Human Supervisory Control of Swarming Networks

M. L. Cummings
Massachusetts Institute of Technology
77 Massachusetts Avenue, Room 33-305
Cambridge, MA 02139
Phone (617) 252-1512
Fax: (617) 253-4196
email: MissyC@mit.edu

Abstract

The introduction of autonomous swarming command and control networks will introduce new layers of human decision making complexity never before experienced in command and control environments. One of the primary advantages of swarming networks of autonomous vehicles is the ability of these networks to process large amounts of information in relatively short periods of time to more optimally achieve a mission goal, all while protecting humans from high risk and potentially hostile environments. However, even though the human may be taken out of the physical control loop in these systems, it will still be critical to include the human at some level of decision making within these swarming networks both as a safety check and also to ensure that the automation is truly supporting overall mission goals. Because of the revolutionary nature of swarming technology, futuristic human interaction with complex autonomous vehicle networks is not well understood, and more research emphasis needs to be placed on human requirements, strengths, and limitations for supervisory control of swarming networks of multiple autonomous vehicles. This research should include an investigation of the interaction of increasing vehicle autonomy on human supervisory control, the effect of increased levels of automation in decision-making, and how situation awareness is impacted by these increasing levels of vehicle autonomy and decision making automation.

Introduction

There is significant growing interest in many government agencies to design and build networks of unmanned vehicles that will have the ability of operating autonomously, which will in effect, take the human out of the loop at certain levels of tasking. For military command and control agencies, autonomous vehicle operations represent a significant leap in leveraging technology to achieve successful mission completion. However, command and control is a term not just representative of military architectures. Any system that involves the coordination of resources by a designated authority to meet a common mission objective is a command and control system. The act of planning, directing, coordinating, and controlling resources to include personnel, equipment, communications, and facilities to accomplish a mission goal applies to many domains such as first response systems, law enforcement, space missions etc.

Unmanned autonomous systems of the future will require less intensive manual control than present day systems. While currently unmanned aerial vehicles (UAVs) require relatively concentrated human input for flight control, in the future, it is likely that the human role for direct flight control will diminish and the need for supervisory control, to include higher level cognitive reasoning, will become much more substantial. This is likely to be true for all autonomous vehicles, not just UAVs, which is exemplified by the current explorations of the Spirit and Opportunity rovers on Mars. In addition to the increased need for a human in the supervisory role in this futuristic use of multiple autonomous vehicles, intra-vehicle collaboration in which vehicles make decisions within their own dedicated network without involving a human adds an entirely new layer of complexity, both cognitively and operationally. This paper will discuss these complex autonomy issues as they relate to human decision making and supervisory control of swarming networks, and outline what areas of research are needed to support the human in the swarming command and control loop.

Swarming and Human Supervisory Control

The concept of human supervisory control is depicted in Figure 1. In supervisory control, a human operator monitors a complex system and intermittently executes some level of control on a process, always acting though some automated agent. During supervisory control, an operator plans an activity that is mediated by the computer, instructs the computer through a series of commands to perform the desired plan, the human then monitors to ensure the action is executed, intervenes when the computer either makes a mistake or requires assistance, and then the human learns from the experience (Sheridan, 1992). Supervisory control is ubiquitous in automated domains, and can be found in air traffic control, process control plants, emergency response coordination such as ambulance dispatch, remote control of robotic vehicles as well as in medical systems such as remote surgery. Networks of autonomous vehicles will include elements of supervisory control, although this area has not been adequately researched. The critical supervisory element for swarming networks will be the feedback mechanism that allows the human to understand what, how, and why a swarm behaves like it does.

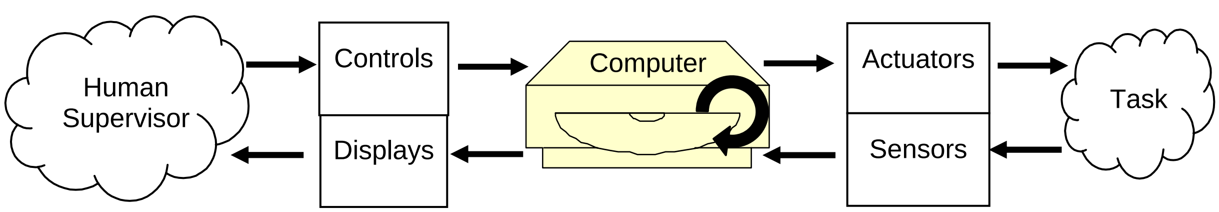

Figure 1: Human Supervisory Control (adapted from Sheridan, 1992)

In supervisory control systems, various levels of automation can be introduced into decision support systems from fully automated where the operator is completely left out of the decision process to minimal levels of automation where the automation only provides recommendations and the operator has the final say. The various levels of automation that can be incorporated into a decision support system are depicted in Table 1, and are known as the SV levels as they were originally proposed by Sheridan and

Table 1: Sheridan & Verplank (1978) Levels of Automation

Automation Level	Automation Description
1	The computer offers no assistance: human must take all decision and actions.
2	The computer offers a complete set of decision/action alternatives, or
3	narrows the selection down to a few, or
4	suggests one alternative, and
5	executes that suggestion if the human approves, or
6	allows the human a restricted time to veto before automatic execution, or
7	executes automatically, then necessarily informs humans, and
8	informs the human only if asked, or
9	informs the human only if it, the computer, decides to.
10	The computer decides everything and acts autonomously, ignoring the human.

Verplank (1978). Much research has been conducted to determine what levels of automation promote effective human computer interaction (Billings, 1997; Endsley & Kiris, 1995; Hancock & Scallen, 1998; Moray, Inagaki, & Itoh, 2000; Sarter & Schroeder, 2001; Scerbo, 1996), however, in the command and control domain, the research has been limited and is virtually non-existent for human interactions with swarming autonomous vehicles. The effect of increasingly higher levels of automation on human control of remotely piloted vehicles has been studied for a few specific military projects (Cummings & Guerlain, in review; Howitt & Richards, 2003; Ruff, Narayanan, & Draper, 2002). These studies were primarily focused on the development of decision aids for human interaction with the vehicles, which were assumed to be operating independently of one another.

For rigid tasks that require no flexibility in decision making and with a low probability of system failure, higher levels of automation often provides the best solution (Endsley & Kaber, 1999; Kaber & Endsley, in press) However, in systems like those that deal with decision-making in dynamic environments with many external and changing constraints, higher levels of automation are not advisable because of the risks and the complexity of both the system and the inability of the automated decision aid to be perfectly reliable (Sarter & Schroeder, 2001; Wickens, 1999). Known as the "brittleness problem," automated decision-support algorithms are typically "hard-wired" in initial design phases, and thus are not able to able to account for and respond to unforeseen problems (Guerlain & Bullemer, 1996; Guerlain, 1995; Smith, McCoy, & C. Layton, 1997). Autonomous vehicles are generally envisioned to operate in areas of uncertainty, which is what makes them so invaluable, however, they will also be subject to the brittleness problem, which is why it is critical that a human be in the supervisory loop.

One primary concern for the interaction of human decision making with automated recommendations is how the automation influences performance. Dixon and Wickens (2003a) have demonstrated that manual control of individual UAVs is impractical and automation is needed for successful mission accomplishment. However, if automation aids do not achieve high reliability, UAV controllers with unreliable automation aiding sometimes perform poorer than those with no automated assistance (Dixon & Wickens, 2003b). In addition, many studies have demonstrated the human tendency to increasingly rely on computer-based recommendations, even though the recommended solutions are not always correct. Known as automation bias, Mosier and Skitka (1996) demonstrated that humans have a tendency to rely upon automated recommendations and pay less attention to contradictory information. In a study examining the potential for automation bias for in-flight retargeting of Tomahawk missiles, when compared to controllers using level 3 automation, controllers with a level 4 recommendation were significantly biased by an automated recommendation (Cummings & Guerlain, in review).

In another study examining the effectiveness of computer-generated recommendations on pilots' decisions to counter in-flight icing problems, the results were ambiguous. When the computer provided accurate advice, pilots with the aid performed better than pilots without the aid. However, when the computer's advice was erroneous, those people without a decision aid outperformed those with one (Sarter & Schroeder, 2001) In addition, as illustrated in a flight-planning tool, just because a human is in the loop for decision making does not mean that better solutions will be found (Layton, Smith, & McCoy, 1994), so the trade-off between the human and automation is not always clear.

Levels of Automation Versus Levels of Autonomy

The previous discussion of levels of automation focused on the human-computer interaction in a decision making task but in the case of human supervisory control of swarming vehicles, another layer of possible automation levels exists between the airborne vehicles. Figure 2 exemplifies how intra-vehicle autonomy could increase, which is not the same as increasing automation for decision support. At the minimum network autonomy level, there is essentially no collaboration between airborne vehicles and at maximum network autonomy, vehicles are in full collaboration and need no human intervention for emergent situations.

Levels of automation for decision do not necessarily increase with the increasing levels of intra-vehicle autonomy. For example, if in a military battlefield a target suddenly emerges such as a surface-to-air missile (SAM) site and unmanned aerial vehicles do not have collaborative capabilities (minimum level of autonomy), the levels of decision support can still vary from 1-10 in that the decision of which vehicle to use to destroy the target can be made by the human, the computer, or some combination of the two. In the case of full intra-vehicle collaboration (maximum autonomy), in the example of the SAM site, the network of unmanned vehicles would determine the best candidate to attack the target. In this case, the levels of automation for decision support can vary from levels 7-10 in determining how involved the human should be in observing the sequence of actions and possibly kept in the loop strictly in a monitoring role.

This duality in levels of automation for both decision-making and intra-vehicle collaboration presents a difficulty supervisory control problem. In current command and control domains without swarming vehicles, the problem space for supervisory control is limited to 10 discrete levels (Table 1.) As networks of vehicles are created that communicate with both humans and one another, the problem space becomes much larger and perplexing. Thus when attempting to design a decision support system for humans interacting with swarming vehicles, not only will it be essential to determine the impact of automation levels for decision making, but equally important to examine the effects of various levels of collaboration between vehicles and the interactions between the two automated systems.

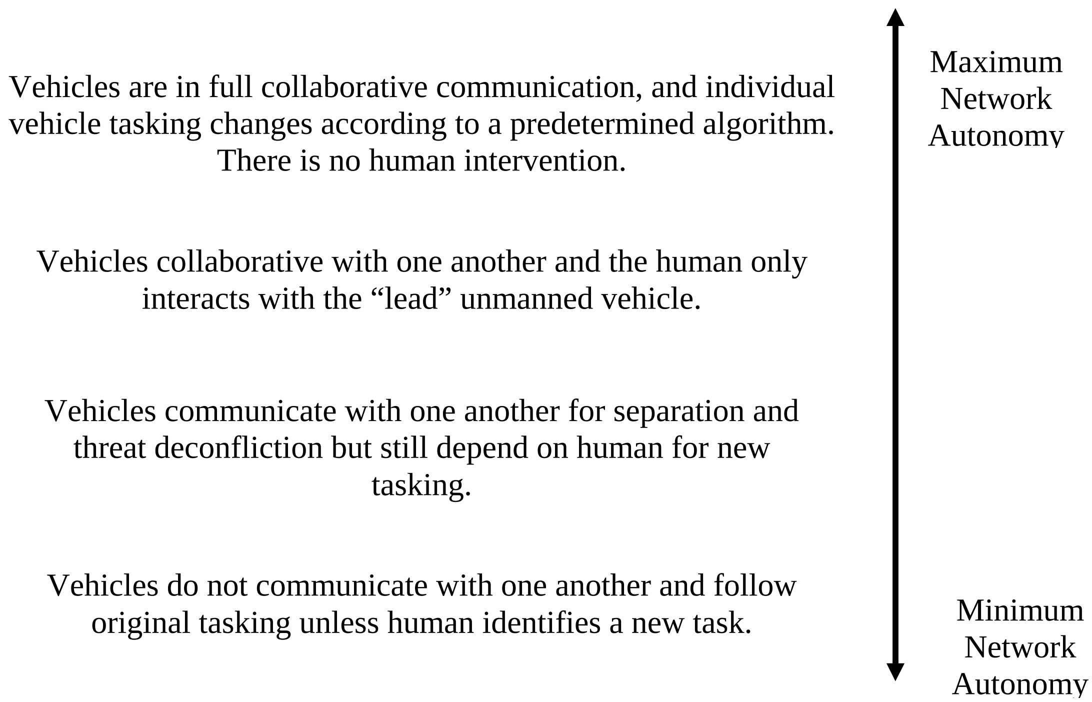

Figure 2: Examples of Intra-vehicle Levels of Autonomy

Situation Awareness and Swarming

In a recent statement, the Office of the Secretary of Defense identified ten primary goals for unmanned aviation in support of the Department of Defense's larger goal of force transformation. Of the ten primary goals, one conveyed the recognition of the critical role that the human plays in a network of autonomous unmanned vehicles, which declared the need for a standard UAV interface that provides critical situational awareness data and precise location data to support airspace integration (Office of the Secretary of Defense, 2002). Indeed the military's vision of Network Centric Warfare hinges on the ability of networks to provide information to support shared situation awareness between operators and decision makers (DoD, 2001). Situation awareness (SA) has long been recognized as a critical human factor in military command and control systems. Military command and control centers must attempt to assimilate and reconstruct the battle picture based on information from a variety of sensor sources such as weapons, satellites, and voice communications. In fact, Klein (2000) determined that SA was a key element for those naval personnel engaged in tasks associated with air target tracking aboard Navy AEGIS cruisers. Because of the complexity and dynamic nature of the command and control environment, maintenance of SA is considered to be of utmost importance (Santoro & Amerson, 1998). Thus understanding how human supervisory control is influenced by swarming networks of autonomous vehicles is critical in developing these networks as well as future operational concepts.

Situation awareness is generally defined as having three levels, which are: 1) the perception of the elements in the environment, 2) the comprehension of the current situation, and 3) the projection of future status (Endsley, 1988; Endsley, 1995). Advanced autonomy schemes that will be used in swarming networks will be important operationally, however, advanced automation and autonomy levels can create problems in the development and maintenance of SA. There are measurable costs to human performance when higher levels of automation are used, such as skill degradation complacency, automation bias, and loss of situational awareness, (Parasuraman, Sheridan, & Wickens, 2000).

Especially important in the development of supervisory control interfaces that promote SA will be the presentation of enormous amount of incoming information generated by the swarming networks. This information must be classified, filtered, and synthesized in such a way that a human supervisor can quickly assess the status of the swarm and determine if any intervention is needed. In general, swarming networks will optimize the aggregate performance against a set of multi-objective cost functions. It will be critical to communicate this information to human decision makers who will be supervising both the global progress of a swarm, and possibly individual elements if a mission dictates. A central issue in maintaining appropriate levels of SA in the supervision of swarming networks will be how to develop visualization tools for multi-objective cost functions that should either be minimized or maximized.

In addition, it will be important to develop human-computer interactive sensitivity analysis tools to determine how human-determined adjustments of variables could change an overall cost function. Allowing human supervisors active participation in decision-making processes provides not only safety-critical benefits and the promotion of SA, but additionally, permitting humans to "tweak" computer-generated solutions protects against automation brittleness and provides for a more robust system that can respond to uncertain and unexpected events in a flexible manner. It will be important in this sensitivity analysis focus to develop visualization tools that convey both the severity of cost function change, but also still convey critical qualitative information such as how an overall balance in resource allocation would shift and affect the mission goals. Because command and control scenarios generally occur under time-pressure and require rapid decision making, the problem of information visualization in the supervision of swarming networks is one not easily solved.

Conclusion

One of the primary advantages of swarming networks of autonomous vehicles is the ability of these networks to process large amounts of information in relatively short periods of time to more optimally achieve a mission goal, all while protecting humans from high risk and potentially hostile environments. Even though the human may be taken out of the physical control loop in these systems, it will still be critical to include the human at some level of decision making within these swarming networks both as a safety check and also to ensure that the automation is truly supporting overall mission goals. However, one consequence of higher levels of autonomy is the trade-off in human understanding of present and future system states as well as the increase in cognitive complexity that is introduced in attempting to assimilate and understand a plethora of incoming information from a network of autonomous vehicles.

Because of the potential problems that will be introduced into command and control scenarios by human interactions with complex autonomous vehicle networks, there is a clear need to explore the human requirements, strengths and limitations for supervisory control of swarming networks of multiple autonomous vehicles. This research should include an investigation of the interaction of increasing vehicle autonomy on human supervisory control, the effect of increased levels of automation in decision-making, and how situation awareness is impacted by these increasing levels of vehicle autonomy and decision making automation.

References

Billings, C. E. (1997). *Aviation Automation: The Search For A Human-Centered Approach*. Hillsdale, N.J: Lawrence Erlbaum Associates.

Cummings, M., & Guerlain, S. (in review). The Decision Ladder as an Automation Planning Tool. *Cognition, Technology, and Work*.

Dixon, S. R., & Wickens, C. D. (2003a). *Control of Multiple-UAVs: A Workload Analysis.* Paper presented at the 12th International Symposium on Aviation Psychology, Dayton, OH.

Dixon, S. R., & Wickens, C. D. (2003b). *Imperfect Automation in Unmanned Aerial Vehicle Flight Control* (AFHD-03-17/MAAD-03-2). Savoy, IL: Institute of Aviation.

DoD (2001). *Network Centric Warfare: Department of Defense Report to Congress*. Washington DC: Office of the Secretary of Defense.

Endsley, M. (1988). *Design and Evaluation for situation awareness enhancement.* Paper presented at the Human Factors Society 32nd Annual Meeting.

Endsley, M. (1995). Measurement of situation awareness in dynamic systems. *Human Factors, 37*(1), 65-84.

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

Endsley, M. R., & Kaber, D. B. (1999). Level of automation effects on performance, situation awareness and workload in a dynamic control task. *Ergonomics, 42*(3), 462-492.

Guerlain, S., & Bullemer, P. (1996). *User-initiated notification: A concept for aiding the monitoring activities of process control operators.* Paper presented at the Human Factors and Ergonomics Society 40th Annual Meeting, Santa Monica, CA.

Guerlain, S. A. (1995). *Using the critiquing approach to cope with brittle expert systems.* Paper presented at the Annual Meeting of the Human Factors and Ergonomics Society.

Hancock, P. A., & Scallen, S. F. (1998). Allocating functions in human-machine systems. In R. Hoffman & M. Sherrick & J. Warm (Eds.), *Viewing psychology as a whole: The integrative science of William N. Dember* (pp. 509-539). Washington DC: American Psychological Association.

Howitt, S. L., & Richards, D. (2003). *The Human Machine Interface for Airborne Control of UAVs.* Paper presented at the 2nd AIAA Unmanned Systems, Technologies, and Operations, San Diego, CA.

Kaber, D. B., & Endsley, M. R. (in press). 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*.

Klein, G. (2000). Analysis of Situation Awareness from Critical Incident Reports. In M. Endsley & D. J. Garland (Eds.), *Situation Awareness Analysis and Measurement* (pp. 51-71). Mahwah, NJ: Lawrence Erlbaum Associates.

Layton, C., Smith, P. J., & McCoy, E. (1994). Design of a cooperative problem-solving system for en-route flight planning: An empirical evaluation. *Human Factors, 36*, 94-119.

Moray, N., Inagaki, T., & Itoh, M. (2000). Adaptive Automation, Trust, and Self-Confidence in Fault Management of Time-Critical Tasks. *Journal of Experimental Psychology: Applied, 6*(1), 44-58.

Mosier, K. L., & Skitka, L. J. (1996). Human Decision Makers and Automated Decision Aids: Made for Each Other? In R. Parasuraman & M. Mouloua (Eds.), *Automation and Human Performance: Theory and Applications* (pp. 201-220). Mahwah, New Jersey: Lawrence Erlbaum Associates, Inc.

Office of the Secretary of Defense. (2002). *Unmanned Aerial Vehicles Roadmap 2002-2027*. Washington DC: Department of Defense.

Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A Model for Types and Levels of Human Interaction with Automation. *IEEE Transactions on Systems, Man, and Cybernetics, 30*(3), 286-297.

Ruff, H., Narayanan, S., & Draper, M. H. (2002). Human Interaction with Levels of Automation and Decision-Aid Fidelity in the Supervisory Control of Multiple Simulated Unmanned Air Vehicles. *Presence*, *11*(4), 335-351.

Santoro, T. P., & Amerson, T. L. (1998). *Definition and Measurement of Situation Awareness in the Submarine Attack Center.* Paper presented at the Command and Control Research and Technology Symposium.

Sarter, N. B., & Schroeder, B. (2001). Supporting decision making and action selection under time pressure and uncertainty: The case of in-flight icing. *Human Factors, 43*, 573-583.

Scerbo, M. W. (1996). Theoretical perspectives on adaptive automation. In R. Parasuraman & M. Mouloua (Eds.), *Automation and human performance: Theory and applications*. Mahwah, NJ: Erlbaum.

Sheridan, T. B. (1992). *Telerobotics, Automation, and Human Supervisory Control*. Cambridge, MA: MIT Press.

Sheridan, T. B., & Verplank, W. (1978). *Human and Computer Control of Undersea Teleoperators*. Cambridge, MA: Man-Machine Systems Laboratory, Department of Mechanical Engineering, MIT.

Smith, P., McCoy, E., & C. Layton. (1997). Brittleness in the design of cooperative problem-solving systems: The effects on user performance. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans, 27*, 360-371.

Wickens, C. D. (1999). Automation in Air Traffic Control: The Human Performance Issues. In M. W. Scerbo & M. Mouloua (Eds.), *Automation Technology and Human Performance: Current Research and Trends*. Mahwah, NJ: Lawrence Erlbaum Associates.