Large-scale parallel simulations of earthquakes at high frequency: the SPECFEM3D project

Dimitri Komatitsch, University of Pau, Institut universitaire de France and INRIA Magique3D, France

Jeroen Tromp et al., Caltech, USA

Jesús Labarta, Sergi Girona, BSC MareNostrum, Spain

David Michéa, Nicolas Le Goff, Roland Martin, University of Pau, France

20 years of CERFACS, October 10, 2007

Global 3D Earth

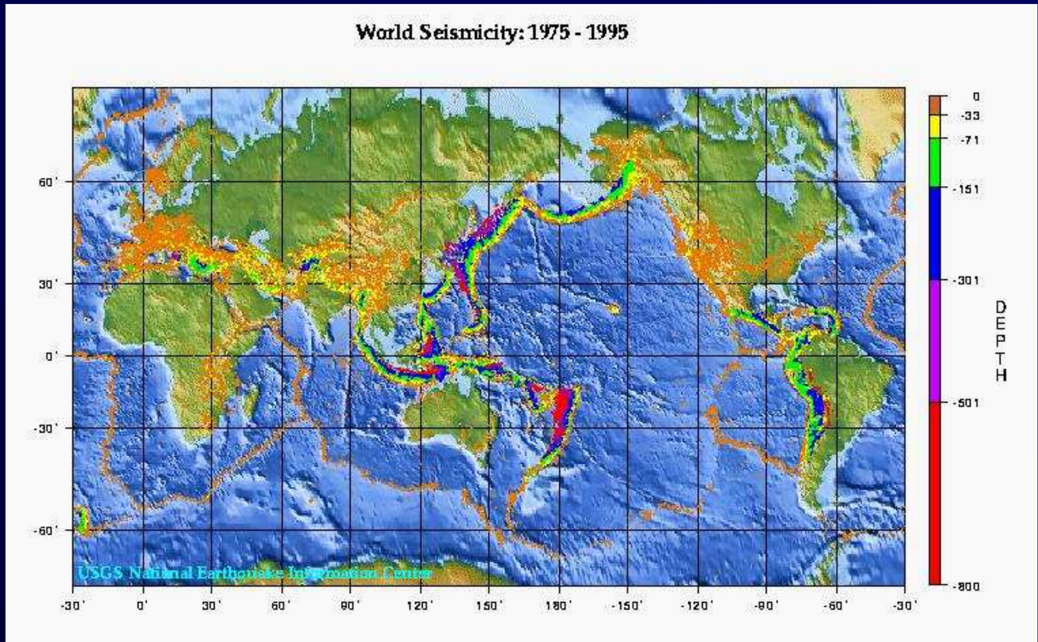

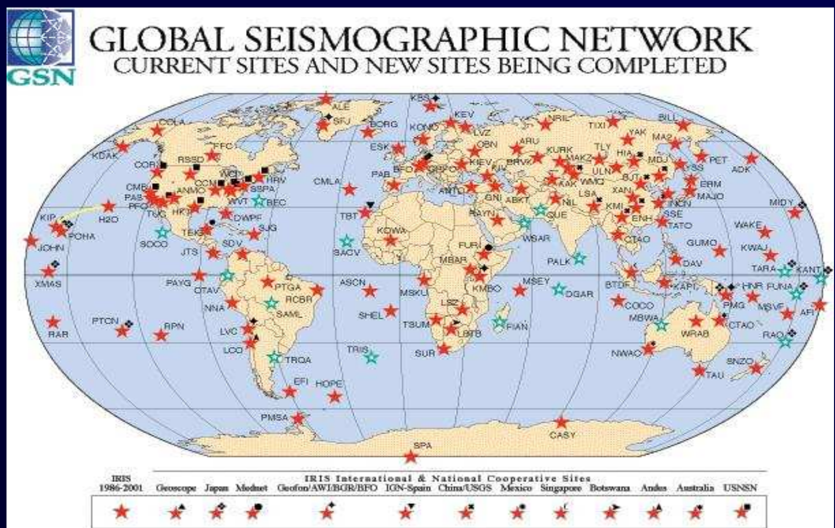

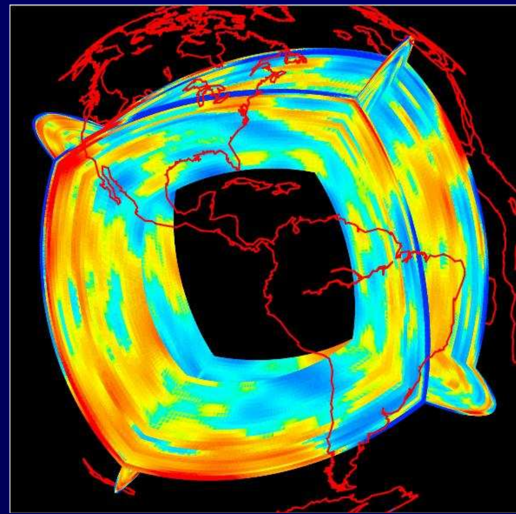

S20RTS mantle model (Ritsema et al. 1999)

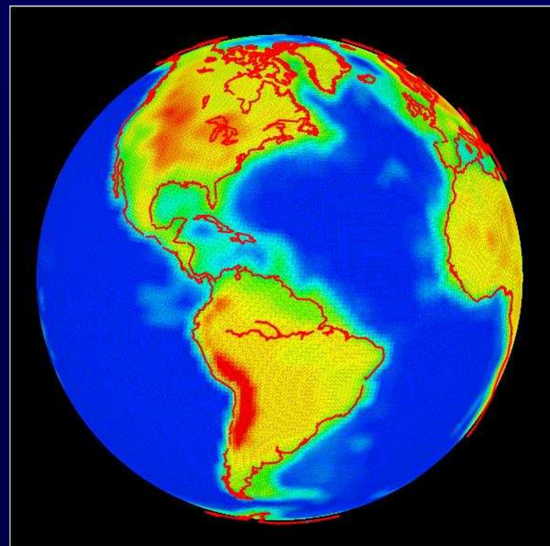

Crust 5.2 (Bassin et al. 2000)

Brief history of numerical methods

Seismic wave equation : tremendous increase of computational power
⇒ development of numerical methods for accurate calculation of synthetic seismograms in complex 3D geological models has been a continuous effort in last 30 years.

Finite-difference methods : Yee 1966, Chorin 1968, Alterman and Karal 1968, Madariaga 1976, Virieux 1986, Moczo et al, Olsen et al..., difficult for boundary conditions, surface waves, topography, full Earth

Boundary-element or boundary-integral methods (Kawase 1988, Sanchez-Sesma et al. 1991) : homogeneous layers, expensive in 3D

Spectral and pseudo-spectral methods (Carcione 1990) : smooth media, difficult for boundary conditions, difficult on parallel computers

Classical finite-element methods (Lysmer and Drake 1972, Marfurt 1984, Bielak et al 1998) : linear systems, large amount of numerical dispersion

Spectral-Element Method

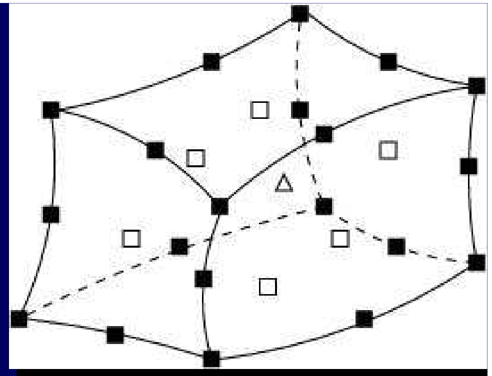

- Developed in Computational Fluid Dynamics (Patera 1984)
- Accuracy of a pseudospectral method, flexibility of a finite-element method
- Extended by Komatitsch and Tromp, Chaljub et al.
- Large curved “spectral” finite-elements with high-degree polynomial interpolation
- Mesh honors the main discontinuities (velocity, density) and topography
- Very efficient on parallel computers, no linear system to invert (diagonal mass matrix)

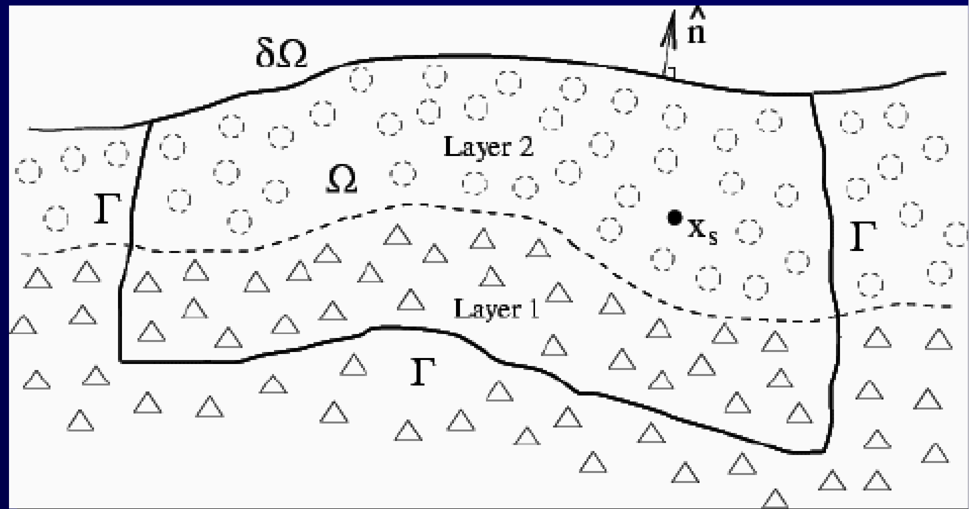

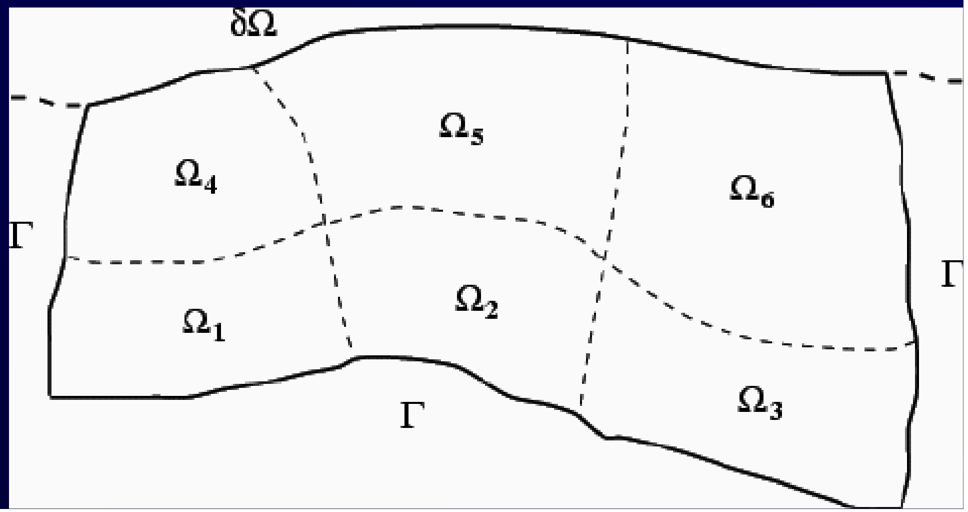

Equations of Motion (solid)

Differential or *strong* form (e.g., finite differences):

$$\rho \partial_t^2 \mathbf{s} = \nabla \cdot \mathbf{T} + \mathbf{f}$$

We solve the integral or *weak* form:

$$\int \rho \mathbf{w} \cdot \partial_t^2 \mathbf{s} \, \mathrm{d}^3\mathbf{r} = -\int \nabla \mathbf{w} : \mathbf{T} \, \mathrm{d}^3\mathbf{r}$$

$$+ \mathbf{M} : \nabla \mathbf{w}(\mathbf{r}_s) S(t) - \int_{\mathrm{F-S}} \mathbf{w} \cdot \mathbf{T} \cdot \hat{\mathbf{n}} \, \mathrm{d}^2\mathbf{r}$$

+ attenuation (memory variables) and ocean load

Equations of Motion (Fluid)

Differential or *strong* form:

$$\rho \partial_t \mathbf{v} = -\nabla p \qquad \partial_t p = -\kappa \nabla \cdot \mathbf{v}$$

We use a generalized velocity potential χ $\quad p = \partial_t x$

the integral or *weak* form is:

$$\int \kappa^{-1} w \partial_t^2 x \mathrm{d}^3\mathbf{r} = -\int \rho^{-1} \nabla w \cdot \nabla x \mathrm{d}^3\mathbf{r}$$

$$+\int_{\mathrm{F-S}} w \hat{\mathbf{n}} \cdot \mathbf{v} \, \mathrm{d}^2\mathbf{r}$$

⇒ 3 times cheaper (scalar potential)

⇒ natural coupling with solid

Finite Elements

- High-degree pseudospectral finite elements with Gauss-Lobatto-Legendre integration
- N = 5 to 8 usually
- *Exactly* diagonal mass matrix
- No linear system to invert

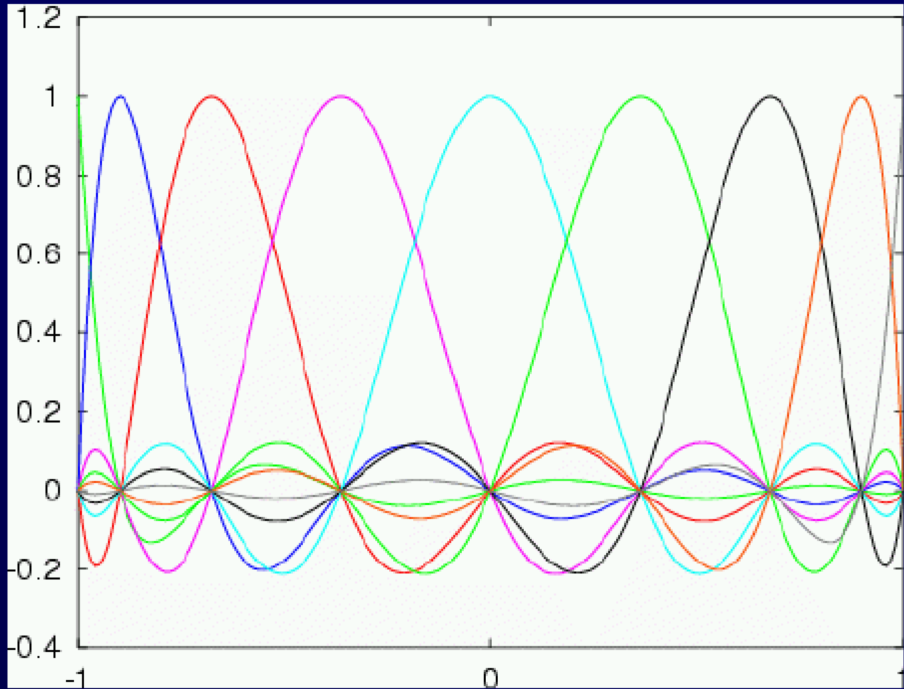

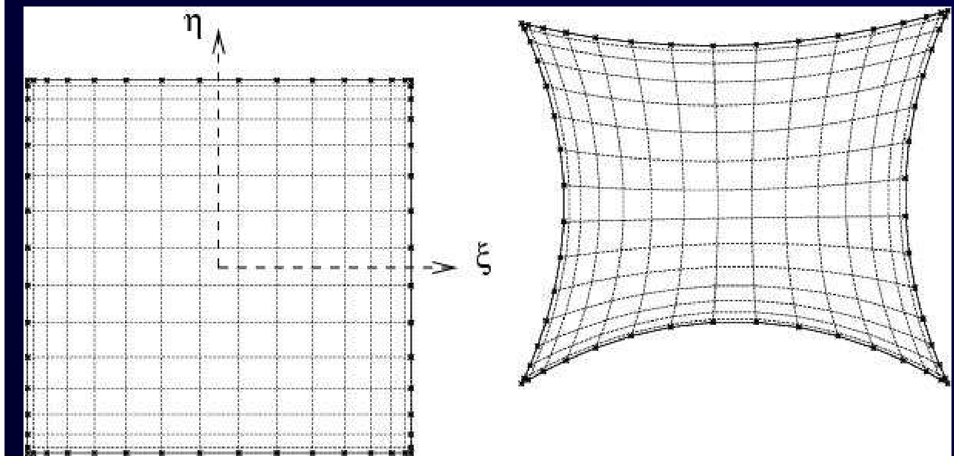

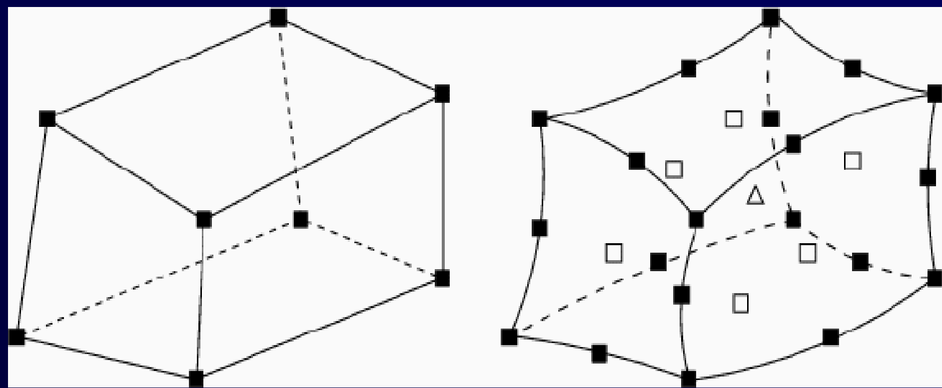

The Challenge of the Global Earth

- A slow, thin, highly variable crust
- Sharp radial velocity and density discontinuities
- Fluid-solid boundaries (outer core of the Earth)
- Anisotropy
- Attenuation
- Ellipticity, topography and bathymetry
- Rotation
- Self-gravitation
- 3-D mantle and crust models (lateral variations)

The Cubed Sphere

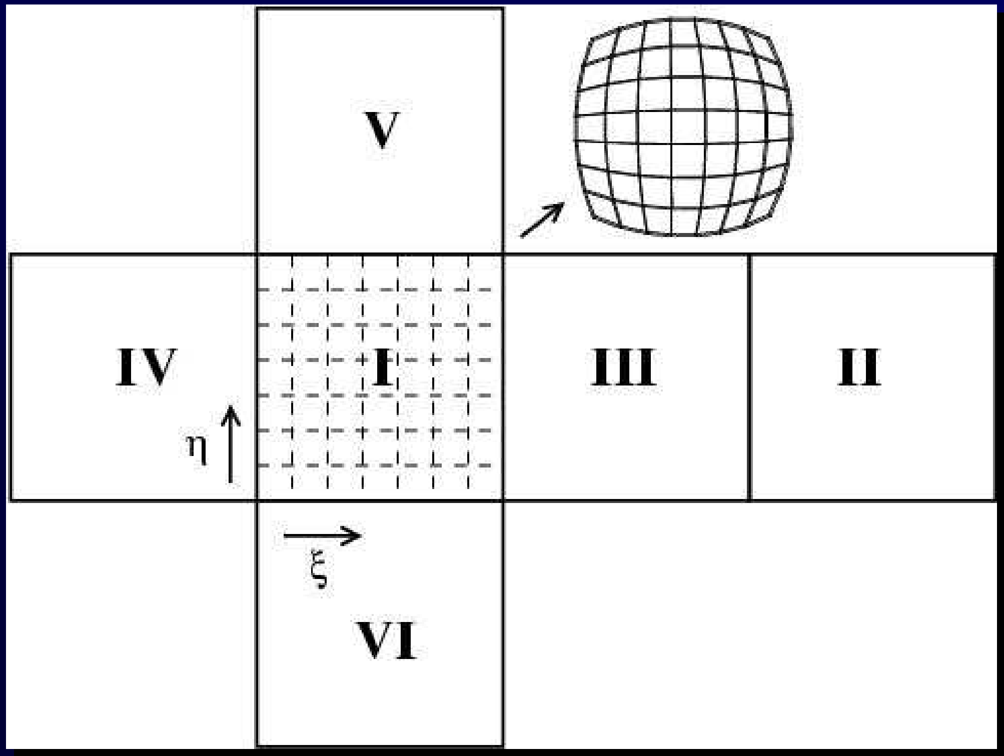

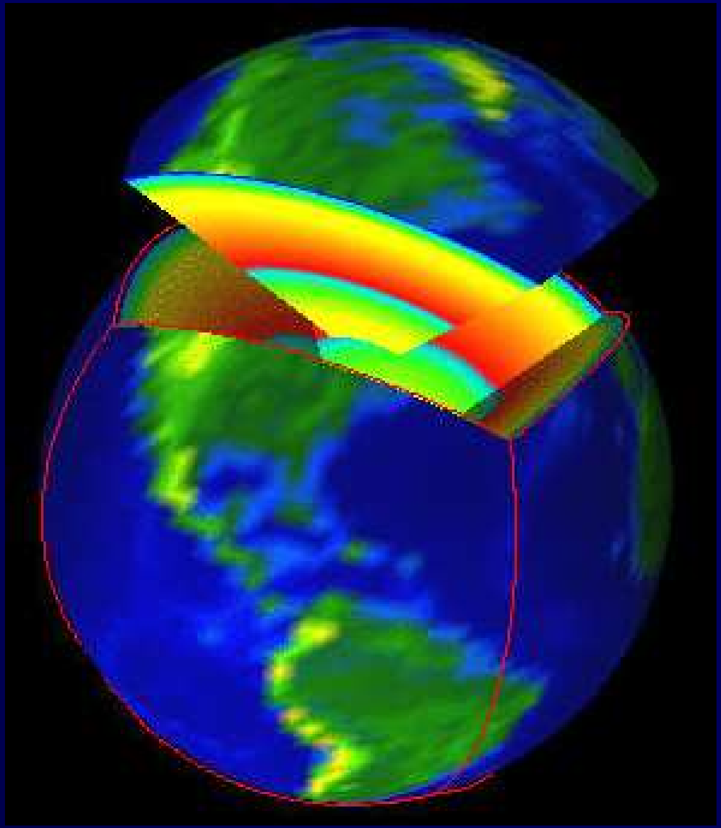

- “Gnomonic” mapping (Sadourny 1972)
- Ronchi et al. (1996), Chaljub (2000)
- Analytical mapping from six faces of cube to unit sphere

Final Mesh

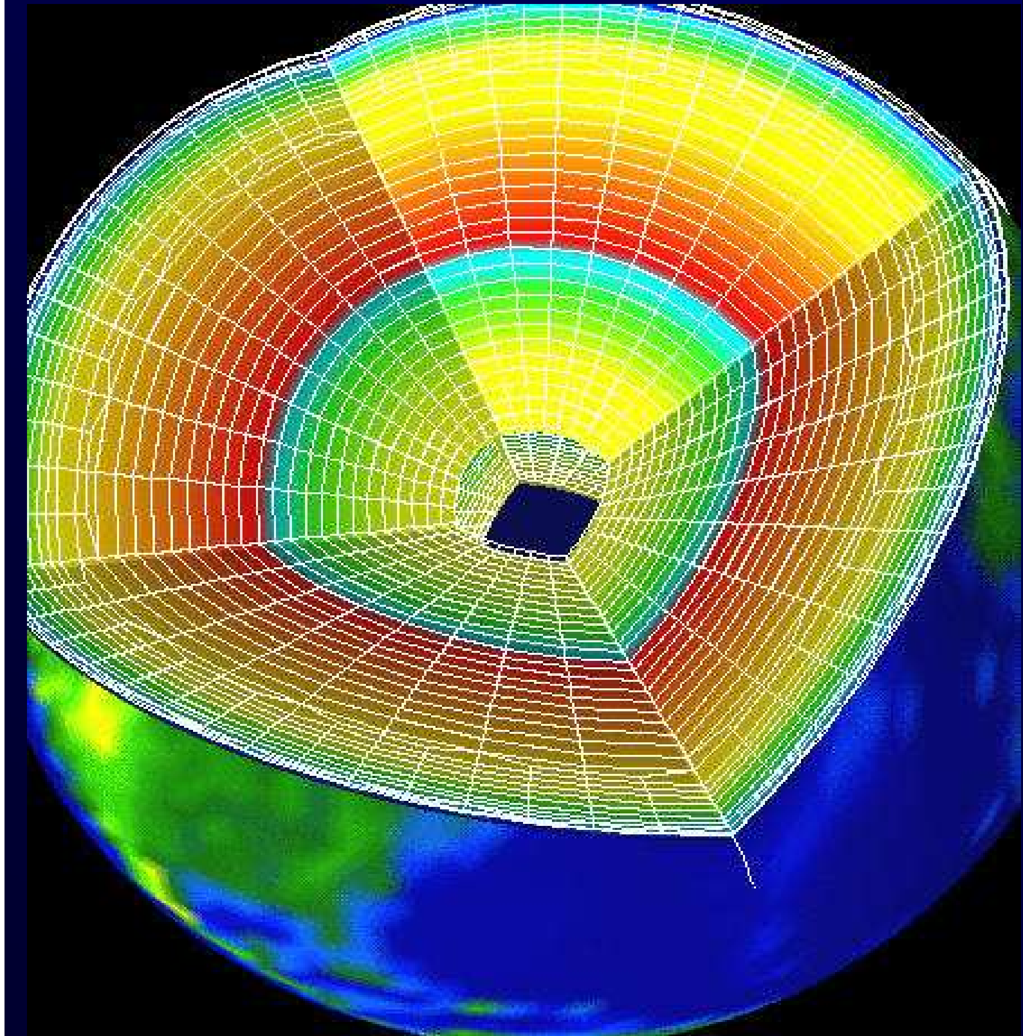

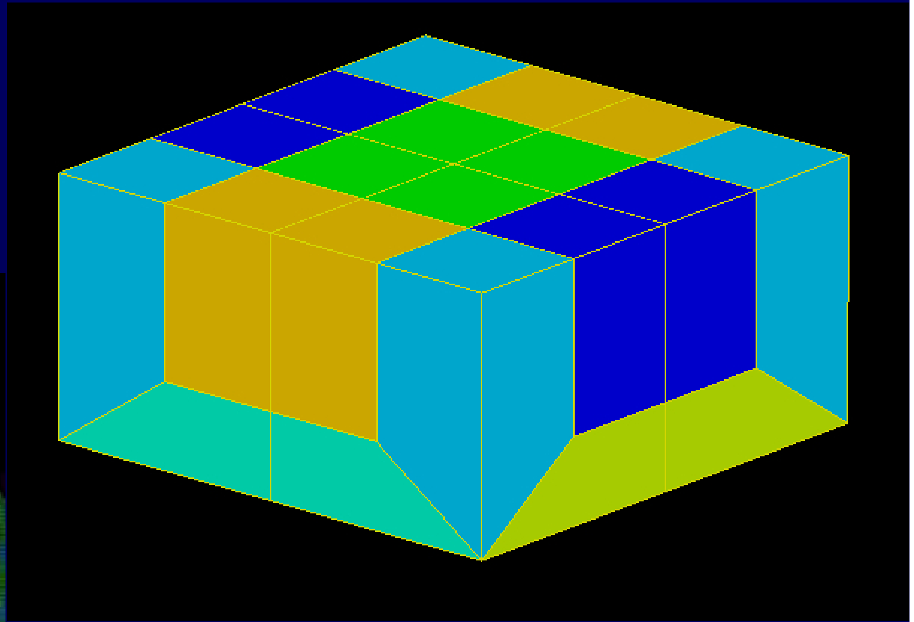

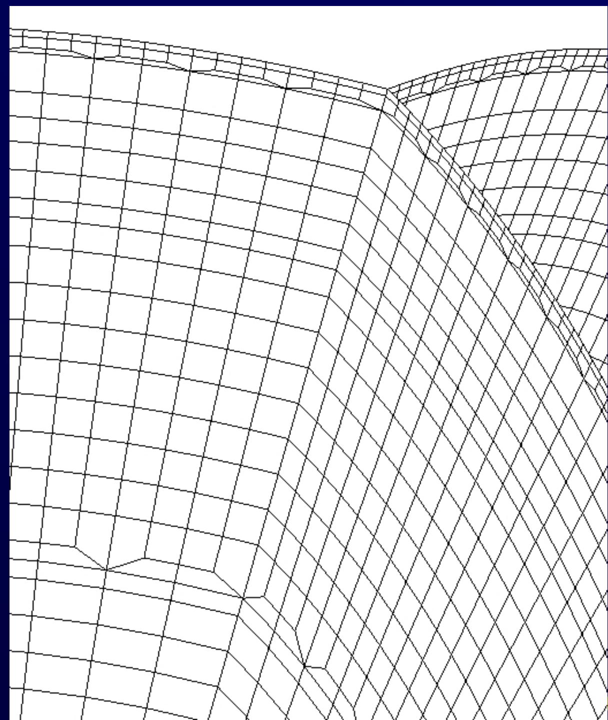

Global 3-D Earth

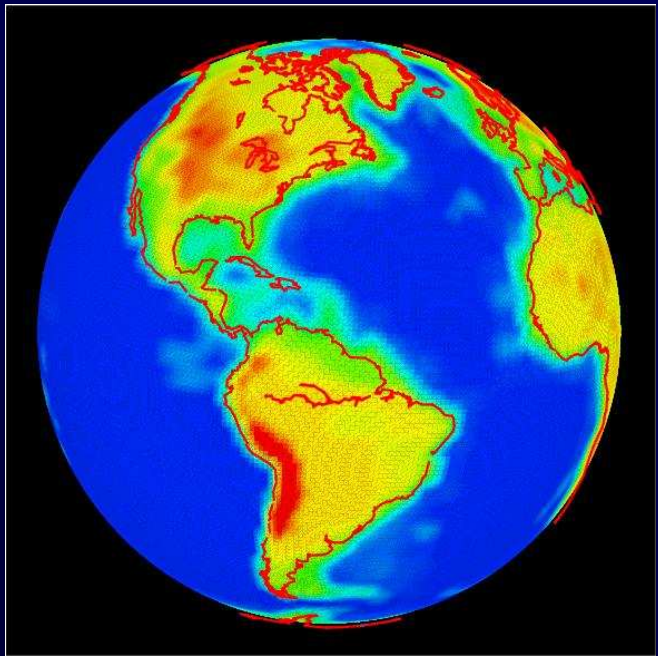

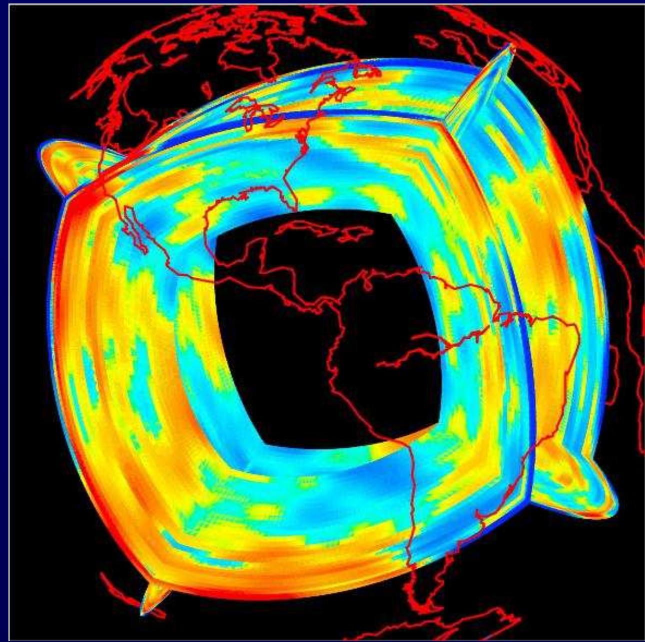

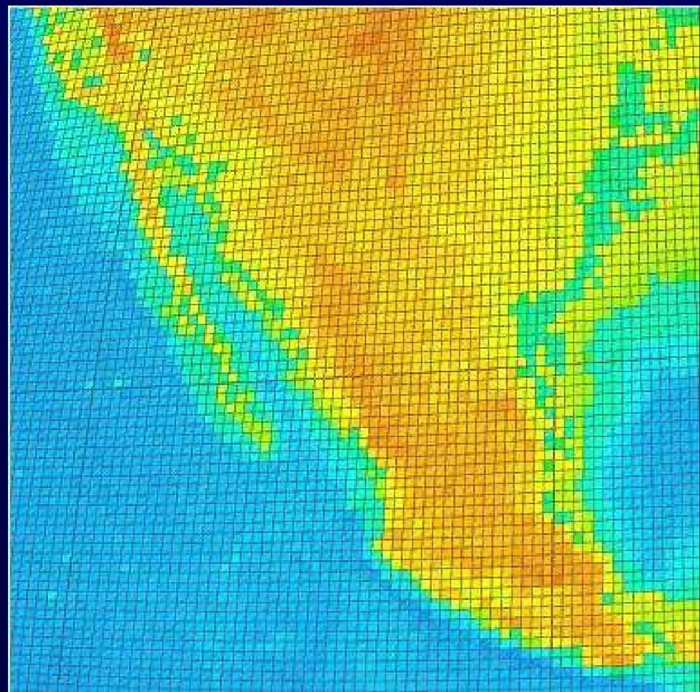

Crust 5.2 (Bassin et al. 2000)
Mantle model S20RTS (Ritsema et al. 1999)

Ellipticity and topography

Small modification
of the mesh, no problem

Topography

- Use flexibility of mesh generation
- Accurate free-surface condition

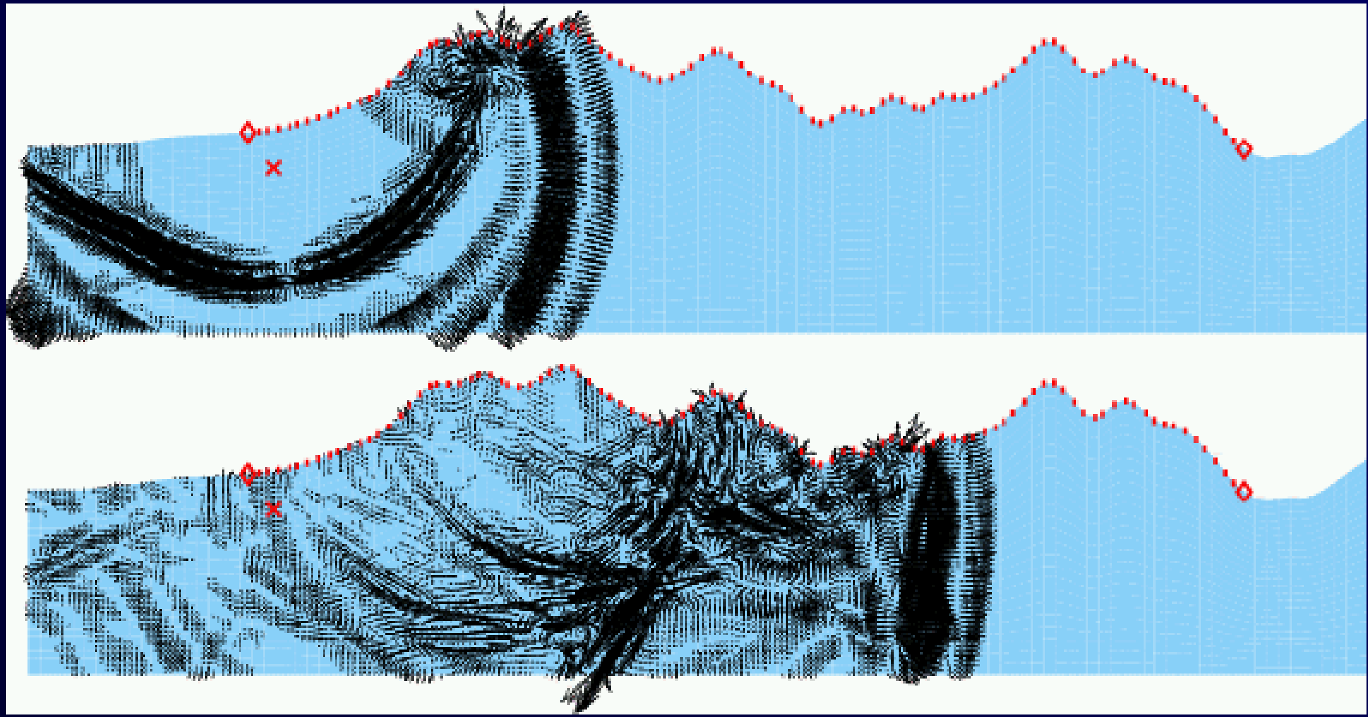

Fluid / solid

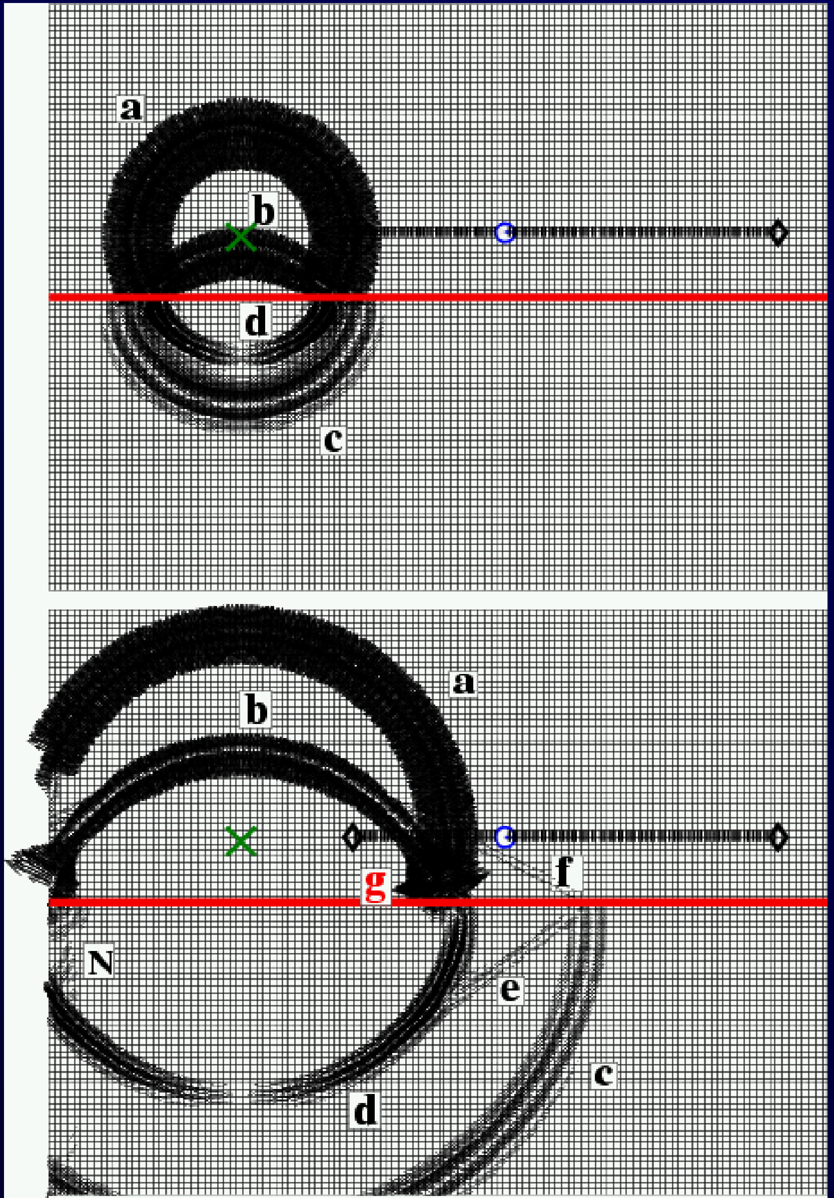

a
b
d
c
a
b
g
f
N
e
c
d

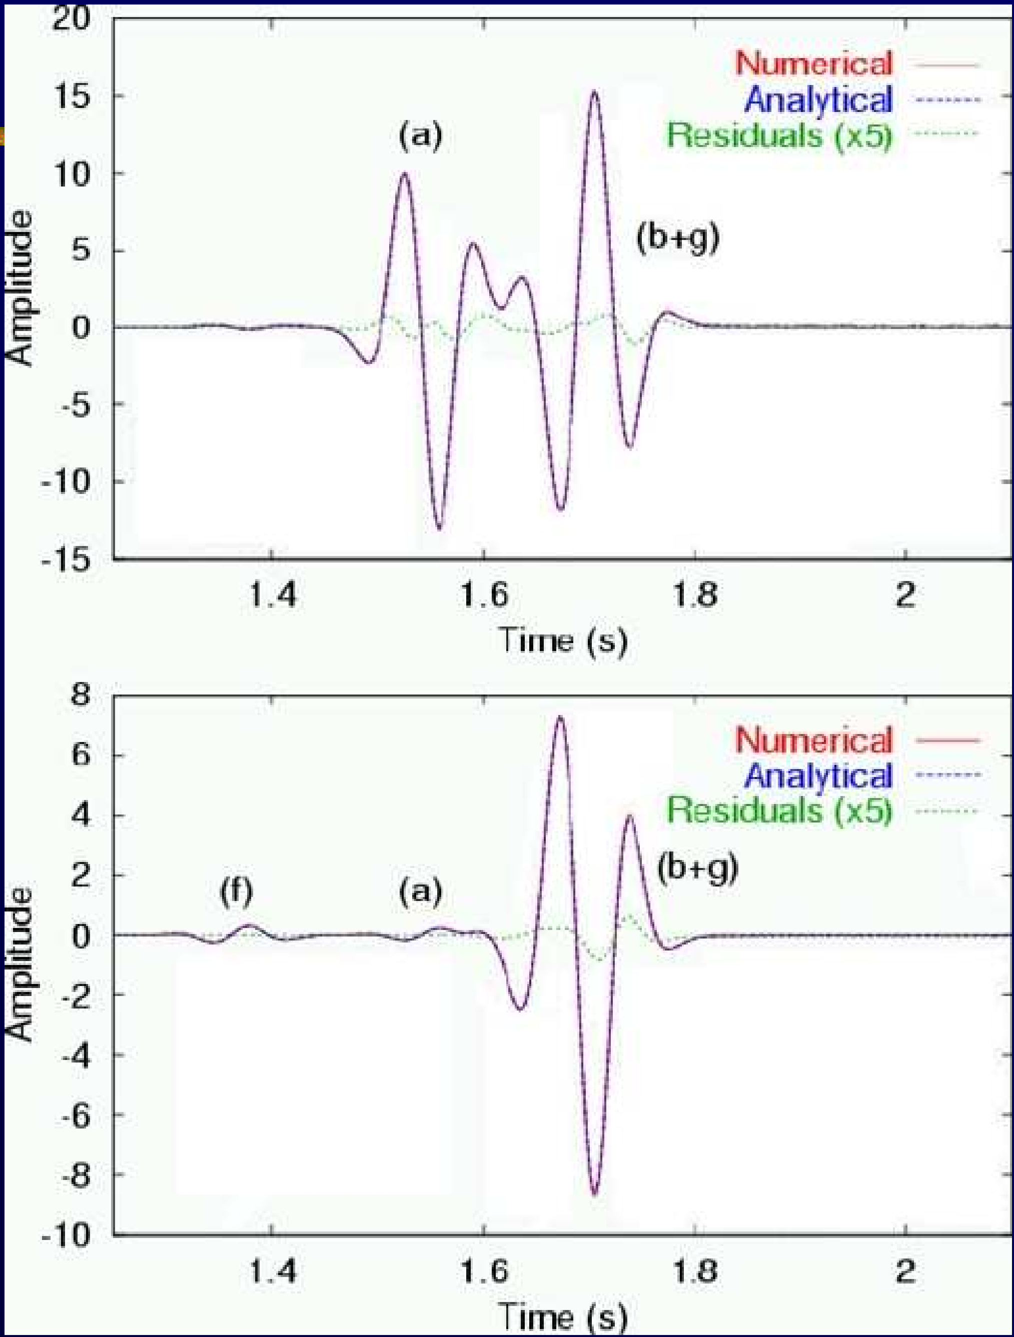

Numerical
Analytical
Residuals (x5)
(a)
(b+g)
Amplitude
Time (s)
Numerical
Analytical
Residuals (x5)
(f)
(a)
(b+g)
Amplitude
Time (s)

Bathymetry

- Use flexibility of mesh generation process
- Triplications
- Stoneley

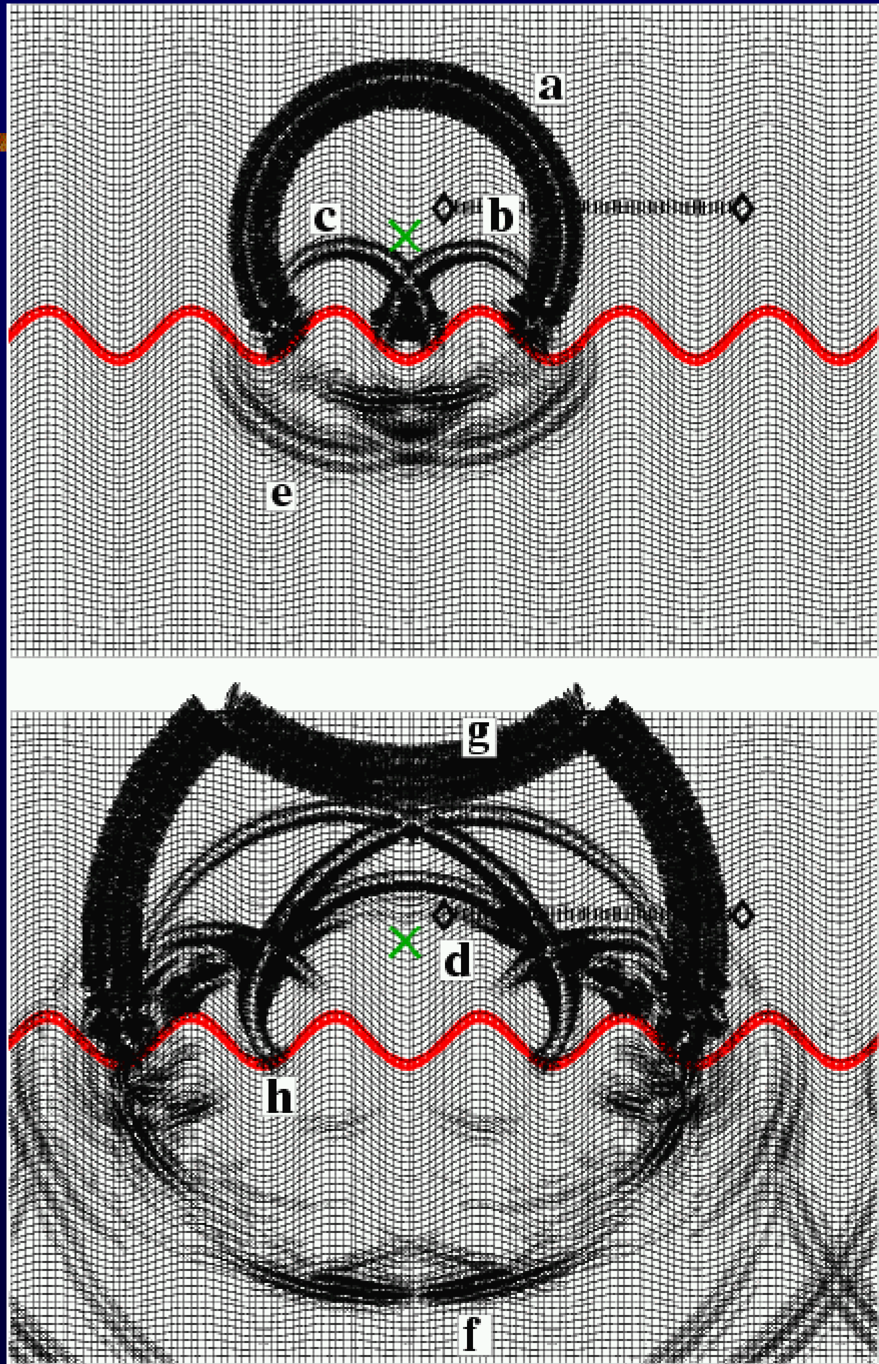

Anisotropy

- Easy to implement up to 21 coefficients
- No interpolation necessary
- Tilted axes can be modeled

Cobalt

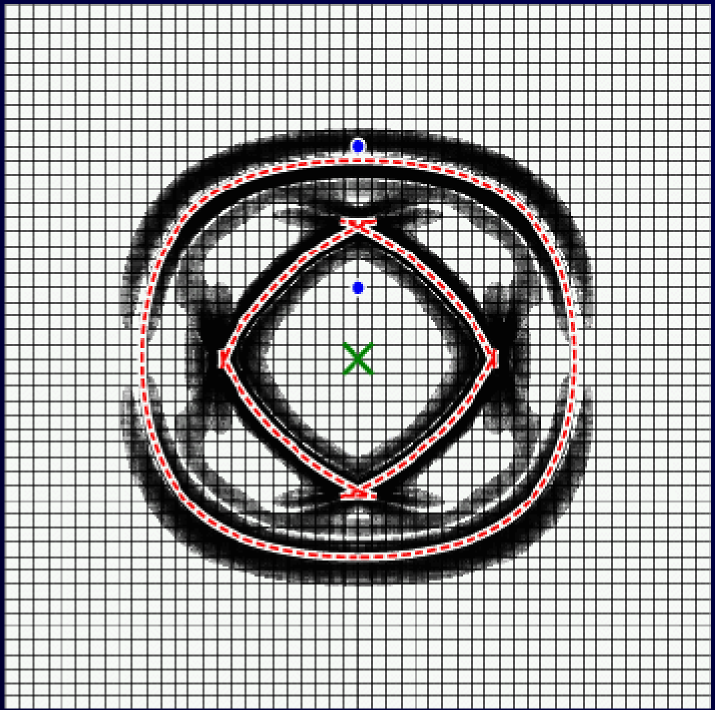

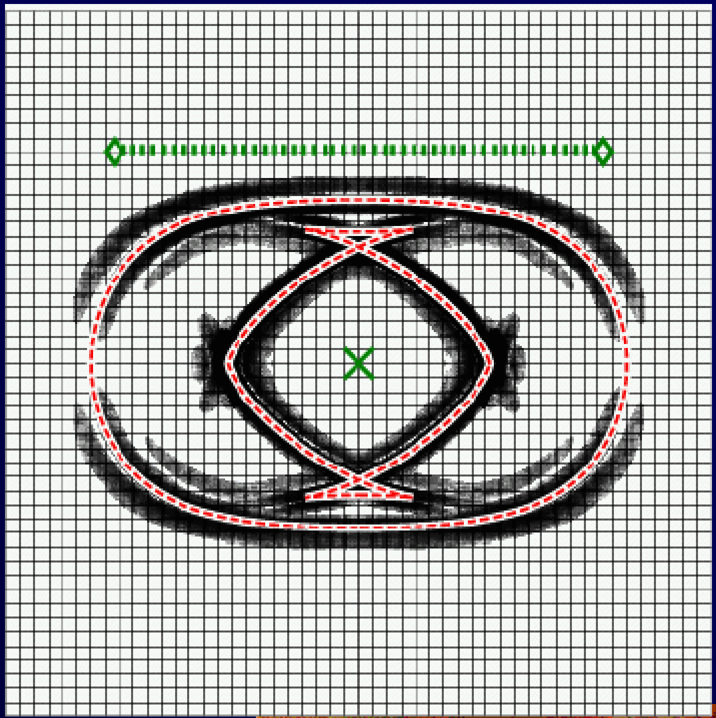

Zinc

Effect of Attenuation

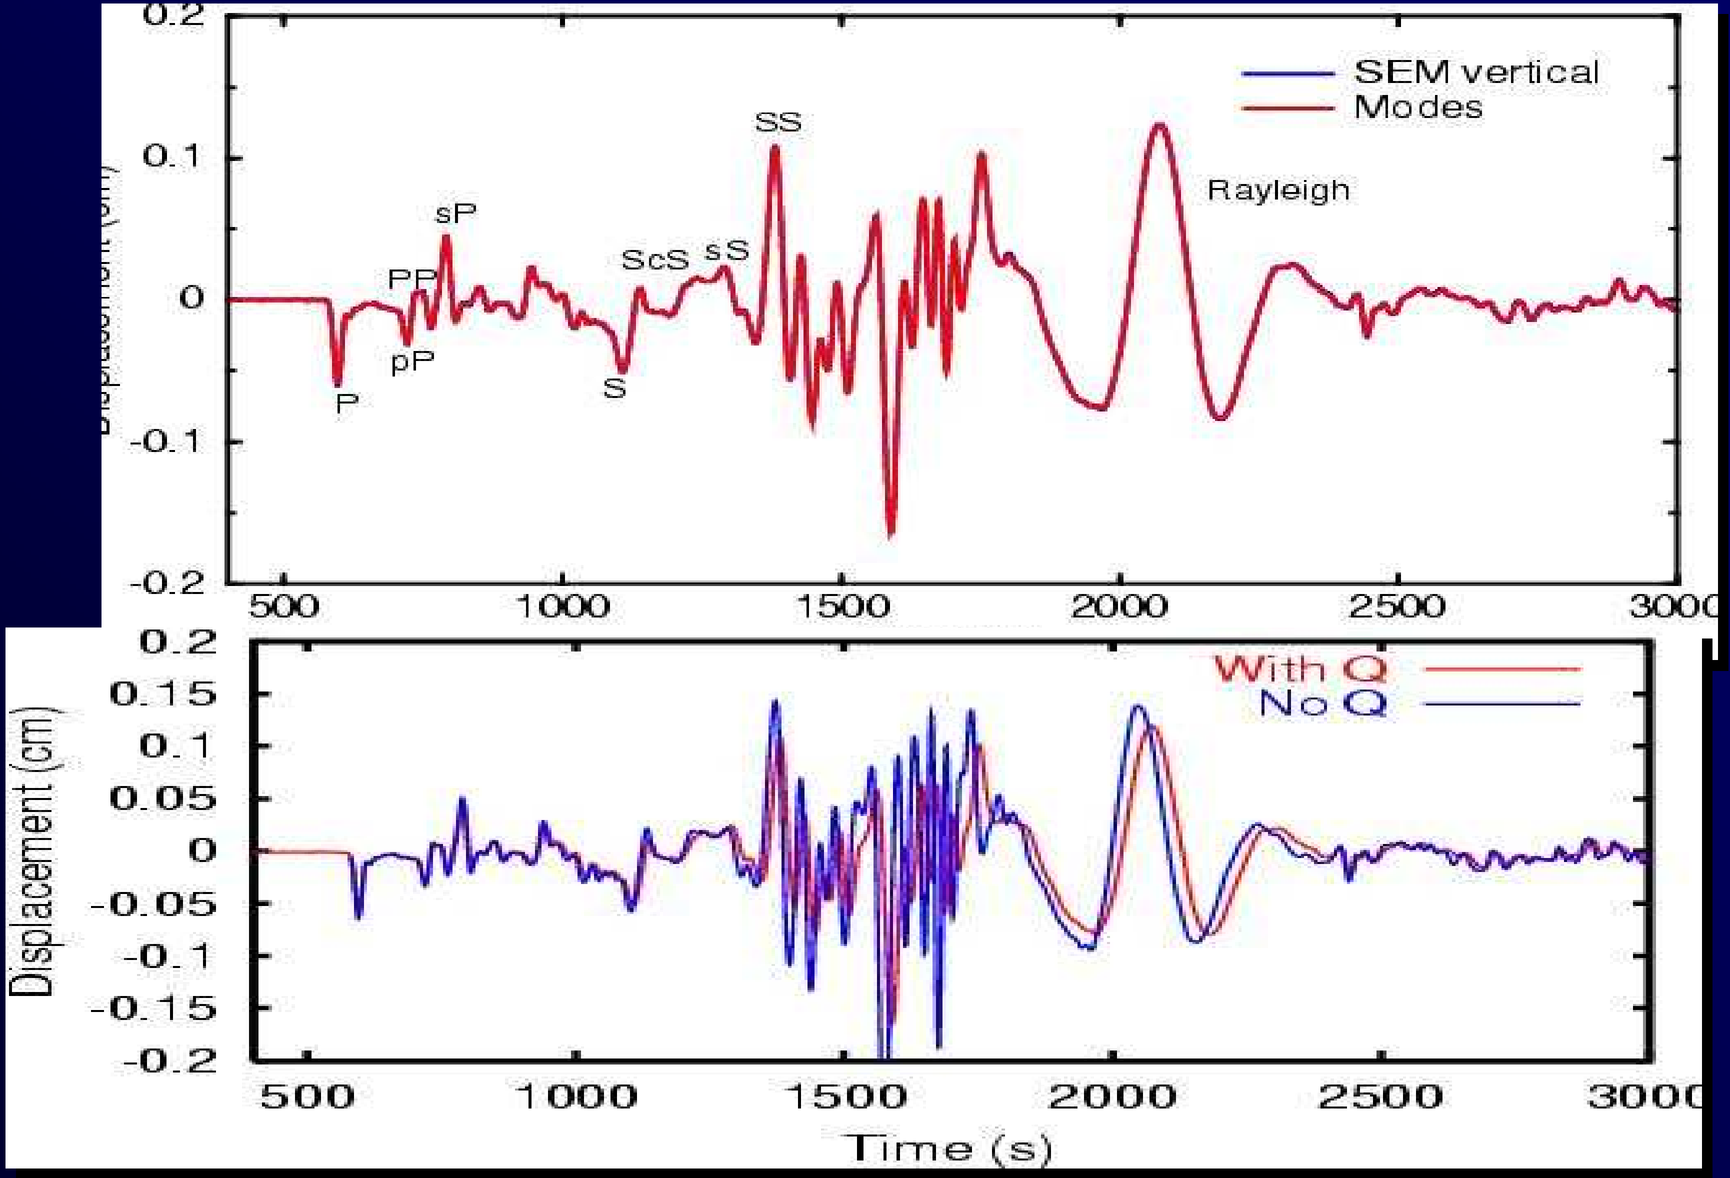

Accurate surface waves

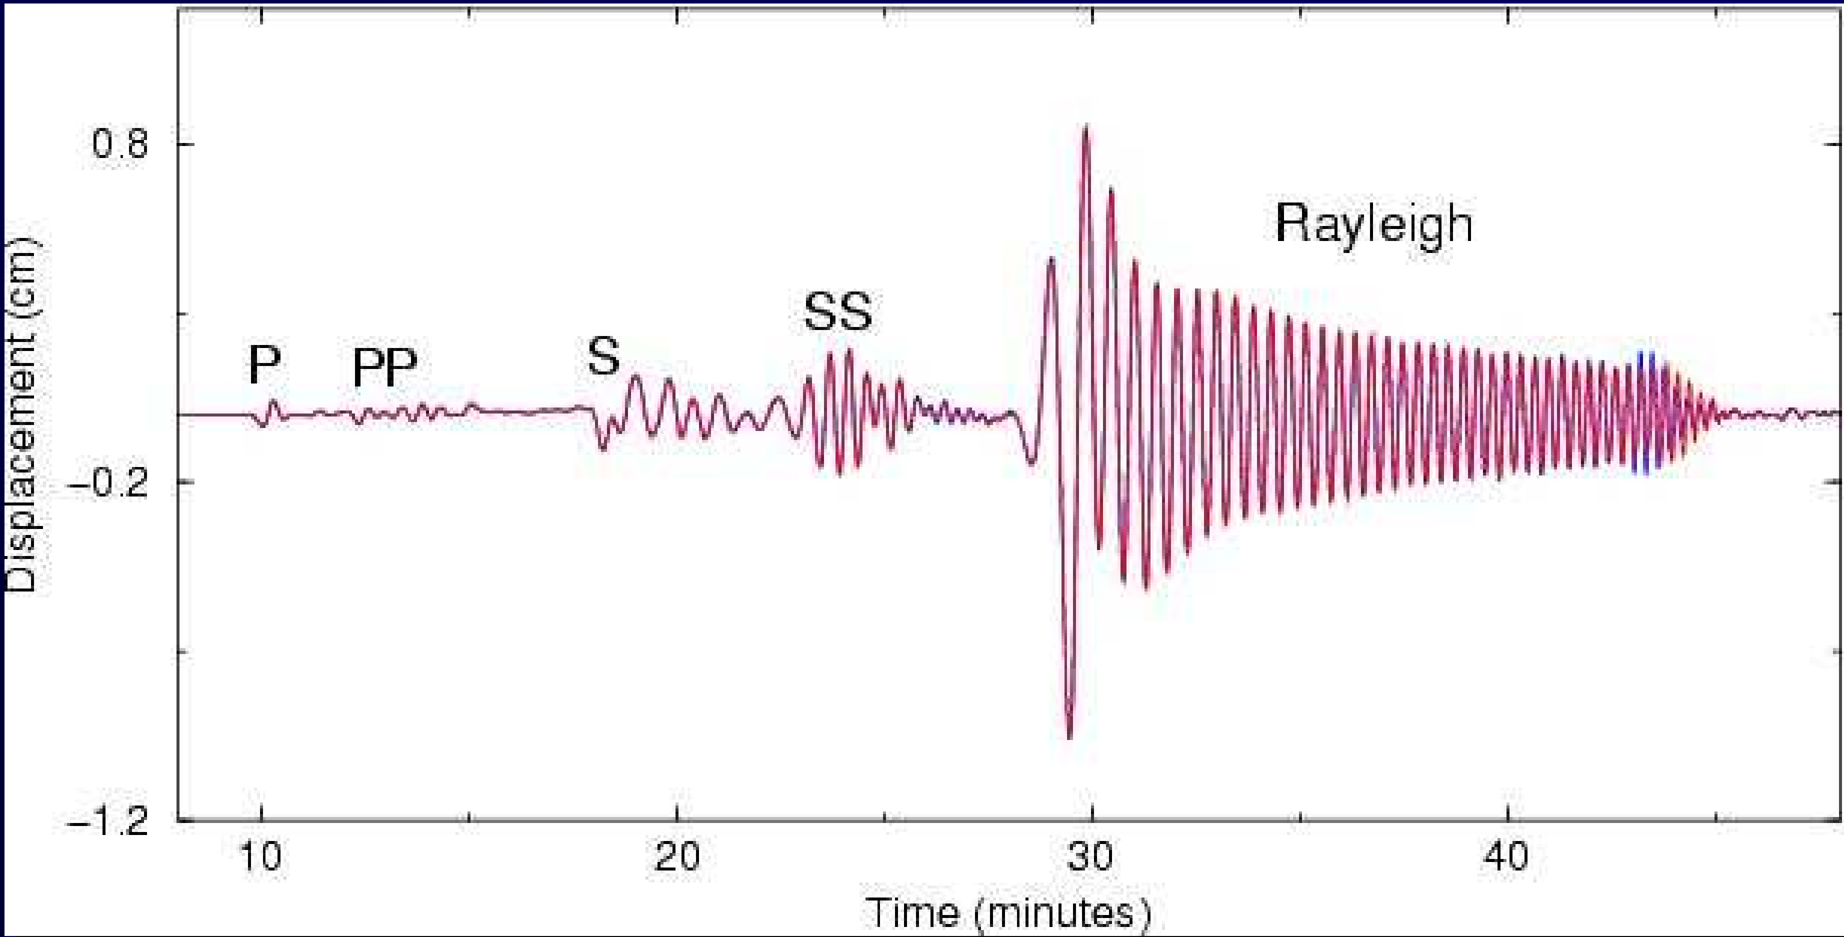

Excellent agreement with normal modes – Depth 15 km
Anisotropy included

Vanuatu

Depth 15 km

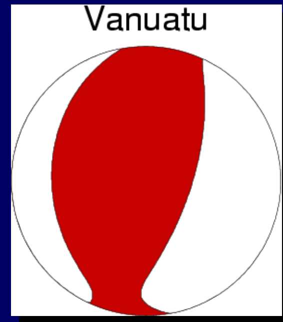

Composante verticale (onde de Rayleigh), trajet océanique, retard 85 s à Pasadena, meilleur fit au Japon

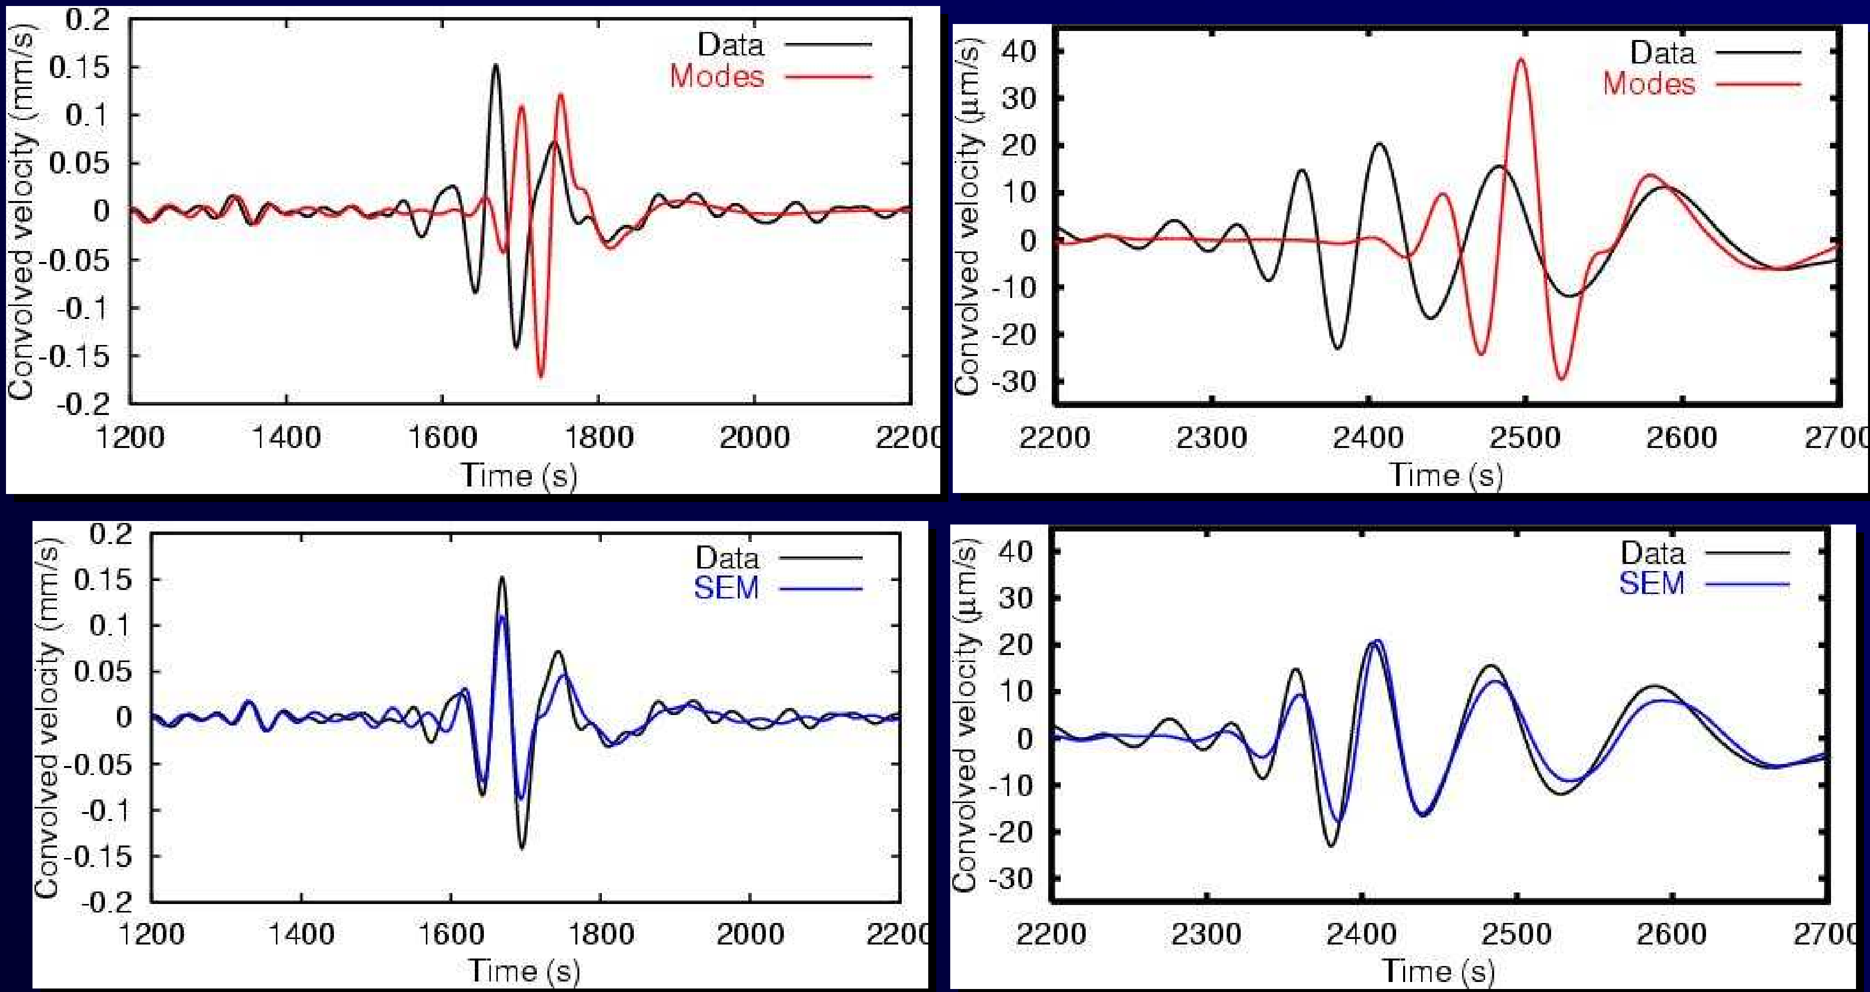

Parallel Implementation

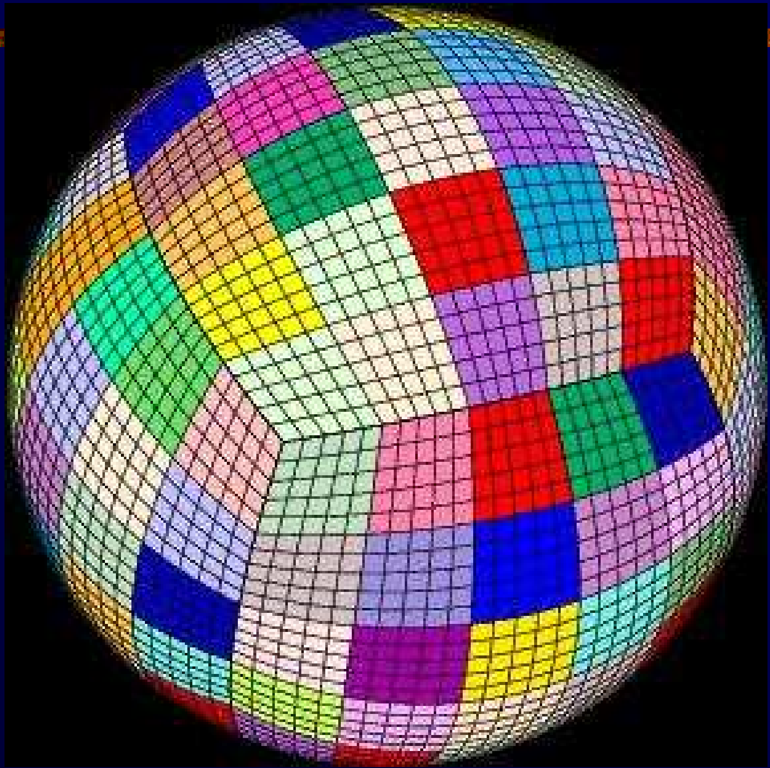

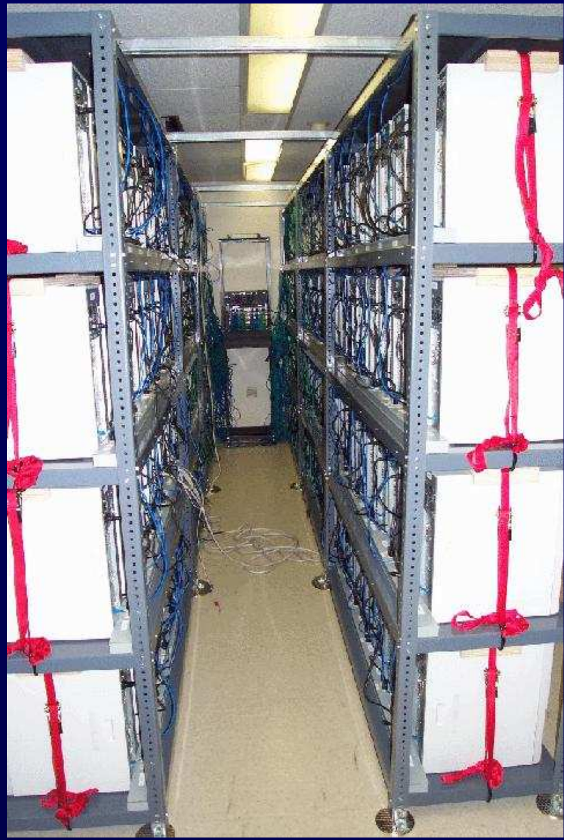

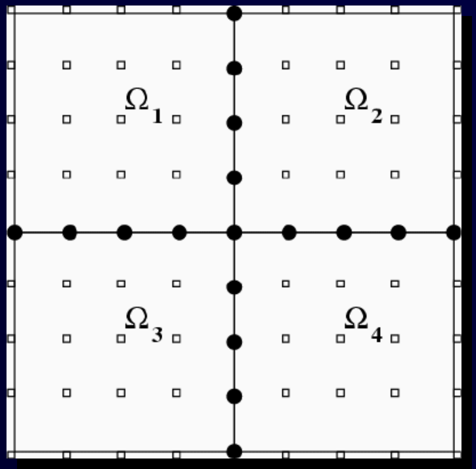

- Mesh decomposed into 150 slices
- One slice per processor – MPI communications
- Mass matrix *exactly* diagonal – no linear system
- Central cube based on Chaljub (2000)

Non-blocking MPI

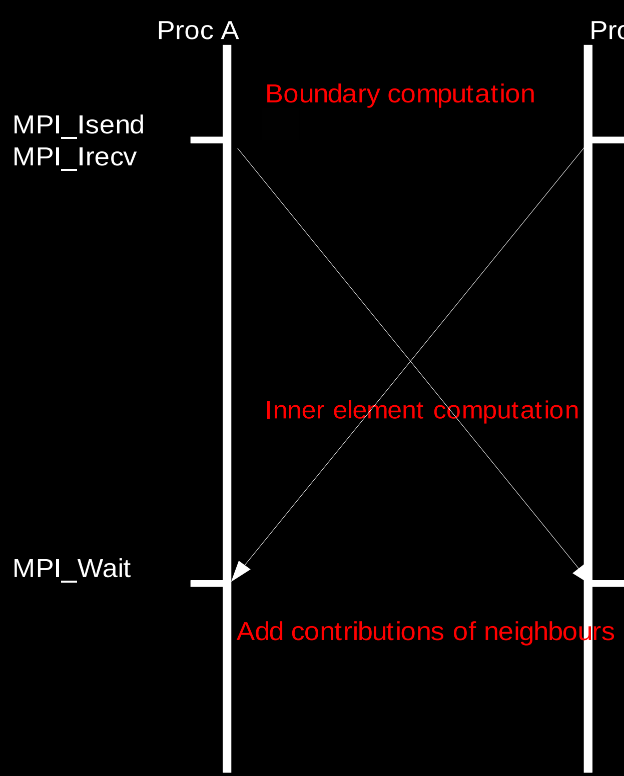

Collaboration with Roland Martin and Nicolas Le Goff (Univ of Pau, France)

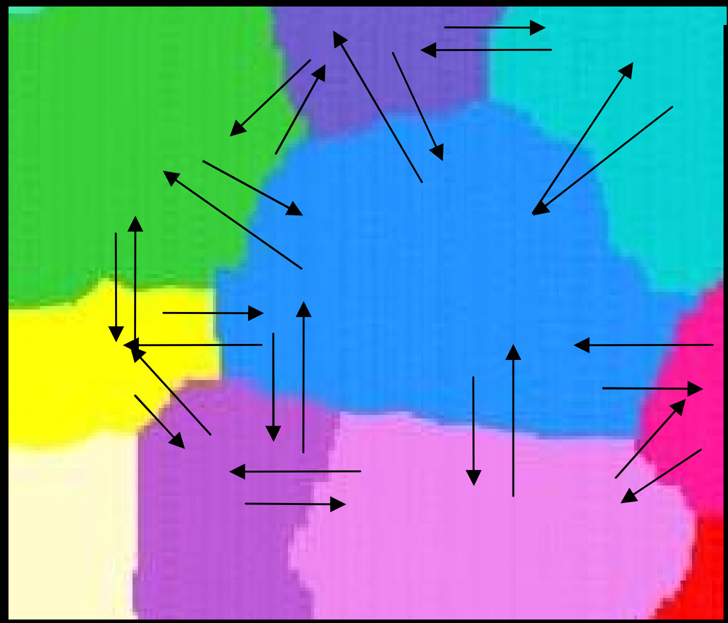

Another way to optimize MPI code is to overlap communications with computations using non-blocking MPI. But, for our code, the overall cost of communications is very small (< 5%) compared to CPU time.

Also, looping on boundary elements contradicts Cuthill-McKee order and therefore causes cache misses.

=> No need to use non blocking MPI because potential gain is comparable to overhead

=> Tested in 2D, and we did not gain anything significant

Collaboration with the oil industry

UMR tripartite

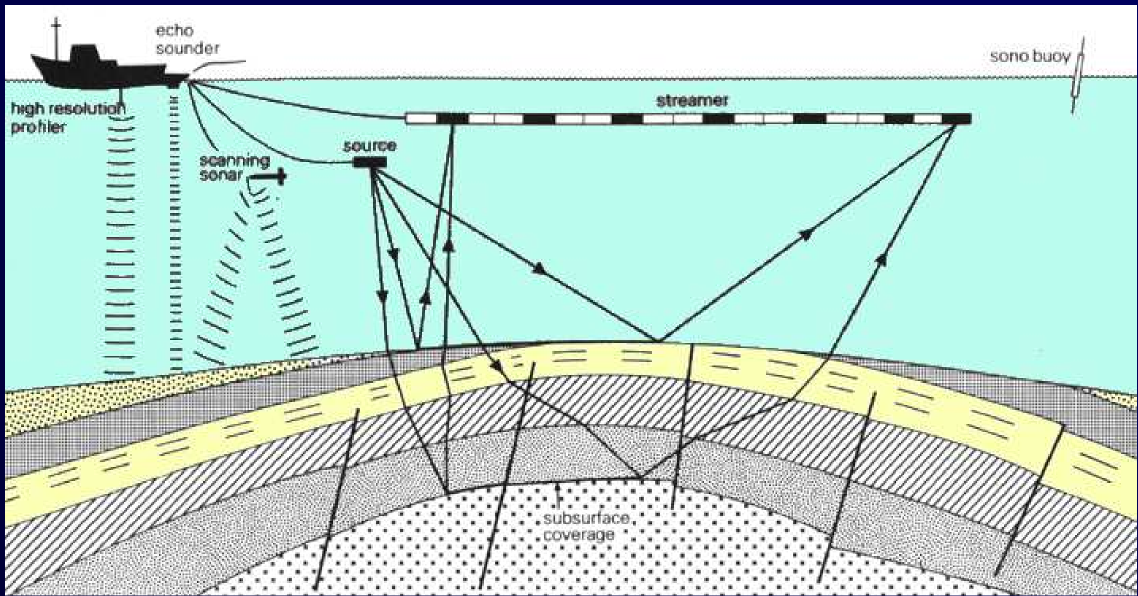

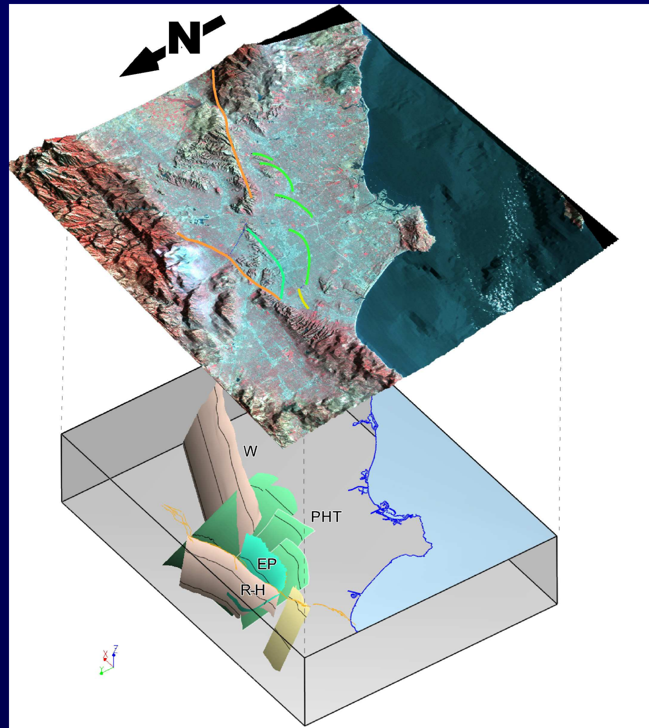

Dynamic geophysical technique of imaging subsurface geologic structures by generating sound waves at a source and recording the reflected components of this energy at receivers.

The Seismic Method is the *industry standard* for locating subsurface oil and gas accumulations.

Meshing an oil industry model

- Méthode d'éléments finis d'ordre élevé développée en **dynamique des fluides** (Patera 1984), en **sismique** 3Dpar Komatitsch et coll. (1998, 2002), Chaljub et coll. (2001).
- **SPECFEM:**Parallélisation MPI d'un Code F90 de 20000 lignes

  → **mailleur professionnel** (GiD-UPC/CIMNE)

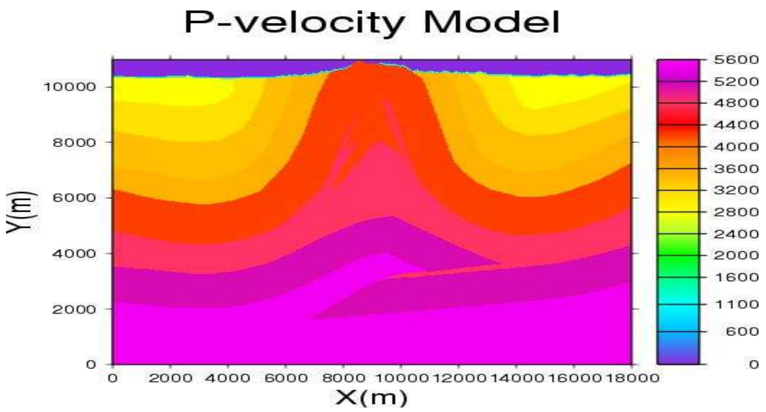

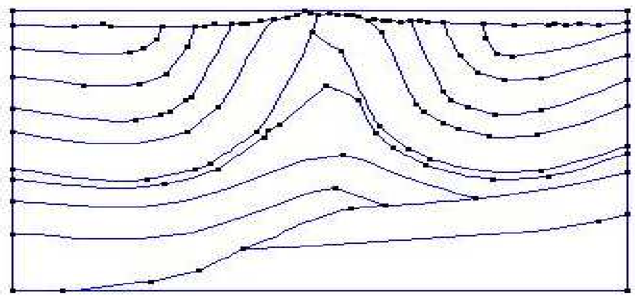

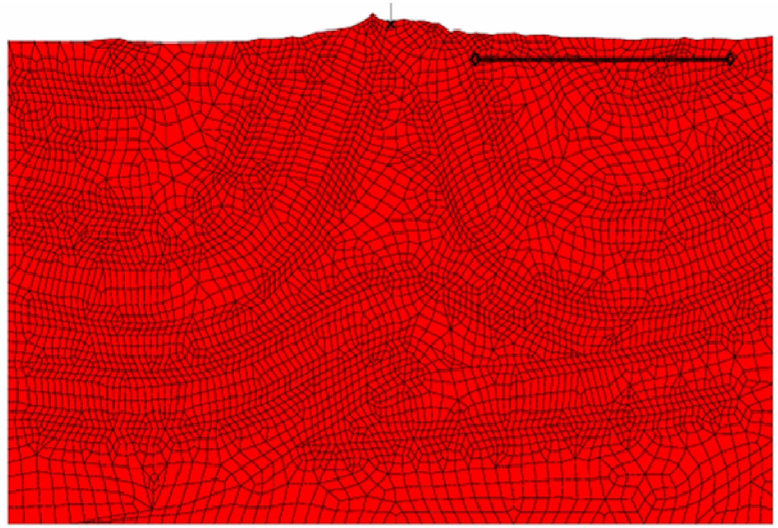

- 5.3 millions de points à 10 Hz.
- Générateur GiD automatique de maillage (UPC/ CIMNE). 98% des angles $45^o < \theta < 135^o$. Pires angles: 9.5^o and 172^o

- Structures géologiques dans les Andes (Pérou)
- Couche fine altérée en surface

→ Problème de dispersion en surface (Freq0 > 10 Hz).

Sandrine Fauqueux
Thèse INRIA/IFP (2003)

Partitionneur de domaine
(METIS or SCOTCH)

METIS or SCOTCH

Interface: gestion des communications MPI

Zone Tampon
Irecv, Isend
non bloquants

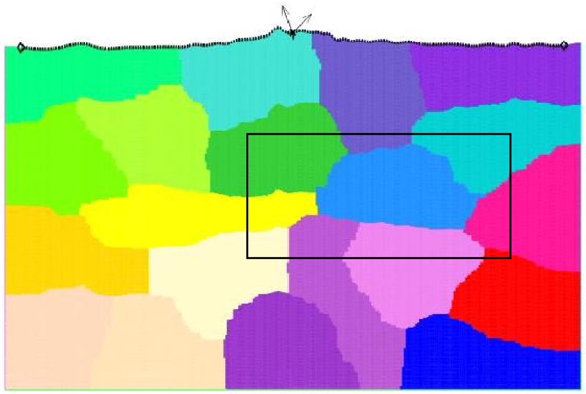

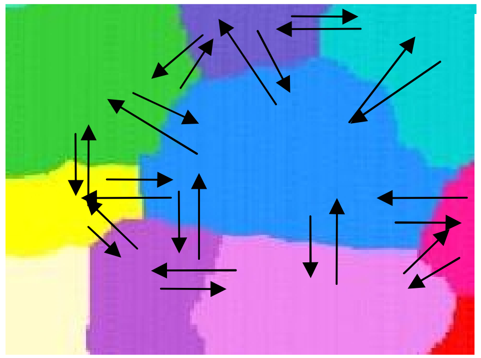

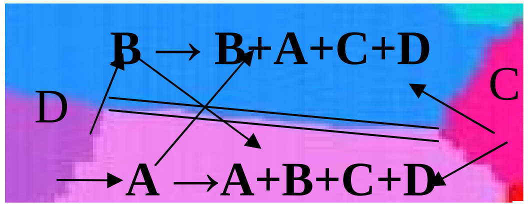

Maillage (GiD, <u>Cubit</u>)

Avantages: Nb éléments non multiple des partitions
t=∑(calcVol +calcFront+comm)
t=max[∑calcVol,∑(calcFront+comm)]
Gain -15%

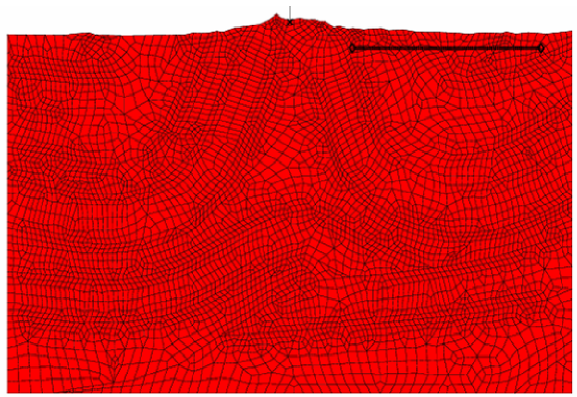

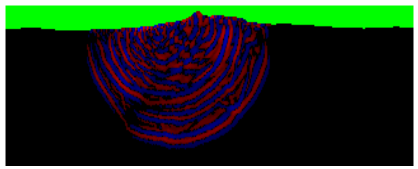

Séquentiel (10Hz)

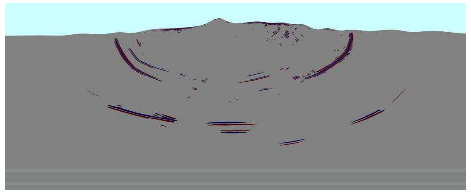

Optimization of global addressing

In 3D and for NGLL=5 (Q4), for a regular hexahedral mesh there are:

- 125 GLL integration points in each element
  - 27 belong only to this element (21.6%)
  - 54 belong to 2 elements (43.2%)
  - 36 belong to 4 elements (28.8%)
  - 8 belong to 8 elements (6.4%)

=> 78.4% of the GLL integration points belong to at least 2 elements

=> it is crucial to reuse these points by keeping them in the cache

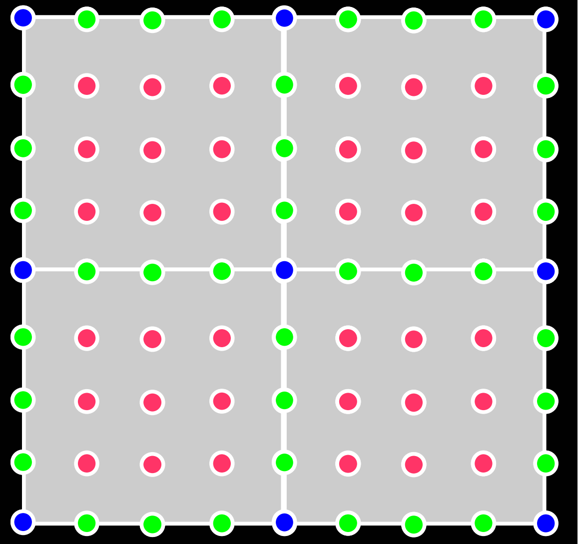

We use the classical reverse Cuthill-McKee (1969) algorithm, which consists in renumbering the vertices of the graph to reduce the bandwidth of the adjacency matrix

We gain a factor of 1.55 in CPU time on Intel Itanium and on AMD Opteron, and a factor of 3.3 on Marenostrum (the IBM PowerPC is very sensitive to cache misses)

Results for load balancing: instructions

Analysis of parallel execution performed with Prof. Jesús Labarta in Barcelona (Spain) using his Paraver software package

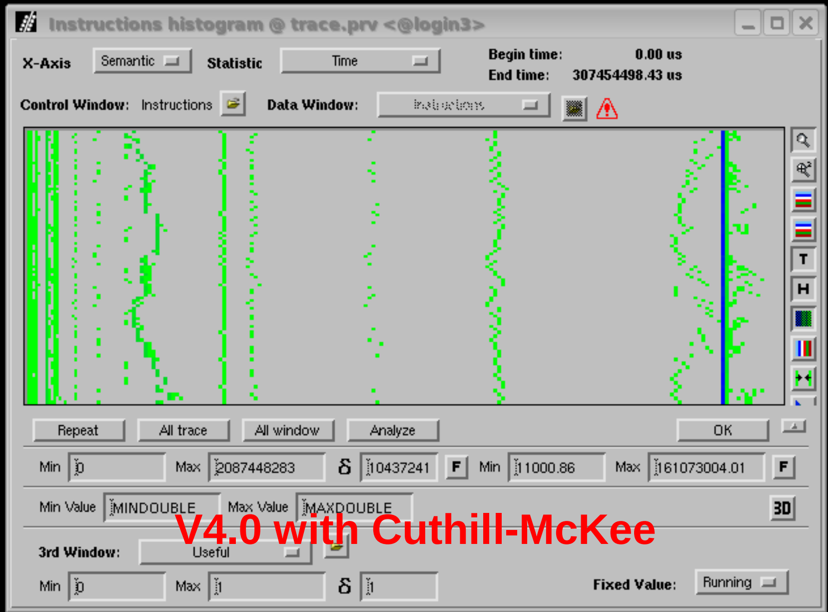

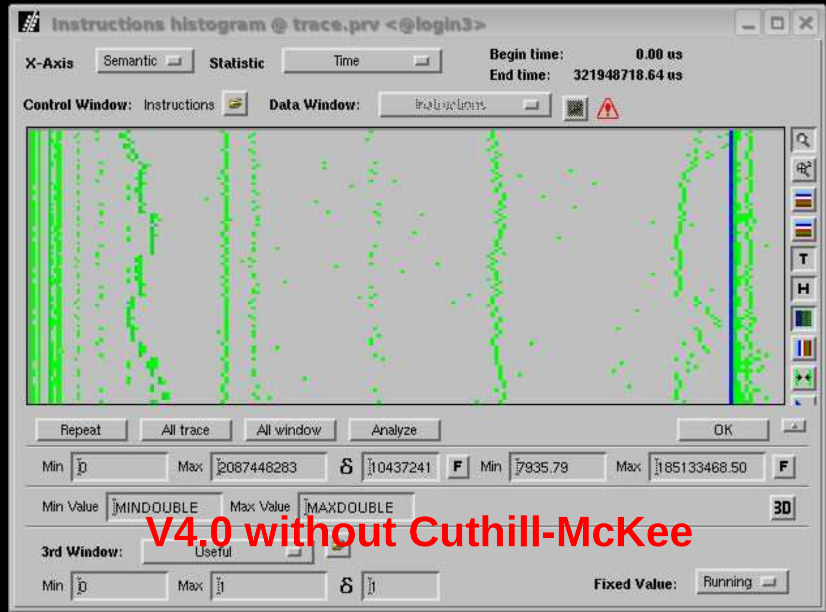

- Number of instructions executed in each slice is well balanced

- Cuthill-McKee has almost no effect on that because we use high-order finite elements (of Q4 type), each of them fits in the L1 cache and for any such element we perform a very large number of operations using data that is already in L1

Results for load balancing: cache misses

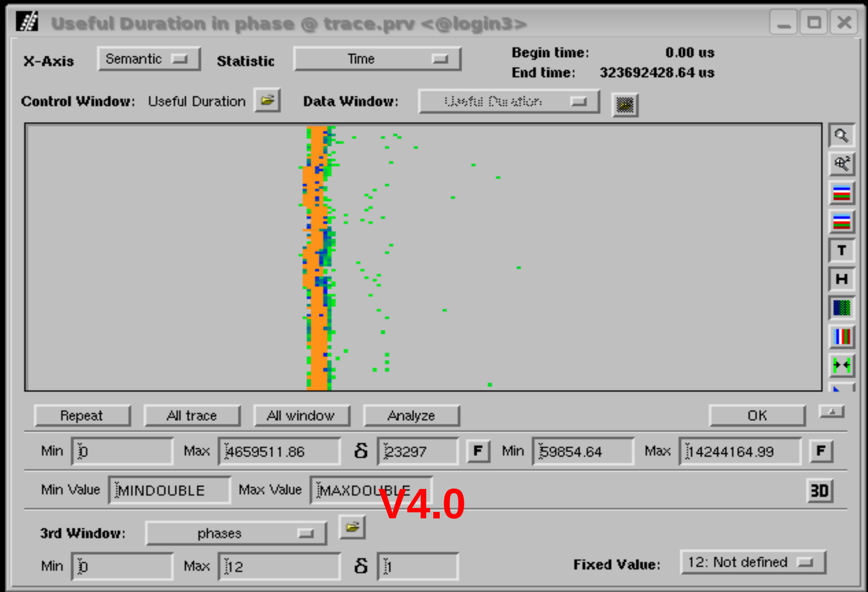

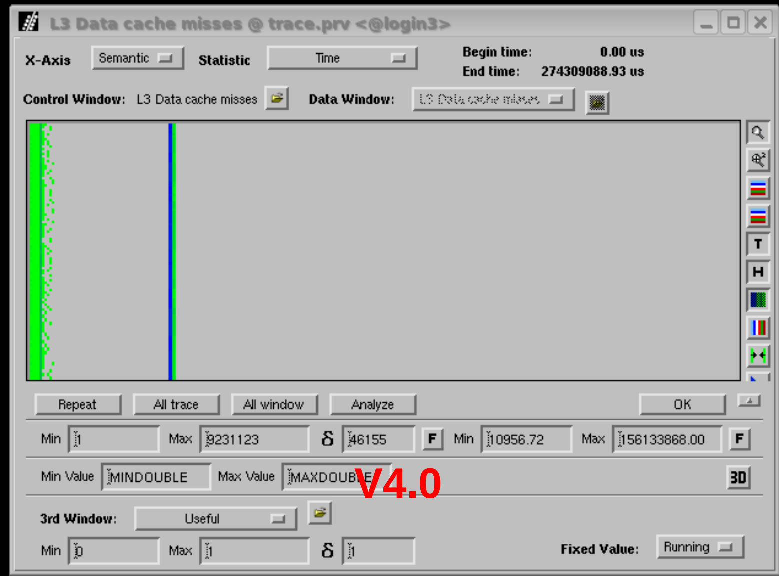

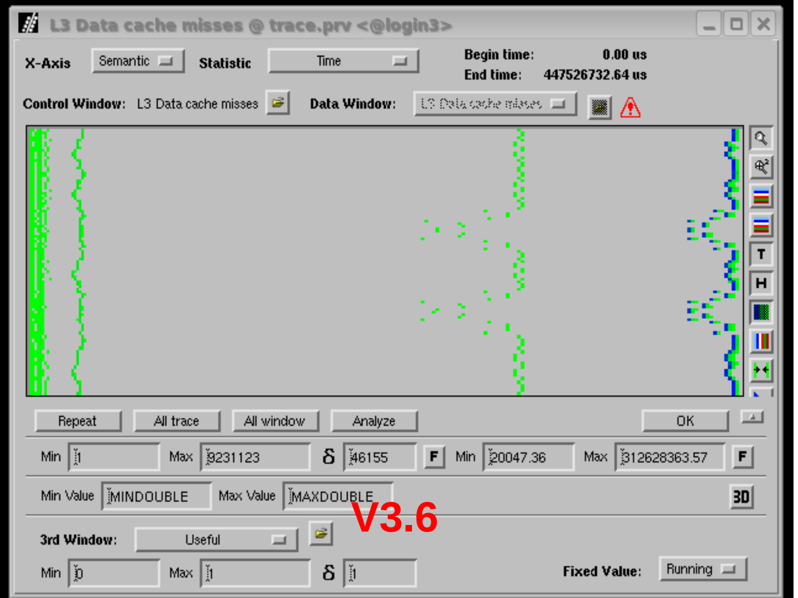

After adding Cuthill-McKee sorting, global addressing renumbering and loop reordering we get a perfectly straight line for cache misses, i.e. same behavior in all the slices and also almost perfect load balancing.

The total number of cache misses is also much lower than in v3.6

CPU time (in orange) is also almost perfectly aligned

BLAS 3
(Basic Linear Algebra Subroutines)

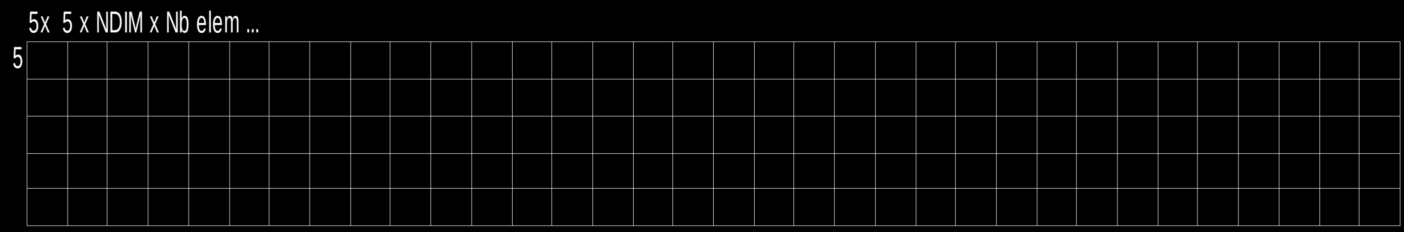

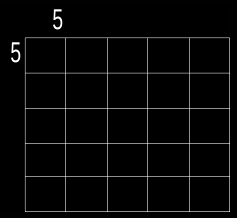

Collaboration with Nicolas Le Goff (Univ of Pau, France)

Can we use highly optimized BLAS matrix matrix products (90% of computations)?

- For one element: matrices (5x25, 25x5, 5 x matrices of (5x5)), BLAS is not efficient: overhead is too expensive for matrices smaller than 20 to 30 square.
- If we build big matrices by appending several elements, we have to build 3 matrices, each having a main direction (x,y,z), which causes a lot of cache misses due to the global access because the elements are taken in different orders, thus destroying spatial locality.
- Since all arrays are static, the compiler already produces a very well optimized code.

=> No need to, and cannot easily use BLAS

=> Compiler already does an excellent job for small static loops

A very large run for PKP phases at 2 seconds

- The goal is to compute differential effects on PKP waves (collaboration with Sébastien Chevrot at OMP Toulouse, France)
- Very high resolution needed (2 to 3 seconds typically)
- Too big for current big machines (pangu: 4000 processors, Marenostrum 10000 processors) therefore we convert the mantle to acoustic instead of elastic and remove the crust because we are only interested in differential effects on P phases
- We keep an elastic anisotropic medium in the inner core only
- We can then design a mesh that is accurate down to periods of 2 seconds for P waves and that fits on 2166 processors (6 blocks of N x N slices, with N = 19)

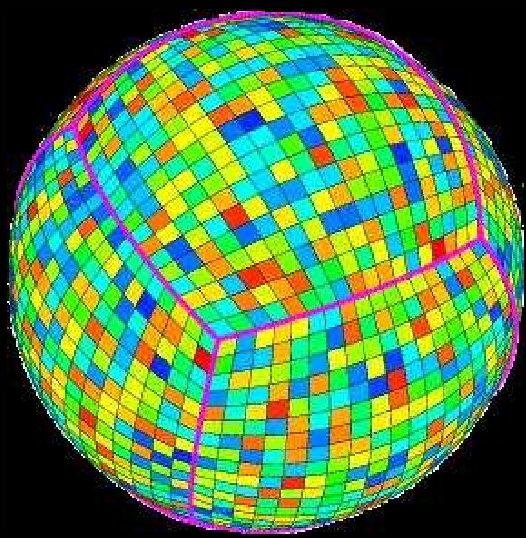

- The mesh contains 126 billion points (the "equivalent" of a 5000 x 5000 x 5000 grid); We use 50000 time steps in 60 hours of CPU on 2166 processors on MareNostrum in Barcelona. Total memory used is 3.5 terabytes.
- Calculations with v4.0 (acoustic mantle, no crust) are finished, the code performed well and performance levels were very satisfactory; the analysis of the seismograms is under way

Dec 26, 2004 Sumatra event

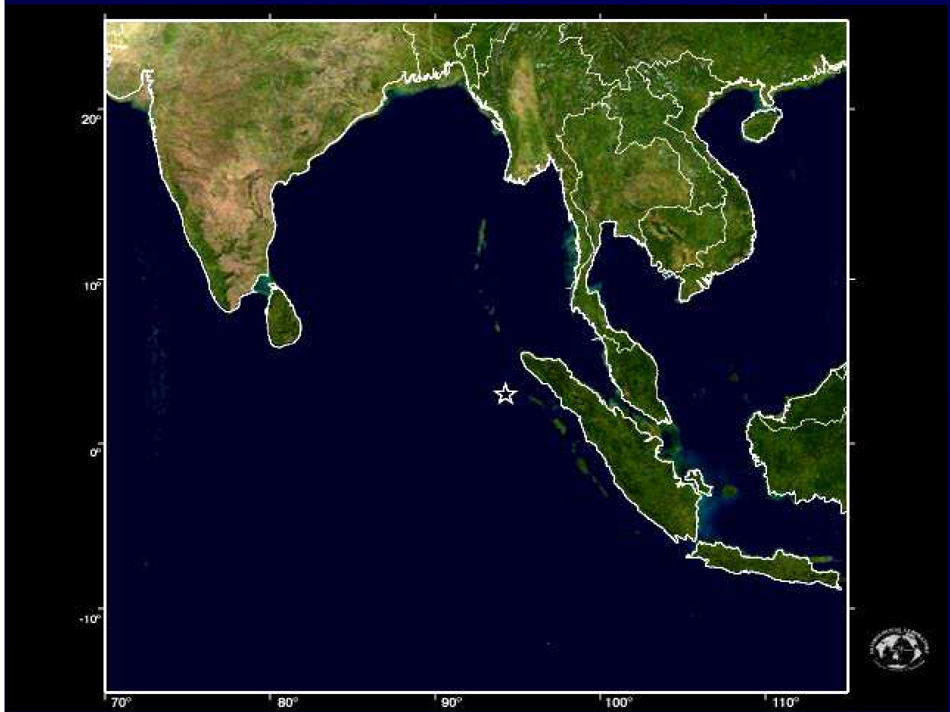

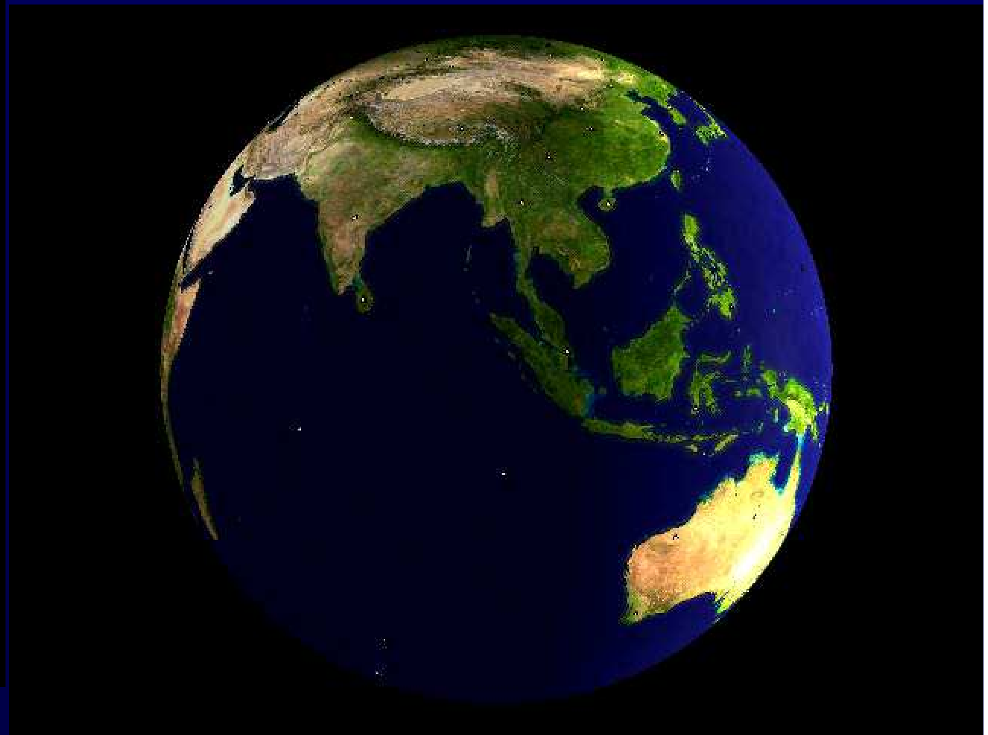

From Tromp et al., 2005

vertical component of velocity at periods of 10 s and longer on a regional scale

Hollywood Earthquake

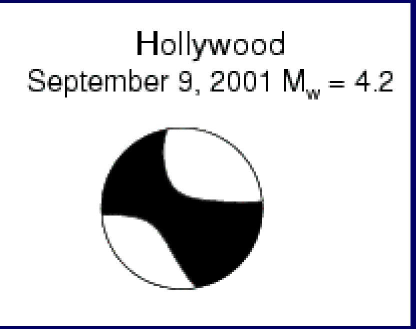

Small M 4.2 earthquake on Sept 9, 2001

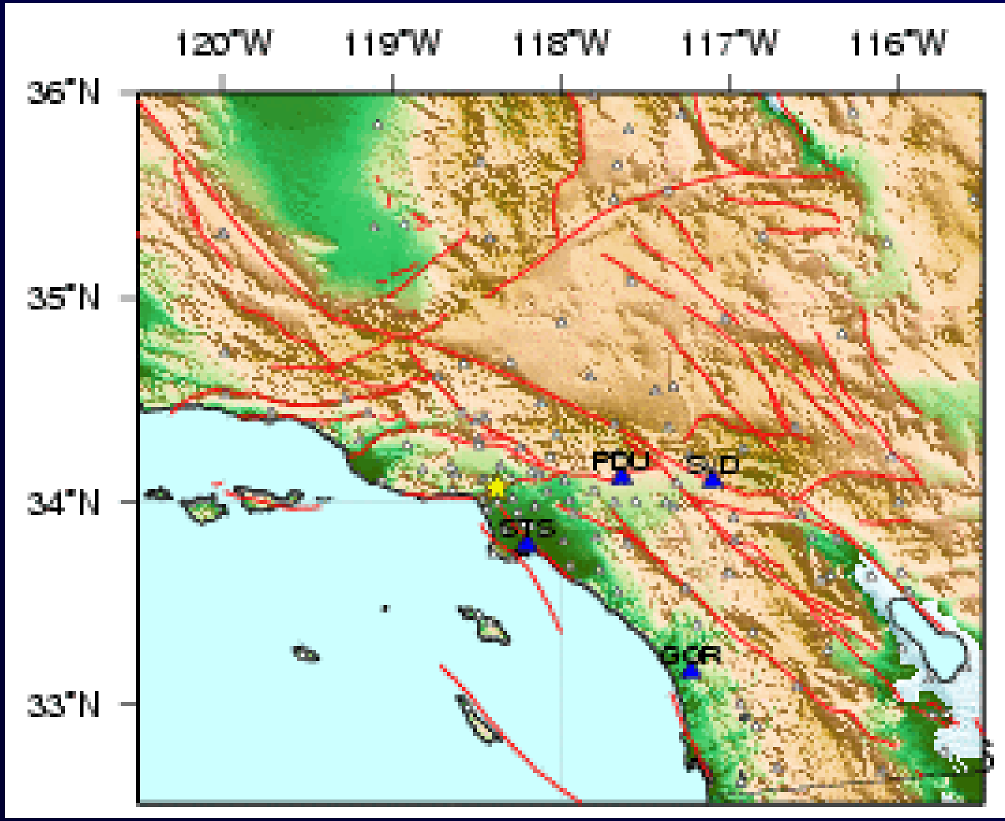

Amplification in basin

San Andreas – January 9, 1857

Vertical scale approximately 1 km

Carrizo Plain, San Andreas Fault, California, USA

Earthquakes at the regional scale

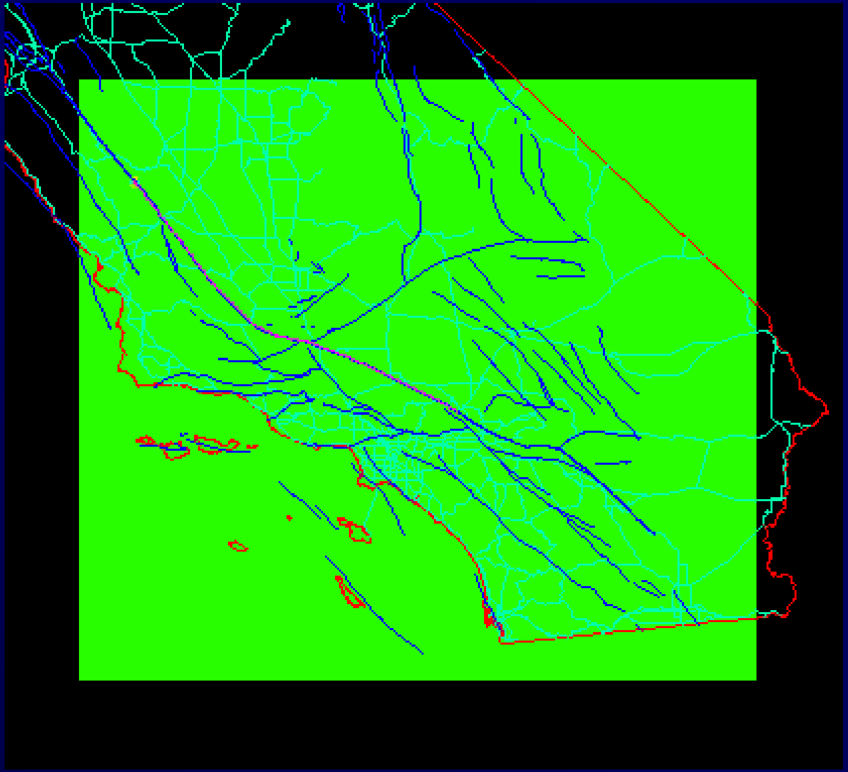

Carrizo Plain, USA, horizontal scale ≅ 200 m

Scale approximately 500 km

3D spectral-element method (SEM)

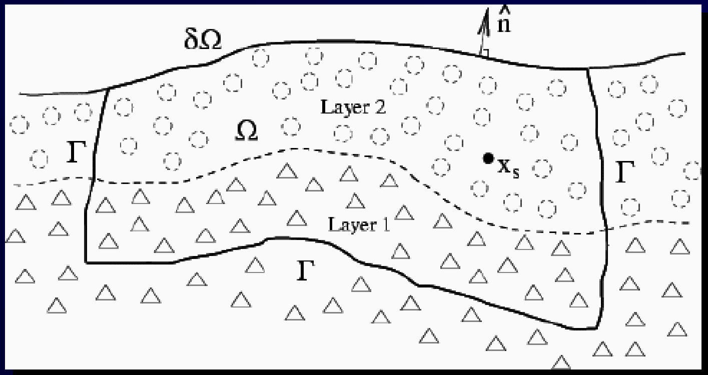

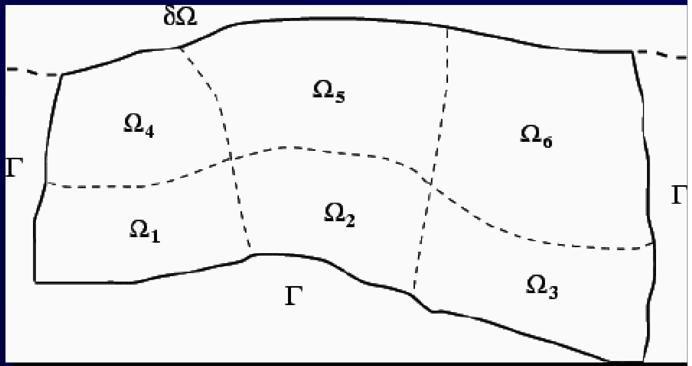

Absorbing conditions

- Used to be a big problem
- Bérenger 1994
- INRIA (Collino, Cohen)
- Extended to second-order systems by Komatitsch and Tromp (2003)

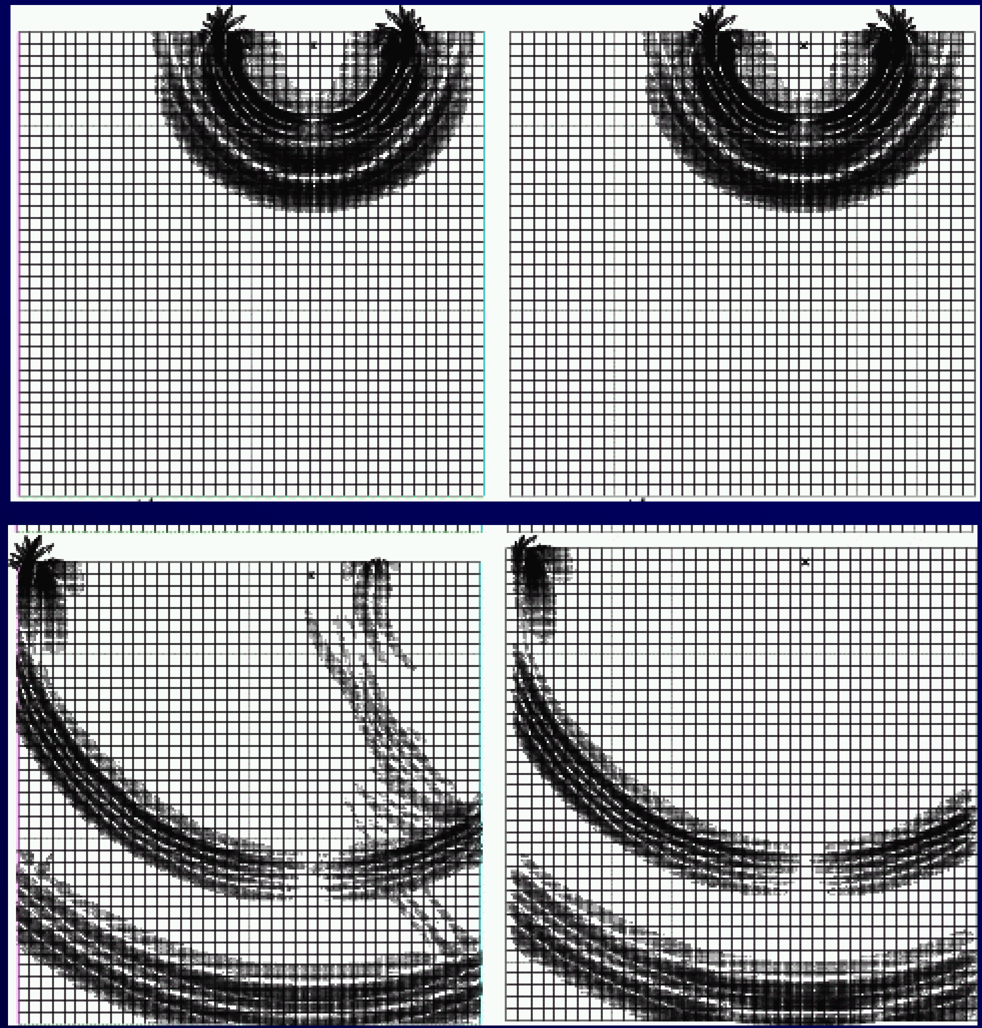

PML (Perfectly Matched Layer)

Convolution-PML in 3D for seismic waves

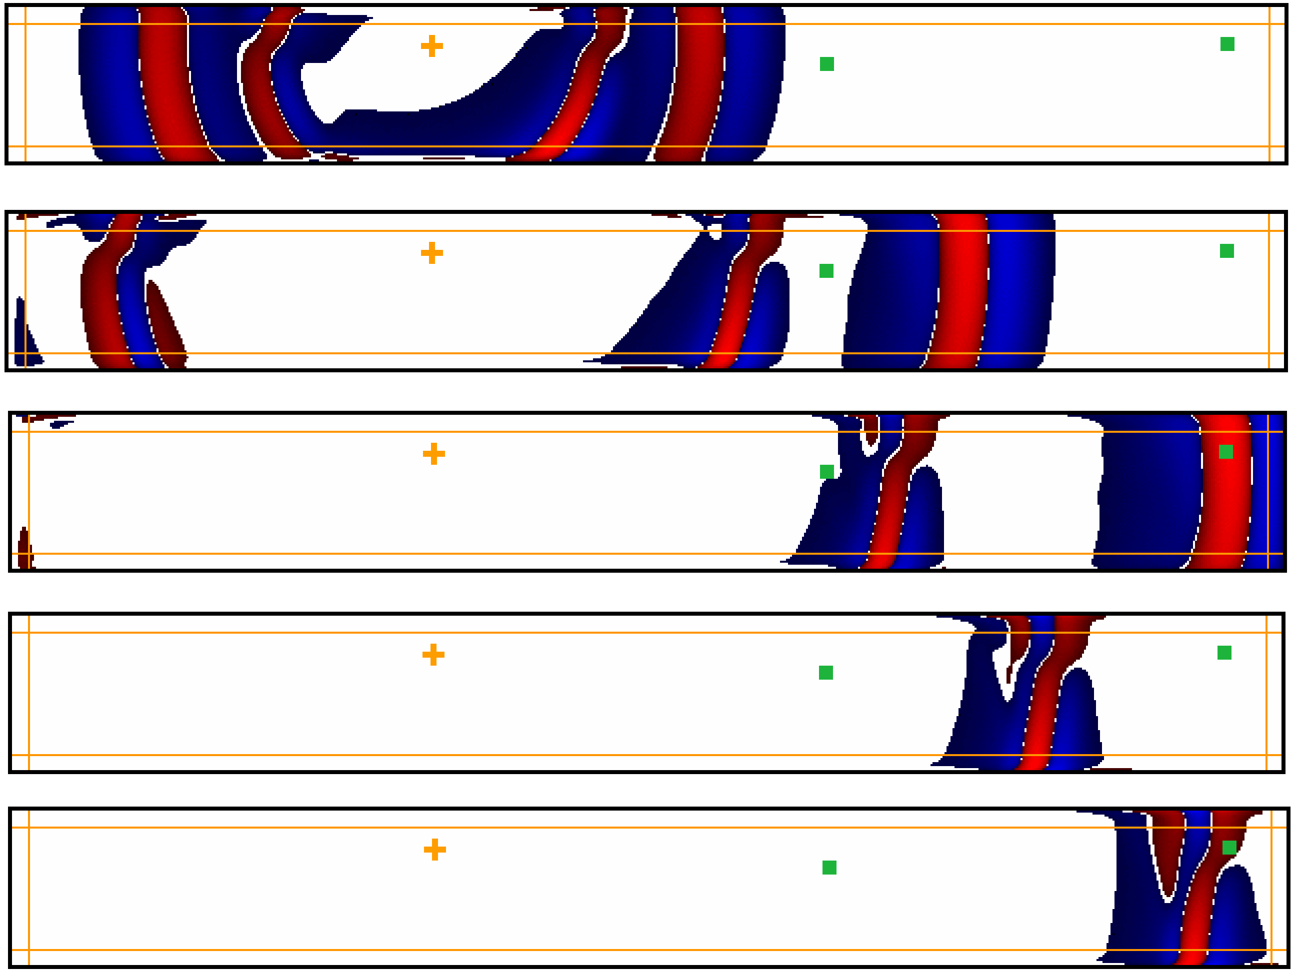

Finite-difference technique in velocity and stress:
staggered grid of Madariaga (1976), Virieux (1986)

- **Optimized for grazing incidence**
- ***Not* split**
- **Use recursive convolution based on memory variables (Luebbers and Hunsberger 1992)**
- **« 3D at the cost of 2D »**

Future work

- ANR NUMASIS (2006-2009): optimize SPECFEM3D (among other codes) on NUMA machines (e.g. CEA Bull Tera10)
- Inverse problems (already done for the source by Liu et al 2004, but not for the model yet)
- Operto et al. for finite-differences $\Rightarrow$ SEM in frequency? but full stiffness matrix $\Rightarrow$ use MUMPS (ANR Solstice project)
- Altivec, CELL SuperScalar (Barcelona), GPGPU for the SPECFEM3D kernel (ENSEEIHT internship scheduled this year)
- Out-of-core for large-scale problems (talk by Jennifer Scott this morning, Abdou Guermouche in MUMPS etc)
- Grid computing (successful attempts on Egée): talk by Ronan Guivarch