論文の内容の要旨

論文題目　Replication Middleware for High Performance Database Systems
（高性能データベースシステムのためのレプリケーションミドルウェア）

氏名　三島健

Recently, several middleware-based approaches have been proposed. If we implement all functionalities of database replication only in a middleware layer, we can avoid the high cost of modifying existing database servers or scratch-building. However, it is a big challenge to propose middleware which can enhance performance and scalability without modification of database servers because the restriction may cause extra overhead. Unfortunately, many existing middleware-based approaches suffer from several shortcomings, i.e., some cause a hidden deadlock, some provide only table-level locking, some rely on total order communication tools, and others need to modify existing database servers.

In this paper, we propose a new correctness criterion for eager replicated database systems called global snapshot isolation (GSI). GSI not only guarantees snapshot isolation to clients but also maintains consistency between database servers. We also propose two middlewares called DV-SI and Pangea to meet GSI. DV-SI provides table-level concurrency control and higher throughput by exploiting the property of snapshot isolation. Pangea provides tuple-level concurrency control and much higher throughput by exploiting the property of the first-updater-wins rule. We have implemented the prototypes of DV-SI and Pangea on top of PostgreSQL servers without modification. Because GSI also presents the guide of minimal implementation, DV-SI and Pangea use less than 2000 lines of C code.

Our experimental results with the TPC-W benchmark reveal that, compared to existing middlewares without modification of database servers, DV-SI and Pangea provide better performance in terms of throughput and scalability. More surprisingly, Pangea outperformed any other middleware in any workload.