Slides for Webinar on Multicore

Bjorn Andersson

Dionisio de Niz

Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

DM22-0340

DoD Systems Interact with Their Physical Environment

- software controls physical environment
- has real-time requirements

"The trick there, when you're processing flight critical information, it has to be a deterministic environment, meaning we know exactly where a piece of data is going to be exactly when we need to — no room for error," Langhout says. "On a multi-core processor there's a lot of sharing going on across the cores, so right now we're not able to do that."

- Jeff Langhout, Acting Director, U.S. Army Aviation and Missile Research Development and Engineering Center (AMRDEC)

Source: "Army still working on multi-core processor for UH-60V," May 2017, Available at https://www.flightglobal.com/news/articles/army-still-working-on-multi-core-processor-for-uh-6-436895/.

Commonality of DoD Systems

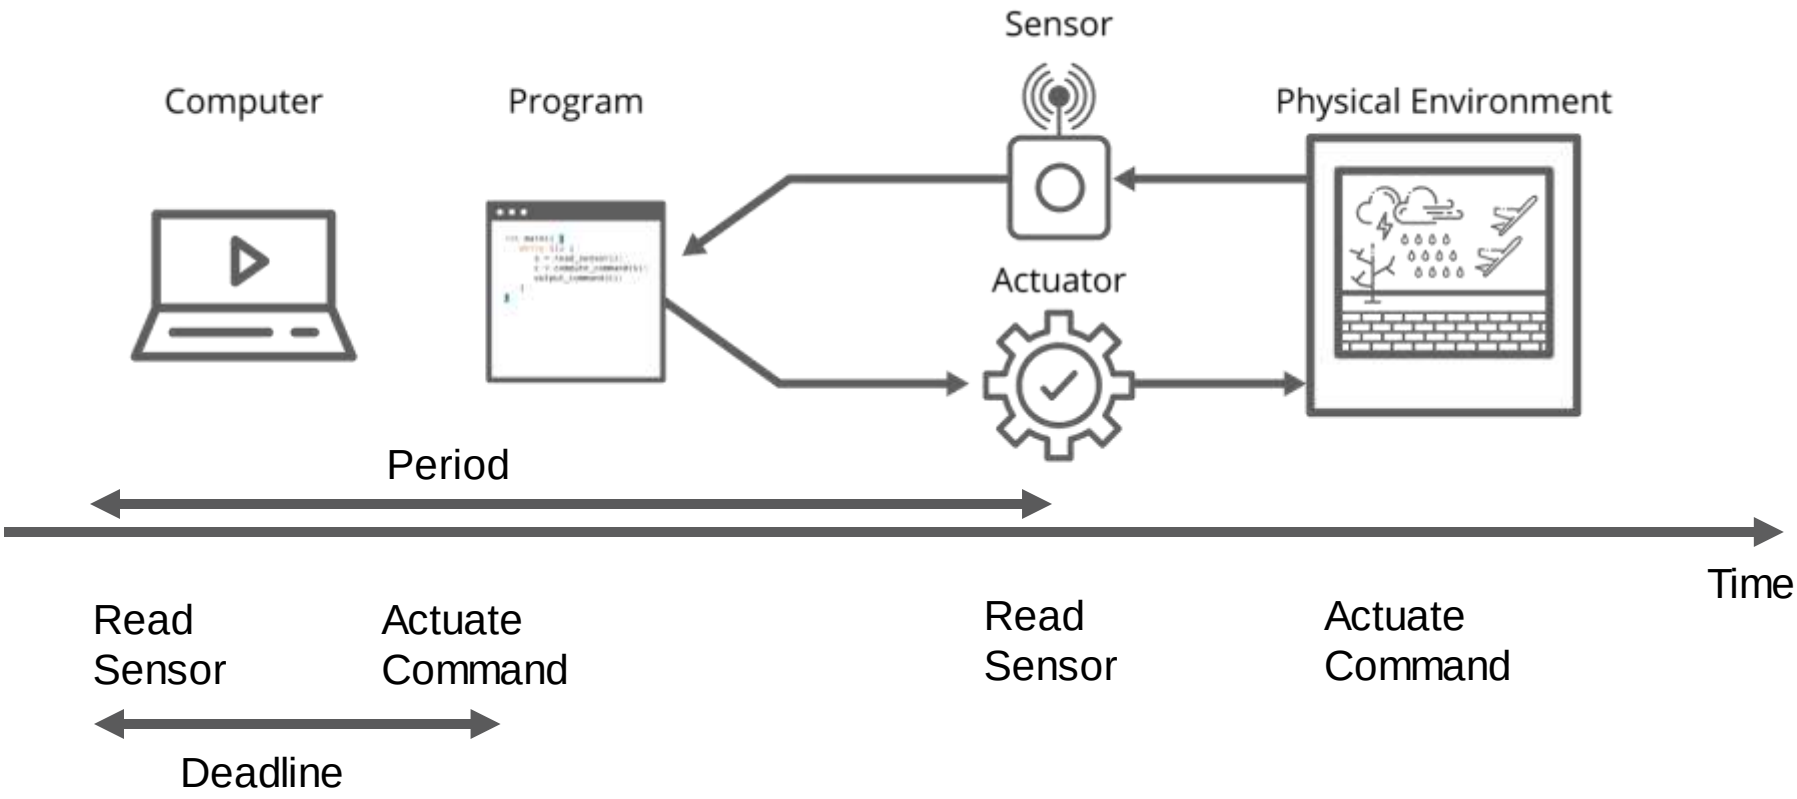

Commonality of DoD Systems

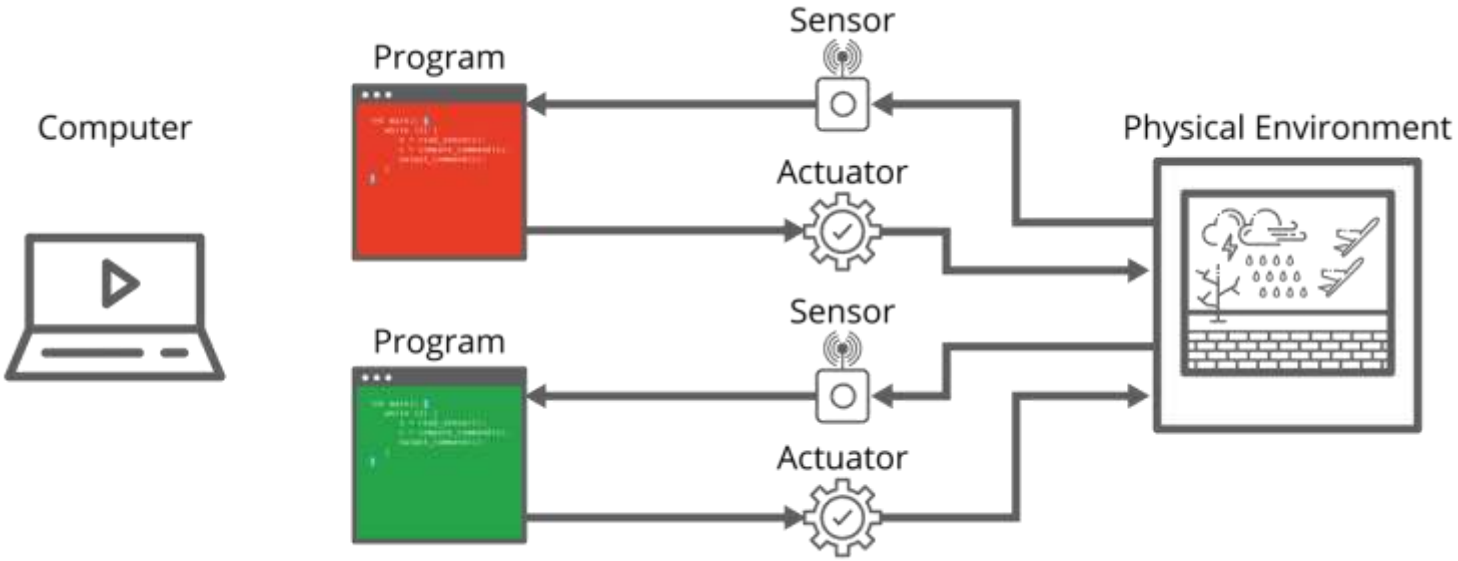

Why is Satisfying Real-Time Requirements Challenging?

What Causes Delay of Software?

What Causes Delay of Software?

What Causes Delay of Software?

What Causes Delay of Software?

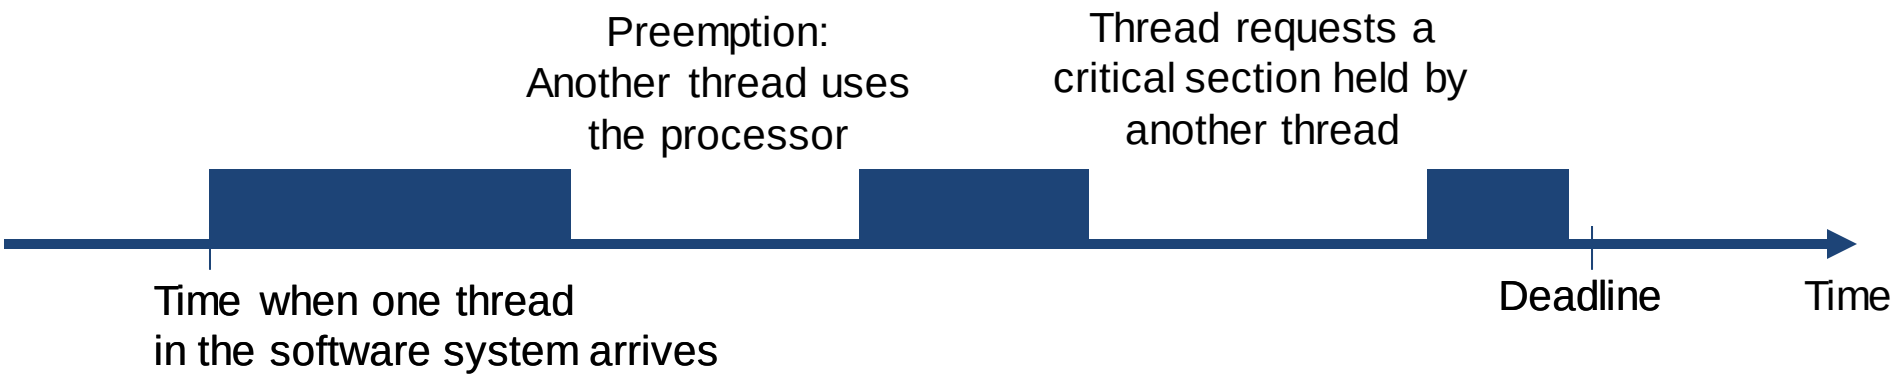

Real-Time Requirements of Software Executing on a Multicore Processor

Hardware Trends

- *All computers are multicores*
- *Most chip makers do not offer single core.*
- *Most multicores have shared memory.*

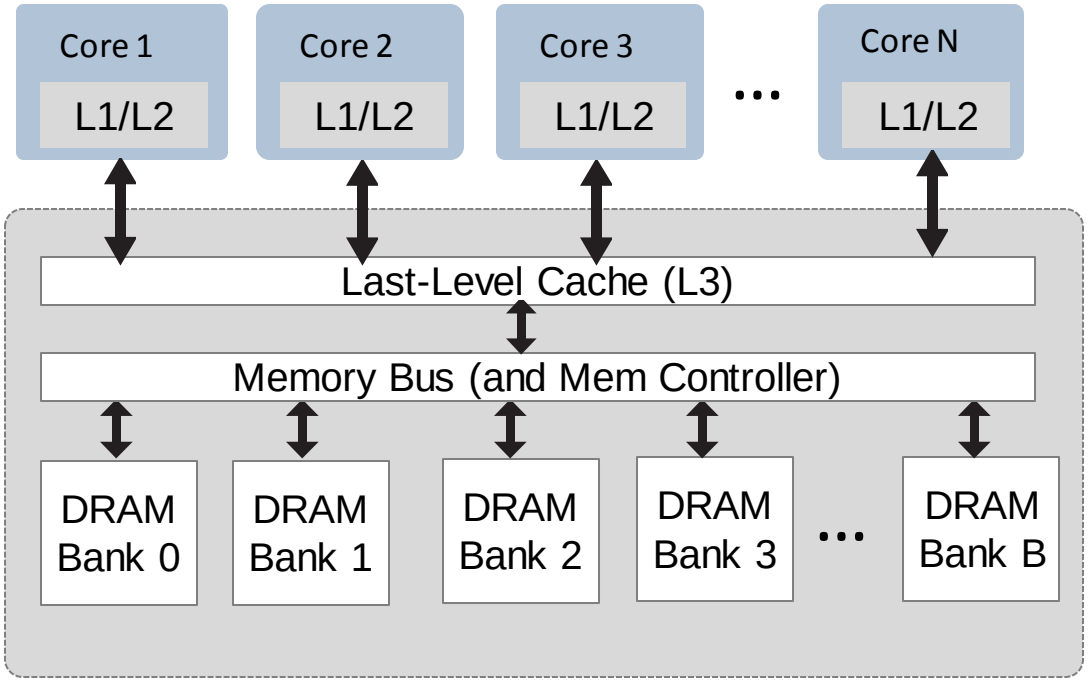

Problem: For each process, compute an upper bound on its response time.

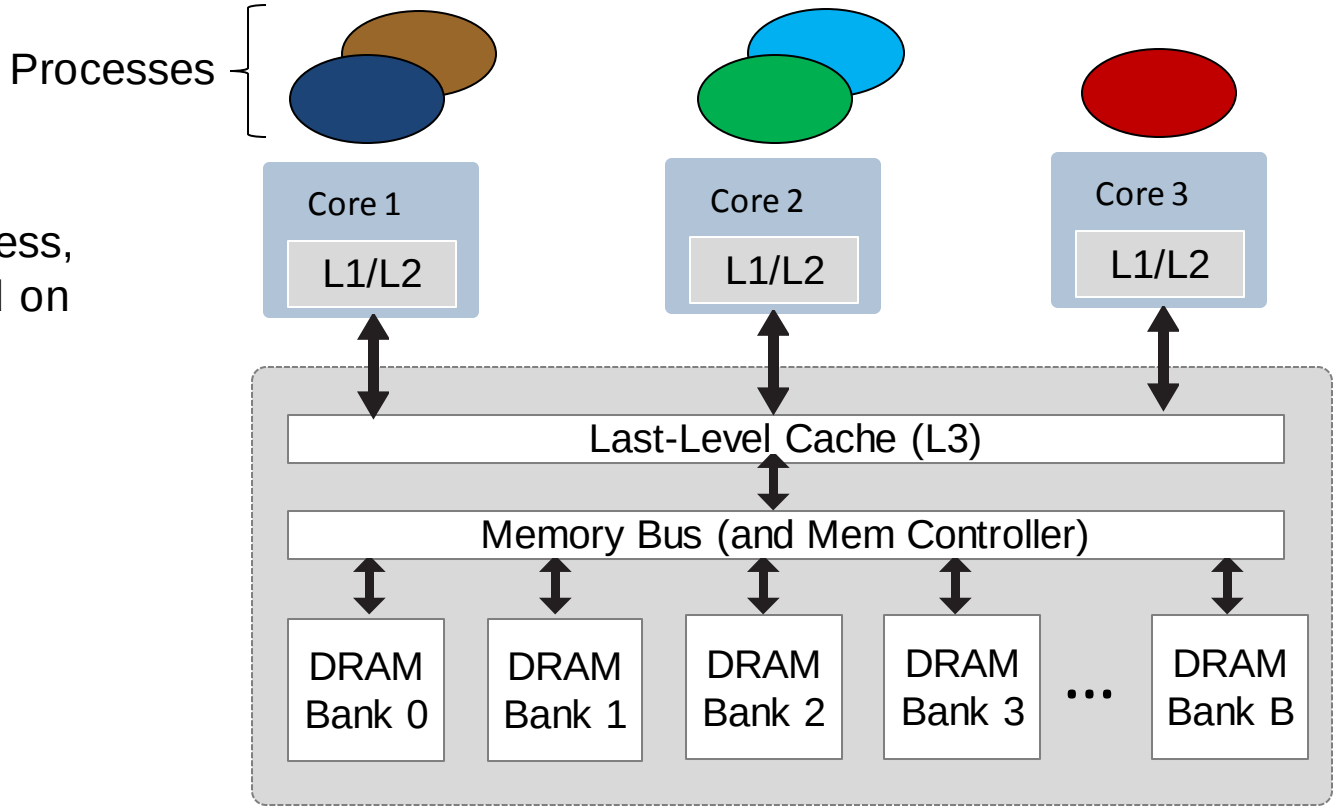

How Co-Runners Impact Speed of Execution

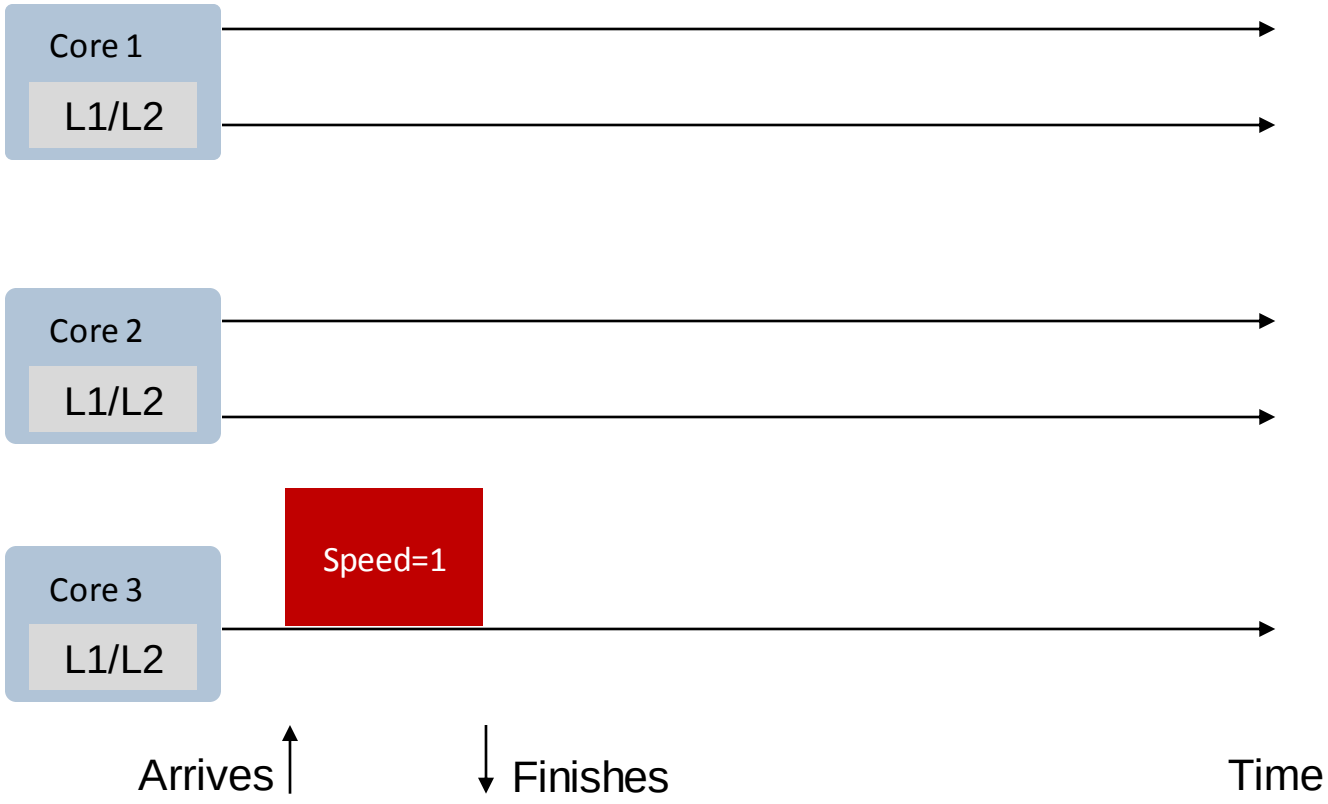

How Co-Runners Impact Speed of Execution

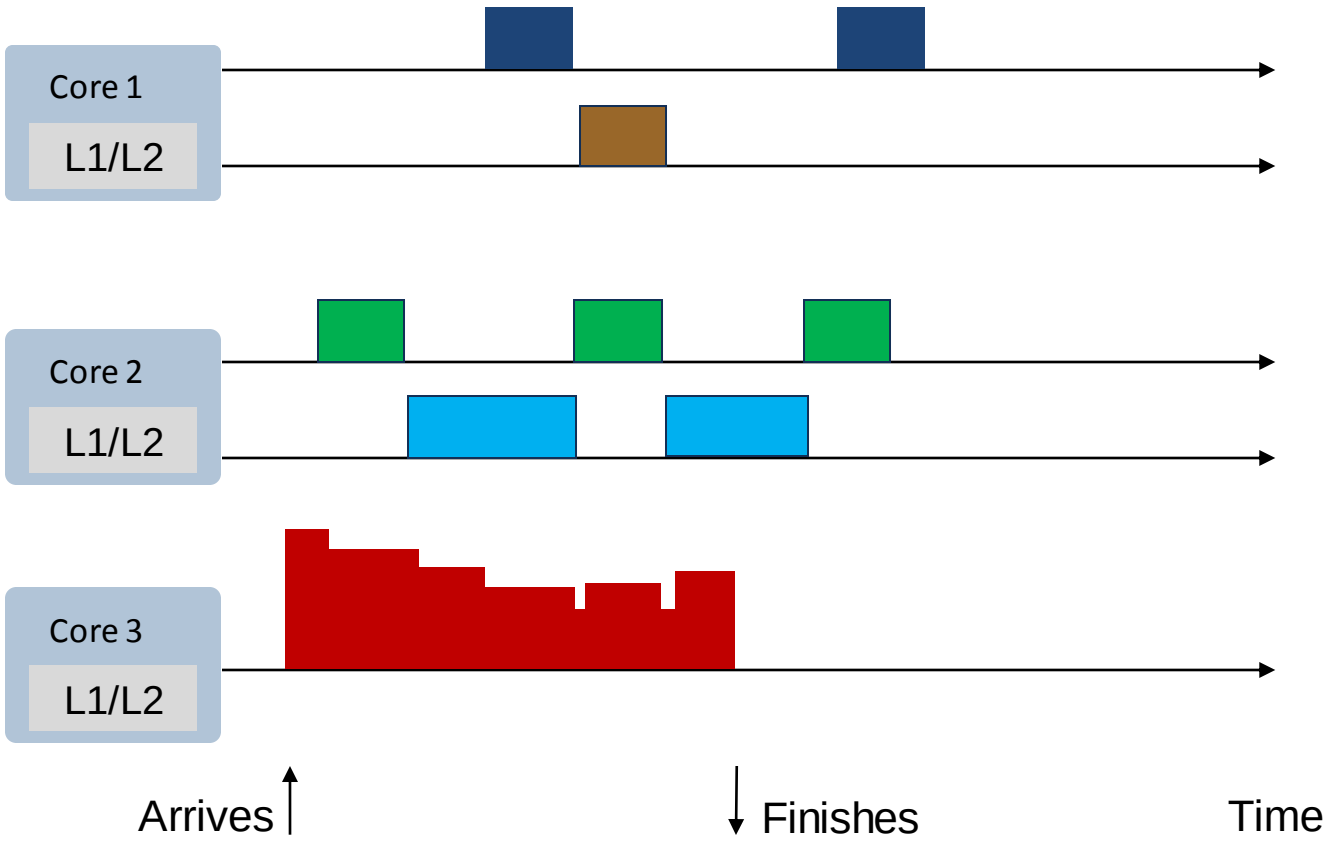

Issues

- Shared hardware resources impact timing.
- 103 times slowdown has been observed [Yun15].
- Current methods cannot deal with undocumented resources.
- Even when resources are documented, current methods can only analyze/manage a small set of them.
- The problem is getting worse:
  * Slowdown increasing
  * More undocumented h/w

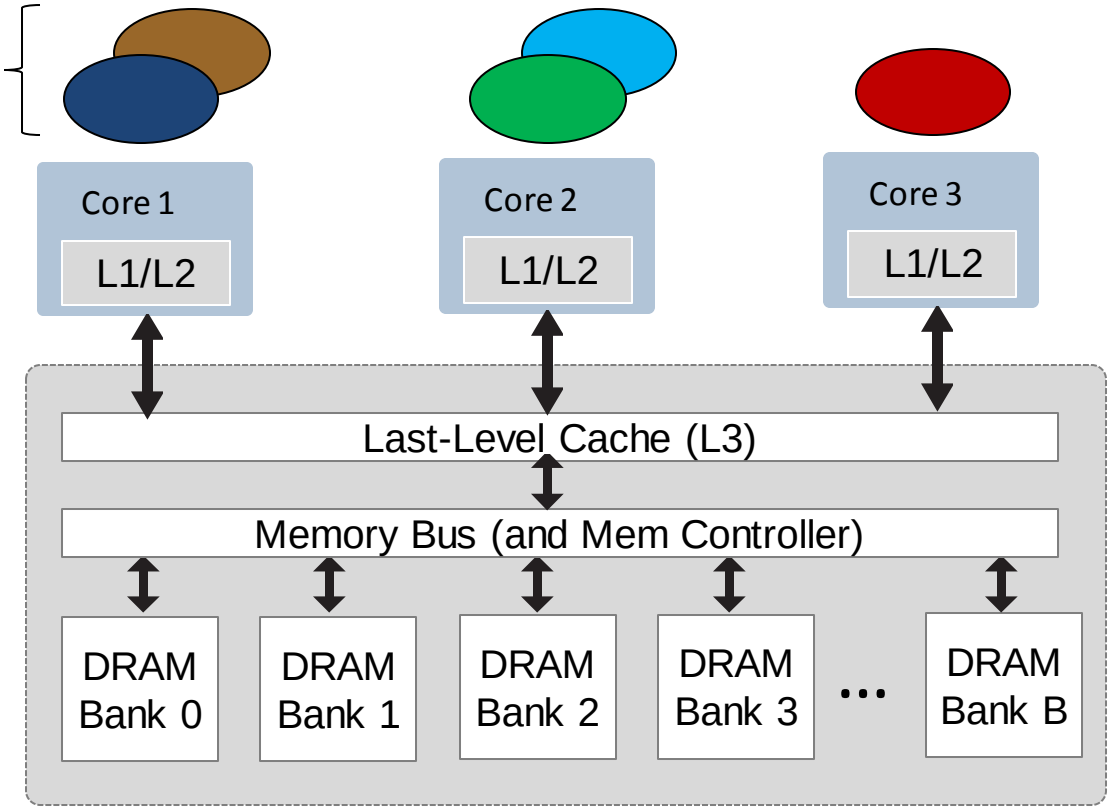

This slide is intentionally left blank

Multicore Becomes Mainstream

Power Wall: May 2004 Intel cancels the *Tejas* CPU

- Projected to run at 7GHz
- Generated 150 Watt of heat at 2.7GHz
- Intel starts with dual core

Memory Wall: Sandia Labs

- 2007/2008 study on memory bandwidth in data intensive applications
- With conventional memory
  - More than 8 cores performance decreases
  - **More than 16 cores worse than 1 core**

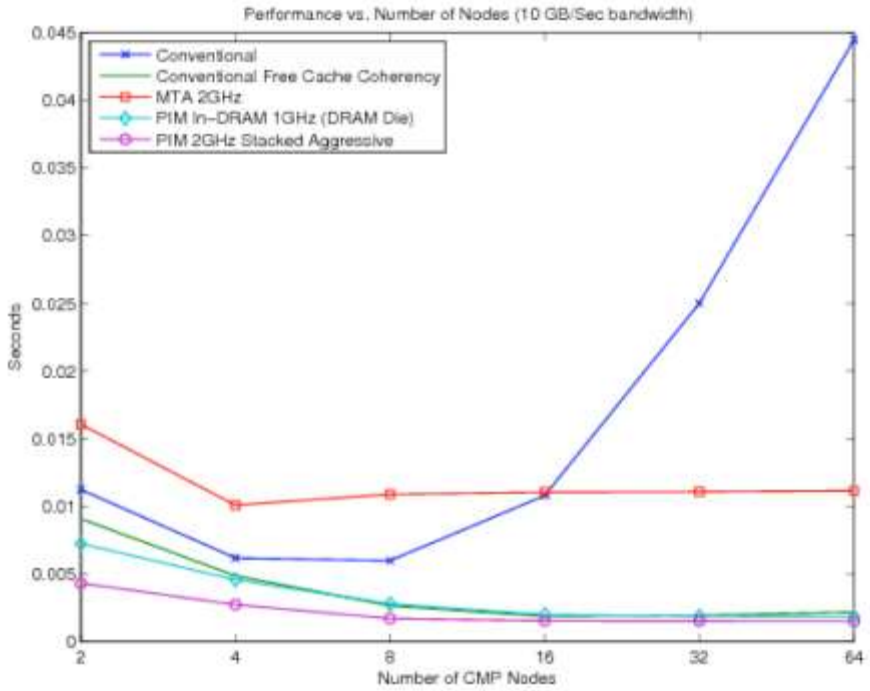

Good News !

Working solutions for the memory wall

Growing number of cores – more computing power!

- Intel :
  - Intel® Core™ i9-10980XE : 18 cores
- AMD:
  - AMD Ryzen™ 9 5950X : 16 core, 32 hyperthreads

- Apple M1 Max
  - 10 cores

Bad News! ☹

Memory wall solutions focused on throughput

- Can make predictable worst-case execution time (WCET) worst

Hardware solutions for throughput hinders timing guarantees / predictability

Hardware solutions for throughput hinders timing analysis

Hardware solutions are poorly documented

Shared Hardware Resources

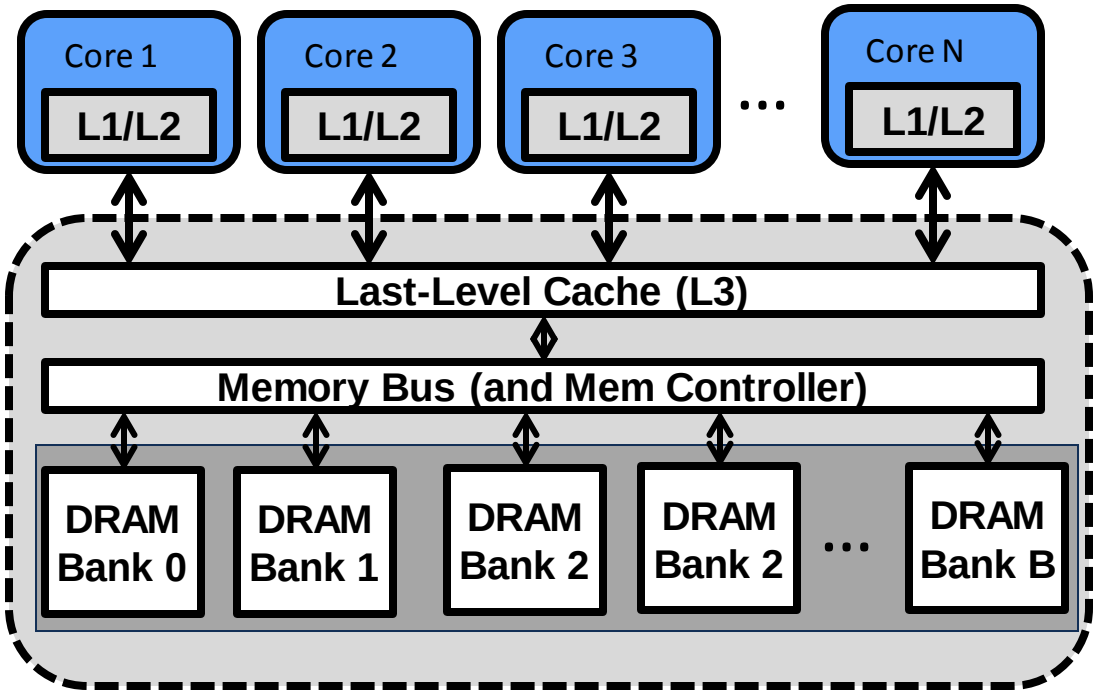

Shared Hardware Delays: Shared Cache

Core 1

Core 2

Shared Cache

Addr	Contents	Set #
		0
		1
		2
		3

Memory

	Addr:
A	0
B	1
C	2
D	3
E	4
F	5
G	6
H	7
I	8
J	9
K	10
L	11
M	12
N	13
O	14
P	15

Icons credit: http://www.doublejdesign.co.uk

Shared Hardware Delays: Shared Cache

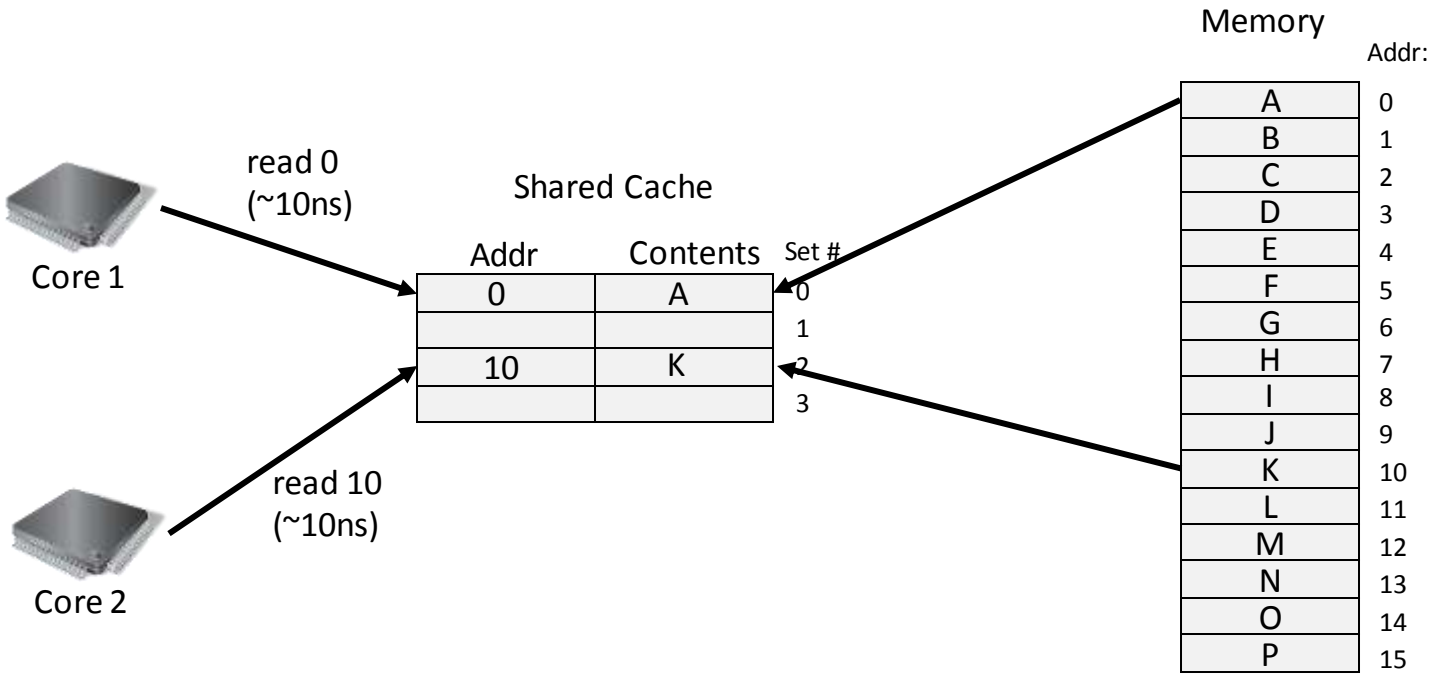

Icons credit: http://www.doublejdesign.co.uk

Shared Hardware Delays: Shared Cache

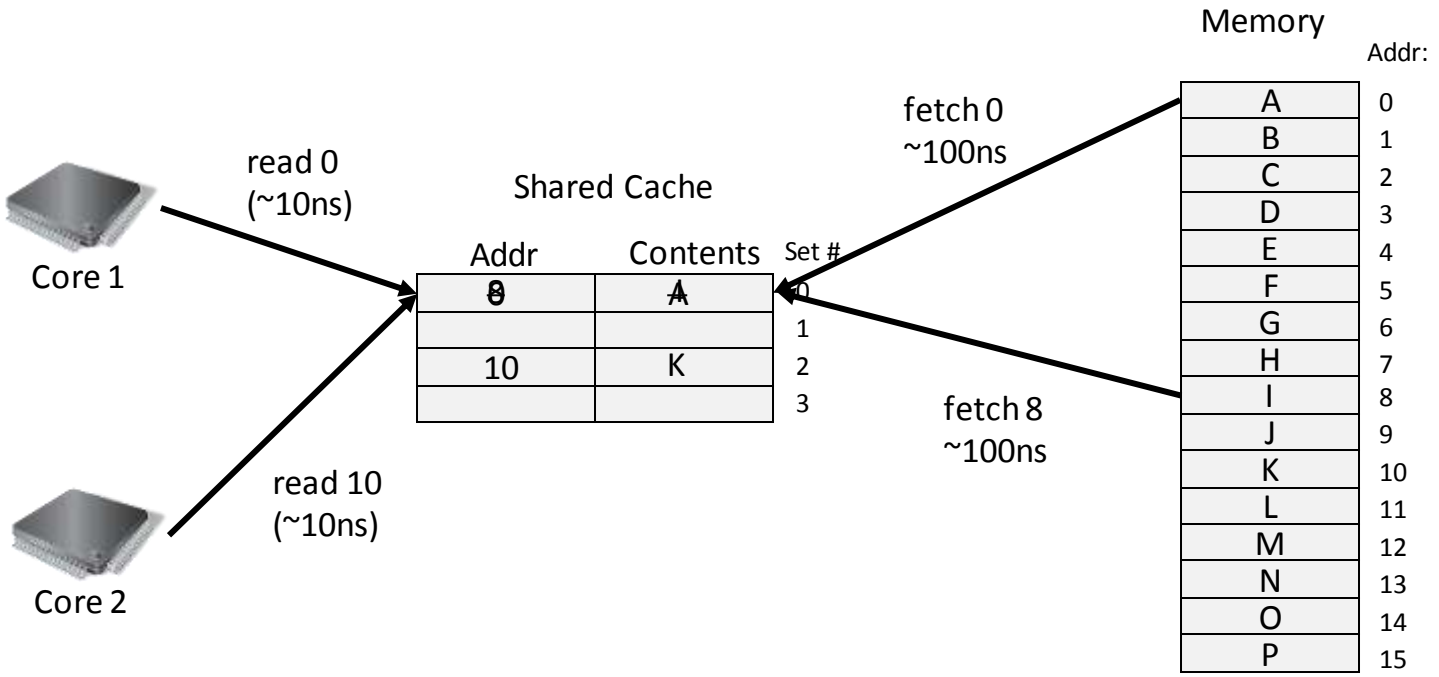

Icons credit: http://www.doublejdesign.co.uk

Shared Hardware Delays: Shared Cache

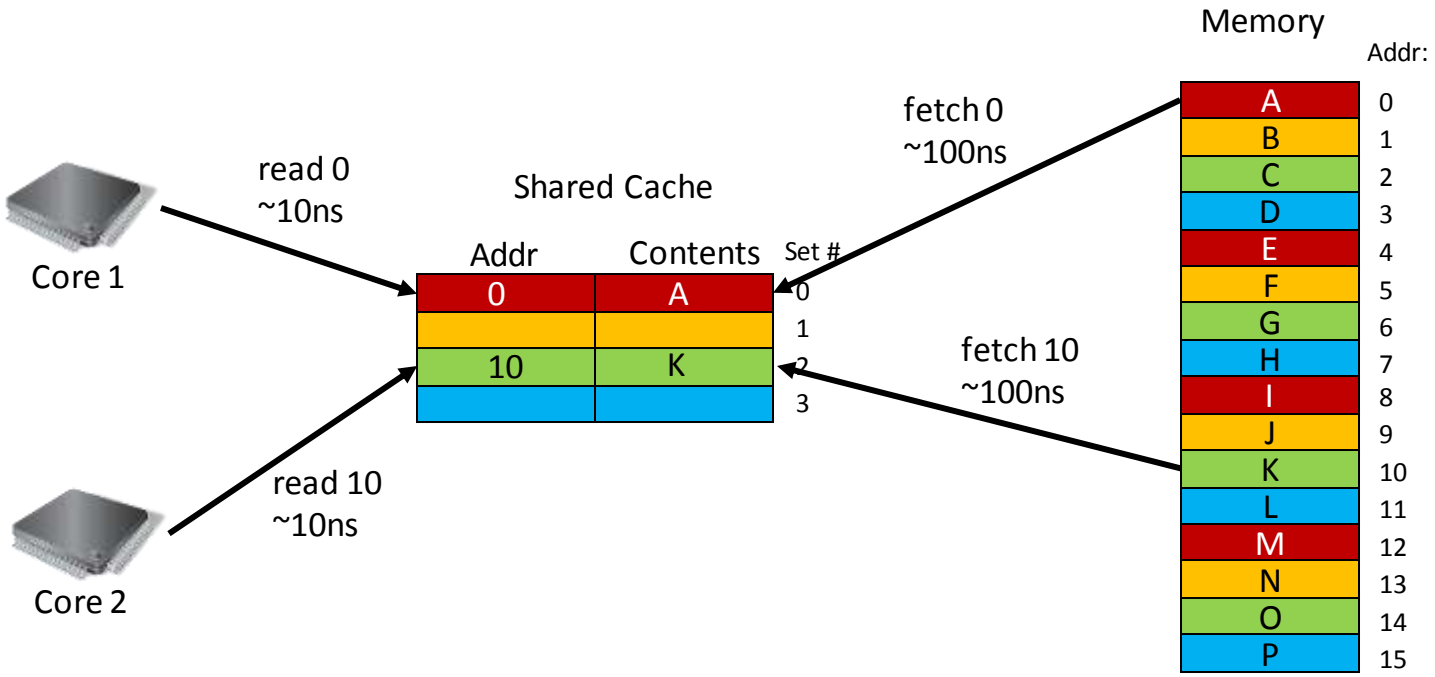

Icons credit: http://www.doublejdesign.co.uk

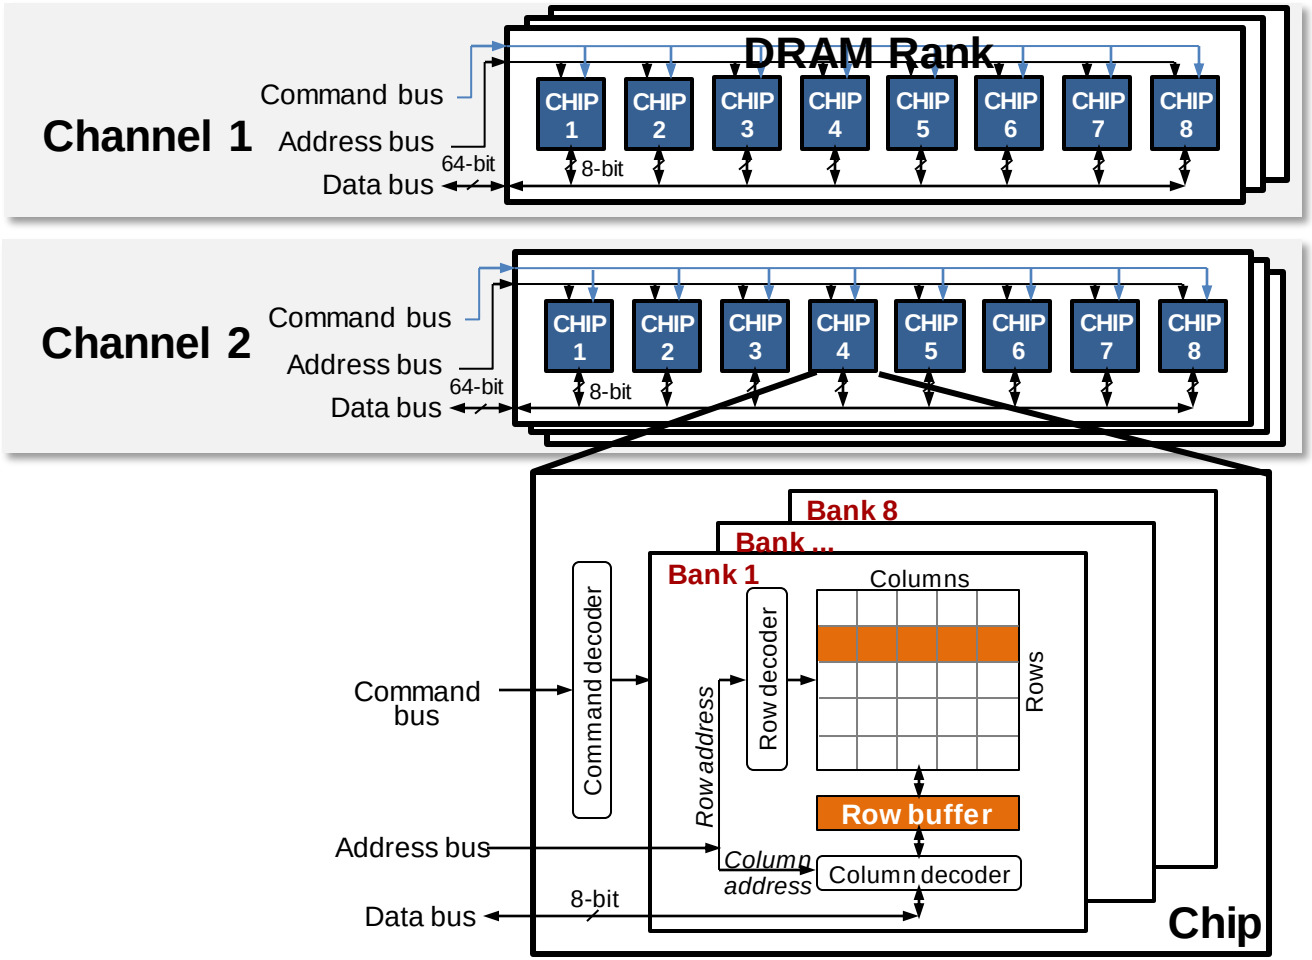

Memory Banks Delays

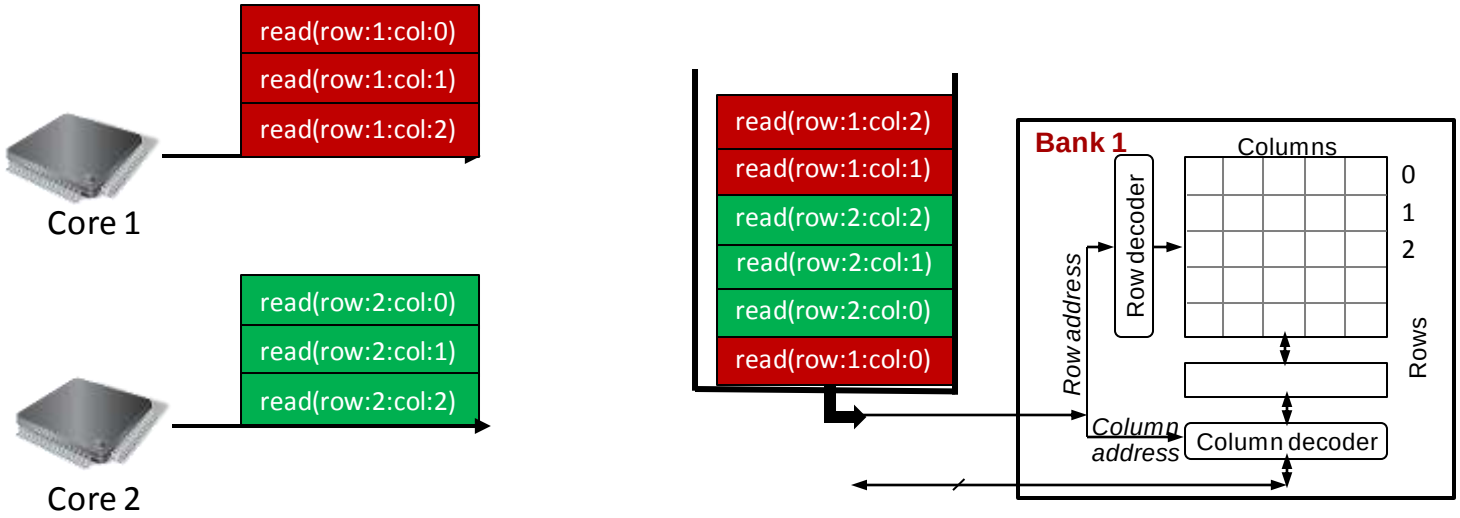

Icons credit: http://www.doublejdesign.co.uk

Memory Banks Delays

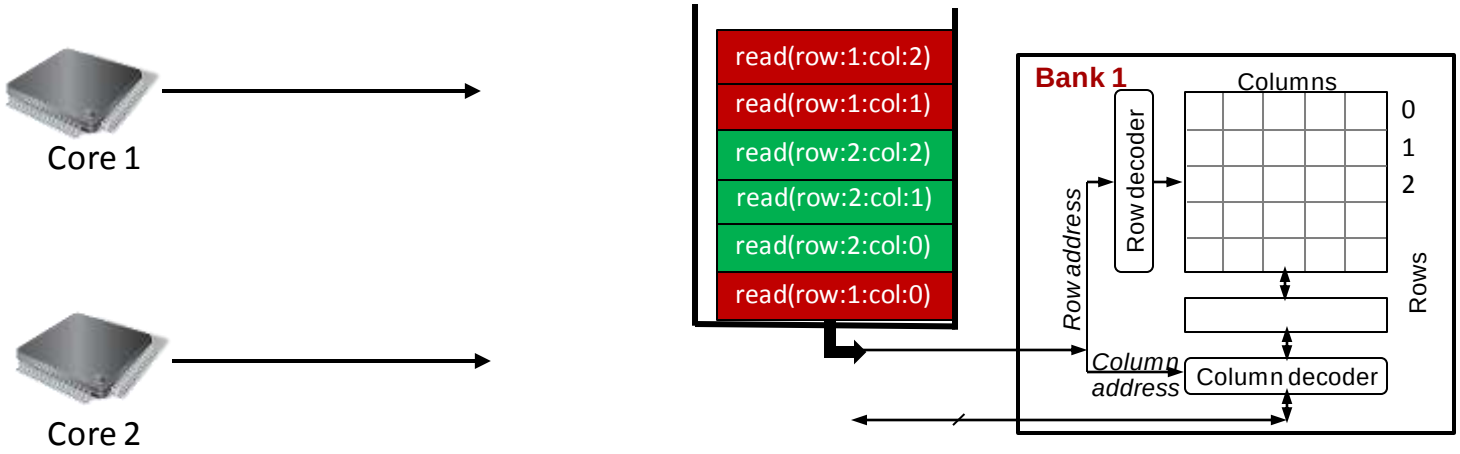

Memory Banks Delays

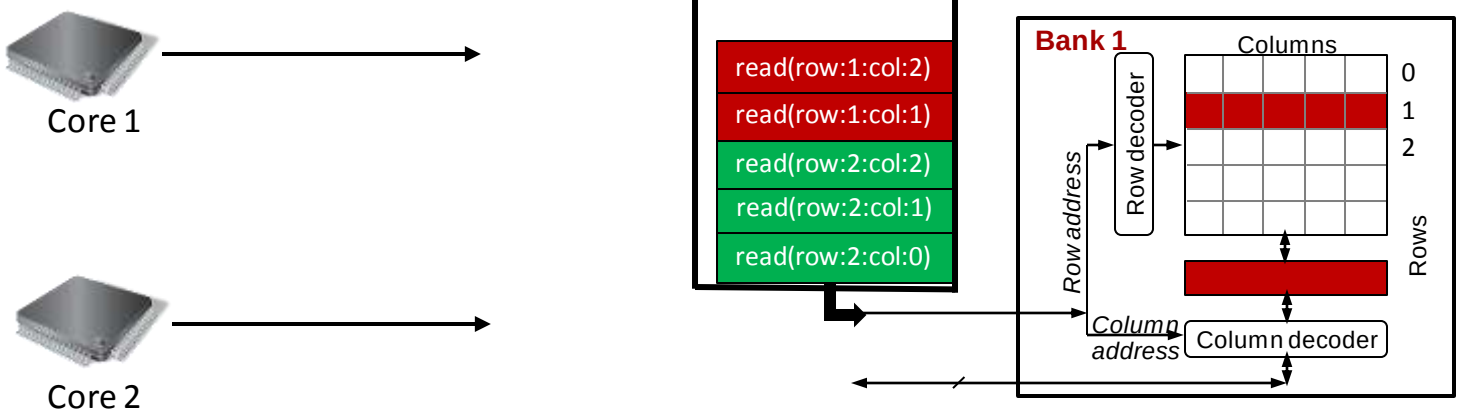

Icons credit: http://www.doublejdesign.co.uk

Memory Banks Delays

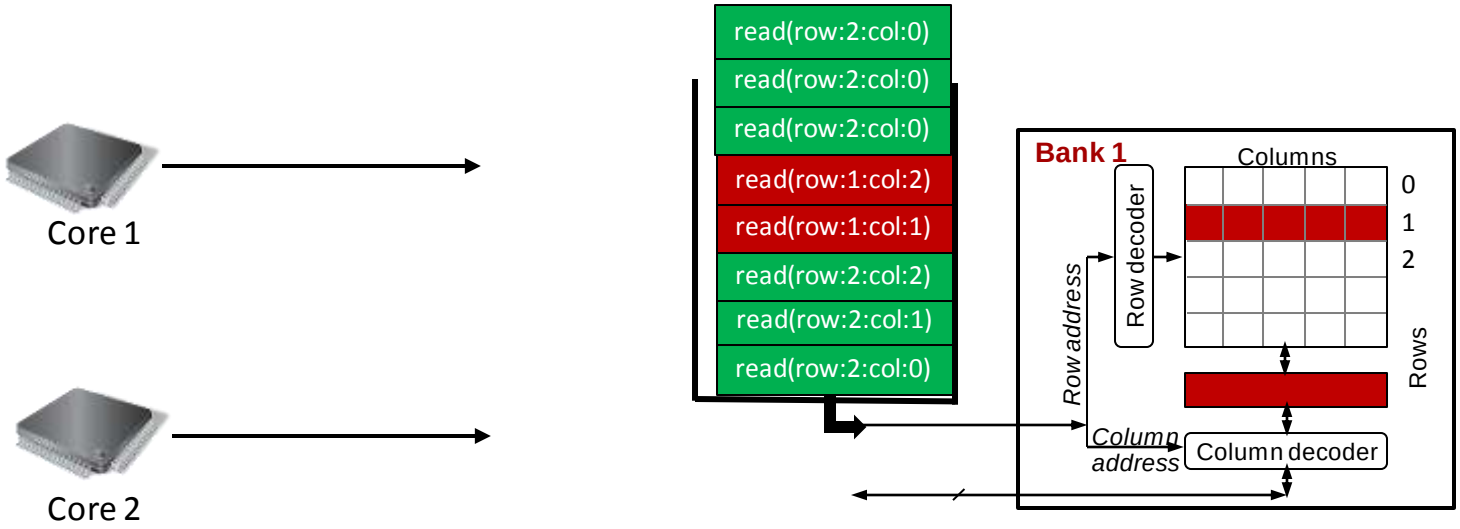

Icons credit: http://www.doublejdesign.co.uk

Memory Banks Delays

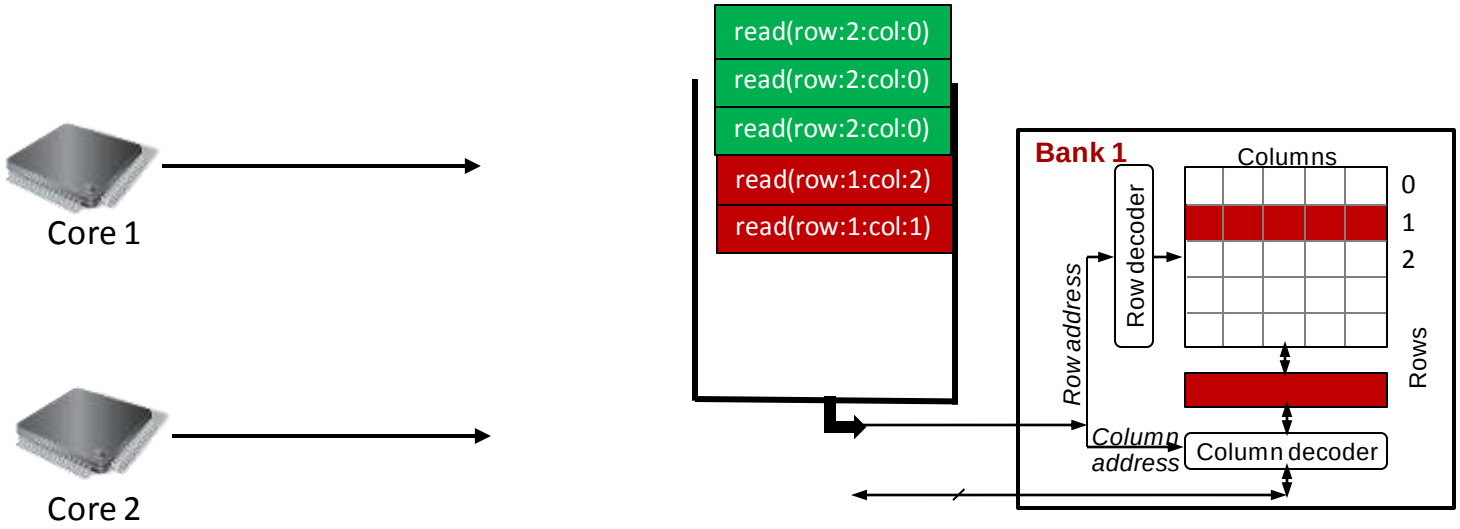

Icons credit: http://www.doublejdesign.co.uk

Memory Bank Coloring

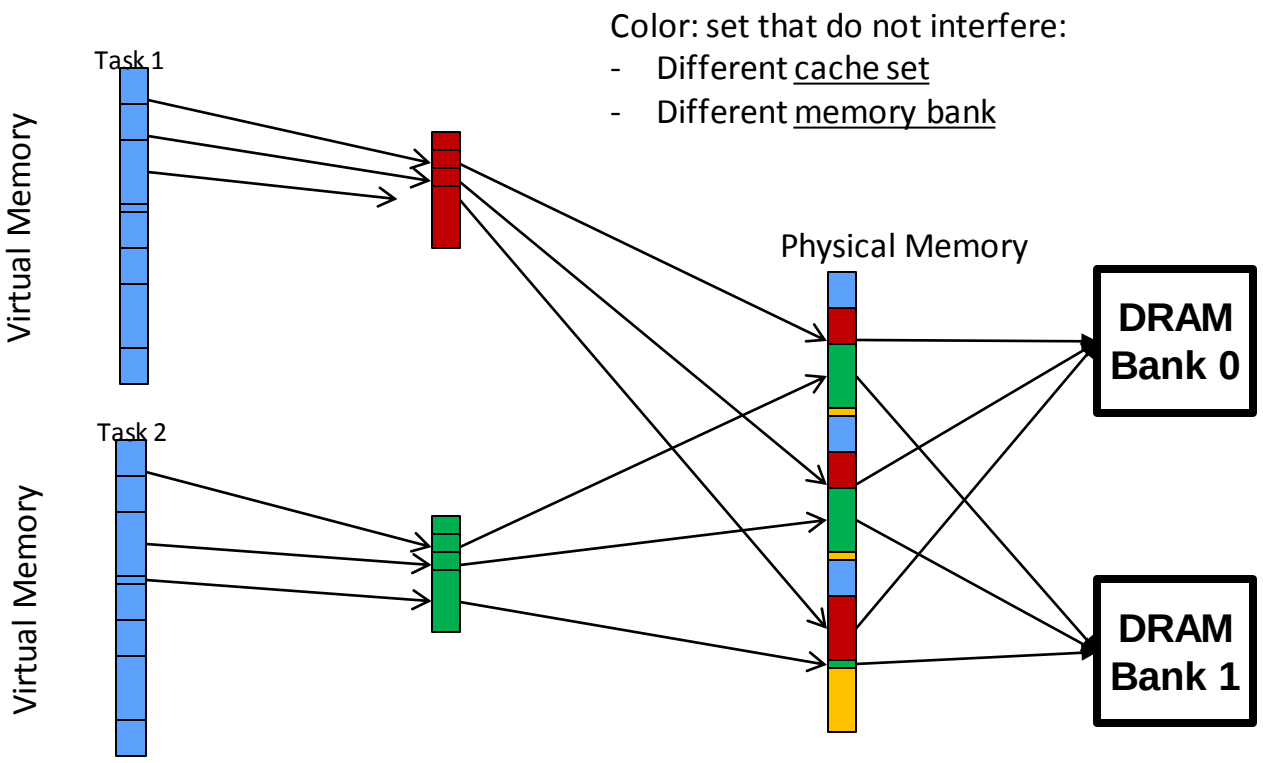

Conflicting Mappings

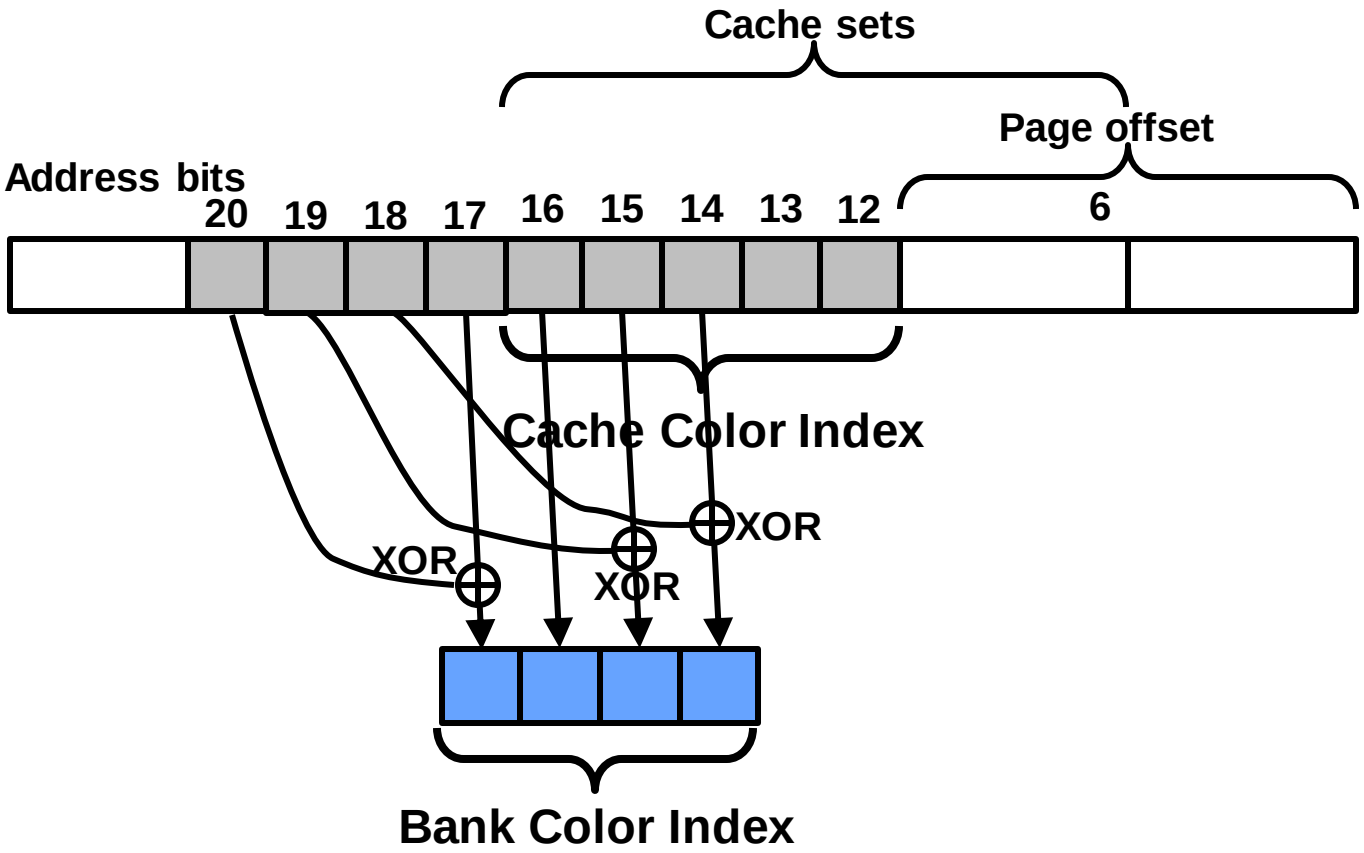

Other Shared Hardware

Direct Memory Access (DMA) I/O: Disk, Video cards, etc.

- Behave as another core that access memory

Translation Lookaside Buffer (TLB)

- Cache of virtual-to-physical page for virtual memory
- Similar effect to shared cache but for virtual-to-physical memory translations

Backup Slides

Other Architectures: NoC / Tiles

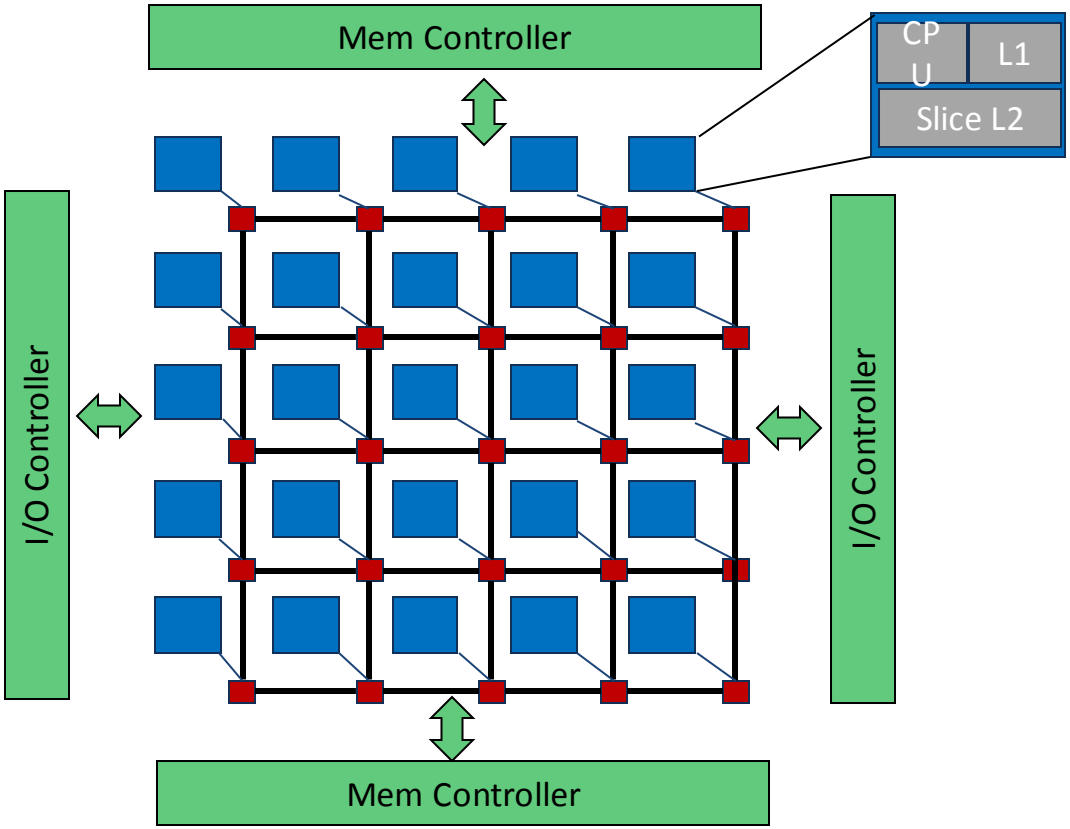

Final Remarks

Adding core is the only way to increase speed

- To avoid power wall

But leads to memory wall

Current solutions/mechanisms improve throughput

- But worsen predictable WCET

Need to understand current mechanisms to

- Prevent unpredictable behavior
- Analyze impact on predictable real-time behavior

Current mechanisms poorly documented

Real-Time Systems: Predictable Periodic Execution

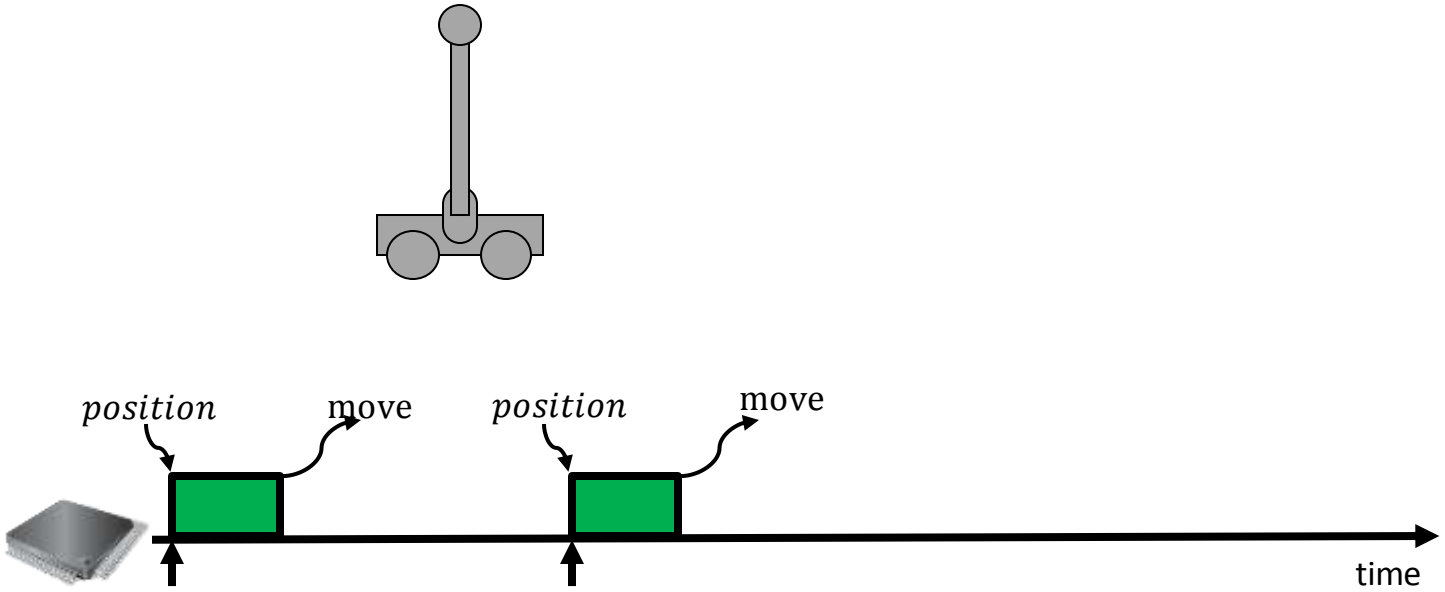

Real-Time Systems: Predictable Periodic Execution

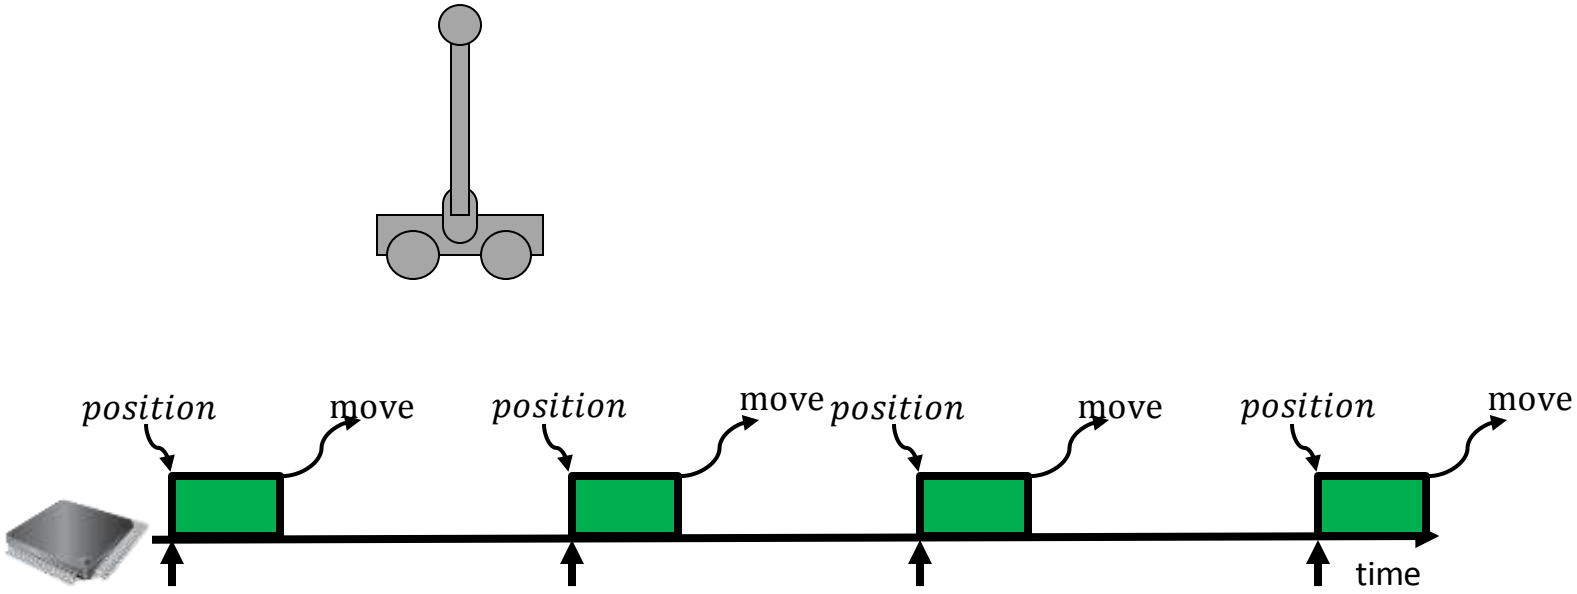

Periodically: **Senses** physical process and **actuates** on it to keep it under control.

Periodic Execution Needs in Avionics

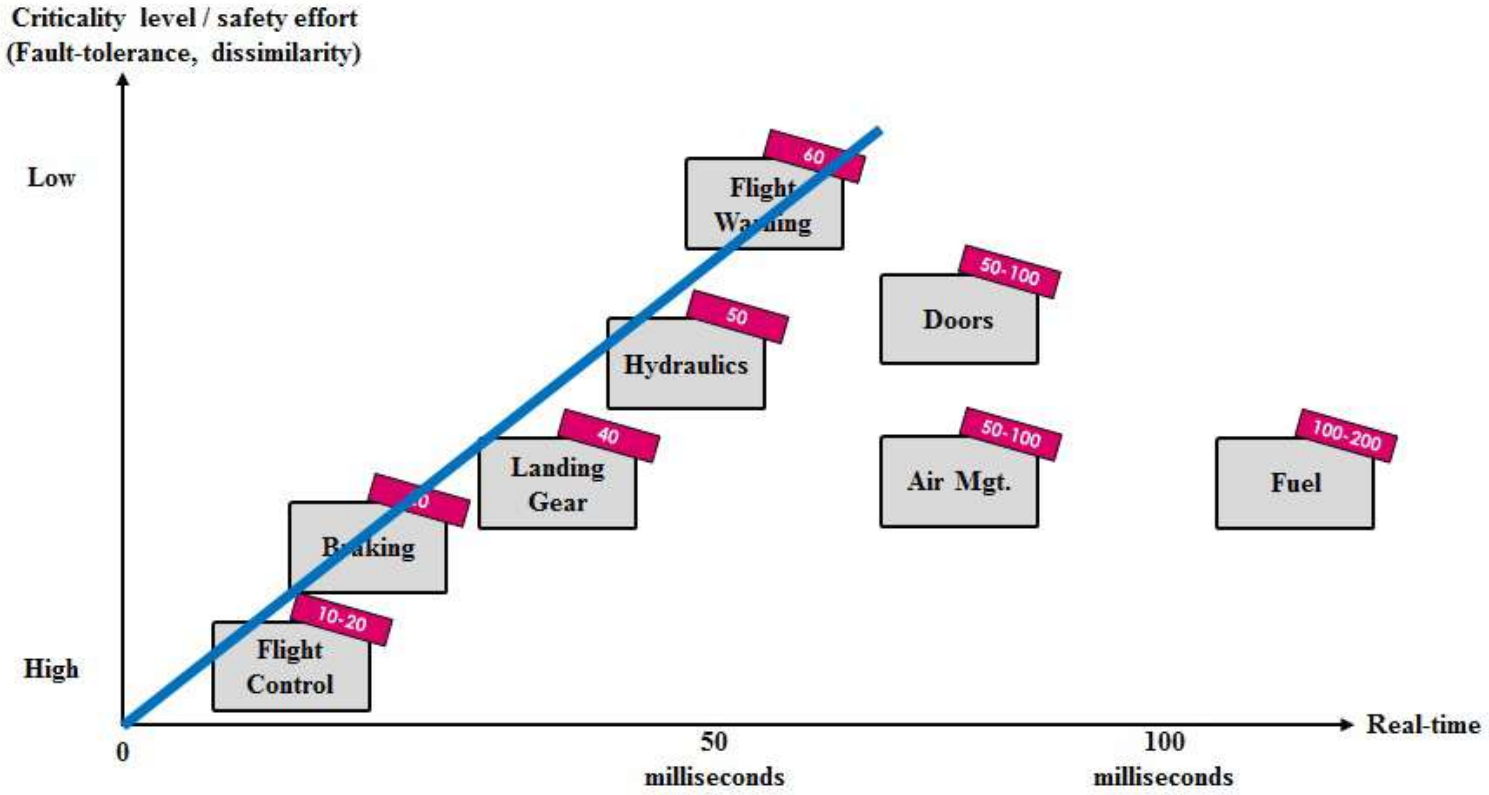

DOT/FAA/TC-16/51: Assurance of Multicore Processors in Airborne Systems

Single-Core Fixed-Priority Scheduling + Rate Monotonic

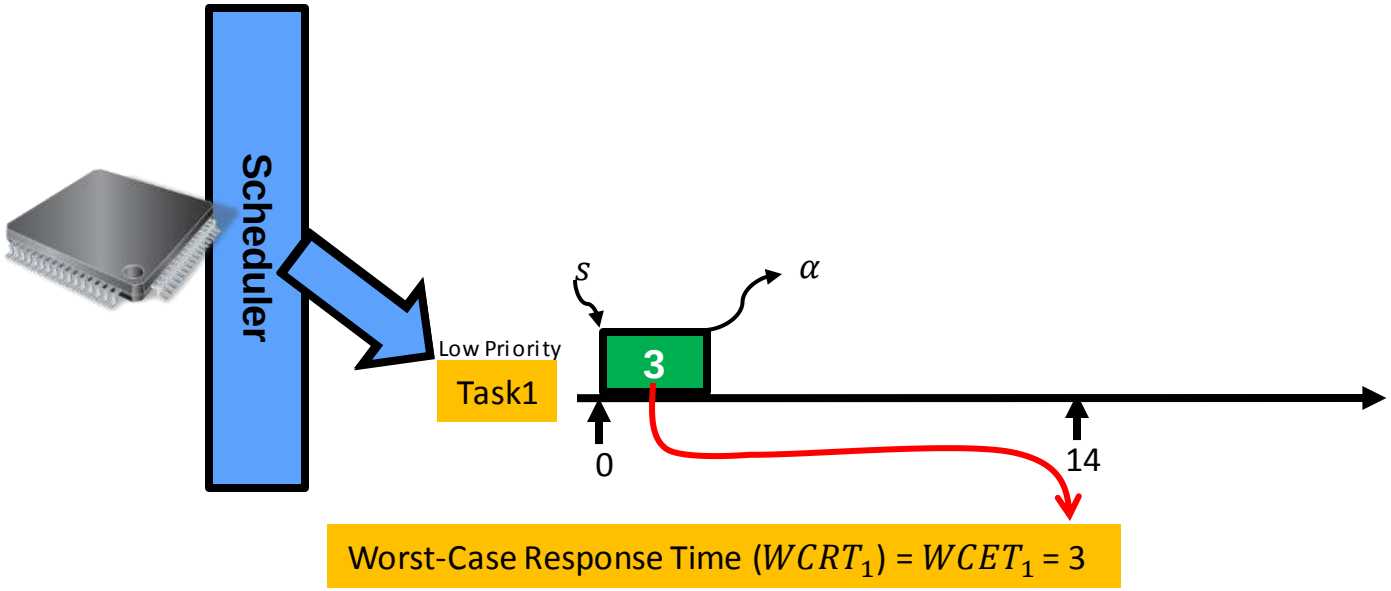

Icons credit: http://www.doublejdesign.co.uk

Single-Core Fixed-Priority Scheduling + Rate Monotonic

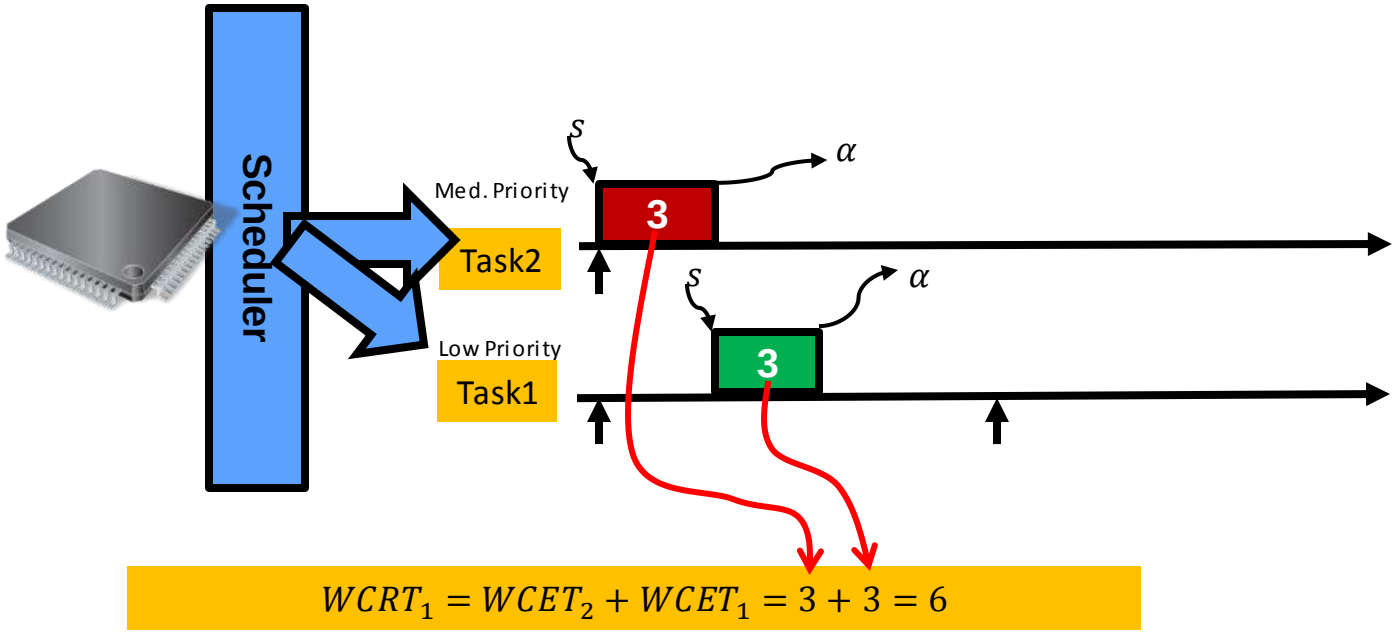

Icons credit: http://www.doublejdesign.co.uk

Single-Core Fixed-Priority Scheduling + Rate Monotonic

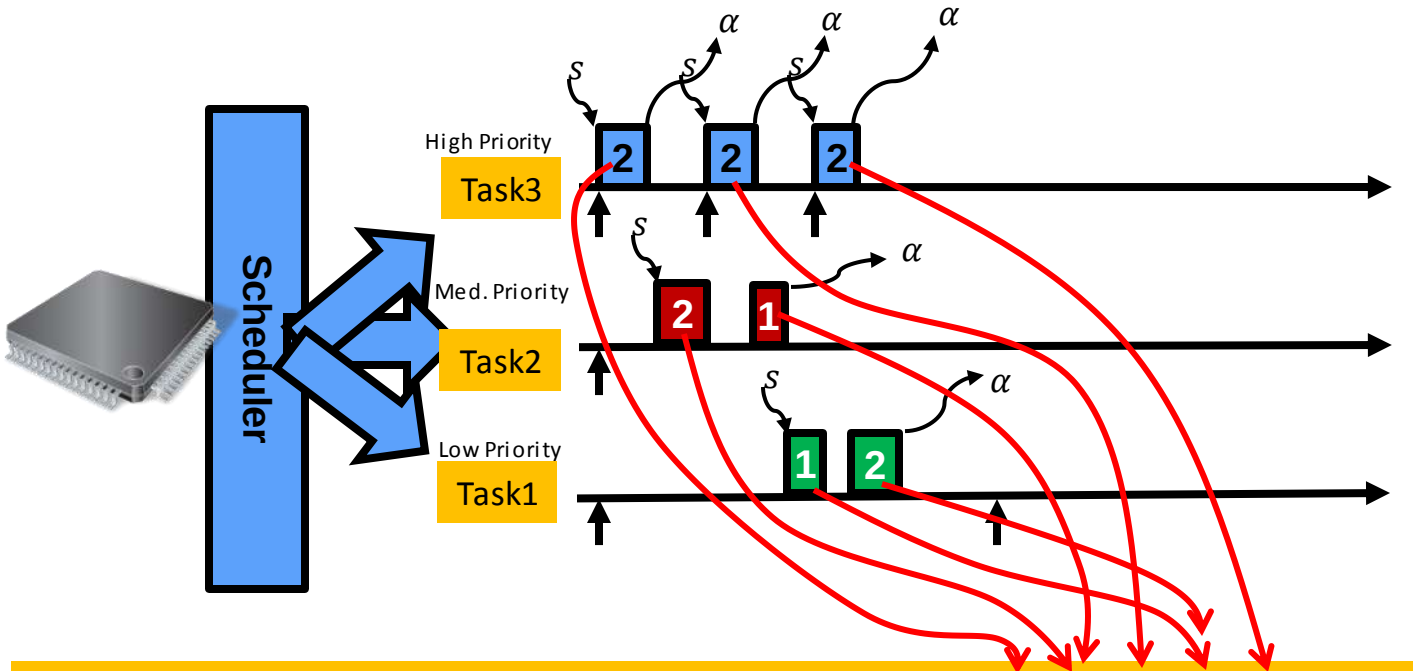

$$WCRT_1 = WCET_3 + WCET_2 + WCET_3 + WCET_1 + WCET_3 = 2 + 3 + 2 + 3 + 2 = 12$$

$$R_i = WCET_i + \sum_{j \in hp(i)} \left\lceil \frac{R_i}{T_j} \right\rceil WCET_j$$

Icons credit: http://www.doublejdesign.co.uk

Dual Core: Prevent delays from Task3?

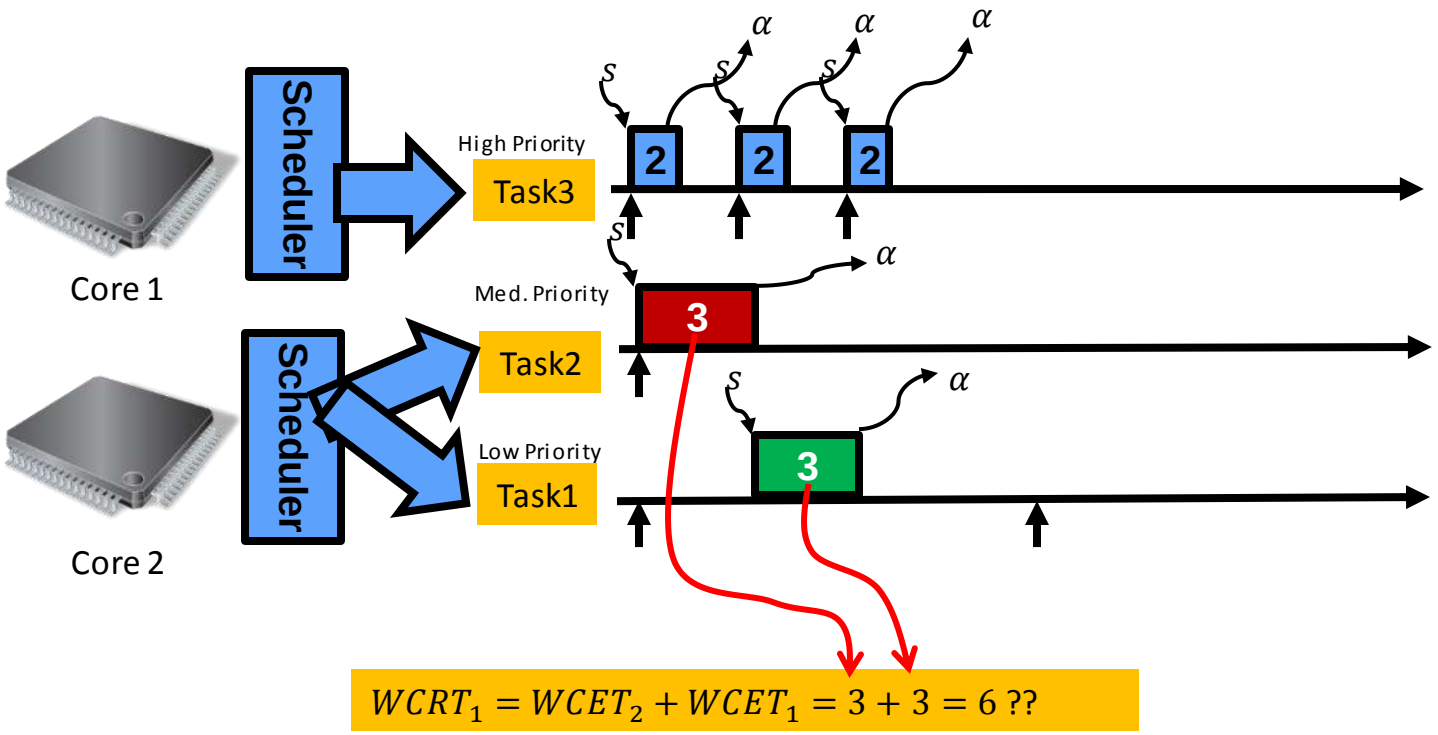

Icons credit: http://www.doublejdesign.co.uk

Dual Core: Prevent delays from Task3?

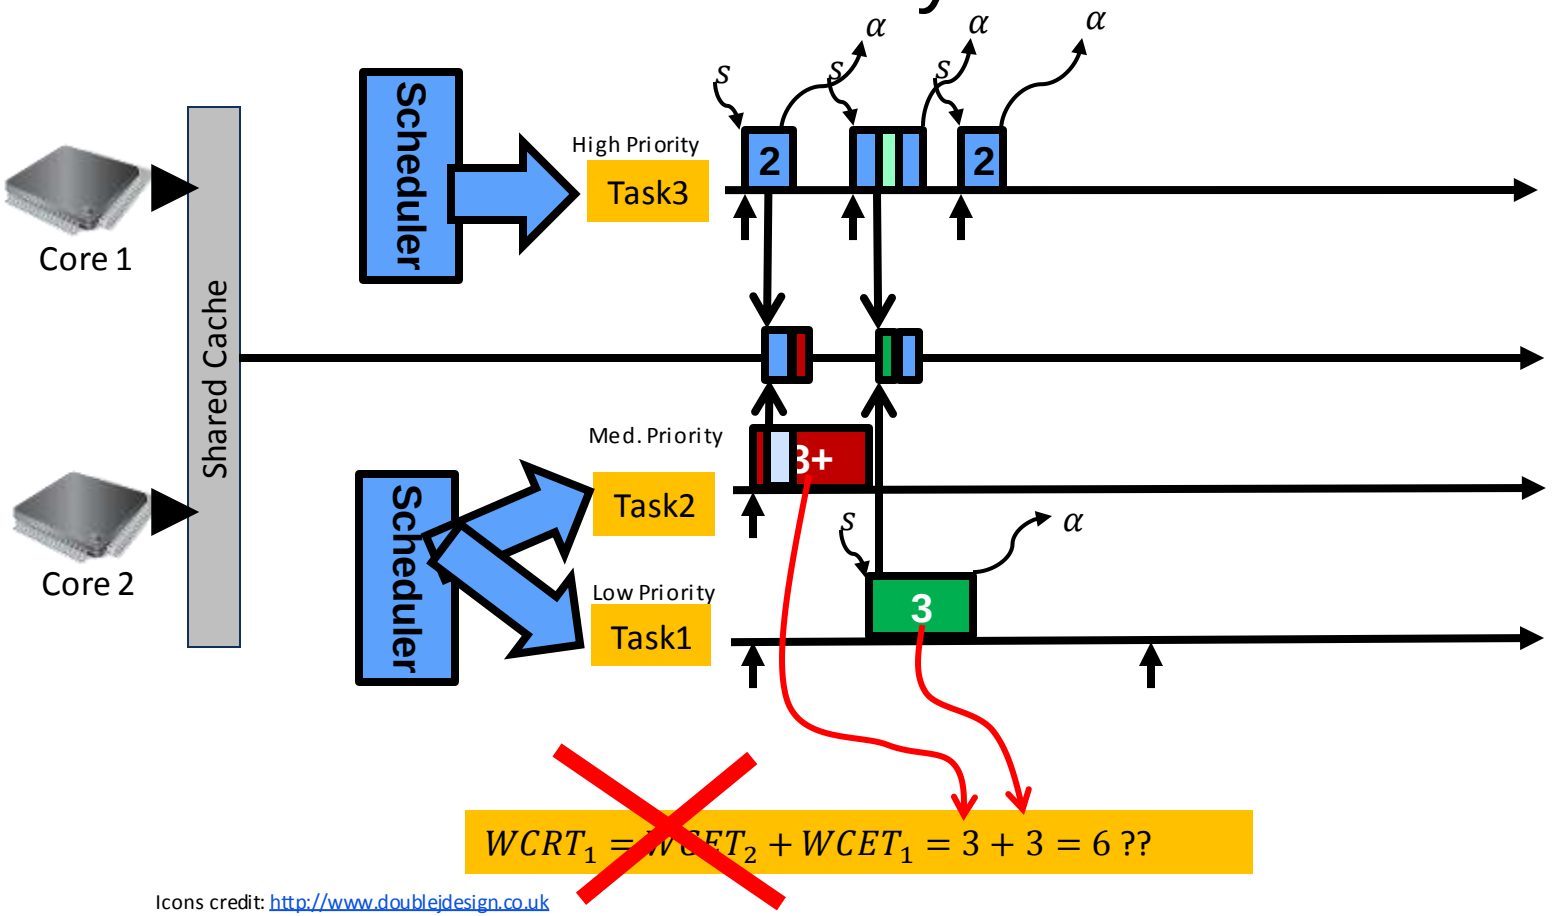

Icons credit: http://www.doublejdesign.co.uk

Dual Core: Prevent delays from Task3?

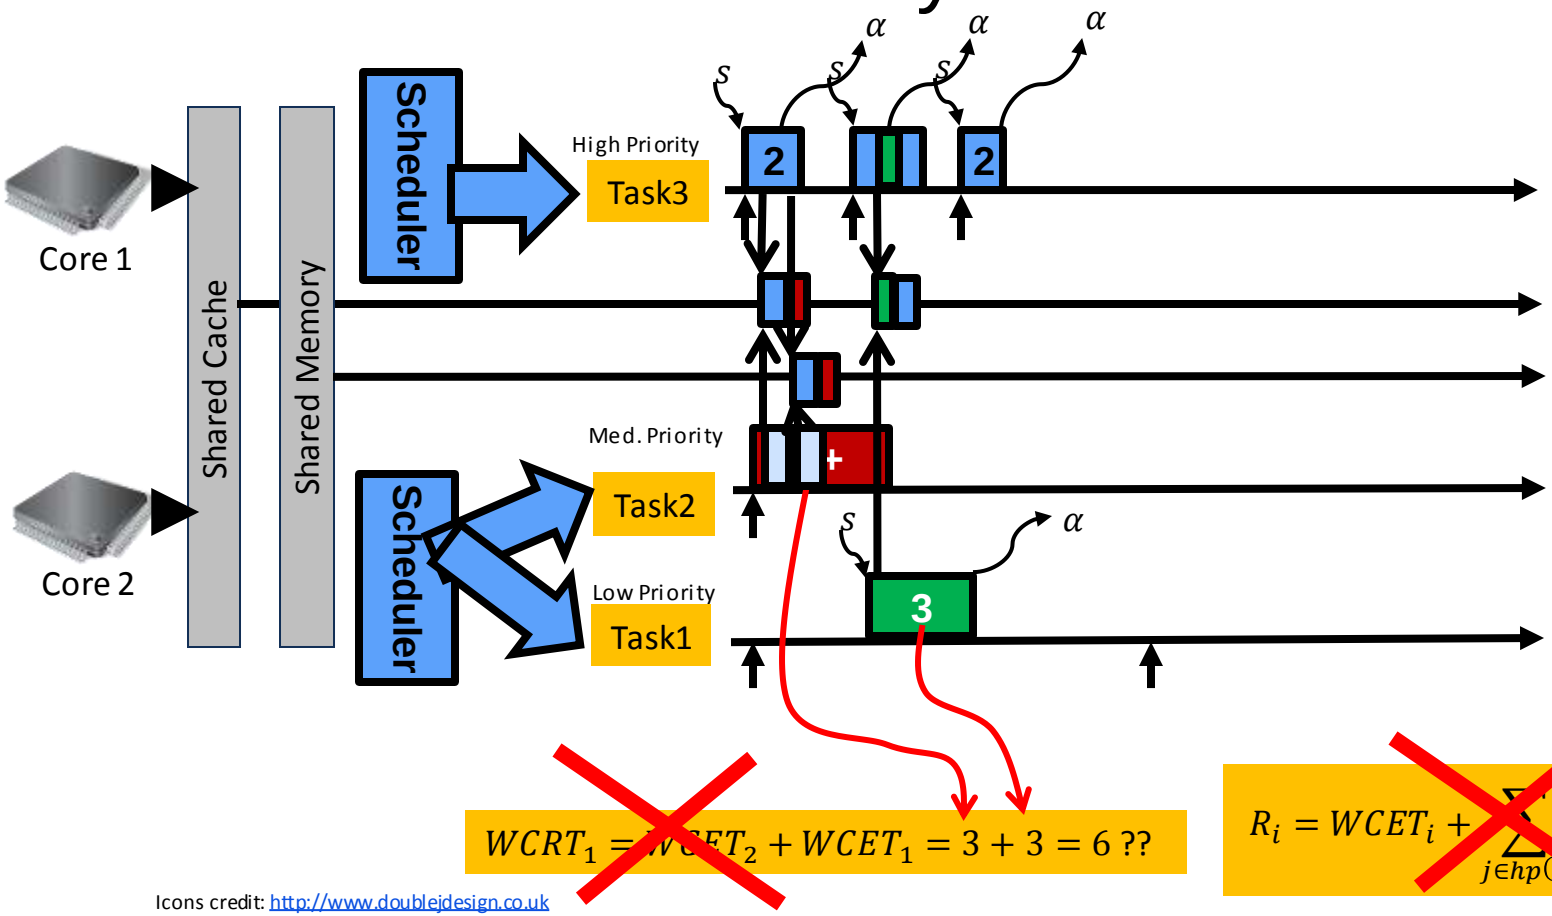

Icons credit: http://www.doublejdesign.co.uk