ETH zürich

Enzian: a Research Computer

Timothy Roscoe, David Cock
Systems Group
ETH Zurich Department of Computer Science

ENZIAN
DINFK
Systems@ETH zürich

What This Talk is About

- The challenge of ***relevant*** academic ***systems software research*** in an age of custom and reconfigurable hardware.
- Our response: build a computer ***designed for academic research***
- A variety of use-cases
- Methodological issues and wider ramifications

David Cock, Systems Group
Department of Computer Science
ETH Zurich

Typical Rack-Scale Architecture (in research, at least)

FDR ~ 56Gb/s

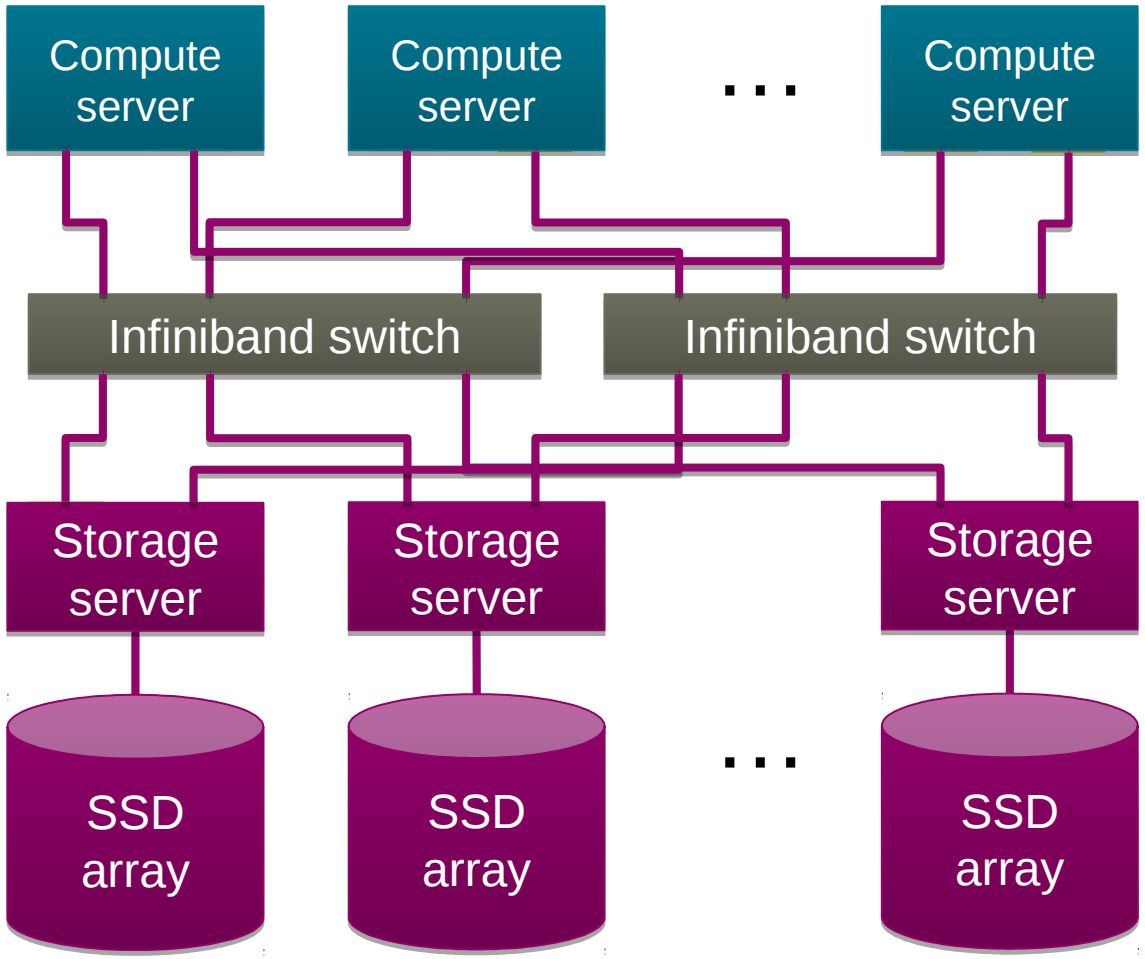

Q. Message latency in this network?

A. ~ 0.7μs, or ~10 LLC misses.

⇒ need to think of this as ***one big machine***

David Cock, Systems Group
Department of Computer Science
ETH Zurich

But it's not typical!

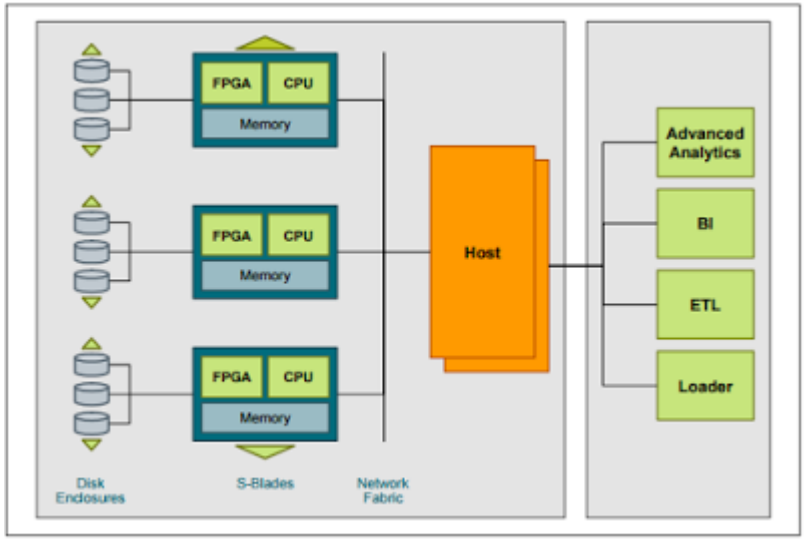

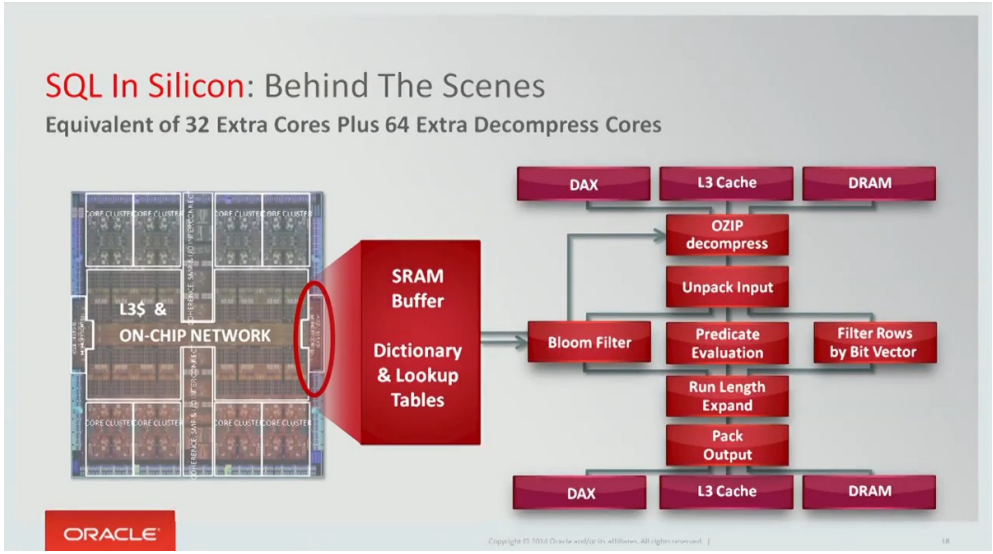

David Cock, Systems Group
Department of Computer Science
ETH Zurich

The Challenge for Systems Software Research

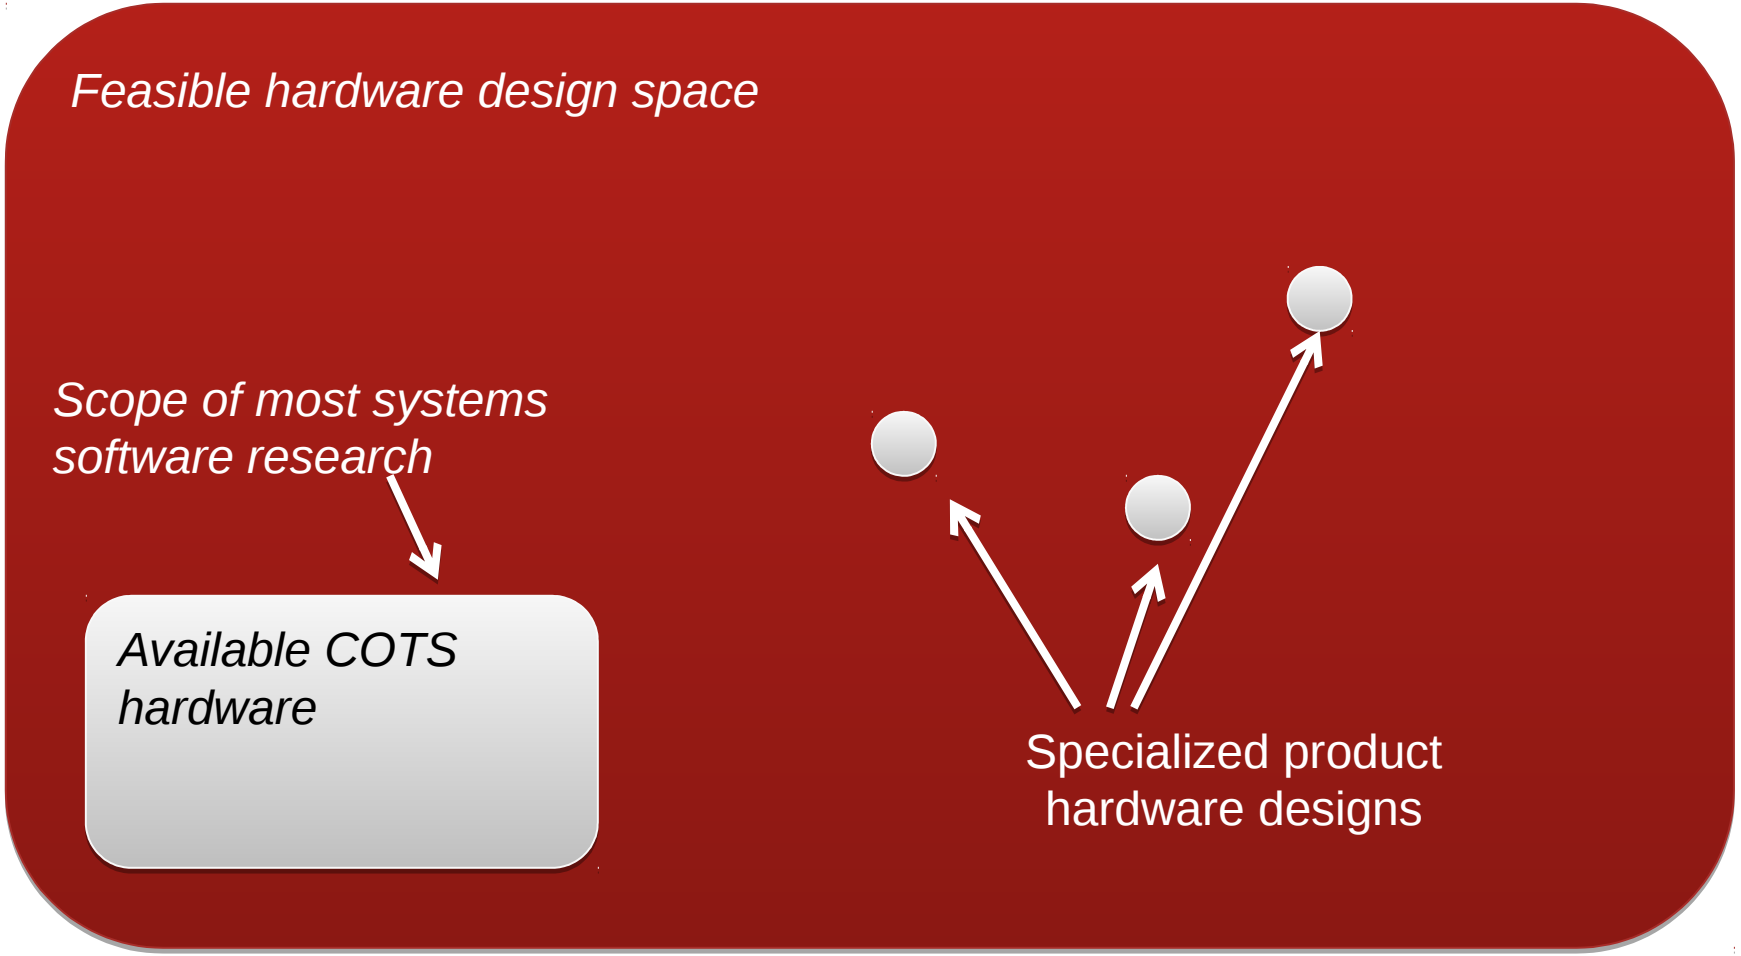

David Cock, Systems Group
Department of Computer Science
ETH Zurich

Challenges

Challenge 1: *Getting the hardware*

- Often developed by *aaS companies
- Fear of support costs

Challenge 2: *Getting the documentation*

- Documentation non-existent or proprietary
- Lending is legally easier than releasing documentation

Challenge 3: *The big decisions have already been made*

- Custom hardware is built with an application and mode of use already in mind
- Hardware probably never tested outside a narrow envelope

David Cock, Systems Group
Department of Computer Science
ETH Zurich

Enzian v2

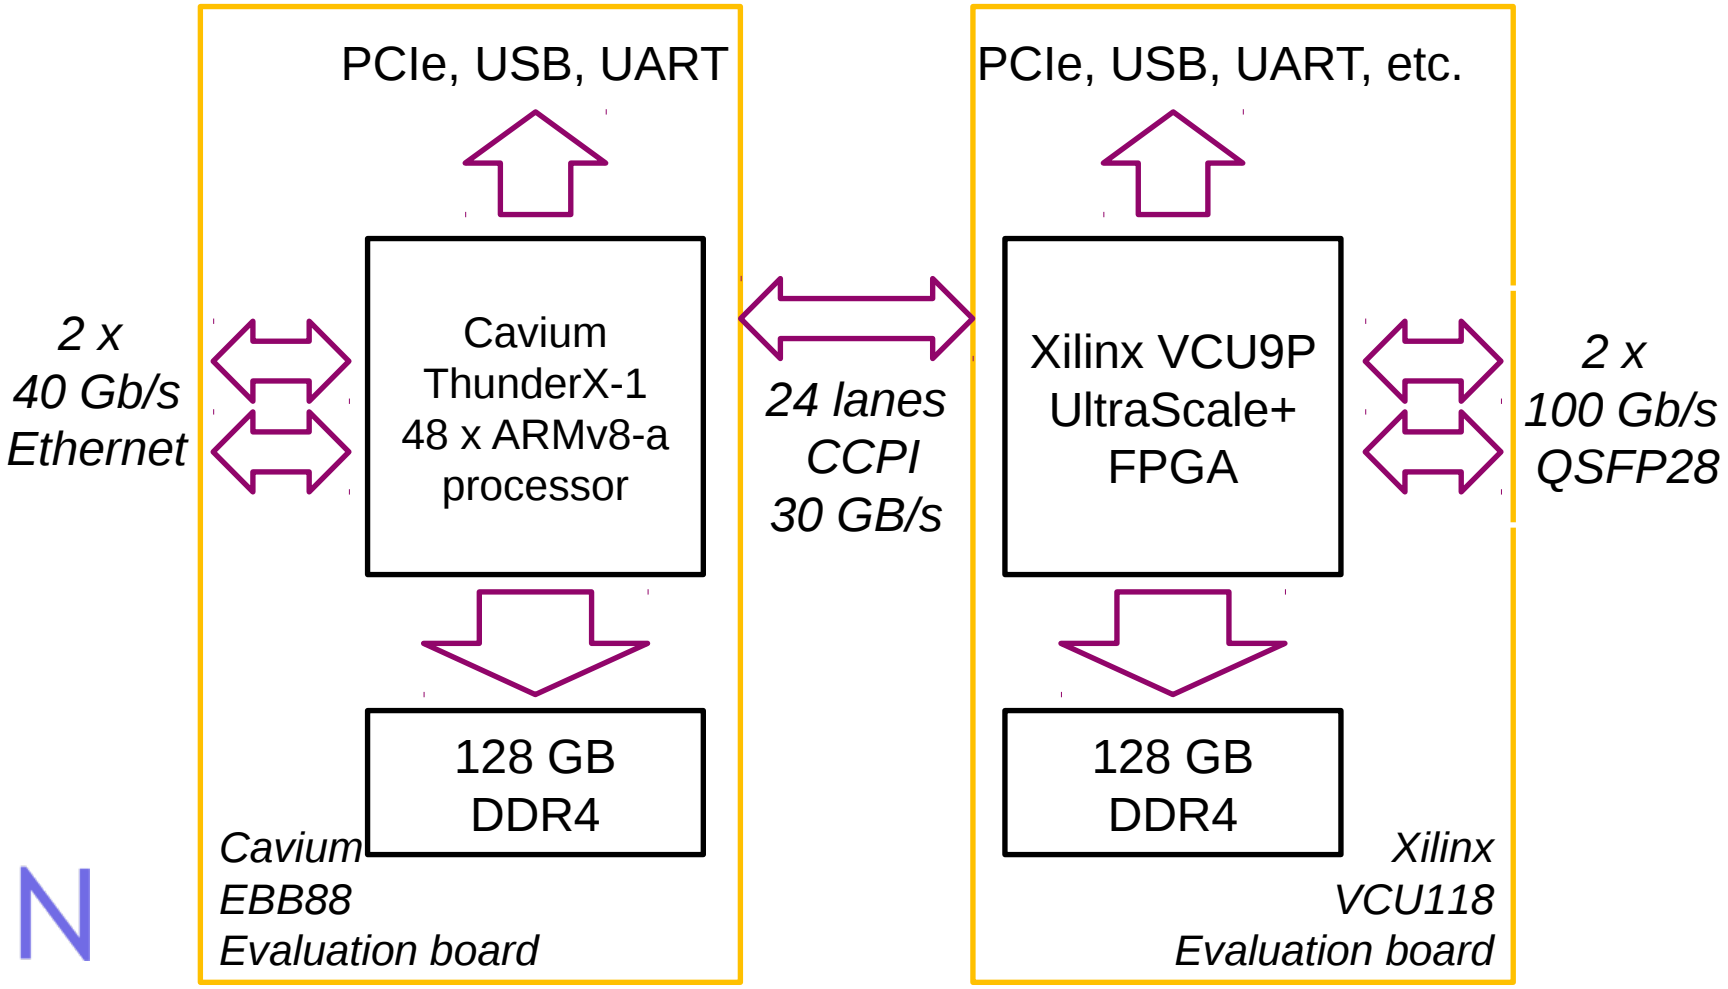

ENZIAN

David Cock, Systems Group
Department of Computer Science
ETH Zurich

Enzian v2

ENZIAN

David Cock, Systems Group
Department of Computer Science
ETH Zurich

Enzian v2

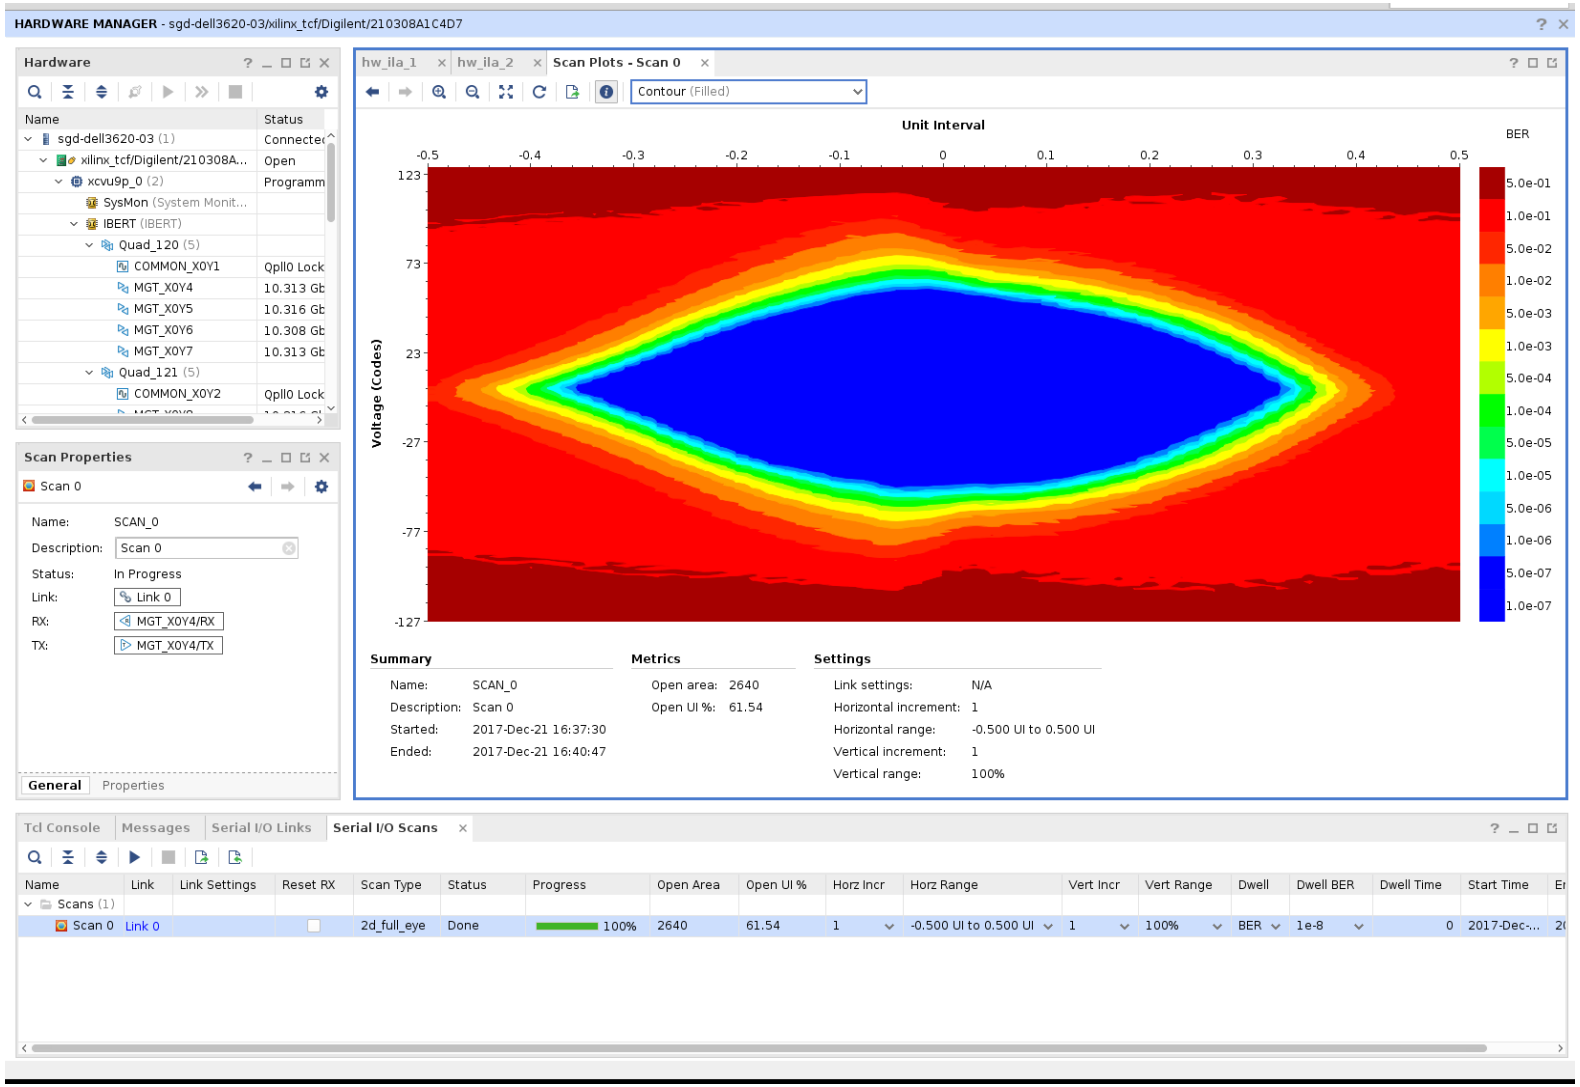

ENZIAN

David Cock, Systems Group
Department of Computer Science
ETH Zurich

Enzian v3 (2018, Design Contract Awarded)

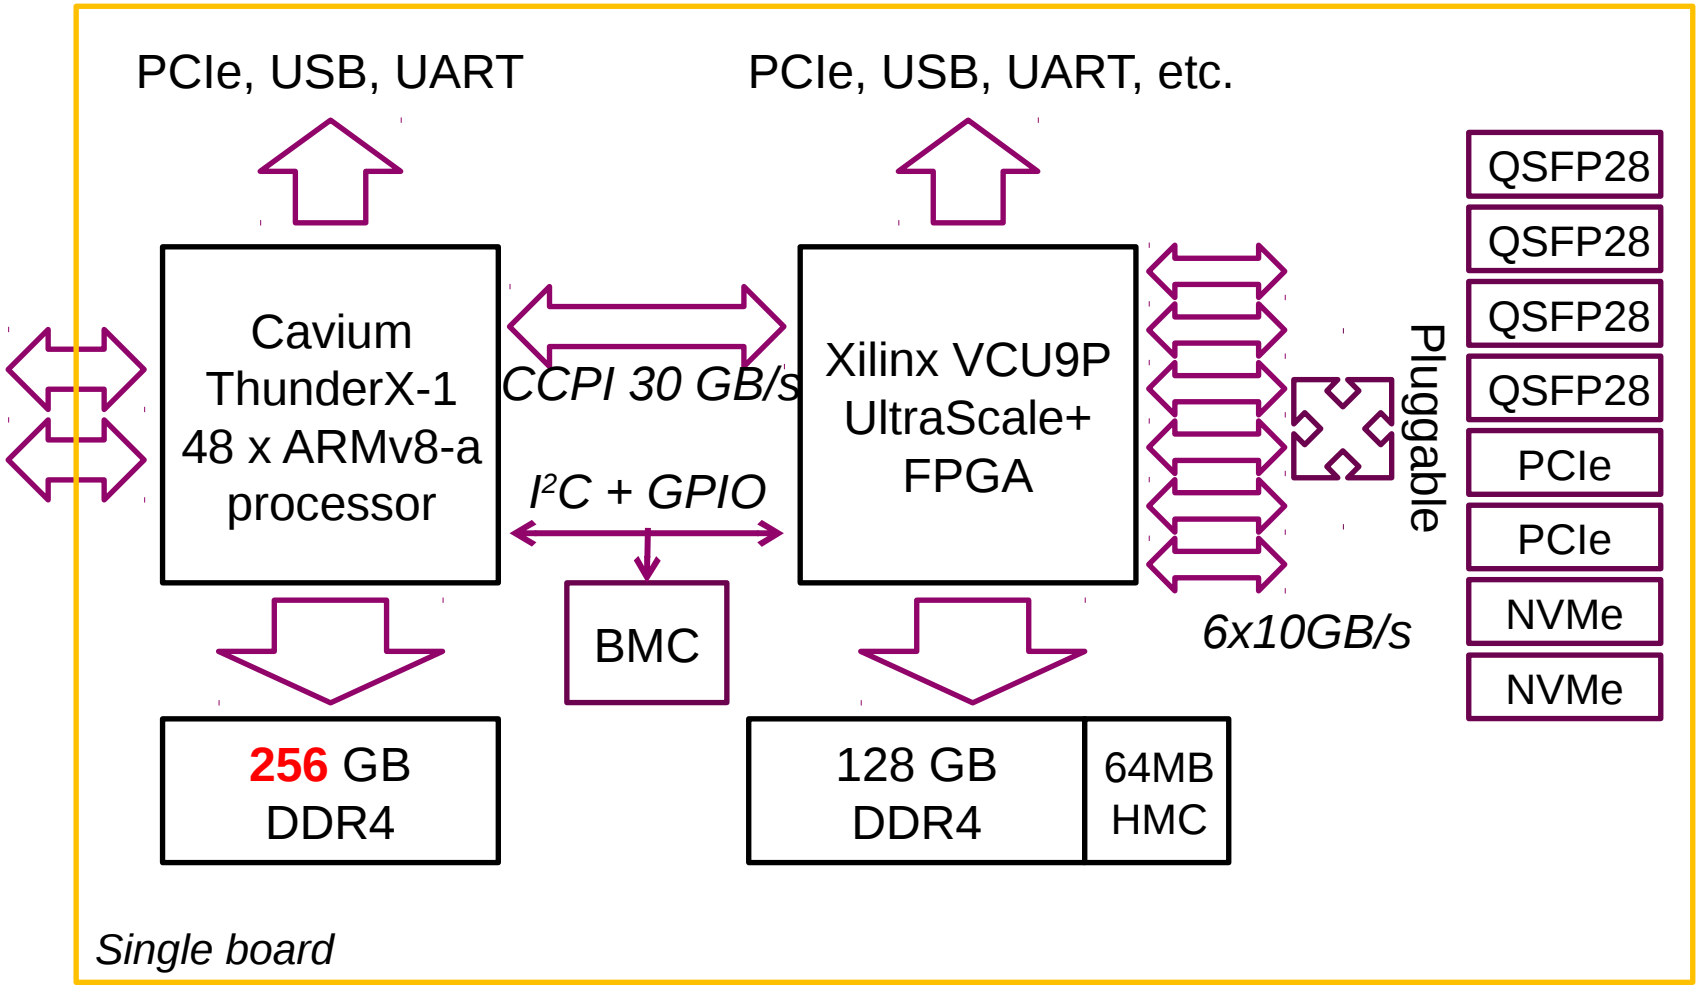

ENZIAN

David Cock, Systems Group
Department of Computer Science
ETH Zurich

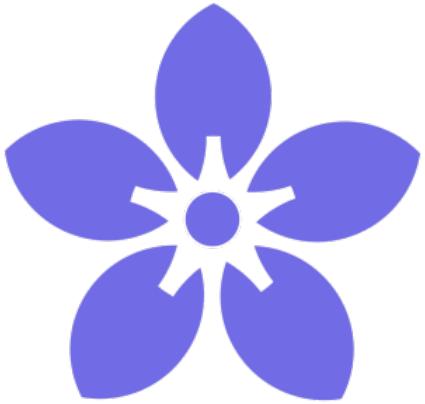

ENZIAN

RESEARCH APPLICATIONS (THAT WE'VE THOUGHT OF)

David Cock, Systems Group
Department of Computer Science
ETH Zurich

Accelerator Design

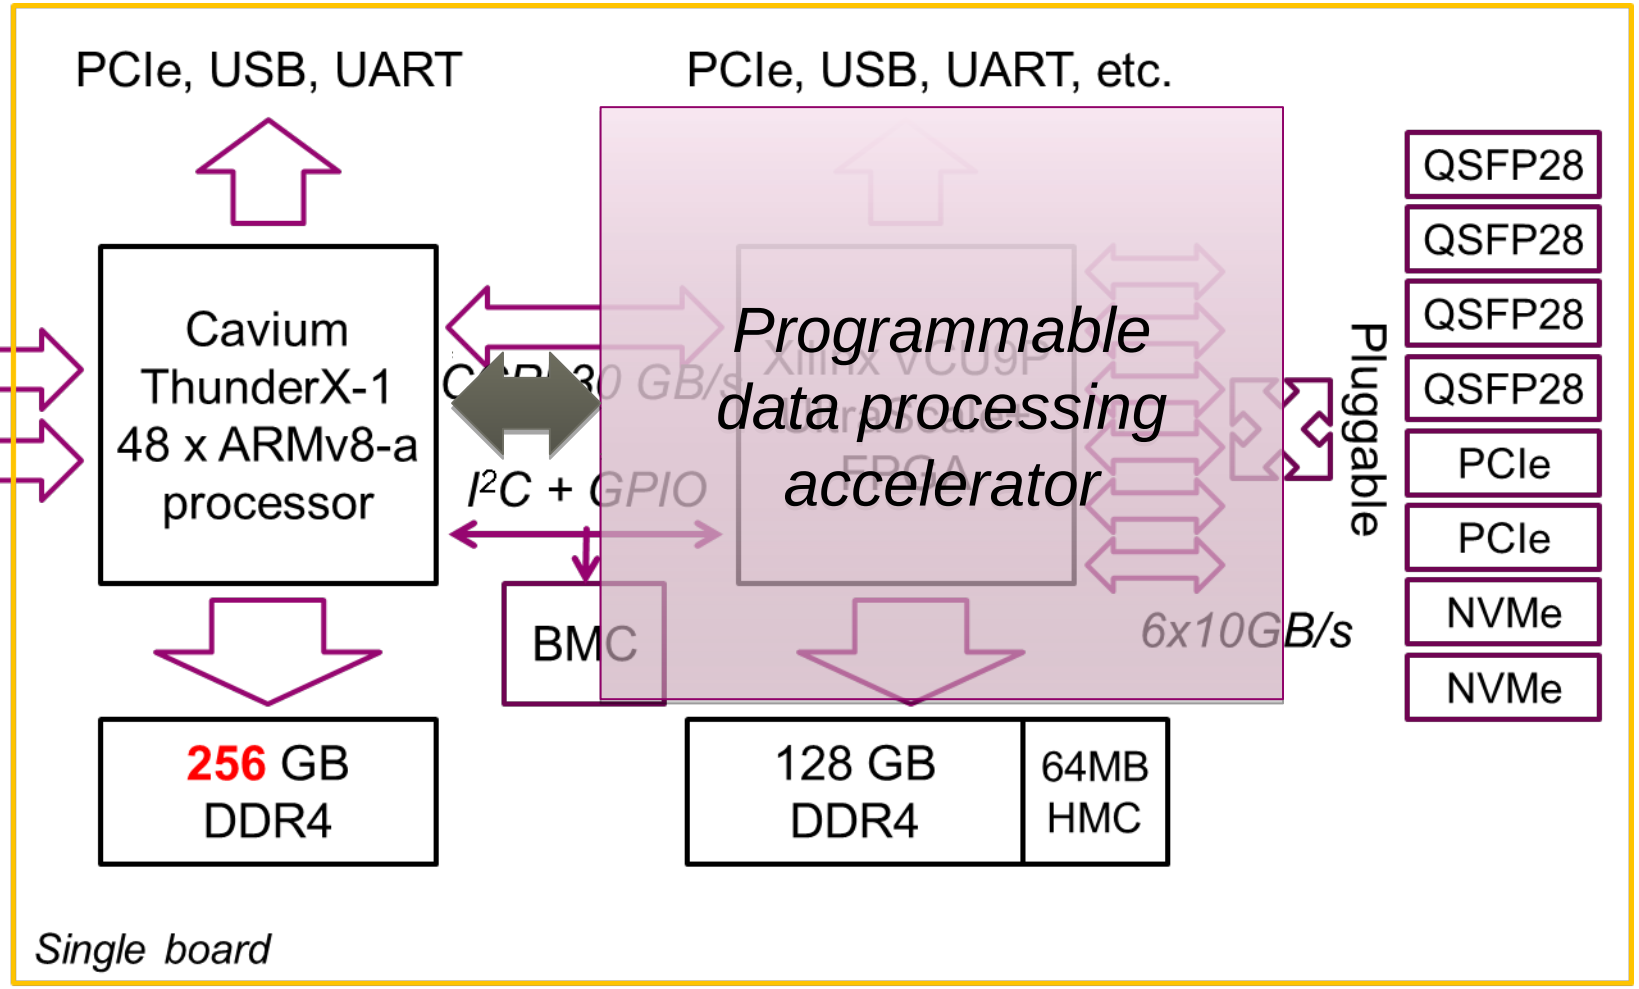

ENZIAN

David Cock, Systems Group
Department of Computer Science
ETH Zurich

The Next NetFPGA

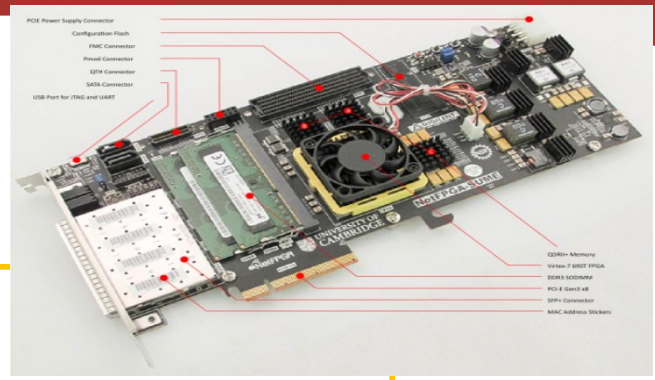

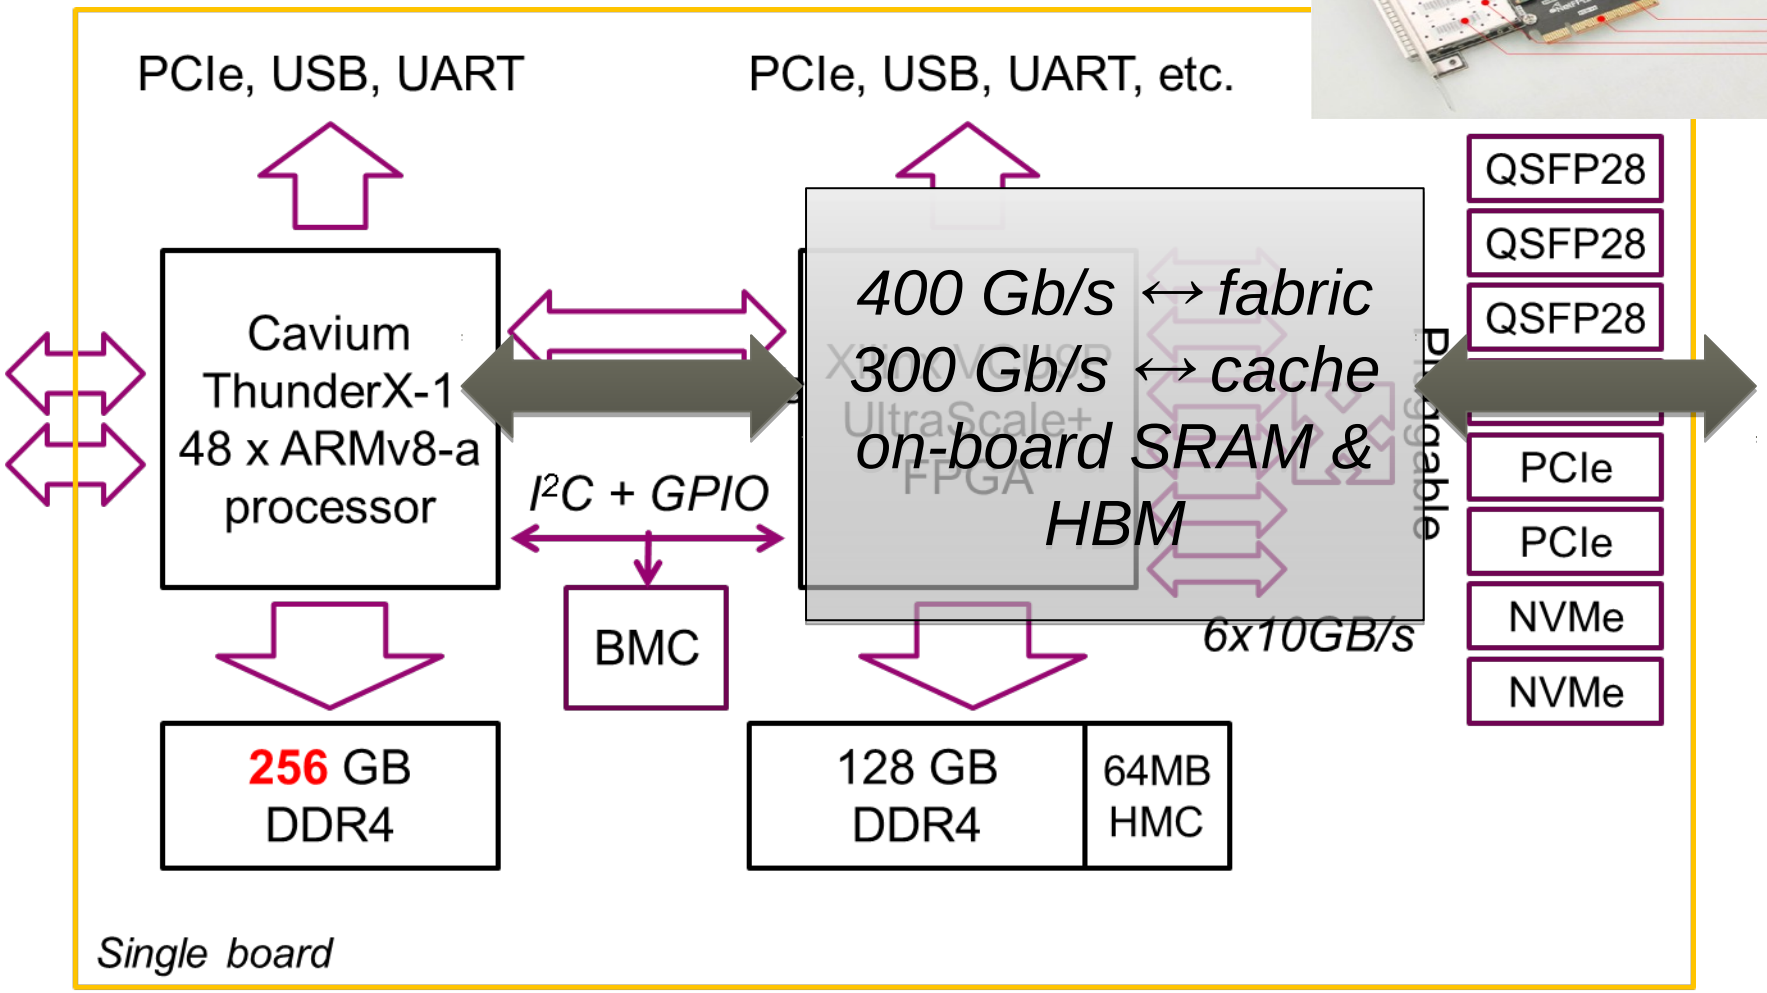

ENZIAN

David Cock, Systems Group
Department of Computer Science
ETH Zurich

Remote Memory

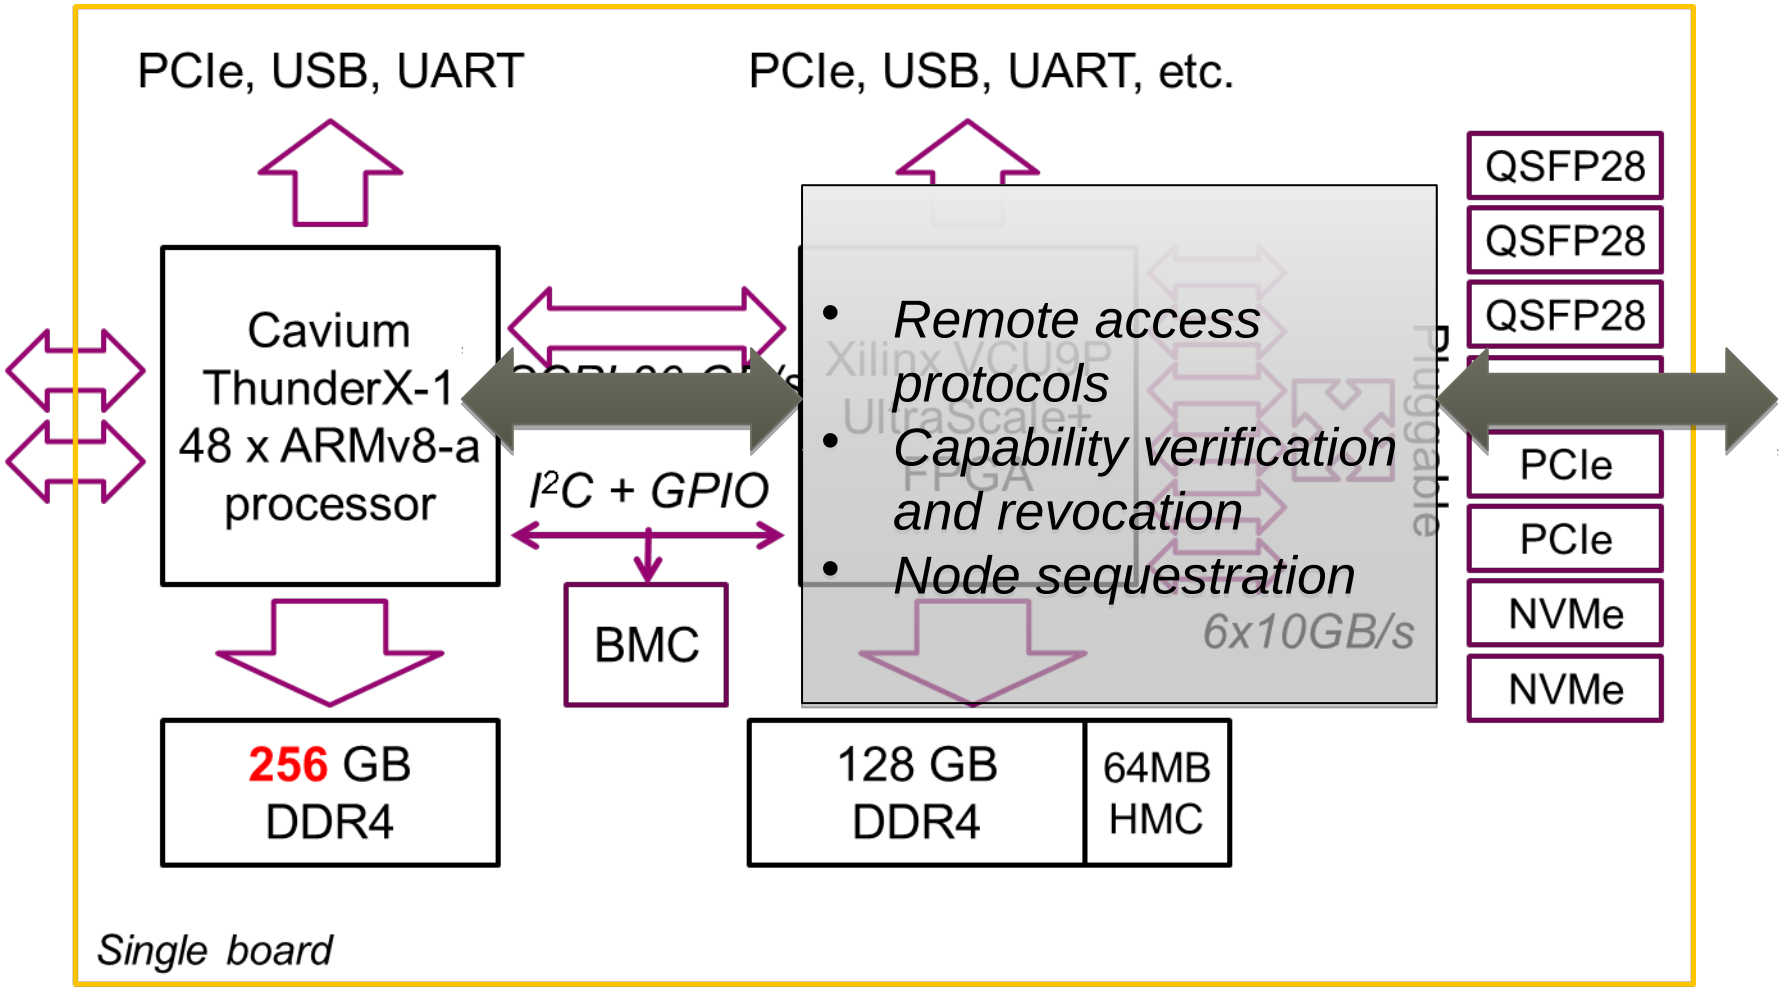

ENZIAN

David Cock, Systems Group
Department of Computer Science
ETH Zurich

POSSIBLE IMPLICATIONS

David Cock, Systems Group
Department of Computer Science
ETH Zurich

Claiming Validity

1. This ***is*** the platform
   - Why is this platform a valid design point?
   - How to meaningfully compare with others?

2. Emulating a ***possible new*** platform design
   - What would this mean if built from scratch?
   - What is the cost/benefit?

David Cock, Systems Group
Department of Computer Science
ETH Zurich

Methodological Questions

- ***Application*** use-cases, e.g:
  - NFV applications & control/data plane design
  - Database design exploiting hardware accelerators

- ***Systems*** use-cases, e.g.:
  - Accelerator interface design
  - OS support for closely-coupled FPGAs

David Cock, Systems Group
Department of Computer Science
ETH Zurich

Higher Goal: Research Amplification

- Seed the research community
  - Remove major barrier to innovation in academia
- Precedents:
  - PlanetLab
  - Berkeley Unix
  - NetFPGA
  - Intel's DPDK
  - …

DPDK
DATA PLANE DEVELOPMENT KIT

NetFPGA

PLANETLAB
An open platform for developing, deploying, and accessing planetary-scale services

David Cock, Systems Group
Department of Computer Science
ETH Zurich

Summary: COTS Hardware is Unrealistic

- All the product action uses **custom** hardware
- Vendors can build almost **anything**
- What to build is an **economic** question
  - ⇒ not something we can answer
- We needs **overengineered research platforms**
  - Our goal should be to deliver **options** and **techniques**

ENZIAN

www.enzian.systems

David Cock, Systems Group
Department of Computer Science
ETH Zurich

Questions?

David Cock, Systems Group
Department of Computer Science
ETH Zurich

Xilinx Ultrascale+ MPSoC

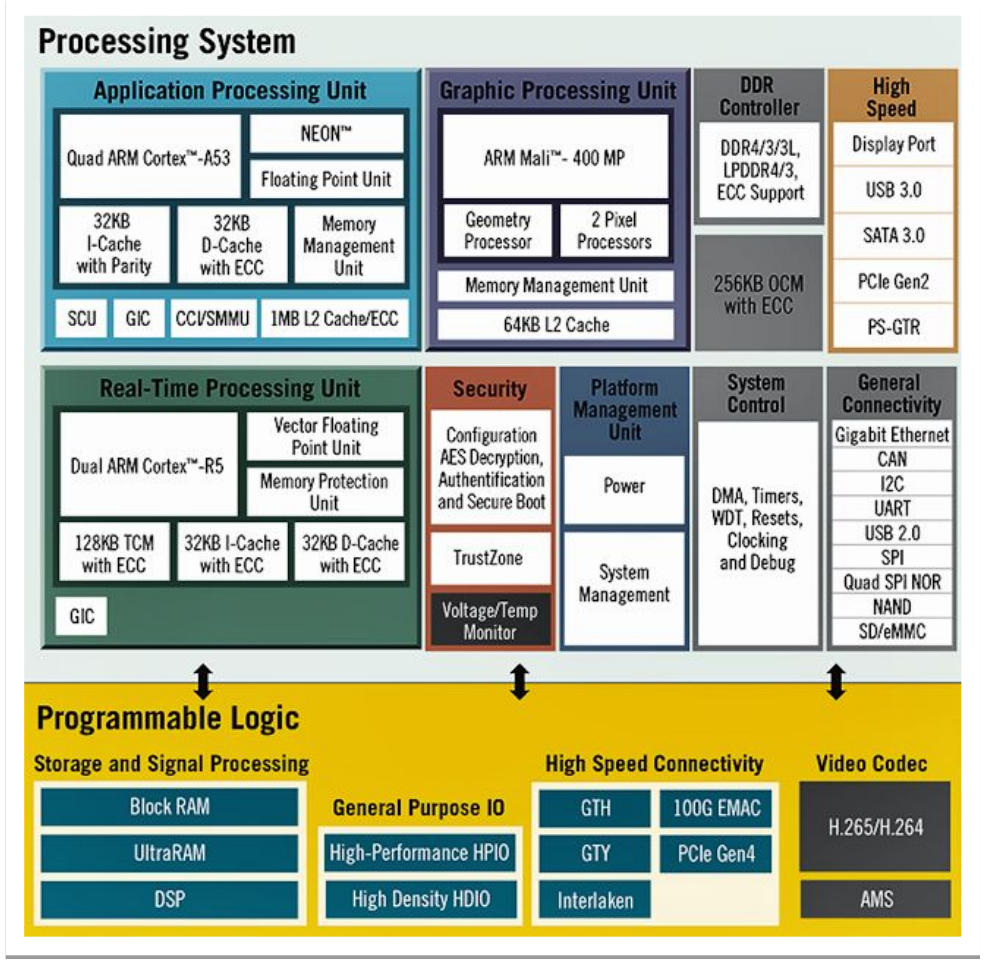

David Cock, Systems Group
Department of Computer Science
ETH Zurich

Evaluation Boards

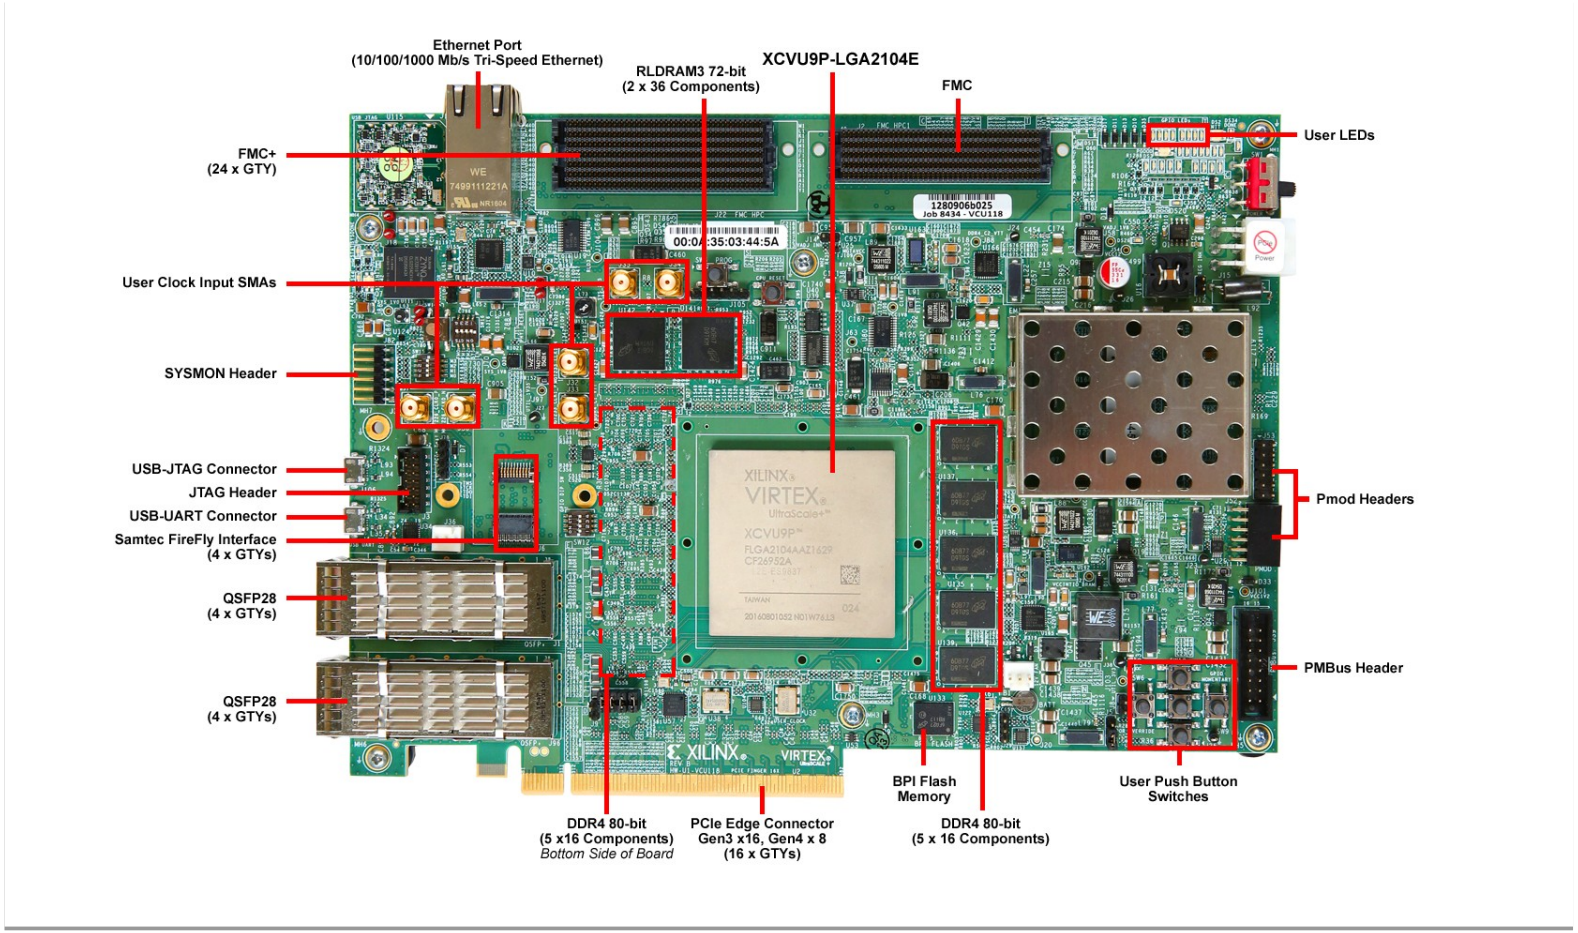

Can't take full-system approach

David Cock, Systems Group
Department of Computer Science
ETH Zurich

Intel HARP Program

IvyTown Xeon + Stratix V FPGA

Accelerating Workloads using Xeon and coherently attached FPGA in-socket

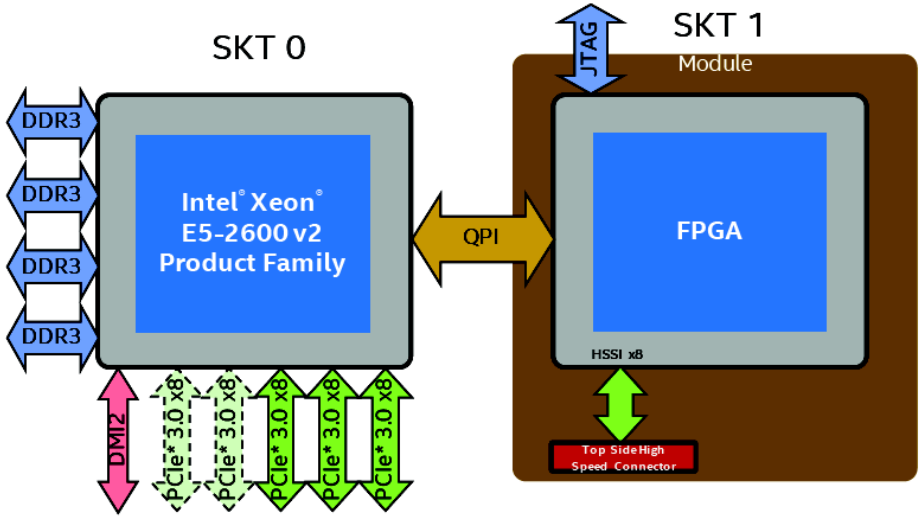

Processor	Intel® Xeon® E5-26xx v2 Processor
FPGA Module	Altera Stratix V
QPI Speed	6.4 GT/s full width
Memory to FPGA Module	2 channels of DDR3 (not used on HARP platform)
Features	Configuration Agent, Caching Agent,, (optional) Memory Controller
Software	Accelerator Abstraction Layer (AAL) runtime, drivers, sample applications

intel | 4