RETHINKING SERIALIZABLE MULTIVERSION CONCURRENCY CONTROL

Jose M. Falerio & Daniel J. Abadi Yale University

Presenter: Suyash Gupta, Purdue University

January 24, 2017

Suyash Gupta

CS59000BDS

January 24, 2017 1 / 31

Initial Thoughts

FOOD FOR THOUGHT



Can we design a concurrency control protocol for multi-version database systems that is scalable?



(日)

< □ > < 同 >

Keywords

Lets Refresh !!!

- Serializability
- Multiversion
- Timestamps
- Anti-dependency



< □ > < 同 > <



< ∃⇒

Basics I

BASICS I: SINGLE VS. MULTI-VERSION SYSTEMS



Suyash Gupta

January 24, 2017 4 / 31

Image: A marked black

ા પશ્ચ સં

Basics

BASICS I: SINGLE VS. MULTI-VERSION SYSTEMS



Basics I

BASICS I: SINGLE VS. MULTI-VERSION SYSTEMS

Multi-versioning buys more concurrency! Right?



Suyash Gupta

CS59000BDS

January 24, 2017 6 / 31

<ロト < 同ト < ヨト < ヨト

BASICS II: MULTI-VERSIONING EXAMPLE



BASICS II: MULTI-VERSIONING EXAMPLE



<ロト < 同ト < ヨト < ヨト

HIVERSI H

BASICS II: MULTI-VERSIONING EXAMPLE

Use of conservative isolation techniques, similar to single-version systems can avoid such inconsistencies.



Image: A matrix and a matrix

BASICS III: MULTI-VERSIONING EXAMPLE II



BASICS III: MULTI-VERSIONING EXAMPLE II





PURDUE

イロト イボト イヨト イヨト

BASICS III: MULTI-VERSIONING EXAMPLE II



BASICS III: MULTI-VERSIONING EXAMPLE II



Savings: 25		Savings: 75	
begin: 0	end: 10	begin: 10	end: inf

PURDUE

<ロト < 同ト < ヨト < ヨト

BASICS III: MULTI-VERSIONING EXAMPLE II



Lets Motivation

MOTIVATION I – NON-SCALABILITY



Lets Motivation

MOTIVATION II – MANAGEMENT ISSUES



Lets Motivation

MOTIVATION III – GUARANTEEING SERIALIZABILITY



Figure 1: Non-serializable interleaving, and corresponding serialization graph of T_r and T_w . $r[x_1]$ denotes to a *read* of version 1 of record x, correspondingly, $w[x_1]$ denotes a *write* to record x, which produces version 1. A record's subscript corresponds to the version read or written by the transaction.

- Separate concurrency control from transaction execution.
- Concurrency control determines transaction order and version visibility.
- Execution performs logic given concurrency control ordering



< ∃⇒

System at a Glance

- Transactions handed over to a single thread that maintains a log.
- Position of a transaction in the log is its timestamp.
- *m* concurrency threads, each owning a logical partition, analyze each transaction in the log, and create a space "placeholder" if the transaction writes to any record in its partition.
- *n* separate execution threads, each execute a batch of transactions, and fill the pre-allocated spaces (write operation).

Representation

BOHM DIAGRAMMATICALLY



January 24, 2017 20 / 31

▶ < ∃ ▶</p>

<

PURDUE

TIMESTAMP ASSIGNMENT

- Prior Multi-version CC systems assign each transaction, two timestamps, *t_{begin}* and *t_{end}*.
 - *t_{begin}* determines which versions of pre-existing records are visible.
 - *t_{end}* determines the time at which writes become visible to other transactions.
- BOHM assigns each transaction a single timestamp, t_s.
- Each transaction appears to execute atomically at time ts.

<ロト < 同ト < ヨト < ヨト

INSERTING PLACEHOLDERS

- Several threads contribute to the processing of a single transaction's write-set.
- BOHM assigns the responsibility of each record to one concurrency control thread.
- When concurrency control layer receives a transaction, every concurrency control thread examines its write-set to determine whether any records belong to their partitions.



VERSION FIELD'S

- Start timestamp set to the timestamp of the transaction that creates the version.
- End timestamp is set to infinity,
- Txn pointer is set to the transaction that creates the version.
- Data is left uninitialized.
- Prev pointer is set to the preceding version of the record.



BATCHING

- Once a transaction has been processed by all threads it can be handed off to the transaction execution layer.
- One way use synchronization barriers.
- BOHM avoids global coordination.
- Amortizes the cost of coordination across large batches of transactions.
- Each concurrency control thread receives an ordered batch of transactions

Concurrency Control

CONCURRENCY CONTROL LAYER

- Partition data across multiple threads
- For every write, create a new version









▶ ∢ ⊒ ▶

EVALUATING TRANSACTION LOGIC

- Computation of version data occurs in this phase.
- Read dependencies may cause a transaction to wait as data may not be available.
- Write dependencies allow parallel execution
- Read-Modify-Write may require similar wait as Read.
- Versions can be Garbage Collected, as transactions are ordered in batches.

EXECUTION LAYER

- Begins executing a batch after concurrency control completes
- Perform txn logic, write out data





a

▶ < ∃ ▶</p>

EXECUTION LAYER



FUTURE THOUGHTS

- BOHM's biggest disadvantage is its need to pre-determine the write-sets of the transaction, prior to its execution.
- Interesting thought can be to design an approach on similar lines for on-line or real-time systems, with obvious tradeoffs.
- Batching transactions entering at same instant.

Questions





Suyash Gupta

January 24, 2017 30 / 31

・ロト ・四ト ・ヨト ・ヨト





Suyash Gupta

January 24, 2017 31 / 31

◆□ ▶ ◆圖 ▶ ◆ 圖 ▶ ◆ 圖 ▶