

The Semantic Reef: Managing Complex Knowledge to Predict Coral Bleaching on the Great Barrier Reef

Trina S. Myers, Ian M. Atkinson and William J. Lavery

School of Maths, Physics and Information Technology
James Cook University
Townsville, Queensland

(Trina.Myers, Ian.Atkinson, Bill.Lavery)@jcu.edu.au

Abstract

A semantically driven Virtual Organisation (VO) model for predicting important events for coral reefs is presented – the Semantic Reef. The model is an application of a Semantic Grid that encapsulates services and resources to produce a complex higher-level knowledge base in order to predict coral bleaching on the Great Barrier Reef. The goal is to enable the synergistic interplay between different e-Research tools to produce a model that can evaluate complex hypotheses. Built on a Grid infrastructure and invoking an ontology to map to currently non-communicative datasets, the Semantic Reef model is designed to address specific scientific problems such as the prediction of coral bleaching events. The outcome is not only a model for exploring a diverse range of e-Research challenges, but also potentially a re-usable interoperable knowledge base can then be accessed as a resource in other VOs.

Keywords: Semantic Grid, ontology, knowledge base, OWL.

1 Introduction

The research process often seeks to synthesize new knowledge by building on existing but disparate knowledge, and to do so in a well managed and well structured way. In practise, this process can be prohibitively difficult when the existing knowledge resources are unstructured and unrelated, and the operating processes and structures, management systems and information systems are inflexible and non-communicative. To address these issues, we propose to structure a research community's operational processes and information as well as computational resources on the principles of a semantically driven Virtual Organization (VO). The aim is to create a model to manage a complex

knowledge base for hypothetically driven research in ecology.

An example implementation of this model is focused on knowledge from the domain of coral reef ecologies. The Great Barrier Reef (GBR) is a World Heritage region that is one of Queensland's largest and most valuable environmental assets generating annually \$5.8 billion gross value for the Australian economy (Access-Economics, 2005). Unfortunately, due to increasing sea temperatures, coral bleaching poses a serious threat to the long term sustainability of the GBR (Berkelmans et al., 2004). An early warning system for predicting likely occurrences of coral bleaching *may* allow preventative and corrective measures to be taken.

The Semantic Reef Project aims to utilise existing coral reef databases augmented by real time sensor data to predict the possibility of such coral bleaching occurrences. It will map static and dynamic data to an ecosystem ontology that will make the information machine-understandable, enabling intelligent decisions based on inference rules and description logics.

Following a brief examination of the actual coral bleaching process, this paper outlines the proposed architecture for the model, and then describes the mechanics involved in developing the Semantic Reef model.

2 The Bleaching Problem

The GBR is the largest coral reef system in the world, covering an area of 348,000 square kilometres and spanning ~2300 kilometres of Queensland's coast from the northern tip to Bundaberg. The GBR is a large contributor to Australia's economy supporting fishing and tourism activities. The Great Barrier Reef is sometimes referred to as the single largest organism in the world. In reality it is made up of many billions of tiny coral formations (GBRMPA, 2006). The impact of global warming, specifically the rise in ocean temperatures, is an issue that threatens the viability of the GBR, as it results in the phenomenon of coral bleaching.

Coral bleaching is a stress condition in reef corals that involves a breakdown of the symbiotic relationship between corals and unicellular algae (zooxanthellae). This symbiotic relationship is essential to both parties survival as the zooxanthellae reside in every cell of the coral animal's tissue providing food for the coral. Like



Figure 1 – Coral bleaching - Photo by Ray Berkelmans, AIMS.

all forms of plant life, the zooxanthellae survive by photosynthesis producing energy-rich compounds and some are passed on to the coral.

Reef corals are very sensitive to sea temperatures outside their normal range. Elevated temperatures of 1⁰ Celsius above the long term monthly summer averages are enough to cause the stress factors that result in coral bleaching in many dominant coral species. When temperatures exceed threshold levels for long enough, the symbiotic relationship between the zooxanthellae and the corals breaks down. Energy from the sun, that is normally used to produce food, begins producing oxygen radicals that are highly corrosive and damage both the zooxanthellae and the coral. Ultimately, the coral expunge the zooxanthellae, and bleaching results. These algae give corals their characteristic brownish colour and once they have been expelled, the white skeleton shows through a coral's transparent tissue, giving it a bleached white appearance as shown in Figure 1 (Jones et al., 1998).

If stressful conditions prevail for long enough, the corals bleach and die. However, if stressful conditions abate, then the bleached corals can recover their symbiotic algae and return to their normal, healthy colour. The severity of bleaching can vary substantially according to water depth, location and species of corals (GBRMPA, 2006).

3 The Architecture

The model employs a range of existing technologies including agent-oriented technologies, Grid computing and the Semantic Web technologies to allow for virtually organized operation and the management of knowledge bases. The combination of these technological communities creates a synergy to address separate problems in the model.

The architecture consists of a bottom-up hierarchy that begins with an agent invoking services, which will in-turn access the lowest level raw resources (e.g. raw data, image repositories, compute cycles, etc). The agent will then map the disparate data to an ontology for hypothesis testing.

This base level architecture, shown in figure 2, depicts the hierarchical requirements to answering an initial

hypothetical question. This will entail identifying the goal and purpose of the project and accessing the necessary resources and services. In defining the goals, the exact questions (hypotheses) will come from collaborating with domain experts, in this case coral reef ecologists and marine biologists.

An agent can take advantage of services that can promote or discover resources, and then distribute them to the requestors (Jennings, 2001). Agent-based technologies are adaptive programs that perform a process or invoke services, typically based on artificial intelligence and often in dynamically changing environments. For example, autonomically mapping a dataset to populate a class with instances within an ontology, as the data is streaming in real time directly from the source (e.g. reef sensor data). The useful information that results from any data processing will become a possible reusable resource to foster an increased participation in collaboration.




Figure 2 –Components of the base level of the Semantic Reef.

3.1 The Grid Infrastructure

Building the model on a grid foundation provides a strong interoperable infrastructure and the tools for secure and reliable resource sharing. The Grid computing paradigm allows for the creation and maintenance of VOs and ultimately a virtual marketplace in which multiple VOs will be able to exchange commodities or collaborate in order to solve problems (Foster et al., 2001). Grid technologies allow for the decentralised management of resources that can be simultaneously accessed from a number of geographically separate locations.

The Semantic Reef will utilise an underlying Grid infrastructure to create a VO whose members are the various stakeholders that maintain, or own, the datasets required. The initial parties include the James Cook University¹ (JCU) precinct for sensor data, the Australian Institute of Marine Science² (AIMS) for spatial and environmental data, and the Great Barrier Reef Marine

¹ www.jcu.edu.au

² www.aims.gov.au

Park Authority (GBRMPA)³ for domain specific reef data and information. The Grid Middleware will provide the capabilities required for these distributed data resources to be of maximum benefit, such as resource accessibility, management and processing. The Grid foundation will address many other issues posed by the complexity of this project, such as security services, data movement controls, resource monitoring and discovery services, and many other available Web and Grid services needed to form this VO. Some services, such as workflow, discovery and composition, data streaming for large volumes of live data are not required at this stage of the project. In future works, the integration of data collected from sensor technology will be essential to the main goal of the Semantic Reef, namely, predicting coral bleaching events. Using the real-time sensor data will require the additional computational power made accessible through Grid technologies. This is being addressed in the DART⁴ project, as well as several other programs such as the integration of distributed sensor data in the Environmental Sensor Networks (ESNs) (Hart and Martinez, 2006).

Grid technologies alone will not be enough to handle the range of issues involved in this trial, for instance it will enable accessing data and storage but not the semantic integration of that data (Foster et al., 2004). The knowledge of the *semantics* (i.e. meaning) of the data and operations are not expressed explicitly in machine-understandable form, it is only explicitly hard-coded into the programs (Uschold, 2003). At present, the machine (i.e. a computer or computer program, such as a software agent, that performs tasks on the web or within a grid infrastructure) can not understand information because the information has no well-defined meaning. In order to allow the machine to make judgments followed by informed decisions on the data it is processing, we need to incorporate Semantic Web technologies to make the dataset inputs computer-comprehensible (De Roure et al., 2005).

The Open Grid Services Architecture (OGSA) based grid infrastructure allows interoperable and uniform access to all grid services (Foster et al., 2002). The Semantic Grid is an extension of the Grid where applying semantically rich information to current Grid resources can create a more intelligent Grid Service (Goble and Bechhofer, 2005). A proposal for a Semantic Grid Reference Architecture, Semantic-OGSA (S-OGSA), is being developed (Corcho et al., 2006). The S-OGSA defines a model that extends the OGSA via lightweight mechanisms to incorporate semantic and knowledge services.

The OntoGrid project⁵ is an eight-partner EU FP6 undertaking to investigate fundamental issues in Semantic Grids and is an implementation of the S-OGSA (Goble and Bechhofer, 2005). One of the case studies within the

OntoGrid project involves satellite data management – Quality Analysis of Satellite Missions (QUARC) and has many commonalities with the Semantic Reef Project. The QUARC system involves complex processes where disparate data belonging to different autonomous systems must be queried, processed and transferred (Sánchez-Gestido et al., 2006). The challenges and lessons learned from the QUARC project will be an invaluable example to use throughout the development of the Semantic Reef.

3.2 The Semantic Web Component

Using cutting edge Semantic Web technologies allows for the semantic integration of knowledge, information and data. These technologies are designed to share and process data independent of the programs or formats they eventuated from by using well-defined ontologies to give meaning to the data (Goble and De Roure, 2004).

The Web Ontology Language (OWL) (McGuinness and van Harmelen, 2004) is a new formal language for representing ontologies in Semantic Web applications, developed by the World Wide Web Consortium (W3C) Web Ontology Working Group.

As stated by McGuinness and van Harmelen (2004) OWL is being designed “in order to provide a language that can be used for applications that need to understand the content of information instead of just understanding the human-readable presentation of content”. The Semantic Reef Project will utilise OWL, as it is the adopted standard for engineering ontologies and can be built using the Protégé-OWL API (Knublauch et al., 2004), which includes a Semantic Web Rules Language (SWRL) plugin for applying inference rules that can bridge to a rule engine such as Jess (O'Connor et al., 2005). Reasoning over the ontology with Description Logic (DL) rules can be achieved using a classifier such as Racer PRO (RacerPRO, 2006) or FaCT (FaCT, 2006).

A number of steps are involved in realising the overall Semantic Reef Project, these can be summarised as follows:

- a) Define the actual questions for the Semantic Reef that will be the outcome of the correlation. E.g. a meteorological hypothesis – Can coral bleaching be predicted, by location, by coral species and by sea surface temperatures (SST)? Are there particular locations prone to bleaching due to foreseeable, predictable temperature changes given temporal data (e.g. time of year, time of day, etc)? This will entail semantic descriptions of the coral ecosystem and descriptions of the state of the environment. It could involve a number of data sets (e.g. SST, sensor data of water temperatures at depth, coral types, depth and location).
- b) Decide on particular locations of the GBR that are well studied with the background datasets readily available.
- c) Build the ontologies for relevant domains. The reef ontology will require defining the sets of terms and the taxonomy, the objects and classes and the relationships between them. Asserting the axioms

³ www.gbrmpa.gov.au

⁴ dart.edu.au

⁵ www.ontogrid.net

that add constraints and restrictions on how these entities interact

- d) Define the relationships between all the concepts involved in the coral bleaching process and forming queries using description logics and clearly defined inference rules. For example, IF a hard coral is of type hermatypic, AND the sea surface temperature is high AND water depth is less then a prescribed amount, it will infer a risk of bleaching, and the possible response may be to monitor that location.
- e) Test the accuracy of the ontologies by doing a comparison analysis using historical data sets and known knowledge to validate results.
- f) Architect a model for continuous correlation of the incoming live datasets in conjunction with inference rules and the semantic grid ontologies to produce an answer to the hypothesis.

4 The Description Logic and Inference Rules of the Reef Ontology

In order to understand how the use of description logics and inference rules can give the machine the ability to infer information about relationships between individual data objects in the ontology one must understand Semantic search techniques, which involve Open World Assumption (OWA), the Unique Name Assumption (UNA) and Axioms.

4.1 Open World Assumption

Most DL systems (e.g. RacerPro, FaCT, etc) use OWA for reasoning. OWA means what cannot be proven to be true is not automatically false (Horrocks et al., 2003), it assumes that its knowledge of the world is incomplete. In the OWA, what is not stated is considered unknown, rather than wrong; it simply assumes the extra information needed has not been added to the knowledge




Figure 3 –IS_A relationships in a segment of the Semantic Reef ontology.

base (Rector et al., 2004). For example, if the statement ‘carnivores eat some herbivores and some mammals’ was added as a restriction in the Semantic Reef Ontology, the query ‘do carnivores eat seagrass?’ will return an ‘unknown’. In a Closed World Assumption (used predominantly with relational databases), where it is assumed the information that is there is everything, the answer would be ‘false’. However, in OWA, because it has not been explicitly stated that carnivores do not eat seagrass it is assumed there is a possibility they do because there is not enough known to suppose otherwise (also, some carnivores are omnivorous). The open world assumption is considered implicit in RDF and OWL, as every tuple not explicitly contained in the ontology is implicitly assumed to represent a fact that is unknown, rather than false.

4.2 Unique Name Assumption

Another significant factor in the reasoning process is the fact that OWL does not employ the UNA, which means in OWL two different names could actually refer to the same individual. This is an important aspect when querying the ontology across a variety of different data stores, what might be named hermatypic in one repository might be named zooxanthellate in another, both referring to the symbiotic relationship between the reef building coral and the algae zooxanthellae

4.3 Axioms

A Description Logic knowledge base consists of sets of axioms, or statements of truisms. Axioms can infer automatically one class is a subclass of another, or that an individual is an instance of an inferred class as well as its asserted class. An axiom is a sentence or proposition that is taken for granted as true, and serves as a starting point for deducing other truths (Wikipedia, 2006). The axioms most common in OWL are disjoint, class, domain and range and closure axioms:

4.3.1 Disjoint Axioms

All OWL classes are assumed to overlap, i.e. an individual of one class can also belong to another class simultaneously, unless explicitly stated that these classes are disjoint from each other – an individual from one cannot possibly belong to the disjoint class. In the Semantic Reef all sibling classes are declared as disjoint except those that may share individuals, for example, carnivorous crustaceans may have the same individuals as herbivorous crustaceans if they are omnivorous. Conversely, if it were not explicitly stated that individuals from the ‘Algae’ class can not possibly belong to the ‘Coral’ class they could overlap and create incorrect inferences by the Reasoner.

OWL utilizes *is a* relationships and disjoint relationships to allow the Reasoner to determine which individuals can belong to which classes. OWL has an implication that ‘everything’ *is a* ‘everything else’ unless explicitly stated with a disjoint axiom. Also, every subclass should have an explicit ‘*is a*’ relationship with its superclass. A classified version of a segment of the Semantic Reef

Ontology, showing all ‘is a’ relationships within that segment, is depicted in Figure 3.

4.3.2 Class Axioms

There are two types of classes in OWL; Primitive and Defined, both having conditions that determine membership to a particular class. Primitive classes are described by *necessary* conditions, that is, it is necessary to meet the condition to be a member of this class. A defined class is one that is described by at least one *necessary and sufficient* condition, these statements signify that if an entity fulfils these *necessary and sufficient* conditions then it must be, by default, a member of that class (Rector, 2003).

Being able to use a Reasoner to automatically compute the class hierarchy is a major advantage of building an ontology using the OWL Description Logics (OWL-DL) sub-language. The job of the Reasoner not only includes dynamically computing and maintaining multiple inheritances but also automatically discovering anonymous classes. As opposed to the asserted classes that are manually defined, anonymous classes are inferred by the logic reasoning of the inference rules, axioms and restrictions. In order to take advantage of these benefits of using a Reasoner the ontology must use OWL-DL and have Defined Classes, as Primitive classes will not be classified.

In the Semantic Reef Ontology, many of the first level of classes can be defined with *necessary and sufficient* conditions. For example, by definition a herbivore is any entity that, among other things, eats plant life. It is *necessary* that they eat plant life, but it is also *sufficient* to belong to the Herbivore class if an individual from any other class eats plant life. So the OWL statement:

Class Herbivores Defined
restriction (eats \exists someValuesFrom
(Phytoplankton OR Algae OR Coral))

Will result in any class or individual that fits these conditions, will be classified a member of the Herbivore class.

4.3.3 Domain and Range Axioms

Domain and Range Axioms, also known as ‘global axioms’ as they are not set on particular classes but instead are specified for object properties, that link individuals from one class to individuals from another class. In other words, properties link individuals from the domain to individuals from the range. For example, the property ‘is_eaten_by’ in the reef ontology would link individuals belonging to, among others, the class ‘Herbivores’ to individuals belonging to, among others, the class ‘Carnivores’. In this case the domain is ‘Herbivores’ and the range is ‘Carnivores’ for the ‘is_eaten_by’ object property.

4.3.4 Closure Axioms

Because OWL uses the Open World Assumption, descriptions of classes should be ‘closed off’ where appropriate, these are known as closure axioms. Closure

axioms are a way of disambiguating a concept, leaving no opportunity for a wrong assumption. For example, when paraphrased, one could describe the concept of a carnivore, for the sake of the reef ontology, as “a carnivore is an animal that, *among other things*, eats *some* herbivores and also *some* carnivores and eats *only* herbivores *and/or* carnivores”. This translates to the OWL restriction statement:

Class Carnivores eats \exists someValuesFrom
Herbivores
Class Carnivores eats \exists someValuesFrom
Carnivores
Class Carnivores eats \forall allValuesFrom
(Herbivores OR Carnivores)

OWL’s existential quantifier ‘someValuesFrom’ (some) and universal quantifier ‘allValuesFrom’ (only) property restrictions used in this example close off the possibility of further additions for a given property (Rector et al., 2004). To clarify the OWL meaning of these restrictions ‘someValuesFrom’ equates to ‘at least one value of the property must be of a certain type but others might exist’ whereas ‘allValuesFrom’ equates to ‘All values of the property must only be of a certain type or have null values’. Without the final closure axiom ‘allValuesFrom’ (only values from) it could be possible, due to OWA, that carnivores could eat plants because there are no statements to say otherwise and the OWA assumes if it is not there it is unknown, not false.

4.4 Querying the Ontology

When making queries of the ontology the ‘is a’




Figure 4 –Unclassified ‘tangled’ segment shows only manually asserted IS A relationships.




Figure 5 – The segment after the Reasoner classification - ‘untangled’; now showing the inferred subsumptions.

relationships that have been refined through the Reasoner will allow the Semantic technologies to find complex relationships and aspects that would be overlooked by both humans and relational databases (DB). This is because Semantic technologies do not use keywords as such; they allow the keyword to be used as a guideline to all implicit possibilities. For example, a simple query of a relational DB searching for Algae would return only directly asserted Algae objects (as shown in Figure 4), alternatively, when using semantic support a subclass of Plankton (i.e. Phytoplankton) would also be returned as it has been classified as having an inferred *is a* relationship with the superclass Algae. Figure 5 is the ‘untangled’, or normalised, version of the ontology after the classification process, which shows, among other implied relationships, Phytoplankton is an inferred subclass of Algae. The implied relationships are significant as they apply to all individuals that are members of the affected classes.

4.5 Inference Rules

Developing the inference rules is another powerful technique that works in conjunction with description logic. As mentioned, DL is a set of logical statements used in OWL DL to normalise classes and individuals by assumption and subsumption, it allows for negation/complement of classes, disjoint information and existential and universal quantification. The Semantic Web Rules Language (SWRL) manages inference using horn-like logic, which is a subset of predicate logic (first-order logic), and is orthogonal to description logic, for monotonic and non-monotonic rules.

A SWRL inference rule, which is based on the RuleML format, is atom centric and contains antecedents and consequences, or the *body* and *head* respectively. Where the antecedent (body) of the rule represents the information supplied in order to draw a conclusion, and the consequence (head) is the implication that is ultimately drawn. An SWRL rule has the form:

$$\text{Body}(X_1 \wedge X_2 \wedge X_n) \rightarrow \text{Head}(Y_1 \wedge Y_2 \wedge Y_n)$$

Both the body and the head can consist of conjunctions of atoms (X and Y). These atoms can be in the form of an OWL class, an OWL property or a declaration of *different from* or *same as* that refers to OWL individuals or OWL data values.

One SWRL rule to determine whether a bleaching warning is to occur would be:

$$\begin{aligned} &\text{Coral_Reef}(?z) \wedge \text{Coral}(?x) \wedge \text{hermatypic}(?x,\text{true}) \\ &\wedge \text{Sea_Surface_Temperature}(?y, ?z) \wedge \\ &\text{swrlb:greaterThan}(?y, 32) \rightarrow \text{bleaching}(?x,\text{true}) \end{aligned}$$

When translated the antecedent (body) states the conjunction between all factors included in the bleaching process: if an individual (x) in the class Coral is hermatypic and the Coral Reef (z) has an SST (y) above 32, will result in the consequence (head) that the ‘bleaching’ data property be true.

These SWRL rules will be utilised to infer areas where the information is either incomplete or requires the decisive regulation that predicate logic offers knowledge representation (KR). If a conclusion can be drawn and it will remain valid, even after new knowledge is formed, it is a monotonic rule. For example, if a coral is ahermatypic it will not bleach, which can be deduced by the fact that bleaching is caused by the symbiont relationship between the coral and the algae, if a coral is ahermatypic it does not have zooxanthellae, therefore cannot bleach. However, a monotonic or deductive logic cannot handle all reasoning issues because there are many occasions that not all knowledge of a concept is known.

Non-monotonic rules are true until new knowledge can prove otherwise. Also classed as defeasible logic, which can be ‘defeated’ by other rules because scientific reasoning is fallible, in many cases we can never be 100% conclusive in science as hypotheses can be proven wrong in the next experiment (Antoniou et al., 2001). Priorities, or an order of preference, can be applied to sets of rules, that is, if there is a contradiction in the consequence of two rules, one will take precedence over the other (Antoniou and van Harmelen, 2004). For example, if rule (R₁) defines a factor that makes up the bleaching process to be ‘SST is greater then 32 degrees’ and a second rule (R₂) states ‘SST is greater then the maximum summer average in a region’, R₂ can be given precedence over R₁ as a more desirable conclusion because it is more specific and is an exception to the general rule that is R₁.

Inference rules can be applied dynamically using a rule engine like Jess or Algernon and used to infer new knowledge from the existing OWL knowledge base (Horrocks et al., 2004).




Figure 6 – The Semantic Reef architecture

5 Testing the Semantic Reef Ontology

Creating sets of inference rules to allow the machine to determine whether bleaching has occurred will involve giving it explicit guidelines to characterise what bleaching actually means and what circumstances need to be in place to imply it might happen. Circumstances such as the temperature rising and characteristics such as the coral is of type hermatypic.

Initially we will create a simple trial to test the validity of the ontology and check the accuracy of the inference rules. Comparing the outcome to the 1998 and 2002 bleaching research will offer an element of ground validation where there is actual historical factual data and research available for reverse hypothesis testing. The preliminary analysis will be uncomplicated and use singular scale static SST spatial data, which will come from the same source as used in the previously mentioned comparison done by Berkelmans et al (2004), namely the spatial SST data taken from the radiometer sensors aboard the NOAA14 and NOAA16 satellites. At this stage, a basic approach is warranted in order to validate the accuracy of the ontology; therefore, a singular scale will be employed at a spatial resolution of 1 square kilometre matching available data sets. Choosing a particular location in which to directly test the logics and functionality of individual relationships in the ontology requires similar historical knowledge. The work and studies carried out on the 1998 and 2002 coral bleaching events focused around Davies Reef, among other areas. Davies Reef will be ideal for this experiment, as all relevant information in which to populate the ontology are pre-existing, such as the quantity and type of coral species and algae, etc.

The disparate data sets mentioned will be mapped to the ontology using an agent invoking an XML schema layer and XSLT to transform the data into XML form making the data independent of the originating databases. Once in XML form they can be mapped to the ontology either as individuals of a particular class (Coral class, Algae class, etc) or as a value of an OWL data type property. Once integrated with the ontology we allow the Reasoner to classify the individuals, setting the inferred relationships along with the asserted ones. Finally using the SWRL inference rules to query the ontology and find any positive bleaching occurrences, a comparison would then be made against the historical research for a match (or mismatch).

If the outcome of the hypothesis posed, when compared to the historical findings, proves correct, it will establish the reef ontology as a practicable working model (Figure 6).

6 Future Work and Conclusions

In this paper, we have explored the methodology of a Semantic Grid application for predicting coral bleaching. We are using a simple experiment initially to test the ontology itself, and to prove the accuracy of the prediction tool. In the simple test, only a limited number of factors are included in the query, specifically SST in correlation with location and time, as these can be directly compared to the findings of Berkelmans et al (2004). Using this ground validating method for testing the Reef Ontology will result in a prototype with quantifiable accuracy within the scope of this single hypothesis. Subsequently the prototype will be built on by integrating Grid services to form and manage VOs and Semantic Grid services to bind the Grid entities and resources with the knowledge services, entities and

resources (e.g. ontology service, reasoning service, semantic binding service, etc). These Semantic Grid services are part of the S-OGSA and are currently in an alpha stage of development and testing as part of the OntoGrid project (Corcho et al., 2006).

Although the stress factor most commonly associated with bleaching is elevated sea temperature there are a number of other causal factors. These additional stresses, such as high light intensity, low salinity and pollutants, are known to exacerbate coral bleaching (Hughes et al., 2003). In future work, once the accuracy of the ontology and inference rules is known, other factors that contribute to the bleaching phenomenon can be incorporated in the inference rules and the queries in order to produce increasingly useful and realistic predictions.

Upon further development of the Semantic Reef Project, we will see additions in the hypotheses incorporating different types of datasets, for example, live sensor data. The sensor data will not only add the benefit of information contained in a smaller area (i.e. 1 square metre instead of 1 square kilometre), but also be real-time instead of static historical snapshots. Only when using real-time streaming data (e.g. temperature, etc) and the compute power made available via Grid Middleware, will the true possibility of predicting a bleaching event become an actuality.

In conclusion, this paper describes the Semantic Reef Project and the methodology used to prove the theory – can coral bleaching be predicted using this semantically-driven hypothesis-testing model?

7 Acknowledgements

We thank Professor Jeremy Frey, Professor Carole Goble, Oscar Corcho and the OntoGrid team and Associate Professor Bruce Litow, for helpful discussions and useful feedback. We also thank the anonymous reviewers for their valuable comments and advise.

8 References

Access Economics Pty Limited (2005), *Measuring the Economic and Financial Value of The Great Barrier Reef*, Technical report, Great Barrier Reef Marine Park Authority, Townsville.

Antoniou, G., Billington, D., Governatori, G. and Maher, M. J. (2001): Representation results for defeasible logic. *ACM Trans. Comput. Logic*, **2**(2):255-287.

Antoniou, G. and van Harmelen, F. (2004): *A Semantic Web Primer (Cooperative Information Systems)*, The MIT Press.

Berkelmans, R., De'ath, G., Kininmonth, S. and Skirving, W. J. (2004): A comparison of the 1998 and 2002 coral bleaching events on the Great Barrier Reef: spatial correlation, patterns, and predictions. *Coral Reefs*, **23**(1):74-83.

Corcho, O., Alper, P., Kotsiopoulos, I., Missier, P., Bechhofer, S. and Goble, C. (2006): An overview of S-OGSA: A Reference Semantic Grid Architecture. *Web*

Semantics: Science, Services and Agents on the World Wide Web, **4**(2):102-115.

De Roure, D., Jennings, N. R. and Shadbolt, N. R. (2005): The Semantic Grid: past, present, and future. *Proceedings of the IEEE*, **93**(3):669-681.

FaCT: Fast Classification of Terminologies Description Logic classifier.
<http://www.cs.man.ac.uk/~horrocks/FaCT/>. Accessed June 2006.

Foster, I., Jennings, N. R. and Kesselman, C., (2004): Brain meets brawn: Why Grid and agents need each other. *Proc. 3rd Int. Conf. on Autonomous Agents and Multi-Agent Systems*, New York, USA, **01**:8-15, AAMAS.

The Physiology of the Grid: An open grid services architecture for distributed systems integration (2002), *The Physiology of the Grid: An open grid services architecture for distributed systems integration*. Technical report, Open Grid Service Infrastructure WG, Global Grid Forum.

Foster, I., Kesselman, C. and Tuecke, S. (2001): The Anatomy of the Grid: enabling scalable virtual organizations. *International Journal Supercomputer Applications*, **15**(3):200-222.

GBRMPA: Great Barrier Reef Marine Park Authority.
<http://www.gbrmpa.gov.au/>. Accessed May 2006.

Goble, C. and Bechhofer, S. (2005): OntoGrid: A Semantic Grid Reference Architecture. *CTWatch Quarterly*, **1**(4).

Goble, C. and De Roure, D., (2004): The Semantic Grid: Myth busting and bridge building. *In 16th European Conference on Artificial Intelligence*, Valencia, Spain, 1129-1135, ECAI-2004.

Hart, J. K. and Martinez, K. (2006): Environmental Sensor Networks: A revolution in the earth system science? *Earth-Science Reviews*, **78**(3-4):177-191.

Horrocks, I., Patel-Schneider, P. F., Boley, H., Tabet, S., Grosz, B. and Dean, M.: SWRL: A Semantic Web Rule Language - Combining OWL and RuleML.
<http://www.w3.org/Submission/SWRL/>. Accessed July 2006.

Horrocks, I., Patel-Schneider, P. F. and van Harmelen, F. (2003): From SHIQ and RDF to OWL: The making of a web ontology language. *Journal of Web Semantics*, **1**(1):7-26.

Hughes, T. P., Baird, A. H., Bellwood, D. R., Card, M., Connolly, S. R., Folke, C., Grosberg, R., Hoegh-Guldberg, O., Jackson, J. B. C., Kleyvas, J., Lough, J. M., Marshall, P., Nyström, M., Palumbi, S. R., Pandolfi, J. M., Rosen, B. and Roughgarden, J. (2003): Climate Change, Human Impacts, and the Resilience of Coral Reefs. *Science*, **301**(5635):929 - 933.

Jennings, N. R. (2001): An agent-based approach for building complex software systems. *Communications of the ACM*, **44**(4):35-41.

Jones, R., Hoegh-Guldberg, O. and Larkum, A. (1998): Temperature-induced bleaching of corals begins with impairment of the CO2 fixation mechanism in

- zooxanthellae. *Plant Cell and Environment*, **21**(12):1219-1230.
- Knublauch, H., Musen, M. A. and Rector, A. L., (2004): Editing description logics ontologies with the Protege OWL plugin. *In the International Workshop on Description Logics*, Whistler, BC, Canada, IWDL-2004.
- McGuinness, D. and van Harmelen, F.: OWL Web Ontology Language overview. <http://www.w3.org/TR/owl-features/>. Accessed May 2006.
- O'Connor, M. J., Knublauch, H., Tu, S. W., Grossof, B., Dean, M., Grosso, W. E. and Musen, M. A., (2005): Supporting Rule System Interoperability on the Semantic Web with SWRL. *Fourth International Semantic Web Conference*, Galway, Ireland, ISWC-2005.
- RacerPRO: RacerPRO Reasoner. <http://www.racer-systems.com>. Accessed June 2006.
- Rector, A., (2003): Modularisation of domain ontologies implemented in description logics and related formalisms including owl. *Proceedings of the 2nd International Conference on Knowledge Capture*, Sanibel Island, FL, USA, **K-CAP '03**:121-128, ACM Press.
- Rector, A., Drummond, N., Horridge, M., Rogers, J., Knublauch, H., Stevens, R., Wang, H. and Wroe, C., (2004): OWL Pizzas: Practical Experience of Teaching OWL-DL: Common Errors & Common Patterns. *Proceedings of the European Conference on Knowledge Acquisition (EKAW 2004)*, Northampton, England, **3257**:63-81, Springer-Verlag.
- Sánchez-Gestido, M., Pérez-Hernández, M. S., González-Cabero, R. and Gómez-Pérez, A., (2006): Improving a satellite mission system by means of a semantic grid architecture. *GGF16 Semantic Grid Workshop*, Athens, Greece, Global Grid Forum.
- Uschold, M. (2003): Where are the semantics in the semantic web? *AI Mag.*, **24**(3):25-36.
- Wikipedia: Wikipedia, the free encyclopedia. <http://en.wikipedia.org/wiki/>. Accessed June 2006.