

On the Correctness of Transactional Memory

Rachid Guerraoui Michał Kapałka

EPFL, Switzerland

PPoPP 2008

Fundamental Question

What is a TM precisely?

Fundamental Question

What is a TM precisely?

When is a TM correct?

Serializability, linearizability, . . . do not answer those!

Important Consequences

- Hard to compare TMs
- Hard to prove a TM correct
- Hard to show inherent limitations of TMs

Our Take

Opacity – a correctness condition for TMs
(ensured by many TMs)



ACID / serializability for DBs

Opacity

- Compare TMs
opacity is implementation-agnostic
- Prove a TM correct
graph interpretation of opacity
- Show inherent limitations of TMs
time complexity lower bound

Outline

- 1 TM semantics: intuition
- 2 Existing correctness criteria: examples
- 3 Opacity: step-by-step

TM Semantics: Perspective

1-lock semantics

```
atomic {  
   $a = x$   
   $c . add(5)$   
   $s . push(a)$   
}
```


TM Semantics: Perspective

```
atomic {  
   $a = x$   
   $c . add(5)$   
   $s . push(a)$   
}
```

1-lock semantics



transaction

TM Semantics: Perspective

```
atomic {  
   $a = x$   
   $c . add(5)$   
   $s . push(a)$   
}
```

1-lock semantics



transaction

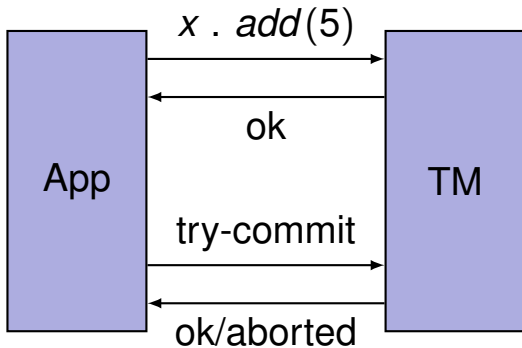


opacity

TM Semantics

- Committed: instantaneous
- Aborted: never visible
- All: observe consistent state

Model of Interaction



Outline

- 1 TM semantics: intuition
- 2 Existing correctness criteria: examples
- 3 Opacity: step-by-step

Serializability + Recoverability

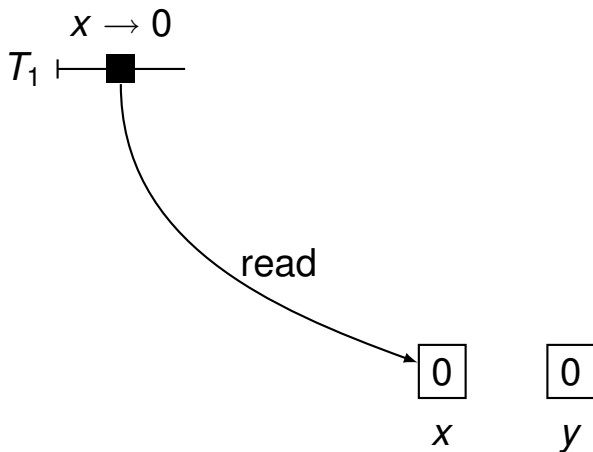
Serializability: [Papadimitriou '79]

- Committed: instantaneous



Recoverability: [Hadzilacos '88]

- If T updates x ,
no tx accesses x
until T commits/aborts.



Ser. / Recov. Too Weak




Ser. / Recov. Too Weak






Ser. / Recov. Too Weak




Ser. / Recov. Too Weak





Ser. / Recov. Too Weak



Ser. / Recov. Too Strict



Ser. / Recov. Too Strict



Other Desirable Properties

- Arbitrary objects
- Multiple versions of each object
- Updates at any time
- Real-time order
- User's perspective

Outline

- 1 TM semantics: intuition
- 2 Existing correctness criteria: examples
- 3 **Opacity: step-by-step**

Opacity: Question to Answer

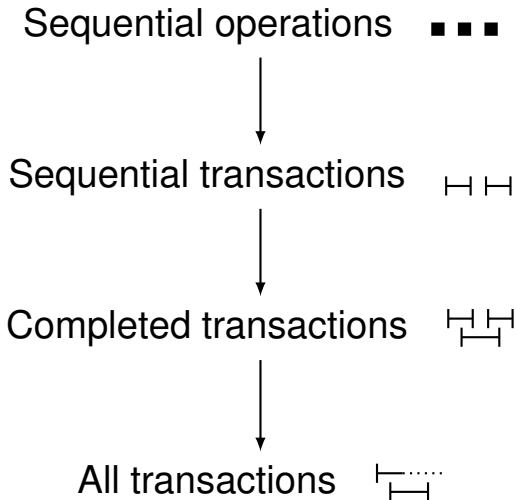
Given:

transactional execution
(sequence of invocations & responses)

Answer:

is the execution correct?

Opacity: Step by Step




Sequential Operations

$s . push(3)$ $x = 5$ $s . pop \rightarrow 3$ $x \rightarrow 5$
■ ■ ■ ■


Sequential specification of s and x



Sequential Transactions



commit



commit
(abort)



Every transaction is **legal**

Concurrent, Complete Transactions


$s . push(3) \quad x = 5$
|-----■-----■-----| com.

$s . pop \rightarrow 3$
|-----■-----| com.

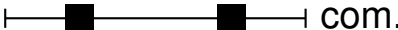
$s . pop \rightarrow 3$ $x \rightarrow 5$
|-----■-----■-----| abort

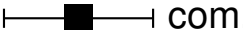



Concurrent, Complete Transactions

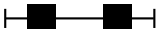


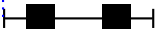
Concurrent, Complete Transactions

$s . push(3) \quad x = 5$
 com.

$s . pop \rightarrow 3$
 com.

$s . pop \rightarrow 3$ $x \rightarrow 5$
 abort









Concurrent, Complete Transactions

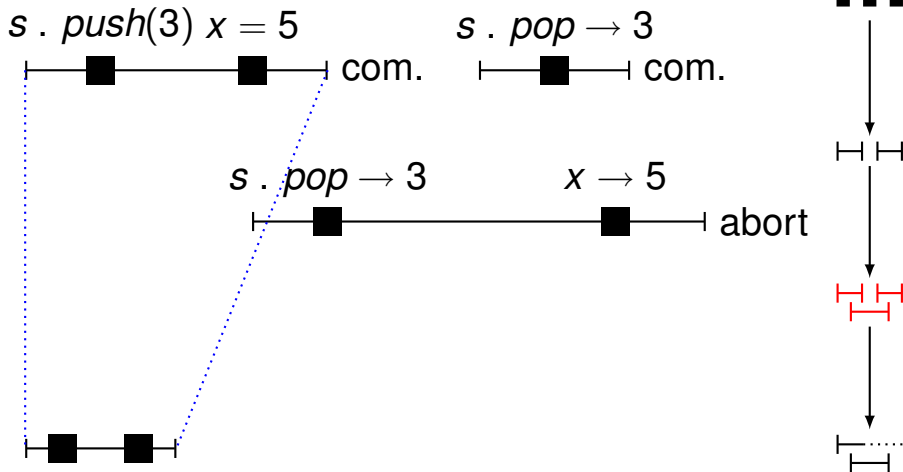
$s . push(3) \quad x = 5$
 com.

$s . pop \rightarrow 3$
 com.

$s . pop \rightarrow 3$
 abort
 $x \rightarrow 5$





Concurrent, Complete Transactions



Concurrent, Complete Transactions




$s . push(3) \quad x = 5$

$s . pop \rightarrow 3$

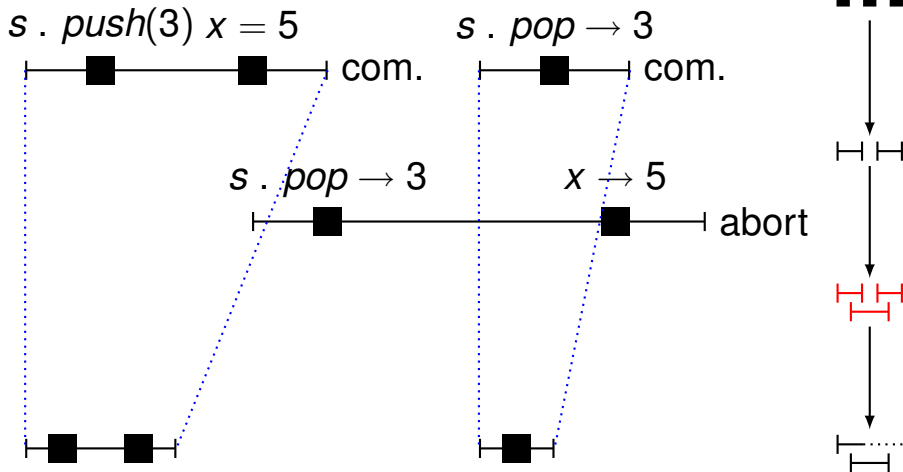


$s . pop \rightarrow 3$

$x \rightarrow 5$




Concurrent, Complete Transactions




Live Transactions

$s . pop \rightarrow 3$




Live Transactions

$s . pop \rightarrow 3$




=

$s . pop \rightarrow 3$



Live Transactions

$s . pop \rightarrow 3$



=

$s . pop \rightarrow 3$



The Definition

A history H ensures **opacity**

if there exists a sequential history S
equivalent to some history in set $Complete(H)$,

such that:

- 1 S preserves the real-time order of H ,
- 2 every transaction $T_i \in S$ is legal in S .

Opacity: Captures TM Semantics

- Committed: instantaneous
- Aborted: never visible
- All: observe consistent state

Opacity: Benefits

- Compare TMs
opacity is implementation-agnostic
- Prove a TM correct
graph interpretation of opacity
- Show inherent limitations of TMs
time complexity lower bound