Ontology-Driven Management of Semantic Spaces

Reto Krummenacher

Digital Enterprise Research Institute, University of Innsbruck, Austria
reto.krummenacher@deri.at

Abstract. Recent work in the field of middleware technology proposes semantic spaces as a tool for coping with the *scalability*, *heterogeneity* and *dynamism* issues of large scale distributed environments. Reflective middleware moreover offers answers to the needs for adaptivity and self-determination of systems. Based on experiences with traditional middleware we argue that ontology-driven management is a major advancement for semantic spaces and provides the fundamental means for reflection. In this research we look at ontology-based metadata to turn semantic spaces into reflective middleware.¹

1 Introduction

With the transformation of today's Web to a mobile and even ubiquitous web of interactive computers and small physical devices the coordination of large numbers of autonomous nodes gets particularly challenging, and further *scalability*, *heterogeneity* and *dynamism* issues arise.² In consequence middleware solutions are required that can adapt to dynamic changes in application requirements and environmental conditions and customize its service to various end user devices. Reflective middleware is considered to provide essential answers in this respect [1,2,8].³

Self-representation – an explicit representation of the internal structure of the implementation that the middleware maintains and manipulates – is important in order to support reflection in form of inspection and adaptation [1]. A system is reflective when it is able to manipulate and reason about itself [8]. A critical concept in this respect, in particular in absence of central control, is the concept of metadata. Metadata is data about the properties, capabilities, and requirements of system elements to enable their coordination [6]. It is thus essential to develop metadata infrastructures for the self-representation of the space middleware.

Our proposed solution follows a recent trend in the field of middleware for large scale open systems: *semantic spaces*. Spaces are a powerful concept for the coordination of autonomous nodes. Instead of explicitly exchanging messages or performing remote procedure calls, communication is done by reading and writing distributed data structures in a shared space. Semantic space middleware is then particularly tailored to the Semantic Web and Semantic Web service environments. Consequently, ontology-based modeling of metadata becomes the natural choice. Moreover, ontologies provide the grounds for formal reasoning about the middleware implementation – a prerequisite for reflection.

¹ This work is supported by the projects TSC (tsc.deri.at) and TripCom (www.tripcom.org).

² cf. W3C Ubiquitous Web Domain (www.w3.org/UbiWeb/).

³ cf. Workshops on Adaptive and Reflective Middleware (www.ics.uci.edu/~arm06/).

2 Research Problem

Different semantic space proposals (Section 3) brought the interaction primitives and data models to maturity. However, the issues of network dynamics and scalability (use at Web scale) are barely addressed, although they are crucial to any installation. In other words, semantic spaces lack so far solutions to the increasing complexity of middleware's non-functional properties: decentralization, availability, reliability, fault-tolerance, scalability, and security. Moreover, it is important to take the use and application context into account in order to tailor the delivered service to the current needs [6]. The application of ontologies that describe the published data, the spaces and their interrelationships and characteristics is expected to be an effective instrument in this respect. Ontology-based metadata is thus essential for the management of the middleware and provides the formal grounds for reasoning, and hence reflection. This leads us directly to the main questions of our research.

Main Question: *what does ontology-based metadata have to model and how in order to turn semantic spaces into reflective and self-adaptive middleware for large scale open systems?*

Problem 1: *how to ontologize the space middleware and the data in order to provide reflective management of core non-functional properties, in particular distribution?*

Problem 2: *what requirements result from ontology-driven adaptation algorithms and which additional metadata modules are needed?*

Problem 3: *how can metadata be acquired and provided to participating nodes for adaptation of their routines and to improve the data distribution in particular?*

Improving the data distribution does not only decrease the discovery overhead, but influences the availability and scalability of the middleware (distribution by reasoning about the network and middleware) and the efficiency and quality of retrieval routines (distribution by reasoning about the usage and needs).

3 State of the Art and Related Activities

Within this paper we look at related activities in two domains: (1) the emerging field of semantic space middleware and their approaches to ontology-driven management (reflection is to our knowledge not addressed so far at all), and (2) the use of ontologies and metadata for the management of middleware in general.

Semantic spaces: First ideas for triplespace computing (a particular approach to semantic spaces) have been proposed in [3], as middleware for Semantic Web services. This and subsequent work mostly neglected the use of ontologies for the management tasks. As example, the TSC project [4] only knows limited meta-information about the published data. It fails however to define the way the space is structured and behaves; thus the middleware remains a black box without self-representation.

Semantic Web Spaces [11] has conceived a generic lightweight coordination middleware for sharing and exchanging semantic data on the Web. This project includes a brief outline of an ontology-based metadata model, however without solid investigation of more comprehensive application areas and use cases.

As a joint successor of TSC and Semantic Web Spaces the TripCom project was established [14]. TripCom spends significant effort on ontology and rule-driven management of spaces. Rules model the adaption strategies and decisions for the optimization of triplespaces. The work presented in this paper contributes to TripCom.

Ontologies for middleware management: Over the past years semantic technologies and tools have gained maturity and its applicability for the management of middleware [12] and distributed information systems [6] is acknowledged.

Based on the argument that most available ontologies for middleware suffer from conceptual ambiguity, poor axiomatization, loose design, and a too narrow scope, [12] developed a management ontology as extension to DOLCE [5], a foundational ontology. While [12] targets software components and (Web) services and thus has a different domain, some of the ideas and decisions provide a fruitful basis for our space management ontologies and self-representation strategies.

There are other relevant vocabularies to describe people and agents (FOAF), online communities (SIOC) or documents (Dublin Core).⁴ The Ontology Metadata Vocabulary [7] could be used to describe the schemes and topics a space adheres to, which is relevant for the discovery and clustering of semantically adjacent nodes. Such small and simple vocabularies are expected to provide many of the ground terms of our management ontologies, also because we emphasize on the importance of ontology reuse.

Furthermore, ontologies have gained momentum for modeling context [10]. Information about devices, network status and user intentions are important for adaptation. One early example of ours in this respect is the modeling of context rules with ASC [9,15]. ASC describes contextual data by taking into account quality of information, provenance, validity and satisfiability of information.

4 Research Methodology

The evaluation of the needs for metadata and the development of the formal models is an important part of this work. Requirements are deduced from our experience with semantic space middleware. First ideas for a semantic space ontology were outlined as part of TripCom [14]. Additionally we will consider the Semantic Web specific requirements outlined in [12], amongst others: awareness for Semantic Web languages, ontology modularization, interoperability, and verification. In [10] we compiled a further set of evaluation criteria dedicated to context modeling ontologies: traceability, comparability, logging and quality of data, and scalability of the ontology. Similar criteria form the basis for the development and theoretical evaluation of our models.

The metadata is first of all used to cluster data and to formally describe the decision rules that lead to the choice of the appropriate clustering algorithm. Distribution through semantic clustering provides the chosen 'proof-of-concept'. By applying the developed ontologies to a semantic space middleware implementation, we can evaluate our solutions against an existing infrastructure.⁵

⁴ FOAF: www.foaf-project.org, SIOC: www.sioc-project.org, and DC: www.dublincore.org

⁵ A demonstrator is available from TSC; a first prototype from TripCom.

The distribution algorithms for structuring the virtual spaces and for clustering the published data will be evaluated qualitatively. First, we compare the management effort and retrieval quality with a non-reflective implementation, by help of use cases. Moreover, we measure the quality of the clusters (e.g. divergence of members, size of clusters and difference of content), the management overhead of the algorithms, and the number of hops needed to resolve user requests - a more P2P-based criterion.

5 Expected Contribution

As outcome of this research we expect a metadata infrastructure consisting of an integrated set of ontologies. The ontologies will be developed alongside the criteria shortly mentioned in the previous section and tailored to the management processes present in semantic space middleware; in particular to the ones that profit from reflection and self-adaptation with a primary focus on distribution. As [13] pointed out it is advisable to approach the Semantic Web by small and simple ontologies that are easier adopted and reused in the large. Consequently we will develop a collection of well-integrated, but distinct ontologies for different aspects of our system: the representation of spaces and data, users and their context, descriptions of kernels, their functionality and also their connection to other kernels. We like to highlight one particular issue that is mostly neglected in distributed information systems [6] and that is, as pointed out before, essential for the installation of reflection:

In many scenarios the relevance of information depends on the usage purpose and its interpretation on the context at hand. This becomes especially eminent in distributed settings where data producers and consumers have different backgrounds. Context-awareness and personalization are indispensable instruments to increase the sensitivity of middleware.

In summary we develop ontologies for the management of semantic space middleware implementations for large-scale, open and dynamic systems with the goal to enable reflection. Therewith we expect to significantly contribute to the success of semantic space technology: our ontology-driven management procedures go beyond the ones of Trip-Com and other related projects by being adaptive with respect to changing user needs and environmental conditions. The potential adaptations to the structure and behavior of the middleware are encoded in decision rules and the management processes are thus at least partly automated. Improvement is in particular expected in settings where semantic spaces are applied to applications in large mobile or ubiquitous web settings.

6 Conclusion

While the conceptual work for semantic spaces becomes mature and the data models and interaction primitives are well defined, the given approaches to ontology-driven management are very much in its infancy. There are ontological vocabularies released for the modeling of spaces and the description of data. However, these ontologies do clearly neither address the requirements of distributed and dynamic systems nor the

additional needs of reflection and adaptation algorithms: "an explicit representation of the internal structure of the middleware implementation that the middleware maintains and manipulates" [8].

Our next steps include further analysis of the management processes and metadata vocabularies of existing distributed information systems and reflective middleware solutions. A number of reflective middleware implementations are discussed in [2]. Based thereon we will develop the set of ontologies necessary for our purposes and apply it to an existing semantic space implementation. This provides a solid proof of concept and allows showcasing our approach with respect to the non-functional property distribution: considering in particular the management overhead, availability of nodes and scalability. The findings with respect to data distribution are then easily generalized, as the use of reflection is known to improve other non-functional aspects too.

References

1. Coulson, G.: What is Reflective Middleware? *IEEE Distr. System Online* 2(8) (2001)
2. Curry, E.: Adaptive and Reflective Middleware. In: Mahmoud, Q. (ed.) *Middleware for Communications*, Wiley, Chichester (2004)
3. Fensel, D.: Triple-Space Computing: Semantic Web Services Based on Persistent Publication of Information. In: *IFIP Int'l. Conf. on Intelligence in Communication Systems*, pp. 43–53 (November 2004)
4. Fensel, D., Krummenacher, R., Shafiq, O., Kuehn, E., Riemer, J., Ding, Y., Draxler, B.: TSC - Triple Space Computing. *e&i Elektrotechnik und Informationstechnik* 124(1/2) (2007)
5. Gangemi, A., Guarino, N., Masolo, C., Oltramari, A., Schneider, L.: Sweetening Ontologies with DOLCE. In: *13th Int'l. Conf. on Knowledge Engineering and Management*, pp. 166–181 (October 2002)
6. Haase, P.: Semantic Technologies for Distributed Information Systems. PhD thesis, Universität Karlsruhe (2006)
7. Hartmann, J., Palma, R., Sure, Y., Suarez-Figueroa, M.C., Haase, P.: OMV - Ontology Metadata Vocabulary. In: *Workshop on Ontology Patterns for the Semantic Web* (November 2005)
8. Kon, F., Costa, F., Blair, G., Campbell, R.H.: The Case for Reflective Middleware. *Communications of the ACM* 45(6), 33–38 (2002)
9. Krummenacher, R., Lausen, H., Strang, T.: On the Modeling of Context-Rules with WSML. In: *Workshop Contexts&Ontologies: Theory, Practice and Applications* (August 2006)
10. Krummenacher, R., Lausen, H., Strang, T.: Analyzing the Modeling of Context with Ontologies. In: *Int'l. Workshop on Context-Awareness for Self-Managing Systems* (May 2007)
11. Nixon, L., Simperl, E., Antonenko, O., Tolksdorf, R.: Towards Semantic Tuplespace Computing: The Semantic Web Spaces System. In: *22nd ACM Symp. on Applied Computing*, ACM Press, New York (March 2007)
12. Oberle, D.: *Semantic Management of Middleware*. Springer, Heidelberg (2006)
13. Rousset, M.-C.: Small Can Be Beautiful in the Semantic Web. In: *3rd Int'l. Semantic Web Conf.* (November 2004)
14. Simperl, E., Krummenacher, R., Nixon, L.: A Coordination Model for Triplespace Computing. In: *9th Int'l. Conference on Coordination Models and Languages* (June 2007)
15. Strang, T., Linnhoff-Popien, C., Frank, K.: CoOL: A Context Ontology Language to Enable Contextual Interoperability. In: *4th Int'l. Conf. on Distributed Applications and Interoperable Systems* (November 2003)