MADHAN ARUMUGAM
The Peer-to-Peer Semantic Web: A Distributed Environment for Sharing Semantic
     Knowledge on the Web
(Under the Direction of DR. AMIT SHETH)

The goal of Semantic Web is to build a web of meaning. This vision can be demonstrated if people and applications can create and discover new and interesting knowledge and share this knowledge in a transparent manner similar to the way data is exchanged today. Therefore, it is believed that Semantic Web will consist of a distributed environment of shared and interoperable ontologies, which have emerged as common formalisms for knowledge representation. Users will need to discover new knowledge that is not known to them before and use it to either annotate the content or to formulate their information requests. This requires an environment that supports creating, maintaining, and controlled sharing of ontologies. A Peer-to-Peer (P2P) infrastructure can enable such capabilities. This thesis proposes an approach called P2P Semantic Web (PSW) with capabilities of finding relevant sets of ontologies, facilitating reuse of existing ontologies to create additional ontologies, and advertising the resulting ontologies. It discusses an ontology driven search of concepts and exploration of inter-ontological relationships over a P2P infrastructure. The prototyping of this approach in the InfoQuilt system is also discussed.

INDEX WORDS: Semantic Web, Knowledge Sharing, Ontology, Peer-to-Peer, Semantic
               Search, Information Correlation, InfoQuilt.

THE PEER-TO-PEER SEMANTIC WEB: A DISTRIBUTED ENVIRONMENT FOR

SHARING SEMANTIC KNOWLEDGE ON THE WEB

by

MADHAN ARUMUGAM

B.E. Anna University, India, 1999

A Thesis Submitted to the Graduate Faculty of The University of Georgia in Partial

Fulfillment of the Requirements for the Degree

MASTER OF SCIENCE

ATHENS, GEORGIA

2001

THE PEER-TO-PEER SEMANTIC WEB: A DISTRIBUTED ENVIRONMENT FOR

SHARING SEMANTIC KNOWLEDGE ON THE WEB

by

MADHAN ARUMUGAM

Approved:

Major Professor: Amit Sheth

Committee: Krzysztof Kochut
    Daniel Everett

Electronic Version Approved:

Gordhan L.Patel
Dean of the Graduate School
The University of Georgia
December 2001

DEDICATION

To my beloved parents Ramakrishnan and Chandrakala, brother Vijayan, and friend

Preethi for their love, support, encouragement and belief in me.

ACKNOWLEDGEMENTS

TABLE OF CONTENTS

CHAPTER

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

1.1   Knowledge and its role in today's web

Definitions are the building blocks of knowledge. A definition, by definition describes the meaning of something, as a collection of substantiated facts, or as a relative behavior of something already known. For instance, Webster defines "knowledge" as *the fact or condition of having information*. For the user to understand the meaning of "knowledge" the user should understand the meaning of "fact", "condition" and "information", and if he/she does not understand those meanings, he/she should delve into their definitions. So a definition of A can look like,

A *is a* B *or* C *having* D

where, B, C and D are other definitions/facts agreed upon or already known and *is a*, *or* and *having* are binary verbs that define the relationship between two information entities(fact or definition). Thus the entire knowledge space is built with definitions, one defining another. A knowledge model is the view of the knowledge space from the perspective of the user. The counterpart of definitions in a knowledge model is ontology. Ontology according to Gruber is defined as specification of a conceptualization [GRU93]. It basically details a concept, its properties and relationships with other concepts in the knowledge model.

In today's Web, knowledge plays the key role of transforming terabytes of data into meaningful information. This knowledge, however, needs to be made machine processable in order for the machines to *understand* it. Semantic Web is one of the initiatives to build an environment, where techniques are developed to make today's web a web of meaning.
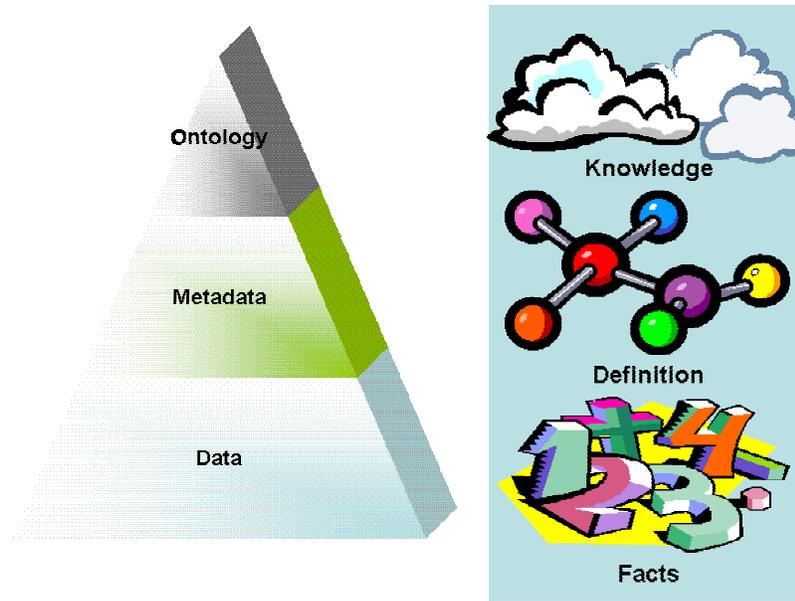
Figure 1.1: Knowledge Model and Knowledge Space

1.2   Semantic Web and its Vision

Semantic Web is the next step in the evolution of the World Wide Web [B99, BHL01, FM01] with the promise that the users will be able to use programs that can understand semantics of the data. In support of such a vision, a good amount of recent work has focused on annotating data using ontologies to make data machine-understandable. From a broader perspective, Semantic Web is intended to make the Web more intelligent by improving the value of, and interactivity with the vast amount of data on the Web. One approach for such a support for semantics is the use of ontologies to provide the context for capturing the meaning of data as well as the user's intention in a query. Such an approach has been advocated for some time [e.g., KS94, MKSI96]. Eventually, this can help users answer complex questions like *find a correlation between earthquakes and nuclear tests*, from data of different types, formats, and media from multiple resources**.**

1.3 Challenges in building the Semantic Web

There are many challenges that need to be addressed in order to realize the above vision. We need an environment where people and applications can share knowledge, locate new knowledge, and use this knowledge in a transparent manner similar to the way data is exchanged today. Difficulty in creating and maintaining large ontologies and the need for knowledge sharing and ontology interoperation has been presented in [MWD01, KS94]. One challenge is that of advertising knowledge and ontologies for different information domains, which are maintained by different persons, groups and organizations on the Web. For example one needs to easily locate ontologies for *earthquake* and *nuclear test* domains in the above example. With a wide variety of communities and user groups on the web, it is difficult to find most relevant set of ontologies. A semantic search mechanism is necessary to find most relevant set ontologies using users' context for information request and his profile. Once the ontologies are located there is a need for introducing some relationships across ontologies and supporting techniques for ontology interoperation [MWD01, MIKS00]. Finally, users need tools that would allow them to define information requests (*e.g., find a site for a new landfill near the main source of the municipal refuse*, where the earthquake's impact must also be evaluated) or knowledge discovery involving hypothesis validation (*e.g., do nuclear tests effect earthquakes? If so find the correlation*), and to collaborate with other users involved in similar knowledge-creation activities.

1.4 Knowledge Sharing in InfoQuilt

The InfoQuilt System developed at the LSDIS Lab. provides a framework for formulating complex information requests, which can capture the semantics of user's request involving multiple ontologies [LSA+01], and support a form of knowledge discovery [TSP01]. This thesis involves integrating InfoQuilt's semantic capabilities with technologies more appropriate for developing a distributed and collaborative Semantic

Web platform. As a step towards this objective, the use of Peer-to-Peer (P2P) computing is investigated as a possible infrastructure. The new concept is called Peer-to-Peer Semantic Web (PSW). In realizing PSW, DAML+OIL provides a specification framework for independently creating, maintaining, and interoperating ontologies while preserving their semantics, and P2P is used to provide a distributed architecture which can support sharing of independently created and maintained ontologies. The PSW facilitates:

- Distributed and autonomous creation and maintenance of local ontologies,

- Advertisement (*i.e.* registry) of local ontologies,

- Introducing inter-ontological relationships between relevant ontologies on an as-needed basis once they are located,

- Controlled sharing of knowledge base components among users in the network,

- Ontology-driven semantic search of concepts and services,

- Knowledge discovery and exploration of inter-ontological relationships.

The two important problems the PSW approach intends to address are 1) what knowledge to share, in other words semantic search of relevant ontologies and 2) how to share, meaning a platform to share independently created and maintained ontologies.

1.5 Thesis Organization

The thesis is organized as follows. Chapter 2 is the paper "The Peer-to-Peer Semantic Web: A Distributed Environment for Sharing Semantic Knowledge on the Web" submitted to the eleventh international World Wide Web conference (WWW2002) to be held in May 2002. Section 2.2 briefly summarizes the related work. Section 2.3 discusses the background of InfoQuilt system and its semantic capabilities as essential components of Semantic Web. It also discusses P2P as a suitable infrastructure for integrating those capabilities in Semantic Web. Section 2.4 explains in detail the architecture and modules

in the InfoQuilt System that facilitates sharing of autonomously maintained ontologies. It also explains knowledge space construction and navigation. Section 2.5 details the semantic search of relevant ontologies with examples. Chapter 3 presents our conclusions and future directions of further research.

1.6 References

[B99]   Tim Berners-Lee (with Mark Fischetti), Weaving the Web, The original design and ultimate destiny of the World Wide Web, Harper, 1999.

[BHL01]   T. Berners-Lee, J. Hendler, and O. Lassila, "The Semantic Web: A new form of Web content that is meaningful to computers will unleash a revolution of new possibilities," Scientific American, May 2001.

[FM01]   D. Fensel and M. Musen, Eds. "The Semantic Web: A Brain for Humankind," IEEE Intelligent Systems, March/April 2001

[GRU93]   Gruber, T.R. 1993. A translation approach to portable Ontology specifications. Knowledge Acquisition, 5:166-220.

[KS94]   V. Kashyap and A. Sheth, "Semantic based Information Brokering," Proceedings of the 3[rd] Intl. Conference on Information and Knowledge Management, November 1994, pp. 363-370.

[LSA+01]   Tarcisio Lima 1,2 , Amit Sheth 3 , Naveen Ashish, Mukesh Guntamadugu, Sriram Lakshminarayan 4 , Narayanan Palsena, and Dilpreet Singh, Digital Library Services Supporting Information Integration over the Web, WIIW 2001

[MKSI96]   E. Mena, V. Kashyap, A. Sheth, A. Illarramendi, "OBSERVER: An Approach for Query Processing in Global Information Systems Based on Interoperation Across Pre-existing Ontologies", Proceedings of the 1st IFCIS Intl. Conference on Cooperative Information Systems (CoopIS'96), Brussels, Belgium, June 1996, pp. 14-25.

[MIKS00]   E. Mena, A. Illarramendi, V. Kashyap, A. P. Sheth. OBSERVER: An approach for query processing in global information systems based on interoperation across pre-existing ontologies. International Journal on Distributed and Parallel Databases, Vol. 8, No. 2, pp. 223-271, April 2000.

[MWD01]   Prasenjit Mitra, Gio Wiederhold and Stefan Decker: A Scalable Framework for Interoperation of Information Sources. Proceedings of the 1st International Semantic Web Working Symposium (SWWS `01), Stanford University, Stanford, CA, July 29-Aug 1, 2001, Jul. 2001

[TSP01]   S. Thacker, A. Sheth and S. Patel, "Complex Relationships for the Semantic Web," Creating the Semantic Web, D. Fensel, J. Hendler, H. Liebermann, and W. Wahlster (eds.), MIT Press, 2001 (to appear).

CHAPTER 2

THE PEER-TO-PEER SEMANTIC WEB: A DISTRIBUTED ENVIRONMENT FOR SHARING SEMANTIC KNOWLEDGE ON THE WEB[1]

---

[1] Arumugam, M. Sheth, A. and Arpinar, B. 2001. Submitted to Eleventh International World Wide Web Conference (WWW2002), 11/13/2001.

Abstract

The real value of Semantic Web vision can be demonstrated if people and applications can create and discover new and interesting knowledge and share this knowledge in a transparent manner similar to the way data is exchanged today. Therefore, we believe that Semantic Web will consist of a distributed environment of shared and interoperable ontologies, which have emerged as common formalisms for knowledge representation. The users will need to discover new ontologies, that are not known to them before and use them to either annotate the content or to formulate their information requests. This requires an environment that supports creating, maintaining, and controlled sharing of ontologies. We believe that a Peer-to-Peer (P2P) infrastructure can enable such capabilities. We call our approach P2P Semantic Web (PSW) with capabilities of finding relevant set of ontologies, facilitating reuse of existing ontologies to create additional ontologies, and advertising the resulting ontologies. We discuss an ontology driven search of concepts and services and exploration of inter-ontological relationships over a P2P infrastructure. The prototyping of this approach in the InfoQuilt system is also discussed.

2.1 Introduction

Semantic Web has been presented as the next step in the evolution of the World Wide Web [B99, BHL01, FM01]. Central to the current Semantic Web vision is the promise that the users will be able to use programs that can understand semantics of the data. In support of such a vision, a good amount of recent work has focused on annotating data using ontologies to make data machine-understandable. From a broader perspective, Semantic Web is intended to make the Web more intelligent by improving the value of, and interactivity with, vast amount of data on the Web. One approach for such a support for semantics is that of using ontologies to provide the context for capturing the meaning of data as well as the user's intention in a query. Such an approach has been advocated for some time [e.g., KS94, MKSI96]. Eventually, this can help users answer complex questions like *find a correlation between earthquakes and nuclear tests*, from data of different types, formats, and media from multiple resources.

There are many challenges we need to address in order to realize the above vision. We see a need for an environment where people and applications can share knowledge, locate new knowledge and use this knowledge in a transparent manner similar to the way data is exchanged today. Difficulty in creating and maintaining large ontologies and the need for knowledge sharing and ontology interoperation has been presented in [MWD01, KS94]. One challenge is that of advertising knowledge and ontologies for different information domains, which are maintained by different persons, groups and organizations on the Web. For example one needs to easily locate ontologies for *earthquake* and *nuclear test* domains in the above example. With a wide variety of communities and user groups on the web, it is not trivial to find the most relevant set of ontologies. A semantic search mechanism is needed to find the most relevant set ontologies using users' context for information request and his profile. Once the ontologies are located there is a need for introducing some relationships across ontologies and supporting techniques for ontology interoperation [MWD01, MIKS00]. Finally, users

need tools that would allow them to define information requests (*e.g.*, *find a site for a new landfill near the main source of the municipal refuse*, where the earthquake's impact must also be evaluated) or knowledge discovery involving hypothesis validation (*e.g.*, *do nuclear tests effect earthquakes? If so find the correlation*), and to collaborate with other users involved in similar knowledge-creation activities.

The InfoQuilt System developed at the LSDIS Laboratory at the University of Georgia. provides a framework for formulating complex information requests, which can capture the semantics of user's request involving multiple ontologies [LSA+01], and support a form of knowledge discovery [TSP01]. Our current work involves integrating InfoQuilt's semantic capabilities with technologies more appropriate for developing a distributed and collaborative Semantic Web platform. As a step towards this objective, we investigate the use of Peer-to-Peer (P2P) computing as a possible infrastructure. We call the new concept a Peer-to-Peer Semantic Web (PSW). In realizing PSW, DAML+OIL provides a specification framework for independently creating, maintaining, and interoperating ontologies while preserving their semantics, and P2P is used to provide a distributed architecture which can support sharing of independently created and maintained ontologies. Our PSW facilitates:

- Distributed and autonomous creation and maintenance of local ontologies,

- Advertisement (*i.e.* registry) of local ontologies,

- Introducing inter-ontological relationships between relevant ontologies as-needed basis once they are located,

- Controlled sharing of knowledge base components among users in the network,

- Ontology-driven semantic search of concepts and services,

- Knowledge discovery and exploration of inter-ontological relationships.

The two important problems the PSW approach intends to address are 1) what knowledge to share, in other words semantic search of relevant ontologies and 2) how to share, meaning a platform to share independently created and maintained ontologies.

The rest of the paper is organized as follows: Section 2 briefly summarizes related work. Section 3 discusses the background of the InfoQuilt system and its semantic capabilities as essential components of Semantic Web. It also discusses P2P as a suitable infrastructure for integrating those capabilities in Semantic Web. Section 3 explains in detail the architecture and modules in the InfoQuilt System that facilitate sharing of autonomously maintained ontologies. It also explains knowledge space construction and navigation. Section 4 details the semantic search of relevant ontologies with examples. Section 5 presents our conclusions and future directions of further research.

2.2 Related Work

The related work on knowledge sharing primarily focuses on development of a multi-ontology environment and either integration or interoperation of these ontologies. As discussed in the literature, since ontology interoperation supports development and maintenance of locally developed independent ontologies it is superior to the integration approach. Our PSW approach is also based on the principle of multi-ontology interoperation but also facilitates an efficient search technique for new ontology discovery. In this way a proper set of ontologies are automatically located and relevant inter-ontological relationships are introduced, and this is an advancement which does not exist in the previous multi-ontology integration or interoperation approaches.

Some of the past research on knowledge sharing in a multi-ontology environment motivates our PSW approach. For example, [KS94] talked about supporting information resource discovery in a system with multiple ontologies, and [MKSI96, MIKS00] discussed processing of information requests, again involving multiple ontologies. [EFG+99, PB99] talked about approaches in using ontologies for knowledge re use and

sharing. Several issues of knowledge sharing, especially those related to ontology interoperation are discussed in [MWD01]. [STP01] presented the concept of information requests involving complex relationships and multiple ontologies, and use of parameterized information requests for knowledge discovery through what-if analysis. [MD00] proposed an approach for interoperating different information resources on a syntactic level. [GHI+01] focused on the data placement problem of how to distribute data and work among peers cost effectively. The above efforts primarily addressed the infrastructure and syntactical issues of knowledge sharing.[DHM+01] detailed about accessing information and services in a DAML enabled web, although it lacked an infrastructure to support knowledge sharing. [ERN01] proposed using P2P to support the semantic web to link semantic definitions, but lacks support for complex relationships and knowledge discovery. Our PSW approach addresses the logical and semantic level of knowledge sharing with support for exploration of inter-ontological relationships and therefore knowledge discovery.

2.3 Knowledge Discovery and Sharing in InfoQuilt

In this section, we discuss the InfoQuilt System, and its semantic capabilities like support for complex information requests and knowledge discovery, and also motivate the need for knowledge sharing support in the evolving Semantic Web.

2.3.1 Information Requests (IScapes)

InfoQuilt is an agent-based system that allows users to semantically request information, semantically correlate data from different sources and of heterogeneous type or representation, and analyze data available from diverse autonomous and heterogeneous sources [STP01]. Furthermore, it also supports knowledge discovery through interactive what-if analysis or information hypotheses including conjectures and relationships within and across domains. InfoQuilt system includes: (a) language and tools to specify IScapes

(*i.e.,* semantic information requests), and (b) tools and algorithms to perform what-if analyses to search the information space of semantically related data. IScapes allow parameterized specification of information requests and correlation that utilizes the domain ontologies, inter-ontological relationships and user defined functions to accurately describe a user's information need.

The InfoQuilt system provides a framework to support knowledge discovery and learning from a multitude of diverse autonomous distributed resources. It allows the user to put pieces of information together and understand their interactions and form complex information requests on data available on heterogeneous sources. IScapes are more than a traditional query in that they can understand user's request by embedding semantic information [LSA+01]. For example consider the following information request:

"*Find all earthquakes with epicenter in a 5000 miles radius of the location at latitude 60.790 North and longitude 97.570 East and find all tsunamis that they might have caused.*"

In addition to the known explicit constraints, the system needs to comprehend what the user means by saying *find all tsunamis that might have been caused due to the earthquake*. An IScape is specified in terms of the components of the knowledge base of the system such as ontologies, inter-ontological relationships (or articulations), and user-defined operations. Representing that *an earthquake caused a tsunami* requires a complex inter-ontological relationship, involving ontologies that define earthquake and tsunami using different ontologies, and inter-ontological relationships with additional parameters and functions for spatial and temporal proximity between a pair of earthquake and tsunami events. Thus we define IScape as a computing paradigm that allows users to query and analyze the data available from diverse autonomous sources, gain better understanding of the domains and their interactions as well as discover and study

relationships [STP01]. Building upon previous work, the current work changes the representation to DAML+OIL and DAML-S (for service discovery), which provides more features than our earlier XML and RDF based specification. More importantly, the earlier system could execute an IScape on distributed resources, but the IScape construction was done using tools that used a single knowledge base. So a key objective of our current work is to support distributed knowledge construction and sharing of a distributed knowledge base.

Most of the past work has been based on simple and well-defined relationships like "is-a", "is-part-of", "is-similar-to", "is-associated-with" which are hierarchical or similarity based. However, the several real-world relationships between entities are much more complex and it is not possible to express those using simple relational and logical operators. For example, consider the relationships "earthquake causes tsunami", "air-pollution affects vegetation", and "nuclear test causes earthquake". These are some examples of the complex relationships that typically span across multiple classes or ontologies. To express such relationships one needs to be able to model the semantics involved in the interaction between the domains. InfoQuilt already provides an infrastructure to model such complex relationships and use them to specify information requests [STP01]. But the knowledge, more specifically ontologies, used in constructing such IScapes is not shareable, hence not interoperable in Semantic Web. Our PSW approach solves this problem by providing a platform for creating, maintaining, controlled sharing and using ontologies.

2.3.2 Knowledge Construction and Sharing

Knowledge base construction remains an expensive task, even as new tools are being developed for knowledge acquisition. This is primarily because a new system that is built requires the knowledge base to be constructed from scratch. As a result, several information systems in a given domain (such as medicine), have significant duplication

or overlap in their knowledge bases. The cost of this duplication is prohibitive for large systems. This problem is solved by sharing, reusing and building on existing knowledge bases [NFF+91]. For example consider an automobile company that tracks customer information in all its stores. Although each store is located in a different location, and franchised by different retailers, some amount of knowledge regarding the make and model of this brand is shared among all the stores. Instead of duplicating this information in all the stores, this knowledge when shared saves considerably. However, not all overlap or duplication can be avoided, in which case we need to promote interoperation. Knowledge is not a substance that can be moved from one place to another place, rather a specification of what the representation can do [MUS92]. Knowledge that is represented as symbols and structures delineates the knowledge model that the user designed and constructed. In trying to share the knowledge stored as symbols and structures, we encounter a number of challenges. First of all, symbols and structures representing knowledge are tightly related to a user's context, and without context, the semantics is lost while sharing. The context is sensitive to spatial and temporal co-ordinates (see Figure 2.1). For instance, the symbol "book" can have two different interpretations (book as in *I lost my physics book* and book as in *I need to book my tickets today*) depending on the context. Similarly, the notion of a symbol "September 11" before and after the World Trade Center attacks, is different. Before, it was just another date, but after the attacks it has some added significance, and when shared between users, the knowledge has to be shared with the context. Otherwise each user will get a different notion of the same knowledge resulting in anomalies.

The four main challenges for knowledge sharing are heterogeneous representations, syntactic and semantic heterogeneities within a knowledge representation, lack of communication or mapping protocol, lack of shared vocabulary and domain terminology [NFF+91]. The problem of heterogeneous representations can be solved by using knowledge base frameworks supporting formal definitions of the terms that can be used

to model a domain or class of domains and assertions that govern the requirements and constraints for creating valid models. For instance DAML+OIL is a framework that allows the users or communities to define their own ontologies, extending basic ontologies. This allows the bottom up design of meaning and also sharing of high -level concepts [COV01]. Knowledge bases with inference engines can resolve the syntactic and semantic differences in a knowledge representation system. The use of ontologies, solve the issue of interoperability among knowledge base systems.
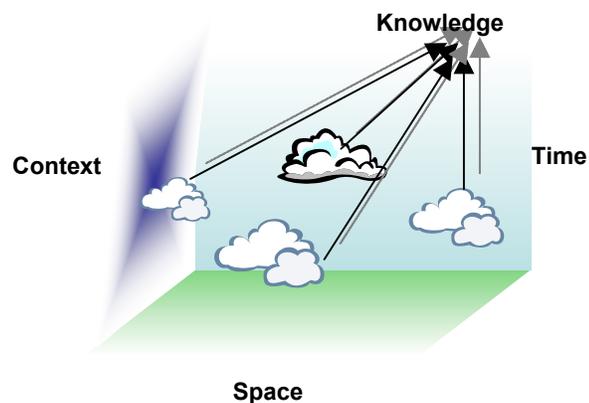

Figure 2.1: Knowledge as a 3-Dimensional Entity

2.3.3 Benefits of Peer-to-Peer Architecture

Peer-to-peer computing is the sharing of computer resources and services by direct exchange between systems. These resources and services include the exchange of information, processing cycles, cache storage, and disk storage for files [Pee]. The classical model of P2P involves no use of central server. Peers are connected directly to one another (*e.g.*, Gnutella). The hybrid model of P2P involves a central server which controls the services and peers having control over their resources. The popular P2P systems like Napster [Nap] and Morpheus [Mus] have attracted enormous attention

towards Peer-to-Peer computing. The features in P2P that makes it desirable to be an infrastructure for knowledge sharing in Semantic Web are that it

- encourages distributed architecture,

- supports decentralization of control and access to information and services,

- provides a way to harness the computing power and knowledge of millions of computers in the web,

- allows user to access semantic information published by several independent content providers, and to create personalized semantic relationships among those information

- provides support for publishing peer definitions and relationships to other peers and software agents,

- offers user-centered, data-centered and computing centered models [GAR01]provides a suitable architecture for distributed content management
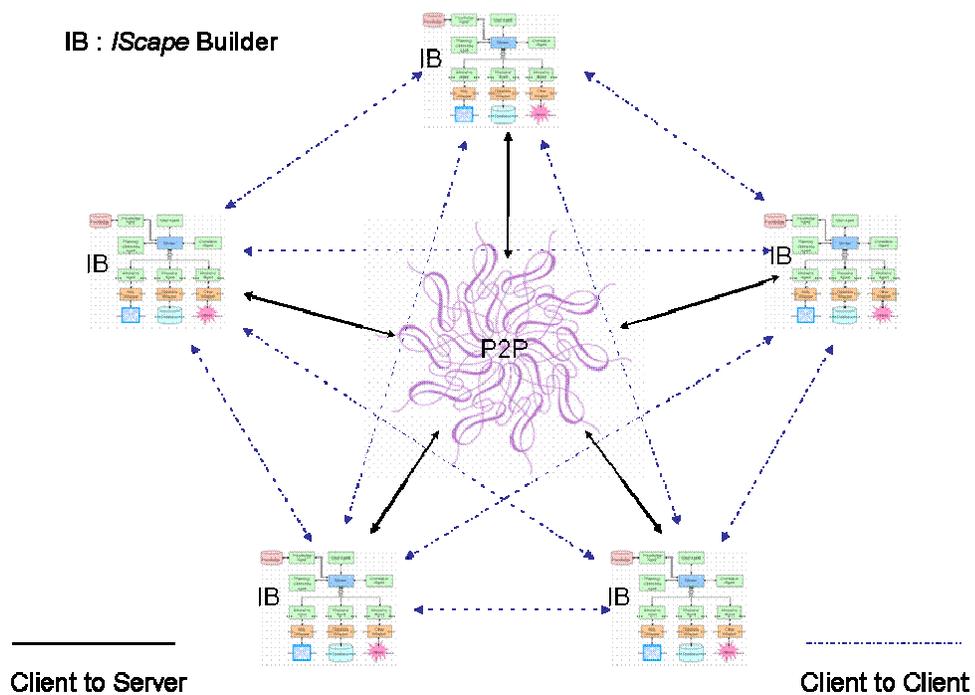

Figure 2.2: InfoQuilt's P2P Knowledge Sharing Network

Global Web Content Management (GWCM) is one of the biggest challenges of enterprise web. GWCM solutions focus on processes and controls for localizing content for language, culture and commerce, requiring multiple content repositories in different locations. These solutions can use distributed content management frameworks like P2P to assuage the replication of content [GAR01]. Decentralization is another desirable feature for knowledge sharing and Semantic Web [ERN01] as a whole because knowledge is autonomously created and maintained in diverse locations, yet there is a need for collaboration and consensus making, esp. for reaching ontological commitment. We will next discuss the evolution of the InfoQuilt system architecture using P2P infrastructure, which allows us to prototype its advanced semantic capabilities in a form suitable for Semantic Web.

2.4   Semantic Knowledge Sharing in InfoQuilt

2.4.1   InfoQuilt's Peer-to-Peer Semantic Network

The InfoQuilt System modified as part of the current work has users connected in a P2P network, accessing a shared directory of information and services. The P2P network in InfoQuilt is user centered [GAR01], where a client registers itself to a directory or directories, and searches the directory for peers providing *semantically relevant* information and services. The directory also provides the necessary contact information to connect to a specific peer, so that a direct peer-to-peer connection can be established (see Figure 2.2). Currently we are using a proprietary directory service but are also investigating UDDI [UDD00] for a possible use for advertising and registering different ontologies. Each user in the P2P network runs a multi-agent Information Brokering System whose architecture is built upon the work done in [LSA+01]. The information and services that a user is interested in are semantically identified using the Knowledge Space Navigation algorithm, which provides means for locating relevant ontologies, hence relevant content. Users can build IScapes using the IScape Builder [TPS01]

utilizing the most relevant knowledge about concepts, which are semantically identified and retrieved from the P2P Knowledge Sharing network.

An earlier work focused on supporting complex inter-ontological relationships and knowledge discovery in InfoQuilt. This work focuses on knowledge sharing in a P2P driven Semantic Web using DAML+OIL ontologies. The ontologies define concepts, their properties and their relationship with other concepts. Thus with sharing ontologies, we also share inter-ontological relationships. PSW offers the users more semantic and personalized choices in formulating IScapes, in terms of ontologies created by other peers in the network. To realize the above feature, the IScape Builder now has the capability to access and add new ontologies and user-defined inter/intra-ontological relationships to the knowledge space in the P2P network. Therefore, a global knowledge space within the scope of this paper is assumed to be composed of local ontologies, which are created and maintained by independent peers and connected by some inter-ontological relationships. The user searches for a relevant ontology and the results are re-ranked based on the user profiles and query relationships by the personalization agent. The architecture of an InfoQuilt peer is shown in Figure 2.3:
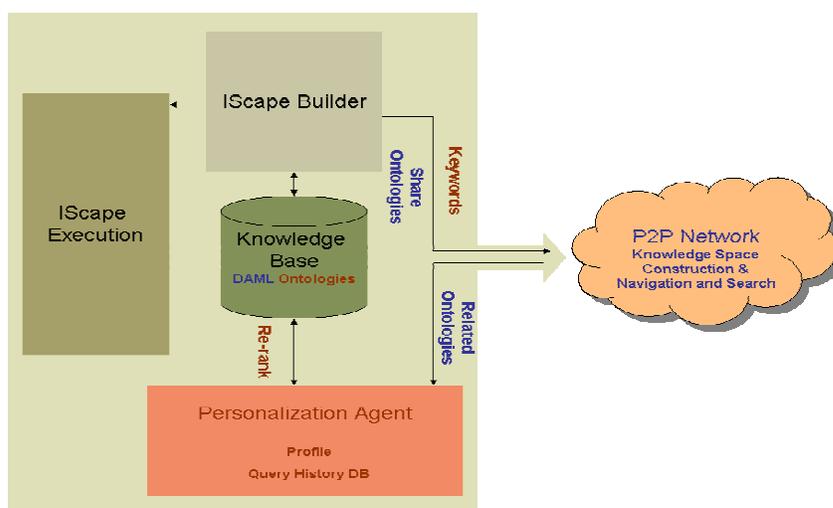

Figure 2.3: Architecture of an InfoQuilt Peer

2.4.2 Knowledge Space Construction and Navigation in InfoQuilt

The key issue while implementing the Semantic Knowledge Sharing is to construct a globally accessible knowledge space. Users must be able to do more than mere access of knowledge space. Each user must be able to add or remove concepts, without having to re-engineer the knowledge space [MUS92]. In a PSW, users can create their own ontologies subscribing to the definitions they refer to. So the users can construct their own ontologies and at the same time conform to the Global Knowledge Space. All ontologies defined in DAML+OIL are qualified by namespaces. Conceptually namespaces are identifiers or references in the knowledge space. When a person defines his own concept or notion based on a pre-defined or agreed upon concepts, it is marked up in the knowledge space by these references. It does not necessarily mean that we have to agree to all the definitions existing in the knowledge space, but DAML+OIL helps to create a democratic and ecological knowledge space where different concepts and definitions can co-exist. The definition that survives is the one that is most referenced, other definitions go by unnoticed. In using DAML+OIL, the only concepts everybody agrees upon are the basic classes like *Thing* [Thi]. These are the predefined most or least general classes of all the classes that are defined in DAML+OIL [Dam]. With the use of DAML+OIL the knowledge space is constructed from roots. When a new ontology is created, it is already hooked up in the knowledge space because of the use of imports and namespaces [Dam]. PSW's support for peer registration of ontologies is discussed later in this section.

In order for the programs to access this knowledge space programmatically, a data structure is used as described. A DAML+OIL ontology is a collection of RDF triples [Ref]. Each RDF statement is a resource reifying a triple. Such a resource must have at least three properties *rdf:subject, rdf:object* and *rdf:predicate,* valued by the corresponding resources. Each RDF statement in an ontology can define a class, its properties and its relationships with other classes.

For example a statement would look like

<#boy> <#drinks> <#coffee>.

Now all of the triples, *boy*, *drinks* and *coffee* are qualified by use of URIs. The data structure stores subject (*boy*), the object (*coffee*) and the verb (*drinks*) that relates them. So the core components of the data structure are

- Concepts: Concepts or KObjects as we refer to contain information about each class. In the example, *boy* and *coffee* qualify to be KObjects.
- Links: Links or relationships contain information about the predicate and the KObjects it relates. In our example *drinks* qualifies as a Link.

For each KObject, pointers to the Links that have this KObject as a subject or object and the ontology which defines it along with the user information are maintained. A sample knowledge space representation is shown in Figure 2.4. In creation of the knowledge space the following steps are involved:

1. Retrieve every RDF triple (subject, predicate, and object) from each source ontology,
2. For every assertion of a fact or a definition made in the ontology, recursively trace its link to the most general class of the knowledge space (#Thing),
3. Repeat 1 and 2 till all the ontologies are hooked into the knowledge space.

For knowledge space navigation, we can start with the KObject Thing and then traverse through the Links in the KObject.

2.4.3 Ontology Registration

In PSW, Ontologies (DAML+OIL) are defined as collections of definitions and assertions. The peers can create and maintain their own ontologies. They have the control as in whether or not to share an ontology. Any peer can create his or her own definitions conforming to DAML+OIL formalisms. Additionally peers can define ontologies through assertions that refer to definitions created by other peers and also can let other peers refer to their definitions. When a peer decides to share an ontology in the global knowledge space, s/he has to upload the ontology into the PSW. This process is called registration of ontologies. New concepts (KObjects) and relationships (Links) in the uploaded ontology are created appropriately. Once an ontology is registered, other peers can refer to the definitions in this ontology for their use. This approach enables an environment for controlled sharing of ontologies. For example, a peer can protect his definition by not sharing, but use assertions that refer to shared definitions. In case a peer decides to remove or deregister his ontology, all the definitions and the assertions that refer to these definitions become invalid in the knowledge space. Tools like DAML Validator can be used to check for obsolete definitions and stale links.
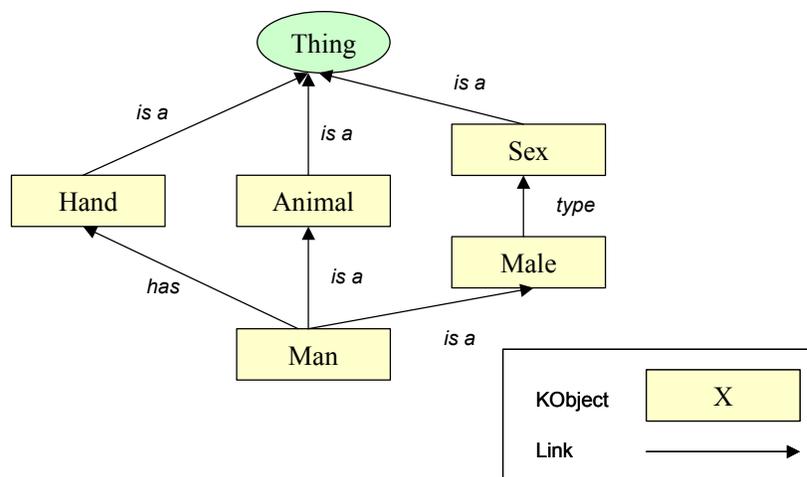
Figure 2.4: A Sample Knowledge Space Representation

2.5   Semantic Search

One of the key advantages of constructing a knowledge space is semantic search. In the IScape Builder, the user specifies the keywords (usually common nouns) used in the information request. The next step in the construction of the IScape is selection of ontologies that help the system understand the user's query [STP01]. In the previous version of InfoQuilt, on which our PSW approach is built, the users need to create either their own ontologies from scratch or manually select some of the previously created local ontologies. However, the PSW approach enables InfoQuilt to search the knowledge space for the relevant ontologies.

The data structure representing the knowledge space is a collection of KObjects and Links. The process of searching involves the following steps. The input is a set of keywords and the output is a list of ontologies.

1. Take each keyword and run a basic keyword match on the subject, object and the predicate (in that particular order) in the entire knowledge space,

2. Retrieve the name of the ontologies that satisfy the above match along with the ownership details,

3. If the keywords result in a number of ontologies, compare the ontologies for common parents and eliminate the ontologies without any common links,

4. If there is more than one ontology describing the same keyword, perform search with more keywords or compare the resulting ontologies to help user select the ontology.

Even though this search is invoked using keywords, we call it semantic, because the keyword matches are done with the KObjects qualified in the knowledge space, which by its very construction preserves the semantics of the knowledge it describes. Unlike ordinary keyword searches, this kind of search has specific direction, because we believe

that the keywords used in a valid information request are semantically woven in the holistic meaning they represent. One other utility we have developed using the knowledge space is the ability to compare two ontologies. This involves the following steps.

1. Identify the KObjects (concepts) used in each ontology in knowledge space,
2. Find a common parent KObject that links two KObjects that are defined in each of the compared ontologies; in other words find a connecting link (relationship) between the two ontologies and trace it for the user.

With DAML+OIL used in defining the ontologies, any two concepts will be linked together by the fact that both of them are *Thing*s. The scope of defining how close a common KObject can be in terms of the number of Links to be defined as a certain relationship is an interesting research problem and is beyond the scope of this paper. This utility can thus help the user explore inter-ontological relationships, and hence knowledge discovery. Now we will demonstrate the search with an example.

Example 1: Consider the information request

"*Find all earthquakes with epicenter in a 5000 mile radius of the location at latitude 60.790 North and longitude 97.570 East and find all tsunamis that they might have caused.*"

The keywords in the above information request are *earthquake, epicenter, radius, location, latitude, longitude,* and *tsunamis.* Just viewing these keywords, any reader might guess the verbs and the semantics that connect the meanings of these keywords. The challenge is to make the system understand with the use of the knowledge space. Let

us assume the following results for the keyword matches on the subject, object and the predicate of all the triples in the DAML+OIL ontologies

Table 2.1: Example 1

Keyword	Ontologies
Earthquake	earthquake.daml, damage.daml
latitude, longitude, location	location.daml, weather.daml, earthquake.daml
Epicenter, radius	earthquake.daml, circle.daml
Tsunami	tsunami.daml

Figure 2.5 shows the simplified versions of ontologies, which are used in this example. After the above results are obtained, we have to arrive at the semantically relevant set of ontologies. This is done by the comparing utility where every KObject in the ontology is compared with every other KObject in the other ontologies. If they have a common KObject linked by both the KObjects, they are related (e.g., *tsunamis* and *earthquakes* are related because they have a common parent, i.e., a KObject, namely *disaster*). In our example we have six ontologies matching seven keywords. Of these ontologies earthquake.daml, location.daml and tsunami.daml. are linked with KObjects latitude, longitude. So the relevant set of ontologies will be earthquake.daml, location.daml, tsunami.daml discarding damage.daml, weather.daml and circle.daml. The ontology weather.daml is discarded even though it has a reference to the definition of *location*. Because, although the definition of *weather* involves *country,* which is a sub-class of *location* it is not related in the sense it does not have a common parent with *earthquake* and *tsunami*. Thus, the system considers earthquake, tsunami, and location ontologies as relevant because this is the minimal set of ontologies with all keyword matches and has at least one common KObject linked.

This example also demonstrates the compare function, which compares two ontologies to trace the relation between them. For instance, consider ontologies earthquake.daml and tsunami.daml. Their links in knowledge space are latitude, longitude and disaster. This function can be effectively used for exploring inter-ontological relationships and to discover knowledge on that basis. The diagram below illustrates the ontologies and their relationships and how they are linked up in the knowledge space. This example is also illustrative of how the system can navigate between concepts in the knowledge space.
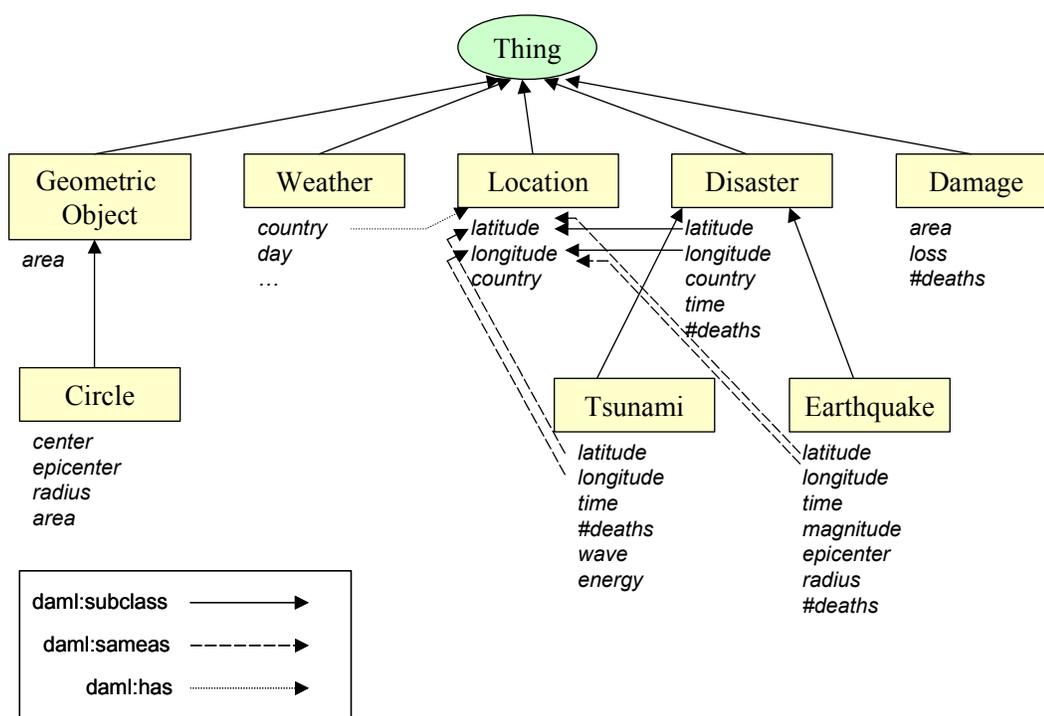

Figure 2.5: Example of Semantic Search and Exploration of

Inter-Ontological Relationships

Example 2: Consider the information request

"List the venues of all games played by the Bulldogs this fall"[1].

The keywords in the above information request are venue, game, bulldog and fall. Let us assume the following results for the keyword matches

Table 2.2: Example 2

Keyword	Ontologies
venue	Event.daml
game	Game.daml
bulldog	Dog_varieties.daml, uga_sport.daml
fall	Seasons.daml, waterfall.daml

.

The relevant set of ontologies will be Event.daml, Game.daml, uga_sport.daml and Seasons.daml, because all these ontologies have at least one common KObject linked. The listing of the example ontologies shows the common KObjects that match the keywords.

Event.daml

<rdf:RDF

xmlns:rdf ="http://www.w3.org/1999/02/22-rdf-syntax-ns#"

xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"

xmlns:daml="http://www.daml.org/2001/03/daml+oil#"

---

[1] Bulldog is the official mascot of the University of Georgia (UGA).

```
 xmlns     ="http://localhost/event#"

>


<daml:Ontology rdf:about="">
 <daml:versionInfo>$Id:   event.daml,v  1.0   2001/11/02   16:38:38   madhan   Exp
$</daml:versionInfo>
 <rdfs:comment>
   An event ontology
 </rdfs:comment>
 <daml:imports rdf:resource="http://www.daml.org/2001/03/daml+oil"/>
</daml:Ontology>


<daml:Class rdf:ID="Event">
 <rdfs:label>Event</rdfs:label>
 <subClassOf resource="http://www.daml.org/2001/03/daml+oil#Thing"/>
</daml:Class>


<daml:Property ID="venue">
 <rdfs:label>occured at</rdfs:label>
 <domain resource="#Event"/>


<rangeresource="http://www.cs.umd.edu/projects/plus/DAML/onts/general1.0.daml#Lo
cation"/>
</daml:Property>


<daml:Property ID="participant">
 <rdfs:label>had participant</rdfs:label>
```

```
    <daml:domain resource="#Event"/>
    <daml:range

resource="http://www.cs.umd.edu/projects/plus/DAML/onts/general1.0.daml#Agent"/>
</daml:Property>


</rdf:RDF>
```

Gams.daml
```
<rdf:RDF
  xmlns:rdf ="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
  xmlns:daml="http://www.daml.org/2001/03/daml+oil#"
  xmlns    ="http://localhost/game#"
>


<daml:Ontology rdf:about="">
  <daml:versionInfo>$Id:  game.daml,v  1.0  2001/11/02  16:38:38  madhan  Exp
$</daml:versionInfo>
  <rdfs:comment>
    A Game ontology
  </rdfs:comment>
  <daml:imports rdf:resource="http://www.daml.org/2001/03/daml+oil"/>
</daml:Ontology>


<daml:Class rdf:ID="Game">
  <rdfs:label>Game</rdfs:label>
  <subClassOf resource="http://localhost/event#Event"/>
```

```
</daml:Class>


<daml:Property ID="Teams">
 <rdfs:label>has Teams</rdfs:label>
 <daml:domain resource="#Event"/>
 <daml:range resource="http://localhost/event#Agent"/>
</daml:Property>


</rdf:RDF>
```

Uga_sport.daml

```
<rdf:RDF
 xmlns:rdf ="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
 xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
 xmlns:daml="http://www.daml.org/2001/03/daml+oil#"
 xmlns:xsd ="http://www.w3.org/2000/10/XMLSchema#"
 xmlns    ="http://localhost/uga_sport#"
>


<daml:Ontology rdf:about="">
 <daml:versionInfo>$Id: uga_sport.daml,v 1.0 2001/11/02 16:38:38 madhan Exp
$</daml:versionInfo>
 <rdfs:comment>
   A UGA sports ontology
 </rdfs:comment>
 <daml:imports rdf:resource="http://www.daml.org/2001/03/daml+oil"/>
</daml:Ontology>
```

```
<daml:Class rdf:ID="football">
  <rdfs:label>Football</rdfs:label>
  <subClassOf resource="http://localhost/game#Game"/>
</daml:Class>


<daml:Property ID="Mascot">
  <rdfs:label>bulldog is the Mascot</rdfs:label>
<mascot_name><xsd:string rdf:value="bulldog"/></mascot_name>
</daml:Property>


<daml:DatatypeProperty rdf:ID="mascot_name">
  <rdfs:comment>
    mascot name is a DataType Property whose range is mascot_name
  </rdfs:comment>
  <rdf:type rdf:resource="http://www.daml.org/2001/03/daml+oil#UniqueProperty"/>
  <rdfs:range                rdf:resource="http://www.daml.org/2001/03/daml+oil-ex-
dt#mascot_name"/>
</daml:DatatypeProperty>


</rdf:RDF>
```

After the relevant ontologies are obtained from the P2P network, the personalization techniques re-rank the results according to the user's profile and query history. The user then continues with the selection of ontologies and other tasks in the IScape Builder. The IScape Builder is a stand-alone Java application that provides a graphical interface to create and execute IScapes. In the Figure 2.6, a screen shot of Step 1 of this builder is

provided. It enables users to provide the keywords for searching ontologies. Then the user can select the ontologies s/he wants to use in the IScape. The next step for constructing the IScapes is to define user-defined relationships based on the selected ontologies. The rest of the steps (Steps 2, 3, etc.) involve choosing user-defined functions (*e.g.,* operators such as finding distance between two physical locations), and setting run-time parameters for the information request. The tool has been built over the IScape Builder in our previous version of InfoQuilt [STP01]. The IScape query optimization and execution details are discussed in our previous work [PS01].

Figure 2.6: IScape Builder Tool

2.6 Conclusions

The web is an overdose of scattered data. The goal of Semantic Web is to build a web of meaning. When knowledge is applied to data, we make data meaningful and get useful information. Knowledge is also important for making data understandable by programs. We need a platform for creating, exchanging and sharing of knowledge so that we can obtain more useful information. We provide one initial step towards addressing the above challenge in two parts: bring semantic capabilities of defining complex relationships, inter-ontological relationship, and knowledge discovery into the Semantic Web framework, and investigate how P2P can help in knowledge sharing. The key issues we have discussed include the following:

- Knowledge space construction and maintenance,

- Knowledge discovery based on semantic searching for relevant ontologies,

- Multi-ontology based semantic knowledge sharing, and exploration of new relationships between different ontologies,

- A peer-to-peer infrastructure in realizing ontology sharing and dissemination of knowledge on the Semantic Web.

A prototype implementation has been completed for the capabilities discussed in this paper, and several examples can be demonstrated now. Looking ahead, we have the following agenda that can enhance the contribution made by the current work:

- The algorithm for knowledge navigation paired with an effective reasoning tool can help us explore complex relationships with ease.

- When semantically searching for a related concept, instead of simple keyword matches in the knowledge space, use of fuzzy logic can bring more meaningful results.

- Knowledge space, just like today's web can become too big and explosive, so that there might be some interesting concepts, resources, or services that a user is not

aware of. The research can be focused on the knowledge space exploration based on user interests.

- In searching for a concept in knowledge space, the notion of a concept being related to another concept is defined by how close their common parents are to each other. The future research can be focused on choosing a criterion of closeness.

2.7 References

[B99]   Tim Berners-Lee (with Mark Fischetti), Weaving the Web, The original design and ultimate destiny of the World Wide Web, Harper, 1999.

[BHL01]   T. Berners-Lee, J. Hendler, and O. Lassila, "The Semantic Web: A new form of Web content that is meaningful to computers will unleash a revolution of new possibilities," Scientific American, May 2001.

[COV01]   R.Cover, DARPA Agent Mark Up Language (DAML) by http://xml.coverpages.org/daml.html

[Dam]   http://www.daml.org/2001/03/daml+oil-index.html

[DHM+01]   G. Denker, J. R. Hobbs, D. Martin, S. Narayanan, R. Waldinger, Accessing Information and Services on the DAML-Enabled Web Proceedings of the Second International Workshop on the Semantic Web SemWeb'2001 Hongkong, China, May 1, 2001

[EFG+99]   E. Motta, D. Fensel, M. Gaspari and R. Benjamins, Specifications of Knowledge Components for Reuse, in Proc. of The 11 th Int. Conf. on Software Engineering and Knowledge Engineering, Kaiserslautern, Germany (June 1999), KSI Press, pp. 36-43.

[ERN01]   J. Ernst, Peer-to-Peer infrastructure supporting the semantic web , Semantic Web Working Symposium, 2001.

[FM01]   D. Fensel and M. Musen, Eds. "The Semantic Web: A Brain for Humankind," IEEE Intelligent Systems, March/April 2001

[GAR01]   Gartner Consulting, The Emergence of Distributed Content Management and Peer-to-Peer Content networks, January 2001

[GHI+01]   S. Gribble, A. Halevy, Z. Ives, M. Rodrig, D. Suciu, What Can Peer-to-Peer Do for Databases, and Vice Versa? WebDB Workshop on Databases and the Web, June 2001.

[KS94]   V. Kashyap and A. Sheth, "Semantic based Information Brokering," Proceedings of the 3rd Intl. Conference on Information and Knowledge Management, November 1994, pp. 363-370.

[LSA+01]   Tarcisio Lima 1,2 , Amit Sheth 3 , Naveen Ashish, Mukesh Guntamadugu, Sriram Lakshminarayan 4 , Narayanan Palsena, and Dilpreet Singh, Digital Library Services Supporting Information Integration over the Web, WIIW 2001

[MKSI96]   E. Mena, V. Kashyap, A. Sheth, A. Illarramendi, "OBSERVER: An Approach for Query Processing in Global Information Systems Based on Interoperation Across Pre-existing Ontologies", Proceedings of the 1st IFCIS Intl. Conference on Cooperative Information Systems (CoopIS'96), Brussels, Belgium, June 1996, pp. 14-25.

[MIKS00]   E. Mena, A. Illarramendi, V. Kashyap, and A. P. Sheth. OBSERVER: An approach for query processing in global information systems based on interoperation across pre-existing ontologies. International Journal on Distributed and Parallel Databases, Vol. 8, No. 2, pp. 223-271, April 2000.

[Mus]   http://www.musiccity.com

[MUS92]   Dimensions of Knowledge Sharing an reuse Mark A. Musen COMPUTERS AND BIOMEDICAL RESEARCH 25, 435–467 (1992)

[MWD01]   Prasenjit Mitra, Gio Wiederhold and Stefan Decker: A Scalable Framework for Interoperation of Information Sources. Proceedings of the 1st International Semantic Web Working Symposium (SWWS `01), Stanford University, Stanford, CA, July 29-Aug 1, 2001, Jul. 2001

[Nap]   http://www.napster.com

[NFF+91]   R. Neches, R. Fikes, T. Finin, T. Gruber, R. Patil, T. Senator, and W. R. Swartout, Enabling Technology for Knowledge Sharing, AI Magazine, Volume 12, No. 3, Fall 1991

[PB99]   A.G. Perez and R. Benjamins (1999). Overview of Knowledge Sharing and Reuse Components: Ontologies and Problem-Solving Methods. IJCAI-99 workshop on Ontologies and Problem-Solving Methods (KRR5), Stockholm, Sweden, August 2, 1999.

[Pee]   http://www.peer-to-peerwg.org/whatis/index.html

[PS01]   Patel, S. and Sheth, A. Planning and optimizing Semantic Information Requests using domain modeling and resource characteristics, Sixth International Conference on Cooperative Information Systems (CoopIS 2001)2001.

[Ref]   DAML+OIL [March 2001] reference description http://www.daml.org/2001/03/reference.html

[STP01]   A. Sheth, S. Thacker, and S. Patel, Complex relationships and Knowledge discovery support in the InfoQuilt System, Technical Report, LSDIS Lab, Univ. of Georgia, September 2001 (submitted for publication).

[TSP01]   S. Thacker, A. Sheth and S. Patel, "Complex Relationships for the Semantic Web," Creating the Semantic Web, D. Fensel, J. Hendler, H. Liebermann, and W. Wahlster (eds.), MIT Press, 2001 (to appear).

CHAPTER 3

CONCLUSION AND FUTURE WORK

3.1 Summary and Contributions

The web is an overdose of scattered data. We discussed the motivation and the need for knowledge in order to realize the goal of Semantic Web. We also discussed the challenges that lie in building Semantic Web and how we require a platform for creating, exchanging and sharing of knowledge so that we can obtain more useful information. We provided one initial step towards addressing the above challenge in two parts: bringing semantic capabilities of defining complex relationships, inter-ontological relationship, and knowledge discovery in the Semantic Web framework, and investigating how P2P can help in knowledge sharing. The key issues we have discussed include the following:

- Knowledge space construction and maintenance,

- Knowledge discovery based on semantic searching for relevant ontologies,

- Multi-ontology based semantic knowledge sharing, and exploration of new relationships between different ontologies,

- A peer-to-peer infrastructure in realizing ontology sharing and dissemination of knowledge on the Semantic Web,

- Details about the prototype implementation that has been completed for the capabilities discussed in this thesis, and examples for demonstration.

3.2 Future Work

Looking ahead, we have the following agenda that can enhance the contribution made by the current work:

- The algorithm for knowledge navigation paired with an effective reasoning tool can help us explore complex relationships with ease.

- When semantically searching for a related concept, instead of simple keyword matches in the knowledge space, use of fuzzy logic can bring more meaningful results.

- Knowledge space, just like today's web can become too big and explosive, such that there might be some interesting concepts, resources, or services that a user is not aware of. The research can be focused on the knowledge space exploration based on user interests.

- In searching for a concept in knowledge space, the notion of a concept being related to another concept is defined by how close their common parents are to each other. The future research can be focused on choosing a criterion of closeness.

- Knowledge space is extremely large with more than one route to access the relevant information. Research can identify techniques for shortest, fastest and most efficient ways of discovery user-interested knowledge.

- Web Services define the dynamic behavior of knowledge. The knowledge spcace construction and maintenance algorithms can be enhanced to accommodate the dynamic behavior of knowledge with limited resources and context

- Natural Language Processing techniques can be used to automatically extract keywords from the user's information request

- With the use of DAML+OIL, ontologies now accommodate both assertions and definitions. Hence all the resources are now annotated in DAML+OIL describing their content. DAML+OIL Enabled agents can *understand* this content and hence the query processing and optimization modules need to be reprogrammed to support evolving standards in Semantic Web.

- The semantic searching of ontologies can be enhanced if we use the relationships between different terms involved in the information request.