Kalray MPPA Many-Core Processors

Stefan Kosnac

22.12.2015

Picture from Kalray Website, [1].

What is Many-Core?

Many-Core is more but simpler.

- No definition, but starting somewhere between 32 to 64 cores.
- Core properties:
  - low frequency
  - low complexity
- Uses on chip interconnection networks.

A. Vajda: *"Programming Many-Core Chips"*, [2].

Why Many-Core?

Multi-Core → Many-Core

- It is not possible to switch every transistor at the same time (dark silicon).
- $P = \alpha C_{Load} V_{DD}^2 f$
- V_{DD} can be reduced when the frequency is reduced.
- Cache coherence is not scalable.
  →Private memory for groups of cores.
- Increasing unreliability of hardware at smaller node sizes.
  →Perform redundant calculations.

A. Vajda: *"Programming Many-Core Chips"*, [2].

Contents

Overview

- Andey (32-bit)
- 28 nm
- 400 MHz
- 70 GFLOPS DP

- Bostan (64-bit)
- 28 nm
- 600/800 MHz
- 422 GFLOPS DP

- Coolidge (64-bit)
- 16/14 nm
- 1000 MHz
- 527 GFLOPS DP

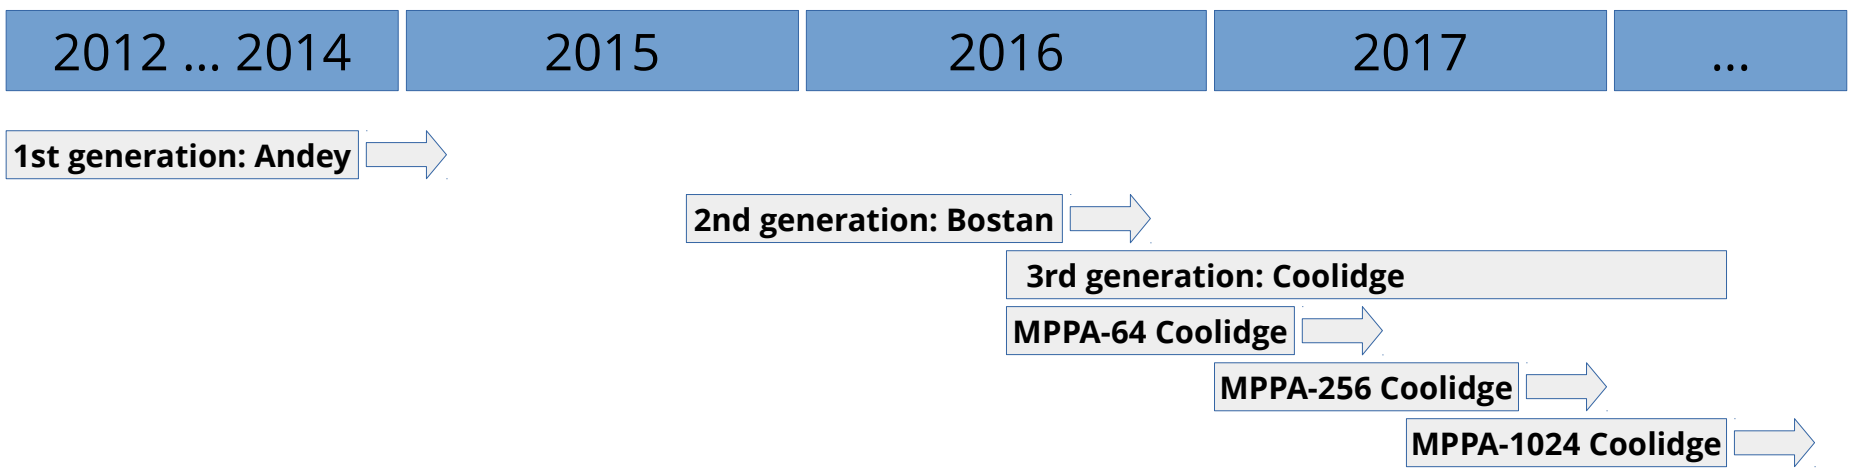

Based on information from February 2015, [5][6].

Processor

Many-Core Architecture

- Quad-Cores running Linux or a Real Time OS (RTEMS).
- Ethernet, PCIe, DDR & Interlaken interfaces.
- NoCX extends the on chip network to other MPPA chips.
- Message passing between clusters.

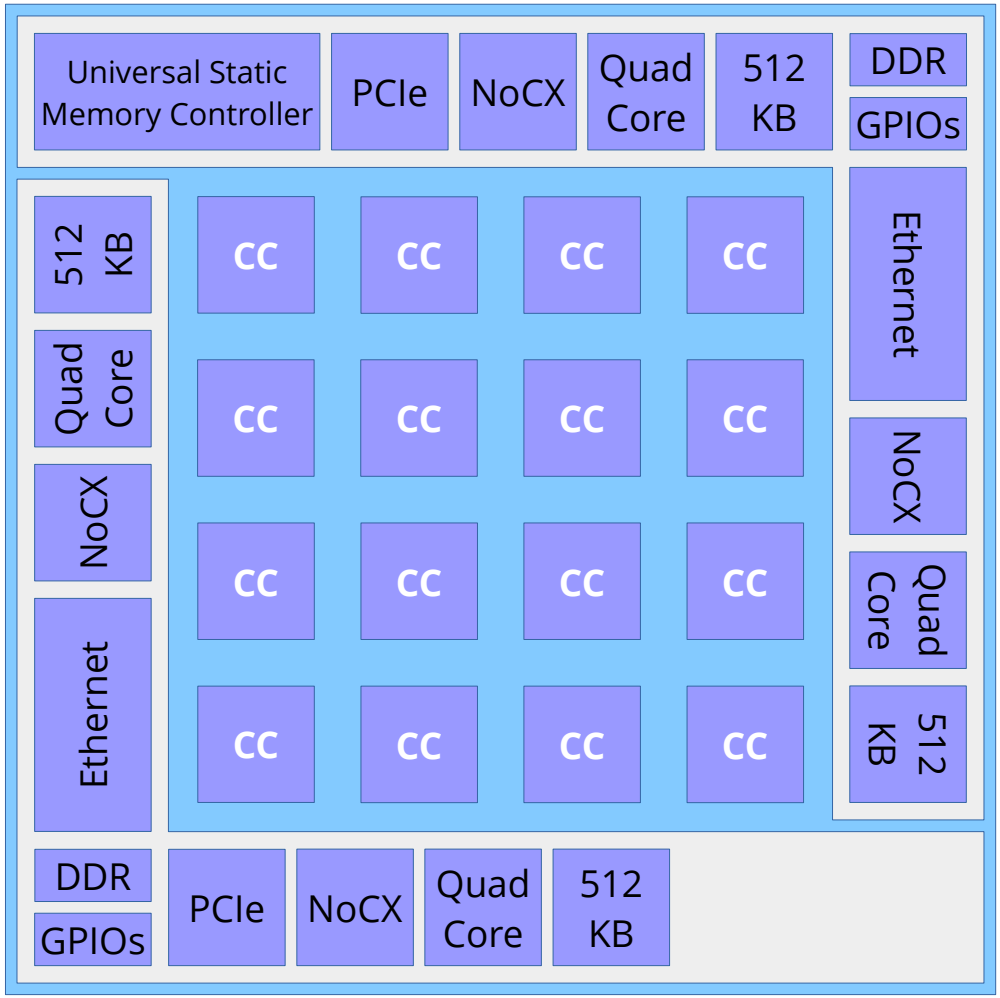

Picture based on [7].

NoC

Network on Chip

- 2D torus topology.
- D-NoC: Optimized for bulk data transfer.
- C-NoC: Optimized for small messages at low latency.
- Wormhole switching.

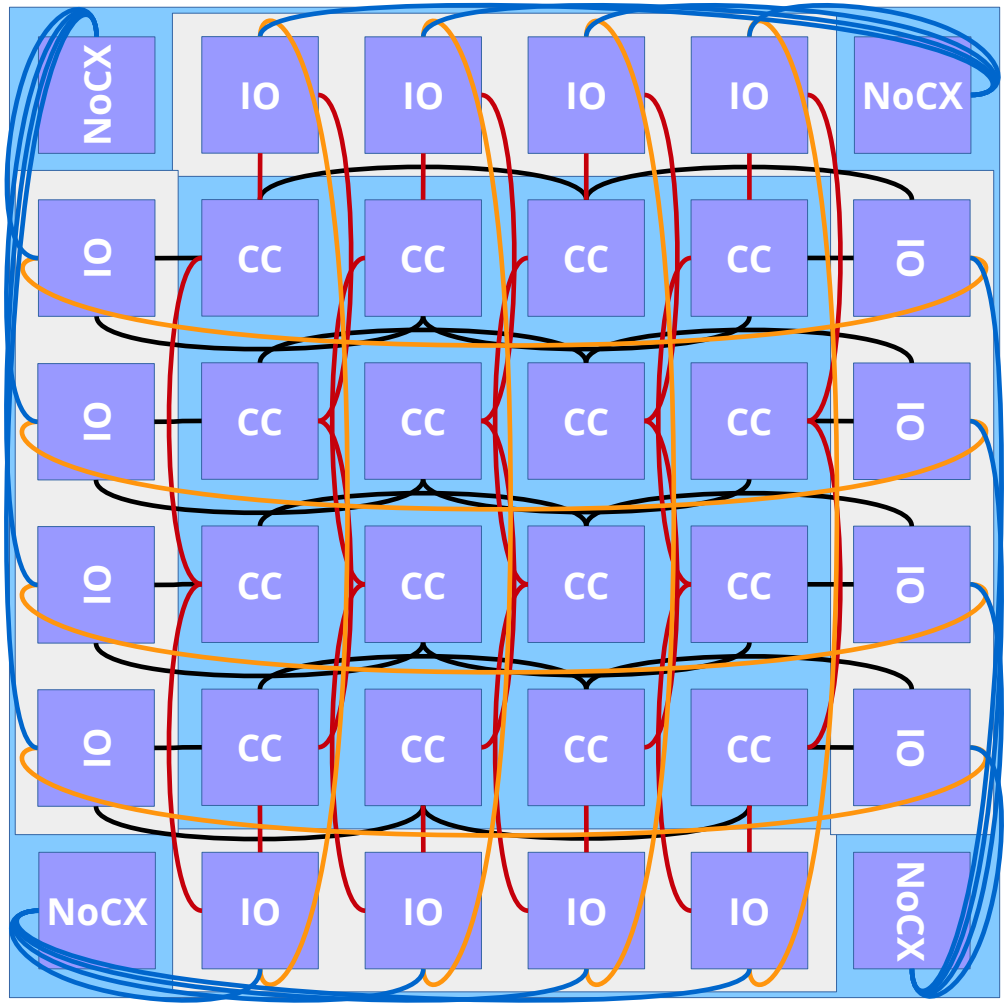

Picture based on [5].

Cluster

Compute Clusters

- 16 VLIW cores running user threads.
- 2 MB of shared memory (16 banks).
- 1 System Core.
- DMA transfers $3.2\,\mathrm{GB\,s^{-1}}$ (full-duplex).
- Debug & System Unit supports JTAG.
- *NodeOS* conforming to POSIX API.

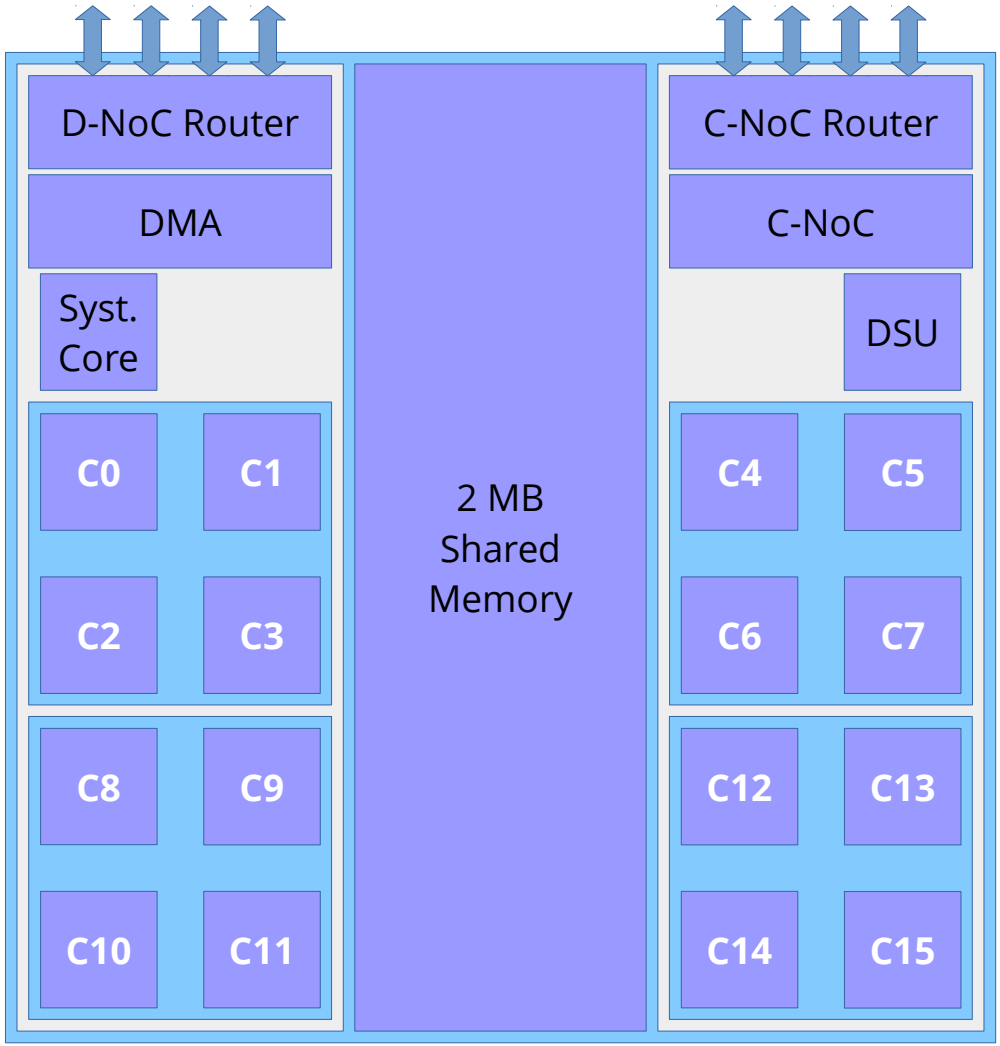

Picture based on [3].

VLIW

Very Long Instruction Word

- Scalar processors (pipelined) have $IPC \leq 1$.
- m-issue superscalar processors ($IPC \leq m$) are dynamically scheduled.
- m-issue VLIW processors ($IPC \leq m$) are statically scheduled.

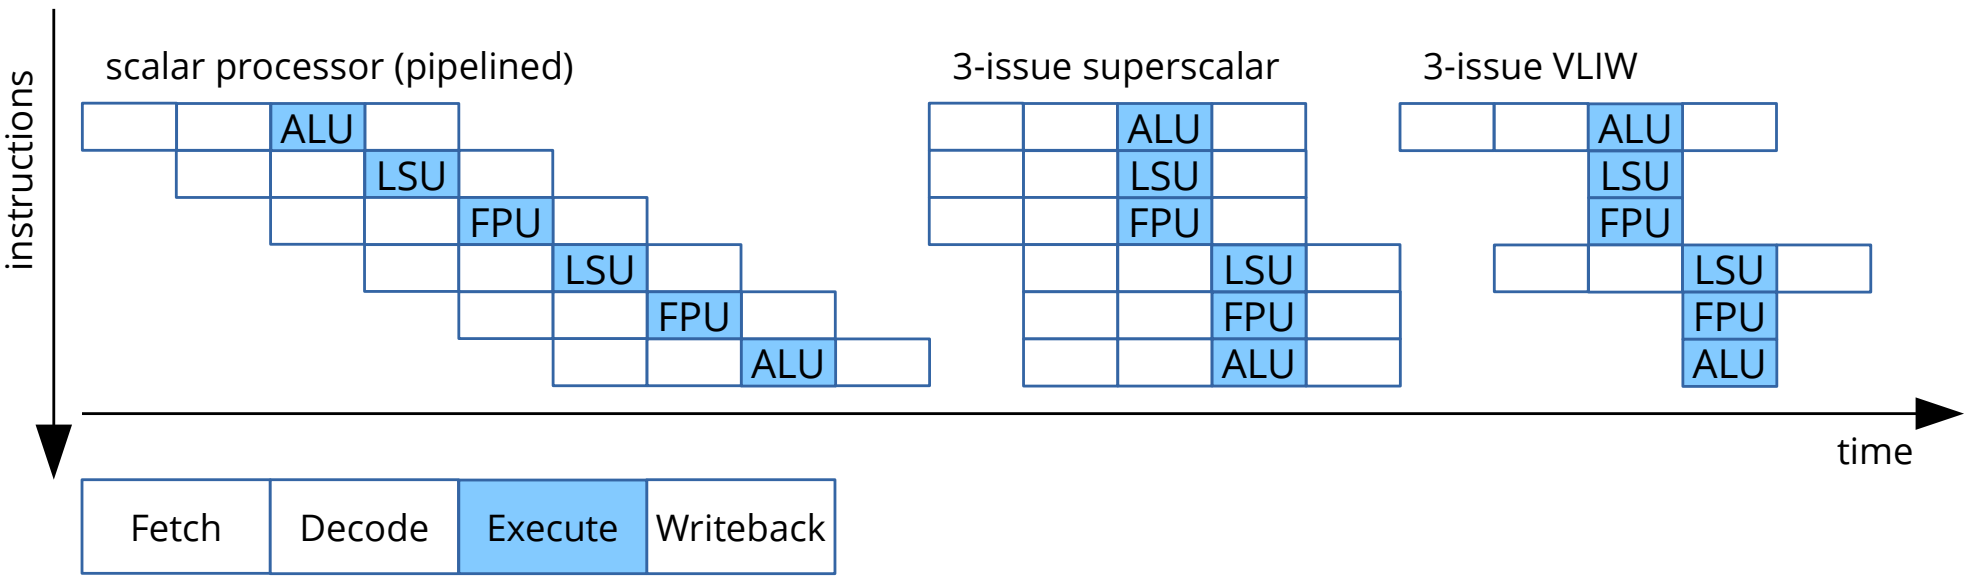

K. Hwang: *"Advanced Computer Architecture: Parallelism, Scalability, Programmability"*, 1992, [12].

VLIW-Core

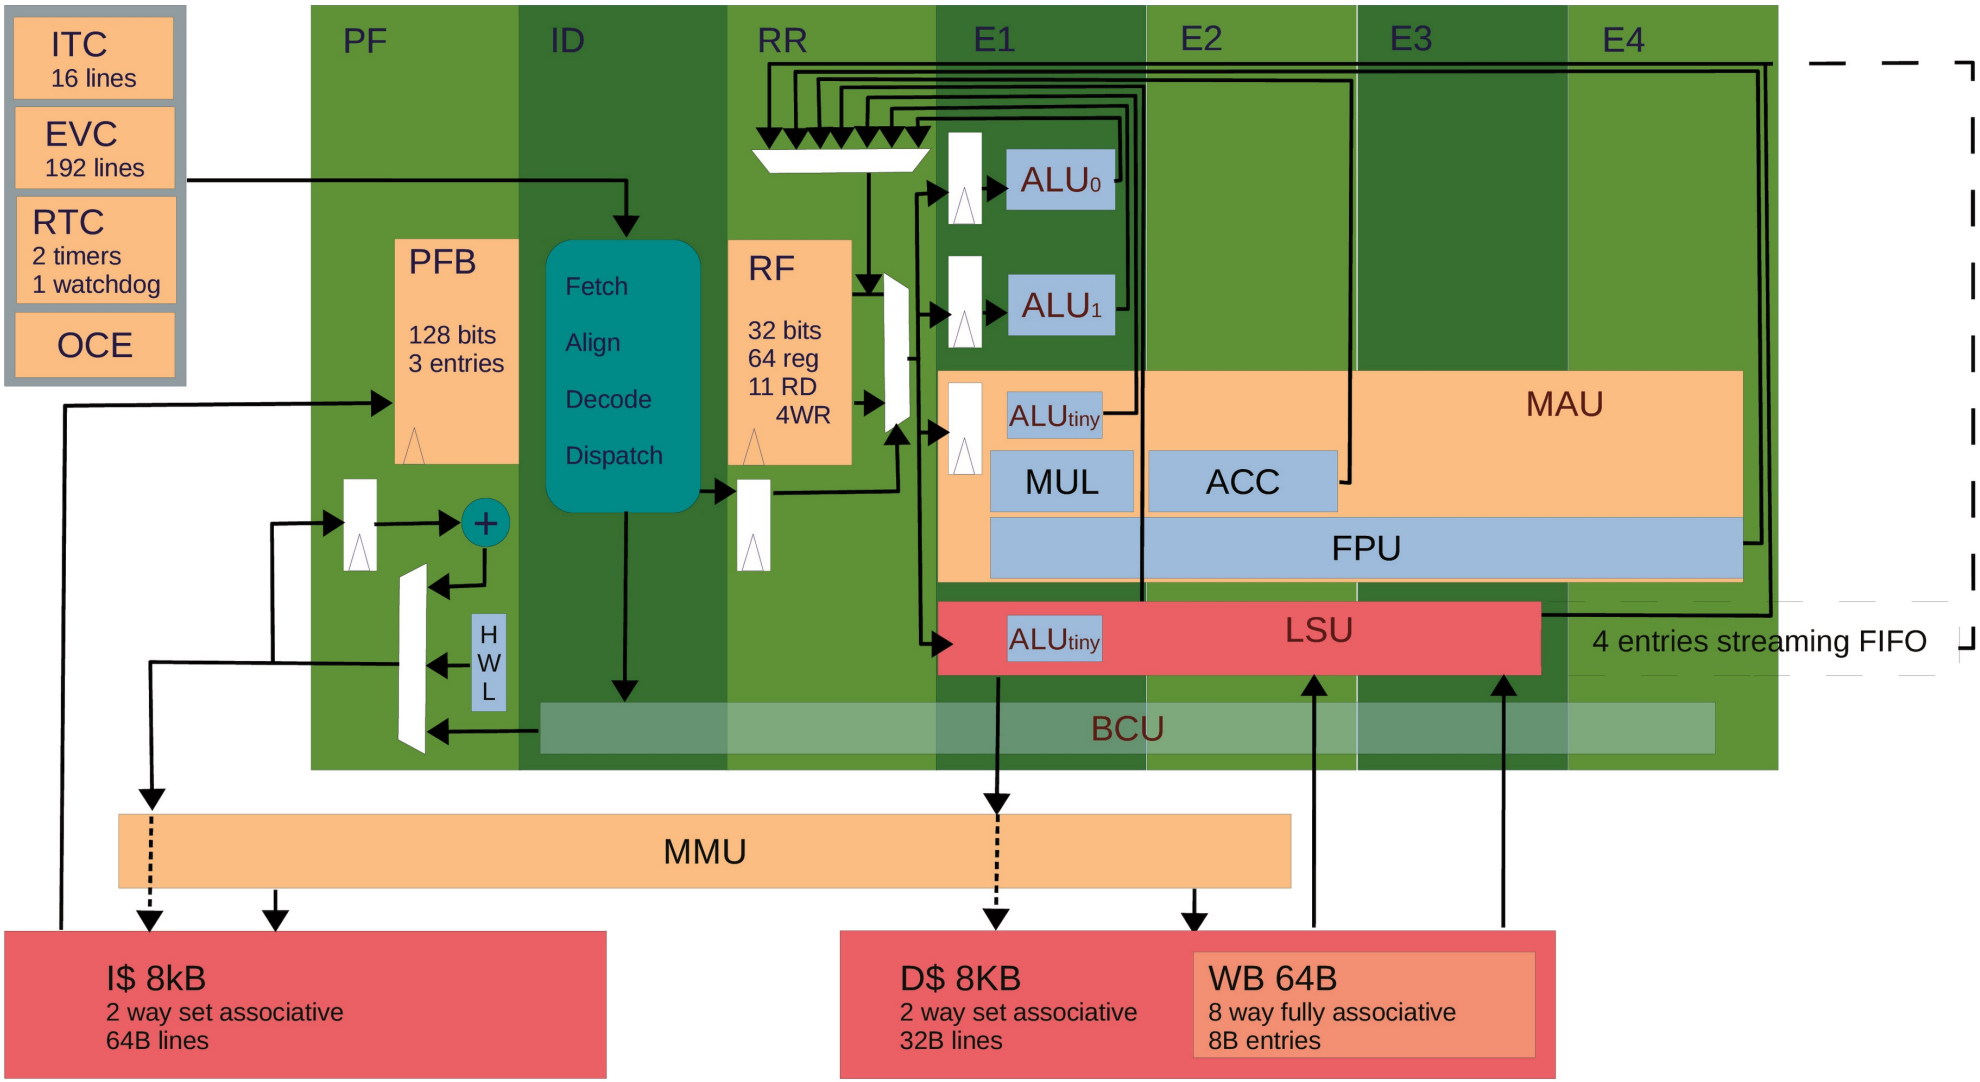

B. D. de Dinechin et al.: *"A Clustered Manycore Processor Architecture for Embedded and Accelerated Applications"*, 2013, [3].

Programming & Software

Overview

- Linux and POSIX API are available.
- Two programming models are currently supported:
  - A cyclostatic dataflow language based on C syntax called ΣC.
  - POSIX threads (and OpenMP) on compute clusters.
- Support for the Eclipse IDE.
- There are simulation, power measurement, etc. tools.
- Code transfer:
  - Adapt code for compute cluster architecture.
  - Recompilation for VLIW ISA.

Intel MIC

Many **I**ntegrated **C**ore Architecture

- Codename: *Knights Corner*
- 61 cores running at 1.238 GHz.
- Peak Double Precision Performance: 1208 GFLOPS
- Thermal Design Power (TDP): 300 W
- Code transfer from Xeon to Xeon Phi involves adding some directives but **no** general changes.

The Intel® Xeon Phi™ Product Family PRODUCT BRIEF, [8].

Tilera

Example: TILE-Gx72

- Origins at MIT (RAW research project).
- 72 RISC cores running at 1.2 GHz.
- Typical Power: 65 W
- 64-bit architecture.
- NoC has a 2D mesh topology.
- 18 MB coherent L3 cache.

TILE-Gx72 Processor PRODUCT BRIEF, [9].

GPUs

Example: NVIDIA Tesla K40

- Uses a GK110B chip.
- 2880 cores running at 745/875 MHz.
- Peak Double Precision Performance: 1430/1660 GFLOPS
- Board Power: 235 W
- Mapping code to GPUs is a bigger task.

TESLA K40 GPU ACCELERATOR, Board Specification, [10].

Seismic Wave Propagation

Overview

- *"Simulations of large scale seismic wave propagation are very important for risk mitigation, assessment of damage in future hypothetical earthquake scenarios, and oil and gas exploration."* [11]
- Equations (Einstein notation, $i, j, k \in 1, 2, 3$):

$$\begin{aligned} \rho \frac{\partial v_i}{\partial t} &= \frac{\partial \sigma_{ij}}{\partial x_j} + F_i \\ \frac{\partial \sigma_{ij}}{\partial t} &= \lambda \delta_{ij} \frac{\partial v_k}{\partial x_k} + \mu \left(\frac{\partial v_i}{\partial x_j} + \frac{\partial v_j}{\partial x_i} \right) \end{aligned}$$

Seismic Wave Propagation

Algorithm

1. Two grid point overlap due to fourth order stencil code.
2. Vertical tiling requieres good overlap between communication and computation.
3. This is done by prefetching planes to the compute clusters while computing.
4. OpenMP is used within the compute clusters.

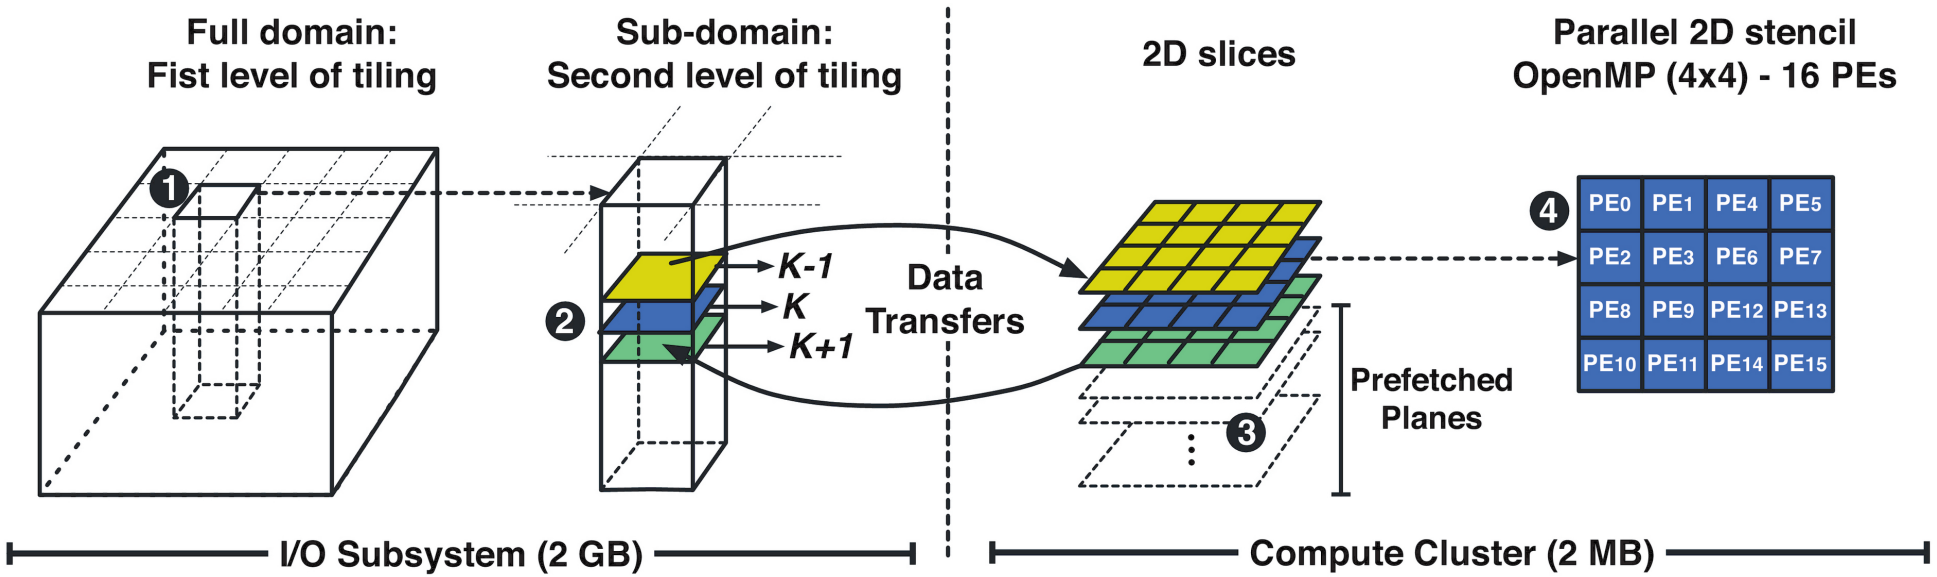

M. Castro et. al: *"Energy Efficient Seismic Wave Propagation Simulation on a Low-power Manycore Processor"*, 2014, [11].

Seismic Wave Propagation

Results

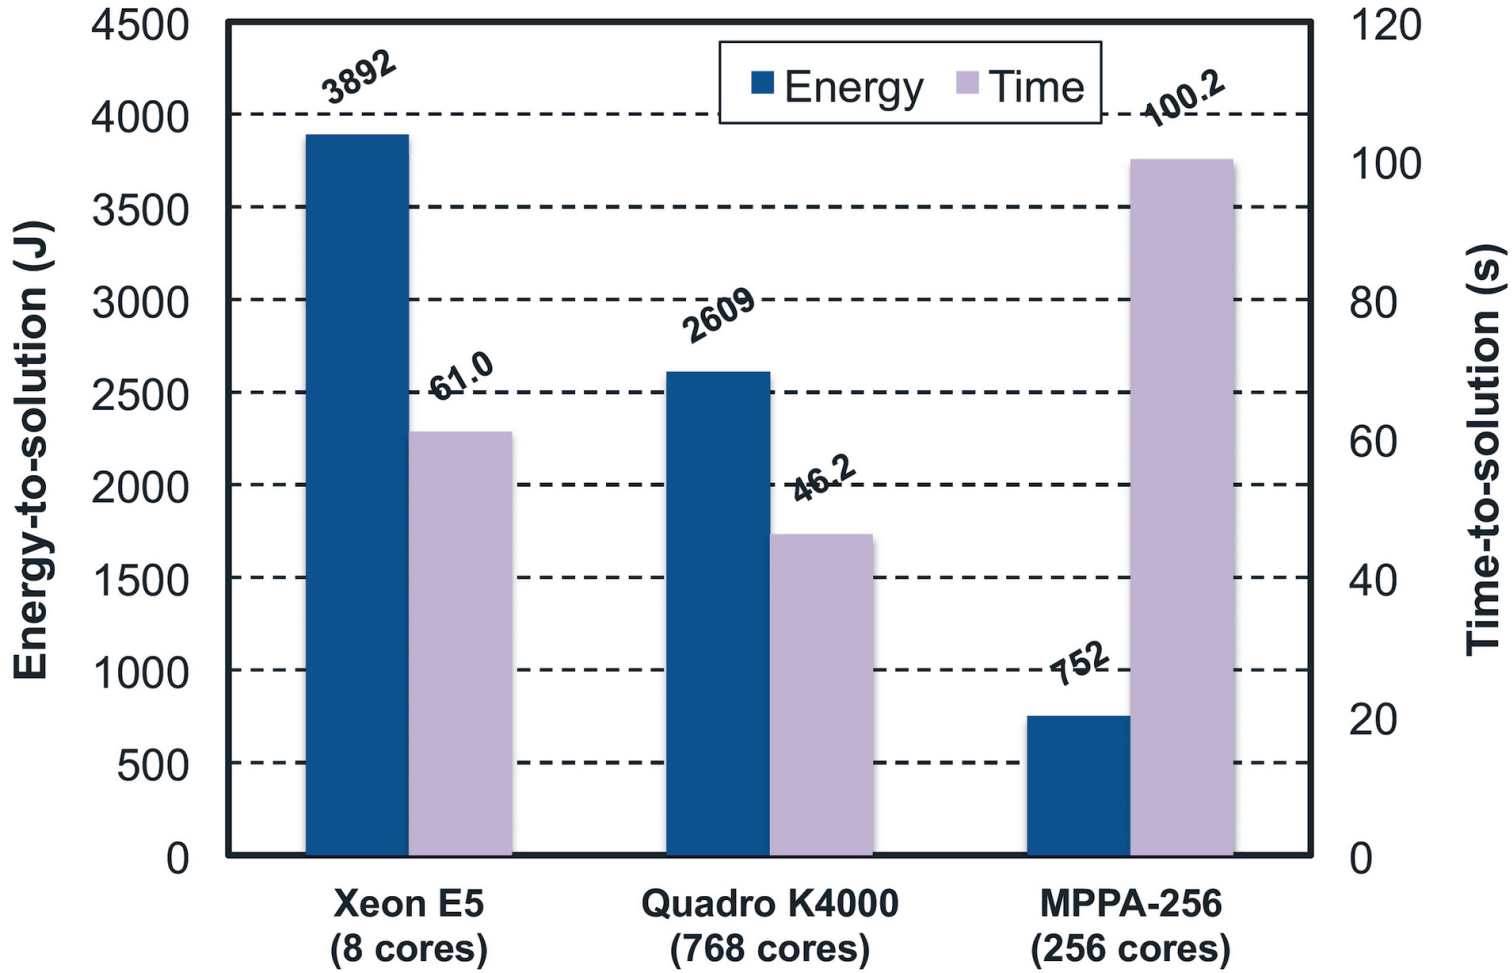

M. Castro et. al: *"Energy Efficient Seismic Wave Propagation Simulation on a Low-power Manycore Processor"*, 2014, [11].

H.264 (AVC) Encoding

Energy efficiency

- ΣC implementation compared to x264 library.
- Quality comparison with SSIM (structural similarity), PSNR and file size.
- MPPA version showed better results, because many motion vectors could be tested in parallel.

Processor	Performance	Energy efficiency
Intel Core i7-3820	49 fps	2.60 W/fps
Kalray MPPA1-256	52 fps	0.14 W/fps

B. D. de Dinechin et al.: *"A Clustered Manycore Processor Architecture for Embedded and Accelerated Applications"*, 2013, [3].

Monte Carlo Option Pricing

Overview

- *"Options are financial derivative instruments and represent a contract where the holder has the right but not the obligation to buy (or sell) an underlying asset for a determined price at the determined date."* [4]
- Black-Scholes equation:

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} = rV - rS\frac{\partial V}{\partial S}$$

$V(S,t)$: price of the option
S: stock price
r: risk-free interest rate
σ: volatility of the stock

Monte Carlo Option Pricing

Energy efficiency

Processor	Time [s]	Performance	Energy [J]
Intel Core i7-3820	13.86	0.17	1802.2
NVIDIA Tesla C2075	2.37	1.00	531.7
Kalray MPPA1-256	5.75	0.41	86.3

B. D. de Dinechin et al.: *"A Clustered Manycore Processor Architecture for Embedded and Accelerated Applications"*, 2013, [3].

Further Application Areas

- Krypto: Cryptography solutions.
- Storage applications.
- H.265 (HEVC) Encoding.
- Kalray Open Network Interface Card 80GbE.
- Time critical real time applications.

Conclusion

or what you may want to remember...

- The MPPA Many-Core architecture uses a NoC instead of relying on cache coherence.
- MPPA processors are an energy-efficient alternative to CPUs or GPUs for parallel applications.
- VLIW architectures exploit ILP and need a good compiler to make use of it.

Thank you for your attention.

Merry Christmas!

References I


Kalray Processor (picture) [25.10.2015]
http://prod.kalray.eu/wp-content/uploads/2014/11/KALRAY-Visuel-HDef-1.png


A. Vajda: *"Programming Many-Core Chips"*,
2011, Springer, chapter 2


B. D. de Dinechin et al.: *"A Clustered Manycore Processor Architecture for Embedded and Accelerated Applications"*,
2013, IEEE Conference Publications, pp. 1-6


V. Cvetanoska, T. Stojanovski: *"Using high performance computing and Monte Carlo simulation for pricing american options"*,
CIIT Conference, April 2012, Bitola Macedonia


D. Kanter and L. Gwennap: *"Kalray Clusters Calculate Quickly"*,
2015, Linley Group (article)


B. D. de Dinechin: *"Next-Generation Accelerated Computing"*,
2012, Kalray (presentation), p. 25
http://www.anciens-amis-cnrs.com/pdf/Dupont_de_Dinechin.pdf
[13.12.2015]


Product Brief: *"MPPA1®-256 Andey generation"*,
http://www.kalrayinc.com/kalray/downloads/ [04.12.2015]

References II


Product Brief: *"The Intel® Xeon Phi™ Product Family"*, http://www.intel.com/content/www/us/en/high-performance-computing/high-performance-xeon-phi-coprocessor-brief.html [14.12.2015]


Product Brief: *"TILE-Gx72 Processor"*, http://www.tilera.com/files/drim__TILE-Gx8072_PB041-04_WEB_7683.pdf [20.12.2015]


Board Specification: *"TESLA K40 GPU ACCELERATOR"*, http://international.download.nvidia.com/tesla/pdf/tesla-k40-passive-board-spec.pdf [14.12.2015]


M. Castro et. al: *"Energy Efficient Seismic Wave Propagation Simulation on a Low-power Manycore Processor"*, 2014, IEEE 26th International Symposium on Computer Architecture and High Performance Computing, pp. 57-64


K. Hwang: *"Advanced Computer Architecture: Parallelism, Scalability, Programmability"*, 1993, McGraw-Hill Publishing, chapter 4