Predicting Controller Capacity in Supervisory Control of Multiple UAVs

Mary L. Cummings, *Member, IEEE*, and Paul J. Mitchell

Abstract—In the future vision of allowing a single operator to remotely control multiple unmanned vehicles, it is not well understood what cognitive constraints limit the number of vehicles and related tasks that a single operator can manage. This paper illustrates that, when predicting the number of unmanned aerial vehicles (UAVs) that a single operator can control, it is important to model the sources of wait times (WTs) caused by human–vehicle interaction, particularly since these times could potentially lead to a system failure. Specifically, these sources of vehicle WTs include cognitive reorientation and interaction WT (WTI), queues for multiple-vehicle interactions, and loss of situation awareness (SA) WTs. When WTs were included, predictions using a multiple homogeneous and independent UAV simulation dropped by up to 67%, with a loss of SA as the primary source of WT delays. Moreover, this paper demonstrated that even in a highly automated management-by-exception system, which should alleviate queuing and WTIs, operator capacity is still affected by the SA WT, causing a 36% decrease over the capacity model with no WT included.

Index Terms—Fan-out (FO), multiple unmanned vehicles, operator capacity, supervisory control.

I. INTRODUCTION

WITH the recognition that intelligent autonomy could allow a single operator to control multiple vehicles (including air, ground, and water), instead of the converse which is true today, there is an increasing interest in predicting the maximum numbers of autonomous vehicles that an operator can control. Indeed, the Office of the Secretary Defense's Roadmap for unmanned aircraft systems specifically calls for such future architectures [1]. Because of the increased number of sensors, the volume of information, and the operational demands that will naturally occur in a multiple-vehicle control environment, excessive cognitive demands will likely be placed on operators. As a result, efficiently allocating attention between a set of dynamic tasks will be critical to both human and system performance. To this end, this paper will discuss recent efforts to model human capacity for the management of multiple

M. L. Cummings is with the Aeronautics and Astronautics Department, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: missyc@mit.edu).
P. J. Mitchell is with General Electric Energy in Power Generation Technology, Greenville, SC 29615 USA (e-mail: paulj.mitchell@ge.com).

unmanned/uninhabited aerial vehicles (UAVs), outline an extension to a previous operator capacity model, and demonstrate, using experimental evidence, how the upper bound predictions can significantly change with the new model.

II. BACKGROUND

While a number of research efforts have experimentally demonstrated in simulations that, under various levels of autonomy, operators can control anywhere from 1 to 12 unmanned air vehicles [2]–[4] and one to eight ground vehicles [5], [6], there has been very little research in developing a model for predicting controller capacity. In terms of operator intervention, control can range from manual control (e.g., actually flying the UAV) to supervisory control (higher level planning and goal state changes). Operators that are engaged in manual control will necessarily be able to control fewer vehicles than in supervisory control since manual control requires significantly higher levels of dedicated attention to lower level skill-based cognitive activities, with fewer resources available for higher level planning tasks. Indeed, the tasks and their required attentional resources of each vehicle will likely limit the human capacity as well as the number of vehicles.

In one of the few attempts to theoretically predict an upper bound for an operator controlling multiple independent homogeneous unmanned vehicles, it has been proposed that the number of ground robots that a single individual could control can be represented by (1) [7]–[10]. In this equation, fan-out (FO) is equal to the number of robots that a human can effectively control, neglect time (NT) is the expected amount of time that a robot can be ignored before its performance drops below some predetermined threshold, and interaction time (IT) is the time it takes for a human to interact with the robot to raise the performance to an acceptable level. Thus, the total capacity is the summation of all NT and IT divided by the IT. While originally intended for ground-based robots, this paper has direct relevance to other unmanned systems in the air or on/under the water, as these systems are becoming more autonomous and will move from the manual control domain to that of multiple-vehicle supervisory control. Because these vehicles are assumed to be homogenous, the assumption is that the operators are responsible for a homogeneous set of tasks

$$\text{FO} = \frac{\text{NT} + \text{IT}}{\text{IT}} = \frac{\text{NT}}{\text{IT}} + 1. \qquad (1)$$

The ratio of operator IT to an overall mission time, like NT, is similar to another metric known as the utilization (UT) or

Fig. 1.   Relationship between IT, NT, and WTs.

percent busy time. UT is the ratio of time that an operator is actively engaged in a task (including the three major elements of information processing: perception, cognition, and action) to the overall operating time of a system. Several studies have demonstrated that UT is a metric that can be used to evaluate human–automation interaction and that operators working beyond 70% utilization experience significantly degraded performance [2], [11], [12]. However, UT is a gross measure of human–automation IT in that it does not discriminate among the different kinds of IT (which will be discussed below).

Equation (1) is an improvement over the more general UT metric in that it specifically compares NT and IT and is a useful measure to establish an ideal upper bound for operators. However, it suffers from the same limitations as UT in that it does not capture the different components of IT. Furthermore, we propose that there remains an additional critical variable that must be considered when modeling human control of multiple vehicles, regardless of whether they are on the ground or in the air, and that is the concept of wait time (WT). In complex problem-solving tasks, humans are serial processors in that they can only solve a single complex problem or task at a time [13], and while they can rapidly switch between tasks, any sequence of tasks requiring complex cognition will form a queue, and consequently, WTs for these tasks will be built. In the context of a system of multiple vehicles in which two or more vehicles will likely require attention simultaneously from a human operator, WTs are significant in that as they increase, the actual number of vehicles that can be effectively controlled decreases. Moreover, while task and human–computer interface dependently, if humans are cognitively saturated or have a low situation awareness (SA), they may not realize that a particular vehicle is in need of attention, thus inducing another WT until the task is recognized.

Fig. 1 shows the relationships between IT, NT, and WT in terms of overall multiple-vehicle system performance. The role of the operator, in general, is to ensure that all vehicles operate at some predetermined minimally acceptable level of performance. Because of the intermittent supervisory control aspect of multiple-vehicle control, operators will only attend to those vehicles in need (IT) to bring them to this acceptable performance level and, hence, the increase in performance

during IT in Fig. 1. During NT periods, the vehicles operate without the need for operator intervention until such time that the performance drops below the threshold. Point A in Fig. 1 represents a discrete event that terminates NT, which causes the vehicle to require operator assistance such as a system failure (e.g., an engine loss or loss of a sensor). However, the termination of an NT state is not necessarily caused by singular discrete events. Point B represents performance degradation during NT, which causes the vehicle performance to eventually drop below the performance threshold, e.g., a slow degradation of an inertial navigation system. In both NT cases, once the performance has dropped below an acceptable level, the vehicle must wait until the human recognizes the problem, solves the problem internally, and then communicates that goal to bring the vehicle to an acceptable state so that it can move into an NT state. As shown in Fig. 1 at point C, if the problem is not addressed immediately, the vehicle must wait and operate at some nonoptimal performance level. As will be discussed in depth in the next section, this delay can be caused by either a loss of operator SA (the operator does not recognize that the vehicle is in a degraded state) or the operator is overloaded with more critical tasks from other vehicles. Moreover, since operators cannot generally and instantly bring the performance above the threshold, there is an additional interaction WT (WTI) while the operator is actively attending to a vehicle (e.g., commanding it back to the correct course) to achieve the desired performance level.

IT and NT are important in predicting operator capacities for handling multiple vehicles, particularly for those domains that are time-critical and have high risk, like UAVs, as WT could become a critical point for possible system failure. Even for unmanned underwater vehicles that must surface for communications, waiting is not only suboptimal but also can be extremely hazardous. Moreover, ground-based robots (or unmanned ground vehicles) engage in time-critical missions, such as search and rescue, which would be negatively impacted if the problem of WT was not addressed. While most robots and vehicles can be preprogrammed to follow some predetermined contingency plan if they do not receive the required attention, mission success will likely be significantly degraded if WTs grow unexpectedly.

A. WTs

As outlined previously, from a robot or vehicle perspective, WT imposed by human interaction (or lack thereof) can be decomposed into three basic components: 1) WT in the human decision-making queue (WTQ); 2) WTI, and 3) WT due to loss of situation awareness (WTSA). For example, suppose that an operator is controlling two robots on a semiautonomous navigation task (much like the Mars Rovers). While typical operations involve human interaction with a single vehicle, there will be times when both vehicles require attention near-simultaneously. When this occurs, once the human operator begins assisting the first robot, the second robot must wait while the operator solves the problem and then issues the commands to it (WTI_1). For the second robot, the time it waits in the queue (WTQ_2) is effectively WTI_1.

IT is the time during which a human's attention is focused on a single vehicle in order to solve a problem or induce some change to improve the performance above a specified threshold. From the human perspective, IT includes the time required to determine the nature of the problem, to solve the problem, and to communicate that solution with some type of feedback. Thus, the vehicle must wait some period of time during the "interaction" due to the human decision-making process. In teleoperation, where the human is directly controlling the robot's movements and positions, WTIs might be very small and occur in rapid succession as the controller gets the sensor feedback and adjusts commands accordingly. However, in other supervisory control scenarios that require minimal manual control but significant cognitive input, such as the need to provide a new mission to a UAV, WTI can be quite large depending on the complexity of the problem.

WTSA is perhaps the most difficult WT component to model because it represents how cognitively engaged an operator is in a task. SA is generally defined as having three levels, which are as follows: 1) the perception of the elements in the environment; 2) the comprehension of the current situation; and 3) the projection of future status [14], [15]. While SA can decrease under a high workload due to the competition for attentional resources [16], it can also decrease under a low workload due to boredom and complacency [17]. If an operator does not realize that a vehicle needs attention (and, thus, experiences a loss of SA), the time from the initial onset of the event to the actual operator recognition of the problem could range from seconds to minutes.

Equation (2) categorizes general WT that is not part of human–robot interaction as the summation of WTs that result from queues due to the near-simultaneous arrival of problems plus the WTs due to the operator loss of SA

$$\text{WT} = \sum_{i=1}^{X} \text{WTQ}_i + \sum_{j=1}^{Y} \text{WTSA}_j \qquad (2)$$

$$\text{FO} = \frac{\text{NT}}{\text{IT(WTI)} + \text{WT}} + 1 \qquad (3)$$

where

WT   wait time;
WTQ   queuing WT;
WTSA   WT caused by a loss of situation awareness;
WTI   WT due to human decision making nested in overall IT;
X number of human–automation interaction queues that build;
Y number of time periods in which a loss of SA causes WT.

Equation (3) demonstrates how the FO equation would change as a result of the inclusion of WTs. For the remaining discussion in this paper, WTI will be considered to be subsumed in IT, which includes the time needed for the human to make decisions about a new vehicle problem and then to communicate the solution to the vehicle.

While not explicitly linked to IT, WT as a function of WTSA and WTQ occurs in the denominator of (3) because it represents the time that should have been spent in the interaction task, which is due to a degraded vehicle state but not due to human attention inefficiency (either a loss of SA or an inefficient sequencing of tasks in a queue). Thus, WTs occur only after a vehicle has moved from the NT period to an IT period but must wait until the operator attends to it to achieve a performance increase. WTs, in effect, increase the overall IT.

The original FO (1) represents a theoretically perfect system with an instantaneous system response, as well as humans who induce no system delays. The modified FO in (3) represents a more conservative upper bound which accounts for inefficiencies in human attention allocation (WTSA), limits on the human ability to attend to multiple complex tasks simultaneously (WTQ), as well as the inherent delay that will always accompany the human information processing loop of perception, cognition, and action that occur in every decision-making process (WTI). It should be noted that the modified FO (3) is not intended to make an accurate prediction of exactly how many vehicles a person could control but merely to set an upper limit given human and system limitations.

B. Levels of Automation

One of the primary variables that will influence the operator capacity in the supervision of multiple vehicles is the level of system autonomy. The challenge in achieving the one-controlling-many goal for the management of multiple unmanned vehicles in the future is to determine how automation can be used to reduce human workload. Higher levels of system autonomy will increase NT, which should decrease IT and the associated WTs. In terms of aiding the operator in complex decisions, automation decision support can range from fully automatic, where the operator is completely left out of the decision process, to minimal levels where the automation offers basic data filtering or recommendations for the human to consider [18]. For rigid tasks that require no flexibility in decision making and have a low probability of system failure, higher levels of automation (LOAs) often provide the best solution in terms of operator workload [19]. However, even partially automated systems can result in measurable costs in human performance, such as loss of situational awareness, complacency, skill degradation, and decision biases [20].

In the context of managing multiple vehicles, increasing the LOAs should reduce the workload and WTs by effectively reducing IT, but there are several measurable costs for both operators and the system in general. The loss of SA and the propensity for automation bias are significant problems that can result when automation authority increases. For example, at high LOAs where the system takes over the execution of some or all functions, overall WTs are expected to decrease as automation can generally make faster decisions than humans, and the number of opportunities for lower level human errors should be reduced. Superficially, it seems that the system should perform well at higher LOAs, but under abnormal and unexpected conditions, automation could fail, possibly causing a catastrophic event to occur. This is particularly problematic under uncertain or novel conditions because the human operator may not understand the situation well enough to
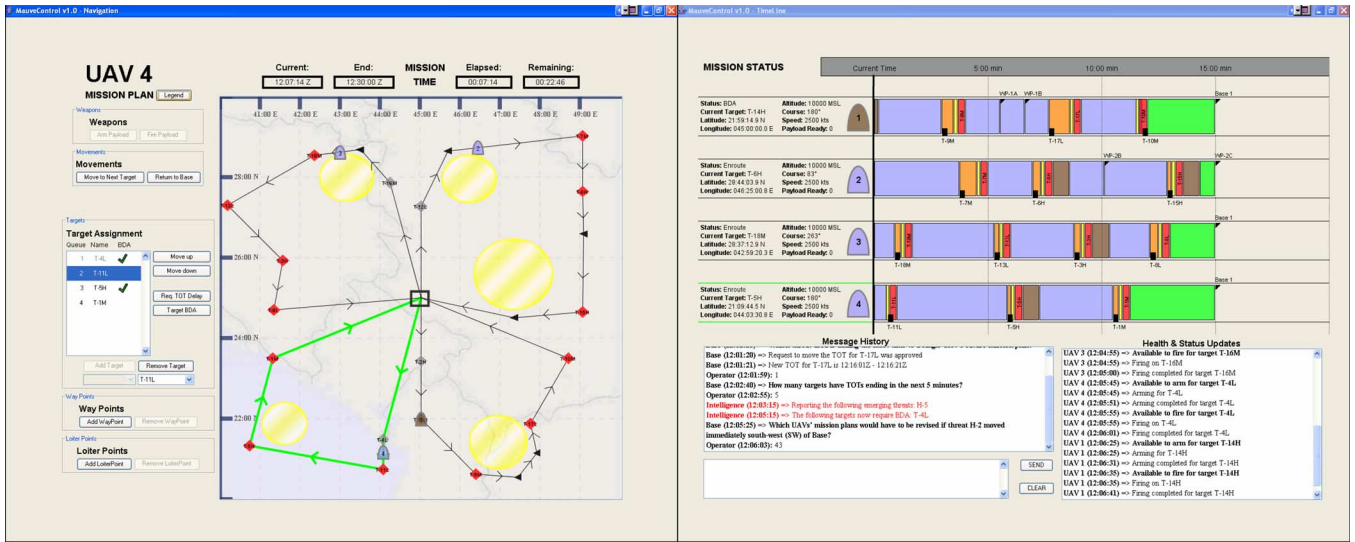
Fig. 2.   MAUVE interface.

determine if the automation is working correctly and if it requires intervention.

III. MAUVE: The Experimental Test Bed

In order to study how the increasing autonomy influences the proposed different components of WT in the multiple unmanned vehicle control domain, a dual-screen simulation test bed, which is named the Multi-Aerial Unmanned Vehicle Experiment (MAUVE) interface, was developed (Fig. 2). This interface allows an operator to supervise four homogeneous and independent UAVs simultaneously and to intervene as the situation requires. In this simulation, operators are responsible for supervising four UAVs that are tasked to destroy a set of time-sensitive targets in a suppression of enemy air defenses mission. Because the simulated UAVs are highly autonomous, they only require operators to provide high-level mission planning and execution actions as inputs to the UAVs. The UAVs launch with a predetermined mission plan; therefore, initial target assignments and routes are already completed. The operator's job is to monitor each UAV's progress, to replan aspects of the mission in reaction to unexpected events, and, in some cases, to manually execute mission critical actions such as arming and firing of payloads.

MAUVE UAVs are capable of six high-level actions: traveling enroute to targets, loitering at specific locations, arming payloads, firing payloads, performing battle damage assessment (BDA), and returning to base, generally in this order. BDA is the postfiring phase of combat where it is determined whether the weapon(s) hit the target and if the desired effect was achieved. In MAUVE, BDA is semiautomated in the sense that operators are responsible for scheduling BDA in advance, but the UAV performs it automatically after firing, if scheduled.

A. Navigation Display

The left-hand side of the MAUVE interface in Fig. 2 is known as the navigation display, and it consists of a map display

and a mission planning and execution panel. Both the elapsed time and the time remaining in absolute and relative terms are shown on the top right of the map display. The map display represents a 2-D spatial layout of the battle space, which is updated in real time. Threat or hazard areas, which are circular in shape, have a striped yellow pattern and can be dynamic throughout scenarios, changing in size and locations, disappearing entirely, or emerging as time progresses.

The UAVs, which are always held constant at four, independently change colors according to their current action which includes arm payload, fire payload, move to next target, and return to base. Arming and firing are only enabled if the preestablished rules of engagement (RoE) of the simulation are met. For arming, the UAV must be directly on top of a target within predetermined arming or firing windows. For firing, the UAV should be armed at the correct target. The move-to-next-target button allows operators to bring UAVs out of the loiter patterns in case of scheduling problems, and the return-to-base button causes all future targets, waypoints, and loiter points to be deleted from the mission plan. Subsequently, a straight-line path is planned directly back to base and is intended for emergency scenarios.

Targets are designated by a diamond-shaped icon and are assigned a priority of high (H), medium (M), or low (L). Active targets are differentiated from the inactive targets by their color, which is either red or gray on the display, respectively. Waypoints, which are shown on the map display with black triangle icons, represent UAV turning points. In addition, UAVs can be loitered at specific points, and typically, a UAV loiters for a user-specified amount of time before moving. However, the departure from a loiter pattern must be commanded by the operator. UAV routes can be changed in minor ways by selecting a particular waypoint or loiter point and by dragging it to the desired location. More significant routing changes, such as the addition or removal of waypoints, loiter points, or targets, can be accomplished using the mission planning and execution panel. Routing changes are typically only required as a result of unexpected scenarios and represent real-time replanning.

Operators are provided with a "request time-on-target (TOT) delay" button which gives them limited opportunities to manipulate the TOTs for those targets assigned. Operators can request a TOT delay for a given target for two reasons: 1) According to the current mission plan, they are predicted to arrive late to a target and, therefore, will miss their deadline; or 2) for workload purposes, i.e., if operators feel that they need to spread out a very high workload time block to manage the UAVs more effectively. However, this function should be used with care because moving back one target's deadline likely changes the UAV's arrival time at all subsequent targets. It is important to highlight this change of TOT as a request, not a command, and the operators' requests can be approved or denied. In MAUVE, the probability of approval is a function of how far, in advance of the deadline, the request is sent, as would likely be the case in true military situations. When a TOT deadline is immediately approaching, the chance of approval is zero but nearly 1.0 when requested 15 min in advance, which is the limit for the decision support—to be discussed in the next section. In MAUVE, a request always takes 5 s for response, and during this intervening time, no other TOT requests can be made.

B. Decision Support Display

The right-hand side of the MAUVE simulation in Fig. 2 provides decision support for task management, and it consists of a UAV status window, chat box, UAV health and status updates, and the decision support window. The status window at the top left of the decision support display gives operators a low-level detailed information for each UAV, such as current target, current action being performed, latitude and longitude, course, and weapons information. Speed and altitude are also shown, although they are not directly controllable by operators.

The bottom left of the decision support display has a text-based datalink communication tool known as a chat box that contains a time history of all human communication interactions. The chat box window displays various notification messages that appear in response to scenario events or actions taken by users, as well as periodic task-relevant questions for operators to answer from a simulated superior officer. One message that is particularly important to operators is a notification that a TOT request is accepted or denied. The bottom right of the decision support display contains a UAV health and status notification window which separates human communications in the simulation from system communications and only contains messages from individual UAVs.

The manipulation of the appearance and functionality of the actual task management decision support (the timeline block) is the primary independent variable for the experiment that will be discussed in the subsequent section. The basic premise of the decision support is to simplify standard air tasking order (ATO) data and to combine them in a single interface with an up-to-date mission planning information. An ATO provides a schedule of events for a mission as well as the required resources. The information contained in an ATO includes which aircraft has been assigned to certain strikes, times on targets, waypoints, and call signs that are to be used on those missions. Since
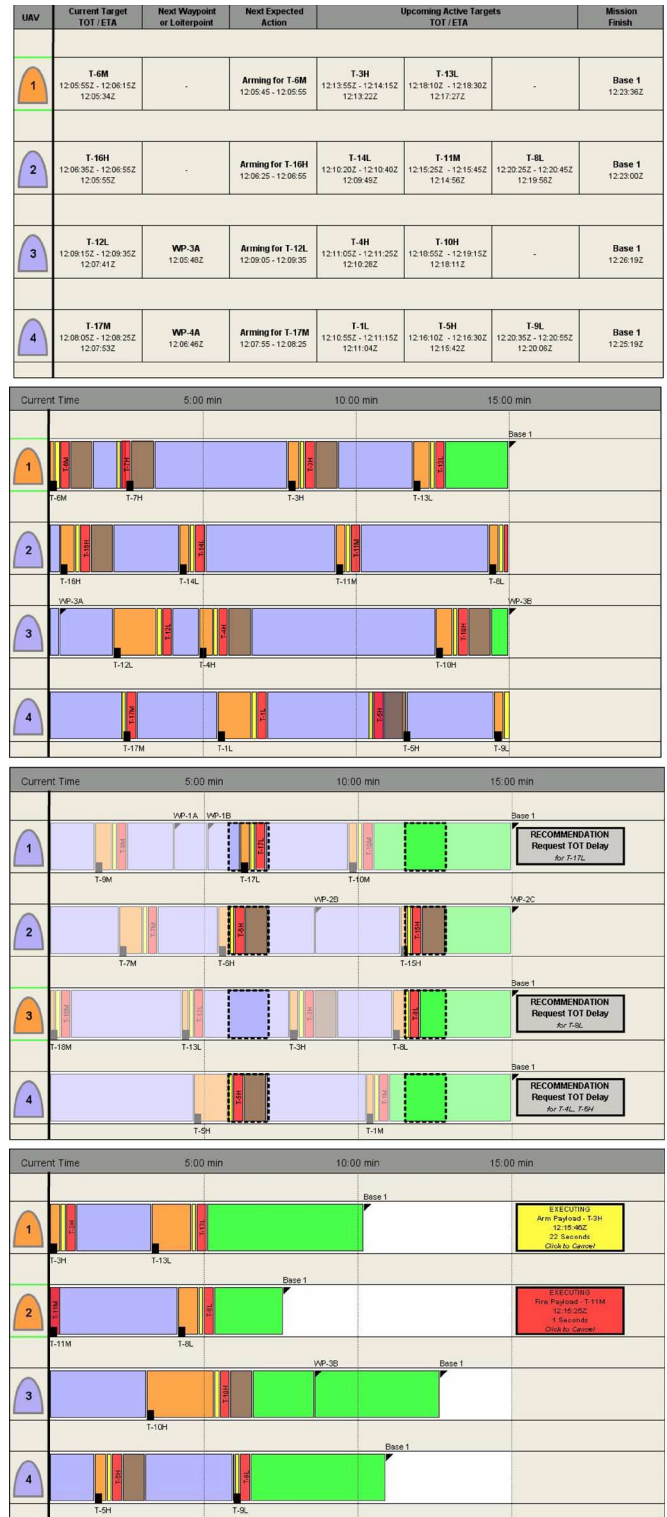

Fig. 3. Increasing levels of mission autonomy. (a) Manual. (b) Passive. (c) Active. (d) Superactive.

the focus of this paper was to determine how the increasing automation would affect WTs, four versions of ATO decision support were created in MAUVE: termed manual, passive, active, and superactive, respectively. These will be described in detail below.

The manual LOA level of the decision support [Fig. 3(a)] presents all the required ATO and mission planning information

in a text-based table format. Current TOT windows and estimated times of arrivals (ETAs) for up to the next four targets in a UAV's timeline are presented. The next waypoint or navigation point on the current route segment is provided, as well as the final ETA for arrival at base in the last column. A low level of automated assistance is provided to the user through the "Next Expected Action" column, which tells the user what they should be doing next and at what time, according to the ATO. This information is updated dynamically to reflect the changing ATO requirements and mission planning changes initiated by the user.

The passive LOA [Fig. 3(b)] assimilates all of the ATO and current mission information contained in the manual level and transforms it into a horizontal timeline format, which is color coded by action. The major difference between the passive and the manual levels is the entire schedule integration for users using a graphical format versus a text-based format. The visual timelines are relative, with the left side representing the predicted UAV actions in the near future and the right side up to 15 min into the future. Target ETAs are represented by black rectangles on the bottom of each timeline, and waypoint, loiter point, and base arrival times are marked by black triangles on the top of each timeline. The TOT windows, arming windows, and BDA periods are represented by red, yellow, and brown blocks of time.

All of the information in the graphical timeline is available under manual automation, but it cannot be obtained without a substantial cross-checking of table entries. Moreover, the manual table format does not facilitate schedule comparisons across different vehicles. With this visual representation, recognizing problems with the current schedule is perceptually based, which allows users to visually compare the relative location of display elements instead of specific times to one another. This level of decision support is termed passive because the automation is not performing any tasks except transforming the basic ATO and mission planning information into a graphical format.

The active LOA [Fig. 3(c)] uses the same horizontal timeline format as the passive automation level but provides additional help from the computer. In the active version, an algorithm searches for periods of time predicted to cause high workload for the operator, and then directs the operator's attention towards them. The computer identifies a high workload area or "bottleneck" as a period of time during which multiple UAVs are scheduled to simultaneously execute mission critical actions that require human interaction, including arming, firing, or performing BDA. The automation draws attention to these areas of concern by a reverse-shading technique, in which the "bottlenecks" are highlighted while the rest of the timeline's colors are muted but still visible. As no information is hidden, but only made less salient, the operator's attention can be directed to the appropriate areas of the schedule while allowing him to maintain SA for the overall mission.

In addition to identifying areas of high workload, the computer also recommends a course of action to alleviate the high workload areas, such as moving a particular TOT. Computer recommendations appear in gray boxes to the right of each relevant UAV's timeline, and subjects have several options:

1) they can acknowledge a high workload area but take no action; 2) they can follow the recommendation to relieve the projected high workload area by shifting a TOT; or 3) they can make other mission planning changes to ensure that the high workload area does not occur, such as deleting a target from a UAV's plan. While the automation makes locally optimal recommendations, the algorithm is not globally optimal. Following the computer's recommendation to relieve a high workload area removes that particular schedule conflict but sometimes creates another in the process.

The reverse-shading technique, in conjunction with the recommendations, permits operators to make local changes to alleviate workload and to immediately see the effect on the global plans of all the UAVs. The purpose of the active LOA is to help operators identify future time periods of potential high workload further in advance so that they can avoid them or at least be better prepared to handle them. This level of decision support is termed active because the automation actively aids operators by narrowing down a set of possible solution alternatives for high workload problems.

The superactive LOA [Fig. 3(d)] also builds upon the passive-level visual timeline, but instead of making recommendations to the operator, as in the active LOA, a management-by-exception (MBE) approach is taken. MBE occurs when automation notifies a human that he is going to take some action and gives the operator a limited time to veto the automation. In this experiment, MBE occurs in the superactive condition when the computer automatically executes the arming and firing actions for all UAVs at each target when the RoE are met. Operators are given 30 s to intervene with the automation. For example, in order to fire, a UAV has to be located at the particular target it was due to fire on, already armed, and within the TOT window for that target. While the automation handles the actual execution of tasks, the operator is still responsible for determining if the arming and firing actions are appropriate, as well as replanning actions and manipulating routes to ensure that the UAVs arrive at the correct targets on time. For the 30 s prior to every arming and firing action, exception boxes appear to the right of the timeline, which allow the operator to veto these actions. The color of the box indicates which action the UAV is preparing to perform, red for firing and yellow for arming.

IV. EXPERIMENTAL VALIDATION

Fig. 4 shows our hypotheses in terms of how WTs would be influenced by the MAUVE's different LOAs. Because increasing LOAs should theoretically reduce the operators' workload, we proposed that WTI would be progressively lower for the increasing levels. As previously discussed, WTQ depends heavily on WTI, which is expected to decrease with increasing LOAs in MAUVE. Therefore, WTQ should follow the same decreasing trend as LOA increases.

For WTSA, the highest LOA (superactive) should theoretically eliminate any WTSA but only for expected events. When the unexpected events occur, operators may have low SA due to complacency, and therefore, they may incur WTSA by not noticing that the mission plan needs adjustment or that a UAV
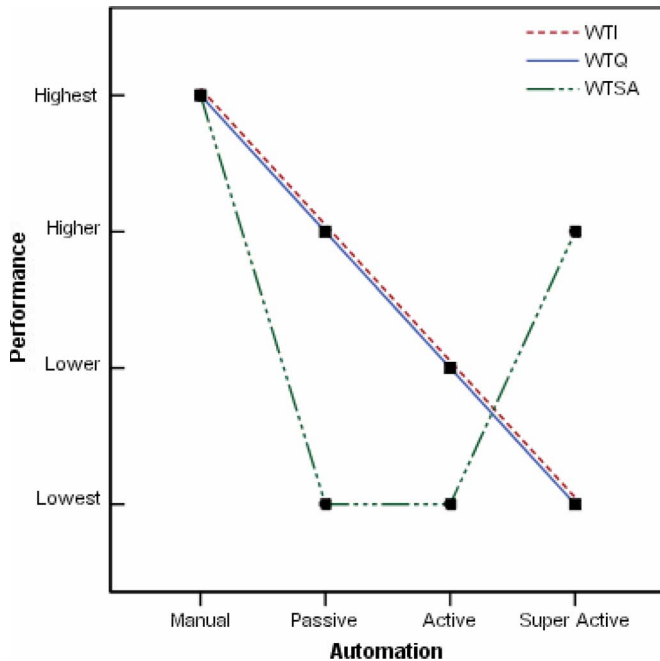
Fig. 4. Experimental hypotheses.

needs attention. In contrast, given the lack of decision support, the manual level should produce the largest total amount of WTSAs of all LOAs. We hypothesized that the active and passive levels should have the lowest accumulated WTSA because operators are continually interacting with the vehicles and are thus "in-the-loop" as opposed to "on-the-loop" for those operators with a superactive control. The hypotheses in Fig. 4 are ranked on an ordinal scale both to normalize the variables and because no prior evidence exists to formulate more quantitative hypotheses. In order to validate these hypotheses, an experiment with the MAUVE simulation interface was conducted, which will be described next.

A. Apparatus, Participants, and Procedures

Training and testing were conducted on a four-screen system called the multimodal workstation [21]. A total of 12 subjects took part in this experiment, ten men and two women. Subjects were recruited based on whether they had UAV experience, military and/or pilot. The subject population consisted of undergraduates and graduate students, as well as those from the local reserve officer training corps (ROTC) and active duty military personnel. All were paid $10/h. In addition, a $50 incentive prize was offered for the best performance. The age range of participants was 20–42 years with an average age of 26.3 years. Nine participants were members of the ROTC or active duty officers. While no subjects had operational UAV experience, nine participants had piloting experience. The average number of flight hours among this group was 120.

Subjects had two main objectives in this experiment: 1) to guide each UAV so that together, all UAVs properly executed the required missions, which changed over time; and 2) to answer periodic questions about the situation from commanders, which was a secondary task. All subjects received between

90 and 120 min of training which culminated in a 15-min training scenario that is similar to, but easier in difficulty than, the testing sessions. Subjects were required to achieve a predetermined minimum level of performance in this scenario, making no more than two replanning or targeting mistakes, to move onto the actual test sessions.

Following the training, participants tested on two consecutive 30-min sessions, which represented low and high workload scenarios. These were randomized and counterbalanced to prevent a possible learning effect. The low replanning condition contained seven replanning events, whereas the high replanning condition contained 13 replanning events. These numbers were based on objective and subjective feedbacks gained through pilot testing. Each simulation was run several times faster than real time so that an entire strike could take place over 30 min (instead of several hours as is commonplace in actual strikes).

B. Experimental Design

The primary independent variable of interest in this experiment was the level of decision support (Fig. 3). A secondary-factor level of replanning was included to simulate the low and high workload scenarios. The level of decision support was between subjects, and the level of replanning was within subjects. The statistical model used for the majority of the analyses was a $2 \times 4(3)$ repeated measure linear mixed model. Three subjects were nested within each automation level, and both independent factors were fixed while subjects represented a random factor. Age was used as a covariate in all analyses and for all reported results $\alpha = 0.05$ unless stated otherwise. The dependent variables were the different components of WT (WTI, WTQ, and WTSA) and, as will be described in more detail next, were captured through mouse clicks and data file time stamps of screen objects.

V. RESULTS

A. WTI

In this experiment, WTI was recorded from the time an operator selected a degraded UAV to the time a solution was communicated to the UAV (mean = 81.5 s; SD (standard deviation) = 57.6 s). WTI was significant for LOA ($F(3, 13) = 8.08; p = 0.003$) and not significant for the level of replanning ($F(1, 12.1) = 2.13; p = 0.170$). From posthoc analyses, two homogeneous subsets of automation levels were found, superactive/active and passive/manual, which means that the superactive and active LOAs produced statistically the same WTs, so did the passive and manual. Fig. 5 shows how these results were compared to predictions, which were generally in line with expectations. It is worth noting for these and the remaining results that the relatively low subject number limits the generalizability of these results.

B. WTQ

WT spent in the queue was recorded whenever two or more vehicles required operator attention simultaneously, and the operator moved immediately to the vehicle(s) waiting in the

Fig. 5.   WTIs results.

Fig. 6.   WT in the queue results.

queue after finishing a task with another vehicle (mean $=$ 35.0 s; SD $=$ 43.4 s). Fig. 6 shows the results being compared with the predictions. The amount of WT spent in the queue was significant for the level of replanning ($F(1, 12.3) = 18.08$; $p = 0.001$) and marginally significant for LOA ($F(3, 13.2) = 3.10; p = 0.063$). The unusual spike in WTQ for the active LOA under high replanning in Fig. 6 warranted further investigation which showed that for the high replanning condition, the active level had significantly more WTQ than both the superactive and passive levels (A (active) versus SA (superactive), $p = 0.009$; A versus P (passive), $p = 0.074$). In addition, those

Fig. 7.   WTSA results.

subjects with an active decision support experienced an equal amount of WTQ as those with the manual level.

C. WTSA

WTSA was recorded when at least one vehicle required attention, but the operator did not realize it (mean $=$ 263.7 s; SD $=$ 239.7 s). This time interval was marked by operator inactivity, i.e., if a problem was present, and the operator did not attend to it and was not engaged in any other task. Common situations, where WTSA was incurred, included when subjects forgot to arm and fire on a target and left a UAV loitering unnecessarily or when a subject flew a UAV into a threat area without attempting to redirect it to a safe area. WTSA was significant for the level of replanning ($F(1, 12.3) = 18.70; p = 0.001$) but not for LOA ($F(3, 13.2) = 2.14$; $p = 0.144$). However, as shown in Fig. 7, there was a significant difference in cell means under the high replanning condition between the active and superactive levels ($p = 0.046$). Interestingly, under manual automation, there was no significant difference in WTSA between the high and low replanning conditions.

VI. DISCUSSION

In general, there was a decreasing trend of WTI with increasing LOAs, which was reasonably consistent with expectations (Fig. 5). Similarly, the WTQ results approximately followed the predicted trend of decreasing WTQ with an increasing automation level, except for the relatively high-average WTQ for active automation (Fig. 6). A quantitative analysis of WTQ accumulation in the active high replanning test sessions showed that the majority were accumulated in several large queues that formed late in the scenarios when multiple difficult replans were required from the operator.
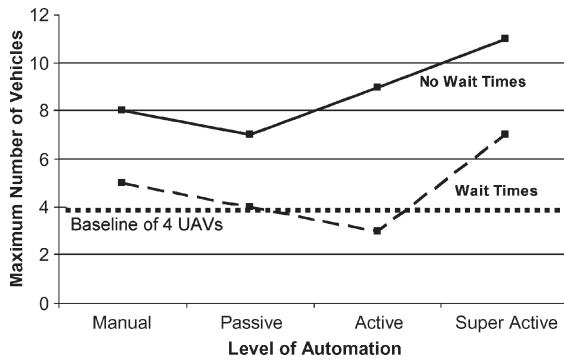
Fig. 8. WT proportions.

Fig. 9. Operator capacity predictions. Original FO (1) versus revised FO (3) including WTs.

Further investigation of the anomalous WTQ active condition demonstrated that this problem can be attributed to the subjects' inappropriate use of the "Request TOT change" function. It appeared that for those operators in the active condition, when not engaged in a specific task, they fixated on resolving possible areas of predicted high workload [22]. Because of this fixation, the WTSA predictions did not match the actual results (Fig. 7). Subjects using the active LOA spent significantly more time attempting to achieve an optimal schedule based on computer recommendations, which led to a loss of SA and subsequent significantly higher queuing times.

As can be seen in Fig. 8, of all of the proposed WT components, WTSA contributes significantly more than the other elements under both low and high workload conditions. There was no significant difference in the relative proportions of WTI, WTQ, or WTSA across the different LOAs. The proportion of WTSA is made up of 63% of all WTs in the low replanning condition and 72% in the high replanning condition, showing a trend of an increasing proportion of WTSA with an increasing workload.

Fig. 9 shows how the predictions from (1) would change for the high workload condition using the data from this experiment if the proposed WTs are included, as in (3). By including WTs in the prediction model, the upper bound of the number of homogeneous independent unmanned vehicles that a single controller can manage drops by 36%–67% based on the effectiveness of the decision support. More importantly, even under a high LOA represented by the superactive MBE decision support, subjects still experienced WTs due to a loss

of SA that then led to a reduction in predicted capacity (36%). This significant result demonstrates the importance of including human decision-making limitations in an upper bound prediction model, even for a highly automated system. In addition, the increasing WTSA for the lowest three LOAs indicates that operators were approaching their workload limit, culminating in the cognitive saturation of operators under the active condition, who were not effectively controlling the number they were assigned (Fig. 9). In the active condition of the experiment, the operators clearly struggled with the multiple-vehicle management task driven primarily by WTSA (Figs. 7 and 8). This relationship was captured in the upper limit predictions in Fig. 9, which was not captured in the original FO (1) because it does not capture inherent inefficiencies in human decision making and action.

These results illustrate how critical it is to ensure that when predicting an upper limit for operator capacity based on temporal attributes, the impact of specific automation decision support designs on decision times and SA is critical. These results also highlight the fact that any model that attempts to capture temporal measurements of human–automation interaction is both subject to the task and the context. The significance of the level of replanning across the WTs demonstrates that one external environmental variable can dramatically influence the results; therefore, caution should be exercised in generalizing specific results to other domains.

One additional source of WT that deserves further investigation not addressed in this paper is WT due to cognitive reorientation (WTCR), which is a component of WTI. WTCR is the time it takes for an operator to regain the correct mental model, and the SA is needed to solve a problem once recognized. Thus, WTCR is primarily a component of WTI but could be related to WTSA in that operators with lower WTSAs will likely have lower WTCR. WTCR represents a switching cost that occurs when an operator is cognitively engaged in one task but then must spend some period of time reorienting to a new problem. This environment of switching attention between multiple vehicles means that operators must not only concentrate their attention on a primary task, such as higher level mission planning, but also be prepared for alerts for other events, such as a vehicle system failure. This need to concentrate on one task yet maintain a level of attention for alerts/information from other potential tasks causes operators to have a conflict in mental information processing.

WTCR and the associated switching costs are very difficult to capture as the performance-based simulations, such as the one reported in this UAV study, cannot accurately and consistently categorize WTCR. These difficulties have also been shown in the multiple control of ground robots [9]. By using a software that tracked the users' cursor movements and activation of control devices, we were able to determine when a subject was engaged in a particular UAV, but we were unable to determine at what point the cognitive reorientation occurred, which is a subtle transition. This inability to partition WTCR from the overall category of WTI could possibly be addressed through some psychophysiologic measure or perhaps a more carefully crafted scenario. This area deserves more attention because it is unclear how much WTCR contributes to the overall WTs

and how technological interventions such as intelligent decision support could possibly mitigate this and other WTs.

VII. CONCLUSION

This paper extends a previous work attempting to predict the number of homogeneous and independent unmanned vehicles that a single operator can control. We propose that any predictive model of operator capacity that includes human-in-the-loop remote interaction should include various sources of WTs which include WT due to human–computer interactions (including cognitive reorientation), WTQ, and WTSA. By using the data from a simulation examining control of multiple homogeneous and independent UAVs, capacity predictions that included these sources of delay dropped by up to 67%, with a loss of SA as the primary source of WT delays. Furthermore, this paper demonstrates that both interface design components, as well as the level of decision support automation, can be a significant contributor to the overall WTs. Moreover, even in a highly automated MBE system which should alleviate queuing and ITs, this paper demonstrated that operator capacity is still affected by WTSA, causing a 36% decrease over the capacity model with no WT. Further work is needed to more accurately capture the cognitive reorientation times which represent a switching cost as well as loss of WTSAs.

ACKNOWLEDGMENT

REFERENCES

[1] Office of the Secretary of Defense, *Unmanned Aircraft Systems (UAS) Roadmap, 2005–2030*, Washington, DC: DoD, 2005.

[2] M. L. Cummings and S. Guerlain, "Developing operator capacity estimates for supervisory control of autonomous vehicles," *Hum. Factors*, vol. 49, no. 1, pp. 1–15, Feb. 2007.

[3] S. R. Dixon, C. D. Wickens, and D. Chang, "Comparing quantitative model predictions to experimental data in multiple-UAV flight control," in *Proc. 47th Annu. Meeting Hum. Factors Ergonom. Soc.*, Denver, CO, 2003, pp. 104–108.

[4] H. A. Ruff, S. Narayanan, and M. H. Draper, "Human interaction with levels of automation and decision-aid fidelity in the supervisory control of multiple simulated unmanned air vehicles," *Presence*, vol. 11, no. 4, pp. 335–351, Aug. 2002.

[5] R. Parasuraman, S. Galster, P. Squire, H. Furukawa, and C. Miller, "A flexible delegation-type interface enhances system performance in human supervision of multiple robots: Empirical studies with RoboFlag," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 35, no. 4, pp. 481–493, Jul. 2005.

[6] S. A. Rehfeld, F. G. Jentsch, M. Curtis, and T. Fincannon, "Collaborative teamwork with unmanned ground vehicles in military missions," presented at the Human Computer Interaction Int., Las Vegas, NV, 2005.

[7] D. R. Olsen and M. A. Goodrich, "Metrics for evaluating human-robot interactions," presented at the Performance Metrics Intelligent Syst., Gaithersburg, MD, 2003.

[8] D. R. Olsen and S. B. Wood, "Fan-out: Measuring human control of multiple robots," in *Proc. CHI*, Vienna, Austria, 2004, pp. 231–238.

[9] M. A. Goodrich, M. Quigley, and K. Cosenzo, "Task switching and multi-robot teams," presented at the 3rd Int. Multi-Robot Systems Workshop, Washington, DC, 2005.

[10] J. W. Crandall, M. A. Goodrich, D. R. Olsen, Jr., and C. W. Nielsen, "Validating human–robot interaction schemes in multitasking environments," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 35, no. 4, pp. 438–449, Jul. 2005.

[11] W. B. Rouse, *Systems Engineering Models of Human-Machine Interaction*. New York: North Holland, 1983.

[12] D. K. Schmidt, "A queuing analysis of the air traffic controller's work load," *IEEE Trans. Syst., Man, Cybern.*, vol. SMC-8, no. 6, pp. 492–498, Jun. 1978.

[13] A. Chapanis, F. C. Frick, W. R. Garner, J. W. Gebhard, W. F. Grether, R. H. Henneman, W. E. Kappaif, E. B. Newman, and A. C. Williams, *Human Engineering for an Effective Air Navigation and Traffic Control System*, P. M. Fitts, Ed. Washington, DC: Nat. Res. Council, 1951.

[14] M. R. Endsley, "Design and evaluation for situation awareness enhancement," in *Proc. 32nd Annu. Meeting Hum. Factors Soc.*, Santa Monica, CA, 1988, pp. 97–101.

[15] M. Endsley, "Measurement of situation awareness in dynamic systems," *Hum. Factors*, vol. 37, no. 1, pp. 65–84, Mar. 1995.

[16] A. Andre and C. Wickens, "When users want what's not best for them," *Ergonom. Des.*, vol. 3, no. 4, pp. 10–14, Oct. 1995.

[17] M. D. Rodgers, R. H. Mogford, and B. Strauch, "Post hoc assessment of situation awareness in air traffic control incidents and major aircraft accidents," in *Situation Awareness Analysis and Measurement*, M. Endsley and D. J. Garland, Eds. Mahwah, NJ: Lawrence Erlbaum, 2000, pp. 73–112.

[18] T. B. Sheridan and W. L. Verplank, "Human and computer control of undersea teleoperators," MIT, Cambridge, 1978. Man-Machine Systems Lab. Rep.

[19] M. R. Endsley, "Situation awareness in aviation systems," in *Handbook of Aviation Human Factors*, D. J. Garland, J. A. Wise, and V. D. Hopkin, Eds. Mahwah, NJ: Lawrence Erlbaum, 1999.

[20] R. Parasuraman, T. B. Sheridan, and C. D. Wickens, "A model for types and levels of human interaction with automation," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 30, no. 3, pp. 286–297, May 2000.

[21] G. Osga, K. Van Orden, N. Campbell, D. Kellmeyer, and D. Lulue, *Design and Evaluation of WarfighterTask Support Methods in a Multi-Modal Watchstation*. San Diego, CA: SPAWAR, 2002. 1874.

[22] M. L. Cummings and P. J. Mitchell, "Automated scheduling decision support for supervisory control of multiple UAVs," *AIAA J. Aerosp. Comput., Inf., Commun.*, vol. 3, no. 6, pp. 294–308, 2006.

Mary L. Cummings (S'02–M'03) received the B.S. degree in mathematics from the U.S. Naval Academy, Annapolis, MD, in 1988, the M.S. degree in space systems engineering from the Naval Postgraduate School, Monterey, CA, in 1994, and the Ph.D. degree in systems engineering from the University of Virginia, Charlottesville, in 2003.

From 1988 to 1999, she was an officer and pilot for the U.S. Navy, where she was one of the Navy's first female Fighter Pilots. She is currently an Assistant Professor with the Aeronautics and Astronautics Department, Massachusetts Institute of Technology, Cambridge. Her previous teaching experience includes being an Instructor with the U.S. Navy, Pennsylvania State University, University Park, and an Assistant Professor with the Virginia Tech Engineering Fundamentals Division. Her research interests include human supervisory control, human-uninhabited vehicle interaction, bounded collaborative human–computer decision making, decision support, information complexity in displays, and the ethical and social impact of technology.

Paul J. Mitchell received the B.S. degree in mathematics and engineering from Queen's University, Kingston, ON, Canada, in 2002 and the M.S. degree in aeronautics and astronautics from the Massachusetts Institute of Technology, Cambridge, in 2005.

He is currently a Control Development Engineer with General Electric Energy in Power Generation Technology, Greenville, SC. His research and professional interests include human supervisory control, human–computer interfaces and decision support, particularly for uninhabited vehicles, and the application of advanced control methods to aerospace and power generation systems.