

# Teams Organization and Performance Analysis in Autonomous Human-Robot Teams

Huadong Wang  
School of Information Sciences  
University of Pittsburgh  
Pittsburgh, PA 15260 U.S.A.  
+1 (412) 624-9426  
huw16@pitt.edu

Michael Lewis  
School of Information Sciences  
University of Pittsburgh  
Pittsburgh, PA 15260 U.S.A.  
+1 (412) 624-9426  
ml@sis.pitt.edu

Shih-Yi Chien  
School of Information Sciences  
University of Pittsburgh  
Pittsburgh, PA 15260 U.S.A.  
+1 (412) 624-9426  
shc56@pitt.edu

## ABSTRACT

This paper proposes a theory of human control of robot teams based on considering how people coordinate across different task allocations. Our current work focuses on domains such as foraging in which robots perform largely independent tasks. The present study addresses the interaction between automation and organization of human teams in controlling large robot teams performing an Urban Search and Rescue (USAR) task. We identify three subtasks: perceptual search-visual search for victims, assistance-teleoperation to assist robot, and navigation-path planning and coordination. For the studies reported here, navigation was selected for automation because it involves weak dependencies among robots making it more complex and because it was shown in an earlier experiment to be the most difficult. This paper reports an extended analysis of the two conditions from a larger four condition study. In these two “shared pool” conditions Twenty four simulated robots were controlled by teams of 2 participants. Sixty paid participants (30 teams) were recruited to perform the shared pool tasks in which participants shared control of the 24 UGVs and viewed the same screens. Groups in the manual control condition issued waypoints to navigate their robots. In the autonomy condition robots generated their own waypoints using distributed path planning. We identify three self-organizing team strategies in the shared pool condition: joint control operators share full authority over robots, mixed control in which one operator takes primary control while the other acts as an assistant, and split control in which operators divide the robots with each controlling a sub-team. Automating path planning improved system performance. Effects of team organization favored operator teams who shared authority for the pool of robots.

## Categories and Subject Descriptors

I.2.9 [Artificial Intelligence]: Robotics—operator interfaces

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PerMIS'10, September 28-30, 2010, Baltimore, MD, USA.

Copyright © 2010 ACM 978-1-4503-0290-6-9/28/10...\$10.00

## General Terms

Human Factors, Measurement, Experimentation

## Keywords

Human-robot interaction, metrics, evaluation, multi-robot system, autonomy, team organization

## 1. INTRODUCTION

Unmanned vehicle systems (UVSs), whether in the air or ground, are intrinsically complex systems and rely on remote operator guidance to accomplish different missions. Specifically, applications for multi-robot systems (MRS) such as interplanetary construction or cooperating uninhabited aerial vehicles will require close coordination and control between human operator(s) and teams of robots in uncertain environments. Human supervision will be needed because humans must supply the perhaps changing, goals that direct MRS activity. Robot autonomy will be needed because the aggregate decision making demands of a MRS are likely to exceed the cognitive capabilities of a human operator. Envisioned missions such as search and rescue or underwater construction, however, will require multiple UV operators to work as teams to control a much larger team of UVs. Yet today, a UV such as the Predator requires a team made up of three operators to be operational [1].

Controlling multiple robots substantially increases the complexity of the operator’s task because attention must be shared among robots in order to maintain situation awareness (SA) and exert control. In the simplest case an operator controls multiple independent robots interacting with each as needed. A foraging task [2] in which each robot searches its own region would be of this category although minimal coordination might be required to avoid overlaps and prevent gaps in coverage especially if robots are in close proximity. Control performance at such tasks can be characterized by the average demand of each robot on human attention [3]. Because robots are operated independently an additional robot imposes only an additive demand on cognitive resources. Under these conditions increasing autonomy for individual robots should allow them to be neglected for longer periods of time making it possible for a single operator to control more robots.

For dependent tasks the round-robin control strategy used for controlling individual robots would force an operator to plan and predict actions needed for multiple joint activities and be highly

susceptible to errors in prediction, synchronization or execution. For highly dependent tasks such as teleoperating robots to push a box, coordination demand for even two robots completely occupies an operator's attention excluding any other task [4]. A multi-UV, multi-operator discrete event simulation model has shown a similar steep increase in difficulty when teams of operators must control interacting robots as a team. [5]

If robots are not rigidly assigned to small teams under the control of a single operator, then each event requiring cooperation will either occupy more of an operator's attention than corresponding independent tasks or require the operator to find the controller of another robot and assume the communication and coordination overhead needed to coordinate with him. These interaction times are likely to be highly variable making it difficult to schedule interactions without introducing excessive idle times. Since even moderate variability in neglect time (NT) has been shown [6] capable of having an operator spending 90% of the time waiting, avoiding such bottlenecks is crucial.

We are developing a general architecture for controlling robot teams based on these observations. We begin by considering how operators self-organize to control robot teams and the effects of the different strategies they adopt. Two possible ways to impose organization on operators are through assigning a subset of robots to each or through a Shared Pool [7] in which operators service robots from the full population as needed. Robot assignment has the advantage of reducing the number of robots the operator must monitor and control. The shared pool offers the scheduling advantage of load balancing in that a pool of operators are available as robots need servicing eliminating situations in which one operator is overloaded and the other idle. Efficiencies such as improved SA that might result from controlling a dedicated team at a particular locale, however, must be sacrificed if operator attention is switched among robots following FIFO (first in first out) or similar discipline.

For monitoring, shared pool offers the redundant observer advantage in that a second observer with partially overlapping perceptual judgments may detect things missed by the first. We expect the effects of these advantages to interact with the types of autonomy possessed by the controlled robots. If navigation and path planning were fully autonomous, we would expect benefits to accrue to shared pool operators due to both scheduling and redundant observer advantages. Autonomous path planning should additionally lessen the effects of loss of SA due to switching between robots because only the victim-marking subtask would be affected. If robots were able to self-reflect and report when they need assistance we might expect to see a stronger scheduling advantage for shared pool. We would additionally expect to see substantial differences between types of autonomy in the numbers of robots that could be adequately controlled.

The present experiment compares performance of robot teams navigating either autonomously or using operator supplied waypoints. The teams were controlled by pairs of operators organized through assigned robots or as a shared pool. In recent experiments [10] we have found that participants performing an Urban Search And Rescue (USAR) foraging task using waypoint control were at or over their limits when controlling 12 robots each. Participants who were asked merely to explore showed very similar performance in area covered and reported similar

levels of workload on the NASA-TLX. Participants in a perceptual search condition in which the foraging task was performed without the requirement to navigate found twice the victims when monitoring 12 robots and reported substantially lower workload.

This paper continues this line of research by allowing robots to autonomously explore, while operators focus purely on the perceptual task. Specifically, we use autonomous path planning for 24 robots and require operators to focus on the perceptual task. Operators can also teleoperate robots that become stuck. There are distinct qualitative differences between the paths taken by autonomous robots and those laid out by human operators. Humans are able to use camera feedback and an intuitive understanding of the environment to reason about angles and perspectives that will give them the most information. The autonomous planning on the other hand, relied exclusively on occupancy grids generated from laser scan data. However, robots have the advantage of being able to focus exclusively on their path planning and work cooperatively to provide the most coverage. One of the question asked in this paper is whether the gain in search efficiency due to autonomous path planning outweighs the less natural paths and, presumably, poorer video angles that result.


The present study uses the same robots and environment but with teams of two operators assigned to control 24 robots. These operators controlled teams of 12 robots in the assigned robots condition. In the shared pool condition operators shared control of the 24 robots. Robots were navigated by operator assigned waypoints as in [10] in the manual condition and by an autonomous path planner in the autonomy condition. Participants were told they were a team and would share a joint score for the experiment. Because team organization was not dictated by the experimenter, operators were free to choose their own strategies for accomplishing their tasks. Strategies could vary over the course of the task. Operators might act with relative independence in the initial path planning phase, for example, while dividing monitoring duties later on. One of the goals of our experiment was to identify the coordination strategies that emerged and find the relationship between these strategies and performance.

## 2. Methods

### 2.1 USARSim and MrCS

The experiment reported in this paper was conducted using the USARSim robotic simulation with 24 simulated Pioneer P2-AT robots performing Urban Search and Rescue (USAR) foraging tasks. USARSim is a high-fidelity simulation of urban search and rescue (USAR) robots and environments developed as a research tool for the study of human-robot interaction (HRI) and multi-robot coordination. USARSim supports HRI by accurately rendering user interface elements (particularly camera video), accurately representing robot automation and behavior, and accurately representing the remote environment that links the operator's awareness with the robot's behaviors. USARSim uses Epic Games' UnrealEngine2 [11] to provide a high fidelity simulator at low cost and also serves as the basis for the Virtual Robots Competition of the RoboCup Rescue League. Other sensors including sonar and audio are also accurately modeled. Validation data showing close agreement in detection of walls and

associated Hough transforms for a simulated Hokuyo laser range finder are described in [13]. The current UnrealEngine2 integrates MathEngine’s Karma physics engine [14] to support high fidelity rigid body simulation. Validation studies showing close agreement in behavior between USARSim models and real robots being modeled are reported in [15,16,17,18,19] as well as agreement for a variety of feature extraction techniques between USARSim images and camera video are reported in Carpin et al. [12]. MrCS (Multi-robot Control System), a multi-robot communications and control infrastructure with accompanying user interface, developed for experiments in multirobot control and RoboCup competition [20] was used in many experiments. MrCS provides facilities for starting and controlling robots in the simulation, displaying multiple camera and laser output, and supporting inter-robot communication through Machinetta which is a distributed multi-agent coordination infrastructure.



**Figure 1. The MrCS user interface with 24 robots for shared pool condition of both autonomy and manual groups.**

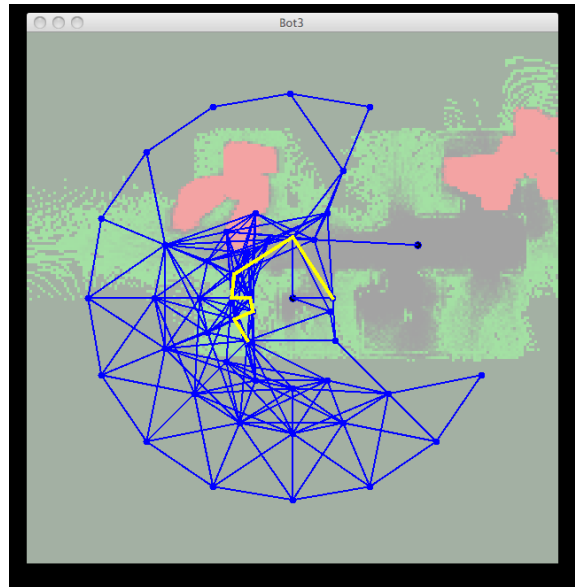
Figure 1 shows the elements of the MrCS. The operator selects the robot to be controlled from the colored thumbnails at the top of the screen. To view more of the selected scene shown in the large video window the operator uses pan/tilt sliders to control the camera. The current locations and paths of the robots are shown on the Map Data Viewer (bottom right). Under manual control, robots are tasked by assigning waypoints on a heading-up map on the Map Viewer (bottom right) or through a teleoperation widget (upper right). In the autonomous condition robots were equipped with autonomous path planning and could explore autonomously. In the shared pool condition the participants have equal authority to control every robot and modify marked victims.

## 2.2 Path Planning

Autonomous path planning was performed by a deterministic roadmap planner [20] developed using the Carnegie Mellon Robot Navigation Toolkit (CARMEN) [21] for these experiments. As input, the planner used the current occupancy grid representing the joint robot team knowledge of the environment and available information about the planned paths of other robots. Possible locations are generated and accepted or rejected based on the expected information gain for being at that location. The expected information gain was a function of the uncertainty in the occupancy grid around that point and whether or not another robot was known to be planning to go near that point. Edges were generated between nodes if the occupancy grid indicated a

sufficiently high probability of being able to move between the locations. Finally, a branch-and-bound search was performed across the network of locations and edges for the path that maximized the expected information gain. Plans were allowed to backtrack, but no additional value was received for visiting a location multiple times. When a robot finished planning, it shared its planned path with some nearby robots to allow them to both avoid collisions and search distinct areas.

Figure 2 shows a screenshot of the path planning debugging interface. Green background shows unexplored areas, which brighter green representing higher uncertainty. Red background is the proposed path of another robot. Blue lines connecting blue circles show the possible locations and edges. The yellow line shows the planned path from the center of the window.



**Figure 2. Screen shot for path planning debugging interface**

## 2.3 Experimental Conditions

A large USAR environment previously used in the 2006 RoboCup Rescue Virtual Robots competition [20] was selected for use in the experiment. The environment was an office like hall with many rooms and full of obstacles like chairs, desks, and bricks. Victims were evenly distributed within the environment. The experiment followed a between groups design with 24 robots. Each task was performed by a team of 2 participants. The teams of participants in the autonomous condition were assigned 24 robots with autonomous path planning capability but they could also control one robot each time via teleoperation. Participants were instructed to use teleoperation only for helping stuck robots, not for exploration. The teams of participants in the manual control condition were assigned 24 robots, for which participants could issue sequences of waypoints as well as teleoperate any robot.

The users were seated at separate interfaces, were able to control the same robots and watch the same video and were able to communicate freely with one another. They were not given any specific instructions on how to coordinate, although they were told it was a cooperative task. Generally, the participants

informally either divided the area or the robots between them. No difference was noticed in the different modalities.

## 2.4 Participants

60 paid participants (30 teams) were recruited from the University of Pittsburgh community balanced among conditions for genders. None had prior experience with robot control although most were frequent computer users.

## 2.5 Procedure

After providing demographic data participants read standard instructions on how to control robots via MrCS. In the following 30 minute training session, participants in all conditions practiced control operations. Participants were encouraged to find and mark at least one victim in the training environment under the guidance of the experimenter. After the training session, participants then began the experimental session (25 minutes) in which they performed the search task controlling 24 robots in teams. After the task, the participants were asked to complete the NASA-TLX workload survey.

Table 1 Performance Measures

Variables	Autonomy		Manual		T-value	P
	(N = 15)		(N = 15)			
	$\bar{x}$	SD	$\bar{x}$	SD		
Victim Found	15.86	2.538	12.33	3.200	3.270	.003
Region Explored	729.69	69.404	638.89	105.433	4.603	.000
Victim/Region ratio	0.020	0.0026	0.193	0.0046	0.473	.640
RMS Errors	0.521	0.1341	0.670	0.0988	-3.424	.002
Missing Sequence	69.71	18.898	43.07	15.714	4.140	.000
Select to Mark Time	31.35	11.593	18.58	9.196	3.298	.003

## 3. Result

### 3.1 Performance Measures

Table 1 shows the T-test results of performance measures between autonomy and manual conditions. Overall participants were successful in searching the environment in both conditions finding as many as 21 victims per team on a trial. The average number of victims found was 15.86 in the autonomous condition but only 12.33 for the manual control condition. A T-test showed that there was a significant difference for victims found in the two conditions ( $t(28) = 3.270$ ;  $p = .003$ ).

The region explored also showed a significant advantage ( $t(28) = 4.603$ ;  $p < .001$ ) for the autonomous condition. The extra exploration was likely due to the autonomous robots moving almost constantly, while in the manual case, an average of 7.66 robots were left after being given a single set of waypoints, with an average of 4.26 receiving no waypoints at all. It is clear from this result that taking the cognitively and time demanding task of exploration away from the operator and automating it helped overall system performance.

While participants enjoying automated path planning found more victims for which areas explored were comparable, their overall advantage in finding victims might have resulted simply from the greater opportunity afforded by exploring larger areas. To examine this possibility we tested the adjusted measure victims/region explored. A T-test was used to test the difference

in the Victim/Region ratio among the autonomous and manual control conditions. The victims found per square meter had no significant difference across the two conditions ( $t(28) = 0.473$ ;  $p = .640$ ). This suggests that the difference in the number of victims found was exclusively due to the larger area searched.

Operator actions, robot states, and artifacts such as laser generated maps were collected throughout the experiment. Later, measures of operator behavior linked to victim observations were annotated in the data. After the user has successfully selected a robot, a series of actions need to be performed to develop sufficient situation awareness to perform the victim marking task, such as stopping the robot, viewing the map and locating the robot which may cost a certain period of time. Otherwise, when the victim passed out of the field of view (FOV) of a robot's camera, it was counted as a missed victim. T-tests showed participants in the Autonomy condition missed more victims,  $t(28)=4.140$ ,  $p<.001$ . The select to mark time showed a similar pattern,  $t(28)=3.298$ ,  $p=.003$ , with participants following manual control marking victims more quickly. However, the related issue of accuracy in marking victims on the laser generated map favored the autonomy conditions. T-tests showed smaller RMS errors in marking for autonomy participants,  $t(28)=-3.424$ ,  $p=.002$ .

### 3.2 Team Organization

Since participants shared control of the 24 UGVs and viewed the same screens in the shared pool conditions, a responsibility allocation procedure was a necessary part of the task. In this study

we observed three patterns of self-organization in the shared pool condition. Some of the participants practiced joint control controlling all 24 robots together, while others followed a Split control strategy by splitting the robots with each controlling a sub-team. A third group of mixed strategy teams failed to settle on an identifiable strategy and instead alternated between strategies suggesting joint or split control. As a result, the shared pool conditions in both autonomy and manual condition could be divided into subgroups according to the way the robots were controlled. (Table 2)

Table 2. Team Organization for shared Pool

Condition	Split	Mixed	Joint
Auto	3	4	8
Manual	9	3	3

A chi-square test of independence was performed to examine the relation between autonomy and team organization (Figure 3). The relation between these variables was significant,  $\chi^2(2, N = 23) = 5.239, p = .022$ , Manual condition participants were less likely to choose joint control than were Autonomy participants.

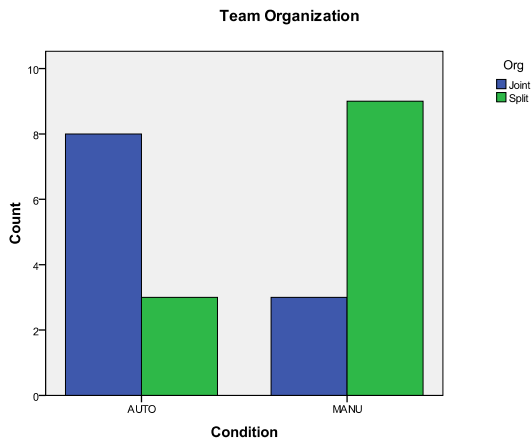


Figure 3. Team Organization Counts

A two-way ANOVA was conducted comparing strategies (joint, mixed, split) and the level of autonomy (auto vs. manual) for all the performance measures. A main effect for team organization was found for victim found per region explored,  $F_{1,26} = 3.627, p = .042$  (Figure 4).

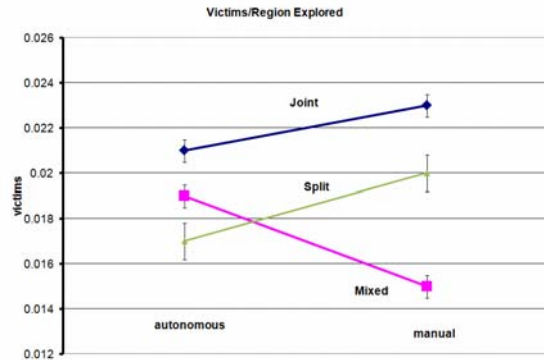


Figure 4. Victims per Region Explored for shared pool

Considering only teams with clearly discernible strategies (joint vs. split) T-tests show joint control participants found more victims,  $t(21) = -2.764, p = .012$  (Figure 5). RMS error showed a similar pattern,  $t(21) = 2.134, p = .045$ , with teams following joint strategies marking victims more accurately (Figure 6). However, joint control participants also missed more victims,  $t(21) = -3.836, p = .001$ .




Figure 5. Victims Found for shared pool

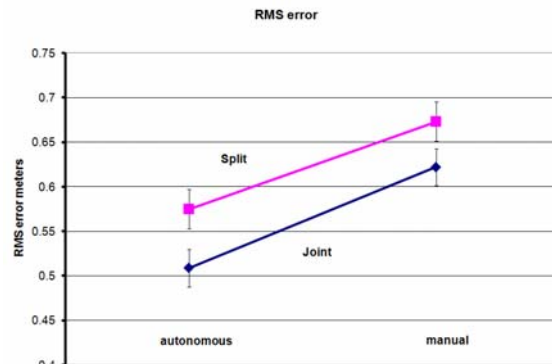


Figure 6. RMS for shared pool

## 4. Discussion

### 4.1 Performance Measures

The current experiment with teams of two operators replicates the effects of automated path planning found in an earlier single operator experiment [10]. In both experiments, relieving operators of the need to perform path planning and manually controlling robots led to finding more victims and marking their locations more accurately. A deterministic roadmap planner using a distributed information gain algorithm was used in the current experiment to drive the robots unlike the previous study in which paths were played back from those generated by earlier participants. Operators appeared to have little difficulty in following these algorithmically generated paths and identified approximately the same numbers of victims as in [10] following human generated paths.

The extra exploration appears principally due to autonomous robots being able to move more or less continuously with only brief pauses. In the manual condition, by contrast, an average of 6.19 robots were left after being given a single set of waypoints, while an average of 3.13 received no waypoints at all.

Replication of the accuracy advantage for automated path planning was also reassuring because studies such as [23] suggested that this advantage might go in the other direction. Peruch et al. [23] demonstrated that self-controlled viewers tended to develop a rich survey knowledge more quickly than passive observers. Because operators in the manual control condition needed to match landmarks between camera views and the laser map, the active exposure to the environment offered by path planning and entering waypoints might have been expected to provide them a more detailed knowledge of the environment and hence greater accuracy in marking victims. Our replicated finding that automated path planning improves accuracy suggests that either the advantage in reduced cognitive load masks poorer survey knowledge in the autonomous condition or that the frequent switching between robots and viewpoints common to the two conditions allows autonomous participants to develop equivalent or superior survey knowledge.

Since avoiding missed targets is crucial to many foraging tasks such as de-mining or search and rescue, thoroughness may be more important than other performance gains such as widening the search area. The analysis of victims per region explored shows that in the assigned robot condition participants using automated path planning found twenty-two percent more victims. This gain is particularly significant because this group was exploring 67% of the map and coming close to matching the actual density of victims of .029/m<sup>2</sup>. Similar improvements in RMS error and reduction in reported workload suggest that substantial cognitive resources were required for navigation and became available for other subtasks improving overall performance when navigation was automated.

The process measures provide a somewhat different account. Operators in the autonomous conditions found more victims, explored regions more thoroughly, and marked victims more accurately, superiorities in performance of the sort often attributed to better SA. An examination of process measures, however, show the opposite may be true. Autonomous operators, however, miss almost twice as many of the victims that appear in their

cameras, perhaps the result of attempting to monitor continuously moving robots which may have multiple victims in view. Select-to-mark time is much shorter in the manual conditions going as low as 18 seconds, approximately half of the 31 seconds required for autonomous operators controlling dedicated robots. These data suggest that operators in the autonomous path planning condition had, in fact, poorer SA than those choosing paths themselves. Although missing fewer of the victims appearing in their thumbnails operators in the manual condition had fewer opportunities as their robots were often idle at terminal waypoints while those in the autonomous condition moved continuously explaining their advantage on the overall performance measures. This account supports our earlier conjecture that reduced cognitive load may mask poorer survey knowledge (SA) in the autonomous condition. Confusions in marking in the shared pool condition are an additional finding with nearly half of the markings in the shared pool condition leading to deletions while only ~16% did so when controlling dedicated robots.

### 4.2 Team Organization

A premise of our research is that if the supervisory control task can be made more similar to conventional alarm driven control, teams of operators will be able to control increasing numbers of robots due to advantages in load balancing and redundant observations. An overall advantage for Area Explored was not observed. We attribute this lack of effect to the weak contribution the assistance subtask made to team performance. While we had expected the office environment to provide many opportunities in which robots would require human assistance our results show that the robots (autonomous) managed to explore an equivalent area without any human assistance. Because the assistance subtask was the task we expected to benefit most from load balancing, its effects were muted. Suggestions of such effects, however, were found within the shared pool conditions. We hypothesized that increasing automation would improve shared pool performance which allows load balancing and redundant observations to a greater extent than assigned robot performance which does not.

This shared pool advantage was found but only for teams choosing the joint control strategy which allowed it. Joint control teams performance bettered in accuracy as well finding .021 and .023 victims/m<sup>2</sup> respectively in auto and manual conditions. A similar advantage for joint control participants was found in locating victims more accurately (Low RMS error).

These results add to a growing picture of the complex problem of controlling multiple robots with human teams. Such tasks are invariably a mixture of subtasks of varying difficulties and contributions. In this experiment the difficulty of the monitoring task dominated and human interventions of the sort described by the neglect tolerance model had little impact. In future studies we hope to examine a range of task/autonomy combinations to develop a more comprehensive theory of team for teams HRI.

## 5. ACKNOWLEDGMENTS

This research has been sponsored in part by AFOSR FA955008-10356 and ONR Grant N0001409-10680.

## 6. REFERENCES

- [1] Cummings, M.L., Brzezinski, A.S. and J.D. Lee. 2007. The impact of intelligent aiding for multiple unmanned aerial vehicle schedule management. *IEEE Intelligent Systems: Special Issue on Interacting with Autonomy* 22(2): 52-59.
- [2] Cao, Y., Fukunaga, A. and Kahng, A. 1997. Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots*, 4, 1-23.
- [3] Crandall, J., Goodrich, M., Olsen, D. and Nielsen, C. 2005. Validating human-robot interaction schemes in multitasking environments. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 35(4):438-449.
- [4] Wang, J. and Lewis, M. 2007. Assessing coordination overhead in control of robot teams. *Proceedings of 2007 IEEE International Conference on Systems, Man, and Cybernetics*, 2645-2649.
- [5] Mekdeci, B. and Cummings, M.L. 2009. Modeling Multiple Human Operators in the Supervisory Control of Heterogeneous Unmanned Vehicles. *9th Conference on Performance Metrics for Intelligent Systems (PerMIS'09)*, Gaithersburg, MD.
- [6] Goodrich, M., Morse, B., Engh, C., Cooper, J. and Adams, J. 2009. Towards Using UAVs in Wilderness Search and Rescue: Lessons from Field Trials. *Interaction Studies, Special Issue on Robots in the Wild: Exploring Human-Robot Interaction in Naturalistic Environments*.
- [7] Lewis, M., Polvichai, J., Sycara, K. and Scerri, P. 2006. Scaling-up Human Control for Large UAV Teams. *Human Factors of Remotely Operated Vehicles, Advances in Human Performance and Cognitive Engineering Research*, Volume 7, 237-250.
- [8] Gerkey, B. and Mataric, M. 2004. A formal framework for the study of task allocation in multi-robot systems. *International Journal of Robotics Research*, 23(9):939-954.
- [9] Harchol-Balter, M., Crovella, M. and Murta, C. 1998. On Choosing a Task Assignment Policy for a Distributed Server System. *Proceedings of the 10th International Conference on Computer Performance Evaluation: Modelling Techniques and Tools*, 231-242.
- [10] Wang, H., Lewis, M., Velagapudi, P., Scerri, P. and Sycara, K. 2009. How search and its subtasks scale in N robots. *Human Robot Interaction Conference, ACM*.
- [11] (UE 2) UnrealEngine2. 2008. <http://udn.epicgames.com/Two/rsrc/Two/KarmaReference/KarmaUserGuide.pdf>.
- [12] Carpin, S., Stoyanov, T., Nevatia, Y., Lewis, M. and Wang, J. 2006. Quantitative assessments of USARSim accuracy. *Proceedings of PerMIS*.
- [13] Carpin, S., Wang, J., Lewis, M., Birk, A. and Jacoff, A. 2005. High fidelity tools for rescue robotics: Results and perspectives. *Robocup 2005 Symposium*.
- [14] MathEngine Karma User Guide. 2005. <http://udn.epicgames.com/Two/KarmaReference/KarmaUserGuide.pdf>.
- [15] Carpin, S., Lewis, M., Wang, J., Balakirsky, S. and Scrapper, C. 2006. Bridging the gap between simulation and reality in urban search and rescue. *Robocup 2006: Robot Soccer World Cup X*, Springer, Lecture Notes in Artificial Intelligence.
- [16] Lewis, M., Hughes, S., Wang, J., Koes, M. and Carpin, S. 2005. Validating USARsim for use in HRI research. *Proceedings of the 49th Annual Meeting of the Human Factors and Ergonomics Society*, Orlando, FL.
- [17] Pepper, C., Balakirsky, S. and Scrapper, C. 2007. Robot Simulation Physics Validation. *Proceedings of PerMIS'07*.
- [18] Taylor, B., Balakirsky, S., Messina, E. and Quinn, R. 2007. Design and Validation of a Whegs Robot in USARSim. *Proceedings of PerMIS'07*.
- [19] Zaratti, M., Fratarcangeli, M. and Iocchi, L. 2006. A 3D Simulator of Multiple Legged Robots based on USARSim. *Robocup 2006: Robot Soccer World Cup X*, Springer, LNAI.
- [20] Balakirsky, S., Carpin, S., Kleiner, A., Lewis, M., Visser, A., Wang, J. and Zipara, V. 2007. Toward heterogeneous robot teams for disaster mitigation: Results and performance metrics from RoboCup Rescue. *Journal of Field Robotics*.
- [21] Latombe, J.C. 1991. *Robot motion planning*. Springer.
- [22] Montemerlo, M., Roy, N. and Thrun, S. 2003. Perspectives on standardization in mobile robot programming: the carnegie mellon navigation (carmen) toolkit. *In Proc. of IROS*.
- [23] Peruch, P., Vercher, J. and Guthier, G. 1995. Acquisition of Spatial Knowledge through Visual Exploration of Simulated Environments. *Ecological Psychology* 7(1): 1-20.
- [24] Montemerlo, M., Roy, N. and Thrun, S. 2003. Perspectives on standardization in mobile robot programming: the carnegie mellon navigation (carmen) toolkit. *In Proc. of IROS*, Vol. 3, 2436- 2441.
- [25] Lewis, M., Wang, H., Chien, S., Scerri, P., Velagapudi, P., Sycara, K. and Kane, B. (to appear). Teams for teams: Performance in multi-human/multi-robot teams, *Proceedings of the 54th Annual Meeting of the Human Factors and Ergonomics Society*.