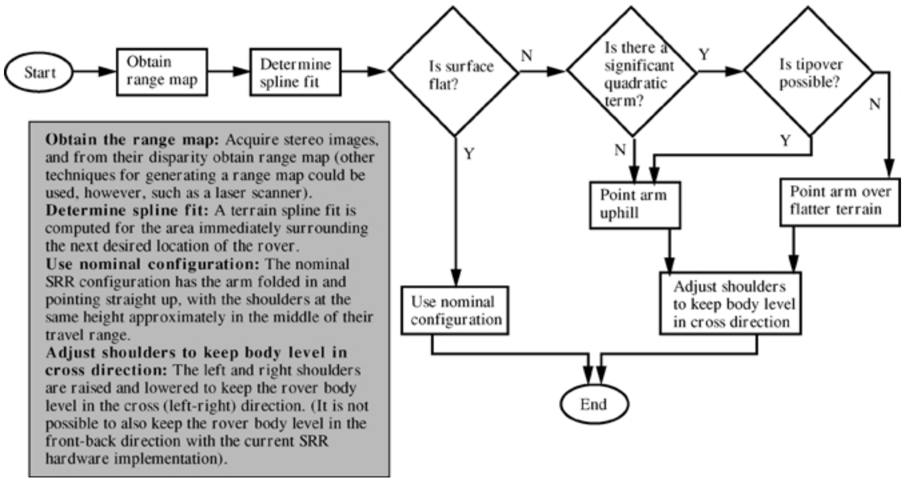

Figure 11. Simple control adaptation scheme for the SRR reconfiguration in response to perceived terrain model.

the manifold of contact forces able to resist the applied wrench without regard to constraints. A linear programming solution uses these inequality and equality constraints to determine if a feasible set of friction and normal forces exists to resist the applied wrench. A binary search algorithm then computes the metric by determining the smallest value of friction coefficient that suffices to resist the applied vehicle wrench.

We are experimenting with different control strategies and levels of modeling detail in implementing this concept for SRR reconfiguration. Figure 11 is one somewhat simplified, approximating approach through which we have achieved promising results to date, in some cases enabling the rover to make stable descents of ∼50 degree slopes and perform ascents and cross-traverses of 30 degrees, per Fig. 12.

We note also our collaboration with MIT colleagues in the development of a complementary approach (Iagnemma et al., 2000, 2001). A stability metric is defined using a quasi-static model and optimized on-line. This method relies on estimation of wheel-terrain contact angles as derivable from simple on-board sensors

Figure 12. SRR in steep descent on hill at Arroyo Seco near JPL. By comparison to a fixed geometry and nominal c.g. (arm stowage position), rover tip-over instability is greatly improved.

alone. Again, due to the slow speed of the rover as operated here (less than 6 cm/sec) only static forces are considered in calculating the rover stability. System stability is expressed in terms of a set of "stability angles." A stability angle is the angle formed between a line originating at the center of mass and normal to the tip-over axis, and the gravitational (vertical) axis. This angle goes to zero at marginal stability. A performance index, Φ, is defined for the SRR from these stability angles $\gamma_j, i = \{1, \dots, 4\}$, and the reconfigurable shoulder degrees of freedom, ψ_1 and ψ_2 as:

$$\Phi = \sum_{j=1}^{4} \frac{K_j}{\gamma_j} + \sum_{i=1}^{2} K_{i+4}(\psi_i - \psi_i')^2$$

where K_i and K_j are positive constants and the stability angles γ_j are functions of the shoulder and manipulator degrees of freedom (i.e., $\gamma_j = \gamma_j(\psi_1, \psi_2, \theta_1, \theta_2, \theta_3)$). Note that the first term of Φ tends to infinity as the stability at any tip-over axis tends to zero. The second term penalizes deviation from a nominal configuration of the shoulder joints, thus maintaining adequate ground clearance, an important consideration in rough terrain. The goal of this stability-based kinematics reconfigurability optimization problem is to minimize the performance index Φ subject to joint-limit and interference constraints. For rapid computation, and due to the simple nature of Φ a basic optimization technique such as conjugate-gradient search are employed. Overall the approach is computationally simple, does not require or exploit visual terrain knowledge, and focuses its concerns to gravitational versus traction issues.

Returning to earlier-described general analysis, we note that the Locomotion Metric is sensitive to the fact that the most stable configuration (which implies using an infinite coefficient of friction) may not be the most advantageous one from the viewpoint of slip or traction. Indeed, those configurations that concentrate the weight on the "flatter" parts of the terrain are to be preferred, trading stability for slip resistance. If the vehicle is unstable, then even an infinite friction coefficient is unable to generate the resisting forces, resulting in the Locomotion Metric being infinity. A finite value of the Locomotion Metric indicates that sliding (or loss of traction) is inevitable if the terrain/wheel coefficient of friction drops below the value indicated in the metric. Availability of the metric allows the current configuration of the vehicle shoulders and center-of-mass to be compared to adjacent configurations. The configuration with the lowest possible metric is a candidate for vehicle reconfiguration and is recommended to the vehicle on-board controller.

In general, modeling of vehicle-terrain interactions remains a somewhat open problem for more than idealized media, and certainly for the frangible, sandy, variable soils of Mars. We continue to gain related experimental experience, both at JPL and MIT, through instrumented *lateral motion test tracks* like that shown in Fig. 10. As we describe in Section 6—after introducing some necessary background on the enabling CAMPOUT rover architecture—we are now extending the type of system development above to multiple, modular, cooperative robots, wherein collective sensing and control is utilized to effect desired system-wide reconfiguration and mobility functions.

5. Rovers that Cooperate

There are a number of surface mission concepts that could benefit from, and directly motivate, distribution of activity across multiple rover platforms (Huntsberger et al., 2000c, 2001a). We have, in fact, already discussed one limiting case of this in Section 2, where we described a form of "passive" cooperation in which one rover performs precision rendezvous with another rover/robot for purposes of manipulative sample cache pick-up, transfer, etc. A predominant driver for the use of multiple robots is future Mars outposts,

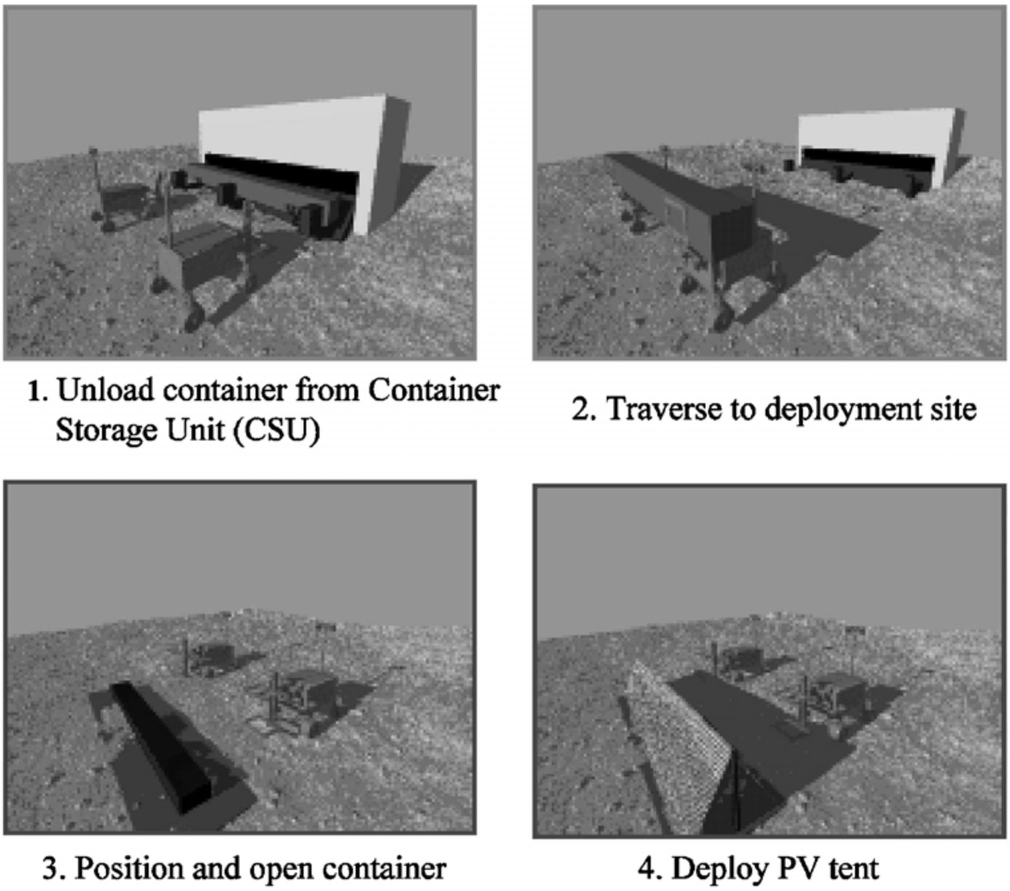

Figure 13. Robots cooperate in deployment of Photovoltaic (PV) tent array on Martian surface (graphic simulation).

wherein robots would act as precursors to human exploration, and once that human presence is achieved, continue to be infrastructure for a sustained habitation. The precursor role would clearly focus to such activities as power station deployment, assembly of habitats and other life-support functions, one such example shown in Fig. 13 below.

There are related roles for cooperating planetary robots, including their surface-based applications in assembly and deployment of science facilities. E.g., multiple, closely coordinated robots may be needed to assemble/service large aperture optical instruments, and/or to distribute and implement "networked" observational/science concepts—ranging from local area activities such as incoherent 3D imaging (long-baseline stereo) to geographically dispersed structures with on-line metrology and/or inter-localization.

What we outline here is our recent development of an architecture for, and experimental implementation of cooperative rover activities focused on shared physical tasks (Schenker et al., 2000a; Huntsberger et al., 2000, 2001, 2002b). The obvious analogy is a human work crew in construction, where, as shown in Fig. 13 above, two or more crew workers are called upon to carry an extended object over sparse, obstructed terrain—object acquisition, transport, and deployment (not a piano mover's problem, but one that does require a high degree of shared state knowledge!). Challenges to this *Robot Work Crew*, as we call it, are major, in that achieving a general performance requires tight, instantaneous coordination of kinematics and force constraints between the two robots over variable terrain, all subject to pre-emptive behaviors that manage obstacles and anomalies, within a non-holonomic space. Much prior related work treats problems of multi-rover cooperation as "sequential interactions", versus actual closed loop real-time kinematics coordination with force constraints; work that does address "tight coordination" is in most cases is restricted to idealized environments, e.g., flat lab floors (see Huntsberger et al.

(2002b) for a comprehensive review). Real terrain operations are significant; we have found in simulation and practice that as little as 2° differential inclination of the rovers/payload can introduce complications. We present our overall approach in Schenker et al. (2000) and Pirjanian et al. (2000, 2001), including the relevant priors, and give details of our underlying architecture in Huntsberger et al. (2002b). Here, we briefly sketch our concept, the major architectural features, and two recent significant field experiments in natural terrain.

5.1. Tight Coordination of Multiple Mobile Robots

Figure 14 is an overview of CAMPOUT; this is basically a hierarchical architecture, one functionally derivable from the types of environments in which planetary rover systems are expected to operate and survive. A long duration mission such as a robotic outpost on a planetary surface has a wide ranging needs—from low-level, reactive components supporting local navigation and manipulator control, to high-level planning of large-area tasks. CAMPOUT spans a range of tactical-strategic requirements via low-level control drivers directly tied to actuators, commanded in turn by a behavior-based control hierarchy, overseen by a higher deliberative task planning layer (we are at present investigating tasks that do not require such higher level planning). CAMPOUT is highly distributed. An example of a CAMPOUT behavior hierarchy for the transport phase of the PV tent deployment scenario (Fig. 13) is shown in Fig. 15.

The advantages of distributed control and coordination include the efficient use of system resources, parallel execution of multiple tasks, reliability and fault-tolerance to failure of individual components, including failure of single robots. Behaviors within a single robot operate in a distributed manner, thus allowing concurrent and/or parallel execution of several tasks. However, each robot can operate on its own, independent of other agents, based on its inherent faculties of perception and action. Cooperation between the multiple robots occurs through active collaboration—there is no centralized planning or decision-making to dictate explicit commands. Note that *reactive behaviors* (Huntsberger et al., 1999b, 2000; Arkin, 1998) facilitate tight perception-action feedback loops that react promptly to unexpected situations, guided by *deliberative plans* for efficient use of global resources. In effect, such plans guide, but do not dictate, the control of reactive components.

CAMPOUT provides coordination mechanisms that are specifically tailored for not only cooperative, but also tightly coordinated tasks. Behaviors are organized in a hierarchy where abstract behaviors are built upon less abstract behaviors and so on. Each behavior has an objective that it pursues by coordinating its subordinate behaviors. Thus, behaviors can have two roles in an agent: as *actions* and as *action selection mechanisms*. With respect to its subordinates, a behavior is an action selection mechanism; with respect to its superior a behavior is viewed as an action to be implemented. This approach is attractive for its low computational and communications overhead.

5.1.1. Primitive Behavior Library. The main architectural substrate in CAMPOUT consists of a *behavior producing module* (commonly known as a *behavior*). A *behavior* is a perception-to-action mapping module

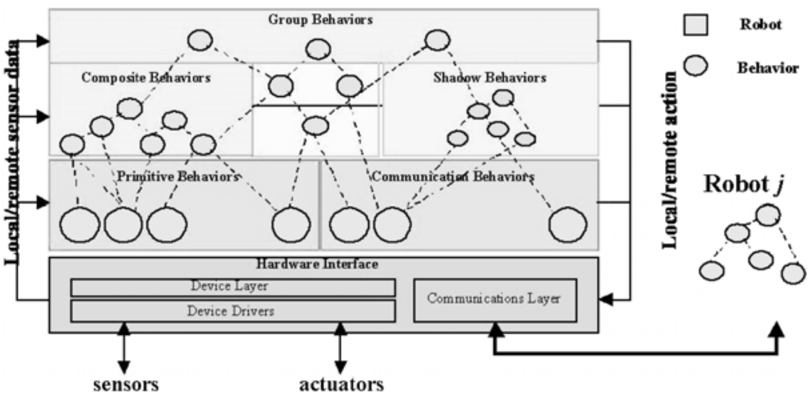

Figure 14. CAMPOUT high-level organization (Pirjanian et al., 2000; Huntsberger et al., 2002).

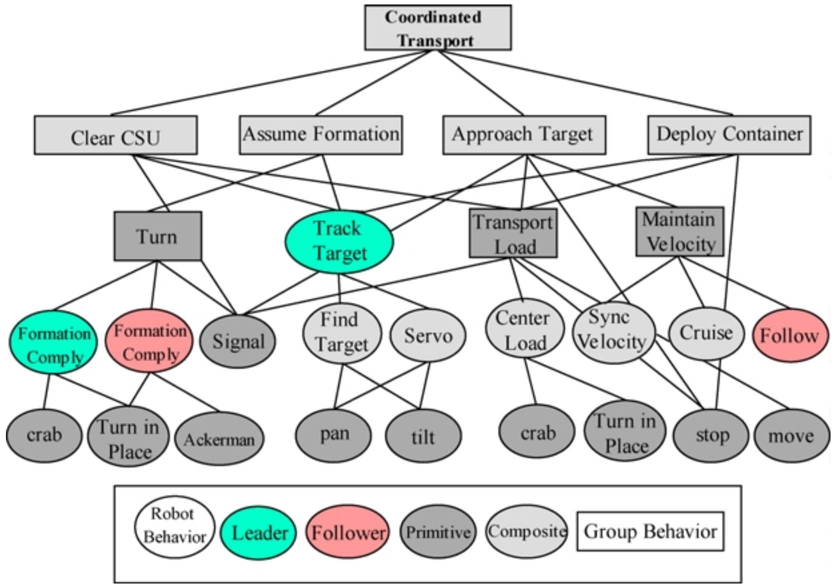

Figure 15. CAMPOUT "behavior hierarchy" describing a coordinated transport task (see Fig. 13, frame 2, graphic). Bubbles represent single robot behaviors; boxes represent multi-robot "group" coordinated behaviors. Higher-level actions, themselves behaviors, are composed from yet lower-level behaviors (Pirjanian et al., 2000).

that based on selective sensory information produces recommendations for actions in order to maintain or achieve a given, well-specified task objective. For example, for safe navigation the system will often require a minimum of two behaviors: *AvoidObstacle* for safety and *GotoTarget* for navigation. Note that the *Avoid Obstacle* behavior is concerned with obstacle avoidance and obstacle avoidance only. Similarly, the *GotoTarget* behavior is only concerned with controlling the robot towards a target and is not concerned with obstacle avoidance at all, nor is it aware of the existence of the obstacle avoidance behavior. Behaviors can have conflicting objectives and hence require efficient *behavior coordination mechanisms* (BCM) to resolve such conflicts—viz., produce a useful combination of the behaviors into higher level behaviors, known as composite behaviors (*c-behaviors*).

5.1.2. Composite Behaviors. Composite behaviors are constructed through careful combination of lower-level behaviors. At the lowest level of such a behavioral hierarchy, there are primitive behaviors (p-behaviors), which constitute a library of core capabilities for a robot. By coordinating the activities of primitive behaviors, one can construct a composite behavior that enhances the skill set of the robot. Composite behaviors can further be constructed from other (lower-level) composite behaviors or a mix of primitive and composite behaviors. For instance a composite behavior, *SafeNavigation*, can be constructed from primitive behaviors *AvoidObstacle* and *GotoTarget* by a simple fixed priority-based coordination of the two that enables the *AvoidObstacle* behavior when the robot is close to obstacles, and the *GotoTarget* behavior when the path of the robot is obstacle free.

5.1.3. Behavior Coordination Mechanisms. Behavior coordination mechanisms (BCMs) provide tools for integration of behaviors to achieve higher-level goals. Priority-based behavior coordination represents a very primitive (but some times useful) type of coordination. CAMPOUT is an open architecture in the sense that any other behavior coordination mechanism can easily be integrated; this is an important property. Basically a BCM will be implemented as an operator (analogous to the logical AND or OR operators) and used to compose behaviors. These operators are also provided in our *Behavior Composition Language* (BCL), in one form or another, as a high-level description language for behavior composition. Note that we have recently done a detailed study of related behavior coordination mechanisms (Pirjanian, 2000).

5.1.4. Communication Behaviors. The primitive and the composite behaviors constitute the skill set that enables a given robot to interact with and accomplish tasks in its environment. For cooperation and interaction with each other, robots are required to communicate, thus they require basic behaviors for communication at various levels of task abstraction. Note that communication is not necessarily limited to explicit exchange of information via a hard data link, but can also include visual, auditory, tactile, and other types of implicit communication. For instance, one robot could visually determine relative position of another, or make inference via shared kinematics and force constraints (a payload, one end carried by each) from purely local information. Alternatively, the other robot could explicitly transmit its position within a global coordinate system. CAMPOUT provides the methodology and infrastructure that support all such approaches, as we indicate by a later experimental example.

5.1.5. Shadow Behaviors. The above communication behaviors provide information necessary to facilitate cooperation between a team of robots. This information is encoded in form of shadow behaviors (*s-behaviors*) that basically represent a remote behavior, including full state information, running on a separate robot. This allows physical behaviors to have a collective access to remote, distributed sensing/actuation capabilities which is equivalent to treating the system as a single entity, and ultimately can be the basis for aggregate and evolutionary *networked robotics* properties (McKee and Schenker, 2000; Fryer et al., 1997).

5.1.6. Group Behaviors. In order to cooperate and collectively contribute to a common task objective, the robots must cooperate and coordinate activities. S-behaviors facilitate composition of high-level behaviors that can achieve such coordination and formally need not be treated any differently than p- or c-behaviors. Indeed their interfaces are exactly the same. Thus, coordination among several robots can be achieved by a simple coordination of the activities of (a subset of) s-behaviors and the robot's own behaviors using suitable behavior coordination mechanisms (including fusion, arbitration and several other mechanisms discussed and compared in depth in Pirjanian (1999, 2000)).

5.2. Architecture Implementation

In this sub-section, we briefly describe the current implementation of CAMPOUT components and infrastructure; we illustrate this within the scope of the cooperative transport task earlier noted—two robots transporting an extended payload across uneven terrain. This task involves full use of CAMPOUT's facilities. As noted earlier, in its current implementation, CAMPOUT implements no high-level planning capabilities, and we focus on behaviors here. Further details of the CAMPOUT implementation can be found in Pirjanian et al. (2000, 2001) and Huntsberger et al. (2002b).

CAMPOUT provides these resources for coordination and control of physically distributed robots:

- *Behavior Representation*: a set of abstract data types known as objective functions and related operations to describe the output of a general behavior as a multi-valued preference.
- *Behavior Prototyping Toolkit*: provides a set of tools for rapid-prototyping of primitive as well as composite behaviors, i.e., facilities that can be used to easily develop behaviors. The general behavior representation used in CAMPOUT does not suggest—or prohibit—any particular approach to behavior implementation; CAMPOUT does, however, provide specific tools for developing behaviors, currently built around rule-based and state-machine representations. The CAMPOUT toolkit for synthesizing behaviors utilizes fuzzy control. Using standard fuzzy inference models, e.g., *max-prod*, we combine rules into a multi-valued output that encodes the (grade of) desirability of each action from the behavior's point of view.
- *Behavior Coordination Mechanisms*: provides a repertoire of mechanisms that can be used to coordinate the activities of lower-level behaviors to form higher-level composite behaviors. CAMPOUT provides a set of coordination mechanisms that can be used for action sequencing, conflict resolution, priority-based behavior invocation, and context-dependant behavior invocation—activities related to composition of higher level behaviors. Specifically, CAMPOUT provides a set of complementary behavior arbitration mechanisms including finite-state machines, and subsumption-style arbitration. CAMPOUT also provides command fusion mechanisms based in multi-valued logic (*MVL*) and

Figure 16. Coordinated transport of extended container (2.5 meters) by SRR and SRR2K, as performed in Arroyo Seco near JPL. (Left) row transport formation; (Right): column (leader-follower) transport formation.

multiple objective decision-making (*MODM*) approaches (Pirjanian, 2000).

- *Communications Infrastructure*: provides tools and functions for interconnecting a set of robots and/or behaviors in order to share resources (e.g., sensors or actuators), exchange information (e.g., state, percepts), implement synchronization, etc. CAMPOUT provides diverse facilities for this within its behavioral communications infrastructure, the current implementation of which is based on UNIX-style sockets. These facilities currently include communication behaviors for synchronization, data exchange, and behavior exchange.

5.3. *CAMPOUT in Action*

Objects that are four to five times the length of a single mobile platform are extremely difficult to manipulate and transport. The *Robot Work Crew* (*RWC*) concept assumes use of multiple rovers for coordinated operations on such an extended payload, with examples of *row* and *column* transport being shown in Fig. 16. These tightly coordinated multi-robot operations are implemented on the SRR platforms, described in Section 3. The baseline SRR design is reported in Schenker et al. (1998) and Huntsberger et al. (1999a), wherein it incorporated skid steering and basic functions for stereo-based obstacle detection, continuous motion visual traverse (10+ cm/sec), visually-servoed manipulation, and in-field visual object detection, tracking, and rendezvous. More recently, as summarized in Table 2, we augmented the SRR design with 4-wheel steering, improved computational resources, the CAMPOUT behavioral control architecture, and gimbaled grippers that support compliant payload handling (Fully actuated approaches to transport of extended structures may not be realistic for planetary surface operations due to mass/power constraints). We initially are investigating a fully instrumented passive gripper design per Fig. 17.

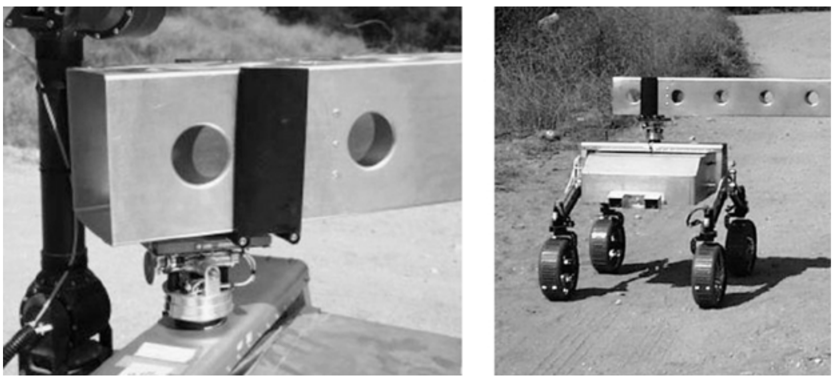

Figure 17. Instrumented gimbal (SRR2K close-up at left).

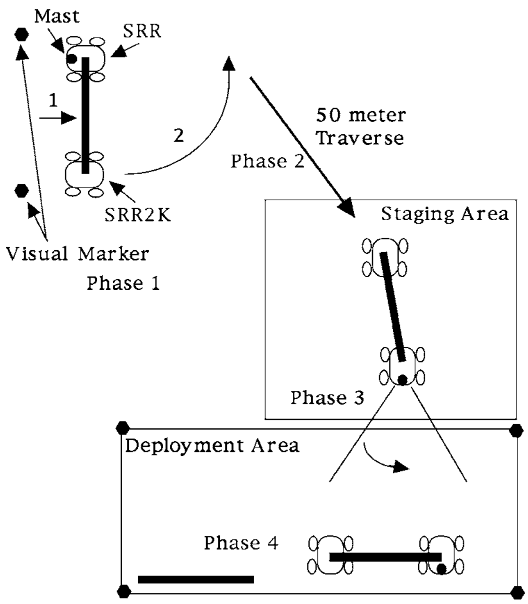

Figure 18. (1) Initiate transport configuration, (2) move to staging area, (3) initiate site survey, and (4) dock into site.

The gimbal is attached to a cross brace that spans the shoulders of the SRR and has 3 DOF force sensors and potentiometers for monitoring movement of the container relative to the rover body. Our goal for the experimental study was the transport of an extended container (12.5 cm × 12.5 cm × 250.0 cm) by two rovers (SRR and SRR2K, the latter being a minimalist mechanization of the first) from a pickup point to a deployment zone that is up to 50 meters away, over natural terrain. This was accomplished with the four-phase sequence of Fig. 18.

We provide a detailed description of the experimental implementation using CAMPOUT in Pirjanian et al. (2000, 2001), Trebi-Olennu et al. (2002), and Huntsberger et al. (2002b), including the specific sensory-control behaviors and their higher level compositions (see also Arkin (1998) and Huntsberger et al. (1999b, 2000)). As a general strategy, we attempt to minimize explicit communication between the rovers, as reflects possible operational constraints (i.e. power use) during an actual mission. This is facilitated by using the shared container as an implicit means of communication—e.g., relative positions of the rovers are known through the yaw gimbal angle on each rover. Also, we are exploiting natural design constraints of the task where possible to assess useful trades of mechanized cooperation versus explicit closed loop controls (as one example, the use

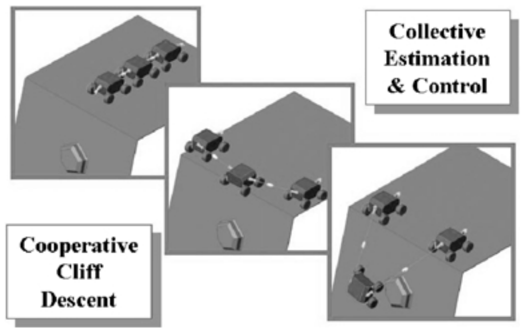

Figure 19. Multiple, modular robots reconfigure to perform a cooperative descent for analysis of cliff stratigraphy.

of passive compliance in both grippers along the beam axis).

In a second experimental validation of the utility of CAMPOUT for distributed control we designed a system for access to cliff-faces using a modified version of SRR2K as a "cliff-bot" and two anchored rover analogs as "anchor-bots" at the top of the cliff. The conceptual design is shown in Fig. 19, where a modular ensemble of robots reconfigures itself at the top of the cliff, anchoring two members at the top and sending a tethered robot over the edge onto the cliff-face. The tethered robot actively traverses the cliff-face using way point navigation, with stability being maintained by the actively controlled tethers from the anchor-bots. This system is configured under CAMPOUT using the behavior network shown for one of the anchor-bots in Fig. 20.

The BCM under CAMPOUT for this system uses a priority weighting of the four primary behaviors *Stability*, *Maintain Tension*, *Match Velocity*, and *Haul*, with *Stability* being given the highest priority. The *Stability* behavior minimizes the risk of tip-over, the *Maintain Tension* behavior keeps a constant tension on the tethers, the *Match Velocity* behavior controls the tether playout rates to match those of the active agent on the cliff face, and the *Haul* behavior gives the active agent a pull if it has insufficient torque to get moving at the start of a traverse. We ran numerous trials on the mesa overlooking JPL, an example of which is shown in Fig. 21. The system successfully performed way point navigation in any direction over the cliff face with slopes greater than 70° under the distributed control of CAMPOUT. Further details can be found in Pirjanian et al. (2002) and Huntsberger et al. (2002b).

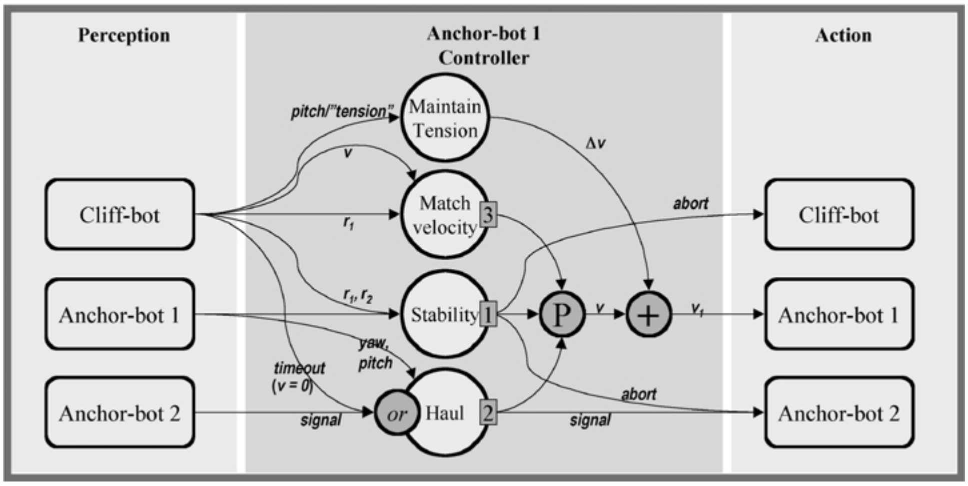

Figure 20. A subset of the behavior network for collective cliff-descent illustrating sub-system for controlling the velocity of anchor-bot 1. The arrows represent data links between local blocks as well as remote components (behaviors, sensors, actuators) thus spanning a behavior network across the team of robots.

Figure 21. Cliff-bot performing active way point navigation on a cliff face with closely coordinated active tether control by two anchor-bots at the top of the cliff.

6. Summary and Conclusions

We have described our recent work on planetary rovers. This is a partial overview, in several senses. It is part, not all that we are investigating, as further noted below. It is representative, but not encompassing, of possibilities for Mars surface missions and their technology requirements. And finally, it is only indicative of work at large as many other institutions have parallel efforts, as noted in the introduction. Hopefully the reader has found significant threads of continuity in the presentation, as these were intended. First, we note that as NASA mission planning currently sits, there is a strategic line to exploration, one linked to discovering Mars history

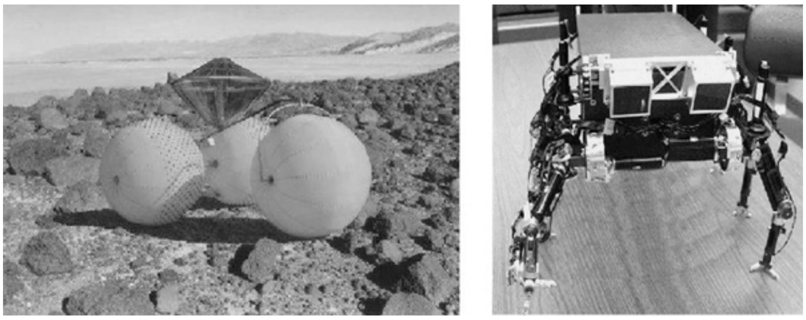

Figure 22. (L) Inflatable Rover, solar array; (R) LEMUR.

and resources ("follow the water"), characterizing those resources in detail—and, given indication that those resources could sustain human life—extended visitation and human-robot cohabitation. Second, we have described a progressive program of robot system architectures aimed at science, sample return and characterization, and outposts. As always, it is convenient to impose taxonomies—Mars "mobile science platforms, all terrain rovers, sample return rovers, robot work crews, etc." Underlying this operational breakout are deeper, less discrete threads of technology development in areas such as perception, control, planning, mobility mechanization, manipulation and actuation, robotic software architecture, simulation and the like. These technology paths reflect an upward evolution of on-board autonomy and robot system complexity. Yet in some cases there is also suggestion of a more reactive, distributed intuitive robotic "intelligence" that can be implemented in a modular fashion—robots that perform collective estimation, share control functions, allocate networked resources to global system objectives, etc. as we discuss further in McKee and Schenker (2000) and Fryer et al. (1997) and references therein.

As to our immediate future plans—we are extending the work reported here in several directions: First, FIDO goes to another desert field trial which will serve as one of the final tests of the MER mission concept prior to launch in mid-2003. The technology concepts developed for FIDO (long range navigation, single command science target approach and instrument placement, auto-focusing, etc.) (Huntsberger et al., 2002a) will be evaluated for infusion into the Mars Science Laboratory (MSL) 2009 mission. We also plan this year to implement a science mission concept for sample acquisition on a cliff-face using the combination of autonomy described earlier for FIDO and the cliff-bot system running under CAMPOUT.

Similarly, we are exploring other classes of robots for improved mobility on both natural and artificial surfaces. One interesting concept for Mars mobility, quite different in design from LSR but having common motivations, is the "Inflatable Rover" (Jones and Wu, 2000)—light, resilient, collapsible—also, potentially fast and long ranging. Another is the LEMUR (Legged Excursion Mechanical Utility Robot) (Hickey et al., 2000), conceived for a class of possible structural assembly, inspection, and maintenance activities. These two concepts are shown in Fig. 22.

Acknowledgment

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. The work involves important contributions from many colleagues at both JPL and collaborating institutions. We gratefully acknowledge these interactions and note many of the specific developments in references that follow.

References

Arkin, R.C. 1998. *Behavior-Based Robotics*, Intelligent Robotics and Autonomous Agents Series. The MIT Press: Cambridge, MA.

Arkin, R.C. and Bekey, G.A. (Eds.). 1997. Special issue on robot colonies. *Autonomous Robots*, 4(5).

Arvidson, R.E., Squyres, S.W., Baumgartner, E.T., Schenker, P.S., Niebur, C.S., Larsen, K.W., Seelos IV, F.P., Snider, N.O., Jolliff, B.L. 2002. FIDO prototype Mars rover field trials, Black Rock Summit, Nevada, as test of the ability of robotic mobility systems to conduct field science. *Journal of Geophysical Research-Planets*, 107(8), doi: 10.1029/2000JE001464.

Arvidson, R.E., Squyres, S., Baumgartner, E.T., Dorsky, L., and Schenker, P. 2000. Rover trials for Mars sample return mission prove successful. *EOS Transactions, American Geophysical Union*, 81(7):65–72.

Backes, P.G., Tso, K.S., and Tharp, G.K. 1999. The web interface for telescience. *Presence*, 8(5):531–539.

Baumgartner, E.T., Aghazarian, H., Trebi-Ollennu, A., Huntsberger, T.L., and Garrett, M.S. 2000. State estimation and vehicle localization for the FIDO Rover. In *Proc. Sensor Fusion and Decentralized Control in Robotic Systems III*, SPIE Vol. 4196, Boston, MA.

Fryer, J.A., McKee, G.T., and Schenker, P.S. 1997. Configuring robots from modules: An object oriented approach. In *Proc. 8th IEEE International Conference on Advanced Robotics (ICAR'97)*, Monterey, CA, pp. 907–912.

Hickey, G., Kennedy, B., and Ganino, A. 2000. Intelligent mobile systems for assembly, maintenance, and operations for space solar power. In *Proc. ASCE Robotics 2000 Conference*, Albuquerque, NM.

Hickey, G.S. and Kennedy, B.A. 2001. Six legged experimental robot. *NASA Tech Brief* NPO-20897.

Hoffman, B.D., Baumgartner, E.T., Huntsberger, T., and Schenker, P.S. 1999. Improved rover state estimation in challenging terrain. *Autonomous Robots*, 6(2):113–130.

Huntsberger, T.L., Aghazarian, H., Baumgartner, E., and Schenker, P.S. 2000. Behavior-based control systems for planetary autonomous robot outposts. In *Proc. IEEE Aerospace 2000*, Big Sky, MT.

Huntsberger, T.L., Aghazarian, H., Cheng, Y., Baumgartner, E.T., Tunstel, E., Leger, C., Trebi-Ollennu, A., and Schenker, P. 2002a. Rover autonomy for long range navigation and science data acquisition on planetary surfaces. In *Proc. 2002 IEEE International Conf. on Robotics and Automation (ICRA2002)*, Washington, DC, May 11–15, 2002, pp. 3161–3168.

Huntsberger, T.L., Baumgartner, E.T., Aghazarian, H., Cheng, Y., Schenker, P.S., Leger, P.C., Iagnemma, K.D., and Dubowsky, S. 1999a. Sensor-fused autonomous guidance of a mobile robot and applications to Mars sample return operations. In *Proc. Sensor Fusion and Decentralized Control in Robotic Systems II*, SPIE Vol. 3839, Boston, MA, pp. 2–8.

Huntsberger, T.L., Mataric, M., and Pirjanian, P. 1999b. Action selection within the context of a robotic colony. In *Proc. Sensor Fusion and Decentralized Control in Robotic Systems II*, SPIE Vol. 3839, Boston, MA.

Huntsberger, T.L., Pirjanian, P., and Schenker, P.S. 2001. Robotic outposts as precursors to a manned Mars habitat. In *Proc. Space Technology and Applications International Forum (STAIF-2001)*, Albuquerque, NM, 46–51.

Huntsberger, T.L., Pirjanian, P., Trebi-Ollennu, A., Das, H., Aghazarian, H., Ganino, A.J., Garrett, M.S., Joshi, S.S., and Schenker, P.S. 2002b. CAMPOUT: Tightly-coupled coordination of multi-robot systems for planetary surface exploration. *IEEE Transactions on Systems, Man, & Cybernetics*, in review.

Huntsberger, T.L., Rodriguez, G., and Schenker, P.S. 2000c. Robotics challenges for robotic and human Mars exploration. In *Proc. ROBOTICS2000*, Albuquerque, NM, pp. 340–346.

Iagnemma, K., Rzepniewski, A., Dubowsky, S., Huntsberger, T., and Schenker, P. 2000. Mobile robot kinematic reconfigurability for rough-terrain. In *Proc. Sensor Fusion and Decentralized Control in Robotic Systems III*, SPIE Vol. 4196, Boston, MA.

Iagnemma, K., Rzepniewski, A., Dubowsky, S., and Schenker, P. 2001. Control of robotic vehicles with actively articulated suspensions in rough terrain. *Autonomous Robots*, in review.

Jones, J.A. and Wu, J.J. 2000. Inflatable technology for robotics. In *Proc. ASCE Robotics 2000 Conference*, Albuquerque, NM.

McKee, G.T., Fryer J.A., and Schenker, P.S. 2001. Object-oriented concepts for modular robotics systems. *IEEE Transactions on Robotics and Automation*, in review.

McKee, G.T. and Schenker, P.S. 2000. Networked robotics. In *Proc. Sensor Fusion and Decentralized Control in Robotic Systems III*, SPIE Vol. 4196, Boston, MA.

Pirjanian, P. 1999. Behavior Coordination Mechanisms—State-of-the-art. Tech-report IRIS-99-375, Institute for Robotics and Intelligent Systems, School of Engineering, University of Southern California.

Pirjanian, P. 2000. Multiple objective behavior-based control. *Journal of Robotics and Autonomous Systems*, 31(1/2):53–60.

Pirjanian, P., Huntsberger, T.L., and Schenker, P.S. 2001. Development of CAMPOUT and its further applications to planetary rover operations: A multirobot control architecture. In *Proc. SPIE Sensor Fusion and Decentralized Control in Robotic Systems IV*, Vol. 4571, Newton, MA.

Pirjanian, P., Huntsberger, T.L., Trebi-Ollennu, A., Aghazarian, H., Das, H., Joshi, S., and Schenker, P.S. 2000. CAMPOUT: A control architecture for multi-robot planetary outposts. In *Proc. Sensor Fusion and Decentralized Control in Robotic Systems III*, SPIE Vol. 4196, Boston, MA.

Pirjanian, P., Leger, C., Mumm, E., Kennedy, B., Garrett, M., Aghazarian, H., Schenker, P.S., and Farritor, S. 2002. Distributed control for a modular, reconfigurable cliff robot. In *Proc. 2002 IEEE International Conf. on Robotics and Automation (ICRA2002)*, Washington, DC, pp. 3136–3141.

Schenker, P.S., Bar-Cohen, Y., Brown, D.K., Lindemann, R.A., Garrett, M.S., Baumgartner, E.T., Lee, S., Lih, S.S., and Joffe, B.E. 1997a. A composite manipulator utilizing rotary piezoelectric motors: New robotic technologies for Mars in situ planetary science. In *Proc. Smart Structures and Integrated Systems*, SPIE Vol. 3041, San Diego, CA.

Schenker, P.S., Baumgartner, E.T., Dorsky, L.I., Backes, P.G., Aghazarian, H., Norris, J.S., Huntsberger, T.L., Cheng, Y., Trebi-Ollennu, A., Garrett, M.S., Kennedy, B.A., Ganino, A.J., Arvidson, R.E., and Squyres, S.W. 2001b. FIDO: A field integrated design & operations rover for Mars surface exploration. In *Proc. 6th Intl. Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS-'01)*, Montreal, Canada.

Schenker, P.S., Baumgartner, E.T., Lee, S., Aghazarian, H., Garrett, M.S., Lindemann, R.A., Brown, D.K., Bar-Cohen, Y., Lih, S.S., Joffe, B., Kim, S.S., Hoffman, B.H., and Huntsberger, T.L. 1997b. Dexterous robotic sampling for Mars in situ science. In *Proc. Intelligent Robotics and Computer Vision XVI*, SPIE Vol. 3208, Pittsburgh, PA.

Schenker, P.S., Baumgartner, E.T., Lindemann, R.A., Aghazarian, H., Zhu, D.Q., Ganino, A.J., Sword, L.F., Garrett, M.S., Kennedy, B.A., Hickey, G.S., Lai, A.S., Matthies, L.H., Hoffman, B.D., and Huntsberger, T.L. 1998. New planetary rovers for long range Mars science and sample return. In *Proc. Intelligent Robotics and Computer Vision XVII*, SPIE Vol. 3522, Boston, MA.

Schenker, P.S., Blaney, D.L., Brown, D.K., Bar-Cohen, Y., Lih, S.S., Lindemann, R.A., Paljug, E.D., Slostad, J.T., Tharp, G.K., Tucker, C. E., Voorhees, C.J., Weisbin, C., Baumgartner, E.T., Singer, R.B., and Reid, R. 1995. Mars lander robotics and machine vision capabilities for in situ planetary science. In *Proc. Intelligent Robots and Computer Vision XIV*, SPIE Vol. 2588, Philadelphia, PA.

Schenker, P.S., Huntsberger, T.L., Pirjanian, P., and McKee, G.T. 2001a. Robotic autonomy for space: Closely coordinated control of networked robots. In *Proc. 6th International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS-'01),* Montreal, Canada.

Schenker, P.S., Huntsberger, T.L., Pirjanian, P., Trebi-Ollennu, A., Das, H., Joshi, S., Aghazarian, H., Ganino, A.J., Kennedy, B.A., and Garrett, M.S. 2000a. Robot work crews for planetary outposts: Close cooperation and coordination of multiple mobile robots. In *Proc. Sensor Fusion and Decentralized Control in Robotic Systems III*, SPIE Vol. 4196, Boston, MA.

Schenker, P.S., Pirjanian, P., Balaram, B., Ali, K.S., Trebi-Ollennu, A., Huntsberger, T.L., Aghazarian, H., Kennedy, B.A., Baumgartner, E.T., Iagnemma, K., Rzepniewski, A., Dubowsky, S., Leger, P.C., Apostolopoulos, D., and McKee, G.T. 2000b. Reconfigurable robots for all terrain exploration. In *Proc. Sensor Fusion and Decentralized Control in Robotic Systems III*, SPIE Vol. 4196, Boston, MA.

Schenker, P.S., Sword, L.F., Ganino, A.J., Bickler, D.B., Hickey, G.S., Brown, D.K., Baumgartner, E.T., Matthies, L.H., Wilcox, B.H., Balch, T., Aghazarian H., and Garrett, M.S. 1997c. Lightweight rovers for Mars science exploration and sample return. In *Proc. Intelligent Robotics and Computer Vision XVI*, SPIE Vol. 3208, Pittsburgh, PA.

Schenker et al. 2000c. *In Situ Exploration of Mars: The FIDO Rover*, 6:30 min narrated NASA video, JPL-Photographic Imaging Group.

Shirley, D.L. and Matijevic, J.R. 1997. Mars rovers: Past, present, and future. In *Proc. Princeton Space Studies Inst. 20th Anniversary Conf.*

Shirley, D.L. et al. 1997. The Pathfinder Microrover. *Journal of Geophysical Research*, 102(E2):3989–4001.

Trebi-Ollennu, A., Das, H., Aghazarian, H., Ganino, A., Pirjanian, P., Huntsberger, T., and Schenker, P. 2002. Mars rover pair cooperatively transporting a long payload. In *Proc. 2002 IEEE International Conf. on Robotics and Automation (ICRA2002)*, Washington, DC, May 11–15, 2002, pp. 3136–3141.

Volpe, R.A., Baumgartner, E.T., Schenker, P.S., and Hayati, S.A. 2000. Technology development and testing for enhanced Mars rover sample return operations. In *Proc. of the IEEE Aerospace Conference 2000*, Big Sky, MT.

Weisbin, C.R., Rodriguez, G., Schenker, P.S., Das, H., Hayati, S.A., Baumgartner, E.T., Maimone, M., Nesnas, I.A., and Volpe, R.A. 1999. Autonomous rover technology for Mars sample return. In *Proc. 5th International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS '99)*, Noordwijk, The Netherlands.

Paul S. Schenker is Manager, Mobility Systems Concept Development Section, Jet Propulsion Laboratory (JPL), California Institute of Technology, Pasadena, and as such, responsible for JPL R&D in planetary mobility and robotics. His published research includes topics in robotic perception, robot control architectures, telerobotics and teleoperation, multi-sensor fusion, and most recently, multi-robot cooperation. He has led the development of robotic systems that include the Field Integrated Design & Operations Rover (FIDO), Planetary Dexterous Manipulator (MarsArm, microArm), Robot Assisted Microsurgery System (RAMS), Robotic Work Crew (RWC), and All Terrain Explorer (ATE/Cliff-bot), with resulting technology contributions to NASA missions. Dr. Schenker is active in the IEEE, Optical Society of America, and SPIE. He has served as an elected Board member and 1999 President of the last; he currently serves as an elected member of the NAS/United States Advisory Committee to the International Commission for Optics.

Terry Huntsberger is a Senior Member of the Technical Staff in the Mechanical and Robotics Technologies Group at NASA's Jet Propulsion Laboratory in Pasadena, CA, where he is the Manager for numerous tasks in the areas of multi-robot control systems, rover systems for access to high risk terrain, and long range traverse algorithms for rovers. He is a member of the planning teams for the Mars Science Laboratory (MSL) mission planned for launch in 2009. He is an Adjunct Professor and former Director of the Intelligent Systems Laboratory in the Department of Computer Science at the University of South Carolina. His research interests include behavior-based control, computer vision, neural networks, wavelets, and biologically inspired system design. Dr. Huntsberger has published over 100 technical articles in these and associated areas. He received his PhD in Physics in 1978 from the University of South Carolina. He is a member of SPIE, ACM, IEEE Computer Society, and INNS.

Paolo Pirjanian was formerly a Member of the Engineering Staff in the Mechanical and Robotics Technologies Group at NASA's Jet Propulsion Laboratory in Pasadena, CA, and he is currently the Chief Scientist for Evolution Robotics in Pasadena, CA. During his tenure at JPL, he was the System Engineer for the Reconfigurable Robotic Surface Systems/All Terrain Explorer and a System Architect for the Robot Work Crews tasks. He was the primary developer of CAMPOUT, and is working towards the development of new, sustainable, and aggressive mobility systems for autonomous access to highly variable terrain, cliffs etc. using state-of-the-art multi-functional, modular, and reconfigurable sensing, controls, hardware and robotic systems. He received the M.Sc. CS. and the Ph.D. degrees in 1994 and 1998, respectively, both from Aalborg University, Denmark. Until December 1998 he held an assistant research professor position at the Laboratory of Image Analysis, Aalborg University. From January 1999 until November 1999 he was a postdoctoral research associate at the USC Robotics Research lab. working on multiple robot coordination with Maja Mataric. He has over 30 technical publications in the fields of robotics, control, and computer vision. Dr. Pirjanian serves as a part-time lecturer at the Computer Science department of University of Southern California.

Eric T. Baumgartner is the Group Supervisor of the Mechanical and Robotics Technologies Group and a Senior Member of the Engineering Staff at NASA's Jet Propulsion Laboratory in Pasadena, CA. At JPL, he serves in a systems engineering capacity for the development of advanced planetary rovers and also contributes to technology developments in the areas of robotic sensing and control. He is currently the Cognizant Engineer for the Instrument Deployment Device (IDD) for the upcoming Mars Exploration Rover (MER) dual rover mission to be launched in 2003. Prior to his tenure at JPL, he was an Assistant Professor in the Mechanical Engineering-Engineering Mechanics Department at Michigan Technological University in Houghton, MI. He has published over 50 technical articles in the area of robotics, controls, and state estimation and is active in the SPIE and ASME. He received his B.S. degree in Aerospace Engineering from the University of Notre Dame, the M.S. degree in Aerospace Engineering from the University of Cincinnati in 1990, and the Ph.D. in Mechanical Engineering from the University of Notre Dame in 1993.

Edward W. Tunstel is a Senior Member of the Engineering Staff in the Robotic Vehicles Group at NASA's Jet Propulsion Laboratory in Pasadena, CA., where he directs engineering activities for robotics technology development, and also develops autonomous control and navigation algorithms and software for planetary rover research and flight projects. He received the B.S. and M.E. degrees in Mechanical Engineering, with a concentration in robotics, from Howard University. His thesis addressed the use of symbolic computation for automated modeling of robotic manipulators. In 1989 he joined the Robotic Intelligence Group at JPL supporting research and development activities on NASA planetary rover projects. As a JPL Fellow he received the Ph.D. in Electrical Engineering from the University of New Mexico, where he became affiliated with the NASA Center for Autonomous Control Engineering (ACE). His dissertation addresses distributed fuzzy control of adaptive hierarchical behavior-based systems and its application to mobile robot navigation. He presently serves as Lead Systems Engineer for the Field Integrated Design and Operations (FIDO) rover technology task, and as a Mobility & Navigation System Engineer for the 2003 Mars Exploration Rovers (MER) flight project. His research interests include mobile robot navigation, autonomous control, cooperative robotics, and applications of fuzzy logic and other soft computing techniques to intelligent control systems synthesis. He has published over 50 technical articles in these and related areas. Dr. Tunstel is a Senior Member of IEEE and member of NSBE, AAAI, Sigma Xi Scientific Research Society, and the New York Academy of Sciences. He also serves as the NASA Technical Monitor for the NASA PURSUE University student research program administered by ACE.