



School of Computer Science & Engineering
Trustworthy Systems Group

Clearing roadblocks to Time Protection verification for seL4

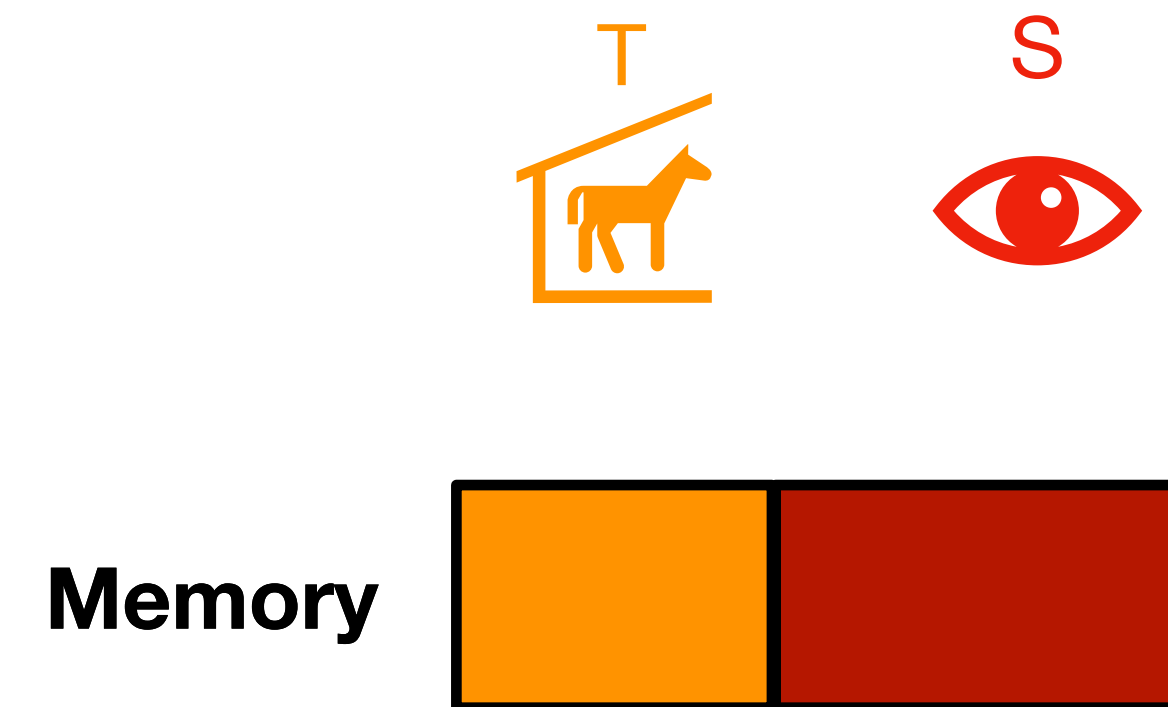
Dr Rob Sison

Senior Research Associate, UNSW Sydney

r.sison@unsw.edu.au



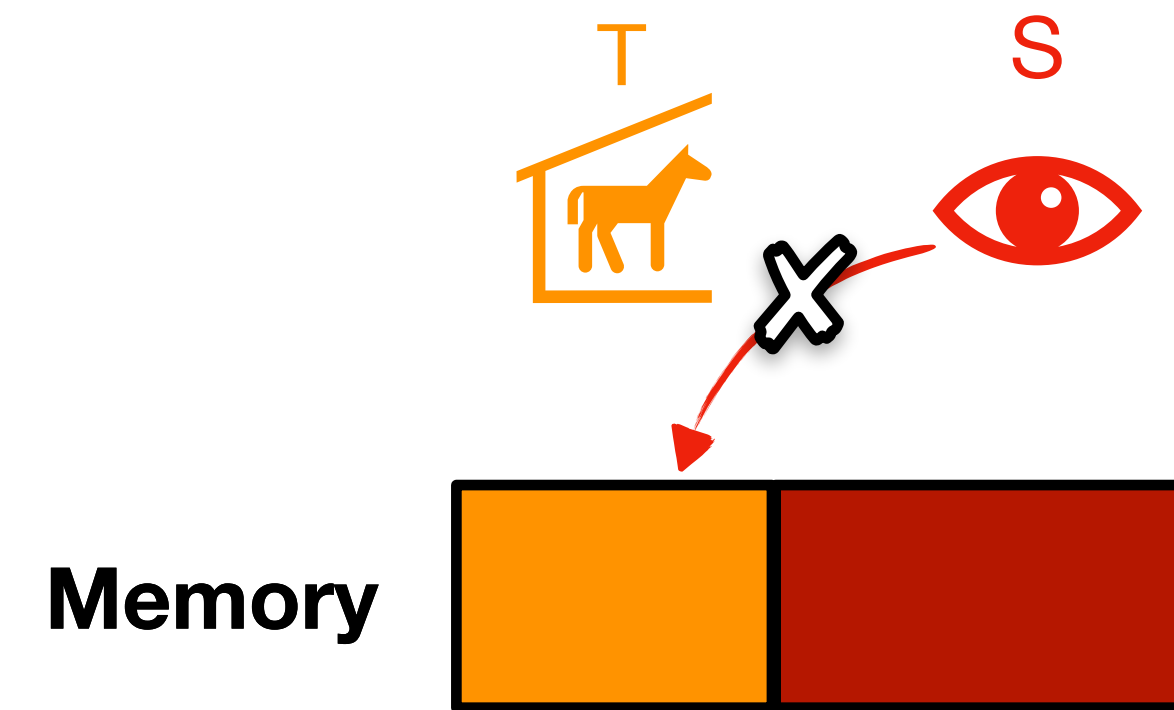
What is Time Protection?



What is Time Protection?



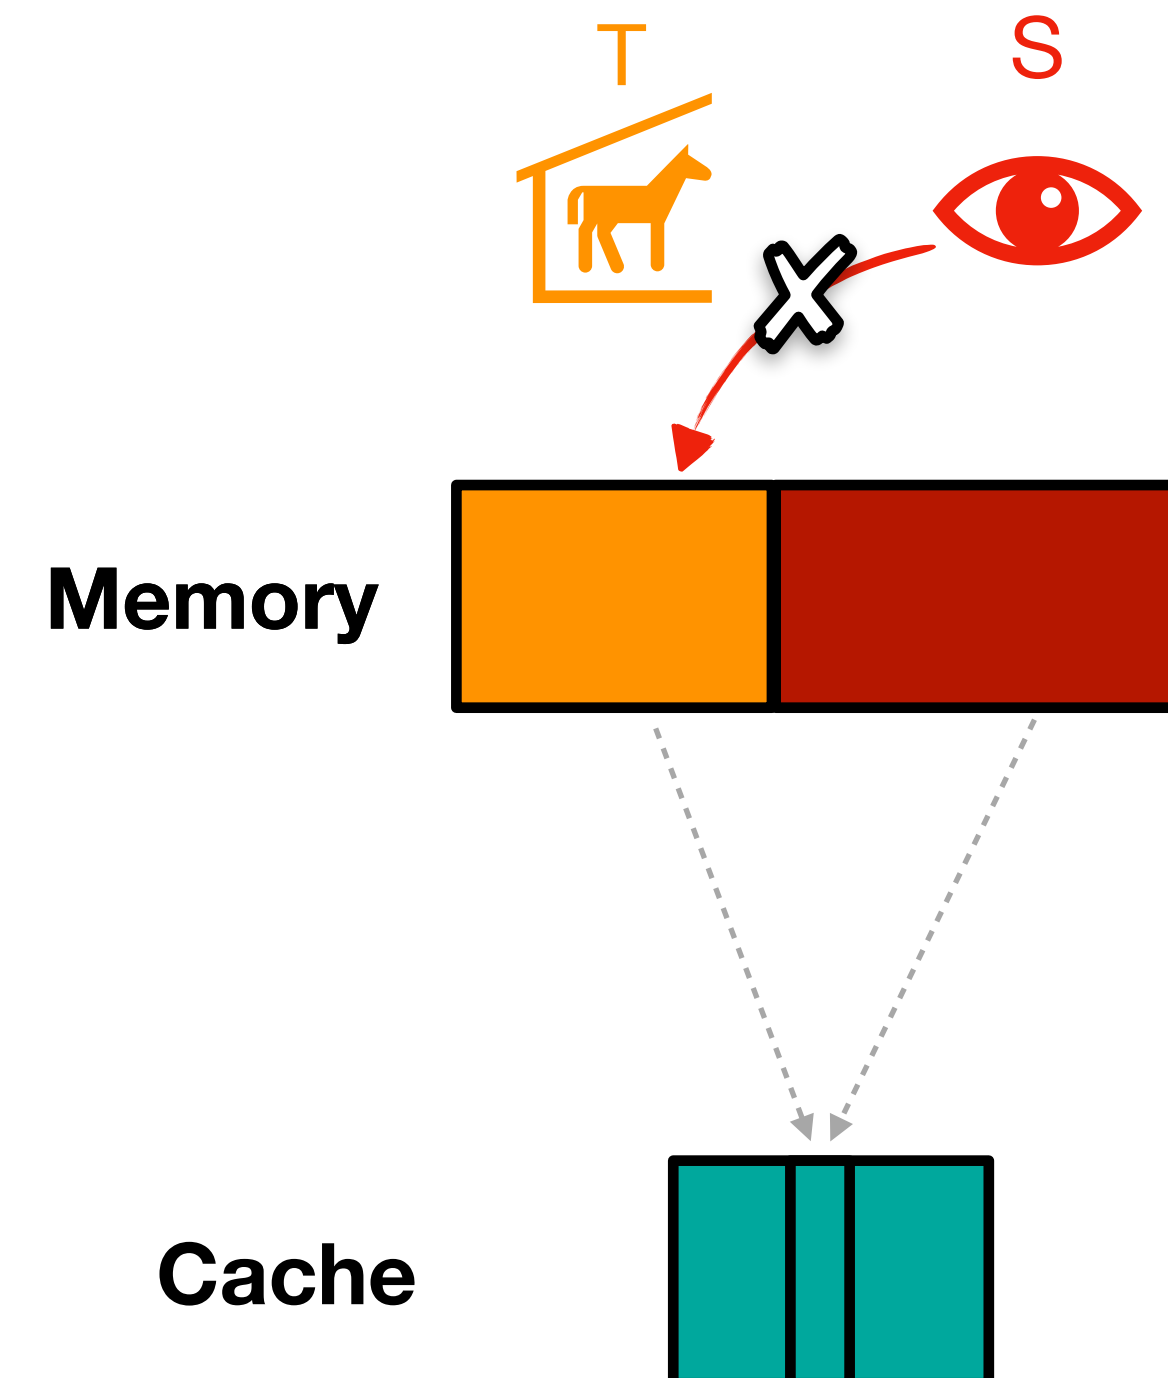
- OSes typically implement *memory protection*.



What is Time Protection?



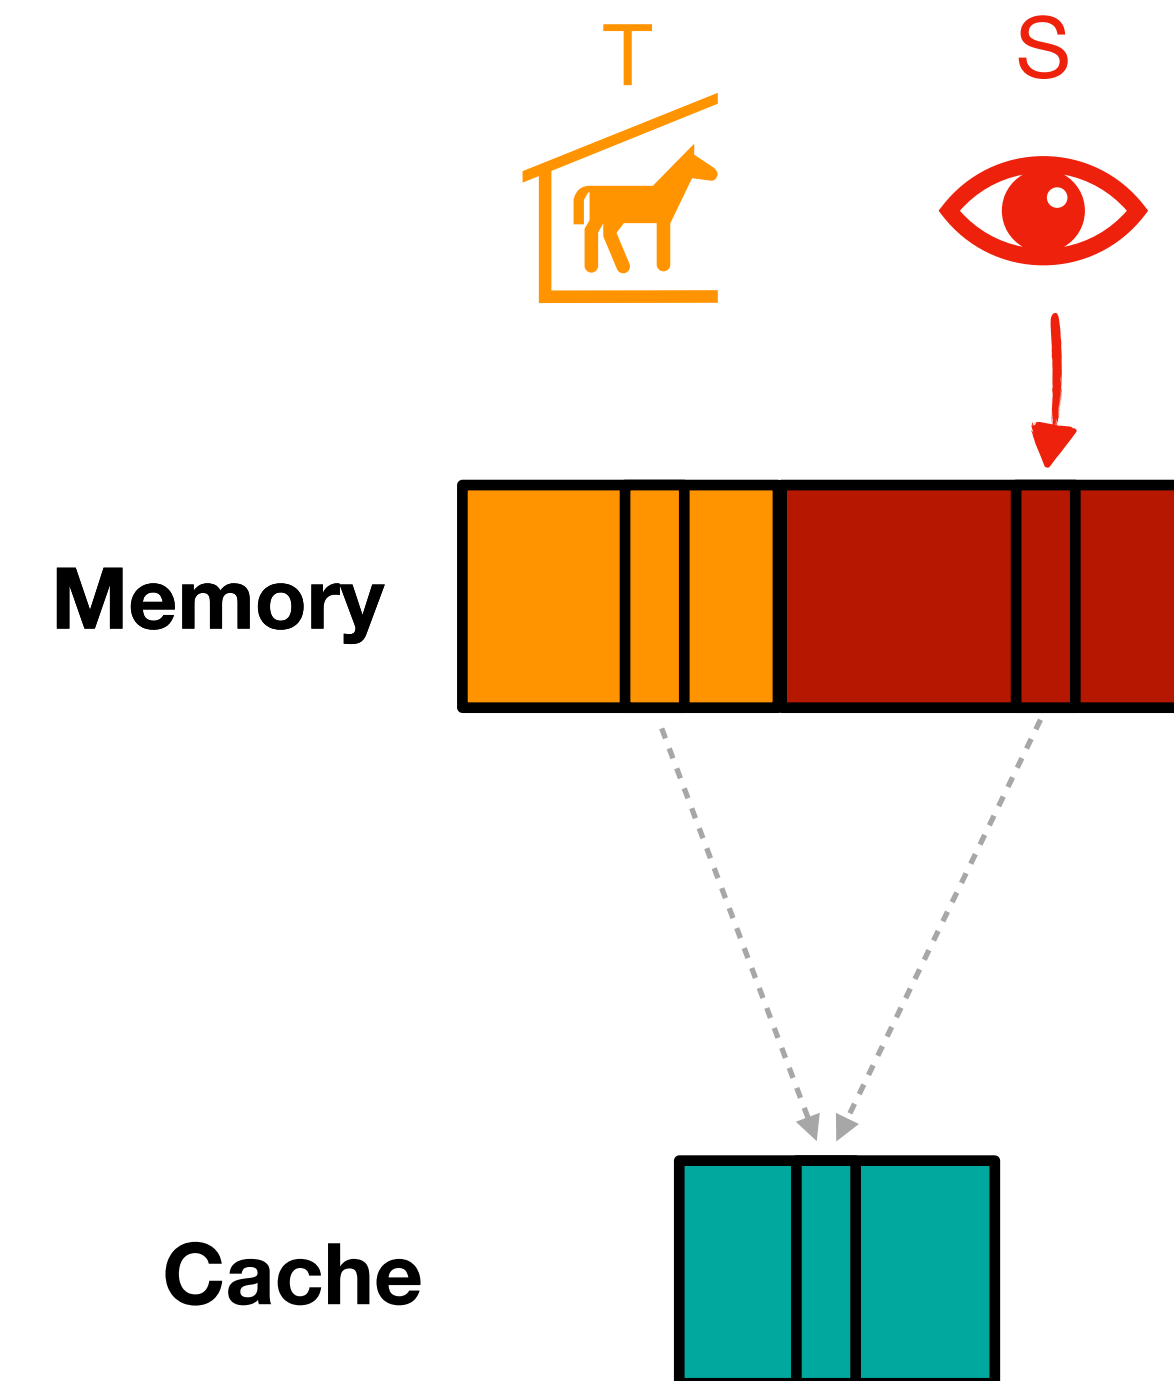
- OSes typically implement *memory protection*.
- But: Mere memory access changes microarchitectural state — this affects timing.



What is Time Protection?



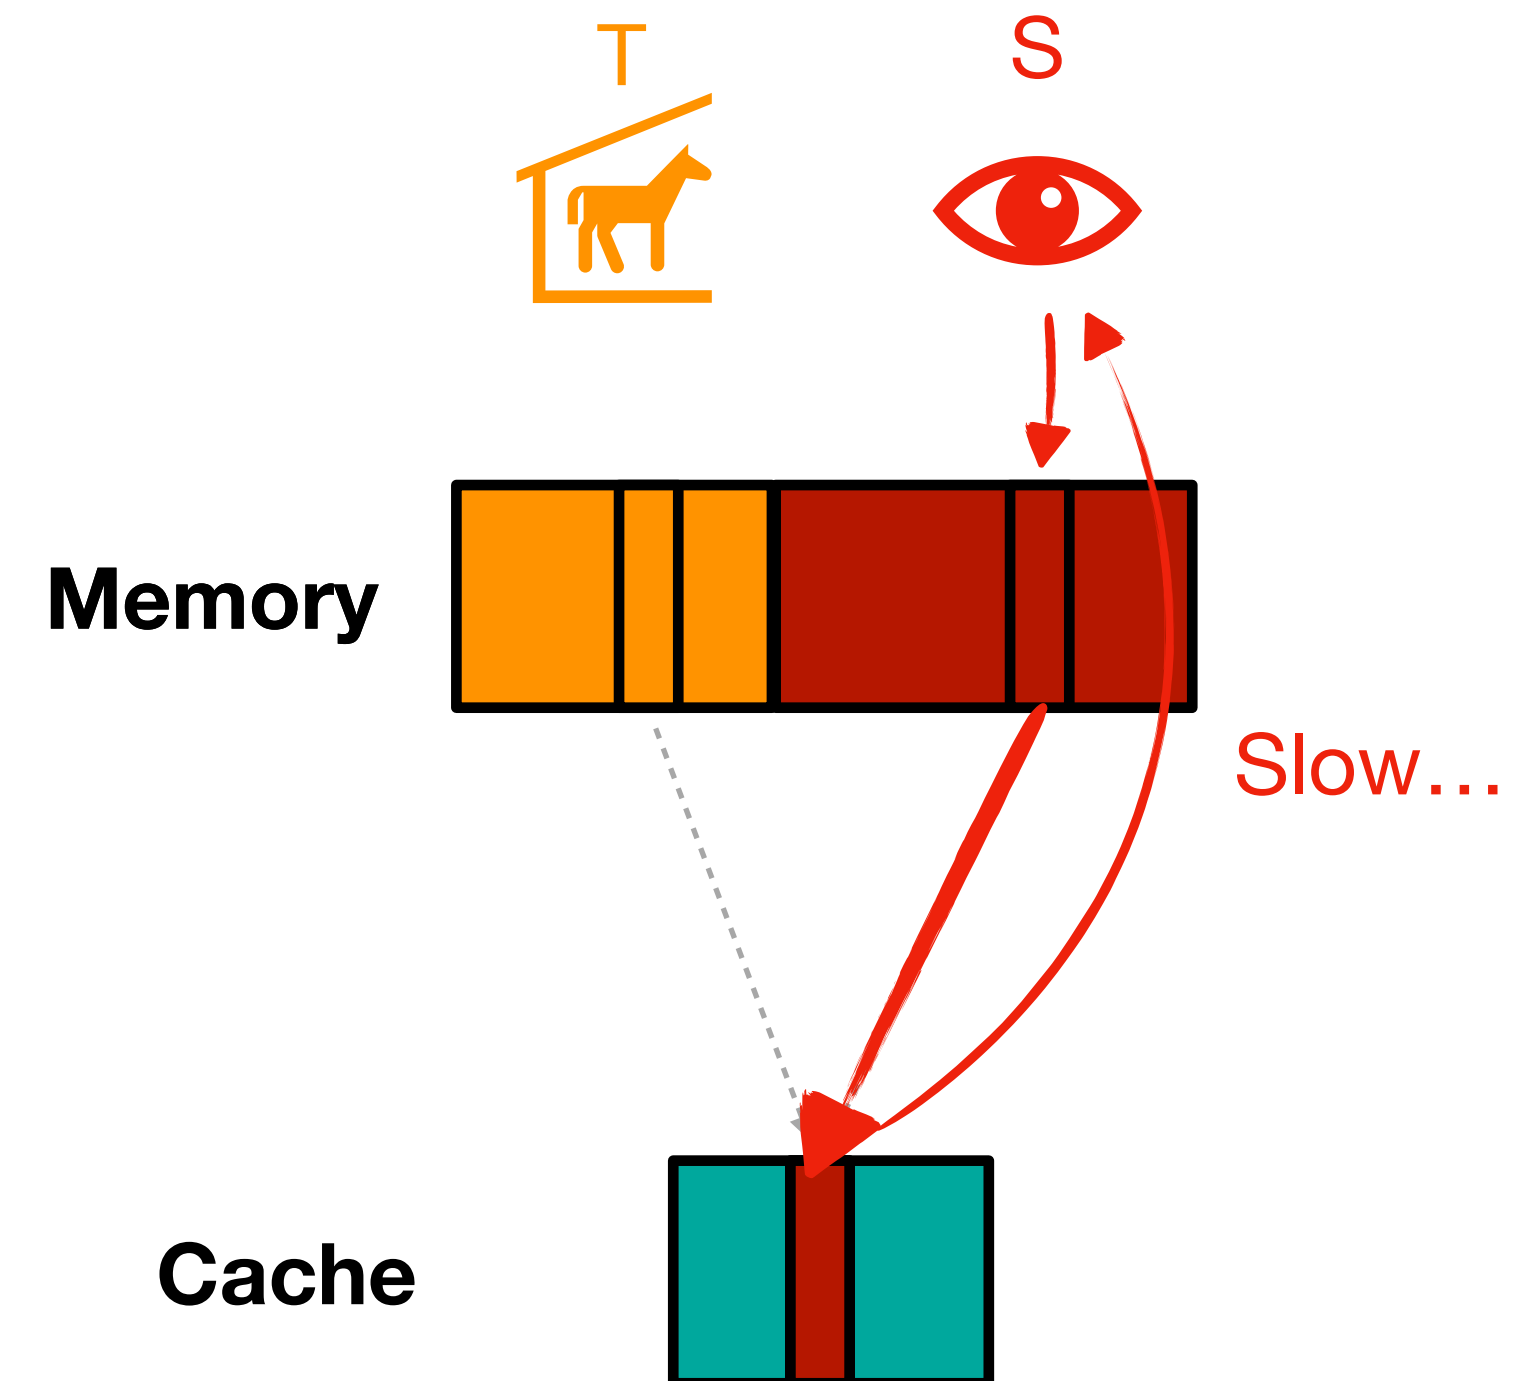
- OSes typically implement *memory protection*.
- But: Mere memory access changes microarchitectural state — this affects timing.



What is Time Protection?



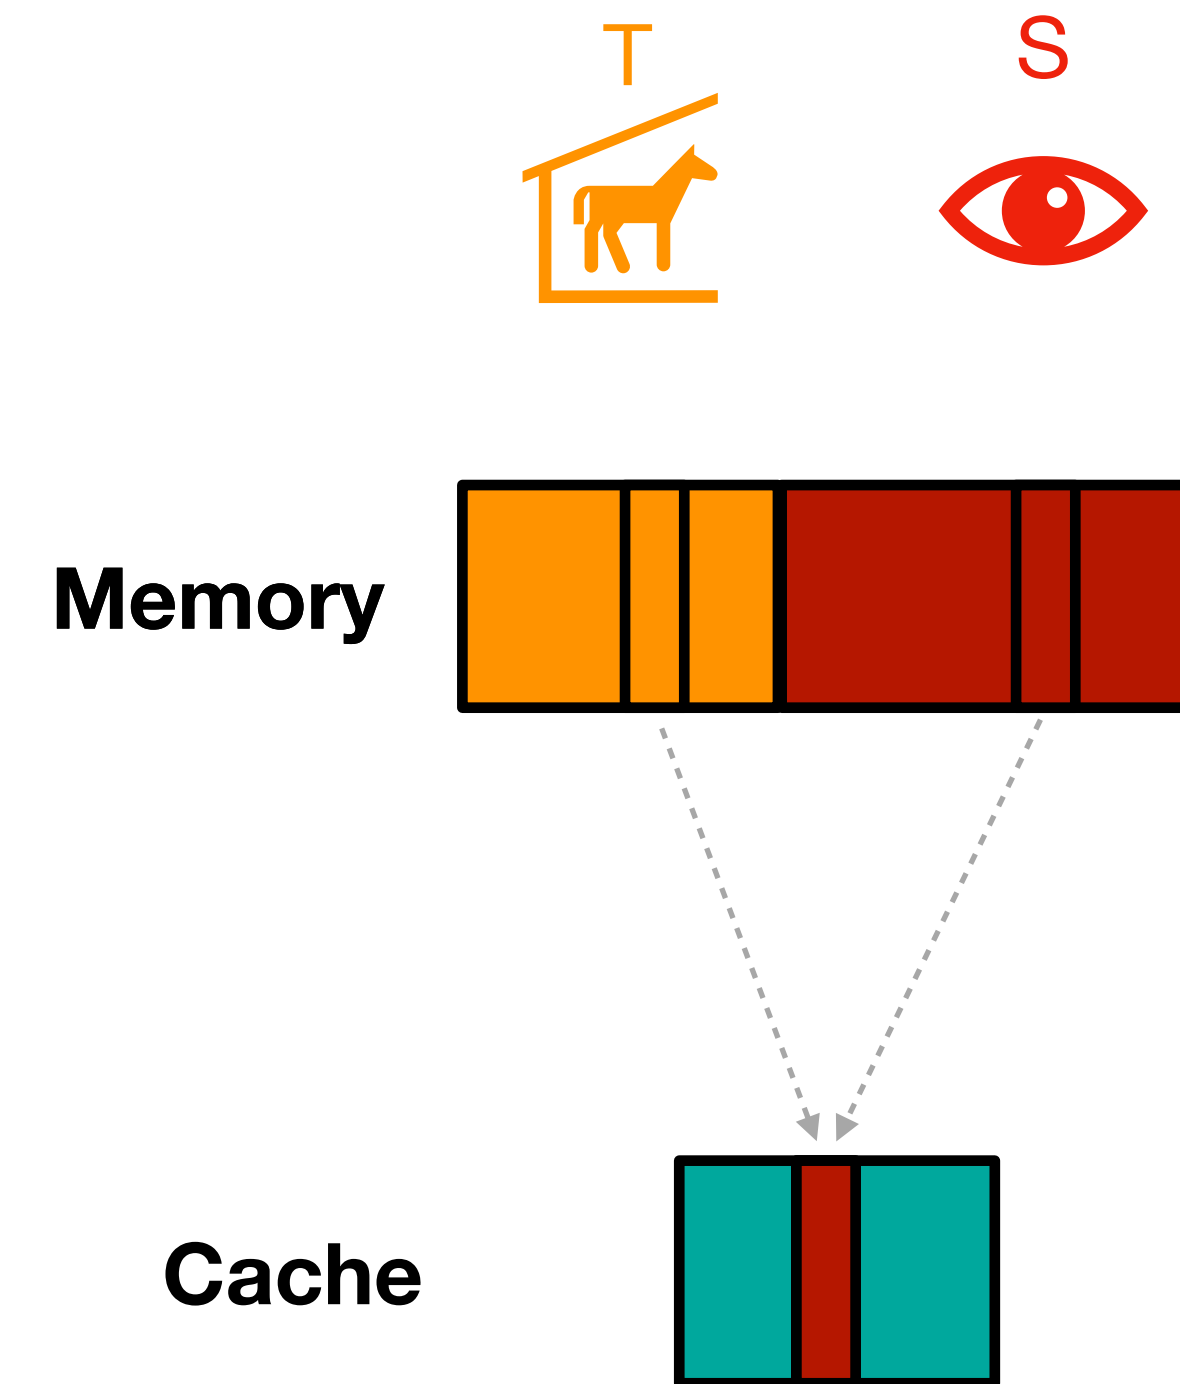
- OSes typically implement *memory protection*.
- But: Mere memory access changes microarchitectural state — this affects timing.



What is Time Protection?



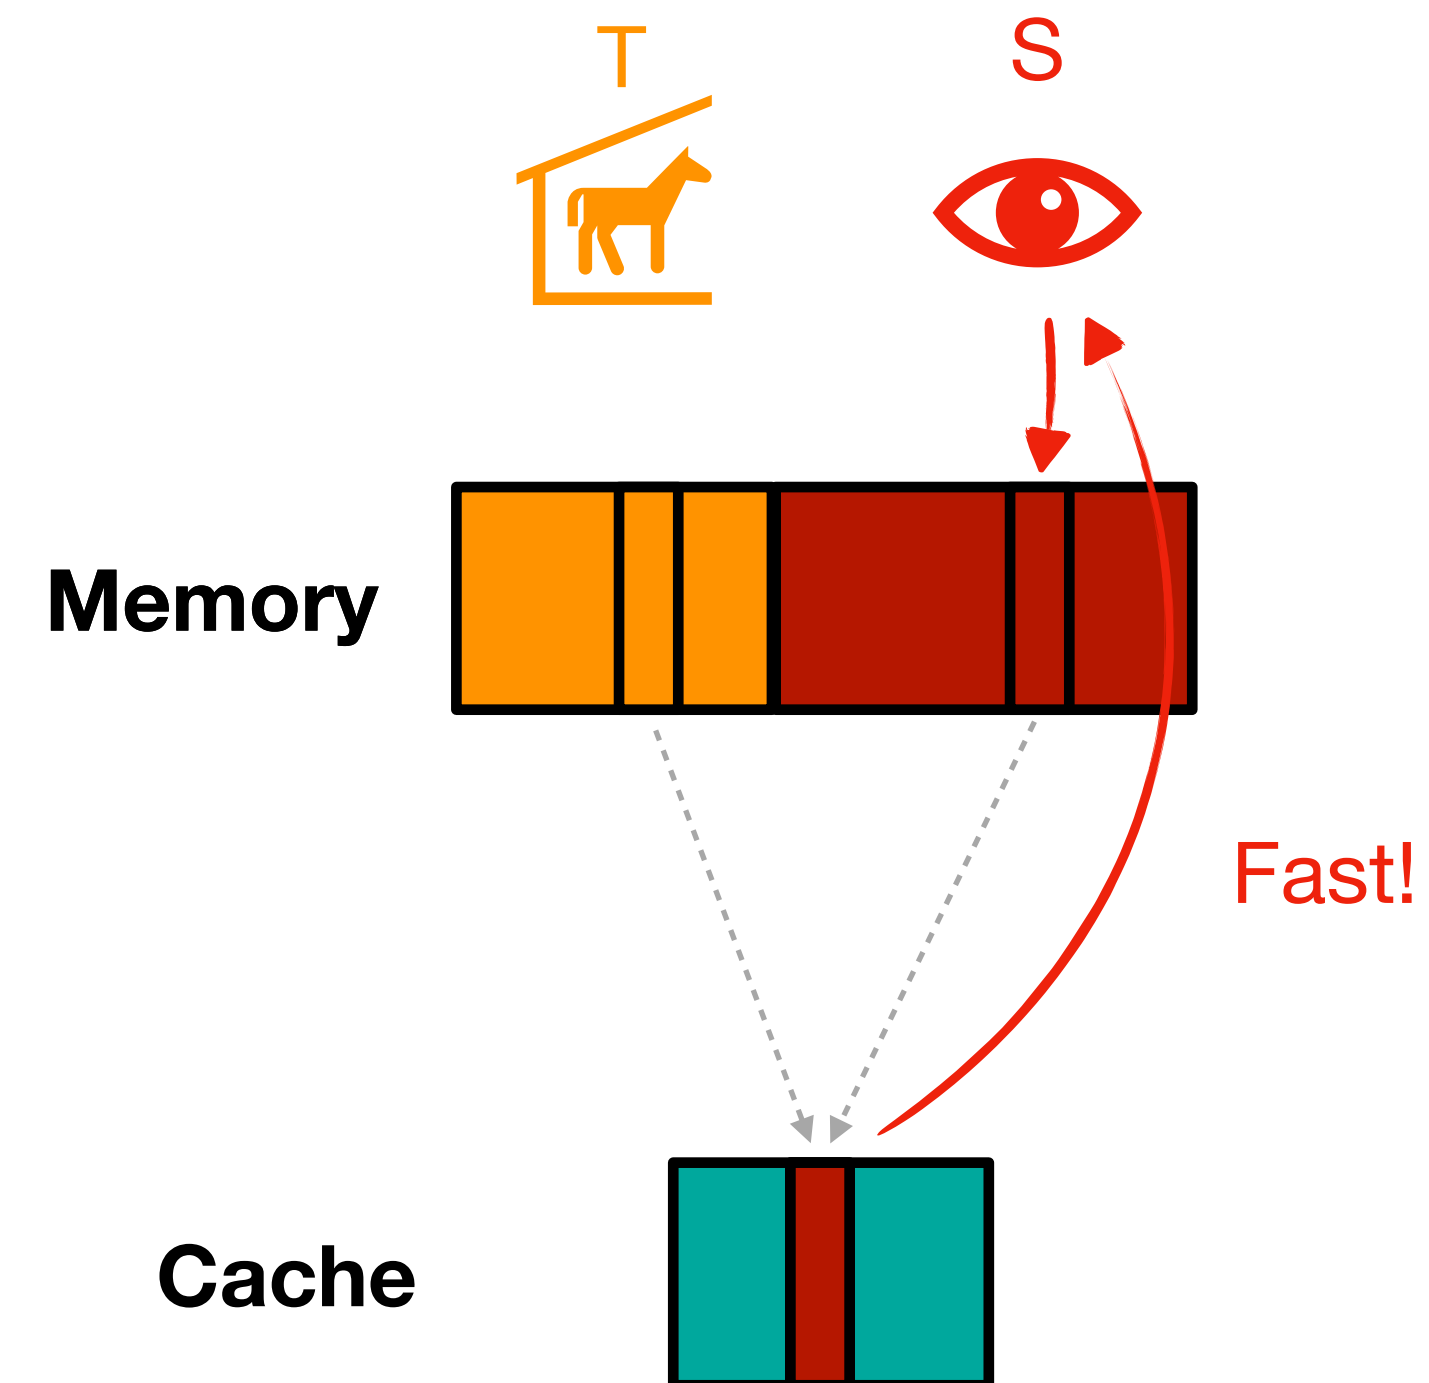
- OSes typically implement *memory protection*.
- But: Mere memory access changes microarchitectural state — this affects timing.



What is Time Protection?



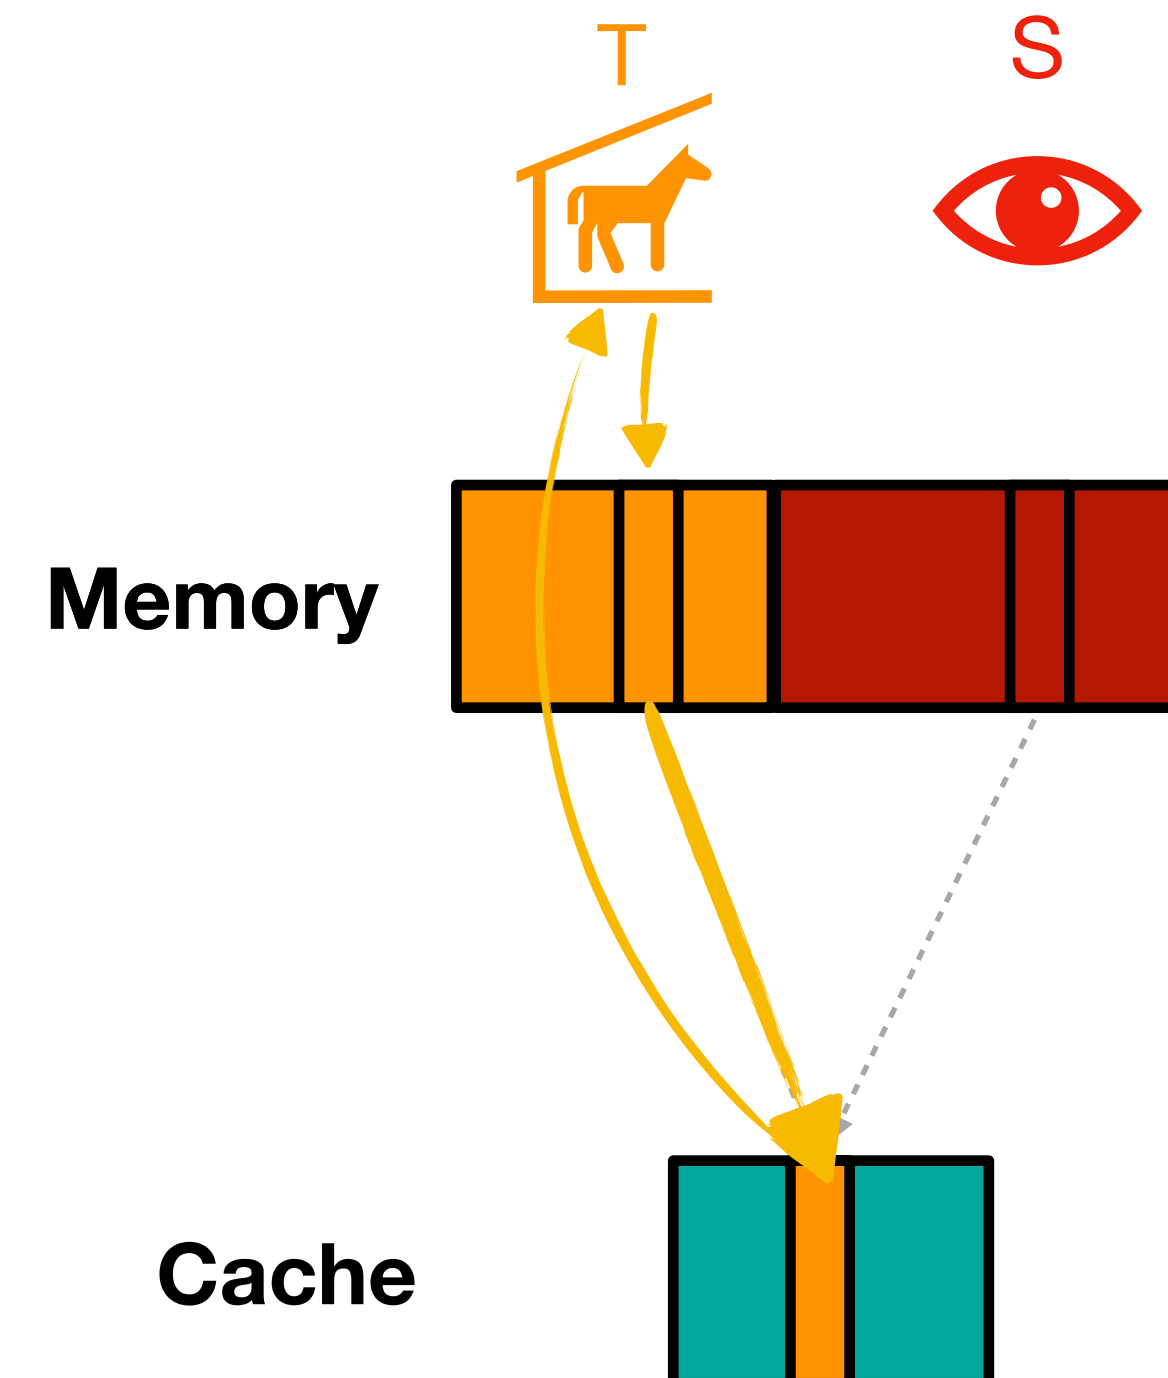
- OSes typically implement *memory protection*.
- But: Mere memory access changes microarchitectural state — this affects timing.



What is Time Protection?



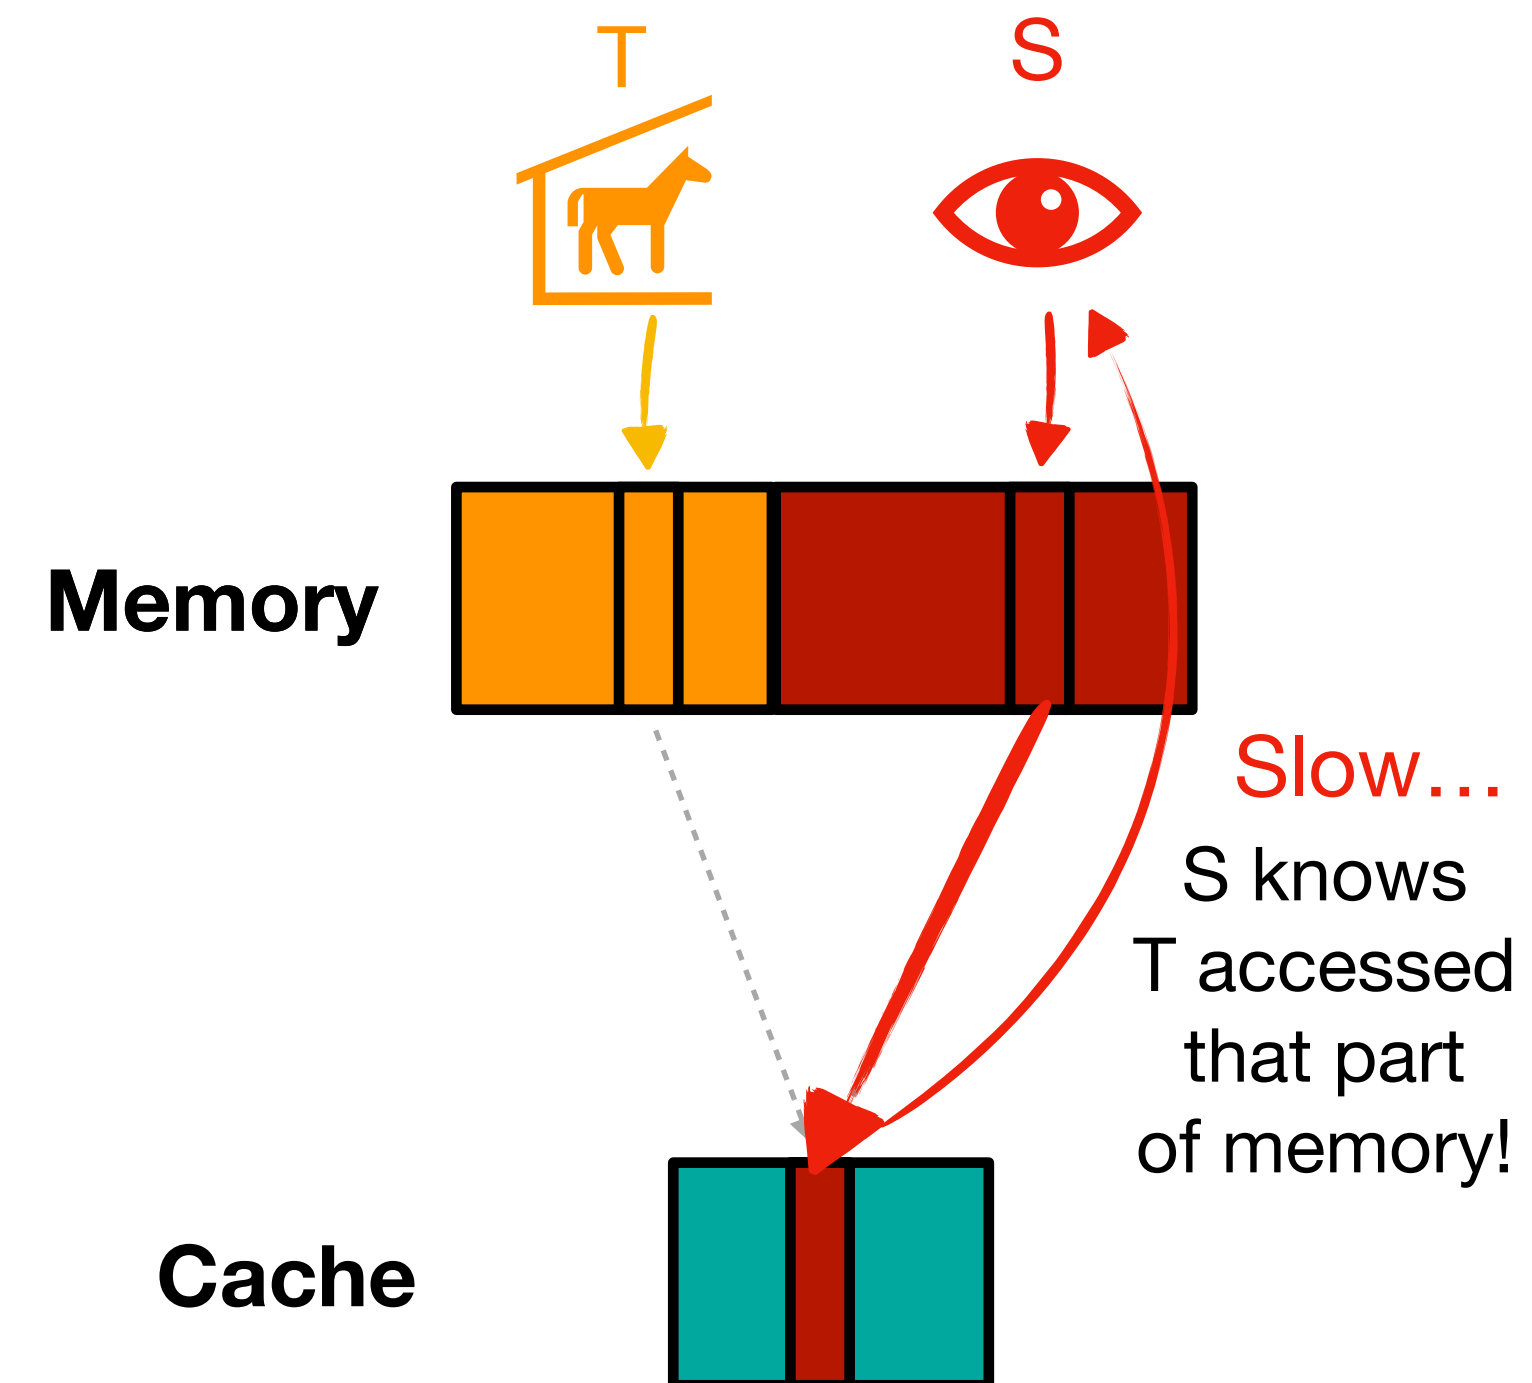
- OSes typically implement *memory protection*.
- But: Mere memory access changes microarchitectural state — this affects timing.



What is Time Protection?



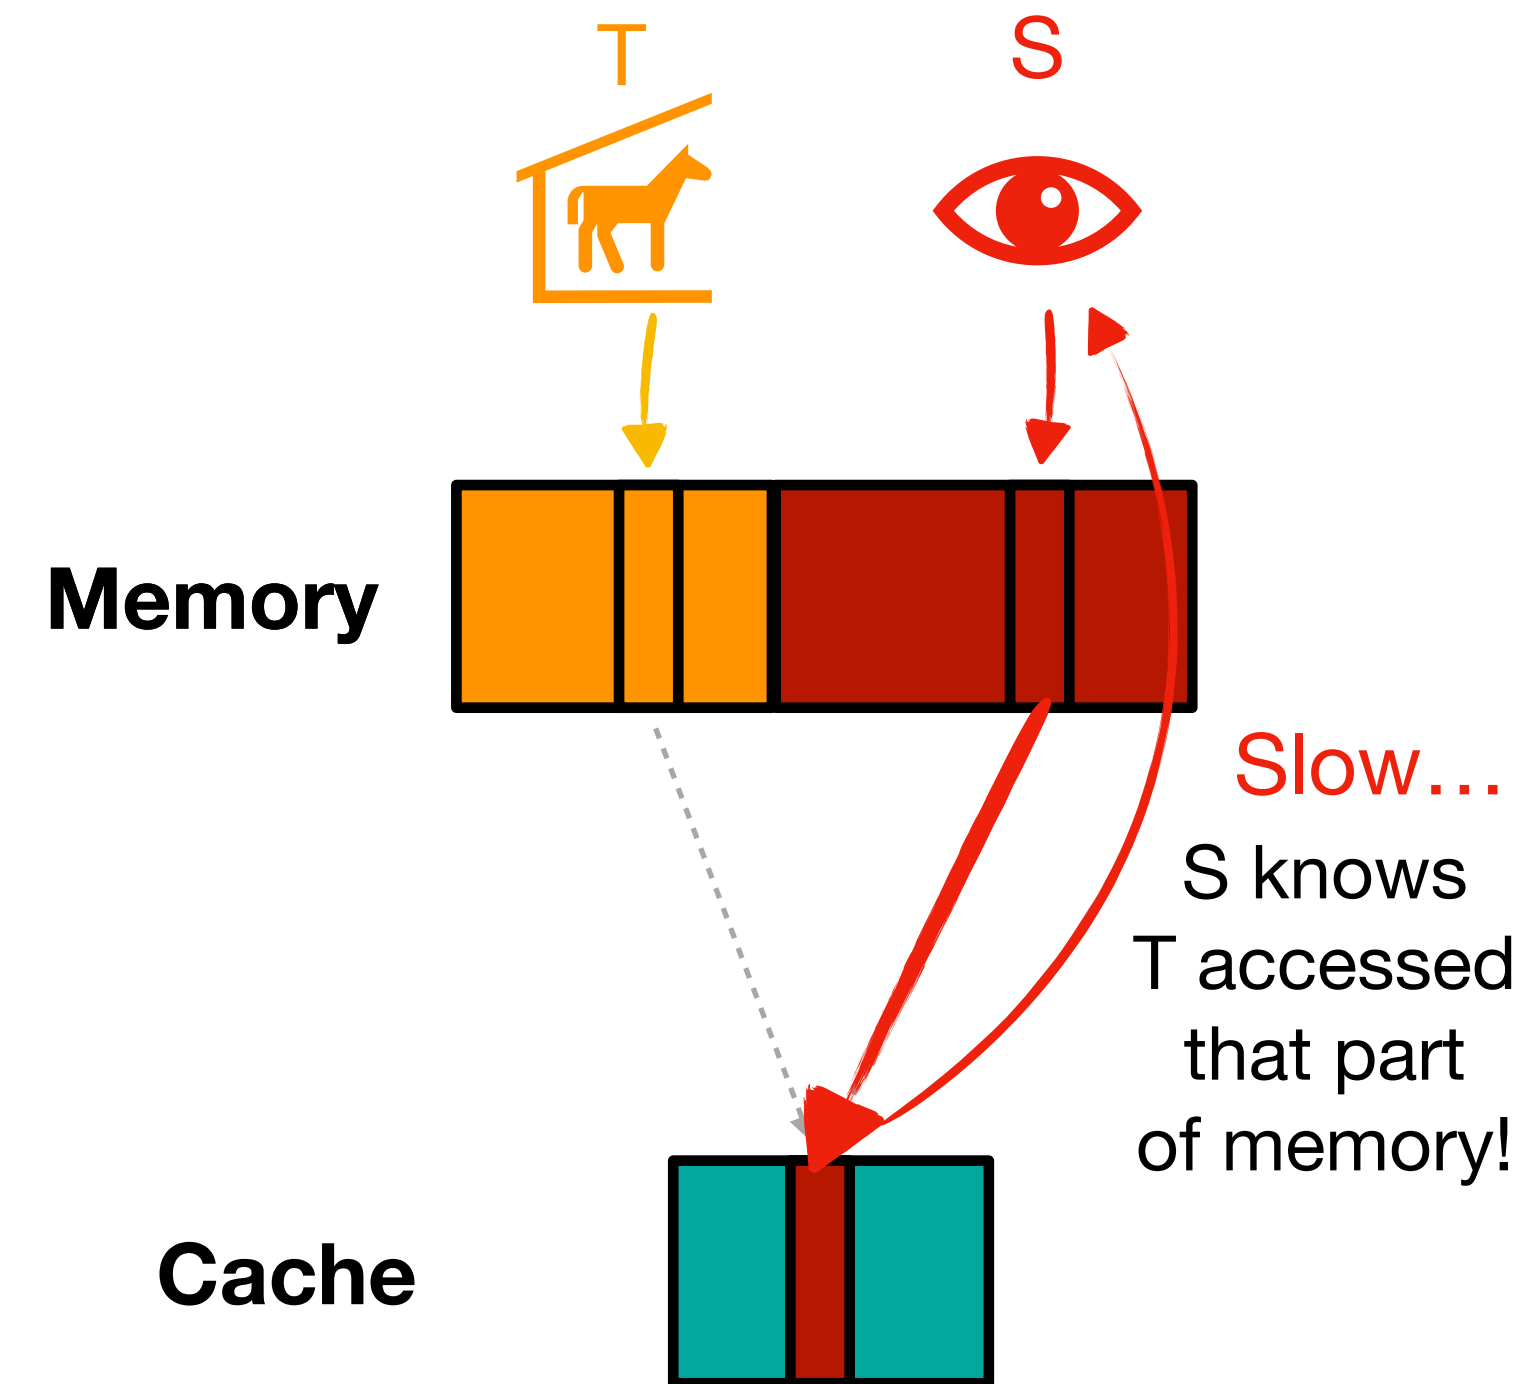
- OSes typically implement *memory protection*.
- But: Mere memory access changes microarchitectural state — this affects timing.



What is Time Protection?



- OSes typically implement *memory protection*.
- But: Mere memory access changes microarchitectural state — this affects timing.
- To prevent these *timing channels*, OSes can implement *time protection*:
See EuroSys: [Ge et al. 2019]



Qian Ge



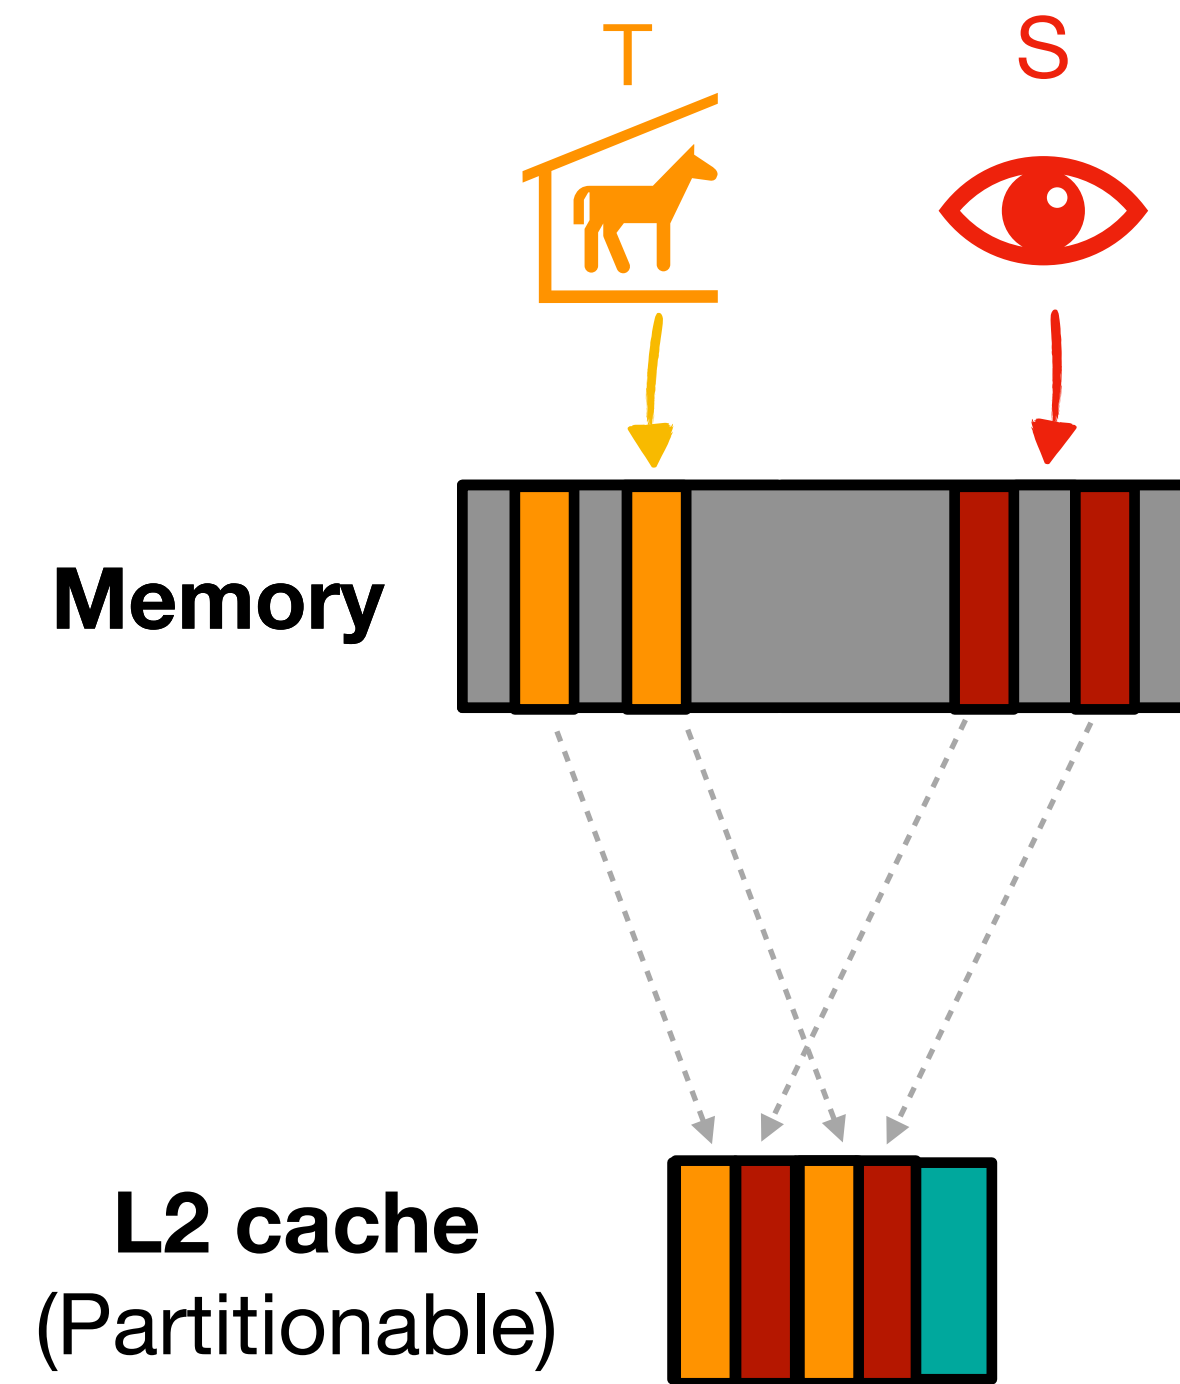
Gernot Heiser

with
Yuval Yarom (U. Adelaide)
Tom Chothia (U. Birmingham)

What is Time Protection?



- OSes typically implement *memory protection*.
- But: Mere memory access changes microarchitectural state — this affects timing.
- To prevent these *timing channels*, OSes can implement *time protection*:
See EuroSys: [Ge et al. 2019]
- *Partition* off-core memory caches



Qian Ge



Gernot Heiser

with
Yuval Yarom (U. Adelaide)
Tom Chothia (U. Birmingham)

What is Time Protection?



- OSes typically implement *memory protection*.
- But: Mere memory access changes microarchitectural state — this affects timing.
- To prevent these *timing channels*, OSes can implement *time protection*:
See EuroSys: [Ge et al. 2019]
- *Partition* off-core memory caches
- *Flush* on-core and non-architected state and *pad* time on context switch

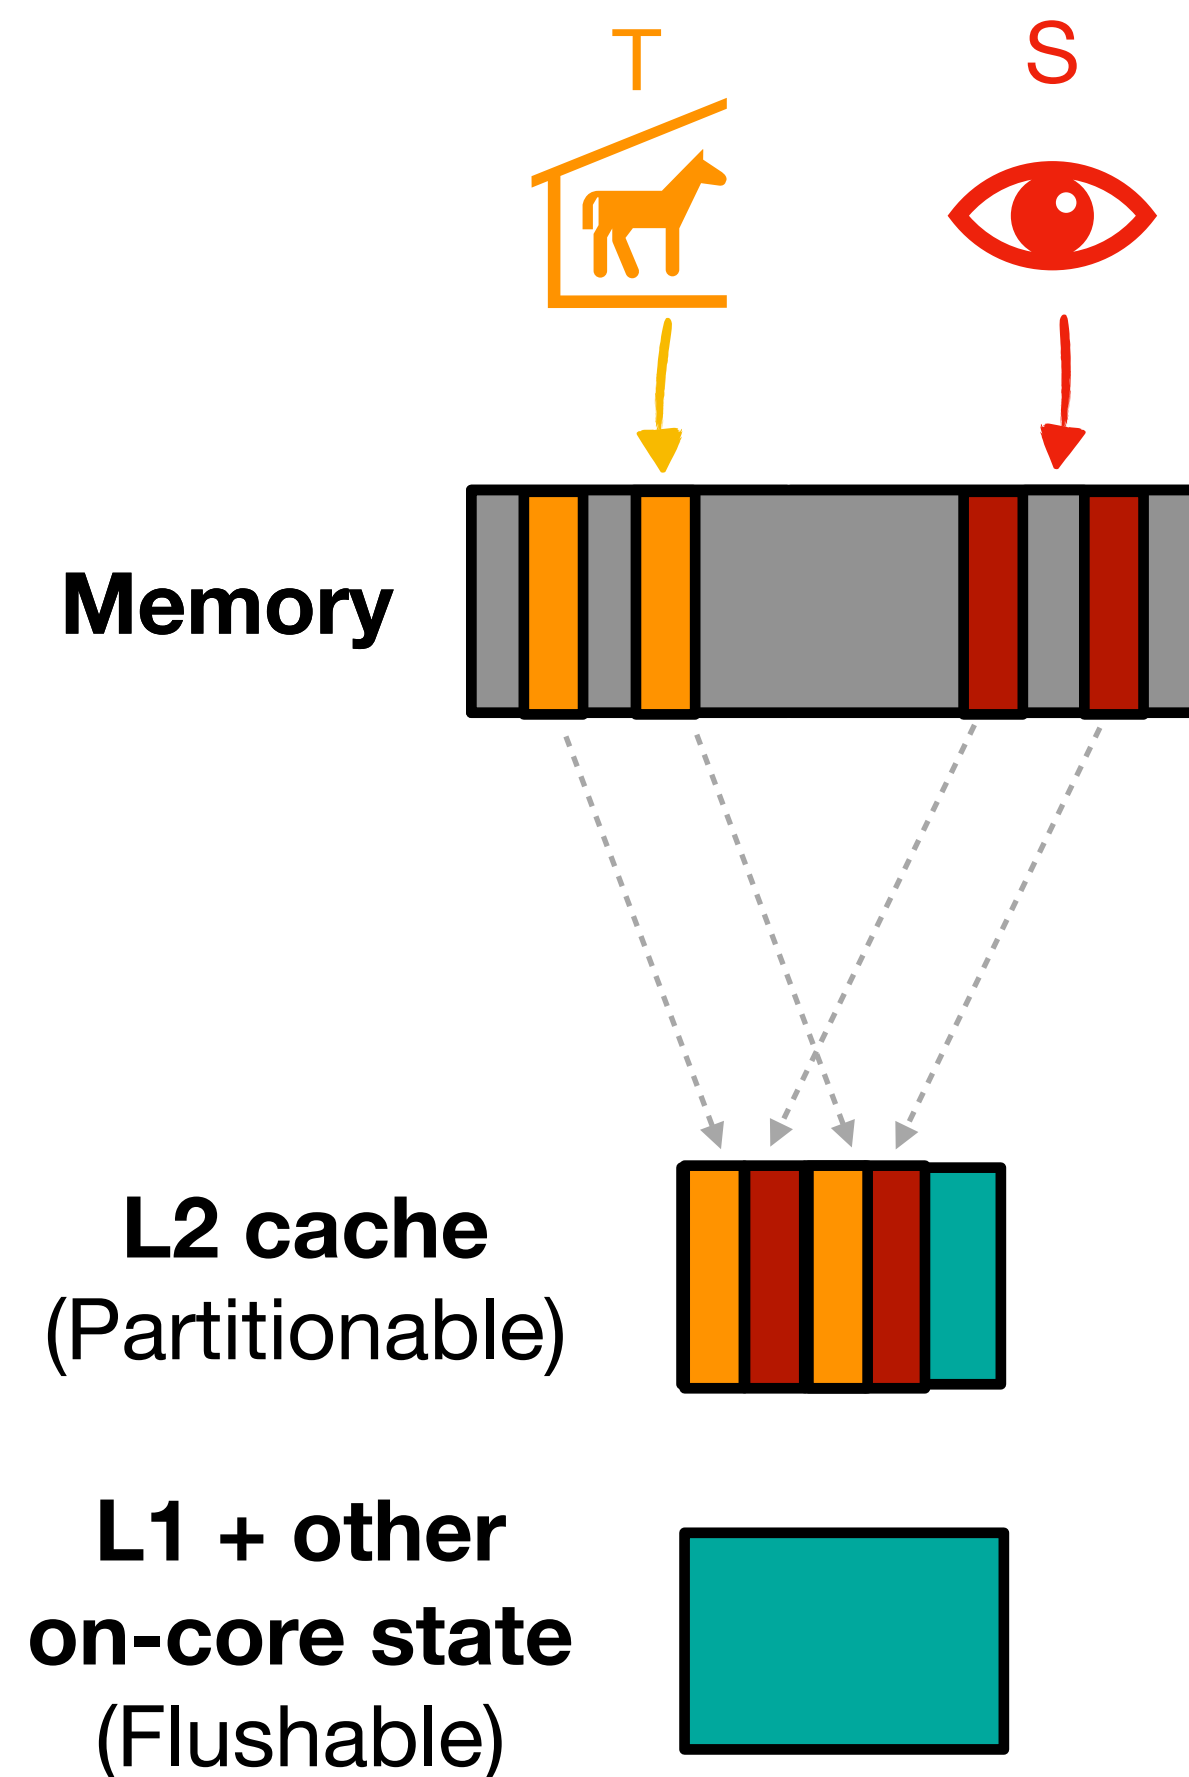


Qian Ge



Gernot Heiser

with
Yuval Yarom (U. Adelaide)
Tom Chothia (U. Birmingham)



What is Time Protection?



- OSes typically implement *memory protection*.
- But: Mere memory access changes microarchitectural state — this affects timing.
- To prevent these *timing channels*, OSes can implement *time protection*:
See EuroSys: [Ge et al. 2019]
- *Partition* off-core memory caches
- *Flush* on-core and non-architected state and *pad* time on context switch

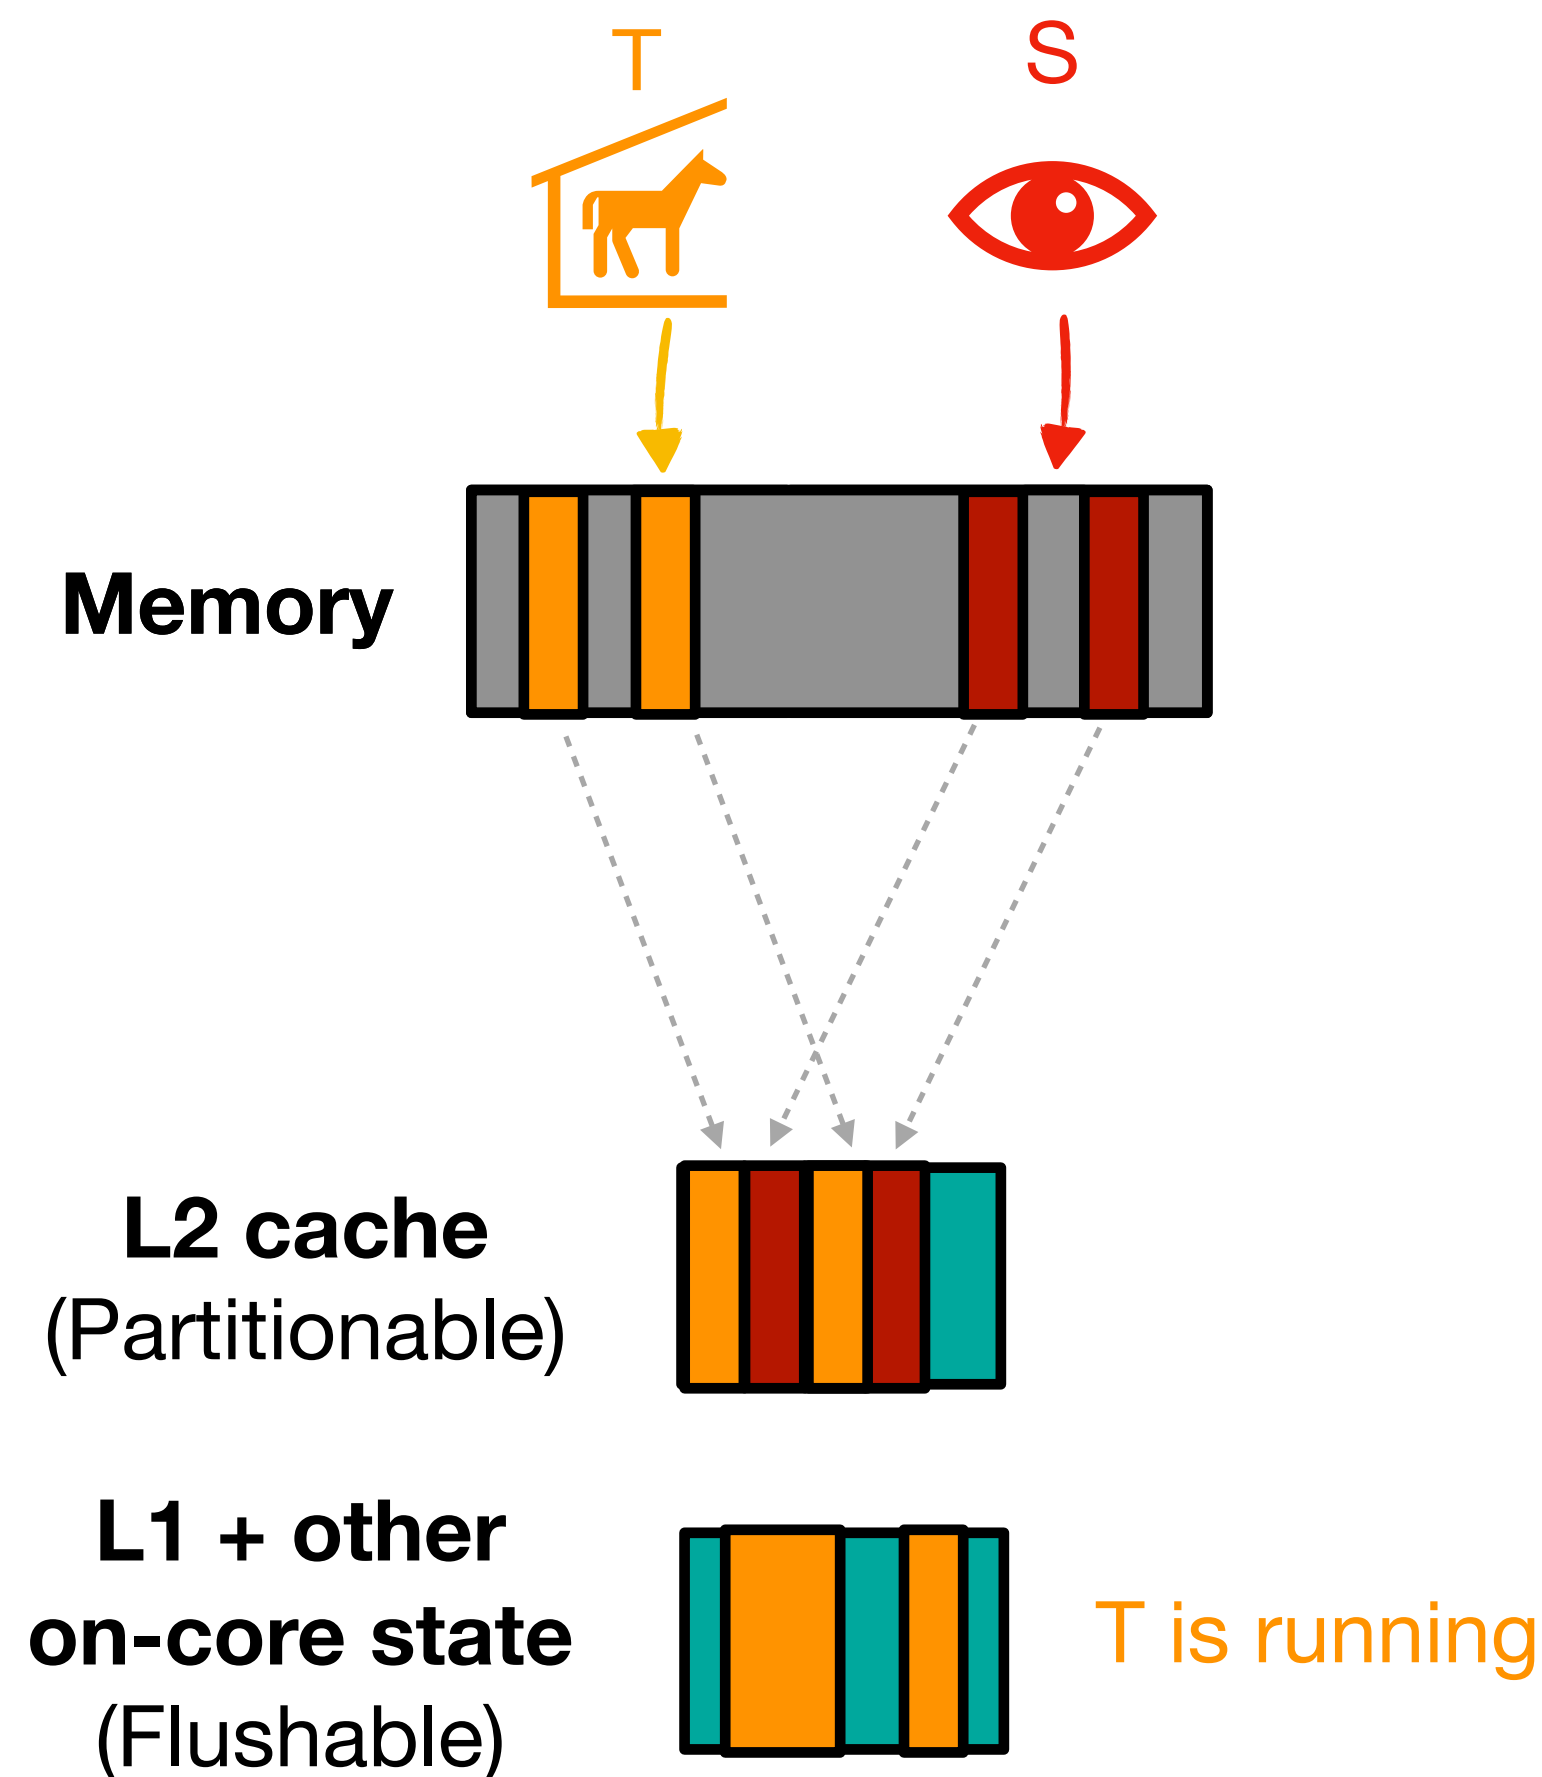


Qian Ge



Gernot Heiser

with
Yuval Yarom (U. Adelaide)
Tom Chothia (U. Birmingham)



What is Time Protection?



- OSes typically implement *memory protection*.
- But: Mere memory access changes microarchitectural state — this affects timing.
- To prevent these *timing channels*, OSes can implement *time protection*:
See EuroSys: [Ge et al. 2019]
- *Partition* off-core memory caches
- *Flush* on-core and non-architected state and *pad* time on context switch

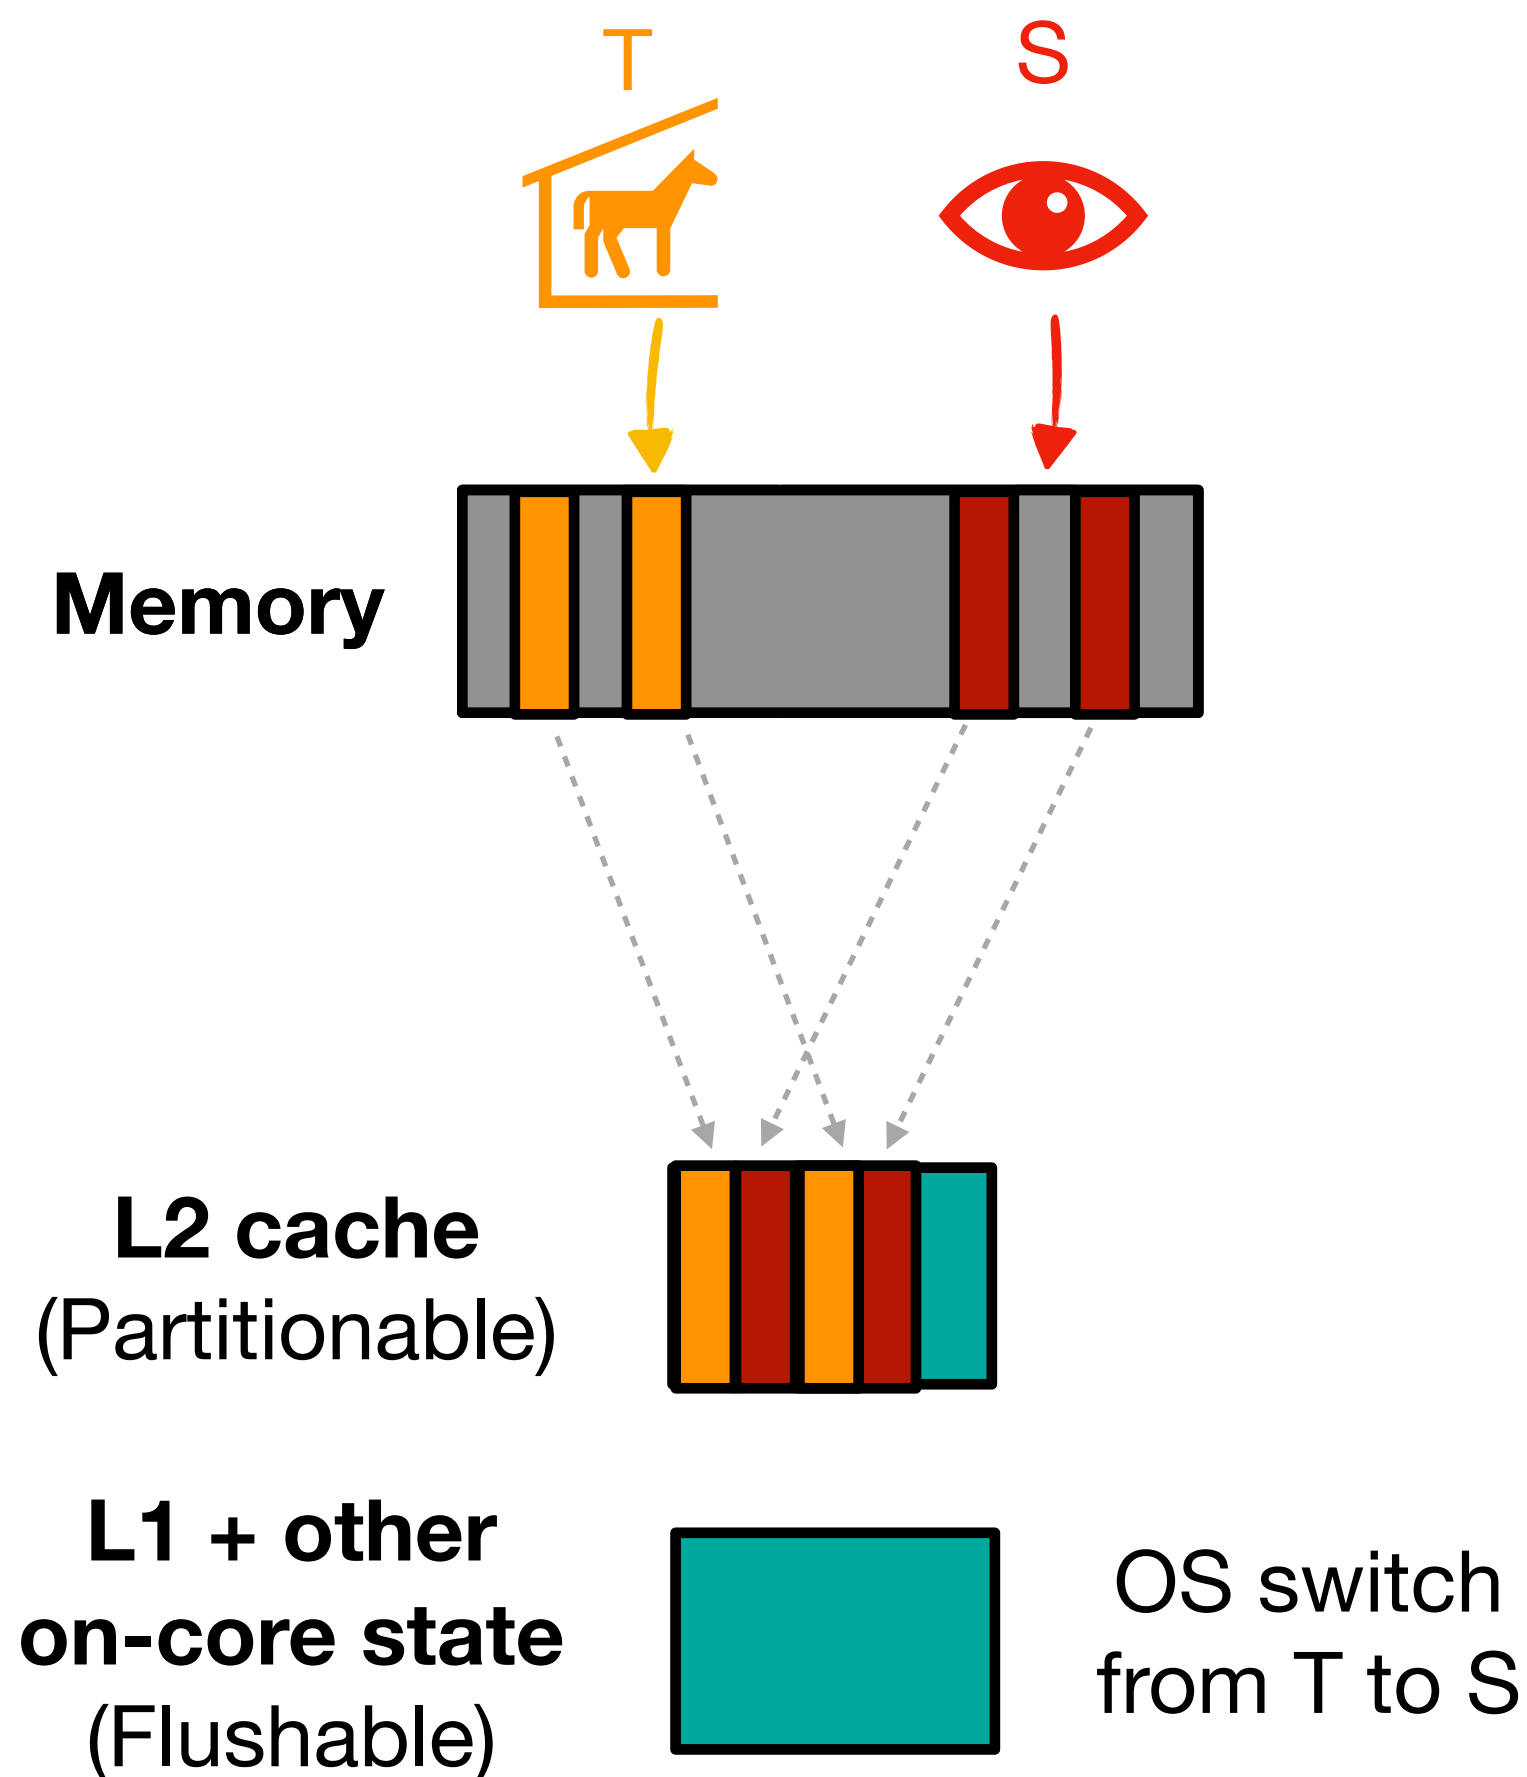


Qian Ge



Gernot Heiser

with
Yuval Yarom (U. Adelaide)
Tom Chothia (U. Birmingham)



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

What is Time Protection?



- OSes typically implement *memory protection*.
- But: Mere memory access changes microarchitectural state — this affects timing.
- To prevent these *timing channels*, OSes can implement *time protection*:
See EuroSys: [Ge et al. 2019]
- *Partition* off-core memory caches
- *Flush* on-core and non-architected state and *pad* time on context switch

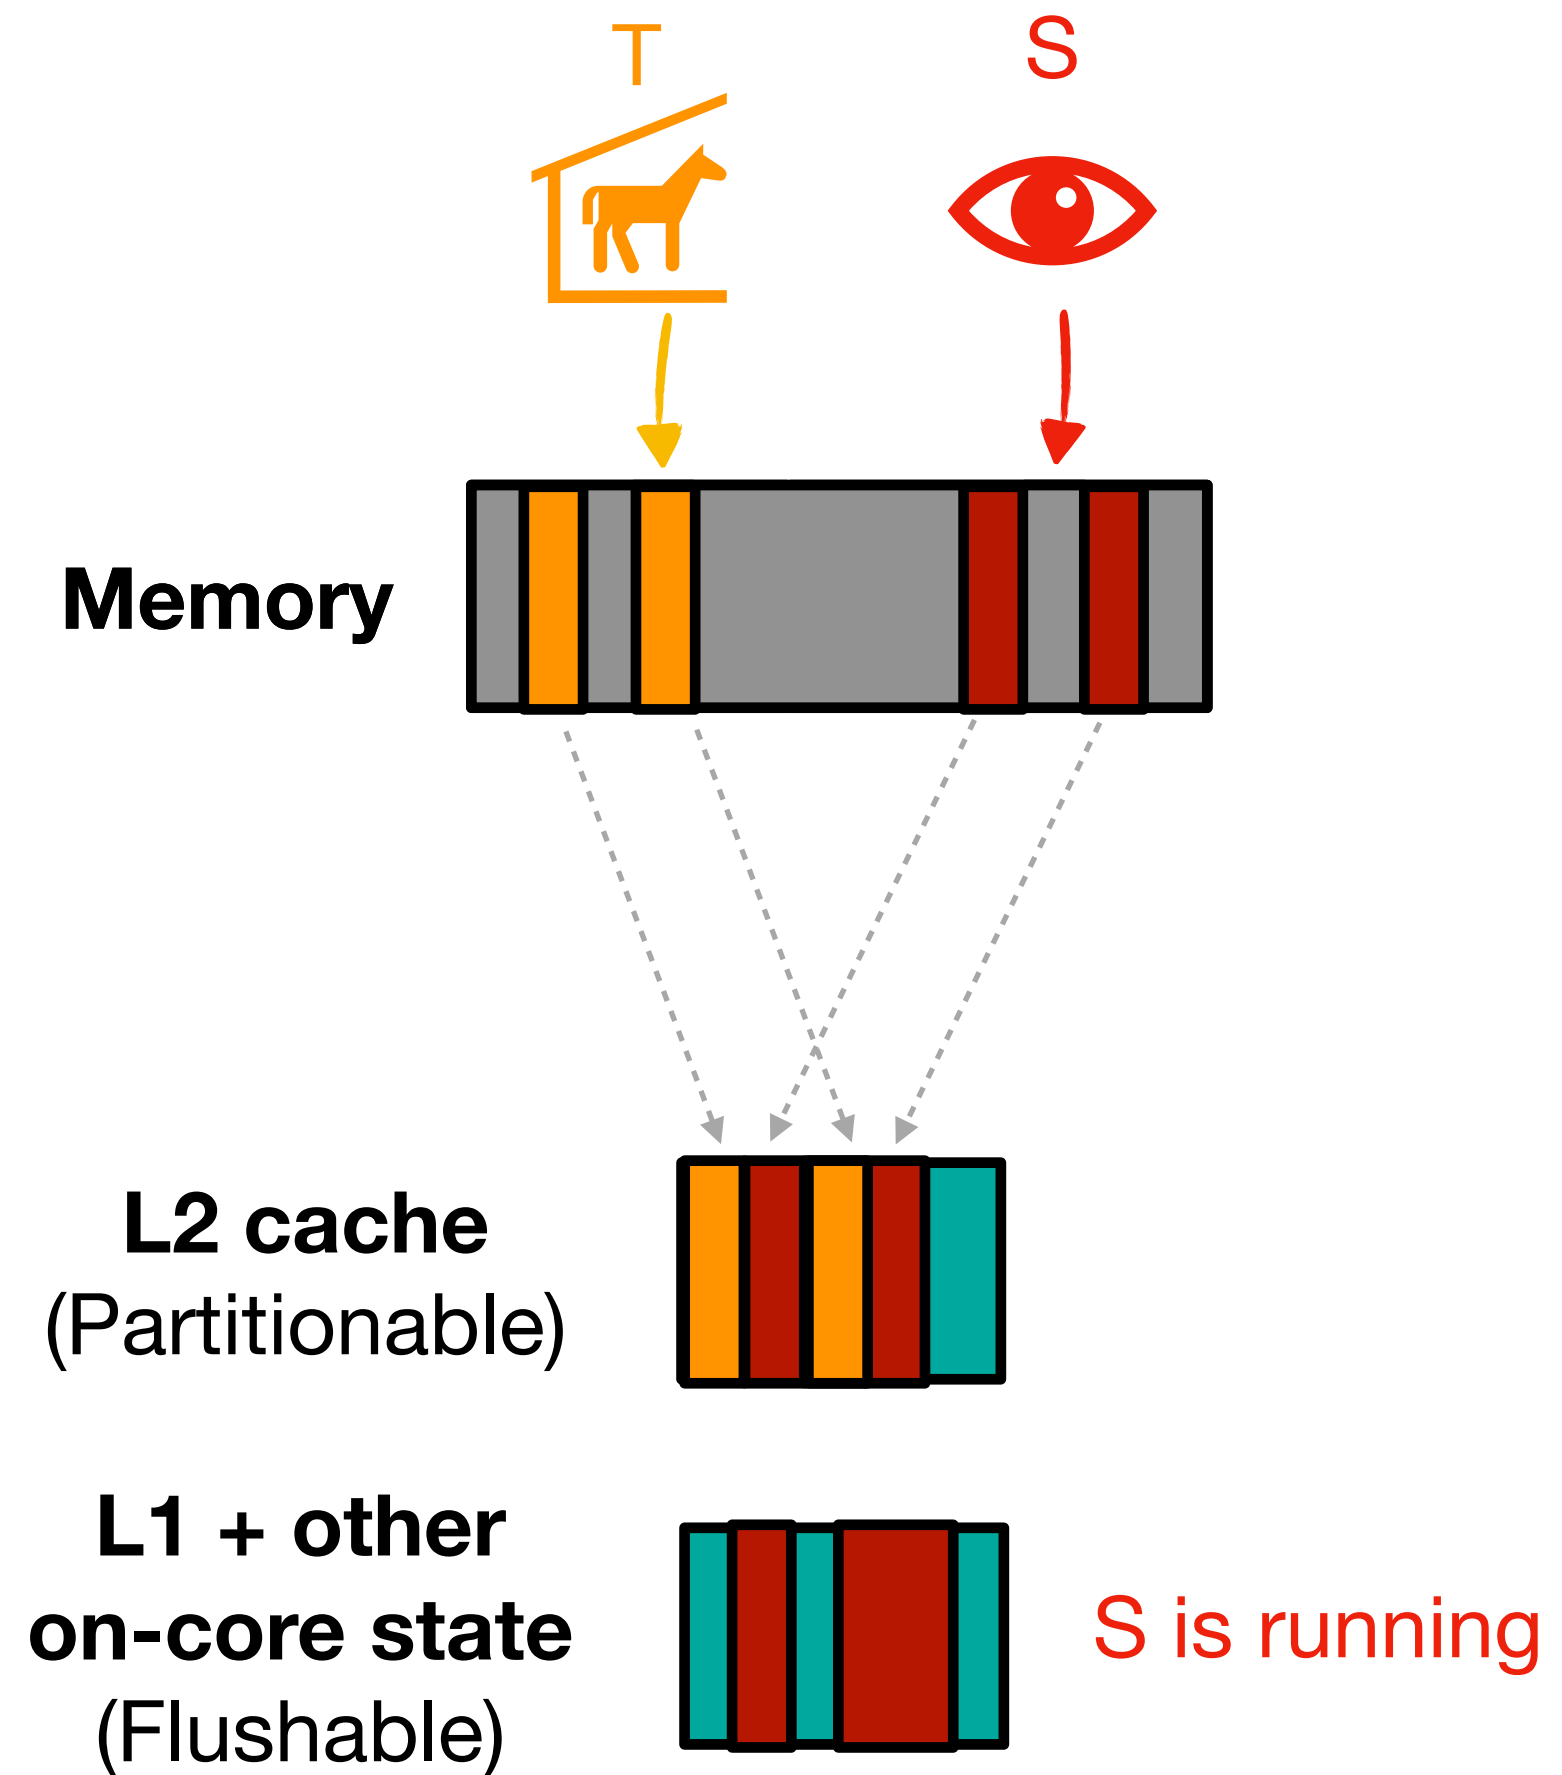


Qian Ge



Gernot Heiser

with
Yuval Yarom (U. Adelaide)
Tom Chothia (U. Birmingham)



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

What is Time Protection?



- OSes typically implement *memory protection*.
- But: Mere memory access changes microarchitectural state — this affects timing.
- To prevent these *timing channels*, OSes can implement *time protection*:
See EuroSys: [Ge et al. 2019]
- *Partition* off-core memory caches
- *Flush* on-core and non-architected state and *pad* time on context switch

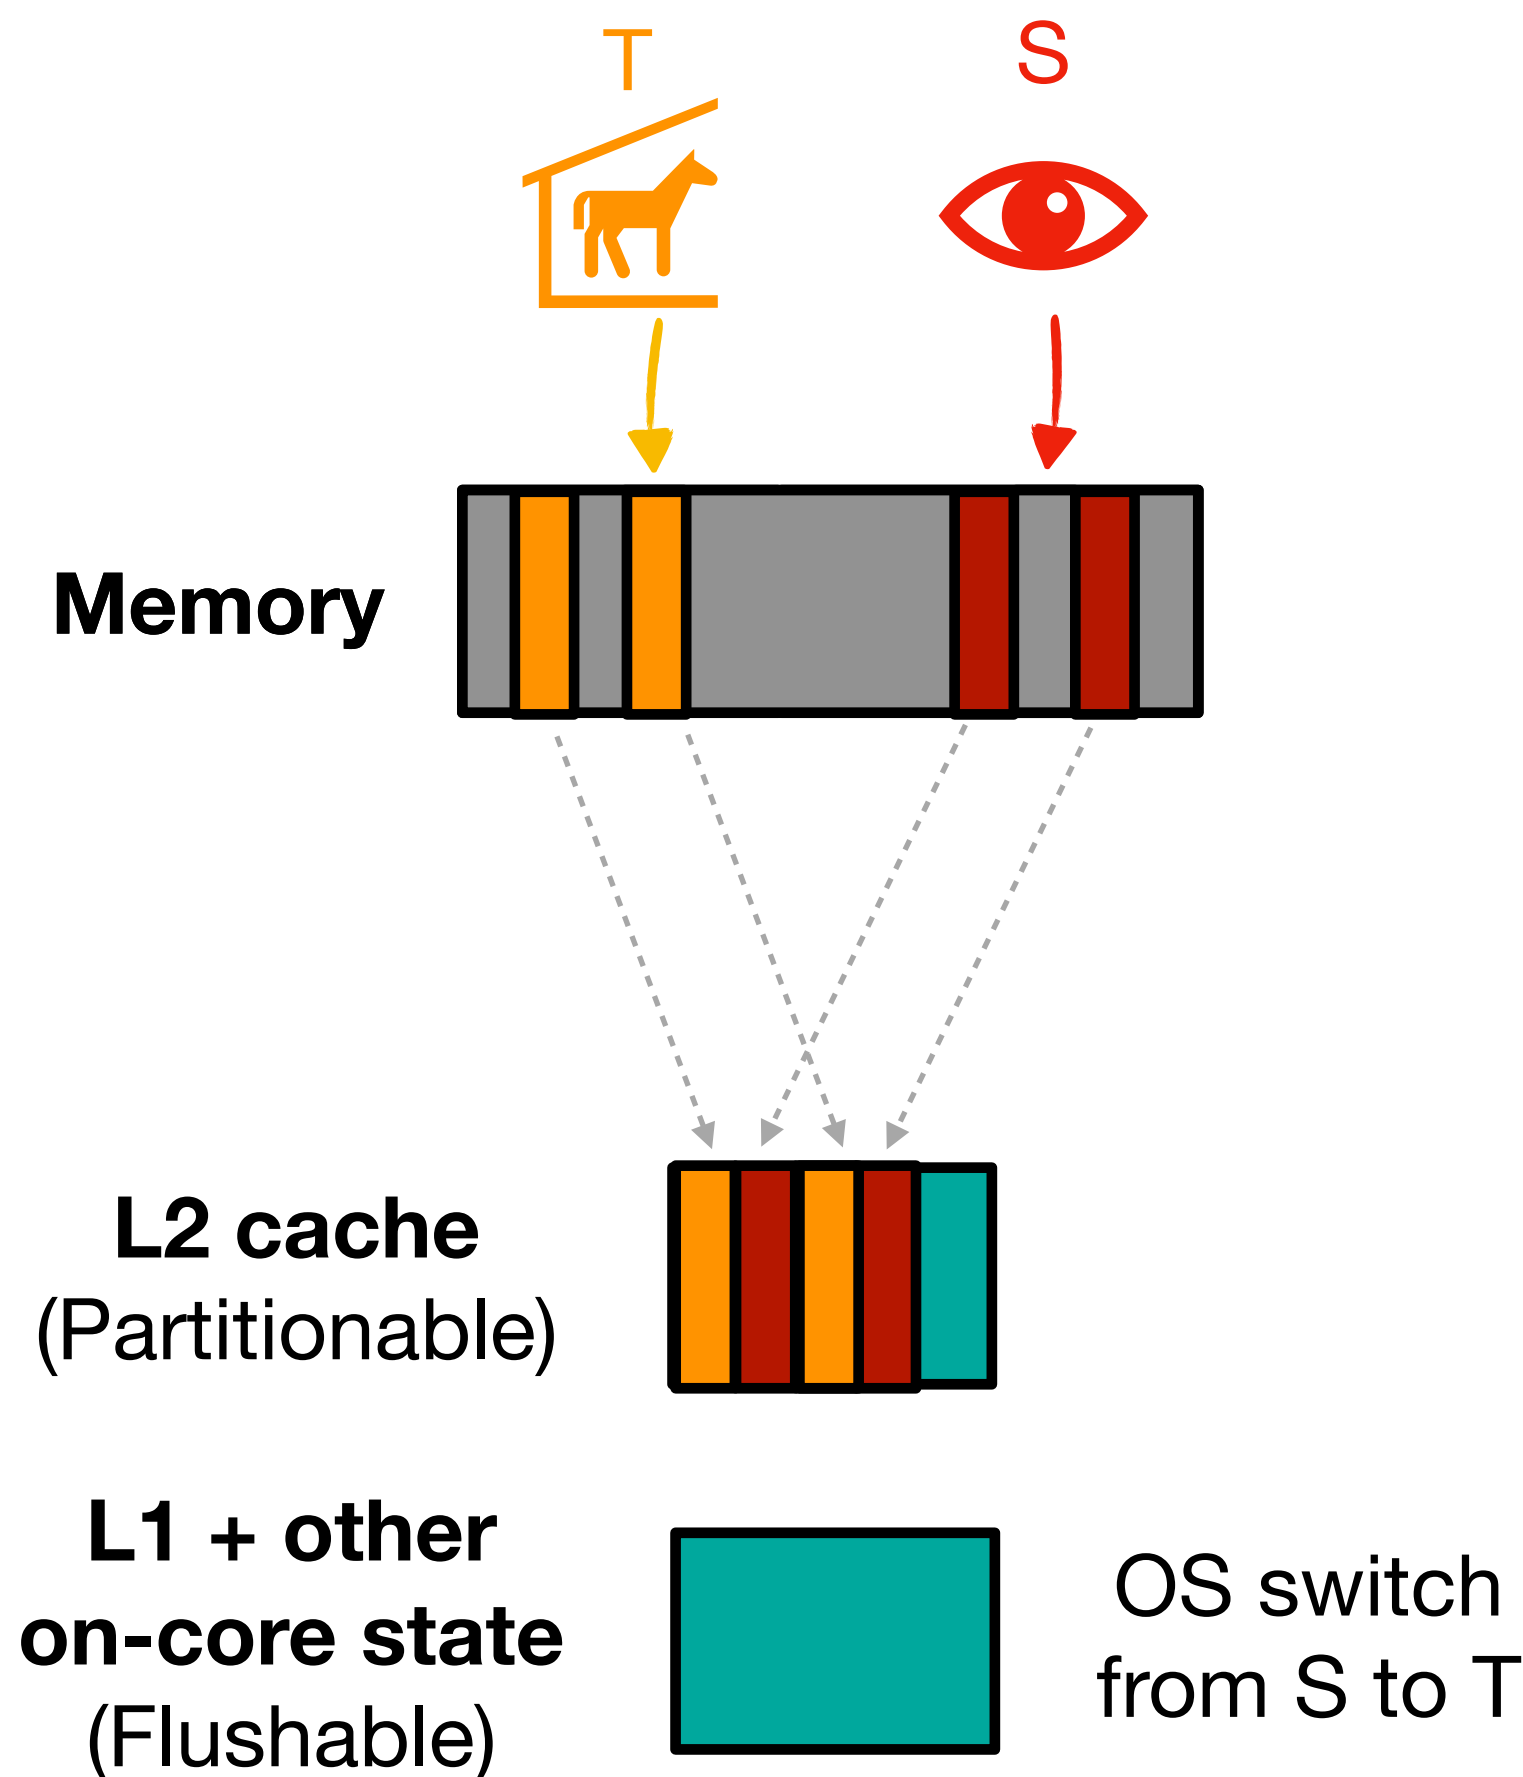


Qian Ge



Gernot Heiser

with
Yuval Yarom (U. Adelaide)
Tom Chothia (U. Birmingham)



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

seL4 The seL4 kernel and Time Protection



- To prevent these *timing channels*, OSes can implement *time protection*:
See EuroSys: [Ge et al. 2019]
- *Partition* off-core memory caches
- *Flush* on-core and non-architected state and *pad* time on context switch
- For the seL4 OS kernel:

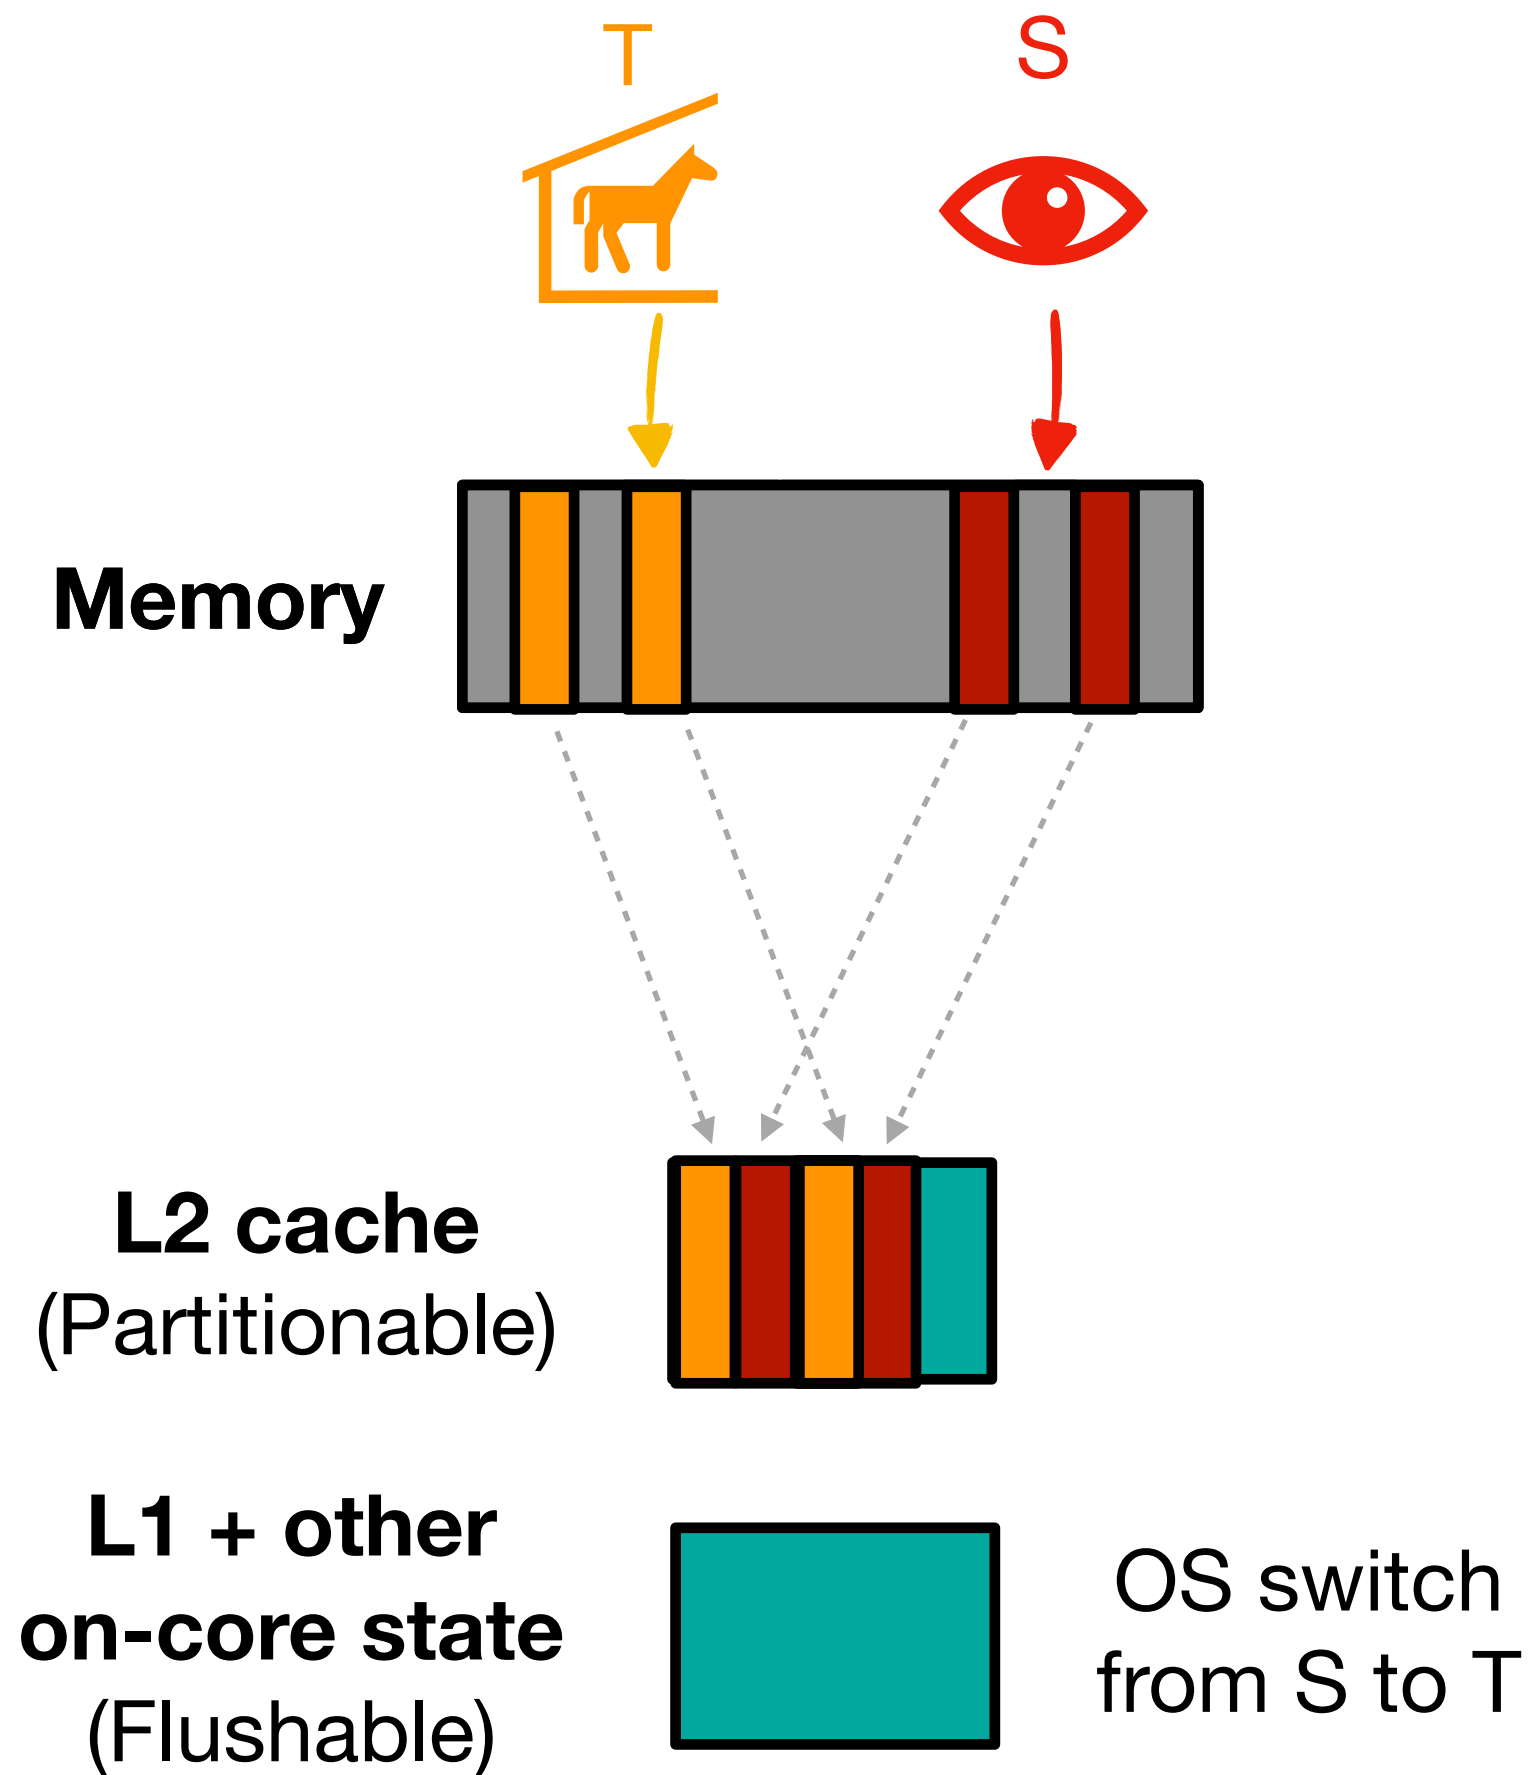


Qian Ge



Gernot Heiser

with
Yuval Yarom (U. Adelaide)
Tom Chothia (U. Birmingham)



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

seL4 The seL4 kernel and Time Protection



- To prevent these *timing channels*, OSes can implement *time protection*:
See EuroSys: [Ge et al. 2019]
- *Partition* off-core memory caches
- *Flush* on-core and non-architected state and *pad* time on context switch

- For the seL4 OS kernel:

seL4: World's first OS kernel with correctness proof!

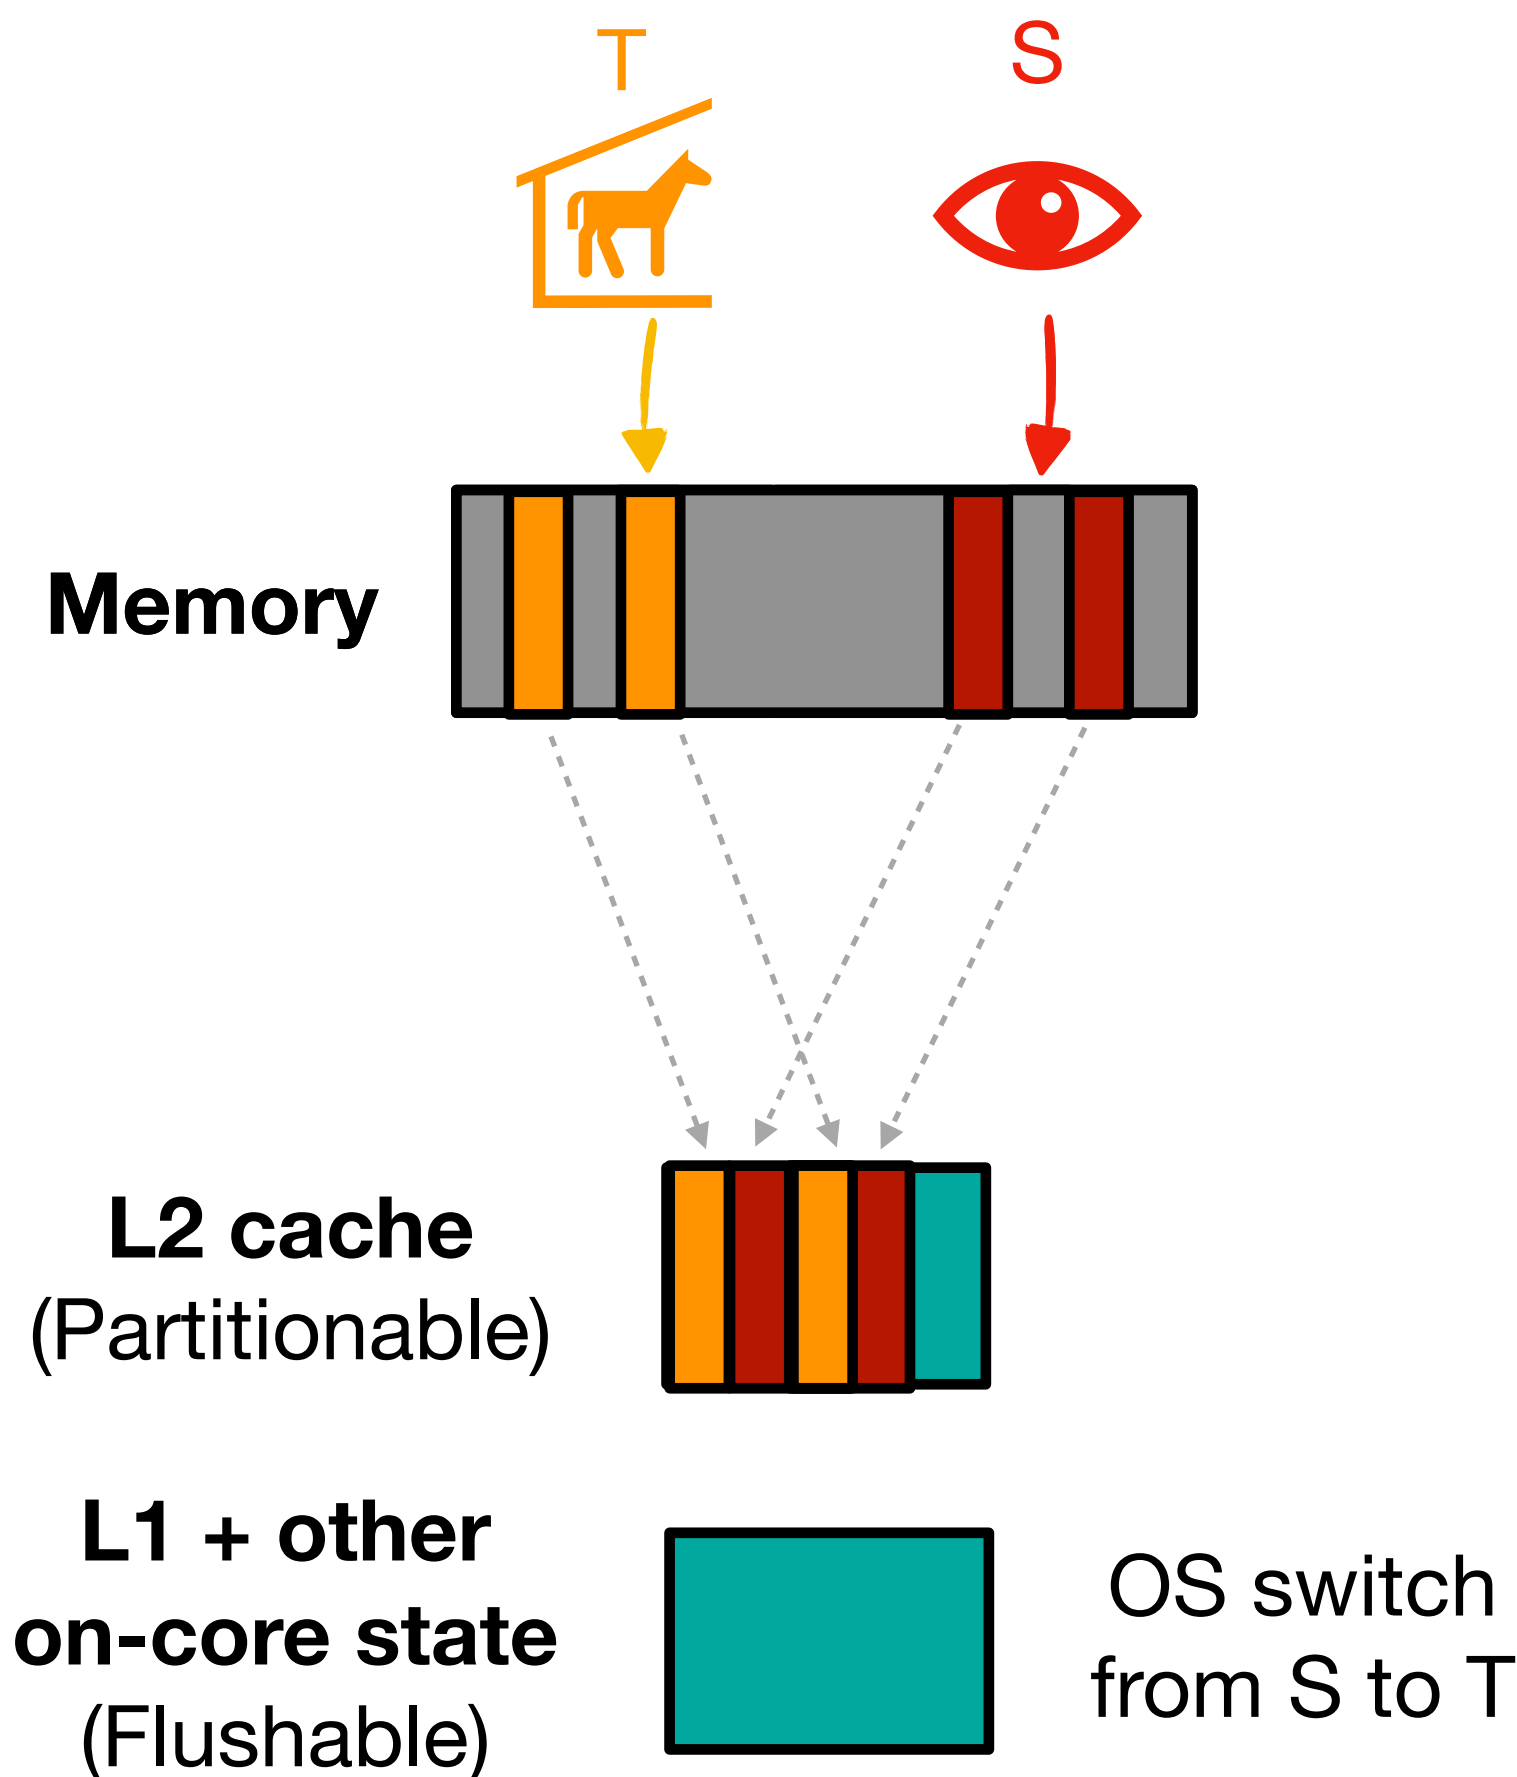


Qian Ge



Gernot Heiser

with
Yuval Yarom (U. Adelaide)
Tom Chothia (U. Birmingham)



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

seL4 The seL4 kernel and Time Protection



- To prevent these *timing channels*, OSes can implement *time protection*:
See EuroSys: [Ge et al. 2019]
- Partition* off-core memory caches
- Flush* on-core and non-architected state and *pad* time on context switch

- For the seL4 OS kernel:
 - Implemented, evaluated empirically on ARM, x86
See EuroSys: [Ge et al. 2019]

seL4: World's first OS kernel with correctness proof!

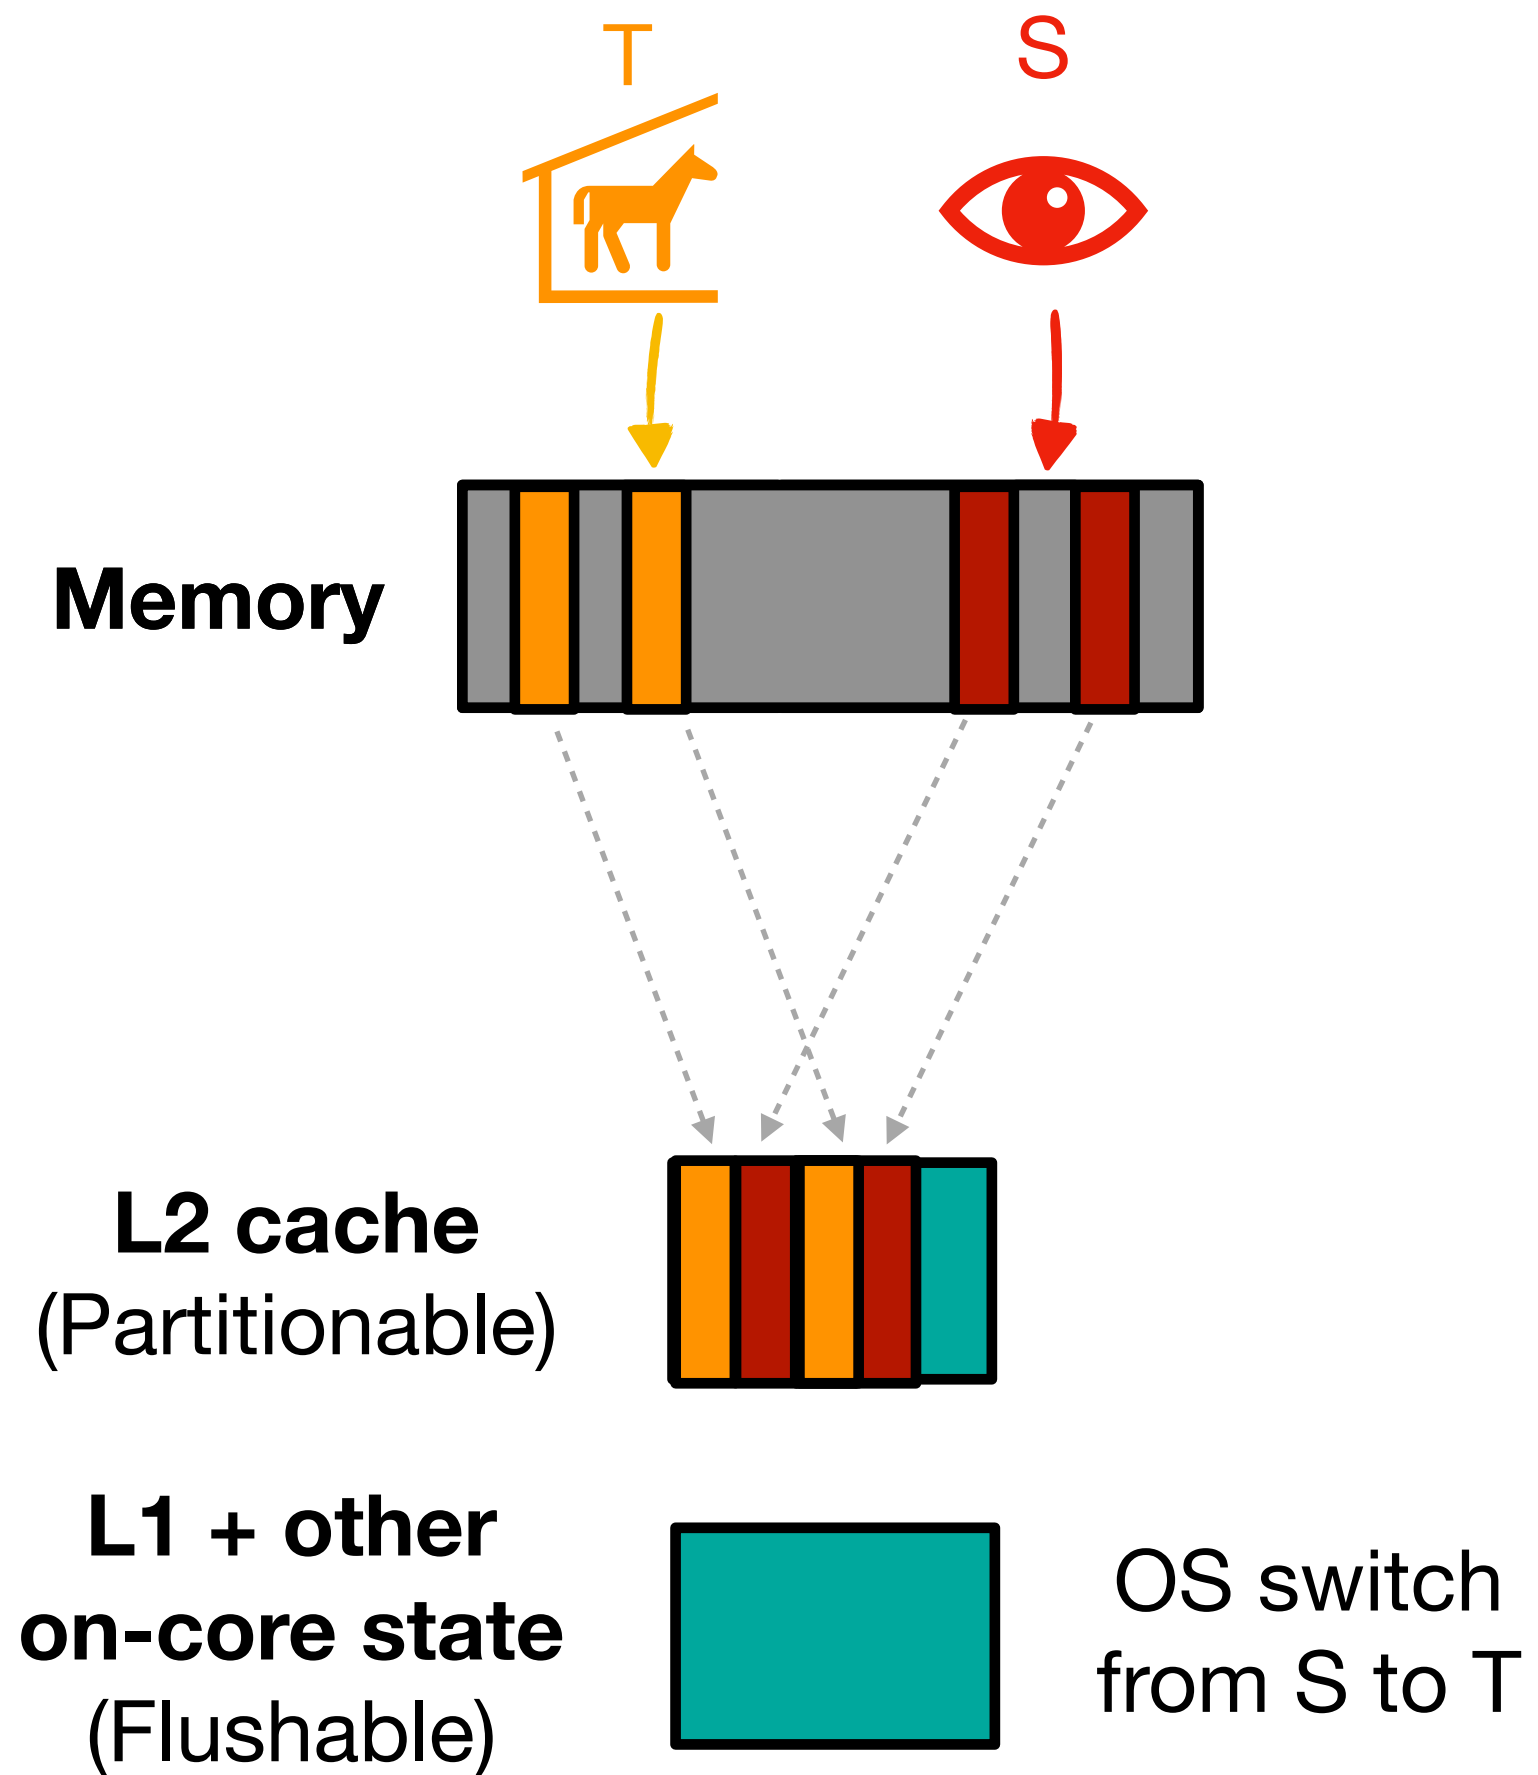


Qian Ge



Gernot Heiser

with
Yuval Yarom (U. Adelaide)
Tom Chothia (U. Birmingham)



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

seL4 The seL4 kernel and Time Protection



- To prevent these *timing channels*, OSes can implement *time protection*:
See EuroSys: [Ge et al. 2019]
- *Partition* off-core memory caches
- *Flush* on-core and non-architected state and *pad* time on context switch
- For the seL4 OS kernel:
 - Implemented, evaluated empirically on ARM, x86
See EuroSys: [Ge et al. 2019]
 - Ported to RISC-V with hardware support
See arXiv preprint: [Buckley, Sison et al. 2023]
(HW support by [Wistoff et al. 2023])

seL4: World's first OS kernel with correctness proof!



Scott Buckley



Rob Sison



Nils Wistoff



Curtis Millar



Toby Murray



Gerwin Klein



Gernot Heiser

Thanks to funding from Australian Research Council

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

seL4 The seL4 kernel and Time Protection



- To prevent these *timing channels*, OSes can implement *time protection*:
See EuroSys: [Ge et al. 2019]
- *Partition* off-core memory caches
- *Flush* on-core and non-architected state and *pad* time on context switch
- For the seL4 OS kernel:
 - Implemented, evaluated empirically on ARM, x86
See EuroSys: [Ge et al. 2019]
 - Ported to RISC-V with hardware support
See arXiv preprint: [Buckley, Sison et al. 2023]
(HW support by [Wistoff et al. 2023])
- **This talk:** How do we verify it works?

seL4: World's first OS kernel with correctness proof!



Scott Buckley



Rob Sison



Nils Wistoff



Curtis Millar



Toby Murray



Gerwin Klein



Gernot Heiser

Thanks to funding from
Australian Research Council

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



How is seL4 verified?



- **This talk:** How do we verify Time Protection on seL4 works?

- Implemented, evaluated empirically on ARM, x86
See EuroSys: [Ge et al. 2019]

- Ported to RISC-V with hardware support
See arXiv preprint: [Buckley, Sison et al. 2023]
(HW support by [Wistoff et al. 2023])



Thanks to funding from
Australian Research Council

seL4: World's first
OS kernel with
correctness proof!



How is seL4 verified?



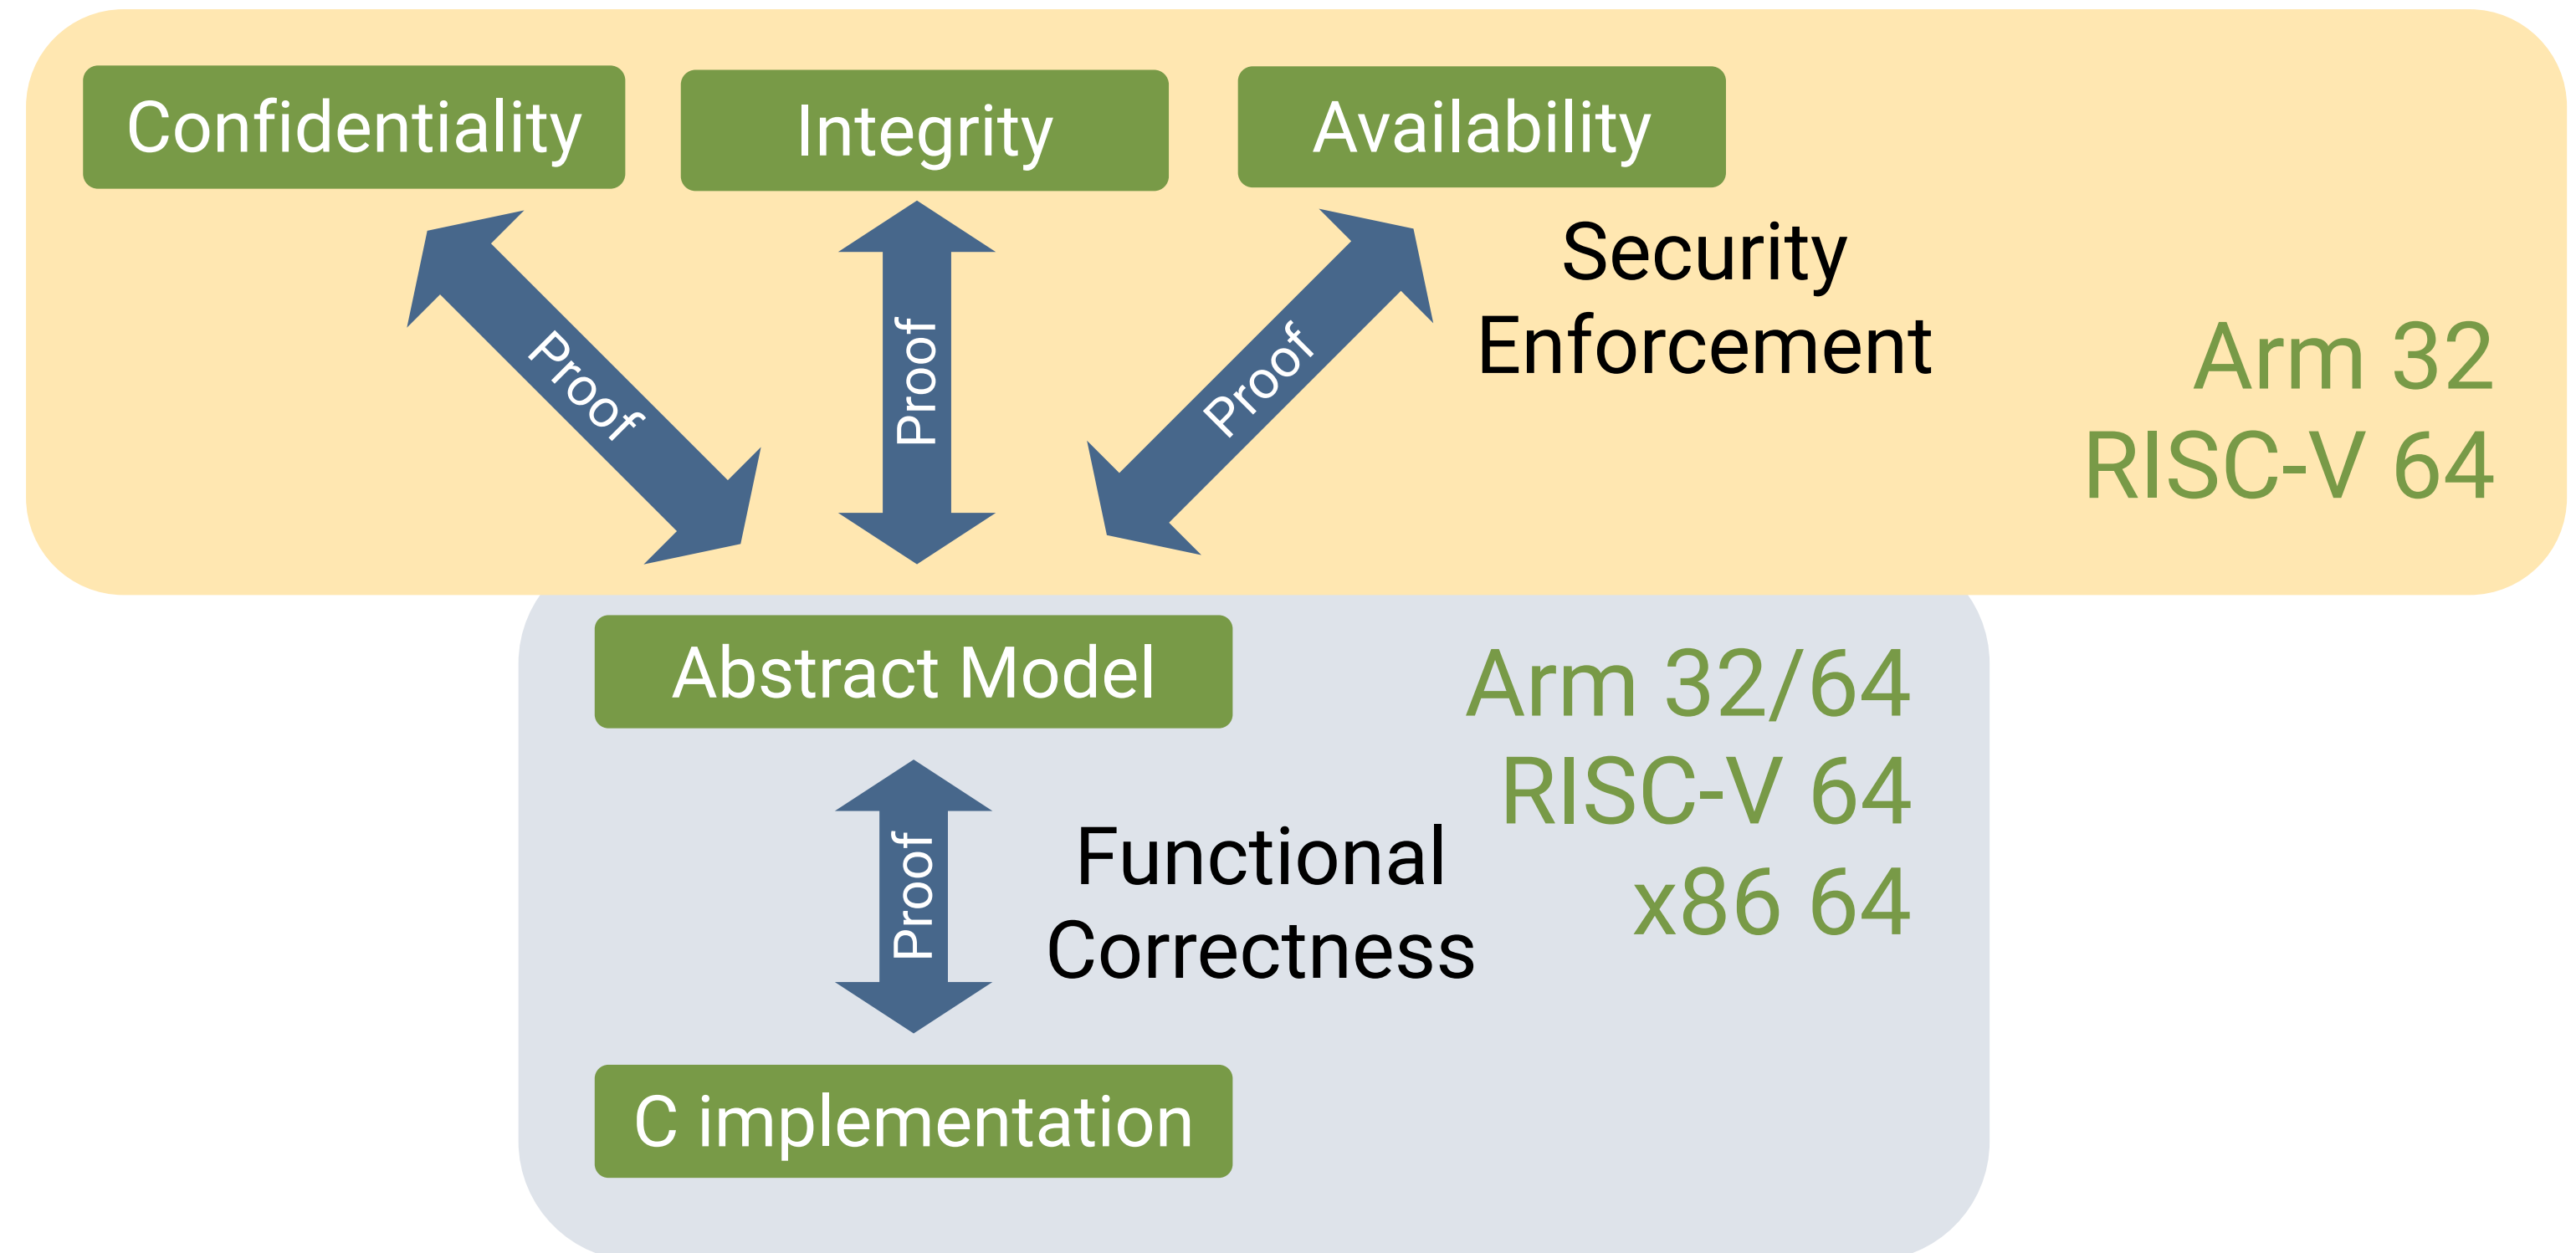
- **This talk:** How do we verify **Time Protection** on seL4 works?

- Implemented, evaluated empirically on ARM, x86
See EuroSys: [Ge et al. 2019]

- Ported to RISC-V with hardware support
See arXiv preprint: [Buckley, Sison et al. 2023]
(HW support by [Wistoff et al. 2023])

Thanks to funding from Australian Research Council

seL4: World's first OS kernel with correctness proof!





How is seL4 verified?

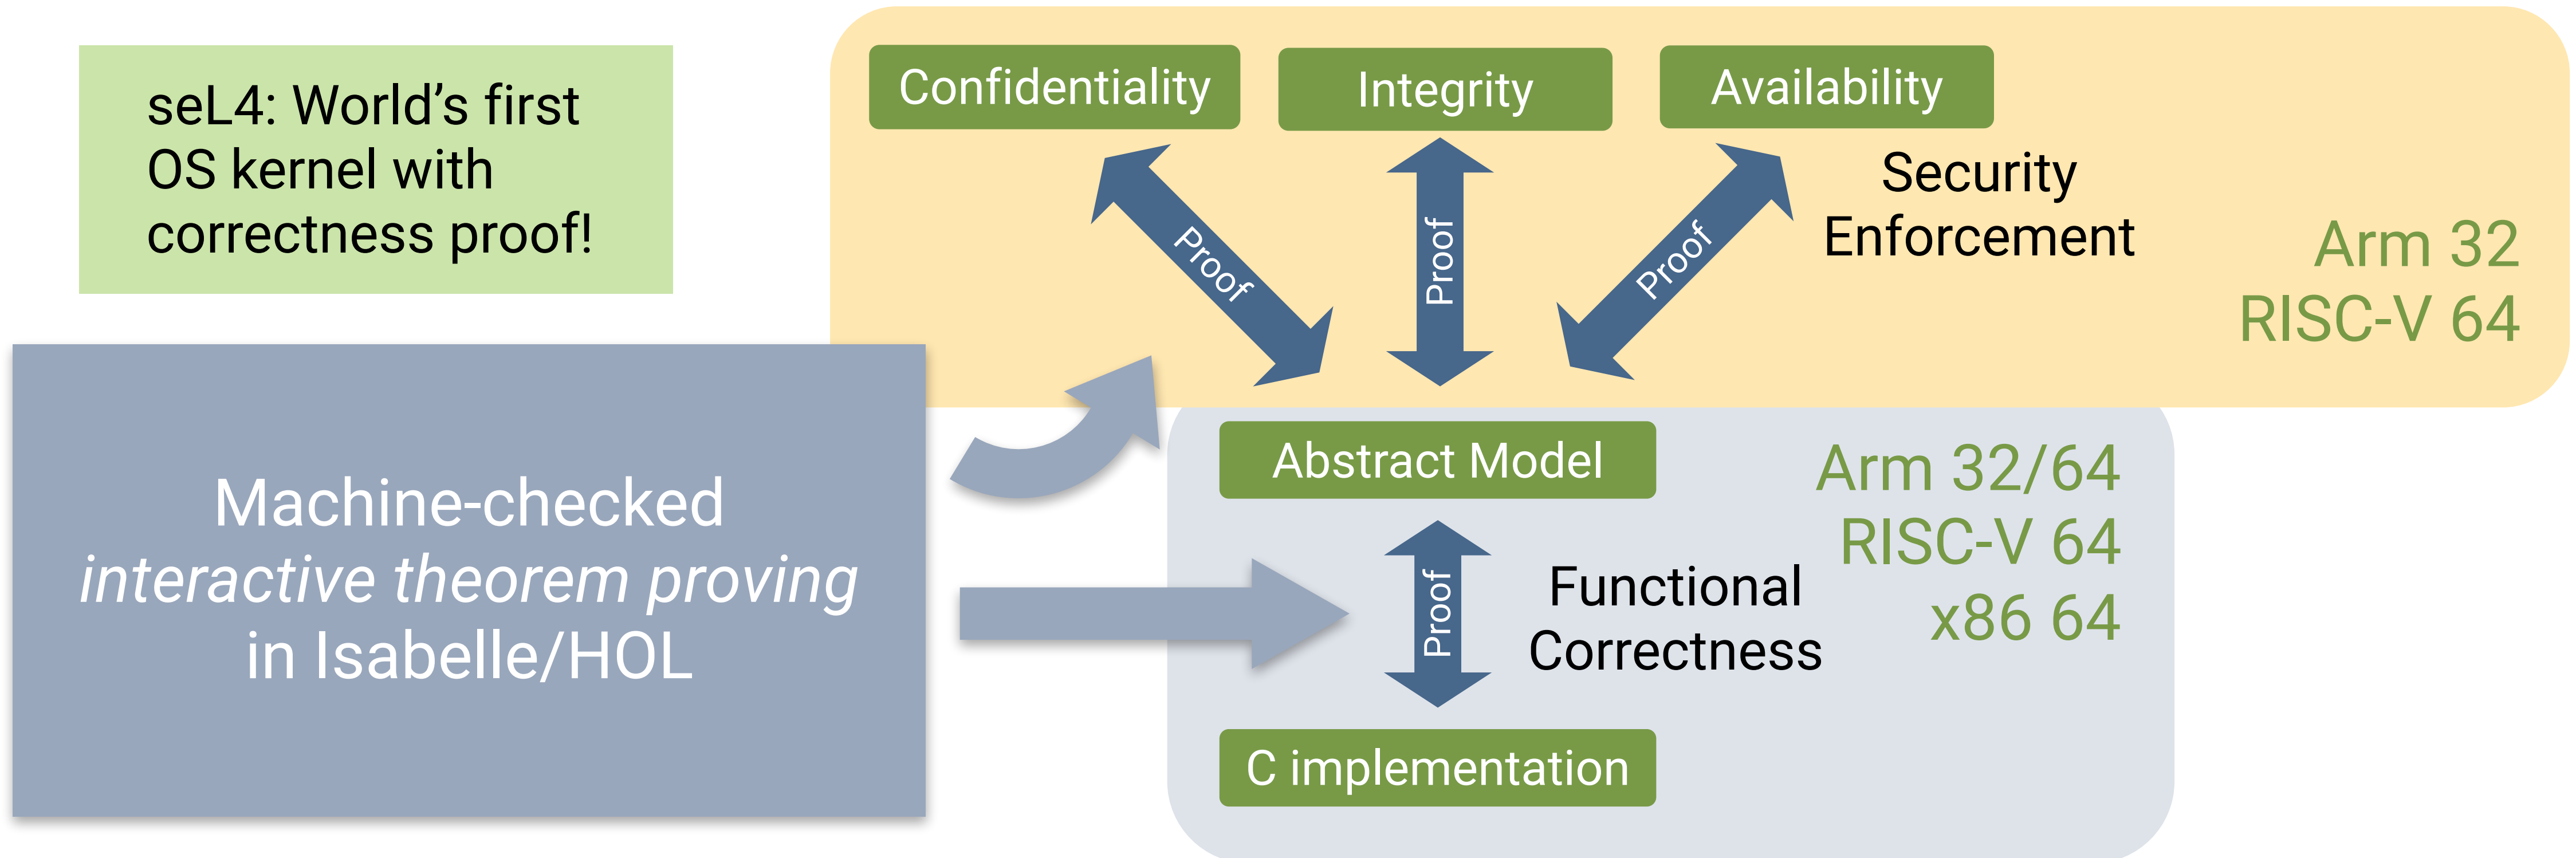


- **This talk:** How do we verify **Time Protection** on seL4 works?

- Implemented, evaluated empirically on ARM, x86
See EuroSys: [Ge et al. 2019]

- Ported to RISC-V with hardware support
See arXiv preprint: [Buckley, Sison et al. 2023]
(HW support by [Wistoff et al. 2023])

Thanks to funding from Australian Research Council



- **This talk:** How do we verify **Time Protection** on seL4 works?

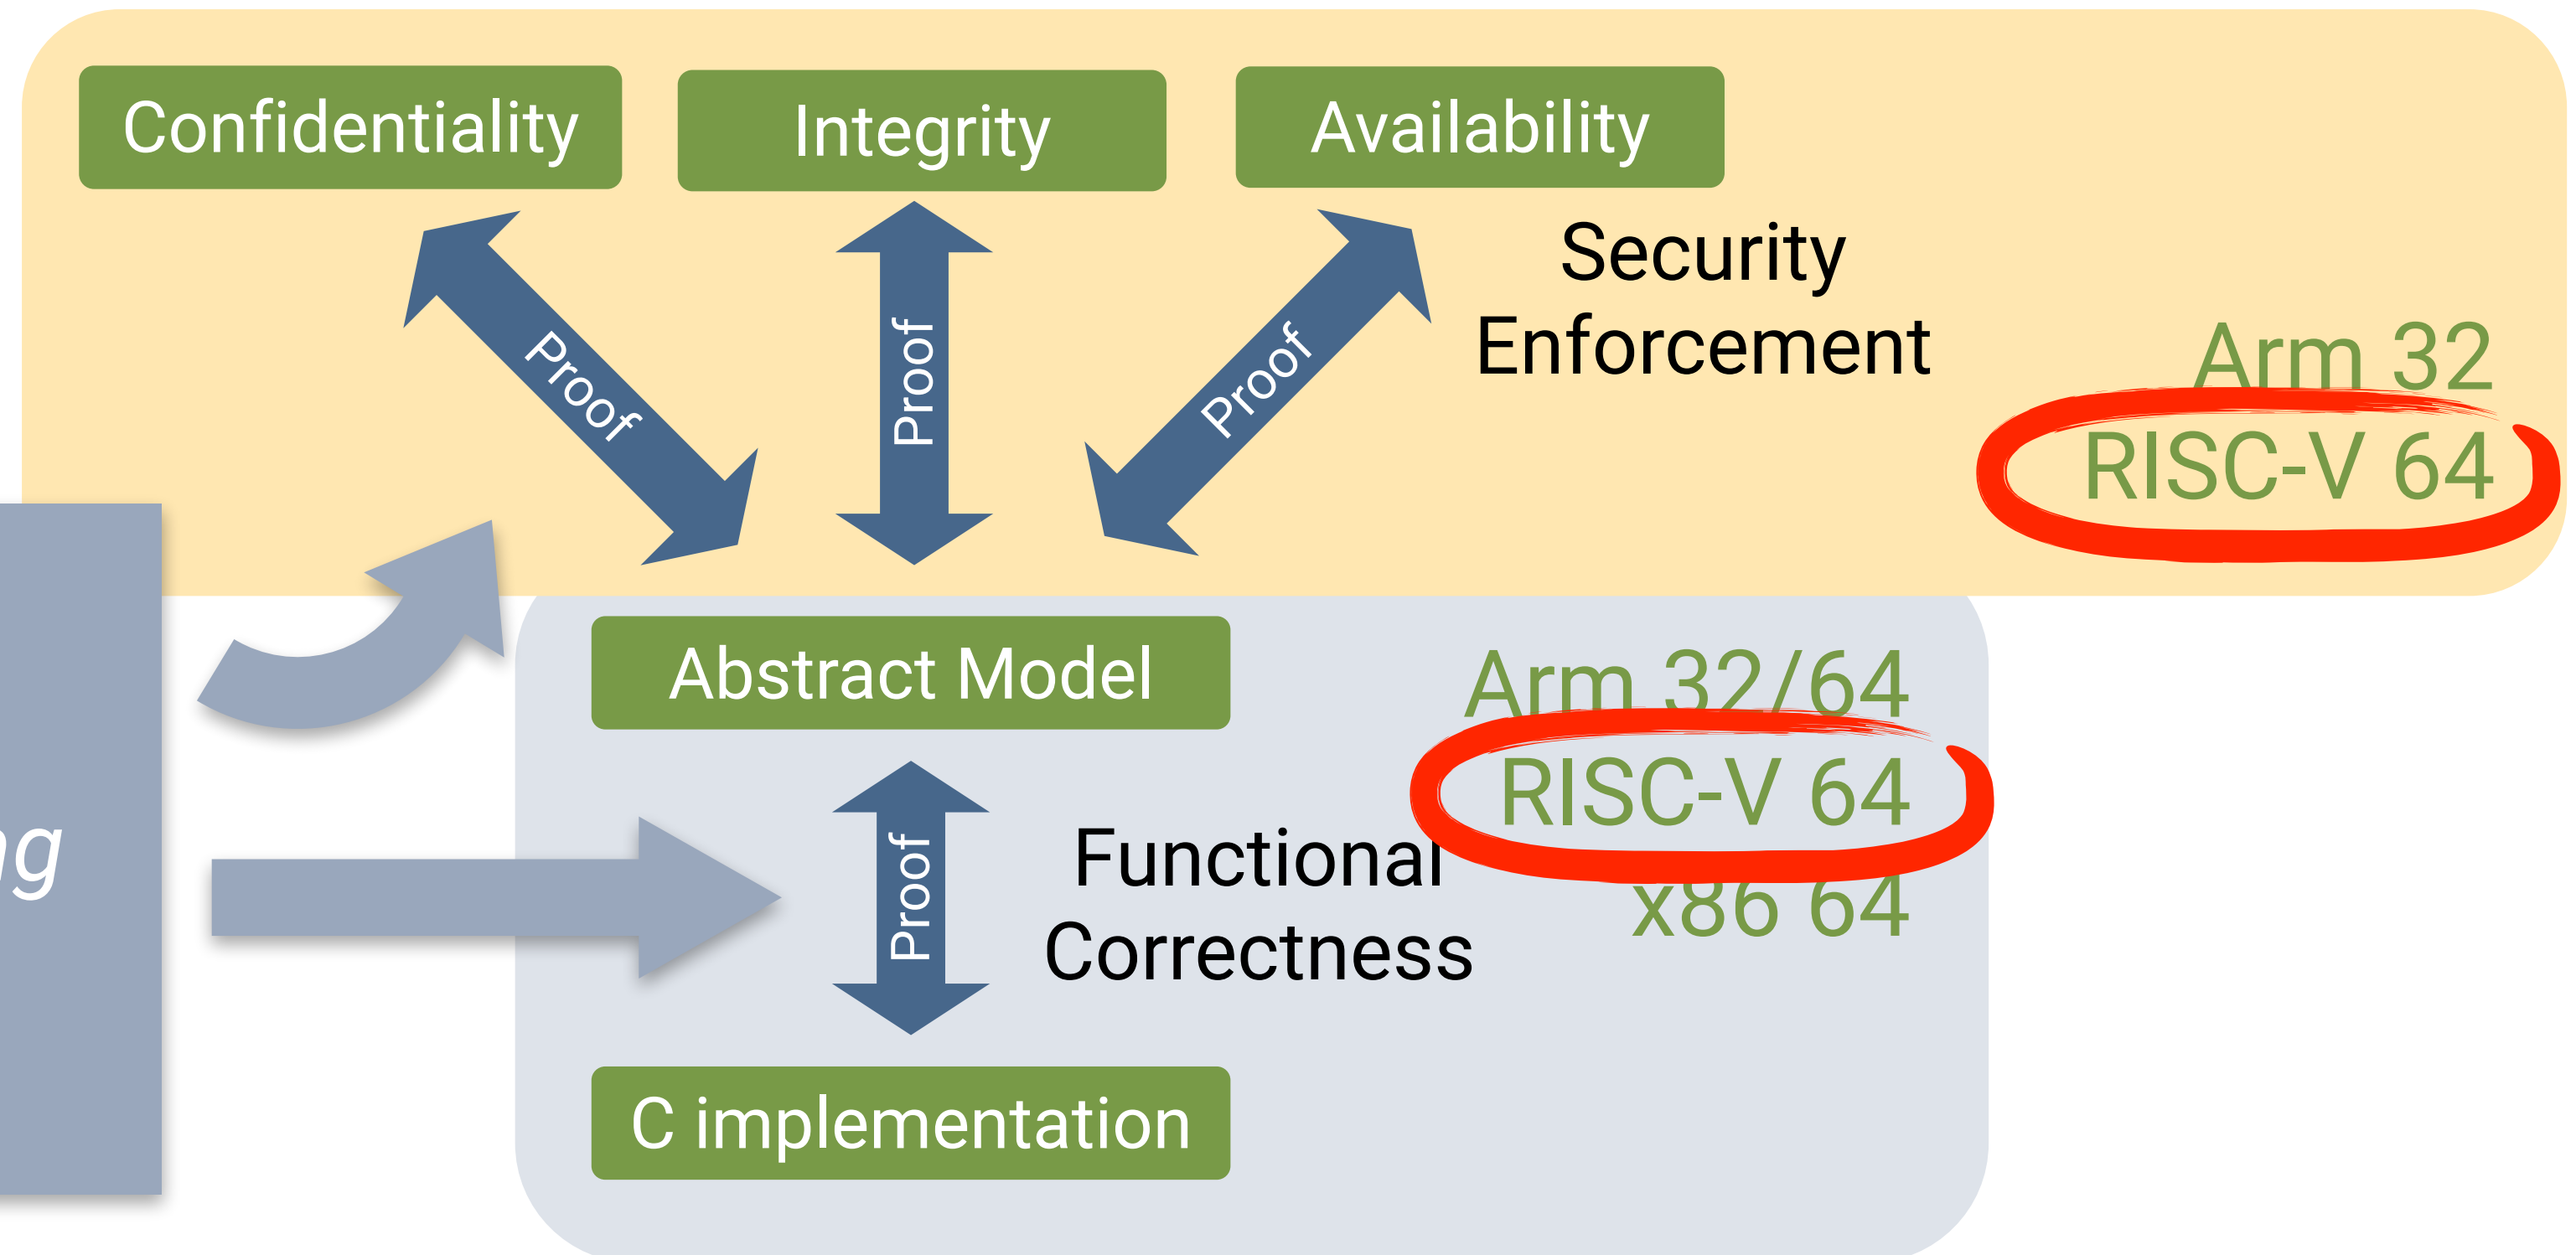
- Implemented, evaluated empirically on ARM, x86
See EuroSys: [Ge et al. 2019]

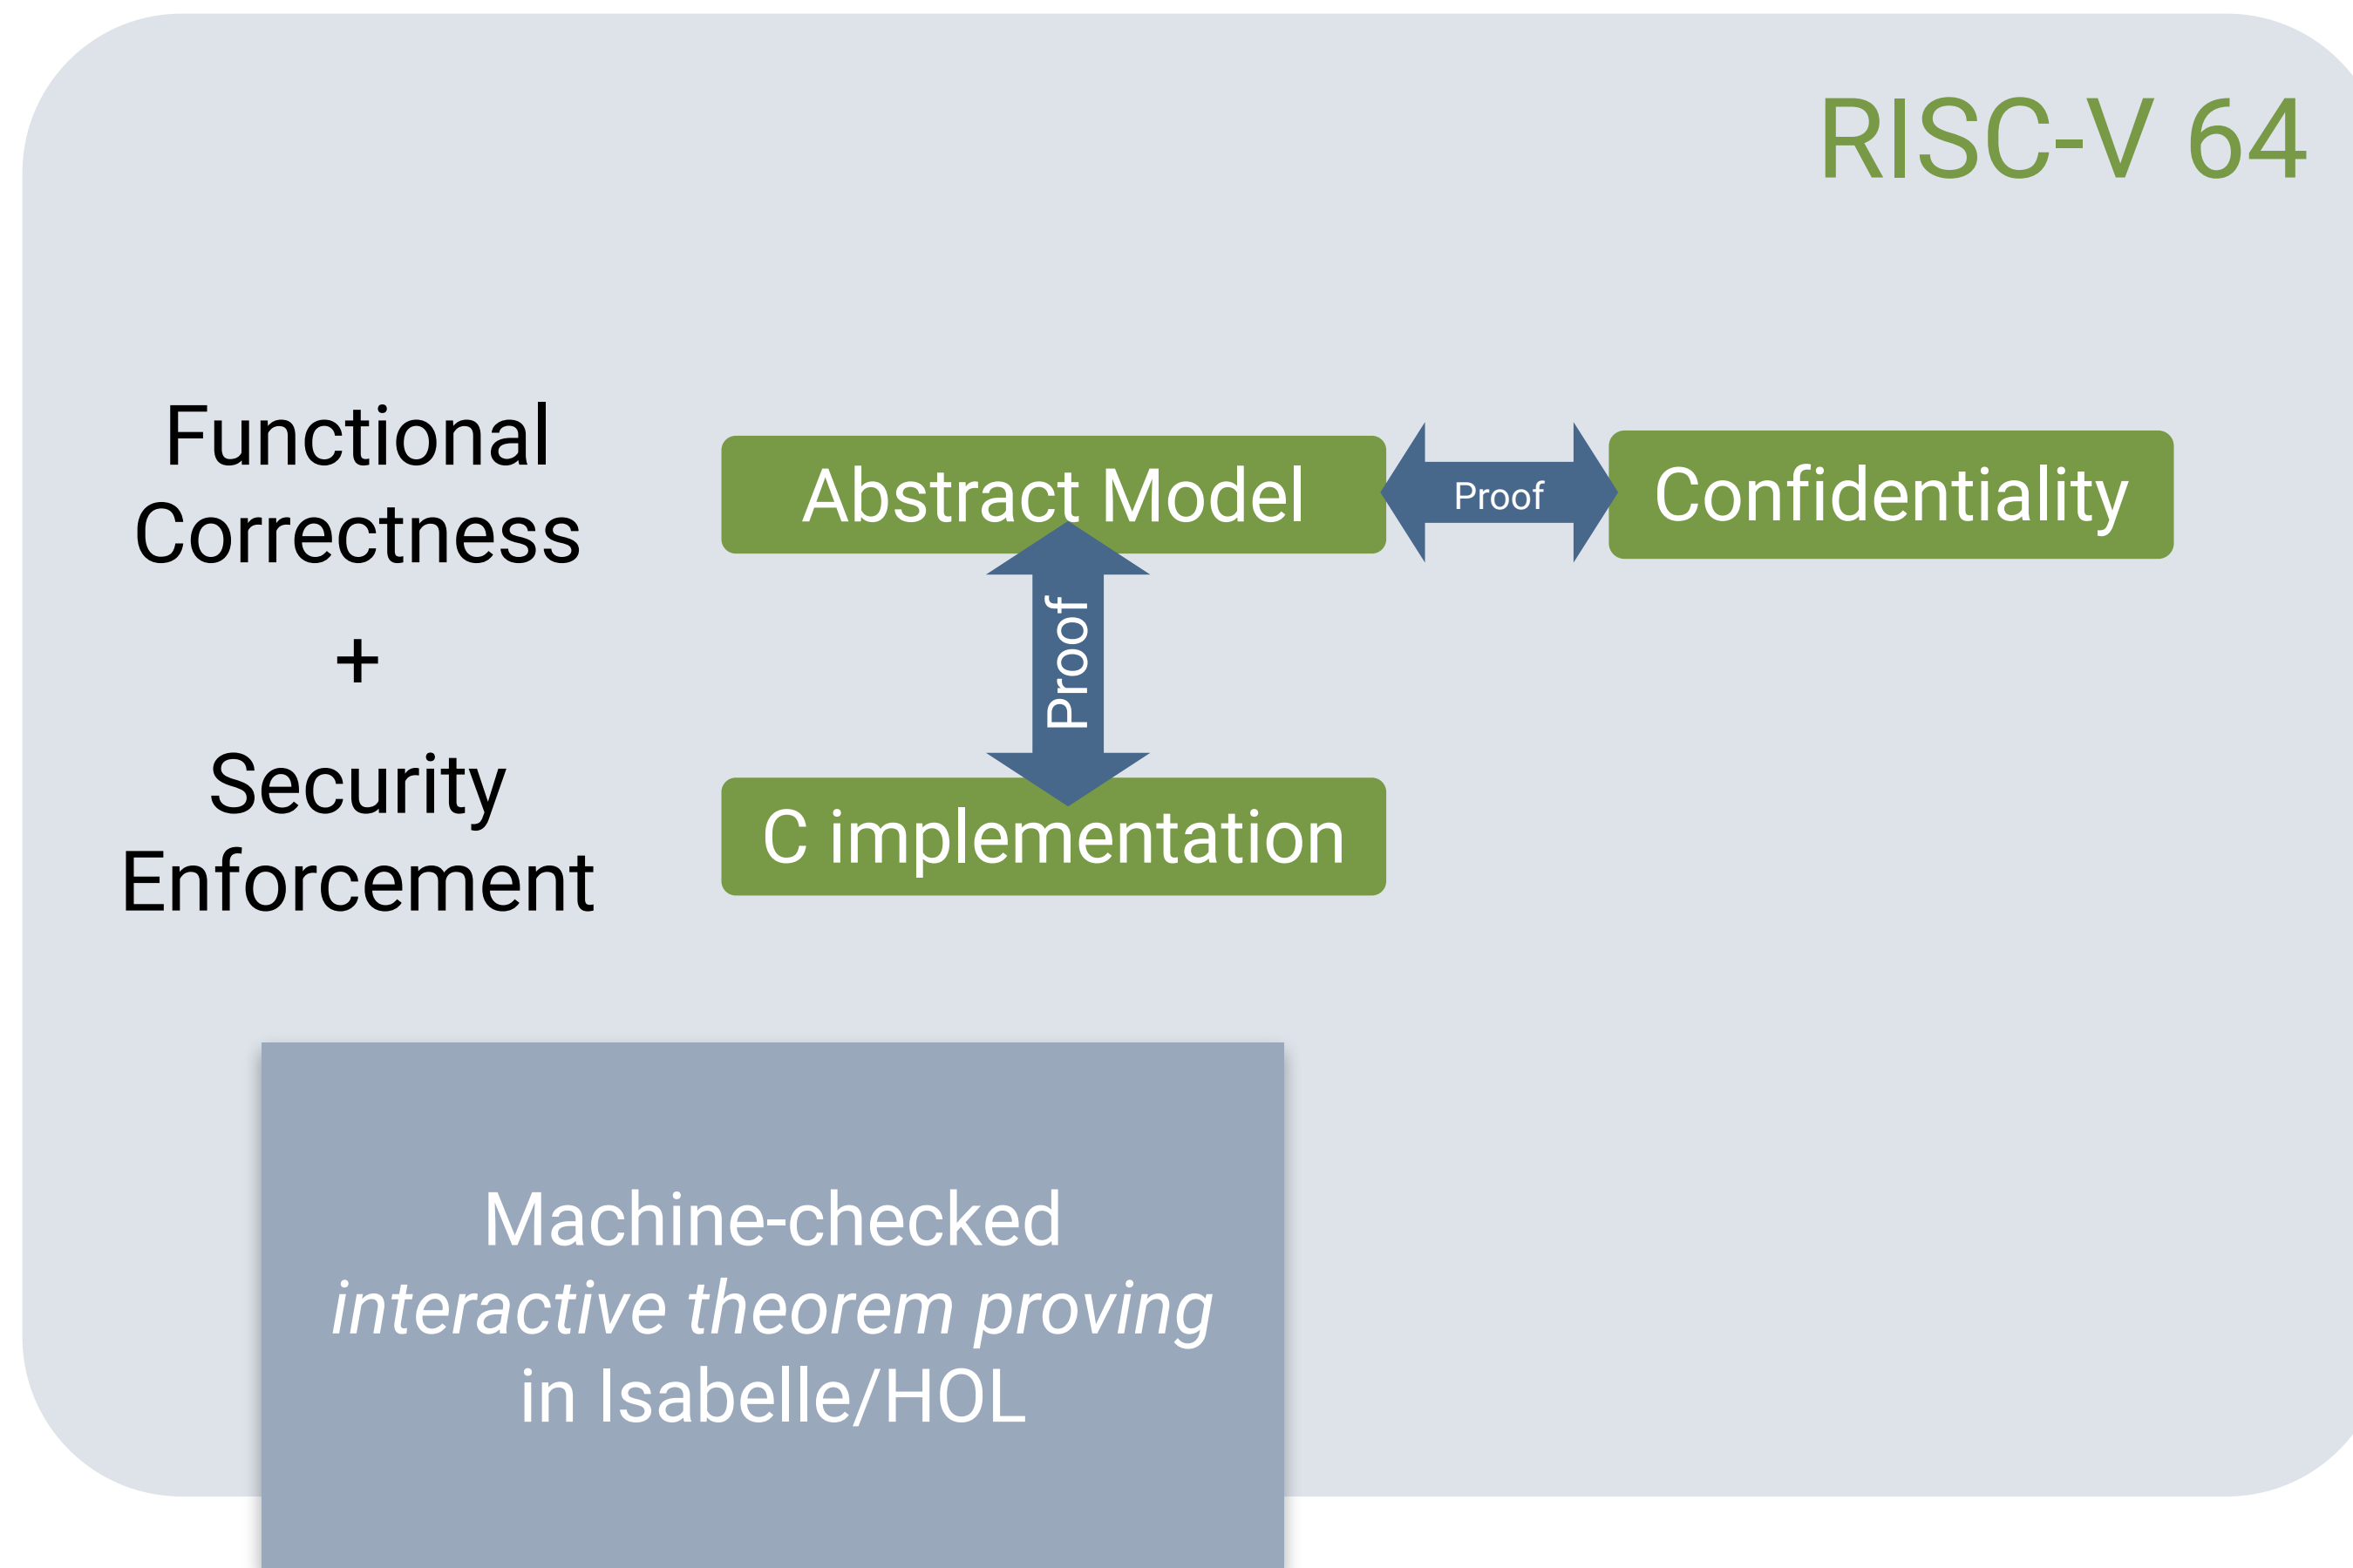
- Ported to RISC-V with hardware support
See arXiv preprint: [Buckley, Sison et al. 2023]
(HW support by [Wistoff et al. 2023])

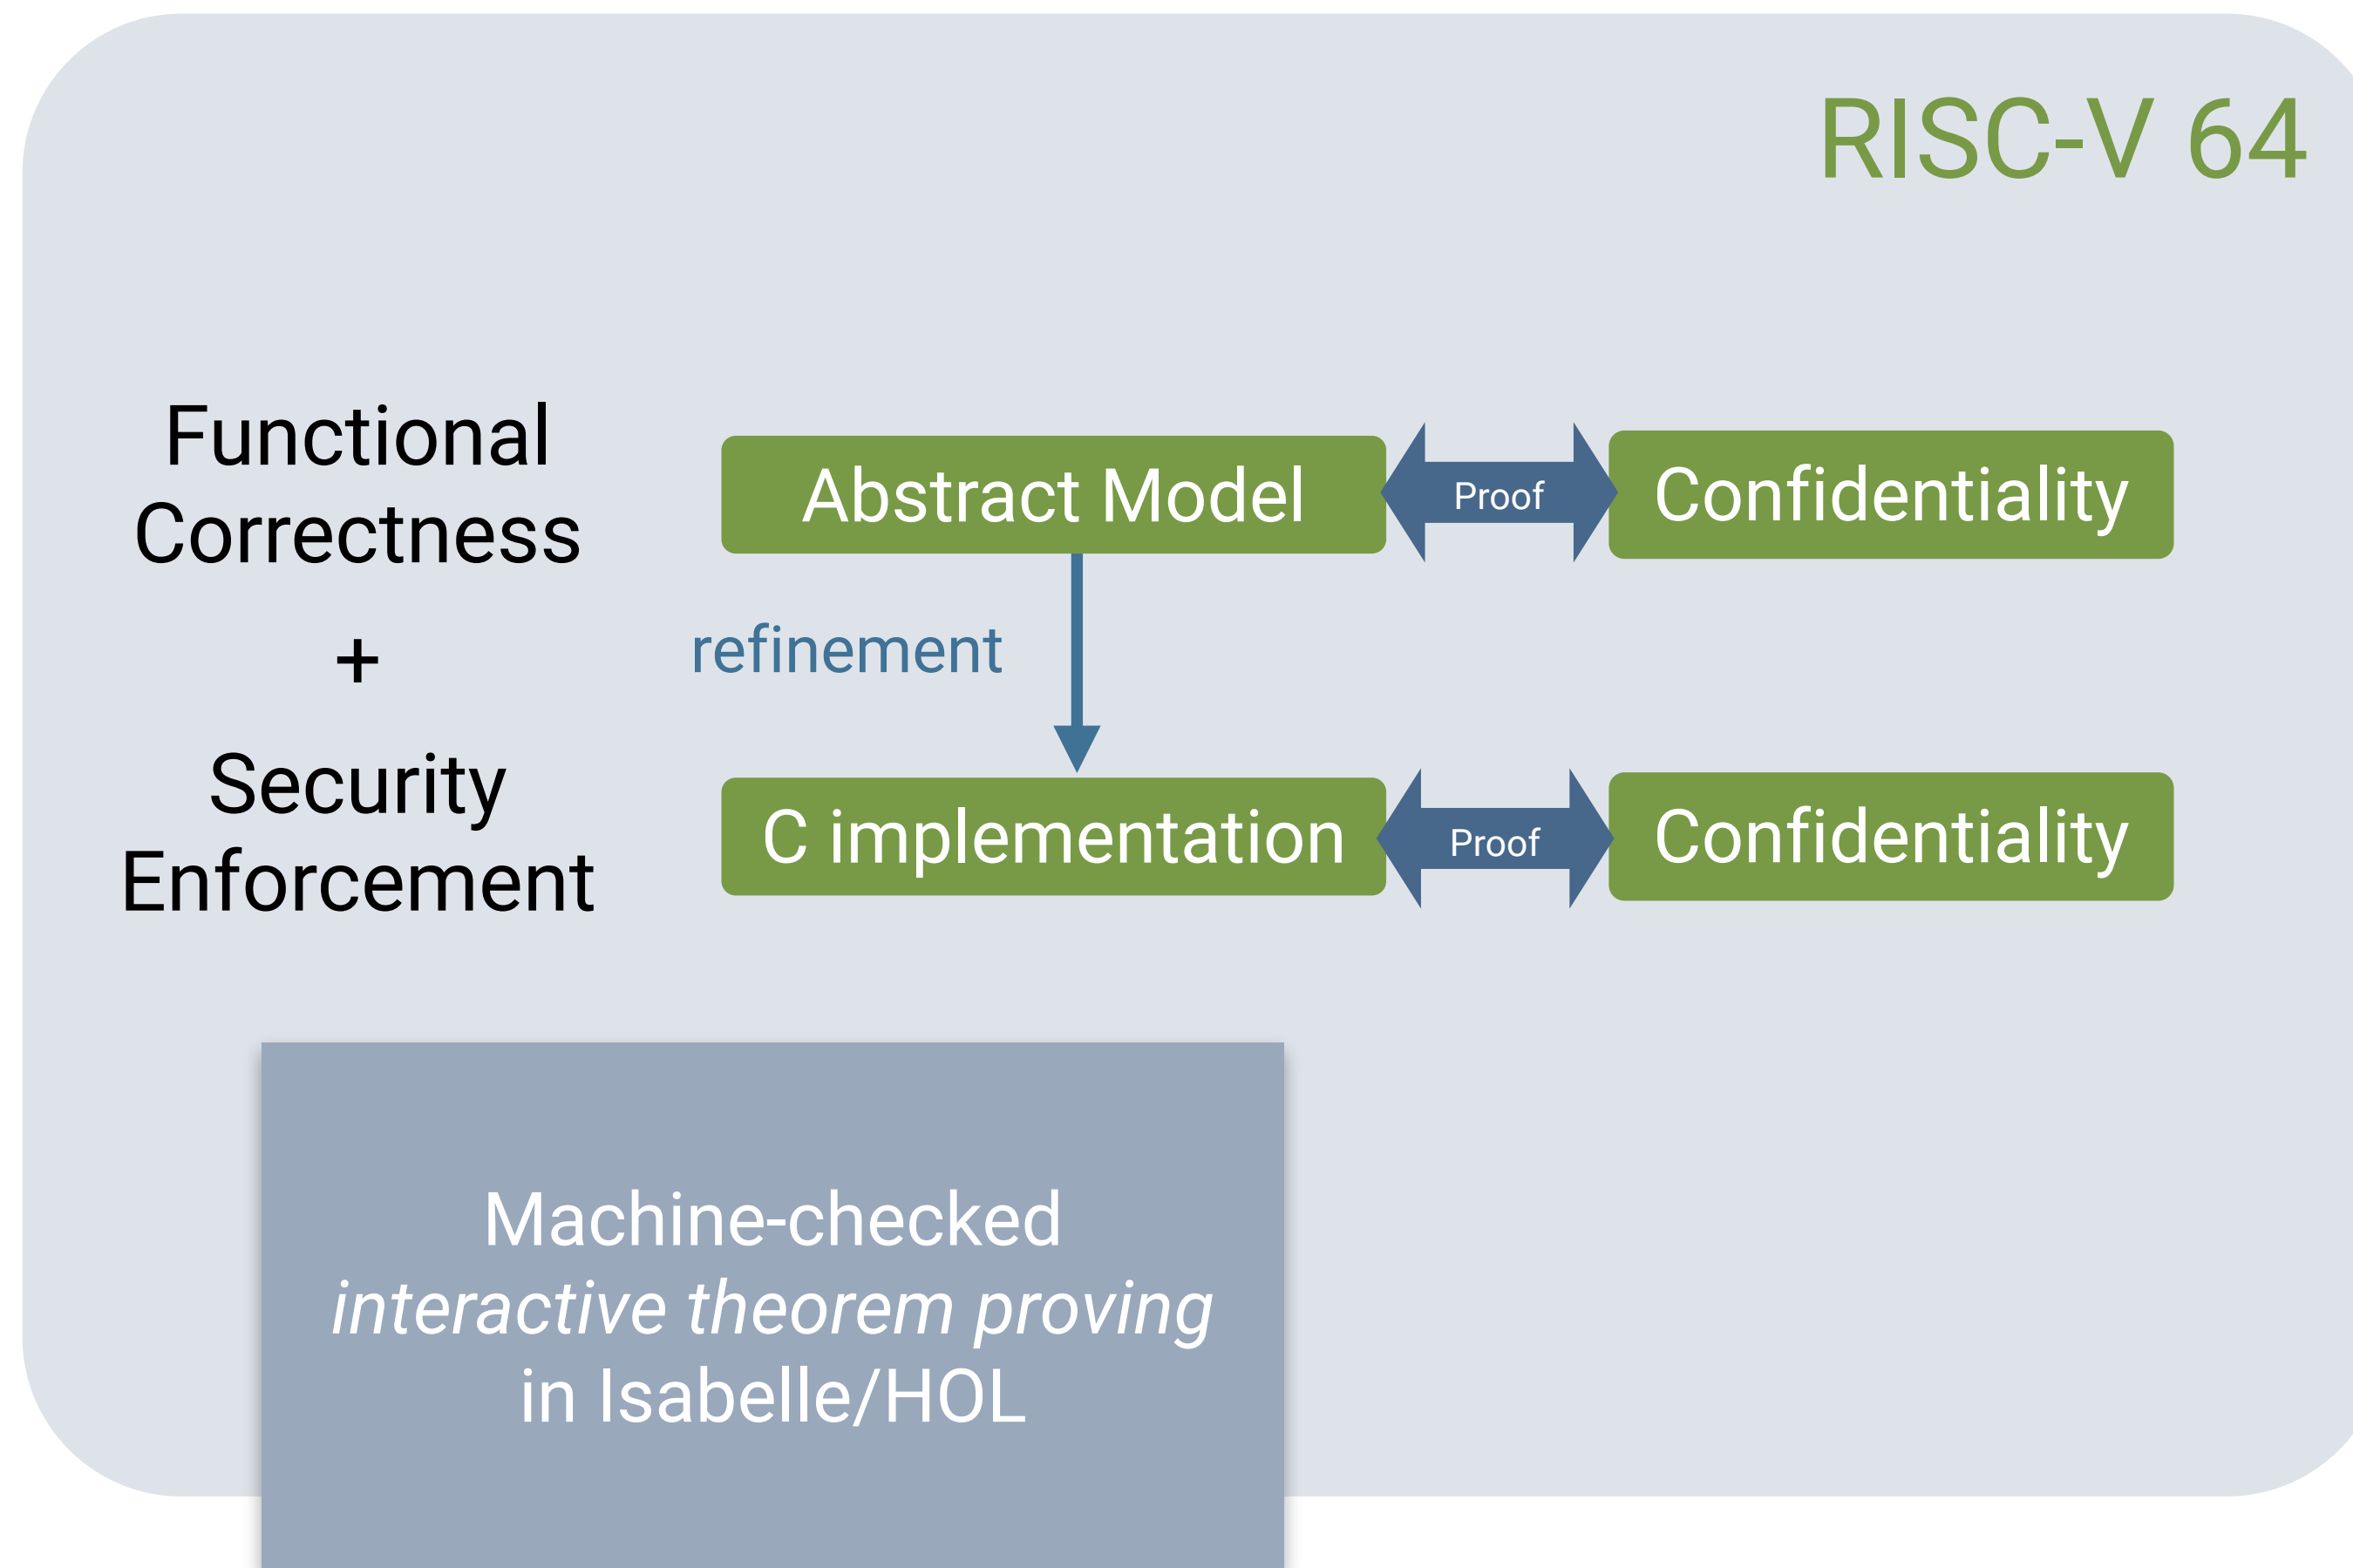
Thanks to funding from Australian Research Council

seL4: World's first OS kernel with correctness proof!

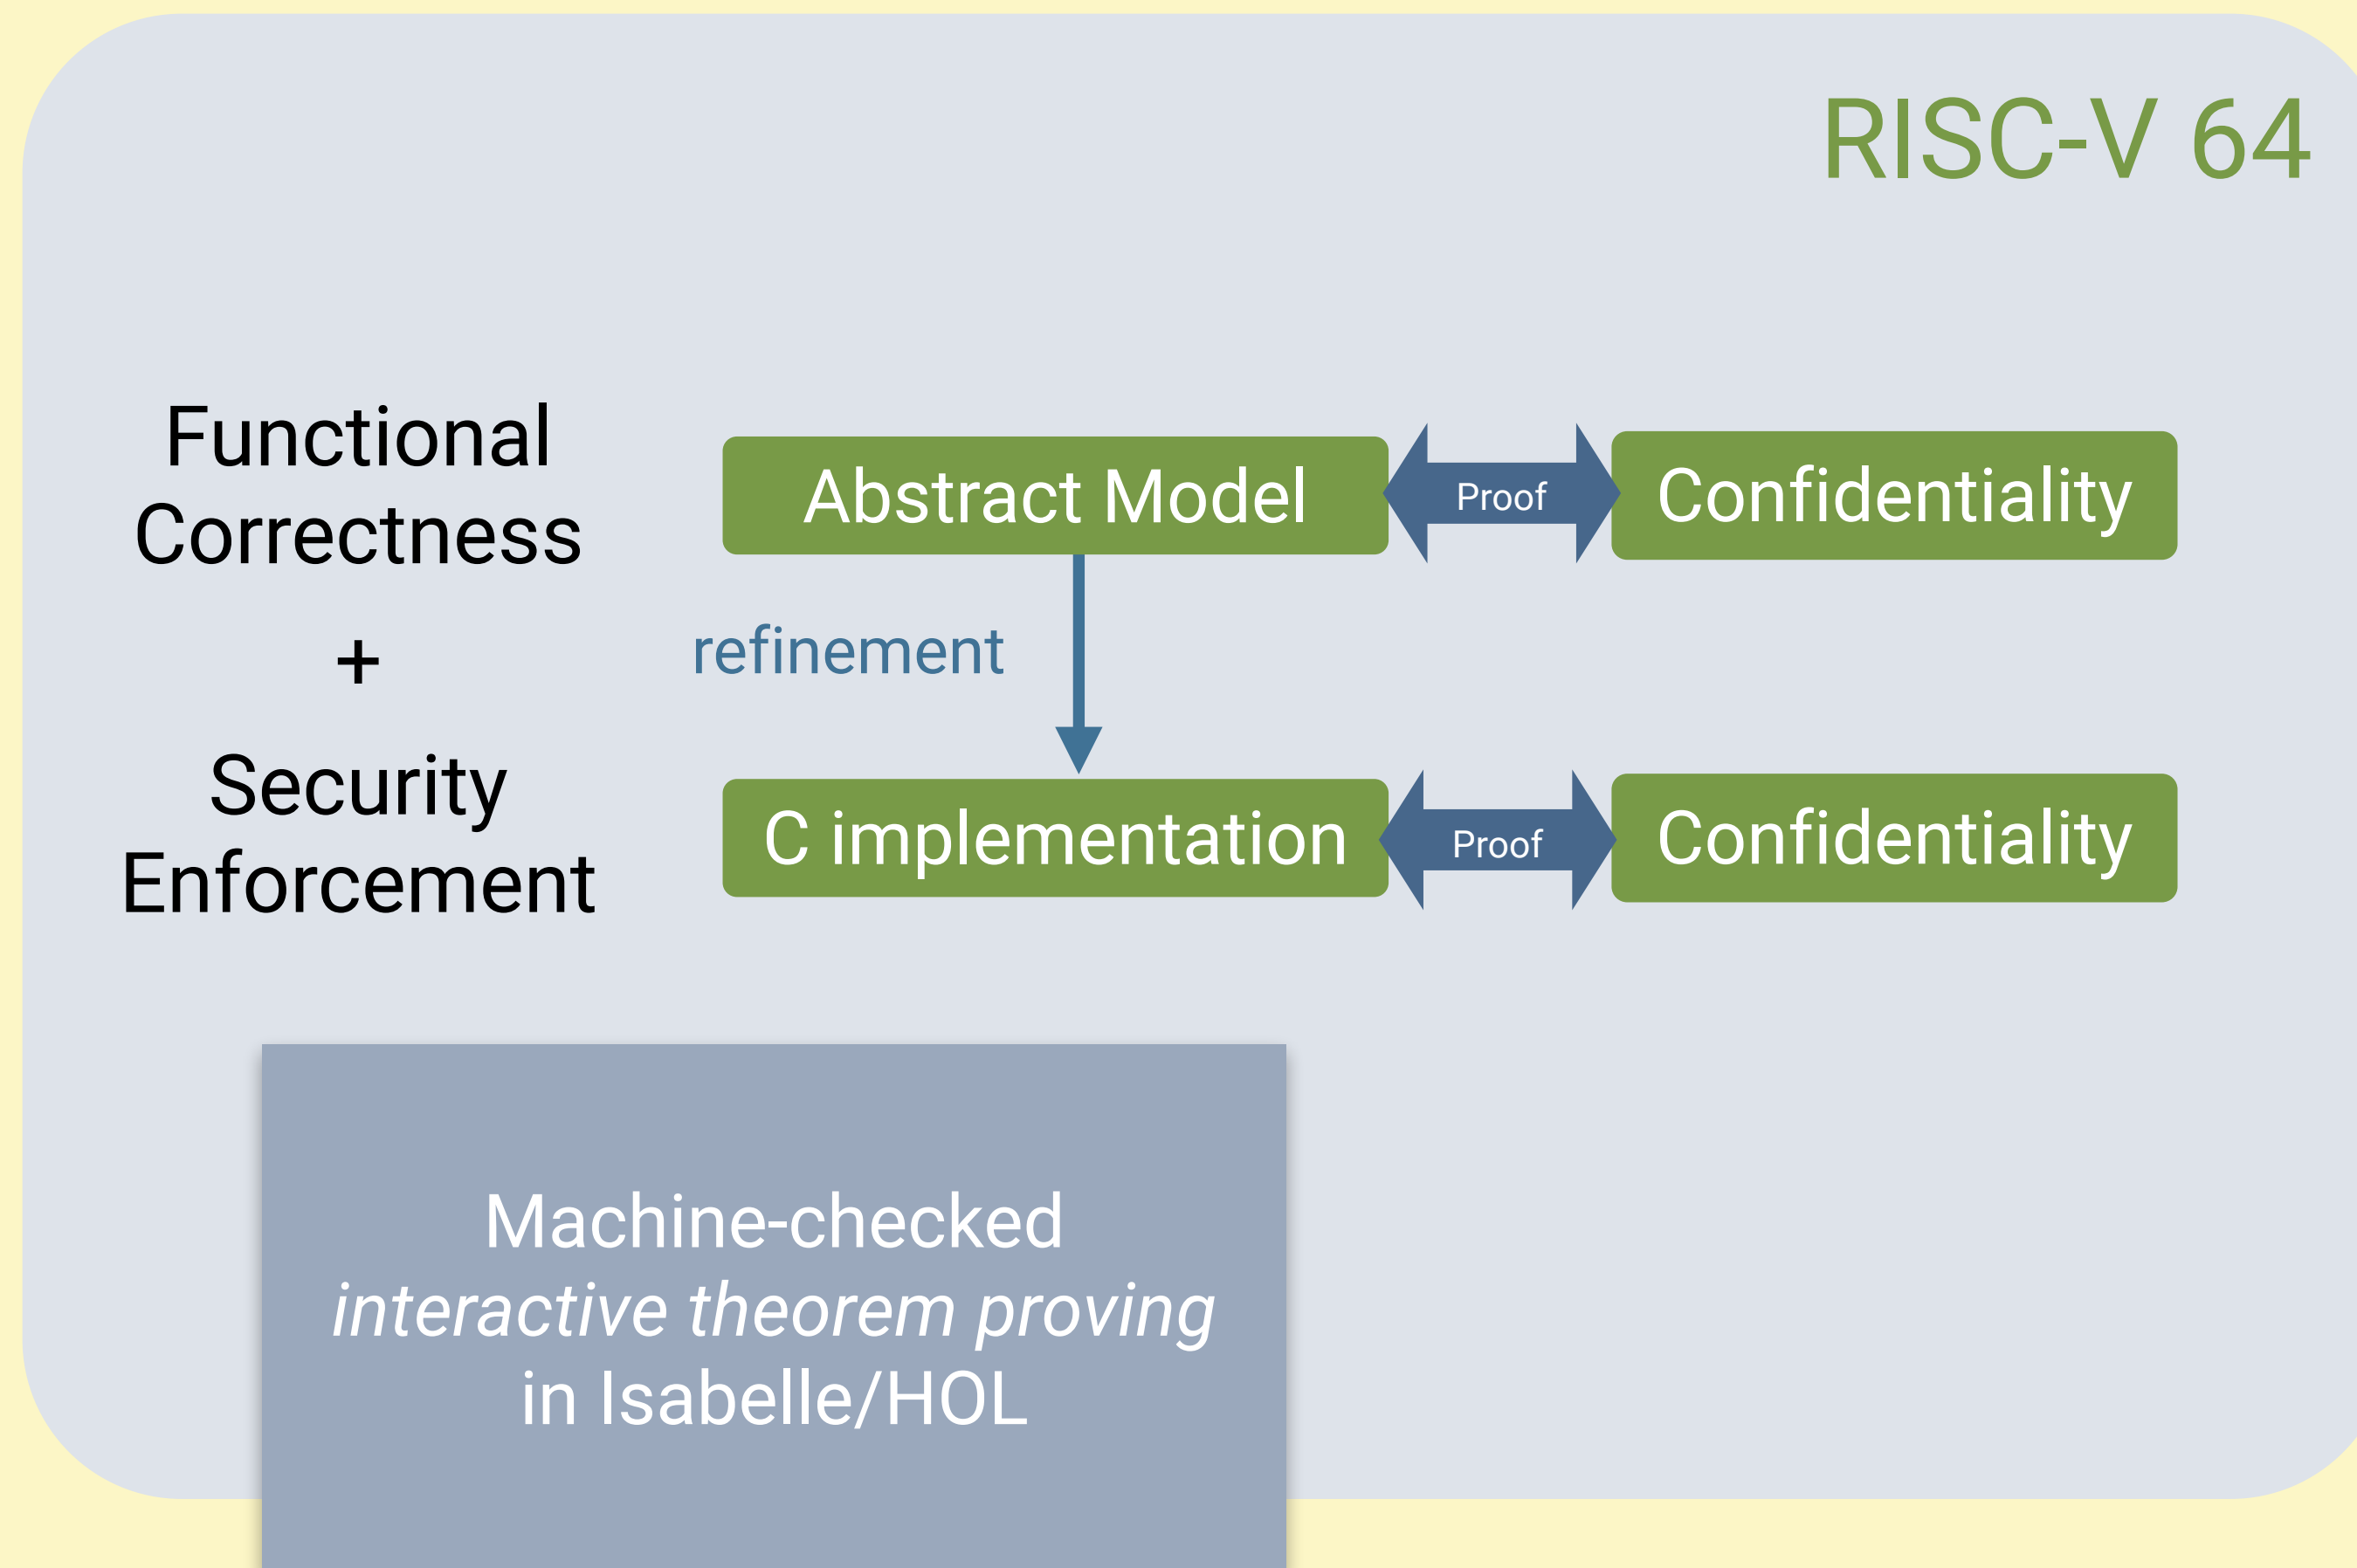
Machine-checked *interactive theorem proving* in Isabelle/HOL





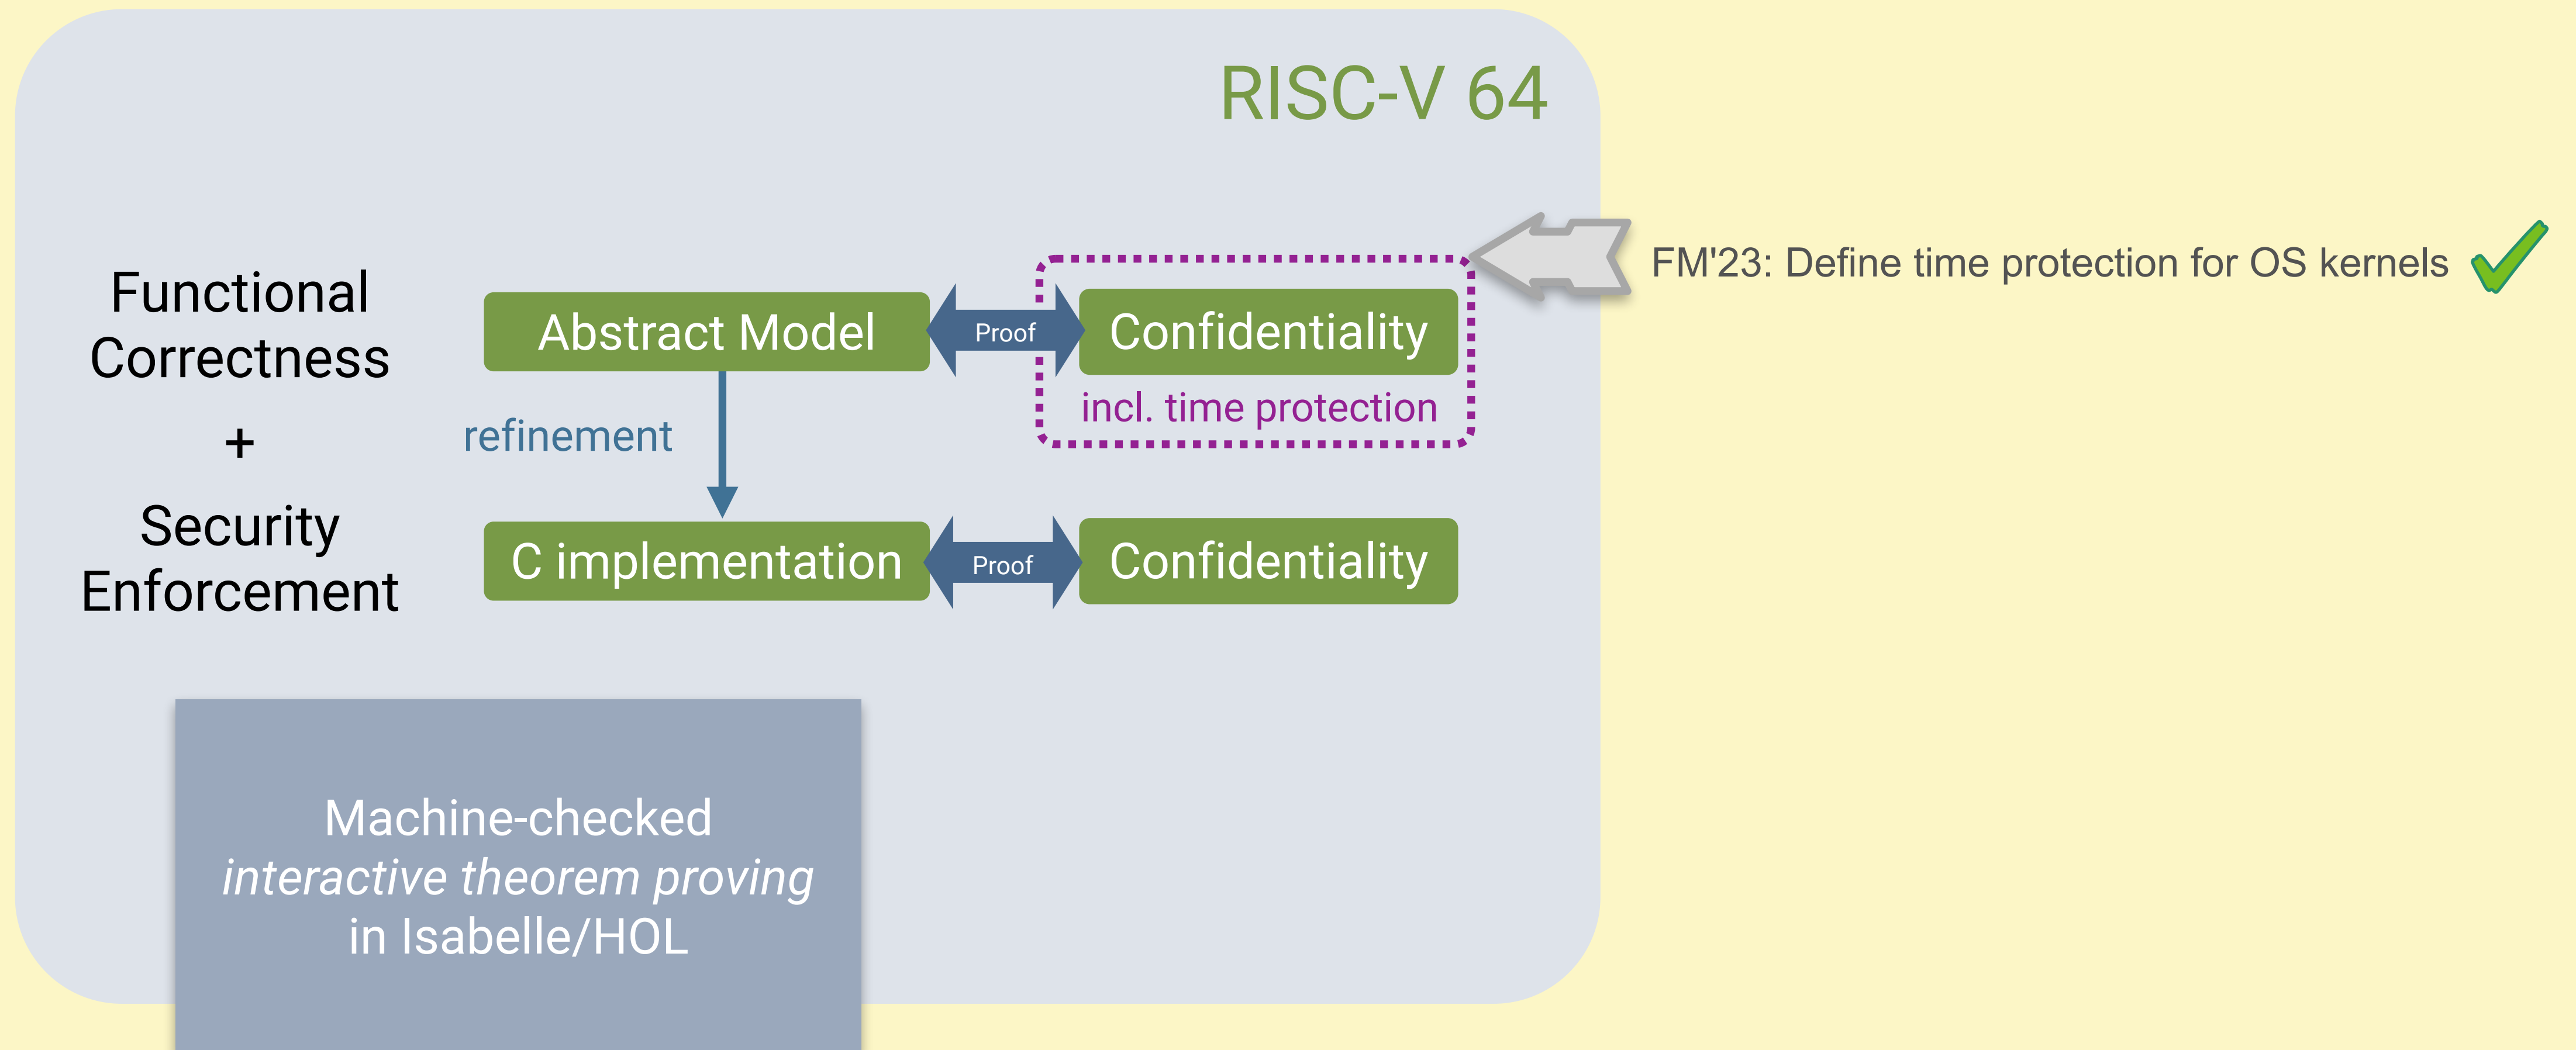


The plan 2 years ago... (presented at FM-OZ'24)



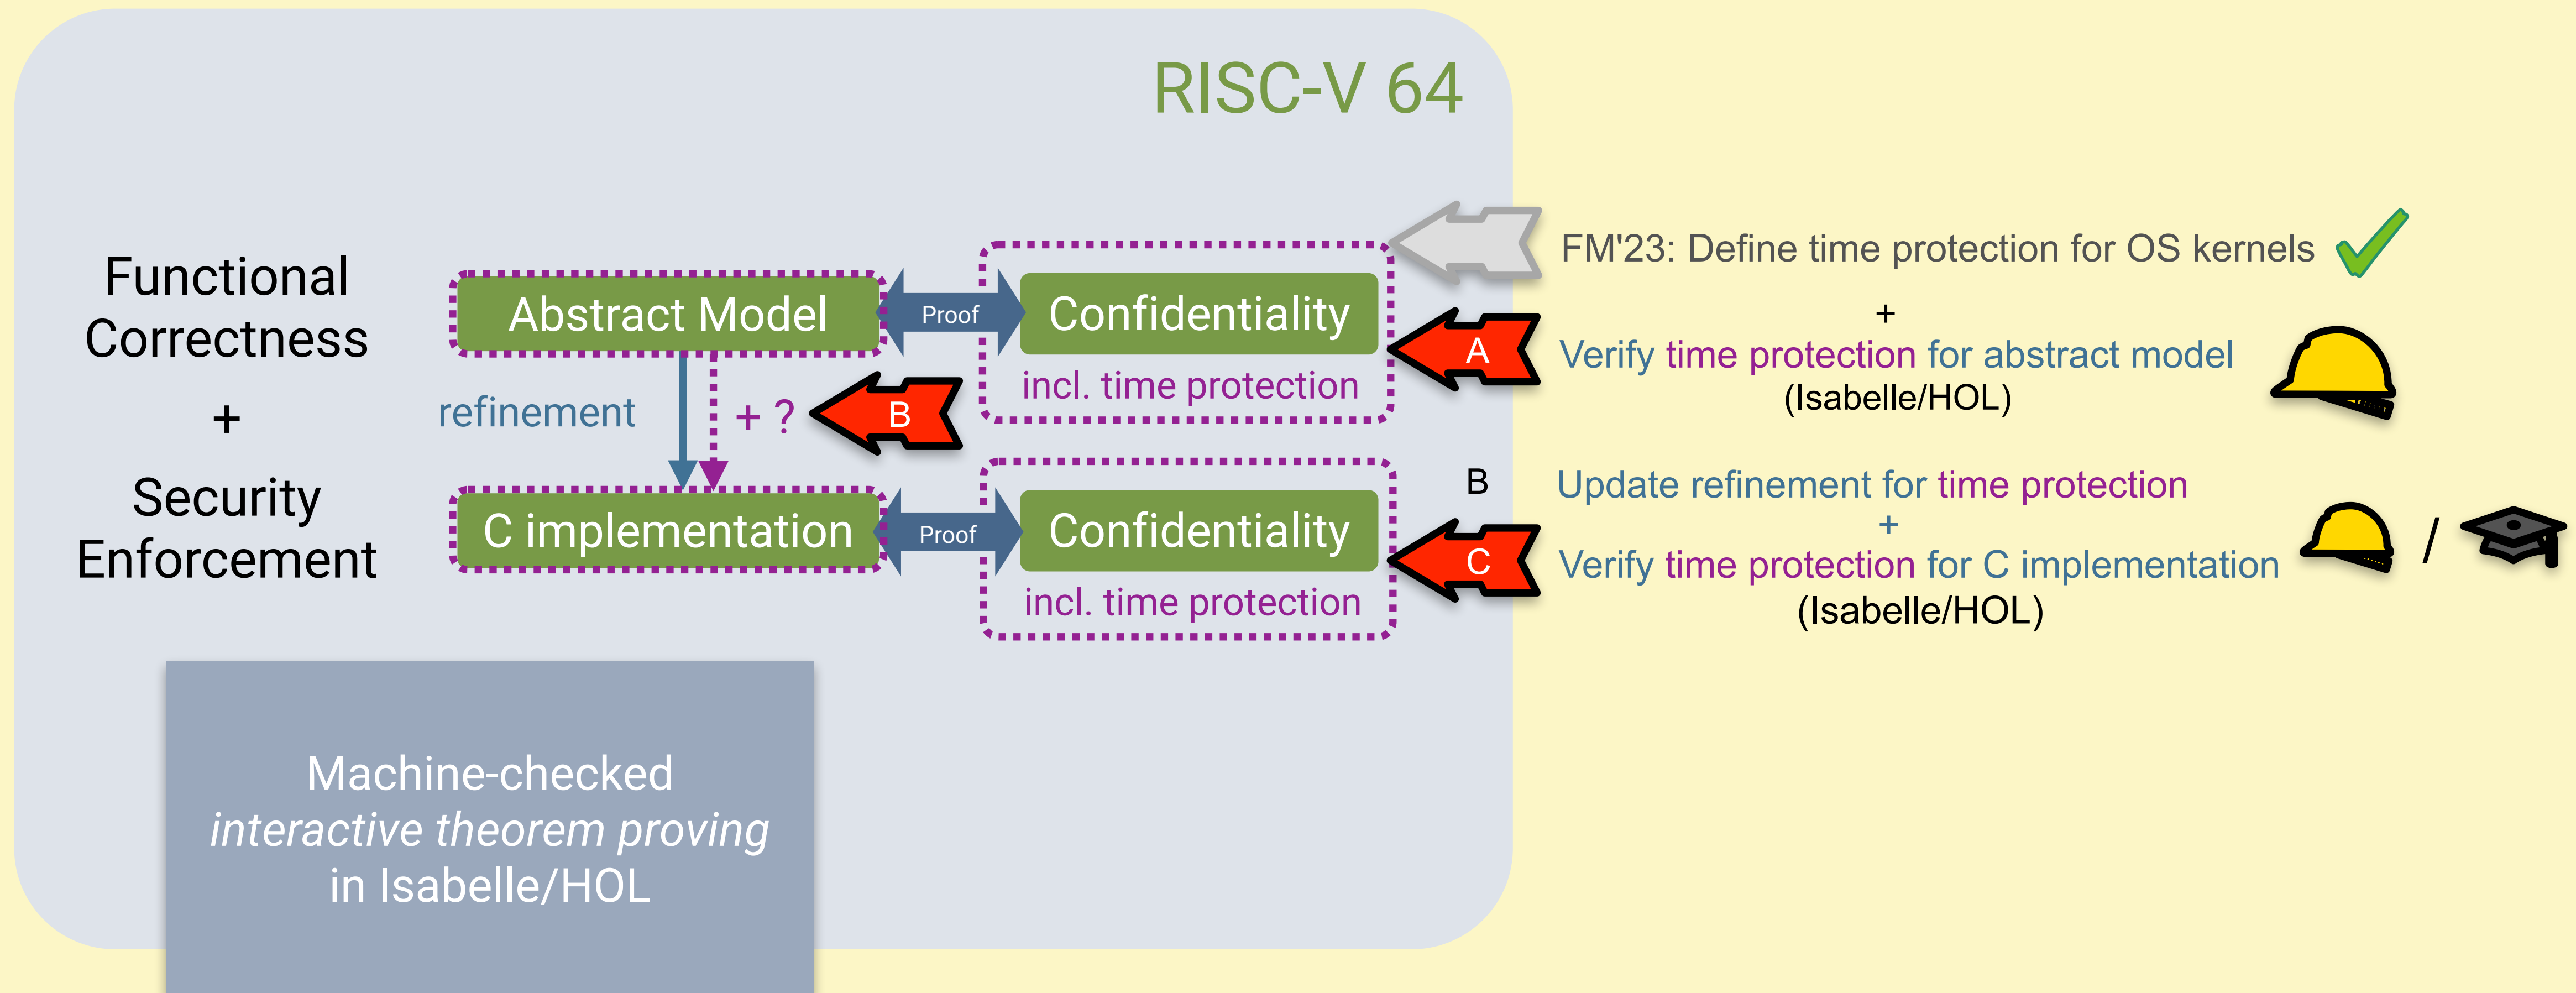
The plan 2 years ago... (presented at FM-OZ'24)

Thanks to funding from Australian Research Council



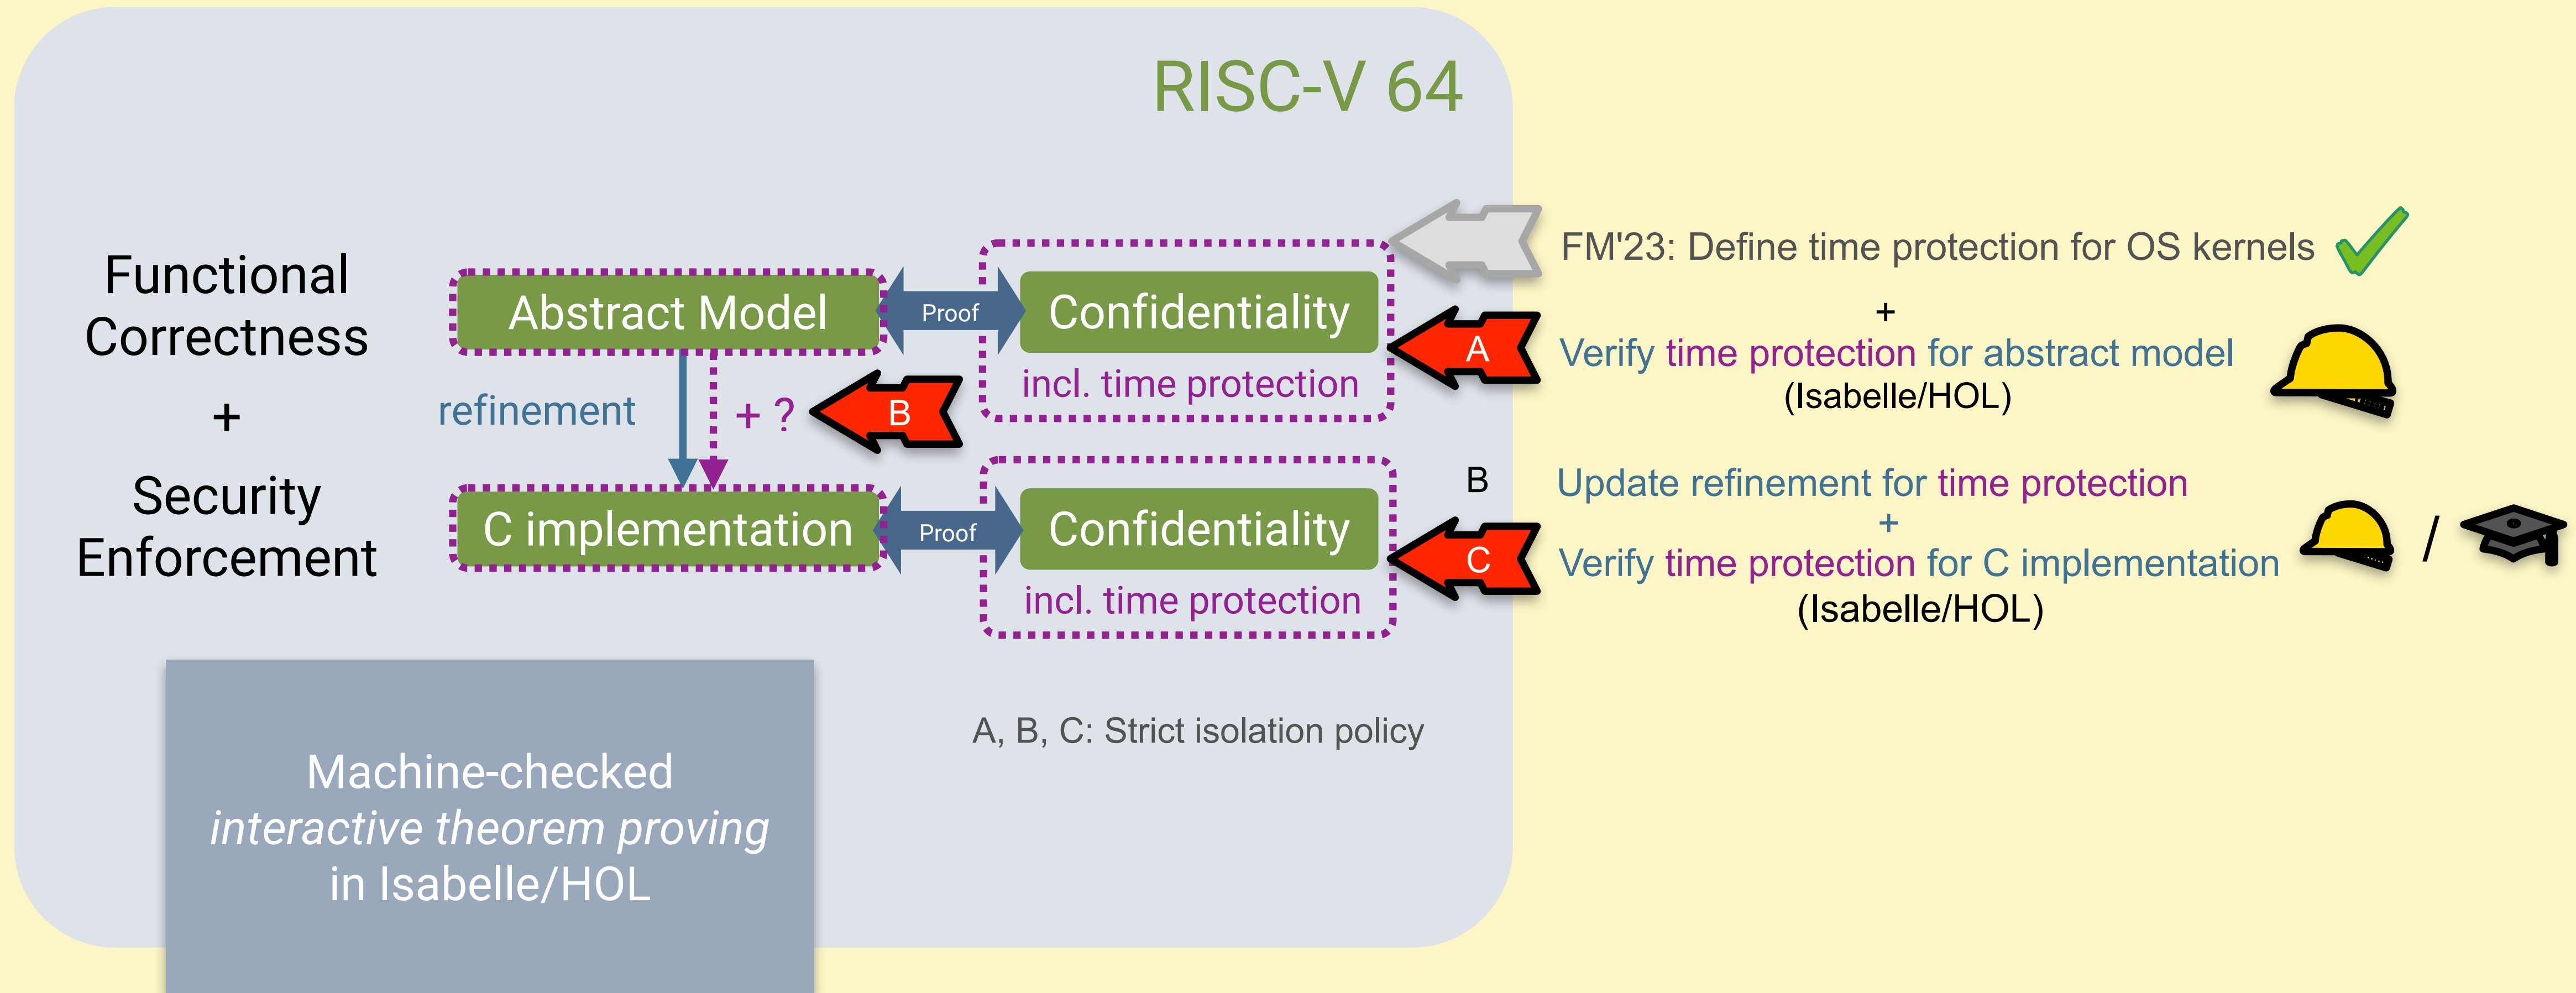
The plan 2 years ago... (presented at FM-OZ'24)

Thanks to funding from Australian Research Council



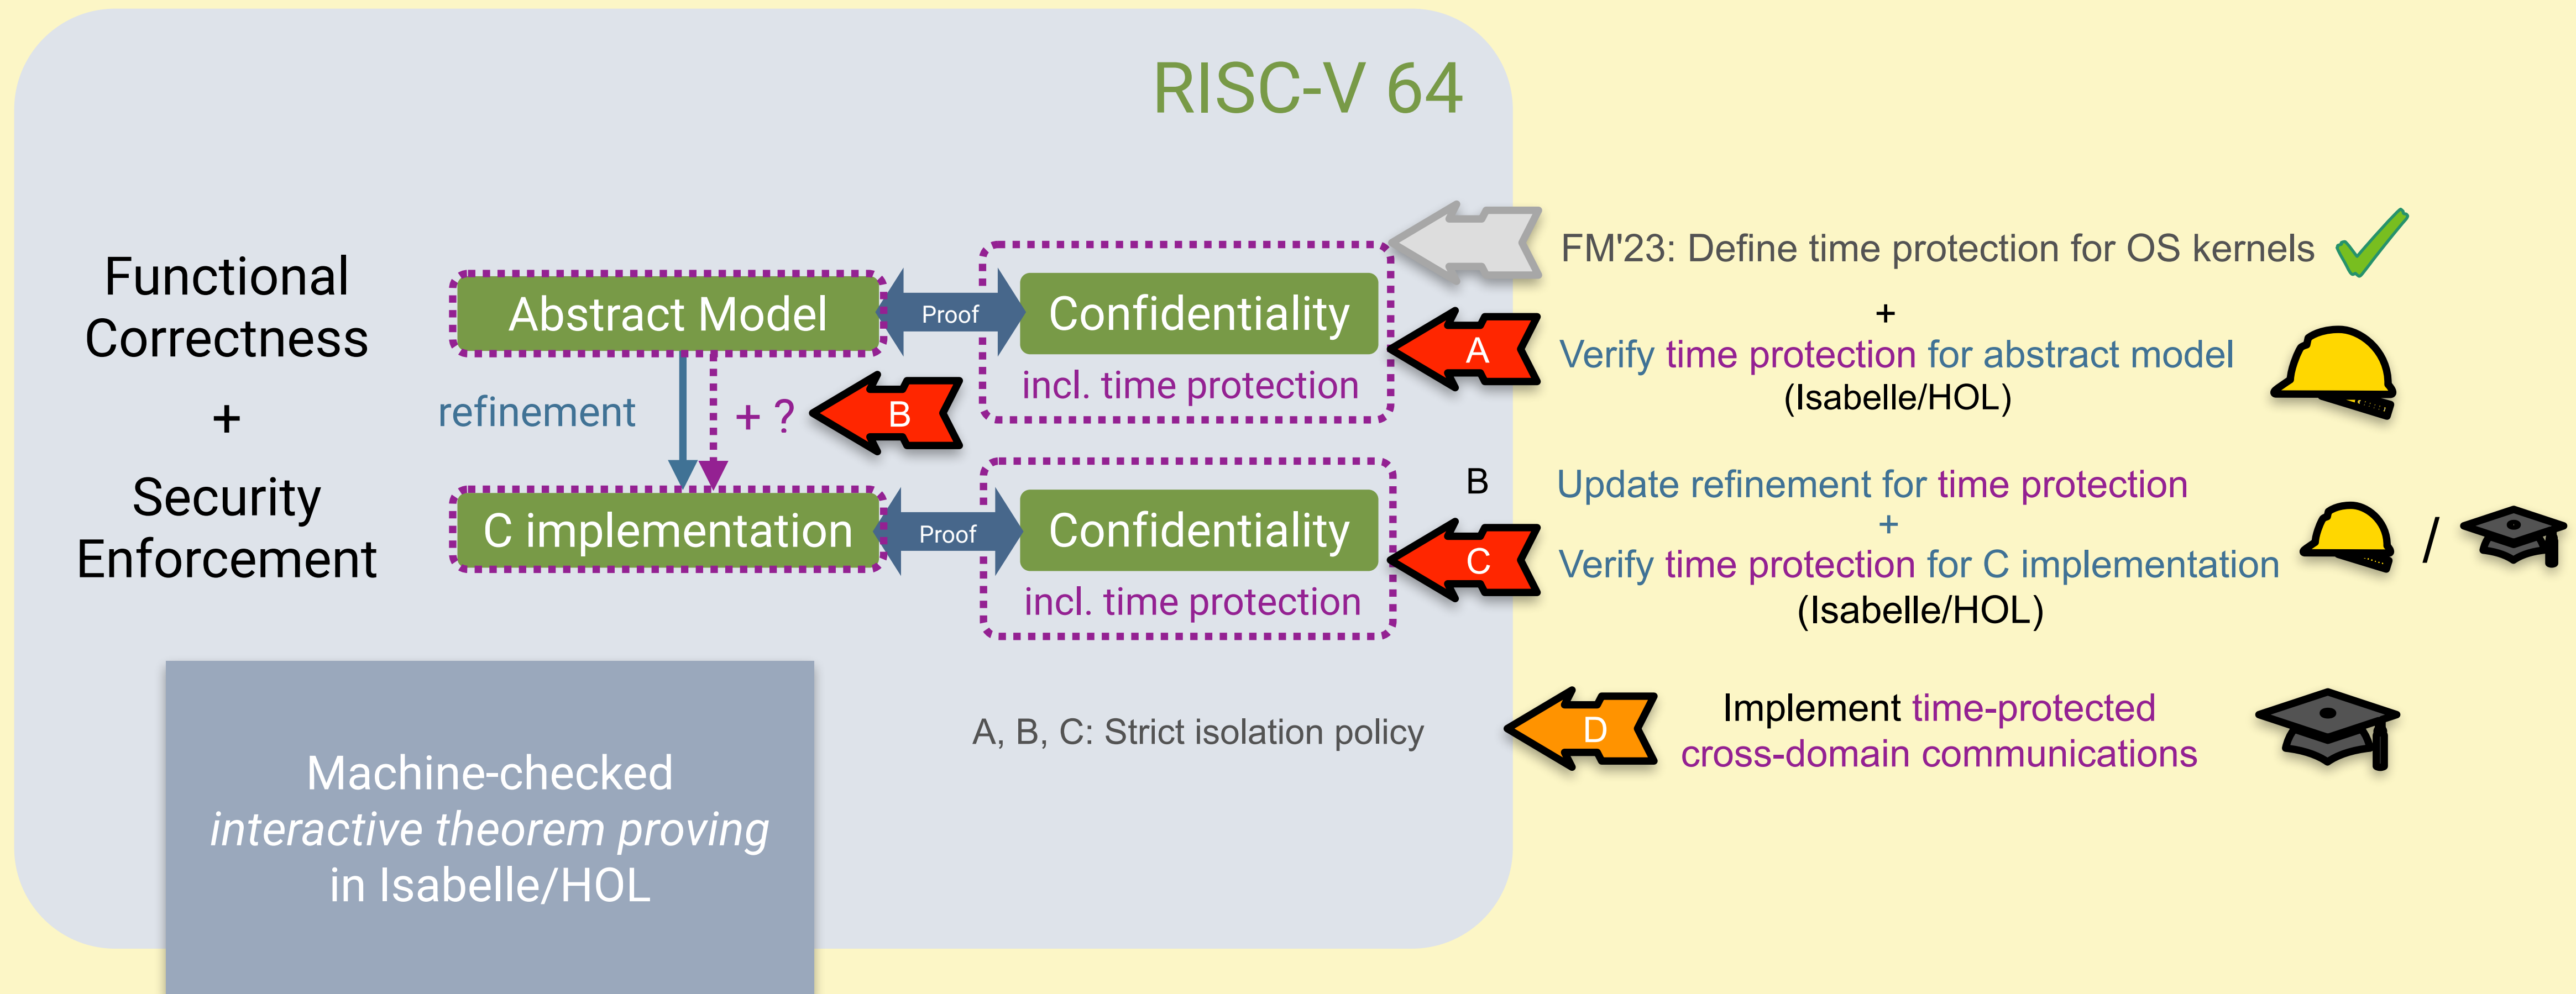
The plan 2 years ago... (presented at FM-OZ'24)

Thanks to funding from Australian Research Council



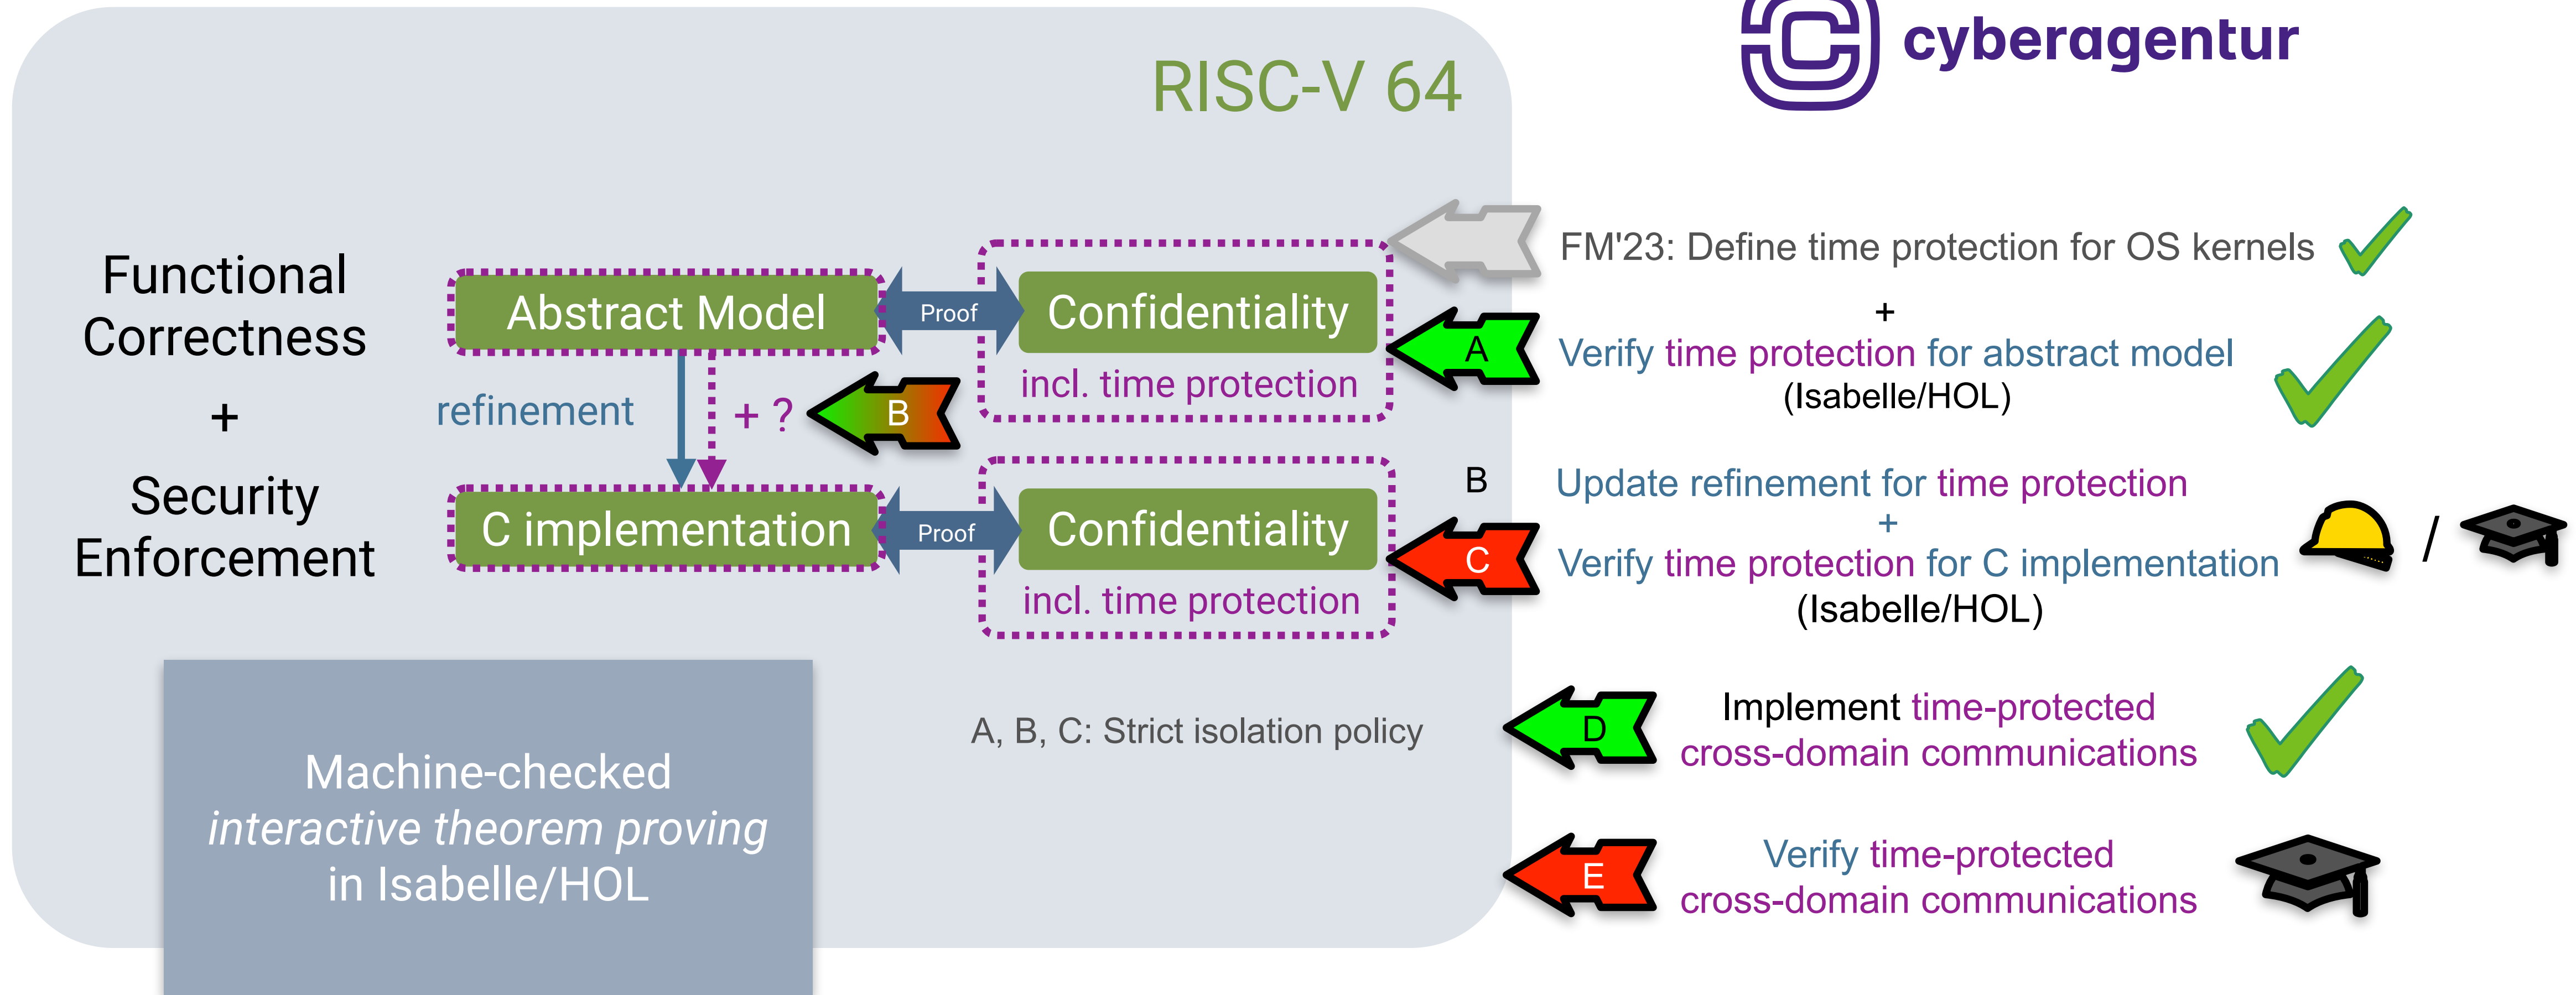
The plan 2 years ago... (presented at FM-OZ'24)

Thanks to funding from Australian Research Council



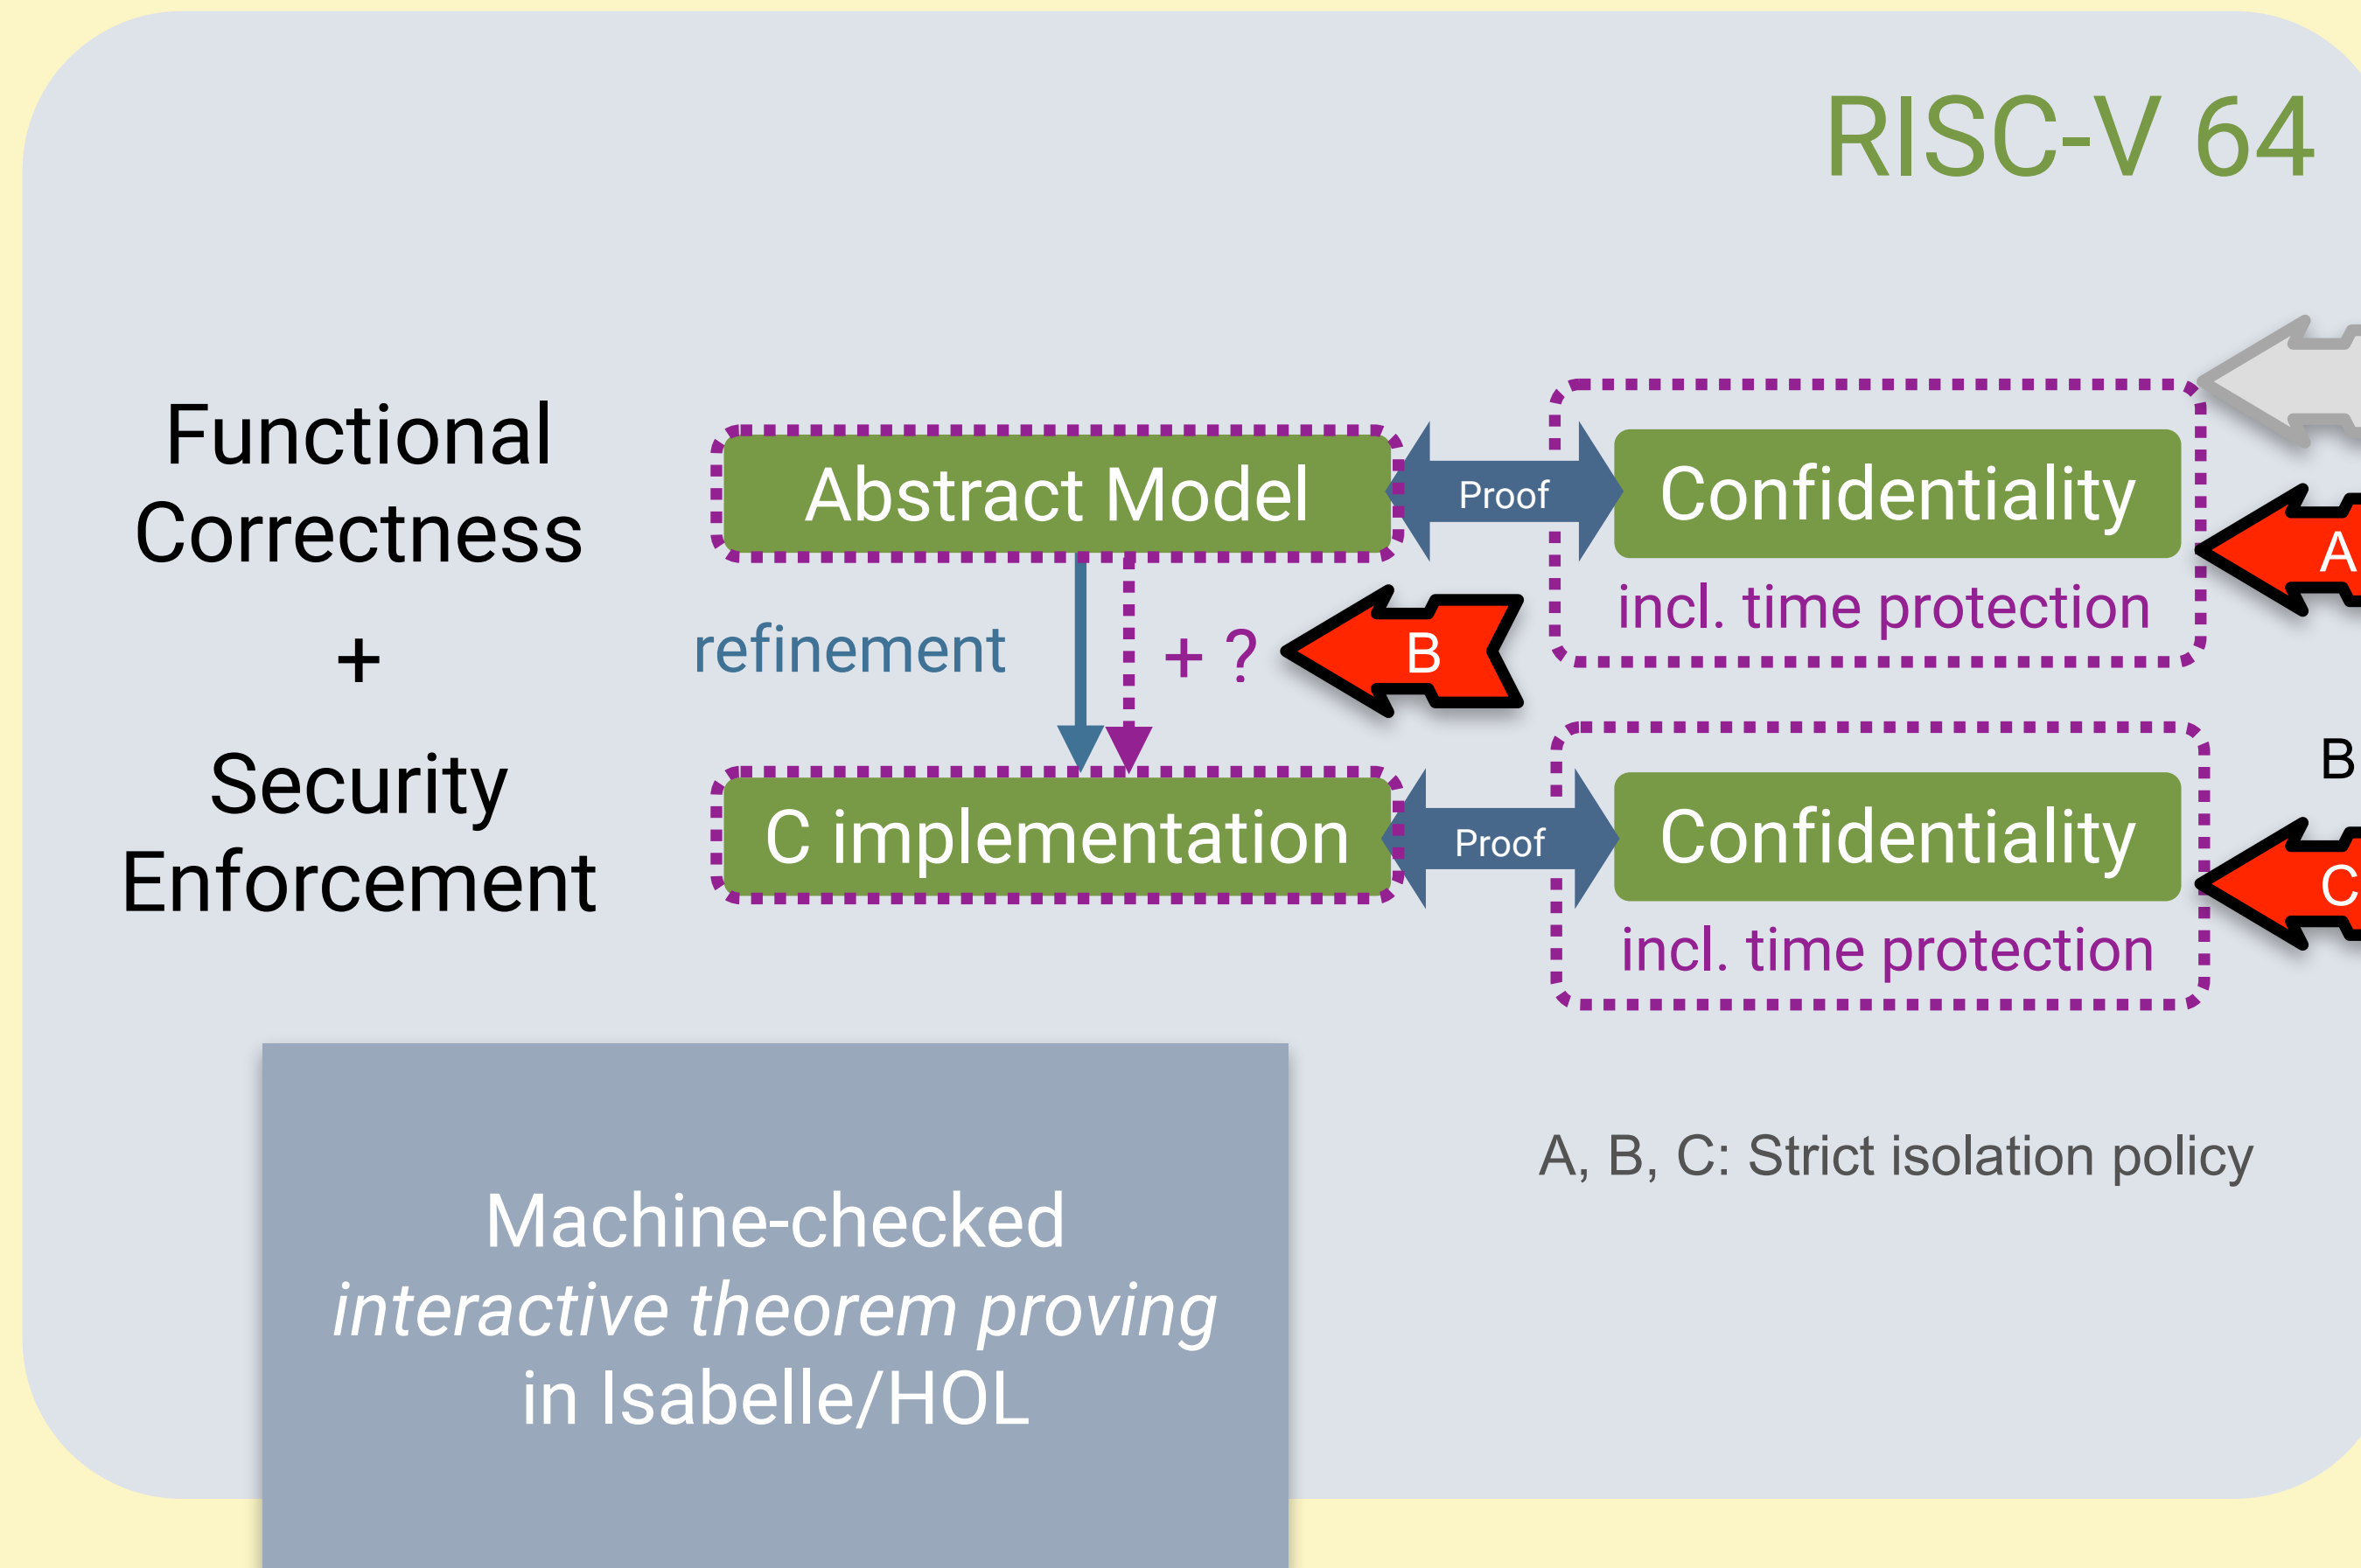
The status today:

Thanks to funding from



Since the plan 2 years ago...

Thanks to funding from



FM'23: Define time protection for OS kernels ✓

+
Verify time protection for abstract model (Isabelle/HOL)



B
Update refinement for time protection

+
Verify time protection for C implementation (Isabelle/HOL)



A, B, C: Strict isolation policy

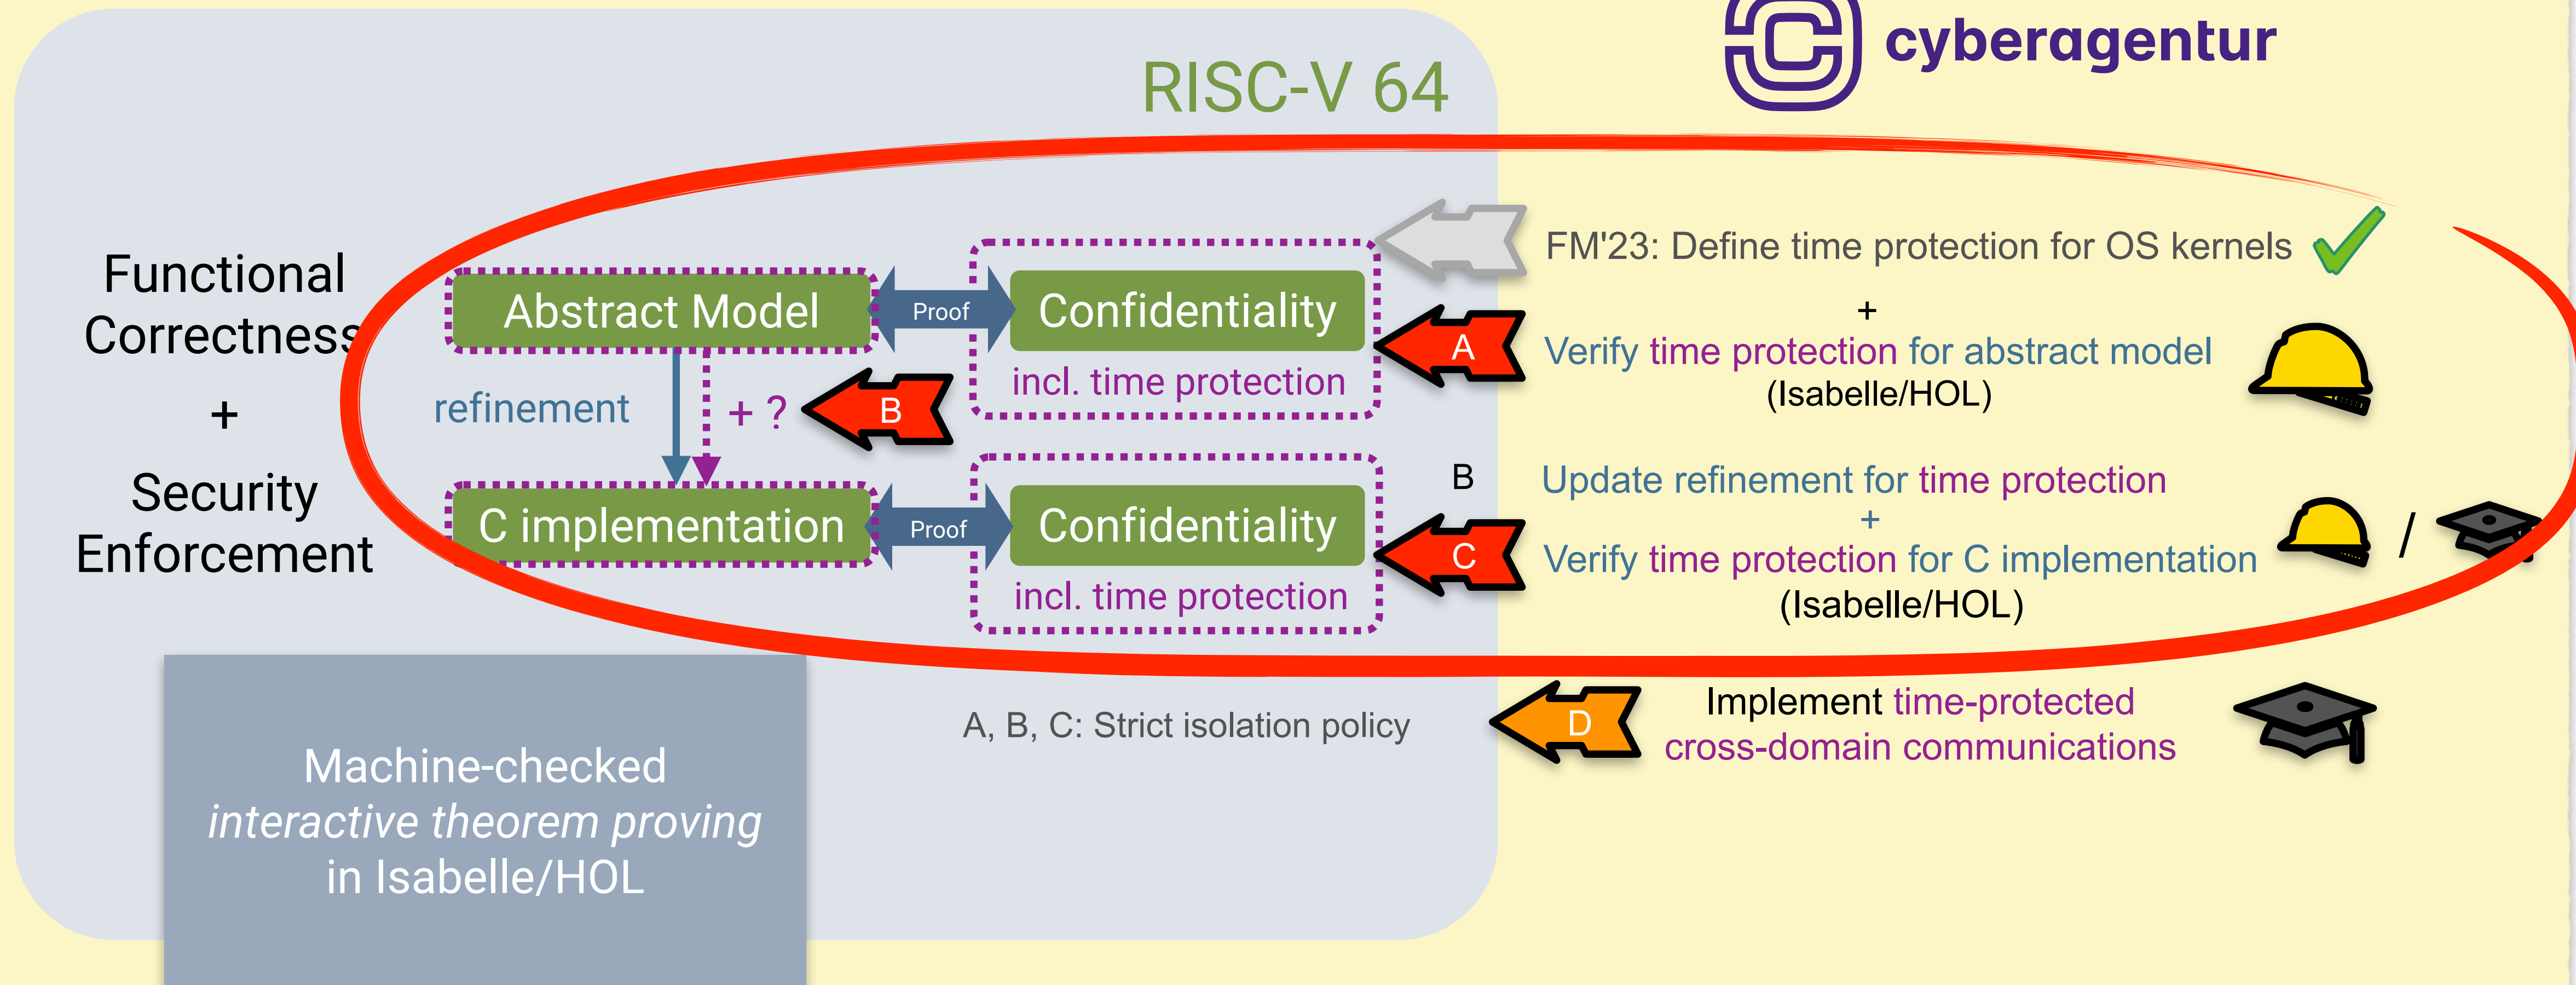


D
Implement time-protected cross-domain communications

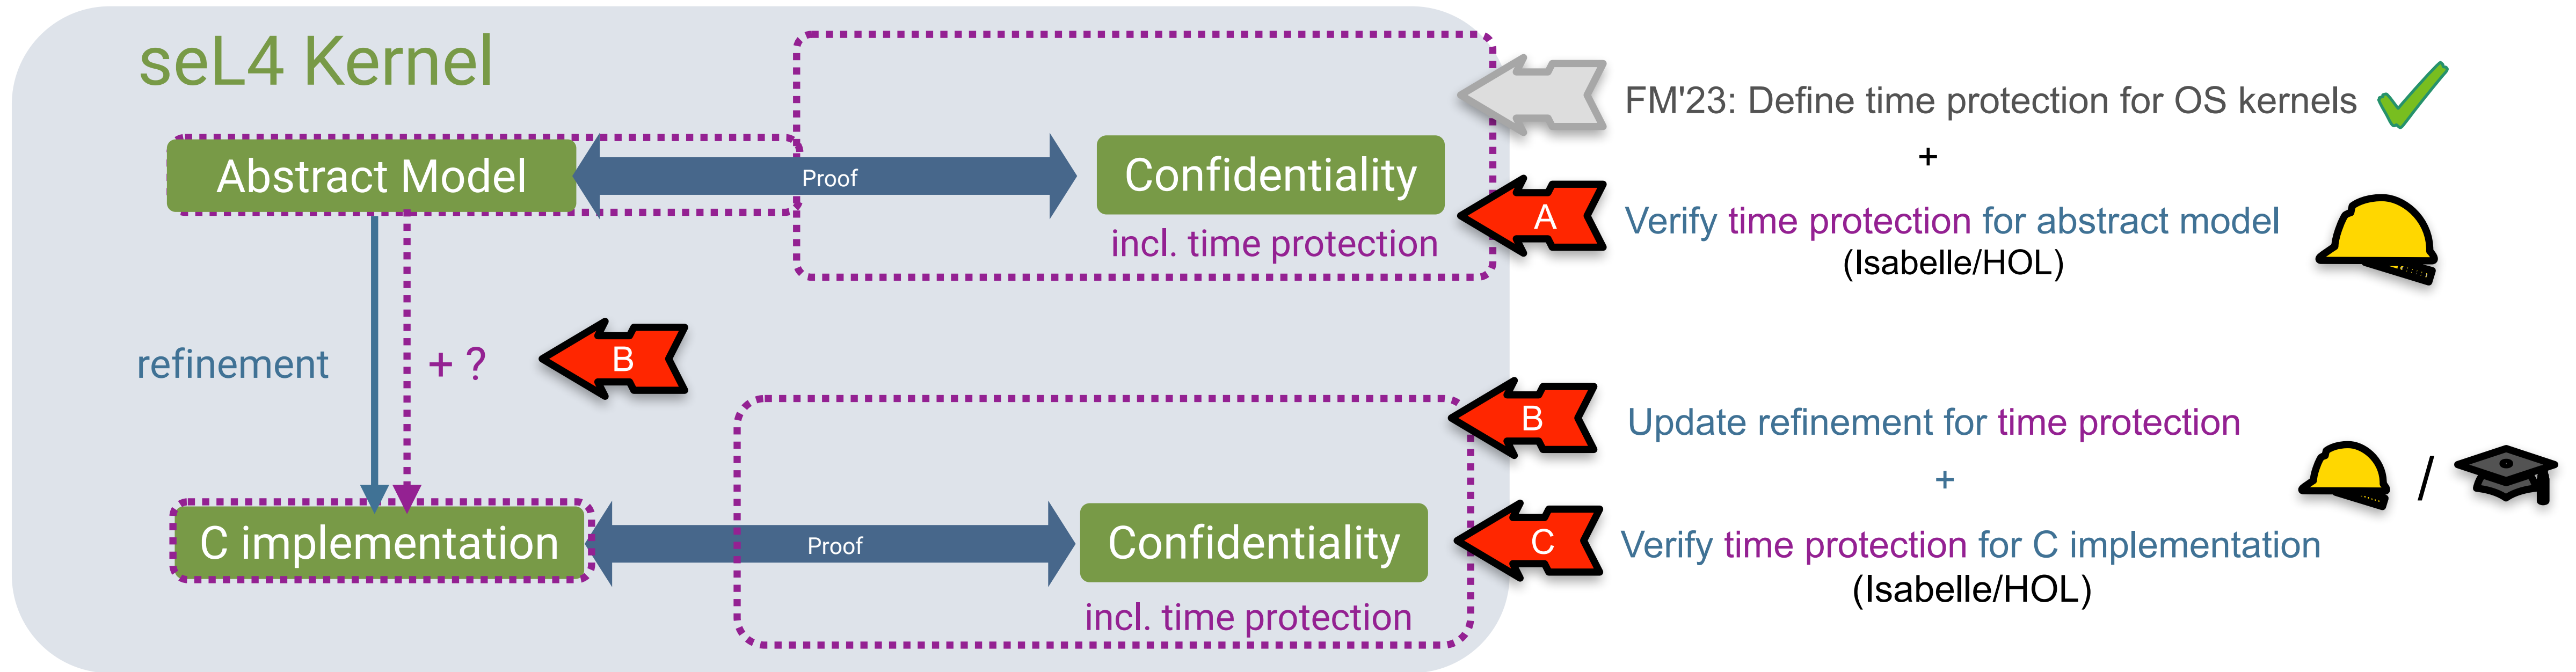


Since the plan 2 years ago...

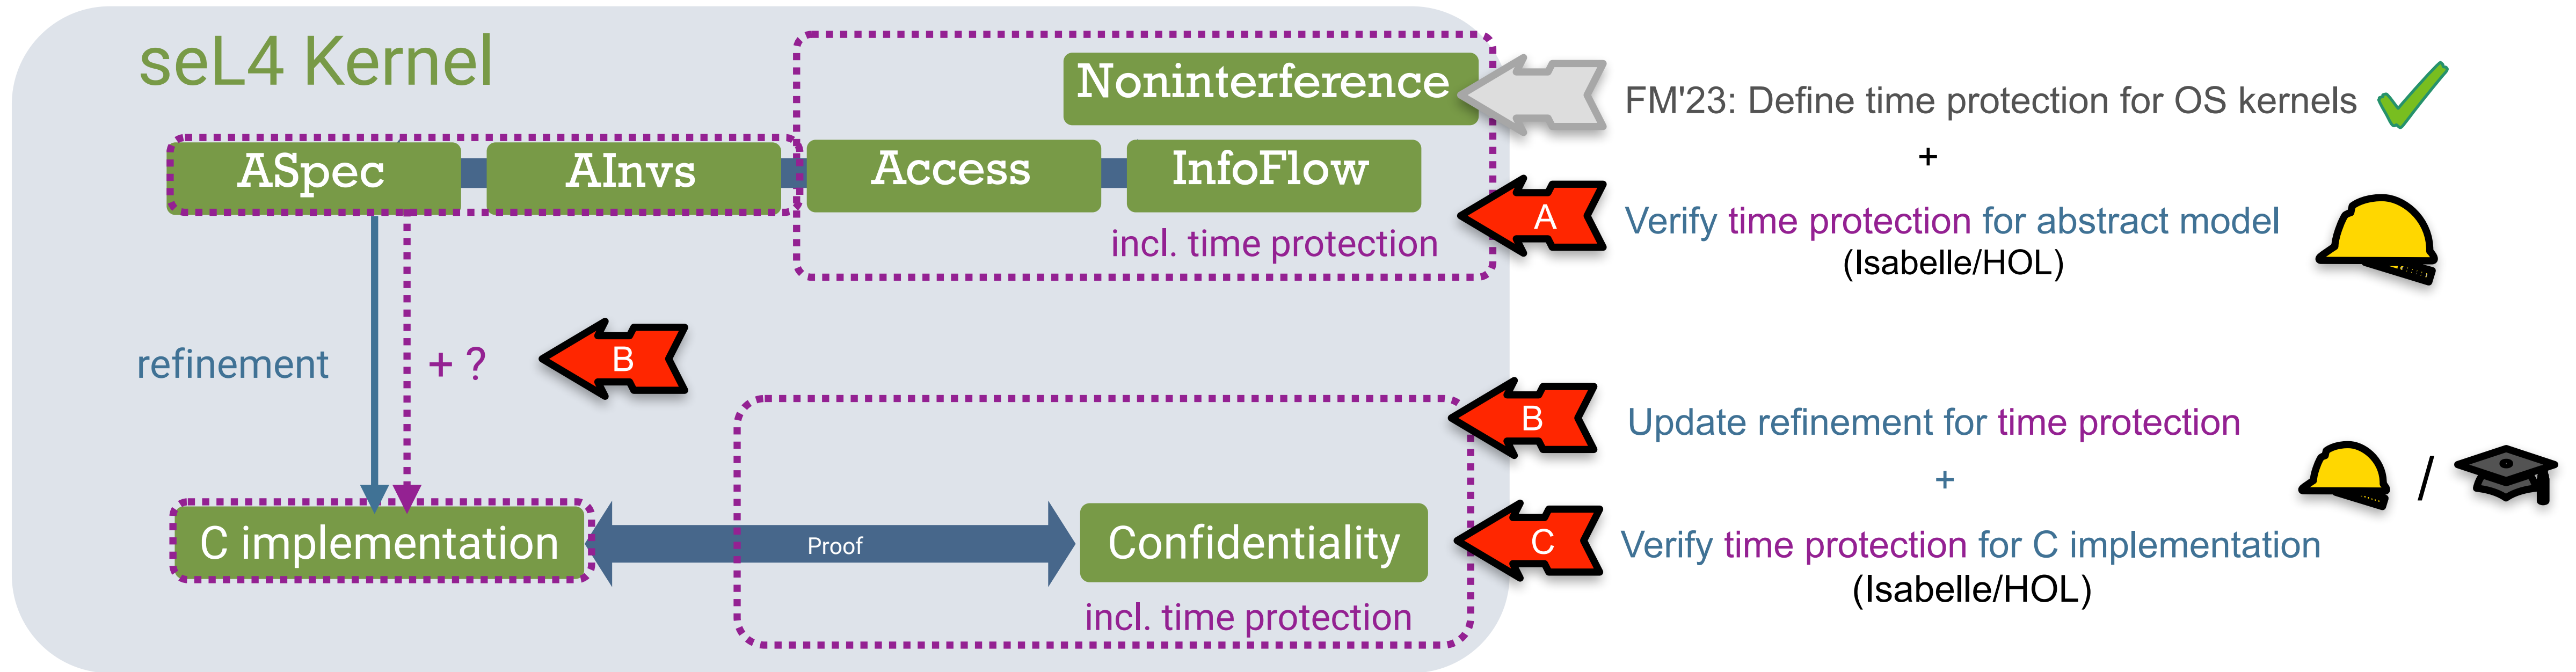
Thanks to funding from



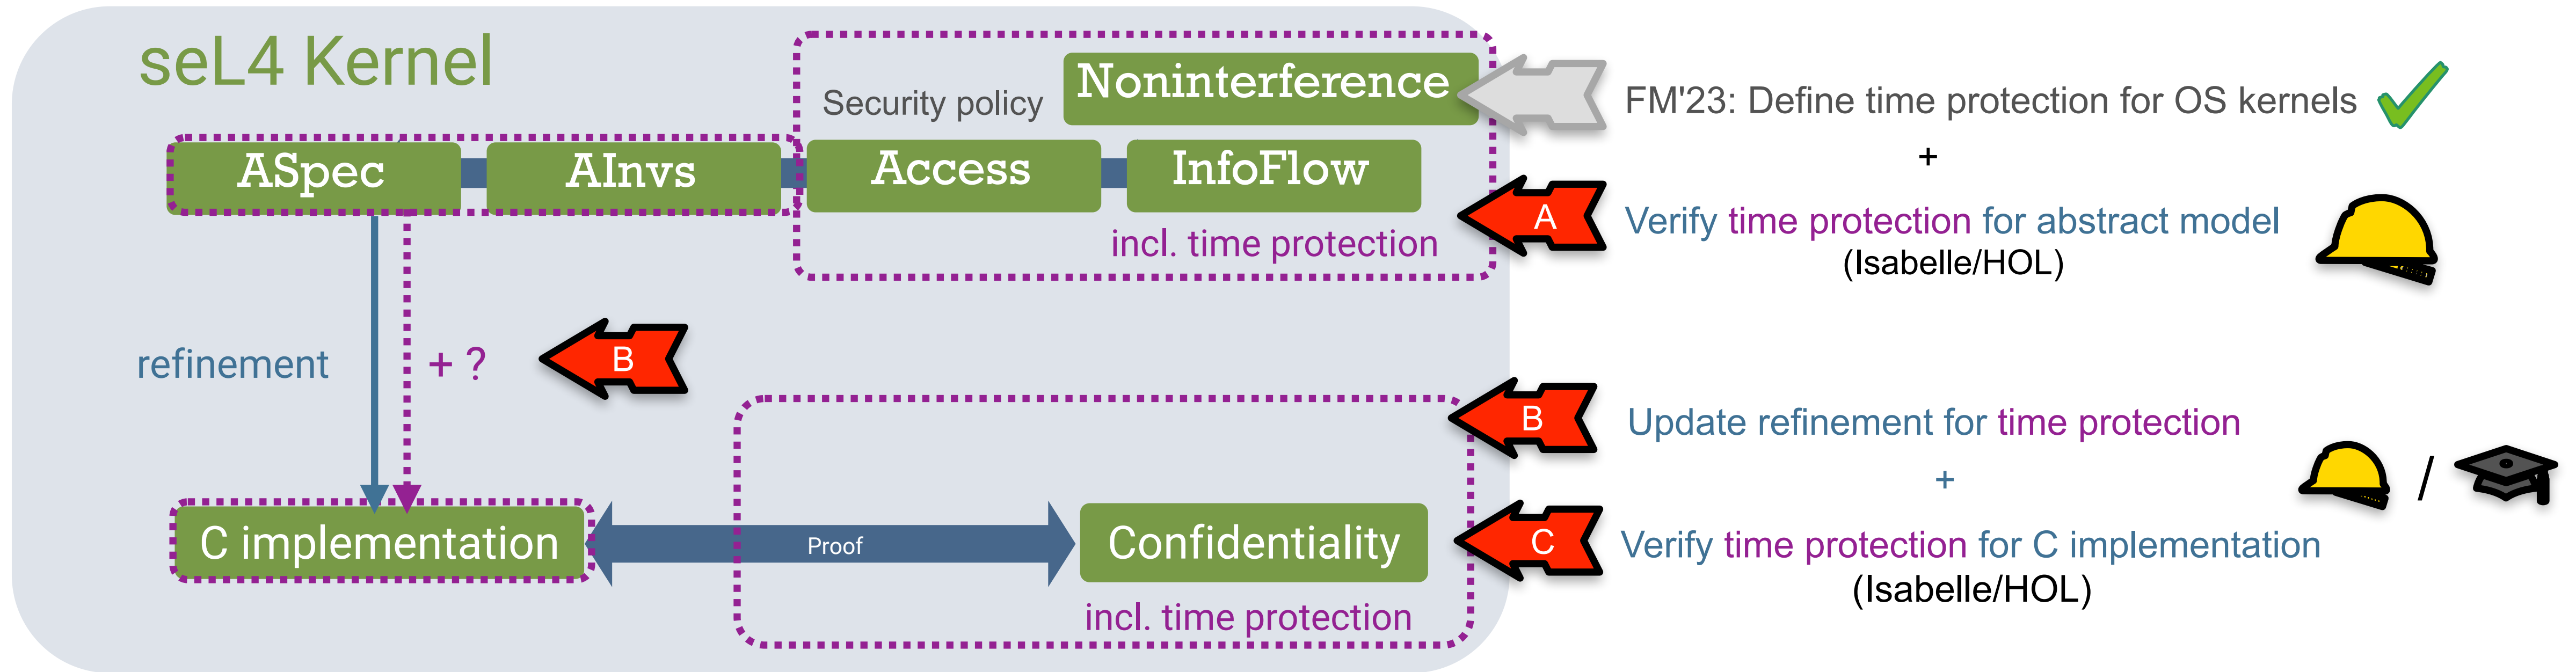
Originally planned approach... (FM-OZ'24)



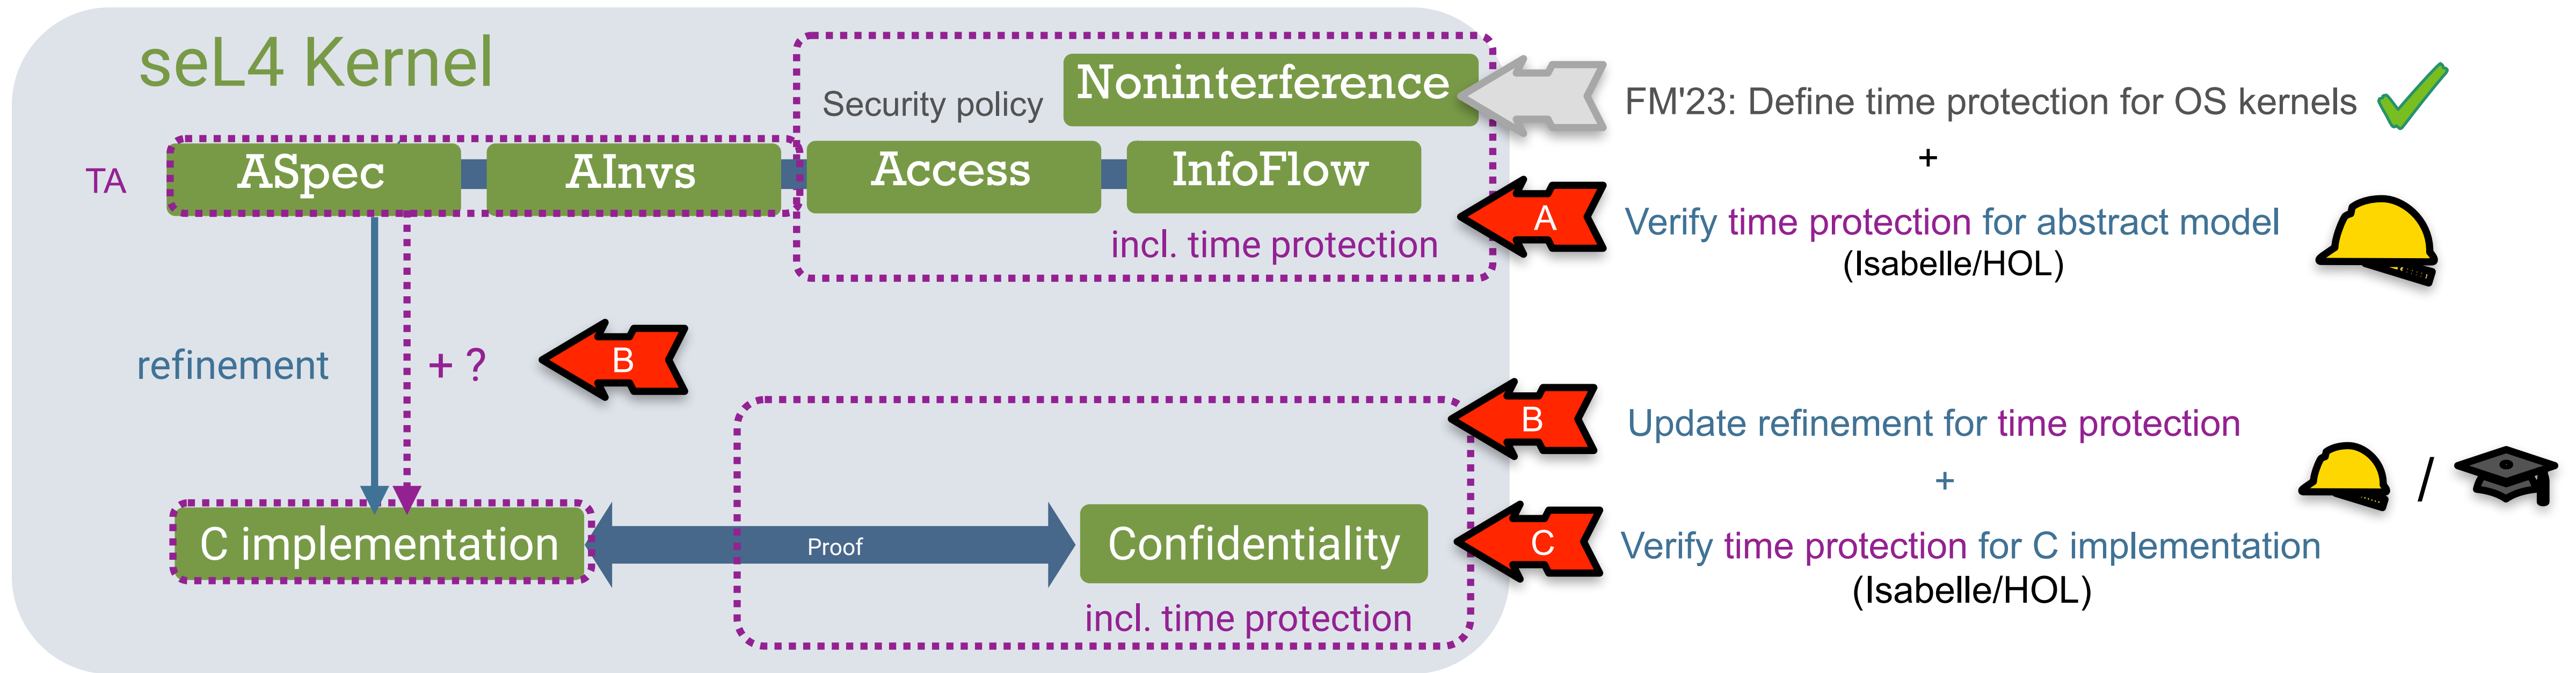
Originally planned approach... (FM-OZ'24)



Originally planned approach... (FM-OZ'24)

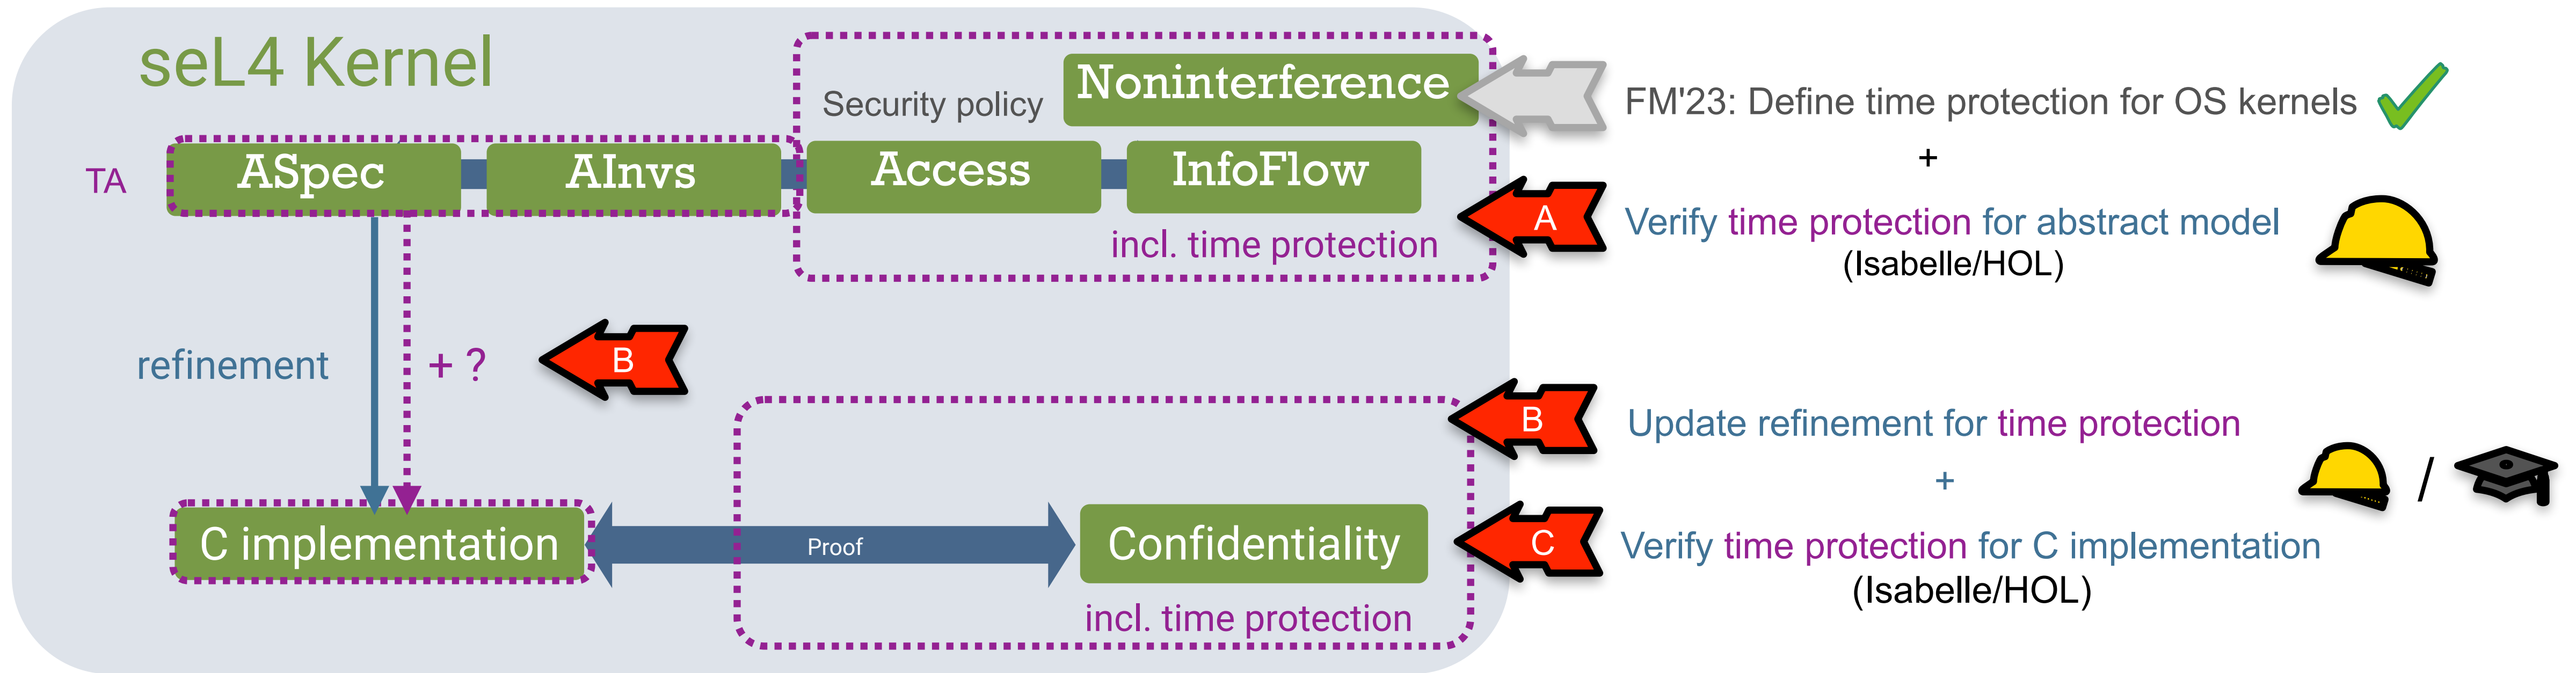


Originally planned approach... (FM-OZ'24)



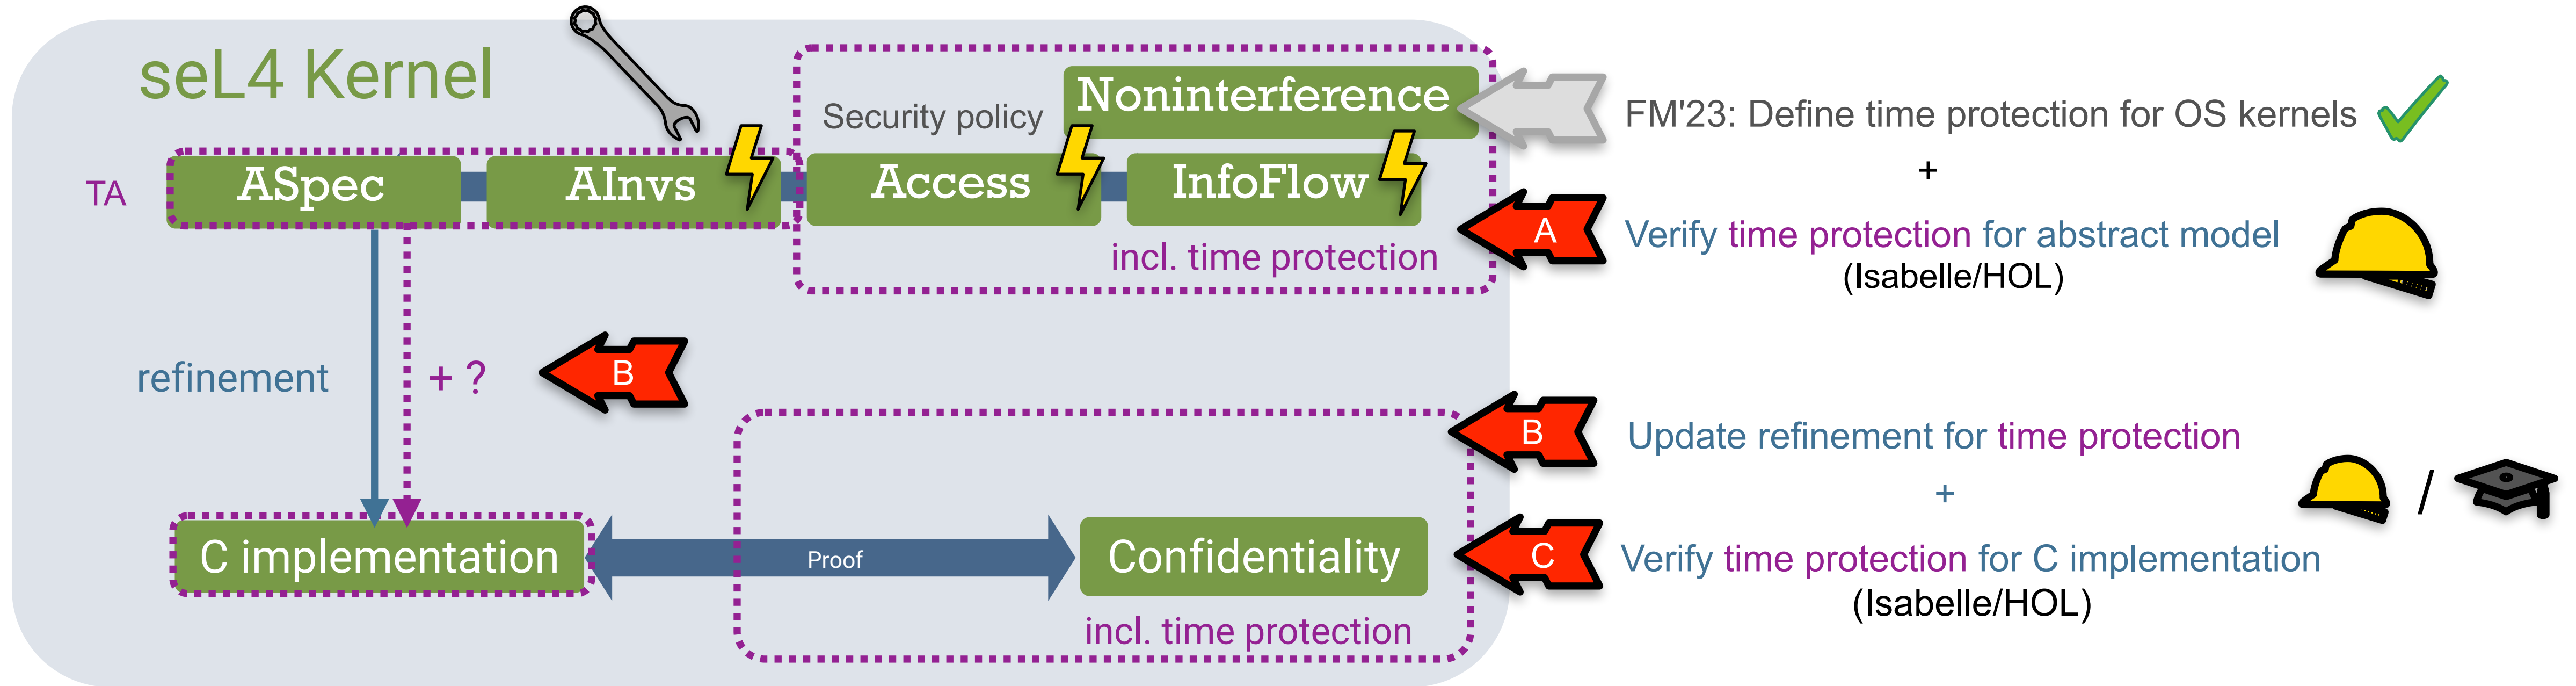
- Abstract model (ASpec) checks dynamic *touched addresses* (TA) set

Originally planned approach... (FM-OZ'24)



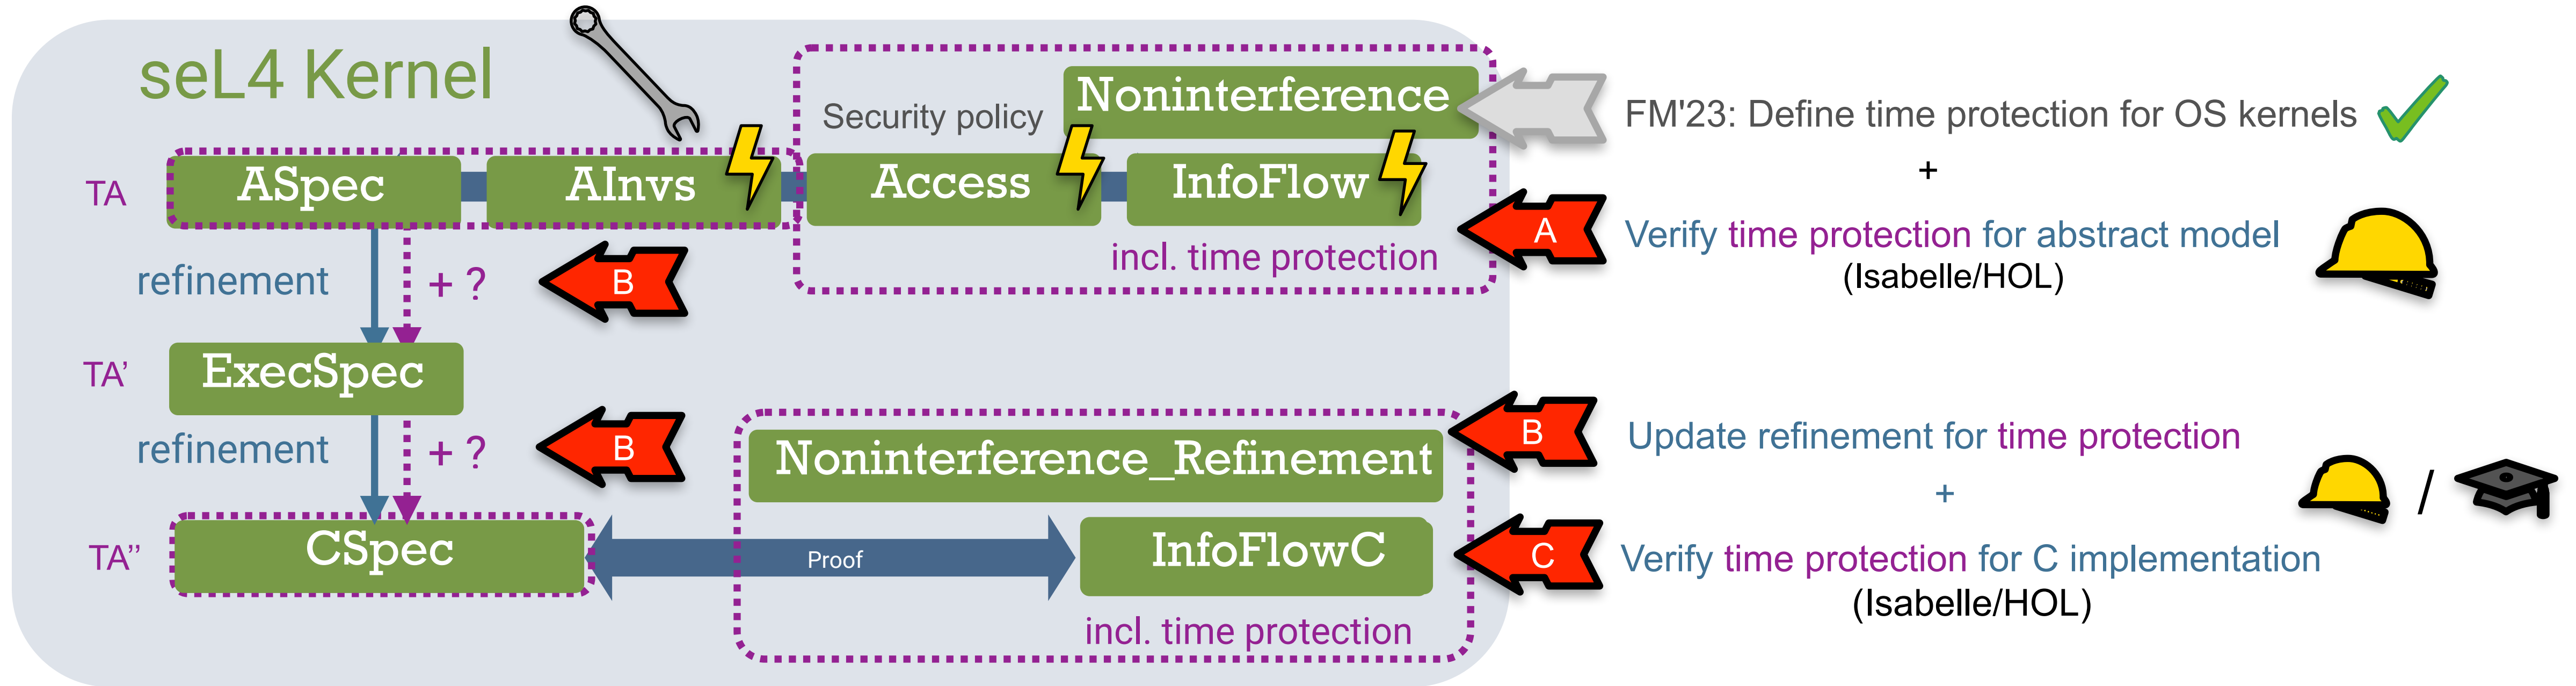
- Abstract model ($\overline{A}Spec$) checks dynamic *touched addresses* (TA) set
- (A) Key $\overline{A}Spec$ property: $TA \subseteq domain's\ addresses$ (according to *security policy*)

Originally planned approach... (FM-OZ'24)



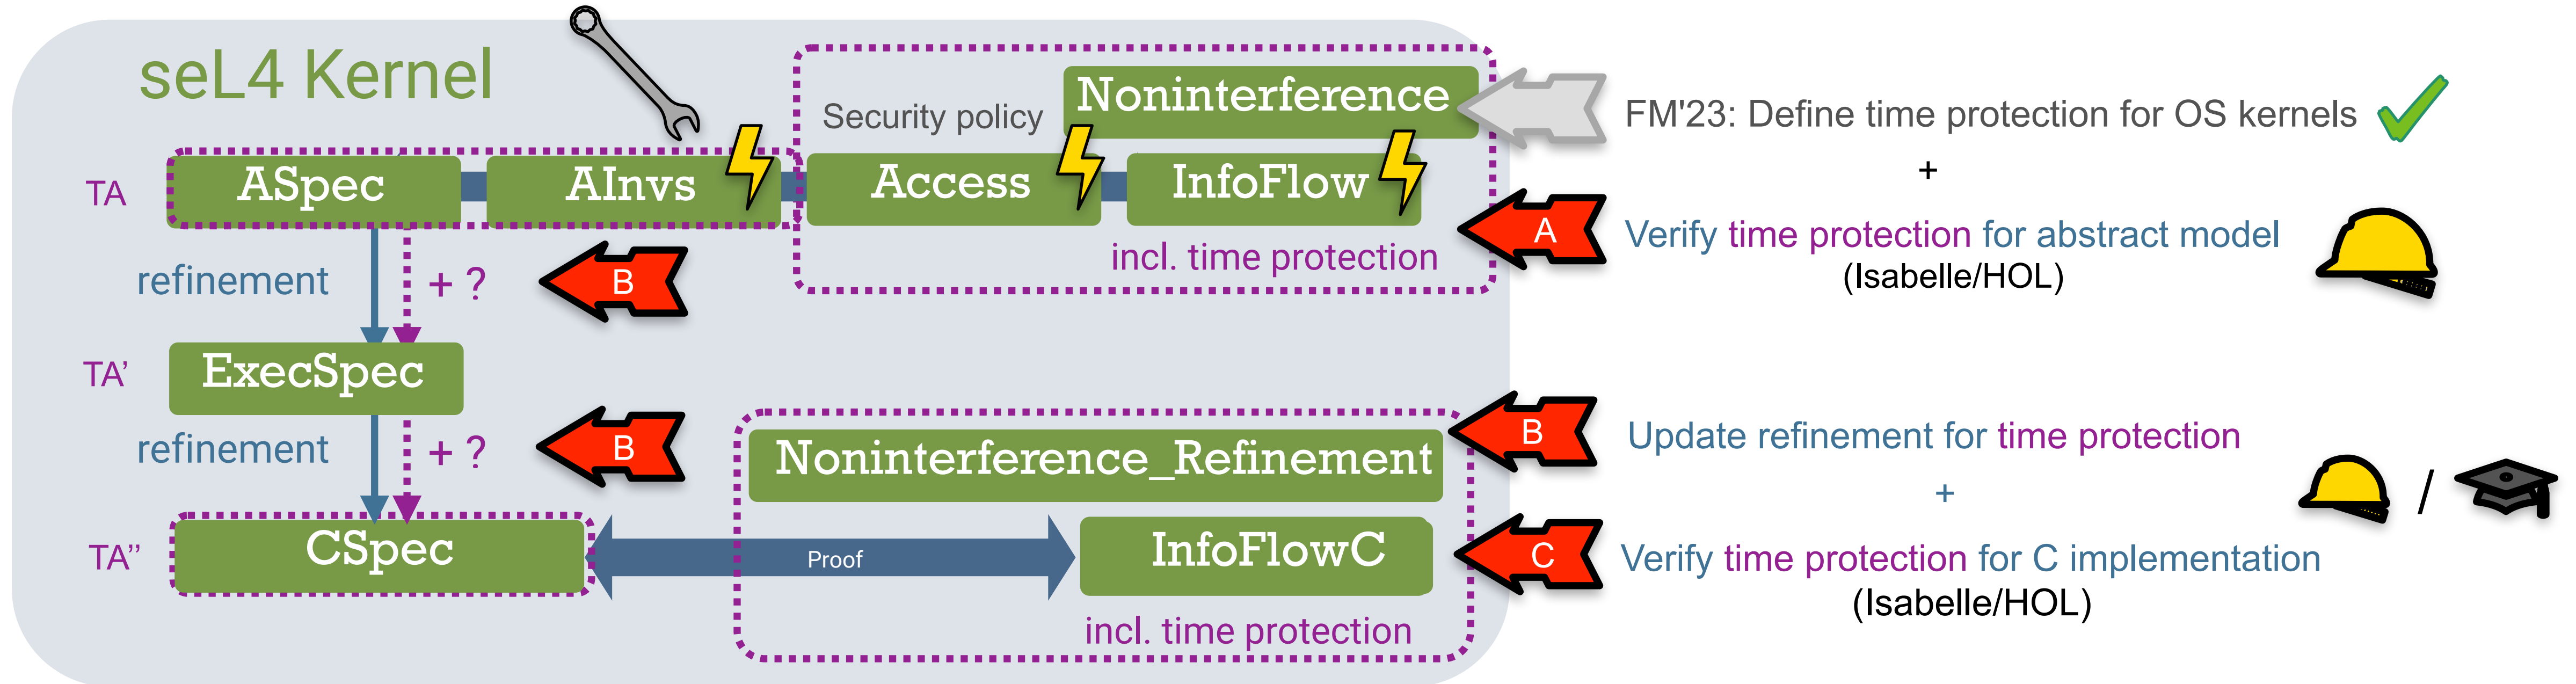
- Abstract model (**ASpec**) checks dynamic *touched addresses* (TA) set
- (A) Key **ASpec** property: $TA \subseteq \text{domain's addresses}$ (according to *security policy*)
 - Also: **AInvs**, **Access**, **InfoFlow** need repairs

Originally planned approach... (FM-OZ'24)



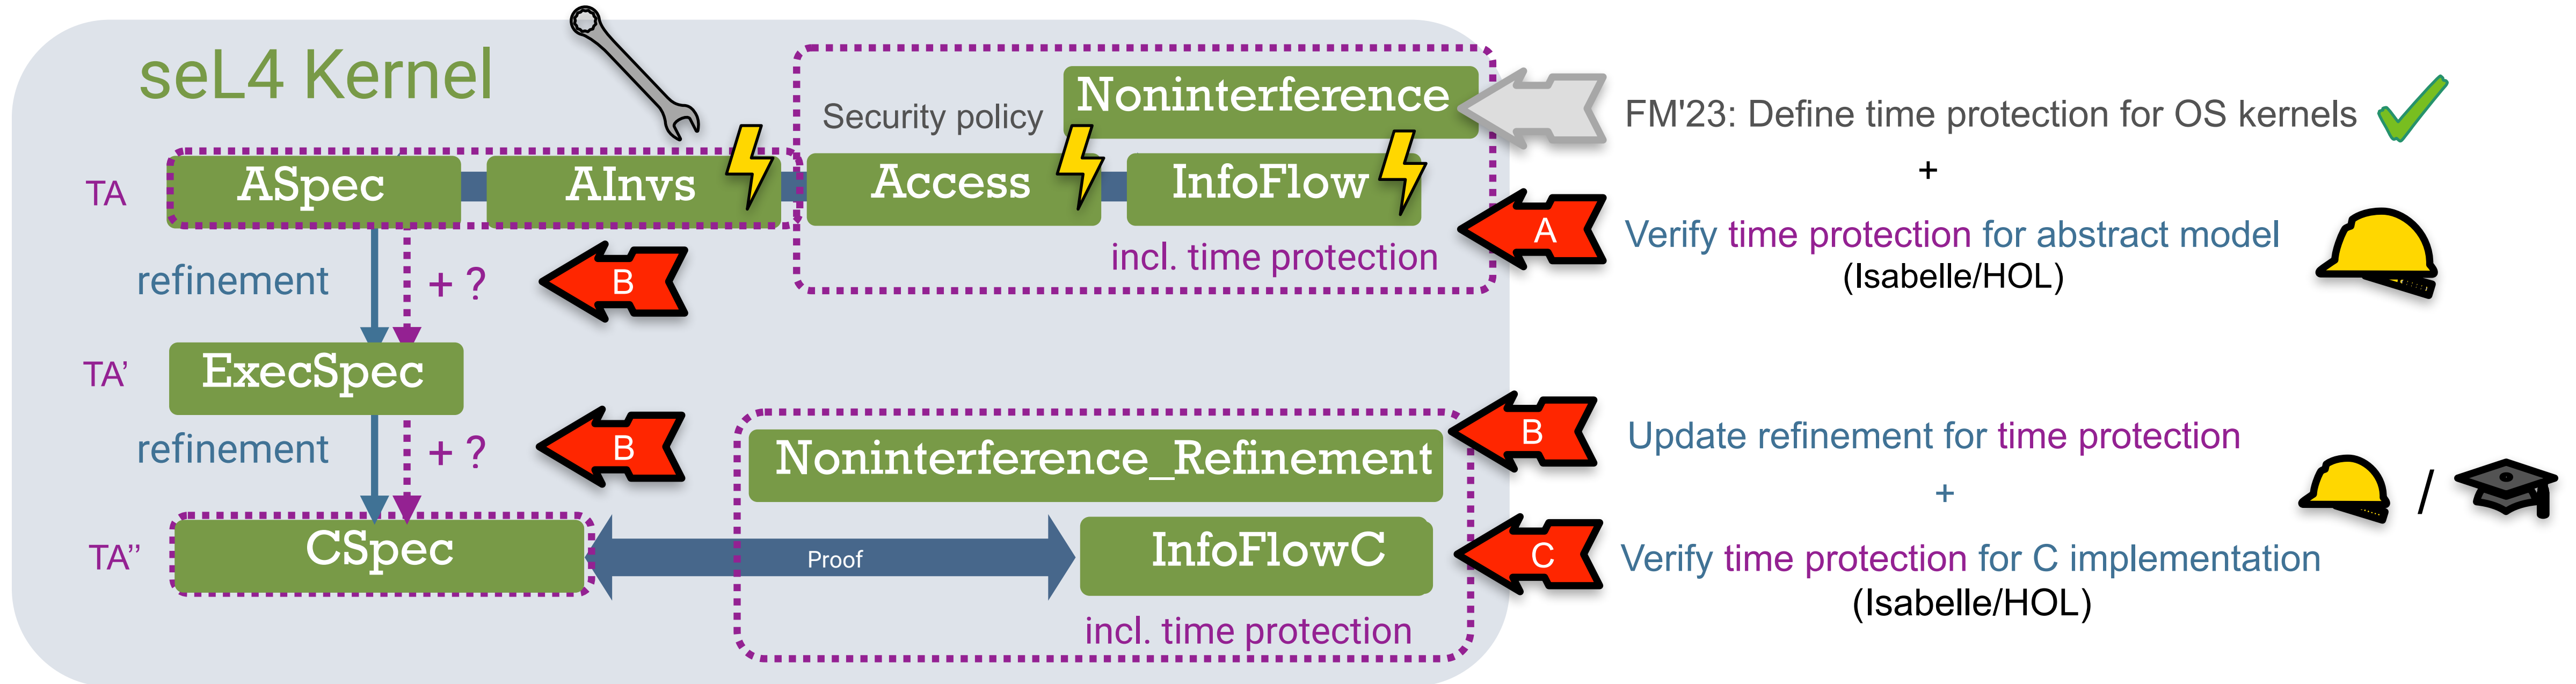
- Abstract model (**ASpec**) checks dynamic *touched addresses* (TA) set
- (A) Key **ASpec** property: $TA \subseteq \text{domain's addresses}$ (according to *security policy*)
 - Also: **AInvs**, **Access**, **InfoFlow** need repairs
- (B + C) Refinement: $TA'' \subseteq TA' \subseteq TA$? (via **ExecSpec**)

Originally planned approach... (FM-OZ'24)



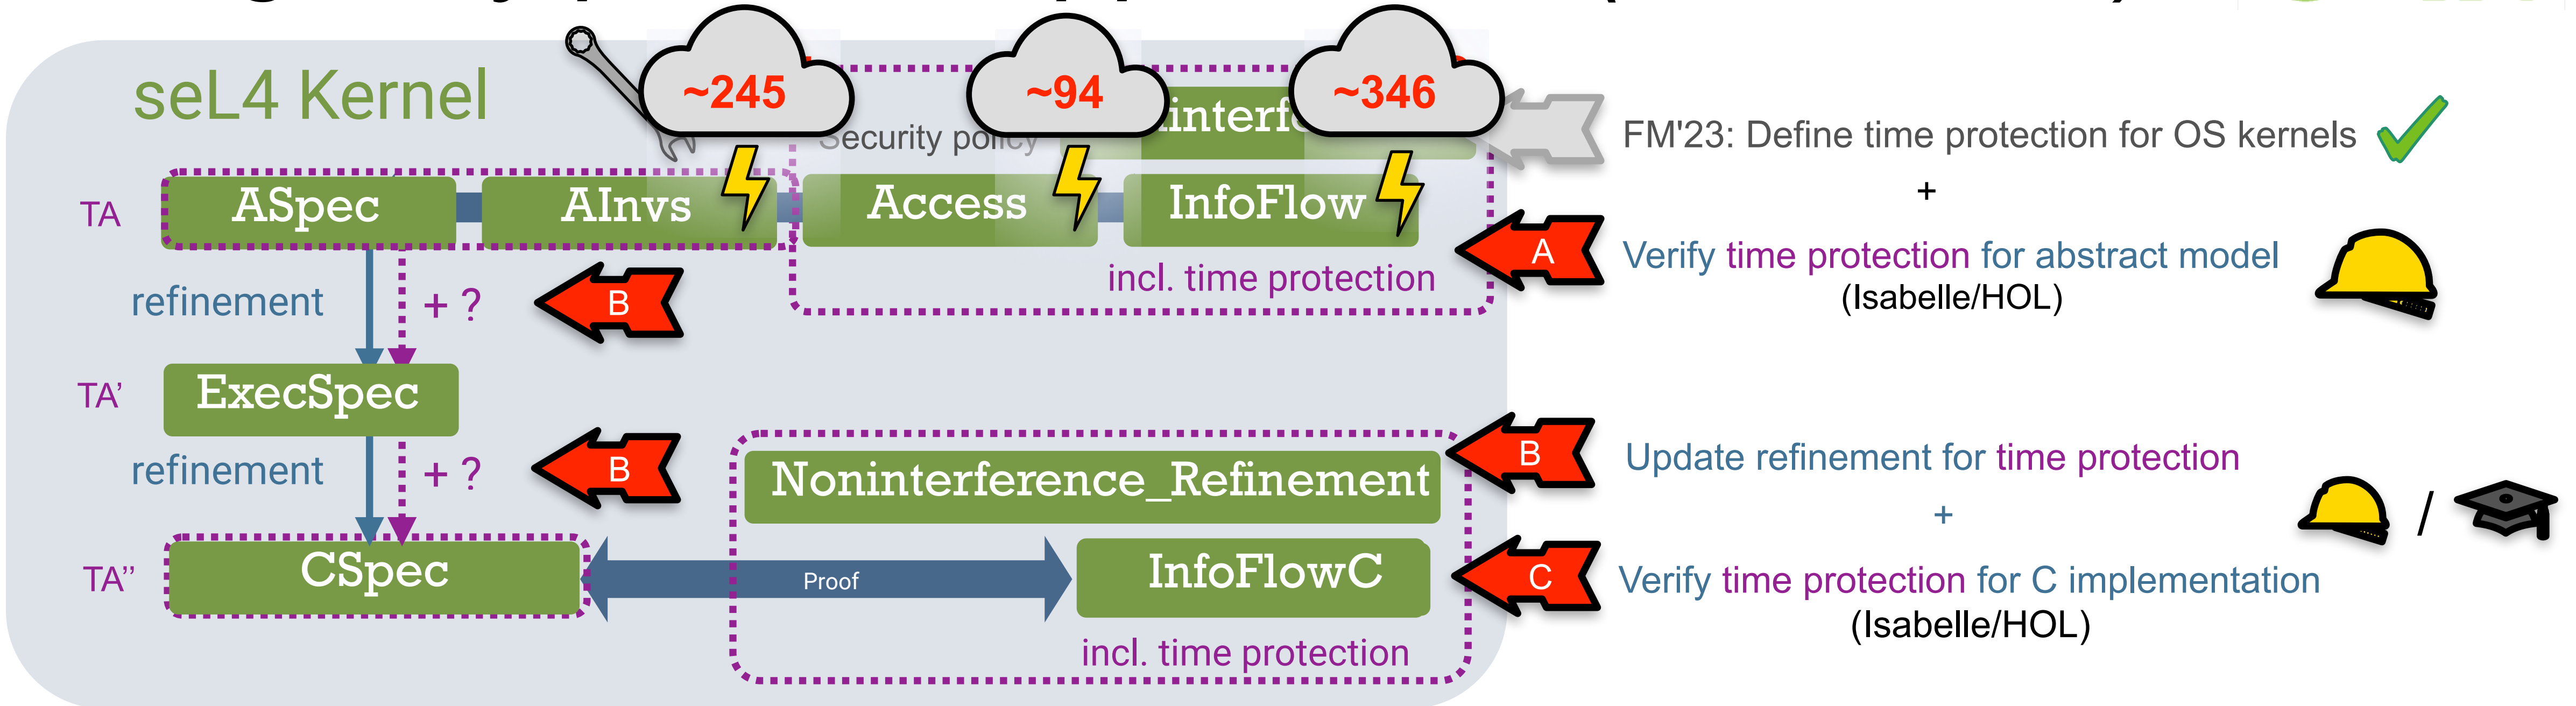
- Abstract model (**ASpec**) checks dynamic *touched addresses* (TA) set
- (A) Key **ASpec** property: $TA \subseteq \text{domain's addresses}$ (according to *security policy*)
 - Also: **AInvs**, **Access**, **InfoFlow** need repairs
- (B + C) Refinement: $TA'' \subseteq TA' \subseteq TA$? (via **ExecSpec**)
 - Tricky part #1: Refinement to **CSpec** accesses new addresses
 - Tricky part #2: Making TA checks scalable at **CSpec** level

Originally planned approach... (FM-OZ'24)



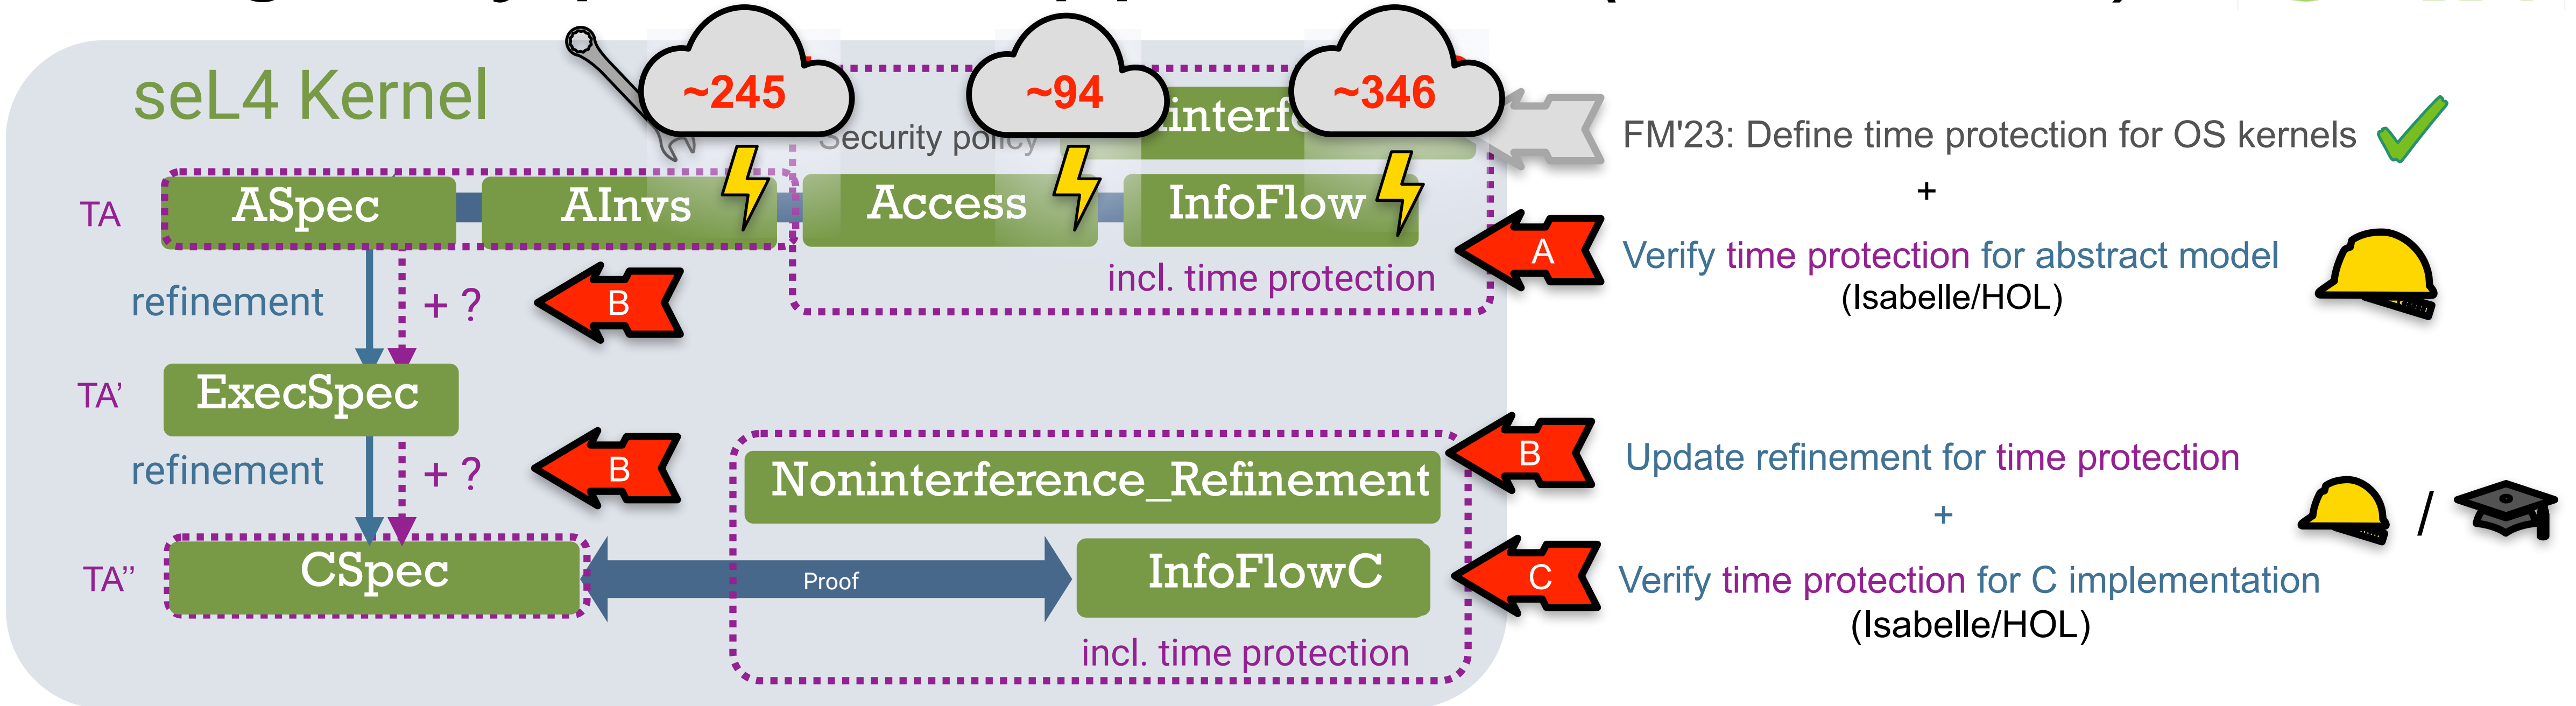
- Abstract model (**ASpec**) checks dynamic *touched addresses* (TA) set
- (A) Key **ASpec** property: $TA \subseteq \text{domain's addresses}$ (according to security policy)
 - Also: **AInvs**, **Access**, **InfoFlow** need repairs **How many breakages?**
- (B + C) Refinement: $TA'' \subseteq TA' \subseteq TA$? (via **ExecSpec**)
 - Tricky part #1: Refinement to **CSpec** accesses new addresses
 - Tricky part #2: Making TA checks scalable at **CSpec** level

Originally planned approach... (FM-OZ'24)



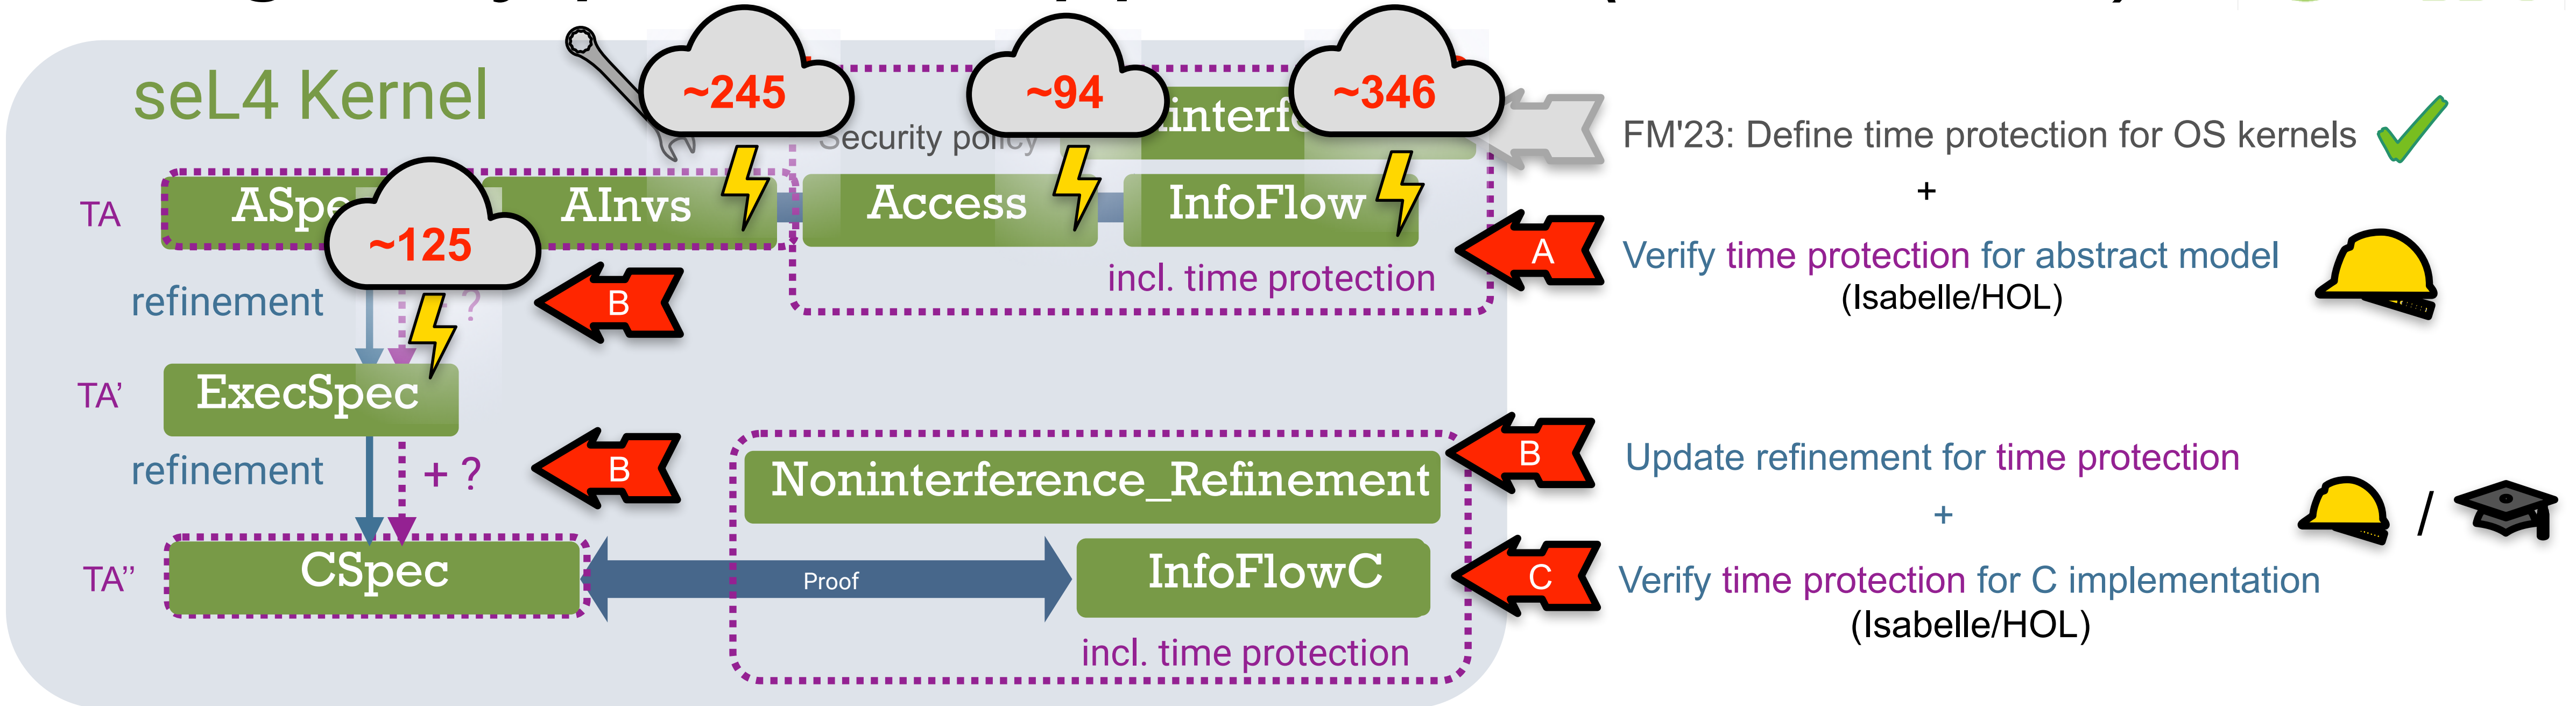
- Abstract model (**ASpec**) checks dynamic *touched addresses* (TA) set
- (A) Key **ASpec** property: $TA \subseteq \text{domain's addresses}$ (according to *security policy*)
 - Also: **AInvs**, **Access**, **InfoFlow** need repairs **How many breakages?** **> 685, after the ones we already fixed**
- (B + C) Refinement: $TA'' \subseteq TA' \subseteq TA$? (via **ExecSpec**)
 - Tricky part #1: Refinement to **CSpec** accesses new addresses
 - Tricky part #2: Making TA checks scalable at **CSpec** level

Originally planned approach... (FM-OZ'24)



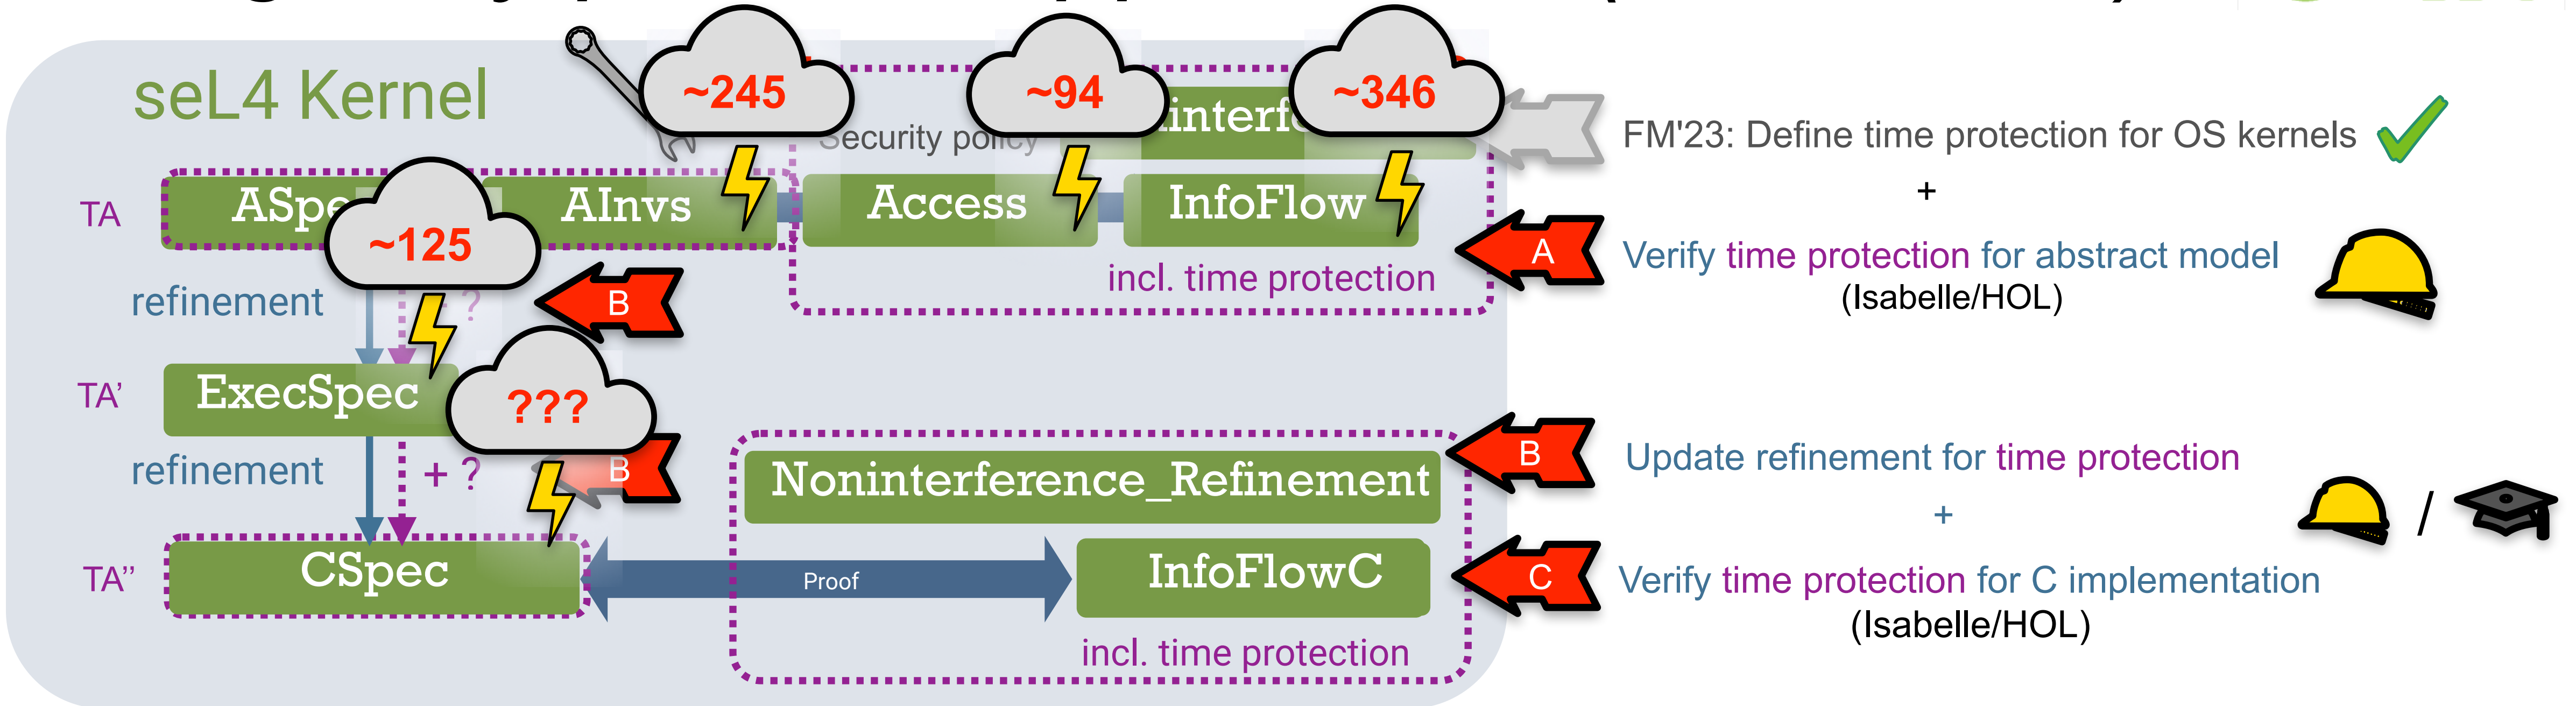
- Abstract model (ASpec) checks dynamic touched addresses (TA) set \rightarrow 125 times
- (A) Key ASpec property: $TA \subseteq \text{domain's addresses}$ (according to security policy)
 - Also: AInvs, Access, InfoFlow need repairs How many breakages? \rightarrow > 685, after the ones we already fixed
- (B + C) Refinement: $TA'' \subseteq TA' \subseteq TA$? (via ExecSpec)
 - Tricky part #1: Refinement to CSpec accesses new addresses
 - Tricky part #2: Making TA checks scalable at CSpec level

Originally planned approach... (FM-OZ'24)



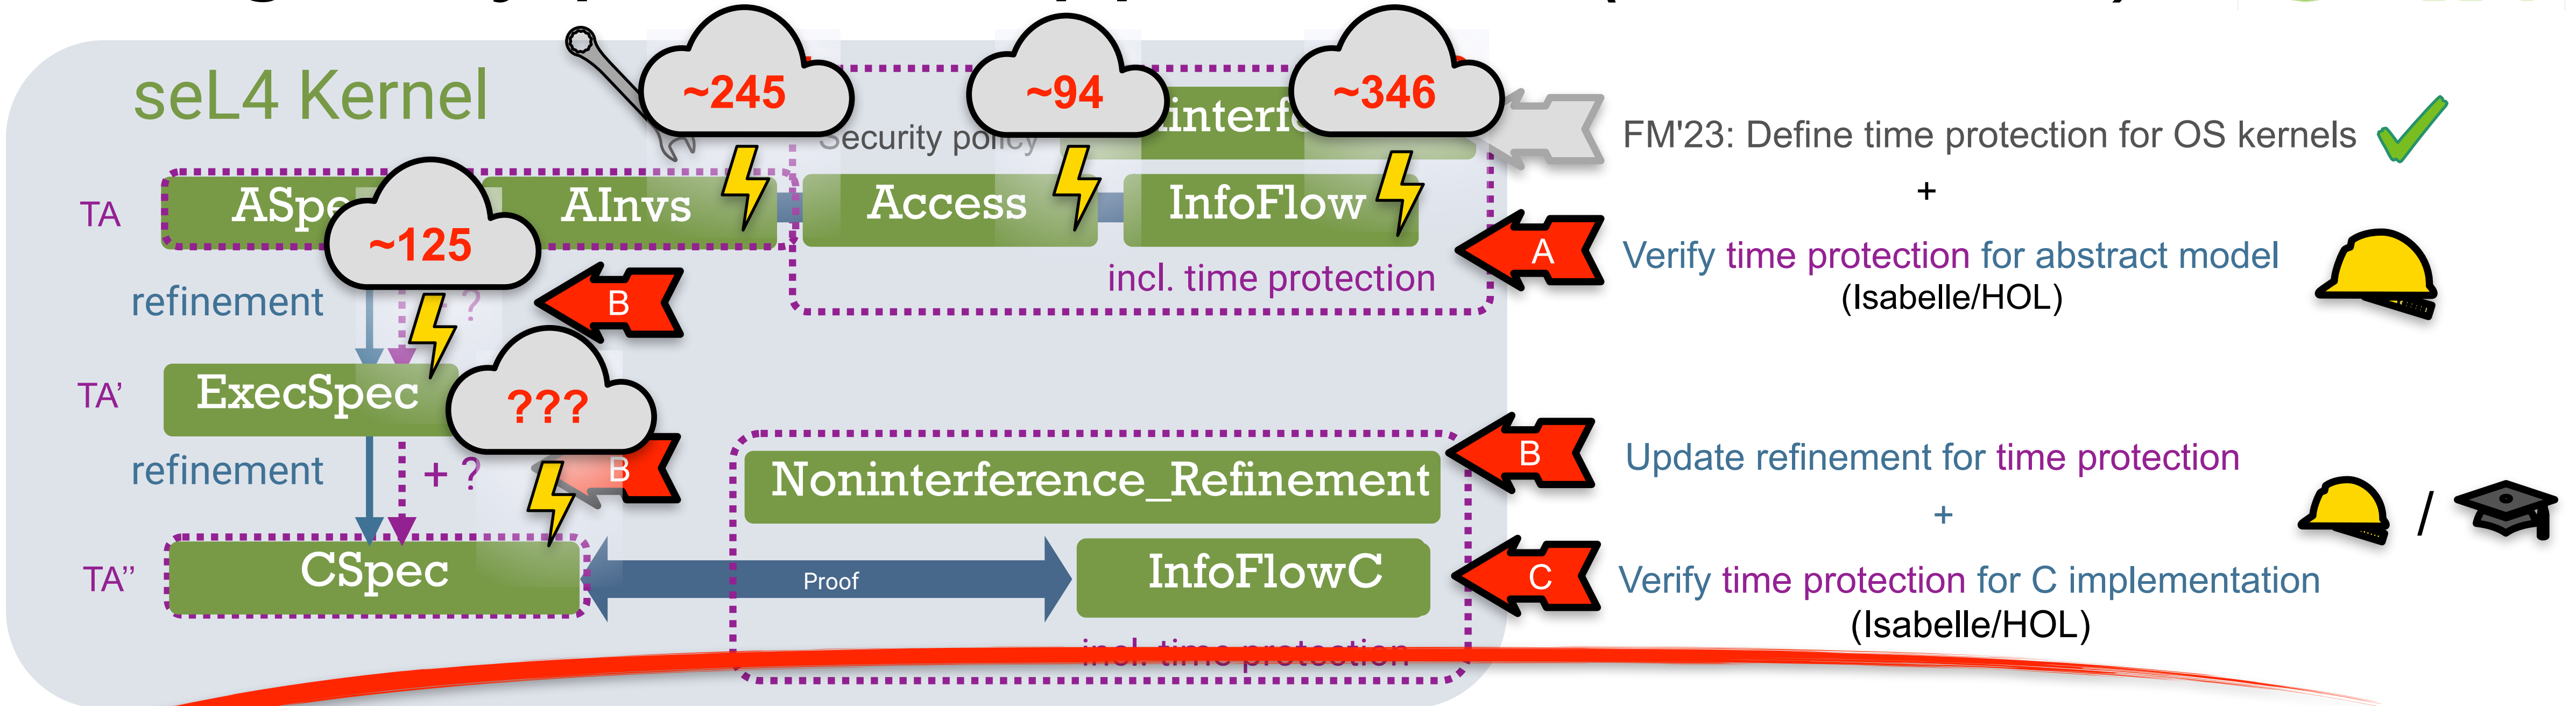
- Abstract model (ASpec) checks dynamic touched addresses (TA) set → 125 times
- (A) Key ASpec property: $TA \subseteq \text{domain's addresses}$ (according to security policy)
 - Also: AInvs, Access, InfoFlow need repairs How many breakages? → > 685, after the ones we already fixed
- (B + C) Refinement: $TA'' \subseteq TA' \subseteq TA$? (via ExecSpec)
 - Tricky part #1: Refinement to CSpec accesses new addresses
 - Tricky part #2: Making TA checks scalable at CSpec level

Originally planned approach... (FM-OZ'24)



- Abstract model (ASpec) checks dynamic touched addresses (TA) set → 125 times
- (A) Key ASpec property: $TA \subseteq \text{domain's addresses}$ (according to security policy)
 - Also: AInvs, Access, InfoFlow need repairs How many breakages? → > 685, after the ones we already fixed
- (B + C) Refinement: $TA'' \subseteq TA' \subseteq TA$? (via ExecSpec)
 - Tricky part #1: Refinement to CSPEC accesses new addresses
 - Tricky part #2: Making TA checks scalable at CSPEC level

Originally planned approach... (FM-OZ'24)



FM'23: Define time protection for OS kernels ✓

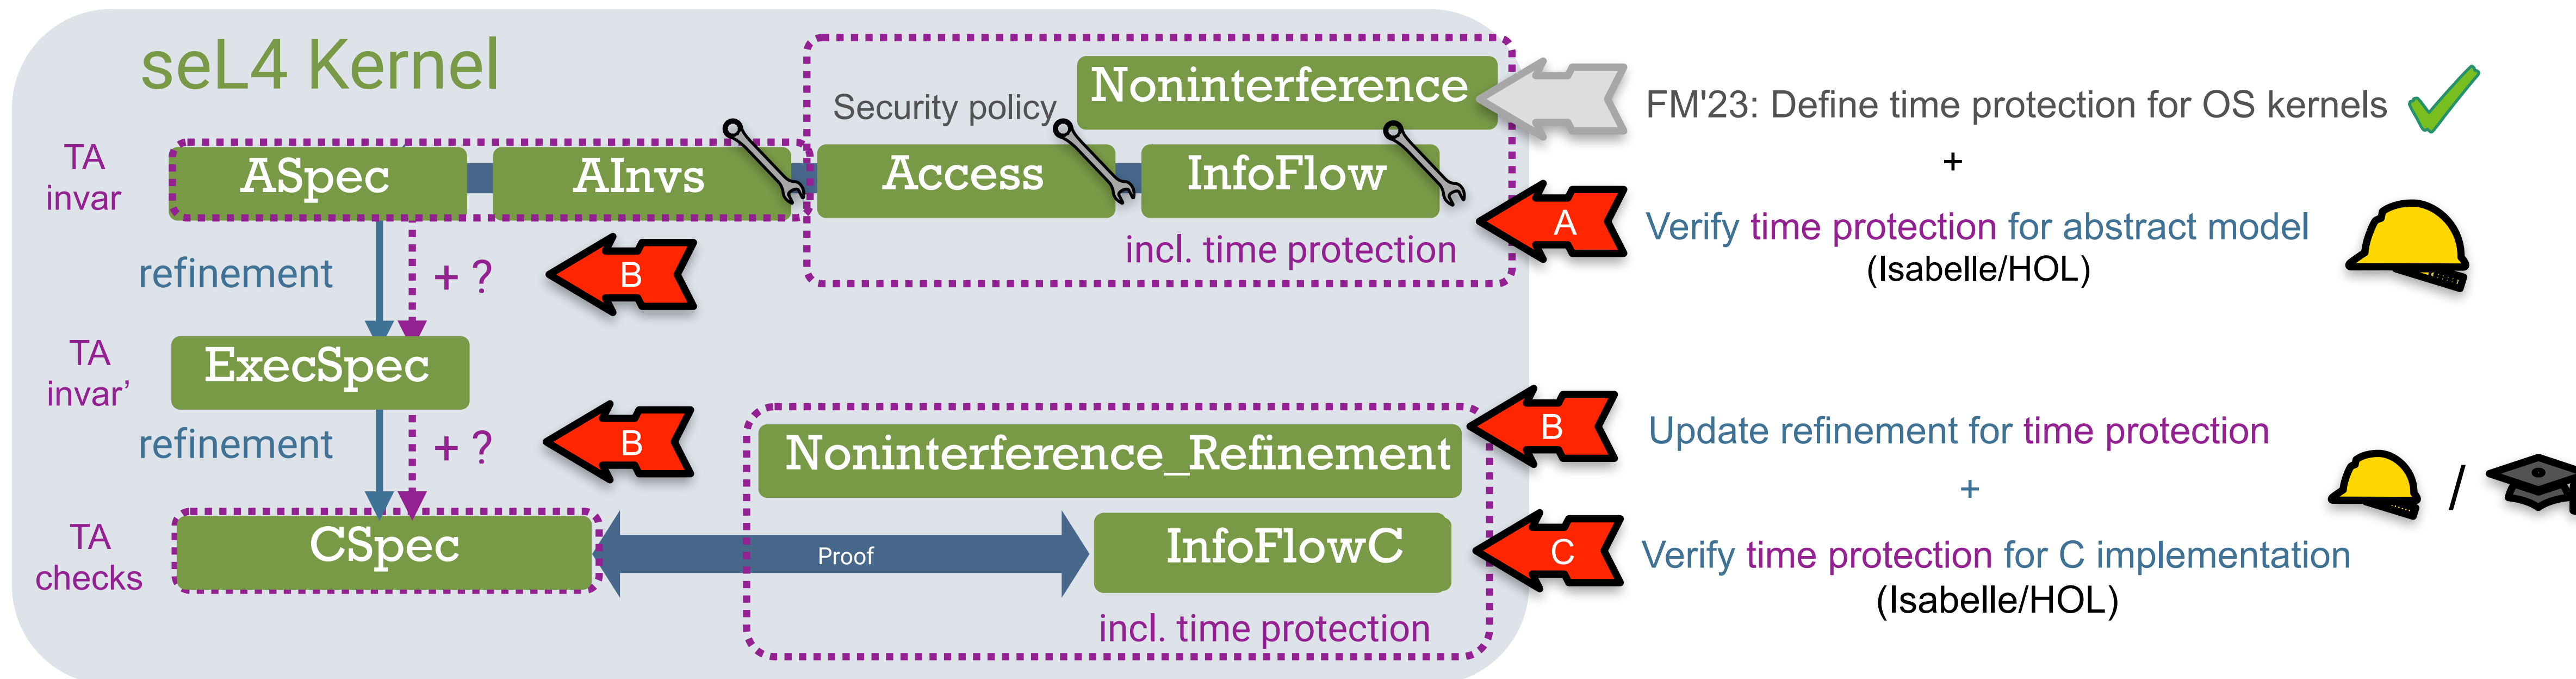
Verify time protection for abstract model (Isabelle/HOL) 🛡️

Update refinement for time protection + 🛡️ / 🎓

Verify time protection for C implementation (Isabelle/HOL)

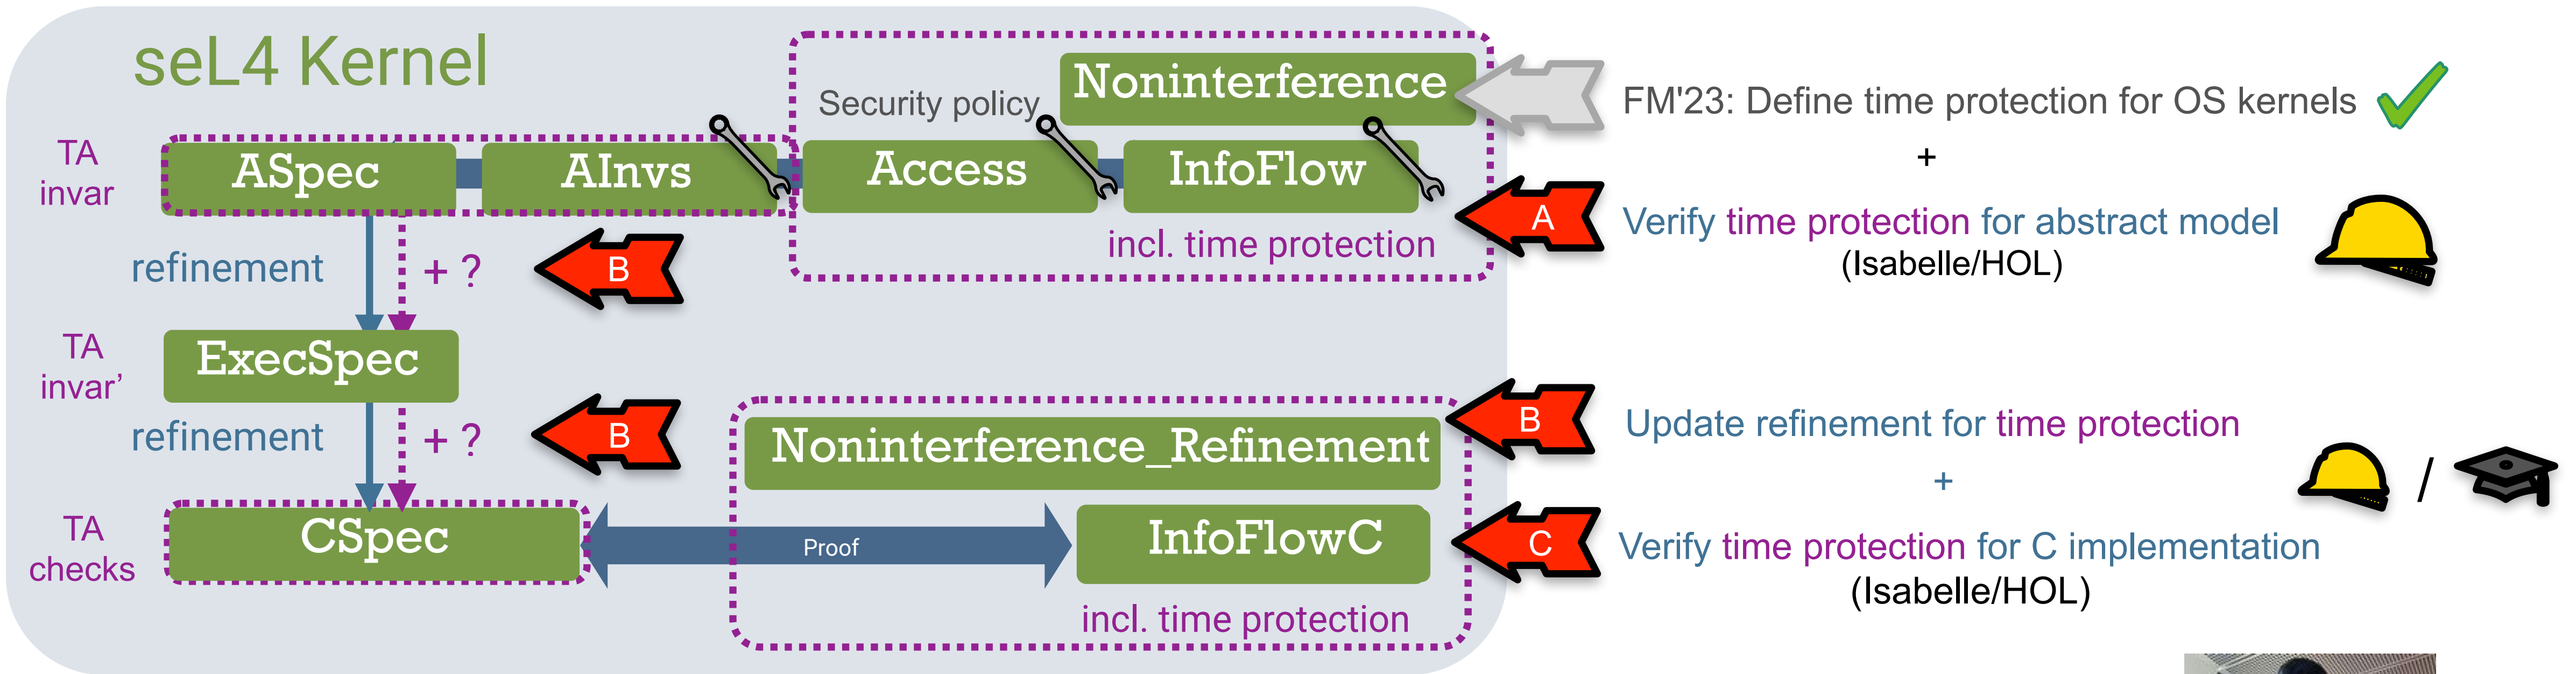
- Abstract model (ASpec) checks dynamic touched addresses (TA) set ➡️ 125 times
- (A) Key ASpec property: $TA \subseteq \text{domain's addresses}$ (according to security policy)
 - Also: AInvs, Access, InfoFlow need repairs 🛠️ ⚡ How many breakages? ➡️ > 685, after the ones we already fixed
- (B + C) Refinement: $TA'' \subseteq TA' \subseteq TA$? (via ExecSpec)
 - Tricky part #1: Refinement to CSpec accesses new addresses
 - Tricky part #2: Making TA checks scalable at CSpec level

New approach



- ~~Abstract model (ASpec) checks dynamic touched addresses (TA) set~~
- ~~(A) Key ASpec property: $TA \subseteq \text{domain's addresses}$ (according to security policy)~~
 - ~~Also: AInvs, Access, InfoFlow need repairs~~
- ~~(B + C) Refinement: $TA'' \subseteq TA' \subseteq TA$? (via ExecSpec)~~
 - Tricky part #1: Refinement to CSpec accesses new addresses
 - Tricky part #2: Making TA checks scalable at CSpec level

New approach



FM'23: Define time protection for OS kernels ✓

+
Verify time protection for abstract model (Isabelle/HOL) 🛑

+
Update refinement for time protection 🛑 / 🎓

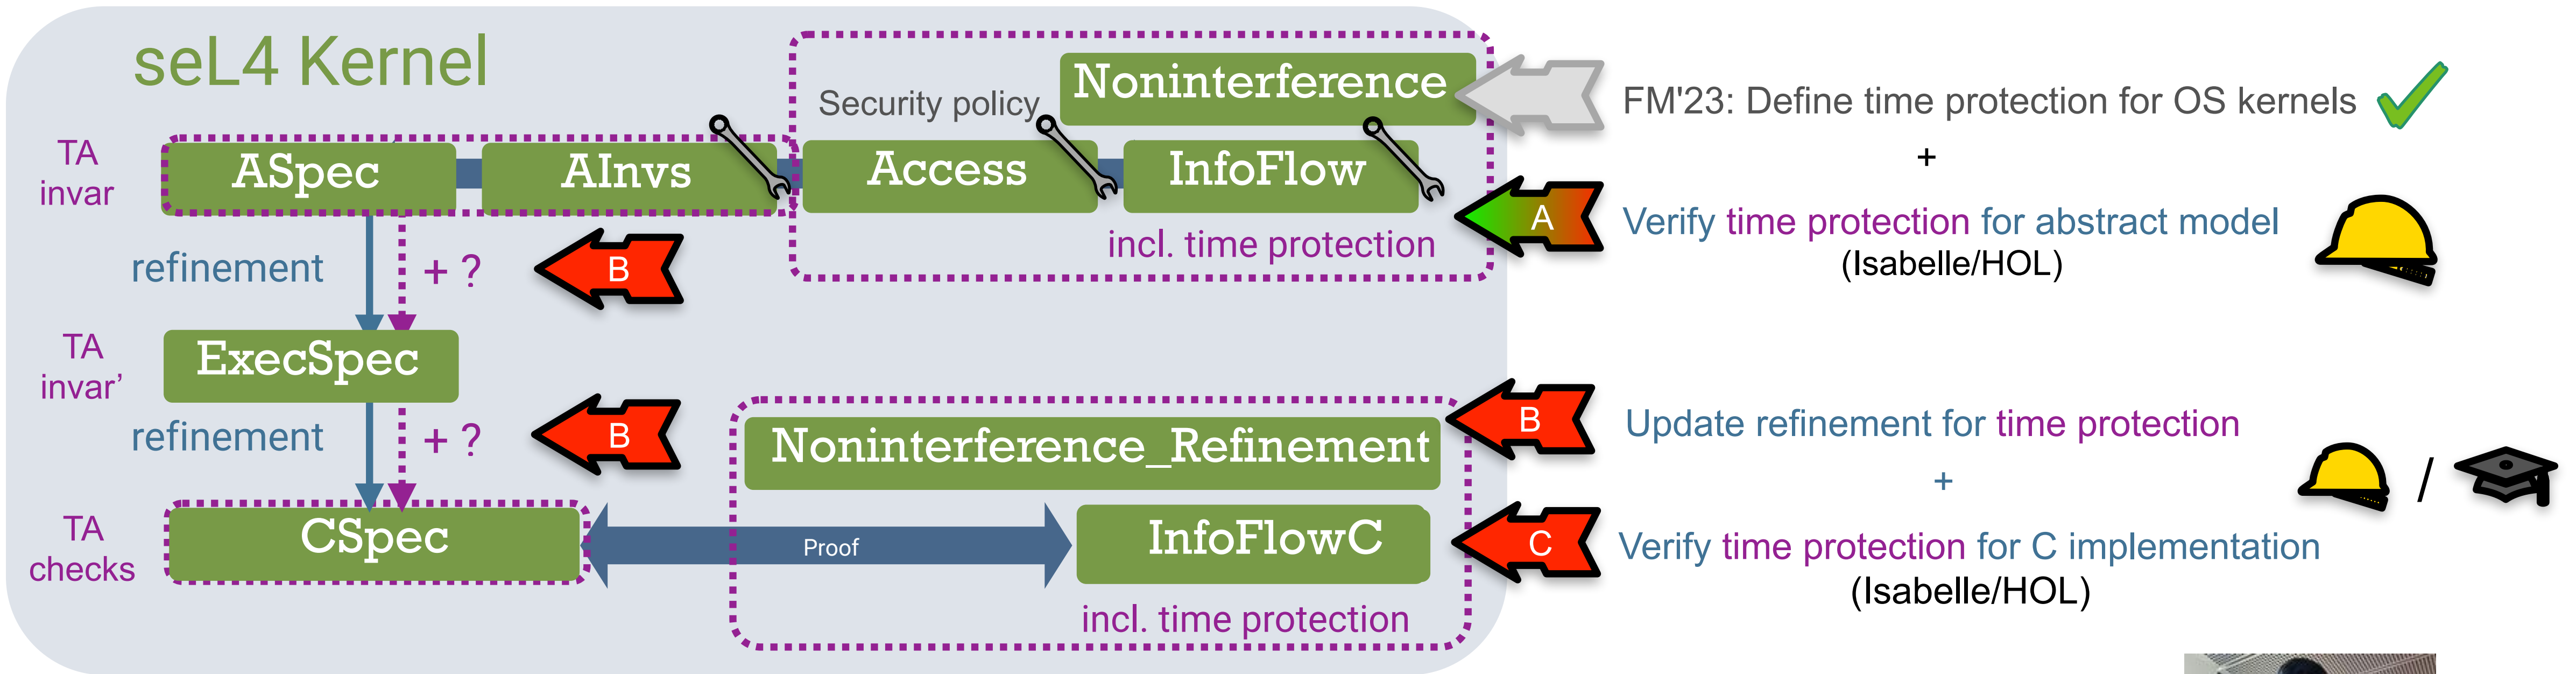
+
Verify time protection for C implementation (Isabelle/HOL)

- ~~Abstract model (ASpec) checks dynamic touched addresses (TA) set~~
- ~~(A) Key ASpec property: $TA \subseteq \text{domain's addresses}$ (according to security policy)~~
 - ~~Also: AInvs, Access, InfoFlow need repairs~~
- ~~(B + C) Refinement: $TA'' \subseteq TA' \subseteq TA$? (via ExecSpec)~~
 - Tricky part #1: Refinement to CSpec accesses new addresses
 - Tricky part #2: Making TA checks scalable at CSpec level

Sai Nair
BSc(Hons)
thesis



New approach

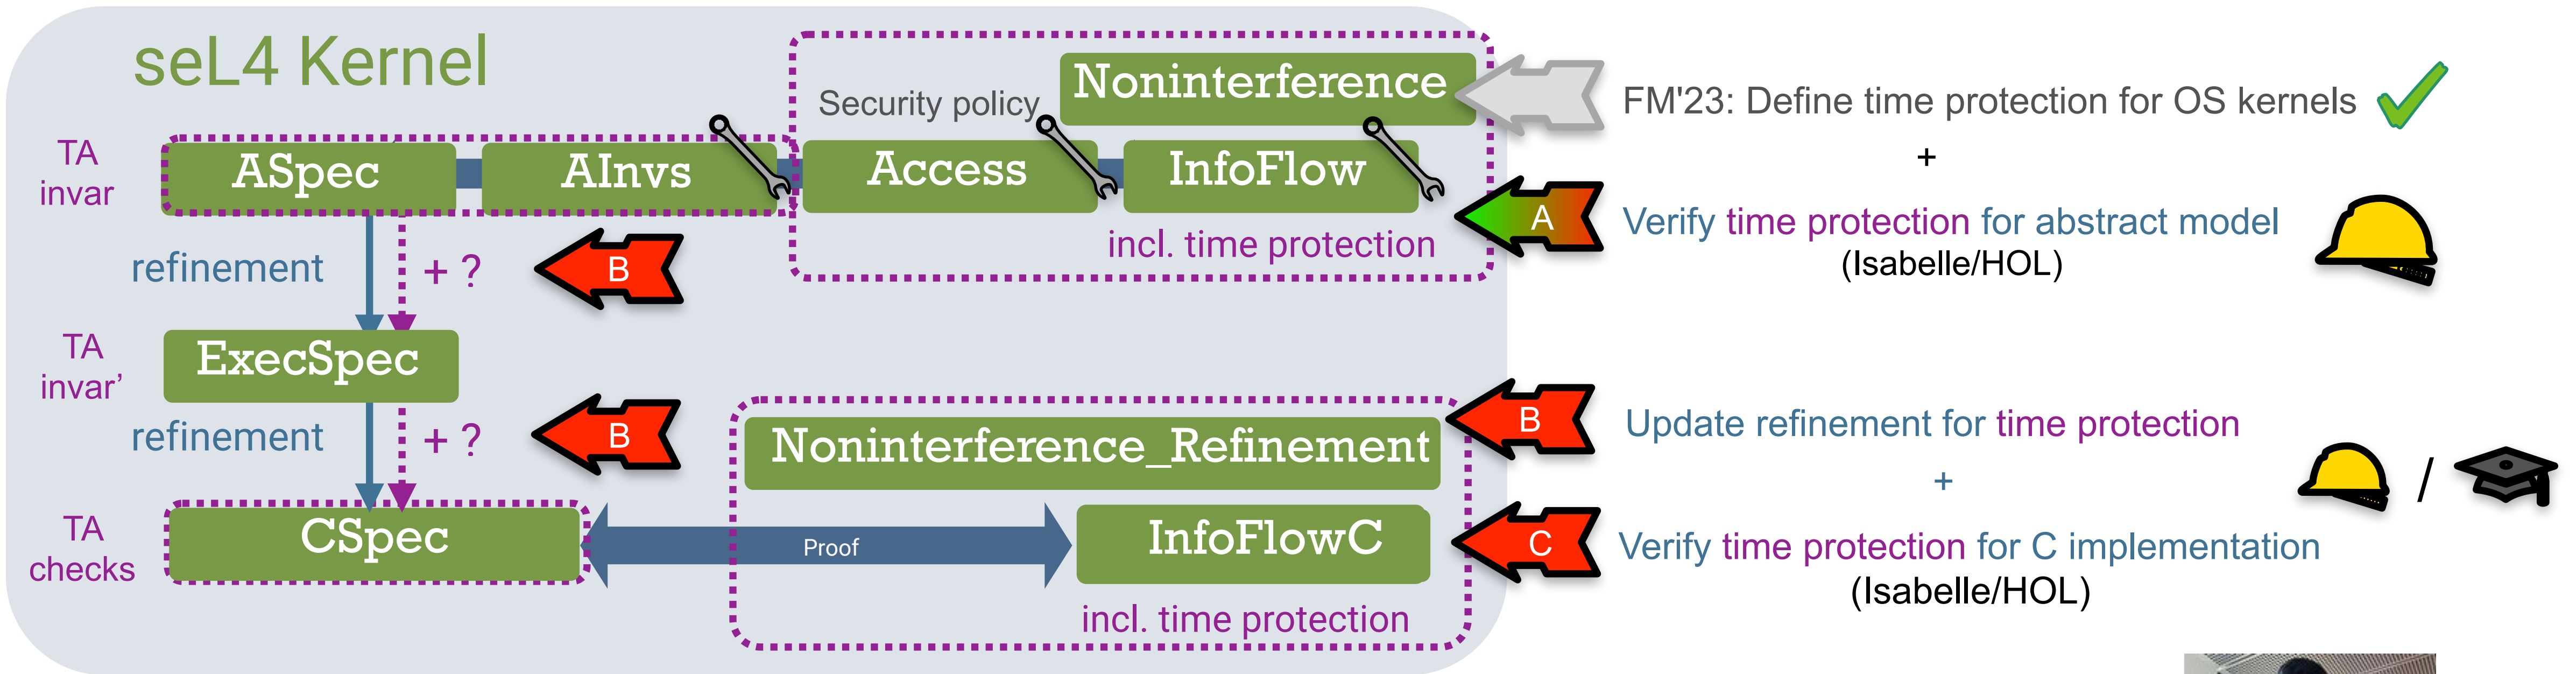


- ~~Abstract model (ASpec) checks dynamic touched addresses (TA) set~~
- (A) Key ASpec property: “TA invariant”
 - ~~Also: AInvs, Access, InfoFlow need repairs~~
- ~~(B + C) Refinement: TA'' ⊆ TA' ⊆ TA ? (via ExecSpec)~~
 - Tricky part #1: Refinement to CSpec accesses new addresses
 - Tricky part #2: Making TA checks scalable at CSpec level

Sai Nair
BSc(Hons)
thesis



New approach



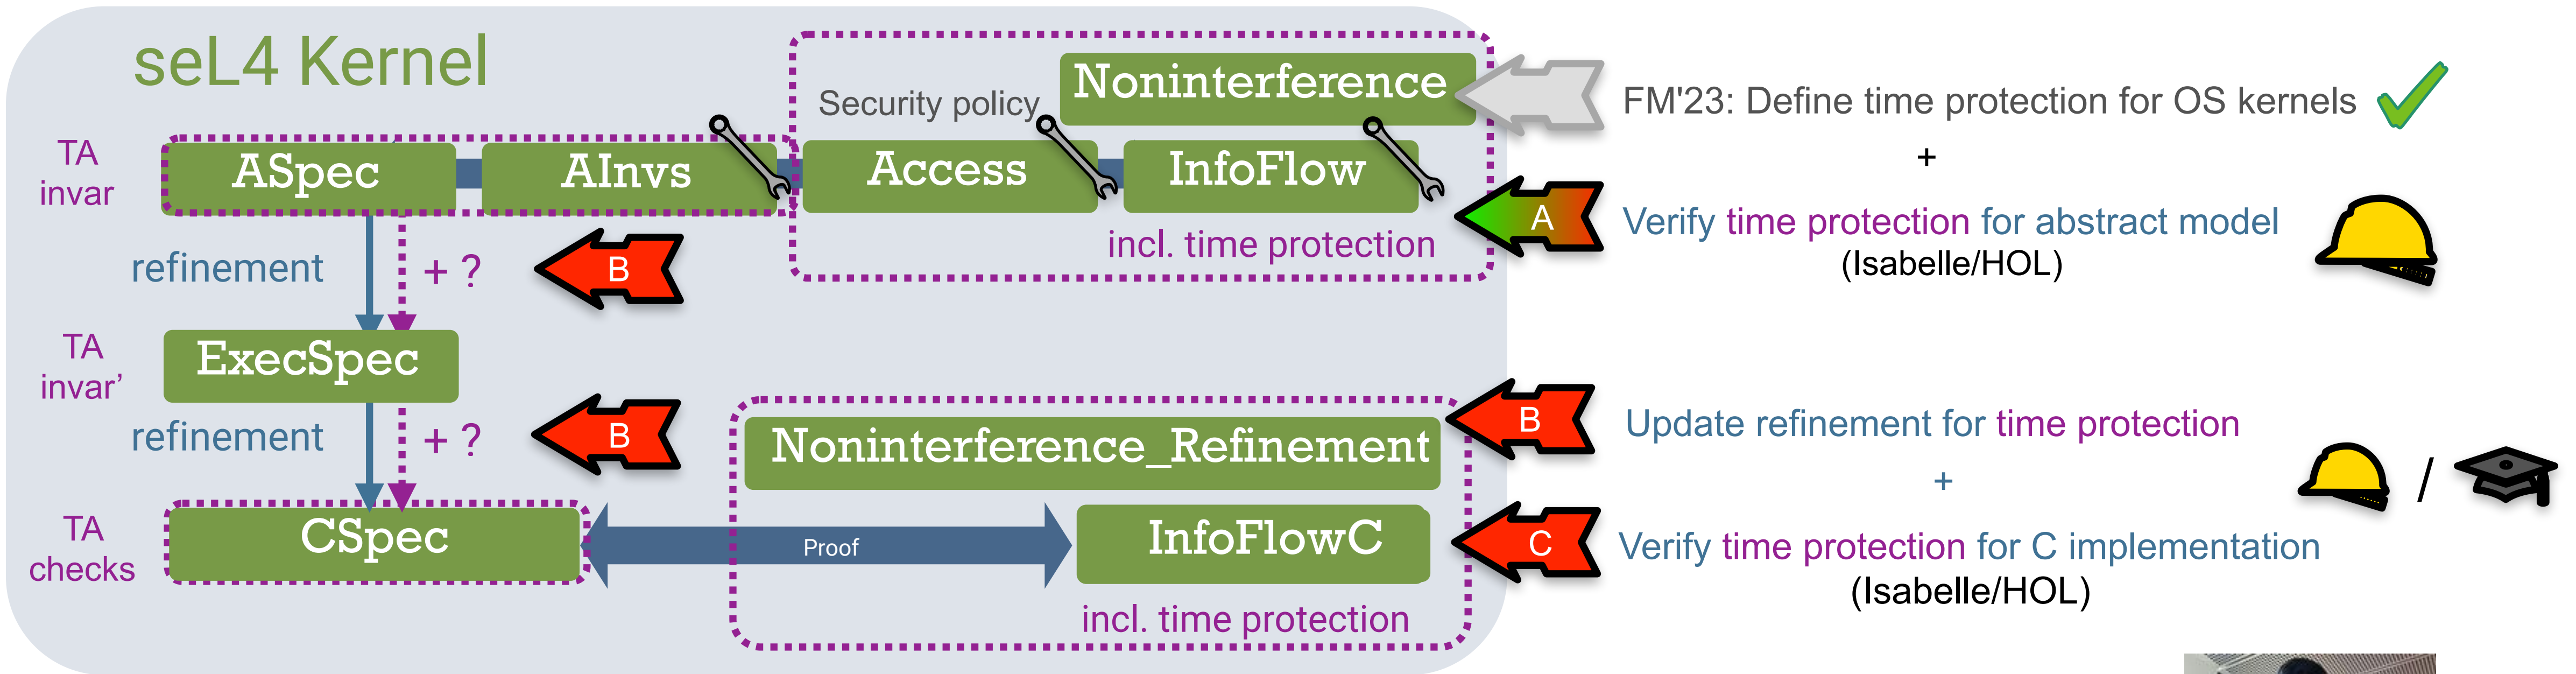
"obj refs don't point outside static TA set"

- (A) Key ASpec property: "TA invariant" ↻
 - Also: AInvs, Access, InfoFlow need repairs
- ~~(B + C) Refinement: TA'' ⊆ TA' ⊆ TA ? (via ExecSpec)~~
 - Tricky part #1: Refinement to CSpec accesses new addresses
 - Tricky part #2: Making TA checks scalable at CSpec level

Sai Nair
BSc(Hons)
thesis



New approach



“obj refs don't point outside static TA set”

- (A) Key ASpec property: “TA invariant”
 - Also: ~~AInvs, Access, InfoFlow~~ need repairs
- ~~(B + C) Refinement: TA'' ⊆ TA' ⊆ TA ? (via ExecSpec)~~
 - Tricky part #1: Refinement to CSpec accesses new addresses
 - Tricky part #2: Making TA checks scalable at CSpec level

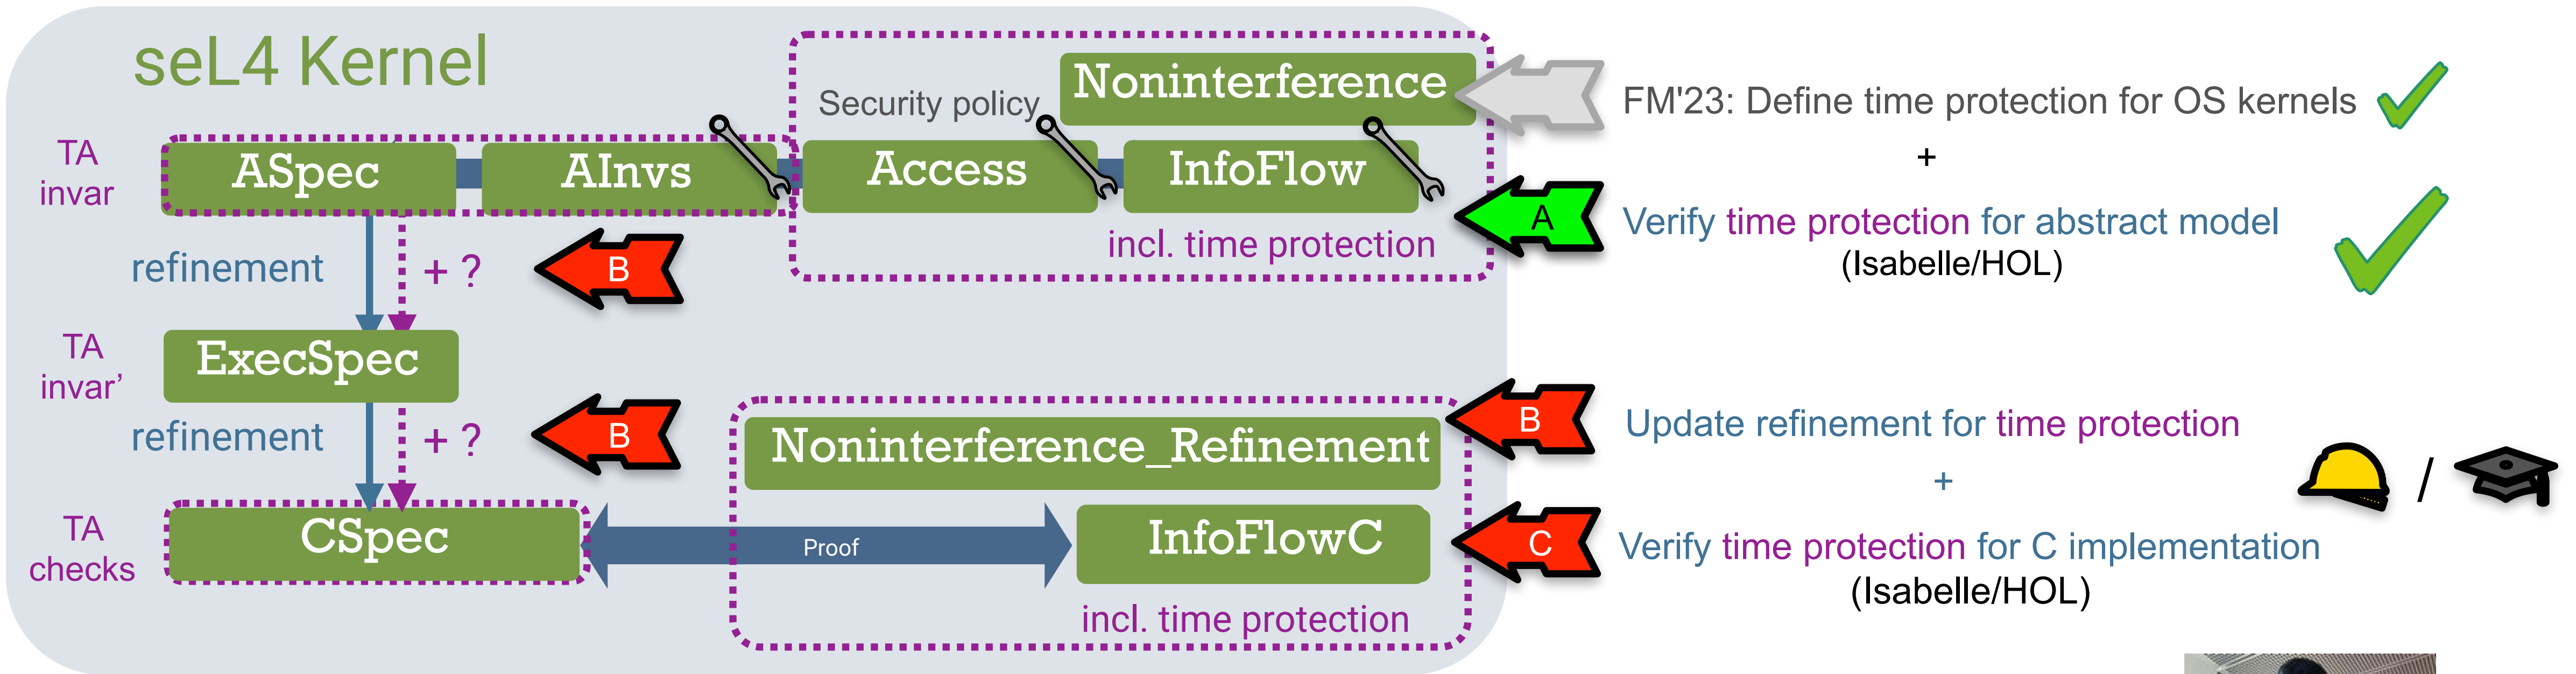
• Proof across ASpec in progress

Continuing on from

Sai Nair
BSc(Hons)
thesis



New approach



“obj refs don't point outside static TA set”

- (A) Key ASpec property: “TA invariant”
 - ASpec otherwise integrated into time protection theory ✓
- ~~(B + C) Refinement: $TA'' \subseteq TA' \subseteq TA$? (via ExecSpec)~~
 - Tricky part #1: Refinement to CSpec accesses new addresses
 - Tricky part #2: Making TA checks scalable at CSpec level

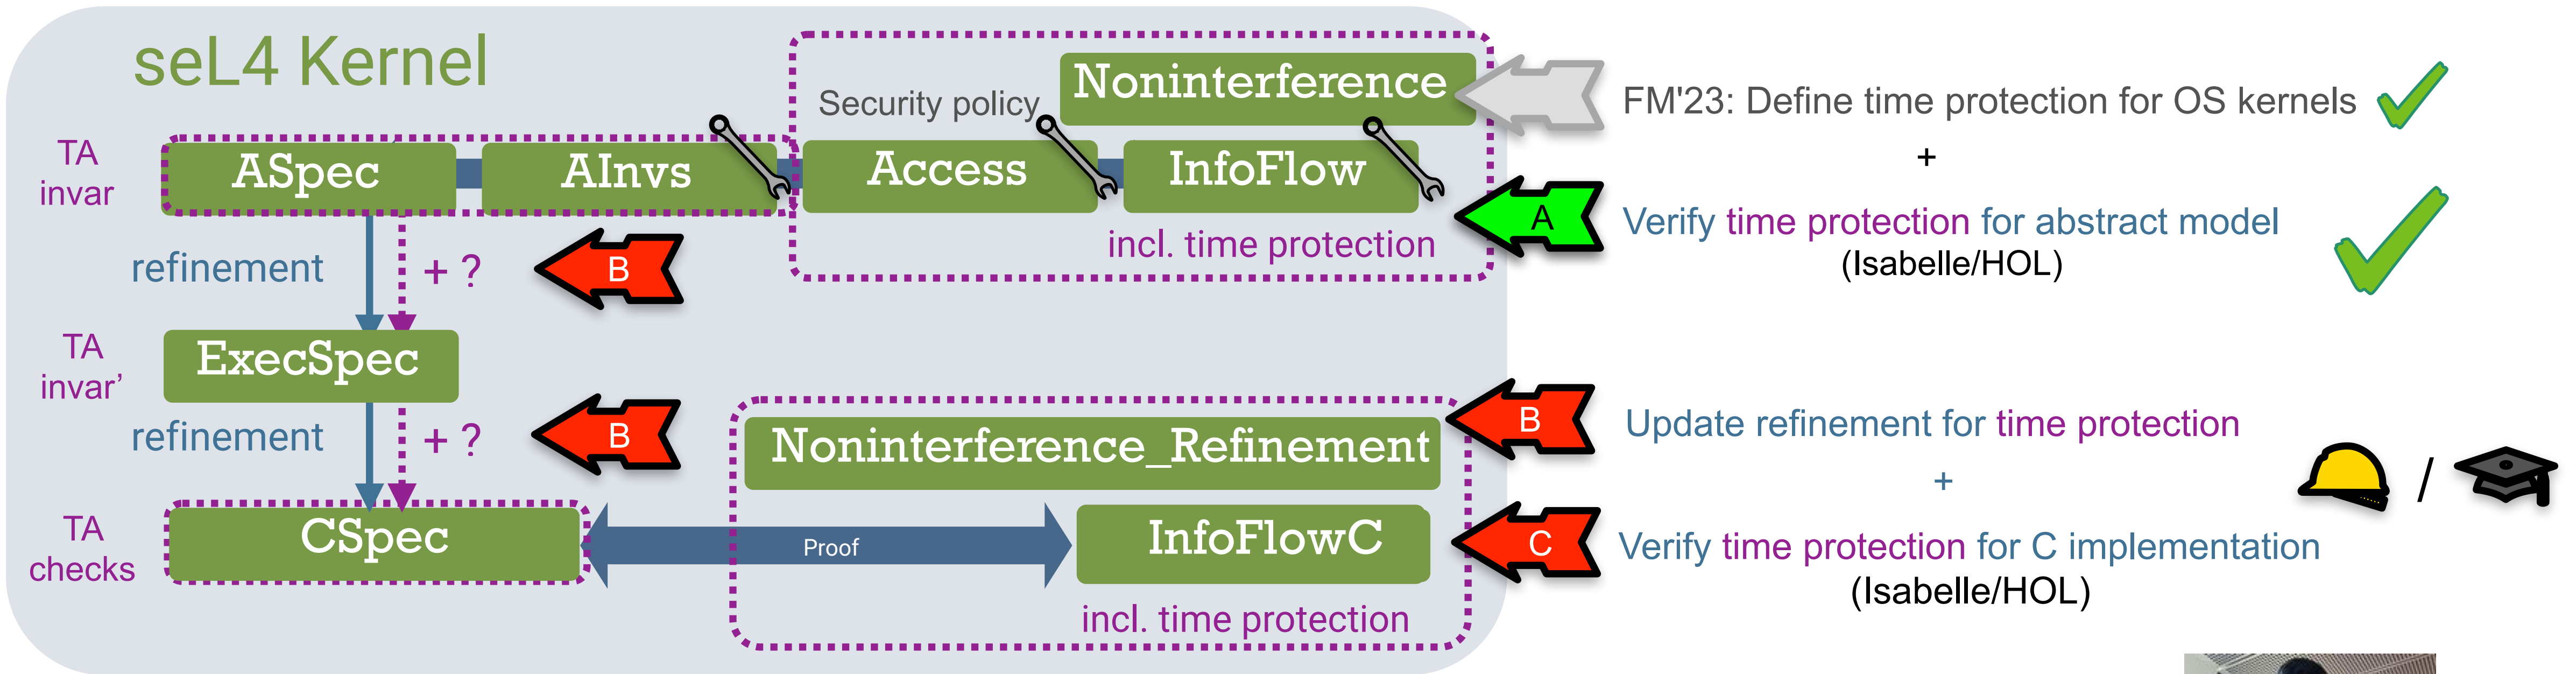
• Proof across ASpec in progress

Sai Nair
BSc(Hons)
thesis



Continuing on from

New approach



“obj refs don't point outside static TA set”

- (A) Key ASpec property: “TA invariant”
 - ASpec otherwise integrated into time protection theory ✓
- ~~(B + C) Refinement: $TA'' \subseteq TA' \subseteq TA$? (via ExecSpec)~~
 - Tricky part #1: Refinement to CSpec accesses new addresses
 - Tricky part #2: Making TA checks scalable at CSpec level

• Proof across ASpec in progress

Continuing on from

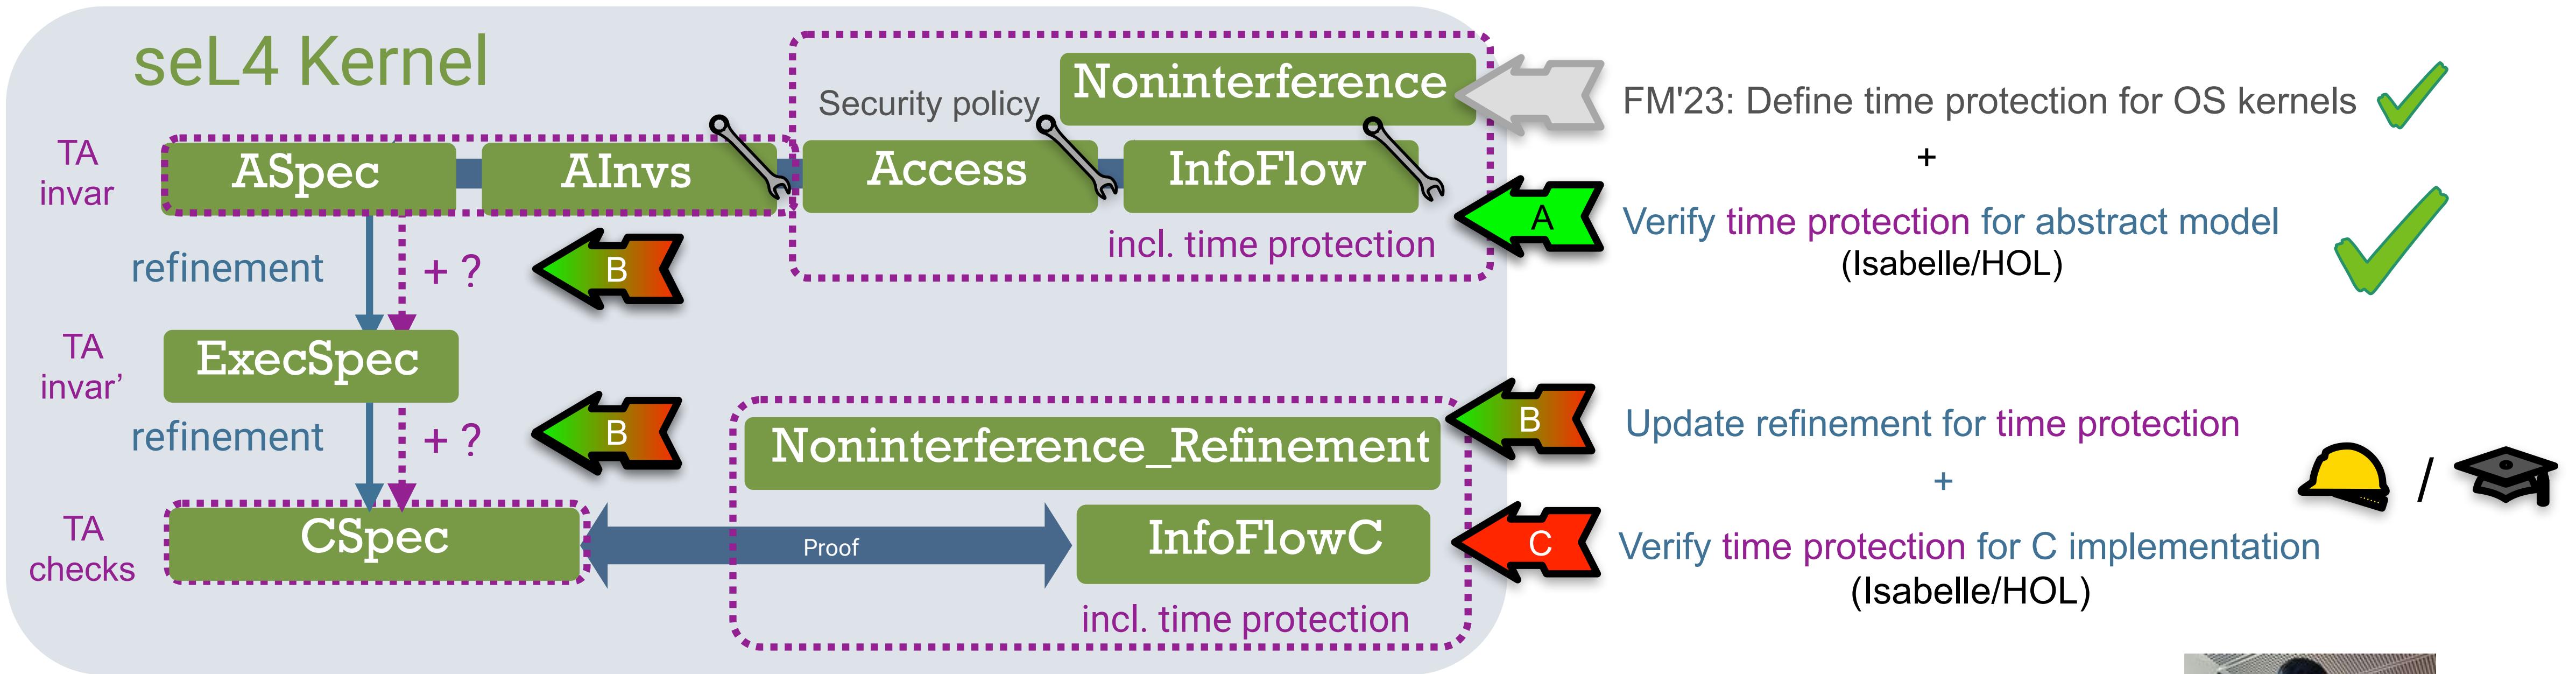
Sai Nair
BSc(Hons)
thesis



Thomas Liang
BSc (Hons)
Thesis



New approach



“obj refs don't point outside static TA set”

- (A) Key **ASpec** property: “TA invariant” • Proof across **ASpec** in progress
- **ASpec** otherwise *integrated* into time protection theory ✓
- (B + C) Refinement: TA invariant (via **ExecSpec**) \Rightarrow TA checks pass (**CSpec**)
- Tricky part #1: Refinement to **CSpec** accesses new addresses
- Tricky part #2: Making TA checks scalable at **CSpec** level

Continuing on from

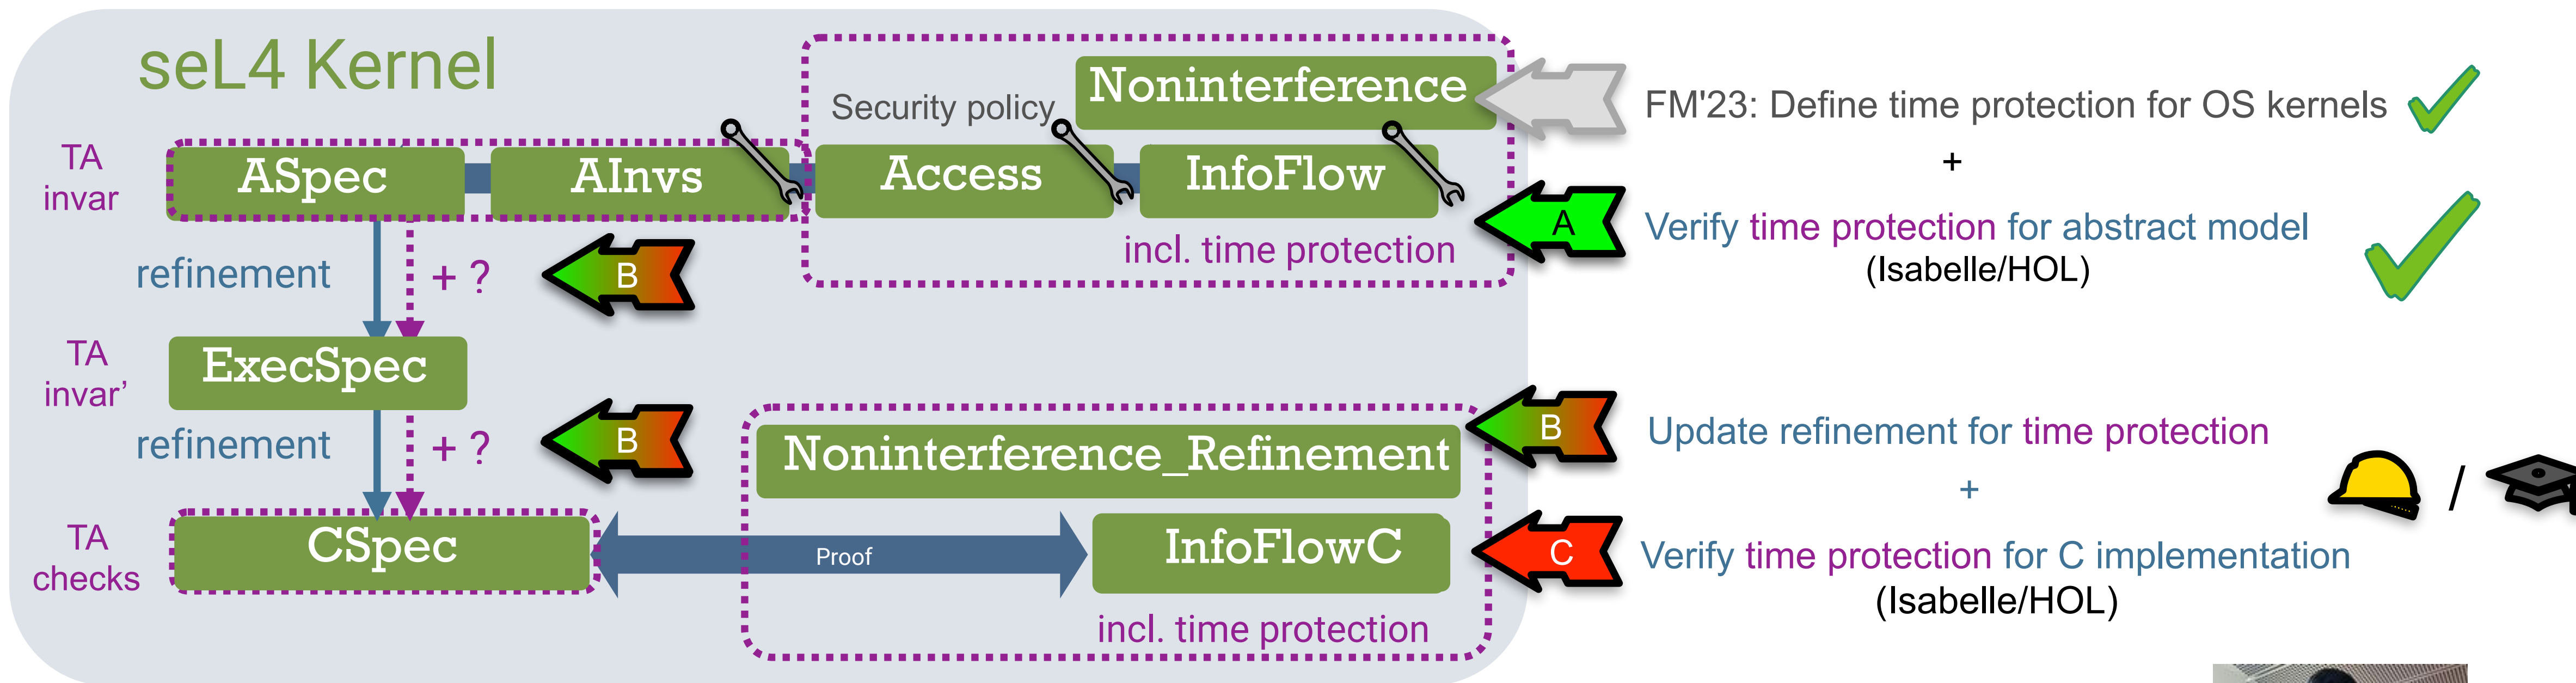
Sai Nair
BSc(Hons)
thesis



Thomas Liang
BSc (Hons)
Thesis



What's next?



“obj refs don't point outside static TA set”

- (A) Key **ASpec** property: “TA invariant” • Proof across **ASpec** in progress
- **ASpec** otherwise *integrated* into time protection theory
- (B + C) Refinement: TA invariant (via **ExecSpec**) \Rightarrow TA checks pass (**CSpec**)
- Tricky part #1: Refinement to **CSpec** accesses new addresses
- Tricky part #2: Making TA checks scalable at **CSpec** level

Continuing on from

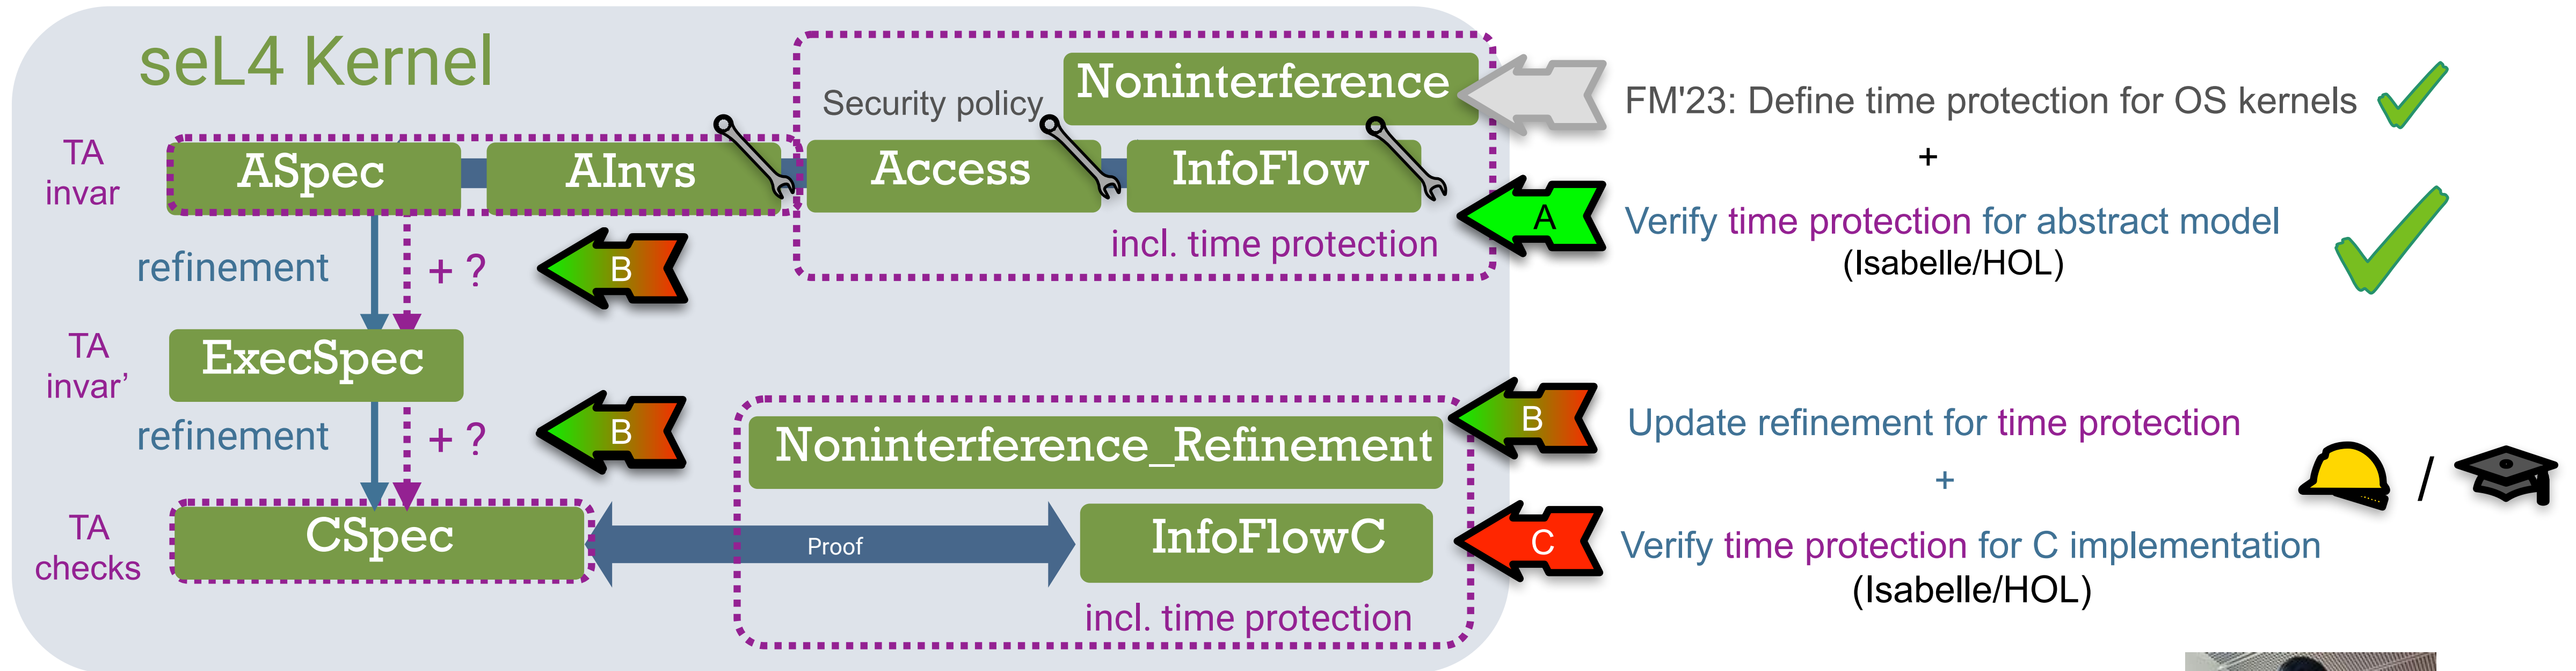
Sai Nair
BSc(Hons)
thesis



Thomas Liang
BSc (Hons)
Thesis



What's next? Refinement



“obj refs don't point outside static TA set”

- (A) Key **ASpec** property: “TA invariant” • Proof across **ASpec** in progress
- **ASpec** otherwise *integrated* into time protection theory
- (B + C) Refinement: TA invariant (via **ExecSpec**) \Rightarrow TA checks pass (**CSpec**)
- Tricky part #1: Refinement to **CSpec** accesses new addresses
- Tricky part #2: Making TA checks scalable at **CSpec** level

Continuing on from

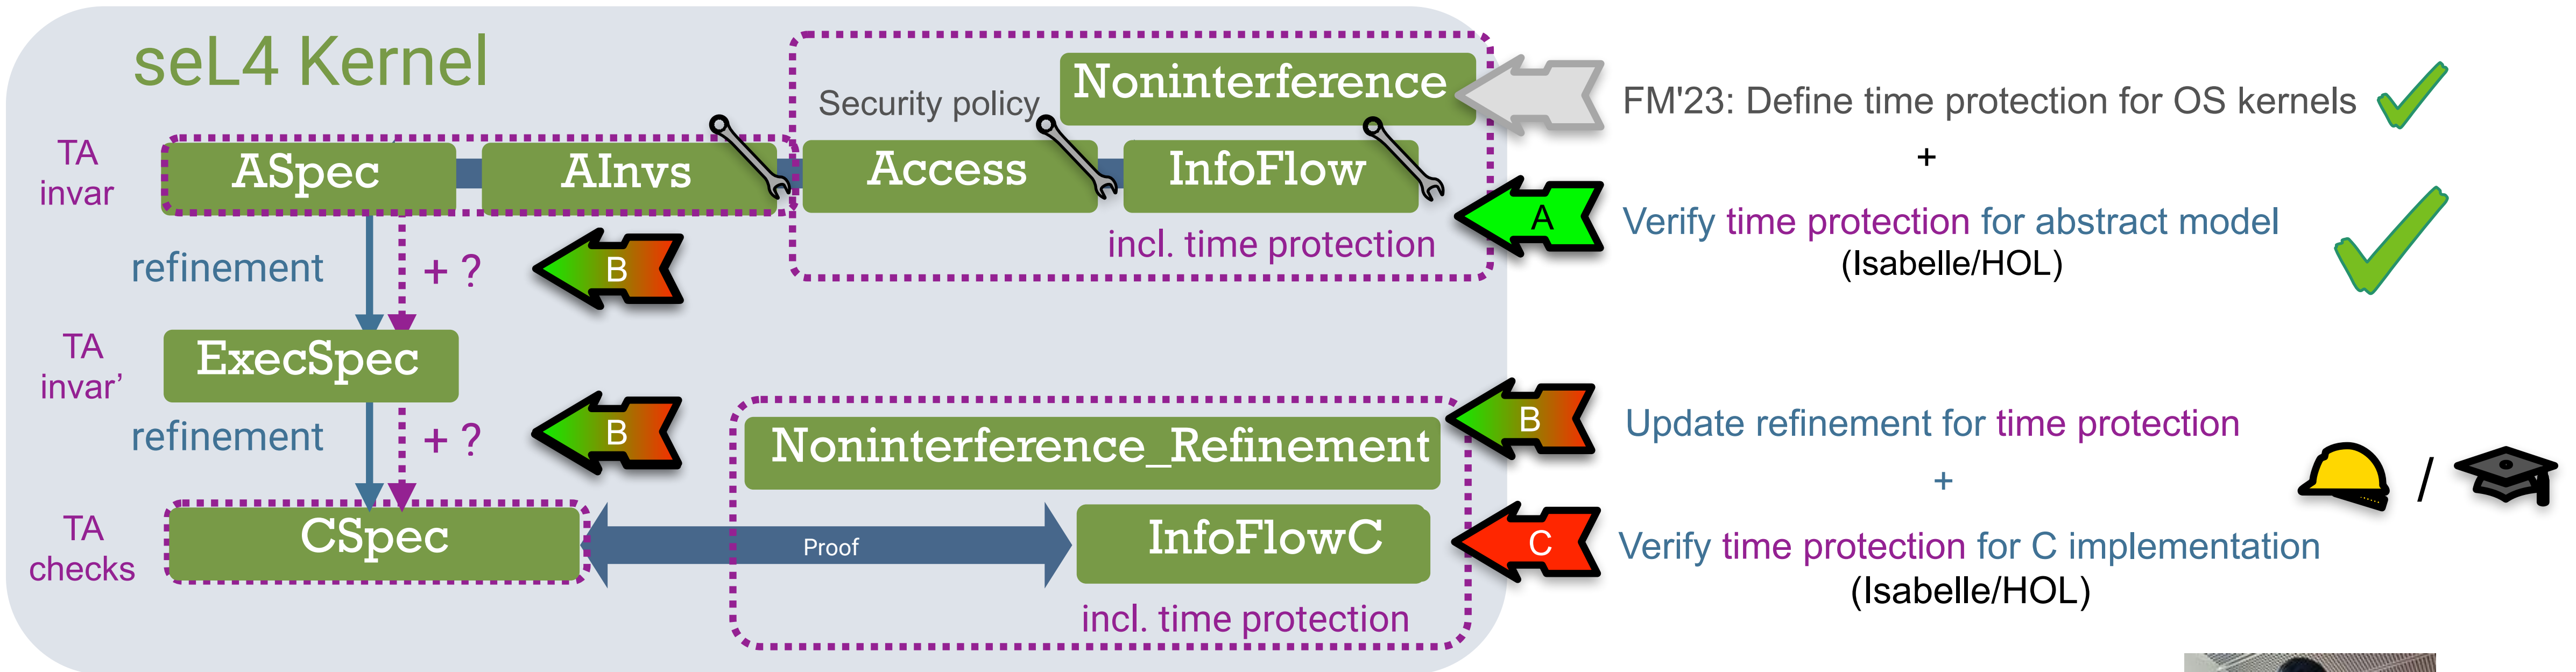
Sai Nair
BSc(Hons)
thesis



Thomas Liang
BSc (Hons)
Thesis



What's next? Refinement



“obj refs don't point outside static TA set”

- (A) Key ASpec property: “TA invariant” • Proof across ASpec in progress
- ASpec otherwise *integrated* into time protection theory
- (B + C) Refinement: TA invariant (via ExecSpec) \Rightarrow TA checks pass (CSpec)
 - Tricky part #1: Refinement to CSpec accesses new addresses
 - Tricky part #2: Making TA checks scalable at CSpec level

Sai Nair
BSc(Hons)
thesis

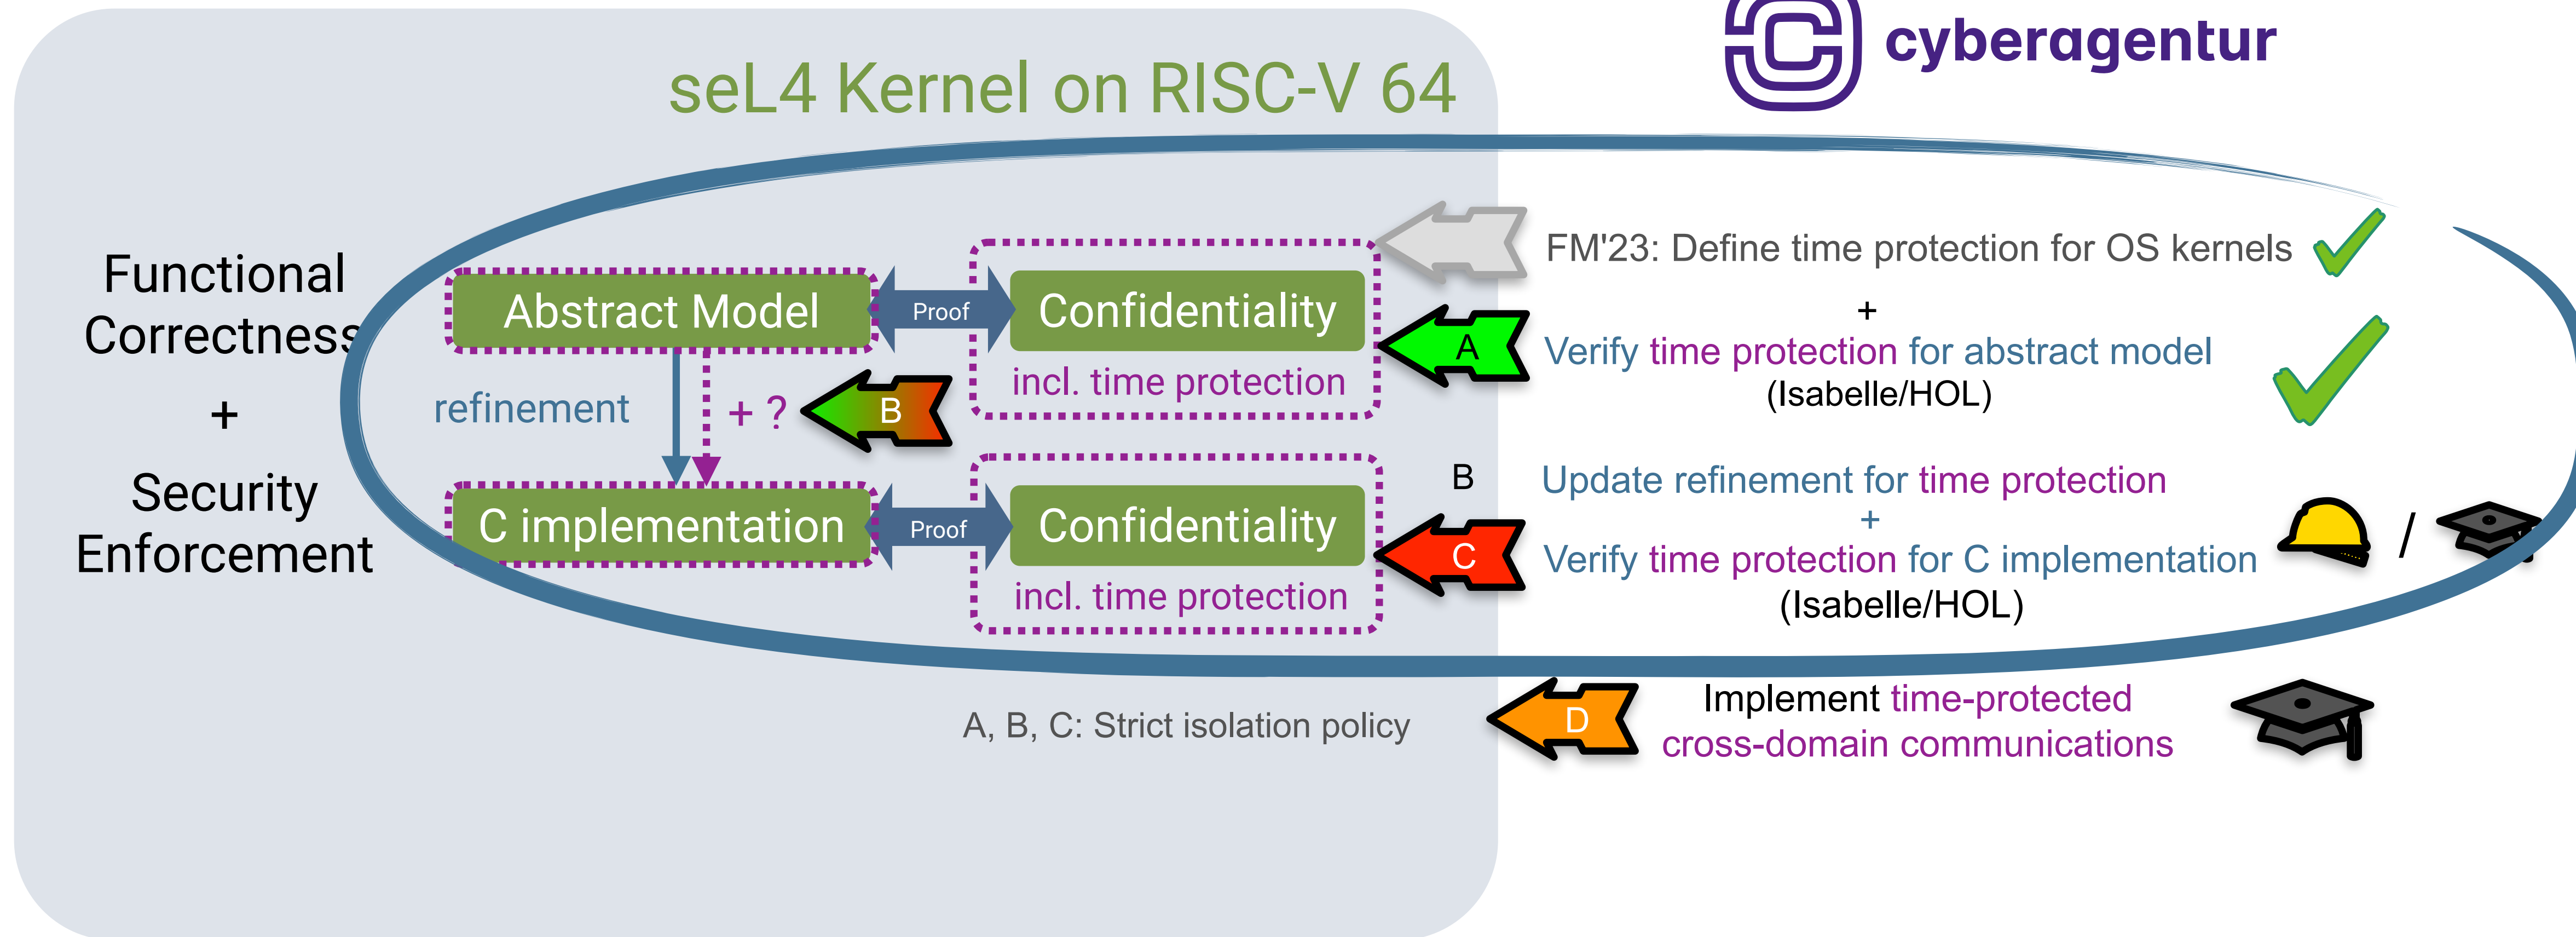


Thomas Liang
BSc (Hons)
Thesis



Since the plan 2 years ago...

Thanks to funding from





Time protection verification for seL4



Since the plan 2 years ago...

Thanks to funding from



seL4 Kernel on RISC-V 64

Functional Correctness
+
Security Enforcement

Abstract Model

Confidentiality
incl. time protection

refinement

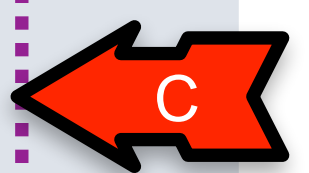
+ ?

C implementation

Confidentiality
incl. time protection

Proof

Proof



FM'23: Define time protection for OS kernels ✓

+
Verify time protection for abstract model (Isabelle/HOL) ✓

+
Update refinement for time protection

+
Verify time protection for C implementation (Isabelle/HOL) 📎 / 🎓

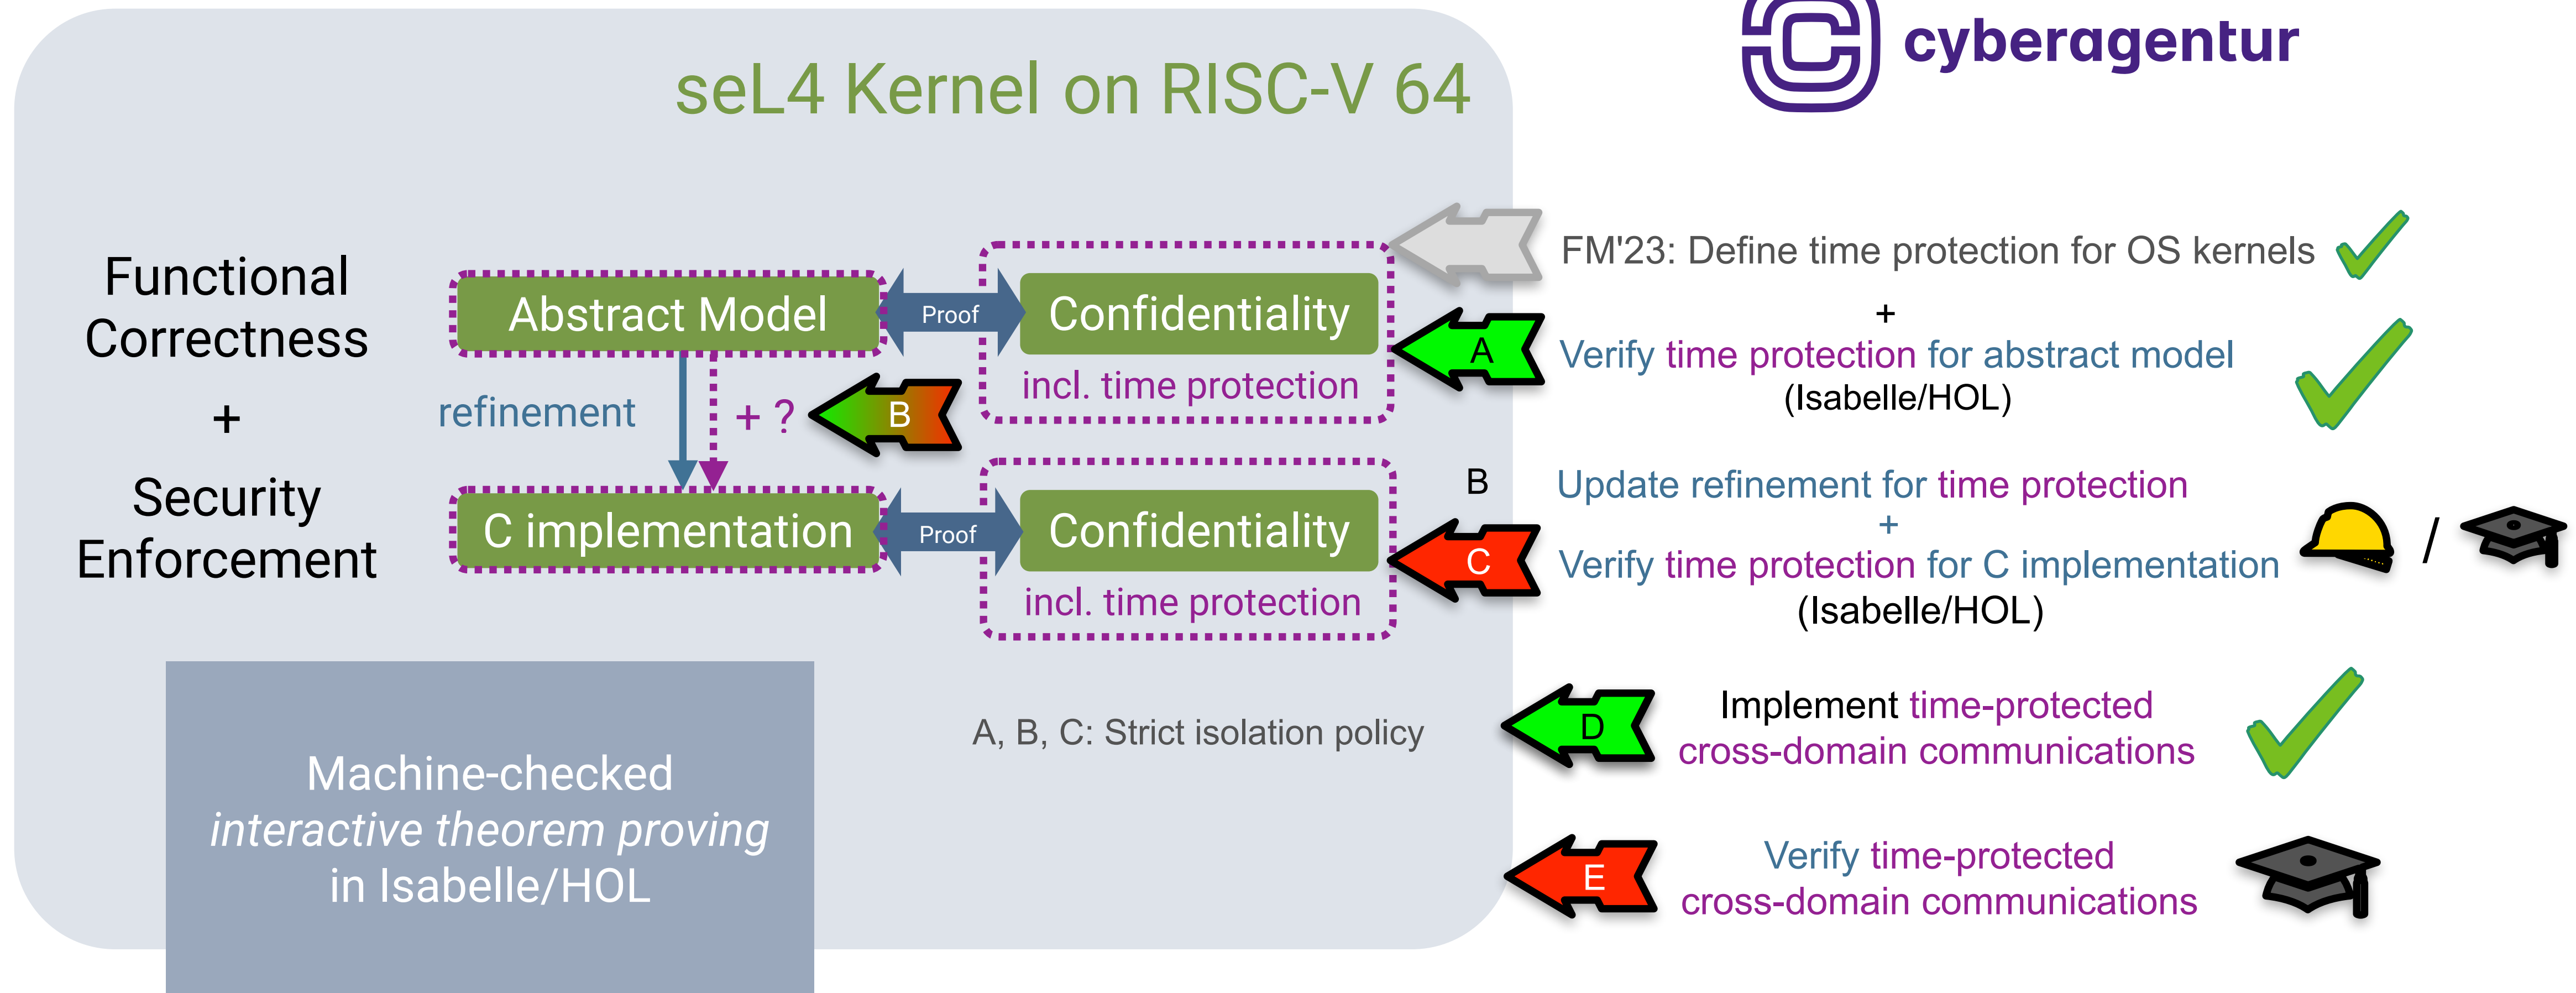
Implement time-protected cross-domain communications 🎓

A, B, C: Strict isolation policy

Machine-checked
interactive theorem proving
in Isabelle/HOL

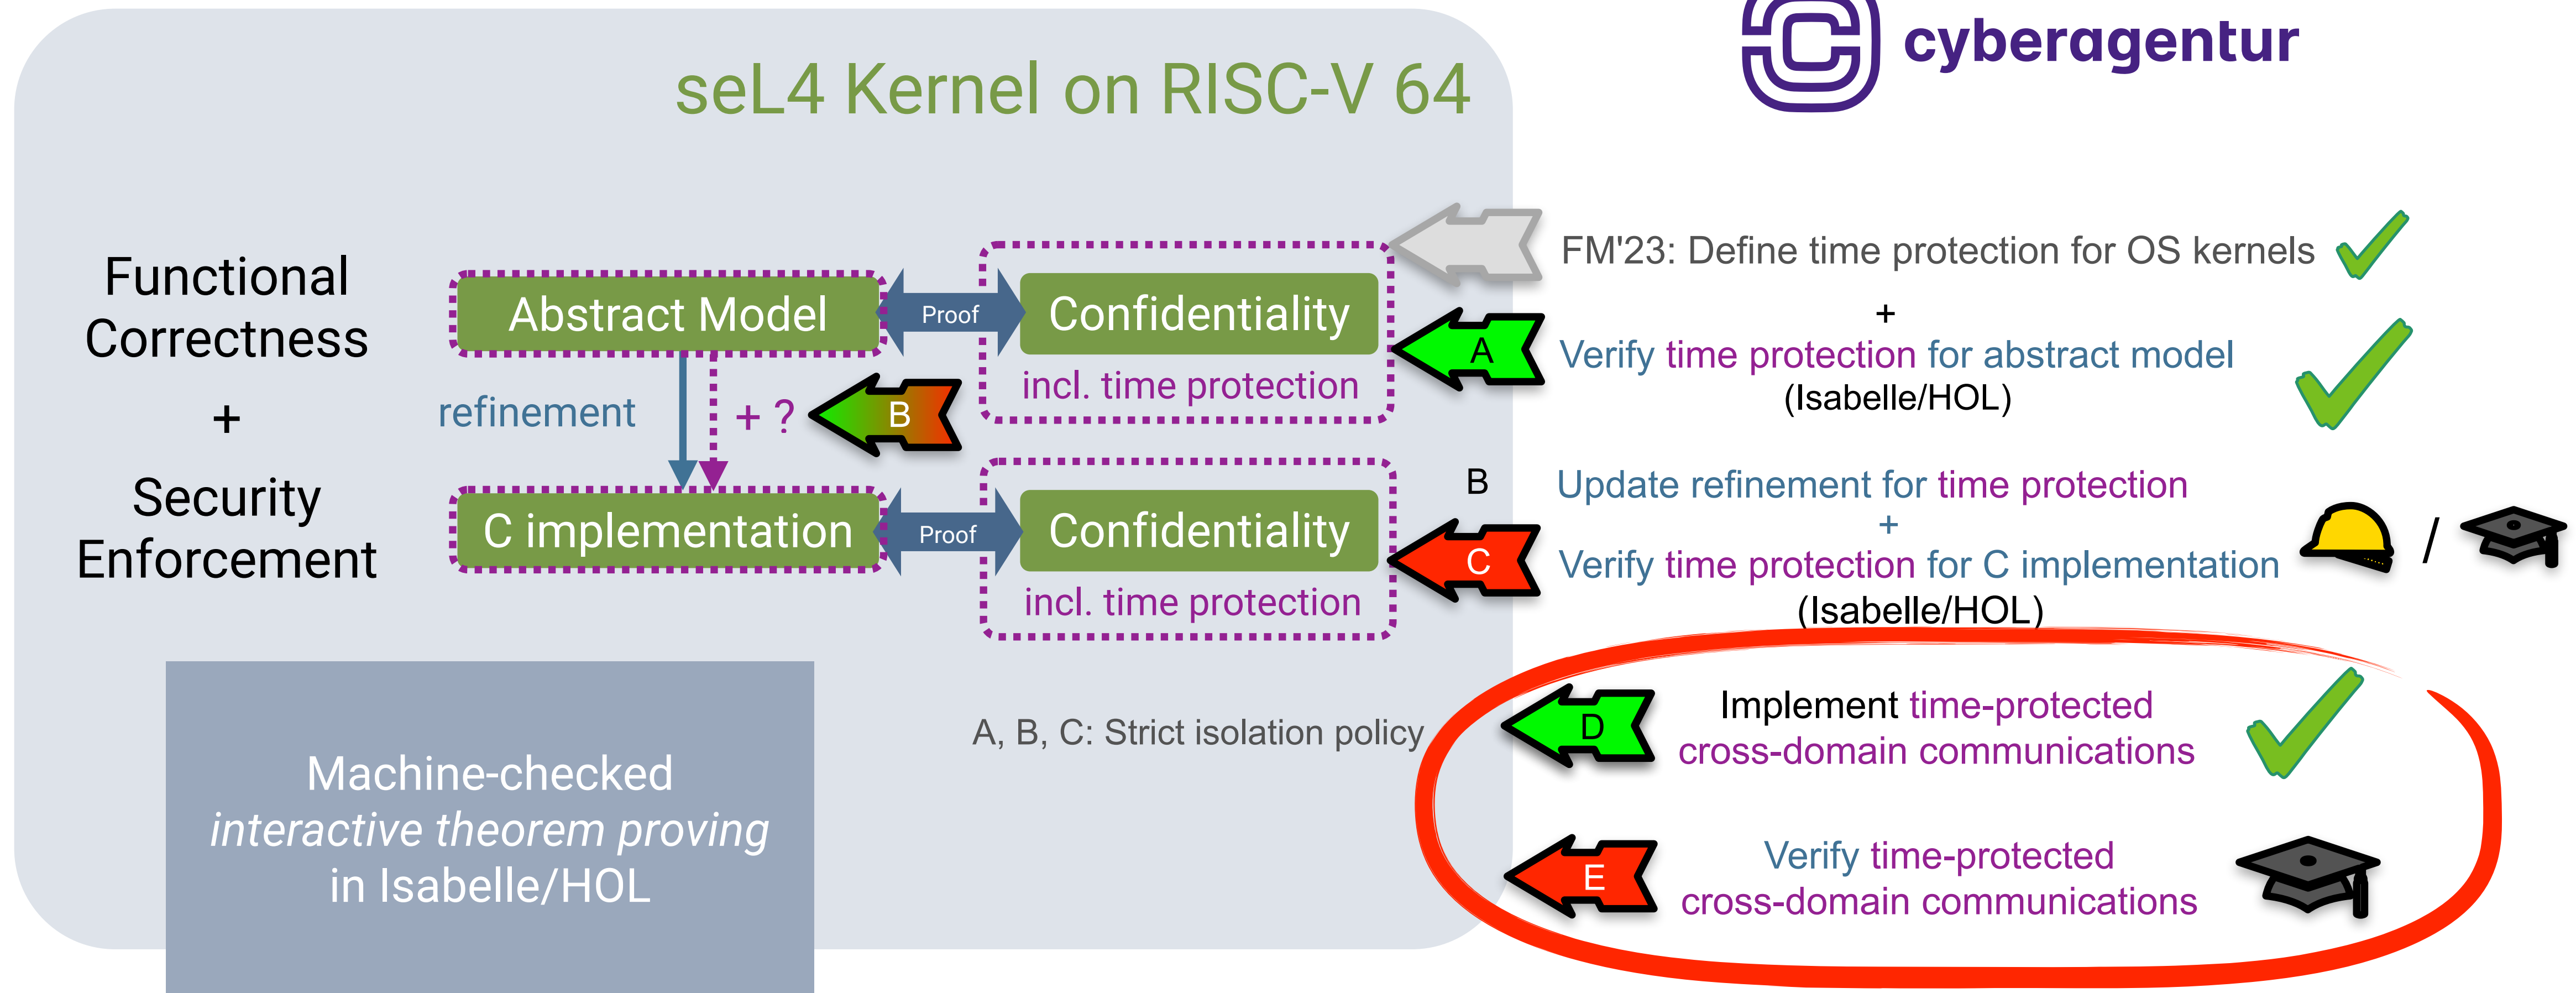
The status today:

Thanks to funding from

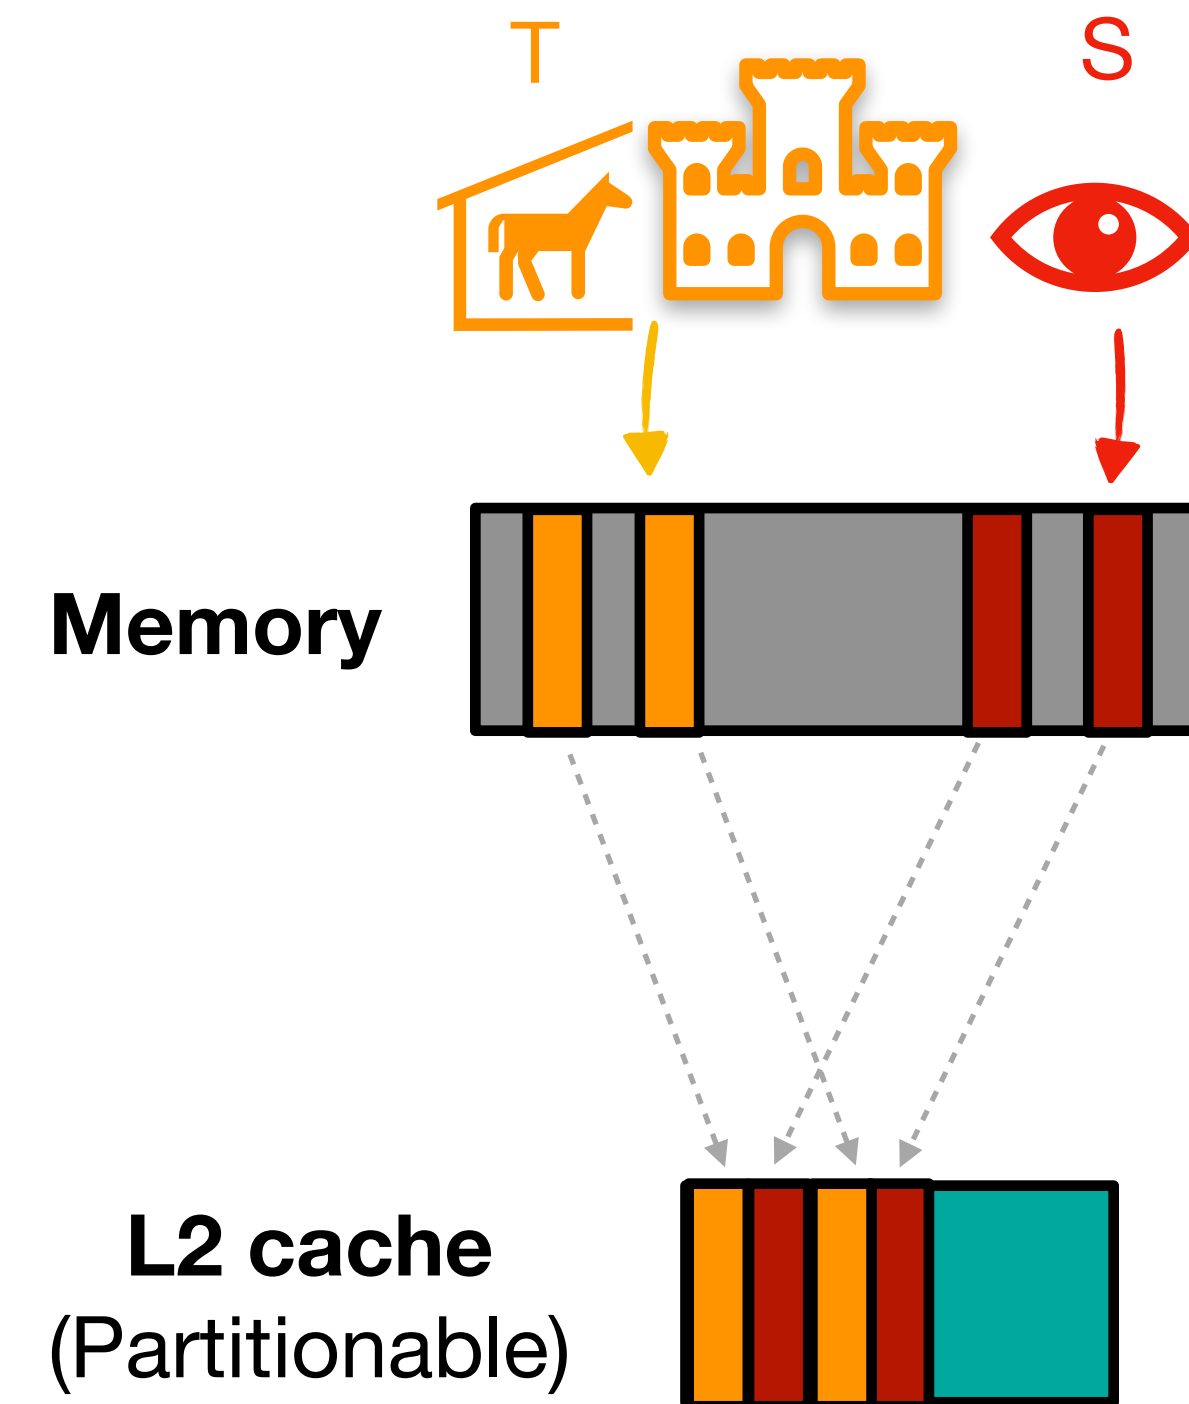


The status today:

Thanks to funding from



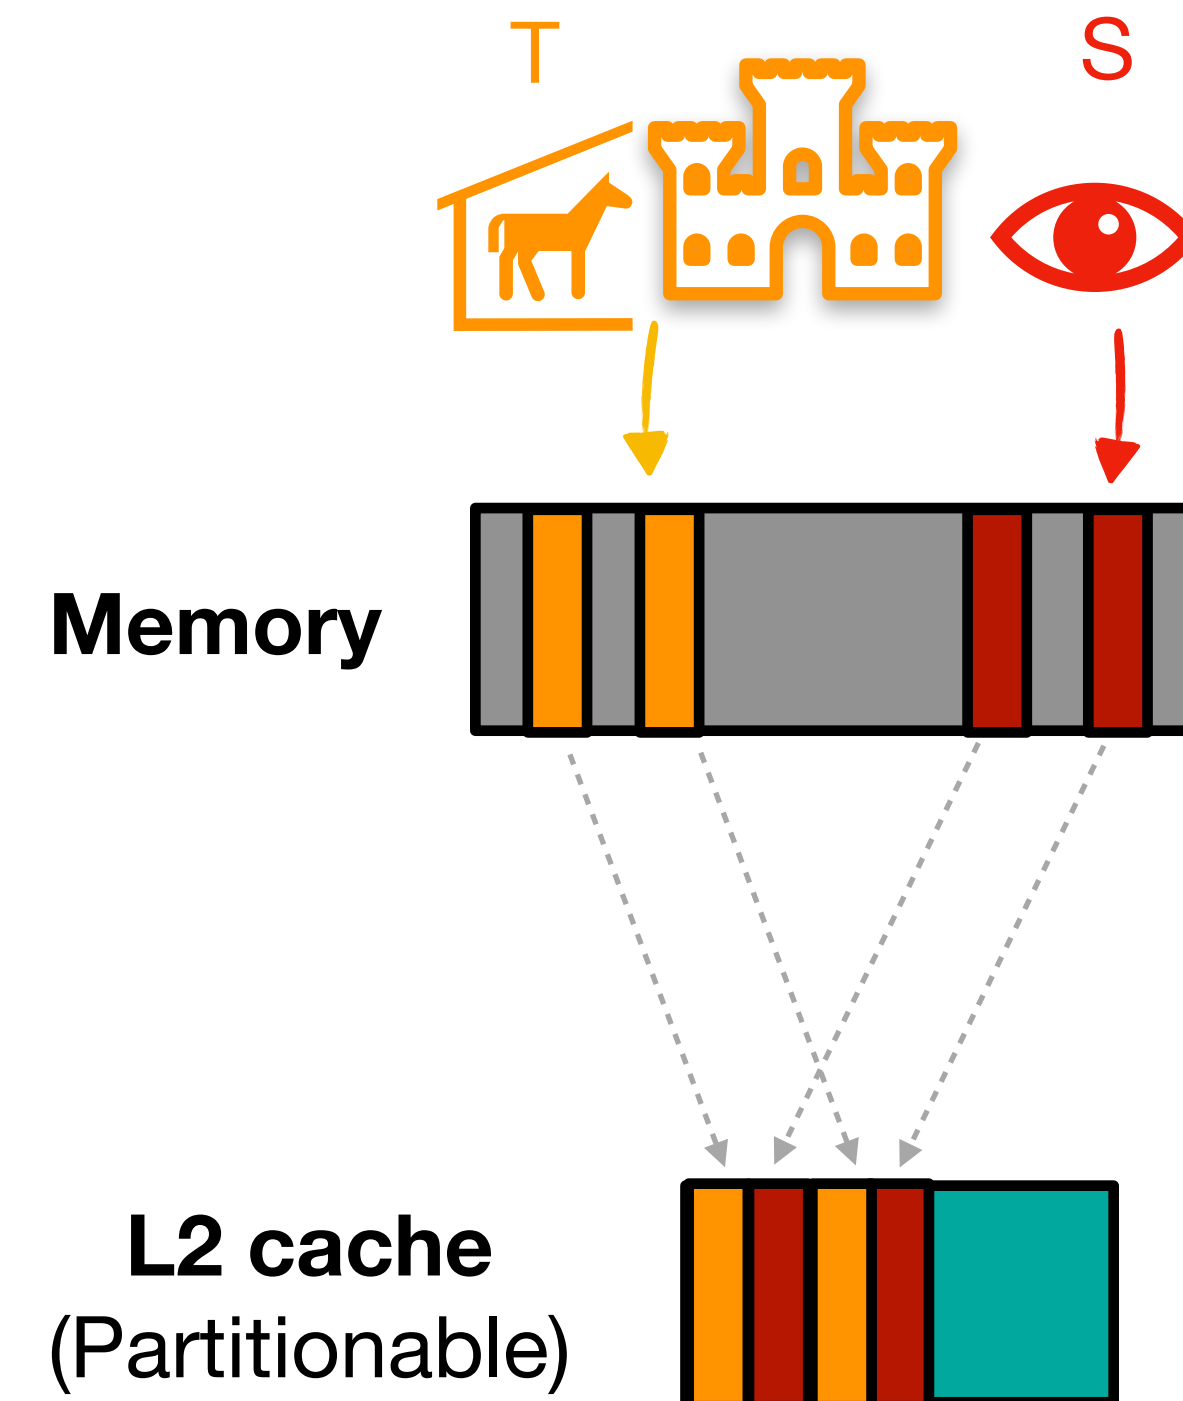
What are time-protected communications?



What are time-protected communications?



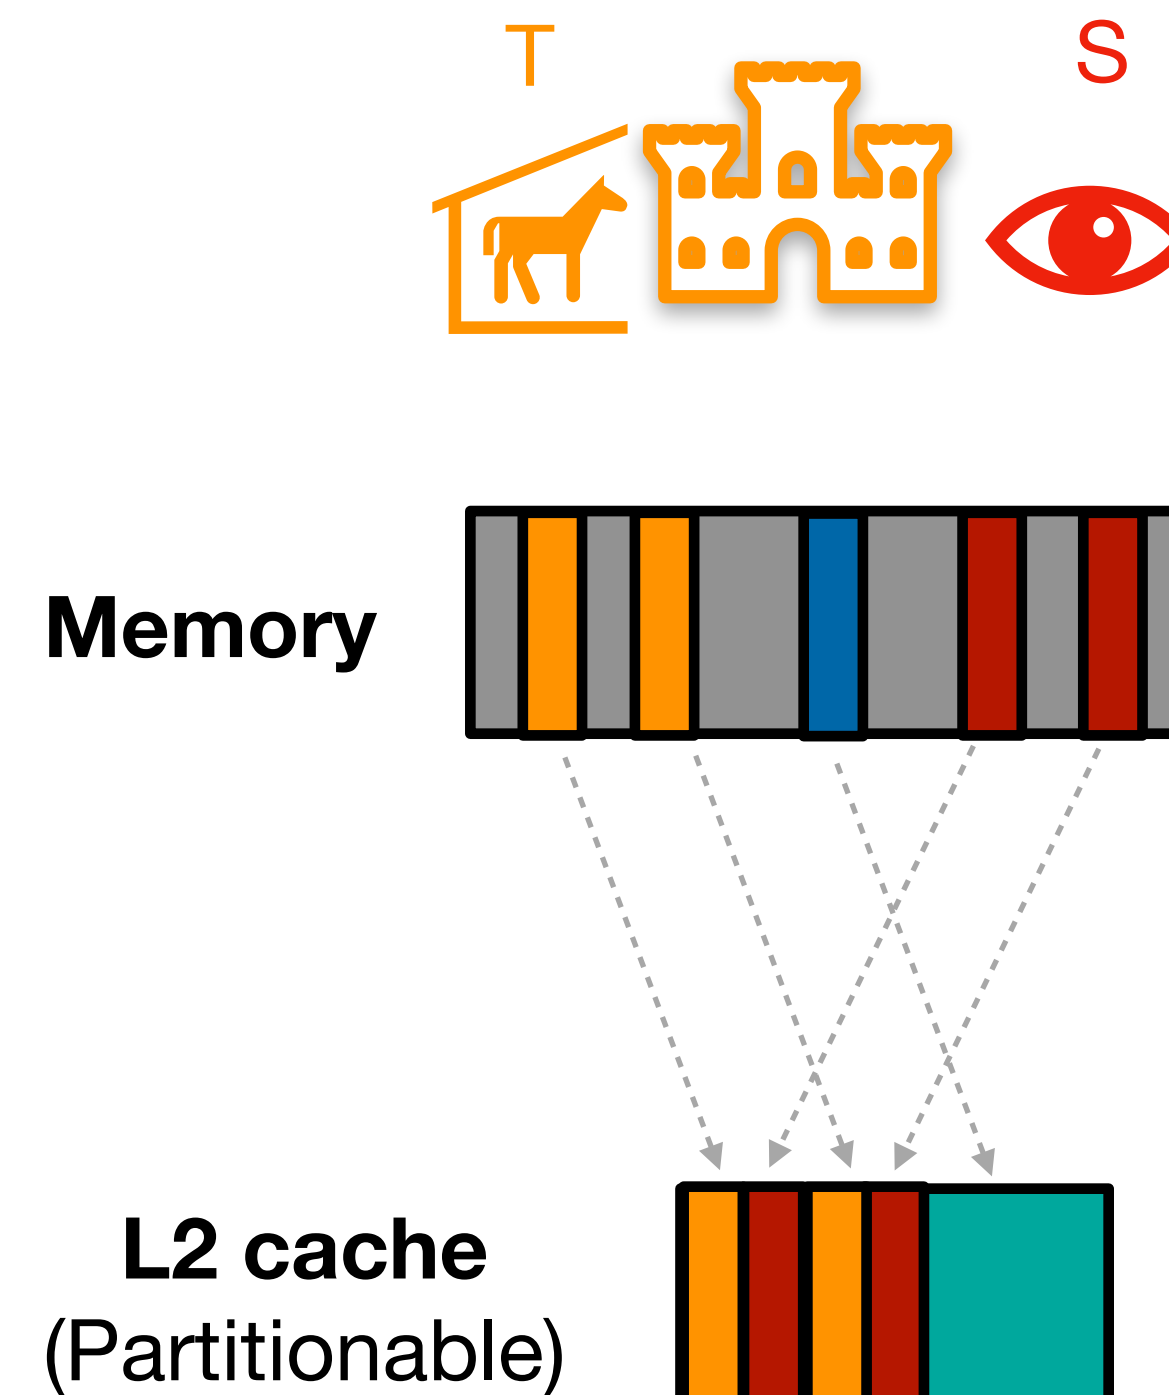
- So far, we've discussed verifying time protection in seL4 with *strict isolation* policies,
i.e. $\forall A B. A \sim /> B$



What are time-protected communications?



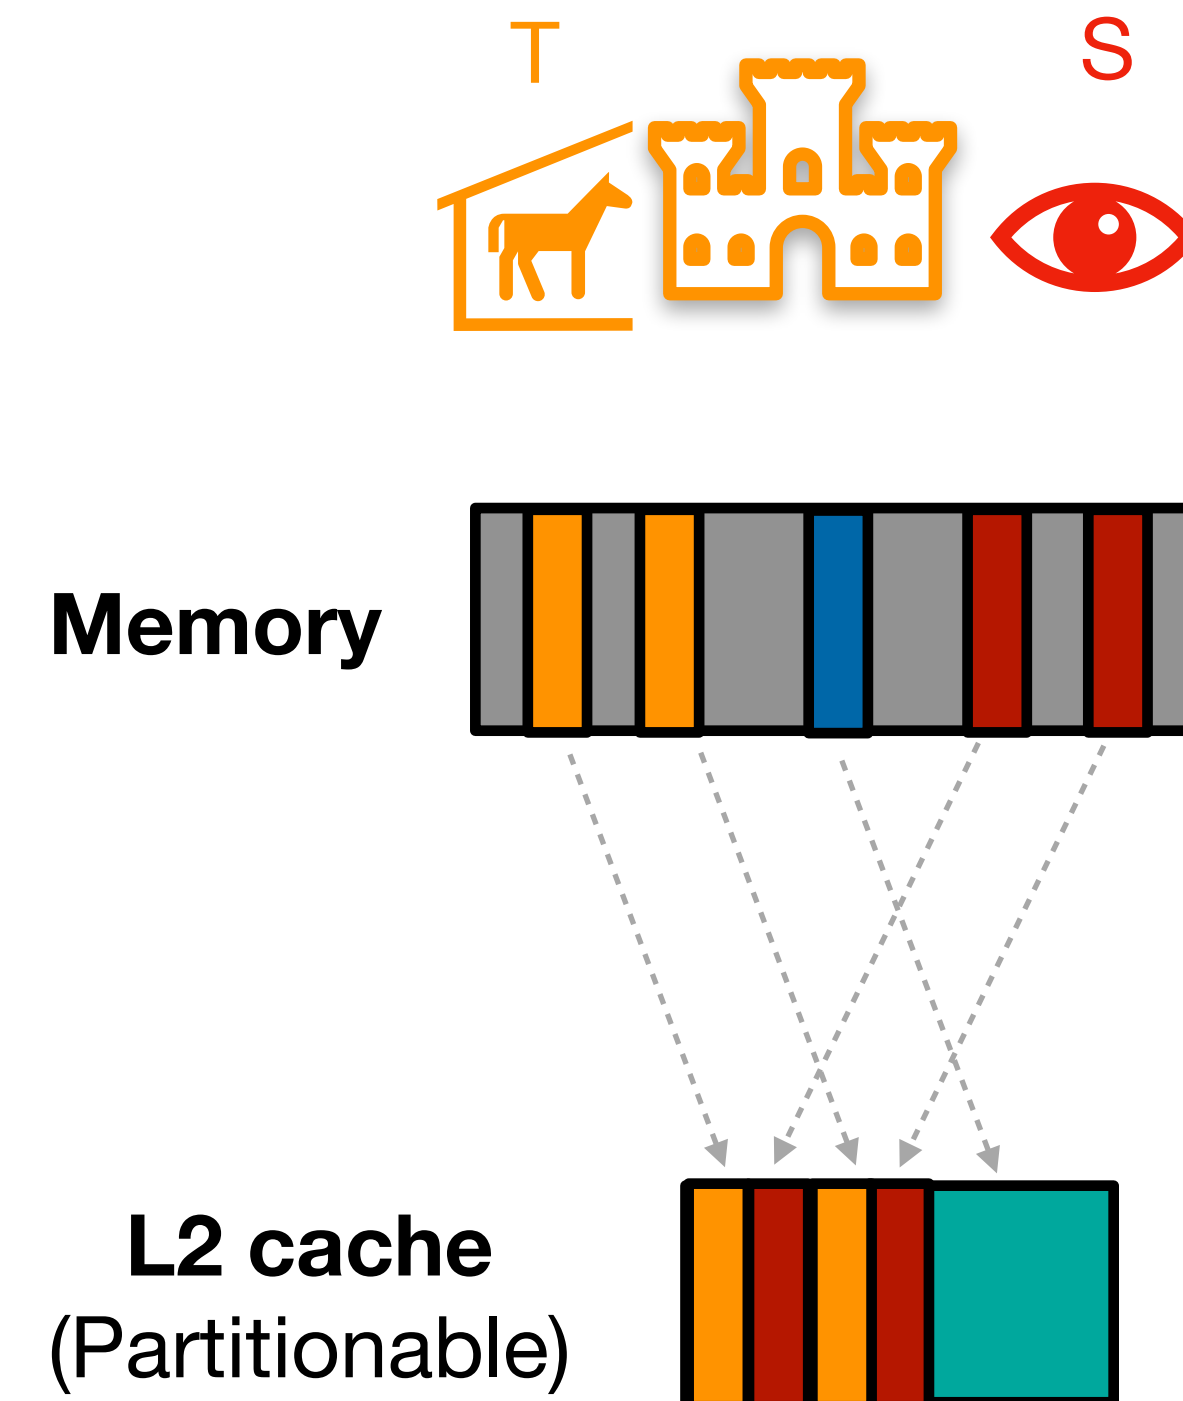
- So far, we've discussed verifying time protection in seL4 with *strict isolation* policies,
i.e. $\forall A B. A \sim /> B$
- Let's consider a *data diode* policy on infowflow via a page of **shared memory**:



What are time-protected communications?



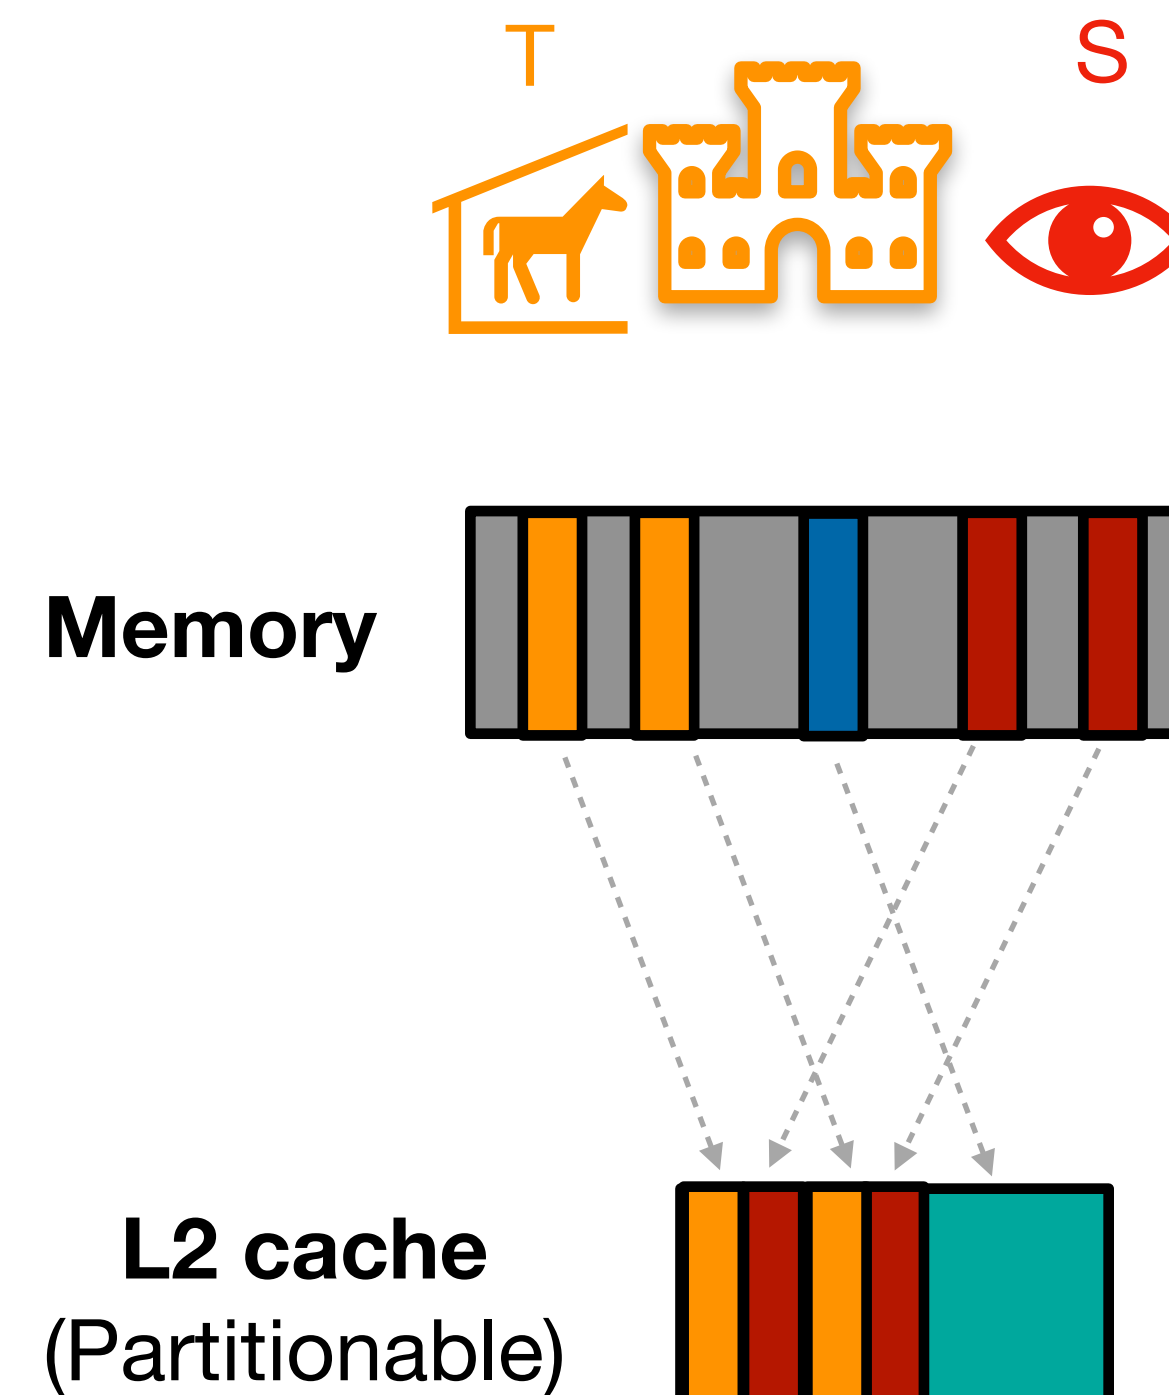
- So far, we've discussed verifying time protection in seL4 with *strict isolation* policies,
i.e. $\forall A B. A \sim /> B$
- Let's consider a *data diode* policy on infowflow via a page of **shared memory**:
 - One direction allowed: $T <\sim S$



What are time-protected communications?



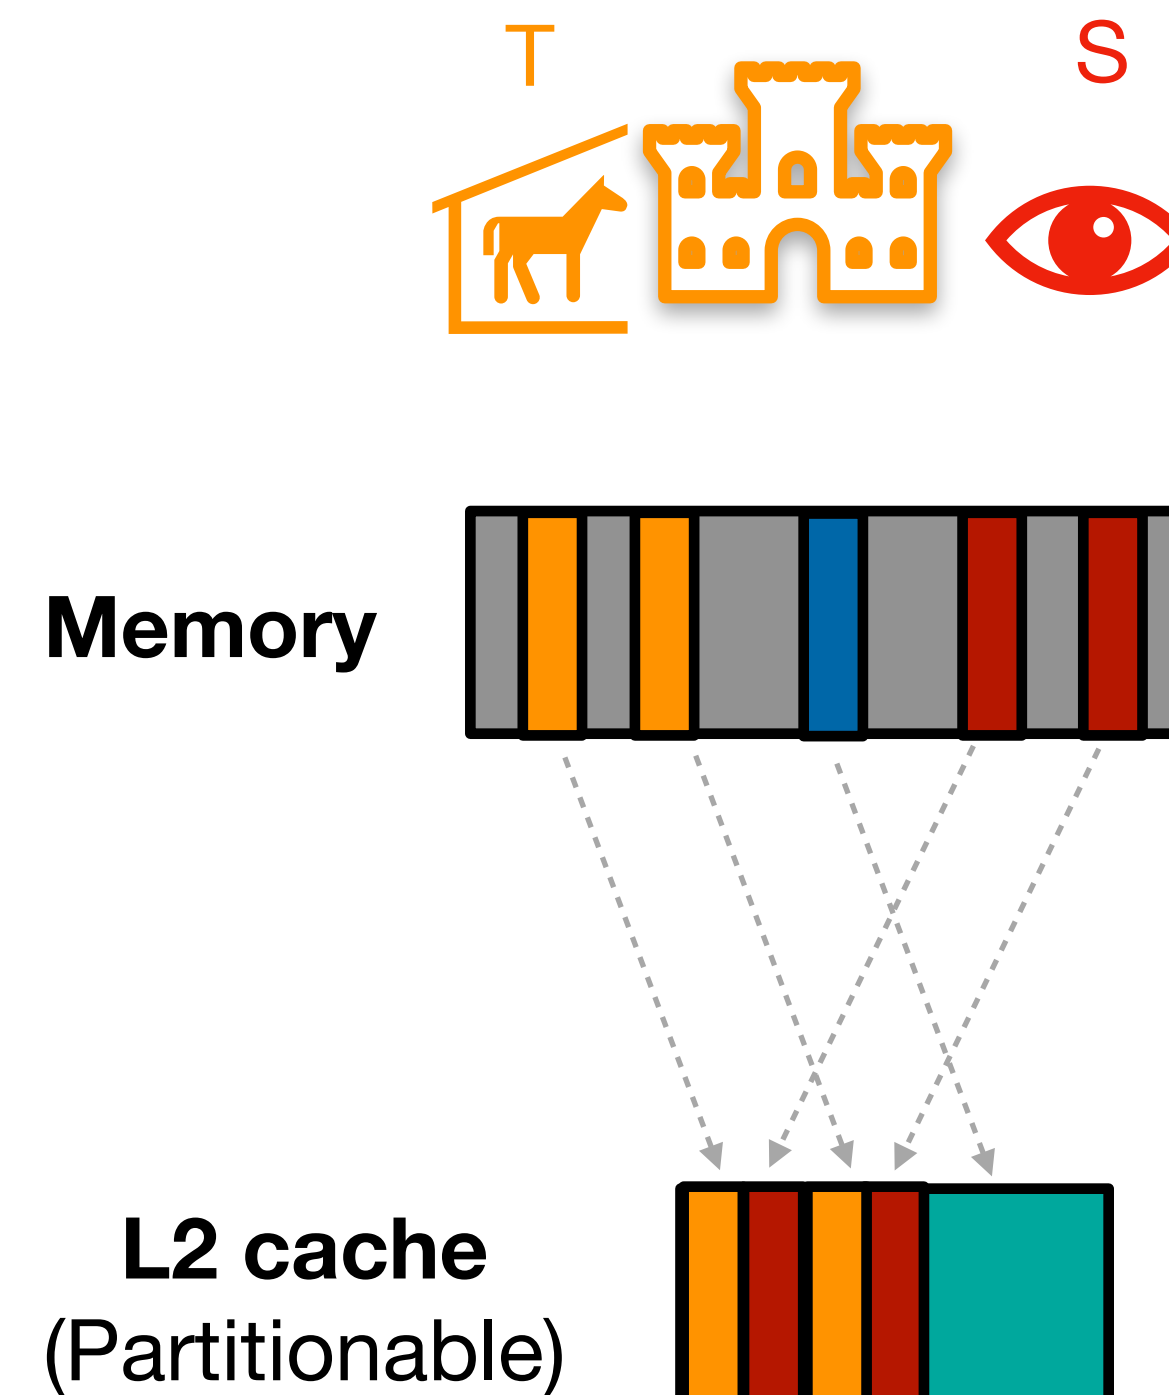
- So far, we've discussed verifying **time protection** in seL4 with *strict isolation* policies,
i.e. $\forall A B. A \sim /> B$
- Let's consider a *data diode* policy on infowflow via a page of **shared memory**:
 - One direction allowed: $T <\sim S$
 - Other direction *disallowed*: $T \sim /> S$



What are time-protected communications?



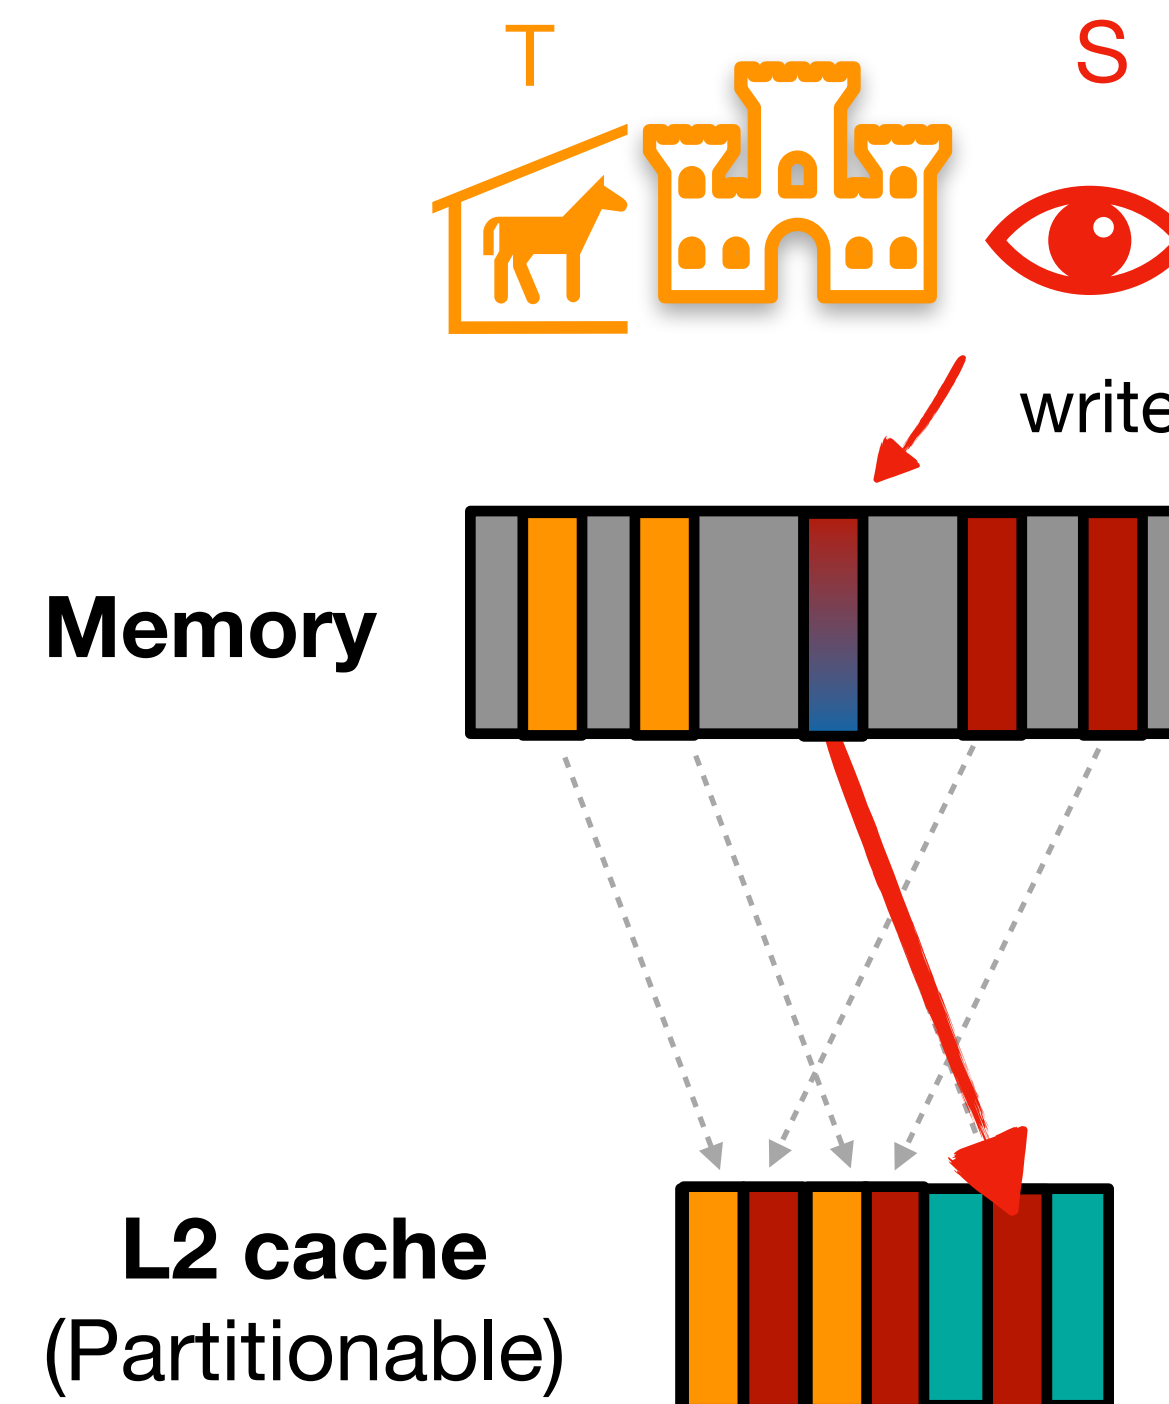
- So far, we've discussed verifying time protection in seL4 with *strict isolation* policies,
i.e. $\forall A B. A \not\sim /> B$
- Let's consider a *data diode* policy on infowflow via a page of **shared memory**:
 - One direction allowed: $T <\sim S$
 - Other direction *disallowed*: $T \not\sim /> S$
- Can we do it with a **partitionable** L2 cache?



What are time-protected communications?



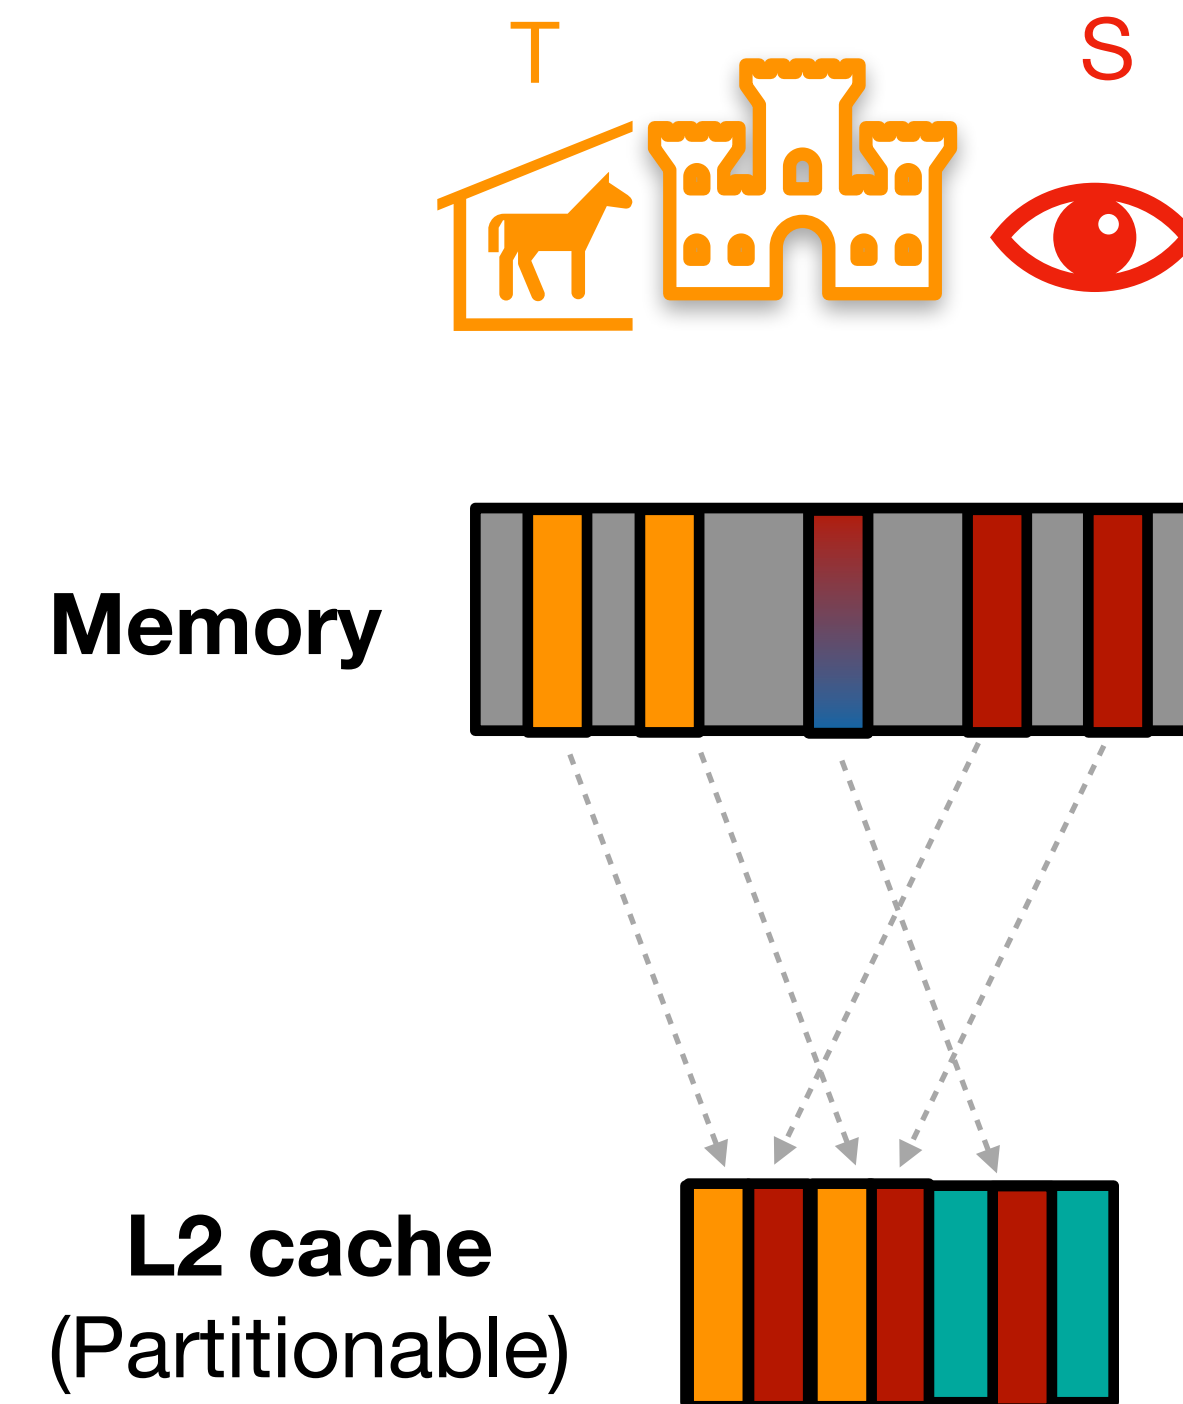
- So far, we've discussed verifying time protection in seL4 with *strict isolation* policies,
i.e. $\forall A B. A \sim /> B$
- Let's consider a *data diode* policy on infowflow via a page of **shared memory**:
 - One direction allowed: $T < \sim S$
 - Other direction *disallowed*: $T \sim /> S$
- Can we do it with a **partitionable** L2 cache?



What are time-protected communications?



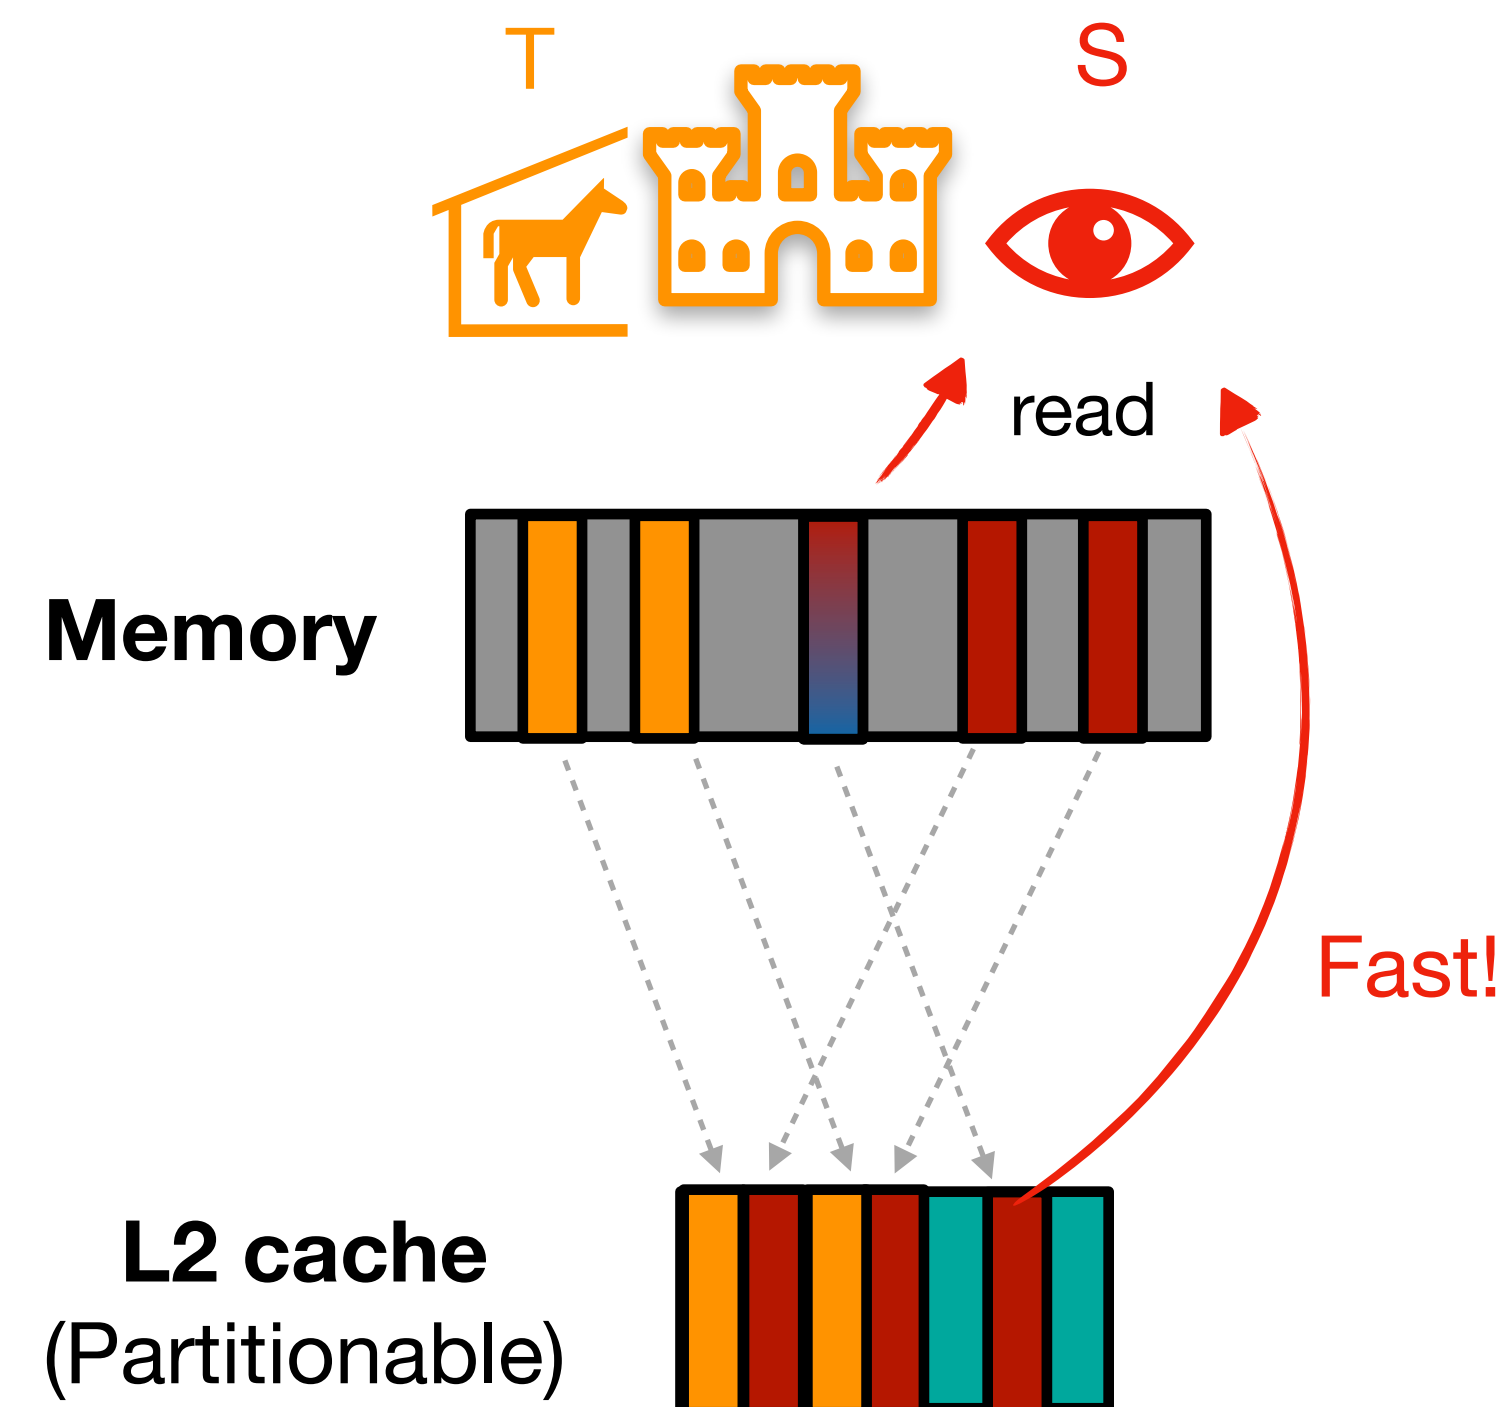
- So far, we've discussed verifying **time protection** in seL4 with *strict isolation* policies,
i.e. $\forall A B. A \sim /> B$
- Let's consider a *data diode* policy on infoflow via a page of **shared memory**:
 - One direction allowed: $T <\sim S$
 - Other direction *disallowed*: $T \sim /> S$
- Can we do it with a **partitionable** L2 cache?



What are time-protected communications?



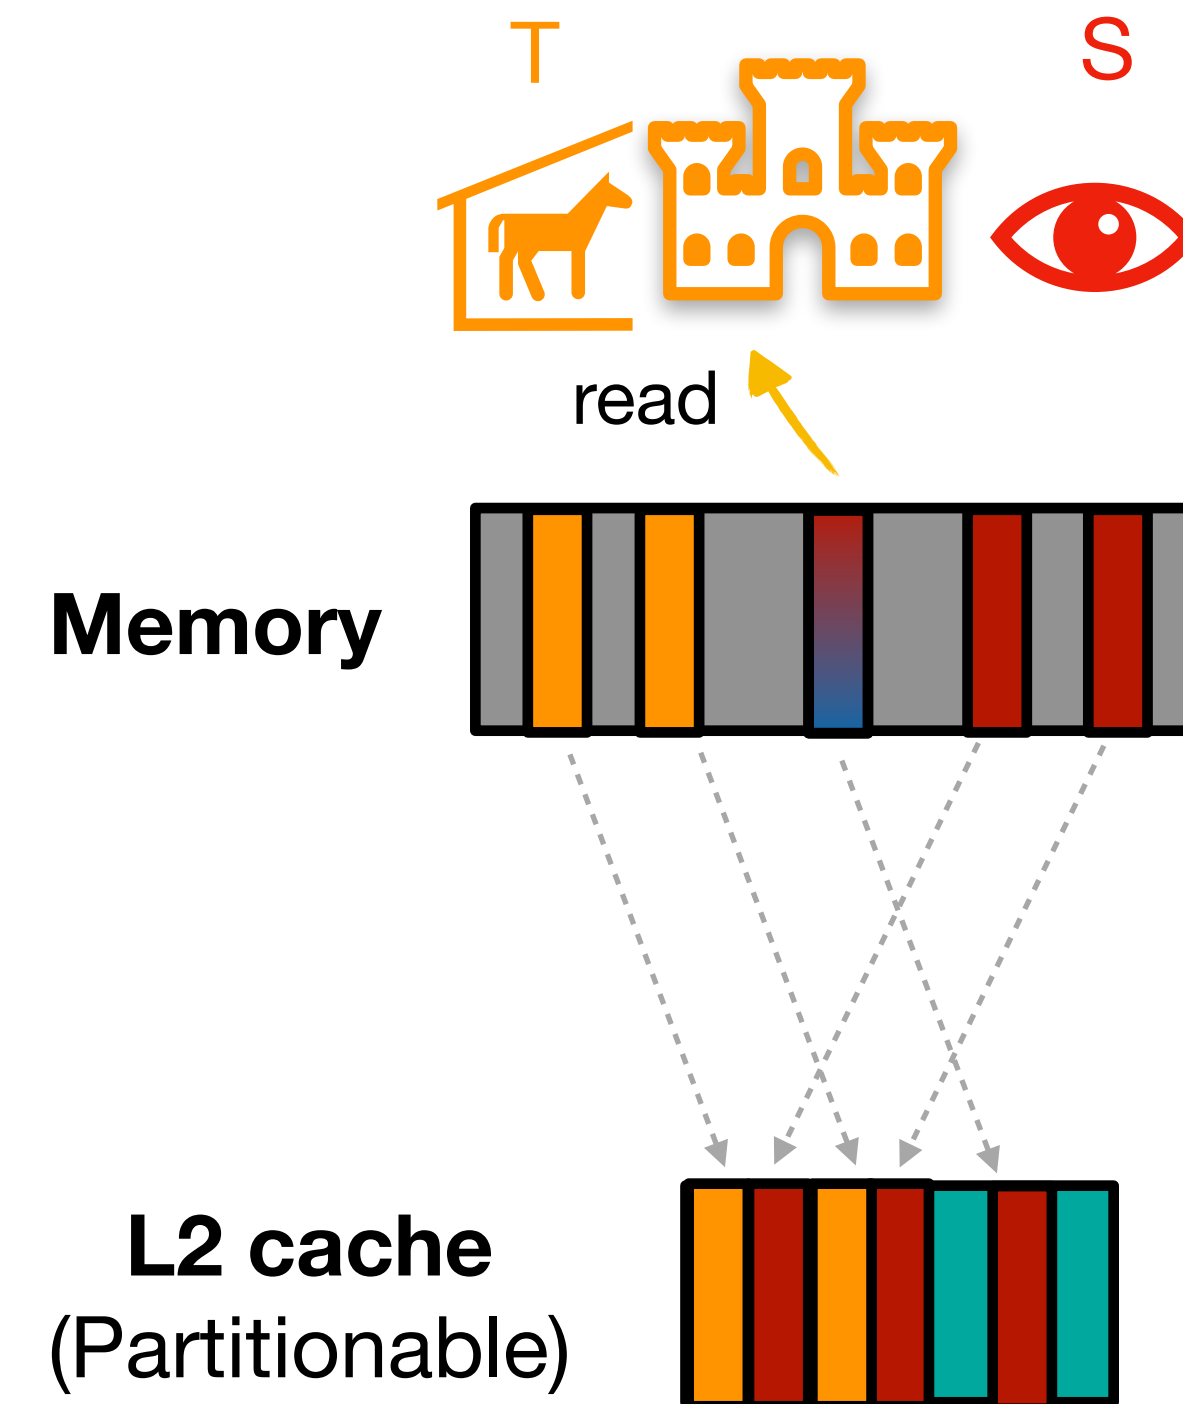
- So far, we've discussed verifying time protection in seL4 with *strict isolation* policies,
i.e. $\forall A B. A \not\sim /> B$
- Let's consider a *data diode* policy on infoflow via a page of **shared memory**:
 - One direction allowed: $T \llsim S$
 - Other direction *disallowed*: $T \not\sim /> S$
- Can we do it with a **partitionable** L2 cache?



What are time-protected communications?



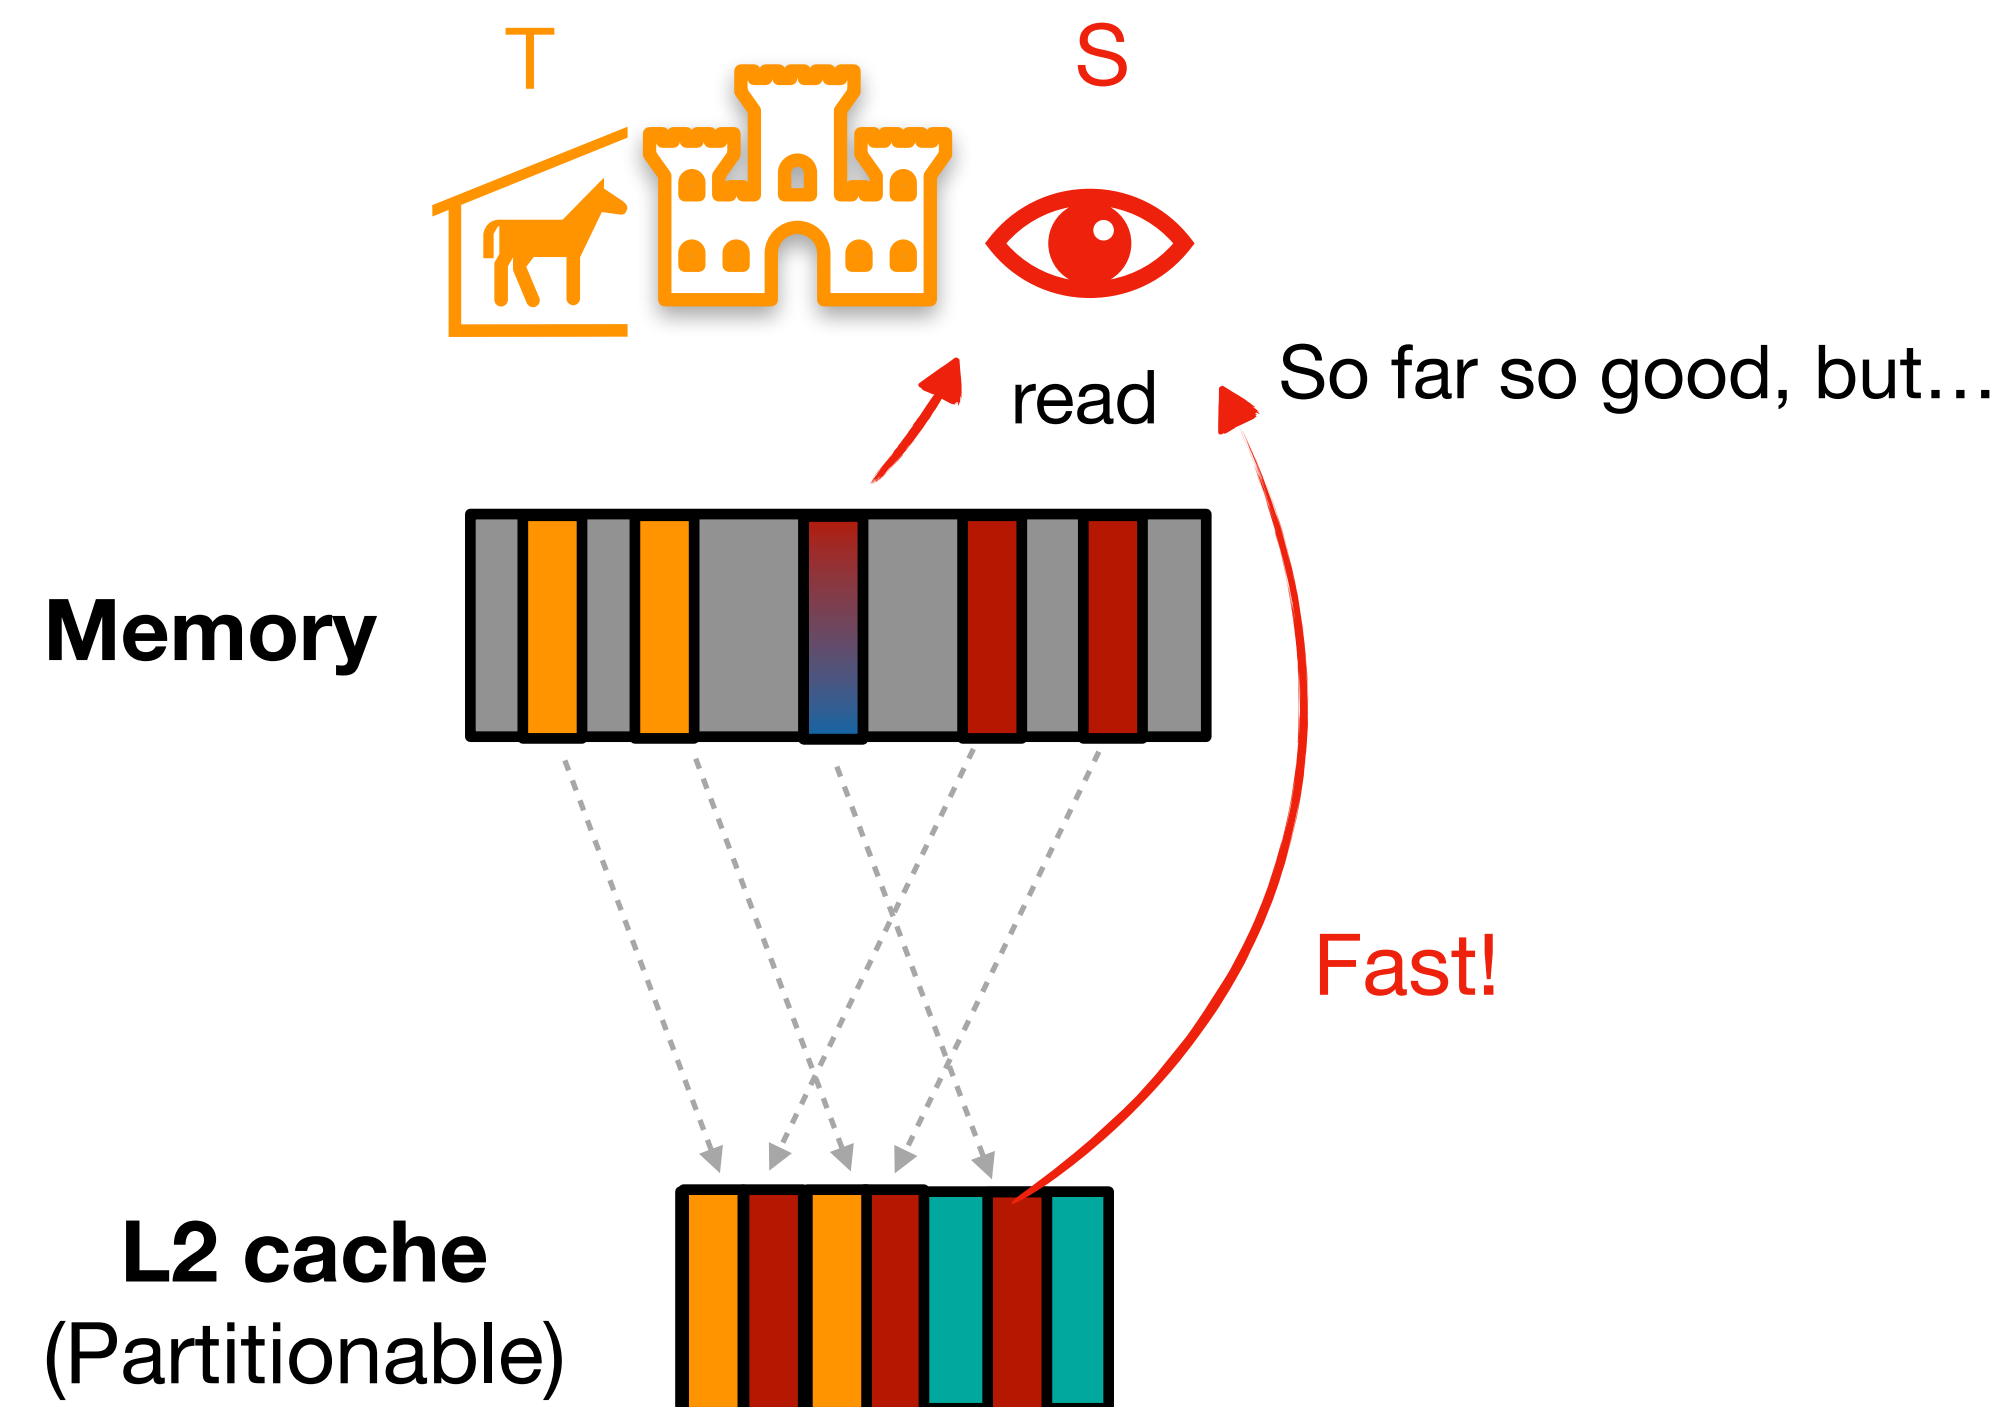
- So far, we've discussed verifying time protection in seL4 with *strict isolation* policies,
i.e. $\forall A B. A \sim /> B$
- Let's consider a *data diode* policy on infowflow via a page of **shared memory**:
 - One direction allowed: $T <\sim S$
 - Other direction *disallowed*: $T \sim /> S$
- Can we do it with a **partitionable** L2 cache?



What are time-protected communications?



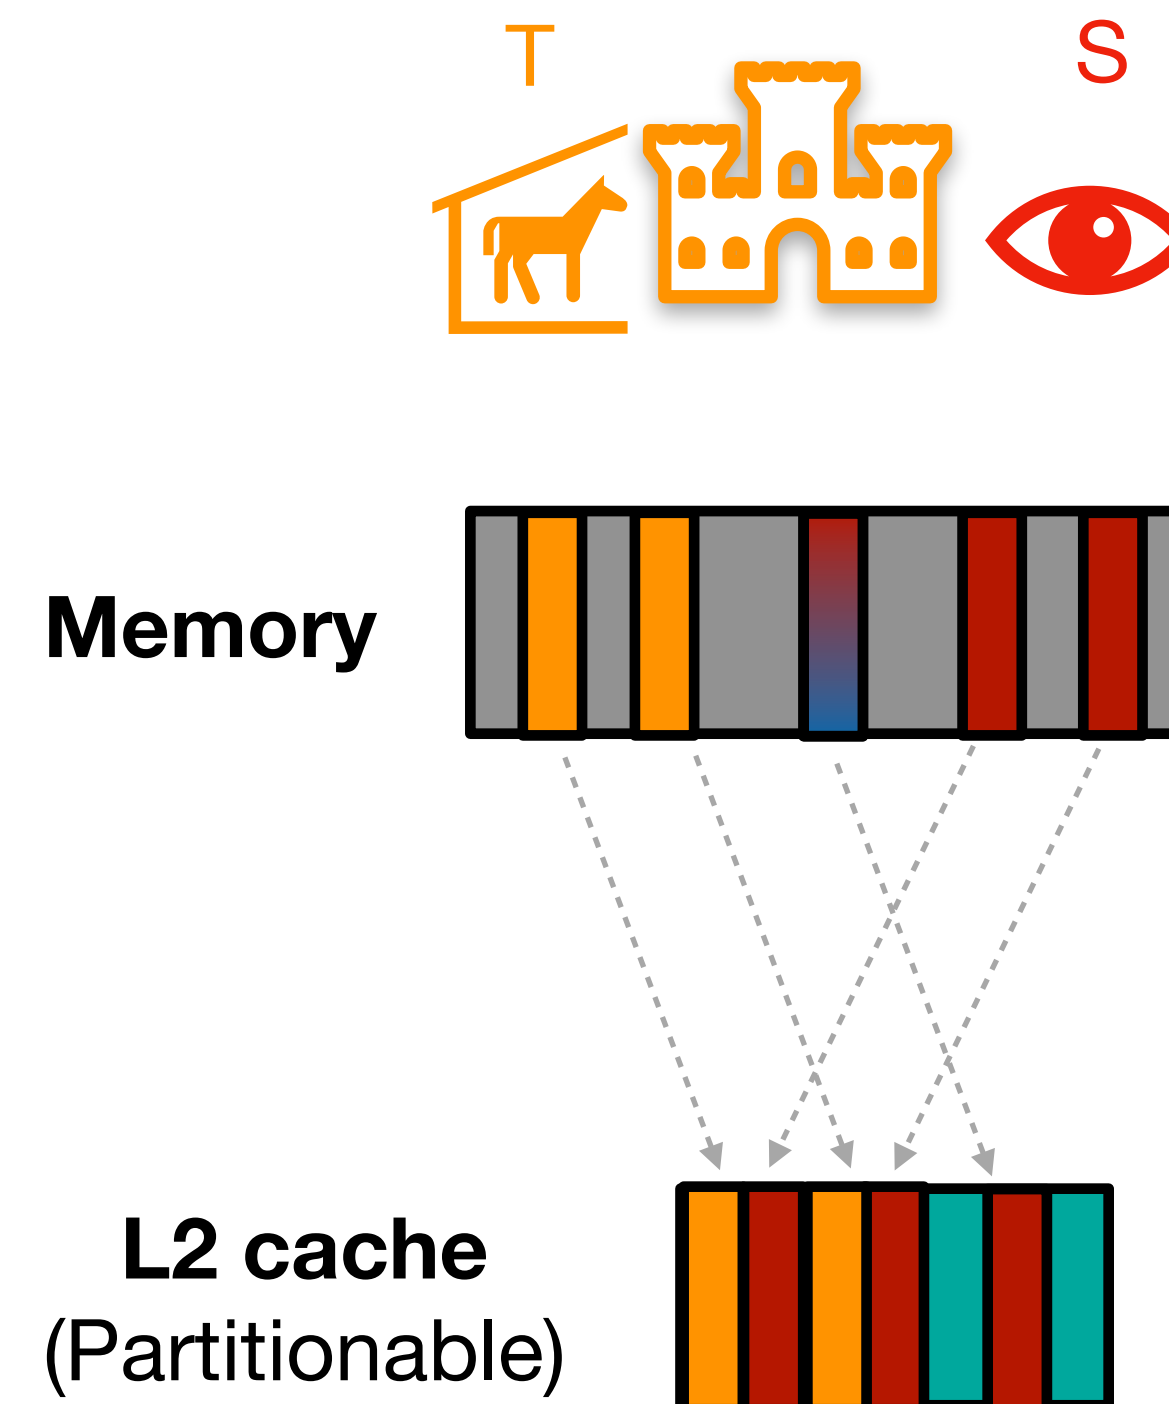
- So far, we've discussed verifying time protection in seL4 with *strict isolation* policies,
i.e. $\forall A B. A \not\sim /> B$
- Let's consider a *data diode* policy on infowflow via a page of **shared memory**:
 - One direction allowed: $T <\sim S$
 - Other direction *disallowed*: $T \not\sim /> S$
- Can we do it with a **partitionable** L2 cache?



What are time-protected communications?



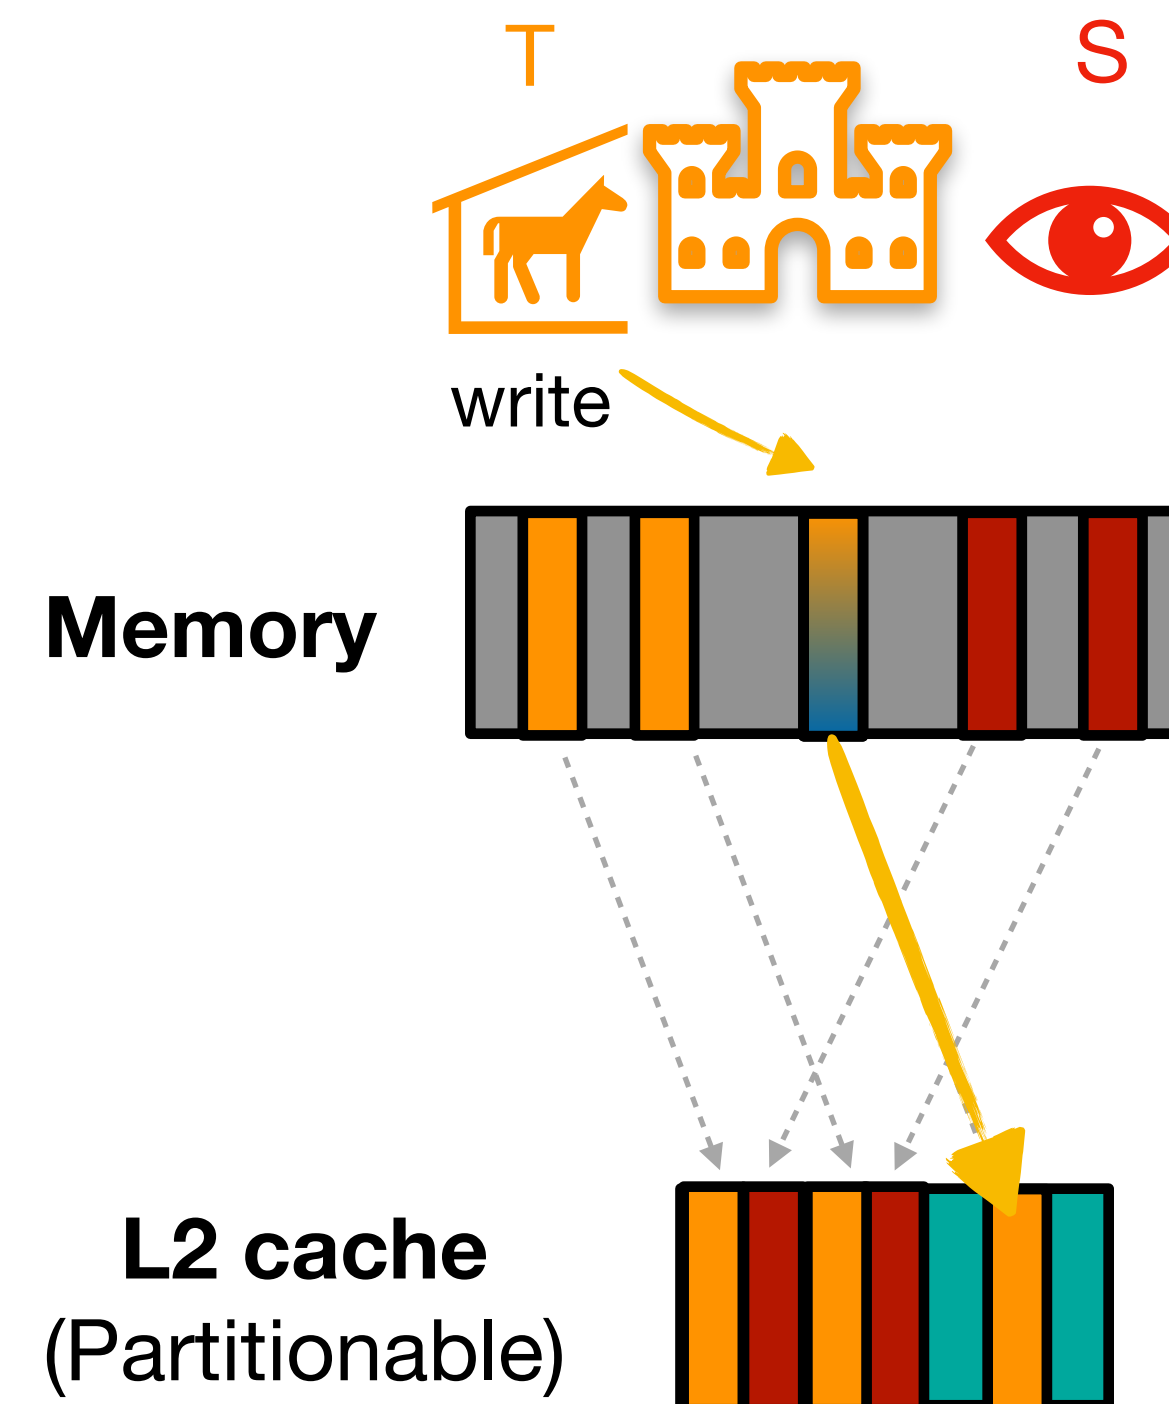
- So far, we've discussed verifying time protection in seL4 with *strict isolation* policies,
i.e. $\forall A B. A \sim /> B$
- Let's consider a *data diode* policy on infowflow via a page of **shared memory**:
 - One direction allowed: $T <\sim S$
 - Other direction *disallowed*: $T \sim /> S$
- Can we do it with a **partitionable** L2 cache?
S sees a timing difference on:



What are time-protected communications?



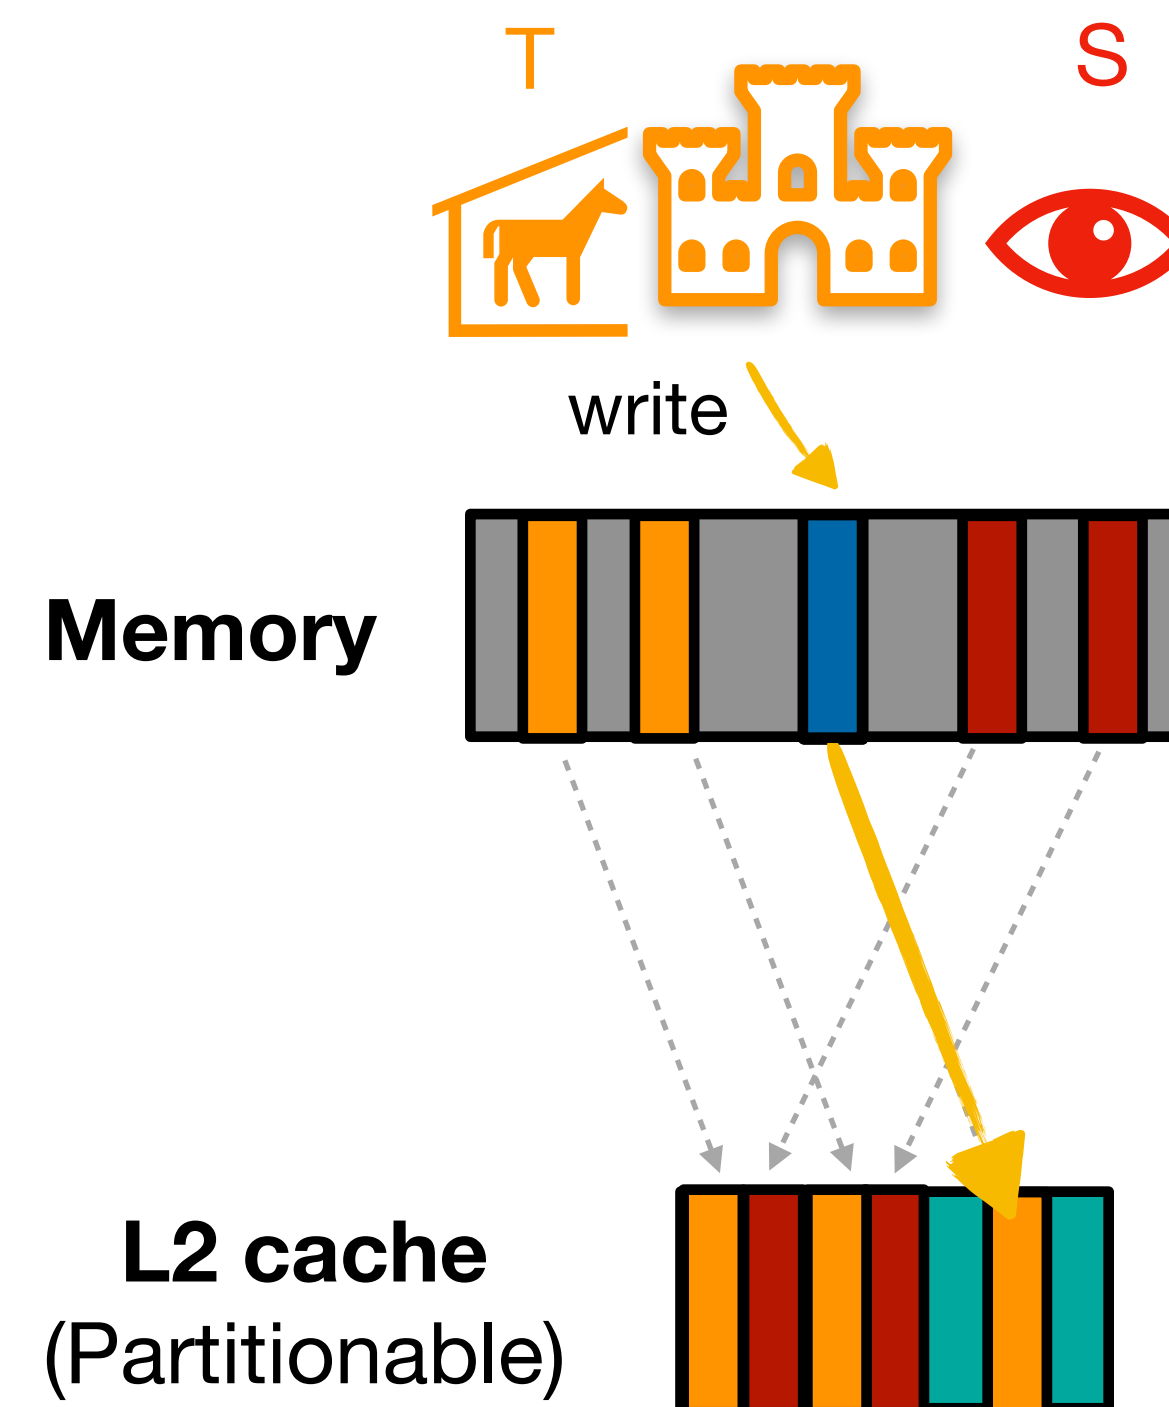
- So far, we've discussed verifying time protection in seL4 with *strict isolation* policies,
i.e. $\forall A B. A \not\sim /> B$
- Let's consider a *data diode* policy on infowflow via a page of **shared memory**:
 - One direction allowed: $T <\sim S$
 - Other direction *disallowed*: $T \not\sim /> S$
- Can we do it with a **partitionable** L2 cache?
S sees a timing difference on:
 - Any writes by T



What are time-protected communications?



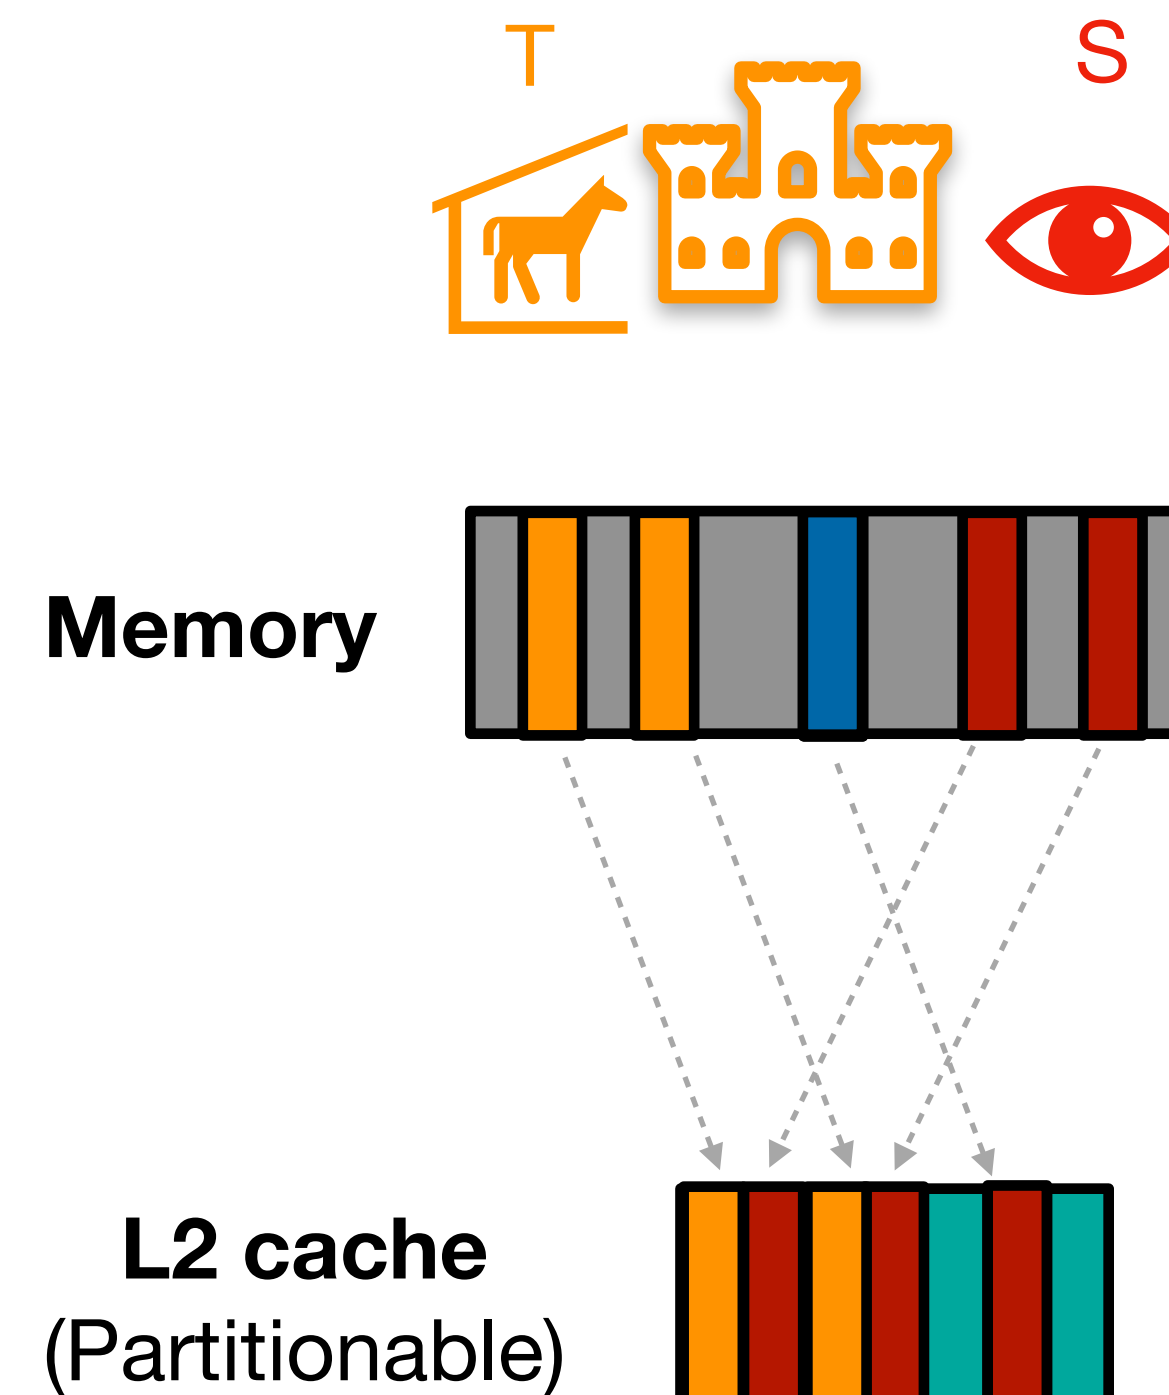
- So far, we've discussed verifying **time protection** in seL4 with *strict isolation* policies,
i.e. $\forall A B. A \not\sim /> B$
- Let's consider a *data diode* policy on infowflow via a page of **shared memory**:
 - One direction allowed: $T <\sim S$
 - Other direction *disallowed*: $T \not\sim /> S$
- Can we do it with a **partitionable** L2 cache?
S sees a timing difference on:
 - Any writes by T (even to wipe the data!)



What are time-protected communications?



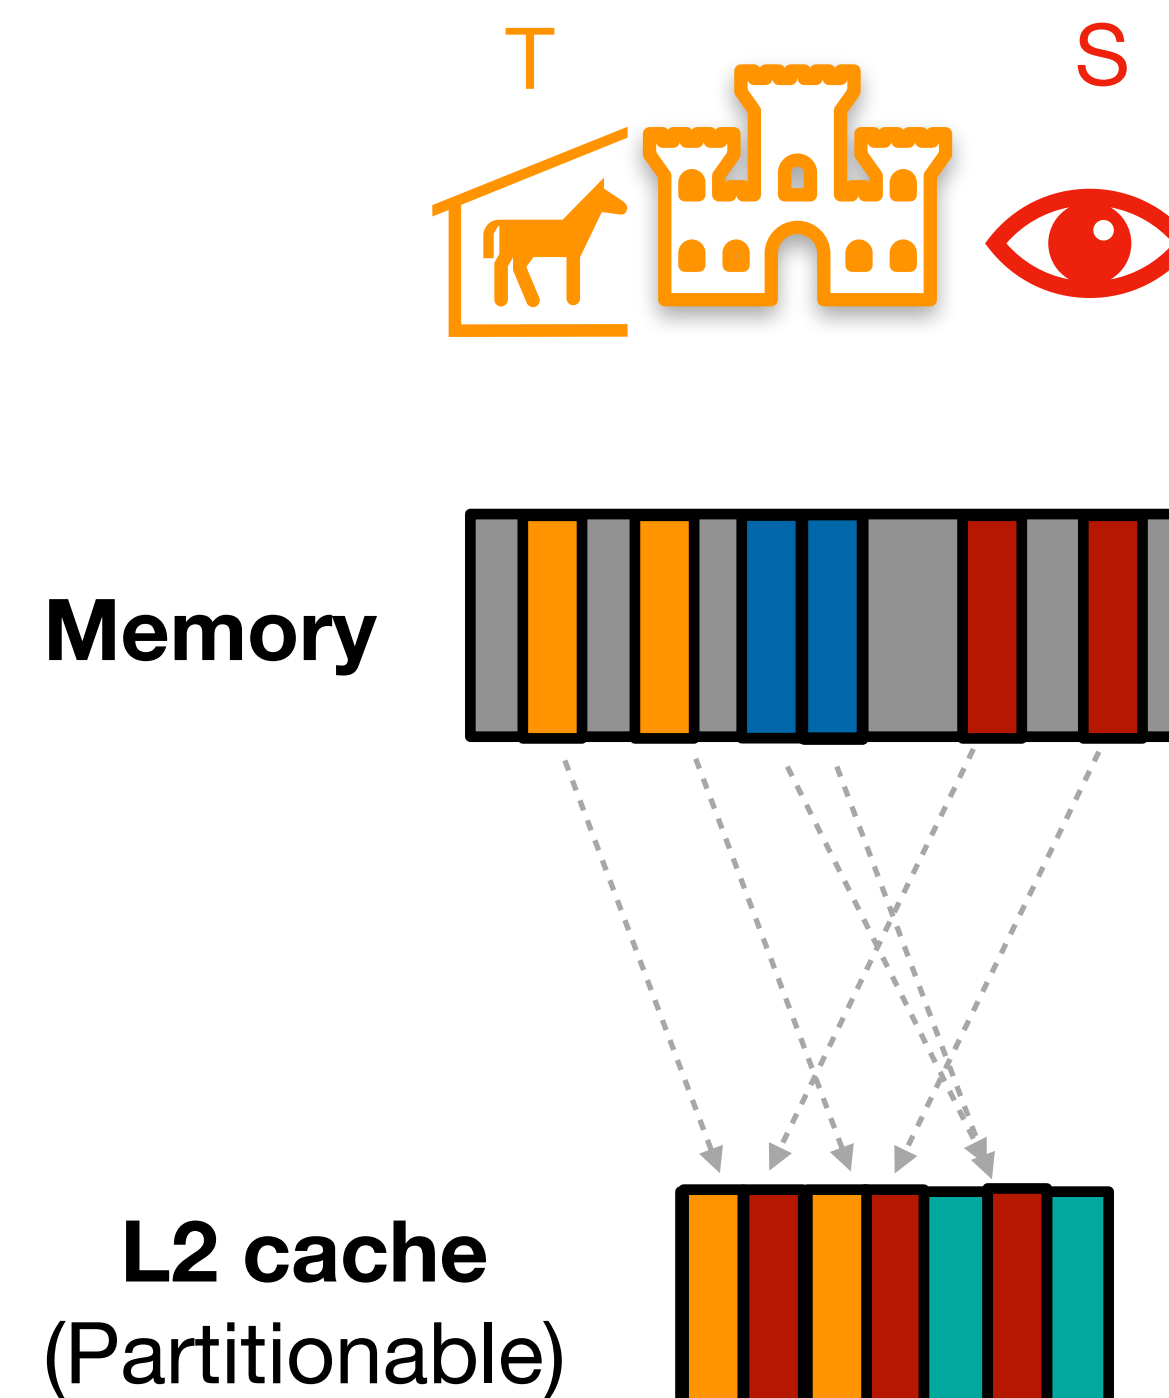
- So far, we've discussed verifying **time protection** in seL4 with *strict isolation* policies,
i.e. $\forall A B. A \sim /> B$
- Let's consider a *data diode* policy on infowflow via a page of **shared memory**:
 - One direction allowed: $T <\sim S$
 - Other direction *disallowed*: $T \sim /> S$
- Can we do it with a **partitionable** L2 cache?
S sees a timing difference on:
 - Any writes by T (even to wipe the data!)
 - Any reads by T that evict lines (many per page!)



What are time-protected communications?



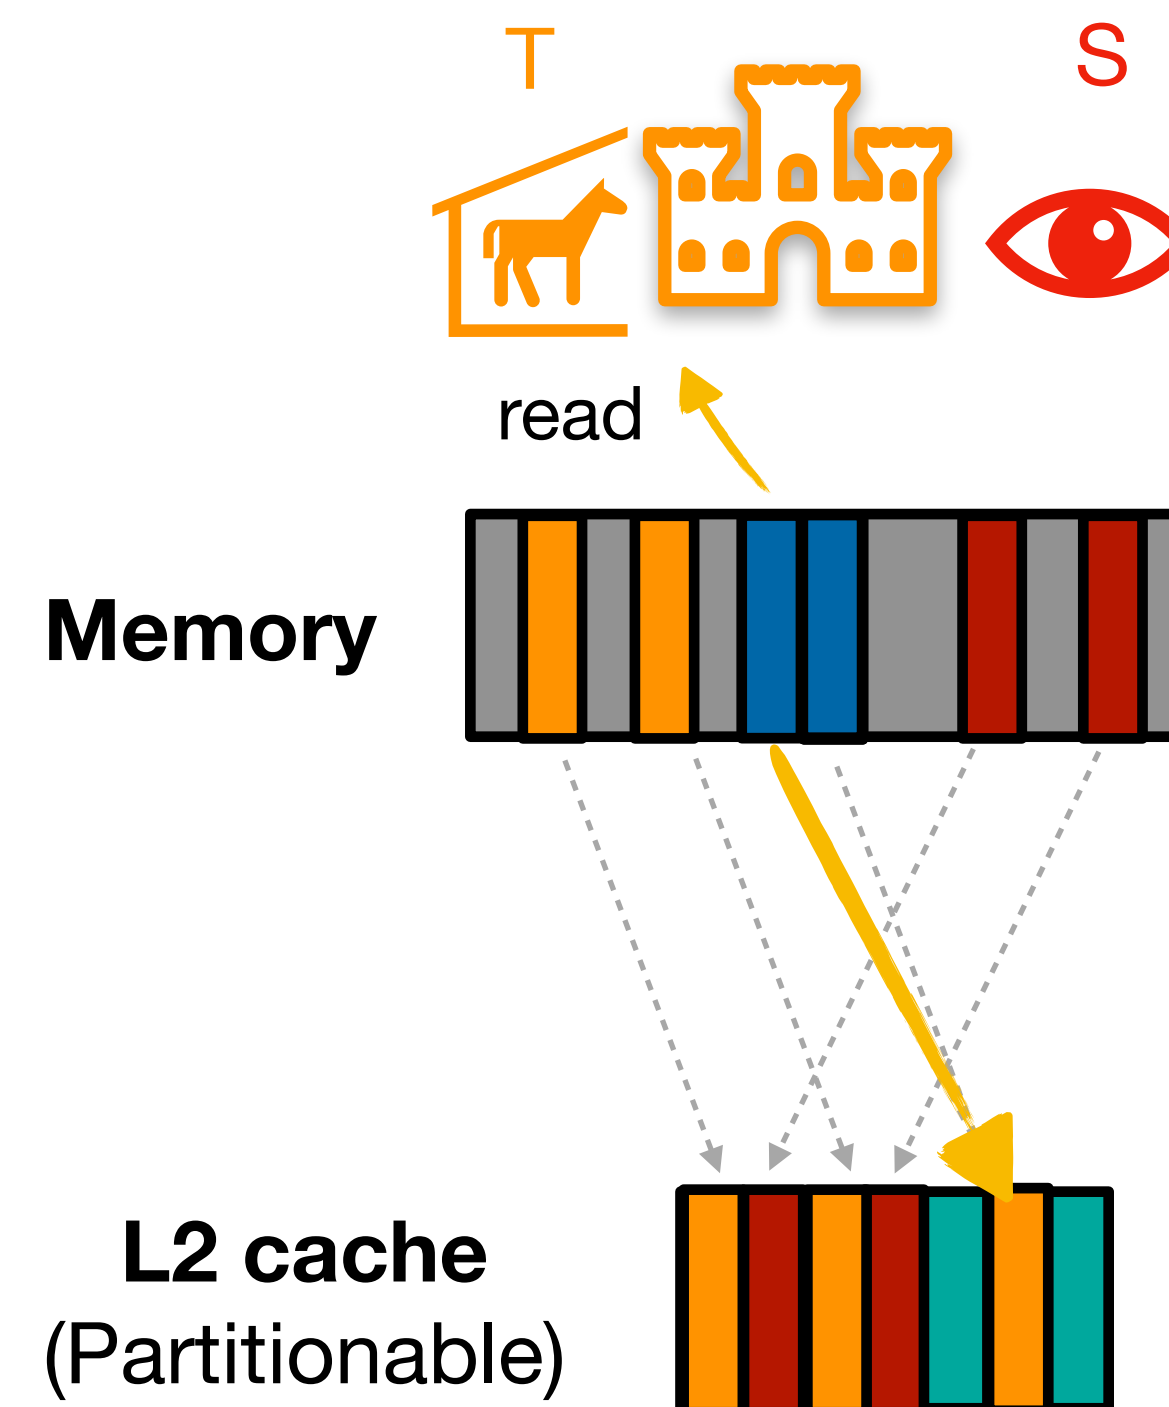
- So far, we've discussed verifying **time protection** in seL4 with *strict isolation* policies,
i.e. $\forall A B. A \sim /> B$
- Let's consider a *data diode* policy on infowflow via a page of **shared memory**:
 - One direction allowed: $T <\sim S$
 - Other direction *disallowed*: $T \sim /> S$
- Can we do it with a **partitionable** L2 cache?
S sees a timing difference on:
 - Any writes by T (even to wipe the data!)
 - Any reads by T that evict lines (many per page!)



What are time-protected communications?



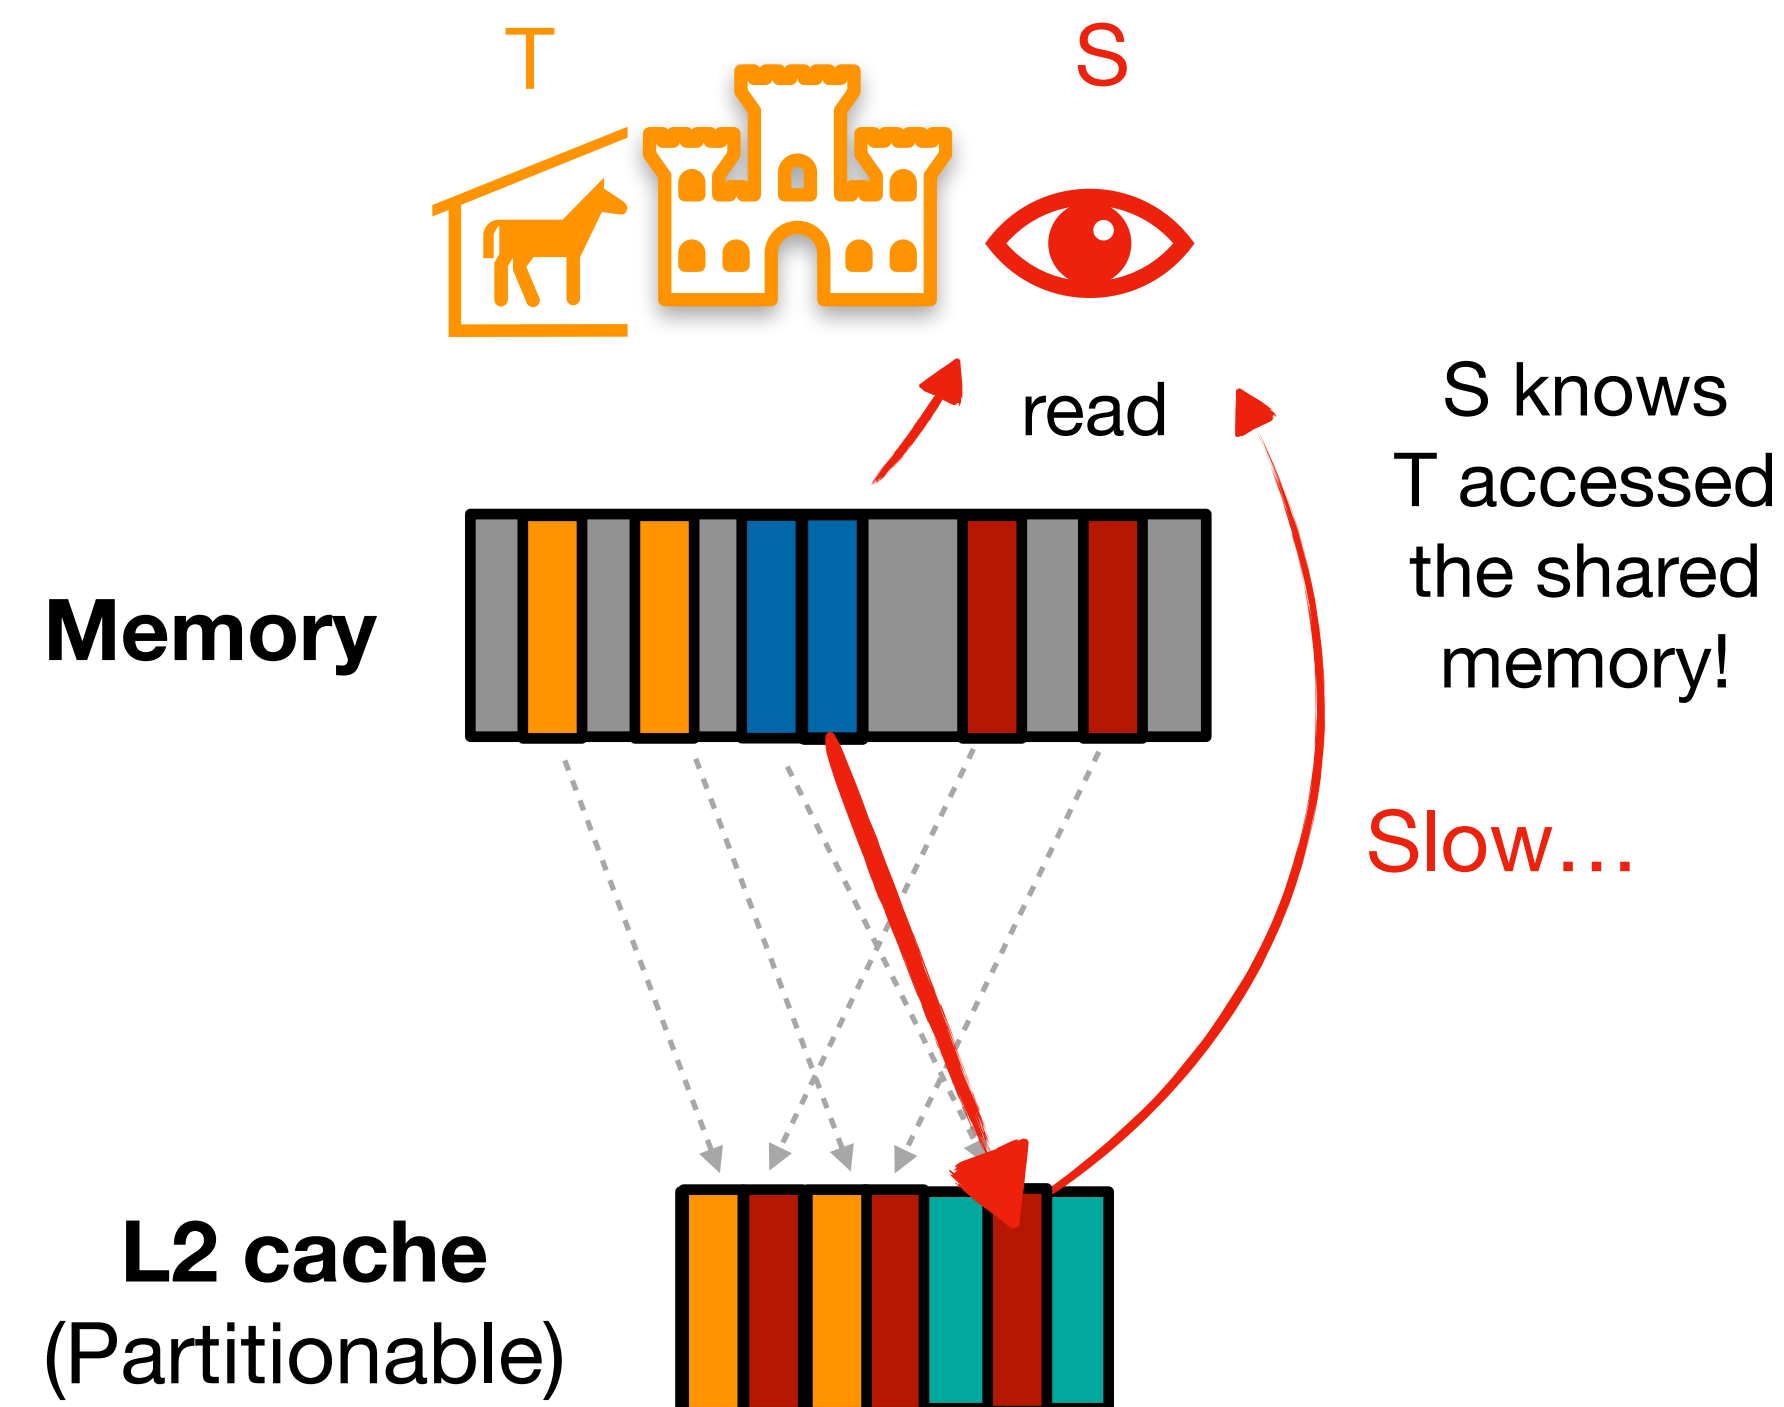
- So far, we've discussed verifying **time protection** in seL4 with *strict isolation* policies,
i.e. $\forall A B. A \sim /> B$
- Let's consider a *data diode* policy on infowflow via a page of **shared memory**:
 - One direction allowed: $T <\sim S$
 - Other direction *disallowed*: $T \sim /> S$
- Can we do it with a **partitionable** L2 cache?
S sees a timing difference on:
 - Any writes by T (even to wipe the data!)
 - Any reads by T that evict lines (many per page!)



What are time-protected communications?



- So far, we've discussed verifying **time protection** in seL4 with *strict isolation* policies,
i.e. $\forall A B. A \not\sim /> B$
- Let's consider a *data diode* policy on infowflow via a page of **shared memory**:
 - One direction allowed: $T <\sim S$
 - Other direction *disallowed*: $T \not\sim /> S$
- Can we do it with a **partitionable** L2 cache?
S sees a timing difference on:
 - Any writes by T (even to wipe the data!)
 - Any reads by T that evict lines (many per page!)



What are time-protected communications?



- So far, we've discussed verifying **time protection** in seL4 with *strict isolation* policies,

i.e. $\forall A B. A \not\sim /> B$

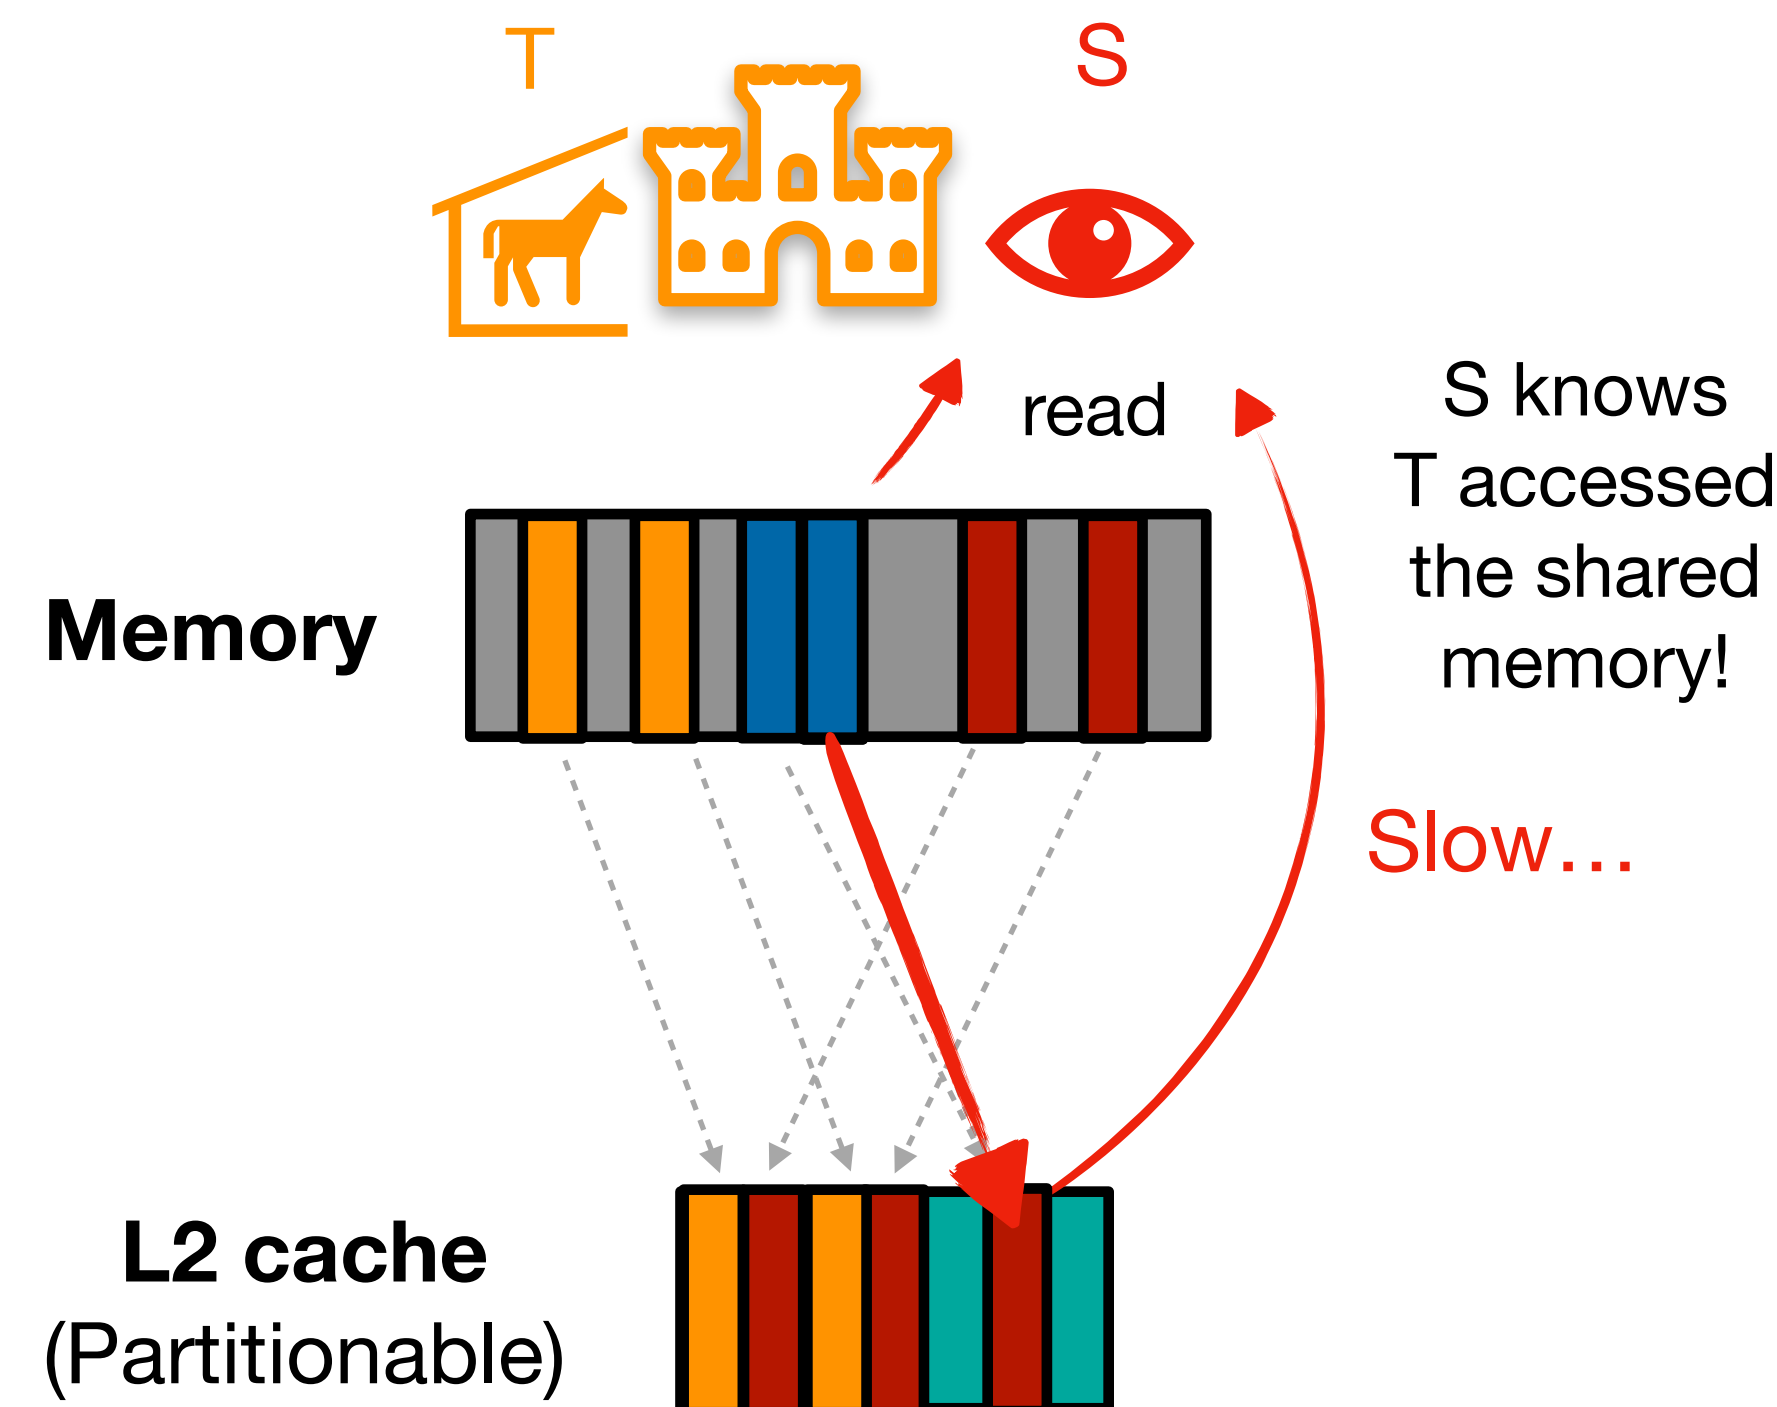
- Let's consider a *data diode* policy on infowflow via a page of **shared memory**:

- One direction allowed: $T <\sim S$
- Other direction *disallowed*: $T \not\sim /> S$

- Can we do it with a **partitionable** L2 cache?
S sees a timing difference on:

- Any writes by T (even to wipe the data!)
- Any reads by T that evict lines (many per page!)

We need **targeted flush** for partitionable L2 caches to enforce $T \not\sim /> S$ while allowing $T <\sim S$



What are time-protected communications?



- So far, we've discussed verifying **time protection** in seL4 with *strict isolation* policies,

i.e. $\forall A B. A \sim /> B$

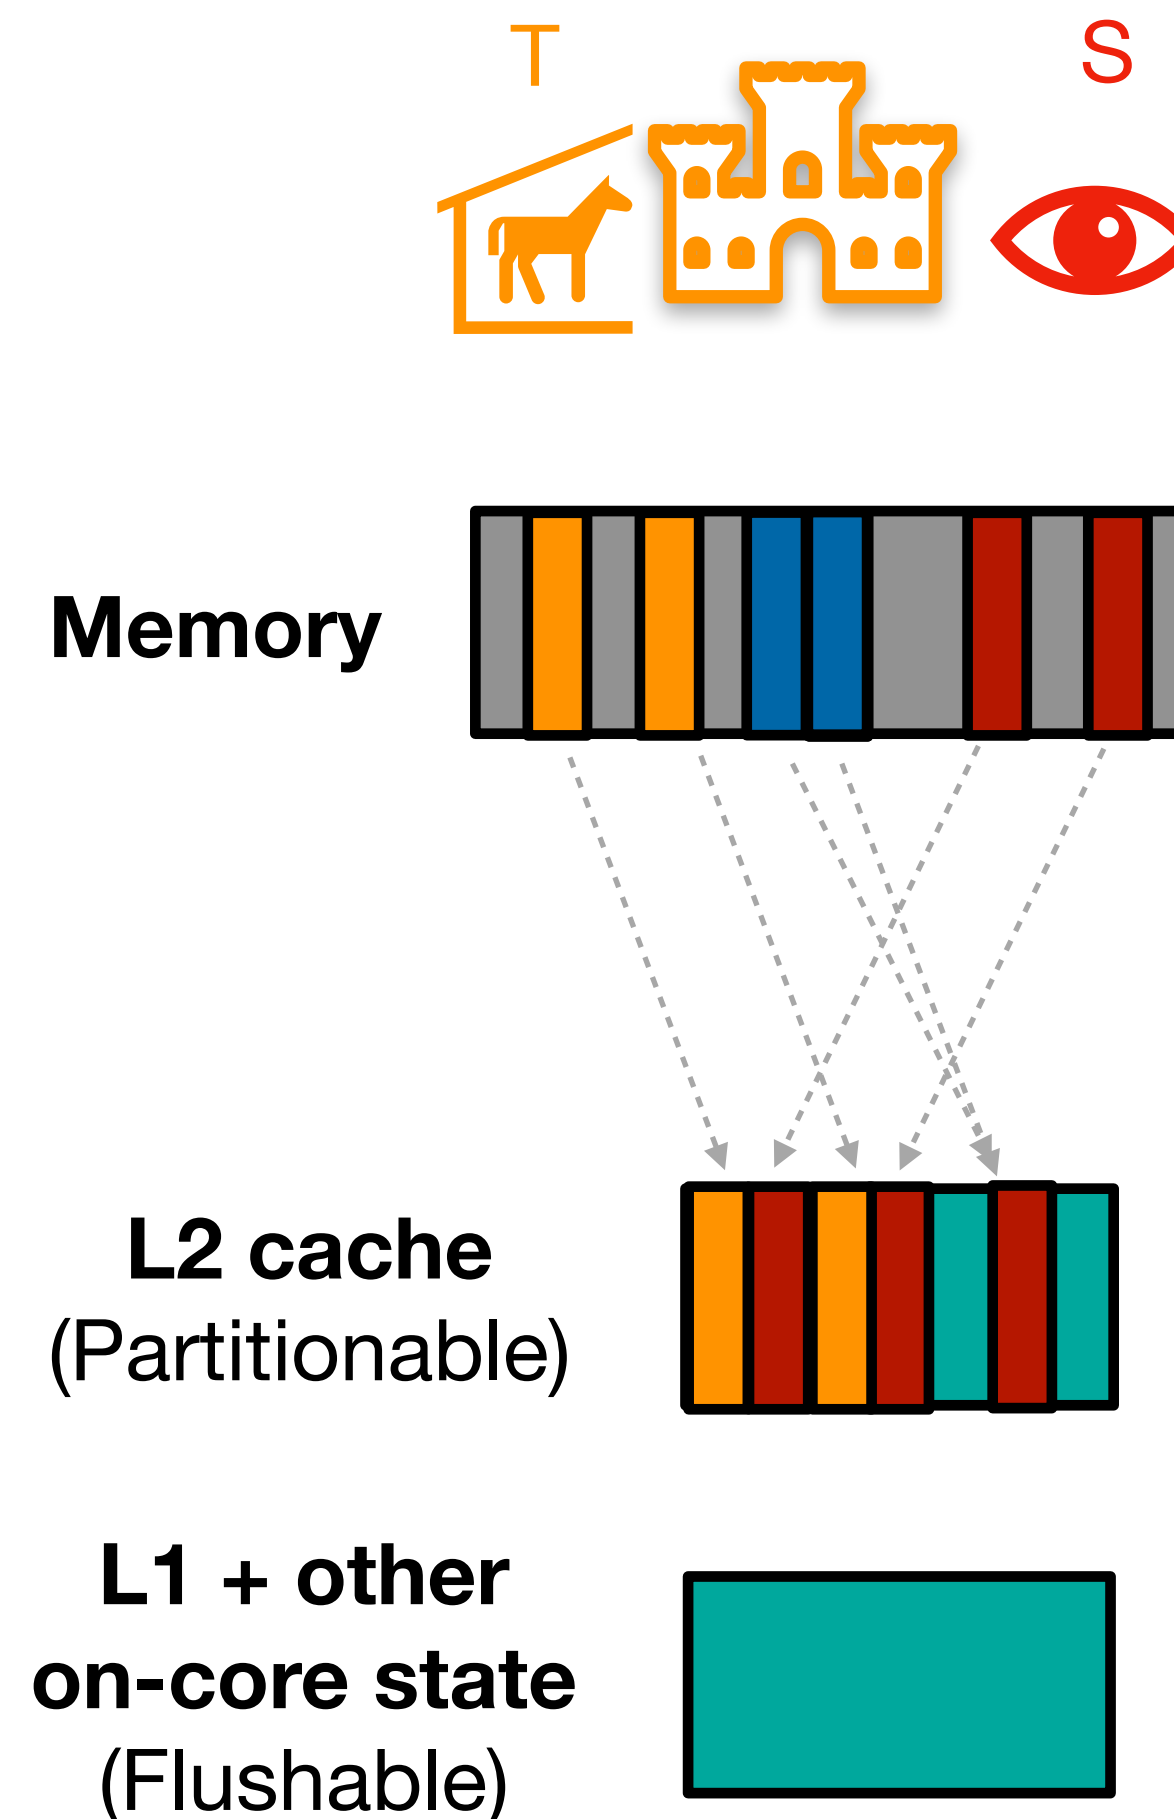
- Let's consider a *data diode* policy on infowflow via a page of **shared memory**:

- One direction allowed: $T < \sim S$
- Other direction *disallowed*: $T \sim /> S$

- Can we do it with a **partitionable** L2 cache?
S sees a timing difference on:

- Any writes by T (even to wipe the data!)
- Any reads by T that evict lines (many per page!)

We need *targeted flush* for partitionable L2 caches to enforce $T \sim /> S$ while allowing $T < \sim S$



What are time-protected communications?



- So far, we've discussed verifying **time protection** in seL4 with *strict isolation* policies,

i.e. $\forall A B. A \sim /> B$

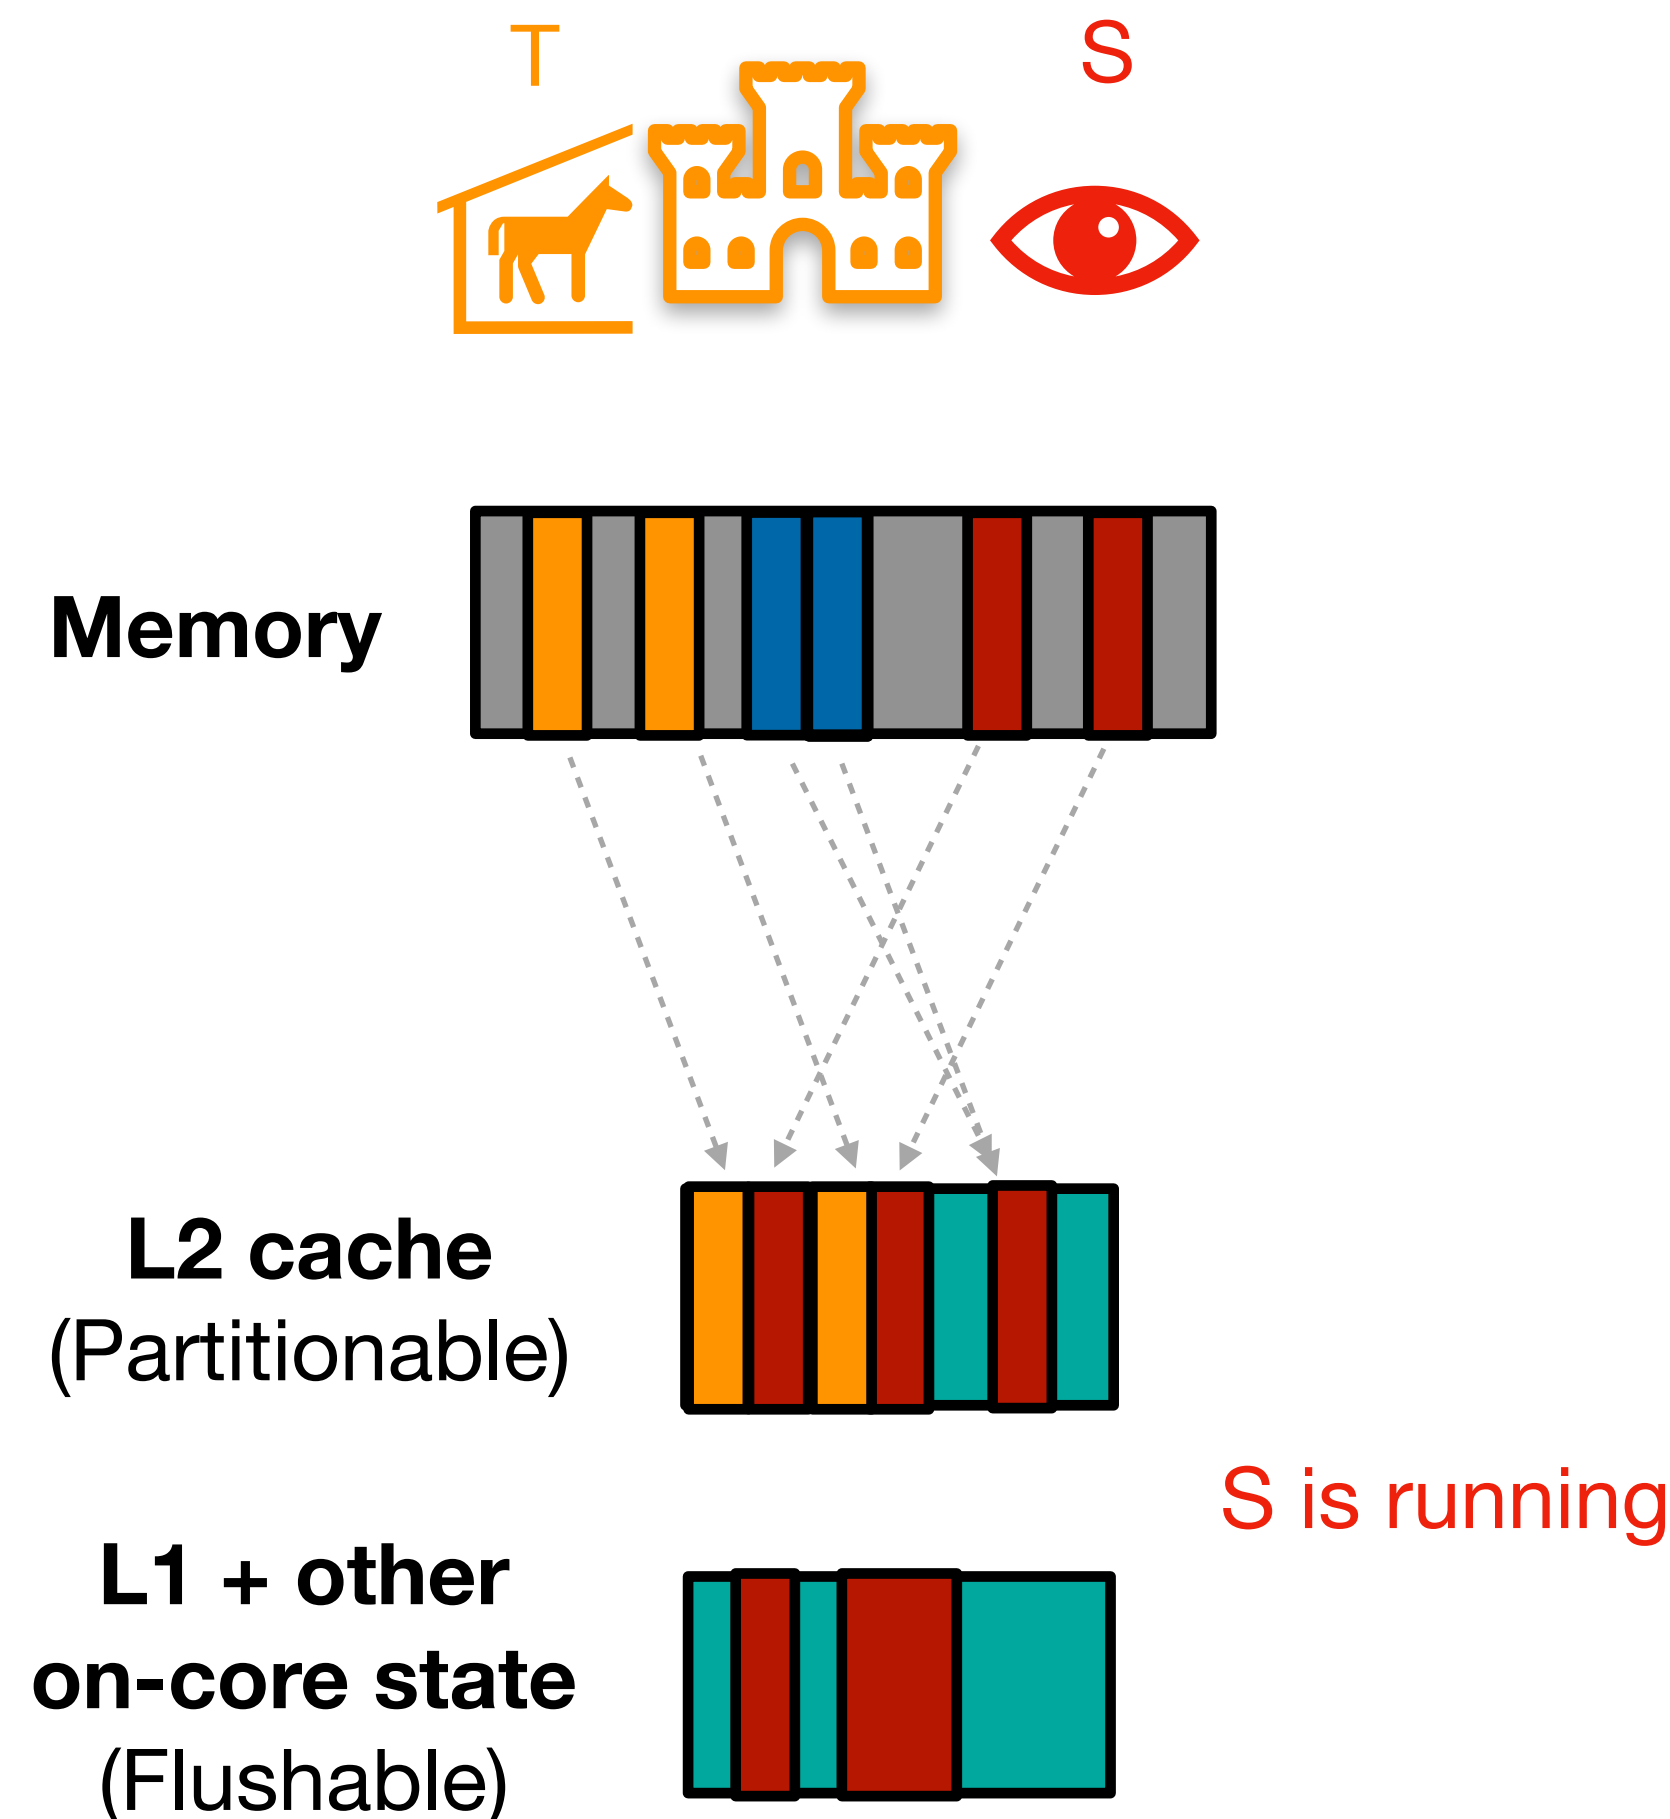
- Let's consider a *data diode* policy on infowflow via a page of **shared memory**:

- One direction allowed: $T < \sim S$
- Other direction *disallowed*: $T \sim /> S$

- Can we do it with a **partitionable** L2 cache?
S sees a timing difference on:

- Any writes by T (even to wipe the data!)
- Any reads by T that evict lines (many per page!)

We need *targeted flush* for partitionable L2 caches to enforce $T \sim /> S$ while allowing $T < \sim S$



What are time-protected communications?



- So far, we've discussed verifying time protection in seL4 with *strict isolation* policies,

i.e. $\forall A B. A \sim /> B$

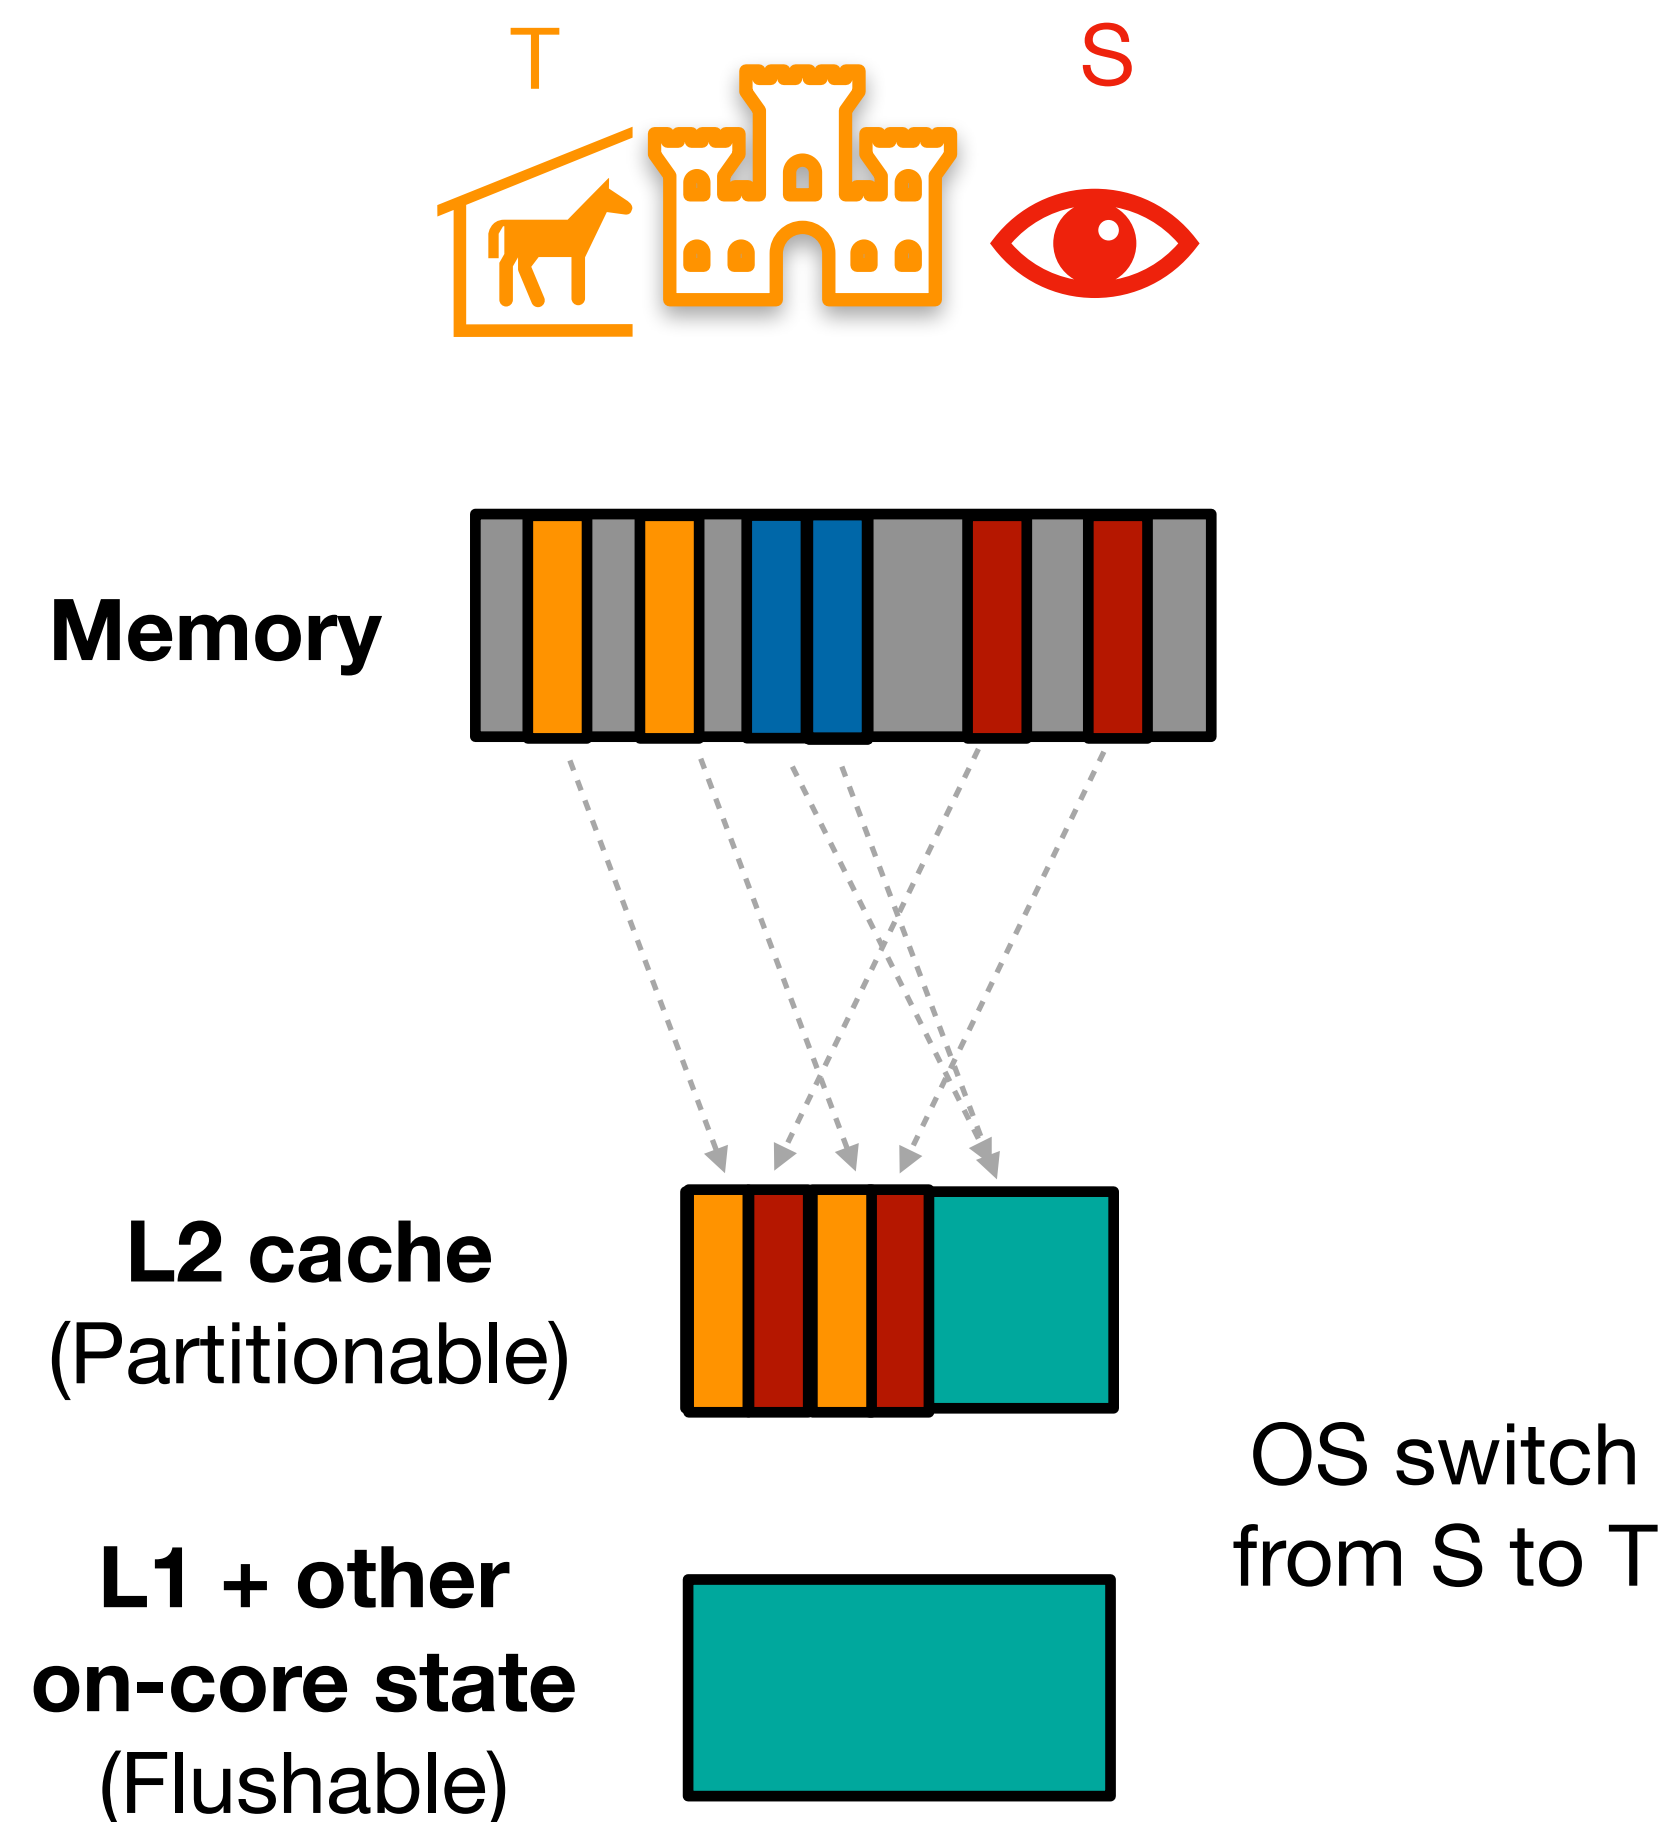
- Let's consider a *data diode* policy on infowflow via a page of **shared memory**:

- One direction allowed: $T < \sim S$
- Other direction *disallowed*: $T \sim /> S$

- Can we do it with a **partitionable** L2 cache?
S sees a timing difference on:

- Any writes by T (even to wipe the data!)
- Any reads by T that evict lines (many per page!)

We need *targeted flush* for partitionable L2 caches to enforce $T \sim /> S$ while allowing $T < \sim S$



What are time-protected communications?



- So far, we've discussed verifying time protection in seL4 with *strict isolation* policies,

i.e. $\forall A B. A \not\sim /> B$

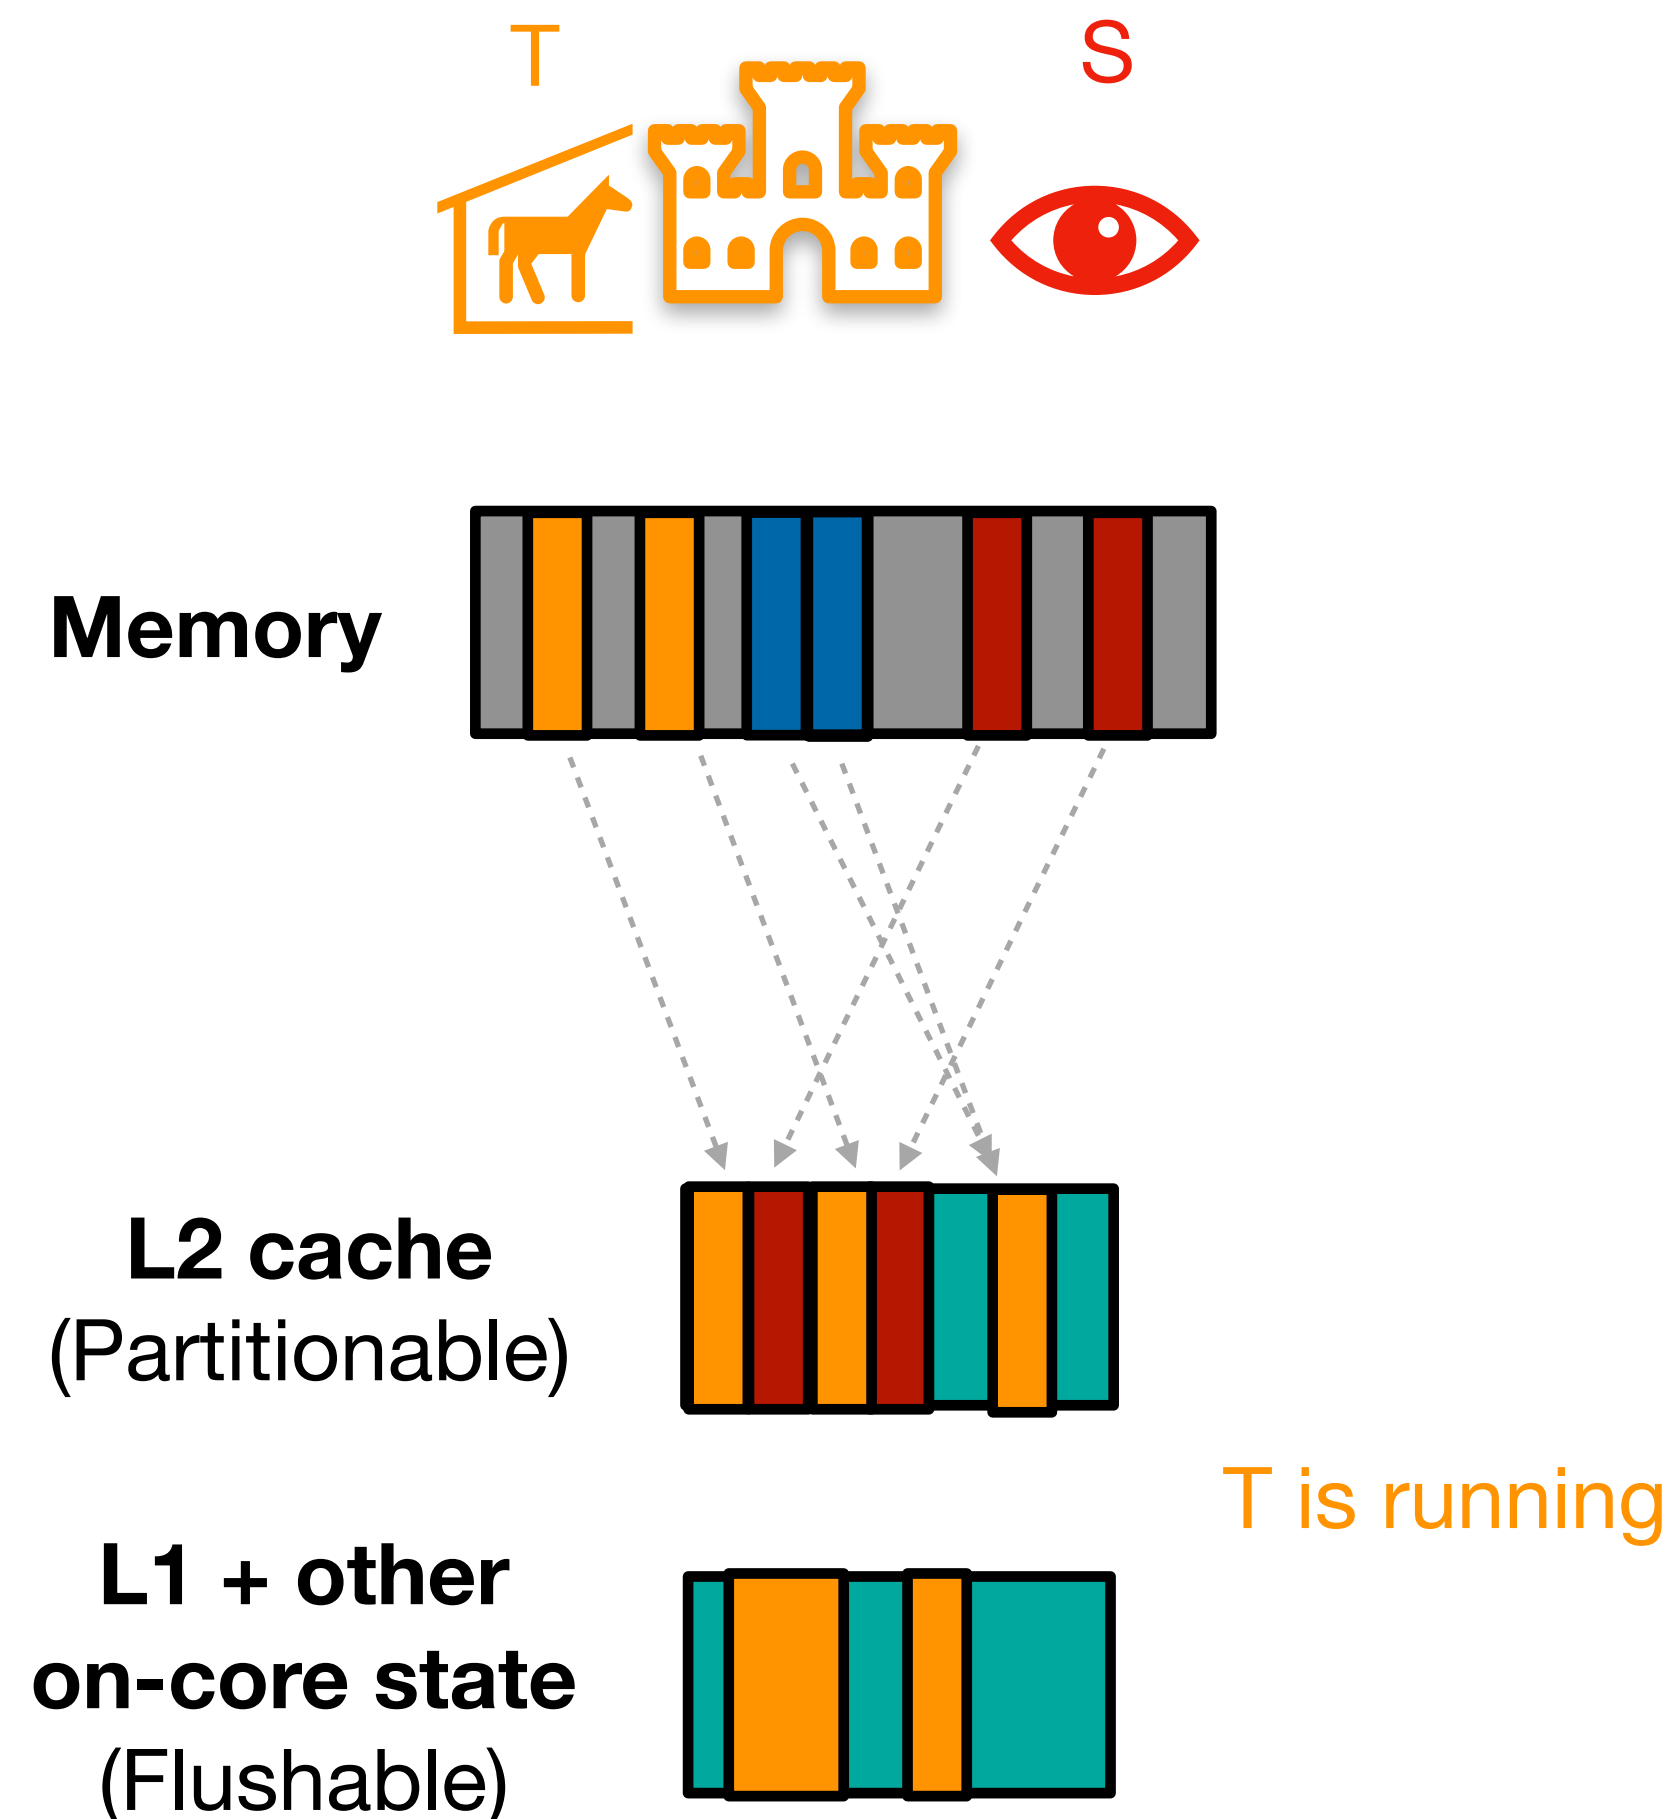
- Let's consider a *data diode* policy on infowflow via a page of **shared memory**:

- One direction allowed: $T <\sim S$
- Other direction *disallowed*: $T \not\sim /> S$

- Can we do it with a **partitionable** L2 cache?
S sees a timing difference on:

- Any writes by T (even to wipe the data!)
- Any reads by T that evict lines (many per page!)

We need *targeted flush* for partitionable L2 caches to enforce $T \not\sim /> S$ while allowing $T <\sim S$



What are time-protected communications?



- So far, we've discussed verifying **time protection** in seL4 with *strict isolation* policies,

i.e. $\forall A B. A \sim /> B$

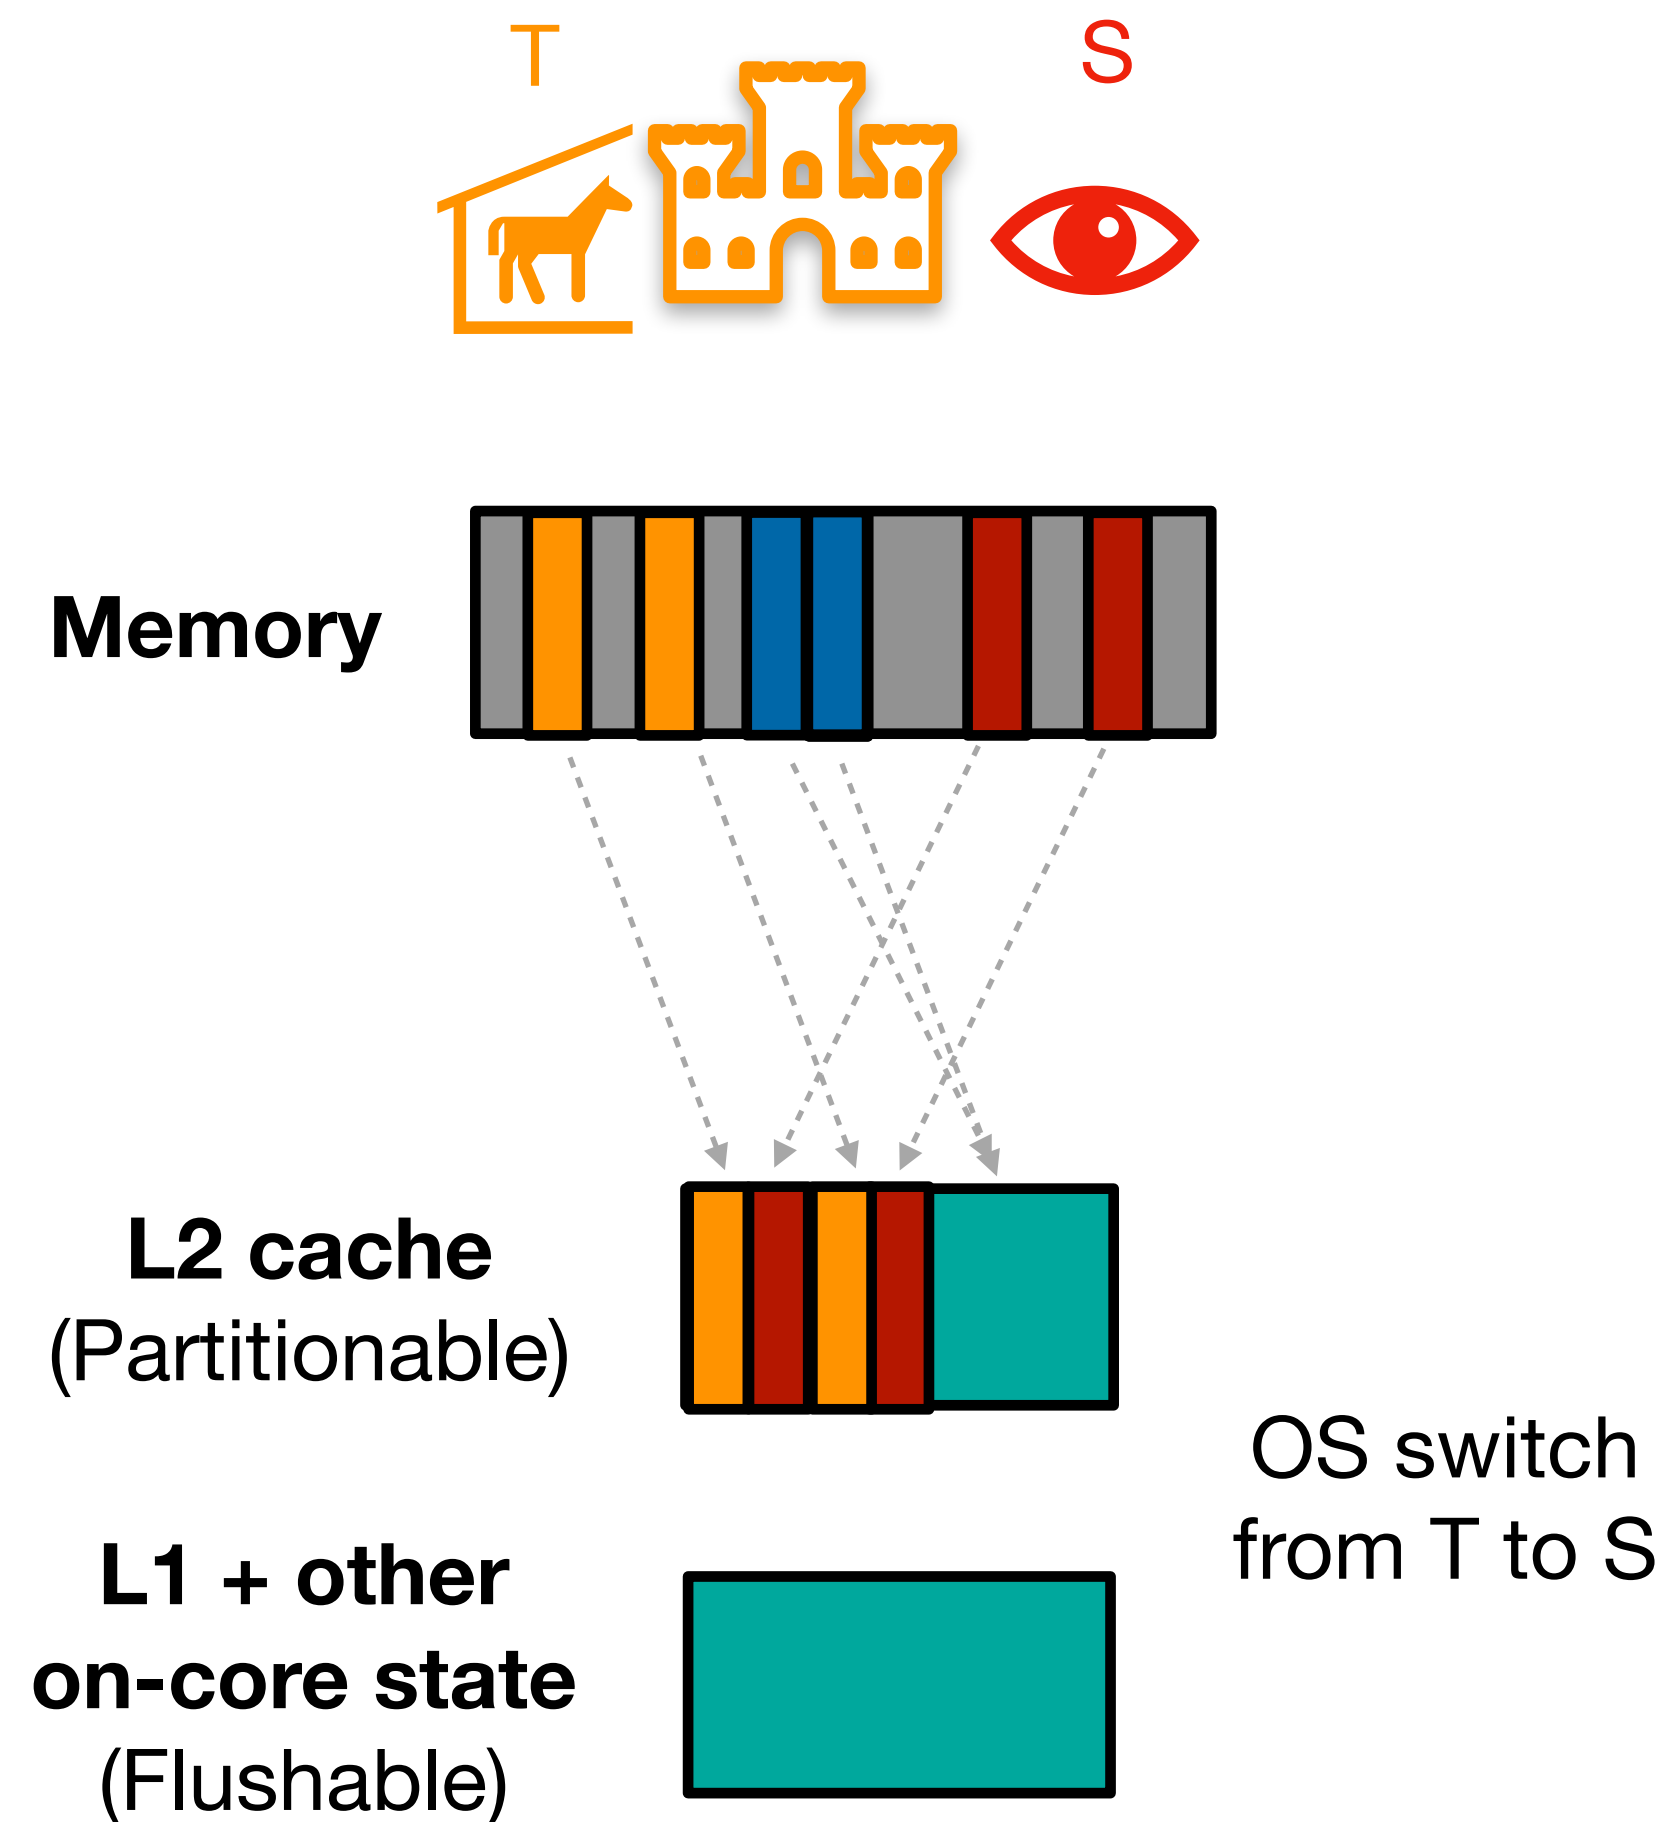
- Let's consider a *data diode* policy on infowflow via a page of **shared memory**:

- One direction allowed: $T < \sim S$
- Other direction *disallowed*: $T \sim /> S$

- Can we do it with a **partitionable** L2 cache?
S sees a timing difference on:

- Any writes by T (even to wipe the data!)
- Any reads by T that evict lines (many per page!)

We need **targeted flush** for partitionable L2 caches to enforce $T \sim /> S$ while allowing $T < \sim S$

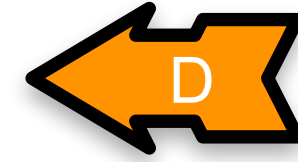




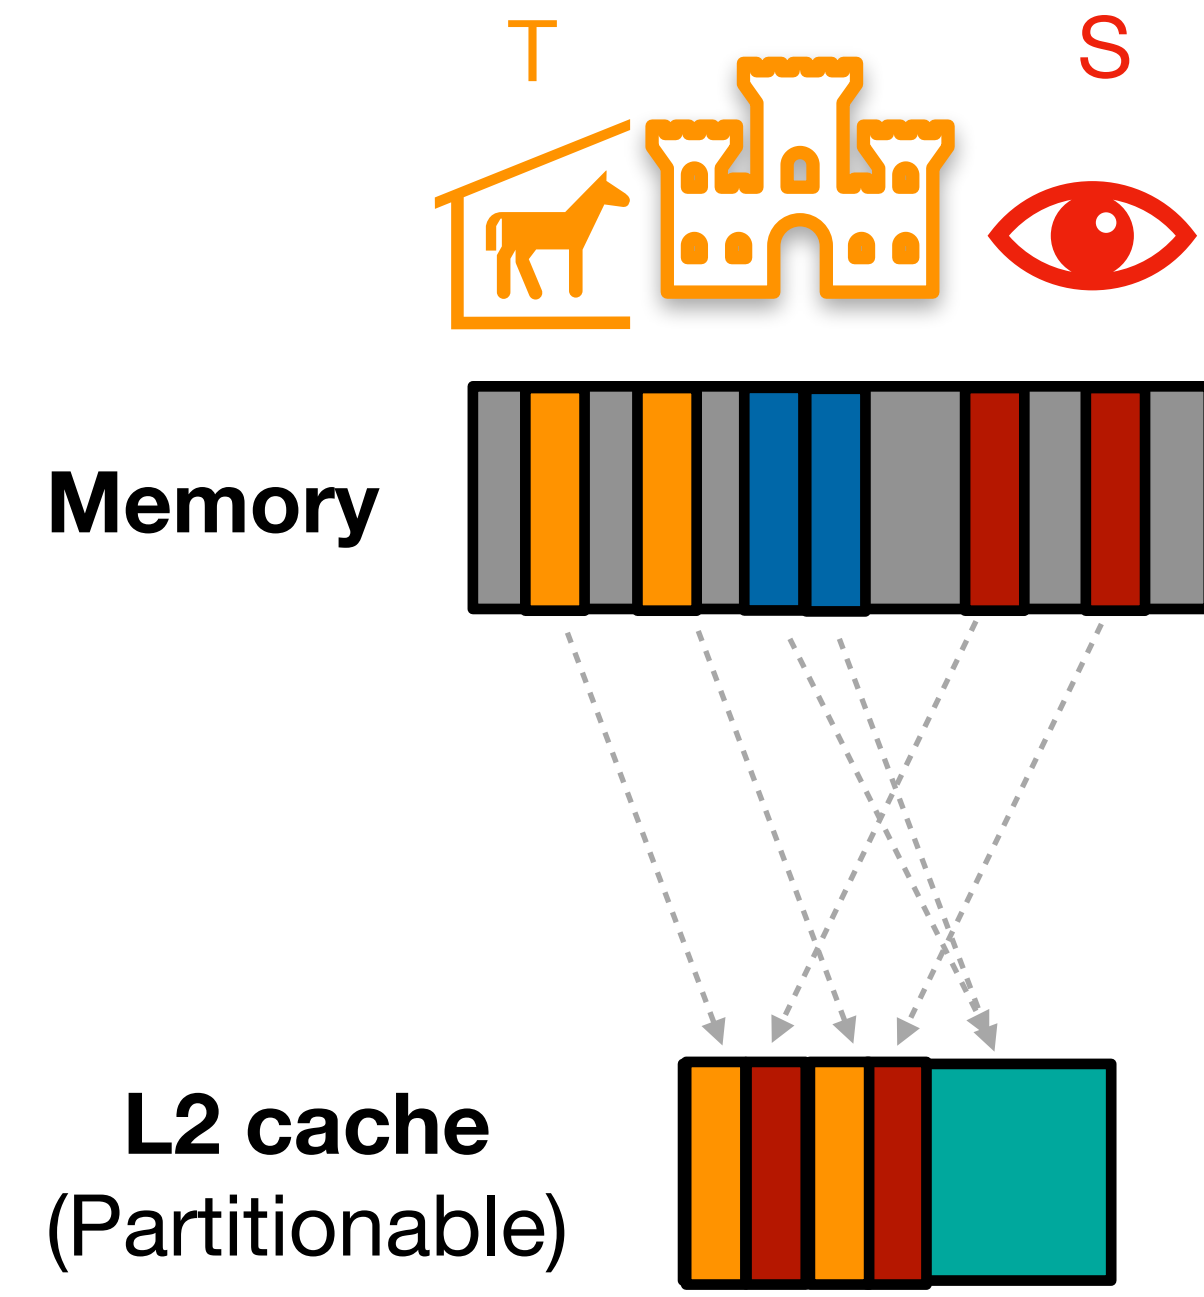
Time-protected communications in seL4



seL4 on RISC-V



Implement time-protected cross-domain communications



We need *targeted flush* for partitionable L2 caches to enforce $T \not\sim S$ while allowing $T \sim S$



Time-protected communications in seL4



seL4 on RISC-V



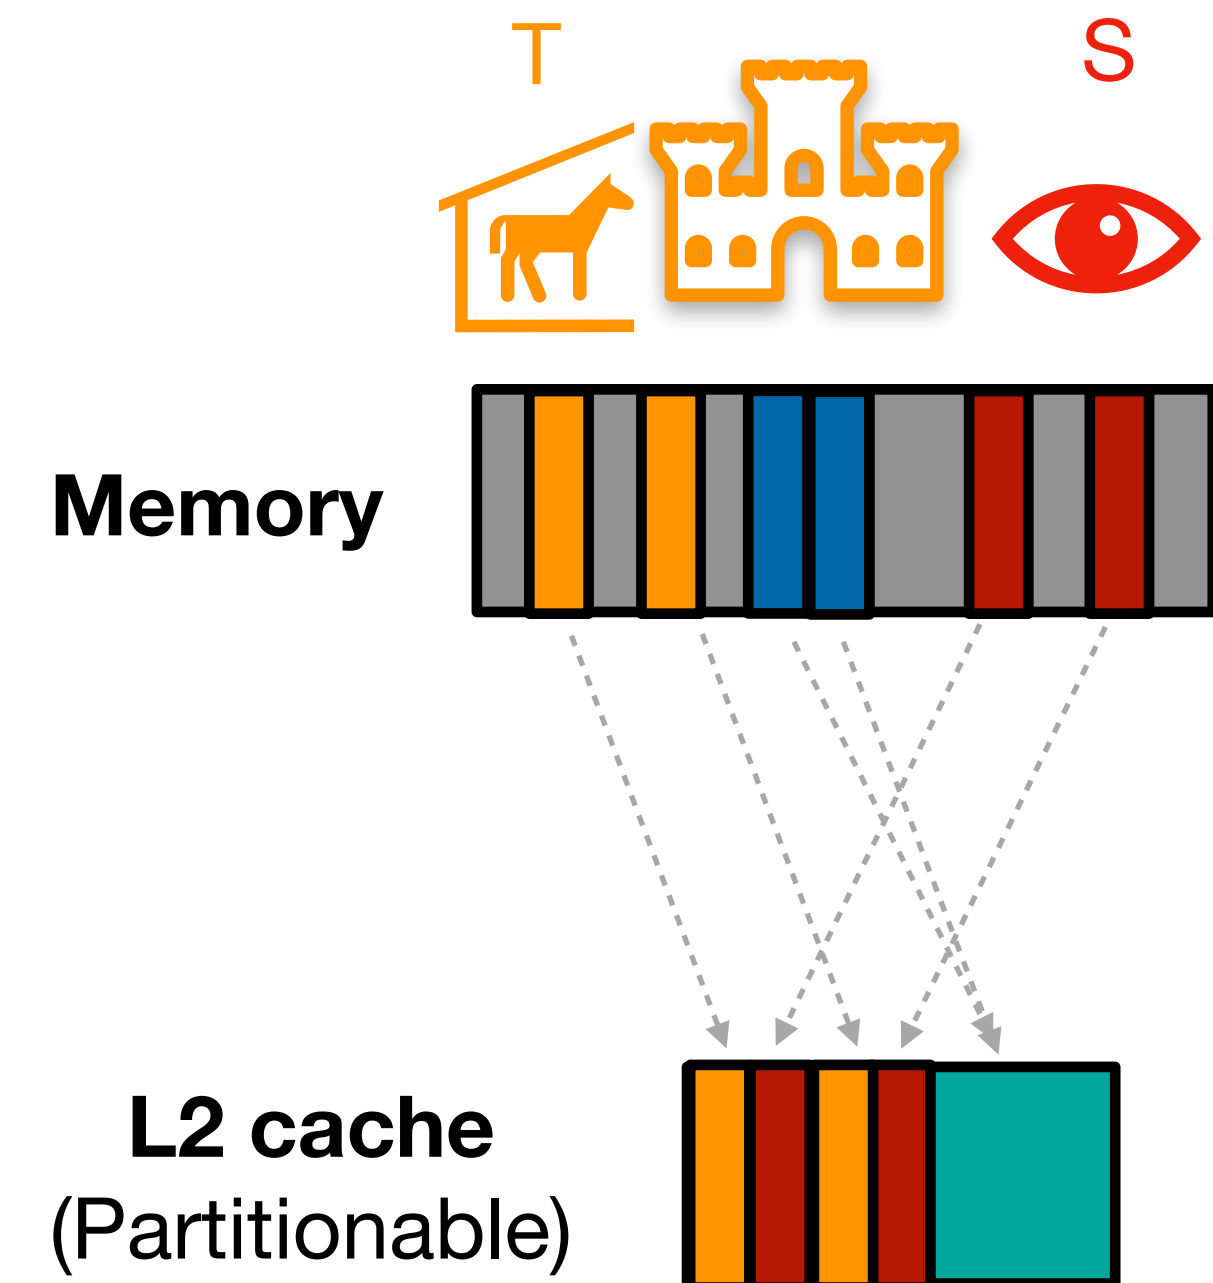
Implement time-protected cross-domain communications



- [Sethu 2025] comprehensively evaluated “legacy” implementation options on seL4, with the L2 partitioned via *cache colouring*:



Varun Sethu
BE Thesis



We need *targeted flush* for partitionable L2 caches to enforce $T \not\sim S$ while allowing $T \sim S$



Time-protected communications in seL4



seL4 on RISC-V



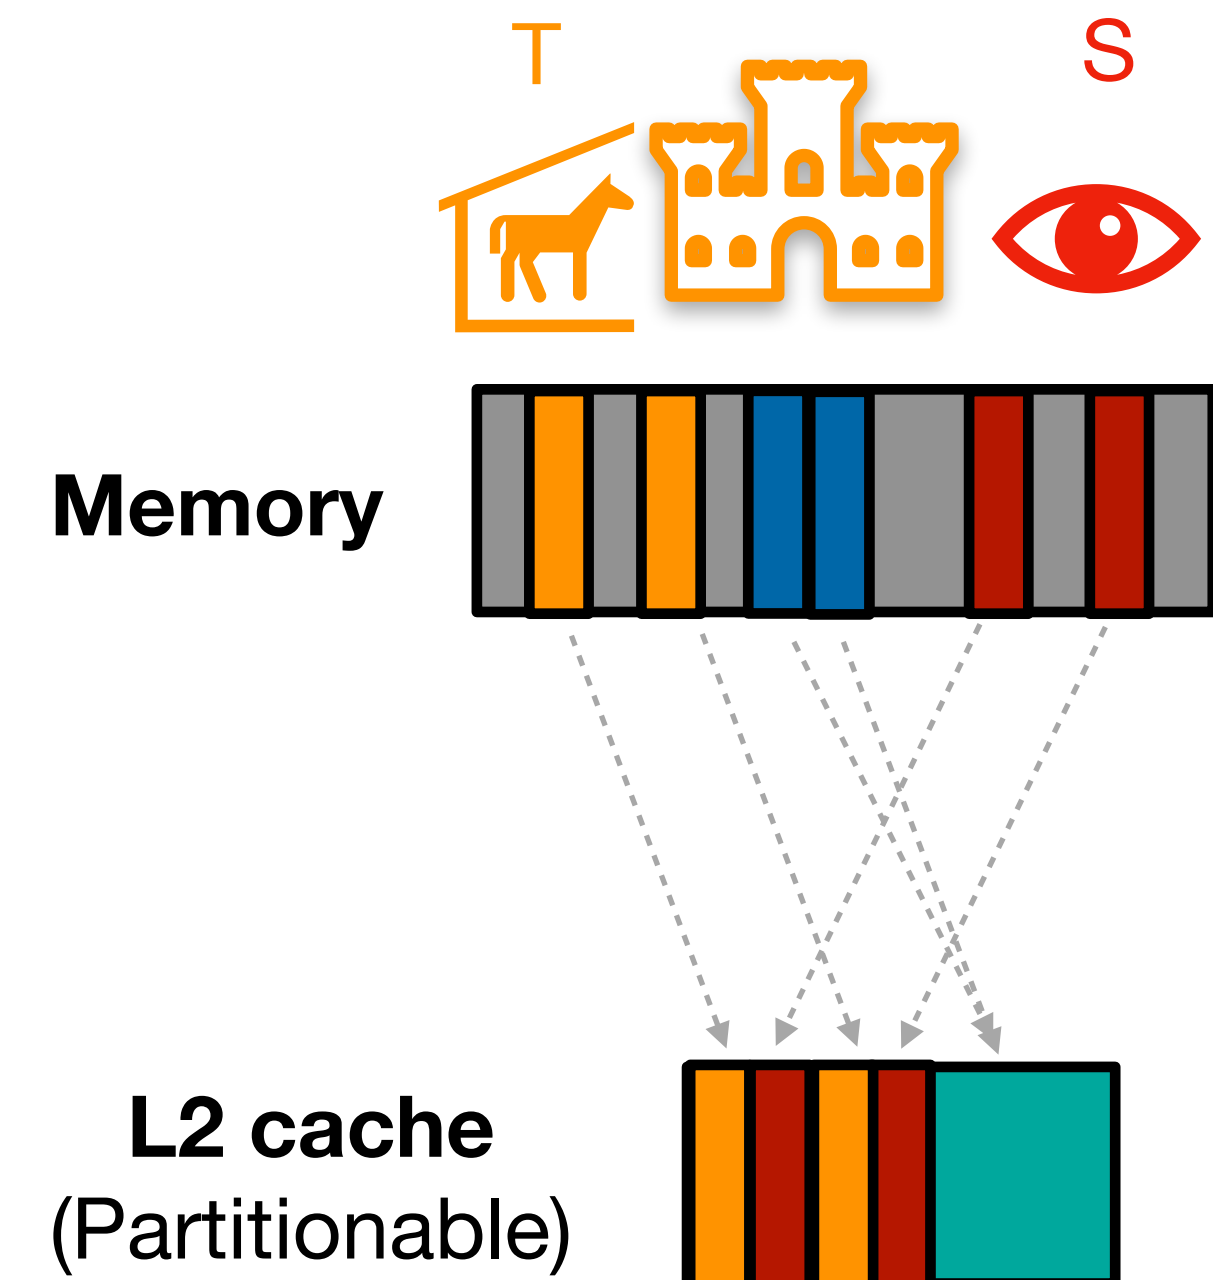
Implement time-protected cross-domain communications



- [Sethu 2025] comprehensively evaluated “legacy” implementation options on seL4, with the L2 partitioned via *cache colouring*:
 - Copy between non-shared memory (slow)



Varun Sethu
BE Thesis



We need *targeted flush* for partitionable L2 caches to enforce $T \not\sim S$ while allowing $T \sim S$



Time-protected communications in seL4



seL4 on RISC-V



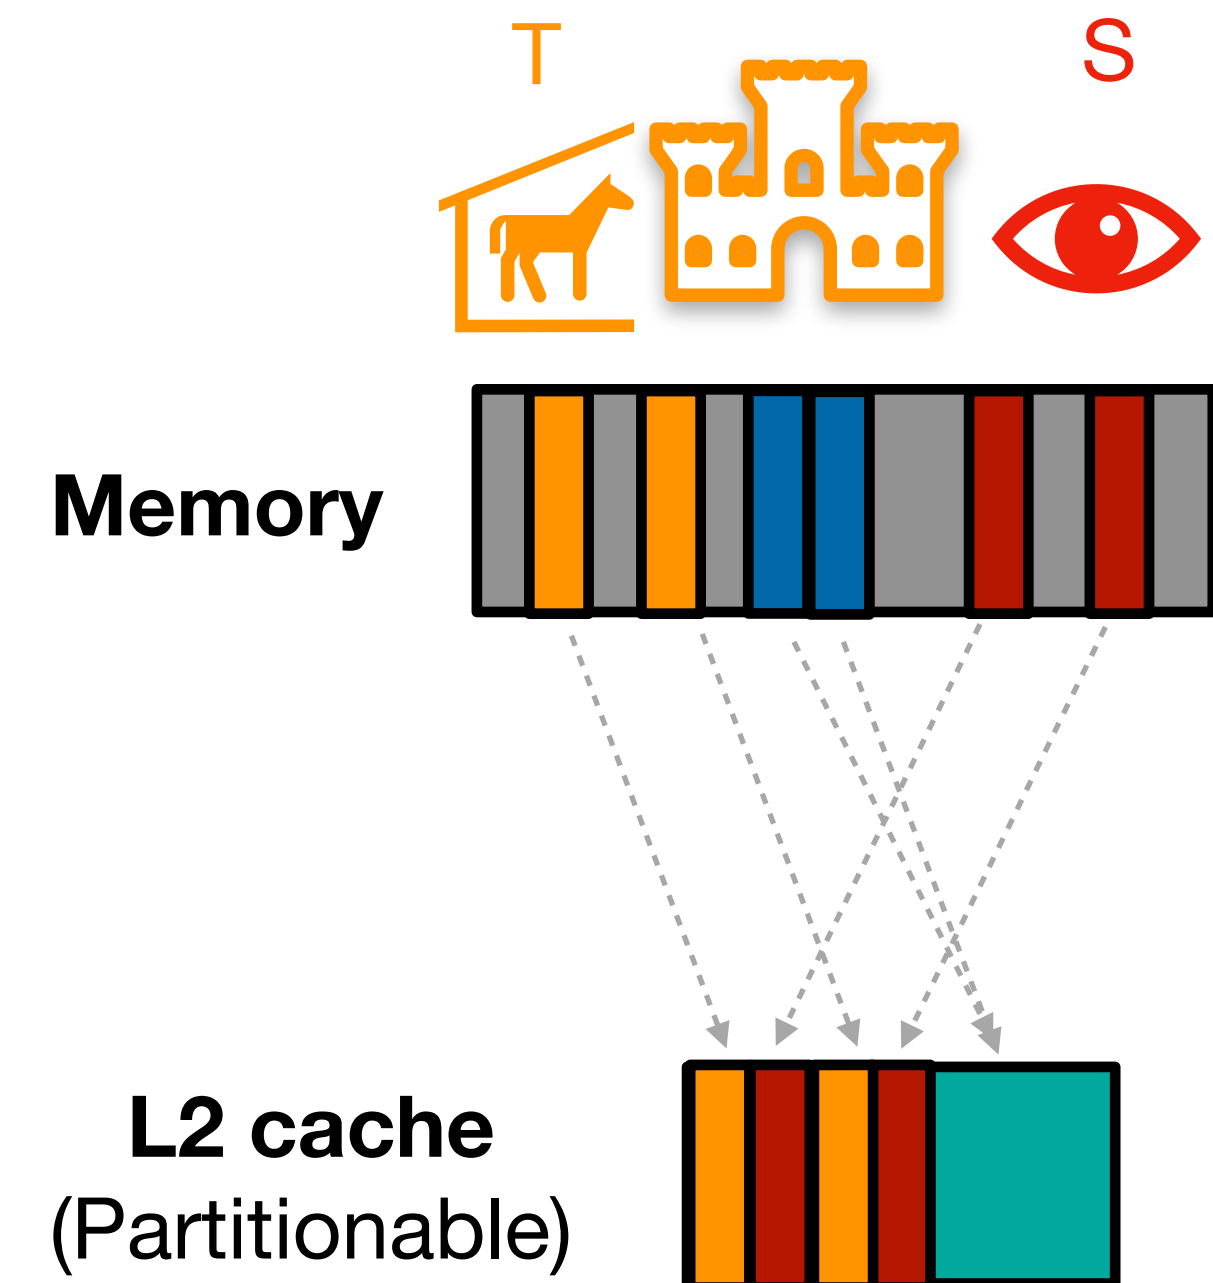
Implement time-protected cross-domain communications



- [Sethu 2025] comprehensively evaluated “legacy” implementation options on seL4, with the L2 partitioned via *cache colouring*:
 - Copy between non-shared memory (slow)
 - Dedicate colours for shared memory (wastes space)



Varun Sethu
BE Thesis



We need *targeted flush* for partitionable L2 caches to enforce $T \not\sim S$ while allowing $T \sim S$



Time-protected communications in seL4



seL4 on RISC-V



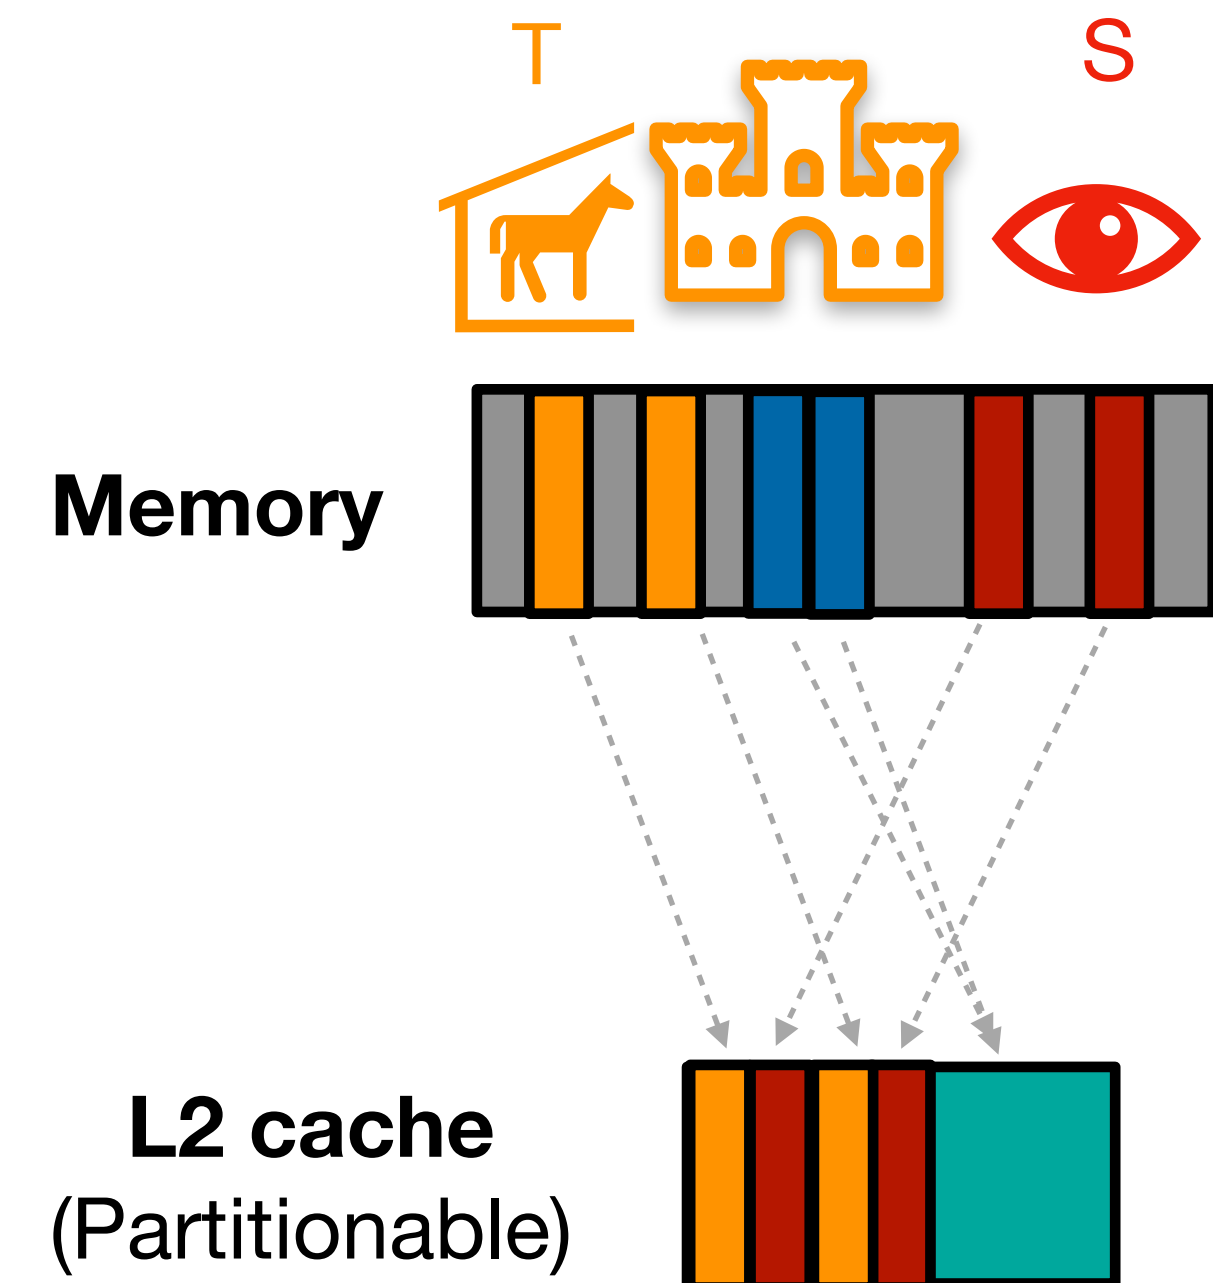
Implement time-protected cross-domain communications



- [Sethu 2025] comprehensively evaluated “legacy” implementation options on seL4, with the L2 partitioned via *cache colouring*:
 - Copy between non-shared memory (slow)
 - Dedicate colours for shared memory (wastes space)
 - Targeted “flush” of L2 via prefetching (still leaks)



Varun Sethu
BE Thesis



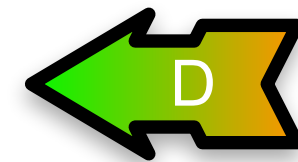
We need *targeted flush* for partitionable L2 caches to enforce $T \not\sim S$ while allowing $T \sim S$



Time-protected communications in seL4



seL4 on RISC-V



Implement time-protected cross-domain communications

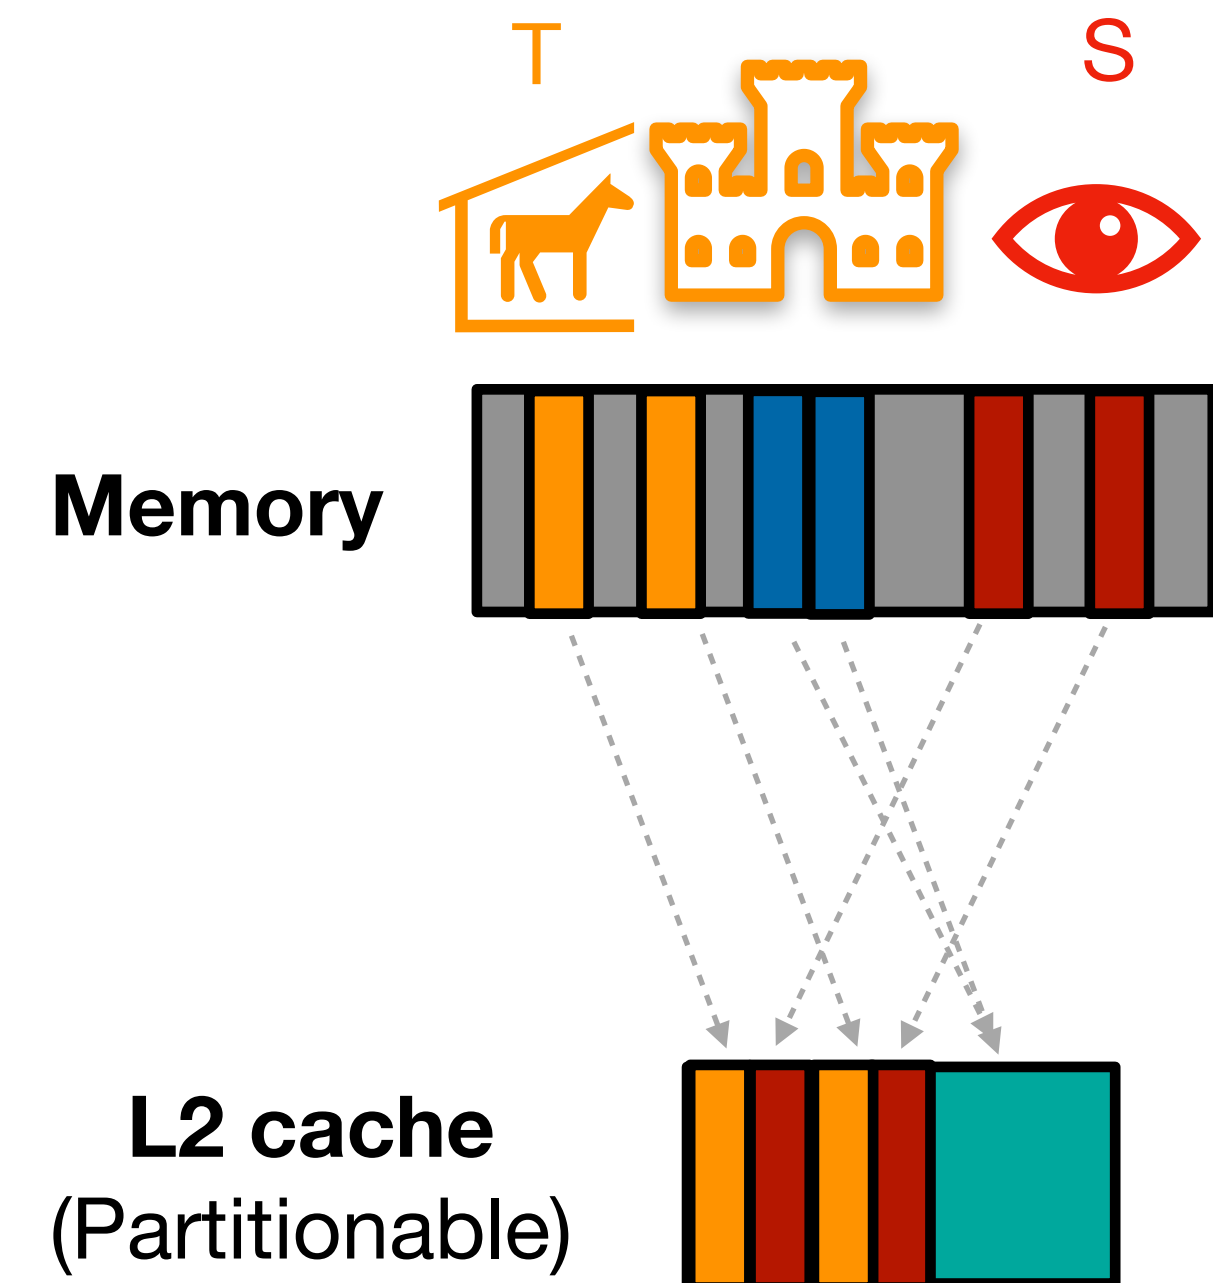


- [Sethu 2025] comprehensively evaluated “legacy” implementation options on seL4, with the L2 partitioned via *cache colouring*:
 - Copy between non-shared memory (slow)
 - Dedicate colours for shared memory (wastes space)
 - Targeted “flush” of L2 via prefetching (still leaks)



Varun Sethu
BE Thesis

Continuing on from



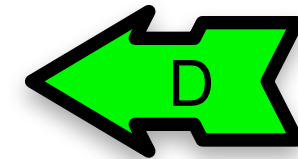
We need *targeted flush* for partitionable L2 caches to enforce $T \not\sim S$ while allowing $T \sim S$



Time-protected communications in seL4



seL4 on RISC-V



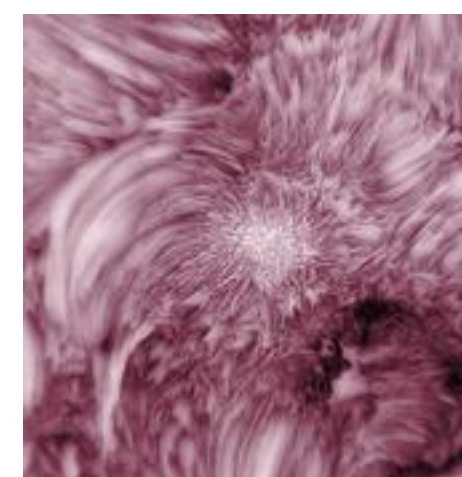
Implement time-protected cross-domain communications



- [Sethu 2025] comprehensively evaluated “legacy” implementation options on seL4, with the L2 partitioned via *cache colouring*:
 - Copy between non-shared memory (slow)
 - Dedicate colours for shared memory (wastes space)
 - Targeted “flush” of L2 via prefetching (still leaks)
- Now we're using *dynamically partitionable last-level cache* (DPLLC) i.e. dedicated HW support on RISC-V Cheshire



Varun Sethu
BE Thesis



Julia Vassiliki

Continuing on from



Memory



L2 cache (Partitionable)



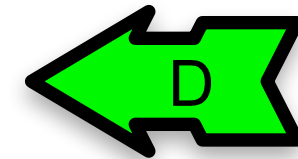
We need *targeted flush* for partitionable L2 caches to enforce $T \not\sim S$ while allowing $T \sim S$



Time-protected communications in seL4



seL4 on RISC-V



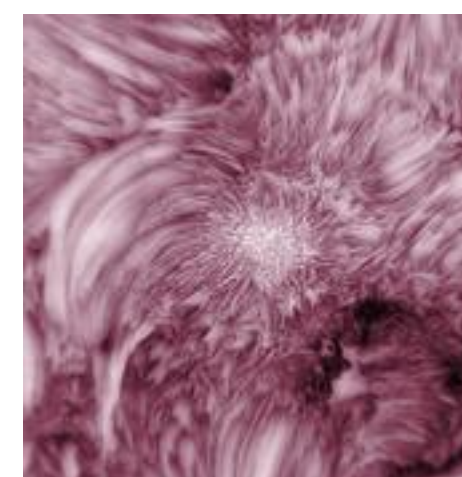
Implement time-protected cross-domain communications



- [Sethu 2025] comprehensively evaluated “legacy” implementation options on seL4, with the L2 partitioned via *cache colouring*:
 - Copy between non-shared memory (slow)
 - Dedicate colours for shared memory (wastes space)
 - Targeted “flush” of L2 via prefetching (still leaks)
- Now we're using *dynamically partitionable last-level cache* (DPLLC) i.e. dedicated HW support on RISC-V Cheshire
 - Supports *targeted flush* of partitions



Varun Sethu
BE Thesis



Julia Vassiliki

Continuing on from



Memory



L2 cache
(Partitionable)



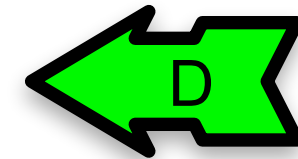
We need *targeted flush* for partitionable L2 caches to enforce $T \not\sim S$ while allowing $T \sim S$



Time-protected communications in seL4



seL4 on RISC-V



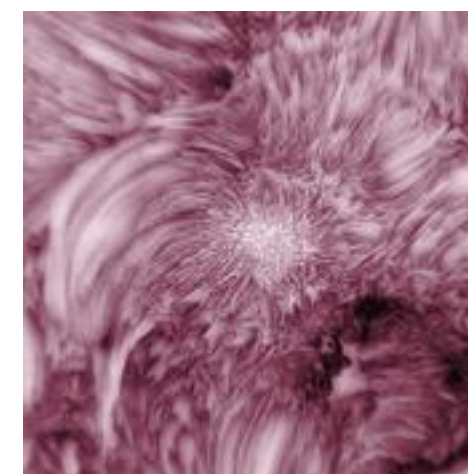
Implement time-protected cross-domain communications



- [Sethu 2025] comprehensively evaluated “legacy” implementation options on seL4, with the L2 partitioned via *cache colouring*:
 - Copy between non-shared memory (slow)
 - Dedicate colours for shared memory (wastes space)
 - Targeted “flush” of L2 via prefetching (still leaks)
- Now we're using *dynamically partitionable last-level cache* (DPLLC) i.e. dedicated HW support on RISC-V Cheshire
 - Supports *targeted flush* of partitions
 - Adding needed *time padding* on flush

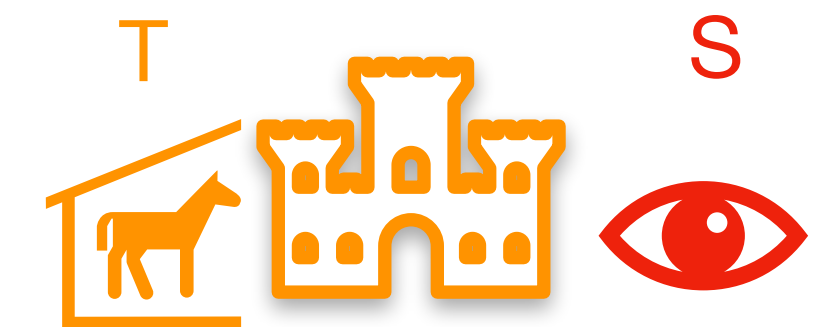


Varun Sethu
BE Thesis



Julia Vassiliki

Continuing on from



Memory



L2 cache
(Partitionable)



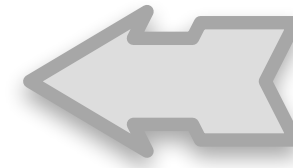
We need *targeted flush* for partitionable L2 caches to enforce $T \not\sim S$ while allowing $T \sim S$



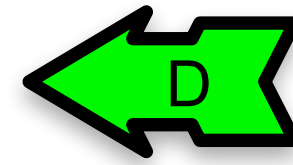
Verifying seL4 time-protected communications



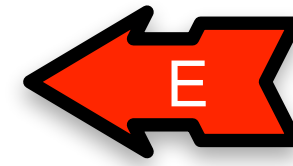
seL4 on RISC-V



FM'23: Define time protection for OS kernels

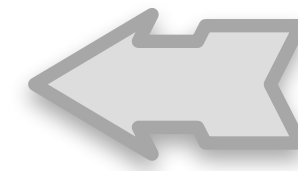


Implement time-protected cross-domain communications

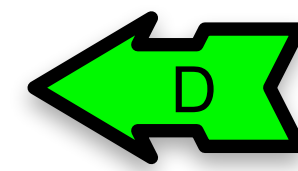


Verify time-protected cross-domain communications

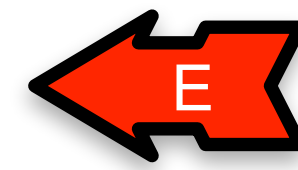




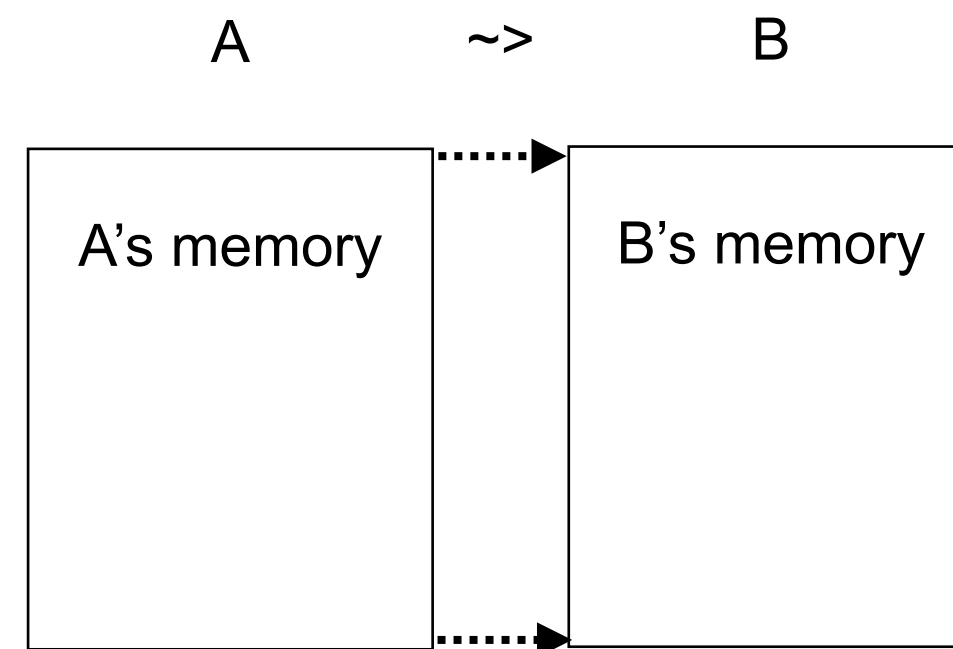
FM'23: Define time protection for OS kernels



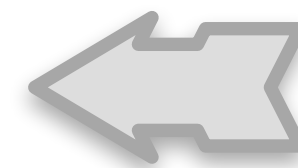
Implement time-protected cross-domain communications



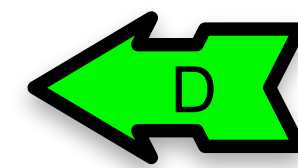
Verify time-protected cross-domain communications



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



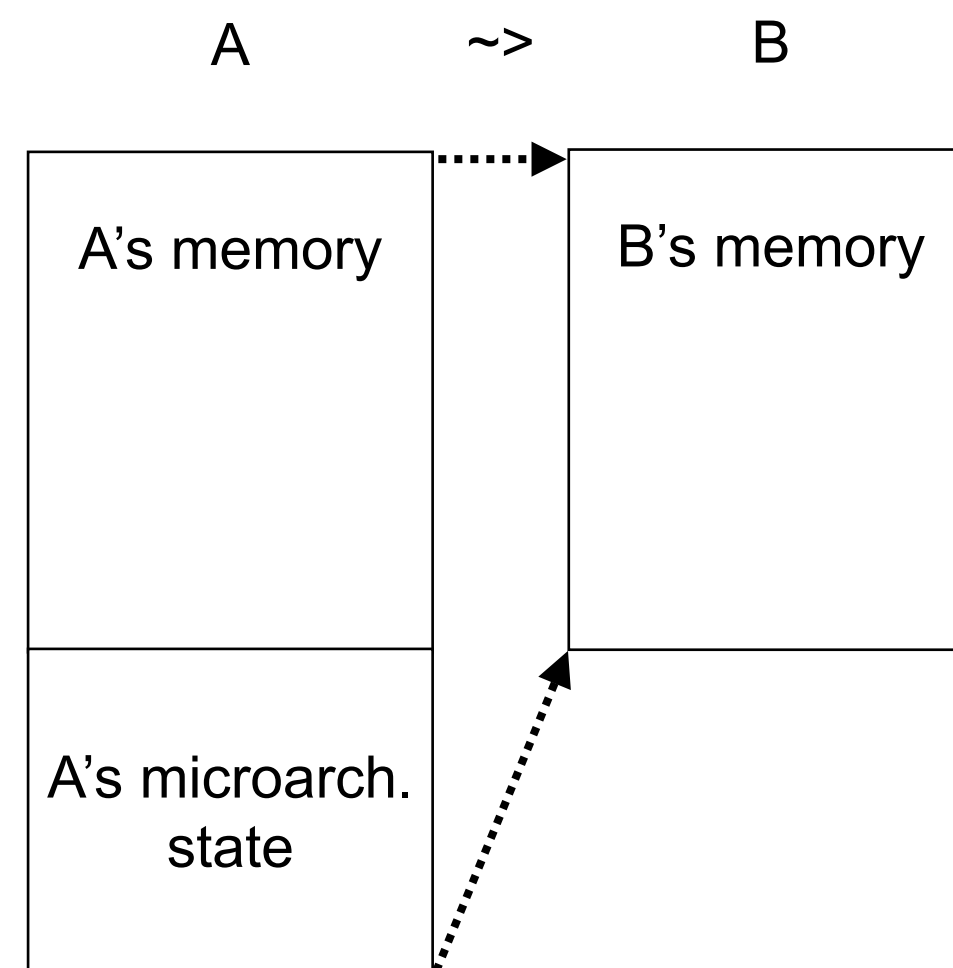
FM'23: Define time protection for OS kernels



Implement time-protected cross-domain communications



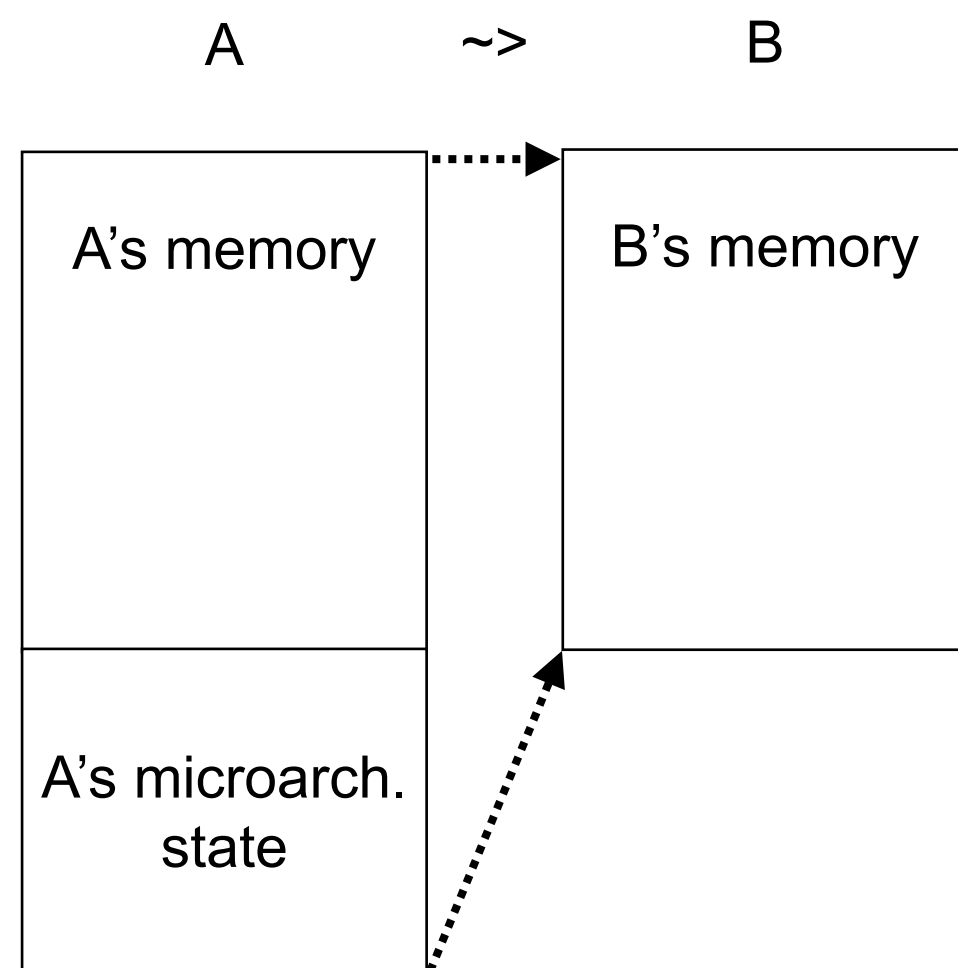
Verify time-protected cross-domain communications



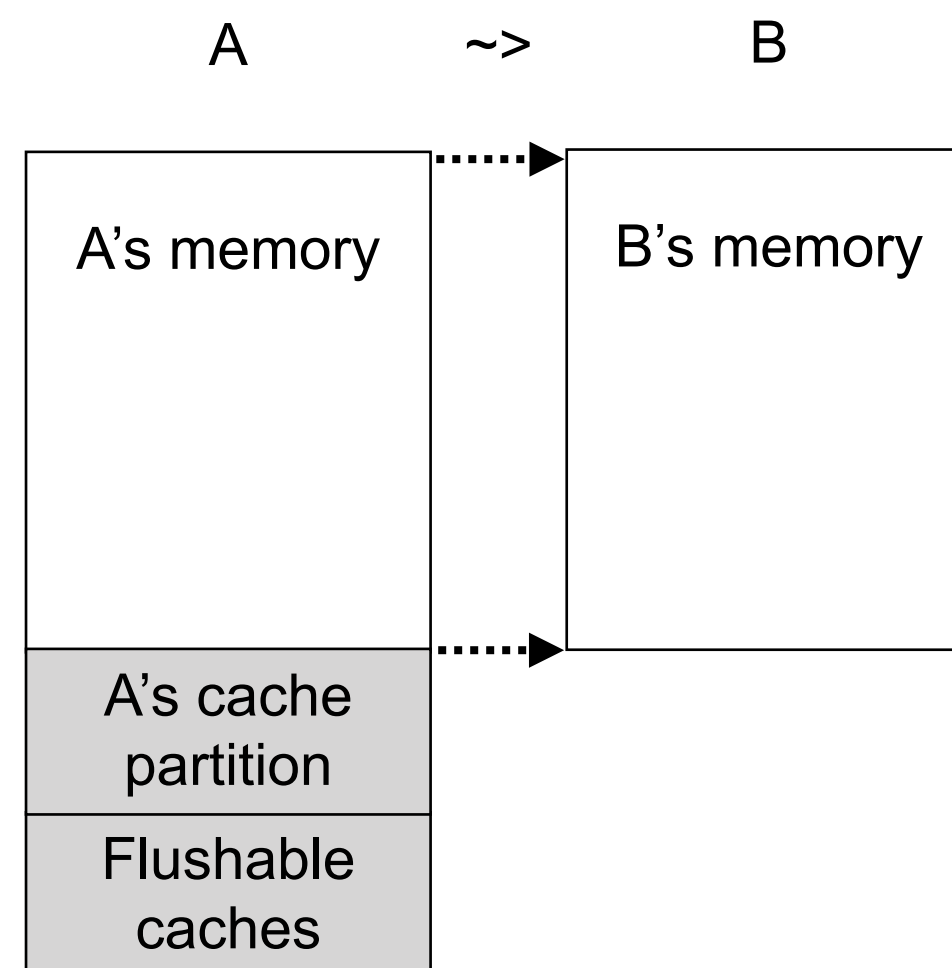
From prior seL4 infoflow proofs
[Murray et al. 2012, 2013]:
"all or nothing" policies



- FM'23: Define time protection for OS kernels
- Implement time-protected cross-domain communications
- Verify time-protected cross-domain communications

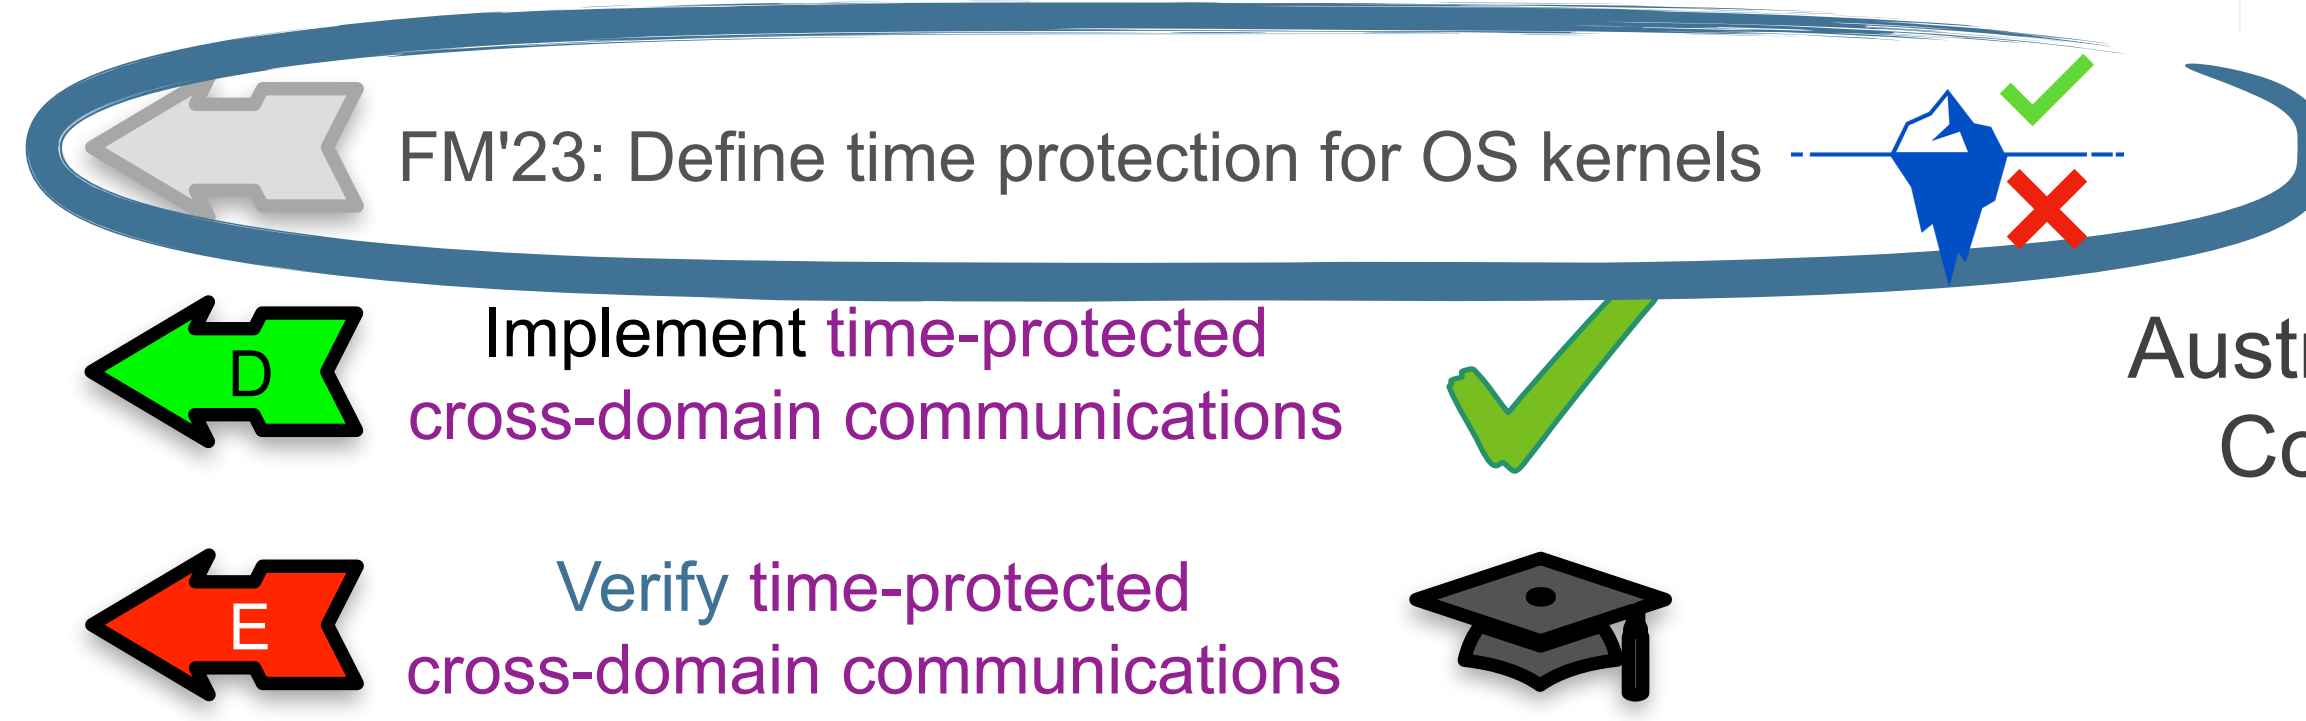


From prior seL4 infowflow proofs
 [Murray et al. 2012, 2013]:
 “*all or nothing*” policies

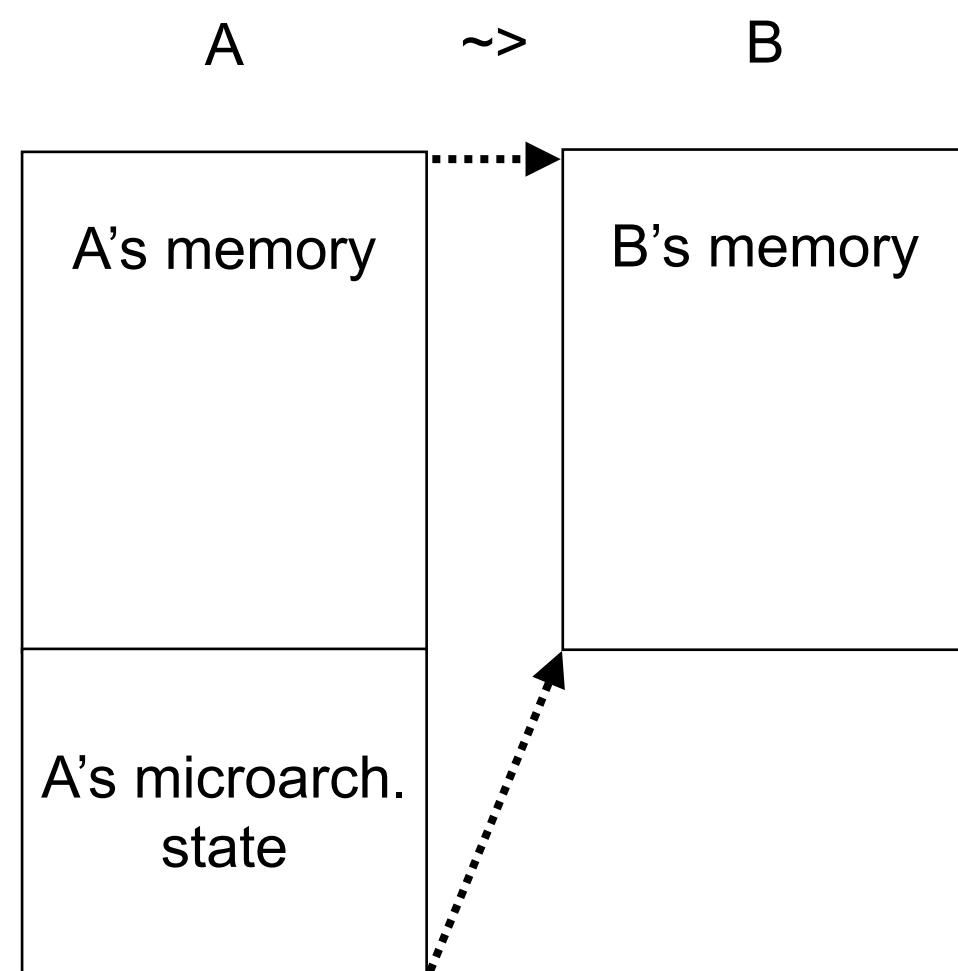


**Need *spatial precision*
 to *allow some flows*
 but *exclude others***

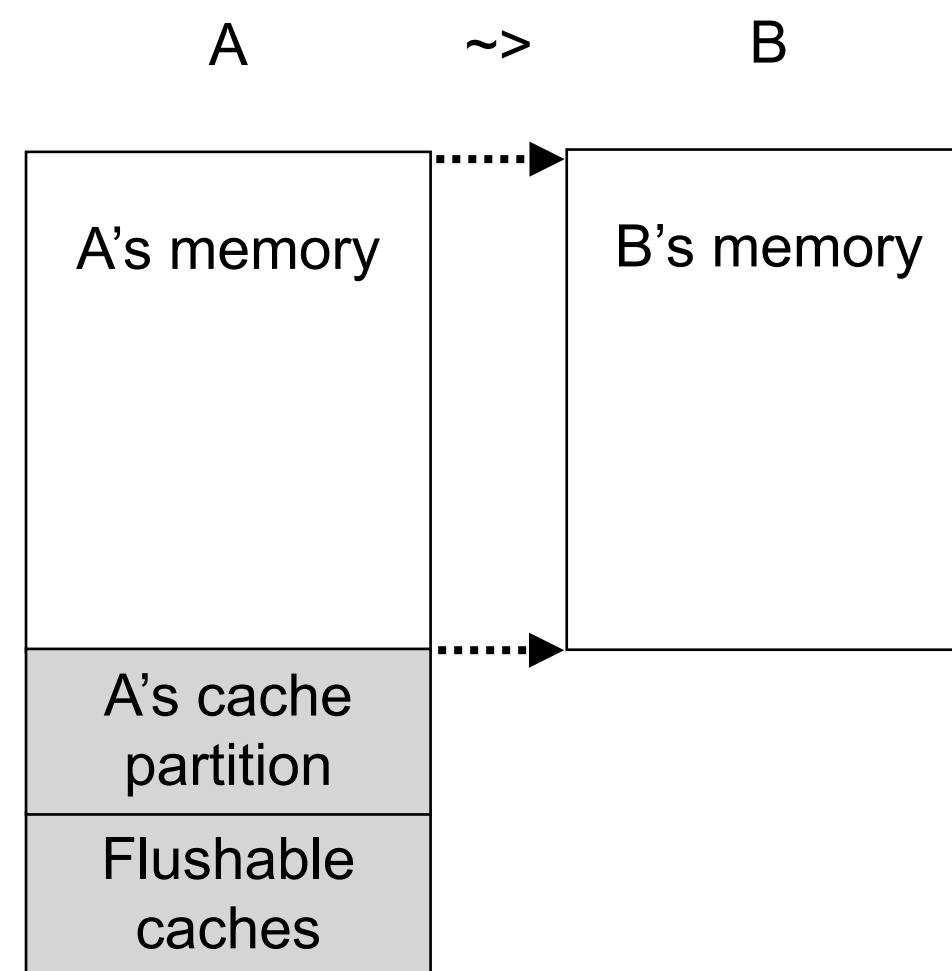
seL4 on RISC-V



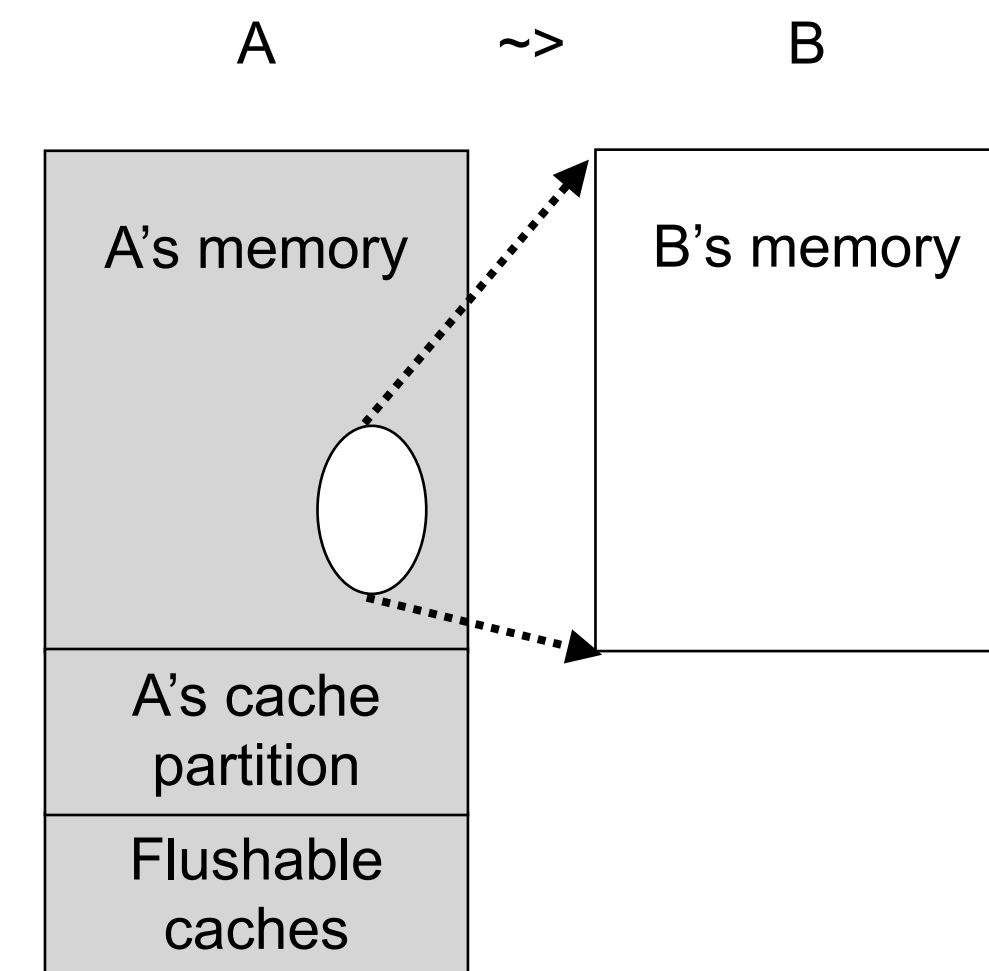
Thanks to Australian Research Council funding



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

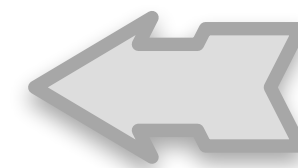


Need spatial precision to allow some flows but exclude others

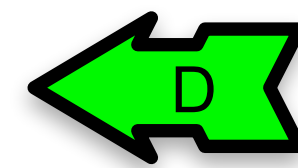


Our FM'23 paper [Sison et al. 2023] allows infowall policies with arbitrary spatial and temporal precision

seL4 on RISC-V



FM'23: Define time protection for OS kernels



Implement time-protected cross-domain communications



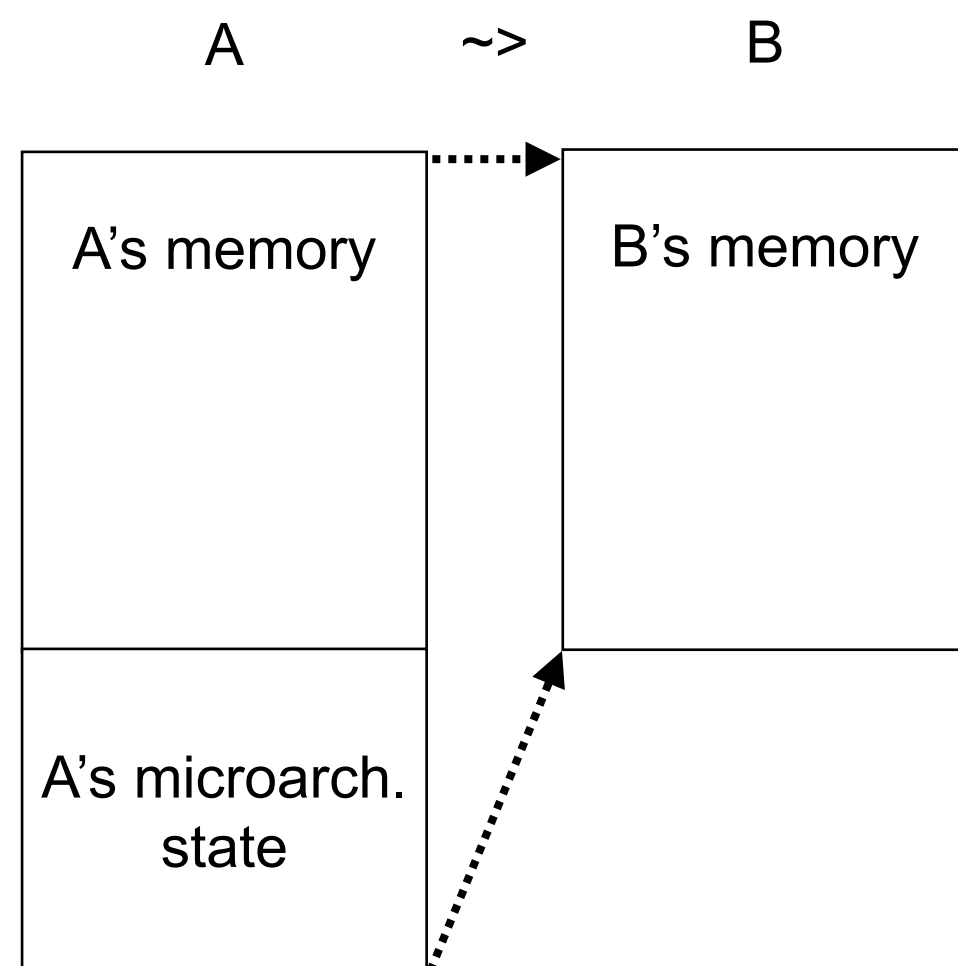
Thanks to funding from



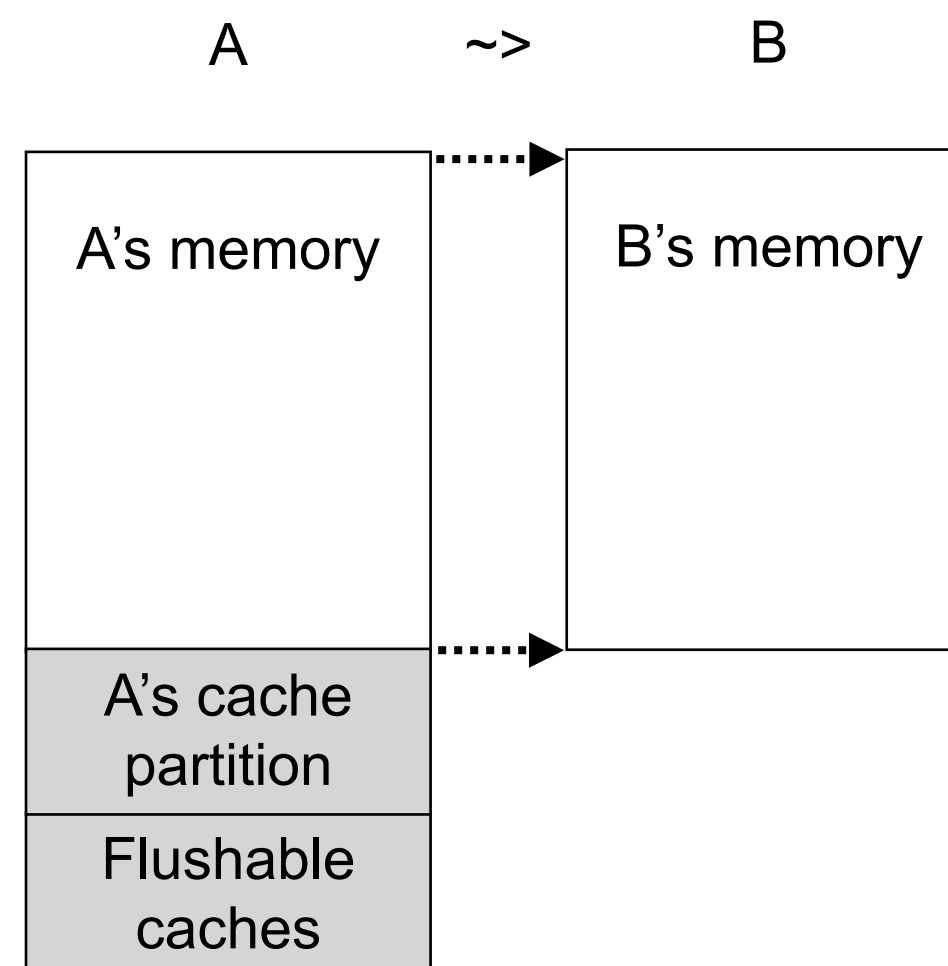
Verify time-protected cross-domain communications



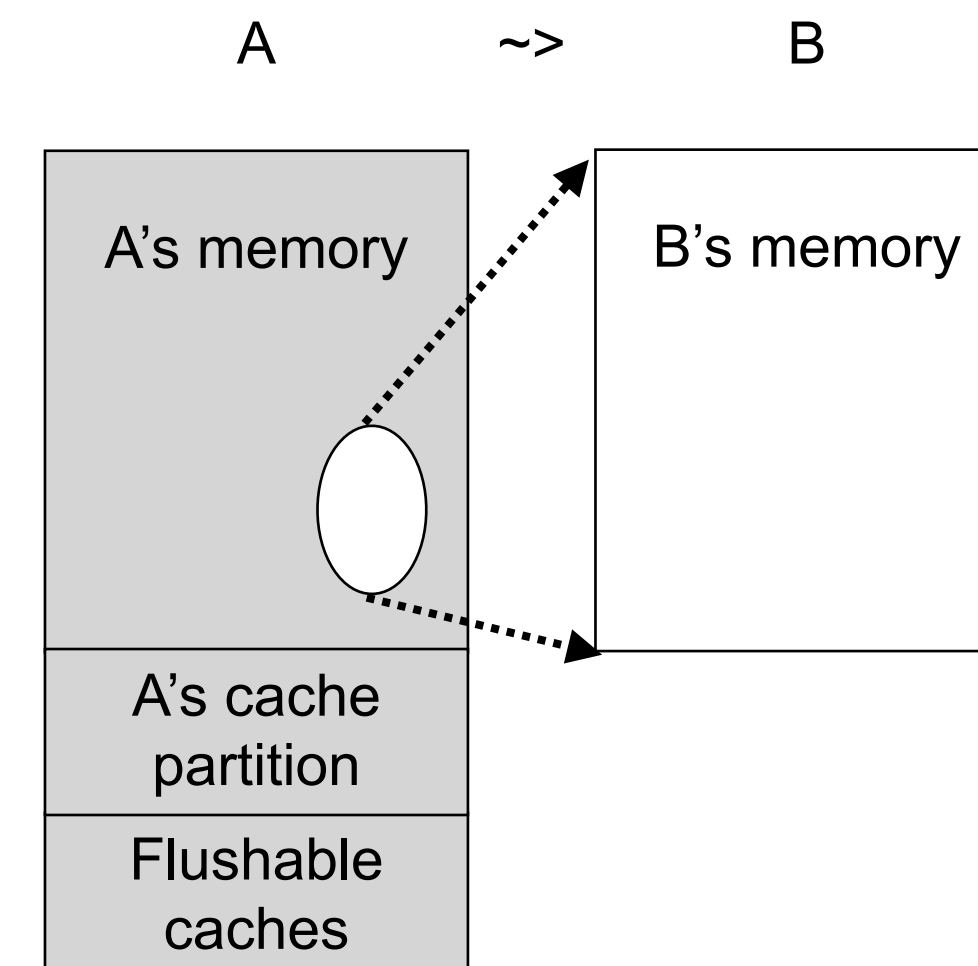
cyberagentur



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



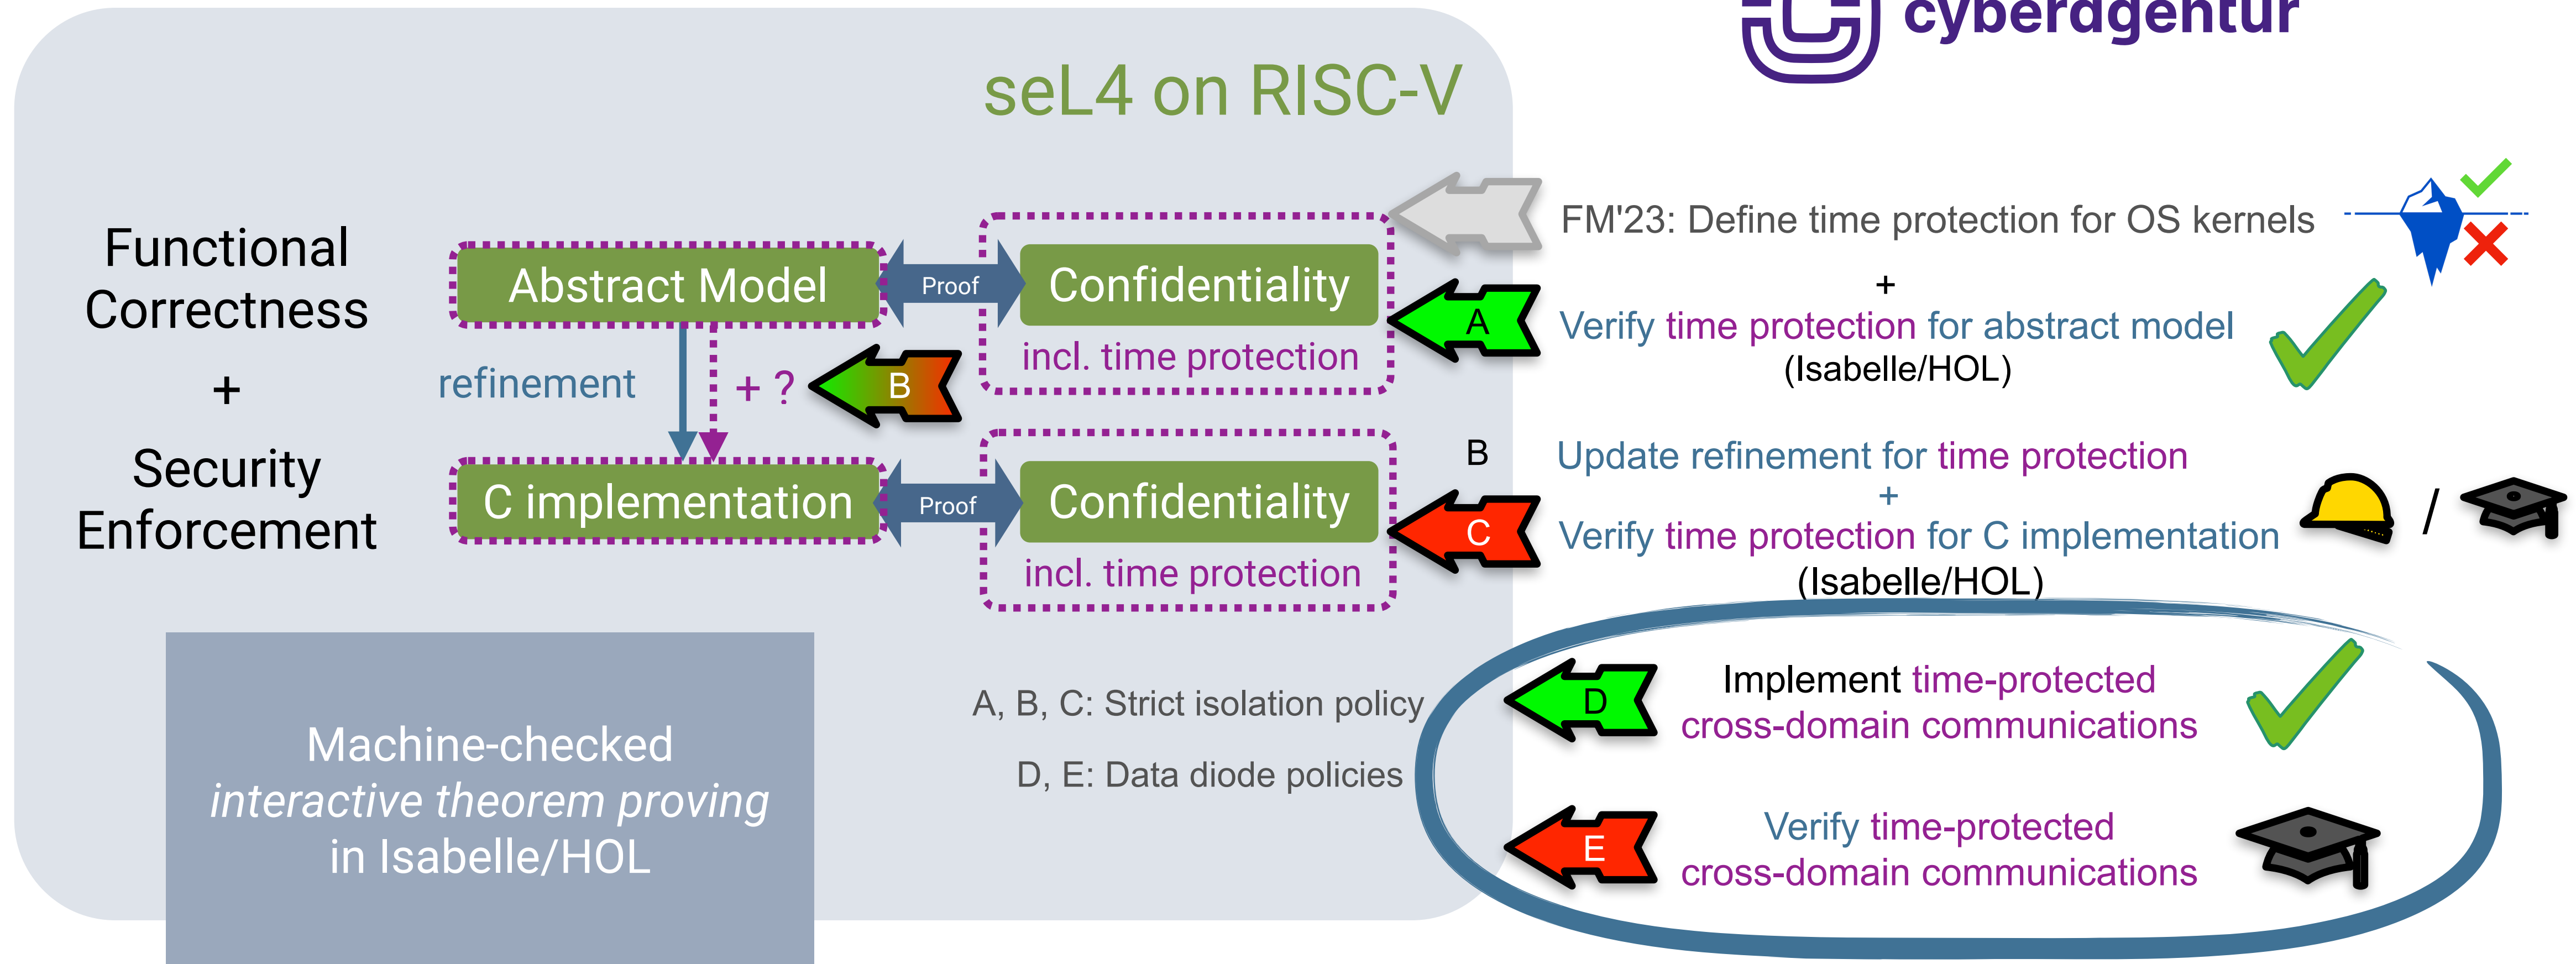
Need spatial precision to allow some flows but exclude others



Our FM'23 paper [Sison et al. 2023] allows infowall policies with arbitrary spatial and temporal precision

The status today:

Thanks to funding from



The status today:

Thank you!
Q&A

Thanks to funding from

