Managing Autonomy in Robot Teams: Observations from Four Experiments

Michael A. Goodrich
Computer Science Dept.
Brigham Young University
Provo, Utah, USA
mike@cs.byu.edu

Timothy W. McLain, Jeffrey D. Anderson, Jisang Sun
Mechanical Engineering Dept.
Brigham Young University
Provo, Utah, USA
tmclain@et.byu.edu

Jacob W. Crandall*
Dept. of Aeronautics and Astronautics
MIT
Cambridge, Mass., USA
jwcrandall@gmail.com

ABSTRACT

It is often desirable for a human to manage multiple robots. Autonomy is required to keep workload within tolerable ranges, and dynamically adapting the type of autonomy may be useful for responding to environment and workload changes. We identify two management styles for managing multiple robots and present results from four experiments that have relevance to dynamic autonomy within these two management styles. These experiments, which involved 80 subjects, suggest that individual and team autonomy benefit from attention management aids, adaptive autonomy, and proper information abstraction.

Categories and Subject Descriptors

H.5.m [**Miscellaneous**]

General Terms

Human Factors

Keywords

Human-Robot Interaction, Teams, Adjustable Autonomy, Dynamic Autonomy

1. INTRODUCTION

Much current research focuses on allowing a single human to manage a team of robots. From this research, there appear to be two paradigms that are most frequently discussed: the sequencing management style and the playbook management style.

*J. W. Crandall was a graduate student at BYU during this research

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 8-11, 2007, Arlington, Virginia, USA.

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

In the *sequencing style*, the human sequentially attends to individual robots giving them instructions and then neglecting them for a period of time. Each robot takes a turn receiving full attention from the human, and the human must schedule neglect and interaction intervals to avoid conflicts [9, 11]. At the extreme of very sophisticated autonomy, the problem becomes similar in spirit to air traffic control [18]; at this extreme, the human's role is less about directing robot behavior and more about deconflicting and sequencing activities [24]. At the other extreme of very simple autonomy, the problem is one of teleoperating one robot while other robots wait for attention.

In the *playbook-style*[1], the human manages clusters or subteams of agents, and issues high level directives that the agents implement in a coordinated manner. The human must plan and select relevant plays, identify agents to cluster into subteams, and act in a supervisory role to determine failures and intervene when necessary. At one extreme, the human gives high-level directives that are performed by a swarm of low-intelligence robots [2]. At the other extreme, the human identifies one play from a set of possible plays; this play requires close coordination among or strong role distinctions between the robots.

Another area of current research is addressing how adjustable autonomy can be utilized to support a human who must manage a team of robots under conditions of dynamic workload in a changing environment [13, 6]. The motivations for doing so are twofold. First, the realization that robots often act in a volatile environment suggests that a static "autonomy level" will not be sufficient to maintain high levels of performance. Second, adjusting autonomy has known benefits for increasing the engagement level of humans, as compared to purely supervisory interaction [13].

Typically, research in adjustable or dynamic autonomy addresses only the sequential management style. We believe that it is desirable to apply dynamic autonomy to support each of the two management styles described above. This concept is illustrated in Figure 1. On the x-axis, the autonomy level of a single robot can be shifted. By moving to the right, the robot can be neglected for a longer period of time but may not allow as high a level of performance and may require longer interaction intervals [21, 20, 7]. On the

[1]The use of the term *playbook* in HRI-related domains was coined by Miller [15]. Because his term is so descriptive, we use it to describe the general class of management styles.

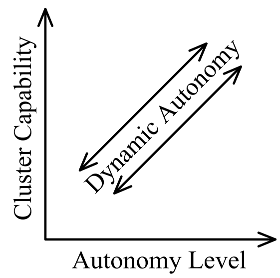

Figure 1: Dynamic autonomy under sequential and playbook autonomy paradigms.

y-axis, the capabilities of a cluster of robots can be shifted. By moving up, more robots can be included in a cluster and the coordination required among these robots becomes more critical. Such cluster capability can include the ability to retask robots from sequential to playbook-style interactions and the capability of performing team autonomy. This latter concept is operationally defined as the ability of team members to autonomously coordinate their activities to perform some shared goal. Dynamic autonomy for large multi-robot teams should consider shifting individual autonomy and cluster capabilities to balance workload and maintain acceptable levels of performance. When designing technology to support dynamic autonomy within this framework, the way information is presented to a human must also be considered. We refer to this information as *Information Support* indicating the need to support both playbook and sequential management styles.

The goal of this paper is to explore how people manage robots under different autonomy levels and different coordination capabilities. To accomplish this goal, we describe four experiments that address fundamental aspects of managing a team of robots. The relevance of the experiments to dynamic autonomy with information support are diagrammed in Figure 2. The first experiment primarily addresses the need to provide explicit support for attention management when autonomy increases and sequential management is used. The second experiment addresses the need to allow playbook management style when robot autonomy is insufficient to allow humans to meet workload demands. The third experiment explores how team autonomy can improve performance and how it benefits from the adaptive autonomy of individual robots. The fourth experiment explores how team-level management may allow an operator to neglect certain types of information that appear critical in a sequential-only style interaction.

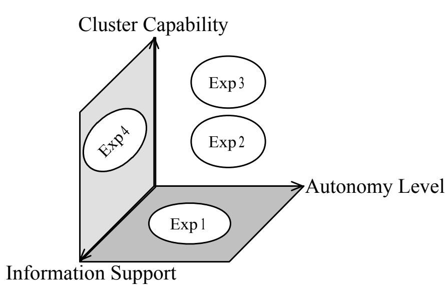

Figure 2: Projecting experiments onto the dynamic autonomy design space.

2. RELATED LITERATURE

Developing autonomy for multiple robots has received careful attention in the literature. Examples include task sharing [23], sensor network coverage [5], and coordinated exploration [4]. When humans are introduced to manage such teams, there are a variety of approaches that can be used. At the swarm level of control, one approach blends behavior-based autonomy with human input [1]. At the other extreme, there is a large body of work on air traffic control [19].

There are a number of studies that explore autonomy between these two extremes. Wang et al. explored how humans manage multiple robots in a search and rescue domain under either manual or coordinated control; their results strongly favored coordinated control [26]. Sellner et al. performed a similar study where four autonomy configurations, including two variations of sliding autonomy, were managed by a human working on a construction task with a team of heterogeneous robots [12]. In this study, the tradeoffs between time to completion, quality of behavior, and operator workload were strongly evident. This result emphasizes the importance of using dynamic autonomy when the world is complex and varies over time. Kaminka and Elmaliach explored how making coordination between robots explicit can reduce failures and improve consistency, in contrast to traditional interfaces [14]. This result complements the observations in our fourth experiment where presenting information relevant to coordination quality improves performance. Finally, Barber et al. discuss how allowing the decision-making framework of a team of agents can benefit from dynamic adaptive autonomy [3].

There is an extensive literature on dynamic autonomy. In the interest of space, we note the three concepts that have most influenced this paper. The first concept is Sheridan's definition of 10 levels of automation [25], extended more recently to include various aspects of human machine interaction [22]. A study that applies such levels in a problem where humans must manage multiple assets using adjustable autonomy is presented in [13]. Miller et al. have written extensively on playbook-style interactions; see for example [15].

The work in this paper relies heavily on descriptions of neglect tolerance and interface efficiency [7]. Applying these concepts to multi-robot control has produced a mathematical bound on the number of robots that a single operator can manage [21, 20]. This model has been extended to include factors that affect situation awareness and attentional conflicts [16].

3. EXPERIMENT 1: ATTENTION

In the first experiment, we explore how different attention aids affect how a human manages a team of robots under sequential-style management when the robots can use different levels of autonomy. Attention management is critical because achieving maximum fan-out [20, 7] requires timely attention to avoid penalties due to wait times [16].

The experiment was a secondary task experiment. The primary task was for a human to guide a robot through a maze to a goal using the map-based interface shown in Figure 3. When the goal was reached, that goal disappeared

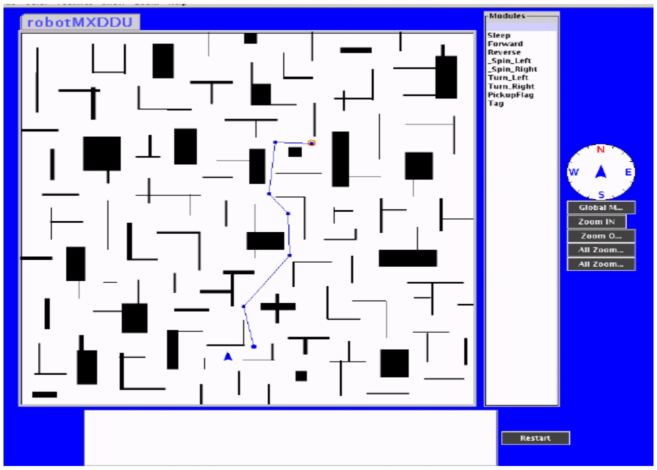

Figure 3: Simple interface for guiding a robot through a maze.

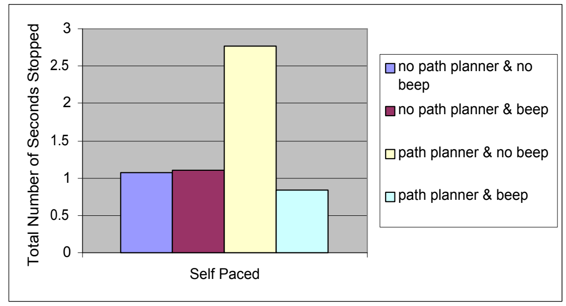

Figure 4: Time required to notice that a robot is stopped.

The secondary task was to solve two-digit addition and subtraction problems that appeared to the side of the screen by selecting the correct answer from four choices. Subjects were informed whether they selected the correct answer or not. The experiment was performed with simulated robots, but all autonomy levels were validated[2] on physical robots.

Subjects were instructed to guide the robot to the series of goals as quickly as possible while solving as many math problems as possible. The entire map with all obstacles was presented to the operator at all times. The experiment was 2x2 corresponding to two types of autonomy and two types of attention support. The two types of autonomy were: (a) teleoperation and (b) path-based, where the operator dropped a series of waypoints that led the robot to the goal. Importantly, the autonomy failed occasionally meaning that the robot would get stuck and require human attention. When the robot became stuck or when it reached a goal, it required operator assistance. Two conditions were explored: (a) an audible alarm paired with a flashing icon indicating a stuck robot, and (b) no notification.

This experiment design is an abstract type of sequential-style management. The primary robot control task represents periods where an individual robot requires complete attention from the operator. The secondary math task represents periods where a robot was neglected because the operator turns attention to other tasks. Average results from 16 subjects are illustrated in Figure 4. The data represent the amount of time that the robot is stopped, either because it is stuck or because it has reached its terminal waypoint. Note that when the level of autonomy is high (waypoint instead of teleoperation), in the absence of a prompt from the interface, subjects are significantly slower ($p < 0.05$) at detecting and managing a stuck robot. When the interface prompts the operator, this effect disappears.

The conclusion from this experiment is that higher levels of autonomy make it possible for an operator to become engaged in another task. This result, which replicates similar findings from the general human factors literature [17], is not surprising but it has implications for adjustable autonomy. In Experiment 3, we will see that allowing the robot to adapt autonomy when it gets stuck contributes to improved performance. A caution is, however, in order. Similar to the way that false alarms can substantially reduce the effectiveness of warnings [10], automatically changing autonomy introduces the risk of making unwarranted adaptations and thereby reducing or negating the benefit of this technology.

[2]For this and subsequent experiments, “validation” means that the algorithms were applied to physical robots and demonstrated to produce analogous behaviors.

4. EXPERIMENT 2: WORKLOAD

In the second experiment, we observe that people sometimes adapt to high-workload situations by employing playbook-style management. Some of the results from this experiment were originally presented in [8], but the emphasis of that paper was on developing a model for predicting the performance of a team. In this paper, we discuss when the model failed and why.

In the experiment, neglect impact and interface efficiency curves were identified for two types of autonomy: (a) a form of teleoperation (P) and (b) region-of-interest (R), which allows an operator to specify a general target location of the robot but requires the robot to autonomously perform obstacle avoidance and path planning to reach this location. Expected performance declines as the amount of time the robot is neglected grows, and expected performance increases as the amount of time the human interacts with the robot increases. Thus, by controlling the *interaction rate* (including both frequency and duration of interaction), it is possible to find that interaction rate which maximizes the expected team performance. This is very important under dynamic autonomy because some autonomy levels experience a large performance hit when either world complexity or operator workload increase, and others do not. For example, the performance of teleoperation in this experiment was negatively impacted by operator workload (see Figure 5), but region of interest control was not. The data were obtained for simulated robots, but the autonomy was validated on physical robots. Under the assumption that humans will select an interaction rate that will maximize the expected total of individual robot performances, it is possible to predict the behavior of a team of robots.

This model was tested by having humans control a team of three robots with mixtures of the two autonomy levels (all the same level, or two of one level and one of the other). When a robot reached a goal, a new goal would appear at and a new goal appeared.

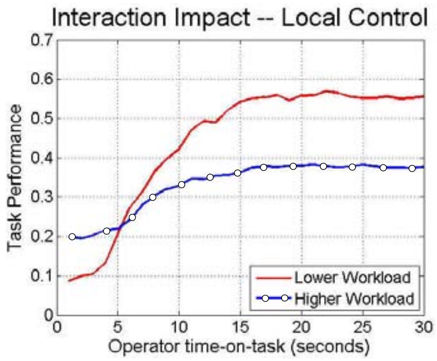

Figure 5: Under teleoperation, more time on task improves performance more when workload is low. Performance is measured as the fraction of the optimal rate of approaching the goal.

a random location in the environment; any robot can go to any goal at any time. Three levels of world complexity were considered. Results for the highest world complexity are

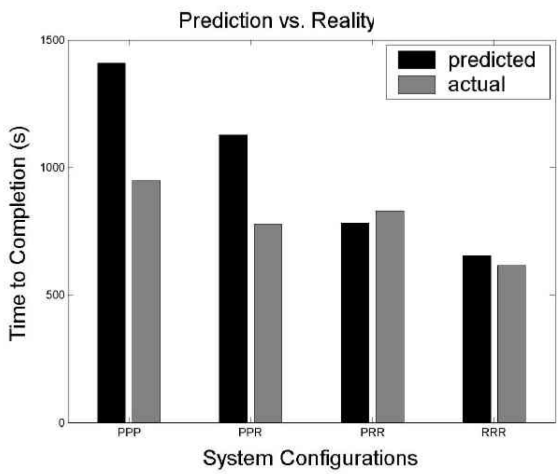

Figure 6: Time required to accomplish the task for a high world complexity.

shown in Figure 6, which is replicated from [8]; results for low and moderate world complexities are not shown because predictions are close to observed behavior. These results use time to complete the task as the inverse measure of team performance; high bars are bad, and low bars are good. 23 subjects participated in the experiment.

The key observation is that the predictions for low levels of autonomy, those with two or more teleoperated robots (denoted by the P's), suggest that people will do worse than they actually do. The reason for this difference is that people do not use sequential-style management as was intended. The experiment was designed under the assumption that the operator would sequentially cycle through robots, guiding them toward a goal. At the highest world complexity, neglect tolerance for the low autonomy levels was not sufficient to allow the operator to sequentially service the robots during the available neglect times.

To compensate, operators qualitatively changed their management paradigm from sequential style to a simple playbook-style; operators abandoned the "one robot to one goal" model and adopted a "zone defense" approach. In this approach, operators would distribute the robots throughout the world. If a goal was near a robot, that robot received guidance toward that goal; goals far from a robot were ignored. When a robot reached a goal and the new goal appeared, the operator evaluated the set of robots and the set of goals, and (re-)assigned a subset of the robots to a subset of goals.

The fact that operators switched strategies when workload became too high yields two observations. First, it confirms the long-identified tendency of real people in real situations to use autonomy in ways that differ from its intended use. Second, it suggests that operators may shift management styles to cope with changing workload conditions. This implies that adjustable autonomy should not restrict itself to changing autonomy levels for individual robots, but should also consider changing management styles as needed.

5. EXPERIMENT 3: COORDINATION

In the third experiment, we explored how people utilize various individual levels and team levels of autonomy to accomplish a mission that cannot be performed by an individual robot alone. This experiment was also performed in simulation, but all team and robot behaviors were verified on omnidirectional physical robots.

Individual robot behaviors were developed from basic move-to-point and path-planning capabilities. These behaviors include: Wander, Follow, and Unstick. The Wander behavior caused a robot to wander around the world randomly. The Follow behavior caused a robot to track the nearest enemy robot within sonar range. This behavior was very sensitive to malfunctions in the sonar sensors. The Unstick behavior was designed to guide the robot in the opposite direction of the nearest obstacle if the robot stopped moving.

In addition to individual robot behaviors, a Team-Move behavior was created. The entire team of robots could be commanded to move to a point in the maze. When this command was issued, the current configuration of the robots became a type of loose formation. The centroid of this formation was used to represent the location of the formation, and the team-move command was interpreted as a desired location for the centroid. Each robot communicated its position to the others and proceeded to a point that would maintain as closely as possible the original formation. Each robot moved through the maze toward its destination along a path specified by the path planner, thus the formation was not maintained while robots were moving. For mazes of at least medium complexity the team-move command caused robots to converge on a point from all different directions.

It is possible to transition between autonomy levels using either adjustable and adaptive autonomy. *Adjustable* autonomy requires the operators to invoke a change in automated behavior. In this experiment, this meant that the operator could choose to change from one mode to another. In contrast, *adaptive* autonomy means that the robot can change behavior without input from the operator. In this experiment, adaptive autonomy included (a) initiating a Wander if no command is issued, (b) switching to Following if an opponent is encountered when Wandering, and (c) using the Unstick behavior (i) while going-to-point (continuing to destination), (ii) while Wandering, or (iii) while Following.

The set of adjustable, adaptive, and team-level autonomy configurations that were used in the experiment are shown in Table 1. The large number of possible configurations makes

Mode	Autonomy
1	Teleoperation
2	Adjustable Autonomy
3	Adjustable + Adaptive Autonomy
4	Adjustable Autonomy + Adaptive Autonomy + (Non-adaptive) Team Autonomy
5	Adjustable + Adaptive Autonomy + (Adaptive) Team Autonomy

Table 1: Categories of individual and team behaviors. Team autonomy means that the robots can execute the Team Move behavior. In non-adaptive team autonomy, the robots stop after completing the assigned move; in an adaptive team move, individual robots automatically select a new behavior.

it difficult to test each possibility in an experiment with human subjects. Reasonable configurations were selected in a pilot experiment, but these particular configurations allow only general observations to be made about the effects of autonomy.

These experiments use the autonomous behaviors according to Table 2. If no autonomy is listed, the default is for the robot to wait for the human to issue a command.

Autonomy	Mode
Wander Command	2, 3, 4, 5
Follow Command	2, 3, 4, 5
Automatically wander if no new command	3, 5
Automatically follow while wandering	3, 4, 5
Unstick while going-to-point	3, 4, 5
Unstick while wandering	3, 4, 5
Unstick while following	2, 3, 4, 5
Team Move	4, 5

Table 2: Exact configurations of individual and team behaviors.

The experiment consisted of a team of three robots, managed by a human, that try to tag three autonomous robot opponents. Opponent robots autonomously wander through the world avoiding other robots. An opponent is tagged if two of the pursuer robots are simultaneously within a small threshold distance from them. The game ends when all opponent robots are tagged. A secondary spatial reasoning task (Tetris) interrupted at random intervals and lasted for a variable amount of time. During these interruption times, the human can neither intervene nor monitor the robots.

Eleven subjects participated in the experiment. As an estimate of performance, we measured the amount of time that at least one robot on the team was within sonar distance of an opponent. This measure estimates instantaneous performance of an individual robot because the nature of the task requires robots to be within a close proximity to their opponents to be successful. Figure 7 shows the average number of sonar contacts per second across subjects. The figure shows this metric both when the operator is performing the

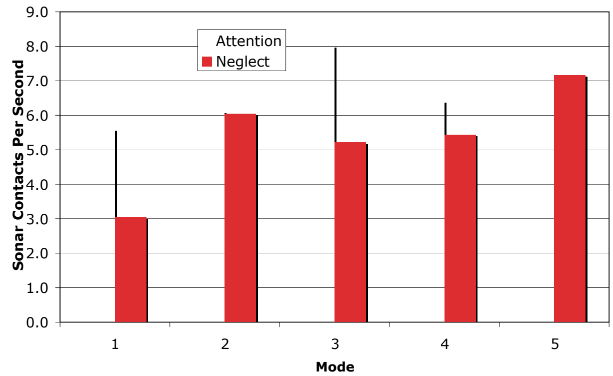

Figure 7: Efficiency of individual autonomy as measured by the frequency that at least one robot is in contact with an opponent.

primary robot-management task and when the operator is performing the secondary task; we refer to these conditions as periods of attention and neglect, respectively.

Three quantitative observations from the experiment are relevant to this paper. First, autonomy tends to improve the performance of individual robots. This is seen by summing the number of sonar contacts per second across the neglect and attention period; modes 5 and 3 are best, followed by 2 and 4, and then 1. Second, good performance by an individual does not necessarily translate into good team performance; in this game, at least two robots must be in proximity to an opponent to tag it. A measure of team performance is the amount of time it takes to complete the game. As shown in Figure 8, adjustable and adaptive autonomy tend to improve team performance.

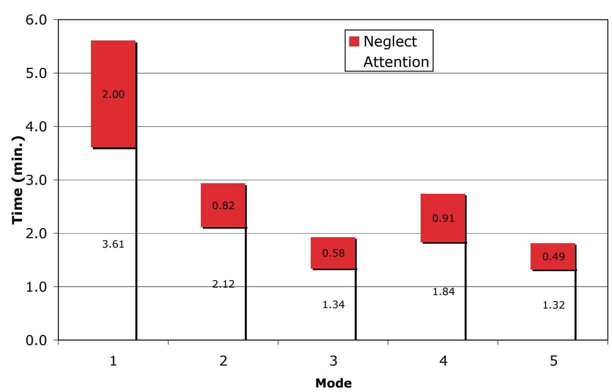

Figure 8: Efficiency of team autonomy as measured by the duration of the game.

Third, autonomy tends to shift what humans do during periods of interaction. This is seen by comparing the relative performance within autonomy modes during attention and neglect periods. Higher relative performance during interaction than during neglect indicates that the autonomy is easy to control but not tolerant to neglect. Higher relative performance during neglect than interaction indicates that the people must work to invoke strategies during periods of interaction, but that these strategies are relatively robust when the team is neglected. Thus, autonomy mode 3, which is a sequential management style, benefits from adjustable and adaptive autonomy by allowing periods of interaction to be maximally efficient. Autonomy mode 5, by contrast, is more of a playbook-style interaction which also includes adaptive autonomy. Interaction periods do not produce as many sonar contacts with the opponents as mode 3, but neglect periods produce many more contacts. This indicates that, although people spend more time planning during interaction periods, the resulting plans are robust during periods of neglect.

Equally important to the objective data, subjective evaluation of the results indicates how people cope with the different autonomy levels. In mode 1, subjects tended to adopt a “zone defense” approach (similar to Experiment 2) by distributing robots uniformly throughout the maze and then driving an enemy toward one of their teammates. In mode 2, subjects tended to adopt a supervisory control strategy of putting all robots into a Wander mode, and then switching a robot’s behavior to a Follow mode when close to an opponent. In mode 3, subjects also tended to use a supervisory control strategy, but waited until a robot automatically began to Follow an opponent before intervening. When a robot began to Follow, they would guide another robot to the area to help. Interestingly, mode 4 tended to induce humans to adopt a very similar behavior to mode 3. Very few subjects used the Team-Move mode; when they did use this mode they were effective during periods of interaction but ineffective during periods of neglect. Mode 5 induced similar behavior, but automatically Wandering or Following allowed more neglect tolerance and caused more subjects to use team autonomy than in Mode 4.

Two conclusions are suggested by these results. First, adjustable and adaptive autonomy, if designed correctly, do allow people to more successfully manage robots under the sequential management style. However, they may perform better if the different agents operate under different autonomy levels, some teleoperated and some using guided autonomy. Second, team-level autonomy may not always improve performance, but it may shift how people spend their time and thus alter workload and neglect tolerance.

6. EXPERIMENT 4: INFORMATION

In the fourth experiment, we explore the type of information that best supports an operator managing a team of robots. More precisely, we explore how timing, cost, and spatial information can be used in playbook-style management.

The type of mission in the experiment is a three-robot coordinated timing mission. Such missions include both (a) simultaneous or sequential rendezvous where robots must meet at a desired point and (b) coordinated strike missions where robots must reach target points at approximately the same time or in a sequence. The mission requires that the robots be deployed from different physical locations and pass through risky terrain to arrive at three different locations simultaneously or in a sequence within some user-specified tolerances. Velocity and path constraints imply that the robots cannot stop and that the robots cannot follow arbitrary paths.

Given a range of velocity and path constraints, a set of individual robot paths from a starting point to an assigned destination are identified. These paths are parameterized by the range of arrival times. For example, suppose that the path planner identifies five paths that satisfy the path constraints. For path 1, the velocity constraints imply that the robot can travel to the goal in no fewer than t_1 seconds and in no more than $\tau_1 > t_1$ seconds. Thus, it is possible for the robot to arrive, using path 1, within the interval (t_1, τ_1). Feasible paths 2-5 have similar intervals. Associated with each path and arrival time is a cost function; paths that expose the robot to risk have a higher cost than paths that do not, and arrival times that expose the robot to risk for longer amounts of time have higher cost than those that do not. If two or more paths can produce the same arrival time, only the path with lowest cost is considered.

The problem is for a human to select a set of paths that satisfy the timing constraint and that minimize team cost. To support the decision-making process of the human, there are three types of information that can be presented: feasible arrival times, path cost, and paths for each robot on the team. This is a playbook-style of interaction since the task of the human is to select a plan that minimizes total cost to the team. The human is given discretion in selecting the total cost and path, since there may be tradeoffs involved in the choice. For example, one robot may be disposable or one target less important, so it may be useful to allow this robot to have greater cost and other robots to have less cost even if the accumulated team cost is higher than another choice.

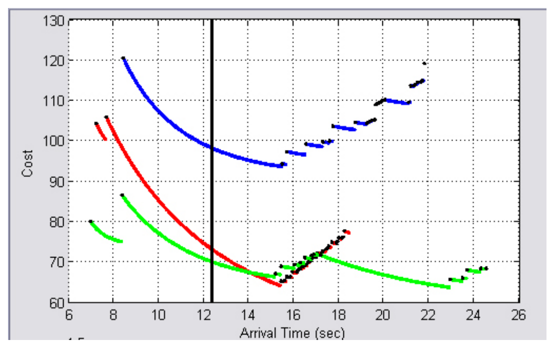

Figure 9: GUI that displays individual costs as a function of feasible arrival times. Discontinuities occur when paths change to satisfy constraints.

We explored the relative benefits of four types of information-presentation configurations. Configuration 1, shown in Figure 9, used McLain’s notion of a coordination function and presented the user with a graph that represented cost versus arrival time. The user can select candidate arrival times and explore the total and relative costs to each robot. Configuration 2, shown in Figure 10, used a standard spatial display that showed areas of risk and possible paths. The user can select or modify a possible path for an individual robot and see the paths of other robots automatically change to paths that satisfy the timing constraints. Configuration 3, which is not shown in a figure because it is so simple, presents only the set of feasible arrival times and allows the operator to select from these times. Configuration 4, shown in Figure 11, presents only information on individual and total team cost. The operator can change the amount of cost for a single robot or the total team costs by dragging the top of the cost bar; team cost and the cost to the other robots automatically change to the minimum feasible path that satisfies the timing constraints.

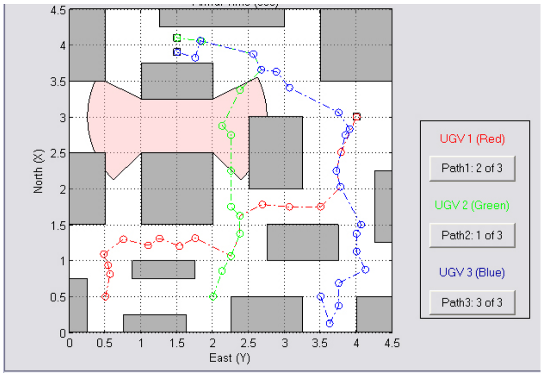

Figure 10: GUI that displays path information and allows the operator to cycle between feasible arrival times by testing different paths for an individual robot.

Thirty subjects participated in the experiment, which was a $4 \times 2 \times 3$ design corresponding to the four types of interface, two maze complexities, and three types of constraints: (i) none, (ii) constrained path in which some regions of the world were forbidden and (iii) constrained cost in which certain vehicles were required to have minimum cost. A full factorial experiment was performed. Dependent measures included response time, workload (both the number of mouse clicks and NASA TLX), and solution quality (in terms of overall team cost). Each subject performed each of the tests which were presented in pseudo-random order.

An ANOVA was performed to evaluate effects. Interestingly, the type of interface yielded no significant response-time difference among four of the six complexity-constraint combination, only making a difference when there was a path constraint. Under the path constraint condition, the spatial-presentation interface reduced response times compared to the other interfaces. Importantly, although there was not generally a significant difference in response time, there was a significant difference both in workload and in performance across the interfaces. Figure 12 indicates that the spatial interface was perceived as significantly harder to

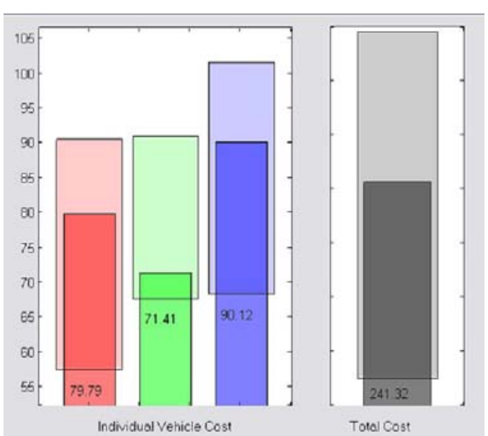

Figure 11: GUI that provides only individual and team costs.

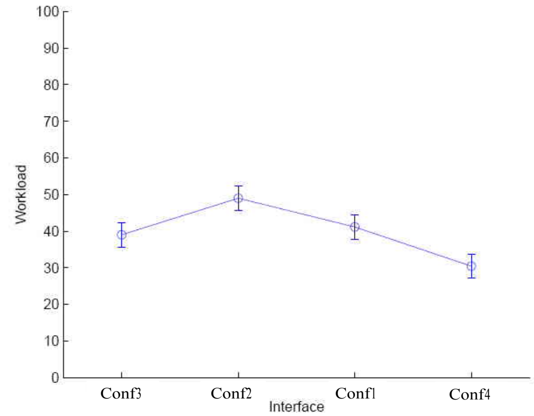

Figure 12: NASA TLX measurements across interface types.

use than the other interfaces as measured by NASA TLX; mouse clicks produced similar results. Although the spatial display required more work, it did not produce better decisions. Recall that decision quality measures the amount of risk exposure experienced by the robots. The spatial interface produced significantly more costly solutions than the other interfaces, even when controlling for maze complexity and path constraints.

The results from this experiment have implication for presenting information to operators using playbook management styles. If there are factors that determine the quality of coordination (such as cost), it is sometimes better to allow operators to explore tradeoffs in this space without giving them direct control over robot behaviors. In this experiment, allowing an operator to control individual robots by altering their paths allowed the operator to alter the execution of the play, but at the cost of worse paths and more workload; this holds even though the other robots automatically adapted their behavior to satisfy the timing constraint. Thus, we conclude that managing clusters of robots may require managing the parameters of coordination rather than the behavior of robots within a particular set of constraints imposed by the selected play. For the coordinated timing mission discussed herein, we suggest that each type of information be available and used as required for the specifics of the task.

7. CONCLUSIONS

It is often desirable to have a human manage multiple robots, but increasing the number of robots can increase the human's workload. One way to mitigate this is to allow adjustable and adaptive autonomy to support both individual and team autonomy. In this paper, we presented four experiments that help focus attention on important issues that arise when we seek to create such systems. Experiment 1 suggests that high robot autonomy frees the operator to become engaged in another task, which might require information to be presented in a way that helps the human manage attention. Experiment 2 suggests that adjusting autonomy should include the ability to change management styles between sequential and playbook (and back) to cope with shifts in workload. Experiment 3 suggests that adjustable and adaptive autonomy have the power to improve both sequential and playbook management. Experiment 4 suggests that team-level performance may be improved by allowing the control of quality parameters (e.g., cost and timing) rather than robot behaviors. An important area of future work is studying how information should be presented when either individual-level or team-level autonomy changes in response to shifts in workload.

8. ACKNOWLEDGEMENTS

Portions of this material are based upon work supported by the National Science Foundation under Grant No. 0534736, by DARPA under Grant No. NBCH1020013, and by ARL under contract DAAD19-02-D-001. The authors thank Joshua Johansen for conducting the first experiment.

9. REFERENCES

[1] K. Ali and R. C. Arkin. Multiagent teleautonomous behavioral control. *Machine Intelligence and Robotic Control*, 1(2):3–10, 2000.

[2] R. C. Arkin and K. Ali. Integration of reactive and telerobotic control in multi-agent robotic systems. In *Proc. 3rd Int'l Conf. on Simulation of Adaptive Behavior*, pages 473–478, 1994.

[3] K. S. Barber, A. Goel, and C. E. Martin. Dynamic adaptive autonomy in multi-agent systems. *Journal of Experimental and Theoretical Artificial Intelligence*, 12(2):129–147, April 2000.

[4] M. A. Batalin and G. S. Sukhatme. Spreading out: A local approach to multi-robot coverage. In *Proc. of the 6th Int'l Symp. on Distributed Autonomous Robot Systems*, pages 373–382, 2002.

[5] M. A. Batalin and G. S. Sukhatme. The analysis of an efficient algorithm for robot coverage and exploration based on sensor network deployment. In *Proc. of IEEE Int'l Conf. on Robotics and Automation*, 2005.

[6] J. M. Bradshaw, M. Sierhuis, A. Acquisti, P. Feltovich, R. Hoffman, R. Jeffers, D. Prescott, N. Suri, A. Uszok, and R. V. Hoof. Agent autonomy. In H. Hexmoor, R. Falcone, and C. Castelfranchi, editors, *Adjustable Autonomy and Human-Agent Teamwork in Practice: An Interim Report on Space Applications*. Kluwer, 2002.

[7] J. W. Crandall, M. A. Goodrich, D. R. O. Jr., and C. W. Nielsen. Validating human-robot interaction schemes in multi-tasking environments. *IEEE Trans. on Systems, Man, and CyberneticsPart A: Systems and Humans*, 35(4), 2005.

[8] J. W. Crandall, C. N. Nielsen, and M. A. Goodrich. Towards predicting robot team performance. In *Proc. of IEEE Int'l Conf. on Systems, Man, and Cybernetics*, pages 906–911, 2003.

[9] M. L. Cummings and S. Guerlain. Developing operator capacity estimates for supervisory control of autonomous vehicles. In review, Presence, 2005.

[10] S. R. Dixon, C. D. Wickens, and D. Chang. Unmanned aerial vehicle flight control: False alarms versus misses. In *Proc. HFES Annual Meeting*, September 20-24 2004.

[11] M. A. Goodrich, M. R. Quigley, and K. A. Cosenzo. Task switching and multi-robot teams. In *Proc. NRL Workshop on Multi-Robot Systems*, 2005.

[12] B. S. F. W. H. L. M. Hiatt, R. Simmons, and S. Singh. Coordinated multi-agent teams and sliding autonomy for large-scale assembly. *Proc. of the IEEE*, 94(7), July 2006.

[13] D. B. Kaber and M. R. Endsley. The effects of level of automation and adaptive automation on human performance, situation awareness and workload in a dynamic control task. *Theoretical Issues in Ergonomics Science*, 5(2):113–153, March-April 2004.

[14] G. A. Kaminka and Y. Elmaliach. Experiments with an ecological interface for monitoring tightly-coordinated robot teams. In *Proc. IEEE Int'l Conf. on Robotics and Automation*, pages 200–205, 2006.

[15] C. A. Miller, H. B. Funk, M. Dorneich, and S. D. Whitlow. A playbook interface for mixed initiative control of multiple unmanned vehicle teams. In *Proc. Digital Avionics Systems Conference*, volume 2, pages 7E4–1 – 7E4–13, 2002.

[16] P. M. Mitchell and M. L. Cummings. Management of multiple dynamic human supervisory control tasks. In *Proc. Int'l Command and Control Research and Technology Symp.*, 2005.

[17] N. Moray. Designing for transportation safety in the light of perception, attention, and mental models. *Ergonomics*, 33(10/11):1201–1213, 1990.

[18] A. Nunes and H. Mogford. Identifying controller strategies that support the 'Picture'. In *Proc. HFES Annual Meeting*, 2003.

[19] A. Nunes and H. Mogford. Identifying controller strategies that support the 'Picture'. In *Proceedings of the 47th Annual Meeting of the Human Factors and Ergonimics Society*, Santa Monica, CA, 2003.

[20] D. R. Olsen and M. A. Goodrich. Metrics for evaluating human-robot interactions. In *Proc. PERMIS*, 2003.

[21] D. R. Olsen, Jr. and S. B. Wood. Fan-out: Measuring human control of multiple robots. In *Proc. Human Factors in Computing systems*, pages 231–238, 2004.

[22] R. Parasuraman, T. B. Sheridan, and C. D. Wickens. A model for types and levels of human interaction with automation. *IEEE Trans. on Systems, Man, and Cybernetics – Part A: Systems and Humans*, 30(3):286–297, May 2000.

[23] L. E. Parker. ALLIANCE: An architecture for fault tolerant, cooperative control of heterogeneous mobile robots. In *Proc. IROS*, volume 2, pages 776–783, 1994.

[24] T. B. Sheridan. *Telerobotics, Automation, and Human Supervisory Control*. MIT Press, 1992.

[25] T. B. Sheridan and W. L. Verplank. Human and computer control of undersea teleoperators. Technical report, MIT Man-Machine Systems Laboratory, 1978.

[26] J. Wang, M. Lewis, and P. Scerri. Cooperating robots for search and rescue. In *Proc. AAMAS*, 2005.