Fourth Brazil-France Workshop on High Performance Computing and Scientific Data Management Driven by Highly Demanding Applications

Load Balancing Analysis for Seismic Model -Ondes3Don multi GPU platforms

Víctor Martínez Philippe Navaux Fabrice Dupros - BRGM

Parallel and Distributed Processing Group (GPPD) Federal University of Rio Grande do Sul (UFRGS)





Schedule

- * Ondes 3D
- Motivation
- * Experiments
- * Measures
- * Results
- * TODO





Ondes 3D - The Model

$$\rho \ddot{\mathbf{u}} = \nabla \cdot \boldsymbol{\sigma} + \mathbf{f} ,$$

$$\boldsymbol{\sigma} = \mathbf{c} : \boldsymbol{\varepsilon} ,$$

$$\boldsymbol{\varepsilon} = \frac{1}{2} [\nabla \mathbf{u} + (\nabla \mathbf{u})^T],$$

$$\rho \partial_t \mathbf{v} = \nabla \cdot \boldsymbol{\sigma} ,$$

$$\partial_t \boldsymbol{\sigma} = \mathbf{c} : \nabla \mathbf{v} .$$

$$\partial_x u\left(i + \frac{1}{2}, j, k\right) \simeq \frac{9}{8} \frac{u(i+1, j, k) - u(i, j, k)}{\Delta x} - \frac{1}{24} \cdot \frac{u(i+2, j, k) - u(i-1, j, k)}{\Delta x}$$



Stress - Velocity





Ondes 3D - The Model

$$\rho \ddot{\mathbf{u}} = \nabla \cdot \boldsymbol{\sigma} + \mathbf{f} ,$$

$$\boldsymbol{\sigma} = \mathbf{c} : \boldsymbol{\varepsilon} ,$$

$$\boldsymbol{\varepsilon} = \frac{1}{2} [\nabla \mathbf{u} + (\nabla \mathbf{u})^T],$$

$$\rho \partial_t \mathbf{v} = \nabla \cdot \boldsymbol{\sigma} ,$$

$$\partial_t \boldsymbol{\sigma} = \mathbf{c} : \nabla \mathbf{v} .$$

$$\partial_x u\left(i + \frac{1}{2}, j, k\right) \simeq \frac{9}{8} \frac{u(i+1, j, k) - u(i, j, k)}{\Delta x} - \frac{1}{24} \cdot \frac{u(i+2, j, k) - u(i-1, j, k)}{\Delta x}.$$



Stress - Velocity





Ondes 3D - The Model

$$\rho \ddot{\mathbf{u}} = \nabla \cdot \boldsymbol{\sigma} + \mathbf{f} ,$$

$$\boldsymbol{\sigma} = \mathbf{c} : \boldsymbol{\varepsilon} ,$$

$$\boldsymbol{\varepsilon} = \frac{1}{2} [\nabla \mathbf{u} + (\nabla \mathbf{u})^T],$$

$$\rho \partial_t \mathbf{v} = \nabla \cdot \boldsymbol{\sigma} ,$$

$$\partial_t \boldsymbol{\sigma} = \mathbf{c} : \nabla \mathbf{v} .$$

$$\partial_x u\left(i + \frac{1}{2}, j, k\right) \simeq \frac{9}{8} \frac{u(i+1, j, k) - u(i, j, k)}{\Delta x} - \frac{1}{24} \cdot \frac{u(i+2, j, k) - u(i-1, j, k)}{\Delta x}.$$



Stress - Velocity





Motivation







Motivation



Is *performance loss* caused by *poor load balancing*? Is *the unbalanced problem* correlated with *communication time* and *memory usage*?





Experiments

- * Guane-1 (Colombia)
 - * 16 nodes
 - * 84 Intel Xeon E5640.
 - * 128 GPUs TESLA FERMI M2050.
- Increasing the number of GPUs, until to get unbalanced load.
- * Executed with MPI Charm++









Timeloop







Speedup



























































Results (Correlation)









Results (Correlation)

- * Memory:
 - GPU load Memory utilization: 0.88







Results (Correlation)

- * Memory:
 - * GPU load Memory utilization: 0.88
- * Communication:
 - * GPU load Communication time: **-0.30**.
 - STD load Communication
 time: -0.71





Results

- * Unbalanced load:
 - * GPUs with lower load are using less memory than the others with higher load.
 - * GPUs have high rates of communication.
 - * GPUs are solving the inner space on the grid.
 - * GPUs are solving the lowest values on Y axis.
- * Charm++
 - * Same behavior: We have only information of CPUs load.





TODO

- * We need to increment memory utilization and to reduce communication: it is necessary to send fragments of the 3D grid instead to a 2D tile, in the sense that it will have more data and less communication to find the nearest values to calculate the stencil (12 points).
- * To exploit the advantages of process virtualization on load balancing will be necessary. This optimization could be done with concurrent execution (Fermi) and dynamic parallelism (Kepler).
- * StarPU implementation for better scheduling.







Merci! Obrigado! Thanks! ¡Gracias!



Fourth Brazil-France Workshop on High Performance Computing and Scientific Data Management Driven by Highly Demanding Applications

Load Balancing Analysis for Seismic Model Ondes3D on multi GPUs platforms

Víctor Martínez Philippe Navaux Fabrice Dupros

Parallel and Distributed Processing Group (GPPD) Federal University of Rio Grande do Sul (UFRGS)





