

# TicToc: Time Traveling Optimistic Concurrency Control

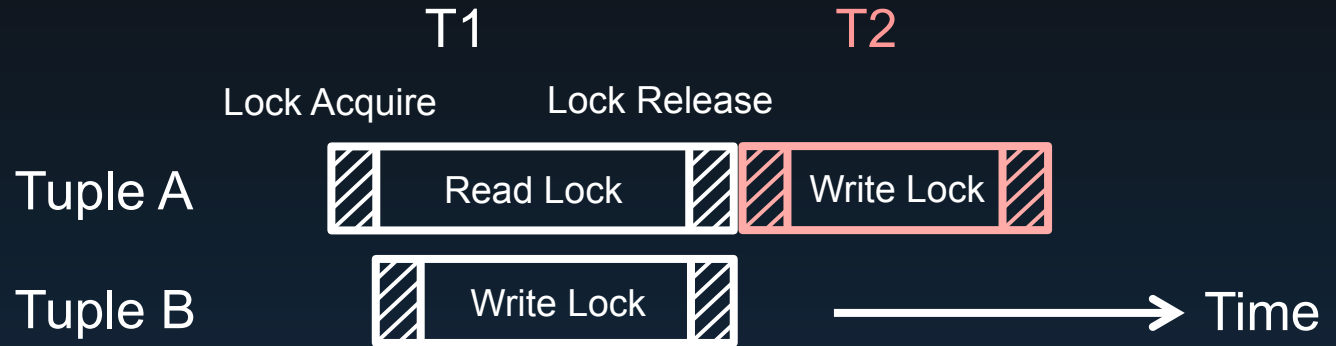
Xiangyao Yu<sup>1</sup>, Andrew Pavlo<sup>2</sup>, Daniel Sanchez<sup>1</sup>, Srinivas Devadas<sup>1</sup>

Massachusetts Institute of Technology<sup>1</sup>  
Carnegie Mellon University<sup>2</sup>

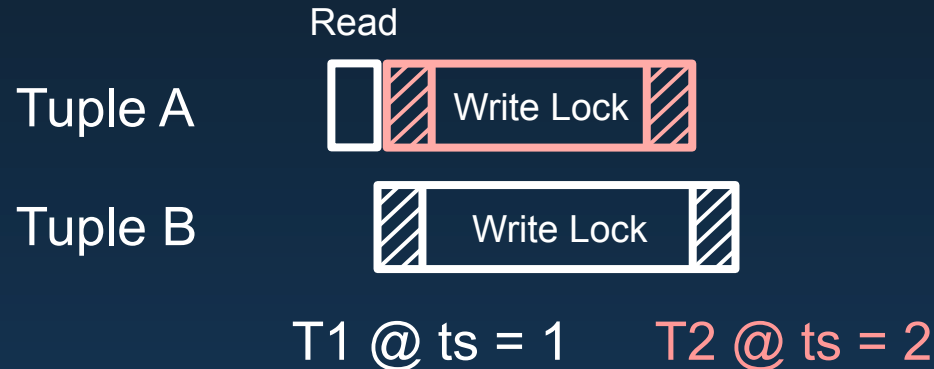


# PHYSICAL VS LOGICAL TIME

Two-Phase Locking (2PL)



Timestamp Ordering (T/O)



# TIMESTAMP ORDERING

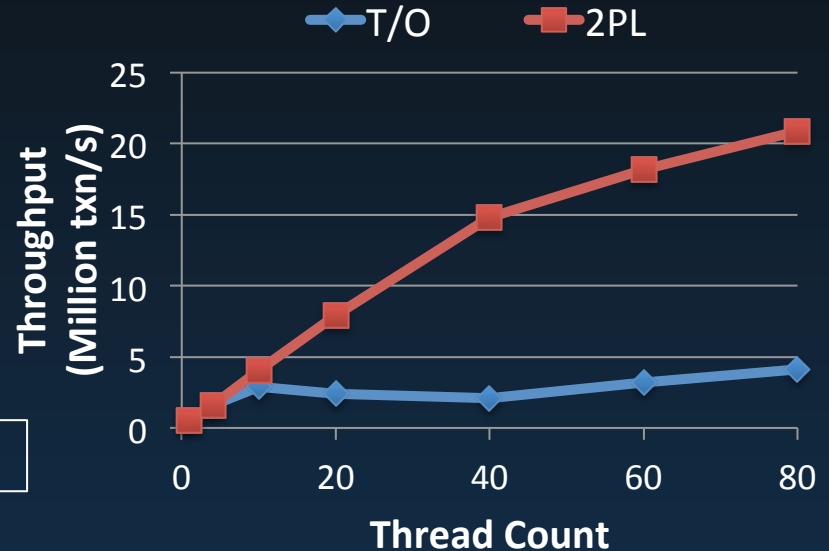
## Timestamp Allocation

(txns have unique timestamps)

- Centralized Allocator
  - Timestamp allocation is a scalability bottleneck

```
ts = __sync_fetch_and_add(&glob_ts, 1)
```

- Synchronized Clock
  - Clock skew causes unnecessary aborts

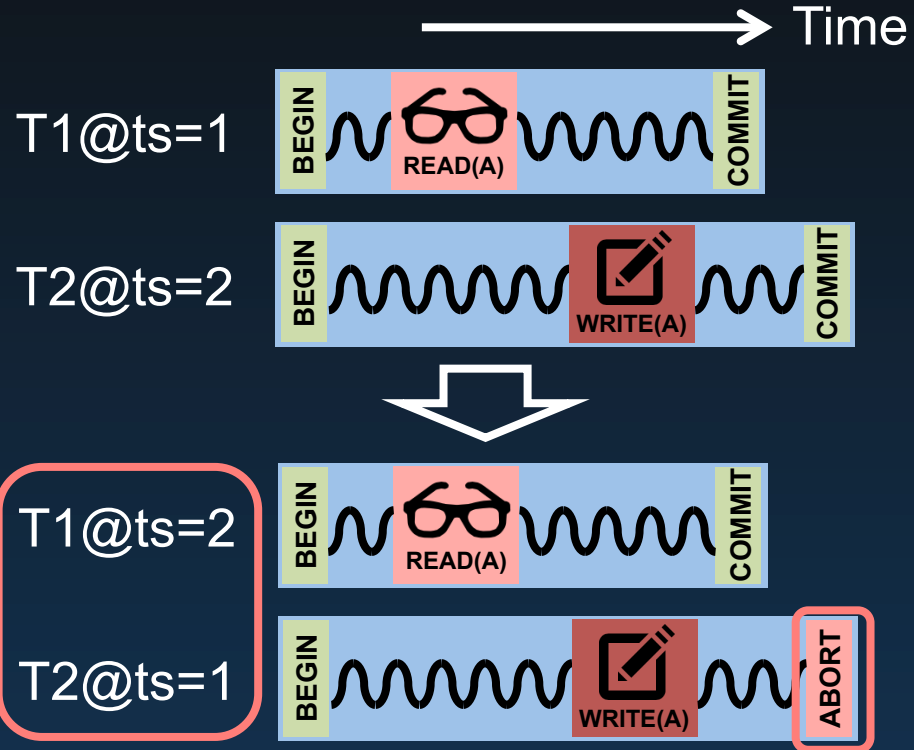


# TIMESTAMP ORDERING

Timestamp Allocation  
(txns have unique timestamps)

## Static Timestamp Assignment

- Txns abort due to bad timestamp assignment



# DATA DRIVEN TIMESTAMP MANAGEMENT

## Traditional T/O

1. Acquire timestamp (TS)
2. Determine tuple visibility using TS

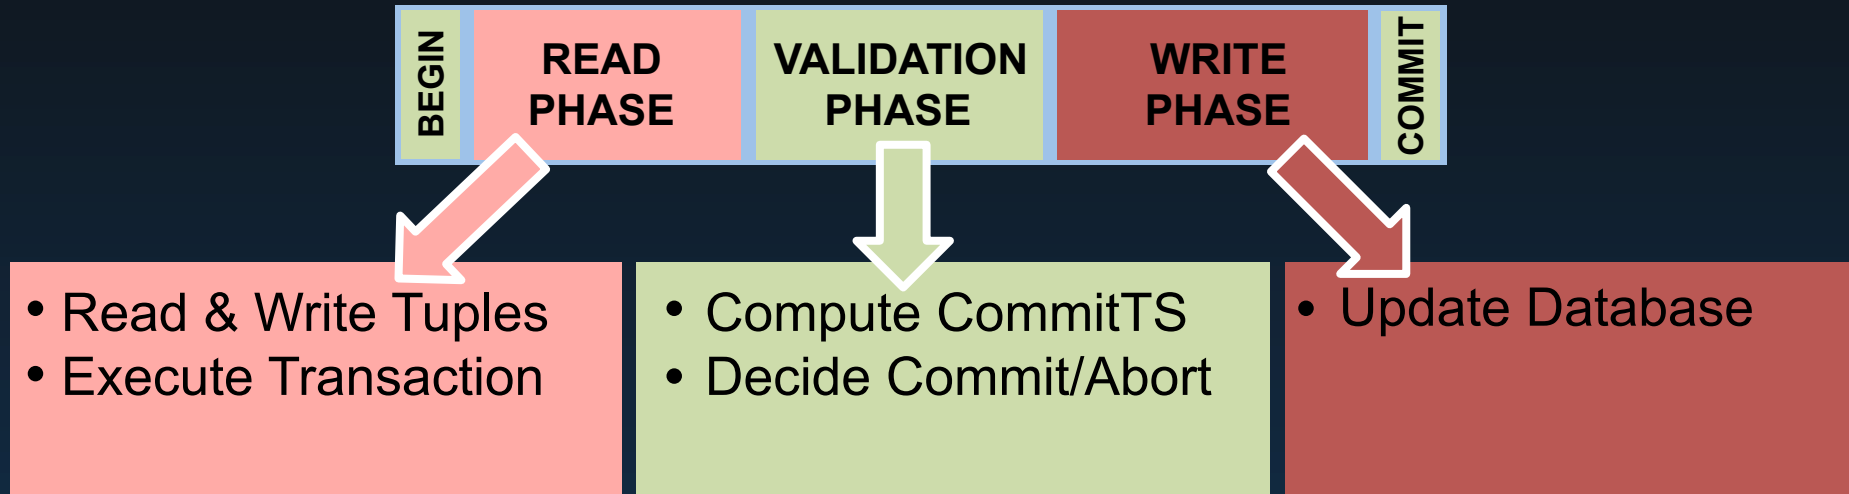
-  Timestamp Allocation
-  Static Timestamp Assignment

## TicToc


1. Access tuples and remember their timestamp info.
2. Compute commit timestamp (CommitTS)

-  No Timestamp Allocation
-  Dynamic Timestamp Assignment

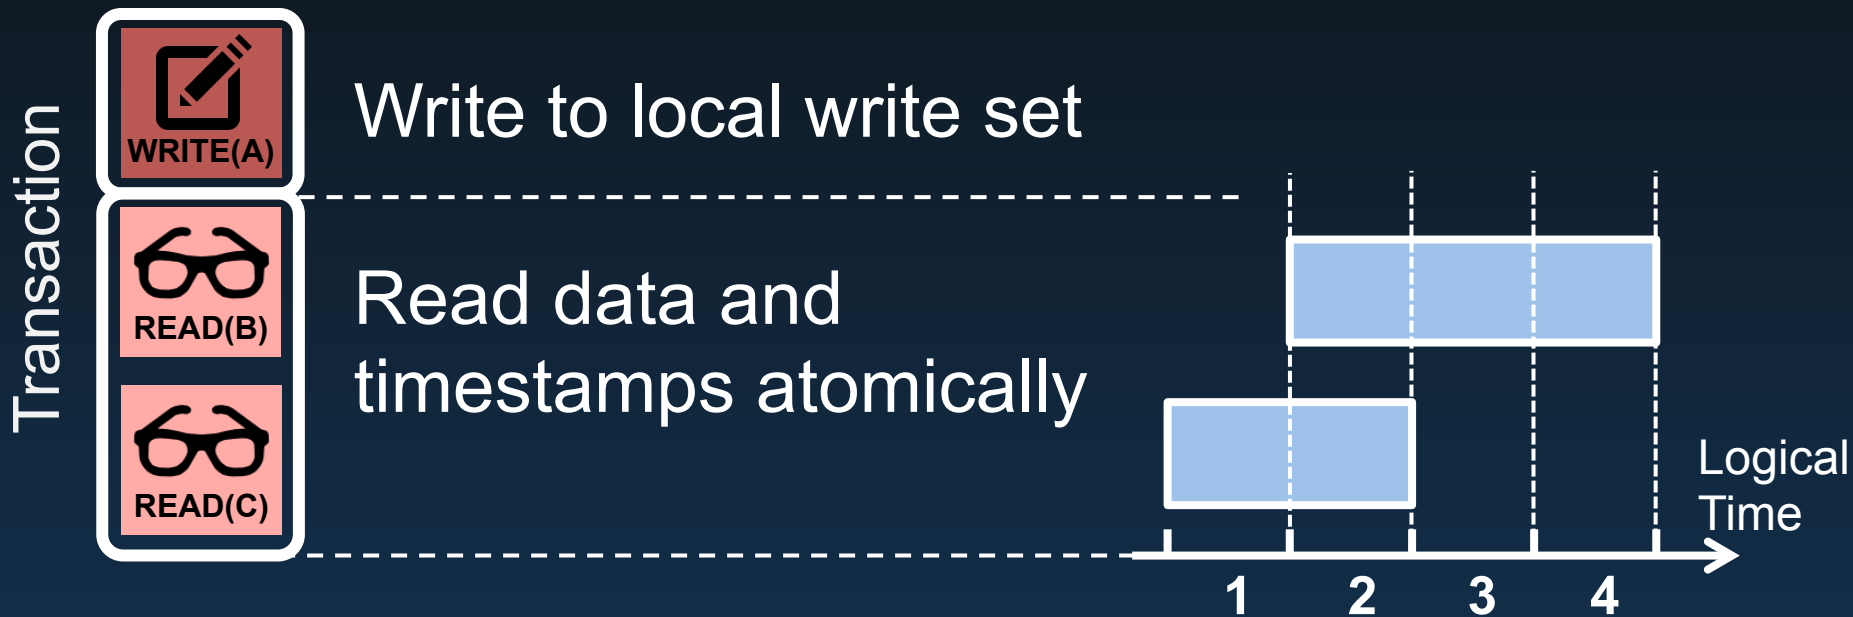
# OPTIMISTIC CONCURRENCY CONTROL (OCC)



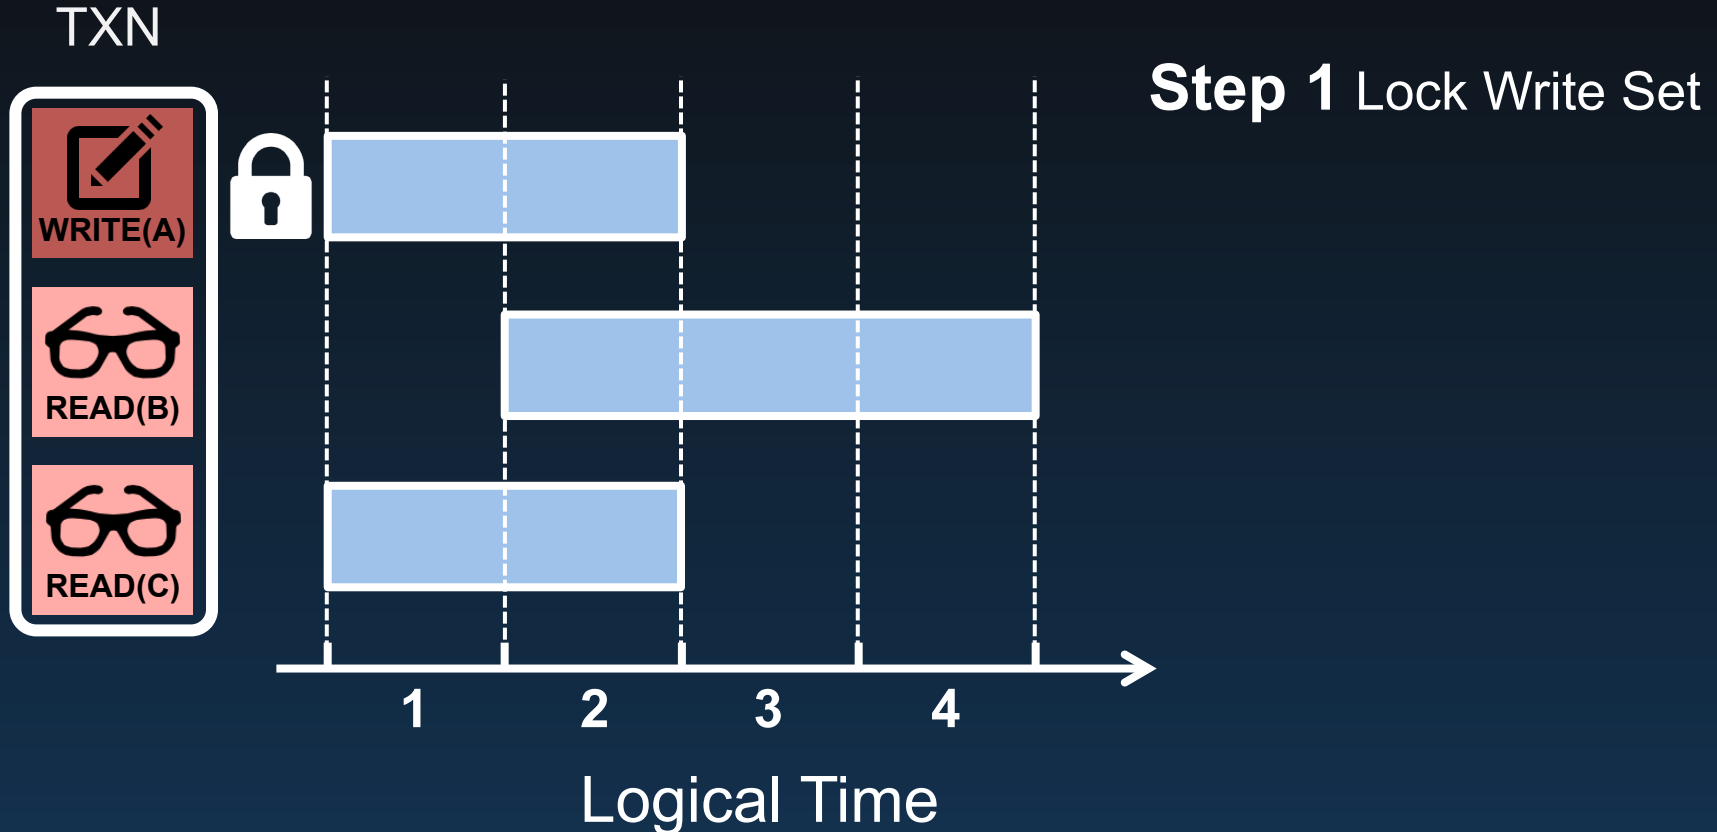
Tuple Format



# READ PHASE

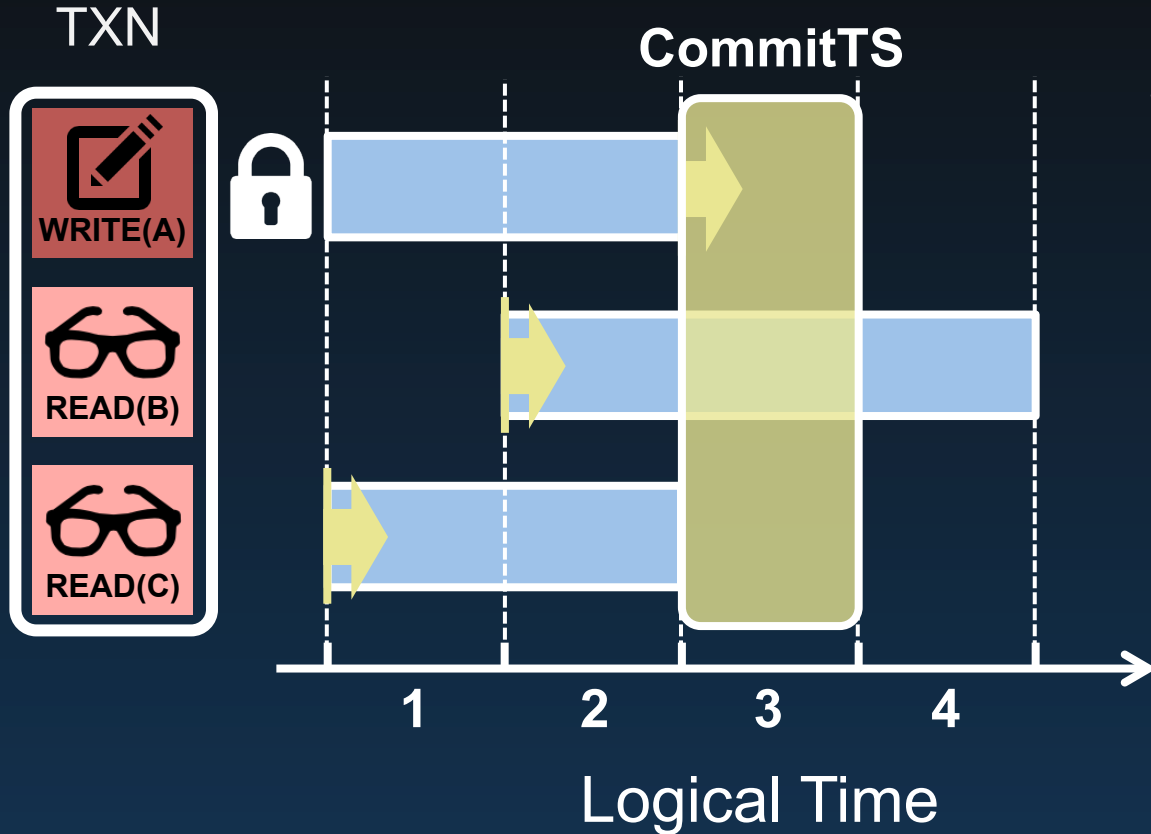


# VALIDATION PHASE





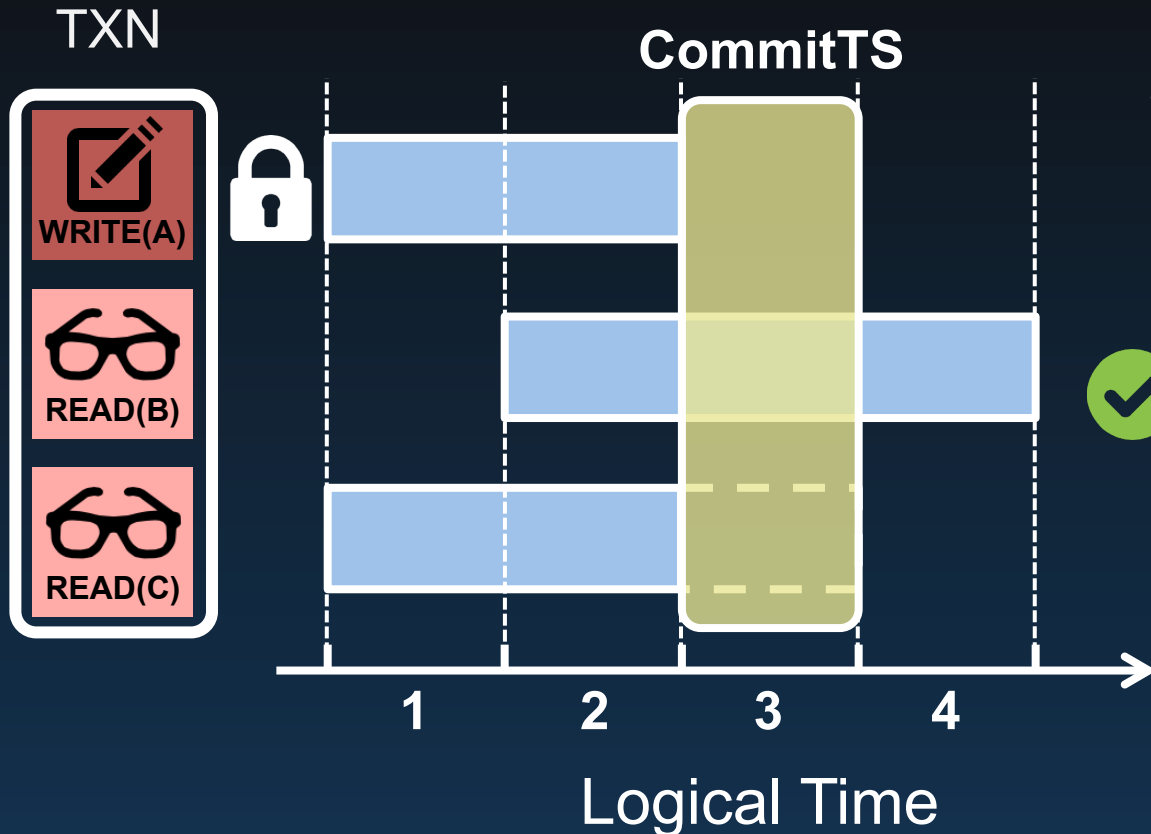
# VALIDATION PHASE



Step 1 Lock Write Set

Step 2 Compute CommitTS

# VALIDATION PHASE



**Step 1** Lock Write Set


**Step 2** Compute CommitTS

**Step 3** Validate Read Set



Case 1: latest version

# VALIDATION PHASE



**Step 1** Lock Write Set

**Step 2** Compute CommitTS

**Step 3** Validate Read Set

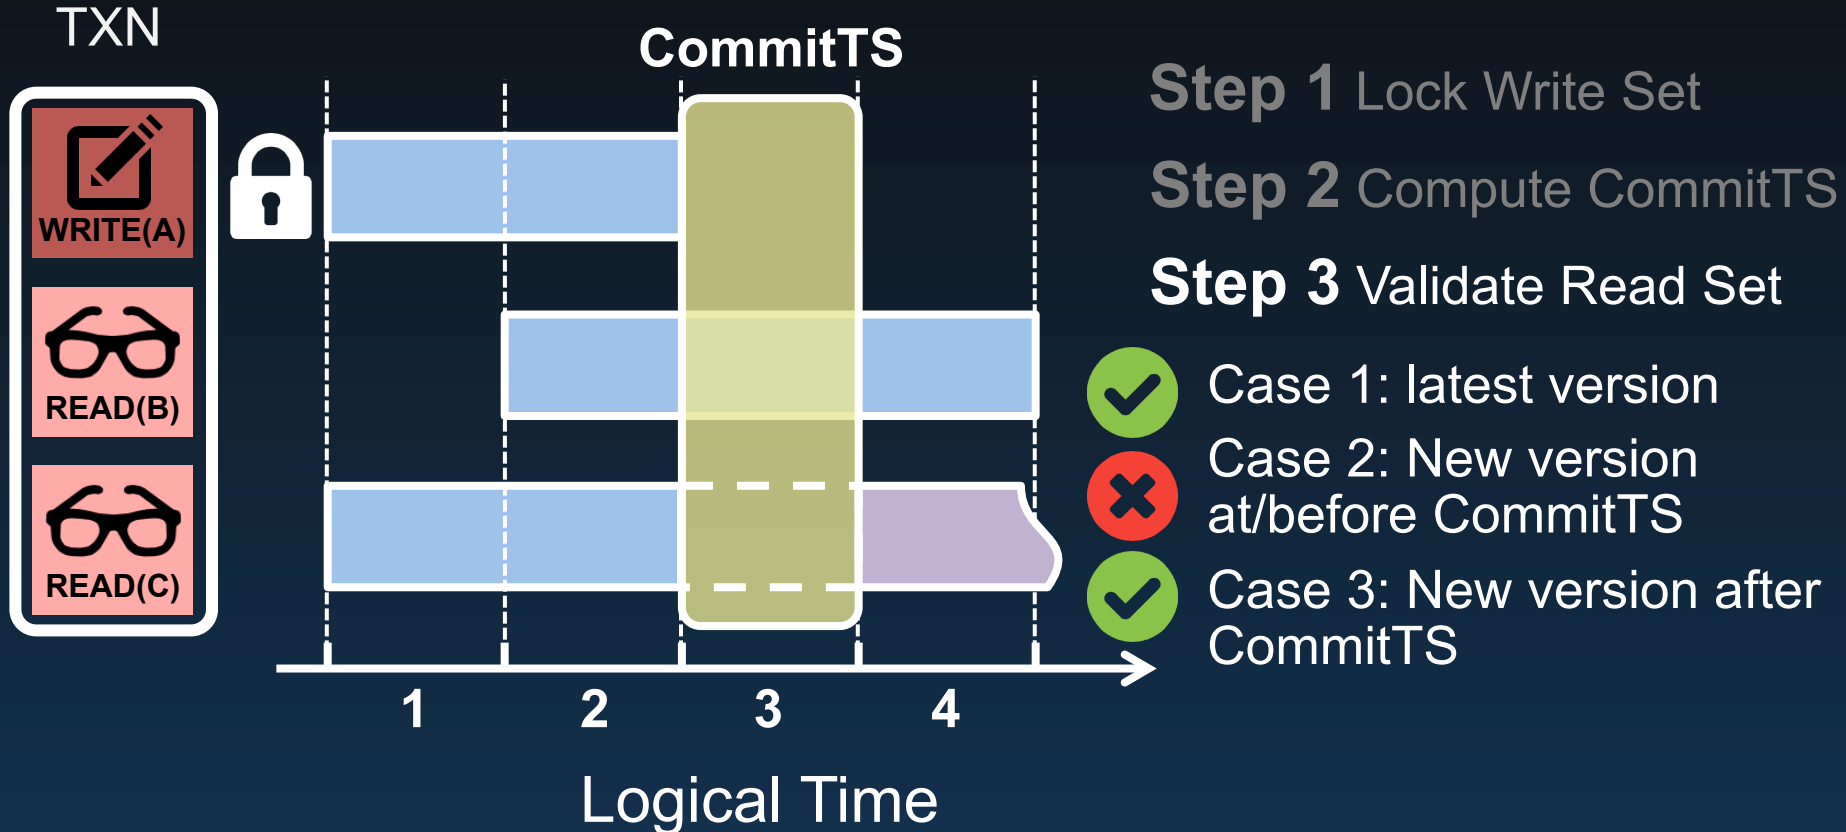


Case 1: latest version

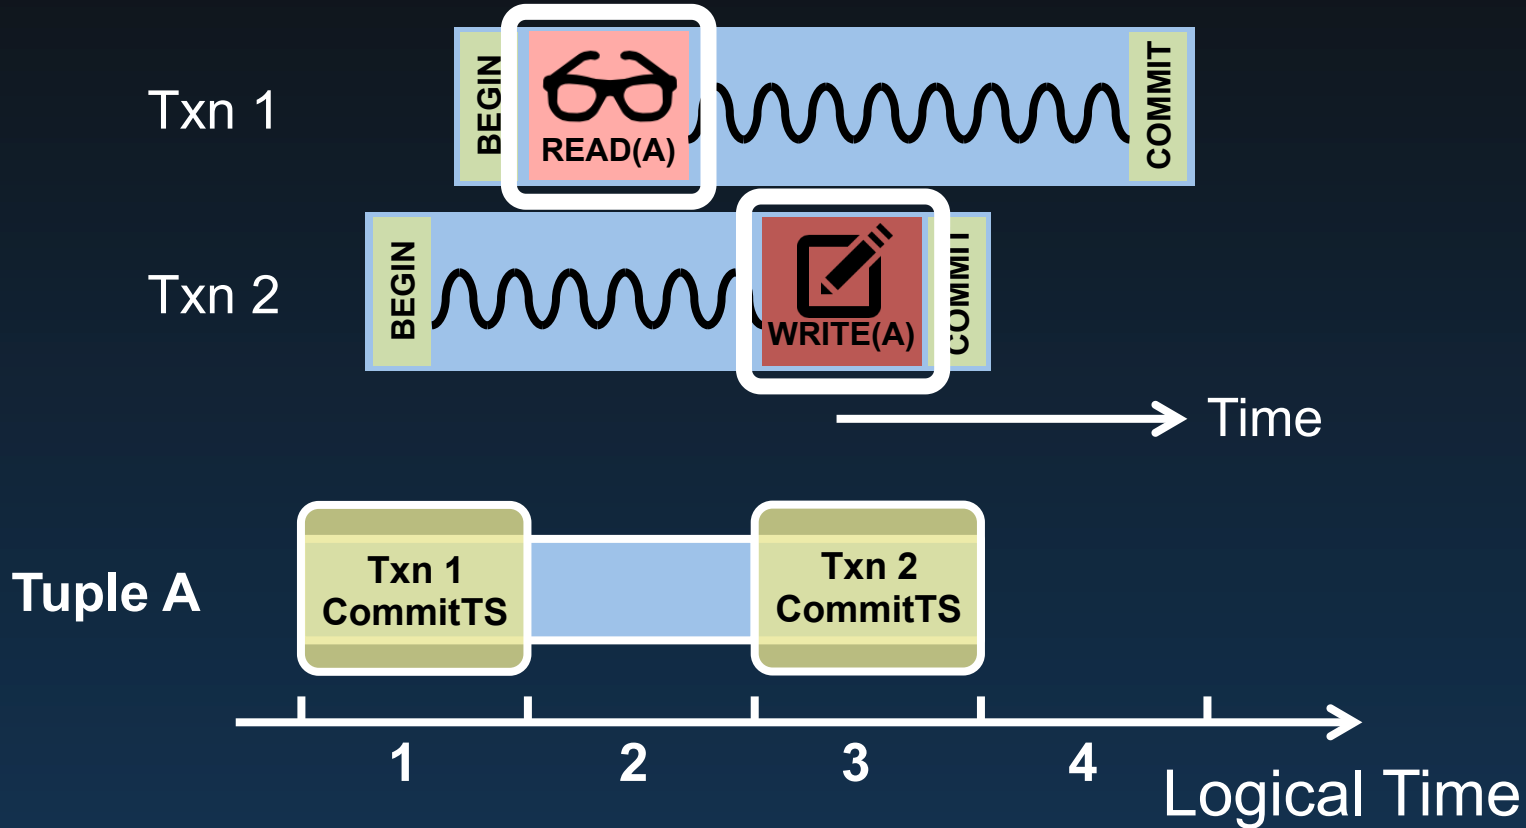


Case 2: New version  
at/before CommitTS

# VALIDATION PHASE



# TIME TRAVELING



# TICTOC OPTMIZATIONS

- No-wait locking in validation phase
- Preemptive aborts
- Timestamp history
- Lower isolation levels


# EXPERIMENTAL SETUP

- 4-socket, 40-core Machine (80 threads w/ hyper-threading)
- Main Memory DBx1000
  - No logging
  - No B-tree (Hash indexing)
- Concurrency Control Algorithms
  - MVCC: HEKATON
  - OCC: SILO
  - 2PL: DL\_DETECT, NO\_WAIT


# YSCB Medium Contention

TICTOC HEKATON DL\_DETECT NO\_WAIT SILO

## Throughput



## Abort Rate

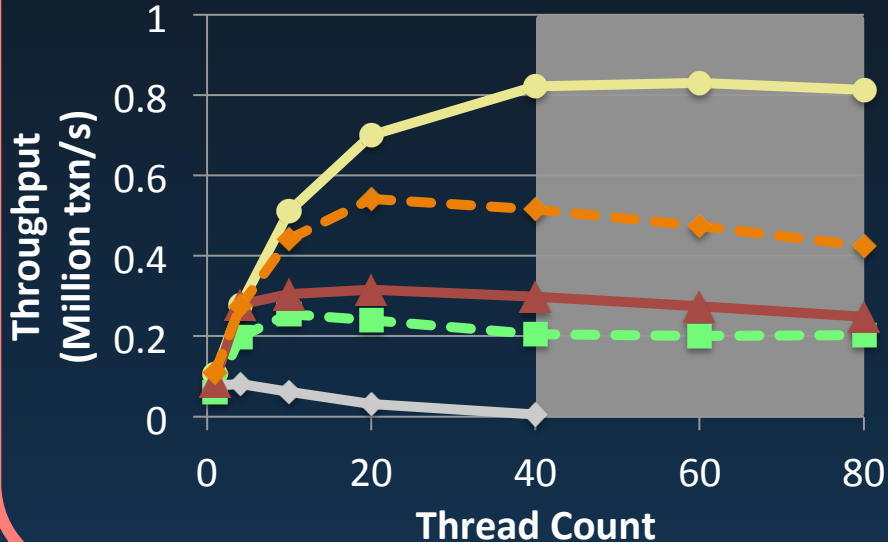




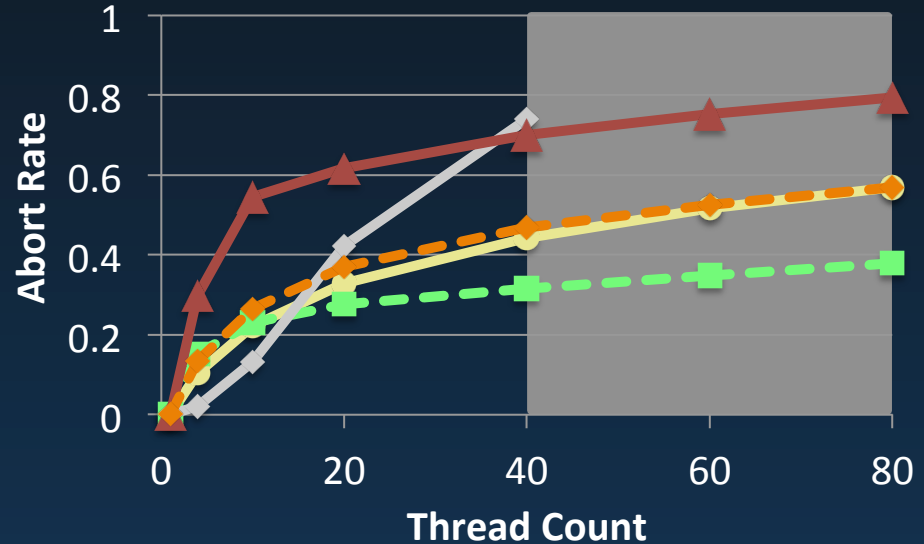
# YSCB High Contention

TICTOC HEKATON DL\_DETECT NO\_WAIT SILO

## Throughput



## Abort Rate





Xiangyao Yu  
xyx@mit.edu



Andy Pavlo  
pavlo@cs.cmu.edu



Daniel Sanchez  
sanchez@mit.edu



Srinivas Devadas  
devadas@mit.edu



DBx1000 (<https://github.com/yxymit/DBx1000>)

