

New Projects at ETH Systems (and why you should come do a Ph.D.)



The Systems Group at ETH



About ETH



- Swiss Federal Polytechnic
- Founded 1854 (with EPFL)
- In central Zürich
- CHF1.8B/year
- Famous Alumni:
 - Albert Einstein
 - Wolfgang Pauli
 - Gernot Heiser ;)

Industrial Collaboration and Support

ORACLE®

ARM



A Selection of New/Ongoing Projects


- Barrelfish
- Runtime Verification
- Formal Hardware Modelling
- Strymon
- Enzian

The State of the Fish



- 7 architectures: OMAP44xx, ARMv7/GEM5, X-Gene 1, ARMv8/GEM5, Xeon Phi, x86-64, x86-32
- 42 applications + 51 test apps
- 9 languages
- 32 committers
- 9 years old
- > 1.1M lines of code

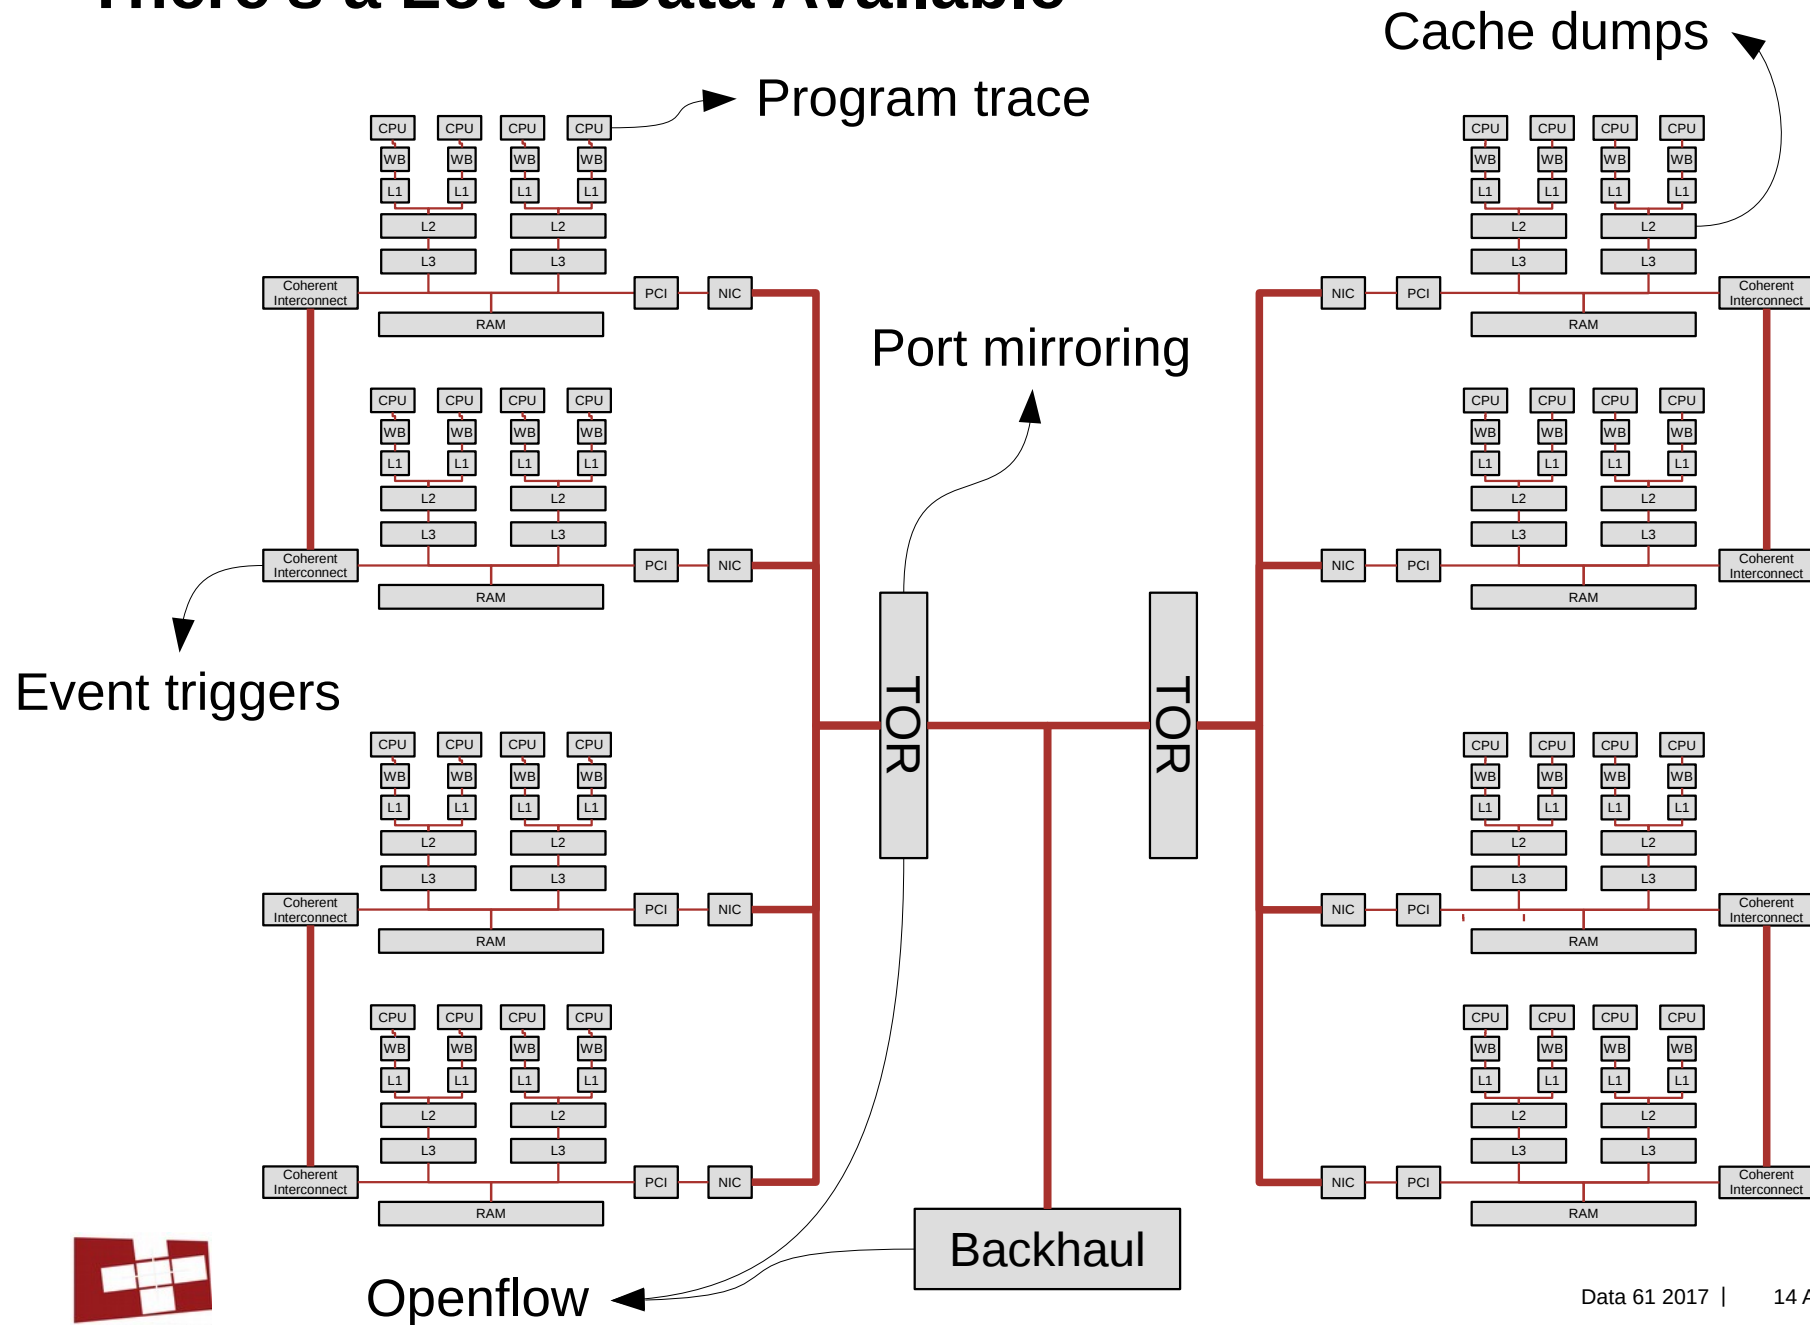
We're Building a *Large Program Collider*



Images: CERN; Chaix & Morel et associés

Collide *instructions* at $0.99c$, and observe the decay products.


There's a Lot of Data Available




ARM High-Speed Serial Trace Port

- Streams from the *Embedded Trace Macrocell*.
- Cycle-accurate control flow + events @ 6GiB/s+
- Compatible with FPGA PHYs.
- Well-documented protocol.
 - Aurora 8/10
- Available on ARMv8


Image: Teledyne Lecroy



HSSTP Testbench



Hardware Tracing for Correctness



Are HW operations right?


$\exists va.va \rightarrow pa$

```
unmap(pa);
cleanDCache();
flushTLB();
```


$\nexists va.va \rightarrow pa$

- Real time pipeline trace on ARM.
- Can halt and inspect caches.
- HW has “errata” (bugs).
- Check that it actually works!
- Catch transient and race bugs.

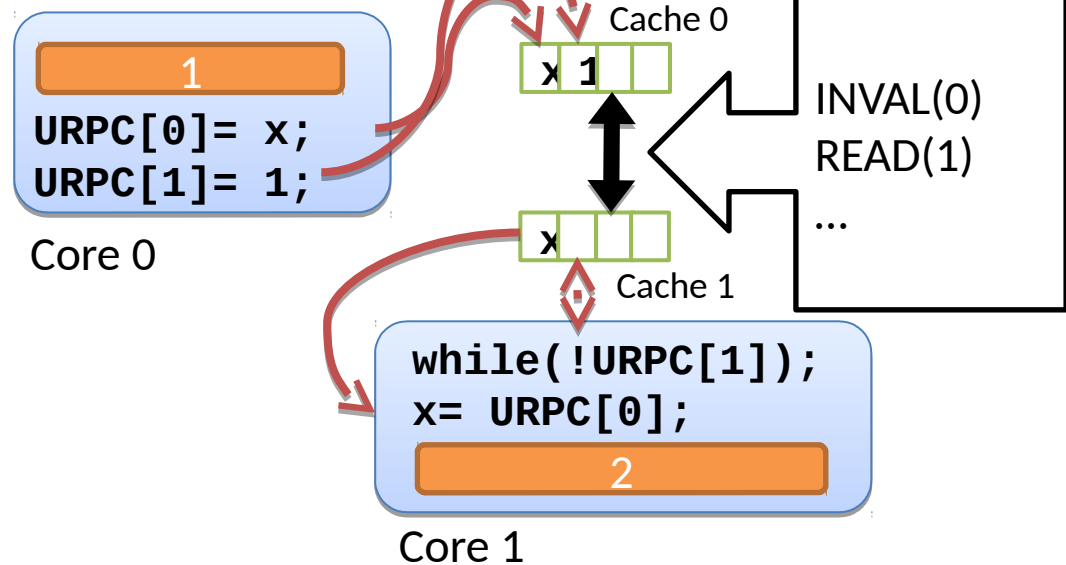

Hardware Tracing for Performance



- Should see N coherency messages.
- Do we?
 - The HW knows!




Is URPC optimal?



Properties to Check: Security

- Runtime verification is an established field.
- Lots of existing work to build on.
- What properties could we check efficiently?
- How could we map them to the filtering pipeline?



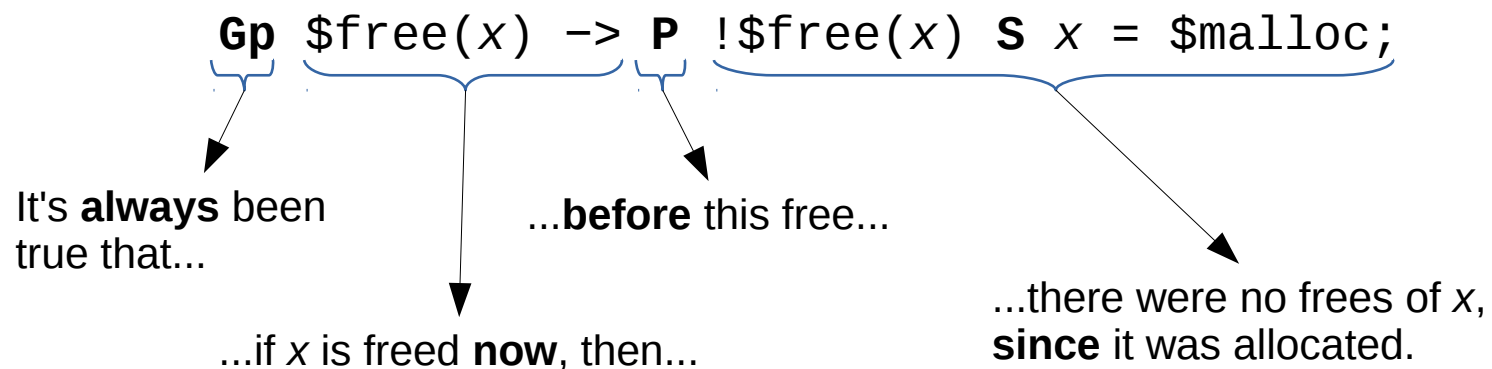
```
/* A very simple TESLA assertion. */  
TESLA_WITHIN(example_syscall,  
    previously(security_check(ANY(ptr),  
        o, op) == 0));
```

<http://www.cl.cam.ac.uk/research/security/ctsrd/tesla/>

Properties to Check: Memory Management

- Could we check this?

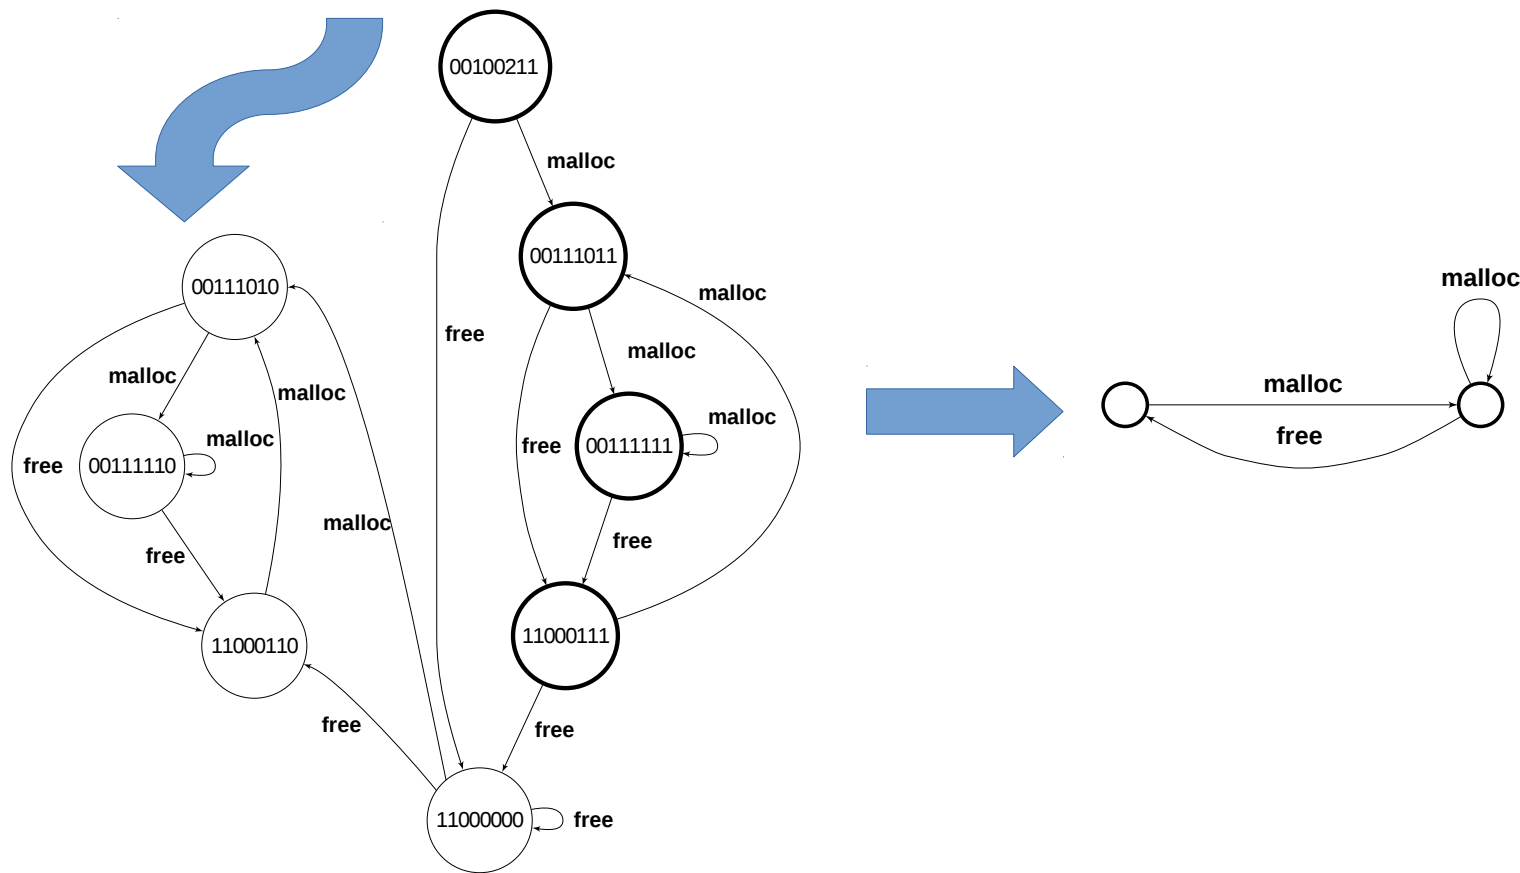
```
void *a = malloc();
...
{a is still allocated}
free(a);
```




Checking LTL with Automata

This is a well-studied problem, and standard algorithms exist:

Gp \$free(x) \rightarrow **P** !\$free(x) **S** x = \$malloc;



A Streaming Verification Engine



Sources

HSSTP


Packet
Capture



Capture

ETM
Sequencer


FPGA
Capture



Processing

Dataflow
Engine

FPGA
Offload



Properties

TESLA

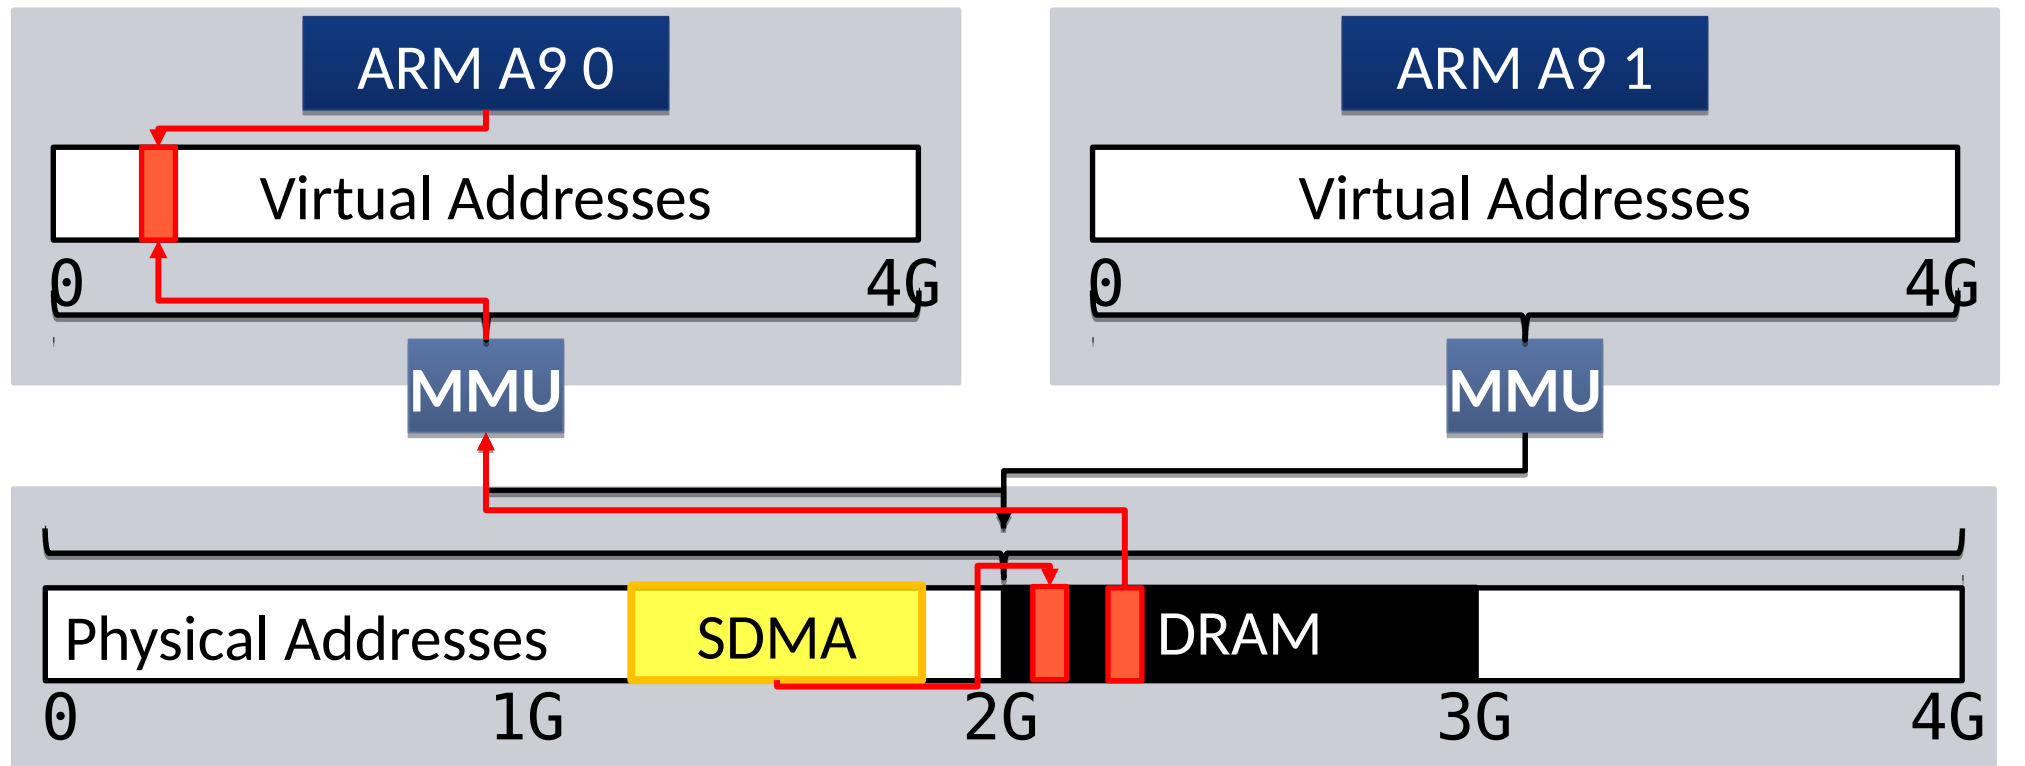
malloc ()
pairing

Coherence
correctness

Constraints

Requirements

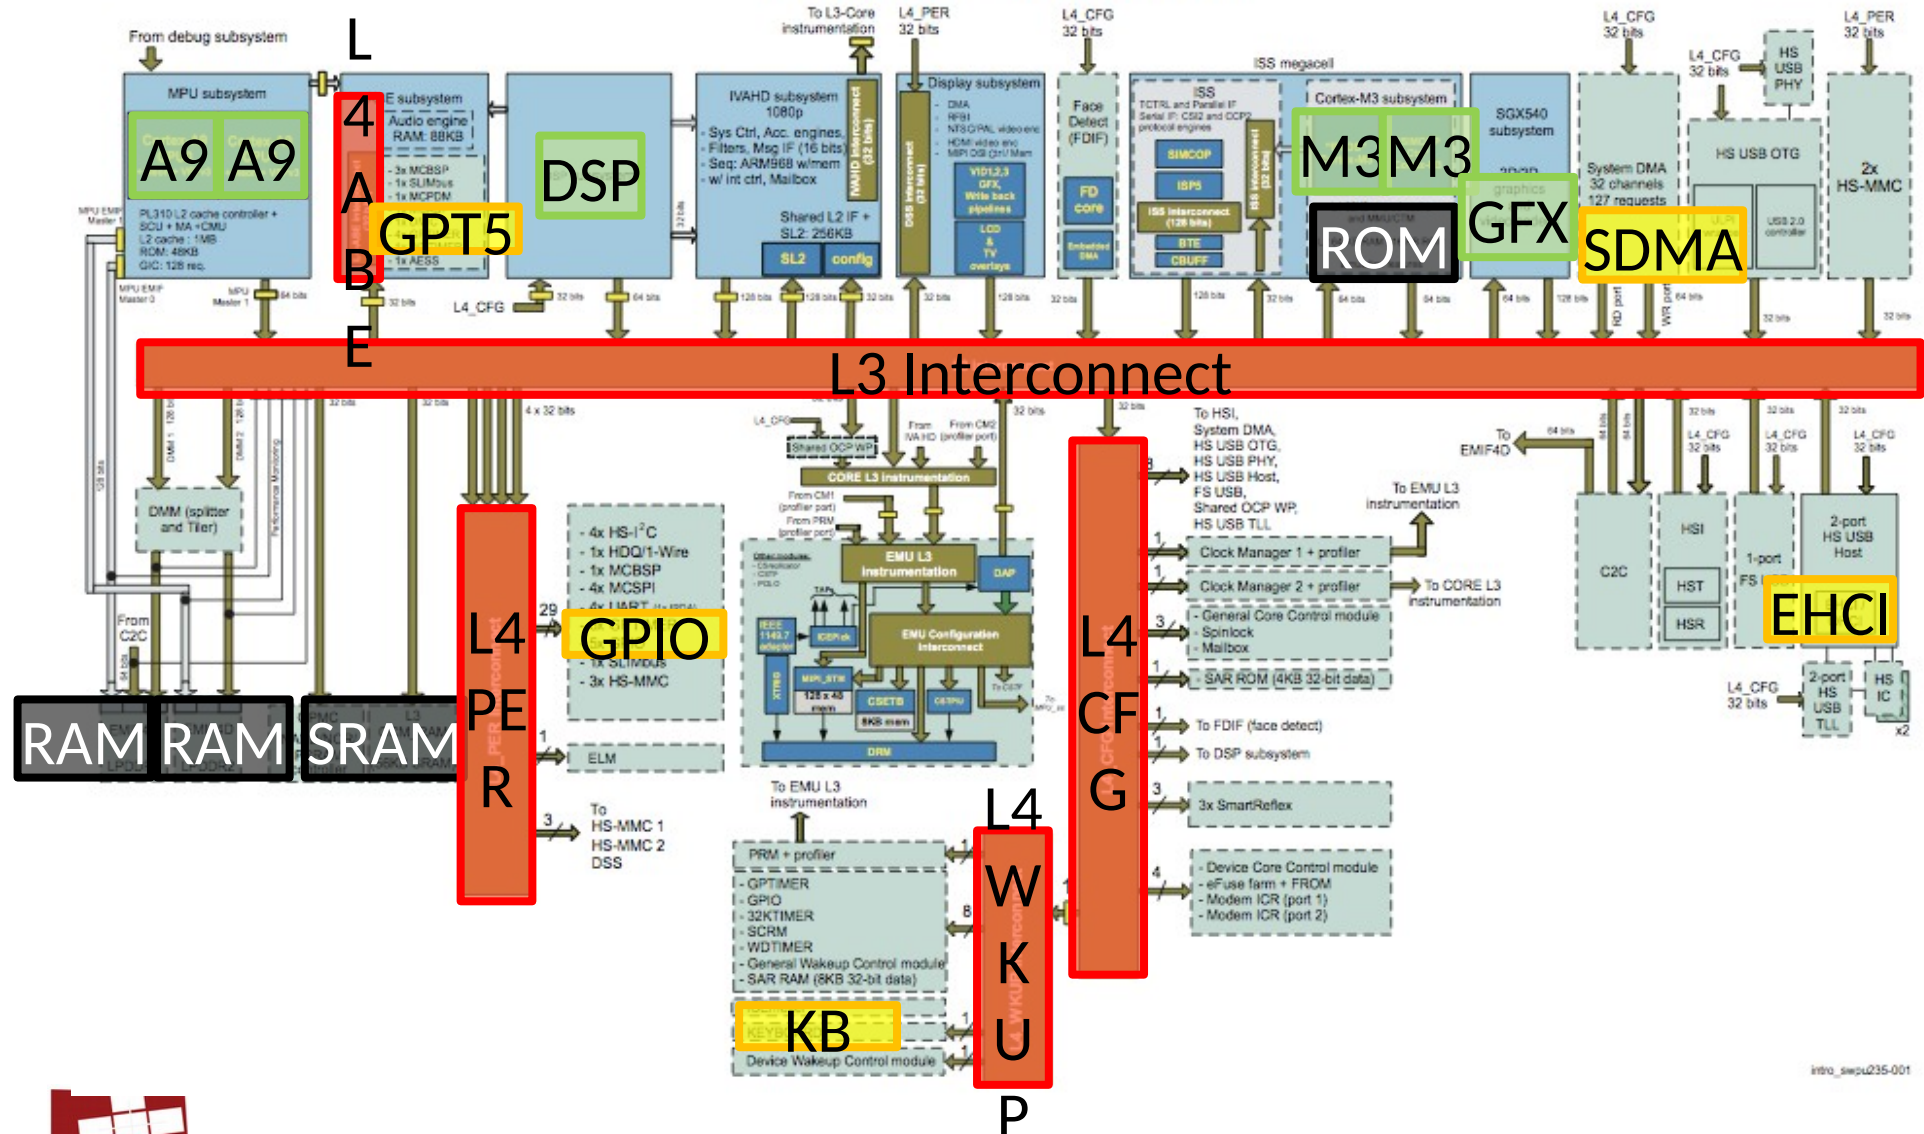
Formal Hardware Modelling



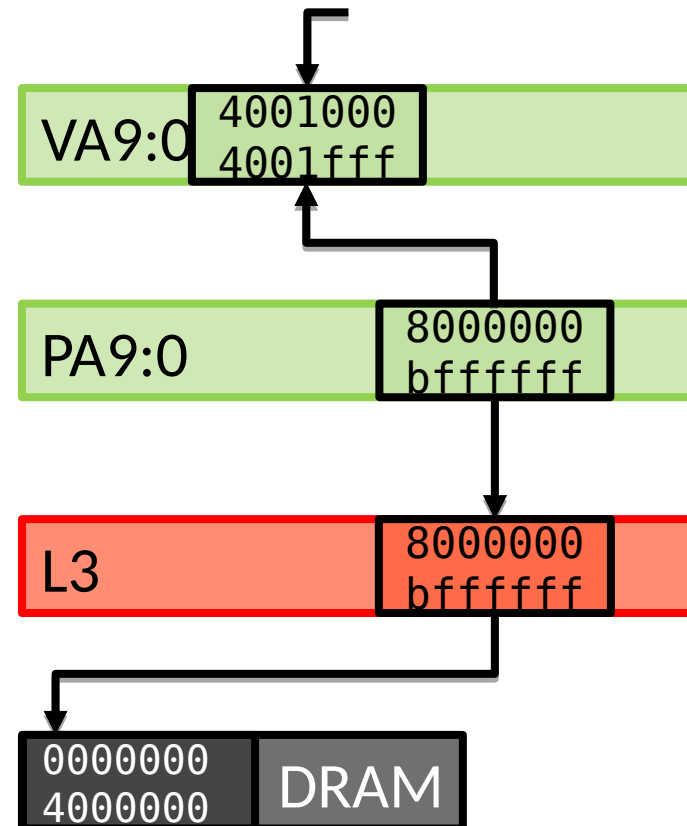

Ti OMAP 4460 SoC

How the Computer *Actually* Looks

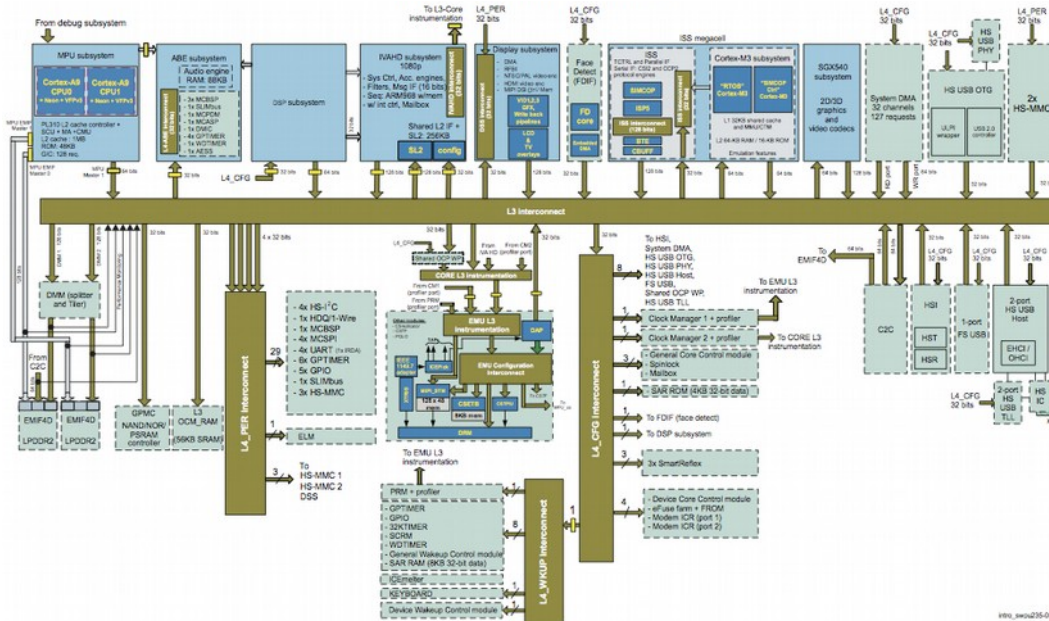
Your mobile phone... 5-10 years ago!



Decoding Nets



The OMAP4460 Decoding Net



$V_{A9:0}$ is map $[20000_3/12 \text{ to } P_{A9:0} \text{ at } 80000_3]$

$P_{A9:0}, P_{A9:1}$ are map $[40138_3/12 \text{ to } GPT \text{ at } 0]$ over $L3$

P_{DSP} is map $[1d3e_3/12 \text{ to } GPT \text{ at } 0]$ over $L3$

V_{M3}, V_{M3} are over $L1_{M3}$

RAM_{M3} is accept $[55020_3/16]$

ROM_{M3} is accept $[55000_3/14]$

MIF is map $[0 - 5ffffff \text{ to } L2_{M3}, 55000_3/14 \text{ to } RAM_{M3}, 55020_3/16 \text{ to } ROM_{M3}]$

$L3$ is map $[49000_3/24 \text{ to } L4 \text{ at } 40100_3, 55000_3/12 \text{ to } MIF]$ accept $[80000_3/30]$

$V_{A9:1}$ is map $[20000_3/12 \text{ to } P_{A9:1} \text{ at } 80000_3]$

V_{DSP} is over P_{DSP}

$L2_{M3}$ is map $[0_{30} \text{ to } L3 \text{ at } 80000_3]$

$L1_{M3}$ is map $[0_{28} \text{ to } MIF]$

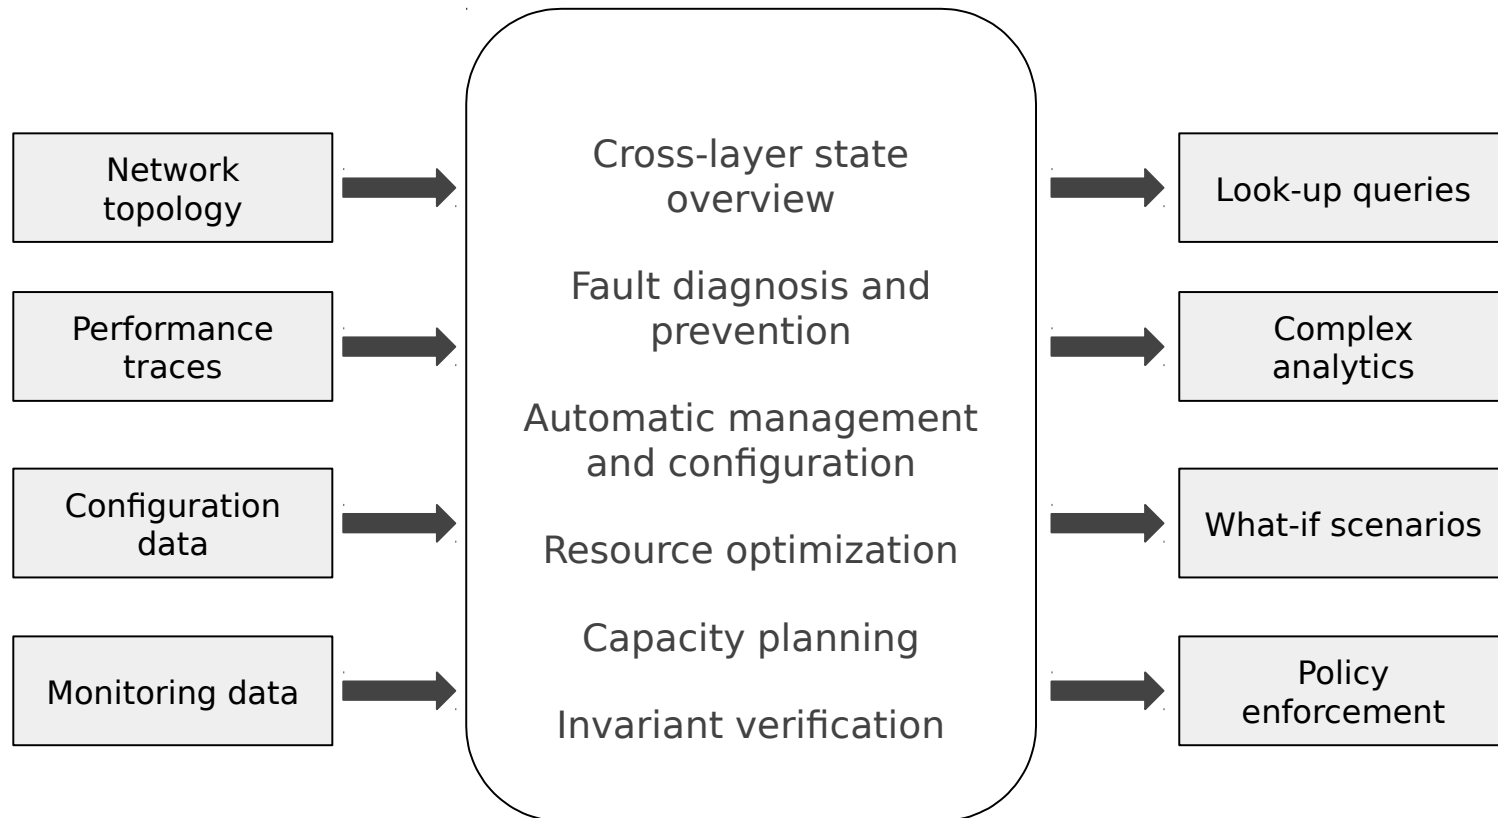
$L4$ is map $[49038_3/12 \text{ to } GPT \text{ at } 0]$

GPT is accept $[0/12]$

Using the model

- Static Configuration:
 - We can now generate the kernel page tables directly from the formal spec.
- Dynamic Discovery and Reconfiguration:
 - The SKB can be populated at runtime – extend the model as hardware is discovered.
- Scheduling:
 - We collaborate with the DB research group on operator scheduling – this work needs the model data.

Strymon: Online Datacenter Analytics and Management



Support:

amadeus



FONDS NATIONAL SUISSE
SCHWEIZERISCHER NATIONALFONDS
FONDO NAZIONALE SVIZZERO
SWISS NATIONAL SCIENCE FOUNDATION



More info: <http://strymon.systems.ethz.ch>

Data
Management

Distributed
Stream
Processing

Graph
Algorithms

Real-time sessionization

Incremental network routing

Online critical path analysis

Provenance and interactive explanations

Streaming fault-tolerance and data durability

Automatic re-scaling of distributed dataflows




Apache Flink

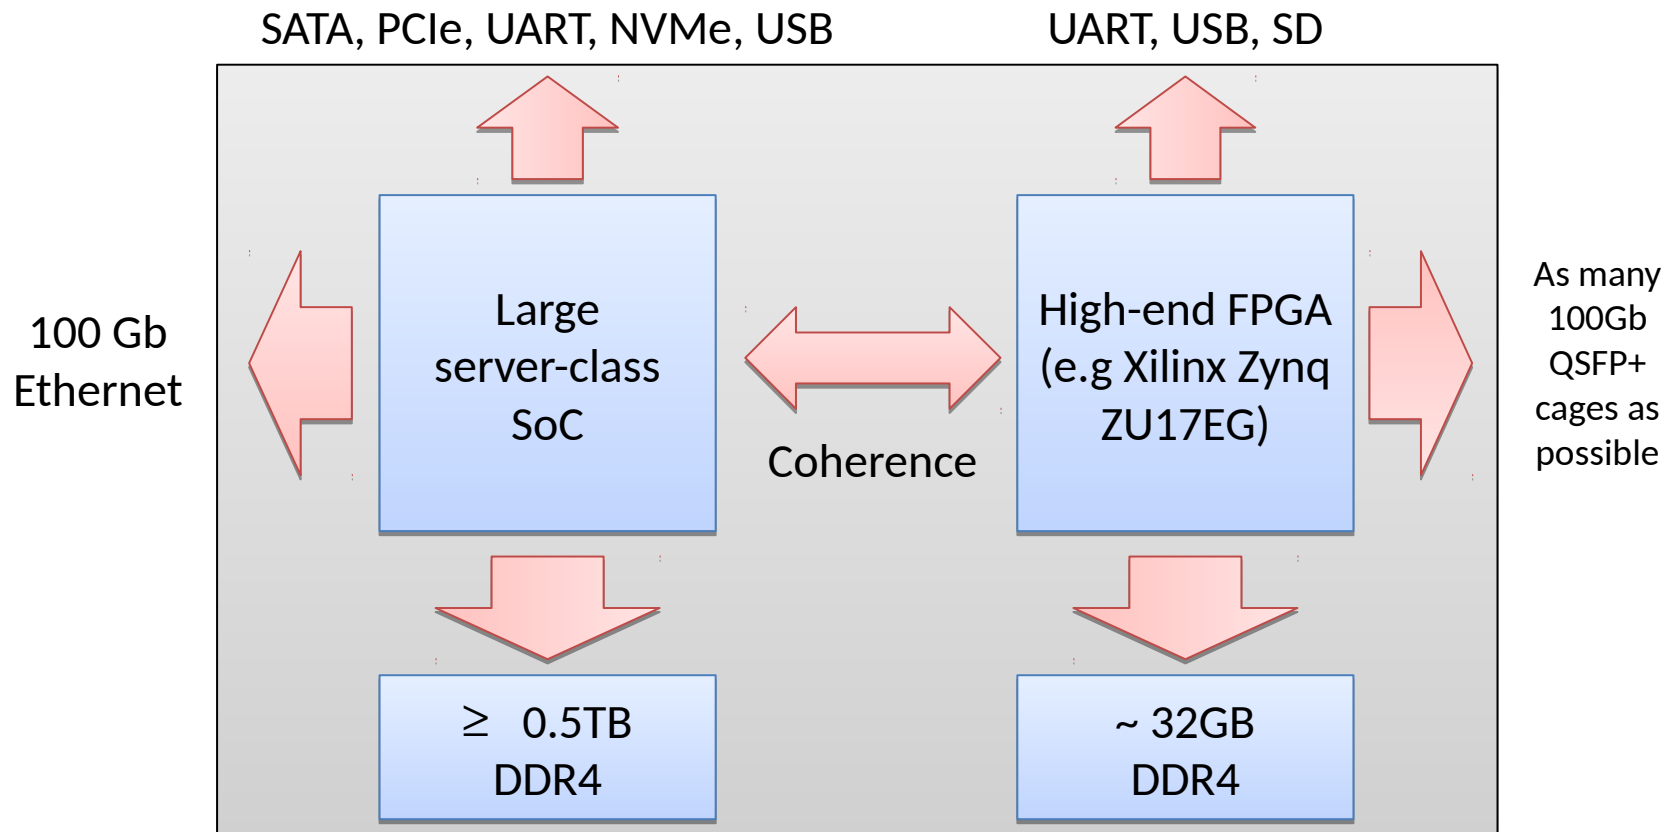


Contact: Vasia Kalavri (kalavriv@inf.ethz.ch)

Enzian: Building Understandable Hardware



Enzian: Sketch



Living and Working in Zürich



Living and Working in Zürich

