DE LA SALLE UNIVERSITY

A RELIABLE AND CONSISTENT FULLY DISTRIBUTED DATABASE

A Thesis Presented to the Faculty of the
Computer Science Graduate Program
College of Computer Studies
De La Salle University

In Partial Fulfillment
of the Requirements for the Degree
Master of Science in Computer Science

by

Patricia Maria R. Claudio

August 1989

DE LA SALLE UNIVERSITY

A B S T R A C T

A number of studies on algorithms for concurrency control and recovery mechanisms have been conducted during the past decade. Performance evaluations were made on these algorithms but it has been noted that an integrated approach for concurrency control and recovery for distributed database systems (DDBS) have not been given much significance. Moreover, most of these studies have treated concurrency control and recovery as two separate problems regardless of the fact that they are closely related.

This paper presents an integrated concurrency and recovery algorithm, strict timestamp ordering was used for concurrency control and for recovery, incremental log with immediate updates. Strict timestamp ordering delays each write operation on a data item until there is no active transaction that issued a write on that data item.

Reliability is achieved in the DDBS through replication and and the recovery technique adopted. The algorithm allows a transaction to access any of the available copies of a data item. Since any copy of the data item can be accessed, data are available even in cases of site failures. The recovery technique ensures that each transaction would eventually terminate either by commiting or aborting. Moreover, it also takes care of restoring the database to a consistent state during recovery of a failed site.

Consistency is achieved by controlling the accesses made to the database. This is solved by the concurrency control mechanism. Transaction restarts were minimized since strict timestamp ordering was used. The synchronization of the locals clocks was also included in the algorithm.

The correctness of the algorithm was also proven and the message complexity of the algorithm was found to be of $O(\alpha n)$, where α is a function of size of the queue used in the algorithm.

A closed queueing network model was developed to evaluate the performance of the integrated algorithm. The performance measurements taken include response time, throughput and system utilization. The results indicate that the performance exhibited by the system is good even when subjected to heavy load.

Index Terms : Integrated concurrency control and recovery algorithm, distributed database systems, strict timestamp ordering, incremental log with immediate updates

Table of Contents