Heterogeneous Shared-Memory Multicore Processors

H. Peter Hofstee
IBM Austin / TU Delft

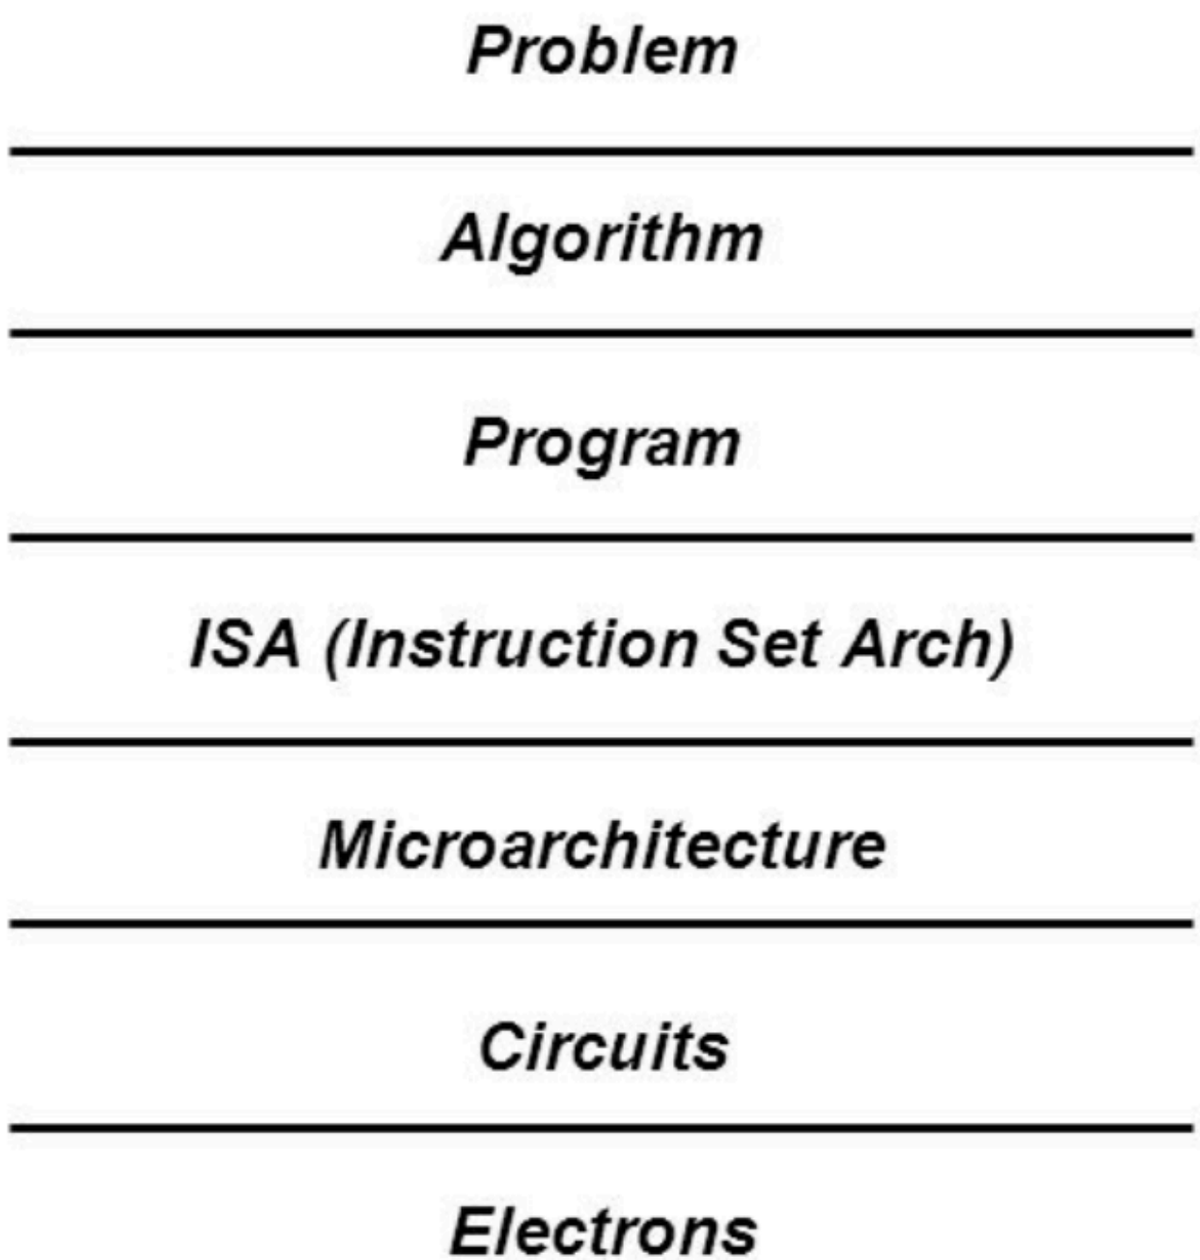

Transformation Hierarchy. Yale Patt

Playing at the program/ISA boundary ….

E peròsappia ciascuno che nulla cosa per legame musaico armonizzata si può de la sua loquela in altra transmutare, sanza rompere tutta sua dolcezza e armonia.

Dante (Convivio)

And yet each of them knows that nothing by a harmonized mosaic can be told of his talk in another transmutation without breaking all his sweetness and harmony.

… and Google Translate

Playing at the program/ISA boundary ….

E peròsappia ciascuno che nulla cosa per legame musaico armonizzata si può de la sua loquela in altra transmutare, sanza rompere tutta sua dolcezza e armonia.

Dante (Convivio)

And so everyone should know that nothing harmonized according to the Muse’s rules can be translated from its native speech into another without breaking all its sweetness and harmony.

Overview

- It is (indeed) all about the memory
- Programming Heterogeneous Multicore
- Apache Arrow and Fletcher

CPU

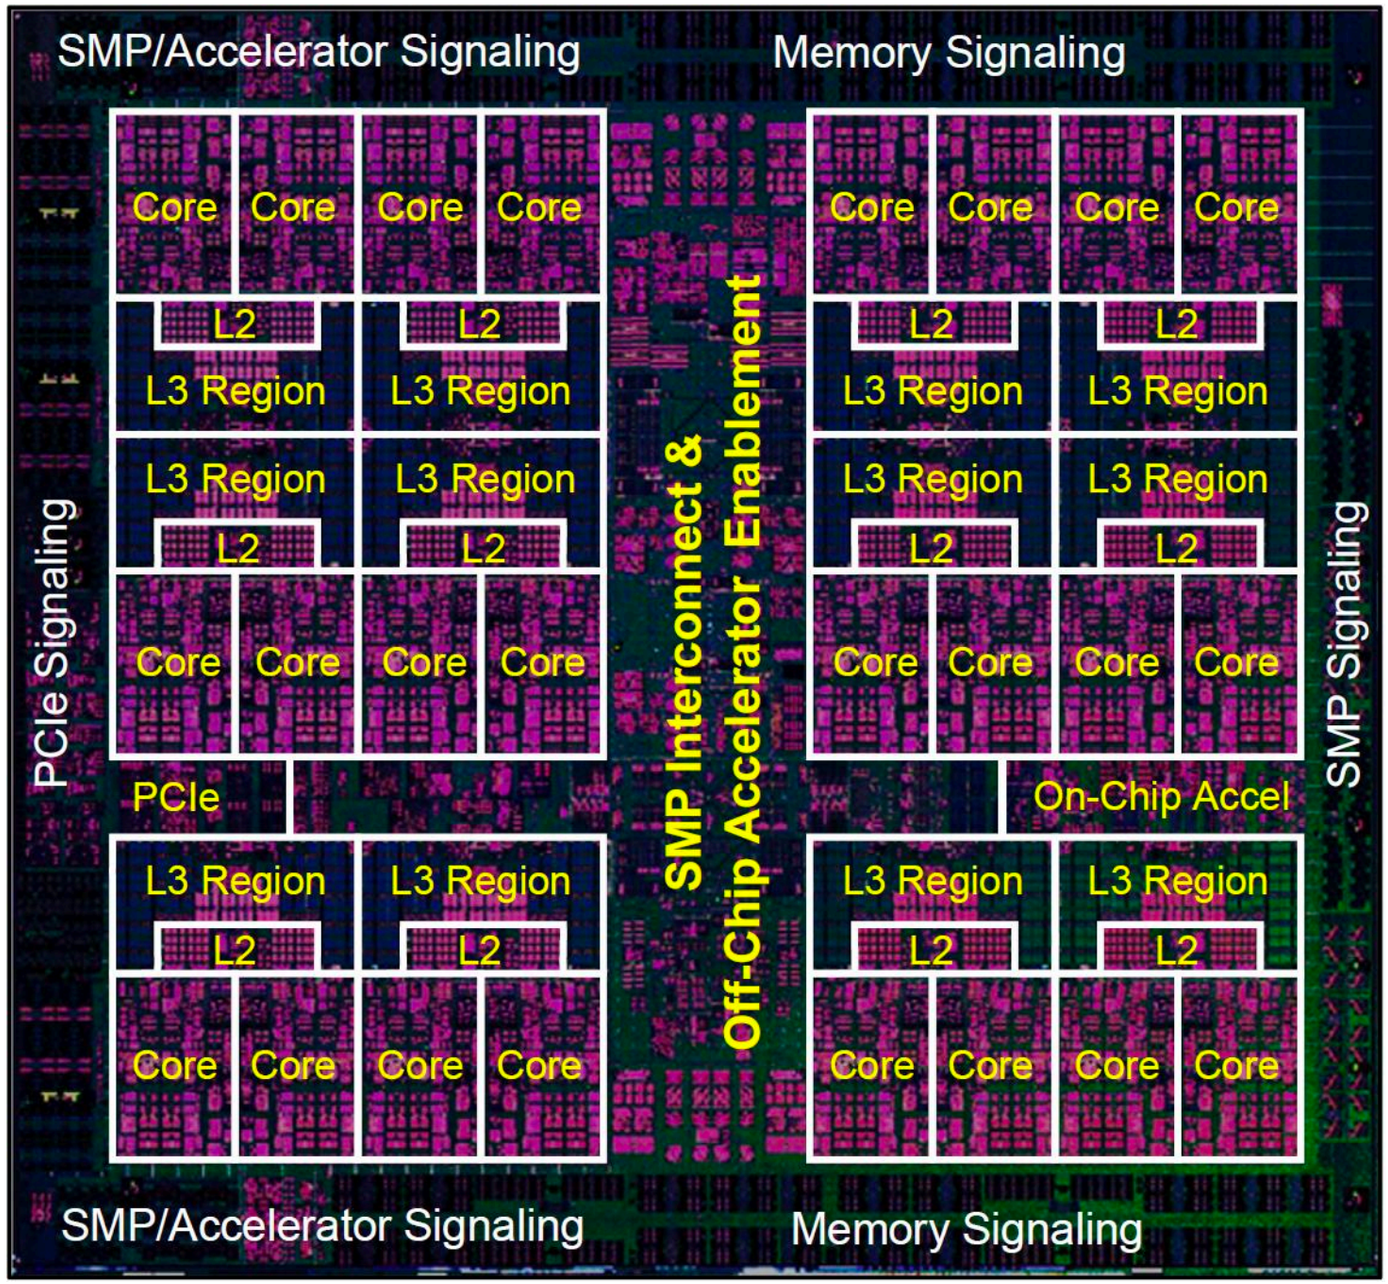

POWER9

Source: "POWER9 Processor for the Cognitive Era". IBM presentation by Brian Thompto. Hot Chips 28 Symposium, October 2016

A(nother) story for a 10-year old …

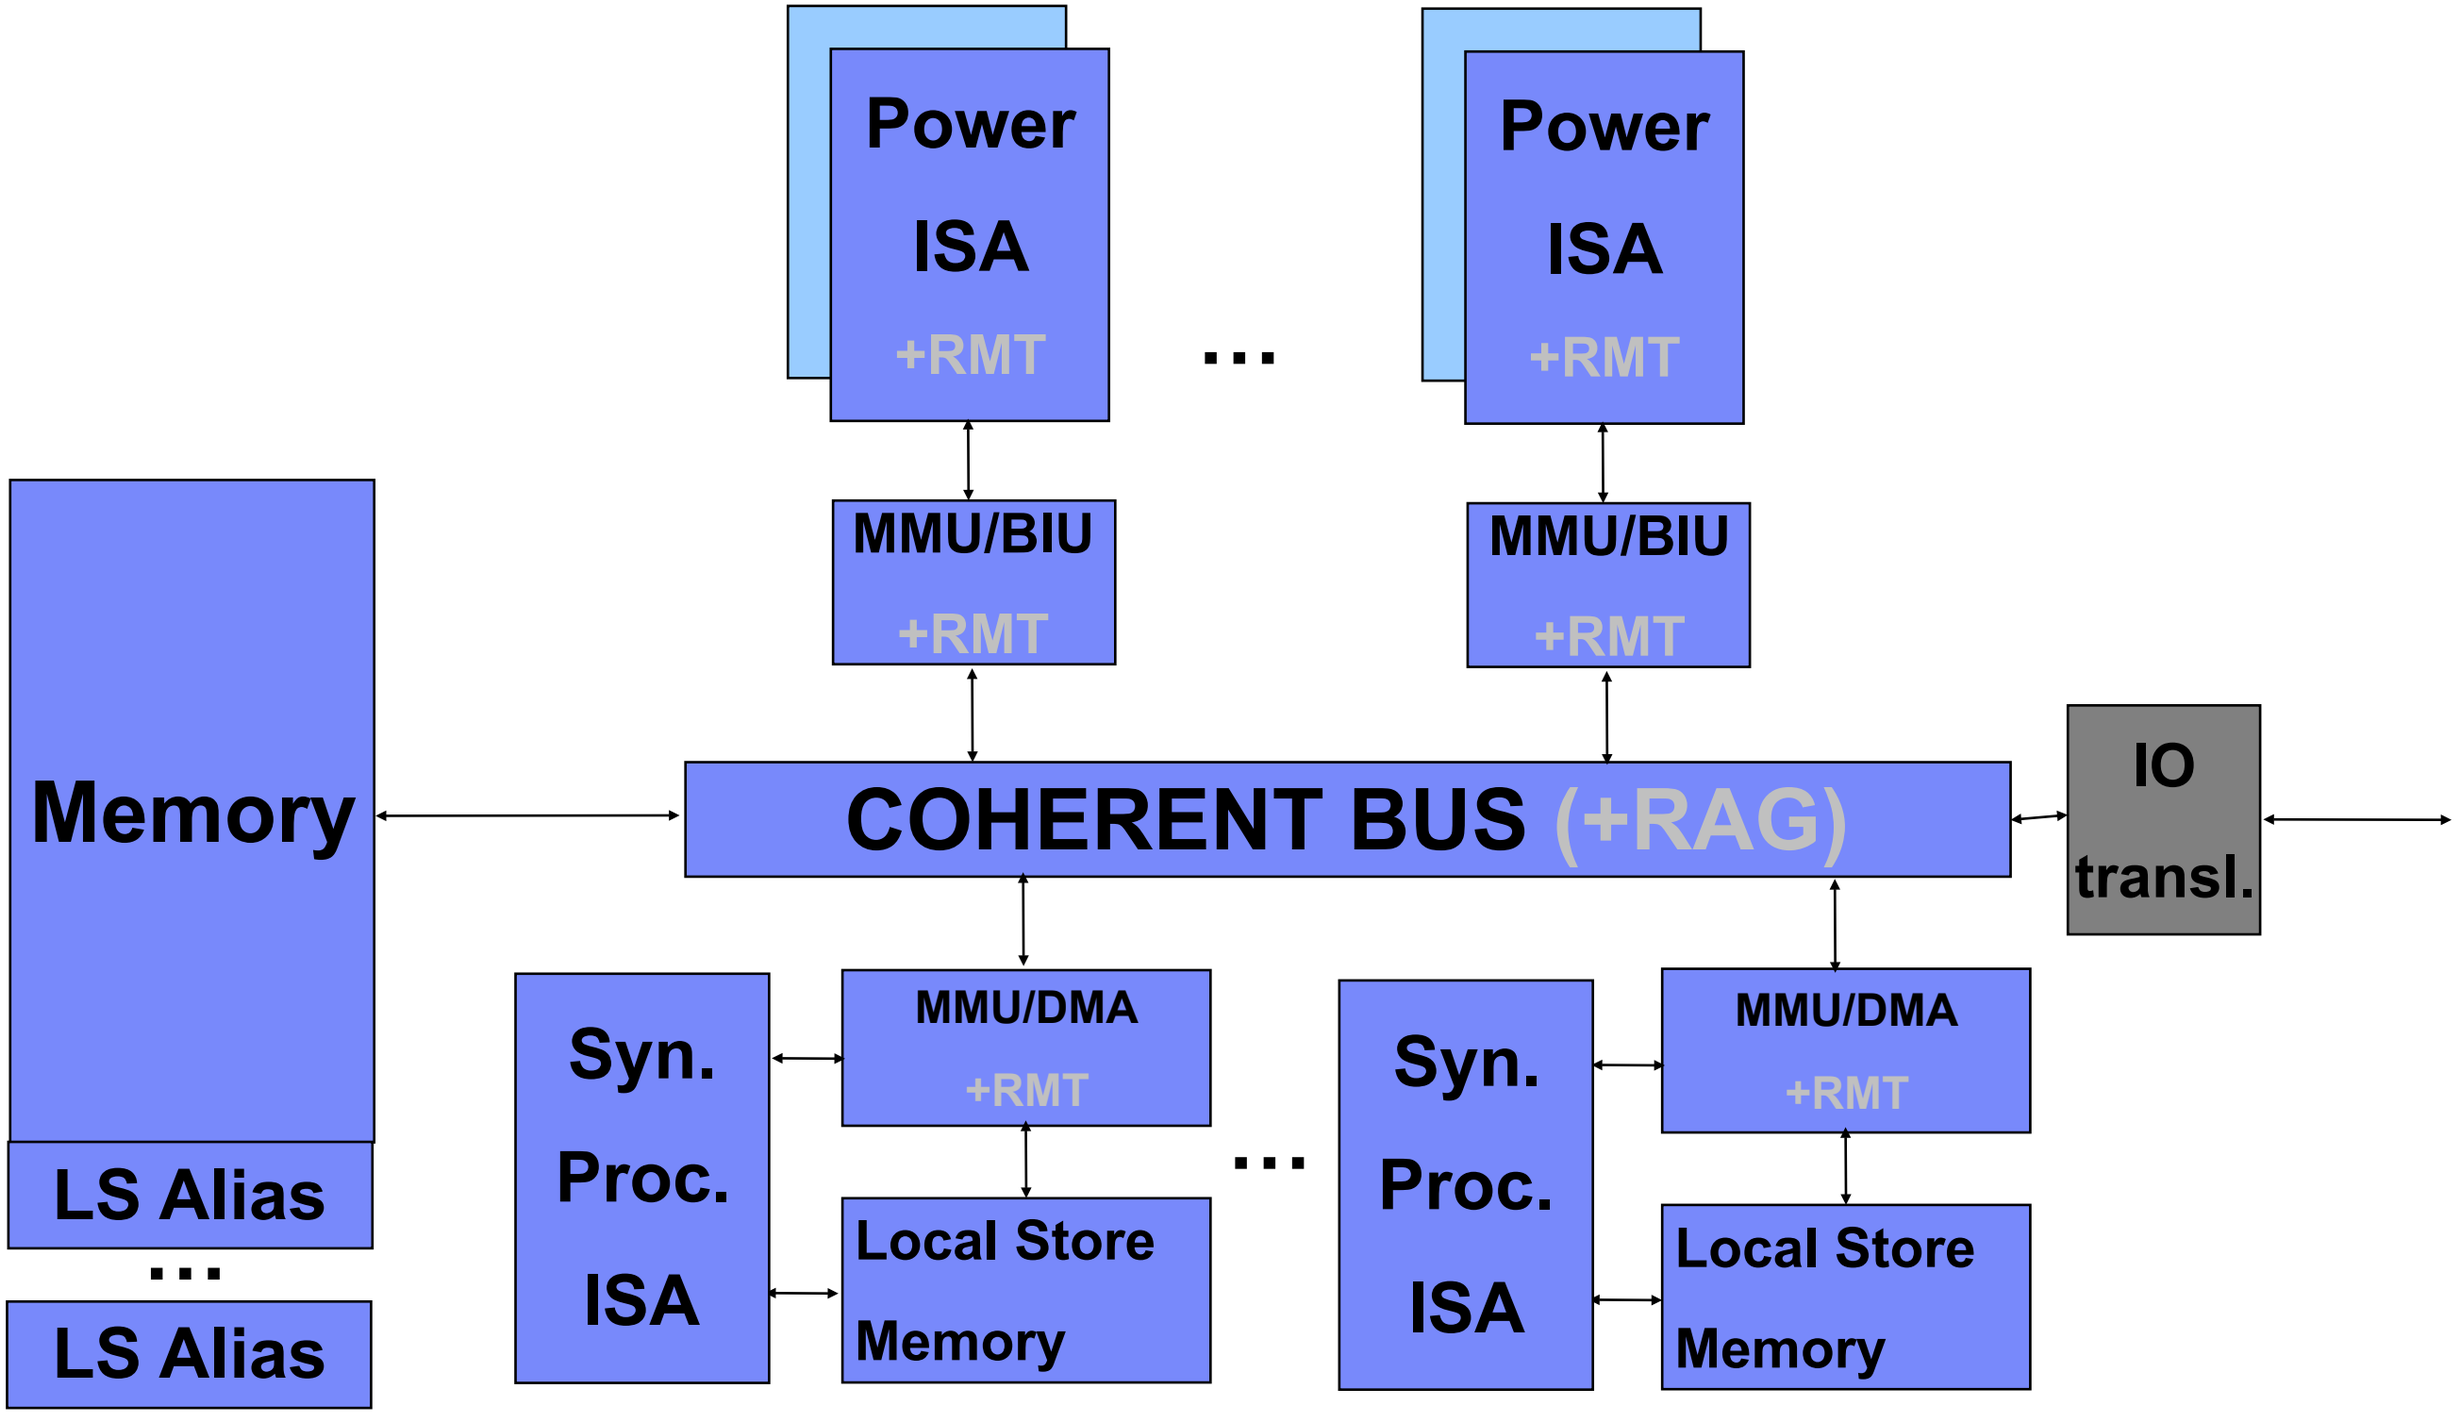

Memory Managing Processor vs. Traditional General Purpose Processor

Cell BE

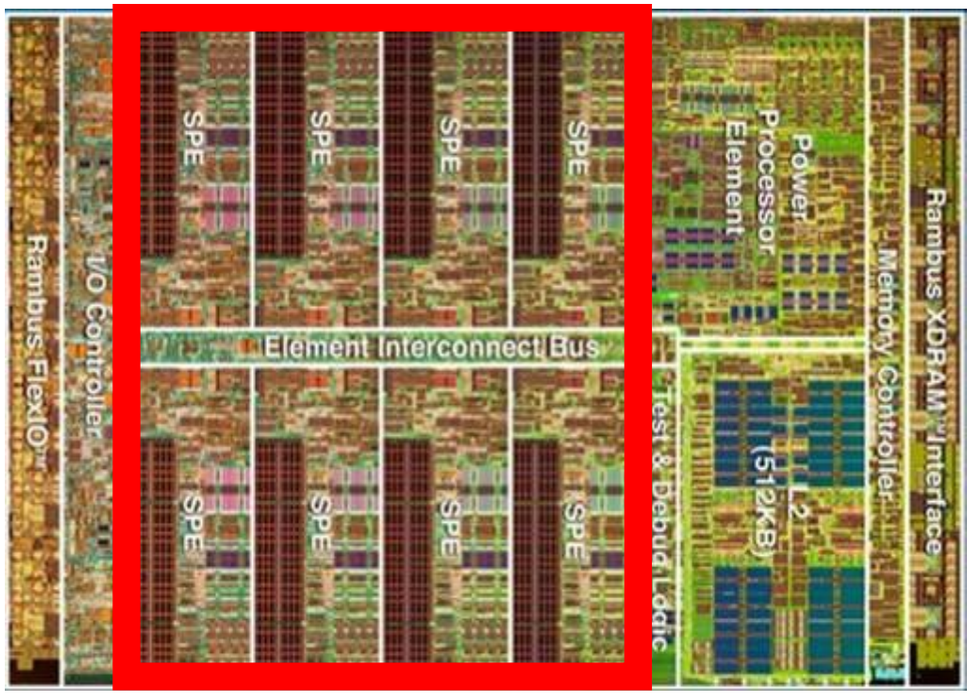

AMD

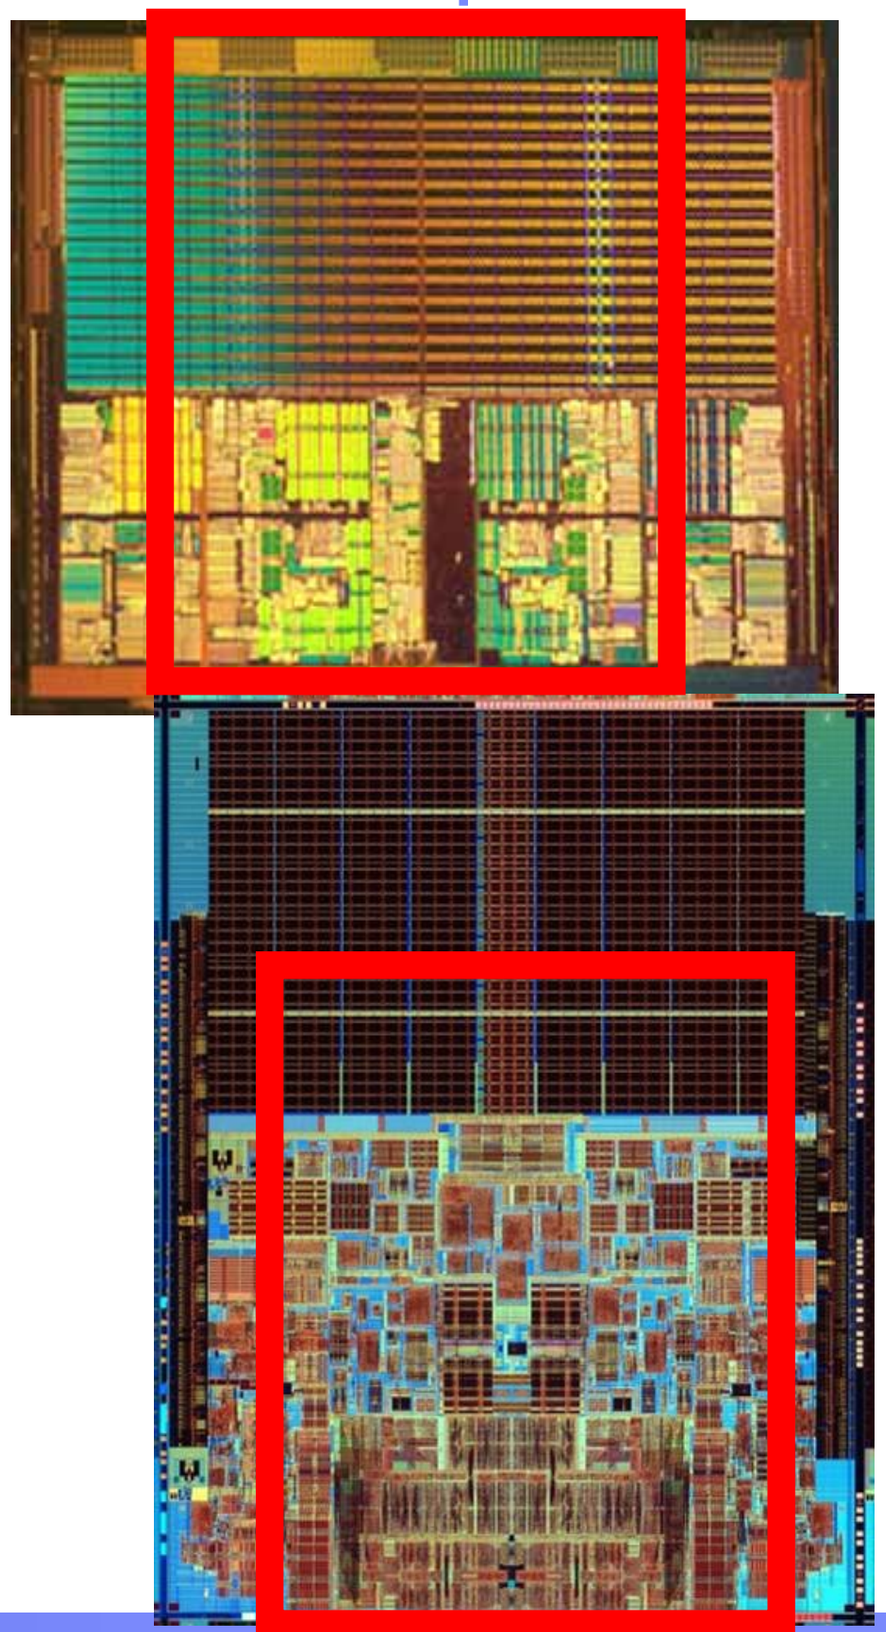

IBM

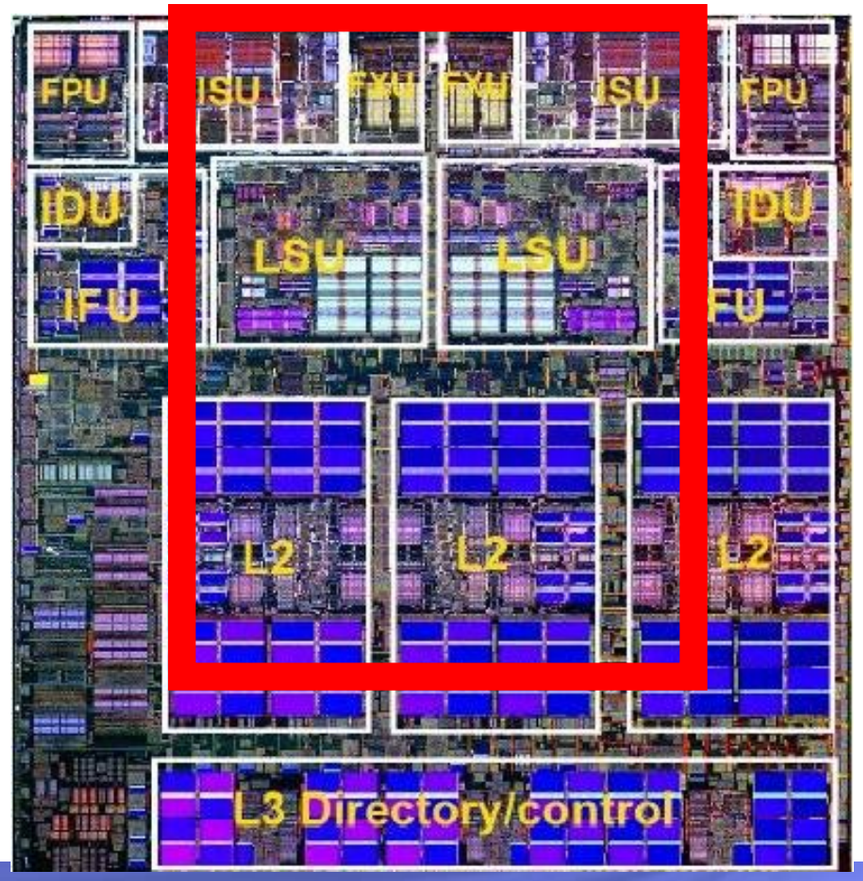

Intel

Cell

What do you most associate Cell with?

Cell

PlayStation 3 ?

Cell

Roadrunner Supercomputer?

Cell

Difficult to program?

Programming Cell

- Cell memory model
- A near miss?
- Tasks
- CellSuperScalar
- OpenCL
- OpenMP

Is shared coherent memory enough? …

Apache Arrow

```
Schema X {
   A: Float (nullable)
   B: List<Char>
   C: Struct{
         E: Int16
         F: Double
      }
}
```

Arrow terminology:

Schema:
Description of data types in a *RecordBatch*

RecordBatch:
Tabular structure containing *arrays*

Arrays:
Combination of *buffers*, can be nested

Buffers:
Contiguous C-like arrays

Index	A	B	C
0	1.33f	beer	{1, 3.14}
1	7.01f	is	{5, 1.41}
2	∅	nice	{3, 1.61}

Arrow in-memory example

Index	Data
0	1.33f
1	7.01f
2	X

Index	Valid
0	1
1	1
2	0

Index	Offset
0	0
1	4
2	6
3	10

Offset	Data
0	b
1	e
2	e
3	r
4	i
5	s
6	n
7	i
8	c
9	e

Index	Data
0	1
1	5
2	3

Index	Data
0	3.14
1	1.41
2	1.61

Johan Peltenburg e.a., TU Delft

Old Way

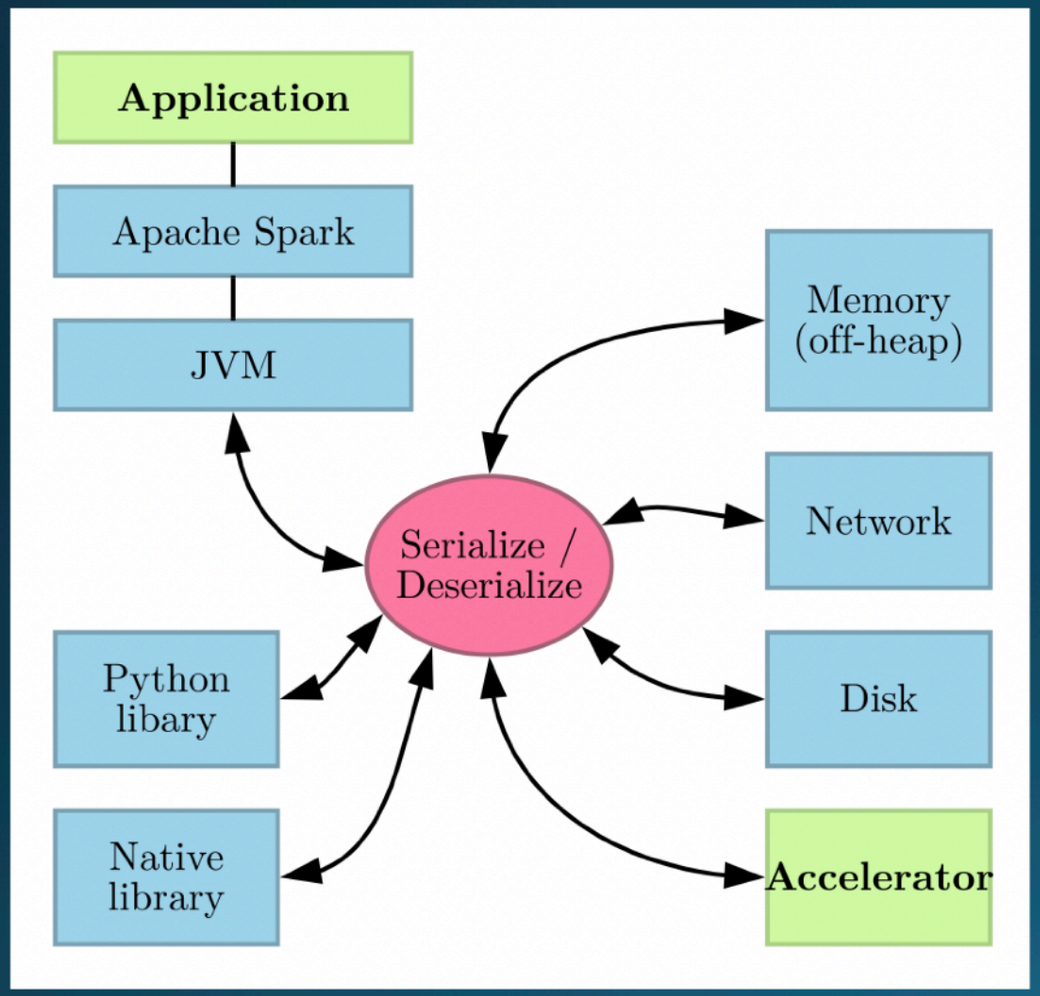

Apache Arrow & Fletcher

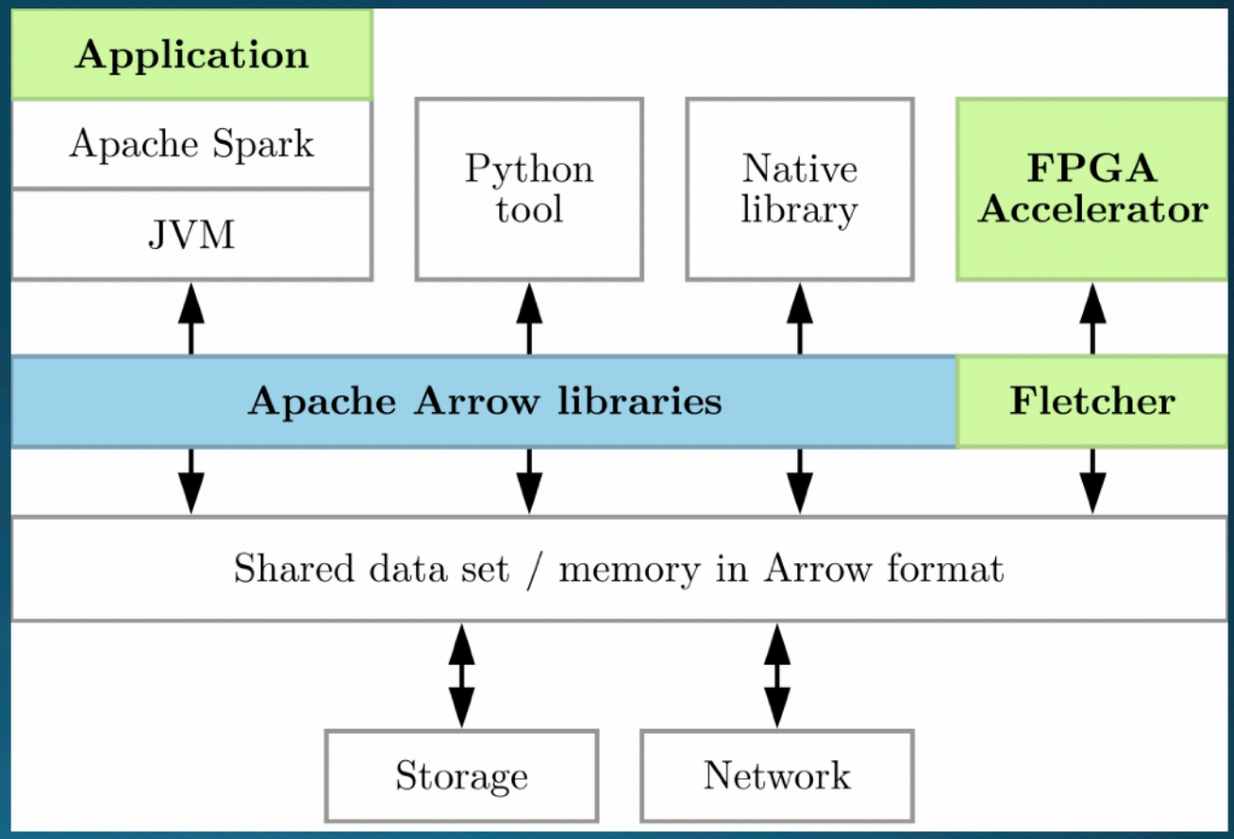

J. Peltenburg, e.a., TU Delft (OpenPOWER Summit USA 2018)

Fletcher in Action

R=16 different regular expressions per unit

AWS EC2 F1:
- Virtex Ultrascale+
- N=16 regex units
- 256 regexes being matched in parallel

POWER8 CAPI (Supervessel, & soon at Nimbix):
- AlphaData KU3 (Kintex Ultrascale)
- N=8 regex units
- 128 regex being matched in parallel

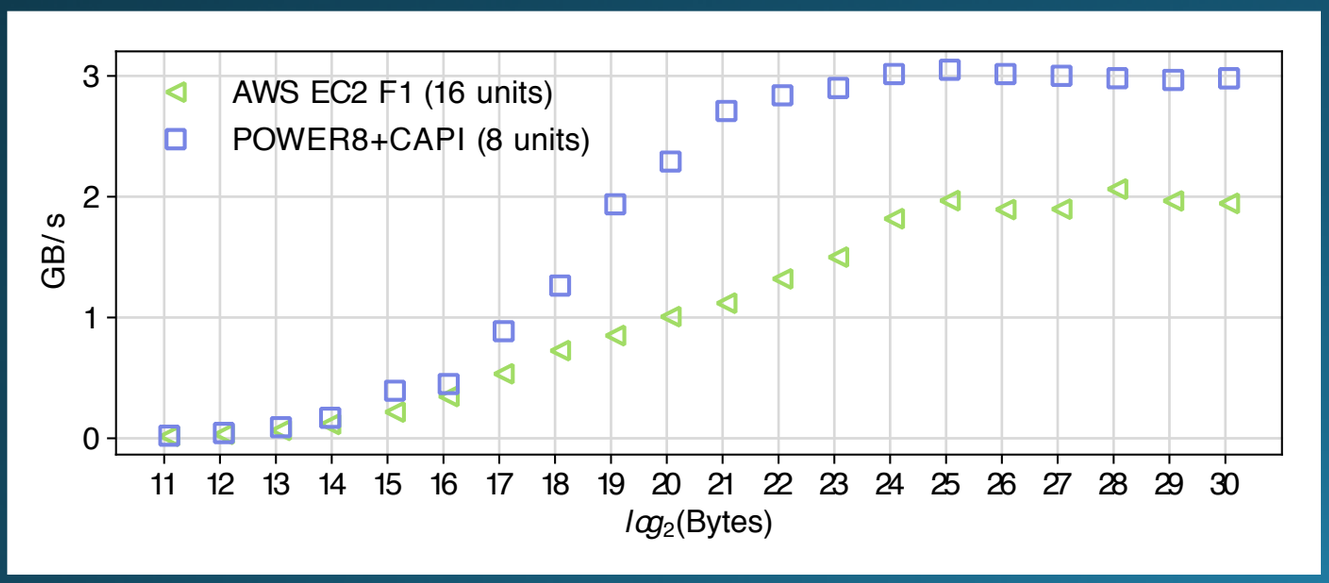

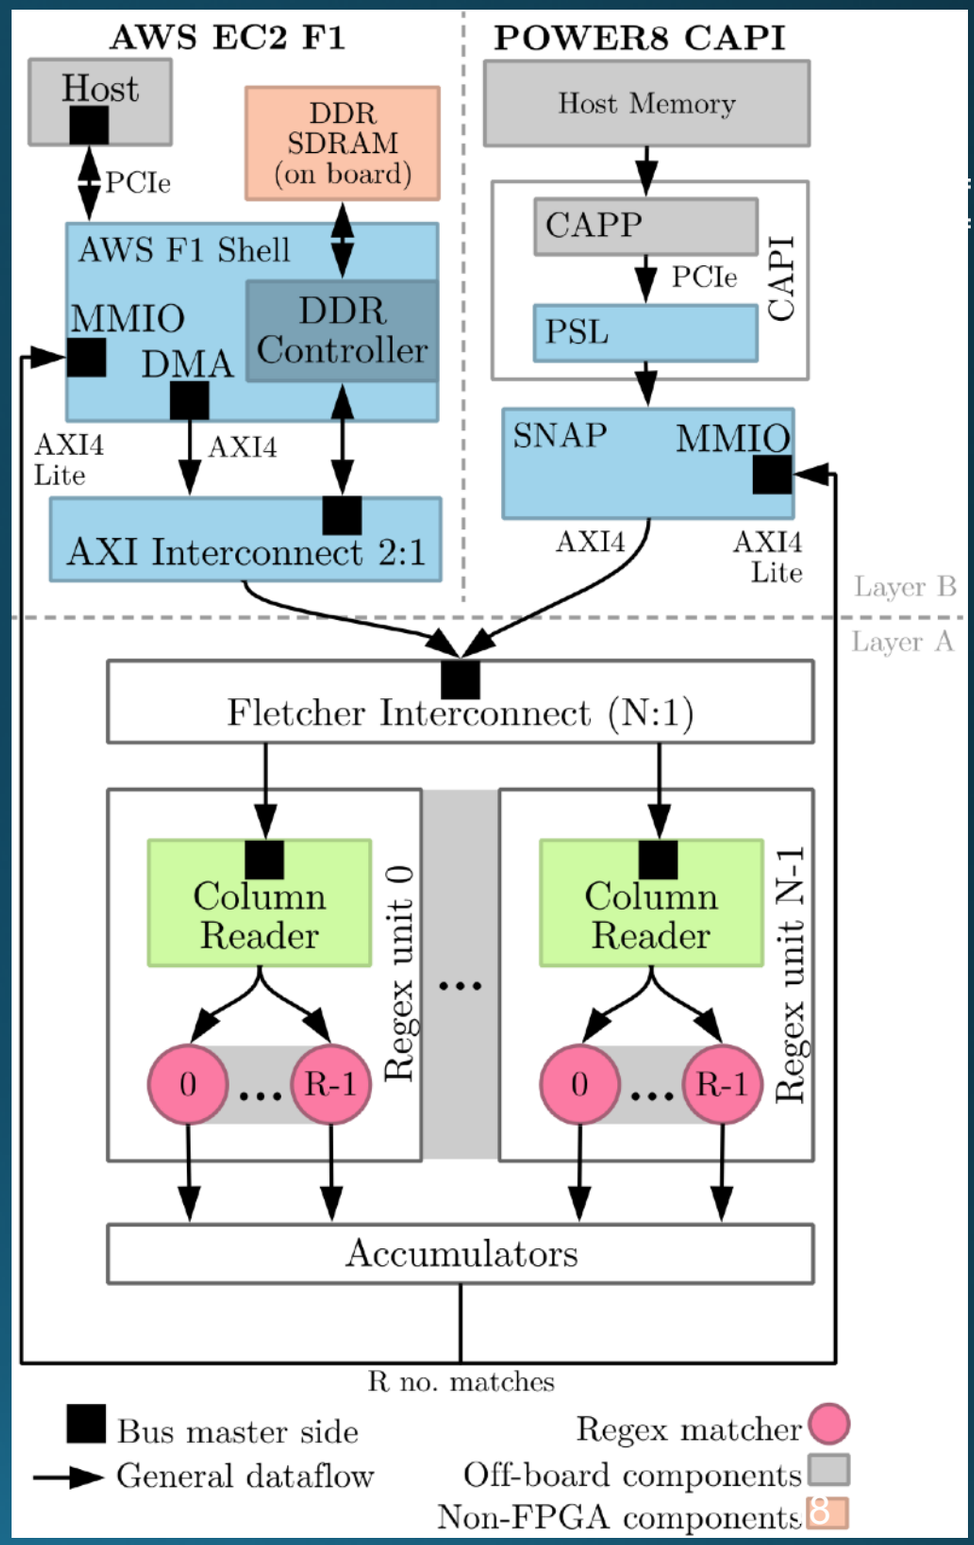

Johan Peltenburg e.a., TU Delft

Sharing Memory

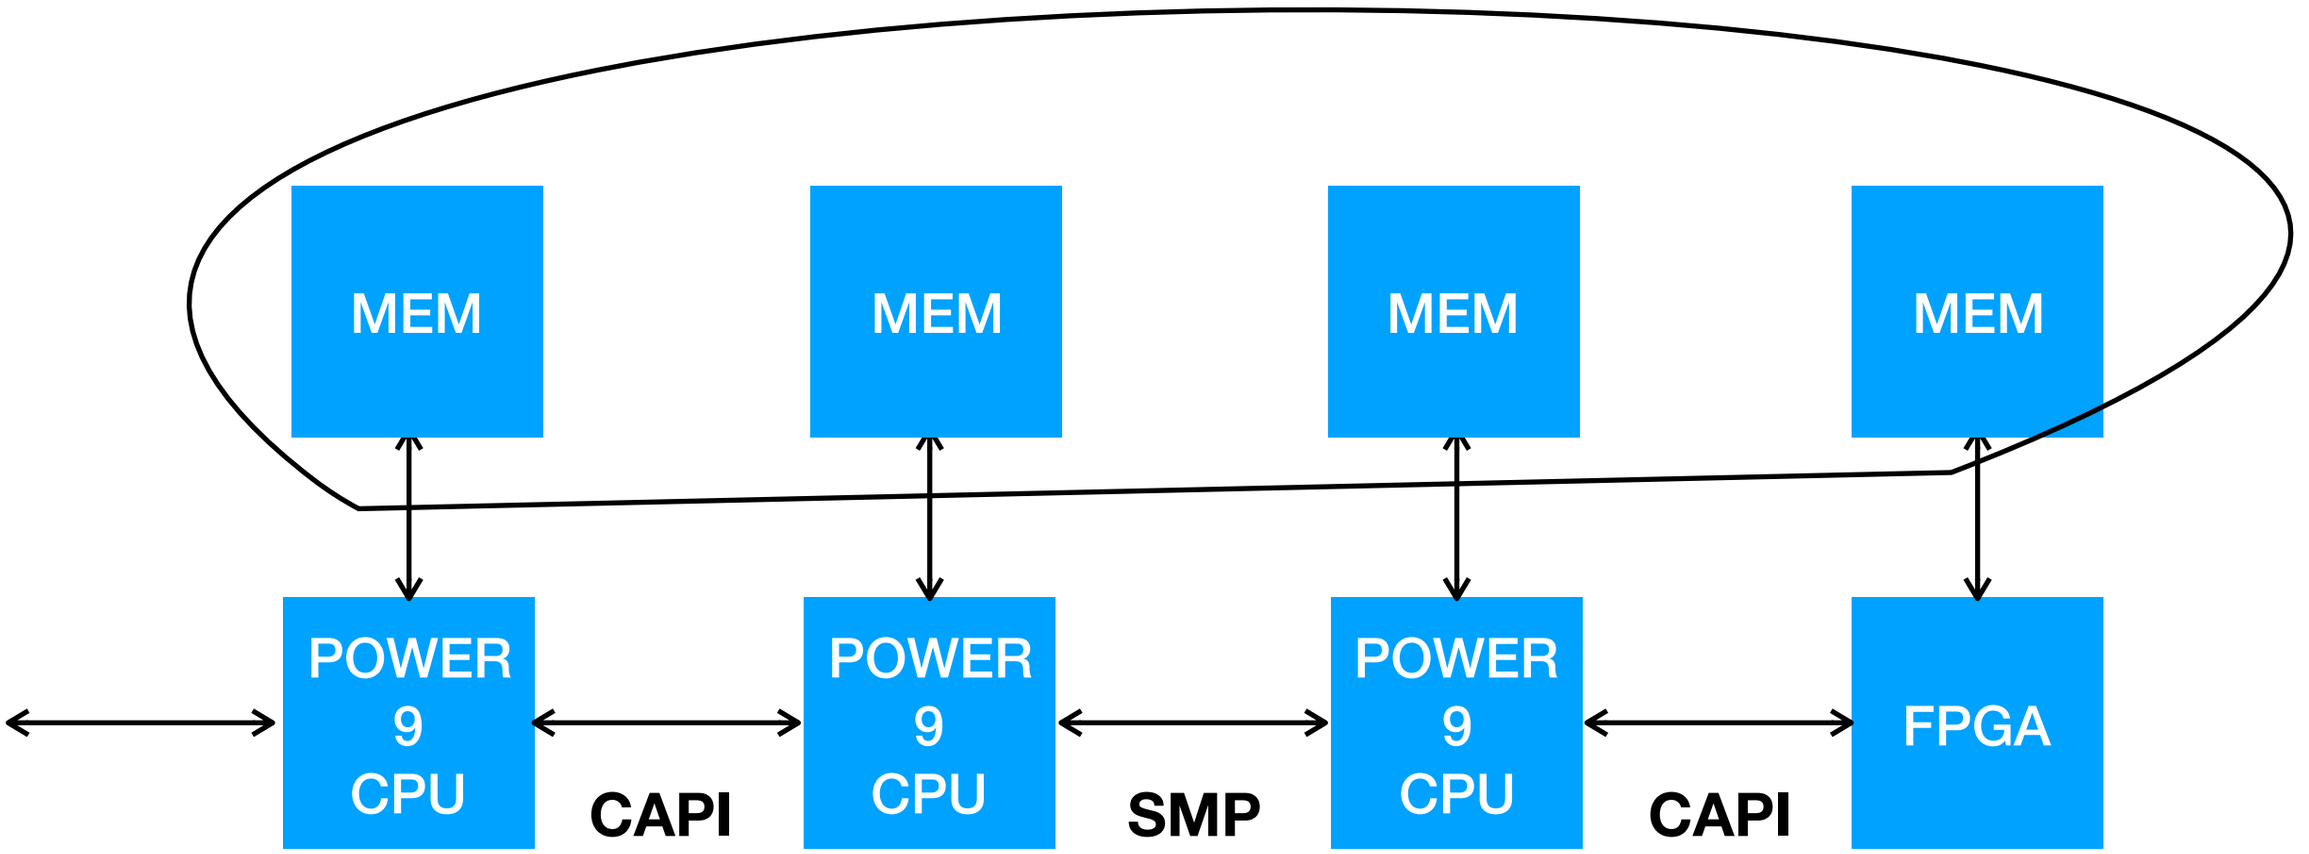

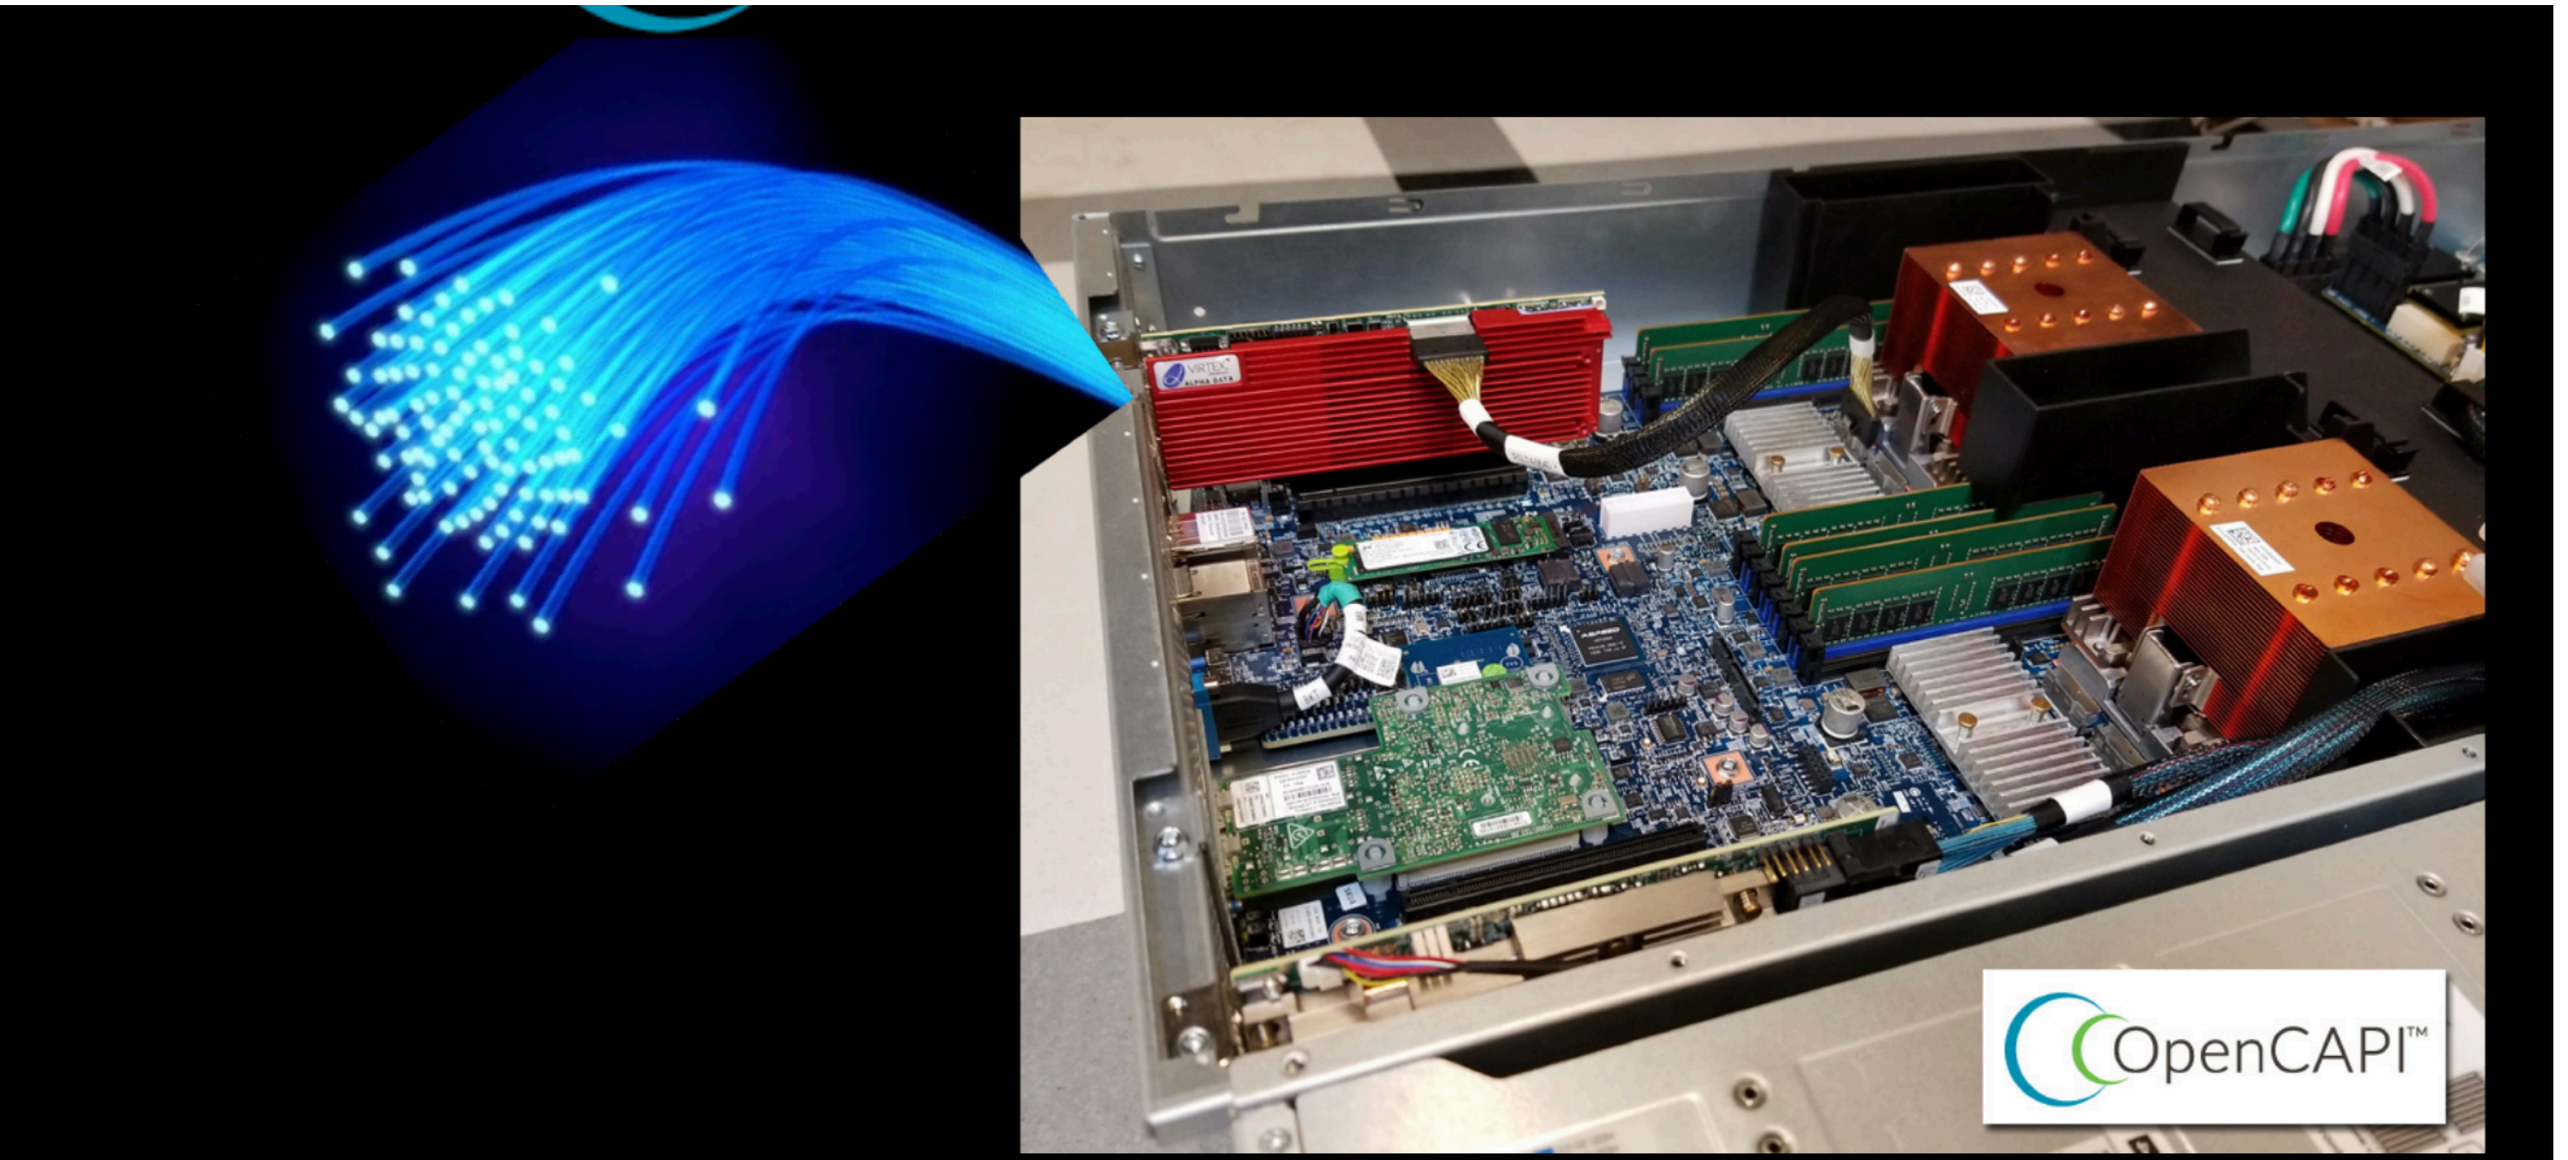

Dimitris Syrivelis, IBM Research - Ireland

OpenCAPI Extended Main System Memory over OpenCAPI
OpenPOWER Summit Europe

NVIDIA RAPIDS

- A set of libraries that operate on Apache Arrow-based data in GPU memory

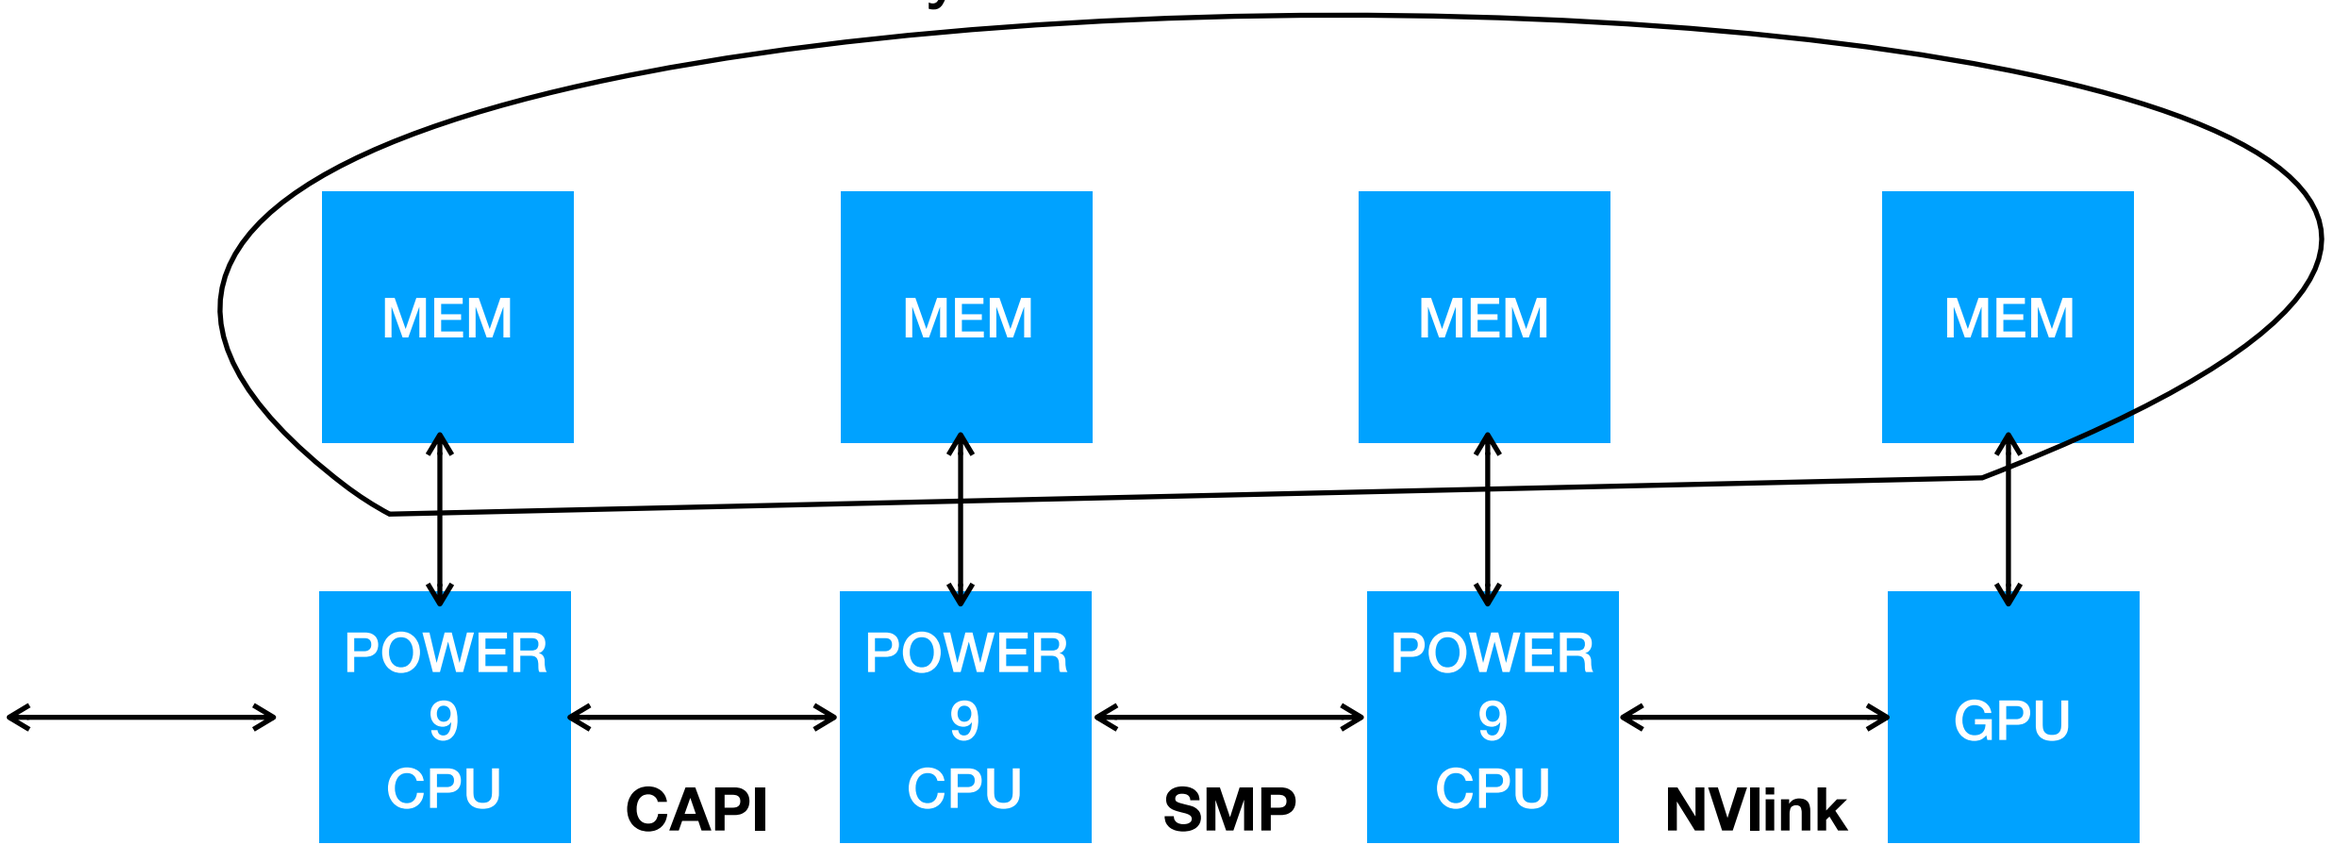

Conclusions

- The curve is flattening … we need some (new) ideas
- Preserve more of the programmer’s information
    - Tasks
    - Defined layout for data types (across languages)
- Ideas have made their way into mainstream
    - OpenMP
    - Apache Arrow