Auton Robot (2006) 20:113–123
DOI 10.1007/s10514-006-5943-4

Sustainable cooperative robotic technologies for human and robotic outpost infrastructure construction and maintenance

Ashley Stroupe · Avi Okon · Matthew Robinson · Terry Huntsberger · Hrand Aghazarian · Eric Baumgartner

Published online: 22 April 2006

Abstract Robotic Construction Crew (RCC) is a heterogeneous multi-robot system for autonomous acquisition, transport, and precision mating of components in construction tasks. RCC minimizes use of resources constrained by a space environment such as computation, power, communication, and sensing. A behavior-based architecture provides adaptability and robustness despite low computational requirements. RCC successfully performs several construction related tasks in an emulated outdoor environment despite high levels of uncertainty in motions and sensing. This paper provides quantitative results for formation keeping in component transport, precision instrument placement, and construction tasks.

Keyword Multi-robot systems · Cooperative transport · Space robotics · Robotic construction

A. Stroupe (✉) · A. Okon · M. Robinson · T. Huntsberger · H. Aghazarian · E. Baumgartner
Jet Propulsion Laboratory/California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA, 91109
e-mail: Ashley.Stroupe@jpl.nasa.gov

A. Okon
e-mail: Avi.Okon@jpl.nasa.gov

M. Robinson
e-mail: Matthew.Robinson@jpl.nasa.gov

T. Huntsberger
e-mail: Terry.Huntsberger@jpl.nasa.gov

H. Aghazarian
e-mail: Hrand.Aghazarian@jpl.nasa.gov

E. Baumgartner
e-mail: Eric.Baumgartner@jpl.nasa.gov

1. Introduction

The current NASA vision for space exploration calls for sustained human and robotic presence in the solar system, beginning with a return to the Moon by 2020 (NASA Office of Exploration Systems, 2004). A safe, sustainable human presence requires construction of infrastructure (habitat, power, atmosphere control, etc.) prior to human arrival, which can only be accomplished with robotic technologies. These same capabilities are required to build other types of self-sustaining robotic colonies for in-situ resource utilization, power production, mining, orbital facilities, and autonomous maintenance and repair facilities. Flight constraints limit these robotic technologies to be low power, mass and volume, and computing power; yet, space operations constraints due to communications delays and blackouts and uncertainty may require systems to be highly autonomous, reliable, and self-sustaining. Additionally, while habitat designs are not finalized, components are typically expected to be at least four times rover length, which will require cooperative transport and manipulation.

JPL's Planetary Robotics Lab (PRL) is currently developing robotic technologies with the capabilities necessary to support infrastructure construction under hard constraints on power, mass, communications, and computing. Research has focused on the issues of reliability in the presence of noisy sensing and multiple agents. Early work in cooperative transport was presented as part of Robotic Work Crew (Trebi-Ollennu et al., 2002). Current results for Robotic Construction Crew (RCC) include demonstrations of end-to-end component acquisition, transport, and precision deployment into a modular structure by a heterogeneous robotic team in irregular terrain. To minimize the need for computationally expensive processes such as high-level task allocation and planning, RCC uses a distributed multi-robot behavior-based

architecture with distributed sensing and explicit and implicit communication to coordinate agents in executing a series of high-level behaviors designed to accomplish the task.

2. Related work

Current state of the art in autonomous construction provides for simple mating of marked components in a laboratory setting with a flat floor. Carnegie Mellon University has demonstrated multiple component mating using three specialized robots (vision, coarse manipulation, fine manipulation) (Brookshire et al., 2004; Simmons et al., 2000). The vision robot aligns where it can see the necessary work space and tells the manipulator robots how to adjust for alignment. Others looked at conceptual design of robotic systems for construction (Howe, 2000). Work in cooperative transport has primarily focused on cooperative pushing behaviors on flat floors (Brown and Jennings, 1995; Parker, 1994; Rus et al., 1995; Wawerla et al., 2002), including some approaches that use the transported object to communicate implicitly. Most work in cooperative manipulation with force feedback focuses on fixed-based manipulators (Mukaiyama et al., 1996; Qingguo and Payandeh, 2003). A robot has also stacked masonry blocks (Pritschow et al., 1994).

In order to transport rigid components cooperatively with tight grasp, the robots must simultaneously maintain a formation. Most formation work applies potential fields using vision and/or explicit communication (Balch and Arkin, 1998; Carpin and Parker, 2002; Desay et al., 1999). These efforts do not require maintaining a grasp on a cooperatively carried component or tight bounds on permitted formation errors. JPL has demonstrated cooperative pick up, transport, and put-down of large components in an outdoor environment (Huntsberger et al., 2003; Trebi-Ollennu et al., 2002).

Instrument placement has relied on stereo vision with forward kinematics, visual servoing, and camera-space manipulation (CSM). Stereo with forward kinematics (Squyres et al., 2003) can run near real-time on flight systems but is susceptible to camera model errors and, for manipulation, to manipulator kinematics errors. Visual servoing uses high-rate images to eliminate relative errors (Allen et al., 1992; Feddema and Mitchell, 1989; Hutchinson et al., 1996; Nelson et al., 1996; Papanikolopoulos et al., 1991; Urmson et al., 2001). While robust to frame errors, delays between data or missing data can cause controller instability and large positioning errors (Feddema and Mitchell, 1989; Hutchinson et al., 1996; Papanikolopoulos et al., 1991), and it is not robust to noisy sensors (Allen et al., 1992). CSM uses widely space cameras to determine the relationship between joint space and image space and provides high positioning accuracy (less than 1 mm) (Baumgartner and Schenker, 1996; Chen et al., 1994; Skaar et al., 2001). However, CSM wide camera spacing requirements are impractical for rovers.

To date, no work has demonstrated end-to-end grasping, transport, and precision placement of a rigid component into a fixed structure by an autonomous team. In particular, work has not addressed such performance in a natural environment with motion and kinematics errors, variable lighting, and high-precision requirements. RCC demonstrates these capabilities in an outdoor-like laboratory setting using iterative processes, force-torque feedback, and Hybrid Image Plane Stereo (HIPS) to provide robustness despite uncertainty (Stroupe et al., 2005).

3. Task domain

3.1. Environment

The task domain consists of two heterogeneous robot platforms, a structure to be constructed, and a large sand pit in Planetary Robotics Laboratory at JPL that simulates benign outdoor terrain. The robots must move to the storage area, acquire the component, cooperatively transport the component to the structure, and place the component into the structure on the next layer.

Structure components are of two types: Rigid beams and flexible panels. Both types are four times the rover width in length. These components interlock when stacked. Each component has two grasping points (one left and one right) at which robots can hold the component. On the left and right of each component is a set of three fiducials that provide position and orientation of the component and, using a model of the component, the position and orientation of the grasping points. Figure 1 depicts the structure and beam components.

3.2. Platforms

The RCC team consists of the two heterogeneous platforms shown in Fig. 2, SRR and SRR2K. The robots are flight relevant in terms of computation. Each is holonomic and has a four degree-of-freedom arm. Additionally, each rover is

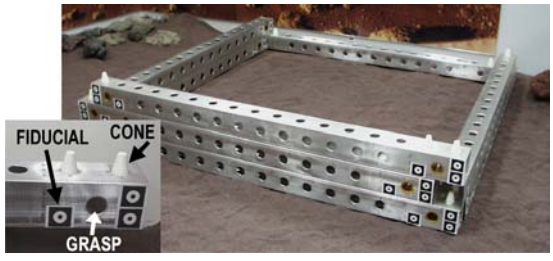

Fig. 1 The structure of interlocking beams. *Inset*: Component end with three fiducials and two interlocking cones

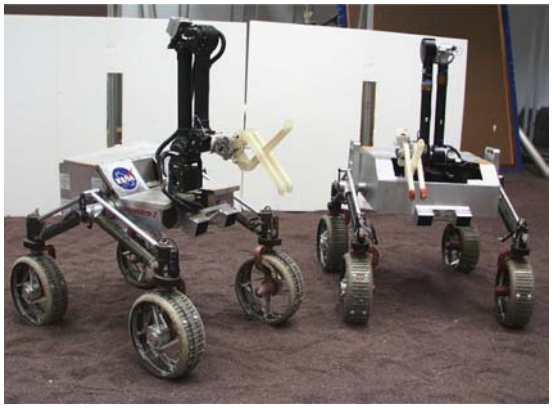

Fig. 2 Rovers SRR (left) and SRR2K (right) in the laboratory sandpit in front of component storage

equipped with a forward-facing stereo pair of cameras and a 3-axis force-torque sensor mounted at the base of the gripper. Power, computing, and wireless Ethernet are all on-board.

3.3. Challenges

There are several primary challenges to performing robust construction tasks in a natural environment that this research specifically addresses. First is the need to perform these tasks in real-time using flight relevant processing and sensing. Another primary challenge is cooperative transport of components which require the team to remain in a rigid formation to avoid damaging themselves or the components. After transport, the robots must position themselves accurately enough so that they can cooperatively place the components into the structure with the required high level of precision. Additionally, sensing is very limited (due to power and mass constraints for spacecraft) and is noisy; thus the team must perform robustly despite limited and inaccurate information. In some cases, target positions may not be visible during all stages. Lastly, operating on natural terrain exacerbates any small errors in motion due to slippage.

To address these challenges, several techniques are applied. First, a behavior-based architecture is highly adaptive to changing environments, thus RCC applies a behavior-based architecture to make the system more robust to errors and uncertainty. To mitigate errors in sensing and motions, RCC uses an iterative process for positioning. Robots refine position based on additional relative visual information. Additionally, for cooperative actions the team fuses all available sensing information in order to select that most appropriate action (Section 4). All cooperative actions execute in small parallel steps and force-torque sensing provides indirect feedback on the relative formation through the component allowing the team to identify and correct errors. (Section 5) Lastly, for precision placement, the camera models are defined relative to the arm frame to account for small errors in kinematics and models (Section 6). The team identifies catastrophic failures, those the team cannot handle autonomously, and reports them to the human operator.

4. Behavior-based control

The overall architecture is CAMPOUT, a behavior-based multi-robot control architecture. An overall view of the architecture is provided in Huntsberger et al. (2003). CAMPOUT gives commands to a real-time control system performing low-level actuator and sensor control. Generally, complex CAMPOUT behaviors are composed of simple platform-specific control and sensing behaviors. Task-specific behaviors are built using hierarchical combinations of generic complex and simple behaviors.

For the construction task, several high-level task specific behaviors were developed, as illustrated in Fig. 3. Figure 3 also illustrates the dependence of each task-specific behavior on the complex CAMPOUT behaviors such as drive, turn, look for fiducials, and communicate. Each of these task-specific behaviors is implemented as an iterative process in order to accommodate the many sources of error.

Figure 4 illustrates an example of how the iterative process for rover positioning cycles through crab, drive, and turn. Figure 5 illustrates how this process iteratively aligns the team for component placement. To maintain formation when operating cooperatively, the team moves jointly. At the start of each move the robots synchronize to ensure simultaneous

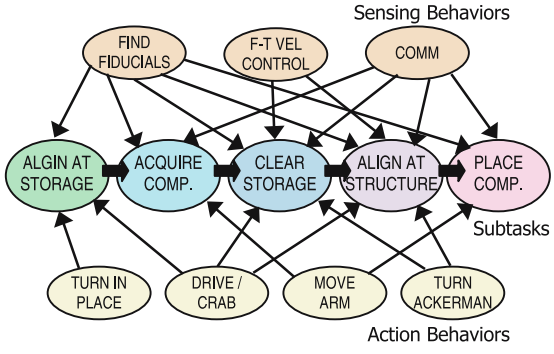

Fig. 3 Execution begins with Align at Storage and completes after Place Component. Two levels of the behavior hierarchy are shown. Large ovals are subtasks and small ovals are complex behaviors (sensing and action)

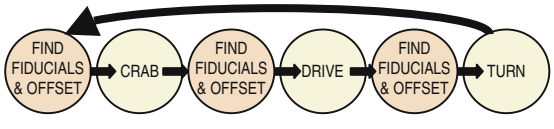

Fig. 4 Alignments cycle through crab, drive, and turn. Before each action, robots compute magnitude and direction from visible fiducials. When no further correction is necessary or a timeout occurs, the iteration completes

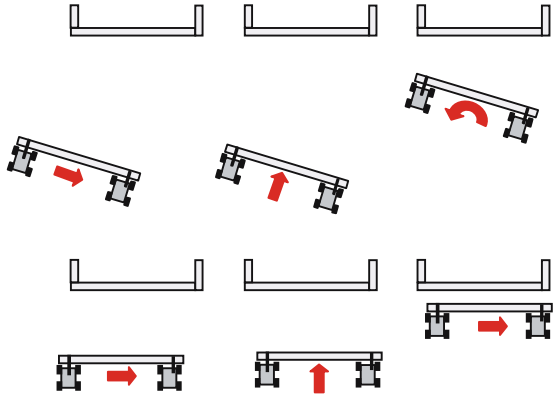

Fig. 5 An example of team alignment cycles through *crab*, *drive*, and *turn* as needed to finally align

actions. Theleader determines the action and sends it to the follower for simultaneous execution.

Prior to action selection, the robots exchange all available data so that the team selects the proper action. The team fuses data to ensure that both robots are selecting the same actions and that actions are based on all available information. The fusion method depends on the type of action. Typically, both robots will move the same magnitude and direction. For precision, however, the team must correct small errors in formation. Therefore, the robots incorporate small differences in desired motion by moving independently the desired distances.

For driving (forward or backward), the desired distance D_X for each robot is the distance to the structure G_X (as determined by the fiducials) minus the desired offset distance X_D. This distance is computed as in (1). If both robots see the structure fiducials, the team base drive distance is the individual distance with the smallest magnitude. If one robot sees the structure, the team base distance is that robot's individual distance.

$$D_X = G_X - X_D \tag{1}$$

For crabbing (left or right), the desired distance D_Y for each robot is the lateral distance to the grasp point G_Y (as determined by the fiducials) minus the desired offset Y_D to accommodate the robot's arm position relative to the robot's frame. This distance is computed as in (2). If both robots see the structure, the team base drive is the average distance. If one robot sees the structure, the team base distance is that robot's individual distance.

$$D_Y = G_Y - Y_D \tag{2}$$

For turning, the team must do complementary Ackermann turns around the center of the formation. The radius of the turn is half the distance between the grasp points

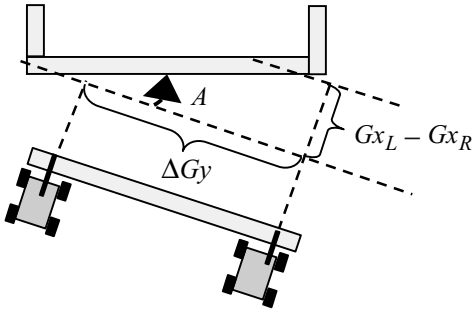

Fig. 6 Computing turn angle from two sets of observed fiducial range and the known distance between grasp points (ΔY_G)

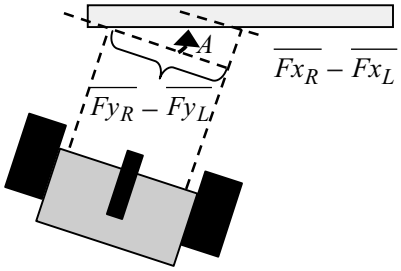

Fig. 7 Computing turn angle from one set of observed fiducial positions (left L and right R)

of the component. If both robots see the structure, the angle is determined by the distance between the left and right grasp points ΔY_G and the difference in distance from the left and right structure grasp points to the left and right robot, $\overline{Fx_R} - \overline{Fx_L}$. This computes the angle of the team relative to the structure and reduces the effects of small vision errors. Figure 6 illustrates this angle, which is computed as in (3). If one robot sees the structure, the positions of fiducials seen by the robot determine the angle: the difference in lateral distance between left and right fiducials $Fy_r - Fy_l$ and difference in range between left and right fiducials $Fx_r - Fy_l$. This angle is illustrated in Fig. 7 and computed as in (4).

$$A = \pm \tan^{-1}(\Delta Y_G, \overline{Fx_R} - \overline{Fx_L}) \tag{3}$$

$$A = -\tan^{-1}(Fy_r - Fy_l, Fx_r - Fx_l) \tag{4}$$

If at any time neither robot can see the structure (due to lighting conditions, for example), the team performs a small drive forward to attempt to reacquire the fiducials.

5. Force-torque feedback velocity control

When cooperatively carrying a component, the component transmits forces and torques between the robots. The magnitude and direction of these forces and torques provide information regarding the relative position of the teammate. If the

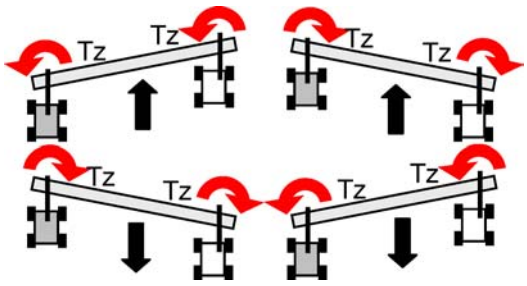

Fig. 8 The relationship of formation errors and torque (leader gray, follower white). *Left:* Leader is lagging behind. *Right:* Leader is moving ahead

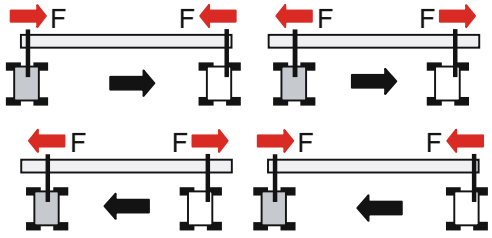

Fig. 9 Examples of the relationship of formation errors and lateral force. *Left:* Leader is falling behind. *Right:* Leader is moving ahead

robots are in the correct relative formation, these forces and torques are minimized. If one robot is front of the other, as shown in Fig. 8, this will result in a torque in the direction of the lagging teammate about the vertical axis (Tz). As shown in Fig. 9, if the robots are at the improper separation, they will experience lateral forces (Fy). Inward forces indicate they are too close together and outward forces indicate they are too far apart.

Once the team detects a formation error, they can correct it. In this implementation, only one robot (the follower, shown in white) takes corrective action to avoid overcompensation. If during a drive, the follower detects that the leader is ahead (Fig. 8 right column), it will speed up. Conversely, if it detects the leader is behind (Fig. 8 left column) it will slow down or stop. If during a crab, the follower determines that the leader is moving ahead (pushing as in Fig. 9 top right, or pulling as in Fig. 9 bottom right) it will speed up. If the leader is behind (pushing as in Fig. 9 top left, or pulling as in Fig. 9 bottom left) it will slow down or stop.

Due to disturbances from soil, motors, and noisy force-torque sensing, simple controllers bound forces and torques rather than trying to fully cancel them. The follower sets its velocity (slowing down or speeding up relative to the base speed) according to the simple PI controller shown in (5) to (7).

$$E = \begin{cases} Tz - Tz_0 - T_{Fy}, & \text{Drive} \\ Fy - Fy_0, & \text{Crab} \end{cases} \tag{5}$$

$$V_{NEW} = V_B + K_P E + K_I \sum_t E \tag{6}$$

$$\text{if } |V_{NEW} - V| > \Delta V_{MAX},\ V_{NEW} = V \pm \Delta V_{MAX} \tag{7}$$

The error E used to determine the desired velocity is relative to a reference point set when the robots are in the correct nominal formation (Tz_0 and Fy_0). T_{Fy} is the computed torque about z due to lateral force. The change in velocity for each step is bounded to ΔV_{MAX} to prevent jerky motion or exceeding acceleration capabilities. The mapping from force and torque to formation error was empirically calibrated through placing the robots at different formation offsets and measuring forces and torques. The selected control parameters values are based on empirical performance. V_B is the base (nominal) velocity of 6 cm/s. For drive, $K_P = 0.4$ and $K_I = 0.0$. For crab, $K_P = 1.5$ and $K_I = 0.05$. If the force or torque exceeds a threshold for a sustained period of time, the formation has failed, and the team stops. These thresholds, experimentally determined, are 4.0 N-m in torque and 50 N in lateral force sustained over 0.5 seconds. Velocity control is active after the initial acceleration (about 2 seconds) and until final deceleration.

6. Hybrid image plane stereo

Hybrid Image Plane Stereo (HIPS) provides high precision manipulator placement despite errors and uncertainty in kinematics. The core of the approach is the generation of camera models in the frame of the manipulator's end effector. These models are generated through comparing observed position of a fiducial on the end effector and the reported kinematics position of the end effector. HIPS continually updates models to account for any changes to the kinematics, as well as account for other types of errors. Thus, computed goal positions based on image coordinates match with arm configuration (rather than to ground truth) and improve manipulator position accuracy. These models, while quite precise for manipulator positioning, may be highly inaccurate in terms of world coordinates.

HIPS uses an 18-parameter CAHVOR model, a pinhole camera with symmetrical radial distortion. Full camera model estimation requires two steps. The initial calibration may be computationally expensive and is done offline ahead of time. The initial model accounts for any systematic errors including frame transformation errors and kinematics model errors in link lengths or offsets. The second model estimation step occurs online and readapts the models to time-varying errors and run-time uncertainties. Types of errors include flexion and droop (which may be orientation-dependent), joint resolution limitations, effects due to wear, finite image-plane cue detection, and additional camera modeling errors.

Initial estimation of this model begins by placing the manipulator at a sequence of predetermined poses that covers the portion of the manipulator work space that overlaps

with the visual field of view. At each pose, the stereo pair of cameras takes images. The fiducial on the end effector is located in the images and the "actual" position of the fiducial on the end effector is computed using the measured joint angles and forward kinematics. Using the data from each of the preplanned poses, the CAHVOR model parameters are estimated using a least-squares error method, providing a mapping from image coordinates to rover coordinates.

Once initial models are generated, they can be refined online during nominal operations. At points along manipulator trajectories, current images and joint angles provide new data. HIPS makes an incremental change to the camera models using this new data. The weight of the new sample in the recalibration may vary, and in this implementation does increase as the manipulator approaches the goal position to improve precision.

Fiducials for model estimation and for identifying goal locations in images (shown in Fig. 1) are identified in the image plane using gradient techniques to search for black-white-black transitions.

7. Experimental results

A number of experimental studies were conducted to determine the effectiveness of force-torque feedback to maintain formations by controlling velocity, of HIPS for precision manipulation operations, and of RCC to perform complex construction-related tasks. All studies were performed in the PRL large sand pit shown in Fig. 1.

7.1. Formations with force-torque feedback

To test formation keeping using force-torque feedback to control velocity, the team carried a beam for:

- 300 cm drives with varying follower start time offsets (−2 to +5 seconds)
- 80 cm crabs (left and right) with a −2 second follower start time offset.

Each experiment type was repeated without using velocity control. The 80 cm fixed velocity crab tests were also repeated without a start delay. The follower is on the formation's right. Failure conditions arise if forces or torques on either rover arm exceed thresholds as in Section 5.

Tables 1 and 2 summarize quantitative results for these experiments, including mean and standard deviation for each run. Statistical analysis covers only data from during the active control (2 seconds into the drive until stopping).

These results demonstrate the effectiveness of force-torque feedback velocity control for formation keeping. Using velocity control, the mean torque and force are consistently lower than they are without velocity control, reducing the stress on the manipulators. Mean torque during drive is (on average for 0 start time offset) 25 times greater without velocity control. Additionally, the standard deviation is much smaller, keeping the range of forces and torques much lower and nearer to zero. This reflects the robots remaining much nearer to the nominal formation and correcting errors that arise due to slippage and control differences. Without velocity control, the team reaches failure conditions much more frequently. Figure 10 shows an example of compared torque profiles for a drive and Fig. 11 shows compared force profiles for a crab. Note that the velocity control line (solid) stays near zero while the fixed velocity line (dotted) produces large errors, even in cases with large initial errors.

Velocity control can also correct for errors arising from initial offsets in position or start time. If the follower starts ahead or behind it will adjust its speed

Table 1 Mean torque in cooperative drive (N-m)

Start delay	Fixed velocity	Velocity control
0 sec	−0.43 ± 0.25	−0.06 ± 0.22
	−0.26 ± 0.43	−0.01 ± 0.20
	−0.95 ± 0.43	−0.03 ± 0.14
	−0.06 ± 0.69	−0.02 ± 0.13
	−1.76 ± 0.79	−0.02 ± 0.13
2 sec	1.22 ± 0.50	−0.03 ± 0.18
	1.94 ± 0.53	−0.01 ± 0.21
	1.09 ± 0.63	−0.03 ± 0.21
	1.40 ± 0.91	0.00 ± 0.21
	1.46 ± 0.67	0.02 ± 0.16
−2 sec	−2.09 ± 1.02	−0.04 ± 0.16
	−2.22 ± 0.80	−0.05 ± 0.22
	Failure	−0.04 ± 0.21
	−2.03 ± 0.79	0.02 ± 0.20
	−2.30 ± 0.81	−0.03 ± 0.20
−5 sec	Failure	0.07 ± 0.21
	Failure	0.06 ± 0.15
	Failure	0.03 ± 0.17
	Failure	0.08 ± 0.16
	Failure	0.06 ± 0.18

Table 2 Force profiles in cooperative crab (*N*)

Drive	Fixed velocity (0 sec)	Fixed velocity (−2 sec)	Velocity control (−2 sec)
80 cm	3.91 ± 8.67	Failure	−3.76 ± 7.16
	−1.48 ± 8.67	−14.04 ± 7.59	0.05 ± 5.38
	−5.54 ± 6.66	−7.30 ± 9.42	−0.44 ± 4.59
−80 cm	−6.73 ± 12.38	5.27 ± 9.90	1.61 ± 6.15
	6.98 ± 11.06	Failure	−0.02 ± 3.93
	15.56 ± 9.86	−3.56 ± 9.34	−0.78 ± 4.45

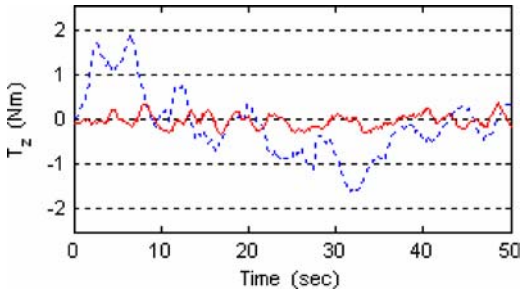

Fig. 10 With velocity control during a drive (solid), torques remain near zero with little variance. Without velocity control (dotted), torques vary greatly and reach high values

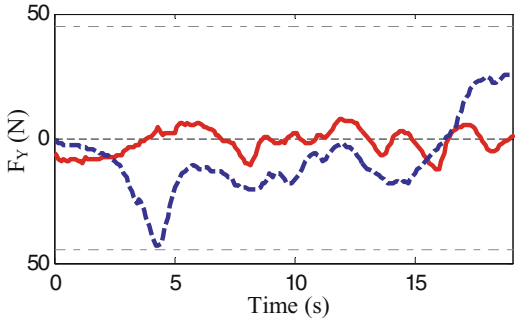

Fig. 11 With velocity control during a crab (solid), forces remain near zero and have little variance. Without velocity control (dotted), the forces reach near failure (4.5 sec) and have large variance

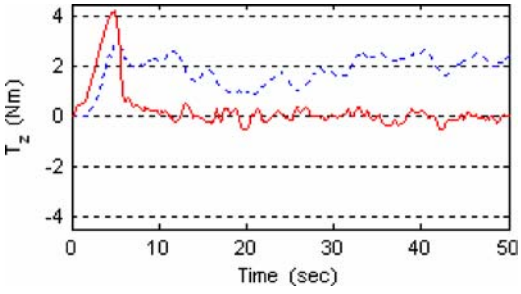

Fig. 12 With velocity control (solid), a formation error causing high torques is corrected and torques return to near zero. With fixed velocity (dotted), initial error remains

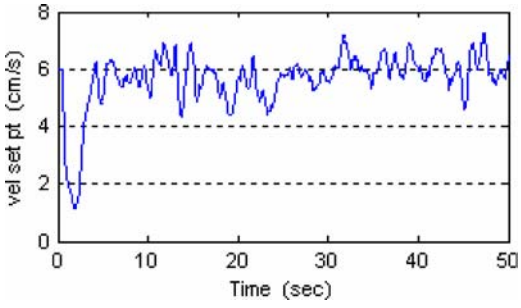

Fig. 13 The follower starts moving while the leader is stationary (0 sec). The follower observes increasing torque indicating a lagging leader and stops (2 sec). As the leader starts to move and torque reduces, the follower gradually increases speed but remains below nominal speed. When the leader catches up (5 sec) the follower reaches nominal speed and remains near nominal with minor corrections

until the forces return to nominal. If the follower (using velocity control) starts before the leader, it starts velocity control and becomes ready to respond to any formation errors. As a result, it slows down and waits for the leader to catch up before returning to nominal speed. In this way, the formation can accommodate large head-starts by the follower. If the follower starts after the leader, it can speed up to catch up. This can accommodate offsets limited by the follower's ability to accelerate. Larger offsets will reach failure condition before velocity control can fully compensate, but the recognition of such a failure condition stops the team prior to damaging the manipulators and component.

Figures 12 and 13 show examples of correcting initial offsets. In Fig. 12, the follower starts 2 seconds after the leader. This results in a build-up of force while the follower is stationary. Once the follower starts moving, it increases speed to catch up and reduces the forces to nominal. Figure 13 shows an example velocity profile in which the follower compensates for a late-starting leader. Initially, the follower starts to move and experiences a negative force. The velocity quickly reduces to zero while the leader is stationary (1 second). As the leader starts to move (2 seconds) and the negative force slowly reduces, the follower begins to increase velocity but remains below nominal speed. As the leader finally catches up at 5 seconds, putting forces near nominal, the follow reaches and maintains near-nominal velocity.

In summary, force-torque feedback velocity control greatly increases the robustness of formation keeping. It corrects for small errors introduced due to slippage and control differences and compensates for many initial formation errors. Finally, it allows the team to stop motion if forces or torques exceed safe bounds to prevent catastrophic failure.

7.2. Manipulator positioning accuracy with HIPS

Several types of experiments were conducted in simulation and using several research robotic platforms. These experiments include:

- *Simulation Verification:* This set of simulation experiments compares positioning accuracy of HIPS to that of traditional stereo and forward kinematics. Large kinematics errors were added after camera model estimation (combined total of 2.0 cm change in link lengths) to test robustness.
- *Manipulator Positioning Accuracy:* For these experiments, SRR places its manipulator on an observed target. Position error is measured.

Figure 14 illustrates the simulation results for 50 runs with each model. Using only initial HIPS camera models, performance improves relative to traditional stereo, with a mean error reduction of 60%. However, performance using initial models degrades as range increases as the errors

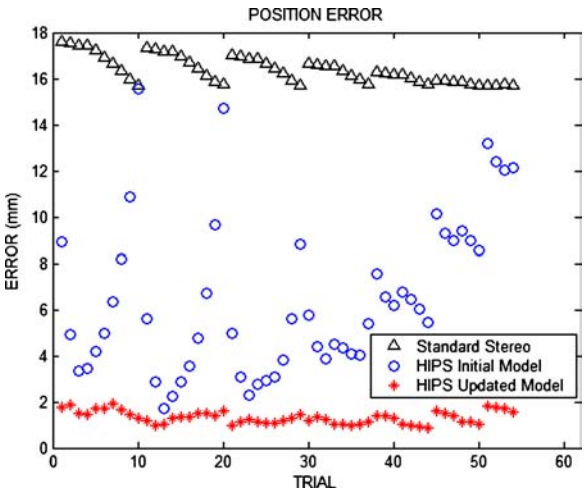

Fig. 14 Average error was over 1.0 cm for standard stereo. Using only the initial model, HIPS shows an improvement to 6.4 mm. With model updates average error dropped to 1.3 mm

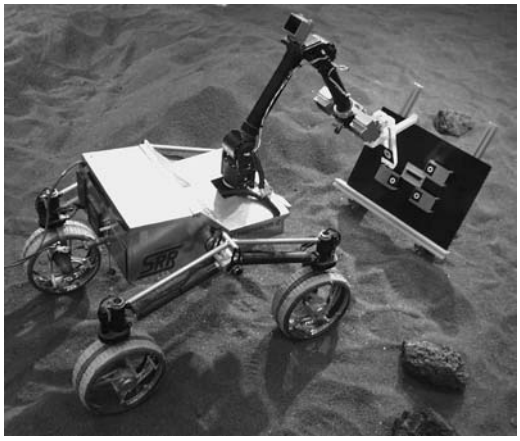

Fig. 15 SRR demonstrates manipulator placement at a goal point on a target board

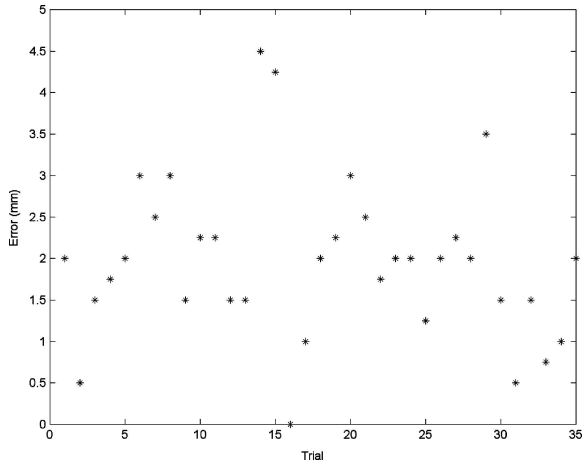

Fig. 16 Results of 35 positioning tests show an average error of less than 2.0 mm

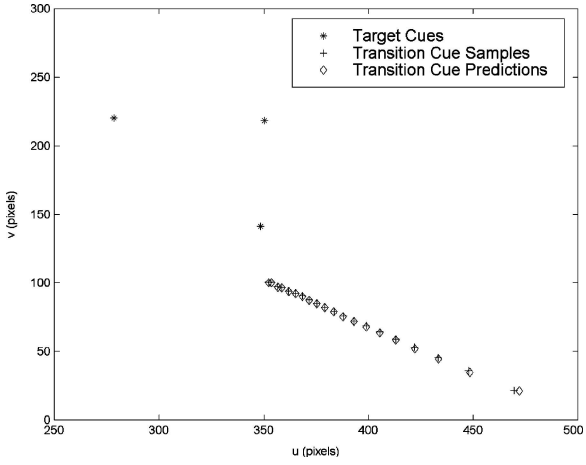

Fig. 17 Updating the model as the arm approaches the goal at (350,100) improves accuracy during approach

in kinematics have greater impact at greater range. Using updates during operation reduces mean error by more than 90%. Additionally, updates correct the errors at larger range seen in the initial models.

For the experiments using SRR, initial camera models were generated using 35 manipulator poses. SRR attempted to position the manipulator to a goal on the target board shown in Fig. 15 (in the middle of three fiducials). A total of 35 positioning tests were conducted.

Figure 16 shows the position test results. The average error in these tests was 1.96 mm with a 3σ error bound of ± 2.9 mm. The terminal positioning error in nearly 85% of these tests was less than or equal to 2.5 mm. This is an improvement compared to the approximately 10 mm accuracy provided by traditional stereo on other flight-relevant platforms.

The improvement in accuracy achieved by updating the model online during operation is apparent in Fig. 17. Note that as the manipulator moves from the initial position (lower right) toward the goal point (center) the errors reduce.

In summary, HIPS improves the accuracy of manipulator placement by computing estimated camera models relative to manipulator kinematics rather than to ground truth. This provides improved repeatability and a reduction of up to 90% in positioning error compared to traditional stereo and forward kinematics approaches.

7.3. Construction task performance

The test environment for each of the construction task experiments is a large sand pit that provides benign outdoor-like terrain. In all cases, the structure foundation is placed at an unknown location and orientation but known direction. The team places the component into the structure on the first face of the structure they encounter. Four classes of experiment were conducted:

- *Acquire Beam:* The rovers begin at arbitrary locations where a component is storage is visible. The rovers independently align at a beam in storage. Then the rovers independently reach out to place their grippers into the grasp points. Finally, the rovers cooperatively lift the beam into a carrying position.
- *Align at Structure:* The rovers begin holding a beam, in formation, at an arbitrary orientation relative to the structure and such that at least one rover can see the fiducials on the structure. The rovers then align with the structure within 1-cm position accuracy for placement.
- *Place Beam:* The rovers begin immediately after completion of an Align at Structure experiment. The rovers cooperatively move the beam over the structure and then lower it into place. Finally, the rovers release the beam and stow their arms.
- *End to End:* The rovers begin as for an Acquire Beam experiment. They proceed to acquire the beam. After acquisition they back away from the structure, turn to face the direction of the structure, and drive until the structure becomes visible to at least one rover. The team then proceeds to complete Align at Structure and Place Beam.

Each experiment type was run several times. Any failure to achieve the desired result is noted, along with the cause of the failure.

Table 3 summarizes the construction task results. Figure 18 shows a series of photographs from several different experimental tests, in operational sequence.

The Align at Structure failure occurred when both robots reached a location from which they could not see the structure fiducials. This prompted the addition of the small bump forward described in Section 4, and no further failures of this kind occurred. The overall success rate of RCC for construction tasks (breaking end to end up into component parts) is 100% for acquiring beams, 95% for aligning at the structure (100% with bump mode for reacquiring the structure), and 100% for placing beams.

Table 3 HIPS simulation results

Experiment	Mean error	3σ
Stereo	16.5 mm	1.7 mm
Initial HIPS	6.4 mm	9.3 mm
Updated HIPS	1.3 mm	0.8 mm

Table 4 Construction results

Experiment	Runs	Failures
Acquire beam	24	0*
Align at structure	19	1
Place beam	18	0
End to end	5	0

*Excludes a non-algorithmic failure due to a poorly caliberated wrist.

Fig. 18 *Top left*: Align at Storage brings rovers into grasping position. *Top Right*: Align Grasp places hand in grasping position. *Second Left*: Team lifts the component out of storage in Grasp Component. *Second Right*: Team turns away from storage in Clear Storage. *Third Left*: Rovers begin Align at Structure. *Third Right*: Rovers complete Align at Structure at correct relative position. *Bottom Left*: Rovers begin Place Component. *Bottom Right*: Rovers complete Place Component

8. Additional preliminary results

8.1. Flexible panel components

In addition to the beams shown in the construction experiments presented previously, some work has investigated handling flexible panel components, as shown in Fig. 19. The RCC has successfully aligned with the structure while carrying a panel and deployed the panel into the structure in several preliminary trials.

The adaptation of the behaviors from handling beams to carrying panels requires only providing several parameters representing the model of the panel (the side of the panel and the location of its grasp point relative to fiducials). Using force-torque feedback velocity control will require

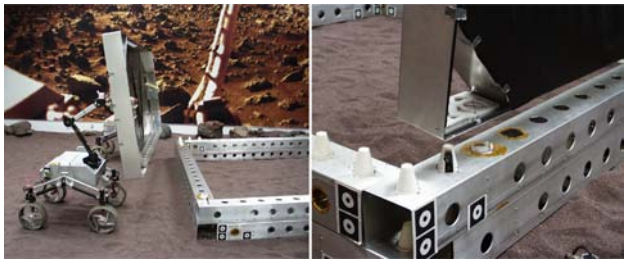

Fig. 19 *Left*: The team carries a panel and aligns with the structure. *Right*: The team places the panel into the structure

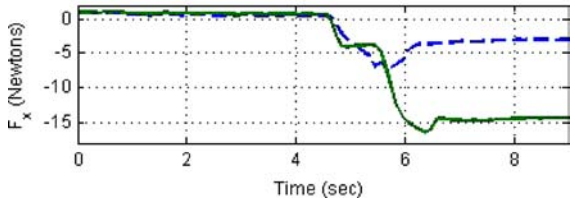

Fig. 20 Examples of an arm approaching a component. In a nominal grasp (dotted) the gripper enters the grasp point but rubs along the grasp bottom. The robot sees small friction forces. In a missed grasp (solid), the gripper hits the component and pushes it. The robot experiences large forces and detects failure

a calibration of parameters relating forces and torques to formation errors.

8.2. Force feedback for position error detection

In order to make grasping and placement more robust, the force-torque sensing capabilities are being applied to detect errors. In approaching a desired position (for grasp or for placement), the robots monitor manipulator forces. Grasp misalignment results in the finger pushing the component rather than entering the grasp point, which increases axial forces on the manipulator. Similarly, lowering a misaligned component (such that it contacts cones) increases vertical forces. If forces increase above nominal an error condition is present. Figure 20 compares forces during a successful and unsuccessful beam grasps.

Recently, detecting and correcting these errors during grasping has been implemented. Upon detecting an error as an increase in axial force, the rover performs a local search, retracting and re-extending the arm at several points in an attempt to find the grasp (an extension with nominal forces). The rover has successfully compensated for small misalignments in several preliminary tests.

9. Future work

Current efforts in RCC are focusing on increasing robustness and expanding the construction capabilities.

Currently, force-torque feedback velocity control can only correct forces or torques reflecting formation errors in the direction of motion. This may allow other forces to grow and lead to failure. To prevent this, in addition to compensating in velocity, the robot can adjust steering direction to minimize off-axis forces. This would also enable velocity control during Ackermann turns.

Force-torque feedback will be expanded to eliminate more manipulator positioning failures, taking advantage of the failure detection capabilities described in Section 8.2. The force-torque sensor detects such increases, and the team could take corrective action.

Fiducial detection using current gradient methods is highly sensitive to lighting conditions. To improve robustness, this sensitivity needs to be reduced to handle the highly variable lighting outdoors.

The current acquire-transport-deploy scenario is only one step in the construction of a full structure. This work will expand to a larger-scale construction mission including more realistic types of habitats, multiple heterogeneous components and larger transport distances.

10. Conclusion

The Robotic Construction Crew project has demonstrated autonomous multi-robot construction and assembly capabilities in simulated natural terrain. Construction tasks include acquisition, manipulation, transport, and precision placement of two types of construction components (Section 7.3). A behavior-based system that tightly couples current state and sensor information with action within a hand-tuned task decomposition and sequencing structure provides performance reliability (Section 4). The robustness of maintaining formations required for successful cooperative transport improves by using force-torque feedback for velocity control (Section 5). Applying Hybrid Image Plane Stereo (HIPS) to directly calibrate cameras in the manipulator's frame also improves manipulator positioning robustness (Section 6). Performance results demonstrate a nearly 100% success rate in our single construction sequence task in an indoor natural environment. Future work will integrate more realistic habitat components, more complex construction scenarios, and greater fault detection and recovery.

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.

We thank Dr. Paul Schenker and Dr. Neville Marzwell for supporting the RCC project and Brett Kennedy, Tony Ganino, Lee Magnone, and Mike Garrett for supporting the RCC platforms.

REFERENCES

Allen, P.K., Timcenko, A., Yoshimi, B., and Michelman, P. 1992. Real-time visual servoing. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 1850–1856.

Balch, T. and Arkin, R.C. 1998. Behavior-based formation control for multi-robot teams. *IEEE Transactions on Robotics and Automation*, 14(6):926–939.

Baumgartner, E.T. and Schenker, P.S. (1996). Autonomous image-plane robot control for martian lander operations. In *Proceedings of the International Conference on Robotics and Automation,* pp. 726–731.

Brookshire, J., Singh, S., and Simmons, R. 2004. Preliminary results in sliding autonomy for coordinated teams. In *Proceedings of the 2004 Spring Symposium Series*.

Brown, R.G. and Jennings, J.S. 1995. A pusher/steerer model for strongly cooperative mobile robot manipulation. *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 3, pp. 562–568.

Carpin, S. and Parker, L.E. 2002. Cooperative leader following in a distributed multi-robot system. In *Proceedings of the IEEE International Conference on Robotics and Automation*, vol. 3: pp. 2994–3001.

Chen, W.Z., Korde, U., and Skaar, S.B. 1994. Position-control experiments using vision. *International Journal of Robotics Research*, 13(3):199–208.

Desay, J.P., Kumar, V., and Ostrowski, P. 1999. Control of change in formation for a team of mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation*, vol. *2*, pp. 1556–1561.

Feddema, J.T. and Mitchell, O.R. 1989. Vision-guided visual servoing with feature-based trajectory generation. *IEEE Transactions on Robotics and Automation*, 5(5):691–700.

Howe, A.S. 2000. Designing for automated construction. *Automation in Construction*, 9(3):259–276.

Huntsberger, T., Pirjanian, P., Trebi-Ollennu, A., Nayar, H.D., Aghazarian, H., Ganino, A., Garrett, M., Joshi, S. S. and Schenker, P. S. 2003. CAMPOUT: A control architecture for tightly coupled coordination of multi-robot systems for planetary surface exploration. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans, Collective Intelligence*, 33(5):550–559.

Hutchinson, S., Hager, G., and Corke, P. 1996. A tutorial on visual servo control. *IEEE Transactions on Robotics and Automation*, 12(5):651–670.

Mukaiyama, T., Kyunghwan, K., and Hori, Y. 1996. Implementation of cooperative manipulation using decentralized robust position/force control. In *Proceedings of the 4th International Workshop on Advanced Motion Control*, vol. 2: pp. 529–534.

NASA Office of Exploration Systems. Human and Robotic Technology (H and RT) Formulation Plan. Version 3.0, May 14, 2004.

Nelson, B., Papanikolopoulos, N. P., and Khosla, P. 1996. Robotic visual servoing and robotic assembly tasks. *IEEE Robotics and Automation Magazine*, 3(2):23–31.

Papanikolopoulos, N., Khosla, P. K., and Kanade, T. 1991. Vision and control techniques for robotic visual tracking. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 857–864.

Parker, L.E. 1994. ALLIANCE: an architecture for fault tolerant, cooperative control of heterogeneous mobile robots. In *Proc.eedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 2: pp. 776–683.

Pritschow, G., Dalacker, M., Kurz, J., and Zeiher, J. 1994. A mobile robot for on-site construction of masonry. In *Proceedings of the IEEE/RSJ/GI International Conference on Intelligent Robots and Systems*, vol. 3, pp. 1701–1707.

Qingguo, L. and Payandeh, S. 2003. Multi-agent cooperative manipulation with uncertainty: a neural net-based game theoretic approach. In *Proc. IEEE International Conference on Robotics and Automation*, vol. 3, pp. 3607–3612.

Rus, D., Donald, B., and Jennings, J. 1995. Moving furniture with teams of autonomous robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 1, pp. 235–242.

Simmons, R., Singh, S., Hershberger, D., Ramos, J., and Smith, T. 2000. First Results in the coordination of heterogeneous robots for large-scale assembly. In *Proceedings of the International Symposium on Experimental Robotics*.

Skaar, S.B., Seelinger, M.J., Robinson, M.L., and Gonzalez-Galvan, E.J. 2001. Means and method of robot control relative to an arbitrary surface using camera-space manipulation. U.S. Patent 6 304 050, October 16.

Squyres, S. et al., 2003. Athena investigation overview. *Journal of Geophysical Research*.

Stroupe, A., Huntsberger, T., Okon, A., and Aghazarian, H. 2005. Precision manipulation with cooperative robots. multi-robot systems: From Swarms to Intelligent Automata Volume III. Schultz et al (Eds).

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. *Proc. IEEE International Conference on Robotics and Automation*.

Urmson, C., Shamah, B., Teza, J., Wagner, M.D., Apostolopoulos, D., and Whittaker, W.L. 2001. A sensor arm for robotic antarctic meteorite search. In *Proceedings of the 3rd International Conference on Field and Service Robotics*, Helsinki, Finland.

Wawerla, J., Sukhatme, G.S., and Matarić, M.J. 2002. Collective construction with multiple robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 3, pp. 2696–2701.