

SYNCHRONIZATION AND RECOVERY IN DATABASE SYSTEMS

Rudolf Bayer

Institut für Informatik
Technische Universität München
D-8000 München 2, Germany

Abstract*

Transactions are processes, which must have an atomic effect - from the user's point of view - on the state of a database. To provide this property, a database system must be prepared to undo any effects of a transaction on the state of the database, if for some reason the transaction cannot be completed successfully. Copies of data objects - shadows - are kept to be able to recover the state of the database before the transaction was started.

Single transactions are assumed to be correct, i.e. to transform a correct database state into a correct database state, if they run by themselves without any disturbances by other transactions. For several reasons, however, a viable database system must be able to perform transactions in parallel, still guaranteeing the same net effect on the state of the database as if the transactions were run in some serial order one after the other. This is achieved by proper consistent synchronization, which should be supplied automatically by the database system. An important class of synchronization protocols uses locks of various types. Locks, however, tend to decrease the parallelism between transactions.

Recovery techniques based on shadows and synchronization based on locking can be combined to substantially increase the potential parallelism between transactions. Several special types of locks are introduced and the corresponding synchronization protocols are developed leading to a unified approach to synchronization and recovery.

One of these protocols has the property that transactions, which only read the database, never have to wait for other transactions. This property allows new longterm recovery techniques

*Most of the material of this talk is contained in the paper:

R. Bayer, H. Heller, A. Reiser: Parallelism and recovery in database systems. Technische Universität München, Report TUM-INFO-7832, November 1978. To be published in the ACM Transactions on Database Systems.

requiring a complete copy of the state of the database for archival purposes.

The generalization of these techniques to distributed database systems will be discussed.