Policy Management and Negotiation: Enabling Effective Human Autonomy Teaming

Michael Thomas
QinetiQ
Farnborough, GU14 0LX
UNITED KINGDOM

mthomas1@qinetiq.com

Rob Cottrell
QinetiQ
Farnborough, GU14 0LX
UNITED KINGDOM

rjcottrell@qinetiq.com

Antony Grabham
Dstl
Portsmouth, PO17 6AD
UNITED KINGDOM

apgrabham@mail.dstl.gov.uk

ABSTRACT

Autonomy technology is being developed and deployed in support of military unmanned systems. However issues persist, particularly in relation to future aspirations for unmanned aircraft operations, on how to balance requirements for flight certification and attendant needs for transparency and the trust of machine autonomy by human operators and team members, with the need for increasingly complex autonomy to achieve mission objectives. With autonomy being essential to unmanned aircraft operations during certain mission phases or under specific circumstances, human autonomy teaming cannot be effectively achieved by merely managing levels of autonomy. Effective human autonomy teaming needs to include paradigms for ensuring that autonomy technologies are employed that are appropriate to the prevailing mission and operational context and which satisfy the applicable policies. Furthermore, human autonomy teaming needs to include provision for autonomy technology to electronically negotiate with controlling authorities where such negotiations are expected such that routine adaptations can be accommodated. In such cases human intervention is only required (by policy) when human authority or decision complexity demands it. This paper discusses research into a Configurable Operating Model for Policy Automation and Control of Tasks (COMPACT) which seeks to employ higher-integrity policy management technology to enable lower integrity with more complex autonomy technology invoked according to the prevailing context and with provision for electronic negotiation between operating tiers of command. The paper also summarises Live Virtual Constructive trials conducted in 2017 to test elements of the research with military participants in militarily representative scenarios.

1.0 INTRODUCTION

This paper provides an overview of a conceptual adaptable autonomy architecture that is currently the subject of a UK MOD research programme undertaken by QinetiQ on behalf of Dstl. The paper also describes a recent trial during which the architecture was tested using military personnel against the backdrop of a militarily representative scenario.

The Configurable Operating Model for Policy Automation and Control of Tasks (COMPACT) was originally conceived [1] as a means of addressing the problems associated with the employment of advanced and novel computing techniques that can deliver the complex, and often ‘intelligent’, software that underpins automated and autonomous systems.

The role of automation and autonomy software in aircraft, particularly unmanned aircraft, has significantly increased and has led to more complex sets of requirements to provide the necessary functionality to enable human operators and pilots’ greater scope and capacity for undertaking missions.

In many cases Tthese complex requirements can only be delivered through provision of novel and non-traditional software development concepts, such as Artificial Intelligence (AI) and Machine Learning (ML). This can lead to software components that are difficult to certify and therefore exploit on operational platforms [2]. In addition, the increased software complexity can sometimes be difficult to interpret by the operators and pilots who use them, leading to issues of trust and comprehension.

For a number of decades, this potential division of roles and responsibilities between the human and the system has been subject to much research. Early efforts focussed on the role of a system providing pilots of manned systems some level of assisted, associate or coach support [3]. Examples include the US PA [4]and RPA [5], and the UK MMA [6] & [7], FOAEW TDSS [8] and Cognitive Cockpit [9].

In many of these cases, the final responsibility for decision making was retained by the human operator as they were ultimately in control of their platform. But today, with operators of unmanned systems potentially expected to control multiple platforms, and multiple platform types, the human component has much less cognitive capacity to effectively manage their responsibilities without the application of novel software able to provide the required levels of decision support, automation and autonomy. Furthermore, it is expected that there will be mission phases or specific circumstances when autonomy is routinely required as an essential component of future unmanned aircraft operations, rather than being required as a crew aid as was often the case for the research focussed on manned platforms. The COMPACT concept provides a potential means to bridge many of these issues and provide a robust mechanism for ensuring complex software exploitation and trusted decision making.

2.0 THE COMPACT CONCEPT

The COMPACT concept is intended to provide a means for monitoring and controlling Unmanned Air Systems (UAS) that necessarily employ various autonomy and automation technologies, based on a set of pre-determined (and configurable) rules.

When characterising the COMPACT concept, it is worth looking back to the human machine teaming concepts reported in [3]; the assistant, associate and coach. Within these definitions, the COMPACT concept has a role similar to that of a coach, in that its purpose is to monitor both the performance of the system and the human (in terms of how the human interacts with the system), and determines whether this performance is consistent and compliant with applicable rules and constraints.

In many respects, the COMPACT concept architecture (see Figure 1) shares many of the capabilities of the Situational Assessment Support System (SASS) and certain elements of the Task Interface Manager (TIM) developed under the Cognitive Cockpit Programme [9]. The SASS comprised a rule based system which, by understanding the pilot's cognitive state, the platform status in relation to the mission plan (and associated task model), and the environment, could determine the appropriate course of action in response to dynamic events and most appropriate means of conveying information to the pilot for the current level of autonomy (using the Pilot Authority and Control of Tasks taxonomy – PACT). In relation to COMPACT, the UAS operator's state could be an analogue of pilot's cognitive state, however the initial research focus has been on the assessment of the UAS in relation to its mission goals, mission plan and the prevailing context.

In order to ensure that COMPACT can achieve the desired aims, it requires information and knowledge that relates to the platform/system state, mission and operational context, mission plans and goals, and the task models that describe tactics, techniques and procedures (TTPs). Using this information COMPACT is will be able to monitor the platform state against the platform model, task model, mission and task plans and identify discrepancies. COMPACT will have knowledge of courses of action when discrepancies in the platform state against the platform or task model, or the sortie, mission and task plan occur. COMPACT will expedite defined actions to ensure consistency is maintained in the platform state against the platform and task models and the sortie, mission and task plans unless this conflicts with the control state.

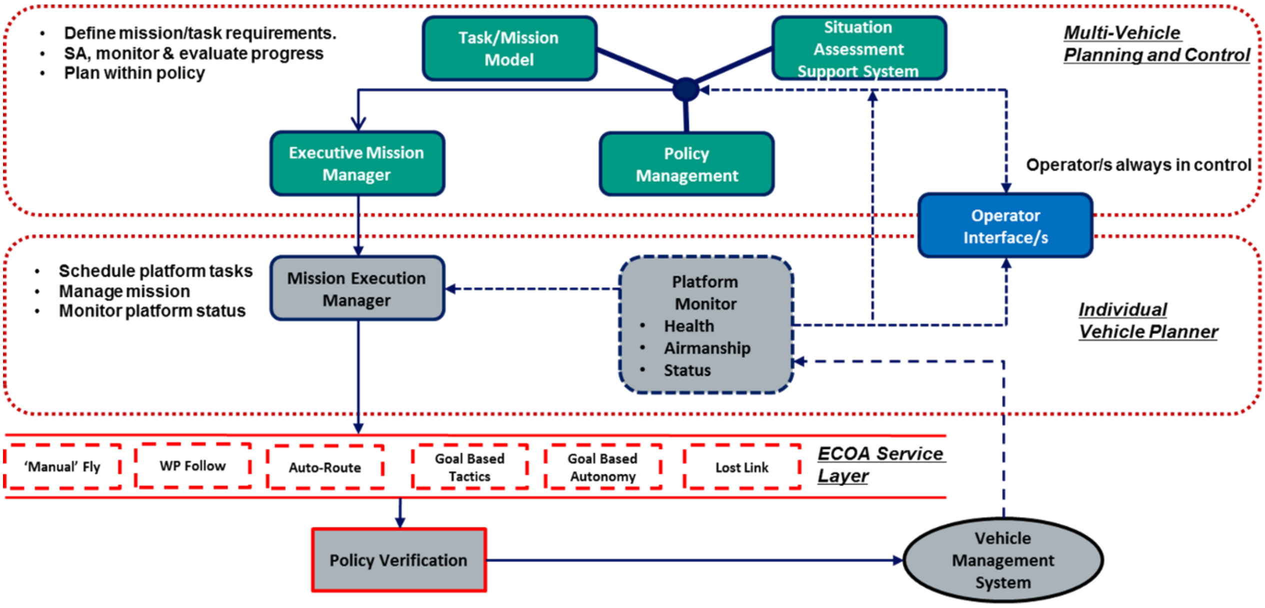

Figure 1 COMPACT conceptual architecture

COMPACT is predicated on the idea that there are a collection of software agents that can invoke various types of software components (including novel and AI techniques) able to provide platform or system decision support, automation and autonomy functions (which may be implemented using the European Component Orientated Architecture – ECOA, see [10]). The invocation of these agents is controlled by a Task or Mission Management function in response to evolving mission requirement and situation assessment.

The Task Management component continuously monitors which agents can be used through Situation Assessment (SASS) and policy management. The SASS Module will continuously monitor current status. To include:

- Environmental: terrain, landscape, weather
- Platform: position, altitude, speed, direction, fuel,
- Mission Phase– current and future (as advised by Task Manager)
- Policy Manager - Airspace / Mission Compliance; rules, constraints and Rules Of Engagement (ROE) as defined by Policy data

The Task Manager Module has knowledge of how the mission should be carried out (based on a priori knowledge) and will 'decide' which behaviours are to be conducted in order to achieve the task. These behaviours may include those delivered by the agents, the operator, or combinations thereof.

2.1 Policy Management and Negotiation

Policy Management is the means by which the performance and configuration of a system is monitored by COMPACT to ensure its operation is compliant with current policies. The term policy refers to the constraints and guidelines within which military operations are to be conducted. At the highest level, different nations apply laws and regulations that must be adhered to, such as (in the UK) the Law of Armed Conflict. At the operational level, the ROE provide a set of constraints that determine how military force should be applied. Tactical operations are managed and directed using a multitude of policies and in the case of aircraft; airspace management and the Airspace Control Order (ACO) are of particular relevance as is the

Air Tasking Order (ATO). At the platform level, Tactics, Techniques and Procedures (TTP) shape behaviours that maximise platform performance and the mission plan to ensure the desired military effect (see Figure 2).

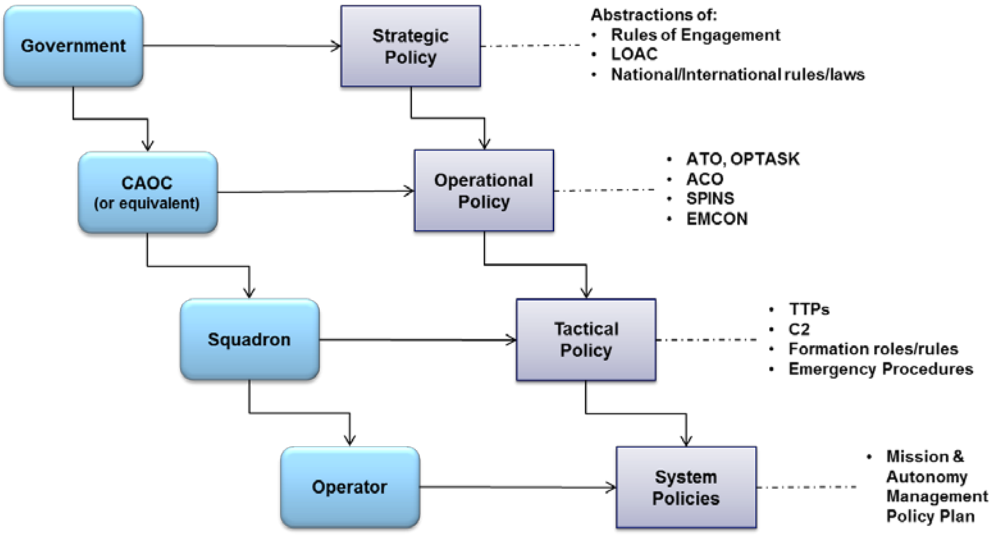

Figure 2 Policy hierarchy

As is the case in manned military aviation, there are occasions when, as a mission progresses, policies come into conflict. A typical example includes the re-planning or re-tasking of a UAS asset that requires it to enter airspace it is not permitted to enter. For manned aviation the pilot will request access to the required airspace from the relevant authority. This interaction between the pilot and the authority represents a negotiation, which is a natural consequence of policy management.

An initial model for the application of negotiations within the policy management framework is depicted in Figure 3. The detection, assessment and classification of a policy conflict is made and the determination as to whether the adjustment will require a change in Goal, Means or specific Acts will determine the process to resolve the conflict.

- A Goal adjustment is a change to the mission/task itself, and may require a negotiated resolution and approval from a higher C2 authority
- An adjustment of a Means, may require either;
  - A resolution requiring approval (from a GCS operator or Mission Commander)
  - A resolution requiring authorisation from a higher C2 authority
- An adjustment of an Act may require either;
  - An internal resolution (i.e. a minor change to the route plan)
  - A resolution requiring negotiation and approval from another agent (i.e. from a GCS operator or Mission Commander)

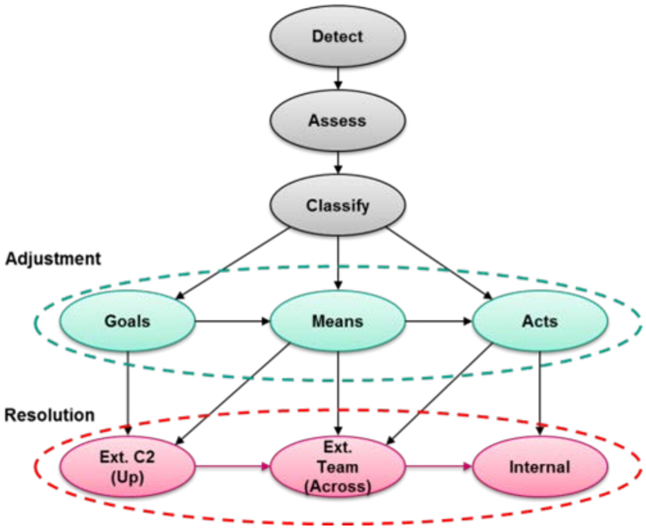

Figure 3 Policy conflict identification and resolution.

To aid the development and implementation of negotiations, as well as their validation and verification, formal methods have been applied in the form of formal agreement patterns [11]. Agreement patterns define patterns for sharing and reusing previous experience in the application of agreements (see [11] & [12]) and provide an abstract way of defining negotiations. Within COMPACT, each policy resolution can be considered to be a type of "negotiation". For instance 'platform 1' talks to 'platform 2' and exchanges some kind of information transaction forming a negotiation.

There are three parts to an agreement pattern; informal representation using classifications to direct the description, reasoning diagrams to show the elements of the agreement, and formal representations identifying the preconditions and post conditions for entering and concluding the agreement. For example, the informal representation of an Airspace Negotiation is defined below:

NAME: Auto Airspace Change
DURATION: Short Term
UAV Operator/autonomy system needs access to airspace for which there is no current authorisation (for example UAV is tracking a target that is heading outside of its currently allocated airspace and needs to relocate to maintain track)
FORCES: Pattern applicable when task requiring access to new airspace is within mission bounds (i.e. a new mission plan is not required), UAV operator/autonomy can handle any restrictions authoriser may give (such as area/height/time bounds), there is an identified authoriser (e.g. Air Battlespace Manager) and there are no other known circumstances which negate the validity of the request (e.g. emergency recovery plan invalidated due to insufficient fuel)
SOLUTION: Approver authorises the request, if there are restrictions with the approval then the UAV operator/autonomy must confirm compliance with these restrictions before entering the airspace.
EXAMPLES: COMPACT
RESULTING CONTEXT: UAV is authorised to enter the airspace to conduct task
PHASE: Agreement negotiation, Agreement enactment, Agreement renegotiation, Agreement conclusion
RELATED PATTERNS: ---

In this example,

- **Name** is a meaningful name that provides a vocabulary for identifying and discussing the negotiation agreement pattern.
- **Duration** is the time the negotiation is covering.
- **Forces** are a description of the relevant forces and constraints and how they interact with one another and with the goals and considerations to be taken into account to select a solution for a problem.
- **Solution** describes static relationships and dynamic rules describing how to realise the desired outcome.
- **Examples** are one or more sample applications of the pattern which illustrate its application and known occurrences of the pattern which help in verifying that the pattern is a proven solution to a recurring problem.
- **Resulting context** is the state or configuration of the system after the pattern has been applied.
- **Related patterns** are any other compatible patterns which can be combined with the described pattern.
- There can be any of the following **phases** in the agreement pattern:
  - **Provider selection**: Selection of the possible provider before making an agreement
  - **Agreement negotiation**: Process of establishing the conditions of the agreement
  - **Agreement enactment**: Commitments of the agreement being effective
  - **Agreement renegotiation**: Phase of changing the conditions of a previously negotiated agreement
  - **Agreement monitoring**: Process of reviewing the established agreement conditions during the agreement enactment
  - **Agreement conclusion**: Process of concluding the agreement by any party

The reasoning diagram for the Airspace negotiation uses notations defined within the CommonKADS methodology [13], specifically identifying *Knowledge Roles* and *Inferences* (Figure 4). The reasoning flow commences when the air platform needs to enter airspace which it has not been assigned (this is a knowledge role). The system then needs to assess the mission bounds, e.g. by assessing whether failing to get permission to enter this airspace will have an impact on the fuel or emergency recovery/diversion plans. Since the mission bounds need more information this is an inference role. The approver then needs to respond to the platforms request. In the example the approver is usually the airspace controller. Their response can either; allow the platform to enter the airspace, not allow the platform to enter the airspace, or allow the platform to enter the airspace with restrictions. In the example, restrictions could include duration, height, or a subset of the airspace requested. The platform then asses the approvers response and decides if it can handle the restrictions and enter the airspace. Since there is some calculation in the assessment part of the reasoning diagram this becomes and inference role. After assessing the response the platform then decides if it will enter the airspace or not.

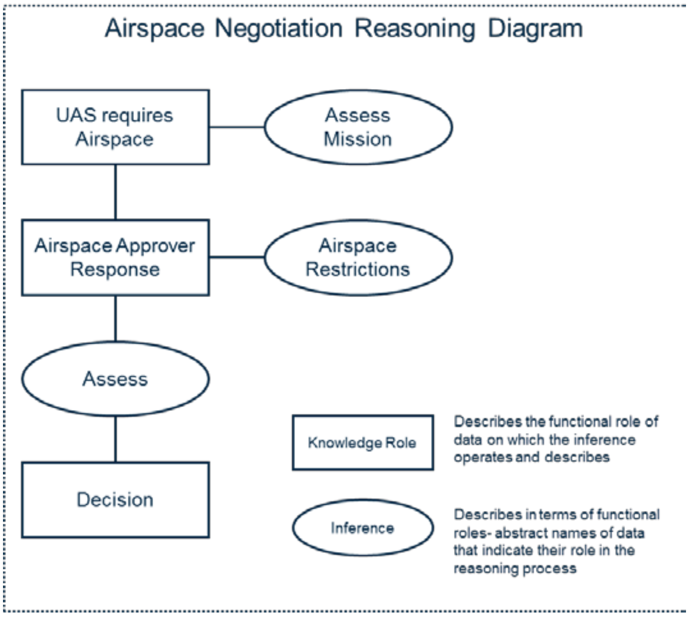

Figure 4 Airspace Negotiation reasoning diagram

The third element, formally defining the pre and post conditions, requires initial definitions of the relevant variables and their types:

variable: TYPE	**Description**
ap : ACTOR	ap is the approver
approved: (ACTOR1* ACTOR2)	Approved is a set of pairs which takes 2 actors meaning ACTOR1 is approved by ACTOR2
m: ACTOR	m is the mission bounds
flying: (ACTOR1* AIRSPACE1)	flying is a set of pairs which takes 2 actors meaning ACTOR1 is flying in AIRSPACE1
u: ACTOR	u is the unmanned air system
a_1 : AIRSPACE	a_1 is the airspace the uav is allowed to fly in
a_2 : AIRSPACE	a_2 is the airspace the uav is trying to get into
res: ACTOR -> {PARAMETERS}	res is a function which takes an actor and outputs a set of parameters (or restrictions)

The first precondition that is represented formally is that before the platform requests entry to new airspace, it must be flying in some kind of airspace already. Therefore the following applies:

$$\underline{\text{flying}(u, a_1) \land (\neg\ \text{flying}(u, a_2))}$$

This equation represents that the platform, in this case UAS (*u*) is *flying* in airspace1 (a_1) **and** the UAS (*u*) is **not** *flying* in airspace2 (a_2). Another precondition that should be identified before entering the negotiation is that the UAS has some restriction it will be able to manage if the approver gives any. This is defined as:

$$\underline{\text{res } u \neq \{\ \}}$$

In this equation the variable *res* takes an actor (the UAS) as an argument. Therefore the first part of the equation *res u* represents the set of restrictions that the UAS (*u*) can manage. The second part "≠ { }", represents that the set is not empty. Therefore the UAS has to have at least some restrictions which it is able to manage. At this abstract level of representation it is not important as to what the restrictions are, just that they exist.

The post condition of the negotiation defines the outcome of the agreement. This is represented as:

$$\underline{\text{flying}(u, a_2) \wedge (\neg\, \text{flying}(u, a_1))}$$

This equation shows that after the agreement, the UAV (u) is flying in airspace2 (a_2) and the UAV (u) is not flying in airspace1 (a_1). This is similar to the first formal statement produced where the UAV was flying in airspace1.

The formal notation taken from the formalising of the Auto Airspace change can be used to define a formal specification. The pre- and post-conditions in the pattern agreement denote the preconditions before entering the pattern and the post conditions after the agreement. Whereas the preconditions and post conditions in the formal specification denote the states before and after the actual process has been conducted. The preconditions used in the formal specification take the preconditions from the agreement pattern and add to them any other preconditions before conducting the action. Various notations can be used within the formal specifications for negotiations. We define the notations as follows:

NameOfSchema
Declarations
Preconditions Postconditions

The **NameOfSchema** gives a name to the process. The declarations include the variables which may be used in the process. This includes options, function names etc. each variable is assigned to a type which it belongs to. The format for a declaration is "*v*:*T*", where *v* is the variable and *T* is the type. There can be none, one or many declarations in each schema.

The preconditions define what is needed before the desired outcome is to be executed. The post-conditions define what state the variables are in after the desired outcome. Additionally, the Ξ is used to define the same variables from a previous schema negotiation.

AutoAirspaceChange
ap : ACTOR approved: (ACTOR1* ACTOR2) m: ACTOR flying: (ACTOR1* AIRSPACE1) u: ACTOR a_1, a_2 : AIRSPACE res: ACTOR -> {PARAMETERS}

$\text{flying}(u, a_1) \land (\neg\, \text{flying}(u, a_2))$
approved(u, m)
approved(u, ap)
$\text{res}\ u \neq \{\ \}$
$\text{res}\ ap \subseteq \text{res}\ u$
$\text{flying}(u, a_2) \land (\neg\, \text{flying}(u, a_1))$

Three new preconditions have been added to the specification which denotes the conditions which need to be satisfied in order for the aircraft to actually change its airspace.

(1) approved(u, m)

(2) approved(u, ap)

(3) $\text{res}\ ap \subseteq \text{res}\ u$

Precondition (1) states that the UAV (*u*) has to been approved by the mission bounds (*m*) for it to enter the new air space. In the formal notation the mission bounds (*m*) are defined as an actor. The mission bounds relate to making sure the air space change does not conflict with the mission. If there are restrictions on the mission bounds then the higher chain of command may recommend that the aircraft should change airspace as that is a priority. Therefore if the higher chain of command makes changing of airspace a priority then it has been approved. At this abstract level there is no need to identify what the mission bounds may be only that the UAS does not change the mission and therefore does not need permission, or that it does change the mission and has been approved by the appropriate party.

Precondition (2) states that the UAV (*u*) must have been approved by the approver (*ap*). The approver would usually be the airspace controller in this example but could be another authority in other examples such as a Mission Commander.

Precondition (3) states that the restriction (*res*) the approver (*ap*) is a subset of the restrictions (res) the UAV (*u*) can manage. Thus the restrictions the UAV can manage can be in a set as follows:

{Can fly between 7000-7400 feet, can fly for 4 hours, can loiter for 20 minutes before entering, … }

The approver may allow the UAV to enter the airspace but only for 4 hours. Therefore the restrictions given by the approver can defined in the following set:

{can fly for 4 hours}

Thus the restrictions given by the approver is a subset of the restrictions the UAV can handle. At this abstract view specific restrictions are unimportant but that the UAV may have some restrictions it can manage which the approver may or may not state.

Since there are 5 preconditions which must satisfied in order for the UAV to successfully change airspace then there must be 5 outcomes which may happen if one of the pre conditions has not been satisfied. A new type OUTCOME is defined to output an outcome for each pre condition if it has not been satisfied.

OUTCOME ::= success | outWithinMissionBounds | approverRefuses |
uCantHandleRestricitions | uNotInCorrectAirspace | uHasNoRestrictions

For each of the preconditions a schema is identified to show what happens if the pre-conditions are compromised.

Precondition	Schema
flying(u, a1) ∧ (¬ flying(u, a2))	**NotFlyingInCorrectAirspace** Ξ AutoAirspaceChange o: OUTCOME ¬flying(u, a_1) o = uNotInCorrectAirspace
approved(u, m)	**NotWithinTheMissionBounds** Ξ AutoAirspaceChange o: OUTCOME ¬approved(u, m) o = outWithinMissionBounds
approved(u, ap)	**NotApprovedByAirspaceHead** Ξ AutoAirspaceChange o: OUTCOME ¬ approved(u, ap) o = approverRefuses
res u ≠ { }	**UAVHasNoRestrictions** Ξ AutoAirspaceChange o: OUTCOME res u = { } o = uHasNoRestrictions
res ap ⊆ res u	**UAVCan’tHandleRestrictions** Ξ AutoAirspaceChange o: OUTCOME res ap ⊄ res u o = uCantHandleRestricitions

Since there are outcomes each time a precondition is compromised, a schema can be defined to identify when no preconditions are compromised and the UAV has successfully entered a new airspace.

SuccessfullyChangesAirspace
Ξ AutoAirspaceChange o: OUTCOME
¬approved(u, m) o = success

Using the functions defined, the entire operation of the negotiation between the UAV and other actors involved can be totalised.

TotalAirspaceChange== (AutoAirspaceChange ∧ SuccessfullyChangesAirspace)
∨ NotFlyingInCorrectAirspace
∨ NotWithinTheMissionBounds
∨ NotApprovedByAirspaceHead
∨ UAVHasNoRestrictions
∨ UAVCan’tHandleRestrictions

For a total negotiation we can have a successful negotiation to enter an airspace (**and** a success outcome) **OR** the change wouldn’t be within the mission bounds **OR** on request to enter the airspace the UAV is denied access **OR** the restrictions given to enter the airspace is more than the UAV can handle. By totalising the operation all the possible outcomes can be viewed. All exceptions have been identified and there should be no unidentified outcomes when conducting this operation.

3.0 TESTING THE CONCEPT

The COMPACT concept has been subject of several years of Dstl funded research. Earlier phases of the research programme demonstrated proof of concept in laboratory environments using representative UAS control stations, C2 and rotary and fast jet mission management systems. The most recent phase of the research programme tested the concept in a Live, Virtual and Constructive (LVC) trial

The trial, called Cardigan Bay 2017 (CB17), was undertaken to address a number of key research requirements identified by UK MOD (specifically, Dstl and DE&S) some of which directly related to Human Autonomy Teaming or Manned and Unmanned Teaming. From these requirements, additional requirements were derived to form the basis for a set of research hypothesis. Use Cases were constructed against which the hypotheses could be tested and success criteria associated with each hypothesis were derived – thereby identifying whether the trial achieved its aims. Those hypotheses that directly related to Human Autonomy Teaming and Manned Unmanned Teaming were;

- More operationally effective MUMT can be obtained when remote users, including airborne users, can exercise in-mission task-level control of UAS and/or their payloads to satisfy dynamically evolving mission requirements. This was compared against one UAS per operator)
- Autonomy provides an enabler for MUMT remote users to exercise task-level control of UAS and/or their payloads
- For autonomy in support of MUMT to be militarily exploitable it will need to be applicable within multi-tier C2 infrastructures
- Automated support for Policy Management and Negotiation is an enabler to exploiting autonomy for MUMT within multi-tier C2 infrastructures

3.1 Cardigan Bay 2017

CB17 was primarily undertaken at the Llanbedr airfield in West Wales, with additional functions undertaken at Aberporth and at the QinetiQ Farnborough Site (for geography, see Figure 5). The high level scenario for CB17 involved a Forward Operating Base (FOB) at Llanbedr, which would comprise small UAS flying (the Live element) as well as host virtual manned and unmanned platforms. A Main Operating Base (MOB) at Aberporth which would include both Constructive and Live elements. Note, Aberporth is a UK MOD range which can track air platforms flying in the in Cardigan Bay range with live track data from fast-jet traffic fused during trials together with pre-recorded data from Watchkeeper UAS sorties undertaken previously. In the case of the constructive element, recorded telemetry data was replayed in the LVC environment. Live data was also be injected into the LVC environment when aircraft are operating in the area. Farnborough acted as the main Joint Air Operations Centre (JAOC) undertaking theatre Air Battle Management. All sites were linked through wide area network that simulated various datalink and satellite communications protocols.

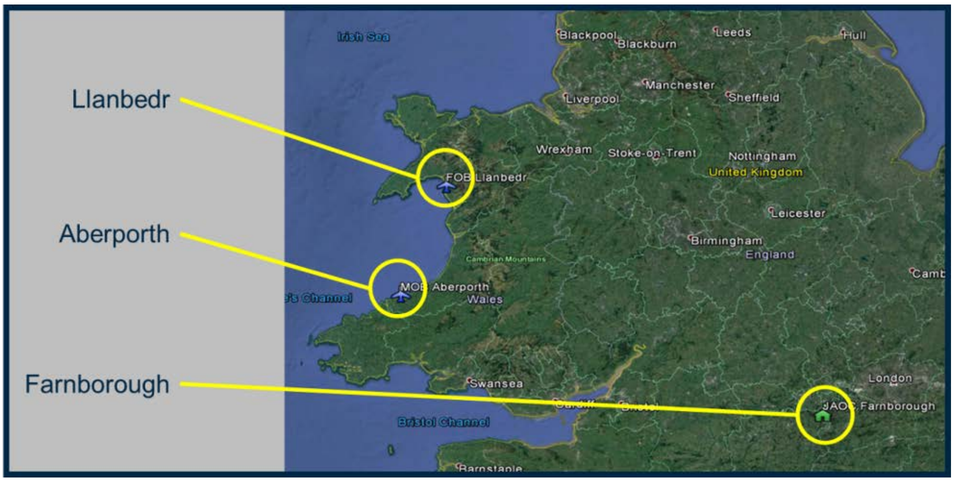

Figure 5 CB17 trial sites

The scenario used during the trial was derived from earlier activities undertaken in the West Wales and Cardigan Bay areas. Principally, it is a littoral scenario that encapsulates both maritime surveillance and land focussed Base Protection. The trials architecture, depicted in Figure 6, involved multiple platforms and stations spread between Farnborough and West Wales. Communications between stations was managed using standard datalink messaging (Link 16, STANAG 4586 and VMF) that would be applied in real operations. Many of the stations were manned by Military Participants (and some manned/operated by the technical team).

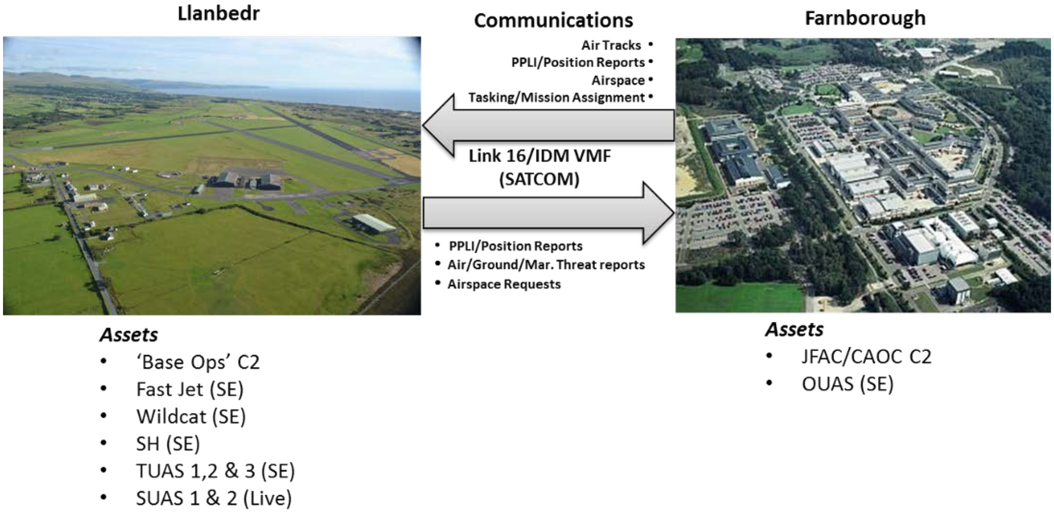

Figure 6 Trial set up

The trial design was focussed on three Use Cases that would address both littoral operations as well as Base Protection (BP). The first Use Case, in Figure 7, addressed the collection and dissemination of ISR from different airborne platforms, some manned, but most unmanned and of different classes (size, weight and endurance). Enemy forces are operating both at sea and in-land, and so for those in tactical command of friendly forces, they must schedule and plan missions for the air assets to ensure maximum coverage of the area of interest, evolving and adapting plans and schedules as the tactical situation unfolds. In order to coordinate air assets, airspace was managed by an Air Battle Manager (ABM). As the Use Case commenced, all Air assets are allocated airspace, within which they would remain unless operators negotiate with the ABM to access to other airspace.

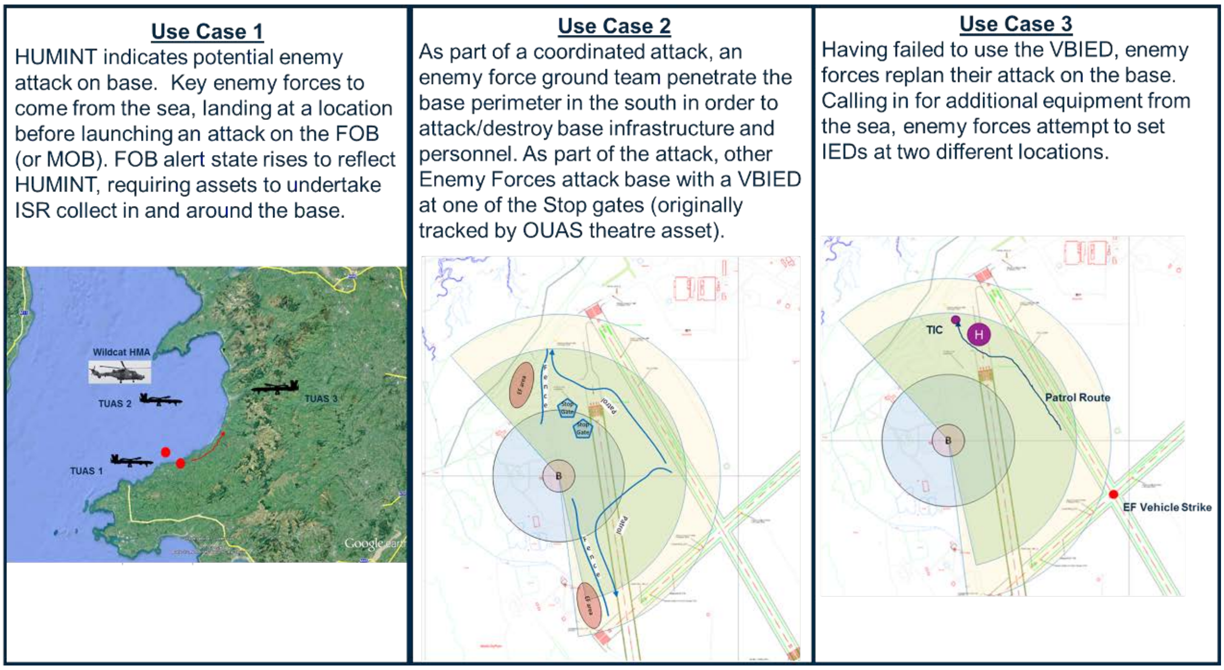

Use Case 1

HUMINT indicates potential enemy attack on base. Key enemy forces to come from the sea, landing at a location before launching an attack on the FOB (or MOB). FOB alert state rises to reflect HUMINT, requiring assets to undertake ISR collect in and around the base.

Use Case 2

As part of a coordinated attack, an enemy force ground team penetrate the base perimeter in the south in order to attack/destroy base infrastructure and personnel. As part of the attack, other Enemy Forces attack base with a VBIED at one of the Stop gates (originally tracked by OUAS theatre asset).

Use Case 3

Having failed to use the VBIED, enemy forces replan their attack on the base. Calling in for additional equipment from the sea, enemy forces attempt to set IEDs at two different locations.

Figure 7 Use Cases for CB17

The second Use Case (Use Case 2 in Figure 7) builds on the first but includes Red Forces preparing a coordinated attack on the FOB. The focus was therefore towards the protection of the FOB through utilisation of ISR assets to locate and identify potential threats and react accordingly. This requires assessment of ISR data from multiple sources, the coordination of which spans the three sites, SUAS at Llanbedr (the Base), TUAS from Aberporth, and the OUAS operated from Farnborough (under the JAOC). The specific threats to the base included small teams attempting to infiltrate the base perimeter and a vehicle borne improvised explosive device (VBIED).

The final Use Case (Use Case 3 in Figure 7) focusses on the management of ISR assets when reacting to multiple events which occur within the base or outside or the base perimeter. This requires assessment of ISR data from multiple sources and coordination by the Base Commander to utilise assets to the best effect, ultimately requiring coordination of strike and ISR platforms to neutralise threats. In this instance, the Base perimeter is infiltrated by red forces leading to troops in contact which the Base Commander must react to.

The trials plan accounted for one assessed run for each Use Case. A training run was undertaken at the beginning of the trial to familiarise the military participants with the systems, which included instantiations of the COMPACT components.

3.2 COMPACT Assessment

The COMPACT concept is realised through provision of architectural components incorporated into research systems that provide policy management functions and support to negotiation (via appropriate digital messaging). For the CB17 trial the systems that included COMPACT were the TUAS and SUAS control stations, manned rotary wing mission system, fixed wing mission systems, the C2 station and JTAC support system. There were 3 Military participants located at the Llanbedr trials site who were able to support the trial and operate the equipment. Their roles and systems they used in the scenario were as follows;

- Military Participant 1 was the Mission Commander at the FOB and operated the C2 and JTAC systems
- Military Participant 2 was the TUAS commander and operated the TUAS control station
- Military Participant 3 was the Rotary Wing aircraft commander (use case 1 & 2) and the SUAS operator (use case 3), and operated the Rotary Wing mission system and SUAS control Station respectively.

Each system had implemented policies and support to negotiations that were suitable for the relevant role. Details of implemented policies and negotiations are described in Table 1

Table 1. Policy Management and Negotiation support implemented on the CB17 systems.

	Policy Checks	**Negotiations**
TUAS CS	• Checking that auto-routeing (and manual routeing) conforms to airspace constraints • Checking that SAR plan conforms to SAR TTPs • Checking that Cooperative Laser Targeting with coordinating strike platform conforms to TTP and ROE • Checking that dynamic mission/route changes do not compromise Emergency Recovery Planning policy, including de-confliction and fuel planning	• Request for Airspace via Digital CHAT message (vertical cross-tier negotiation with C2) • Formatted messages for cooperative target engagement (horizontal peer-to-peer negotiation) • Digital authorisation for platform/sensor Level of Interoperability (LOI) hand-over (horizontal peer-to-peer negotiation)
SUAS CS	• Checking that auto-routeing (and manual routeing) conforms to Base Operating Procedures and airspace constraints (intended to be exercised in all use cases) • Checking that dynamic mission/route changes do not compromise Emergency Recovery Planning policy	• Request for Airspace via Digital CHAT message (cross tier negotiation with C2) • Digital authorisation for platform/sensor LOI hand-over (peer-to-peer negotiation) •

Rotary MM	• Checking that auto-routeing (and manual routeing) conforms to airspace constraints • Checking that coordinated strike behaviour with TUAS buddy-lase conforms to TTP and ROE • Warnings, Cautions and Advisories associated with TUAS platform/sensor control	• Request for Airspace via Digital CHAT message (cross-tier negotiation with C2) • Digital request for UAS platform/sensor LOI hand-over (peer-to-peer negotiation) • Digital authorisation for LOI hand-over of own EO/IR sensor to remote user (peer-to-peer negotiation) • Checking that Cooperative Laser Targeting with coordinating strike platform conforms to TTP and ROE (peer-to-peer negotiation
Fast Jet MM	• Checking that task/routes conform to airspace constraints • Checking that coordinated strike behaviour with JTAC conforms to TTP and ROE in terms of weapons effects and blue force positions	• Request for Airspace via Pre-formatted Digital message (Cross tier negotiation) • Checking that Cooperative Laser Targeting with coordinating JTAC conforms to TTP and ROE (peer-to-peer negotiation)
C2 MM	• Checking that task/routes conform to airspace constraints • Checking that task/routes are de-conflicted • Checking that air-vehicles are de-conflicted • Checking that Positive Identification (PID) is maintained during Track Handover • Checking that range, LOS, airspace and time constraints are met where a platform is employed as a Communications Relay • Checking that UAS tasking have associated emergency recovery plans that meet policy • Checking that when a Communication Relay role is to be withdrawn from a platform, it is safe/appropriate to do so • Checking that DACAS "9-Line" request can be serviced according to TTP	• Allocation of Airspace in response to digital airspace request via pre-formatted digital message (Cross tier negotiation intended to be exercised in all use cases • Track handover management • Checking that Cooperative Laser Targeting with coordinating JTAC conforms to TTP and ROE (peer-to-peer negotiation)

JTAC MM	• Checking that Airspace is allocated/de-conflicted • Checking that digital targeting messages are populated correctly • Checking that coordinated strike behaviour with TUAS buddy-lase conforms to TTP and ROE	• Allocation of airspace using J and/or K Series digital message to effector platforms (peer-to-peer negotiation) • Request for Airspace from C2 via Digital CHAT message (cross-tier negotiation with C2) • Digital request for UAS platform/sensor LOI hand-over (peer-to-peer negotiation) • Checking that Cooperative Laser Targeting with coordinating strike platform conforms to TTP and ROE (peer-to-peer negotiation)

To test the application of COMPACT concepts across these systems, suitably qualified and experienced military operators were required to operate these systems so that appropriate assessment could be undertaken. To assess the COMPACT, military operators were asked to provide a subjective assessment of the systems that they used. Principally, the assessment was concerned with measures of effectiveness and performance (MOE&P) for 4 key aspects of system capability; Autonomy, Policy Management, Negotiation and MUMT. Military operators were provided with MOE&P scoresheets to record their subjective views (Figure 8).

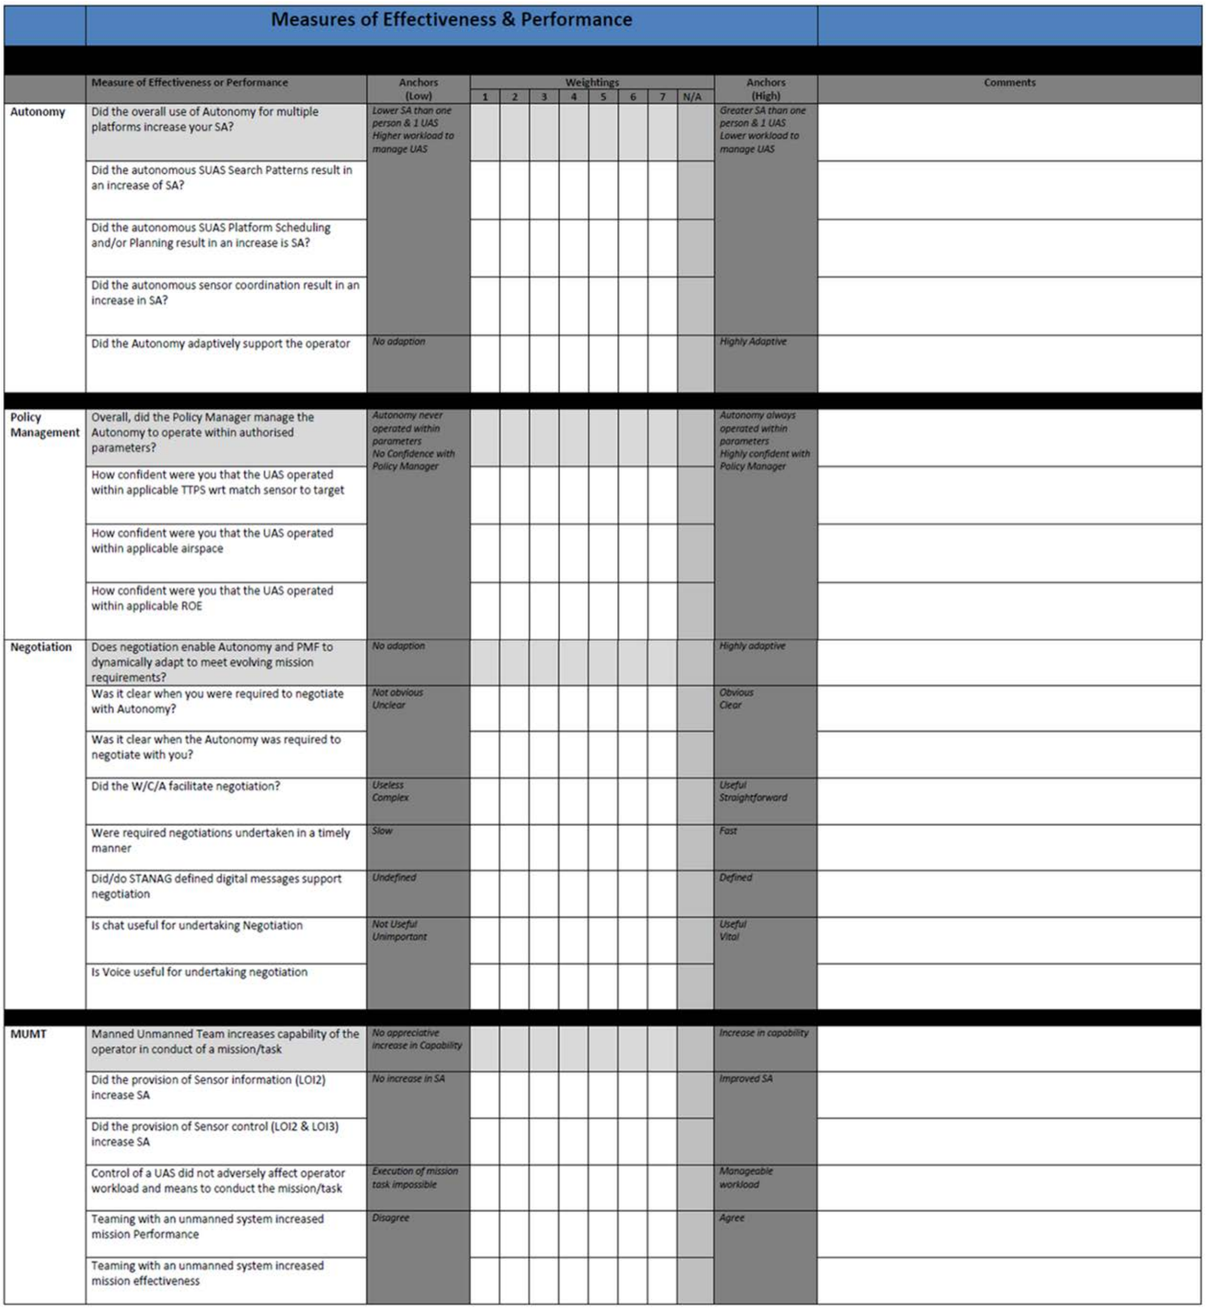

Measures of Effectiveness & Performance

	Measure of Effectiveness or Performance	Anchors (Low)	Weightings								Anchors (High)	Comments
			1	2	3	4	5	6	7	N/A		
Autonomy	Did the overall use of Autonomy for multiple platforms increase your SA?	*Lower SA than one person & 1 UAS Higher workload to manage UAS*									*Greater SA than one person & 1 UAS Lower workload to manage UAS*	
	Did the autonomous SUAS Search Patterns result in an increase of SA?											
	Did the autonomous SUAS Platform Scheduling and/or Planning result in an increase is SA?											
	Did the autonomous sensor coordination result in an increase in SA?											
	Did the Autonomy adaptively support the operator	*No adaption*									*Highly Adaptive*	
Policy Management	Overall, did the Policy Manager manage the Autonomy to operate within authorised parameters?	*Autonomy never operated within parameters No Confidence with Policy Manager*									*Autonomy always operated within parameters Highly confident with Policy Manager*	
	How confident were you that the UAS operated within applicable TTPS wrt match sensor to target											
	How confident were you that the UAS operated within applicable airspace											
	How confident were you that the UAS operated within applicable ROE											
Negotiation	Does negotiation enable Autonomy and PMF to dynamically adapt to meet evolving mission requirements?	*No adaption*									*Highly adaptive*	
	Was it clear when you were required to negotiate with Autonomy?	*Not obvious Unclear*									*Obvious Clear*	
	Was it clear when the Autonomy was required to negotiate with you?											
	Did the W/C/A facilitate negotiation?	*Useless Complex*									*Useful Straightforward*	
	Were required negotiations undertaken in a timely manner	*Slow*									*Fast*	
	Did/do STANAG defined digital messages support negotiation	*Undefined*									*Defined*	
	Is chat useful for undertaking Negotiation	*Not Useful Unimportant*									*Useful Vital*	
	Is Voice useful for undertaking negotiation											
MUMT	Manned Unmanned Team increases capability of the operator in conduct of a mission/task	*No appreciative increase in Capability*									*Increase in capability*	
	Did the provision of Sensor information (LOI2) increase SA	*No increase in SA*									*Improved SA*	
	Did the provision of Sensor control (LOI2 & LOI3) increase SA											
	Control of a UAS did not adversely affect operator workload and means to conduct the mission/task	*Execution of mission task impossible*									*Manageable workload*	
	Teaming with an unmanned system increased mission Performance	*Disagree*									*Agree*	
	Teaming with an unmanned system increased mission effectiveness											

Figure 8 Military Participant assessment score sheet.

Military participants undertook the assessment and completed the scoresheets after each run. The assessment process include an initial 'after action review' which acted as an opportunity to debrief and discuss the execution of the run prior to the formal assessment process. In addition to the MOE&P assessment, a formal assessment was undertaken by Dstl (results of which are reported separately).

4.0 CB17 RESULTS

The subjective results gathered from military operators are summarised in Table 2. The results generally show a steady improvement, with each aspect being rated with either low- or mid- range scores for the first use case run, and mid- to high-range scores for the final use case run.

The scores for the TUAS GCS *Autonomy* were originally rated quite low due to the fact that the military participant using the system was unaware that the policy manager had constrained the application of the autonomy software components due to airspace policy violations. However, this is precisely what the COMPACT concept is aiming to achieve and it was clear that the means by which policy violations are communicated to the operator required additional work.

Table 2 Subjective results summary from the CB17 trial

		Use Case 1	**Use Case 2**	**Use Case 3**
TUAS GCS	*Autonomy*	Low	Low	Mid
	Policy Management	Low	Low	Mid
	Negotiation	Low	Low	Mid
	MUMT	N/A	N/A	Mid
C2/JTAC	*Autonomy*	Mid	Mid	Mid
	Policy Management	Low/Mid	Low/Mid	Mid
	Negotiation	Low	Low/Mid	Mid
	MUMT	Low	Mid	Mid
Wildcat/SUAS	*Autonomy*	Mid	High	High
	Policy Management	Mid	High	High
	Negotiation	Mid	Mid	N/A
	MUMT	Mid	High	N/A

For the first run the results from the MOE&P assessment showed that for the mission commander, airspace policy violations by UAS were not raised at the C2 station and consequently no Policy Management or automated Negotiation functionality was triggered. The ABM functionality between Mission Commander at Base Ops and TUAS operator were therefore conducted by voice, which may have contributed to the TUAS operator assessment that autonomy was not working in relation to airspace, whereas in fact any airspace that was in the TUAS control station and was attributed correctly triggered autonomy and policy management functionality including Warnings, Cautions and Alerts (WCA) – but these were overlooked/ignored by the TUAS operator during the run. MUMT was scored higher by the mission commander reflecting that he could see the platform positions and route plans for both manned and unmanned aircraft that were operating, or tasked to operate, in the airspace delegated to him by the JAOC.

The Wildcat military operator scores for autonomy, policy management, negotiation and MUMT were mid-range to high. It is likely that these reflect the relatively maturity of that particular system and the fact that the operator had experience of using it in previous trials. Functionality exercised included instantiations of remote (TUAS) sensor feeds in integrated with the tactical map display showing remote platform stare-point and Field of View (FOV) for MUMT, together with digital LOI request and approval for remote platforms and sensors, which are a form of digital negotiation.

The TUAS operator scored autonomy low as it was not apparent to them that it was functioning. In actuality autonomy was working in the background in conjunction with policy management to produce WCA in relation to airspace, and also in relation to handling platform/sensor delegation to remote users (such as the Wildcat). The TUAS operator scored Policy Management low and this is believed again to be influenced by the operator not being clear as to what policy checking was being automatically applied behind the scenes, and only being aware of WCA when there was an issue, albeit that these were sometimes ignored. As a consequence, the operator stated that they were "not confident that that UAS operated within applicable policies" – despite this being the case (e.g. UAS staying within allocated airspace). Negotiation was scored low but this was influenced by the need to employ, for technical reasons, Chat messaging for some negotiations intended to be undertaken by preformatted datalink message. The operator found Chat messaging distracting and voice quicker. From the TUAS operator's perspective, MUMT was scored as "not applicable" because, once handed-off, they did not actively participate in MUMT with teaming exercised between RW operator and the handed-off platform. Taken in conjunction with the RW operator's scores, this should probably be taken to mean that MUMT was effective!

For the second run, the Mission Commander scored Autonomy mid to high with good SA of the Air Picture. Policy Management scores were low to mid-range because not all airspace issues were raised correctly to the operator, with the TUAS flying into unauthorised airspace. This was because of errors in the Chat message that should have triggered the automatic negotiation and airspace allocation. Negotiation scores were therefore low to mid-range. When the negotiation was repeated with the error corrected, the automated airspace allocation worked correctly and consequently there was no policy violation. As always, it was noted that voice was useful when autonomy (automatic negotiations in this case) did not work. MUMT was scored mid-range with SA assessed as having increased as a consequence of the TUAS being handed-off to the Wildcat helicopter.

The Wildcat HMA military operator scores for autonomy, policy management, negotiation and MUMT were mid-range to high, with SA assessed as having improved with hand-over of the TUAS, which correlates with the assessment of the Mission Commander. The Wildcat operator noted that control of the TUAS was achievable with simple commands.

During the third run, the TUAS operator scored autonomy higher than in the previous runs. The operator commented that the adaptive support provided by the autonomy sped up decision making, but most actions were still undertaking manually, indicating that there is potential advantage to be had from more automation/autonomy in relation the TUAS functionality exercised (which was search, track and cooperative targeting).

The TUAS operator scored Policy Management slightly higher (low to mid-range) than in previous runs, but still was not clear as to what policy checking was being automatically applied behind the scenes, and only being aware of WCA when there was an issue. Comments this time being "Operator unsure as to whether all actions were policy compliant" although the operator did qualify this with "Actions appeared to point to task (wrt airspace)". The Operator also commented "Messages conveying policies were minimal and could easily be missed" and "Alerts were clear but could also be missed" reinforcing assessments from earlier runs that some HMI development is required in relation to Policy Management.

Negotiation was also scored slightly higher than in previous runs, with the work-around of using a Chat message rather than a formatted message to request airspace now being familiar to the operator. However, the operator commented that "Chat was sometimes helpful but can slow you down depending on situation" which is a recognised limitation of Chat and indicates that its use needs to be carefully controlled, ideally in accordance with TTP such that operators understand and are familiar with when it should be used and for what types of information (the intended mechanism for the negotiation being a computer generated, pre-formatted message).

The TUAS operator scored MUMT for the first time with scores mid-range. The operator assessed that “Target handoff was quicker” and that “Teaming sped up some processes”. This is likely to reflect that the tools in the TUAS control station for supporting coordinated targeting in accordance with policy (TTP), together with the datalink integration which is a step towards automated negotiation, assisted the operator in conducting this time-critical task.

In this third run, the Mission Commander station was exercised both as a C2 station for managing airspace delegated to the Base Commander, but also as a JTAC station with support for digital targeting using the JTAC system. The Mission Commander scored Autonomy mid-range, commenting that there were “Multiple platforms but limited autonomy” which is true from the C2/JTAC perspective. C2 tools/functionality was designed to provide SA and decision support rather than autonomy. However, there were platform-to-platform de-confliction algorithms, route deconfliction algorithms in addition to the airspace de-confliction algorithms running in C2 as part of the Policy Management system.

The fact that these were not mentioned may be indicative of a similar issue to that raised in previous runs with respect to the TUAS control station – where autonomy/automation is being applied as a background task, particularly with respect to policy checking, some indication that this has happened and everything is satisfactory may be required. Notwithstanding the limited autonomy available at C2, the Mission Commander commented that “Assets used successfully” and “Platform and sensor coordination appeared to raise SA”, tending to indicate that whilst more autonomy might be useful, the functionality provided was sufficient.

Policy Management scores were mid-range with the comment that operations were “Policy compliant mainly due to the operators” and that “UAS operated in allocated airspace (but none requested)”. This is again indicative of the autonomy at the TUAS and SUAS control stations that attempts to ensure that all UAS operations remain within authorised (allocated) airspace was not being recognised by either UAS operators or the Mission Commander.

Negotiation scores were mid-range with the comment “Clear when negotiation required but not successful” and “Negotiations did not always work in timely manner” which indicates that there was scope for automated negotiation within clearly recognised circumstances but that the Chat work-around was not a viable solution. When undertaking the JTAC role, the operator commented that Chat was useful but that voice was a good fall-back – consistent with previous runs and comments from other operators.

In relation to MUMT, scores were mid-range with the Mission Commander/JTAC assessing that LOI2 of UAS sensors provided for greater SA and that LOI3 not used, the implication being that it was deemed to offer no additional value over LOI2 in this Use Case. A further comment was that it “Appeared control of more than 2 UAS slowed mission” and this is likely to be influenced by the amount of autonomy available, since a further comment was “If autonomy works, could increase mission performance and effectiveness”.

5.0 CONCLUSIONS

The CB17 trial provided subjective results from suitably qualified and experience military participants. The complexity of the trial was a determining factor when exercising the COMPACT concept. Limited timeframes meant a short training opportunity for the military participants which led, on many occasions, to system behaviours that were unexpected, which in turn led to lower subjective ratings – particularly for applications of Autonomy.

What became apparent was in the case of Autonomy, the policy manager was restricting execution of specific behaviours due to policy violations. This was intended, but provided a clear indication that this was not being communicated to the Military participant during the runs. Therefore, a key outcome from CB17 was the need to ensure that policy violations are more clearly alerted to the operator.

Each system's design reflects Military Advisor advice that operators would not want to be constantly distracted by information when autonomy and policy management was working with no issues/policy conflicts, but alerted when there are issues/policy conflicts. The CB17 trials results may however indicate that a compromise solution is required, perhaps similar to the "traffic lighting" of Survivability, Effectiveness, Timeliness" under previous research with "SET" greyed out when automatic behaviours and checking are not being applied (the SET concept is described in [14]).

Negotiations, particularly airspace negotiations, were deliberately accounted for in the design of the use cases, such that there would always be a need for platforms to require entry into new airspace. A limiting factor in the design of the systems was the need to use a specifically formatted freetext message to convey the negotiation (this is because there is no support for negotiations in standard digital messaging). Military Participants initially found this cumbersome however, as the trial continued, began using the negotiations as intended. What was apparent was that Negotiations were continuously being undertaken, albeit via voice communications, validating the intent to investigate automation of these in future trials of COMPACT.

The CB17 trial was in part designed to support the evaluation and assessment of the current instantiation of the COMPACT concept architecture. The trial demonstrated that a key element of COMPACT, the Policy management function, has the potential to provide a viable means of controlling and managing the application of Autonomy on complex systems and that applying negotiations provides a means to manage and mitigate and potential policy conflicts.

6.0 REFERENCES

[1] Cottrell, R. J. Thomas, M. J, Hybrid Architectures For Adaptable Autonomy, QinetiQ, November 2013.

[2] Stevens, B. Investigation of Advanced Novel Computing Techniques Interfacing with Integrated Modular Avionics, DERA UK MOD, 1998 (Unpublished)

[3] Geddes, N. D. Shalin, V. L. Intelligent Decision Aiding for Aviation, Linkoping, Sweden: Centre for Human Factors in Aviation, 1997

[4] Banks, S. B. Lizza, C. S. Pilot's Associate – A Cooperative, Knowledge-Based System Application, IEEE Expert, June 1991

[5] Miller, C.A Hannen, M. D. The Rotorcraft Pilot's Associate: design and evaluation of an intelligent user interface for cockpit information management, Knowledge Based Systems, 12 (p443-456), 1999

[6] Catford, J. R. The Mission Management Aid – A New Level of Systems Integration in Future Fighter Aircraft, IEEE Colloquium Integrating Issues in Aerospace Control, April 1991

[7] Gibson, C. P. Garrett, A. J. Towards a Future Cockpit – The Prototyping and Pilot Integration of the Mission Management Aid (MMA), AGARD Conference Proceedings No. 478 – Situation Awareness

in Aerospace Operations

[8] Howells, H. Davies, A. Macauley, B. Zanconato, R. Large Scale Knowledge Based Systems for Airborne Decision Support, Knowledge Based Systems, 12 p215-222, 1999

[9] Taylor, R. M. et al. Cognitive Cockpit Engineering: Coupling Functional State Assessment, Task Knowledge Management, and Decision Support of Context-Sensitive Aiding, in Cognitive Systems Engineering in Military Aviation Environments: Avoiding Cogminutia Fragmentosa! (eds McNeese, M.D Vidulicj, M. A) 2002

[10] ECOA – European Component orientated Architecture, www.ecoa.technology, accessed 23/08/2018

[11] Burski, L. A. *Agreement Patterns for Agent Negotiations, 2017, QinetiQ (unpublished)*

[12] A. Iglesias and Mercedes Garijo and Jose Ignacio Fernandez-Villamor and Jose Javier Duran, Agreement Patterns, 2009

[13] Luck M, McBurney P, Computing as interaction: agent and agreement technologies. In: IEEE SMC conference on distributed human-machine systems, pp 1–6, 2008

[14] Schreiber, A, Th. et al (2000), Knowledge Engineering and Management - The CommonKADS Methodology, MIT Press

[15] Cottrell R.J., Dixon D.G., Hope T. and Taylor RM. *Operator in the Loop? Adaptive decision support for military air operations*, In HCI International 2005, Proceedings of the 1st International Conference on Augmented Cognition, Las Vegas, AZ, 22-27 July 2005. Mawhah, NJ: Lawrence Erlbaum, 2005.