

**Paper:** <https://doi.org/jzwj>  
**Artifact:** <https://doi.org/jzwk>

# Formalising the Prevention of Microarchitectural Timing Channels by Operating Systems

Formal Methods (FM), 7 March 2023

Robert Sison<sup>1,2</sup>, Scott Buckley<sup>2</sup>,  
Toby Murray<sup>1</sup>, Gerwin Klein<sup>3,2</sup>, and Gernot Heiser<sup>2</sup>



<sup>1</sup> The University of  
Melbourne, Australia



<sup>2</sup> UNSW Sydney,  
Australia



<sup>3</sup> Proofcraft,  
Sydney, Australia

# Formalising the Prevention of Microarchitectural Timing Channels by Operating Systems

Formal Methods (FM), 7 March 2023

Robert Sison<sup>1,2</sup>, Scott Buckley<sup>2</sup>,  
Toby Murray<sup>1</sup>, Gerwin Klein<sup>3,2</sup>, and Gernot Heiser<sup>2</sup>

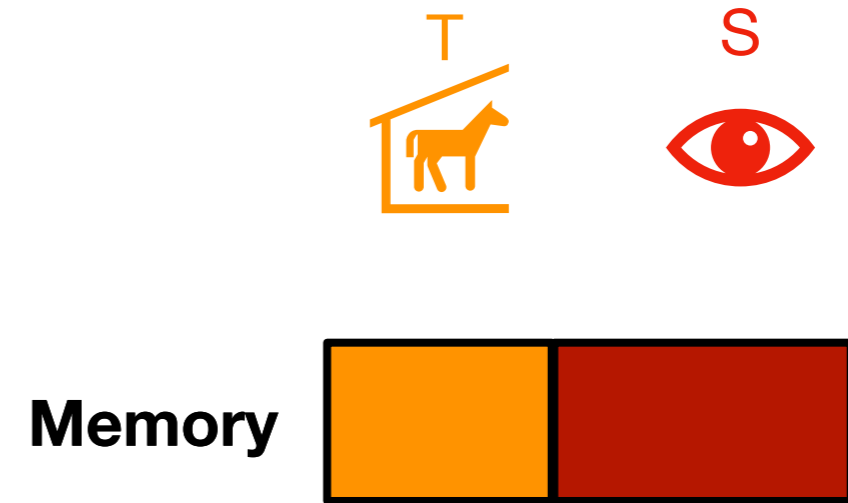


<sup>1</sup> The University of  
Melbourne, Australia

<sup>2</sup> UNSW Sydney,  
Australia

<sup>3</sup> Proofcraft,  
Sydney, Australia

# Threat scenario: *Trojan and spy*



# Threat scenario: *Trojan and spy*

Covert channels



Memory



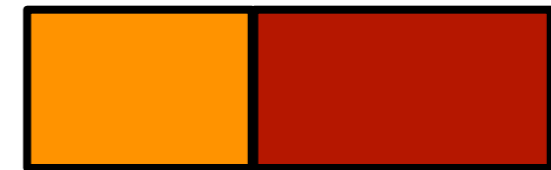
# Threat scenario:

## *Victim*/~~Trojan~~ and spy?

Covert channels  
+  
Side channels



Memory



# Threat scenario:

## *Victim*/~~Trojan~~ and spy?

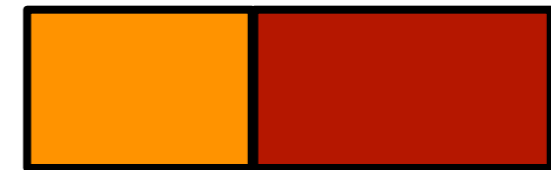
Covert channels ✓

+

Side channels ✓



Memory

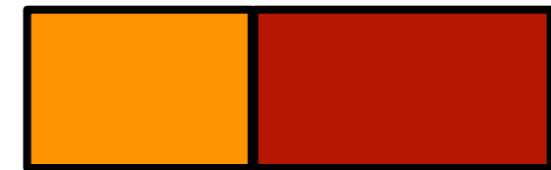


# Threat scenario: *Trojan and spy*

Covert channels  
+  
Side channels



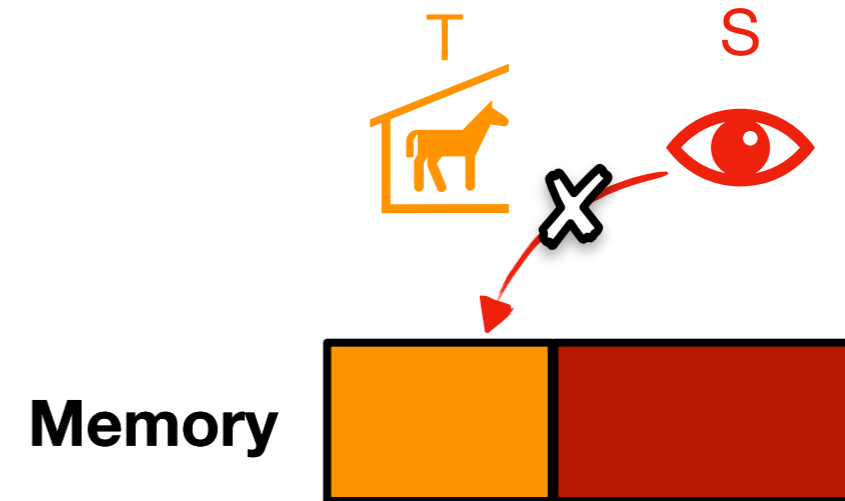
Memory



# Threat scenario: *Trojan and spy*

Covert channels  
+  
Side channels

- OSes typically implement *memory protection*.

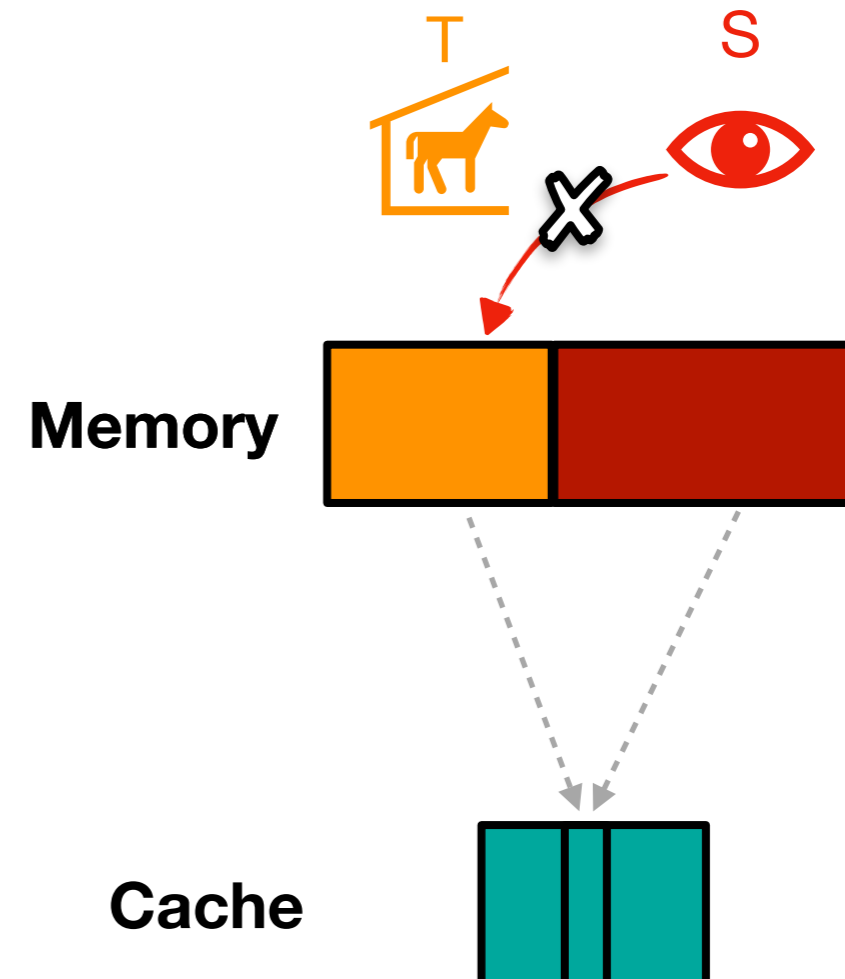




# Threat scenario: *Trojan and spy*

Covert channels  
+  
Side channels

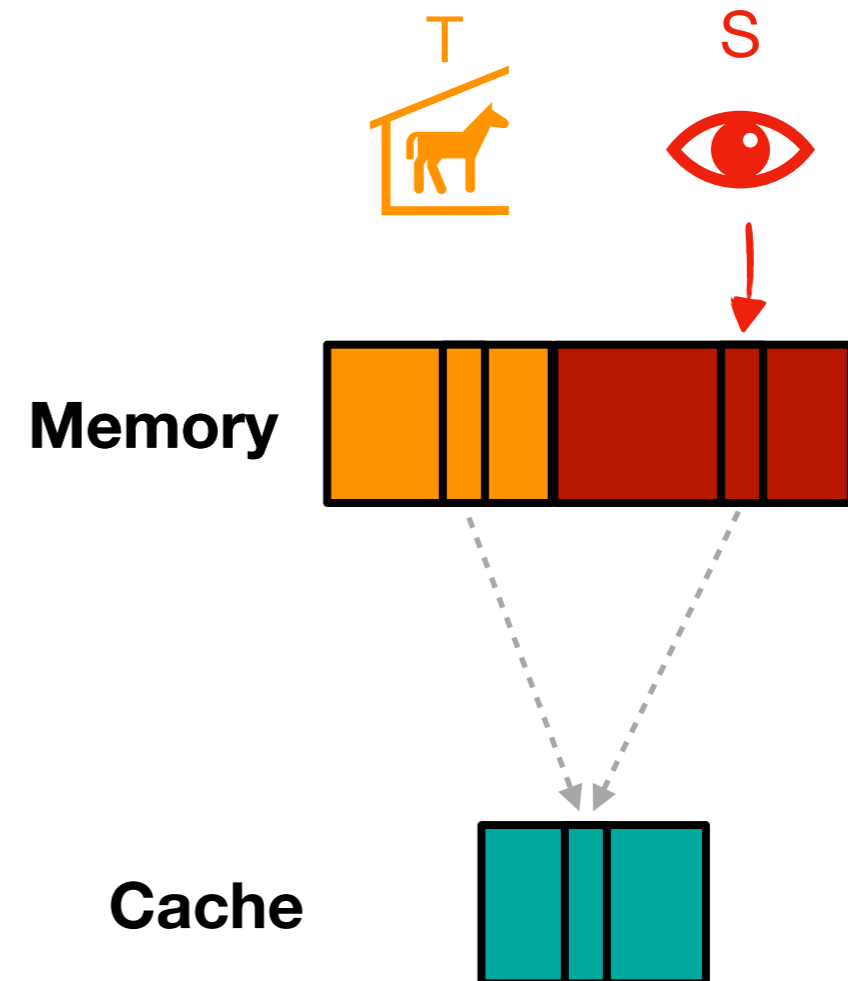
- OSes typically implement *memory protection*.
- But: Mere memory access can change the microarch. state — this affects timing.



# Threat scenario: *Trojan and spy*

Covert channels  
+  
Side channels

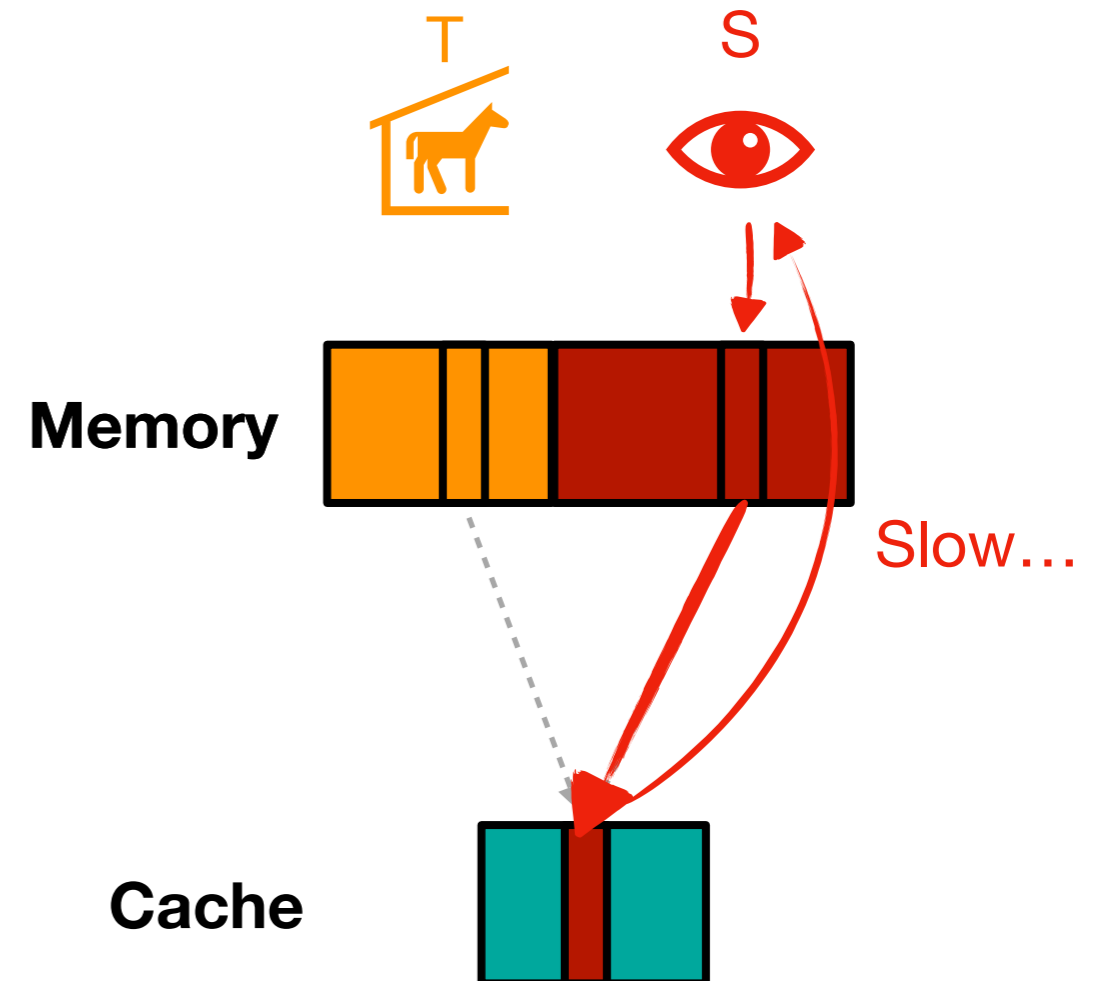
- OSes typically implement *memory protection*.
- But: Mere memory access can change the microarch. state — this affects timing.



# Threat scenario: *Trojan and spy*

Covert channels  
+  
Side channels

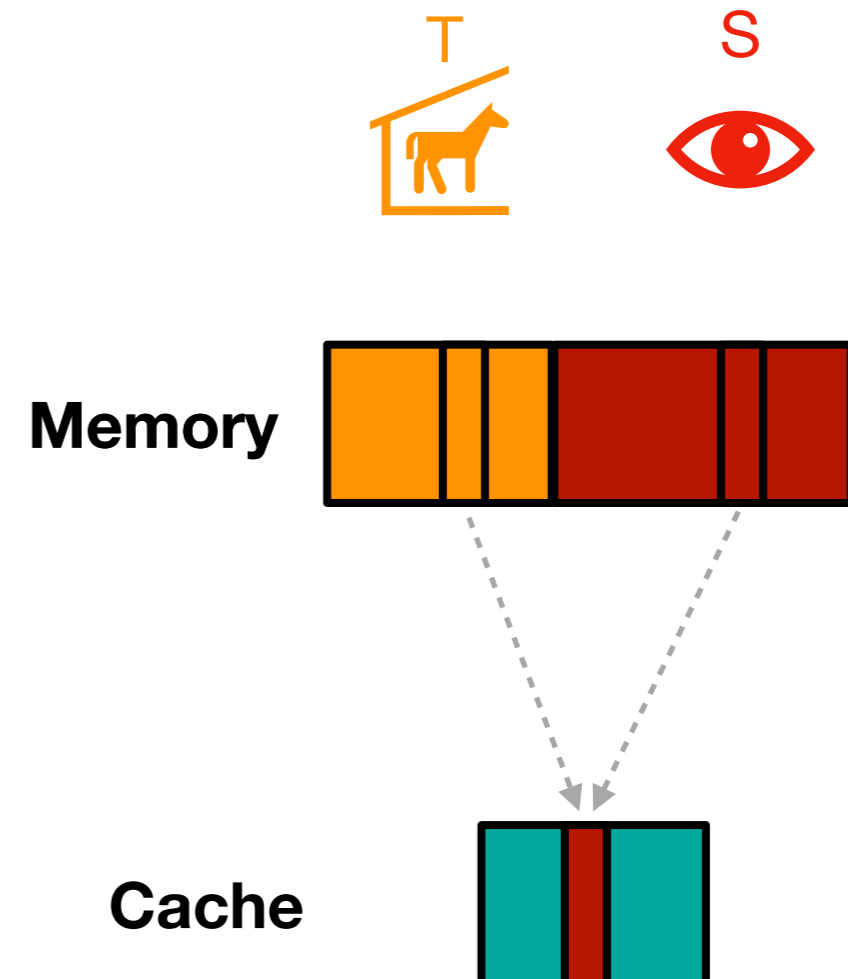
- OSes typically implement *memory protection*.
- But: Mere memory access can change the microarch. state — this affects timing.



# Threat scenario: *Trojan and spy*

Covert channels  
+  
Side channels

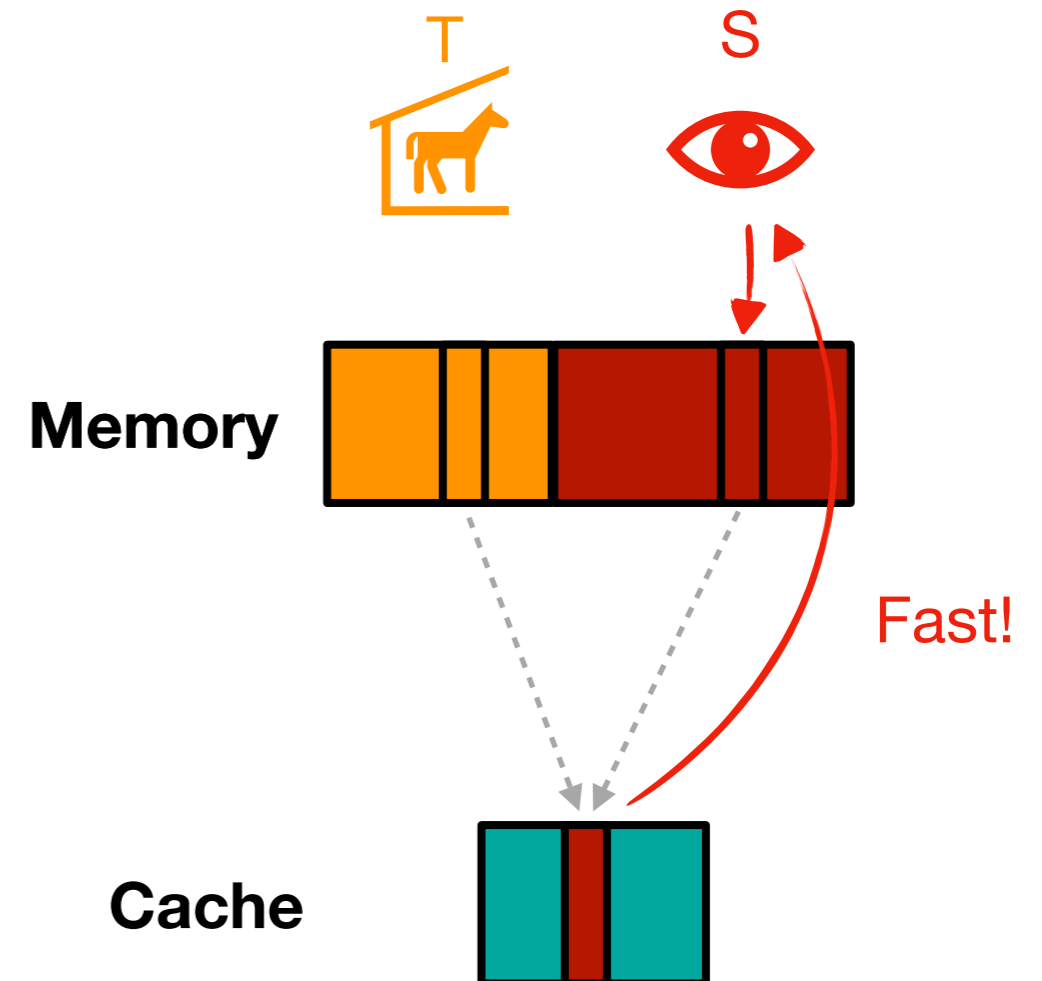
- OSes typically implement *memory protection*.
- But: Mere memory access can change the microarch. state — this affects timing.



# Threat scenario: *Trojan and spy*

Covert channels  
+  
Side channels

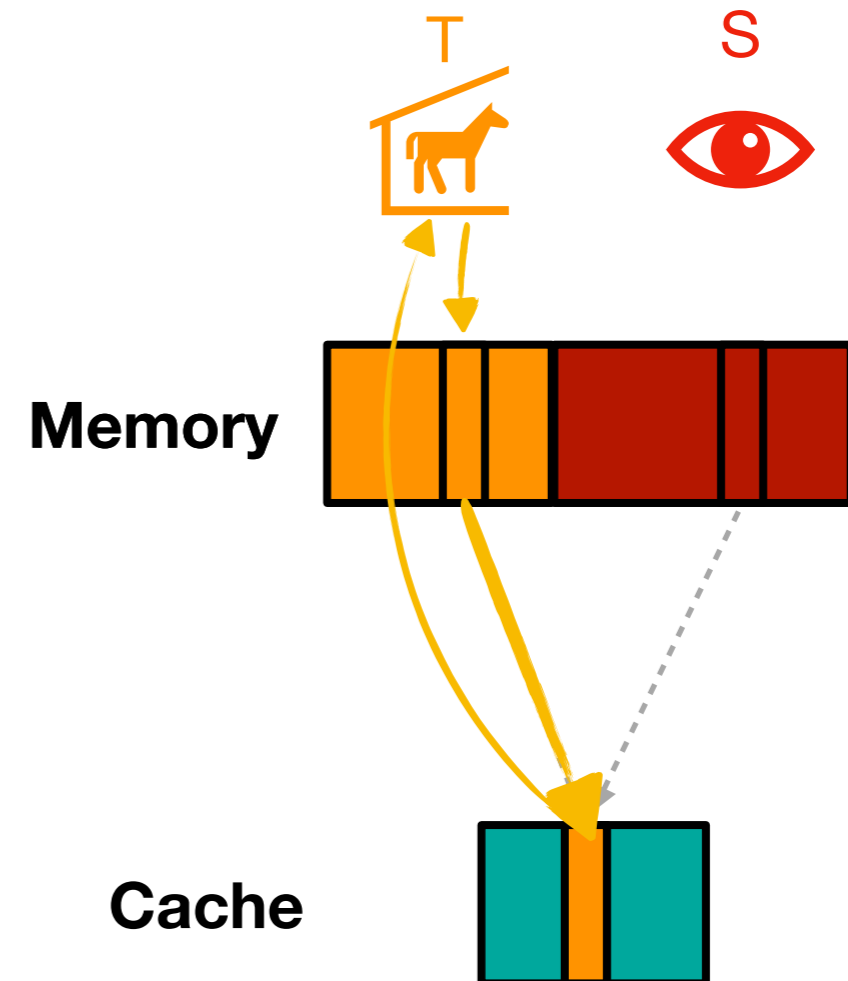
- OSes typically implement *memory protection*.
- But: Mere memory access can change the microarch. state — this affects timing.



# Threat scenario: *Trojan and spy*

Covert channels  
+  
Side channels

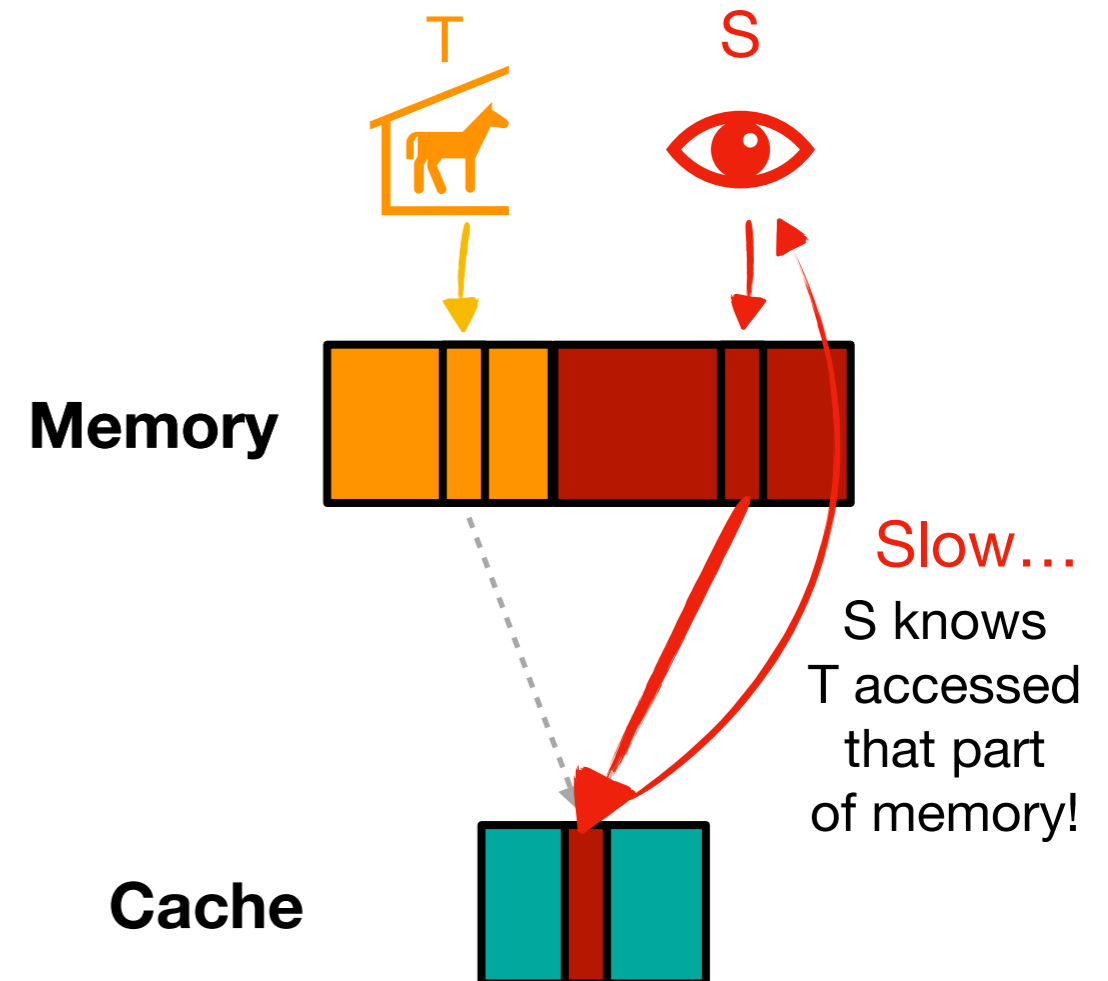
- OSes typically implement *memory protection*.
- But: Mere memory access can change the microarch. state — this affects timing.



# Threat scenario: *Trojan and spy*

Covert channels  
+  
Side channels

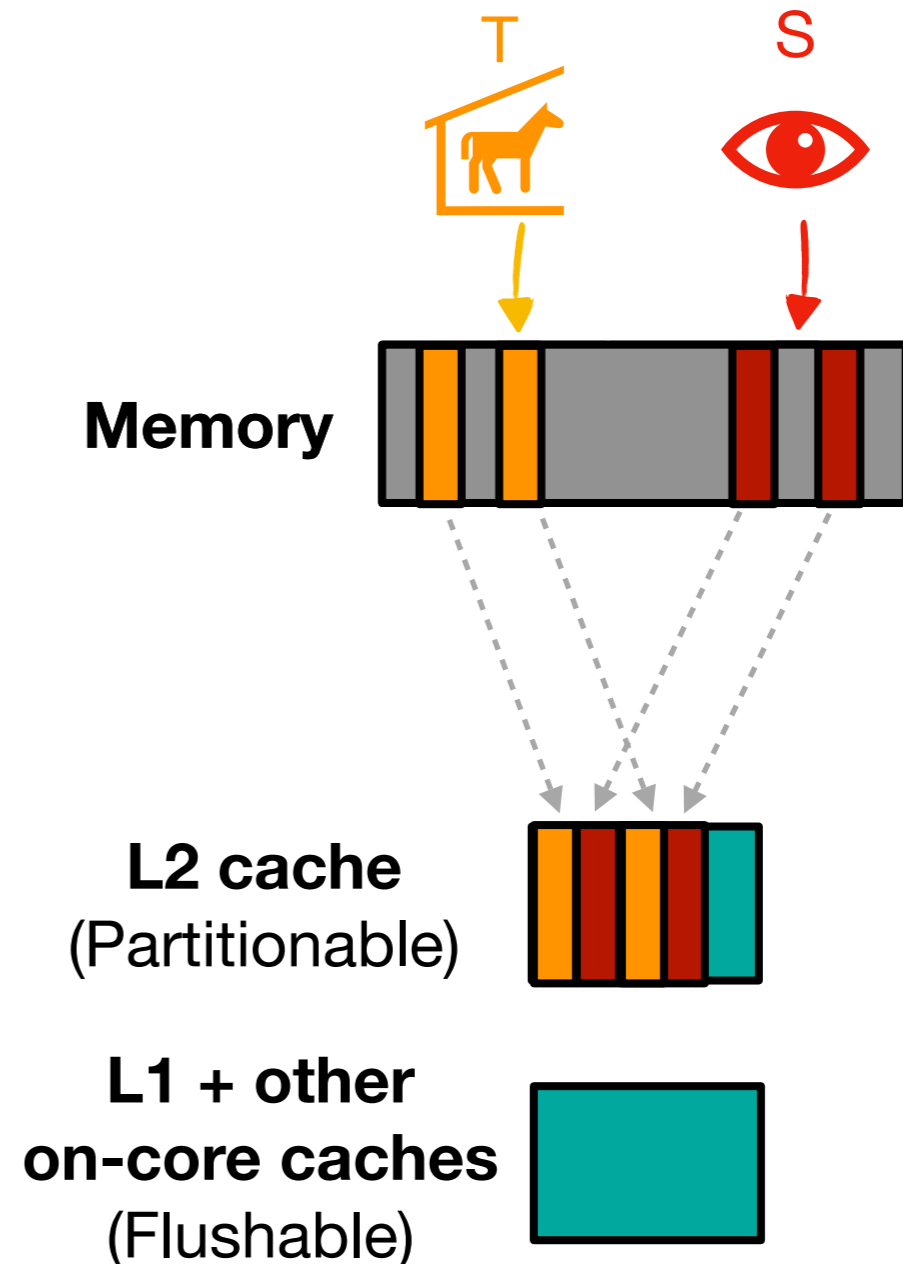
- OSes typically implement *memory protection*.
- But: Mere memory access can change the microarch. state — this affects timing.



# Threat scenario: *Trojan and spy*

Covert channels  
+  
Side channels

- OSes typically implement *memory protection*.
- But: Mere memory access can change the microarch. state — this affects timing.
- To prevent these *timing channels*, OSes can implement *time protection*:  
e.g. [Ge et al. 2019] for seL4 microkernel OS
  - Partition what we can
  - Flush what we can't

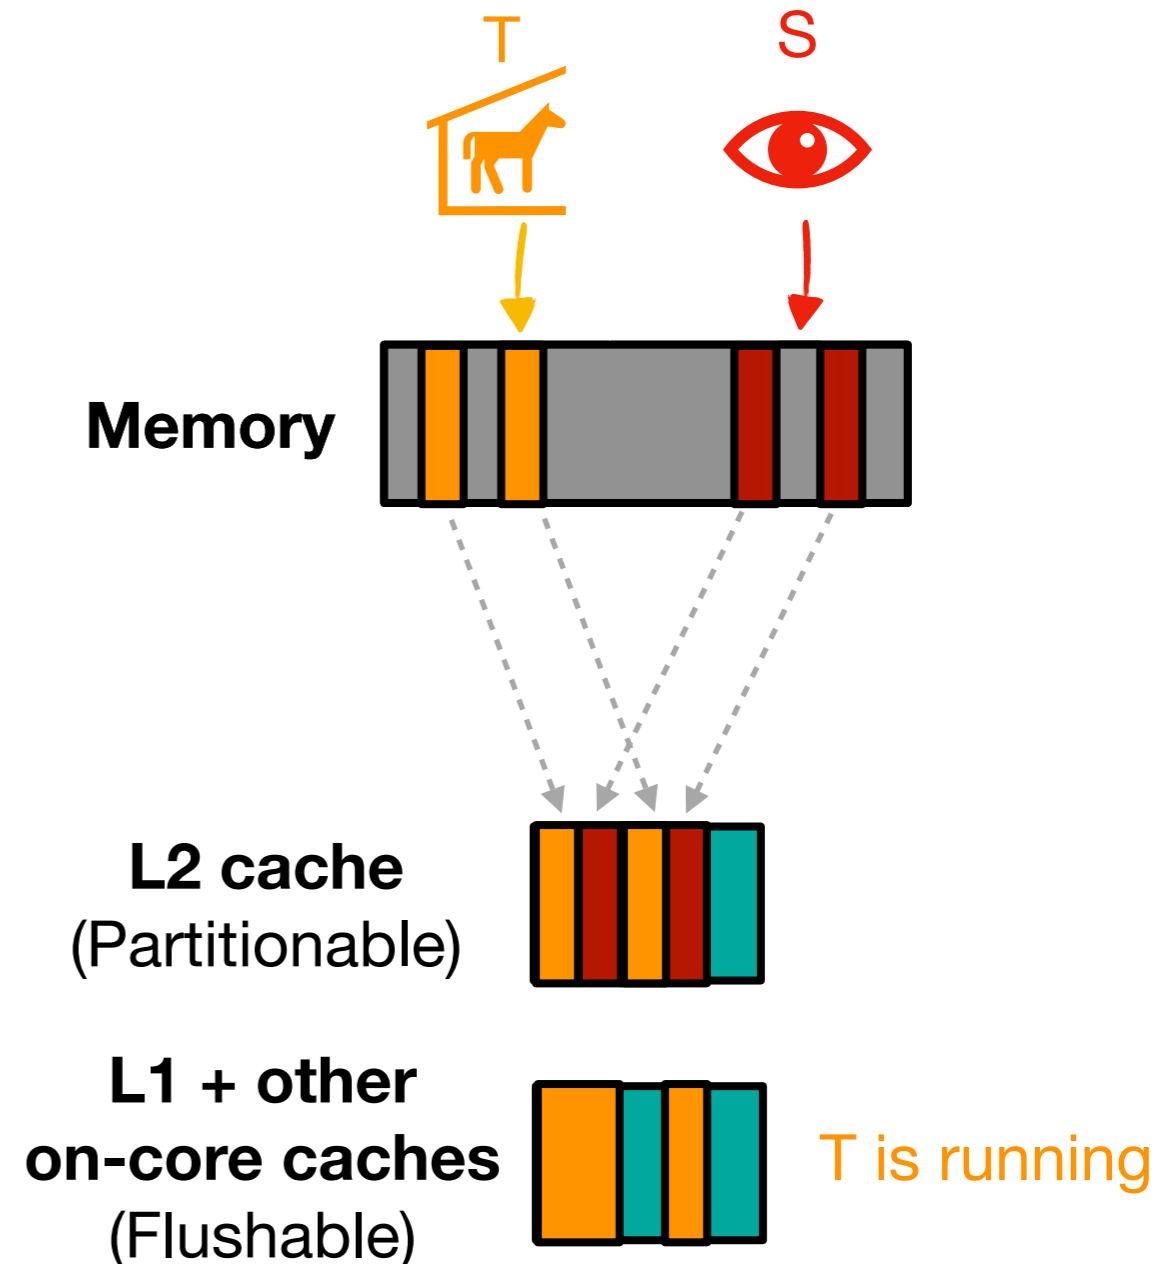




# Threat scenario: *Trojan and spy*

Covert channels  
+  
Side channels

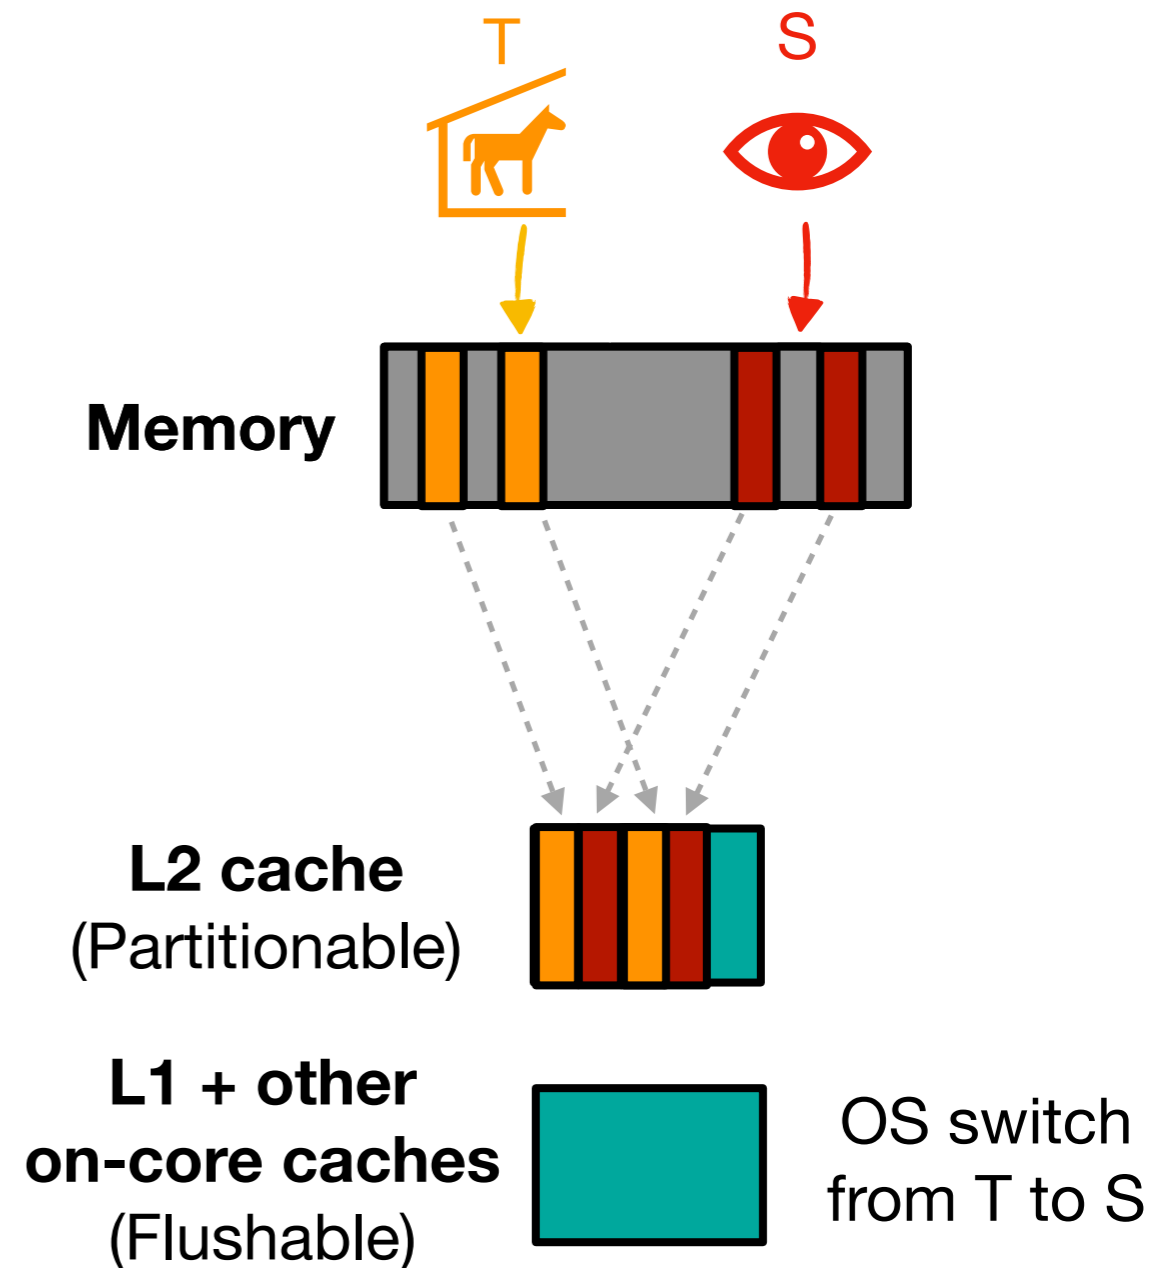
- OSes typically implement *memory protection*.
- But: Mere memory access can change the microarch. state — this affects timing.
- To prevent these *timing channels*, OSes can implement *time protection*:  
e.g. [Ge et al. 2019] for seL4 microkernel OS
  - Partition what we can
  - Flush what we can't



# Threat scenario: *Trojan and spy*

Covert channels  
+  
Side channels

- OSes typically implement *memory protection*.
- But: Mere memory access can change the microarch. state — this affects timing.
- To prevent these *timing channels*, OSes can implement *time protection*:  
e.g. [Ge et al. 2019] for seL4 microkernel OS
  - Partition what we can
  - Flush what we can't

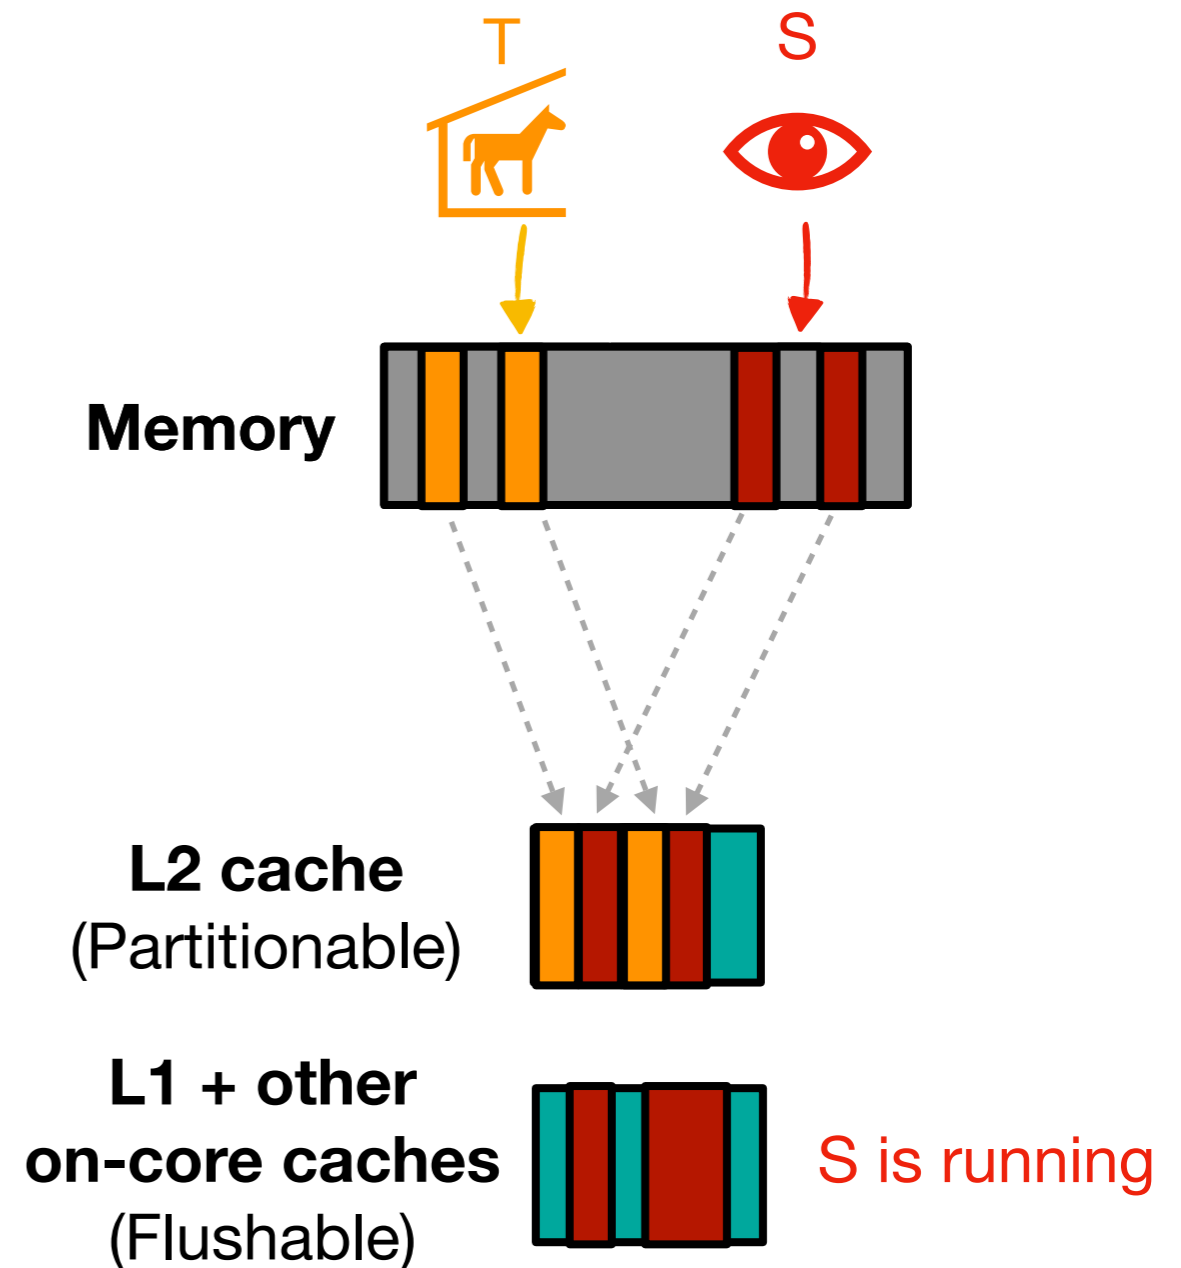


“Flush”: Write fixed content; wait up to fixed time.

# Threat scenario: *Trojan and spy*

Covert channels  
+  
Side channels

- OSes typically implement *memory protection*.
- But: Mere memory access can change the microarch. state — this affects timing.
- To prevent these *timing channels*, OSes can implement *time protection*:  
e.g. [Ge et al. 2019] for seL4 microkernel OS
  - Partition what we can
  - Flush what we can't

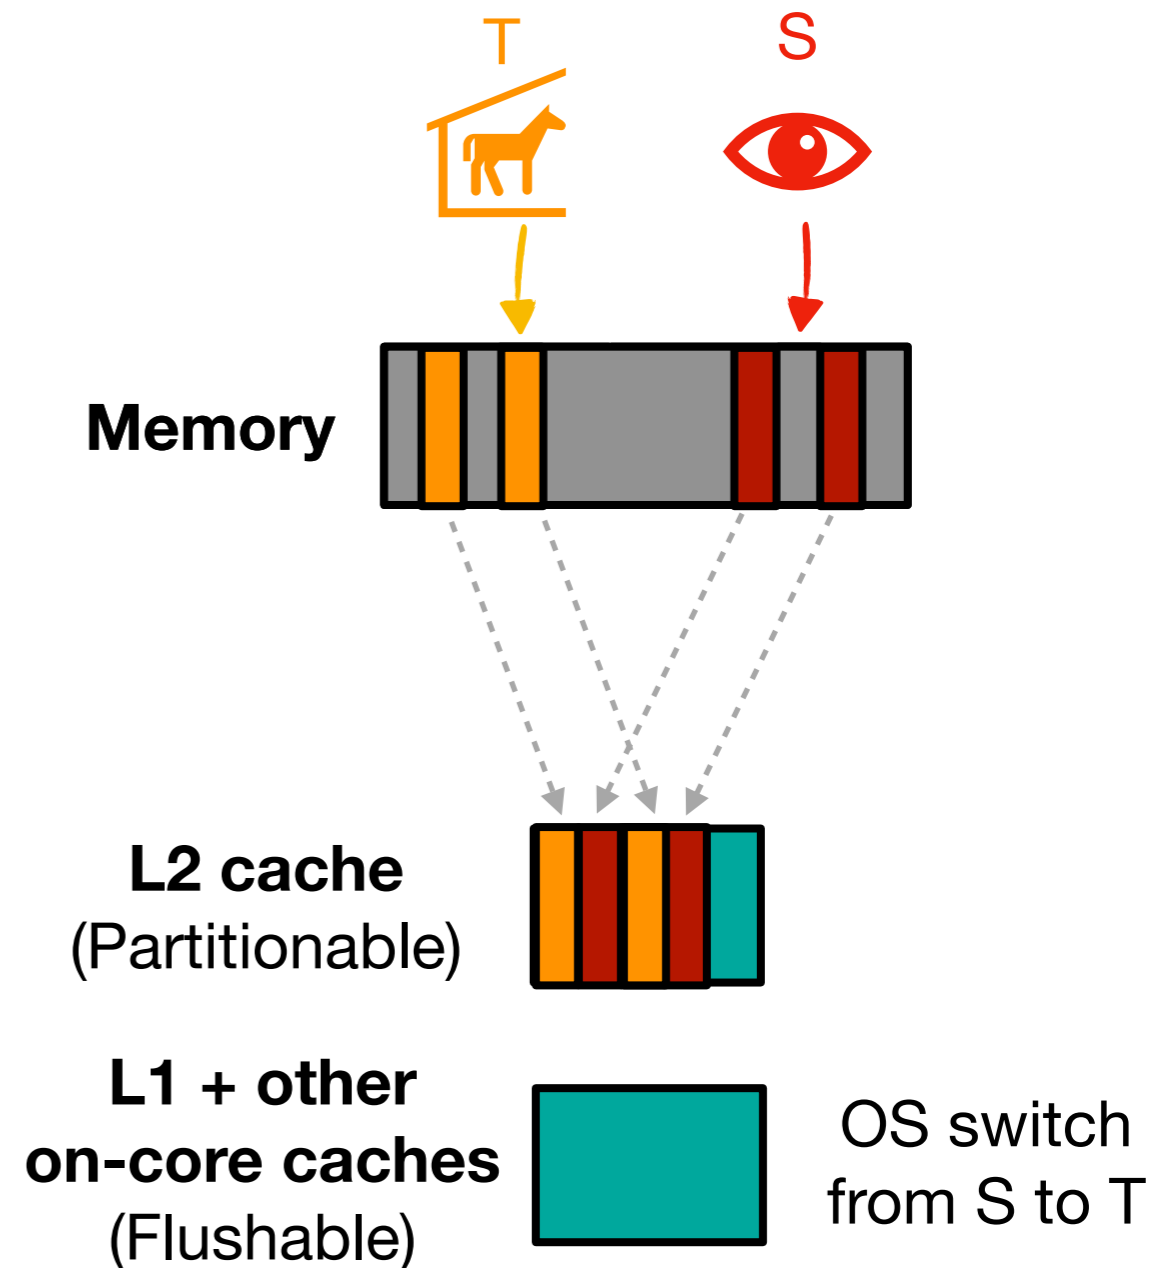


“Flush”: Write fixed content; wait up to fixed time.

# Threat scenario: *Trojan and spy*

Covert channels  
+  
Side channels

- OSes typically implement *memory protection*.
- But: Mere memory access can change the microarch. state — this affects timing.
- To prevent these *timing channels*, OSes can implement *time protection*:  
e.g. [Ge et al. 2019] for seL4 microkernel OS
  - Partition what we can
  - Flush what we can't



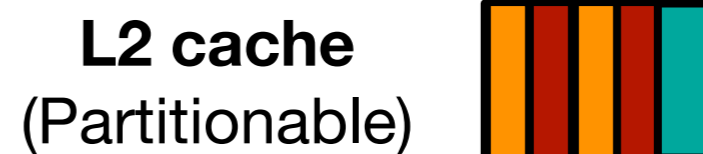
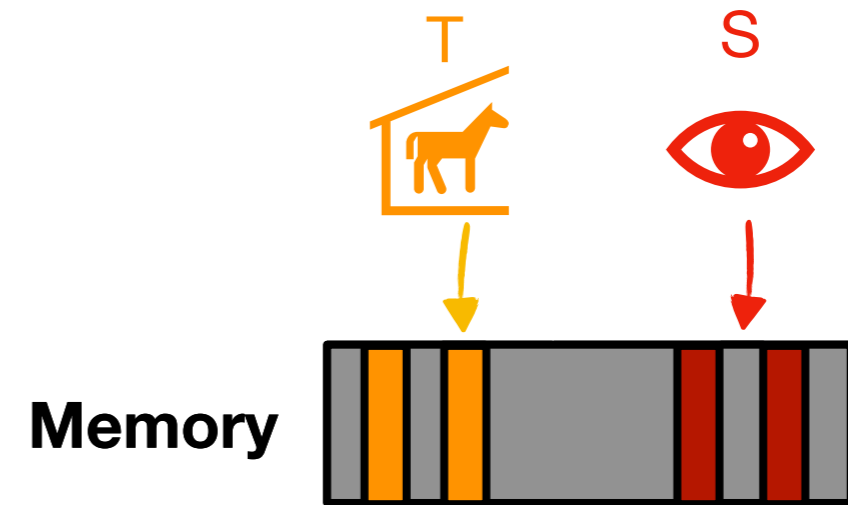
“Flush”: Write fixed content; wait up to fixed time.

# Threat scenario: *Trojan and spy*

Covert channels  
+  
Side channels

- OSes typically implement *memory protection*.
- But: Mere memory access can change the microarch. state — this affects timing.
- To prevent these *timing channels*, OSes can implement *time protection*:  
e.g. [Ge et al. 2019] for seL4 microkernel OS

- Partition what we can *HW-SW (hardware-software) contract*
- Flush what we can't



“Flush”: Write fixed content; wait up to fixed time.

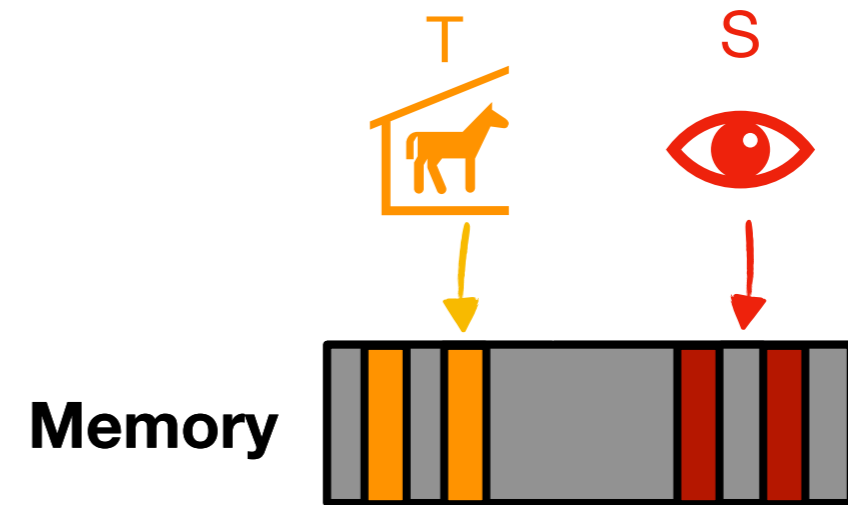
# Threat scenario: *Trojan and spy*

Covert channels  
+  
Side channels

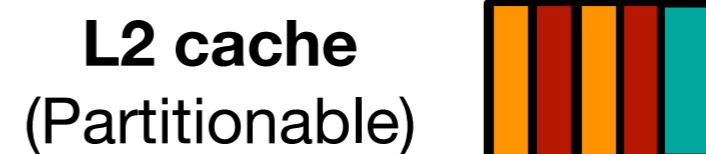
- OSes typically implement *memory protection*.
- But: Mere memory access can change the microarch. state — this affects timing.

- To prevent these *timing channels*, OSes can implement *time protection*: e.g. [Ge et al. 2019] for seL4 microkernel OS

- Partition what we can *HW-SW (hardware-software) contract*
- Flush what we can't



**Today:**  
How to formalise?

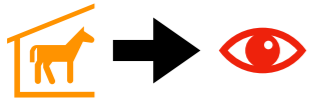


“Flush”: Write fixed content; wait up to fixed time.

# How to formalise an OS enforces *time protection*?

Versus threat scenario:

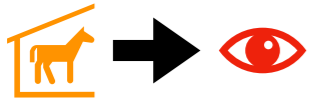
trojan and spy



# How to formalise an OS enforces *time protection*?

Versus threat scenario:

trojan and spy



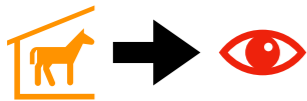
Abstract *covert state* + *time* to reflect  
strategies enabled by HW:

Partition or flush state; pad time.



# How to formalise an OS enforces *time protection*?

Versus threat scenario:  
trojan and spy



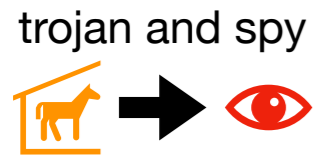
Abstract *covert state* + *time* to reflect strategies enabled by HW:  
Partition or flush state; pad time.



Make security property precise enough to exclude flows from covert state.

# How to formalise an OS enforces *time protection*?

Versus threat scenario:



Abstract *covert state* + *time* to reflect strategies enabled by HW:

Partition or flush state; pad time.

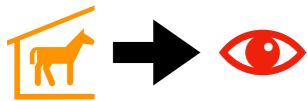
Make security property precise enough to exclude flows from covert state.



Demonstrating these principles, we formalised in Isabelle/HOL:

# How to formalise an OS enforces *time protection*?

Versus threat scenario:  
trojan and spy



Abstract *covert state* + *time* to reflect strategies enabled by HW:  
Partition or flush state; pad time.

Make security property precise enough to exclude flows from covert state.

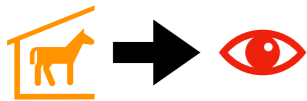


Demonstrating these principles, we formalised in Isabelle/HOL:

1. OS security model imposing requirements on relevant parts of OS.  
(Intended for seL4, but *generic*)

# How to formalise an OS enforces *time protection*?

Versus threat scenario:  
trojan and spy



Abstract *covert state* + *time* to reflect strategies enabled by HW:  
Partition or flush state; pad time.

Make security property precise enough to exclude flows from covert state.



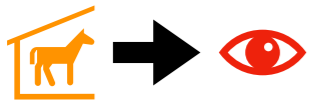
Demonstrating these principles, we formalised in Isabelle/HOL:

1. OS security model imposing requirements on relevant parts of OS.  
(Intended for seL4, but *generic*)

2. OS security property that is dynamic; this makes it observer relative.  
(Improving on seL4's of [Murray et al. 2012])

# How to formalise an OS enforces *time protection*?

Versus threat scenario:  
trojan and spy



Abstract *covert state* + *time* to reflect strategies enabled by HW:  
Partition or flush state; pad time.

Make security property precise enough to exclude flows from covert state.



Demonstrating these principles, we formalised in Isabelle/HOL:

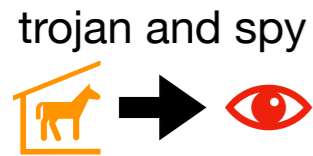
1. OS security model imposing requirements on relevant parts of OS.  
(Intended for seL4, but *generic*)

2. OS security property that is dynamic; this makes it observer relative.  
(Improving on seL4's of [Murray et al. 2012])

3. Proof our security property holds if OS model's requirements hold.

# How to formalise an OS enforces *time protection*?

Versus threat scenario:



Abstract *covert state* + *time* to reflect strategies enabled by HW:

Partition or flush state; pad time.

Make security property precise enough to exclude flows from covert state.



Demonstrating these principles, we formalised in Isabelle/HOL:

1. OS security model imposing requirements on relevant parts of OS.  
(Intended for seL4, but *generic*)

3. Proof our security property holds if OS model's requirements hold.

2. OS security property that is dynamic; this makes it observer relative.  
(Improving on seL4's of [Murray et al. 2012])

4. Basic instantiation of OS model exercising dynamic policy.

# How to formalise an OS enforces *time protection*?

Versus threat scenario:  
trojan and spy



Abstract *covert state* + *time* to reflect strategies enabled by HW:  
Partition or flush state; pad time.



Make security property precise enough to exclude flows from covert state.



Demonstrating these principles,  
we formalised in Isabelle/HOL:

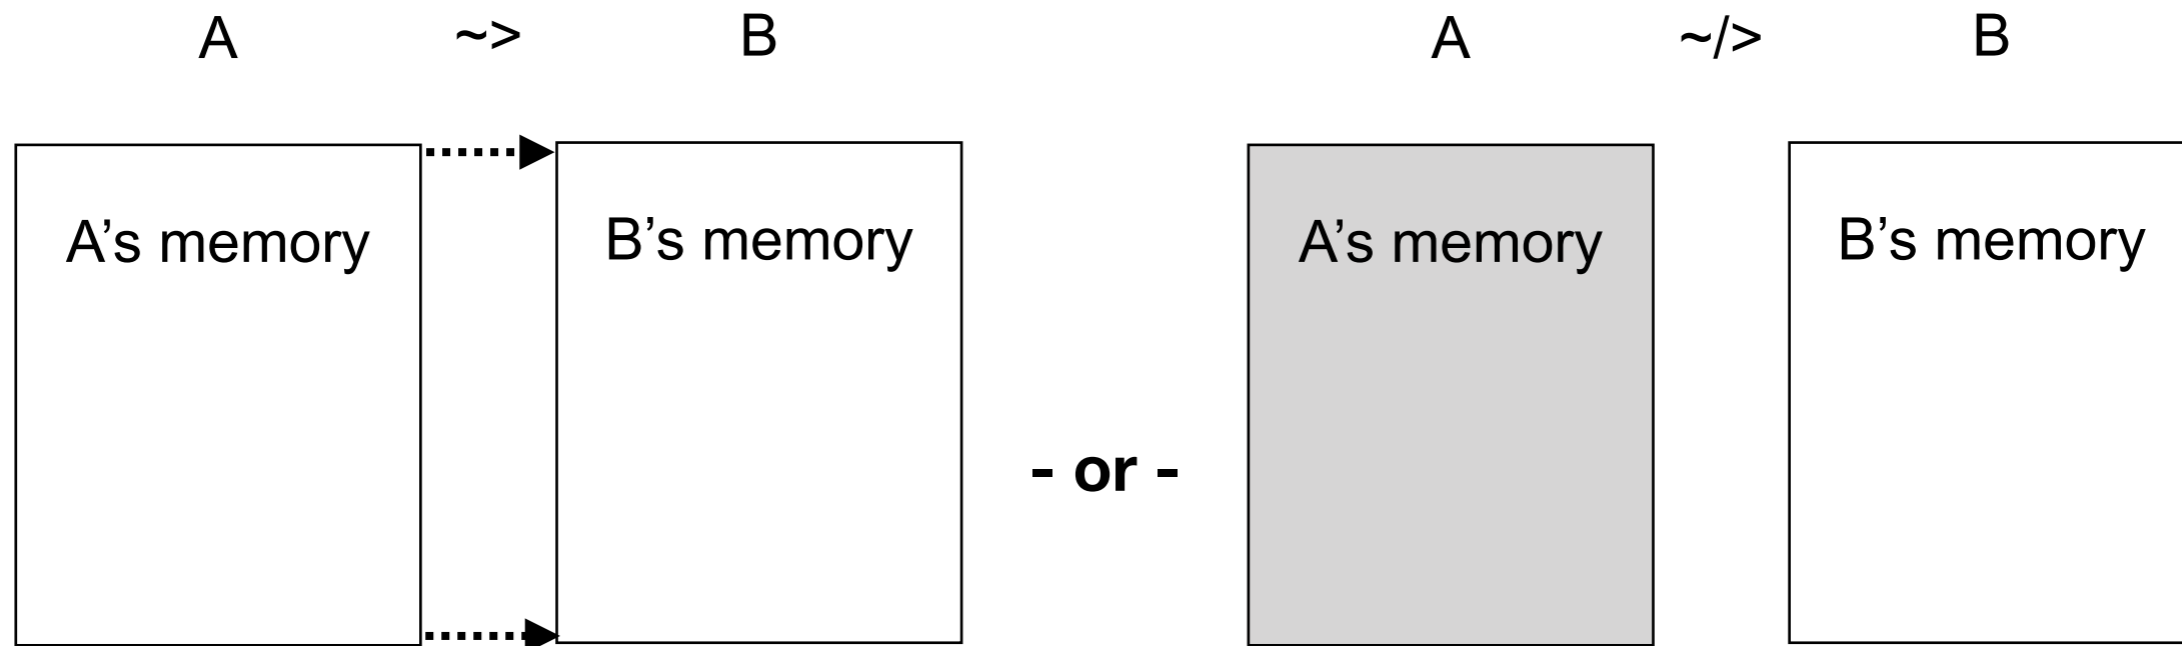
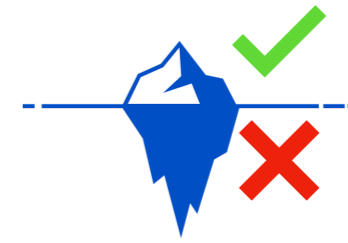
1. OS security model imposing requirements on relevant parts of OS.  
(Intended for seL4, but *generic*)

3. Proof our security property holds if OS model's requirements hold.

2. OS security property that is dynamic; this makes it observer relative.  
(Improving on seL4's of [Murray et al. 2012])

4. Basic instantiation of OS model exercising dynamic policy.

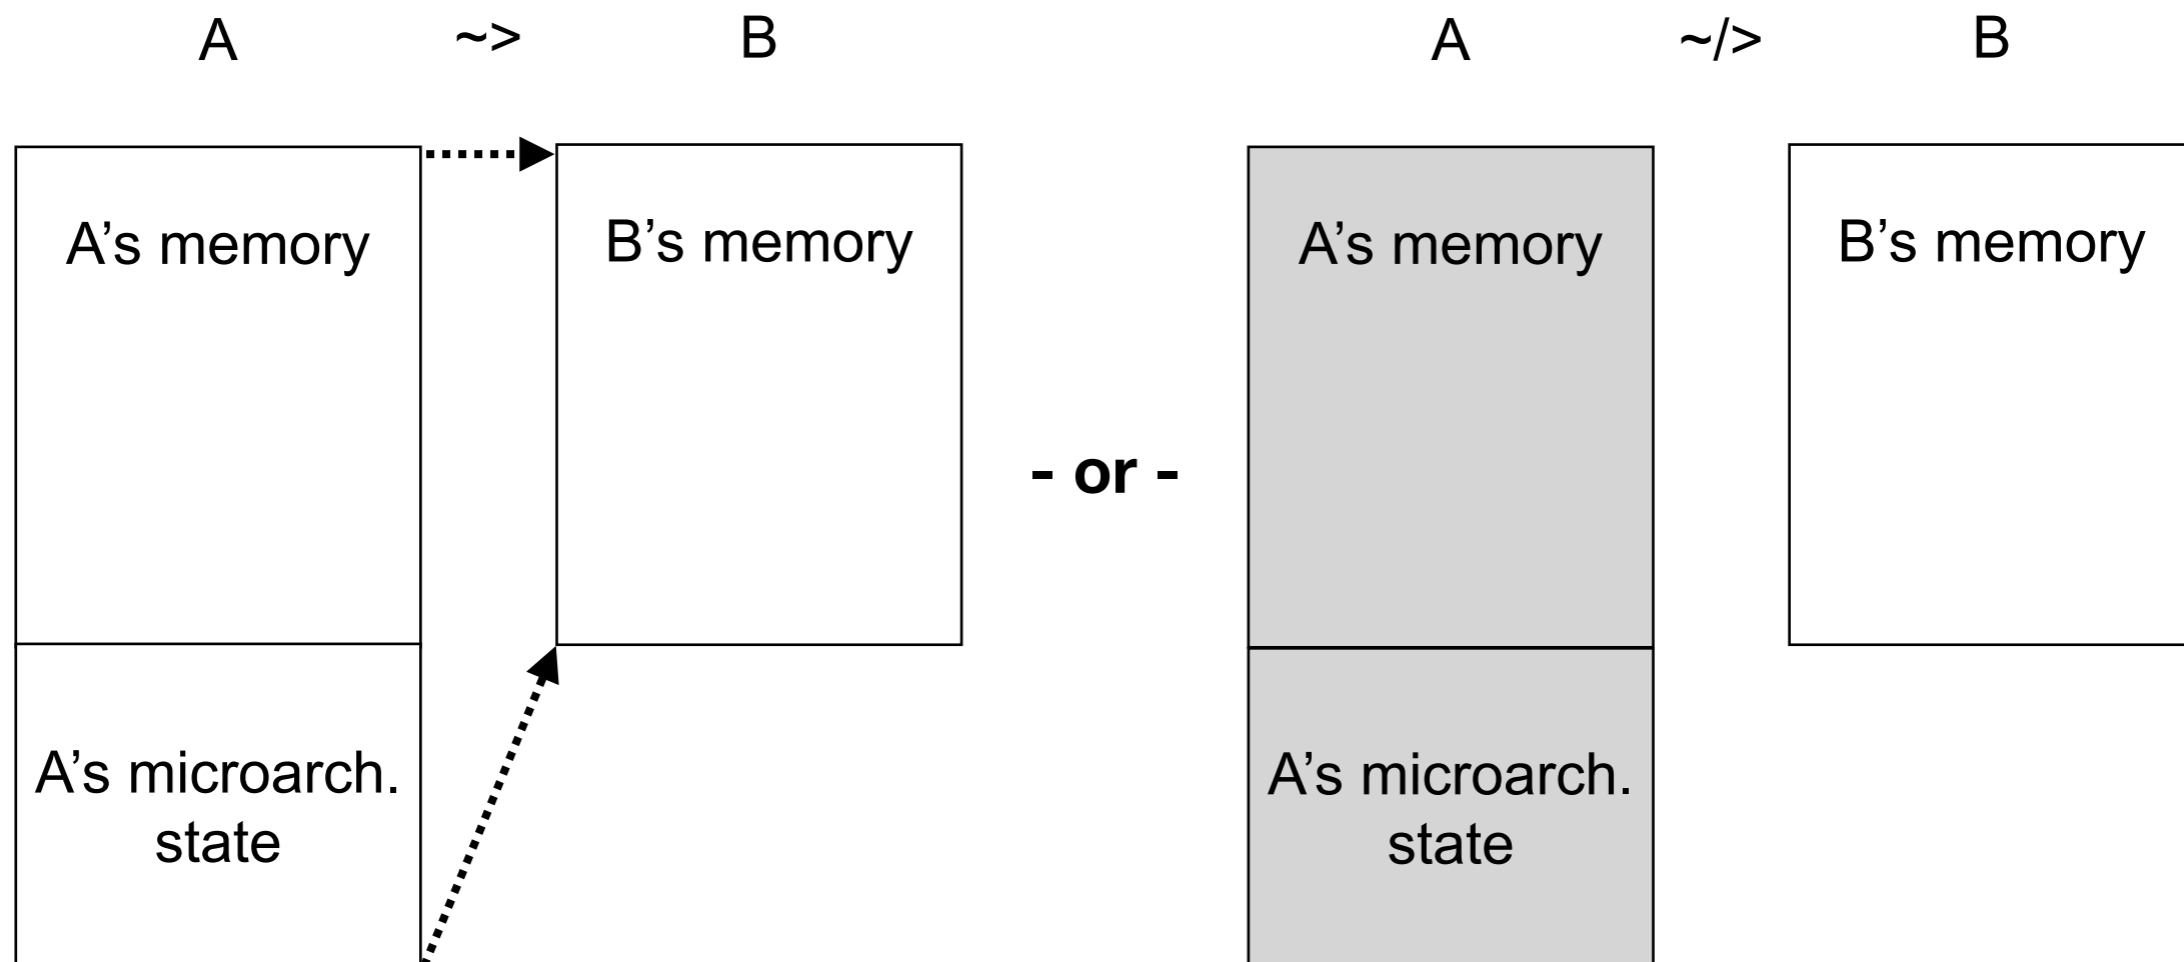
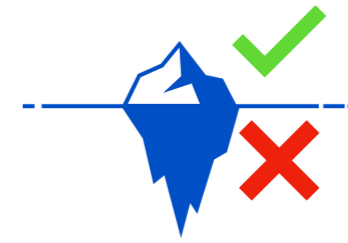
# Overt vs covert state



From prior seL4 infoflow proofs  
[Murray et al. 2012, 2013]:  
*“all or nothing” policies*

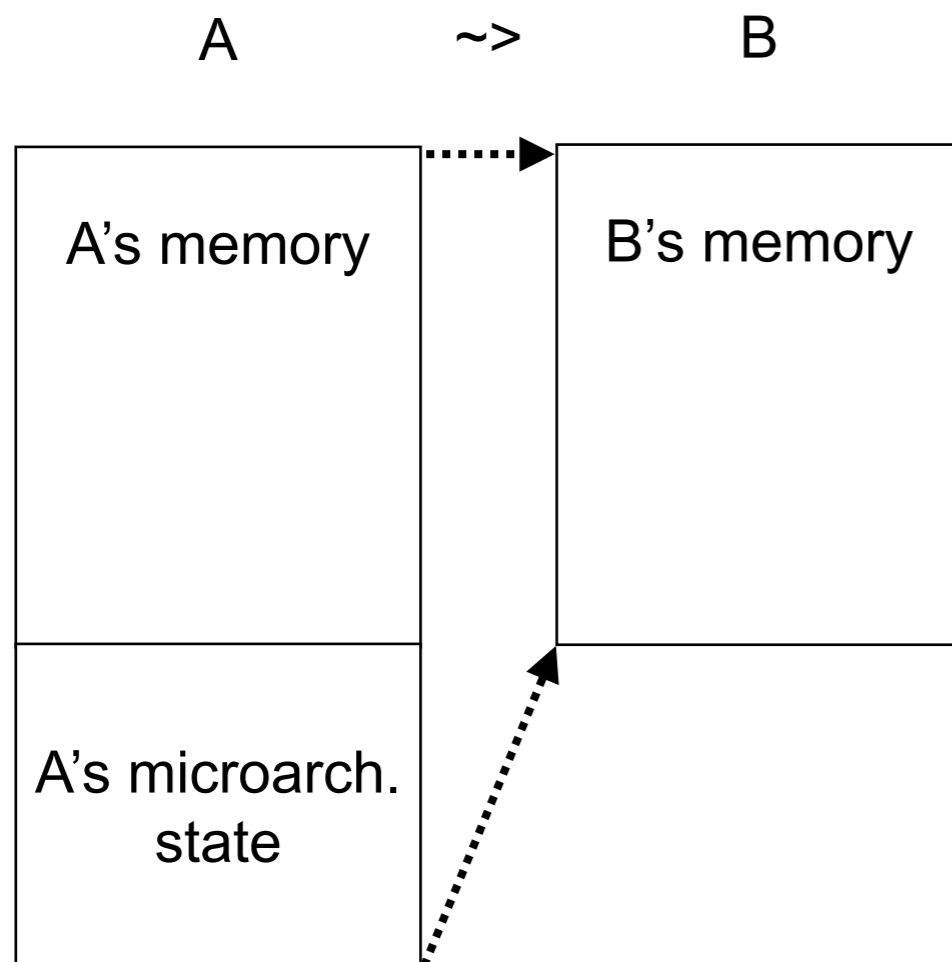
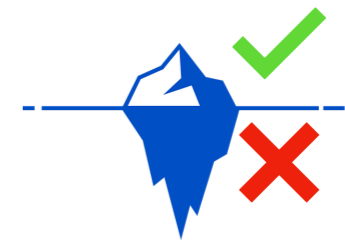


# Overt vs covert state

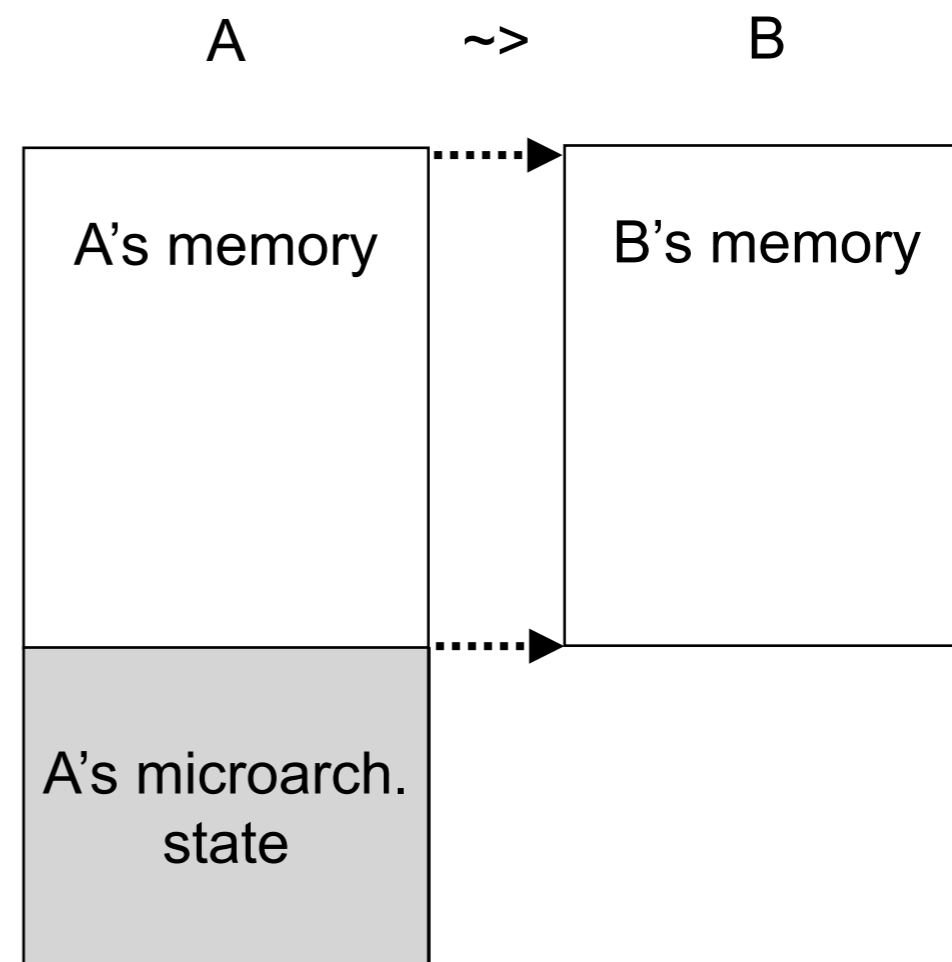


From prior seL4 infoflow proofs  
[Murray et al. 2012, 2013]:  
*“all or nothing” policies*

# Overt vs covert state



From prior seL4 infowall proofs  
[Murray et al. 2012, 2013]:  
*“all or nothing” policies*



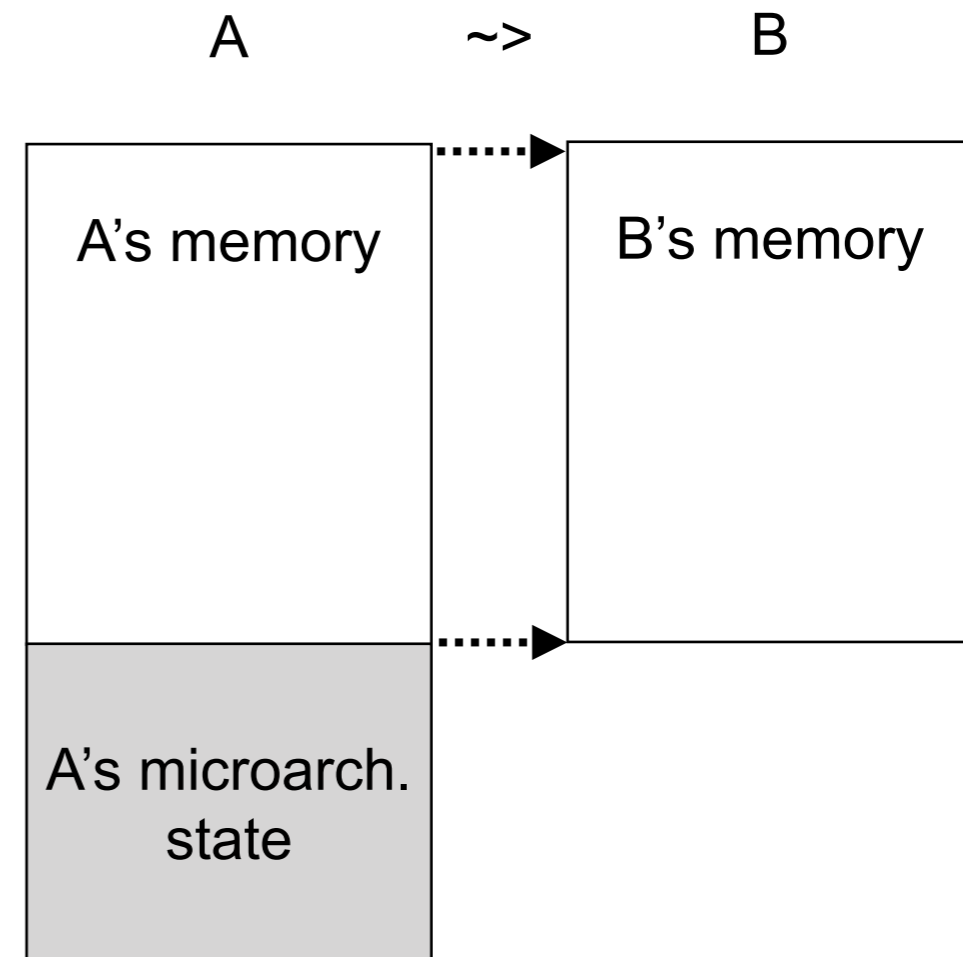
**Principle: Need policies  
to allow some (*overt*) flows  
while excluding other (*covert*) ones**

# Covert state: Partitionable vs flushable



## Principle:

**Model channels as *state elements*  
by their *elimination strategy*  
as per *HW-SW contract***



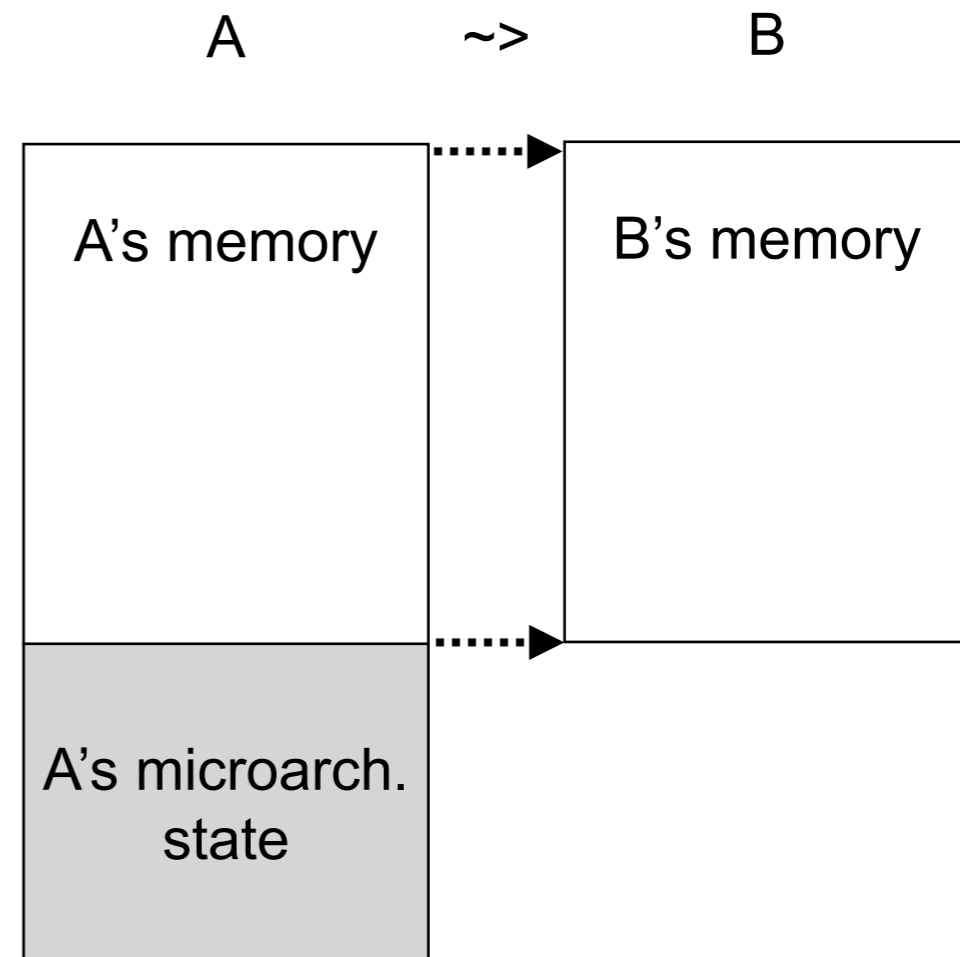
# Covert state: Partitionable vs flushable



## Principle:

**Model channels as *state elements*  
by their *elimination strategy*  
as per *HW-SW contract***

- Strategy for OS:  
*Partition or flush state; pad time.*



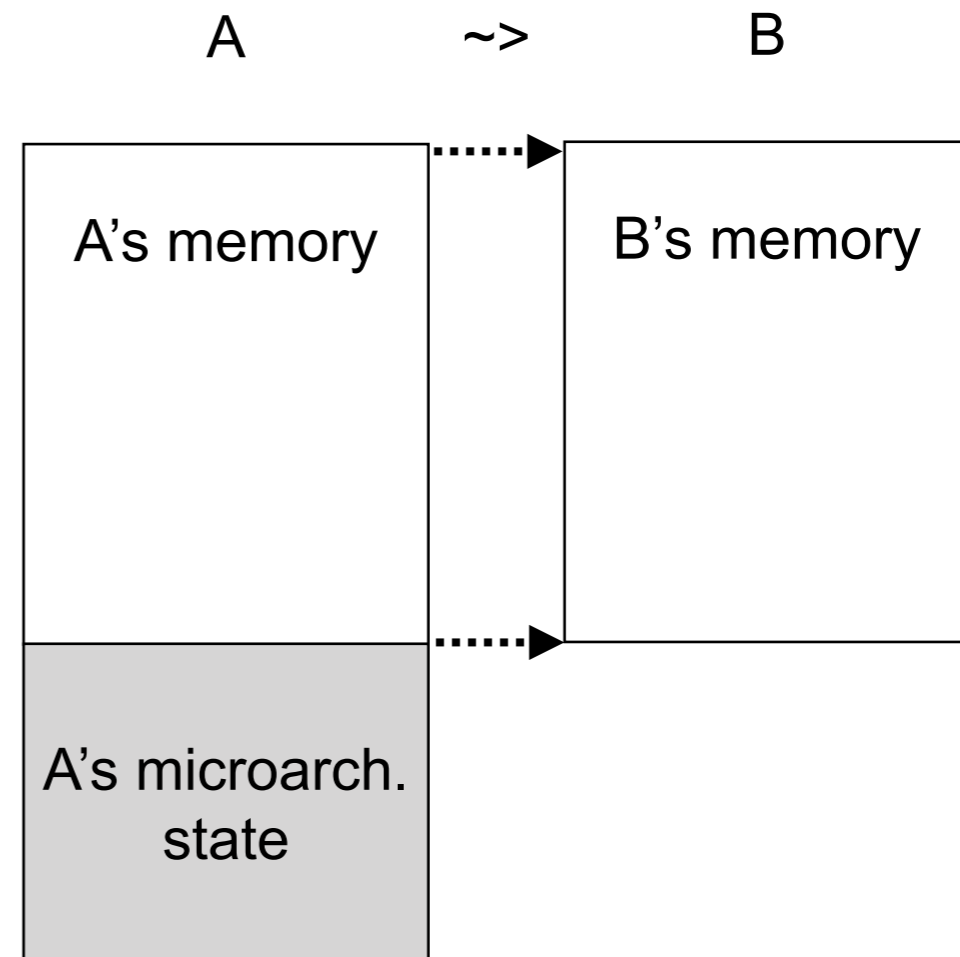
# Covert state: Partitionable vs flushable



## Principle:

**Model channels as *state elements*  
by their *elimination strategy*  
as per *HW-SW contract***

- Strategy for OS:  
*Partition or flush state; pad time.*
- Relies on HW-SW contract:



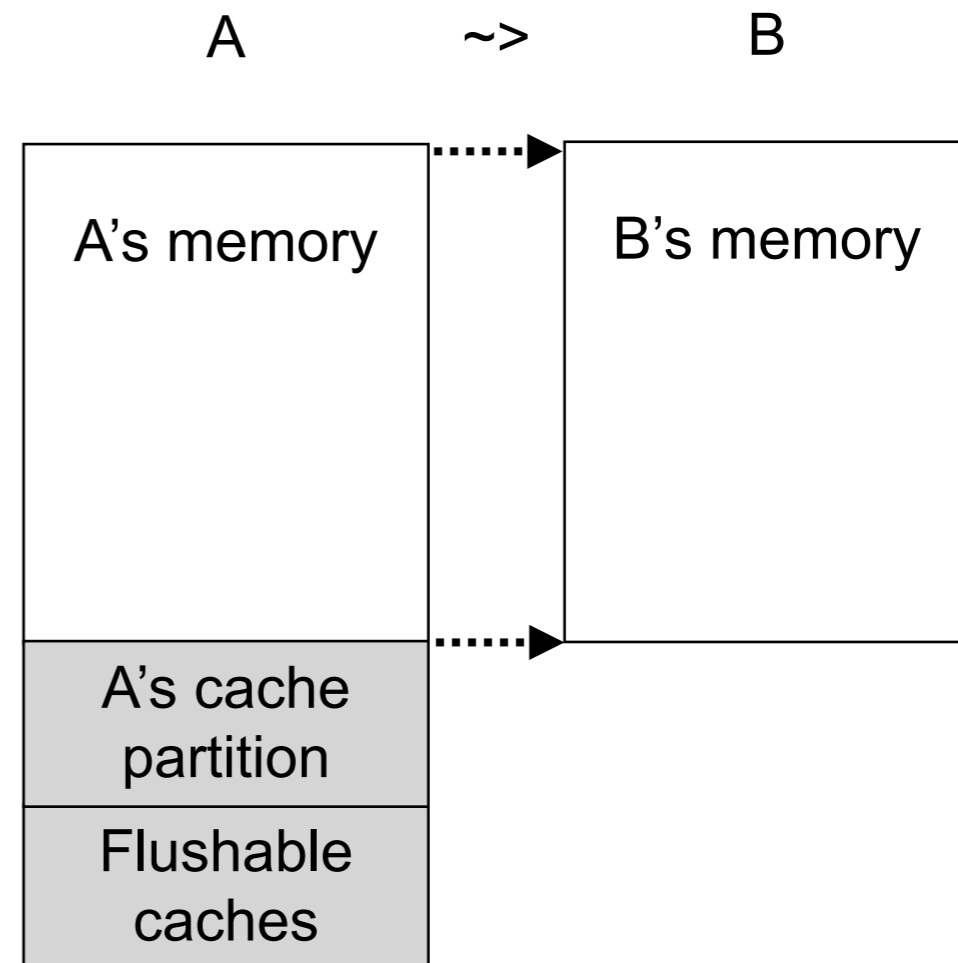
# Covert state: Partitionable vs flushable



## Principle:

**Model channels as *state elements*  
by their *elimination strategy*  
as per *HW-SW contract***

- Strategy for OS:  
*Partition or flush state; pad time.*
- Relies on HW-SW contract:
  - State: Everything must be *partitionable*  
or *flushable*.



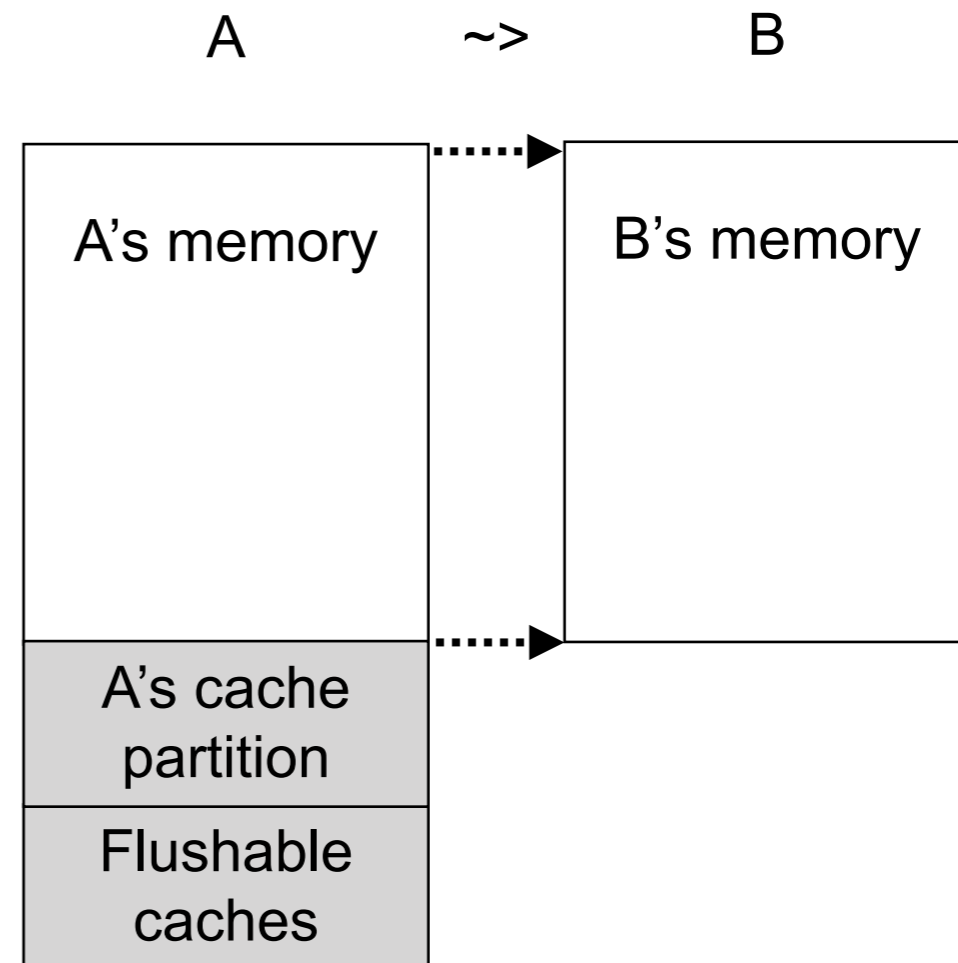
# Covert state: Partitionable vs flushable



## Principle:

**Model channels as *state elements*  
by their *elimination strategy*  
as per *HW-SW contract***

- Strategy for OS:  
*Partition or flush state; pad time.*
- Relies on HW-SW contract:
  - State: Everything must be *partitionable* or *flushable*.
    - e.g. Off-core vs on-core caches.



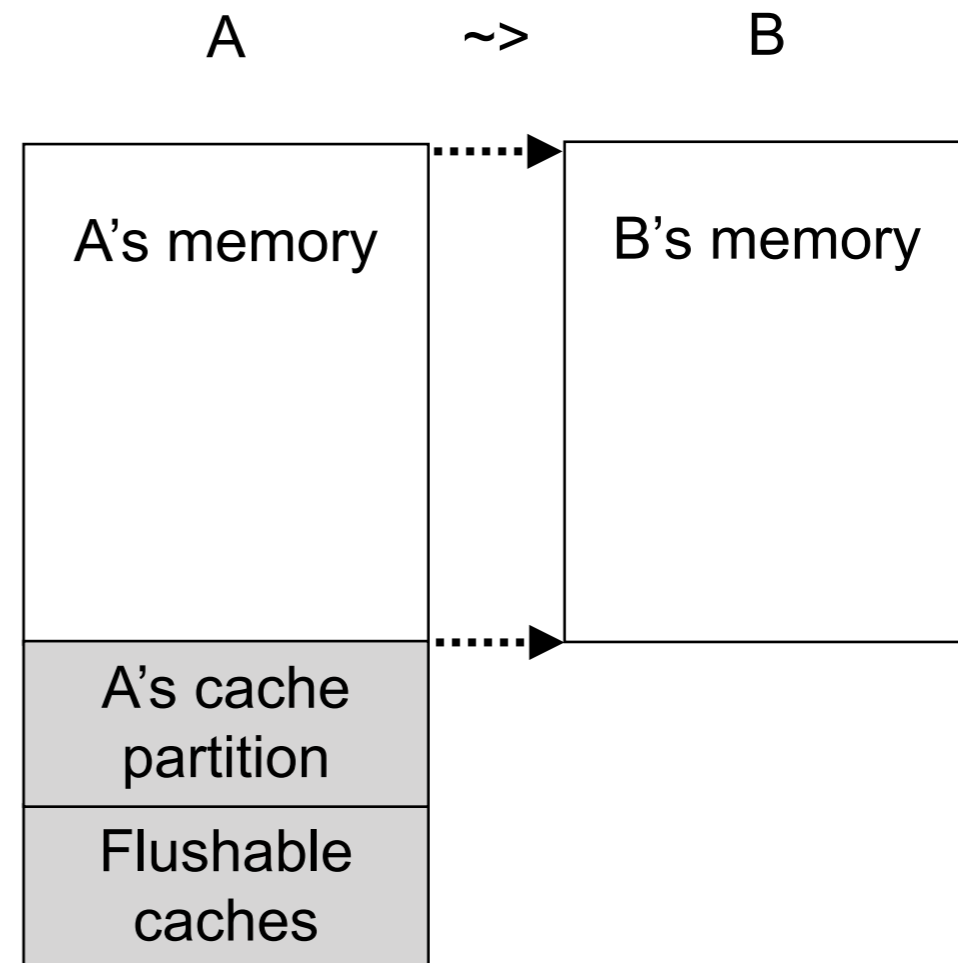
# Covert state: Partitionable vs flushable



## Principle:

**Model channels as *state elements*  
by their *elimination strategy*  
as per *HW-SW contract***

- Strategy for OS:  
*Partition or flush state; pad time.*
- Relies on HW-SW contract:
  - State: Everything must be *partitionable* or *flushable*.
    - e.g. Off-core vs on-core caches.
    - Interrupt-generating devices (partitionable; not pictured).





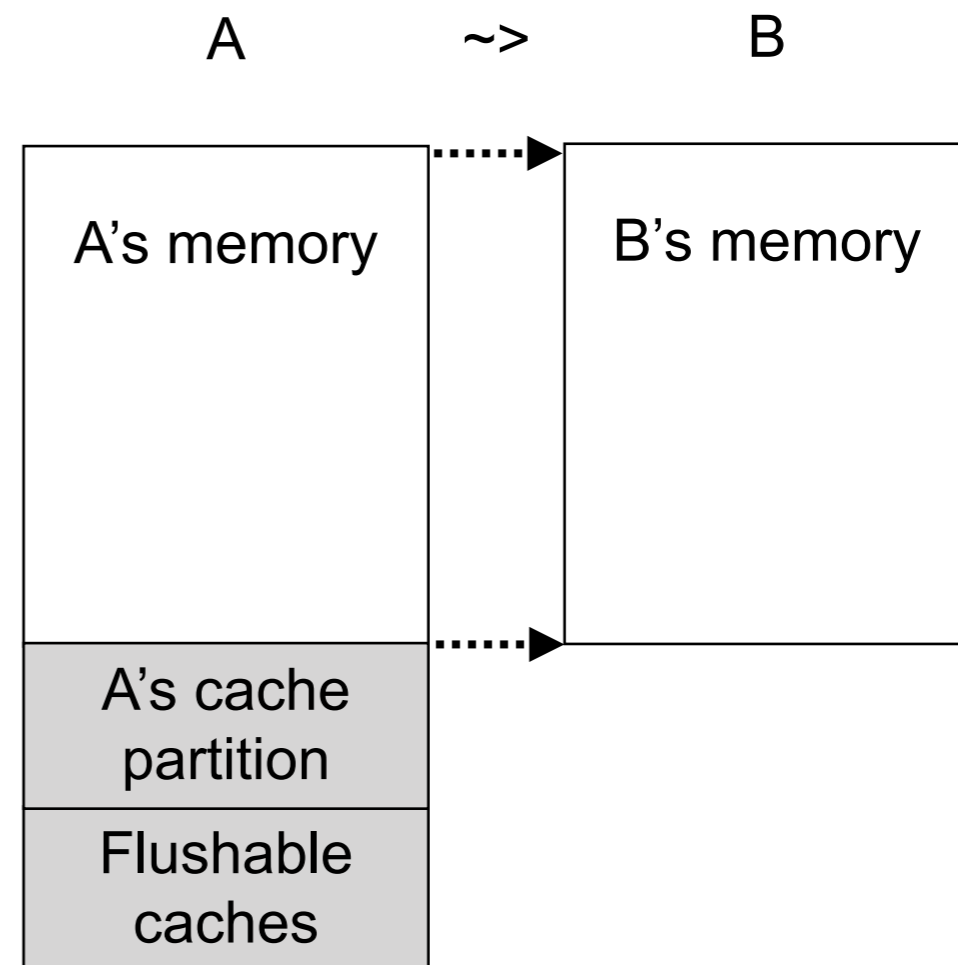
# Covert state: Partitionable vs flushable



## Principle:

**Model channels as *state elements*  
by their *elimination strategy*  
as per *HW-SW contract***

- Strategy for OS:  
*Partition or flush state; pad time.*
- Relies on HW-SW contract:
  - State: Everything must be *partitionable* or *flushable*.
    - e.g. Off-core vs on-core caches.
    - Interrupt-generating devices (partitionable; not pictured).
  - Time: HW must give reliable
    - WCETs (worst-case execution times)



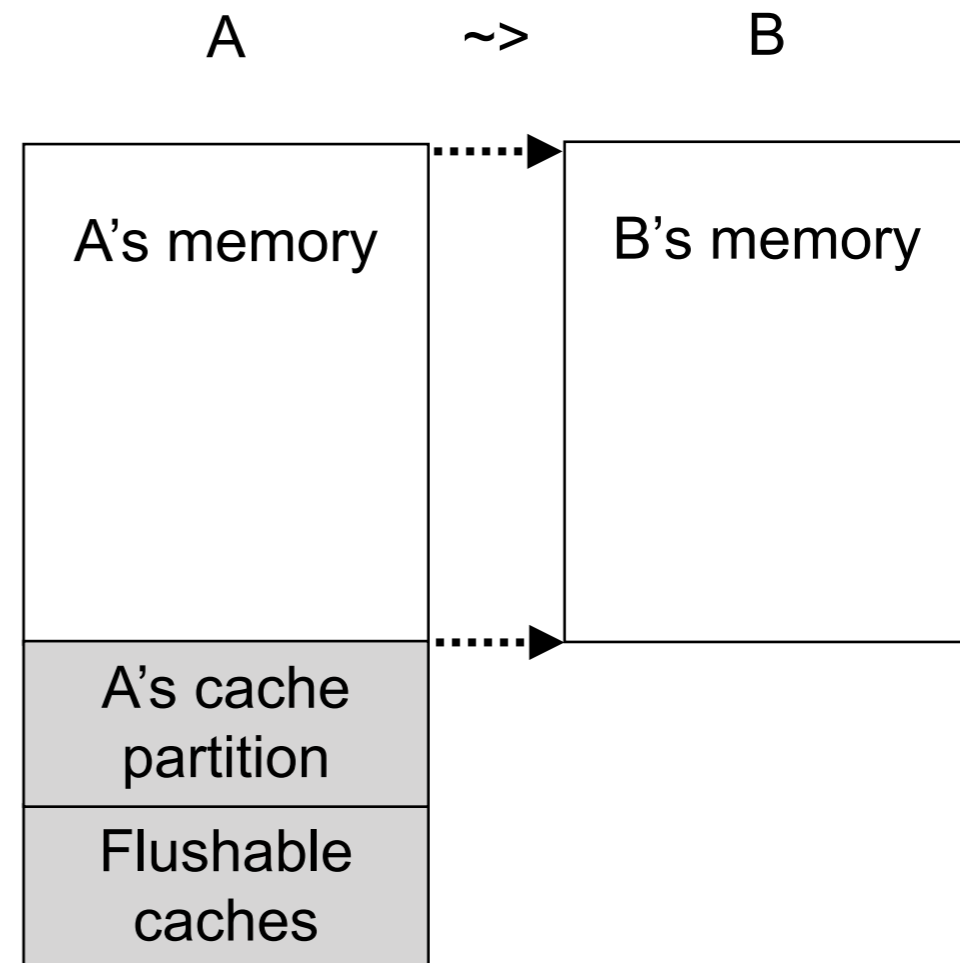
# Covert state: Partitionable vs flushable



## Principle:

**Model channels as *state elements*  
by their *elimination strategy*  
as per *HW-SW contract***

- Strategy for OS:  
*Partition or flush state; pad time.*
- Relies on HW-SW contract:
  - State: Everything must be *partitionable* or *flushable*.
    - e.g. Off-core vs on-core caches.
    - Interrupt-generating devices (partitionable; not pictured).
  - Time: HW must give reliable
    - WCETs (worst-case execution times)
    - method of *padding*.



# How to formalise an OS enforces *time protection*?

Versus threat scenario:  
trojan and spy



Abstract *covert state* + *time* to reflect strategies enabled by HW:  
Partition or flush state; pad time.



Make security property precise enough to exclude flows from covert state.



Demonstrating these principles,  
we formalised in Isabelle/HOL:

1. OS security model imposing requirements on relevant parts of OS.  
(Intended for seL4, but *generic*)

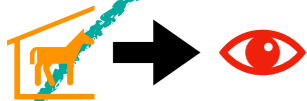
3. Proof our security property holds if OS model's requirements hold.

2. OS security property that is dynamic; this makes it observer relative.  
(Improving on seL4's of [Murray et al. 2012])

4. Basic instantiation of OS model exercising dynamic policy.

# How to formalise an OS enforces *time protection*?

Versus threat scenario:  
trojan and spy



Abstract *covert state* + *time* to reflect strategies enabled by HW:  
Partition or flush state; pad time.



Demonstrating these principles,  
we formalised in Isabelle/HOL:

1. OS security model imposing requirements on relevant parts of OS.  
(Intended for seL4, but *generic*)
3. Proof our security property holds if OS model's requirements hold.



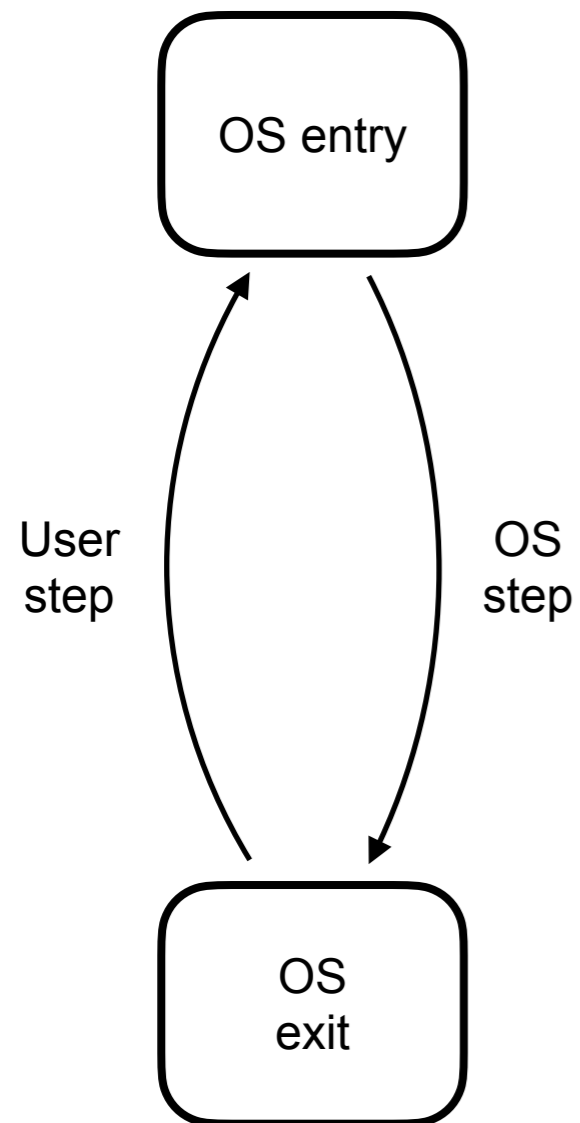
Make security property precise enough to exclude flows from covert state.

2. OS security property that is dynamic; this makes it observer relative.  
(Improving on seL4's of [Murray et al. 2012])
4. Basic instantiation of OS model exercising dynamic policy.

# OS security model



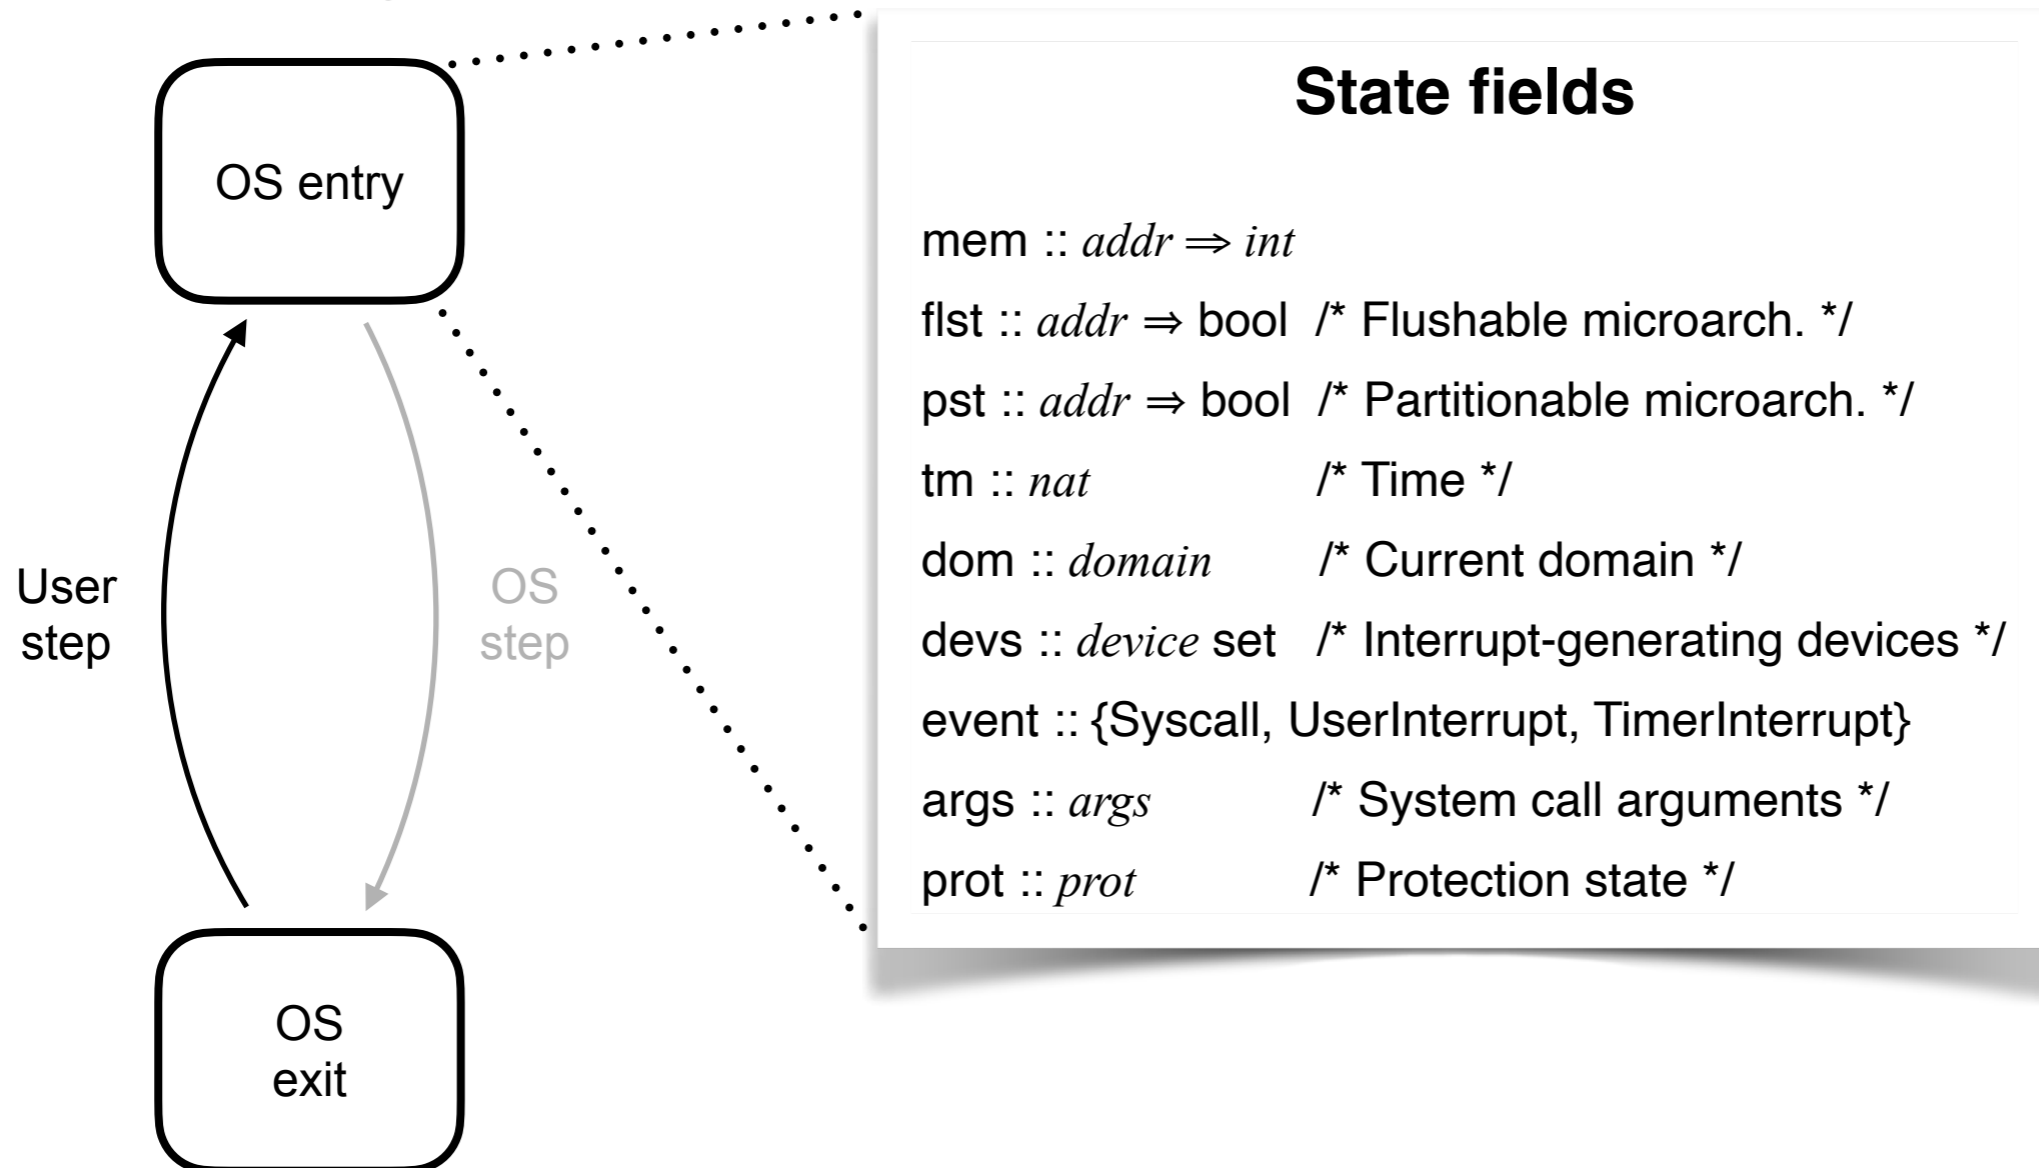
## Transition system



# OS security model



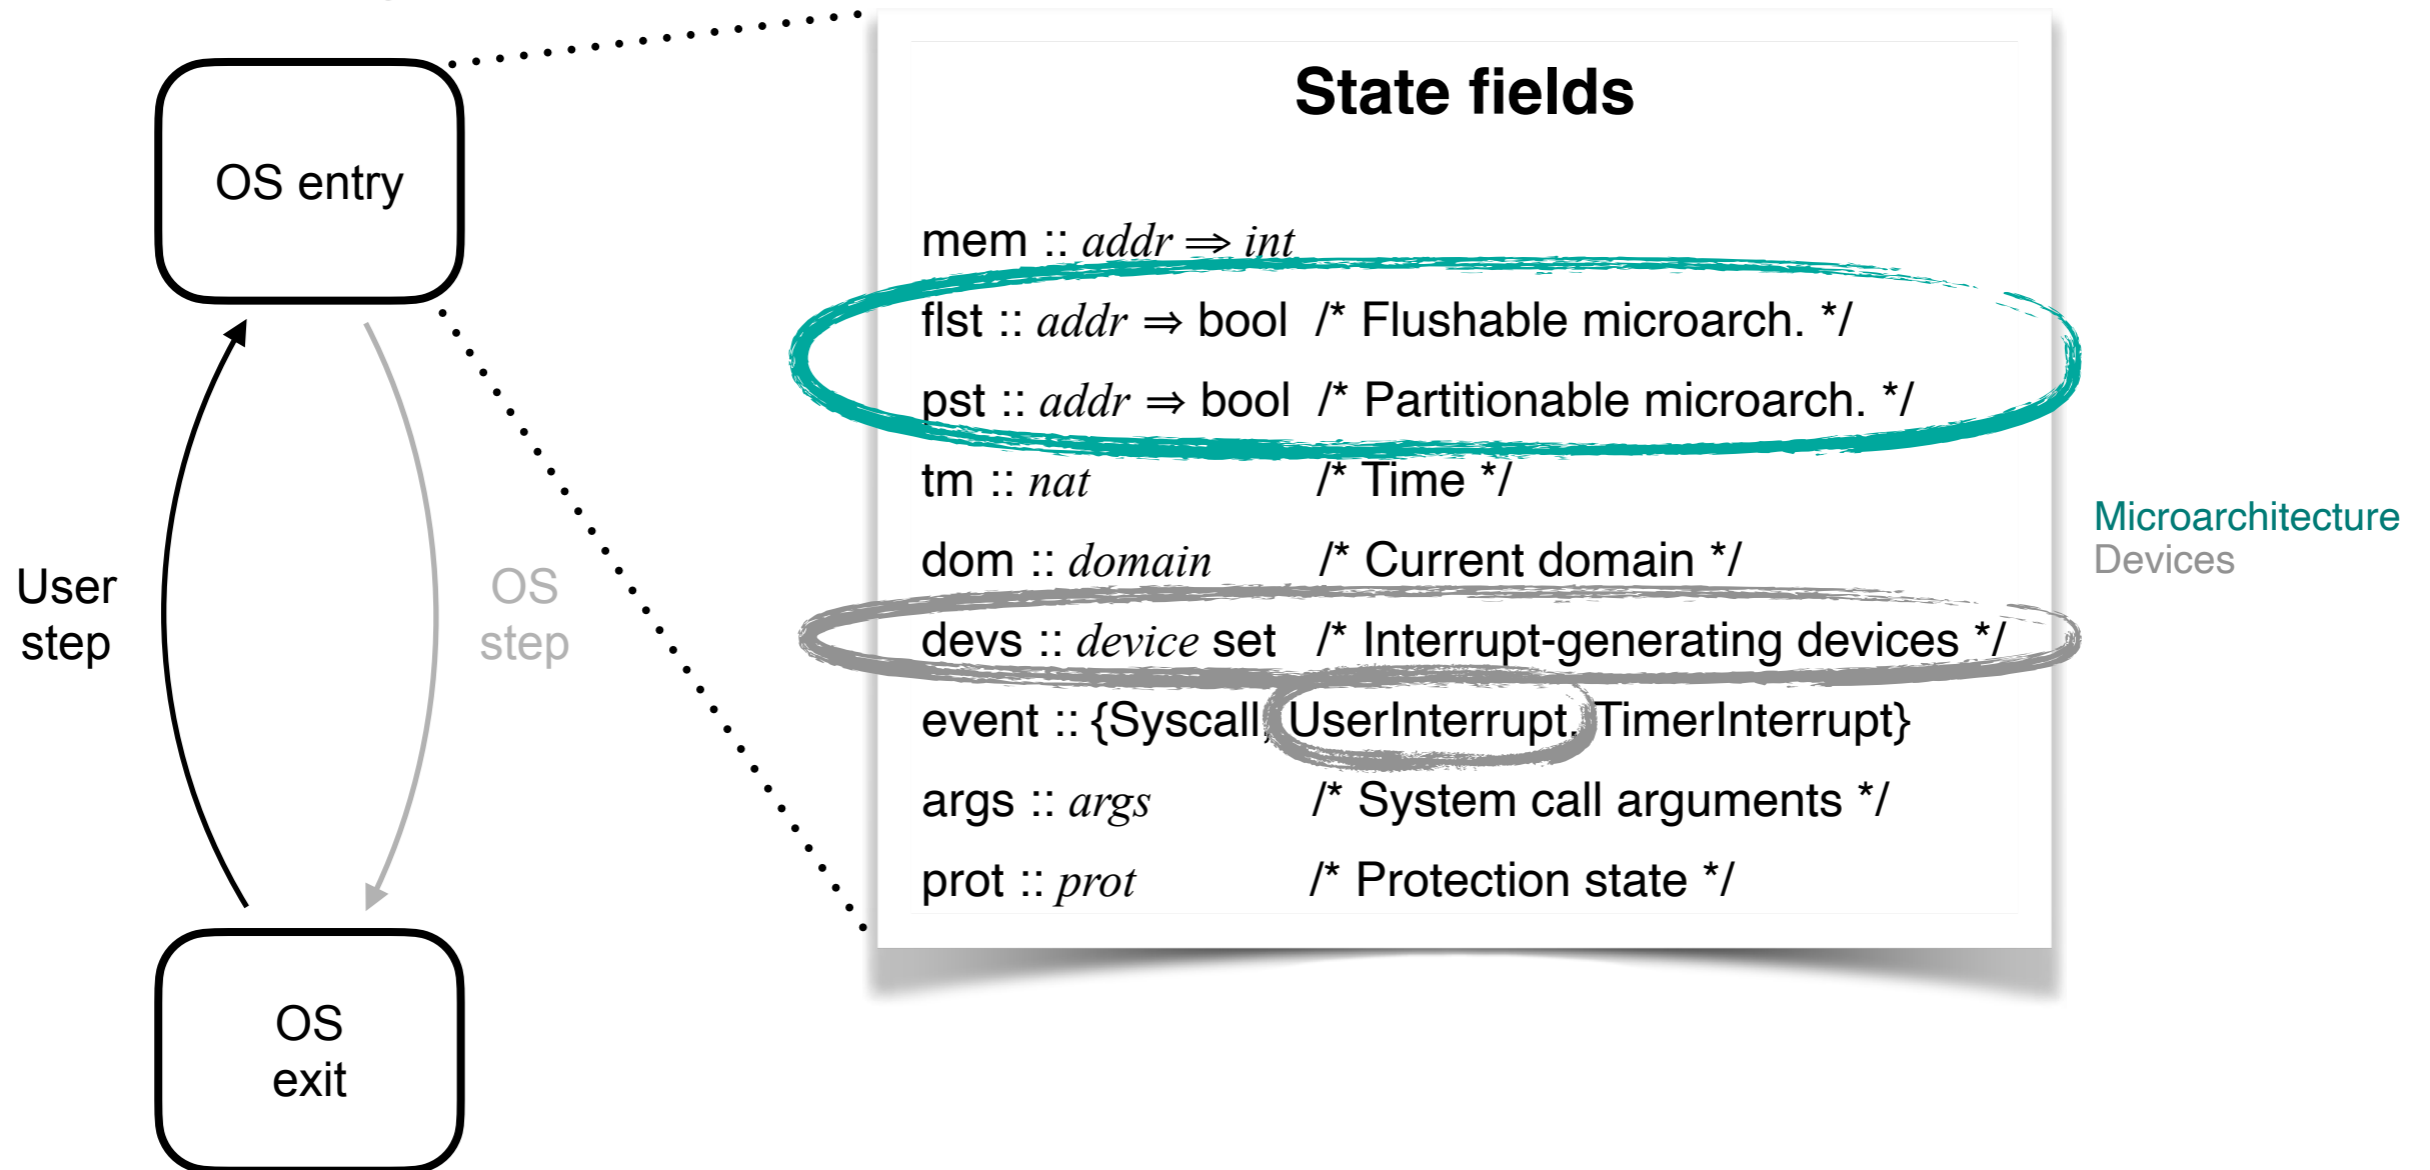
## Transition system



# OS security model



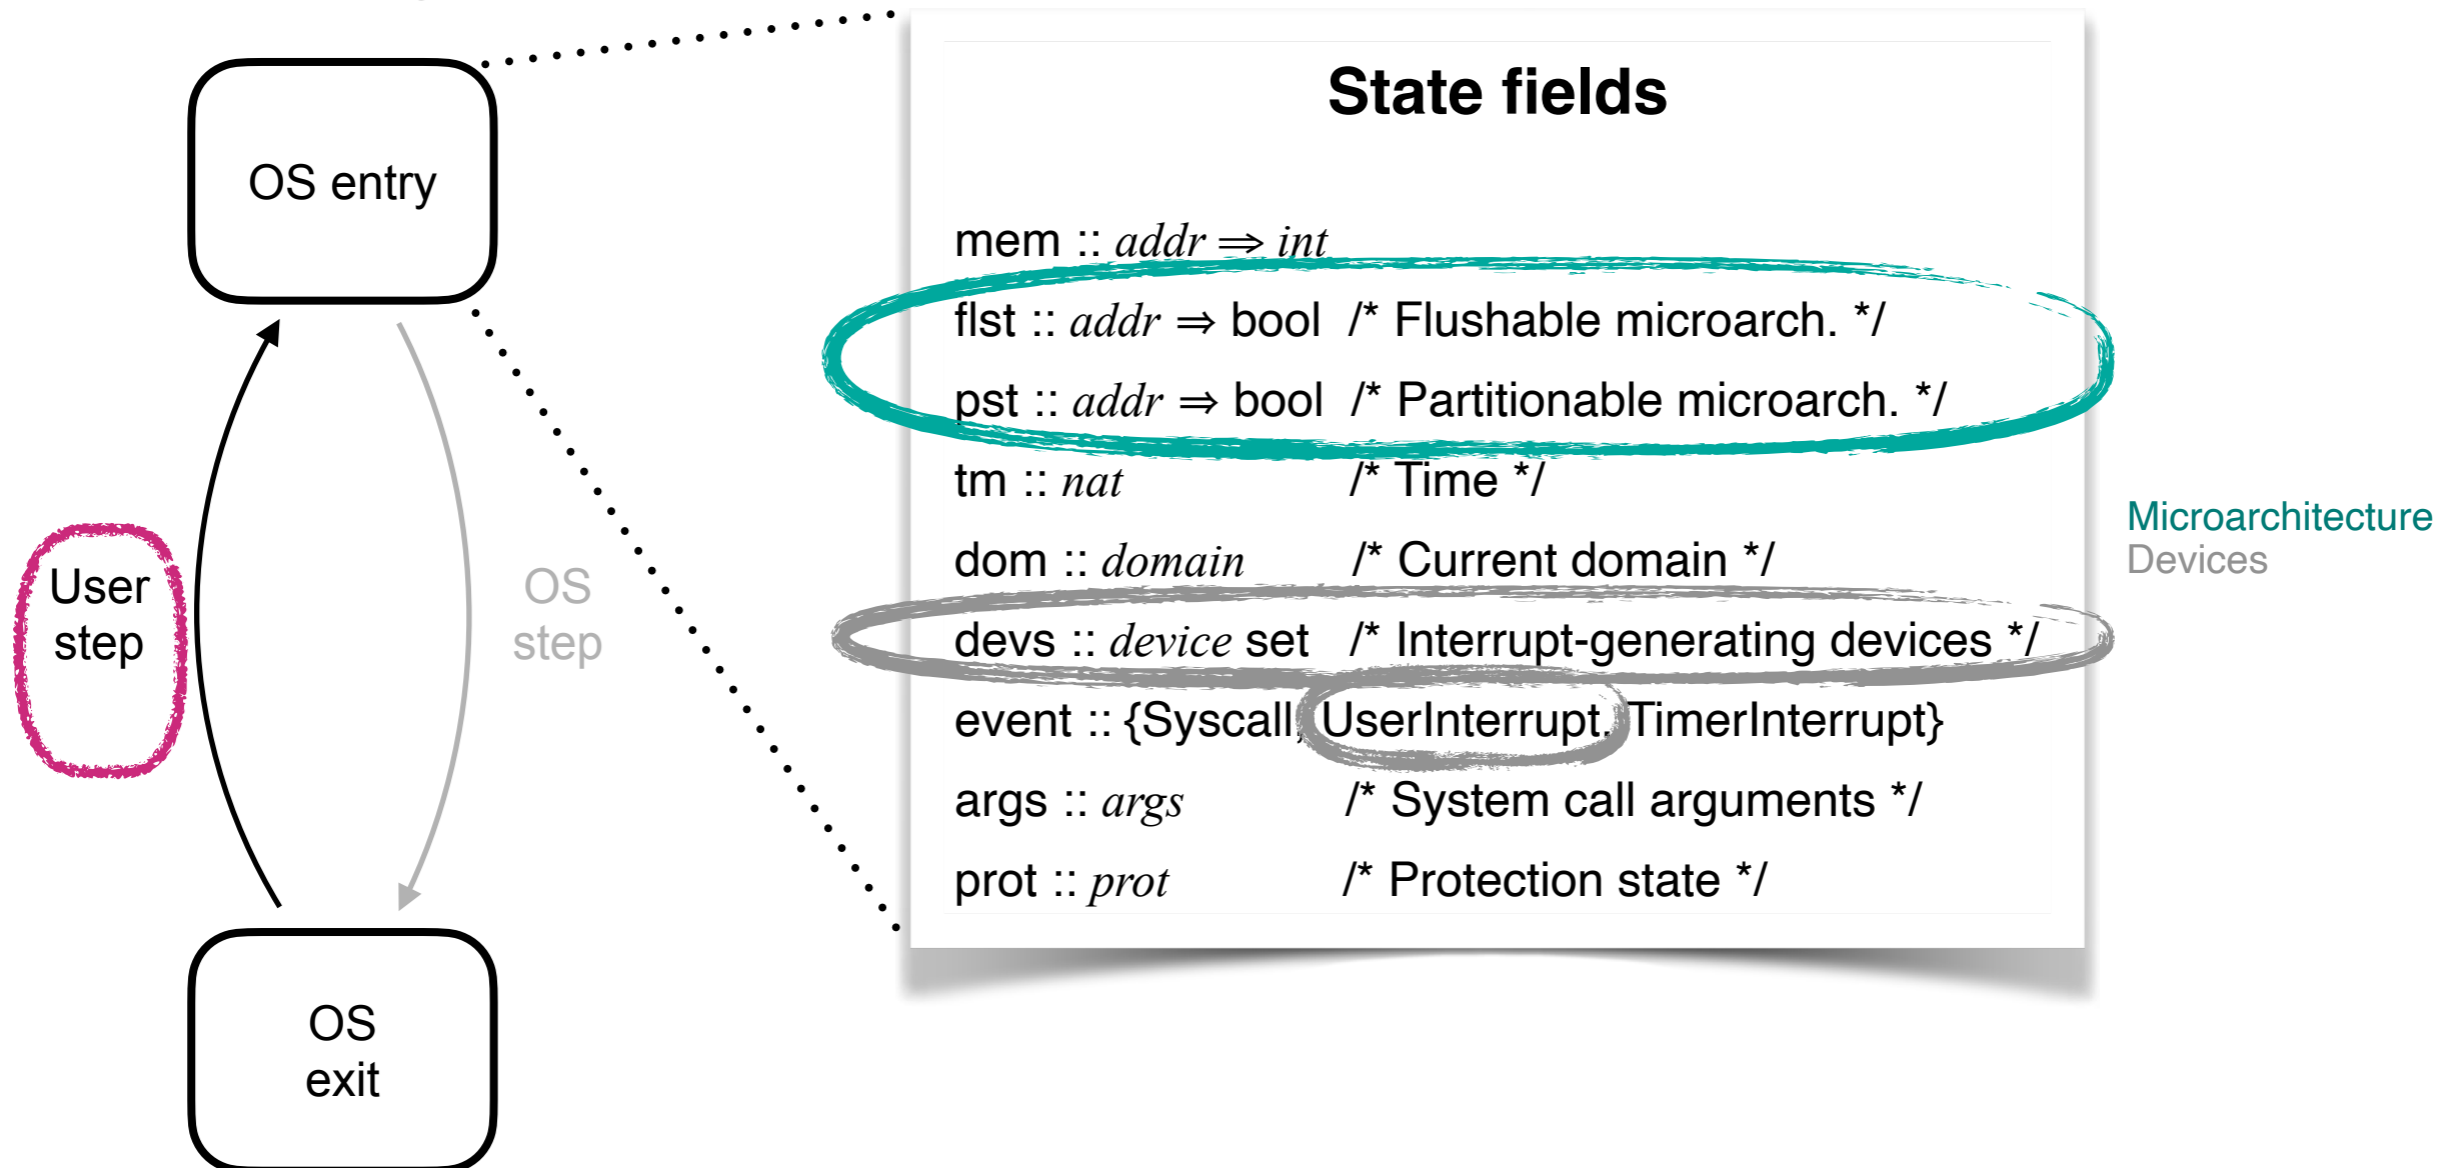
## Transition system



# OS security model



## Transition system

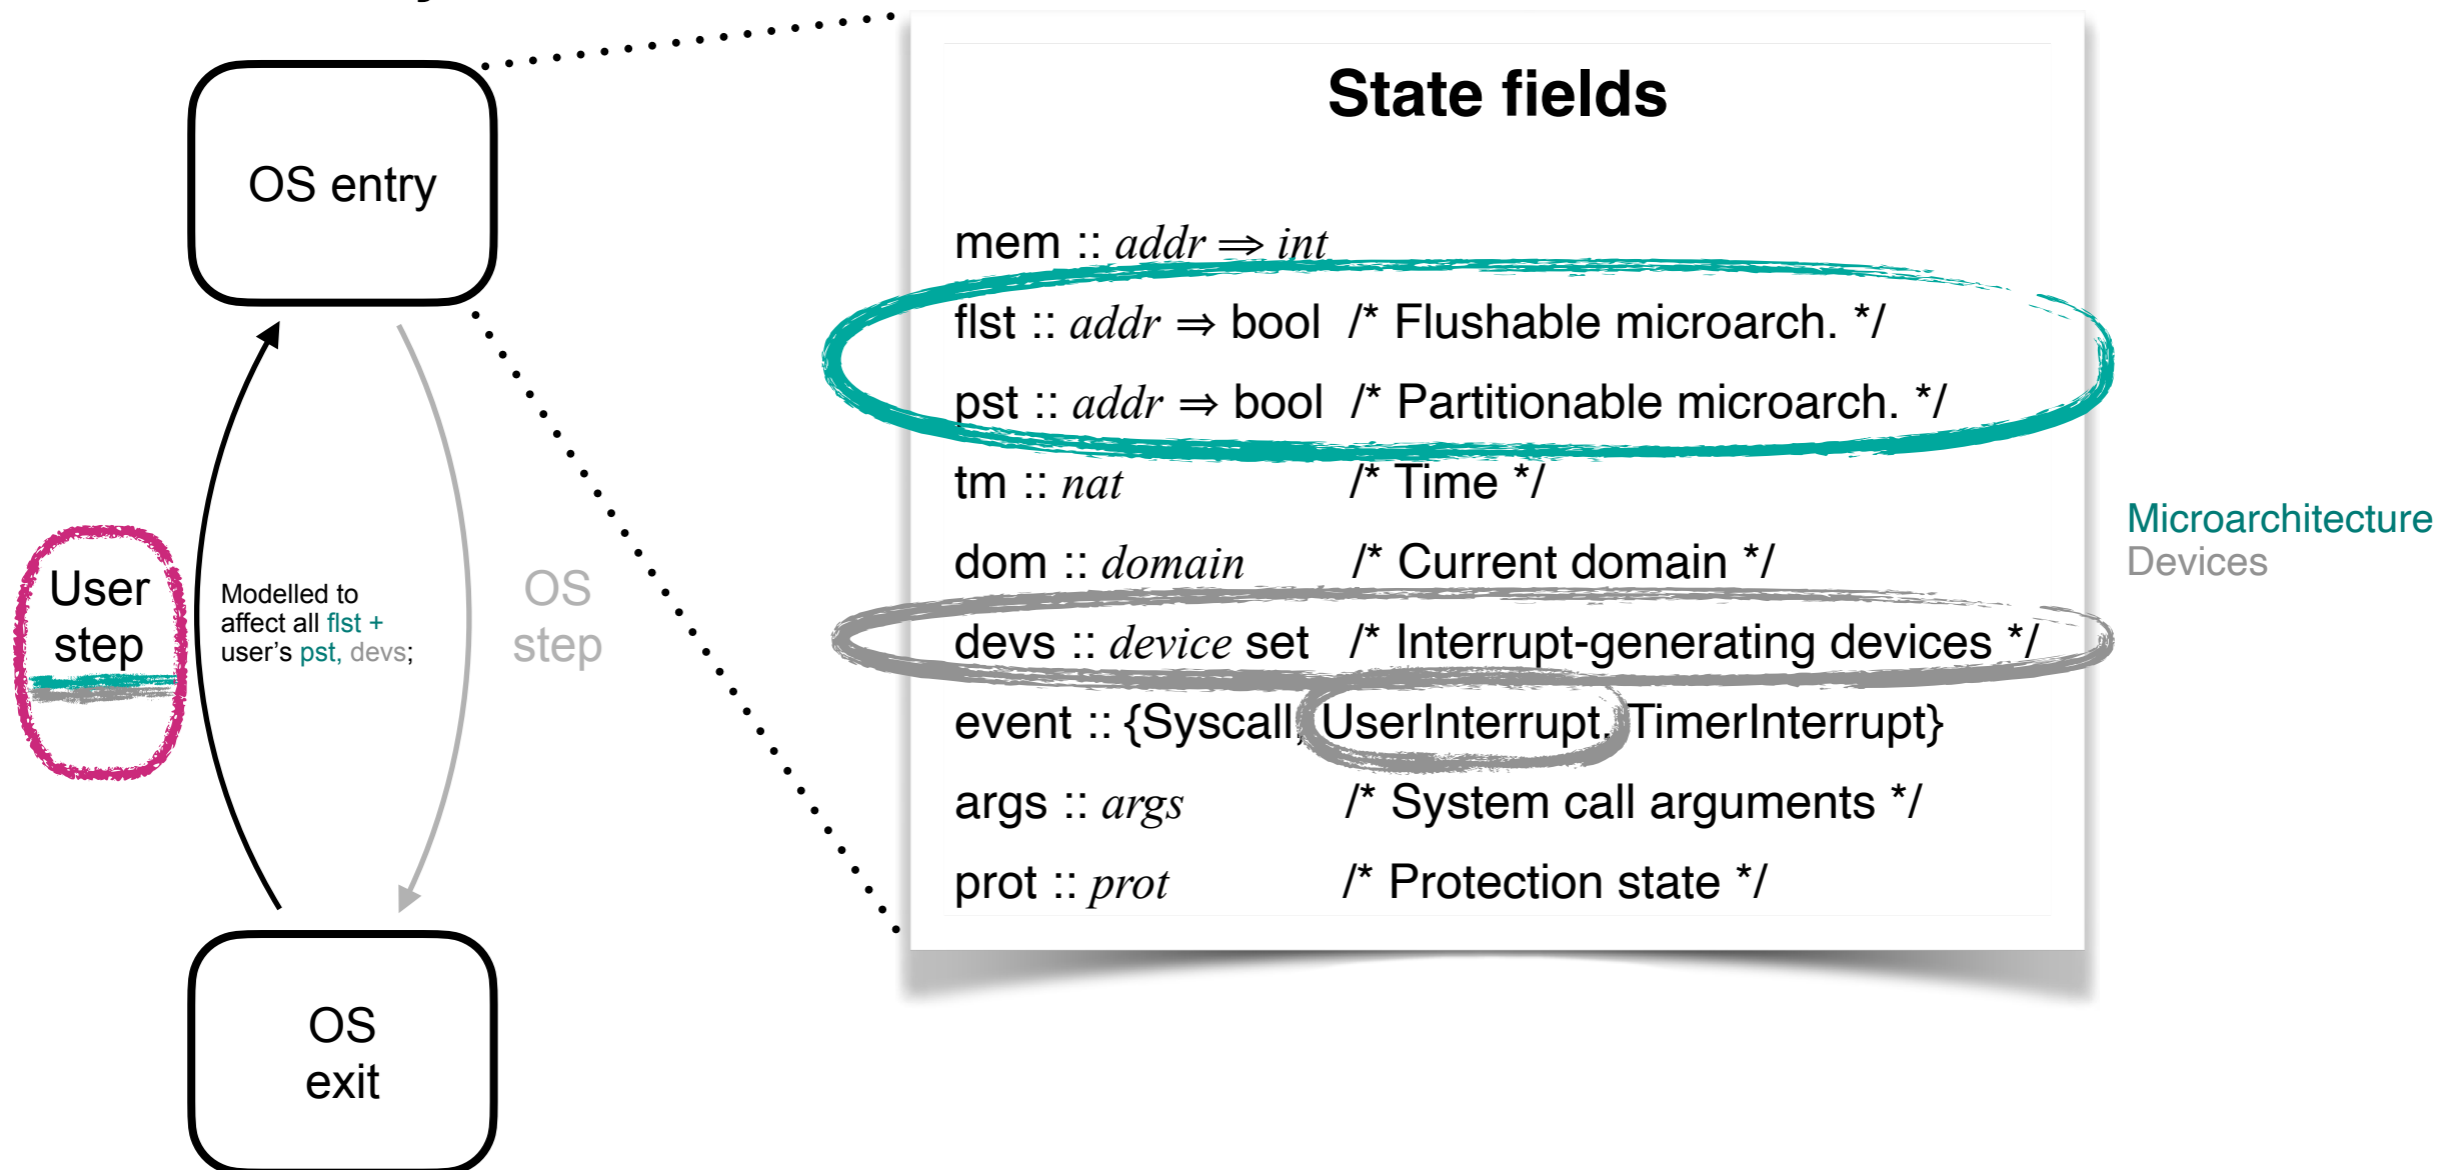




# OS security model



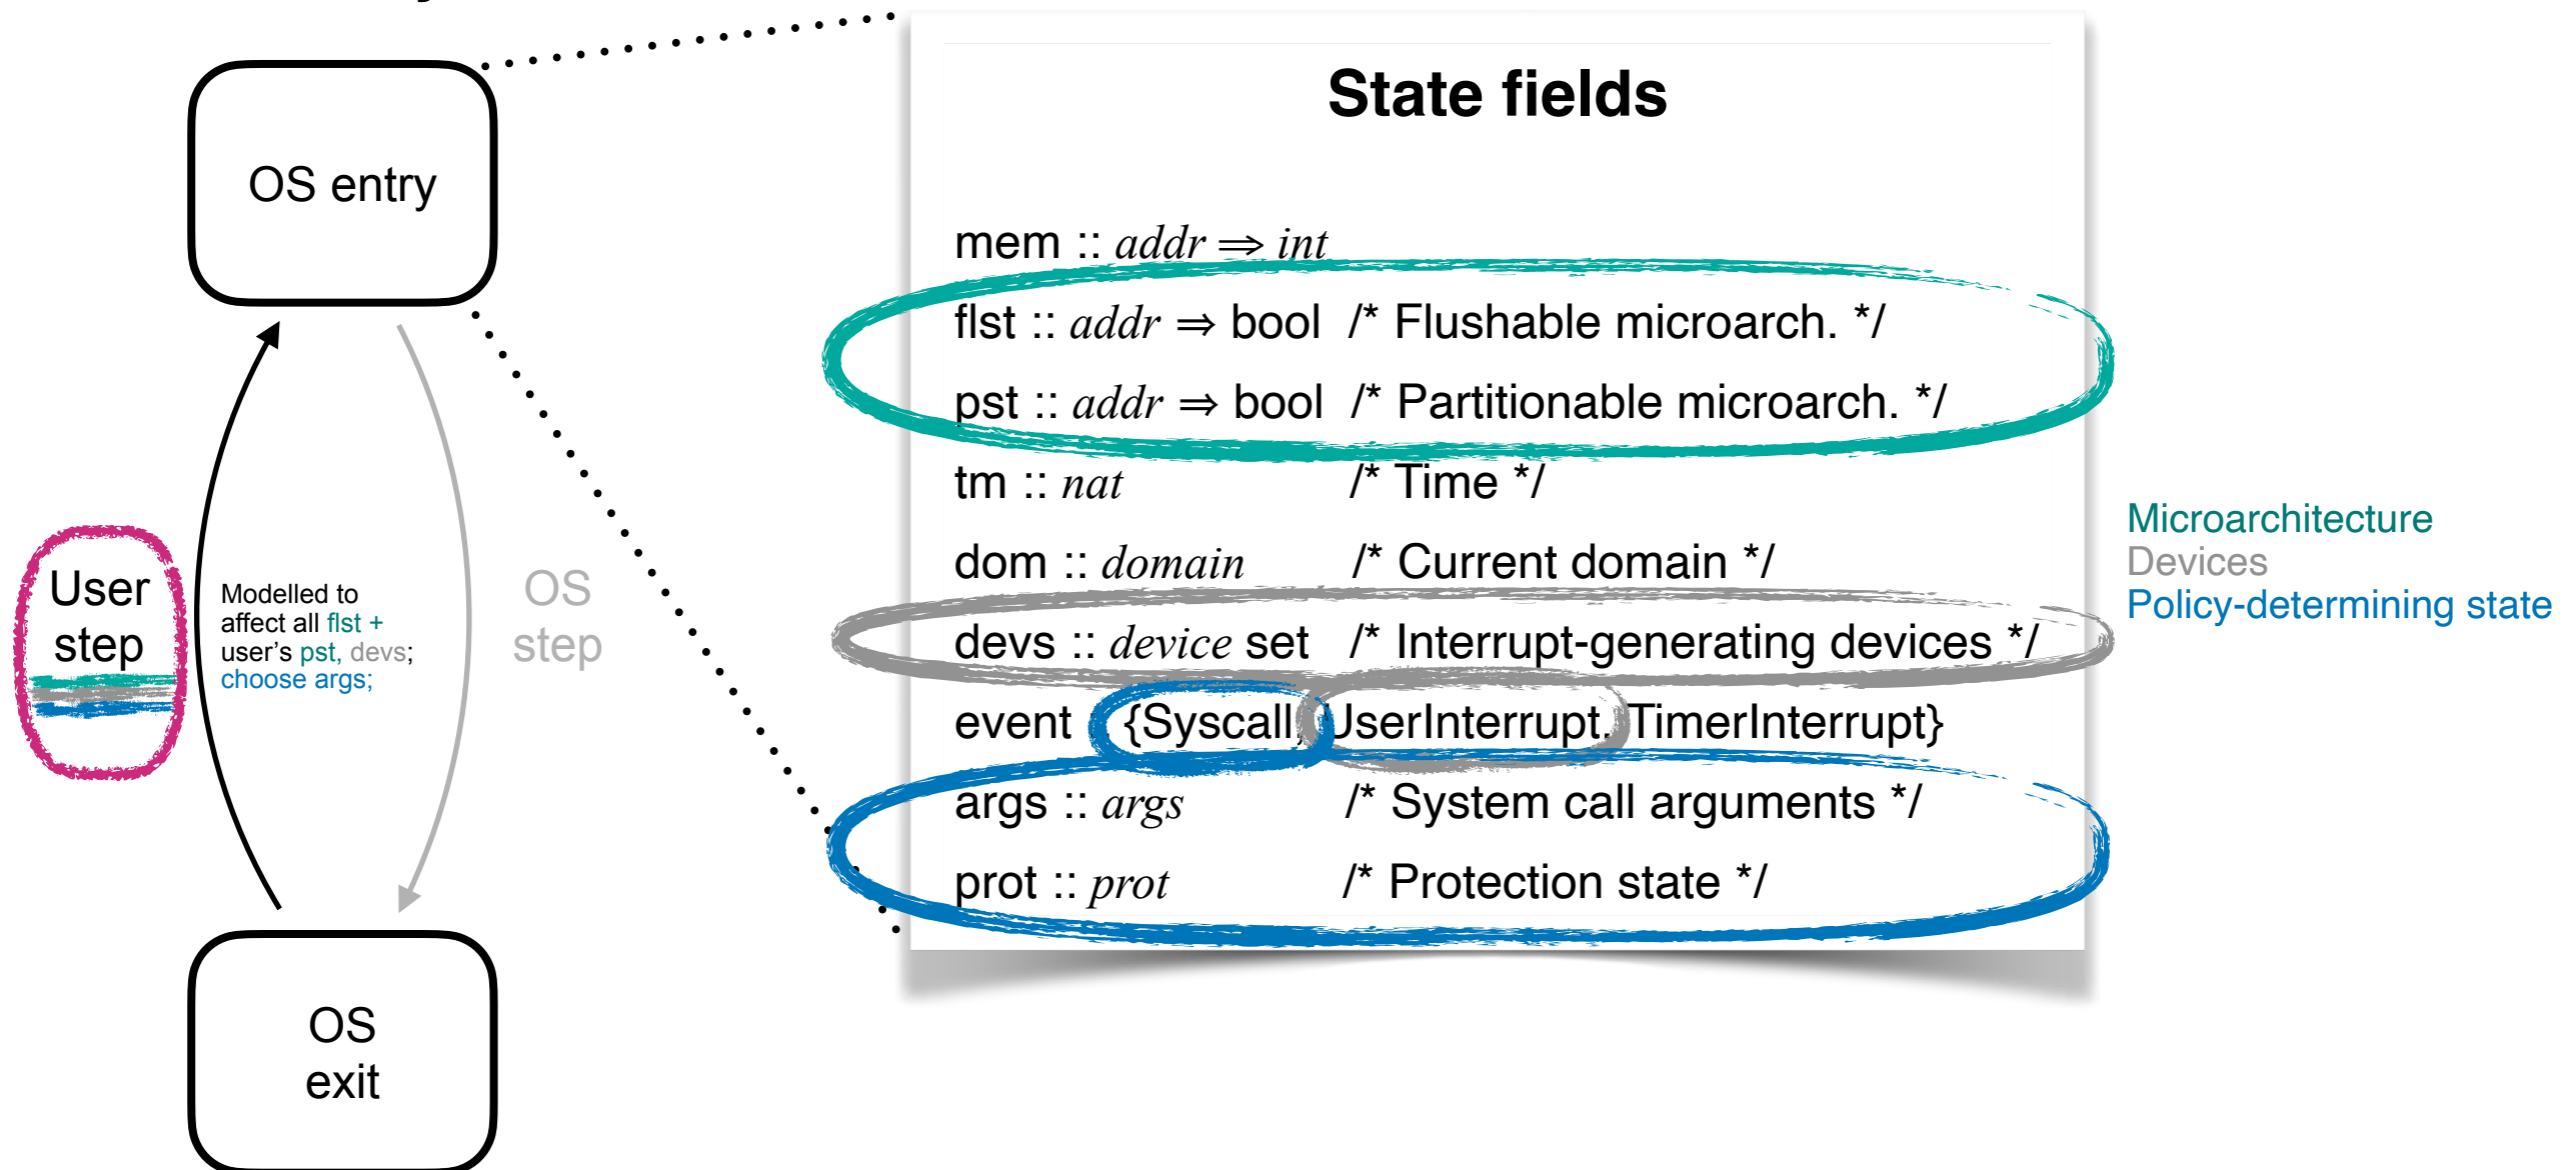
## Transition system



# OS security model



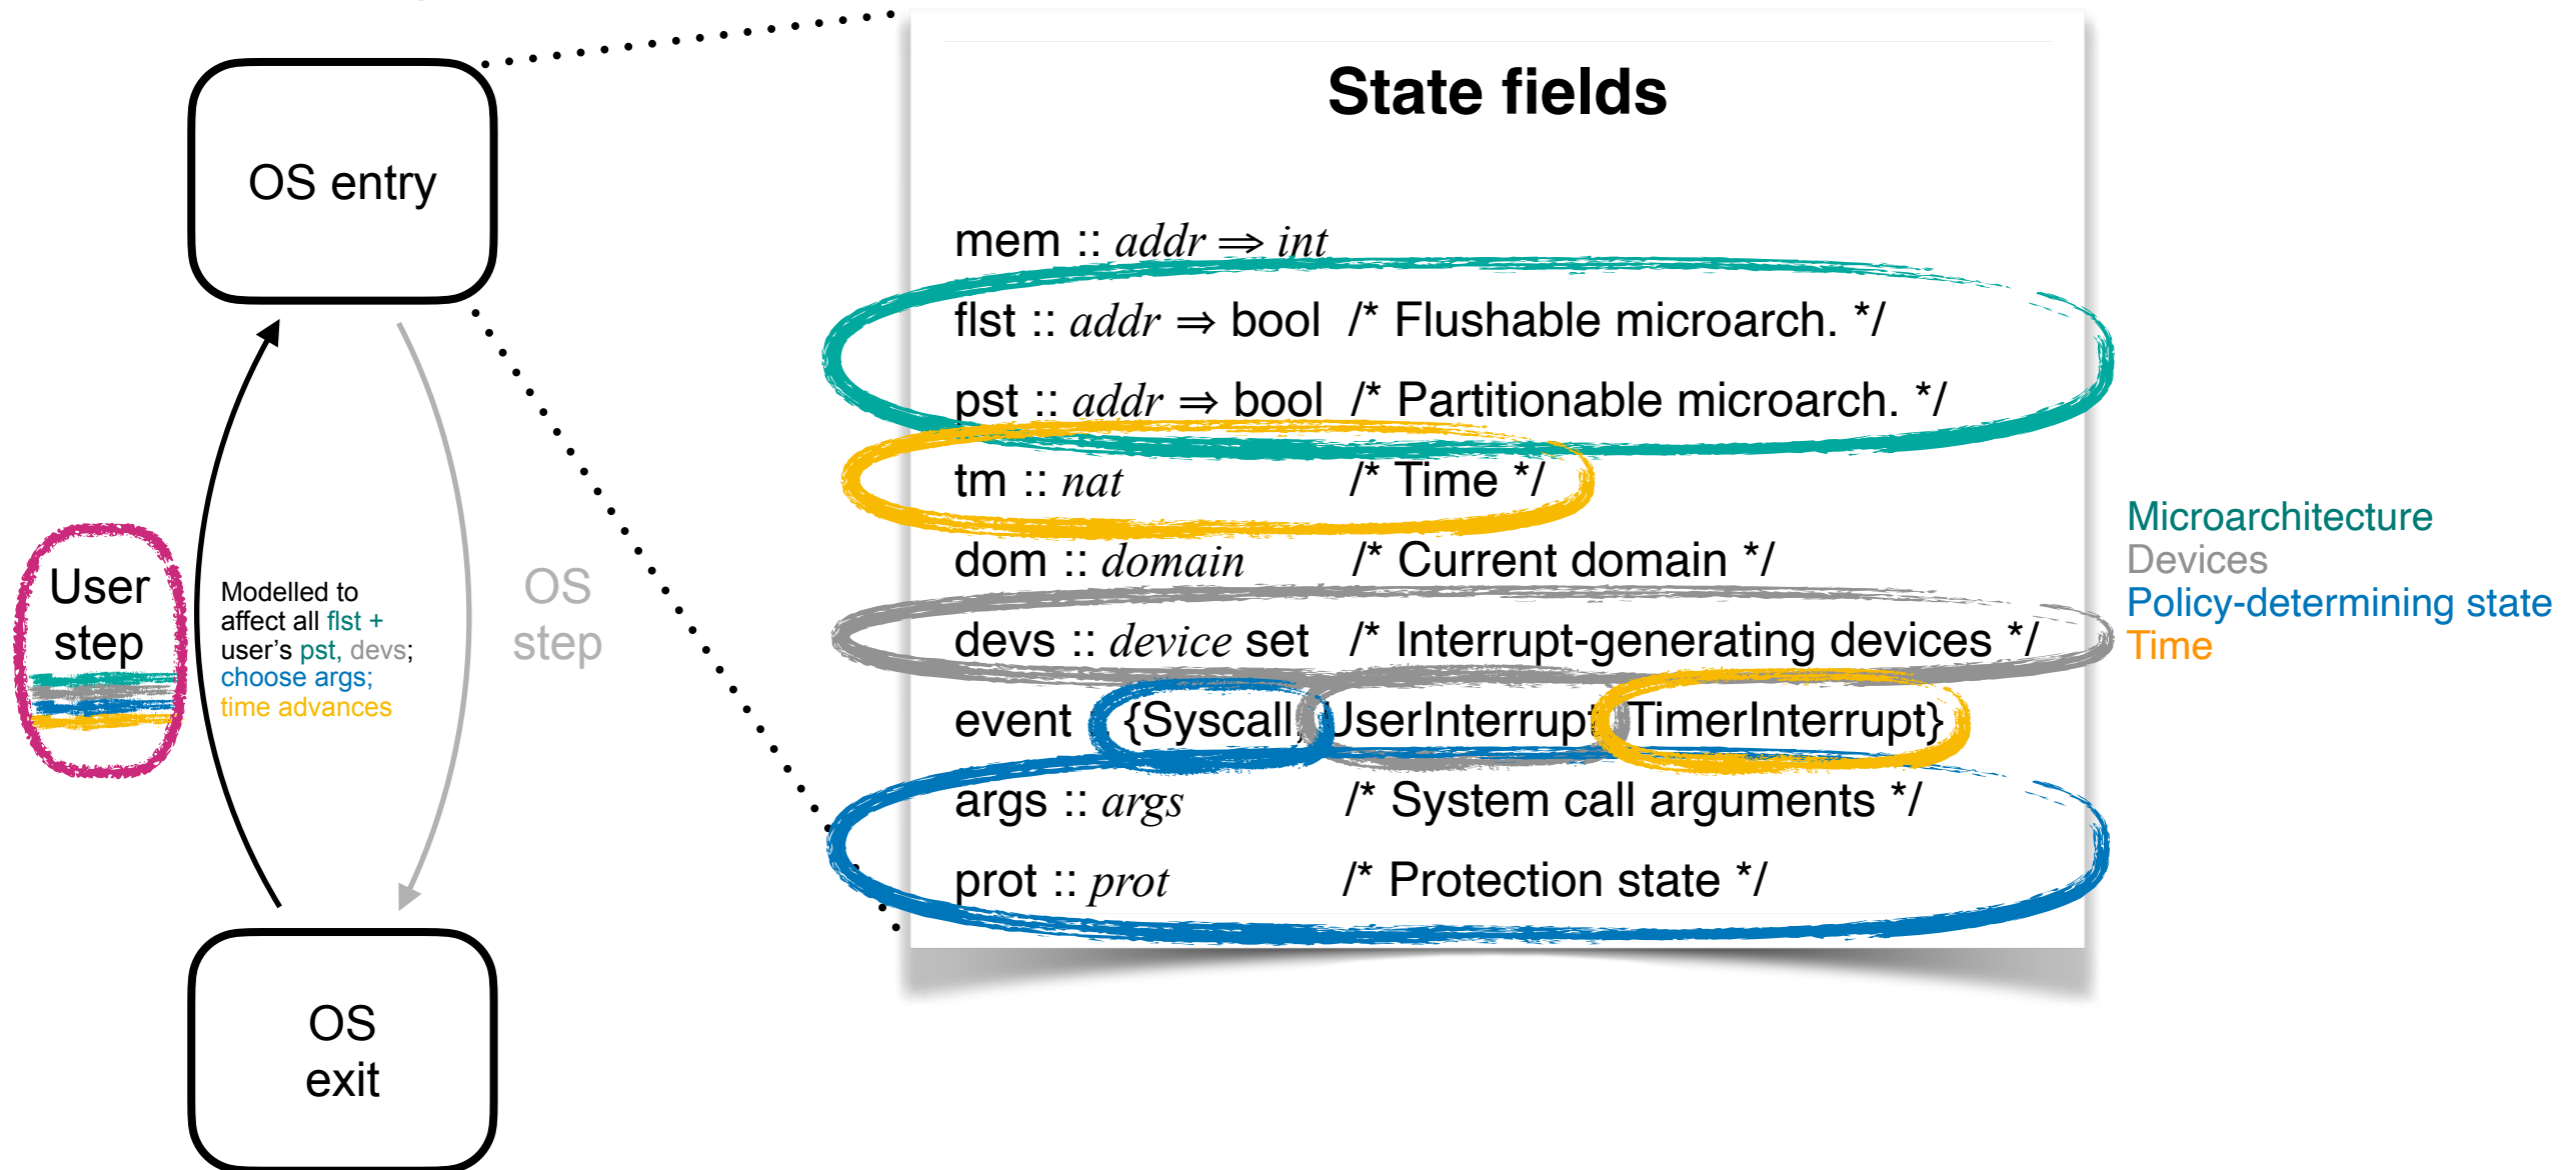
## Transition system



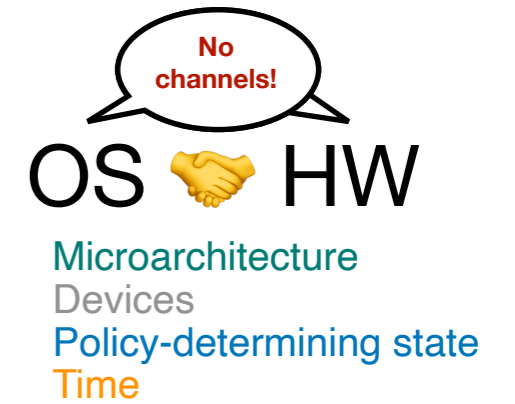
# OS security model



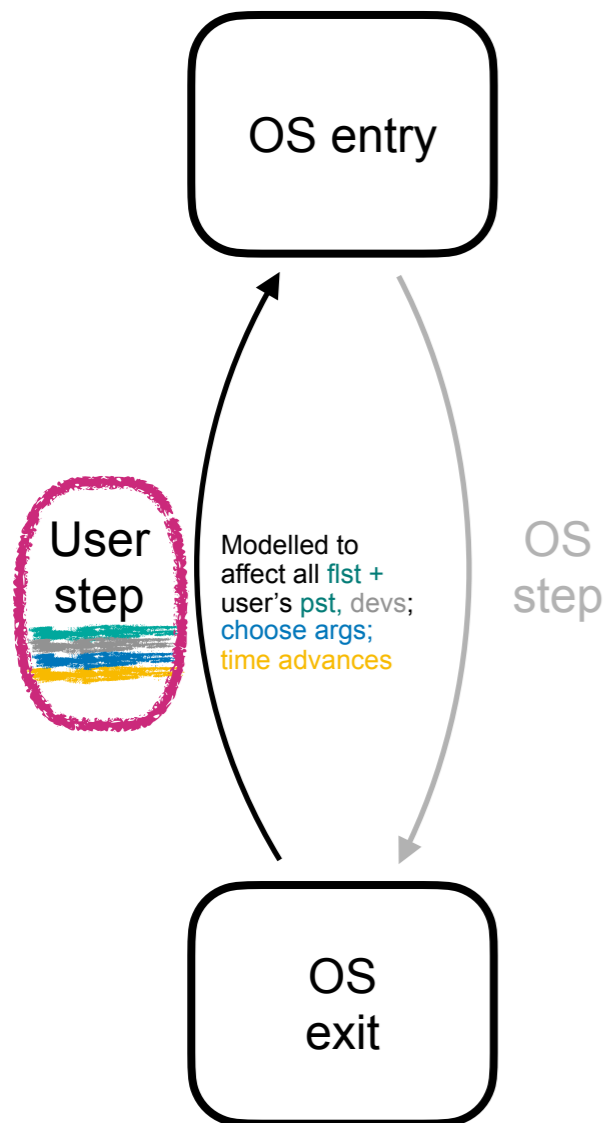
## Transition system



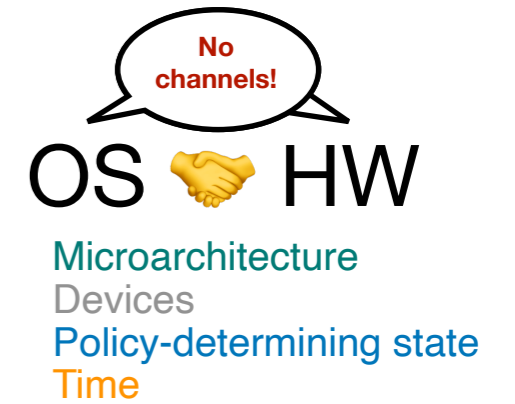
# OS security model



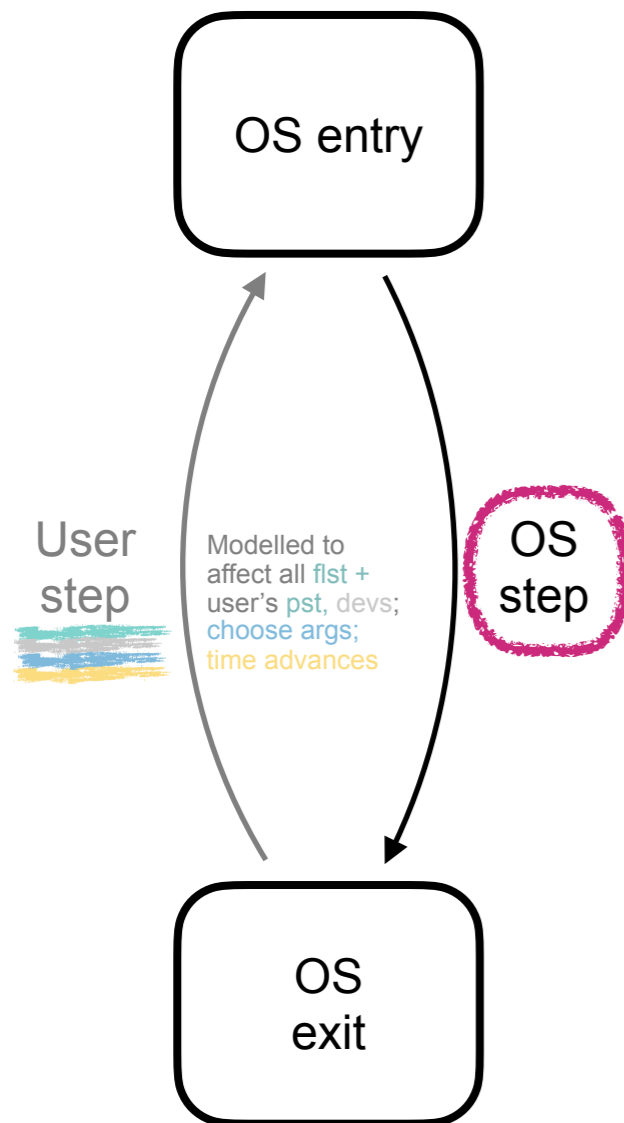
## Transition system



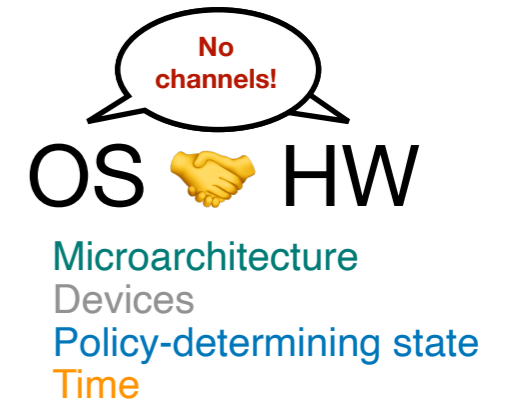
# OS security model



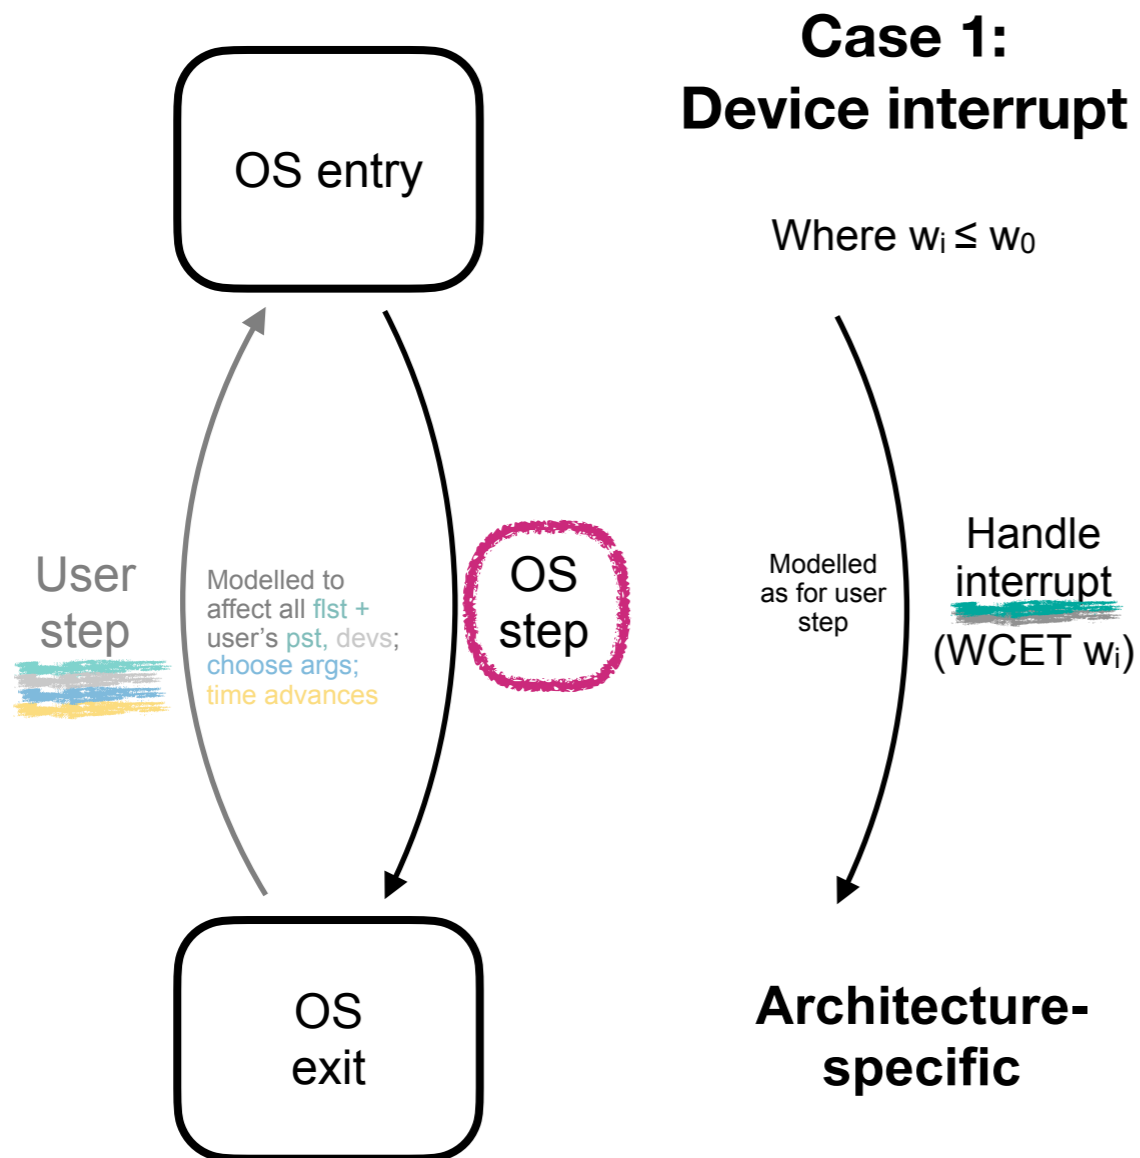
## Transition system



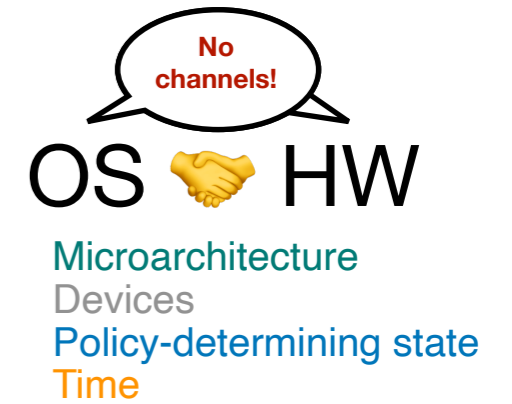
# OS security model



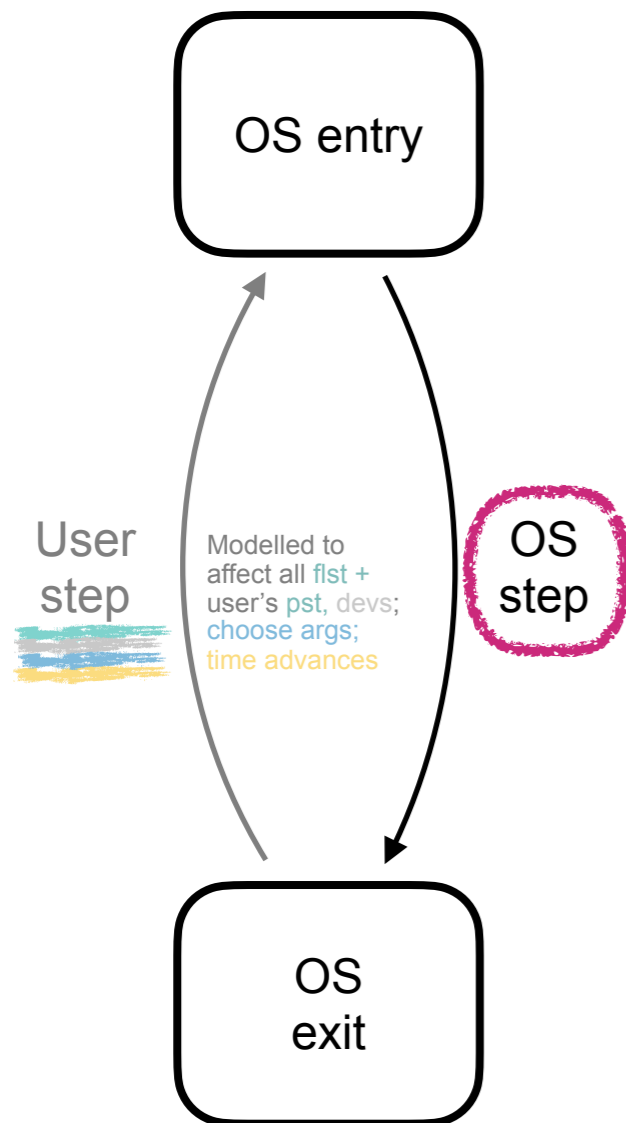
## Transition system



# OS security model

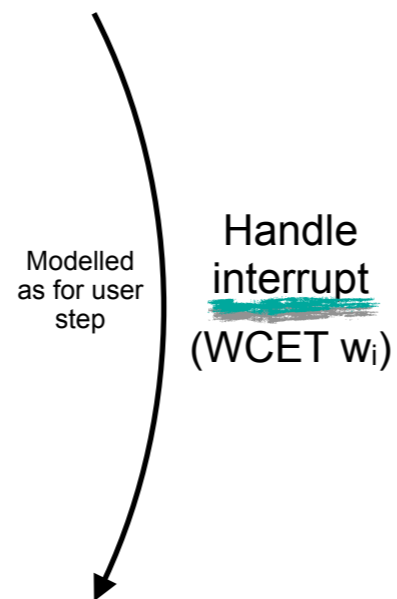


## Transition system



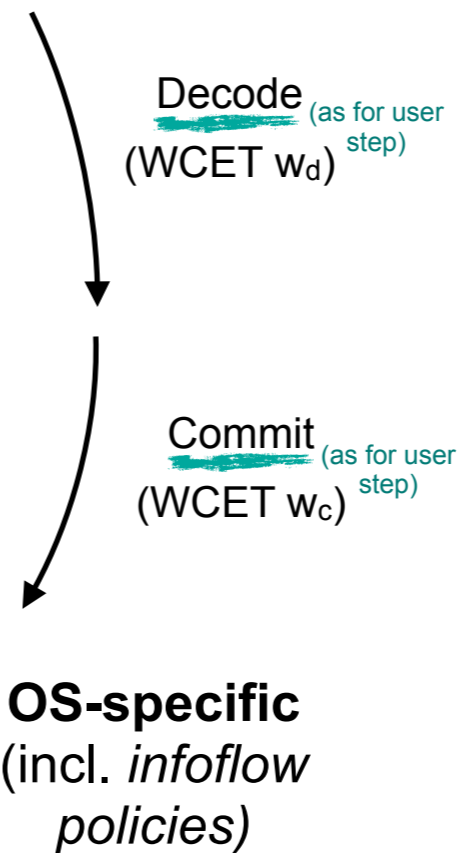
### Case 1: Device interrupt

Where  $w_i \leq w_0$

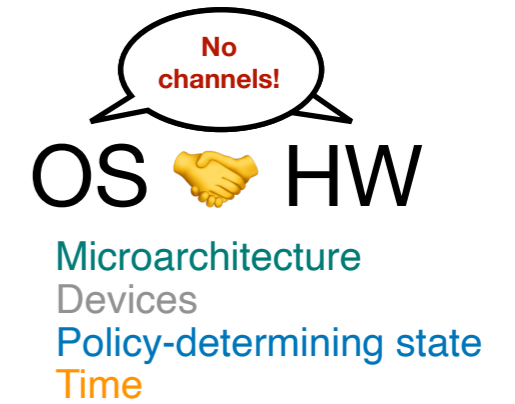


### Case 2: System call

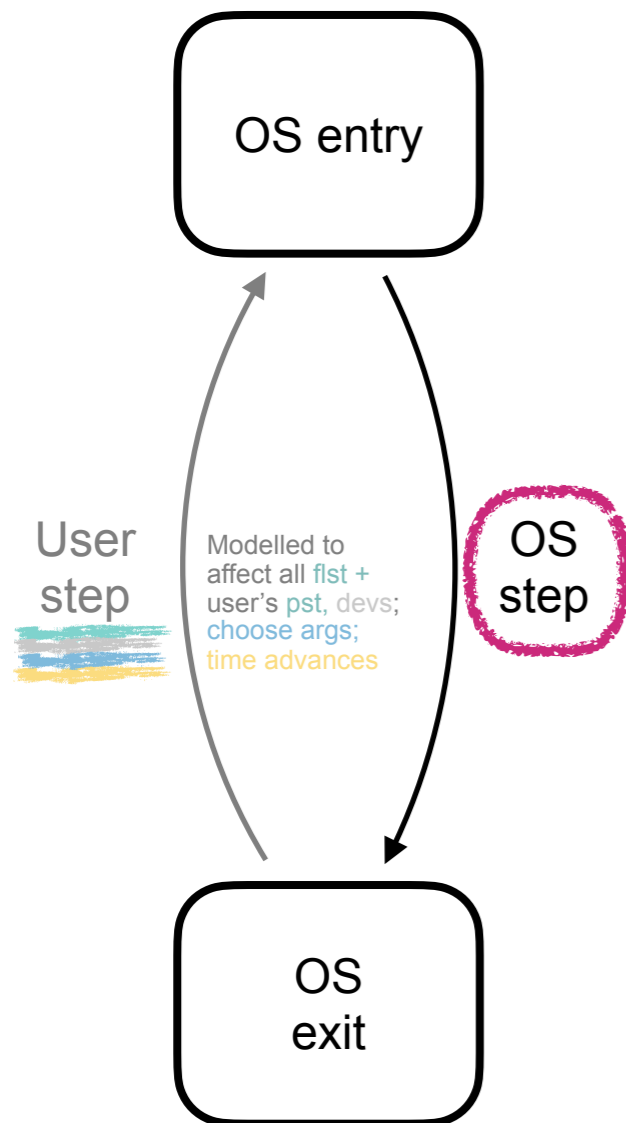
Where  $w_d + w_c \leq w_0$



# OS security model

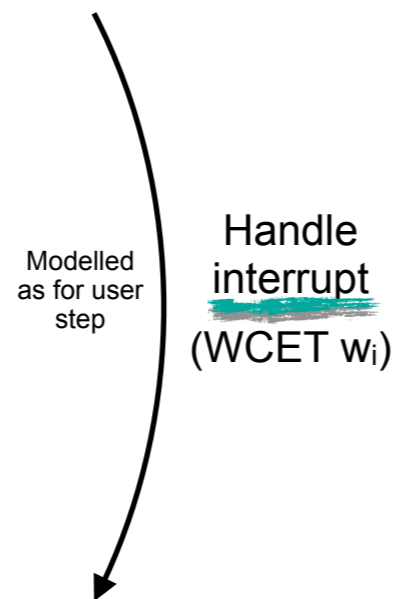


## Transition system



### Case 1: Device interrupt

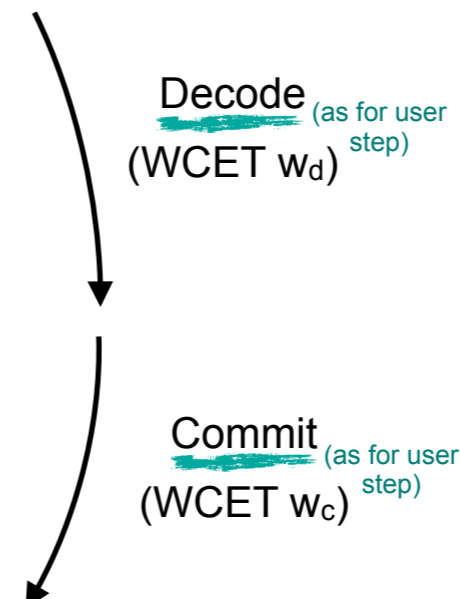
Where  $w_i \leq w_0$



Architecture-specific

### Case 2: System call

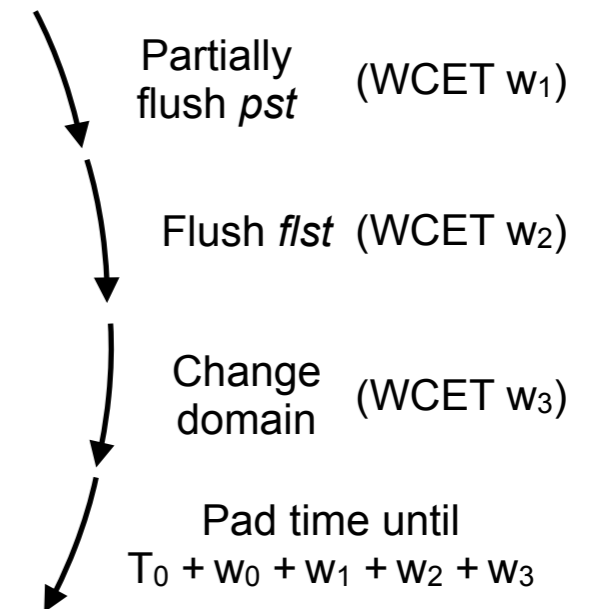
Where  $w_d + w_c \leq w_0$



OS-specific (incl. *inflow policies*)

### Case 3: Domain switch

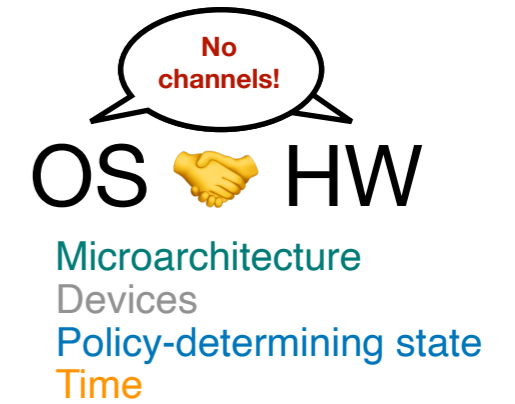
Timer interrupt delivered at (worst-case)  $T_0 + w_0$



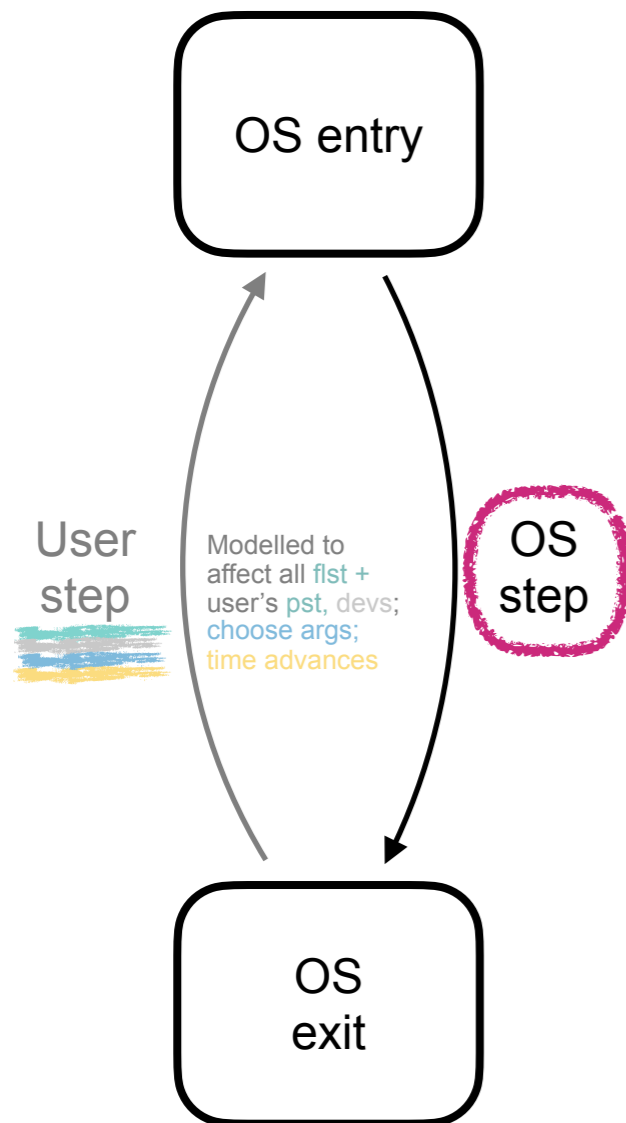
Architecture-specific



# OS security model

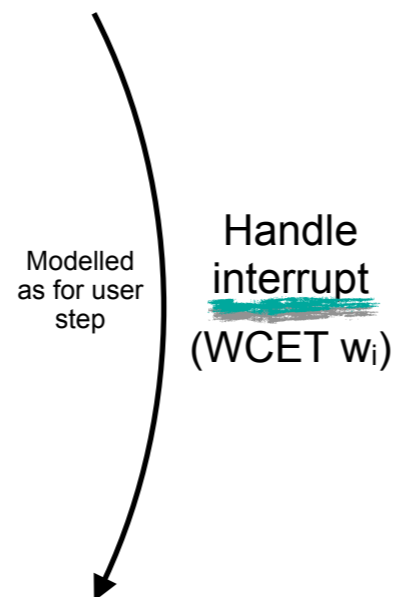


## Transition system



### Case 1: Device interrupt

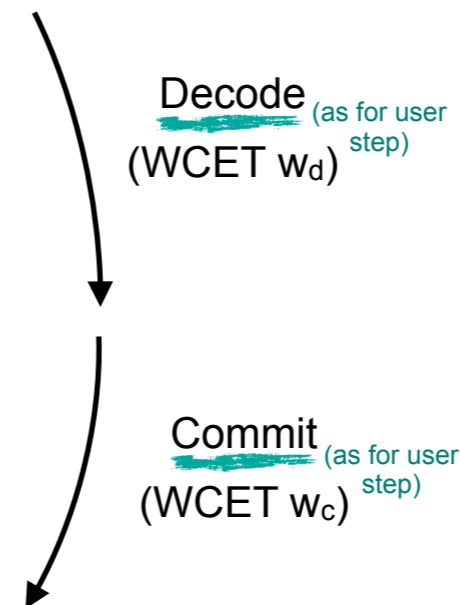
Where  $w_i \leq w_0$



Architecture-specific

### Case 2: System call

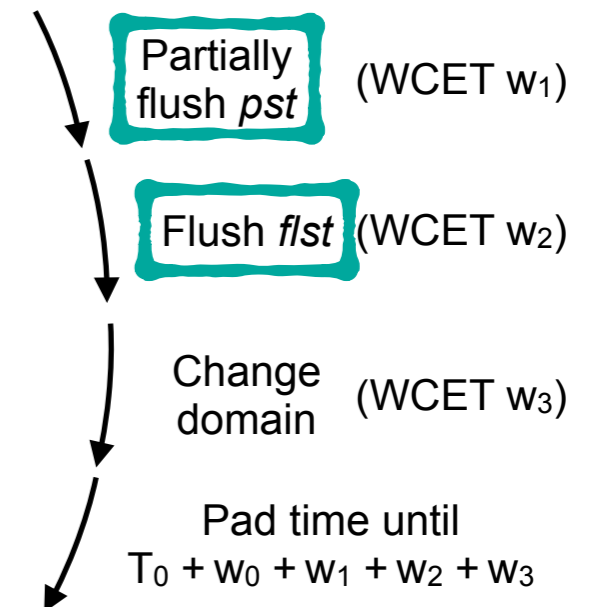
Where  $w_d + w_c \leq w_0$



OS-specific  
(incl. *inflow policies*)

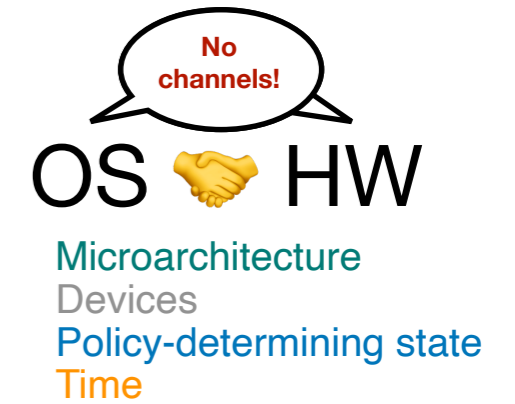
### Case 3: Domain switch

Timer interrupt delivered at (worst-case)  $T_0 + w_0$

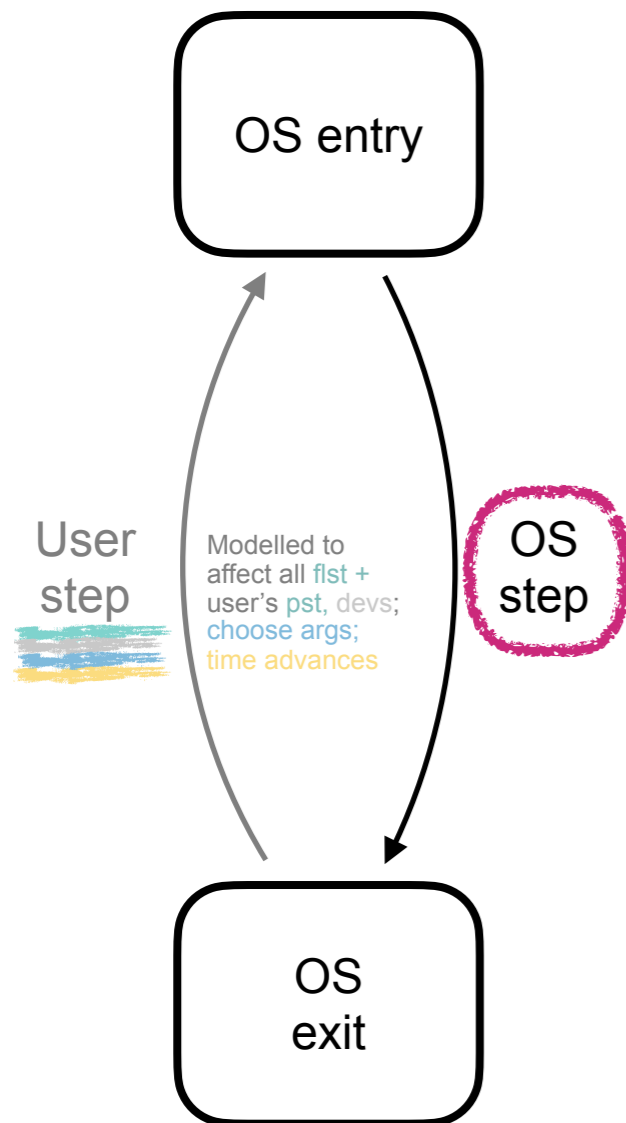


Architecture-specific

# OS security model

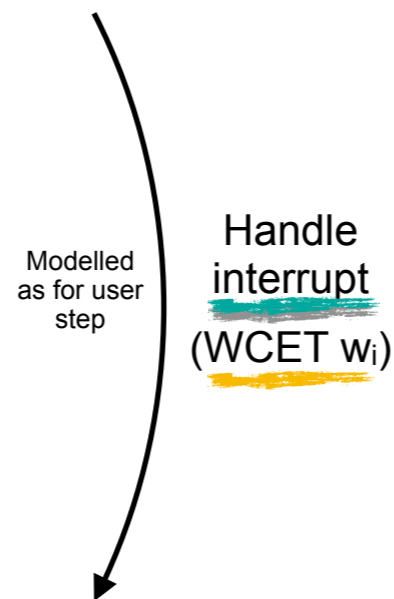


## Transition system



### Case 1: Device interrupt

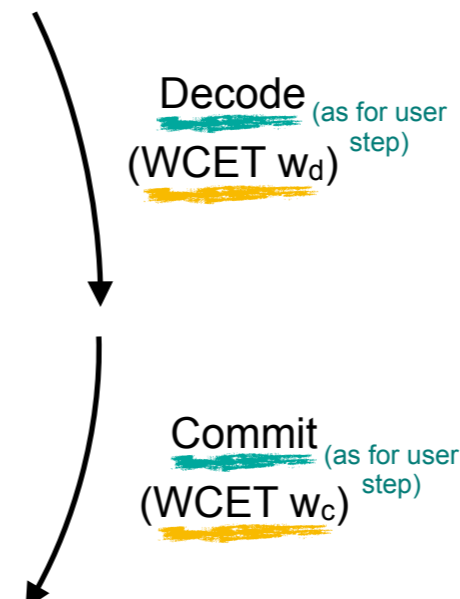
Where  $w_i \leq w_0$



Architecture-specific

### Case 2: System call

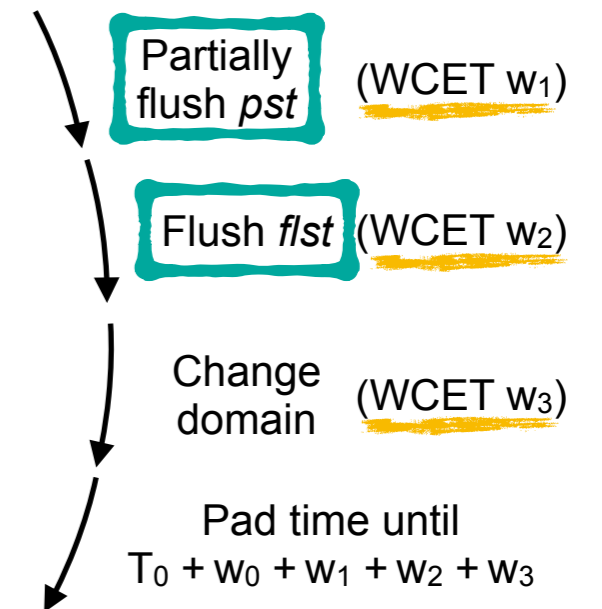
Where  $w_d + w_c \leq w_0$



OS-specific  
(incl. *inflow policies*)

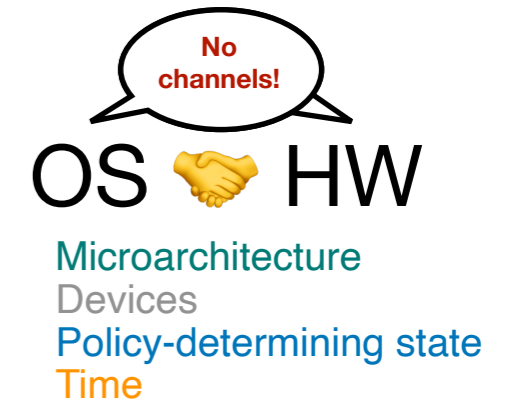
### Case 3: Domain switch

Timer interrupt delivered at (worst-case)  $T_0 + w_0$

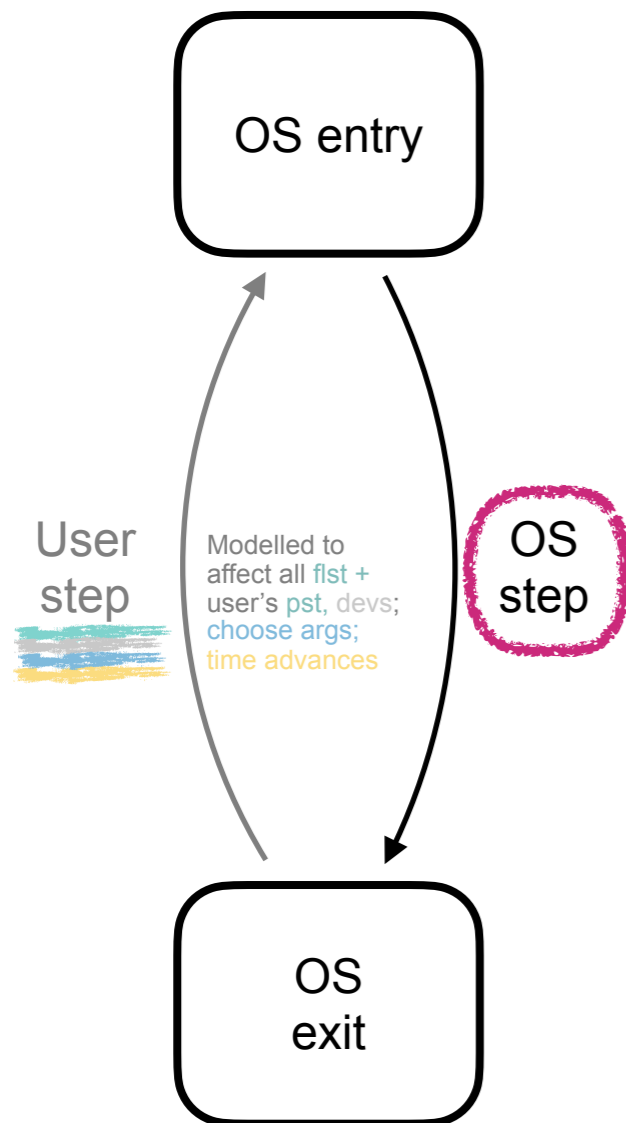


Architecture-specific

# OS security model

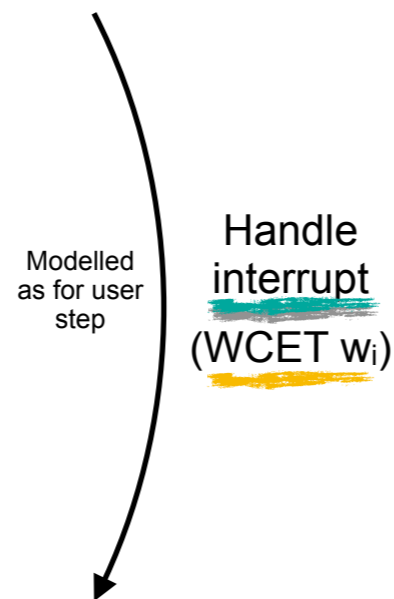


## Transition system



### Case 1: Device interrupt

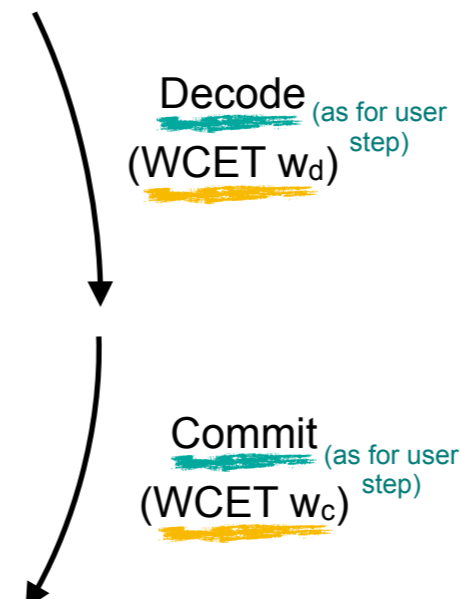
Where  $w_i \leq w_0$



Architecture-specific

### Case 2: System call

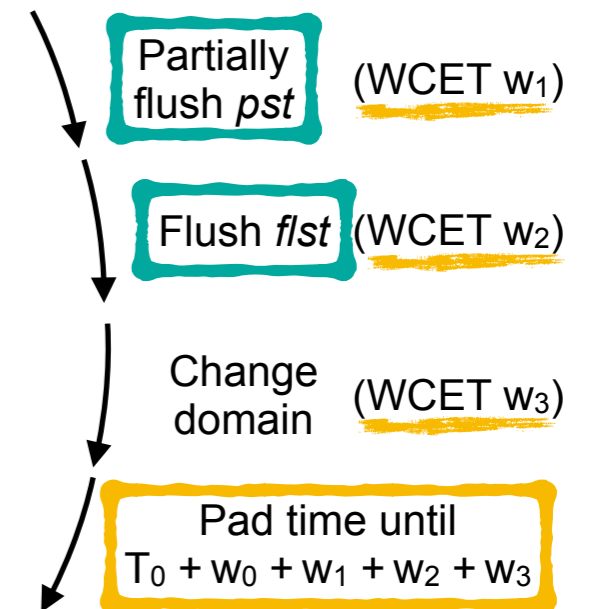
Where  $w_d + w_c \leq w_0$



OS-specific (incl. *inflow policies*)

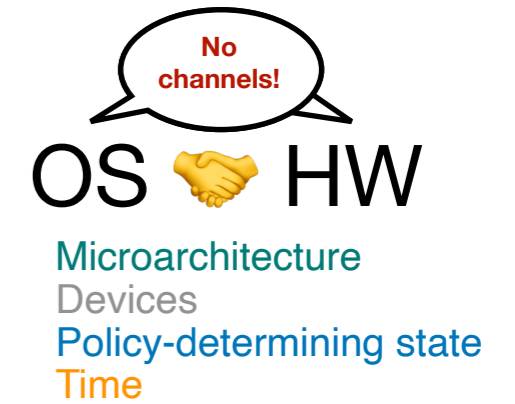
### Case 3: Domain switch

Timer interrupt delivered at (worst-case)  $T_0 + w_0$

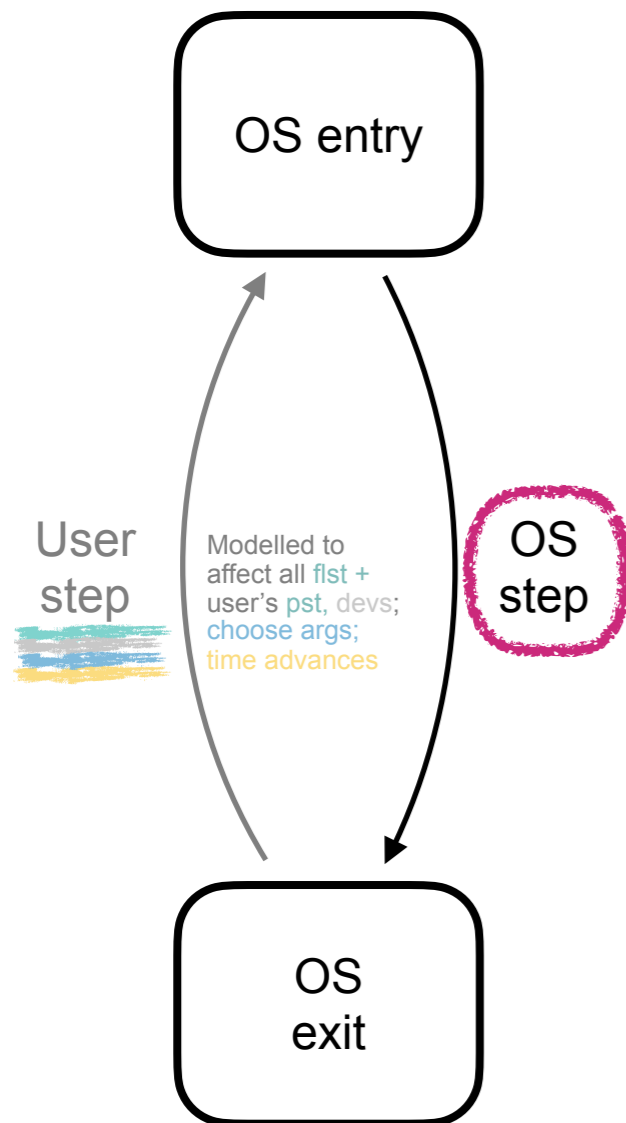


Architecture-specific

# OS security model

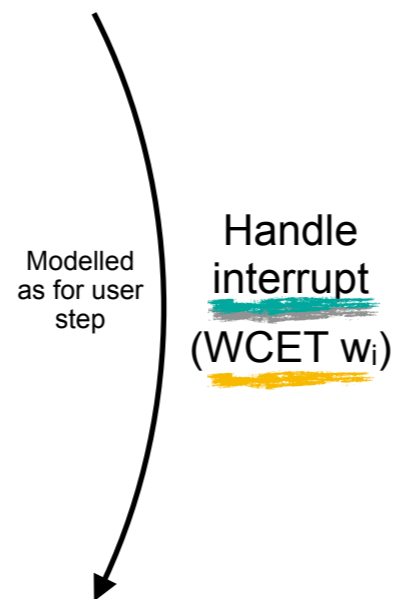


## Transition system



### Case 1: Device interrupt

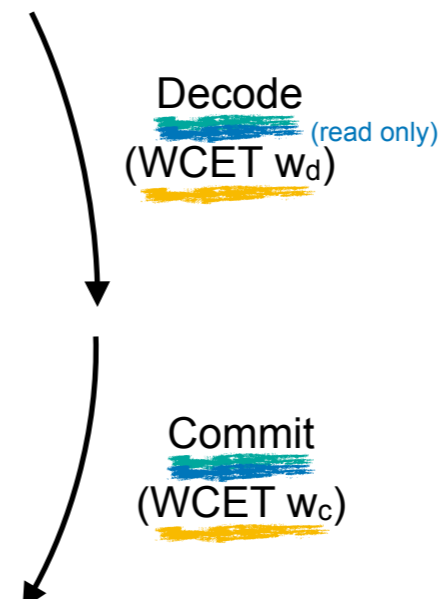
Where  $w_i \leq w_0$



Architecture-specific

### Case 2: System call

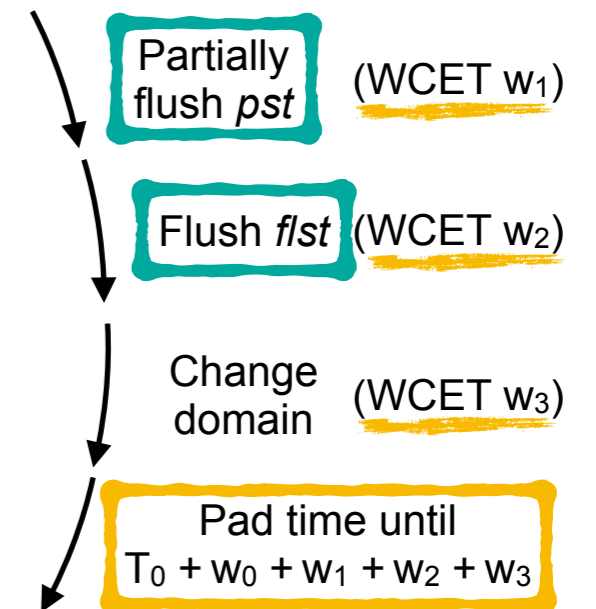
Where  $w_d + w_c \leq w_0$



OS-specific  
(incl. *inflow policies*)

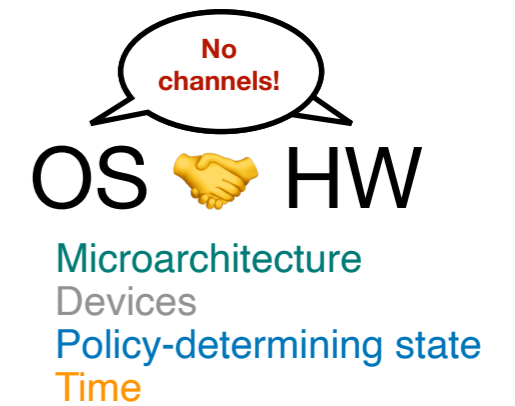
### Case 3: Domain switch

Timer interrupt delivered at (worst-case)  $T_0 + w_0$

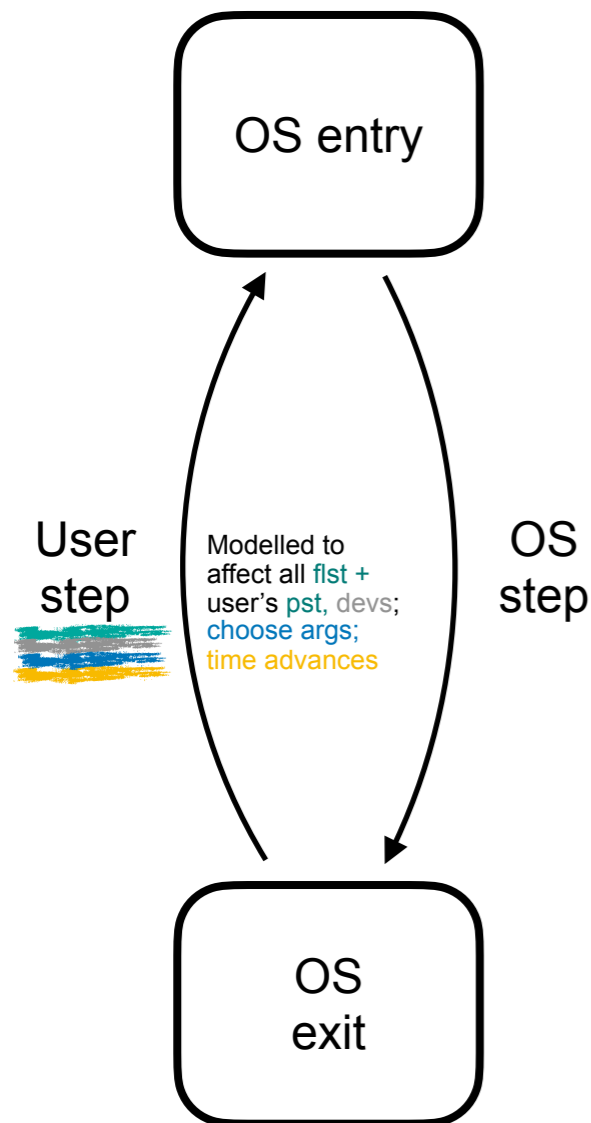


Architecture-specific

# Security proof approach

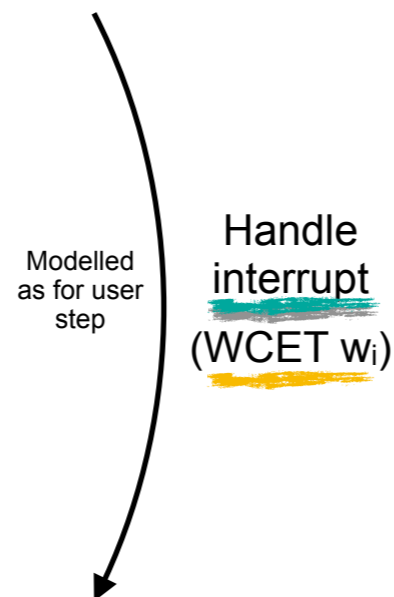


## Transition system



### Case 1: Device interrupt

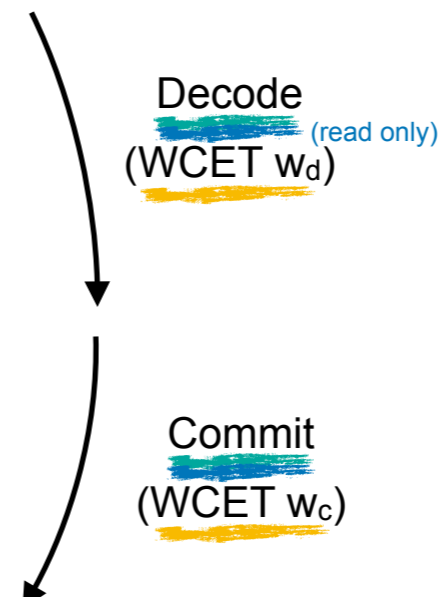
Where  $w_i \leq w_0$



Architecture-specific

### Case 2: System call

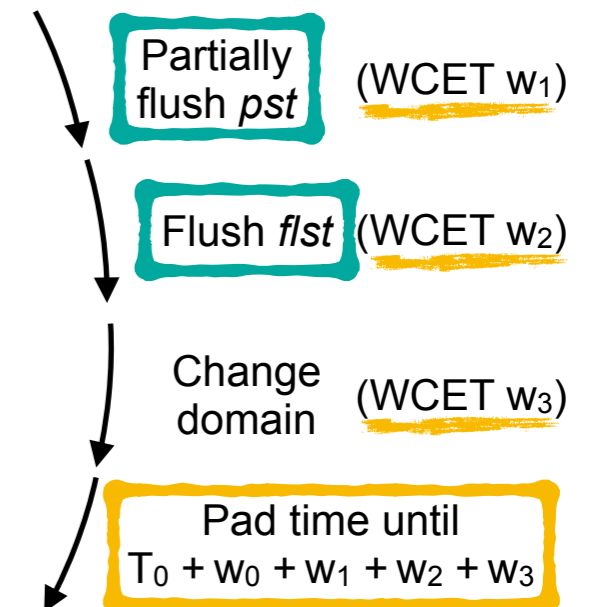
Where  $w_d + w_c \leq w_0$



OS-specific (incl. *inflow policies*)

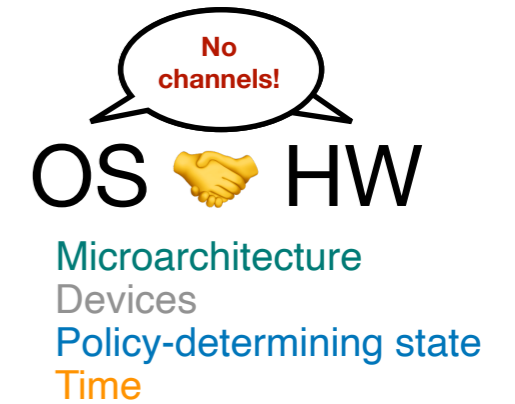
### Case 3: Domain switch

Timer interrupt delivered at (worst-case)  $T_0 + w_0$



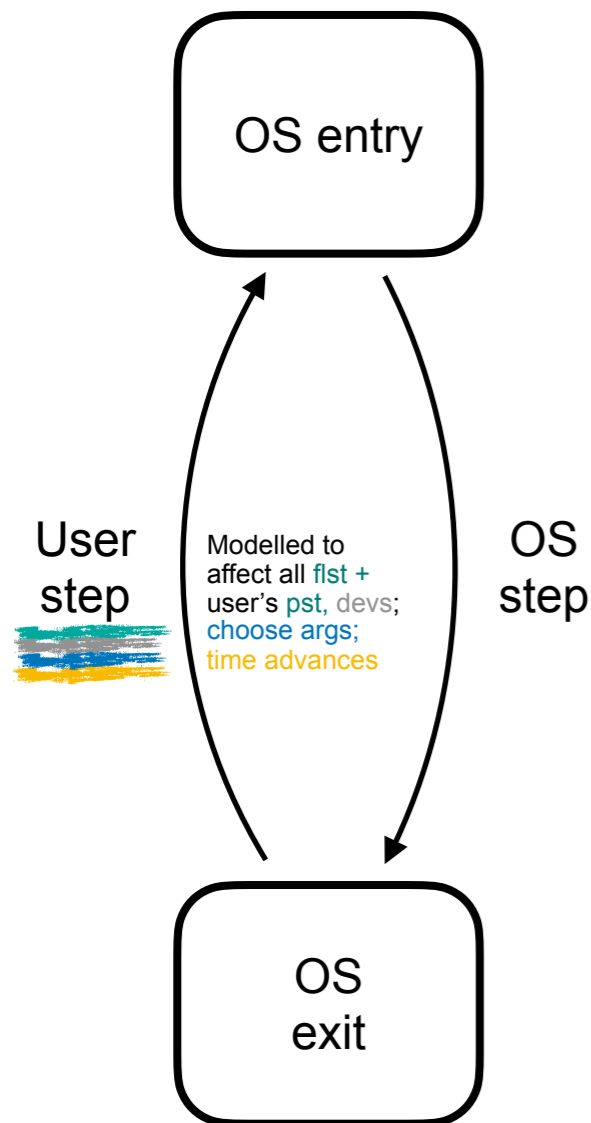
Architecture-specific

# Security proof approach



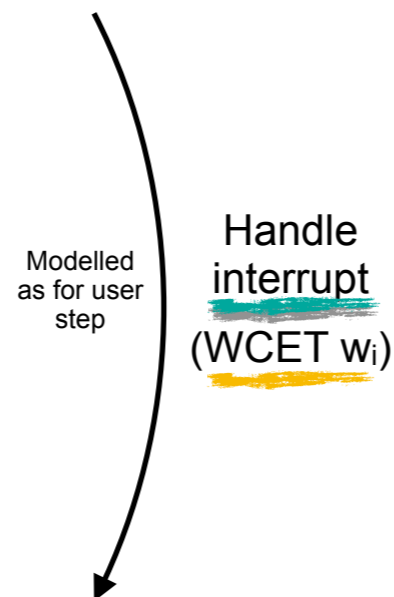
**Requirements**  
(In addition to WCETs)

## Transition system



### Case 1: Device interrupt

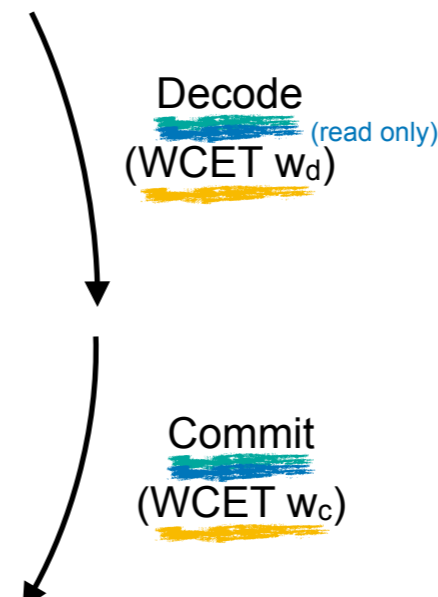
Where  $w_i \leq w_0$



Architecture-specific

### Case 2: System call

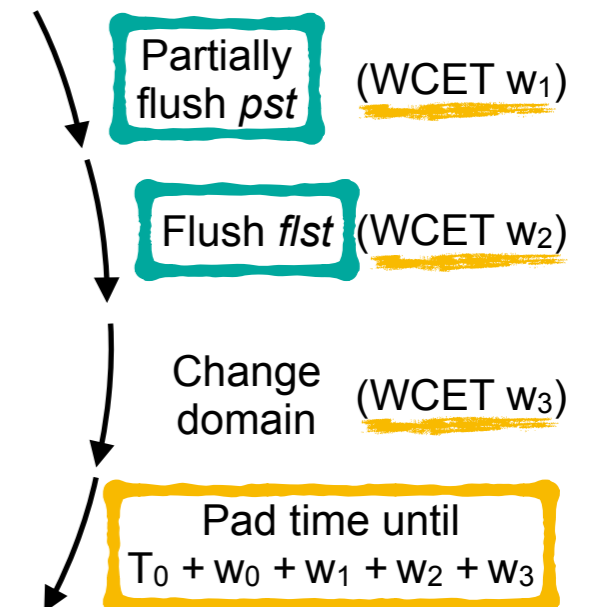
Where  $w_d + w_c \leq w_0$



OS-specific (incl. *inflow policies*)

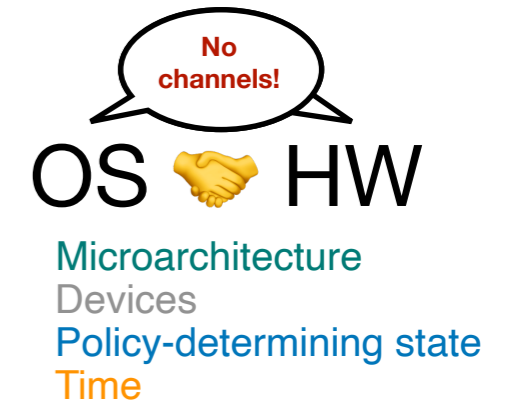
### Case 3: Domain switch

Timer interrupt delivered at (worst-case)  $T_0 + w_0$



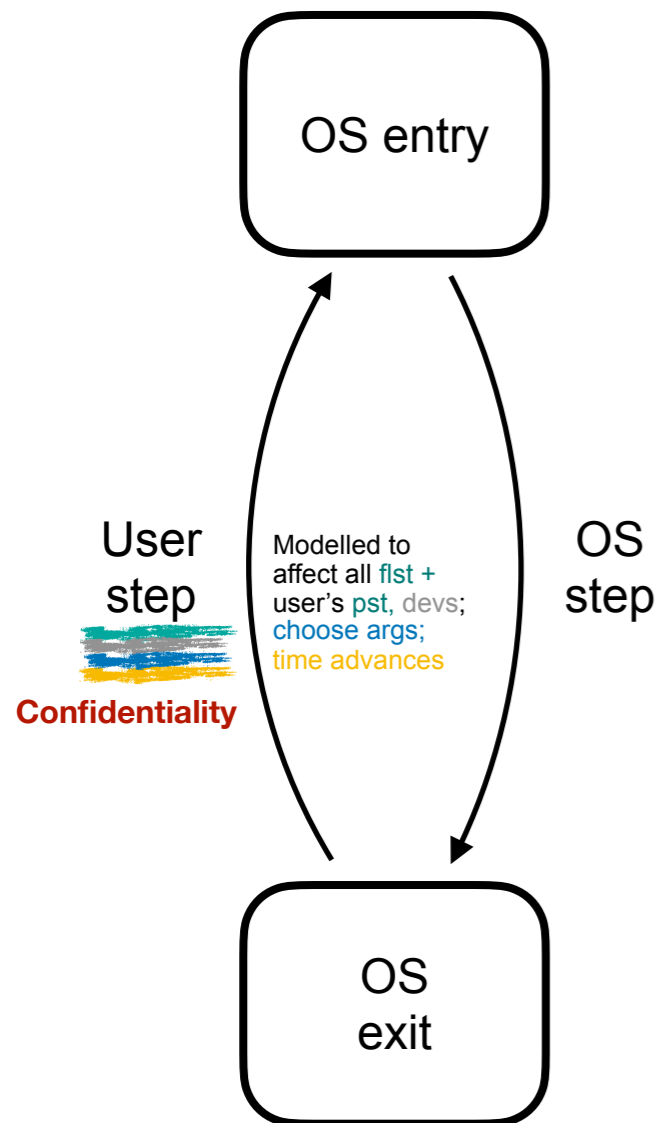
Architecture-specific

# Security proof approach



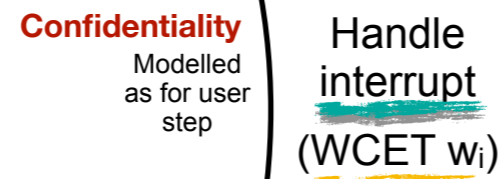
**Requirements**  
(In addition to WCETs)

## Transition system



### Case 1: Device interrupt

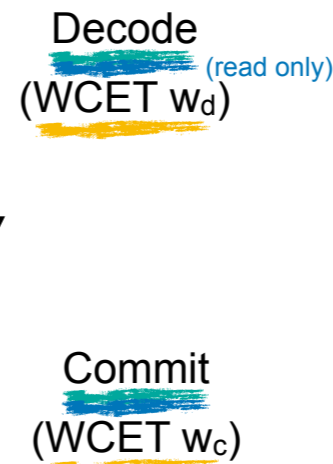
Where  $w_i \leq w_0$



Architecture-specific

### Case 2: System call

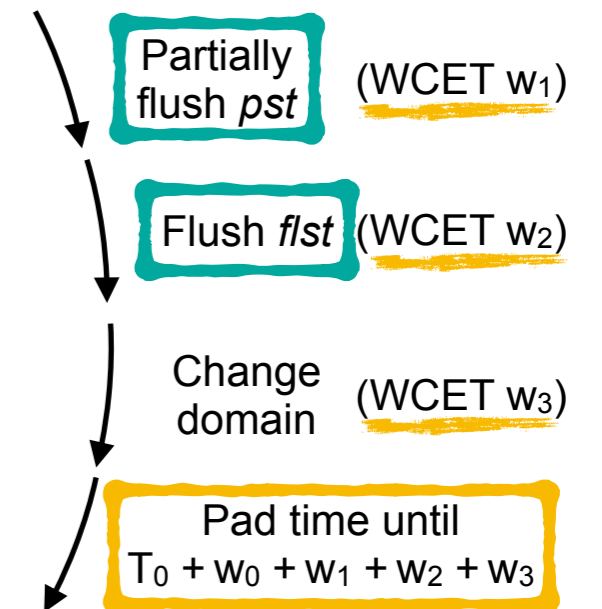
Where  $w_d + w_c \leq w_0$



OS-specific  
(incl. *inflow policies*)

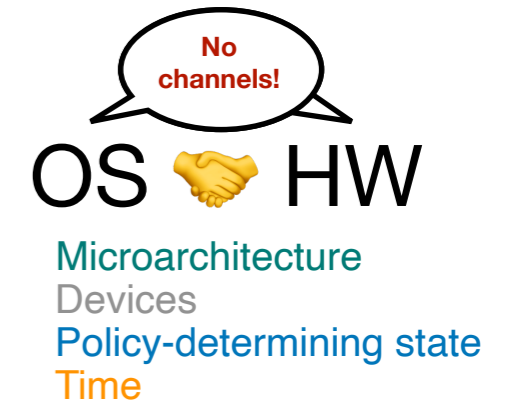
### Case 3: Domain switch

Timer interrupt delivered at (worst-case)  $T_0 + w_0$



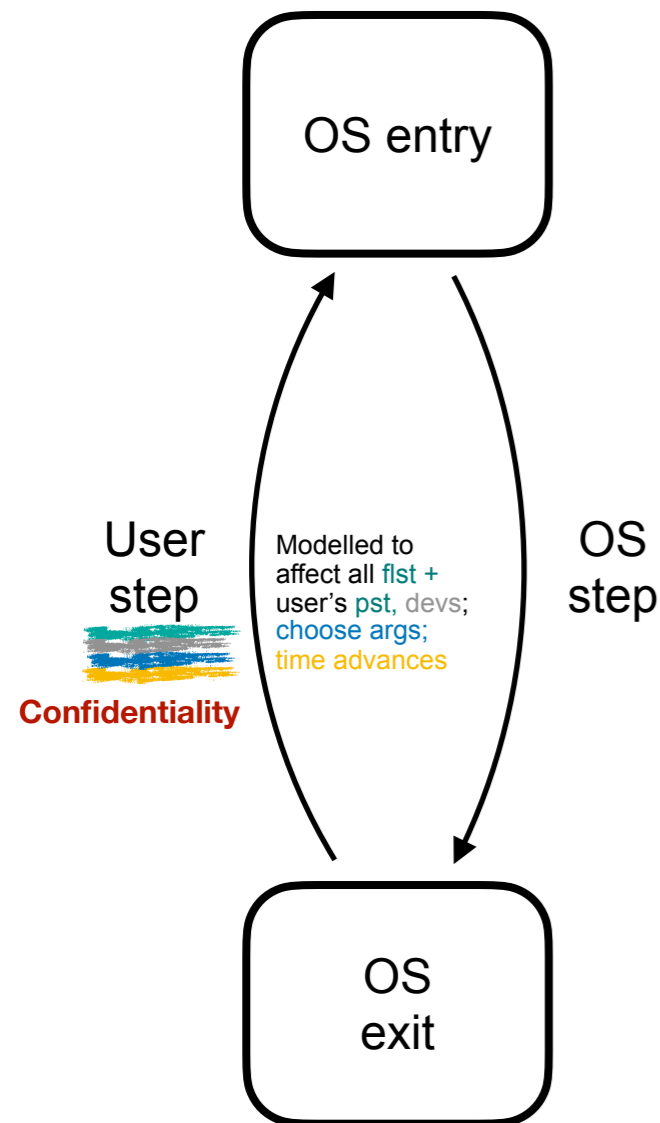
Architecture-specific

# Security proof approach



**Requirements**  
(In addition to WCETs)

## Transition system



### Case 1: Device interrupt

Where  $w_i \leq w_0$

Confidentiality  
Modelled as for user step

Handle interrupt  
(WCET  $w_i$ )

Architecture-specific

### Case 2: System call

Where  $w_d + w_c \leq w_0$

Integrity

Decode (WCET  $w_d$ )  
(read only)

Confidentiality (relative to policy)

Commit (WCET  $w_c$ )

OS-specific  
(incl. *inflow policies*)

### Case 3: Domain switch

Timer interrupt delivered at (worst-case)  $T_0 + w_0$

Partially flush *pst* (WCET  $w_1$ )

Flush *flst* (WCET  $w_2$ )

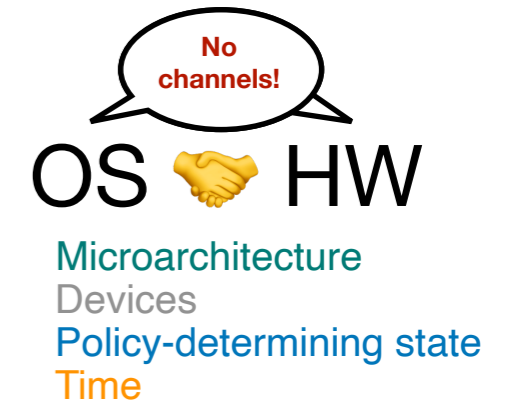
Change domain (WCET  $w_3$ )

Pad time until  $T_0 + w_0 + w_1 + w_2 + w_3$

Architecture-specific

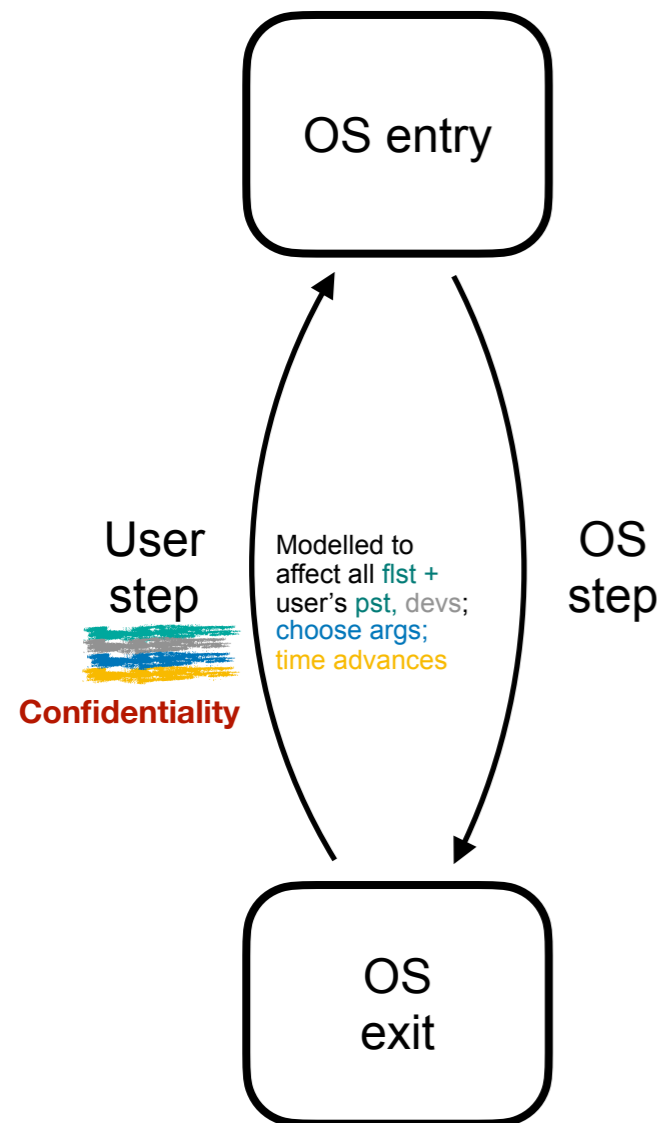


# Security proof approach



**Requirements**  
(In addition to WCETs)

## Transition system



### Case 1: Device interrupt

Where  $w_i \leq w_0$

Confidentiality  
Modelled as for user step

Handle interrupt  
(WCET  $w_i$ )

Architecture-specific

### Case 2: System call

Where  $w_d + w_c \leq w_0$

Integrity

Decode  
(WCET  $w_d$ )  
(read only)

Confidentiality (relative to policy)

Commit  
(WCET  $w_c$ )

OS-specific  
(incl. *inflow policies*)

### Case 3: Domain switch

Timer interrupt delivered at (worst-case)  $T_0 + w_0$

Correctness

Partially flush *pst* (WCET  $w_1$ )

Correctness

Flush *flst* (WCET  $w_2$ )

Correctness

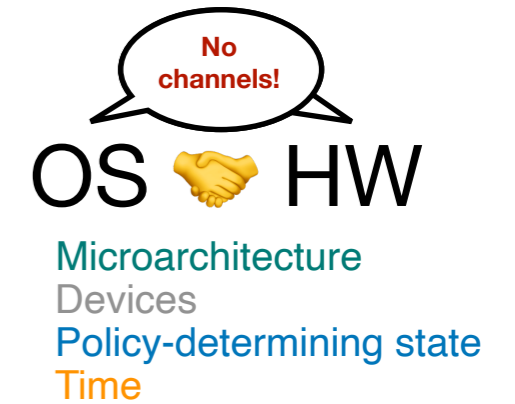
Change domain (WCET  $w_3$ )

Correctness

Pad time until  $T_0 + w_0 + w_1 + w_2 + w_3$

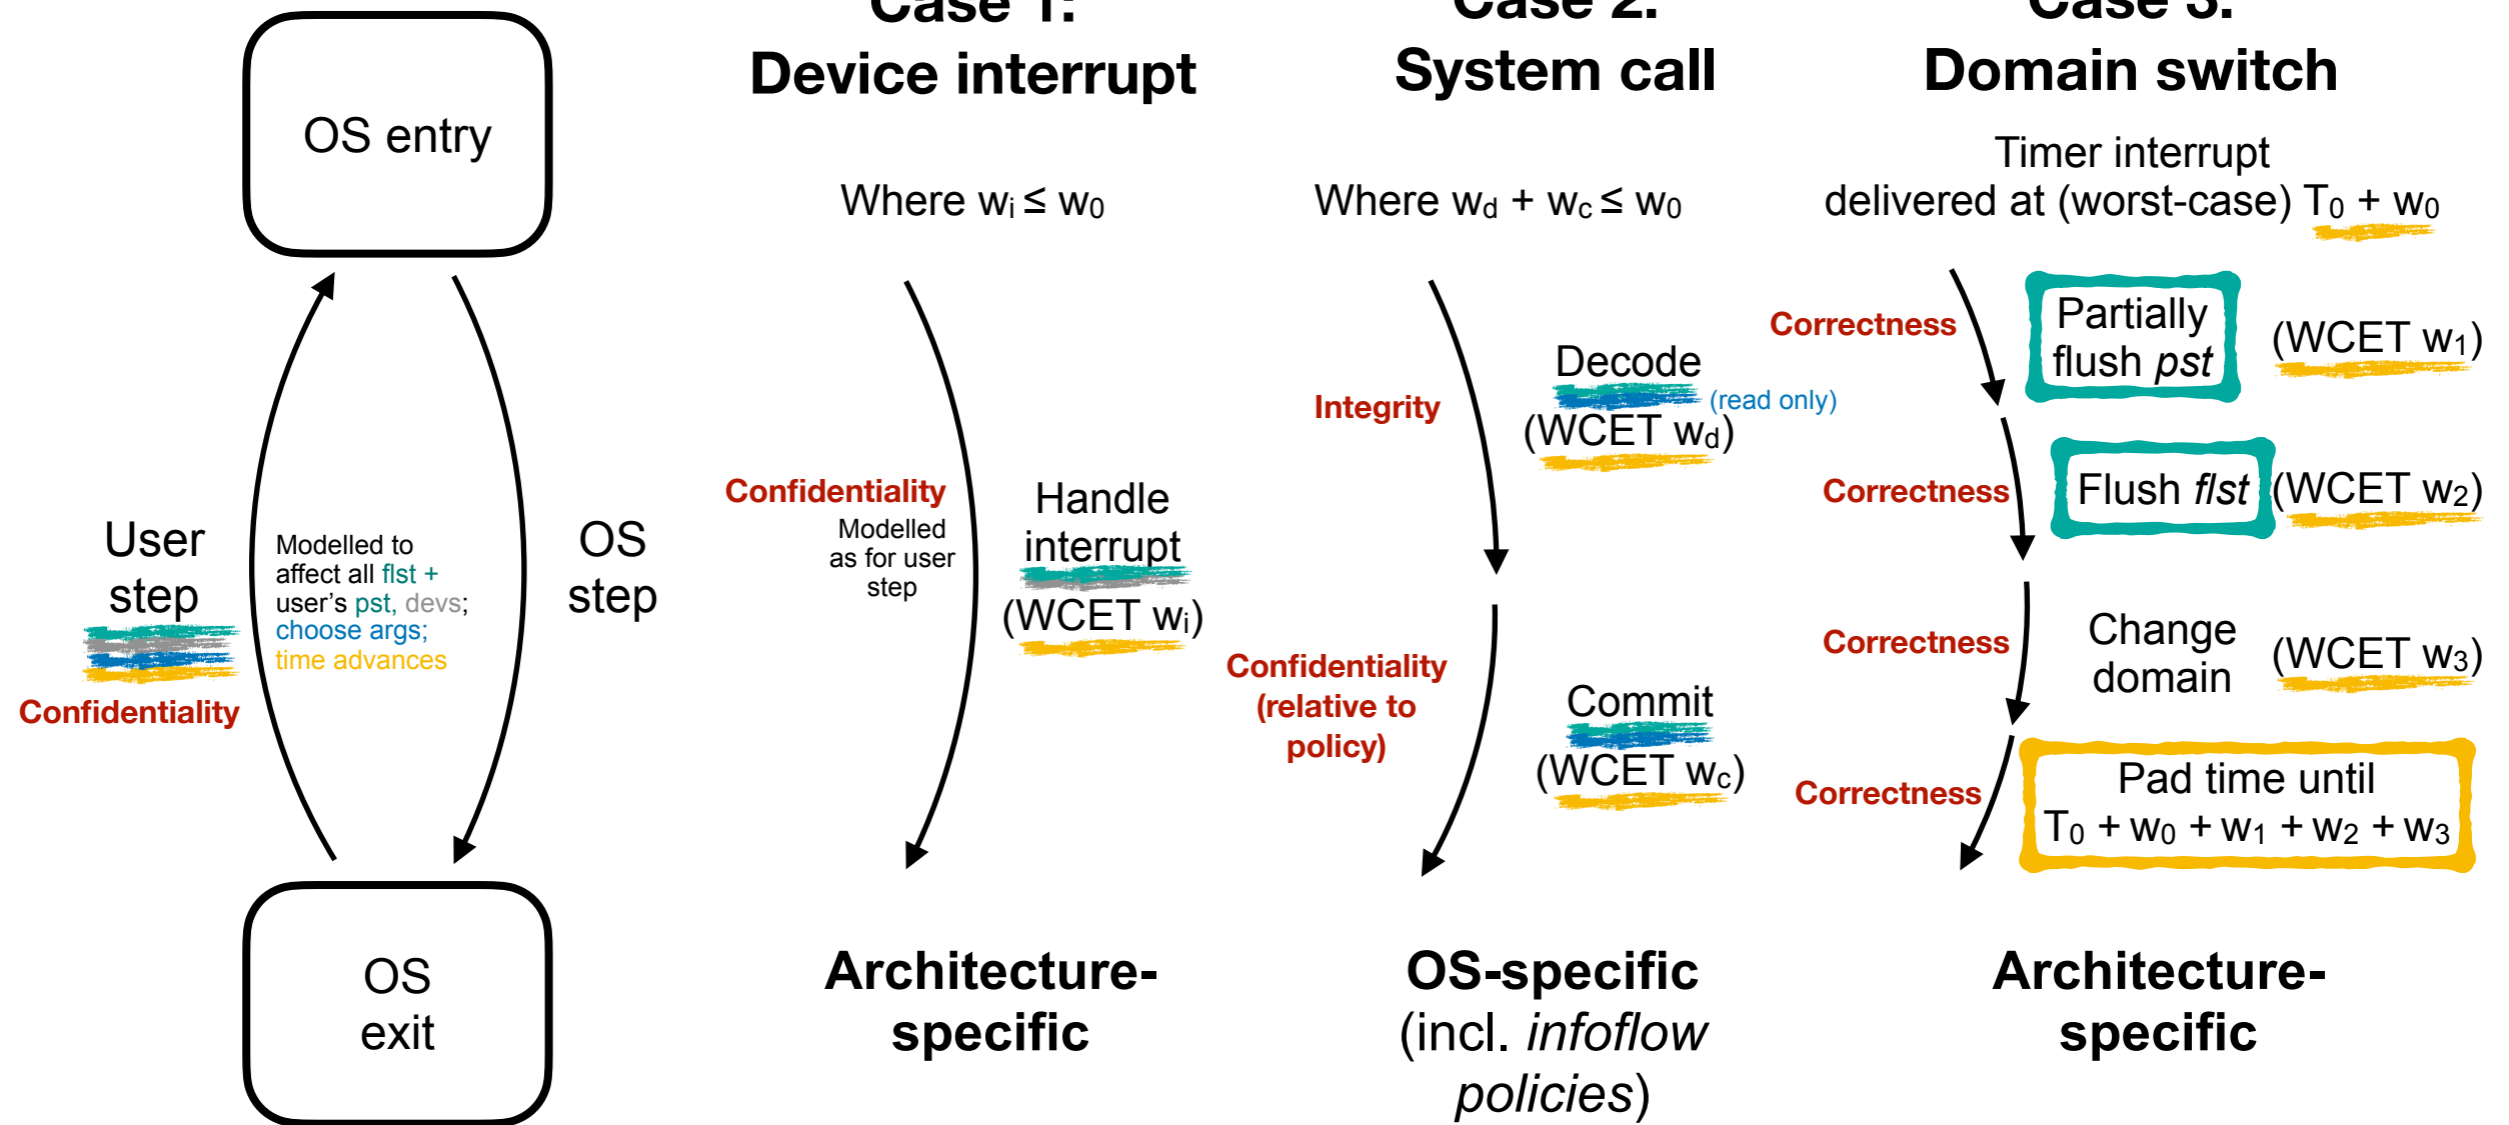
Architecture-specific

# Security proof approach



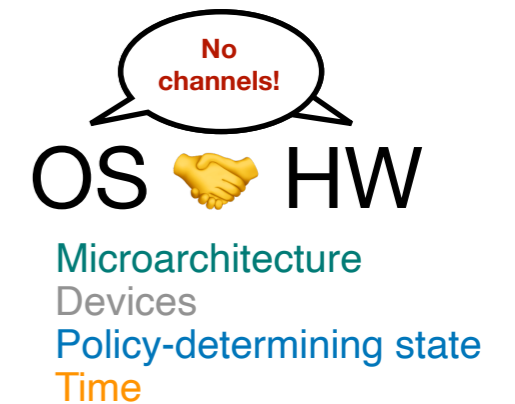
**Requirements**  
(In addition to WCETs)

## Transition system



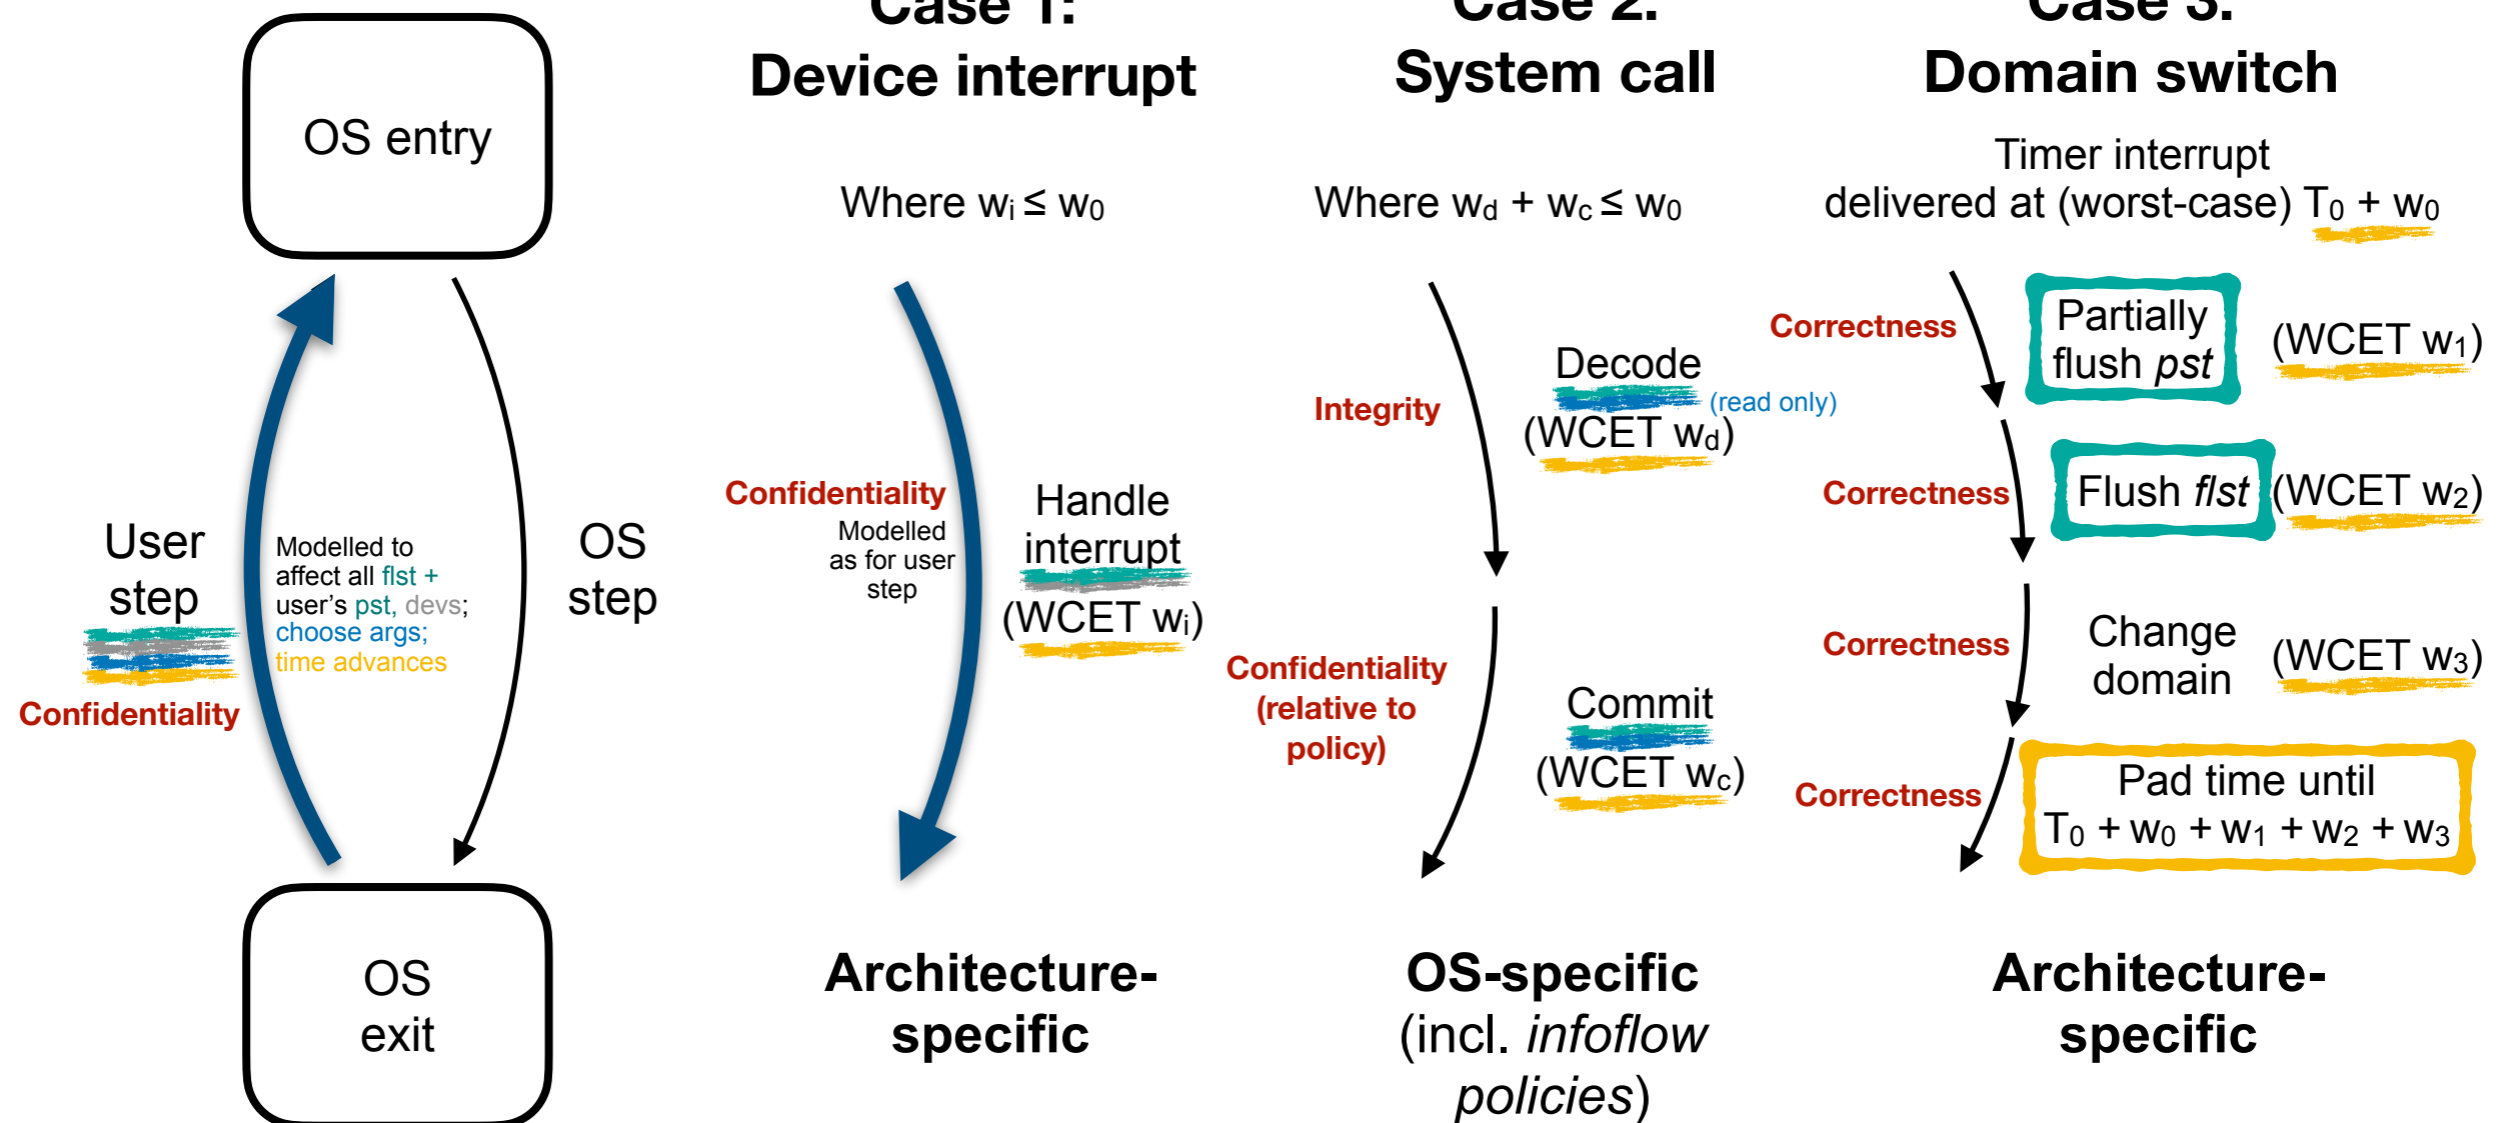
We prove: **Confidentiality property (bisimulation) step lemmas**

# Security proof approach



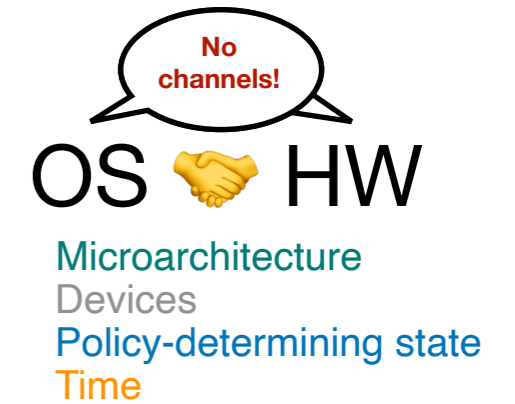
**Requirements**  
(In addition to WCETs)

## Transition system



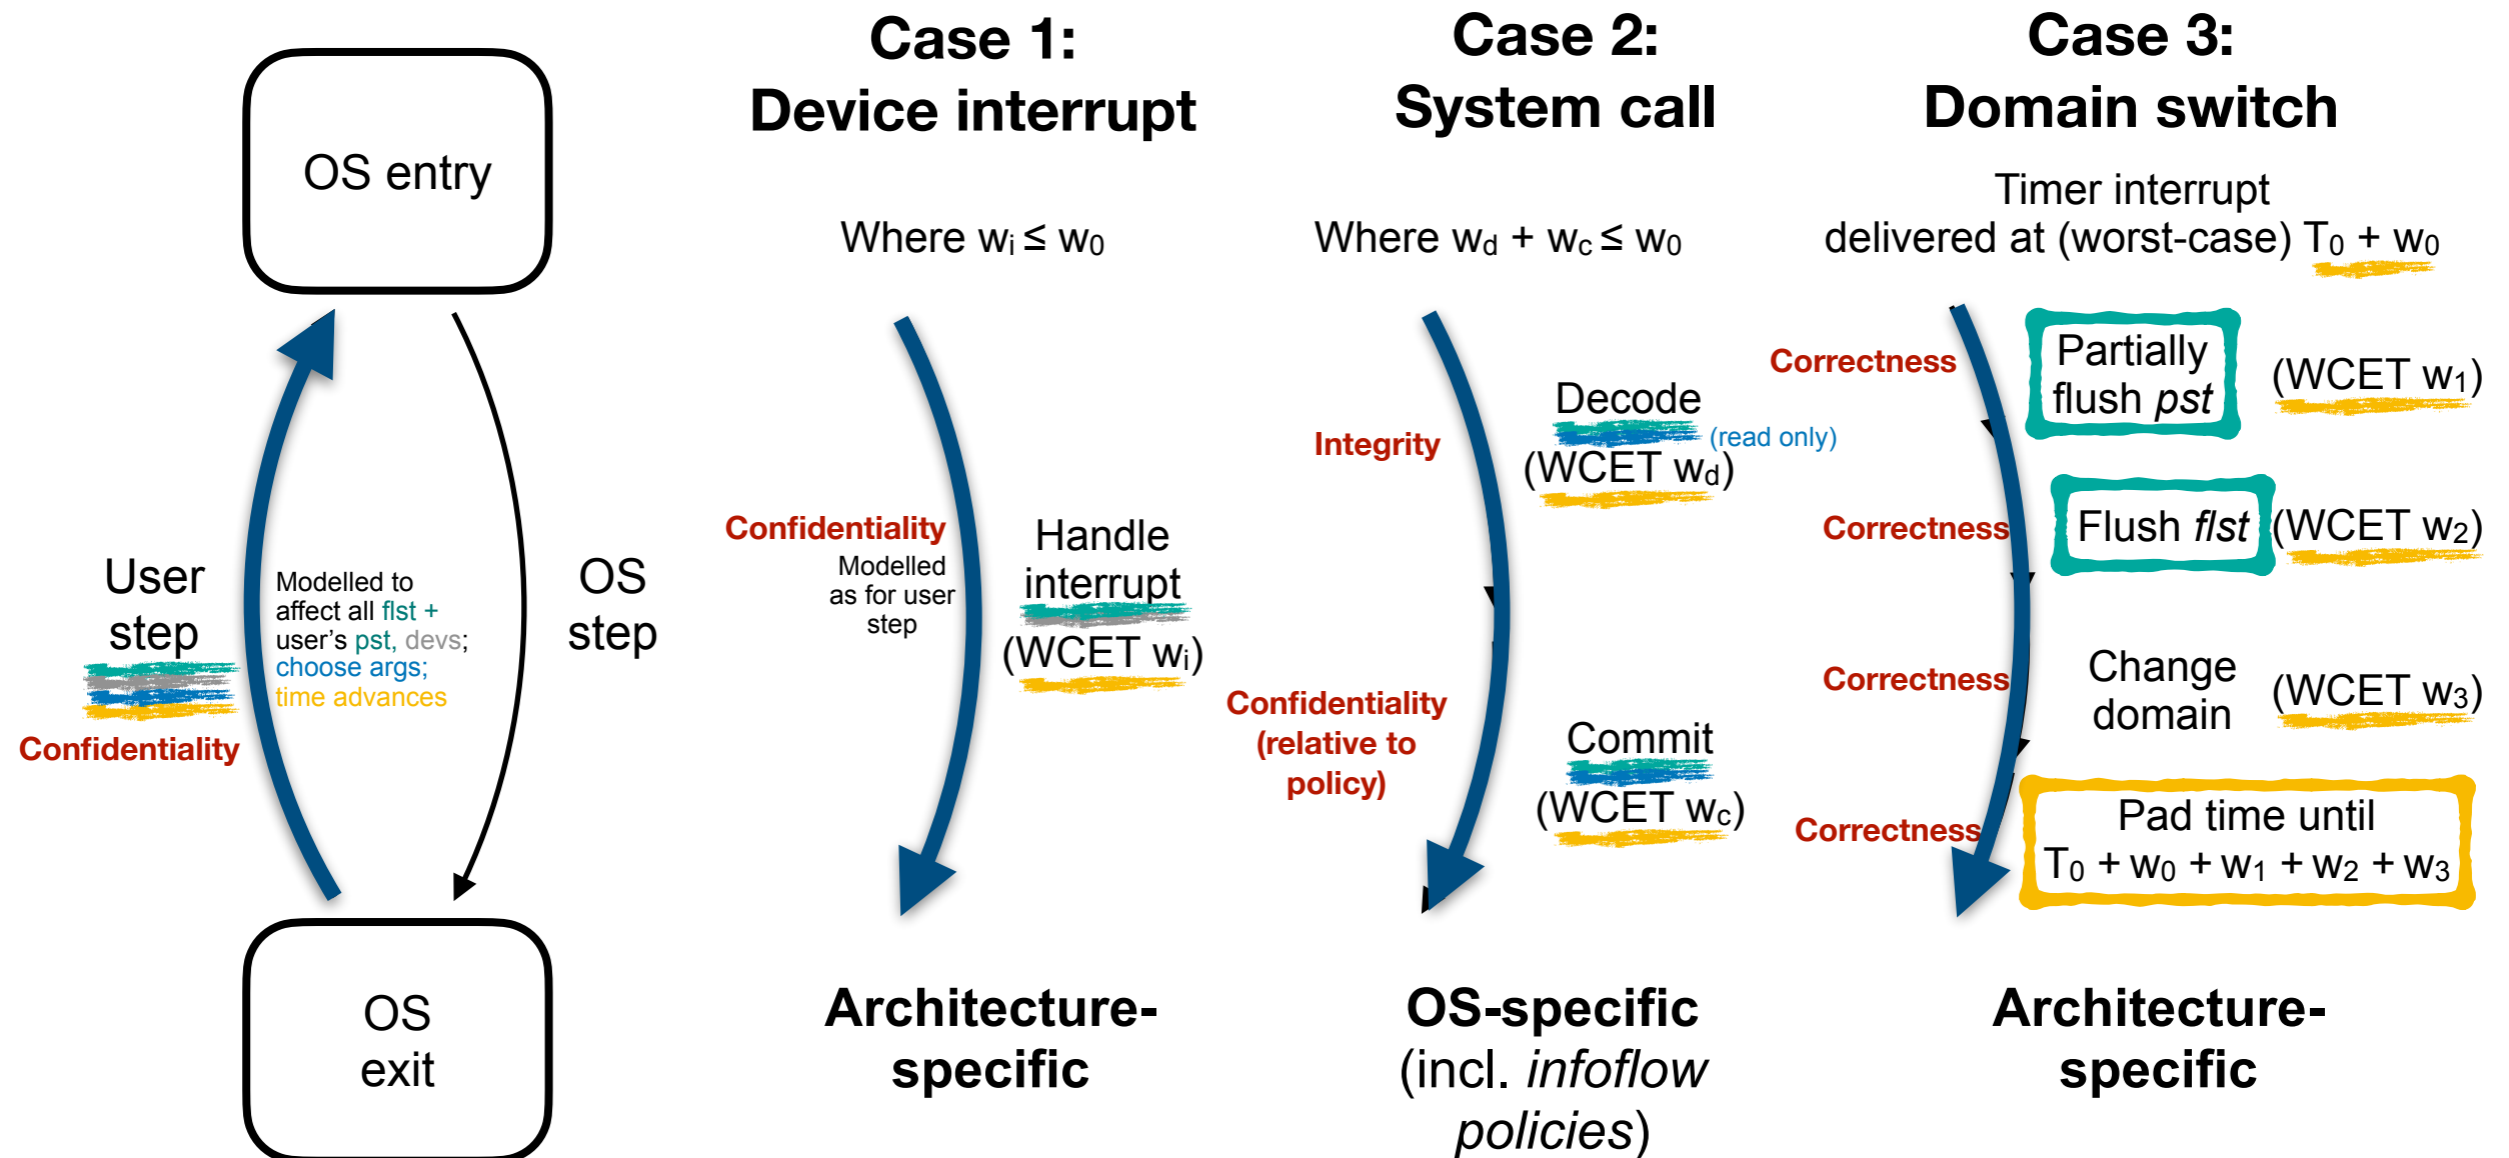
We prove: **Confidentiality property (bisimulation) step lemmas**

# Security proof approach



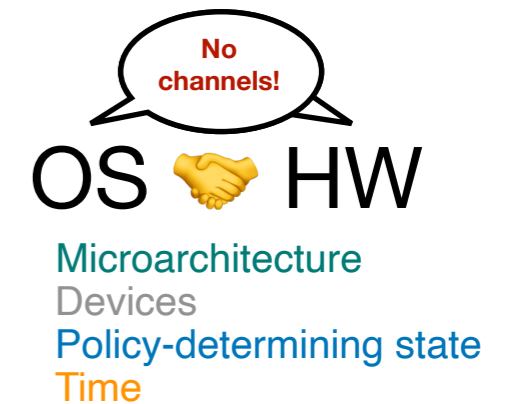
**Requirements**  
(In addition to WCETs)

## Transition system



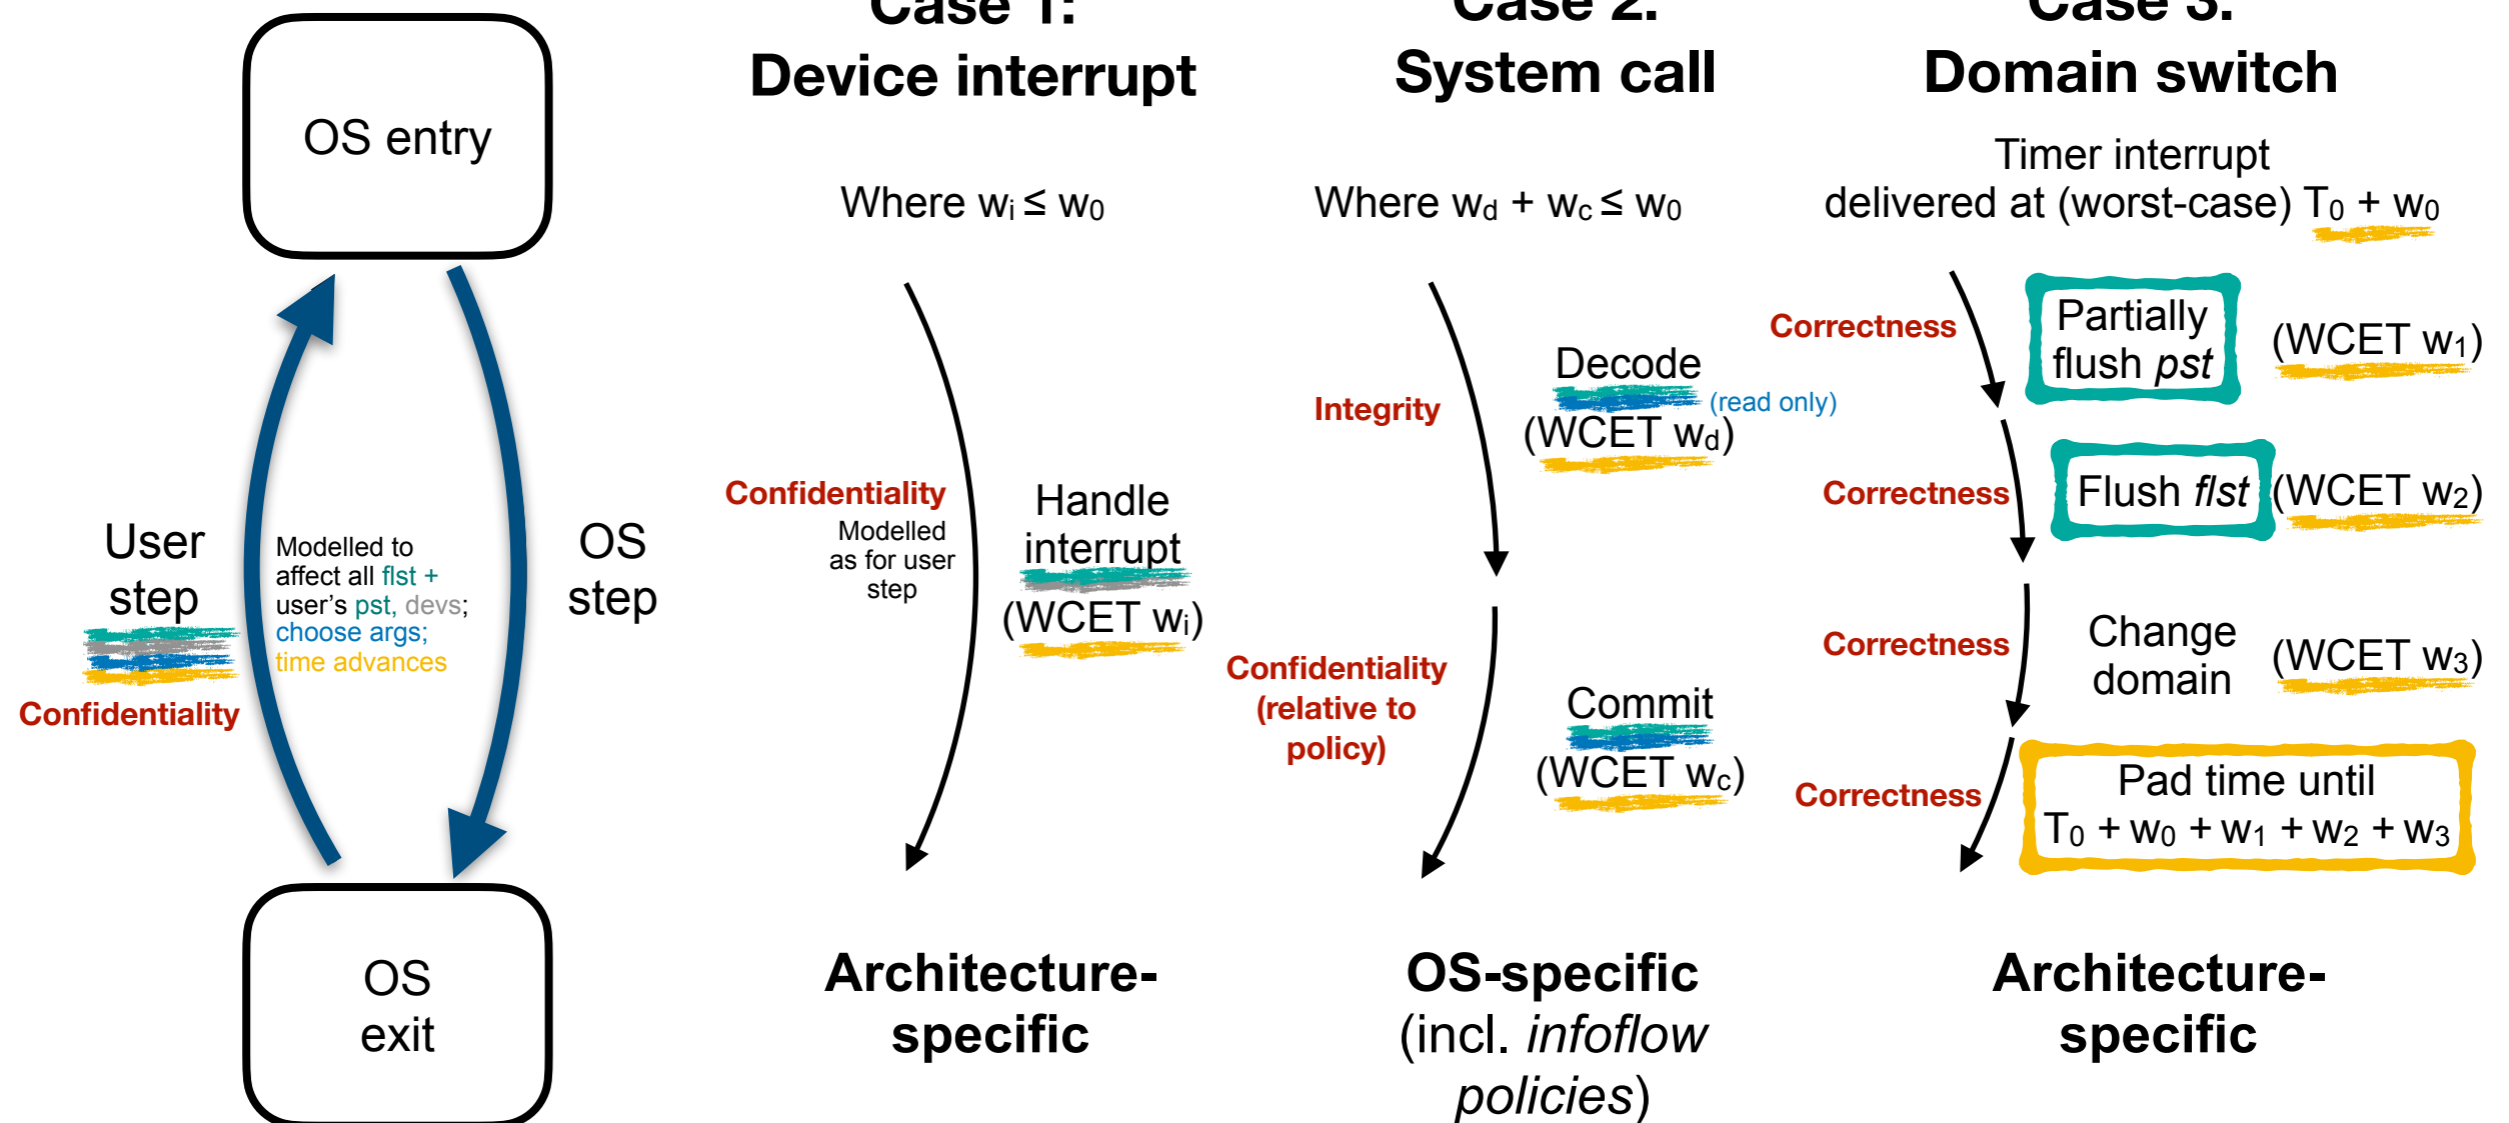
We prove: **Confidentiality property (bisimulation) step lemmas**

# Security proof approach



**Requirements**  
(In addition to WCETs)

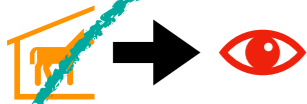
## Transition system



We prove: **Confidentiality property (bisimulation) step lemmas**

# How to formalise an OS enforces *time protection*?

Versus threat scenario:  
trojan and spy



Abstract *covert state* + *time* to reflect strategies enabled by HW:  
Partition or flush state; pad time.



Demonstrating these principles,  
we formalised in Isabelle/HOL:

1. OS security model imposing requirements on relevant parts of OS.  
(Intended for seL4, but *generic*)
3. Proof our security property holds if OS model's requirements hold.

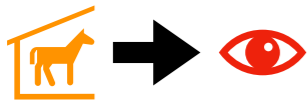


Make security property precise enough to exclude flows from covert state.

2. OS security property that is dynamic; this makes it observer relative.  
(Improving on seL4's of [Murray et al. 2012])
4. Basic instantiation of OS model exercising dynamic policy.

# How to formalise an OS enforces *time protection*?

Versus threat scenario:  
trojan and spy



Abstract *covert state* + *time* to reflect strategies enabled by HW:  
Partition or flush state; pad time.



Demonstrating these principles, we formalised in Isabelle/HOL:

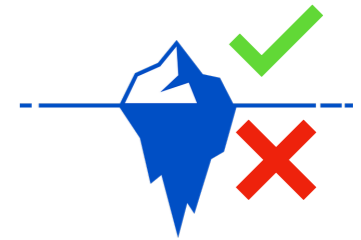
1. OS security model imposing requirements on relevant parts of OS. (Intended for seL4, but *generic*)
3. Proof our security property holds if OS model's requirements hold.



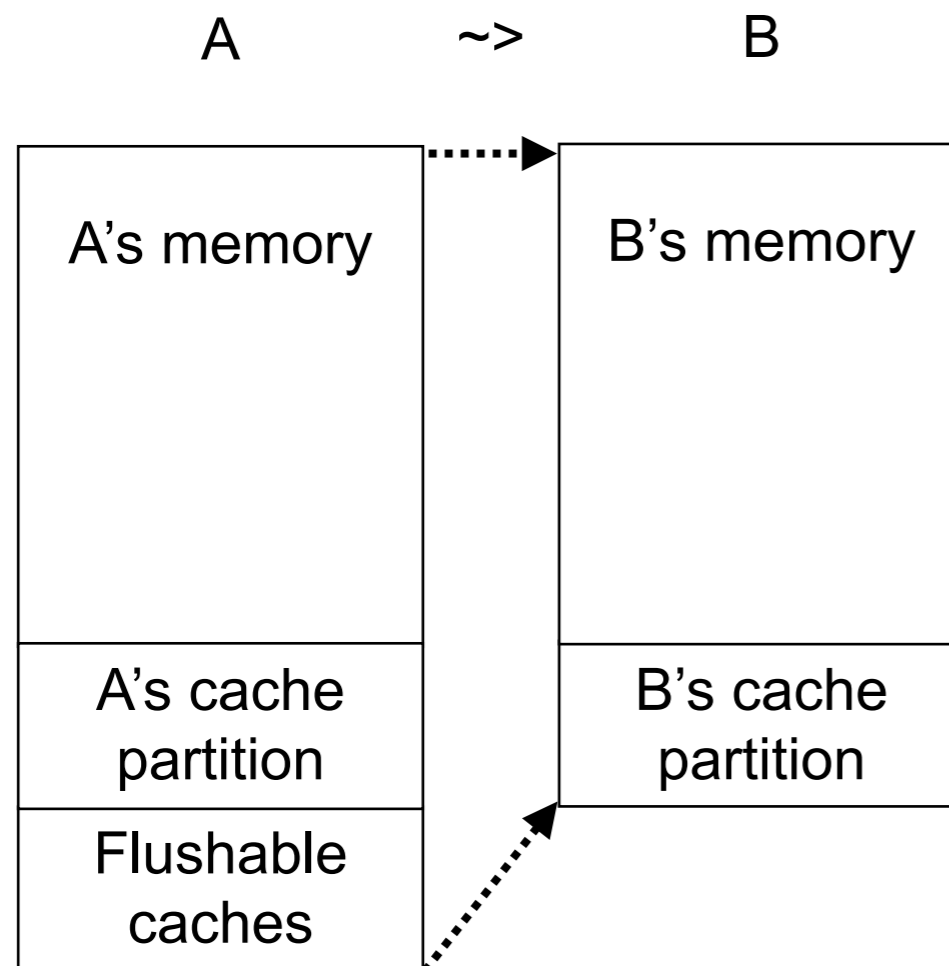
Make security property precise enough to exclude flows from covert state.

2. OS security property that is dynamic; this makes it observer relative. (Improving on seL4's of [Murray et al. 2012])
4. Basic instantiation of OS model exercising dynamic policy.

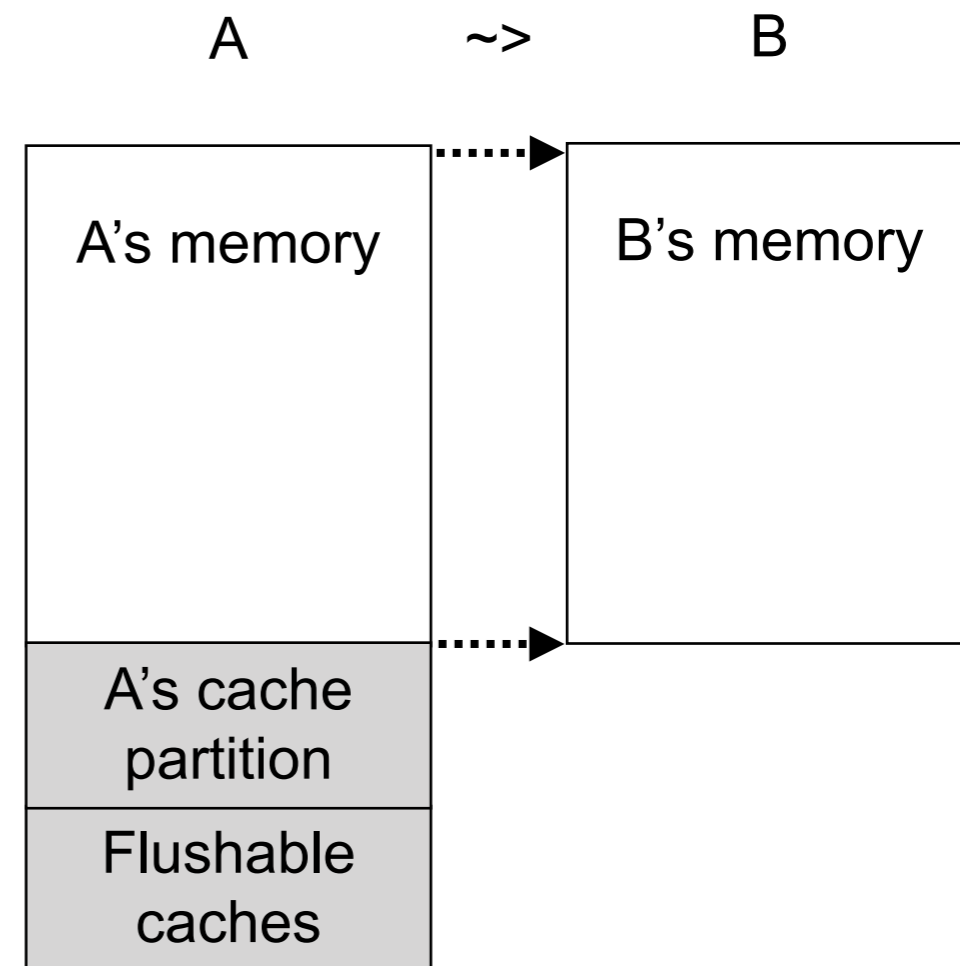
# OS security property



Recall:



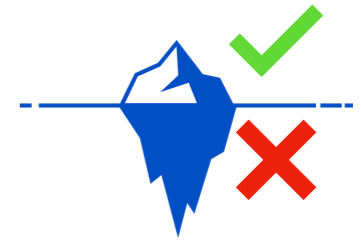
From prior seL4 infowflow proofs  
[Murray et al. 2012, 2013]:  
*“all or nothing” policies*



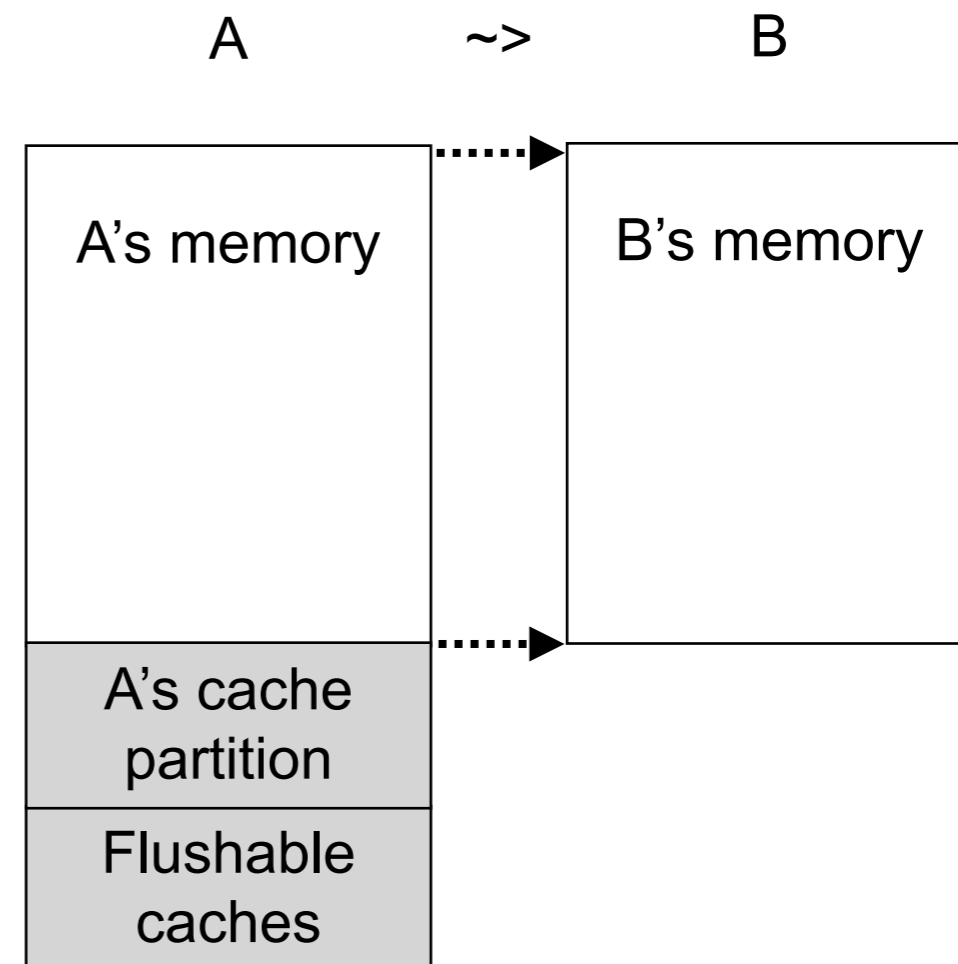
**For time protection, need  
*spatial precision to allow some flows  
but exclude others***



# OS security property

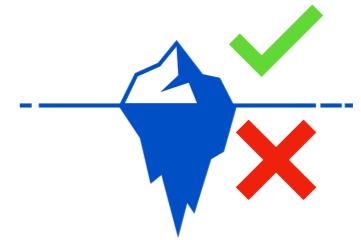


Our infoflow policies:



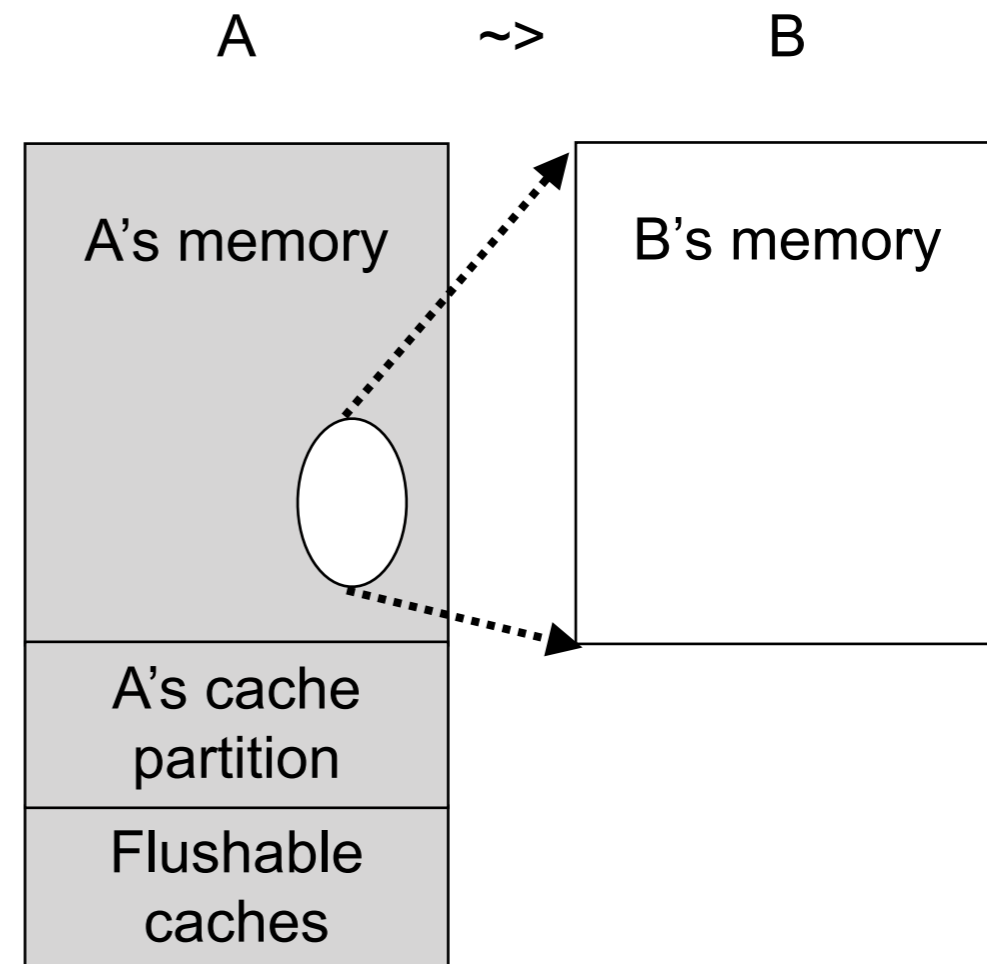
**For time protection, need  
*spatial precision to allow some flows  
but exclude others***

# OS security property



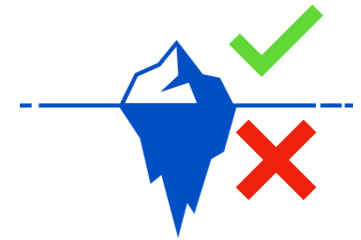
Our infoflow policies:

- *Arbitrary* spatial precision



**For time protection, need  
*spatial precision to allow some flows  
but exclude others***

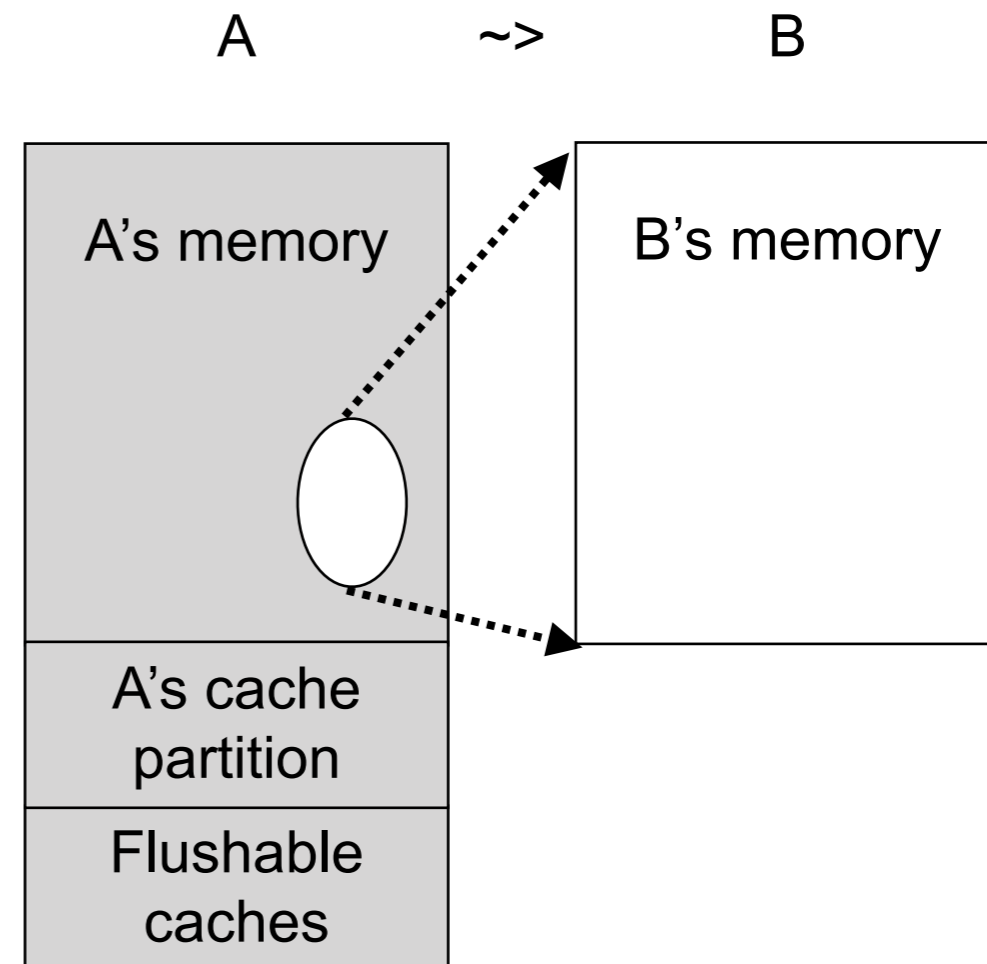
# OS security property



Our infoflow policies:

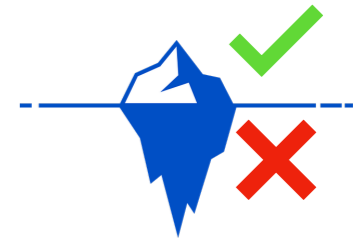
- *Arbitrary* spatial precision
- *Policy channels* specified as *state relations*:  $s \mid A \rightsquigarrow B \mid t$

If  $\mid A \rightsquigarrow B \mid$  equates part of A, then info flow is allowed from there to B.



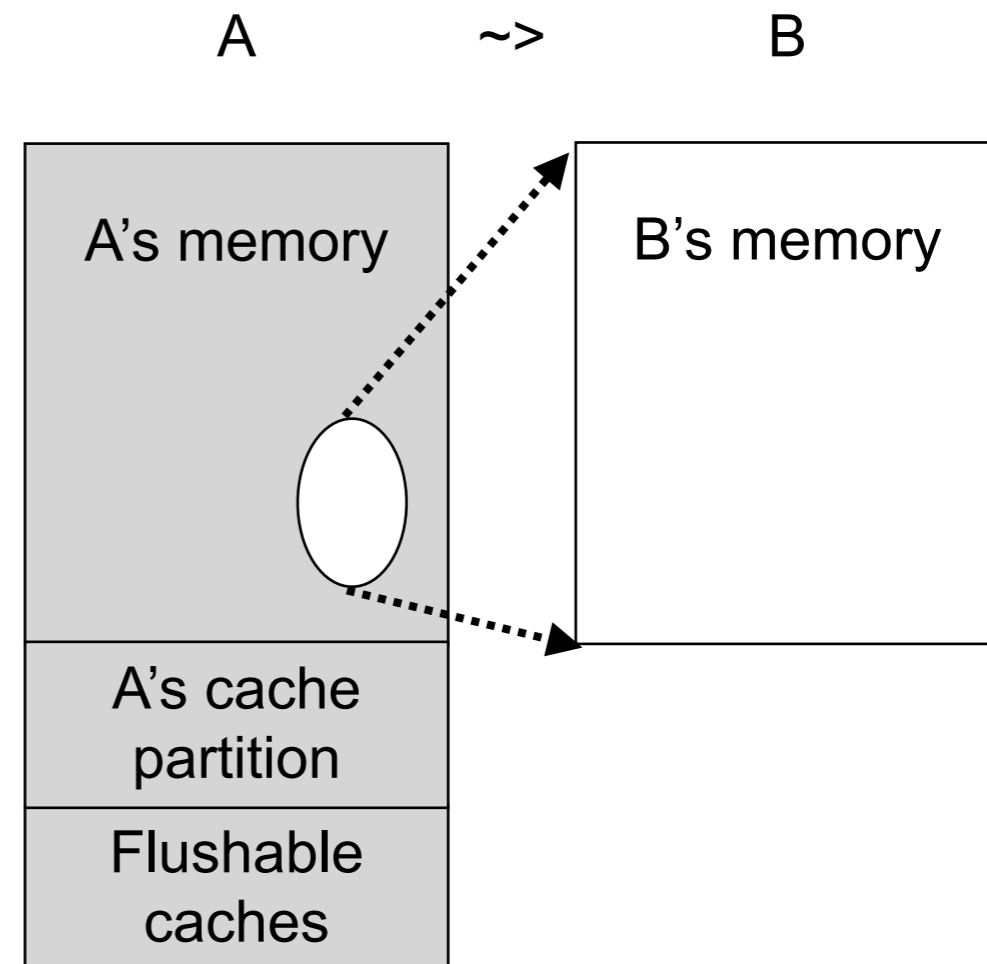
**For time protection, need  
*spatial precision to allow some flows  
but exclude others***

# OS security property



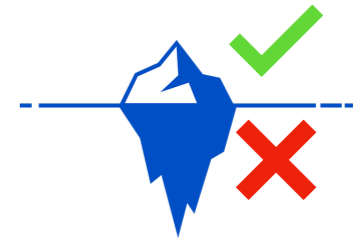
Our infoflow policies:

- *Arbitrary* spatial precision
- *Policy channels* specified as *state relations*:  $s \mid A \rightsquigarrow B \mid t$   
If  $\mid A \rightsquigarrow B \mid$  equates part of A, then info flow is allowed from there to B.
- Also arbitrary *temporal precision*



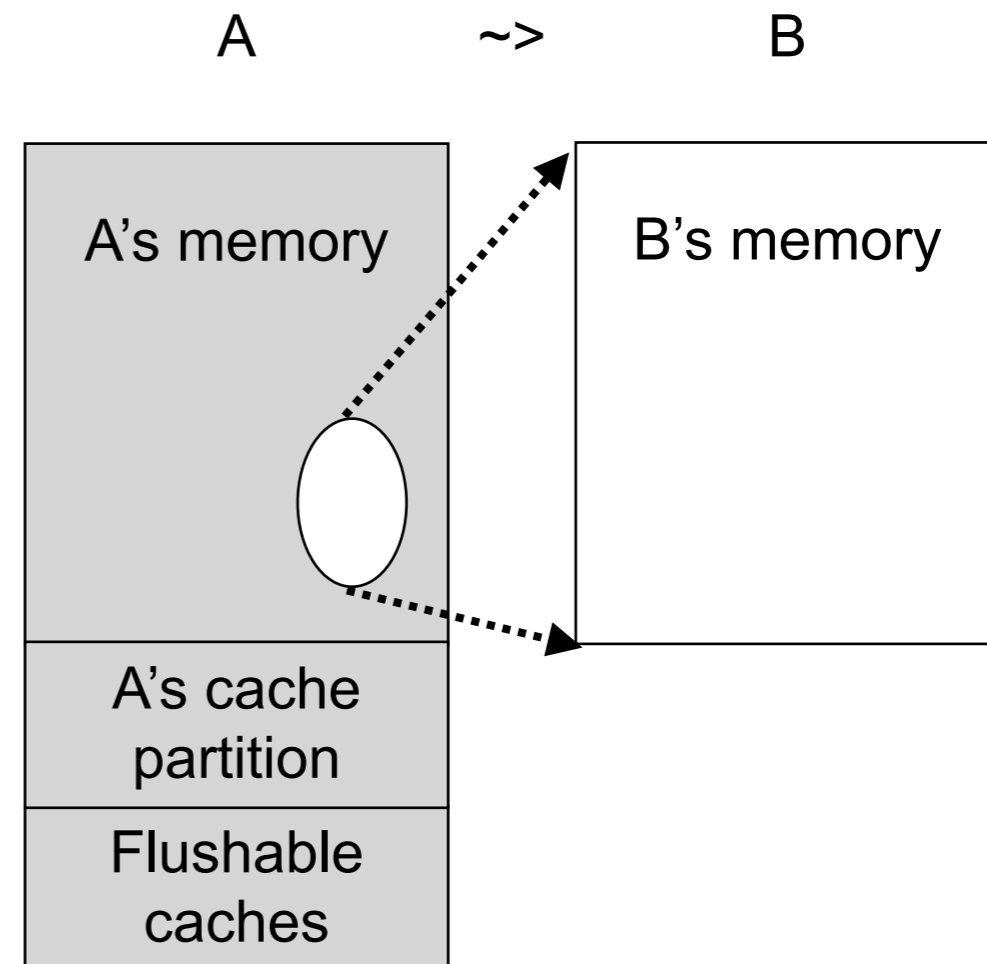
**For time protection, need  
*spatial precision to allow some flows  
but exclude others***

# OS security property



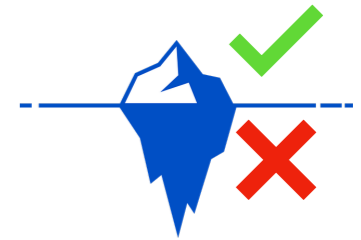
Our infowflow policies:

- *Arbitrary* spatial precision
- *Policy channels* specified as *state relations*:  $s \mid A \rightsquigarrow B \mid t$   
If  $\mid A \rightsquigarrow B \mid$  equates part of A, then info flow is allowed from there to B.
- Also arbitrary *temporal precision*
- The *dynamicity* gives us *observer-relative* properties



**For time protection, need  
spatial precision to allow some flows  
but exclude others**

# OS security property

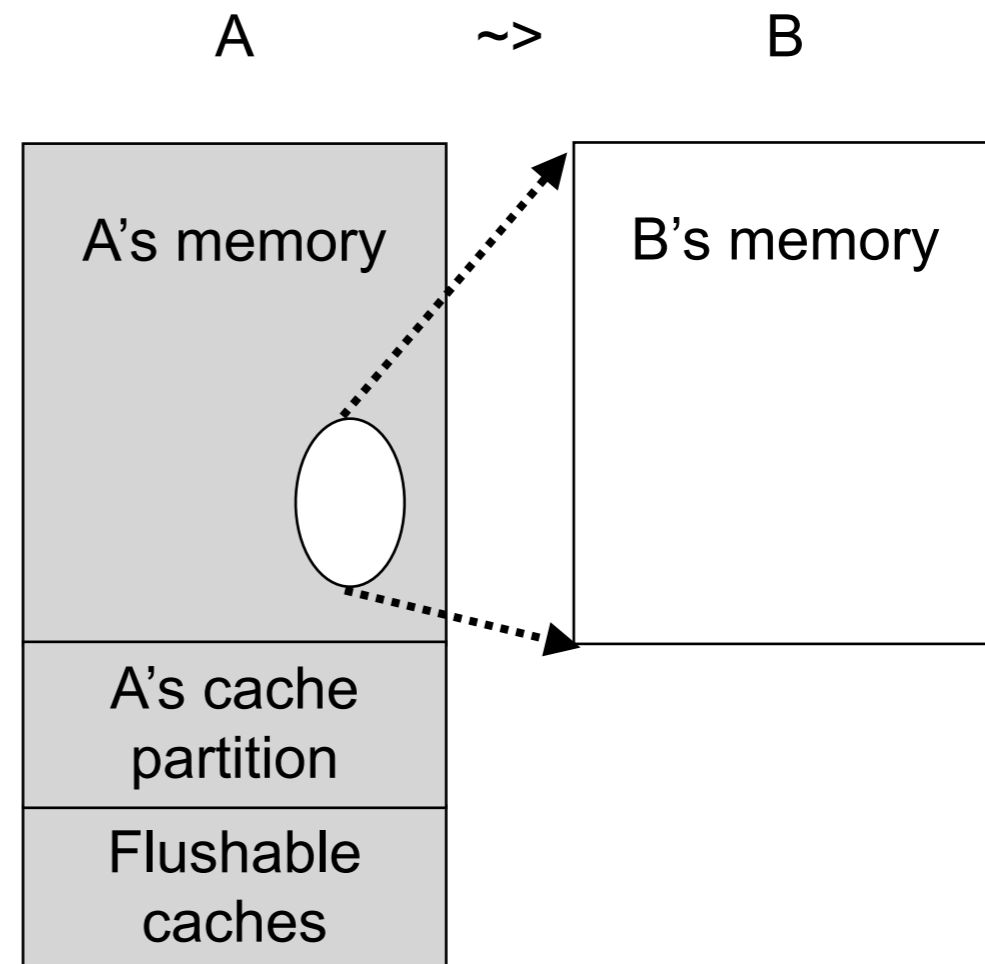


Our infowflow policies:

- *Arbitrary spatial precision*
- *Policy channels* specified as *state relations*:  $s \mid A \rightsquigarrow B \mid t$

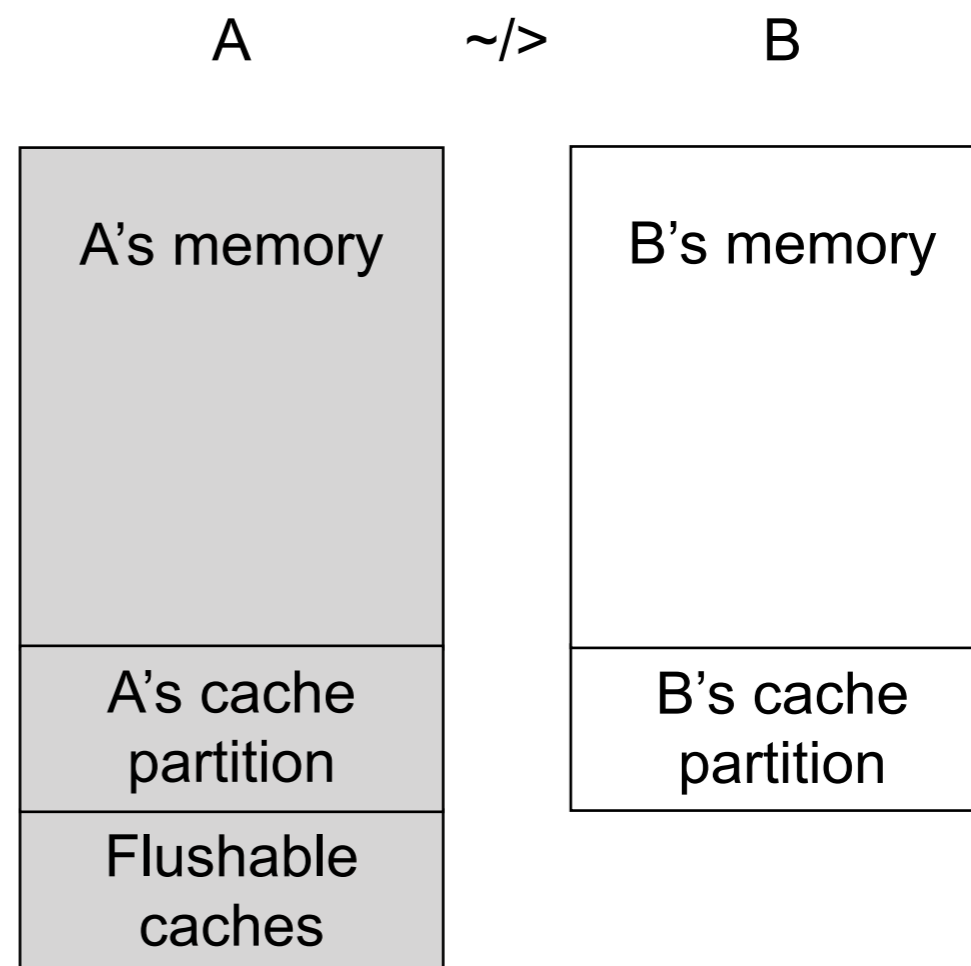
If  $\mid A \rightsquigarrow B \mid$  equates part of A, then info flow is allowed from there to B.

- Also arbitrary *temporal precision* ?
- The *dynamycity* gives us *observer-relative* properties

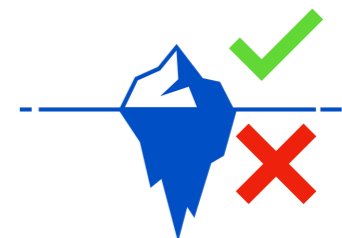


**For time protection, need *spatial precision to allow some flows but exclude others***

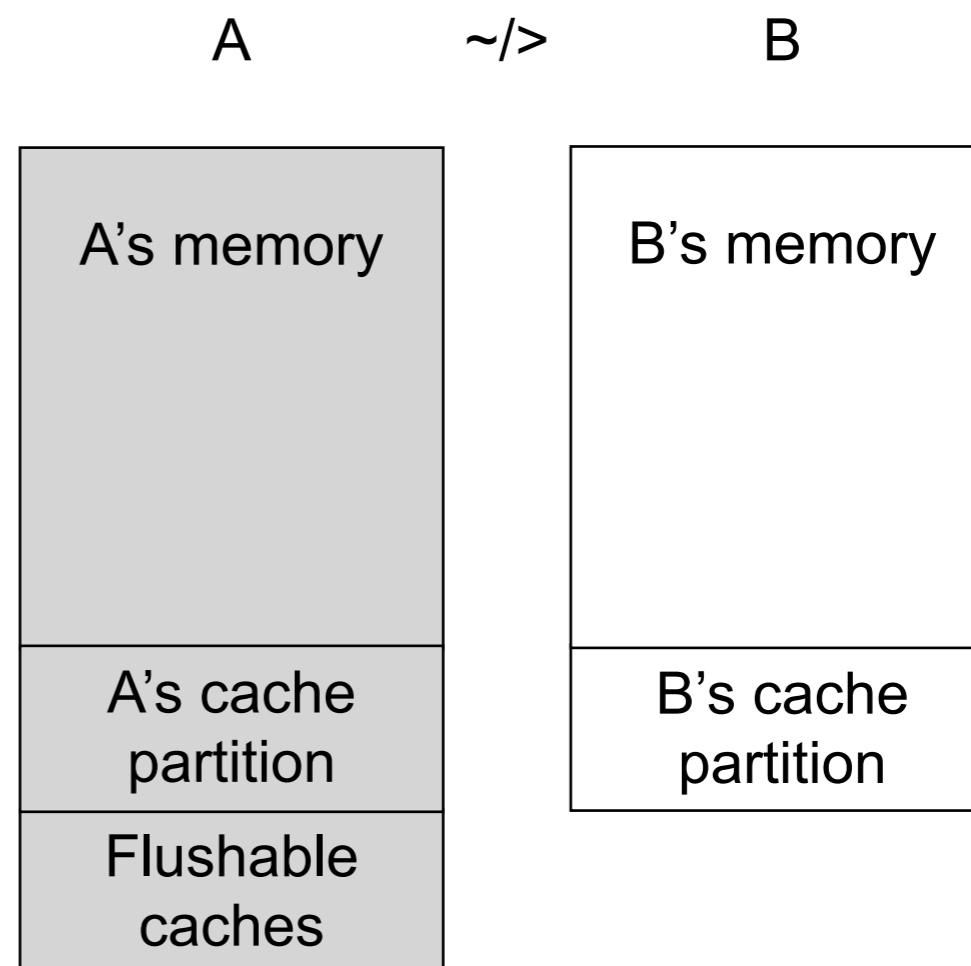
# Dynamic policy, observer relativity



# Dynamic policy, observer relativity

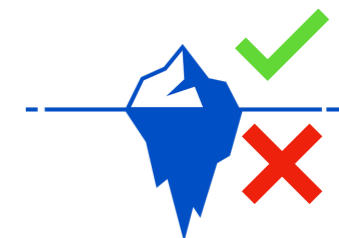


Two basic system calls:  
**Subscribe(*d*), Broadcast()**





# Dynamic policy, observer relativity

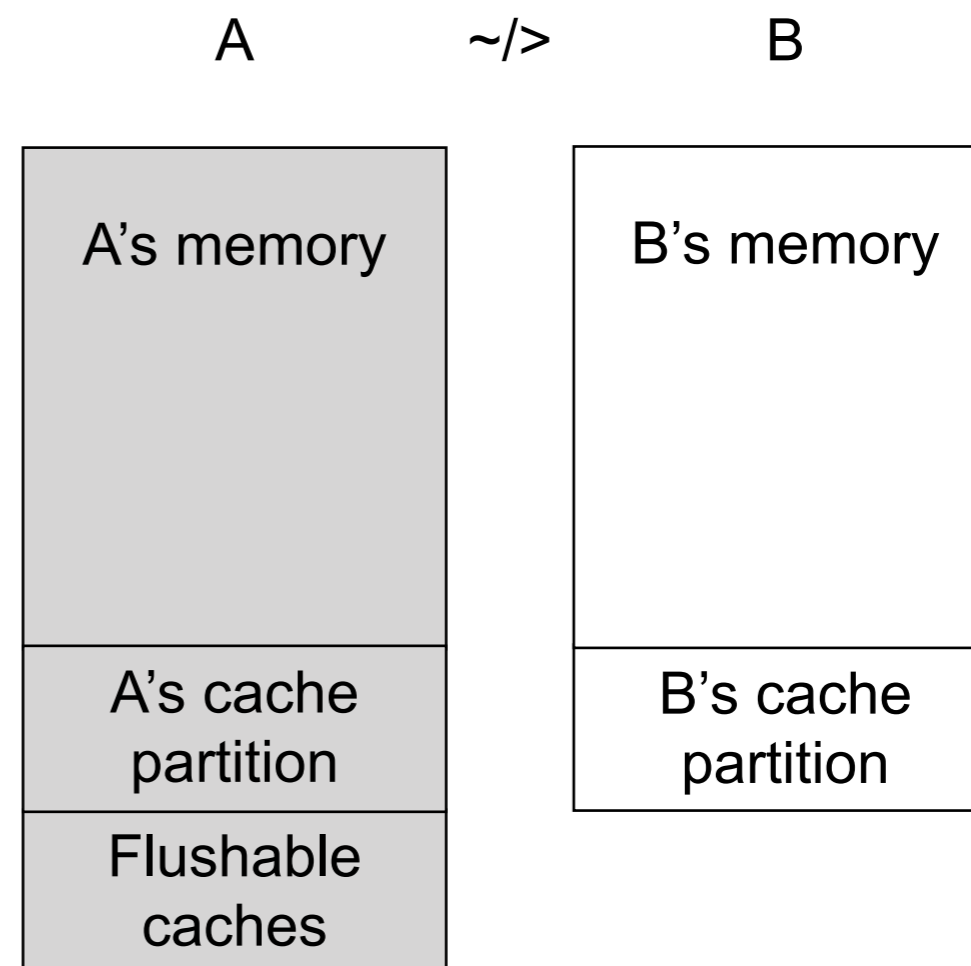


Two basic system calls:  
**Subscribe(d), Broadcast()**

1. Dynamic policy:  $A \rightsquigarrow B$  ?

✓ Only when A calls

- **Subscribe(B)**, or
- **Broadcast()** 1st time after B called **Subscribe(A)**.



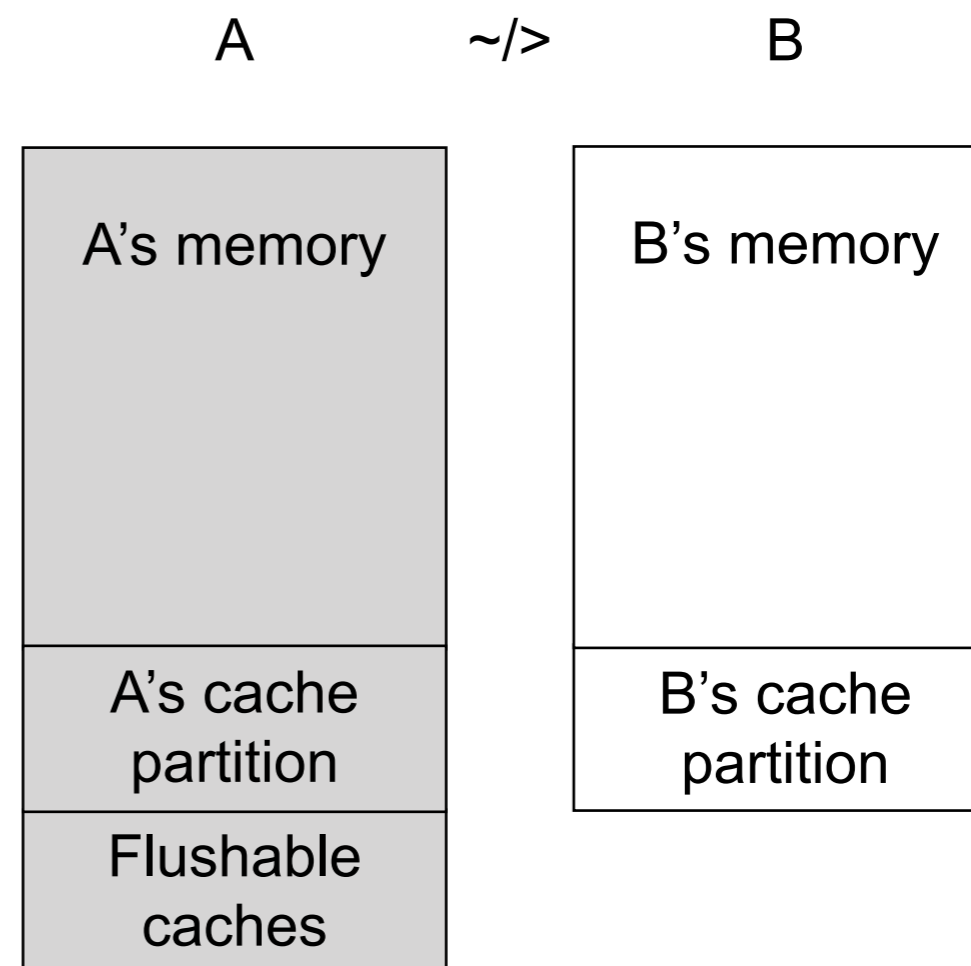
# Dynamic policy, observer relativity



Two basic system calls:  
**Subscribe(d), Broadcast()**

1. Dynamic policy:  $A \rightsquigarrow B$  ?

- ✓ Only when A calls
  - **Subscribe(B)**, or
  - **Broadcast()** 1st time after B called **Subscribe(A)**.
- ✗ Otherwise, no channel.



# Dynamic policy, observer relativity



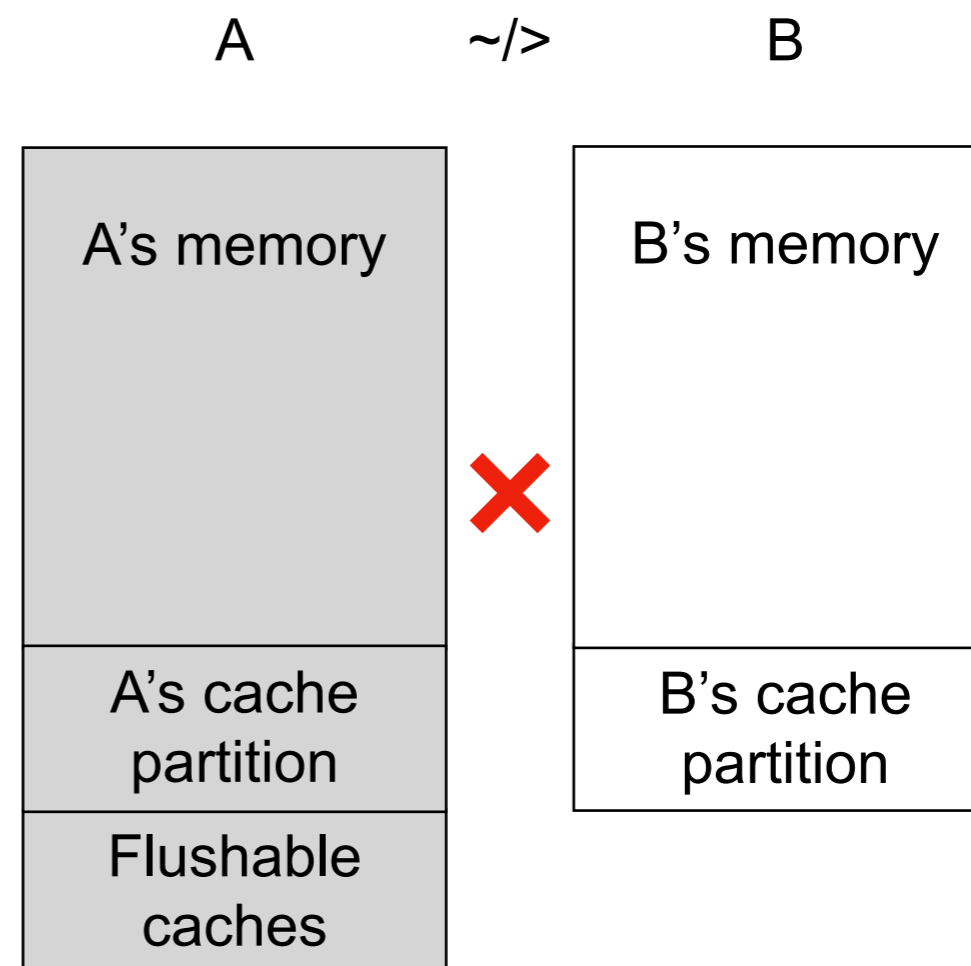
Two basic system calls:  
**Subscribe(d), Broadcast()**

1. Dynamic policy:  $A \rightsquigarrow B$  ?

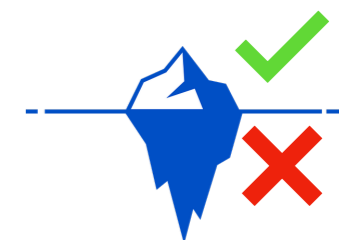
- ✓ Only when A calls
  - **Subscribe(B)**, or
  - **Broadcast()** 1st time after B called **Subscribe(A)**.

✗ Otherwise, no channel.

- Example:



# Dynamic policy, observer relativity



Two basic system calls:  
**Subscribe(d), Broadcast()**

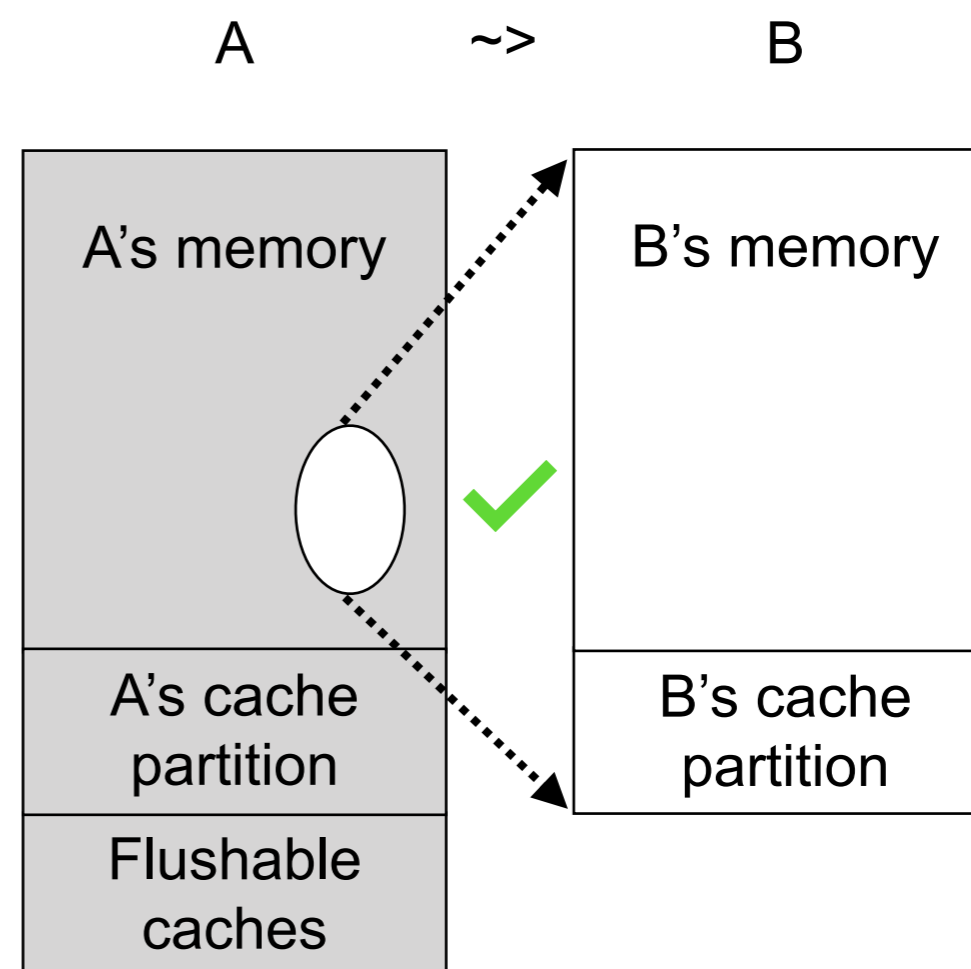
1. Dynamic policy:  $A \rightsquigarrow B$  ?

- ✓ Only when A calls
  - **Subscribe(B)**, or
  - **Broadcast()** 1st time after B called **Subscribe(A)**.

✗ Otherwise, no channel.

- Example:

1. A calls **Subscribe(B)**



# Dynamic policy, observer relativity



Two basic system calls:  
**Subscribe(d), Broadcast()**

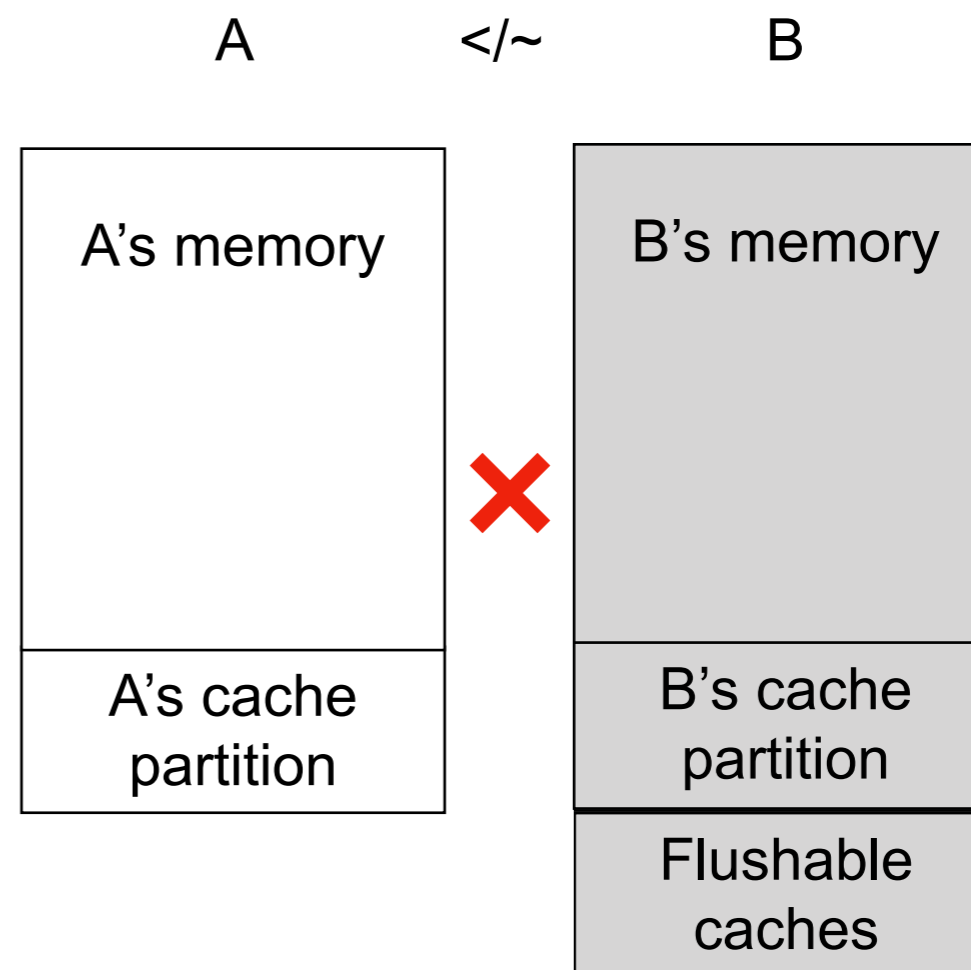
1. Dynamic policy:  $A \rightsquigarrow B$  ?

- ✓ Only when A calls
  - **Subscribe(B)**, or
  - **Broadcast()** 1st time after B called **Subscribe(A)**.

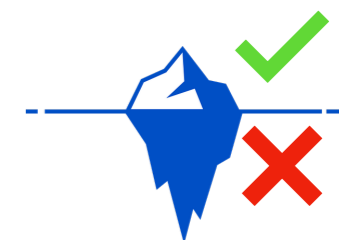
✗ Otherwise, no channel.

• Example:

1. A calls **Subscribe(B)**



# Dynamic policy, observer relativity



Two basic system calls:  
**Subscribe(d), Broadcast()**

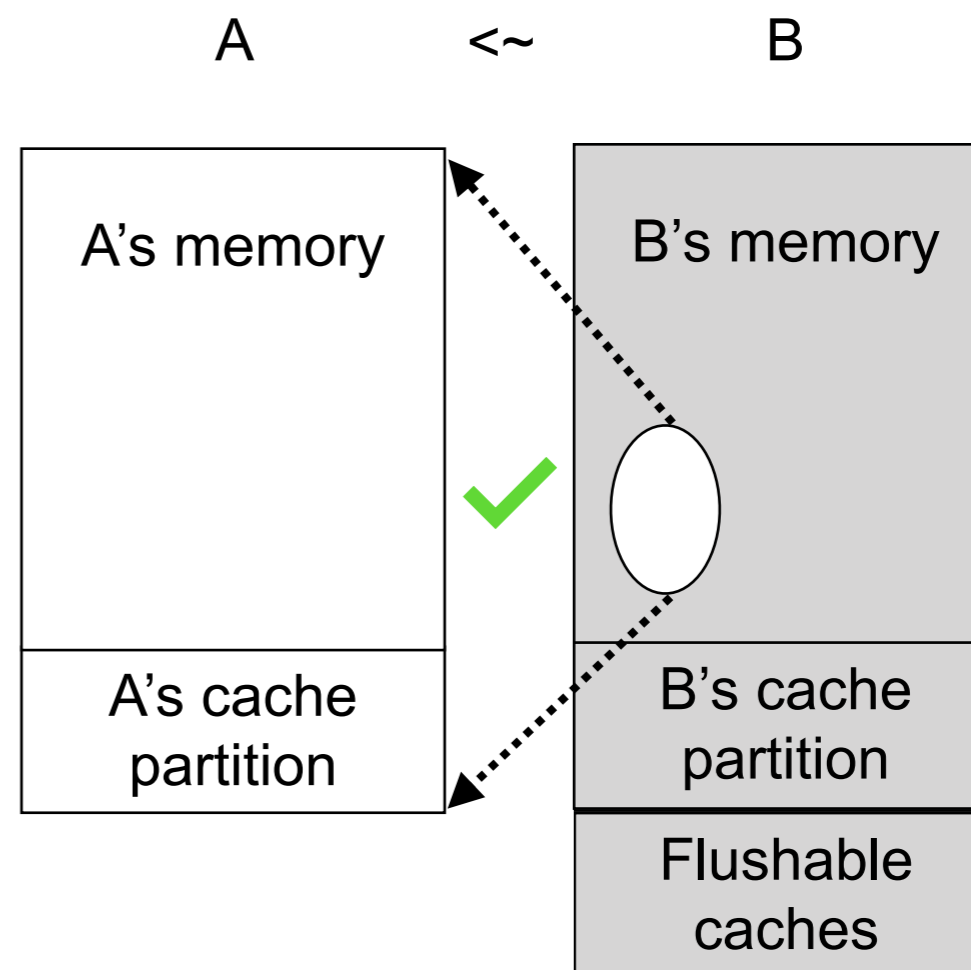
1. Dynamic policy:  $A \rightsquigarrow B$  ?

- ✓ Only when A calls
  - **Subscribe(B)**, or
  - **Broadcast()** 1st time after B called **Subscribe(A)**.

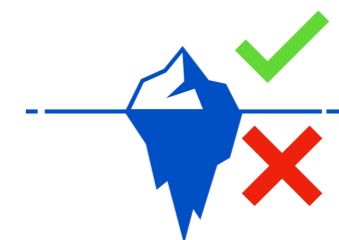
✗ Otherwise, no channel.

• Example:

1. A calls **Subscribe(B)**
2. B calls **Broadcast()**



# Dynamic policy, observer relativity



Two basic system calls:  
**Subscribe(d), Broadcast()**

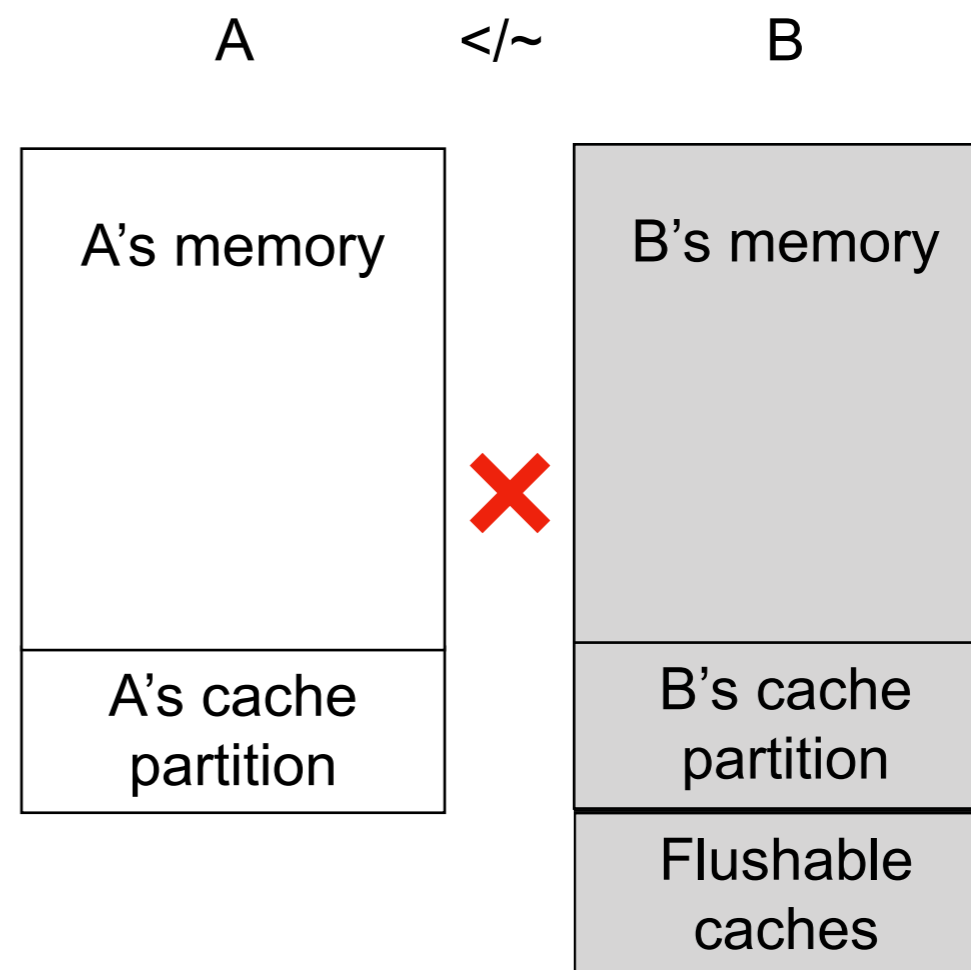
1. Dynamic policy:  $A \rightsquigarrow B$  ?

- ✓ Only when A calls
  - **Subscribe(B)**, or
  - **Broadcast()** 1st time after B called **Subscribe(A)**.

✗ Otherwise, no channel.

• Example:

1. A calls **Subscribe(B)**
2. B calls **Broadcast()**



# Dynamic policy, observer relativity



Two basic system calls:  
**Subscribe(d), Broadcast()**

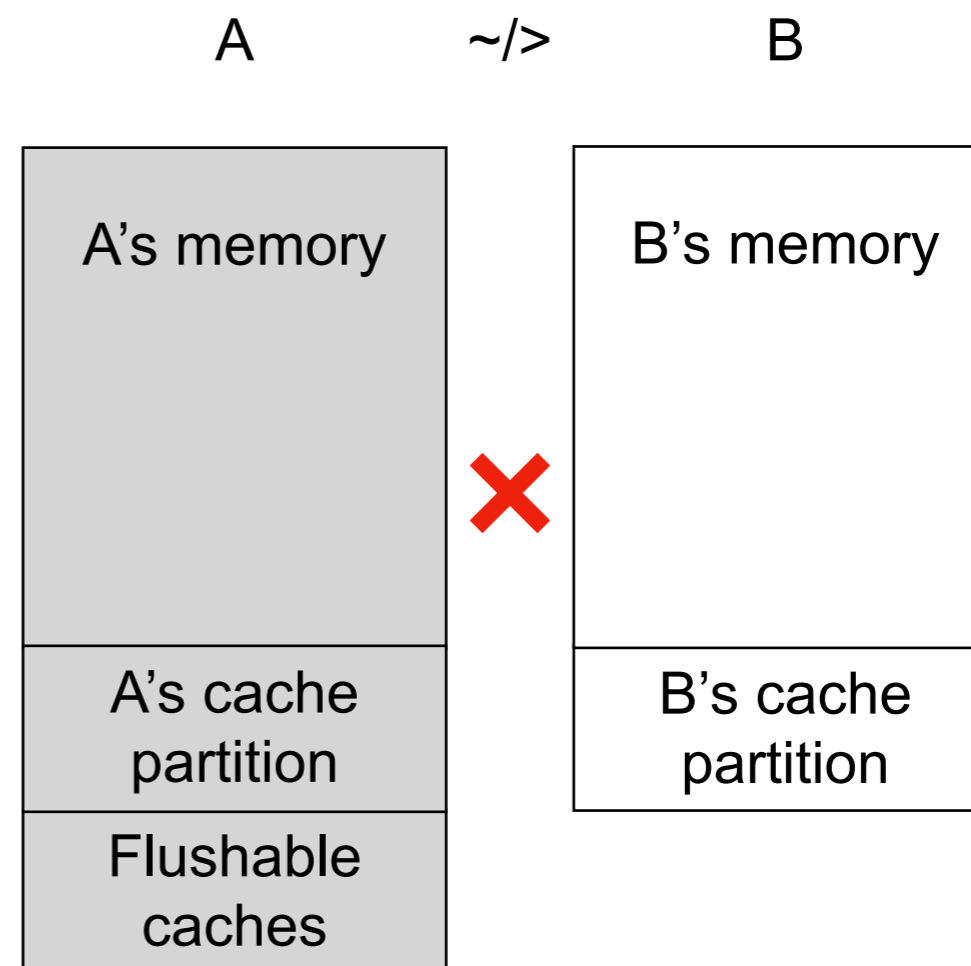
1. Dynamic policy:  $A \rightsquigarrow B$  ?

- ✓ Only when A calls
  - **Subscribe(B)**, or
  - **Broadcast()** 1st time after B called **Subscribe(A)**.

✗ Otherwise, no channel.

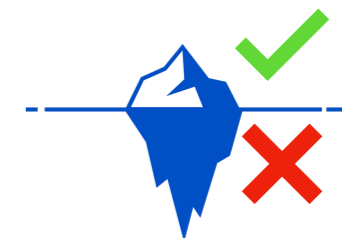
- **Example:**

1. A calls **Subscribe(B)**
2. B calls **Broadcast()**





# Dynamic policy, observer relativity



Two basic system calls:  
**Subscribe(d), Broadcast()**

2. Property must be observer relative!

1. Dynamic policy:  $A \sim > B$  ?

- ✓ Only when A calls
  - **Subscribe(B)**, or
  - **Broadcast()** 1st time after B called **Subscribe(A)**.

✗ Otherwise, no channel.

- Example:

1. A calls **Subscribe(B)**

2. B calls **Broadcast()**

# Dynamic policy, observer relativity



Two basic system calls:  
**Subscribe(d), Broadcast()**

1. Dynamic policy:  $A \sim > B$  ?

- ✓ Only when A calls
  - **Subscribe(B)**, or
  - **Broadcast()** 1st time after B called **Subscribe(A)**.

✗ Otherwise, no channel.

- Example:

1. A calls **Subscribe(B)**

2. B calls **Broadcast()**

2. Property must be observer relative!

- If not, can't prove the (bisimulation) property for unrelated user C!  
(i.e. where  $A \sim /> C$ ,  $B \sim /> C$ )

# Dynamic policy, observer relativity



Two basic system calls:  
**Subscribe(d), Broadcast()**

1. Dynamic policy:  $A \sim > B$  ?

✓ Only when A calls

- **Subscribe(B)**, or
- **Broadcast()** 1st time after B called **Subscribe(A)**.

✗ Otherwise, no channel.

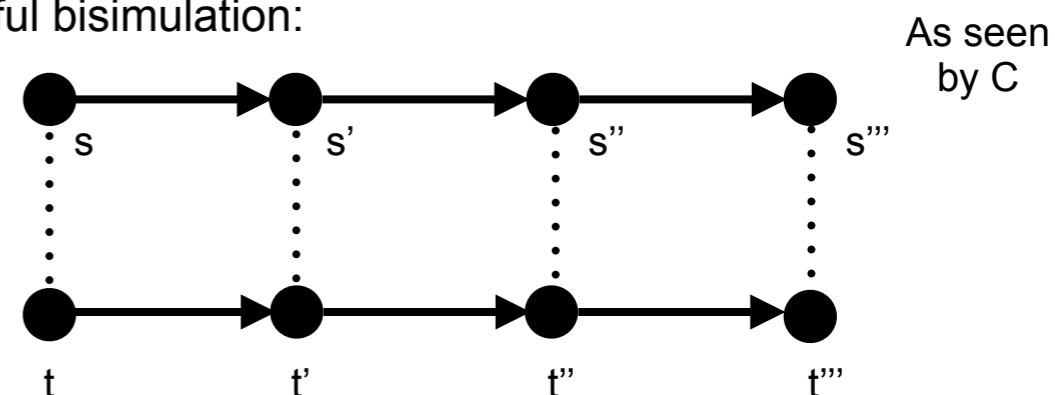
• Example:

1. A calls **Subscribe(B)**
2. B calls **Broadcast()**

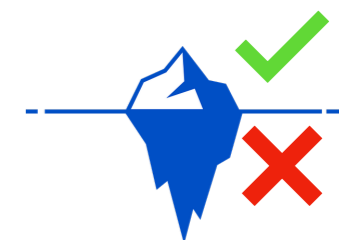
2. Property must be observer relative!

- If not, can't prove the (bisimulation) property for unrelated user C!  
(i.e. where  $A \sim > C$ ,  $B \sim > C$ )

Successful bisimulation:



# Dynamic policy, observer relativity



Two basic system calls:  
**Subscribe(d), Broadcast()**

1. Dynamic policy:  $A \sim > B$  ?

✓ Only when A calls

- **Subscribe(B)**, or
- **Broadcast()** 1st time after B called **Subscribe(A)**.

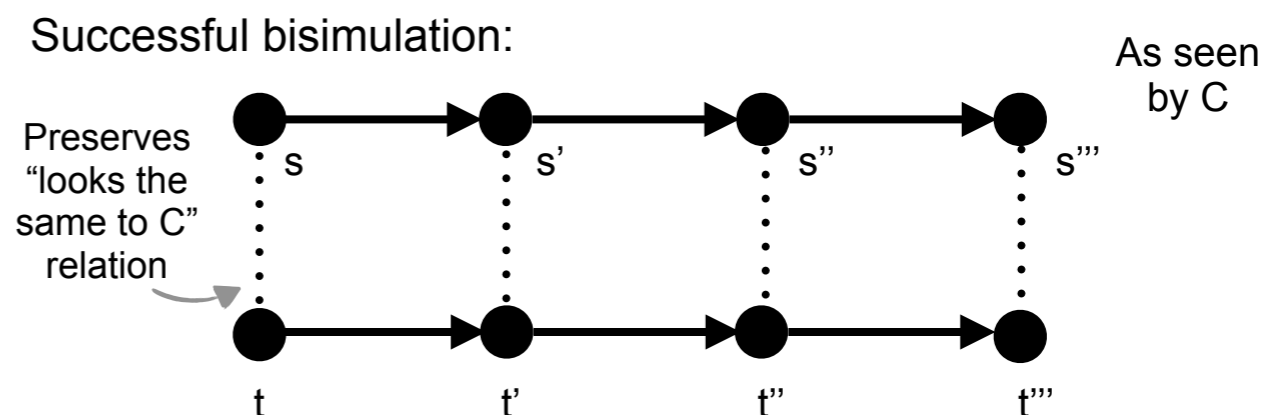
✗ Otherwise, no channel.

• Example:

1. A calls **Subscribe(B)**
2. B calls **Broadcast()**

2. Property must be observer relative!

- If not, can't prove the (bisimulation) property for unrelated user C!  
(i.e. where  $A \sim > C$ ,  $B \sim > C$ )



# Dynamic policy, observer relativity



Two basic system calls:  
**Subscribe(d), Broadcast()**

1. Dynamic policy:  $A \sim > B$  ?

✓ Only when A calls

- **Subscribe(B)**, or
- **Broadcast()** 1st time after B called **Subscribe(A)**.

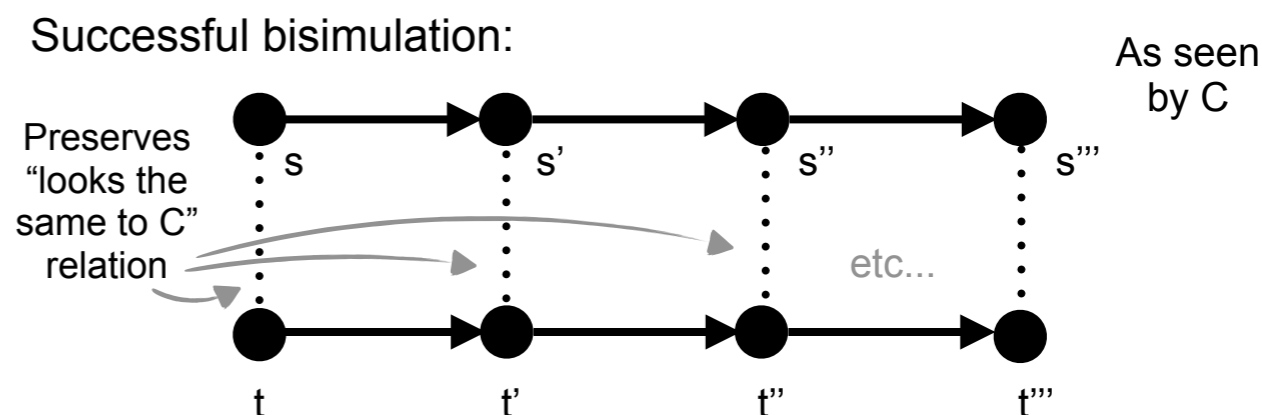
✗ Otherwise, no channel.

• Example:

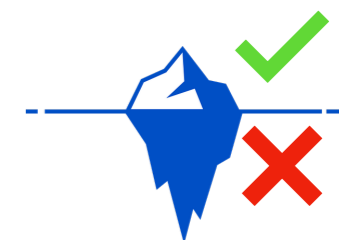
1. A calls **Subscribe(B)**
2. B calls **Broadcast()**

2. Property must be observer relative!

- If not, can't prove the (bisimulation) property for unrelated user C!  
(i.e. where  $A \sim > C$ ,  $B \sim > C$ )



# Dynamic policy, observer relativity



Two basic system calls:  
**Subscribe(d), Broadcast()**

1. Dynamic policy:  $A \rightsquigarrow B$  ?

- ✓ Only when A calls
  - **Subscribe(B)**, or
  - **Broadcast()** 1st time after B called **Subscribe(A)**.

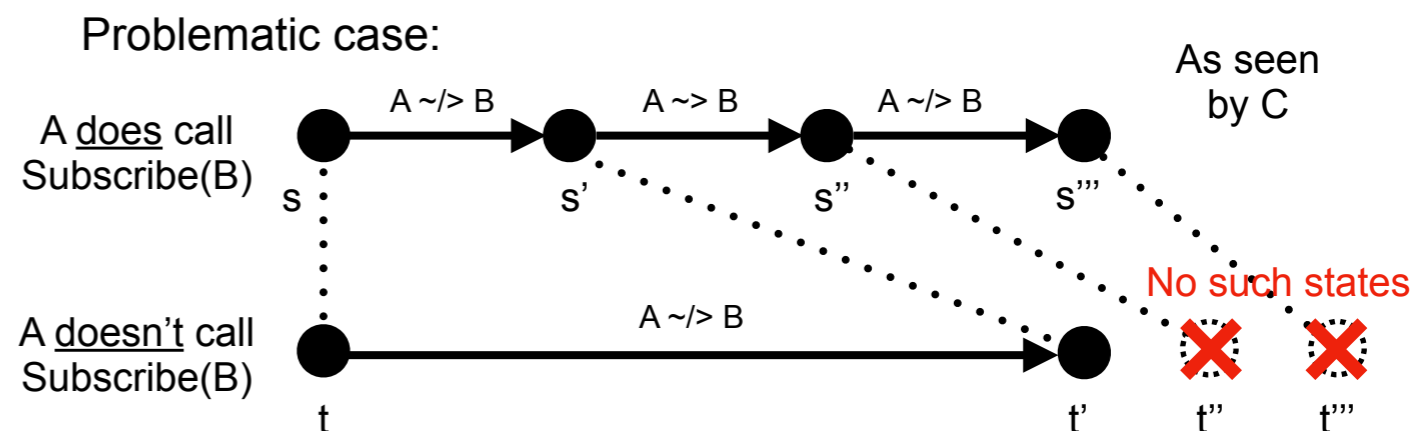
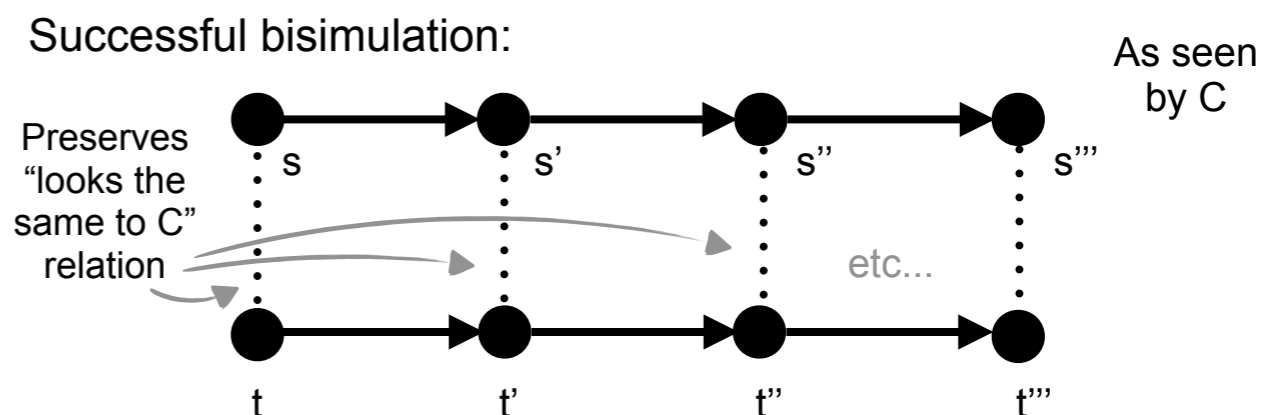
✗ Otherwise, no channel.

• Example:

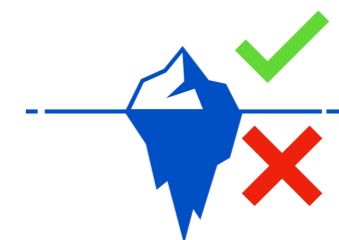
1. A calls **Subscribe(B)**
2. B calls **Broadcast()**

2. Property must be observer relative!

- If not, can't prove the (bisimulation) property for unrelated user C!  
 (i.e. where  $A \rightsquigarrow C$ ,  $B \rightsquigarrow C$ )



# Dynamic policy, observer relativity



Two basic system calls:  
**Subscribe(d), Broadcast()**

1. Dynamic policy:  $A \rightsquigarrow B$  ?

✓ Only when A calls

- **Subscribe(B)**, or
- **Broadcast()** 1st time after B called **Subscribe(A)**.

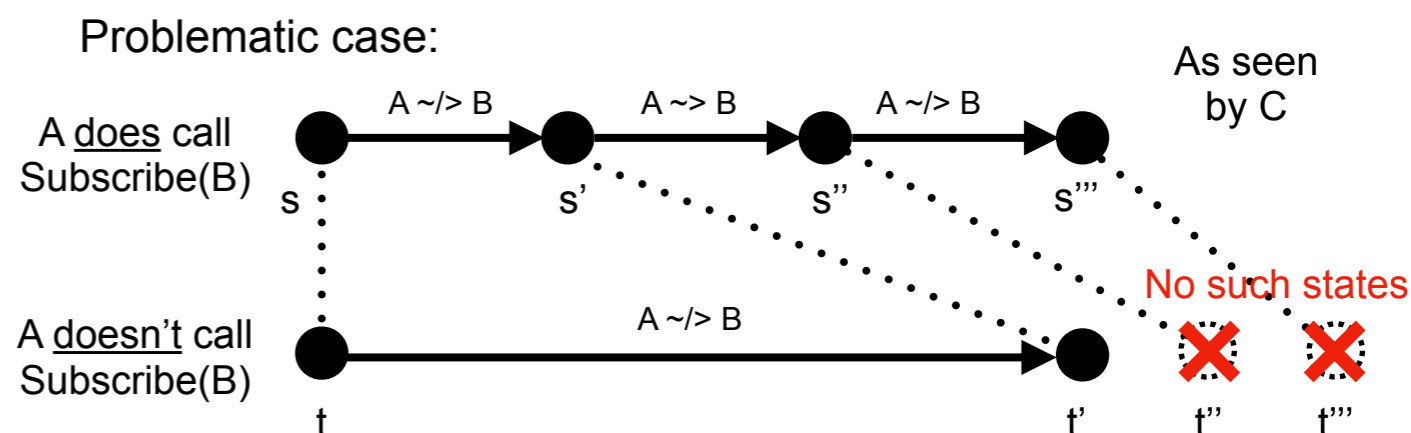
✗ Otherwise, no channel.

• Example:

1. A calls **Subscribe(B)**
2. B calls **Broadcast()**

2. Property must be observer relative!

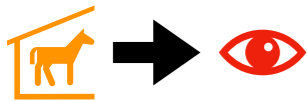
- If not, can't prove the (bisimulation) property for unrelated user C!  
(i.e. where  $A \rightsquigarrow C$ ,  $B \rightsquigarrow C$ )



- *Solution*: C's property must treat *states* (in the state machine) as observable only whenever
  - C is running, or
  - When  $d$  is running,  $d \rightsquigarrow C$ .

# How to formalise an OS enforces *time protection*?

Versus threat scenario:  
trojan and spy



Abstract *covert state* + *time* to reflect strategies enabled by HW:  
Partition or flush state; pad time.



Demonstrating these principles, we formalised in Isabelle/HOL:

1. OS security model imposing requirements on relevant parts of OS. (Intended for seL4, but *generic*)
3. Proof our security property holds if OS model's requirements hold.



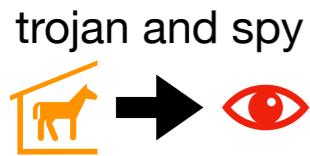
Make security property precise enough to exclude flows from covert state.

2. OS security property that is dynamic; this makes it observer relative. (Improving on seL4's of [Murray et al. 2012])
4. Basic instantiation of OS model exercising dynamic policy.



# How to formalise an OS enforces *time protection*?

Versus threat scenario:



Thank you!  
Q & A

Abstract *covert state* + *time* to reflect strategies enabled by HW:

Partition or flush state; pad time.

Make security property precise enough to exclude flows from covert state.



Demonstrating these principles, we formalised in Isabelle/HOL:

**Paper:** <https://doi.org/jzww>  
**Artifact:** <https://doi.org/jzwwk>

1. OS security model imposing requirements on relevant parts of OS.  
(Intended for seL4, but *generic*)

3. Proof our security property holds if OS model's requirements hold.

2. OS security property that is dynamic; this makes it observer relative.  
(Improving on seL4's of [Murray et al. 2012])

4. Basic instantiation of OS model exercising dynamic policy.