The Philosophy of the Grid: Ontology Theory - From Aristotle to Self-Managed IT Resources

Thomas Sandholm
Dept. of Numerical Analysis and Computer Science and PDC
Royal Institute of Technology
SE-100 44 Stockholm, Sweden
+46-8-7907811

Abstract
The definition of ontologies has occupied philosophers ever since Aristotle's writings. In this paper I take a closer look at how the definition of ontologies has evolved and relate it to the research field of Grid computing, where it has had a recent revival and popularity boost for solving large-scale network resource integration issues. This paper builds on the previous works defining the Anatomy of the Grid, and the Physiology of the Grid. I conclude that there are indeed some very interesting parallels and views on the creation of ontologies, in particular the pragmatic views held by Carnap and Quine can be applied directly to current research, such as the development of the Semantic Web.

1 Introduction

Grid computing is a new paradigm of computer science with the main objective of facilitating *compute resource*[1] sharing across geographically distributed organizations to solve large-scale, and complex computational problems. It has evolved as a result of a revolutionary increase in network communication performance and availability. However, ubiquitous access to compute resources raises new questions on how to fairly share and utilize the qualities of service offered. My research focuses on investigating the approach of using automatically managed electronic contracts, a.k.a. Service Level Agreements (SLAs) to address some of these questions. Automation is driven by formal specification of semantics and ontology. When managing SLAs, it is crucial to agree on taxonomies that are semantically 'complete' in order to negotiate contracts in any deterministic and automated manner.

The philosophical field of ontology dates back to Aristotle and has influenced the view of ontology and formal logic in computer science greatly. Historically, some of the philosophical discussions are of a purely academic and theoretical nature and not directly or easily applicable in 'real-world' scenarios, whereas other discussions provide useful guidelines and recommendations on how to construct a sound ontology.

This paper therefore investigates what a selected set of acclaimed philosophers have concluded about the construction of an ontology, e.g., what it should and should not contain, and how to evaluate concrete instances. Defining exact semantics and the meaning of words, terms, and symbols; and how they refer to phenomena and various objects, may seem like an abstract exercise, but it is fundamental to the field of soft computing, where the human way of thinking is mimicked in software components to make them reason and make decisions like humans. Without a common classification system and a common universe of discourse no automated reasoning can take place. One example of an application of this idea is the use of fuzzy logic to drive the decision making process of automatically managed IT resources. The main objective of this study is to try and pick up ideas presented by the classical thinkers on

[1] Resources with a computational capacity to offer to remote users

ontology, and evaluate their applicability in the research field of Grid Computing in general and management of SLAs in particular.

The paper is structured as follows. In Section 2, the research field of Grid computing and how it relates to ontology research in computer science is defined in more detail. Section 3 presents and discusses ideas of selected philosophers and relates the ideas to the concepts defined in the preceding section. Finally, Section 4 draws some conclusions related to the discussion in Section 3.

2 Ontologies in Grid Computing

An ontology in computer science is a conceptual schema describing and classifying entities and their relationships and rules within a certain domain[2]. When defining an ontology, a foundational or upper ontology which is independent of a specific problem domain is used and extended. A foundational ontology could be seen as a glossary of terms used in a certain descriptive language such as a programming language, or a modelling language. A well-known and widely used upper ontology among computational ontologists is Dublin Core[3], which describes digital objects using metadata elements such as title, creator and subject. Due to the disparity of the field of computer science, and the potential competitive value of having a standard ontology that fits nicely into the specific problem domain of a commercial venture, there has not been any consensus on a single standard upper ontology. Computer scientist take a more pragmatic view of ontologies than philosophers in that a particular problem domain can accept an ontology as long as it describes all the entities that are useful for its applications, such as inductive reasoning, classification and problem solving in the field of artificial intelligence a.k.a. soft computing. So in that sense they are not so interested in specifying 'what exists' in the 'real world' and may adopt imperfect ontologies for the sake of making short-term progress.

Recent years have seen an exponential growth of standards, which could be seen as a community agreeing on an ontology, as a direct consequence of the enormous success of the World Wide Web (WWW) and its ubiquity resulting in a major both commercial and social impact. The key to the success is the simplicity and applicability of its upper ontology, the HyperText Markup Language (HTML). HTML has now been generalized into XML (eXtensible Markup Language) but the idea is the same. In terms of communication protocol design this idea is often referred to as the hourglass design, and its most prominent use is in the definition of the foundation of the Internet protocol (the OSI[4] model). In the *hourglass design*, there is a core fabric of definitions that are so simple that everyone can agree on them and apply them to their problem domain. Towards the two ends of the hourglass there are more specific and refined definitions that may only be applicable in certain contexts and communities. When data is communicated from one end of the hourglass to the other it goes through a number of transformation steps to reformat the data into more and more general ontologies. At the opposite end, the reverse process takes place to repackage the data into a format that can be understood by the receiving party. For this process to work on a large scale, such as the Internet, the definition of each layer and each layer transformation must be agreed on and standardized.

---

[2] See definition at http://en.wikipedia.org/wiki/Ontology_%28computer_science%29

[3] http://dublincore.org/

[4] http://www.iso.org/iso/en/CatalogueListPage.CatalogueList?ICS1=35&ICS2=100&ICS3=1&scopelist=

As communication capacity and performance has increased, it has become feasible to model the network of computers connected to the Internet as a single virtual computer for solving large-scale computational problems. This approach is the paramount driver of what is known as Grid computing. The term was coined after the electricity grid to emphasize the vision of making it as easy to use computational power in the future as it is to use electricity today. The pioneers in this area recognized the success of the Internet and its underlying hourglass design and thus constructed a similar design of the fabric supporting the Grid. For this reason the Grid is sometimes called the next generation Internet. The hourglass design of the Grid was first described by Foster et. al. (2001) in a landmark publication called the *Anatomy of the Grid* . It is based on experience with the Globus Toolkit®[5], which at that time implemented the de-facto standard of Grid computing. However, the authors of the toolkit realized that Globus would never survive as an upper ontology in the emerging Grid, and thus started a process of standardization, which we are in the midst of today, and which has now reached beyond the academic super computer departments (the initial application group of the Grid) to commercial companies as well. The key value proposition of standardization here is the added value of integrating heterogeneous resources and allowing them to be used and utilized more efficiently by a larger community of users such as a geographically distributed group of scientist trying to solve a common problem. This approach of using the Grid as a systems integration platform is the focal point of another landmark publication in the field, the *Physiology of the Grid*, Foster et. al. (2003), where the ontology defined in the anatomy publication is instantiated in a software architecture. This publication is also important from the point of view that it marks the start of the industry involvement in the Grid.

A parallel effort to the Grid is called the *Semantic Web*, and it is a direct extension of the current Web to add explicit representation of knowledge allowing a more sophisticated use of the WWW for solving problems rather than simply publishing and retrieving information, which is its primary model of use today. It is this effort that has brought back the computational ontology research in focus, and this community is revisiting the theories of philosophers such as Quine (discussed in Section 3) when defining their formal ontologies, which is the term used by philosopher for denoting the ontologies constructed by mathematicians and computer scientists. These formal representations are typically described using XML and can be used to automatically translate between representations, classify concepts, reason about and negotiate preferences, interfere decisions based on expert knowledge bases, and do semantic searches, as opposed to text-based, regular expression like searches typical of today's Web. In the future Semantic Web one could for instance ask questions like "What possible travel arrangement can be made for a trip from Stockholm to San Francisco with a budget of $2000"– a question you would typically ask a travel agent. Therefore the Semantic Web effort has also been a revival of the research into autonomous agents common in robotics.

Adding Semantic Web capabilities to the Grid as well is an emerging discipline of Grid computing, coined the *Semantic Grid*. There are many possible applications of knowledge-based problem solving functionality in the Grid. The application that I have focused on in my research is the automatic management (a.k.a. self-management) of IT-resources by the means of electronic contracts of Service Level Agreements (SLAs) as defined in Dan et. al. (2004). When reaching these agreements, ontologies must be agreed upon to define a universe of discourse. Furthermore, there must be a translation of ontologies from specifications of high-level user or business oriented preferences and goals to low-level resource policy

[5] http://www.globus.org/

configurations. Finally, SLA fulfilment must be assessed by analysing resource usage information described using accounting ontologies. Standardizing ontologies on each of these different levels and relating them to each other is thus crucial to the success of such an approach. With the ontologies in place, one can start adding reasoning capabilities to automatically adapt the system according to its usage. It would also allow users as well as resource providers to set their individual goals, such as meeting a budget or optimizing utilization to be fed into autonomous agents to search for and execute the best solutions. A promising model, described by Osborne and Rubinstein (1994), that has been applied to electronic markets is game theory, originally envisioned by John von Neumann and refined by John Forbes Nash, to find equilibrium solutions to combinatorial problems involving autonomous agents with individual preferences. Some experiments presented by Kasabov (1998) involving sensor and controller networks based on fuzzy logic[6] and neural network theory are other successful approaches taken for constructing decision-making infrastructure capable of dynamic adaptation.

3 Philosophical Ontology Theory

In this section I will present the philosophical view of semantics and ontology followed by accounts of the historical discussions related to these topics. In order to scope the discussion in this paper I have selected three philosophers who all have made substantial contributions to the discussion and definition of philosophical ontology. Aristotle (384 BC-322 BC) was selected for his groundbreaking definition and categorization of a universal base ontology. Rudolf Carnap (1891-1970) and Willard Van Orman Quine (1908-2000) were selected for their famous debate on meta-ontology, which has had a major impact on how philosophers view these issues today. Carnap was one of the leading philosophers and proponents of logical positivism in the Vienna Circle in the 1920's, Chalmers (2004, p. 59). Quine, on the other hand, has made major contributions in the field of formal logic, and is therefore frequently quoted in the computer science literature as well.

To frame the discussion it is important to settle working definitions of what is meant by semantics and ontology. 'Semantics'[7] is originally a Greek word (semantikos) meaning "significant meaning". The study of semantics is thus the same as the study of meaning, and it is often opposed to the study of syntax or the structure of expressions. Ontology, on the other hand, is in a general sense the study of the existence of things, i.e. "what there is". More specifically philosophical ontology involves the study of four different aspects, as described by Hofweber (2005):

1. what shared beliefs we have that make us acknowledge, or commit to, certain entities;
2. what there is;
3. features and relations of what there is;
4. and finally the semantics of 1-3, a.k.a. meta-ontology.

Ontology is a branch of metaphysics, the study of the most general features of reality. The usefulness of both metaphysics and ontology are contested among philosophers so it is important to keep them separate in the discussion.

[6] In turn based on the theory of Fuzzy Sets and extending the Boolean logic with the concept of degrees of truth
[7] See definition at http://en.wikipedia.org/wiki/Seman tics

Aristotle's view on ontology is often referred to as the realist view. In his work 'Categories' he provides a list of categories, which can be seen as an inventory of what there is in terms of the most general kinds of entities. These categories can be used to differentiate things as well as to refine specific aspects of things. As such it is a powerful tool to describe nature, i.e. what there is and make statements about nature that can yield affirmation, i.e. be evaluated to true or false. The categories are translated in Aristotle (350 B.C.E.) as:

- Substance (e.g., man, horse)
- Quantity (e.g., two cubits long, three cubits long)
- Quality (e.g., white, grammatical)
- Relation (e.g., double, half, greater)
- Place (e.g., in the Lyceum, in the market-place)
- Time (e.g., yesterday, last year)
- Postition (e.g., is lying, is sitting)
- State (e.g., shod, armed)
- Action (e.g., to lance, to cauterize)
- Affection (e.g., to be lanced, to be cauterized)

Aristotle, further, gives a recipe for how he arrived at the list, which continues to inspire ontologists constructing categories today. The important property of the individual categories is that they are not composite. That is, they need to be composed in order to make statements about the nature that can yield affirmation. One should not focus too much on the language used but rather on the things that the categories try to define. The exhaustiveness of the list is also debateable but as a first delineator for structuring an ontology its precision is striking. Relating it to the field of computer science discussed in the previous section, it is obvious that things like substance, quantity, quality, location and time are all vital components of an SLA describing the obligations and guarantees of service providers and consumers. A typical SLA would, for example, have statements concerning maximum (quantity) bandwidth (substance) in a subnet (place) during peak hours (time) with a certain reliability (quality). The state, action and affection categories on the other hand are vital artefacts of autonomic computing and system management, e.g. policy-based management and sensor/controller feedback loops. To allow for automatic management of an SLA, there must therefore also be policies specifying how the guarantees and obligations are to be monitored (retrieve state) to allow independent validation. Typical for policy-based management systems is that they specify what conditions (affection) should trigger various events causing actions to be taken that modify the state. Relational statements are also very important to SLAs in that contracts exist on many different levels of the computing infrastructure, e.g. group level and individual level, and therefore there must be ways to relate SLAs to each other. The position category is the only one I would argue against having as a top-level category. In the SLA use case it seems to fit better as a subcategory of state, which may be changed by various actions. Having a single uniform high-level list of categories across all ontology definitions is a goal that is, in many cases, counter-productive in the field of computer science where software needs to be developed in parallel by many different communities that may not agree on specifics of categories. It is therefore perfectly acceptable to end up with multiple ontologies as long as they are defined in a way so that they can be related to one another. Note the plural form of ontology, which is used more often in computer science texts than in philosophical texts, which may be an indication of the fact that the infeasibility of coming up with a single base ontology that everyone can agree with is less controversial in the former texts.

This concept of relativity of ontology leads us into the Quine and Carnap debate. Quine's most valued (and debated) contributions to philosophy are his theories about ontological relativity and commitment. These theories were formed as a direct result of the debate with Carnap originating in their, to some extent, opposite views on ontology and how it relates to empiricism and scientific work in general. Carnap (1950) argues that the questions of existence of certain entities can only be answered in terms of how fruitful they are to the aim for which the ontology is intended . The question of what is real from a theoretical point of view (i.e. seeking theoretical justification), which has occupied ontologists in the past, is meaningless at best and may even be confusing and counter-productive according to Carnap. Philosophers should instead, continues Carnap, focus their efforts on constructing linguistic frameworks, which will then be used to construct theories. The frameworks will be accepted or rejected simply based on the results that the theories are able to produce. In Carnap's theory, there are two different kinds of questions one can ask about entities: *internal questions* and *external questions*. Internal questions are questions of the existence of certain entities within the framework, whereas external questions are questions concerning the reality of the system of entities as a whole and they are meaningless and have trivial answers within a particular linguistic framework.

Quine (1951), on the other hand, argues that these trivial answers are invalid but agrees that the traditional ontological philosophy approach of seeking an answer to what is real is meaningless. Quine, also provides a recipe for deciding what entities to include in an ontology, partly based on formal logic. In his article *On what there is* he describes the process as binding values to variables. If one constructs a sentence like "There is something (bound variable) that is…" then one is committed to the entity instantiating the bound variable in the sentence. For example "There is something that is a car used to transport people", commits us to accept 'car' into our ontology. Replacing 'car' with 'bus' would allow us to include 'bus' into our ontology etc. Quine formulates this as "To be is to be the value of a variable". According to Quine, and also applying to scientific theories in general, we should look at the simplest ontology that fits our experience. Which ontology is the simplest will be determined by evaluating conflicting ontological views against each other in some application context. This form of meta-ontology is now widely accepted among philosophers. There are, however, problems with this approach in that it is very bound to linguistic properties of a particular language. That problem is, however, addressed by Quine's theory on ontological relativism presented in Quine (1968). It is based on the doctrines of inscrutability of reference and indeterminacy of translation. Some crude interpretations , for example Hacking (1983, p. 101), of this theory hold that one can never tell what someone else is talking about, thus giving it the bizarre effect that it does not matter what anybody says. The doctrines are based on the naturalistic view holding that apart from observations of the behaviour of other speakers the learner of a language has no data to work with. Identity and other important properties of objects cannot be determined by direct nor indirect observations leading to inscrutability of reference which in turn leads to problems when translating names of entities between languages. According to Quine's theory, the solution is to settle on a frame of reference or coordinate system, which is then used as basis for describing all entities within its scope. In terms of theories this coordinate system is called a background theory. The absolute truth of a theory and the absolute existence of an ontological entity are thus not possible to determine in isolation. Theories must always be related to and evaluated within the framework of other theories and ontology instances must regress to a well-established base ontology. The cyclic definition of lexical terms in a dictionary is an example of how new terms are explained using well-established terms, which in turn are established using even more fundamental terms.

So, what can we learn from the Carnap vs. Quine debate and ontological commitment. –There are many pragmatic approaches to constructing sound and usable ontological descriptions in both schools. Both, for instance, hold that one should focus on usability of a particular ontology in a particular context of application. Being non-discriminative to new ontological views is just as important as following the pragmatic approach of simplicity according to the principle of Occams Razor. However, spending time on constructing a sound base ontology would allow easier translation between a multitude of views. An SLA will be viewed and interpreted by many stakeholders with different backgrounds and objectives, and it will be interpreted by many different layers of abstraction and granularity in the middleware software. Providing an easy means to translate between ontological representations automatically is thus a necessity. A user X may for instance want to specify in a contract that she needs to transfer a file of size 900TB from machine A to machine B in less than 10 hours at a cost of less than $1 using GGF JSDL[8] (a job submission description language). This request may then be enforced by generating a policy that configures routers according to the router ontology of DiffServ[9] (differentiated services) marking packets from user X to allow expedited forwarding.

4 Conclusions

In this paper, I have shown that many philosophical theories concerning the construction of ontologies can provide useful guidelines for computer scientists. Many of the historical discussions on the existence of entities in the real world and the truth of theories that have occupied philosophers in the past (c.f. the rationalist vs. non-rationalist and the realist vs. non-realist disputes) are, however, to a large extent fruitless for a pragmatic computer science practitioner who needs to deliver programs to solve complex problems before they render businesses obsolete. The traditional philosophical view of ontologies was also rejected both by Carnap and Quine, although they had different reasons for rejecting it. Quine argued for formal ontologies grounded in formal logic relative to a background theory which directly corresponds to upper ontologies, and he also introduce the concept of proxy functions to translates between ontologies– a concept that is key in computational problems such as ubiquitous protocol design compliant with the hourglass principle. These obvious correspondences between Quine's theories and problems faced by today's computer scientists can also explain his recent popularity in the field. The ontological discussions not rooted in, nor related to, background ontologies are meaningless and do not cope with various problems of individualism according to Quine. Carnap's pragmatic approach to validating ontologies is also very much in line with the computer scientist view.

Constructing an ontology means classifying entities and relating them in relationships having the forms '… is a …' and '… is part of …'. Thus it is a process of systematic categorization from primitives to complex compounds. Aristotle took a realist approach to this problem and it is striking how well his envisioned categorization resonates with the categories one would expect to be covered in a Service Level Agreement contract, such as quality, quantity and time.

The conclusion is therefore that all the philosophers investigated, Aristotle, Carnap, and Quine provide valuable insights to computer scientists constructing ontologies to address the

[8] See http://www.ggf.org
[9] http://www.ietf.org/rfc/rfc2475.txt

needs of the next generation Semantic Grid. One should however treat the more theoretical philosophical discussions about the existence and reality of entities, or absolute truth of theories in relation to nature with scepticism, if considering them at all. As a concluding example of how philosophical ideas can contribute to computer science, Osborne and Rubenstein (2001, p. 5) hold Carnap's views on the process of science as systematic observations revealing regularities in the world, as the foundation for game theoretic models.

Bibliography

Aristotle. "Categories." *The Internet Classic Archive* (350 B.C.E).

Carnap, R. "Empricism, Semantics, and Ontology." *Revue Internationale de Philosophie* 4 (1950): 20-40.

Chalmers, A. F. *What Is This Thing Called Science?* Berkshire: Open University Press, 2004.

Dan, A., E. Davis, R. Kearney, A. Keller, R.P. King, D. Kuebler, H. Ludwig, M. Polan, M. Spreitzer, and Youssef. A. "Web Services on Demand: Wsla-Driven Automated Management." *IBM Systems Journal* 43, no. 1 (2004).

Foster, I., C. Kesselman, J. Nick, and S. Tuecke. "The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration.": Global Grid Forum, 2002.

Foster, I., C. Kesselman, and S. Tuecke. "The Anatomy of the Grid: Enabling Scalable Virtual Organizations." *International Journal of Supercomputer Applications* 15, no. 3 (2001).

Hacking, I. *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science*. Cambridge: Cambridge University Press, 1983.

Hofweber, Thomas. "Logic and Ontology." *The Stanford Encyclopedia of Philosophy* Spring 2005 (2005).

Kasabov, N. K. *Foundations of Neural Networks, Fuzzy Systems, and Knowledge Engineering*. London, England: MIT Press, 1998.

Osborne, M. J., and A. Rubinstein. *A Course in Game Theory*. Cambridge: The MIT Press, 2001.

Quine, W. V. "Ontological Relativity." *The Journal of Philosophy* 65, no. 7 (1968): 185-212.

———. "Two Dogmas of Empiricism." *The Philosophical Review* 60 (1951): 20-43.