

AFRL-IF-RS-TR-2005-361
Final Technical Report
October 2005



DARPA AGENT MARKUP LANGUAGE (DAML) UNIFIED MODELING LANGUAGE (UML)-BASED ONTOLOGY TOOLSET (UBOT)

Lockheed Martin Integrated Systems and Solutions

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. K359/00

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

STINFO FINAL REPORT

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2005-361 has been reviewed and is approved for publication

APPROVED: /s/

JOSEPH A. CAROZZONI
Project Engineer

FOR THE DIRECTOR: /s/

JAMES W. CUSACK, Chief
Information Systems Division
Information Directorate

REPORT DOCUMENTATION PAGE			<i>Form Approved</i> OMB No. 074-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE OCTOBER 2005	3. REPORT TYPE AND DATES COVERED Final Aug 00 – Dec 04	
4. TITLE AND SUBTITLE DARPA AGENT MARKUP LANGUAGE (DAML) UNIFIED MODELING LANGUAGE (UML)-BASED ONTOLOGY TOOLSET (UBOT)			5. FUNDING NUMBERS C - F30602-00-C-0188 PE - 62702E PR - DAML TA - 00 WU - 01	
6. AUTHOR(S) Paul Kogut				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Lockheed Martin Integrated Systems and Solutions PO Box 8048 Philadelphia Pennsylvania 19101			8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Defense Advanced Research Projects Agency AFRL/IFSB 3701 North Fairfax Drive Arlington Virginia 22203-1714			10. SPONSORING / MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TR-2005-361	
11. SUPPLEMENTARY NOTES AFRL Project Engineer: Joseph Carozzoni/IFSB/(315) 330-7796/ Joseph.Carozzoni@rl.af.mil				
12a. DISTRIBUTION / AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.				12b. DISTRIBUTION CODE
13. ABSTRACT (Maximum 200 Words) The objective for the UBOT project was to reduce the barriers for adoption of Semantic Web technology. The target adopters were DoD and Intelligence community software practitioners. The UBOT project focused on three major barriers: Ontology engineering is hard, Semantic markup creation is time consuming and expensive, and Semantic Web application engineering trade-offs are not well understood. The approach to reducing these barriers was to apply software engineering principles: automated tools, formal methods and software architecture. There were four research focus areas in this project: 1) Ontology engineering, 2) Automated markup generation, 3) Reasoning architectures and benchmarks, and 4) Semantic Web architectures and design patterns.				
14. SUBJECT TERMS DARPA Agent Markup Language, DAML, W3, RDF, OIL, Software Agents				15. NUMBER OF PAGES 21
				16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

Table of Contents

SUMMARY	1
1. INTRODUCTION.....	3
1.1 IDENTIFICATION	3
1.2 PURPOSE AND SCOPE	3
1.3 AUDIENCE.....	3
1.4 DOCUMENT OVERVIEW.....	3
2. ONTOLOGY ENGINEERING	4
2.1 METHODS, ASSUMPTIONS AND PROCEDURES.....	4
2.2 RESULTS AND DISCUSSION.....	5
2.3 REFERENCES	6
3. MARKUP GENERATION	7
3.1 METHODS, ASSUMPTIONS AND PROCEDURES.....	7
3.2 RESULTS AND DISCUSSION.....	8
3.3 REFERENCES	9
4. REASONING ARCHITECTURES AND BENCHMARKS.....	10
4.1 METHODS, ASSUMPTIONS AND PROCEDURES.....	10
4.2 RESULTS AND DISCUSSION.....	11
4.3 REFERENCES	11
5. SEMANTIC WEB ARCHITECTURES AND DESIGN PATTERNS	12
5.1 METHODS, ASSUMPTIONS AND PROCEDURES.....	12
5.2 RESULTS AND DISCUSSION.....	12
5.3 REFERENCES	13
6. CONCLUSIONS AND RECOMMENDATIONS.....	14
REFERENCES.....	15

List of Figures

FIGURE 1 CONSISTENCY REASONING AND VISUALIZATION ENVIRONMENT	4
FIGURE 2 UML GRAPHICAL INTERFACE FOR ONTOLOGY ENGINEERING	5
FIGURE 3: AEROSWARM	8
FIGURE 4 BENCHMARK TOOLS	10
FIGURE 5 DLDB.....	11
FIGURE 6 DAML-S PROTOTYPE.....	12
FIGURE 7 OWL-S PROTOTYPE	13

Summary

This document is the final technical report for the DARPA Agent Markup Language (DAML) program UML-Based Ontology Toolset (UBOT) project. This work was performed under Air Force Research Lab (AFRL) contract F30602-00-C-0188. The document describes the technical work performed on the UBOT project.

Lockheed Martin Integrated Systems and Solutions (formerly Management & Data Systems) was the lead contractor on the UBOT project. However, the report reflects work performed by the entire team, including participants from the following organizations:

- Versatile Information Systems
- Lehigh University
- Kestrel Institute
- Lockheed Martin Advanced Technology Center
- Telelogic

Last but not least, this work was driven and shaped by the leadership and vision of the DARPA DAML PMs Jim Hendler, Murray Burke, Mark Greaves.

The basic motivation for the UBOT project was to reduce the barriers for adoption of Semantic Web technology. The target adopters were DoD and Intelligence community software practitioners (government and contractor). The UBOT project focused on three major barriers:

- Ontology engineering is hard
- Semantic markup creation is time consuming and expensive
- Semantic Web application architectures and engineering tradeoffs are not well understood

The approach to reducing these barriers was to apply software engineering principles:

- Automated tools
- Formal methods
- Software architecture

There were four research focus areas in this project:

- Ontology engineering

- o UML based graphical tools were developed by Lockheed Martin (LM) and Telelogic to help visualize complex Semantic Web ontologies emphasizing both taxonomic hierarchies and object properties.
 - o Easy to use consistency reasoning tools were developed by Versatile Information Systems (VIS) and Kestrel to ensure that ontologies and markup were semantically valid.
- Automated markup generation tools based on natural language processing were developed by LM to reduce the cost of creating markup from text.
- Reasoning architectures and benchmarks were developed by Lehigh University to understand the engineering tradeoffs involved in Semantic Web applications.
- Semantic Web architectures and design patterns were developed by LM to help adopters get started on real world applications in areas such as knowledge management and net-centric operations.

1. Introduction

1.1 Identification

This document is the final technical report for the DARPA Agent Markup Language (DAML) program UML-Based Ontology Toolset (UBOT) project. This work was performed under Air Force Research Lab (AFRL) contract F30602-00-C-0188.

1.2 Purpose and Scope

This report describes the technical work performed on the UBOT project.

1.3 Audience

This report was written for a technical audience that has a basic understanding of Semantic Web technology. See www.daml.org and www.semwebcentral.org for background information.

1.4 Document Overview

The UBOT project had four separate technical threads. This document is organized based on these threads into the following sections and supplementary material:

- *Section 1: Introduction*
- *Section 2: Ontology Engineering*
- *Section 3: Markup Generation*
- *Section 4: Reasoning Architectures and Benchmarks*
- *Section 5: Semantic Web Architectures and Design Patterns*
- *Section 6: Conclusions and Recommendations*
- *Appendix A: Publications*

2. Ontology Engineering

2.1 Methods, Assumptions and Procedures

The goal of this research thread was to help mainstream software engineers move from object-oriented development to the Semantic Web paradigm. Around 2000 many software engineers were trained to do object-oriented analysis and design using the Unified Modeling Language (UML) which was a standard sponsored by the Object Management Group (OMG). Because UML was familiar to developers the approach was to allow users to model ontologies in UML and automatically generate DAML as shown in figure 1. The front end tool (see figure 2) was based on the COTS product Telelogic Tau which had strong support for large scale collaborative team modeling. To avoid a proliferation of proprietary UML mappings to Semantic Web languages we initiated OMG efforts to define a standard mapping.




Figure 1 Consistency Reasoning and Visualization Environment




Figure 2 UML Graphical Interface for Ontology Engineering

UML was designed for communication between humans unlike Semantic Web languages like DAML and OWL which were designed for communication and reasoning by software agents. Therefore you cannot assume that humans will be available to interpret inconsistent ontologies. Ontologies often quickly become complex and it is difficult for a human to check all of the implications of the ontology. Also, it is important that the ontologies conform to standard semantics so that different reasoners will draw the same inferences. The approach was to do fast reasoning about the conformance of an ontology to the DAML and OWL axiomatic semantics and then use the more powerful Specware/SNARK environment to search for more subtle semantic inconsistencies. ConsVISor evolved as the fast semantics checker. ConsVISor is a web accessible service that allows the user to submit ontologies and get specific explanations of inconsistency symptoms.

2.2 Results and Discussion

The ontology engineering approaches described above came together in the environment shown in Figure 1. The user builds a UML model, checks it for consistency and generates a Semantic Web ontology. We found that importing an existing ontology to visualize and extend it was an important feature.

Efforts to develop an OMG standard mapping between UML and DAML/OWL were more difficult than anticipated. One reason was that UML and DAML/OWL were rapidly evolving. Another reason was that there were some significant semantic mismatches between UML and Semantic Web languages. For example, properties are first class entities in DAML/OWL but in

UML 1.x they cannot exist without being associated with classes. OMG expects to approve a semantic mapping standard in 2005.

The application of the Specware/SNARK environment proved to be not practical for consistency checking of ontologies because it was not easy to use and often ran for minutes without terminating. However, the investigation of this approach yielded unanticipated benefits in that we discovered and fixed inconsistencies in the DAML axiomatic semantics. This formal method approach was later used to check the axiomatic semantics of OWL thus contributing to the W3C OWL standard.

2.3 References

K. Baclawski, M. Kokar, P. Kogut, L. Hart, J. Smith, W. Holmes, J. Letkowski, M. Aronson, "Extending UML to Support Ontology Engineering for the Semantic Web," *Fourth International Conference on UML (UML 2001)*, Toronto, October 2001.

P. Kogut, S. Cranefield, L. Hart, M. Dutra, K. Baclawski, M. Kokar, J. Smith, "UML for Ontology Development," in *Knowledge Engineering Review Journal*, Special Issue on Ontologies in Agent Systems 2002 vol. 17

K. Baclawski, M. Kokar, R. Waldinger, P. Kogut, "Consistency Checking of Semantic Web Ontologies," *1st International Semantic Web Conference (ISWC 2002)*, Sardinia, Italy, June 2002.

K. Baclawski, C. Matheus, M. Kokar, J. Letkowski and P. Kogut, "Towards a Symptom Ontology for Semantic Web Applications." In *Proceedings of Third International Semantic Web Conference*, Hiroshima, Japan, pages 650-667, November, 2004.

3. Markup Generation

3.1 Methods, Assumptions and Procedures

The goal of this research thread is to reduce the cost/effort required to mark up webpages and other text documents. The creation of markup from unstructured text sources such as web pages is tedious and time-consuming. Anyone who produces documents on a regular basis (e.g., intelligence analysts, commanders) or has a large quantity of legacy documents needs some form of automated markup assistance.

The UBOT team experimented with the application of natural language processing technology to reduce the effort required for markup. They have built a tool called AeroSWARM (AeroTextTM Semantic Web Automated Relation Markup) which automatically generates OWL markup for a number of common domain-independent classes and properties. The author can then manually do markup additions and corrections to the output of AeroSWARM. The markup can be posted on the web or an intranet or ingested into a knowledge-base for querying and reasoning.

AeroTextTM, a product of Lockheed Martin Corporation, provides the information extraction capabilities for AeroSWARM tool. AeroSWARM tool is available as an open service on the web. AeroSWARM can also be customized for domain-specific markup generation. Customization is done by adding to the core linguistic rules that come with AeroText. Significant customization was done for the Horus project.

Figure 3 shows the architecture of AeroSWARM. A user can specify the set of web pages to markup via the user interface accessible at <http://ubot.lockheedmartin.com/ubot/hotdaml/aeroswarm.html>. The user then chooses a target ontology and AeroSWARM generates OWL markup. The markup includes entities (e.g., *person*, *place*, *organization*), relations (e.g., *Pinochet hasLocation Santiago*) and co-references (e.g., *Pinochet sameAs Augusto Pinochet*). A table on the AeroSWARM site describes all the entities and relations that are automatically identified and marked-up.




Figure 3: AeroSWARM

There are a number of advanced features that AeroSWARM supports. If the markup creator does not want to use the native AeroSWARM ontology as a target then she can choose a popular upper ontology (e.g., OpenCyc or IEEE SUMO which have predefined mappings) or AeroSWARM provides a drag-and-drop tool to create ontology mappings for user specified ontologies. The user can also semantically check the generated markup against constraints specified in an OWL ontology (e.g., *only one person can be the biological mother of a person*) by sending the output to the ConsVISor tool. Finally, AeroSWARM is a web service, which can be invoked by external tools or applications as part of a more complex workflow.

3.2 Results and Discussion

The processing of raw text is very difficult but sufficient levels of precision and recall are being attained to make this automated assistance approach worthwhile. AeroSWARM can generate markup for 44 common properties with around 80% precision. This is good enough for most human in the loop tasks like semantic filtering of information retrieval results. It is probably not good enough for autonomous financial transactions and safety critical applications. A promising hybrid approach is where AeroSWARM creates draft markup and the author corrects and adds markup. A hybrid approach using Microsoft Word was investigated by LM and Teknowledge.

The UBOT team supported a number of organizations in applying an AeroSWARM approach including UMBC IT talks, Horus project, AT&T Ontology Driven Knowledge Dissemination and other intelligence community applications.

The UBOT team investigated two different applications of AeroSWARM and Semantic Web reasoning to real world knowledge management problems. The first application called GOWLgle uses keywords chosen from an ontology to semantically filter the list of documents returned by

Google. The goal was to improve information retrieval precision by eliminating documents that Google found by chance occurrence of keywords. This approach can be applied by analysts for searching open source or classified document repositories.

The second application involved the use of Semantic Web reasoning to determine the plausibility of cross document co-reference (e.g., person X in document A is the same as person Y in document B). An automated system cannot depend on linguistic cues for cross-document co-reference. It must use corroborating assertions in multiple documents to determine co-reference plausibility. Cross-document co-reference is critical to intelligence analysis tasks such as link analysis.

3.3 References

P. Kogut and W. Holmes, "AeroDAML: Applying Information Extraction to Generate DAML Annotations from Web Pages," *First International Conference on Knowledge Capture (K-CAP 2001) Workshop on Knowledge Markup and Semantic Annotation*, Victoria, BC, October 2001

P. Kogut, Y. Leung, K. Ryan, L. Gohari, M. Kokar, J. Letkowski "Applying Semantic Web Technologies to Counter-Terrorism", to appear in *21st Century Information Technologies and Enabling Policies for Counter-Terrorism* R. Popp and J. Yen, IEEE Press 2005

4. Reasoning Architectures and Benchmarks

4.1 Methods, Assumptions and Procedures

The goal of this research thread was to help practitioners choose appropriate reasoning infrastructures for large OWL applications and investigate scalability approaches for reasoners. It is well known that there is a qualitative tradeoff between size and complexity of the knowledge-base and query response time. Our goal was to establish a quantitative approach to making this engineering tradeoff. This research thread was only partially funded under the DAML program.

To support reasoning architecture choices we developed benchmark processes and tools as shown in figure 4. The data generator and test queries support repeatable benchmark evaluations.




Figure 4 Benchmark Tools

We also investigated an approach to reasoning scalability called Description Logic Database (DLDB). The idea shown in figure 5 was to pre-compute all inferences and store the results in a database for fast query response.




Figure 5 DLDB

4.2 Results and Discussion

The benchmark includes 14 test queries representative of real applications (vs. OWL conformance tests). The data generator has been used to create a data set with 6.8 million RDF triples. The benchmark process was applied to four different reasoning infrastructures including DLDB. The benchmarking results are summarized in a paper that won the best paper award at the 2004 International Semantic Web Conference.

4.3 References

- Guo, Y.; Heflin, J; and Pan, Z. Benchmarking DAML+OIL Repositories *Second International Semantic Web Conference, ISWC 2003*, LNCS 2870. Springer, 2003, pp. 613-627.
- Pan, Z and Heflin, J. DLDB: Extending Relational Databases to Support Semantic Web Queries. In *Workshop on Practical and Scaleable Semantic Web Systems*, ISWC 2003, pp. 109-113.
- Y. Guo, Z. Pan, and J. Heflin. An Evaluation of Knowledge Base Systems for Large OWL Datasets. *Third International Semantic Web Conference*, Hiroshima, Japan, LNCS 3298, Springer, 2004, pp. 274-288.

5. Semantic Web Architectures and Design Patterns

5.1 Methods, Assumptions and Procedures

The goal of this research thread was to answer the following question: *How can the Semantic Web be applied to solve real problems for the DoD and the Intelligence community?* The approach was to explore prototype applications and disseminate design patterns to practitioners. There were two main application areas that were explored:

- Semantic Web services for net-centric operations including command and control (C2) and battlefield intelligence, surveillance and reconnaissance (ISR). ISR involves the allocation of sensors and the fusion of sensor data.
- Knowledge management which includes assistant agents and semantic filtering (which was discussed in section 3).

Some of the prototype applications described in Section 5 originated in the UBOT project and continued on other sources of funding.

5.2 Results and Discussion

We developed a prototype shown in figure 6 which used DAML-S for finding electro-optic imagery and synthetic aperture radar (SAR) imagery relevant to a military course of action. This prototype was integrated into the larger DAML experiment demo. We also experimented with DAML-S in the NASA solar science domain.




Figure 6 DAML-S Prototype

We are continuing our Semantic Web service efforts by developing an OWL-S prototype that illustrates the vision of Semantic Web service based improvisational workflows for effect-based

Net-Centric Operations. Figure 7 shows the approach. The user chooses an effect and prioritizes measures of effectiveness (MOEs) by selecting classes from an OWL ontology. The workflow generator identifies the best set of services that instantiate activities derived from the chosen effect. Workflow generation involves the interaction between service discovery and genetic programming based service composition. The current output is a graphical representation of the workflow. In the future the output should be an OWL-S process model that can be used for monitoring execution.




Figure 7 OWL-S Prototype

We also developed a prototype system that illustrates the vision of Semantic Web-based assistant agents for the Intelligence Community and the Department of Homeland Security analysts. The user defines an email filter profile by selecting classes from an OWL ontology and filling in instance values. The agent uses the profile and reasons about the relevance of incoming messages. Reasoning is accomplished by automatically generated queries to the Network Inference Cerebra server. The user also defines an info gather profile that instructs the agent to automatically gather additional relevant information based on the contents of the message that has passed through the filter. The output is a concise report of this proactive info gathering with hyperlinks to original sources.

5.3 References

P. Kogut, and J. Heflin, "Semantic Web Technologies for Aerospace," *Proceedings of IEEE Aerospace Conference*, Big Sky, MT, March 2003.

P. Bose, M. Woodward, N. Hurlburt, S. Freeland "Information Fusion in the Sensorweb" *Proceeding of Earth Science Technology Conference*, Pasadena, CA June 2002.

P. Kogut, J. Yen, Y. Leung, S. Sun, R. Wang, T. Mielczarek, B. Hellar, "Proactive Information Gathering for Homeland Security Teams," *Communications of the ACM*, March 2004.

6. Conclusions and Recommendations

The UBOT project made significant contributions to the Semantic Web revolution. We applied software engineering principles to the following problems:

- Ontology engineering is hard
 - We developed practical semantic consistency checking tools - ConsVISor
 - We helped bridge OMG (UML) and W3C (OWL)
- Markup creation is time consuming and expensive
 - We pioneered application of NLP for markup generation – AeroSWARM
- Semantic Web architecture and engineering tradeoffs are not well understood
 - We developed and disseminated design patterns
 - Knowledge management and Semantic Web services
 - We established a benchmark framework for choosing appropriate Semantic Web reasoning infrastructures

AeroSWARM, ConsVISor and Lehigh University Benchmark are available on www.semwebcentral.org.

There is much more work that needs to be done related to the problems listed above:

- Ontology engineering is still hard. Various approaches such as machine learning, cognitive science, advanced NLP, intelligent assistants, reasoners with explanation capability should be investigated. The ConsVISor tool should be extended to check Semantic Web Rule Language (SWRL).
- Research in semantic filtering (e.g., GOWLgle) and cross-document co-reference is still in its infancy. Further research in these knowledge management areas would require improvements in automated markup generation.
- The reasoning benchmark approach should be extended to cover SWRL and Semantic Web service matchmakers.
- There are many open questions for Semantic Web services:
 - How do service discovery, service composition and execution monitoring interact?
 - How do we do matchmaking in heterogeneous dynamic hostile environments like DoD net-centric operations?

References

Ontology Engineering

K. Baclawski, M. Kokar, P. Kogut, L. Hart, J. Smith, W. Holmes, J. Letkowski, M. Aronson, "Extending UML to Support Ontology Engineering for the Semantic Web," *Fourth International Conference on UML (UML 2001)*, Toronto, October 2001.

P. Kogut, S. Cranefield, L. Hart, M. Dutra, K. Baclawski, M. Kokar, J. Smith, "UML for Ontology Development," in *Knowledge Engineering Review Journal*, Special Issue on Ontologies in Agent Systems 2002 vol. 17

K. Baclawski, M. Kokar, R. Waldinger, P. Kogut, "Consistency Checking of Semantic Web Ontologies," *1st International Semantic Web Conference (ISWC 2002)*, Sardinia, Italy, June 2002.

K. Baclawski, C. Matheus, M. Kokar, J. Letkowski and P. Kogut, "Towards a Symptom Ontology for Semantic Web Applications." In *Proceedings of Third International Semantic Web Conference*, Hiroshima, Japan, pages 650-667, November, 2004.

Markup Generation

P. Kogut and W. Holmes, "AeroDAML: Applying Information Extraction to Generate DAML Annotations from Web Pages," *First International Conference on Knowledge Capture (K-CAP 2001) Workshop on Knowledge Markup and Semantic Annotation*, Victoria, BC, October 2001

P. Kogut, Y. Leung, K. Ryan, L. Gohari, M. Kokar, J. Letkowski "Applying Semantic Web Technologies to Counter-Terrorism", to appear in *21st Century Information Technologies and Enabling Policies for Counter-Terrorism* R. Popp and J. Yen, IEEE Press 2005

Reasoning Architectures and Benchmarks

Guo, Y.; Heflin, J; and Pan, Z. Benchmarking DAML+OIL Repositories *Second International Semantic Web Conference, ISWC 2003*, LNCS 2870. Springer, 2003, pp. 613-627.

Pan, Z and Heflin, J. DLDB: Extending Relational Databases to Support Semantic Web Queries. In *Workshop on Practical and Scaleable Semantic Web Systems*, ISWC 2003, pp. 109-113.

Y. Guo, Z. Pan, and J. Heflin. An Evaluation of Knowledge Base Systems for Large OWL Datasets. *Third International Semantic Web Conference*, Hiroshima, Japan, LNCS 3298, Springer, 2004, pp. 274-288.

Semantic Web Architectures and Design Patterns

P. Kogut, and J. Heflin, "Semantic Web Technologies for Aerospace," *Proceedings of IEEE Aerospace Conference*, Big Sky, MT, March 2003.

P. Bose, M. Woodward, N. Hurlburt, S. Freeland "Information Fusion in the Sensorweb" *Proceeding of Earth Science Technology Conference*, Pasadena, CA June 2002.

P. Kogut, J. Yen, Y. Leung, S. Sun, R. Wang, T. Mielczarek, B. Hellar, "Proactive Information Gathering for Homeland Security Teams," *Communications of the ACM*, March 2004