Fast Abort-Freedom for Deterministic Transactions

Chen Chen¹, Xingbo Wu², Wenshao Zhong¹, Jakob Eriksson¹ ¹ University of Illinois at Chicago ² Microsoft Research

• Transaction: a logical unit of multiple data access operations

• Transaction: a logical unit of multiple data access operations



• Transaction: a logical unit of multiple data access operations



• Transaction: a logical unit of multiple data access operations



A lot of transactions concurrently







• Pessimistic, locking based: 2-phase locking (2PL)



- Pessimistic, locking based: 2-phase locking (2PL)
- Optimistic concurrency control (OCC)

C COMMIT A ABORT Time	T
1 lock(x) lock(y) EXECUTION C	
<pre>[lock(x) A lock(x) A lock(x) lock</pre>	ck(y) EXECUTION C (wait-die)
COPY EXECUTION VERIFY C	
COPY EXECUTION VERIFY A COPY E	EXECUTION VERIFY C

- Pessimistic, locking based: 2-phase locking (2PL)
- Optimistic concurrency control (OCC)

	C COMMIT A ABORT	
T1	lock(x)lock(y) EXECUTION C	
T2	<pre>lock(x) A lock(x) A lock(x) lock(y) EXECUTION C</pre>	(wait-die)
T1	COPY EXECUTION VERIFY C	
T2	COPY EXECUTION VERIFY A COPY EXECUTION VERIFY C	

- Abort is inevitable if:
 - The data access of a transaction is unknown

- Abort is inevitable if:
 - The data access of a transaction is unknown
- Deterministic transactions bring new opportunities
 - $\circ~$ w/ known read set and write set

- Abort is inevitable if:
 - The data access of a transaction is unknown
- Deterministic transactions bring new opportunities
 - w/ known read set and write set
- Exploit determinism:
 - Ordered locking



- Abort is inevitable if:
 - The data access of a transaction is unknown
- Deterministic transactions bring new opportunities
 - w/ known read set and write set
- Exploit determinism:
 - Ordered locking
 - Batched scheduling

- Abort is inevitable if:
 - The data access of a transaction is unknown
- Deterministic transactions bring new opportunities
 - w/ known read set and write set
- Exploit determinism:
 - Ordered locking
 - Batched scheduling





• Get the best of 2PL, scheduling and OCC

- Get the best of 2PL, scheduling and OCC
 - 2PL: locking
 - Scheduling: centralized scheduler
 - OCC: opportunistic commit



- Get the best of 2PL, scheduling and OCC
 - 2PL: locking queuing
 - Scheduling: centralized scheduler decentralized scheduling
 - OCC: opportunistic commit deterministic commit order

- Get the best of 2PL, scheduling and OCC
 - 2PL: locking queuing
 - Scheduling: centralized scheduler decentralized scheduling
 - OCC: opportunistic commit deterministic commit order
- Each transaction can observe their global order individually, with minimal cost.

- Get the best of 2PL, scheduling and OCC
 - 2PL: locking queuing
 - Scheduling: centralized scheduler decentralized scheduling
 - l scheduling
 - OCC: opportunistic commit deterministic commit order
- Each transaction can observe their global order individually, with minimal cost.
- DecentSched

- Three transactions accessing three variables
 - T1: x, z
 - $\circ~$ T2: x, y
 - \circ T3: y, z

- Three transactions accessing three variables
 - T1: x, z
 - $\circ~$ T2: x, y
 - $\circ~$ T3: y, z

Queuing



- Three transactions accessing three variables
 - T1: x, z
 - \circ T2: x, y
 - $\circ~$ T3: y, z

Queuing

Dependency tracking





- Three transactions accessing three variables
 - T1: x, z
 - T2: x, y
 - $\circ~$ T3: y, z

Queuing







- Three transactions accessing three variables
 - T1: x, z
 - T2: x, y
 - T3: y, z

Queuing



Dependency tracking



- Three transactions accessing three variables
 - T1: x, z
 - T2: x, y
 - \circ T3: y, z

Queuing







- Three transactions accessing three variables
 - T1: x, z
 - T2: x, y
 - T3: y, z

Queuing





Dependency tracking



- Three transactions accessing three variables
 - T1: x, z
 - T2: x, y
 - T3: y, z

Queuing



Dependency tracking



- Three transactions accessing three variables
 - T1: x, z
 - T2: x, y
 - T3: y, z

Dependency tracking



- Scheduling transaction T is done by evaluating each transaction
 T' in its dependency set.
 - 1. Ignore T' if If T' and T are not each other's direct dependencies.
 - 2. Otherwise:
 - a. Unidirectional -> unconditional wait
 - b. Bidirectional (cyclic) -> lower id executes first

Scheduling transaction T is done by evaluating each transaction
 T' in its dependency set.

Scheduling transaction T is done by evaluating each transaction
 T' in its dependency set.



T1: precedes T3 and T2

Scheduling transaction T is done by evaluating each transaction
 T' in its dependency set.



T1: precedes T3 and T2 T2: waits for T1, precedes T3

Scheduling transaction T is done by evaluating each transaction
 T' in its dependency set.



T1: precedes T3 and T2 T2: waits for T1, precedes T3 **T3: waits for T1 and T2**

• Advantages in comparison

- Advantages in comparison
 - Compared to centralized scheduling:
 - No scheduler bottleneck



- Advantages in comparison
 - Compared to centralized scheduling:
 - No scheduler bottleneck
 - Compared to Locking:
 - No lock contentions



- Advantages in comparison
 - Compared to centralized scheduling:
 - No scheduler bottleneck
 - Compared to Locking:
 - No lock contentions
 - Compared to OCC:
 - Deterministic commit order













More Details in the Paper

- Implementation and optimizations
- Memory management
- TPC-C Benchmark Results
- Future works
- Code and artifacts are available on
 <u>https://github.com/decentralized-scheduling/</u>