



Multi-Robot Task Allocation in Uncertain Environments

MAJA J. MATARIĆ, GAURAV S. SUKHATME AND ESSEN H. ØSTERGAARD*

Robotics Research Laboratory, Department of Computer Science, University of Southern California,
Los Angeles, CA 90089-0781, USA

mataric@cs.usc.edu

gaurav@usc.edu

esben@mip.sdu.dk

Abstract. Multiple cooperating robots hold the promise of improved performance and increased fault tolerance for large-scale problems such as planetary survey and habitat construction. Multi-robot coordination, however, is a complex problem. We cast this problem in the framework of multi-robot dynamic task allocation under uncertainty. We then describe an empirical study that sought general guidelines for task allocation strategies in multi-robot systems. We identify four distinct task allocation strategies, and demonstrate them in two versions of the multi-robot emergency handling task. We describe an experimental setup to compare results obtained from a simulated grid world to those obtained from physical mobile robot experiments. Data resulting from eight hours of experiments with multiple mobile robots are compared to the trend identified in simulation. The data from the simulations show that there is no single strategy that produces best performance in all cases, and that the best task allocation strategy changes as a function of the noise in the system. This result is significant, and shows the need for further investigation of task allocation strategies and their application to planetary exploration.

Keywords: task allocation, multiple robots

1. Introduction

Following the success of the Sojourner robot rover (Mishkin et al., 1998), it is likely that future missions to Mars and other planets will rely on mobile robot explorers. In particular, *multiple cooperating robots* hold the promise of improved performance and increased fault tolerance for large-scale problems such as planetary survey and habitat construction (Schenker et al., 2000; Huntsberger et al., 2001). NASA already has tentative plans for multi-robot missions (NASA/JPL, 2002) in the coming decades.

Fielding multiple robots for space exploration is a daunting challenge (Pirjanian et al., 2000), since basic research in the area is still addressing fundamen-

tal problems that are currently unsolved in multi-robot system design and coordination. There has been significant prior research in multi-robot coordination (Arkin, 1992; Mataric, 1995; Balch and Arkin, 1998; Parker, 1998; Cao et al., 1995; Werger and Mataric, 2000; Gerkey and Mataric, 2002a). However, the general problem of *dynamically allocating tasks in a group of multiple robots satisfying multiple goals*, is as yet unsolved. We view this problem as an instance of *dynamic task allocation* under uncertainty.

Presently, there is no general theory of task allocation in uncertain multi-robot domains, but there is work addressing this difficult problem (Gerkey and Mataric, 2002b). In this paper, we empirically derive some guidelines for selecting task allocation strategies for multi-robot systems. The explored strategies are individualistic in that they do not involve explicit cooperation and negotiation among the robots. However, they are part of a large class of approaches that produce

*Current address: The Maersk McKinney Moller Institute for Production Technology, University of Southern Denmark, DK-5230 Odense, Denmark.

coherent and efficient cooperative behavior without explicit coordination.

Given the empirical nature of this work and the scope of the problem addressed, these guidelines are necessarily incomplete, though they provide useful insight. We demonstrate that the choice of task allocation strategy is far from trivial. We also empirically show that no optimal task allocation strategy exists for all domains, and that it can be very difficult to identify the optimal task allocation strategy even for a particular task.

These results are derived through the use of a framework developed for understanding the task allocation problem, which illustrates a common approach to decomposing the problem. Using this framework, we compare four distinct task allocation strategies, in both grid world and real world task allocation experiments, applied to the emergency handling problem domain. We compare the grid world and real world results.

This paper is organized as follows. Section 2 formally presents the problem of dynamic task allocation under uncertainty. Section 3 then applies the framework to present four task allocation strategies chosen to cover a relevant part of the multi-robot coordination parameter space. Section 4 describes our experimental validation of the framework; it presents the emergency handling task and two experimental domains: the simulation grid world and the team of real robots. Section 5 presents the experimental results from the two domains. Section 6 compares and discusses the common trends among the two experimental testbeds. Section 7 concludes the paper.

2. Problem Statement

In the context of multi-robot coordination, dynamic task allocation can be viewed as the selection of appropriate actions (Maes, 1994) for each robot at each point in time so as to achieve the completion of the global task by the team as a whole. From a global perspective, in multi-robot coordination, action selection is based on the mapping from the combined robot state space to the combined robot action space. For homogeneous robots, it is the mapping

$$S^{|R|} \rightarrow A^{|R|}$$

where S is the state space of a robot, $|R|$ is the number of robots, and A is the set of actions available to a robot (Matarić, 1994). In practice, even with a small number of robots, this is an extremely high-dimensional map-

ping, a key motivation for decomposing and distributing control.

Based on the approach introduced in Gerkey and Matarić (2001), we decompose the task allocation problem into the following three steps:

1. each robot bids on a task based on its perceived fitness to perform the task;
2. an auctioning mechanism decides which robot gets the task;
3. the winning robot's controller performs one or more actions to execute the task.

We use the above decomposition to construct a general formulation for the multi-robot coordination problem. In this formulation, a *bidding function* determines each robot's ability to perform a task based on that robot's state. Next, the *task allocation mechanism* determines which robot should perform a particular task based on the bids. Finally, the robot controllers determine appropriate actions for each robot, based on the robot's current task engagement. This partitioning, illustrated in Fig. 1, serves two purposes: it reduces the dimensionality of the coordination problem, and it reduces the amount of inter-robot communication required. Instead of mapping

$$S^{|R|} \rightarrow A^{|R|}$$

we now have the mapping

$$B^{|R||T|} \rightarrow T^{|R|}$$

namely from all robots' bids B for all tasks T to a task assignment for each robot. We call this overall mapping the *Task Allocation Strategy* for the system as a whole. We treat the overall mapping here as a global, centralized process (as depicted in Fig. 1), but distributed auctioning mechanisms (Dias and Stentz, 2000; Gerkey and Matarić, 2001), blackboard algorithms (Corkill, 1991), and cross-inhibition of behaviors (Werger and Matarić, 2000) are some validated methods for distributing the task allocation function. In this paper, we focus on what the task allocation function should be, rather than on how it should be distributed.

The above framework is a general way that dynamic task allocation for multi-robot systems can be formulated. Thus, various prior approaches to multi-robot coordination can be described in the context of this framework. For example, Parker's ALLIANCE architecture (Parker, 1998) determines each robot's ability to perform a task and maps it to a scalar quantity which

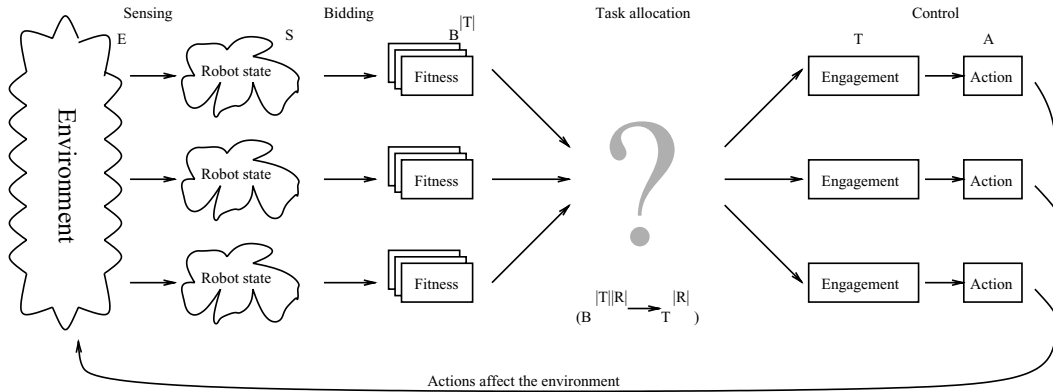


Figure 1. Reducing dimensionality of multi-robot coordination.

is used to assign tasks to robots. In Werger’s BLE approach (Werger and Mataric, 2000), a local eligibility mechanism is used as the robots’ perceived ability to perform a task, i.e., its fitness, and the best robot (computed through a port-based max function) “wins” at each time-step. Continued work by Gerkey and Mataric (2002a), on which this formulation is based, applies and validates this decomposition in different task domains.

3. Four Task Allocation Strategies

The dynamic task allocation problem, i.e., the mapping from bids to tasks, can be performed in numerous ways. We limit our discussion here to Markovian systems, where the task allocation mapping for a given robot is based on the mapping between that robot’s current task assignments and every other robot’s current bid on each task, to the given robot’s new task assignment, as shown in Fig. 2. The problem, then, is: *given each robot’s bid on each task and each robot’s current task engagement, what should each robot’s new task assignment be?*

We focused on exploring the effects of two key aspects of distributed control, *commitment* and *coordination*, have on performance. Given the large space of possibilities, we considered only the extreme cases of each: no commitment and full commitment, and no coordination and full coordination. The combination of these extremes resulted in four task allocation strategies (see Fig. 3). Along the commitment axis, a fully committed strategy meant a robot would complete its assigned task before considering any new engagements, while a fully opportunistic strategy allowed a robot to drop an ongoing engagement at any time in favor of a new one. Along the coordination axis,

Cur. eng.	Bids	A	B	C	D	New eng.
A	R.1	6	4	2	5	?
-	R.2	4	1	0	3	?
C	R.3	7	2	3	2	?

Figure 2. An example task allocation scenario.

Strategies	Coordination		
Commitment	Individ. Mut. Excl.		
	Commit.	Str. 1	Str. 2
	Oppor.	Str. 3	Str. 4

Figure 3. The four task allocation strategies considered are set up as combinations of two variables, the amount of commitment, and the amount of coordination.

the uncoordinated (individualistic) strategy meant each robot performed based on its local information, while a coordinated strategy simply implemented mutual exclusion, so only one robot could be assigned to a task, and no redundancies were allowed. We note that this notion of “coordination” is simple, and not intended to represent explicit cooperation and coordination strategies being explored in other work. Our tasks were structured so that one robot was sufficient for completion of an individual task assignment. Thus, mutual exclusion was the simplest yet effective form of coordination. Figure 2 shows the table that results from listing each robot’s current engagement and each robot’s current bid on each task.

As an example, the fully committed mutually exclusive strategy, one of the four described above, is as follows:

1. If a robot is currently engaged in a task, and its bid on that task is greater than zero, remove the row and column of the bid from the table, and set the robot's new assignment to its current one.
2. Find the highest bid in the remaining table. Assign the corresponding robot to the corresponding task. Remove the row and column of the bid from the table.
3. Repeat from step 2 until there are no more bids.

In case of individualistic (uncoordinated) strategies, the same algorithm is run on a separate table for each robot. In the opportunistic (uncommitted) case, step 1 above is skipped.

4. Experimental Validation

4.1. The Task

In order to study and compare the task allocation strategies described above, we devised a task domain that can be used in simulation and in an indoor building setting, and is also relevant to real-world problems, including those found in space exploration.

We used *emergency handling* (Østergaard et al., 2001) as our problem task domain for evaluation. In it, robots roam around a planar environment, in which *alarms* occur at unpredictable times and in unpredictable locations. The task of the robot team is to detect alarms and “fix problems” indicated by those alarms. There is a variable time-cost associated with traveling to an alarm, depending on the robot's speed and the distance to the alarm. There is also a fixed time-cost for fixing the alarm. In the implementation presented here, we restricted ourselves to the case where any robot can fix any alarm.

This task domain can be generalized to a variety of real-world multi-robot scenarios. For example, alarms can correspond to new incoming goals (e.g., “get that rock”), as well as to failures or cries for help by one or more members of the robot team.

4.2. Grid World Experimental Setup

We implemented a simplified version of the above-described multi-robot emergency handling task in a grid world, as illustrated in Fig. 4, in order to conduct large numbers of experiments that are practically impossible with physical robots.

As the base case of the grid world implementation, we considered a 10×10 grid inhabited by 10 “robots”.




Figure 4. An example 10×10 grid world with four robots and three active alarms.

Robots bid on alarms depending on their distance to those alarms. The bid was set to $20 - d$, where d is the Manhattan distance to the alarm. In each time-step, any robot assigned to a particular alarm moved toward that alarm. When a robot arrived at an alarm, that alarm was instantly put out (i.e., the fixed time-cost was 0). Three new alarms appeared every twelve time-steps at random positions on the grid.

4.3. Physical Experimental Setup

We also implemented the multi-robot emergency handling task in a physical multi-robot testbed, in an indoor office building setting. In our experiments with mobile robots, we used ActivMedia Pioneer 2 DX mobile robots, equipped with 233 MHz Linux PCs, SICK laser range finders, cameras, wireless Ethernet, speakers, and microphones, as shown in Fig. 5. The microphones were made directional by placing them at the bottom of two Styrofoam cups. All control of the robots was done through *Player* (Gerkey et al., 2001), a server and protocol that connects robots, sensors, and control programs through a standard TCP socket. *Player* was developed jointly at the USC Robotics Research Lab and HRL Labs and is freely available under the GNU Public License from <http://playerstage.sourceforge.net>.

In the physical experiments, alarms were speakers placed in the environment (Fig. 6), marked with



Figure 5. Left: A fully equipped Pioneer robot. Right: Close-up of the sensors. The microphone is glued to the bottom of two Styrofoam cups to add directionality.




Figure 6. The environment used. A–D are alarm positions, 1–3 are robot start positions.

brightly colored paper. Each alarm emitted a tone with a unique frequency, which could be detected by each robot. The robots' bids were proportional to the intensity with which the frequency was received. Due to sensor uncertainty and the unknown structure of the environment, the robots could not accurately estimate their distance to the alarms. Instead, they used the perceived absolute alarm intensity to decide which robot would win the bid. When assigned to an alarm, a robot followed its frequency until it visually acquired the brightly colored paper. From that point, the robot relied on visual servoing to approach the alarm. The controller that servoed the robot to the sound source consisted of a repeated two-step process:

1. Make a 360° scan for the frequency corresponding to the robot's engagement,
2. Go forward in the direction of highest intensity of that frequency, until a junction or dead end is detected.

When a robot was close enough to an alarm of the appropriate frequency, it emitted a counter sound, a tone

of half the frequency of the alarm. Thirty seconds after the counter sound was emitted, the alarm turned off (i.e., the fixed time-cost was 30 s).

Robots' motors were controlled by a weighted average between the output from the above described sound-servoing controller, and an obstacle avoidance controller which prevented collisions with objects in the environment. The relative weight of the collision avoidance input was scaled by the distance to perceived obstacles. Further details of the physical setup are found in Østergaard et al. (2001).

5. Experimental Results

This section describes the experimental results obtained in simulation and with the real robots, and compares the two. In both the simulated and physical experiments we measured performance as the sum of the number of active alarms at each time-step; lower scores thus indicate better performance.

5.1. Grid World Experimental Results

Figure 7 shows the results obtained after executing the four task allocation algorithms for 1000 time-steps. These base case data indicate that the combination of mutual exclusion and opportunism produces the best performance. However, this was not the case if key parameters, such as world size, number of robots, noise, and distribution of alarms, were varied. In general, when these parameters were varied, any of the four task allocation strategies we tested could outperform the rest.

We focused our further analysis on noise and uncertainty, since they are key issues in real world robot systems. To simply model sensor noise, we added a random number (from a normal distribution) to the bid of each robot. Actuator uncertainty was modeled by introducing a finite, but small, probability that robots would move in a random direction instead of the intended direction (towards the alarm) at each time-step.

Strategy:	I, O	I, C	M, O	M, C
Results:	980	1045	435	722

Figure 7. Results from the base case grid world run showing "alarm on-time" (lower is better) for the four task allocation strategies. The strategies are obtained by crossing individualism (I) and mutual exclusion (M) with opportunism (O) and commitment (C).




Figure 8. A cut through the parameter space of the grid world multi-robot emergency handling task. The graph shows the best performing strategy for each of 50×50 settings, obtained by varying "sensor noise" and "actuator noise" parameters.

Figure 8 shows the results from varying the sensor noise and actuator uncertainty parameters. The X-axis shows σ for the "sensor noise" varied from 0 to 20. The Y-axis shows "actuator noise" varied from 0% to 50%. As can be seen, for low amounts of noise, mutual exclusion and opportunism performed best, while for larger amounts of noise, commitment and individualism performed best.

5.2. Physical Experimental Results

To validate the simulation results, we performed two sets of experiments with the physical robots. The first used the same base case scenario as in the simulation experiments to evaluate the performance of each of the four task allocation strategies. Then, in a second set of experiments we reduced the sensor noise and actuator uncertainty of the real world system. To achieve this, we placed laser landmarks in the corridors, and provided the robot with a model of the environment. These modifications allowed us to improve the sound-based navigation, resulting in robots almost always turning correctly when servoing on sound. In addition, each robot's perceived distance to an alarm was significantly less noisy, resulting in an increased accuracy of bids.




Figure 9. The two real world setups are shown in relative position to each other on the noise axes. We conjecture that the grid world results correspond to the upper right part of the graph.

These two experiments correspond to two points in the space shown in Fig. 8. The first setup has higher sensor noise and higher actuator uncertainty than the second, as illustrated in Fig. 9. Our hypothesis was that different strategies would perform best in the two setups.

We performed 6 runs for each task allocation strategy for each setup, totaling $6 \times 4 \times 2 = 48$ runs. The results are shown in Fig. 10. In addition, an applet showing a visualization of the experiments is found at <http://robotics.usc.edu/applets/taskalloc>. As expected, the average scores for the reduced noise case are lower than those for the base case, showing that the modifications improved the performance of the system.

The task performance results show that in both real world setups, the opportunistic mutually excluding strategy performed best. This suggests that both setups are in a noise regime corresponding to the lower left part of the space shown in Fig. 8, where mutual exclusion and opportunism are the best alternatives. By extrapolating these tendencies from the two experiments, we can hypothesize that if the noise in the system increases sufficiently, the committed individualistic strategy would perform best, as is the case in the grid world. This extrapolation is shown in Fig. 11. However, since there is a large amount of stochasticity in the measured data, we cannot be certain whether the tendency we found is permanent or transient. The tendency is unfortunately not statistically significant in our data sample.

6. Discussion

It is important to know whether the tendencies derived from the grid world simulation apply to the real world

<i>Base case</i>				
Strategy	I, O	I, C	M, O	M, C
1	1173	1158	1258	1132
2	1567	1297	1172	1074
3	976	708	1238	1014
4	1338	1010	790	1280
5	790	1464	883	1308
6	774	992	996	1016
μ	1103.00	1104.83	1056.17	1137.33
σ	315.52	263.80	195.79	129.19

<i>Reduced noise</i>				
Strategy	I, O	I, C	M, O	M, C
1	1089	892	758	1041
2	977	884	1005	793
3	1094	917	803	862
4	913	797	842	924
5	522	1338	888	901
6	1120	1265	808	876
μ	952.50	1015.50	850.67	899.50
σ	225.48	226.38	87.14	82.37

Figure 10. Quantitative results for the four cases given by combining Individualism (I) or Mutual Exclusion (M) with Opportunism (O) or Commitment (C). The numbers are the sum of on-time for all alarms in each trial given in seconds. Lower is better.

(i.e., the indoor environment with physical robots) in our problem domain. The results from the real world experiments imply that the noise levels correspond to the regime of the lower left of the space shown in Fig. 8, in that the combination of mutual exclusion and opportunism was the best performing strategy. Further experiments could determine whether the correspondence between the grid world and the real world results is a coincidence or a systematic trend.

The grid world results are interesting if they actually represent real world system behavior. The fact that the




Figure 11. Extrapolation of tendencies. Mean values for the two setups, connected with lines.

best performing task allocation strategy changes as we vary noise parameters in the grid world implies that it can be very difficult to decide a priori which task allocation strategy should be used in a given task for any real world implementation. From the grid world results, it seems that the benefit from mutual exclusion is dependent on the total noise in the system, while the benefit of commitment seems to be dependent on the ratio between “actuator noise” and “sensor noise”. Part of this trend is also acknowledged and utilized in Goldberg and Mataric (2002), for a controller that, when noise is increased, degrades gracefully from a mutually excluding to an individualistic strategy.

The four task allocation strategies we examined are *extreme*, in that they take into consideration only the complete presence or absence of commitment and coordination in the given context. Arguably, the best strategy for any particular task would most likely be a carefully-balanced compromise. However, as stated previously, the goal of this work was not to attempt to find the best strategy (which is necessarily task- and parameter-specific), but rather to gain some insight into task allocation in general. The experimental data indicate that the four strategies we explored provide a reasonable span of strategy space and provide leading insights for further study.

7. Conclusion

We have described an empirical study that sought general guidelines for task allocation strategies in systems of multiple cooperating robots. We identified

four distinct task allocation strategies, aimed at studying tradeoffs between commitment and coordination, and demonstrated them in two versions of the multi-robot emergency handling task. We described an experimental setup to compare results obtained from a simulated grid world to the results from real robot experiments. Data resulting from eight hours of real mobile robot experiments are compared to the trend identified in simulation. The data from the simulations show that there is no single strategy that produces best performance in all cases, and that the best task allocation strategy changes as a function of the noise in the system. This result is significant, and shows the need for further investigation of task allocation strategies.

The described work is a small step toward the larger goal of principled analysis and synthesis of multi-robot coordination strategies for complex and uncertain domains, such as space exploration. The task domain of emergency handling we used as the context for experimental validation was chosen because of its potential relevance to real-world distributed robotics terrestrial and space applications.

Acknowledgments

This work was supported in part by DARPA grant DABT63-99-1-0015 under the Mobile Autonomous Robot Software (MARS) program, and in part by the ONR DURIP grant N00014-00-1-0638 and NSF grants ANI-9979457 and ANI-0082498 under the Special Projects in Networking Program.

References

- Arkin, R.C. 1992. Cooperation without communication: Multiagent schema based robot navigation. *Journal of Robotic Systems*, 9(3):351–364.
- Balch, T. and Arkin, R. 1998. Behavior-based formation control for multi-robot teams. *IEEE Transactions on Robotics and Automation*, 14(6):1–15.
- Cao, Y., Fukunaga, A., Kahng, A., and Meng, F. 1995. Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots*, 4(1):7–27.
- Corkill, D.D. 1991. Blackboard systems. *AI Expert*, 6(9):40–47.
- Dias, M.B. and Stentz, A.T. 2000. A free market architecture for distributed control of a multirobot system. In *6th International Conference on Intelligent Autonomous Systems (IAS-6)*, pp. 115–122.
- Gerkey, B. and Matarić, M.J. 2001. Principled communication for dynamic multi-robot task allocation. In *Experimental Robotics VII, LNCIS 271*, D. Rus and S. Singh (Eds.), Springer-Verlag: Berlin, pp. 353–362.
- Gerkey, B. and Matarić, M.J. 2002a. Pusher-watcher: An approach to fault-tolerant tightly-coupled robot coordination. In *Proceedings, IEEE International Conference on Robotics and Automation*, Washington DC.
- Gerkey, B.P. and Matarić, M.J. 2002b. Multi-robot task allocation: Analyzing the complexity and optimality of key architectures. Technical Report Center for Robotics and Embedded Systems Technical Report, CRES-02-005, University of Southern California.
- Gerkey, B.P., Vaughan, R.T., Støy, K., Howard, A., Sukhatme, G.S., and Matarić, M.J. 2001. Most valuable player: A robot device server for distributed control. In *Proc. IEEE/RSJ International Conference on Robots and Systems (IROS)*, Maui, Hawaii, pp. 1226–1331.
- Goldberg, D. and Matarić, M.J. 2002. Design and evaluation of robust behavior-based controllers for distributed multi-robot collection tasks. In *Robot Teams: From Diversity to Polymorphism*, T. Balch and L.E. Parker (Eds.), AK Peters (in press).
- Huntsberger, T., Pirjanian, P., and Schenker, P. 2001. Robotic outposts as precursors to a manned Mars habitat. In *Proc. Space Technology and Applications International Forum (STAIF-2001)*, Albuquerque, NM.
- Maes, P. 1994. Modeling adaptive autonomous agents. *Artificial Life*, 1, 1(2):135–162.
- Matarić, M.J. 1994. Interaction and intelligent behavior. Technical Report AI-TR-1495, MIT Artificial Intelligence Lab.
- Matarić, M.J. 1995. Issues and approaches in the design of collective autonomous agents. *Robotics and Autonomous Systems*, 16(2–4):321–331.
- Mishkin, A., Morrison, J., Nguyen, T., Stone, H., Cooper, B., and Wilcox, B. 1998. Experiences with operations and autonomy of the Mars Pathfinder Microrover. In *1998 IEEE Aerospace Conference Proceedings*, pp. 337–351.
- NASA/JPL. 2002. Sample return and other missions. <http://mars.jpl.nasa.gov/missions/future/2005-plus.html>.
- Østergaard, E.H., Matarić, M.J., and Sukhatme, G.S. 2001. Distributed multi-robot task allocation for emergency handling. In *Proc. IEEE/RSJ International Conference on Robots and Systems (IROS)*, Maui, Hawaii, pp. 821–826.
- Parker, L. 1998. ALLIANCE: An architecture for fault-tolerant multi-robot cooperation. *IEEE Transactions on Robotics and Automation*, 14(2):220–240.
- Pirjanian, P., Huntsberger, T., Trebi-Ollennu, A., Aghazarian, H., Das, H., Joshi, S., and Schenker, P.S. 2000. CAMPOUT: A control architecture for multi-robot planetary outposts. In *Proc. SPIE Symposium on Sensor Fusion and Decentralized Control in Robotic Systems III*, Boston, MA, Vol. 4196.
- Schenker, P., Huntsberger, T., Pirjanian, P., Trebi-Ollennu, A., Das, H., Joshi, S., Aghazarian, H., Ganino, A., Kennedy, B., and Garrett, M. 2000. Robot work crews for planetary outposts: Close cooperation and coordination of multiple robots. In *Proc. SPIE Symposium on Sensor Fusion and Decentralized Control in Robotic Systems III*, Boston, MA, Vol. 4196.
- Werger, B. and Matarić, M. 2000. Broadcast of local eligibility for multi-target observation. In *Proceedings, 5th International Symposium on Distributed Autonomous Robotic Systems (DARS)*, Knoxville, TN, Oct. 4–6, pp. 347–356.



Maja Matarić is an associate professor in the Computer Science Department and the Neuroscience Program at the University of Southern California, Director of the USC Center for Robotics and Embedded Systems, and of the USC Robotics Research Lab. She received her Ph.D. in Computer Science and Artificial Intelligence from MIT in 1994, her M.S. in Computer Science from MIT in 1990, and her B.S. in Computer Science from the University of Kansas in 1987. She is a recipient of the NSF Career Award, the IEEE Robotics and Automation Society Award Early Career Award, the MIT TR100 Innovation Award, and the USC School of Engineering Junior Research Award. She is a editor of three major journals: the IEEE Transactions on Robotics and Automation, the International Journal of Autonomous Agents and Multi-Agent Systems, and Adaptive Behavior, has published over 30 journal articles, 7 book chapters, 65 conference papers, and 20 workshop papers, and has two books in the works with MIT Press. She has been involved with NASA's Jet Propulsion Lab, the Free University of Brussels AI Lab, LEGO Cambridge Research Labs, GTE Research Labs, the Swedish Institute of Computer Science, and ATR Human Information Processing Labs. Her research is in the areas of control and learning in behavior-based multi-robot systems and skill learning by imitation based on sensory-motor primitives.



Gaurav Sukhatme is an Assistant Professor in the Computer Science Department at the University of Southern California (USC)

and the co-director of the USC Robotics Research Laboratory. He received his M.S. and Ph.D. in Computer Science from USC. His research interests include embedded systems, sensor networks, mobile robot coordination and control, sensor fusion for robot fault tolerance, and human-robot interfaces. Dr. Sukhatme has served as PI or CoPI on several NSF, DARPA and NASA grants and contracts. At USC, he directs the Robotic Embedded Systems Lab, which performs research in two related areas: (1) The control and coordination of large numbers of distributed embedded systems, and (2) The control of systems with complex dynamics (hopping robots, robotic helicopters and haptic interfaces). Dr. Sukhatme is a member of AAAI, IEEE and the ACM and has served on several conference program committees. He has published over 60 technical papers, five book chapters and several workshop papers.



Esben H. Østergaard received a masters degree in computer science from University of Aarhus in 2000, worked as a research scientist at the USC Robotics Research Labs from 2000–2001, after which he started his Ph.D. studies at the Maersk Institute, University of Southern Denmark. Research interests include robot soccer, evolutionary robotics, multi-robot coordination and self-reconfigurable robots.