

15-721 DB Sys. Design & Impl.

Optimistic CC

Christos Faloutsos

www.cs.cmu.edu/~christos

Roadmap

- 1) Roots: System R and Ingres
- 2) Implementation: buffering, indexing, q-opt
- 3) Transactions: locking, recovery
 - granularity of locks
 - ➔ **optimistic CC**
 - B-trees
 - ...
- 4) Distributed DBMSs
- 5) Parallel DBMSs: Gamma, Alphasort
- ...

Optimistic CC (Kung&Robinson)

- Assumption: conflicts are rare
- Optimize for the no-conflict case.
- All transactions consist of three phases
 - **Read:** Here, all writes are to private storage.
 - **Validation:** Make sure no conflicts have occurred.
 - **Write:** If Validation was successful, make writes public. (If not, abort!)



Why Might this Make Sense?

Why Might this Make Sense?

- All transactions are readers
- Lots of transactions, each accessing/modifying only a small amount of data, large total amount of data
 - Low probability of conflict, so again locking is wasted
- Fraction of transaction execution in which conflicts “really take place” is small compared to total path length
 - Locks until end of Xact are way too restrictive most of the time

Validation Phase (1)

- Goal: guarantee only serializable schedules
- intuitively: at validation, T_j checks it's ‘elders’ for RW and WW conflicts - specifically:
- Technique:
 - Assign each transaction a TN (transaction number)
 - Require TN order to be the serialization order
 - If TN(T_i) < TN(T_j) ⇒ **ONE** of the following must hold:

CMU SCS

Observations

- When to better assign TN's?
- at beginning of read phase: Tj has to wait...

Tj has to wait
for W(Ti)

15-721 C. Faloutsos 13

CMU SCS

Observations

- When to better assign TN's?
- at beginning of **validation** phase:
 - Tj can start
 - condition (3): automatic!

Tj has to wait
for W(Ti)

15-721 C. Faloutsos 14

CMU SCS

Observations (cont'd)

- BUT: subtle problem: T with very long R!
 - must check ALL T's within its lifetime!!
 - Requires unbounded buffer space. Solution?
 - Bound buffer, toss out when full, abort possibly affected Ts
 - Starvation!
- Serial/Parallel validation – Pros & cons?

15-721 C. Faloutsos 15

CMU SCS

A Serial Validation Technique

Goal: to ensure conditions 1 and/or 2 above.

Requires that write phases be done serially.

15-721 C. Faloutsos 16

CMU SCS

Serial Validation Algorithm

- Record *start_tn* when Xact starts (to identify active Xacts later)
- Obtain the Xact's real Transaction Number (TN) at the start of validation phase
- Record read set and write set while running and write into local copy
- Do validation and write phase inside a critical section

15-721 C. Faloutsos 17

CMU SCS

Serial Validation: Critical Section

```

beginCriticalSection
  finish_tn := currentTN; /* tentatively assign tn */
  valid := true;
  for T from start_tn + 1 to finish_tn do
    if (write set of Xact T intersects read set)
      then valid := false;
  if valid
    then { write phase; currentTN++; tn := currentTN }
  endCriticalSection
  if valid then cleanup() else backup();
  
```

15-721 C. Faloutsos 18

CMU SCS

Serial Validation (cont.)

Optimization: Do not assign TN (TID) unless success!
 Informally,

1. check current TN;
2. check everything from start until current TN;
3. then enter critical region and do the rest.

Read-only Xacts are not assigned TNs; just check write sets of Xacts with $start_tn < TN < finish_tn$

15-721 C. Faloutsos 19

CMU SCS

A Serial Validation Technique

Optimization: move some of the validation outside the critical section.

15-721 C. Faloutsos 20

CMU SCS

A Serial Validation Technique

This can be repeated for 2nd, 3rd etc time!

15-721 C. Faloutsos 21

CMU SCS

Parallel Validation

Only real difference:
 now must check condition 3, using *active*, the set of Xacts that have finished their read phase but have not yet completed their write phases.

Algorithm: in paper
 Subtlety: An aborting Xact may cause another Xact to abort!

15-721 C. Faloutsos 22

CMU SCS

Opt CC vs. Locking

| | |
|---|--|
| <p>Locking:</p> <ul style="list-style-type: none"> • order is of first lock; • wait • on deadlock, abort | <p>Optimistic cc</p> <ul style="list-style-type: none"> • order is of $t(i)$ • abort • on starvation, lock |
|---|--|

15-721 C. Faloutsos 23

CMU SCS

Conclusions

- Analysis [Agrawal, Carey, Livny, '87]:
 - dynamic locking performs very well, in most cases
- All vendors use locking
- optimistic cc: promising for OO systems, or when resource utilization is low.

15-721 C. Faloutsos 24