




Kalray MPPA[®]

Manycore Challenges for the Next Generation of Professional Applications

Benoît Dupont de Dinechin

MPSoC 2013

The End of Dennard MOSFET Scaling Theory



Manycore Challenges on Next Technology Nodes

- Dark Silicon Projection (Esmaeilzadeh et al. CACM 2013)
 - “At 8nm, over 50% of the chip will be dark and cannot be utilized”
 - Based on Device x Core x Multicore models
 - Multicore model assumes x86 CPU or GPU architecture

- Dally on “Future Challenges of Large-Scale Computing” (ISC 2013)
 - Exascale computing requires 1000x improvement in energy efficiency
 - By 2020: technology => 2.2x, circuit design => 3x, architecture => 4x
 - Power goes into moving data around – communication dominates power

- Not considered above
 - Manycore platforms based on low-power CPUs and distributed memory
 - SoC nodes integrating high-speed networking and parallel processing


First MPPA[®]-256 Chips with CMOS 28nm TSMC



Available since November 2012

- High processing performance
700 GOPS – 230 GFLOPS SP
- Low power consumption - 5W
- High execution predictability
- High-level programming models
- PCI Gen3, Ethernet 10G, NoCX


MPPA[®]-256 Processor I/O Interfaces



- DDR3 Memory interfaces
- PCIe Gen3 interface
- 1G/10G/40G Ethernet interfaces
- SPI/I2C/UART interfaces
- Universal Static Memory Controller (NAND/NOR/SRAM)
- GPIOs with Direct NoC Access (DNA) mode
- NoC extension through Interlaken interface (NoC Express)


MPPA[®]-256 Processor Hierarchical Architecture

VLIW Core




**Instruction Level
Parallelism**

Compute Cluster




**Thread Level
Parallelism**

Manycore Processor



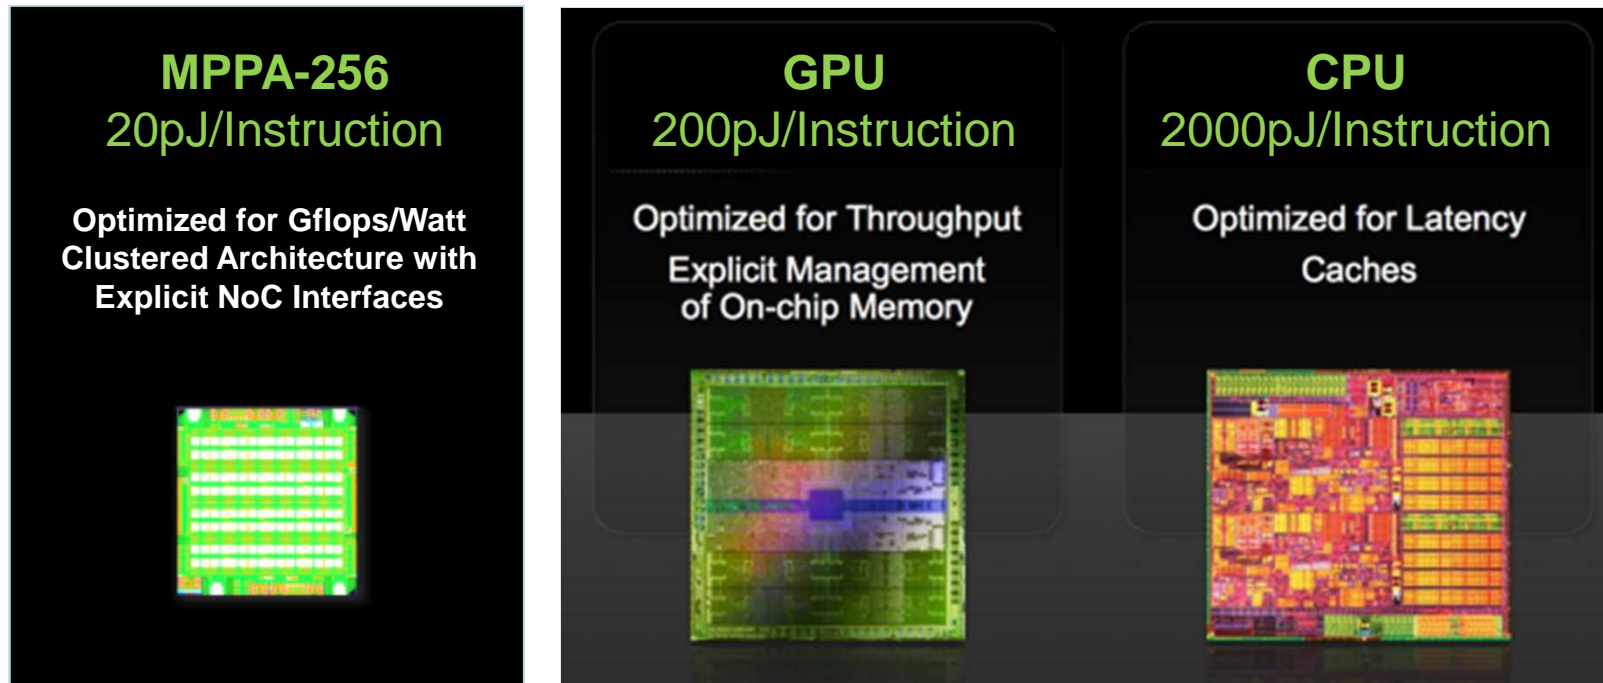
**Process Level
Parallelism**

MPPA[®]-256 Clustered Memory Architecture



- 20 memory address spaces
 - 16 compute clusters
 - 4 I/O subsystems with direct access to external DDR memory
- Dual Network-on-Chip (NoC)
 - Data NoC & Control NoC
 - Full duplex links, 4B/cycle
 - 2D torus topology + extension links
 - Unicast and multicast transfers
- Data NoC QoS
 - Flow control and routing at source
 - Guaranteed services by application of network calculus
 - Oblivious synchronization

MPPA[®] Technology Compared to GPU & CPU



Source: Bill Dally, "To ExaScale and Beyond" - NVidia

From 10x to 100x better energy efficiency vs. CPU, GPU or FPGA

MPPA[®] Architecture Compared to other Manycores

- NVIDIA, ATI, ARM generalize the GPU architecture into GP-GPU
 - Streaming multiprocessors that share a cache and DDR memory
 - Each stream multiprocessor operates multi-threaded cores in SIMT
 - **CUDA or OpenCL data parallel kernel programming models**
- Cavium, Tiler TILE Gx, Intel MIC support shared coherent memory
 - Thread-based parallel programming (POSIX threads, OpenMP)
 - **Non uniform memory access (NUMA) times, challenging cache design**
- Kalray MPPA[®] extends the supercomputer clustered architecture
 - **Clustered memory architecture scales to > 1M cores (BlueGene/Q)**
 - **Low energy per operation, high execution predictability**
 - **Stand-alone configurations, low-latency processing**

Kalray Software Development Kit

MPPA ACCESSCORE – MPPA ACCESSLIB



Today



Q4 2013

**Standard C/C++
Programming
Environment**

**Simulators,
Profilers, Debuggers
& System Trace**

**Operating Systems &
Device Drivers**



**Dataflow
Programming
FPGA Style**

**POSIX-Level
Programming
DSP Style**

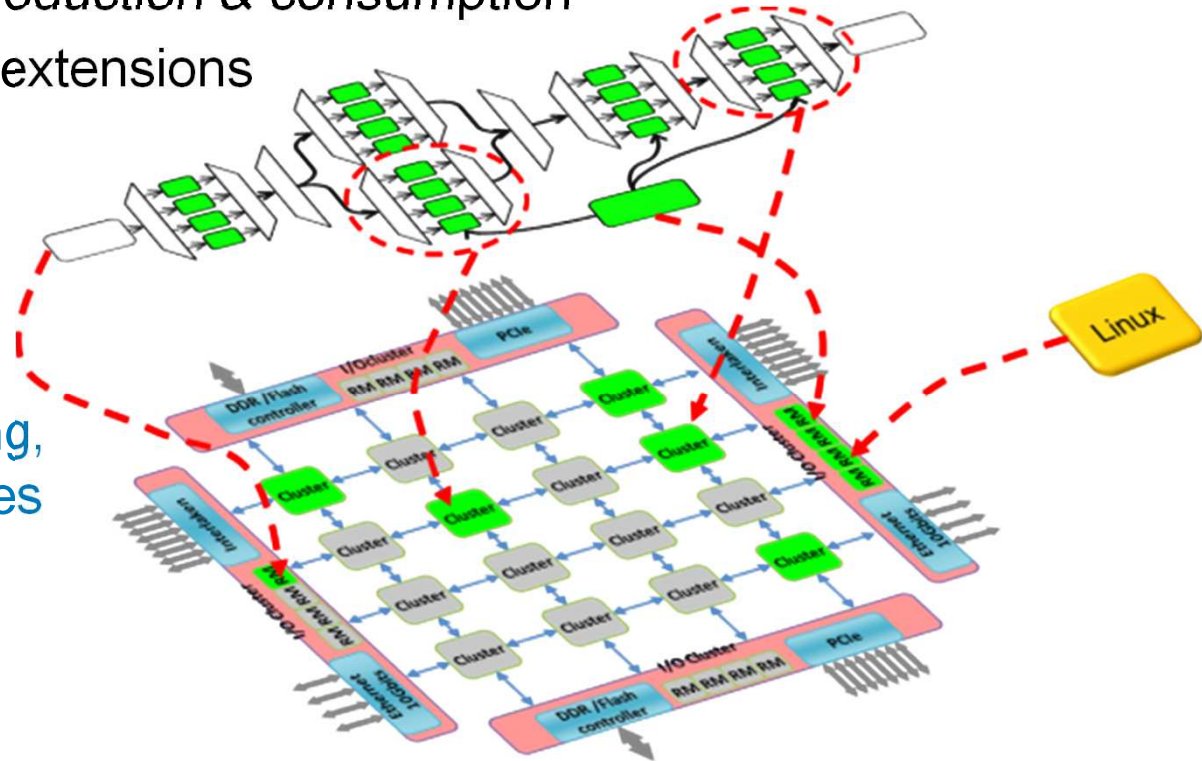
**Streaming
Programming
GPU Style**



Dataflow Programming Environment

- Computation blocks and communication graph written in plain C
- Supports “task parallelism” and “data parallelism”
- Cyclostatic data production & consumption
- Dynamic dataflow extensions

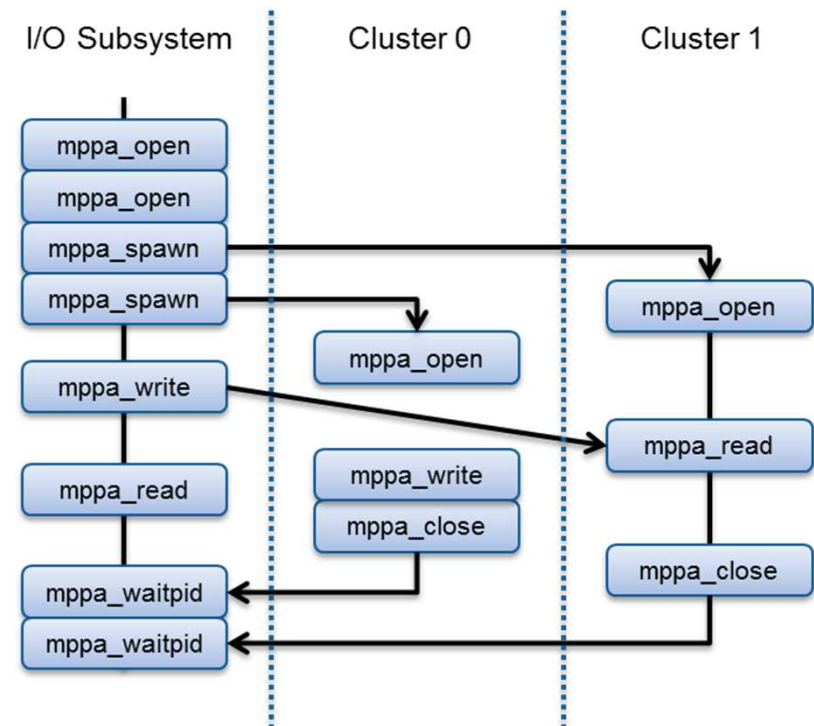
Automatic mapping on MPPA[®] memory, computing, & communication resources



POSIX-Level Programming Environment



- POSIX-like process management
 - Spawn 16 processes from the I/O subsystem
 - Process execution on the 16 clusters start with main(argc, argv) and environment
- Inter Process Communication (IPC)
 - POSIX file descriptor operations on 'NoC Connectors'
 - Extension to the PCIe interface with the 'PCI Connectors'
 - Rich communication and synchronization
- Multi-threading inside clusters
 - Standard GCC/G++ OpenMP support
 - #pragma for thread-level parallelism
 - Compiler automatically creates threads
 - POSIX threads interface
 - Explicit thread-level parallelism



Programming Environments Highlights



- Accommodate cluster memory size
 - Automatic partition, place & route of dataflow programs
 - GCC OpenMP support for the 16 user cores of a cluster

- Various programming models, from Embedded to HPC
 - Cyclostatic dataflow, from the KPN family of programming models
 - Communication by Sampling (CbS) for the time-triggered architecture
 - Bulk Synchronous Parallel (BSP) model based on Oxford BSPlib
 - Lightweight implementation of MPI (Message Passing Interface)

- OpenCL with task parallel model and distributed shared memory
 - Leverage the MMU on each core to paginate DDR memory on chip

MPPA ACCESSLIB optimized building blocks

- Provide Kalray core optimized building blocks for each scope
 - MPPA Core – register file & cache
 - MPPA Cluster – shared memory
 - MPPA Partition – distributed memory
- Delivered as C libraries
 - Dataflow programming
 - POSIX-level programming
- Numerical computing
 - FFT, Filtering and convolution
 - BLAS-level primitives
 - Matrix factorizations
 - libm extensions
- Video and image processing
 - H264, HEVC encode / decode
 - Computer vision

Scope	2013		2014	
	Q3	Q4	Q1	Q2
Core	★ DSPLIB H264 codec	★ HEVC Codec		
Cluster		★ DSPLIB		
Partition	★ Matmul, FFT		★ BLAS	★ OpenCV

Target Application Areas

INTENSIVE COMPUTING

- Finance
- Numerical Simulation
- Geophysics
- Life sciences



IMAGE & VIDEO

- Broadcast
- Medical Imaging
- Digital Cinema
- Augmented reality
- Vision



EMBEDDED SYSTEMS

- Signal Processing
- Aerospace/Defence
- Transport
- Industrial Automation
- Video Protection





TELECOM / NETWORKING

- Packet Switching
- Network Optimisation
- Security Services
- Software Defined Radio
- Software Defined Network



Computational Finance Application


- Option pricing by Monte Carlo method
- Optimized pseudo random generator
- Parallel Map / Reduce scheme across multiple MPPA processors
- Optimized mathematical primitives for Kalray core



Power efficiency 5x better than recent GPU

Audio Processing Application

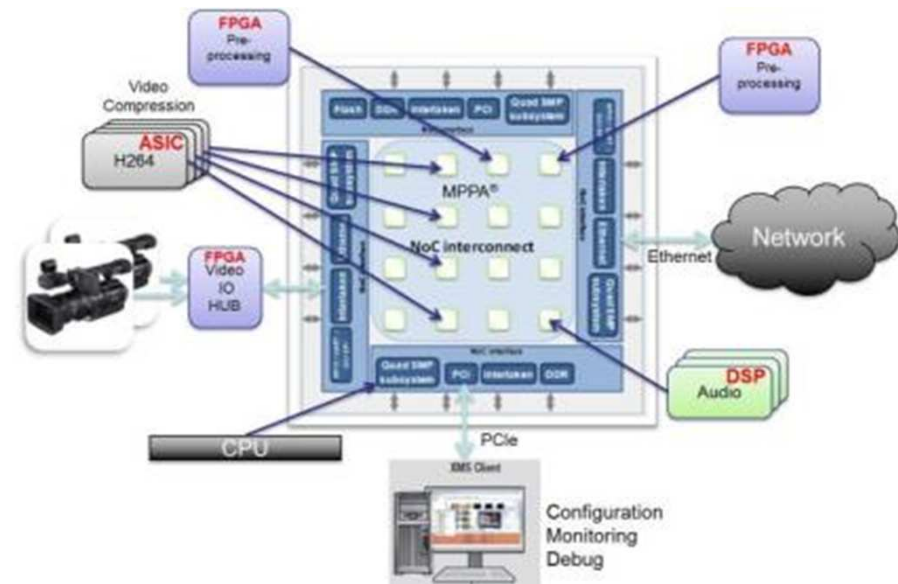

Increase performances and Reduce Audio system cost



- Multi Channel processing
 - 256 VLIW cores ~ 500 Low End DSPs
- Channel routing and control
 - High performance NoC + 32 integrated DMAs
- System integration
 - Up to 8 x Ethernet 1GbE
- Low Latency audio processing
 - 500µs latency from input to output samples
- Cost effective
 - Equivalent to complex multi DSPs + FPGAs system

Video Broadcasting Example

- High definition H264 encoder on one MPPA[®]-256
- System integration, lower power and cost
- Heterogeneous implementation
- Flexibility & scalability




H264 Encoder running on MPPA-256 at less than 6W

Video Protection Example

- Improved Content analysis
 - High resolution camera / low false detection rate
 - Robust algorithms → high performance computing of MPPA
 - Real Time detection
 - More simple infrastructure → Compute power at the source

- System integration: Ethernet input / decode / content analysis / encode

- Multi camera system on single chip



Augmented Reality Example

- Assisted operation & maintenance
 - ARMAR (Augmented Reality for Maintenance and Repair)


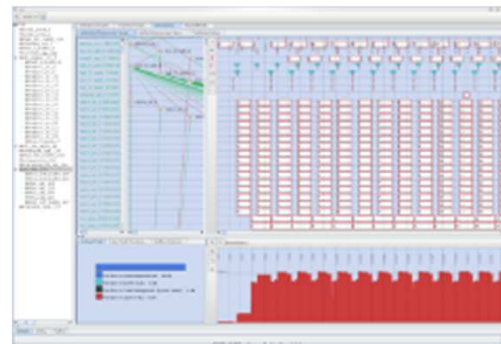



- Assisted conformity control



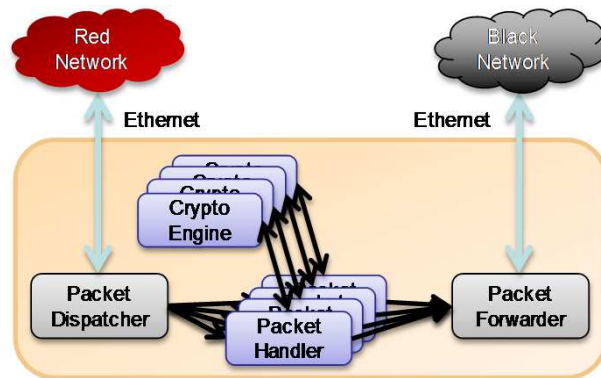
Signal Processing Example

- Radar applications: STAP, ...
- Beam forming : Sonar, Echography
- Software Defined Radio (SDR)
- Dedicated libraries (FFT, FTFR, ...)



Well suited to massively parallel architectures
Alternative of embedded DSP + FPGA platforms

High-Speed VPN Gateway Example

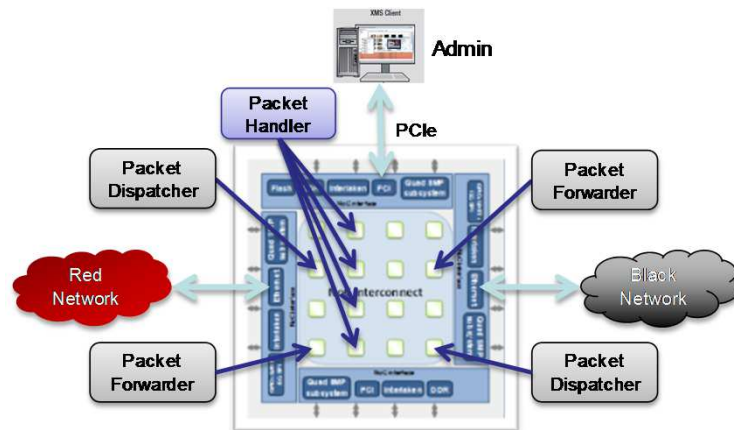


Evaluation for the implementation of a 20 to 40 Gbs VPN gateway

- IP packet processing
- AES cryptography

Exploit key features of the MPPA architecture

- 2 x 40 Gbs Ethernet interfaces (or 8 x 10 Gbs)
- PCIe Gen 3 for integration
- Optimized instructions for efficient cryptography
- NoC extension interface for multi-chip solutions



KALRAY, a global solution



Powerful, Low Power and Programmable Processors




C/C++ based Software Development Kit (SDK) for massively parallel programming



Development platform Reference Design Board



Reference Design board Application specific boards Multi-MPPA or Single-MPPA boards



Kalray Offices

Headquarters – Paris area

86 rue de Paris,
91 400 Orsay
France

Tel: +33 (0)1 69 29 08 16
email: info@kalray.eu



Grenoble office

445 rue Lavoisier,
38 330 Montbonnot Saint Martin
France

Tel: +33 (0)4 76 18 09 18
email: info@kalray.eu



Japan office

CVML, 3-22-1, Toranomom,
Minato-ku, Tokyo 105-0001,
Japan

Tel: 080-4660-2122
email: ksugiyama@kalray.eu

All trademarks, service marks, and trade names are the marks of the respective owner(s), and any unauthorized use thereof is strictly prohibited. All terms and prices are indicatives and subject to any modification without notice.