

# Task Switching and Single vs. Multiple Alarms for Supervisory Control of Multiple Robots

Michael Lewis<sup>1,\*</sup>, Shi-Yi Chien<sup>1</sup>, Siddarth Mehrotra<sup>2</sup>, Nilanjan Chakraborty<sup>2</sup> and Katia Sycara<sup>2</sup>

<sup>1</sup> University of Pittsburgh, School of Information Sciences, Pittsburgh, PA 15260, USA  
ml@sis.pitt.edu, gsechien@gmail.com

<sup>2</sup> Carnegie Mellon University, Robotics Institute, Pittsburgh, PA 15213, USA

siddarthmehrotrall@gmail.com,  
{nilanjan, katia}@cs.cmu.edu

**Abstract.** Foraging tasks, such as search and rescue or reconnaissance, in which UVs are either relatively sparse and unlikely to interfere with one another or employ automated path planning, form a broad class of applications in which multiple robots can be controlled sequentially in a round-robin fashion. Such human-robot systems can be described as a queuing system in which the human acts as a server while robots presenting requests for service are the jobs. The possibility of improving system performance through well-known scheduling techniques is an immediate consequence. Unfortunately, real human-multirobot systems are more complex often requiring operator monitoring and other ancillary tasks. Improving performance through scheduling (jobs) under these conditions requires minimizing the effort expended monitoring and directing the operator's attention to the robot offering the most gain. Two experiments investigating scheduling interventions are described. The first compared a system in which all anomalous robots were alarmed (Open-queue), one in which alarms were presented singly in the order in which they arrived (FIFO) and a Control condition without alarms. The second experiment employed failures of varying difficulty supporting an optimal shortest job first (SJF) policy. SJF, FIFO, and Open-queue conditions were compared. In both experiments performance in directed attention conditions was poorer than predicted. A possible explanation based on effects of volition in task switching is proposed.

**Keywords:** human-robot interaction, neglect tolerance model, scheduling, task-switching

## 1 Introduction

In the simplest case of multirobot control, an operator controls multiple independent robots interacting with each as needed. A foraging task [1] in which each robot searches its own region would be of this category. Control performance at such tasks can be characterized by the average demand of each robot on human attention [2]. Such operator interactions with a robot might be described as a sequence of control episodes in which an operator interacts with the robot for period of time (interaction time, IT) raising its performance above some upper

threshold (UT) after which the robot is neglected for a period of time (neglect time, NT) until its performance deteriorates below a lower threshold (LT) when the operator must again interact with it. In practice the operator's task is even more complex. Humans are additionally included in robotic systems to perform tasks the automation cannot. The most common of these tasks is searching for targets in noisy displays such as remote video or aerial imagery.

Research in robot self-reflection [3] has progressed to the point that it is plausible to presume robots capable of reporting their own off normal conditions such as an inability to move or unsafe attitude. By focusing the operator's attention on robots needing interaction rather than requiring the operator to monitor for the failures, time spent monitoring can be eliminated increasing the number of robots that can be serviced over the intervening interval. With robots informing the operator of their need for interaction the human-robot system becomes more like a queuing system in which the operator acts as the server and robot interaction requests as jobs. Using operations research methods the performance of such a queuing system might be further improved by prioritization of jobs or adjustment of service levels [4] to match current conditions. Deriving full benefit from such aiding, however, would require the ability to focus an operator's attention on a particular robot. We refer to the possibility that human attention might be closely directed in this manner without loss of cognitive efficiency as the *attention scheduling* hypothesis.

Alarms are commonly used in complex human-machine systems to direct human attention but usually in an open and unrestrictive way. Annunciator systems in nuclear power plants or aircraft cockpits typically alarm separately for each setpoint that has been exceeded allowing the human to prioritize and schedule attention among competing demands. Human-multirobot tasks exert similar competing demands on operators frequently requiring them to mix navigation, visual search, and status monitoring to accomplish their objectives. If operators can manage their own attentional resources to avoid damaging interruptions and/or exploit common situational elements among tasks these advantages might outweigh benefits available from externally directed attention.

Experiment I tests the *attention scheduling* hypothesis by comparing operators performing a multirobot foraging task without alarms for robot failures, with all alarms available (Open-queue), or with a first-in-first-out (FIFO) queue making only a single alarm available at a time. Effects were measured for both the primary task of searching for and identifying victims and the secondary task of identifying and restoring failed robots. Because all failures were of the same difficulty, the order in which they were serviced should make no difference so under the *attention scheduling* hypothesis the FIFO and open alarm conditions should produce equivalent performance.


Experiment II extends the test to a condition under which the attention scheduling hypothesis would predict superior performance for directed attention. The shortest job first (SJF) discipline is a provably optimal policy for maximizing throughput in a queuing system [5]. Using this policy to direct human attention, therefore, should lead to superior performance under the *attention scheduling* hypothesis providing the undirected operators did not follow precisely the same policy. This experiment compares SFJ, FIFO, and Open-queue conditions with *attention scheduling* hypothesis

predictions that SJF should produce the best performance followed by Open-queue provided that operators did better than random (FIFO) in selecting robots to be serviced.

## 2 Methods

The reported experiments were conducted using the USARSim robotic simulation with simulated Pioneer P3-AT robots performing an Urban Search and Rescue (USAR) foraging task. USARSim is a high-fidelity simulation of 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' UnrealEngine3 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.

MrCS (Multi-robot Control System), a multi-robot communications and control infrastructure with accompanying user interface, developed for experiments in multi-robot control and RoboCup competition [6] was used in these experiments. MrCS provides facilities for starting and controlling robots in the simulation, displaying multiple camera and laser output, and supporting inter-robot communication.



**Fig. 1.** MrCS Open-queue condition with status bar on left

Figure 1 shows the MrCS user interface in the Open-queue alarm condition. Thumbnails of robot camera feeds are shown on the top, the currently selected video feed of interest in the bottom right. A GUI element in the middle right allows teleoperation and camera pan and tilt. Current locations and paths of the robots are shown on the Map Viewer (middle) which also allows operators to mark victims. The team

status window (left) for the Open-queue condition shows each robot's current status and briefly summarizes any problem. Green indicates the robot is in autonomous condition and functioning safely, yellow indicates an abnormal condition, such as stuck at a corner. When a robot is manually controlled, its tile turns white. The operator selects the robot to be controlled from either the team status window or camera thumbnail.

In forced queue conditions robots in abnormal states are presented one at a time. Additional alarms can only be reviewed after the presenting problem is resolved. To avoid "clogging" the status window with an unrecoverable failure, operators have an alternative in a "Dead" button. Once switched off, the robot will stop reporting and no longer be scheduled. The status panel is removed in the Control condition requiring operators to monitor the Map Viewer and thumbnails to identify malfunctioning robots.

When an operator detects a victim in a thumbnail, a complex sequence of actions is initiated. The operator first needs to identify the robot and select it to see the camera view in a larger window and to gain the ability to stop or teleoperate the robot. After the user has successfully selected a robot, it must be located on the map by matching the window border color or numerical label. Next the operator must determine the orientation of the robot and its camera using cues such as prior direction of motion and matching landmarks between camera and map views. To gain this information the operator may choose to teleoperate the selected robot to locate it on the map, determine its orientation through observing the direction of movement, or simply to get a better viewing angle. The operator must then estimate the location on the map corresponding to the victim in the camera view. If "another" victim is marked nearby, the operator must decide whether the victim she is preparing to mark has already been recorded on the map.

Detecting and restoring a failed robot follow a similar time course: identifying the failed robot on the map and selecting it, then teleoperating it to its next waypoint where the automation can resume control.

The selected USAR environment was an office like hall with many rooms full of obstacles like chairs and desks. Victims were evenly distributed within the environment. Maps were rotated by 90° and each robot entered the environment from different locations on each trial. Because the laser map is built up slowly as the environment is explored and the office like environment provides few distinctive landmarks, there was little opportunity for participants to benefit from prior exposure to the environment. Robots followed predefined paths of waypoints, similar to paths generated by an autonomous path planner [7] to explore the map. All robots traveled paths of the same distance encountering the same number of victims and failures in each designed path. Upon reaching a failure point the operator needed to assume manual control to teleoperate the robot out of its predicament to its next waypoint where autonomous exploration resumed.

### 3 Experiment I

Experiment I reported in [8] compared a Control condition without alarms with two alarm conditions: Open-queue in which all malfunctions were displayed on a status panel and FIFO which displayed alarms one at a time in the order in which they occurred. Because all failures were of the same difficulty the order in which they are serviced should make no difference so according to the attention scheduling hypothesis the FIFO and open alarm conditions should produce equivalent performance. The experiment followed a three condition repeated measures design comparing the conventional MrCS displays with MrCS augmented by alarm panels. Conditions were fully counterbalanced for Map/starting points and display with 5 participants run in each of the six cells

#### 3.1 Participants and Procedure

31 paid participants were recruited from the University of Pittsburgh community balanced among conditions for gender. None had prior experience with robot control although most were frequent computer users. Due to a system crash data was lost for one participant.

After providing demographic data and completing a perspective taking test, participants read standard instructions on how to control robots via MrCS. In the following 15 minute training session, participants 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 began the first 15 minute experimental session in which they performed the search task controlling 6 robots in the first assigned condition. At the conclusion of the session participants were asked to complete the NASA-TLX workload survey [9]. After brief breaks, the next two conditions were run accompanied by repeated workload surveys.

#### 3.2 RESULTS

Data were analyzed using a repeated measures ANOVA comparing search and rescue performance between the control and the two alarmed displays. No difference was found on the overall performance measures areas covered ( $F_{1,29} = .488$ ,  $p = .490$ ), victims found ( $F_{1,29} = .294$ ,  $p = .592$ ), or NASA-TLX workload survey ( $F_{1,29} = 2.557$ ,  $p = .121$ ). Significant effects were found on measures relating to operator strategy and the ways they performed their tasks.

Neglect time (NT) and latency in responding to failures are indicators of operator performance. Long NTs can indicate that some robots may have been ignored while latency in responding to failures can suggest noncompliance with assistance requests or heavy workload at other parts of the task. Robots in the FIFO condition were neglected longer than in the Control condition ( $p = .033$ ,  $SD = 619.507$ ) but did not differ significantly from the Open-queue condition. The neglect times were Open-queue= 1741, FIFO = 1887, and Control = 1629 seconds.

Fault Detection time was defined as the interval between the initiating failure and the selection of the robot involved in that event. Cumulative Fault Detection times were significantly shorter for participants in the Alarm condition,  $p = .021$ , with a cumulative Fault Detection time of 933 seconds. Times for FIFO and Control conditions were 1120, and 1210 seconds respectively. A pairwise T-test shows a significant difference between the Alarm and Control conditions ( $p = .021$ ,  $SD = 607.914$ ).

Victim Delay time was defined as the interval between when a victim first appeared in a robot's camera and the selection of that robot. Victim Delay time again differed across conditions with average times of Open-queue 1303, FIFO 1548, and Control 1559 seconds. A pairwise T-test shows differences between Open-queue and FIFO ( $p = .041$ ,  $SD = 613.725$ ), and Open-queue and Control conditions ( $p = .025$ ,  $SD = 578.945$ ).

A related measure, Select-to-Mark, is defined by the interval between selecting a robot with a victim in view and marking that victim on the map by the process described earlier. Select to mark times can be interpreted as a measure of situation awareness (SA) because they require the operator to orient and interpret the environment. For this measure the results are reversed with users in the Open-queue condition taking the longest times (17.56 sec) and the Control the shortest (14.91 sec) with the FIFO condition (16 sec) again falling in between. There was no overall effect for select to mark time across the three experimental conditions ( $F(1.669,56) = 1.618$ ,  $p = .212$ ). A pairwise T-test, however, shows a difference between Open-queue and Control conditions ( $p = .025$ ,  $SD = 6.02$ ).

## 4 Experiment II

Experiment II reported in [10] extended the investigation begun in Experiment I by introducing multiple types of failures to allow a condition for which the schedule-aiding hypothesis would predict superior performance. Serving the shortest job first (SJF) is a provably optimal policy for maximizing throughput in a queuing system [5]. An alarm system that displayed only the current failure with the shortest time to repair, therefore, should improve the performance of the human-multirobot system over the Open-queue condition unless the unaided human is also following the same SJF policy.


### 4.1 Types of Failures

Recoverable failures were categorized into 4 major types, based on the data for commonly occurring non terminal and field repairable failures for the Pioneer P3-AT [11]. Two of these, camera and map failures, involve loss of display due to communication difficulties. The third, teleoperation lag is a control problem found by [12] to significantly degrade operator performance. The fourth, "stuck", is a common condition in which a robot becomes entangled with obstacles. To resolve encountered failures, the operator needed to manually guide the robot from its current location to the next way-point. Because each of the failure types imposed different difficulties for recovery,

they took varying amounts of time to resolve. In order to estimate typical resolution times for different failures, a pretest using 10 participants was conducted as shown in Table 1 and Figure 2.

**Table 1.** Error Types

Failure	Description
Stuck	Robot was stopped by approaching obstacles
Teleoperation Lagged	Robot executed operator's command with 2~3 seconds delay
Camera Sensor Failed	Robot's video feed will be frozen right before the failure happened
Map Viewer Failed	Robot's position on the map viewer will be unable to update



**Fig. 2.** Time to repair

In the training session, participants practiced control operations for different types of failures for 5 minutes each. Participants were instructed that their goal was to resolve failures by teleoperating to the next indicated waypoint as rapidly as possible. To avoid unrelated delays, such as those associated with switching attention among robots, participants controlled a single designated robot at a time. Because teleoperating the robot to its next waypoint was most easily accomplished by locating both on the map, loss of map indication proved to be the lengthiest failure to repair. The stuck condition which required extensive manual maneuvering and the camera failure that made obstacle avoidance more difficult were the easiest to overcome, with teleoperation delay falling in the middle. This ordering of estimated interaction times allowed failures to be presented to the operator in a priority queue following a shortest job first (SJF) discipline, known to maximize throughput [5].

#### 4.2 Participants and Procedure

The experiment followed a three condition repeated measures design comparing the MrCS augmented by a status panel (Open-queue) with presentation of single alarms following either a FIFO or SJF policy. Thirty paid participants were recruited from the University of Pittsburgh community balanced among conditions for gender. None had prior experience with robot control although most were frequent computer users. Participants read standard instructions on how to control robots via MrCS. In the following 20 minute training session, 5 minutes for each type of failure, participants practiced control operations by resolving failures, three times for each type. 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 began the first 15 minute experimental session in which they performed the foraging task controlling 8 robots in their first assigned condition. Participants had

been told the main task was to locate victims with detecting and resolving robot failures as a secondary task. At the conclusion of the session, participants were asked to complete the NASA-TLX workload survey [9]. After brief breaks, the next two conditions were run accompanied by repeated workload surveys.

### 4.3 RESULTS

#### Victims Found & Distance Traveled.

No difference was found for the number of victims identified ( $F(2,58)=1.10$ ,  $p=.896$ ). Each victim marking was compared to ground truth to determine whether there was actually a victim near the location. If a mark was made further than 2 meters away from any victim or multiple marks for a single victim were found, the marks were counted as false positives. The number of false positives showed a main effect for queue condition ( $F(2,58)=4.637$ ,  $p=.014$ ). A pairwise T-test found a significant difference between Open-queue (1.13 false) and FIFO (2 false) conditions ( $p=.030$ ), as well as a difference between SJF (1.2 false) and FIFO ( $p=.012$ ). No differences were found between Open-queue and SJF.

Unmarked victims that had appeared within a robot's FOV (field of view) without being marked were counted as false negatives (misses). Operators in the Open-queue condition missed the most victims (15) and FIFO the fewest (11) with SJF falling in between (13). A repeated measures ANOVA shows a main effect among queue conditions,  $F(2,58)=20.5$ ,  $p<.001$ . Pairwise T-tests revealed differences between Open-queue and FIFO ( $p<.001$ ), Open-queue and SJF ( $p=.006$ ), and SJF and FIFO ( $p=.003$ ).

No difference was found for the distance traveled ( $F(2,58)=1.73$ ,  $p=.186$ ) although Open-queue (321m) appears slightly better than FIFO (293m) with SJF again in the middle (310m).

Neglect time (NT) and latency in responding to failures again served as indicators of operator performance. Long NTs can indicate that some robots may have been ignored while latency in responding to failures can suggest noncompliance with assistance requests or heavy workload at other parts of the task. NT ( $F(2,58)=1.66$ ,  $p=.20$ ) and the latency in responding to failure ( $F(2,58)=1.75$ ,  $p=.183$ ) were not significantly different among the three conditions. Pairwise T-tests found no difference between Open-queue and FIFO in either Neglect Time ( $p=.086$ ) or fault detection time, ( $p=.079$ ) although Experiment I had found longer NT in the FIFO condition.

The time to service failed robots, measured as the time between selecting the robot and resolving its problem again showed no difference among conditions ( $F(2,58)=.579$ ,  $p=.507$ ), which suggests the types of pre-designed failures were well distributed among three conditions. Overall, FIFO-queue appears slightly worse in the above three measurements.

Select-to-Mark, is defined by the interval between selecting a robot with a victim in view and marking that victim on the map. Select to mark times can be interpreted as a measure of situation awareness (SA) because they require the operator to orient and interpret the environment. A repeated measures ANOVA shows a significant difference among conditions ( $F(2,58)=5.413$ ,  $p=.011$ ). Operators in the



FIFO condition took the longest time (583 sec) and the Open-queue was the shortest (389 sec) with the SJF falling in between (478 sec). A pairwise T-test showed a significant difference between Open-queue and FIFO conditions ( $p=.002$ ), and a marginal difference between Open-queue and SJF ( $p=.061$ ).

The operator must successfully teleoperate the stopped robot from its current location to the next predefined waypoint to resolve a failure. A repeated measures ANOVA showed a significant difference for the count of resolved failures among experimental conditions ( $F(2,58)=5.5$ ,  $p=.006$ ). Participants in the Open-queue condition solved the most failures (17.8), which was significantly more than FIFO ( $p=.003$ ). A pairwise T-test also revealed a difference between SJF, 17 failures, and FIFO 15.7 failures, ( $p=.057$ ).

As in Experiment I the full-scale NASA-TLX workload measure found no advantage among conditions. To examine effects related to the highly prescriptive aiding in FIFO and SJF, we analyzed the frustration scale separately. Repeated measures ANOVA showed a significant difference ( $F(2,58)=5.159$ ,  $p=.009$ ). Pairwise T-tests revealed differences between Open-queue and FIFO ( $p=.038$ ) and between Open-queue and SJF ( $p=.004$ ).

## 5 Discussion

In Experiment I we found that alerting operators to robots in need of interaction improved performance along a number of dimensions. The study compared a control condition without alerting with experimental conditions corresponding to the Open-queue and FIFO conditions of Experiment II. While alerting was beneficial, FIFO which directed the operator to service a particular robot was less effective than the Open-queue which allowed the operator to choose. This contradicts the predictions of the attention scheduling hypothesis which required that human attention be directed without loss of cognitive efficiency. The advantage for less constrained operators might be explained either by superiority of strategies of Open-queue operators when allowed choice or operator difficulties in complying with automation that prescribed the robot to be serviced.

Experiment II partially supported the premise that operator attention can be directed to interaction with individual robots without degrading performance. Open-queue performed slightly better than SJF on false positives, distance traveled, and failures resolved, but only for select-to-mark times did the difference approach significance. For the primary task of marking victims, FIFO participants proved slightly better, however, SJF participants were significantly superior to Open-queue users yielding a balanced performance which was never poorest. The above results may be due to the differences in allocation of attention. Within limited cognitive capacity of processing information, operators have to selectively dedicate attention to any of the "wanted" targets and filter out the irrelevant information simultaneously [16]. Open-queue operators must devote time and attention to monitoring and selection of robots for servicing as well as the interaction leaving less available for the victim monitoring and marking tasks; whereas operators in the forced queue (Priority-/FIFO-

Queue) conditions, by contrast, do not have to compete with monitoring and selecting robots to service leaving more resources available for victim-related tasks, which leads to the reversed results in unmarked victims among three conditions.

The FIFO-queue condition which directed operator attention suboptimally also led to the greatest loss of SA as reflected in its longest Select-to-Mark victims times and lowest marking accuracy. This may have been exacerbated by the FIFO discipline which did not distinguish between distracting recoveries such as loss of track on map and brief interventions such as maneuvering around an obstacle. For the Priority-queue, the SJF discipline had not only the advantage of allowing operators to work primarily on briefer interventions thereby preserving SA, but by clustering similar types of failures increased opportunities for reducing the cost to switch between recovery strategies and sharing the similar cognitive procedures among failures. However, the Priority-queue operators may have simply devoted more of their time and attention to robot requests than operators using the less efficient FIFO because of their greater payoff, which could be observed from the higher rate of unmarked victims.

Table 2 summarizes effects from the two studies. Performance on the primary victim detection and marking task was poorest in the Open-queue condition with more misses and fewer false alarms suggesting operators may have been devoting less effort to this task. When they did see a victim, however, they were faster to select the robot and mark the victim than those using priority queues indicating better situation awareness. This advantage extended to the secondary task where Open-queue users were faster to address and resolve faults. The performance improvements came at some cost, however, as indicated by the elevated frustration scale of the workload measure.

Taken together these experiments fail to confirm the attention scheduling hypothesis as the FIFO and SJF interfaces that dictated the malfunctioning robot to be serviced led to decreased cognitive efficiency as reflected in poorer performance in direct comparisons. Two possible explanations are that: 1) lack of volition in choice of robot to service led to inefficiencies due to task switching [13] and loss of situation awareness due to shifts of attention to potentially remote locations or 2) characteristics of the priority queue conflicted with operator's intentions leading to disuse [14] and hence poorer performance.

In forced queue conditions operators receive an explicit recommendation for the robot to assist. Under extreme stress or time pressured tasks, humans tend to defer to automation and rely on the system for making decisions [15]. This increased compliance under high workload could be especially beneficial to system performance where optimal strategies such as SJF can be used to steer operator attention. Although automated aids can reduce decisional load, they carry little additional information about other robots in need of assistance or the general state of the system. Operators therefore need to regain SA every time they switch to serve a new robot. While working from a forced queue, operators must match the alarmed robots to the thumbnails and/or maps, which could increase the cost in switching attention among failures and robots.

**Table 2.** Summary of Effects

	<b>Experiment I</b>	<b>Experiment II</b>	<b>Effects</b>
<b>Primary Task Performance Measures</b>			
Area Covered	No Effect	No Effect	No Effect
N of Victims Found	No Effect	No Effect	No Effect
False Positives	Not Tested	FIFO > (Open, SJF)	FIFO > (Open, SJF)
Misses	Not Tested	Open > SJF > FIFO	Open > SJF > FIFO
Victim Delay Time	Open < (Control, FIFO)	Not Tested	Open < (Control, FIFO)
Select to Mark Time	Open < Control	Open < (SJF, FIFO)	Open < (SJF, FIFO, Control)
<b>Secondary Task Performance Measures</b>			
Failures Resolved	Not Tested	(Open, SJF) > FIFO	(Open, SJF) > FIFO
Fault Detection Time	Open < Control		Open < Control
<b>Full Task</b>			
Neglect Time	FIFO > Control	No Effect	FIFO > Control
NASA-TLX Workload	No Effect	No Effect	No Effect
Frustration Subscale	Not Tested	Open > (FIFO, SJF)	Open > (FIFO, SJF)

The study results are promising for the prospects of improving HRI performance through scheduling operator attention. The improvement of performance in queuing discipline shows that forced queue aiding can be effectively used by operators and might even lead to superior performance under more complex conditions where choice among robot requests becomes more difficult.

**Acknowledgments.**

This research has been sponsored in part by ONR Grant N0001409-10680.

**References**

1. Y. U. Cao, A. S. Fukunaga, and A. Kahng, "Cooperative mobile robotics: antecedents and directions," *Autonomous Robots*, vol. 4, no. 1, pp. 7-27, 1997.

2. 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 Systems and Humans*, vol. 35, no. 4, pp. 438-449, 2005.
3. M. Scheutz and J. Kramer, "Reflection and Reasoning Mechanisms for Failure Detection and Recovery in a Distributed Robotic Architecture for Complex Robots," *Components*, no. 1, pp. 3699-3704, 2007.
4. Y. Xu, T. Dai, K. Sycara, and M. Lewis, "Service Level Differentiation in Multi-robots Control," *System*, pp. 2224-2230, 2010.
5. M. R. Garey, D. S. Johnson, and R. Sethi, "The Complexity of Flowshop and Jobshop Scheduling," *Mathematics of Operations Research*, vol. 1, no. 2, pp. 117-129, 1976.
6. S. Carpin, M. Lewis, J. Wang, S. Balakirsky, and C. Scrapper, "Bridging the gap between simulation and reality in urban search and rescue," *Robocup 2006 Robot Soccer World Cup X*, vol. 4434, pp. 1-12, 2007.
7. S. Y. Chien, H. Wang, and M. Lewis, "Human vs . Algorithmic Path Planning for Search and Rescue by Robot Teams," *Human Factors*, vol. 54, no. 4, pp. 379-383, 2010.
8. S.-Y. Chien, H. Wang, M. Lewis, S. Mehrotra, and K. Sycara, "Effects of Alarms on Control of Robot Teams," in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 2011, vol. 55, no. 1, pp. 434-438.
9. S. G. Hart and L. E. Staveland, "Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research," in *Human Mental Workload*, vol. 1, P. A. Hancock and N. Meshkati, Eds. North-Holland, 1988, pp. 139-183.
10. S. Y. Chien, S. Mehrotra, M. Lewis, and K. Sycara, "Scheduling Operator Attention for Multi-Robot Control," *Proceedings of the 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2012)*, 2012, pp. 473-479.
11. J. Carlson, R. R. Murphy, and A. Nelson, "Follow-up analysis of mobile robot failures," in *IEEE International Conference on Robotics and Automation 2004 Proceedings ICRA 04 2004*, 2004, vol. 5, no. April, pp. 4987-4994
12. T. B. Sheridan, "Space teleoperation through time delay: review and prognosis," *IEEE Transactions on Robotics and Automation*, vol. 9, no. 5, pp. 592-606, 1993.
13. A. Kiese., M. Steinhäuser, M. Wendt M. Falkenstein, K. Jost, A. Philipp & I. Koch. (2010). Control and interference in task switching--a review, *Psychological Bulletin*, 36(5), 840-874.
14. A. Kirlik, "Modeling strategic behavior in human-automation interaction: why an 'aid' can (and should) go unused.," *Human Factors*, vol. 35, no. 2, pp. 221-242, 1993.
15. T. Inagaki, "Adaptive Automation : Sharing a Trading of control," in *Handbook of Cognitive Task Design*, E. Hollnagel, Ed. 2003.
16. R. Parasuraman, T. B. Sheridan, and C. D. Wickens, "A model for types and levels of human interaction with automation.," *IEEE transactions on systems, man, and cybernetics. Part A, Systems and humans : a publication of the IEEE Systems, Man, and Cybernetics Society*, vol. 30, no. 3, pp. 286-97, May 2000.