

A Human-Centric Design Process for Highly Autonomous Unmanned Air Systems

Dr J T Platts, J A Whitby, Dr D Richards

QinetiQ Ltd

A7 Bldg, G069, Cody Technology Park
Farnborough, GU14 0LX, UK

jtplatts@QinetiQ.com

ABSTRACT

The paper will propose that for effective operation of highly complex and intelligent systems a human and machine team must be developed. It will detail research addressing delivery of unprecedented levels of unmanned system intelligence and the concomitant need to establish the operator in a teamed arrangement. It will discuss the impact of the teaming on the operator and so build the case for human system integration to be a necessary part of intelligent system design rather than it being treated as an afterthought. A case study will be described whereby technology has been developed over a period of 10 + years to enable 4 unmanned combat air vehicles to be controlled by a single operator by adopting a “teaming” approach. Trials in both the synthetic environment and flight will be outlined.

1.0 INTRODUCTION

This paper is formed from two previous papers of references [1] and [2]. They are combined here to act as the accompanying notes for part of the NATO Lecture Series SCI-208 on “Advanced Automation Issues for Supervisory Control in Manned-Unmanned Teaming Missions”.

Autonomous vehicles may fulfil a developing requirement in many applications and domains. Of particular interest in the defence sector, is the Unmanned Air Vehicle (UAV). UAVs applied in Combat roles (UCAVs) have the potential to reduce significantly the risk to aircrew in military operations and improve mission effectiveness. Moreover, in addition to the reduced risk to aircrew, the attraction of relatively low cost, highly reliable, readily available assets that are not subject to the physical, physiological and training constraints of human pilots, is self-evident. This promise has prompted studies of systems destined to provide future capability to be broadened to research UAV deployment and consider the possibility of using UCAVs in isolation, in swarms and in packages of mixed unmanned and manned platforms.

Existing UAV systems place a large reliance on remotely situated operators with some of the larger vehicles requiring a pilot in the loop for mission phases such as landing and take-off. High levels of operator interaction with the vehicle increase the vulnerability of the UAV due to detection of electromagnetic emissions and the increased likelihood of control datalink jamming. In addition, the already intense competition for scarce bandwidth is exacerbated. Moreover, reliance on a critical communications link has safety implications with regard to the need for system redundancy to protect against single point failure with concomitant cost implications. Consequently, to reduce this interaction, many of the air vehicle sub-systems must be largely autonomous. With military aspirations towards deploying UAVs in combat roles, such vehicles must be able to operate with minimal interaction from remote operators, communicating only when human decisions are a necessity to meet constraints such as those resulting from Rules of Engagement (ROE).

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE SEP 2009		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE A Human-Centric Design Process for Highly Autonomous Unmanned Air Systems				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) QinetiQ Ltd A7 Bldg, G069, Cody Technology Park Farnborough, GU14 0LX, UK				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES See also ADB381581. RTO-EN-SCI-208 Advanced Automation Issues for Supervisory Control in Manned-Unmanned Teaming Missions (Problemes de l'automatisation evoluee pour la supervision de missions associant des systemes habites et non habites)., The original document contains color images.					
14. ABSTRACT The paper will propose that for effective operation of highly complex and intelligent systems a human and machine team must be developed. It will detail research addressing delivery of unprecedented levels of unmanned system intelligence and the concomitant need to establish the operator in a teamed arrangement. It will discuss the impact of the teaming on the operator and so build the case for human system integration to be a necessary part of intelligent system design rather than it being treated as an afterthought. A case study will be described whereby technology has been developed over a period of 10 + years to enable 4 unmanned combat air vehicles to be controlled by a single operator by adopting a teaming approach. Trials in both the synthetic environment and flight will be outlined.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 28	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

A Human-Centric Design Process for Highly Autonomous UAS

The terms “autonomy” and its corollary “autonomous system” are much used and abused. Systems are often described as “semi” or “fully” autonomous, terms which carry little precision. The aim of this paper is to bring some clarity to the terms by describing one particular application and to outline the issues surrounding the achievement of a controllable capability which can be trusted. This paper discusses the meaning of autonomy in the context of autonomous UAVs by means of a conceptual architecture of a UCAV system. It will go on to discuss methods for calibrating autonomy by defining levels. It will also describe the proposed methods for achieving higher levels of autonomy in UAV systems. The barriers to achieving these unprecedented levels of autonomy will be discussed before concluding.

Since 1997, QinetiQ has been engaged in research addressing high levels of autonomy in air vehicles and in particular UCAVs. Initially the work was confined to low cost Synthetic Environment (SE) based experimentation where the focus was on the operator workload aspects of the problem and addresses what were the key human machine interface requirements for effective operator decision making. Later, the addition of machine based decision technology allowed the targeted introduction of autonomy to minimise workload and communication bandwidth requirements. From its early SE based requirements capture work it has now evolved into a populated architecture where the operator and the intelligent Unmanned Combat Air System (UCAS) work in partnership to carry out highly complex tasks.

From the outset the military customer played a key role in developing both the approach and requirements. At each stage of the development, prototype systems were realised and tested within the SE by military operators. One of the key requirements was to understand the trade between the role of the operator in terms of how much control was required for any given mission vignette, and the location of the operator. Two extremes have been examined, a ground (rear-echelon) based operator, and a single seat fast-jet based operator. These positions represent extremes of workload and, as it turns out, not too dissimilar in terms of HMI requirements.

Another key focus of the work was the need to examine the ratio of operators to air-vehicles. The target was for any single operator to control more than one vehicle. Achieving this required the level of command abstraction to be raised to the extent that the operator issues commands such as “[capability] search this area”. That is, the operator issues the command and delineates the area to be searched and the capability – in this case 4 UCAVs – organises the search in some optimal way. The operator is now less concerned about developing and delegating a “bread-crumbs” trail of waypoints for the respective vehicles with concomitant work load benefits.

The work has reached a degree of maturity such that the system has been trialed in flight using a Surrogate UAS and it is the purpose of this paper to give an overview of recent flight trials. Both the Human Machine Interface (HMI) and system intelligence must be increased to achieve suitable autonomy. Accordingly, Sections 2-5 will set out some of the theoretical background to the work whilst Section 6 will explain some key aspects of the machine intelligence technologies and the HMI. Section 7 will discuss SE based requirements capture and design progressing to the Surrogate Unmanned Air System (UAS) and the implementation of the autonomous system architecture on board. The paper will conclude by discussing some top level results and outlining future work.

2.0 WHAT ARE AUTONOMOUS UAVS?

The introduction outlined the need for autonomy in UAVs. To aid discussion this section will define the terms introduced and postulate the likely system architecture of an autonomous UAV enabled capability.

Turning first to some working definitions of commonly abused terms:

- Automatic - Performed from force of habit or without conscious thought [3].

A Human-Centric Design Process for Highly Autonomous UAS

- Autonomy - Ability to initiate or modify actions in the light of ongoing events. The freedom to determine one's own actions or behaviour. Self-governing. [3]
- Intelligence - Capacity for understanding or the ability to exploit knowledge to achieve goals. [3]

Autonomy and intelligence are terms often confused but Clough [4] makes the point that autonomy and intelligence are not necessarily linked. For example, the Amoeba is a very simple organism that can be described as being autonomous but is also possesses very little intelligence. So are high levels of intelligence required of UAVs?




Figure 1. Capability vs Certainty

White [5, 6] introduces the following argument. Figure 1 shows a number of systems plotted on a graph of flexibility & capability versus certainty. In essence, as one travels down the page and to the right, the assumption is that we are more certain of the environment and situation such that less sophisticated weapons can be deployed. Conversely as we move up the graph and to the left, we are less certain of the target and situation and so must field a more intelligent system able to deal robustly with the level of uncertainty and hence the addition of the crew to provide what humans currently do best. In the case of the UAV, the intelligence to cope with uncertainty is now shared between the on-the-spot vehicle and the remotely situated human operator. White goes on to describe how the system autonomy is therefore achieved via an appropriate partnership between the UAV operator and the intelligent vehicle system. This conclusion is implied in the following statement taken from [7] defining intelligent control:

“control that replaces the human mind in making decisions, planning control strategies, and learning new functions whenever the environment does not allow or does not justify the presence of a human operator.”


So how will this decision making partnership be embodied and what is the work share? The system will contain a control station incorporating the interface to the vehicle, system-monitoring functionality, and possibly some in-built mission planning capability and decision support functionality. The system will contain one or more datalinks to the vehicle or vehicles. The vehicles themselves will possess a systems manager to control aspects of flight management, system health, sensor payload etc. High levels of autonomy may require an additional system element performing the autonomous reasoning, planning, conflict-resolution and decision making functionality. Decisions will be shared across the system (partnership) as a function of system intelligence and communications bandwidth. Critically some

A Human-Centric Design Process for Highly Autonomous UAS

decisions will have to be made by human beings due to regulatory and liability issues regardless of technological capability and this implies that “total autonomy” in the context of lethal UAVs is a *non sequiter*.

Breaking down further the notional architecture into levels of hierarchy aids discussion of the various elements of technology that combine to deliver the overall autonomous UAV system. Consider first Figure 2 showing a single vehicle structure. A single operator communicates with the vehicle over a datalink. The vehicle autonomy is to be delivered using three conceptual hierarchical tiers. A description of the tier functionality is best achieved through use of an example. Whilst the structure encompasses all vehicle sub-systems the example focuses purely on the flight control aspects. In this respect Tier 1 is what is colloquially known as the “inner-loop” and concerned with stabilisation and control of the body rates in pitch roll and yaw of the air vehicle. Tier 2 would be the “outer-loop” and deals with such things as air vehicle trajectory control.

Tier 3 is the functionality traditionally brought to the air-vehicle by the aircrew and partially delivers the flexibility and robustness discussed already.



Optimised partnership between operator and intelligent platform

Figure 2. Operator vehicle partnership.

For example, focussing on the human contribution to the operation of a twin seat ground attack aircraft, whilst there are levels of automation within the cockpit for both the pilot (automatic pilot modes) and the navigator (Navigation systems), the scheduling and initiation of these systems is still carried out by the crew. Moreover, real-world uncertainties in the kinds of environment in which this aircraft operates frequently mean that sub-system failures and damage must be compensated for such that, ideally, the mission continues. In the case of the UAV this Tier 3 functionality is now shared between a remotely situated human and the UAV based capability. In order to formulate a view on who does what in the partnership it is necessary to understand how human beings think.

Rasmussen [8, 9] proposes a model of human cognition known as the Skills-Rules-Knowledge model. In this model low level reflexive and motor skills are generally executed without thought (changing gear in a car for example) and initiated according to established rules (approaching a junction so change down

A Human-Centric Design Process for Highly Autonomous UAS

through the gear-box). When unfamiliar situations arise, humans are adept at building new rules from knowledge to cope with the new situation (skidding on the approach to the junction for example). In the case of the UAV/human decision making partnership then it would seem logical to automate the lower level skills involved (Tiers 1 and 2) and then share as appropriate the higher level rules and knowledge levels. Unfortunately there is no clearly defined “decision surface” between these levels. Therefore in the case of UAVs, having removed the aircrew from the vehicle, the functionality lost must now be shared between any on-board intelligent capability (Tier 3) and the remotely situated operator. This sharing of the decision making between the operator and vehicle partnership will be discussed in greater detail later in what follows and is referred to in the literature as a mixed-initiative system.

Having discussed the single vehicle system architecture the aspiration is for a system where a single operator now controls a number of UAVs. The level of abstraction with which the operator interacts with the UAVs is increased by the addition of a Tier 4. Tier 4 carries out the “pooling” tasks and resource allocation. An assumption is that the operator now deals with an uninhabited capability comprising a number of platforms and is no longer concerned with the individual platforms. The intelligent system is defined such that the mission is designed around a pool of UAV resources by default and only operator intervention will take resources from the pool to divert them to additional tasks. Synthetic Environment based human in the loop trials have shown that this assumption is sound up to a point. There are instances where the autonomous behaviour exhibited by the “pool” is not exactly that required by the operator and so he or she will interrupt and begin to draw single platforms from the pool to carry out additional tasks. Clearly this requirement begins to indicate the level of sophistication required of Tier 4.




Figure 3. Operator capability partnership.

In order to understand the relative sophistication required of various autonomous UAV systems it is necessary to first explore more fully potential degrees of autonomy.

A Human-Centric Design Process for Highly Autonomous UAS

3.0 HOW DO WE MEASURE AUTONOMY IN UAVS?

The Autonomous Control Level (ACL) matrix was developed at the US Air-Force Research Laboratory Air Vehicles Directorate at Wright-Patterson Air Force Base as part of the Fixed-Wing Vehicle Initiative by researchers across US Services and industry [4]. The matrix is an attempt to place arbitrary levels on an autonomy scale but with the added dimensions of the "OODA" loop. The Observe Orient Decide Act loop was adopted in an attempt to map the autonomy level to a capability. Moreover, a number of technologies exist to increase autonomy levels whose strengths and weaknesses constrain their application. The added OODA dimension allows for comparison of such methods and subsequent judgements to be made in respect of their capabilities and likely contribution to increasing the level of autonomy. Indeed, it is understood that the matrix has been used as a Balance of Investment tool.

Strens [10] pointed out that the decision column of the ACL Matrix essentially provided definitions for the level descriptors. i.e. it was an input to the table. For example, ACL 7 (Battlespace Knowledge) is defined to mean capability for 'group accomplishment of strategic goal with minimal supervisory assistance'. The orient and decide requirements are the fundamental elements of autonomy, whereas the observe requirements must support these processes by providing the necessary information. Continuing his critique in the context of theories for optimal interaction used in robotics research he advanced alternative dimensions to autonomy having first defined an autonomous agent being one that can, without intervention, sustain robust behaviour for a period of time or until a goal is reached. He advanced that high autonomy implies accomplishment of a strategic group goal without intervention, and lower levels of autonomy imply intervention is required to set sub-goals (e.g. for different phases of a mission or different agents in a group). Clearly definitions of strategic goal will differ application to application.

Having first interpreted the ACL in the context of three dimensions of requirements (i) operator interaction; (ii) agent reasoning ability (to support this interaction); and (iii) the level of knowledge needed by the group of agents (to support this reasoning). He also went on to argue for a reduction of the number of levels and proposed the scale of autonomy shown at Table 1.

	Operator interaction	Reasoning	Knowledge	UAV task example
8	Long-term strategic goal	Strategic imperfect information game player. <i>e.g. has expectation about opponent's intention for operating air defences</i>	Broad battlespace understanding <i>e.g. has detailed models for command systems and reasons about their effective deployment</i>	<i>Long-term no-fly zone enforcement.</i>
7	Strategic goal	Battlespace planning or strategic behaviours <i>e.g. plans to achieve the strategic goal and adapts the plan in view of communicated/sensed info.</i>	Broad battlespace knowledge <i>e.g. generic predictive model for all entity types</i>	<i>SEAD with ALUAVs in high-intensity environment.</i>
6	Select from variety of tactical goals	Multiple specialised planning capabilities or behaviours. <i>e.g. separate search, surveillance, attack, BDI skills</i>	Generic and specialised knowledge. <i>e.g. tactical picture plus task-specific co-ordination models</i>	<i>SEAD with operator-intensive mission control.</i>
5	Tactical multi-UAV goals for heterogeneous UAV resource pool	Strong co-ordination <i>e.g. real-time spatio-temporal reasoning</i>	Local co-ordinative picture <i>e.g. relative position information or co-ordination functions</i>	<i>Search and attack of a mobile target given its last known location.</i>
4	Tactical goal for homogeneous UAV resource pool	Weak co-ordination <i>e.g. dynamic allocation of UAVs to tasks</i>	Abstract picture <i>e.g. platforms as resources, high-level task descriptions</i>	<i>Search a large region with an ALUAV swarm.</i>
3	Tactical goal for each UAV (specialised)	On-board re-planning <i>e.g. near real-time flight path planner</i>	Spatial <i>e.g. terrain and pop-up threats</i>	<i>Send ALUAV to reconnoitre a location.</i>
2	Mission planning, event-based monitoring	Adaptive control (self-regulation) <i>e.g. adapts trajectory; refers to operator if outside bounds</i>	Internal histories <i>e.g. distance, fuel against plan</i>	<i>UCAV vs fixed target, in all-weather.</i>
1	Mission planning, continuous monitoring	Plan execution <i>e.g. follow flight plan then deploy a weapon</i>	Specialised spatial <i>e.g. list of waypoints</i>	<i>UCAV vs fixed target in good weather</i>

Table 1. Alternative to the ACL (Strens 2002)

A Human-Centric Design Process for Highly Autonomous UAS

As has been described, at least two scales of UAV system autonomy exist. Focussing on the scale proposed for Table 1, at least two difficulties remain; (i) the mapping of technologies to an appropriate level and (ii) the mapping of levels to the achievable military capability. The former is an issue in that it has been shown that achievable levels of autonomy are a function of a number of different technologies of differing maturity level and the challenge of integration of those technologies. The lack of a well defined application requirement further adds to this challenge. In developing Table 1 it is clear that some assumption has been made about the "context" in which the system will be operated. This is a further dimension to autonomy level that allows meaningful discussion about the user's expectations of the system and comparison with other "autonomous" systems. Often autonomy is assumed to be a function of flight management alone.

In the case of the second difficulty it is not known what level of investment is required to optimise the mix of operator, available bandwidth (a dynamic quantity), technology maturity and the concomitant military capability. This lack of knowledge stems from the nature of current procurement programmes in that the user's requirements are, at this stage, aspirational depending largely on the "art of the possible". Without a true understanding of costs it is not a foregone conclusion that UAVs of the type discussed here are tenable in any future force-mix. The required concepts of use are based on operational assessment studies that are untested other than in SE trials. The nature of such trials makes them constrained and relatively expensive such that the entire mission space can not be explored. Consequently, it is extremely difficult to extrapolate across this mission space to the extent that a clear picture of the technological needs and capabilities is understood such that the cost versus capability relationship is apparent.

The key to use of Table 1 comes in deciding the command level of abstraction which in turn comes from understanding the "context" of the system's application. The UAV task column briefly describes some typical tasks at the appropriate level of abstraction. These tasks can be read across to the operator interaction required and the implied information requirements for the appropriate decision making. Thus giving insight into the required command and control architecture and required technologies.

4.0 HOW DO WE REALISE HIGHLY AUTONOMOUS UAVS?

Having discussed both the nature of UAV autonomy and the measurement of degree of autonomy, some ideas about how these required levels are achieved are presented. The methods consider the hierarchical Tiers already described and the interface to the human operator.

Two concepts associated with safety should be introduced at this stage of the discussion, mission critical and safety critical decisions. These concepts will be returned to later in this paper but for now it is worth pointing out that currently in manned aircraft, a proportion of the risks to safe flight are mitigated merely by the presence of the crew. For example, the "see & avoid" philosophy of collision avoidance, the ability of the crew to spot potential collision hazards is a cornerstone of current operations under Visual Flight Rules – the ability to minimise flight over populated areas or to manoeuvre clear in an emergency for example. Mission critical decisions are those whereby the success of the mission is compromised - due to sensor failure for example - but crew safety or collateral damage is not necessarily a risk. In the case of mission critical case the crew are required to make the appropriate decision to abandon the mission.

Safety criticality is a key issue and barrier to the deployment of highly autonomous UAVs. To this end, and to capitalise on many years research into the automation of aspects of manned flight, the philosophy adopted in the work underpinning this discussion is to use where possible, tried and tested technologies. In this respect Tiers 1 and 2 rely on well-established design methods for inner and outer loop control with regard to flight management aspects of autonomy. Similarly other system aspects of autonomy such as sensor control, communications system control, weapon control and so on may not have specific parallels in the manned flight world but do not, it is felt, represent significant technological challenges for

A Human-Centric Design Process for Highly Autonomous UAS

mainstream automation methodologies. A possible exception is Aided (or Augmented, but not Automatic) Target Recognition (ATR) where novel approaches may well have increased utility and the subject as a whole remains a challenging research issue. However, notwithstanding the case of ATR, the real challenges appear at Tiers 3 and 4, the boundaries between Tiers 3 and 2 and the decision sharing partnership described earlier.

The methods to increase autonomy can be split into those techniques to increase overall system robustness by endowing the vehicle part of the partnership with greater intelligence and those allowing a true mixed-initiative approach.

Turning first to increased intelligence, any survey of research areas associated with autonomy will show a vast range of techniques, often under the "Artificial-Intelligence" umbrella. It is advanced that the techniques are often limited (in the application discussed here) in that they are applicable where so-called classical established techniques are already adequate [11, 12]. Moreover they are often targeted at a very small part of the overall UAV autonomy problem. For example, artificial neural nets may be applicable in image recognition but may not deliver a high level tactical decision maker operating with partially observable information. For example, it is the difference between the methods to allow the UAV to decide that to successfully attack it is going to be vulnerable to attack from a ground threat and reasoning to resolve the evident conflict, and methods to model vehicle characteristics to understand that damage has resulted in a loss of performance. Both methods may be needed to achieve an overall capability level but both will adopt different technologies to realise the functionality. Consequently, greater intelligence in the vehicle will only be achieved by a range of approaches and thus making architecture a key issue. It is believed that the real risk-reward area for research into UAV intelligence is more likely to be found amongst agent oriented methods which provide a useful framework to express task management and execution of asynchronous events in the face of uncertain information.

Being mainly concerned with task execution, agent based systems are not best placed to carry out planning tasks. To carry out planning the agent based system will delegate the task to a subordinate planner. So for example, the overall task might be for the multi-platform based capability to carry out a search for a mobile target. Only the last sighting would be known and the approximate speed of the vehicle and knowledge of the terrain over which the vehicle is moving may be gleaned from a digital map database. The planner's job is to offer a series of routes for each platform such that the search area is efficiently covered given available resources. Boundary conditions on the problem exist in the form of fuel limitations, performance limitations in respect of platform and sensors. The planned solution might be passed in a series of waypoints via Tier 4 over a datalink to each platform Tier 3 and thence to the Tier 2 waypoint following system for each platform. A planner with such capabilities is a non-trivial requirement particularly when one considers that the mobile target is not likely to be co-operative and use tactics such as hiding under bridges, doubling back etc to confound the UAVs.

Again it is clear that architecture is an issue to successfully integrate the different techniques. In this application, the appropriate agent architecture must be developed to cater for mixed-initiative control of the capability made up of multiple UAVs. The architecture needs to robustly cater for the fact that the decision making team is distributed over a series of datalinks and that the team plan may be interrupted and the team split by the human operator into sub-capability teams deployed on different tasks.

Turning to the other aspect of increasing UAV autonomy level, the human interaction aspect is addressed. Having advanced an argument in respect of the mixed initiative approach to UAV autonomy, the Human Machine Interface (HMI) requires discussion. In this context the HMI refers rather to the philosophy of how the operator manages the level of autonomy of the vehicle rather than the construction of the display media and control station.

A Human-Centric Design Process for Highly Autonomous UAS

Norman in Harris [13] captured what makes an automated [autonomous] system hard to control [trust] in terms of two so-called gulfs:

- The gulf of execution – Does the system allow actions that correspond to the intentions of the operator?
- The gulf of evaluation – Can the operator monitor the state of the system and what is the difference in state from that intended?

where the square brackets are suggested alternatives for the UAV application. Crossing these “gulfs” would seem to be the objective of any mixed-initiative decision making system. Given the Skills-Rules-Knowledge based model discussed earlier, a boundary between the human and machine decision making was implied in the rules and knowledge layers of the human-machine system. Intuitively this interface must change position according to the desired level of autonomy. In simulation work ([14, 15, 16]), operators expressed a desire to be able to select an appropriate autonomy level with respect to mission elements and sub-systems. That is, there will be times where the operator is overwhelmed with new information and would be prepared to accept a higher level of autonomy to remain in overall control – still able to bridge the gulfs!

It should be noted that the two scales of autonomy discussed earlier tend to be meant for use as descriptors rather than a practical implementation of a variable level of autonomy. Some work has been undertaken to define and measure autonomy levels ([17]) as applied to the robotics field where the agents may be less constrained by the need for continual human involvement, for reasons of safety or legality, in the task. However, as has been described, for this application some key decisions may be mandatory for humans. More applicable is the work done on the Pilot’s Associate programme ([18]) that covered some of the earlier Pilot Authority and Control of Tasks (PACT) work that will be discussed next.

This is a pragmatic approach to bridging the gulfs described earlier and has hitherto been applied to increasing pilot capacity within the context of a single seat fast jet [19]. The PACT framework was realised in that application within the UK MoD’s COGPIT programme and showcased the role of adaptive automation and intelligent decision aiding in a future manned military aircraft. The PACT system provides a logical, practical set of levels of automation, ranging from fully manual, assisted, to fully automatic modes, with four levels of automation assistance which, in the combat aircraft implementation, can be changed adaptively or by pilot command. The four assisted levels provide progressive support of pilot situation assessment and decision action. Table 2 below shows the revised PACT system as modified by Taylor and Bonner ([19]).

Primary Levels	Secondary Levels	Operational Relationship	Computer Autonomy	Pilot Authority	Adaptation	Information on performance
AUTOMATIC	5	Automatic	Full	Interrupt	Computer Monitored by pilot	On/off. Failure warnings Performance only if required.
ASSISTED	4	Direct Support	Action unless revoked	Revoking action	Computer backed up by pilot	Feedback on action. Alerts and warnings on failure of action.
	3	In Support	Advice, and if authorised, action	Acceptance of advice and authorising action	Pilot backed up by the computer	Feed-forward advice and feedback on action. Alerts and warnings on failure of authorised action.
	2	Advisory	Advice	Acceptance of advice	Pilot assisted by computer	Feed-forward advice
	1	At Call	Advice only if requested.	Full	Pilot, assisted by computer only when requested.	Feedforward advice, only on request
COMMANDED	0	Under Command	None	Full	Pilot	None, performance is transparent.

Table 2. Pilot Authority and Control of Tasks (PACT)

A Human-Centric Design Process for Highly Autonomous UAS

In the fast jet application the pilot forms a PACT contract with the automation by allocating tasks to PACT modes and levels of automation aiding. This allocation was made either by default settings set prior to flight and/or during mission planning, or by prior authorisation of level changes following agreed onset and triggering conditions or events, or alternatively by making changes to levels by command in flight.

The application of PACT to the control of UAVs is a new development [20]. The PACT system as applied to the UAV application aids the control of the multi-vehicle capability and is related solely to the control of higher level UAV activity. In initial work ([21]) the principal focus was the heavy workload target location and identification tasks. The underlying philosophy of the PACT applied in this work is that:

- skill based tasks should be automated first (PACT levels 4 and 5),
- rule-based tasks should be automated where possible (PACT 3, 4 and 5) and if not able to be automated they should be supported by decision-aiding (PACT 2) and
- knowledge-based tasks should be supported by decision-aiding (PACT 1 and 2).

This philosophy corresponds to the “KRS” paradigm as discussed earlier in this paper. The key point here is that the level of human interaction is variable on a task by task and mission by mission basis. A further point is that the required technologies to achieve PACT levels 2 and a proportion of 1 may well increase significantly the overall software costs of the system. Trials work (*op cit*) has yet to establish a need for this technology. The desired “PACT level” is set as a function of:

- regulation due to safety clearance or legality such as RoE
- technological capability – e.g. do we trust aided target recognition for target identification?
- bridging the “gulfs” to keep the operator engaged at the correct level of arousal and in control
- the command level of abstraction as defined in loose terms in Table 1.

In summary, higher levels of autonomy in UAV systems will be achieved by the appropriate mix of technologies, using an agent based architecture to build intelligence built around a mixed initiative approach that accepts from the outset the partnership between the user and the intelligent machine. A key risk area is the establishment of a suitable architecture to facilitate this approach.

5.0 WHAT ARE THE BARRIERS TO HIGH AUTONOMY IN UAVS?

The previous sections have discussed in some detail what autonomy is and how it might be measured. The argument that the level of autonomy is a function of the degree of human interaction was presented. Furthermore, the robustness of the system in the face of uncertainty together with a workable level of decision sharing is delivered through the increased intelligence of the air vehicle capability and a variable autonomy interface.

To achieve high intelligence in vehicles an agent based software architecture has been argued as the obvious approach and this presents a number of barriers to wide adoption. Agent based methods, where modern Java type development environments are used, represent a significant departure for the generally conservative aerospace domain. The ability to use the agent based approaches will make a huge difference to the life-cycle costs of UCAVs as they are designed specifically for this representation of a problem and the rapid development of appropriate behaviours. Using such means tends to fly in the face of the accepted approaches of using methods such as ADA and its sub-sets. However, whilst it is almost certainly possible to achieve high intelligence using traditional ADA based approaches, it is argued that it becomes prohibitively expensive to do so.

A Human-Centric Design Process for Highly Autonomous UAS

It should be noted however that UAV tactics will change on a mission by mission basis as target sets, weapon & sensor fits change. This may mean that tactical behaviours will also need to be changed to reflect system changes on a daily basis. Moreover, as doctrine changes throughout any campaign then so, it may be inferred, must the UAV behaviours. Therefore, it is a key requirement that both architecture and software can accommodate this without compromising the overall system safety case and without limiting system flexibility throughout the system life. The process must be rapid, safety clearable and cost effective.

Implementation of the UAV behaviours (derived from tactics) using an agent based task execution framework, as has been discussed, represents a key challenge to the overall design philosophy in terms of architecture and software engineering, if system flexibility is to be maintained.

The barriers to use of agent based approaches become therefore:

- Trust – will the community come to accept these novel approaches?
- Clearance – can methods be found to clear these approaches in a safety critical sense to the satisfaction of the authorities?
- Expertise – is there sufficient expertise in these methods within the UAV/aerospace domain?

The situation is exacerbated by the need to use an array of methods to achieve appropriate parts of the overall requirement. The road search example highlighted the need for at least two distinct approaches. Clearly when mapped across the foreseeable mission space the required specialised planners may be key. Furthermore, novel methods to fully realise the Tier 3 and 2 elements of the decision hierarchy may need to be developed.

The likely decision space (the range of all possible decisions that would be shared between operator and UAV capability) for a range of missions has, thus far in this document, been considered largely in the reactive sense in that decisions are made on-line based on current information. One must also consider the a-priori decision making (currently embodied in the mission planning process) and the exact interplay between decisions to be made a-priori and those on-line in real time. The nature of this relationship makes the difference between capability represented by systems such as that in the bottom right of Figure 1 and those of the top left and reflects why it is we need UCAVs at all. Another example might be Inter-planetary space based vehicles. These may appear to be the apotheosis of autonomous technology but one must consider the scale and length of the pre-planning and support during missions to cater for contingency, but also need high levels of intelligence to cater for significant system latency. Whilst UCAVs may arguably enjoy less latency they can not afford the scale of a-priori planning found in space projects.

In order to integrate the various software components of the system a suitable architecture is required. The architecture must consider the philosophy of autonomy described in this document but must also recognise:

- the technological constraints imposed – sensor fit and performance for example
- concepts of use and tactical behaviours – do we need to fly a figure of eight and possibly linger too long within a threat area
- doctrine – what are the decisions deemed critical that may necessitate a human decision maker for reasons of safety or legality
- safety – that the distinction is made between safety critical and mission critical functionality such that inevitable changes to the mission critical aspects can not impact on the safety critical aspects

A Human-Centric Design Process for Highly Autonomous UAS

- flexibility – the ability to grapple with these issues as they change on a regular basis and having accepted that almost by definition the fielded system will not meet requirements on day-one of any conflict

Clearly a suitable architecture, a process by which its requirements are captured and the system developed becomes a significant barrier to the realisation of highly autonomous UAVs.

Finally, the key to targeted autonomy such that the “gulfs” are bridged is the Knowledge Acquisition (KA) process. This is the process whereby the user’s requirements are captured in a form such that they are readily understood by the developer of the architecture and the intelligent software. It is a skill traditionally associated with the developers of Knowledge Based Systems but in the context used here refers to the whole process of the UAV system design. Given that in essence no operational UCAV is currently in service then the KA process can not be applied in the traditional sense. KA is usually applied to a known process where an expert exists and the aim is to capture that human’s expertise or a sub-set of it such that an artificial process can duplicate the functionality. In this UAV application the vehicle concepts, technological capabilities, mission types, role of the operator and UAV tactical behaviours are all the subject of research and the “art of the possible” being developed in parallel. Consequently no “expert” exists and a new process is required that embodies more of the overall system design process than hitherto within intelligent systems design. This too represents a barrier to realising highly autonomous UCAVs.

This paper will now go on to illustrate how some of the limitations outlined above have been addressed using case study information from both flight and SE trials.

6.0 ACHIEVING AUTONOMY

The previous sections described the process and the barriers to developing high levels of autonomy. This section of the paper will describe using case study how the intelligence of the system was increased to allow communication between the operator and the system at high levels of abstraction. It will also describe how the design of the human machine interface is fundamental to managing the operator workload.

6.1 Machine Intelligence Technologies

6.1.1 Rationale

To achieve high levels of autonomy in teams of UAVs it is necessary to endow the package with sufficient powers of decision making to cope with the high levels of uncertainty present in the modern battle space. To do this a number of technologies have been developed to carry out the overall mission execution, drive behaviours and optimise the resource allocation of the UAV based capability. These technologies can equally be applied to autonomy and decision support problems in any wider enterprise where time critical decisions are required in the context of overwhelming levels of information. The higher level task execution framework allows the ready integration of specialised planners which are used to address particular planning issues. The next sections will describe both the task execution framework as well as select specialised planners.

6.1.2 Agent-based task execution

The task execution framework is an agent-based approach to group behaviours based on Joint Intentions theory. It allows an operator to control a pool of unmanned vehicles (such as UAVs) to carry out a number of task-specific behaviours in realistic scenarios. Each agent uses its beliefs about the world (from sensors and messages from other agents) together with its current orders to determine how to act, what

A Human-Centric Design Process for Highly Autonomous UAS

specialised sub-planners are required to generate plans and for how long such plans remain valid. Different types of agents within the system have different roles, such as controlling an individual vehicle platform.

The current implementation can:

- control individual and groups of UAVs to achieve goals set by a human operator;
- organise and re-organise group behaviours to achieve tasks;
- achieve robust responses to individual losses and failures;
- accommodate a variable level of autonomy given to the system
- provide timely plans to decision makers.

The agent-based execution framework provides the architecture for the integration of a number of different planning and reasoning techniques into one system (through specialist planning and scheduling agents).




Figure 4. Agent Task Execution Framework


The system has been demonstrated in flight allowing an operator in a Tornado to control a search and destroy mission carried out by a pool of UAVs.

6.1.3 Specialised Team Planners

The approach to autonomous searching and monitoring is two-layered. Sophisticated recursive Bayesian filtering methods are used for belief estimation and sequential decision theory is then applied to model optimal behaviour. Belief compression techniques are used to link the output of the estimation layer with the input of the decision-making layer. Approximate solutions to the sequential decision problem are obtained using dynamic programming and/or reinforcement learning. This architecture leads to highly efficient search strategies that maximise the chance of detecting a target in scenarios where conventional behaviour-based techniques would fail badly.

A Human-Centric Design Process for Highly Autonomous UAS

A real-time system has been developed that controls a team of UAVs to acquire mobile targets. The targets may be known to be present as a result of a missile launch signature, intelligence report or stand-off detection. The estimation layer is able to track the possible locations of all such targets, taking into account their movement capability on roads and across open terrain. The decision layer uses a dynamic programming (short-term planning) approach to decide routes for each UAV in the team, maximising the chance of acquisition from the UAV's EO imaging sensors.



Searching a road network with a team of UAVs (triangles): black = UAV plans; green = searched area; red = possible target location; blue = known no target.

Figure 5. Search Planner.

6.1.4 Heterogeneous Task Scheduling

Given the unscheduled arrival of tasks for a capability consisting of many disparate assets it is not necessarily obvious which asset(s) should be tasked to carry out the demand. Constraints on individual assets and future tasking requirements need to be taken into account. The heterogeneous task scheduler automatically and dynamically allocates resources to tasks accommodating these concerns.

A real-time module for autonomous allocation of a UAV team to tasks dispersed over a wide area has been developed. The system is driven by a stream of surveillance and monitoring task requests. The hybrid planning/off-line learning system ensures not only efficient short-term planning, but also tactical and strategic awareness. The tactical awareness ensures that the UAV team is always positioned to respond rapidly to new requests. The strategic awareness ensures conservation of resources (fuel or weapons) and minimisation of risk to assets. The component-based architecture allows rapid reconfiguration for UAVs with differing performance envelopes and for differing task characteristics. In a complex scenario, performance of the hybrid planning/learning system was shown to be 85% higher than a simple rule-based system: force multiplication without increased equipment or operating costs. The same planner can be readily adapted to other resource allocation tasks.

6.2 Managing Operator Workload

As was argued in the earlier part of this paper increasing machine intelligence forms only part of the solution. This section will detail key facets of the Human Machine Interface (HMI) that allow the effective partnership of human operator and intelligent system. The main focus for the HMI is the ability to offer sufficient information to the pilot in order for him/her to carry out the mission effectively. Simply presenting information may indeed present the pilot all the information available pertaining to the mission and context of achieving mission success, but a cognitive engineering approach to HMI enables the design of the visual display to be managed in relation to the user's mental workload and situational awareness.

The ability of a user to absorb information from a visual display depends on where and how that information is presented on the display. Previous studies have shown that the random presentation of information on a screen can lead to mental overload and a decrease in performance and usability. For example, Ververs and Wickens [22] presented pilots with different interfaces, and found that the random cluttering of information led to a loss in situation awareness. Their findings suggested that the location (or placement) of information on displays was an important determinant of usability. Indeed, models of visual search (see [23]) have indicated that – in rapid parallel search - information about the location of a visual target pattern amongst distracters and information about its identity was not available simultaneously to the operator, unless the target location was presented at earlier stages of visual processing.

Previous SE trials results implied that when an operator was instructed to command and task multiple UCAVs he or she would become overloaded causing degraded human performance during the mission. However, by introducing a degree of supervisory autonomous control, the pilot was then enabled to command and monitor multiple UCAVs and still maintain good situational and tactical awareness.

By designing the cockpit displays (under cognitive engineering principles) it was possible to map the system design with the pilot's cognitive abilities that therefore brought the interaction between system and human to be more cognitively compatible [24].

A number of Multi-Function Displays (MFDs) were designed in order to allow the pilot to interact with the autonomous system (see Figure 6). The main aspects of the display allowed the operator to view (a) asset and task information; (b) PACT (Pilot Authority and Control of Tasks – The ability to tag decisions with the level of the human authority ceded to the machine whilst cueing the corresponding operator response) messages; and (c) the moving map. Other displays were included to provide more in depth information if the pilot so requested. By providing information across these MFDs the system was able to present the pilot with specific information related to the mission whilst providing feedback to engage the pilot in the interaction with the autonomous system under his/her control. The moving map MFD (c) was viewed as the primary display as it was upon this that track information was relayed to the pilot and provided pertinent tactical awareness as the mission progressed. The other displays were concerned with assigning UCAVs to specific tasks (a) and providing feedback to actioned (or requested) autonomy messages that had been passed to the pilot (b).

A Human-Centric Design Process for Highly Autonomous UAS




Figure 6. Example set-up of Multi-function Displays

In order to establish trust with the system, it was important that the pilot was provided with cues as to what the system was doing (even after being tasked by the pilot). During phases of the mission the pilot could expect to be prompted (as advisory messages) by the system as to when to perform actions. For example, when a target has been identified (and scheduled for attack) the system will advise the pilot as to the window of opportunity for releasing a weapon against the target. The pilot may accept or reject this advice, but is aware of the temporal context of the decision to be made (see [25]).

Due to the nature of controlling multiple UCAVs, it is important that the pilot is assisted in viewing the mission in terms of tasks, as opposed to assets available. By introducing a task-based display (as shown in (a)), it is possible to allow the pilot to view a top-level picture of available assets and what actions those assets have been assigned and are currently executing. This allows the operator to concentrate on the mission task as opposed to being swamped with potentially too much information in terms of planning and controlling all assets available. This visualisation also allows the operator to quickly ascertain how many assets are available for tasking, as shown in Figure 7.




Figure 7. Visual representation of Asset display

The use of visualisation to display the composition of the assets available to the operator provides an overview at any given time as to how the UCAVs are grouped together to perform co-operative tasks. In the current form the display being developed by QinetiQ also enables the operator to interrogate the individual package on the display and receive a breakdown as to the information pertaining to each singleton UCAV.

7.0 AUTONOMY DEVELOPMENT IN PRACTICE

7.1 Use of the SE to increase UAV autonomy

Having discussed the level of autonomy, and made the case for the intimate involvement of the operator in a mixed-initiative approach to UAV control this section will describe the process used to accurately capture the user's requirements using a Synthetic Environment (SE), de-risk and then carry out flight trials.

QinetiQ have been addressing UCAV autonomy issues since 1997 when an SE was assembled to simulate various UCAV based concepts and in particular their interaction with manned assets ([5], [6], [14], [15], [16], [21], [26]). Since then the tools, process and programme aims have all matured to the extent that the current imperative is to develop flight-trial-ready software whilst still addressing the broader de-risking aims of the research programme. Specifically, the process has evolved in a series of SE trials conducted annually since 1998 covering a range of missions (Suppression of Enemy Air Defences, deep strike, Attack of High Value Mobile Targets...) and force mix combinations of UAV and manned platforms. Modern sensors, communications systems and weapons have been modelled. A critical part of the SE has been the implementation of autonomy options, the HMI and the human/vehicle intelligence partnership required to achieve airborne control of unmanned strike systems.

Latterly the increased focus on a narrower range of mission types has highlighted the highly sensitive nature of the knock on effects resulting from apparently small changes in system specification. For example, field of regard of sensors can impact the required turning performance of the air vehicle with concomitant effects on mission timelines.

A Human-Centric Design Process for Highly Autonomous UAS

Figure 8 shows the process by which the autonomy level is increased in the work described. The process has evolved as a result of the trials described above. The process begins (1) with the definition of the likely system capabilities in terms of both performance and composition. For example the airframe may be of limited manoeuvre performance, different sensors may be distributed across different members of the package, sensor performance will be defined and sensors selected, weapon capabilities such as stand-off and targeting requirements will be defined and self-defence capabilities listed.

The operational analysis teams (2) then define concepts of use. This involves developing appropriate timelines, distributing the notional work share across the resources such that the appropriate platforms are in position to carry out their task at the correct time. For example, two platforms will stand-off and use Synthetic Array Radar to search an area whilst the other two will penetrate to carry out any required attack.

At this point (3) the required behaviours must be articulated down to individual component behaviour. For example, sensor coverage is achieved by a swathe of a predefined size and dictated by a racetrack pattern carried out at a certain slant range from the target. This sensor swathe must be established before other assets penetrate the target area and require the sensor information to carry out their part of the mission. Accordingly we can describe the optimum position and duration of any platform racetrack requirements. Further example behaviours could be to position the non-attacking UAVs in support to carry out battle damage assessment and then re-attack if necessary?

Having completed stage 3 of the process the system decision making capabilities must be assessed (4). In this discussion the term decision making is loosely applied and used to encompass a number of capabilities. For example, the flight management system of the vehicle can produce a range of loiter patterns but may not have the intelligence to decide on which is most applicable to any given situation. It relies on an external decision making authority to decide which to use. The sensor may have the sophistication to be able to gather or attempt to gather data on a particular target location and to continue to do so until data of sufficient quality is obtained and could be said to have a decision making capability in that it knows when it can move on to the next target. As a further example reliance and trust of specific Aided Target Recognition algorithms may allow us to delegate image sorting and gathering to decision making within the sensor management components but not allow the decision to attack a track selected by such means.

In the light of knowledge of the system decision making capabilities the level of trust we can place in the system to make the correct decision will become apparent. Accordingly, armed with knowledge of constraints for reasons of legality, Rules Of Engagement and safety clearance, we can mandate which of the decisions are critical and must be made by the human being in the loop (5). For example, association of track information to the target, identification of the target from suitable imagery and weapon release would almost certainly be mandatory human decisions.

Having obtained the critical decisions it is now possible to allocate the decision authority according to a variable autonomy levels, controlling the human versus machine authority, to each of them (6) and decide on the appropriate technology to achieve each level. For example, it may require decision support technology for the lower levels but the higher levels can rely on the task execution components ([20]).

Step 7 and 8 constitute a lower level iteration process as they require the agent based task execution software used here to be developed in full knowledge of, and integrated with, the variable autonomy interface and HMI. Depending on the authority level set then different actions are permissible. A decision where the machine has full autonomy can be carried out entirely by the system. Whereas, a decision requiring operator acceptance requires operator interaction and an implied delay in task execution until the decision is made. This delay requires the system to make the UAVs do something sensible whilst they await the human decision such that they are able to continue to address mission aims. This may throw a

A Human-Centric Design Process for Highly Autonomous UAS

tight timeline into disarray and cause significant on-line re-planning. Moreover, the selected level of autonomy also has an impact on the information displayed to the human operator such that he or she remains fully engaged and cognisant of system state.

Finally, more than one iteration will be needed to achieve true compliance with the user's needs and a key attraction of the use of SE's becomes apparent at this point. The real time human in the loop trial affords many opportunities for demonstration and peer review of the work. Many trials can be run with a broad range of operators. Indeed this work, focussed mostly on the airborne control of UAVs, has been reviewed by a number of UK pilots as well as pilots from the US, France and Germany. Another expedient has been allowing aircrew customers the opportunity of flying in the system. This conveys much more than a series of progress reports and presentations and has been shown to help build a compelling case for the continuation of the work as well as endowing the customer with a true feeling of ownership.




Figure 8. The autonomy development process

7.2 The Surrogate UAS

The purpose of the SE in the early stages was to extract the requirements for a candidate system and to converge on a workable architecture. However, as maturity levels increase the SE also serves to de-risk flight trials. The transition of simulation based technologies into a flight environment provides an opportunity for software previously exposed only to SE functionality tests to encounter real-world problems. In addition the use of a real air vehicle poses its own challenges in terms of its variable turn and speed performance. The ability of an autonomous system to deal with these effects gives an indication of the level of technological maturity of the system. Accordingly the UAS Surrogate was conceived to continue the maturation of the autonomous mission system, its architecture and its constituent components.

The UAS Surrogate comprises a number of elements, the UAV Surrogate itself, a Control Station and a ground vehicle.

7.2.1 UAV Surrogate

The UAV Surrogate aircraft is a passenger jet BAC1-11 that has been converted into a research laboratory aircraft, see Figure 9. A phased approach of installation and concept demonstration trials has taken place over the past four years gradually enhancing the capability.



Figure 9. BAC1-11 UAV Surrogate Aircraft

Although the BAC1-11 behaves as a UAV in that the flight path is determined by the autonomy system which drives the aircraft autopilot, the pilots use a flight-deck display to closely monitor the progress of the aircraft throughout the flight (Figure 11) to maintain safe operation. The flight deck display displays the roll and speed demands and the aircraft response as well as the refined route on a moving map (Figure 10). The presence of the safety pilots allows the BAC1-11 to operate in UK airspace by following normal procedures and processes. The pilots also perform the takeoff and landing functions, and engage the autonomy system once within suitable airspace.



Figure 10. Flight Deck Display



Figure 11. Pilot monitoring Flight Deck Display

The UAV Surrogate system was installed in a series of four phases with flight trials to demonstrate new capability after each phase:-

- Phase 1 involved the installation of the BAC1-11 control station, integration of the autopilot roll channel and installation of the flight deck display for the pilots to monitor. The BAC1-11 was controlled via the roll autopilot interfaces using a generic vehicle control system running on flight capable hardware

A Human-Centric Design Process for Highly Autonomous UAS

directly coupled to the autopilot. This allowed the waypoint following ability, a fundamental part of the system, to be demonstrated early on in the project.

- Phase 2 saw the installation of the rest of the onboard computer hardware. This included a Control Station, Synthetic Environment (SE) station and additional synthetic UAV hardware as well as the expansion of the autopilot functionality to include pitch and throttle. Multiple UAVs, simulated ground entities and autonomous technology could now be generated and a more complex scenario demonstrated in a flight environment.
- Phase 3 was concerned with the installation of the external communications links. The Control Station interface, in this phase a Tornado (TIARA – Tornado (F2) Integrated Avionics Research Aircraft) fast-jet, used a short-range TALON radio link. Surrogate UAV communications with the ground for SE ground truth purposes, was achieved using a satellite communications link. TIARA is a twin-seat research aircraft that can act as a single-seat aircraft, in this particular trial the second pilot's role was to act as a safety pilot.
- Phase 4 involved the installation of a Wescam MX-15 Electro-Optical (EO) and Infra-Red (IR) turret. A master control unit, hand controller and system interface were also installed to allow manual control during the initial integration and testing of the Wescam sensor.




Figure 12. Layout of Equipment onboard the BAC1-11 aircraft

A Human-Centric Design Process for Highly Autonomous UAS

The layout of the equipment onboard the BAC1-11 is shown in Figure 12. Rack 1 contains the hardware controlling the BAC1-11 aircraft and interfaces directly to the autopilot. The hardware generating the synthetic UAVs and their associated capabilities such as sensors and weapons is located in Racks 3 and 4. Other onboard hardware runs the Control Station (Rack 2), communications links (Racks 2 and 5), monitoring tools for the trials team (Rack 3) and generates simulated ground target entities (Rack 3). Rack 3 is also the location for the manual control of the real Wescam sensor using a hand controller.

7.2.2 Control Station

The BAC1-11 and the onboard synthetic UAVs were controlled using a form of operator control station. Flight trials in March 2007 culminated in the control station being transferred from Rack 2 on the BAC1-11 to the TIARA. This allowed the programme to determine the feasibility of a single-seat aircraft pilot controlling multiple UAVs by using appropriate levels of autonomy, in addition to their own platform.



Figure 13. Tornado Integrated Avionics Research Aircraft (TIARA)

7.2.3 Ground Vehicle

During the Phase 4 trials in March 2008, the aim also included the integration of a real sensor to image ground targets. A ground vehicle was equipped with satellite communications capability and injected the ground truth vehicle position into the on-board SE. This entity was then replicated onboard the BAC1-11 and injected into the SE to allow the vehicle to be located using both the onboard synthetic sensors and the real Wescam sensor as if cued from longer range RF sensors.



Figure 14. Ground Target Vehicle used during March 2008 Trials

7.2.4 Wescam Installation

The most recent addition to the UAV Surrogate aircraft was a Wescam MX-15 turret (Figures 15 and 16). This was the first time a real capability in terms of the UAVs' sensors and weapons had been used.




Figure 15. Representation of installed Wescam MX-15 turret



Figure 16. Wescam MX-15

7.2.5 Installation Process

The software to be trialled in flight was developed by iteration through the SE process outlined previously. The product of this process was then issued for integration into the Surrogate UAS. In order to integrate and test the software to be installed on the BAC1-11 aircraft a simulation rig with duplicate hardware was created. This allowed the software to be developed elsewhere and then integrated and tested on identical hardware. In addition, prior to trials involving the TIARA aircraft the Surrogate rig integrated with a rig version of the equipment on the Tornado was utilised.

A Human-Centric Design Process for Highly Autonomous UAS

7.3 Flight trials

7.3.1 Tornado Trials – March 2007

The BAC1-11 acted as one of multiple UAVs within the SE generated onboard the aircraft. The UAVs were commanded to perform tasks by the pilot flying the TIARA Tornado. These tasks take the form of goals and commands transmitted from the Control Station to the BAC1-11 and acted upon by the UAVs including the BAC1-11. Several synthetic ground entities were generated, these were to be located by the UAVs' sensors and identified from sensor images passed to the Control Station operator's displays during the course of the mission.

Initially the UAVs were commanded to Ingress into the area of interest. A Search task was then commanded which used a short range sensor to locate synthetic ground entities within the pilot designated search area having been cued by long range RF sensors. The Control Station operator then classified from imagery the entities located with the sensors and ordered a Destroy task. The Destroy task performed a simulated weapon attack on the designated target. At certain decision gates, such as weapon release and track classification, the pilot was prompted for confirmation. This allowed the pilot to retain high level control of the scenario without having to manually fly or plan routes for the UAVs.

7.3.2 Wescam Trials – March 2008

As in previous trials the BAC1-11 acted as one of multiple UAVs, the other UAVs being generated onboard the aircraft. The Control Station in these latter trials was a specially designed HMI running on a work station onboard the BAC1-11 and more suited for a ground control station or wide bodied aircraft based application. Previously, the control station had been confined to integration with the existing TIARA cockpit displays. The addition of the Wescam sensor shifted the focus from external communications links to the control of the real sensor by the autonomous system. A synthetic version of the Wescam was created to be used by the synthetic UAVs to allow all the UAVs to have the same sensor capabilities. The same tasks of Search and Destroy were used with the real and synthetic Wescam sensors in both phases. All targets, both real and synthetic, were present in the simulation environment, and consequently, to the autonomy system, there was no difference between respective sensor product and all were treated exactly the same way. The BAC1-11 UAV needed both a real Wescam to image real entities and a synthetic Wescam to image synthetic entities, see Figure 17. To allow this a router was required between the autonomy system component controlling the sensor and the two sensors.




Figure 17. BAC1-11 WESCAM sensor command route

7.4 Results

In rapid prototyping terms the essential nature of SE based work up was re-affirmed. The interfaces between the “real” software and both the SE based models and flight hardware being identical meant that trial work-up was carried out almost exclusively in the SE. This meant that flights themselves are almost all productive with very little time spent “de-bugging” whilst airborne. The presence of a simulation rig replicating the capability and interfaces present onboard the BAC1-11 is invaluable for performing thorough testing before transitioning technology into a flight environment. It also allows rapid deployment of software modifications during the integration phase and greatly reduces the costs associated with installing and testing software on aircraft.

Another key result is the impact of flight demonstration on perceptions. The difference as has been explained, between flight software and that used in SE based experimentation was minimal in this instance, yet the impact of the flight trials was significantly greater.

In terms of operator workload and whether a fast jet operator can be expected to fulfil the role explored here, results are best summarised with the following quotation from Flight International 10-16 April 2007

[Using it is] “... no more difficult a demand than operating a sensor like a targeting pod. We have proved the system from the most difficult type”– TIARA Test Pilot

Following the trial the ease with which the pilot used the displays in the TIARA cockpit was noted. The level of information conveyed sufficient information to the pilot that enabled him to supervise the UCAS mission effectively. Sufficient trust was engendered that at no time did the operator need to question the engineers on board the BAC1-11 as to what the autonomous system was doing.

Due to the smaller real estate afforded to a cockpit environment there is only so much information that can be conveyed to a pilot controlling multiple UCAVs. QinetiQ have also designed a Ground/Wide-body Control Station, to explore the opposite extreme where screen size is less restrictive. It was this Control Station that took part in the March 2008 and 2009 trials and addressed questions regarding the provision of further information to the operator in order to convey intent of the UCAV behaviour. This will help to build trust between the human and machine and the results of experimentation will form the basis of a future paper.

8.0 SUMMARY

8.1 Conclusions

This document has described a view on what is meant by an autonomous UAV. It has also defined and discussed the degree of autonomy, whilst proposing a measurement matrix in order to bring some clarity to the imprecise terms such as highly autonomous and semi-autonomous. The paper then discussed a hierarchical autonomous system architecture to aid discussion before going on to recommend how high autonomy may be achieved and finally what might be the barriers to achieving high autonomy.

In conclusion, from the literature, it is clear that there is a need for highly autonomous UAVs. High autonomy is a complex function of the:

- overall system architecture,
- available technologies and their capability,
- nature and level of the decision sharing between the human and intelligent machine,

A Human-Centric Design Process for Highly Autonomous UAS

- desired capability.

Autonomy level is a function of the system context, the level of human interaction, the system reasoning capability and its knowledge.

Three key barriers to successful development of highly autonomous UAVs are

1. the need to adopt agent based software methods
2. the process by which the software architecture is defined to meet foreseen and unforeseen needs
3. the knowledge acquisition process in the absence of “experts”

It is clear that research must continue to:

1. Find ways to allow the adoption of agent based software in safety critical applications.
2. Continue to refine knowledge acquisition processes to allow seamless requirements capture and system realisation in a rapid manner.

In developing a highly autonomous UAS system a multi-disciplinary team is needed if the solution is to be acceptable. The end user must be intimately involved with development and requirements capture. Technologists from a range of backgrounds are needed along with human factors experts. No knowledge should be assumed. The SE in this respect allows the “art of the possible” to be explored in an incremental and consensual manner across the stakeholder community. As a general rule, when compared to flight demonstration, SE testing tends to be less directed and allows relatively unconstrained experimentation being less constrained by real implementations.

By contrast, flight trials form a necessary part of maturing the system and certain technologies. They are characterised by longer lead times, greater cost and are more heavily constrained. Accordingly, they tend to be very directed in their aims and used only where necessary to introduce real world effects in specific areas so informing system design.

8.2 Future Work

As is implicit in the demonstrations described here, elements of the current approach to the autonomous system design have shown the concept of low operator to platform ratios. Parts of the system sit at a number of maturity levels with some elements having been taken to flight whilst less mature components are still in development. This range of maturity levels is likely to be constant as system design continually tracks emerging requirements.

This mode of development is planned to continue by steadily increasing fidelity in certain parts of the system by introducing real sub-systems such as multiple vehicle fusion and Aided Target Recognition and other image processing. Integration of these elements and broad test in the SE will be followed by targeted demonstration in further flights throughout 2008 and 2009.

9.0 REFERENCES

- [1] Platts, J. T., Autonomy in UAVs, Royal Aeronautical Society Journal, Feb 2006
- [2] Platts, J. T., Kemsley, J. A., Richards, D., Fast Jet Control of Multiple UAVs, Proceedings of the AUVSI North America Conference 2008, San Diego, CA, Jun 2008
- [3] OED, 2004, Oxford English Dictionary, Oxford University Press.

A Human-Centric Design Process for Highly Autonomous UAS

- [4] Clough, B May 2002, Metrics, Schmetrics! How Do You Track a UAV's Autonomy?, 1st AIAA UAV Systems, Technologies, and Operations Conference and Workshop, Portsmouth VA, USA
- [5] White, A., 2002 (a), The Human-machine Partnership in UCAV Operations, 17th Unmanned Air Vehicle Systems Conference, Bristol, UK.
- [6] White, A., 2002 (b), The role of the operator in UCAV operations, Unmanned Systems Conference 2002, Disney's Coronado Springs Resort, FL USA.
- [7] Lin, Ching-Fang, 1994, Advanced Control Systems Design, Prentice Hall Series in Advanced Navigation, Guidance, and Control, and their Applications, Prentice-Hall Inc, USA.
- [8] Rasmussen, J., 1983, Skills, Rules, and Knowledge; Signals, Signs, and Symbols, and Other Distinctions in Human Performance Models, IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-13, No. 3, May/June 1983.
- [9] Rasmussen, J., 1986, Information Processing and Human-Machine Interaction, An Approach to Cognitive Engineering, Vol 12, Elsevier Science Publishing Company.
- [10] Strens, M., J., 2002, e-mail from Strens M (mjstrens@qinetiq.com) to Platts J T (jtplatts@qinetiq.com), dated 30 May 02, Subject: Levels Comment.
- [11] Platts, J. T., Forsythe W., July 2001 (a), The Application of Self-Organising Fuzzy Control to Station-Keeping in Unmanned Air Vehicles, The Aeronautical Journal, Journal of the RAeS, UK,
- [12] Platts, J. T., 2001 (b), A Self-Organising Fuzzy Logic Supervisor, PhD Thesis, Loughborough University.
- [13] Norman, D. (1991), in Harris, D., The Human Factors of Fully Automatic Flight, Measurement & Control, the journal of the Institute of Measurement and Control, Vol 36/6, July 2003.
- [14] Smith, P. R. S., Mayo, E., O'Hara, J., Griffith, D., 1999, Combat UAV Real-Time SEAD Mission Simulation, AIAA Flight Mechanics Conference 1999, AIAA-99-4185, Baltimore, USA.
- [15] Howitt, S. L., Mayo, E., Platts, J. T., 2001, Simulating the Attack of High Value Mobile Targets Using Combat UAVs, Bristol RPV Conference, Bristol, UK.
- [16] Howitt, S. L., Platts, J. T., 2002, Real-Time Deep Strike Mission Simulation Using Air-Launched UAVs, AIAA Conference on UAVs, Portsmouth VA, USA.
- [17] Barber, K. S., Martin, C. E., 1999, Agent Autonomy: Specification, Measurement, and Dynamic Adjustment. In Proceedings of the Autonomy Control Software at Autonomous Agents 1999 (Agents '99), 8-15 May, 1999, Seattle, WA
- [18] AGARD, 1995, Knowledge Based Guidance and Control Functions, AGARD Report No 325.
- [19] Taylor, R.M., Abdi, S., Dru-Drury, R., Bonner M.C., 2001, Cognitive cockpit systems: information requirements analysis for pilot control of cockpit automation, Ch. 10, p81-88 in 'Engineering psychology and cognitive ergonomics Vol. 5, Aerospace and transportation systems, ed. Harris, D., published by Ashgate, Aldershot, 2001.

A Human-Centric Design Process for Highly Autonomous UAS

- [20] Platts, J. T.; Application of a Variable Autonomy Framework to the Control of Multiple Air Launched UAVs; Proceedings of 16th Association of Unmanned Vehicle Systems International Conference, Orlando, FL; July 2002.
- [21] Bonner, M., Diethe, T., Mathews, S., Aug 2001, Scoping Study for the insertion of Cognitive Cockpit adaptive automation taxonomy and control system into RTAVS for control of UCAV autonomy, QinetiQ internal study report.
- [22] Ververs May, P., Wickens, C.D. (1998). Head-up displays: Effects of clutter, display intensity, and display location on pilot performance. *The International Journal of Aviation Psychology*, 8(4) 377-403.
- [23] Saarinen, J. (1994) Visual search for global and local stimulus features . *Perception*, 23, 237-243.
- [24] Richards, D., Howitt, S. (2006) Cognitive Engineering and HMI design of a UAV Ground Control Station. Third Annual UAV and Human Factors Conference. Arizona:USA.
- [25] White, A., Richards, D. (2006) Building Operator trust in semi-autonomous UCAV systems, Moving Autonomy Forward Conference 2006. Lincoln: United Kingdom.
- [26] Platts, J. T, 2004, The use of Synthetic Environments for Autonomous UAV Development, AUVSI, Anaheim, CA, Aug 2004.