A novel ontological approach to semantic interoperability between legacy air defence command and control systems

Chris Partridge*

BORO Solutions,
25 Hart Street, Henley on Thames,
Oxfordshire, RG9 2AR, UK
E-mail: partridgec@borogroup.co.uk
*Corresponding author

Mike Lambert and Mike Loneragan

QinetiQ,
Room 60 CSL, Portsdown Technology Park,
Southwick Road, Cosham,
Portsmouth, Hants, PO6 3RU, UK
E-mail: mjlambert@qinetiq.com
E-mail: mjloneragan@qinetiq.com

Andrew Mitchell

BORO Solutions,
25 Hart Street, Henley on Thames,
Oxfordshire, RG9 2AR, UK
E-mail: mitchella@borogroup.co.uk

Pawel Garbacz

John Paul II Catholic University of Lublin,
Al. Racławickie, 14, 20-950 Lublin, Poland
E-mail: garbacz@kul.pl

Abstract: In common with many other government defence departments, the UK Ministry of Defence (MoD) has realised that it has a plethora of legacy systems that were procured as domain specific with little emphasis given to integration requirements. In particular, it realised that the lack of integration between a significant number of the legacy air defence command and control (AD-C2) systems meant it could not deliver the increased agility needed for joint force AD and that current approaches to integration were unlikely to resolve the problem. They realised that they needed a new approach that demonstrably worked.

This paper describes a programme initiated by the MoD to address this problem through the formulation of a novel solution and its demonstration in the tactical AD-C2 environment using a sample of these existing legacy systems. It describes the ontological solution deployed to resolve the 'hard' semantic interoperability challenge. It outlines the physical and semantic architecture that was developed to support this approach and describes the implemented planning and collaborative execution (PACE-based) and semantic interoperability engine (SIE) solution.

Keywords: tactical air defence; command and control; C2; legacy systems; ontology; BORO; planning and collaborative execution; PACE; joint tactical air defence integration system; J-TADIS; semantic interoperability; intelligent defence support systems; know-what.

Reference to this paper should be made as follows: Partridge, C., Lambert, M., Loneragan, M., Mitchell, A. and Garbacz, P. (2011) 'A novel ontological approach to semantic interoperability between legacy air defence command and control systems', *Int. J. Intelligent Defence Support Systems*, Vol. 4, No. 3, pp.232–262.

Biographical notes: Chris Partridge is the Chief Ontologist at BORO Solutions and a Senior Researcher at Brunel University. He holds an MA in Mathematics and Philosophy from Oxford University. He has been working on the pragmatic application of ontology to information systems since the late 1980s. His work focuses on semantic interoperability and semantic application modernisation, particularly the mining of ontologies from legacy systems. He is the author of a number of articles and the book *Business Objects: Reengineering for Reuse*.

Mike Lambert is the Chief Software Engineer for the Maritime Systems Group at QinetiQ. He is a Chartered Engineer and member of the British Computer Society, with a BSc in Electrical and Electronic Engineering from UMIST and a Post Graduate Diploma in Computing. He is responsible for various QinetiQ systems, notably operational combat system analysis tools for the type 23 frigates and type 45 destroyers. He has an interest in distributed systems, collaborative planning frameworks and particularly systems integration strategies.

Mike Loneragan is the Chief Engineer for the Maritime Systems Group at QinetiQ. He holds a BSc in Electrical and Electronic Engineering, MSc in Guided Weapon Design and an MA in Defence History. He currently leads the design and acceptance into service of new equipment into type 23 frigates for the MoD. He has significant experience in communications, interfaces and information management.

Andrew Mitchell is a Principal Ontology Consultant for BORO Solutions. He has a First-class honours degree in Engineering Studies from Portsmouth University. His current work involves visualising, validating and verifying semantic models; these models are used in projects such as software modernisation and semantic interoperability. Previous work has focused on analysis for the military domains of littoral manoeuvre and command and control within air defence; including system integration and decision support projects. He has five years' experience of ontology and 15 in the defence science industry.

Pawel Garbacz is an Associate Professor of Philosophy at John Paul II Catholic University of Lublin. His research interests cover applied ontology for CAD/CAM systems, application of epistemic logic to automatic deduction systems, database design for history, and philosophical foundations of formal logic.

1 Introduction

In common with many other government defence departments, the UK Ministry of Defence's (MoD) need for network enabled capabilities has brought about a realisation that it has a plethora of legacy systems that were procured as domain specific with little emphasis given to integration requirements. In particular, it realised that the lack of integration between a significant number of the air defence (AD) command and control (C2) (AD-C2) legacy systems meant it could not deliver the increased agility needed for current joint force[1] AD. What made the situation more difficult was that current approaches to integration[2] were unlikely to resolve the problem. It was clear that they needed a new approach that demonstrably worked and asked QinetiQ[3] to provide it. Specifically, they asked QinetiQ to devise a novel solution and demonstrate it worked in the AD-C2 environment using a representative sub-set of the existing legacy systems. This paper describes the demonstrated solution and some key lessons learnt in its implementation.

The two strategic lessons learnt are firstly the usefulness of an ontological approach to semantic interoperability[4] and secondly the need for an architecture that is capable of using the ontology and meeting the demanding requirements of the C2 environment.

2 Air defence context

2.1 The military situation

According to UK military doctrine, control of the air is fundamental to the success of joint operations and is normally achieved through a mix of defensive counter air (DCA) and offensive counter air (OCA) operations. DCA is primarily executed as AD operations defined as "all measures designed to nullify or reduce the effectiveness of hostile air action"[5].

The AD commander needs timely access to AD related information across the battlespace from AD capable sensors to be able to employ effectors[6] for the protection of friendly and neutral forces. The spectrum of potential AD threats in deployed operations is increasing as the world security situation changes. It ranges from conventional high technology air systems such as cruise missiles to improvised weapons employed asymmetrically. These threats include a number of difficult air targets which present particular detection and engagement challenges to the AD system. Currently, the systems for managing the sensors and effectors are not particularly dynamic, and improving that dynamism is an obvious way to counter these threats.

2.2 Current technology

In recognition of its importance, substantial resources have been committed to AD-C2 capability in recent years. However, both current and planned systems have, in general, been procured as environment specific and so have an environment specific architecture (including connectivity, communications systems and quality of service). This restricts the extent to which dynamic changes to extant plans and current operations can be implemented. In particular, changes to plans and airspace control procedures that require coordination between AD assets in the land and either the maritime or air environments are generally inefficient. The UK MoD perceived a need to remedy this situation and deliver increased flexibility and dynamism in joint force AD in the near term through improved interoperability within current and planned systems.

3 J-TADIS programme context

3.1 Programme initiation

In March 2007 the UK MoD asked QinetiQ to research novel techniques for system integration in the AD domain to improve interoperability between operational systems. The main focus of the research was to investigate software solutions capable of significantly improving the joint AD environment efficiency by enhancing interoperability. QinetiQ proposed a programme of work called the Joint Tactical Air Defence Integration System (J-TADIS).

The QinetiQ proposal resulted in a substantial two year applied research programme. This programme involved the demonstration of the integration of existing operational AD systems from major defence companies using representative data in a realistic environment. The final system achieved a technology readiness level (TRL) 6. TRL 6 is defined as:

> "Prototype system, which is well beyond that of TRL 5, is tested in a relevant environment. Represents a major step up in a technology's demonstrated readiness. Examples include field testing a prototype in a high fidelity laboratory environment or in a simulated operational environment operating under proposed protocols[7]".

3.2 Scope of the demonstrator

A key criterion for assessing the success of the demonstrator was whether it showed the increased agility needed for joint AD operations. Subject matter expert (SME) feedback identified that improving the integration of 'last minute' collaborative planning and execution (CP&E), which depends upon a number of legacy systems interoperating in near real time, would significantly increase agility in this domain. Hence, this area of joint AD operations was selected for the demonstrator.

In this selected area semantic quality is a central concern. Poor semantic quality can, literally, be fatal. This increased the importance of the semantic requirement that the various systems shared a sufficiently detailed common understanding of a shared real world. Implementing this requirement involved subjecting the legacy system data to a high level of semantic quality assurance. The detailed common understanding brought a number of important benefits for joint AD, a couple of which are illustrated in the paper with simple examples.

3.3 Ontology enabling semantic interoperability

It was recognised from the start of the programme that the most challenging requirement was semantic interoperability. In the early stages of the programme, research uncovered a novel technique – BORO – that used ontology to enable semantic interoperability and provide semantic quality assurance for this type of legacy systems integration. This technique was incorporated into the programme in stages. It influenced the architecture of the final system leading to two architectural components; planning and collaborative execution (PACE) and semantic interoperability engine (SIE) – in addition to the legacy system components.

4 Ontological context

Ontology is a relatively new technique in the AD-C2 operational environment. It is described below.

4.1 What is an ontology?

For the purposes of semantic interoperability, the traditional philosophical (metaphysical) notion of ontology is useful (Kuśnierczyk, 2006) – where this is “the set of things whose existence is acknowledged by a particular theory or system of thought”[8]. This view was famously summarised by Quine, who claimed that the question ontology asks can be stated in three words ‘What is there?’ – and the answer in one ‘everything’. Not only that, but tongue in cheek, he also said ‘everyone will accept this answer as true’ though he admitted that there was some more work to be done as ‘there remains room for disagreement over cases’ (Quine, 1948). As these clarifications show, ontology in this sense is directly concerned with what exists – in what business modellers call the ‘real world’[9].

From the perspective of semantic interoperability, each system’s data can be regarded as a ‘theory’ that acknowledges the existence of a set of objects – its ontology. The task of the ontological analysis process is then to work out what ‘real world’ objects this data commits to. This is not an easy task as the surface structure of the data is usually a very poor guide[10] (which helps to explain why the data for the same domain from different systems can have quite startlingly different structures). A key tool in the analysis is a top ontology to provide a framework with which the analysis is structured.

4.2 The BORO approach

The BORO approach was originally developed in the late 80s and early 90s to meet a requirement to reengineer legacy systems (Partridge, 1996, 2005). The prime challenge of the reengineering was to clarify the underlying ontology of the systems, and the work focused on developing a process for mining ontologies and a top ontology tailored to form its foundation[11]. Early feedback on the top ontology established that a key factor was to make a series of clear ‘metaphysical choices’[12] to provide a solid (metaphysical) foundation. A key choice was for an extensional (and hence, four-dimensional) ontology which provided neat criteria of identity. Another conscious choice was for a form of ontological realism – which assumes for engineering reasons that the world which science describes exists and also that the types used in this description also exist[13]. Using this top ontology as a basis, a systematic process for reengineering legacy systems was developed. From a software engineering perspective, a key feature of this process was the identification of common general patterns, under which the legacy system was subsumed. It has been substantially developed since then. Although much of the work (such as this) is in the private domain, elements of it have appeared in the public domain, including a number of standards. The ISO standard, ISO 15926 – industrial automation systems and integration – was heavily influenced by an early version. The IDEAS (International Defence Enterprise Architecture Specification for exchange) standard is based upon BORO, which in turn was used to develop DODAF 2.0[14].

4.3 Aspects of the BORO approach

Some analysis processes for a common ontology will start with SMEs and a 'clean sheet of paper'. BORO works from the premise that experience seems to show that humans (including SMEs) are not particularly good at specifying the ontology with the degree of accuracy needed for computer systems. A better source is working operational systems – these plainly work to a sufficient degree of accuracy for current operations[15]. The issue when working with the legacy systems is how to mine the ontology when the surface structure is often seriously misleading.

One of the approaches BORO uses to investigate the underlying structure is a focus on the legacy data (Partridge, 2005) (where many other analysis processes focus on the data schema). The issue here is that in most operational systems users will have had to work around the schema to get the system to perform as required. With the result that the schema is misleading. Hence, the data is a far more trustworthy source for the analysis than the data schema.

BORO has developed a variety of ways of testing the semantic quality. One task is the devising of a suitable semantic stress test. In this programme, the collaborative planning and execution environment required data from a number of legacy systems with quite different data structures which needed to be consolidated, manipulated and then sent back to the source system. In this situation, feedback on the semantic quality issues shows up quickly as operational problems. The high quality of the stress test raises confidence in the semantic quality of the overall system.

5 Setting the scope of the demonstrator

As a first stage in setting the scope, the requirements for a fully fielded capability were drawn up. These were reviewed by SMEs to establish a scope that would provide a reasonable test for feasibility of the proposed solution.

The SMEs divided the overall scope into three phases: pre-tactical planning, tactical planning and engagement. In the pre-tactical planning phase collaboration is primarily involved with shaping the battlespace; activities such as the positioning of sensors and effectors in response to intelligence surveillance target acquisition and reconnaissance (ISTAR) updates. The tactical planning phase is 'last minute' collaborative planning and execution CP&E; that is, in the final 35–5 minutes of an engagement. In the engagement phase, that is in the final few minutes, interventions would be more concerned with such things as target tracking, sensor cueing, engagement veto decisions and the like.

The SMEs determined that a system that could support tactical planning would have the greatest impact on the final outcome. Hence, the J-TADIS programme focused on collaborative planning during 'last minute' tactical planning phase and collaborative execution of an engagement.

5.1 Flexibility and agility

This put a requirement on the J-TADIS capability to support flexibility and agility, and reduce tempo drag at the tactical level. There are currently limitations in being able to support fast-changing tactical operations under the UK's manoeuvrist doctrine. The limiting factors include both technical infrastructure and procedures. One straight-forward way to achieve flexibility is to maximise the principle enshrined in Joint Warfare Publication (JWP) 3–63 of de-centralised execution. This demands that control is delegated to the lowest practicable level commensurate with the requirements implicit in tactical level operations. It was proposed that J-TADIS should make significantly more de-centralised control practicable through enabling a collaborative approach to both planning and execution.

5.2 Selecting representative legacy systems

Generally, the legacy systems that are used in the AD domain were not designed to share data with other systems. Furthermore, they usually were not designed to use event driven technologies. This means that currently the frequency of information transfer is limited to that supported by the legacy system and is often a manual process. This makes it difficult for such legacy systems to play a role in event driven systems of systems, such as those designed to support interactive collaborative planning.

The NATO integrated command and control software for air operations (ICC), BAE's ground-based AD, battlefield information systems application (GBAD BISA) and Thales' MPlanIt mission planning system were chosen for the programme as they were good examples of this issue.

For the final TRL6 demonstration, a complex and challenging scenario was devised that illustrated how integrating these operational legacy systems would enable a level of agility that had not previously been possible.

6 Crystallising the requirements for the demonstrator

To overcome current limitations of a largely procedural joint tactical air defence command and control (JTAD C2) infrastructure, J-TADIS sought to expose or externalise (make available) stove piped AD-C2 systems making their functionality and information more accessible and dynamic. This would provide a significantly improved set of services to the wider AD community.

Within this scope, the requirement for J-TADIS crystallised into two main areas. First, techniques had to be developed to enable legacy operational systems to share their information with the wider planning community within timescales relevant to the tactical planning phase. Secondly, the environment which would share this information must allow distributed users to collaborate as they manipulated this information.

Specifically, this involved a requirement for two main system functions:

- provision of a common set of information to support identified joint force AD activities
- access to (existing or new) services that allow users to make effective use of this common information.

J-TADIS proposed to meet this via two main mechanisms. The first supports the identified joint AD activities through provision of services that allow wider access to a common set of information, for example airspace control means (ACMs). The second provides access to existing or new services that allow users to manipulate and make effective use of this common information, for example, by allowing users access to collaborative airspace management functionality.

The term ‘collaborative planning’ is widely used but has no clearly agreed upon definition. When the term is used here it refers to activities supported by systems that enable distributed users to collaborate on the manipulation of a plan, in such a way that any proposed modifications made by a planner are immediately available to his colleagues. It will generally be the case that multiple variants of a plan will exist simultaneously as different groups propose different solutions. These proposals can be thought of as ‘what if’ scenarios as planners search to determine optimal responses. Military planners are often in competition for such resources as airspace and assets and regularly have to coordinate activities with other teams, leading to inter-plan constraints. To successfully manage these types of complex plans at the required tempo, automated consistency checking is vital to highlight the impact, both within and across plans, of proposed changes.

It became apparent that an enterprise level framework, capable of managing a wide spectrum of information types, would be required. Furthermore, this management would have to include elements of collaborative planning and execution in some sort of structured workspace, i.e., a workspace capable of dynamically consistency checking operator actions.

6.1 Increased information fidelity

The requirement for increased agility implies that more C2 detail should be made available for exploitation. In airspace control terms, for example, there could be a greater number of smaller high-density airspace control zones (HIDACZs) with more and shorter time of validity slots, greater differentiation of targets, more fidelity in the rules of engagement (ROE) and hostile act criteria. This information can be used, for example, to identify when to relax such things as weapon control states (WCS), which is difficult to achieve with the currently enforced procedural controls. Increasing the information fidelity in this way would impact the information modelling and the software architectures used to disseminate it, as explained later in the paper, in the architecture section.

Under normal circumstances, the forces attempting to perform a task at the tactical level know what they wish to achieve and their immediate concerns. J-TADIS would provide this information in a collaborative working environment that would enable these forces to utilise a far more bottom-up pro-active request process. They could, therefore, approach the appropriate superior authority and indicate their level of knowledge and their ability to execute a task in a particular area, over a certain time frame.

Currently, fratricide is prevented by measures designed mainly to separate functions using space and time factors with large margins designed to reduce potential errors. J-TADIS could improve on this situation as it has the ability to provide rapid feedback on remote system status, such as WCS. This type of information can be used to provide additional factors to assist with maximising airspace usage safely.

6.2 Preserving individual communities' semantics

Military communities and their systems often emerge with a focus on a particular goal. They develop their own particular culture with its own terminology and specialised world view.

This creates a dilemma in the joint force environment, how to continue to support specialised planning and operational staff, with their legacy systems that directly support their ways of working, while at the same time provide infrastructure that allows an enterprise view and all of the benefits that would bring. This is particularly acute in the case of CP&E, which cannot be effective unless information can not only be shared with sufficient timeliness, but also that the implications of proposed changes in plans, necessary for agility, can be made apparent to all stakeholders. Effective CP&E requires, inter alia, cross-plan consistency checking, determination of the implications of plan change on other plans, decision support and the visualisation of cross-plan constraints.

This translates into a requirement to preserve the semantics of the different military communities (and their systems) in the joint AD environment while simultaneously providing a single common semantics for the totality of the information involved.

7 Overall architecture

There are well-known difficulties involved with invasive architectures that involve modifying legacy systems from multiple suppliers. The problems are both commercial and technical. Commercial concerns may place restrictions on the ability to change the legacy systems and the access to required data as file formats or database schemas may be propriety. Technical problems include the lack of suitable documentation for the objects that need changing.

Accordingly, the first major architectural decision was to adopt a non-invasive approach and treat the legacy systems as components. The second major architectural decision was to segregate out the semantic interoperability function into a single component – driven by separation of concerns (Dijkstra, 1982) considerations – leading to this three layer architecture:

- common information layer
- semantic interoperability layer
- legacy system layer.

The legacy systems are as given. The common information layer was implemented using PACE and the semantic interoperability layer was implemented using SIE.

8 Common information layer: information modelling and dissemination strategies

Within the common information layer, it was necessary to address the current limiting factors in both information modelling and dissemination.

8.1 Key physical architecture requirements

There are three main architectural requirements that impacted the decisions regarding information modelling and dissemination:

- the trade-offs between fixed models and generic configurable models
- the complexity of managing a variety of plans, each with layers of alternative options, in a dynamic event driven environment, constrain the architecture in particular ways
- the performance implications of disseminating generic configurable layered models of this type.

8.1.1 Configurable generic models versus fixed models

The choice of when to use a generic model over a fixed model is influenced by two main criteria. The first criterion is the relationship between the scope of the information to be modelled and its potential to be generalised. If the scope is small then the additional effort required to build a generic system may not be cost effective. If the scope is large but the information has little potential to be generalised, then the resulting generic model may not offer much compaction. If, however, both the scope and the potential for generalisation is large then the generic approach will scale better than a fixed model.

The improvement is a function of processing that can be applied at the higher abstract levels of the model and so does not require reinventing each time at the detailed level. C2 information concerned with multi-environment planning and execution has considerable potential to be generalised. The same common patterns of asset allocation, tasking authority, temporal and geographic space allocation endlessly repeat and lend themselves to generalisation.

The second criterion is the expected stability of the model. If the scope of the information to be modelled is well understood and is unlikely to be subject to major changes over the lifetime of the system, a fixed model may be preferable. If, at design time, the final scope is unknown, which is the norm for IT systems, a configurable generic approach will allow the scope to be increased through changes to the configuration rather than the software. There are considerable advantages in modelling many warfare environments using the same generic model. Often planners from different environments, trying to achieve different goals, compete for assets and geographic space allocation. This has the effect of creating constraints between their various plans. Modelling those constraints, within a common generic model, allows the impact of proposed changes to be better understood by the commander resulting in a safer and more efficient operation.

J-TADIS uses information from all three environments air, land and maritime. Its scope encompasses long term planning through tactical planning to execution. In this context, a fixed information model is likely to unnecessarily constrain the potential development of the system.

8.1.2 Layered information

To achieve the required agility, J-TADIS needed to support collaborative working over multiple variants of plans and courses of action. Consider the case when a number of different alternatives to a plan are to be maintained. One strategy is to create complete copies of the plan and then modify them. This is inefficient for storage and for dissemination. If different users are working on different parts of the plan it can become extremely complex to manage. A better solution is to implement each alternative as a set of changes to the main plan. For the purposes of this paper we shall refer to this as layered information.

Different user groups would need to work on different alternatives. From a technological point of view any provider of this information would need to know which consumer, required which 'slice' through this layered information. The appropriate information must be constructed for each subscribing user group and this is not just a filtering problem. This is further complicated when the collaborative planning is considered. As users manipulate the information, the changes must be distributed to only those subscribers who subscribe to the alternative being changed.

8.1.3 Efficient dissemination

CP&E requires the sharing of plans and situation awareness data dynamically as the information changes. If rapid decision making is the goal, CP&E must occur within, at least, the time limits of a typical human conversation. The military networks that will host J-TADIS have bandwidth and latency constraints and it is crucial that information distribution strategies minimise the network footprint. Current military messages such as the air control order (ACO) and the air tasking order (ATO) are effectively entire plans in themselves. Updates to these plans tend to be delayed, and then managed in bulk, particularly for detached forces. The information models described above, however, are implemented as interconnected webs of objects so possibilities exist to disseminate only the changes (deltas) to individual objects. These deltas can be relatively small. Changing a waypoint on a route, for example, may take a few tens of bytes.

8.2 Review of potential architectures

Having defined the key requirements on the architecture, a review was carried out into traditional architectural frameworks, particularly, real time distributed architectures and message-based architectures such as enterprise service bus (ESB)[16].

Real time data distribution architectures can be constructed from technologies such as the common object request broker architecture (CORBA)[17] or the more modern data distribution service (DDS)[18]. Both excel at the rapid dissemination of messages where the message structures and the type of subscription is well defined and unlikely to change. They are not so well suited to the type of functionality required for the common information layer. In fact, for DDS, the only option is to disseminate all layers and force the client to construct the layer it requires. This is unsatisfactory from both performance and security perspectives.

An ESB requires an enterprise message model (EMM) which is an enterprise wide set of message formats usually defined according to some meta-model such as the XML[19] schema definition language (XSDL).

These EMMs are generally implemented using a fixed set of message formats; see Section 8.1.1 ‘configurable generic models versus fixed models’ above.

ESBs are usually constructed using frameworks which provide the basic functionality for distribution, persistence and the addition of services. With these frameworks there is little opportunity for optimising the information distribution techniques used, particularly for layered information.

8.2.1 PACE/GIM

8.2.1.1 Overview

As the traditional architectures were not suited to the J-TADIS requirements, a new framework was developed and it became known as PACE.

The PACE framework was designed to support distributed planning teams through interactive group working sessions. Its generic information model (GIM) can be configured to manage C2 information from different warfare domains. Its strength lies in its ability to support near real time decision making through its event driven architecture. The PACE framework employs a service-based approach specifically tailored to provide event driven publish subscription over rapidly changing complex C2 information. It was designed to manage the further complication of different views of this C2 information representing alternative courses of action available to the commander.

8.2.1.2 Plug-in architecture

The PACE framework allows a variety of plug-in modules, such as user interfaces, business logic or interfaces to external systems, to dynamically interact with the GIM. Any proposed changes to the GIM can be consistency checked, as they happen, so immediate feedback of plan consistency or useful advice is relayed back to the user.

PACE differs from the traditional ESB in two main ways. It employs a configurable generic approach for its enterprise data modelling and it uses that genericity to optimise the dissemination, visualisation and processing of that data.

8.2.1.3 Information handling – the generic information model

The heart of the PACE framework is the GIM. The GIM is a framework that supports ontologies. This separates technical aspects, such as performance which are built into the framework, from the details of the content, which are loaded as needed. This is a classical separation of concerns architecture.

The GIM has a top ontology that captures some of the core military planning patterns; including the relationships between temporal activities, geospatial information, resource allocation, ownership and control.

For a particular application, the GIM is configured with the background ontology for the relevant domains. This means that PACE-GIM sets no constraints on what C2 domains it can support and to what level of detail.

When deployed, users and legacy systems can add further information to the GIM through the various PACE user interfaces. The GIM is then the source for common C2 information for both plan and execution time. Figure 1 shows a simple example GIM fragment configured as a scheduling and resource plan for aircraft sorties. The light grey rectangles show classes defined as part of the configuration, the white rectangles show objects created by the user at plan time. The fragment shows a Tornado GR4[20] sortie is planned using the sub-class ‘GR4 Mission’ which requires a resource template modelled as a tree of resource slots. At plan time specific resources are chosen to fill the resource slots tree, comprising in this simple example as a GR4, a pilot and a weapon. The mission itself is planned to follow a particular route ‘GR4 Route’.

Figure 1 Example GIM structure

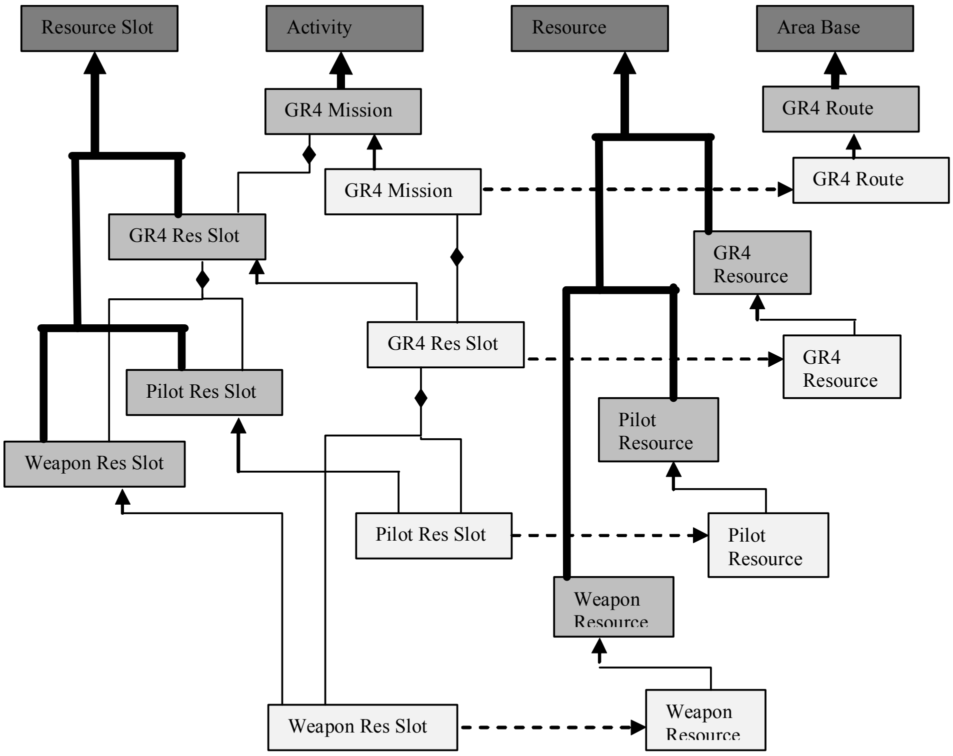

8.2.1.4 *Information dissemination – the generic approach*

The PACE framework is an event-driven publish/subscribe system, its GIM is optimised for object distribution through a delta mechanism technique which only pushes change. The resulting reduction in bandwidth footprint is a crucial advantage for many military situations.

The genericity inherent in the GIM is utilised in techniques used by the PACE framework when disseminating changes. The object distribution modules, which manage the delta mechanism, are implemented to work with GIM root classes and so can accommodate any configured classes without software change.

It is possible to finely tune the amount of information disseminated to any particular client and which users have permission to access or manipulate data using collections of objects called information sets (InfoSets), which are, effectively, the unit of subscription and the unit of permission.

8.2.1.5 Visualisation – the generic approach

Users can be given a visualisation of geospatial and temporal inter-plan relationships in the context of the whole plan through the use of generic visualisation tools that access the plans (from different warfare domains) kept in the common information layer.

9 Ontological deployment

The ontological deployment can be divided into two aspects: architectural and content. The architectural aspect looks at where the ontology is being deployed within the programme. The content looks at what is being deployed.

9.1 Deployment architectural aspects

Architecturally, the ontological work is deployed in two main areas: the configuration of the GIM and the configuration of the SIE.

In the case of the GIM, the ontological model that arises from the analysis is used to build the configuration of the GIM. In the case of the SIE, the BORO ontological analysis of the legacy systems produces a mapping from the legacy system data (and so the legacy system API) to the ontology. This can be used to configure the SIE, which will convert messages from the common PACE format to the legacy system format and back – as required.

9.2 Deployment content aspects

The hardest problem in systems integration is probably semantic interoperability – how one ensures that the messages between systems work to a common semantics, share sufficiently similar meaning to interoperate (Partridge and Stefanova, 2001; Partridge, 2002b, 2002d, 2002e, 2002g; Lycett and Partridge, 2009). Without some way of working with semantics, it is not possible to exchange and use data such that the meaning of the data sent by one system is sufficiently well understood by the receiving system that it can process it safely and properly. Currently, there are few rigorous methodologies for undertaking this task.

An alternative and much more rigorous approach is to use ontology to enable the necessary semantic understanding. This provides a robust solution and is the approach adopted for the J-TADIS programme. An ontological approach has benefits here as it works directly with the semantics. The BORO approach selected for this work has the additional advantage that the ontological analysis of the legacy system is undertaken in a way that provides the mapping between the systems.

9.2.1 The ontological approach to semantic integration

The ontological approach assumes that the data structures in the different legacy systems refer to a common 'real world' and that the purpose of the ontological analysis is to identify the 'real world' objects that are being referred to. Then, the mappings between the legacy systems are intermediated by these 'real world' objects. In practice, a mapping between the legacy system data structures and the ontology model is kept and used as the requirements specification[21].

9.2.2 Semantic architecture for AD legacy systems integration

A key semantic architecture decision for AD legacy system integration is whether to be 'point to point' or 'hub and spoke'. In a semantic point to point architecture, semantic mappings are made between the interfaces of individual systems as required. In a semantic hub and spoke architecture, each application has a semantic mapping from its interface to the hub semantics – and communications to other systems involve a mapping from the hub semantics to that system's interface. In general, a hub and spoke architecture requires fewer semantic mappings when there are more than three or four interfaces/systems – significantly less when there are significantly more systems. This suggests that AD systems integration is more suited to a hub and spoke semantic architecture. Furthermore, the ontological approach suits a hub and spoke approach.

Note that this is a semantic architecture and places no constraints on the physical network architecture which could still be point to point – in fact, PACE is implemented as a distributed network of information servers, there is no single central physical server. In this particular scenario, where the PACE/GIM data structure is configurable and the home for the collaborative planning data which requires the consolidation of the information from a range of legacy systems – PACE/GIM is a natural choice for a physical implementation of the semantic hub. Hence, the J-TADIS programme selected a semantic hub and spoke architecture, with the distributed PACE at its semantic hub.

In a hub and spoke architecture, the responsibility for conforming to the hub semantics can be either placed upon the spoke interface or integrated outside the spoke system. In the case of legacy systems, where an interface already exists, there is a good case for managing it outside: this avoids the need to re-open the development of each of the systems, typically from different suppliers, and enables the integration work to be consolidated within a single development project, taking advantage of economies of scale. Taking the ontological approach helps to consolidate these economies of scale. Hence, the J-TADIS programme developed a general SIE framework within which the individual mappings from spoke legacy system to hub were implemented.

9.3 Semantic interoperability layer – SIE

The SIE comprises parsers for the various military messaging [such as the ACO, ATO and common route definition (CRD)].The parsers use the mappings from the legacy systems to the common AD ontology produced by the ontological analysis to translate the messages to and from the common format used by PACE/GIM.

The eventual transformation results in GIM objects where the legacy system information is now in a form to be processed and visualised with information generated both within PACE and from other legacy systems.

For J-TADIS the SIE was implemented as a web service accessed via a PACE plug-in which sent it legacy system messages for conversion to GIM objects.

The general message translation process is shown in Figure 2. In the specific configuration used messages were routed through PACE and so one of the translations was trivial as it was configured with the common AD ontology. This was used to safely import, for example, ACOs and the ATOs into PACE's configurable GIM.

The SIE's configurable structure and production of the configuration from the ontological analysis of the legacy systems provide a framework for radically simplifying application interface semantic complexity. The approach also reduces the number of interfaces that would be required if a one to one integration approach was employed.

Figure 2 SIE message translation process

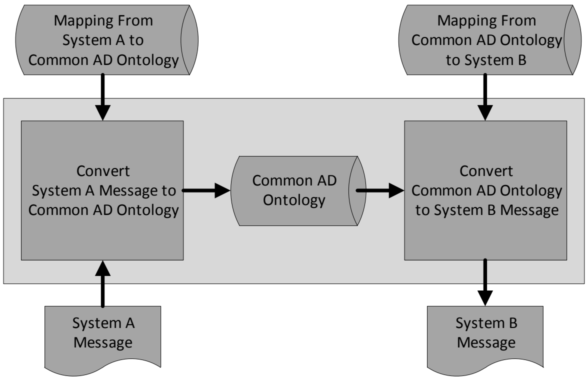

9.4 *Architecture of the final system*

Figure 3 shows the final J-TADIS architecture – and the legacy system, SIE and PACE layers.

Figure 3 J-TADIS the architecture

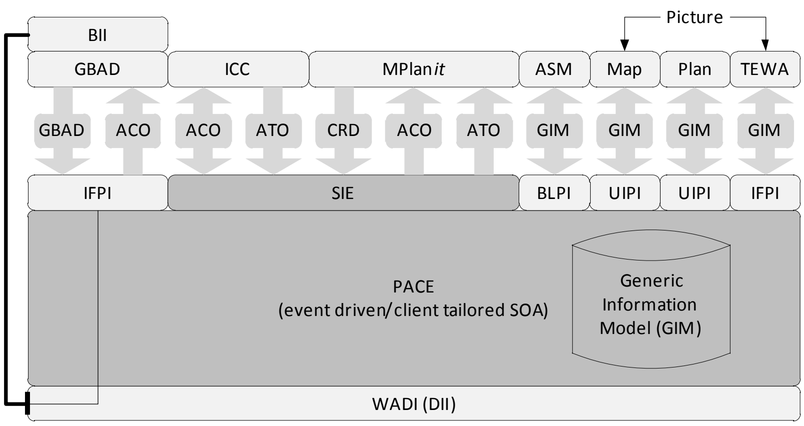

The PACE framework runs on the wide area distributed infrastructure (WADI) which is used to simulate the defence information infrastructure (DII) UK's main military network. Partners from the defence industry, BAE and Thales, provided their own operational systems to ensure that the final system was realistic in terms of the information to be shared, the type of plan consistency checking and the timeliness required. BAE provided their GBAD BISA system. The GBAD BISA is the operational C2 system for the ground-based AD assets. It runs on the battlefield information infrastructure (BII) over the BOWMAN communication bearers. Thales provided the MPlanIt mission planning system used operationally to generate low level detailed air mission plans. NATO ICC was used for airspace management as it is the operational tool used within the joint force air component commander (JFACC). Plans initially generated using operational systems were imported into the PACE framework though the SIE, manipulated during collaborative working sessions then exported again via the SIE to multiple legacy systems.

In further stages plans were modified collaboratively in real time as a Link16 tactical picture feed showed a change in the tactical situation. These modifications were immediately displayed on the legacy systems in both the JFACC command centres and land-based AD HQs.

10 Ontologically driven improvements

10.1 Examples of ontological improvements

The J-TADIS ontological analysis exhibited the practical qualities of a good information model. A significant number of underlying general patterns were discovered and this led to a corresponding simplification and reduction in size of the model. It also led to an increase in unity and explanation as the integrated structure of real world being represented was made clear. Within the scope of this paper it is not possible to show the consistent high level of semantic quality that was achieved across the domain[22].

The SMEs identified a number of surprising (to them) semantic issues that analysis revealed, which they felt were important enough to need communicating more widely. We include two examples below – amended to stay within the constraints imposed by security classifications. These provide a useful illustration of what the SMEs regarded as significant improvements.

10.2 A general context

The team found it useful to have a framework to classify the improvements and benefits. For this kind of work, it is important to provide a practical engineering rather than a theoretical scientific perspective. From a software systems perspective, the main benefits that an ontological approach aims to provide are [Partridge, (2002g), pp.4–5]:

- improved inter-operability
- reducing complexity
- increasing longevity
- technology proofing.

From this ontological modelling perspective, the qualities of a good model (ones that bring practical benefits) are [Partridge, (2002g), pp.22–26]:

- relevant precision and sufficient formality
- sufficient simplicity and relevant generality
- appropriate unity and explanation

- relevant fruitfulness
- relevant repeatability – re-usability.

Interestingly, these are generally recognised as the qualities that characterise good scientific theories.[23, 24]

10.2.1 Real world semantics driving relevant precision

Within the programme, there were a number of cases where insisting on a clear real world semantics under a well-defined top ontology led to a clear relevant increase in accuracy. One telling example was the analysis of exactly what the ACMs data in ACOs referred to.

Essentially, an ACM is a portion of airspace – one that is used for some purpose. Prior to the introduction of computers, these were described using simple geometric figures drawn on maps with standard ruler, compass and protractor and then associated with a minimum and maximum height – as shown in Figure 4.

Figure 4 Representing an ACM on a map

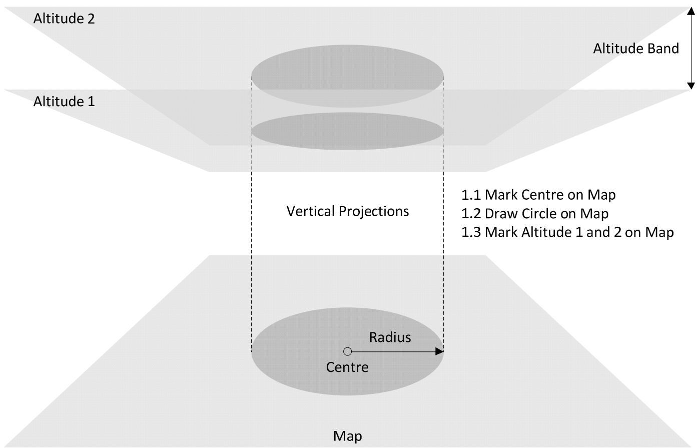

For example, a restricted operating zone (ROZ) around an artillery battery would be drawn as a circle on the map, with the battery as its centre and the radius based upon the battery's range. The minimum height would be ground level and the maximum height determined by the battery's range.

A standard set of figures was codified and these evolved into the data structures in the current messages – which provide the inputs[25] to an algorithm to calculate the extent of the airspace. Interestingly, the ontological analysis of these data structures, their associated documentation and existing working practices clearly showed their map-based roots. More importantly, the analysis also revealed a number of possible different algorithms and so interpretations of the airspace boundaries[26].

The map as a representation can be regarded as a Euclidean plane in which the constructed geometric figures are regular. However, when interpreting these figures in terms of the Earth's surface [or some idealisation of it – for example, World Geometric System (WGS) 84], one needs to invert the original projection (which could be of a variety of types) and this distorts the neat geometric figures drawn on the map. The current data structures have resolved this map projection issue by stipulating idealised reference ellipsoids (such as WGS 84) as their starting point.

However, we found that there are additional sources of semantic indeterminacy that have not been resolved. We concentrate upon one of these here – the algorithms for estimating the exact extent of the airspace when projecting the surface figures (of whatever shape) onto the minimum and maximum heights given. The analysis unearthed several possibilities, each of increasing accuracy. We developed an example using this ACM data:

Shape: circle-centre 24° 01’ 15”N, 055° 53’ 42”E, radius 15 km.

Height: minimum 15,000 ft, maximum 35,000 ft.

This data is shown graphically in Figure 5 (making some simplifying assumptions). This shows clearly that the data does not directly represent the ACM air space, rather it is the input to an algorithm that calculates where the ACM airspace is.

Figure 5 Graphical view of the ACM data

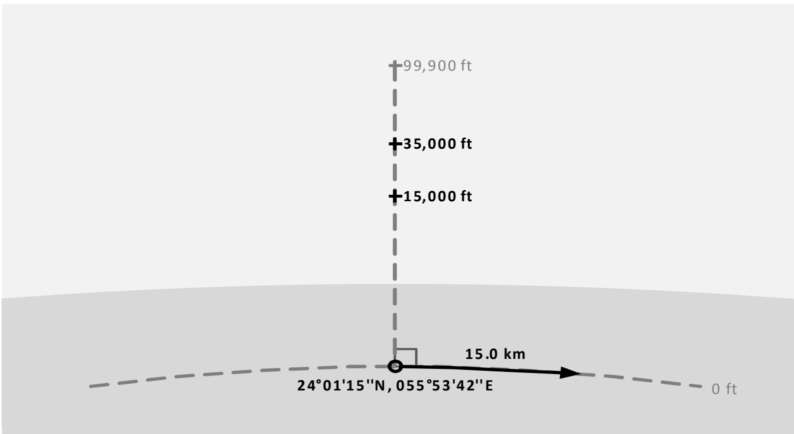

10.2.1.1 Lack of supporting documentation

The ontological analysis process requires the identification of the airspace being represented. One of the issues we faced was that there was nothing in the data structure or associated documentation that gave any hint of what algorithm we should use. Furthermore, the SMEs had not considered the issue and had no intuition of which was intended. However, when we proposed particular interpretations and found practical issues, the SMEs offered useful guidance. Lack of documentation is a very common issue when dealing with legacy systems, and an ontological approach, such as BORO, is an excellent way of compensating for this (Daga et al., 1995).

10.2.1.2 Cylindrical projection interpretation

When starting the analysis, we first assumed that the airspace is an extrusion of the surface circle, in other words, a cylinder. And that the minimum and maximum heights are cross-sections at right angles to the height axis. (To simplify the explanation, this description assumes the height axis is a normal to the idealised surface and not, for example, through a notional centre. It also assumes that the circle is drawn on a plane at a tangent to the idealised surface.). The resultant extruded cylinder and its airspace segment is shown in Figure 6.

Figure 6 Cylindrical extrusions

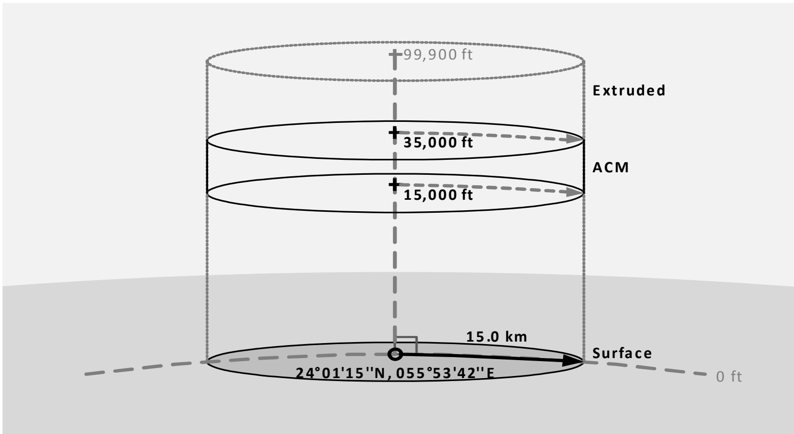

However, we soon realised that practical considerations suggested that this was not an ideal interpretation. If two surface circles whose edges touch are drawn on the surface, then there will be a gap between their extruded airspace segments, as their normal height axis (and so cylinders) will not be parallel (this applies to any contiguous figures) – see Figure 7. In practice, one needs airspaces that are contiguous to enable air traffic to pass from one airspace into another.

Figure 7 Incompletely contiguous cylindrical extrusions

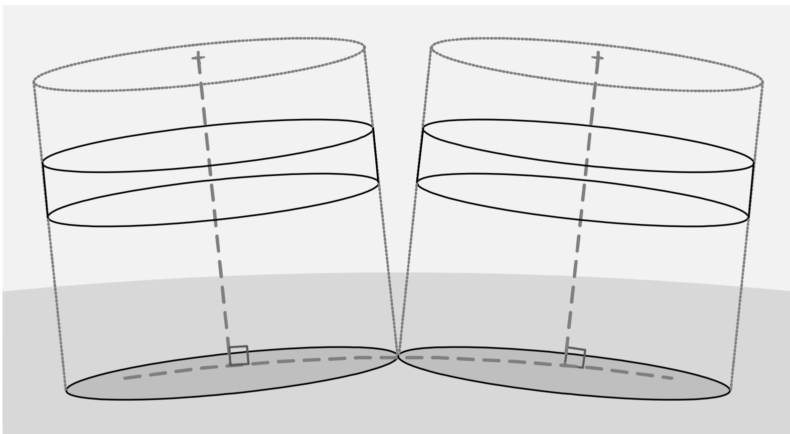

10.2.1.3 Conical projection

To resolve this, we considered an interpretation that extrudes the edges of the circles along normal at that point on the surface. As before, the minimum and maximum heights are cross-sections at right angles to the central height axis. This results in a conical shape shown in Figure 8. Under this interpretation, if two figures are contiguous at the surface, they are contiguous at all heights.

Figure 8 Conical extrusion

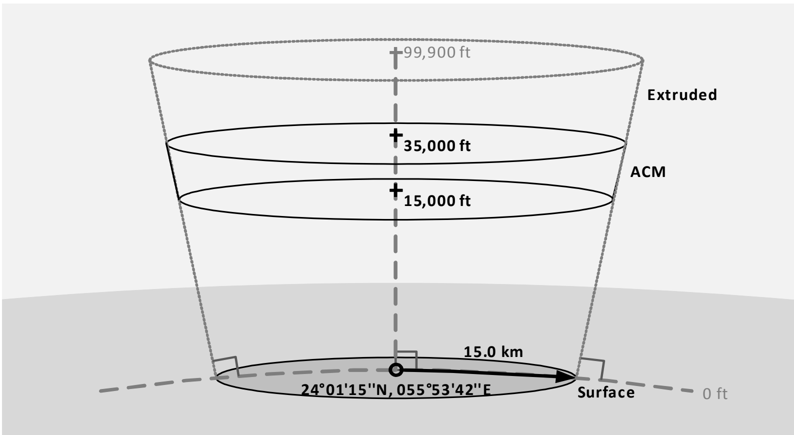

However, after further consideration, we realised that this interpretation also has practical issues. The height of the airspace varies as the height cross-section is not parallel to the (idealised) surface of the earth – as shown in Figure 9. If an aircraft were to fly along the lowest level of the airspace from the centre of one of two touching circles to the other centre, then it would have to gain height as it flew towards the edge of one circle and then lose height as it flew to the centre of the other – a less simple manoeuvre than just maintaining height. One would expect that it should be able to just maintain its height.

Figure 9 Abutting conical extrusions

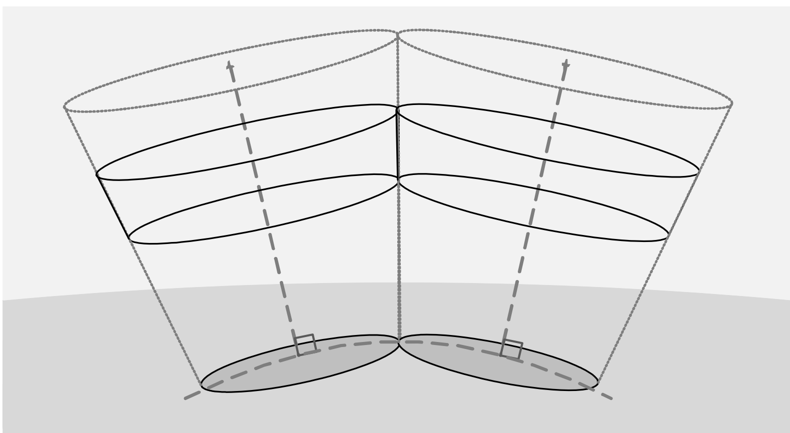

10.2.1.4 Elliptical segment

We resolved this by shifting our interpretation of the maximum and minimum heights. We considered a height to be an extruded surface, where at each point, the distance between the surface and the extruded surface was the same. Then, the top (or bottom) of the airspace was the intersection of the cone (considered above) and the relevant extruded height – illustrated in Figure 10. This ensured that the height at the top and bottom of the airspace remained consistent throughout the airspace.

Figure 10 Extruded ellipsoidal altitudes

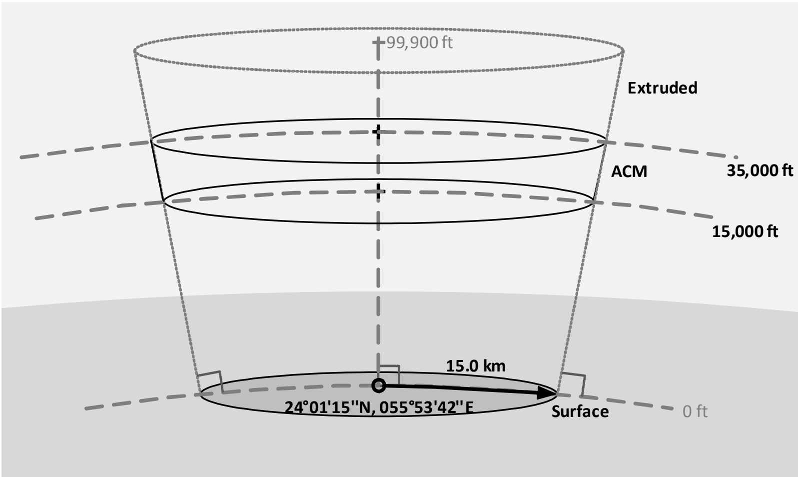

The issue is perhaps clearer if one considers an airspace that is sufficiently big such that the curvature of the idealised Earth has a significant effect.

10.2.2 Increases in accuracy

In most cases, the increases in accuracy here are relatively small, measured in tens of metres – as shown by calculations in Figure 11. One could argue the accuracy was (and, to some extent, still is) sufficient, but it is becoming more of an issue now and will increasingly become one in the future. As the management of airspace becomes more automated, there is a corresponding need for this management to be more reliable. Where there are a range of possible interpretations, it is likely that different equipment will use different interpretations. In this case, they are likely to show different results. For example, one threat assessment system will show an incursion into an airspace and the other will not – and there is no way of determining which is right. Where, as here, these differences are endemic then the trustworthiness of the system is undermined. Furthermore, with the introduction of new types of air vehicles, such as UAVs, it is likely that the required degree of accuracy will increase to avoid air collisions. This is not feasible without stipulating an interpretation.

Figure 11 Error between cone and cylinder projection

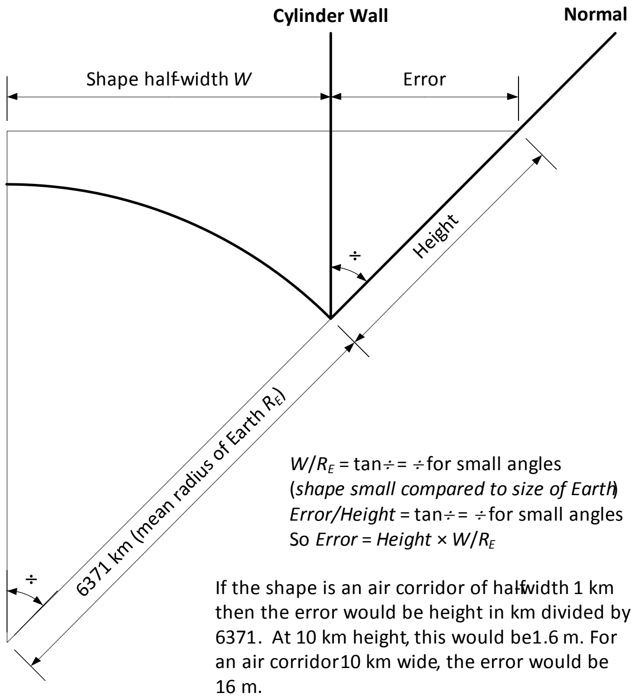

10.2.2.1 Ontologically vague intentional objects

These airspaces are intentionally constructed, and so the correct interpretation is determined by the intention of the relevant authority. We spent some time investigating what the intention was. As far as we could determine, there was no consideration of the issues raised here. As often happens, the data structures evolved without any detailed consideration of their accurate interpretation. So, as we understand it, there is currently no intentional support from the relevant authority for any of the interpretations. Hence, from a (theoretical) ontological perspective, it would seem that the airspaces are vague or indeterminate objects [under whatever account of vagueness or indeterminacy one holds (Williamson, 1994)]. However, we had a more practical perspective. The ontological analysis has identified this vagueness and there is now a need for the relevant authority to decide upon which interpretation they intend (or would have intended) that will provide a suitably accurate foundation for agile air systems.

One of the reasons this issue may not have been spotted earlier is that unless one takes an ontological stance, one can proceed without trying to directly identify the airspace. For example, when making a threat assessment, one only needs an algorithm that determines whether given the ACM input data and the threat position, whether the threat is inside the ACM airspace. The airspace does not need to be directly calculated. Quality assurance may then check that this algorithm is working properly. As only one algorithm with its implicit assumptions about the interpretation is being used, the vagueness will not be discovered. It is only when multiple systems, using different algorithms, are tested that an issue is likely to be found.

10.3 Example of fruitfulness

The use of a four-dimensional ontology threw up some interesting examples of fruitfulness. We developed an example that relates to the efficient use of airspace, which is described below.

In the ontology model, the shape of the airspace was separated from the other mechanisms of air control measure. This modularisation (separation of concerns) meant that the introduction of new airspace shapes, which happens from time to time, could be easily accommodated without changes elsewhere. Furthermore, the uses of a four-dimensional ontology mean that all the current shapes were captured from a four-dimensional perspective – and so facilitate the introduction of new four-dimensional shapes.

What is not immediately obvious to many users of the ACMs is the constraints that the current allowable airspace shapes have. Our analysis showed quite clearly their map/ruler/protractor origins – each shape is easily drawn using those instruments (even though they have not been used to draw the shapes for decades) and the current data structures reflect these origins. These shapes are two-dimensional in origin and extruded (as the discussion of interpretations above shows) to a third-dimension.

This is reflected in the data structures, where the surface 2D shape data is stored separately from their 1D height data. This segregation places constraints on the possible shapes. In the case of the battery ROZ, a spheroid shape that more accurately captured the range limits of the artillery would be better than a cylinder or cone – the two shapes are shown in outline in Figure 12 – where the savings in airspace are also clear. However, the 3D spheroid shape cannot be described within the 2D + 1D constraints of the current data structure.

Figure 12 3D shape in 2D extrusion

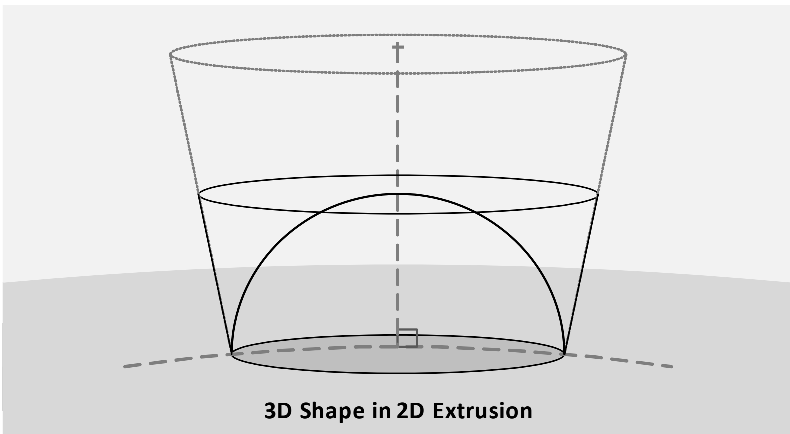

However, there is a further historical constraint that this 2D + 1D brings (and that a 3D approach does not solve), that is not so immediately obvious. The use of a 4D top ontology in the analysis encourages a 4D perspective, and from this perspective the constraint can be clearly seen. From this perspective, each of these 2D + 1D / 3D shapes can be seen as persisting unchanging though time. This happens because the space shape is described separately from its 1D time element – in data terms, it is described using a start and end time. In other words, the data structure has a 2D + 1D + 1D structure. This example shows a constraint upon the use of airspace that this 2D + 1D + 1D structure brings with it.

When reviewing the data sample, we came across entries in the text section of the message that described how airspaces were to be used (divided between) a number of airspace users. From an analysis perspective, the use of text fields often signifies a situation where the information cannot be fitted into the structured data. This turned out to be the case here. The typical situation was an airspace corridor where traffic was in one direction. Normally aircraft would book a time slot in the airspace and each aircraft uses its timeslot – as shown in the first space-time map in Figure 13. However, in this situation, it is a safe and more efficient use of airspace for aircraft to enter one end of the airspace before the previous aircraft had left the other end. The text advises the ad hoc rules for managing this. Clearly this is a workaround made necessary by the constraints inherent in the data structure.

Figure 13 More economic use of airspace

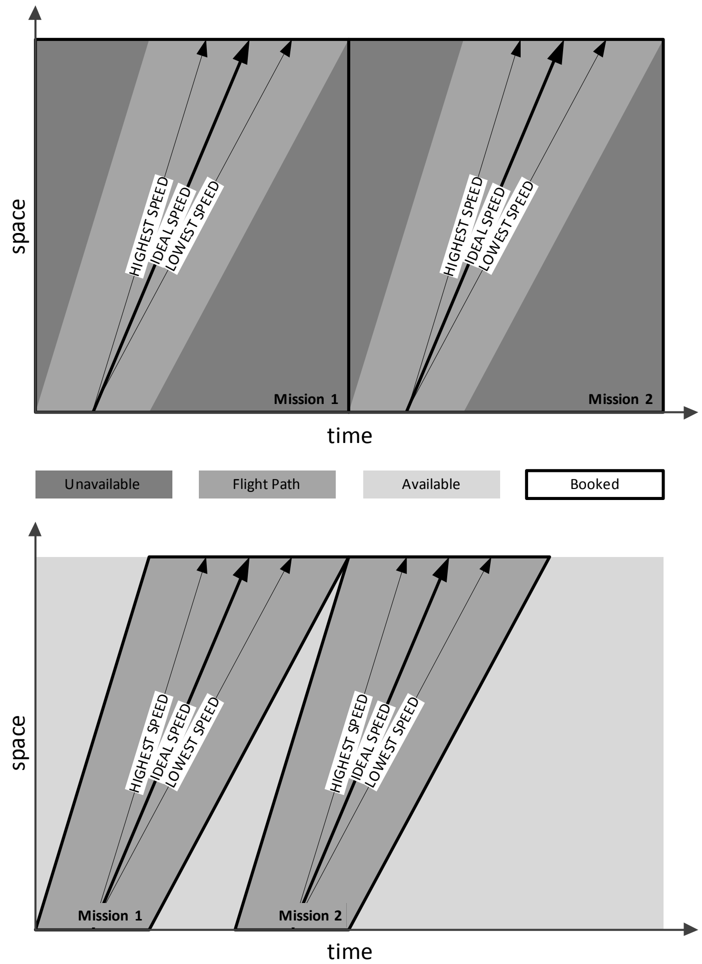

The constraint that is being worked around here is the separation of space and time – and only allowing booking in time slots. From a four-dimensional perspective, if one is booking airspace for an aircraft, it would make sense to book a shape that reflected the aircraft's flight path with a tolerance for any deviations – to consider its path in four-dimensions. The resulting airspace is shown in the second space-time map in Figure 13 – the increase in airspace utilisation is clearly visible. Note also that the shape from a 3D perspective (the 'space' axis in the diagram) changes over time – increasing in size (length) to reflect the tolerance needed to handle the potential variations in speed. As there is no persisting 3D shape – one needs to describe the shape in four-dimensions.

When the SMEs had time to reflect upon this, they found a number of situations where this kind of more efficient use of airspace would be useful. One example they provided was long flights – such as from Diego Garcia to Europe. These could not be planned using a single corridor, as then the whole airspace would need to be commandeered for the time of the whole flight. This is not practical when airspace is busy – as it is in Europe. Instead a large number of smaller spatially and temporally aligned corridors had to be booked. This turned out to be time consuming, especially during the planning phase where a change to one corridor had to be rolled out across the rest of the chain to keep them aligned.

10.4 The nature of SME expertise

The SMEs played a vital role in the ontological analysis. However, this turned out to be different in some key respects from the role they are traditionally expected to play; where they are expected to technically verify the final deliverable is appropriate for their domain. Implicit in this practice is the assumption that they have the expertise to make this validation. We found that the SMEs had great difficulty in making any kind of independent technical assessment of the ontology – this could only be carried out by a team including both SMEs and ontologists. Furthermore, the SMEs had great difficulty in articulating their knowledge in a form that could be directly represented and comprehending the representations when these were produced. It appeared that their expertise was not easily translated to or from the ontological representations – and that this was not due to the nature of the representations.

Our experience on this programme reinforced experiences on previous ontological analysis projects and enabled us to articulate the underlying issue. We have come to the view that the traditional approach is seriously flawed. Indeed, we think most experienced practitioners realise this and most successful projects only pay lip service to the practice. We have done some initial research to characterise the problem and document it here. This should help to make the case for more analysis.

Typically, an SME is an individual who exhibits the highest level of expertise in performing a specialised job, task, or skill within an organisation. However, expertise in performing a task is not the same as expertise in understanding and articulating that task – this is captured in the distinction between know-how (Polanyi, 1966) and know-that (or know-what) (Ryle, 1949). Indeed, becoming an expert may involve letting go the conscious understanding of what one is doing. Searle (1983, 1995) describes how becoming expert in a task one moves from conscious control to unconscious action, where one has no access to a picture in one's mind of what one is doing. Searle calls this the ‘background’. One key aspect of Searle’s analysis is that the more expertise one has, the less one has an internal representation of that expertise (or conscious access to that representation). This accorded well with our experience on the programme.

Further evidence for this analysis comes from situations where experts need to provide an explanation of their expertise. In some case, they, post hoc, rationalise one. As the expert has no access to his/her tacit knowledge, there is no guarantee that this rationalisation will be correct. Shaffer and McBeath (2005) provide a good example: where expert baseball players provide a completely false rationalisation of how they catch a fly ball.

This characterisation of expert knowledge has implications for the use of SMEs during analysis. It implies that they are not a good source for determining what exists in the domain when aiming for a fine grained ontological model and that they are not best equipped to independently technically verify the final deliverable as a good picture of the domain. The methodology needs to harness their expertise in more appropriate ways.

This also has implications for the nature of the analysis. When the SME is regarded as having a picture (representation) of the domain in his/her head, the analysis consists in extracting the details of the picture. However, if no such picture exists (or it is not accessible), then a different approach is required (Partridge, 2001, 2002a, 2002d, 2002f, 2002g). The analysis then becomes a kind of ‘rational reconstruction’[27] of the model that would have been constructed if such explicit knowledge was available from the evidence provided by the SME and other sources. A key element in the analysis is ‘inference to the best explanation’[28] or perhaps more accurately ‘inference to the best interpretation’.

11 Further work

As often happens with novel work, the implementation of the solution suggested a number of further areas of work.

11.1 PACE and SIE frameworks

For non-generic and non-configurable systems the cost of creating distribution and persistence infrastructure is a function of the number of message types to be distributed. For generic configurable systems, the work is primarily concerned with the configuration process. This had a large manual element in the PACE/SIE implementation. In future research, it will be worth looking at how this could be automated.

11.2 The generality of the solution

As far we could tell, there is nothing in the nature of the solution that is specific to AD, or indeed the military environment. It should be possible to apply the approach to any domain where there are a significant number of legacy systems that need to be integrated. The team have been involved in using elements of the solution in a variety of projects, so have confidence in it working elsewhere, it would be interesting to see how generally the specific PACE and SIE solutions developed here could be applied.

11.3 Coding reduction impact of the general patterns in the ontology

One of the software engineering benefits of the BORO ontological analysis is the identification of general patterns. Several members of the development team noted that this led to a significant reduction in code[29]. Further work could be done to confirm that this happens and investigate how this works, the factors that affect it and the potential scale of reduction.

12 Conclusions – what did the research show?

The UK MoD asked for a demonstration at TRL6 of a novel cost-effective way to improve the agility of AD-C2 legacy systems. This was demonstrated using 'last minute' CP&E as the test case. The two interlinked key factors in this were the configurable PACE and SIE application architecture and a novel ontological analysis that together allowed rapid, semantically assured, information dissemination.

From a systems development perspective, the research demonstrated the symbiosis of the ontological approach and system configurability in a situation where data structures can evolve. The ontological analysis provided the configuration; the systems provided the capability to consume it. It also demonstrated the benefits of a configurable application in an environment where the data structures are evolving.

From an interoperability perspective, the research demonstrated the feasibility of building a system of legacy AD-C2 systems with a level of semantic quality sufficient to support challenging applications such as 'last minute' CP&E.

From an ontological perspective, the research demonstrated an ontological approach that could mine a collection of legacy systems for a common AD ontology and produce mappings from the legacy system data structures to this ontology with a consistently high level of semantic quality assurance. Also, that the programme highlighted the usefulness of an ontological realism as a strong basis for producing an integrated picture across the range of legacy systems.

Acknowledgements

We are grateful to Cdr Jerry Woods RN, the MoD sponsor for this programme, for his substantial help and support. We are grateful to Bob Curtis (QinetiQ), our prime SME, for the invaluable help he gave us. We are grateful for the support of our collaborators, BAE Systems and Thales. We are grateful to QinetiQ and BORO Solutions for allowing us to use material from their reports.

References

Carnap, R. (1928) *Der Logische Aufbau Der Welt*, Weltkreis-verlag, Schlachtensee, Berlin.

Daga, A. et al. (2005) 'An ontological approach for recovering legacy business content', Daga, A., de Cesare, S., Lycett, M. and Partridge, C. (Eds.): *HICSS'05 BURA.*

Dijkstra, E.W. (1982) 'On the role of scientific thought', in Dijkstra, E.W. (Ed.): *Selected Writings on Computing: A Personal Perspective*, pp.60–66, Springer-Verlag New York, Inc., New York, NY, USA, ISBN 0-387-90652-5.

Frege, G. (1892) *Über Sinn und Bedeutung*, in Zeitschrift für Philosophie und philosophische Kritik, Vol. 100, pp.25–50, translated as *On Sense and Reference*, by M. Black in translations from the philosophical writings of Frege, G., Geach, P. and Black, M. (Eds. and Trans.), 3rd ed., 1980, Blackwell, Oxford.

Habermas, J. (1998) *On the Pragmatics of Communication*, The MIT Press, Massachusetts, USA.

Kuhn, T. (1962) *The Structure of Scientific Revolutions*, University of Chicago Press, Chicago, ISBN 0-226-45808-3.

Kuhn, T. (1977) *Objectivity, Value Judgment, and Theory Choice in the Essential Tension: Selected Studies in Scientific Tradition and Change*, pp.320–239, University of Chicago Press, Chicago, USA.

Kuśnierczyk, W. (2006) 'Nontological engineering', *Proceedings of the 2006 Conference on Formal Ontology in Information Systems: Proceedings of the Fourth International Conference* (*FOIS 2006*) ISBN: 1-58603-685-8.

Lipton, P. (1991) *Inference to the Best Explanation*, Routledge, London.

Lycett, M. and Partridge, C. (2009) 'The challenge of epistemic divergence in IS development', *Communications of the ACM*, June, Vol. 52, No. 6, pp.127–131.

McCarthy, J. (1980) 'Circumscription – a form of non-monotonic reasoning', *Artificial Intelligence*, April, Vol. 13, Nos. 1/2, pp.27–39.

Partridge, C. (1996) *Business Objects: Re-engineering for Re-use*, Butterworth Heinemann, Oxford.

Partridge, C. (2002a) *LADSEB-CNR – Technical Report 04/02 – What is Pump Facility PF101?*, Padova, LADSEB-CNR, Italy.

Partridge, C. (2002b) *LADSEB-CNR – Technical Report 05/02 – The Role of Ontology in Integrating Semantically Heterogeneous Databases*, Padova, LADSEB-CNR, Italy.

Partridge, C. (2002c) *LADSEB-CNR – Technical Report 06/02 – Note: A Couple of Meta-Ontological Choices for Ontological Architectures*, Padova, LADSEB-CNR, Italy.

Partridge, C. (2002d) *LADSEB-CNR – Technical Report 08/02 – STPO – A Synthesis of a TOVE Persons Ontology*, Padova, LADSEB-CNR, Italy.

Partridge, C. (2002e) 'The role of ontology in semantic integration', *Second International Workshop on Semantics of Enterprise Integration*, OOPSLA 2002, Seattle.

Partridge, C. (2002f) 'What is a customer? The beginnings of a reference ontology for customer', *11th OOPSLA Workshop on Behavioral Semantics*, Seattle, Washington, Northeastern.

Partridge, C. (2002g) 'The CEO project: an introduction', *LADSEB-CNR – Technical Report 07/02*, Padova, LADSEB-CNR, Italy.

Partridge, C. and Stefanova, M. (2001) 'A synthesis of state of the art enterprise ontologies: lessons learned', D'Atri, A., Solvberg, A. and Willcocks, L. and Luiss Edizioni (Eds.): *Open Enterprise Solutions: Systems, Experiences, and Organizations* (*OES-SEO 2001*), pp.130–133, Centro di Ricerca sui Sistemi Informativi, Rome.

Partridge, C. (2005) *Business Objects: Re-engineering for Re-use*, 2nd ed., BORO Centre, London.

Peirce, C.S. (1906) 'Prolegomena to an apology for pragmaticism', *The Monist*, October, Vol. 16, No. 4, pp.492–546, The Open Court Publishing Co., Chicago, IL.

Polanyi, M. (Ed.) (1966) *The Tacit Dimension*, Routledge, London.

Quine, W.V. (1948) 'On what there is', *Review of Metaphysics*, Vol. 2, No. 5, reprinted in from a logical point of view (1961).

Russell, B. (1905) 'On denoting', *Mind*, Vol. 14, pp.479–493.

Ryle, G. (1949) *The Concept of Mind*, Hutchison & Company, London.

Searle, J. (1983) *Intentionality: An Essay in the Philosophy of Mind*, Cambridge University Press, Cambridge.

Searle, J. (1995) *The Construction of Social Reality*, Free Press, New York.

Shaffer, D.M. and McBeath, M.K. (Eds.) (2005) 'Beliefs in baseball: systematic distortion in perceived time of apex for fly balls', *Journal of Experimental Psychology. Learning, Memory, and Cognition*, Vol. 31, No. 6, pp.1492–1501.

Smith, B. and Ceusters, W. (Eds.) (2010) 'Ontological realism: a methodology for coordinated evolution of scientific ontologies', *Applied Ontology*, Vol. 5, Nos. 3–4, pp.139–188, IOS Press.

Smith, B. Ceusters, W. and Temmerman, R. (Eds.) (2005) *Wüsteria. Proceedings Medical Informatics Europe 2005*, Geneva, and *Stud Health Technol Inform*, Vol. 116, pp.647–652.

Williamson, T. (1994) *Vagueness for One Recent Account of Vagueness*, Routledge, London.

Notes

1 Joint Force involves one or more of land, sea and air forces.

2 These typically involved techniques such as infrequent fixed format messaging or database replication.

3 www.QinetiQ.com – a leading defence technology and security company.

4 Where this is defined as the ability of computer systems to communicate information and have that information properly interpreted by the receiving system in the same sense as intended by the transmitting system.

5 AAP-6 NATO Glossary of Terms. 2009

6 In the AD context an effector is any system that can be used to 'effect' the enemy, an example is a weapon system.

7 See TRL Definitions.pdf in http://www.aof.mod.uk/aofcontent/tactical/techman/content/trl_applying.htm.

8 E.J. Lowe in the Oxford Companion to Philosophy.

9 Barry Smith also makes this point, most recently in Smith and Ceusters (2010) – note the criticisms of the 'The concept orientation' in this paper.

10 Philosophers make the same point about the true logical form beneath the 'surface grammar' of natural language – see Bertrand Russell's Theory of Descriptions in (Russell, 1905).

11 For a survey of approaches to reengineering legacy systems see Daga et al. (2005).

12 Partridge (2002c) has some relevant examples. Most introductory textbooks on metaphysics will also provide a number of examples, though these may not all be of engineering interest.

13 See Smith and Ceusters (2010) for a robust defence of a similar position. See also Smith et al. (2005) for an attack on an alternative position.

14 The Department of Defence Architecture Framework (DODAF) official version can be found here: http://cio-nii.defence.gov/sites/dodaf20/index.html.

15 These opposing positions can be seen as rational and empirical. Where the rationalist assumes that he/she can arrive at the answer by speculation; whereas the empiricist believes that she/he needs to work by observation.

16 For more details on ESB, see David Chappell, 'Enterprise Service Bus' (O'Reilly: June 2004, ISBN 0-596-00675-6).

17 The official CORBA standard from the Object Management Group can be found here: http://www.omg.org/spec/CORBA/3.1/.

18 The official catalog of the Object Management Group's DDS specifications can be found here: http://www.omg.org/technology/documents/dds_spec_catalog.htm.

19 The Extensible Markup Language (XML) specification can be found here: http://www.w3.org/TR/REC-xml/.

20 The Tornado GR4 is a variable geometry, two-seat, day or night, all-weather attack aircraft, capable of carrying a wide variety of weapons.

21 See Partridge (2005) Figure 11.4: Reengineering our conceptual patterns for picture of this mapping.

22 In addition, the contents of the ontology are classified and so cannot be made publicly available.

23 See the similar list in the penultimate chapter of Kuhn (1962) and Kuhn (1977) – where, interestingly, he says "I am suggesting, of course, that the criteria of choice with which I began function not as rules, which determine choice, but as values, which influence."

24 The BORO ontological analysis was designed as a systematic process that provides a level of assurance that these semantic qualities have been scrutinised – providing a form of semantic quality assurance.

25 Interestingly, these seem to be a good example of Fregean senses (see Frege, 1892) – a way to identify the object.

26 As the algorithm is not made explicit anywhere in the documentation – or known to the SMEs – there is no real basis to determine which was originally intended. A reasonable conclusion would be that this was left vague.

27 For more details of rational construction see, for example, Carnap (1928) and Habermas, J. (1998) who says that "we can distinguish between know-how, the ability of a competent subject who understands how to produce or accomplish something, and know-that, the explicit knowledge of how it is that he is able to do so" and

"…[In rational reconstruction] the distinction between drawing on a-priori knowledge and drawing on a-posteriori knowledge becomes blurred. On the one hand, the rule consciousness [i.e. intuitive know-how] of competent subjects is for them an a-priori knowledge; on the other hand, the reconstruction of this calls for inquiries undertaken with empirical [methods]".

28 See, for example, Lipton (1991) for more details. See also McCarthy's (1980) discussion of circumscription. Earlier Peirce (1906) called this abduction – saying that "Long before I first classed abduction as an inference it was recognised by logicians that the operation of adopting an explanatory hypothesis – which is just what abduction is – was subject to certain conditions. Namely, the hypothesis cannot be admitted, even as a hypothesis, unless it be supposed that it would account for the facts or some of them." He also more light-heartedly said "Abduction is no more nor less than guessing".

29 See Partridge (2005) Ch.6, Section 3 – An environment that encourages compacting for more background details.