

# Multi-robot Task Allocation in the Light of Uncertainty

Esben H. Østergaard<sup>2,\*</sup>, Maja J Matarić<sup>1</sup>, and Gaurav S. Sukhatme<sup>1</sup>

*esben@mip.sdu.dk, {mataric|gaurav}@robotics.usc.edu*

<sup>1</sup>Robotics Research Laboratories  
Department of Computer Science  
University of Southern California  
Los Angeles, CA 90089-0781

<sup>2</sup>Maersk Mc-Kinney Moller Institute  
for Production Technology  
University of Southern Denmark  
Campusvej 55 DK-5230, Odense

## Abstract

*We 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 the results from real world 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.*

## 1 Introduction

There has been significant prior research in multi-robot coordination [1, 12, 2, 14, 3, 15, 7]. We view this problem as an instance of *dynamic task allocation*. Presently, a general theory of task allocation for multi-robot domains remains elusive. This paper empirically derives general guidelines for selecting task allocation strategies for multi-robot systems. The guidelines are necessarily incomplete given the empirical nature of the work. 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.

## 2 Problem Statement

In the context of multi-robot coordination, dynamic task allocation can be viewed as the selection of appropriate actions [10] 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 [11]. In practice, even with a small number of robots, this is an extremely high-dimensional mapping, a key motivation for decomposing and distributing control. In [6], a system is described that decomposed the task 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 a sequence of actions to execute the task. In [14], each robot's ability to perform a task is mapped to a scalar quantity, which is used to assign tasks to robots. In [15], a local eligibility mechanism is described as the robots' perceived ability to perform a task.

We use the approach from [6] 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 Figure 1, serves two purposes:

---

\*Contact author. This work was performed during the first author's one-year stay at the USC Robotics Research Labs.




Figure 1: Reducing dimensionality of multi-robot coordination.

it reduces the dimensionality of the coordination problem, and it reduces the amount of inter-robot communication required. Instead of a mapping  $S^{|R|} \rightarrow A^{|R|}$  we now have the mapping  $B^{|R||T|} \rightarrow T^{|R|}$  (all robots' bids for all tasks to a task assignment for each robot). We call this mapping the *Task Allocation Strategy* for the system as a whole. We treat it here as a global, centralized process (as depicted in Figure 1), but distributed auctioning mechanisms [5, 6], blackboard algorithms [4], and cross-inhibition of behaviors [15] 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 how it should be distributed.

### 3 Four Task Allocation Strategies

The mapping from bids to tasks can be done in many different ways. Here we consider a Markovian system, where the task allocation mapping for a given robot is from that robot's current task assignments and every robot's current bid on each task, to the given robot's new task assignment, as shown in Figure 2. 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 explored the effects that *commitment* and *coordination* have on performance in the context of four task allocation strategies. These four were derived from the combination of two variables: the amount of commitment to a given task engagement, and the amount of coordination among the robots (see Figure 3). Along the commitment axis we examined a fully committed strategy and a fully opportunistic strategy. Along the coordination axis we examined an (uncoordinated) individualistic strategy and a (highly-coordinated) mutually exclusion strategy, where no two robots were allowed to be engaged in the same task at the same time. Fig-

<i>Cur.eng.</i>		<i>Bids</i>	A	B	C	D	<i>New eng.</i>
A	×	<b>R1</b>	6	4	2	5	?
-		<b>R2</b>	4	1	0	3	?
C		<b>R3</b>	7	2	3	2	?

Figure 2: An example task allocation scenario.

<i>Strategies</i>	Coordination		
		Individ.	Mut. Excl.
Commitment	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.

ure 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, one of the four algorithms we tested, the fully committed mutually exclusive strategy, looks 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 (i.e., uncoordinated) strategies, the same algorithm is run on a separate table for each robot. In the opportunistic (i.e., uncommitted) case, step 1 above is skipped.




Figure 4: An example  $10 \times 10$  grid world with four robots and three active alarms.

## 4 Experimental Validation

### 4.1 The Task

We used *emergency handling* [13] as our problem task domain for evaluation. In it, *alarms* occur at unpredictable times in an office (thus planar) environment. 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 this implementation we restrict ourselves to the case where any robot can fix any alarm.

### 4.2 Grid World Experimental Setup

We implemented a simplified version of the multi-robot emergency handling task in a grid world, as illustrated in Figure 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  $10 \times 10$  grid inhabited by 10 “robots”. 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.




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.




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

### 4.3 Physical Experimental Setup

In our experiments with real robots, we used ActivMedia Pioneer 2 DX mobile robots, equipped with 233MHz Linux PCs, SICK laser range finders, cameras, wireless Ethernet, speakers, and microphones, as shown in Figure 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* [8], a server and protocol that connects robots, sensors, and control programs through a standard TCP socket.<sup>1</sup>

In the physical experiments, alarms were speakers placed in the environment (Figure 6), marked with 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 realistically estimate their distance to the alarms. Instead, they simply used the absolute alarm intensity to decide which robot would win the bid. When assigned to

<sup>1</sup>Player was developed at the USC Robotics Research Lab and is freely available under the GNU Public License from <http://robotics.usc.edu/player/>

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).

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, and 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 30s). 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 [13].

#### 4.4 Grid World Experimental Results

In both the simulated and physical experiments we measured performance as the sum of the number of active alarms at each time-step. After executing the base case in simulation for 1000 time-steps, we obtained the data shown in Figure 7. The combination of mutual exclusion and opportunism gave the best performance. However, this was not always the case, since parameters such as world size, number of robots, noise, and distribution of new alarms could change. When these parameters were varied, any of the four task allocation strategies we tested could outperform the rest. We chose to focus our analysis on noise and uncertainty, since they are key issues in real world robotic 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.

Figure 8 shows the results from varying these two parameters. The Y-axis shows “actuator noise” varied from 0% to 50%. The X-axis shows  $\sigma$  for the “sensor noise” varied from 0 to 20. It can be seen that for low amounts of noise, mutual exclusion and opportunism work best,




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  $50 \times 50$  settings, obtained by varying “sensor noise” and “actuator noise” parameters.

whereas for larger amounts of noise, commitment and individualism work best.

#### 4.5 Physical Experimental Results

We performed two sets of experiments with the physical robots. The first used the setup described above to test the performance of each of the four task allocation strategies. In the second set of experiments we reduced the sensor and actuator noise of the real world system. 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 a higher number of correct bids.

These two experiments correspond to two points in the space shown in Figure 8. The first setup has higher sensor noise and higher actuator uncertainty than the second, as illustrated in Figure 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 Figure 10. In addition, an applet showing a visualization of the experiments is found at <http://robotics.usc.edu/applets/taskalloc>. As expected, the




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.

Base case				
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
$\mu$	<b>1103.00</b>	<b>1104.83</b>	<b>1056.17</b>	<b>1137.33</b>
$\sigma$	315.52	263.80	195.79	129.19

Reduced noise				
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
$\mu$	<b>952.50</b>	<b>1015.50</b>	<b>850.67</b>	<b>899.50</b>
$\sigma$	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.




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

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 Figure 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 Figure 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.

## 5 Discussion

In general, it is desirable to know whether the tendencies derived from the grid world simulation apply to the real world 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 Figure 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. One alternative is to perform more trials with the existing setups. Another is to design a third setup where further noise is added to the system, as suggested in Figure 11.

The grid world results are interesting if we believe that

they actually represent real world system behavior. The fact that the 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 [9], for a controller that, when noise is increased, degrades gracefully from a mutually excluding to an individualistic strategy.

Note that the four task allocation strategies we examined are in a sense *extreme*. Presumably the best strategy for any particular task would be a compromise. As stated previously, the goal of this work was not to attempt to find the best strategy, but rather to gain some insight into task allocation in general. It seems, however, that the four chosen strategies provide a reasonable span of the space of possible strategies.

## 6 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, 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 world 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.

## Acknowledgments

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

## References

- [1] Ronald C. Arkin. Cooperation without communication: Multiagent schema based robot navigation. *Journal of Robotic Systems*, 9(3):351–364, April 1992.
- [2] T. Balch and R. Arkin. Behavior-based formation control for multi-robot teams. *IEEE Transactions on Robotics and Automation*, 14(6):1–15, 1998.
- [3] Y. Cao, A. Fukunaga, A. Kahng, and F. Meng. Cooperative mobile robotics: Antecedents and directions. *Autonomous robots*, 4(1):7–27, 1995.
- [4] Daniel D. Corkill. Blackboard systems. *AI Expert*, 6(9):40–47, September 1991.
- [5] M Bernardine Dias and Anthony (Tony) Stentz. A free market architecture for distributed control of a multirobot system. In *6th International Conference on Intelligent Autonomous Systems (IAS-6)*, pages 115–122, July 2000.
- [6] Brian Gerkey and Maja J Matarić. Principled communication for dynamic multi-robot task allocation. In D. Rus and S. Singh, editors, *Experimental Robotics VII, LNCIS 271*, pages 353–362. Springer-Verlag Berlin Heidelberg, 2001.
- [7] Brian Gerkey and Maja J. Matarić. Pusher-watcher: An approach to fault-tolerant tightly-coupled robot coordination. In *Proceedings, IEEE International Conference on Robotics and Automation*, Washington DC, May 2002.
- [8] Brian P. Gerkey, Richard T. Vaughan, Kasper Støy, Andrew Howard, Gaurav S. Sukhatme, and Maja J Matarić. Most valuable player: A robot device server for distributed control. In *Proc. IEEE/RSJ International Conference on Robots and Systems, (IROS)*, Maui, Hawaii, October 2001 (to appear).
- [9] Dani Goldberg and Maja J Matarić. Design and evaluation of robust behavior-based controllers for distributed multi-robot collection tasks. In Tucker Balch and Lynne E. Parker, editors, *in Robot Teams: From Diversity to Polymorphism*. AK Peters (in Press), 2002.
- [10] P. Maes. Modeling adaptive autonomous agents. *Artificial Life, I*, (1&2)(9), 1994.
- [11] Maja J Matarić. Interaction and intelligent behavior. Technical Report AI-TR-1495, MIT Artificial Intelligence Lab, 1994.
- [12] Maja J Matarić. Issues and approaches in the design of collective autonomous agents. *Robotics and Autonomous Systems*, 16(2–4):321–331, December 1995.
- [13] Esben H. Østergaard, Maja J. Matarić, and Gaurav S. Sukhatme. Distributed multi-robot task allocation for emergency handling. In *Proc. IEEE/RSJ International Conference on Robots and Systems, (IROS)*, pages 821–826, Maui, Hawaii, 2001.
- [14] L. Parker. Alliance: An architecture for fault-tolerant multi-robot cooperation. *IEEE Transactions on Robotics and Automation*, 14(2):220–240, 1998.
- [15] B. Werger and M. Matarić. Broadcast of local eligibility for multi-target observation. In *Proceedings, 5th International Symposium on Distributed Autonomous Robotic Systems (DARS)*, Knoxville, TN, Oct 4-6, pages 347–356, Oct. 2000.