

KNOWLEDGE SERVICES ON THE SEMANTIC WEB

Developing infrastructures for trading knowledge services using semantic technologies.

The first phase of knowledge management (KM), in which companies institutionalized knowledge creation, storage, and sharing through internal KM initiatives, is almost complete. A central tenet of this article is that the time is

BY GREGORIS MENTZAS, KOSTAS KAFENTZIS,
AND PANOS GEORGOLIOS

ripe for companies to revise their knowledge agenda and start planning how they will externally exploit their knowledge assets and codified intellectual capital in order to exploit new opportunities and enter the second phase of KM, which is commonly called the “knowledge commerce” era [8].

The current focus of internal KM initiatives in organizations is one of identifying, leveraging, and sharing knowledge more widely. Better management of this knowledge has been used to improve business processes, increase productivity, reduce new product development times, and has led into the massive accumulation and storage of knowledge produced through organizational operations [1, 6]. As companies are becoming aware of the fact that they are part of a complex network of connections with their partners and customers their focus has shifted toward expanding the KM concept externally: they explore new ways to cultivate and exploit knowledge sharing with customers, suppliers, and partners.

Beyond knowledge sharing, organizations are increasingly converting internally generated knowledge into viable knowledge-based products and services that can be provided and traded externally. This trend has been attributed to reasons such as the preference of managers from knowledge-seeking companies for outsider knowledge and is justified on cost grounds [12]. This is especially true in cases when knowledge can be packaged and hence becomes portable and migrant and when it facilitates problem solving and addresses efficiency problems [9].

Interorganizational knowledge exchanges may take many different forms depending on the tradability of the streams of knowledge that flow among organizational entities and on how open the exchange is for new members [7]. Here, we focus on knowledge exchanges that treat knowledge assets as tradable goods, that is, in knowledge supplies and knowledge markets; see Table 1.

In this article we present a Semantic Web-enabled architecture for trading knowledge assets. We develop



Table 1. Categories of knowledge sharing and trading exchanges.

WEB-ENABLED KNOWLEDGE SERVICES

Knowledge movement within and across organizations is powered by market forces similar to those that animate markets for other, more tangible goods [1]. Like markets for goods and services, the knowledge market has buyers, sellers, and brokers, as well as market pricing and exchange mechanisms, even though money is rarely the form of payment [2]. An in-depth review of existing knowledge markets identified three critical points for their development [4].

First, knowledge has manifold complex context and content features, which determine its applicability and usefulness in a given situation; thus knowledge assets can not be described and retrieved with simple keywords; more rich representations are required that take into account these features. Second, in electronic knowledge trading one cannot simply copy

THE FACT THAT A KNOWLEDGE MARKET IS A MEETING POINT OF AGENTS WITH DIFFERENT LANGUAGES, MENTAL MODELS, AND WORLD PERCEPTIONS SIGNIFICANTLY INCREASES ITS COMPLEXITY.

the notion of “knowledge services” as an approach to commercially exploit a company’s knowledge assets, offer competitive advantage and extend market reach. Such knowledge services may seem to be initially applicable to organizations that have traditionally based their business models to information and/or knowledge-centric activities, such as publishing, software, education, research, or consulting companies [9]. However, even traditional service and manufacturing organizations have a wealth of knowledge assets to expose and trade externally in order to radically improve their competitive position [8].

ways of working that are already known from traditional business, but should exploit the strength of multiple synchronous and asynchronous communication means. Third, the technical, business, and organizational mechanisms for managing and maintaining an electronic knowledge market cannot be derived from a simple adaptation of conventional e-commerce paradigms; rather they should explicitly take into account the need for supporting shared dialogues between participants and focus on the need for long-term knowledge partnerships, rather than short-term buyer-supplier relationships.

These issues are crucial for knowledge trading. The fact that a knowledge market is a meeting point of agents with different languages, mental models, and world perceptions significantly increases its complexity. Sharing a common understanding of the needs and the knowledge that meet these needs depends on the degree of formalization of the language used within the specific community served by a knowledge market. A rich representation of the problem and solution space is necessary.

In this article, we address these several points by developing the concept of Web-based “knowledge services,” proposing an ontology-based framework and outlining a service-oriented implementation infrastructure that adopts the Semantic Web services paradigm as the underlying technological architecture [5].

The following is a simple example of a knowledge service. Consider a business development manager in the automotive industry who examines the opportunity of setting up a representative office in China, since she believes that through a representative office she can obtain useful information such as consumer patterns and behavior, market demand, market practices, local laws and regulations, local costs, and so forth. However, setting up a representative office raises multiple legal and financial issues: Should the office be a joint venture and if yes, which of the two types allowed in China (equity joint venture or cooperative joint venture)? Which are the local administrative processes for setting up such a joint venture? Which is the appropriate local authority for registering the office? Which are the tax and labor management regulations for such a venture? The manager seeks solutions to these problems. Turning within the internal organizational surroundings she does not find any support-

ive evidence so she turns to external knowledge sources. The needed knowledge may come in the form of a roadmap to setting up new offices in China (such knowledge is usually provided by chambers of commerce), best practices from other firms that have undertaken similar endeavors in the past (such knowledge is usually provided by specialized consultancies either in document form or as personalized advice), or in the form of pointers to the contact data of knowledgeable freelance consultants that have experience in setting up businesses in China. What

the manager really needs is a “knowledge service” that will orient him or her toward the appropriate knowledge objects. These knowledge objects need to be discovered, retrieved, evaluated, selected, their acquisition has to be negotiated and their delivery monitored.

The most suitable environment for technologically supporting Web-enabled knowledge provision services is the use of Semantic Web services. The semantic descriptions of such services allow external agents to understand their functionality and internal structure so that they can discover, compose, and invoke such services. In this area, we

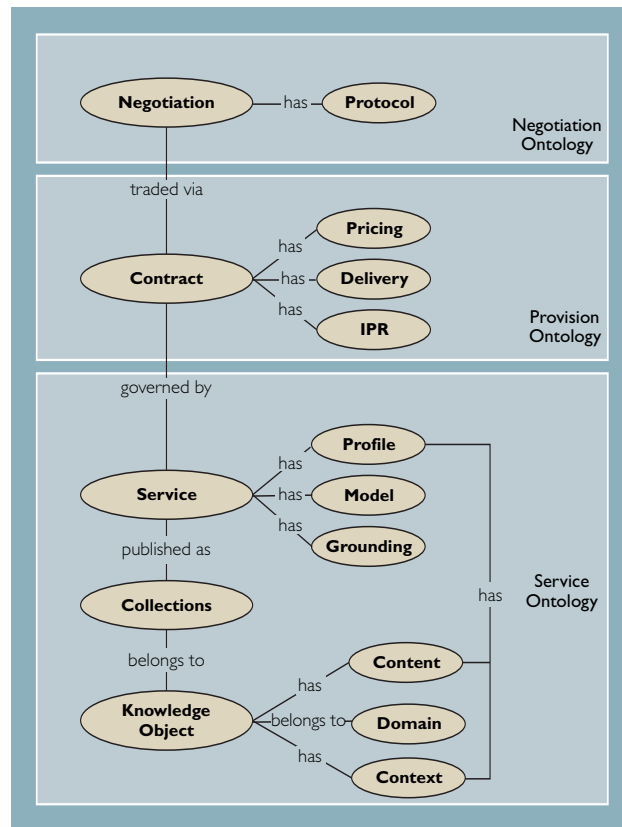


Figure 1. Ontological structure of tradable knowledge services.

should note the recent work of the Semantic Annotations for WSDL (SAWSDL) Working Group of the W3C, whose objective is to develop a mechanism to enable semantic annotation of Web services descriptions [11]. In our work we developed multifaceted ontological structures in order to define the necessary modeling primitives that are important for describing knowledge provision services that go beyond common Web services like a flight booking or book selling.

Let us revisit the example of our business development manager to pinpoint our approach more eloquently. The central question is as follows: how could the manager use a knowledge service that would retrieve all the necessary knowledge objects and help to resolve a particular problem? A highly expressive way that respects the contextual information of the required knowledge is necessary. In this article, we

argue that this could be provided by a knowledge object ontology that describes the specific concepts to be used during knowledge retrieval. Furthermore, the codification of knowledge about human expertise is necessary (for example, for the retrieval of knowledgeable people that could provide advice on the problem).

In our example, the manager would ask from the service a document of the type ‘roadmap’ or ‘best practice’ that could be used to set up a ‘new representative office’ in the country ‘China’ and the ‘Automobile’ industry. All the parameters in single quotes should be part of the knowledge represented in the knowledge object ontology and could be used by the knowledge provision service to retrieve the proper knowledge object(s). The same parameters could be used to infer which expert (such as a specialized business analyst or a lawyer with knowledge of Chinese business law) could provide useful advice. Let us further assume that the required knowledge has been discovered. How would the manager negotiate either the one-off purchase of some critical knowledge objects, or initiate a request for quotes for buying consulting time from experts in the field? In order to address these issues we introduce the knowledge service ontology.

A SEMANTIC WEB INFRASTRUCTURE FOR KNOWLEDGE TRADING

This article introduces a multifaceted reference ontology to model the semantics of knowledge services and objects (see Figure 1). The use of ontologies allows for a formal and shared understanding of the critical concepts and the relations between them that are vital for creating integrative views of knowledge trading. We have developed three interconnected ontologies that describe the various entities that participate in

Ontology	Facets	Description
Service	Content	• Describes the content of a knowledge object and content attributes like the kind of content, what it is about, how it is physically manifested, and so forth.
	Context	• Describes the application context in which a particular knowledge object can be used (user organization, organizational roles that may apply a knowledge object).
	Domain	• Provides vocabularies about concepts of the application domain (medical, law, engineering) and their relationships.
Provision	IPR	• Defines the legal framework that is required to accomplish legally correct transactions with regard to the complex issues of digital rights in the inter-organizational, and often multinational setting of knowledge trading.
	Pricing	• Specifies pricing and payment issues. It includes all information required for supporting the trading aspects of knowledge services, namely information on applicable pricing schemes and negotiation schemes, payment mechanisms as well as preconditions for applying the different possibilities. • In case trading does not involve negotiation, a fixed price or a subscription model can be followed.
	Delivery	• Contains information about the user community that may enjoy the use of a service and the way that the service is delivered (which includes attributes mainly featured in service-level agreements).
Negotiation	Protocol	• Provides a general framework that permits negotiation participants to reach agreement by establishing a shared understanding of the negotiation “mechanics.”

Table 2. Ontology-based modeling of knowledge services.

knowledge transactions: a service ontology; a provision ontology; and a negotiation ontology.

We model the semantics of knowledge objects with two facets/sub-ontologies; the content facet that describes the content itself (it provides information about what and how) and the context facet that represents the application context of the knowledge object (it provides information about when, where, by whom, and so forth). Domain-specific ontologies

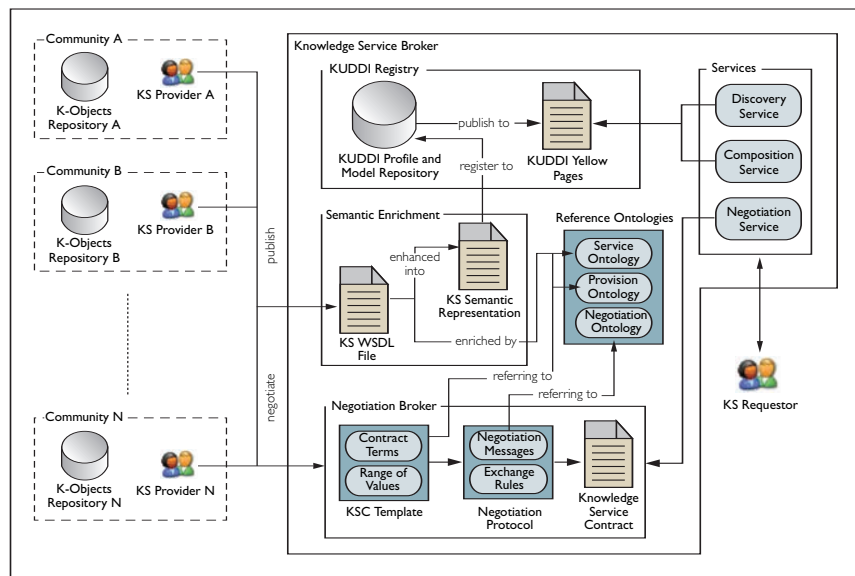


Figure 2. Architecture for trading knowledge services (KS: knowledge service; KSC: knowledge service contract).

are used to further structure the particular application domain by providing domain-dependent semantics that allow for discrimination and classification of the knowledge objects that a knowledge service will provide access to. A set of knowledge objects that is accessed by a single point is called a collection. A collection can be the portion of a database, a folder in a file system, or other elements that can be wrapped and published by a knowledge service that in turn provides unified access to

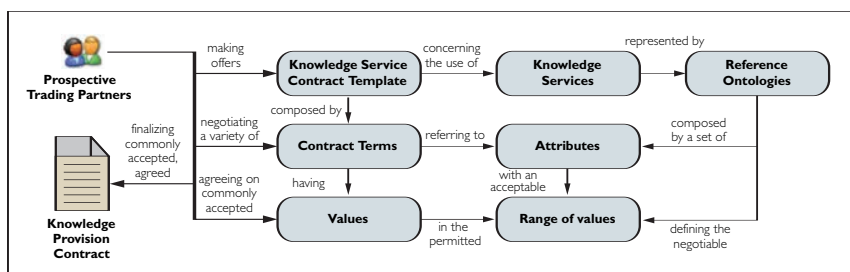
the knowledge objects of the collection.

A knowledge service can then be conceptualized as a means to provide access to a collection of knowledge objects, and is accompanied by a set of operations that can be performed on this resource. These operations include discovery, navigation, retrieval, and interaction. The knowledge service utilizes the content and context ontology for a twofold purpose: to discover knowledge objects within a collection and to be discovered as a service, namely to determine its identity. The latter is achieved by extending the profile element of a Semantic Web service by incorporating concepts

of reference ontologies that semantically enrich the Web services infrastructure [3].

We incorporate k-UDDI within an architecture we call the Knowledge Service Broker. Besides the k-UDDI, the broker includes a service publishing component that allows knowledge providers to annotate their services based on the reference ontologies and a negotiation broker that applies a negotiation ontology in order to facilitate providers and requesters to conclude service provision contracts (see Figure 2).

With the help of the knowledge service reference ontology we have a means of publishing knowledge repositories through Web services. Since more than one knowledge repository (by various service providers) may be published, a means of discovering the proper knowledge service is required. In the traditional Web services environment this role is played by UDDI, while in our approach the k-UDDI embeds ontology-based metadata in the tModel keys of the traditional UDDI. The k-UDDI



from the content and context ontologies. Therefore, a knowledge service instance is aware of the knowledge objects that it “represents” by inheriting their top-level content and context attribute values.

A knowledge service is traded on the basis of a contractual agreement. A contract is the product of negotiation among the entities that are involved in the transaction and defines the terms of usage of the knowledge service. The contract is conceptually a container of the entire set of elements that define a legal transaction of a knowledge service. The elements that are negotiated in a knowledge service contract are defined in relevant ontologies, which include: the IPR ontology, where intellectual property issues are defined; the Delivery ontology, where access and delivery issues are handled; and the Pricing ontology, which specifies pricing and payment issues (see Table 2).

In order to enable knowledge providers (individuals, groups, and organizations) to trade or share the knowledge that is stored in various forms in distributed repositories within their organizations, our architecture employs a Web services infrastructure and ontologically enhances it so the life cycle of a knowledge transaction is supported. We have specified an enhanced universal discovery, description, and integration (UDDI) platform known as k-UDDI, which enables the discovery, negotiation, and invocation of knowledge services with the incorpora-

Figure 3. Negotiation protocol.

holds all reference ontologies that allow a common understanding of services and facilitate semantically enhanced service discovery, IPR and business specific issues and finally negotiation processes generating sound contracts. Knowledge service discovery is provided by the discovery service of the registry, which is exposed via a Web service interface. Queries to the registry are ontology-enabled and made feasible by using the same conceptual model as when querying single repositories for knowledge objects [3].

As knowledge services will be traded, mechanisms are needed to support negotiation and contracting tasks. We make use of our negotiation ontology and develop a flexible negotiation mechanism that enables bargaining between the service provider and requester concerning the terms and conditions of use of a knowledge service. Our approach follows a one-to-one negotiation protocol enabling the bargaining of a variety of terms of a predefined contract template. Our protocol defines a set of suggested Negotiation Messages and Documents, exchanged in the frame of the negotiation process and associated with the negotiation of the contents of the knowledge service contract (KSC); see Figure 3. Our approach is based on [10], while the core of our schema is based on the “Automated negotiation of collaboration-protocol agreements,” which was proposed by OASIS, the Organization for the Advancement of Structured Information Standards.

A scenario that shows the functionality of our solu-

tion is as follows. Knowledge service providers annotate their services with semantic information following the reference ontology that is defined in OWL and then publish them in the k-UDDI. The service requester creates a query to the registry describing the type of service he or she is interested in. The discovery mechanism of the registry searches the metadata of the published services that follow the reference ontologies. The registry proposes the retrieved services to the requester who ends up with a specific service to invoke. If the service is negotiable, the requester initiates a negotiation process that follows the specific protocol the provider has selected and they both tune the contract attributes, defined in the provision ontology. The contract is validated against the contracting schema and if accepted the service is invoked.

As internal knowledge services have been shown to be asymmetrical and resource intensive, especially in speedy problem-solving project-based environments [9], it is expected that external knowledge provision will be quite favorable. This will justify the move away from pay-as-you-go and subscription-based membership models to flexible pricing structures in which negotiation mechanisms like those described here will be indispensable.

CONCLUSION

Sustaining enterprise development and retaining competitive advantage in the knowledge-based economy necessitates that firms focus on their knowledge assets. This vanguard is impossible to be ensured when an organization functions as a closed system. Importing knowledge from sources lying outside the organizational boundaries and harnessing knowledge across cross-organizational networks is critical. Related business models, together with emerging technologies like Web services and the Semantic Web, provide ample opportunities to develop appropriate infrastructures for online trading of knowledge goods and associated services.

We have introduced the concept of knowledge trading and analyzed an ontology-based approach for trading knowledge services using semantic technologies. This approach addresses the highly context-dependent nature of knowledge transactions and the need to support flexible negotiation mechanisms for online trading. Our insights have implications for practitioners, as various models of exposing internal knowledge can be adopted, especially in information-centric environments. For example, the choice between alternative partnership options such as short-term buyer-supplier relationships, long-term knowledge partnerships, or multi-directional knowledge

exchanges is highly critical. For researchers, the vision of semantically enabled trading of commoditized knowledge raises many questions, such as the need to further support community modalities of exchange and develop secure infrastructures for publishing internal organizational knowledge. Organizations that extend their KM agenda to include the incorporation of external knowledge provision and acquisition can reap significant benefits and realize first-mover advantages in the coming knowledge commerce era. **C**

REFERENCES

1. Davenport, T.H. and Prusak, L. *Working Knowledge: How Organizations Manage What They Know*. Harvard Business School Press, Boston, MA, 1998.
2. Desouza, K.C., Awazu, Y., Yamakawa, S., and Umezawa, M. Facilitating knowledge management through market mechanisms. *Knowledge and Process Management* 12, 2 (2005), 99–107.
3. Georgiolos, P., Kafentzis, K., and Mentzas, G. Knowledge provision with intelligent e-services. *International Journal of Intelligent Systems, Special Issue on E-service Intelligence* 22 (2007), 501–518.
4. Kafentzis, K., Mentzas, G., Apostolou, D., and Georgiolos, P. Knowledge marketplaces: Strategic issues and business models. *Journal of Knowledge Management* 8, 1 (2004), 130–146.
5. McIlraith, S., Son, T.C., and Zeng, H. Semantic Web services. *IEEE Intelligent Systems* (Mar./Apr. 2001), 46–53.
6. Mentzas, G., Apostolou, D., Abecker, A., and Young, R. *Knowledge Asset Management: Beyond the Process-centred and Product-centred Approaches*. Springer, London, 2002.
7. Mentzas, G., Apostolou, D., Kafentzis, K., and Georgiolos, P. Interorganizational networks for knowledge sharing and trading. *Information Technology and Management, Special Issue on Information Sharing Across Multiple Organizations* 7, 4 (2006), 259–276.
8. Skyrme, D.J. *Capitalizing on knowledge: From e-business to k-business*. Butterworth-Heinemann, London, 2001.
9. Smyth, H. and Longbottom, R. External provision of knowledge management information services: The case of the concrete and cement industries. *European Management Journal* 23, 2 (Apr. 2005), 247–259.
10. Tamma, V., Wooldridge, M. and Dickinson, I. An ontology for automated negotiation. In *Proceedings of the Workshop on Ontologies in Agent Systems* (Bologna, Italy, July 2002).
11. Verna, K. and Sheth, A. Semantically annotating a Web service. *IEEE Internet Computing* (Mar.–Apr. 2007), 83–85.
12. Zucker, L., Darby, M.R., and Armstrong, J.S. Commercializing knowledge: University science, knowledge capture, and firm performance in biotechnology. *Management Science* 48, 1 (Jan. 2002), 138–153.

GREGORIS MENTZAS (gmentzas@mail.ntua.gr) is a professor in the School of Electrical and Computer Engineering at the National Technical University of Athens and the director of the Information Management Unit at the Institute of Communication and Computer Systems, Athens, Greece.

KOSTAS KAFENTZIS (kkafe@softlab.ntua.gr) is a doctoral candidate in the School of Electrical and Computer Engineering at the National Technical University of Athens, Greece.

PANOS GEORGIOLOS (pgeorgol@softlab.ntua.gr) is a doctoral candidate in the School of Electrical and Computer Engineering at the National Technical University of Athens, Greece.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright of Communications of the ACM is the property of Association for Computing Machinery and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.