## Energy Efficient Seismic Wave Propagation Simulation on a Low-power Manycore Processor

Márcio Castro, Fabrice Dupros, Emilio Francesquini, Jean-François Méhaut, Philippe O. A. Navaux

email: marcio.castro@ufsc.br
web: www.marciocastro.com



HOSCAR Workshop'14: High Performance Computing and Scientific Data Management Driven by Highly Demanding Applications

# MOTIVATION

- Simulations of large scale seismic wave propagation
  - Risk mitigation
  - Study of future hypothetical earthquakes
  - Oil and gas exploration
- Realistic simulations  $\rightarrow$  complex models
  - Intensive computations on large amounts
     of data
  - Need of HPC platforms to achieve reliable results in feasible time

- Until the last decade
  - Performance of HPC platforms has been quantified by their processing power (Flops)
- Nowadays
  - Energy efficiency (Flops/Watt) is as important as processing power
  - Critical aspect to the development of scalable systems

- Example: data-centers
  - Power and cooling costs largely dominate the operational costs
    - **30%** of the energy is used for **cooling**
    - 10-15% is lost in power conversions and distribution losses
- Defense Advanced Research Projects Agency, EUA (DARPA) report
  - Acceptable energy efficiency for Exascale systems → 50 GFlops/Watt
  - TSUBAME-KFC: number one platform in Green500 performs 4.5 GFlops/Watt

- New alternatives for low-power HPC
  - Light-weight manycore processors
  - Examples: Tilera TILE-Gx, Kalray MPPA-256

- Light-weight manycores vs. GPUs
  - Autonomous cores
  - Cores can be used to accomplish both data and task parallelism
  - Low power consumption: few tens of watts

- Developing efficient scientific parallel applications for light-weight manycores is challenging
  - Built and optimized for specific classes of embedded applications
  - Memory constraints
  - Absence of coherent caches
  - Peculiar Networks-on-Chip (NoCs)

### Our work

 Adapt a main kernel of a seismic wave propagation simulator (Ondes3D - BRGM) to a recent light-weight manycore

 Kalray MPPA-256

 Compare the performance and energy efficiency of our solution with optimized solutions for GPUs and general-purpose multicores

### Outline

- Seismic wave propagation kernel
- MPPA-256 overview
- Proposed solution
- Results
- Conclusions

## **SEISMIC WAVE PROPAGATION**

### Seismic wave propagation

In our case, the earthquake process is described as elastodynamics

• Finite-differences scheme is used for solving the wave propagation problem

### Seismic wave propagation

- Simulation composed by time steps
- In each time step (3D simulation)
  - The first triple nested loop computes the velocity components
  - The second loop reuses the velocity results of the previous time step to update the stress field





### Seismic wave propagation

- Current parallel implementations
  - Multicores: OpenMP to parallelize the triple nested loops (3D domain)
  - GPUs: sliding-window algorithm that relies on a two-dimensional tiled decomposition of the 3D domain

## **MPPA-256 OVERVIEW**

• Kalray



- French semiconductor and software company (Grenoble and Paris) developing and selling a new generation of manycore processors for HPC
- MPPA-256



- Multi-Purpose Processor Array (MPPA)
- Manycore processor: 256 cores in a single chip
- Low power consumption (5W 11W)



- 256 cores (PEs) @ 400 MHz: 16 clusters, 16 PEs per cluster
- PEs share 2 MB of memory
- Absence of cache coherence protocol inside the cluster
- Network-on-Chip (NoC): communication between clusters
- 4 I/O subsystems: 2 connected to external memory





 A master process runs on an RM of one of the I/O subsystems



- The master process spawns slave processes
- One slave process per cluster



**MPPA-256** 

- The slave process runs on the PEO and may create up to 15 threads, one for each PE
  - POSIX or OpenMP
- Threads share 2 MB of memory



- Communications take the form of remote writes
- Data travel through the **NoC**

## **PROPOSED SOLUTION**

### **Proposed solution**

#### Challenge 1: memory

- Real simulation data don't fit in 2 MB per cluster
- Data transfers from/to the DDR explicitly managed by the programmer

#### Challenge 2: data transfers

- Specific API to perform data movements
- Assynchronous data transfers to overlap communications with computations

#### • Challenge 3: NoC

- Data transfers should match the NoC topology to reduce communication costs
- Send few data transfers containing large amounts of data is better than several data transfers containing few data

### **Proposed solution**

### Two-level tiling scheme to exploit the memory hierarchy of MPPA-256



## RESULTS

### Results

#### • Xeon E5

- Sandy Bridge-EP with 8 cores at 2.4 GHz
- 32 GB of DDR3
- SGI Altix UV 2000
  - ccNUMA (24 NUMA nodes)
  - NUMA node: Xeon E5 (8 cores at 2.4 GHz)
  - 192 cores and 768 GB of main memory in total
- NVIDIA Quadro K5000 GPU
  - Kepler architecture
  - 768 CUDA cores at 800 MHz
  - 3 GB of main memory (GDDR5)







#### MPPA-256 vs. SGI Altix UV 2000

Platform	Time-to-Solution	<b>Energy-to-Solution</b>
MPPA-256 Altix UV 2000	100.2 s <b>3</b> 2.9 s <b>3</b>	<b>4.5x</b> 752 J <b>5.8x</b> 4418 J

#### Benefits of the prefetching scheme



## CONCLUSIONS

### Conclusions

- Light-weight manycores
  - Opportunity to perform highly-parallel energy-efficient computations
- Seismic wave propagation on MPPA-256
  - Several architecture peculiarities
  - Multi-level tiling scheme to deal with the cluster's limited memory size
  - Explicit software prefecting mechanism
     to overlap communications/computations

### Conclusions

- Multi-MPPA co-processor
  - Kalray recently announced a multi-MPPA solution that features four MPPA-256 processors on the same board with less than 50 W of power consumption

#### Future works

- Adapt our multi-level tiling and prefetching scheme to exploit multi-MPPA solutions
- Deal with input problem sizes of 32 GB or more

### Obrigado!

### Latest publications on manycores/MPPA-256

- Márcio Castro, Emilio Francesquini, Thomas M. Nguélé, and Jean-François Méhaut. Multicoeurs et Manycoeurs: Une Analyse de la Performance et l'Éfficacité Énergétique d'une Application Irrégulière. In: Conférence de recherche en informatique (CRI). Yaoundé, Cameron, 2013.
- Márcio Castro, Emilio Francesquini, Thomas M. Nguélé, and Jean-François Méhaut. Analysis of Computing and Energy Performance of Multicore, NUMA, and Manycore Platforms for an Irregular Application. In: Workshop on Irregular Applications: Architectures & Algorithms (IA3) - Supercomputing Conference (SC). Denver, USA: ACM, 2013.
- Márcio Castro, Fabrice Dupros, Emilio Francesquini, Jean-François Méhaut, and Philippe O. A. Navaux. Energy Efficient Seismic Wave Propagation Simulation on a Low-power Manycore Processor. In: International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD). Paris, France: IEEE Computer Society, 2014, (accepted).
- Emilio Francesquini, Márcio Castro, Pedro H. Penna, Fabrice Dupros, Henrique C. de Freitas, Philippe O. A. Navaux, and Jean-François Méhaut. On the Energy Efficiency and Performance of Irregular Application Executions on Multicore, NUMA and Manycore Platforms. Journal of Parallel and Distributed Computing (JPDC) - 2nd round review.