

Experiments in Coordination Demand for MultiRobot Systems

Jijun Wang, Huadong Wang, Michael Lewis
University of Pittsburgh
Pittsburgh, PA 15260, USA

Paul Scerri, Prasanna Velagapudi, Katia Sycara
Carnegie Mellon University
Pittsburgh, PA 15213, USA

Abstract—Conventional models of multirobot control assume independent robots and tasks. This allows an additive model in which the operator controls robots sequentially neglecting each until its performance deteriorates sufficiently to require new operator input. This paper presents a model of coordination demand, CD, and experiments intended to extend the neglect tolerance model to situations in which robots must cooperate to perform dependent tasks. In the first experiment operators controlled 2 robot teams to perform a box pushing task under high cooperation demand, teleoperation, moderate demand (waypoint control/ heterogeneous robots), and low demand (waypoint control/homogeneous robots) conditions. In the second experiment participants performed a search and rescue task requiring cooperation between robots creating maps and others carrying cameras. Measured demand and performance were largely consistent with the CD model's predictions.

I. INTRODUCTION

The performance of human-robot teams is complex and multifaceted reflecting the capabilities of the robots, the operator(s), and the quality of their interactions. Recent efforts to define common metrics for human-robot interaction [1] have favored sets of metric classes to measure the effectiveness of the system's constituents and their interactions as well as the system's overall performance. In this paper we present new measures of the demand coordination places on operators of multirobot systems and two experiments to test the usefulness of the measures.

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 likely be needed because it is these activities [2] that theoretically impose the greatest decision making load.

Controlling multiple robots substantially increases the complexity of the operator's task because attention must constantly be shifted among robots in order to maintain situation awareness (SA) and exert control. In the simplest case an operator controls multiple independent robots interacting with each as needed. A search task in which each robot searches its own region would be of this category although minimal coordination might be required to avoid overlaps and prevent gaps in coverage. Control performance at such tasks can be characterized by the average demand of

each robot on human attention [3]. Under these conditions increasing robot autonomy should allow robots to be neglected for longer periods of time making it possible for a single operator to control more robots.

For more strongly cooperative tasks and larger teams individual autonomy alone 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. Estimating the cost of this coordination, however, proves a difficult problem. Established methods of estimating MRS control difficulty, neglect tolerance and fan-out [3] are predicated on the independence of robots and tasks. In neglect tolerance the period following the end of human intervention but preceding a decline in performance below a threshold is considered time during which the operator is free to perform other tasks. If the operator services other robots over this period the measure provides an estimate of the number of robots that might be controlled. Fan-out works from the opposite direction, adding robots and measuring performance until a plateau without further improvement is reached. Both approaches presume that operating an additional robot imposes an additive demand on cognitive resources. These measures are particularly attractive because they are based on readily observable aspects of behavior: the time an operator is engaged controlling the robot, interaction time (IT), and the time an operator is not engaged in controlling the robot, neglect time (NT).

II. MEASURING COORDINATION DEMAND

To separate coordination demand (CD) from the demands of interacting with independent robots we have extended Crandall's [3] neglect tolerance model by introducing the notion of occupied time (OT) as illustrated in Figure 1.

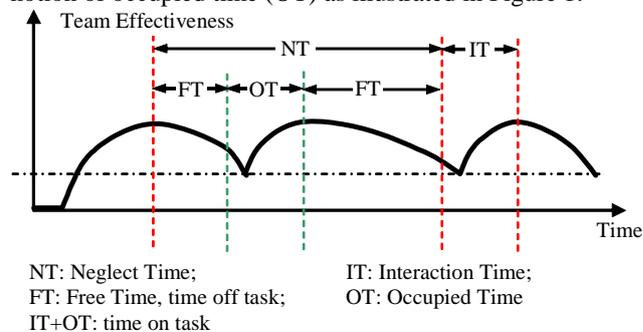


Figure 1. Extended neglect tolerance model for cooperative robot control. The neglect tolerance model describes an operator's interaction with multiple robots as a sequence of control

episodes in which an operator interacts with a robot for period IT raising its performance above some upper threshold after which the robot is neglected for the period NT until its performance deteriorates below a lower threshold when the operator must again interact with it. To accommodate dependent tasks we introduce OT to describe the time spent controlling other robots in order to synchronize their actions with those of the target robot. The episode depicted in Figure 1 starts just after the first robot is serviced. The ensuing FT preceding the interaction with a second dependent robot, the OT for robot-1 (that would contribute to IT for robot-2), and the FT following interaction with robot-2 but preceding the next interaction with robot-1 together constitute the neglect time for robot-1. Coordination demand, CD , is then defined as:

$$CD = 1 - \frac{\sum FT}{NT} = \frac{\sum OT}{NT} \quad (1)$$

Where, CD for a robot is the ratio between the time required to control cooperating robots and the time still available after controlling the target robot, i.e.; the portion of a robot's free time that must be devoted to controlling cooperating robots. Note that OT_n associated with robot_n is less than or equal to NT_n because OT_n covers only that portion of NT_n needed for synchronization. A related measure, team task demand (TAD), adds IT 's to both numerator and denominator to provide a measure of the proportion of time devoted to the cooperative task; either performing the task or coordinating robots.

Most MRS research has investigated homogeneous robot teams where additional robots provide redundant (independent) capabilities. Differences in capabilities such as mobility or payload, however, may lead to more advantageous opportunities for cooperation among heterogeneous robots. These differences among robots in roles and other characteristics affecting IT , NT , and OT introduce additional complexity to assessing CD . Where tight cooperation is required as in the box-pushing experiment, task requirements dictate both the choice of robots and the interdependence of their actions. In the more general case requirements for cooperation can be relaxed allowing the operator to choose the subteams of robots to be operated in a cooperative manner as well as the next robot to be operated. This general case of heterogeneous robots cooperating as needed characterizes the types of field applications our research is intended to support. To accommodate this more general case the Neglect Tolerance model must be further extended to measure coordination between different robot types. This leads to a definition of CD_i , coordination demand for robots of type i as:

$$CD_i = \frac{\sum_{type \neq i} IT}{\frac{\sum_{j=1}^{m_i} l_{ij}}{\max_{j=1}^{m_i} (l_{ij})} T_i - \sum_{type=i} IT} \quad (2)$$

where l_{ij} is the number of times an interaction occurs for robot j of type i and T_i is the total control time for robots of type i .

One approach to investigating coordination demand is to design experiments that allow comparisons between "equivalent" conditions with and without coordination demands. Wang and Lewis [4], for example, compared search performance between a team of autonomously coordinating robots, manually (waypoint) controlled robots, and mixed initiative teams with autonomously coordinated robots that accepted operator inputs. The impact of coordination demand was observable through the difference in performance between the manually controlled teams and the mixed initiative ones. The fully automated teams provided a control ensuring that the benefits in the mixed initiative condition were not due solely to the superior performance of the automation. While [4] examined coordination demand indirectly by comparing performance between conditions in which it was filled either manually or through automation, the present experiments attempt to manipulate and measure coordination demand directly. In the first experiment robots perform a box pushing task in which CD is varied by control mode and robot heterogeneity. The second experiment attempts to manipulate coordination demand by varying the proximity needed to perform a joint task in two conditions and by automating coordination within subteams in the third.

III. USARSIM AND MRCS

Both experiments were conducted in the high fidelity USARSim robotic simulation environment we developed as a simulation of urban search and rescue (USAR) robots and environments intended as a research tool for the study of human-robot interaction (HRI) and multi-robot coordination. USARSim is freely available and can be downloaded from www.sourceforge.net/projects/usarsim. USARSim uses Epic Games' UnrealEngine2 to provide a high fidelity simulator at low cost. 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. MrCS (Multi-robot Control System), a multirobot communications and control infrastructure with accompanying user interface developed for experiments in multirobot control and RoboCup competition [5] was used with appropriate modifications in both experiments. MrCS provides facilities for starting and controlling robots in the simulation, displaying camera and laser output, and supporting inter-robot communication through Machinetta [6] a distributed multiagent system. The distributed control enables us to scale robot teams from small to large.

IV. EXPERIMENT 1

A. Experimental Design

Finding a metric for cooperation demand (CD) is difficult because there is no widely accepted standard. In this

experiment, we investigated CD by comparing performance across three conditions selected to differ substantially in their coordination demands. We selected box pushing, a typical cooperative task that requires the robots to coordinate, as our task. We define CD as the ratio between occupied time (OT), the period over which the operator is actively controlling a robot to synchronize with others, and FT+OT, the time during which he is not actively controlling the robot to perform the primary task. This measure varies between 0 for no demand to 1 for maximum demand. When an operator teleoperates the robots one by one to push the box forward, he must continuously interact with one of the robots because neglecting both would immediately stop the box. Because the task allows no free time (FT) we expect CD to be 1. However, when the user is able to issue waypoints to both robots, the operator may have FT before she must coordinate these robots again because the robots can be instructed to move simultaneously. In this case CD should be less than 1. Intermediate levels of CD should be found in comparing control of homogeneous robots with heterogeneous robots. Higher CD should be found in the heterogeneous group since the unbalanced pushes from the robots would require more frequent coordination. In the present experiment, we measured CDs under these three conditions.



Figure 2. Box pushing task

Figure 2 shows our experiment setting simulated in USARSim [7]. The controlled robots were either two Pioneer P2AT robots or one Pioneer P2AT and one less capable three wheeled Pioneer P2DX robot. Each robot was equipped with a GPS, a laser scanner, and a RFID reader. On the box, we mounted two RFID tags to enable the robots to sense the box's position and orientation. When a robot pushes the box, both the box and robot's orientation and speed will change. Furthermore, because of irregularities in initial conditions and accuracy of the physical simulation the robot and box are unlikely to move precisely as the operator expected. In addition, delays in receiving sensor data and executing commands were modeled presenting participants with a problem very similar to coordinating physical robots.

We introduced a simple matching task as a secondary task to allow us to estimate the FT available to the operator. Participants were asked to perform this secondary task as

possible when they were not occupied controlling a robot. Every operator action and periodic timestamped samples the box's moving speed were recorded for computing CD.

A within subject design was used to control for individual differences in operators' control skills and ability to use the interface. To avoid having abnormal control behavior, such as a robot bypassing the box bias the CD comparison, we added safeguards to the control system to stop the robot when it tilted the box.

The operator controlled the robots using a distributed multi-robot control system (MrCS) shown in Figure 3. On the left and right side are the teleoperation widgets that control the left and right robots separately. The bottom center is a map based control panel that allows the user to monitor the robots and issue waypoint commands on the map. On the bottom right corner is the secondary task window where the participants were asked to perform the matching task when possible.

B. Participants and Procedure

14 paid participants, 18-57 years old were recruited from the University of Pittsburgh community. None had prior experience with robot control although most were frequent computer users.

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 using the MrCS. In the following 8 minutes training session, the participant practiced each control operation and tried to push the box forward under the guidance of the experimenter. Participants then performed three testing sessions in counterbalanced order. In two of the sessions, the participants controlled two P2AT robots using teleoperation alone or a mixture of teleoperation and waypoint control. In the third session, the participants were asked to control heterogeneous robots (one P2AT and one P2DX) using a mixture of teleoperation and waypoint control. The participants were allowed eight minutes to push the box to the destination in each session. At the conclusion of the experiment participants completed a questionnaire about their experience.

C. Results

Figure 4 shows a time distribution of robot control commands recorded in the experiment. As we expected no free time was recorded for robots in the teleoperation condition and the longest free times were found in controlling homogeneous robots with waypoints. The box speed shown on Figure 4 is the moving speed along the hallway that reflects the interaction effectiveness (IE) of the control mode. The IE curves in this picture show the delay effect and the frequent bumping that occurred in controlling heterogeneous robots revealing the poorest cooperation performance.

None of the 14 participants were able to perform the secondary task while teleoperating the robots. Hence, we uniformly find TAD=1 and CD=1 for both robots under this condition. Within participants comparison found that under waypoint control the team attention demand in heterogeneous robots is significantly higher than the demand in controlling

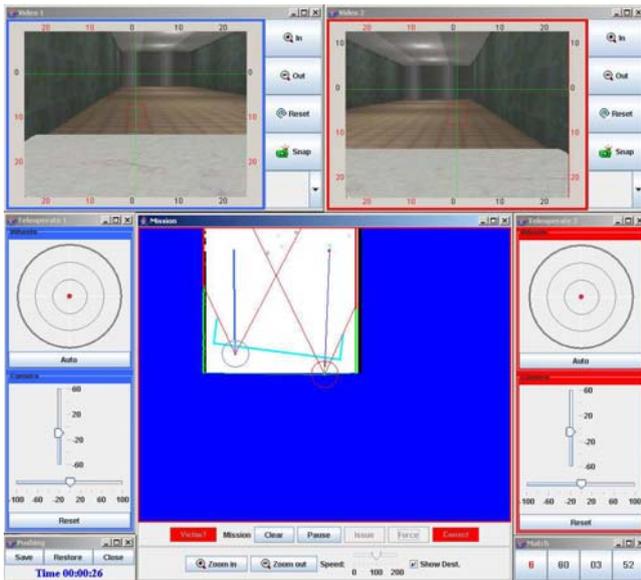


Figure 3. GUI for multirobot control

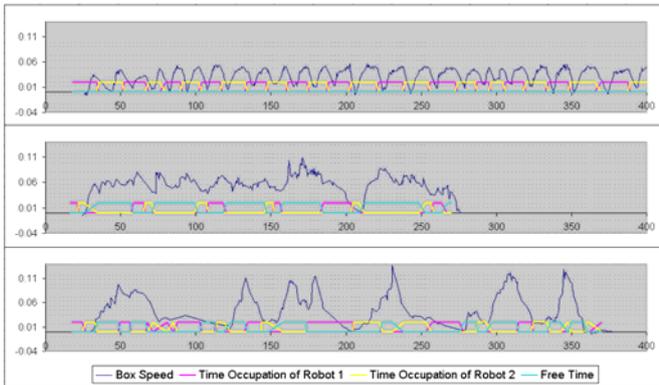


Figure 4 The time distribution curves for teleoperation (upper) and waypoint control (middle) for homogeneous robots, and waypoint control (bottom) for heterogeneous robots

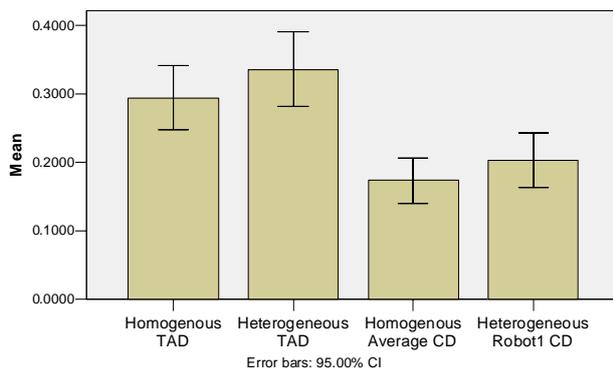


Figure 5. Team task demand (TAD) and Cooperation demand (CD)

homogeneous robots, $t(13)=2.213$, $p=0.045$ (Figure 5). No significant differences were found between the homogeneous P2AT robots in terms of the individual cooperation demand ($P=0.2$). Since the robots are identical, we compared the average CD of the left and right robots with the CDs measured under heterogeneous condition. Two-tailed t-test shows that

when a participant controlled a P2AT robot, lower CD was required in homogeneous condition than in the heterogeneous condition, $t(13)=-2.365$, $p=0.034$. The CD required in controlling the P2DX under heterogeneous condition is marginally higher than the CD required in controlling homogenous P2ATs, $t(13)=-1.868$, $p=0.084$ (Figure 5). Surprisingly, no significant difference was found in CDs between controlling P2AT and P2DX under heterogeneous condition ($p=0.79$). This can be explained by the three observed robot control strategies: 1) the participant always issued new waypoints to both robots when adjusting the box's movement, therefore similar CDs were found between the robots; 2) the participant tried to give short paths to the faster robot (P2DX) to balance the different speeds of the two robots, thus we found higher CD in P2AT; 3) the participant gave the same length paths to both robots and the slower robot needed more interactions because it tended to lag behind the faster robot, so lower CD for the P2AT was found for the participant. Among the 14 participants, 5 of them (36%) showed higher CD for the P2DX contrary to our expectations.

V. EXPERIMENT 2

To test the usefulness of the CD measurement for a weakly cooperative MRS, we conducted an experiment assessing coordination demand using an Urban Search And Rescue (USAR) task requiring high human involvement [8] and of a complexity suitable to exercise heterogeneous robot control. In the experiment participants were asked to control explorer robots equipped with a laser range finder but no camera and inspector robots with only cameras. Finding and marking a victim required using the inspector's camera to find a victim to be marked on the map generated by the explorer. The capability of the robots and the cooperation autonomy level were used to adjust the coordination demand of the task. The experiment was conducted in simulation using USARSim and MrCS.

A. Experimental design

Three simulated Pioneer P2AT robots and 3 Zergs [5], a small experimental robot were used. Each P2AT was equipped with a front laser scanner with 180 degree FOV and resolution of 1 degree. The Zerg was mounted with a pan-tilt camera with 45 degree FOV. The robots were capable of localization and able to communicate with other robots and control station. The P2AT served as an explorer to build the map while the Zerg could be used as an inspector to find victims using its camera. To accomplish the task the participant must coordinate these two types robot to ensure that when an inspector robot finds a victim, it is within a region mapped by an explorer robot so the position can be marked.

Three conditions were designed to vary the coordination demand on the operator. Under condition 1, the *explorer* had 20 meters detection range allowing inspector robots considerable latitude in their search. Under condition 2, scanner range was reduced to 5 meters requiring closer

proximity to keep the inspector within mapped areas. Under condition 3, explorer and inspector robots were paired as subteams in which the explorer robot with a sensor range of 5 meters followed its inspector robot to map areas being searched. We hypothesized that CDs for explorer and inspector robots would be more even distributed under condition-2 (short range sensor) because explorers would need to move more frequently in response to inspectors' searches than in condition-1 in which CD should be more asymmetric with explorers exerting greater demand on inspectors. We also hypothesized that lower CD would lead to higher team performance. Three equivalent damaged buildings were constructed from the same elements using different layouts. Each environment was a maze like building with obstacles, such as chairs, desks, cabinets, and bricks with 10 evenly distributed victims. A fourth environment was constructed for training. Figure 6 shows the simulated robots and environment.

A within subjects design with counterbalanced presentation was used to compare the cooperative performance across the three conditions. The same control interface shown in Figure 7 allowing participants to control robots through waypoints or teleoperation was used in all conditions.

B. Participants

19 paid participants, 19-33, years old were recruited from the University of Pittsburgh community. None had prior experience with robot control although most were frequent computer users. 6 of the participants (31.5%) reported playing computer games for more than one hour per week.

C. Procedure

After collecting demographic data the participant read standard instructions on how to control robots via MrCS. In the following 15~20 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 three testing sessions in counterbalanced order with each session lasting 15 minutes. At the conclusion of the experiment participants completed a questionnaire.

D. Results

Overall performance was measured by the number of victims found, the explored areas, and the participants' self-assessments. To examine cooperative behavior in finer detail, CDs were computed from logged data for each typerobot under the three conditions. We compared the measured CDs between condition 1 (20 meters sensing range) and condition 2 (5 meters sensing range), as well as condition 2 and condition 3 (subteam). To further analyze the cooperation behaviors, we evaluated the total attention demand in robot control and control action pattern as well. Finally, we introduce control episodes showing how CDs can be used to identify and diagnose abnormal control behaviors.



Figure 6 The robots and environment

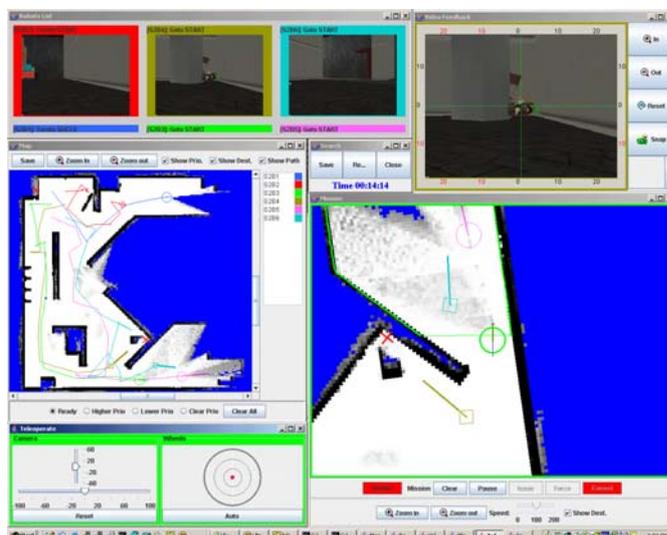


Figure 7 The GUI

1) Overall performance

Examination of data showed two participants failed to perform the task satisfactorily. One commented during debriefing that she thought she was supposed to mark inspector robots rather than victims. After removing these participants a paired t-test shows that in condition-1 (20 meters range scanner) participants explored more regions, $t(16) = 3.097$, $p = 0.007$, as well as found more victims, $t(16) = 3.364$, $p = 0.004$, than under condition-2 (short range scanner). In condition-3 (automated subteam) participants found marginally more victims, $t(16) = 1.944$, $p = 0.07$, than in condition-2 (controlled cooperation) but no difference was found for the extent of regions explored.

In the posttest survey, 12 of the 19 (63%) participants reported they were able to control the robots although they had problems in handling some interface components, 6 of the 19 (32%) participants thought they used the interface very well, and only one participant reported it being hard to handle all the components on the user interface but still maintained she was able to control the robots. Most participants (74%)

thought it was easier to coordinate inspectors with explorers with long range scanner. 12 of the 19 (63%) participants rated auto-cooperation between inspector and explorer (the subteam condition) as improving their performance, and 5 (26%) participants though auto-cooperation made no difference. Only 2 (11%) participants judged team autonomy to make things worse.

2) Coordination effort

During the experiment we logged all the control operations with timestamps. From the log file, CDs were computed for each type robot according to equation 2. Figure 8 shows a typical (IT,FT) distribution under condition 1 (20 meters sensing range) in the experiment with a calculated CD for the explorer of 0.185 and a CD for the inspector of 0.06. The low CDs reflect that in trying to control 6 robots the participant ignored some robots while attending to others. The CD for explorers is roughly twice the CD for inspectors. After the participant controlled an explorer, he needed to control an inspector multiple times or multiple inspectors since the explorer has a long detection range and large FOV. In contrast, after controlling an inspector, the participant needed less effort to coordinate explorers.

Figure 9 shows the mean of measured CDs. We predicted that when the explorer has a longer detection range, operators would need to control the inspectors more frequently to cover the mapped area. Therefore a longer detection range should lead to higher CD for explorers. This was confirmed by a two tailed t-test that found higher coordination demand, $t(18) = 2.476$, $p = 0.023$, when participants controlled explorers with large (20 meters) sensing range.

We did not find a corresponding difference, $t(18)=.149$, $p=0.884$, between long and short detection range conditions for the CD for inspectors. This may have occurred because under these two conditions the inspectors have exactly the same capabilities and the difference in explorer detection range was not large enough to impact inspectors' CD for explorers. Under the subteam condition, the automatic cooperation within a subteam decreased or eliminated the coordination requirement when a participant controlled an inspector. Within participant comparisons shows that the measured CD of inspectors under this condition is significantly lower than the CD under condition 2 (independent control with 5 meters detection range), $t(18) = 6.957$, $p < 0.001$. Because the explorer always tries to automatically follow an inspector, we do not report CD of explorers in this condition.

As auxiliary parameters, we evaluated the total attention demand, i.e. the occupation rate of total interaction time in the whole control period, and the action pattern, the ratio of control times between inspector and explorer, as well. Paired t-test shows that under long sensing conditions, participants interacted with robots more times than under short sensing which implies that more robot interactions occurred. The mean action patterns under long and short range scanner conditions are 2.31 and 1.9 respectively. This means that with

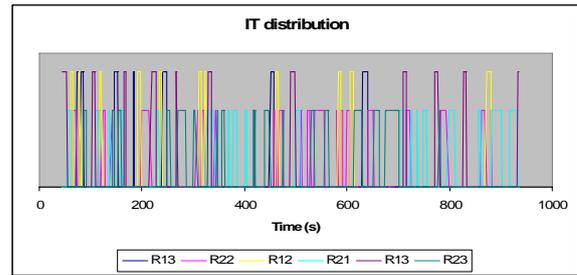


Figure 8 Typical (IT,FT) distribution (higher line indicates the interactions of explorers).

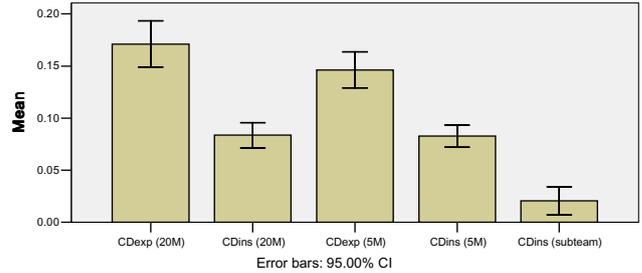


Figure 9 CDs for each robot type

20 and 5 meters scanning ranges, participants controlled inspectors 2.31 and 1.9 times respectively after an explorer interaction. Within participant comparisons shows that the ratio is significantly larger under long sensing condition than under short range scanner condition, $t(18) = 2.193$, $p = 0.042$.

3) Analyzing Performance

As an example of applying CDs to analyze coordination behavior, Figure 10 shows the performance over explorer CD and total attention demand under the 20 meters sensing range condition. Three abnormal cases A, B, and C can be identified from the graph. Associating these cases with recorded map snapshots, we observed that in case A, one robot was entangled by a desk and stuck for a long time, for case B, two robots were controlled in the first 5 minutes and afterwards ignored, and in case C, the participant ignored two inspectors throughout the entire trial.

VI. DISCUSSION

We proposed an extended Neglect Tolerance model to allow us to evaluate coordination demand in applications where an operator must coordinate multiple robots to perform dependent tasks. Results from the first experiment that required tight coordination conformed closely to our hypotheses with the teleoperation condition producing $CD=1$ as predicted and heterogeneous teams exerting greater demand than homogenous ones. The CD measure proved useful in identifying abnormal control behavior revealing inefficient control by one participant through irregular time distributions and close CDs for P2ATs under homogeneous and heterogeneous conditions (0.23 and 0.22), a mistake with extended recovery time (41 sec) in another, and a shift to a satisficing strategy between homogeneous and heterogeneous

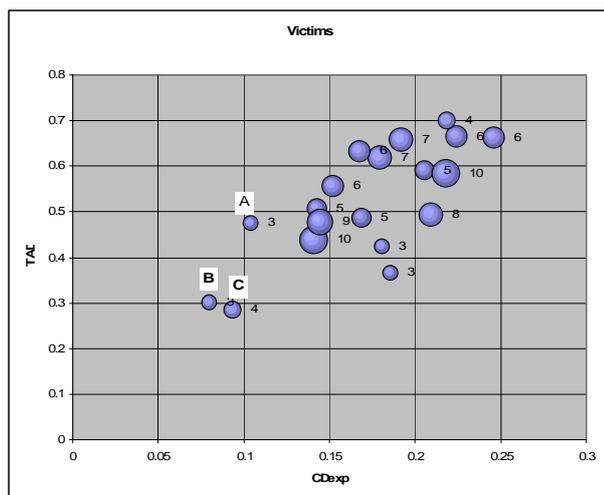


Figure 10. distribution of victims found over CDexp and TAD (total attention demand).

conditions revealed by a drop in CD (0.17 to 0.11) in a third. As most target applications such as construction or search and rescue require weaker cooperation among heterogeneous platforms the second experiment extended NT methodology to such conditions. Results in this more complex domain were mixed. Our findings of increased CD for long sensor range may seem counter intuitive because inspectors would be expected to exert greater CD on explorers with short sensor range. Our data show, however, that this effect is not substantial and provide an argument for focused metrics of this sort which measure constituents of the human-robot system directly. Moreover, this experiment also shows how CD can be used to guide us to identify and analyze aberrant control behaviors.

We anticipated a correlation between found victims and the measured CDs. However, we did not find the expected relationship in this experiment. From observation of participants during the experiment we believe that high level strategies, such as choosing areas to be searched and path planning, had a significant impact on the overall performance. The participants had few problems in learning to jointly control *explorers* and *inspectors* but needed time to figure out effective strategies for performing the task. Because CD measures control behaviors not strategies these effects were not captured. These experiments have demonstrated the utility of measuring the process of human-robot interaction as well as outcomes to diagnosing operator performance and identifying aspects of the task, particularly for multiple robots, that might benefit from automation. While our approach to measuring CD was supported in both experiments the second experiment suggests the need for more sophisticated measures that can take into account strategies and patterns of actions as well as their durations.

ACKNOWLEDGMENT

This work was supported in part by the Air Force Office of Scientific Research under Grant FA9550-07-1-0039 and Office of Naval Research Grant N00014-03-1-0580.

REFERENCES

- [1] Steinfeld, A., Fong, T., Kaber, D., Lewis, M., Scholtz, J., Schultz, A. and Goodrich, M. (2006). Common Metrics for Human-Robot Interaction, *2006 Human-Robot Interaction Conference*, ACM, March, 2006.
- [2] 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.
- [3] 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.
- [4] J. Wang and M. Lewis (2007) Human control of cooperating robot teams, *2007 Human-Robot Interaction Conference*, ACM.
- [5] Balakirsky, S., Carpin, S., Kleiner, A., Lewis, M., Visser, A., Wang, J., and Zipara, V. Toward heterogeneous robot teams for disaster mitigation: Results and performance metrics from RoboCup Rescue, *Journal of Field Robotics*, 2007, in press.
- [6] Scerri, P., et al., Coordinating large groups of wide area search munitions, in *Recent Developments in Cooperative Control and Optimization*, D. Grundel, R. Murphey, and P. Pandalos, Editors. 2004, Singapore: World Scientific. p. 451-480.
- [7] Lewis, M., J. Wang, and S. Hughes, USARSim : Simulation for the Study of Human-Robot Interaction. *Journal of Cognitive Engineering and Decision Making*, 2007. 1(1): p. 98-120.
- [8] Murphy, R.R. and J.L. Burke, Up from the Rubble: Lessons Learned about HRI from Search and Rescue, in *Proceedings of the 49th Annual Meetings of the Human Factors and Ergonomics Society*. 2005: Orlando.