Loop-Lifted XQuery RPC With Deterministic Updates

Ying Zhang* Peter Boncz†

Centrum voor Wiskunde en Informatica‡
P.O.Box 94079, 1090 GB
Amsterdam, the Netherlands

Abstract

XRPC is a minimal XQuery extension that enables distributed query execution, combining the Remote Procedure Call paradigm with the existing concept of XQuery functions. By calling out of a `for`-loop to multiple destinations, and by calling functions that themselves perform RPC calls, complex P2P communication patterns can be achieved with XRPC. We further propose the use of SOAP as the protocol for XRPC, which allows seamless integration with web services and Service Oriented Architectures (SOA).

XRPC is implemented in the open source MonetDB/XQuery system. We show that the technique of *loop-lifting*, that executes all expressions inside a `for`-loop in a single bulk operator – pervasively applied in MonetDB/XQuery to obtain efficient relational query plans – nicely extends to RPC. Loop-lifting enables us to send *bulk* RPC requests, dramatically reducing the number of SOAP messages, and thus the performance impact of network latency.

The XRPC extension is orthogonal to all XQuery language features, including the XQuery Update Facility (XUF). The XUF W3C Draft proposal does not define the order in which multiple update actions to the same node must be applied. We instead choose to make this order deterministic, and show how distributed updates can be made deterministic using a small protocol extension.

*e-mail: Y.Zhang at cwi.nl
†e-mail: P.Boncz at cwi.nl
‡Center of Mathmatics and Computer Science