

Relaxed

Virtual

Memory

— ARMv8 edition —

Ben Simmer¹, Alasdair Armstrong¹, Thibaut Pérami¹

Jean Pichon-Pharabal², Christopher Pulte¹

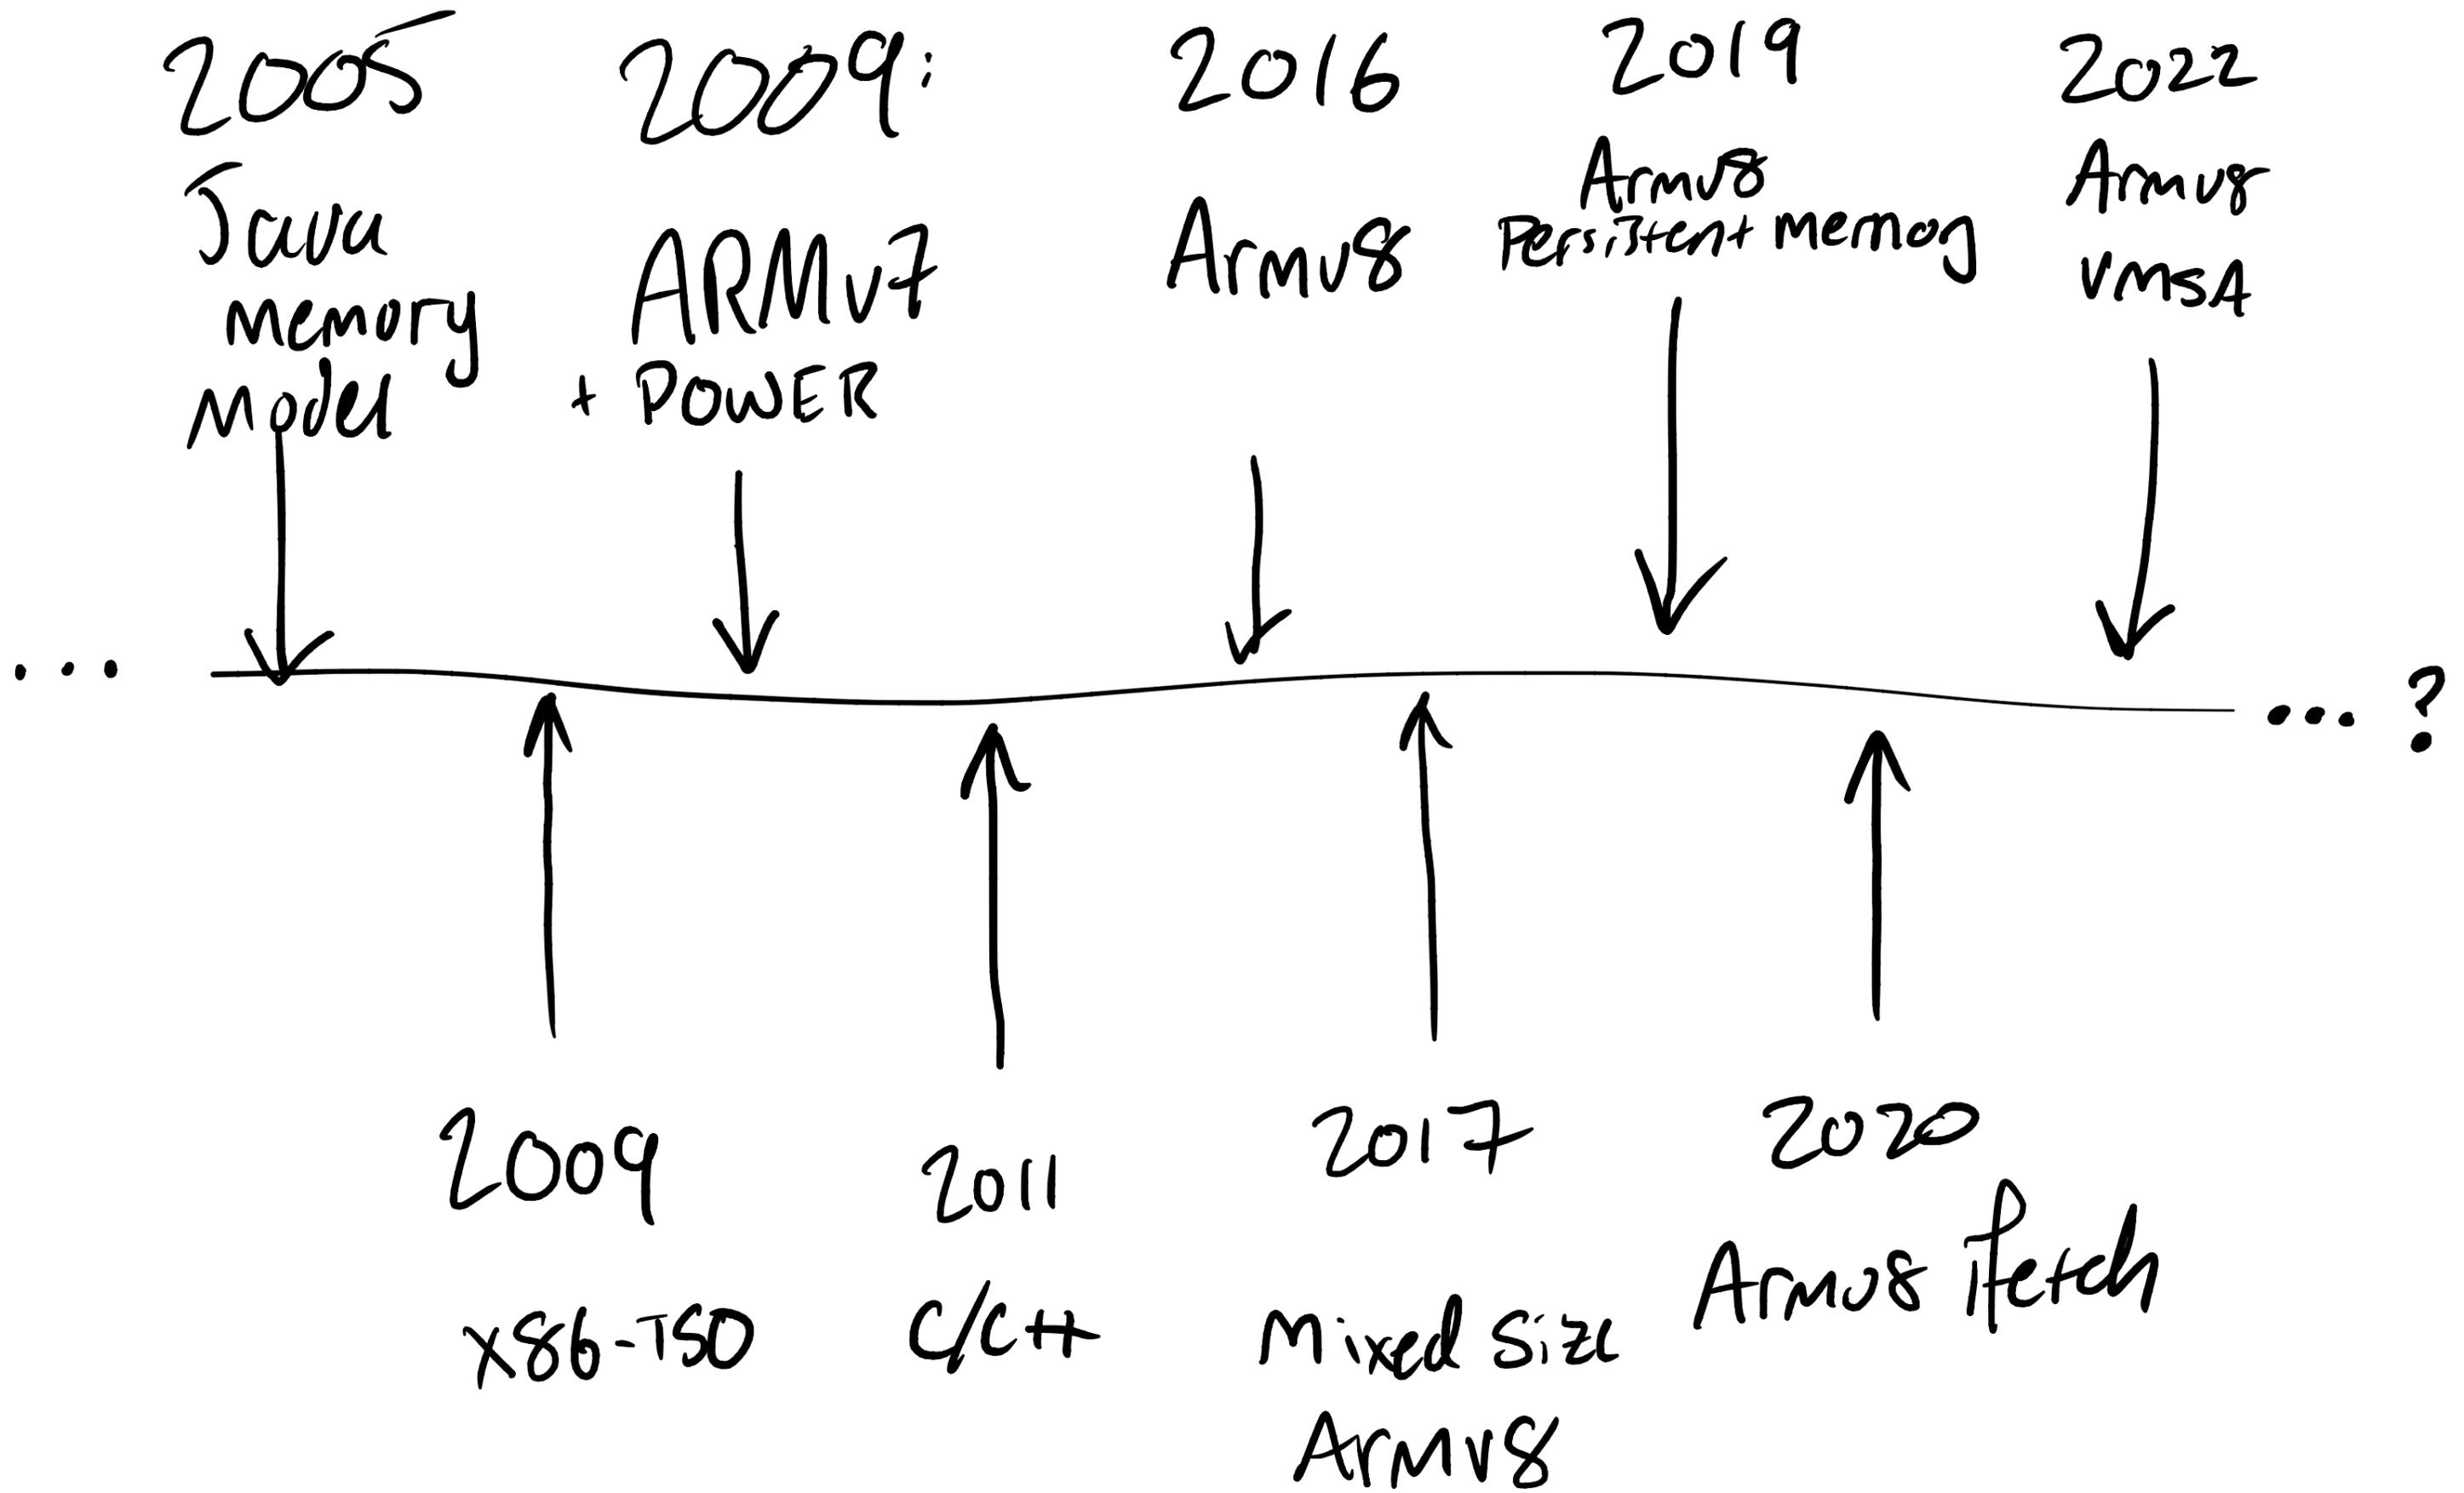
Richard Grisenthwaite³, Peter Sewell¹

¹ Cambridge

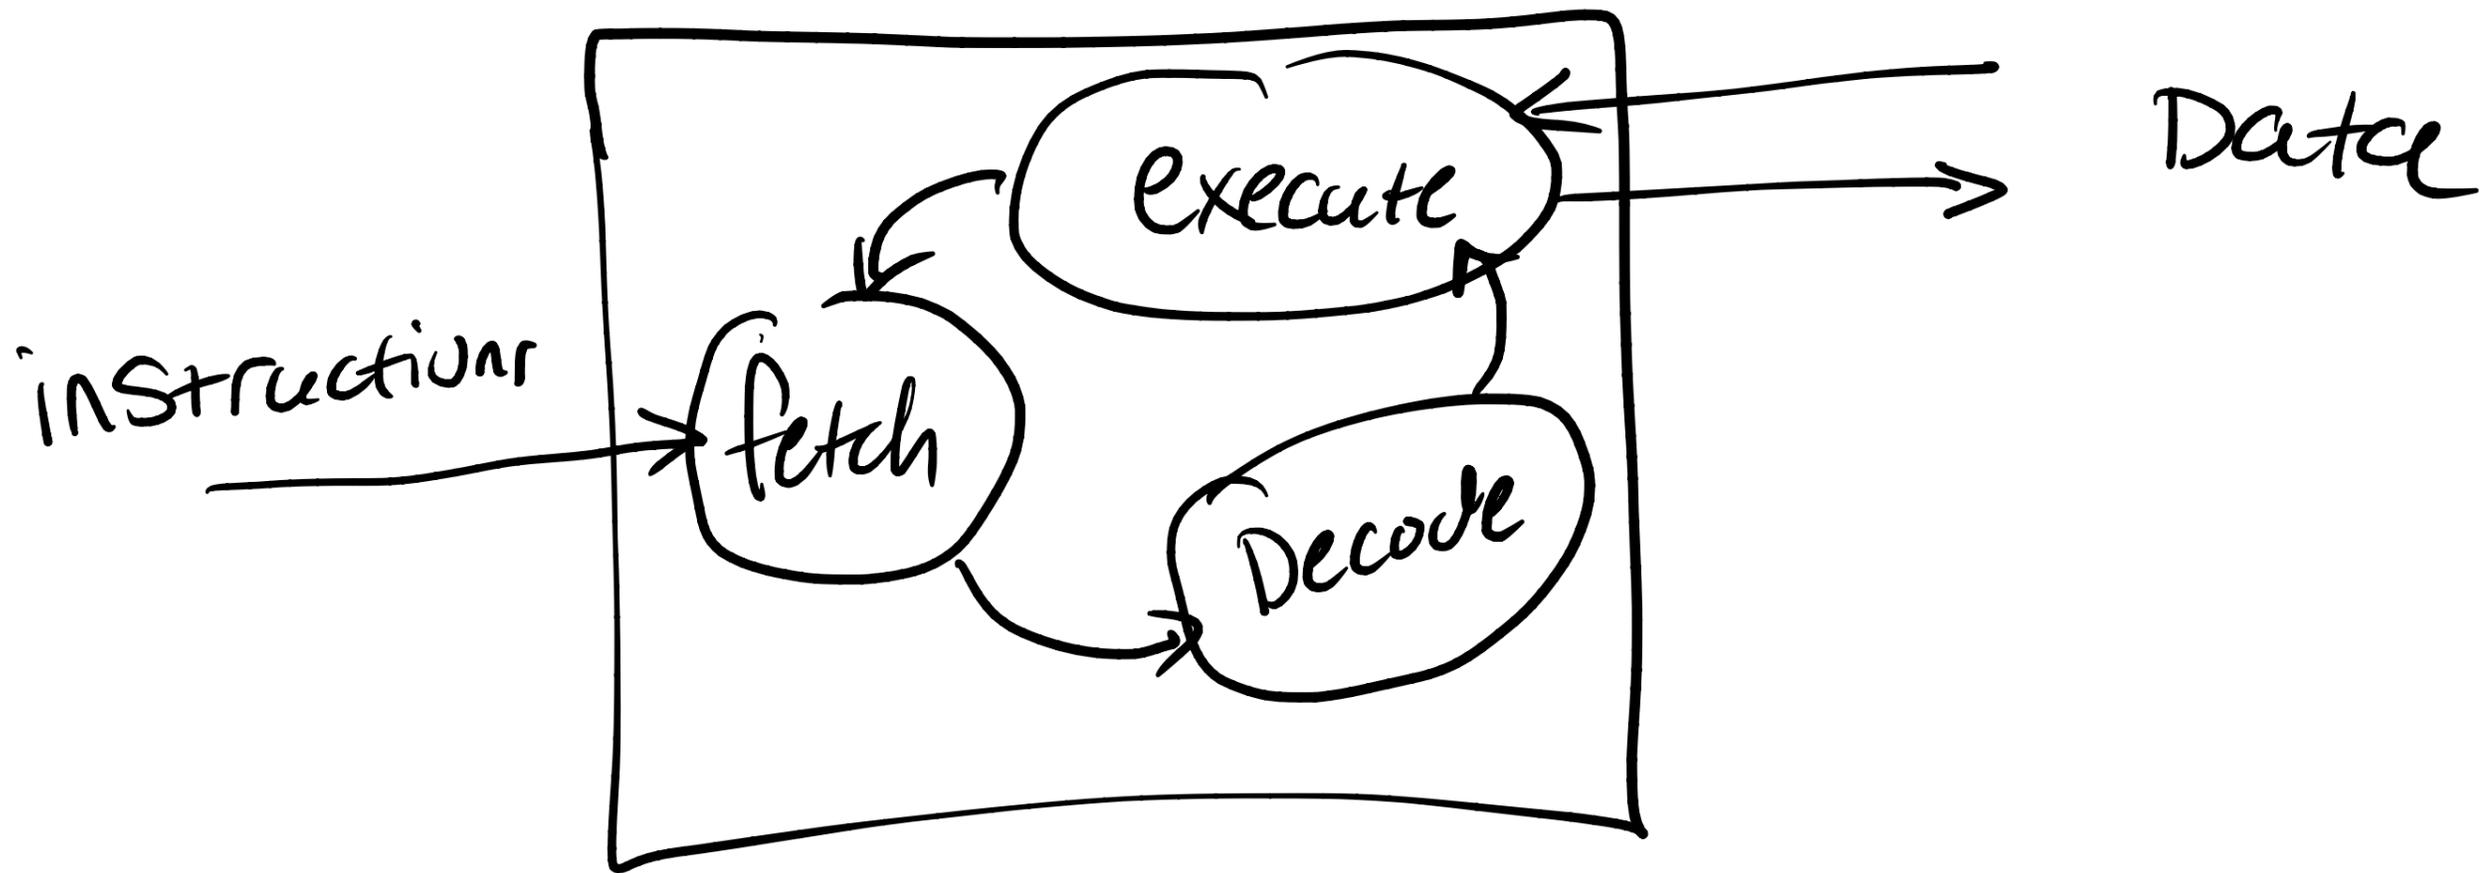
² Aarhus

³ ARM

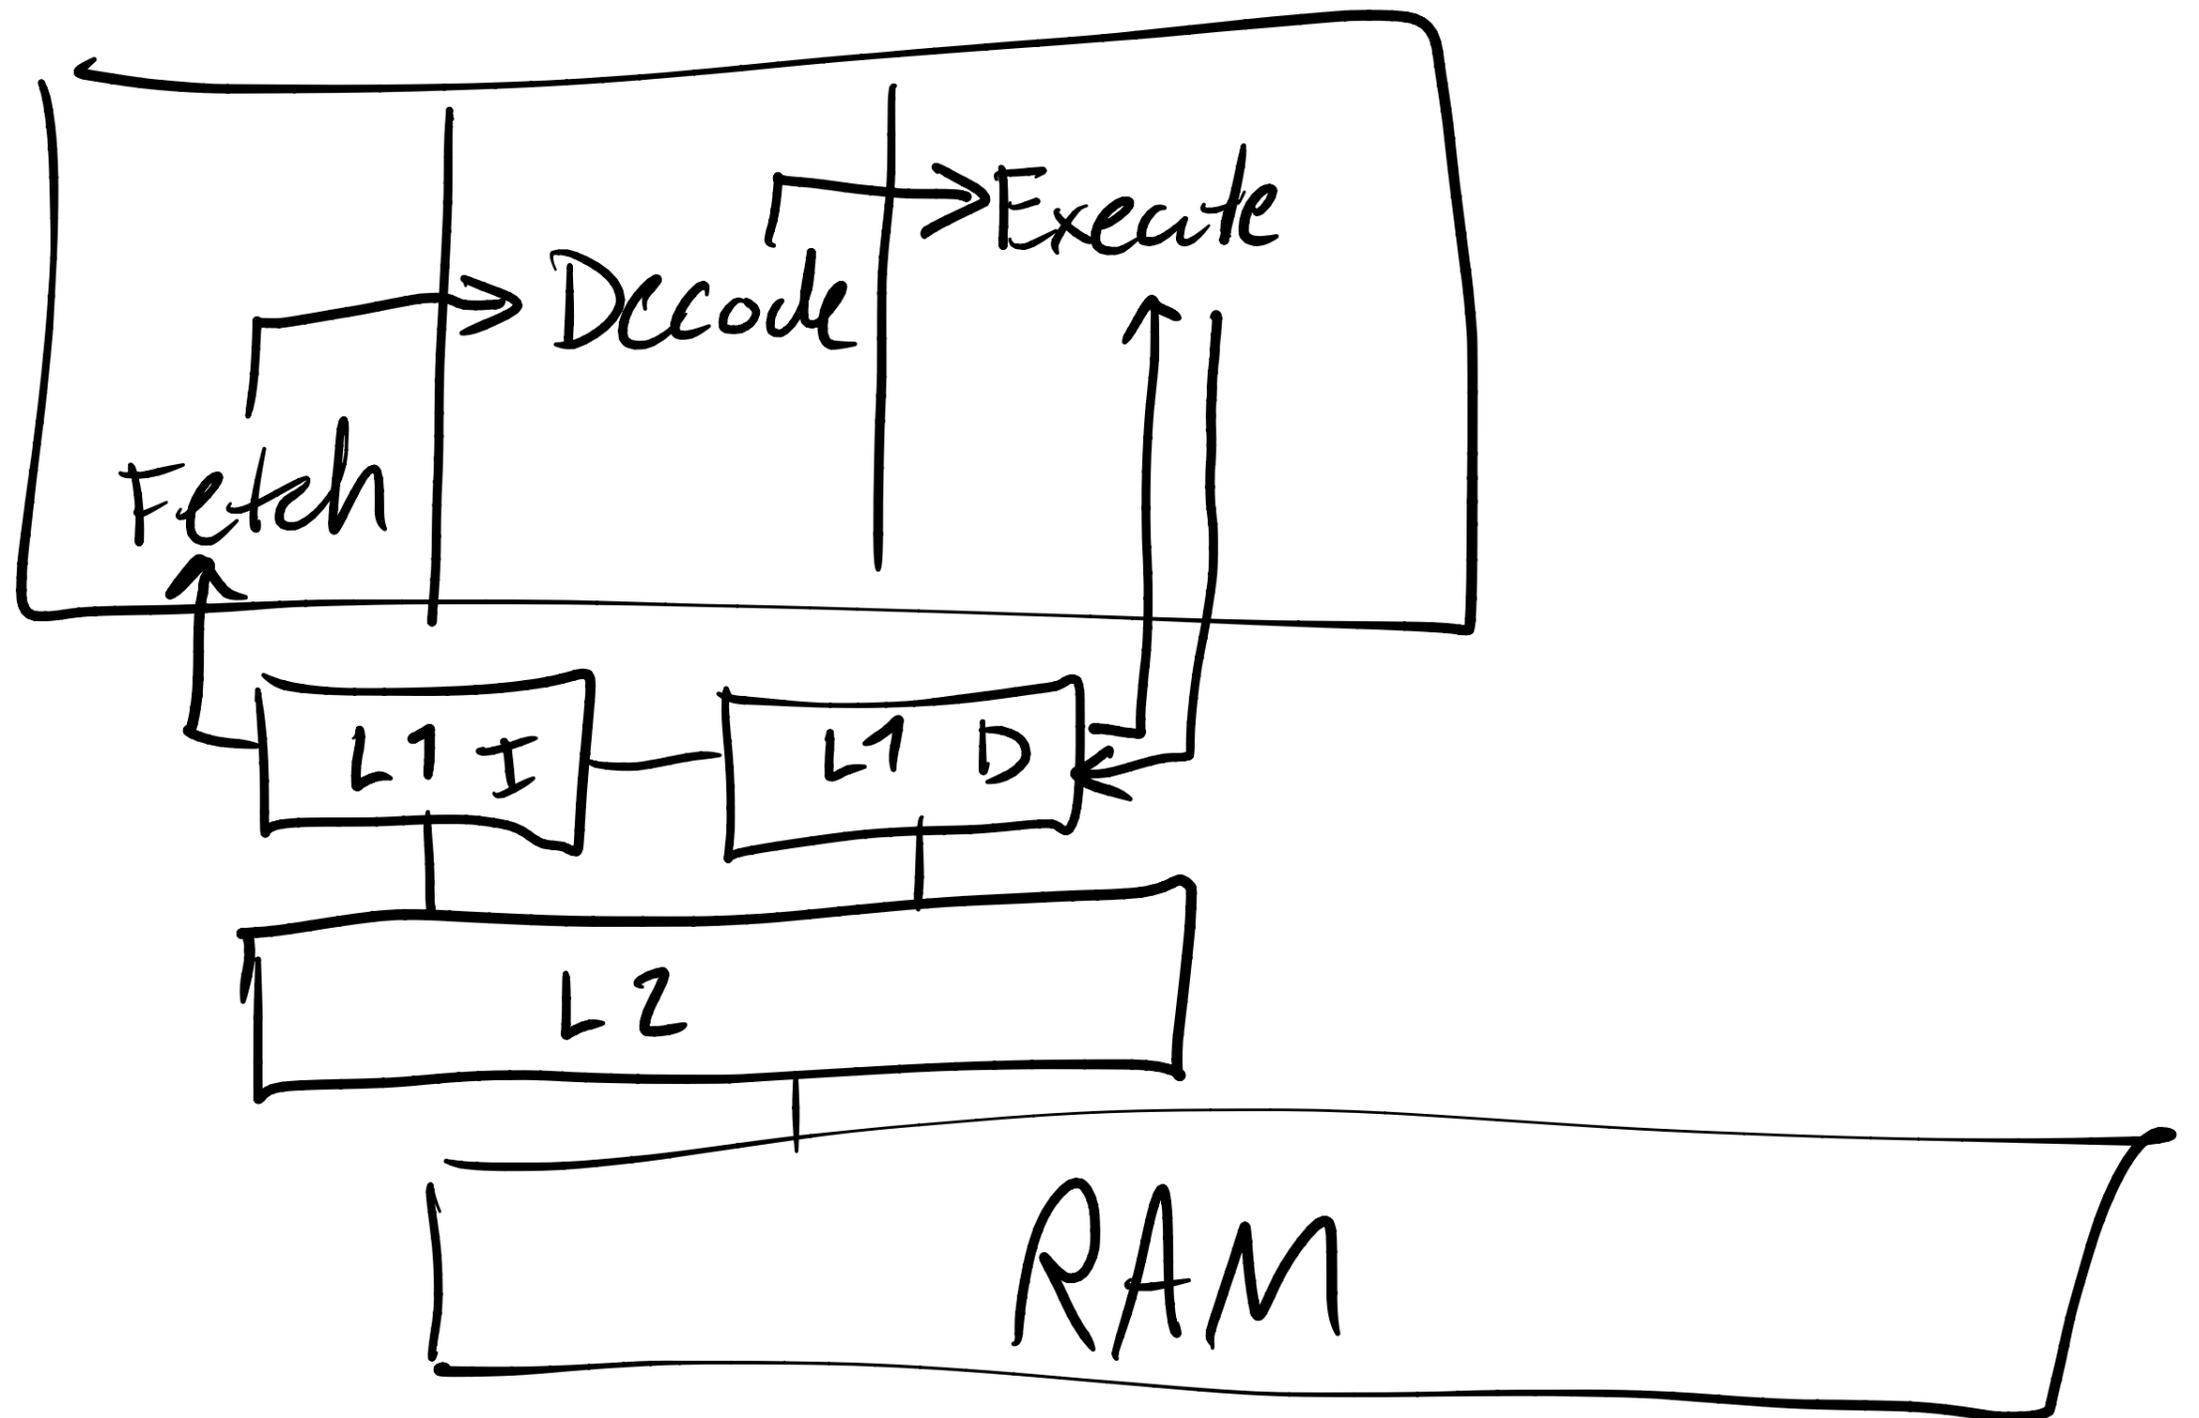
relaxed memory?

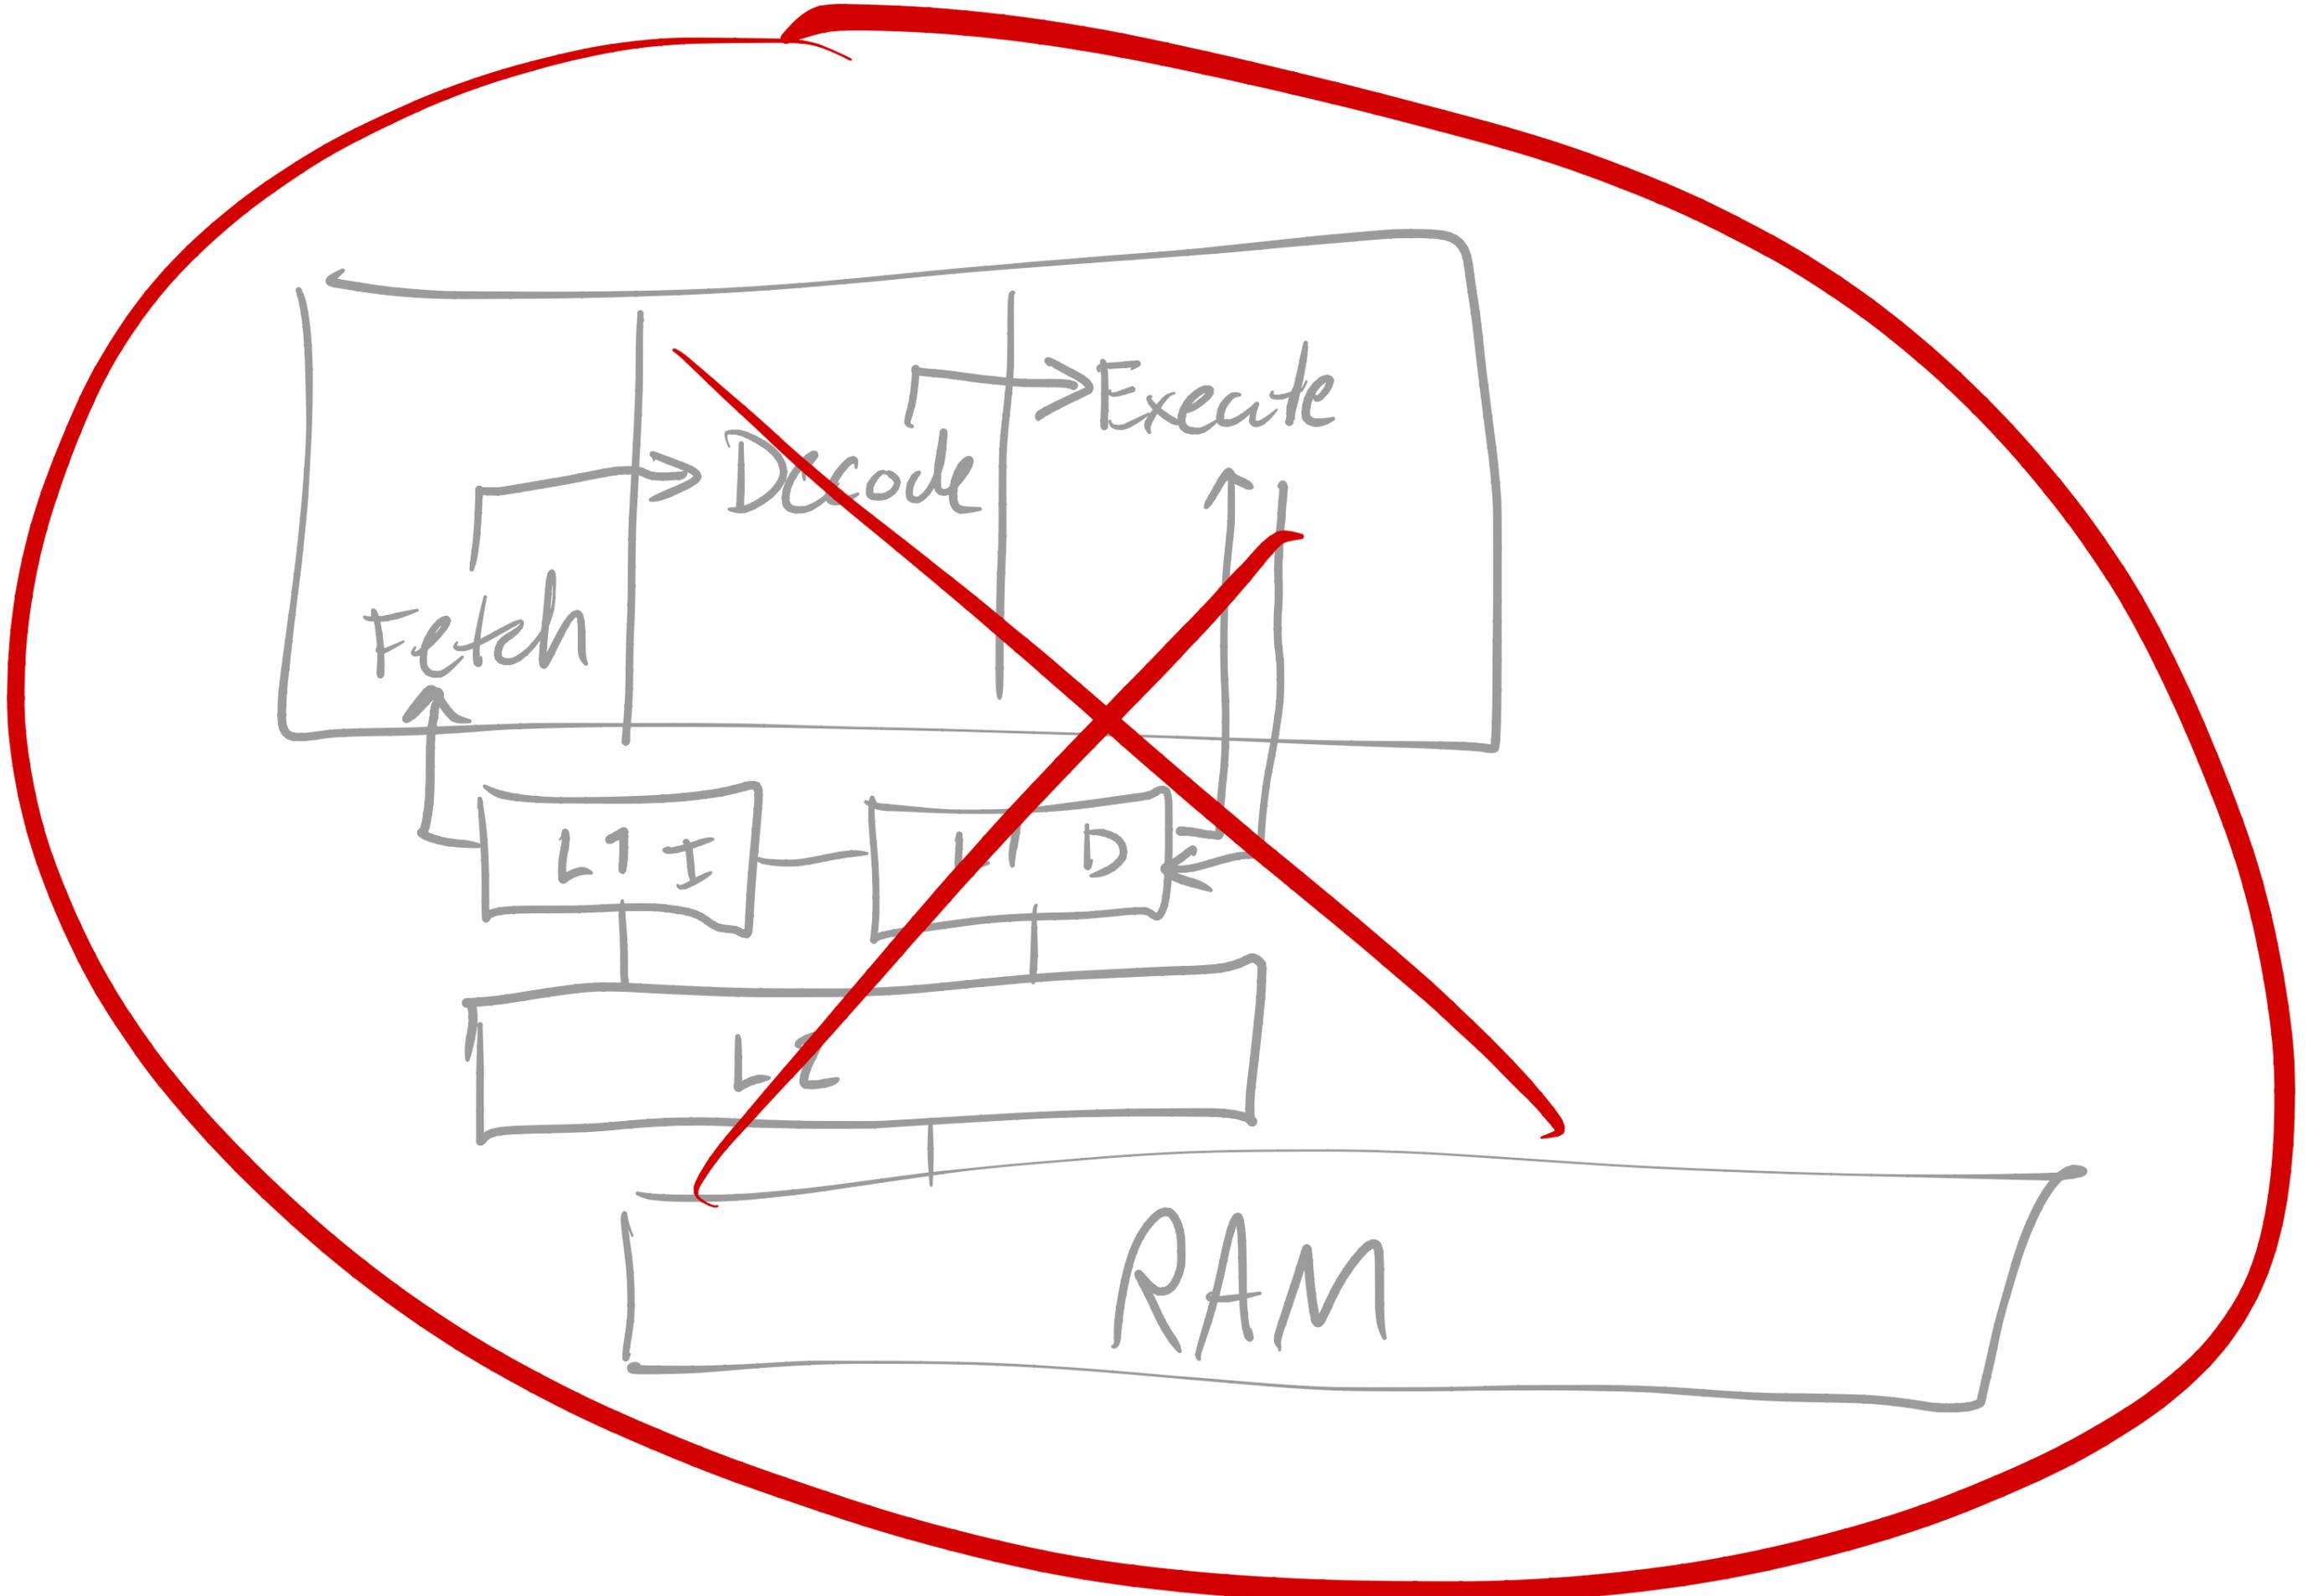


Hardware



A Better Cartoon?





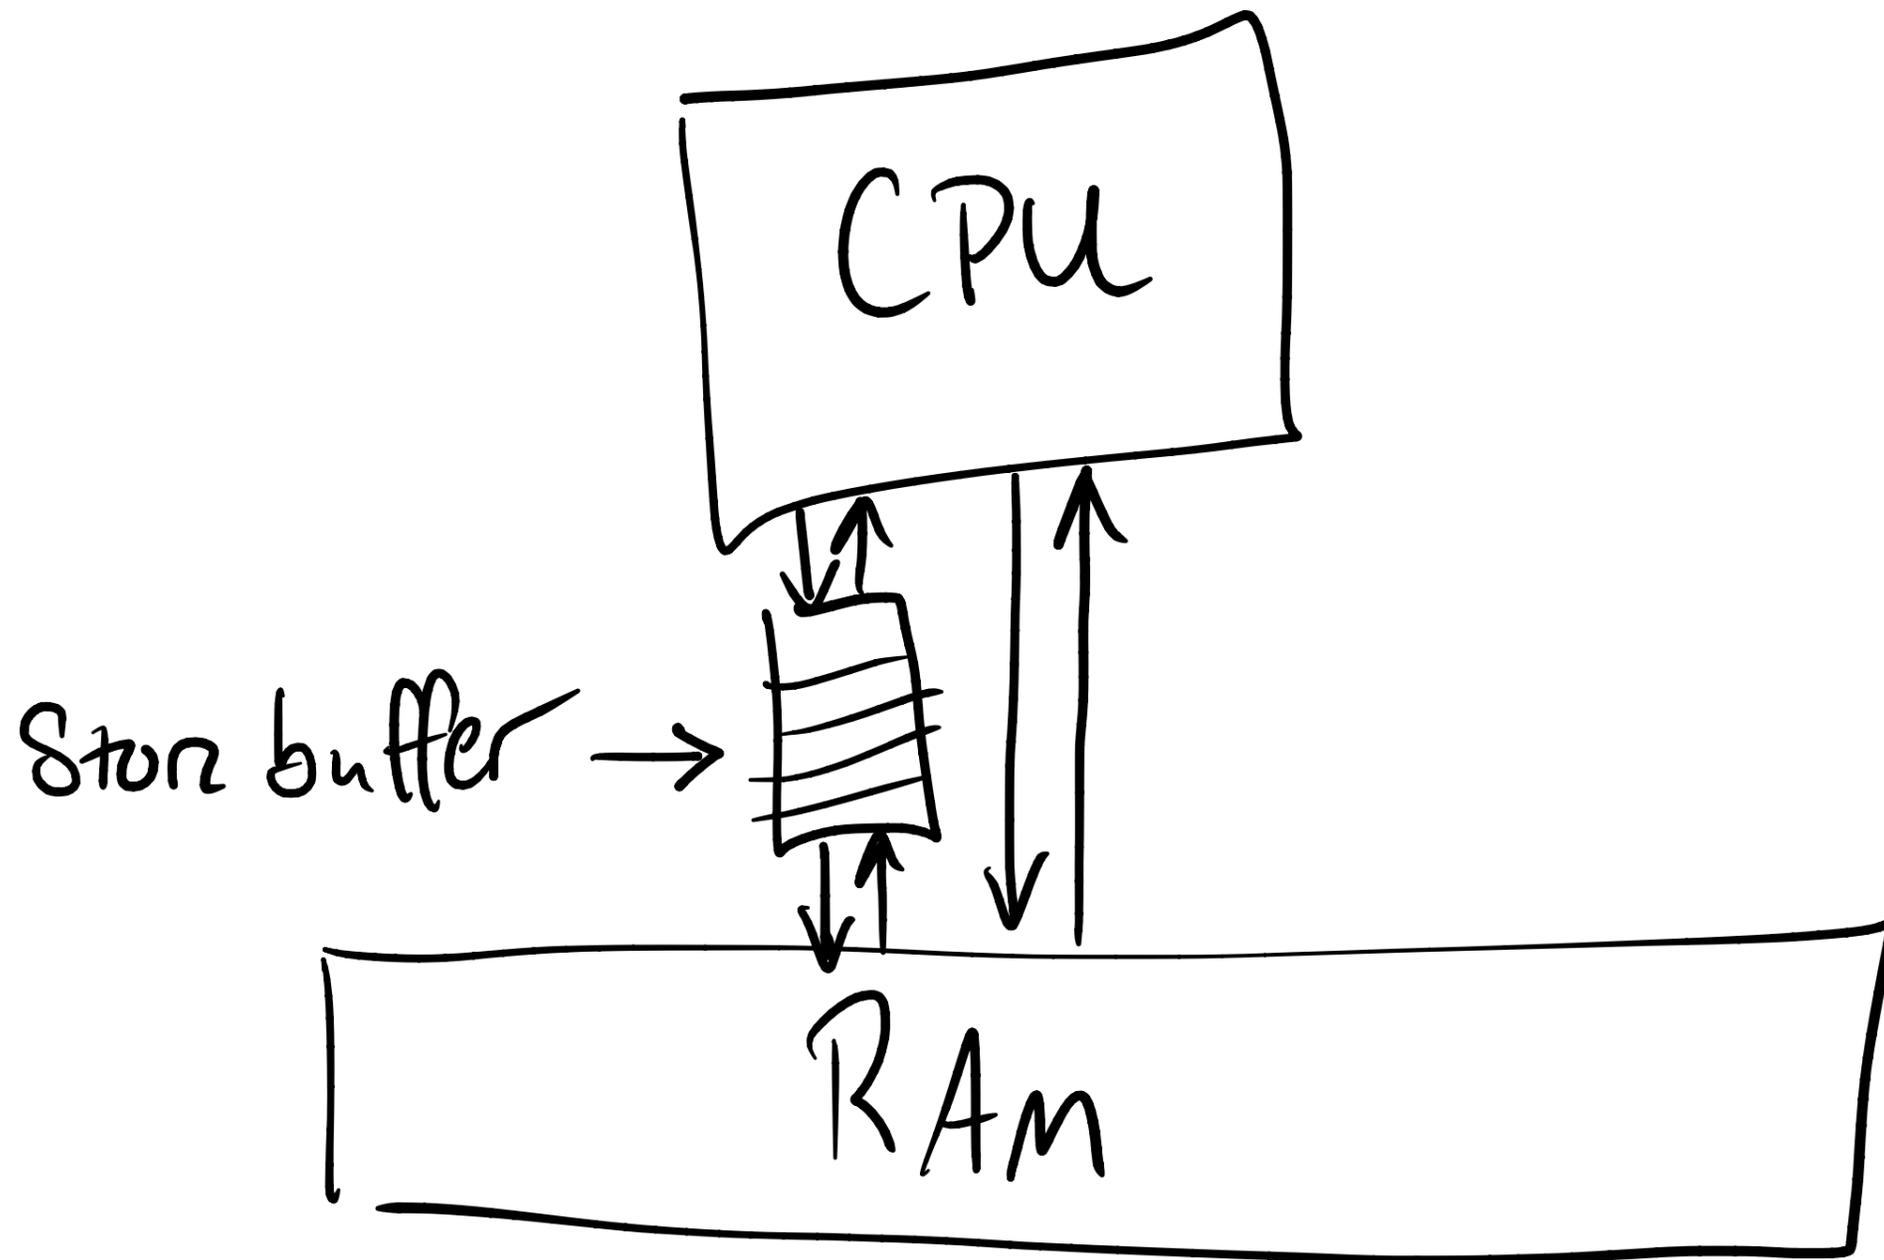
``

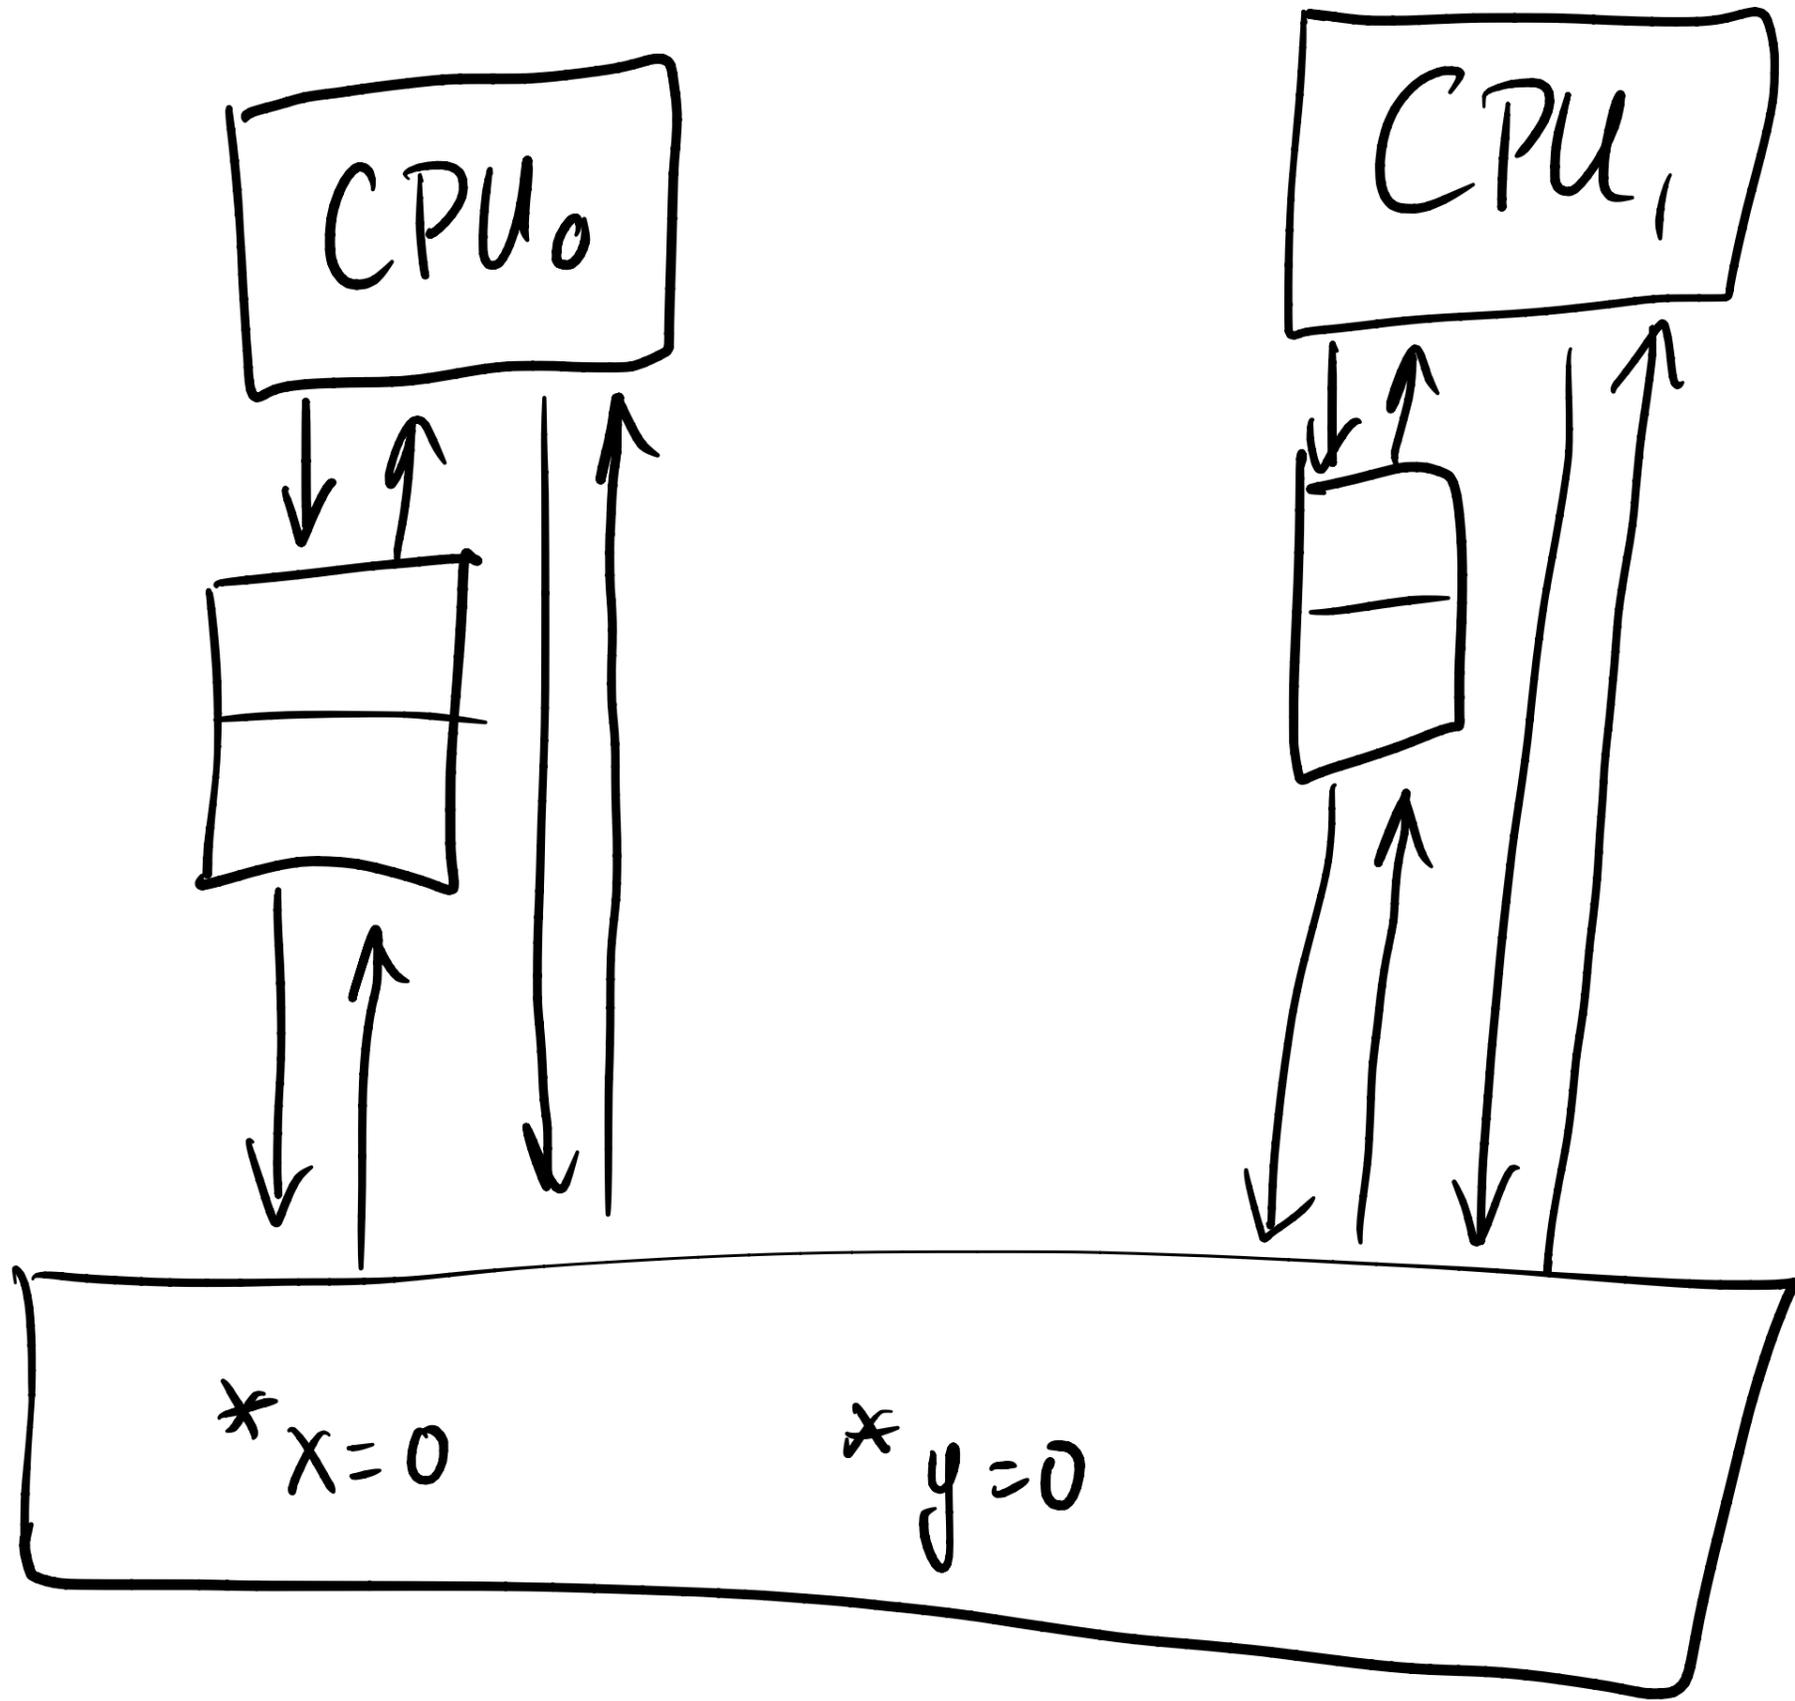
~~~~

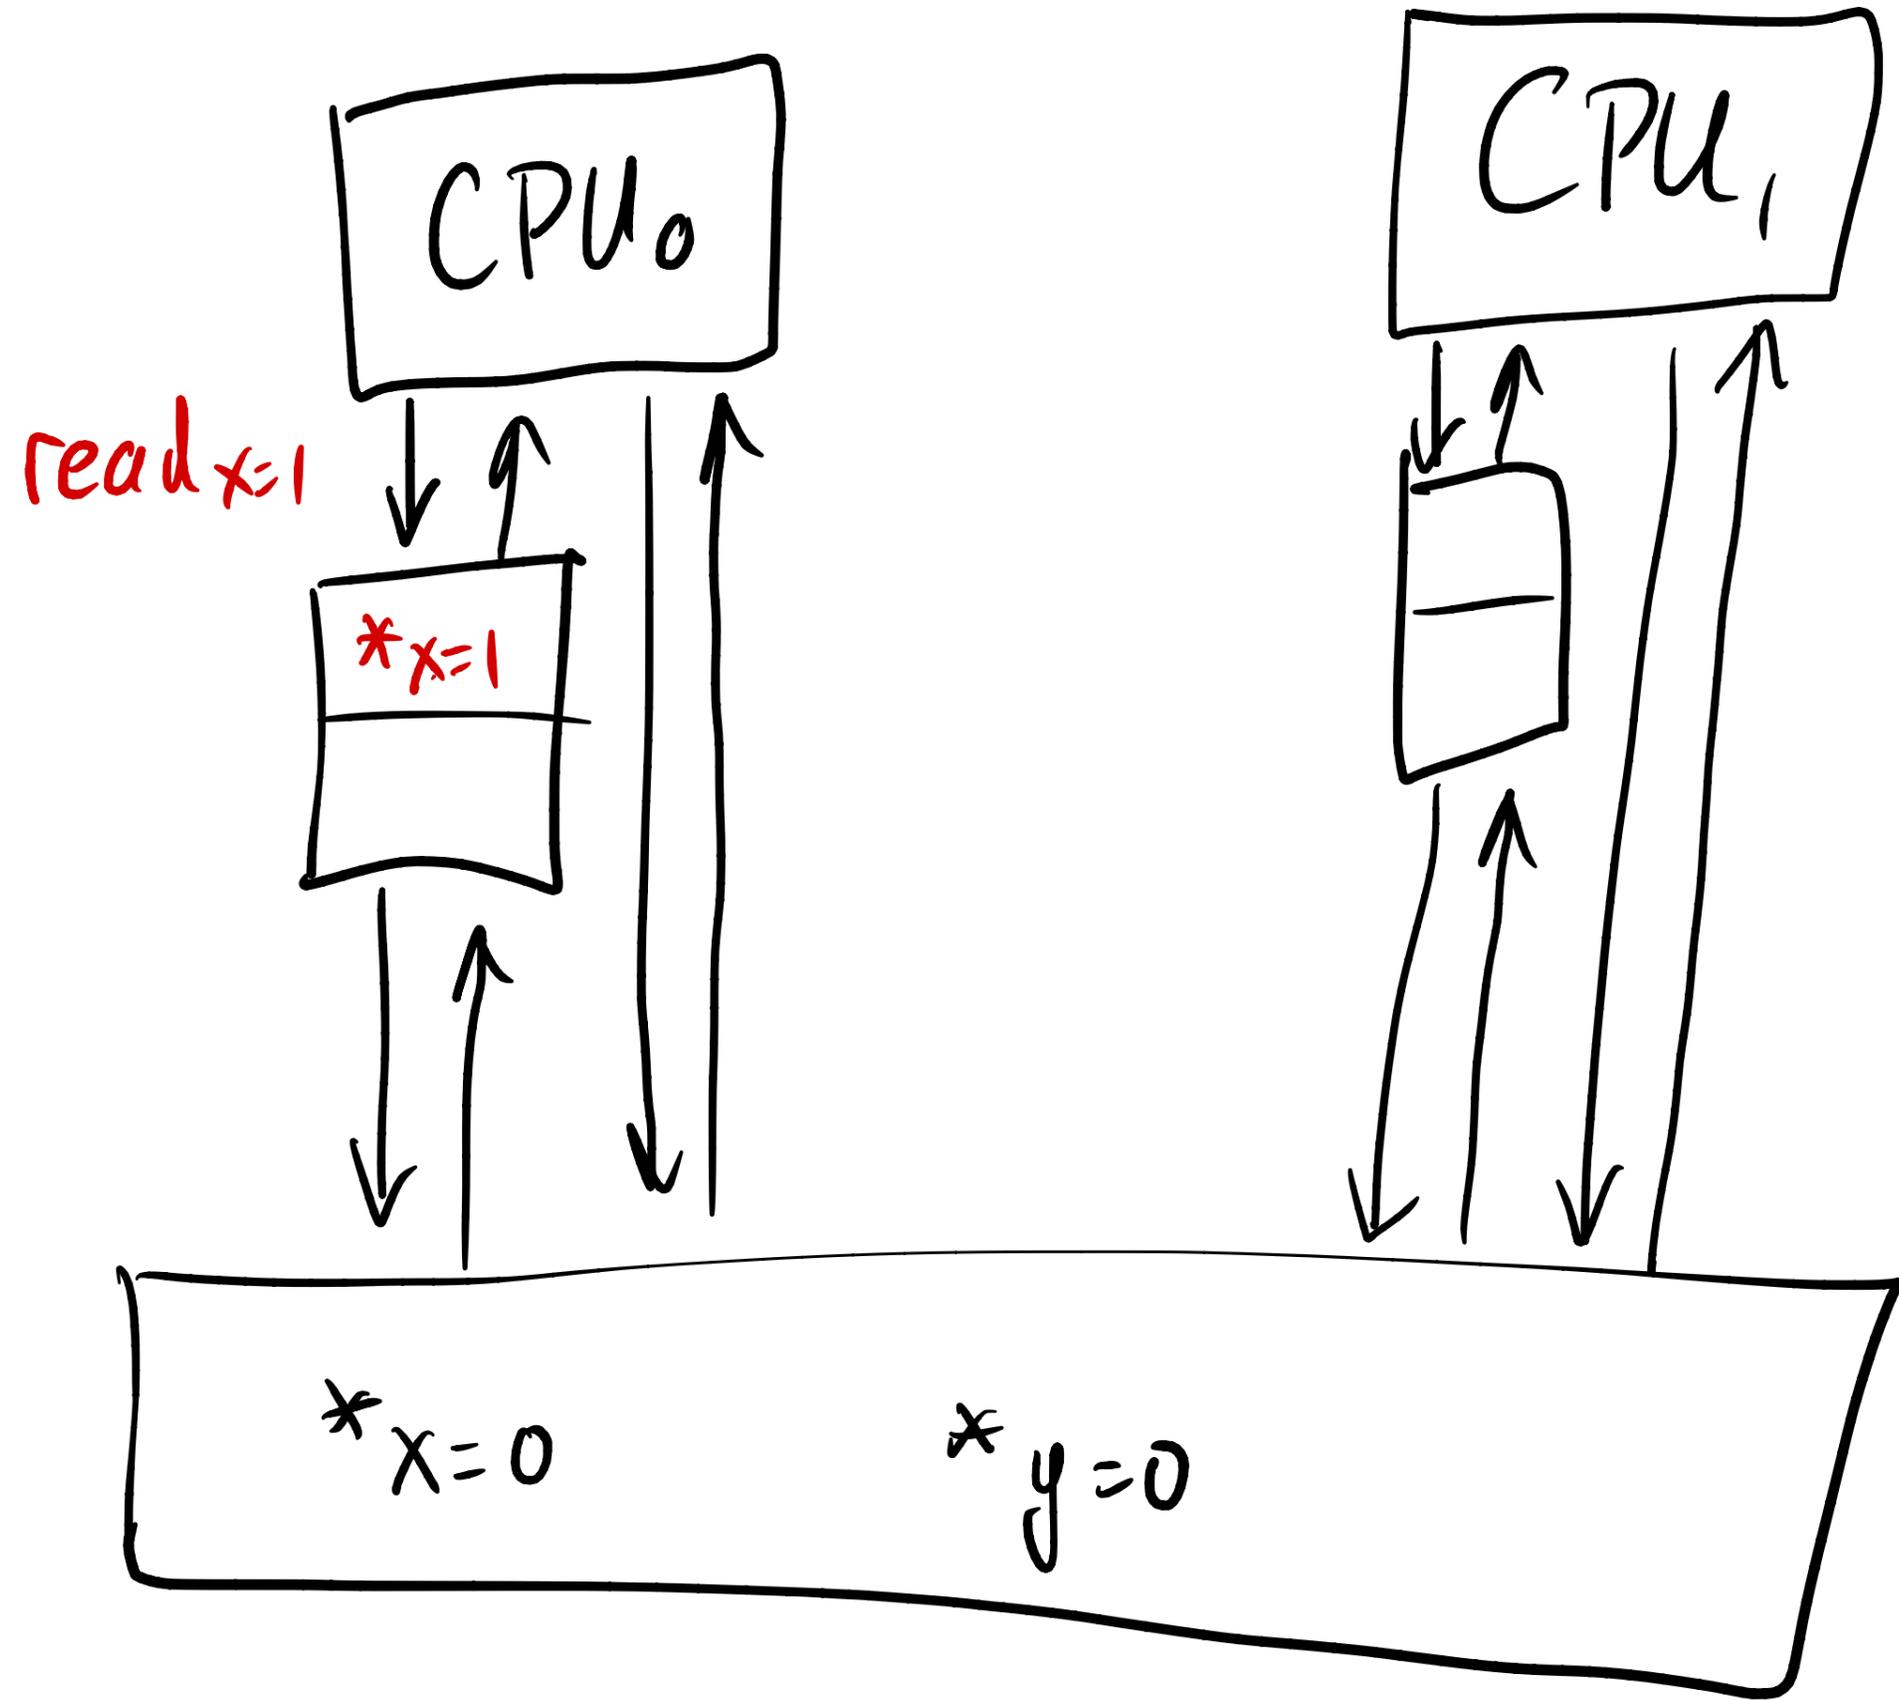
# Programmer-visible behaviour

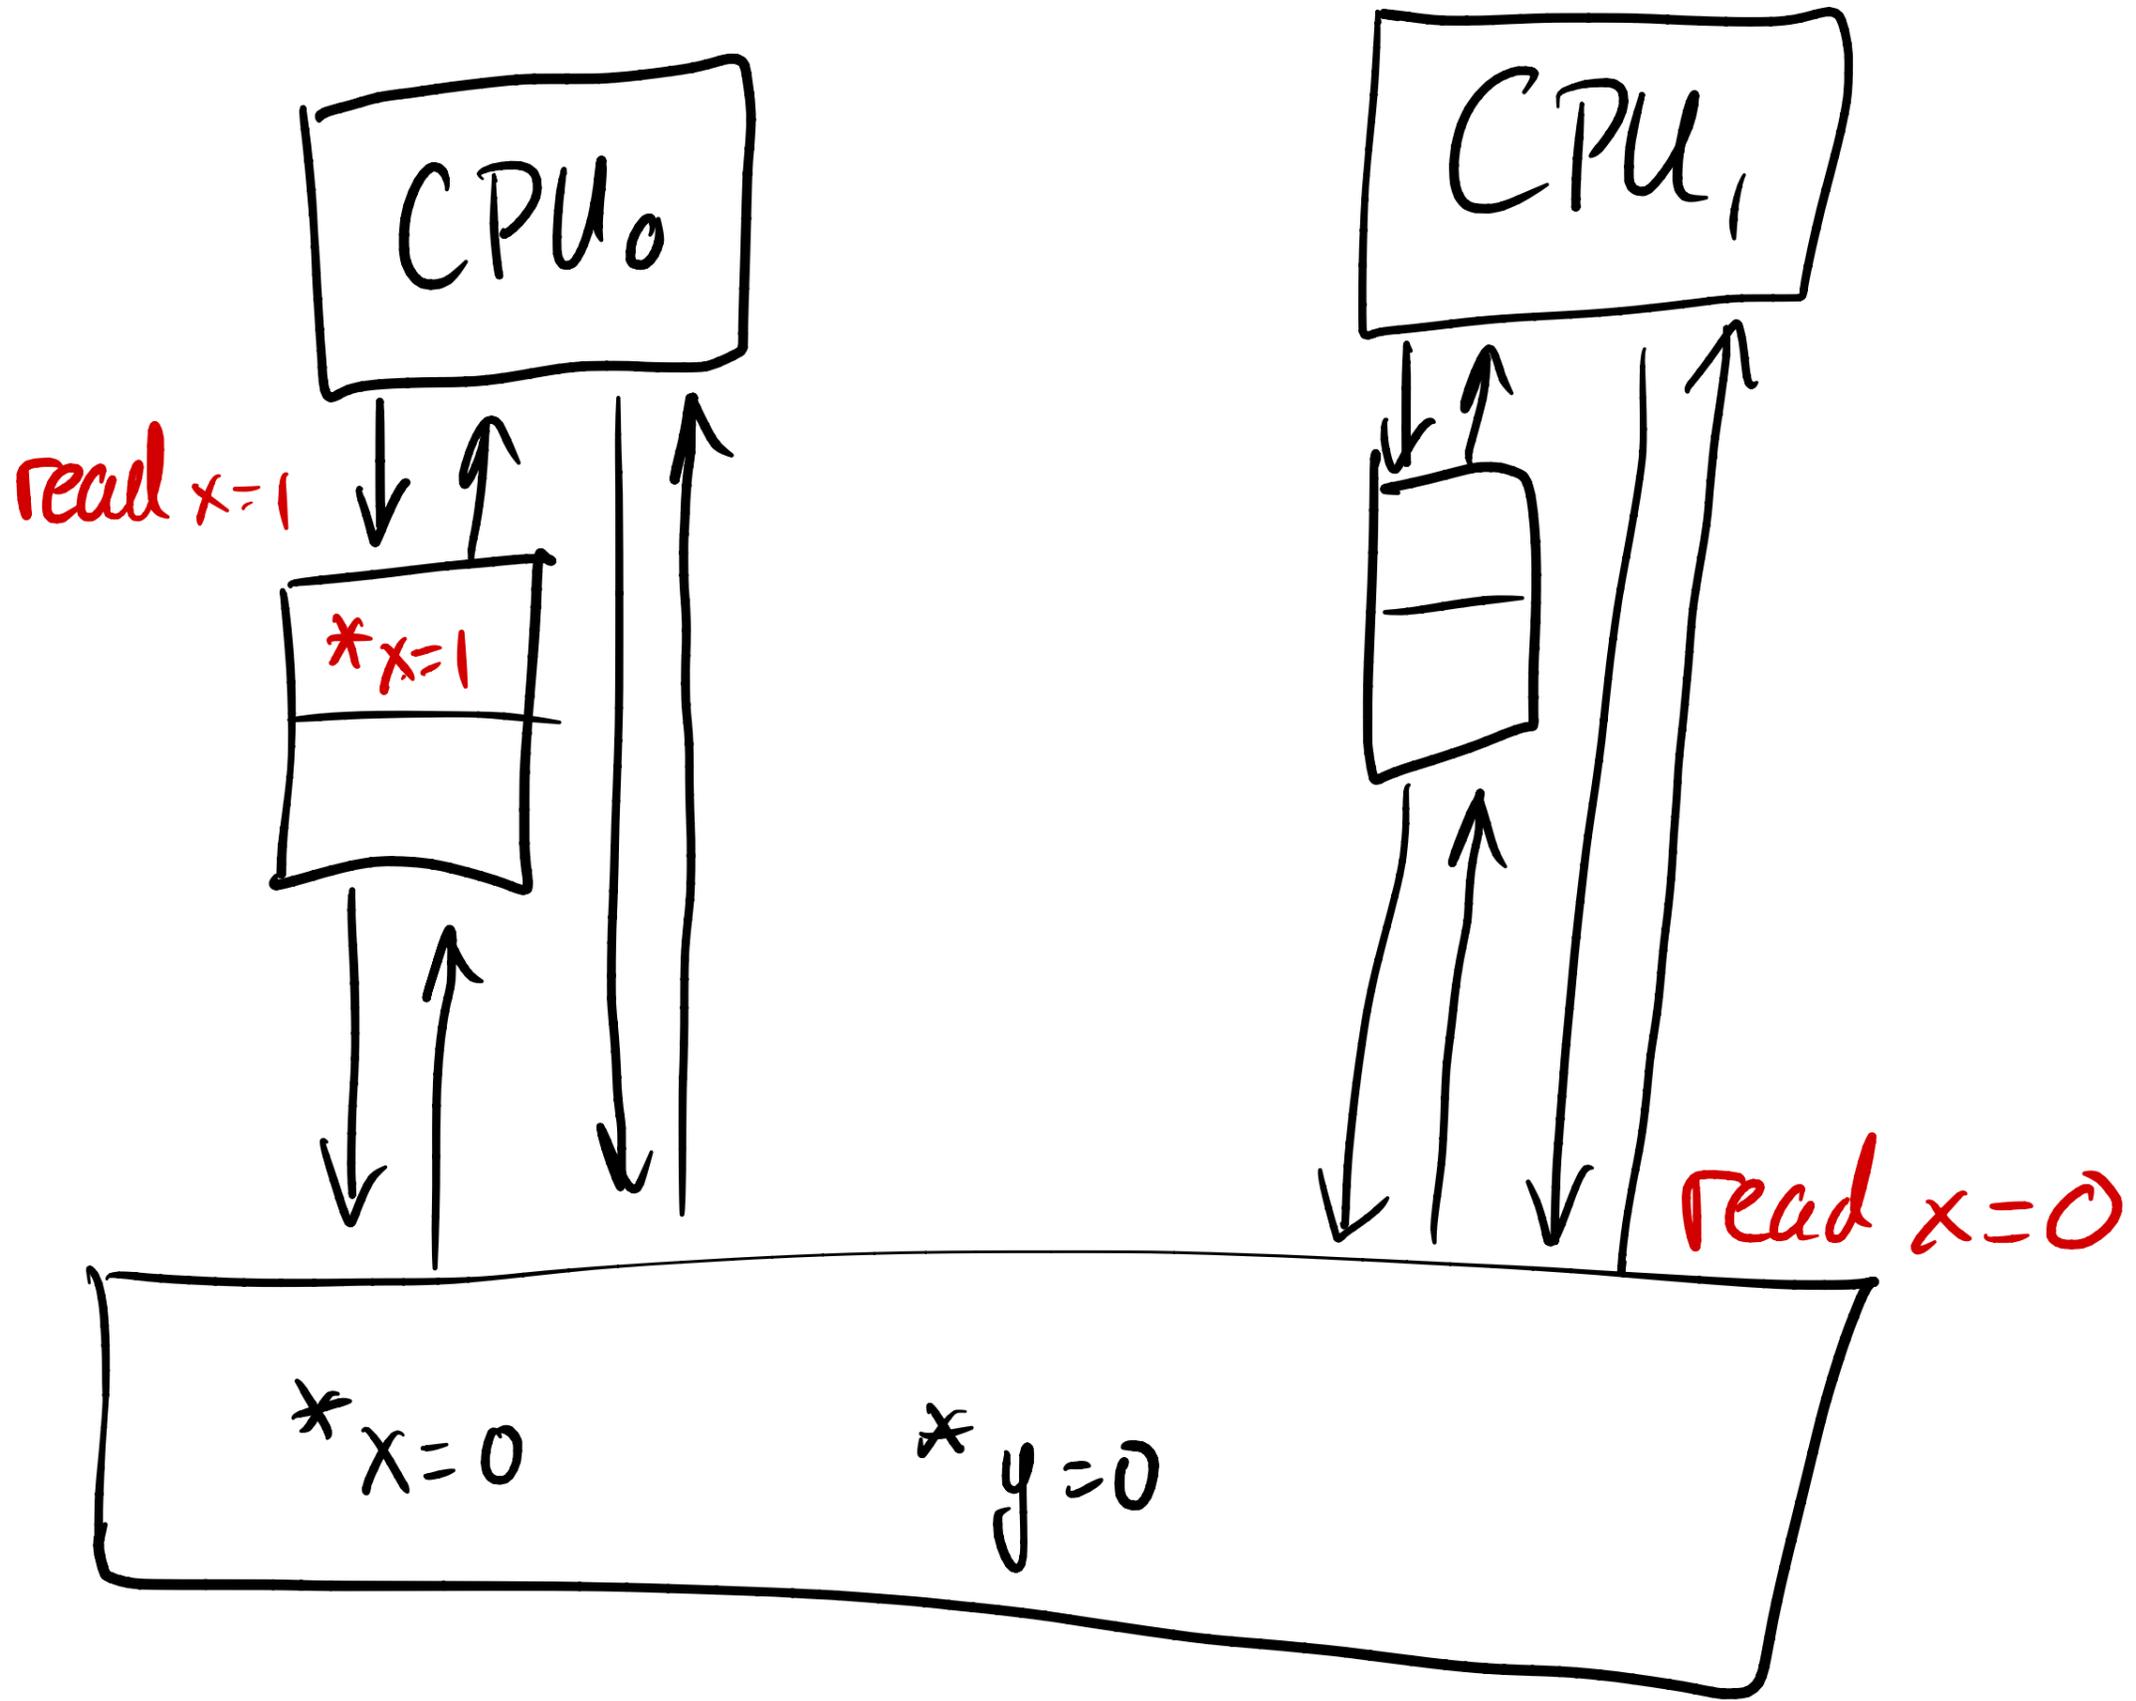
```
for (int i=0; i < N; i++) {  
    buf[i] = ...;  
}
```

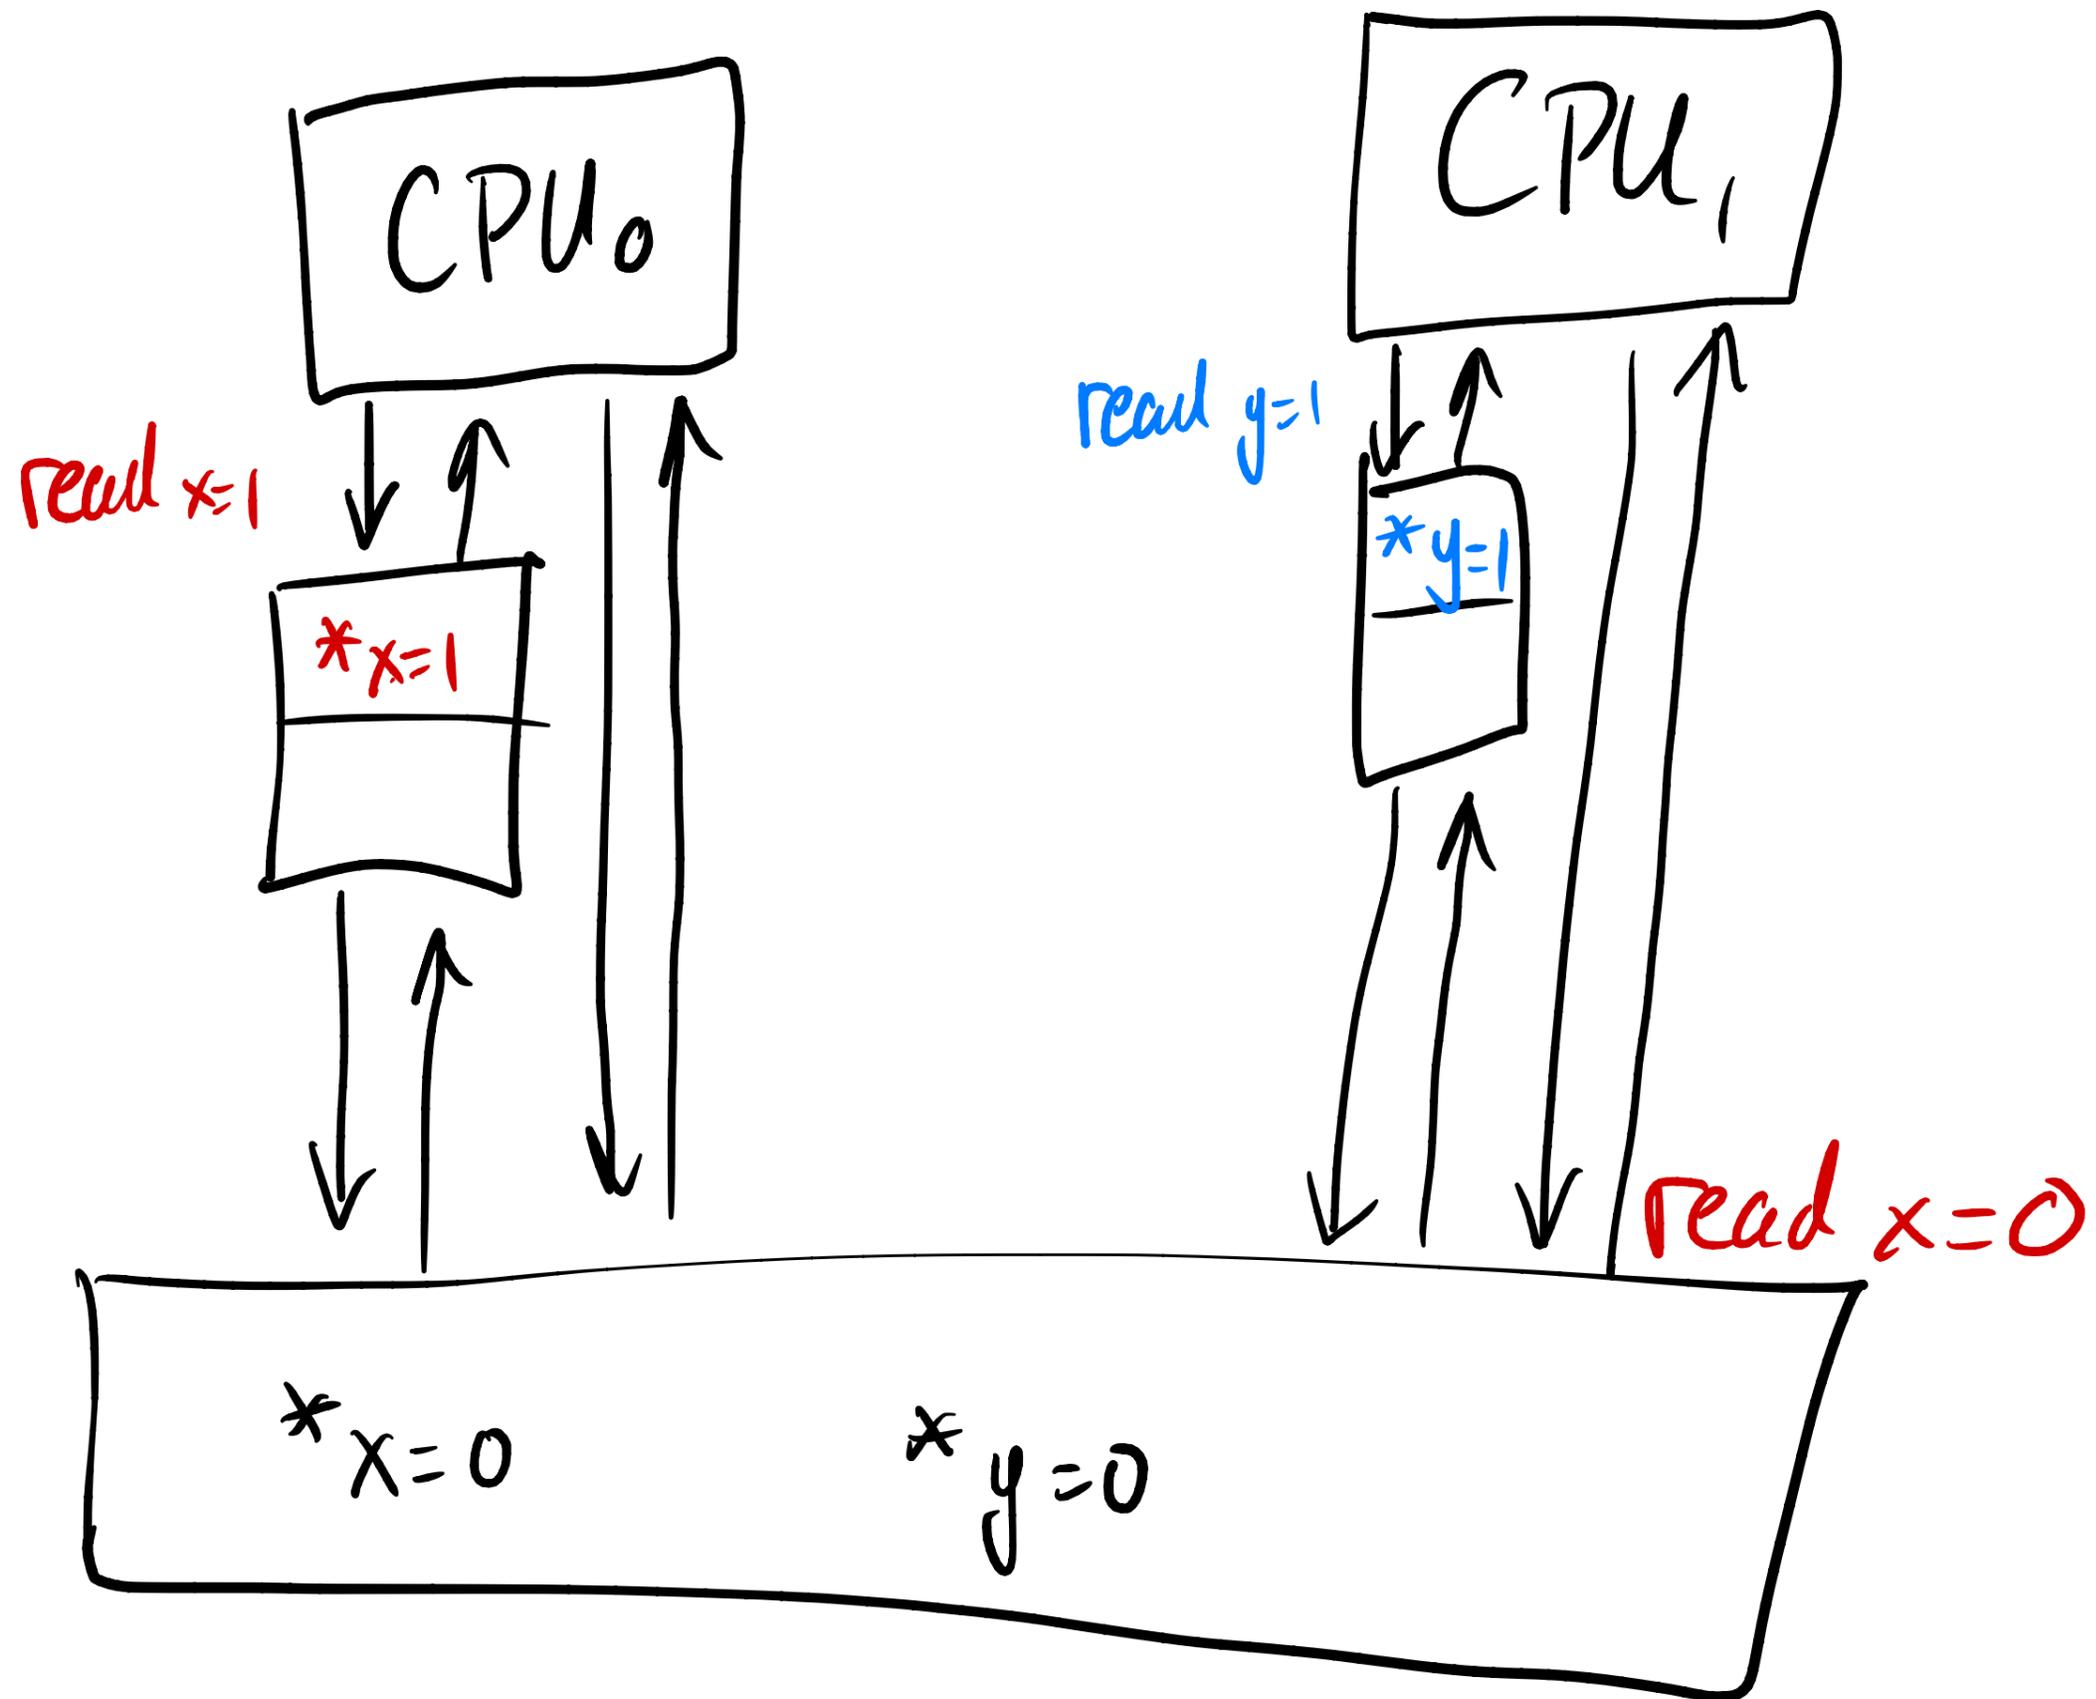
# Programmer-visible behaviour

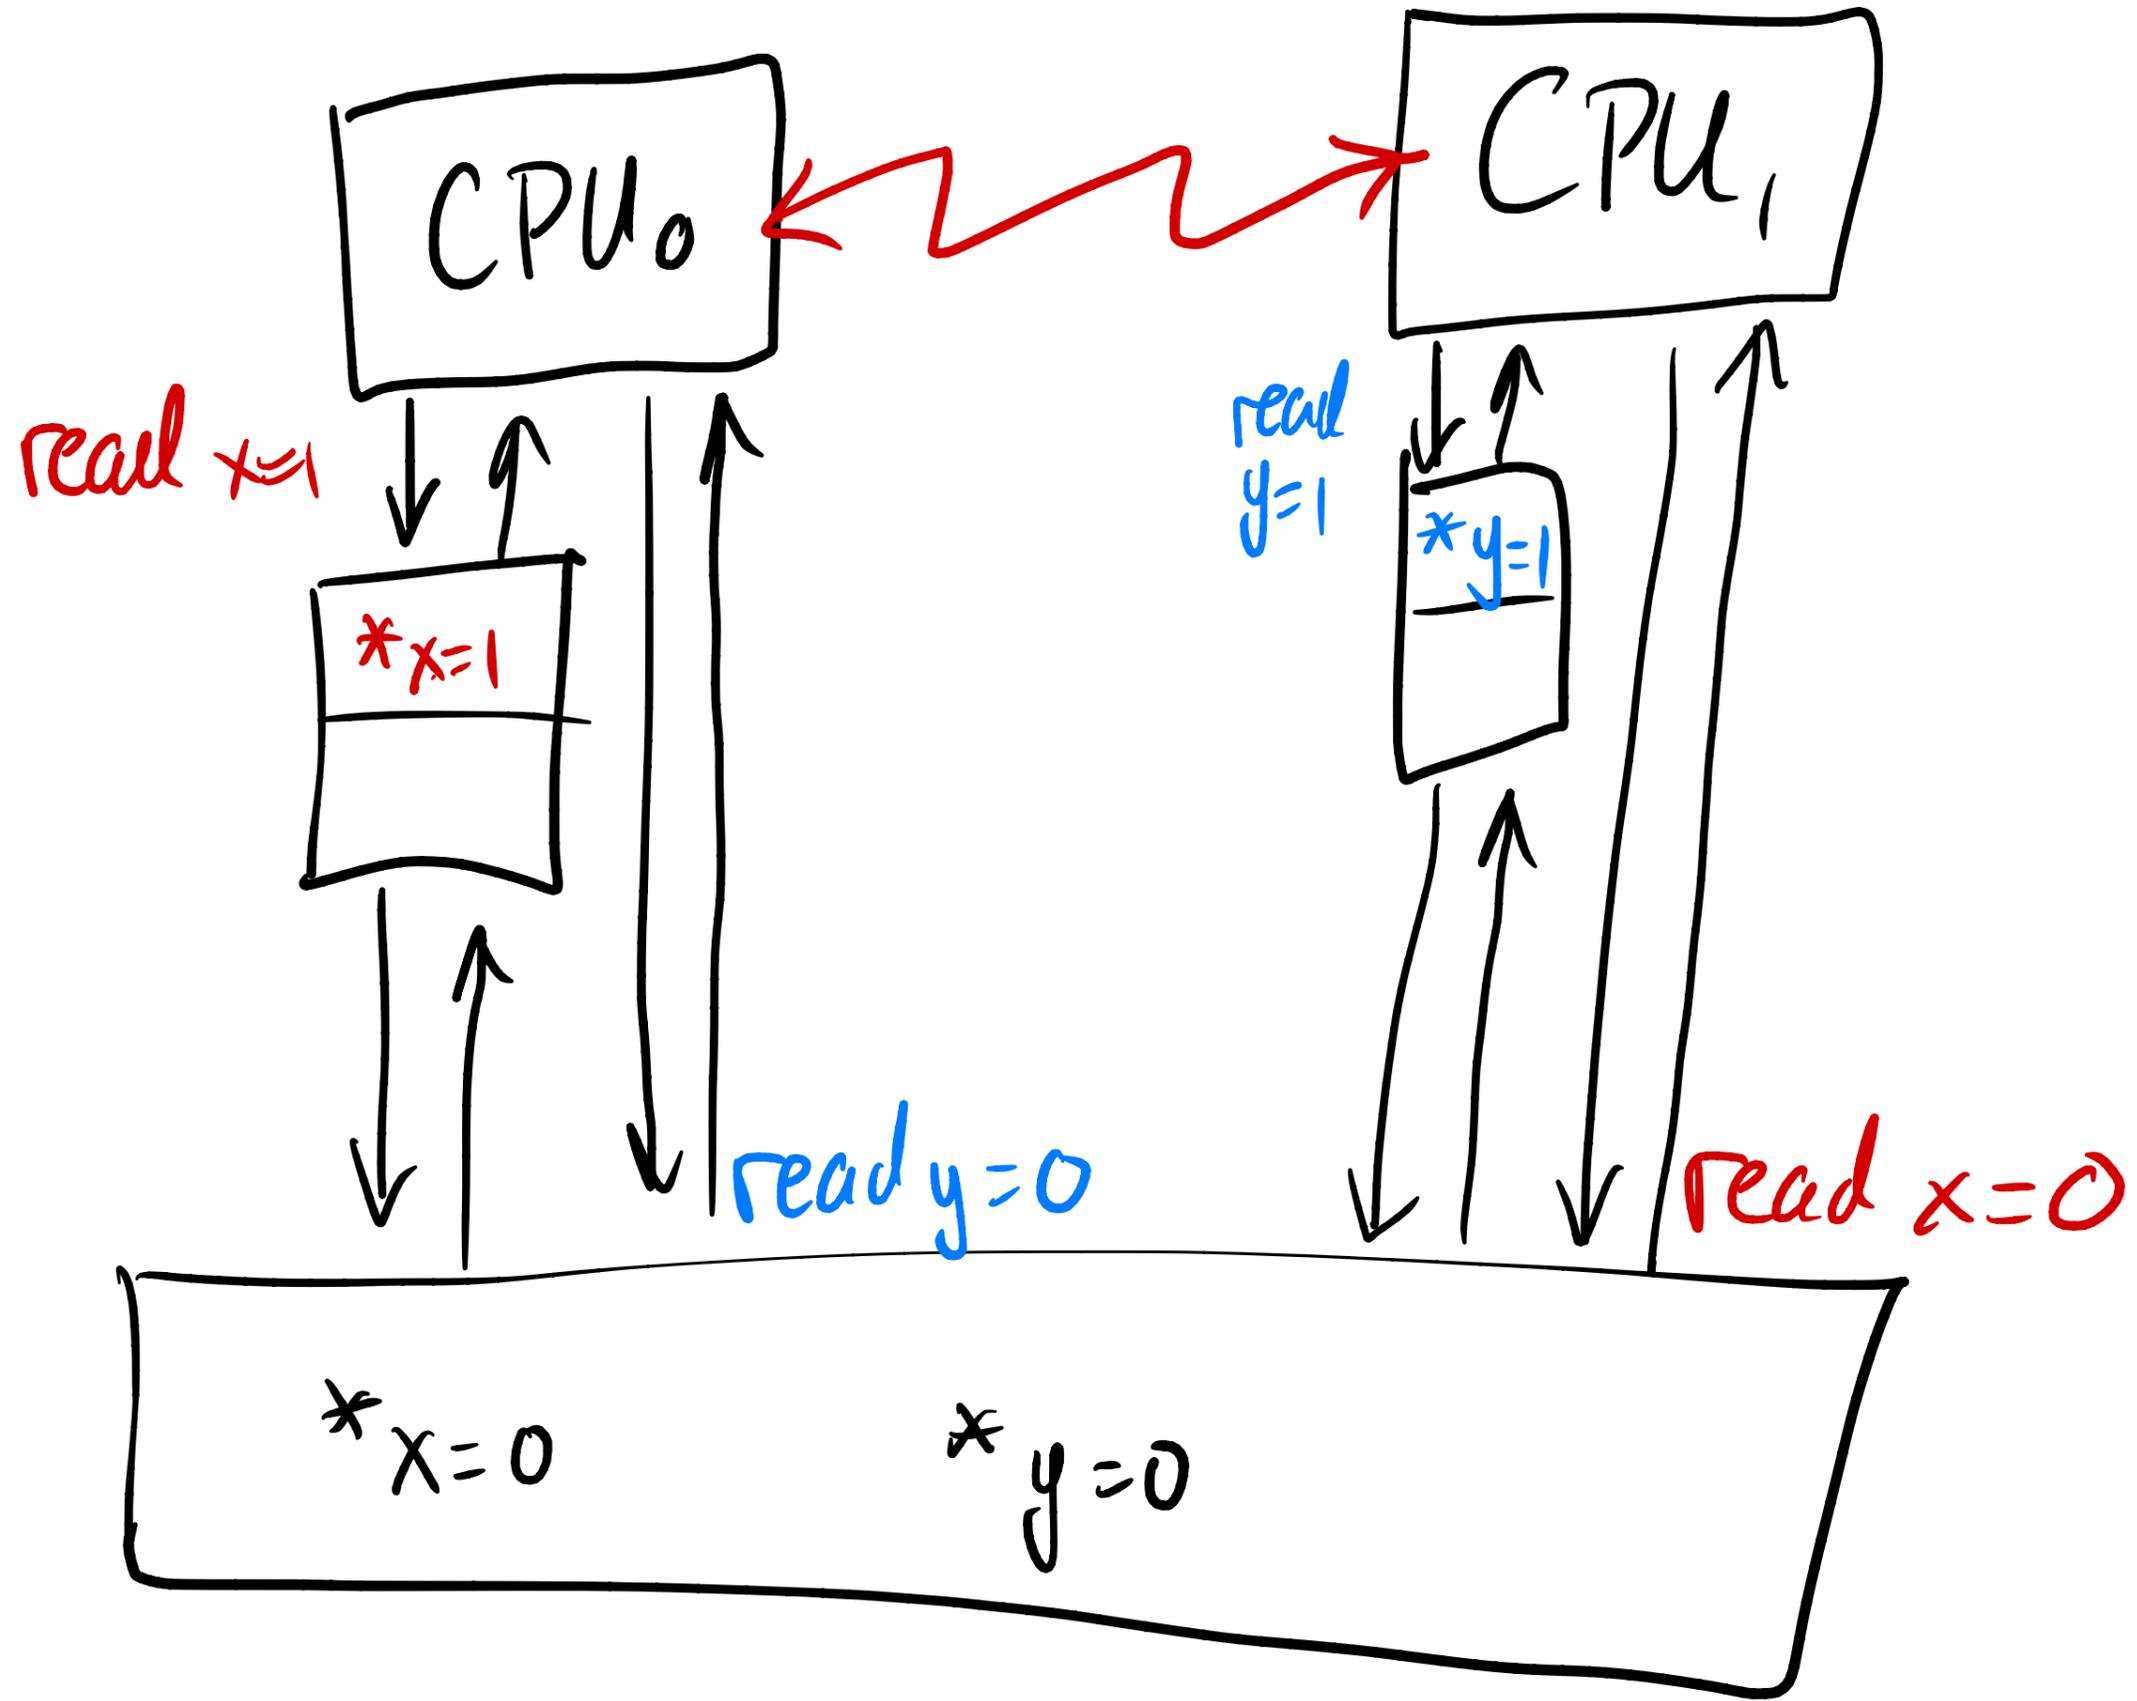












... 2022:

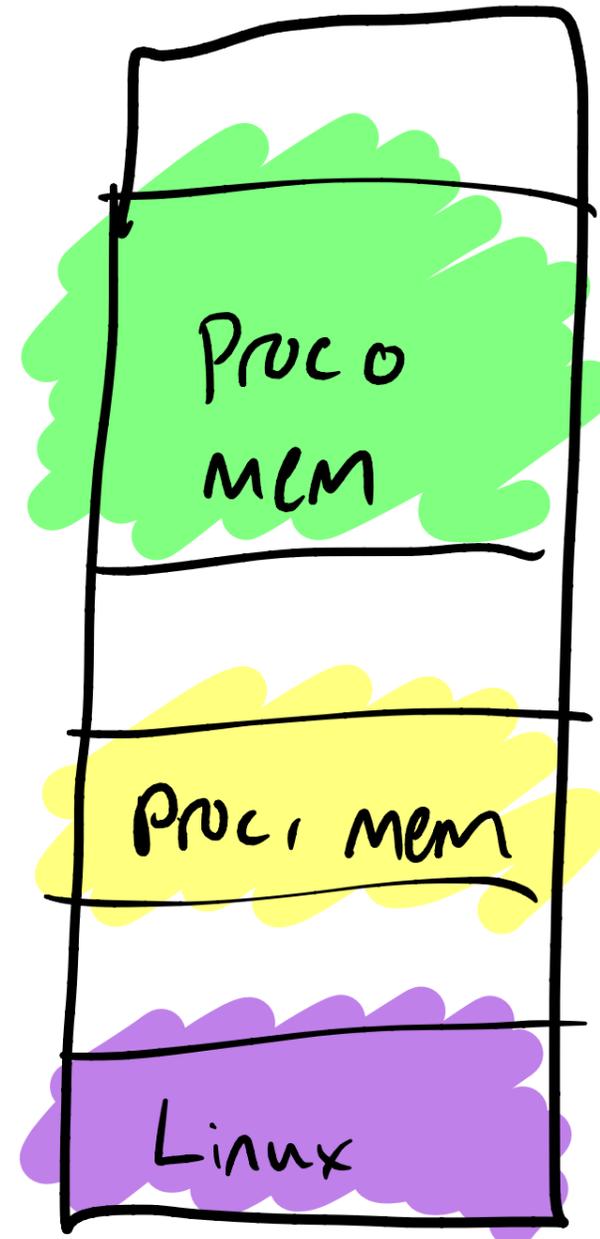
Virtual

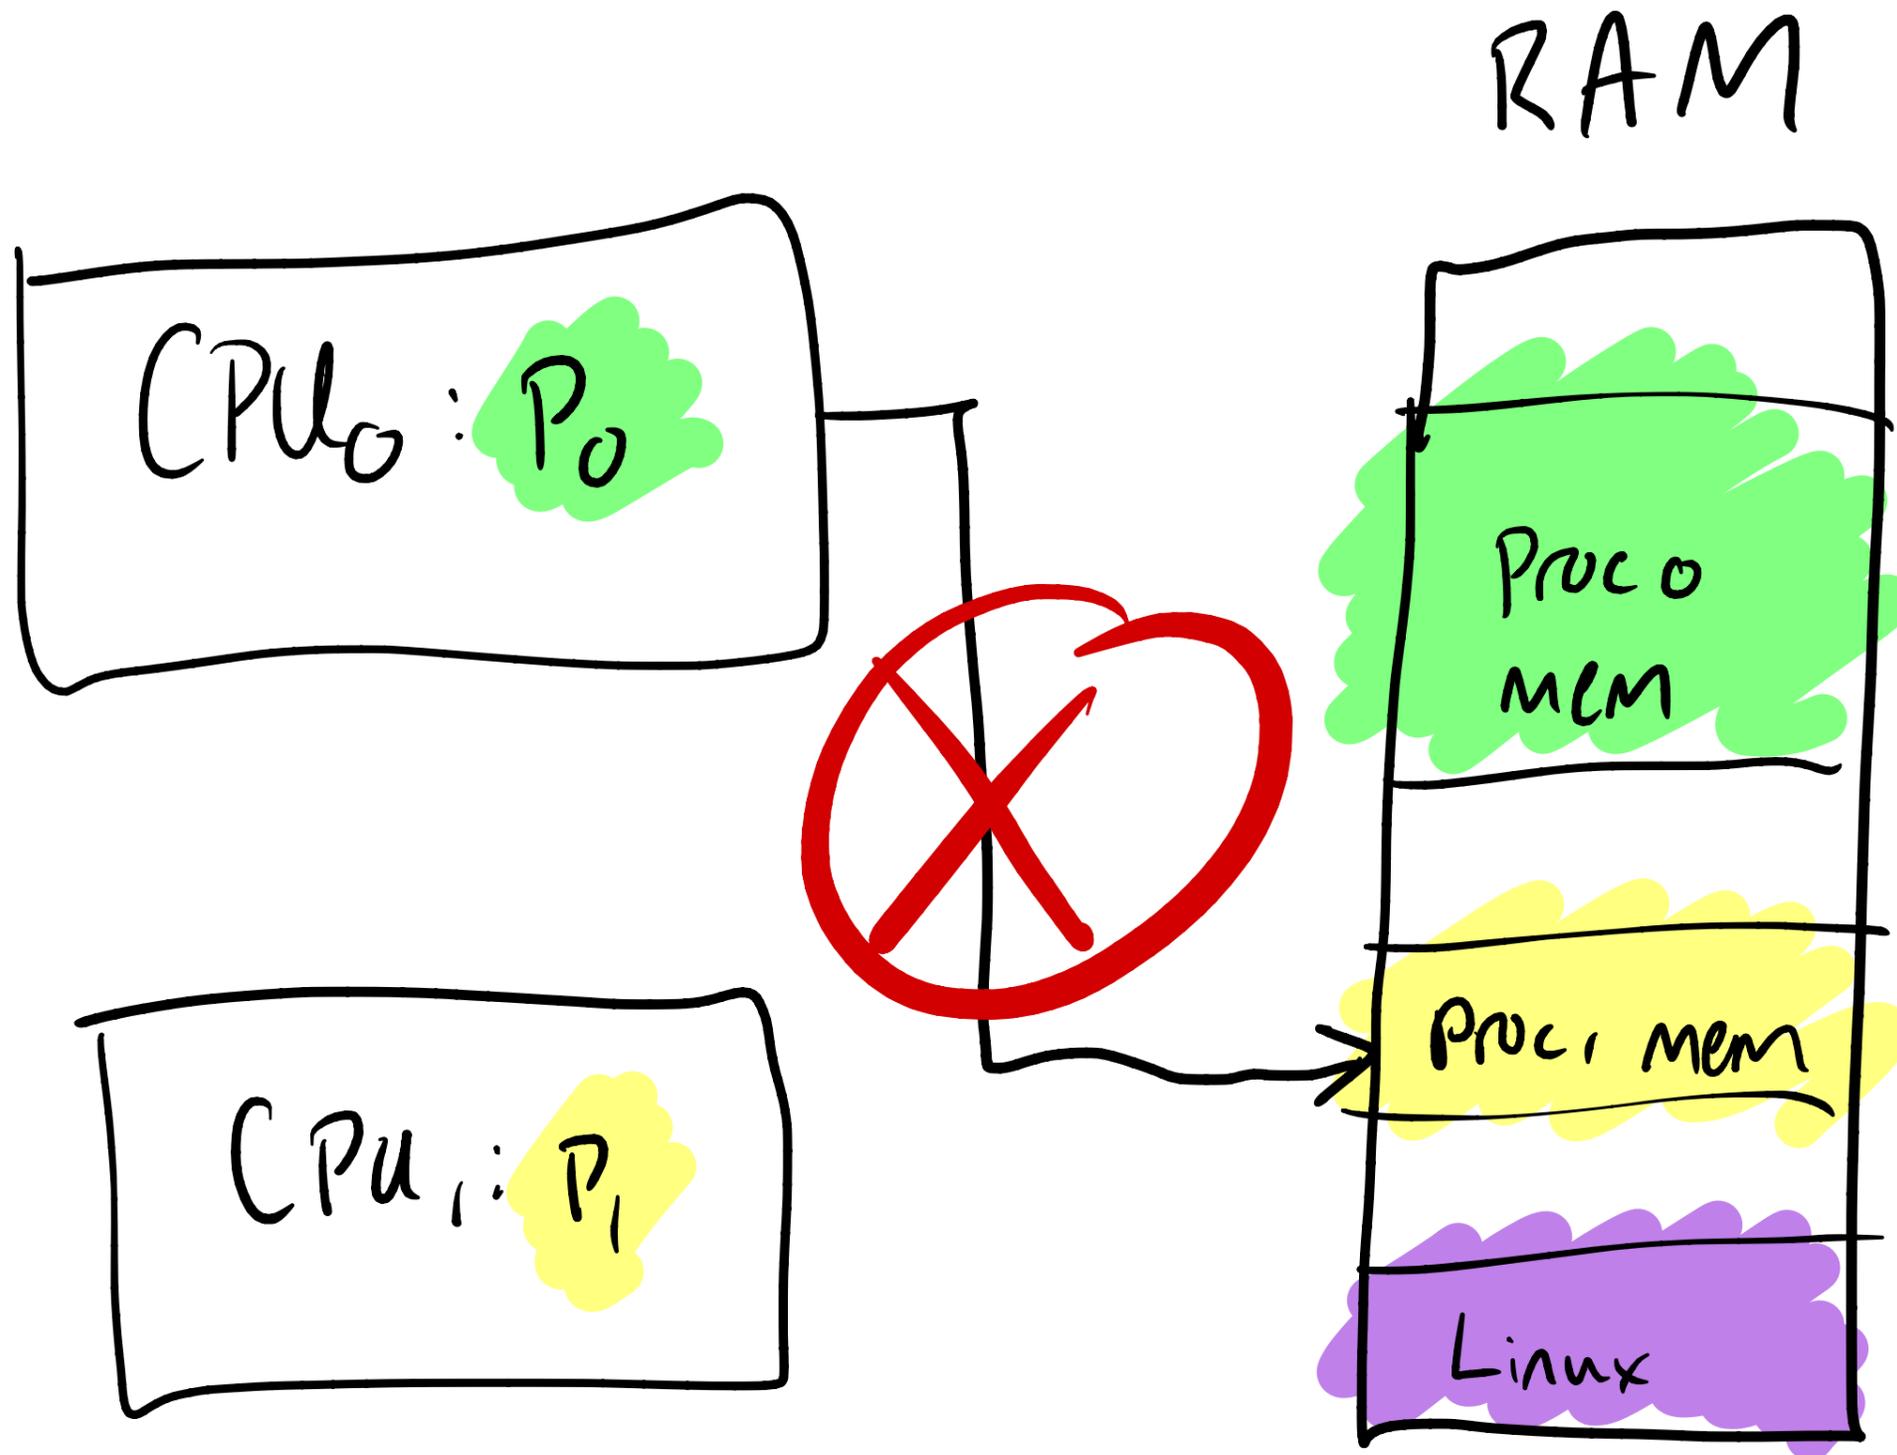
Memory

CPU<sub>0</sub> : P<sub>0</sub>

CPU<sub>1</sub> : P<sub>1</sub>

RAM

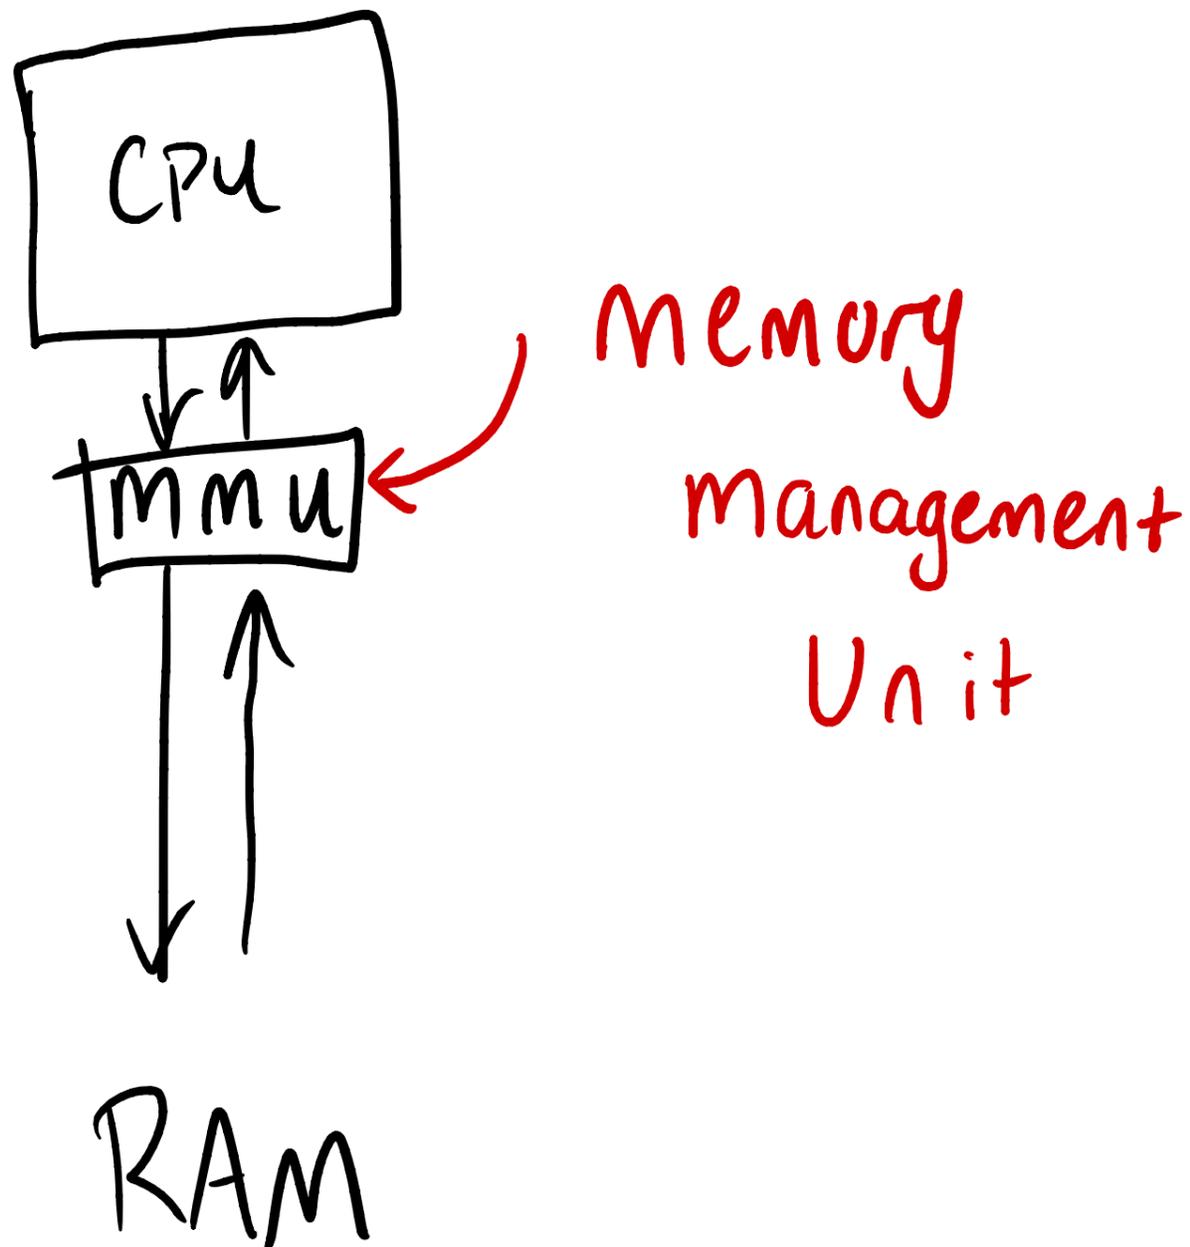




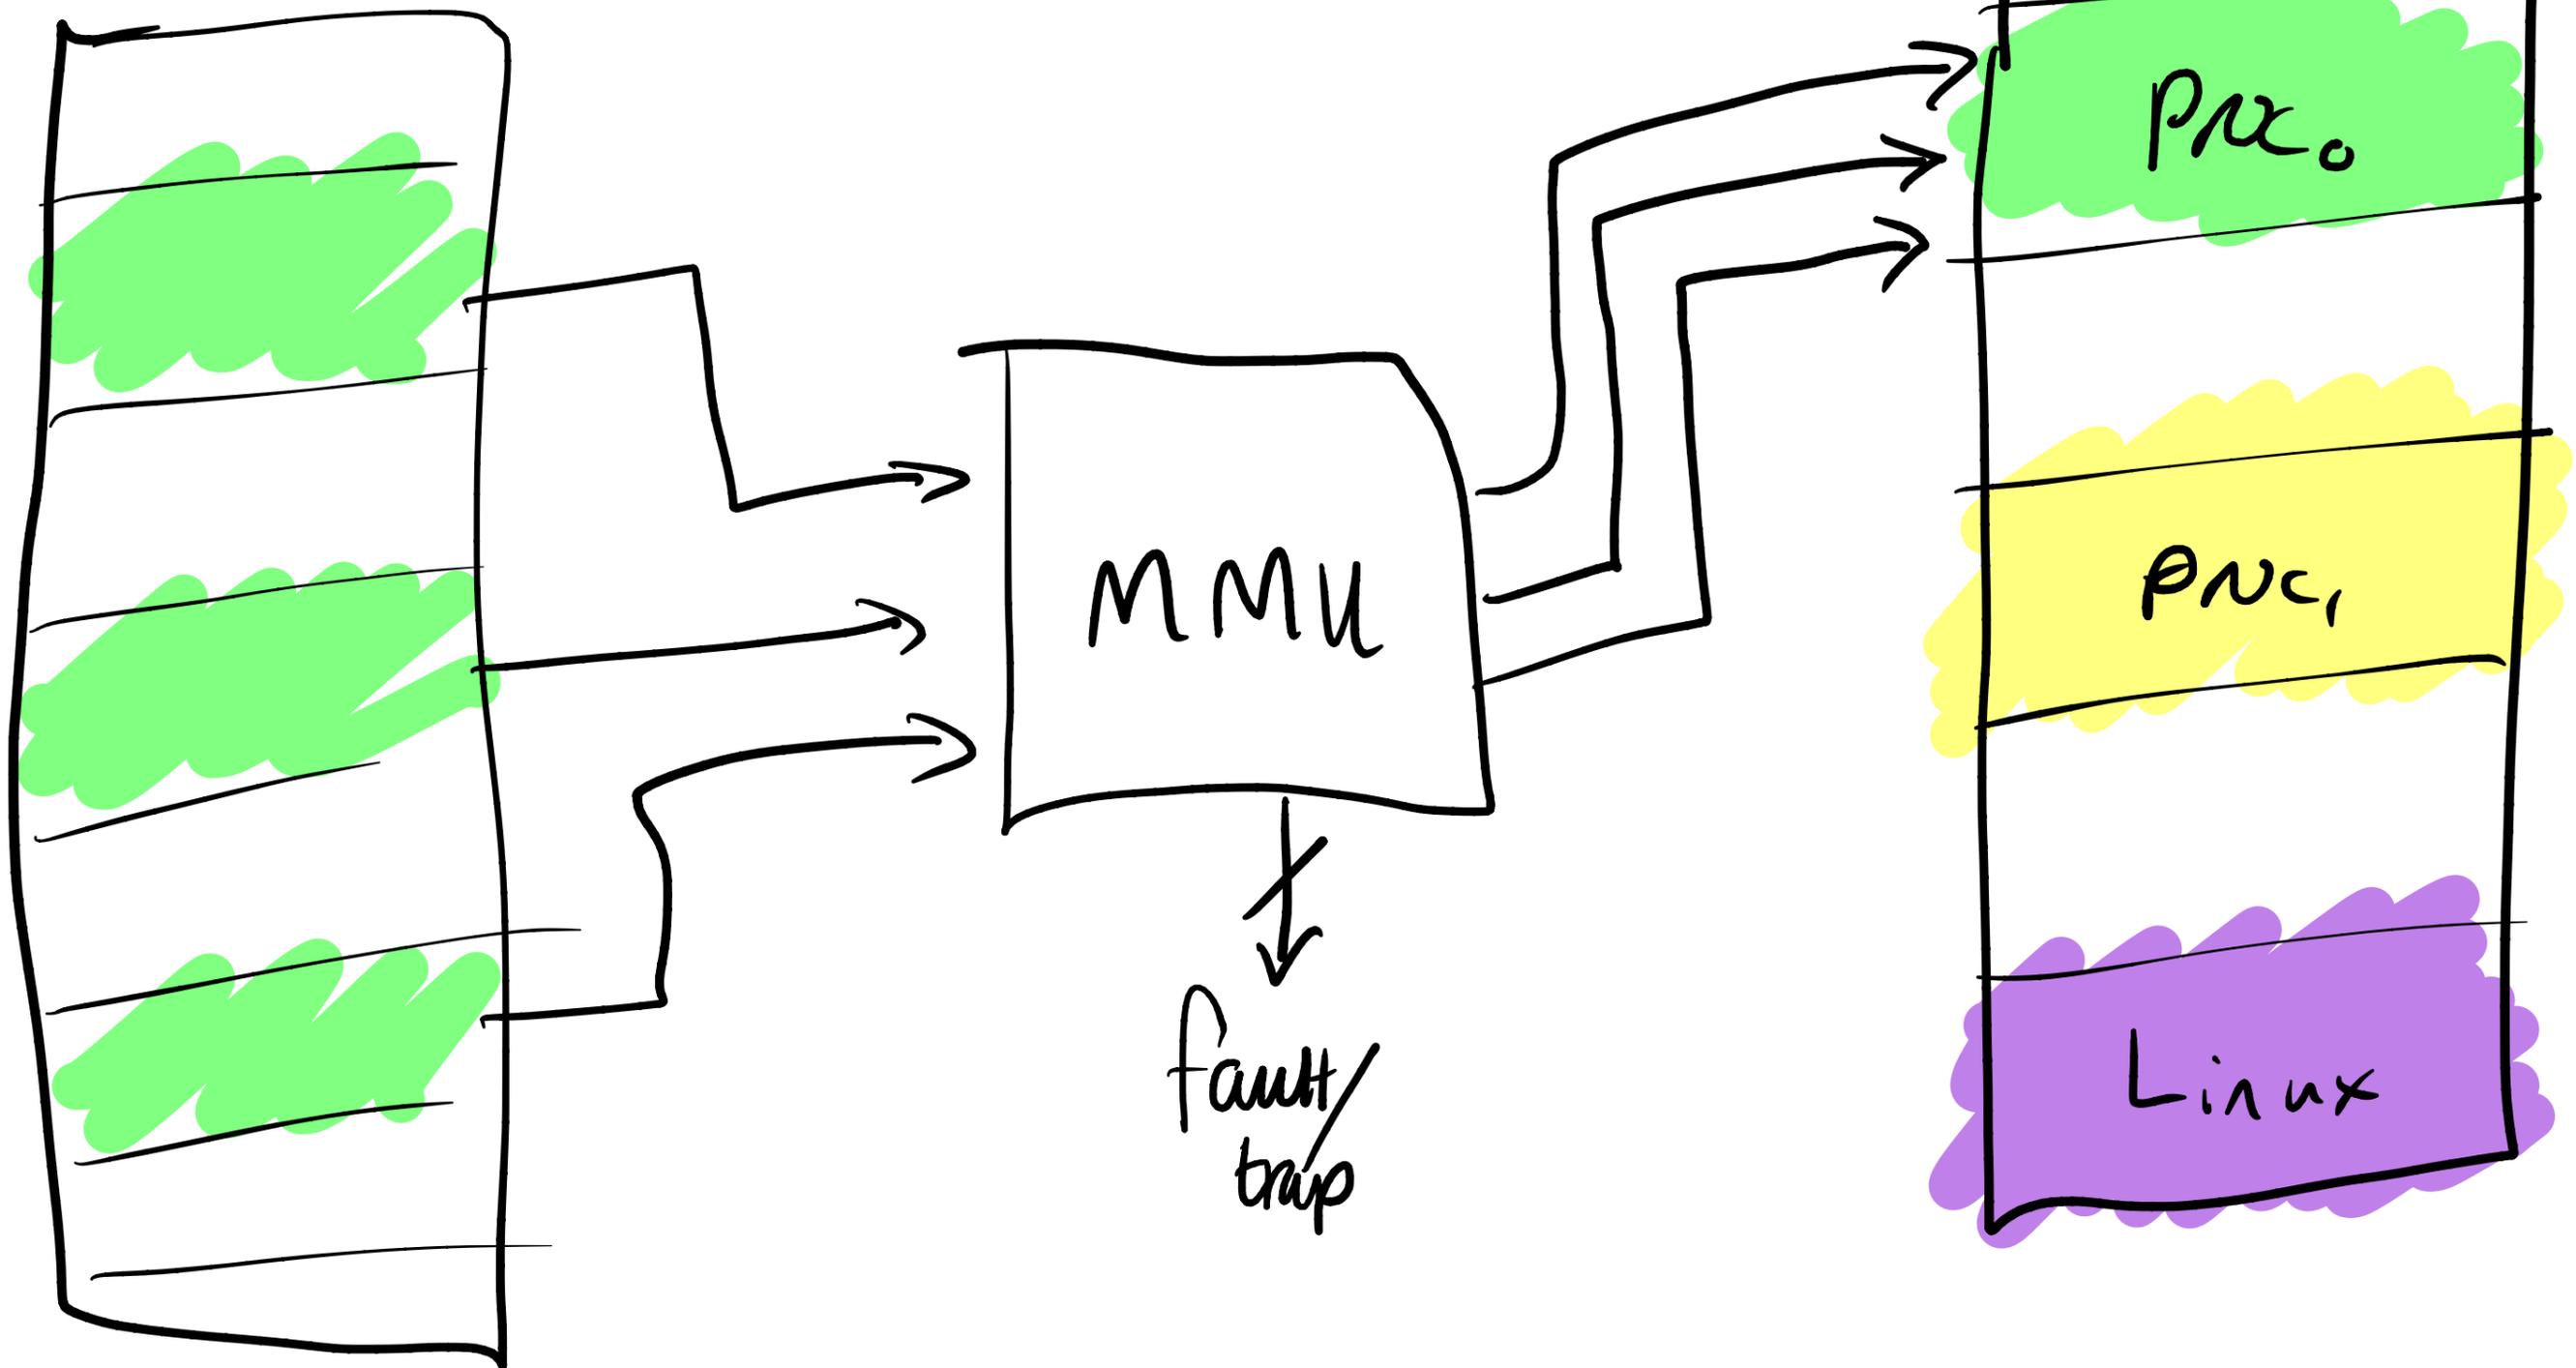
# Aims

1. formal Semantics of Programmer-visible behaviours
2. build Support tooling
3. Use Semantics to reason about code

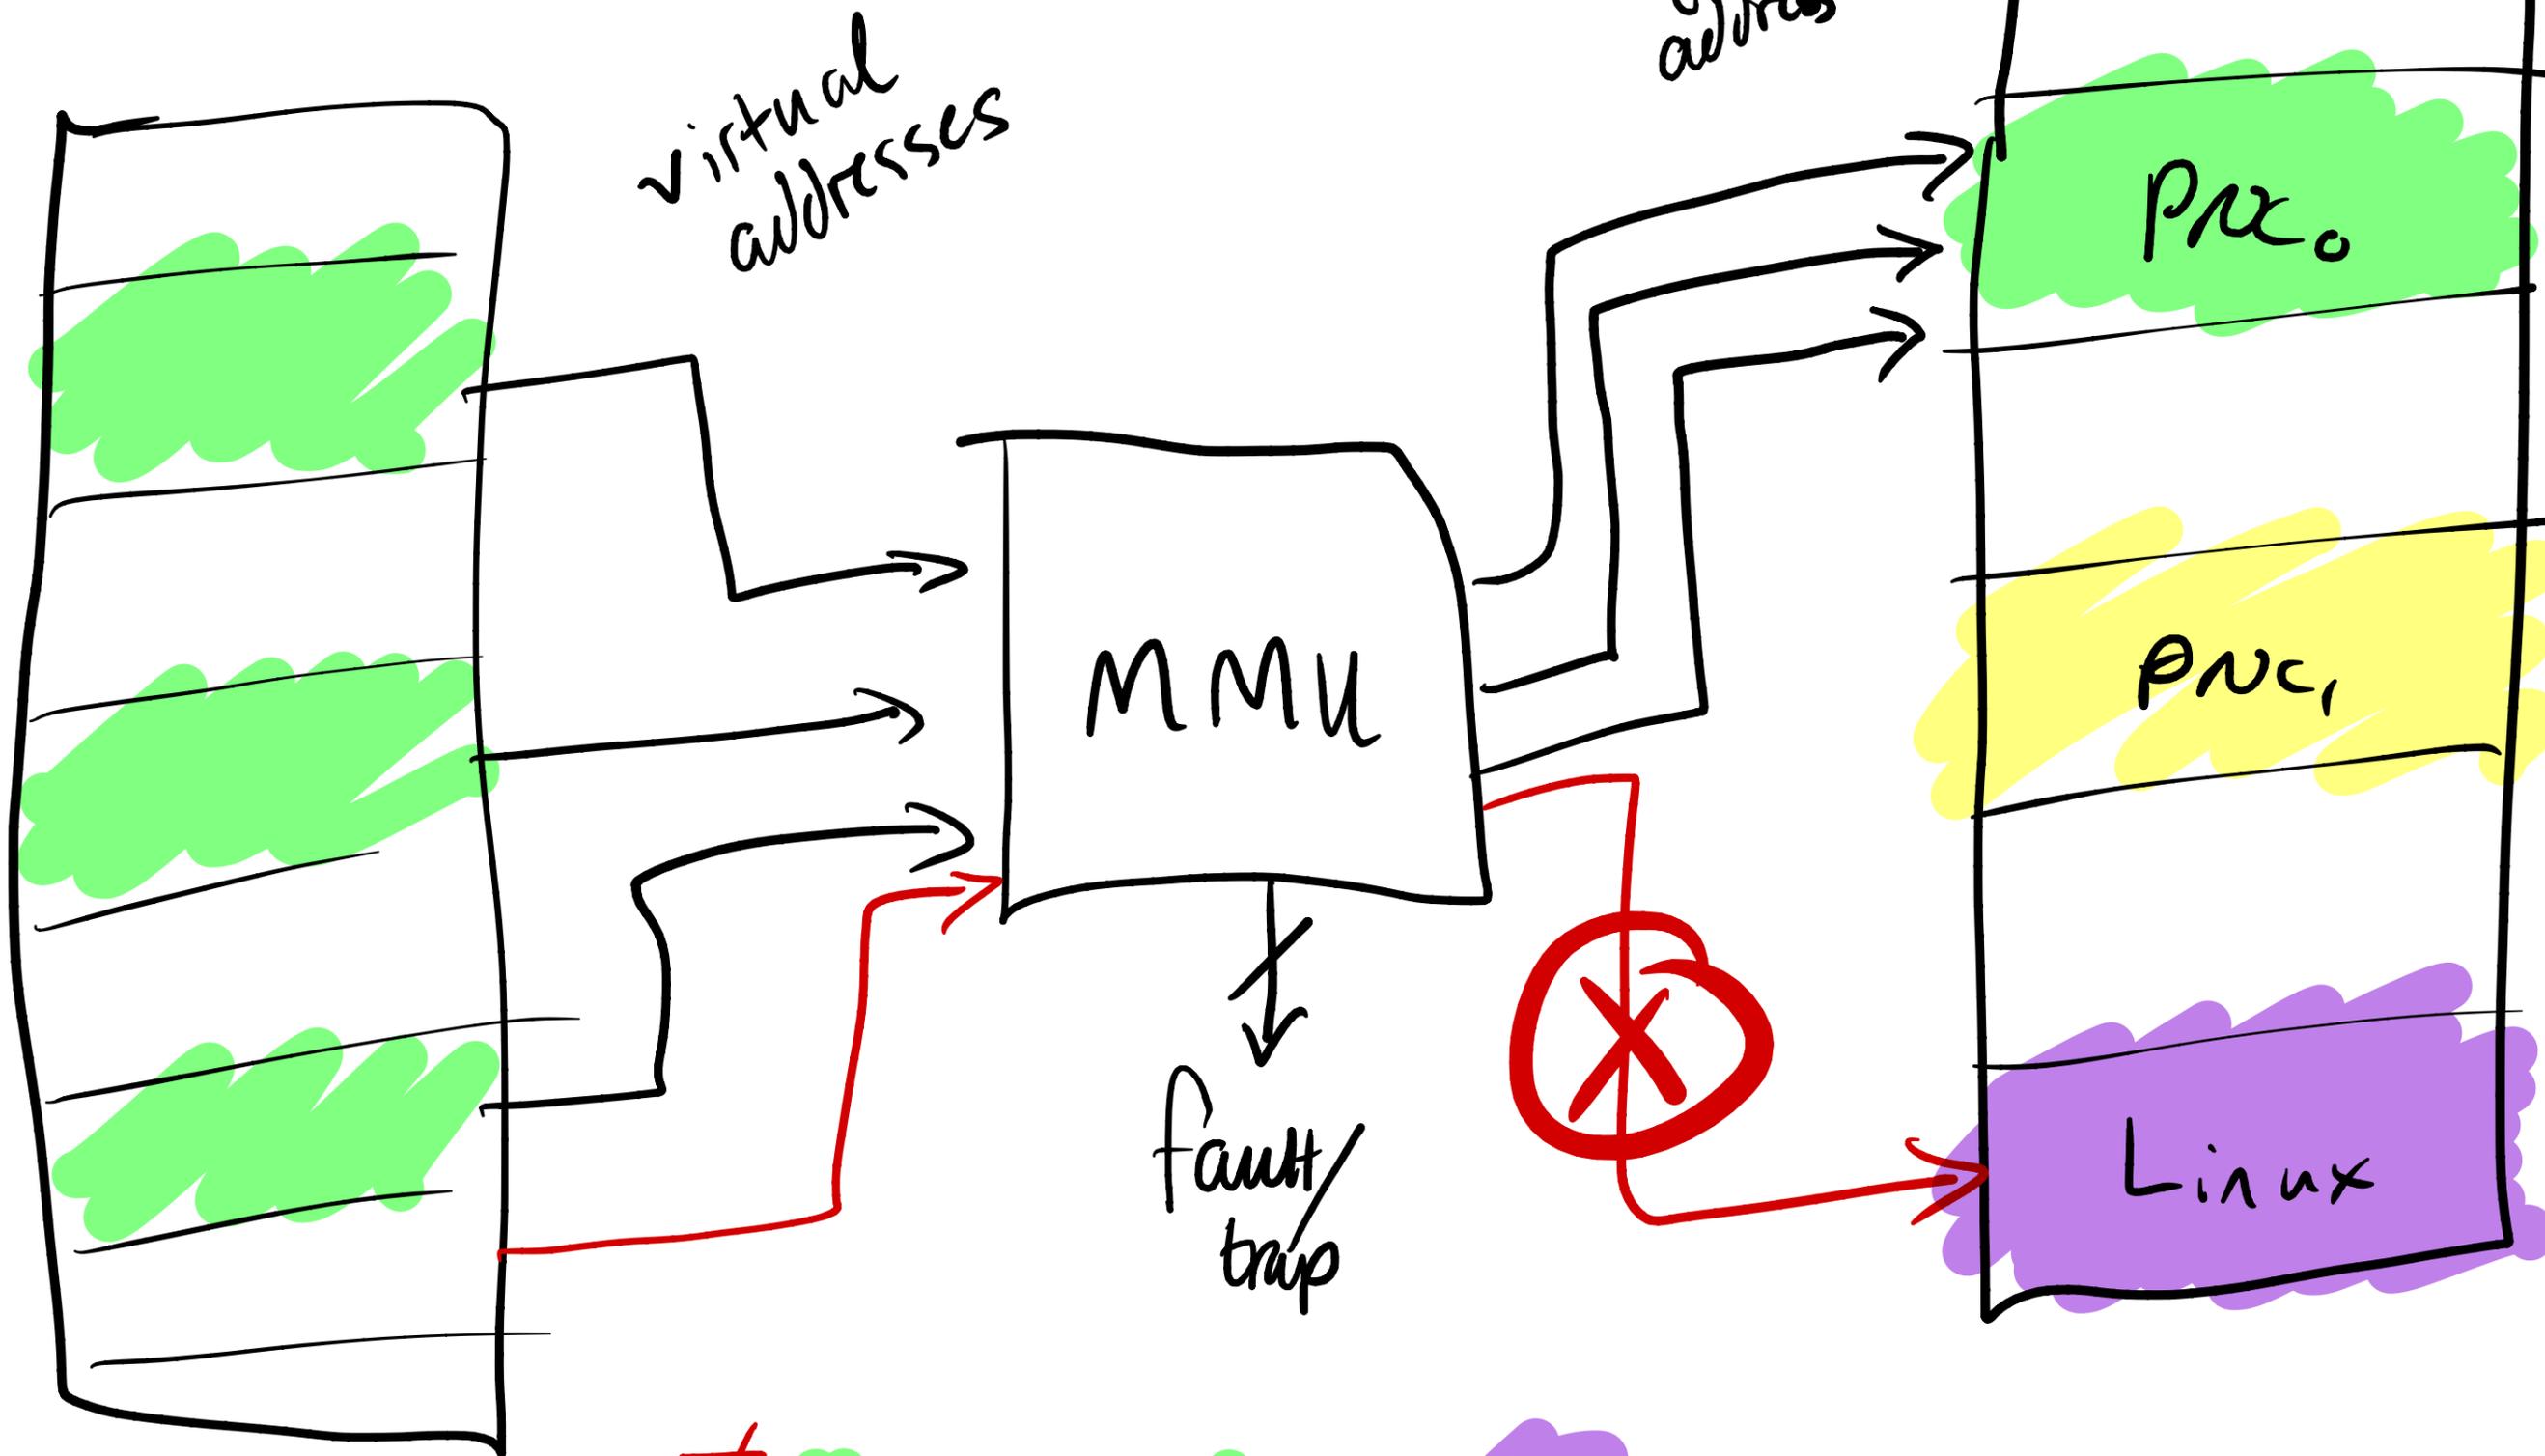
# Memory Management



Proco view

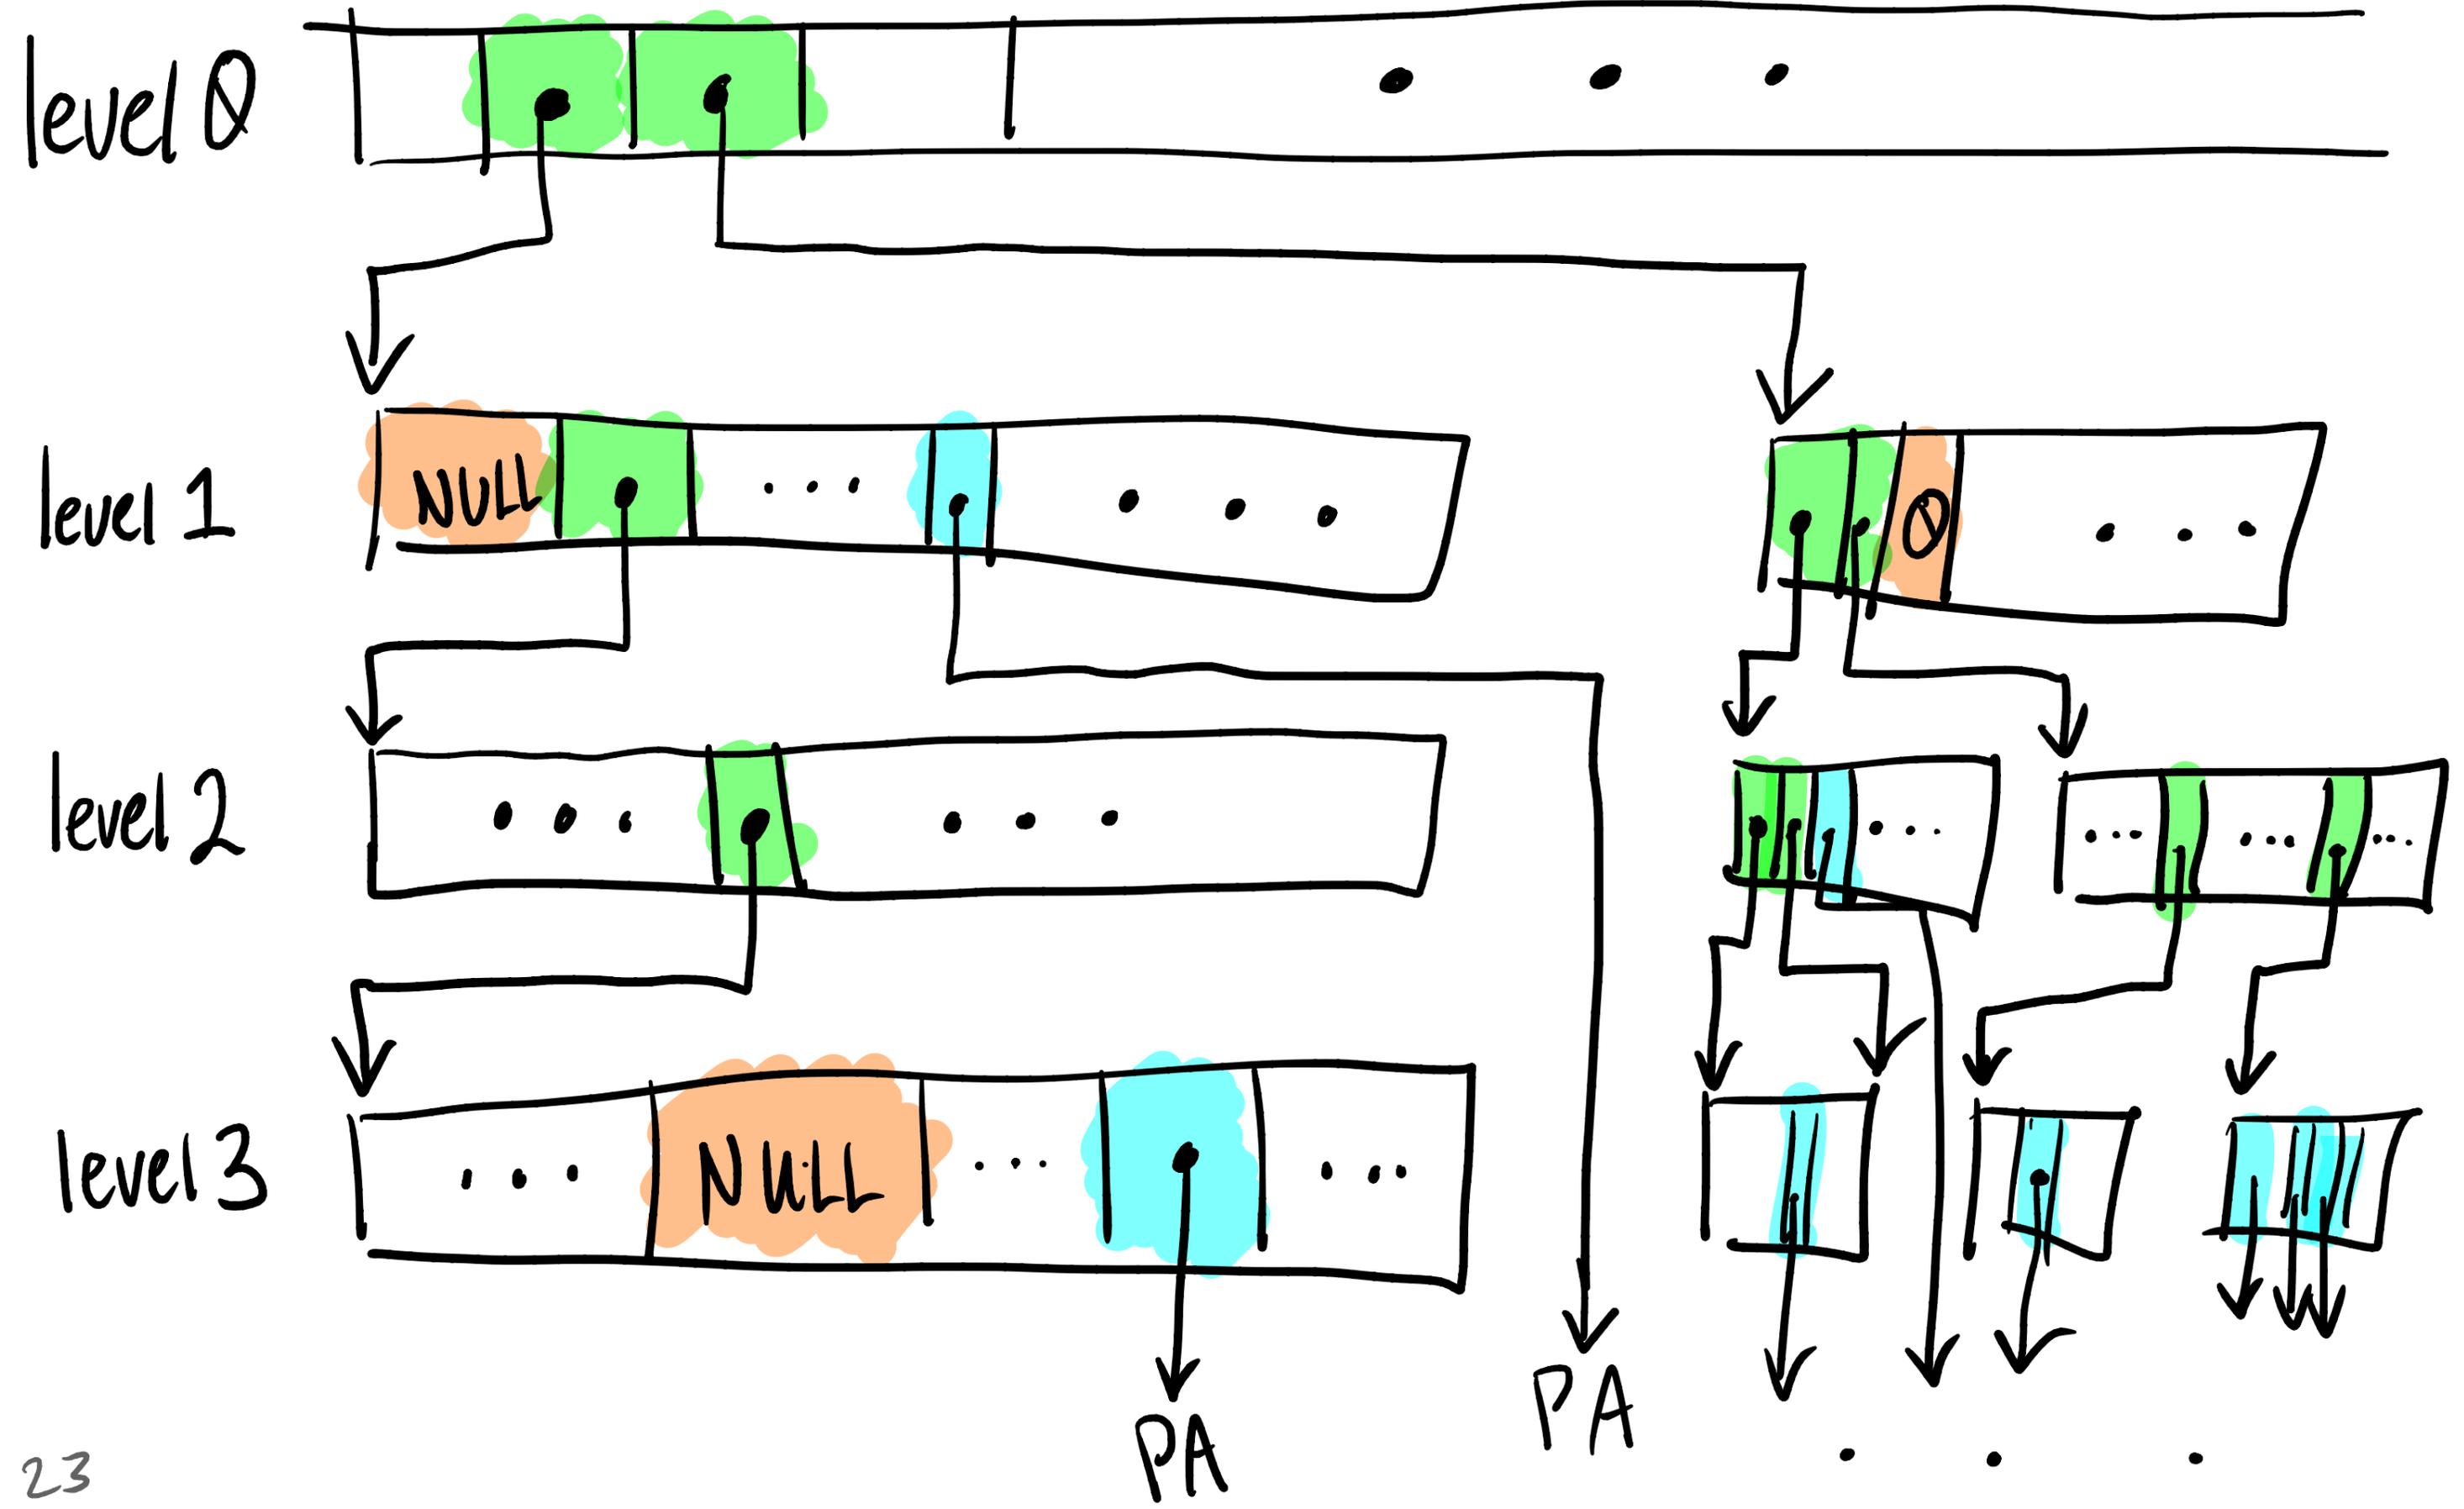


Proco view



$$\nexists va. mmu(va) \in Linux$$

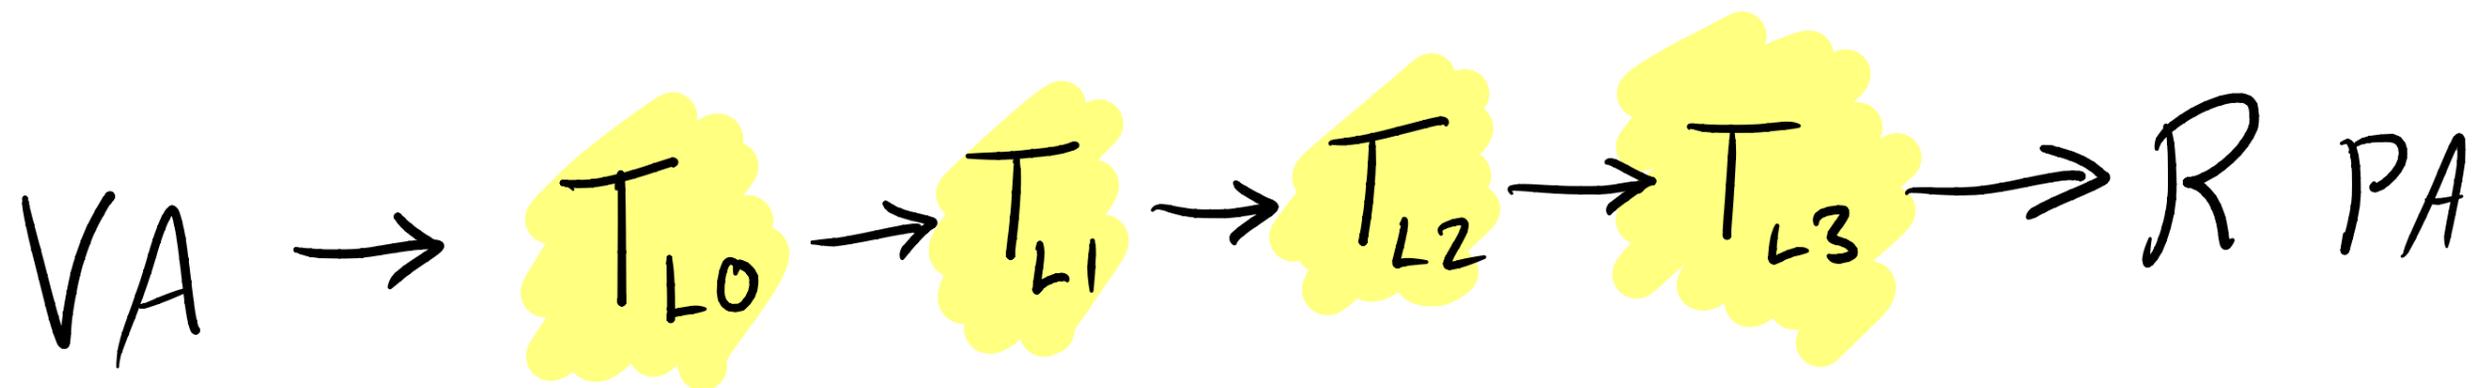
# In-memory Data Structure

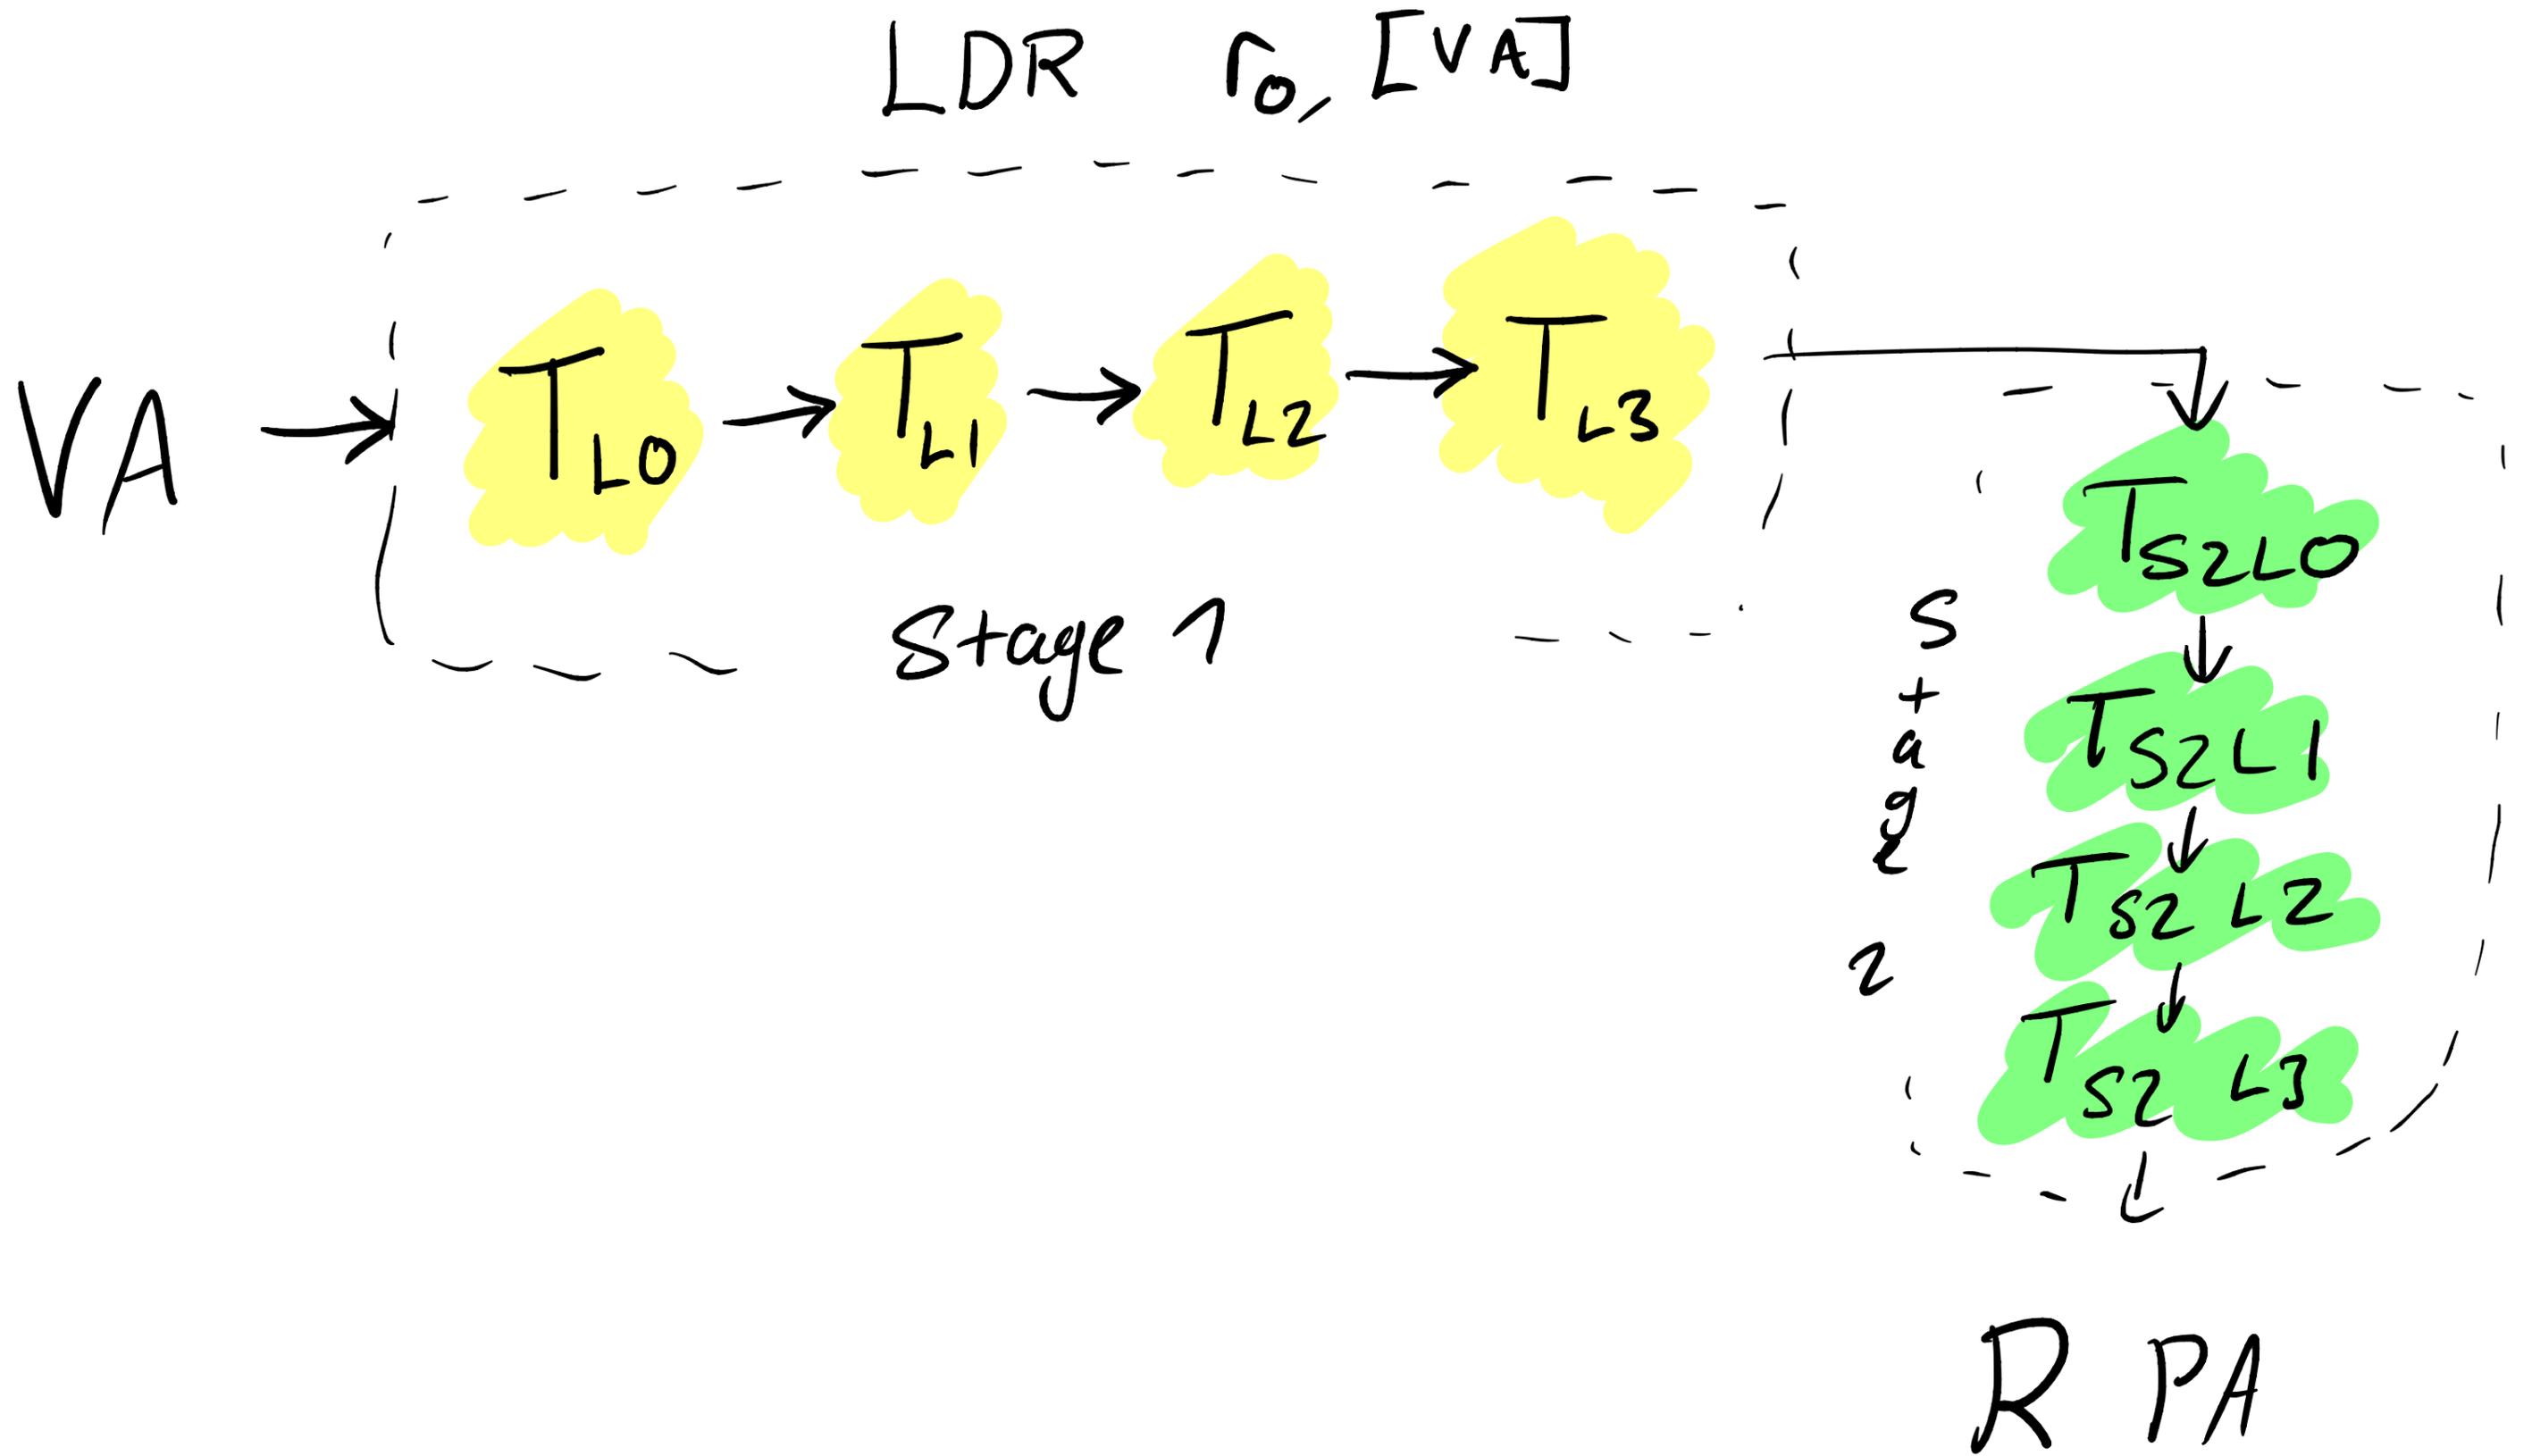


LDR  $\tau_0, [VA]$

VA  $\longrightarrow$  R PA

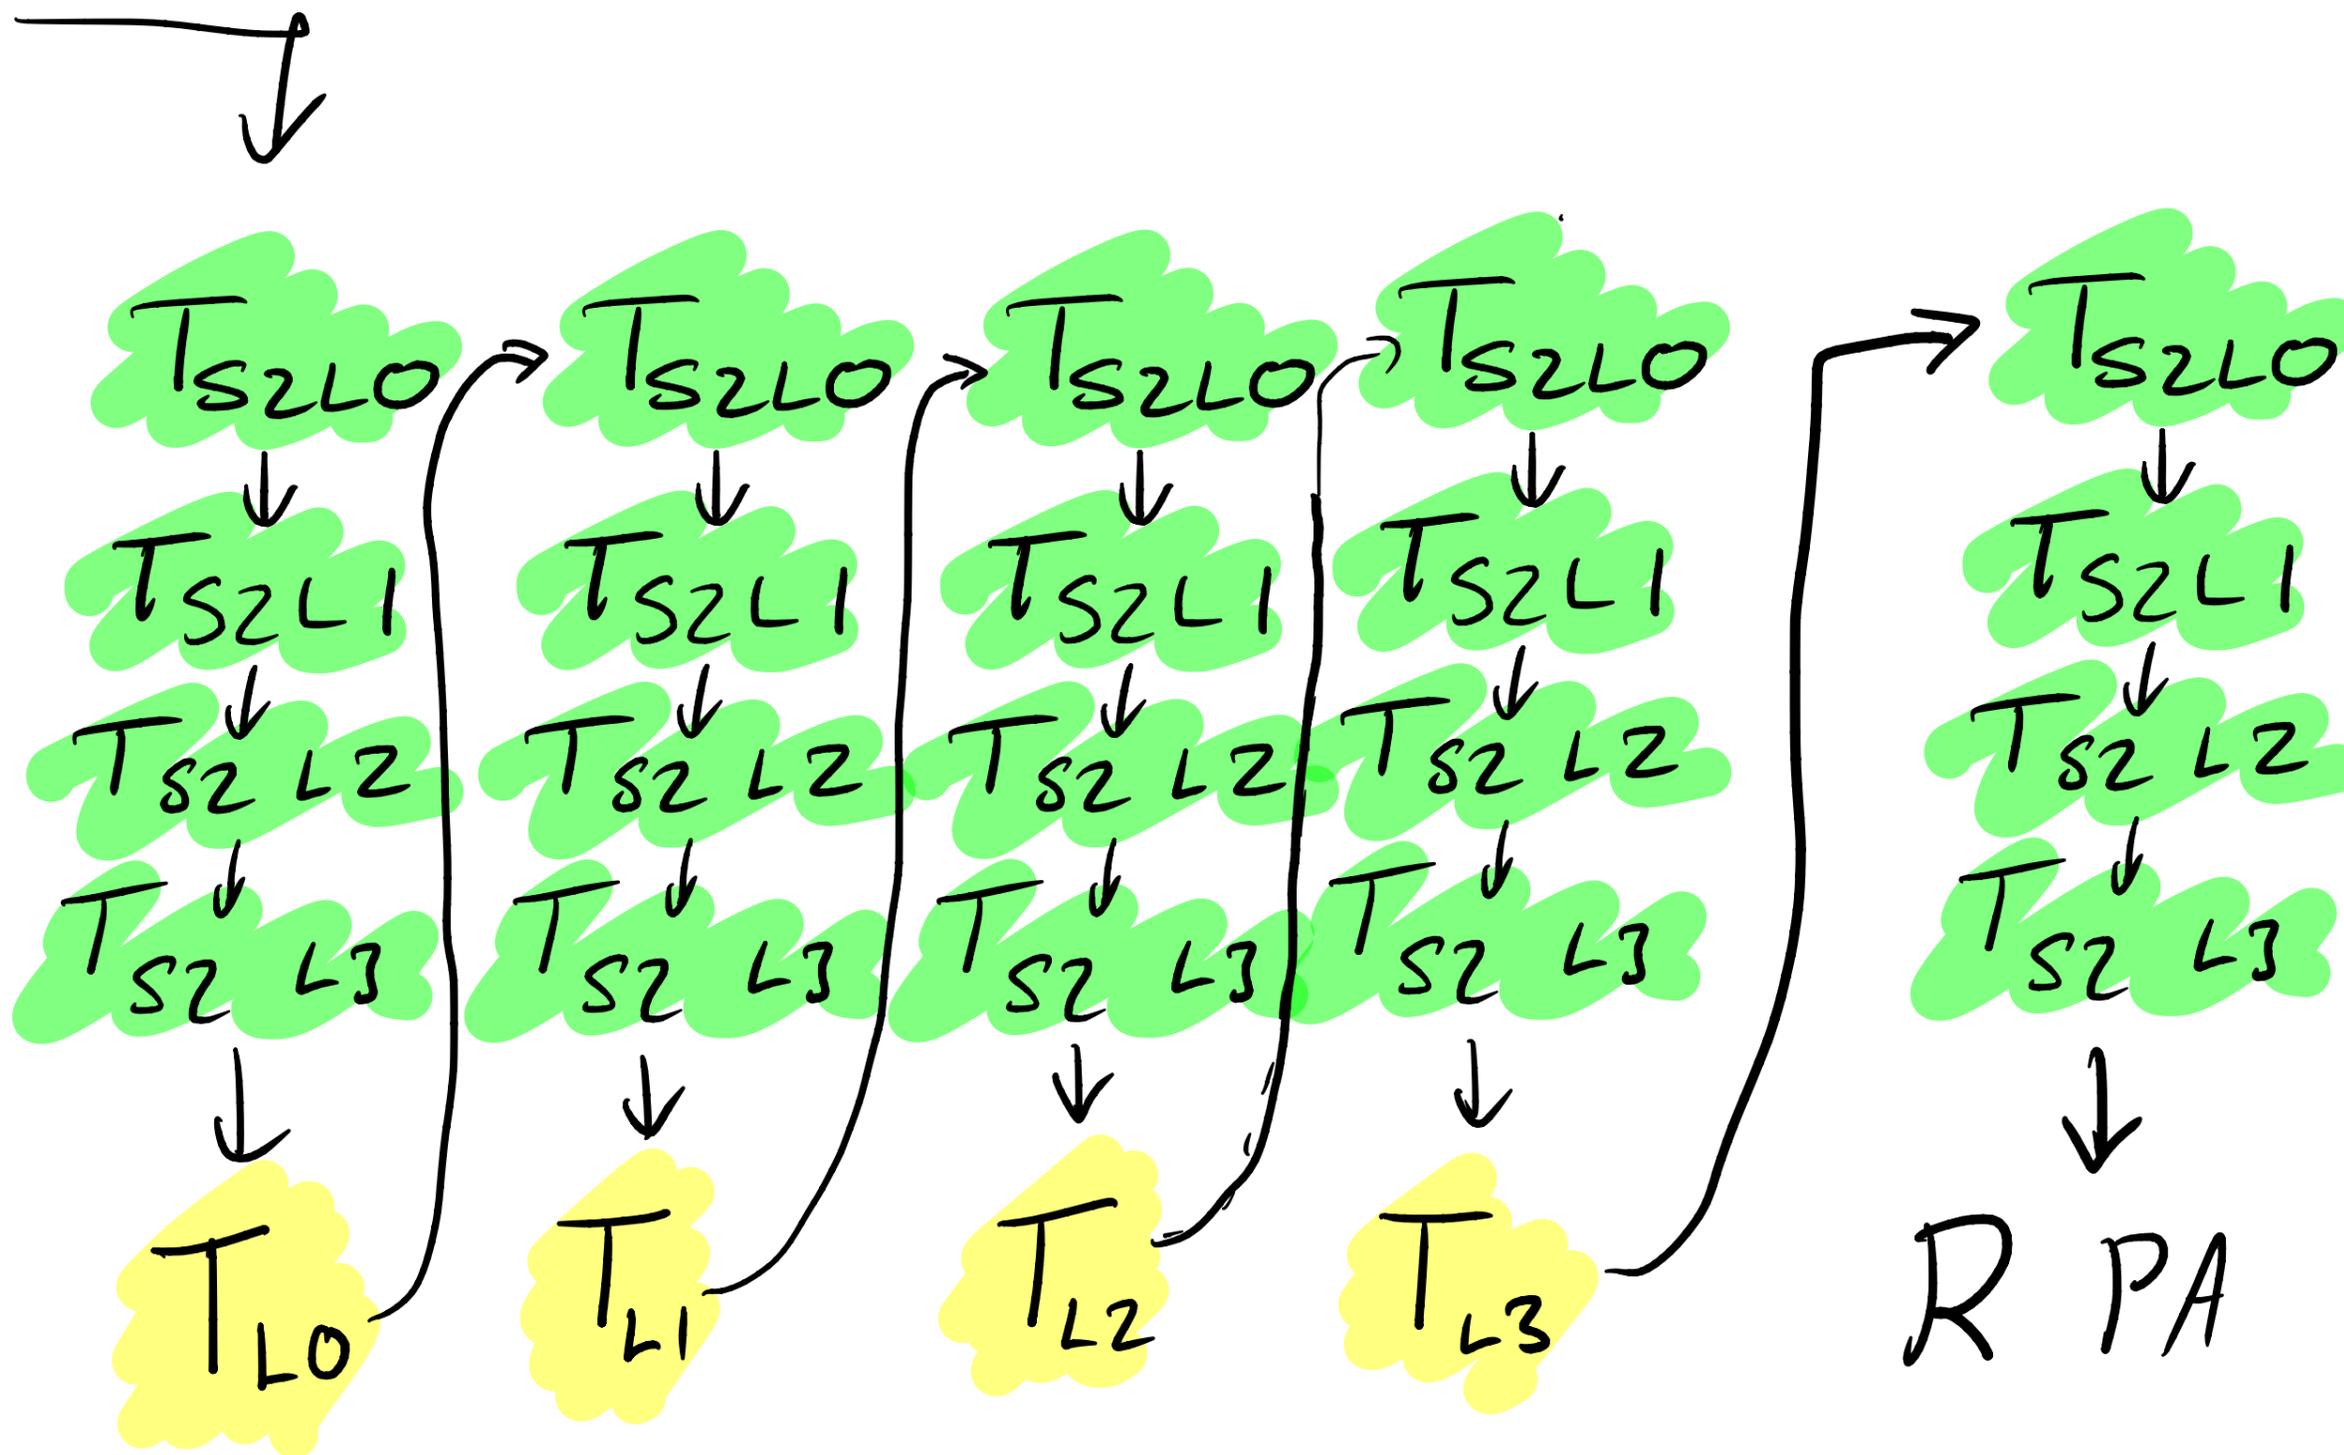
LDR  $\tau_0, [VA]$





LDR  $r_0, [VA]$

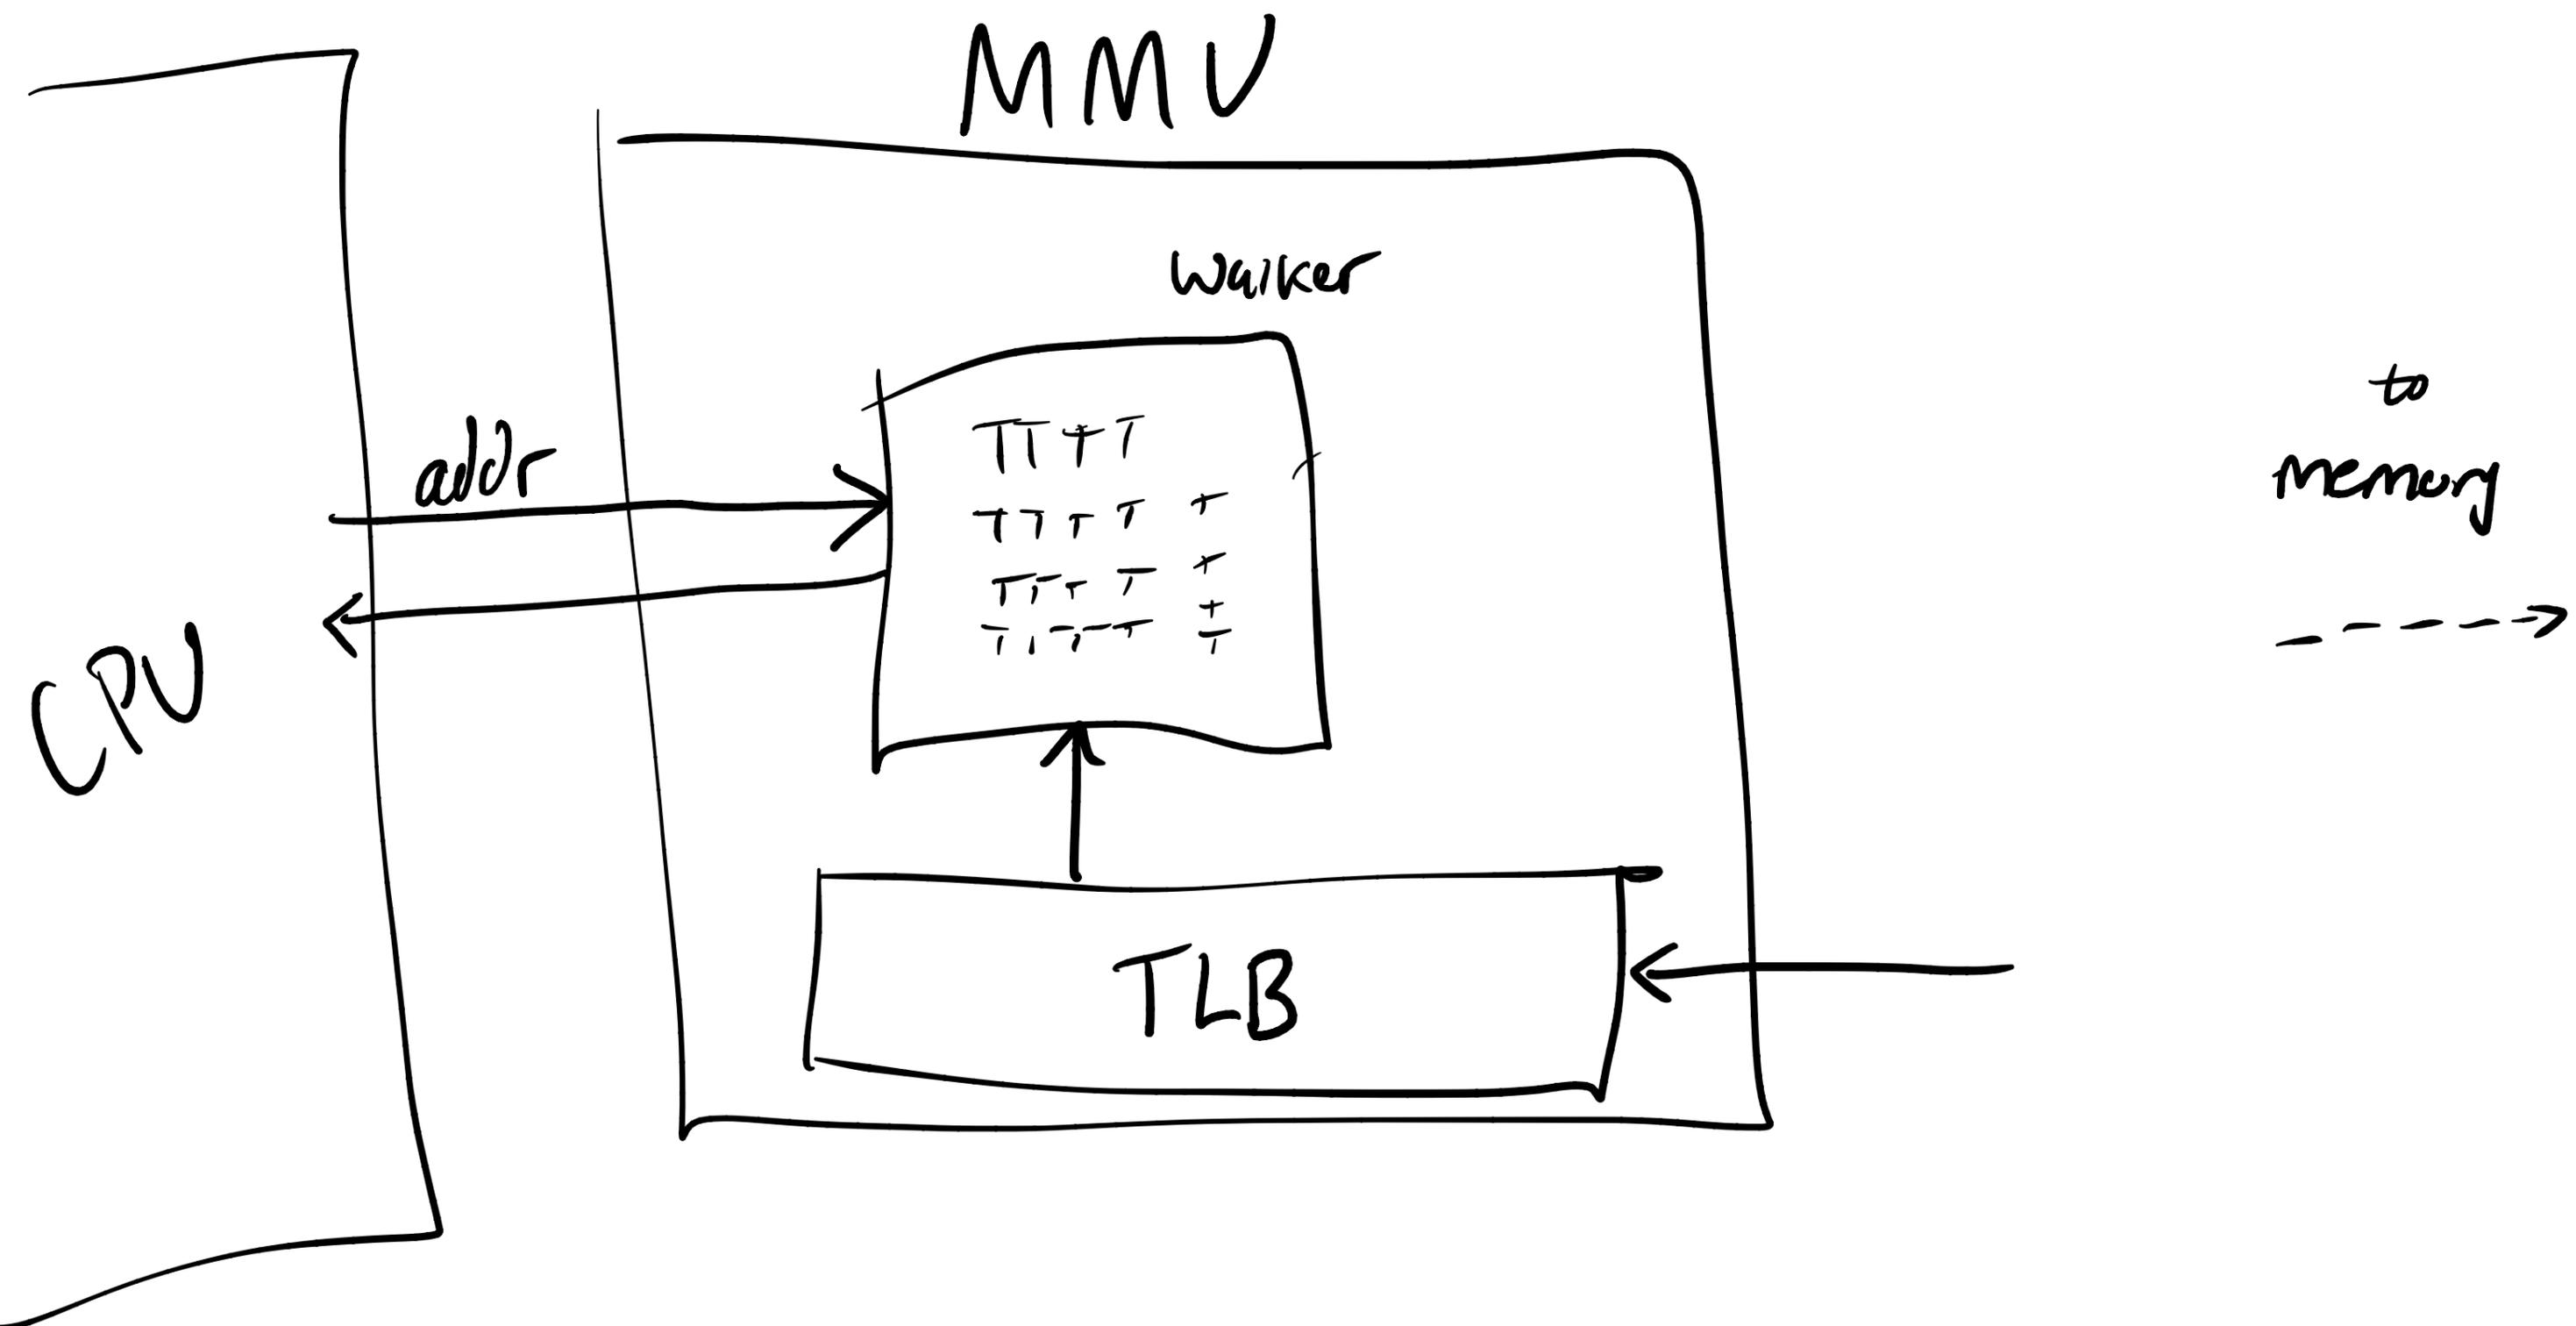
VA



LDR  $r_0, [VA]$

VA





# Modelling: TLBs and walkers

1. Clarity Architecture

2. Extended Arm Axiomatic Model

3. isla - axiomatic tool

# Architectural Design Questions

- Physical/virtual Coherence
- Spontaneous translation table walks
- Coherent TLB fills
- Re-ordered translations
- In-order walks
- Speculative translations
- Micro TLBs
- Multi-copy atomility
- ETS
- Break-before-make
- Write forwarding
- TLB Shutdown
- ASIDs and VMIDs

# Example: instruction re-ordering

init:  $x \mapsto \text{NULL}$

Thread 0

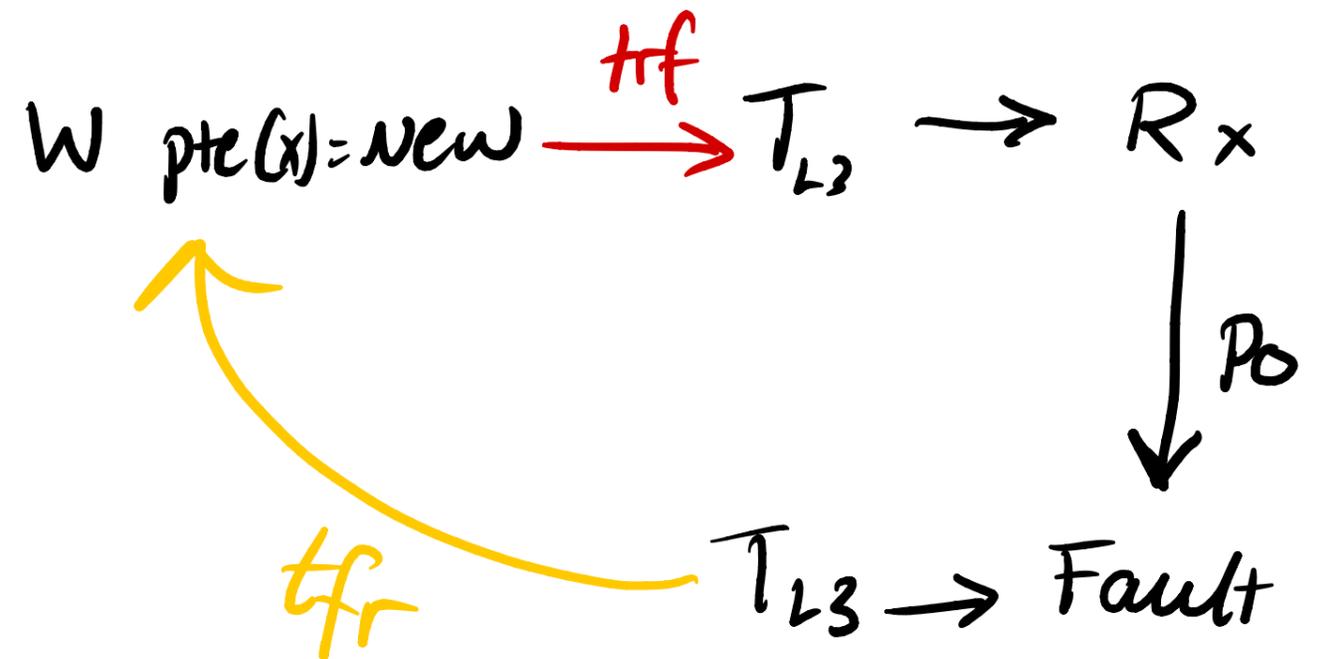
Thread 1

store  $\text{ptr}(x) = \text{new}$

$r_0 \leftarrow \text{load } x$

$r_1 \leftarrow \text{load } x$

Outcome:  $r_0$  sees new  
 $r_1$  sees old  
allowed



# Axiomatic Model

```

let tlb-affects =
  ...

let TLB_barrier =
  ([TLBI] ; tlb-affects ; [T] ; tfr ; [W])^-1 & wco

let maybe_TLB_cached =
  ([T] ; trf^-1 ; wco ; [TLBI-S1]) & tlb-affects^-1

let tcache1 = [T & Stage1] ; tfr ; TLB_barrier
let tcache2 = [T & Stage2] ; tfr ; TLB_barrier

let speculative =
  ctrl
  | addr; po
  | [T] ; instruction-order
  (* translation-ordered-before *)
let tob =
  [T_f] ; tfre
  | ([T_f] ; tfri) & (po ; [DSB.SY] ; instruction-order)^-1
  | [T] ; iio ; [R|W] ; po ; [W]
  | speculative ; trfi
  (* observed by *)
let obs = rfe | fr | wco
  | trfe
  (* ordered-before TLBI and translate *)
let obtlbi_translate =
  tcache1
  | tcache2 & (iio^-1 ; [T & Stage1] ; trf^-1 ; wco^-1)
  | (tcache2 ; wco? ; [TLBI-S1]) & (iio^-1 ; [T & Stage1] ;
  maybe_TLB_cached)

  (* ordered-before TLBI *)
let obtlbi =
  obtlbi_translate
  | [R|W|Fault] ; iio^-1 ; (obtlbi_translate & ext) ; [TLBI]

  (* context-change ordered-before *)
let ctxob =
  speculative ; [MSR]
  | [CSE] ; instruction-order
  | [ContextChange] ; po ; [CSE]
  | speculative ; [CSE]
  | po ; [ERET] ; instruction-order ; [T]

  (* ordered-before a translation fault *)
let obfault =

```

```

data ; [Fault & IsFromW]
| speculative ; [Fault & IsFromW]
| [dmbst] ; po ; [Fault & IsFromW]
| [dmbld] ; po ; [Fault & (IsFromW|IsFromR)]
| [A|Q] ; po ; [Fault & (IsFromW | IsFromR)]
| [R|W] ; po ; [Fault & IsFromW & IsReleaseW]

(* ETS-ordered-before *)
let obETS =
  (obfault ; [Fault]) ; iio^-1 ; [T_f]
  | ([TLBI] ; po ; [dsb] ; instruction-order ; [T]) & tlb-affects

  (* dependency-ordered-before *)
let dob =
  addr | data
  | speculative ; [W]
  | addr; po; [W]
  | (addr | data); rfi
  | (addr | data); trfi

  (* atomic-ordered-before *)
let aob = rmw
  | [range(rmw)]; rfi; [A | Q]

  (* barrier-ordered-before *)
let bob = [R] ; po ; [dmbld]
  | [W] ; po ; [dmbst]
  | [dmbst]; po; [W]
  | [dmbld]; po; [R|W]
  | [L]; po; [A]
  | [A | Q]; po; [R | W]
  | [R | W]; po; [L]
  | [F | C]; po; [dsbsy]
  | [dsb] ; po

  (* Ordered-before *)
let ob = (obs | dob | aob | bob | iio | tob | obtlbi | ctxob |
  obfault | obETS)^+

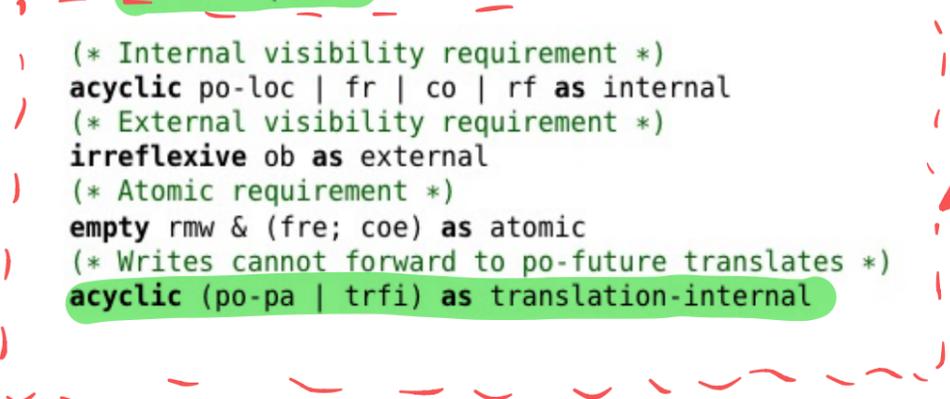
  (* Internal visibility requirement *)
acyclic po-loc | fr | co | rf as internal
  (* External visibility requirement *)
irreflexive ob as external
  (* Atomic requirement *)
empty rmw & (fre; coe) as atomic
  (* Writes cannot forward to po-future translates *)
acyclic (po-pa | trfi) as translation-internal

```

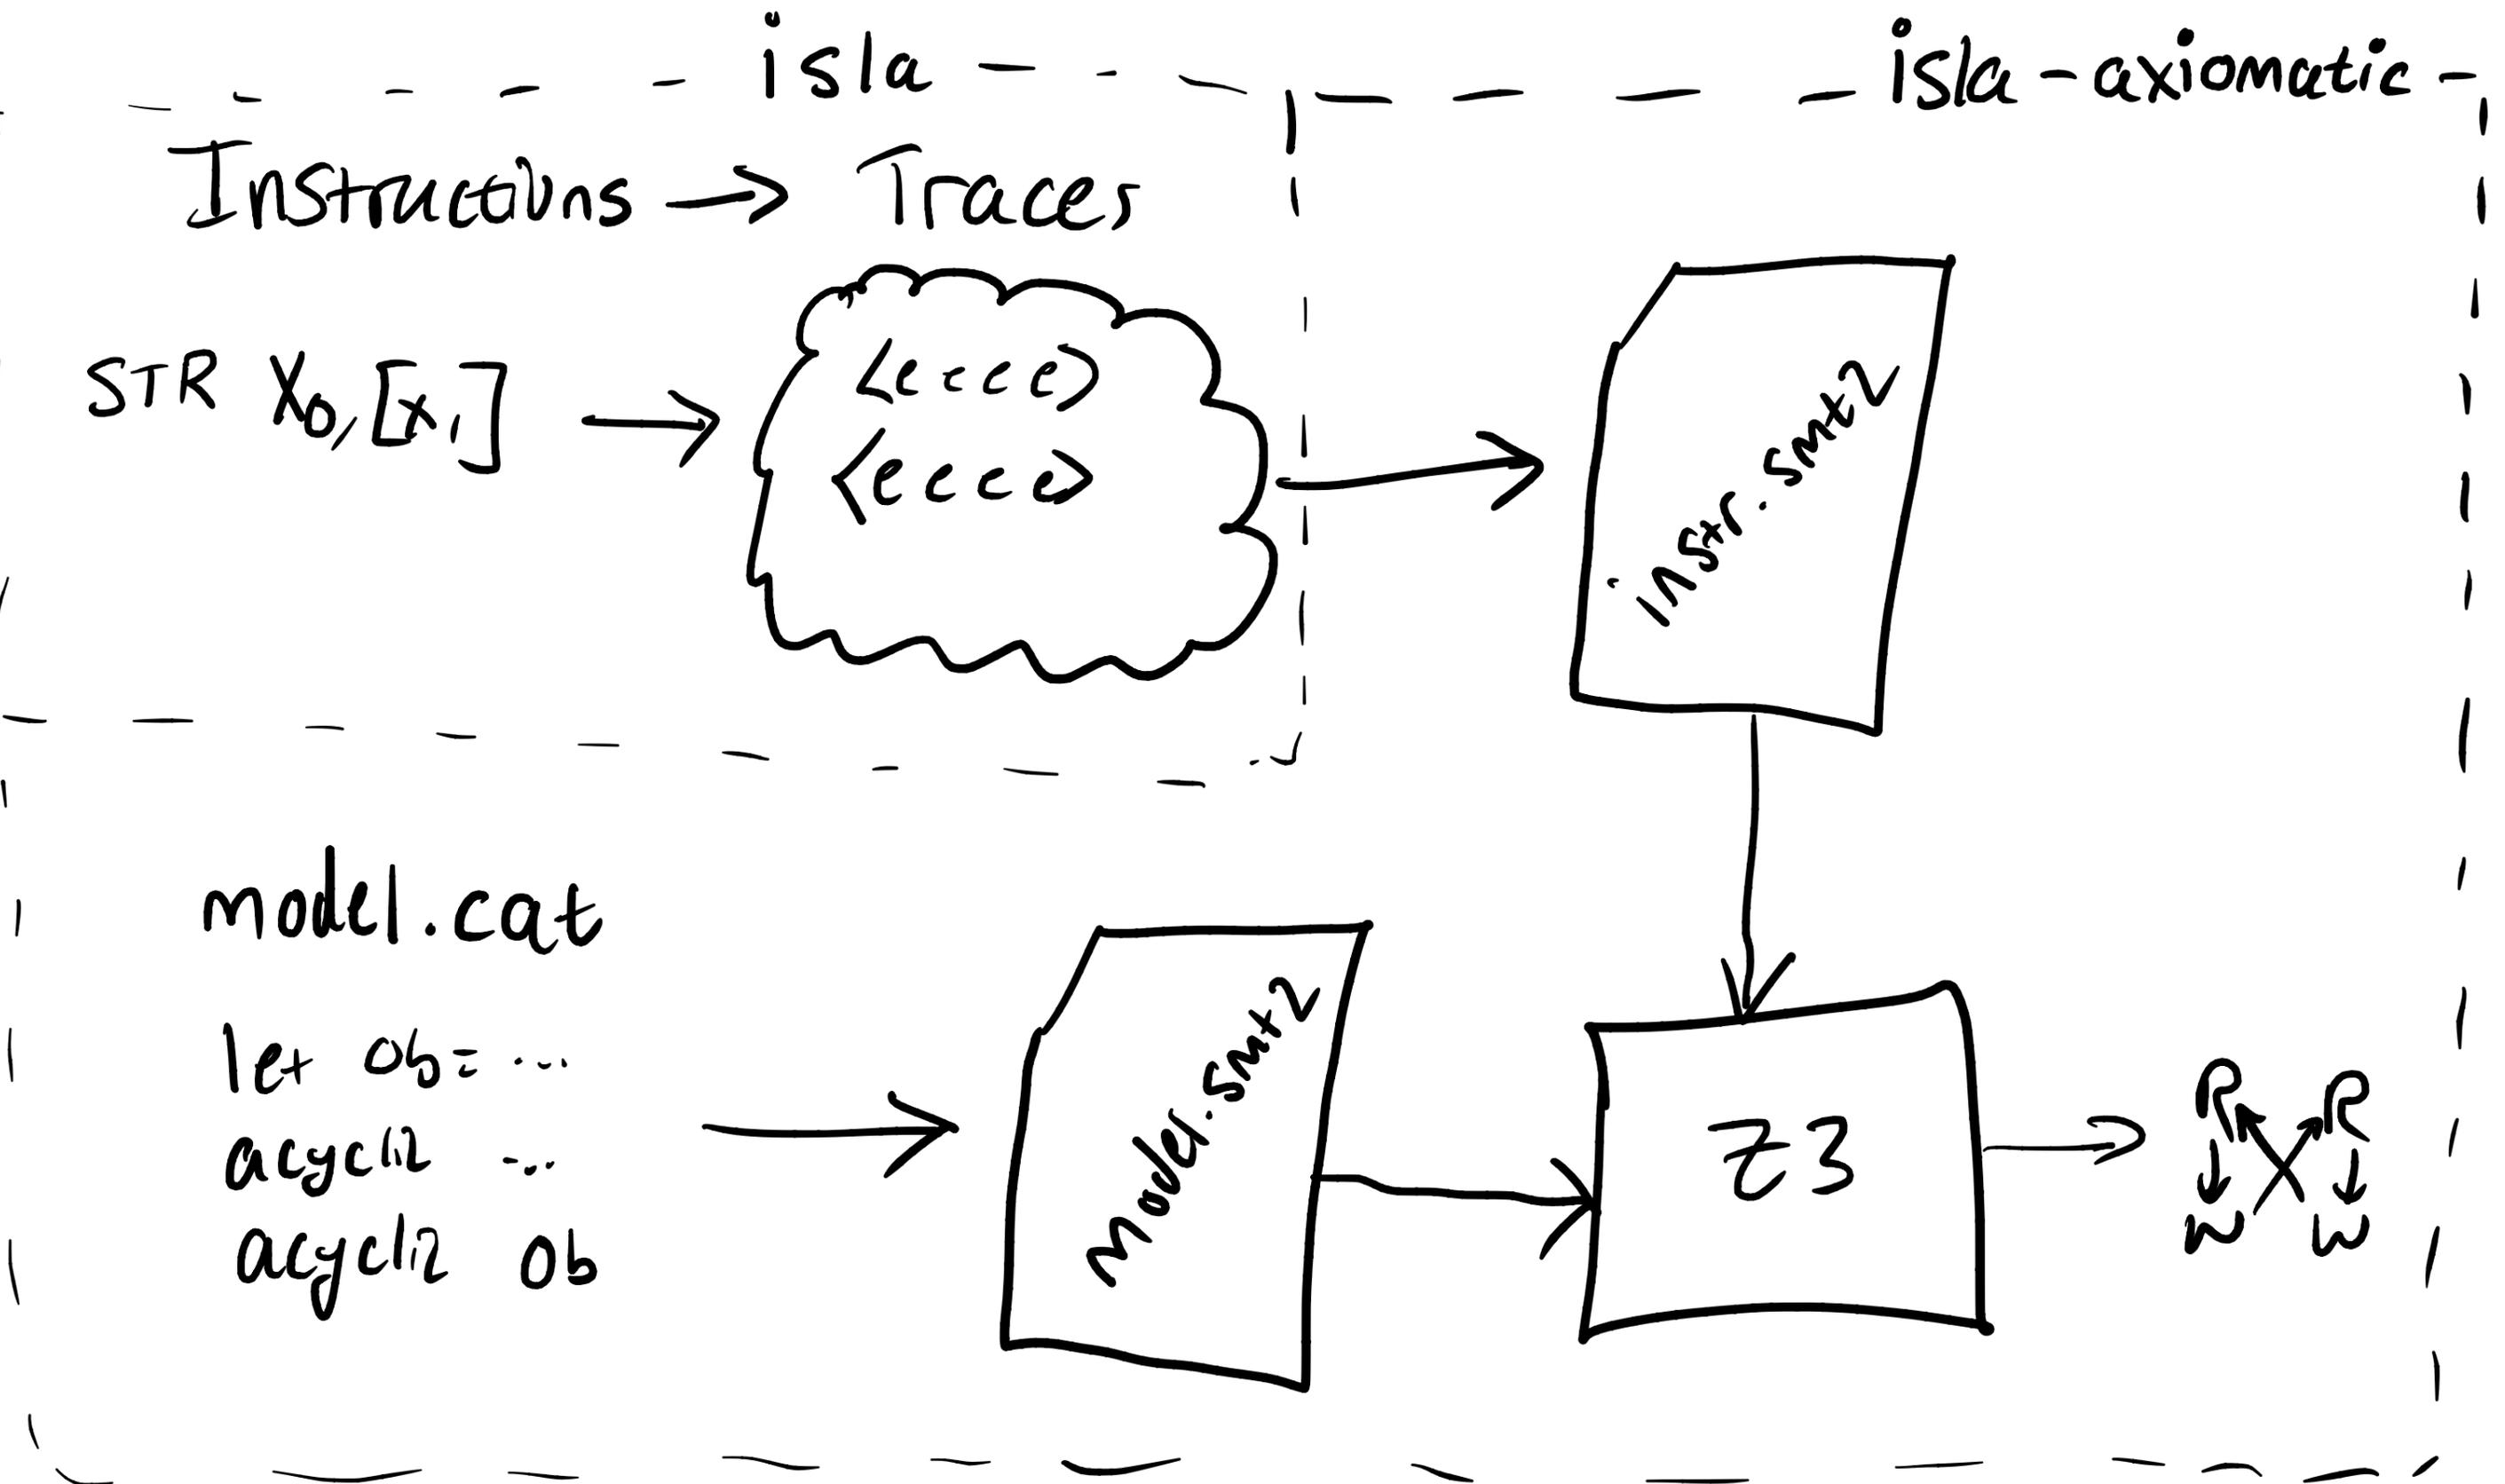
Relations



Actions



# Executable Model



# isla - axiomatic

isla-axiomatic.cl.cam.ac.uk/#

Isla W / aarch64-vmsa-strong Litmus file Memory model Sail architecture (aarch64-vmsa) Run test Options Share

W.toml Memory model EventGraph x

```
11 identity 0x1000 with code;
12 ""
13
14 [thread.0]
15 init = {}
16 code = ""
17 STR X0,[X1]
18 ""
19
20 [thread.0.reset]
21 R0 = "desc3(y, page_table_base)"
22 R1 = "pte3(x, page_table_base)"
23
24 [thread.1]
25 init = {}
26 code = ""
27 LDR X2,[X1]
28 MOV X0,X2
29 LDR X2,[X3]
30 ""
31
32 [thread.1.reset]
33 R1 = "x"
34 R3 = "x"
35 VBAR_EL1 = "extz(0x1000, 64)"
36
37 "PSTATE.SP" = "0b0"
38 "PSTATE.EL" = "0b00"
39
40 [section.thread1_el1_handler]
41 address = "0x1400"
42 code = ""
43 MOV X2,#0
44
45 MRS X13,ELR_EL1
46 ADD X13,X13,#4
47 MSR ELR_EL1,X13
48 ERET
49 ""
50
51 [final]
52 assertion = "1:X0=1 & 1:X2=0"
```

Initial State

Thread 0

a: str x0, [x1]: W s1:l3pte(x) = s1:l3desc(y)

Thread 1

a1: T s2:l3pte(x) iiio a2: ldr x2, [x1]: R pa1 = 0x1

b1: T s1:l3pte(x) iiio b2: ldr x2, [x3]: Fault

c: eret

# More examples: TLBs

Init:  $x \mapsto \text{old}$

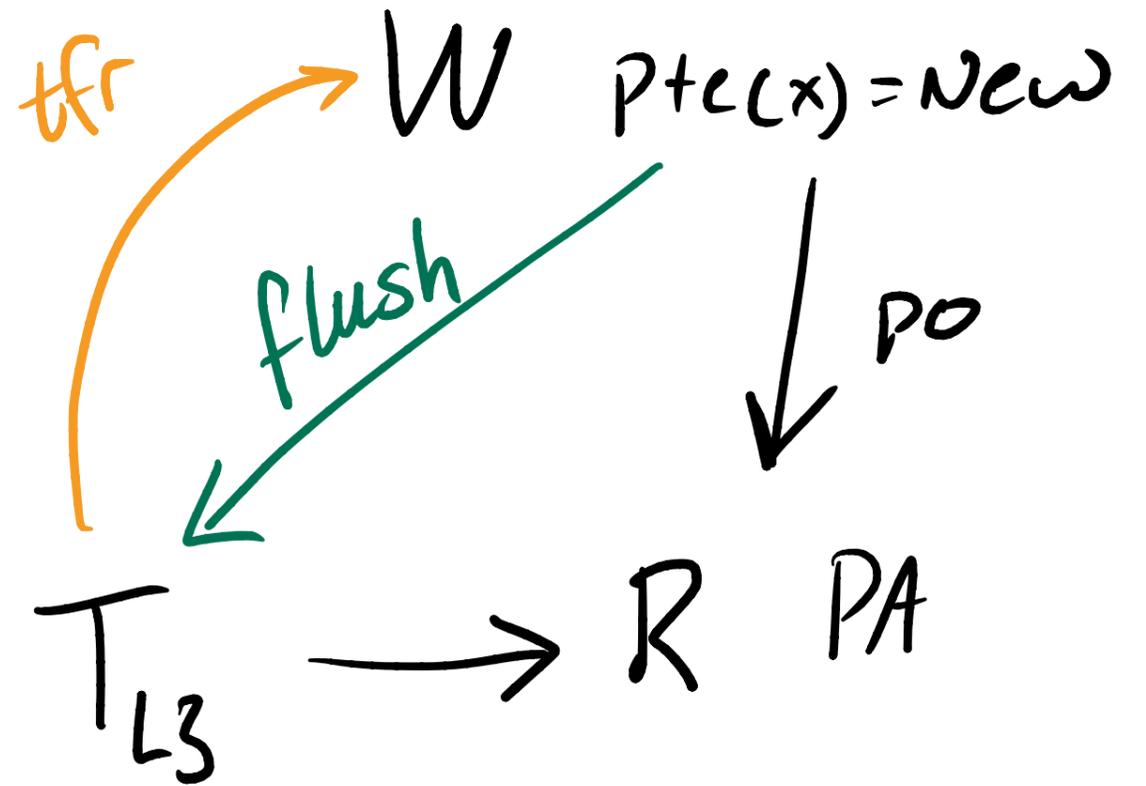
Store  $\text{pte}(x) = \text{New}$

DSB.SY

ISB

$r_0 \leftarrow \text{Load } x$

Outcome  $r_0$  sees old  
*allowed*



# More examples: TLB Maintenance

Init:  $x \mapsto \text{old}$

Store  $\text{pte}(x) = \text{New}$

DSB.SY

TLBI x

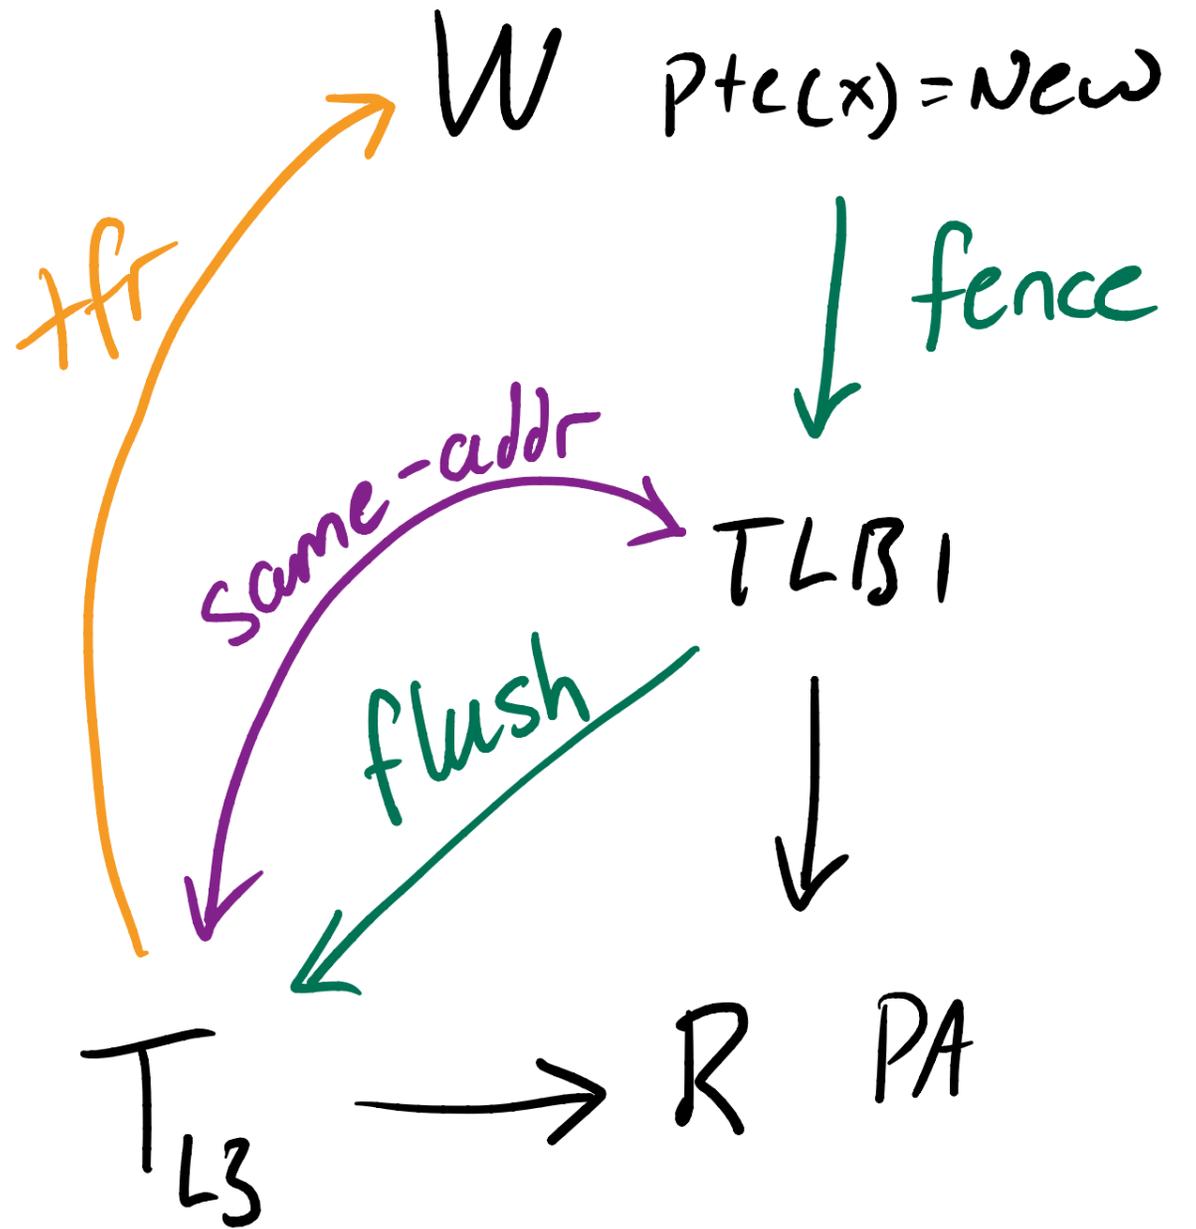
DSB.SY

ISB

$r_0 \leftarrow \text{Load } x$

Outcome:  $r_0$  sees old

*Forbidden*



# More examples: Breaking

init  $x \mapsto p$   
 $\wedge$   $y \mapsto p$

Store  $pte(x) = \text{NULL}$

DSB.SY

TLBI X

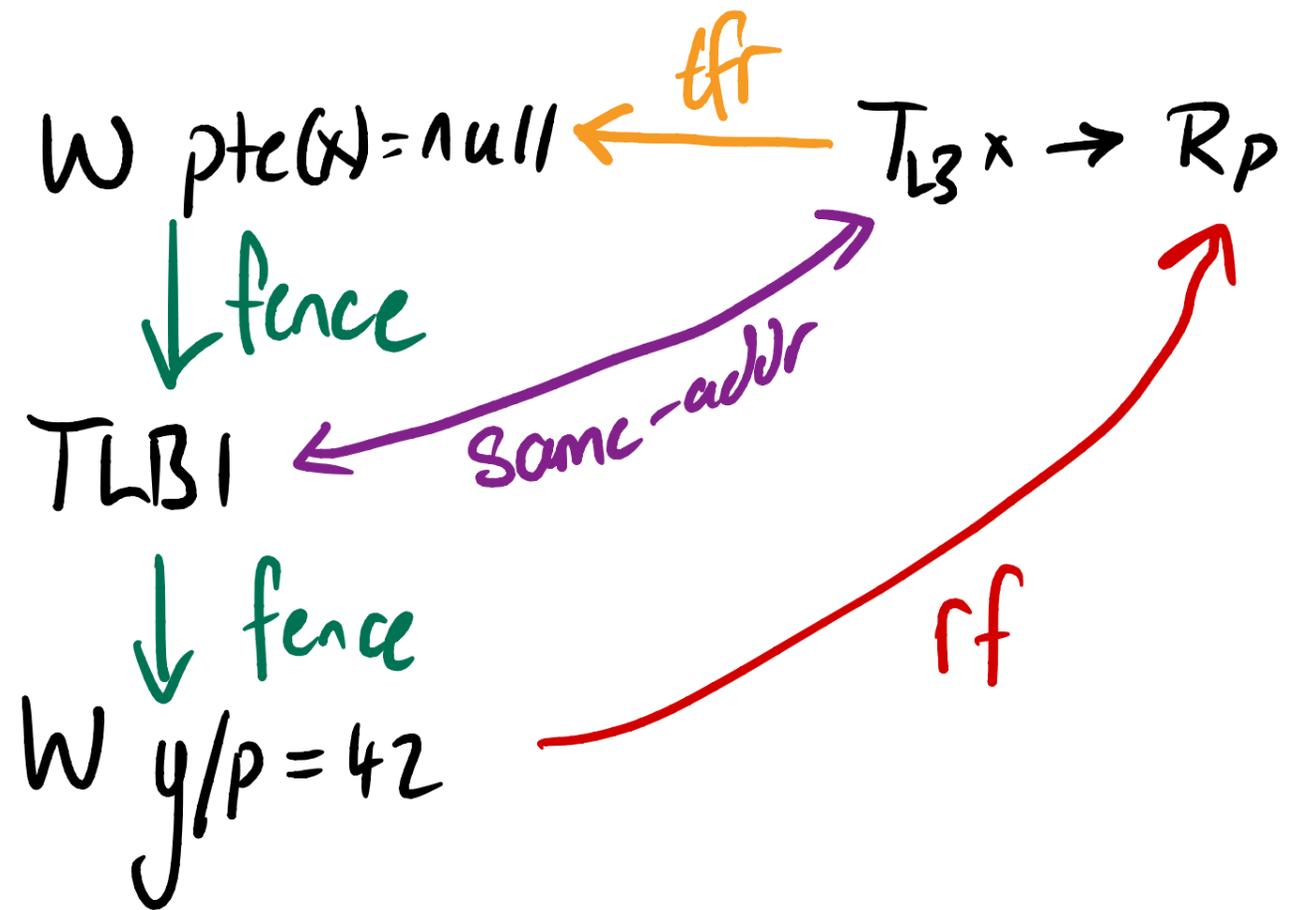
DSB.SY

Store  $y = 42$

$r_0 \leftarrow \text{load } x$

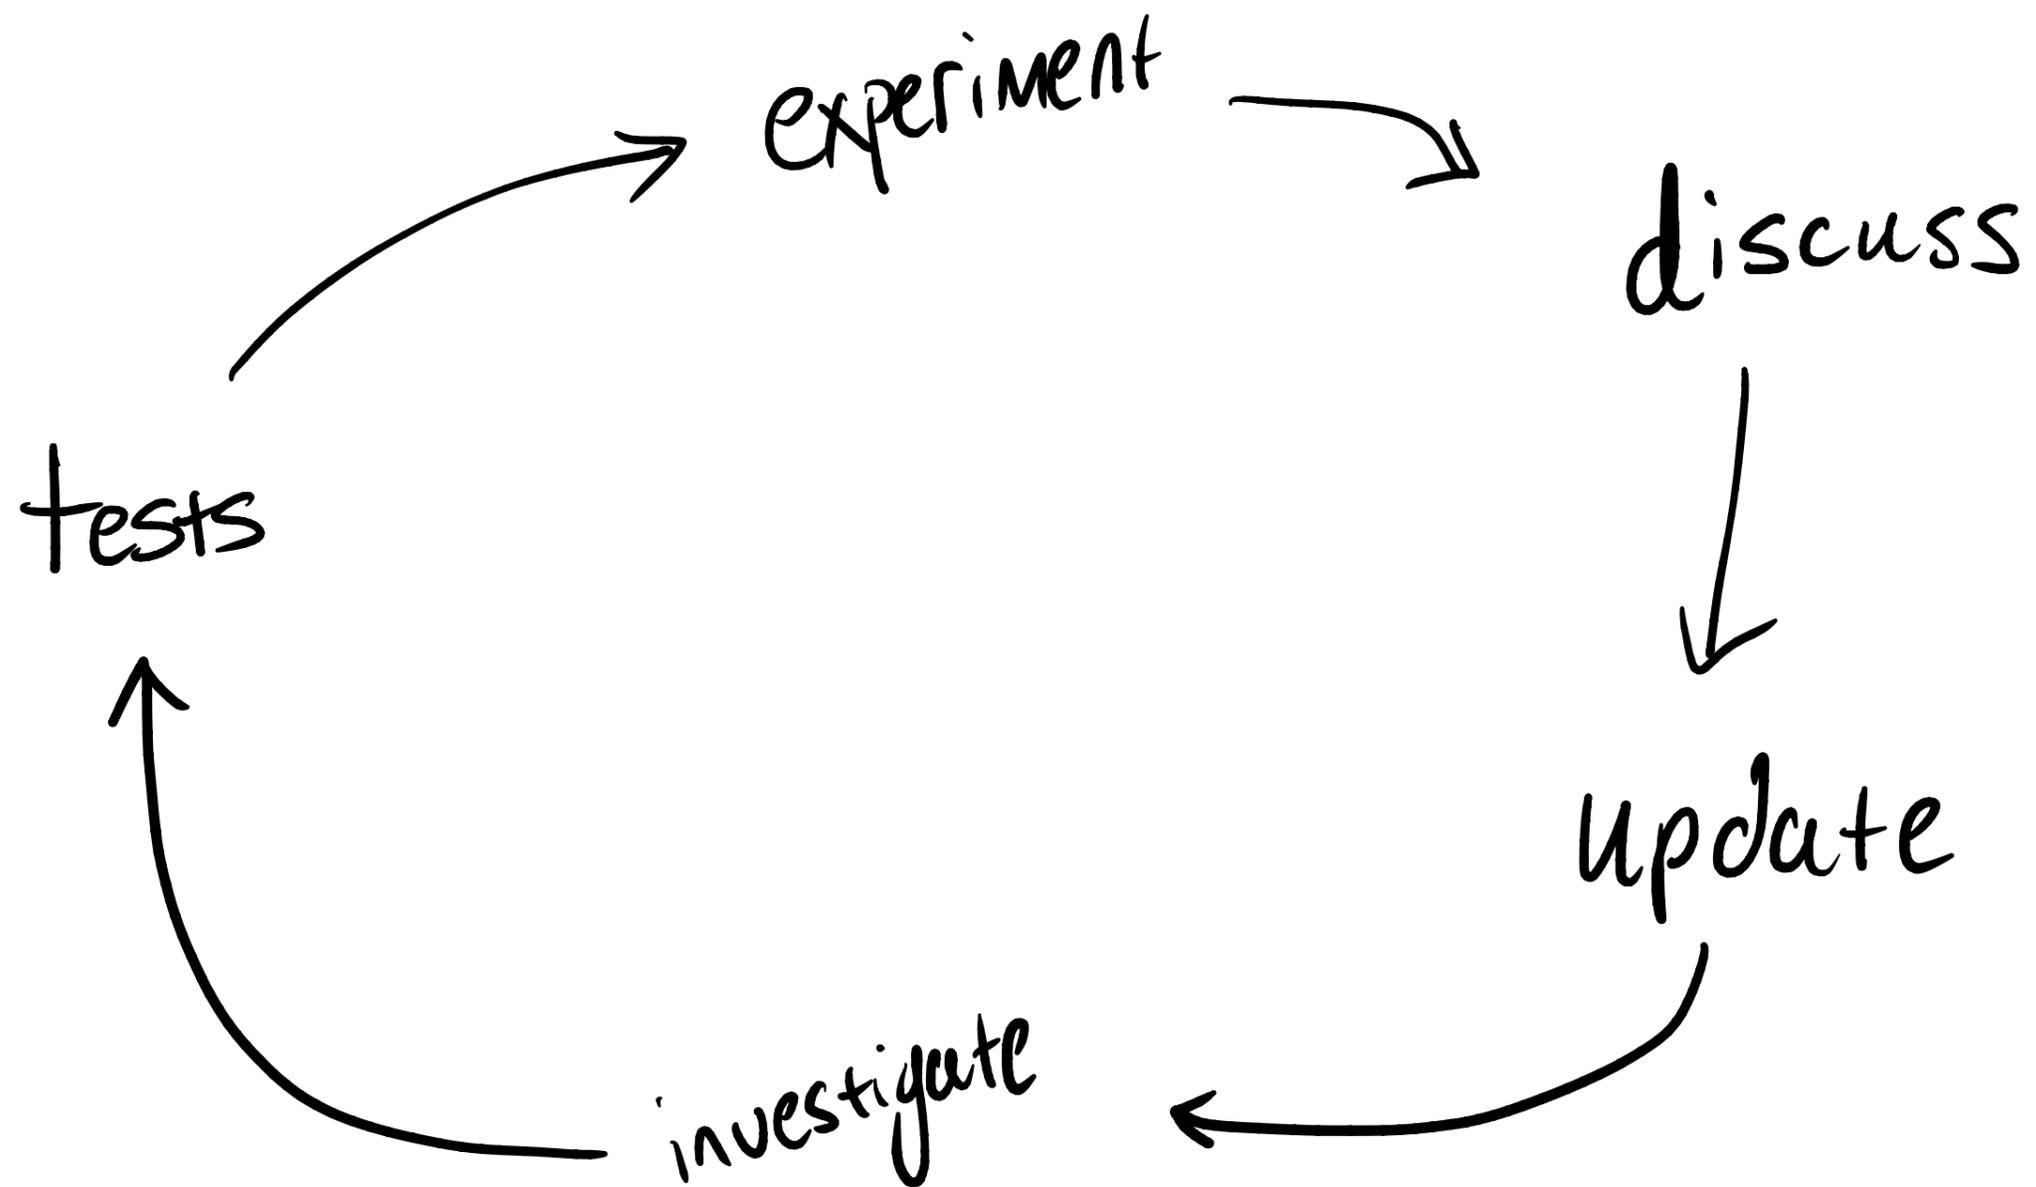
Outcome  $r_0 = 42$

Forbidden

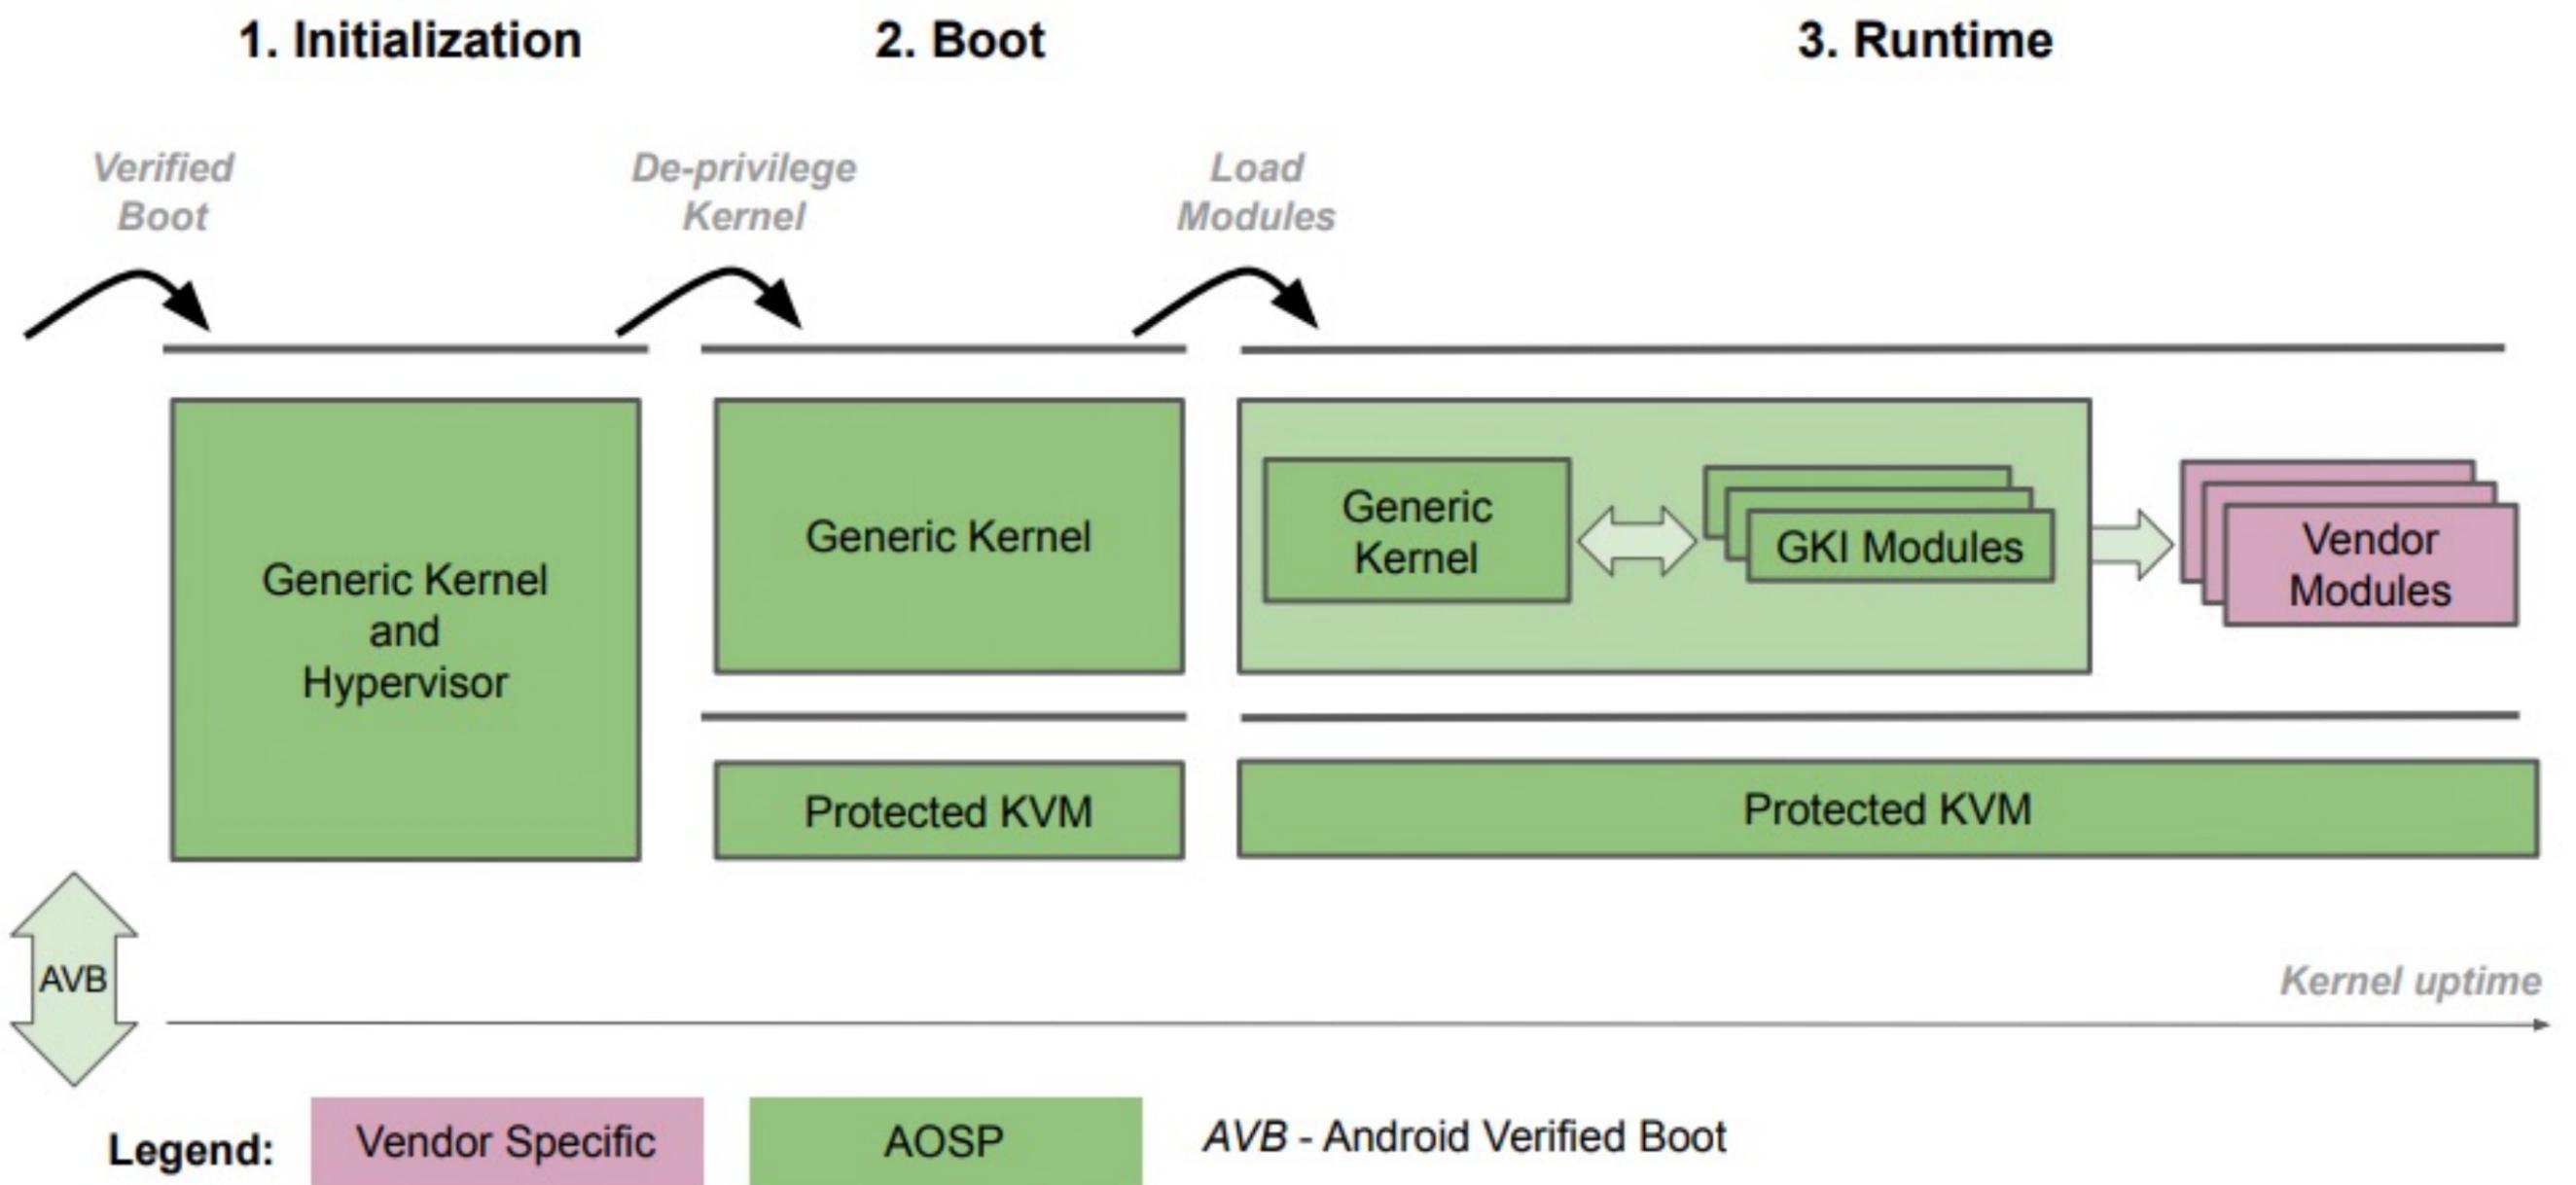


... and  
much  
more ...

# An aside: Clarifying the Architecture

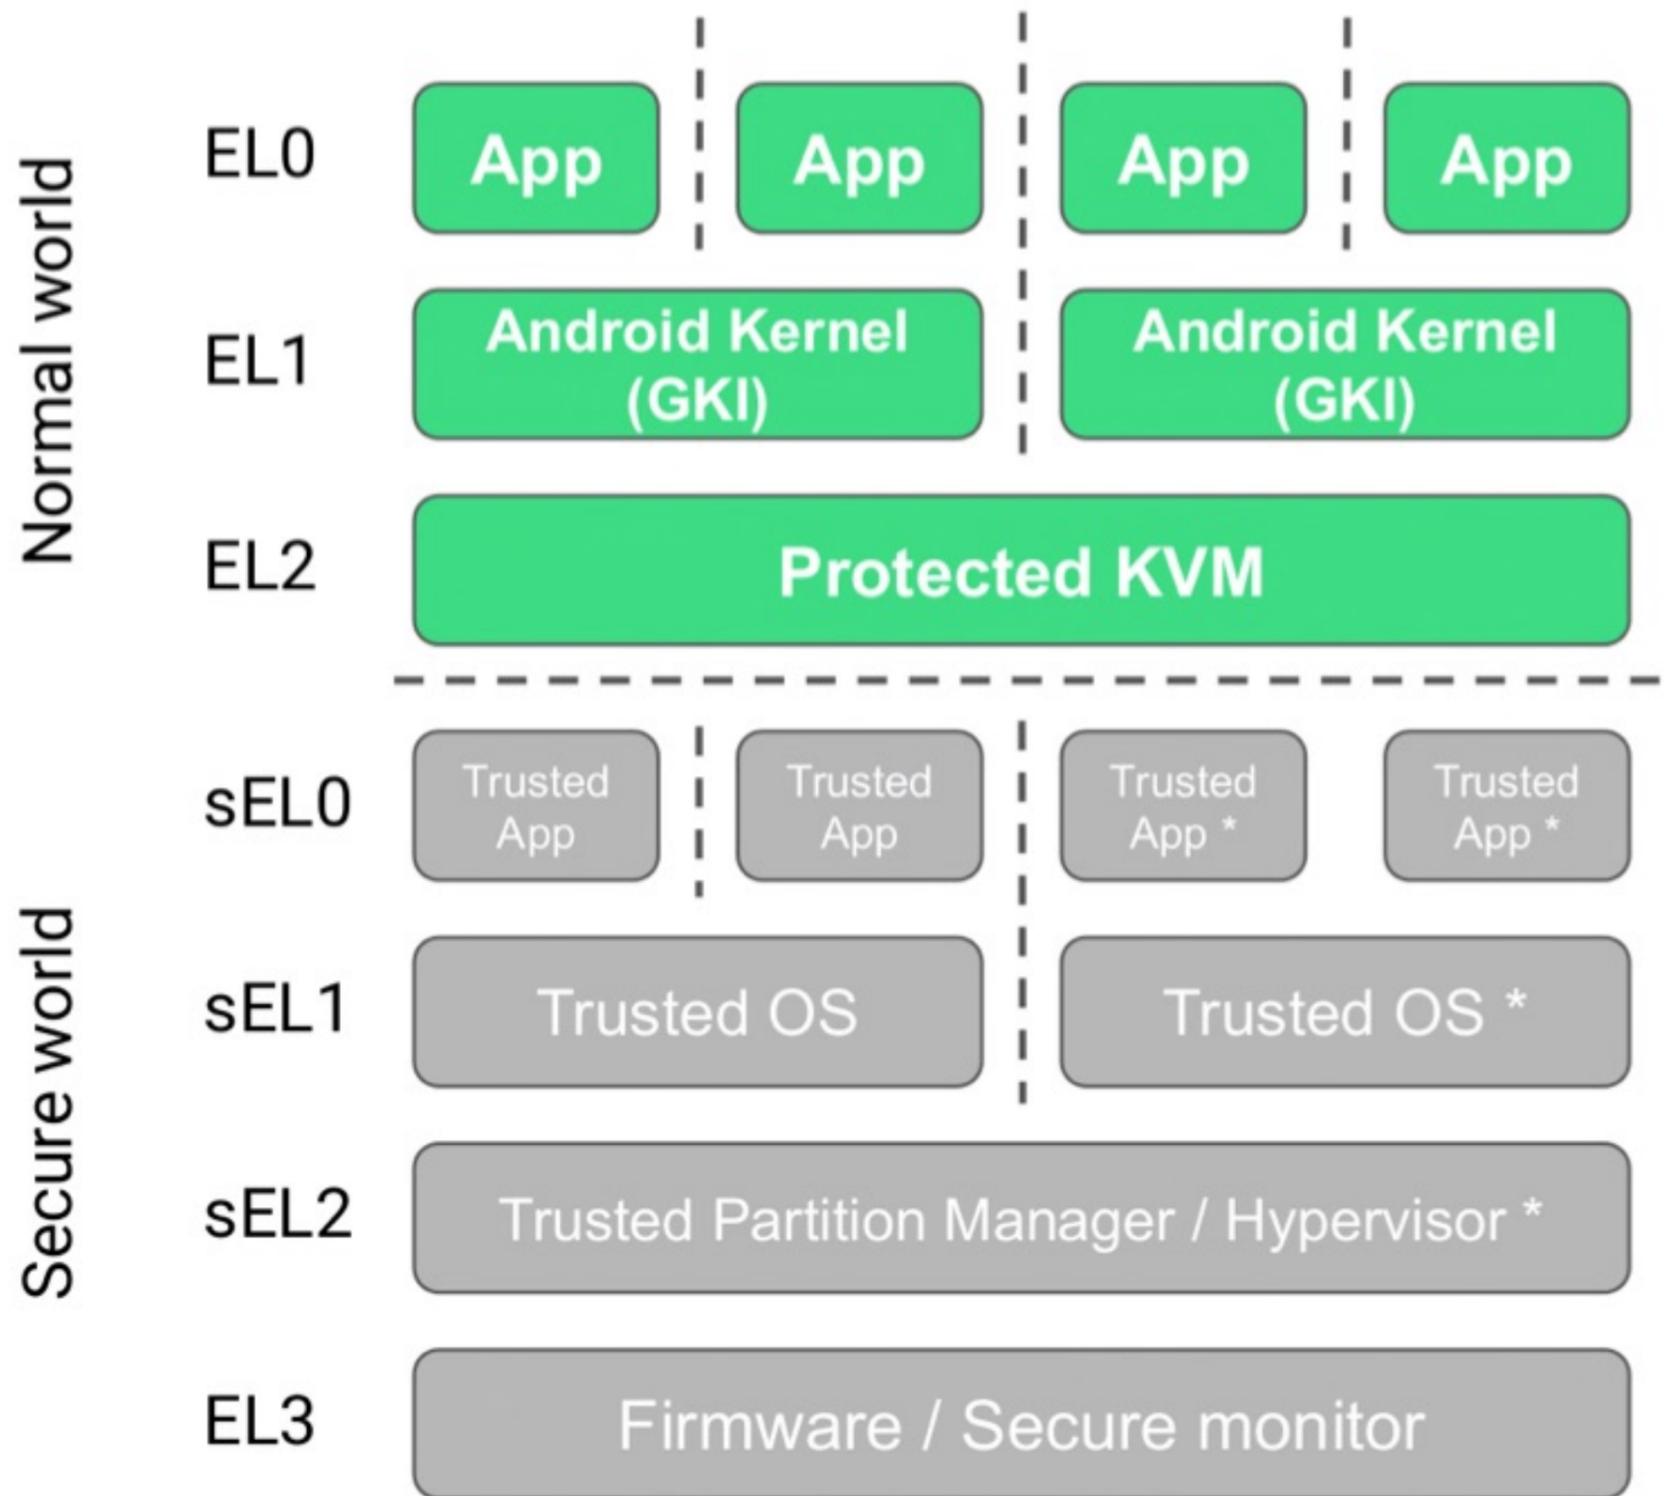


PKVM



android

Credit: Will Deacon, Google 2020 @ KVM Forum



Credit: Will Deacon KVM Forum 2020

# Recap

Clarify architecture

Formal semantics

Tooling