

Time-stamp protocol – ensures that conflicting read and write operations occur in timestamp order.

If a read (Q) is issued by T_i :

- 1) If $TS(T_i) < W\text{-timestamp}(Q)$ then T_i needs to read a value of Q that was already overwritten. The read operation is rejected and T_i is rolled back.
- 2) If $TS(T_i) \geq W\text{-timestamp}(Q)$ then the read operation is executed and $R\text{-timestamp}(Q)$ is set to the maximum of $R\text{-timestamp}(Q)$ and $TS(T_i)$.

If a write(Q) is issued:

- 1) If $TS(T_i) < R\text{-timestamp}(Q)$ then the value of Q that T_i is producing was needed previously, and the system assumed that the value would never be produced. The write is rejected and T_i is rolled back.
- 2) If $TS(T_i) < W\text{-timestamp}(Q)$ then T_i is attempting to write an obsolete value of Q. The write is rejected and T_i is rolled back.
- 3) Otherwise the write is accepted and $W\text{-timestamp}(Q)$ is set to $TS(T_i)$.

Time-stamp protocol ensures conflict serializability – because conflicting operations are processed in timestamp order. It also ensures freedom from deadlock since no transaction ever waits. There is the possibility of starvation if a sequence of short transactions continue to cause repeated restarting of a long transaction.

The protocol can generate schedules that are not recoverable. It can be extended to include recoverability by:

- 1) (Recoverable and cascadeless) Forcing all the writes together at the end of the transaction. The writes are atomic in that while the writes are in progress, no transaction is permitted to access any of the data items being written.
- 2) (Recoverable and cascadeless) Using a limited form of locking where the reads of uncommitted items are postponed until the transaction that updated the item commits.
- 3) (Recoverable) Tracking uncommitted writes and allowing T_i to commit only after the commit of any transaction that write a value the T_i read.

Thomas write rule:

This involves the case that obsolete writes can be ignored under certain circumstances. Read rules remain unchanged, but writes are slightly different.

If a write(Q) is issued:

- 1) If $TS(T_i) < R\text{-timestamp}(Q)$ then the value of Q that T_i is producing was needed previously, and the system assumed that the value would never be produced. The write is rejected and T_i is rolled back.
- 2) If $TS(T_i) < W\text{-timestamp}(Q)$ then T_i is attempting to write an obsolete value of Q. The write is ignored.
- 3) Otherwise the write is accepted and $W\text{-timestamp}(Q)$ is set to $TS(T_i)$.

The Thomas write rule makes use of view serializability by deleting obsolete write operations from the transactions that issue them.

Multiversion schemes

Multiversion schemes keep old versions of a data item to increase concurrency.

Multiversion timestamp ordering

Each successful write results in the creation of a new version of the data item written
Use timestamps to label versions

When a $R(Q)$ operation is issued, select the appropriate version of Q based on the timestamp of the transaction and return that value. Reads never have to wait as the appropriate version is available immediately.

Each data item Q has a sequence of versions $\langle Q_1, Q_2, Q_3, \dots, Q_n \rangle$

Each version contains the value, $W\text{-timestamp}(Q_k) - TS$ of the transaction that wrote Q_k , and $R\text{-timestamp}(Q_k) - \text{largest } TS$ of a transaction to successfully read version Q_k .

When T_i creates a new version Q_k of Q , Q_k 's $W\&R$ -timestamps are initialized to $TS(T_i)$.

R -timestamp of Q_k is updated whenever a transaction T_j reads Q_k and $TS(T_j) > R\text{-timestamp}(Q_k)$.

Suppose T_i issues a $R(Q)$ or a $W(Q)$ operation. Let Q_k denote the version of Q whose W -timestamp is the largest W -timestamp $\leq TS(T_i)$.

- 1) If T_i issues a $R(Q)$, then the value returned is the content of of version Q_k . Reads always succeed.
- 2) If T_i issues a $W(Q)$
 - a. If $TS(T_i) < R\text{-timestamp}(Q_k)$, then transaction T_i is rolled back. Some other transaction T_j (as serialization is defined by the timestamp values) should read T_j 's write and has already read a version created by a transaction older than T_i .
 - b. If $TS(T_i) = W\text{-timestamp}(Q_k)$, the contents of Q_k are overwritten; Q_k was written before by T_i .
 - c. If $(TS(T_i) > W\text{-timestamp}(Q_k))$ a new version of Q is created.
 - d. If $(TS(T_i) < W\text{-timestamp}(Q_k))$, write is ignored (Thomas write rule)

Validation protocols

When the majority of the transactions are read-only, rate of conflict is low. In this situation, you may not want to incur the cost of a concurrency control scheme. Instead, we look at a scheme that does monitoring rather than locking.

Each transaction is assumed to execute in three phases whether it is read-only or an update transaction. The phases are:

- 1) Read phase – all values to be read are read into variable local to T_i . All writes are performed on these local variables rather than the actual database.
- 2) Validation phase – T_i performs a validation test to determine whether it can copy to the database the results of the writes without causing a violation of serializability.
- 3) If the validation succeeds, the system applies the changes – otherwise the transaction is rolled back.

Each transaction must go through the three phases, but the three phases can be interleaved by concurrently executing transactions. To perform the validation phase we need a couple of pieces of data. Three timestamps are associated with the transaction.

- 1) Start (T_i) – the time when T_i started its execution.
- 2) Validation (T_i) – the time when T_i finished its read phase and started its validation phase.
- 3) Finish (T_i) – the time when T_i finished its write phase.

The serializability order is determined by the timestamp Validation – i.e. $TS(T_i) = \text{Validation}(T_i)$. The validation test for T_j is as follows: For all transactions T_i where $TS(T_i) < TS(T_j)$ one of the two following must hold.

- 1) $\text{Finish}(T_i) < \text{Start}(T_j)$. Since T_i completes before T_j starts, the serializability order is maintained.
- 2) The set of data items written by T_i does not intersect with the set of data items read by T_j ; and T_i completes its write phase before T_j starts its validation phase. ($\text{Start}(T_j) < \text{Finish}(T_i) < \text{Validation}(T_j)$). This condition ensures the writes of T_i and T_j do not overlap. Since the writes of T_i do not affect the read of T_j and since T_j cannot affect the read of T_i – the serializability order is maintained.

This is an example of an optimistic concurrency control.