

On Obstruction-Free Transactions

Rachid Guerraoui Michał Kapałka

EPFL, Switzerland

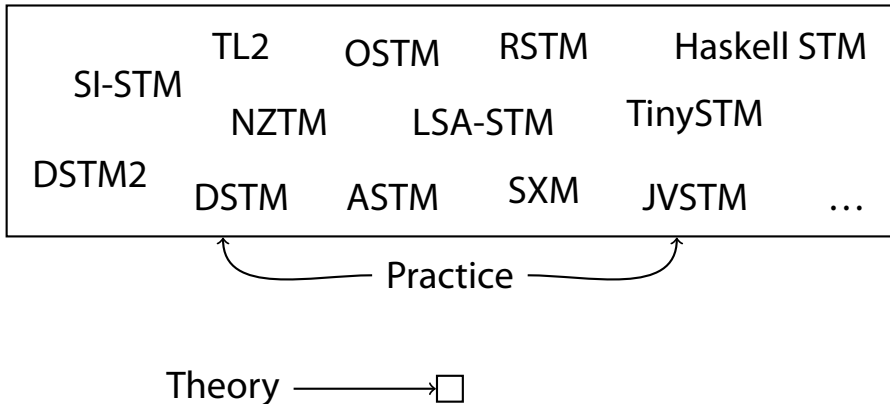
Transactional Memory

```
atomic {  
    accountA . debit(sum)  
    accountB . credit(sum)  
}
```

TM Research


	TL2	OSTM	RSTM	Haskell STM	
SI-STM					
	NZTM	LSA-STM		TinySTM	
DSTM2	DSTM	ASTM	SXM	JVSTM	...

TM Research

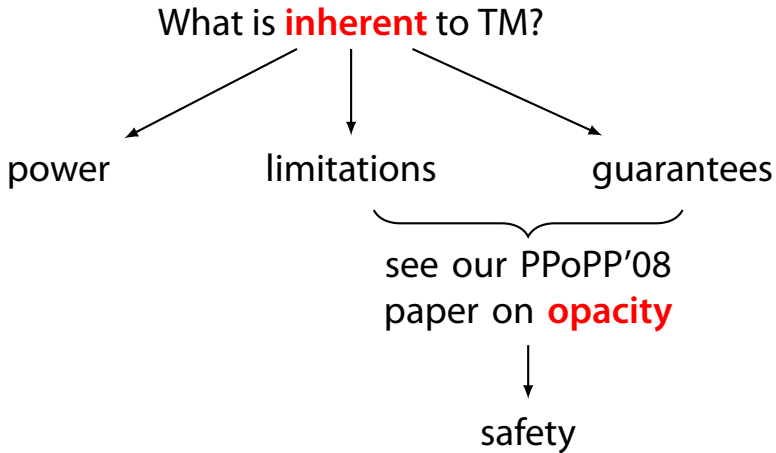


What is inherent to TM?

Fundamental Questions



Fundamental Questions



Focus

Obstruction-Free TM (OFTM)

DSTM, ASTM, SXM, RSTM, NZTM, ...

Why OFTM?

Advantages: \Rightarrow real-time, OS


- No priority inversion
- Fault tolerance
- Can provide strong guarantees

Additional overheads:

- Do not matter in complex workloads
(see our Transact'08 paper)
- Can be reduced (see NZTM)

Contributions

Obstruction-free TM (OFTM)



Power



Consensus
number = 2

Contributions

Obstruction-free TM (OFTM)

Power

Consensus
number = 2


Not universal

Cannot be implemented
from registers

Do not need CAS to be
implemented

Contributions

Obstruction-free TM (OFTM)



Power

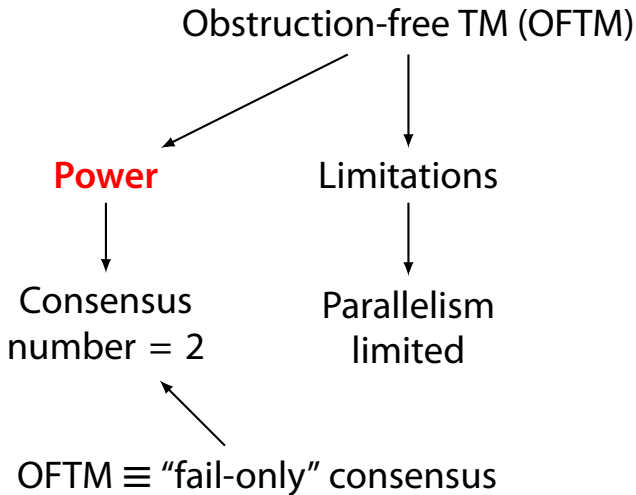


Consensus
number = 2

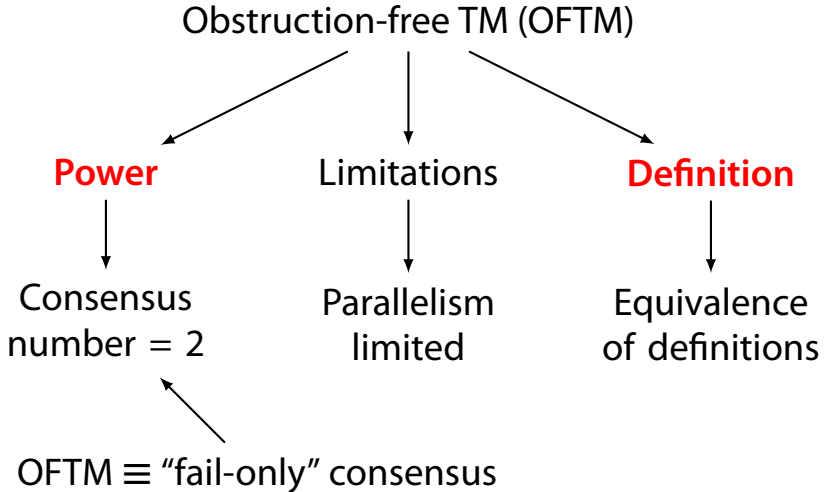


OFTM \equiv "fail-only" consensus

Contributions



Contributions



The Rest of This Talk

- 1 Defining OFTM
- 2 The power of an OFTM


Definition

Intuitive Definition

"A synchronization mechanism is obstruction-free if any thread that runs by itself long enough makes progress (...)"


[Herlihy et al. 03]

Basic Definition




OFTM: if T encounters no **step contention**,
 T cannot be aborted

Other Definitions



ic-OFTM: if T encounters no **interval contention**,
 T cannot be aborted

Other Definitions



eventual if T encounters no **interval contention**,
ic-OFTM: T **eventually** cannot be aborted

Definition Equivalence

In an asynchronous system:

$$\text{OFTM} = \text{ic-OFTM}$$

ic-OFTM equivalent to eventual ic-OFTM

(proof uses “fail-only” consensus)

Power

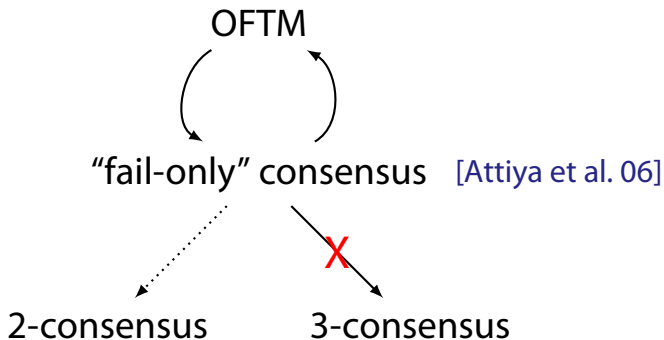
Consensus Number

Object X has consensus number K

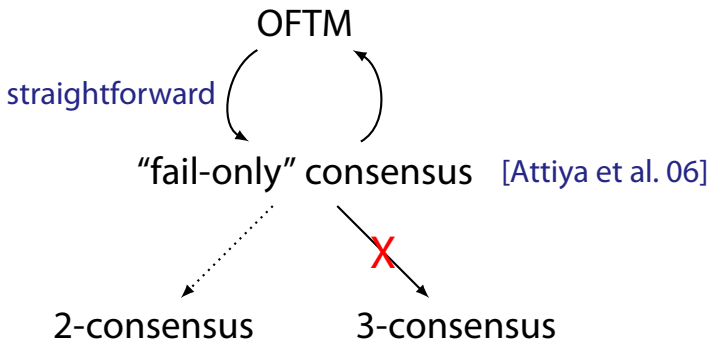
=

One can implement wait-free consensus from X
for at most K processes

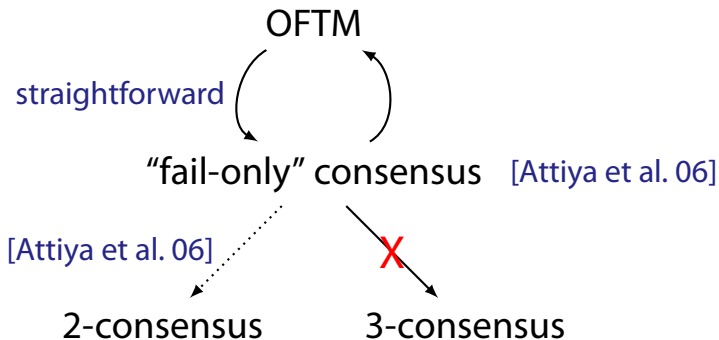
The Power of an OFTM



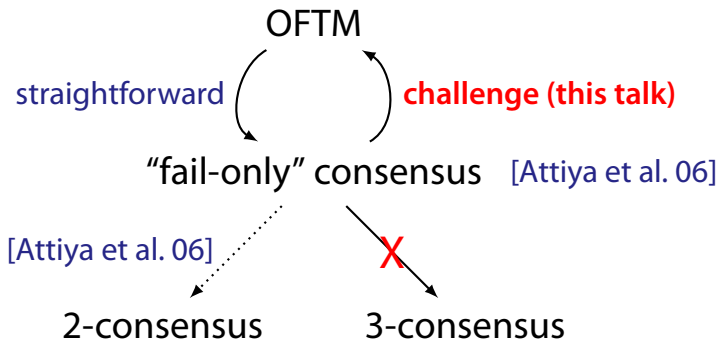
The Power of an OFTM



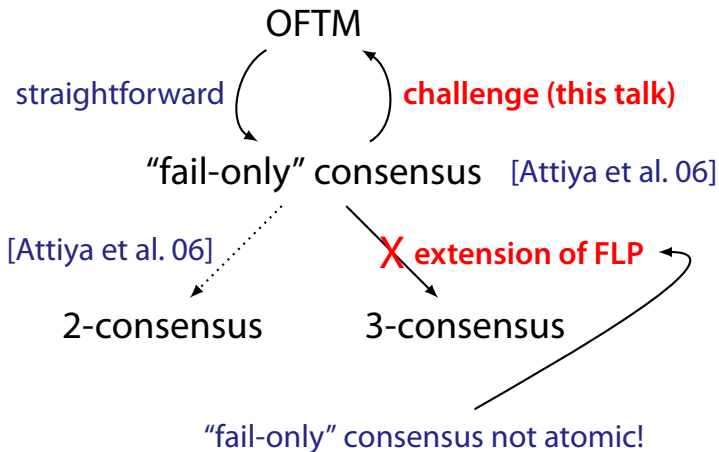
The Power of an OFTM



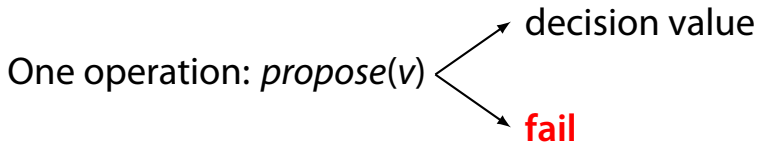
The Power of an OFTM



The Power of an OFTM



“Fail-Only” Consensus




- No two processes decide different values.
- A value decided must be a value proposed by some *propose* **that does not fail**.
- **Fail** – only on step contention.


“Fail-only” consensus \rightarrow OFTM

Use of CAS

1. Object acquisition



2. Committing/aborting a transaction



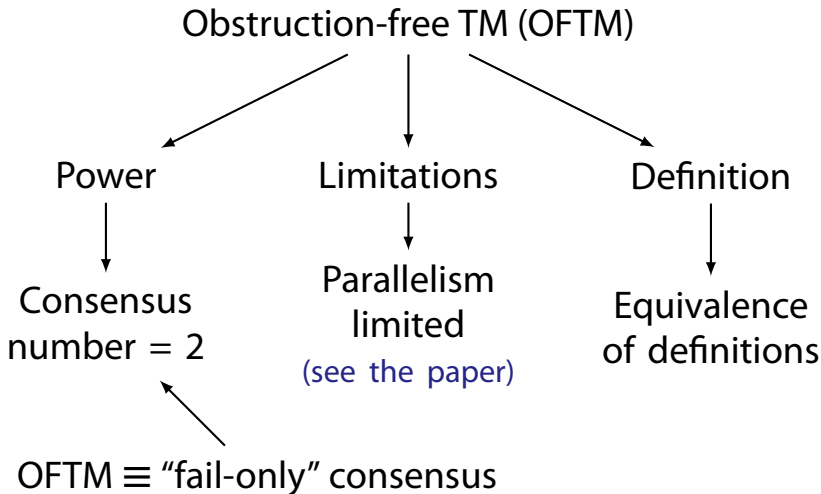
Challenge

“Fail-only” consensus is:

One-shot

Not readable

Summary



What is inherent to TM?