

Challenges in Empirically Testing Memory Persistency Models


V. KLIMIS, A. RAAD, V. VAFEIADIS, J. WICKERSON, A. F. DONALDSON


Why Bother?

Ensure Hardware Vendor Accountability

Ascertain the Soundness and Strength
of Persistency Semantics

Memory Persistency: The Sneaky Game of Hide and Seek



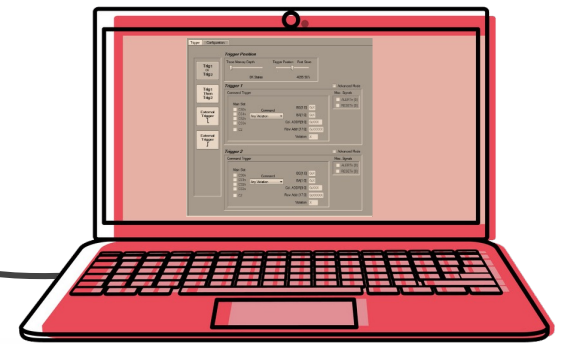


NVRAM



DDR Detective

Probe Manager



Litmus Test

```

x := 1;      | a := y;
wb x;       | DMBfull;
DSBfull;  | if (a)
y := 1;     |   z := 1;

```


rec: (y=1 ∨ z=1) ⇒ x=1




possible values of x, y, z upon recovery

Persistence Domain

Intel-x86



Arm




Persistence Domain

We currently do not support building fine grained persistent memory platform solutions with the Ampere Altra SOCs. [...] That given, the Ampere Altra SOC does not have a Point of Persistence.

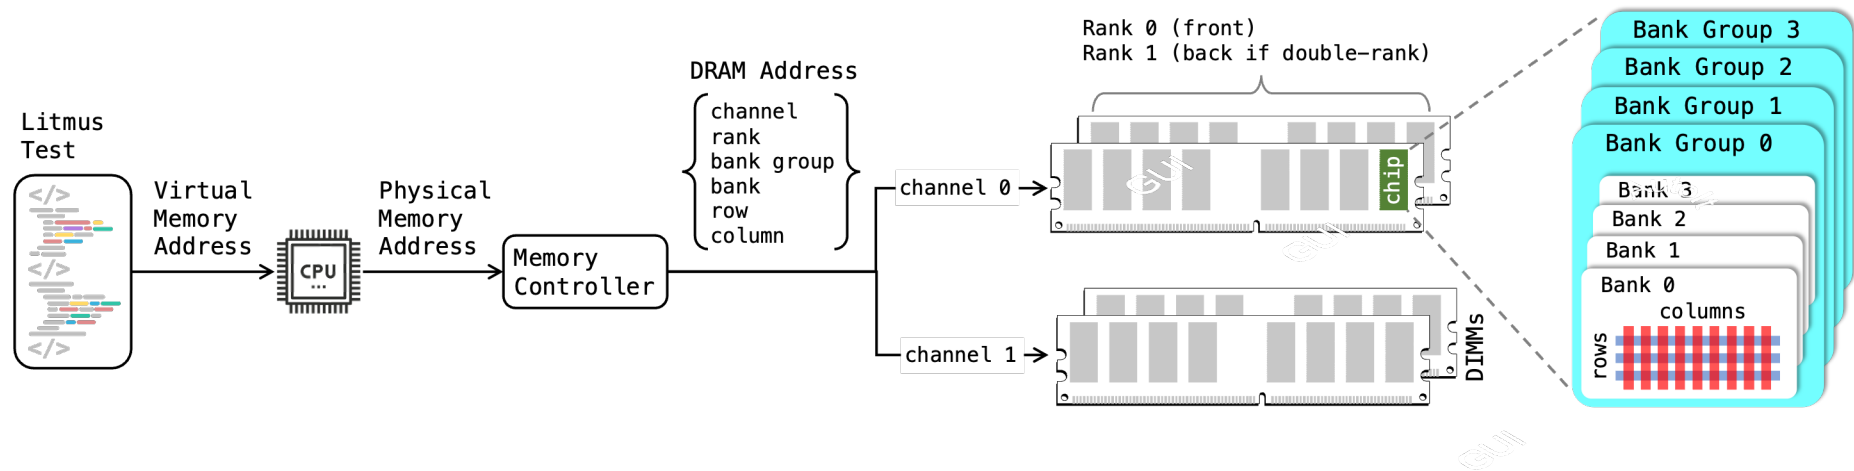
If the memory system does not support the Point of Persistence, a data cache clean to the PoP, DC_CVAP, behaves as a data cache clean to the PoC, DC_CVAC

Arm




Challenges Posed by the DDR Detective

1. Virtual → Geometric Address



Challenges Posed by the DDR Detective

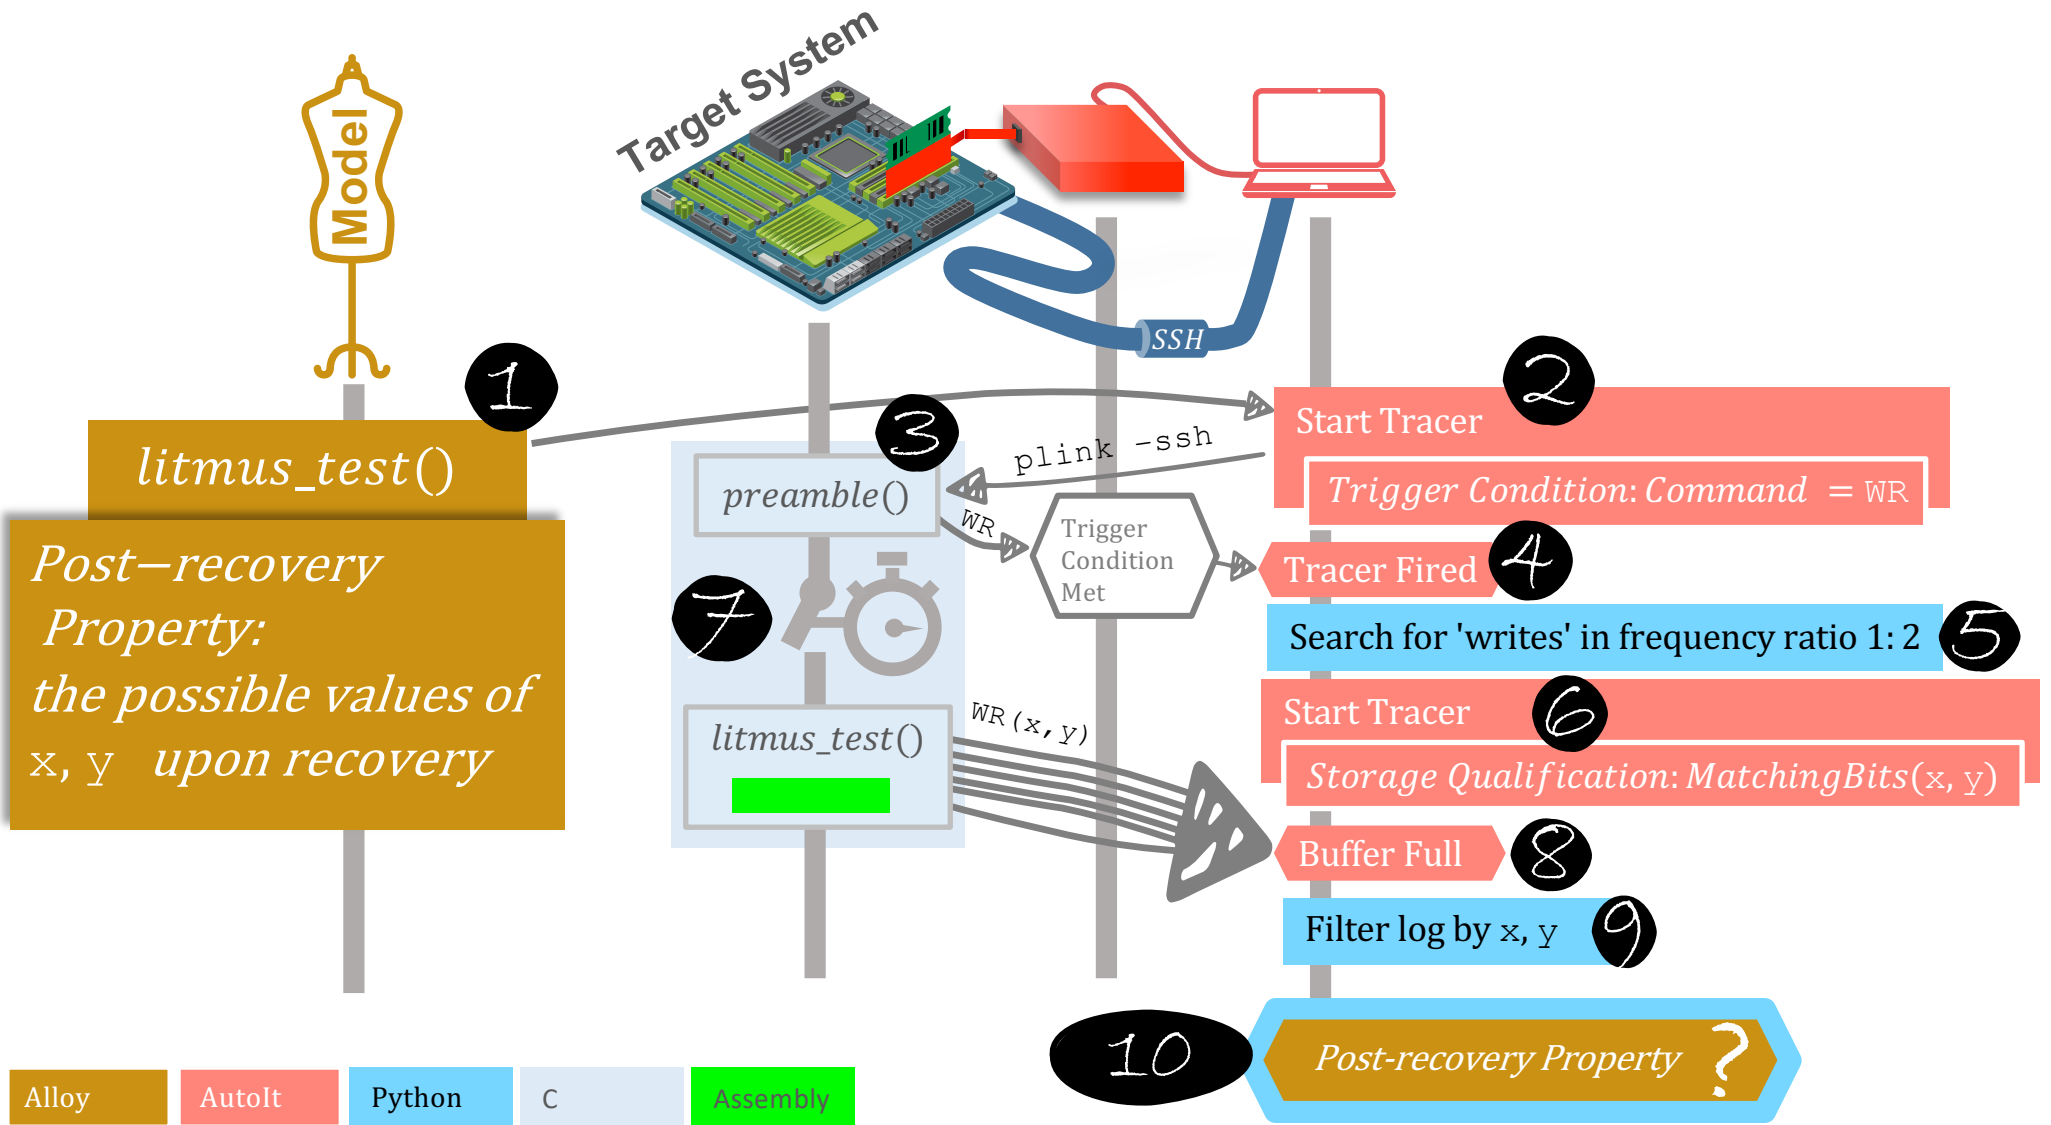
1. Virtual → Geometric Address
2. Address Logging, Data Omission



Challenges Posed by the DDR Detective

1. Virtual → Geometric Address
2. Address Logging, Data Omission
3. Automating DDR Detective: Lack of Command-Line Support





A Litmus Test


```
1: x ← posix_memalign(L1 Cache ≤ size ≤ L1 Cache, CACHE_LINE_SIZE)
2: y ← posix_memalign(L1 Cache ≤ size ≤ L2 Cache, CACHE_LINE_SIZE)
3: z ← posix_memalign(L1 Cache ≤ size ≤ L1 Cache, CACHE_LINE_SIZE)
4: reg ← 0
5: while time < 5 do
6:   x ← 1; dc_cvap(x); dsb(sy);
7:   y ← 1; dc_cvap(y); dsb(sy);
8:   y ← 1; dc_cvap(y); dsb(sy);
9:   z ← 1; dc_cvap(z); dsb(sy);
10:  z ← 1; dc_cvap(z); dsb(sy);
11:  z ← 1; dc_cvap(z); dsb(sy);
12: end
13: sleep(20)
14: while time < 60 do
15:   ++reg
16:   // Thread 1 || // Thread 2
17:   x ← reg;    || a ← y //thread-local register
18:   dc_cvap(x); || dmb(sy);
19:   dsb(sy);
20:
21:   y ← reg;    || if(a)
22:   dc_cvap(y); ||   | z ← reg;
23:   dsb(sy);    ||   | dc_cvap(z);
                ||   | dsb(sy);
24:   // persistency property
25:   //  $y_{pd} = reg \vee z_{pd} = reg \Rightarrow x_{pd} = reg$ 
26: end
```

preamble

litmus test

Curiosity-driven experiments


Does DC CVAP degenerate to DC CVAC?



No sufficient statistical evidence to indicate that DC CVAP defaults to a different behaviour than DC CVAC

Curiosity-driven experiments


Does the distance between the memory locations affect reorderings?



Larger memory chunks exhibit increased reorderings

Curiosity-driven experiments

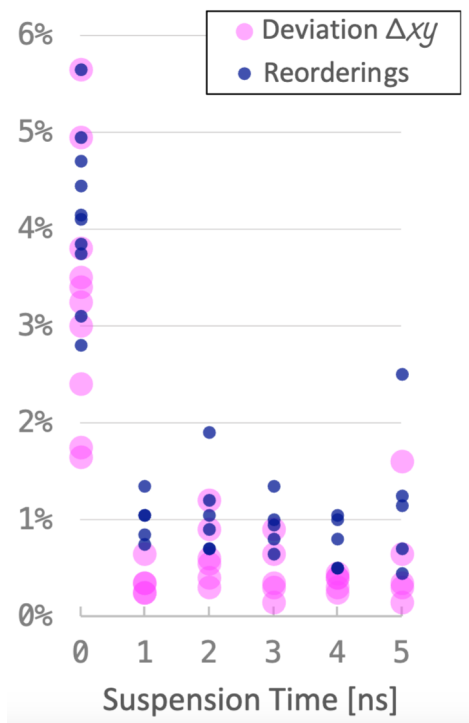
Do repeated persists inhibit reorderings?



As the persist repetitions grow, reorderings decrease

Curiosity-driven experiments

Does suspending the processor between writes inhibit reorderings?



Delays following persist in the litmus test decrease reorderings
The presence of a delay, rather than its specific duration, is crucial for this shift.

Curiosity-driven experiments

Does the value written make a difference?

No noticeable difference

Conclusions

Emphasises the significance of a Tailored Validation Approach

Showcased discrepancies between observed behaviours and vendors' specifications, underscoring the need for hands-on testing over mere reliance on documentation

A reusable methodology to validate persistency guarantees if vendors pledge new promises in their forthcoming releases



Thanks