Data Ring: Let us turn the net into a database!

Serge Abiteboul
INRIA-Futurs & LRI-Univ. Paris 11
firstname.lastname@inria.fr

Neoklis Polyzotis
UCSC
alkis@cs.ucsc.edu

June 19, 2006

Abstract

Because of information ubiquity, one observes an important trend towards transferring information management tasks from database systems to networks. We introduce the notion of Data Ring that can be seen as a network version of a database or a content warehouse. A main goal is to achieve better performance for content management without requiring the acquisition of explicit control over information resources. We discuss the main traits of Data Rings and argue that Active XML provides an appropriate basis for such systems.

The collaborating peers that form the Data Ring are autonomous, heterogeneous and their capabilities may greatly vary, e.g., from a sensor to a large database. To support effectively this paradigm of loose integration, the Data Ring enforces a seamless transition between data and metadata and between explicit and intentional data. It does not distinguish between data provided by web pages and data provided by web services, between local (extensional) data and external data obtained via a Web service call. This is achieved using the Active XML technology that is based on exchanging XML documents with embedded service calls both for the logical and physical data model.

A Brief Introduction to Active XML. An XML document is an unbounded, labeled, ordered tree, and web services are protocols for distributed computation. Active XML (AXML for short) extends the XML model by allowing documents to contain embedded calls to web services. An XML document thus contains both extensional information, i.e., the specific XML structure, as well as intensional information from the results of service calls. AXML documents may be viewed as objects that can be exchanged between peers. Persistency is supported in the spirit of object databases [6]. Furthermore, an AXML document is at the same time a client for web services and possibly a service provider. The service calls embedded in the document provide both intensional (in the sense of deductive databases [3]) and active (in the sense of active databases [10]) data. A survey of the AXML project may be found in [5]. Papers on AXML as well as the open-source code of an AXML peer may be found from [1].

Data Ring. Imagine a system of collaborating peers, where each peer exposes a set of resources according to its capabilities. For instance, some peers may expose data sources and/or services, while others may offer computational resources. A *Data Ring* is a middleware system that enables the integrated management of the information that is present in the P2P system, with performance, reliability, robustness and more generally, quality of service tailored to different applications. By information, we mean first data, traditional as in relational databases, and less so as in content management systems (mails, letters, reports, etc.). By information, we also mean metadata about the data as well as knowledge (e.g., ontologies) that are used to interpret data and metadata. Last but not least, we mean web services providing access to data sources. All this information is published through the Data Ring and accessed by a range of mechanisms that are tailored to the needs of different applications.

The collaborating peers that form the Data Ring are autonomous, heterogeneous and their capabilities may greatly vary, e.g., from a sensor to a large database. The goal of the Data Ring, however, is to hide this inherent diversity and export a homogeneous and unified layer for data management. Some of the mechanisms supplied by a Data Ring are typical DBMS functionalities, e.g.: persistence, query/update, concurrency control, fault tolerance, recovery, access control. Others are more specific, e.g.,: information discovery and enrichment.

A Data Ring is primarily a data integration system. It is a middleware layer by which applications can access this diverse information efficiently and in a unified manner. In order to support effectively this paradigm of loose integration, we insist on two aspects regarding the logical model of a Data Ring:

- Seamless transition between data and metadata. The use of metadata is essential for linking data resources and in particular classifying them.
- Seamless transition between explicit and intentional data. One should not have to distinguish between data provided by web pages and data provided by web services, between local (extensional) data and external data obtained via a Web service call.

Foundations of Data Ring. A fundamental property of our proposed architecture is the use of AXML for the logical and physical model of the Data Ring. At the logical level, AXML is well suited for representing the "intentional" and "active" information that is managed by the Data Ring. Essentially, the intensional component enables the sharing of information and knowledge in a distributed context, while the active component provides support for aspects such as pub/sub, monitoring, synchronization/reconciliation, network connectivity, awareness. Besides the logical level, our thesis is that AXML can also serve as the basis for the physical model. Consider for instance the key issue of distributed query optimization, and in particular the description and exchange of distributed query execution plans. As recent work [4] has shown, this can be achieved using AXML expressions describing standard algebraic XML operations and send/receive operators over XML streams.

Another equally crucial aspect is the use of self-activation (or, self-tuning) in all layers of the Data Ring architecture. The P2P architecture of the Data Ring makes self-tuning a requirement, since the manual tuning of the system by an administrator is impractical, in many applications impossible. Note that in this new setting, self-tuning is a distributed process and thus, has to be entirely reconsidered. The importance of self-tuning is accentuated by the lack of direct control over the data sources exported by the peers. In a nutshell, the efficiency and reliability of accessing a particular data source is initially constrained by its physical organization in its owning peer. Since that organization cannot be changed, the Data Ring may decide to replicate, reorganize, and re-index the data source (or parts of it) in different peers, in order to enhance its accessibility. Such decisions will be driven of course by the characteristics of the query load, and by its evolution. Thus, system statistics play an important role in this physical organization both at the peer level and in the Data Ring considered as a whole.

Related Work and Conclusion. Many ideas found here are influenced by previous works on P2P content warehouse, see [2]. Similar ideas have been promoted in [8]. The authors of the present paper have been strongly influenced by on-going works in the Gemo team (notably with Ioana Manolescu and Tova Milo) and at Santa Cruz.

Some of the underlying technology has already been developed in related software systems, e.g.: structured p2p network such as Chord or Pastry, XML repositories such as Xyleme or DBMonet, file sharing systems such as BitTorrent or Kazaa, distributed storage systems such as OceanStore or Google File System, content delivery network such as Coral or Akamai, multicast systems such as Bullet or Avalanche, Pub/Sub system such as Scribe or Hyper, application platform suites as proposed by Sun or Oracle for integrating software components, data integration as provided in warehouse or mediator systems. Of course, all the work in distributed database systems [9] is relevant.

Finally, the development of data rings leads to a number of research issues, notably the distribution of certain functionalities such as query optimization (with self tuning), or access control. An important milestone towards the realization of this vision is the completion of "the glue" between the components and in particular:

- The data/knowledge model that is used.
- The API for the various functionalities.
- The format for publishing logical resources as well as physical ones (AXML is just a starting point).

References

[1] Active XML web site, activexml.net

[2] S. Abiteboul: Managing an XML Warehouse in a P2P Context. CAiSE, 4-13, 2003

[3] S. Abiteboul, R. Hull and V. Vianu: Foundations of Databases, Addison-Wesley, Reading-Massachusetts, 1995

[4] S. Abiteboul, I. Manolescu, E. Taropa: A Framework for Distributed XML Data Management. EDBT, 1049-1058, 2006

[5] S. Abiteboul, T. Milo, O. Benjelloun: The Active XML Project: an overview, INRIA Internal Report, ftp://ftp.inria.fr/INRIA/Projects/gemo/gemo/GemoReport-331.pdf

[6] Edited by R. Cattell: The Object Database Standard: ODMG-93, Morgan Kaufmann, San Mateo, California, 1994

[7] The Edos Project, http://www.edos-project.org/

[8] A. Halevy, M. Franklin, D. Maier: Dataspaces: A New Abstraction for Information Management. DASFAA, 1-2, 2006

[9] T. Özsu, P. Valduriez: Principles of Distributed Database Systems. 2nd Edition, Prentice Hall, Englewood Cliffs, New Jersey, 666 pages, 1999

[10] N. Paton, O. Diaz, Active database systems, ACM Computing Surveys, Vol. 31, No. 1, March 1999

[11] W3C - The World Wide Web Consortium, http://www.w3.org/