MODELING AND SIMULATION FOR EXPLORING
HUMAN-ROBOT TEAM INTERACTION REQUIREMENTS

Donald D. Dudenhoeffer
David J. Bruemmer
Midge L. Davis

The Human-System Simulation Laboratory
Idaho National Engineering and Environmental Laboratory
P.O. Box 1625
Idaho Falls, ID 83415-3779, U.S.A.

ABSTRACT

Small-sized and micro-robots will soon be available for deployment in large-scale forces. Consequently, the ability of a human operator to coordinate and interact with large-scale robotic forces is of great interest. This paper describes the ways in which modeling and simulation have been used to explore new possibilities for human-robot interaction. The paper also discusses how these explorations have fed implementation of a unified set of command and control concepts for robotic force deployment. Modeling and simulation can play a major role in fielding robot teams in actual missions. While live testing is preferred, limitations in terms of technology, cost, and time often prohibit extensive experimentation with physical multi-robot systems. Simulation provides insight, focuses efforts, eliminates large areas of the possible solution space, and increases the quality of actual testing.

1   INTRODUCTION

Advances in robotics will soon give rise to the development of a diverse array of small to miniature robots capable of autonomous travel through air, sea, and on land. Coupled with advanced sensor and transmission technologies, these units have tremendous potential for intelligence gathering applications, especially in filling current gaps in intelligence collection in times of peace and of conflict.   A key element, however, in the transition from tabletop development to field deployment is the role of the human operator and the necessary interaction between the robotic force and human during the mission. The successful use of large-scale numbers of robots in field applications depends on the ability of human operators to exchange information, provide direction, and gain an understanding of force intent and operations at both micro

and macroscopic levels. The introduction of large-scale forces of autonomous/semi-autonomous robots adds a new dimension and, likewise, new challenges to effective human-machine interaction. Some of the challenges which need to be addressed include operator situation awareness, data representation, and system automation.

This paper begins by presenting a vision for future micro-robotic deployments and examines some relevant human-robot interaction issues. Next the paper describes the role of simulation and modeling in developing concepts for a prototype command and control system for robotic force deployment. Finally, the paper describes the development of *AgentSim*, a simulation framework for evaluating command and control (C2) architectures for human-robot interactions developed under a joint Idaho National Engineering and Environmental Laboratory (INEEL) and Defense Advanced Research Projects Agency (DARPA) project.

2   BACKGROUND

2.1 Research Motivation

Although great strides have been made in technology, the introduction of truly autonomous robotic forces into military applications has yet to be realized.   Radio controlled, tether controlled, and semi-autonomous robotic platforms have been used by the military for surveillance and intelligence. The interaction with human operators has been primarily on a one-to-one or a one-to-several (<10) robot level. The potential for large-scale numbers (e.g.; hundreds to thousands) of robots deployed as a collective force represents tremendous capability in terms of area coverage, redundancy, and time savings.   By the same token, however, it presents a nightmare in terms of control and monitoring of the collective.

This problem was identified some time ago and remains an ongoing concern (Gage 1992, Lee 2000). Although much work in the past few years has explored the utility of distributed control concepts, little research has been conducted to develop robust, scalable command and control tools for interfacing humans with large-scale numbers of robots. If distributed approaches are to be deployed across military, humanitarian and commercial domains, there is an acute need for further consideration of human factors. Recent studies performed for the Department of Defense indicate that one of the greatest obstacles to the inclusion of autonomous mobile surrogates within the battle-space is the need for operator confidence (Halbert et al. 2001, Dudenhoeffer and Bruemmer 2001). Systems that cannot provide appropriate task awareness and system understanding will not be successfully deployed. A simple example of this need is in the use of a team of robots to search a minefield and clear or breach a path for troop movement. After the robots report that the mission is complete, how confident is the squad leader in leading his people along that path?

The monitoring and control of hundreds to thousands of mobile robots demands significant effort in terms of cognitive workload, specifically in the area of maintaining battle-space or *situation awareness* (SA). Within the greater sphere of command and control, much research has been done to understand the need for situation awareness: "the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning and the projection of their status in the near future" (Endsley 1987). The operator must not only understand the robotics force in terms of where they are, but s/he must understand what they are doing now and what they will likely do in the near future. SA is a critical element in decision making, especially in highly dynamic situations that are outside of normal operations. One specific goal of the command and control system should be to provide tools designed to minimize the operator's cognitive workload in developing and maintaining situation awareness of the battle-space. A loss of situation awareness will likely result in slower detection and reaction times as the operator struggles to re-orient him /herself with the current situation.

The envisioned robotics force can act autonomously and yet must be responsive to user control at a variety of levels. Mission planning is the most important aspect of their successful deployment. However, once deployed, the majority of operator interaction is devoted to monitoring their status and conducting minor changes to the original mission plan as the system operates autonomously. Research has indicated that in activities with high level of automation in which an operator serves mainly in a monitoring role, situation awareness may be negatively impacted. It has been hypothesized that this may result from: (1) a loss of vigilance as the operator assumes a monitoring role, (2) the shift from the operator being an active processor of information to that of a passive recipient, or (3) a loss or change in system feedback concerning the state of the system (Endsley and Kiris 1995). As the degree of automation increases, it becomes more difficult for the operator to understand the underlying state of the system. This ability to track and anticipate the status and behavior of the automated system is referred to as *mode awareness* (Sarter and Wood 2000). In the case of the robots, this lack of understanding is evident by the "Now why are they doing that?" response.

Studies on situation awareness and the effects of automation have focused on air traffic control crews, airline pilot crews, and nuclear power plant crews. Gawron (1998) provided some of the first data on these issues in regards to robotics with her research on the human-machine problems associated with the deployment of Uninhabited Aerial Vehicles (UAVs). Some of the relevant human-user interface problems identified are:

1. Data link drop-outs were not always apparent to the operator and the UAVs traveled beyond the data link and control range of the operators
2. Operators had trouble maintaining vigilance over long periods of time during UAV missions of 3.5 to 40 hours
3. Humans could process imagery exploration on only a single data stream at a time, but several UAVs collected two simultaneous data streams
4. Operators had difficulty controlling vehicles when the systems possessed significant time delays in the control system.

The focus of our research is the exploration and development of mechanisms and user tools, which can alleviate these problems. Modeling and simulation have played a key role in this research and development.

2.2 Required C2 Functionality

Command and control for large numbers of autonomous agents represents a unique situation for a human operator. In some instances, it resembles an air traffic controller trying to monitor and coordinate the movements of a large number of aircraft. In other cases, the operator assumes a role much like that of a sonar operator on a submarine, who in monitoring a vast array of sensors is constantly trying to optimize the sonar system performance to identify that one piece of information in a vast ocean of noise.

The functional requirements for autonomous systems control were discussed at a 1998 national technical workshop sponsored by the DOE and the DOD. Figure 1 illustrates the roles / functions that a supervisor of an autonomous system must meet. It also indicates some of the functional elements for command and control systems (DOE 1998).

Of these functional requirements, planning is the most critical due to the autonomous nature of the robots. After

deployment, however, most of an operator's time will be spent in a monitoring mode. It is also in the planning phase that the impact of simulation is greatly seen. Through this planning phase, decision-makers build expectations about the system's execution, including how the robots should perform and the anticipated response to unexpected stimuli. Additionally, through simulation, emergent behaviors become evident which are exceedingly
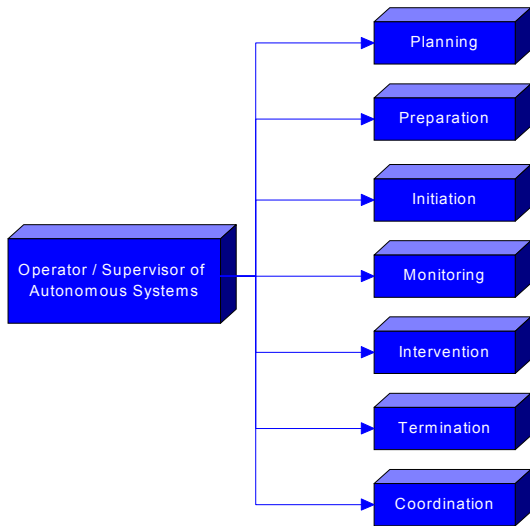

Figure 1: Functional Requirements

difficult even for the system designer, much less an operator, to predict. Emergent effects are dependent on the environment, the number of robots and the elements of the task at hand. They arise from a multitude of subtle interactions and yet can have a devastating effect on the overall system performance. Before we even began implementation of a real robot team, simulation alerted us to fundamental problems that we would face.

As stated earlier, SA plays a large part in the decision process. SA is closely tied to a person's *mental model* of the systems operation. More precisely mental model can be defined as "the mechanism whereby humans are able to generate descriptions of system purpose and form, explanations of system functioning and observed system states, and predictions of future system states." (Rouse and Morris 1986, p.351) The simulation allows the operator the opportunity to develop a mental model. This mental model provides a basis for the operator to recognize and deal with unplanned or unexpected behavior. Three generic situations may emerge from the actual deployment. The first is performance of the robots as planned. In the second situation, the robots do not perform as desired, but the cause is known and the behavior can be altered to mitigate the effects. The third and most devastating case, especially in terms of its effect in eroding user confidence, is when the robots act in a totally unpredicted manner and the cause cannot be determined. Effective simulation can help alleviate this third pos-

sibility, removing uncertainty and confusion both for a developer in the process of implementing a robotic system and for an operator who must plan and coordinate a robotic task.

2.3 Existing C2 Architectures for Distributed Robot Control

As stated earlier, this effort represents one of the first research projects which examines and attempts to develop a multi-agent human command and control tool for large-scale numbers of robots. Dr. Douglas Gage (1992) initially explored the area of large-scale multi-agent command and control, but despite a great deal of recent work with multi-robot systems, little emphasis has been placed on command and control for large-scale robot populations. Several command and controls systems have been developed to support small forces of robots (4 -15). The most notable systems include MissionLab, Demo III, and the United States Air Force's Uninhabited Combat Air Vehicle (UCAV) command system.

Developed at Georgia Tech under the auspices of the Tactical Mobile Robotics (TMR) program, MissionLab is able to greatly facilitate the process of designing, developing and deploying robotic systems (Arkin 1998). Demo III is a DOD sponsored program to develop and demonstrate small autonomous ground vehicles. The primary operator interface is through Operator Control Units (OCU). The system facilitates mission planning, task execution, and re-tasking through a Windows-like map based interface (Morgenthaler et al. 2000). The objective of the Uninhabited Combat Air Vehicle (UCAV) command system developed by the United States Air Force is to provide an interface mechanism for controlling multiple UCAVs in coordinated mission execution. The operator is responsible for establishing mission goals, monitoring system status, and refining task execution (Barbato 2000).

At the present time, there is still a great deal of work to be done in order to enable effective command and control for large numbers of robots. While the systems described above are effective with small numbers of robots, (e.g., 1–4), they do not support levels of force abstraction necessary to control large numbers of robots. The ability to create abstract levels of control, i.e. groups and units, is a key element in permitting human operator "one to many" control of large numbers of robots. Another key ability for an operator to have is an adjustable level of influence over the deployed force. Flexible interfaces and architectures must keep the operator in the loop, provide understanding, cognizance and opportunities for involvement while at the same time filtering and coalescing information so as not to overload the operator or require interaction with each individual robot.

2.4 Hierarchical Structures

Drawing both from the dominance and caste societies that exist in biological systems, a hierarchical system of command and control was selected as the first model for evaluation. The hierarchical system consists of an organizational structure with various levels of control between units within the organization. A "chain of command" exists within the organization, which dictates the relationship between levels units. The military is the most common example of a hierarchical structure. Figure 2 illustrates the basic construct of the hierarchy. In this case, the individual soldiers consist of the base element of the structure.


Figure 2: Hierarchical Command Structure

Soldiers are grouped into units, which possess a unit leader, units into squads, etc. Command and control using a hierarchical framework reduces system complexity by allowing the user to interface with individual soldiers at a high level of abstraction. Command and control is achieved along organizational lines and not by a one-to-one correspondence. Two methods exist for implementation of a hierarchical-based structure. The first involves the specific designation of group membership for individuals. The second involves the designation and control through a leader around which a following develops.

2.4.1 Group Abstraction

Hierarchical structures can be developed through operator designation of groups. Here the operator designates the hierarchy by explicitly choosing group membership. The operator may make this choice based on proximity or based

upon functionality. The operator may want to establish a certain functional capability among a non-homogenous collection of robots. The key is that the operator specifies group membership.

This type of hierarchical structure has strong roots in biological societies that exhibit a dominance ordering among members. This dominance order consists of a set of sustained aggressive-submissive behaviors among members of the society. In the simplest form, this consists of rule by one individual; despotism. In many cases, however, it consists of a hierarchy among members with rank distinction. Here an alpha individual dominates the society; a beta dominates all but the alpha, down to the omega at the bottom of the line. In these societies, dominance is normally based on size, strength, and aggressiveness (Wilson 2000).

2.4.2 Group Leadership

The second method of implementing a hierarchical structure is through dynamic group formation. In this model, the human user does not select the group, but rather selects individual leaders from among the masses. Commands are issued to these leaders who in turn invoke a following among the collective. Group membership is not predefined, but is a function of the "charisma" of the leader.

Leadership in the animal kingdom commonly refers to the simple act of physically leading other group members during movement from one location to another. In this case, the movement leader is not necessarily the dominant member of the group. The leader of the group may change as circumstances warrant such as the discovery of a predatory threat (Wilson 2000). This type of control has applications when combined with a subsumption-type layering of behavior. In this regard, a leader may attract a following of "unemployed" robots. Other robots engaged in meaningful activity would not be compelled to follow the leader. Additionally, followers could break off from the group if stimulated to perform a higher level action.

Once a leader is designated, the question is how to communicate and instill group action? Nature again provides some interesting examples. Birds commonly use a combination of body gestures and audible sounds to signal intent. The honeybee does the waggle dance to indicate the direction and distance to a target. Another form of leadership found in honeybees produces an autocatalytic reaction. This form of action initiation is called the buzzing run, the breaking dance, or Schwirrlauf which honeybees use to induce swarming. In this method "...one or several bees begin to force their way through the throngs with great excitement, running in a zigzag pattern, butting into other workers, and vibrating their abdomens and wings..." (Wilson 2000 p.213). This action incites other worker bees to perform in the same manner and soon most of the collective is effected. After about 10 minutes, the bees nearest the opening depart and the frenzied collective follows.

The ways by which a leader can influence the collective can therefore be grouped into at least the following four styles.

1. Leadership by example – The imitation of the leaders actions by the collective, i.e., follow the leader.
2. Tasking by explicit order – The issuance of direct communication from the leader to subordinates to signify action.
3. Tasking by a preprogrammed response to a leader's or other member's actions – Not necessarily an imitation behavior, but a response to actions by the leader which results in a cascading effect.
4. Any combination of the above.

Just as in the animal kingdom, multi-robot systems require hierarchical control architectures and some means for this hierarchical structure to be realized through an implementation of "influence" as described above. The following section discusses how modeling, simulation and real-world implementation have provided a means to explore social potential fields as an answer to these issues. The ability to modulate social behavior and instill hierarchical structure provides the key to enabling appropriate command and control for a team of multiple robots.

3 SIMULATION DEVELOPMENT AND VALIDATION

3.1 System Design

As a means to study human interface requirements and to prototype command structures, the INEEL developed a command and control suite called *AgentCommand*. Figure 3 represents the major elements behind *AgentCommand* and the basic concept.

AgentCommand represents a modular command and control system for the deployment of autonomous robots. It consists of three major elements:

1. *AgentSim* is a simulation driver that can be utilized for robot behavior development, deployment planning and strategy, and course of action (COA) evaluation. It presents a global view of all simulated agents. It can also used as a source of simulated input for testing and training with *AgentHQ* and *AgentCDR*.
2. *AgentCDR* is an operator control unit (OCU) for in field use. It is the main human robot interface mechanism for monitoring and C2 for individual or large-scale numbers of agents. AgentSim can also provide simulated input.
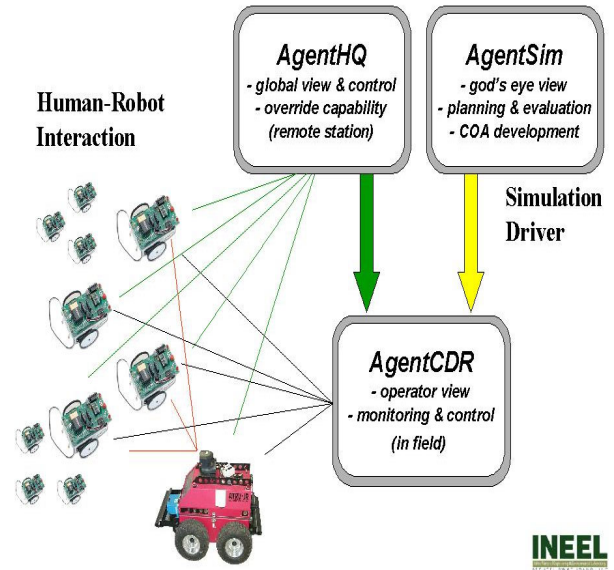3. *AgentHQ* is a centralized command center that permits a high level overview of agent operations.

Figure 3: *AgentCommand* C2 Suite

It interfaces with in field *AgentCDR* modules. AgentSim can also provide simulated input.

A team of "Growbot" robots by Parallax served as the basis for simulation in *AgentSim* and was also the test-platform for integration with *AgentCDR*. Figure 4 shows part of our team of robots. Each robot possesses a behavior set layered in a subsumption style layering. The robots react to the environment and retain very little state knowledge.


Figure 4: The INEEL Robot Team

3.2 C2 System

Confident that the simulation, while not perfect, provided useful insight, the next step was to develop and evaluate command and control structures for robot force deployment. Specifically, we implemented a hierarchical structure, which permitted operator control at an individual robot level, but also allowed control through abstract groupings. AgentSim then served as the test bed for development and evaluation.

3.2.1 Control at the Individual Level

Figure 5 illustrates the ability to query and view individual robots. In this case robots 287 and 340 are selected. The individual robot window displays the system states including type, position, heading, wander direction, and state. The time of death is part of the simulation based on an exponential life expectancy function.
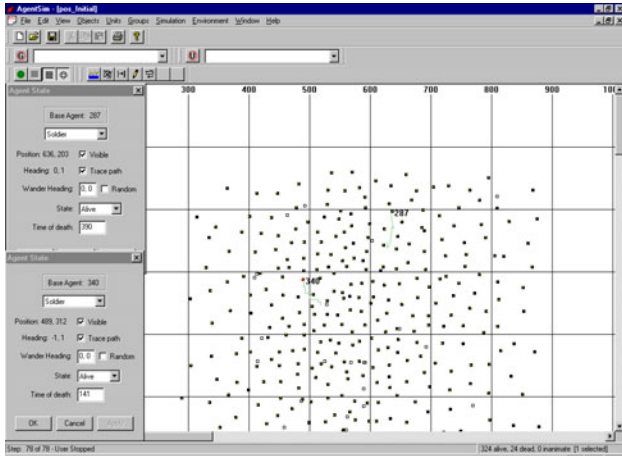

Figure 5: Individual Robot Interaction

The operator is presented two sets of controls. The first type of control does not affect the robot, but provides a means for the operator to view and organize information. Two separate controls of this type exist:

1. A control which allows the operator to remove (hide) information from the viewing screen. This allows the operator to customize a view for focus on a certain area or task while minimizing distracting information.
2. A control which traces the movement path of the robot to assist the operator in understanding behavior and to also track coverage.

The second type of control consists of commands to the robot for behavior modification. These include:

1. Robot type selection. As stated earlier, hierarchical structures can be implemented in multiple ways. One method is by operator grouping; the second is by leadership designation. This control allows the operator to designate a robot as a soldier or as a commander. The change of state to a commander modifies the way neighboring robot react. In this case, a commander imparts a greater attractive force than that exerted by a soldier. Thus a commander is able collect a following of robots.

2. Wander heading. This displays the robot's desired heading subject to the effects of external stimuli. This may be predetermined or a random function.
3. Robot state. This allows the operator to suspend or activate an agent.

Figure 5 also illustrates two visual aids designed to enhance the operator's situation awareness and understanding. These include an adjustable grid overlay to provide the operator a geospatial sense of individual and collective movements. This can be useful in alerting the operator when "out of area" conditions are about to occur. Another feature is the display of "dead" robots. In this case a dead robot represents a robot known to been terminated, or with which communication is lost and cannot be re-established. The open squares on the display represent "dead" robots. The operator may wish to investigate or avoid an area with a high mortality rate. Although this system emphasizes the need for appropriate user input, the system is designed to support a variety of machine learning approaches that can permit autonomous adaptation. For instance, online learning capabilities could permit members of a group to self-adjust their behavior, responding online to significant events such as catastrophic loss of members or physical areas of high mortality.

3.2.2 Control at the Group Level

The C2 level of abstraction immediately above individual control is group control. The interface for Group control and designation is illustrated in Figure 6. Groups permit the designation, selection, tasking and re-tasking of multiple agents by a single operator. A colored ring surrounding the robot identifies group membership. The operator uses the
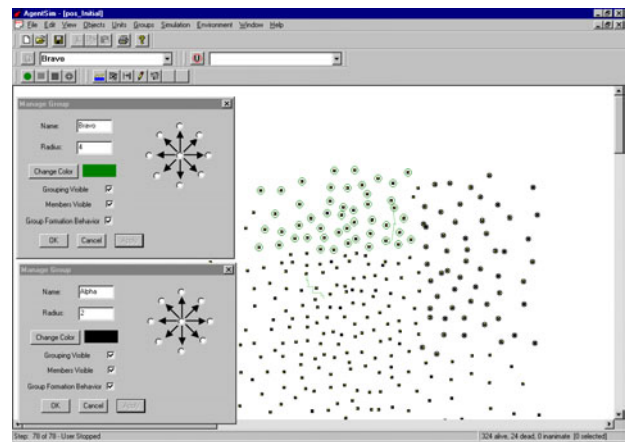

Figure 6: Group Level Interaction

mouse to designate membership in a group. Selection is made individually or by circling a collection of agents. Group controls are similar to those for individual robots and include display (visualization) and operational controls.

Visualization controls allow the operator to custom-configure the information presentation to best fit his/her needs. These controls include:

1. Group name. The operator can give the group a meaningful name beyond the default name.
2. Group radius and color. The operator can modify the appearance of the group by specifying the ring color and size. This supports SA by incorporating pattern recognition into the display.
3. Group visibility. The operator can remove a group representation from the viewing screen. This permits the operator to focus attention while temporarily removing potentially distracting data from the display.

Behavior modification controls include:

1. Group formation behavior. The operator has the ability to enable or disable this behavior among group members. Removal of this behavior removes the potential field effect between neighboring robots. Instead of being influenced by adjacent robots, individual motion is along the goal heading, unless otherwise modified.
2. Goal heading designation. A spin wheel allows the user to specify the group's goal heading by selecting a direction arrow. The geospatial alignment is North at the top of screen.

It should also be noted that robots can be members of multiple groups simultaneously, as denoted by multiple colored rings around the robot representation. This ability promotes flexibility for groups of multiple functionality and also for distributed control among multiple operators.
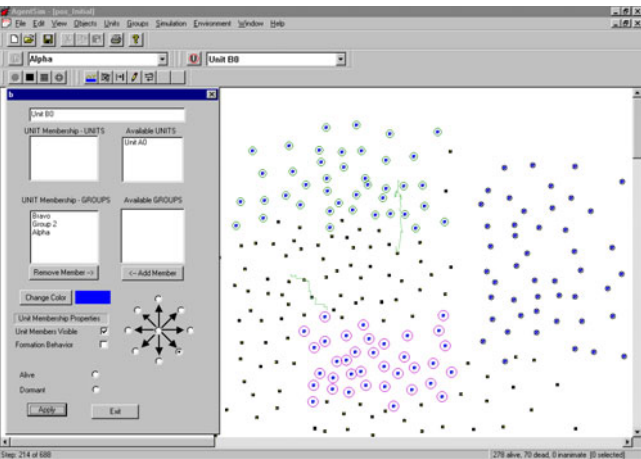

Figure 7: Unit Level Interaction

3.2.3 Control at the Unit Level

Unit designation represents the highest level of control abstraction, Figure 7. This level of abstraction gives commanders flexibility in designating and re-assigning assets to meet specific mission requirements. Membership can consist of groups and/or other units. The display and operational controls are similar to those for group and individual control and include.

1. Unit name.
2. Unit color.
3. Unit visibility.

Behavior modification controls include:

1. Unit formation behavior. The ability to enable or disable this behavior among Unit members.
2. Goal heading designation.

Where as groups may be based on a collection of robots with similar capability, the Unit may designed to contain a specific capability mix of groups to support specific mission profiles. Within the interface is easy to dynamically create and modify the Unit's composition.

3. 3 Simulation Validation

While discussed primarily as a tool for operator planning in this paper, the simulation driver, *AgentSim*, also provided system designers with valuable insight into the development of individual and group behaviors for the team of robots.

Initial simulation trials demonstrated that social potential fields could augment the general robustness of the system; however, it also demonstrated the byproducts of incomplete and imperfect sensing. One of the most pertinent insights was that motion efficiency under social potential force control is highly dependent on accurate neighbor detection. Motion efficiency is defined as the ratio of net distance traveled to total motion. The simulation showed that unless the social potential fields could be appropriately maintained, motion efficiency would greatly degrade. In the case of imperfect sensors the robots alternately lost and regained nearest neighbor detection resulting in an oscillation that wasted time and energy. Dudenhoeffer and Jones (2000) contain these results as well as a detailed explanation of the model's construction in the paper.

When social potential field behavior was implemented in the Growbot robots, some of the same behavior was observed. As in simulation, social potential fields provided a benefit in terms of movement efficiency and significantly augmented performance on a real-world area coverage task. However, the undesirable oscillatory behavior predicted in simulation also manifested as a significant feature of emergent group behavior. Specifically, if the density of

robots rose too high in a given region, the robots began to exhibit the following undesirable emergent effects:

1. Interference:

   - Physical Interference: When density increases, robots sometimes collided or became physically entangled with one another.
   - Chattering: A phenomenon whereby robots hem each other in and, given sufficient obstacle and population density, spin in place. Chattering wastes time and energy, hindering exploration of new ground.

2. Redundancy: Robots tended to cover the same ground as their peers and fall into "ruts."

The field testing of the robots was necessary for model validation, but it also demonstrates the drawbacks of such testing. Naked-eye observations told us little about the effects of ongoing interaction. This drove the need for some empirical, objective means to capture the benefits of our adjustable social potential field interaction. However, one of the problems for gathering empirical data on the behavior of robots has been the difficulty and cost associated with using an accurate indoor positioning system to capture displacement for multiple robots over long periods of time. GPS is not suitable for such indoor purposes and other indoor positioning systems either require costly instrumentation (magnetic field beacon technology, radio, etc.) or are vulnerable to drift (e.g. dead reckoning). Besides, our robots were neither physically large enough nor computationally powerful enough to support sophisticated instrumentation or dead reckoning. Rather than instrument a positioning system, we decided to simply capture robot movement onto the environment.

For our purposes, we constructed an eight by eight foot walled enclosure around a floor covering consisting of large sheets of white-board. Each robot was instrumented

Figure 8: Test Environment

with a Velcro sponge pad, which allowed us to securely attach a dry erase marker to the rear of the robot. Figure 8 illustrates the test-bed environment.

The marker provided a means to track the movements and area coverage for the robots as they explored the environment. When trials involved teams of robots, each was fitted with a different color marker to differentiate its path from the other robots. This novel approach provided valuable "ground-truth" feedback on the precise movements of each robot and the cumulative effect on the resulting area coverage. While effective, this testing consumed tremendous time to conduct the trial runs. Additionally, this setup was sufficient for the nine-member robot team we tested, but would not be suited for large teams of fifty to several hundred robots.

This experiment demonstrates the benefits of incorporating simulation. While the simulation did not reproduce every detail of the actual robot testing, it provided crucial insight, such as a preliminary understanding of the hard problems like how to provide accurate sensing of neighbors and how to avoid the problems of interference that occur when social interaction is not properly controlled.

Another crucial benefit of simulation is that once behavior was observed, it was quickly understood. Especially when multiple robots are concerned, the development process does not lend itself to astute observations. To make matters worse, it is difficult to capture data on large-scale robot interaction. For all of these reasons, simulation can serve a useful role as harbinger, alerting us to effects which are not obviously apparent to the observer of a real robot team. The understanding of social potential fields which we had gained in simulation provided a bootstrap with which we could attack the implementation problem. Even before we turned to real robots, we already knew that our real-world implementation must include some strategy to counter the negative effects of social interaction and quickly found that online learning could provide the necessary means.

4 CONCLUSIONS

The paper has presented the some of the relevant research issues that must be addressed in the area of human-robot interaction. Next the paper has discussed the roll of simulation in exploring these issues. Simulation presents a valuable tool, not only in the development of robotics systems, but as an ongoing human-user aid for planning and providing system understanding. Finally, this utility was demonstrated through its implementation in AgentSim, which served as a framework for developing robot behaviors and prototyping command and control designs for a team of small robots.

The next step is to evaluate the effectiveness of this framework in supporting human-user requirements. Specific areas of concern include situation awareness, the effects of

automation (mode awareness), and mental model formation. Human subject testing and the implementation of the control scheme in actual robots are planned for the near future.

ACKNOWLEDGMENTS

REFERENCES

Arkin, R. C. 1998. *Behavior-Based Robotics*. Cambridge, Massachusetts: The MIT Press.

Barbato, G. 2000, Uninhabited Combat Air Vehicle Controls and Displays for Suppression of Enemy Air Defenses, *CSERIAC Gateway*, XI (1): 1-4.

Department of Energy. 1998. Final Report: Guidance Document and National Technical Workshop Report: Final Report – Guidelines for the Development of Training Programs for the Operation and Maintenance of Robotic Equipment. DOE Office of Science and Technology, Orlando, Florida.

Dudenhoeffer, D. D. and M. P. Jones, 2000. A Formation Behavior for Large-Scale Micro-Robot Force Deployment. In *Proceedings of the 2000 Winter Simulation Conference*, eds. J. A. Joines, R. R. Barton, K. Kang, and P. A. Fishwick, 972-983, Orlando, Florida: Institute of Electrical and Electronics Engineers.

Dudenhoeffer, D. D. and D. J. Bruemmer. 2001. Command and Control Architectures for Autonomous Micro-Robotic Forces. INEEL/EXT-2001-00232, Idaho National Engineering and Environmental Laboratory, Idaho Falls, ID.

Endsley, M.R. 1987. SAGAT: A methodology for the measurement of situation awareness. NOR DOC 87-83, Northrop Corporation, Hawthorne, CA.

Endsley, M.R. and E.O. Kiris. 1995. The out-of-the-loop performance problem and level of control in automation. *Human Factors* 37(2): 381-394.

Gage, D. W. 1992. Command and Control for Many-Robot Systems. In *Proceedings of AUVS-92*. Reprinted in *Unmanned Systems* 10 (4): 28-34.

Gawron, V.J. 1998. Human Factors Problems Associated with Uninhabited Aerial Vehicles (UAVs), In *Proceeding of the Human Factors and Ergonomics Society 42nd Annual Meeting.*1600. Santa Monica, California: Human Factors and Ergonomics Society.

Hallbert, B. P., D. D. Dudenhoeffer, D. J. Bruemmer, M. L. Davis, and G. J. Khoury. 2001. Human Interface Concepts for Autonomous/Distributed Robot Control, INEEL/EXT-2001-00232, Idaho National Engineering and Environmental Laboratory, Idaho Falls, ID.

Lee, J. 2000. Cognitive engineering challenges of managing swarms of self-organizing agent-based automation, In *Proceedings of the XIVth Triennial Congress of the International Ergonomics Society and the 44th Annual Meeting of the HFES* . 1.568-1.571. Santa Monica, California: Human Factors and Ergonomics Society.

Morgenthaler, M., A. Dickinson, and B. Glass. 2000, XUV/Demo III Multi-Vehicle Operator Control Unit, In *Proceedings of SPIE, Unmanned Ground Vehicle Technology II*, 4024. eds. G. R. Gerhart, R. W. Gunderson, and C. M. Shoemaker, 148-158.Bellington, Washington: SPIE- The International Society for Optical Engineering.

Rouse. W. B. and N. M. Morris. 1986. On looking into the black box: Prospects and limits in the search for mental models, *Psychological Bulletin* 100: 349-363.

Sater, N. B. and D. D. Woods. 2000. Team Play with a Powerful and Independent Agent: A Full-Mission Simulation Study. *Human Factors* 42(3): 390-401.

Wilson, E. O. 2000. *Sociobiology: The New Synthesis (Twenty-Fifth Anniversary Edition)*. Cambridge, Massachusetts: The Belknap Press of Harvard University Press.

AUTHOR BIOGRAPHIES

DONALD D. DUDENHOEFFER is research scientist for the Idaho National Engineering and Environmental Laboratory (INEEL) in Idaho Falls, ID. He received his Masters of Science degree in Operations Research from the Naval Postgraduate School in 1994. His research interests include simulation, robotics, military operations and command and control. His email address is `<dudded@ inel.gov>`.

DAVID J. BRUEMMER is a research scientist at the Idaho National Engineering and Environmental Laboratory (INEEL) in Idaho Falls, ID. He has a B.A. in Computer Science and a B.A. in Religious Studies from Swarthmore College. He has worked as a technical advisor for robotics programs at the Defense Advanced Research Projects Agency (DARPA) and as Principal Investigator on a

DARPA project to develop an autonomous landmine detection robot. His research interests include autonomous agents, distributed robotics, adaptive behavior, and human-robot interaction. His email address is `<bruedj@ inel.gov>`.

MIDGE L. DAVIS is a research scientist at the Idaho National Engineering and Environmental Laboratory (INEEL) in Idaho Falls, ID. She received her B.A. in Computer Information Systems from Idaho State University in 1987. Midge's interests are in distributed command and control systems and digital signatures. Her email address is `<mge@inel.gov>`.