

OPERATOR INTERACTION WITH CENTRALIZED VERSUS DECENTRALIZED UAV ARCHITECTURES

Affiliation

M.L. Cummings
Massachusetts Institute of Technology
Cambridge, MA
USA

Abstract

There has been significant recent research activity attempting to streamline Unmanned Aerial Vehicle (UAV) operations and reduce staffing in order to invert the current many-to-one ratio of operators to vehicles. Centralized multiple UAV architectures have been proposed where a single operator interacts with and oversees every UAV in the network. However, a centralized network requires significant operator cognitive resources. Decentralized multiple UAV networks are another, more complex possible architecture where an operator interacts with an automated mission and payload manager, which coordinates a set of tasks for a group of highly autonomous vehicles. While a single operator can maintain effective control of a relatively small network of centralized UAVs, decentralized architectures are more scalable, particularly in terms of operator workload, and more robust to single points of failure. However, in terms of operator workload, the ultimate success of either a centralized or decentralized UAV architecture is not how many vehicles are in the network per se, but rather how many tasks the group of vehicles generates for the operator and how much autonomy is on board these vehicles. Task-based control of UAV architectures with higher degrees of autonomy (i.e., decentralized networks) can mitigate cognitive overload and reduce workload. Mutually exclusive boundaries for humans and computers in multiple UAV systems should not be the goal of designers for either centralized or decentralized architectures, but rather more effort needs to be spent in defining mutually supportive roles such that humans and computers complement one another.

Introduction

The use of unmanned aerial vehicles (UAVs), often referred to as drones, has recently revolutionized military operations worldwide and holds similar promise for commercial settings. The U.S. Air Force now has more UAVs than manned aircraft, small UAVs are now used worldwide by various first response and police units, and they can be found fighting forest fires, monitoring wildlife and possible poachers, in cargo missions, and even in entertainment.

UAVs require human guidance to varying degrees and often through several operators. Most military and government UAVs require a crew of two to be fully operational. Conventional stick-and-rudder skills have been replaced by point-and-click control so that traditional pilots are no longer needed to control such systems. Onboard automation currently determines the most efficient control response, which is true in many commercial aircraft. While one operator supervises the actual flight activity of the UAV (the 'pilot'), the other operator typically monitors the UAVs sensors, such as a camera, and coordinates with the 'pilot' so that he or she can maneuver the UAV for the best system response.

There has been significant recent research activity attempting to streamline UAV operations and reduce staffing in order to invert the current many-to-one ratio of operators to vehicles. This is important not just

for military operations, but also for future commercial operations where air traffic controllers will direct both manned and unmanned aircraft. This chapter will discuss the implications of this staffing inversion for multiple UAV control, particularly in terms of two UAV control architectures, centralized and decentralized. It should be noted that these two architectures are not mutually exclusive and that there really exists a continuum of architectures in between these two bookends. Moreover, while there are many aspects of control architectures that are critical to consider, this chapter focuses on the human implications of such control architectures.




Figure 1: Human Supervisory Control [1]

Operator Interaction In Centralized UAV Architectures

The shift from stick and rudder to point and click control in UAVs represents a shift in the role of humans from the need for highly rehearsed skill sets to more knowledge-based reasoning inputs. For UAVs and for fly-by-wire military and commercial aircraft, pilots are less in direct manual control of systems, but more involved in the higher levels of planning and decision-making, particularly for remote operations. This shift in control from lower level skill-based behaviors to higher-level knowledge-based behaviors is known as human supervisory control (HSC). HSC is the process by which a human operator intermittently interacts with a computer, receiving feedback from and providing commands to a controlled process or task environment, which is connected to that computer [1] (Figure 1).

In a centralized UAV control architecture, human supervisory control in UAV operation is hierarchical, as represented in Figure 2. The innermost loop of Figure 2 represents the basic guidance and motion control

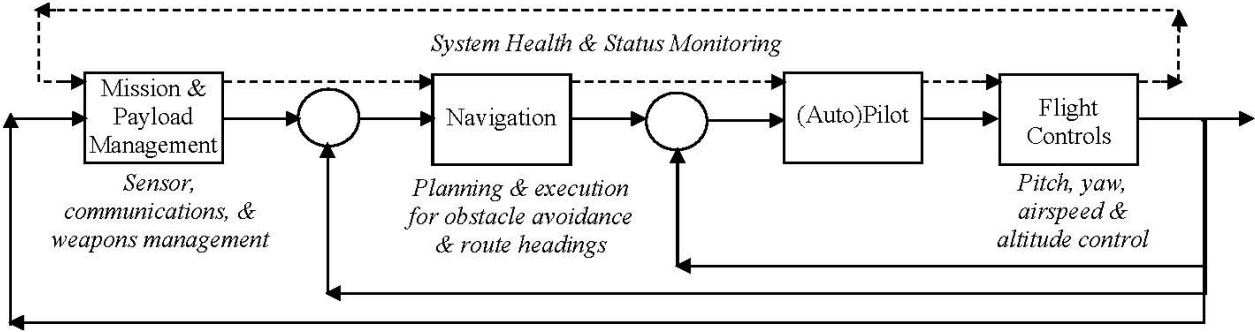


Figure 2: Hierarchical Control Loops for a Single UAV

loop, which is the most critical loop that must obey physical laws of nature such as aerodynamic constraints for UAVs. In this loop, the autopilot optimizes local control (keeping the aircraft in stable flight), and while UAV pilots could theoretically take control in this loop, with inherent time latencies that can cause pilot-induced instabilities, this loop is generally left to the automation.

The second loop, the navigation loop, represents the actions that some agent, whether human or computer, must execute to meet mission constraints such as routes to waypoints, time on targets, and avoidance of

threat areas and no-fly zones. In most current systems, humans enter GPS coordinates as waypoints, and then the system automatically flies to these waypoints. Only now are more advanced automated path planners that generate entire missions instead of serial waypoints starting to appear in operationally deployed UAV systems.

The outermost loop of Figure 2 represents the highest levels of control, that of mission and payload management. In this loop, sensors must be monitored and decisions made based on the incoming information to meet overall mission requirements. In the mission management loop, human operators provide the greatest benefit since their decisions require knowledge-based reasoning that include judgment, experience, and abstract reasoning that in general cannot be performed by automation.

Finally, the system health and status monitoring loop on the right side of Figure 2 represents the continual supervision that must occur, either by a human or automation or both, to ensure all systems are operating within normal limits. The control loop is dashed as it represents a highly intermittent loop in terms of the human, i.e., if the human is engaged in another task, with the highest priority given to the innermost loop, health and status monitoring becomes a distant, secondary task.

From the human-in-the-loop perspective, if the inner loops fail, then the higher (outer) loops will also fail. The dependency of higher loop control on the successful control of the lower loops drives human limitations in control of a single, and especially so, for multiple UAVs. If humans must interact in the guidance and motion control loop (manually fly a UAV), the cost is high because this effort requires significant cognitive resources. What little spare mental capacity is available must be divided between the navigation and mission management control loops. Violations of the priority scheme represented in Figure 2 have led to numerous crashes [3]. When operators become cognitively saturated or do not correctly allocate their cognitive resources to the appropriate control loops in the correct priorities, they violate the control loops constraints, potentially causing catastrophic failure.

Operator Capacity in Centralized UAV Architectures

In centralized UAV systems supervised by a human, the primary consideration for system design is how many vehicles a single controller can effectively supervise. Since this supervisor has to interact with each vehicle individually, just how many vehicles can be effectively and safely controlled will primarily be driven by the amount of autonomy on board the aircraft, which is subsumed across the four loops as shown in Figure 2.

By increasing UAS autonomy, operator workload will theoretically be reduced as it could reduce the number of tasks for the operator, and it should reduce the level of interaction even at the highest levels of control in Figure 2. For example, those UAVs that are flown in an autopilot mode relieve the operator from the manual flying tasks that require significant cognitive resources. This frees the operator to perform other critical tasks like mission planning and imagery analysis.

While there have been many studies that have attempted to experimentally derive the number of UAVs a single operator can control in a given setting (e.g., [4-6]), model-based approaches are generally more useful in determining not only an upper bound, but also provide insight into how much autonomy will be needed if a certain number of UAVS in a system is desired.

One such approach is the Fan-Out approach [7], which predicts the upper bound of the number of vehicles a single operator can control given the amount of autonomy in a vehicle as represented through its Neglect time (NT) and the required human-computer interaction called Interaction Time (IT). The Fan Out model was later modified to account for wait times that would inevitably be experienced by the system due to human inefficiencies. This adjustment provides a more realistic and lowered upper bound [2]. For example, one UAV can operate in a period of NT, and during that period of NT, the operator can attend to other vehicles. However, if an operator fails to notice that a UAV needs assistance (such as needing a new goal once a

waypoint has been achieved), or becomes engrossed in a mission plan for a UAV with a system malfunction, one or more UAVs can wait for the operator’s attention, causing a delay for one or more vehicles.

Figure 3 represents how NT, IT, and Wait Times (WT) interrelate. Point A represents a discrete event that occurs after a period of neglect time, which causes the vehicle to require immediate operator assistance such as an engine loss. NTs may not be so clearly observable, as exemplified by Point B, which represents performance degradation causing vehicle performance to drop below the NT performance threshold, e.g., a slow degradation of an inertial navigation system. In both NT cases in centralized UAV architectures, once performance has dropped below an acceptable level requiring human interaction, the UAV must wait until the operator recognizes and solves the problem, and so that the UAV can move to another NT state. Point C illustrates the system time delay if the problem is not addressed at the appropriate time.

Equation 1 represents the Fan Out mathematical relationship where NT and IT are as defined above. However, it should be noted that IT should account for not just the time an operator inputs commands to a UAV but also it should include delays where an operator has entered a command and waits for a system’s response. Wait times that add delays to the necessary ITs are accounted for in a separate term as defined by Equation 2.

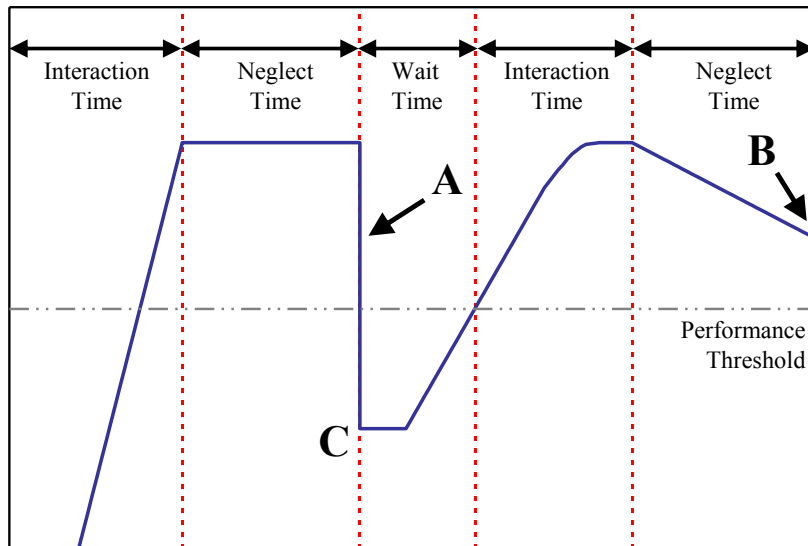


Figure 3: The relationship between interaction, neglect, and wait times [2]

$$FO = \frac{NT}{IT + WT} + 1 \quad (1)$$

$$WT = \sum_{i=1}^X WTQ_i + \sum_{j=1}^Y WTSA_j \quad (2)$$

- WT: Wait Time
- WTQ: Queuing wait time
- WTSA: Wait time caused by a loss of situation awareness
- X = Number of human-automation interaction queues that build
- Y = Number of time periods in which a loss of SA causes wait time

In Equation 2, WTQ, or Wait Time due to Queue results when multiple UAVs require attention, but the operator can only serially attend to them, effectively causing a queue to form for the operator’s attention. For example, if an operator is controlling two UAVs on a search mission and both require the operator to insert

waypoints near-simultaneously, the second UAV may have to loiter in place while the operator attends to the first. Assuming the operator can switch attention quickly once the first UAV is back to NT, the time the second UAV waits in the queue (WTQ) is effectively the IT for the first vehicle.

Table 1. Levels of Autonomy

LOA	Automation description
I	The computer offers no assistance: human must take all decision and actions.
II	The computer offers a complete set of decision/action alternatives.
III	The computer offers a selection of decisions/actions.
IV	The computer suggests a plan, and executes that suggestion if the human approves (management by consent)
V	The computer suggests a plan and allows the human a restricted time to veto before automatic execution (management by exception).
VI	The human is not involved in the decision-making process, the computer decides and executes autonomously.

near-simultaneous arrival of events that require human intervention plus the wait times due to the operator loss of SA. Wait times increase overall IT and reduce the number of vehicles a single operator can supervise.

In terms of operator capacity, Equation 2 demonstrates that as a UAV's degree of autonomy increases (expressed as an increase in NT), holding all other parameters equal mean that a single operator could control more vehicles. Consequently, for a fixed NT (or degree of autonomy), operator capacity could be increased by making the ground control station interactions more streamlined (i.e., lower IT), or ensuring all the correct alarms are in place so that operators do not miss critical points of interventions (i.e., lower WTSA).

A meta-analysis of several previous studies looking at various levels or degrees of automation/autonomy in centralized single operator control of multiple UAVs, particularly at the mission management level (the outermost loop in Figure 2), demonstrates that as NT increases (meaning the degree of autonomy increases), the maximum number of unmanned vehicles a single operator can control increases [2] (Tables 1 and 2, and Figure 4). The Levels of Autonomy (LOAs) in Table 1 are loosely modeled on Sheridan's levels of automation framework [1]. While there have been numerous other levels of automation and autonomy proposed for UAVS, as recently highlighted by a 2012 Department of Defense report [18], such frameworks

WTSA, or Wait time due to a loss of Situation Awareness, is perhaps the most difficult wait time component to model because it represents how effectively an operator can manage his or her attention. Situation Awareness (SA) 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 [15]. While SA can decrease under high workload due to competition for attentional resources [16], it can also decrease under low workload due to boredom and complacency [17]. If an operator does not realize a UAV needs attention (and thus experiences a loss of SA), the time from the initial onset of the need for IT to actual operator recognition of the problem could range from seconds to minutes.

Thus, Equation 2 categorizes system wait times as the summation of wait times that result from queues due to

Table 2. Multiple UAV Study Comparison

	Experiment	LOA	Max UV#
1	Dixon et al. [8] (baseline)	I	1
2	Dixon et al. [8] (auto-pilot)	I	2
3	Dixon et al. [8] (auto-alert)	IV	2
4	Ruff et al. [9, 10]	IV	4
5	Dunlap [11]	IV	4
6	Cummings, et al. [2, 12]	III-IV	5
7	Lewis et al. [13]	IV-V	8
8	Cummings & Guerlain [5]	IV	12
9	Hilburn et al. [14] (ATC)	N/A	11

pose many problems. These LOAs used here only illustrate increasing degrees of autonomy and thus NT, and are not meant to be normative.

As shown in Figure 4, research has previously demonstrated that with very low levels of mission management automation, a single operator can supervise at best only two UAVs. However, given high neglect times enabled by higher degrees of autonomy (i.e., UAVs plan their own routes and obstacle avoidance, only seeking high level mission plan approval), experimentally operators have successfully




Figure 4: Meta-Analysis of Previous Experiments Demonstrating Increased Operator Capacity for Increasing Neglect Time (as expressed by Levels of Automation)

controlled up to 12 Tomahawk Land Attack Missiles (TLAM), which are highly automated missiles that navigate on their own. These vehicles are effectively one-way UAVs, controlled in much the same way as UAVs through GPS commands. WASMs, or Wide Area Search Munitions, are also similar weapons that are launched from another aircraft, and then fly themselves until they find their assigned target. Curiously the number of WASMs and TLAMs, which are highly automated and operate in centralized control systems, a single operator can simultaneously control (8-12) is very similar to the number of airplanes a single air traffic controller can handle (~11, [14]). Arguably, manned aircraft have similar degrees of autonomy since the pilot on board is expected to obey all the high levels commands of the controller.

One of the limitations common across the studies in Table 2 is the lack of measurable *system-level* performance metrics. In general for the studies in Table 2, the performance of the operators was deemed acceptable as a function of expert observation, which is a valid method for performance assessment [19] but is not generalizable across domains and only useful as a descriptive and not predictive metric. Thus a system-level performance metric should capture both aspects of human and automation performance, which indicates an objective level of goodness and/or satisficing [20] (i.e., a “good enough” solution as opposed to optimal.) Such system-level metrics are often referenced as key performance parameters (KPPs) [21].

Towards this end of developing more comprehensive KPPs for multiple UAV systems, a recent study demonstrated that the number of UAVs that a single operator can control in a centralized architecture is not just a function of the level of decision support automation, but is inextricably tied to both mission complexity and overall system performance [12]. Using human experimentation in a multiple UAV simulation test bed and a simulated annealing (SA) technique for heuristic-based optimization, operator performance was predicted to be significantly degraded beyond approximately five UAVs with approximately levels 3-4 of autonomy as defined in Table 1. The optimal range was predicted to be between 2-4 vehicles (Figure 5).

Interestingly, in a different single operator-multiple UAV study with an entirely different test bed but similar levels of autonomy and centralized architecture, the optimal number was experimentally determined to be ~ 4 UAVs [2].

The KPP in Figure 5 is cost, which takes into account not just operational costs such as fuel, but also the cost of missed targets and cost in terms of mission delays introduced by inefficient human interactions. The solid curve in Figure 5 represents a theoretically perfect human operator, and the dotted line represents more realistic human performance that accounts for delays due to inefficient decision making, communication problems, cognitive load, etc. Thus, the performance of the system (the automation *and* the operator) can vary both as a function of the operator, but also can vary due to the operational constraints such as number of targets, operational costs, etc. This variation is why it is important to explicitly link system performance to operator capacity.




Figure 5: Operator Capacity as a Function of Mission Constraints

Operator Interaction in Decentralized UAV Architectures

While Figure 2 demonstrates supervisory control at the single vehicle level, which for centralized multiple UAV control is simply replicated for each vehicle under control, Figure 6 represents a notional system architecture that will be required for single operator control of multiple decentralized UAVs. In order to achieve this futuristic system, operators will need to interact with an overall automated mission and payload manager, which coordinates a set of tasks for a group of vehicles, instead of individually tasking each vehicle. This effectively represents a decentralized architecture, where operators convey high level goals to an automated mission manager (such as requesting that an area be searched), which then allows the UAVs to coordinate across the group to determine how to assign particular tasks, which may be dynamic. In a decentralized architecture, navigation and motion control tasks are necessarily subsumed by automation.

The decentralized architecture provides a substantial benefit in that the operator and his or her ground control station does not become a single point of failure, i.e., if the operator has intermittent or loss of communications with the vehicles, the system can still function. For example, because the network of

vehicles communicate with one another, if one vehicle breaks down, another can take its place. Another advantage is that the system is robust to lapses in operator situation awareness and delays since vehicles do not necessarily have to wait for commands. However, emergent UAV behavior in such systems can be complex and confusing for an operator, and if the system operates in a sub-optimal fashion, it could be difficult for operators to correct problems unless they have the ability to understand and then execute the necessary commands to correct the system.

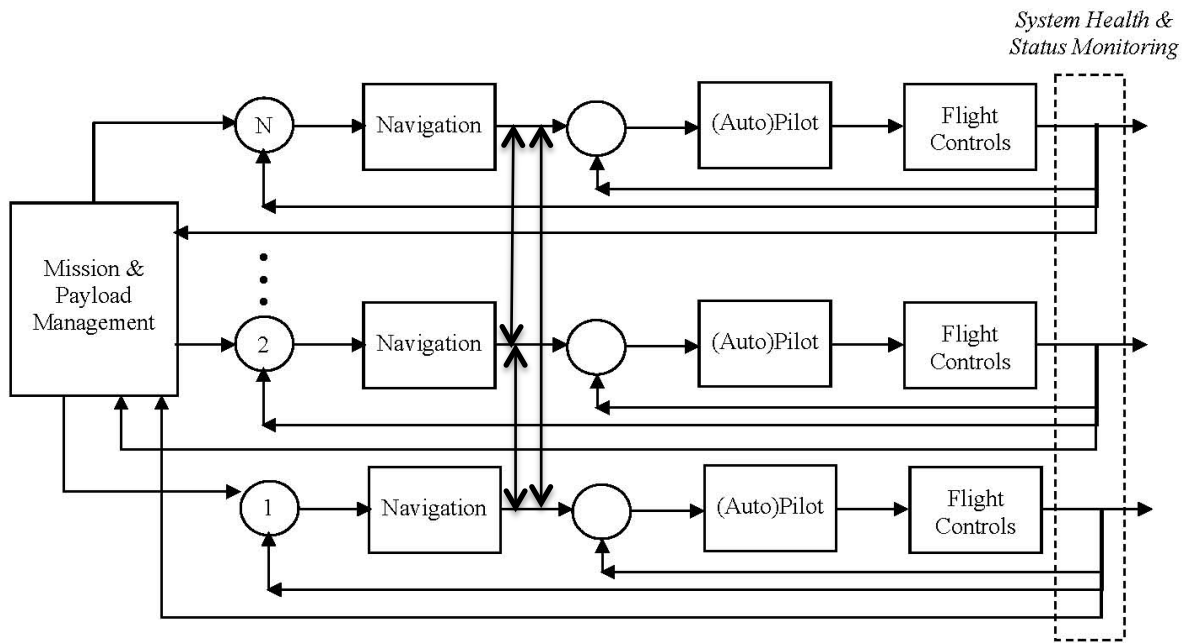


Figure 6: Decentralized Control for Multiple UAVs

For operators supervising a decentralized system, the Fan Out approach as depicted in Equation 1 cannot be used to estimate operator capacity, since for centralized systems, the assumption is that the vehicles have their own independent NTs and ITs, which drives the overall number of vehicles that can be controlled. Since operators only provide high level goals at the mission and payload management level, they do not have an IT for each vehicle, but rather an IT for high level interaction with the team. Similarly for NT, there is not distinct per vehicle NT, since they work together.

In control of a decentralized UAV network, the question of operator capacity is driven by how many *tasks* an operator can handle instead of how many *vehicles*. Under task-based, decentralized control, a human operator provides high-level control by approving which tasks should be completed by the team of vehicles without directly tasking a particular vehicle. Then the decentralized network of vehicles chooses how to allocate the approved tasks among themselves and can make tactical-level changes on their own, such as switching tasks.

In controlling a network of collaborative, decentralized UAVs, the operator could control, for example, 2, 20, 200 or even 2000 UAVs, as long as the tasks generated by the group of UAVs was manageable by a single operator. Determining the task load manageable by a single operator can roughly be thought of as the number of tasks that can be successfully accomplished over the course of the mission. While this number will, of course, vary widely across different missions and with different mixes of vehicles and people, there is one proxy metric that can be used across any decentralized UAV system for workload comparison, which is the concept of utilization.

Utilization refers to the “percent busy time” of an operator, i.e., given a time period, the percentage of time that person is busy. In supervisory control settings, this is generally meant as the time an operator is directed

by external events to complete a task (e.g., replanning the path of a UAV because of an emergent target), or attending to internal tasks like responding to text messages. What is not included in this measurement is the time spent monitoring the system, i.e., just watching the displays and/or waiting for something to happen. The concept of utilization as a mental workload measure has been used in numerous studies examining supervisory controller performance [5, 22-24]. These studies generally support that when tasked beyond 70% utilization, operators' performances decline.

In terms of the previously discussed IT and NT terms, utilization can generally be thought of as $IT/(NT+IT)$ on the aggregate task level. It can be used to describe an operator's response to task load in centralized UAV architectures as well as decentralized, but it is especially useful for decentralized system analysis given the operator's interaction at the meta-level instead of the individual vehicle level, which reflects the architecture of Figure 6.

Table 3: Task Load, Workload, and Performance for Three Multiple UAV Architectures with Increasing Autonomy

Experiment	Architecture	Primary Task Load	Average Tasks/Min	Performance	Average Utilization
Unmanned Ground Vehicles [25]	Mostly centralized, with some navigation sharing	2 vehicles	1.70	Low	.59
		4 vehicles	4.02	Medium	.68
		6 vehicles	5.34	Medium	.70
		8 vehicles	8.09	Medium	.78
Unmanned Aerial and Underwater Vehicles [26]	Somewhat centralized with some autonomous path planning and goal selection, 5 vehicles	Same type vehicles	4.90	Medium	.52
		2 different types	5.59	High	.64
		All different vehicles	8.00	Low	.70
Unmanned Ground and Aerial Vehicles Medium [27] and High [28] Workload Studies	Decentralized, 5 vehicles	23.3 tasks	3.96	High	.44
		30 tasks	4.05	High	.48
		43.3 tasks	7.30	Medium	.65
		50 tasks	7.90	Medium	.67

In order to determine whether decentralized systems provide any utility in terms of reduced workload and improved performance as compared to centralized UAV architectures, a set of studies was compared that

span increasing degrees of autonomy and increasing task loads, summarized in Table 1. In the first experiment at the lowest degree of autonomy [25], mostly centralized, operators controlled 2-8 Unmanned Ground Vehicles (UGVs), which resulted in four different experiment levels. For this experiment set, the individual vehicles relied on the human for goal setting, but had some ability to share local navigation information with one another.

In the second experiment with three experimental levels [26], a single operator controlled various mixes of 5 multiple unmanned aerial and underwater vehicles, with slightly more autonomy in the sense that vehicles would not only path plan themselves, but if the operator did not assign an ultimate goal within a pre-specified time, the vehicles would assign themselves to the nearest target. However, the operator could override any individual vehicle and redirect not only its path, but its ultimate goal as well.

The last experiment with four increasing task load levels represented the highest degree of decentralization, in that operators could only specify a task list to a group of 5 unmanned aerial and unmanned surface ships. The vehicles negotiated amongst themselves through a consensus-based bundled algorithm which vehicle would be assigned to which task, and each vehicle determined its own route. The operator could only insert and reprioritize tasks, but never direct and individual vehicle. Two different studies were included that used this test bed which focused on medium [27] and high [28] levels of task loading.

In order to directly compare these different studies, the average tasks per minute were listed, which shows how many tasks in each experiment the operator was expected to complete over an average one-minute time interval. Performance scoring was aggregated into low, medium, and high categories since the performance metrics used in each of the sets of experiments could not be directly compared. Low means that for the specific study, that condition resulted in the worse performance, and respectively for the high performance ranking. Lastly, the average utilization, which is the percentage of time the operator was busy performing tasks required by the system, was also listed.

Figure 7 illustrates that for each of the 3 sets of experiments, utilization increased with an increasing task load, which is expected and also an internal validity check. In addition, Figure 7 also illustrates that as the degree of decentralization increases, (i.e., more autonomy across a network of vehicles and less direct control by a human operator), utilization decreases. Interestingly all three experiments had an experimental level of ~8 tasks per minute and the most decentralized architecture allowed the average operator to work less by

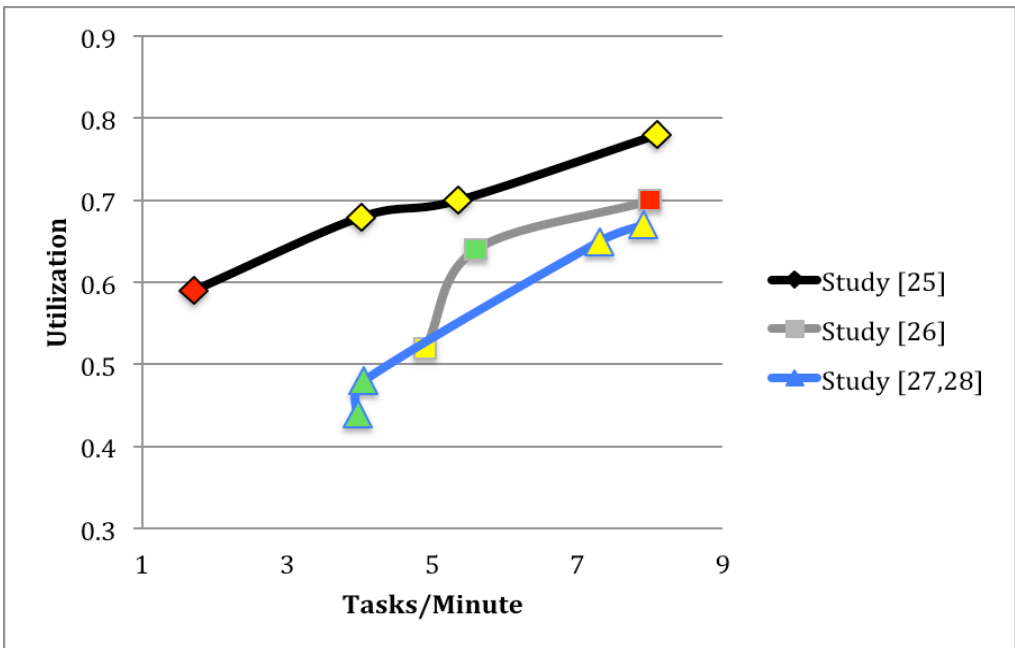


Figure 7: Utilization for the Average Tasks/Minute for the Experiments in Table 3

11% as compared to the more centralized architecture (and 3% less than the somewhat centralized architecture [26]). And in no case did the centralized architectures produce lower utilizations for similar task loads.

In terms of performance, each of the observed data points in Table 3 are color-coded in Figure 7 to reflect the relative performance score, with red indicating the worst performance, yellow demonstrates moderate performance, and green represents the best performance. Recall that each of these scores is a relative ranking so caution is advised in interpretation. In general the lower utilizations produced the best performances, with the caveat that little work has been done in terms of the possible negative impact of low task load on performance (in one exception, see [29]). For all studies, performance suffered when the 70% utilization threshold was exceeded, with only the somewhat centralized study participants exhibited markedly worse performance.

For the somewhat decentralized [26] and more decentralized experiments [27, 28], there appears to be a non-linear relationship in the 6 tasks/minute and below regions, which is not evident in the centralized experiment [25]. More work is needed to determine why such non-linear relationships exist, the nature of the critical point for the sharp rise, and how a system could be better designed to reduce the sharp increases in workload.

Conclusion

Taken together, the results from both the centralized and decentralized sets of experiments demonstrate that decentralized, task-based control of UAV architectures with higher degrees of autonomy can mitigate cognitive overload and reduce workload. While a single operator can maintain effective control of a relatively small network of centralized UAVs, decentralized architectures are more scalable, since adding additional agents also adds computational capability (assuming the tasks generated by the system do not linearly increase). Moreover, the decentralized UAV framework is robust to a single point of failure, since no single agent is globally planning for the fleet.

In terms of workload for a supervising operator, the ultimate success of either a centralized or decentralized UAV architecture is not how many vehicles are in the network per se, but rather how many tasks the group of vehicles generates for the operator and how much autonomy is on board these vehicles so that neglect time can be increased. And while this increasingly autonomy in a decentralized network of UAVs can mean reduced workload for the operator, it also adds significant more complexity to the system, which not only means increased developmental costs than a centralized network, but also one that is much harder to certify as safe.

Another caveat to the use of increased autonomy to mitigate workload across a UAV network is that such increased autonomy and increased neglect times can exacerbate a loss of operator situation awareness, as well as promote complacency and skill degradation [30]. Management-by-exception architectures, which occur when automation takes action based on some set of pre-determined criteria and only gives operators a chance to veto the automation's decision, have been shown to improve operator performance [31]. However, in such control schemes, operators are also more likely to exhibit automation bias, a decision bias that occurs when operators become over-reliant on the automation and do not check to ensure automated recommendations are correct [32]. Automation bias is a significant concern for command and control systems so it will be critical to ensure that when higher levels of automation are used, especially at the management-by-exception level, that this effect is minimized.

Lastly, while it is critical to consider the mental workload of a supervisor of multiple UAVs, the ability of the human to add value to the performance of a team of UAVs cannot be overlooked. In one study that

examined whether human supervisors added value in a search and track task for a decentralized, highly autonomous network of UAVs, results in a controlled study shows that a 30-50% increase in overall system performance, particularly in the search task, could be achieved by letting humans coach the automation [33]. Thus, instead of attempting to dictate mutually exclusive boundaries for human and computers in multiple UAV systems, either centralized or decentralized, more effort needs to be spent in trying to define mutually supportive roles such that humans and computers complement one another.

Acknowledgments

This work is a compilation and meta-analysis of data from many papers, which largely represent the thesis work of several MIT Humans and Automation Laboratory (HAL) students including Paul Mitchell, Carl Nehme, and Andrew Clare. However, many other MIT HALiens, including undergraduate, graduate and visiting students, were instrumental in framing this work and are the real engineers behind this work.

References

- [1] T. B. Sheridan and W. Verplank, "Human and Computer Control of Undersea Teleoperators," Man-Machine Systems Laboratory, Department of Mechanical Engineering, MIT, Cambridge, MA 1978.
- [2] M. L. Cummings and P. J. Mitchell, "Predicting controller capacity in supervisory control of multiple UAVs," *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, vol. 38, pp. 451-460, 2008.
- [3] K. W. Williams, "A Summary of Unmanned Aircraft Accident/Incident Data: Human Factors Implications," Civil Aerospace Medical Institute, Ed., Washington DC: Federal Aviation Administration, 2004.
- [4] H. 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, pp. 335-351, 2002.
- [5] M. L. Cummings and S. Guerlain, "Developing Operator Capacity Estimates for Supervisory Control of Autonomous Vehicles," *Human Factors*, vol. 49, pp. 1-15, 2007.
- [6] S. R. Dixon, C. D. Wickens, and D. Chang, "Comparing quantitative model predictions to experimental data in multiple-UAV flight control," in *Human Factors and Ergonomics Society 47th Annual Meeting*, Denver, 2003, pp. 104-108.
- [7] D. R. Olsen and S. B. Wood, "Fan-out: Measuring Human Control of Multiple Robots," presented at the SIGCHI conference on Human factors in Computing Systems, Vienna, Austria, 2004.
- [8] S. Dixon, C. Wickens, and D. Chang, "Mission Control of Multiple Unmanned Aerial Vehicles: A Workload Analysis," *Human Factors*, vol. 47, pp. 479-487, 2005.
- [9] 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, pp. 335-351, 2002.
- [10] H. A. Ruff, G. L. Calhoun, M. H. Draper, J. V. Fontejon, and B. J. Guilfoos, "Exploring Automation Issues in Supervisory Control of Multiple UAVs," in *2nd Human Performance, Situation Awareness, and Automation Conference (HPSAA II)*, Daytona Beach, FL, 2004.
- [11] R. D. Dunlap, "The evolution of a distributed command and control architecture for semi-autonomous air vehicle operations," in *Moving Autonomy Forward Conference*, Grantham, UK, 2006, pp. 107-123.
- [12] M. L. Cummings, C. E. Nehme, ., and J. Crandall, "Predicting Operator Capacity for Supervisory Control of Multiple UAVs," in *Innovations in Intelligent Machines* vol. 70, J. S. Chahl, L. C. Jain, A. Mizutani, and M. Sato-Ilic, Eds., 2007.
- [13] M. Lewis, J. Polvichai, K. Sycara, and P. Scerri, "Scaling-up Human Control for Large UAV Teams," in *Human Factors of Remotely Operated Vehicles*, N. Cooke, H. Pringle, H. Pedersen, and O. Connor, Eds., New York: Elsevier, 2006, pp. 237-250.

- [14] B. Hilburn, P. G. Jorna, E. A. Byrne, and R. Parasuraman, "The Effect of Adaptive Air Traffic Control (ATC) Decision Aiding on Controller Mental Workload," in *Human-automation Interaction: Research and Practice*, Mahwah, NJ: Lawrence Erlbaum, 1997, pp. 84-91.
- [15] M. R. Endsley, "Toward a Theory of Situation Awareness in Dynamic Systems," *Human Factors*, vol. 37, pp. 32-64, 1995.
- [16] A. Andre and C. Wickens, "When Users Want What's Not Best for Them," *Ergonomics in Design*, vol. October, 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 Associates, 2000, pp. 73-112.
- [18] Defense Science Board, "The Role of Autonomy in DoD Systems," Department of Defense 2012.
- [19] M. R. Endsley and D. J. Garland, *Situation Awareness Analysis and Measurement*. Mahwah, NJ: Lawrence Erlbaum Associates, Inc., 2000.
- [20] H. A. Simon, R. Hogarth, C. R. Piott, H. Raiffa, K. A. Schelling, R. Thaler, *et al.*, "Decision Making and Problem Solving," in *Research Briefings 1986: Report of the Research Briefing Panel on Decision Making and Problem Solving*, Washington D.C., 1986, p. 16.
- [21] Joint Chiefs of Staff, "Chairman of the Joint Chiefs of Staff Instruction 6212.01D," ed: DoD, 2007.
- [22] D. K. Schmidt, "A queuing analysis of the air traffic controller's workload," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 8, pp. 492-498, 1978.
- [23] W. B. Rouse, *Systems Engineering Models of Human-Machine Interaction*. New York: North Holland, 1983.
- [24] B. Donmez, C. Nehme, and M. L. Cummings, "Modeling Workload Impact in Multiple Unmanned Vehicle Supervisory Control " *IEEE Systems, Man, and Cybernetics, Part A Systems and Humans*, vol. 99, 2010.
- [25] C. Nehme, B. Mekdeci, J. W. Crandall, and M. L. Cummings, "The Impact of Heterogeneity on Operator Performance in Future Unmanned Vehicle Systems," *The International C2 Journal*, vol. 2, 2008.
- [26] C. E. Nehme, "Modeling human supervisory control in heterogeneous unmanned vehicle systems," Doctor of Philosophy, Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA, 2009.
- [27] M. L. Cummings, A. Clare, and C. Hart, "The Role of Human-Automation Consensus in Multiple Unmanned Vehicle Scheduling," *Human Factors*, vol. 52, 2010.
- [28] A. S. Clare and M. L. Cummings, "Task-Based Interfaces for Decentralized Multiple Unmanned Vehicle Control," *AUVSI Unmanned Systems North America*, Washington DC, 2011.
- [29] M. L. Cummings, C. Mastracchio, K. M. Thornburg, and A. Mkrtchyan, "Boredom and Distraction in Multiple Unmanned Vehicle Supervisory Control," *Interacting with Computers*, vol. 25, pp. 34-47, 2013.
- [30] R. Parasuraman, T. B. Sheridan, and C. D. Wickens, "A Model for Types and Levels of Human Interaction with Automation," *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, vol. 30, pp. 286-297, 2000.
- [31] M. L. Cummings and P. J. Mitchell, "Automated Scheduling Decision Support for Supervisory Control of Multiple UAVs," *ALAA Journal of Aerospace Computing, Information, and Communication*, vol. 3, pp. 294-308, 2006.
- [32] K. L. Mosier and L. J. Skitka, "Human Decision Makers and Automated Decision Aids: Made for Each Other?," in *Automation and Human Performance: Theory and Applications, Human Factors in Transportation*, R. Parasuraman and M. Mouloua, Eds., Mahwah, New Jersey: Lawrence Erlbaum Associates, Inc., 1996, pp. 201-220.
- [33] M. L. Cummings, J. How, A. Whitten, and O. Toupet "The Impact of Human-Automation Collaboration in Decentralized Multiple Unmanned Vehicle Control," *Proceedings of the IEEE*, vol. 100, pp. 660-671, 2012.