ABSTRACT

Fault Tolerant Optimistic Concurrency Control In A Distributed System With Validation Through Pessimism

Michael Anthony Hosein

Concurrency is the facility to accommodate many of the same type of activity at the same time. For example, several banking type transactions may want to access one account all at once. Concurrency control therefore refers to the activity of coordinating the actions of processes that operate in a parallel manner, access common data and which may possibly conflict with one another. Fault Tolerance is the activity of ensuring that software and hardware failures do not corrupt persistent data. The goal of concurrency control and fault tolerance is to ensure that transactions execute atomically that is:

1. each transaction accesses shared data without interfering with other transactions and
2. if a transaction terminates normally, then all of its effects are made permanent; otherwise it has no effect at all.

One of the major issues in distributed computing today is availability. Availability refers to the percentage of transactions submitted to a system that complete. The aim of this system is to increase availability by implementing

techniques which increase concurrency. Of utmost importance is the fact that integrity of the database must always be preserved. Recovery mechanisms must therefore be in place to ensure that the database contains the effects of committed transactions and no effects of aborted ones.

Distributed concurrency control algorithms can be broadly classified into two groups : optimistic and pessimistic schemes. Optimistic schemes allow transactions to execute and then check for database integrity. Pessimistic schemes check for integrity of the database before determining if to allow a transaction to proceed. This work is a development of that done by previous researchers in which optimistic and pessimistic schemes were combined to increase availability in a one object distributed transaction processing system. It must always be remembered that nodes may fail, communication links may be broken and so on. The aim of the proposed system is to make a more robust system model that would withstand failures in a real world situation.

Optimistic processing is controlled by a value called a cost bound. A transaction will be allowed to proceed if the change that it is about to make is less than or equal to the cost bound. Where an optimistic run has occurred, pessimistic processing must take place later (permanent updates), to ensure that database integrity is preserved.

Data is also fully replicated in order to increase availability.

Keywords: distributed system; concurrency control; fault tolerance.