Behavior-Based Multi-Robot Collaboration for Autonomous Construction Tasks

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

***Abstract* – We present a heterogeneous multi-robot system for autonomous construction of a structure through assembly of long components. Placement of a component within an existing structure in a realistic environment is demonstrated on a two-robot team. The task requires component acquisition, cooperative transport, and cooperative precision manipulation. For adaptability, the system is designed as a behavior-based architecture. For applicability to space-related construction efforts, computation, power, communication, and sensing are minimized, though the techniques developed are also applicable to terrestrial construction tasks.**

Index Terms* – *Multi-robot teams, autonomous construction, cooperative transport.

I. Introduction

The current NASA roadmap calls for a human Lunar presence by 2020, followed by human Martian exploration [3][4][5]. For safety, these missions require prior placement of infrastructure (habitats, power, oxygen, etc) using robotic technologies. Construction tasks include site clearing and leveling, component transport, placement, and docking, and structure inspection and repair. Due to communication delays and blackouts, much of the robotic construction must 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 Robotic Construction Crew, early results in a robotic construction scenario, including cooperative component grasping, transport, and precision placement by a heterogeneous team.

To accommodate space-driven constraints, sensing, computation, and communication are minimized. As a result, the processing intense tasks high-level planning and task decomposition are provided within a hand-designed distributed behavior-based control approach that tightly couples current sensing with execution. The control approach is completely distributed during independent operations. During team tasks (cooperative transport and manipulation) a synchronized leader-follower approach with centralized decision making and distributed execution is applied to maintain tight coupling of team members.

With the Robotic Construction Crew reliable capabilities of component acquisition, transportation, manipulation, and placement has been demonstrated in a simulated outdoor environment. These capabilities, while developed for planetary habitat construction, are also applicable to Terrestrial construction activities.

This work is funded by NASA's something or other program.

II. Background and Related Work

The 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, find manipulation) [1][10]. In previous work, JPL has demonstrated cooperative transport and manipulation of large components and deployment of a simple pull-out structure in an outdoor or outdoor-like environment [2][11]. Additional work in cooperative transport has primarily focused on cooperative pushing behaviors on flat floors [6][7] [9][12]. A single-robot applied to stacking masonry block has also been demonstrated [8].

Previous work has not yet demonstrated end-to-end component acquirement, transport, and precision placement within a rigid structure. This work demonstrates all of these capabilities in an outdoor-like setting (a large sand pit in the Planetary Robotics Lab) in continuous autonomous operation.

III. Task Description

The habitat mock-up used for these experiments consists of a set interlocking components. The team must obtain components from a storage unit, transport them to the construction site and place them in the structure.

Structure components are aluminum beams (180 x 12.5 x 12.5 cm) with regularly spaced holes on each face and cones at each end that interlock with the beam above. At each end is a grasping point, a front-back pair of holes with a cylindrical guide. Each grasping point is marked by three fiducials, T0 (top), T1 (right), and T2 (left). The end of one component in the storage unit is shown in Fig. 1.

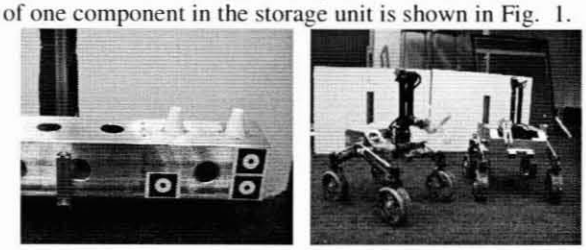

Fig. 1. *Left*: Component end with three fiducials and two interlocking cones. *Right*: Rovers SRR (left) and SRR2K (right).

Two rovers are used as the construction team, SRR and SRR2K (Fig. 1). Each rover has four steered wheels and a four degree-of-freedom manipulator with a three-fingered gripper, though the configurations are heterogeneous. Rovers are equipped with wireless modems and on-board computing and battery power.

IV. Approach

A. *Overview*

The behavior architecture is CAMPOUT [Ref], a behavior-based multi-robot control architecture. 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. As an example, two layers of the behavior hierarchy are shown in (Fig. 2).

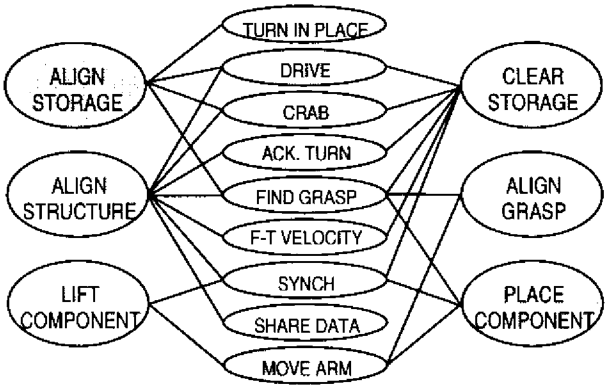

Fig. 2. Two-level behavior hierarchy. Large outer ovals are subtasks and center small ovals are complex behaviors.

B. *Behavioral Architecture*

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. This sequence is illustrated in (Fig. 3) and details of each subtask and high-level complex behavior are provided below.

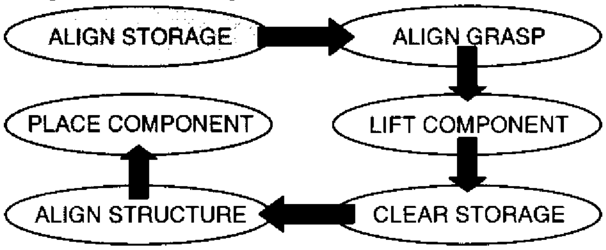

Fig. 3. Flow of subtasks begins with Align at Storage and completes after Place Component.

1. Subtasks

Align at Storage: This subtask drives the rovers independently to place them in position to grasp the components within component storage. This is done using a cycle of behaviors: *drive*, *crab*, and *turn in place* until the rover is aligned within tolerance, as illustrated in Fig. 4. The *find grasp* behavior determines the grasp position. Drive distances and turn angles are determined independently for each robot as in Eq 1 – Eq. 3. *D* is drive distance, *A* is turn angle, G is grasp point position, D_d is desired distance from grasp point, M is manipulator offset in the rover frame, and T1 and T2 are the positions of targets 1 and 2.

$$D_x = G_x - D_d \qquad \text{Eq. 1}$$

$$D_y = G_y + M_y \qquad \text{Eq. 2}$$

$$A = -\tan^{-1}\big((T1_x - T2_x),(T1_y - T2_y)\big) \qquad \text{Eq. 3}$$

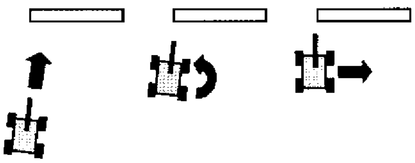

Fig. 4. Align at Storage cycles through drive to approach, turn in place to align perpendicularly and crab to align laterally. This subtask is distributed and independent.

Align Grasp: This subtask places the gripper in position to close on the component. *Find Grasp* determines the component grasp location. Then, *move arm* first positions the hand at the same lateral and vertical position as the grasp point but at minimum extension and then moves the hand forward into the grasp.

Grasp Component: This subtask grasps the component and lifts it into the carrying position for each robot. This is done using *move arm* to first lift the component up a specified distance, then pull the component back, and then adjust to the carrying position. At the beginning of each stage, the robots *synchronize* to keep motions parallel.

Clear Storage: This subtask moves the robots to the structure. This is done by using *drive* to back away, using *Ackerman turn* to face the structure, and then using *drive* to move the team to where the structure is clearly visible by *find grasp* (Fig. 5). Robots *synchronize* at each stage. The leader sets the goal and the follower adjusts velocity to maintain formation using *force-torque velocity update*.

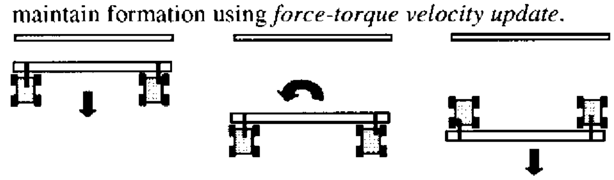

Fig. 5. Clear Storage backs up (left) to make room to turn (center) and then drives toward the structure (right).

Align at Structure: This subtask places the team in position to put the component into the structure. Before each move, robots *share data* and *synchronize*. This subtask has two modes: both robots see the component and one robot sees the component. If only one robot sees the component, the team cycles through *Ackerman turn*, *drive*, and *crab* to bring the component fully into view as shown in Fig. 6. The turn is a small angle (5°) in the direction to make the robot not seeing the component closer to the expected location of the structure. The forward drive distance is the linear distance corresponding to that angle (Eq. 4) and the crab distance is as in Eq 3.

$$D_x = a\frac{(2\pi r)}{360} \qquad \text{Eq. 4}$$

If both robots see the component, the team cycles through *drive*, *Ackerman turn*, and *crab* until the team is within tolerance (Fig. 7). Drive distance is the minimum magnitude drive for both robots (Eq. 1), crab distance is the average for both robots (Eq. 3), and the Ackerman turn magnitude moves the robots to the same distance (Eq. 5):

$$A = -\tan^{-1}(DSx - DKx, B) \quad \text{Eq. 5}$$

where B is the length of the component and DSx and DKx are the distances for SRR and SRR2K, respectively. For precision, when distances are small (<5cm), drive distances are independently determined using Eq. 1. After aligning at the structure grasping points, the team moves laterally to offset the component for the next layer. During each *drive* and *crab*, the follower maintains formation using *force-torque velocity update*.

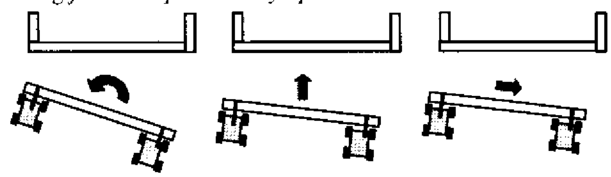

Fig. 6. Align at Structure when only one robot sees the structure (left rover) cycles through Ackerman turn, drive and crab to bring the second rover closer to the structure.

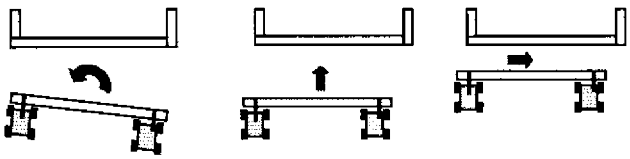

Fig. 7. Align at Structure when both rovers see the structure cycles through Ackerman turn, drive, and crab based on all observations to bring robots into alignment.

Place Component: This subtask places the component into the structure on top of previous components. This is done by obtaining the relative position of the components in the structure using *find grasp*. Using *move arm*, the rovers move the component directly above the position in the structure into which it will be placed, and then lower the component into the structure.

2. Complex Behaviors

Each task is composed of several complex behaviors, which are briefly described here:

Drive: The rover moves forward or backward a specified distance based on odometry.

Crab: The rover turns the wheels to a specified angle, drives (moving at an angle relative to heading) based on odometry, and straightens the wheels.

Turn in Place: The rover turns the wheels to a circular configuration and drives (based on odometry) to turn the specified amount.

Ackerman Turn: The robot turns each wheel to an angle to set desired turning radius, drives forward or backward (based on odometry) on that arc the specified angle, and straightens the wheels.

Move Arm: The robot moves the manipulator to the specified end effector or joint configuration.

Find Grasp: The robot uses stereo vision to find three-dimensional positions of fiducials on the components. The black and white fiducials are identified in images using gradients. Each component has six fiducials, three near each grasping point, and a model of the relationship of these fiducials and the grasping point is known. The three points provide position and orientation of the grasping point relative to the robot. This behavior returns the grasp point position and fiducial positions.

Synchronize: The leader sends ready message and waits for acknowledgment; the follower waits for ready message and sends an acknowledgement. After acknowledgement, rovers are synchronized.

Share Data: Rovers *synchronize* with a message containing data on visual position of the grasping points on visible components (in storage or in the structure).

Force-Torque Velocity Update: The rover (follower) sets its velocity to return load forces and torques to nominal and maintain formation. The mapping of force and torque to relative formation position was experimentally determined during static and driving conditions. During a *drive*, while the direction of the force imparted by the component on the manipulator may vary, the direction and magnitude of the torque is directly correlated with the orientation of the component within the gripper. Thus, a leading or lagging partner (pulling the component off 90 degrees) can be detected and corrected with speeding up or slowing down. During a *crab* (lateral drive), the lateral force on the manipulator is directly correlated to the separation between the rovers. Thus, a leading/lagging partner (pulling/pushing the component) can be detected and corrected by changing speed.

C. Sensing

1. Stereo Vision

One pair of stereo cameras is positioned at the front of each robot to provide three-dimensional visual sensing. The stereo pair is calibrated using 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. During operation, continued estimation and 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.

2. Force-Torque Sensing

Each rover manipulator is equipped with a three-axis force-torque sensor positioned at the base of the gripper. This senses forces and torques imparted by the component on the manipulator. This provides passive communication between the rovers about relative position of the team and the load during cooperative transport, which cannot be observed through the forward-facing cameras.

V. Results: Formation Keeping

A. Experimental Setup

As described in IV.B.2, the follower keeps formation by adjusting velocity during *drive* and *crab* based on force-torque feedback provided by a force-torque sensor in the manipulator wrist. Rovers begin in formation; the follower rover starts at a variable time after the leader.

B. *Results*

Table I compares mean torque about vertical observed during forward and backward driving; Table II compares lateral forces observed during crabbing. In each case, the mean and number of failures are shown. Failures occur if the torque or force moves out of safety bounds.

Table I: Formation Keeping: Drive

Start Offset	No Feedback	Feedback
0 sec	11 (0)	2 (0)
2 sec	17 (2)	2.5 (0)
4 sec	22 (7)	3.2 (0)

Table II: Formation Keeping: Crab

Start Offset	No Feedback	Feedback
0 sec	3.2 (0)	0.3 (0)
0.5 sec	5.1 (5)	0.4 (0)
1 sec	7.2 (8)	0.8 (0)

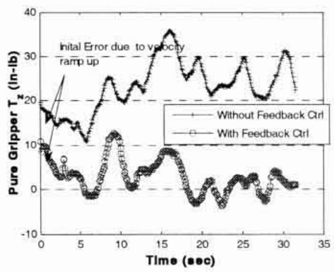

Fig. 8. Using force-torque feedback to set velocity, the lagging follower speeds up to restore formation.

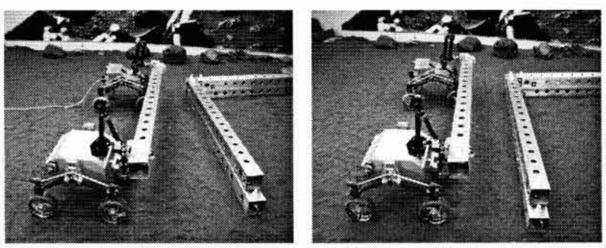

Fig. 9. Using force-torque feedback to set velocity, the lagging follower speeds up to restore formation.

By adjusting velocity based on force-torque feedback, the rovers are able to compensate for initial formation mismatch and keep forces well within operational range despite offsets in start formation. A comparison of force profiles resulting from driving with and without velocity updating is shown in (Fig. 8). An example of regaining formation is shown in (Fig. 9).

VI. Results: Construction Tasks

A. *Experimental Setup*

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

Component Grasp: The rovers are individually positioned randomly but such that they can see the fiducials on the component in storage. Subtasks *Align at Storage*, *Align Grasp*, and *Grasp Component* are run consecutively. A failure occurs is a rover fails to align, a rover manipulator fails to align, or if grasping fails.

Component Placement: The rovers, holding a component, are positioned randomly but such that one or more can see the fiducials on the component on the structure foundation. Subtasks *Align at Structure* and *Place Component* are run consecutively. A failure occurs if the team fails to align or if the component is not properly placed into the structure.

End to End: The rovers are individually positioned as in the Component Grasp experiments. All subtasks are run consecutively. Failures include all those mentioned above.

B. *Results*

Quantitative results are shown for each type of experiment in Table III. Fig. 10 illustrates each step in the construction task.

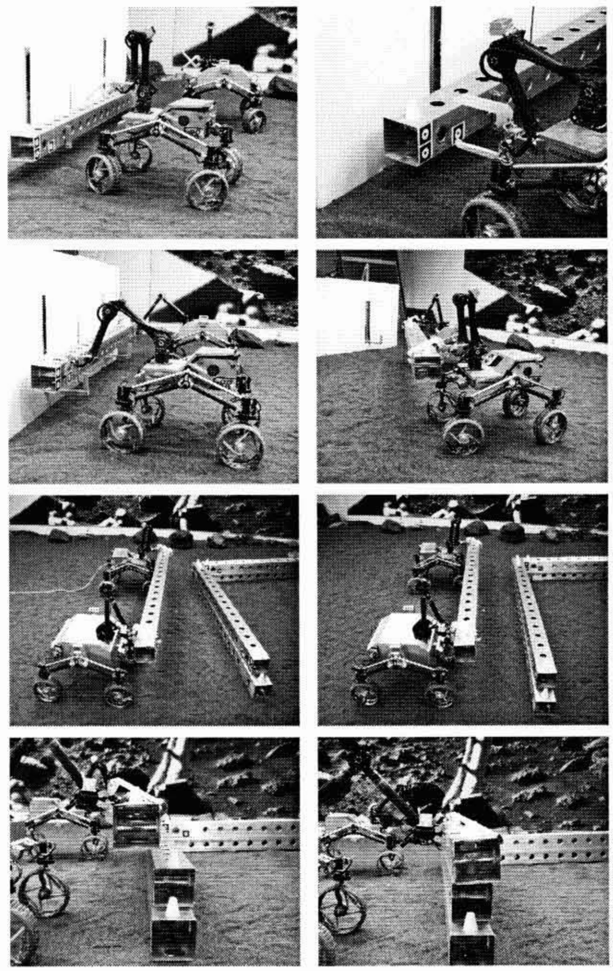

Fig. 10. *Top left*: Align at Storage brings rovers into grasping position. *Top Right*: Align Grasp places hand in grasping position. *Second Left*: Team lifts the component out of storage in Grasp Component. *Second Right*: Team turns away from storage in Clear Storage. *Third Left*: Rovers begin Align 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

Type	Runs	Failures
CG	50	1
CP	50	0
EE	10	0

The team is able to successfully complete the construction tasks with a very low failure rate. Means of addressing these remaining failure conditions are discussed in VII, Future Work.

VII. Future Work

Two primary areas of research are in current development for future addition to the Robotic Construction Crew: robustness and expanding the current task. To eliminate failures of manipulator positioning, force-torque feedback will be used. A misalignment during grasp will result in the finger pushing the component rather than entering the grasp point, which will increase forces on the manipulator. Similarly, when lowering a misaligned beam, forces will increase. These increased forces are detected by the force-torque sensor, and corrective action can be introduced, such as moving the manipulator to reduce forces to nominal. Manipulator placement may additionally be improved by observing the manipulator in the images to confirm its position.

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 distances of traverse. 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. Reliability of performance 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.

Acknowledgment

References

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

[2] 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 Trans. Systems, Man & Cybernetics, Part A: Systems and Humans, Collective Intelligence Special Issue, 33(5): 550-559, 2003.

[3] NASA. "Vision for Space Exploration." February2004.

[4] NASA. "2005 NASA Strategic Planning." 2005.

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

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

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

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

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

[10] 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.

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

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