

Distributed Multi-Robot Task Allocation for Emergency Handling

Esben H. Østergaard^{2*}, Maja J Matarić¹, and Gaurav S. Sukhatme¹

esben|matarić|gaurav@robotics.usc.edu

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

²Maersk Mc-Kinney Møller Institute
for Production Technology
University of Southern Denmark
Campusvej 55 DK-5230, Odense

Abstract

We describe a new prototype task, emergency handling, for multi-robot coordination. The experiments reported here measure the effects of individualism and opportunism in a physically-implemented multi-robot system. We use sound at multiple frequencies to simulate emergencies by producing several locally-sensable gradients in the environment. Our results show that opportunism affords a significant performance improvement over individualism. Our experiments also demonstrate the viability of sound for producing detectable local gradients in the environment.

1 Introduction

This paper contributes to research into multi-robot coordination[3] in three ways: it describes emergency handling, a new prototype task for multi-robot coordination; it introduces the use of sound for modeling locally-sensable gradient information in the environment; and it presents results from experiments which examine the role of opportunism in multi-robot coordination.

We introduce the *Emergency Handling Task*, motivated as follows. Consider a typical indoor environment in which “alarms” occur at unpredictable times. We use the term “alarm” to describe any event in the environment that requires the robots’ intervention. An alarm can be caused by hazardous material spills, a water leak, or a fire, characterized by varying severity and unpredictability. The task of the emergency handling robots is to patrol the environment, and upon detecting alarms, to use the right tools and take appropriate action.

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

Definition 1: The Emergency Handling Task consists of:

- An environment, E
- A set of robots, R
- A set of alarms, A
- A set of tools, T
- A capability function, $c : R \rightarrow T$
- A requirement function, $s : A \rightarrow T$

One robot can carry $|c(r_i)|$ tools, where $0 \leq |c(r_i)| < |T|$, $0 \leq i < |R|$. Each alarm can require $|s(a_i)|$, $0 \leq |s(a_i)| < \min(|R|, |T|)$, $0 \leq i < |A|$ tools to be fixed. We require that all alarms can be handled with one or more of the available tools. Robots are heterogeneous if they are equipped with different tools or have different capabilities. Otherwise, the robots are homogeneous.

As defined, the emergency handling task provides room for variation. By varying the parameters $|R|$, $|A|$, $|T|$ and the functions c and s , a spectrum of multi-robot coordination problems is generated. Further, by varying availability of local communication, global communication, the rate at which alarms appear, time needed to fix alarms in relation to the speed of the robots, and so on, several test cases are generated. In all cases, deciding which robot should go where and when is a scheduling problem. Since the robots do not know when and where the alarms will appear, the scheduling problem is dynamic. Furthermore, we aim to design a system robust to failure, which suggests a distributed architecture. Thus we seek a *dynamic, distributed scheduling*

<i>Experiments</i>	Coordination		
		Individ.	Mut. Excl.
Commitment	Commit.	Exp. 1	Exp. 2
	Oppor.	Exp. 3	Exp. 4

Table 1: *The four experiments are set up as combinations of two variables.*

algorithm that performs well in all of the above cases, as the parameters are varied.

In multi-robot coordination, deciding which robot should go and fix an alarm, when one is detected, is a key problem. The solution depends on the robots’ ability to handle the alarm, their metric distance to each alarm, interference along the route, density of other robots in the area, level of confidence as to where the alarm is, and so on. This paper does not focus on developing an optimal strategy for solving the emergency handling problem, but instead on designing a decentralized approximation algorithm capable of performing robustly and adequately.

The rest of this paper is organized as follows. We discuss the use of sound as a model for any locally-detectable gradient in the environment. Next, a brief discussion of the motivation behind the experimental design is given. Then, the implemented multi-robot system is described, and the results are presented. Finally we discuss related work and conclude with the contribution of the paper.

2 Sound Servoing

The task of the emergency handling robots is to detect and fix “alarms”. Many of the example alarms mentioned in the Introduction have the property that a robot can detect them before it can localize them. For example a robot might be able to detect water on the floor before it detects the leak, or smoke before it detects the fire. To simulate this property in our experimental setup, we caused the alarms to emit sound that could be detected by the robots. Sound has some of the same properties as smoke in the air and water on the floor. It is detectable before the robots are within visual range of the source of the alarm, and it enforces locality in that sound emerging from the source of the alarm generates a gradient. Many real alarms might even be detectable directly from the sound they generate, like the sound of fire or people calling for help. All that we require is that the alarms generate some sort of gradient which robotic sensors can measure. Sound is a convenient synthetic alarm, however the gradient generated by a sound

source is far from being perfectly smooth. Sound reflects from walls in ways that are hard to predict without detailed knowledge of the environment. Nonetheless the overall gradient is present and useful.

3 Experimental Design

To explore the parameter space of the task, we focused on *commitment* and *coordination*. In the context of emergency handling, commitment means that robots stay focused on a single alarm, until the alarm can no longer be detected. The opposite, *opportunism*, means that robots can switch tasks, if for example another alarm is detected with greater intensity or priority. In our experiments, coordination is linked to communication, namely the ability of robots to communicate about who should service which alarms, as opposed to *individualism*, where robots have no awareness of each other. Communication is used to prevent multiple robots from trying to fix the same alarm; robots inhibit others from engaging in the same alarm. The goal is to reduce interference among robots, and to prevent loss of coverage in some areas because all the robots rush to fix an alarm in another area. Deciding the level of commitment and collaboration are key aspects of the multi-robot task allocation problem.

We designed four experiments resulting from the combinations in varying the two parameters, coordination and commitment, as shown in Table 1. On one axis we test commitment versus opportunism, and on the other we test individualism versus mutual exclusion.

4 The Implemented System

Here we briefly describe the parts of the system implementation used to perform the experiments.

4.1 The Robots

We used Pioneer 2 DX robots from ActivMedia, equipped with 233MHz Linux PCs, SICK laser range finders, color cameras, wireless Ethernet, speakers and microphones, as shown in Figure 1 and 2. The microphone is made directional by placing it at the bottom of two Styrofoam cups. All control of the robots is done through *Player* [6], a server and protocol that connects robots, sensors, and control programs through a standard TCP socket.¹

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




Figure 1: A fully equipped Pioneer robot.

4.2 The Sound

Each Pioneer robot is equipped with a sound card, speaker and microphone. A separate thread analyzes the input from the microphone. The analog input is sampled at 8000Hz in 8-bit samples and put into a 800-byte buffer. When 800 bytes are read, the buffer is Fourier transformed to produce frequency data in the range 0-4000Hz. The data are smoothed and differentiated to find peaks, and the height and frequency of the five highest peaks is written to a shared buffer, as the current output from the audio system. The speakers are ordinary passive computer speakers, and the microphones are ordinary monitor microphones.

This system makes it possible to distinguish pure sine wave sounds with reasonable accuracy, with a resolution of approximately ± 5 Hz. We used the frequencies 1700, 1800, 1900 and 2000 Hz, each corresponding to a different alarm. These frequencies were chosen empirically, based on the speaker and microphone performance and low background noise in that range.

4.3 The Test Environment

We studied commitment vs. opportunism with a set of three homogeneous robots and four alarms. To ensure homogeneity, each robot is equipped with all four tools, in this case counter sounds, so each robot is capable of fixing any alarm. The environment chosen was a bounded part of our building, containing three offices and a copier room (Figure 3). Alarms could occur at selected points A-D and a speaker was placed at each location. Next to each speaker we placed brightly colored paper markers for visual servoing. When the alarm was “on”, the speaker emitted a sine wave at its unique frequency. The alarms were turned “on” in a fixed pattern, as follows: 30 seconds from the start of the experiment



Figure 2: A close-up of the sensors. The microphone is glued to the bottom of two Styrofoam cups.

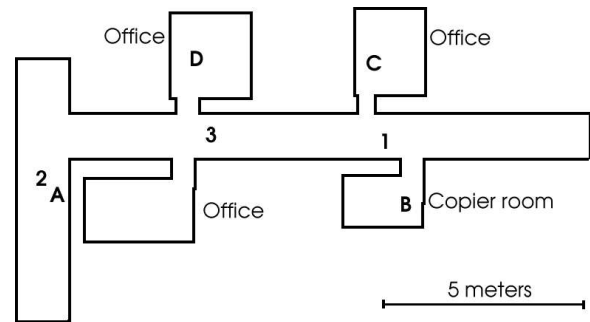


Figure 3: The environment used. A-D are alarm positions, 1-3 are robot start positions.

alarm A goes on. After 5 minutes all 4 alarms go on. Next, alarms D, B and C turn on after 12, 13 and 15 minutes, respectively. After 20 minutes, the experiment is finished regardless of whether all alarms have been fixed.

The volume of sound was such that robots could detect an alarm from 8-16 meters away. The range resulted from factors like background noise, structure of the environment (e.g., number of turns in the sound path), etc. For example, when the robots were in the copier room, they could not hear anything because of the noise from the copying machine, while robots in the corridor could hear any alarm. In addition, the noise from other robots moving could influence the range at which sound could be heard.²

²Turning up the volume would make the sounds audible further away but would be badly received by other users of the office building.

4.4 The Controller

Each robot ran the same control program, producing the following behavior: whenever a robot gets engaged in fixing an emergency, it performs a 360° scan of the intensity of the frequency corresponding to its engagement, and turns towards the highest intensity. Then it goes forward until a junction is detected, performs a new scan, and so forth, until it is within visual range of the alarm, whereupon it relies on visual servoing. When the robot is sufficiently close to the alarm, it emits a counter sound at half the frequency of the alarm. Given the right proximity and right frequency, this counter sound turns off the alarm. Note that no knowledge of the environment in the form of building maps or positions of the alarms was coded into the robot controllers. Thus they can perform in any office-like environment.

4.4.1 The Blackboard Algorithm

To ensure reasonable scalability and robustness, communication among the robots was done through a “blackboard”[4]. To simulate experiments with inter-robot communication, each robot sends its relevant state information to the blackboard at 10Hz, and the blackboard information is read by all the robots at 1Hz. In the case of no communication, the blackboard just contains information from one robot (itself). The information on the blackboard is the current engagement of each robot, and the intensity for each frequency heard by the robots, as shown in Table 2. Intuitively, if all robots have the same blackboard information available and execute the same algorithm, they should all come to the same conclusion as to which robot should pursue which alarm.

To facilitate validation of the four experiments, all parameters were held constant, except the way the information on the blackboard was handled. The algorithm assumes that sensing a higher-intensity of an alarm implies a better fitness to pursue and fix the alarm. The algorithm for calculating which robot should do what is as follows:

- Step 1** All robots engaged in an alarm they cannot sense have their engagement set to ‘none’.
- Step 2** In the case of commitment, all entries on the blackboard for robots already pursuing an alarm is set to zero, along with all entries for alarms already being pursued. In the case of opportunism, this step is skipped.
- Step 3** The highest non-zero score in the table is found, and the robot corresponding to this entry is assigned to the alarm corresponding to this entry.

<i>Blackboard</i>		Alarms			
Robot	Engagement	A	B	C	D
1	none	1400	20	0	0
2	A	800	80	0	0
3	B	900	40	0	0

Table 2: *Example Information on the Blackboard*

<i>Results</i>	Individ.		Mut. Exclusion	
Commitment	2063	1	2325	2
	2016	2	1919	1
	1786	2	2008	1
Opportunism	1087	0	2061	2
	928	0	1406	1
	1917	0	1078	0
			1322	0

Table 3: *Quantitative results. The larger numbers are total sum of on-time for alarms in each trial. The smaller numbers are alarms left on at the end of the trial.*

Then all entries of the robot and the alarms are set to zero. Step 3 is repeated with the new table, until no new assignments are made.

This algorithm has the effect that in the case of commitment robots keep themselves engaged in pursuing an alarm until it is fixed, while in the case of opportunism, robots keep switching engagements.

5 Results

Each experiment was performed three times, with the exception of mutual exclusion/opportunism case, for which 4 trials were performed. Figure 4 shows the total on-time for alarms A-D in each of the 13 trials. The amount of black on the figure corresponds to the sum of alarm on-time. Short black lines indicate that a robot happened to be ready to fix the alarm shortly after it occurred. The pattern at which the alarms were turned on can be distinguished in the data, for example after 300

	Individ.		Mut. Exclusion	
	μ	σ	μ	σ
Commitment	1955	148	2084	213
Opportunism	1311	531	1314	433

Table 4: *Mean and variance of the total alarm on-time in each experiment. Alarms remaining unfixed were not taken into consideration.*

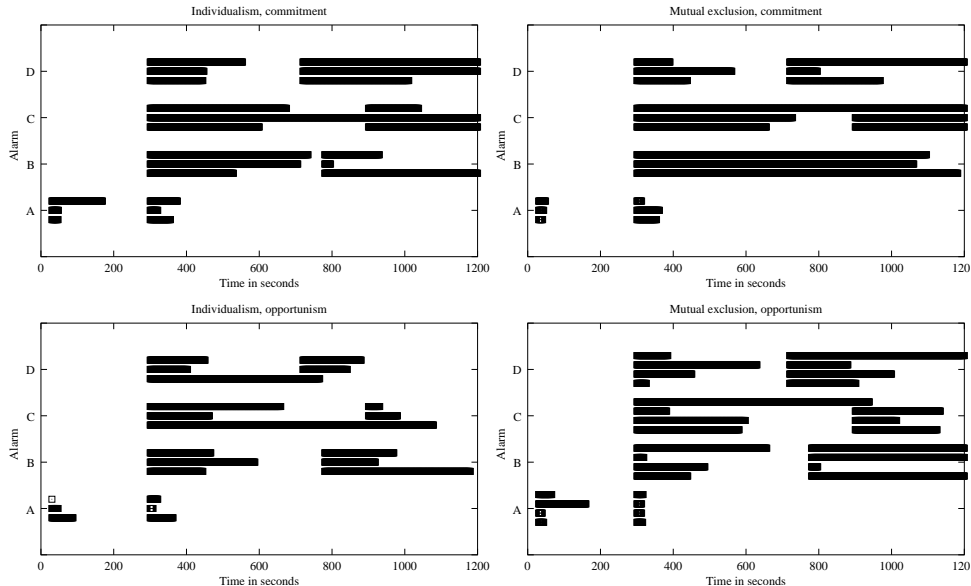


Figure 4: State of each alarm during each of the 15 trials. The x-axis shows the elapsed time during each trial, the y-axis the four alarms, A-D. For each alarm there are three or four lines; the top line for each alarm corresponds to the first trial of the experiment, the second line to the second trial, and so on.

	Individ.		Mut. Exclusion	
	μ	σ	μ	σ
Commitment	16.89	3.44	15.00	6.14
Opportunism	108.78	58.45	125.67	41.62

Table 5: Mean and variance of number of times robots changed their engagement through each trial.

seconds, where all four alarms go on at the same time. Just by examining the amount of black in the graph, it can be seen that opportunism outperforms commitment. Table 3 shows the sum of alarm on-time for each trial, and the number of remaining alarms at the end of each trial. Mean values and variances are shown in Table 4. Statistical analysis shows that there is a significant difference between the commitment row and the opportunism row, and that the difference between the individualism column and the mutual exclusion column is insignificant. Thus the presence of communication in our scenario plays a negligible role. Table 5 shows the average number of times each robot changed engagement through the trials. As expected, in the approach with opportunism robots changed engagement significantly more than in the approach with commitment.

To get a better understanding of the experiments, we created an applet that displays data logged from an

example from each of the four cases graphically. The applet can be seen on: robotics.usc.edu/applets/emergency

6 Related work

There has been significant prior research in multi-robot collaboration [1, 9, 2]. The ALLIANCE architecture [10] presents a robust, multi-robot, task allocation system. An opportunistic approach based on mutual inhibition is presented in [12]. An approach based on commitment is presented in [5], where a task allocation strategy using a market-based auction system commits the robots to their tasks until success or failure.

Previous work on sound-based servoing in robotics includes [7] in which a single robot was able to servo toward a clicking sound in a structured environment, and [8] in which a single robot was shown to servo toward a single sound-emitter based on frequency (corresponding to a cricket chirp) in an open field environment. A survey of robotic phonotaxis (the ability to approach a sound source) is given in [11]. To our knowledge, there is no previous research on several robots servoing on sounds of different frequencies at the same time in a structured environment. Our experiments show that this is viable with inexpensive hardware. The result-

ing setup is an interesting and novel test-bed for mobile robot research.

7 Discussion & Conclusion

Our experiments clearly show that the opportunistic strategy worked significantly better than the commitment-based strategy. This might be because the time to reach an alarm was significantly larger than the time to fix an alarm, once a robot was there. This choice of parameters favors opportunism over commitment since the former effectively uses the presence of robots near emergencies by harnessing them immediately. In other regions of the parameter space of the emergency handling task (e.g., where the ratio of time-to-reach-alarm to time-to-fix-problem is small) opportunism might not be as effective. The present study excluded the case where several robots would be required to fix an alarm in a cooperative fashion, a regime in which performance might improve with commitment. Furthermore, the results do not show any performance improvement by using communication in our version of the emergency-handling task. This does not mean, of course, that communication is not extremely useful in various distributed coordination problems, but it is interesting that it does not have a helpful effect in our system. To test the generality of this result we plan in the future to experiment with the cooperative case requiring multiple robots to fix an alarm.

Our results also demonstrate the effectiveness of using multiple sound frequencies in a noisy environment. During all the trials, most of the alarms were fixed. In 1/3 of the trials, the robots managed to find and put out all the alarms within the available time, in spite of the fact that the way sound propagates through the environment is very complex and highly dependent on the structure of the environment, which makes phonotaxis difficult. We view this as a blessing in disguise, since it provides a realistic simulation of servoing on any non-smooth gradient, like smoke or water on the floor. The resulting setup is thus an interesting and useful test-bed for mobile robotics research.

Our continuing work focuses on large scale exploration of the parameter space of the emergency handling task in order to gain insight into general trends, and on further development of the use of blackboards for multi-robot communications.

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, together with the Maersk Mc-Kinney Moller Institute for Production Technology, University of Southern Denmark.

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] 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.
- [6] 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. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, Wailea, Hawaii, October 2001.
- [7] J. Huang, N. Ohnishi, and N. Sugie. Mobile robot and sound localization. In *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS'97)*, pages 683–689, September 1997.
- [8] H. H. Lund, B. Webb, and J. Hallam. A robot attracted to the cricket species *gryllibus bimaculatus*. In *Proceedings of 4th European Conference on Artificial Life, 1997*, pages 246–255, 1997.
- [9] Maja J Matarić. Issues and approaches in the design of collective autonomous agents. *Robotics and Autonomous Systems*, 16(2–4):321–331, December 1995.
- [10] L. Parker. Alliance: An architecture for fault-tolerant multi-robot cooperation. In *IEEE Transactions on Robotics and Automation*, 14(2):220–240, 1998.
- [11] B. Webb. A robot cricket. *Scientific American*, 275(6):62–67, 1996.
- [12] B. Werger and M. Matarić. Broadcast of local eligibility for multi-target observation. In *Proceedings, of DARS.*, pages 347–356, Oct. 2000.