USING THE SENSOR WEB TO DETECT AND MONITOR THE SPREAD OF WILD FIRES

Deshendran Moodley[a], Andrew Terhorst[b], Ingo Simonis[c], Graeme McFerren[b] and Frans van den Bergh[b]

[a]School of Computer Science
The University of Kwazulu Natal
Durban, South Africa
moodleyd37@ukzn.ac.za

[b]ICT for Earth Observation Research Group
Meraka Institute
Pretoria, South Africa
[aterhorst|gmcferren|fvdbergh]@meraka.org.za

[c]Institute for Geoinformatics
University of Muenster
Muenster, Germany
simonis@uni-muenster.de

KEY WORDS: wild fire, disaster, monitoring, decision support, interoperability, architecture

ABSTRACT:

Key concepts in disaster response are level of preparedness, response times, sustaining the response and coordinating the response. Effective disaster response requires a well-developed command and control framework that promotes the flow of information. The Sensor Web is an emerging technology concept that can enhance the tempo of disaster response. We describe how a satellite-based system for regional wild fire detection is being evolved into a fully-fledged Sensor Web application.

1 INTRODUCTION

Most disasters are of short duration and require a fixed amount of consequence management. Examples include earthquakes, tsunamis and storm events. Other disasters are more complex and unfold in a non-linear fashion over an extended period. Such disasters require ongoing and adaptive consequence management. Examples include the outbreak contagious diseases (*e.g.* bird flu) and wild fires.

Key concepts in disaster response are level of preparedness, response times, sustaining the response and coordinating the response (Annoni et al., 2005). Time is critical and the three primary challenges in the race against time are *uncertainty*, *complexity* and *variability* (Rosen et al., 2002). Dealing with these challenges requires a well-designed command and control framework that promotes the flow of information. Many consider the foundation for command and control to be the Observe-Orient-Decide-Act (OODA) loop (Osinga, 2005). The time it takes to complete an OODA cycle is what determines the tempo of the disaster response (Figure 1).

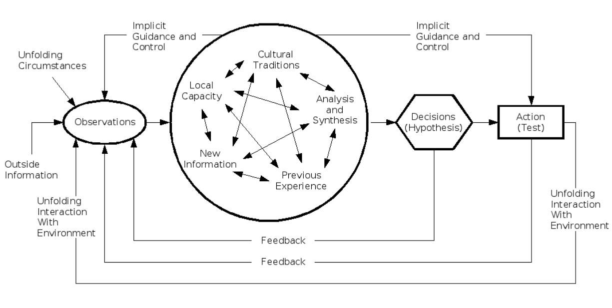

Figure 1: The OODA loop

Advances in sensor technology and distributed computing, coupled with the development of open standards that facilitate sensor/sensor network interoperability, are contributing to the emergence of a phenomenon known as the 'Sensor Web'(Liang and Tao, 2005). This phenomenon can be described as an advanced Spatial Data Infrastructure (SDI) in which different sensors and sensor networks are combined to create a sensor-rich feedback control paradigm (Zibikowski, 2004). In this paper, we describe how the Sensor Web can enhance the tempo of disaster response in context of wild fires.

2 SENSOR WEB ENABLEMENT

Sensor Web Enablement (SWE) is an Open Geospatial Consortium (OGC) initiative that extends the OGC web services framework (OGC, 2005) by providing additional services for integrating web-connected sensors and sensor systems. SWE services are designed to enable *discovery* of sensor assets and capabilities, *access* to these resources through data retrieval and subscription to alerts, and *tasking* of sensors to control observations (Botts et al., 2006).

2.1 SWE Information Model

The SWE initiative has developed draft specifications for modelling sensors and sensor systems (SensorML, TransducerML), observations from such systems (Observations and Measurements) and processing chains to process observations (SensorML) (Botts, 2005, Cox, 2005). The draft specifications provide semantics for constructing machine-readable descriptions of data, encodings and values, and are designed to improve prospects for plug and play sensors, data fusion, common data processing engines, automated discovery of sensors, and utilisation of sensor data.

2.2 SWE Services Model

SWE provides four types of web services: Sensor Observation Service (SOS) (Na and Priest, 2006), Sensor Alert Service (SAS) (Simonis, 2006), Sensor Planning Service (SPS) (Simonis, 2005) and Web Notification Service (WNS) (Simonis and Wytzisk, 2003). The SOS provides a standard interface that allows users to retrieve raw or processed observations from different sensors, sensor systems and observation archives. The SAS provides a mechanism for posting raw or processed observations from sensors, process chains or other data providers (including a SOS) based on user-specified alert/filter conditions. When subscribing to a SAS, users not only define the alert conditions but also the communication protocol for disseminating alerts via the WNS.

The WNS provides a standard interface to allow asynchronous communication between users and services and between different services. A WNS is typically used to receive messages from a SAS and to send/receive messages to and from a SPS. The SPS provides a standard interface to sensors and sensor systems and is used to coordinate the collection, processing, archiving and distribution of sensor observations. Discovery of OGC and SWE services is facilitated by the Sensor Web Registry Service – an extended version of the OGC Catalog Service.

3 THE ADVANCED FIRE INFORMATION SYSTEM

The Advanced Fire Information System (AFIS) is the first near real-time satellite-based fire monitoring system in Africa. It was originally developed for the South African electrical power utility, ESKOM to mitigate the impact of wild fires on regional electricity supply (Fleming et al., 2005, Frost and Vosloo, 2006). AFIS was first implemented using propriety GIS technology but has now been re-engineered as an OGC compliant Sensor Web application based on open source software.

3.1 Hotspot Detection

AFIS currently relies on a contextual algorithm for hotspot detection using the MODIS sensor aboard the polar orbiting TERRA and AQUA satellites and the SEVERI sensor aboard the geostationary METEOSAT-8 satellite (CEOS, 2002). The hotspot update rate for MODIS is every six hours compared to every 15 minutes for SEVERI. Though the SEVERI provides almost near real-time hotspot detection, it can only resolve large hotspots (five hectares or more in extent) unlike MODIS, which can resolve hotspots less than a hectare in size.

The hotspot detection algorithm was originally developed for the AVHRR sensor (Flasse and Ceccato, 1996, Giglio et al., 2003). The algorithm uses 3.9μm and 10.8μm bands to discriminate fire pixels from background pixels. The algorithm first classifies a pixel according to a fixed threshold, *e.g.* $T > 300K$, to identify potential fire pixels – the remaining pixels are called *background* pixels. The neighbourhood of this pixel is then searched for background pixels, growing the neighbourhood if necessary to ensure that at least 25% of the neighbourhood pixels are background pixels. From this set of background pixels, the mean and standard deviation statistics are calculated from the 3.9μm and the 3.9μm $- 10.8\mu$m band difference data. The pixel under consideration is then classified as a hotspot if its 3.9μm value exceeds the background mean by some multiple of the standard deviation – a similar test is performed on the 3.9μm $- 10.8\mu$m band difference.

3.2 Current Sensor Web Architecture

Figure 2 depicts the current OGC compliant architecture for AFIS. The contextual algorithm for hotspot detection has been implemented in dedicated image processing chains for MODIS and SEVERI observations. The processing chains generate hotspots using pre-processed MODIS and SEVERI data retrieved from an image data store via a SOS. Hotspot data is sent to the AFIS Application Server for archiving and spatial processing. End-users are able to query the hotspot archive via a high-level SOS and can subscribe to a SAS that uses the output of the spatial processing to generate fire alerts. At present, spatial processing is limited to intersecting hotspot events with features of interest. However, the aim is to enrich fire alerts by populating hotspot events with additional attribute data such as surface wind vectors and fire danger index provided by other OGC web services *e.g.* Web Coverage Service (WCS) and Web Feature Service (WFS). When subscribing to the SAS, users must specify the preferred medium for receiving fire alerts (*e.g.* Simple Message Service (SMS) or Email) and what parameters to pass through to the spatial process chain. Parameters may include what features of interest to intersect with hotspots (in the case of ESKOM, this would be buffer zones around high-voltage transmission lines).

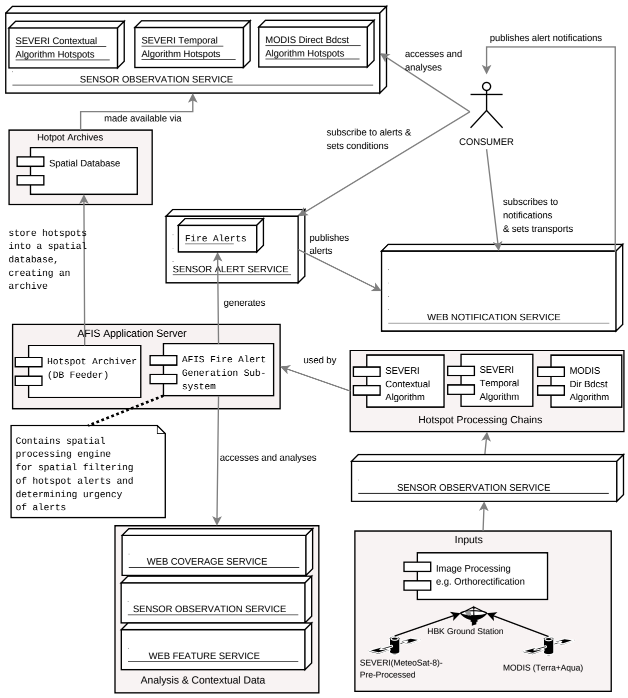

Figure 2: OGC Compliant AFIS Architecture

4 RESULTS AND DISCUSSION

4.1 Hotspot Detection Success Rate

In the first two years of operation, AFIS detected 44% and 46% of the fires that disrupted electricity supply with MODIS and SEVERI respectively. The combined detection succession rate was 60% (Frost and Vosloo, 2006). Most of these fires were detected well in advance of the incidents giving ESKOM controllers ample time to respond to this threat.

The detection success rate is considered too low so we are implementing a non-contextual hotspot detection algorithm for the SEVERI sensor that is more sensitive. The basic approach is to build a general model of the diurnal cycle for the 3.9μm band, and then to fit this model to the observed data of the last 24 hours. This model can then be used to generate accurate estimates of the expected background temperatures. If a statistically significant difference between the current observed temperature and the predicted background temperature is observed, then the pixel in question is classified as a hotspot. The first implementation of this algorithm relied on a Kalman filter to provide the estimates of the background temperature. Initial results indicate that this method is significantly more sensitive, particularly in cases where the background temperature is below 300K (*e.g.* early morning)(van den Bergh and Frost, 2005).

4.2 Extending AFIS Functionality

The intention is to shift the emphasis from simple fire detection to more sophisticated fire risk management. This requires a good understanding of what controls wild fire behaviour. We are currently building a domain ontology (Gruber, 1993, Struder et al., 1998) for wild fires. The ontology will capture key concepts in the wild fire domain such as combustion properties, fuel load, burning regime, fire weather, fire suppression methods, and topographical controls. The aim is to use the Sensor Web to observe specific fire-related phenomena described in the wild fire ontology and employ machine reasoning to determine fire risk *i.e.* automate the *observe* and *orient* parts of the OODA loop and issue more useful fire alerts.

4.3 Sensor Web Agent Platform

To achieve the desired level of automation requires a more intelligent architecture than what the current SWE framework provides. We are advocating an open, service-oriented, multi-agent system architecture for the Sensor Web known as the Sensor Web Agent Platform (SWAP). This architecture is a hybrid of the Foundation for Intelligent Physical Agents (FIPA)(FIPA, 2002) and Open Geospatial Consortium (OGC) standard architectures. SWAP incorporates the following concepts: Ontologies, process models, choreography, service directories and facilitators, service-level agreements and quality of service measures (Huhns et al., 2005). Ontologies will provide explicit descriptions of all components within SWAP i.e. sensors and sensor data, simulation models, algorithms and applications, and how these components can be integrated and used by software agents.

Users can improve or alter the behaviour of SWAP by editing any of the underlying ontologies. SWAP uses the Web Ontology Language (OWL) as the ontology representation language (Bechhofer et al., 2004). The ideal is to re-use existing ontologies wherever possible. We are using the Suggested Upper Merged Ontology (SUMO) to ground all the ontologies in SWAP (Niles and Pease, 2001). In the case of sensors and sensor systems, we are using OntoSensor (Russomanno et al., 2005) and the NASA Semantic Web for Earth and Environmental Terminology (SWEET) ontologies (Raskin, 2006) for building domain and application ontologies.

The SWAP abstract architecture is split into three layers: *Sensor Layer*, *Coordination Layer* and *Decision Layer* (Figure 3). The Sensor Layer is populated by sensor agents that encapsulate individual sensors, sensor systems and archived observations. They expose sensor data in a uniform way and deal with any sensor-dependant processing. Data from sensor agents form input to agents in the second, Coordination Layer, which consist of work flow, tool and simulation agents. Work flow agents receive data from sensor agents and pass this data through a combination of tool and simulation agents and aggregate the results. Tool agents provide feature extraction and image processing functionality, while simulation agents store real-world models and can provide projections and analysis of data. The processed data stored by work flow agents form input to application agents in the Decision Layer. Application agents combine higher level features provided by work flow agents and provide different views of this data to different end users and advanced users would be allowed to compose and deploy new work flow agents. Application agents would also allow users to specify alerts in the system.

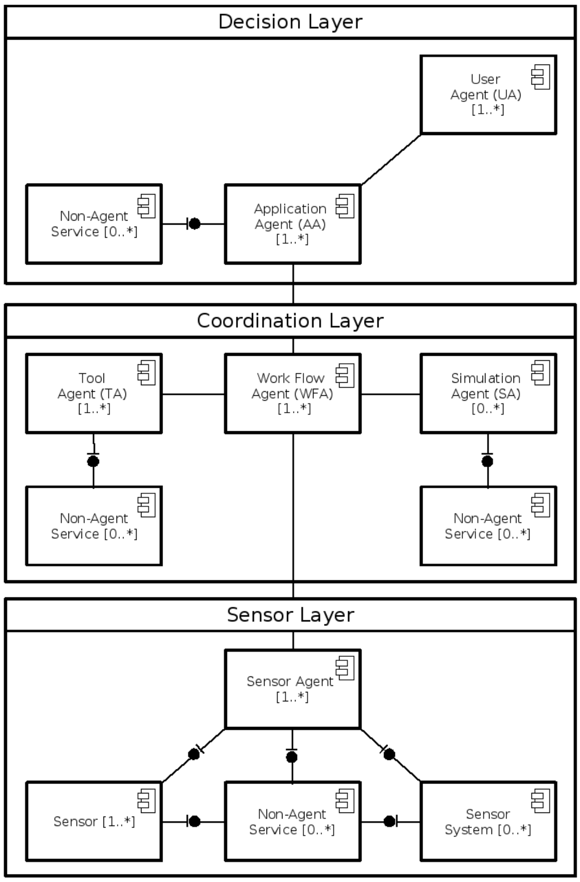

Figure 3: SWAP Abstract Architecture

A prototype implementation of SWAP is currently being developed for AFIS (Figure 4). The prototype will use the non-contextual hotspot algorithm described earlier. A sensor agent will access a SOS offering SEVERI blackbody temperature and expose this to a work flow agent configured to detect hotspots. The work flow agent is responsible for tasking a simulation agent to predict the current background temperature from the diurnal cycle, sending both the predicted background temperature and observed temperature to a tool agent that checks for hotspots, and communicating detected hotspots to the AFIS application server. The simulation agent will retrieve historical SEVERI blackbody temperature needed for building the diurnal cycle from an image data store via the same SOS used by the sensor agent. The AFIS application server (an application agent) uses hotspots retrieved from the workflow agent to issue fire alerts using a SAS and WNS. The AFIS application server can be accessed either via an AFIS client installed on a a user's computer, or via a web interface for configuring fire alerts.

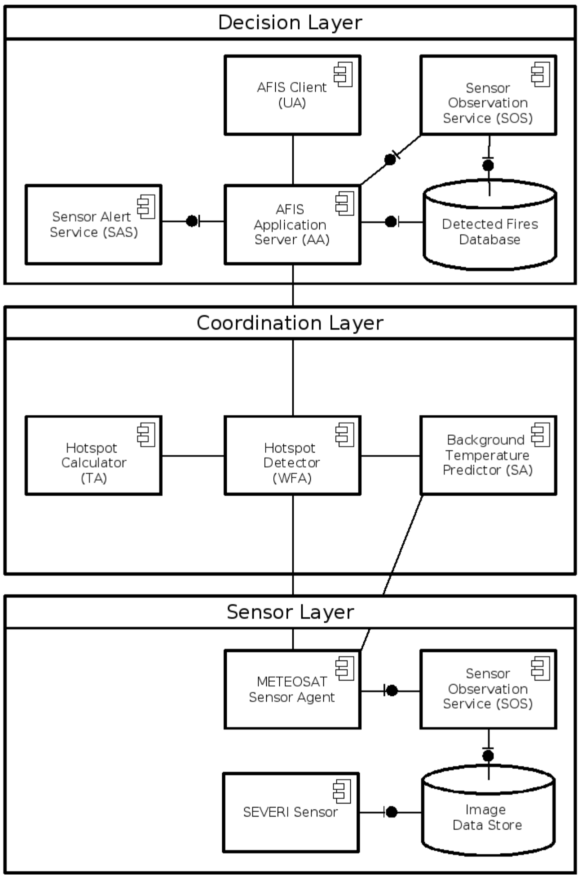

Figure 4: SWAP Architecture for AFIS Prototype

5 CONCLUSIONS AND FUTURE WORK

AFIS is a good example of what the Sensor Web can do to facilitate effective disaster response. The Sensor Web enhances the OODA loop by providing several mechanisms for sensor-rich feedback control. Our proposed service-oriented, multiagent systems architecture for the Sensor Web extends the current SWE framework. SWAP facilitates technical and semantic interoperability and promotes re-usability. The architecture is flexible and extensible: New agents can be deployed or swapped out in a 'plug and play' fashion. Work flows can be easily updated or built from scratch by editing existing or creating new application ontologies. Parts of the OODA loop can be offloaded onto software agents that use machine-reasoning to automatically generate hypotheses. Speeding up the OODA cycle in this way should enhance the tempo of disaster response.

A number of research questions must be addressed before SWAP can be fully realised. These include questions relating to the internal model of agents, communication between agents and between agents and non-agent services, message structure and message payload structure, framework for building ontologies, how to handle contradictory knowledge, how to integrate different types of ontologies into the agent paradigm, maintenance of ontologies, data fusion, dynamic configuration of process chains and appropriate agent development framework.

We intend using the implementation of the SWAP prototype for AFIS to refine these research questions and expose others. Our plan is to work in close collaboration with standards generating bodies such as OGC, FIPA, IEEE and W3C, other research partners and the Open Source Software (OSS) development community to ensure SWAP becomes a future standard open architecture for the Sensor Web.

ACKNOWLEDGEMENTS

AFIS is an initiative of the Meraka Institute (part of the Council for Scientific and Industrial Research (CSIR)). ESKOM funded the initial development of AFIS. The re-engineering of AFIS as an OGC compliant platform and research and development of SWAP is being funded by the CSIR. SWAP is the brainchild of a global research partnership known as the Sensor Web Alliance (www.sensorweb-alliance.org). Research partners include the ICT for Earth Observation Research Group at the Meraka Institute; Institute for Geoinformatics at the University of Muenster; School of Computer Science at the University of Kwazulu-Natal; Bagley College of Engineering at Mississippi State University; TNO Defence, Safety and Security; and 52°North.

REFERENCES

Annoni, A., Atkinson, M., Denzer, R., Hecht, L., Millot, M., Pichler, G., Sassen, A.-M., Couturier, M., Alegre, C., Sassier, H., Coene, Y. and Marchetti, P. G., 2005. Towards an Open Disaster Risk Management Service Architecture for INSPIRE and GMES. Version No. 9.

Bechhofer, S., van Harmelen, F., Hendler, J., Horrocks, I., McGuinness, D. L., Patel-Schneider, P. F. and Stein, L. A., 2004. OWL Web Ontology Language Reference. `http://www.w3.org/TR/2004/REC-owl-ref-20040210/`, Accessed 11 July 2006.

Botts, M., 2005. OpenGIS Sensor Model Language (SensorML). Document 05-086, Wayland, Massachussets, USA. `http://portal.opengeospatial.org/files/?artifact_id=12606`, Accessed Febuary 2006.

Botts, M., Robin, A., Davidson, J. and Simonis, I., 2006. OpenGIS Sensor Web Enablement Architecture Document. Document 06-021r1, Wayland, Massachussets, USA. `http://portal.opengeospatial.org/files/?artifact_id=14140`, Accessed 4 July 2006.

CEOS, 2002. Earth Observation Handbook. Committee on Earth Observation Satellites. `http://www.ceos.org/pages/pub.html#handbook`.

Cox, S., 2005. Observations and Measurements. Document 05-087r5, Wayland, Massachussets, USA. `http://portal.opengeospatial.org/files/?artifact_id=14034`, Accessed 5 May 2006.

FIPA, 2002. FIPA Abstract Architecture Specification. Technical Specification SC00001L, Foundation for Intelligent Physical Agents, Geneva, Switzerland. `http://www.fipa.org`.

Flasse, S. P. and Ceccato, P., 1996. A Contextual Algorithm for AVHRR Fire Detection. International Journal of Remote Sensing 17(2), pp. 419–424.

Fleming, G., van den Bergh, F., Claudel, F. and Frost, P., 2005. Sensor Web Enabling the Advanced Fire Information System. In: Proceedings of the 2005 International Symposium for Environmental Software Systems (ISESS 2005), Hotel do Mar, Sesimbra, Portugal. `http://www.isess.org/documents/2005/presentations/`, Accessed 4 July 2006.

Frost, P. and Vosloo, H., 2006. Providing Satellite-Based Early Warnings of Fires to Reduce Fire Flashovers on South African Transmission Lines. In: Proceedings of the 10th Biennial Australasian Bushfire Conference: Bushfire 2006, ICMS. `http://www.fireandbiodiversity.org.au/`.

Giglio, L., Descloitres, J., Justice, C. and Kaufmann, J., 2003. An Enhanced Contextual Fire Detection Algorithm for MODIS. Remote Sensing of Environment 87, pp. 273–282.

Gruber, T. R., 1993. Towards Principles for the Design of Ontologies Used for Knowledge Sharing. Kluwer Academic Publishers, Deventer, The Netherlands. citeseer.ist.psu.edu/gruber93toward.html.

Huhns, M., Singh, M., Burstein, M., Decker, K.and Durfee, K., Finin, T., Gasser, T., Goradia, H., Jennings, P., Lakkaraju, K., Nakashima, H., Van Dyke Parunak, H., Rosenschein, J., Ruvinsky, A., Sukthankar, G., Swarup, S., Sycara, K., Tambe, M., Wagner, T. and Zavafa, L., 2005. Research Directions for Service-Oriented Multiagent Systems. IEEE Internet Computing pp. 65–70.

Liang, S. H. L. C. and Tao, C. V. ., 2005. A Distributed Geospatial Infrastructure for the Sensor Web. Computers & Geosciences 31(2), pp. 221–231.

Na, A. and Priest, M., 2006. Sensor Observation Service. Document 05-088r1, Wayland, Massachussets, USA. http://portal.opengeospatial.org/files/?artifact_id=12846, Accessed 1 April 2006.

Niles, I. and Pease, A., 2001. Towards a Standard Upper Ontology. In: C. Welty and B. Smith (eds), Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS-2001), Ogunquit, Maine.

OGC, 2005. The OGC Abstract Specification - Topic 0: Abstract Specification Overview. Document 04-084, Open Geospatial Consortium, Wayland, Massachussets, USA. http://portal.opengeospatial.org/files/?artifact_id=7560, Accessed 14 April 2006.

Osinga, F., 2005. Science, Strategy and War: The Strategic Theory of John Boyd. Eburon Academic Publishers, Delft, Netherlands.

Raskin, R., 2006. Guide to SWEET Ontologies for Earth System Science. Technical report. http://sweet.jpl.nasa.gov/guide.doc, Accessed 11 July 2006.

Rosen, J., Grigg, E., Lanier, J., McGrath, S., Lillibride, S., Sargent, D. and Koop, C. E., 2002. The Future of Command and Control for Disaster Response. IEEE Engineering in Medicine and Biology 21(5), pp. 56–68.

Russomanno, D. J., Kothari, C. and Thomas, O., 2005. Building a Sensor Ontology: A Practical Approach Leveraging ISO and OGC Models. In: Proceedings of the 2005 International Conference on Artificial Intelligence, Las vegas, Nevada.

Simonis, I., 2005. Sensor Planning Service. Document 05-089r1, Wayland, Massachussets, USA. http://portal.opengeospatial.org/files/?artifact_id=12971, Accessed 1 April 2006.

Simonis, I., 2006. Sensor Alert Service. Document 06-028, Wayland, Massachussets, USA. http://portal.opengeospatial.org/files/?artifact_id=13921, Accessed 4 July 2006.

Simonis, I. and Wytzisk, A., 2003. Web Notification Service. Document 03-008r2, Wayland, Massachussets, USA. http://portal.opengeospatial.org/files/?artifact_id=1367, Accessed 6 Febuary 2006.

Struder, R., Benjamins, V. R. and Fensel, D., 1998. Knowledge Engineering: Principles and Methods. IEEE Transactions on Data and Knowledge Engineering 25(1), pp. 161–197.

van den Bergh, F. and Frost, P., 2005. A Multi-Temporal Approach to Fire Detection using MSG Data. In: Proceedings of Multitemp05, Biloxi, Mississippi, USA.

Zibikowski, R., 2004. Sensor-Rich Feedback Control: A New Paradigm for Flight Control Inspired by Insect Agility. Instrumentation and Measurement Magazine 7(3), pp. 19–26.