

# Hybrid image plane/stereo (HIPS) manipulation for robotic space applications

Matthew L. Robinson · Eric T. Baumgartner ·  
Kevin M. Nickels · Todd E. Litwin

Received: 23 August 2006 / Accepted: 13 March 2007 / Published online: 30 May 2007  
© Springer Science+Business Media, LLC 2007

**Abstract** Manipulation systems for planetary exploration operate under severe limitations due to power and weight restrictions and extreme environmental conditions. Typically such systems employ carefully calibrated stereo cameras and carefully calibrated manipulators to achieve precision on the order of ten millimeters with respect to instrument placement activities. The environmental and functional restrictions under which these systems are used limit the operational accuracy of these approaches. This paper presents a novel approach to stereo-based manipulation designed to robustly achieve high precision levels despite the aforementioned limitations. The basic principle of the approach, known as Hybrid Image Plane/Stereo (HIPS) Manipulation, is the generation of camera models through direct visual sensing of the manipulator's end-effector. The HIPS method estimates and subsequently uses these models to position the manipulator at a target location specified in the image-planes of a stereo camera pair using stereo correlation and triangulation. In-situ estimation and adaptation of the manipulator/camera models in this method accounts for changes in the system configuration, thus ensuring consistent precision for the life of the mission. The end result is a increase in positioning precision by a factor of approximately two for a limited version of HIPS, and an order of magnitude increase in positioning precision for the full on-line version of HIPS.

**Keywords** Vision-based manipulation · Stereo triangulation · Stereo vision · Space robotics

## 1 Introduction

Manipulation systems for remote planetary exploration operate under severe restrictions not present with terrestrial counterparts (Matijevic et al. 2001; Erickson et al. 2002). For example, communication limitations such as low bandwidth and latency restrict data transfer, while power and weight limitations affect all aspects of system design by restricting robot size and workspace, actuator selection, the number and types of sensors, etc. Additionally hazardous environments require the selection of radiation-hardened components thus severely restricting computer processor speed, memory capacity, and camera frame rate. Finally, the system must be designed to operate under severe conditions, e.g. daily thermal cycling from  $-100$  to  $50^{\circ}\text{C}$  on the Martian surface, in an a priori unknown environment without maintenance and limited user intervention for the life of the mission.

These constraints and restrictions limit the precision, and more importantly, the accuracy of standard non-adaptive approaches to hand-eye coordination. The reliable performance of complex, high precision manipulation operations requires a vision guided manipulation strategy that accommodates these constraints.

This paper describes and demonstrates a novel approach to vision-based control of remote robotic manipulators designed to robustly achieve high precision levels despite the aforementioned limitations. Supporting evidence will include both simulations and experimental results. The basic principle of the approach, known as Hybrid Image

---

M.L. Robinson (✉) · T.E. Litwin  
Jet Propulsion Laboratory, Pasadena, CA 91109, USA  
e-mail: Matthew.L.Robinson@jpl.nasa.gov

E.T. Baumgartner  
Ohio Northern University, Ada, OH 45810-1599, USA

K.M. Nickels  
Trinity University, San Antonio, TX 78212-7200, USA

Plane/Stereo (HIPS) Manipulation, is the generation of camera models through direct visual sensing of the manipulator's end-effector using estimation and the subsequent use of these models to position the manipulator at a target location specified in the image-planes of a stereo camera pair using stereo correlation and triangulation. In-situ estimation and adaptation of the manipulator/camera models in this method accounts for changes in the system configuration, thus ensuring consistent precision for the life of the mission.

### 1.1 History of remote, space-based manipulation systems

Manipulation systems for planetary exploration have evolved over the past 30 years from the early use of a telescoping sampling device on the Viking Landers in the 1970s to a single degree-of-freedom mechanism that was used to deploy the Alpha Proton X-Ray Spectrometer (APXS) from the Sojourner rover during the Mars Pathfinder mission in 1997 (Matijevic 1998). In an effort to increase the dexterity and available work volume for the placement of multiple science instruments, the Mars Polar Lander mission carried a four degree-of-freedom robot arm to be used for soil trenching and digging as well as placement of the Robotic Arm Camera (RAC) (Bonitz 1997).

Currently, the Mars Exploration Rover (MER) vehicles carry a five degree-of-freedom robot arm (known as the Instrument Deployment Device, or IDD) that is used to place three in-situ instruments (the APXS, a Mössbauer spectrometer, and a microscopic imager) as well as place and hold a Rock Abrasion Tool (RAT) in order to abrade the weathered surface of a rock (Squyres et al. 2003). An artist's illustration of the MER rover with the IDD deployed is shown in Fig. 1.

Both the upcoming Phoenix Lander (launch in 2007) and Mars Science Laboratory (MSL, launch in 2009) carry robot arms, and will utilize these manipulators to accomplish in-situ science of increasing complexity. Phoenix will carry an oven and a portable laboratory on-board, and will utilize the robotic arm to scrape samples from the Martian surface and transport them to the instruments (Smith 2004). MSL will study Martian soil and rocks in unprecedented detail, utilizing the robot arm to carry samples from the surface to on-board test chambers for chemical analysis (Savage and Cook-Anderson 2004).

With the increasing demand for a higher level of science return in future surface missions, lander and rover-mounted robotic arms must exhibit a higher level of performance over current capabilities. As the capabilities and dexterity of each instrument deployment approach have evolved and improved over the last 30 years, the requirements associated with instrument placement precision have also increased. As a recent example, the MER mission requirements set for the IDD includes a precision placement requirement of 10 mm



**Fig. 1** An artist's concept of the Mars Exploration Rover. Image Courtesy NASA/JPL-Caltech

in position and 10 degrees in orientation with respect to a science target when the IDD is deployed from a stationary rover base (Baumgartner et al. 2005).

### 1.2 Flight state-of-the-art: stereo triangulation/calibration

The current state-of-the-art in manipulation for planetary exploration relies solely on the use of a well-calibrated system to achieve the required precision with respect to instrument placement activities (Baumgartner et al. 2005). Typically the manipulation process is separated into two independent steps: the determination of the target range using a stereo camera pair and the subsequent control of the manipulator to the measured 3-D location. This process involves two separate calibrations. The first calibration step determines the intrinsic and extrinsic parameters of a camera model relating the mapping between the 2-D image plane of each of the stereo cameras and a physical 3-D "world" coordinate frame. The second step involves kinematic calibration that is concerned with the determination of an accurate relationship between the manipulator joint angles and the 3-D location of points on the manipulator's end-effector by way of the manipulator kinematic model. This step is necessary due to differences between the geometric model of the manipulator and actual dimensions as a result of manufacturing tolerances and deformations of various kinds such as link flexure under load (Baumgartner et al. 2005). Included in this step is the transformation from the manipulator 3-D base frame to the 3D "world" frame.

The difficulty associated with the standard approach is that sources of error tend to accumulate, ultimately reducing positioning precision, due to the separation of the stereo vision and manipulator location processes. Sources of error in the manipulator location process include kinematic uncertainties in manipulator link lengths and reference frame

transformations, unmodeled flexure of the manipulator links and base, and joint position uncertainties due to sensor resolution and backlash in the drive train. Additional errors involve imprecise stereo calibration and ranging accuracy that can be significant. For example, with a baseline of ten centimeters and a nominal target range of one meter, error analysis predicts that range accuracy errors alone account for nearly six millimeters (Maki et al. 2003).

Briefly the vision-guided manipulation process works as follows (Baumgartner et al. 2005). Each camera model is generated by identifying locations on a calibration fixture. Using these models the 3-D range to an identified target then can be determined via stereo correlation and triangulation. From this 3-D range information the joint rotations that position the manipulator at the desired location in 3-D space are determined using the arm's inverse kinematics. In this way the vision and manipulation steps are separated.

The results of terrestrial operations such as the Field Integrated Design and Operations (FIDO) rover field trials, which consisted of tests designed to closely simulate Mars mission operation scenarios with a fully-instrumented rover in Martian-like terrain as shown in Fig. 2, have demonstrated the ability of such techniques to yield precision levels on the order of 10 mm (Tunstel et al. 2002). However, on planetary missions, degradations in positioning precision may result from a decline in calibration fidelity due to changes in the system configuration as a result of environmental factors such as vibration during launch and landing, extreme thermal cycling, and inclement weather conditions. As will be described in Sect. 5.3, this type of degradation was seen on the MER rovers, yielding insufficient instrument placement accuracy. While the nominal level of precision is sufficient for current mission operations, future mission directives call for tasks of significantly increased complexity requiring millimeter level precision. In particular, the MSL Mars exploration mission proposes a long-term, long-range rover with



**Fig. 2** The Field Integrated Design & Operations (FIDO) rover performs manipulator operations during a field trial designed to closely emulate Mars mission operation scenarios in Martian-like terrain

soil/rock sampling and precision placement of samples into a contained science laboratory for in-situ processing and analysis as a precursor to a planned Mars sample return mission.

### 1.3 A comparison of hand-eye coordination approaches

Significant effort has been devoted to the problem of autonomous hand-eye coordination both in the laboratory and the factory floor, with varying degrees of success. One such approach that has received much attention is image-based visual servoing (Hutchinson et al. 1996; Feddema et al. 1993). In image-based visual servoing, a feedback control loop comprising the difference between the current and goal manipulator states as measured in the camera image-plane is used to drive the manipulator to a zero image-plane error state.

While visual servoing has achieved some success in unstructured environments outside the laboratory (Urmson et al. 2001), there are several major advantages of using HIPS in space applications. Image-based visual servoing usually relies on continuous updates of the manipulator-target image error (Hutchinson et al. 1996). Low-update solutions have been proposed, but usually require models of the well-known tracked target (Feddema et al. 1992). In space-based applications, camera-rate control is impractical due to the limited processor speed and available camera frame rate. The resulting long delays could create controller instability and/or final positioning error (Hutchinson et al. 1996). In addition, the target is chosen to achieve some specific science return, which often results in occlusions of targets in the image-plane near the manipulator terminus due to the limited choice of camera and manipulator configurations. Although vision as a sensing modality provides capabilities beyond other types of sensors, it is well known that vision sensors are prone to intermittent failure (Hager et al. 1998). Such events may lead to catastrophic failure in visual servoing systems.

Conversely, the HIPS approach is limited by neither frame rate nor constant access to image-plane error. As new samples become available the camera models are updated and the goal position is refined. However, the manipulator can still be controlled to the target in the absence of new information. This is similar to an iterative look-then-move structure (Hutchinson et al. 1996), but provides an option to update the target as the move progresses.

Flight systems also require predictability and repeatability in vision guided manipulation. The HIPS approach allows rover operations planners to verify planned manipulator motions before sending them to the rover, and allows the use of the robust and well-tested manipulator control hardware and software, while providing the significant improvements in manipulator performance.

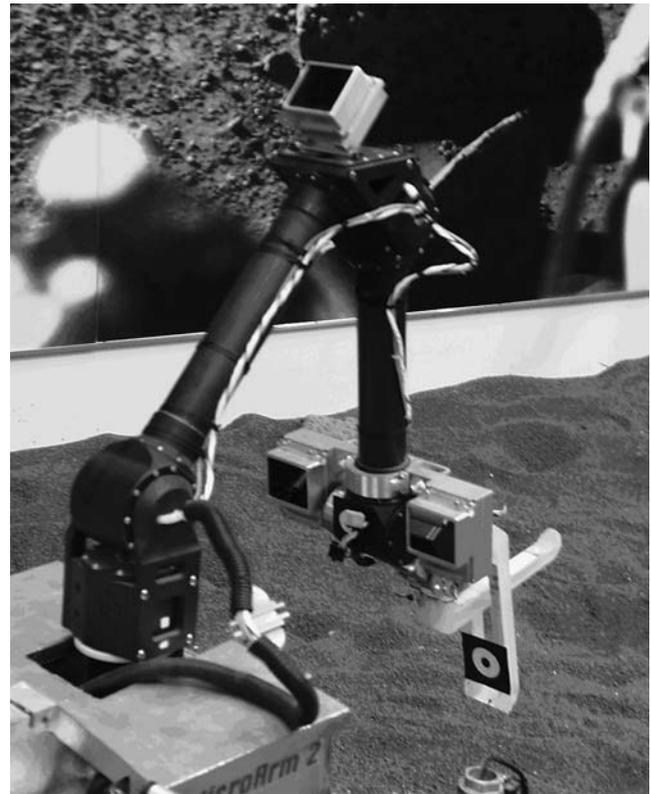
Visual servoing utilizes only the most recent image (or more typically, a set of a few recent images) to compute the image error, i.e. the control variable. A fundamental limit on terminal precision then is the error associated with target extraction from a single image. It is well known that discretization of possibly noisy images can lead to significant errors, possibly as high as one pixel (Allen et al. 1992). Alternatively, the HIPS approach yields improved precision by estimation of the manipulator-generated camera models based on a history of image-plane appearances and internal joint angles of the manipulator. Thus, precision in the HIPS manipulation technique is not limited by the image processing errors of a single image.

HIPS is most similar to an alternative technique for hand-eye coordination known as Camera-Space Manipulation (CSM) (Skaar et al. 1987). CSM is an open-loop, estimation-based, approach to the control of robotic manipulators. In the CSM approach widely separated cameras are utilized to determine the direct relationship between the joint rotations of a manipulator and the image-plane appearance of fiducial marker(s), called cue(s), on the manipulator's end-effector. This approach is based on the nominal kinematic model and the orthographic camera model. The inverse problem then is solved to determine the joint rotations that, when realized, will locate the end-effector at the desired location in the image-plane of each participant camera without regard to any physical reference frame.

The CSM method has been shown to achieve excellent terminal precision (less than 1 mm position and  $1.0^\circ$  orientation) when the participating cameras are widely separated with a vergence of greater than 60 degrees. Unfortunately, the placement of widely spaced cameras on a rover platform is difficult to achieve due to the finite size of a rover and the use of existing platform cameras configured as stereo pairs for rover navigation. The CSM approach becomes unstable as the camera vergence approaches zero due to the use of the orthographic camera model (Chen et al. 1994). This is precisely the case for stereo camera pairs. Since the addition of manipulator-specific cameras is unattractive due to strict mass, volume, power, and stowage volume constraints placed on planetary explorers, an alternative approach must be undertaken. The motivation for the development of HIPS was to achieve the precision capabilities of CSM using stereo camera systems.

The HIPS manipulation technique uses the basic principle of the CSM approach—the generation of camera models through visual sensing of fiducial marker(s) on the manipulator's end-effector as seen in Fig. 3 and the subsequent use of these models to position the manipulator at a target location specified in the image-plane of each camera of the stereo pair. The generation of camera models using parameter estimation will be detailed in Sect. 2.

CSM and HIPS diverge in the solution of the inverse problem. In the CSM approach the inverse problem is solved



**Fig. 3** Fiducial Marker Mounted on Manipulator. A fiducial marker mounted on the end-effector of the manipulator facilitates fast and reliable image processing

directly using a least-squares minimization to determine the joint rotations that will locate the end-effector at the desired pose in the image-plane of each participant camera. Alternatively, the HIPS approach uses stereo correlation and triangulation with the manipulator-generated camera models to determine the range to the target. Thus, with the range computed by stereo triangulation, the inverse kinematic model is used to solve for the joint rotations of the manipulator that place the end-effector at the desired target location.

The actual target range may be quite different from the computed target range. The important point is that the range computed by the manipulator-generated models is accurate with respect to the manipulator's coordinate system. Said another way, *the image-plane to joint-space mapping is highly accurate*. Since the target is specified in the image-plane of each stereo camera, the accuracy of this (invertible) mapping yields precise location of the manipulator.

HIPS is also similar in philosophy to Ruf's work on non-metric visual servoing (Ruf and Horaud 1999). Ruf concentrates on servoing in projective space, and dynamically estimates the joint-space to image-space mapping in much the same way as HIPS, but requires specific motions of the manipulator (termed trial motions) to create the projective mapping. HIPS can utilize any observed motion of the end-effector, and the resulting models will generalize locally to

the region of the workspace where data were gathered. This feature is of great benefit when the manipulator is a highly contested resource on another planet, as is the IDD on MER.

The control of the manipulator using the estimated camera models will be examined in Sect. 2.2. Finally, a review of results from both simulations and actual experiments on multiple platforms will examine the efficacy of the technique in Sects. 4 and 5.

## 2 Parameter estimation for HIPS

The HIPS manipulation technique generates camera models through visual sensing of the manipulator's end-effector then uses these models to position the end-effector at a desired image-plane target. In this way HIPS combines the two independent steps of the standard flight approach by removing the intermediate transformations from the camera frame to vehicle frame and vehicle frame to manipulator frame.

The HIPS manipulation technique has been implemented utilizing both the eighteen-parameter CAHVOR model (a pin-hole camera model with radial distortion (Gennery 2001)) and the twenty-one parameter generalization called CAHVORE, which incorporates the ability to model both perspective projection optics (as in CAHVOR), fish-eye optics, or intermediate geometries (Gennery 2006). The particular camera model used is not crucial to the technique, but these have a long history of use at the Jet Propulsion Laboratory (JPL), and have well-understood performance characteristics and calibration and modeling tools associated with them. Without loss of generality CAHVORE models will be described below since, with particular model settings, a CAHVORE model will degenerate into a CAHVOR model.

### 2.1 The preplanned trajectory

For each camera the CAHVORE parameters are initially estimated using a predetermined set of typically forty to fifty manipulator poses. This preplanned trajectory constitutes a broad sample of both the image-space of the vision sensor and the joint-space of the manipulator and is completely unrelated to a specific task. At each pose the stereoscopic camera pair acquires a set of images of the end-effector and the realized joint angles are recorded. The image data are processed to extract the location of some reliable point on the manipulator, or *fiducial*. Note that while in the laboratory a simple fiducial such as shown in Fig. 3 is used, any portion of the manipulator that can be reliably extracted from image data in the harsh lighting of planetary operations could be used. For example, in Sect. 5.3, the Mössbauer contact plate is utilized as a "fiducial". The 3D location of the fiducial is computed using the nominal forward kinematics of the

robot, and the set of 5D vectors (3D position and 2D image plane location) are used to compute a new set of CAHVORE parameters. See (Gennery 2001) or (Gennery 2006) for details on the least-squares minimization utilized for this step. Briefly, the following equation is minimized:

$$\begin{aligned}
 J(\text{CAHVORE}) &= \sum_{i=1}^n W_i [\{u^i - f_x(\mathbf{P}(\Theta^i), \text{CAHVORE})\}^2 \\
 &\quad + \{v^i - f_y(\mathbf{P}(\Theta^i), \text{CAHVORE})\}^2]. \quad (1)
 \end{aligned}$$

In this equation,  $n$  is the number of poses in the preplanned trajectory. The point  $\mathbf{P}(\Theta)$  is the position of the fiducial in the workspace of the arm, when the arm is at joint angles  $\Theta$  or, more simply, the forward kinematics of the manipulator. Point  $\mathbf{P}(\Theta)$  projects to the image plane location  $(u^i, v^i)$ . The 3D to 2D mapping function  $f()$  is determined by the relevant camera model. Finally, the weighting factor  $W_i$  is typically set to unity, but could be used to bias the model fit, for example, towards the expected work volume of a particular instrument.

In the previous section, it was noted that this estimation process is done without regard to any "real" physical reference frame. This statement is true because the three dimensional coordinates of the fiducial marker are obtained from the nominal forward kinematics of the manipulator arm under the assumptions that the manipulator arm links are rigid (or with a fixed gravity sag model), that the arm lengths are constant, and that there is no joint backlash.

It is understood from the outset that these nominal three-dimensional coordinates are likely to contain, possibly significant, errors. A least-squares minimization of (1) may produce CAHVORE parameters that are significantly different from the parameters that would be generated by an accurate calibration fixture. However, the use of the two CAHVORE parameter sets, referred to as the "Static HIPS Models", to model the relationship between self-reported joint angles and image plane appearance leads to a robust, accurate, and reasonably generalizable (for interpolation) model.

The preplanned trajectory, image acquisition, and parameter fitting is conducted only after dramatic events that might alter either the camera and/or manipulator calibrations during the life of the mission such as high impact vibrations experienced by the vehicle. To expedite the procedure or to allow manual verification of fiducial detection, the data may be processed on ground-based computers and the resulting camera models transmitted to the remote manipulator system.

### 2.2 Manipulation with static HIPS models

Once the Static HIPS models of each of the stereo cameras have been determined, a target in the image-plane is selected

and the models are utilized to determine the target range via stereo correlation and triangulation.

One convenient method for baseline tests is to make use of a fiducial marker, or “cue”, for target acquisition and correlation. For exercises such as soil sample acquisition, this process can be accomplished using natural feature correlation or manual selection by scientists.

Given this target correlated in the image-plane of each camera, stereo triangulation is utilized to determine the target range. The process of triangulation with CAHVORE models is described in (Gennery 2006)—the process for Static HIPS remains unchanged, only using a modified parameter set.

With the computed range, the inverse kinematics are used to solve for the joint rotations of the manipulator that place the end-effector at the desired target location. Again, the computed 3-D target location may be significantly different from the actual 3-D physical location. However, the 3-D target location is accurate with respect to the manipulator kinematic model. Therefore the solution of the inverse kinematic model will position the end-effector at the desired image-plane targets.

### 2.3 On-line HIPS camera models

The aforementioned approach to manipulator control addresses the systematic errors that are present in the standard flight approach. These include separate camera calibrations, transformations between various frames, and manipulator kinematic errors such as link lengths and joint offsets. However, stochastic errors that occur due to finite image-plane cue detection, camera modeling errors, and inaccurate knowledge of joint angles arising from sensor resolution, orientation-dependent droop, joint backlash, etc., are not necessarily accounted for so far.

A solution to this problem is to divide the trajectory to the target, or “transition” trajectory, into a series of intermediate steps. At each intermediate goal the camera models are updated by use of data from newly acquired images of the end-effector, if visual access of the fiducial marker is available, and of the corresponding nominal coordinates in the manipulator reference frame. This is accomplished by a modification of (1) according to

$$\begin{aligned}
 J_2(\text{CAHVORE}) &= \sum_{i=1}^n W_i [\{u^i - f_x(\mathbf{P}(\Theta^i), \text{CAHVORE})\}^2 \\
 &\quad + \{v^i - f_y(\mathbf{P}(\Theta^i), \text{CAHVORE})\}^2] \\
 &\quad + \sum_{j=1}^m W_j [\{u^j - f_x(\mathbf{P}(\Theta^j), \text{CAHVORE})\}^2 \\
 &\quad + \{v^j - f_y(\mathbf{P}(\Theta^j), \text{CAHVORE})\}^2] \quad (2)
 \end{aligned}$$

where  $n$  and  $m$  are the number of poses currently available in the preplanned and transition trajectories respectively. Typically the weight applied to the most recent sample in the transition trajectory,  $W_j$  for  $j = m$ , is increased as the end-effector approaches the target. The effect is to improve the local joint-space to image-space mapping.

Because this approach utilizes continually updated camera models, it is referred to as “On-Line HIPS manipulation”. The target range is re-computed as before and then used to refine the necessary manipulator joint angles to position the end-effector at the target. As the end-effector/target distance decreases, the computed coordinates of the end-effector and target become more nearly affected by the same errors, so that the differences between their coordinates become increasingly small. Therefore precise location of the manipulator’s end-effector is achieved.

### 3 Alternative target generation

For many reasons it is either impractical or undesirable to use a cue for each target position. For example, Sect. 5.1 describes a bolt-fastening task for which the target bolt could not be a cue. If stereoscopic data are used to manually select targets, or natural feature correlation is used to determine targets, the surface normal can be used, but many tasks explicitly specify autonomous operation (Huntsberger et al. 2002) or scientists are involved in selecting the science target (Powell et al. 2005).

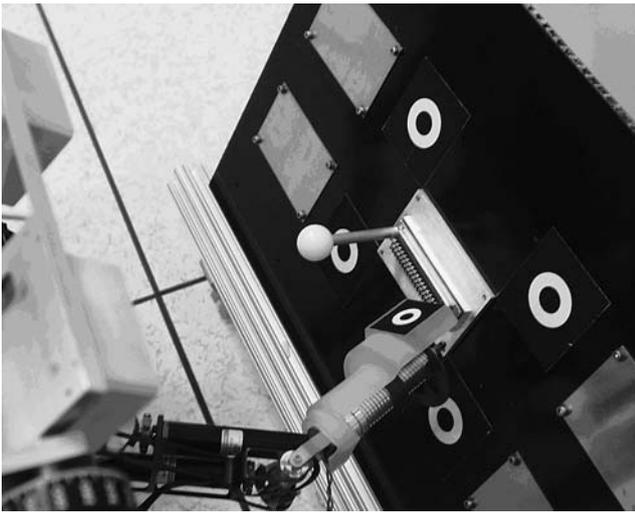
If cues are to be used, a single target cue provides only enough information to solve the positioning problem, as opposed to both position and orientation of the end-effector. An efficient and viable solution to these problems is to create a local coordinate frame using three or more cues, all of which are non-collinear but not necessarily co-planar as displayed in Fig. 4. The cues are placed precisely relative to the origin of the target frame; however, the position and orientation of the target frame relative to the end-effector are unknown a priori.

The location,  $(x_i, y_i, z_i)$ , of each cue is determined using the manipulator-generated camera models. In the case of  $n$  co-planar cues the local tangent plane to the surface,  $\psi$ , is defined as

$$\psi(x, y, z, \mathbf{A}) = a_1x + a_2y + a_3z - 1 = 0, \quad (3)$$

where the parameters  $\mathbf{A} = [a_1, a_2, a_3]^T$  are determined by performing a least-squares minimization of

$$J_3(\mathbf{A}) = \sum_{i=1}^n [\psi(x_i, y_i, z_i; \mathbf{A})^2]. \quad (4)$$



**Fig. 4** Close-up Image of LEMUR inserting the rotary tool into the bolt head. The target board creates a local coordinate system for efficient target generation

Having determined the parameters of the surface model,  $\psi$ , the unit normal to the surface is computed according to:

$$\hat{\mathbf{e}}_n = \frac{\nabla \psi}{\|\nabla \psi\|}. \quad (5)$$

The target is then specified with respect to the local coordinate system according to:

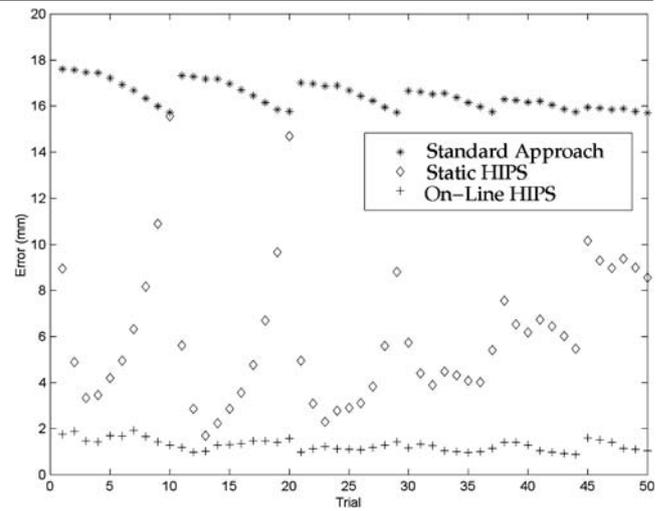
$$[x_t, y_t, z_t] = a_x \hat{\mathbf{e}}_t + a_y \hat{\mathbf{e}}_s + a_z \hat{\mathbf{e}}_n \quad (6)$$

where  $\hat{\mathbf{e}}_t$  and  $\hat{\mathbf{e}}_s$  are the in-plane unit-vectors computed from the cue locations and  $[a_x, a_y, a_z]^T$  are the measured offsets with respect to the local target frame.

Finally, these 3-D targets are mapped back into the image-plane using relevant camera models and the end-effector positioning proceeds. As the end-effector is moved toward the target and the camera models are updated, new image-plane targets are computed. The target range and local surface normal are re-computed and then used to determine the joint angles to position and orient the end-effector at the target.

#### 4 Simulation results

Testing of the HIPS manipulation technique included both a simulation study and a hardware implementation. The simulation study involved a series of positioning exercises of a four degree-of-freedom manipulator (outstretched length of approximately 0.75 m) relative to a target board consisting of three simulated fiducial markers as described in Sect. 3.



**Fig. 5** Results of Simulated Positioning Tests using the Standard Flight Approach (Inverse Kinematic Model), Static HIPS camera models and On-line HIPS models. The average error was over 15 mm for the standard approach. An improvement in the average error of greater than 60% to 6.4 mm was achieved using the Static HIPS camera model. With On-line updates to the model the average error dropped to less than 1.3 mm, an order of magnitude reduction in the positioning error from the standard flight approach

For these positioning simulations a significant error was introduced into the nominal arm kinematics (a combined total of over 20 mm change in the link lengths) without modifying the nominal kinematic model of the manipulator. This “truth” kinematic model is not available to the positioning system and will be called the perturbed kinematic model. Flight hardware is typically calibrated to at least a millimeter level of precision, so this error represents a more challenging scenario. This simulation also verifies that any improvements generalize to a larger than “local” area of the parameter space. A nominal camera model was used to map target locations from three-dimensional space into each image plane of a simulated stereo camera pair.

For the study, three sets of positioning exercises were performed. In each exercise, the target board was “placed” at a workspace position 40 to 70 cm (6 steps) in front of the rover,  $-20$  to  $20$  cm (5 steps) to either side of the manipulator base, and 15 or 30 cm below the manipulator base, resulting in 60 nominal locations. Some extremal positions were not reachable by the modeled manipulator and are thus not simulated, leading to 6 runs of 6–10 positions, and a total of 50 positions, as shown in Fig. 5. The vertical dimension was indexed first, followed by the horizontal, followed by the range.

In the first set of tests the standard flight approach was used to position the manipulator as described in Sect. 1.2. First, the range to the target was determined using the nominal camera model. The nominal inverse kinematic model then was solved to determine the joint angles that should locate the end-effector at the target. Finally, the positioning

accuracy was measured by substituting the final joint angles into the perturbed kinematic model and comparing the “true” terminal position to the target location. As displayed in Fig. 5, the terminal root sum of squares (RSS) error using the standard approach had a mean of 16.5 mm with a  $3\sigma$  error bound of  $\pm 1.72$  mm. This means, for this example, that 99.7% of the errors would fall between 14.78 mm and 18.22 mm (if the errors were Gaussian). The terminal error depends on range, as shown in Fig. 5 (for example, the first ten samples correspond to targets at the closest range—the last six samples correspond to targets at the farthest range) for the standard approach.

In the final two series of tests, HIPS was used to control the manipulator. In these tests a fiducial marker was “mounted” on the end-effector and used to execute a pre-planned trajectory to establish new camera models as described in Sect. 2. Briefly, the manipulator was moved through a set of forty predefined joint poses. At each pose the fiducial marker was mapped into each camera plane using the nominal camera model. In addition, errors in the detection of the fiducial marker were simulated by including zero-mean, Gaussian noise (0.5 pixel standard deviation) to the image plane coordinates. A least-squares fit of the image plane and three-dimensional coordinates then produced Static HIPS camera models.

In the second series of tests the manipulator was positioned using these Static HIPS camera models. The range to the target was computed from the HIPS models and the nominal inverse kinematic model was solved to provide the terminal joint position. Again, the positioning accuracy was measured by substituting the final joint angles into the perturbed kinematic model and comparing the “true” terminal position to the target location. As shown in Fig. 5 the mean terminal error using this approach was 6.39 mm with a  $3\sigma$  error bound of  $\pm 9.3$  mm resulting in a 60% improvement from the standard flight approach. The terminal error for the Static HIPS case does depend on range, but has minima in the middle of the workspace sampled in the preplan step described in Sect. 2.1.

In the final series of tests the Static HIPS camera models were updated using samples during the approach to the target, as described in Sect. 2.3. A total of twenty samples were acquired during the transition trajectory and given a linearly increasing weight. In doing so the camera models are expected to be locally accurate. As before, zero-mean, Gaussian noise (0.5 pixel standard deviation) was applied to the image plane coordinates of the additional samples. As shown in Fig. 5 the mean terminal error using this approach was 1.30 mm with a  $3\sigma$  error bound of  $\pm 0.80$  mm. This represents an order of magnitude increase in accuracy from the standard flight approach. In addition the dependence (seen in the previous two approaches) of the terminal error on range disappears. A comparison of the CAHVOR camera model

**Table 1** CAHVOR left camera model parameters

Description	Model parameters
Truth model (units)	
C (cm)	(9.147, -5.154, 13.049)
A (unit vector)	(0.932, 0.050, 0.359)
H (pixels)	(265.423, 355.507, 110.901)
V (pixels)	(36.251, 2.523, 367.422)
O (unit vector)	(0.920, 0.048, 0.388)
R (no dim.)	(0.000, -0.275, 0.0482)
Static HIPS Model	
C	(8.094, -4.763, 15.015)
A	(0.918, 0.073, 0.390)
H	(263.382, 365.738, 119.973)
V	(45.606, 6.331, 376.911)
O	(0.922, 0.070, 0.382)
R	(0.000, -0.278, 0.038)
On-Line HIPS model (Representative)	
C	(7.757, -4.716, 14.830)
A	(0.909, 0.079, 0.409)
H	(261.297, 370.249, 126.900)
V	(40.250, 6.851, 386.620)
O	(0.914, 0.063, 0.401)
R	(0.000, -0.286, 0.053)

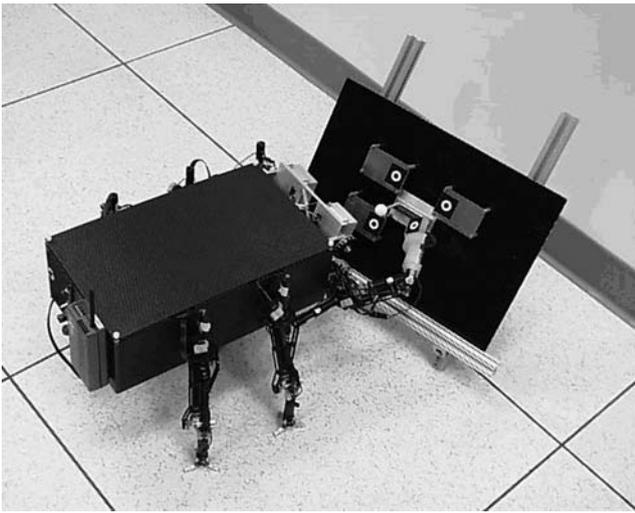
parameters for the nominal left camera model, the Static HIPS camera model for the left camera, and a representative updated left camera model from one of the trials is provided in Table 1.

## 5 Experimental results

HIPS has been used on many platforms at the Jet Propulsion Laboratory since 1998, including the Limbed Excursion Mobile Utility Robot (LEMUR) (Kennedy et al. 2001), LEMUR-II (Nickels et al. 2006; Kennedy et al. 2002), SRR (Sample Return Rover) (Baumgartner et al. 1998), and SRR2K (Stroupe et al. 2005), as well as on the MER vehicles. In this paper, results from the use of HIPS on three of these platforms will be shown: a small six-limbed robot (LEMUR) and a small wheeled rover with a flexible robotic arm (SRR). Finally, the use of HIPS on the MER mission will be discussed.

### 5.1 LEMUR implementation and results

HIPS was initially implemented and evaluated on the LEMUR platform as shown in Fig. 6. LEMUR is a small (approximately 5 kg), agile robot developed at the Jet

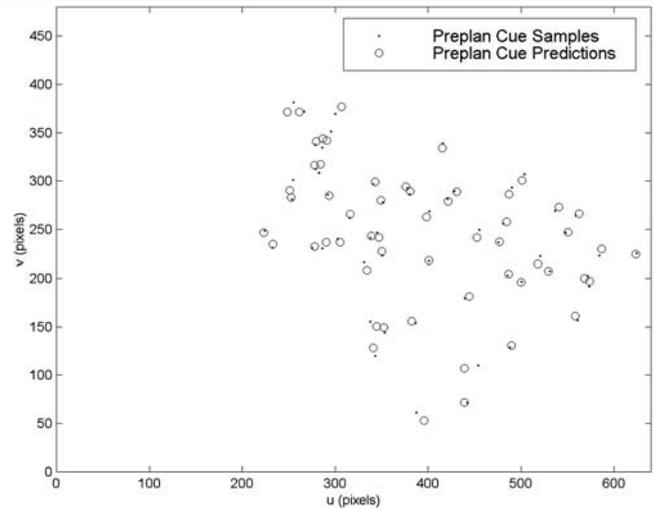


**Fig. 6** The Limbed Excursion Mobile Utility Robot (LEMUR) is designed to perform dextrous small-scale assembly, inspection and maintenance of space facilities such as the displayed mockup

Propulsion Laboratory to perform dextrous small-scale assembly, inspection and maintenance of macro space facilities (Kennedy et al. 2001). The LEMUR platform is equipped with six, three degree-of-freedom limbs for mobility and a forward-mounted stereo camera pair for mobility and manipulation purposes. In addition the front two limbs incorporate “quick connect” end-effector capability to facilitate the swapping of tools such as a rotary driver to perform manipulation operations. With the rotary tool the outstretched length of the manipulator limbs is approximately 20 cm. An important consideration for manipulation purposes is the fact that the joints contain a significant amount of backlash, approximately 3 degrees per joint or 20 mm of error at the tool tip, for which the manipulation approach must account.

The task chosen for the LEMUR implementation was the removal of a worn or damaged connector, as demonstrated in Fig. 6. In particular the goal was the insertion of the tip of the rotary tool into the head of a bolt. In the task the bolt was mounted on the target board relative to the three fiducial markers. Although this setup requires precise knowledge of the bolt location relative to the fiducials, the position and orientation of the target board is unknown a priori as described in Sect. 3.

The first step of the implementation was the initial camera parameter estimation as detailed in Sect. 2. The eighteen parameter CAHVOR models were used in this implementation. This involved moving the manipulator through a series of sixty predefined poses and sampling both the joint positions and the image-plane appearance of the fiducial marker mounted on the end-effector. The initial camera model then was a least-squares fit of the subsequent data as displayed for the left camera in Fig. 7. Typically, the root-mean square



**Fig. 7** Initial Camera Model Fit for Left Camera for LEMUR. Positioning precision based on the initial model alone would be on the order of 6 mm

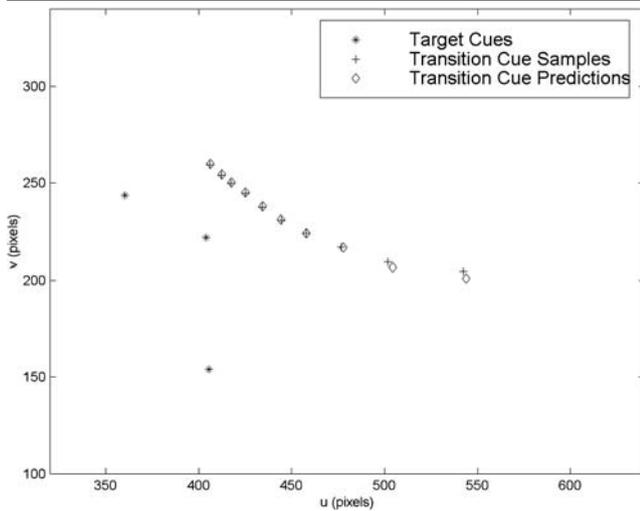
residual was approximately 6 pixels. Since the image resolution is roughly 1 mm per pixel the resulting positioning precision based on the initial camera model alone (Static HIPS) would be on the order of 6 mm. This level of precision is insufficient to achieve the desired task goal. It should be noted that a significant portion of this error is due to the orientation dependent backlash in the joints.

During the approach to the target, ten additional samples were acquired to refine the camera models and, thus account for stochastic errors, a major source of which is orientation dependent backlash. This process was detailed in Sect. 2.3. The application of an increasing weight to the most recent samples accounts for such errors. Figures 8 and 9 display the results of a typical test. As shown the initial camera-space residual in the left camera is approximately four pixels. With the application of additional samples of increasing weight the residual drops to a final value of less than one-half of a pixel.

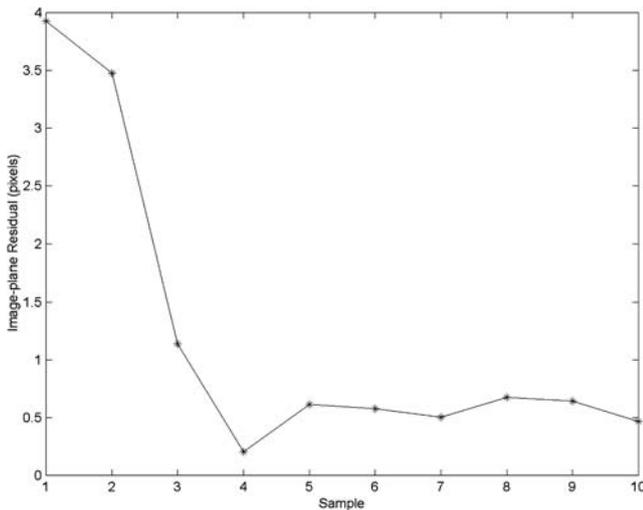
The test scenario involved placing the target board approximately 10–15 cm in front of LEMUR at various orientations. Due to compliance in the arm successful completion of the task required positioning precision within 2 mm of the center of the bolt. In a battery of over 100 tests, successful insertion of the rotary tool was achieved in approximately 90% of the attempts. A close-up image of the task is shown in Fig. 4.

## 5.2 SRR implementation and results

The HIPS implementation on the Sample Return Rover (SRR) is interesting primarily due to the increased length and flexibility of the robotic arm used on the rovers, as shown in Fig. 10.

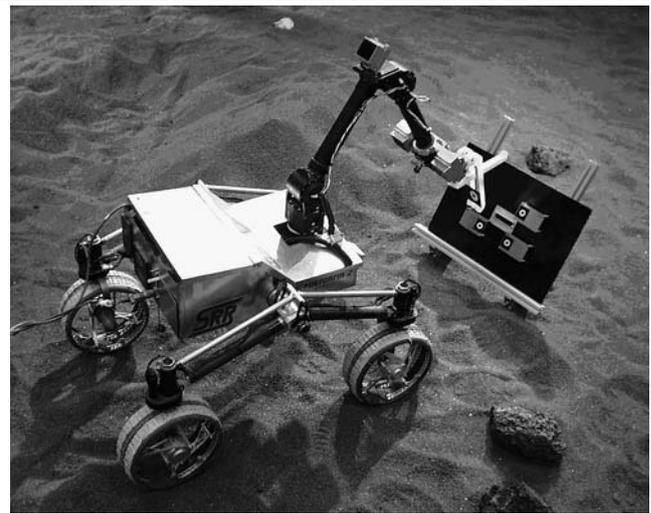


**Fig. 8** Acquisition of 20 Additional Samples During Approach to Target for LEMUR. As the fiducial moves from its initial location on the right, the predicted location of the fiducial (◊) converges with the sampled location of the fiducial (+) as given by the updated camera model

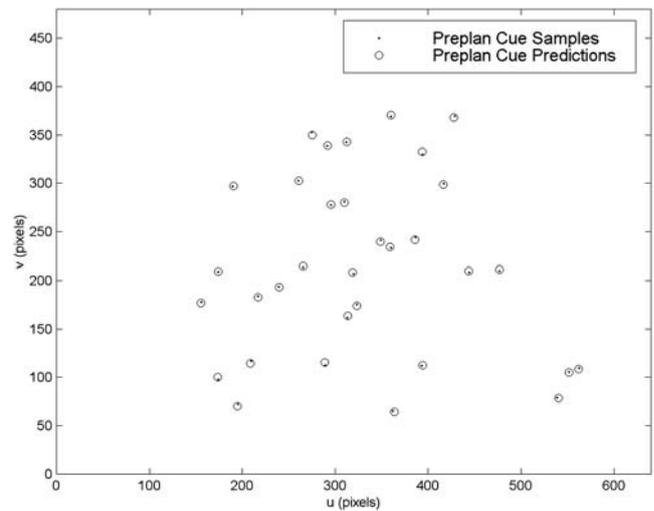


**Fig. 9** Residual as a Function of Additional Samples during Target Approach for LEMUR. Note the decrease in the residual, the difference between the actual camera-space location of the fiducial marker with that predicted by the updated camera model, with additional samples and an increased weighting factor as the terminus is approached. The residual decreases from an initial value of approximately 4.0 pixels to a final value of less than 0.5 pixels

The SRR vehicle is a four-wheeled rover platform designed for research and technology development and demonstration (Baumgartner et al. 1998). The vehicle has a four degree-of-freedom manipulator with a gripper end-effector. When fully extended the nominal length of each manipulator is approximately 0.75 m and is located 0.4 m above the ground. In addition each vehicle is equipped with hazard avoidance cameras for manipulation and mobility.



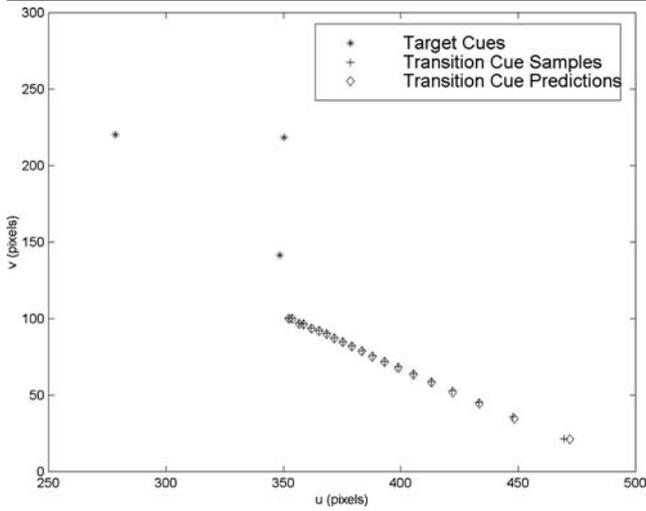
**Fig. 10** The Sample Return Rover (SRR) uses HIPS to locate the end-effector at the target board



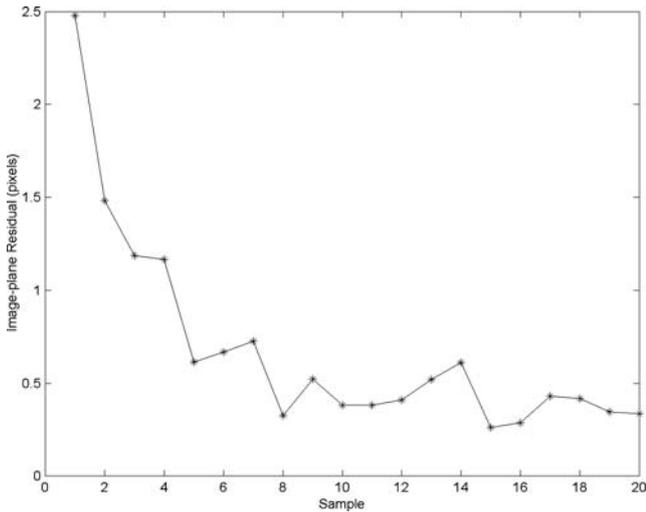
**Fig. 11** Initial Camera Model Fit for Left Camera for SRR. Positioning precision based on the initial model alone would be on the order of 2 cm

For these cameras, the eighteen parameter CAHVOR model was used.

To test HIPS on SRR, a fiducial marker was mounted on the gripper. The positioning tests then consisted of locating the tip of the gripper with respect to a target board with three fiducial markers similar to the one in Fig. 10. In each test the target board was placed approximately 0.4 to 0.7 m from the base of the manipulator at a variety of azimuthal positions with an a priori unknown position and orientation. The tip of the gripper then was commanded to locate itself just above one of the fiducial markers. To record the terminal precision of each test the realized location of the tip was marked on the fiducial marker using a pen.

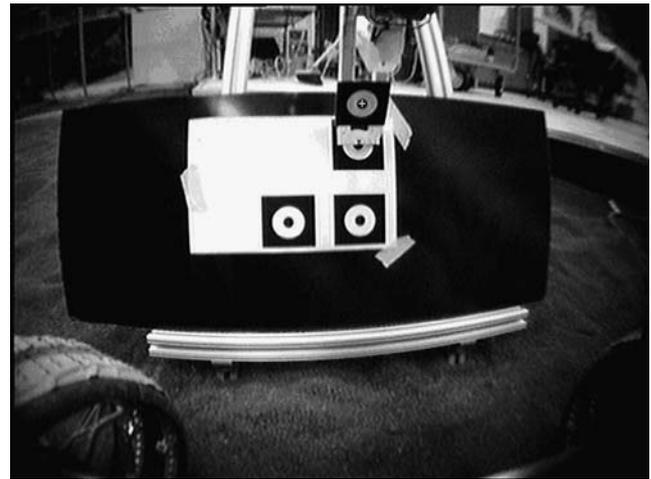


**Fig. 12** Acquisition of 20 Additional Samples During Approach to Target for SRR. As the fiducial moves from its initial location on the right, the predicted location of the fiducial (◊) converges with the sampled location of the fiducial (+) as given by the updated camera model

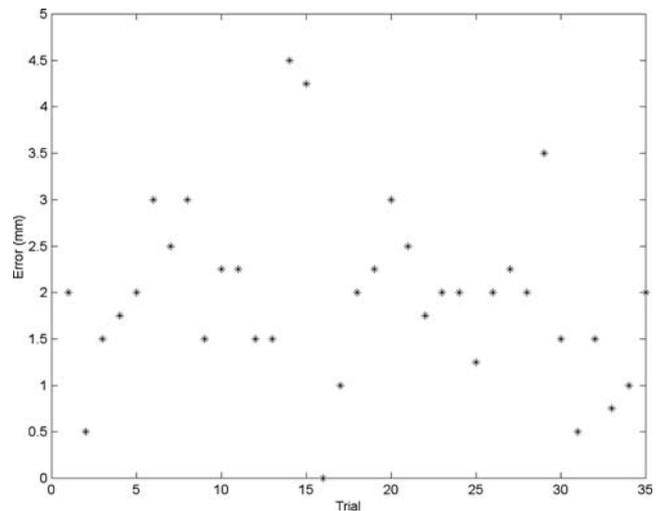


**Fig. 13** Residual as a Function of Additional Samples during Target Approach for SRR. Note the decrease in the residual, the difference between the actual camera-space location of the fiducial marker with that predicted by the updated camera model, with additional samples and an increased weighting factor as the terminus is approached. The residual decreases from an initial value of approximately 2.5 pixels to a final value of less than 0.3 pixels

The initial camera models were determined from thirty-predefined poses as shown in Fig. 11. Using these data the root-mean-square residual of the resulting fit was 1.58 pixels. The better model fit, as compared to the tests on LEMUR, was due primarily to the lack of backlash in the joints since SRR has harmonic gears at each joint. As reported in (Baumgartner et al. 1998), positioning using the flight standard technique would result in 1.5–2 cm of error.



**Fig. 14** Actual 640 × 480 Greyscale Image of Terminal Position from Left Stereo Camera. Note that the radius of the black fiducial marker is 50 mm. In this test the straight black mark on the gripping end-effector is aligned within 1 mm of the center of the target fiducial



**Fig. 15** Results of 35 Positioning Tests. The average error was less than 2.0 mm

The use of these Static HIPS models reduced this to a terminal error of approximately 7.8 mm.

During each positioning exercise twenty additional samples were acquired to further refine the camera models locally. The results of a typical test are displayed in Fig. 12. In this test the initial image-plane residual in the left camera is 2.48 pixels. The residual is the difference between where the kinematics and current CAHVOR models predict the fiducial to appear and where it is eventually found. With the application of additional samples the residual drops to approximately one-third of a pixel as evidenced in Fig. 13. The terminal pose as captured by the left stereo camera is shown in Fig. 14. Note that the goal is to align the black line on the gripper with the center of the top fiducial target. In this case the terminal error is less than 1 mm.

Following this procedure, a total of 35 positioning tests were conducted on SRR to baseline the positioning accuracy as recorded in Fig. 15. The mean error in these tests was 1.96 mm with a  $3\sigma$  error bound of  $\pm 2.9$  mm.

### 5.3 Use in flight—MER implementation and results

In addition to the hardware platforms mentioned above, the HIPS approach has been used to mitigate some positioning degradation of the five degree-of-freedom robot arm known as the Instrument Deployment Device, or IDD, on the Mars Exploration Rover (MER) vehicles. Recall that the MER mission requirements include an IDD placement requirement of 10 mm in position and 10 degrees in orientation (Baumgartner et al. 2005).

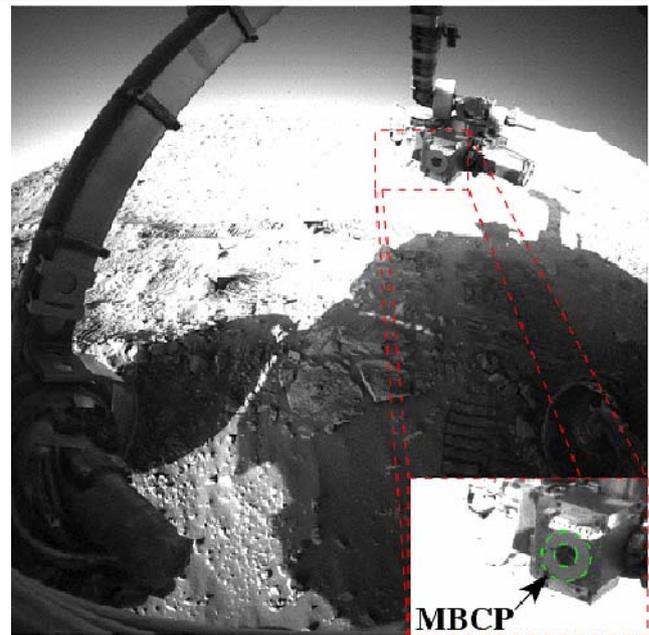
When a science target is to be studied utilizing the IDD, stereo images of the target are acquired and relayed to Earth. The specific target is selected using these images, as interpreted in stereo using CAHVORE models (Gennery 2006). Inverse kinematics can be used to compute the desired joint angles to position the relevant tool of the IDD at the desired location, or the 3D position can be commanded, with the rover performing the inverse kinematics using the on-board processor.

After some time on the surface of Mars, it was noted that the IDD positioning performance on the Spirit rover was not performing as well as the 10 mm specification. The Static HIPS approach described above was used to create updated CAHVORE models for the front Hazard Avoidance Cameras (hazcams) on Spirit.

(1) *Updating the camera models:* A set of poses were commanded, attempting to span the work volume of the IDD, the relevant joint space work volume, and the image plane of the front hazcams. This is exactly the “preplan trajectory” described in Sect. 2.1. A typical image is shown in Fig. 16, with the inset clearly showing the Mössbauer contact plate, or MBCP, that was used as a fiducial as described above.

Comparing the 3D positions of the MBCP as reported by the forward kinematics of the rover (a fairly accurate model that includes arm sag due to Martian gravity, among other nonstandard features) to the 3D position of the MBCP as measured by stereo triangulation based on the pre-flight CAHVORE models yields a mean discrepancy of 16.42 mm with a  $3\sigma$  error bound of  $\pm 20.01$  mm for the preplan trajectory data.

Based on the observed data (3D position from forward kinematics and 2D image-plane locations), new Static HIPS camera models were derived as explained in Sect. 2, with the following exceptions. The pre-flight CAHVORE camera models were used to begin the minimization, and only the camera extrinsic parameters were updated from pre-flight values. Utilizing the new camera models, the mean discrepancy between the 3D positions of the MBCP reported by



**Fig. 16** An example pose from the Spirit MER vehicle with the MBCP indicated

forward kinematics and stereo triangulation (for the same preplan trajectory data) was reduced to 7.69 mm with a  $3\sigma$  error bound of  $\pm 16.28$  mm.

Based on the success with Spirit, a set of poses was commanded on the Opportunity rover, bringing the residual mean from 11.74 mm with a  $3\sigma$  error bound of  $\pm 9.43$  mm for the pre-flight camera models to a mean of 3.34 mm with a  $3\sigma$  error bound of  $\pm 10.95$  mm for the updated models, again for the same set.

(2) *Testing the new models:* A representative sampling of hazcam images from Mars (not those used to generate the models above) that show the MBCP were then analyzed, to test the ability of the Static HIPS models to generalize to other locations in the workspace. As shown in Table 2, they demonstrate an improvement from a 10.10 mm mean residual with a  $\pm 8.86$  mm  $3\sigma$  error bound with the pre-flight camera models to a 6.11 mm mean residual with a  $\pm 7.20$  mm  $3\sigma$  error bound with the HIPS camera models for 215 image pairs from Spirit and an improvement from a 8.83 mm mean residual with a  $\pm 6.62$  mm  $3\sigma$  error bound with the pre-flight camera models to a 3.82 mm mean residual with a  $\pm 7.88$  mm  $3\sigma$  error bound with the HIPS camera models for 146 image pairs from Opportunity.

The HIPS camera models are used in the rover operations planning sequence, relatively seamlessly, as follows. Science target selection is done on Earth, based on the previous day’s hazcam imagery. The HIPS models are used to interpret these hazcam images and to compute 3D IDD targets. While these 3D positions may or may not reflect the actual position of the MBCP on Mars, with respect to the

**Table 2** Summary of experimental results.

Platform	Method	Mean (RSS) Residual	3 $\sigma$ error bound
Simulation	Flight-Standard	16.5 mm	1.72 mm
Study	Static HIPS	6.39 mm	9.30 mm
	On-Line HIPS	1.30 mm	0.80 mm
LEMUR	Flight-Standard	N/A	N/A
	Static HIPS	6 mm	N/A
	On-Line HIPS	$\leq 2$ mm	(90% success)
SRR	Flight-Standard	15–20 mm [27]	N/A
	Static HIPS	7.8 mm [27]	N/A
	On-Line HIPS	1.96 mm	2.9 mm
MER	Pre-Launch	<10 mm	N/A
Spirit	Flight-Standard	10.10 mm	8.86 mm
	Static HIPS	6.11 mm	7.20 mm
MER	Pre-Launch	<10 mm	N/A
Opportunity	Flight-Standard	8.83 mm	6.62 mm
	Static HIPS	3.82 mm	7.88 mm

Some data are not available (N/A)

rover, they do cause the IDD to move to the desired location, to within 10 millimeter. The use of the HIPS-generated camera models has therefore resulted in hundreds of Martian days of improved science return.

## 6 Conclusion

Although the use of fiducial markers for targets is possible in many applications such as the construction and maintenance of satellites or remote stations, sample cache container retrieval, or rendezvous and docking, there are many applications for which the use of fiducials is either impracticable or impossible. One such application is soil or rock sample acquisition. In such cases target acquisition is made using manual feature selection. This target selection process fits in well with science operations as currently practiced in flight projects.

This paper described a novel approach to the control of manipulators using estimation, stereo correlation and triangulation for space exploration applications. The results of simulations and implementations on three distinct hardware platforms, a walking robot and two rovers, demonstrated the ability of the method to achieve average positioning precision of a manipulator with respect to a target board on the order of one to two millimeters. Such precision represents an order of magnitude improvement over the standard calibrated stereo approach. A variant of this approach has been

used to improve the IDD positioning accuracy of the MER vehicles to less than 10 millimeters.

HIPS meshes well with the constraints of robotic space applications: it does not require additional onboard image processing, is safe and leads to easily verifiable workspace motion, and improves instrument positioning accuracy by a factor of two over the current flight-standard approach.

**Acknowledgements** The work described in this publication 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

- Allen, P. K., Timcenko, A., Yoshimi, B., & Michelman, P. (1992). Trajectory filtering and prediction for automated tracking and grasping of a moving object. In *Proceedings of the 1992 IEEE international conference on robotics and automation* (Vol. 2, pp. 1850–1856), May 1992.
- Baumgartner, E. T., Leger, P. C., Schenker, P. S., & Huntsberger, T. L. (1998). Sensor-fused navigation and manipulation from a planetary rover. In *Proceedings of the SPIE symposium on sensor fusion and decentralized control in robotics systems*, Boston, November 1998.
- Baumgartner, E. T., et al. (2005). The Mars Exploration Rover instrument positioning system. In *Proceedings of the IEEE aerospace conference*, Big Sky, MT, March 2005.
- Bonitz, R. G. (1997). Mars surveyor '98 lander MVACS robotic arm control system design concepts. In *Proceedings of the 1997 IEEE international conference on robotics and automation* (pp. 2465–2470), Albuquerque, NM.
- Chen, W. Z., Korde, U., & Skaar, S. B. (1994). Position-control experiments using vision. *International Journal of Robotics Research*, 13(3), 199–208.
- Erickson, J., et al. (2002). Mars Exploration Rover surface operations. In *Proceedings of the 53rd international astronomical congress, the world space congress*, Houston, TX, October 2002. paper No. IAC-02-Q.3.1.03.
- Feddema, J. T., Lee, C. S. G., & Mitchell, O. (1992). Model-based visual feedback control for a hand-eye coordinated robotic system. *IEEE Computer*, 25(8), 21–31.
- Feddema, J. T., Lee, C. S. G., & Mitchell, O. (1993). Visual servoing: Real-time control of robot manipulators based on visual sensory feedback. In K. Hashimoto (Ed.), *Visual servoing* (Vol. 7, pp. 105–138). Singapore: World Scientific.
- Gennery, D. (2001). Least-squares camera calibration including lens distortion and automatic edition of calibration points. In A. Grun & T. Huang (Eds.) *Calibration and orientation of cameras in computer vision* (pp. 123–136). Berlin: Springer.
- Gennery, D. B. (2006). Generalized camera calibration including fish-eye lenses. *International Journal of Computer Vision*, 68(3), 239–266.
- Hager, G. D., Kriegman, D. J., & Morse, A. S. (1998). The block island workshop: A summary. In D. J. Kriegman, G. D. Hager, & A. S. Morse (Eds.) *The confluence of vision and control* (pp. 273–281). London: Springer.
- Huntsberger, T., et al. (2002). Rover autonomy for long range navigation and science data acquisition on planetary surfaces. In *Proceedings of the 1992 IEEE international conference on robotics and automation* (Vol. 32, pp. 3161–3168).

- Hutchinson, S., Hager, G., & Corke, P. (1996). A tutorial on visual servo control. *IEEE Transactions on Robotics and Automation*, 12(5), 651–670.
- Kennedy, B., et al. (2001). LEMUR: limbed excursion mechanical utility rover. *Autonomous Robots*, 11, 201–205.
- Kennedy, B., et al. (2002). Limbed excursion mechanical utility rover: LEMUR II. In *Proceedings of the 53rd annual international astronomical congress*, Houston, TX.
- Maki, J. N., et al. (2003). Mars Exploration Rover engineering cameras 12. *Journal of Geophysical Research*, 108(E12), 8071.
- Matijevic, J. R. (1998). The Pathfinder mission to Mars: Autonomous navigation and the Sojourner microrover. *Science*, 5362, 454–455.
- Matijevic, J. R., Goldstein, B. G., & Welch, R. V. (2001). The Mars Exploration Rover: An in situ science mission to Mars. In *Proceedings of the 31st international conference on environmental systems*, Orlando, FL, July 2001. paper No. 2001-01-2136.
- Nickels, K., et al. (2006). Vision-guided self-alignment and manipulation in a walking robot. In *Proceedings of the 2006 IEEE international conference on system of systems engineering*, Los Angeles, CA, USA, April 2006.
- Powell, M. W., et al. (2005). Scientific visualization for the Mars Exploration Rovers. In *Proceedings of the 2005 IEEE international conference on robotics and automation*.
- Ruf, A., & Horaud, R. (1999). Visual servoing of robot manipulator, Part I: Projective kinematics. *International Journal of Robotics Research*, 11, 1101–1118.
- Savage, D., & Cook-Anderson, G. (2004). *NASA selects investigations for the Mars Science Laboratory*. Jet Propulsion Laboratory, December 2004, Press Release 04-398.
- Skaar, S. B., Brockman, W. H., & Hanson, R. (1987). Camera space manipulation. *International Journal of Robotics Research*, 6(4), 20–32.
- Smith, P. (2004). The Phoenix mission to Mars. In *Proceedings of the IEEE aerospace conference* (Vol. 1, p. 342), March 2004.
- Squyres, S. et al. (2003). Athena Mars rover science investigation. *Journal of Geophysical Research*, 108(E12), 8062.
- Stroupe, A., et al. (2005). Behavior-based multi-robot collaboration for autonomous construction tasks. In *Proceedings of the IEEE/RSJ international conference on intelligent robots and systems*, Edmonton, AB, Canada, August 2005.
- Tunstel, E., et al. (2002). FIDO rover field trials as rehearsal for the NASA 2003 Mars Exploration Rovers Mission. In *Proceedings of 9th international symposium on robotics & applications, 5th world automation congress* (pp. 320–327), Orlando, FL, June 2002.
- Urmson, C., et al. (2001). A sensor arm for robotic Antarctic meteorite search. In *Proceedings of the 3rd international conference on field and service robotics*, Helsinki, Finland, July 2001.



**Matthew L. Robinson** received the B.S. in 1996 and the Ph.D. in 2001, both in mechanical engineering from the University of Notre Dame. He is currently a member of the Engineering Staff in the Mechanical and Robotic Technologies Group at the Jet Propulsion Laboratory in Pasadena, CA. He currently serves on the 2007 Phoenix Lander mission as the robotic arm flight software lead engineer and robotic arm systems engineer. His recent research ef-

forts have included the development of calibration-free, stereo-based manipulation techniques (HIPS) and control systems for holonomic and hybrid nonholonomic/holonomic systems. His research interests include machine vision, simulation, manipulation, mechatronic systems and controls.



**Eric T. Baumgartner** received the B.S. degree in aerospace engineering from the University of Notre Dame in 1988, the M.S. degree in aerospace engineering from the University of Cincinnati in 1990, and the Ph.D. degree in mechanical engineering from the University of Notre Dame in 1993. Currently, he is the Dean of the T. J. Smull College of Engineering at Ohio Northern University in Ada, OH. Prior to his current position, Dr. Baumgartner was a Senior Engineer at NASA's Jet Propulsion Laboratory in Pasadena, CA, and contributed to the 2003 Mars Exploration Rover mission as the lead test and operations engineer for the 5 degree-of-freedom robotic manipulator carried by the rover. His research interests include sensor fusion and state estimation techniques for autonomous robotic systems.



**Kevin M. Nickels** is an associate professor in the Department of Engineering Science at Trinity University. He received the B.S. degree in Computer and Electrical Engineering from Purdue University (1993), and received the M.S. degree (1996) and the Ph.D. (1998) in Electrical Engineering from The University of Illinois at Urbana-Champaign. He is currently working in the areas of computer vision, pattern recognition, and robotics. Dr. Nickels spent the 2005–2006 academic year as a Senior Engineer in Avionics Systems and Space Technology at NASA's Jet Propulsion Laboratory, concentrating on Vision Guided Manipulation.



**Todd E. Litwin** received the B.S. degree in applied physics from Harvey Mudd College in 1979. He began work at JPL in 1979, at first writing ground software to support navigating spacecraft. In the 1980s he moved into robotics research, concentrating on camera calibration, object tracking, and stereo ranging. In the late 1990s he started writing flight software for unmanned spacecraft while still keeping active in ground-based robotic software.