

# TicToc: Time Traveling Optimistic Concurrency Control

Authors: Xiangyao Yu, Andrew Pavlo, Daniel Sanchez, Srinivas Devadas

Presented By: Shreejit Nair

# Background: Optimistic Concurrency Control

---

- Read Phase: Transaction executes on a private copy of all accessed objects.
- Validation Phase: Check for conflicts between transactions.
- Write Phase: Transaction's changes to updated objects are made public.

# Background: Timestamp Ordering Algorithm

- A schedule in which the transactions participate is then serializable, and the equivalent serial schedule has the transactions in order of their timestamp values. This is called **timestamp ordering (TO)**.
- The algorithm associates with each database item  $X$  two timestamp (**TS**) values:
  - **Read\_TS( $X$ )**: The **read timestamp** of item  $X$ ; this is the largest timestamp among all the timestamps of transactions that have successfully read item  $X$ —that is,  $\text{read\_TS}(X) = \text{TS}(T)$ , where  $T$  is the *youngest* transaction that has read  $X$  successfully.
  - **Write\_TS( $X$ )**: The **write timestamp** of item  $X$ ; this is the largest of all the timestamps of transactions that have successfully written item  $X$ —that is,  $\text{write\_TS}(X) = \text{TS}(T)$ , where  $T$  is the *youngest* transaction that has written  $X$  successfully.

# Background: Timestamp Ordering Algorithm (Contd)

- Whenever some transaction  $T$  tries to issue a  $\text{read\_item}(X)$  or a  $\text{write\_item}(X)$  operation, the **basic TO** algorithm compares the timestamp of  $T$  with  $\text{read\_TS}(X)$  and  $\text{write\_TS}(X)$  to ensure that the timestamp order of transaction execution is not violated.
- The concurrency control algorithm must check whether conflicting operations violate the timestamp ordering in the following two cases:
  1. Transaction  $T$  issues a **write\_item(X)** operation:
    - a. If  $\text{read\_TS}(X) > \text{TS}(T)$  or if  $\text{write\_TS}(X) > \text{TS}(T)$ , then abort and roll back  $T$  and reject the operation, else execute **write\_item(X)** & set  $\text{write\_TS}(X)$  to  $\text{TS}(T)$ .
  2. Transaction  $T$  issues a **read\_item(X)** operation:
    - a. If  $\text{write\_TS}(X) > \text{TS}(T)$ , then abort and roll back  $T$  and reject the operation, else if  $\text{write\_TS}(X) \leq \text{TS}(T)$ , then execute the  $\text{read\_item}(X)$  operation of  $T$  and set  $\text{read\_TS}(X)$  to the larger of  $\text{TS}(T)$  and the current  $\text{read\_TS}(X)$ .

# Why TicToc?

- Basic T/O (*Timestamp-Ordering*) -based concurrency algorithm involves assigning a unique and monotonically increasing timestamp as serial order for conflict detection.
- This centralized timestamp allocation involves implementing an allocator via a global atomic add operation.
- Actual dependency between two transactions may not agree with the assigned timestamp order causing transactions to unnecessarily abort.
- TicToc computes a transaction's timestamp lazily at commit time based on the data it accesses.
- TicToc timestamp management policy avoids centralized timestamp allocation bottleneck and exploits more parallelism in the workload.

# TicToc Timestamp Management Policy

➤ Consider a sequence of operations

1.  $\mathcal{A}$  read(x)
2.  $\mathcal{B}$  write(x)
3.  $\mathcal{B}$  commits
4.  $\mathcal{A}$  write(y)

What happens when  $TS(\mathcal{B}) < TS(\mathcal{A})$  in basic T/O?

# TicToc Timestamp Commit Invariant

- Every data version in TicToc has a valid range of timestamps bounded by the write timestamp ( $wts$ ) and read timestamp ( $rts$ )

Version	Tuple Data	Timestamp Range
V1	Data	$[wts_1, rts_1]$
V2	Data	$[wts_2, rts_2]$

*Transaction T writes to the tuple*

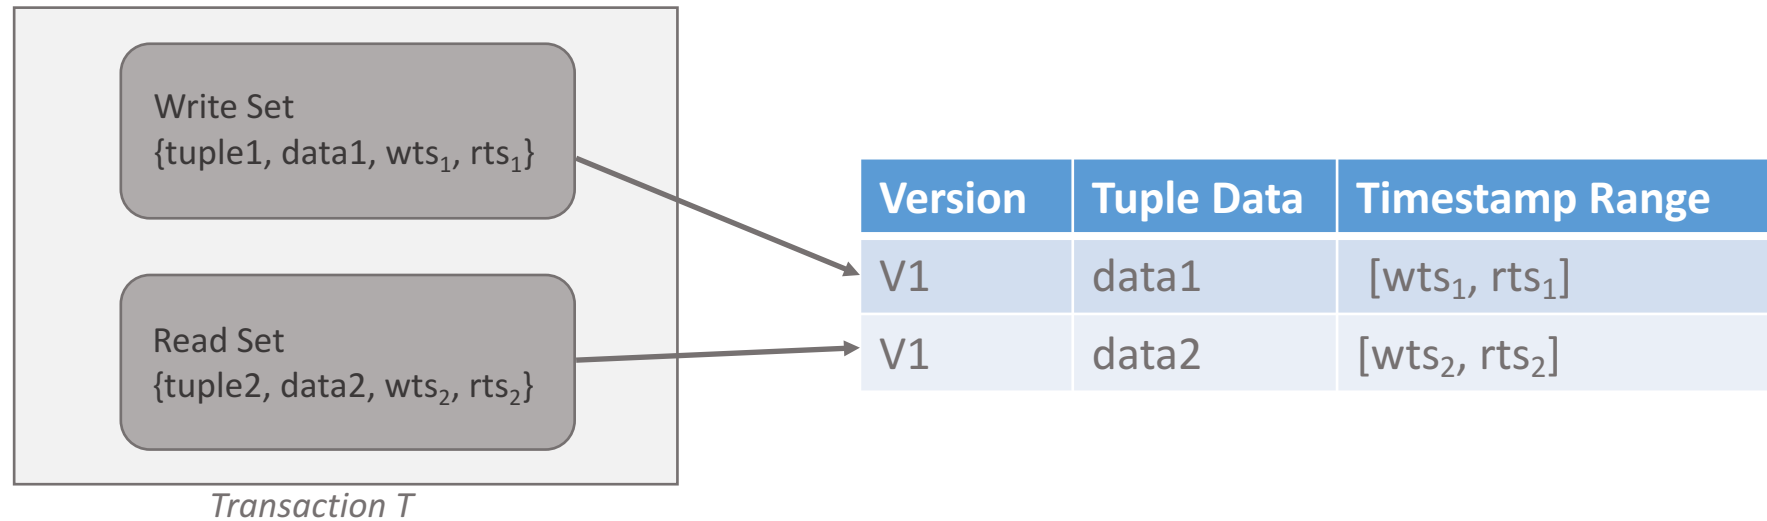
Version	Tuple Data	Timestamp Range
V1	Data	$[wts_1, rts_1]$
V2	Data	$[wts_1, rts_2]$

*Transaction T reads from the tuple*

- Commit timestamp invariant
  - ❖ For all versions read by transaction T,  $v.wts \leq \text{commit\_ts} \leq v.rts$
  - ❖ For all versions written by transaction T,  $v.rts < \text{commit\_ts}$

# TicToc Algorithm

## ➤ Read phase





# TicToc Algorithm (Contd)

## ➤ Validation phase

1. Lock all tuples in the transaction write set
2.  $\text{Commit\_ts} = \max(\max(\text{wts}) \text{ from read set}, \max(\text{rts}) + 1 \text{ from write set})$

---

**Algorithm 2: Validation Phase**

---

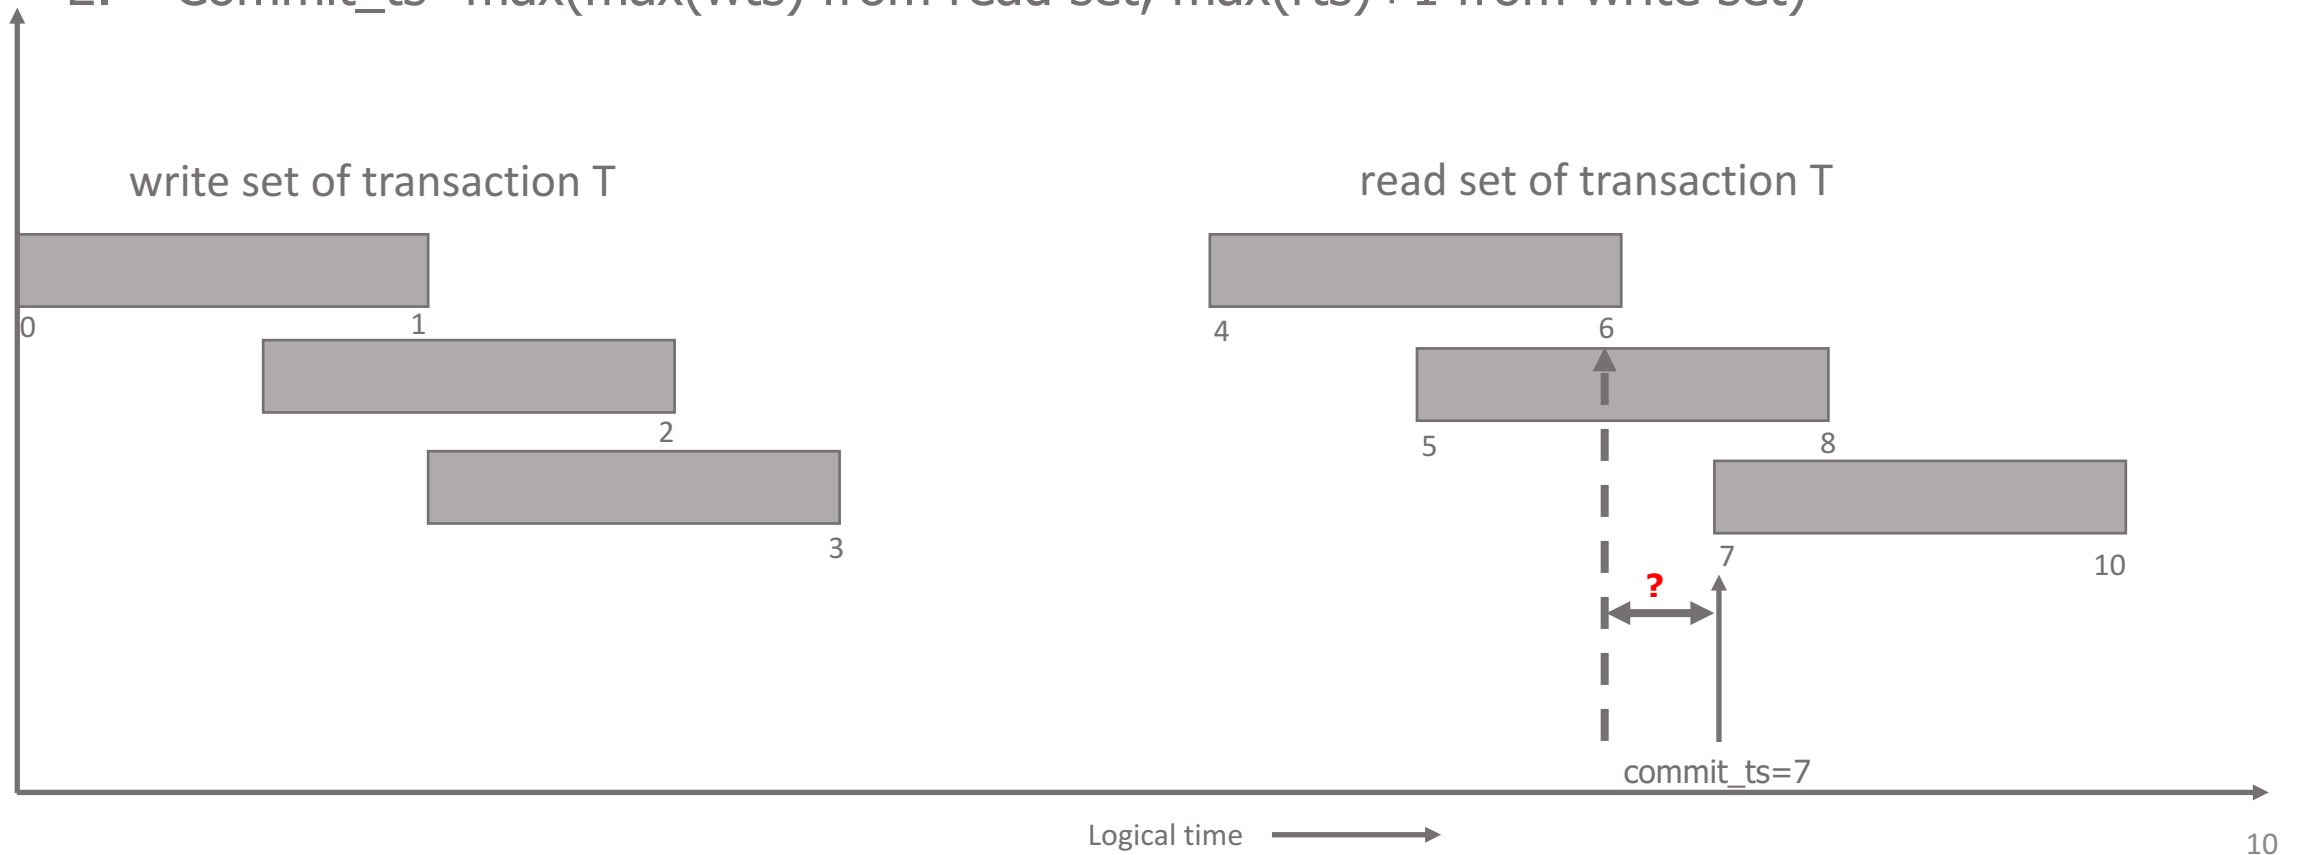
```
Data: read set RS, write set WS  
# Step 1 – Lock Write Set  
1 for w in sorted(WS) do  
2   | lock(w.tuple)  
3 end  
# Step 2 – Compute the Commit Timestamp  
4 commit_ts = 0  
5 for e in WS ∪ RS do  
6   | if e in WS then  
7     | commit_ts = max(commit_ts, e.tuple.rts + 1)  
8   | else  
9     | commit_ts = max(commit_ts, e.wts)  
10  | end  
11 end  
# Step 3 – Validate the Read Set  
12 for r in RS do  
13   | if r.rts < commit_ts then  
14     | # Begin atomic section  
15     | if r.wts ≠ r.tuple.wts or (r.tuple.rts ≤ commit_ts and  
16     | isLocked(r.tuple) and r.tuple not in W) then  
17       | abort()  
18     | else  
19       | r.tuple.rts = max(commit_ts, r.tuple.rts)  
20     | end  
21     | # End atomic section  
22   | end  
23 end
```

---

# TicToc Algorithm (Contd)

## ➤ Validation phase checks

1. Lock all tuples in the transaction write set
2.  $\text{Commit\_ts} = \max(\max(\text{wts}) \text{ from read set}, \max(\text{rts}) + 1 \text{ from write set})$



# TicToc Algorithm (Contd)

---

## ➤ Write phase

For all tuples in  $WS$ (write set) do:

1. *commit* updated values to database
2. *overwrite* tuple.wts = tuple.rts = commit\_ts
3. *unlock(tuple)*

# TicToc Working Example

- Step 1: Transaction A reads tuple x

Version	Tuple Data	Timestamp Range
V1	x	[wts=1, rts=3]
V1	y	[wts=1, rts=2]

Read set A = {x,1,3}

- Step 2: Transaction B writes to tuple x and commits at timestamp 4

Version	Tuple Data	Timestamp Range
V2	x	[wts=4, rts=4]
V1	y	[wts=1, rts=2]

Read set A = {x,1,3}

- Step 3: Transaction A writes to tuple y

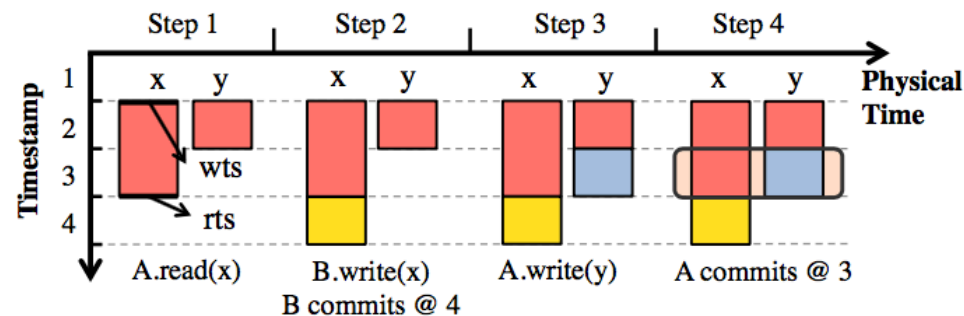
Version	Tuple Data	Timestamp Range
V2	x	[wts=4, rts=4]
V1	y	[wts=1, rts=2]

Read set A = {x,1,3}  
Write set A = {y,1,2}

- Step 4: Transaction A enters validation phase

Version	Tuple Data	Timestamp Range
V2	x	[wts=4, rts=4]
V2	y	[wts=3, rts=3]

Read set A = {x,1,3}  
Write set A = {y,1,2}  
Tran A commit\_ts = 3  
**Tran A COMMITS**



**Figure 1:** An example of two transactions executing using TicToc.

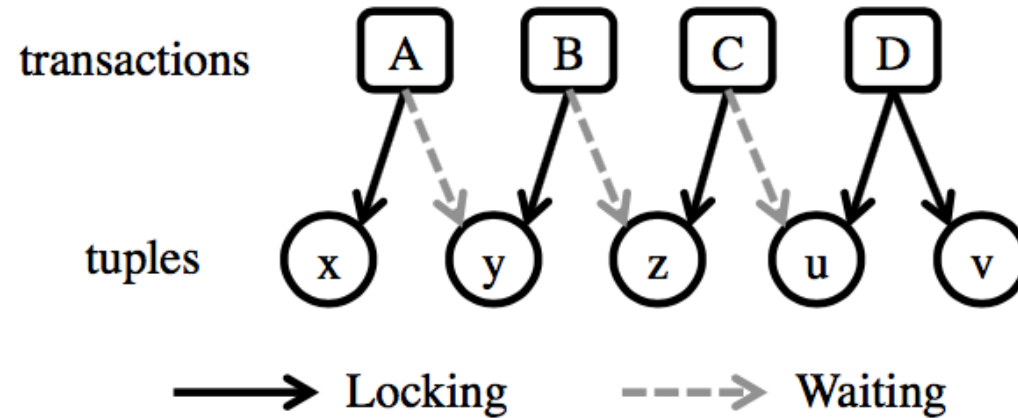
# TicToc Serializability Order

$$A <_s B \triangleq A <_{lts} B \vee (A =_{lts} B \wedge A \leq_{pts} B)$$

- LEMMA 1 : Transactions writing to the same tuples must have different commit timestamps (lts).
- LEMMA 2: Transactions that commit at the same logical timestamp and physical timestamp do not conflict with each other (e.g. Read-Write or Write-Read operations on the same tuples by different transactions).
- LEMMA 3: A *read* operation from a committed transaction returns the value of the latest *write* to the tuple in the serial schedule.

# TicToc Optimizations

- No-Wait locking in validation phase



**Figure 2:** An example of lock thrashing in a 2PL protocol.

# TicToc Optimizations (Contd)

---

## ➤ Preemptive Aborts

- ❖ Validation phase causes other transactions to potentially block unnecessarily.
- ❖ Guessing an approximate commit timestamp to observe if transactions would lead to aborts.

# Timestamp History Buffer

- Step 1: Transaction A reads tuple x

Version	Tuple Data	Timestamp Range
V1	x	[wts=1, rts=2]

Read set A = {x,1,2}

- Step 2: Transaction B extends x's rts.

Version	Tuple Data	Timestamp Range
V2	x	[wts=1, rts=3]

Read set A = {x,1,2}

- Step 3: Transaction C writes to tuple x

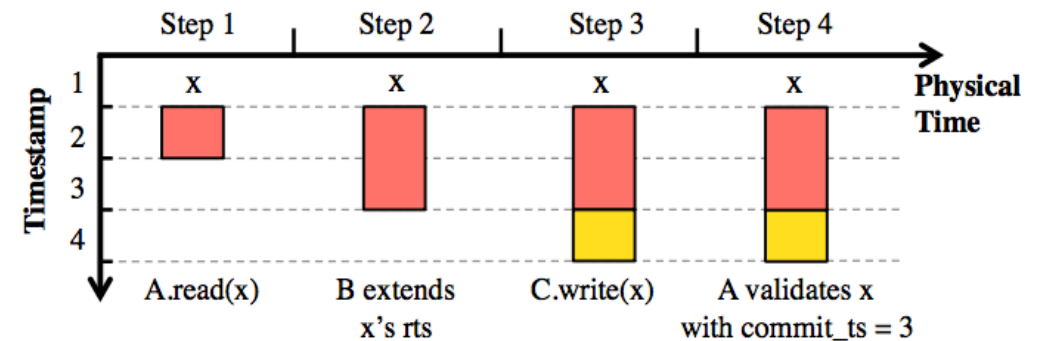

Version	Tuple Data	Timestamp Range
V3	x	[wts=4, rts=4]

Read set A = {x,1,2}  
Tran C commit\_ts = 4

- Step 4: Transaction A enters validation phase

Version	Tuple Data	Timestamp Range
V3	x	[wts=4, rts=4]

Read set A = {x,1,2}  
Tran A commit\_ts = 3



**Figure 3:** Using a tuple's timestamp history to avoid aborting.



# Experimental Evaluation

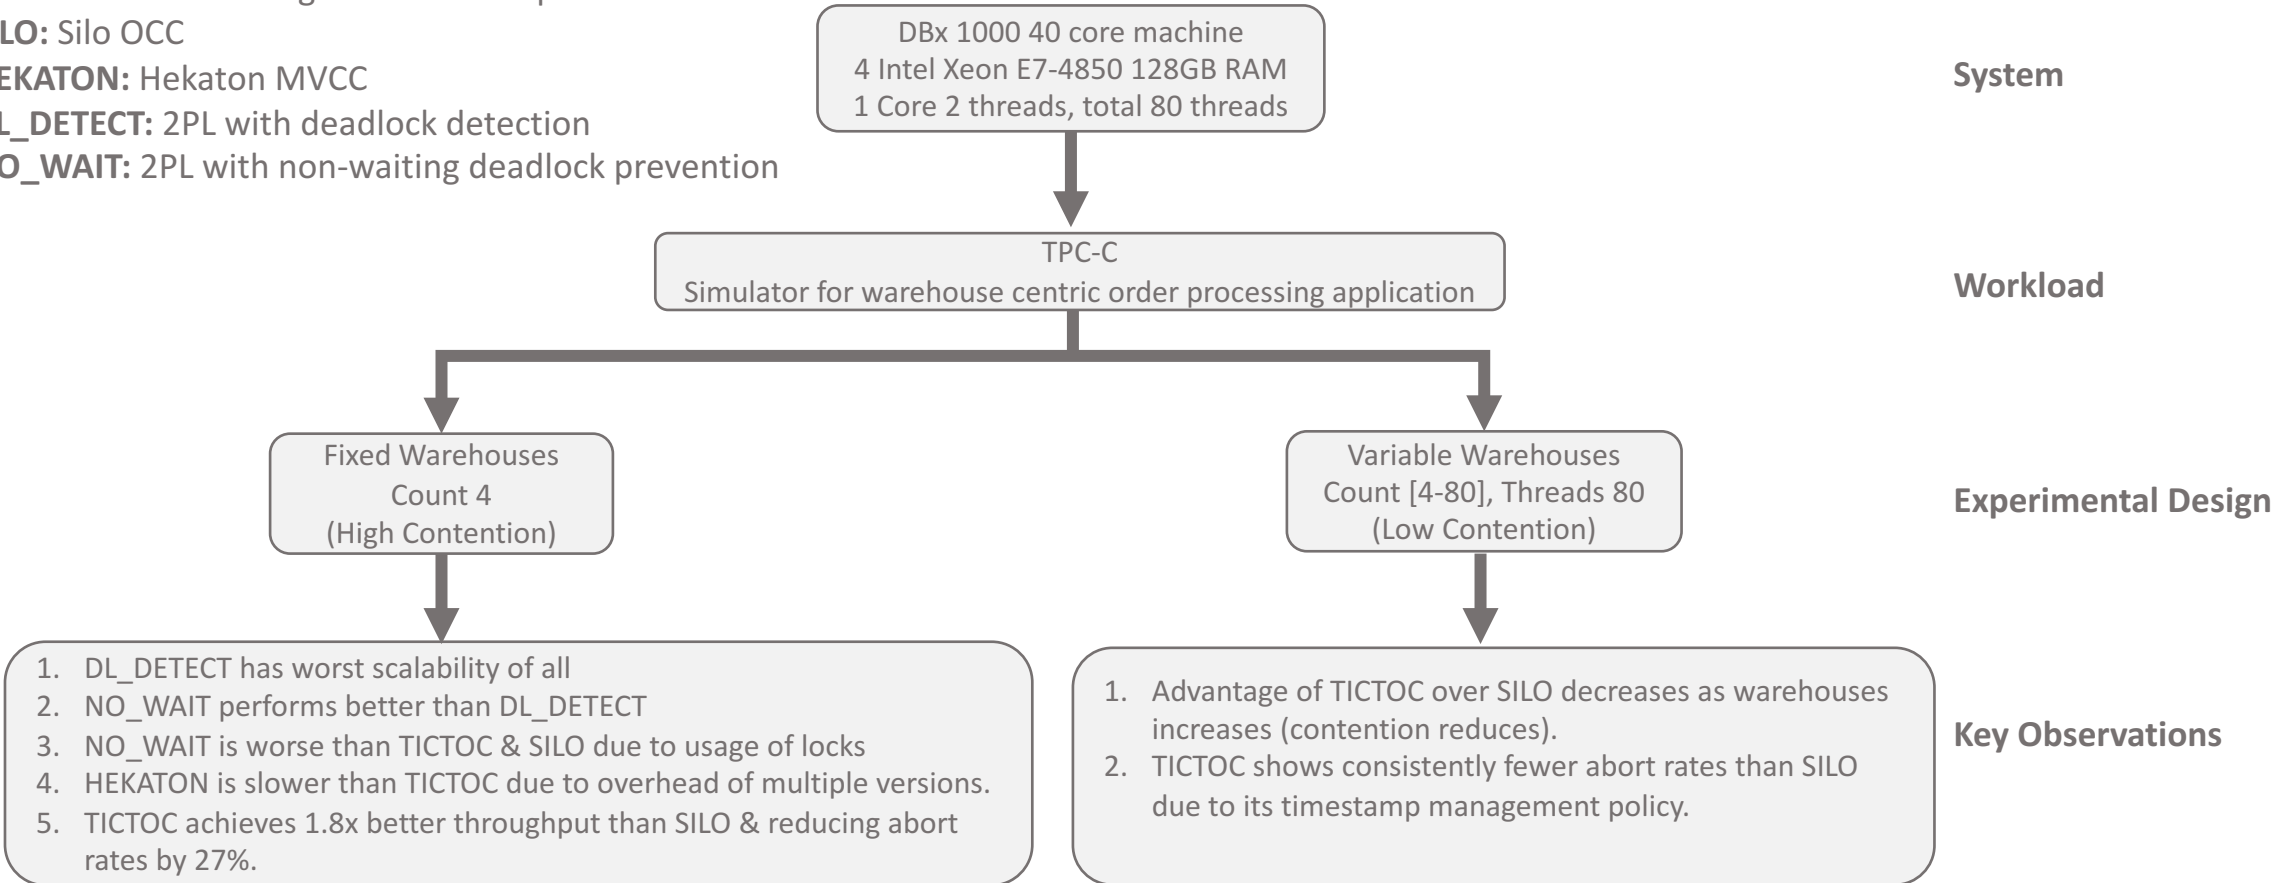
**TICTOC:** Time traveling OCC with all optimizations

**SILO:** Silo OCC

**HEKATON:** Hekaton MVCC

**DL\_DETECT:** 2PL with deadlock detection

**NO\_WAIT:** 2PL with non-waiting deadlock prevention



**System**

**Workload**

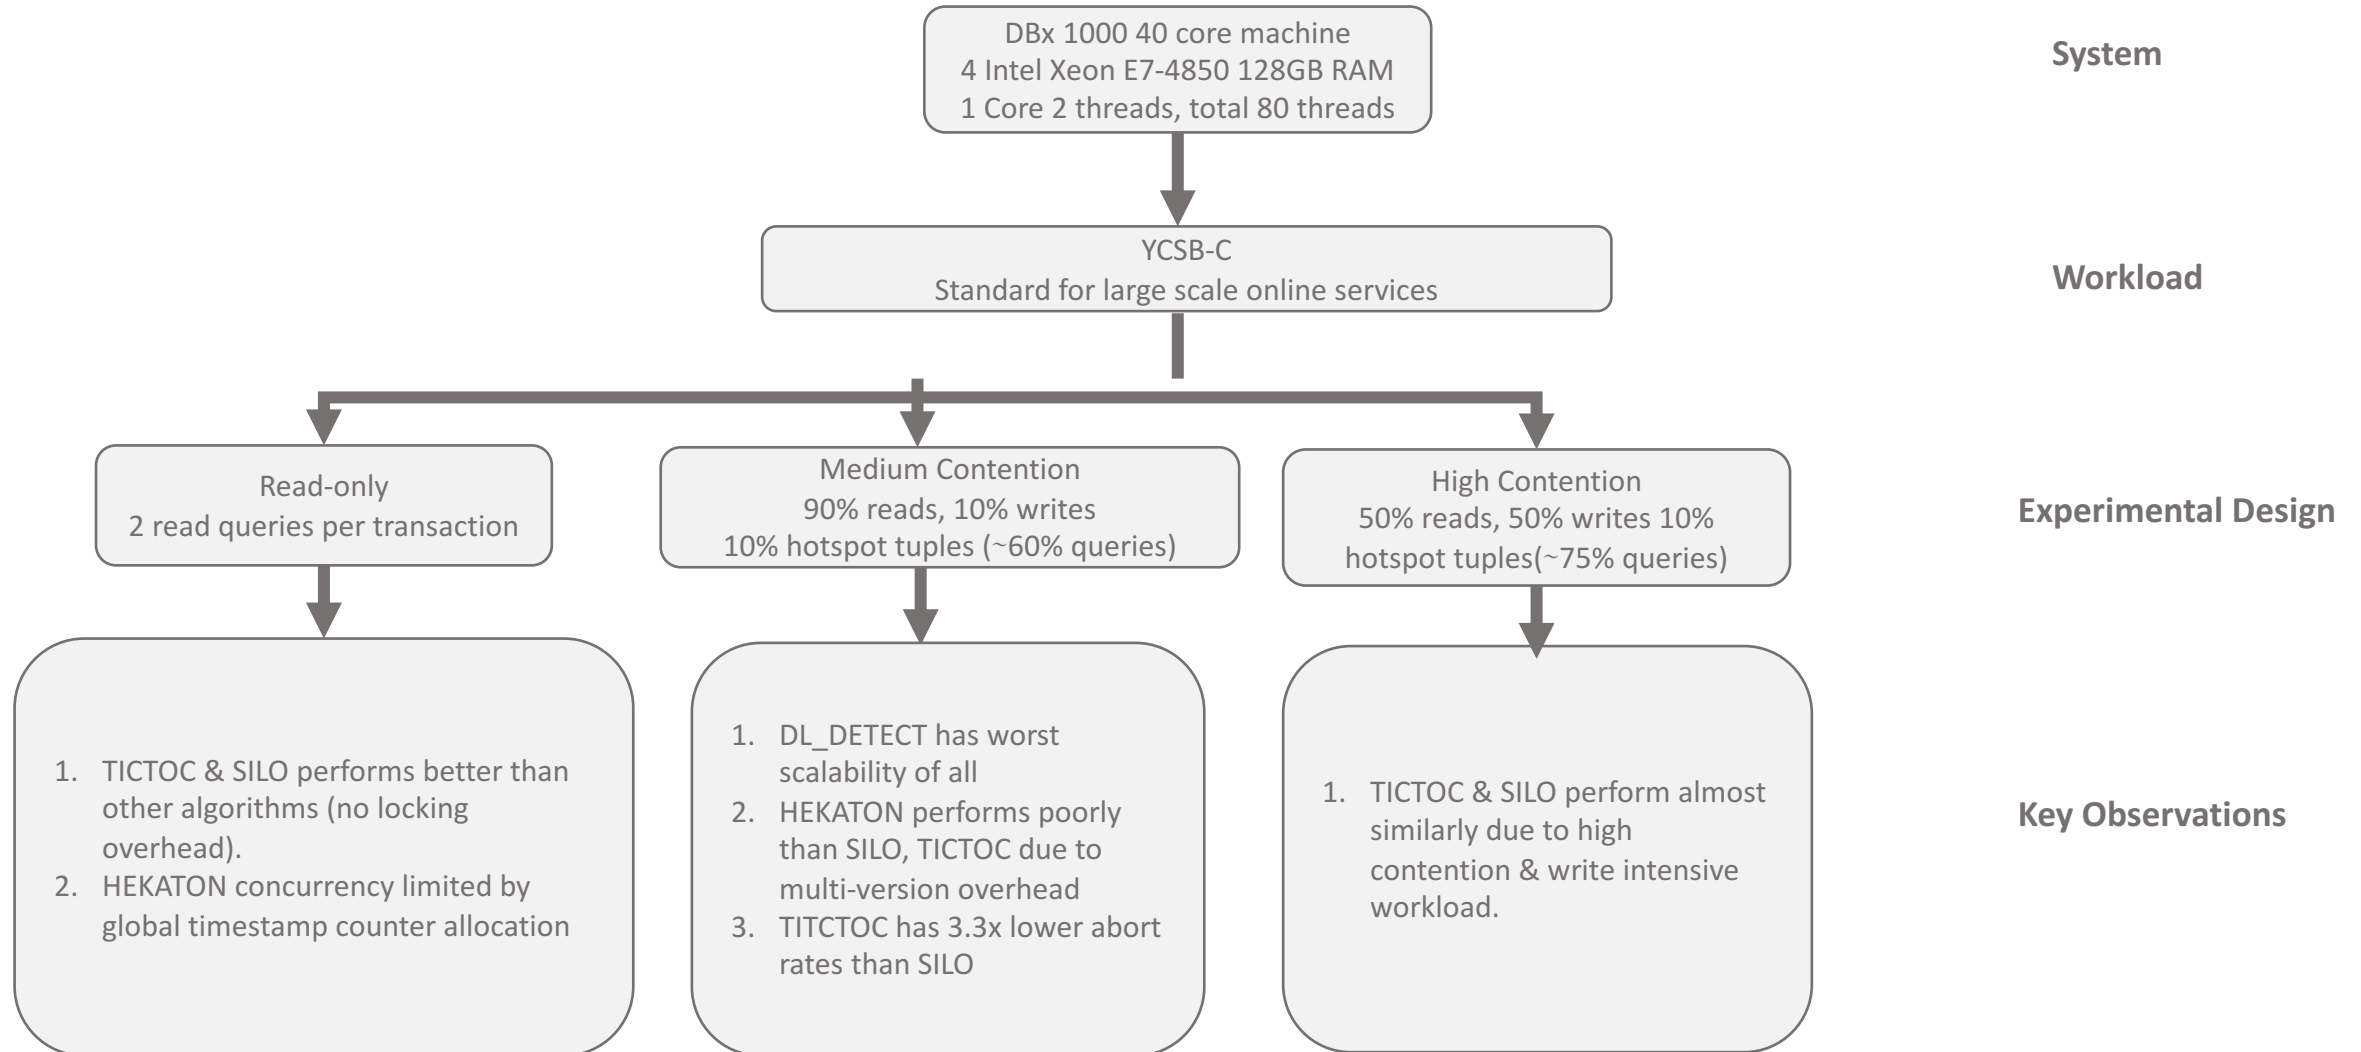
**Experimental Design**

**Key Observations**

1. DL\_DETECT has worst scalability of all
2. NO\_WAIT performs better than DL\_DETECT
3. NO\_WAIT is worse than TICTOC & SILO due to usage of locks
4. HEKATON is slower than TICTOC due to overhead of multiple versions.
5. TICTOC achieves 1.8x better throughput than SILO & reducing abort rates by 27%.

1. Advantage of TICTOC over SILO decreases as warehouses increases (contention reduces).
2. TICTOC shows consistently fewer abort rates than SILO due to its timestamp management policy.

# Experimental Evaluation (Contd)



# Conclusion

---

- The paper presented TicToc, a new OCC-based concurrency control algorithm that eliminates the need for centralized timestamp allocation.
- TicToc decouples logical timestamps and physical time by deriving transaction commit timestamps from data items.
- Key features include exploiting more parallelism and reducing transaction abort rates.
- TicToc achieves up to 92% higher throughput while reducing transaction abort rates by up to 3.3x under different workload conditions.

# Thoughts...

---

- TicToc is definitely one of the better performing OCC algorithm.
- Reducing contention within the validation phase?
- Need for write set validation in the validation phase?