

Simulation of seismic wave propagation in an asteroid based upon an unstructured MPI spectral-element method: blocking and non-blocking communication strategies

Roland Martin, Dimitri Komatitsch, Céline Blitz, Nicolas Le Goff

▶ To cite this version:

Roland Martin, Dimitri Komatitsch, Céline Blitz, Nicolas Le Goff. Simulation of seismic wave propagation in an asteroid based upon an unstructured MPI spectral-element method: blocking and nonblocking communication strategies. Lecture Notes in Computer Science, 2008. inria-00339890

HAL Id: inria-00339890 https://inria.hal.science/inria-00339890

Submitted on 19 Nov 2008 $\,$

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Simulation of seismic wave propagation in an asteroid based upon an unstructured MPI spectral-element method: blocking and non-blocking communication strategies

Roland Martin¹, Dimitri Komatitsch^{1,2}, Céline Blitz¹, and Nicolas Le Goff¹

¹ Université de Pau et des Pays de l'Adour, CNRS & INRIA Sud-Ouest Magique-3D, Laboratoire de Modélisation et d'Imagerie en Géosciences UMR 5212, Avenue de l'Université, 64013 Pau Cedex, France {roland.martin,dimitri.komatitsch,celine.blitz, nicolas.legoff}@univ-pau.fr http://www.univ-pau.fr
² Institut universitaire de France, 103 boulevard Saint-Michel, 75005 Paris, France {dimitri.komatitsch}@univ-pau.fr http://www.cpu.fr/Iuf/

Abstract. In order to better understand the internal structure of asteroids orbiting in the Solar system and then the response of such objects to impacts, seismic wave propagation in asteroid 433-Eros is performed numerically based on a spectral-element method at frequencies lying between 2 Hz and 22 Hz. In the year 2000, the NEAR Shoemaker mission to Eros has provided images of the asteroid surface, which contains numerous fractures that likely extend to its interior. Our goal is to be able to propagate seismic waves resulting from an impact in such models. For that purpose we create and mesh both homogeneous and fractured models with a highly-dispersive regolith laver at the surface using the CUBIT mesh generator developed at Sandia National Laboratories (USA). The unstructured meshes are partitioned using the METIS software package in order to minimize edge cuts and therefore optimize load balancing in our parallel blocking or non-blocking MPI implementations. We show the results of several simulations and illustrate the fact that they exhibit good scaling.

Key words: Non-blocking MPI, load balancing, scaling, mesh partitioning, seismic wave propagation, asteroids.

1 Introduction

In the context of seismic exploration, it is of crucial importance to develop efficient numerical tools to model seismic wave propagation in complex structures with great accuracy at scales of at least tens of kilometers. For this purpose, the seismic wave equation can be solved numerically in heterogeneous media using $\mathbf{2}$

different methods, for instance the finite-difference technique or the boundary integral technique. In the last decade, the spectral-element method (SEM), which is more accurate and flexible, has been used extensively in the context of regional or global seismology. The SEM was introduced twenty years ago in computational fluid mechanics by [1]. It is a high-order variational method that retains the ability of finite elements to handle complicated geometries while keeping the exponential convergence rate of spectral methods. Complex topographies, dipping or curved interfaces, interface and surface waves, and distorted meshes can be easily taken into account. Indeed, the SEM provides a more natural context to describe the free surface of the model thanks to the weak formulation of the equations that is used, for which the free surface condition is a natural condition. The formulation of the SEM based on the displacement vector and on Gauss-Lobatto-Legendre numerical integration which is implemented in our software package called SPECFEM has the property to retain a diagonal mass matrix and is therefore easier to implement than classical low-order finiteelement methods. Applications of the SEM to two-dimensional (2D) (e.g., [3,4]) and three-dimensional (3D) elastodynamics (e.g., [5,6]) have shown that high accuracy (i.e., small numerical dispersion) is obtained. The time discretization is an explicit conditionally-stable second-order centered finite-difference scheme. Because of the diagonal mass matrix and of the standard explicit time scheme, no inversion of a linear system is needed and therefore the method can be efficiently implemented in parallel and large 3D models can be handled (e.g., [5–7]). In this article, we use the SEM to simulate seismic wave propagation resulting from an impact at the surface of a 2D cut plane in an asteroid.

Asteroids are metallic or rocky bodies orbiting in the Solar System. Very little is known about their internal structure and therefore several models of their interior have been proposed, for instance monoliths (objects of low porosity, good transmitters of elastic stress) or rubble pile (shattered bodies whose pieces are grouped into a loose and porous packing [8]). Depending on its structure, the response of an asteroid to an impact (for instance when it is hit by a meteor or by an artificial impactor sent on it at high velocity) can be very different [9]: a rubble pile would be harder to disrupt than a monolith [8]. Thus, to develop mitigation techniques (to prevent a potential collision of a hazardous object with the Earth) it is important to be able to perform simulations of body disruption on reliable models of an asteroid interior [10]. The geophysical knowledge of asteroid interiors can be improved by space missions sent to comets or asteroids. In the year 2000, the NEAR Shoemaker mission to 433-Eros has provided images of the asteroid surface, which displayed numerous fractures and evidence of a coherent but fractured interior [11]. The study of the crater distribution at the surface of Eros has highlighted a deficit in the distribution of small crater sizes. To explain this observation, [12] have proposed impact-induced vibrations as a possible source of downslope movements on crater walls. The related mobilized material could fill the smallest craters, leading to their erasure. In this study we therefore simulate wave propagation in different models of the interior of asteroid 433-Eros.

2 Spatial and temporal discretization of the governing equations

We consider a linear isotropic elastic rheology for the heterogeneous solid, and therefore the seismic wave equation can be written in the strong form as:

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

$$\boldsymbol{\sigma} = \mathbf{C} : \boldsymbol{\varepsilon} = \lambda \operatorname{tr}(\boldsymbol{\varepsilon}) \mathbf{I} + 2\mu \boldsymbol{\varepsilon} ,$$

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

(1)

where **u** denotes the displacement vector, $\boldsymbol{\sigma}$ the symmetric, second-order stress tensor, $\boldsymbol{\varepsilon}$ the symmetric, second-order strain tensor, **C** the fourth-order stiffness tensor, $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$ the two Lamé parameters, $\boldsymbol{\rho}$ the density, and **f** an external force. The trace of the strain tensor is denoted by $\operatorname{tr}(\boldsymbol{\varepsilon})$, **I** denotes the identity tensor, the tensor product is denoted by a colon, and a superscript T denotes the transpose. A dot over a symbol indicates time differentiation. The physical domain is denoted by Ω , and its outer boundary by Γ . The material parameters of the solid, **C** (or equivalently $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$) and $\boldsymbol{\rho}$, can be spatially heterogeneous. We can then rewrite the system (1) in a variational weak form by dotting it with an arbitrary test function **w** and integrating by parts over the whole domain as:

$$\int_{\Omega} \rho \mathbf{w} \cdot \ddot{\mathbf{u}} \, \mathrm{d}\Omega + \int_{\Omega} \nabla \mathbf{w} : \mathbf{C} : \nabla \mathbf{u} \, \mathrm{d}\Omega = \int_{\Omega} \mathbf{w} \cdot f \, \mathrm{d}\Omega + \int_{\Gamma} (\boldsymbol{\sigma} \cdot \hat{\mathbf{n}}) \cdot \mathbf{w} \, \mathrm{d}\Gamma \,.$$
(2)

The free surface (i.e. traction free) boundary condition on Γ is easily implemented in the weak formulation since the integral of traction along the boundary simply vanishes (e.g., [5]) when we set $\boldsymbol{\tau} = \boldsymbol{\sigma} \cdot \hat{\mathbf{n}} = 0$ at the free surface.

This formulation is solved on a mesh of quadrangular elements in 2D, which honors both the free surface of the asteroid and its main internal discontinuities (for instance its fractures). The unknown wave field is expressed in terms of high-degree Lagrange polynomials on Gauss-Lobatto-Legendre interpolation points, which results in an exactly diagonal mass matrix that leads to a simple time integration scheme (e.g. Komatitsch et al., 2005). Let \mathbf{w}_N , \mathbf{u}_N denote the piecewise-polynomial approximations of the test functions and the displacement, respectively. Making use of (2), the discrete variational problem to be solved can thus be expressed as: for all time t, find \mathbf{u}_N such that for all \mathbf{w}_N we have

$$\int_{\Omega} \rho \mathbf{w}_N \cdot \ddot{\mathbf{u}}_N \,\mathrm{d}\Omega + \int_{\Omega} \nabla \mathbf{w}_N : \mathbf{C} : \nabla \mathbf{u}_N \,\mathrm{d}\Omega = \int_{\Omega} \mathbf{w}_N \cdot f \,\mathrm{d}\Omega \,\,. \tag{3}$$

We can rewrite this system (3) in matrix form as:

$$M\hat{d} + Kd = F , \qquad (4)$$

where M is the diagonal mass matrix, F is the source term, and K is the stiffness matrix. For detailed expression of these matrices, the reader is referred for instance to [6].

Time discretization of the second-order ordinary differential equation (4) is achieved using the explicit Newmark central finite-difference scheme [5, 6], which is second order accurate and conditionally stable :

$$M\mathbf{d}_{n+1} + K\mathbf{d}_{n+1} = F_{n+1} , \qquad (5)$$

where

4

$$\mathbf{d}_{n+1} = \mathbf{d}_n + \Delta t \dot{\mathbf{d}}_n + \frac{\Delta t^2}{2} \ddot{\mathbf{d}}_n , \qquad (6)$$

and

$$\dot{\mathbf{d}}_{n+1} = \dot{\mathbf{d}}_n + \frac{\Delta t}{2} [\ddot{\mathbf{d}}_n + \ddot{\mathbf{d}}_{n+1}] .$$
(7)

At the initial time t = 0, null initial conditions are assumed i.e., $\mathbf{d} = \mathbf{0}$ and $\mathbf{d} = \mathbf{0}$. The time step Δt and the distribution of mesh sizes h(x, y) are chosen such that the Courant-Friedrichs-Lewy (CFL) stability condition and a numerical dispersion condition are satisfied. The CFL condition is :

$$max\left(\frac{c_p(x,y)}{h(x,y)}\right)\Delta t \le \alpha , \qquad (8)$$

where $c_p(x, y)$ is the pressure velocity distribution in the model under study, and the numerical dispersion condition is

$$\min\left(\frac{c_s(x,y)}{h(x,y)}\right)/(2.5f_0) \ge n_\lambda , \qquad (9)$$

where f_0 is the dominant frequency of the seismic source, $c_s(x, y)$ the shear velocity distribution in the model and n_{λ} the minimum number of grid points per seismic wavelength. In practice α is taken lower than 0.5, and n_{λ} about 5 [13].

3 Model of the asteroid

The size of Eros is approximately 34 km (length) by 17 km (height) and we have designed its model according to the dataset provided by the NEAR Shoemaker spacecraft. These data display a regolith blanket at the surface (formed by crushed rocks during impacts) of thickness ranging from tens to hundreds of meters [11]. We therefore added a regolith blanket around the asteroid, with a thickness ranging from 50 to 150 m. The images of the surface of Eros also display numerous craters and fractures. These fractures are thought to be formed by impacts and the regolith depression found around them suggests regolith infiltration [11]. This regolith depression can be up to 300 m wide, for instance near the Rahe Dorsum ridge. This long fracture, striking on Eros images, probably crosses the asteroid interior and is therefore included in our 2D model. Except Rahe Dorsum, all fractures have been designed under the main craters in our 2D model and filled with the same material as the regolith. Because of technical constraints, we designed simple fracture shapes, and to avoid low angles in the mesh elements, we tried to draw the fracture geometry with angles close to 90 degrees. We extended the fracture network to a depth of one crater radius or more according to [16].

We define two models of the interior of Eros: one that is homogeneous (it includes the topography of Eros and an elastic material characterized by a pressure wave velocity $c_p = 3000 \text{ m.s}^{-1}$, a shear wave velocity $c_s = 1700 \text{ m.s}^{-1}$, and a density $= 2700 \text{ kg.m}^{-3}$), and another one that in addition comprises a regolith layer as well as fracture networks. The interior of the fractures and the regolith layer have the same material properties: $c_p = 900 \text{ m.s}^{-1}$, $c_s = 500 \text{ m.s}^{-1}$, and a density of 2000 kg.m⁻³. Seismic attenuation (i.e., loss of energy by viscoelasticity) can be ignored in this problem because seismic studies performed on the Moon have shown that it is negligible.

The two models are meshed with CUBIT quadrangular and hexahedral mesh generator developed at Sandia National Laboratories (http://cubit.sandia.gov, USA). Table 1 shows that the quality of the meshes obtained for the two models displayed in Figure 1 is good: the angles of the elements are acceptable (comprised between approximately 30 and 90 degrees) with a related skewness lying between 0 and 0.65 (0 corresponds to the best case of a perfectly cubic element and 1 corresponds to the worst case of a flat element). For speed-up scaling purposes we perform different calculations by increasing the number of mesh points and elements while we increase the dominant frequency of the seismic source according to the dispersion relation. Subsequently the number of partitions increases accordingly, each mesh partition being mapped to one processor core. Dominant frequencies are chosen from 2 Hz to 22 Hz and element sizes take values of 18 m (in fractures and regolith layers) to 522 m (in the bedrock) at 2 Hz, and 1.68 m (fractures and regolith) to 47.5 m (bedrock) at 22 Hz. Chosing higher frequencies is not relevant because the available data resolution (300 m) as well as the interior structure are not known accurately enough from a physical point of view.

4 Mesh partitioning and non-blocking MPI implementation

A 2D mesh designed for a very high-resolution simulation is too large to fit on a single computer. We therefore implement the calculation in parallel based upon MPI. We first partition the mesh using the METIS graph partitioning library [14], which focuses on balancing the size of the different domains and minimizing the edge cut. Figure 2 shows how the mesh for the homogeneous model is partitioned by METIS into 8 or into 80 domains. Balancing the size of the domains ensures that no processor core will be idle for a significant amount of time while others are still running at each iteration of the time loop, while a small edge cut reduces the number and the size of the communications. Contributions between neighboring elements that are located on different processor cores are added using non-blocking MPI sends and receives and following a similar

implementation for low-order finite element method described in [15]. The communication scheme used is the following: contributions (i.e., mechanical internal forces) from the outer elements of a given mesh slice (i.e., elements that have an edge in common with an element located on a different processor core) are computed first and sent to the neighbors of that mesh slice using a non-blocking MPI send. Similarly, each processor core issues non-blocking MPI receives. This allows for classical overlapping of the communications by the calculations, i.e. each process then has time to compute mechanical forces in its inner elements while the communications travel across the network, and if the number of outer elements is small compared to the number of inner elements in all the mesh slices we can be confident that the messages will be arrived when the internal calculations are finished. Once the contributions from neighbors have been received, they are added to the corresponding elements locally. We ran the code on a Dell PowerEdge 1950 cluster with Intel EM64T Xeon 5345 (Clovertown) processors (2333 MHz clock frequency, 9.330 Gigaflops peak performance) and Myrinet network, located at the California Institute of Technology (USA). Each processor is dual core, so we prefer to speak here in terms of processes or processor cores with one process per processor core.

5 Numerical simulations and scaling of the code

In all the simulations we use a polynomial of degree N = 4 to integrate variables at the $(N + 1)^2 = 25$ Gauss-Lobatto-Legendre points along each direction of any spectral element. A series of simulations to measure speedup and scaling are performed at a dominant frequency of 2 Hz for the homogeneous model first and then for the fractured model. A third series of tests are computed to study the weak scaling of the code. By weak scaling we mean how the time to solve a problem with increasing size can be held constant by enlarging the number of processes used and by maintaining a fixed system size per process: when one doubles the number of processes one also doubles the system size.

Moderate mesh sizes are considered for the homogeneous model, and moderate to large mesh sizes for the fractured model (Table 1). The seismic source is an impact represented by a force normal to the surface and located at point (x = -9023 m, y = 6131 m) in the mesh. The time variation of the source is the second derivative of a Gaussian. The time step used is $\Delta t = 1$ ms and the signal is propagated for 150000 time steps (i. e. 150 s). In the case of the homogeneous model, Figure 3 shows pressure, shear and surface waves propagating inside Eros and along its surface for 65 seconds. One can see in particular the pressure wave front that is reflected on the lower free surface and the shear wave front that closely follows the pressure wave front, as well as many reflections and conversions of waves generated along the whole surface of the object. On the contrary, in the fractured model (figure 4), waves travel mostly inside the dispersive surface layer of regolith as well as the fractures, which act as wave guides. The waves then reflect off all the boundaries and are trapped for a while in the left part of the model, where the source is located, and then progressively move to the right part, the large transversal fracture acting as a geological barrier.

Figure 5 shows the scaling of the code for 1 to 80 processes in the case of the mesh for the homogeneous model with 3744 elements with sizes from 147 m to 470 m. As expected, scaling is very good (close to the straight line obtained for hypothetical perfect scaling) as long as the number of outer elements is small compared to the number of inner elements in all the mesh slices, in which case overlapping of communications with calculations using non-blocking MPI works very well. Here for the relatively small mesh used this is true when we use a number of processes lower than 32. When we use a number of processes that is larger and for which this assumption ceases to be true, scaling quickly becomes very poor because communications are no longer overlapped. Indeed, for more than 32 processes the total measured time goes through a minimum and then starts to increase. The three other curves (medium and tiny dashed lines and filled square lines), which represent the lowest, average and highest values of the sum of the time spent in calculations and communications in the outer elements of all the partitions, increase when we use more than 32 processes, which explains the poor scaling observed in this case. It is observed that total time spent in communications and calculations is not really different between blocking and non-blocking communication procedures for this specific application.

In the fractured case, scaling is performed for 57275 spectral elements (919709 grid points) at 2 Hz by increasing the number of processes from 1 to 121. Figure 6 shows that the scaling is much better than for the homogeneous case. The poor scaling that appears beyond 32 processes in the homogeneous case will appear for a much higher number of processes in this heterogeneous case. This is not surprising because the number of elements is higher inside each partition.

Now, instead of increasing the number of processes for a given fixed mesh, a weak scaling study is performed for the complex and highly unstructured meshes generated for the fractured model. The dominant frequency of the source is increased from 2 Hz to 22 Hz while we refine the mesh in the same ratio and increase the number of processor cores consequently from 1 to 121. By construction of the cracks in the asteroid and using a mesh decimation technique in each element, skewnesses and angles are reasonably preserved when the number of elements increases. Time steps and mesh sizes are respectively decreased from 0.45 ms (2 Hz) to 0.04 ms (22 Hz). At 2 Hz, mesh sizes take values varying from 18 m (regolith) to 522 m (bedrock), while at 22 Hz they vary from 1.68 m to 47 m. In Figure 7 we observe that the total computational time per process is nearly constant as the number of processes increases beyond more than 10 processes approximately. This is observed for both blocking or non-blocking communication strategies, which give similar results. The tests correspond to 57275 elements (nearly 1 million points at 2 Hz) for 1 processor core and to nearly 6 million elements (nearly 111 million points at 22 Hz) for 121 processor cores. The number of degrees of freedom (the two components of the displacement vector) is exactly twice the number of points, i.e., close to 222 millions. Let us finally mention that a mixed MPI/OpenMP model could also be used but Komatitsch et al. [7] have shown that it does not bring any significant gain in performance for this particular application.

6 Conclusions

We have simulated wave propagation in a homogeneous or a fractured model of an asteroid represented by a non-structured mesh. A mesh with good skewness has been developed with CUBIT. For both blocking and non-blocking communication strategies using METIS, similar scalings are obtained and mesh configurations of 110 million points can be computed using 121 processor cores and dominant seismic frequencies of up to 22 Hz. In future work it would be interesting to extend the simulations to 3D at high resolution and to apply L2 cache misses reduction techniques [17].

Acknowledgments. The authors would like to thank Philippe Lognonné for fruitful discussion about asteroids, Emanuele Casarotti and Steven J. Owen for fruitful discussion about meshing with CUBIT, Jean Roman and Jean-Paul Ampuero for fruitful discussion about overlapping communications with nonblocking MPI. Calculations were performed on the Division of Geological & Planetary Sciences Dell cluster at the California Institute of Technology (USA). This material is based in part upon research supported by European FP6 Marie Curie International Reintegration Grant MIRG-CT-2005-017461 and by the French ANR under grant NUMASIS ANR-05-CIGC-002.

Model	Homogeneous (3744 elements)	Fractured (57275 elements)
Average angle	84.14	81.98
Standard deviation angle	5.57	7.44
Min Angle	59.83	29.76
Max angle	89.94	89.97
Average skew	0.0606	0.0861
Standard deviation skew	0.0575	0.0925
Min skew	0	0
Max skew	0.356	0.647

Table 1. The quality of the quadrangle elements can be defined by their angles θ (rad) or equivalently by their skewness, which is defined as $|(2\theta - \pi)/\pi|$. Ideal angles are right angles (90 degrees) with skewness of 0, and poor angles are lower than typically 30 degrees with skewness beyond approximately 0.65. Here the number of poor elements is small and their influence on the global calculations is therefore expected to remain reasonable.

References

1. Patera A. T.: A spectral element method for fluid dynamics: laminar flow in a channel expansion. Journal of Computational Physics (1984), 54, pp. 468-488.



Fig. 1. Meshes created using the CUBIT mesh generator for an homogeneous model of asteroid Eros (top, 3744 elements) and a more complex model with a regolith layer and a network of fractures (bottom, 52275 elements) for simulations performed at a central frequency of 2 Hz. Close-ups on the white frame are also shown.



Fig. 2. Partitioning of the mesh for the homogeneous model (Figure 1) obtained with METIS in the case of 8 (left) and 80 (right) domains. We observe that the number of elements along the interface of the partitions is small compared to the number of elements inside each partition in the case of 8 domains, and therefore overlapping of communications with calculations is expected to work fine. But in the case of 80 domains, this number becomes comparable or even higher than the number of inner elements, in which case overlapping is expected to fail, and poor performance should result.



Fig. 3. Snapshots of the propagation of simulated seismic waves in the asteroid for a total duration of 65 seconds in the case of the mesh for a homogeneous model displayed in Figure 1. Snapshots are shown at 10 s, 25 s (top) and 45 s, 65 s (bottom). We represent the vertical component of the displacement vector in red (positive) or blue (negative) at each grid point when it has an amplitude higher than a threshold of 1% of the maximum, and the normalized value is raised to the power 0.30 in order to enhance small amplitudes that would otherwise not be clearly visible.



Fig. 4. Snapshots of the propagation of simulated seismic waves in the asteroid for a total duration of 65 seconds in the case of the mesh for the more complex model with a regolith layer and networks of fractures displayed in Figure 1. Snapshots are shown at 10 s, 25 s (top) and 45 s, 65 s (bottom). We represent the vertical component of the displacement vector in red (positive) or blue (negative) at each grid point when it has an amplitude higher than a threshold of 1% of the maximum, and the normalized value is raised to the power 0.30 in order to enhance small amplitudes that would otherwise not be clearly visible.



sented in Figure 1 for blocking (top) and non blocking communications (bottom). For both strategies the measured total time spent in the calculations and communications (long dashed line with single crosses) is compared to the theoretical straight line obtained assuming perfect scaling (solid line) for different numbers of processes (one per mesh partition): 1, 2, 4, 8, 32, 64 and 80. The two curves are in very good agreement as long as the number of outer elements is small compared to the number of inner elements in all the mesh slices, in which case overlapping of communications with calculations using non-blocking MPI works very well. Here for the relatively small mesh used this is true when we use a number of processes lower than 32. When we use a larger number of processes scaling quickly becomes very poor because communications are no longer overlapped. One can also observe that the total time spent in communications and

calculations is similar between blocking and non-blocking communication strategies.



Scaling in the fractured case (non blocking MPI)

Fig. 6. Scaling of the code in the case of the mesh for the fractured model represented in Figure 1 for blocking (top) and non blocking communication (bottom) strategies. For both strategies the measured total time spent in the calculations and communications (long dashed line with single crosses) is compared to the theoretical straight line obtained assuming perfect scaling (solid line) for different numbers of processor cores (one per mesh partition): 1, 2, 4, 8, 32, 64, 80. The two curves are in very good agreement as long as the number of outer elements is small compared to the number of inner elements in all the mesh slices, in which case overlapping of communications with calculations using non-blocking MPI works very well. Here for the relatively small mesh used this is always the case. We do not observe the same poor scaling as in the homogeneous case but it should start to appear if we used more than 80 processes. Total time spent in communications and calculations is not really different between blocking and non-blocking communication procedures for this specific application, as observed in the homogeneous case.

Simulation of seismic wave propagation in an asteroid using MPI 13



Weak scaling in the fractured case (non blocking MPI)

Fig. 7. Weak scaling of the code for the fractured model represented in Figure 1 in the case of non blocking (top) and blocking (bottom) communications. The measured total time spent in the calculations and communications (long dashed line with single crosses) is compared to the theoretical straight line obtained assuming perfect scaling (solid line) for different numbers of processor cores (one per mesh partition) and different dominant frequencies : 1 (2 Hz and 57275 elements), 4, 9, 16, 25, 36, 49, 64, 81, 121 (22 Hz and around 110 million points) processes. As we increase the number of processes and the corresponding seismic frequency, the number of elements increases until reaching more than 110 million points for 121 processes. It is interesting to note that the total computational time spent per process remains constant for more than 10 processes in the non-blocking case. Similar results are obtained for blocking communications.

14 R. Martin, D. Komatitsch, C. Blitz, N. Le Goff

- Cohen, G., Joly, P., Tordjman, N.: Construction and analysis of higher-order finite elements with mass lumping for the wave equation. Proceedings of the second international conference on mathematical and numerical aspects of wave propagation (1993), R. Kleinman Editor, pp. 152-160, SIAM, Philadelphia, Pennsylvania, USA.
- Priolo, E., Carcione, J. M., Seriani, G.: Numerical simulation of interface waves by high-order spectral modeling techniques. Journal of Acoustical Society of America (1994), 95, 2, pp. 681-693.
- 4. Komatitsch, D., Martin, R., Tromp, J., Taylor, M.A., Wingate, B.A.: Wave propagation in 2-D elastic media using a spectral element method with triangles and quadrangles. Journal of Computational Acoustics (2001), 9, 2, pp. 703-718.
- Komatitsch, D., Vilotte, J. P.: The spectral-element method: an efficient tool to simulate the seismic response of 2D and 3D geological structures. Bull. Seis. Soc. Am. (1998), 88, 2, pp. 368-392.
- Komatitsch, D., Tromp, J.: Introduction to the spectral-element method for 3-D seismic wave propagation. Geophy. J. Int (1999), 139, 3, pp. 806-822.
- Komatitsch, D., Tsuboi, S., Ji, C., Tromp, J.: A 14.6 billion degrees of freedom, 5 teraflops, 2.5 terabyte earthquake simulation on the Earth Simulator. Proceedings of the ACM/IEEE Supercomputing SC'2003 conference (2003), published on CD-ROM and at www.sc-conference.org/sc2003.
- Asphaug, E.: Interior structures for asteroids and cometary nuclei. In: Belton, M. J. S., Morgan, T. H., Samarasinha, N., Yeomans, D. K. (eds): Mitigation of Hazardous Comets and Asteroids. Cambridge University Press (2004), pp. 66-103.
- Holsapple, K. A.: About deflecting asteroids and comets. In: Belton, M. J. S., Morgan, T. H., Samarasinha, N., Yeomans, D. K. (eds): Mitigation of Hazardous Comets and Asteroids. Cambridge University Press (2004), pp. 113-140.
- 10. Michel, P., Benz, W., Richardson, D. C.: Disruption of fragmented parent bodies as the origin of asteroids families. Nature (2003), 421, pp. 608-611.
- Robinson, M. S., Thomas, P. C, Veverka, J., Murchie, S. L., Wilcox, B. B.: The geology of 433 Eros. Meteoritics and planetary sciences (2002), 37, pp. 1651-1684.
- Richardson, J. E., Melosh, H. J., Greenberg, R. J. and O'Brien, D. P.: The global effects of impact-induced seismic activity on fractured asteroid surface morphology. Icarus (2005), 179, pp. 325-349.
- De Basabe, J. D. and Sen, M. K. :Grid dispersion and stability criteria of some common finite-element methods for acoustic and elastic wave equations. Geophysics (2007), 72, 6, pp. 81-95.
- Karypis, G., Kumar, V.: Multilevel k-way Partitioning Scheme for Irregular Graphs. J. Parallel Distrib. Comput. (1998), 48, 1, pp. 96-129.
- Danielson, K. T. and Namburu R. R.: Nonlinear dynamic finite element analysis on parallel computers using Fortran90 and MPI. Advances in Engineering Software (1998), 29, 3-6, pp. 179-186.
- Ahrens, T. J., Xia, K. and Coker, D.: Depth of cracking beneath impact craters: new constraint for impact velocity, in: Shock-Compression of Condensed Matter, edited by Furnish, M. D., Thadhani, N. N., and Y. Horie. American Institute of Physics, New York (2002), pp. 1393-1396.
- 17. Komatitsch, D., Labarta, J., Michea, D.: A simulation of seismic wave propagation at high resolution in the inner core of the Earth on 2166 processors of MareNostrum. Proceedings of the VecPar'2008 8th International Meeting on High Performance Computing for Computational Science (2008), Toulouse, France, 24-27 June 2008, in press.