Behavior-Based Multi-Robot Collaboration for Autonomous Construction Tasks*

Ashley Stroupe, Terry Huntsberger, Avi Okon, Hrand Aghazarian, Matthew Robinson
Jet Propulsion Laboratory / California Institute of Technology
4800 Oak Grove Drive, Pasadena, CA 91109
{Ashley.Stroupe, Terry.Huntsberger, Avi.Okon, Hrand.Aghazarian, Matthew.Robinson}@jpl.nasa.gov

***Abstract* – The Robot Construction Crew (RCC) is a heterogeneous multi-robot system for autonomous construction of a structure through assembly of long components. The two-robot team demonstrates component placement into an existing structure in a realistic environment. The task requires component acquisition, cooperative transport, and cooperative precision manipulation. A behavior-based architecture provides adaptability. The RCC approach minimizes computation, power, communication, and sensing for applicability to space-related construction efforts, but the techniques are applicable to terrestrial construction tasks.**

Index Terms* – *space robotics, cooperative mobile manipulation, multi-robot systems.

I. INTRODUCTION

The current NASA roadmap calls for a human Lunar presence by 2020, followed by human Martian exploration [8]. For safety, these missions require prior placement of infrastructure (habitats, power, oxygen, etc). Robotic technologies must perform site clearing and leveling, component transport, placement, and docking, and structure inspection and repair. Communication delays and blackouts require that much of the robotic construction be autonomous. Launch and space operations constraints require systems to minimize mass, volume, and power while remaining robust to uncertainty and errors. This paper presents early results for Robotic Construction Crew (RCC) in a construction scenario including cooperative component grasping, transport, and precision placement by a heterogeneous team.

To accommodate space-driven constraints, our approach minimizes sensing, computation, and communication. By design, the distributed behavior-based control approach that tightly couples current sensing with execution eliminates the most processing-intense tasks such as mission planning and optimal task-decomposition. The rovers use a completely distributed control approach during independent operations. During team tasks (cooperative transport and manipulation) rovers use a synchronized leader-follower approach with centralized decision making and distributed execution to maintain tight coupling of team members.

The Robotic Construction Crew has demonstrated reliable capabilities of component acquisition, transportation, manipulation, and placement in an outdoor-like laboratory environment. While primarily developed for planetary habitat construction, these capabilities are also applicable to Terrestrial construction activities.

* This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

II. BACKGROUND AND 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) [2][13]. The vision robot aligns where it can see the necessary work space and tells the manipulator robots how to adjust for alignment. Work in cooperative transport has primarily focused on cooperative pushing behaviors on flat floors [3][9][10][12][16], including some approaches that use the transported object to communicate implicitly. Most work in cooperative manipulation with force feedback focuses on fixed-based manipulators [7]. A robot has also stacked masonry blocks [11]. 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 [1][4][5]. JPL has demonstrated cooperative transport and deployment of large components in an outdoor environment [6][15].

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. RCC demonstrates these capabilities in an outdoor-like laboratory setting.

III. TASK DESCRIPTION

A. Environment

The structure is a set of interlocking components. The team must obtain components from a storage unit, transport them to the construction site and place them into the structure. Structure components are aluminum beams (180 x 12.5 x 12.5 cm) with cones that interlock with the beam above. At each end is a grasping point, a cylindrical channel through the component. Three fiducials mark each grasp's location. Fig. 1 shows the structure and a component end in detail.

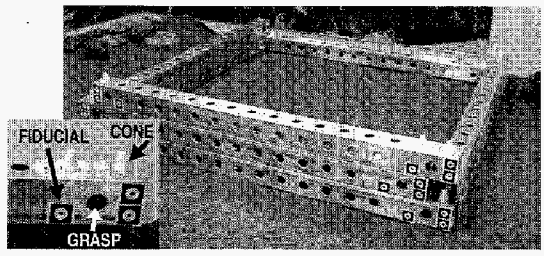

Fig. 1. The structure of interlocking beams. *Inset*: Component end with three fiducials and two interlocking cones.

The SRR and SRR2K rovers (Fig. 2) comprise the heterogeneous team. Each has four steered wheels and a unique four-joint manipulator arm. Rovers are equipped with on-board wireless modems, computing, and battery power.

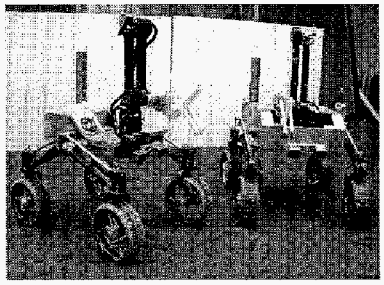

Fig. 2. SRR (left) and SRR2K (right) rovers.

B. *Sensing*

One pair of stereo cameras is at the front of each robot to provide three-dimensional visual sensing. Stereo calibration uses Hybrid Image Plane/Stereo (HIPS) [In preparation for publication]. HIPS generates camera models through direct visual sensing of the manipulator's end-effector in conjunction with end-effector position estimation by manipulator kinematics. By correlating manipulator kinematic position with three-dimensional position, manipulator placement accuracy improves by approximately a factor of 2.5 over traditional calibrated stereo. If run online (not done in this work), continued adaptation of the manipulator/camera models improves placement by up to an additional factor of five and can account for changes in system configuration and ensure consistent precision for the life of the mission.

Each manipulator has a three-axis force-torque sensor positioned at the wrist to sense forces and torques imparted by the component on the manipulator. This provides passive communication between the rovers about relative team and load positions during cooperative transport, which is not observable through the forward-facing cameras.

IV. Approach

A. *Overview*

The overall architecture is CAMPOUT, a behavior-based multi-robot control architecture [6]. CAMPOUT provides commands to a real-time control system performing low-level actuator and sensor control. The overall construction task is decomposed by hand into a series of subtasks. These subtasks are in turn composed of general, reusable complex COMPOUT behaviors which are composed of simple platform-specific control and sensing behaviors.

B. *Subtasks*

The Construction task calls a series of subtasks, each designed to execute one stage of the construction task. Successful completion of each subtask by the team triggers transition to the next subtask. Fig. 3 illustrates this sequence and shows the partial behavior hierarchy. For cooperative tasks, leader-follower control with action synchronization ensures common action selection, parallel action execution, and team recognition of success and failure conditions.

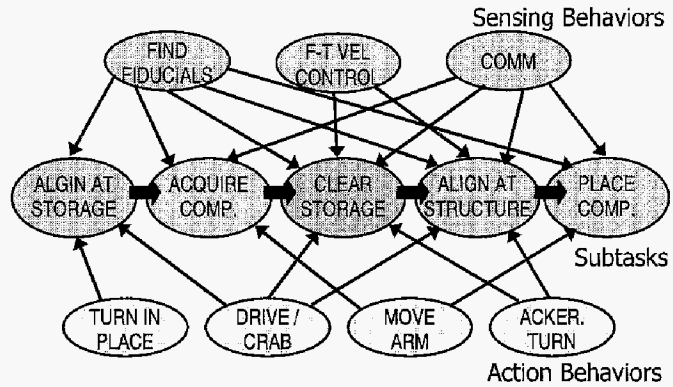

Fig. 3. Execution begins with *Align at Storage* and completes after *Place Component*. Two-levels of the behavior hierarchy are shown. Large ovals are subtasks and small ovals are complex behaviors (sensing and action).

Several subtasks require that the team (independently or cooperatively) align relative to the visual cues. Due to sensing uncertainty and motion uncertainty, executed motions do not match desired. To provide accuracy despite such errors, RCC uses the iterative approach to alignment illustrated in Fig 4. This process iteratively corrects range, lateral offset, and relative angle until no corrections remain or a timeout occurs.

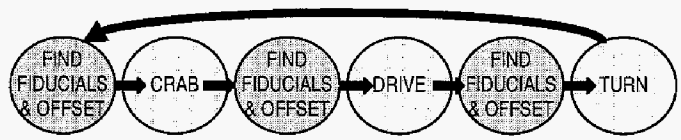

Fig. 4. Alignments cycle through *crab*, *drive*, and *turn*. Before each action, robots compute magnitude and direction from visible fiducials When no further correction is necessary or a timeout occurs, the iteration completes.

Align at Storage: The rovers independently place themselves in position to grasp components in storage. After determining the necessary correction using *find grasp* before each move, the behaviors *drive*, *crab*, and *turn in place* cycle until the rover aligns within tolerance or timeout occurs, as shown in Fig. 5.

Fig. 5. Alignment, as needed, independently crabs to the correct lateral offset, drives forward to the correct range, and turns in place to the correct heading. The iterative process continues until alignment or timeout.

Each robot independently determines drive distance, crab distance, and turn angle using (1), (2), and (3). D_x is drive distance, D_y is *crab* distance, and A is *turn in place* angle. (G_x, G_y, G_z) is the computed grasp point position, X_d and Y_d are the desired distance and lateral offset from the grasp point, and (F_x, F_y, F_z) are the fiducial positions. If the average of the right fiducials (subscript R) is closer than the average left (subscript L), the turn is positive/clockwise to counter the angle (example shown in Fig. 6).

$$D_x = G_x - X_d \tag{1}$$

$$D_y = G_y + Y_d \tag{2}$$

$$A = -\tan^{-1}\left(\overline{F_{yR}} - \overline{F_{yL}}, \overline{F_{xR}} - \overline{F_{xL}}\right) \tag{3}$$

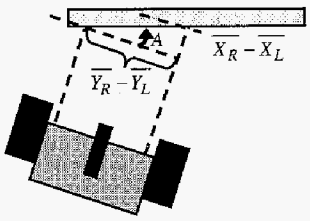

Fig. 6. Computing angle from observed fiducial positions (left L and right R).

Acquire Component: Rovers place grippers in position to grasp the component. *Find Grasp* determines the grasp location. *Move arm* positions the hand at the same lateral and vertical position as the grasp in front of the grasp, and then moves the hand forward into the grasp. The grasp requires rover pose within 1 cm and 1.5° and arm position within 1 cm. Once in position, both rovers grasp the component and lift it and then move it into carrying position using *move arm*. At each stage of the lift, the robots *synchronize* to ensure simultaneous motions and that the component remains level.

Clear Storage: The team moves to where the structure is visible. The team uses *drive* to back away, uses *Ackermann turn* to face the structure, and then uses *drive* to move the team to where the structure is visible by *find grasp* (Fig. 7). Robots *synchronize* at each stage. The follower adjusts velocity to maintain formation using *force-torque velocity control*.

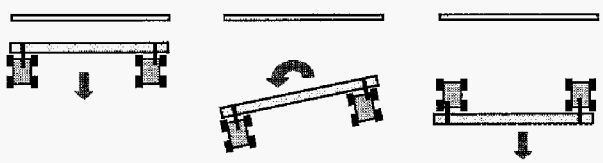

Fig. 7. *Clear Storage* backs up (left) to make room to turn (center) and then drives toward the structure (right).

Align at Structure: The team moves into position to put the component into the structure. The iterative process of *crab*, *drive*, and *Ackermann turn* continues until aligned or timeout occurs (Fig. 8). Before each move, robots *share data* from *find grasp*. The leader selects an action based on all available information, from either robot (coarse align) or both robots (fine align). The leader sends the action with *share data* and team *synchronizes*. For *drive* and *crab*, the follower maintains formation via *force-torque velocity control*.

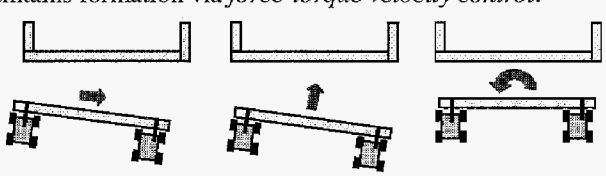

Fig. 8. *Align at Structure* cycles through *crab*, *drive*, and *Ackermann turn* to bring the team into alignment.

For coarse alignment, with data from only one robot, *drive* distance, *crab* distance, and *Ackermann turn* angle are determined as in (1)-(3). For fine alignment, with data from both robots, the data is fused. For *drive* and *crab*, each rover computes distance as in (1) and (2). The final distance is the minimum magnitude for *drive* (to minimize chance of collision) and the average magnitude for *crab* (to satisfy both offsets). For small drives (<5cm), each rover moves independently to the desired range. For *Ackermann turn*, the difference in F_x to the left and right (averaged over the three fiducials) and the lateral separation of left and right grasp points ΔY_G determine the angle (4). If the left side is further than the right side, the turn is clockwise to counter the angle (Fig. 9). If at any point neither robot can obtain visual data, the team nudges forward to try to regain sight of the structure.

$$A = \mp \tan^{-1}\left(\Delta Y_G, \overline{F_{xR}} - \overline{F_{xL}}\right) \tag{4}$$

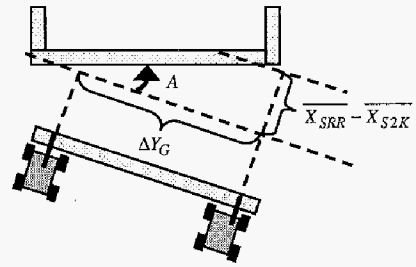

Fig. 9. Computing angle from observed fiducial range and the known distance between grasp points (ΔY_G).

Place Component: The team places the component into the structure. The rovers independently obtain the relative position of the components in the structure using *find grasp*. Using *move arm*, the rovers move the component directly above the desired position in the structure, and then lower the component into the structure. Each step is *synchronized*. Placement requires arm range and rover lateral position accuracy within 1 cm, and rover heading accuracy within 1.5°.

C. Complex Behaviors

Complex behaviors used by subtasks are described here:

Drive/Crab: A rover moves a specified distance based on odometry. For *crab*, the rover first turns wheels to a specified angle and straightens wheels after completion.

Turn in Place: A rover turns wheels inward, moves to turn the specified angle (with odometry), and straightens wheels.

Ackermann Turn: A rover turns wheels to the desired turning radius, moves the specified angle on that arc (based on odometry), and straightens the wheels.

Move Arm: A rover moves the manipulator to the specified end-effector or joint configuration.

Find Fiducials: Stereo vision finds three-dimensional fiducial positions. The rover validates observations using a geometric model of the fiducials and makes multiple attempts to find valid fiducials before failing. Three fiducials provide position and orientation of the grasp point relative to the robot.

Comm: The leader sends a message and waits for a response; the follower waits for a message and then responds. The message may be a ready message (*synchronize*) or a message containing visual data or actions (*share data*).

Force-Torque Velocity Control: The follower sets its velocity to return load forces and torques to nominal and maintain formation. Leader and follower monitor forces and torques for failure conditions. The mapping of force and torque to relative formation was experimentally determined during static and driving conditions. Due to the complex grip-component interaction, the direction of the force imparted by the component on the manipulator may vary during motion, making approaches like [3] ineffective. The torque about the vertical axis does, however, correlate with the component

orientation in the gripper. Thus, the rover detects and corrects for a leading or lagging partner (twisting the component off 90°) by speeding up or slowing down. During *crab*, lateral force on the manipulator correlates with rover separation. The rover detects and corrects for a leading/lagging partner (pulling/pushing the component) by changing speed (Fig. 10).

Fig. 10. The relationship between formation errors and velocity control.

Due to disturbances from soil, motors, and noisy force-torque sensing, simple controllers bound forces and torques rather than trying to fully damp them. (5) and (6) show the driving controller. ΔV_{max} is maximum velocity change (1 cm/s), T_Z is torque about the vertical axis, T_{Z0} is reference torque (set after grasping), and T_{FY} is torque due to the lateral force (rather than angular misalignment). V_B is base velocity, ±6 cm/s, and K_P is ±0.4. Control parameters are selected based on empirical performance. Crab control is similar.

$$V_{new} = V_B + K_P\left(T_Z - T_{Z0} - T_{FY}\right) \tag{5}$$

$$V = \begin{cases} V_{last} + \Delta V_{max} & \text{if } V_{new} - V_{last} > \Delta V_{max} \\ V_{last} - \Delta V_{max} & \text{if } V_{new} - V_{last} < \Delta V_{max} \\ V_{new} & \text{otherwise} \end{cases} \tag{6}$$

V. Results: Adaptive Velocity Control

A. Experimental Setup

As described in IV.C, the follower keeps formation by adjusting *drive* and *crab* velocity based on force-torque feedback from a sensor in the manipulator. Rovers begin in formation with the reference torque set to the initial condition; the follower starts at a variable time offset.

B. Results

Two sets of experiments were conducted: a 250 cm drive and an 80 cm crab. Both experiments were run with and with velocity control. Statistics are computed on data during the velocity-controlled part of the drive only. Failure occurs if the arm experiences torques or forces exceeding experimentally determined thresholds (4.0 N-m or 50 N) for 0.5 seconds.

Table I shows results of the drive experiments as mean and standard deviation of torque. Table II shows results for the crab experiments as mean and standard deviation of force.

By adjusting velocity based on force-torque feedback, torques and forces on the arms remain well within operational range despite offsets in start time (Fig. 11 and 12). Velocity control keeps mean torque and force near zero. Without velocity control, torques and forces become large and have high variance. If the follower starts before the leader, it can sense the increase in torque or force, reduce velocity to 0, and wait; as the leader catches up, it increases velocity and the rovers regain formation. Thus, velocity control accommodates any delay short of timeout. If the follower starts after the leader, its ability to compensate is limited to delays that do not produce failure conditions before velocity control begins; this depends on control parameters and maximum velocity change.

Table I: Torque Profiles in Cooperative Transport: *Drive*

Start Offset	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 Transport: Crab

Drive	Fixed Velocity (No Delay)	Fixed Velocity (Delay)	Velocity Control (Delay)
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

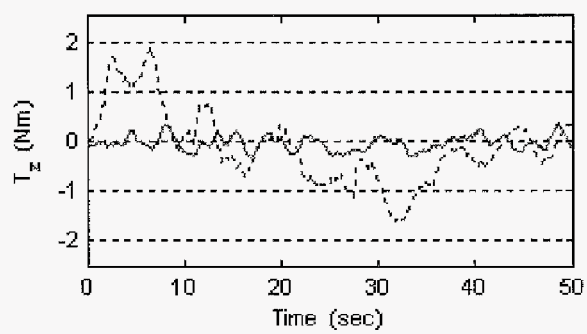

Fig. 11. With velocity control, the follower (solid) quickly compensates for initial errors and keeps torque bounded and rovers in formation. Without control (dotted), large errors with large variance are induced and rovers remain out of formation.

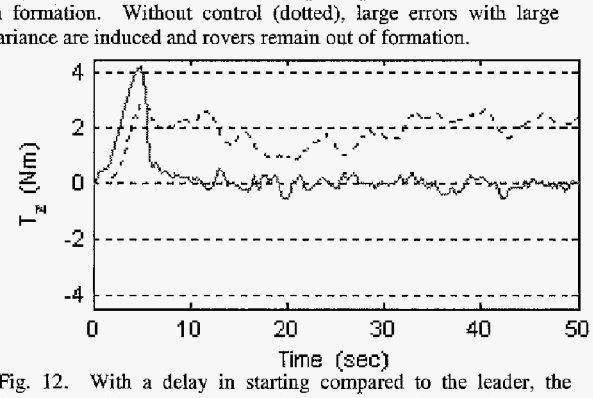

Fig. 12. With a delay in starting compared to the leader, the follower is able to speed up to reduce torques and regain formation (solid). Without velocity control (dotted), robots remain out of formation causing large torques, nearly failing.

Fig. 11 and Fig. 12 compare torques from *drives* with and without velocity control. Velocity control damps any initial offsets and reduces bias and variance. Without velocity control, torques are highly variable and commonly reach failure condition. *Crab* provides similar results (Fig. 13) [14].

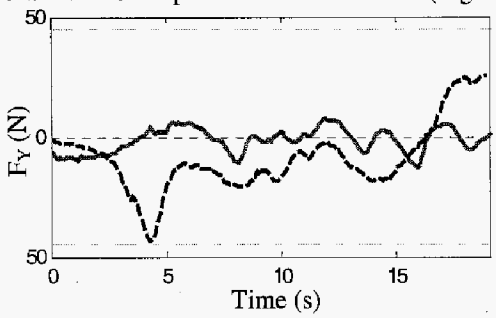

Fig. 13. Force profiles of an -80 cm crab with the follower on the right. Note high forces and large variance without control (dotted) compared to velocity control (solid). Force bounds shown dashed.

Fig. 14 shows follower velocity for the experiment shown in Fig. 11. Velocity is zero while waiting for the leader (starting 2 seconds later) to catch up. Velocity increases as leader starts to move, and adjusts to maintain nominal torque.

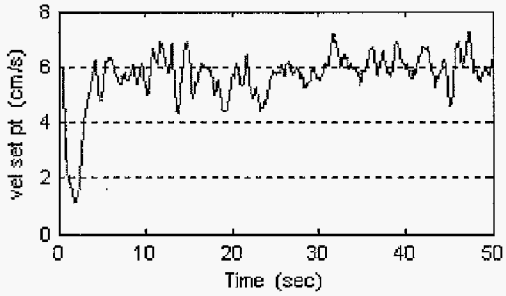

Fig. 14. The follower reduces velocity to near 0 until the leader begins to move (t=2). The follower then increases velocity gradually allowing the leader to catch up (t=5) and then remains near base velocity once formation is regained.

By adjusting velocity based on force-torque feedback, torques and forces on the arms remain well within operational range despite initial offsets. Velocity control keeps mean torque near zero. Without velocity control, torques become large and have high variance. If the follower starts before the leader, it can sense the increase in torque, reduce velocity to 0, and wait; as the leader catches up, it increases velocity and the rovers regain formation. Thus, velocity control accommodates any delay short of timeout. If the follower starts after the leader, its ability to compensate is limited to delays that do not produce failure conditions before velocity control begins; this depends on control parameters and maximum velocity change.

VI. Results: Construction Tasks

A. Experimental Setup

Experiments are conducted in a large sand pit that provides benign outdoor-like terrain for operations. In each of the three experiment types, the structure foundation is in place at an unknown location and orientation but known direction.

Acquire Component: The rovers are individually positioned and oriented randomly but such that at least one component fiducial is visible for each robot. Subtasks *Align at Storage*, *Align Grasp*, and *Grasp Component* run consecutively. A failure occurs is a rover fails to align, a rover manipulator fails to align, or if grasping and lifting fails.

Component Placement: The rovers, holding a component, are positioned randomly but such that one or both can see the fiducials on the structure foundation. Subtasks *Align at Structure* and *Place Component* run consecutively. Failure occurs if the team fails to align or to place the component properly into the structure.

End to End: The rovers are individually positioned as in the Acquire Component experiments. All subtasks run consecutively. Failure modes include those mentioned above as well as failures in formation during *Clear Storage*.

B. Results

Table III shows quantitative results for each experiment. Fig. 15 shows snapshots of the construction task.

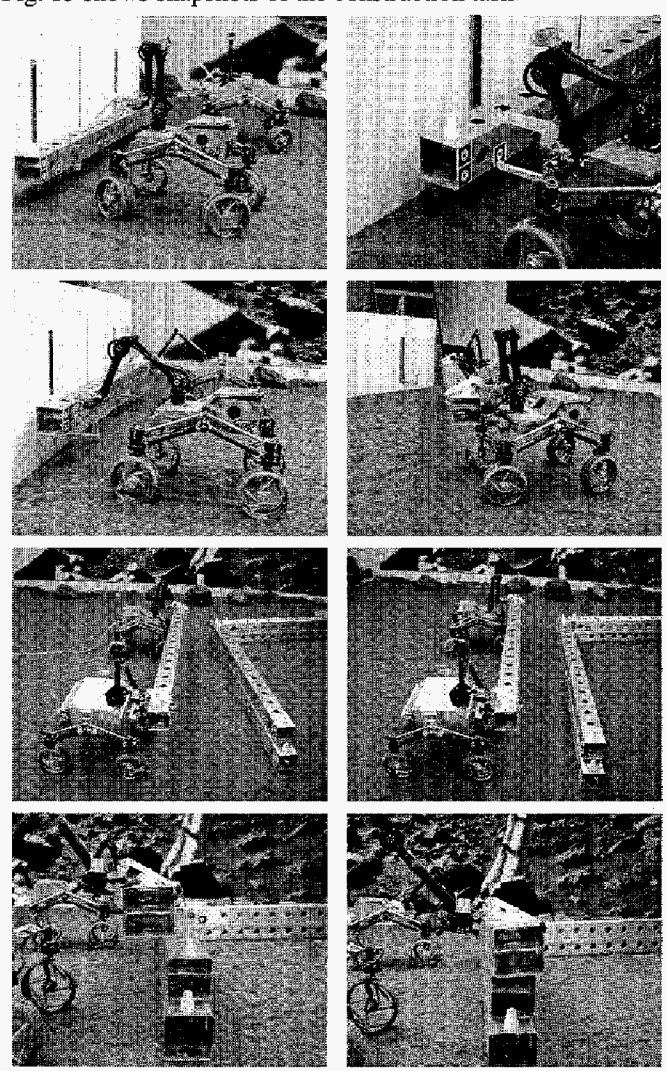

Fig. 15. *Top left*: *Align at Storage* brings rovers into grasping position. *Top Right*: *Align Grasp* places hand in grasping position. *Second Left*: Team lifts the component out of storage in *Grasp Component*. *Second Right*: Team turns away from storage in *Clear Storage*. *Third Left*: Rovers 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*.

Table III: Construction Results

Experiment	Runs	Failures
AC	25	1
CP	17	1
EE	5	0

The failure in Acquire Component occurred due to poor wrist calibration, causing the hand to align below the grasp and push the component; thus, this is not an algorithm failure. Component Placement failure occurred when vision did not return reliable fiducial locations due to lighting variability; noisy sensing is the primary cause of failure. Performance improved by adding the *nudge* behavior that changes position and attempts to reacquire the fiducials.

The team successfully completes all the construction tasks with a very low failure rate. Section VII, Future Work, discusses means of addressing remaining failure conditions.

VII. Future Work

Two primary areas of research are in current development for improvement in RCC: robustness and expanding the task.

Current force-torque feedback only adjusts velocity along the direction of motion; this may allow forces to grow along other dimensions 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 manipulator positioning failures. 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 beam increases vertical forces. The force-torque sensor detects such increases, and the team could then take corrective action.

The fiducial detection is highly sensitive to lighting conditions. To improve robustness, this sensitivity needs to be reduced to handle variable lighting outdoors.

The current task is only one step in the construction of a full structure. This work will be integrated into a larger-scale construction task including multiple heterogeneous components and larger transport distances. Current efforts are under way to transport and place large panels. Later, techniques applied in the Robotic Construction Crew will be applied to constructing/assembling realistic habitats.

VIII. Summary and Conclusions

The Robotic Construction Crew has demonstrated autonomous multi-robot construction and assembly capabilities in simulated natural terrain. Construction tasks include acquisition, manipulation, transport, and precision placement of construction components. Performance reliability is provided by using a behavior-based system that tightly couples current state and sensor information with action within a hand-tuned task decomposition and sequencing structure. The robustness of maintaining formations required for successful cooperative transport improves by using force-torque feedback for velocity control.

Like all behavior-based systems, behaviors are hand tuned. However, individual tuning of a single behavior is simple and once tuned behaviors can be very quickly adapted for new and different tasks. Other objects, for example, have been located using the same fiducial markers and manipulated (using the same arm control) simply by defining new geometric models of the object, its fiducials, and the desired relationship between the object and its environment.

Acknowledgment

We thank Dr. Paul Schenker for initiating Robot Work Crew and Dr. Neville Marzwell and Dr. Eric Baumgartner for supporting RCC. We thank Brett Kennedy, Tony Ganino, Lee Magnone, and Mike Garrett for supporting RCC platforms.

References

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

[2] J. Brookshire, S. Singh, and R. Simmons. "Preliminary Results in Sliding Autonomy for Coordinated Teams." Proc. 2004 Spring Symposium Series, March, 2004.

[3] R.G. Brown and J.S. Jennings. "A pusher / steerer model for strongly cooperative mobile robot manipulation." Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems, 3:562-568, 1995.

[4] S. Carpin and L.E. Parker. "Cooperative leader following in a distributed multi-robot system." Proc. IEEE International Conference on Robotics and Automation, 3:2994-3001, 2002.

[5] J.P. Desay, V. Kumar, and P. Ostrowski. "Control of Change in Formation for a Team of Mobile Robots." Proc. IEEE International Conference on Robotics and Automation, 2:1556-1561, 1999.

[6] T. Huntsberger, P. Pirjanian, A. Trebi-Ollennu, H. D. Nayar, H. Aghazarian, A. Ganino, M. Garrett, S.S. Joshi, P.S. Schenker. "CAMPOUT: A Control Architecture for Tightly Coupled Coordination of Multi-Robot Systems for Planetary Surface Exploration." IEEE Transactions on Systems, Man & Cybernetics, Part A: Systems and Humans, Collective Intelligence, 33(5): 550-559, 2003.

[7] T. Mukaiyama, K. Kyunghwan, and Y. Hori. "Implementation of cooperative manipulation using decentralized robust position/force control." Proc. 4th International Workshop on Advanced Motion Control, 2:529-534, 1996.

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

[9] L.E. Parker. "ALLIANCE: an architecture for fault tolerant, cooperative control of heterogeneous mobile robots." Proc. IEEE/RSJ Int. Conference on Intelligent Robots and Systems, 2:776-683, 1994.

[10] G. Pritschow, M. Dalacker, J. Kurz, and J. Zeiher. "A mobile robot for on-site construction of masonry." Proc. IEEE/RSJ/GI International Conference on Intelligent Robots and Systems 3:1701-1707, 1994.

[11] L. Qingguo and S. Payandeh. "Multi-agent cooperative manipulation with uncertainty: a neural net-based game theoretic approach." Proc. IEEE Int. Conference on Robotics and Automation, 3:3607-3612, 2003.

[12] D. Rus, B. Donald, and J. Jennings. "Moving furniture with teams of autonomous robots." Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems, 1:235-242, 1995.

[13] R. Simmons, S. Singh, D. Hershberger, J. Ramos, and T. Smith. "First Results in the Coordination of Heterogeneous Robots for Large-Scale Assembly." Proc. Int. Symp. on Experimental Robotics, 2000.

[14] A. Stroupe, T. Huntsberger, A. Okon, and H. Aghazarian. "Precision Manipulation with Cooperative Robots." Multi-Robot Systems: From Swarms to Intelligent Automata Volume III. Schultz et al (Eds), 2005.

[15] A. Trebi-Ollennu, H. Das, H. Aghazarian, A. Ganino, P. Pirjanian, T. Huntsberger, and P. Schenker. "Mars Rover Pair Cooperatively Transporting a Long Payload." Proc. IEEE International Conference on Robotics and Automation, 2002.

[16] J. Wawerla, G.S. Sukhatme, and M.J. Mataric. "Collective construction with multiple robots." Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems, 3:2696-2701, 2002.