

Human Control for Cooperating Robot Teams

Jijun Wang
School of Information Sciences
University of Pittsburgh
Pittsburgh, PA 15260
jjw1@pitt.edu

Michael Lewis
School of Information Sciences
University of Pittsburgh
Pittsburgh, PA 15260
ml@sis.pitt.edu

ABSTRACT

Human control of multiple robots has been characterized by the average demand of single robots on human attention or the distribution of demands from multiple robots. When robots are allowed to cooperate autonomously, however, demands on the operator should be reduced by the amount previously required to coordinate their actions. The present experiment compares control of small robot teams in which cooperating robots explored autonomously, were controlled independently by an operator or through mixed initiative as a cooperating team. 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.

Categories and Subject Descriptors

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

General Terms

Human Factors, Measurement, Experimentation

Keywords

Human-robot interaction, metrics, evaluation, multi-robot system

1. INTRODUCTION

Applications for multirobot 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. Autonomous cooperation among robots, in particular, will be needed because it is these activities [6] that impose the greatest decision making load. In addition to this form of high-level supervision, humans are likely to be called upon to assist with a variety of low-level problems such as sensor failures or obstacles that robots cannot solve on their own [5].

Multiple robots substantially increase the complexity of the operator's task because attention must constantly be shifted among robots in order to maintain situation awareness and exert control. In the simplest case an operator controls multiple independent robots interacting with each as needed. Control performance at this task can be characterized by the average demand of each robot on human attention [4] or the distribution of demands coming from multiple robots [13]. 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 [10, 11] 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 [5] 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 MRS has been more restrictive. Because of the need to share attention between robots, teloperation has only been used for one robot out of a team [15] or as a selectable mode [17]. Some variant of waypoint control has been used in all MRS studies reviewed [15, 4, 22, 21, 17, 20] with differences arising primarily in behavior upon reaching a waypoint. A more fully autonomous mode has typically been included involving things such as search of a designated area [15], travel to a distant waypoint [22], or executing prescribed behaviors [17]. In studies in which robots did not cooperate and had varying levels of individual autonomy [15, 4, 22, 21] (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 [16] fan-out study which found highest performance and lowest (per robot activity) imputed workload for the highest levels of autonomy.

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.

HRI'07 March 10–12, 2007, Arlington, Virginia, USA.

Copyright 2007 ACM 978-1-59593-617-2/07/0003 ...\$5.00.

Table 1: Recent MRS studies

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

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 [17, 20] 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 [10, 11] 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 [19] 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 [4, 16] that human control always contributes may not hold. For these more complex MRS control regimes it will be necessary to account for the arguments of Woods et al. [25] and Kirlik’s [9] 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 behavior. The cognitive effort involved in shifting attention between levels of automation and between robots reported by [20] seems a particularly salient problem for MRS.

The present study investigates human interaction with a cooperating team of robots performing a search and rescue task. It compares performance between autonomous teams, manually controlled robots, and operators interacting with a cooperating team in order to identify the contributions of each to system performance. Table 1 organizes details of recent MRS studies. All were conducted in simulation and most involve navigation rather than search. 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 [1, 14]. Our experiment uses USARsim [23], a high fidelity game engine-based robot simulator we developed to study HRI and multi-robot control. USARsim provides a physics based simulation of robot and environment that accurately reproduces mobility problems caused by uneven terrain [24], hazards such as rollover [23], and provides accurate sensor models for laser rangefinders [3] and camera video [2]. This level

of detail is essential to posing realistic control tasks likely to require intervention across levels of abstraction. Previous studies have not addressed the issues of human interaction with cooperating robot teams within a realistically complex environment. Results from 2D simulation [17, 20], for example, are unlikely to incorporate tasks requiring low-level assistance to robots, while experiments with noncooperating robots [15, 4, 22, 21] miss the effects of this aspect of autonomy on performance and HRI.

2. THE SIMULATOR AND MULTI-ROBOT SYSTEM

The present study used three simulated Activemedia P2-DX robots equipped with Sick laser range finder and ptz camera. We built the MrCS (Multi-robot Control System), a multi-robot communications and control infrastructure with accompanying user interface to conduct these studies. MrCS provides facilities for starting and controlling robots in the simulation, displaying camera and laser output, and supporting inter-robot communication through Machinetta [18]. Machinetta is a distributed multiagent system with state-of-the-art algorithms for plan instantiation, role allocation, information sharing, task deconfliction and adjustable autonomy [18]. 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. MrCS provides Machinetta proxies for robots, human interaction (control), and user interface (display). The robot proxy 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. It also communicates between the simulated robot and other proxies to enable the robot 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 user interacts with the system through the user interface which sends messages to robot proxies and reacts to their responses. Sensor outputs from the camera and laser go directly to the interface without passing through any proxy. The user interface built for MrCS is reconfigurable enabling the user to resize and layout the components. A typical interface configuration is shown in Figure 1. On the left side are the global information components: the Robots List (the upper panel) that shows each team member’s execution state

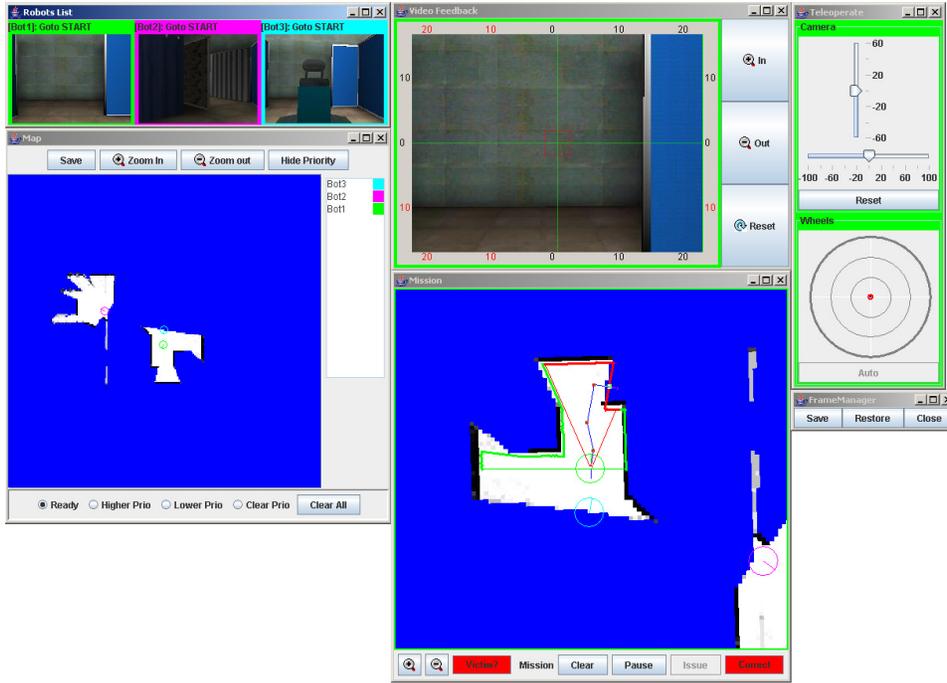


Figure 1: The Graphic User Interface.

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. The local map 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 situation awareness at three ranges. The FOV presented in the red cone shows the operator where he is looking through the camera providing close range SA. Combining this information with the range data shown in the red cone, can give the operator better awareness at medium distances. 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 lower 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 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.

3. METHOD

3.1 Participants

14 paid participants, 19-35, years old were recruited from the University of Pittsburgh community. None had prior experience with robot control although most were frequent computer users. Only two reported playing computer games for more than one hour per week.

3.2 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 minute 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 minute session in Arena-1 followed by a short break and a twenty minute session in Arena-2. At the conclusion of the experiment participants completed a questionnaire.

3.3 Experimental Design

In the experiment, participants were asked to control 3 P2DX robots (Figure 2) simulated in USARsim to search for victims in a damaged building. 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. The participant interacted with the robots through MrCS with fixed user interface shown in Figure 1. Once a victim was identified, the participant marked its location on the map. The testing worlds were simulated versions of the NIST Reference Test Arena, Yellow Arena [8]. Two similar testing arenas 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, semi-transparent boards, and wire grid to add difficulty in situation perception. Bricks, pipes, a ramp, chairs, and other



Figure 2: P2DX robot

debris were put in the arena to challenge mobility and SA in robot control. Figure 2 shows a corner of the testing world. 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 behaviors, 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. RESULTS

In this experiment, we intend to study the interaction between single operator and a robot team in a realistic interactive environment when human and robots must tightly work together to accomplish a task. We at first compare the impact of different level autonomies via evaluating the overall performance from the number of found victims, the explored areas, and the participants’ self-subjective assessments. For the small robot team with 3 robots, we expect similar results as reported in the previous studies [15, 4, 22, 21] that although autonomy will decrease workload, it

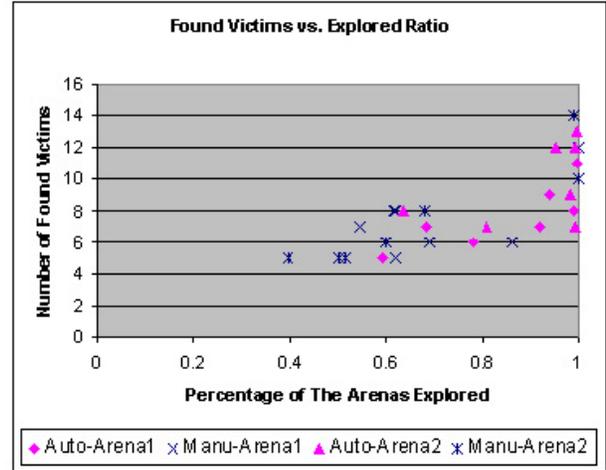


Figure 3: Victims as a function of area explored

will decrease performance as well because of the poorer SA. How human distributes his attention among the robots is an interesting problem especially when human deeply involves in the task via performing some low level functions, such as identifying a victim, in which case human is challenged by balancing between monitoring and control. Therefore, in addition to the overall performance, we assess 1) the distribution of human interactions among the robots and its relationship with the overall performance, and 2) the distribution of control behaviors, 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.

4.1 Overall measurement

All 14 participants found at least 5 of a possible 14 (36%) victims in each of the arenas. The median number of victims found was 7 and 8 for test arena 1 and 2 respectively. Two-tailed t-tests found no difference between the arenas for either number of victims found nor the percentage of the arena explored. Figure 3 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.

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 (Figure 4). 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 (Figure 5). The median number of victims found under full autonomy was 5.

In the posttest survey, 8 of the 14 (58%) participants re-

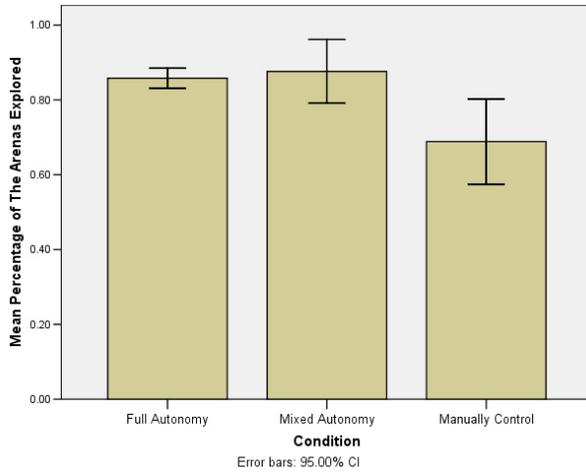


Figure 4: Regions explored by mode

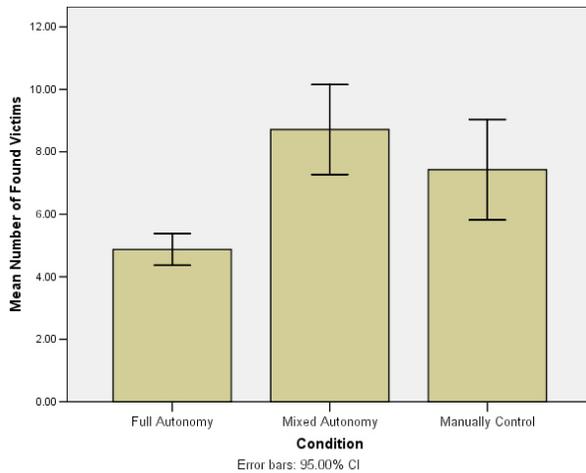


Figure 5: Victims found by mode

ported 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.

4.2 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

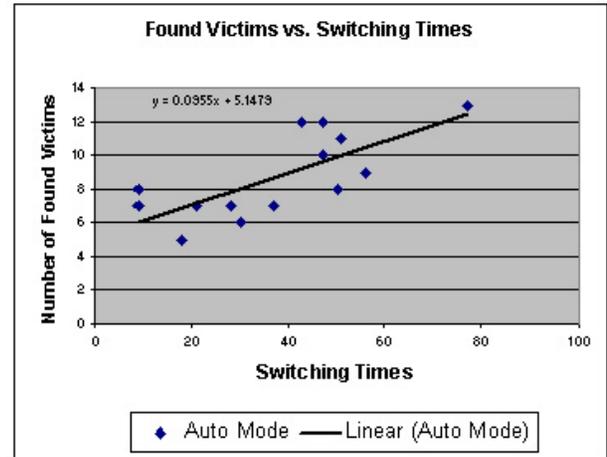


Figure 6: Victims vs. switches under mixed-autonomy mode

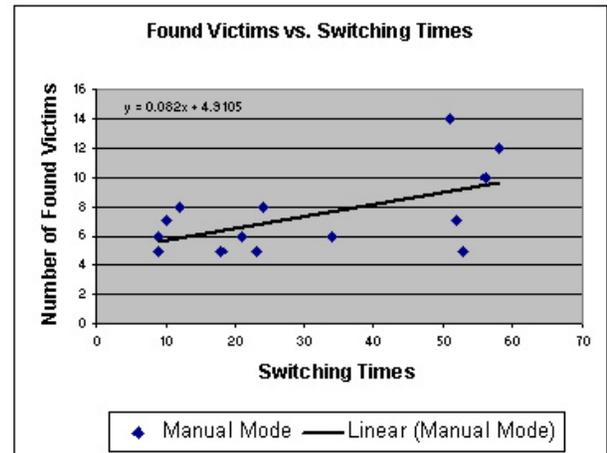


Figure 7: Victims vs. switches under manually control mode

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 (Figures 6 and 7).

An individual robot control episode begins with the pre-observation 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) shows, however, that the proper combination can lead to higher performance.

4.3 Interaction methods

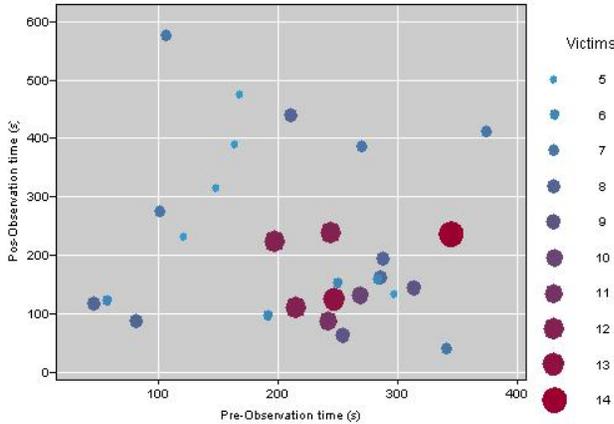


Figure 8: Pre and Post observation time vs. found victims

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 interactions. 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 was used because the operator controls the camera by issuing a desired pose, and monitoring the camera’s movement.

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 4.1).

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. 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

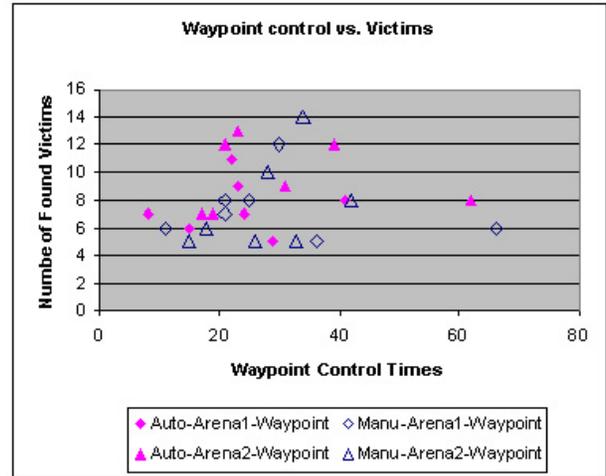


Figure 9: Victims found as a function of waypoint control times

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 [19] 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 [15, 4, 22, 21] 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 [15, 4, 22, 21], 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 [12] have found task switching to impose cognitive costs and switching costs have previously been reported [7, 20] for multirobot 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 com-

ponent 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 [4] 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.

6. REFERENCES

- [1] D. Bruemmer, D. Few, R. Boring, J. Marble, M. Walton, and C. Nielsen. Shared understanding for collaborative control. *IEEE Transactions on Systems, Man, and Cybernetics: A*, 35(4):494–504, July 2005.
- [2] S. Carpin, T. Stoyanov, Y. Nevatia, M. Lewis, and J. Wang. Quantitative assessments of usarsim accuracy. In *Proceedings of PerMIS 2006*, August 2006.
- [3] S. Carpin, J. Wang, M. Lewis, A. Birk, and A. Jacoff. High fidelity tools for rescue robotics: Results and perspectives. In *Robocup 2005: Robot Soccer World Cup IX*, pages 301–311, July 2005.
- [4] J. W. Crandall, M. A. Goodrich, D. R. Olsen, and C. W. Nielsen. Validating human-robot interaction schemes in multitasking environments. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 35(4):438–449, 2005.
- [5] T. W. Fong, C. Thorpe, and C. Baur. Advanced interfaces for vehicle teleoperation: Collaborative control, sensor fusion displays, and remote driving tools. *Autonomous Robots*, 11(1):77–85, July 2001.
- [6] B. Gerkey and M. Mataric. A formal framework for the study of task allocation in multi-robot systems. *International Journal of Robotics Research*, 23(9):939–954, 2004.
- [7] M. Goodrich, M. Quigley, and K. Cosenzo. Switching and multi-robot teams. In *Proceedings of the Third International Multi-Robot Systems Workshop*, March 2005.
- [8] A. Jacoff, E. Messina, and J. Evans. Experiences in deploying test arenas for autonomous mobile robots. In *Proceedings of the 2001 Performance Metrics for Intelligent Systems (PerMIS) Workshop*, Mexico City, Mexico, September 2001.
- [9] A. Kirlik. Modeling strategic behavior in human automation interaction: Why an 'aid' can (and should) go unused. *Human Factors*, 35:221–242, 1993.
- [10] J. Marble, D. Bruemmer, and D. Few. Lessons learned from usability tests with a collaborative cognitive workspace for human-robot teams. In *IEEE International Conference on Systems, Man and Cybernetics*, pages 448–453, October 2003.
- [11] J. Marble, D. Bruemmer, D. Few, and D. Dudenhoeffer. Evaluation of supervisory vs. peer-peer interaction with human-robot teams. In *Proceedings of the 37th Annual Hawaii International Conference on System Sciences*, January 2004.
- [12] N. Meiran, Z. Chorev, and A. Sapir. Component processes in task switching. *Cognitive Psychology*, 41(4):211–253, 2000.
- [13] J. Nickerson and S. Steven. Attention and communication: Decision scenarios for teleoperating robots. In *Proceedings of the 38th Annual Hawaii International Conference on System Sciences*, January 2005.
- [14] C. Nielsen and M. Goodrich. Comparing the usefulness of video and map information in navigation tasks. In *Proceedings of the 2006 Human-Robot Interaction Conference*, Salt Lake City, Utah, March 2006.
- [15] C. Nielsen, M. Goodrich, and J. Crandall. Experiments in human-robot teams. In *Proceedings of the 2002 NRL Workshop on Multi-Robot Systems*, October 2003.
- [16] D. R. Olsen and S. Wood. Fan-out: Measuring human control of multiple robots. In *Proceedings of the 2004 Conference on Human Factors in Computing Systems (CHI 2004)*, pages 231–238, 2004.
- [17] R. Parasuraman, S. Galster, P. Squire, H. Furukawa, and C. Miller. 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*, 35(4):481–493, July 2005.
- [18] P. Scerri, D. Pynadath, L. Johnson, P. Rosenbloom, M. Si, N. Schurr, and M. Tambe. A prototype infrastructure for distributed robot-agent-person teams. In *International Conference on Autonomous Agents*, pages 433–440, Melbourne, Australia, 2003.
- [19] N. Schurr, J. Marecki, M. Tambe, P. Scerri, N. Kasinadhuni, and J. Lewis. The future of disaster response: Humans working with multiagent teams using defacto. In *AAAI Spring Symposium on AI Technologies for Homeland Security*, 2005.
- [20] P. Squire, G. Trafton, and R. Parasuraman. 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*, pages 26–32, Salt Lake City, Utah, March 2003.

- [21] B. Trouvain, C. Schlick, and M. Mevert. 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*, pages 3224–3231, October 2003.
- [22] B. Trouvain and H. L. Wolf. Evaluation of multi-robot control and monitoring performance. In *Proceedings of the 2002 IEEE Int. Workshop on Robot and Human Interactive Communication*, pages 111–116, September 2002.
- [23] J. Wang, M. Lewis, and J. Gennari. A game engine based simulation of the nist urban search and rescue arenas. In *Proceedings of the 2003 Winter Simulation Conference*, pages 1039–1045, December 2003.
- [24] J. Wang, M. Lewis, S. Hughes, M. Koes, and S. Carpin. Validating usarsim for use in hri research. In *Proceedings of the Human Factors and Ergonomics Society 49th Annual Meeting*, pages 457–461, September 2005.
- [25] D. Woods, J. Tittle, M. Feil, and A. Roesler. Envisioning human-robot coordination in future operations. *IEEE Transactions on Systems, Man, and Cybernetics: C*, 34(2):210–218, May 2004.