

## Mixed-initiative multirobot control in USAR

Jijun Wang and Michael Lewis

*School of Information Sciences, University of Pittsburgh  
USA*

### 1. Introduction

In Urban Search and Rescue (USAR), human involvement is desirable because of the inherent uncertainty and dynamic features of the task. Under abnormal or unexpected conditions such as robot failure, collision with objects or resource conflicts human judgment may be needed to assist the system in solving problems. Due to the current limitations in sensor capabilities and pattern recognition people are also commonly required to provide services during normal operation. For instance, in the USAR practice (Casper & Murphy 2003), field study (Burke *et al.* 2004), and RoboCup competitions (Yanco *et al.* 2004), victim recognition remains primarily based on human inspection.

Human control of multiple robots has been suggested as a way to improve effectiveness in USAR. However, multiple robots substantially increase the complexity of the operator's task because attention must be continually shifted among robots. A previous study showed that when the mental demand overwhelmed the operator's cognitive resources, operators controlled reactively instead of planning and proactively controlling the robots leading to worse performance (Trouvain *et al.* 2003). One approach to increasing human capacity for control is to allow robots to cooperate reducing the need to control them independently. Because human involvement is still needed to identify victims and assist individual robots, automating coordination appears to be a promising avenue for reducing cognitive demands on the operator.

For a human/automation system, how and when the operator intervenes are the two issues that most determine overall effectiveness (Endsley 1996). How the human interacts with the system can be characterized by the level of autonomy (LOA), a classification based on the allocation of functions between human and robot. In general, the LOA can range from complete manual control to full autonomy (Sheridan 2002). Finding the optimal LOA is an important yet hard to solve problem because it depends jointly on the robotic system, task, working space and the end user. Recent studies (Squire *et al.* 2003; Envarli & Adams 2005; Parasuraman *et al.* 2005; Schurr *et al.* 2005) have compared different LOAs for a single operator interacting with cooperating robots. All of them, however, were based on simple tasks using low fidelity simulation thereby minimizing the impact of situation awareness (SA). In realistic USAR applications (Casper & Murphy 2003, Burke *et al.* 2004, Yanco *et al.* 2004) by contrast, maintaining sufficient SA is typically the operator's greatest problem.

The present study investigates human interaction with a cooperating team of robots performing a search and rescue task in a realistic disaster environment. This study uses USARSim (Wang *et al.* 2003), a high fidelity game engine-based robot simulator we

developed to study human-robot interaction (HRI) and multi-robot control. USARSim provides a physics based simulation of robot and environment that accurately reproduces mobility problems caused by uneven terrain (Wang *et al.* 2005), hazards such as rollover (Lewis & Wang 2007), and provides accurate sensor models for laser rangefinders (Carpin *et al.* 2005) and camera video (Carpin *et al.* 2006). This level of detail is essential to posing realistic control tasks likely to require intervention across levels of abstraction. We compared control of small robot teams in which cooperating robots exploring autonomously, were controlled independently by an operator, or through mixed initiative as a cooperating team. In our experiment mixed initiative teams found more victims and searched wider areas than either fully autonomous or manually controlled teams. Operators who switched attention between robots more frequently were found to perform better in both manual and mixed initiative conditions.

We discuss the related work in section 2. Then we introduce our simulator and multi-robot system in section 3. Section 4 describes the experiment followed by the results presented in section 5. Finally, we draw conclusion and discuss the future work in section 6.

## 2. Related Work

When a single operator controls multiple robots, in the simplest case the operator interacts with each independent robot as needed. Control performance at this task can be characterized by the average demand of each robot on human attention (Crandall *et al.* 2005) or the distribution of demands coming from multiple robots (Nickerson & Skiena 2005). Increasing robot autonomy allows robots to be neglected for longer periods of time making it possible for a single operator to control more robots. Researchers investigating the effects of levels of autonomy (teleoperation, safe mode, shared control, full autonomy, and dynamic control) on HRI (Marble *et al.* 2003; Marble *et al.* 2004) for single robots have found that mixed-initiative interaction led to better performance than either teleoperation or full autonomy. This result seems consistent with Fong's collaborative control (Fong *et al.* 2001) premise that because it is difficult to determine the most effective task allocation a priori, allowing adjustment during execution should improve performance.

The study of autonomy modes for multiple robot systems (MRS) has been more restrictive. Because of the need to share attention between robots, teleoperation has only been allowed for one robot out of a team (Nielsen *et al.* 2003) or as a selectable mode (Parasuraman *et al.* 2005). Some variant of waypoint control has been used in all MRS studies reviewed (Trouvain & Wolf 2002; Nielsen *et al.* 2003; Squire *et al.* 2003; Trouvain *et al.* 2003; Crandall *et al.* 2005; Parasuraman *et al.* 2005) with differences arising primarily in behaviour upon reaching a waypoint. A more fully autonomous mode has typically been included involving things such as search of a designated area (Nielsen *et al.* 2003), travel to a distant waypoint (Trouvain & Wolf 2002), or executing prescribed behaviours (Parasuraman *et al.* 2005). In studies in which robots did not cooperate and had varying levels of individual autonomy (Trouvain & Wolf 2002; Nielsen *et al.* 2003; Trouvain *et al.* 2003; Crandall *et al.* 2005) (team size 2-4) performance and workload were both higher at lower autonomy levels and lower at higher ones. So although increasing autonomy in these experiments reduced the cognitive load on the operator, the automation could not perform the replaced tasks as well. This effect would likely be reversed for larger teams such as those tested in Olsen & Wood's (Olsen & Wood 2004) fan-out study which found highest performance and lowest (per robot activity) imputed workload for the highest levels of autonomy.

For cooperative tasks and larger teams individual autonomy is unlikely to suffice. 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. A series of experiments using the Playbook interface and the RoboFlag simulation (Squire *et al.* 2003; Parasuraman *et al.* 2005) provide data on HRI with cooperating robot teams. These studies found that control through delegation (calling plays/plans) led to higher success rates and faster missions than individual control through waypoints and that as with single robots (Marble *et al.* 2003; Marble *et al.* 2004) allowing the operator to choose among control modes improved performance. Again, as in the single robot case, the improvement in performance from adjustable autonomy carried with it a penalty in reported workload. Another recent study (Schurr *et al.* 2005) investigating supervisory control of cooperating agents performing a fire fighting task found that human intervention actually degraded system performance. In this case, the complexity of the fire fighting plans and the interdependency of activities and resources appeared to be too difficult for the operator to follow. For cooperating teams and relatively complex tasks, therefore, the neglect-tolerance assumption (Olsen & Wood 2004; Crandall *et al.* 2005) that human control always improves performance may not hold. For these more complex MRS control regimes it will be necessary to account for the arguments of Woods *et al.* (Woods *et al.* 2004) and Kirlik's (Kirlik 1993) demonstration that higher levels of autonomy can act to increase workload to the point of eliminating any advantage by placing new demands on the operator to understand and predict automated behaviour. The cognitive effort involved in shifting attention between levels of automation and between robots reported by (Squire *et al.* 2003) seems a particularly salient problem for MRS.

Experiment	World	Robots	Task	Team
Nielsen <i>et al.</i> (2003)	2D simulator	3	Navigate/build map	independent
Crandall <i>et al.</i> (2005)	2D simulator	3	Navigate	independent
Trouvain & Wolf (2002)	2D simulator	2,4,8	Navigate	independent
Trouvain <i>et al.</i> (2003)	3D simulator	1,2,4	Navigate	independent
Parasuraman <i>et al.</i> (2005)	2D simulator	4,8	Capture the flag	cooperative
Squire <i>et al.</i> (2006)	2D simulator	4,6,8	Capture the flag	cooperative
Present Experiment	3D simulator	3	Search	cooperative

Table 1. Recent MRS Studies

Table 1 organizes details of recent MRS studies. All were conducted in simulation and most involve navigation rather than search, one of the most important tasks in USAR. This is significant because search using an onboard camera requires greater shifts between contexts than navigation which can more easily be performed from a single map display (Bruemmer *et al.* 2005; Nielsen & Goodrich 2006). Furthermore, previous studies have not addressed the issues of human interaction with cooperating robot teams within a realistically complex environment. Results from 2D simulation (Squire *et al.* 2003; Parasuraman *et al.* 2005), for example, are unlikely to incorporate tasks requiring low-level assistance to robots, while experiments with non-cooperating robots (Trouvain & Wolf 2002; Nielsen *et al.* 2003;

Trouvain *et al.* 2003; Crandall *et al.* 2005) miss the effects of this aspect of autonomy on performance and HRI.

This paper presents an experiment comparing search performance of teams of 3 robots controlled manually without automated cooperation, in a mixed-initiative mode interacting with a cooperating team or in a fully autonomous mode without a human operator. The virtual environment was a model of the Yellow Arena, one of the NIST Reference Test Arenas designed to provide standardized disaster environments to evaluate human robot performance in USAR domain (Jacoff *et al.* 2001). The distributed multiple agents framework, Machinetta (Scerri *et al.* 2004) is used to automate cooperation for the robotic control system in the present study.

### 3. Simulator and Multirobot System

#### 3.1 Simulation of the Robots and Environment

Although many robotic simulators are available most of them have been built as ancillary tools for developing and testing control programs to be run on research robots. Simulators (Lee *et al.* 1994; Konolige & Myers 1998) built before 2000 typically have low fidelity dynamics for approximating the robot's interaction with its environment. More recent simulators including ÜberSim (Browning & Tryzelaar 2003) , a soccer simulator, Gazebo (Gerkey *et al.* 2003), and the commercial Webots (Cyberbotics Ltd. 2006) use the open source Open Dynamics Engine (ODE) physics engine to approximate physics and kinematics more precisely. ODE, however, is not integrated with a graphics library forcing developers to rely on low-level libraries such as OpenGL. This limits the complexity of environments that can practically be developed and effectively precludes use of many of the specialized rendering features of modern graphics processing units. Both high quality graphics and accurate physics are needed for HRI research because the operator's tasks depend strongly on remote perception (Woods *et al.* 2004), which requires accurate simulation of camera video, and interaction with automation, which requires accurate simulation of sensors, effectors and control logic.



Figure 1. Simulated P2DX robot

We built USARSim, a high fidelity simulation of USAR robots and environments to be a research tool for the study of 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. It was built based on a multi-player game engine, UnrealEngine2, and so is well suited for simulating multiple robots. USARSim uses the Karma Physics engine to provide physics modeling, rigid-body dynamics with constraints and collision detection. It uses other game engine capabilities to simulate sensors including camera video, sonar, and laser range finder. More details about USARSim can be found at (Wang *et al.* 2003; Lewis *et al.* 2007).



a) Arena-1

b) Arena-2

Figure 2. Simulated testing arenas

In this study, we simulated three Activemedia P2-DX robots. Each robot was equipped with a pan-tilt camera with 45 degrees FOV and a front laser scanner with 180 degree FOV and resolution of 1 degree. Two similar NIST Reference Test Arenas, Yellow Arena, were built using the same elements with different layouts. In each arena, 14 victims were evenly distributed in the world. We added mirrors, blinds, curtains, semitransparent boards, and wire grid to add difficulty in situation perception. Bricks, pipes, a ramp, chairs, and other debris were put in the arena to challenge mobility and SA in robot control. Figure 1 shows a simulated P2DX robot and a corner in the virtual environment. Figure 2 illustrates the layout of the two testing environments.

### 3.2 Multi-robot Control System (MrCS)

The robotic control system used in this study is MrCS (Multi-robot Control System), a multi-robot communications and control infrastructure with accompanying user interface. MrCS provides facilities for starting and controlling robots in the simulation, displaying camera and laser output, and supporting inter-robot communication through Machinetta (Scerri *et al.* 2004). Machinetta is a distributed multiagent system with state-of-the-art algorithms for plan instantiation, role allocation, information sharing, task deconfliction and adjustable autonomy (Scerri *et al.* 2004). The distributed control enables us to scale robot teams from small to large. In Machinetta, team members connect to each other through reusable software proxies. Through the proxy, humans, software agents, and different robots can work together to form a heterogeneous team. Basing team cooperation on reusable proxies allows us to quickly change size or coordination strategies without affecting other parts of the system.

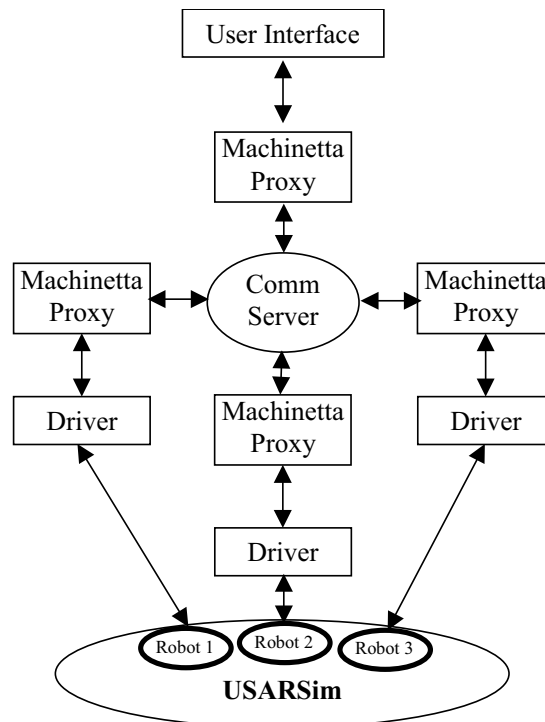


Figure 3. MrCS system architecture

Figure 3 shows the system architecture of MrCS. It provides Machinetta proxies for robots, and human operator (user interface). Each robot connects with Machinetta through a robot driver that provides low-level autonomy such as guarded motion, waypoint control (moving from one point to another while automatically avoiding obstacles) and middle-

level autonomy in path generation. The robot proxy communicates with proxies on other simulated robots to enable the robots to execute the cooperative plan they have generated. In the current study plans are quite simple and dictate moving toward the nearest frontier that does not conflict with search plans of another robot. The operator connects with Machinetta through a user interface agent. This agent collects the robot team's beliefs and visually represents them on the interface. It also transfers the operator's commands in the form of a Machinetta proxy's beliefs and passes them to the proxies network to allow human in the loop cooperation. The operator is able to intervene with the robot team on two levels. On the low level, the operator takes over an individual robot's autonomy to teleoperate it. On the intermediate level, the operator interacts with a robot via editing its exploration plan.

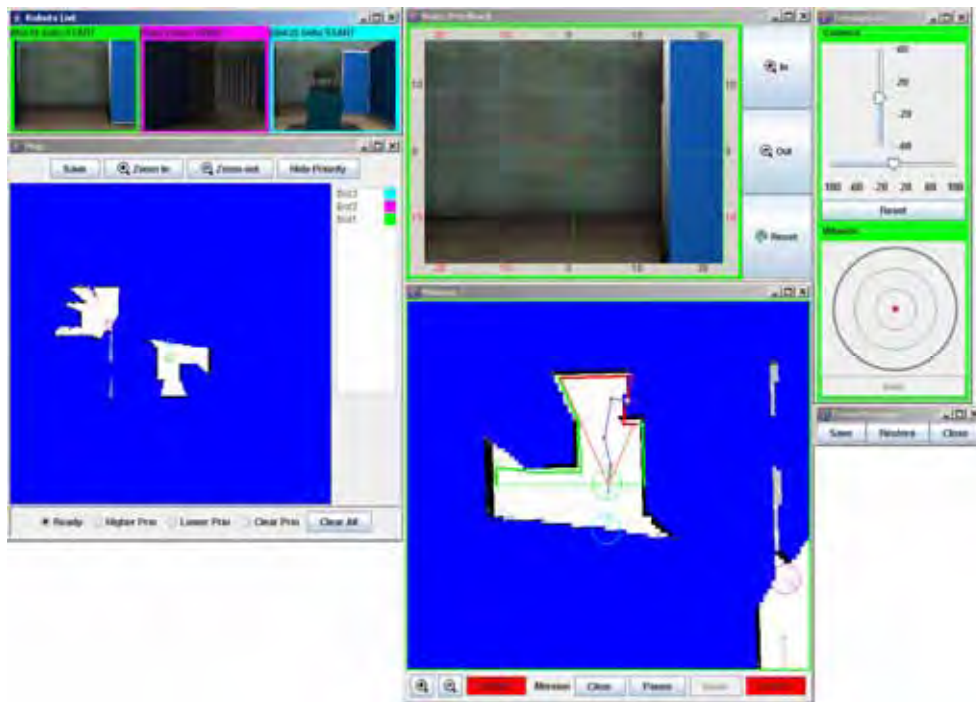


Figure 4. The graphic user interface

In the human robot team, the human always has the highest authority although the robot may alter its path slightly to avoid obstacles or dangerous poses. Robots are controlled one at a time with the selected robot providing a full range of data while the unselected ones provide camera views for monitoring. The interface allows the user to resize the components or change the layout. Figure 4 shows the interface configuration used in the present study. On the left side are the global information components: the Robots List (the upper panel) that shows each team member's execution state and the thumbnail of the individual's camera view; and the global Map (the bottom panel) that shows the explored

areas and each robot's position. From the Robot List, the operator can select any robot to be controlled. In the center are the individual robot control components. The upper component, Video Feedback, displays the video of the robot being controlled. It allows the user to pan/tilt and zoom the camera. The bottom component is the Mission panel that shows the controlled robot's local situation. The local map is camera up, always pointing in the camera's direction. It is overlaid with laser data in green and a cone showing the camera's FOV in red. With the Mission panel and the Video Feedback panel, we support SA at three ranges. The camera view and range data shown in the red FOV cone provide the operator the close range SA. It enables the operator to observe objects through the camera and identify their locations on the map. The green range data shows the open regions around the robot providing local information about where to go in the next step. In contrast, the background map provides the user long range information that helps her make a longer term plan. The mission panel displays the robot's current plan as well to help the user understand what the robot is intending to do. When a marked victim or another robot is within the local map the panel will represent them even if not sensed. Besides representing local information, the Mission panel allows the operator to control a robot by clearing, modifying, or creating waypoints and marking the environment by placing an icon on the map. On the right is the Teleoperation panel that teleoperates the robot or pans/tilts the camera. These components behave in the expected ways.

## 4. Experiment

### 4.1 Experimental Design

In the experiment, participants were asked to control 3 P2DX robots (Figure 1) simulated in USARSim to search for victims in a damaged building (Figure 2). The participant interacted with the robots through MrCS with fixed user interface shown in Figure 4. Once a victim was identified, the participant marked its location on the map.

We used a within subjects design with counterbalanced presentation to compare mixed initiative and manual conditions. Under mixed initiative, the robots analyzed their laser range data to find possible exploration paths. They cooperated with one another to choose execution paths that avoided duplicating efforts. While the robots autonomously explored the world, the operator was free to intervene with any individual robot by issuing new waypoints, teleoperating, or panning/tilting its camera. The robot returned back to auto mode once the operator's command was completed or stopped. While under manual control robots could not autonomously generate paths and there was no cooperation among robots. The operator controlled a robot by giving it a series of waypoints, directly teleoperating it, or panning/tilting its camera. As a control for the effects of autonomy on performance we conducted "full autonomy" testing as well. Because MrCS doesn't support victim recognition, based on our observation of the participants' victim identification behaviours, we defined detection to have occurred for victims that appeared on camera for at least 2 seconds and occupied at least 1/9 of the thumbnail view. Because of the high fidelity of the simulation, and the randomness of paths picked through the cooperation algorithms, robots explored different regions on every test. Additional variations in performance occurred due to mishaps such as a robot getting stuck in a corner or bumping into an obstacle causing its camera to point to the ceiling so no victims could be found. Sixteen trials were conducted in each area to collect data comparable to that obtained from human participants.

#### 4.2 Procedure

The experiment started with collection of the participant's demographic data and computer experience. The participant then read standard instructions on how to control robots via MrCS. In the following 10 minutes training session, the participant practiced each control operation and tried to find at least one victim in the training arena under the guidance of the experimenter. Participants then began a twenty minutes session in Arena-1 followed by a short break and a twenty minutes session in Arena-2. At the conclusion of the experiment participants completed a questionnaire.

#### 4.3 Participants

	Age		Gender		Education			
	19	20~35	Male	Female	Currently Undergraduate	Complete Undergraduate		
<b>Participants</b>	2	12	5	9	10		4	
	Computer Usage (hours/week)				Game Playing (hours/week)			
	<1	1-5	5-10	>10	<1	1-5	5-10	>10
<b>Participants</b>	0	2	7	5	6	7	1	0
	Mouse Usage for Game Playing							
	Frequently			Occasionally			Never	
<b>Participants</b>	8			6			0	

Table 2. Sample demographics and experiences

14 paid participants recruited from the University of Pittsburgh community took part in the experiment. None had prior experience with robot control although most were frequent computer users. The participants' demographic information and experience are summarized in Table 2.

### 5. Results

In this experiment, we studied the interaction between a single operator and a robot team in a realistic interactive environment where human and robots must work tightly together to accomplish a task. We first compared the impact of different levels of autonomy by evaluating the overall performance as revealed by the number of found victims, the explored areas, and the participants' self-assessments. For the small robot team with 3 robots, we expected similar results to those reported in (Trouvain & Wolf 2002; Nielsen *et al.* 2003; Trouvain *et al.* 2003; Crandall *et al.* 2005) that although autonomy would decrease workload, it would also decrease performance because of poorer situation awareness (SA). How a human distributes attention among the robots is an interesting problem especially when the human is deeply involved in the task by performing low level functions, such as identifying a victim, which requires balancing between monitoring and control. Therefore, in addition to overall performance measures, we examine: 1) the distribution of human

interactions among the robots and its relationship with the overall performance, and 2) the distribution of control behaviours, i.e. teleoperation, waypoint issuing and camera control, among the robots and between different autonomy levels, and their impacts in the overall human-robot performance. Trust is a special and important problem arising in human-automation interaction. When the robotic system can't work as the operator expected, it will influence how the operator control the robots and hereby impact the human-robot performance (Lee & See 2004; Parasuraman & Miller 2004). In addition, because of the complexity of the control interface, we anticipated that the ability to use the interface would impact the overall performance as well. At the end of this section, we report participants' self-assessments of trust and capability of using the user interface, as well as the relationship among the number of found victims and these two factors.

### 5.1 Overall Performance

All 14 participants found at least 5 of all possible 14 (36%) victims in each of the arenas. The median number of victims found was 7 and 8 for test arenas 1 and 2 respectively. Two-tailed t-tests found no difference between the arenas for either number of victims found or the percentage of the arena explored. Figure 5 shows the distribution of victims discovered as a function of area explored. These data indicate that participants exploring less than 90% of the area consistently discovered 5-8 victims while those covering greater than 90% discovered between half (7) and all (14) of the victims.




Figure 5. Victims as a function of area explored

Within participant comparisons found wider regions were explored in mixed-initiative mode,  $t(13) = 3.50$ ,  $p < .004$ , as well as a marginal advantage for mixed-initiative mode,  $t(13) = 1.85$ ,  $p = .088$ , in number of victims found. Comparing with "full autonomy", under mixed-initiative conditions two-tailed t-tests found no difference ( $p = 0.58$ ) in the explored regions. However, under full autonomy mode, the robots explored significantly,  $t(44) = 4.27$ ,

$p < .001$ , more regions than under the manual control condition (left in Figure 6). Using two-tailed t-tests, we found that participants found more victims under mixed-initiative and manual control conditions than under full autonomy with  $t(44) = 6.66$ ,  $p < .001$ , and  $t(44) = 4.14$ ,  $p < .001$  respectively (right in Figure 6). The median number of victims found under full autonomy was 5.




Figure 6. Regions explored by mode (left) and victims found by mode (right)

In the posttest survey, 8 of the 14 (58%) participants reported they were able to control the robots although they had problems in handling some components. All of the remaining participants thought they used the interface very well. Comparing the mixed-initiative with the manual control, most participants (79%) rated team autonomy as providing either significant or minor help. Only 1 of the 14 participants (7%) rated team autonomy as making no difference and 2 of the 14 participants (14%) judged team autonomy to make things worse.

## 5.2 Human Interactions

Participants intervened to control the robots by switching focus to an individual robot and then issuing commands. Measuring the distribution of attention among robots as the standard deviation of the total time spent with each robot, no difference ( $p = .232$ ) was found between mixed initiative and manual control modes. However, we found that under mixed initiative, the same participant switched robots significantly more often than under manual mode ( $p = .027$ ). The posttest survey showed that most participants switched robots using the Robots List component. Only 2 of the 14 participants (14%) reported switching robot control independent of this component.

Across participants the frequency of shifting control among robots explained a significant proportion of the variance in number of victims found for both mixed initiative,  $R^2 = .54$ ,  $F(1, 11) = 12.98$ ,  $p = .004$ , and manual,  $R^2 = .37$ ,  $F(1, 11) = 6.37$ ,  $p < .03$ , modes (Figure 7).

An individual robot control episode begins with a pre-observation phase in which the participant collects the robot's information and then makes a control decision, and ends with the post-observation phase in which the operator observes the robot's execution and decides to turn to the next robot. Using a two-tailed t-test, no difference was found in either total pre-observation time or total post-observation time between mixed-initiative and manually control conditions. The distribution of found victims among pre- and post-

observation times (Figure 8) suggests, however, that the proper combination can lead to higher performance.

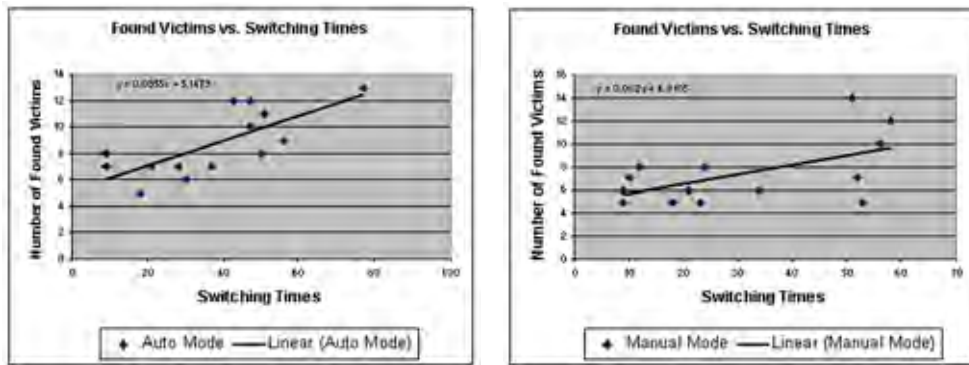


Figure 7. Victims vs. switches under mixed-autonomy (left) and manually control (right) mode

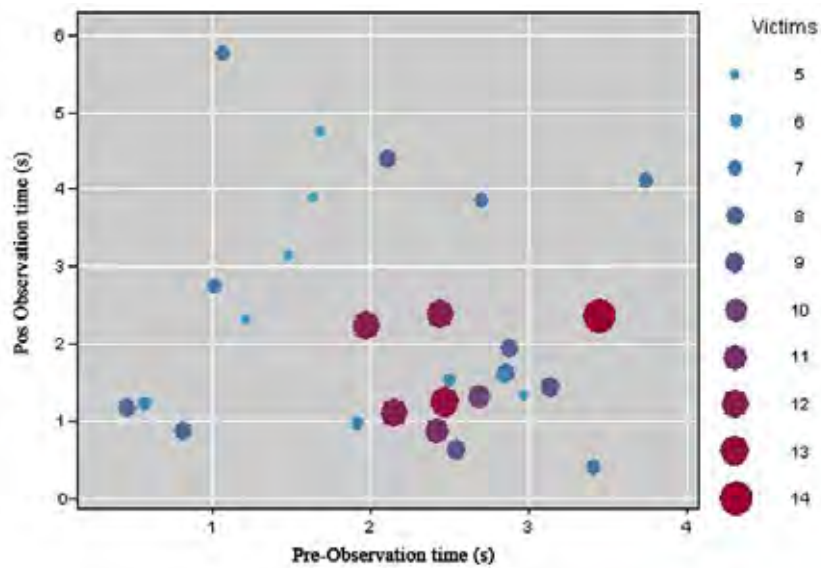


Figure 8. Pre and Post observation time vs. found

### 5.3 Forms of Control

Three interaction methods: waypoint control, teleoperation control, and camera control were available to the operator. Using waypoint control, the participant specifies a series of

waypoints while the robot is in pause state. Therefore, we use the times of waypoint specification to measure the amount of interaction. Under teleoperation, the participant manually and continuously drives the robot while monitoring its state. Time spent in teleoperation was measured as the duration of a series of active positional control actions that were not interrupted by pauses of greater than 30 sec. or any other form of control action. For camera control, times of camera operation were used because the operator controls the camera by issuing a desired pose, and monitoring the camera's movement.

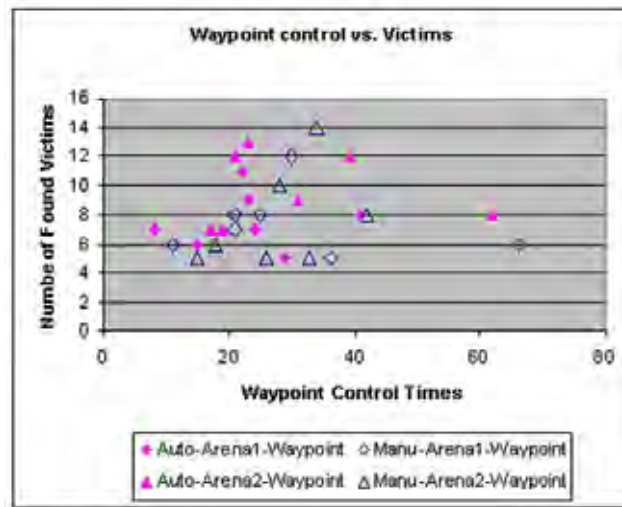


Figure 9. Victims found as a function of waypoint

While we did not find differences in overall waypoint control times between mixed-initiative and manual modes, mixed-initiative operators had shorter,  $t(13) = 3.02$ ,  $p < .01$ , control times during any single control episode, the period during which an operator switches to a robot, controls it and then switches to another robot.

Figure 9 shows the relationship between victims found and total waypoint control times. In manual mode this distribution follows an inverted 'U' with too much or too little waypoint control leading to poor search performance. In mixed-initiative mode by contrast the distribution is skewed to be less sensitive to control times while holding a better search performance, i.e. more found victims (see section 5.1).

Overall teleoperation control times,  $t(13) = 2.179$ ,  $p < .05$  were reduced in the mixed-initiative mode as well, while teleoperation times within episodes only approached significance,  $t(13) = 1.87$ ,  $p = .08$ . No differences in camera control times were found between mixed-initiative and manual control modes. It is notable that operators made very little use of teleoperation, .6% of mission time, and only infrequently chose to control their cameras.

#### 5.4 Trust and Capability of Using Interface

In the posttest we collected participants' ratings of their level of trust in the system's automation and their ability to use the interface to control the robots. 43% of the

participants trusted the autonomy and only changed the robot's plans when they had spare time. 36% of the participants reported changing about half of the robot's plans while 21% of the participants showed less trust and changed the robot's plans more often. A one tail t-test, indicates that the total victims found by participants trusting the autonomy is larger than the number victims found by other participants ( $p=0.05$ ). 42% of the participants reported being able to use the interface well or very well, while 58% of the participants reported having difficulty using the full range of features while maintaining control of the robots. A one tail t test shows that participants reporting using the interface well or very well found more victims ( $p<0.001$ ). Participants trusting the autonomy reported significantly higher capability in using the user interface ( $p=0.001$ ) and conversely participants reporting using the interface well also had greater trust in the autonomy ( $p=0.032$ ).

## 6. Conclusion

In this experiment, the first of a series investigating control of cooperating teams of robots, cooperation was limited to deconfliction of plans so that robots did not re-explore the same regions or interfere with one another. The experiment found that even this limited degree of autonomous cooperation helped in the control of multiple robots. The results showed that cooperative autonomy among robots helped the operators explore more areas and find more victims. The fully autonomous control condition demonstrates that this improvement was not due solely to autonomous task performance as found in (Schurr *et al.* 2005) but rather resulted from mixed initiative cooperation with the robotic team. The superiority of mixed initiative control was far from a foregone conclusion since earlier studies with comparable numbers of individually autonomous robots (Trouvain & Wolf 2002; Nielsen *et al.* 2003; Trouvain *et al.* 2003; Crandall *et al.* 2005) found poorer performance for higher levels of autonomy at similar tasks. We believe that differences between navigation and search tasks may help explain these results. In navigation, moment to moment control must reside with either the robot or the human. When control is ceded to the robot the human's workload is reduced but task performance declines due to loss of human perceptual and decision making capabilities. Search by contrast can be partitioned into navigation and perceptual subtasks allowing the human and robot to share task responsibilities improving performance. This explanation suggests that increases in task complexity should widen the performance gap between cooperative and individually autonomous systems. We did not collect workload measures to check for the decreases found to accompany increased autonomy in earlier studies (Trouvain & Wolf 2002; Nielsen *et al.* 2003; Trouvain *et al.* 2003; Crandall *et al.* 2005), however, eleven of our fourteen subjects reported benefiting from robot cooperation.

Our most interesting finding involved the relation between performance and switching of attention among the robots. In both the manual and mixed initiative conditions participants divided their attention approximately equally among the robots but in the mixed initiative mode they switched among robots more rapidly. Psychologists (Meiran *et al.* 2000) have found task switching to impose cognitive costs and switching costs have previously been reported (Squire *et al.* 2003; Goodrich *et al.* 2005) for multi-robot control. Higher switching costs might be expected to degrade performance, however in this study; more rapid switching was associated with improved performance in both manual and mixed initiative conditions. We believe that the map component at the bottom of the display helped mitigate

losses in awareness when switching between robots and that more rapid sampling of the regions covered by moving robots gave more detailed information about areas being explored.

The frequency of this sampling among robots was strongly correlated with the number of victims found. This effect, however, cannot be attributed to a change from a control to a monitoring task because the time devoted to control was approximately equal in the two conditions. We believe instead that searching for victims in a building can be divided into a series of subtasks involving things such as moving a robot from one point to another, and/or turning a robot from one direction to another with or without panning or tilting the camera. To effectively finish the searching task, we must interact with these subtasks within their neglect time (Crandall *et al.* 2005) that is proportional to the speed of movement. When we control multiple robots and every robot is moving, there are many subtasks whose neglect time is usually short. Missing a subtask means we failed to observe a region that might contain a victim. So switching robot control more often gives us more opportunity to find and finish subtasks and therefore helps us find more victims. This focus on subtasks extends to our results for movement control which suggest there may be some optimal balance between monitoring and control. If this is the case it may be possible to improve an operator's performance through training or online monitoring and advice.

We believe the control episode observed in this experiment corresponds to a decomposed subtask of the team and the linear relationship between switches and found victims reveals the independent or weak relationship among the subtasks. For a multi-robot system, decomposing the team goal into independent or weakly related sub goals allowing the human to intervene into the sub goals is a potential way to improve and analyze human multi-robot performance. From the view of interface design, the interface should fit the sub goal decomposition (or sub goal template) and help the operator in attaining SA. Under mixed-initiative control condition, the number of found victims is less sensitive to waypoint specification than under manually control condition. The relation between found victims and waypoint specification can be generalized to the relationship between performance and human intervention. The potential of extending the present experiment to a generic HRI sensitivity evaluation methodology deserves a further study in the future. Moreover, the control episode can be used as a unit of human intervention, rather than the traditional counting of control actions or durations.

## 7. References

- Browning B. & Tryzelaar E. (2003) UberSim: A Realistic Simulation Engine for Robot Soccer. In: *Proceedings of Autonomous Agents and Multi-Agent Systems, AAMAS'03*, Australia
- Bruemmer D.J., Few D.A., Boring R.L., Marble J.L., Walton M.C. & Nielsen C.W. (2005) Shared Understanding for Collaborative Control. *IEEE Transactions On Systems, Man, And Cybernetics-Part A: Systems And Humans*, 35, 494-504
- Burke J.L., Murphy R.R., Coovert M.D. & Riddle D.L. (2004) Moonlight in Miami: Field Study of Human-Robot Interaction in the Context of an Urban Search and Rescue Disaster Response Training Exercise. *Human Computer Interaction*, 19, 85-116
- Carpin S., Stoyanov T., Nevatia Y., Lewis M. & Wang J. (2006) Quantitative assessments of USARSim accuracy. In: *Proceedings of PerMIS 2006*

- Carpin S., Wang J., Lewis M., Birk A. & Jacoff A. (2005) High fidelity tools for rescue robotics: Results and perspectives. In: *Robocup 2005: Robot Soccer World Cup IX*, pp. 301-311
- Casper J. & Murphy R.R. (2003) Human-robot interactions during the robot-assisted urban search and rescue response at the World Trade Center. *IEEE Transactions on Systems, Man and Cybernetics, Part B*, 33, 367- 385
- Crandall J.W., Goodrich M.A., Olsen D.R. & Nielsen C.W. (2005) Validating human-robot interaction schemes in multitasking environments. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 35, 438-449
- Cyberbotics Ltd. (2006) Webots. URL <http://www.cyberbotics.com/>
- Endsley M.R. (1996) Automation and situation awareness. In: *Automation and human performance: Theory and applications* (eds. Parasuraman R & Mouloua M), pp. 163-181. Erlbaum, Mahwah, NJ
- Envarli I.C. & Adams J.A. (2005) Task Lists for Human-Multiple Robot Interaction. In: *Proceedings of the 14th IEEE International Workshop on Robot and Human Interactive Communication*
- Fong T.W., Thorpe C. & Baur C. (2001) Collaboration, Dialogue, and Human-Robot Interaction. In: *Proceedings of the 10th International Symposium of Robotics Research*. Springer-Verlag, Lorne, Victoria, Australia
- Gerkey B., Vaughan R. & Howard A. (2003) The Player/Stage Project: Tools for Multi-Robot and Distributed Sensor Systems. In: *Proceedings of the International Conference on Advanced Robotics (ICAR 2003)*, pp. 317-323, Coimbra, Portugal
- Goodrich M., Quigley M. & Cosenzo. K. (2005) Switching and multi-robot teams. In: *Proceedings of the Third International Multi-Robot Systems Workshop*
- Jacoff A., Messina E. & Evans J. (2001) Experiences in deploying test arenas for autonomous mobile robots. In: *Proceedings of the 2001 Performance Metrics for Intelligent Systems (PerMIS) Workshop*, Mexico City, Mexico
- Kirlik A. (1993) Modeling strategic behavior in human automation interaction: Why an 'aid' can (and should) go unused. *Human Factors*, 35, 221-242
- Konolige K. & Myers K. (1998) The Saphira Architecture for Autonomous Mobile Robots. In: *Artificial intelligence and mobile robots: case studies of successful robot systems* (eds. Kortenkamp D, Bonasso RP & Murphy R), pp. 211-242. MIT Press, Cambridge, MA
- Lee J.D. & See K.A. (2004) Trust in Automation: Designing for Appropriate Reliance. *Human Factors*, 46, 50-80
- Lee P., Ruspini D. & Khatib O. (1994) Dynamic simulation of interactive robotic environment. In: *Proceedings of International Conference on Robotics and Automation*, pp. 1147-1152
- Lewis M., Wang J. & Hughes S. (2007) USARSim : Simulation for the Study of Human-Robot Interaction. *Journal of Cognitive Engineering and Decision Making*, 1, 98-120
- Lewis, M. and Wang, J. (2007). Gravity referenced attitude display for mobile robots : Making sense of what we see, *Transactions on Systems, Man and Cybernetics Part A*, 37(1), 94-105.
- Marble J.L., Bruemmer D.J. & Few D.A. (2003) Lessons learned from usability tests with a collaborative cognitive workspace for human-robot teams. In: *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, pp. 448-453

- Marble J.L., Bruemmer D.J., Few D.A. & Dudenhoefter D.D. (2004) Evaluation of supervisory vs. peer-peer interaction with human-robot teams. In: *Proceedings of the 37th Annual Hawaii International Conference on System Sciences*
- Meiran N., Chorev Z. & Sapir A. (2000) Component processes in task switching. *Cognitive Psychology*, 41, 211-253
- Nickerson J.V. & Skiema S.S. (2005) Attention and Communication: Decision Scenarios for Teleoperating Robots. In: *Proceedings of the 38th Annual Hawaii International Conference on System Sciences*
- Nielsen C.N., Goodrich M.A. & Crandall J.W. (2003) Experiments in Human-Robot Teams. In: *Proceedings of the 2002 NRL Workshop on Multi-Robot Systems*
- Nielsen C.W. & Goodrich M.A. (2006) Comparing the Usefulness of Video and Map Information in Navigation Tasks. In: *Proceedings of the 2006 Human-Robot Interaction Conference*, Salt Lake City, Utah
- Olsen D.R. & Wood S.B. (2004) Fan-out: measuring human control of multiple robots. In: *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 231-238. ACM Press, Vienna, Austria
- Parasuraman R., Galster S., Squire P., Furukawa H. & Miller C. (2005) A Flexible Delegation-Type Interface Enhances System Performance in Human Supervision of Multiple Robots: Empirical Studies with RoboFlag. *IEEE Systems, Man and Cybernetics-Part A, Special Issue on Human-Robot Interactions*, 33, 481-493
- Parasuraman R. & Miller C.A. (2004) Trust and etiquette in high-criticality automated systems. *Communications of the ACM*, 47, 51-55
- Scerri P., Xu Y., Liao E., Lai G., Lewis M. & Sycara K. (2004) Coordinating large groups of wide area search munitions. In: *Recent Developments in Cooperative Control and Optimization* (eds. Grundel D, Murphey R & Pandalos P), pp. 451-480. Singapore: World Scientific
- Schurr N., Marecki J., Tambe M., Scerri P., Kasinadhuni N. & Lewis J. (2005) The Future of Disaster Response: Humans Working with Multiagent Teams using DEFACTO. In: *Proceedings of AAAI Spring Symposium on AI Technologies for Homeland Security*
- Sheridan T.B. (2002) *Humans and Automation: System Design and Research Issues*. Human Factors and Ergonomics Society and Wiley, Santa Monica, CA and New York.
- Squire P., Trafton G. & Parasuraman R. (2003) Human control of multiple unmanned vehicles: effects of interface type on execution and task switching times. In: *Proceedings of the 2006 Human-Robot Interaction Conference*, pp. 26-32, Salt Lake City, Utah
- Trouvain B., Schlick C. & Mevert M. (2003) Comparison of a map- vs. camera-based user interface in a multi-robot navigation task. In: *Proceedings of the 2003 International Conference on Robotics and Automation*, pp. 3224-3231
- Trouvain B. & Wolf H.L. (2002) Evaluation of multi-robot control and monitoring performance. In: *Proceedings of the 2002 IEEE Int. Workshop on Robot and Human Interactive Communication*, pp. 111-116
- Wang J., Lewis M. & Gennari J. (2003) A game engine based simulation of the NIST urban search and rescue arenas. In: *Proceedings of the 2003 Winter Simulation Conference*, pp. 1039-1045

- Wang J., Lewis M., Hughes S., Koes M. & Carpin S. (2005) Validating USARsim for use in HRI Research. In: *Proceedings of the Human Factors and Ergonomics Society 49th Annual Meeting*, pp. 457-461
- Woods D.D., Tittle J., Feil M. & Roesler A. (2004) Envisioning human-robot coordination in future operations. *IEEE Transactions on Systems, Man & Cybernetics*, 34, 210-218
- Yanco H.A., Drury J.L. & Scholtz J. (2004) Beyond Usability Evaluation: Analysis of Human-Robot Interaction at a Major Robotics Competition. *Journal of Human-Computer Interaction*, 19, 117-149