An Ontology-Based Approach to Sensor-Mission Assignment

Alun Preece, Mario Gomez,
Geeth de Mel, Wamberto Vasconcelos
and Derek Sleeman
Computing Science Department
University of Aberdeen, UK
http://www.csd.abdn.ac.uk/people

Stuart Colley
Defence Science and
Technology Laboratory
(DSTL), UK

Tom La Porta
Computer Science and Engineering
Penn State University, US
http://www.cse.psu.edu/∼tlp

***Abstract*—Effective deployment of limited and constrained intelligence, surveillance and reconnaisance (ISR) resources is seen as a key issue in modern network-centric joint-forces operations. The aim of our work is to enable proactive and reactive deployment of sensors and other information sources to best support the objectives of a task (or mission) being undertaken. In this paper, we consider one aspect of the deployment problem: proactive assignment of sensors and sources to mission tasks. We view this sub-problem as a matchmaking activity: matching the ISR requirements of tasks to the ISR-providing capabilities of available sensors and sources, and the platforms that carry them. A key issue is that of defining sufficiently-rich representations of these various elements — missions, tasks, ISR requirements, ISR capabilities, sensors, sources, and platforms — to support the matchmaking activity. We argue for an approach based on the use of ontologies: formal models of the various elements that can be used with deductive reasoning mechanisms to produce matches that are logically sound. We introduce a new ontology based on the military Missions and Means Framework (MMF), and show that the matchmaking activity is necessarily multidimensional in nature. We indicate how our approach builds on previous work in representing sensors and sources for various purposes, and highlight the role of current Web standards in providing an engineering foundation for our approach.**

I. Introduction and motivation

The work described in this paper is part of the International Technology Alliance project "Task-Oriented Deployment of Sensor Data Infrastructures". The overall aim of this project is to enable proactive and reactive deployment of sensors and other information sources to best support the objectives of a task (or mission) being undertaken. Effective deployment of limited and constrained intelligence, surveillance and reconnaisance (ISR) resources is seen as a key issue in modern network-centric joint-forces operations. For example, the 2004 report *JP 2-01 Joint and National Intelligence Support to Military Operations* states the problem in the following terms: "ISR resources are typically in high demand and requirements usually exceed platform capabilities and inventory. . . . The foremost challenge of collection management is to maximize the effectiveness of limited collection resources within the time constraints imposed by operational requirements."[1]

[1]http://www.dtic.mil/doctrine/jel/new_pubs/jp2_01print.pdf, pages III–10–11, accessed April 27, 2007.

In this paper, we consider one aspect of the deployment problem: proactive assignment of ISR assets to mission tasks. We view this sub-problem as a *matchmaking* activity: matching the ISR requirements of tasks to the ISR-providing capabilities of available assets (sensors and the platforms that carry them) and information sources (e.g. human beings and mass media). The key issue considered in this paper is that of defining sufficiently-expressive representations of these various elements — missions, tasks, ISR requirements, ISR capabilities and ISR assets — to support the matchmaking activity. We argue for an approach based on the use of *ontologies*: formal models of the various elements that can be used with deductive reasoning mechanisms to produce matches that are logically sound. The goal of this paper is to show that such an approach is plausible: that it can be grounded in current best-practice in analysing mission requirements and means, that it can draw on much existing work in representing sensors and platforms for various purposes, and that current Web standards for ontology engineering provide a suitable foundation for such an approach.

II. Semantic Matchmaking of Sensors and Missions

The assignment of ISR assets to multiple competing missions can be seen as a process comprising two main steps:

1) Assessing the fitness for purpose of alternative ISR means to accomplish a mission
2) Allocating available assets to multiple competing missions

We have analyzed current military doctrine to site our work, and as result we have found a framework that has been particularly inspiring: the Missions and Means Framework (MMF) [11] . MMF provides a model for explicitly specifying a military mission and quantitatively evaluating the mission utility of alternative warfighting solutions – the means.

Figure 1 shows a high level diagram of our current thinking of how missions map through to ISR means, based on MMF. Starting from the top left the diagram sketches the analysis of a mission as a top-down process that breaks a mission into a collection of operations (e.g. Search and Rescue), each of which is broken down further into a collection of distinct tasks having specific capability requirements (e.g. Wide Area Surveillance). On the right hand side, the diagram depicts the analysis of capabilities as a bottom-up process that builds up from elementary components (e.g. EO/IR camera) into systems (e.g. camera turret), and from systems up into platforms equipped with or carrying those systems (e.g. a UAV).

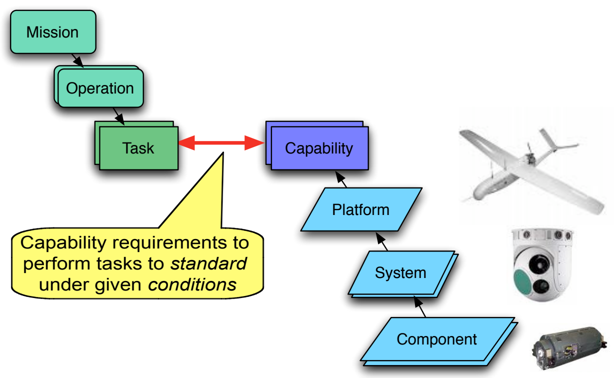

Fig. 1. Mission and Means Framework

The way MMF describes the linking between missions and means naturally fits the notion of matchmaking. Matchmaking is basically the process of discovering, based on a given request (e.g. ISR requirements), promising partners/resources (e.g. sensors) for some kind of purpose (e.g. accomplishing a mission). Important issues arise when the search is not limited to identity matches but, as in real life, when the objective is finding partners/resources suitable at least to some extent, or (when a single partner cannot fulfil the request) to find a pool of cooperating partners (a sensor network, or a platform equipped with several sensors) able to accomplish it. As this process may lead to various possible matches, the notion of ranking becomes central: to provide a list of potential partners ordered according to some criteria. Due to the diversity of frameworks of application, several communities have studied matchmaking through different perspectives and techniques. Recently, semantic matchmaking, which is based on the use of ontologies to specify components, has become a central topic of research in many communities, including Multi-Agent Systems, Web Services and Grid Computing.

In this paper we advocate the use of a semantic matchmaking approach to address the problem of assessing the fitness for purpose of alternative ISR means, which in turn supports the effective allocation of available ISR assets to multiple competing missions. In particular, we propose the use of ontologies as an expressive and logically-sounded way of doing knowledge representation and reasoning. More specifically, our approach uses ontologies in the following activities:

- Specify the requirements of a mission.
- Specify the capabilities provided by ISR assets (sensors, platforms and other sources of intelligence, such as human beings).
- Compare the specification of a mission against the specification of available assets to assess the utility or *fitness for purpose* of available assets; based on these assessments, obtain a set of recommended assets for the mission: either decide whether there is a solution —a single asset or combination of assets— that satisfies the requirements of the mission, or alternatively provide a ranking of solutions according to their relative degree of utility.

There are several approaches to perform matchmaking that do not make use of ontologies, and a natural question is, "Why use ontologies?" The next subsection gives a brief account of what an ontology is and a short motivation for using them to perform matchmaking.

A. *Ontologies for matchmaking*

There is no universally-agreed definition for the term ontology, probably because ontologies have been addressed in a number of contexts and fields, but there are some definitions and authors that are typically referred to. The most cited definition in Computing Science is probably Gruber's: "an ontology is an specification of a conceptualization" [6]. This definition was later modified and extended by Borst: "an ontology is a formal specification of a shared conceptualization" [2]. Borst's definition covers both the nature and purpose of an ontology: it is a formal specification, since it aims at being machine-processable; and it refers to a shared conceptualization, since it aims at mediating among different people and systems. In this paper, we use more precise notion that is fairly well captured by Guarino: "an ontology is a set of logical axioms designed to account for the intended meaning of a vocabulary" [7].

People, organizations and software systems need to communicate and share information, but due to different needs and background contexts, there can be widely varying viewpoints and assumptions regarding what essentially the subject matter is. The lack of shared understanding leads to poor communication between people and their organizations, severely limits systems interoperability and reduces the potential for reuse and sharing. Ontologies aim at solving the former problems. On the one hand, ontologies facilitate communication and knowledge sharing by providing a unifying framework for parties with different viewpoints. On the other hand, ontologies can improve interoperation and cooperation by providing unambiguous semantics in a formal, machine-interpretable way. Matchmaking can benefit from these general properties as far as the elements of the process are distributed or there are several viewpoints; additionally, the use of semantically rich specifications enable the use of specific forms of reasoning that are not available when using a syntactic approach, such as for example subsumption and disjunction. Below we provide a simple motivating example to illustrate on such forms of reasoning.

Figure 2 depicts a partial classification of military Unmanned Aerial Vehicles (UAV). The image shows 6 concepts referring to common UAV categories and specialization ("subclass") relationships among them, represented by arrows. At the top of the classification, the UAV concept encompasses

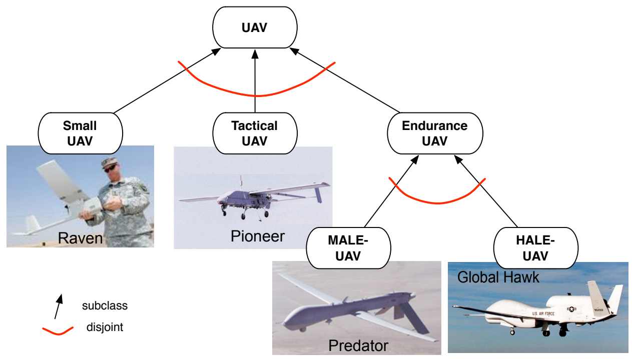

Fig. 2. Partial Classification of Unmanned Aerial Vehicles (UAV)

all kinds of UAV, which may range in cost from a few thousand dollars to tens of millions of dollars, and ranging in capability from Micro Air Vehicles (MAV) weighing less than one pound to aircrafts weighing over 40,000 pounds. In this example we include just three categories that are specializations of the UAV class; these are, from left to right: Small UAV (SUAV), designed to perform "over-the-hill" and "around-the-corner" reconnaissance; Tactical UAV (TUAV), which focuses on the close battle, providing targeting, situation development and battle damage assessment in direct response to the brigade/Task Force commander; and Endurance UAV, aimed at the deep battle, supporting the division to 150 Km and the Corps battle to 300 Km. Further, we have included two categories that specialize the Endurance UAV concept: Medium Altitude Long Endurance (MALE) UAV, designed to operate at altitudes between 5000 and 25000 feet, and High Altitude Long Endurance (HALE) UAV, which are designed to function as Low Earth Orbit satellites. The arcs between subclass relationships indicate a disjoint relationship among subclasses; a disjoint relation among a set of classes entails that an individual cannot belong to more than one of those classes; for example, a UAV that is classified as a Small UAV, can not be classified as being a Tactical UAV. Next, we introduce some basic examples illustrating specific forms of reasoning enabled by the use of ontologies. Let us suppose that we have the following UAVs available for a mission:

- A *Pioneer*, which is a TUAV
- A *Predator*, which is a MALE-UAV
- A *Global Hawk*, which is a HALE-UAV

Now suppose that as part of a given mission a Persistent Surveillance task over a wide area is required to detect any suspicious movement. This kind of tasks is best served by an Endurance-UAV, since it is able to fly for long periods of time. From just the concept definitions we know that: (1) the *Pioneer* is not an endurance UAV (because of the disjoint relationship among Endurance-UAV and TUAV), and (2) both the *Predator* and the *Global Hawk* are Endurance-UAVs (because of the subclass relationships)[2]. Therefore, the matchmaking process will select both the *Predator* and the *Global Hawk* as the assets satisfying the specified mission requirements.

Now, suppose that according to the weather forecast, storms are very likely to occur in the area of operations during the surveillance period. Then, the best option would be to use a HALE-UAV, which has the capability of flying "above the weather". Consequently, the matchmaking process would retrieve the Global Hawk as the only asset satisfying the mission requirements.

The UAV examples introduced above refer to a simple form of matching relationships known as subsumption, but it is possible to devise more complex information containment relationships and even an ordinal ranking scale comprising several degrees of matching. Figure 3 represents graphically the main types of matching that can be established using information containment relationships, using examples from the ISR domain. Q denotes a query which specifies some intelligence requirements to be met, and $S1 - S5$ denote the specification of ISR assets (sensors and sensor platforms) to be matched against Q.

Commencing at the left, our query specifies two basic requirements to be met:

- Provide Infrared (IR) Imagery
- Carry out a Night Reconnaissance task

Going from left to right and up to down, the figure shows the specification for several assets that verify different types of relation in terms of information containment. Below follows a description of these matching relations listed in decreasing strength order:

1) *ExactMatch(S1, Q)*: holds when the specification of a component provides exactly the same type of information described by the query. In the example, $S1$ describes an asset that provides IR vision and is designed to perform night reconnaissance tasks, just as stated in Q. This is represented as $S1 = Q$.
2) *Plugin(S2, Q)*: holds when the class of information described by the query subsumes (i.e. is more general than) the class of information specified by the component. In the example, the asset described by *S2* refers to a Cooled FLIR (forward looking IR), which is a specific type of IR camera. This is represented as $S2 \subseteq Q$.
3) *Subsumes (S3, Q)*: holds when the class of information described by the query is subsumed by the specification of the component, i.e. when the specification of the component is more general than the query. In the example, $S3$ refers to an asset providing night vision capability, which is a more general concept than Infrared Vision, and provides also night reconnaissance. This is represented as $S3 \supseteq Q$.
4) *Overlaps(S4, Q)*: holds when the query and the specification share some information, but neither one subsumes

[2]Note that we only state minimum explicit information about the UAVs (e.g. the Pioneer is-a Tactical-UAV); everything else is inferred from the concept definitions (e.g. the Pioneer is not a HALE-UAV).

the other entirely. In our example, $S4$ describes an asset that provides night reconnaissance as required by Q, but the first requirement is not satisfied, since it carries a radar (SAR stands for Synthetic Aperture Radar) instead of an IR camera, and these two concepts are disjoint. This is represented as $S4 \cap Q$.

5) *Disjoint(S4, Q)*: holds when there is no degree of information containment between the specification of the component and the query. In the example, $S5$ describes an asset that provides TV video and is suited to perform day reconnaissance tasks; radar imagery is disjoint with IR vision, day reconnaissance is disjoint with night reconnaisance, so there is no intersection or information containment between the concepts. This is represented as $S4 \perp Q$

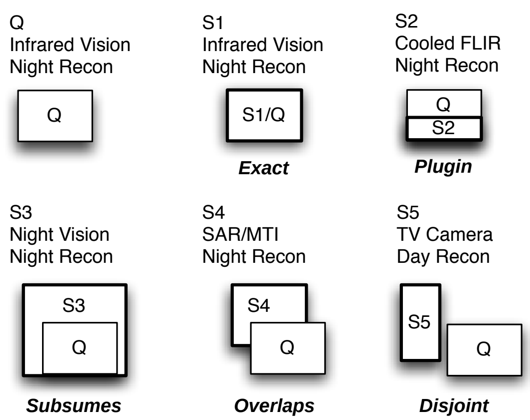

Fig. 3. Basic Matching Relationships

The kind of matching relationships introduced above are typically used to discover software components or services satisfying some specific requirements. Herein we propose to use these kinds of matching relations to discover ISR assets satisfying specific intelligence requirements.

B. Matchmaking abstract architecture

A matchmaking application is not completely characterized by the basic semantic relationships that can be established among the concepts that represent the application domain. An important issue of a matchmaking application is the distinction between the attribute-level and the component-level: a component may be described by different attributes, and so different matching schemas could be applied to each attribute depending on the particular role it plays within the component.

In our application, we have detected two main kinds of components to be matched against the ISR requirements of a mission, each one characterized by different attributes that deserve a separate treatment. Note that the kind of capability requirements that are relevant to select a specific kind of sensor are quite different from the requirements that are relevant to select a platform. For example, in order to assess the utility of different sensors it is very important to consider the kind of intelligence to be produced (IMINT, MASINT, SIGINT[3], etc.), since each type of sensors provide information that supports a different kind of intelligence (eg. infrared cameras support IMINT, while acoustic sensors support MASINT). Besides, to select a specific UAV for a reconnaissance mission there are other factors to consider, such as the range to the targets of interest, the presence or absence of enemy anti-air assets, and so on. In addition, UAVs are limited in the weight and type of sensors they can carry, and the performance of some sensors may be influenced by conditions that depend on the platform they are attached to, such as the altitude. Therefore, one cannot select UAVs and sensors independently; instead, the interaction between these components must also be taken into account.

To address the issues above, we propose an abstract architecture based on three types of components and three kinds of matching relations, as showed in Figure 4.

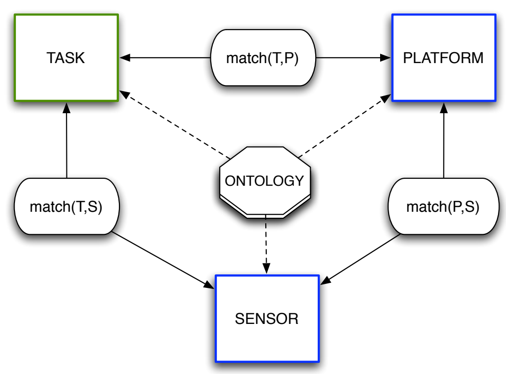

Fig. 4. Abstract Matching Architecture

The three components involved are the following, namely:

- *Tasks*: define the actions to be performed in order to accomplish a mission. A task may have attached environmental conditions (weather, terrain, enemy, etc.) that are expected to impact the performance of a task[4].
- *Sensors*: these are the assets that collect the information required to satisfy the intelligence requirements of a mission. However, sensors do not operate as independent entities, they have to be attached to (or carried by) devices that provide them with energy, protection, mobility, etc.
- *Platforms*: these are the systems to which sensors are attached so as to get energy, protection, mobility, communication, etc. Platforms include both static and mobile systems operating on land, in sea and air.

The three components involved and the dependencies between them result in three different matching relations, as follows:

[3]IMINT stands for Imagery Intelligence, MASINT stands for Measurement and Signature Intelligence, and SIGINT stands for Signals Intelligence

[4]We seek to use standardized catalogues of Tasks and Conditions such as those found in the Universal Joint Task List.

- *Task-Sensor matching*: a sensor S matches a task T, $match(T, S)$, if S provides the information collecting capabilities required to satisfy the intelligence requirements of T.
- *Task-Platform matching*: a platform P matches a task T, $match(T, P)$, if P provides the kind of ISR-supporting capabilities (mobility, survivability, communication) required to perform T.
- *Platform-Sensor matching*: a sensor S matches a platform P, $match(P, S)$, if S can be carried by and is compatible with the characteristics of P.

In order to satisfy the ISR requirements of a mission one needs to select both a platform and a combination of sensors such that the three matching relations of the architecture are satisfied simultaneously. As a proof of concept, we have developed a first software prototype that uses this architecture and the kind of matching relations described in the previous subsection (§II-A).

III. Towards a multidimensional solution

In the last section we introduced and briefly discussed the benefits of a semantic matchmaking approach to address the sensor-mission assignment problem. In this section we describe some issues concerning the development of the ontologies to support matchmaking, and discuss further on the problems faced and the kind of solutions we propose.

Although one can think of a single ontology, actually we adhere to the Semantic Web vision of multiple interlinking ontologies covering different aspects of the domain. On the one hand, we provide an ontology based on the Missions and Means Framework (MMF), which is basically a collection of concepts and properties to reason about the requirements of a mission and the means required to accomplish it (e.g. mission, task, capability, asset, etc.). On the other hand, we provide another ontology that refines some of the generic concepts in the MMF ontology so as to represent the ISR concepts that constitute our particular application domain. This second ontology comprises several areas frequently organized as taxonomies, such as a classification of sensors (acoustic, optical, chemical, radar), a classification of platforms (air, sea, ground and underwater platforms), a classification of mission types, and a classification of ISR capabilities (reconnaissance, surveillance, target acquisition, damage assessment, etc.). Figure 5 sketches an integrated view of the main concepts that are defined in our ontologies. On the left hand side, we have the concepts related to the mission: a mission comprises several tasks that need to be accomplished. On the right hand side we have the concepts related to the means: a sensor is a system that can be carried by or constitute part of a platform; and inversely, a platform can accommodate or have one or more systems; both platforms and systems are assets; an asset provides one or more capabilities; a capability can entail a number of more elementary capabilities, and is required to perform certain type of tasks; and inversely, a task is enabled by a number of capabilities.

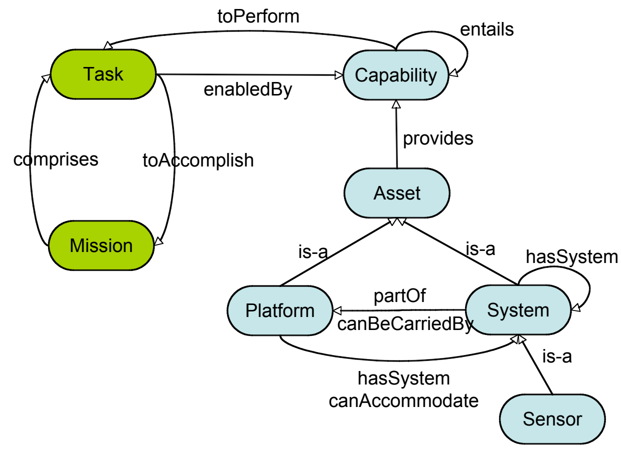

Fig. 5. Main ontological concepts and their relationships

When modelling the ISR domain, we have found several problems that make it difficult to formulate a single taxonomy of concepts. A subset of the problems encountered includes:

- Some issues related to the classification of things:
  - The absence of standardized taxonomies, and the existence of alternative, some times inconsistent classifications for the same concepts.
  - Existing attempts to conceptualize the domain are based on different dimensions, and more usually, several dimensions are mixed: for example, UAV classifications tend to mix dimensions such as size or weight (eg. Micro and Mini UAV), performance (e.g. MALE and HALE UAV), mission type, or certain adhoc features such as having VTOL (Vertical Takeoff and Landing) capabilities, or carrying weapons (Combat UAV)
- There are some fuzzy concepts that are difficult to classify as a single category, for example a LIDAR (LIght Detection and Ranging) is an optical remote sensing technology which measures the time delay between transmission of a light pulse and detection of the reflected signal to find range, so it has properties of both optical sensors and radars (sometimes it is called LADAR, for Laser Radar).
- Some concepts that are supposed to refer to the same aspect of the domain are described at quite different levels of abstraction. Closely related to this issue is the tension between considering a concept as primitive, or as a composition of more basic elements; for example, a reconnaissance capability might be seen as implying a combination of mobility and sensing capabilities.

In order to deal with the challenges introduced above, we propose a compositional and multidimensional approach to conceptualize the ISR domain. Such an approach is well suited to Description Logics (DL) languages such as OWL[5]. One

[5]http://www.w3.org/TR/owl-guide. OWL comprises three sublanguages: Lite, DL, and Full. Herein we are implicitly referring to OWL-DL

of the most powerful features of DL resides in the ability to define classes in terms of sufficient and necessary conditions. New concepts are defined by specifying constraints on existing concepts and properties. As an example, consider the following definitions:

- An Aircraft is-a Platform & has Realm Atmosphere
- An UnmannedVehicle: is-a Platform & has Quality Without-crew-mobility
- A UAV is-a Aircraft & is-a UnmannedVehicle
- Combat UAV: is-a UAV & has Capability Firepower
- A MediumAltitudeLongEndurance (MALE) UAV: is-a UAV & has Capability Endurance & has Capability MediumAltitude

Thus, it is very easy to characterize the same individuals across several conditions that will qualify the same individual as a member of different classes. Having this in mind, OWL enables the use of alternative taxonomies to characterize the same objects across multiple dimensions, and to use those alternative classifications for different purposes. For example at some point, one may be interested in selecting a platform in terms of the type of tasks that it can perform (reconnaissance, surveillance, battle damage assessment, etc.), but in other circumstances one may be interested in selecting a platform according to their takeoff and landing capabilities (catapult, runaway, VTOL, etc.). Below follows a short account of a number of these dimensions:

- For platforms: mobility, realm, performance (range, endurance/dwell time, altitude, speed, etc.), application or mission type (surveillance, reconnaissance, Target acquisition, BDA, decoy, etc.), firepower, landing and takeoff, communications, vulnerability and survivability, availability.
- For missions: target characteristics (e.g. collectable vs observable), range to the target, timeliness, battlespace factors, threat, terrain, contamination, weather.
- For sensors: phenomena detected (type and spectrum), performance (quality of data, accuracy, etc), weather/terrain/contamination influence, vulnerability, availability

Figure 6, adapted from [8], sketches the process followed by intelligence experts to assess and select the intelligence discipline (IMINT, MASINT, SIGINT, etc.) and/or the assets that are more suited to met the intelligence requirements of a mission. On the left hand side, the figure shows some of the key elements resulting from the analysis of the information requirements of a mission, such as range to the mission, geography, weather, or adversary activity. On the right hand side, we can see some features used to characterize the capabilities provided by the available ISR assets, such as range, limitations to weather and terrain masking, or threats the asset survivability. The arrows between the left and right axes represent the matching relations to be established between requirements (left) and asset capabilities (right), so as to select a discipline (IMINT, SIGINT, etc.) and/or a specific system (sensors and platforms). Note that according to this doctrine, there are two main activities to perform: correlate and compare. *Correlate* implies there are numeric variables involved which require a quantitative treatment, such as range and time, while *compare* indicates that the variables involved are qualitative.

The kind of ontology-based matchmaking that we have described so far is based on the semantic relationships between classes, which are inherently qualitative descriptions of the world. In order to tackle the quantitative elements involved in this process we will extend the matchmaking algorithms as part of our future work.

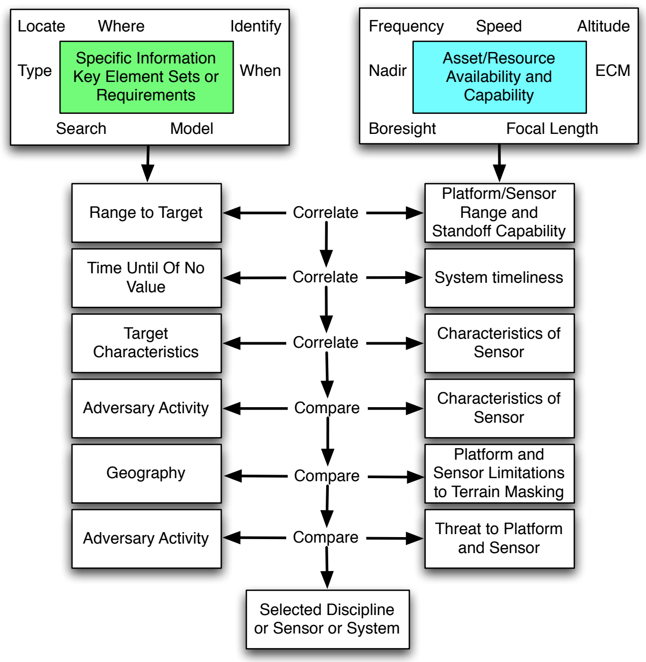

Fig. 6. Dimensions involved in the matchmaking process between information requirements and ISR assets

IV. Related work

So far we have motivated the use of ontologies to support the assignment of ISR means to missions. There is already a sizeable amount of work in providing descriptive schemas for sensors, sensor platforms, and their properties, for example the OpenGeospatial Consortium (OGC)[6] suite of Sensor Web Enablement (SWE)[7] specifications. There are also a number of prototype ontologies for sensors and sensor platforms including OntoSensor [5], [10], CIMA [9], the Marine Platforms Ontology [1], and an ontology for level 1 sensor fusion [4].

One of our goals for this task is to avoid reinventing the wheel, and to build on these existing representations. We believe that the existing sensor and source representations partly cover what we need to model, namely platforms (described in the MMI platforms ontology) and sensors

[6]http://www.opengeospatial.org
[7]http://www.opengeospatial.org/projects/groups/sensorweb

(OntoSensor and CIMA). However, these attempts lack the part that is central to our work, that is the linking of missions and means. In the rest of this section we discuss some aspects from aforementioned schemas and ontologies that are relevant to our work.

SensorML[8] started as a standalone XML-based sensor model language for describing the geometric, dynamic, and radiometric properties of dynamic remote sensors, and has now become an integral part of the SWE initiative of the OGC. The OGC is an international consortium of more than 330 companies, government agencies and universities participating in a consensus process to develop publicly available standards for geospatial and location-based services.

In the SWE approach, sensors are self-describing Web components, which are discoverable and accessible in real time by Web services [3]. Also, sensor information and their observations are made accessible in real-time through Web services, thus enabling sensor systems to find phenomena of immediate interest autonomously. SensorMLs role in SWE is to provide XML Schemas for describing sensor systems and processes, information needed for the discovery of sensors, location of sensor observations, and processing of low-level sensor observations.

The ability of a sensor or a sensor system to represent its capabilities is key to our work. SensorML allows us to represent the capabilities of such systems within an XML element named *capabilities*. Figure 7 shows a code snippet representing sensor capabilities in SensorML.

```
<sml:capabilities>
<swe:DataRecord>
 <swe:field name="Depth Capability"
  xlink:role="urn:x-ogc:def:property:operationalLimit">
  <swe:Quantity
  definition="urn:x-ogc:def:classifier:SBE:depthCapability">
  <swe:uom code="m"/>
 <swe:value>7000</swe:value>
</swe:Quantity>
</swe:field>
    ...
</swe:DataRecord>
</sml:capabilities>
```

Fig. 7. "Depth Capability" description of a MicroCAT CTP Recorder in SensorML

SensorML uses a *definition* attribute (which is used throughout the SWE Common data types *Quantity*, *Category*, *Count*, *DataRecord*, *DataArray*, etc) to uniquely identify terms by pointing to a Uniform Resource Name (URN)[9]. It is important to have a common representation for these definitions; otherwise there will not be any mechanism to guarantee the interoperability of different SensorML descriptions from different organisations.

Currently SensorML's mechanisms to resolve the URNs are not finalised yet. This is where an ontology could be very useful: the *definition* attribute could point to a term in an ontology. So if there is a well-defined ontology for capabilities of sensors and platforms, then one could use these terms in SensorML *definition* attributes to allow that sensor system to be discovered and/or evaluated. There are several activities in public domain to create common vocabularies to address this issue. For example, Marine Metadata Interoperability (MMI)[10] has defined vocabularies for sensor platforms in the marine operations context, and Semantic Web for Earth and Environmental Terminology (SWEET)[11] has defined a list of vocabularies to express Earth science data and information.

OntoSensor aims to create a prototype sensor knowledge repository, which is compatible with the evolving Semantic Web infrastructure [5], [10]. It has been used to mark-up live data from sensors in a network to achieve efficient data fusion. The design approach of OntoSensor has been a compositional one. OntoSensor includes some concepts from SensorML and refers to ISO 19115 for definitions regarding geographic information. Also it has extended IEEE Suggested Upper Merged Ontology (SUMO) upper-level ontology (e.g. *Sensor* and *Platform* concepts of OntoSensor extend *MeasurementDevice* and *TransportationDevice* concepts of the IEEE SUMO ontology.

Though OntoSensor is developed for data fusion some of the concepts and relationships identified by it have influenced our own ontology development efforts.

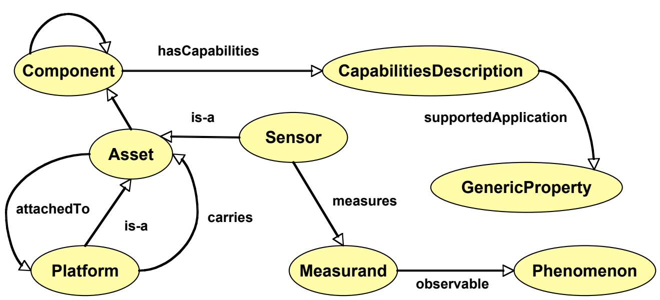

Fig. 8. Selected concepts and relationships from OntoSensor

As Figure 8 shows *Sensors* and *Platforms* are assets, which are of type *Component*. A sensor could be attached to a platform (or platform could carry a sensor) and this is represented by properties *carries* and *attachedTo*, respectively. A sensor can measure multiple measurands; a particular measurand is considered to be an observable instance of some phenomenon. Components (i.e. sensors and platforms) have capabilities (*hasCapabilities*), which are described (*CapabilitiesDescription*) in terms of applications that they support (*supportApplication*). These applications are based on the generic properties (field-of-view (FoV), radio frequency, zero-G output) of the sensors and sensor platforms. Sensors are classified under the energies that they observe (acoustic, chemical, magnetic etc.).

[8]http://vast.uah.edu/SensorML
[9]http://www.w3.org/TR/uri-clarification/#urn-namespaces
[10]http://marinemetadata.org
[11]http://sweet.jpl.nasa.gov/ontology

OntoSensor does not meet our own needs because, under the requirements that it was created, a great deal of emphasis has been put on modelling the data from sensors, but not their functional aspects. We are interested in the functional aspects of sensors and sensor systems so that we can reason about them to match sensing systems to the requirements of missions.

The various initiatives described above are relevant to our work in different ways. SensorML and the other SWE components provide a Service Oriented Architecture (SOA) framework for deploying Web-enabled sensor systems; by annotating SensorML descriptions using our ontology, we can use the ontology to reason about which sensors to deploy, and then use SWE to implement the deployment. OntoSensor and CIMA have features compatible with our approach but they do not satisfy our matchmaking requirements. This is perhaps unsurprising, because these ontologies were both designed to satisfy a different need: data fusion from a sensor / instrument network. They lack the definitions of higher-level capabilities needed to match assets to tasks (Figure 5). However, we expect that it will be possible to align our ontology with OntoSensor, and partially reuse instance data between them.

V. First Software Prototype

As a proof-of-concept of the proposed approach to sensor-mission assignment, we have built a prototype version of matchmaking software that utilizes both the mission scripts and sensor ontology representations. This software allows a commander to specify the information needs of a mission, and obtain recommended assets to accomplish the mission. The prototype is based on a preliminary ontology (Figure 5) that accounts for many of the concepts in MMF.

To support the specification of intelligence requirements we have included three types of requirements:

- *Operational requirements*: refer to the kind of ISR tasks to be performed as part of an operation. These requirements include five main categories (Reconnaissance, Surveillance and Target Acquisition, Damage Assessment and Artillery Adjustment) and more than twenty subcategories (e.g. Maritime Reconnaissance, Constant Surveillance, etc.). This integrates basic knowledge acquired from relevant literature and a more detailed conceptualization based on the CALL Thesaurus[12].
- *Intelligence disciplines*: we include a taxonomy of intelligence disciplines (IMINT, SIGINT, MASINT, etc.) that is based on the Joint Capability Areas[13] terminology and the CALL Thesaurus)
- Some *platform specific capabilities* that allow the commander to further constrain the kind of assets required (or preferred) to accomplish a given mission. This knowledge represents attributes that are specific to certain platform types, such as the altitude and range of an UAV.

[12]http://call.army.mil/products/thesaur-frame.asp
[13]http://www.dtic.mil/futurejointwarfare/cap_areas.htm

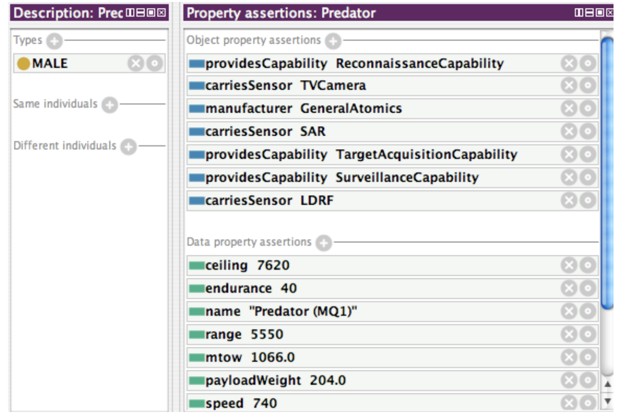

Fig. 9. Example of a platform specification: UAV Predator

In the current prototype we have focused on the UAV domain; the resulting knowledge base includes instances representing a dozen types of UAV and around 20 sensing devices such as microphones and cameras. Figure 9 shows a platform specification example; in particular, we show the specification of a *Predator*, which is an instance of the *MALE-UAV* class. It should be noted that some properties asserted for a given individual are a logical consequence of the class definition; for example, the property of providing constant surveillance capability holds for the *Predator* because the *MALE-UAV* class is a subclass of the *EnduranceUAV* class, which is defined as providing constant surveillance.

Figure 10 shows an screenshot of the software prototype that gives an idea of the kind of requirements used to specify the ISR requirements of a mission. In the example the user has selected IMINT and constant surveillance as the requirements to be met. Figure 11 shows the set of recommended assets suggested by the system to satisfy the former requirements. First, the system recovers all the platforms that provide constant surveillance; second, it retrieves all the sensors that support the production of IMINT; and third, using the results of the former steps, the the system selects valid platform configurations. For example, an *E-Hunter* is UAV that can provide constant surveillance and IMINT because it can carry both a TV camera and also an infrared camera.

VI. Conclusions

In this paper we have argued for an approach to the sensor-mission assignment problem grounded on the use of *ontologies*: we have shown that this approach can be used with deductive reasoning mechanisms to produce matches that are logically sound. We have built on current best-practice in analysing mission requirements and means by creating an ontology based on the military Missions and Means Framework. In addition, we have created an ontology that refines the MMF ontology deal with the ISR domain. We have reviewed existing work in representing sensors and sources, and discussed why these approaches are complementary, rather than alternatives to our work. Finally, we have described a prototype software tool

Fig. 10. Example of requirements

Fig. 11. Example of recommended assets

that is an initial proof-of-concept of our proposed approach.

The current paper has focussed on pre-assignment of sensors and sources to missions: we have not considered real-time and dynamic aspects of the problem here. Future work will expand the approach to consider these issues. Our ontology for sensors and sources can easily be extended to include real-time properties of the sensor, source, and platform instances (for example, power levels, damage sustained, etc); in fact, many of these properties are already defined in ontologies such as those surveyed in the previous section. By incorporating such properties we can reason about the real-time state of a sensor, source, or platform, and consider re-assignment at mission run-time.

References

[1] L. Bermudez, J. Graybeal, and R. Arko. A marine platforms ontology: Experiences and lessons. In *Proceedings of the 2006 Workshop on Semantic Sensor Networks*, Athens GA, USA, 2006.

[2] W. N. Borst. *Engineering Software Library Systems*. PhD thesis, Twente University, 1997.

[3] M. Botts, A. Robin, J. Davidson, and I. Simonis. OpenGIS sensor web enablement architecture document. Technical report, Open Geospatial Consortium Inc, 2006. OpenGIS Discussion Paper.

[4] M. Ceruti. Ontology for level-one sensor fusion and knowledge discovery. In *Proceedings of the 2004 Workshop on Knowledge Discovery and Ontologies*, Pisa, Italy, 2004.

[5] C. Goodwin and D. Russomanno. An ontology-based sensor network prototype environment. In *Proceedings of the 5th International Conference on Information Processing in Sensor Networks*, pages 1–2, Nashville TN, USA, 2006. IEEE.

[6] T. R. Gruber. Toward principles for the design of ontologies used for knowledge sharing. *Journal of Human Computer Studies*, 43(5/6):907–928, 1994.

[7] N. Guarino. Formal ontologies and information systems. In *Proceedings of the 1st International Conference on Formal Ontologies in Information Systems (FOIS-98)*, pages 3–15. IOS Press, June 1998.

[8] Joint Publications. JP 2-01: Joint and national intelligence support to military operations, 2004.

[9] D. McMullen and T. Reichherzer. The common instrument middleware architecture (CIMA): Instrument ontology & applications. In *Proceedings of the 2nd Workshop on Formal Ontologies Meets Industry*, pages 655–663, Trento, Italy, 2006.

[10] D. Russomanno, C. Kothari, and O. Thomas. Building a sensor ontology: A practical approach leveraging ISO and OGC models. In *Proceedings of the 2005 International Conference on Artificial Intelligence*, pages 637–643. CSREA Press, 2005.

[11] J. H. Sheehan, P. H. Deitz, B. E. Bray, B. A. Harris, and A. B. H. Wong. The military missions and means framework. In *Proceedings of the Interservice/Industry Training and Simulation and Education Conference*, pages 655–663, 2003.

Acknowledgements

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.