

Action Selection Within the Context of a Robotic Colony

Terry Huntsberger^a, Maja Mataric^b, and Paolo Pirjanian^b

^aScience and Technology Development Section, Jet Propulsion Laboratory,
California Institute of Technology, Pasadena, CA 91109

^bRobotics Research Lab, Institute for Robotics and Intelligent Systems
University of Southern California, Los Angeles, CA

ABSTRACT

Preparation of planetary surface sites prior to a manned mission can be accomplished through the use of a robotic colony. The tasks of such a colony would include habitat deployment, setup of in-situ fuel and oxygen production plants, and beacons road placement. The colony will have to possess a great deal of autonomy for this ambitious list. BISMARC (Biologically Inspired System for Map-based Autonomous Rover Control) is a behavior based system for the control of multiple rovers on planetary surfaces. During the past few years the system has performed well in multiple cache retrieval simulations, and a certain degree of fault tolerance has been included in the design. In this paper we address the extensions to BISMARC that would be necessary for a robotic colony application. These extensions include a wider array of behaviors, better communication and mapping capabilities, and fault tolerance shared by the colony. The results of some simulations for habitat site preparation are reported.

Keywords: Robotic colonies, planetary rovers, behavior-based control

1. INTRODUCTION

Future manned Mars habitation will require some measure of robotic infrastructure support. This support could be in the form of inspection/maintenance of an existing structure, or on a more ambitious level, the use of robotic precursor missions to set up the habitat. Tasks of such precursor missions might include habitat site preparation and deployment, setup of in-situ fuel and oxygen production plants, and beacons road placement. The site preparation task shown in Figure 1 shares many of the elements of the cooperative box-pushing task.




Figure 1. Beacons delimited site for autonomous clearing operation. Site size is 100 m×50 m.

Cooperative robotic box-pushing has been the subject of numerous recent research studies.¹⁻⁹ A good biological analog for this task is prey or food transport performed by social insects such as ants.^{10,11,4,5} If the steps undertaken

Send correspondence to T. L. Huntsberger, E-mail: Terry.Huntsberger@jpl.nasa.gov.

for the prey transport operation can be duplicated within a robotic control system, an effective although not optimal strategy for cooperative box-pushing or site clearing can be realized. Such an approach was taken by Kube and Bonareau through randomized robot-box interactions in their recent study of the problem.⁵

The series of steps usually undertaken are:⁵

- A single ant will try to carry or drag the item with numerous realignment and regrasping attempts if necessary.
- If unsuccessful, the ant will contact some of its nestmates to assist in the task. These nestmates are recruited from both short range and long range distances, with preference given to the ants in the immediate vicinity.
- The group of ants will collectively start to tug on the item, rearranging themselves prior to actual pickup or dragging. They then cooperatively carry the item back to the nest, once again with realignment and regrasping operations being performed when needed.
- If the item is too large or heavy for all of the recruited nestmates, then ants with specialized mandibles break the item into smaller pieces.

The actual means of assessment of the need for realignment during the collective transport phase is not well understood, but it may be along the lines of the "cooperation without communication" modality. The ants might sense and react to the relative movement of the shared transport item, and thus interact through the environment.

As pointed out by Parker in her study of box-pushing using multiple heterogeneous robots, the problem of action recognition in other robots is extremely difficult using current sensing modalities.¹² This would seem to indicate that explicit communication is needed between robots performing a cooperative task. An alternate approach uses feedback from an object being manipulated, such as in the box-pushing task, to minimize or totally eliminate the need for this communication.^{13-15,4,5} The study by Matarić, et al used a simple communication strategy to enforce turn-taking behavior in legged robots for the box-pushing task.⁶ This effectively neutralized the problems of deadlock and stagnation. Staying within the limited bandwidth of planetary rover communication channels, communication along the lines of the simple broadcast modality of Parker¹² will be used for the recruitment subtask. We use the control strategy of Kube and Bonabeau^{16,5} implemented with temporal penalties¹⁷ to address the realignment/reposition for deadlock and stagnation resolution. This direction was taken in order to eliminate explicit dependence on communication for successful task completion.

We previously developed a control system for autonomous planetary rover control called BISMARC (**B**iolegally **I**nspired **S**ystem for **M**ap-based **A**utonomous **R**over **C**ontrol),^{18,19} that uses a free flow hierarchy (FFH)²⁰ architecture for action selection. The sensor models are based on those used for the Sample Return Rover (SRR) and the Field Integrated Design and Operations (FIDO) rover at the Jet Propulsion Lab.²¹ The system has been used for 800 simulations of a multiple rover, multiple cache recovery operation in a Mars environment with a mission success rate of 98.9%. The cooperative aspects of these missions were limited to the sharing of compressed maps for navigation to the cache sites. BISMARC is currently being ported to SRR for field testing in the Planetary Robotics Lab at JPL.

BISMARC possesses some degree of built-in fault tolerance in the action selection mechanism, and in addition, it handles rover sub-component failures through adaptation of the weights in the FFH.²² Short term memory mechanisms are used for the fault detection process and sensory perception modification is used for fault tolerance. In some sense, the modification can be viewed as a type of learning behavior. This fault tolerant behavior is extremely important for the long duration robot colony missions.

The next section describes the overall organization of the BISMARC architecture. This is followed by a discussion of the anticipated tasks and robotic capabilities needed for a robot colony, including the necessary extensions necessary to BISMARC. In particular, we will concentrate on the site preparation task as an example that includes most of these extensions. The results of some robot colony surface clearing simulations are described next, followed by a final summary section.

2. BISMARC ORGANIZATION

BISMARC is an example of a behavior-based robotic control architecture. A comprehensive review of such architectures can be found in the monograph by Arkin.²³ Sensor inputs are translated into actions through the use of a

process called an action selection mechanism (ASM). A recent review of ASMs can be found in Pirjanian.^{24,25} At the heart of BISMARC is a FFH ASM modeled after that of Rosenblatt and Payton²⁰ with modifications suggested by Tyrrell.¹⁷

The ASM used by BISMARC for the studies reported in this paper is shown in Figure 2. Action nodes are drawn as rectangles, stimulus nodes as ellipses, and those with multi-directional characteristics are indicated using 8 directional bins. The combination rules are additive for a small filled rectangle above the node, multiplicative for a small filled triangle, and a more sophisticated rule is used for plain rectangular nodes.¹⁷ This more sophisticated combination rule was developed by Tyrrell to guarantee the proper transfer of goal and motivational behavior to lower levels of the FFH.

Tyrrell introduced the temporal penalty (T-circle in Figure 2) to control action that will take an inordinate amount of time to complete.¹⁷ The temporal penalty is derived using the assigned value raised to the power of the elapsed time during the current action. Temporal penalty nodes increased the likelihood of satisfying the overall mission goal of totally clearing a designated area. In addition, the uncertainty penalty (U-circle in Figure 2) is used to control actions that are heavily dependent on external sensor inputs, which are usually noisy and imprecise.




Figure 2. FFH system for BISMARC. All weights on the arcs are 1.0 unless otherwise indicated. Notation for symbols is that of Tyrrell.¹⁷ See text for a detailed discussion.

The top level behaviors in BISMARC generally relate to rover health (“Avoid Other Robot”, “Sleep at Night”, “Warm Up”, “Cool Down”, “Get Power”) or the clearing operation (“Scan for Rock”, “Clear Area”, “Keep Variance Low”). These high level behaviors involve a complicated combination of internal control and assimilation of external sensor inputs. Since potentially conflicting behaviors can arise, the FFH offers a better approach to the low level control problem than direct elimination of lower nodes due to inhibition as found in purely reactive control.²⁶ There are only six bottom level actions: “Sleep”, “Move”, “Broadcast”, “Respond to Call”, “Look Around”, and “Rest”. There is some selfishness built into the behaviors in that the weight on the “Respond to Call” node is negative (dozer will be slightly hesitant to stop current behavior).

3. ROBOTIC COLONY OPERATIONS

For the purposes of brevity, this paper will concentrate on the robotic colony operations that are involved in site clearing. A description of a broader class of robotic colony operations can be found in Huntsberger and Rodriguez.²⁷

Site clearing can be accomplished by a solitary robot if all of the significant rocks are within the size and mass constraints that the robot is able to handle. Rocks that are outside these limits will need to be cleared using a cooperative multiple robot strategy. The issues that must be addressed before such a strategy can be implemented are:

- Robot behavior differences in the solitary vs. group surface clearing task.
- Method for solitary robot to decide when portion of task is not doable alone.
- Communication protocol for recruitment.
- Coordination of multiple robots for clearing of difficult terrain areas.
- Determination of correct number of robots for clearing task at hand.
- Method for solving deadlock/stagnation problem.

Each of these items are relatively easy to incorporate into a behavior-based control system. The following subsections discuss each of them in turn.

3.1. Behavior

The possible methods that might be used for solitary vs. group behavior action selection under BISMARC would include the “behavior sets” of the ALLIANCE architecture,^{28,8,12} the direct and/or temporal sequencing of “basic behaviors” in the behavior-based system of Mataric,²⁹⁻³⁴ or the schemas of Arkin.^{13,14,35,23,36} BISMARC uses behavioral cues to trigger the appropriate behavior for the current situation. Behaviors such as “Avoid other Robot” are already built into the hierarchy and invoked automatically when activation of that behavior becomes the dominant input to the arbitration node. For example, weights on the links between levels of the hierarchy give the “Avoid Other Robot” behavior precedence over the “Approach Perceived Rock” behavior.

3.2. Pushing decision

Two sensor inputs that can be used by an earth-moving robot to ascertain its inability to complete the task are the strain gauges on the dozer blade and the wheel odometry. If the wheel odometry indicates a stall mode, then the robot will not be able to move the object. On the other hand, the robot can go into a slip mode, where the odometry data still indicates forward motion. In both cases, the temporal profile of the strain gauges on the dozer blade will give the necessary information for deciding whether to recruit help. If the object is being pushed, the strain gauges will stay at a relatively stable level related to the ground friction. Potential stagnation is detected when the strain gauge inputs increase to a level indicating potential damage to the dozer blade. BISMARC uses the FFH to include temperature and battery level sensors in the action selection process. This information is also combined with relative beacon localization as detailed in the next subsection.

3.3. Communication

If the single robot decides it can't clear a section of ground either at the start of the sequence or midway through, a communication strategy for recruitment needs to be invoked. BISMARC assumes that the boundaries for the clearing site have been previously delimited by radio beacon placement. These beacons are used by the individual robots for their position determination within the site. In the event that a robot has to recruit help, the broadcast communication includes this position information in order for the individual robots to respond based on relative closeness to the robot in need. The first robot to determine that it is within the range for response uses a simple broadcast to notify the other robots that it has accepted the recruitment call, thus allowing the other robots to get on with their current task. We will investigate the potential problems that can arise from the use of such potentially inaccurate information in the experimental studies section.

3.4. Number to recruit

In the case of a swarm model of this process, there is a large number of candidates for recruitment. On the other hand, the number of individuals in a robot colony will necessarily be limited due to mass and power constraints. As such, BISMARC takes a conservative strategy and only tries to recruit a single robot each time it encounters a problem. The team of robots grows one at a time and uses the same methods as the solitary robot to determine the potential need for more help. This behavior has to be tempered with the deadlock/stagnation strategies detailed in the next subsection.

3.5. Deadlock/stagnation

The stagnation problem was explicitly addressed by Mataric, et al using a turn-taking strategy enforced with simple communication.⁶ In the case of more than two robots clearing rough, uneven terrain, this method will not be as effective. The study by Kube and Zhang¹⁶ and Kube and Bonabeau⁵ use small random angle of attack and position changes invoked within a temporal window to solve the deadlock/stagnation problem. BISMARC uses a temporal penalty¹⁷ input to the middle level pushing behavior to implement the random repositioning. Motion towards the goal is defined in terms of the radio beacons.

4. EXPERIMENTAL STUDY

We ran 50 trials each with colonies of from two to six dozers using a randomly generated heightfield. The area encompassed a 50×100 meter rectangle with a grid resolution of 5 cm. Each trial had different placement positions for the rocks, with a statistical profile of the Mars Pathfinder mission site used for mass and number of rocks. None of the rocks were allowed to mass over 175 kilograms. It was assumed that the clearing, staging, and rock pile areas were previously delimited by beacons as shown in Figure 1. Top speed on the dozers was set to 30 cm/sec, with a mass of 100 kilograms, and a size of $2 \times 2 \times 1$ meter (length, width, height). Power use on the dozers varied continuously from 30 watts when idle, 60 watts when traveling over open terrain, to 110 watts when involved in pushing the heaviest rock within the dozer's capability when alone. This capability was set to a maximum of 75 kilograms. In order to simulate wheel slippage, we set a 10% loss of traction when pushing a rock. The dozers were forced to sleep during the night hours of the simulations, since there were no infrared sensors. A collision between two dozers was considered as totally disabling to both, and a dozer that was hit by a rock being pushed by another dozer was also considered totally disabled. Each dozer had a forward-facing set of stereo cameras with a baseline of 25 cm, a spatial resolution of 486×512 pixels, and a 100 degree FOV, a transmitter/receiver with a bandwidth of 56 kbaud, and an 8 channel receiver for beacon monitoring.

Success in the simulation studies was measured in terms of total time for the task, and the number of dozers that were healthy at the end of the run. Figure 3 plots the total average time taken for the task versus the number of dozers that were allowed to participate up to a maximum of six. The graph starts at two dozers, since there was always at least one rock within the clearing area that was larger than a solitary dozer could move.

The graph only includes 48 trials since in two of trials with two dozers, there was a collision that disabled both of the dozers. These were the only trials that caused a total mission failure. The performance was not linear, indicating that there is significant interference between the dozers as the total number increases. This behavior manifested itself through more collisions with rocks being pushed by another dozer, more complicated repositioning operations due to the "avoid other robots" behavior, and travel time delays due to the need to maintain a safe distance during the recruitment phase. The rock-dozer collisions were caused by rocks dynamically approaching the about-to-be damaged dozer from an angle outside the range of the forward-facing stereo hazard avoidance cameras.

5. CONCLUSIONS

This paper extended the rover control system called BISMARC to include behaviors that would be necessary for a site clearing operation in support of a robotic colony. The FFH used for action selection in BISMARC maintained its structure and weights, but the sensor perception models were modified to account for the faults. The results of a total of 250 missions indicated the overall time needed to clear a given area decreased with an increase in the number of available dozers, but did so non-linearly. Collisions between rocks being pushed and dozers, and time delays in repositioning due to the close proximity of other dozers were the main contributing factors.

Although the FFH used in BISMARC was able to solve the control problem for the clearing task, the weights between levels are determined in an ad hoc manner. We are currently examining methods for the learning of these




Figure 3. Average total mission time for 48 simulation trials versus the number of dozers in each trial. Time is in arbitrary simulation units with an average time of 3 weeks, 4 days for two dozers.

weights based on desired behavior for the site clearing task. Ideally, the amount of communication, semi-random exploration of the clearing site, and interference between dozers needs to be minimized.

6. ACKNOWLEDGMENTS

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

1. P. Caloud, W. Choi, J. C. Latombe, C. L. Pape, and M. Yim, "Indoor automation with many mobile robots," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 67–72, 1990.
2. B. Donald, J. Jennings, and D. Rus, "Analyzing teams of cooperating mobile robots," in *IEEE International Conference on Robotics and Automation*, pp. 1896–1903, 1994.
3. C. R. Kube and H. Zhang, "Collective robotic intelligence," in *Second International Conference on Simulation of Adaptive Behavior: From Animals to Animats*, pp. 460–468, MIT Press, (Cambridge, MA), 1992.
4. C. R. Kube and H. Zhang, "Task modelling in collective robotics," *Autonomous Robots* **4**(1), pp. 53–72, 1997.
5. C. R. Kube and E. Bonabeau, "Cooperative transport by ants and robots," Tech. Rep. 99-01-008, Santa Fe Institute, Santa Fe, NM, 1999.
6. M. J. Matarić, M. Nilsson, and K. T. Simsarian, "Cooperative multi-robot box-pushing," in *IEEE International Conference on Intelligent Robots and Systems*, vol. 3, pp. 556–561, 1995.
7. F. R. Noreils, "An architecture for cooperative and autonomous mobile robots," in *IEEE International Conference on Robotics and Automation*, pp. 2703–2710, 1992.
8. L. E. Parker, "ALLIANCE: An architecture for fault tolerant, cooperative control of heterogeneous mobile robots," in *IEEE/RSJ/GI International Conference on Intelligent Robots and Systems*, pp. 776–783, 1994.
9. D. J. Stilwell and J. S. Bay, "Toward the development of a material transport system using swarms of ant-like robots," in *IEEE International Conference on Robotics and Automation*, pp. 766–771, 1993.
10. J. H. Sudd, "The transport of prey by an ant *pheidole crassindoa*," *Behaviour* **16**(3–4), pp. 295–308, 1960.
11. J. H. Sudd, "The transport of prey by ants," *Behaviour* **25**(3–4), pp. 234–271, 1965.
12. L. E. Parker, "Adaptive heterogeneous multi-robot teams," *Neurocomputing*, 1999. to appear.

13. R. C. Arkin, "Cooperation without communication: Multi-agent schema-based robot navigation," *Journal of Robotic Systems* **9**(3), pp. 351–364, 1992.
14. R. C. Arkin, T. Balch, and E. Nitz, "Communication of behavioral state in multi-agent retrieval tasks," in *IEEE International Conference on Robotics and Automation*, vol. 3, pp. 588–594, 1993.
15. C. R. Kube and H. Zhang, "Collective robotics: From social insects to robots," *Adaptive Behavior* **2**(2), pp. 189–219, 1993.
16. C. R. Kube and H. Zhang, "Stagnation recovery behaviors for collective robotics," in *IEEE/RSJ/GI International Conference on Intelligent Robots and Systems*, vol. 3, pp. 1883–1890, IEEE Computer Society Press, (Los Alamitos, CA), 1994.
17. T. Tyrrell, "The use of hierarchies for action selection," *Journal of Adaptive Behavior* **1**(4), 1993.
18. T. L. Huntsberger, "Autonomous multirover system for complex planetary retrieval operations," in *Proc. SPIE Symposium on Sensor Fusion and Decentralized Control in Autonomous Robotic Systems*, P. S. Schenker and G. T. McKee, eds., pp. 221–227, (Pittsburgh, PA), Oct. 1997.
19. T. Huntsberger and J. Rose, "BISMARC: A Biologically Inspired System for Map-based Autonomous Rover Control," *Neural Networks* **11**(7/8), pp. 1497–1510, 1998.
20. J. K. Rosenblatt and D. W. Payton, "A fine-grained alternative to the subsumption architecture for mobile robot control," in *Proc. IEEE/INNS Joint Conf. on Neural Networks*, pp. 317–324, (Washington, DC), June 1989.
21. T. L. Huntsberger, T. Kubota, and J. Rose, "Integrated vision/control system for autonomous planetary rovers," in *IAPR Workshop on Machine Vision in Applications*, (Chiba, Japan), Nov. 1998.
22. T. L. Huntsberger, "Fault-tolerant action selection for planetary rover control," in *Proc. SPIE Symposium on Sensor Fusion and Decentralized Control in Robotic Systems*, P. S. Schenker and G. T. McKee, eds., pp. 150–156, (Boston, MA), Nov. 1998.
23. R. C. Arkin, *Behavior-Based Robotics*, MIT Press, Cambridge, MA, 1998.
24. P. Pirjanian, "Satisficing action selection," in *Proc. SPIE Symposium on Sensor Fusion and Decentralized Control in Robotic Systems*, P. S. Schenker and G. T. McKee, eds., pp. 157–168, (Boston, MA), Nov. 1998.
25. P. Pirjanian, *Multiple Objective Action Selection and Behavior Fusion Using Voting*. PhD thesis, Laboratory of Image Analysis, Department of Medical Informatics and Image Analysis, Aalborg University, Denmark, 1998.
26. R. A. Brooks, "A robust layered control system for a mobile robot," *IEEE Journal of Robotics and Automation* **2**(1), pp. 14–23, 1986.
27. T. Huntsberger and G. Rodriguez, "Robotics challenges for human and robotic Mars exploration," in *Space 2000*, (Albuquerque, NM), Mar. 2000. to appear.
28. L. E. Parker, "Heterogeneous multi-robot cooperation," Tech. Rep. 1865, MIT, Cambridge, MA, 1994.
29. M. J. Matarić, "Integration of representation into goal-driven behavior-based robots," *IEEE Transactions on Robotics and Automation* **8**(3), pp. 304–312, 1992.
30. M. J. Matarić, "Distributed approaches to behavior control," in *Proc. SPIE Sensor Fusion V*, P. S. Schenker, ed., vol. 1828, pp. 373–382, (Boston, MA), Nov. 1992.
31. M. J. Matarić, "Issues and approaches in the design of collective autonomous agents," *Robotics and Autonomous Systems* **16**, pp. 321–331, Dec. 1995.
32. M. J. Matarić, "Designing and understanding adaptive group behavior," *Adaptive Behavior* **4**, pp. 51–80, Dec. 1995.
33. M. J. Matarić, "Behavior-based control: Examples from navigation, learning, and group behavior," *Journal of Experimental and Theoretical Artificial Intelligence, Special Issue on Software Architectures for Physical Agents* **9**(2–3), pp. 323–336, 1997.
34. M. J. Matarić, "Behavior-based robotics as a tool for synthesis of artificial behavior and analysis of natural behavior," *Trends in Cognitive Science* **2**, pp. 82–87, Mar. 1998.
35. R. C. Arkin and D. McKenzie, "Temporal coordination of perceptual algorithms for mobile robot navigation," *IEEE Trans. on Robotics and Automation* **10**(3), pp. 276–286, 1994.
36. D. C. MacKenzie, R. C. Arkin, and J. M. Cameron, "Multiagent mission specification and execution," *Autonomous Robots* **4**(1), pp. 29–52, 1997.