

The MYRIAD Federated Database Prototype

S.-Y. Hwang, E.-P. Lim, H.-R. Yang, S. Musukula
K. Mediratta, M. Ganesh, D. Clements*,
J. Stenoien, J. Srivastava,

Department of Computer Science, University of Minnesota, Minneapolis, MN 55455
{srivasta|ganesh|syhwang|peng}@cs.umn.edu

1 Introduction

Myriad is a federated database system (FDBS) prototype being developed at the University of Minnesota. The main objective behind this prototyping effort is to provide "enterprise-wide" information by integrating independently developed databases while preserving the local autonomy of the component DBMSs and applications. In Myriad, multiple federations can be formed. A federation consists of an integrated database whose schema is represented as a set of integrated relations derived from the export relations provided by the component DBMSs. SQL, mainly due to its simplicity and popularity among database users and vendors, has been adopted to express global queries as well as the queries for the local database gateways. Myriad provides query processing and global transaction management in a federated environment [Myriad93].

2 Current Implementation Status

An implementation of Myriad has been realized in the UNIX environment on a network of Sun SPARCstations. Currently, we have built gateways on two local DBMSs, namely Oracle and Postgres. To demonstrate the essential features of database integration in Myriad, we constructed several example databases on both Oracle and Postgres such that relations from these databases are merged into integrated relations using relational operations as well as user-defined integration functions. All programs are written in C and embedded query languages. The Myriad communication process is implemented using the BSD socket libraries.

The essential functionalities of all query processing components have been implemented. A simple query

optimization strategy has been adopted to evaluate global SQL queries. A full-fledged query optimization is currently being developed and implemented.

We have implemented a Myriad transaction management subsystem that supports the general transaction model. Presently, each integrated local DBMS employs two-phase locking (2PL). The transaction management components support two-phase commit over local DBMSs so as to achieve serializable execution. With respect to the global deadlock problem, Myriad uses a timeout mechanism to resolve it. A timeout period is associated with each local query submitted to the gateway. If the result of a local query does not return within the timeout period, the entire global transaction is assumed to be involved in a global deadlock and is aborted.

At the application tool level, an easy-to-use query interface has been implemented. This allows federation users and DBAs to browse/modify/create federated schemas and pose transaction as well as query requests to the Myriad system.

3 Future Plans

Myriad serves as a testbed for validating and comparing solutions to various FDBS problems such as query optimization and transaction management. In the future, we will adapt Myriad to object-oriented paradigm in which local databases and applications are treated as objects. We will examine the possibilities of constructing a workflow model on top of Myriad. We also seek to build real applications on Myriad.

References

- [Myriad93] D. Clements, M. Ganesh, S.-Y. Hwang, E.-P. Lim, K. Mediratta, J. Srivastav, J. Stenoien, and H.-R. Yang, 'Myriad: Design and Implementation of a Federated Database Prototype,' Technical Report 93-76, Dept. Computer Sci., U. of Minnesota, MN, 1993.

*The address of this author is Data Parallel Systems, Inc., Bloomington, IN 47408.