

Liveness of Transactional Memory

Victor Bushkov

Distributed Programming Laboratory



Part I

Defining transactional memory liveness

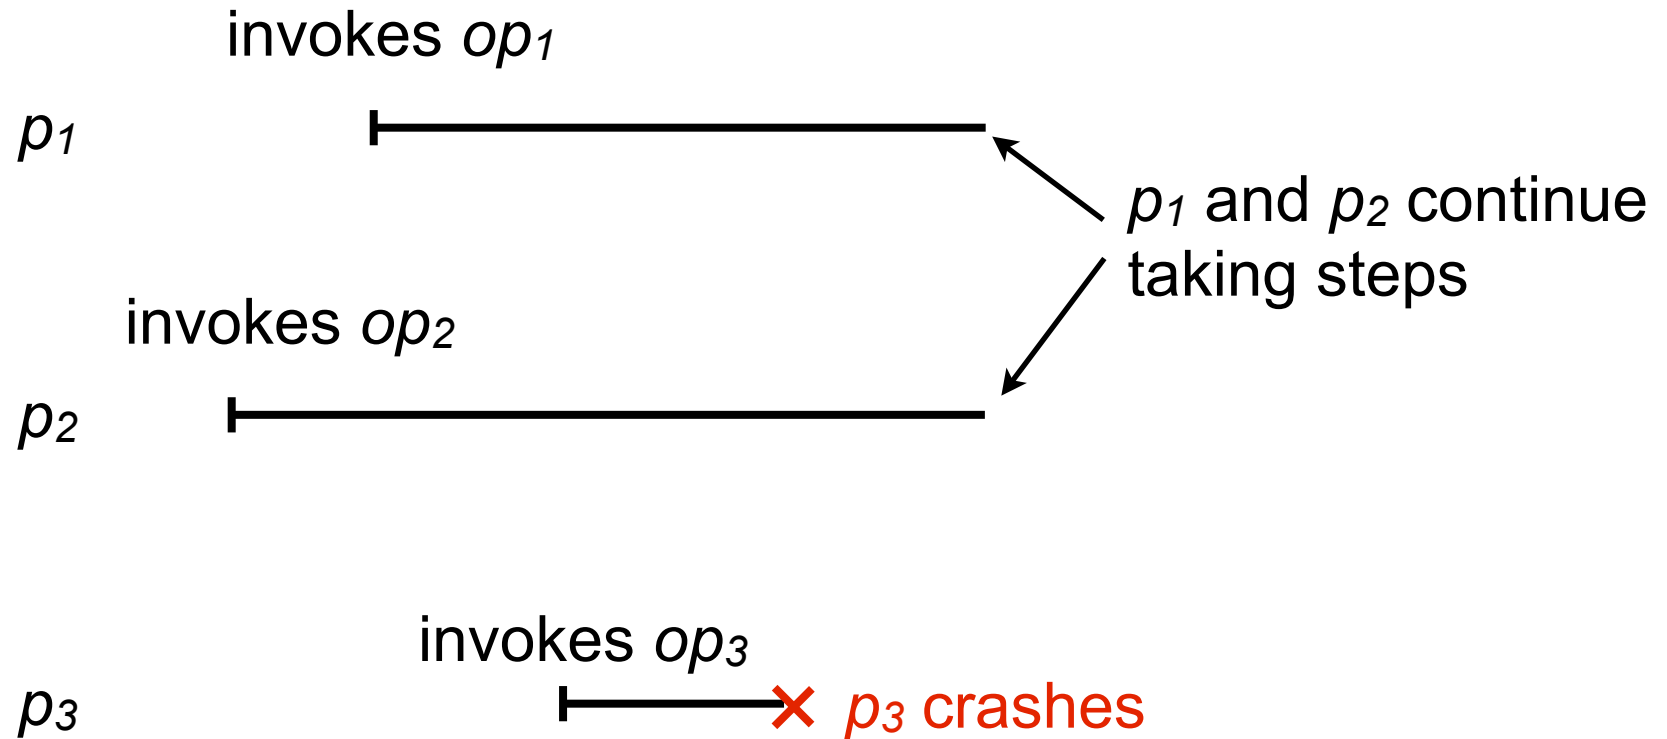
Properties covered so far

- wait-freedom
- lock-freedom
- obstruction-freedom

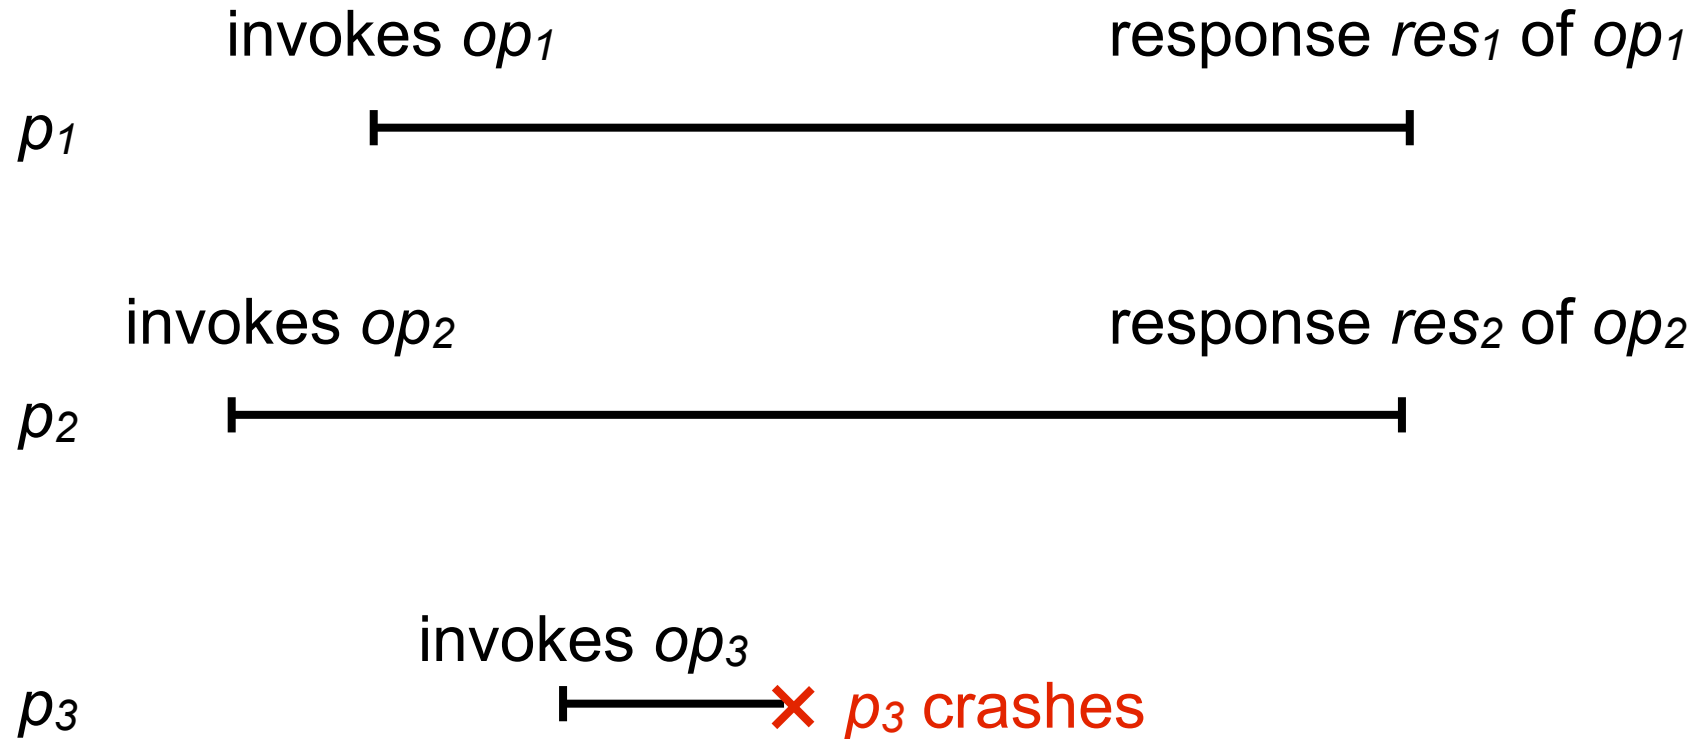
Wait-freedom

Every operation by *every* non-crashed process eventually returns a response

Wait-freedom: example



Wait-freedom: example

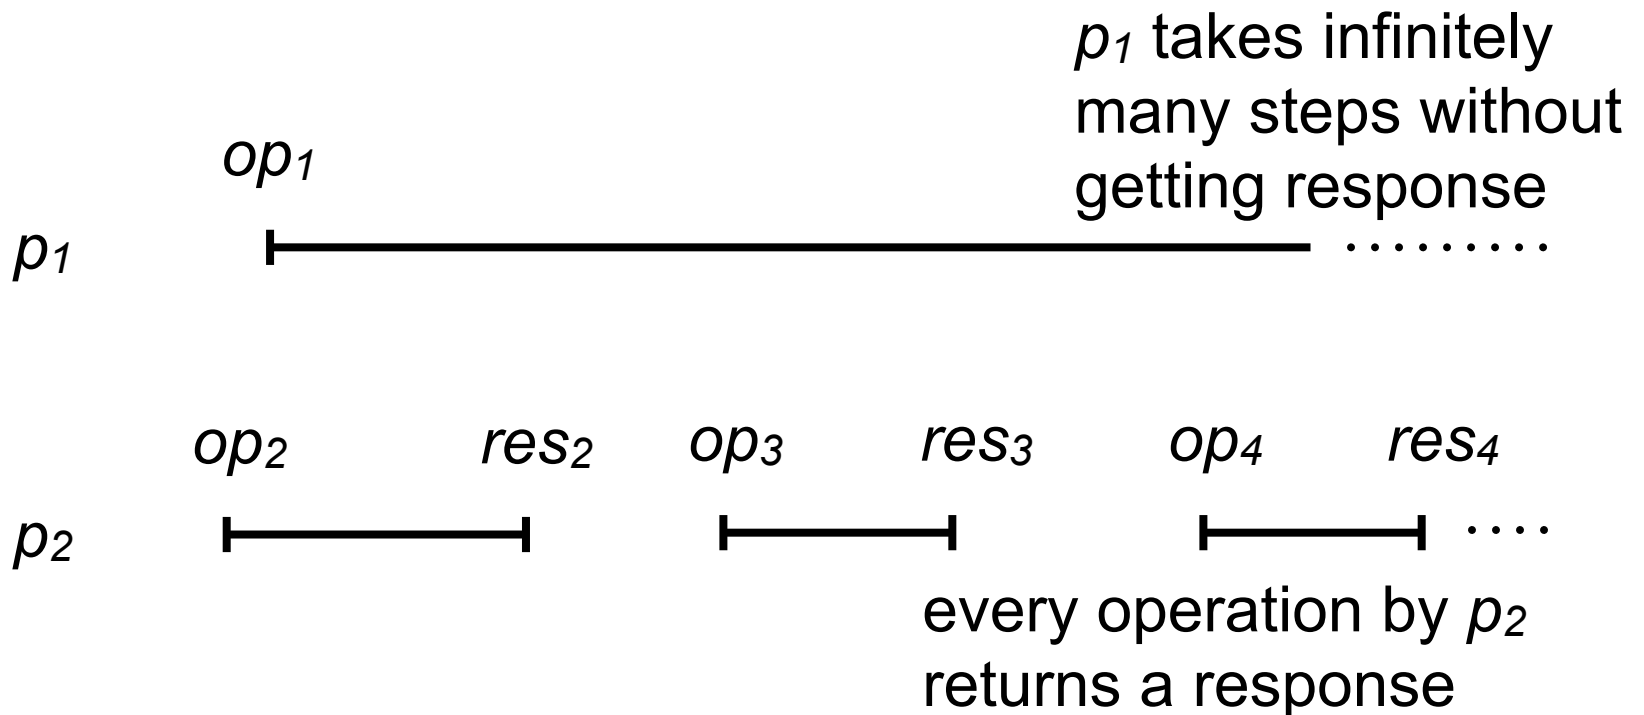


Lock-freedom

Every operation by *some* non-crashed process eventually returns a response

Lock-freedom: example

- execution is **not** wait-free
- but it is lock-free




Obstruction-freedom

If a process p becomes the only process taking steps, then every operation by p eventually returns a response

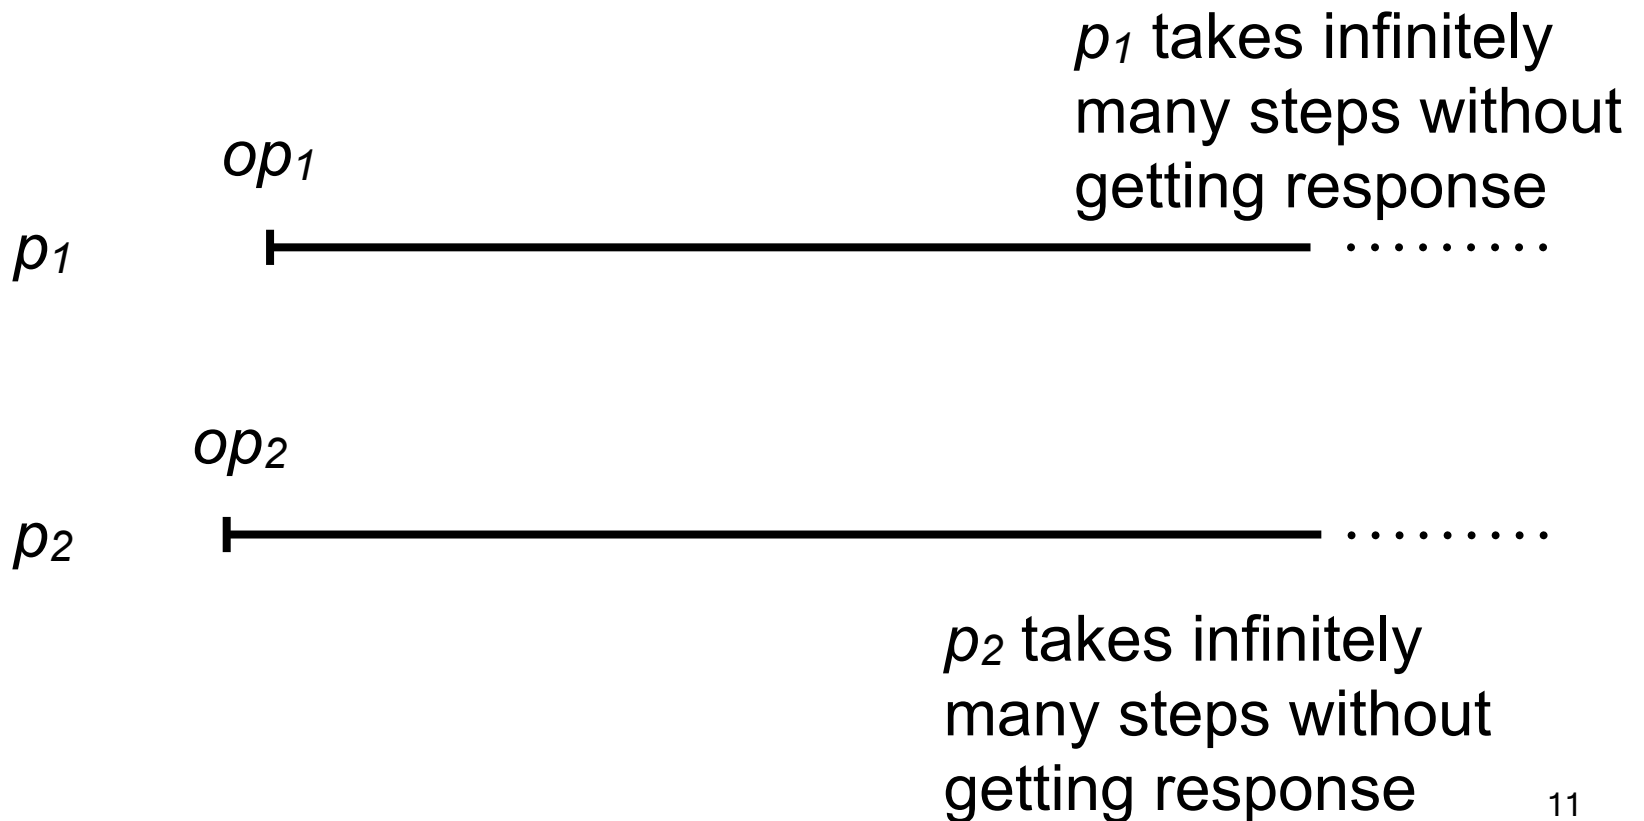
Obstruction-freedom: example

- execution is lock-free
- and it is obstruction-free



Obstruction-freedom: example

- execution is **not** lock-free
- but it is obstruction-free




What is common between these
three properties?


What is common between these three properties?

- state that some good event must *eventually* happen
- i.e. they are liveness properties

Liveness vs Safety




Liveness vs Safety



- wait-freedom (termination)
- lock-freedom
- obstruction-freedom

Liveness vs Safety



- wait-freedom (termination)
- lock-freedom
- obstruction-freedom

- validity and agreement
- regularity of registers
- atomicity (linearizability)
- opacity

Liveness vs Safety

Liveness: some **good** events should *eventually* happen

Safety: some **bad** events should *never* happen

Liveness vs Safety

Liveness: some **good** events should *eventually* happen

Safety: some **bad** events should *never* happen

- violated in finite execution

Liveness vs Safety

Liveness: some **good** events should *eventually* happen

- cannot be violated in a finite execution

Safety: some **bad** events should *never* happen

- violated in finite execution

Liveness of shared objects

- In shared objects **good** events are responses

Liveness of shared objects

- In shared objects **good** events are responses
- In case of wait-freedom, lock-freedom, and obstruction-freedom **any** response is a good event i.e.:

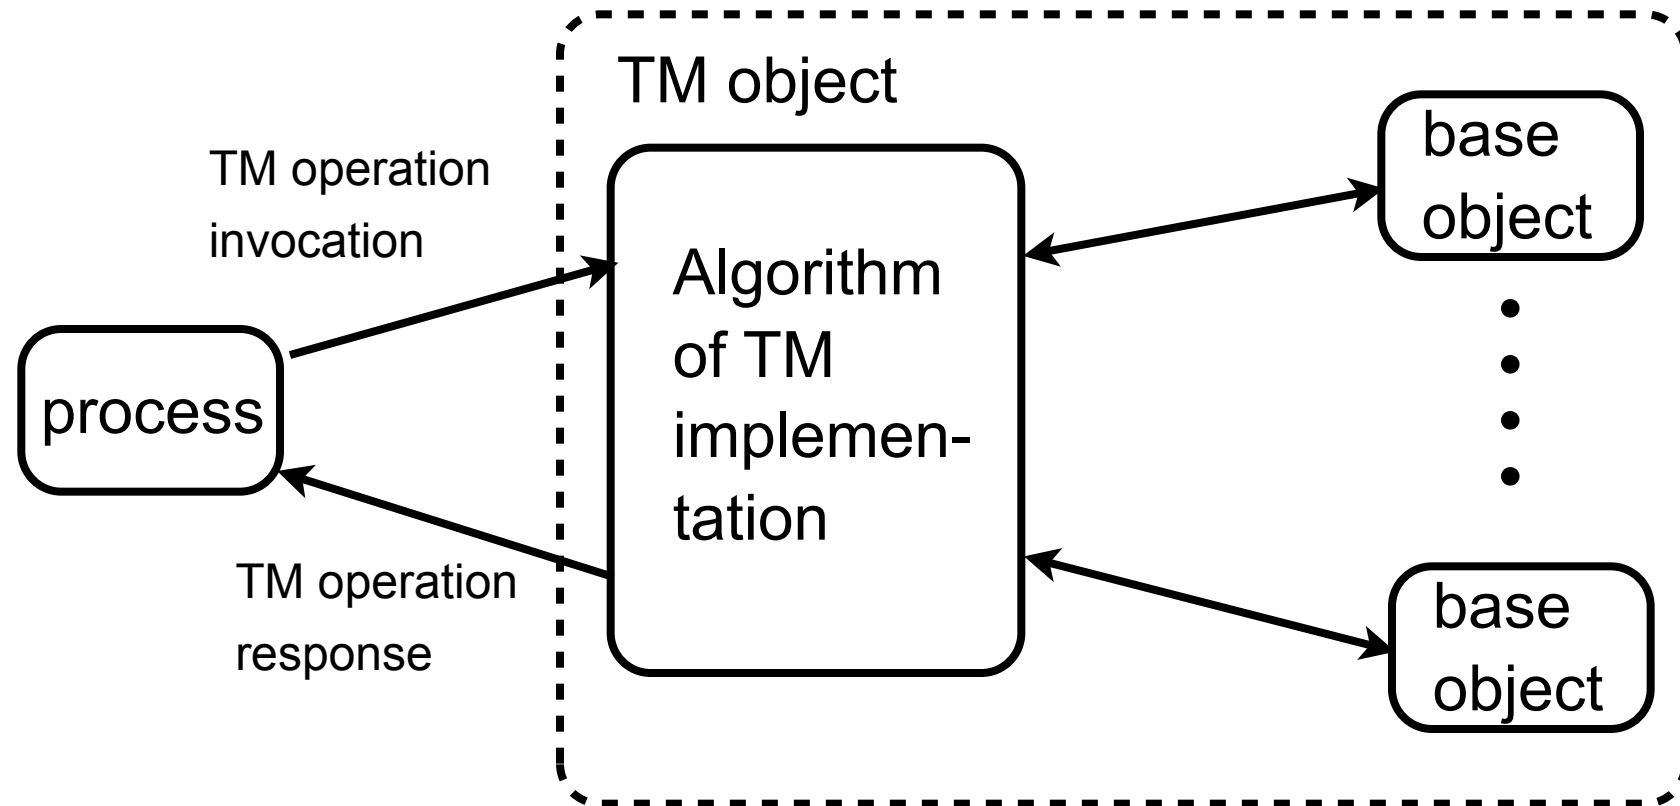
Liveness of shared objects

- In shared objects **good** events are responses
- In case of wait-freedom, lock-freedom, and obstruction-freedom **any** response is a good event i.e.:

e.g. in case of wait-freedom we do not care if we get res_1 or some other response res'_1



Transactional memory (TM) as a shared objects



Transactional memory (TM) as a shared objects

examples of some TM operations

- *x.read()* - returns value of data item *x*
- *x.write(v)* - writes value *v* to data item *x*
- *commit()* - commits current transaction
- *begin_tr()* - starts a transaction

Transactional memory (TM) as a shared objects


examples of some TM operations

- *x.read()* - returns value of data item *x*
- *x.write(v)* - writes value *v* to data item *x*
- *commit()* - commits current transaction
- *begin_tr()* - starts a transaction

- every TM operation can return abort event *A* which aborts current transaction

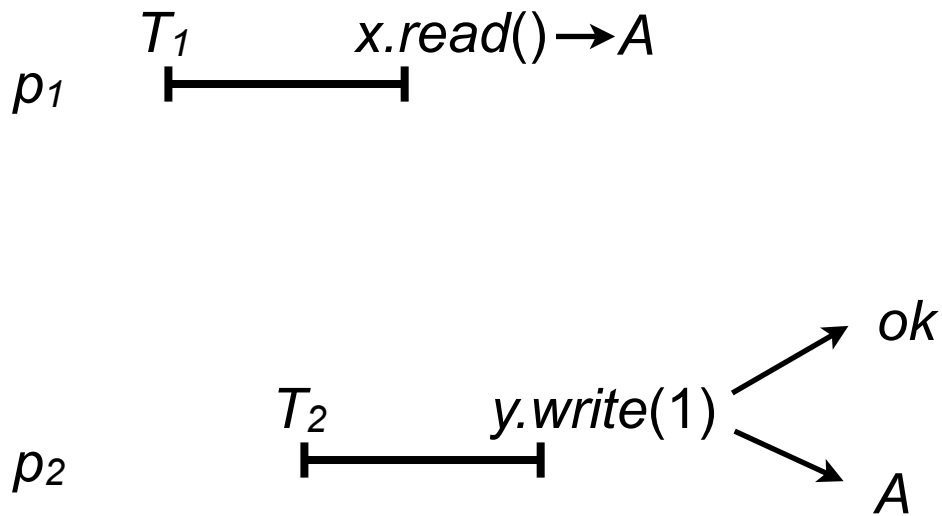
Is wait-freedom enough in TM
context?

Is wait-freedom enough in TM context?




p_2

Is wait-freedom enough in TM context?



Is wait-freedom enough in TM context?

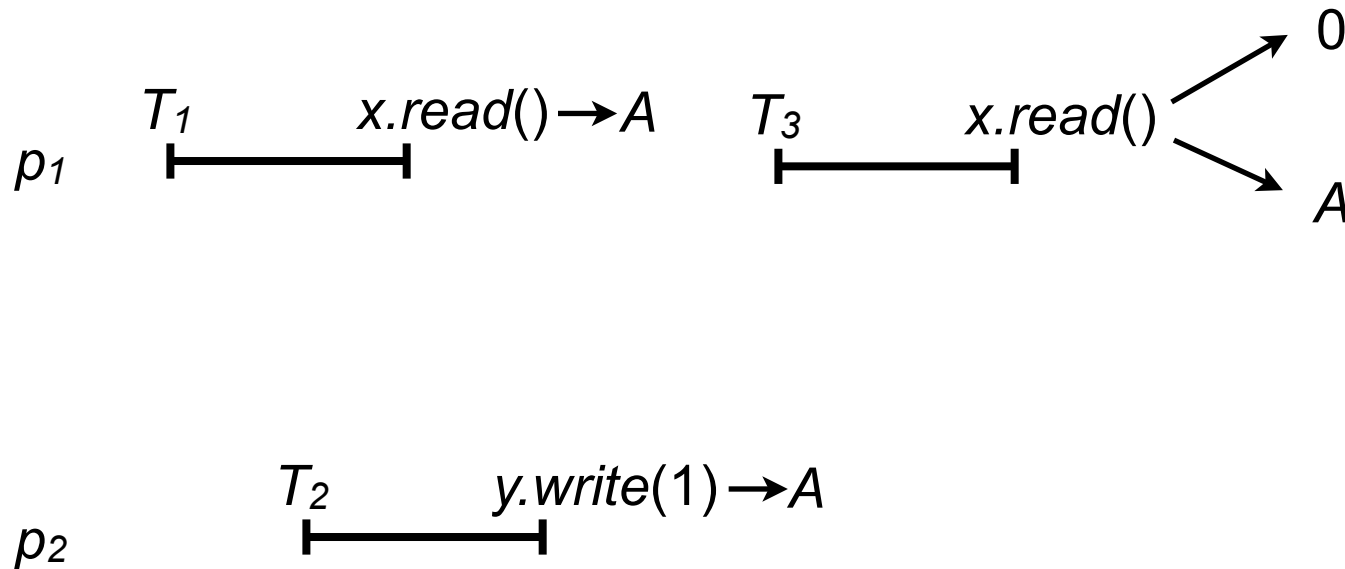
p_1 T_1 $x.read() \rightarrow A$

A horizontal line with vertical end caps represents the execution of process p1. The line is labeled T1 above it and x.read() -> A to its right.

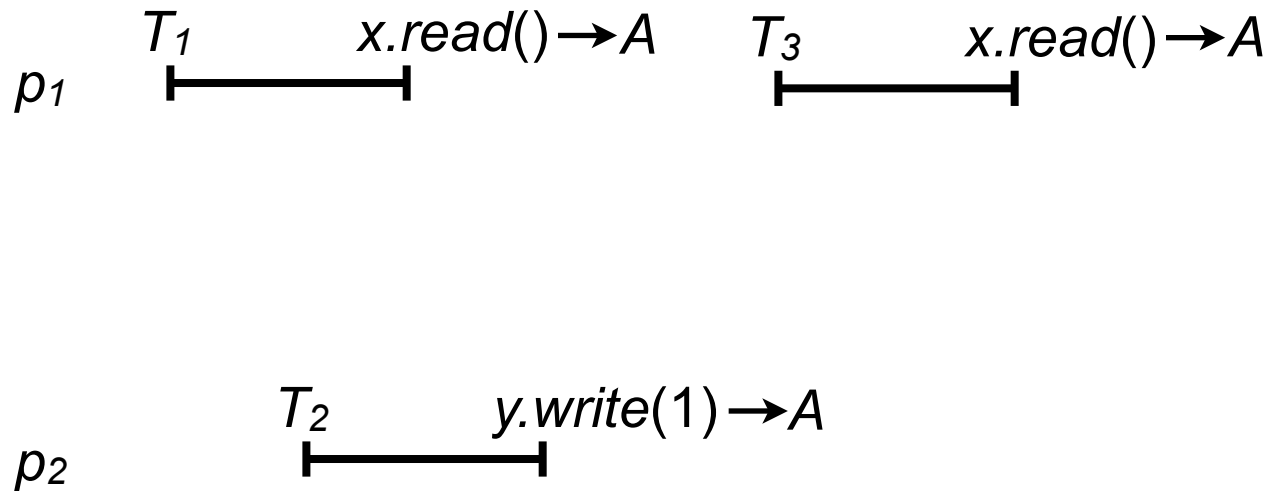
p_2 T_2 $y.write(1) \rightarrow A$

A horizontal line with vertical end caps represents the execution of process p2. The line is labeled T2 above it and y.write(1) -> A to its right.

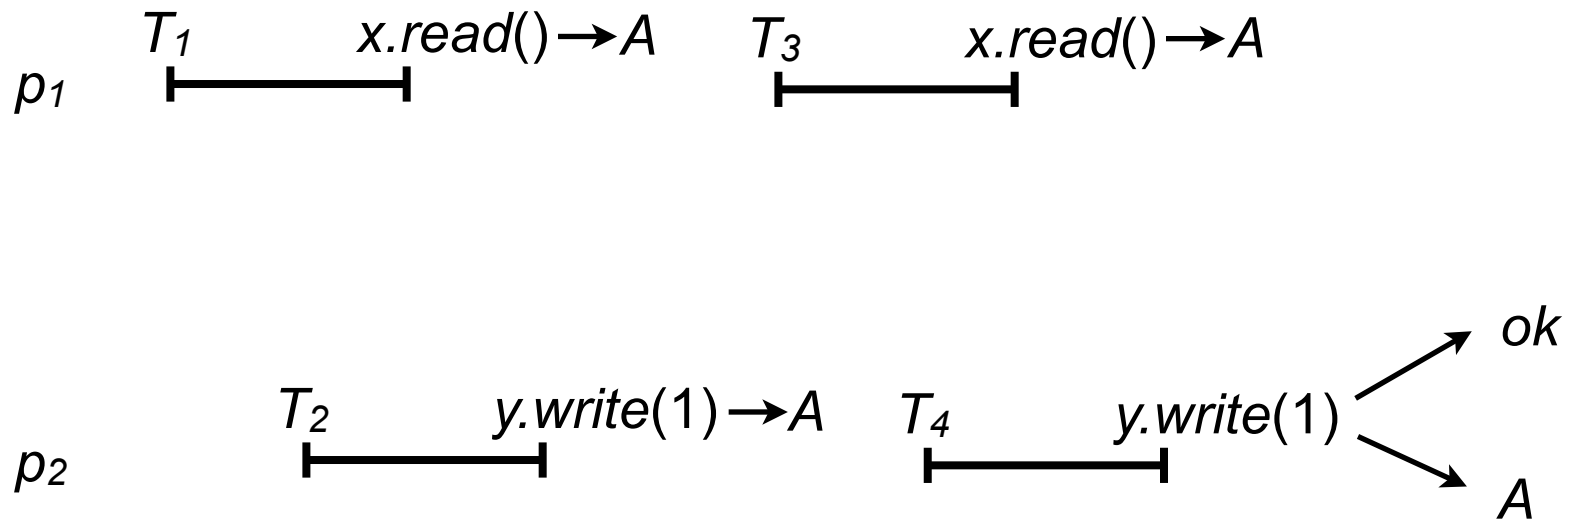
Is wait-freedom enough in TM context?



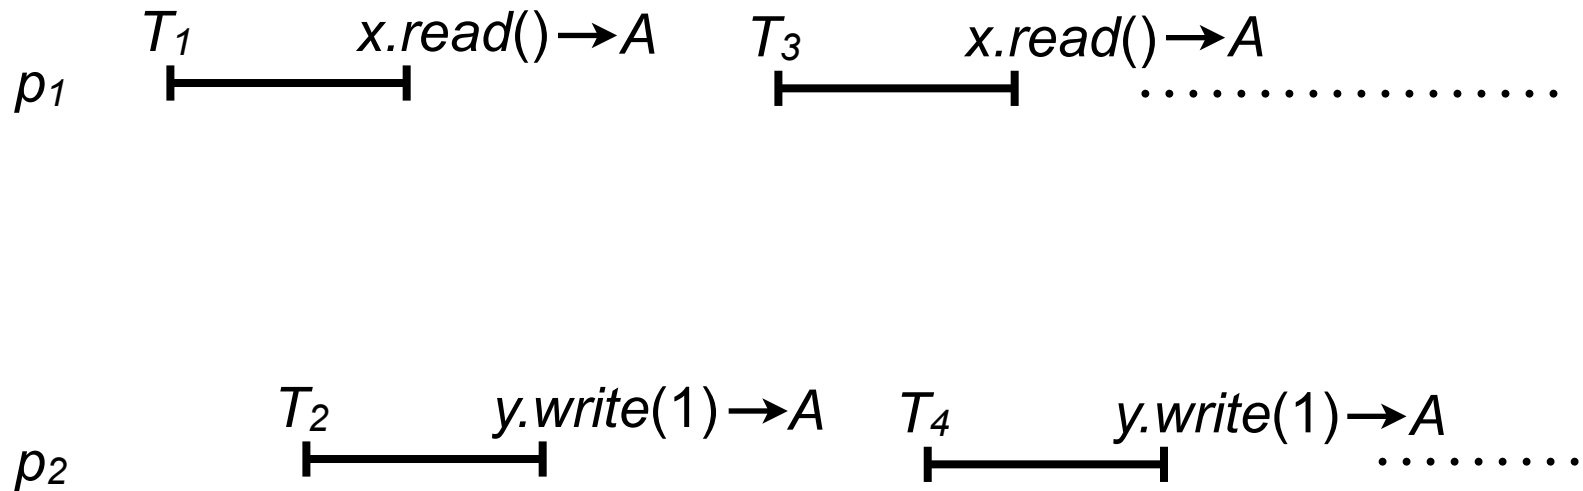
Is wait-freedom enough in TM context?



Is wait-freedom enough in TM context?



Is wait-freedom enough in TM context?



Meaningful progress

- wait-freedom is trivially ensured by aborting every TM operation

Meaningful progress

- wait-freedom is trivially ensured by aborting every TM operation
- operation termination is not enough

Meaningful progress


- wait-freedom is trivially ensured by aborting every TM operation
- operation termination is not enough
- operations need to receive meaningful responses

What about the following property?

- Every TM operation by every non-crashed process eventually returns a response **which is not an abort event**


What about the following property?

- Every TM operation by every non-crashed process eventually returns a response **which is not an abort event**
- It can be violated in a **finite** execution \rightarrow it is not liveness



What about the following property?

- Every TM operation by every non-crashed process eventually returns a response **which is not an abort event**
- It can be violated in a **finite** execution \rightarrow it is not liveness
- TM loses its meaning without ability to abort (TM becomes equivalent to universal construction)



Meaningful progress

TM liveness property should

- allow every transaction to be aborted, and

Meaningful progress

TM liveness property should

- allow every transaction to be aborted, and
- require processes to **eventually** commit some transaction
(**make progress**)


What does eventually committing some transactions mean?

- a process might have some of its transactions aborted

p_1 T_1 A T_2 A T_3 A
┌····┐ ┌····┐ ┌····┐

What does eventually committing some transactions mean?

- a process might have some of its transactions aborted
- but for any point in time of the execution eventually there is a transaction that commits



Eventually there is a transaction that commits

Can we require eventual commitment of *any* process?

```
begin_tr()  
  while(value = i) do {  
    value := x.read( );  
    x.write(value + 1);  
    i := i+1;  
  }  
commit()
```

Initially:
value, *i* = -1
x = 0

Can we require eventual commitment of *any* process?

begin_tr()


```
while(value = i) do {  
  value := x.read( );  
  x.write(value + 1);  
  i := i+1;  
}
```

commit()

Initially:

value, *i* = -1

x = 0



p_1 repeatedly reads and writes x without ever invoking a commit request

Correct processes

We cannot require progress of processes which are **not correct** in a given infinite execution α :

- processes which crash in α , or

Correct processes

We cannot require progress of processes which are **not correct** in a given infinite execution α :

- processes which crash in α , or
- processes which execute a transaction which is not aborted and does not invoke a commit request in α

Correct processes

We cannot require progress of processes which are **not correct** in a given infinite execution α :

- processes which crash in α , or
- processes which execute a transaction **which is not aborted** and does not invoke a commit request **in α**

Correct processes

begin_tr()


```
while(value = i) do {  
  value := x.read( );  
  x.write(value + 1);  
  i := i+1;  
}
```

commit()

Initially:

value, i = -1

x = 0



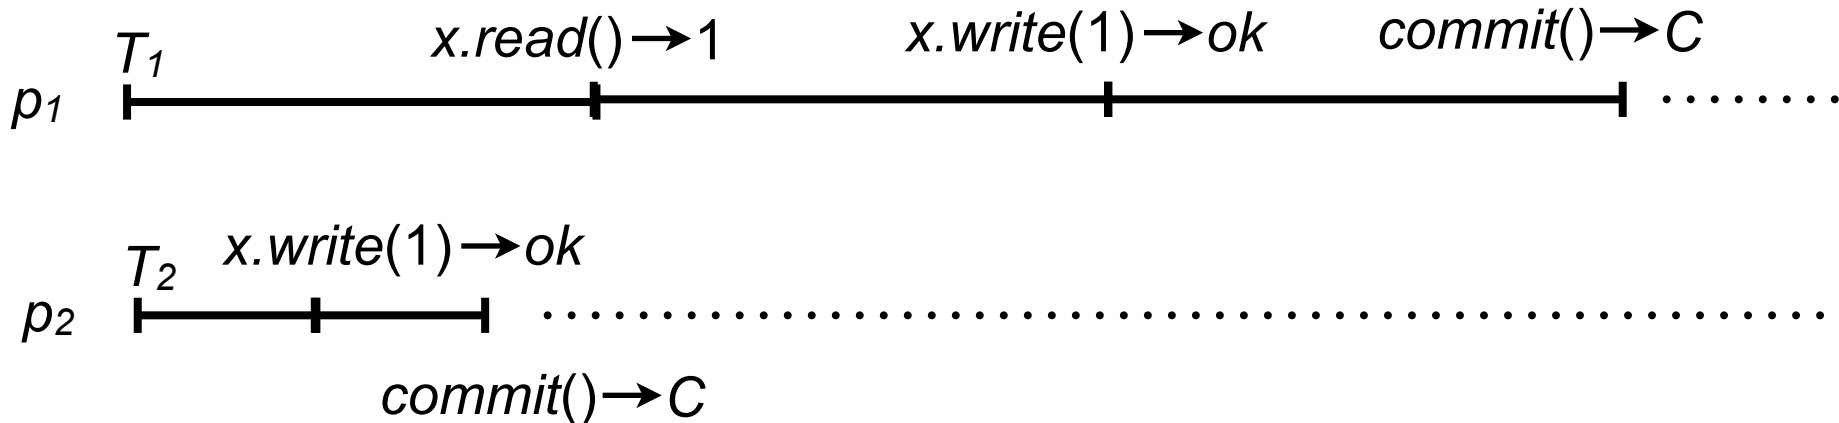
p_1 is not correct in the given execution

Correct processes


```
begin_tr()
  while(value = i) do {
    value := x.read( );
    x.write(value + 1);
    i := i+1;
  }
commit()
```

Initially:
 $value, i = -1$
 $x = 0$

p_1 is correct in the given execution




Correct processes



- p_1 is correct in the given execution

Correct processes



- p_1 is correct in the given execution
- the notion of a correct process depends on an execution

Correct processes

begin_tr()

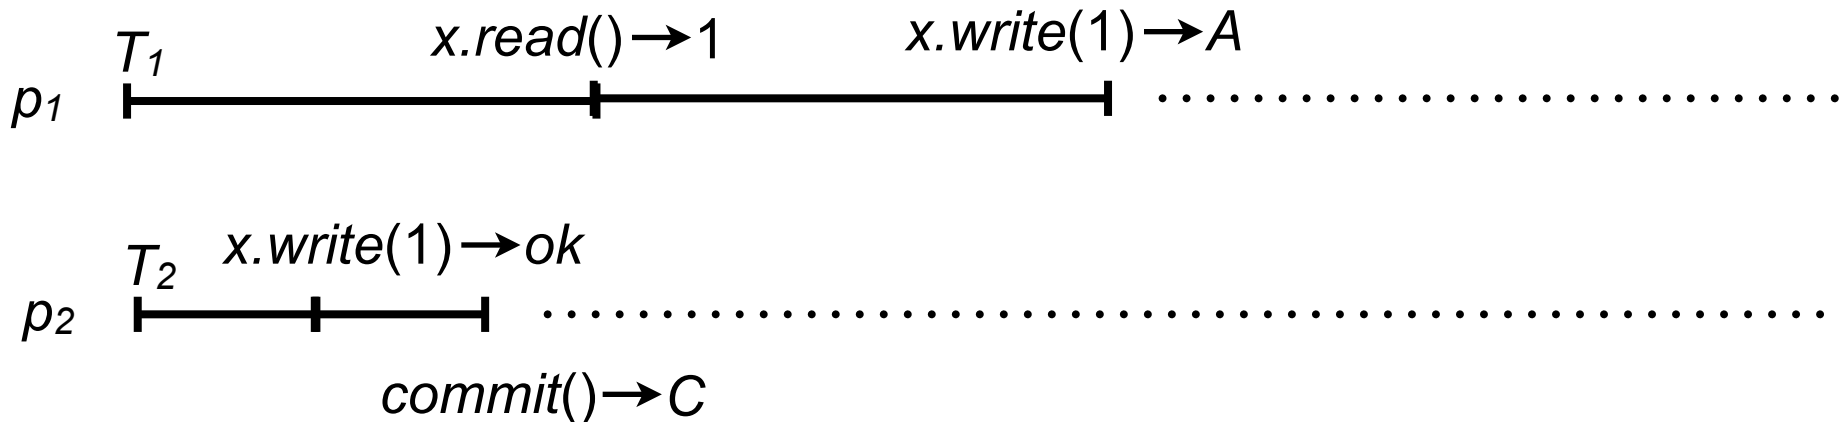
```
while(value = i) do {  
    value := x.read( );  
    x.write(value + 1);  
    i := i+1;  
}
```

commit()


Initially:

value, *i* = -1

x = 0




Correct processes



- p_1 is correct in the given execution

Correct processes



- p_1 is correct in the given execution
- a process which is never given possibility to invoke a commit request is still considered correct

Correct processes



- p_1 is correct in the given execution
- a process which is never given possibility to invoke a commit request is still considered correct


Making progress (in TM context)

A correct process p **makes progress** in an infinite execution α if infinitely many transaction of p commit in α

Making progress (in TM context)

A correct process p **makes progress** in an infinite execution α if infinitely many transaction of p commit in α


- a process might have some of its transactions aborted



Making progress (in TM context)

A correct process p **makes progress** in an infinite execution α if infinitely many transaction of p commit in α

- a process might have some of its transactions aborted
- but for any point in time of the execution eventually there is a transaction that does not abort (and consequently commits)




TM liveness

An infinite execution α is **TM-wait-free** if *every* correct process makes progress in α

TM liveness

An infinite execution α is **TM-wait-free** if every correct process makes progress in α




TM liveness

An infinite execution α is **TM-lock-free** if *some* correct process makes progress in α

TM liveness

An infinite execution α is **TM-lock-free** if *some* correct process makes progress in α




TM liveness

An infinite execution α is **TM-obstruction-free** if for every correct process p in α the following holds: if eventually p becomes the only process taking steps, then p makes progress in α

TM liveness

An infinite execution α is **TM-obstruction-free** if for every correct process p in α the following holds: if eventually p becomes the only process taking steps, then p makes progress in α



Liveness: take home

When arguing about liveness of a shared object implementation, things to keep in mind:

Liveness: take home

When arguing about liveness of a shared object implementation, things to keep in mind:

- depending on the context liveness properties might be defined different ways

Liveness: take home

When arguing about liveness of a shared object implementation, things to keep in mind:

- depending on the context liveness properties might be defined different ways
- specification might include several different kinds of liveness properties (e.g. TM-obstruction-freedom for transactions + wait-freedom for individual TM operations)

Liveness: take home

When arguing about liveness of a shared object implementation, things to keep in mind:

- depending on the context liveness properties might be defined different ways
- specification might include several different kinds of liveness properties (e.g. TM-obstruction-freedom for transactions + wait-freedom for individual TM operations)
- be accurate when specifying which processes should make progress

Part II

The impossibility of TM-wait-freedom

Wait-freedom

- Wait-freedom forms the basis of consensus number hierarchy

Wait-freedom

- Wait-freedom forms the basis of consensus number hierarchy
- In most cases we need to use powerful base objects (like consensus, CAS) to implement wait-freedom

Wait-freedom

- Wait-freedom forms the basis of consensus number hierarchy
- In most cases we need to use powerful base objects (like consensus, CAS) to implement wait-freedom
- Not the case for TM-wait-freedom:
 - it cannot be implemented together with opacity irrespectively of the power of base objects being used

Impossibility

Theorem

- There is no TM implementation that:
 - ensures TM-wait-freedom and

Impossibility

Theorem

- There is no TM implementation that:
 - ensures TM-wait-freedom and
 - opacity

Impossibility

Theorem

- There is no TM implementation that:
 - ensures TM-wait-freedom and
 - opacity
 - in an asynchronous system

Proof

To prove the result

- We use processes and a scheduler as an *adversary*

Proof

To prove the result

- We use processes and a scheduler as an *adversary*
- The *adversary* forces any TM implementation to produce an execution that violates TM-wait-freedom

Proof: processes

- consider a system of two processes p_1 and p_2

Proof: processes


- consider a system of two processes p_1 and p_2
- processes keep executing infinitely many transactions with the following code

begin_tr()

```
value := x.read( );  
x.write(value + 1);
```


commit()

Proof: execution




p_2

Proof: execution



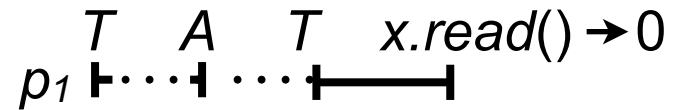
p_2

Proof: execution



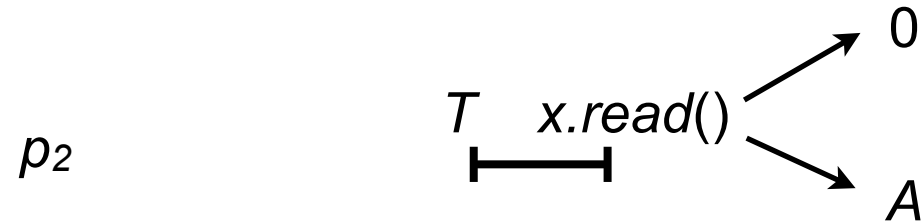
Proof: execution

p_1 T A T $x.read() \rightarrow 0$



The diagram for process p_1 shows a sequence of operations: T , A , T , and $x.read() \rightarrow 0$. Each operation is represented by a horizontal line with vertical end caps. Ellipses between the first three operations indicate that they are interleaved with other operations. An arrow points from the end of the $x.read()$ operation to the value 0.

p_2 T $x.read()$



The diagram for process p_2 shows two operations: T and $x.read()$. Each is represented by a horizontal line with vertical end caps. Two arrows originate from the end of the $x.read()$ operation, pointing to the values 0 and A.

Proof: execution

p_1 $\overbrace{\quad}^T \cdots \overbrace{\quad}^A \cdots \overbrace{\quad}^T \quad \overbrace{\quad}^{x.read() \rightarrow 0}$

p_2 $\quad \quad \quad \overbrace{\quad}^T \quad \overbrace{\quad}^{x.read() \rightarrow A}$

Proof: execution

p_1 T A T $x.read() \rightarrow 0$
|-----|-----|-----|

p_2 T $x.read() \rightarrow A$ T $x.read()$ $\begin{matrix} \nearrow 0 \\ \searrow A \end{matrix}$
|-----|-----|-----|

Proof: execution

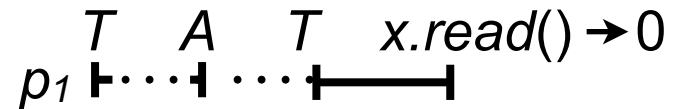
p_1 T A T $x.read() \rightarrow 0$
|-----|-----|-----|

p_2 T A T $x.read() \rightarrow 0$
|-----|-----|-----|

by TM-wait-freedom

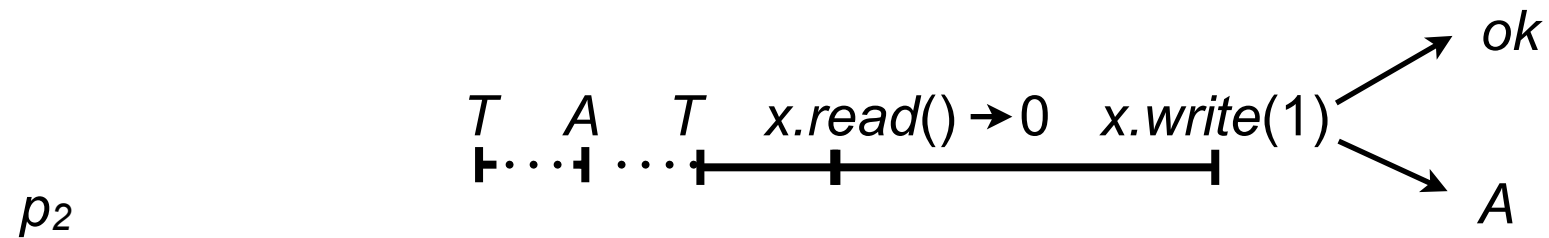
Proof: execution

p_1 T A T $x.read() \rightarrow 0$





The diagram shows a horizontal timeline for process p_1 . It starts with a tick mark, followed by an ellipsis, another tick mark, a second ellipsis, a third tick mark, and a horizontal line segment ending at a fourth tick mark. Above the first, second, and third tick marks are the letters T , A , and T respectively. To the right of the horizontal line segment is the text $x.read() \rightarrow 0$.

p_2 T A T $x.read() \rightarrow 0$ $x.write(1)$ ok A




The diagram shows a horizontal timeline for process p_2 . It starts with a tick mark, followed by an ellipsis, another tick mark, a second ellipsis, a third tick mark, a horizontal line segment ending at a fourth tick mark, a fifth tick mark, and a horizontal line segment ending at a sixth tick mark. Above the first, second, and third tick marks are the letters T , A , and T respectively. To the right of the fourth tick mark is the text $x.read() \rightarrow 0$. To the right of the sixth tick mark is the text $x.write(1)$. From the $x.write(1)$ text, two arrows branch out: one pointing up and to the right to the text ok , and one pointing down and to the right to the text A .

Proof: execution

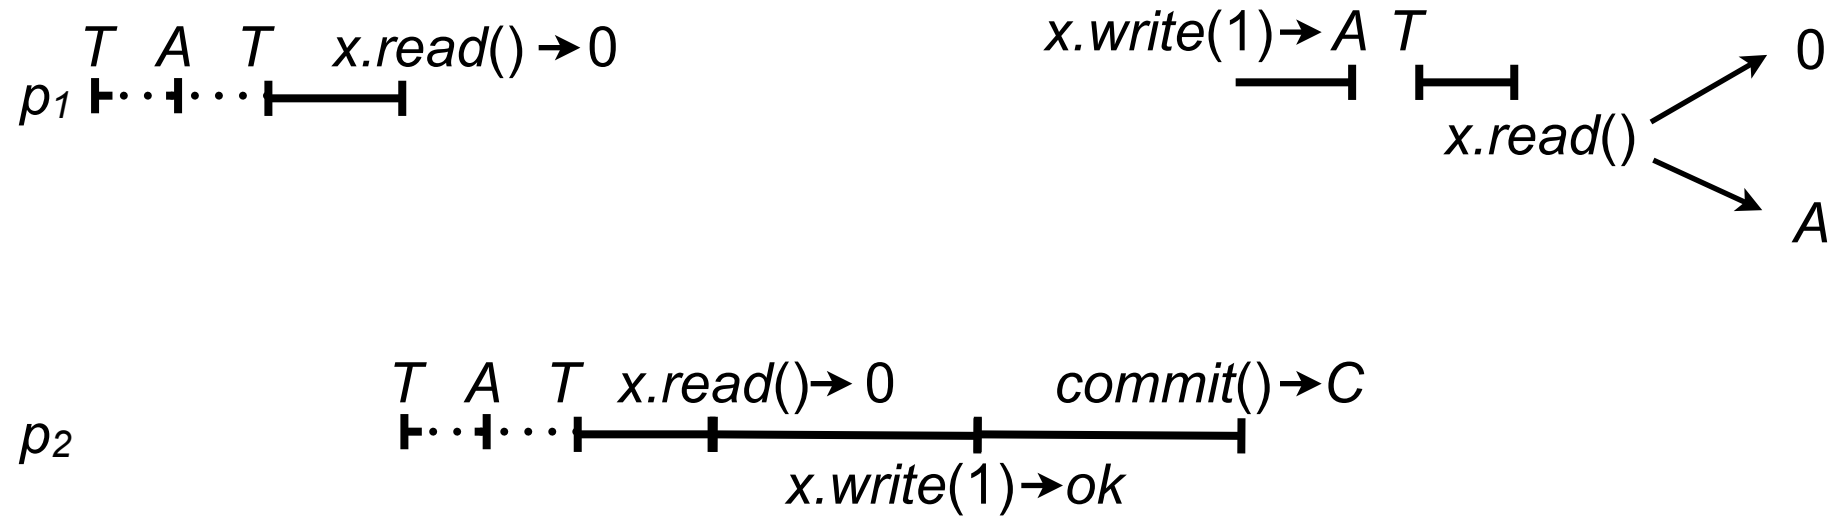


p_2 repeats executing the transaction until eventually the transaction is committed (by TM-wait-freedom)

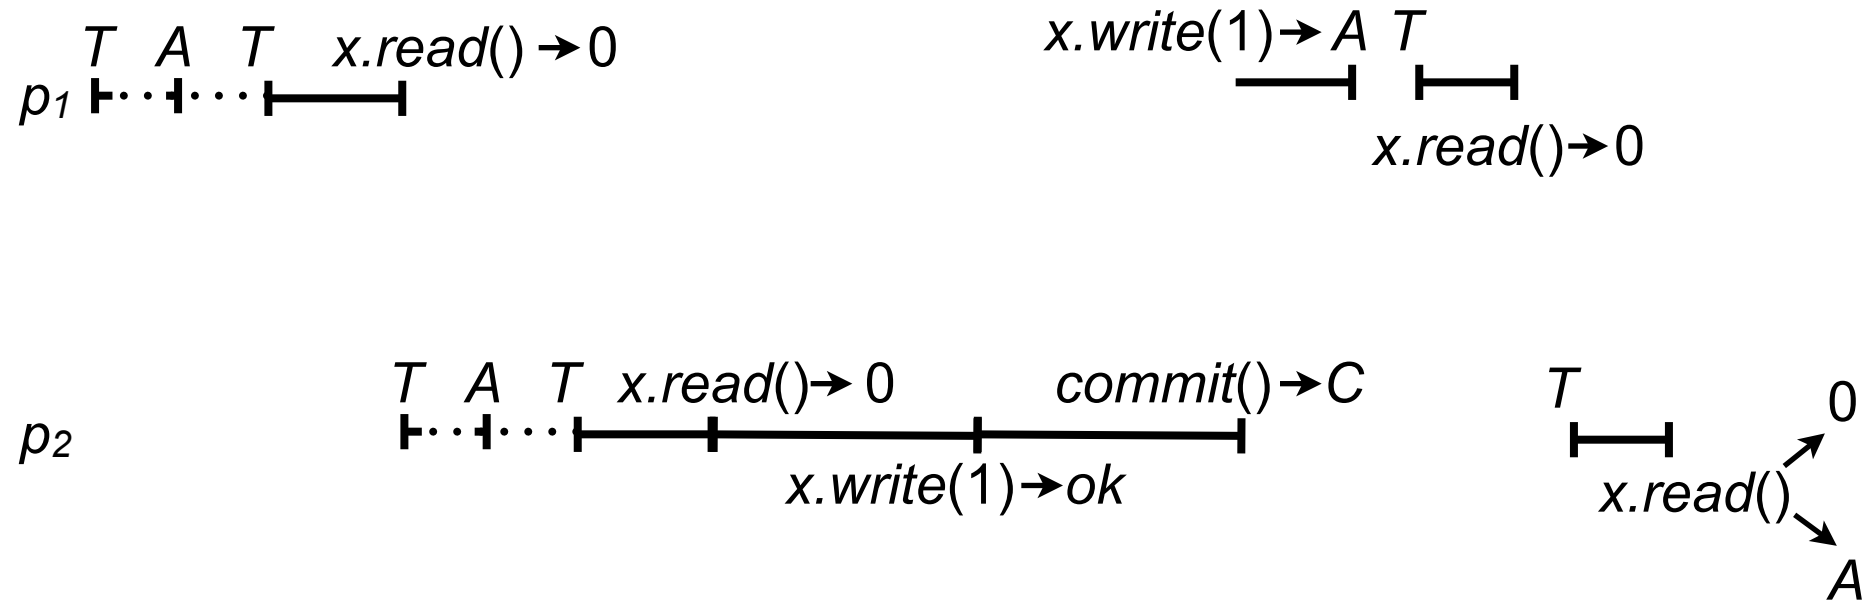
Proof: execution



Proof: execution

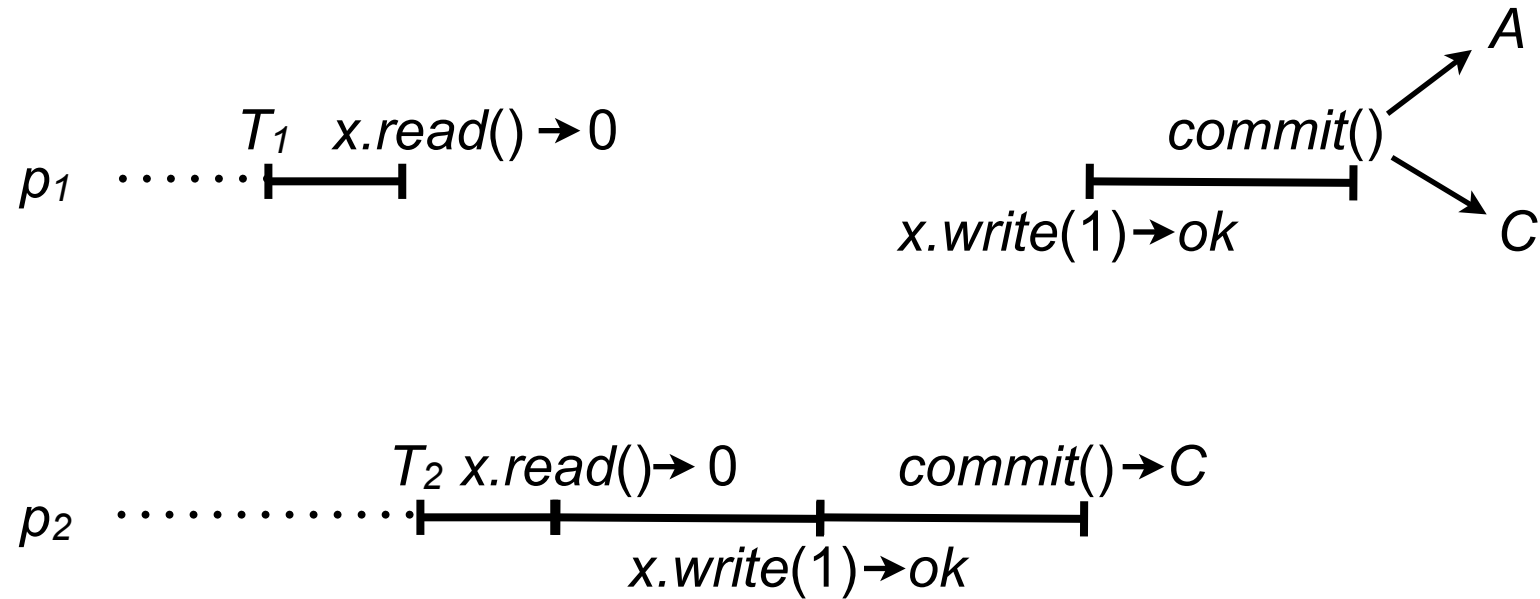


Proof: execution

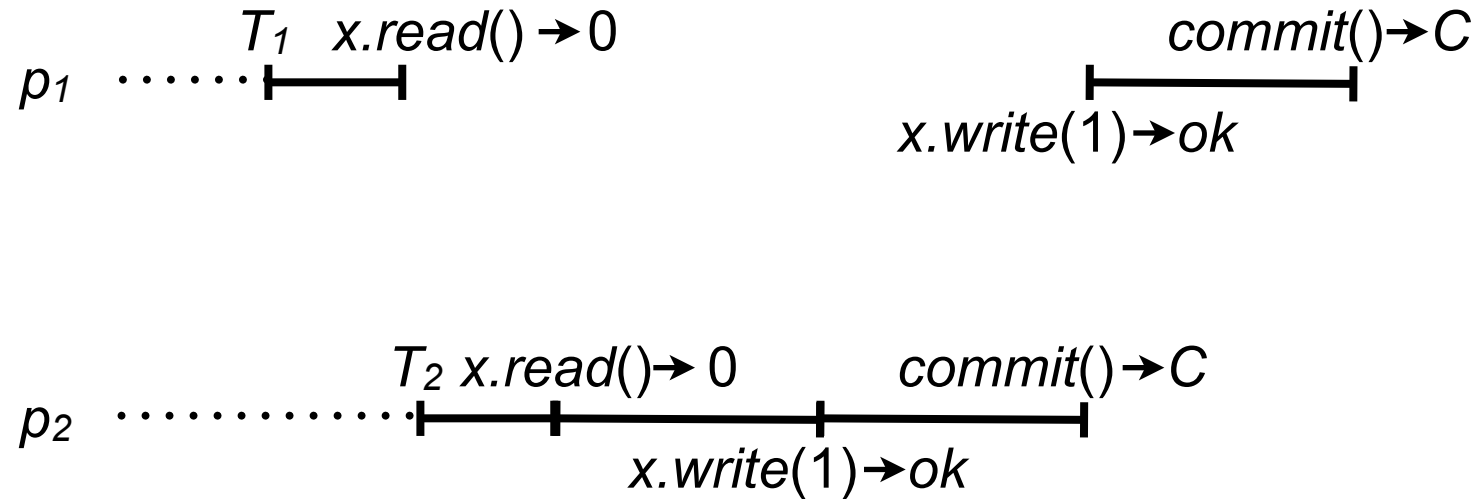


if the write by p_1 aborts we repeat the whole execution again until the write by p_1 is not aborted (by TM-wait-freedom)

Proof: execution

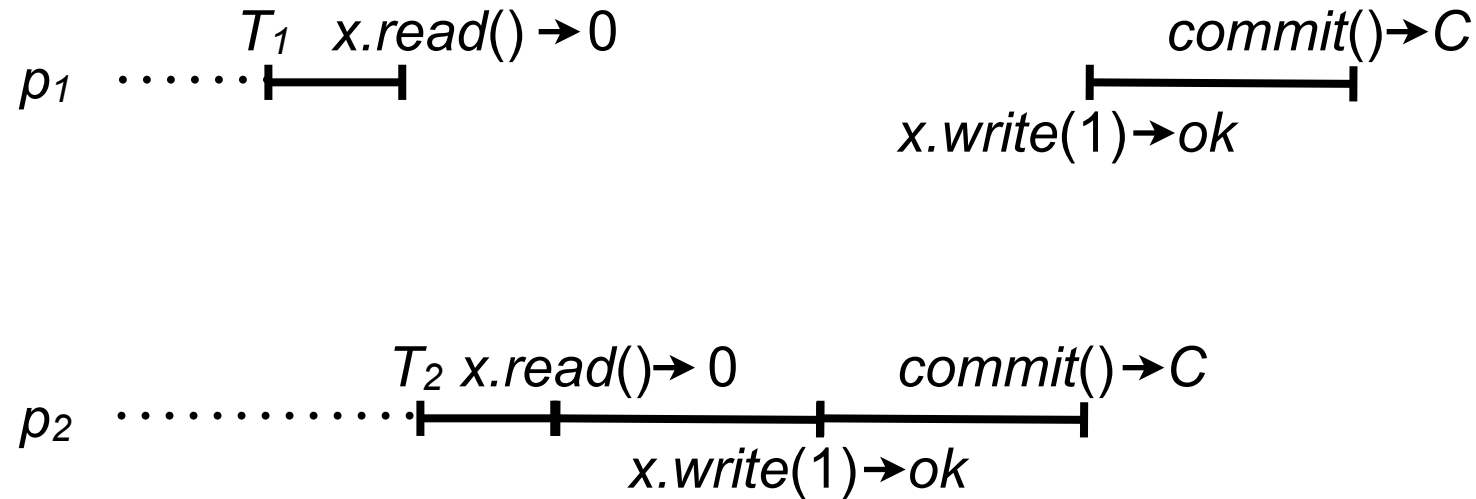


Proof: execution



what happens if T_1 is allowed to commit?

Proof: execution





what happens if T_1 is allowed to commit?

- opacity is violated



Proof: violating opacity

T_1 is serialized before T_2





Proof: violating opacity

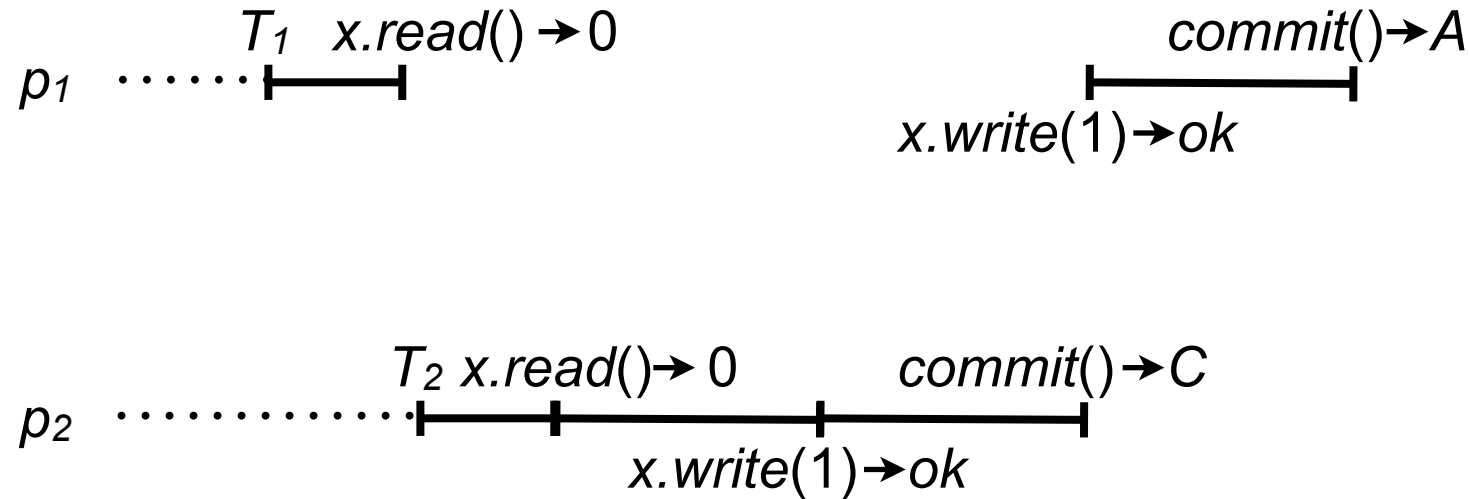
T_1 is serialized before T_2



T_2 is serialized before T_1




Proof: execution



after aborting T_1 we repeat the execution infinitely often

Proof: execution



We get an infinite execution in which:

- p_1 is correct

Circumventing impossibility

To implement TM-wait-freedom

- consider a safety property weaker than opacity

Circumventing impossibility

To implement TM-wait-freedom

- consider a safety property weaker than opacity
- consider a weaker model
 - partially synchronous system in which some process crashes are detectable and no transaction can loop forever without invoking a commit request

Circumventing impossibility

To implement TM-wait-freedom

- consider a safety property weaker than opacity
- consider a weaker model
 - partially synchronous system in which some process crashes are detectable and no transaction can loop forever without invoking a commit request
 - model in which a transaction can be executed by several processes (helping mechanism)

Resources

Overview paper on the liveness of TM:

https://lpd.epfl.ch/site/_media/education/tm_liveness_paper.pdf