



Provably Correct Memory Management

Reto Achermann

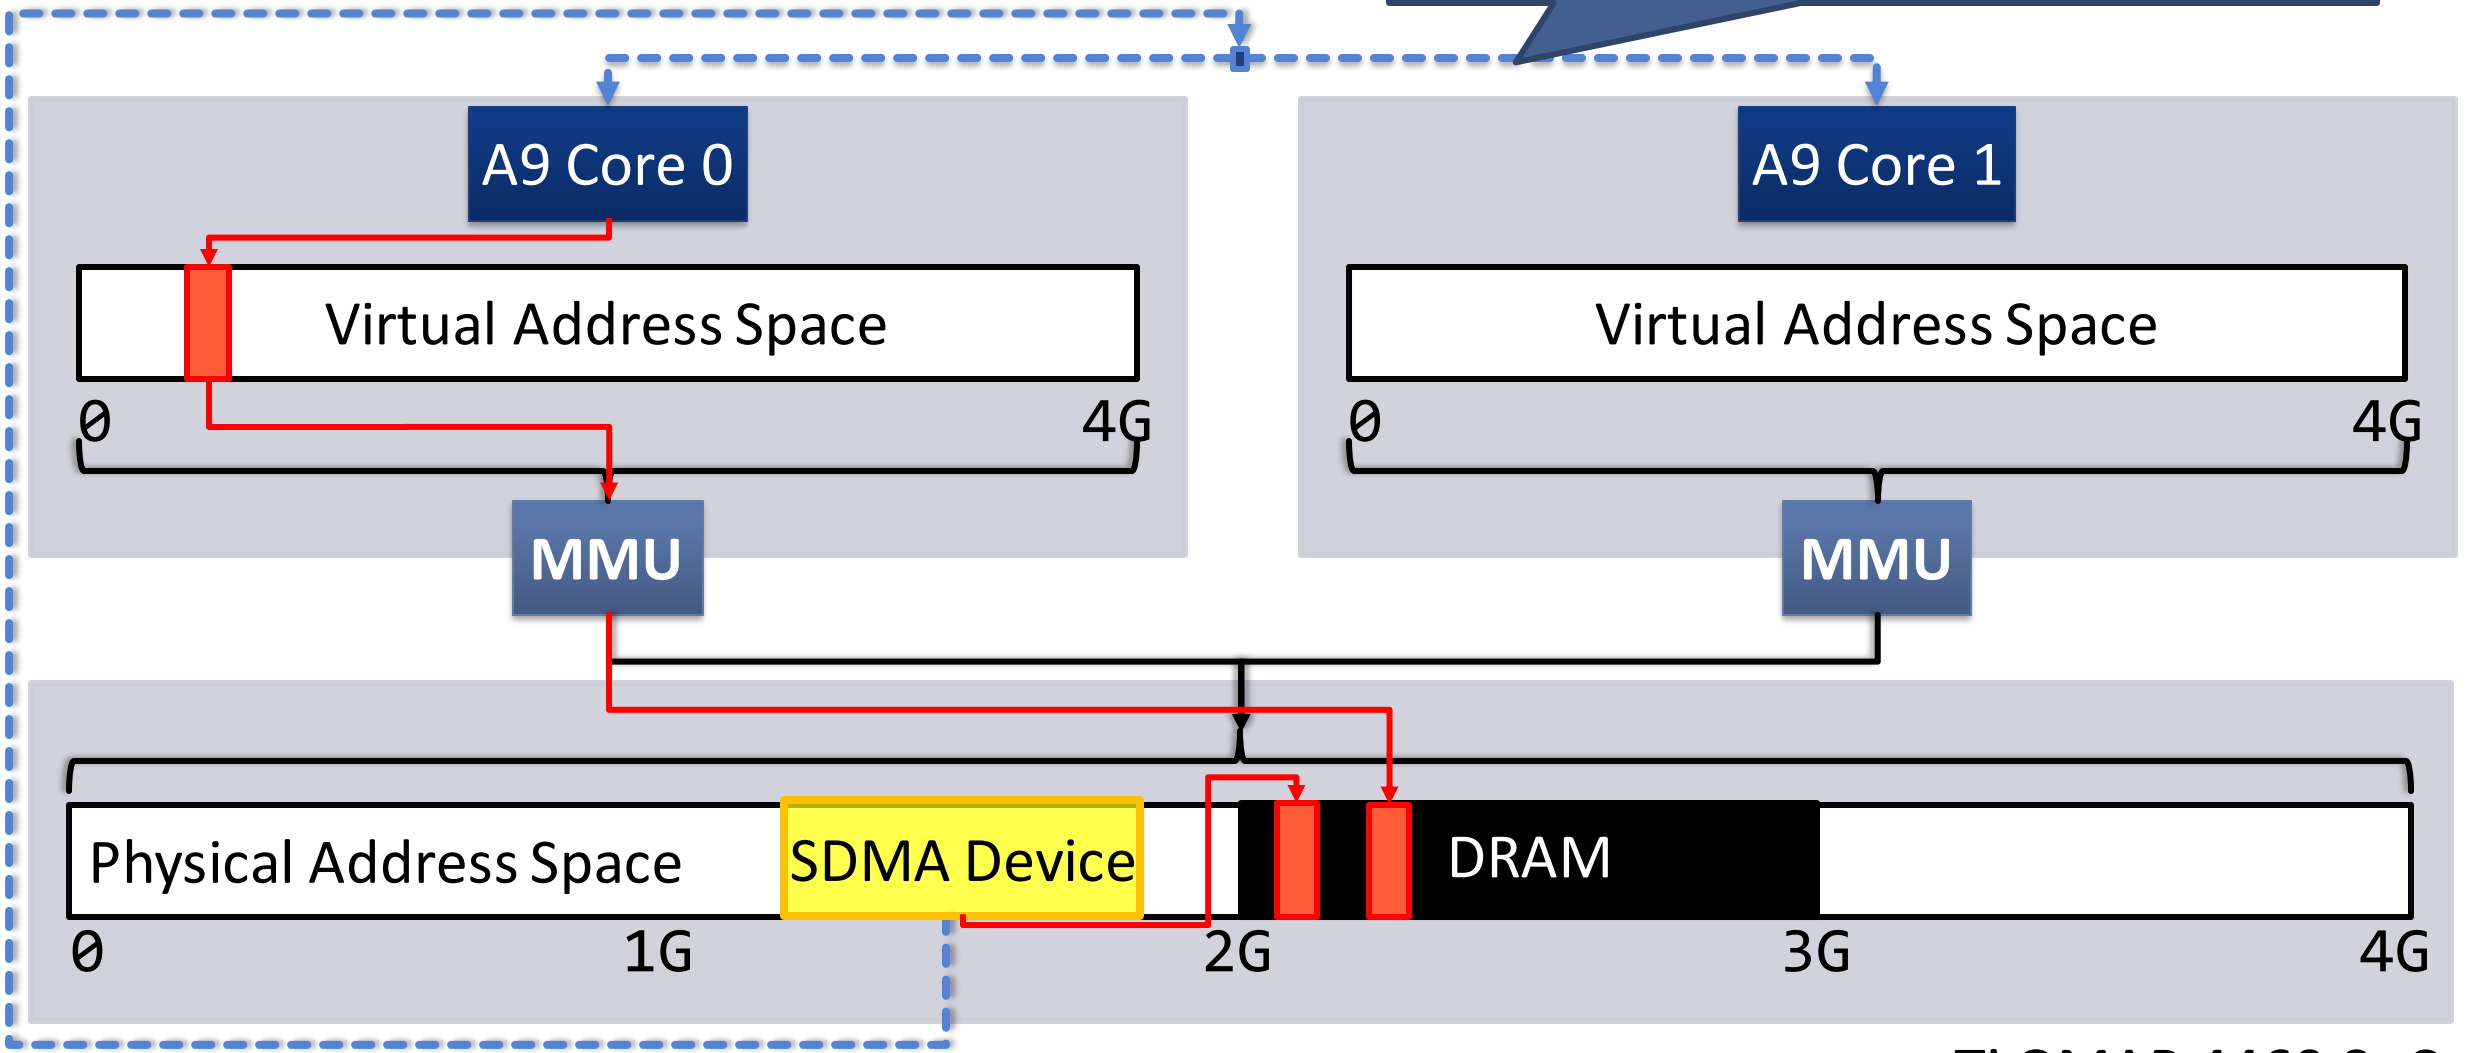
11th EuroSys Doctoral Workshop (EuroDW'17), Belgrade

Systems Group, Department of Computer Science, ETH Zurich



The naïve view of today's systems

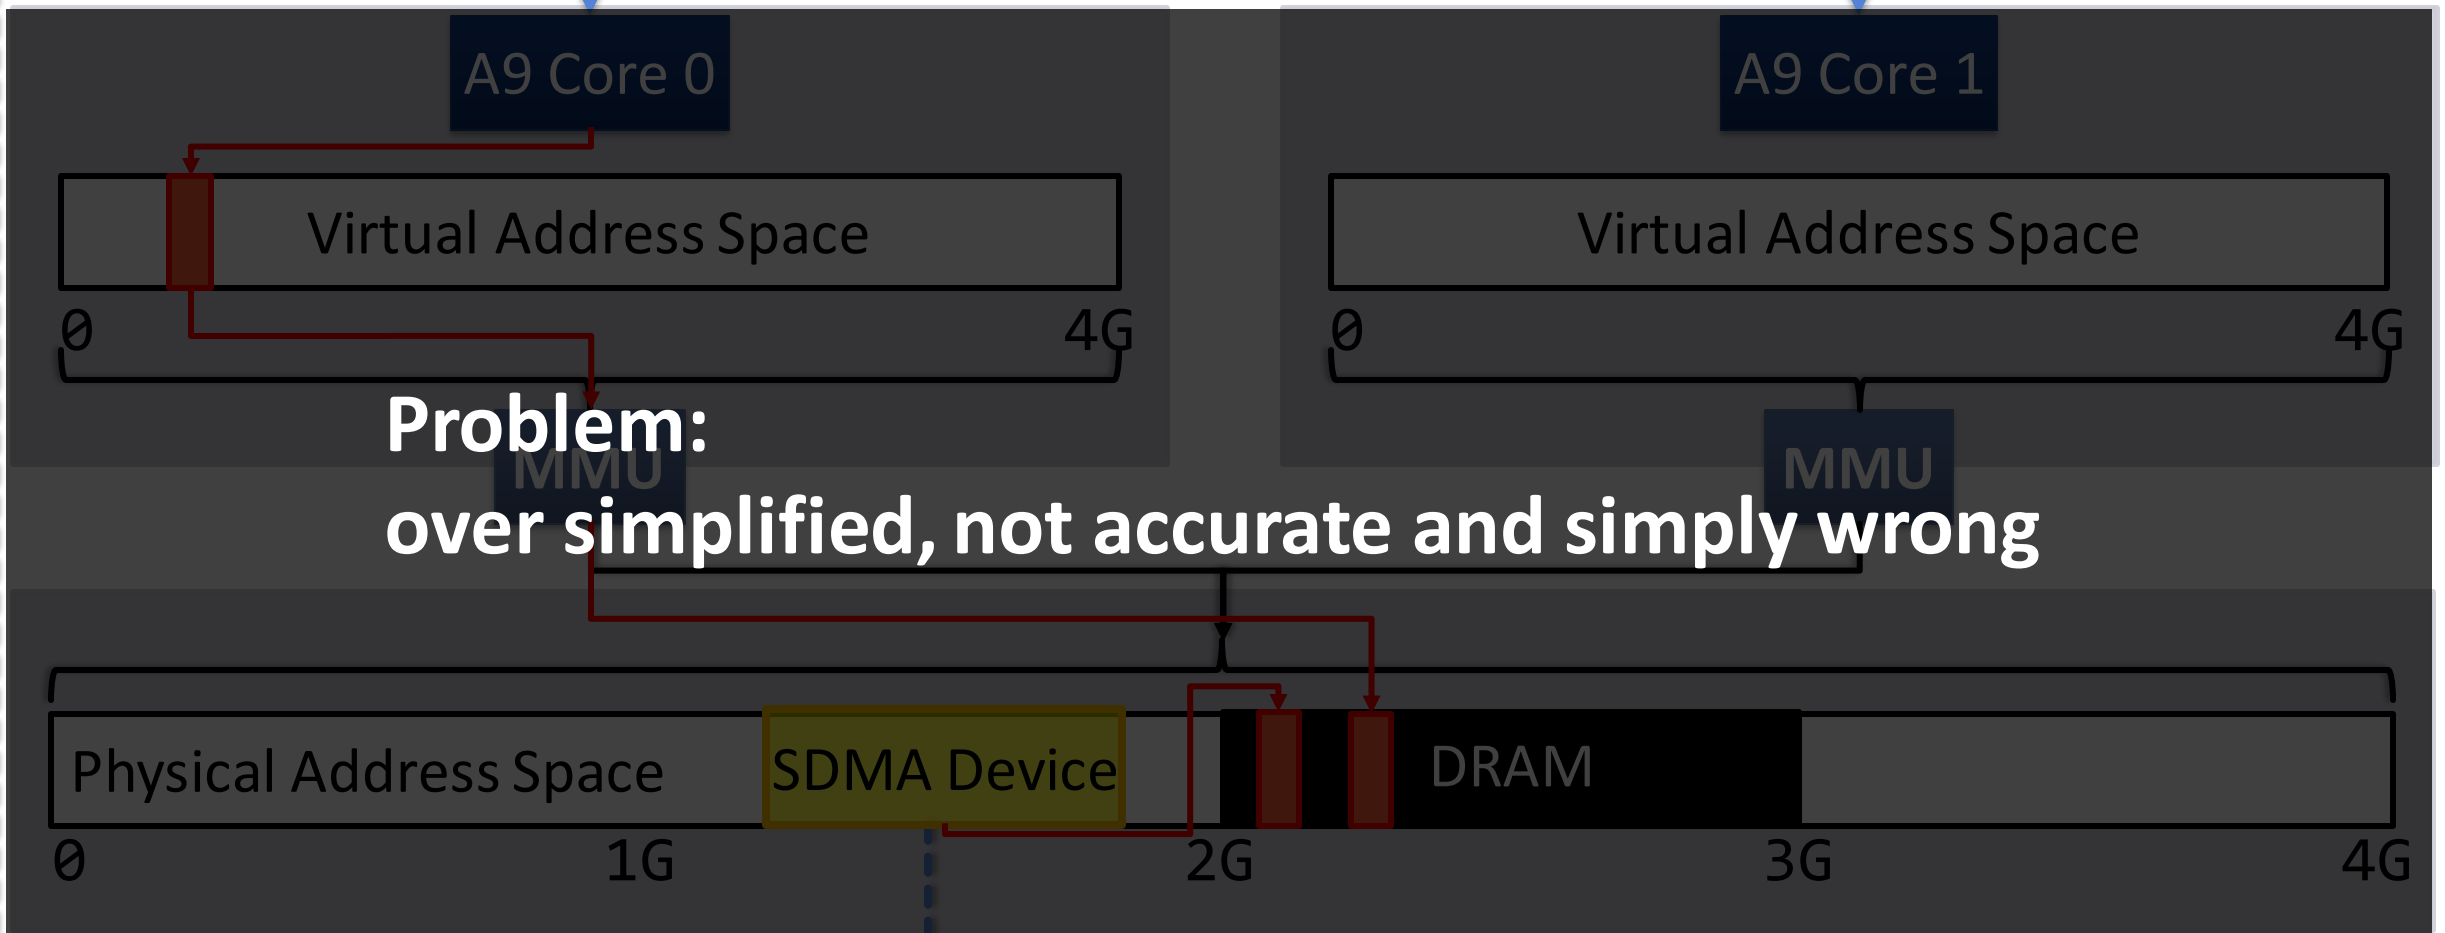
Lukas Humbel "Formalizing Interrupt Routing"



Ti OMAP 4460 SoC

The naïve view of today's systems

Lukas Humbel "Formalizing Interrupt Routing"



Ti OMAP 4460 SoC

Reality: The devil is in the details

Your mobile phone... 5-10 years ago!

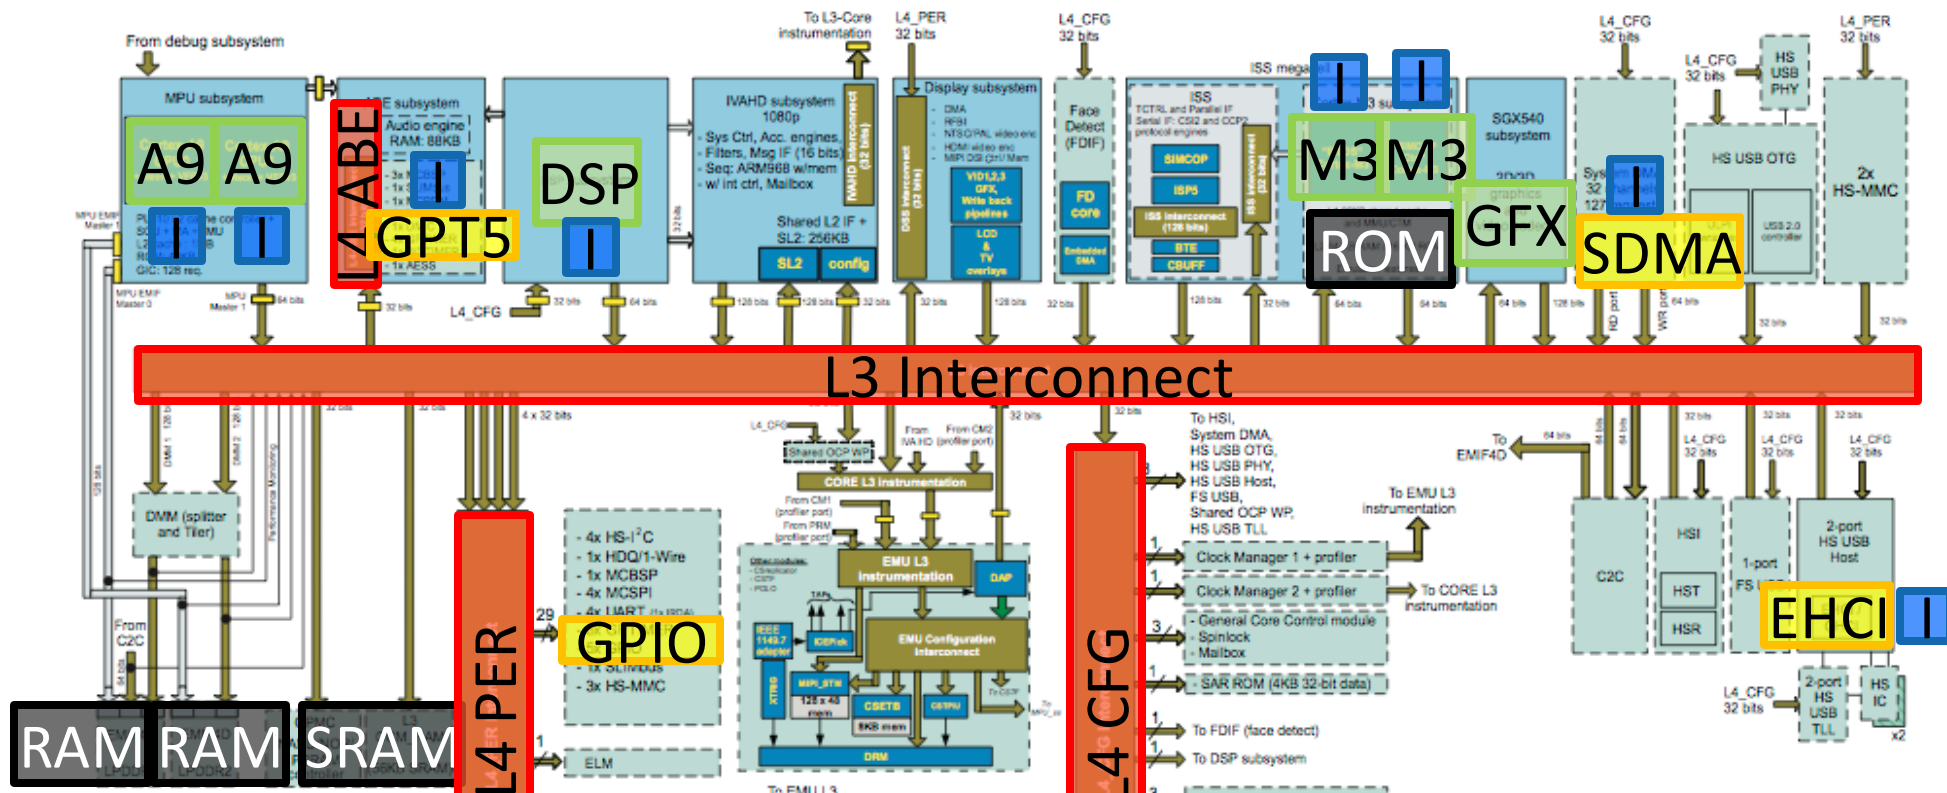
6+ heterogeneous cores

Private and shared memory

5+ Interconnects

Devices attached to different interconnects

Complex interrupt subsystem



RAM RAM SRAM

L4 PER
GPIO


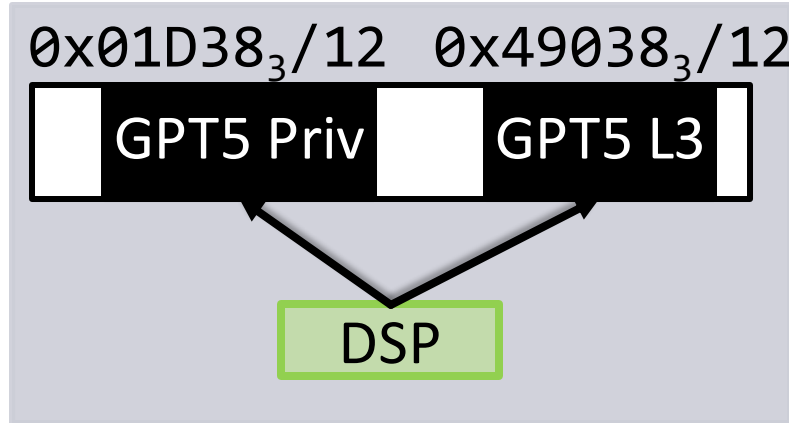
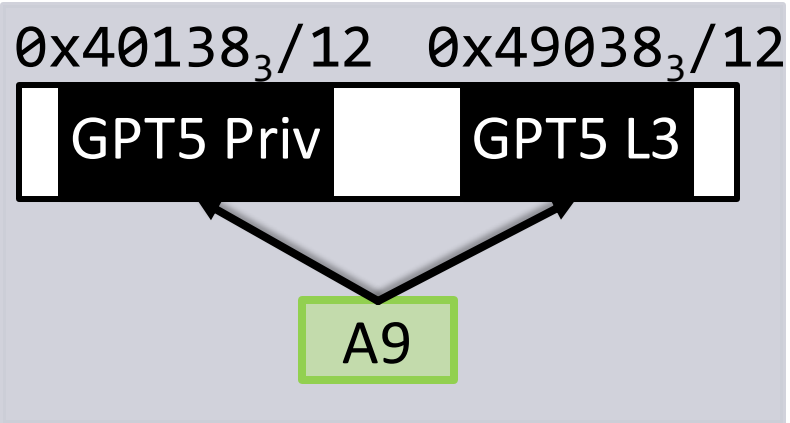
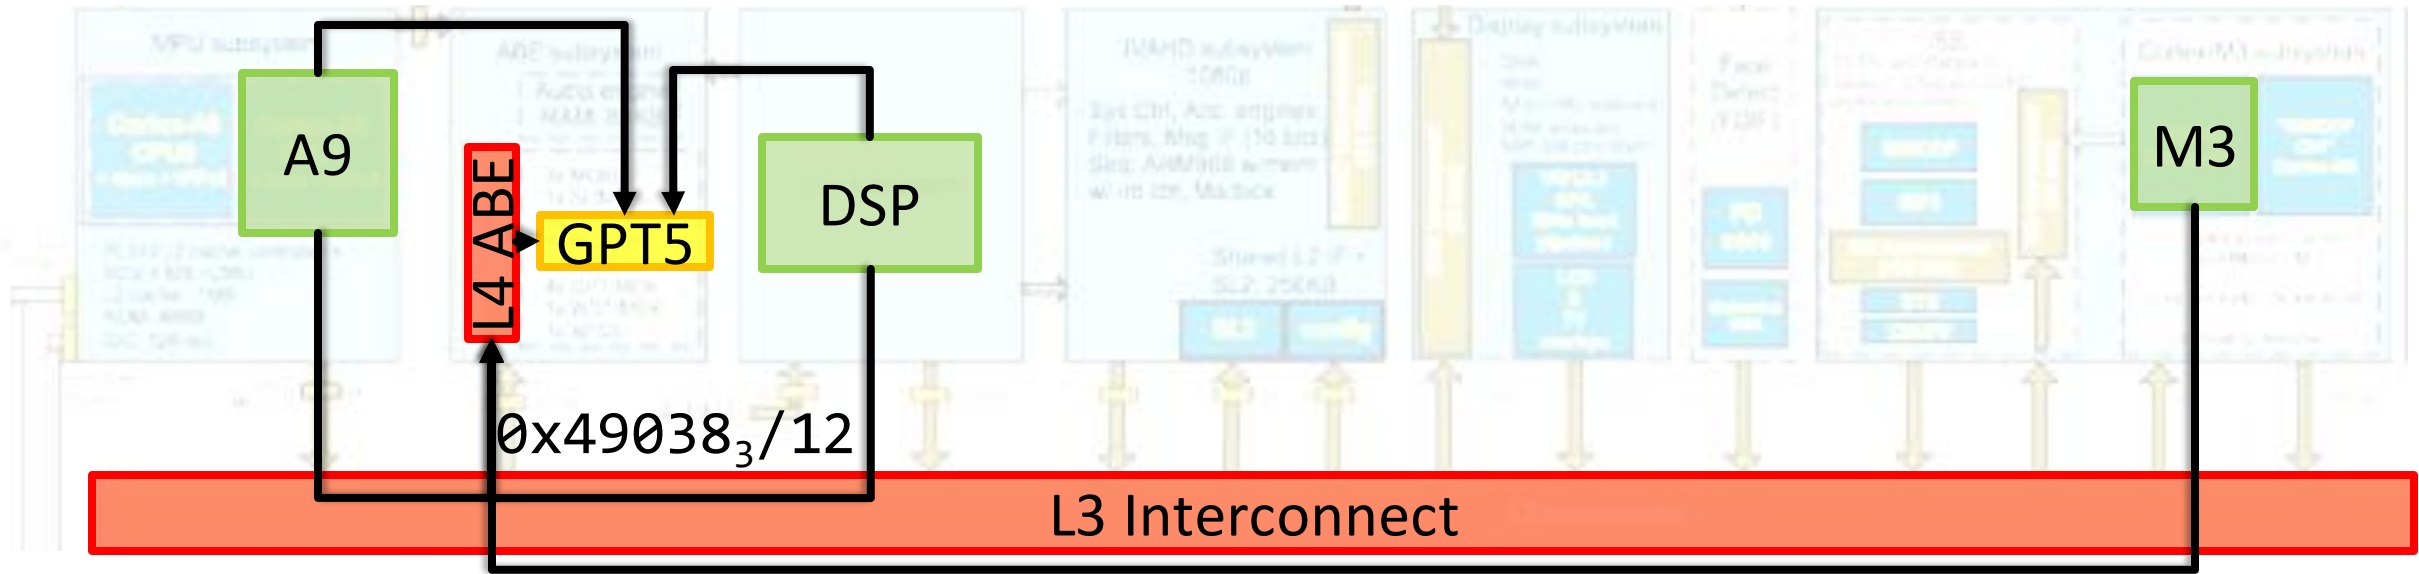
L4 CFG

Takeaway: There are many **details** that are not captured by the naïve representation!


OMAP 4460 SoC,
Technical Reference Manual



There is NO uniform view of the system



There is NO uniform view of the system

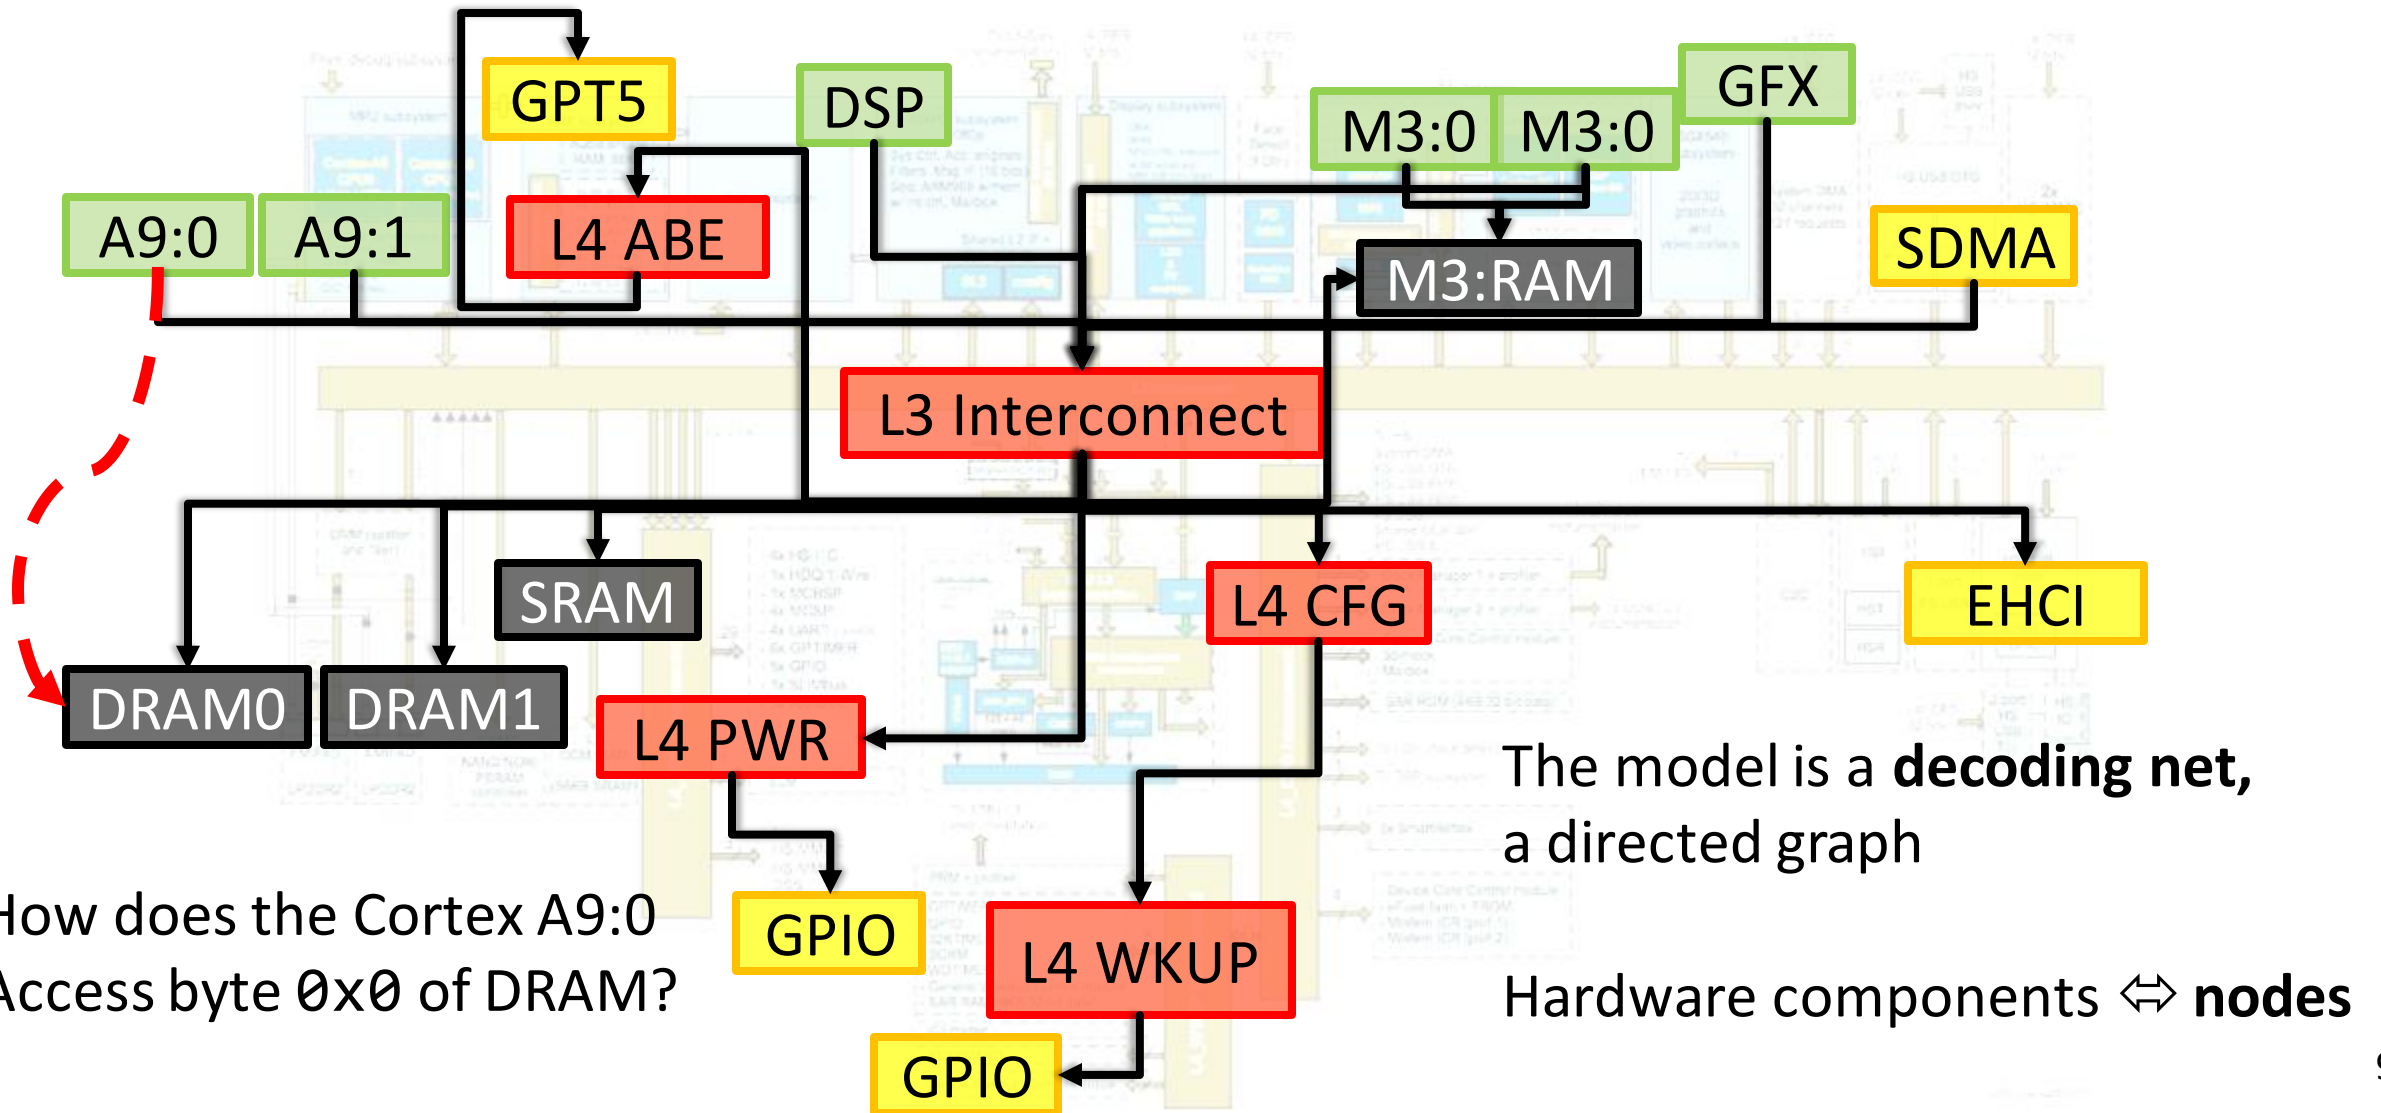


Why do we need a formal model for memory accesses ?

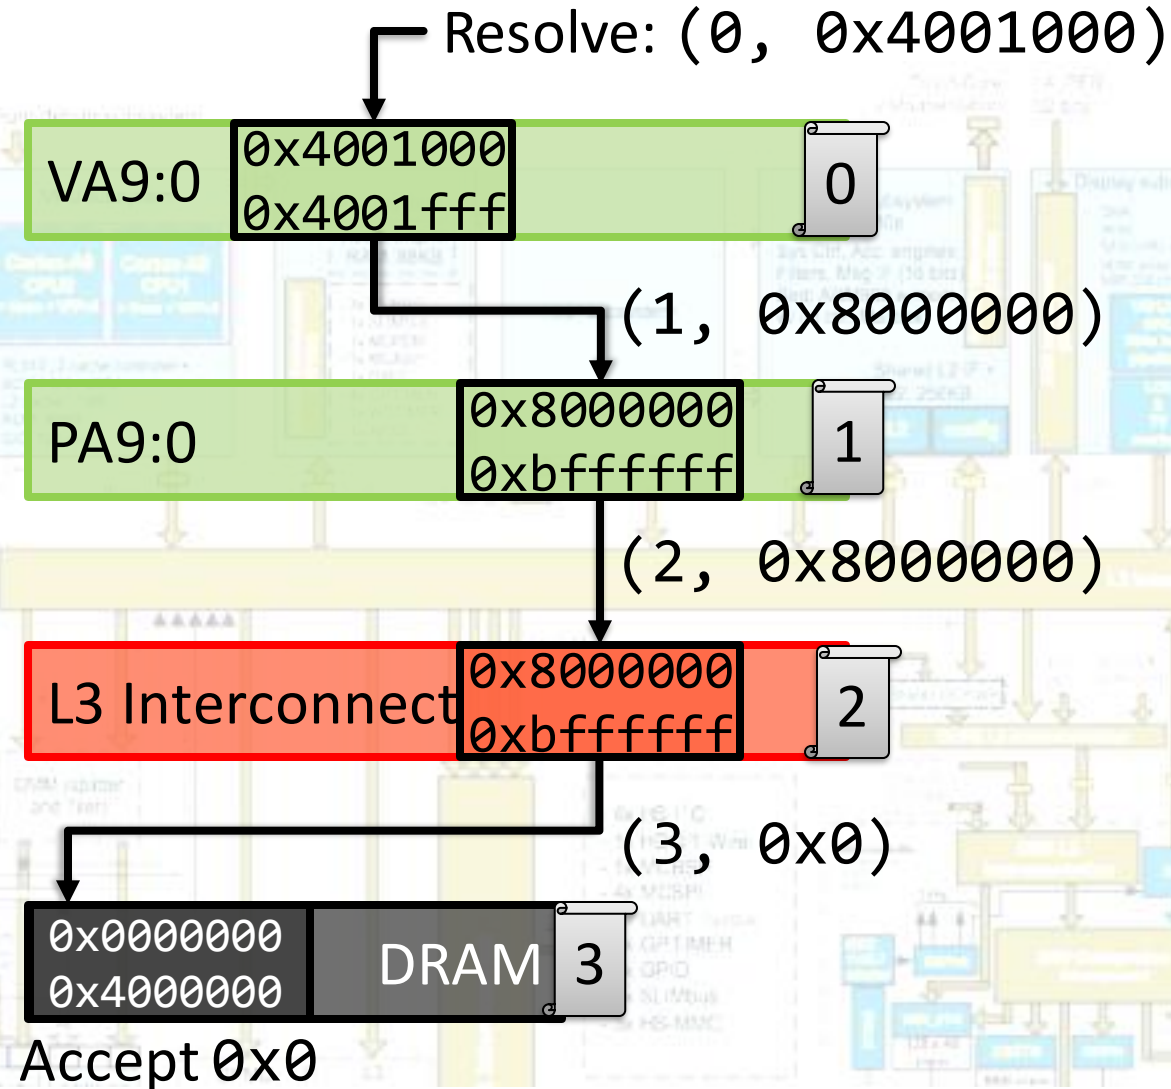
- We **build** systems and want to write **correct** systems code
- Experience from the Barrelfish operating system:
dealing with this complexity every day.
e.g. PCI programming, heterogeneity, resources, devices, new platforms
- Problem:
 - Current abstractions make the **wrong** assumptions
 - System software verification requires a sound system hardware description



A partial decoding net for the OMAP4460



Modelling the access to byte 0x0 of DRAM from an A9 core



Each node has a label

Resolve a **name**
(node, address)

Model of one
particular, static
configuration state




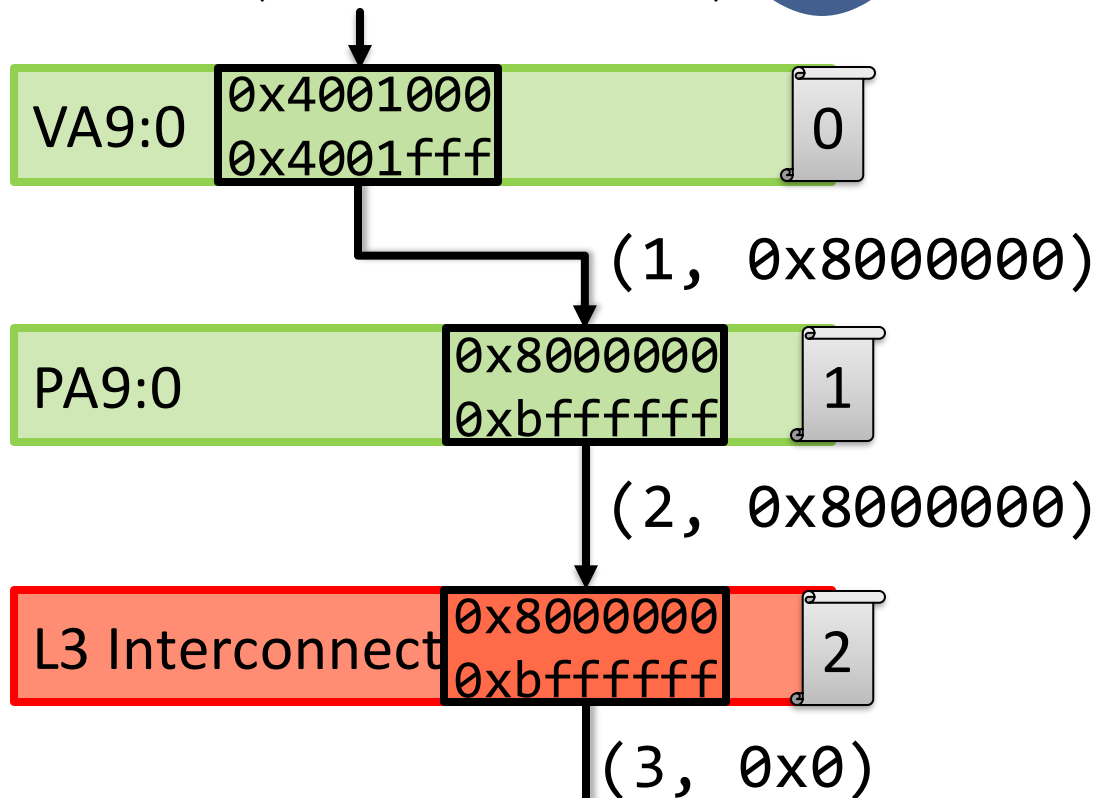
Nodes have **two properties**:


accept: $node \rightarrow \{N\}$

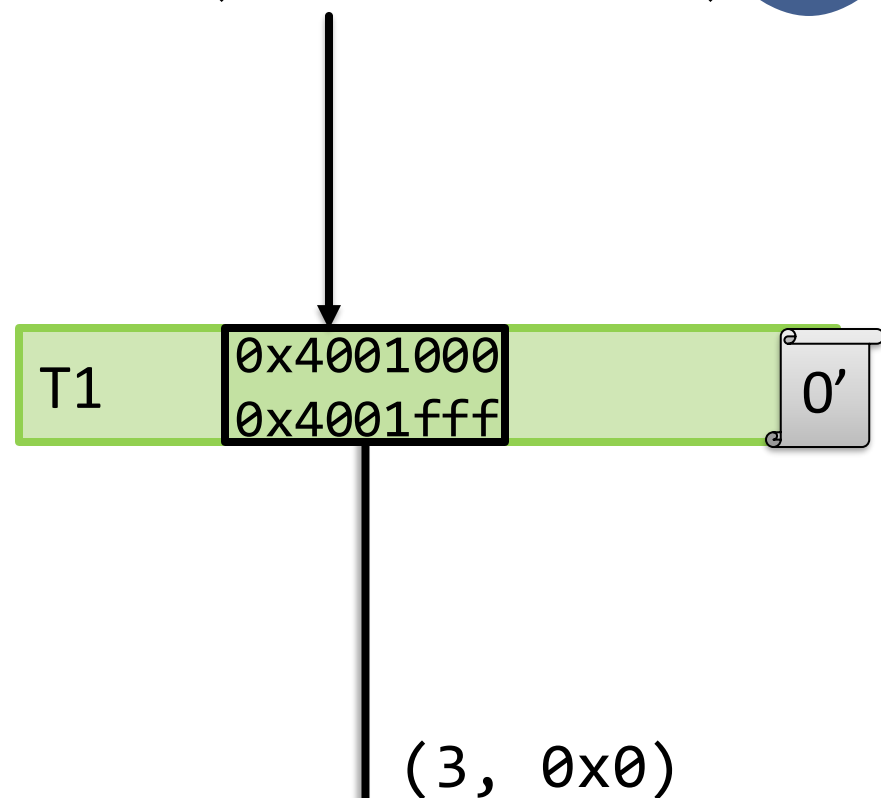
translate: $node \rightarrow N \rightarrow \{name\}$

Flattening using view equivalence preserving operations

Resolve: $(0, 0x4001000)$ 





Resolve: $(0', 0x4001000)$ 

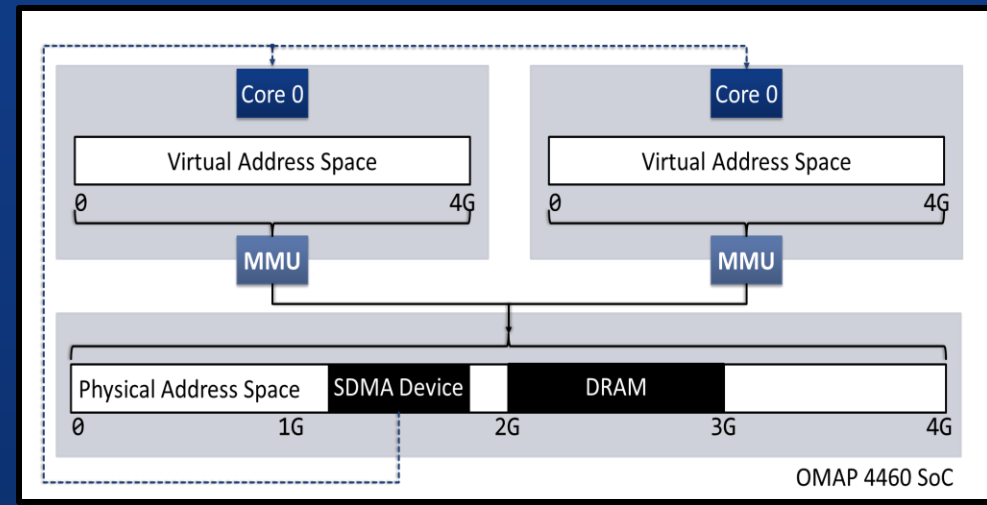
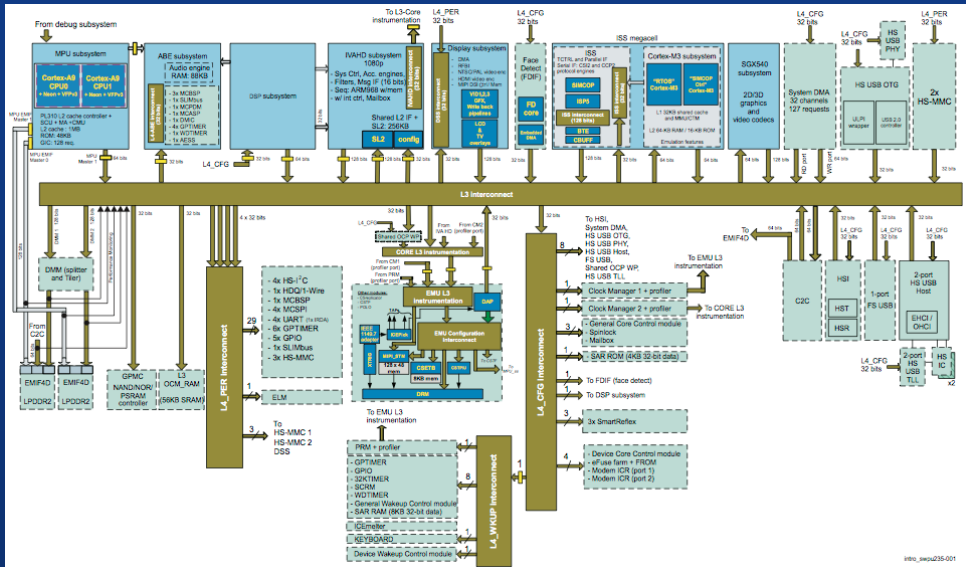


$0x0000000$
 $0x4000000$ **DRAM**
 Accept $0x0$

Flattening using view equivalence preserving operations

Resolve: $(0, 0x4001000)$ 

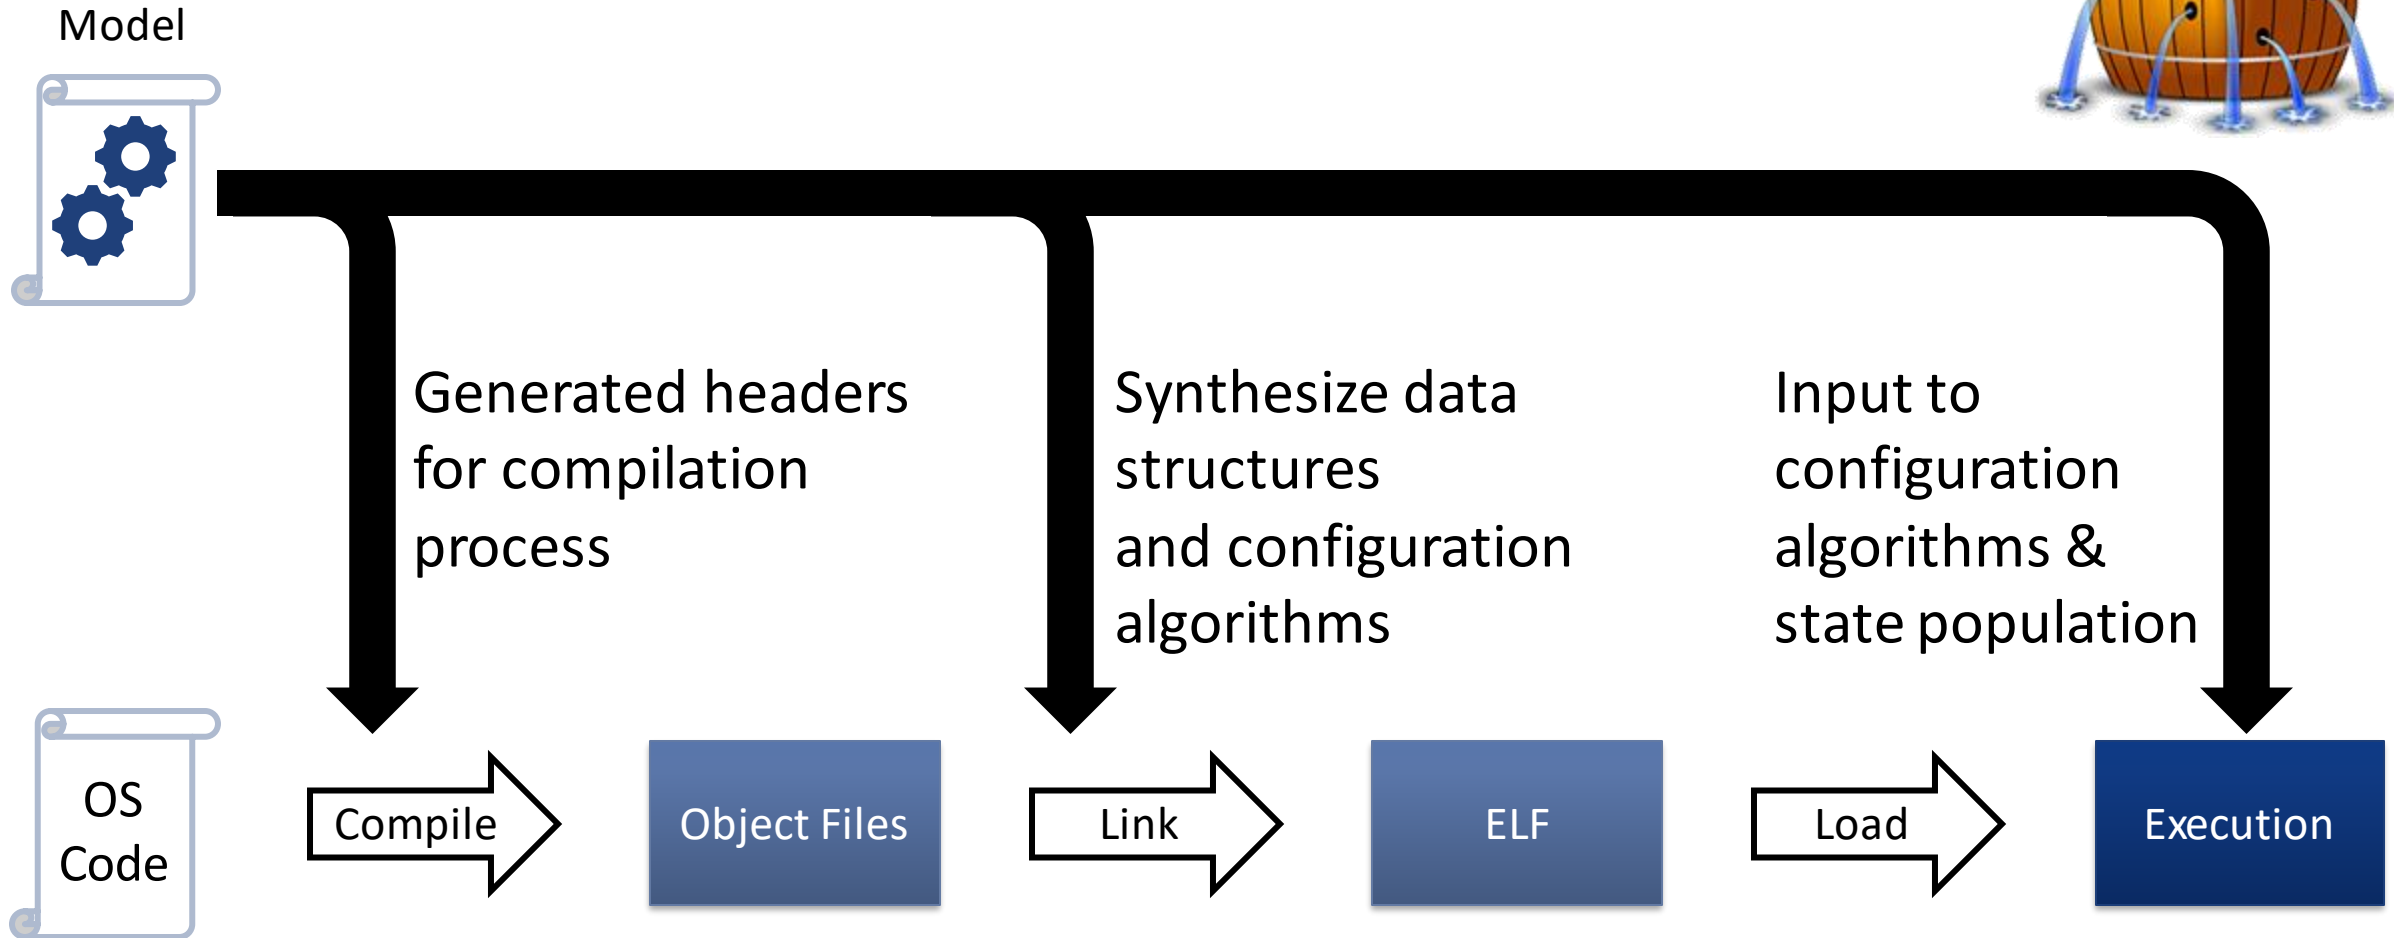
Resolve: $(0', 0x4001000)$ 



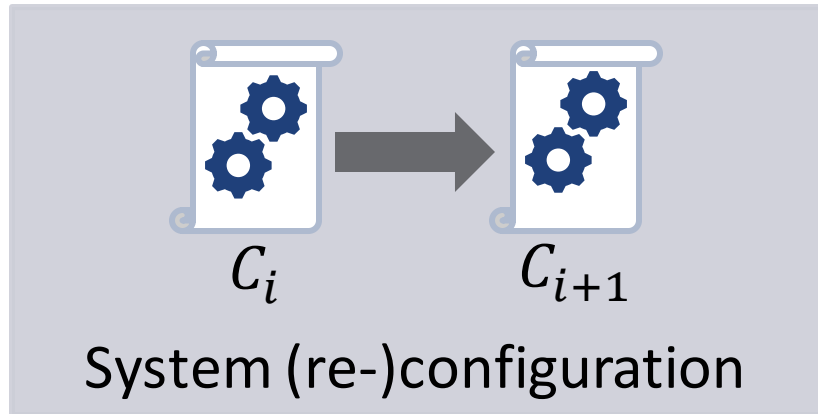
For ONE observer the flattened representation is equivalent to the textbook abstraction

Accept $0x0$

Ongoing work: Using model output at compile and run time



Ongoing Work: Model applications



- Generate system configuration from the model:
 - Kernel page tables
 - Initial capabilities
- **Synthesize** configuration algorithms
- **Transition** between configurations without violation of **invariants**
- **Constraints** on memory accesses

Future work: Model refinements



Distinction of Read/Write
accesses

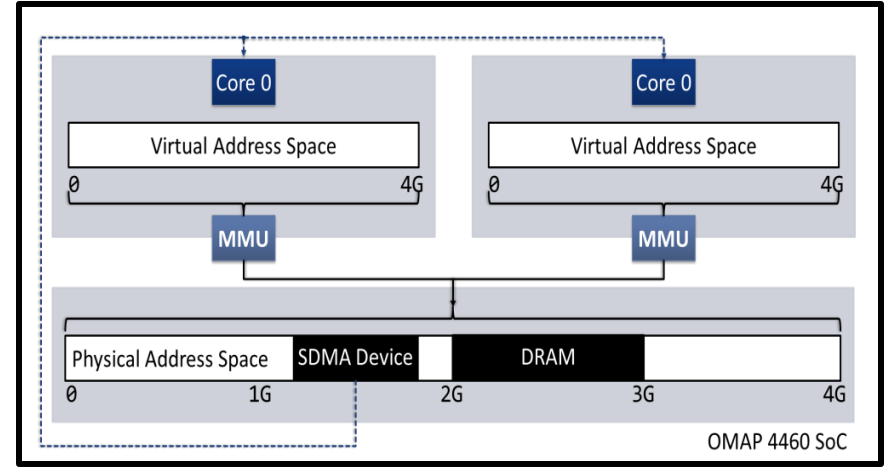
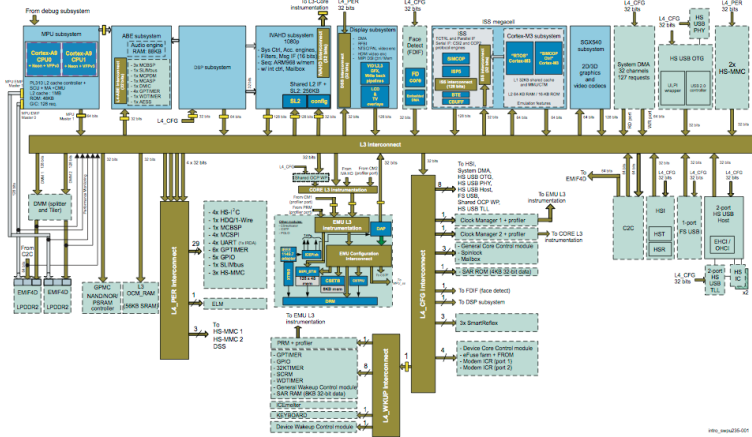
- Reads / writes have different semantics
- Write only / read only regions



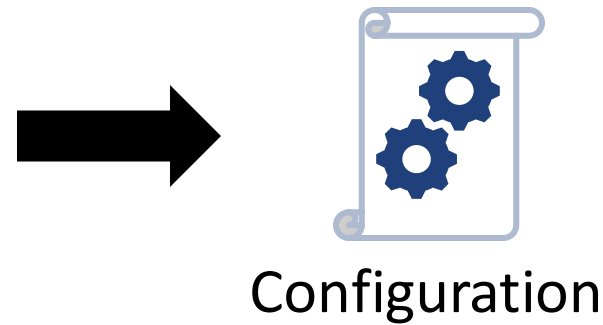
Expressing performance
characteristics

- Basis for a performance model.
- Resource allocation & scheduling

Summary



$V_{A9:0}$ is map [20000₃/12 to $P_{A9:0}$ at 80000₃] $V_{A9:1}$ is map [20000₃/12 to $P_{A9:1}$ at 80000₃]
 $P_{A9:0}, P_{A9:1}$ are map [40138₃/12 to GPT at 0] over L3 V_{DSP} is over P_{DSP}
 P_{DSP} is map [1d3e₃/12 to GPT at 0] over L3 $L2_{M3}$ is map [0₃₀ to L3 at 80000₃]
 V_{M3}, V_{M3} are over $L1_{M3}$ $L1_{M3}$ is map [0₂₈ to MIF]
 RAM_{M3} is accept [55020₃/16] $L4$ is map [49038₃/12 to GPT at 0]
 ROM_{M3} is accept [55000₃/14] GPT is accept [0/12]
 MIF is map [0 – 5fffff to $L2_{M3}$, 55000₃/14 to RAM_{M3} , 55020₃/16 to ROM_{M3}]
 $L3$ is map [49000₃/24 to L4 at 40100₃, 55000₃/12 to MIF] accept [80000₃/30]



Configuration