Human Control Strategies for Multi-Robot Teams

Katia Sycara
Robotics Institute
Carnegie Mellon University
5000 Forbes Ave, Pittsburgh PA 15213
U.S.A.
katia@cs.cmu.edu www.cs.cmu.edu/~sycara

Michael Lewis
School of Information Sciences
University of Pittsburgh
135 N Bellefield Ave, Pittsburgh, PA 15260
U.S.A.
ml@sis.pitt.edu www.sis.pitt.edu/~mlewis

Abstract: - Expanding human span of control over teams of robots presents an obstacle to the wider deployment of robots for practical tasks in a variety of areas. One difficulty is that many different types of human interactions may be necessary to maintain and control a robot team. We have developed a taxonomy of human-robot tasks based on complexity of control that helps explicate the forms of control likely to be needed and the demands they pose to human operators. In this paper we use research from two of these areas to illustrate our taxonomy and its utility in characterizing and improving human-robot interaction.

Key-Words: - swarms, human-robot interaction, USAR, multirobot control

1 Introduction

The basic problem of expanding use of unmanned vehicles (UVs) lies in increasing the span of control of UV operators. As the number of UVs increases, the needs to coordinate activities among UVs and provide operator judgment and assistance at crucial points rapidly exceeds our current capabilities. There have been a variety of proposed answers and experimental implementations addressing the problem of scaling operators' span of control. Unfortunately, these point solutions do little to organize the field or provide guidance on where and how any particular approach is likely to work or fail. Over the past five years we have been developing a comprehensive theory of human interaction with multiple UVs that sheds light on these problems. Our theory, based on the complexity of the operator's task, identifies three types of interactions, namely O(1) control where UVs coordinate autonomously and could be controlled as a group/swarm, O(n) control where UVs can be controlled independently, and O(>n) control where the operator must be directly involved in coordinating UV activities. Each of these forms of control poses its own problems. In O(1) control the swarm of autonomously coordinating UVs is difficult to command because, except in very limited scenarios (e.g. commanding the whole swarm to move from its current area to a new particular location), their goals and behaviors have been predetermined. For O(n) control the problem becomes trying to organize the operators' interactions with UVs for greatest efficiency. Earlier experiments [1] suggest that O(>n) control is likely to be extremely difficult even for small N therefore our research focuses on the two more tractable forms of control.

1.1 A Model of Multi Human-Robot Interaction (M-HRI)

In computer science the notion of computational complexity, the time that must be used to solve a problem as a function of the size of its input, has proved fruitful for weeding out bad algorithms. Algorithms with high complexity may work for small problems, but fail or grow inefficient for even slightly larger ones. The task of controlling multiple robots is similar to an algorithm in that the operator must perform a repetitive sequence of decisions and actions to control a robot. If the robots are performing independent activities, the operator can devote the same attention to each in turn, resulting in a complexity of Order n, written O(*n*), because each of the *n* robots requires the same set of actions and the total operator effort is proportional to the number of robots. Another benefit of independence is that more UVs can be controlled simply by adding more operators. A different form of control, such as designating a region to be searched by drawing it on a map, can command an arbitrary number of robots with a single act. Because the number of actions the operator must take are independent of the number of robots, control of this sort is O(1) and has a constant effort. Dependent tasks such as box pushing, by contrast, can be arbitrarily difficult with command complexity, O(>n), because dependencies among robots create cascading demands. When one robot pushes one corner of a box, for example, the operator must control the other robot to push the other corner to straighten its path, after which the first robot needs attention again.

O(1) tasks require substantial autonomy on the part of the robots but impose only a constant demand on the human operator. In general, O(1) control is appropriate where a large number of UVs must be tightly coordinated with a relatively simple goal such as formation following or area search . O(n) tasks, such as approving targets, or identifying victims, are robot-centric tasks that can be performed independently by one or more operators and impose a predictable additive demand. O(>n) tasks, by contrast, cannot be specified simply and, depending on the task, could require arbitrarily large control effort on the part of the operator. Figure 1 illustrates the hypothesized relation between number of robots and their demand on the operator's cognitive resources.

2 O(n) Sequential Control

A wide class of multirobot control tasks involve operator interactions with individual robots. Where the robots' actions are independent, as for example in some foraging tasks, the operator can interact with robots sequentially in a round robin fashion. The Neglect Tolerance model describes repeated interactions of this sort in which the operator raises a robot's performance above a threshold during an interaction period, IT, and then allows the robot's performance to decline over a neglect period, NT. If the need for interaction can be detected by the robot through self-reflection, the robot could communicate its need for interaction to the operator. The resulting human-robot system would be a queuing system in which the operator is the server and the queue of robots requesting interaction, the jobs. As a queuing system, performance might be optimized using standard techniques providing the operator's attention could be appropriately directed.

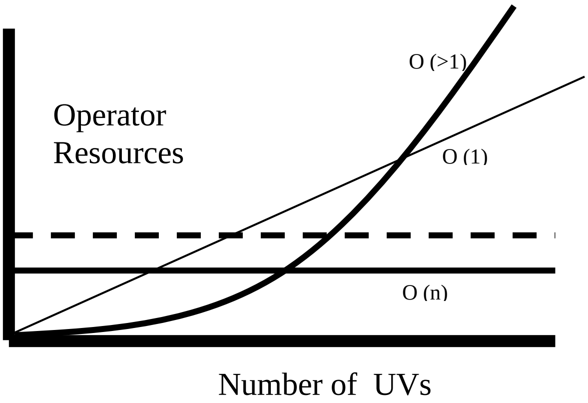

Figure 1 Command Complexity: The figure illustrates the hypothesized relationship between task types and command complexity.

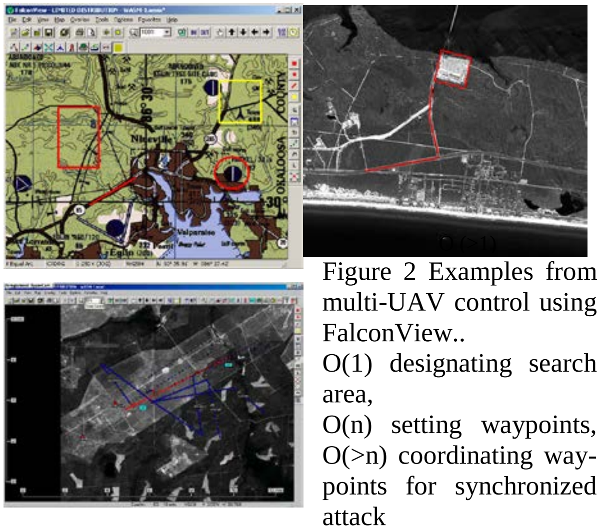

Figure 2 Examples from multi-UAV control using FalconView..
O(1) designating search area,
O(n) setting waypoints,
O(>n) coordinating way-points for synchronized attack

2.1 USAR Experiments

We have conducted a series of studies to identify techniques and conditions necessary for directing operator attention. Our experiments used USARSim, a validated [3,4] high-fidelity simulation of urban search and rescue (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 including video, laser rangefinder, sonar, and audio. In our experiments robots simulated in USARSim were controlled through the MrCS (Multi-robot Control System), a multi-robot communications and control infrastructure with accompanying user interface, developed for experiments in multirobot control and RoboCup competition [5].

Figure 3 shows the MrCS user interface for 8 robots in alarm condition. Thumbnails of robot camera feeds are shown on the top, a 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) in which allows operators to mark victims. The augmented elements, team status window (left) shows each robot's current

condition in different colors and briefly summarizes it. Green color indicates the robot is in autonomous condition, yellow shows the robot is in an abnormal condition, such as stuck at a corner, and when a robot is manually controlled its tile turns white. The operator selects the robot to be controlled from the colored team status window.

The initial study [6] found that HRI performance was improved by communicating requests for interaction to the operator, however, a more directed first-in-first-out (FIFO) display showing only a single request at a time led to poorer performance than one showing the entire (Open) queue of robots reporting difficulties. Because failures were homogeneous and required equal times to repair, the FIFO condition should have produced at least as good performance as the

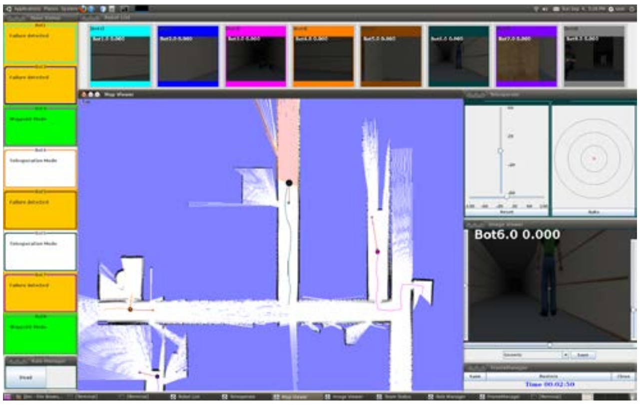

Figure 3. Open queue condition -failures are shown in toolbar on left, In FIFO condition failures are shown one at a time .

Open queue if attention were being efficiently directed. The second experiment incorporated four types of errors with varying times to repair allowing the interface to direct operator attention according to shortest job first (SJF), a discipline proven [7] to maximize throughput.

The results support the conclusion that operator attention can be effectively directed for interaction with individual robots. Based on paired t-tests, SJF tied with Open-queue in besting FIFO performance in false positive identifications ($p=.012$), time to mark victims ($p=.061$), and recovered failures ($p=.057$) and tied with FIFO by missing fewer victims than the Open-queue ($p=.003$).

2.2 Scheduling Algorithms

While our experiments demonstrate that human-multirobot performance can be improved by following SJF, other scheduling policies might improve performance under other conditions. We have developed scheduling algorithms to address two of these situations.

2.2.1 Service Level Differentiation

The neglect tolerance model posits fixed thresholds, yet there is no reason to believe that the mechanism by which interaction improves performance and neglect causes it to decline are limited in this way. An operator, for example, might take more time interacting with a robot during a slack period increasing its supra-threshold performance and neglect it longer usual when busy. We have developed a model [8] relating IT and NT to optimal system performance by allowing the individual thresholds to vary. This increased flexibility not only improves team performance but agrees with human data [9,10] showing performance per robot to decrease smoothly with increasing team size rather than dropping abruptly upon reaching the fan-out threshold.

Our first model for an open queue system allows us to find exact analytic solutions. Our second closed system model is more realistic since it models the inter-dependency between the service process and arrival process. For this model we were only able to find solutions algorithmically. Experimental results comparing system performance for different values of system parameters show that a mixed strategy is a general way to get optimal system performance for a large variety of system parameter settings (e.g.; values of λ, number of robots) and in all cases is no worse than a pure strategy.

2.1.1 Individual Differences & Variation in tasks

Matching limited human resources with multiple heterogeneous UVs is challenging for a variety of reasons: First, human operators staffing different positions may have varying skill levels likely to result in different relations between NTs and ITs. In addition, different tasks may have different service requirements. Second, the human operators may need to make trade-offs between service quality and speed: providing a slower service rate (i.e., a longer IT) increases the service value (i.e., NT) for each UV, but would make UVs wait longer in the queue. Third, task information is stochastic in the sense that the human operator may have no prior knowledge of the types of tasks.

We have addressed these problems through a game-theoretic queuing model in which the robot chooses the time and operator to pose a request [11]. The single-human/multi-robot system is modelled as an open queuing system in which different types of arriving UVs require varying degrees of attention (reservation utility) with differing costs of continuing to operate in their degraded mode (waiting costs). This corresponds to the various forms of degraded performance the operator may need to address. An unmanned ground vehicle (UGV) that has rolled over, for example, may require extensive attention to right and will exert a high waiting cost because it can make no useful contribution while it is immobilized. Another UGV that continues exploring a largely covered area due to difficulty in passing through a narrow aperture may require significantly less operator attention to correct while providing somewhat useful information while operating in its degraded mode. We have found equilibria for this model for both homogeneous and heterogeneous cases varying reservation utilities and waiting costs.

3 O(1) Command of coordinating robots

Automated coordination of UVs is a complex control problem particularly for highly interdependent tasks. The first screen shot shown in Figure 2 illustrates an O(1) command interface for UAVs coordinating using the Machinetta [12] multiagent infrastructure. This form of role-based coordination has been a widely studied in human-multirobot control with systems such as Playbook© [13] and Machinetta. Unfortunately communication requirements slow these architectures prohibitively as the number of robots grows large. While Xu et al. [14] showed that subteams can be effectively formed and coordinated within very large populations; the population itself cannot be effectively coordinated.

The alternative of size-independent swarms whose coordination emerges from interaction among local control laws avoids this bottleneck. Unfortunately, such systems have their behavior “baked in” at design time and are therefore difficult for humans to influence or control at run time.

We have made novel distinctions in order to systematically explore the range of mechanisms proposed for human control of swarms. The first of these is separating the function of connectivity maintenance from algorithmic objective. So flocking behavior [15], for example, which requires matching velocity and heading with neighbors while maintaining a fixed distance would be achieved by combining a primitive connectivity maintenance procedure (maintaining fixed distance) with a consensus algorithm for velocity and heading.

We make a second distinction between influence based on identity, selection, and influence based on location, beacon to accommodate the variety of mechanisms commonly used to influence swarms. Linking influence to identity allows control through splitting the swarm and switching among algorithms and parameters. The beacon mechanism directly implements potential fields and can simulate leader or predator models (actual or virtual) by using beacons that attract or repel nearby robots.

Our initial experiments [16] conducted with 32 participants from the University of Pittsburgh community compared selection and beacon control (distinction 2) for operators who assisted (distinction 3) the swarm in a scavenging task in which the robot team acquired information appearing at random

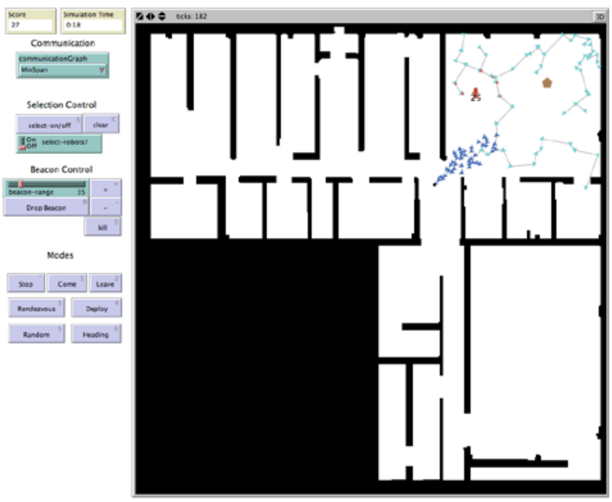

Figure 4. NetLogo interface for swarm experiments

locations. For *beacon* control the operator could place, move, set the mode of, change the range of, and remove beacons. The heading mode requires an additional mouse click to determine the heading. For *selection* control the operator could select a rectangular set of robots, clear the current selection, and set the mode of all robots in the current selection. The come, leave and heading modes require an additional click to determine the target location or direction. Operators were tested across five environments of progressive complexity ranging from completely open to cluttered (many obstacles) and structured (walls and hallways). A task congruent algorithm leaving all robots in the random mode unless they were currently within information range of a source was developed as a performance benchmark.

As Figure 5 shows human assistance switching between generic algorithms and modes of influence was less effective than the task-customized algorithm in simple environments. As environments grew more complex, however, human contributions increased leading to superior performance for control based on selection.

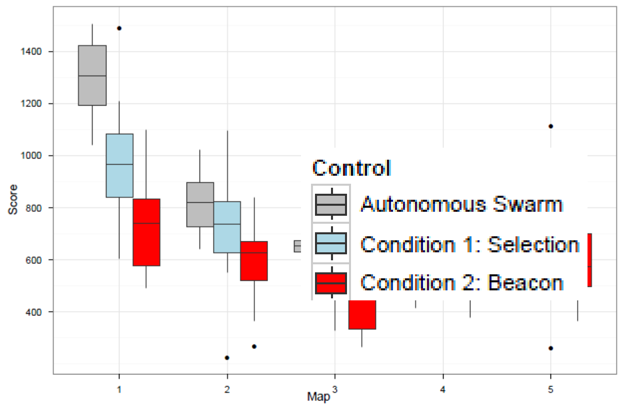

Figure 5. Foraging scores

Participants were free to choose among control algorithms but showed a marked preference for only three. Figure 6 shows frequent use of random walks and attraction and somewhat less frequent use of the deploy dispersion algorithm. Rendezvous, repulsion, and the stop command, by contrast, were rarely used. The key differences between *selection* and *beacons* are their spatial and temporal persistence and the resulting active or passive influence on the robot swarm, enabling different control strategies. Our results showed that novice human operators perform better with selection control. Both types of control enabled human operators to adapt to environments with complex obstacles and their drop in performance was less than that of a simple task congruent algorithm that performed better than human operators in open environments.

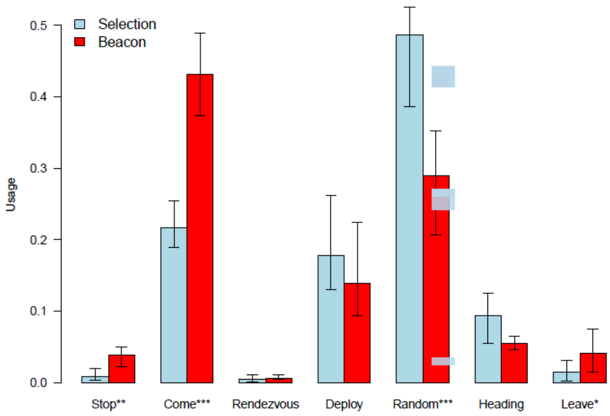

Figure 6. Use of commands

4 Conclusion

In this paper we have presented a taxonomy of human-robot tasks for multi-robot teams and illustrated its utility. In the case of sequential control of independent robots we have shown that human attention can be directed so that sophisticated scheduling algorithms can be used to improve performance. In the case of large teams relying on emergent coordination we have demonstrated that human control using a small set of algorithmic objectives can produce comparable or better performance than a specialized task-congruent algorithm. These examples illustrate both the variety of tasks found in controlling multiple robots and the usefulness of a taxonomy for identifying feasibility and requirements.

References:

[1] J. Wang, M. and Lewis. Assessing coordination overhead in control of robot teams, Proceedings of 2007 *IEEE International Conference on Systems, Man, and Cybernetics*, 2645-2649, 2007.

[2] M. Lewis, K. Sycara, & P. Scerri, Scaling up wide-area-search-munition teams, *IEEE Intelligent Systems*, 24(3), 2009, 10-13.

[3] J. Wang, M. Lewis, S. Hughes, M. Koes, & S. Carpin, Validating USARsim for use in HRI Research, *Proceedings of the Human Factors and Ergonomics Society 49th Annual Meeting*, 2005, 457-461.

[4] S. Carpin, T. Stoyanov, Y. Nevatia, M. Lewis, & J. Wang, Quantitative assessments of USARSim accuracy. *Proceedings of PerMIS* 2006.

[5] S. Balakirsky, S. Carpin, A. Kleiner, M. Lewis, A. Visser, J. Wang, & V. Zipara, Toward hetereogeneous robot teams for disaster mitigation: Results and performance metrics from RoboCup Rescue, *Journal of Field Robotics*, 24(11-12), 2007, 943-967.

[6] S. Chien, S. Mehtra, N. Brooks, K. Sycara & M. Lewis, Effects of Alarms on Control of Robot Teams, *Proceedings of the 55th Annual Meeting Human Factors and Ergonomics Society* (HFES'11), 2011.

[7] M.R. Garey, D.S. Johnson, & R. Sethi, The Complexity of Flowshop and Jobshop Scheduling. *Mathematics of Operations Research*, 1967.

[8] Y. Xu, T. Dai, K. Sycara & MJ. Lewis, Service level differentiation in multi-robots control, *Proceedings of the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS'10), 2010, 2224-2230.

[9] C. Humphrey, and J. Adams, "Assessing the scalability of a multiple robot interface", *Human Robot Interaction Conference* (HRI'07), ACM, 2007, 239-246.

[10] M. Lewis, H. Wang, S. Chien, P. Velagapudi, P. Scerri & K. Sycara. Choosing autonomy modes for multirobot search, *Human Factors*, 52(2), 2010, 225-233.

[11] T. Dai, K. Sycara & M. Lewis, A game theoretic queueing approach to self-reflection in decentralized human-robot interaction systems, *2011 IEEE International Conference on Robotics and Automation* (ICRA'11), 2011.

[12] P. Scerri, D. V. Pynadath, L. Johnson, Rosenbloom P., N. Schurr, M Si, and M. Tambe. A prototype infrastructure for distributed robot-agent-person teams. *The Second International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS'03), 2003.

[13] C. Miller & R. Parasuraman, Designing for flexible interaction between humans and automation: Delegation interfaces for supervisory control. *Human Factors* 49(1), 2007, 57-75.

[14] Y. Xu, P. Scerri, M. Lewis, and K. Sycara (2008). Token-based approach for scalable team coordination, in *Cooperative Networks: Control and Optimization*, Edward Elgar Publishing.

[15] Craig Reynolds, Flocks, herds and schools: A distributed behavioral model, SIGGRAPH '87: Proceedings of the 14th annual conference on Computer graphics and interactive techniques, ACM, 1987, 25–34

[16] A. Kolling, S. Nunnally, M. Lewis, & K. Sycara, Towards human control of robot swarms, *Human Robot Interaction Conference* (HRI'12), ACM, 2012, 239-246.