

Semantic Grid Service Discovery using DHTs

Manolis Koubarakis, Zoi Kaoudi and Iris Miliaraki

Intelligent Systems Laboratory
Dept. of Electronic and Computer Engineering
Technical University of Crete
73100 Chania, Crete, Greece
{manolis,zoi,iris}@intelligence.tuc.gr
<http://www.intelligence.tuc.gr/~{manolis,zoi,iris}>

OntoGrid¹ is a project funded by the Grid Technologies unit of the Information Society Technologies programme under the strategic objective “Grid-based systems for Complex Problem Solving”. The main objective of OntoGrid is to work towards the realization of a Semantic Grid.²

In the context of OntoGrid, the work of the Technical University of Crete concentrates on the implementation of Atlas, a P2P system for the distributed storage and querying of RDF(S) metadata describing OntoGrid resources e.g., services. We use state of the art DHT technology [1] to design and implement a distributed system that will be able to scale to hundreds of thousands of nodes and to large amounts of RDF(S) data and queries.

Nodes in an Atlas network are organized under the Bamboo DHT protocol [5]. Nodes can enter RDF(S) data into the network and pose RQL queries. Two kinds of querying functionality are supported by Atlas: *one-time* querying and *publish/subscribe*. Each time a node poses a one-time query, the network nodes cooperate to find RDF(S) data that form the answer to the query. In the publish/subscribe scenario, a node can *subscribe* with a *continuous* query. A continuous query is indexed somewhere in the network and each time matching RDF(S) data is published, nodes cooperate to *notify* the subscriber. A high level view of the Atlas architecture is shown in Figure 1.

A simple scenario in which Atlas can be used for service discovery in OntoGrid is as follows. When an OntoGrid node has a service that it wishes to make available to the rest of the network, it creates and publishes RDF metadata about this service using a service description language e.g., OWL-S [4]. Later on, this service can possibly be discovered by another OntoGrid node that poses an RQL query to discover services that it needs.

The current implementation of Atlas supports only conjunctive queries expressed in a logic-based query language based on *triple patterns* originally used in [2]. The query processing algorithm we use for one-time queries is also based on the algorithms proposed in [2]. Publish/subscribe scenarios are handled using the algorithms in [3].

Atlas is currently under development at the Technical University of Crete. More information is available on request by the authors of this extended abstract.

¹ <http://www.ontogrid.net>

² <http://www.semantic-grid.org>

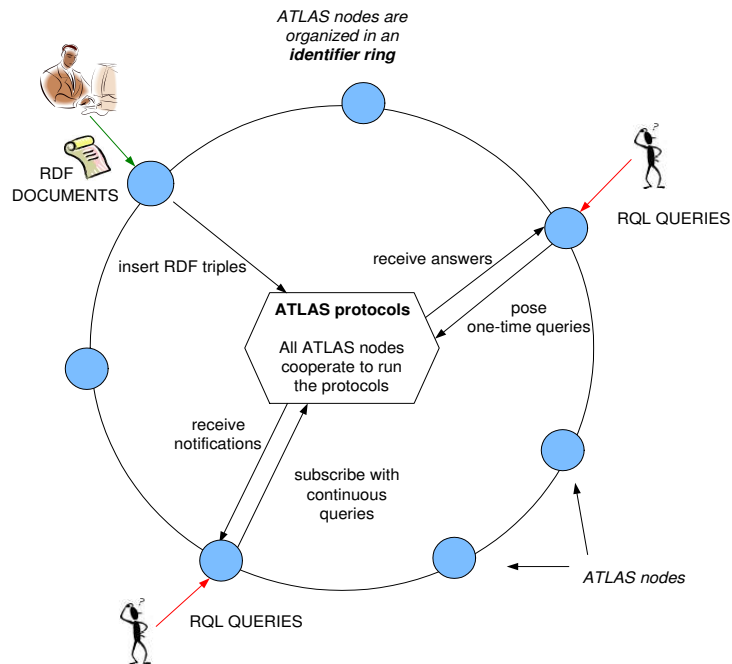


Fig. 1. Atlas architecture

References

1. Hari Balakrishnan, M. Frans Kaashoek, David R. Karger, Robert Morris, and Ion Stoica. Looking up data in P2P systems. *Communications of the ACM*, 46(2):43–48, 2003.
2. Min Cai, Martin R. Frank, Baoshi Yan, and Robert M. MacGregor. A Subscribable Peer-to-Peer RDF Repository for Distributed Metadata Management. *Journal of Web Semantics: Science, Services and Agents on the World Wide Web*, 2(2):109–130, December 2004.
3. Erietta Liarou, Stratos Idreos, and Manolis Koubarakis. Publish-Subscribe with RDF Data over Large Structured Overlay Networks. In *Databases, Information Systems and Peer-to-Peer Computing (DBISP2P 2005) to be held at VLDB 2005 31st International Conference on Very Large Data Bases*, Trondheim, Norway, 28-29 August.
4. The OWL Services Coalition. *OWL-S: Semantic Markup for Web Services*. Available at <http://www.daml.org/services/owl-s/1.0/>.
5. Sean Rhea, Dennis Geels, Timothy Roscoe, and John Kubiatowicz. Handling Churn in a DHT. In *USENIX Annual Technical Conference*, 2004.