

Relaxed

Virtual

Memory

— Armv8 edition —


Ben Simmer¹, Alasdair Armstrong¹, Jean Pichon-Pharabod²,

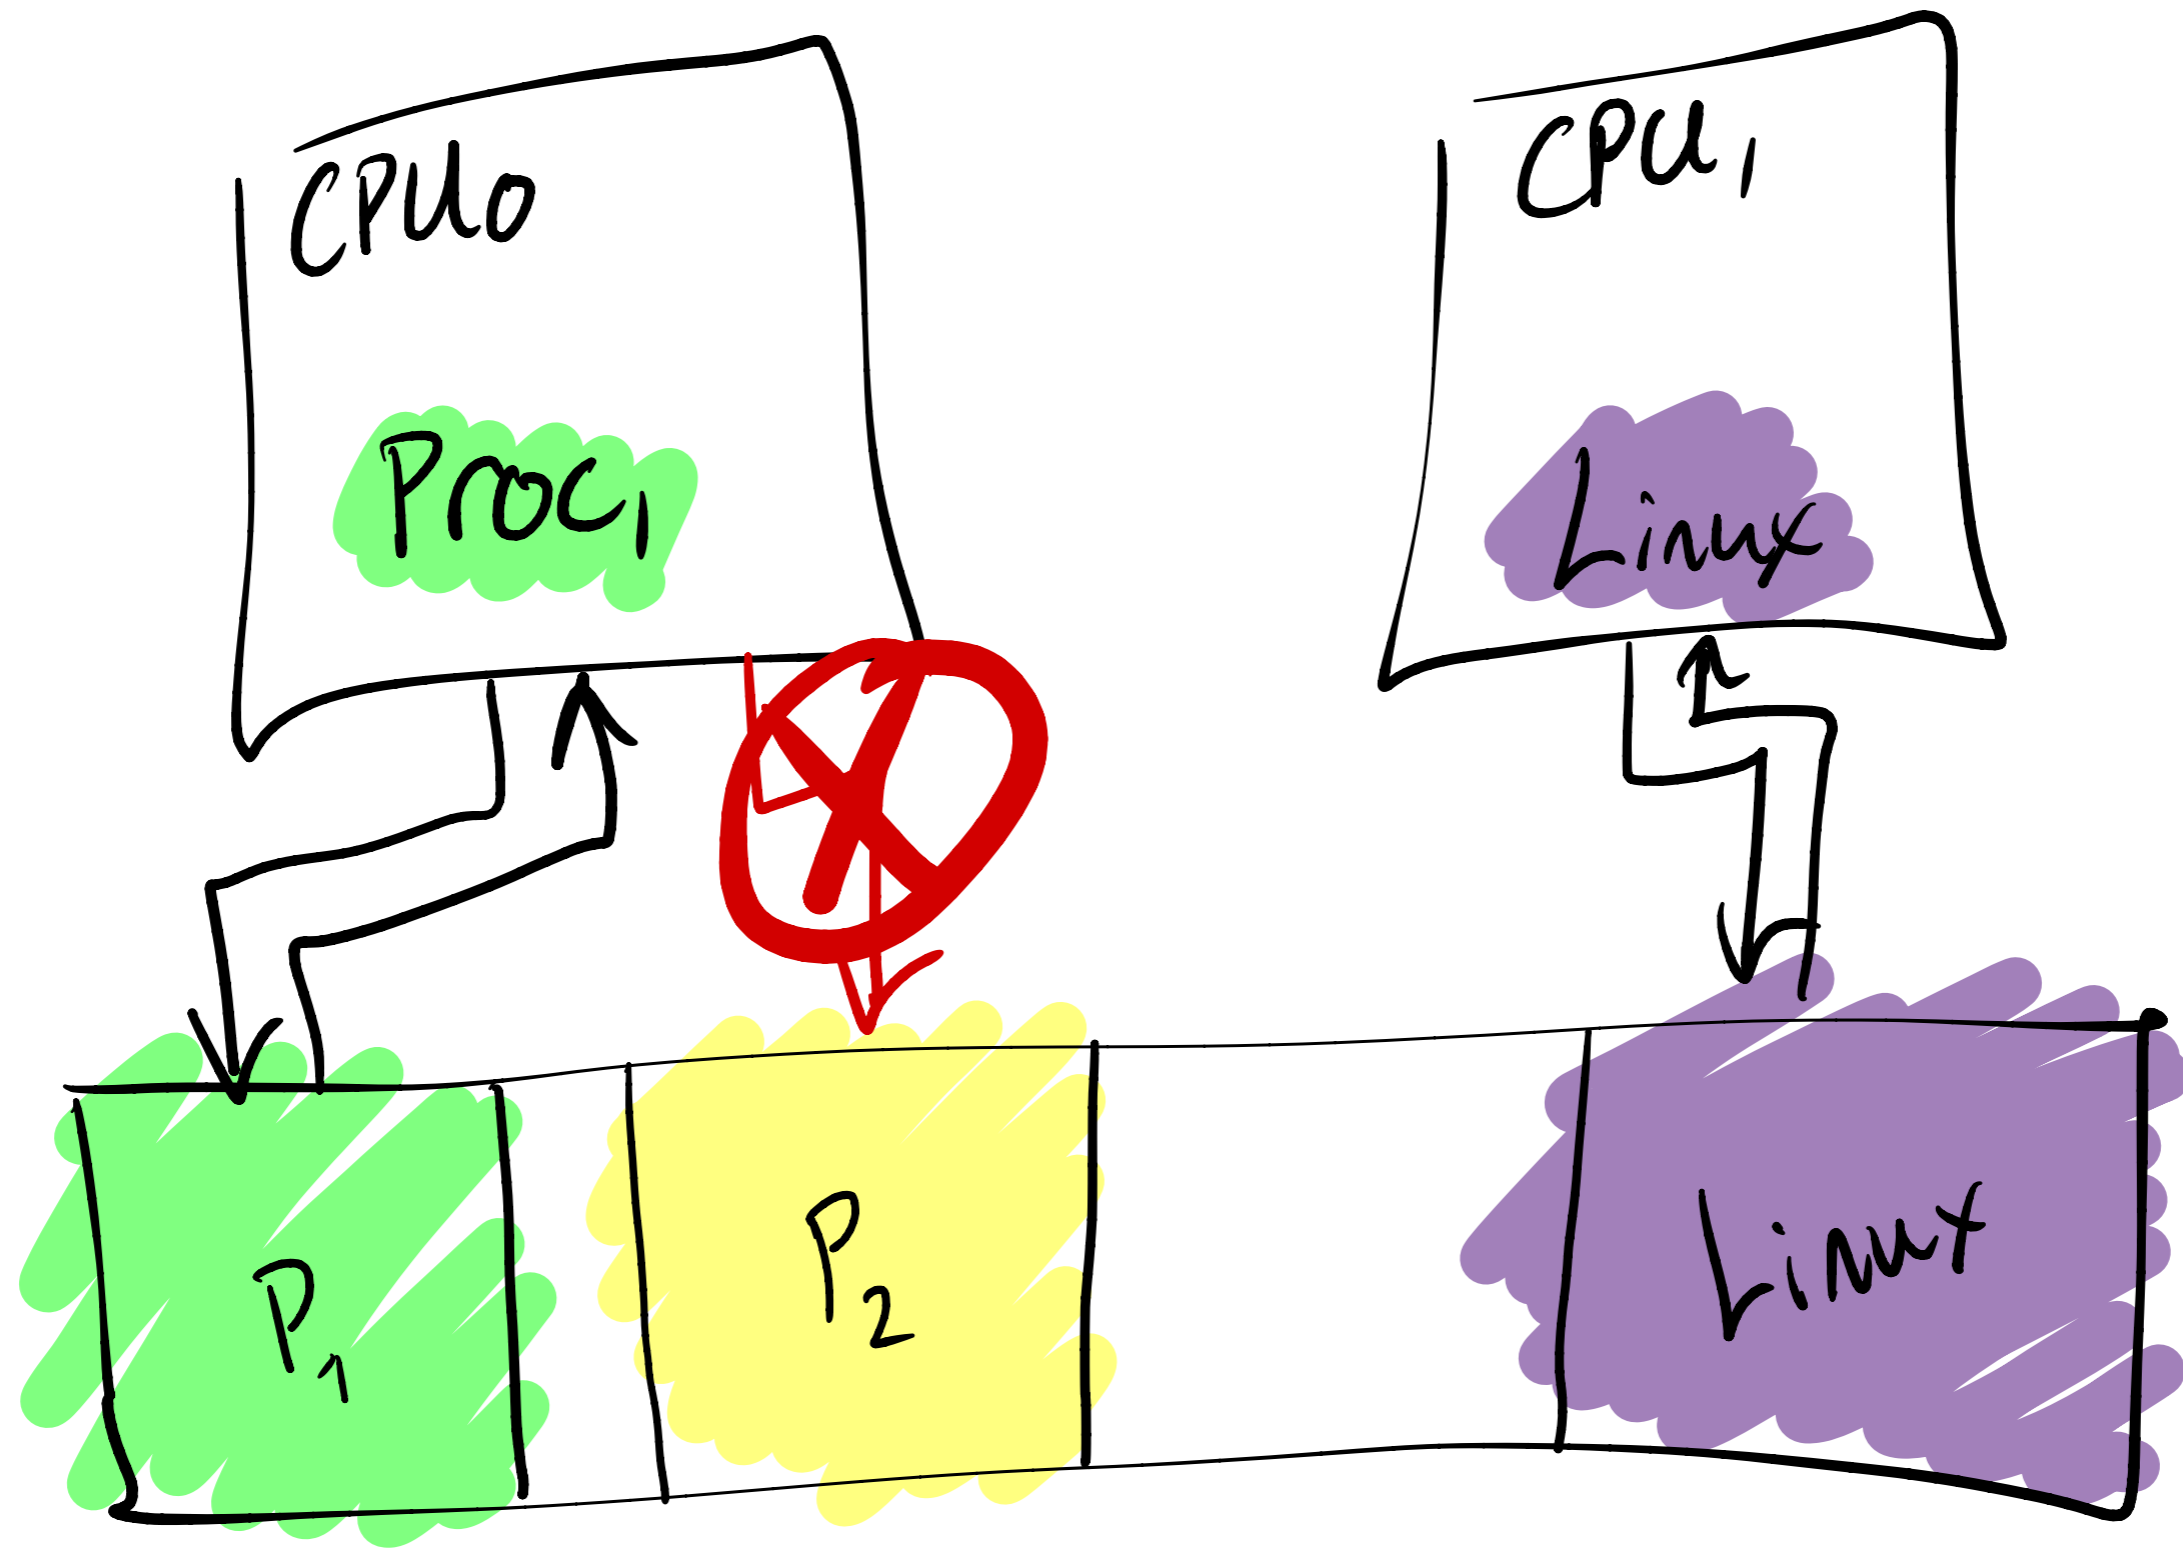
Christopher Pulte¹, Richard Grisenthwaite³, Peter Sewell¹

1. Cambridge, UK

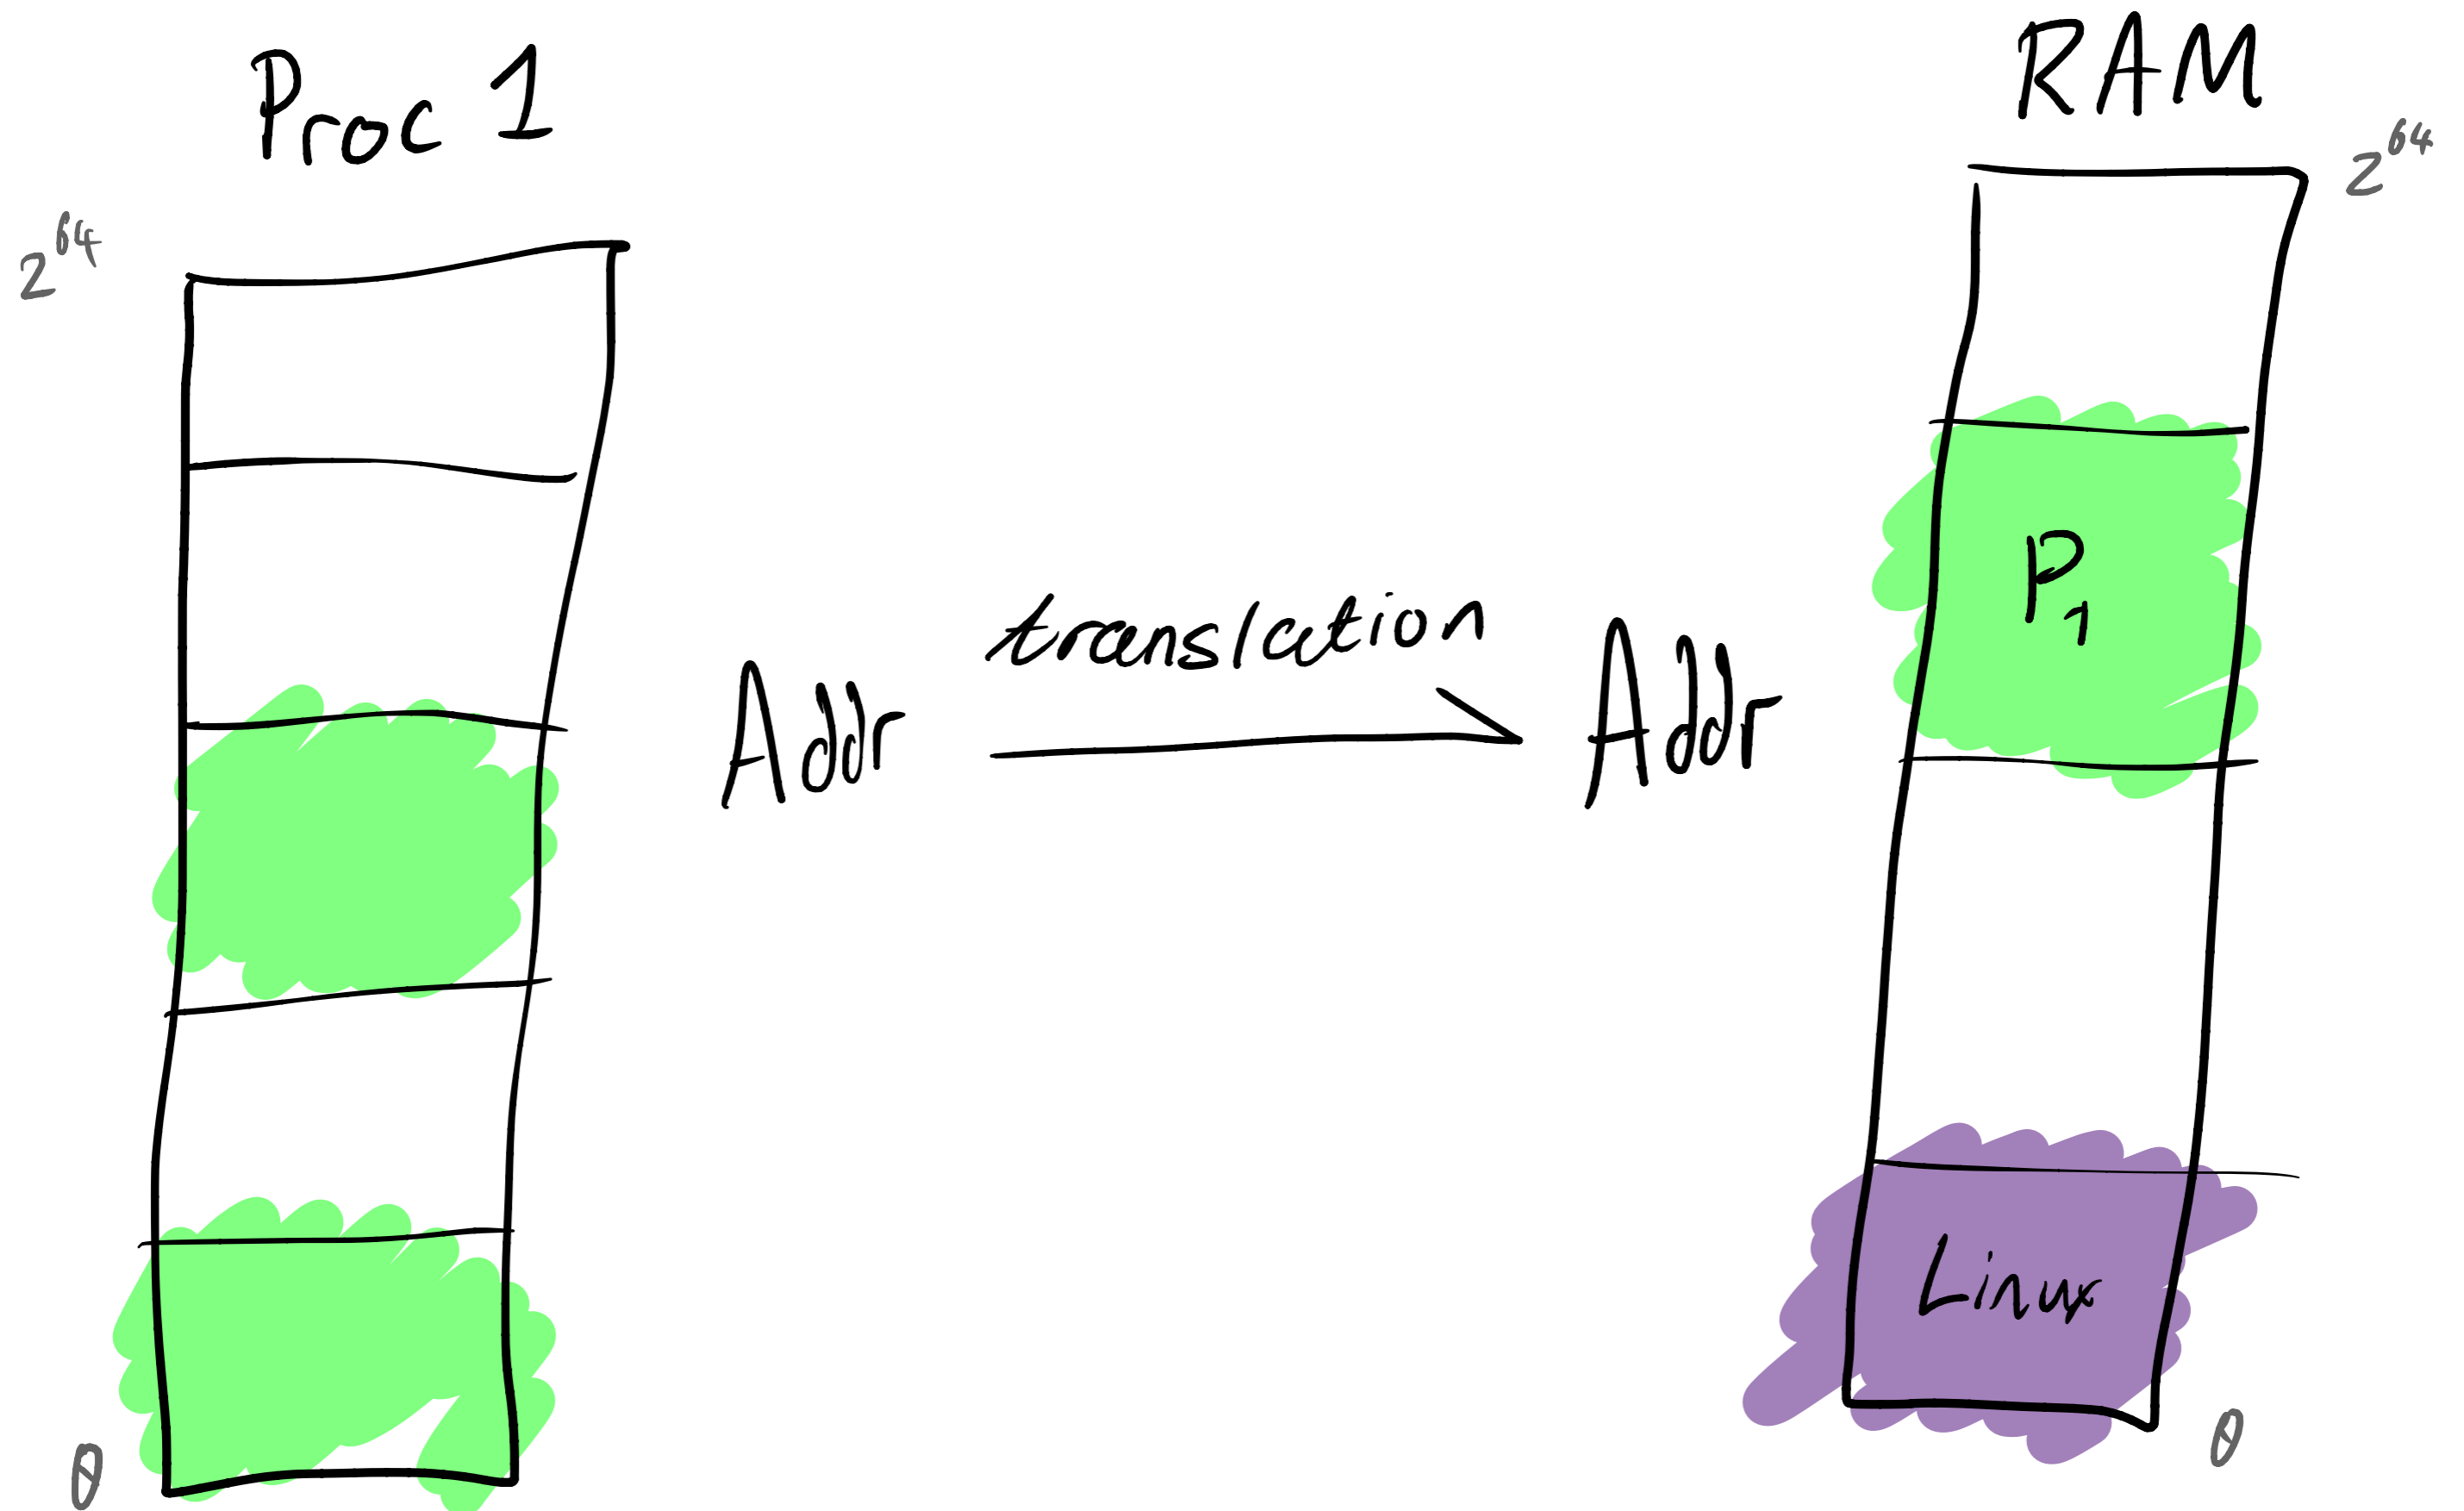
2. Aarhus, DK

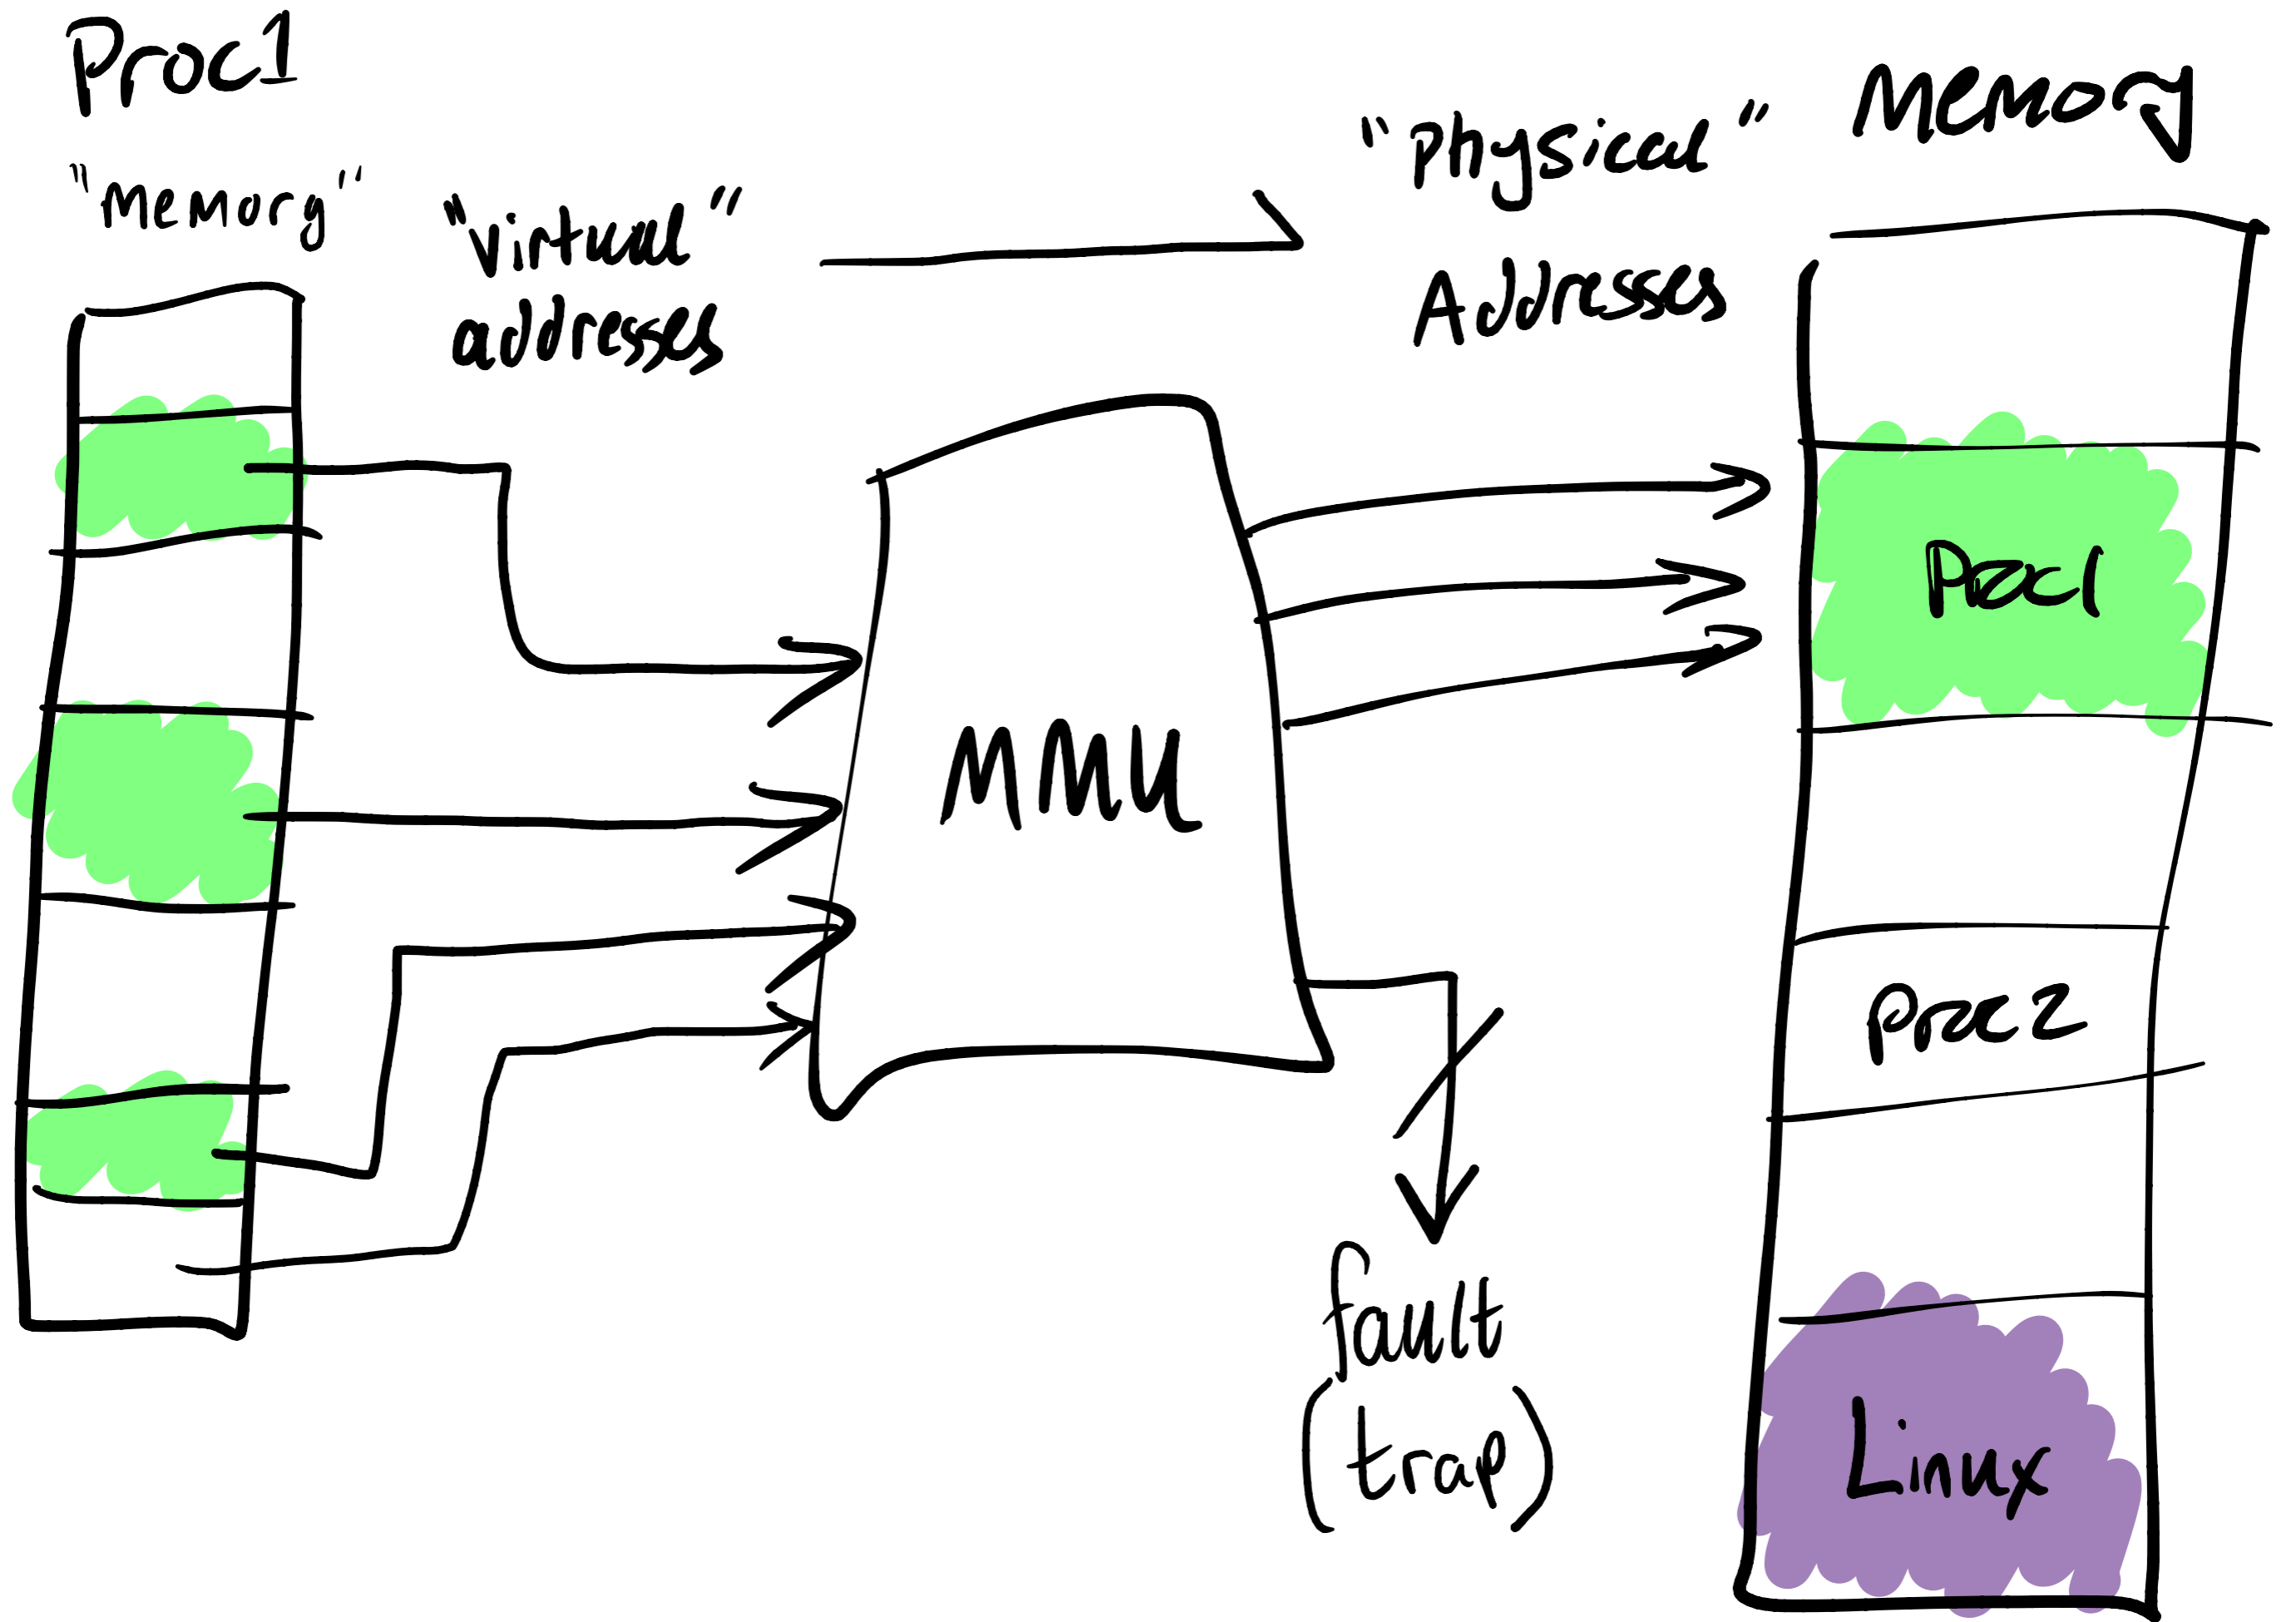
3. Arm Ltd., UK

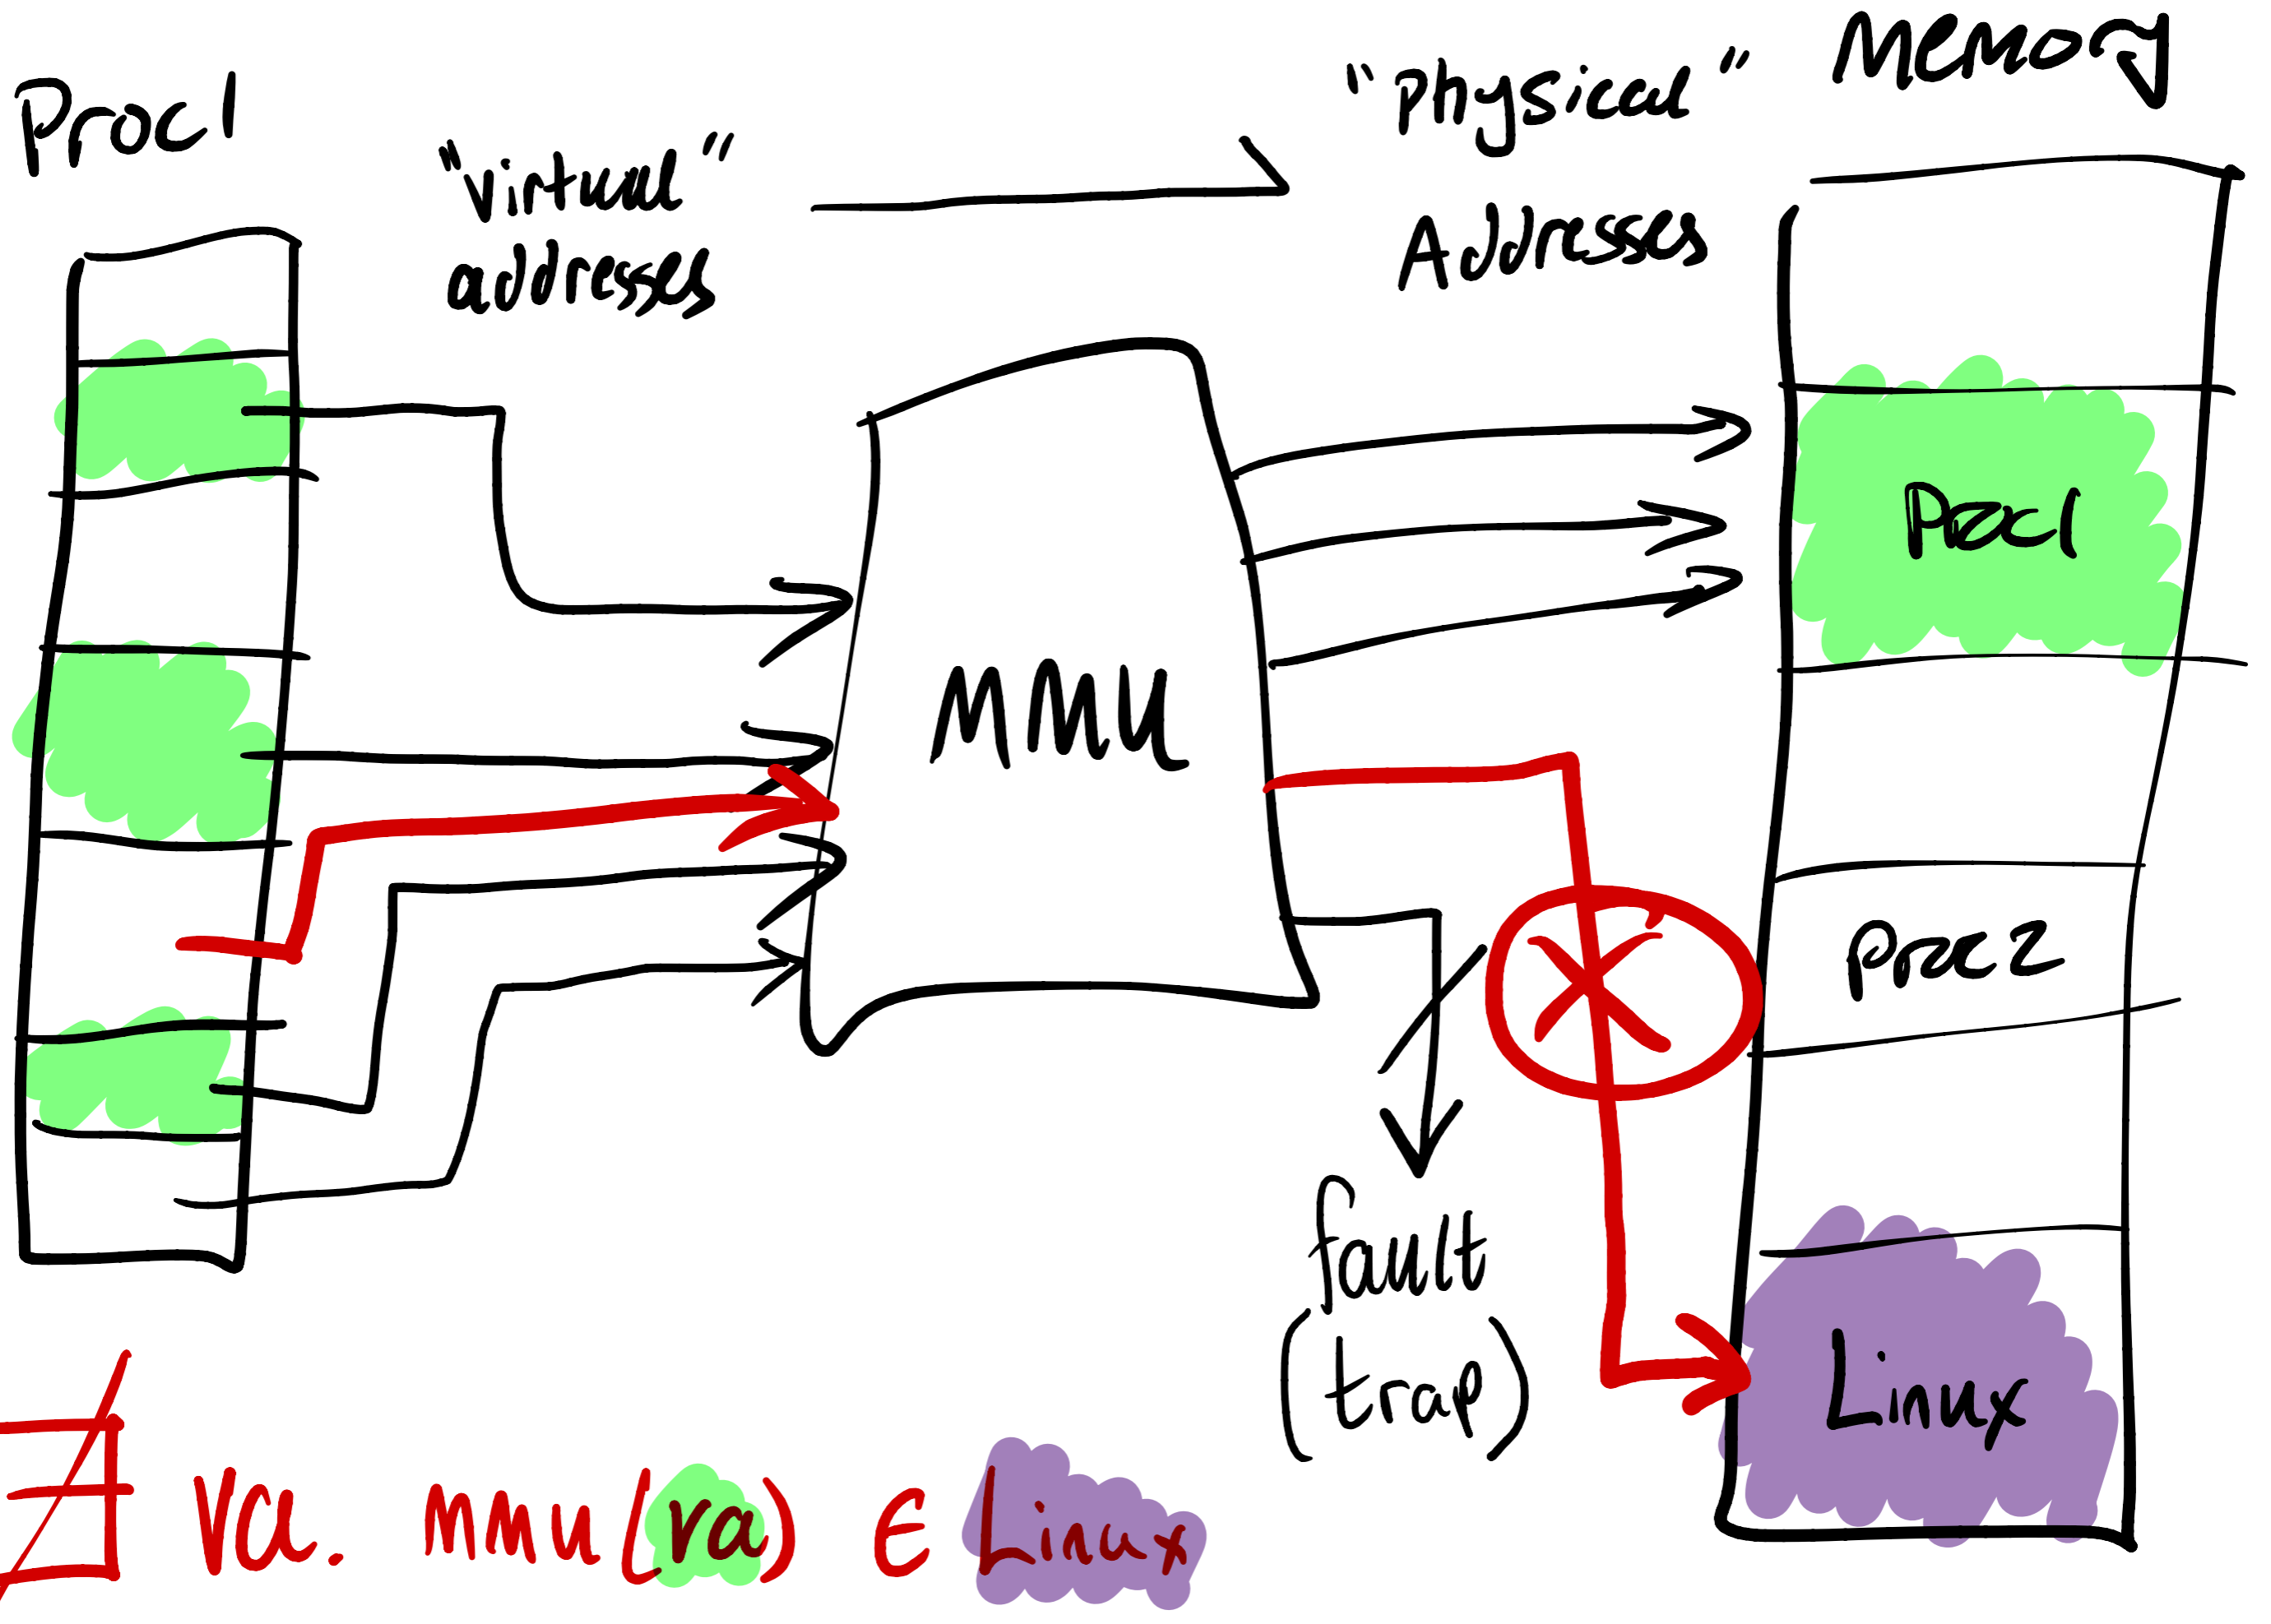




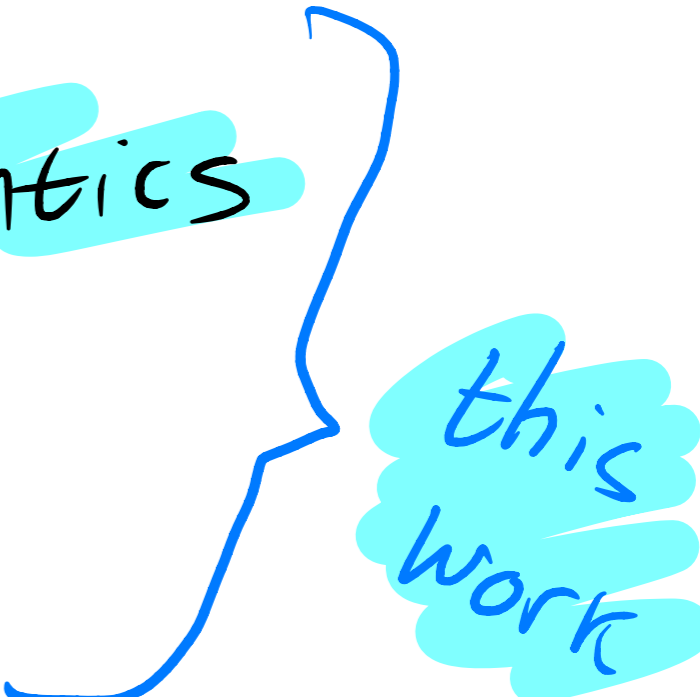
Virtual Memory



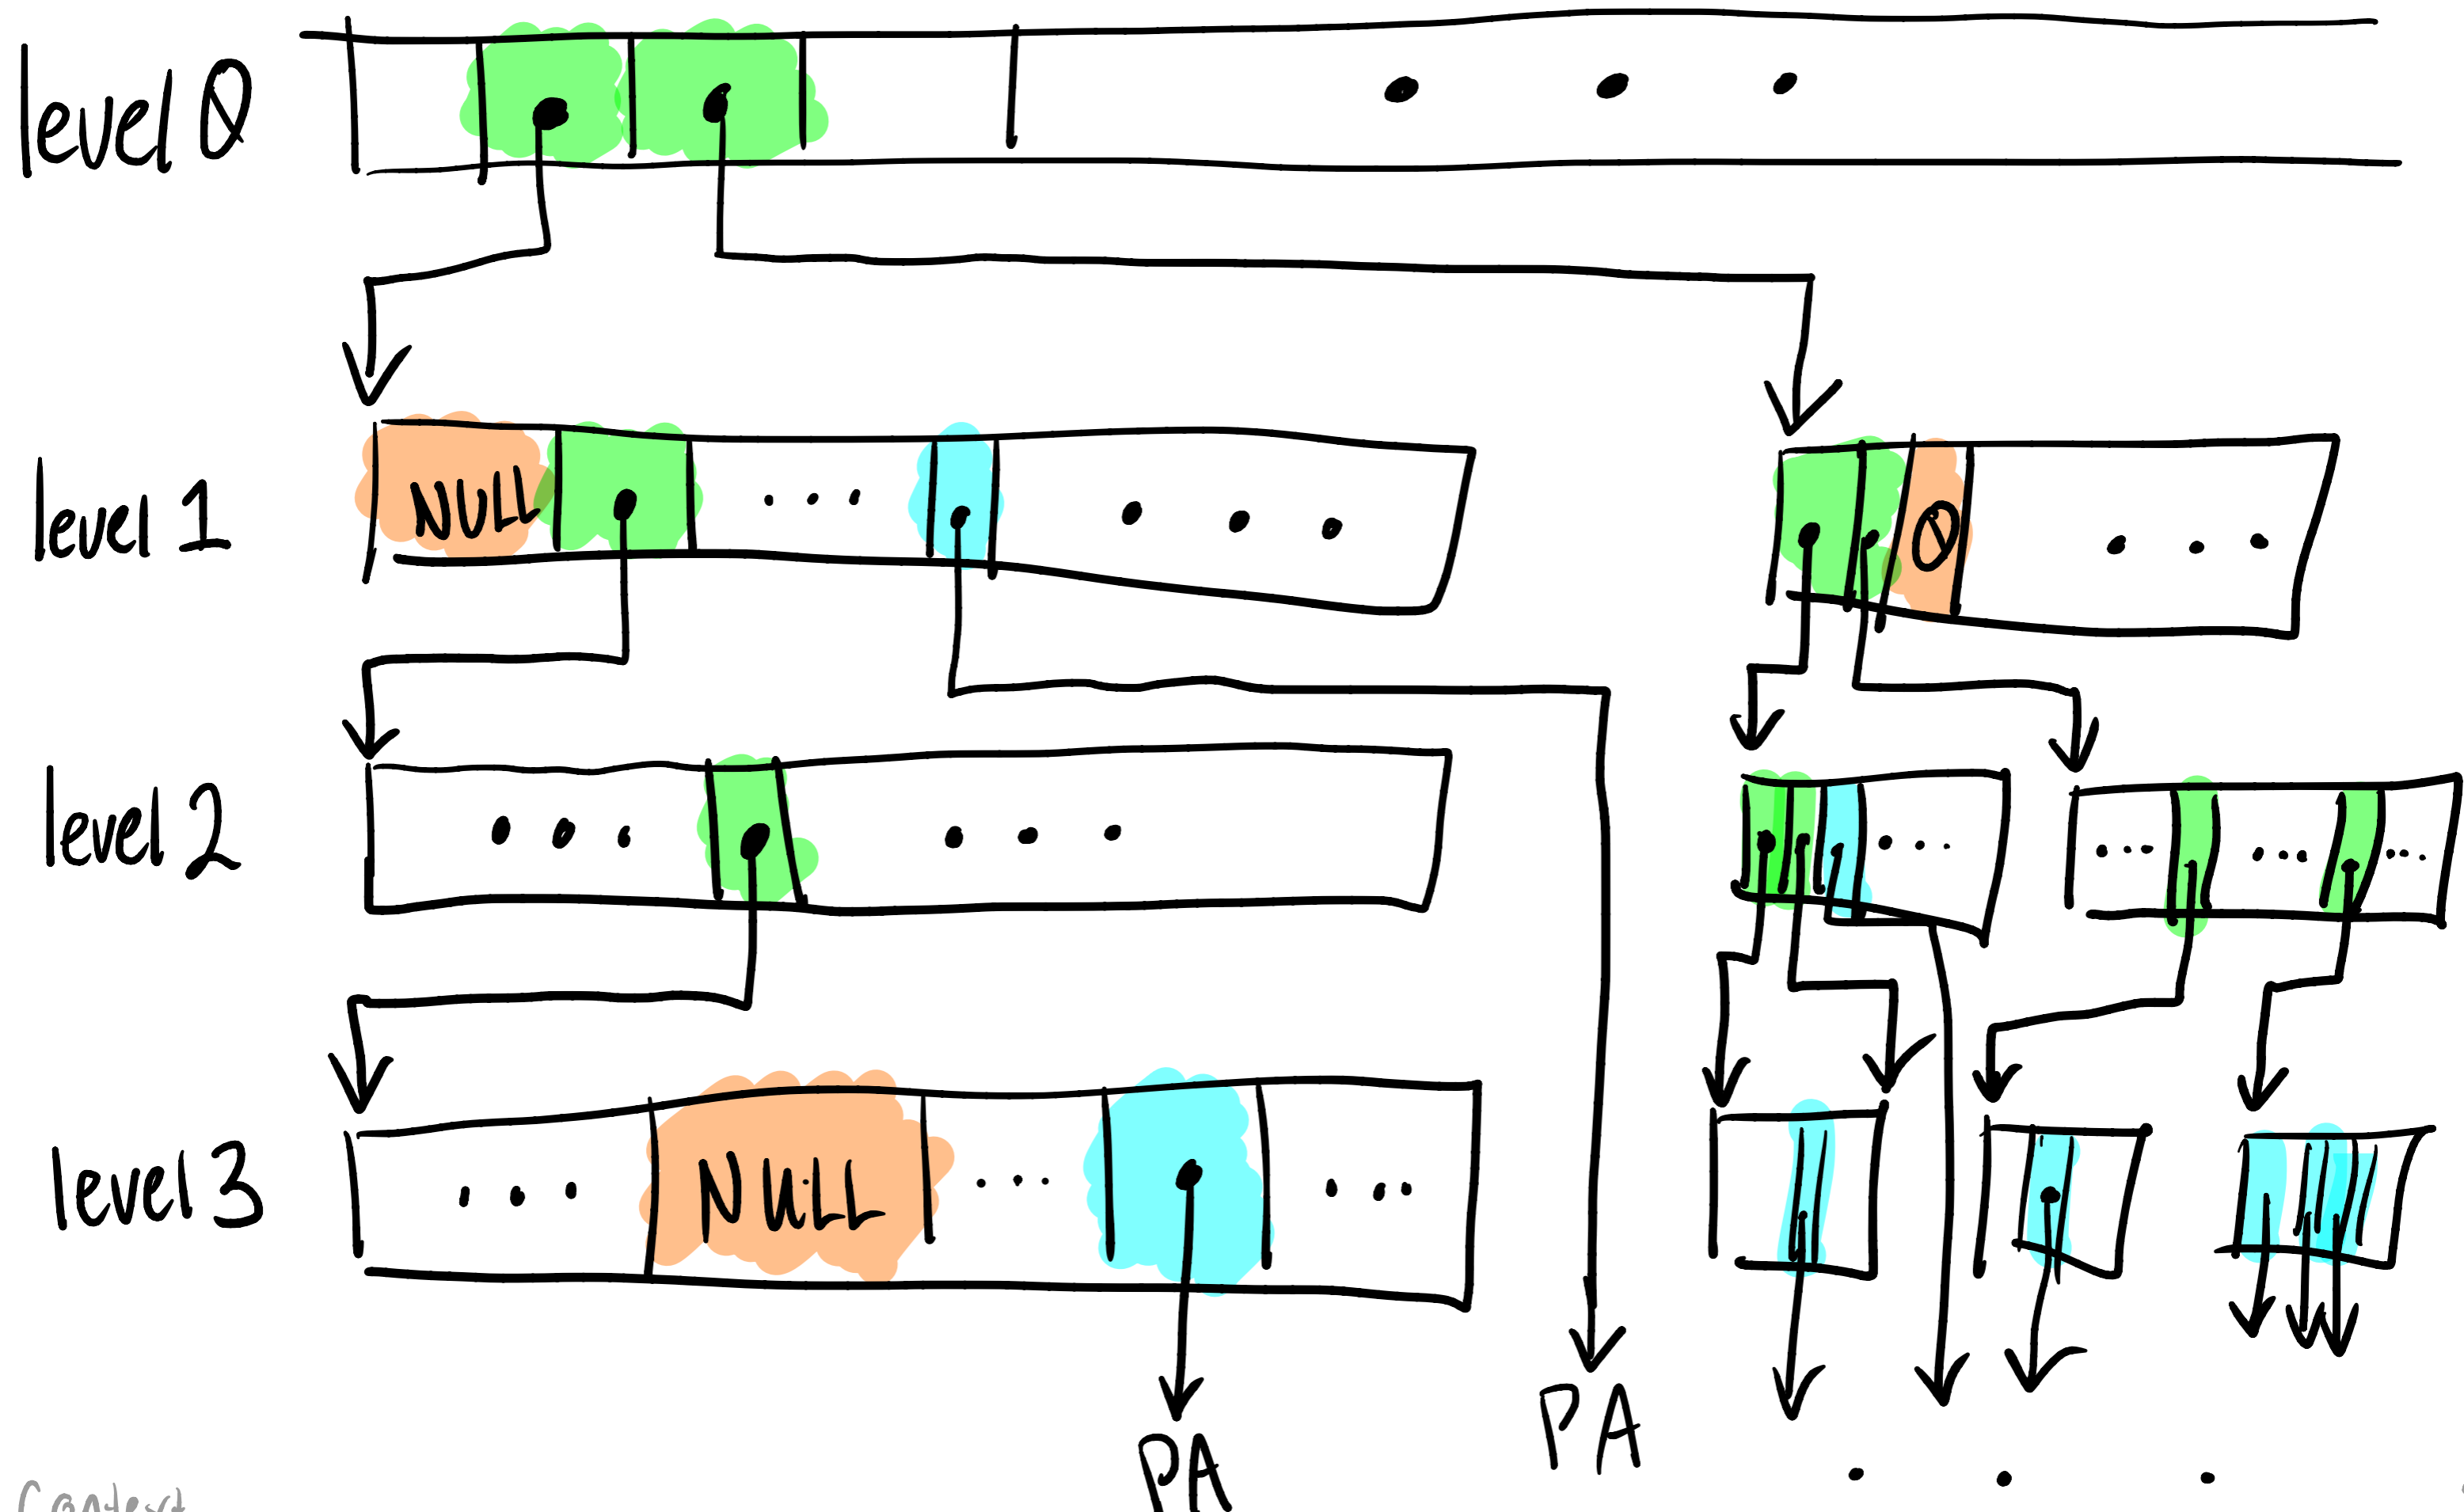





Aims

- Define Formal Semantics
 - Build tooling
 - Reason about Systems
- this work
- 

In-memory Data Structure



LDR $\Gamma_0, [VA]$



VA


LDR $r_0, [VA]$



T_{L0}



T_{L1}



T_{L2}



T_{L3}

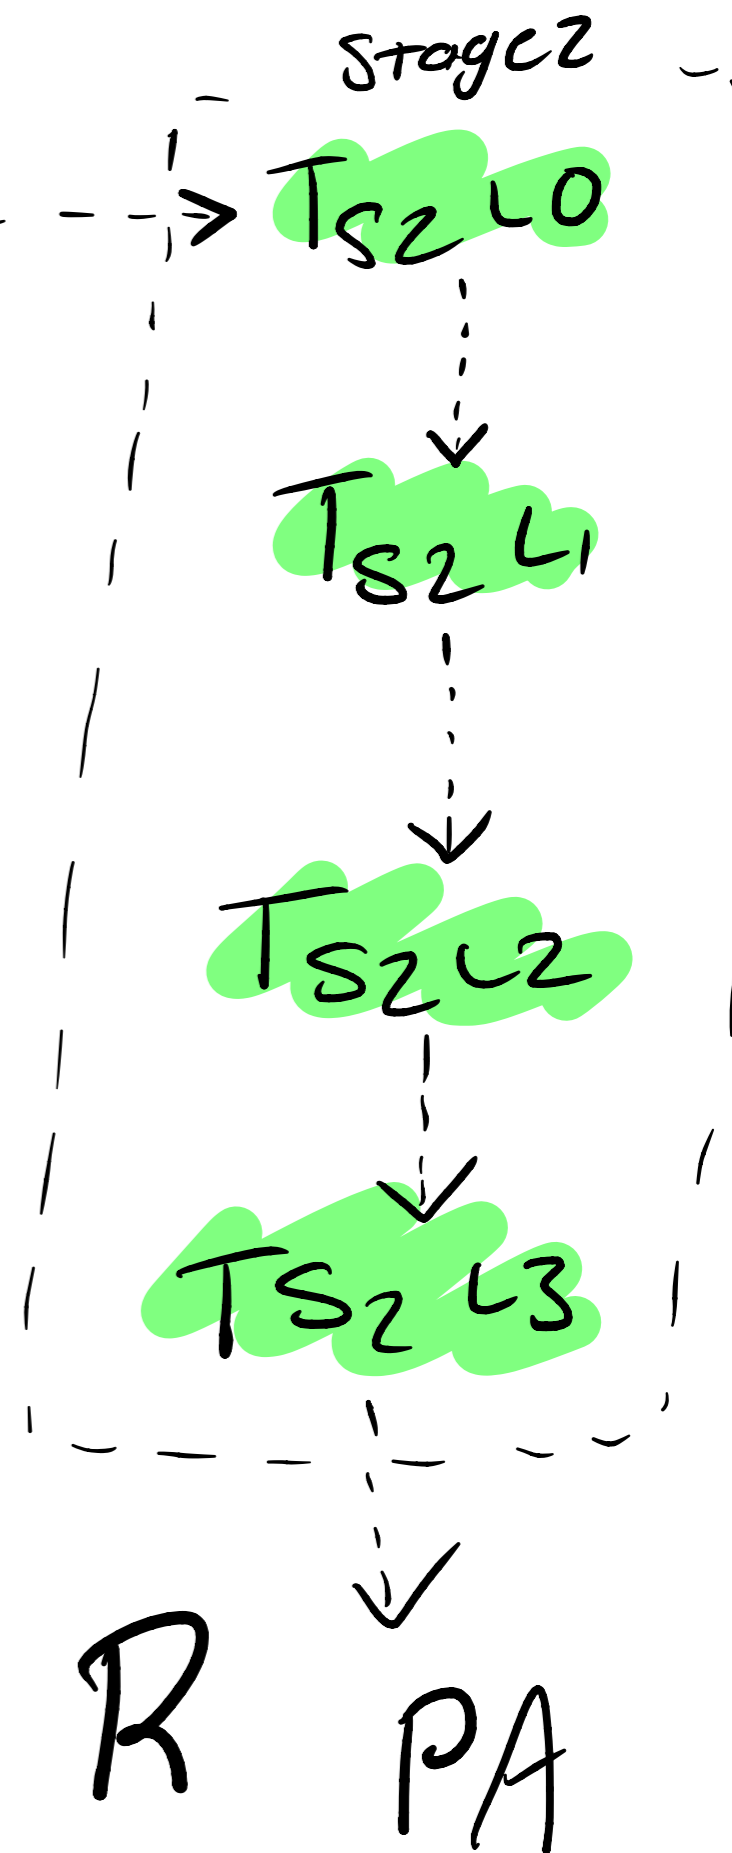



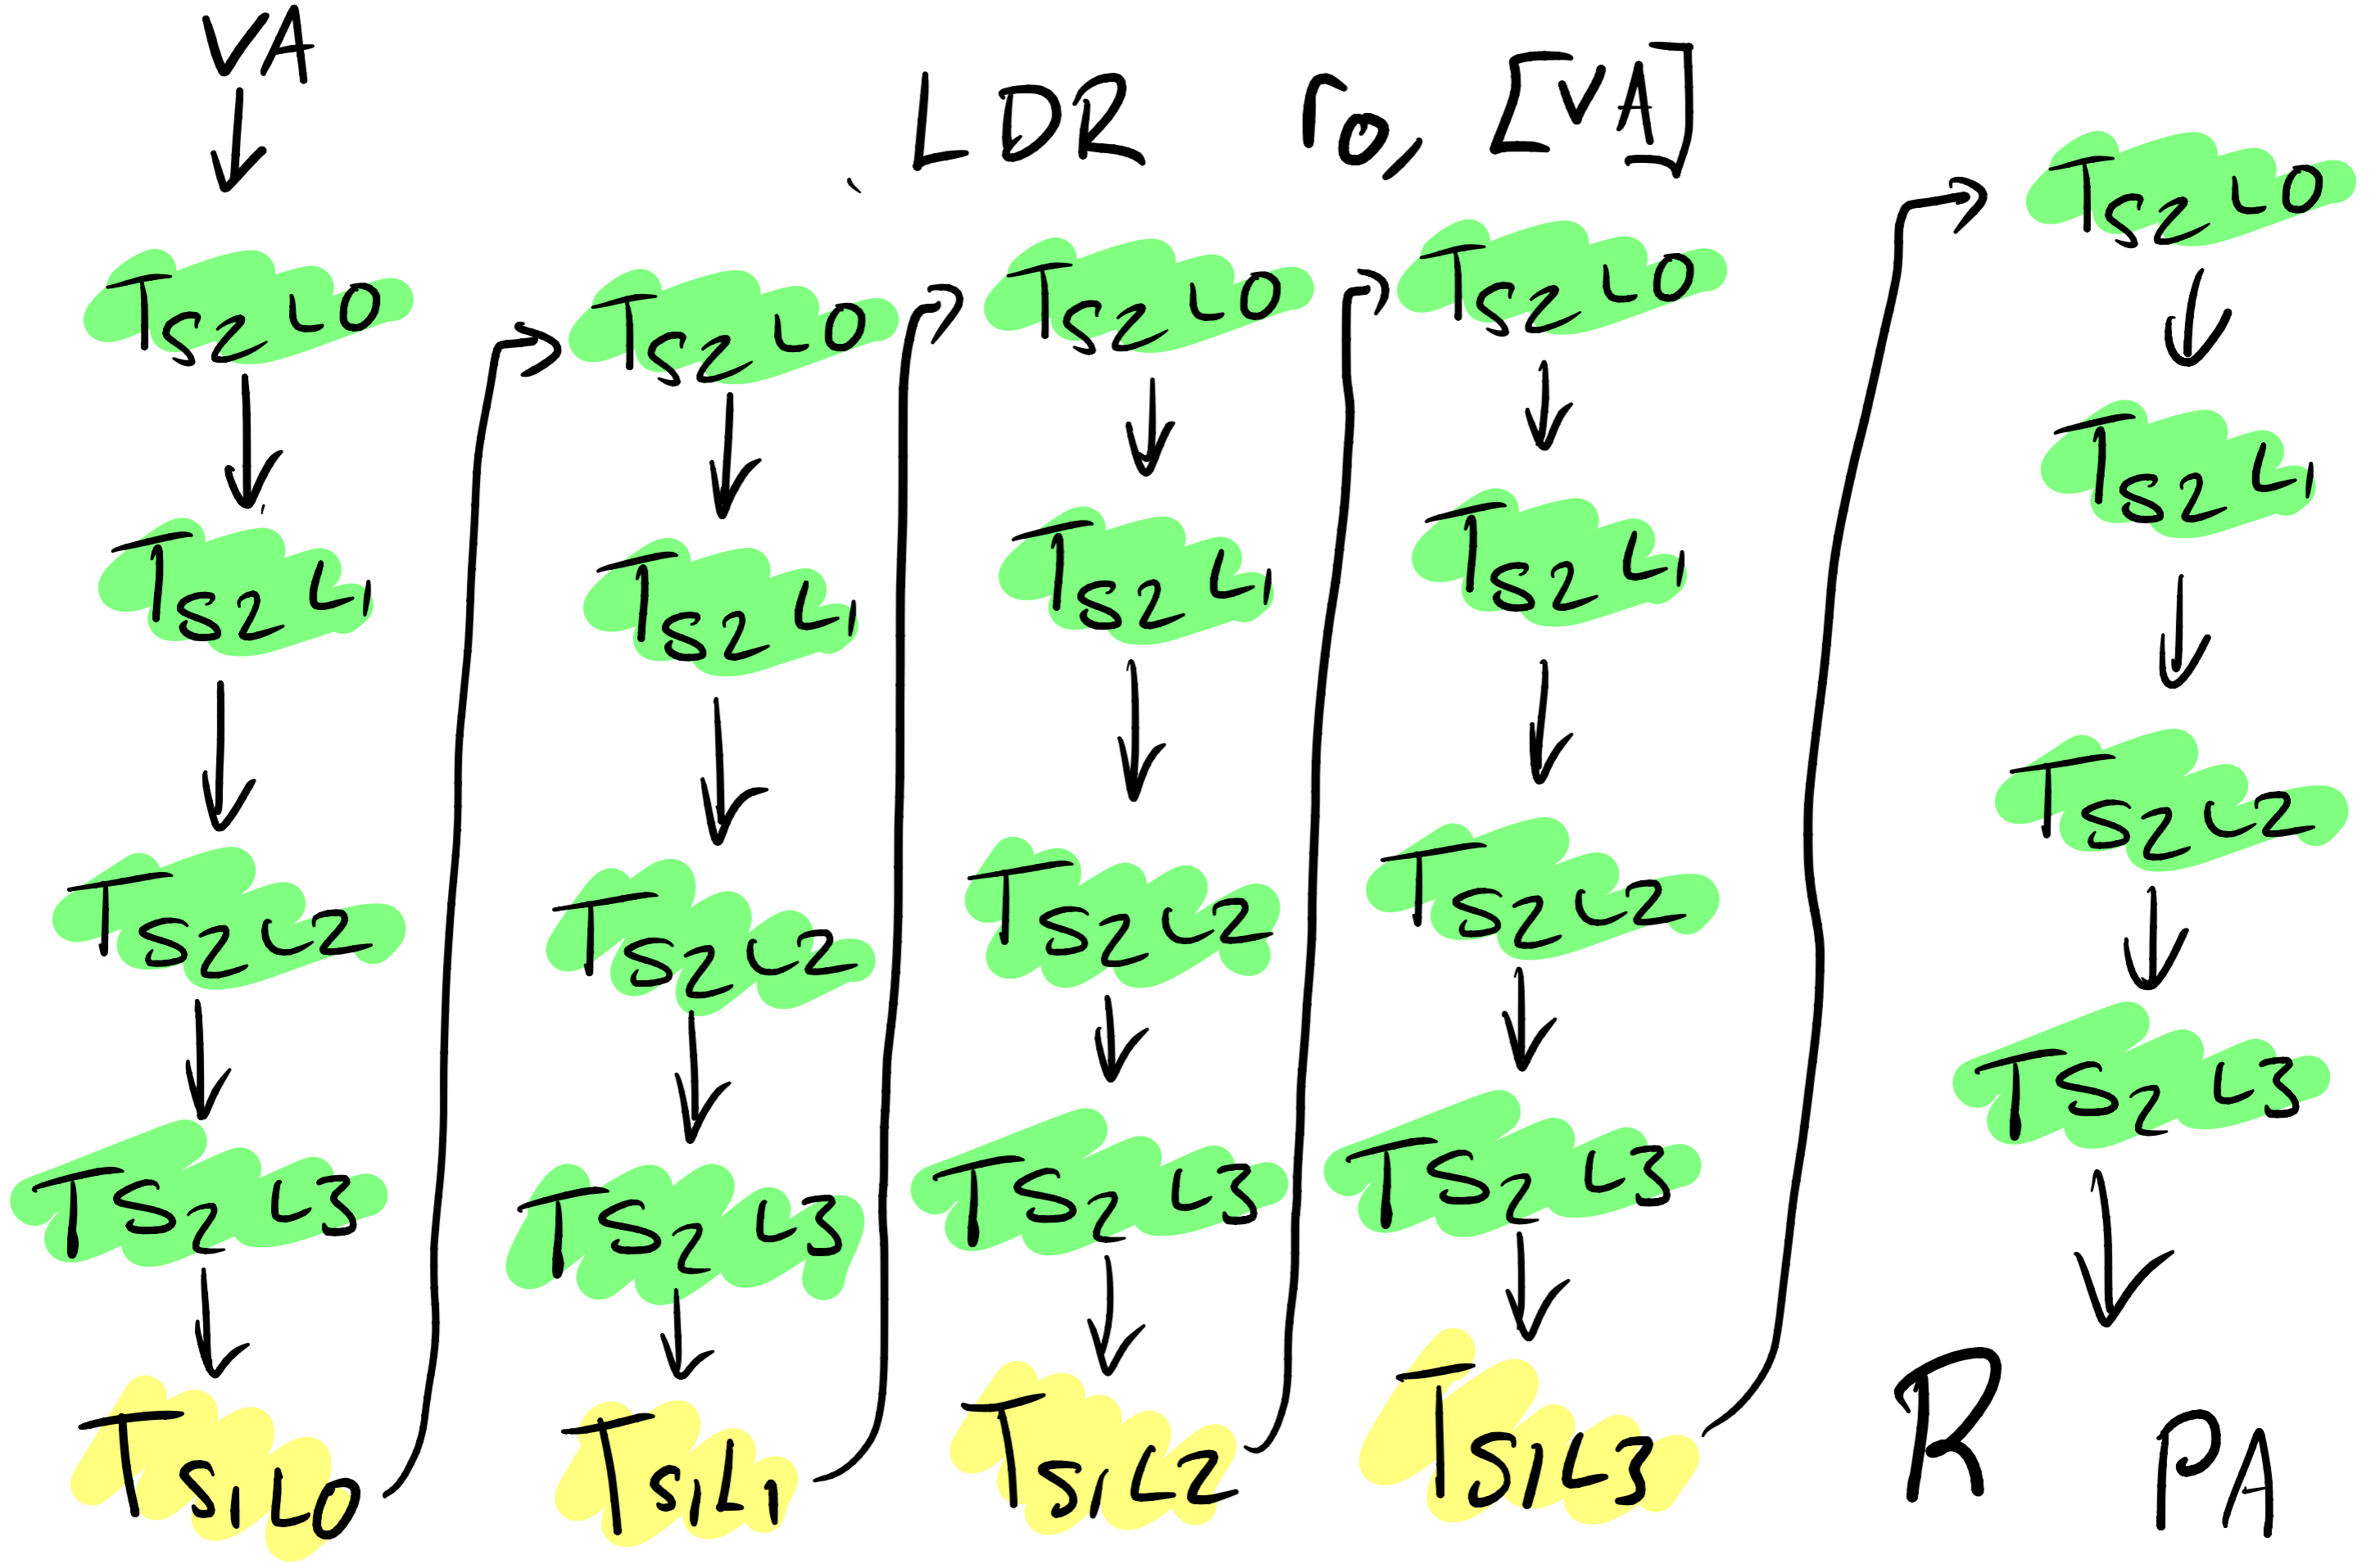
R

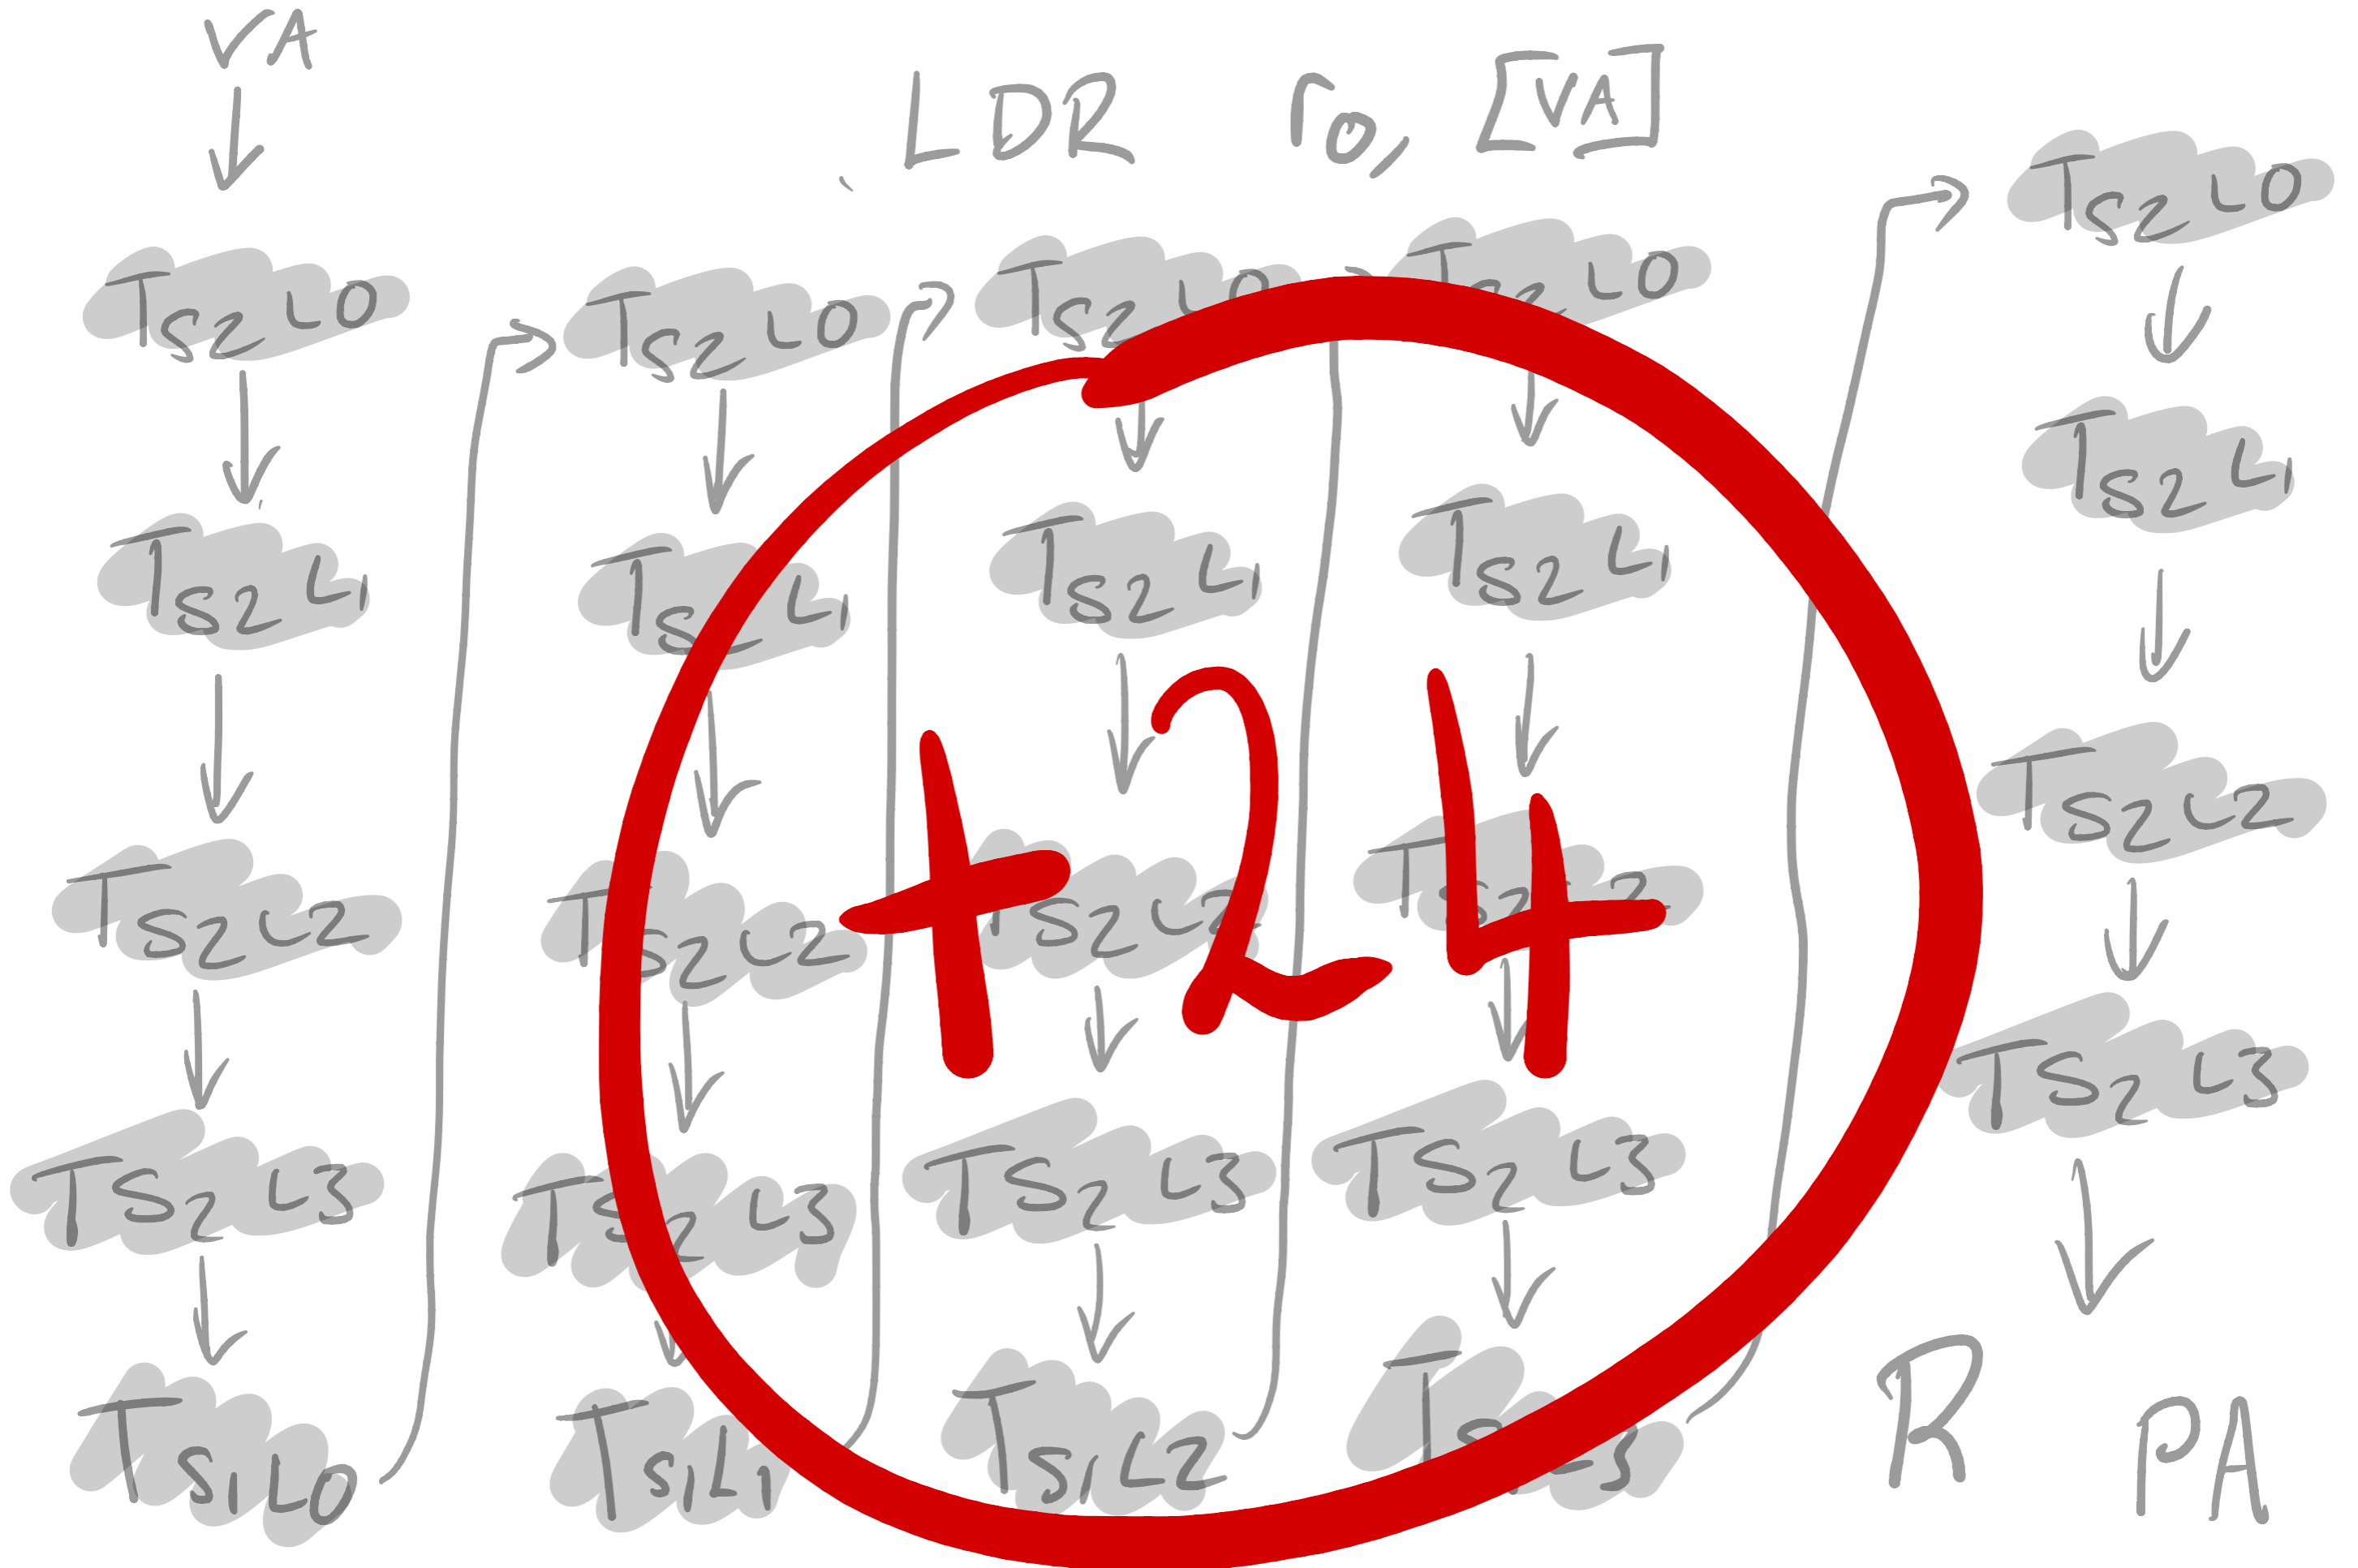
PA

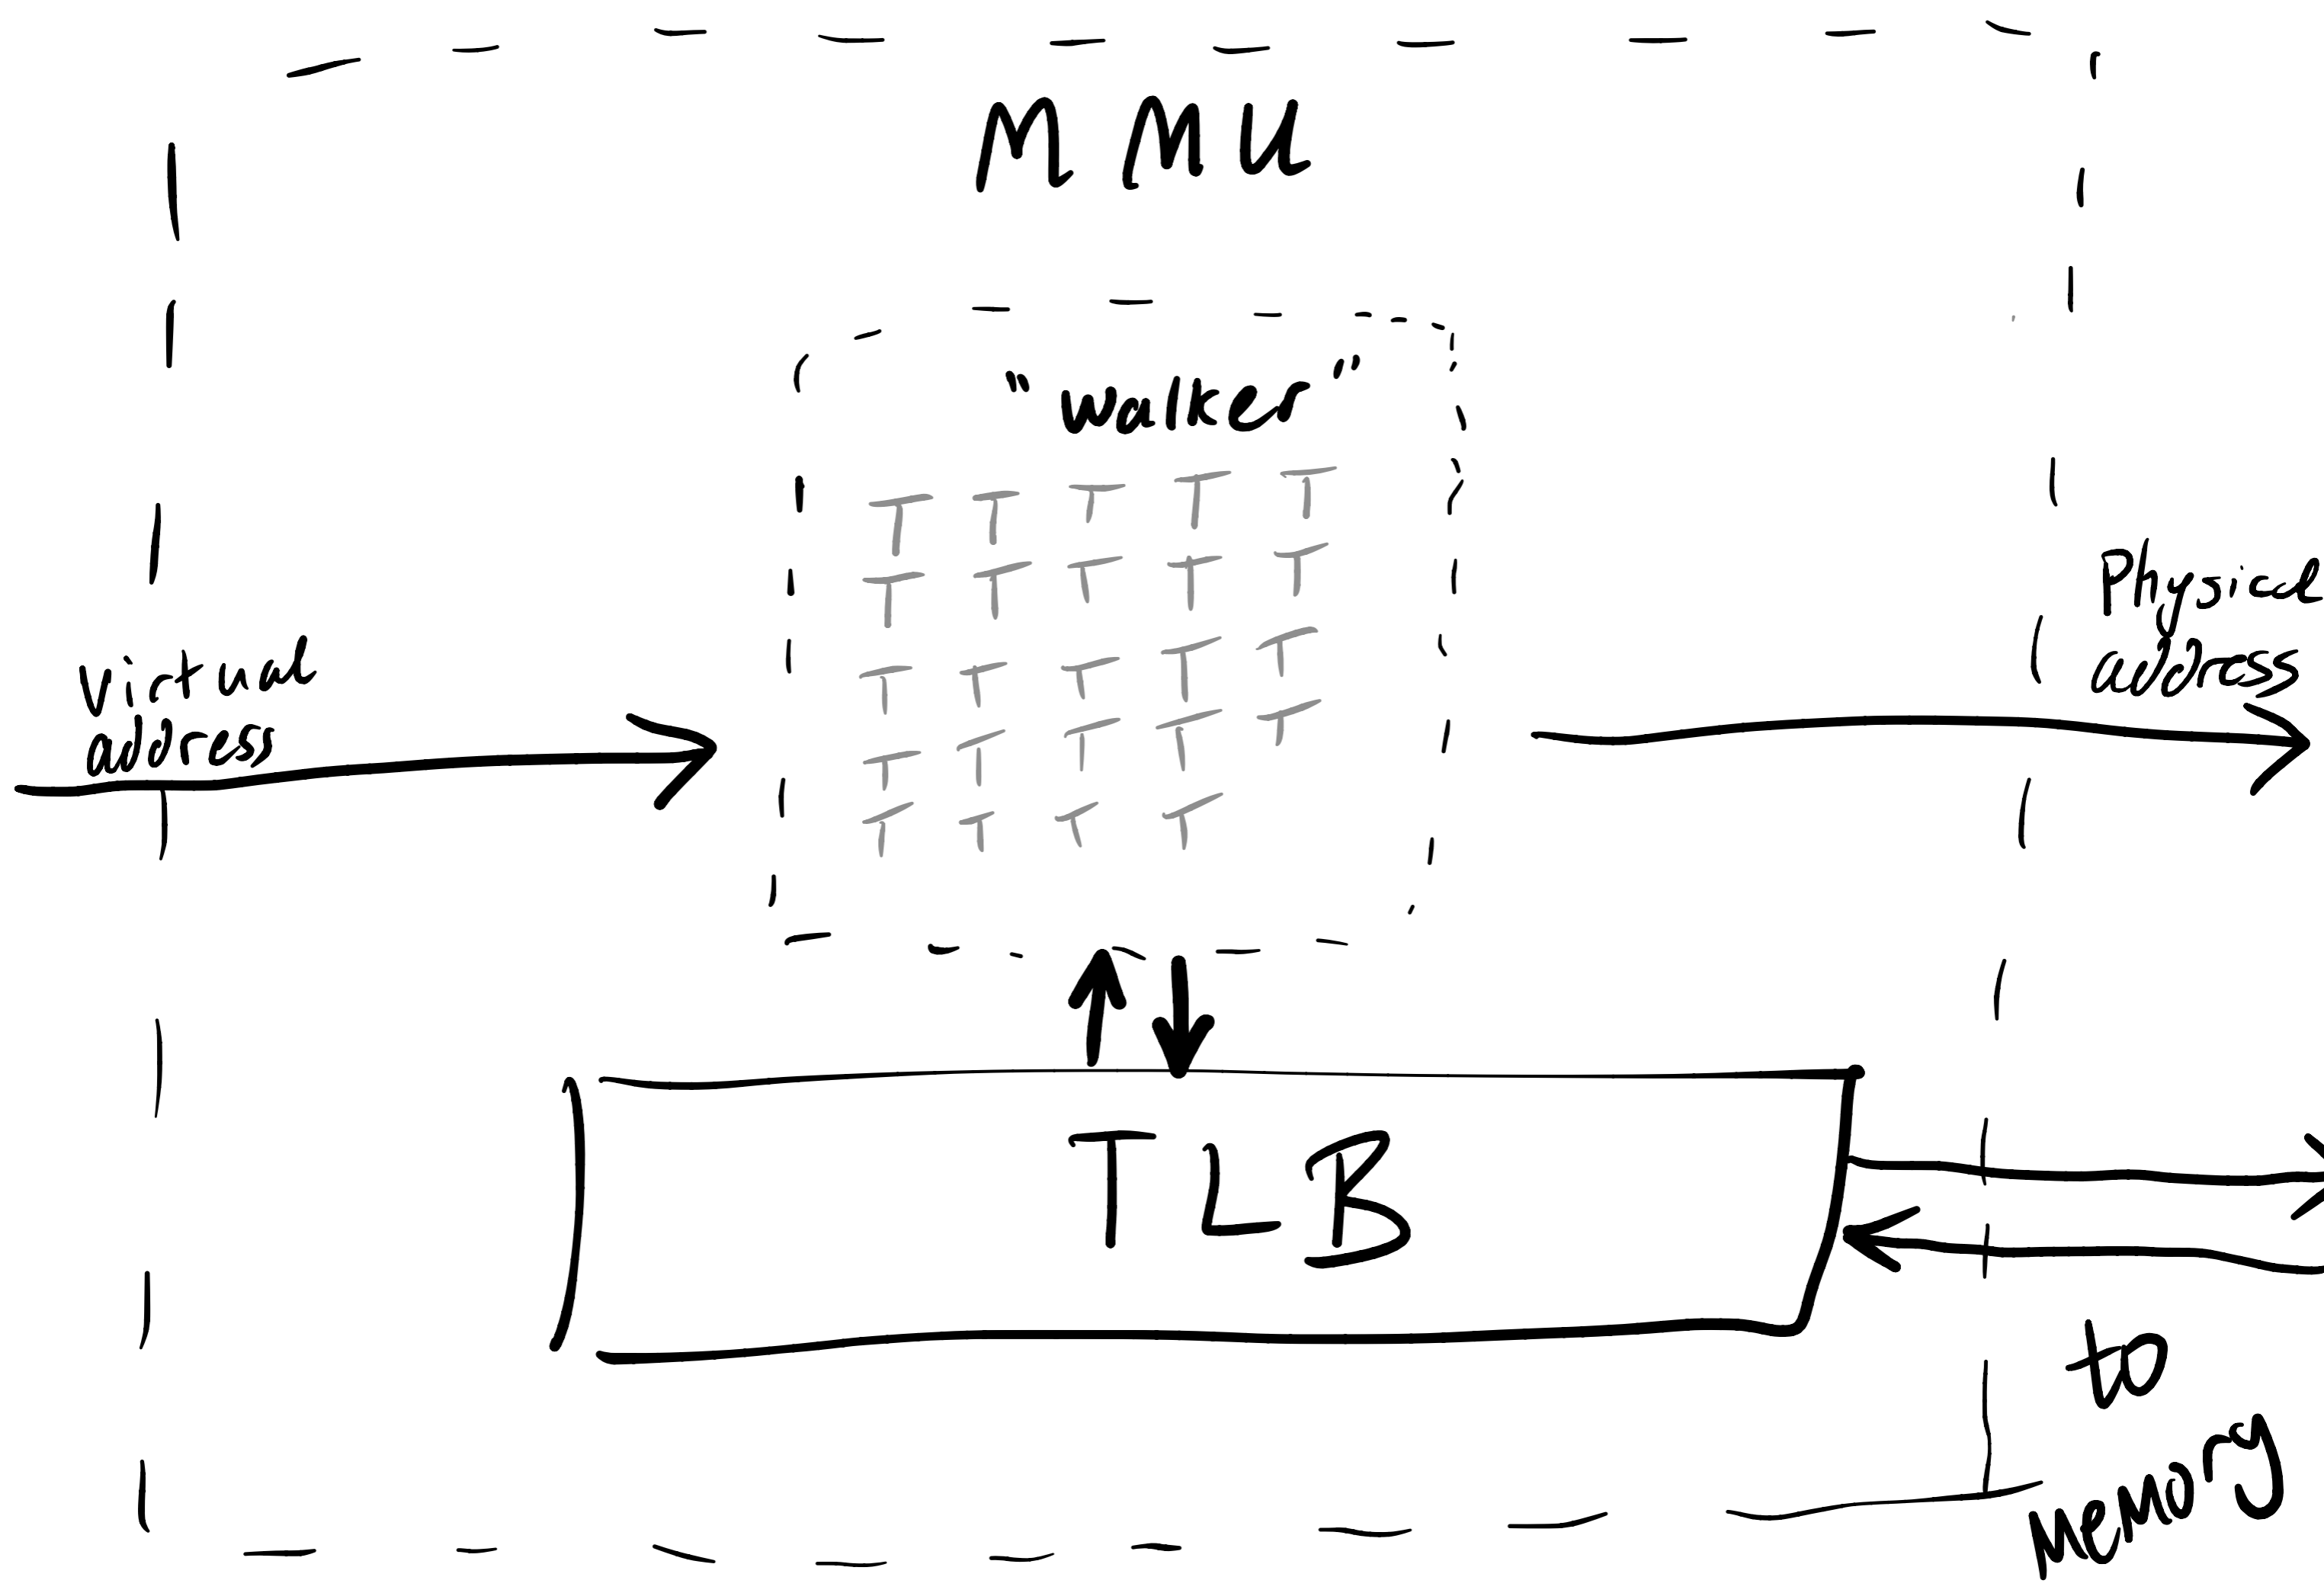
VA

LDR r0, [VA]









MMU

"walker"

T	T	T	T	T
T	T	T	T	T
T	T	T	T	T
T	T	T	T	T
T	T	T	T	T

Virtual address

Physical address

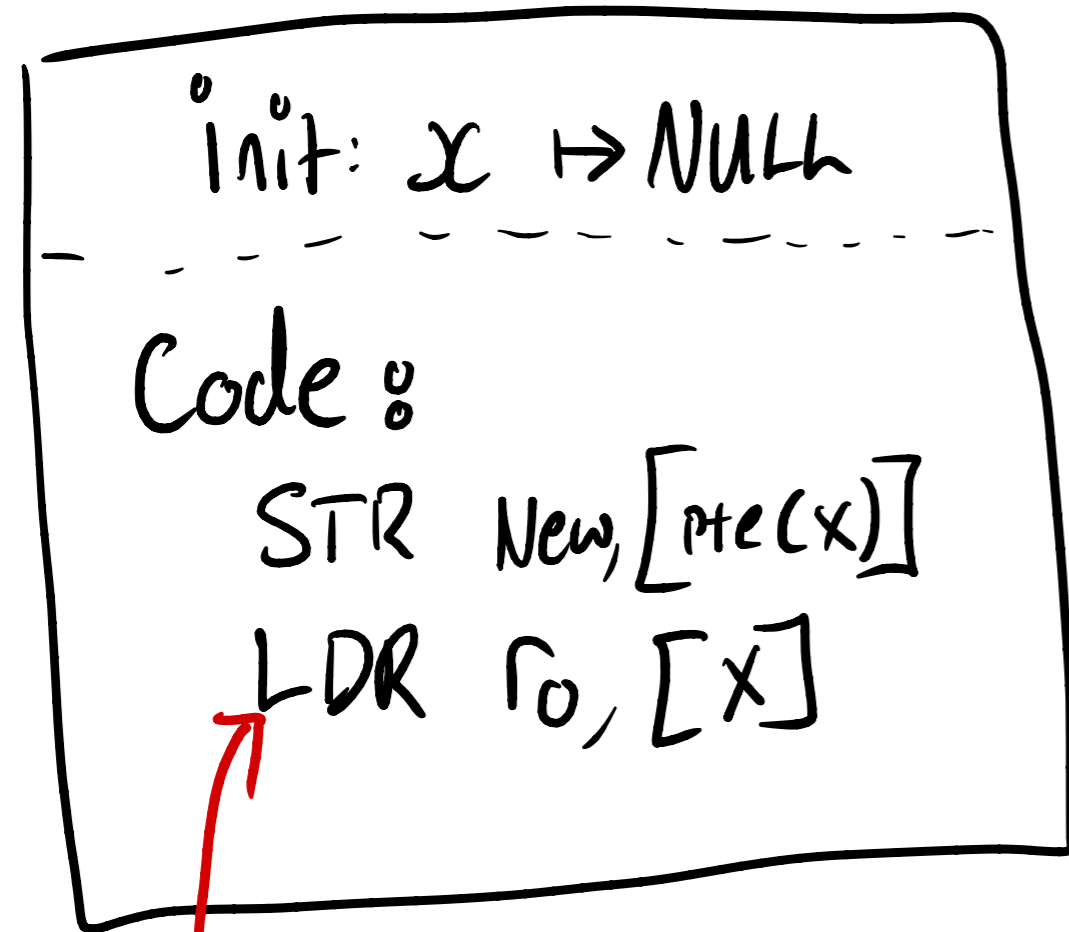
TLB

to memory

Complexity

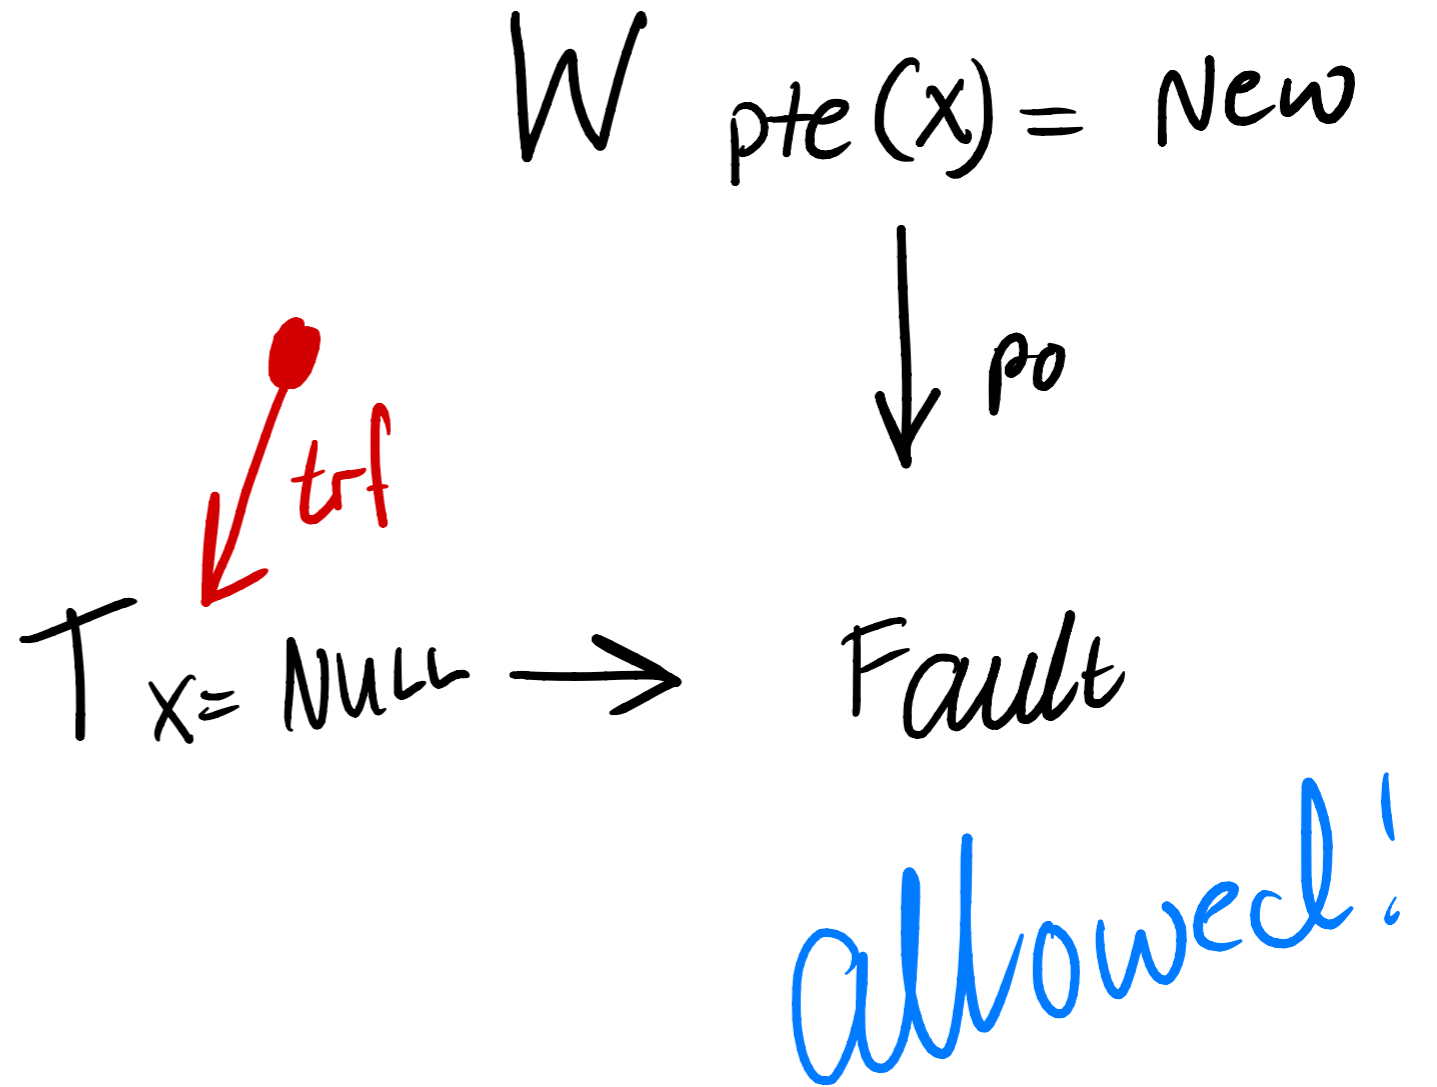
- TLB, Pipeline & cache effects
 - Concurrency
- 11,530 page
Arm Manual

Behaviour — Pipeline



Can see old.

Execution:



Behaviour — Aliasing

init: $x \mapsto P, y \mapsto P,$
 $P = 0$

code:

STR 1, [x]
LDR r0, [y]

Can't see old

$T_{x=P} \rightarrow W_{P=1}$

$\downarrow P_0$

$T_{y=P} \rightarrow R_{P=0}$

(forbidden)

Behaviour — TLB

```
init: x ↦ p
-----
Code:
STR NULL, [pte(x)]
DSB SY
ISB
LDR r0, [x]
```

Thread 0

W pte(x) = NULL

↓ flush-pipe

T_x = p → R old

allowed!

Behaviour → Break-before-make

Thread 0

Thread 1

W pte(x) = NULL



DSB sy



TLBI



DSB sy



W pte(x) = New

$T_x \rightarrow R$

... Lots of Interactions

- Spontaneous walks
- Coherent TLB fills
- Order within a translation table walk
- Speculative translations.
- Micro-TLBs
- Multi-copy Atomicity
- ETS
- Break-before-make
- Write-forwarding
- ASIDs and VMIDs

Axiomatic Model

TLB

```

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] ; tfr
    | ([T_f] ; tfr) & (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 =

```

Walk
+
Speculation

```

  | [T] ; instruction-order
  (* translation-ordered-before *)
  let tob =
    [T_f] ; tfr
    | ([T_f] ; tfr) & (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 =

```

MCA

```

  | [T] ; instruction-order
  (* translation-ordered-before *)
  let tob =
    [T_f] ; tfr
    | ([T_f] ; tfr) & (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 =

```

TLB
maintenance

```

  | [T] ; instruction-order
  (* translation-ordered-before *)
  let tob =
    [T_f] ; tfr
    | ([T_f] ; tfr) & (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 =

```

Pipeline

```

  | [T] ; instruction-order
  (* translation-ordered-before *)
  let tob =
    [T_f] ; tfr
    | ([T_f] ; tfr) & (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 =

```

Conttib - Model

```

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

```

FEAT-ETS

AXIOMS

isla - axiomatic

The screenshot displays the isla-axiomatic web interface. On the left is a code editor with the following content:

```
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"
```

On the right is the EventGraph, showing the execution flow between two threads:

- Initial State** (oval)
- Thread 0** (dashed box):
 - a:** str x0, [x1]: W s1:l3pte(x) = s1:l3desc(y)
- Thread 1** (dashed box):
 - a1:** T s2:l3pte(x) (cyan box)
 - a2:** ldr x2, [x1]: R pa1 = 0x1
 - b1:** T s1:l3pte(x) (cyan box)
 - b2:** ldr x2, [x3]: Fault
 - c:** eret

Relationships between events:

- Thread 0 event **a** has a **trf** (transfer fault) relationship with Thread 1 event **a1**.
- Thread 1 event **a1** has an **iio** (input/output) relationship with Thread 1 event **a2**.
- Thread 1 event **a2** has an **iio** relationship with Thread 1 event **b1**.
- Thread 1 event **b1** has an **iio** relationship with Thread 1 event **b2**.
- Thread 1 event **b2** leads to event **c**.

Recap

- ① Clarified Architecture
- ② Defined Axiomatic model
- ③ Built tooling

Future

USE the model !

more ARM features

... more architectures?