CASPER

CS305: Computer Architecture

Caches: The Virtual World

https://www.cse.iitb.ac.in/~biswa/courses/CS305/main.html

https://www.cse.iitb.ac.in/~biswa/

Virtual World

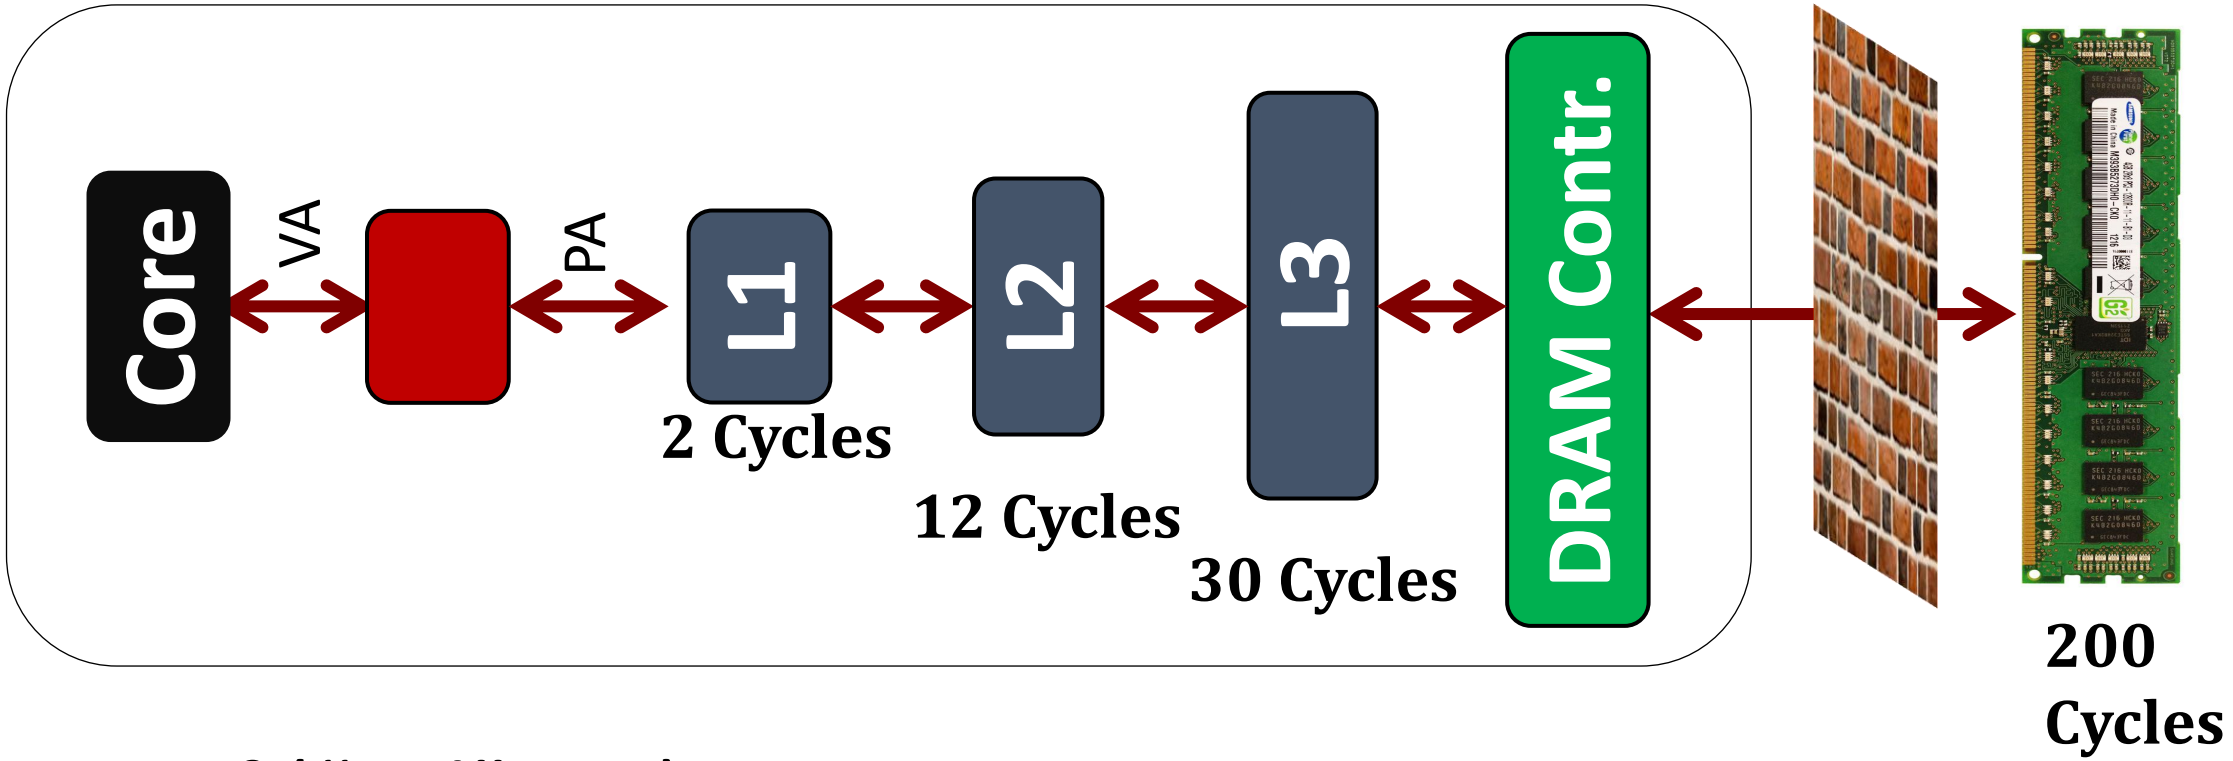

Printf (“%d”, &a);

Virtual address

Virtual Memory

App. 1

Virtual address space

Printf (“%d”, &a);

App. 2

Virtual address space

Printf (“%d”, &a);

Page Table

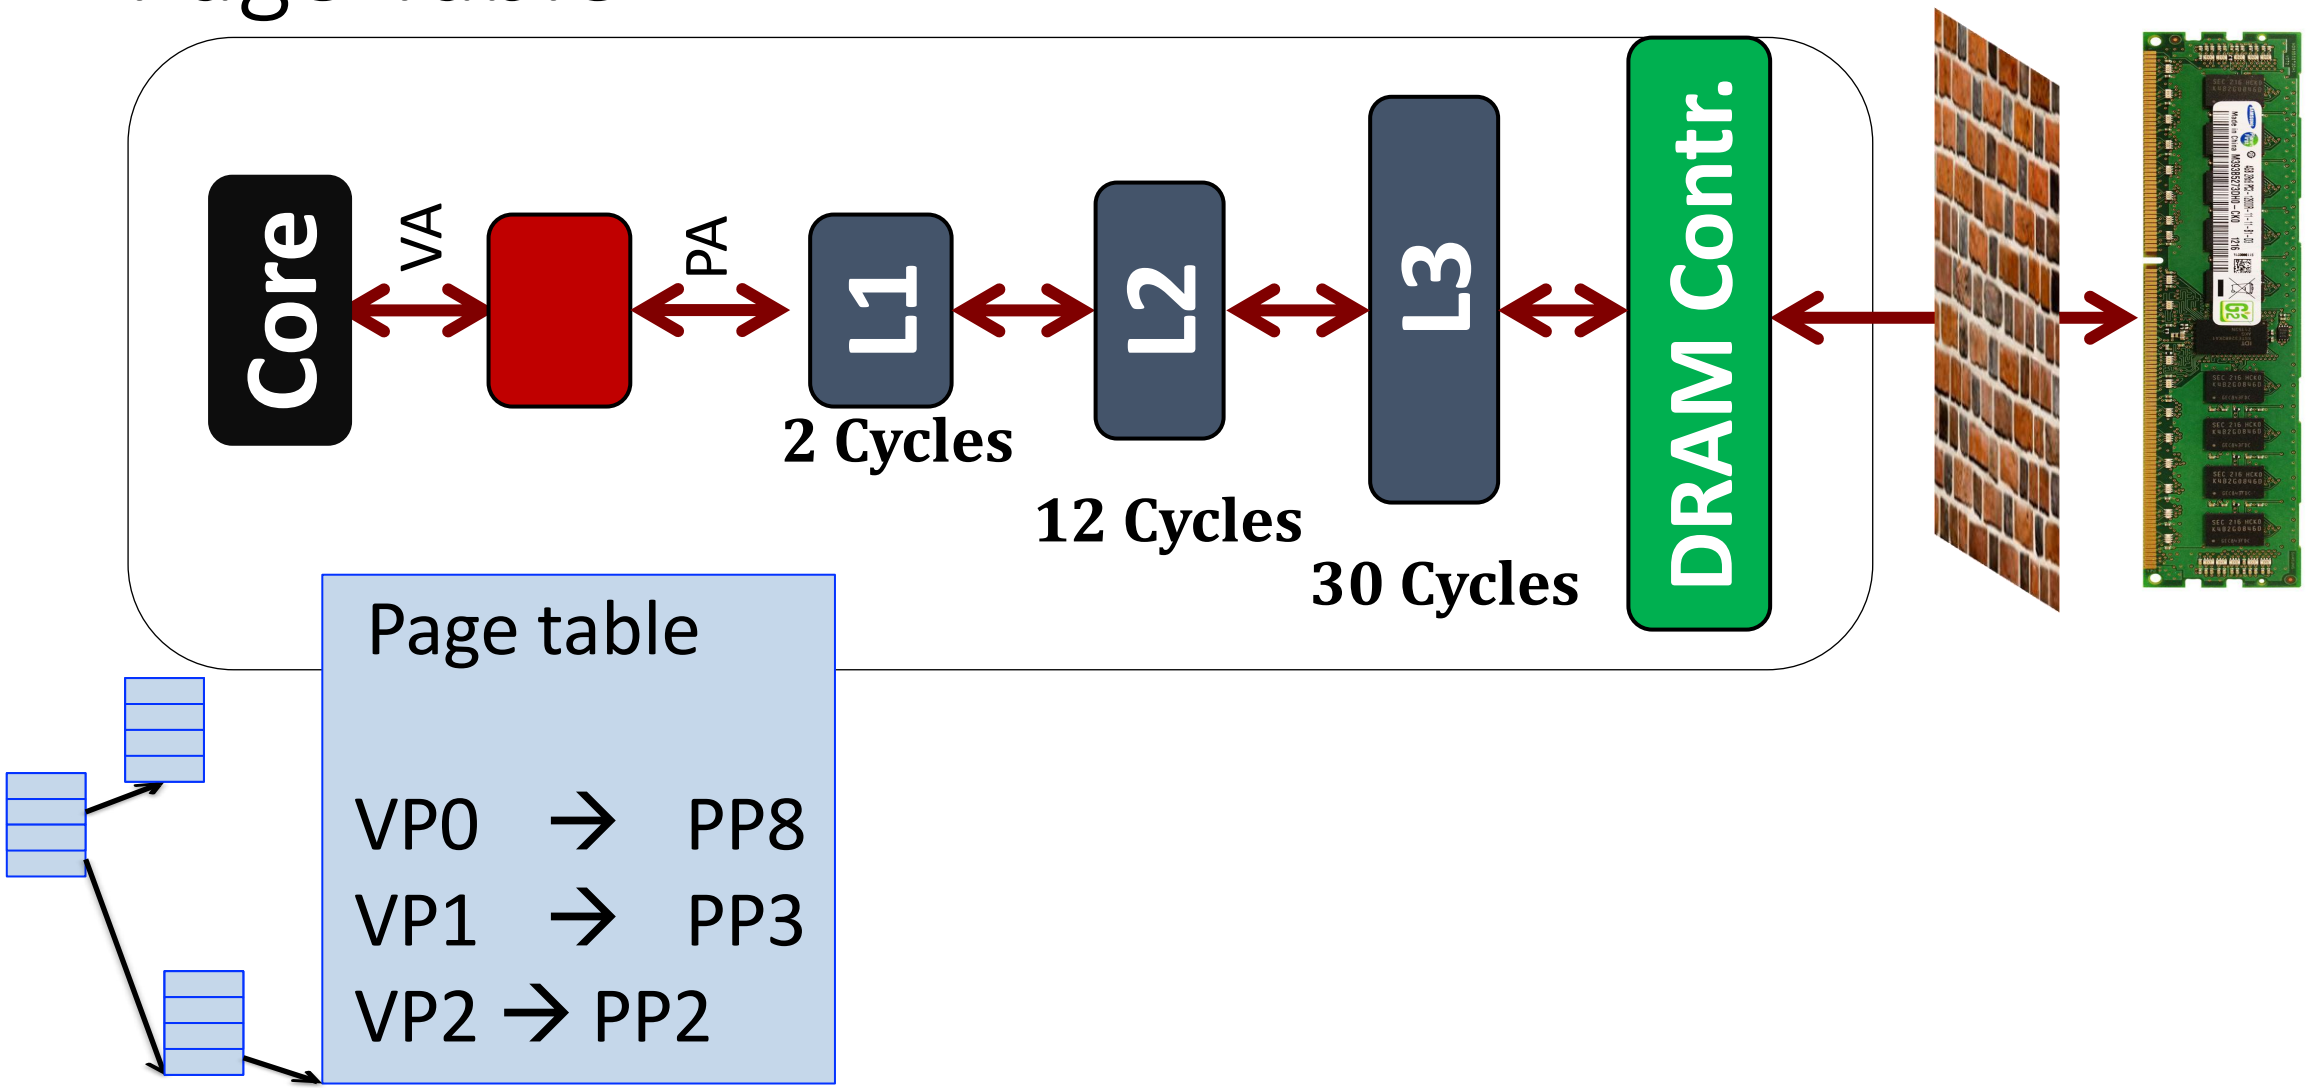
Core
VA
PA
L1
2 Cycles
L2
12 Cycles
L3
30 Cycles
DRAM Contr.
Page table
VP0 → PP8
VP1 → PP3
VP2 → PP2

Page Table Walk

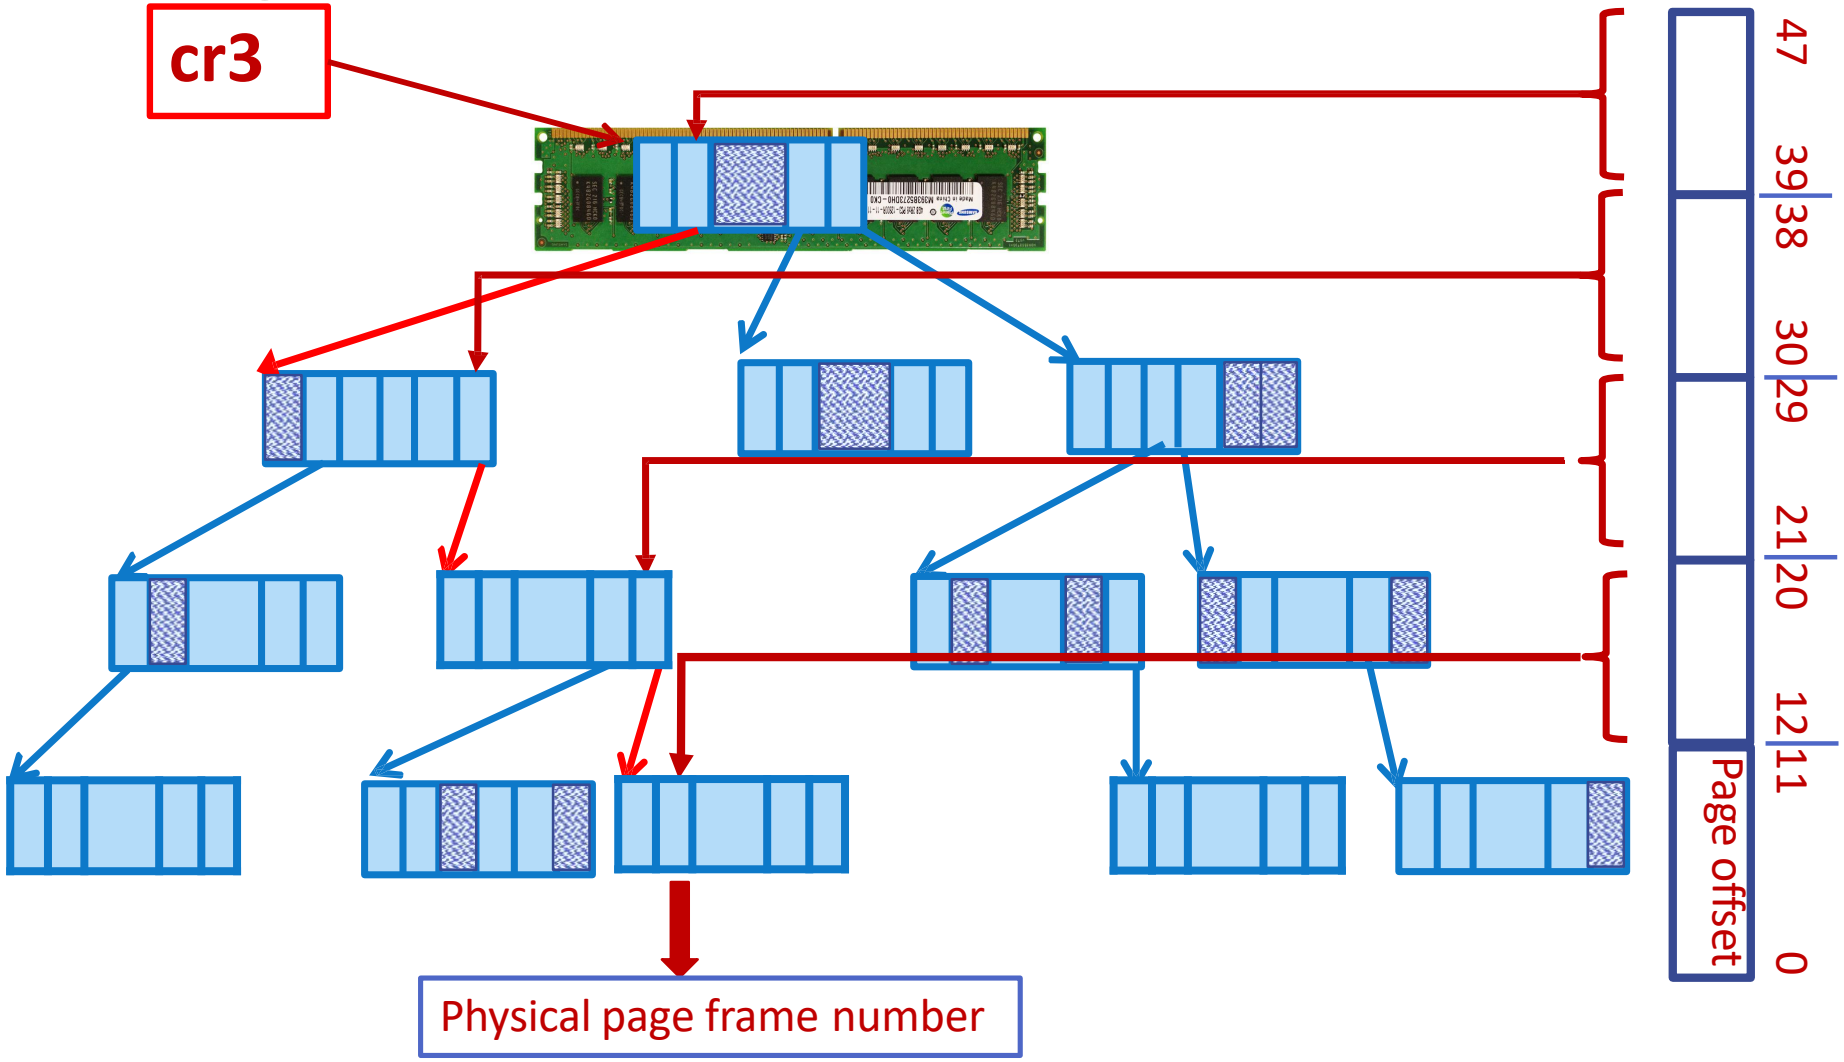

Can We Cache Translations too?

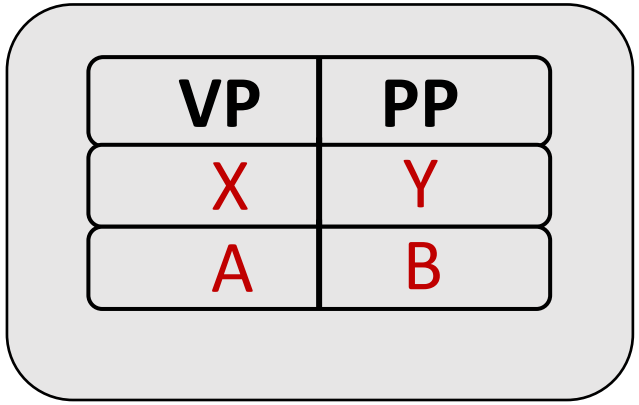

Translation Look-aside Buffers (TLBs)

The Processor Pipeline with the TLBs

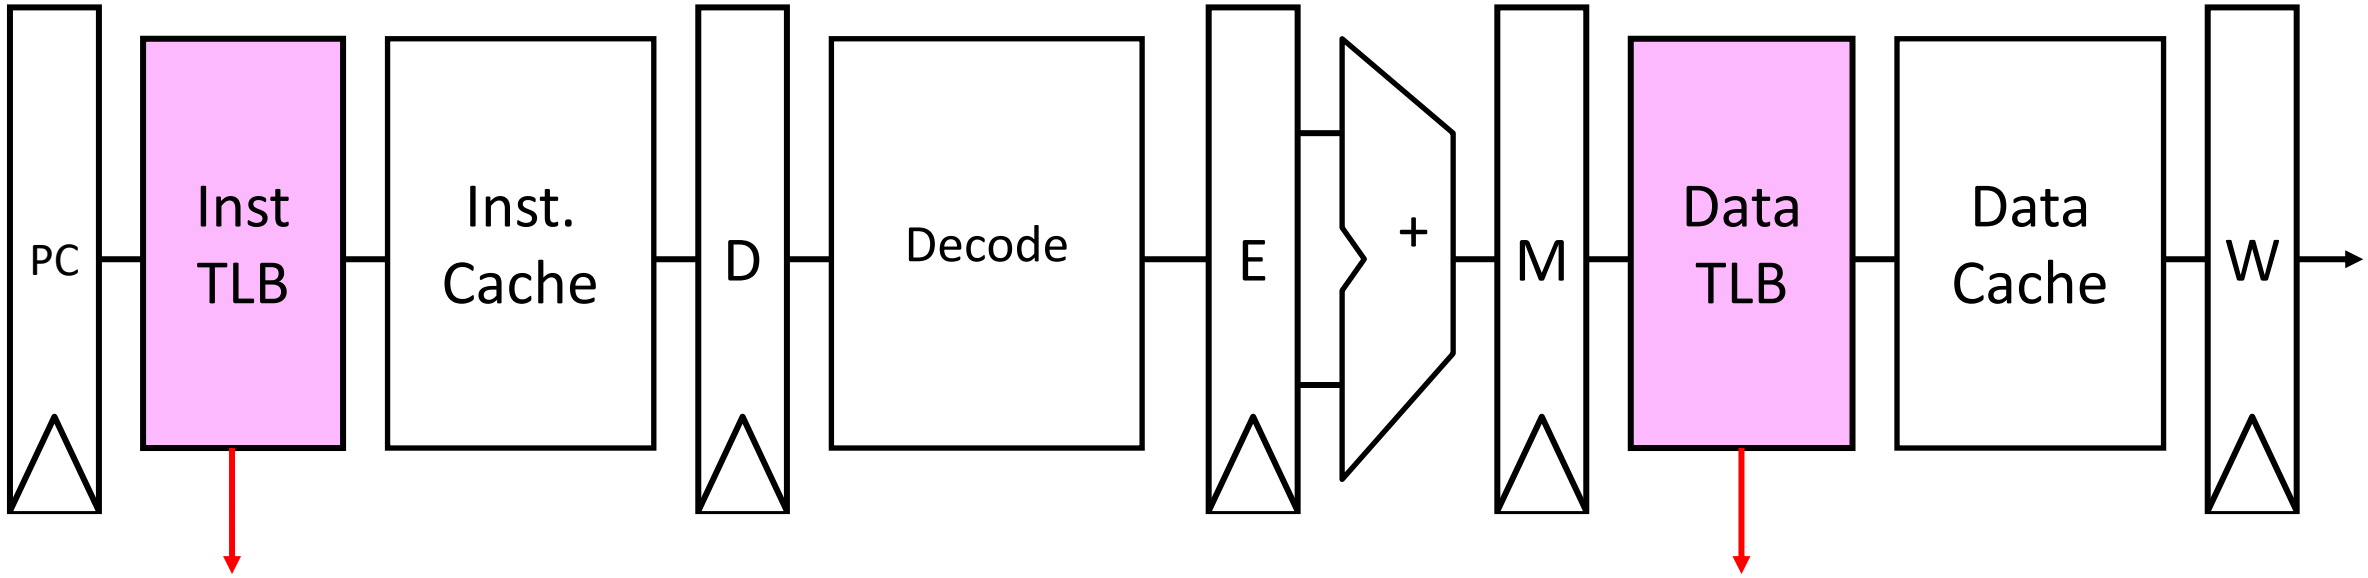

Memory Hierarchy with the TLBs

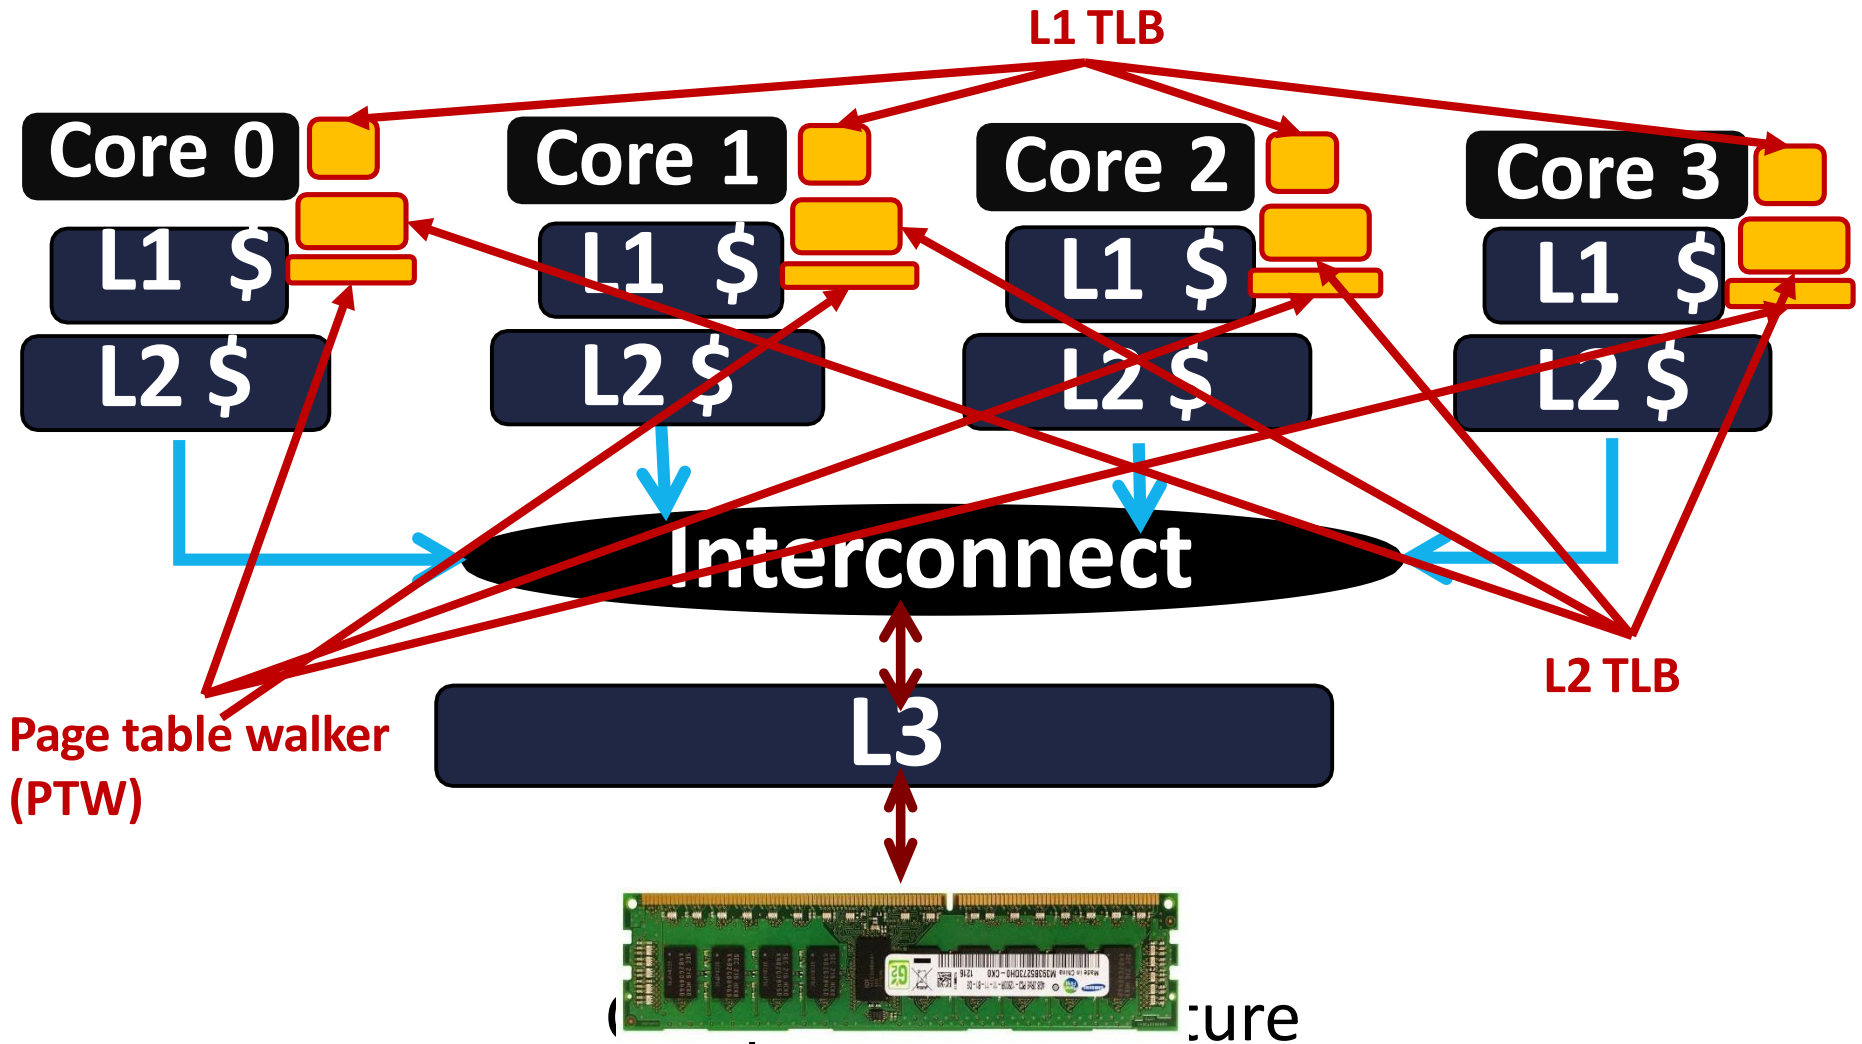

A bit Deeper

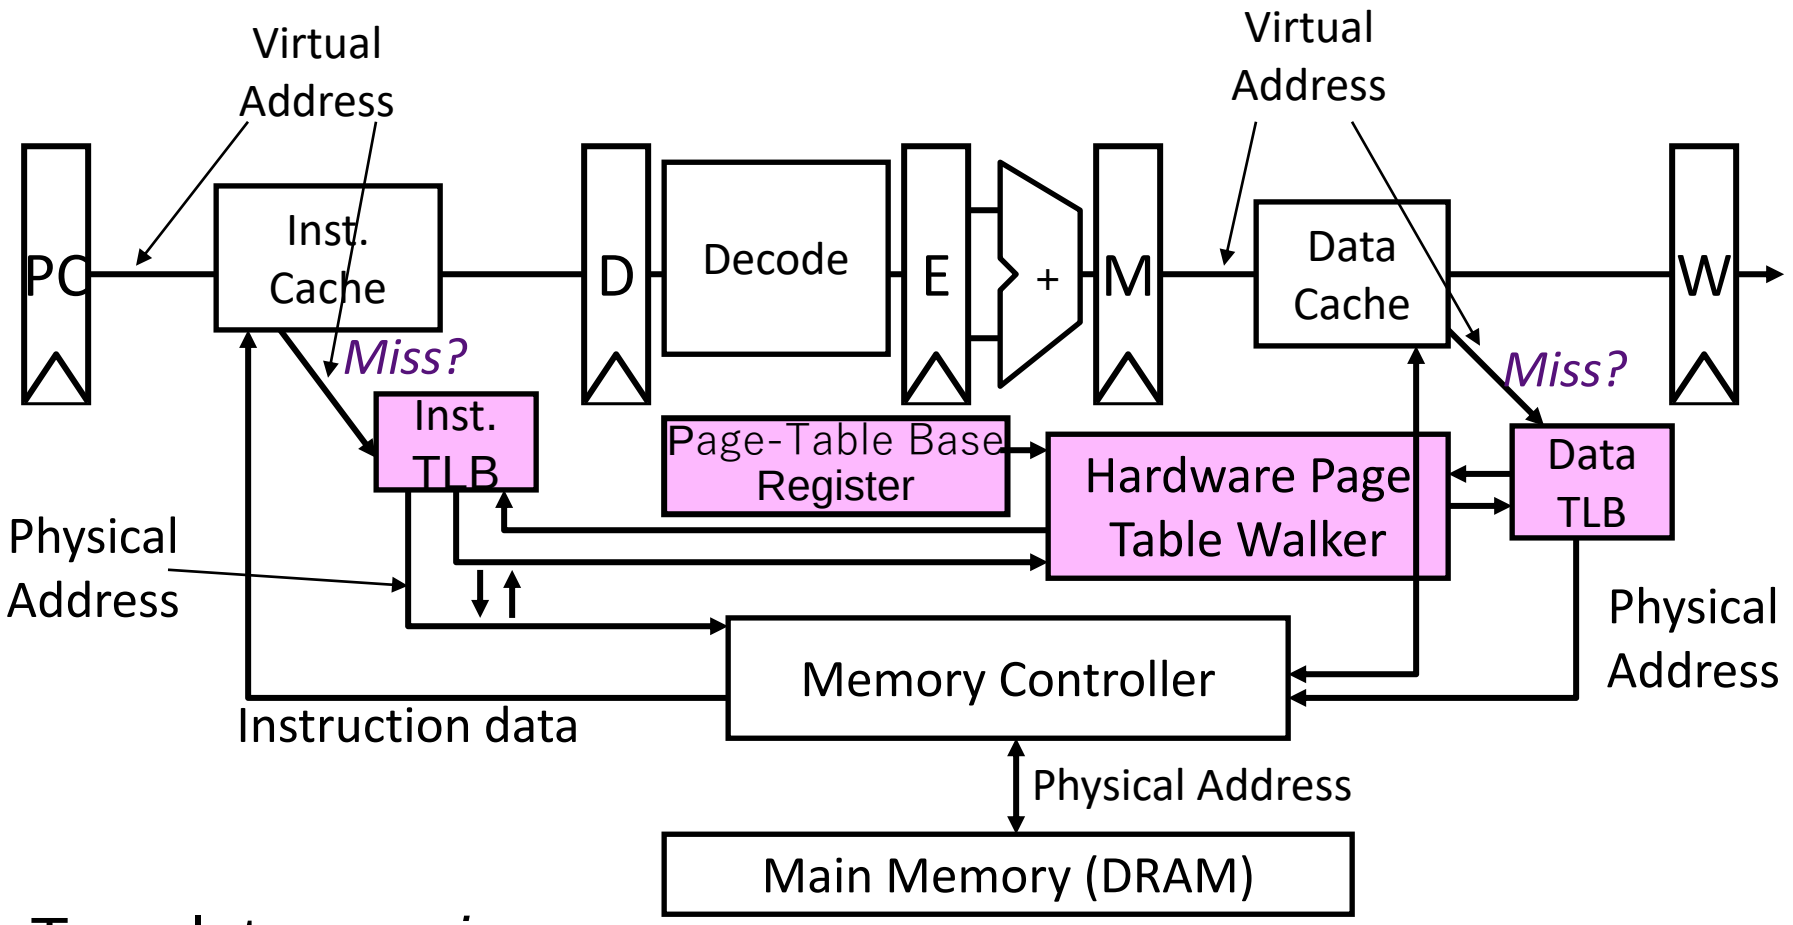

Translate on *miss*

Caches: Virtual or Physical

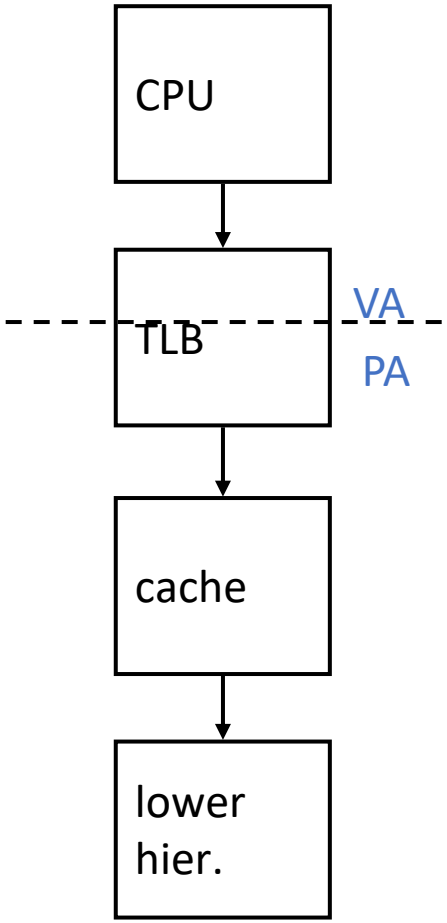

physical cache

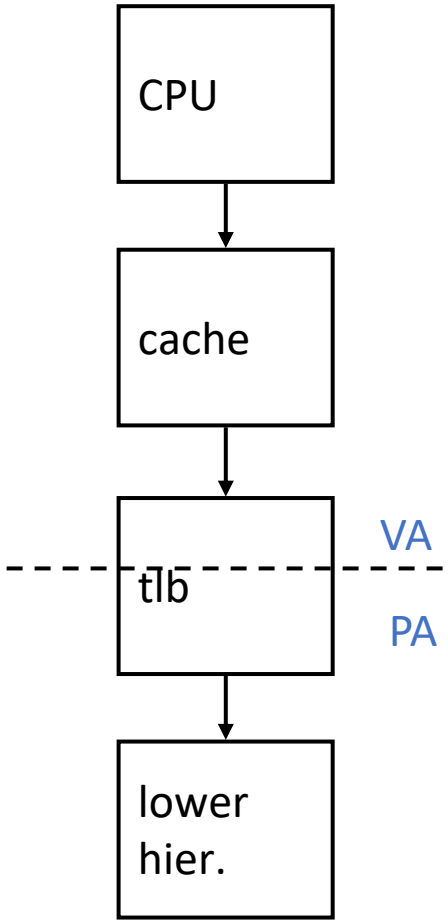

virtual (L1) cache

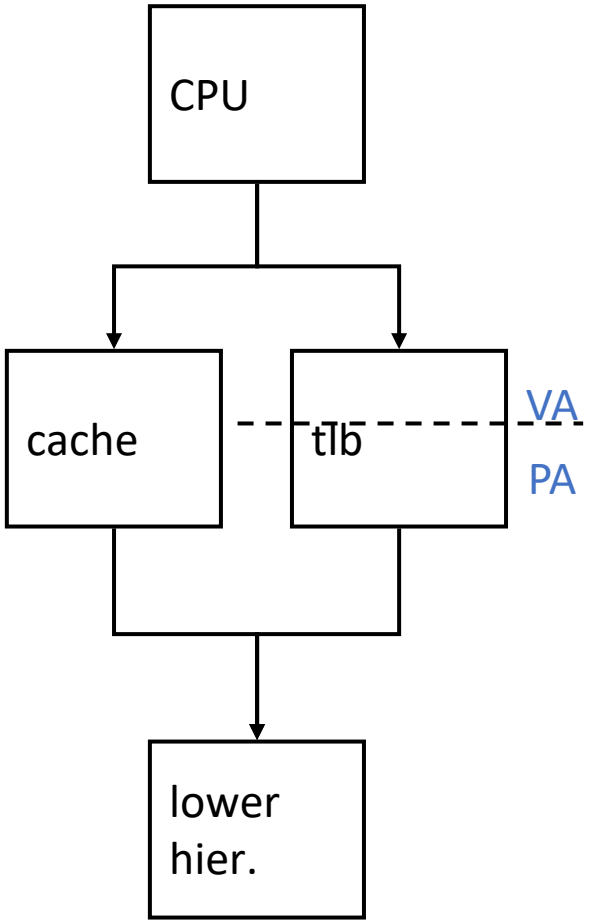

virtual-physical cache

Synonym Problem

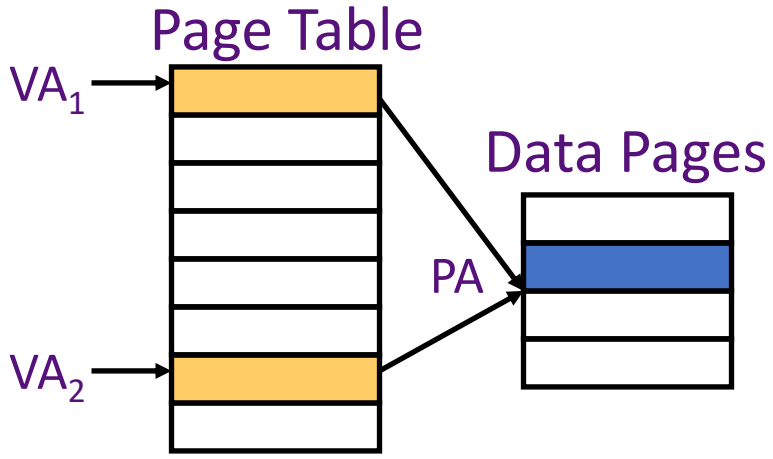

Two virtual pages share one physical page

Tag	Data
VA_1	1st Copy of Data at PA
VA_2	2nd Copy of Data at PA

Virtual cache can have two copies of same physical data. Writes to one copy not visible to reads of other!

General Solution: *Prevent aliases coexisting in cache*

Software (i.e., OS) solution for direct-mapped cache

Homonym Problem

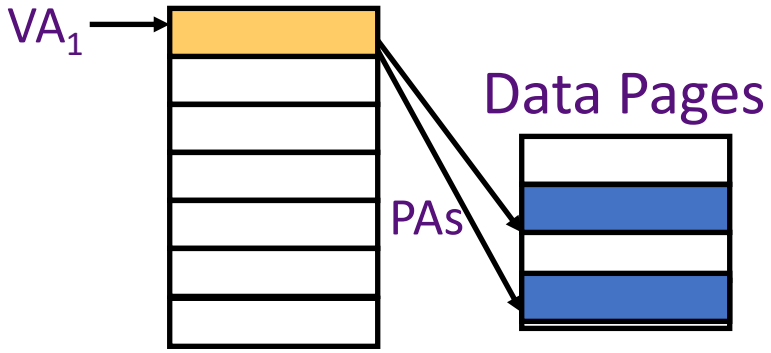

One virtual page maps to two physical pages

Tag may not uniquely identify cache data
Solution: Add ASID with tag
Or
Physical tags
Or
Flush on context switch

VIPT Caches

- If C≤(page_size × associativity), the cache index bits come only from page offset (same in VA and PA)
- If both cache and TLB are on chip: index both arrays concurrently using VA bits, check cache tag (physical) against TLB output at the end

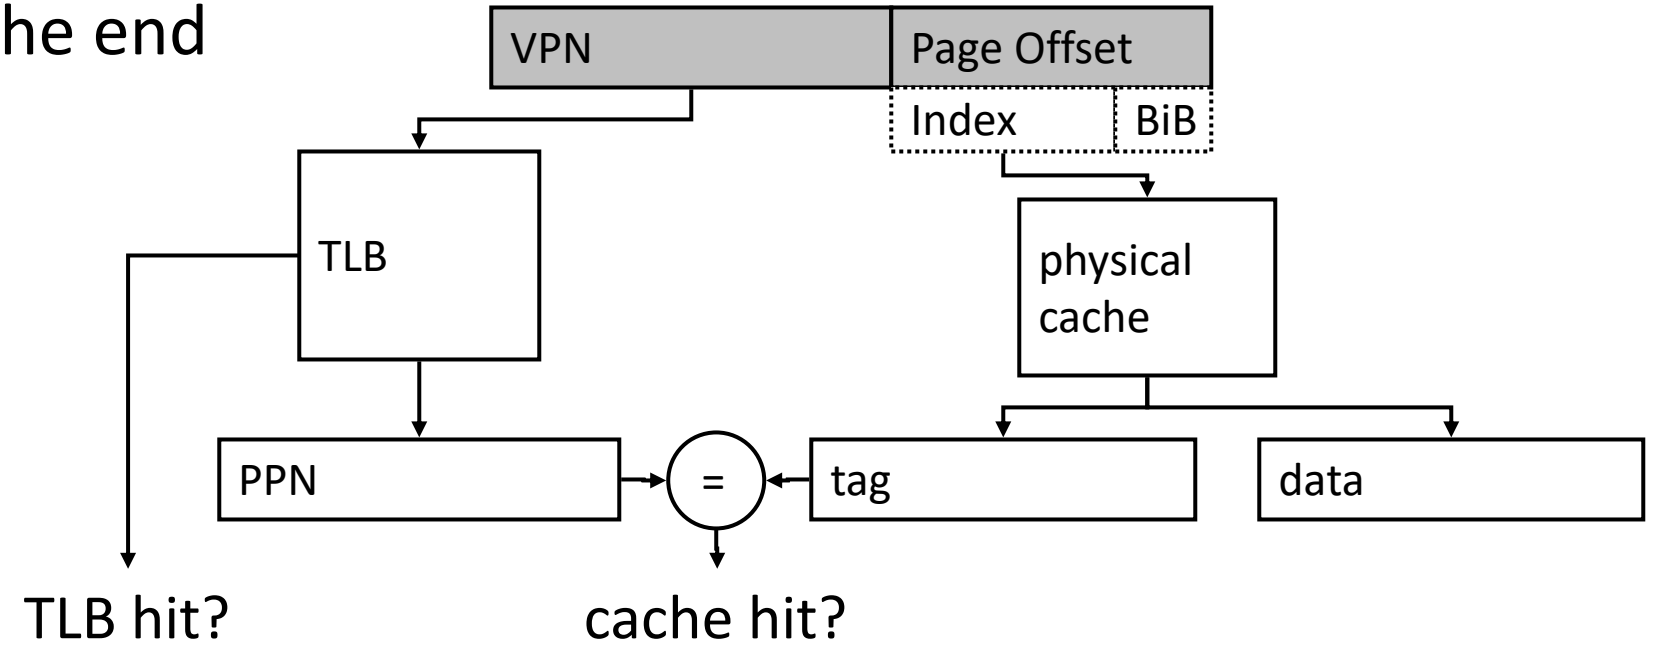

What if? Think about PIPT, VIPT, PIVT, and VIVT

- If C>(page_size × associativity), the cache index bits include VPN ⇒ Synonyms can cause problems
  - The same physical address can exist in two locations
- Solutions?

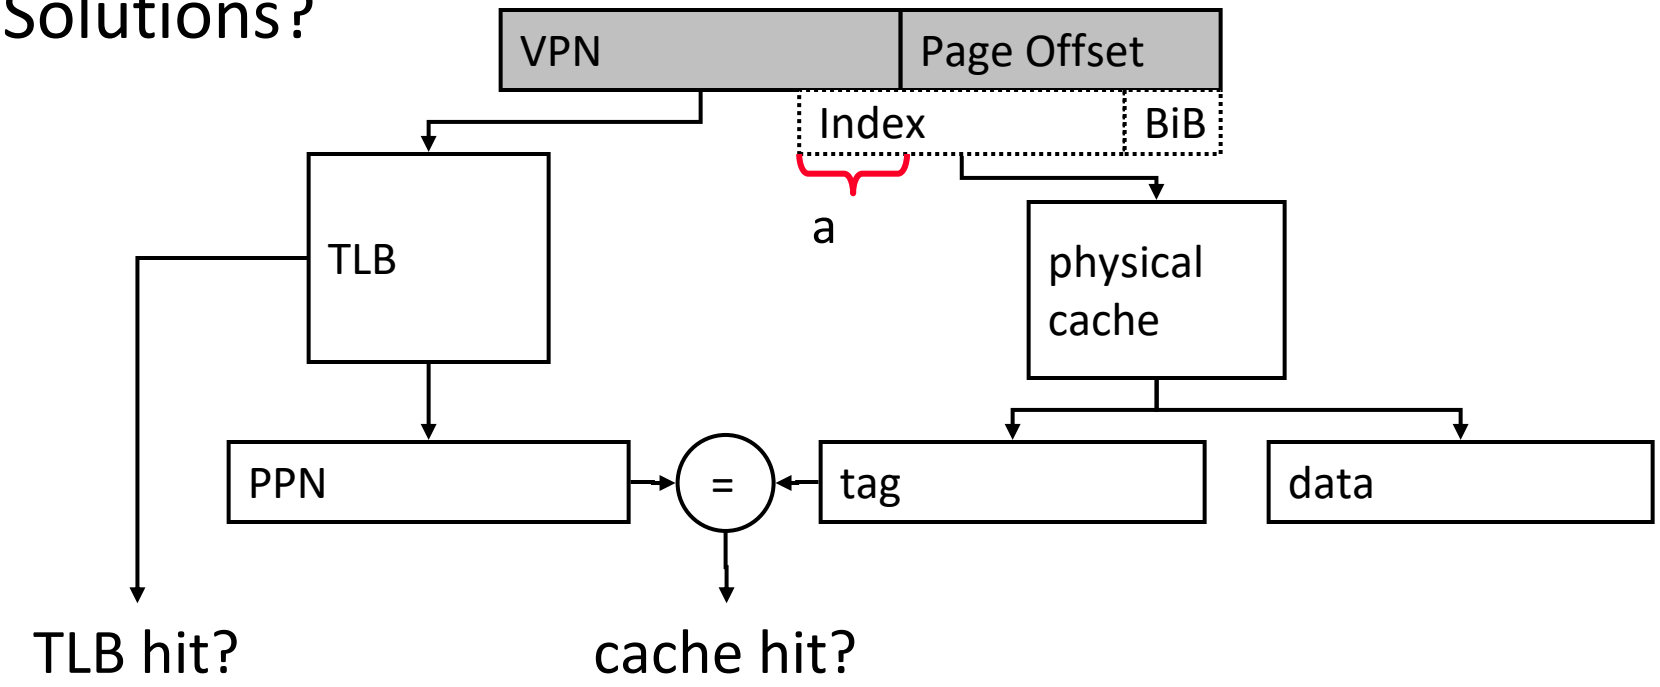

Summary

- VIVT (Virtual cache): Fastest, Synonyms, Homonyms,
- VIPT: Good enough, No Homonyms, mostly in L1 caches
- PIVT: ??
- PIPT (Physical cache): You know it

Real Caches

https://en.wikichip.org/wiki/intel/microarchitectures/sunny_cove

Bohoma Istuti