Application of the MUMPS direct solver

to 3D acoustic wave modeling and seismic imaging

SEISCOPE research group

Email: operto@geoazur.obs-vlfr.fr

Web: http://seiscope.oca.eu

Géoazur - CNRS - IRD - UNSA – OCA, Sophia-Antipolis.

LGIT – CNRS – IRD – UJF, Grenoble.

SEISCOPE project

http://seiscope.oca.eu

• **Main scientific interests:** numerical wave modeling in 2D and 3D elastic/acoustic isotropic/anisotropic media; seismic imaging by full waveform inversion; applications at different scales (near surface, oil exploration, crustal and lithospheric scales from artifical sources and natural sources (i.e., earthquakes).

•**3 permanent people:**

S. Operto (Geoazur), A. Ribodetti (Geoazur), J. Virieux (LGIT)

• **Former and current PhD students:**

H. Ben Hadj Ali (now at Total), R. Brossier (now post-doc at LGIT), V. Etienne (PhD, Geoazur), Y. Gholami (PhD, Geoazur)), G. Hu (PhD, LGIT), Y. Jia (PhD, LGIT), D. Pageot (PhD, Geoazur), V. Prieux (PhD, Geoazur), F. Sourbier (IR, now at CEA)

• **Collaborations** on numerical analysis and computing: MUMPS team, L. Giraud (INRIA Bordeaux), A. Haidar (CERFACS), S. Gratton (CERFACS).

• Industrial and academic fundings: SEISCOPE petroleum consortium (http://seiscope.oca.eu) and ANR

CONTENT

Frequency-domain seismic modeling with MUMPS

- **Introduction: seismic exploration, seismic modeling and inversion**
- **The 3D time-harmonic acoustic wave equation**
- **Finite-difference discretization**
- **Resolution with the MUMPS direct solver**
- **Resolution with a hybrid direct/iterative solver**
- **Conclusion**

Seismic imaging at different scales

Civil engineering

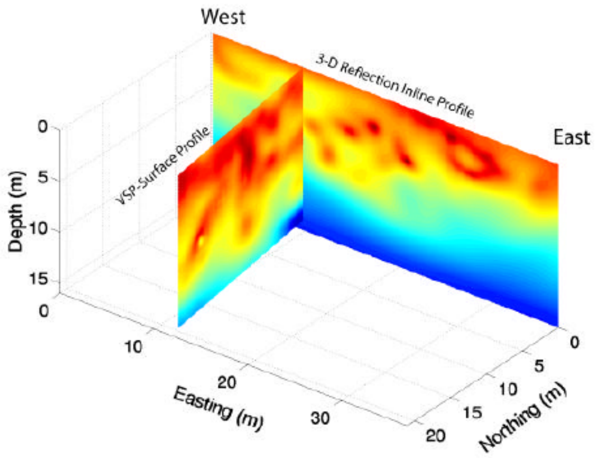

Onshore oil&gas exploration

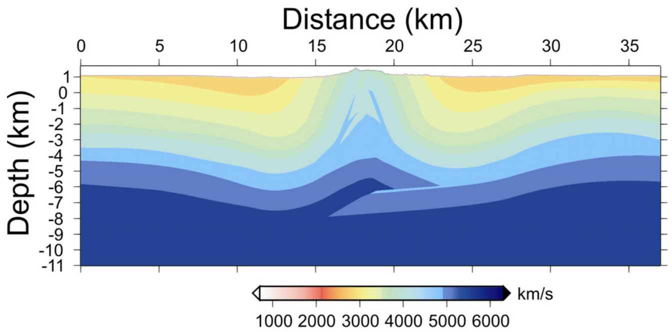

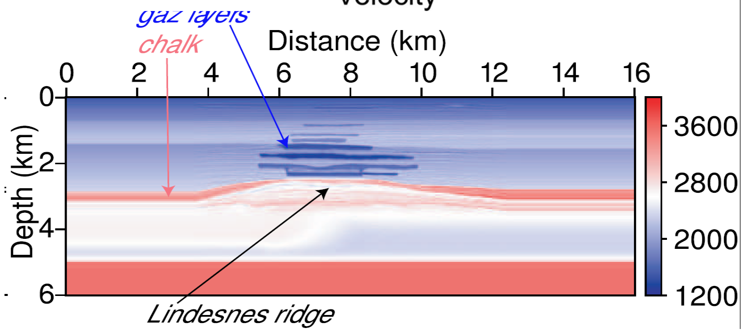

Academic applications

at crustal/lithospheric scales…

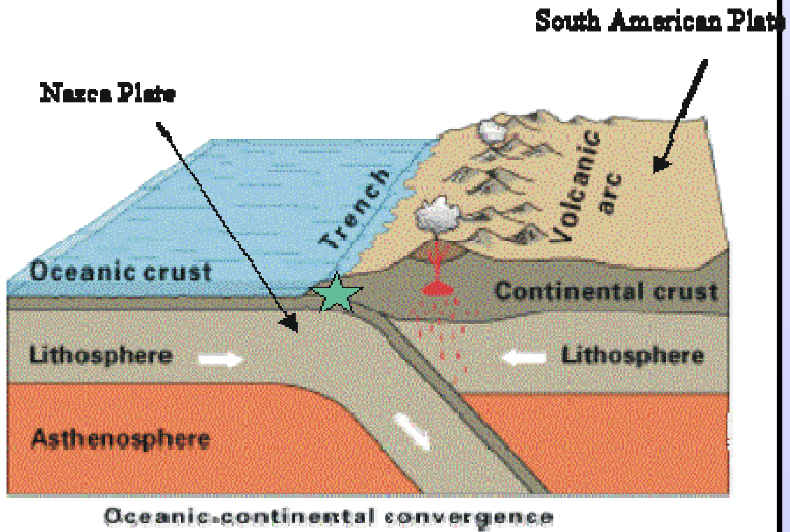

and at the global scale

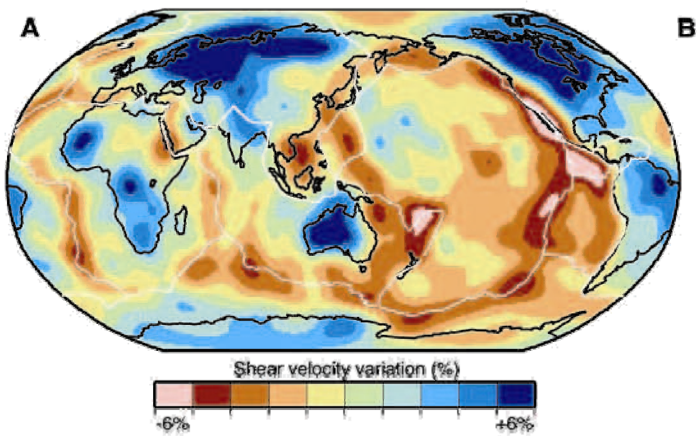

Principle of seismic exploration

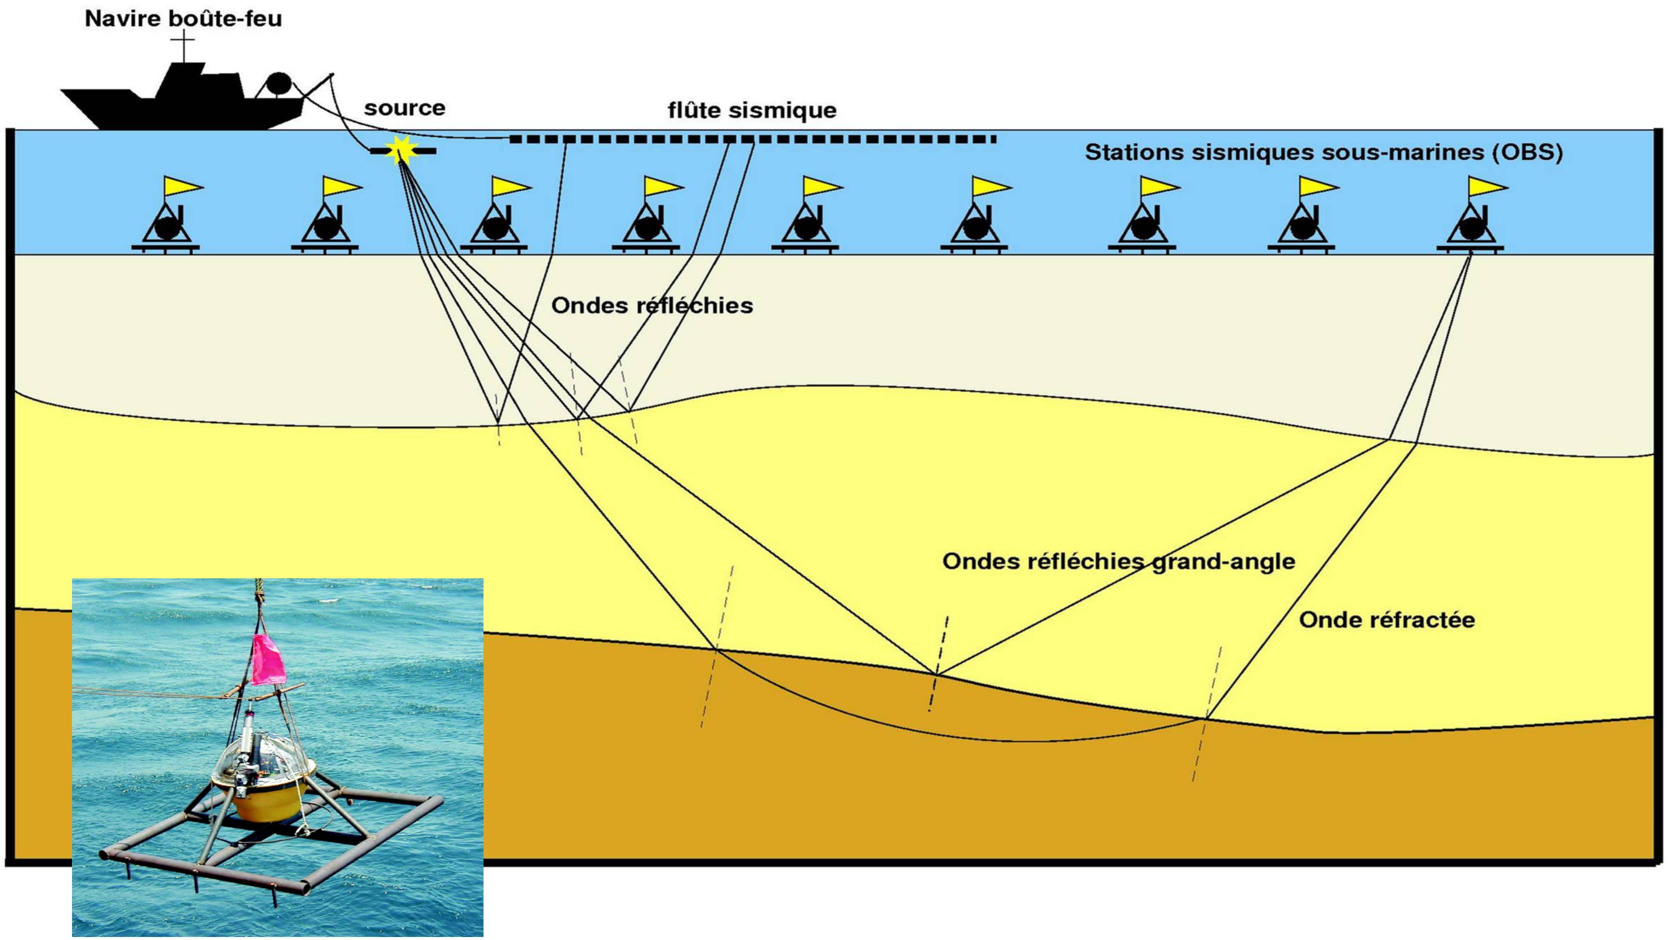

Ocean Bottom Seismometer

Example of a 3D OBC acquisition for oil exploration

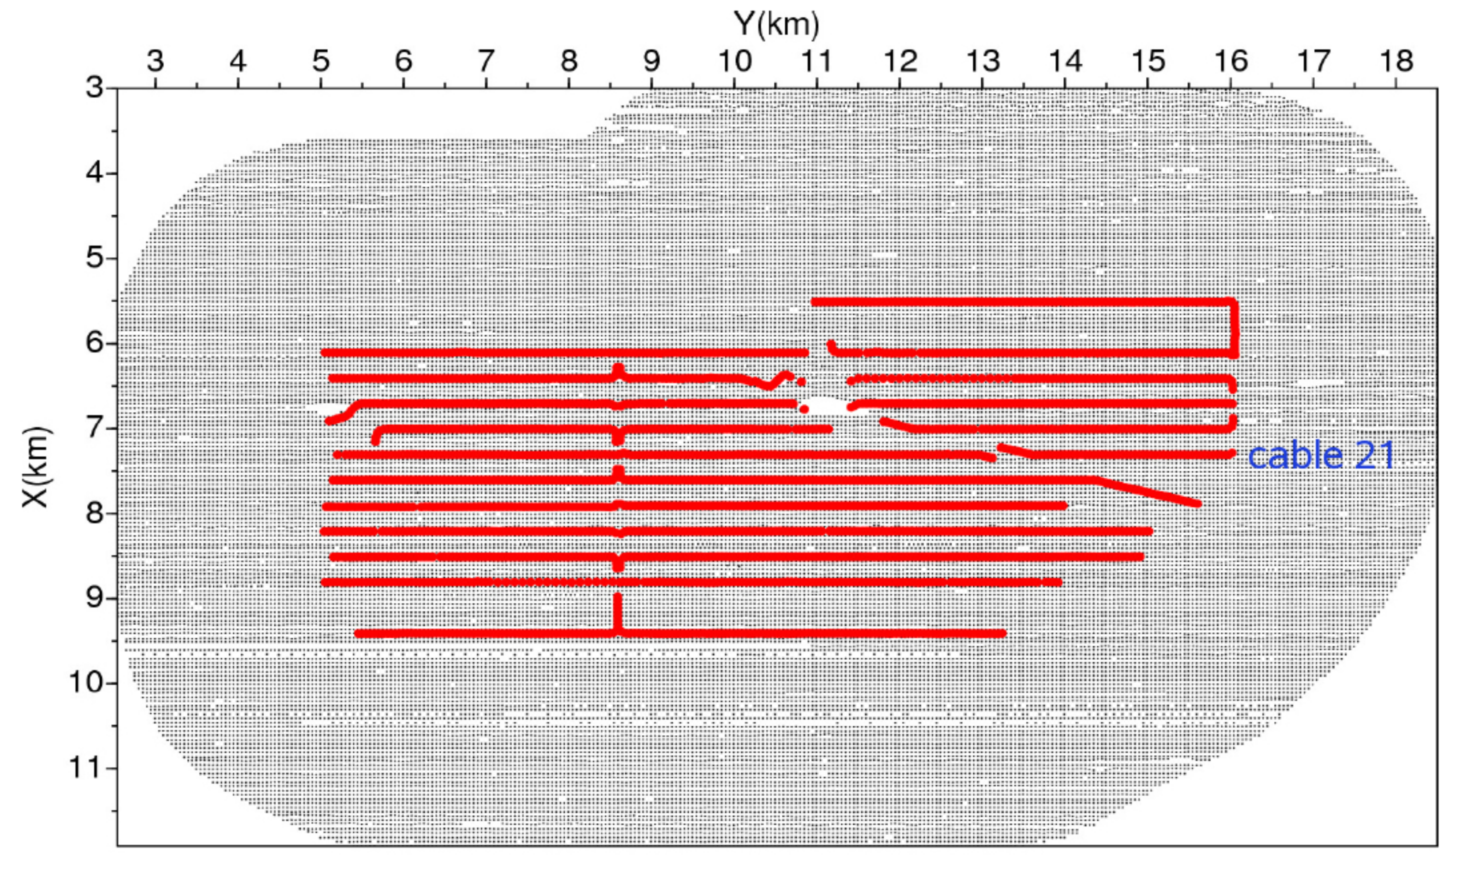

49 954 shots – 12 ocean bottom cables – 2304 receivers

With the courtesy of L. Sirgue (BP, now at Total)

Example of common receiver gather

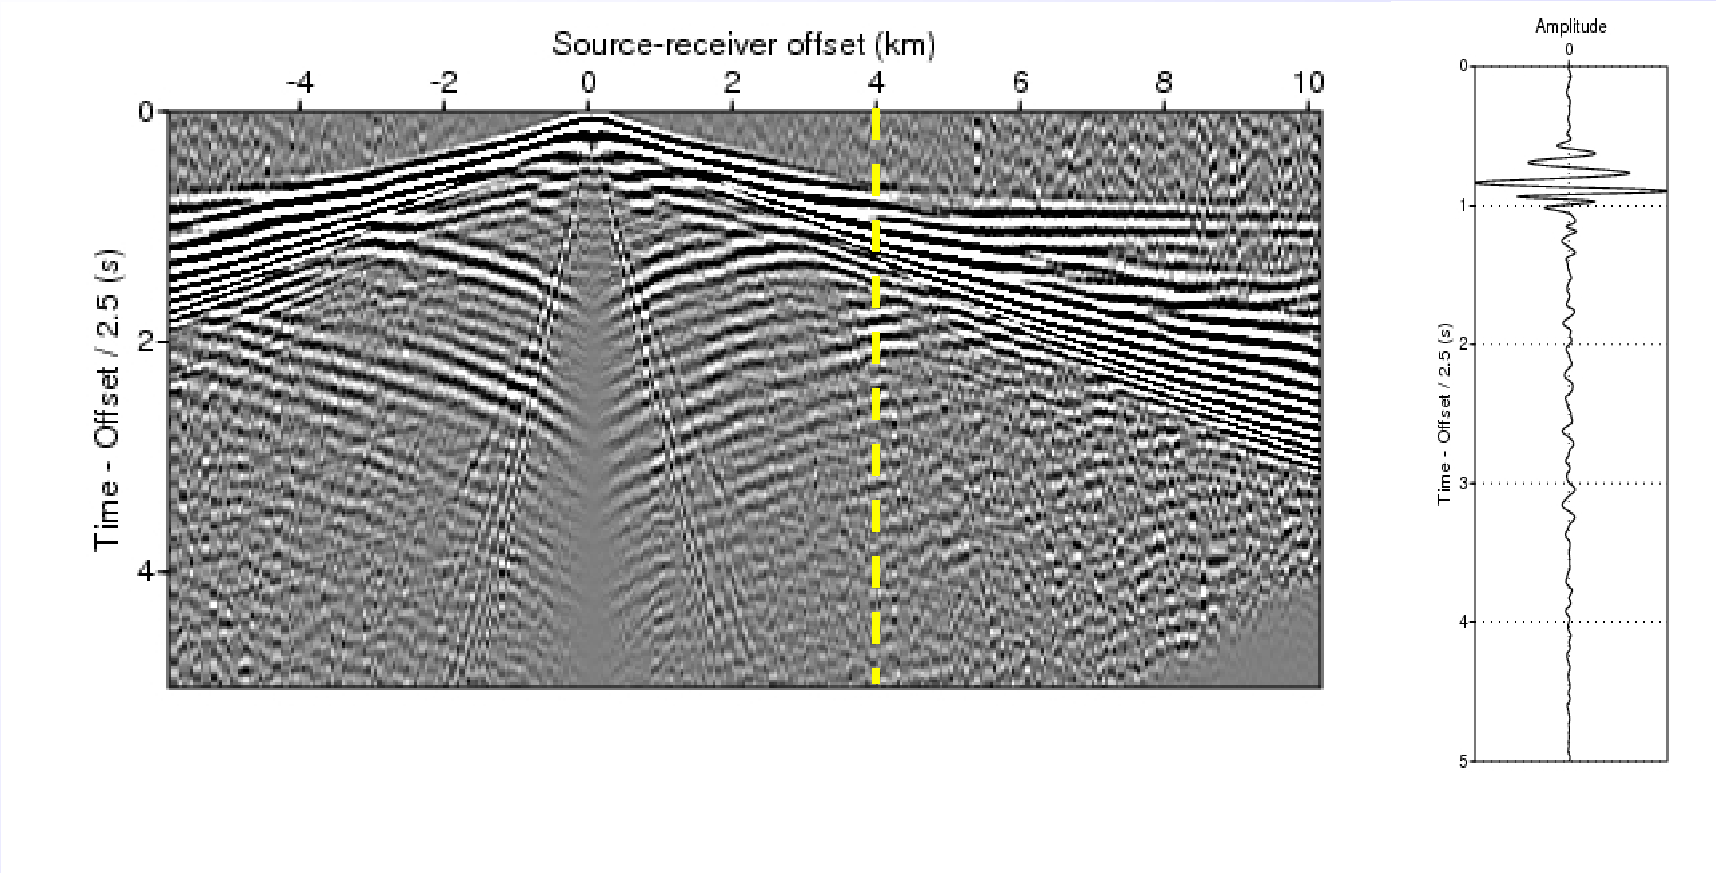

Seismic imaging by full waveform inversion

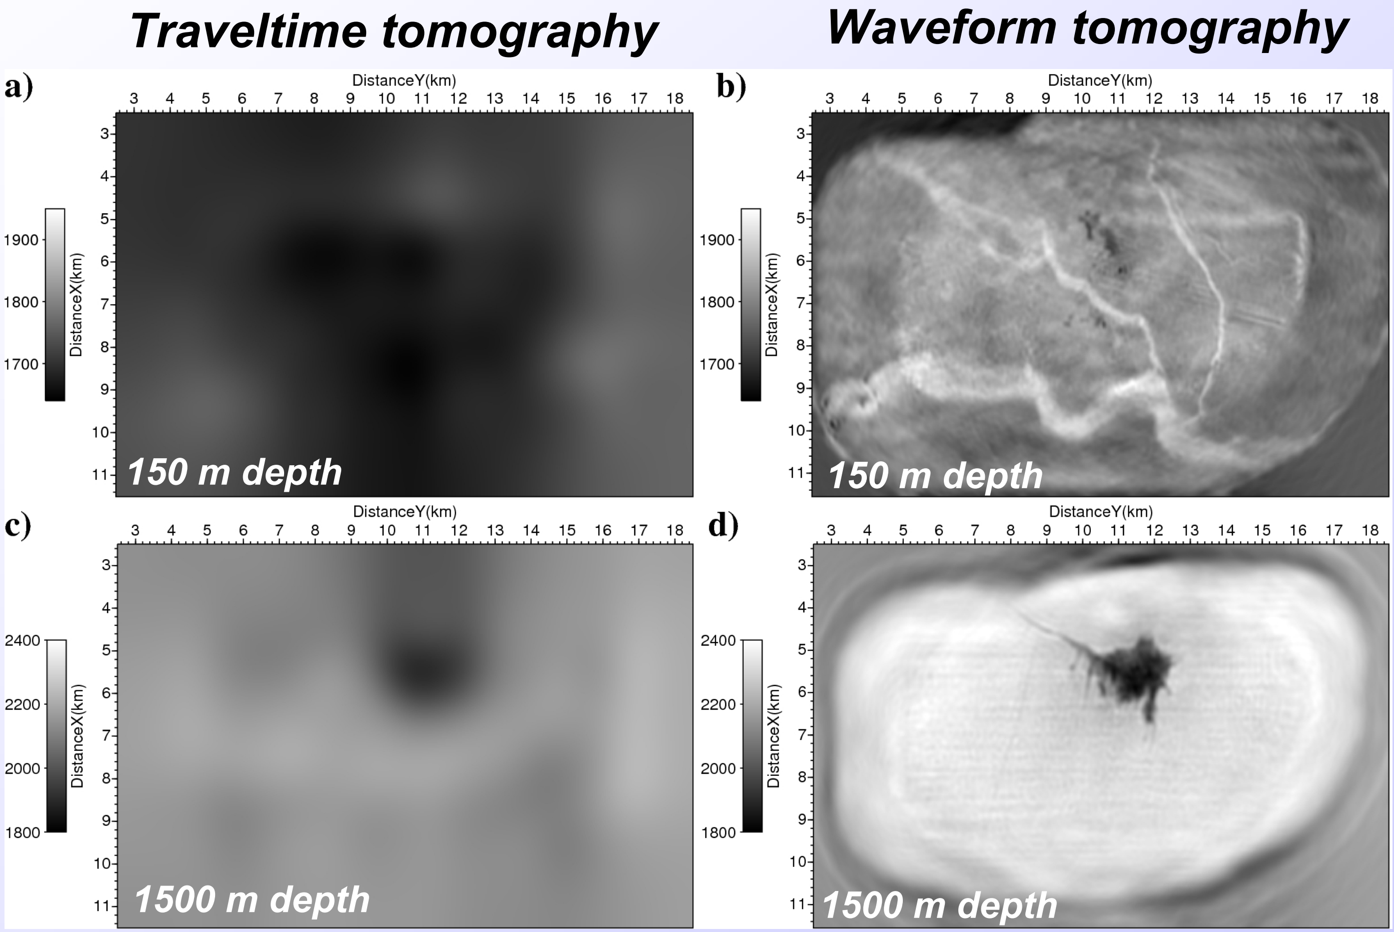

From Sirgue et al., First Break, 2010

Forward problem

Time versus frequency-domain seismic modeling

Time-domain time-marching explicit schemes

1st-order PDEs

(velocity-stress)

$$\mathbf{M}(\mathbf{x})\frac{d\mathbf{w}(\mathbf{x},t)}{dt} = \mathbf{A_1}(\mathbf{x})\mathbf{w}(\mathbf{x},t) + \mathbf{s}(\mathbf{x},t)$$

2nd-order PDEs

(velocity)

$$\mathbf{M}(\mathbf{x})\frac{d^2\mathbf{v}(\mathbf{x},t)}{dt^2} = \mathbf{A_2}(\mathbf{x})\mathbf{v}(\mathbf{x},t) + \mathbf{s}'(\mathbf{x},t)$$

Frequency-domain implicit schemes

$$\mathbf{B}(\mathbf{x},\omega)\mathbf{v}(\mathbf{x},\omega) = \mathbf{s}(\mathbf{x},\omega),$$

Forward problem

Frequency-domain seismic modeling

$$\frac{\omega^2}{\kappa(\mathbf{x})} p(\mathbf{x},\omega) + \nabla\left(\frac{1}{\rho(\mathbf{x})}\nabla p(\mathbf{x},\omega)\right) = -s(\mathbf{x},\omega)$$

$$\left[\frac{\omega^2}{\kappa(x,y,z)} + \frac{1.}{\xi_x(x)}\frac{\partial}{\partial x}\frac{b(x,y,z)}{\xi_x(x)}\frac{\partial}{\partial x} + \frac{1.}{\xi_y(y)}\frac{\partial}{\partial y}\frac{b(x,y,z)}{\xi_y(y)}\frac{\partial}{\partial y} + \frac{1.}{\xi_z(z)}\frac{\partial}{\partial z}\frac{b(x,y,z)}{\xi_z(z)}\frac{\partial}{\partial z}\right] p(x,y,z,\omega) = -s(x,y,z,\omega)$$

Why frequency-domain modeling?

Efficiency of multi-rhs resolution

$$\mathbf{B}(\mathbf{x},\omega)\mathbf{v}(\mathbf{x},\omega) = \mathbf{s}(\mathbf{x},\omega),$$

$$\mathbf{B}\left[\mathbf{v_1 v_2 ... v_N}\right] = \left[\mathbf{s_1 s_2 ... s_N}\right]$$

$$\mathbf{L.U}\left[\mathbf{v_1 v_2 ... v_N}\right] = \left[\mathbf{s_1 s_2 ... s_N}\right]$$

Straightforward implementation of attenuation effects

$$c(\omega) = c_R(\omega) + i c_I(\omega)$$

$$Q = -c_R(\omega)/2c_I(\omega)$$

Possible approaches for frequency-domain solutions

	Direct solver (DSM)	Hybrid solver (DDM)	Iterative solver	Time-domain approach (TDM)
Memory complexity	--	-	++	+
Multi-r.h.s time complexity	++	-	--	--
Scalability	--	+	+	++
Robustness	+	-	--	++

Complexities of DSM modeling

Dimension	Memory complexity	Time complexity
2D	$\mathcal{O}(N^2)Log_2(N)$	$\mathcal{O}(N^3)$
3D	$\mathcal{O}(N^4)$	$\mathcal{O}(N^6)$

Complexities of TDM modeling

2D	$O(N^2)$	$O(N_t\ N^2\ N_{rhs}) \sim O(N^4)$
3D	$O(N^3)$	$O(N^3\ N_{rhs}\ N_t) \sim O(N^6)$

Least-squares local optimization

Data misfit vector $\quad \Delta \mathbf{d} = \mathbf{d}_{obs} - \mathbf{d}_{cal}(\mathbf{m})$ with **d** = $\mathcal{R}$ **v** and $\mathcal{R}$ is a restriction operator

Local search in the vicinity of m$_0$ $\quad \mathbf{m} = \mathbf{m_0} + \Delta \mathbf{m}.$

Cost function $\quad \mathcal{C}(\mathbf{m}) = \frac{1}{2} \Delta \mathbf{d}^{\dagger} \Delta \mathbf{d},$

$$\Delta \mathbf{m} = - \left[\frac{\partial^2 \mathbf{C}(\mathbf{m_0})}{\partial \mathbf{m^2}} \right]^{-1} \frac{\partial \mathbf{C}(\mathbf{m_0})}{\partial \mathbf{m}}.$$

$$\nabla C_l = \underbrace{\sum_{i=1}^{N\omega}}_{\text{New LU facto.}} \underbrace{\sum_{j=1}^{Ns}}_{\text{New solve}} \Re \left[[\mathbf{B_i^{-1} s_j}]^{\mathbf{t}} \left[\frac{\partial \mathbf{B_i}}{\partial \mathbf{m_l}} \right]^{\mathbf{t}} [\mathbf{B_i^{-1}} \left(\mathcal{P} \Delta \mathbf{d}^*_{\mathbf{i,j}} \right)] \right]$$

New LU facto. **New solve**

Forward problem

$$\mathbf{B}(\mathbf{x}, \omega) \mathbf{v}(\mathbf{x}, \omega) = \mathbf{s}(\mathbf{x}, \omega)$$

Multiscale imaging in the frequency domain

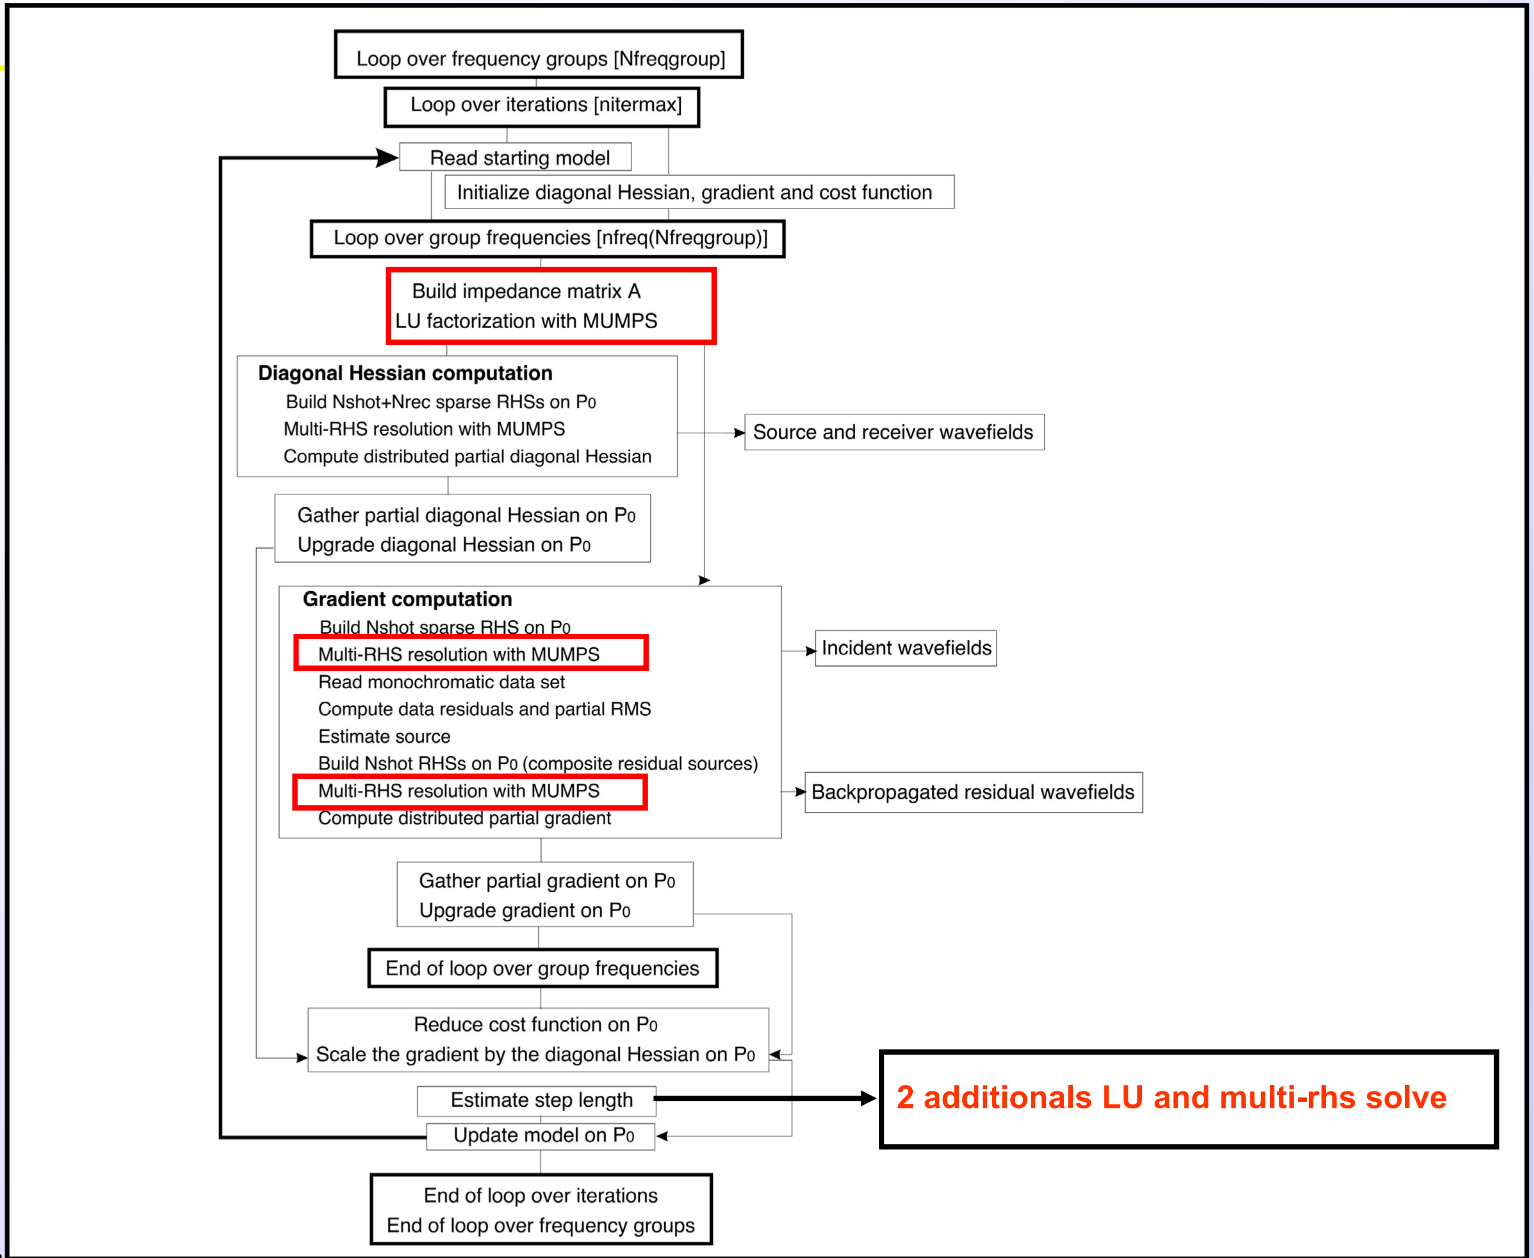

Application to a dip section of the Overthrust model

Efficient multiscale FWI

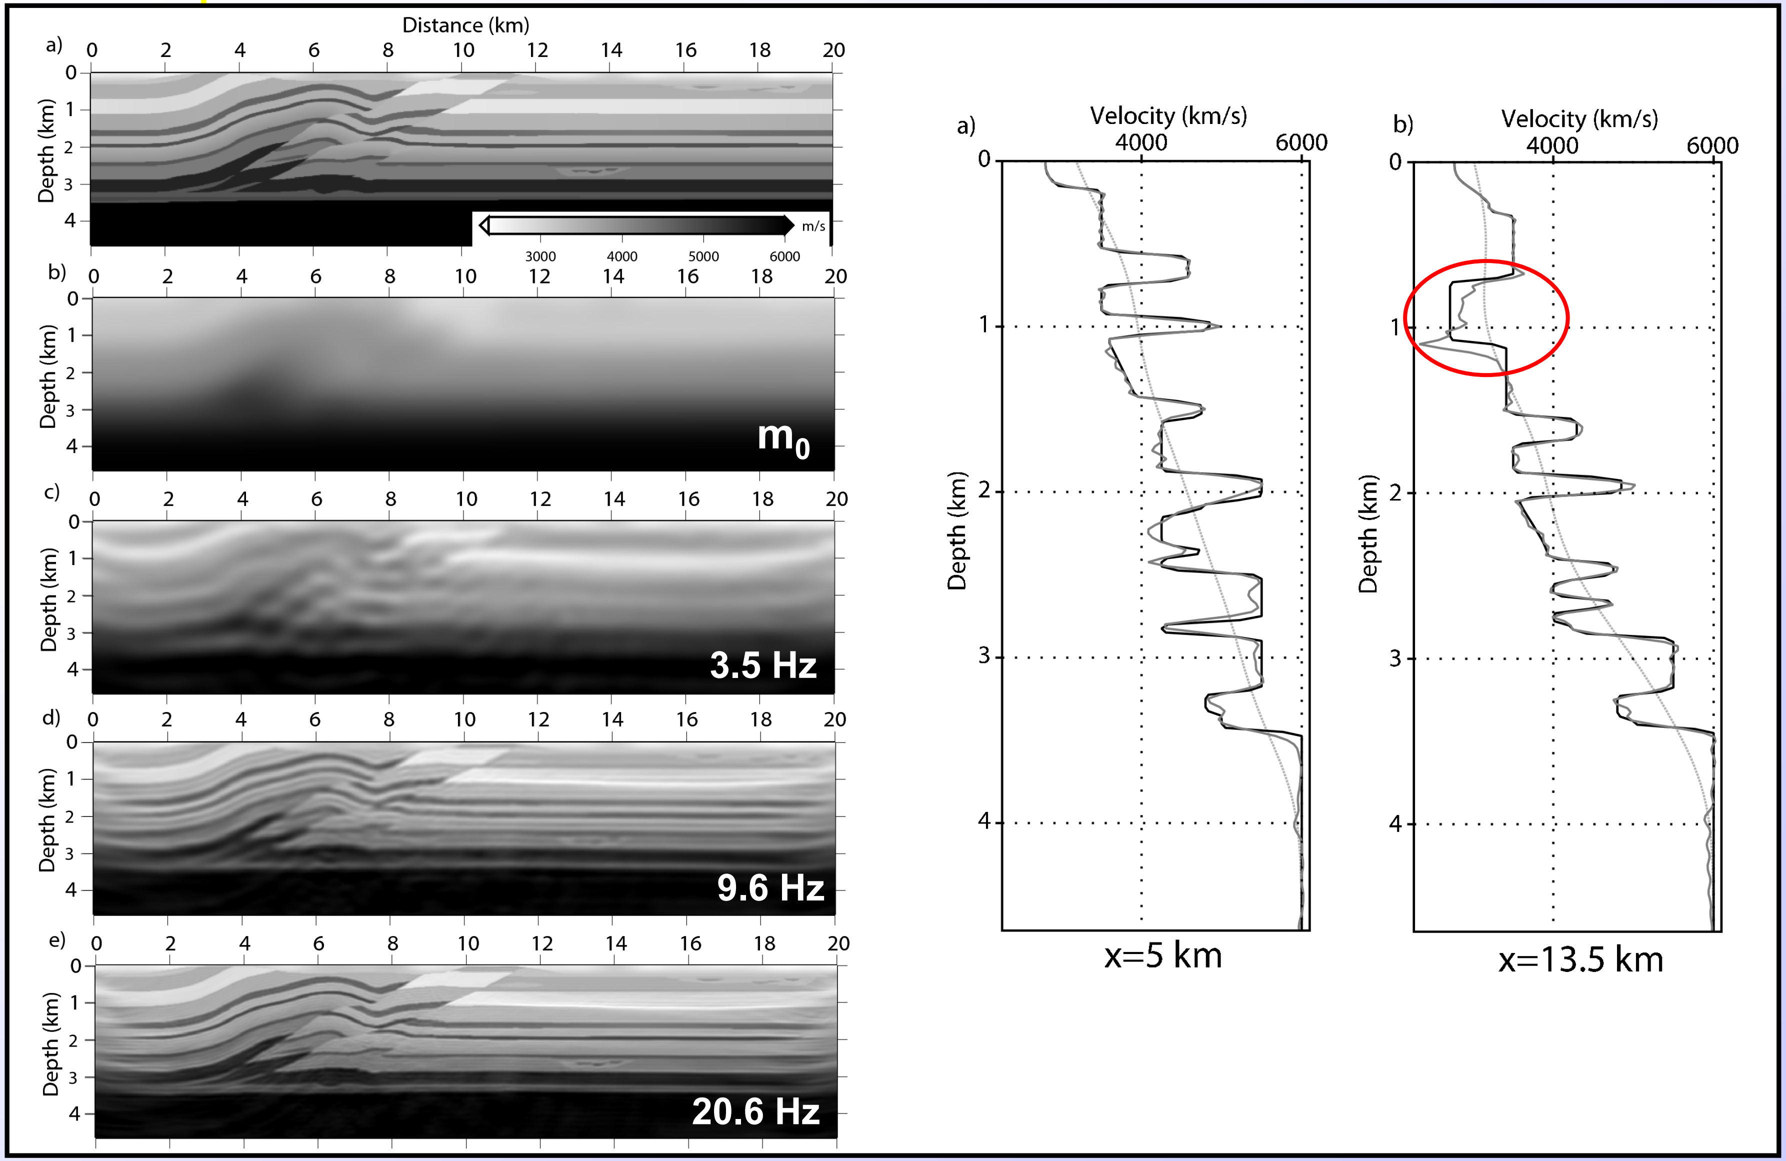

Application to a dip section of the Overthrust model

Bunks multiscale FWI

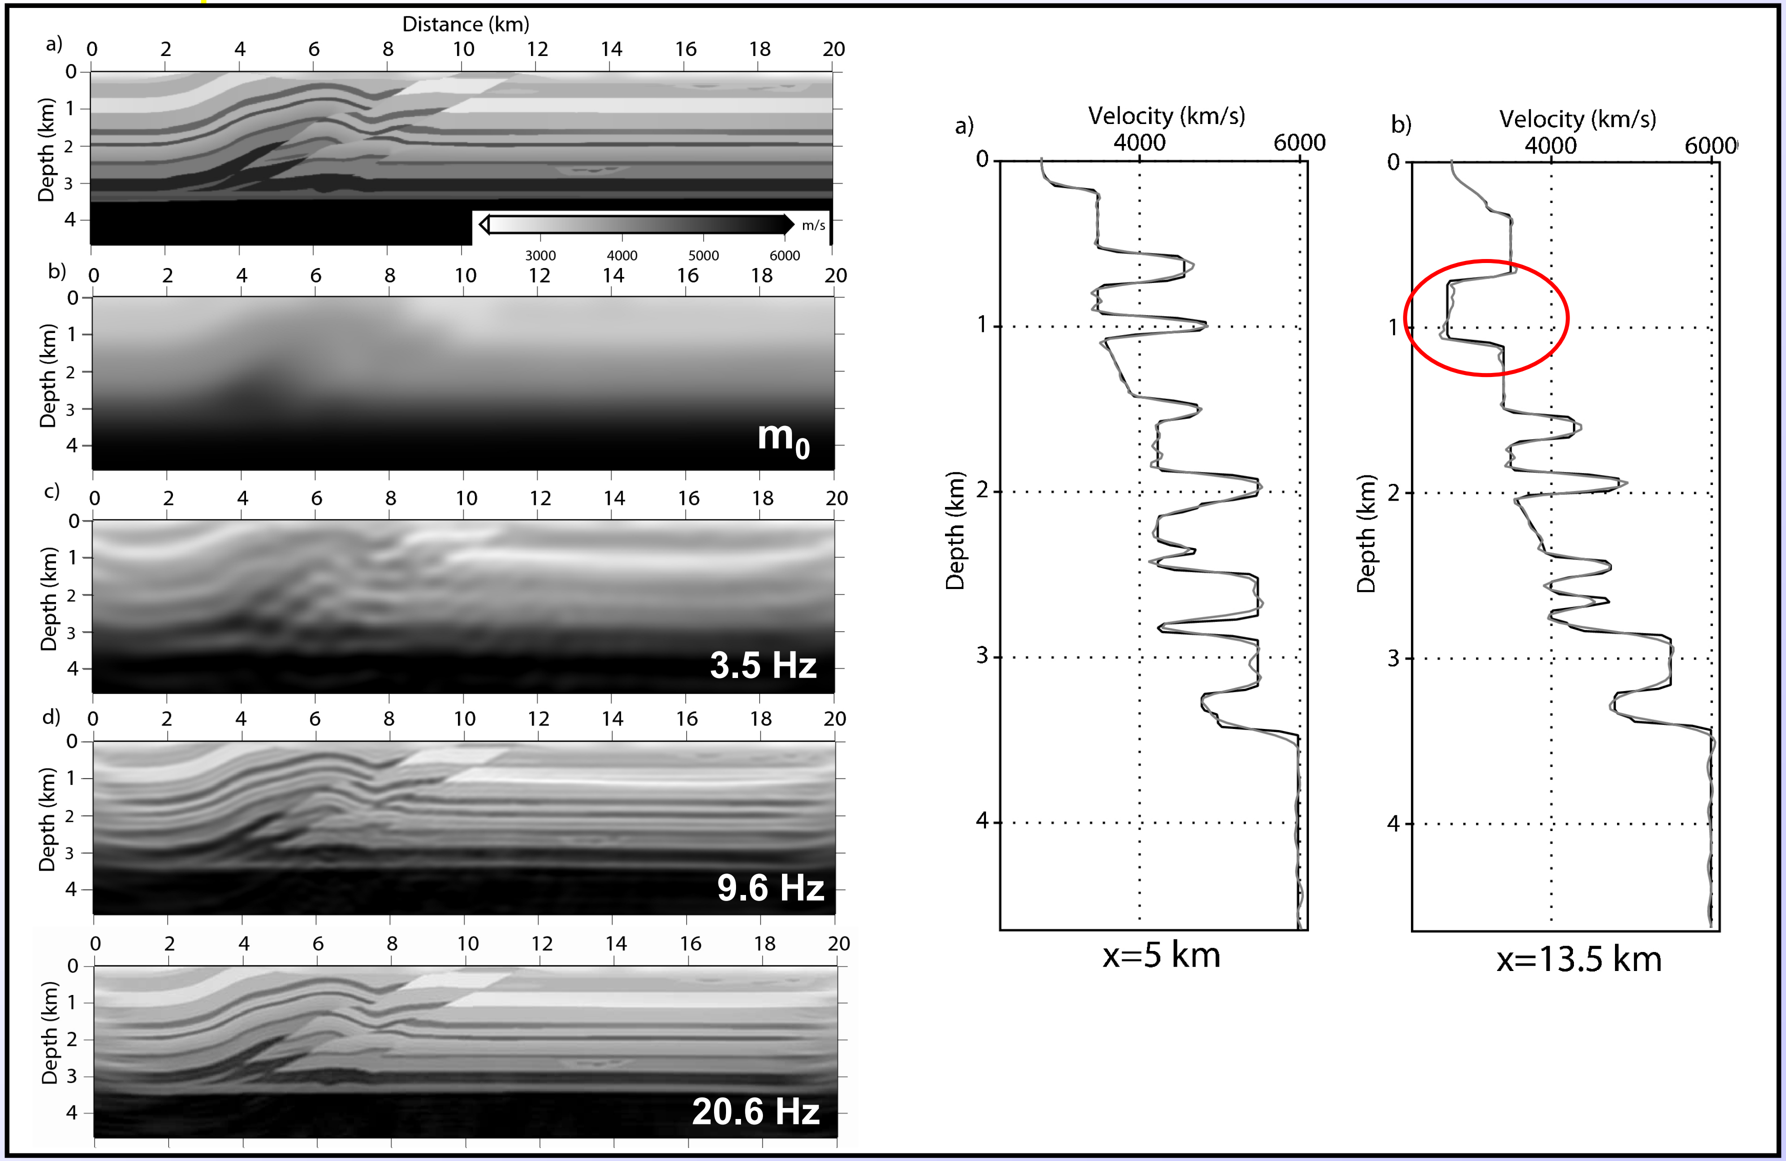

3D frequency-domain seismic modeling with MUMPS

FD discretization

$$\left[\frac{\omega^2}{\kappa(x,y,z)}+\frac{1.}{\xi_x(x)}\frac{\partial}{\partial x}\frac{b(x,y,z)}{\xi_x(x)}\frac{\partial}{\partial x}+\frac{1.}{\xi_y(y)}\frac{\partial}{\partial y}\frac{b(x,y,z)}{\xi_y(y)}\frac{\partial}{\partial y}+\frac{1.}{\xi_z(z)}\frac{\partial}{\partial z}\frac{b(x,y,z)}{\xi_z(z)}\frac{\partial}{\partial z}\right]p(x,y,z,\omega)=-s(x,y,z,\omega)$$

Linear combination of low-order accurate staggered-grid stencils

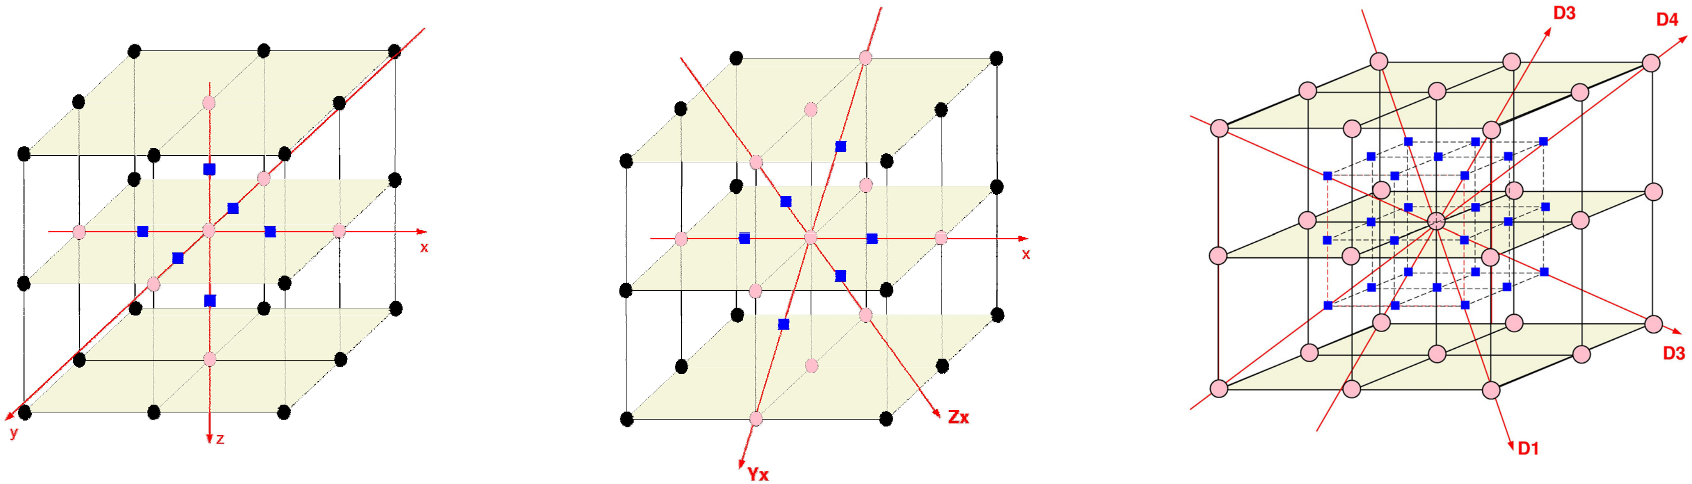

+ anti-lumped mass

Operto et al., Geophysics, 2007

Pattern of the impedance matrix

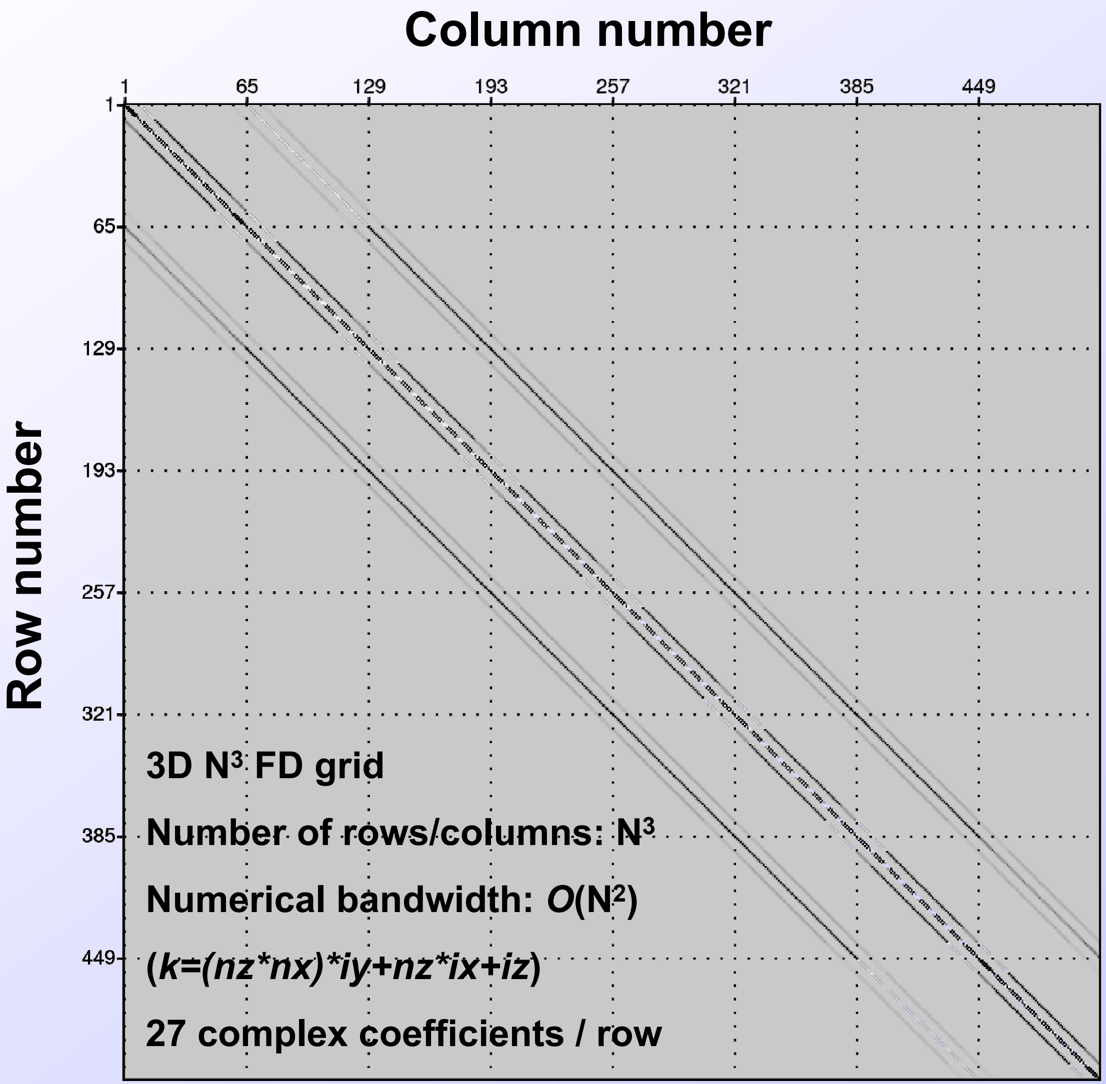

3D N^3 FD grid

Number of rows/columns: N^3

Numerical bandwidth: $O(N^2)$

($k=(nz*nx)*iy+nz*ix+iz$)

27 complex coefficients / row

Accuracy of the mixed grid stencil

Phase velocity dispersion analysis

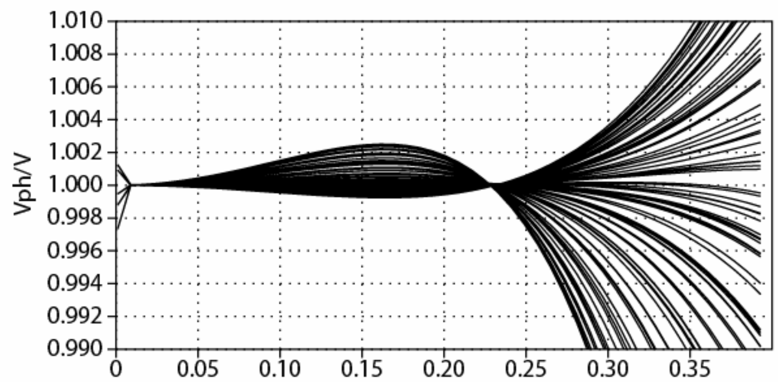

1/G (G: number of grid points per wavelength)

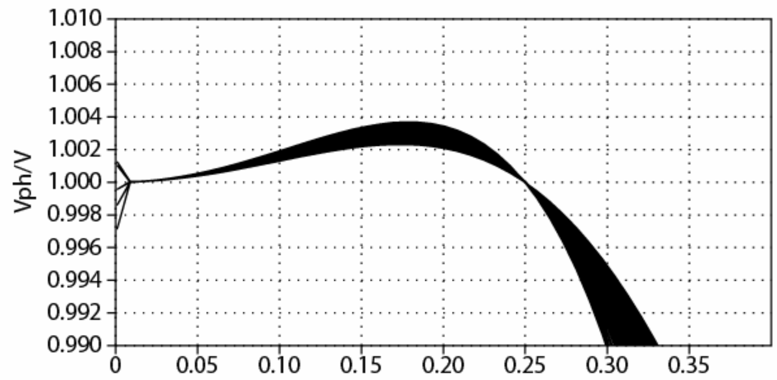

1/G (G: number of grid points per wavelength)

Sensitivity of the dispersion to the weights

Simulation in an infinite homogeneous medium; the grid interval satisfies 4 grid point per wavelength

h = 100 m; n1 = 41, n2 = 41, n3 = 121; freq = 3.75 Hz; Vp = 1.5 km/s

Source position: x1 = 2 km, x2 = 2km, x3 =1 km; **Receiver line:** x1 = 2 km, x2 = 2 km

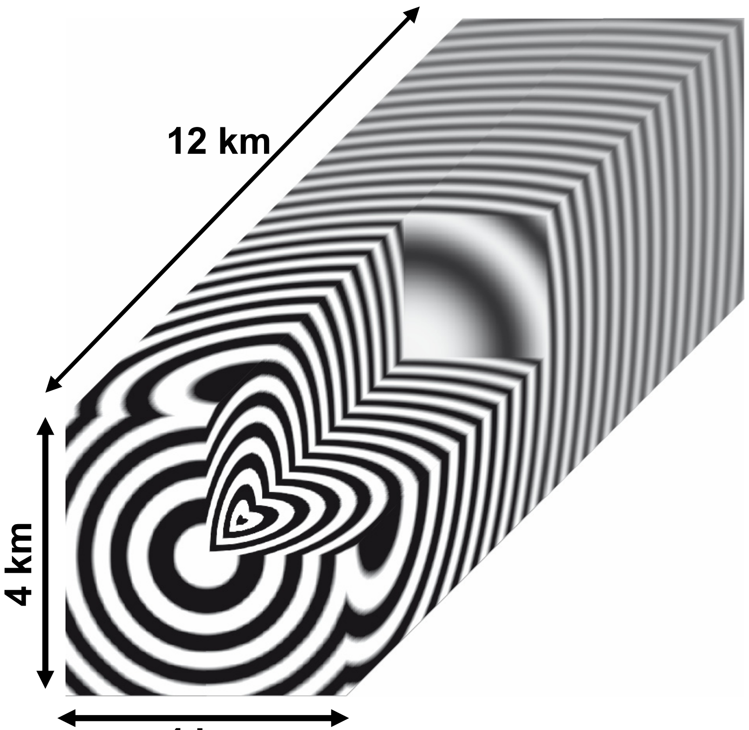

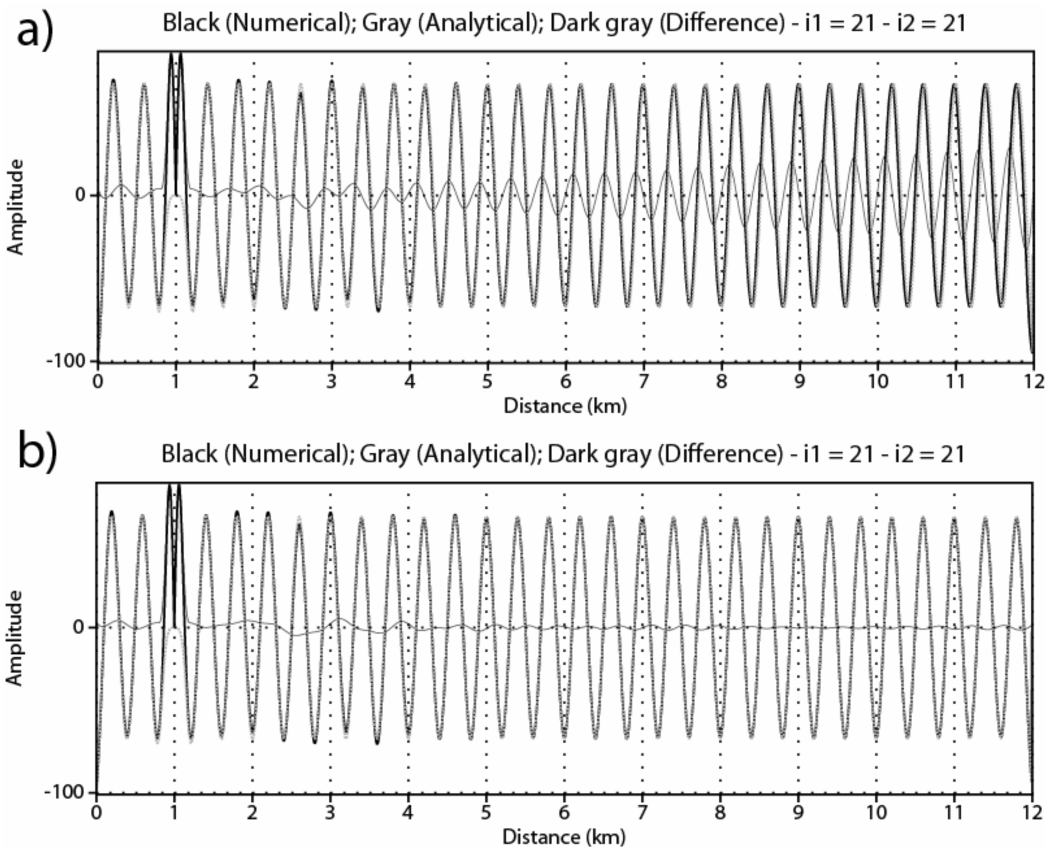

(a) Dispersion minimized for G = 10, 8, 6, 4; (b) Dispersion minimized for G = 4

MUMPS functionalities

• Single precision complex arithmetic

• Sparse right-hand side storage and multiple-rhs solve

• Distributed solutions (although a new domain decomposition is performed for gradient computation)

• Sequential analysis

• METIS ordering

Possible bottlenecks: memory requirement of the sequential (parallel?) analysis; 64-bit integers.

Numerical examples (MUMPS v4.9.2)

Computational platform: SGI Altix ICE 8200 (Jade CINES)

Bi-processor nodes with 30Gb of RAM – Processor Intel Quad-core E5472

Running mode: mpiprocs = 2 (2 MPI process per node)

Nomenclature:

N: number of MPI processes

M: dimension of a M^3 finite-difference grid

N_{LU}: number of LU factors

T_{LU}: elapsed time for factorization

N_u: number of unknowns

Mem_{LU}: working space for factorization

T_{LU}: elapsed time for factorization

T_S: elapsed time for one solve

N_{PML}: number of grid points in PML layers

$M_{PML} = M + 2 \times N_{PML}$

Complexity analysis

Memory

Modeling in a FD M³ grid

M	$N_u(.10^6)$
31	0.10
71	0.66
111	2.05
151	4.66
191	8.87

Number of unknowns as a function of M. $N_{PML} = 8$.

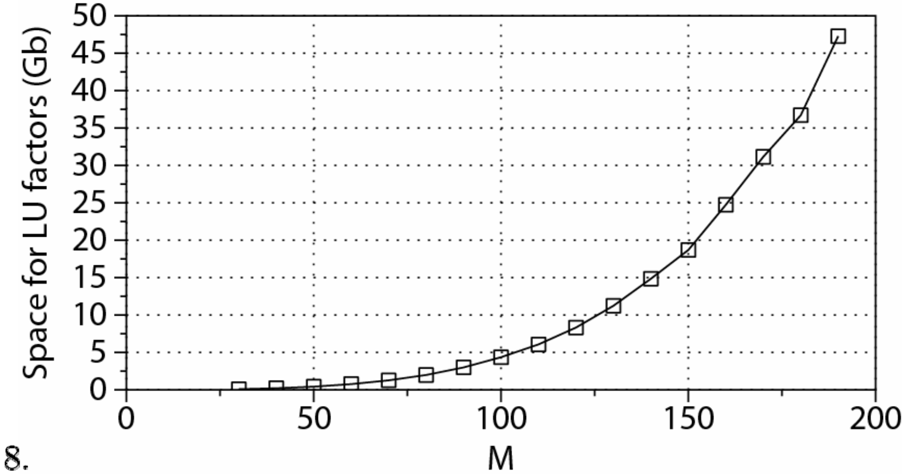

Observed memory complexity

$$\sim O\,(\mathrm{Log}_2\,(M)\;.\;M^{3.9})$$

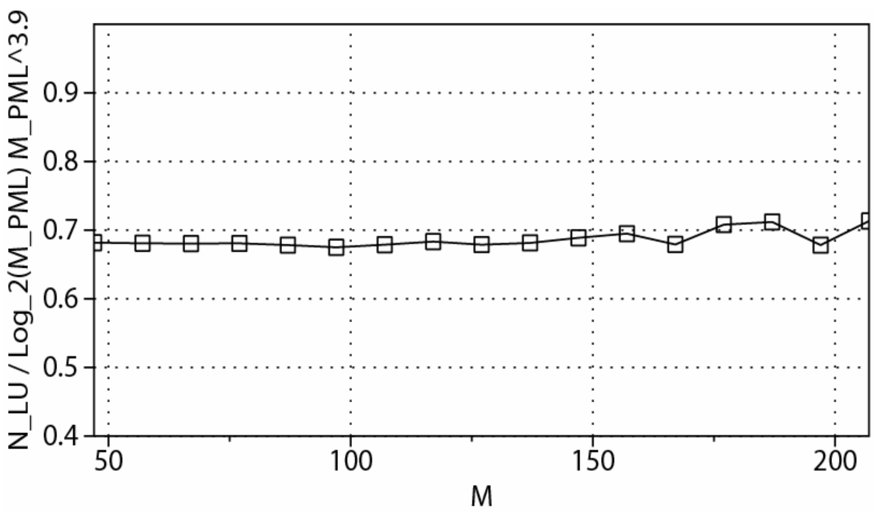

Complexity analysis

Time (LU)

Modeling in a FD M³ grid

M	$N_u(.10^6)$
31	0.10
71	0.66
111	2.05
151	4.66
191	8.87

Number of unknowns as a function of M.

Observed time complexity

~ *O* (M⁶)

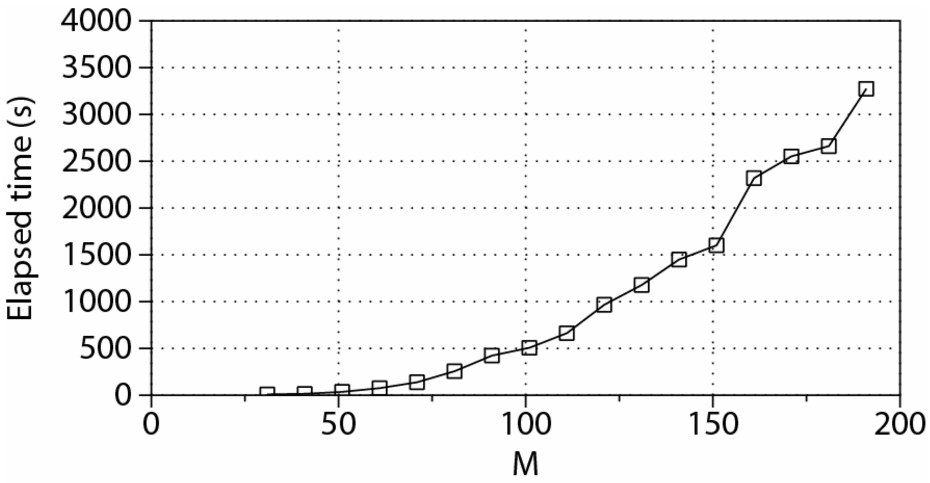

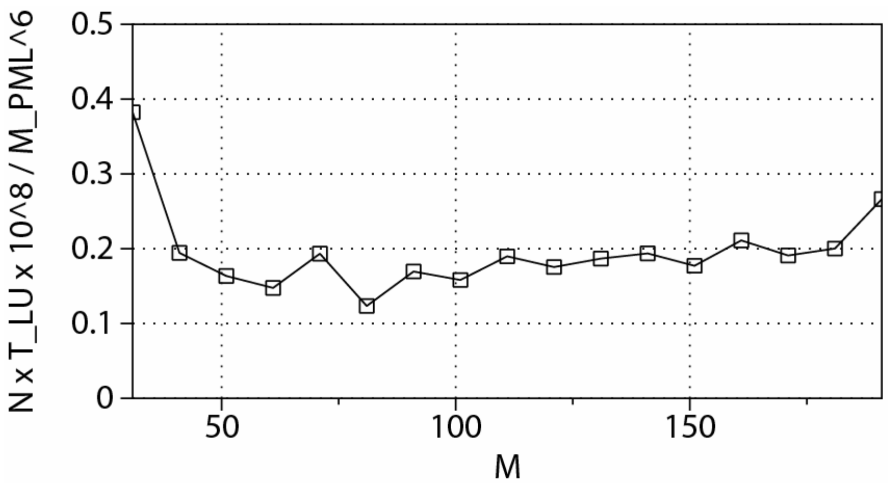

Scalability analysis

Elapsed time for factorization

FD Grid dimensions: $N_1 \times N_2 \times N_3$

N_1	N_2	N_3	N_{PML}	$N_u(10^6)$
46	91	161	8	1.174

Speedup

$$\mathcal{S} = (T_{ref} / T)$$

Efficiency

$$\mathcal{E} = \mathcal{S} \times (N_{ref} / N)$$

N: number of MPI processes

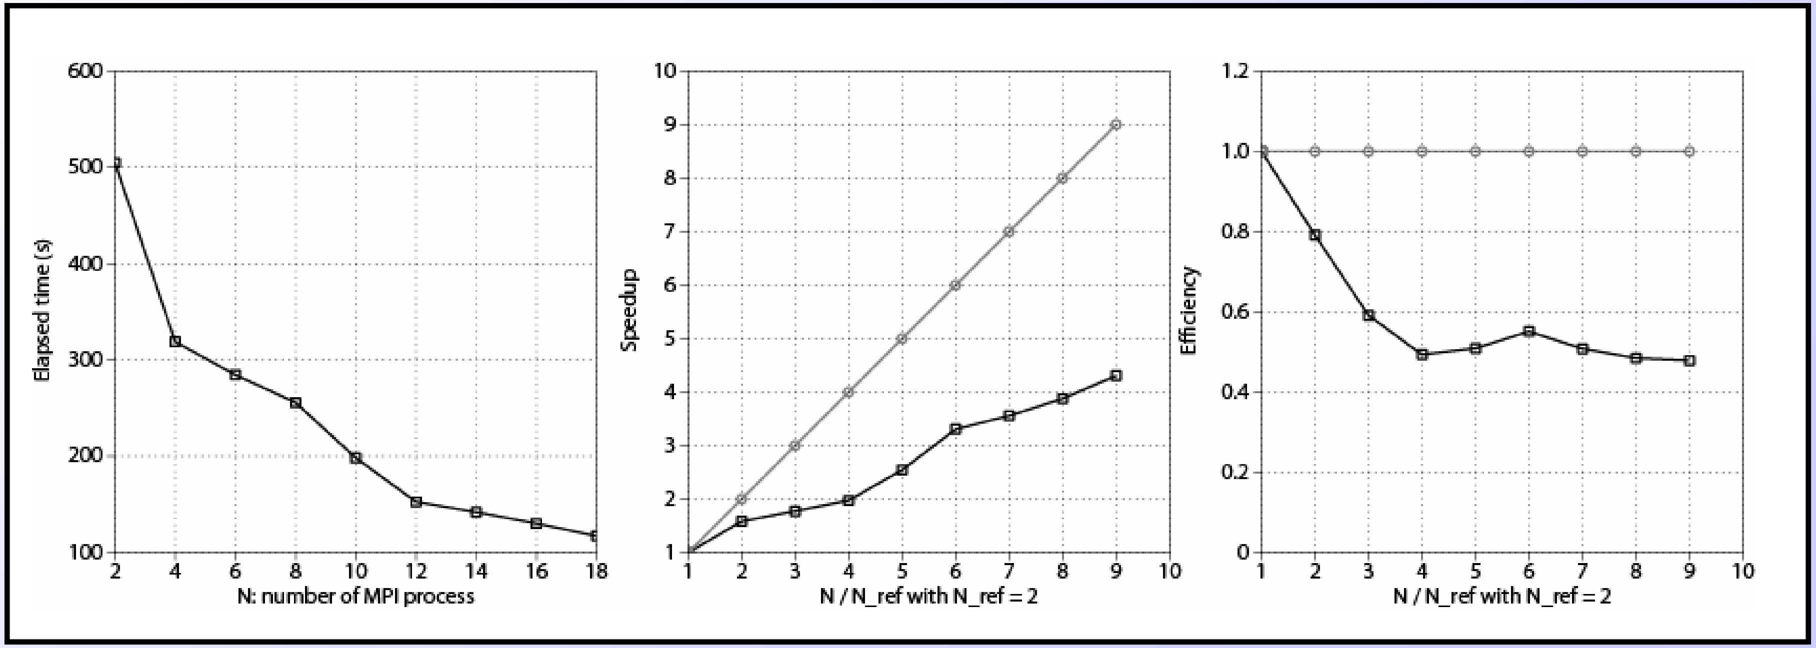

Scalability analysis

Elapsed time for solve

FD Grid dimensions: N_1 x N_2 x N_3

N_1	N_2	N_3	N_{PML}	$N_u(10^6)$
46	91	161	8	1.174

Speedup

$$\mathcal{S} = (T_{ref} / T)$$

Efficiency

$$\mathcal{E} = \mathcal{S} \times (N_{ref} / N)$$

N: number of MPI processes

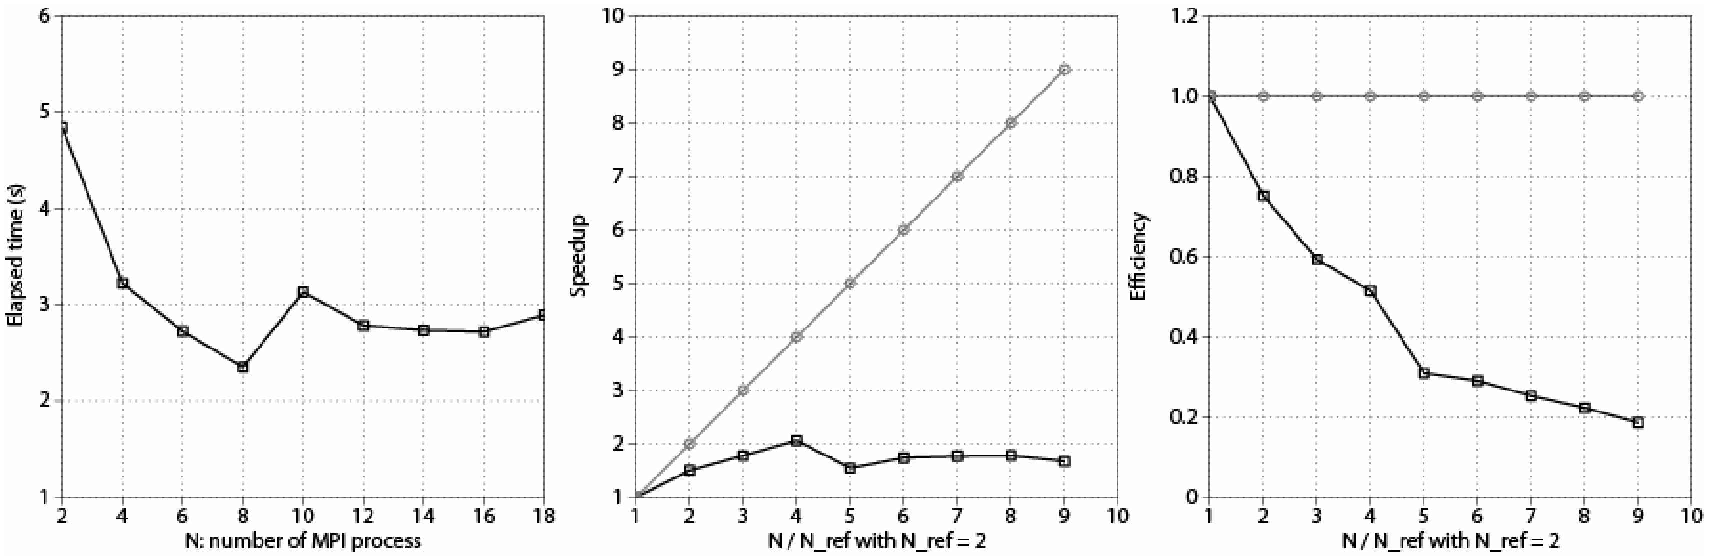

Scalability analysis

Memory overhead

FD Grid dimensions: N_1 x N_2 x N_3

N_1	N_2	N_3	N_{PML}	$N_u(10^6)$
46	91	161	8	1.174

Mem

Average space per working proc. during facto.

Memory overhead:

(N / N_ref) x (Mem / Mem_ref)

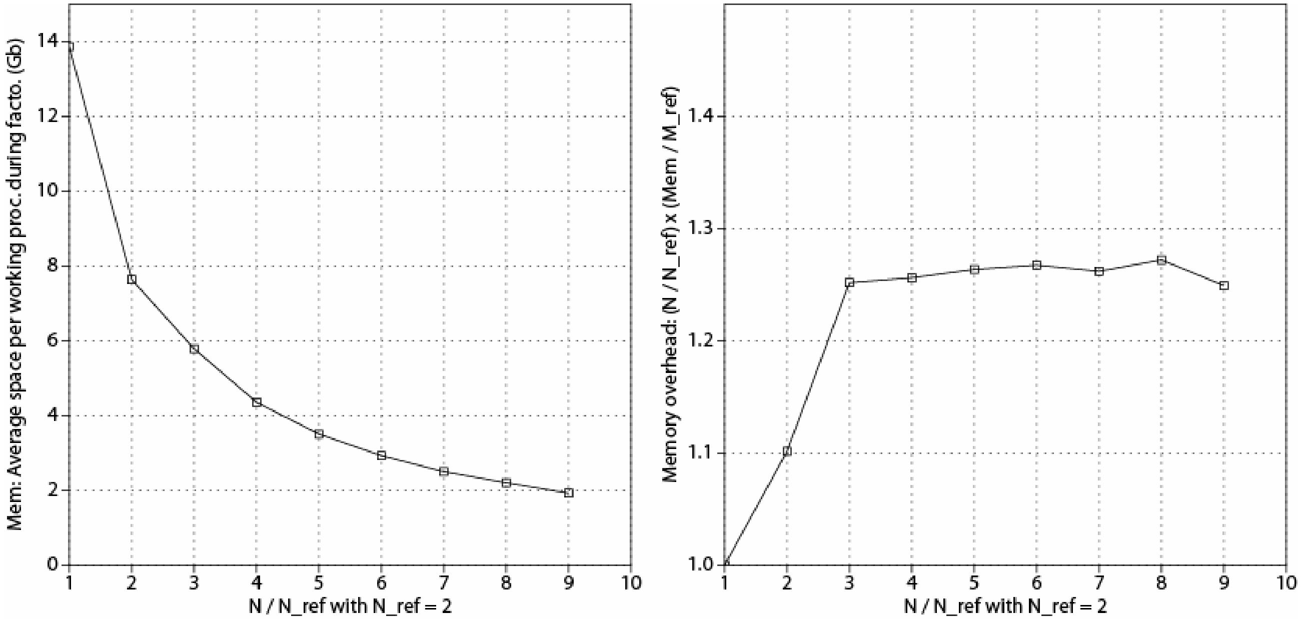

Simulation in the Valhall model

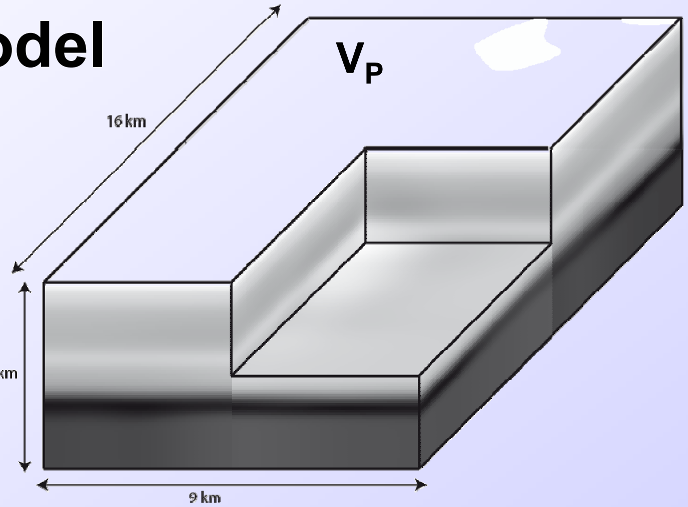

3.5 Hz

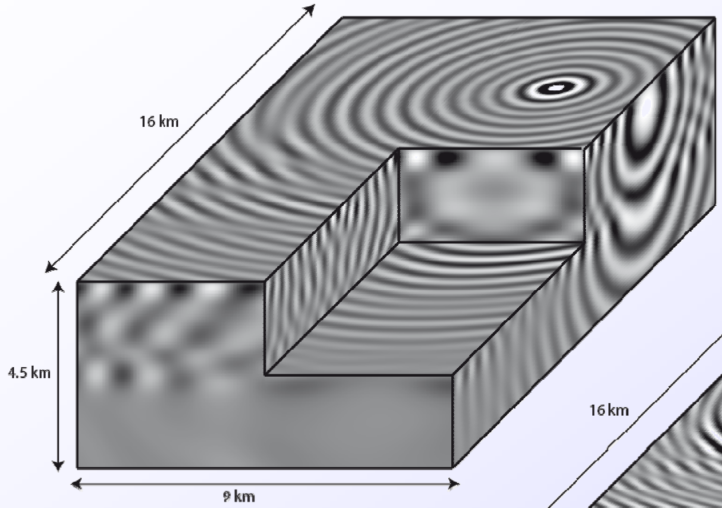

5.0 Hz

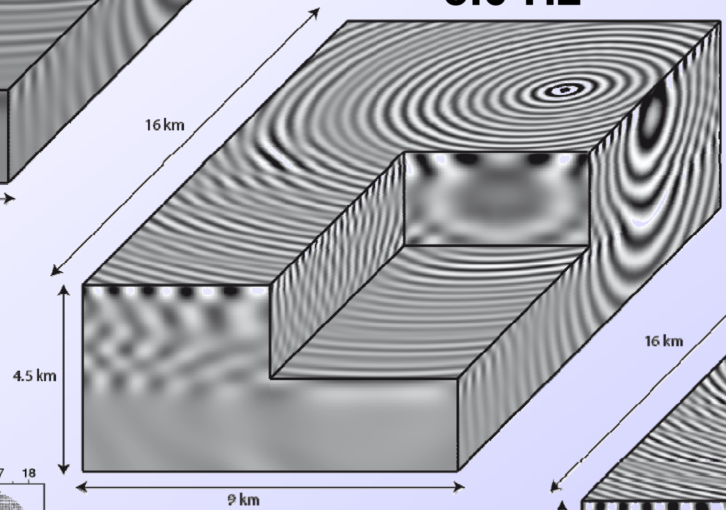

7.5 Hz

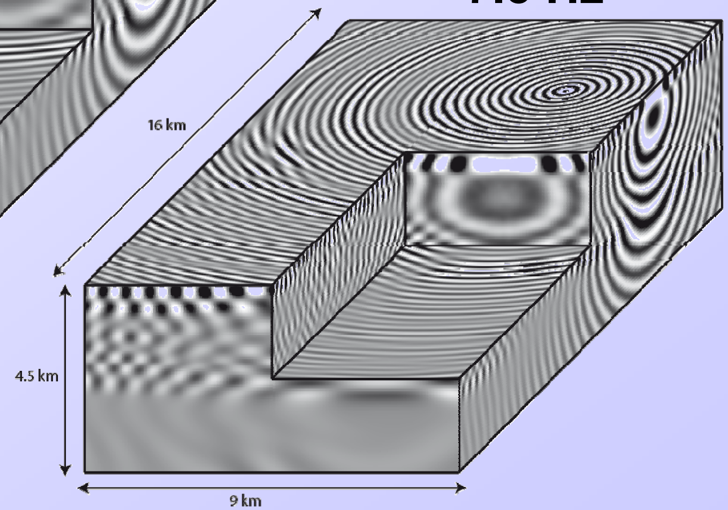

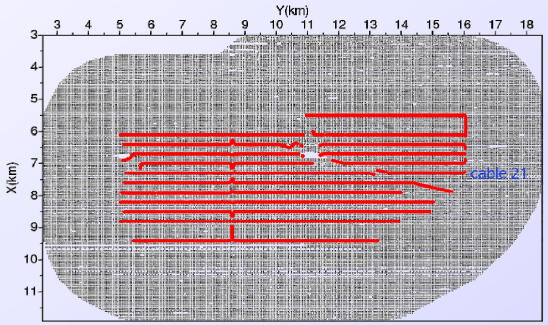

Computational cost of the Valhall case study

Memory per MPI process: 15 Gb

$f(Hz)$	$h(m)$	$n_1 \times n_2 \times n_3$	n_{pml}	$n_u(10^6)$	$mem_{LU}(Gb)$	$t_{LU}(s)$	$t_s(s)$	#nd	#pr	#th	#rhs
3.5	100	46 x 91 x 161	8	1.174	30.2	224	0.32	4	8	2	30
4.7	75	60 x 120 x 213	8	2.37	78.0	327.93	0.53	12	24	2	30
7.5	50	90 x 180 x 320	8	6.98	352	2322.66	1.69	18	36	2	30
3.5	100	46 x 91 x 161	8	1.174	30.2	122	0.45	4	8	4	30
4.7	75	60 x 120 x 213	8	2.37	78.0	191.80	0.65	12	24	4	30
7.5	50	90 x 180 x 320	8	6.98	352	1362.27	2.32	18	36	4	30

Table 1: $f(Hz)$: Frequency; $n_1 \times n_2 \times n_3$: grid dimension. n_{pml}: number of grid points in PML layers. n_u: Number of unknowns. $mem_{LU}(Gb)$: memory for LU factorization. $t_{LU}(s)$: elapsed time for LU factorizarion. $t_s(s)$: elapsed time for one-rhs solve. #nd: number of node. #pr: number of MPI process. #th: number of threads per node. #rhs: number of rhs simultaneously processed during solve phase.

*Elapsed time for 2300 rhs: (**LU + multi-rhs solve**): 960 s (3.5 Hz), 1547 s (5 Hz), 6209 s (7.5 Hz)*

*Elapsed time 1 FWI iteration (**3 x LU + 4 x multi-rhs solve**): 1 h (3.5 Hz), 1.6 h (5 Hz), 6.3 h (7.5 Hz)*

*Elapsed time 1 frequency FWI (**10 FWI iterations**): 10 h (3.5 Hz), 16 h (5 Hz), 63 h (7.5 Hz)*

TDM versus DSM

f = 3.5 Hz - #rhs = 2300

$O(\Delta t^2,\Delta x^4)$ FDTDM - Trace length: 10 s

Sequential elapsed time for 1 source: 27 s.

#proc = 8

Elapsed time for 2300-source modeling

	Preproc. (s)	multi-rhs solve (s)	Total (s)
TDM	0	7762	7762
DSM	224 (LU)	736	960

#proc = 2300

Elapsed time for 2300-source modeling

	Preproc. (s)	multi-rhs solve (s)	Total (s)
TDM	0	27	27
DSM	122 (LU)	3.6	125.6

***TDM**: Finite-Difference Time-Domain modeling; **DSM**: Direct-solver frequency-domain modeling*

TDM versus DSM

f = 7.5 Hz - #rhs = 2300

$O(\Delta t^2, \Delta x^4)$ FDTDM - Trace length: 10 s

Sequential elapsed time for 1 source: 326 s.

#proc = 36

Elapsed time for 2300-source modeling

	Preproc. (s)	multi-rhs solve (s)	Total (s)
TDM	0	20830	20830
DSM	2322 (LU)	3887	6209

#proc = 2300

Elapsed time for 2300-source modeling

	Preproc. (s)	multi-rhs solve (s)	Total (s)
TDM	0	326	326
DSM	1362 (LU)	83	1445

***TDM**: Finite-Difference Time-Domain modeling; **DSM**: Direct-solver frequency-domain modeling*

3D visco-acoustic wave modeling

based on a hybrid direct/iterative solver

Collaboration with L. Giraud and A. Haidar (Cerfacs)

Reference: PhD thesis A. Haidar, CERFACS, 2008

Implementation: F. Sourbier, H. Ben Hadj Ali

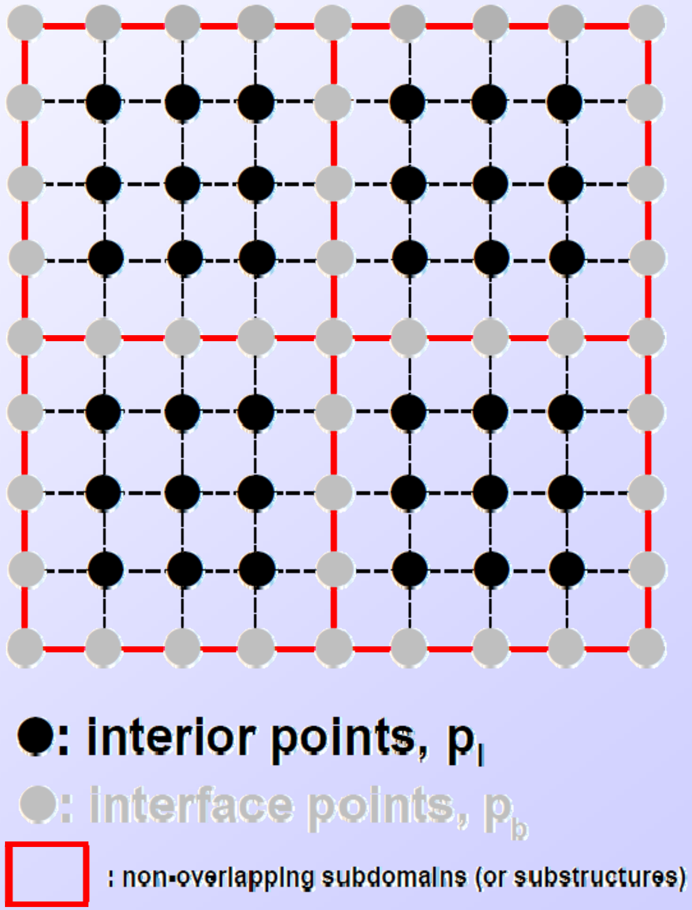

3D visco-acoustic wave modeling

based on a hybrid direct/iterative solver

$$A \, . \, p = s$$

$$\begin{bmatrix} A_{\mathcal{I}\mathcal{I}} & A_{\mathcal{I}\Gamma} \\ A_{\Gamma\mathcal{I}} & A_{\Gamma\Gamma} \end{bmatrix} \begin{bmatrix} p_{\mathcal{I}} \\ p_{\Gamma} \end{bmatrix} = \begin{bmatrix} s_{\mathcal{I}} \\ s_{\Gamma} \end{bmatrix}$$

Schur complement system for interface unknowns

$$\left(A_{\Gamma\Gamma} - A_{\Gamma\mathcal{I}} A_{\mathcal{I}\mathcal{I}}^{-1} A_{\mathcal{I}\Gamma}\right) p_{\Gamma} = s_{\Gamma} - A_{\Gamma\mathcal{I}} A_{\mathcal{I}\mathcal{I}}^{-1} s_{\mathcal{I}},$$

Interior unknowns

$$p_{\mathcal{I}} = -A_{\mathcal{I}\mathcal{I}}^{-1} A_{\mathcal{I}\Gamma} \, p_{\Gamma} + A_{\mathcal{I}\mathcal{I}}^{-1} \, s_{\mathcal{I}}.$$

Block-diagonal structure of the $A_{\mathcal{II}}$ matrix

$$A_{\mathcal{II}} = \begin{pmatrix} A_{\mathcal{I}_1\mathcal{I}_1} & & & \\ & A_{\mathcal{I}_2\mathcal{I}_2} & & \\ & & \ddots & \vdots \\ & & \ddots & \vdots \\ & & \ddots & \vdots \\ & & & A_{\mathcal{I}_{N_P}\mathcal{I}_{N_P}}. \end{pmatrix}$$

Parallel resolution of the Schur complement system *with GMRES*

The Schur complement as the sum of the local Schur complements

$$\mathcal{S} = \sum_{i=1}^{N_P} \mathcal{R}_{\Gamma_i}^T \mathcal{S}_i \mathcal{R}_{\Gamma_i}$$

Right preconditioning

$$SM^{-1}y = s$$

with an additive Schwarz preconditioner

$$M^{-1} = \sum_{i=1}^{N_P} \mathcal{R}_{\Gamma_i}^T \bar{\mathcal{S}}_i^{-1} \mathcal{R}_{\Gamma_i}$$

given by the sum of the local assembled Schur complements (restriction of the Schur complement to the interface i).

Memory complexity of HSM

$$O\ (k^3\ (N/k)^4) = O\ (N^4\ /\ k)$$

k: number of subdomains in 1 direction.

N: dimension of a 3D cubic grid N^3

Stopping criterion of GMRES iterations

$$\epsilon = \|Ap - s\| / \|s\|$$

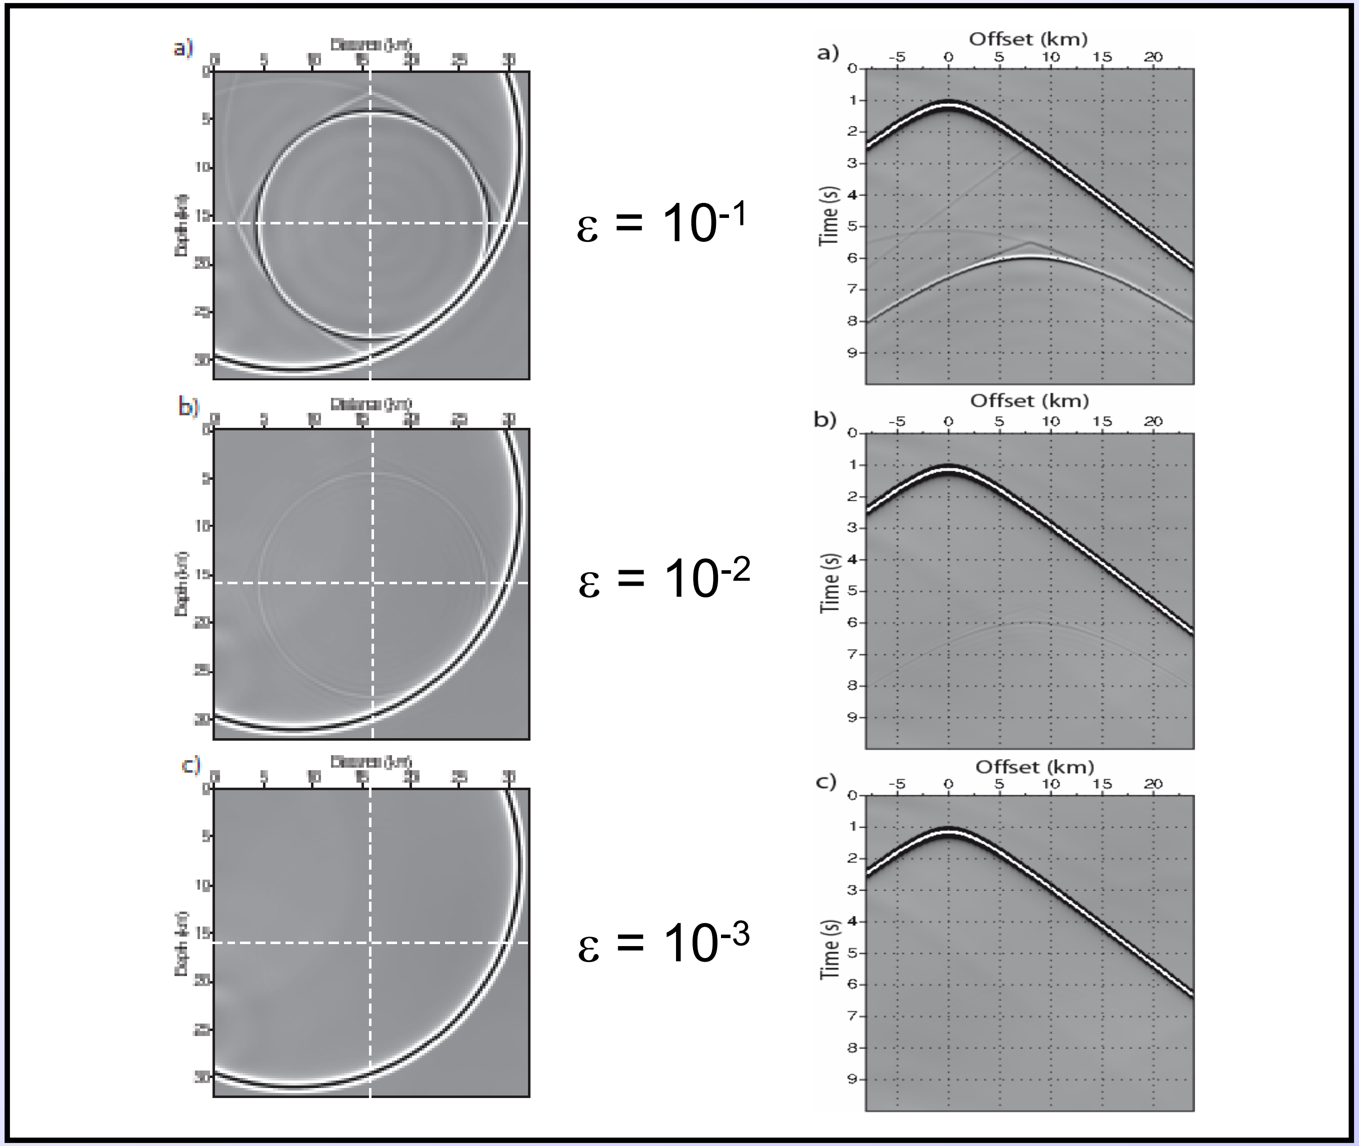

Case studies

SEG/EAGE overthrust and salt models

Comparison between Hybrid Solver and Time Domain Modelings (HSM vs TDM)

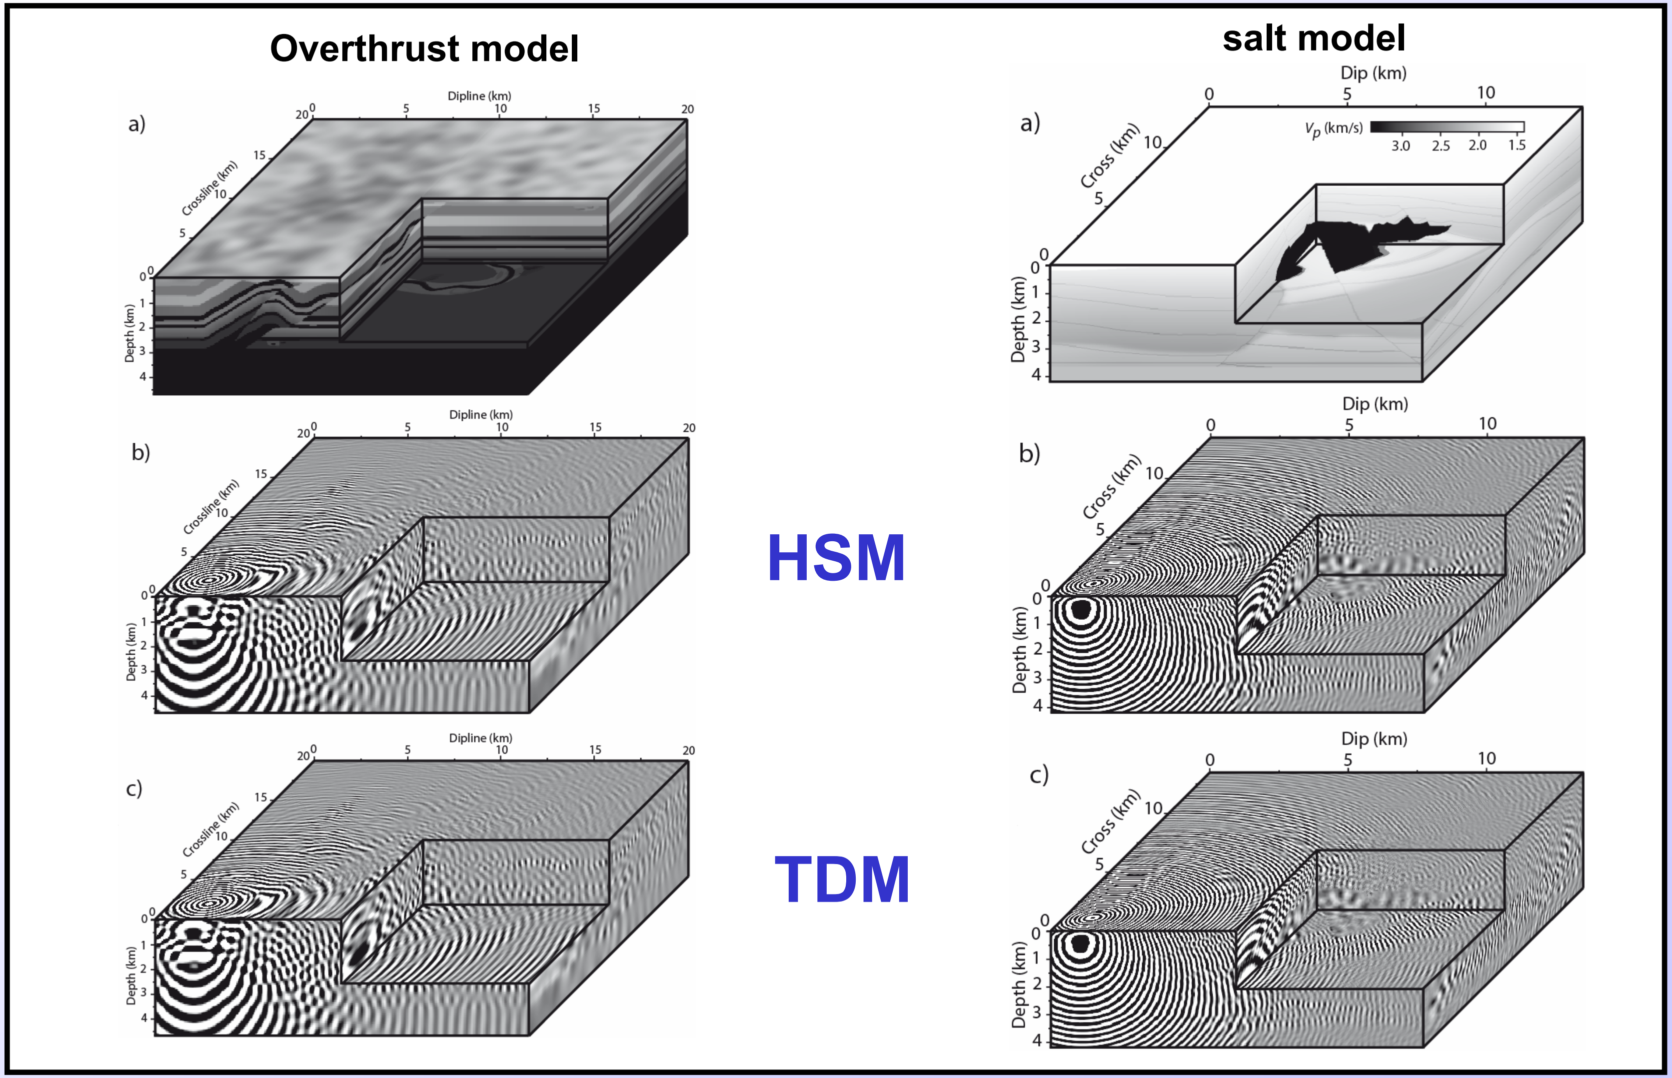

Dimensions of the numerical problem

Model	$L_x \times L_y \times L_z$ (km)	$N_x \times N_y \times N_z$	$h(m)$	$n_x^a \times n_y^a \times n_z^a$	$N_u(10^6)$	$f(Hz)$
Over.	20 x 20 x 4.7	400 x 400 x 94	50	9 x 9 x 8	25.7	10.8
Salt	13.5 x 13.5 x 13.5	450 x 450 x 140	30	6 x 6 x 5	32	12.5

Numerical results on IBM Blue Gene (IDRIS)

Mod.	N_P^{HSM}	k_x x k_y x k_z	n_x x n_y x n_z	$T_{tot}^{HSM}(s)$	$T_{RHS}^{HSM}(s)$	N_P^{TDM}	$T_{RHS}^{TDM}(s)$
Over.	1024	16 x 16 x 4	26 x 26 x 28	425	175	64	1700
Salt	980	14 x 14 x 5	33 x 33 x 30	797	109	64	2115

Scalability analysis of HSM – weak scaling

IBM Blue Gene (IDRIS)

T_{RHS} (s)

106
43
59
89
85

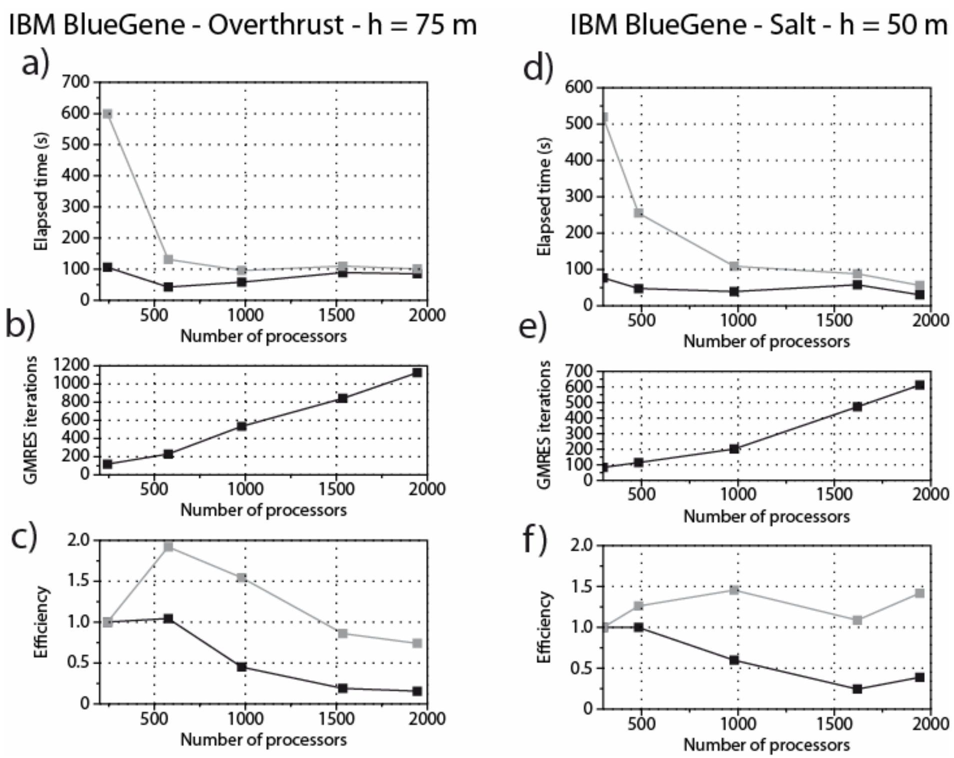

T_{RHS} (s)

77
48
30
58
30

Platform	Model	$L_x \times L_y \times L_z (km)$	$N_x \times N_y \times N_z$	$h(m)$	$f(Hz)$	$N_u(10^6)$
Bab.	Salt	13.5 x 13.5 x 4.15	270 x 270 x 84	50	7.5	8.18
Bab.	Over.	20 x 20 x 4.7	266 x 266 x 62	75	7.5	6.20

Scalability analysis of HSM – weak scaling

IBM Power 6 (IDRIS)

T_{RHS} (s)

15.1
11.2
9.8
10.5
10.1

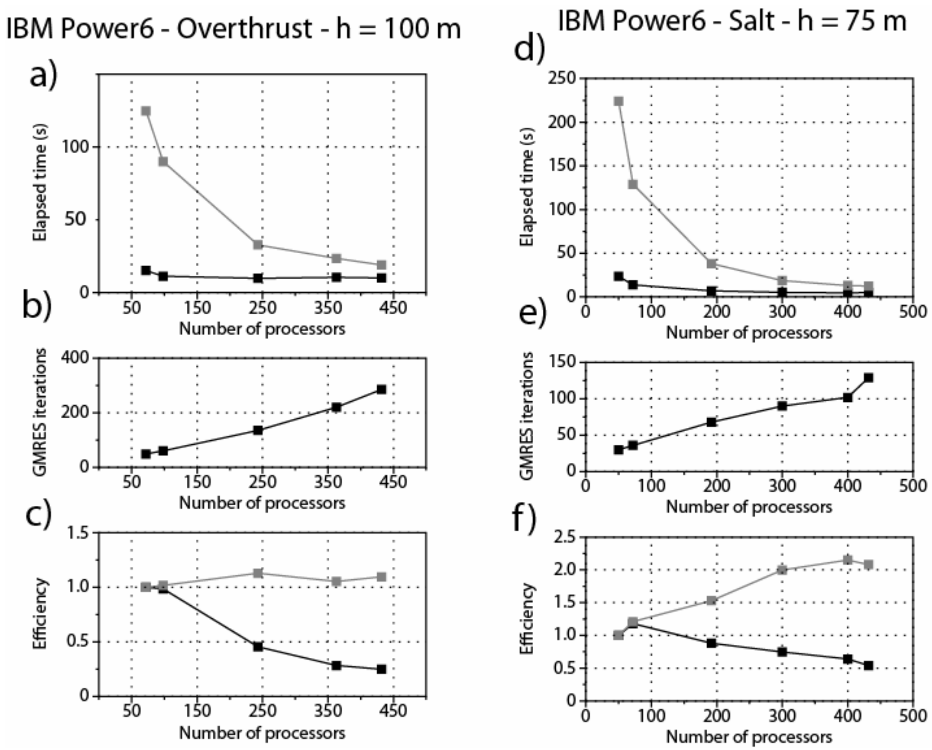

T_{RHS} (s)

23
14
7
4.6
5

Platform	Model	$L_x \times L_y \times L_z (km)$	$N_x \times N_y \times N_z$	$h(m)$	$f(Hz)$	$N_u(10^6)$
Var.	Salt	13.5 x 13.5 x 4.15	180 x 180 x 56	75	5	2.77
Var.	Over.	20 x 20 x 4.7	200 x 200 x 46	100	5.4	2.89

Conclusion and perspectives

• Realistic applications of 3D acoustic FWI at low frequencies is possible today with parallel direct solver as MUMPS.

Application to the 3D OBC Valhall experiment scheduled in the coming year for comparison with FWI results based on time-domain modeling and iterative –ssolver modeling.

• Hybrid solvers allows us to tackle larger problems but the cost of the iterative solver remains prohibitive for a large number of rhs. We use source assembling and phase encoding technics to mitigate the number of rhs in the modeling.

• Combination of the direct-solver approache at low frequencies with the time-domain modeling approach at higher frequencies might be the approach of choice.

• Extension of the direct-solver approach to the elastic wave equation is currently too demanding.

Acknowledgments

Thanks to the sponsor of the SEISCOPE consortium (http://seiscope.oca.eu) for their support.

Thanks to the SIGAMM, IDRIS and CINES computer centers for computational facilities

and thanks to the MUMPS team