

Global seismic tomography and modern parallel computers

Journal Article

Author(s): Soldati, Gaia; Boschi, Lapo; Piersanti, Antonio

Publication date: 2006

Permanent link: https://doi.org/10.3929/ethz-b-000035350

Rights / license: Creative Commons Attribution 3.0 Unported

Originally published in: Annals of Geophysics 49(4-5)

Global seismic tomography and modern parallel computers

Gaia Soldati (1), Lapo Boschi (2) and Antonio Piersanti (1)

(¹) Istituto Nazionale di Geofisica e Vulcanologia, Roma, Italy (²) ETH Zürich, Switzerland

Abstract

A fast technological progress is providing seismic tomographers with computers of rapidly increasing speed and RAM, that are not always properly taken advantage of. Large computers with both shared-memory and distributed-memory architectures have made it possible to approach the tomographic inverse problem more accurately. For example, resolution can be quantified from the resolution matrix rather than checkerboard tests; the covariance matrix can be calculated to evaluate the propagation of errors from data to model parameters; the *L*-curve method can be applied to determine a range of acceptable regularization schemes. We show how these exercises can be implemented efficiently on different hardware architectures.

Key words numerical inverse theory – seismology – global tomography – seismic resolution – Earth's mantle

1. Introduction

Earth tomography requires the solution of inherently large, mixed-determined inverse problems. Since its very beginning, it has always involved the implementation of efficient algorithms on state-of-the-art computers.

In 1984, John Woodhouse and Adam Dziewonski published in Journal of Geophysical Research one of the few articles that defined global seismic tomography. In their conclusions, they noted: «The calculations reported upon here were performed using an array processor (Floating Point, 120B) which, programmed in Fortran, is typically 10 times as fast as, say, a VAX 11/780. The path by path inversions and source determinations, using mantle waves and body waves, occupied the machine for approximately 60 h, and each global iteration took approximately 7 h. The inclusion of more data, and the extension of the method to higher frequencies, will probably require the use of the most advanced 'mainframe' computers.» They are referring to an image of the Earth's upper mantle parameterized in terms of a cubic polynomial (vertically), and spherical harmonics up to 8° (laterally); that is to say, entirely specified by just 324 parameters (Woodhouse and Dziewonski, 1984). The lower mantle was mapped in a separate inversion, as a linear combination of spherical harmonics up to 6°, multiplied by Legendre polynomials up to 4°, resulting in 245 parameters; «the size of the array needed to store the lower (or upper) triangle of the corresponding inner product matrix is 31.035; just a little less than the data memory of our AP120B array processor, without which this study would not be feasible» (Dziewonski, 1984).

In the following years, tomographers took advantage of fast technological progress. The following generation of models published by the Harvard group covered the entire mantle, and were linear combinations of Chebyshev

Mailing address: Dr. Gaia Soldati, Istituto Nazionale di Geofisica e Vulcanologia, Via di Vigna Murata 605, 00143 Roma, Italy; e-mail: soldati@ingv.it

polynomials and spherical harmonics up to degree 12 (*e.g.*, Su *et al.*, 1994).

Other authors preferred a different approach, parameterizing the Earth's mantle with grids of voxels (Hager and Clayton, 1989; Inoue *et al.*, 1990). The «voxel» approach involved a substantially larger number of model coefficients ($\sim 10^4$), but also implied that the integral of data sensitivity multiplied by the basis functions be most often 0. The latter circumstance has important consequences, as we shall briefly illustrate: the tomographic linear inverse problem is typically written

$$A \cdot x = d \tag{1.1}$$

where the entries of the vector x are the coefficients of the solution model (initially unknown), d are the data (*e.g.*, travel times), and A is a matrix whose ij entry equals the integral, over the entire volume of the mantle, of the sensitivity of the *i*-th measurement to the Earth property (typically a seismic velocity or slowness) to be mapped, times the *j*-th basis function used to describe such property.

In the ray theory approximation, and if (as will always be the case here) d are observations of travel time anomaly, the volume integral reduces to an integral along the seismic ray path (sensitivity is 0 everywhere but on the ray path), and

$$A_{ij} = \int_{i-\text{th path}} f_j(r(s)) ds \qquad (1.2)$$

where *s* denotes the incremental length along the ray path, identified by the equation r = r(s), with *r* denoting position; the *N* basis functions f_j are used to describe the slowness $\delta p(r) =$ $= \sum_{k=1}^{N} x_k f_k(r)$ of the seismic phase in question. (It would be equivalent to formulate the problem in terms of velocity, but slowness happens to make algebra simpler.)

If the functions f_j are spherical harmonics, nonzero over the entire surface of the Earth, the integral in eq. (1.2) will be nonzero for all values of j. In the case of a voxel (spline, wavelet or other «local» functions) parameterization, the same integral will be 0, except for values of j whose corresponding voxel is crossed by a ray path. As anticipated, the matrix A will therefore be dense if the model is parameterized in terms of spherical harmonics or other «global» functions; sparse if local functions are used.

The same is consequently true of the matrix $A^T \cdot A$, whose inverse has to be calculated for the least squares solution x_{LS} to (1.1) to be found (a necessary step, as (1.1) in global seismology is strongly mixed-determined and does not have an exact solution),

$$x_{LS} = (A^T \cdot A + D)^{-1} \cdot A^T \cdot d \qquad (1.3)$$

where the matrix *D* depends on the regularization scheme (*e.g.*, Boschi and Dziewonski, 1999).

When A and $A^T \cdot A$ are sparse, eq. (1.3) is most efficiently implemented via an iterative algorithm like CG or LSQR (*e.g.*, Trefethen and Bau, 1997). When they are dense, iterative algorithms become as slow as direct ones: the most efficient approach is then to implement (1.3) via Cholesky factorization of $A^T \cdot A$, and subsequent backsubstitution (*e.g.*, Press *et al.*, 1994; Trefethen and Bau, 1997).

Implementation of LSOR does not require A^{T} A to be calculated, as LSQR operates directly on A. A is bigger but sparser than $A^T A$, and sparse matrices can be stored efficiently (e.g., Press et al., 1994) to minimize the required disc space or RAM; A is therefore often less cumbersome than $A^T A$: it is so, at least, when global body wave travel time databases are inverted to derive global Earth structure. This, and the remarkable speed of LSQR in a regime of sparse A, allowed Grand (1994) and van der Hilst *et al.* (1997) to parameterize the Earth's mantle in terms of as many as N~250000 voxels: a three orders of magnitude increase in nominal resolution, with respect to the early studies of Dziewonski (1984) and Woodhouse and Dziewonski (1984) mentioned above.

Like Woodhouse and Dziewonski some ten years before, Grand and van der Hilst were exploiting available computers to their limit. Although sparse, A was still too large a matrix to be entirely fit on the RAM of a processor; LSQR, however, required only parts of A to be available at one time in the RAM: they could be run without ever storing A entirely in memory, at the expense of massive input from disc at each iteration. Given the number of solution coefficients, at least $\sim 10^2$ iterations were probably needed for LSQR to converge. This made even LSQR a very slow process, and left researchers with relatively little freedom to test the model resolution and the effect on the solution of different regularization schemes.

In view of the exponential growth in CPU speed over the last decade (*e.g.*, Bunge and Tromp, 2003), and the concurrent decrease in the price of RAM, the current generation of global seismic tomographers has the means to approach the discipline in an entirely new fashion. With a fast processor, and enough RAM to store *A* entirely, not only LSQR is sped up enormously, but more time-consuming direct algorithms like Cholesky factorization of $A^T \cdot A$ also become feasible.

2. Cholesky factorization on a multiprocessor, shared-memory computer

It was originally proved by Paige and Saunders (1982), and later confirmed by Nolet (1985) and Boschi and Dziewonski (1999), with applications to mixed-determined tomographic problems, that LSOR converges correctly to the damped least squares solution (1.3), typically after a number of iterations $\ll N$. If A is sparse and sufficient RAM is available to store it. LSOR is therefore the most efficient algorithm to solve an inverse problem in the least squares sense. On the other hand, because it bypasses the calculation of $A^T \cdot A$ and the direct implementation of (1.3), LSOR cannot provide any measure of goodness of resolution and covariance, except by means of resolution, or «checkerboard» tests. The unreliability of the measure of resolution that those tests provide has been pointed out, for example, by Lévêque et al. (1993), and there have been efforts to derive the resolution matrix via an iterative, LSQ-type calculation (Zhang and McMehan, 1995: Mink-off, 1996: Nolet et al., 1999, 2001: Yao et al., 1999, 2001; Vasco et al., 2003).

The resolution matrix *R* can be thought of as the operator that relates «output» and «input» model in any checkerboard test; Menke (1989) shows that

$$R = (A^T \cdot A + D)^{-1} \cdot A^T \cdot A.$$
 (2.1)

Clearly, *R* does not depend on the input model, and its similarity to the identity matrix is a measure of goodness of resolution. Its calculation requires that $A^T \cdot A + D$ be explicitly inverted, and this is most efficiently achieved by Cholesky factorization of this matrix. Once the damped inverse of $A^T \cdot A$ is found, *R* is quickly determined by backsubsitution, applied on the matrix $A^T \cdot A$ instead of the vector $A^T \cdot d$; this endeavour is not significantly more time-consuming than the implementation of (1.3) via Cholesky factorization and backsubstitution.

Boschi (2003) computed R from the global teleseismic P-wave travel time database of Antolik et al. (2003), based upon the ISC Bullettins and including ~600000 summary observations. He parameterized the Earth's mantle in terms of 20 vertical splines and 362 horizontal splines (N=7240). Boschi's (2003) exercise was conducted on an IBM SP2 with 16 processors and 32 Gb of RAM. The IBM SP2 is a «shared-memory» machine: any processor can access at the same speed its entire RAM. This is a very useful feature when large matrices have to be factorized, a process that is inherently hard to parallelize; clusters of PCs are by construction «distributed-memory» computers, and hence more useful for the solution of forward, rather than inverse, problems.

Boschi (2003) notes that the most time-consuming step in deriving *R* is the computation of $A^T \cdot A$, which took about twent-four hours. This process was parallelized by subdividing the database in as many subsets as there were available processors, computing each subset's contribution to $A^T \cdot A$ on a separate processor, and eventually adding up the results. After computing $A^T \cdot A$, which needs to be done only once, x_{LS} and *R* can be derived in a few minutes; Boschi (2003) was thus able to perform numerous inversions, experimenting with the damping scheme and exploring the solution space, calculating each time the associated *R*.

3. Running LSQR repeatedly on a distributed-memory cluster of PCs

The computer on which this article is being written, a Linux PC sitting on the second au-

thor's desk, is equipped with a dual processor and 3 Gb of RAM. One year ago, this much RAM costed just about 800 US Dollars. We concluded Section 1 pointing out that, as one can easily afford enough RAM to store arrays of $\sim 10^5$ elements, iterative algorithms like LSQR become extremely efficient. On this very computer, one LSQR inversion involving some 25000 model coefficients runs to convergence in a matter of seconds.

Let us now show how, on a parallel, distributed-memory computer (a cluster of PCs), the resolution matrix R can also be derived in a reasonable amount of time, without calculating and Cholesky-factorizing $A^T A$.

Implementing eq. (2.1) is equivalent to implementing N times eq. (1.3), replacing each time d with a different column of A (recall that N denotes the number of model parameters, and hence the number of columns of A). R can therefore be derived by means of N independent LSOR inversions of A, without finding $A^T \cdot A$. When only one or few processors are available, and with $N \sim 10^5$ as in some of the experiments mentioned above, this process would be extremely timeconsuming, to the point of not being worthwhile. If a relatively large parallel machine is available, however, the problem can be easily parallelized, by simply subdividing the N inversions into N/n_P subsets. n_P denoting the number of processors. Each subset of inversions is then performed independently on a separate processor, and the time needed to compute R is reduced by a factor n_P .

It should be noted that the most timeconsuming step, input of *A* from disc to RAM, needs to be performed only once per processor, no matter how many inversions are then run on each processor.

Figures 1 and 2 illustrate how this procedure applies to a real inverse problem. We describe the distribution of *P*-velocity heterogeneities in the Earth's mantle in terms of a grid of voxels of constant horizontal extent; voxel functions guarantee that *A* be more sparse than in the case of splines. We invert, again, the *P*-wave travel time database of Antolik *et al.* (2003). Following, *e.g.*, Inoue *et al.* (1990, Section 3.3.1 and fig. 2), we select roughness minimization as our only regularization criterion, and perform a



Fig. 1. Data misfit achieved by a set of solution models, *versus* the integrated roughness of each model. This measure of model complexity is normalized against model RMS. Least squares solutions were found from a wide range of values of the roughness minimization parameter, and no other minimization constraint. We repeated the experiment with voxels of lateral extent $15^{\circ} \times 15^{\circ}$, $10^{\circ} \times 10^{\circ}$, $7.5^{\circ} \times 7.5^{\circ}$, $6^{\circ} \times 6^{\circ}$, $5^{\circ} \times 5^{\circ}$, $3.75^{\circ} \times 3.75^{\circ}$, $3^{\circ} \times 3^{\circ}$, $2.5^{\circ} \times 2.5^{\circ}$, and constant vertical thickness (~200 km) (Soldati and Boschi, 2004): corresponding solutions align on different *L*-curves, and squares of decreasing size correspond to increasingly fine parameterization.

number of preliminary inversions, at different parameterization levels, to assess the dependence of the solution on the regularization parameter. Plotting misfit to the data (defined as 1-the variance reduction) against «total roughness» (the integral of the surface gradient of the model over the entire solid angle is computed for each layer of the model, and then the RMS is taken) in fig. 1, we find, for each parameterization, a set of points aligned along the expected L-shaped curve (e.g., Hansen, 1992; Boschi et al., 2006). Each point on the L-curve corresponds to a model derived at this preliminary stage, and the roughness damping parameter grows monotonically with increasing misfit. The shape of the curve, resembling the letter L, confirms that the data contain coherent and statistically significant information; the decrease in misfit is very fast in an overdamped regime, where a small reduction in the regularization parameter, and therefore a small increase in model complexity, is sufficient to improve the data fit substantially. The white noise that the data necessarily contain, and that regularization is supposed to eliminate, is harder to fit, even with large increases in model complexity: this is why the curve tends to become horizontal in the right part of the plot. Solution models lying in the vertical and horizontal portions of the *L*-



Fig. 2. *Top*: 24840×24840 (5°-voxel grid) resolution matrix *R* for the chosen regularization scheme, averaged (Boschi, 2003) so that it can be plotted in this limited space; here, vertical tradeoffs are most evident. *Bottom*: zooms on *R*, not averaged, at two selected layers (*left*: mid mantle at ~1300 km depth; *right*: lower mantle at ~2600 km).

curve can be discarded as overdamped and underdamped, respectively; preferred models should be chosen near its corner.

The selection of a damping scheme has always been a largely arbitrary process in global seismic tomography. The *L*-curve criterion is a way to reduce this arbitrarity. It is practical so long as a large number of LSQR inversions can be performed in a short time, and we have seen how this is made possible by simultaneous storage of the entire matrix *A* in memory, and/or availability of multiple processors.

After so selecting optimal roughness damping parameters at all parameterization levels. we restrict ourseleves to the case of 5° voxels. We show in fig. 2 the corresponding resolution matrix R as derived with multiple runs of LSOR. As to be expected (Boschi, 2003), R is quite different from the identity matrix: entries smaller than 1 on its diagonal indicate that the amplitude of velocity heterogeneities in the corresponding voxel is underestimated. Entries different from zero away from the diagonal identify episodes of fictitious coupling between model coefficients; naturally, the value of R_{ii} is proportional to the amount of coupling («tradeoff», «smearing», ...) between the *i*-th and *i*-th voxels (entries x_i and x_i of the solution vector).

4. Performance and accuracy of direct *versus* iterative implementations

In analogy with Yao *et al.* (1999), we calculate R associated with one given database and one choice of parameterization and regularization, both in the direct (Cholesky, Section 2 above) and iterative (LSQR, Section 3) approaches, and compare the results. As opposed to Singular Value Decomposition (SVD), the direct algorithm implemented by Yao *et al.* (1999), Cholesky factorization does not involve the cancellation of the smallest singular factors (Press *et al.*, 1994), so that in our experiment regularization is entirely controlled by the matrix D, and is therefore exactly equivalent in the direct and iterative calculations.

We implement eq. (2.1) exactly, by Cholesky factorization of $A^T \cdot A + D$, for the 5°-voxel parameterization described in Section 3, and applying the same regularization scheme that led to *R* in fig. 2. The result is shown in fig. 3, and in fig. 4 two lines of the directly- and iteratively-calculated *Rs* are compared in a geographic view. Differences are everywhere small, and negligible for *i*, *j* such that R_{ij} >0.2. Discrepancies between *R* calculated iteratively and directly (with SVD) by Yao *et al.* (1999, fig. 5) appear to be larger.

R in fig. 2 through 5 is a 24840×24840 matrix, as opposed to the 7240×7240 *R* of Boschi (2003). With 24840 free parameters, Cholesky factorization, backsubstitution (via the Lapack

routines SPOTRF and SPOTRS, respectively), and all necessary input/output from and to disc take about 10 h on a shared-memory Compaq «Alpha» computer (an ES45 with 10 Gb RAM and 4 CPUs at 1250 MHz). To compare this performance with that of repeated LSQR on a PC-cluster, it should be kept in mind that, in the latter architecture, computation time scales perfectly with the number of processors; one LSQR inversion with 24, 840 free parameters, and applying Paige and Saunders' (1982) criterion to evaluate convergence, currently takes ~1 min on a standard PC.



Fig. 3. Same as fig. 2, but R was computed by Cholesky factorization of $A^T \cdot A + D$, as described in Section 4.

Fig. 4. Rows of *R* (figs. 2 and 3) associated with a relatively well resolved voxel *i* located in the mantle under Japan, at 700 km mean depth, from the parallel LSQR (*left*) and Cholesky (*right*) approaches. For each value of *j*, the color of the *j*-th voxel depends on the value of R_{ij} ; R_{ij} is a measure of fictitious tradeoff between *i*-th and *j*-th model parameters.



Fig. 5. Same as fig. 4, but voxel *i*, less well resolved, is located under Central America, at 2200 km mean depth.

5. The covariance matrix

R describes the fictitious coupling between solution coefficients (model parameters); it depends on the geographic distribution of sources and stations, and on the shape of ray paths, but not on the quality of inverted observations. The covariance of solution coefficients depends, instead, on the error and covariance of the initial data, and on the error amplification occuring in the inversion (Menke, 1989, Section 3.11). In the assumption that seismic data be uncorrelated and all have equal variance σ^2 , Menke (1989, eq. 3.48) introduces a covariance matrix

$$C = \sigma^2 (A^T \cdot A + D)^{-1} \cdot A^T \cdot \left[(A^T \cdot A + D)^{-1} \cdot A^T \right]^T$$
(5.1)

(the regularization matrix D was not included explicitly in Menke's (1989) formula). Equation (5.1) can be rewritten

$$C = \sigma^{2} (A^{T} \cdot A + D)^{-1} \cdot A^{T} \cdot A [(A^{T} \cdot A + D)^{-1}]^{T}$$
(5.2)

and making use of (2.1)

$$C = \sigma^2 R \cdot [(A^T \cdot A + D)^{-1}]^T.$$
 (5.3)

After Cholesky factorizing $A^T \cdot A + D$, we find *C* by i) backsubsitution of the $N \times N$ identity matrix, and ii) dotproduct (via the Lapack routine SGEMM) of the (transposed) result with *R*. After *R* is read from disc or calculated again, the process takes about 10 more hours, with N = 24840, on the shared-memory machine described in Section 4 above. Figure 6 shows C, derived in the same parameterization (5°-voxels) and regularization as figs. 2 through 5 above, and assuming for Antolik *et al.*'s (2003)



Fig. 6. 24840×24840 (5°-voxel grid) covariance matrix *C* associated with the same data, parameterization and regularization as *R* above, defined as in Section 5, and derived by Cholesky factorization and backsubstitution.



0.00 0.02 0.04 0.06 0.08 0.10 0.12 0.14

Fig. 7. Absolute error on mapped percent *P*-velocity heterogeneity, calculated from the diagonal entries of *C* (fig. 6) and plotted at each corresponding model voxel. All 15, ~200 km thick layers of the 5°-voxel grid are shown; the shallowest layer is at the top and to the left, the plot below corresponds to the second shallowest layer, and so on; the deepest layer (~2700 km depth to coremantle boundary) is at the bottom and to the right. Constant, uncorrelated variance σ =0.5 is assumed on all traveltime observations.

database a standard deviation $\sigma = 0.5$ s (Antolik, 2005, pers. comm.).

We have not found an effective approach to calculating C on a distributed-memory cluster.

Except for the upper mantle, where the solution is less stable owing to strong nonuniformities in the data coverage, C (fig. 6) is relatively close to diagonal, indicating that errors in model parameters are not strongly correlated. The N diagonal entries of C can be interpreted as squared errors associated with the corresponding model parameters (Menke, 1989; Tarantola, 2005): after taking their square root and multiplying it by 100 (in a voxel parameterization, solution coefficients coincide with velocity heterogeneities in the corresponding voxel, which are typically expressed in percent), we show in fig. 7 each diagonal entry of C at the corresponding voxel. As to be expected, error is smallest in regions of good data coverage, e.g., the upper and midmantle underlying North America and Eurasia, where seismic stations are most densely distributed; it is highest at the top of the upper mantle, where the almost vertical geometry of teleseismic ray paths poses a significant limit to resolution (hence strong «smearing»); it grows with increasing depth in the bottom layers of the lower mantle, sampled more uniformly than shallowest regions, but by a decreasing number of ray paths. Mapped Pvelocity anomalies from the observations considered here (e.g., Boschi and Dziewonski, 1999; Boschi, 2003) range between ±1% in most of the mantle, so that the error of $\pm 0.15\%$ or less that we have derived from C is generally nonnegligible, but small.

6. Summary

We have presented two approaches to the solution of large mixed-determined inverse problems, both exploiting the quickly increasing speed and RAM of modern computers (*e.g.*, Bunge and Tromp 2003). We have verified (Section 4) that the two approaches, applied to the same problem, yield coincident results.

The first approach, described in Section 2, is inherently sequential, and is best applied to shared-memory computers. It rests on the direct implementation of the least squares formula to derive the least squares solution x_{LS} and model resolution matrix R associated with the inverse problem $A \cdot x = d$. x_{LS} and R are computed with one Cholesky factorization of $A^T \cdot A$ and N+1repetitions of the backsubstitution process, N being the number of model coefficients. In this approach, the derivation of R is thus relatively fast. Unfortunately, the number of floating point operations required to Cholesky-factorize A^T : A grows like N^3 , as finer parameterizations are implemented (e.g., Trefethen and Bau, 1997, p. 175). Likewise, as N grows, the size of A^T : A grows like N^2 , and RAM can also become an issue: routines performing Cholesky factorization, available in the literature (e.g., Press et al., 1994) or through optimized libraries, do not allow for efficient storage of $A^T \cdot A$ (which could be quite sparse), and require a comparably large additional amount of RAM to be left free for temporary storage.

The second approach (Section 3) involves the repeated application of an iterative. CG-type algorithm (LSQR in our implementation). x_{LS} is found after one run of LSQR, the calculation of *R* requires *N* runs of the same algorithm. However, we have shown that the problem can be simply parallelized, and it is thus most appropriate for implementation on distributed-memory PCclusters. As N grows, the growth in the number of floating-point operations will not be as fast as in the case of the first approach, and it will be relatively cheap to speed up the process by simply making use of a few more processors. As long as A is sparse, which is always the case in the raytheory approximation and with local-basis-function (voxels, splines, ...) parameterizations, the amount of disc space (and/or RAM) needed to store A also grows more slowly with increasing N than that needed store $A^T \cdot A$: in the experiment discussed in Section 3, A occupies roughly 1 Gb of RAM or disc space; $A^T \cdot A$ would need twice this amount, plus the temporary storage mentioned above.

One drawback of the CG/LSQR multiprocessor approach resides in the difficulty of computing the covariance matrix C. We have shown in Section 5 how C can instead be computed via the first, «sequential» approach, and we have made use of C to evaluate model error in typical, global tomographic inversions of seismic travel time observations (fig. 7).

In summary, both approaches should prove profitable, depending on the available hardware. The optimization of tomographic algorithms for use with modern computers is leading to a better understanding of the tomographic inverse problem, and to more reliable evaluations of model quality and resolution.

Acknowledgements

Our research is part of SPICE (Seismic wave Propagation and Imaging in Complex media: a European network), a Marie Curie Research Training Network in the 6th Framework Program of the European Commission; we are grateful to the coordinators at LMU Munich. We thank Domenico Giardini for his support and encouragement. All figures were done with GMT (Wessel and Smith, 1991).

REFERENCES

- ANTOLIK, M., Y.J. GU, G. EKSTRÖM and A.M. DZIEWONSKI (2003): J362D28: a new joint model of compressional and shear velocity in the mantle, *Geophys. J. Int.*, **153**, 443-466.
- BOSCHI, L. (2003): Measures of resolution in global body wave tomography, *Geophys. Res. Lett.*, **30** (19), 1978, doi: 10.1029/2003GL018222.
- BOSCHI, L. and A.M. DZIEWONSKI (1999): «High» and «Iow» resolution images of the Earth's mantle implications of different approaches to tomographic modeling, J. Geophys. Res., 104, 25,567-25,594.
- BOSCHI, L., T.W. BECKER, G. SOLDATI and A.M. DZIEWONS-KI (2006): On the relevance of Born theory in global seismic tomography, *Geophys. Res. Lett.*, 33, L06302, doi: 10.1029/2005GL025063.
- BUNGE, H.P. and J. TROMP (2003): Supercomputing moves

to universities and makes possible new ways to organize computational research, *Eos, Trans. Am. geophys. Un.*, **84**, 30-33.

- DZIEWONSKI, A.M. (1984): Mapping the lower mantle: determination of lateral heterogeneity in P velocity up to degree and order 6, J. Geophys. Res., 89, 5929-5952.
- GRAND, S.P. (1994): Mantle shear structure beneath the Americas and surrounding oceans, J. Geophys. Res., 99, 11,591-11,621.
- HAGER, B.H. and R.W. CLAYTON (1989): Constraints on the structure of mantle convection using seismic observation, flow models, and the geoid, in *Mantle Convection Plate Tectonics and Global Dynamics*, edited by W.R. PELTIER (Gordon and Breach, Newark, NJ), 657-763.
- HANSEN, P.C. (1992): Analysis of discrete ill-posed problems by means of the *L*-curve, *SIAM Review*, 34, 561-580.
- INOUE, H., Y. FUKAO, K. TANABE and Y. OGATA (1990): Whole mantle P wave travel time tomography, *Phys. Earth Planet. Inter.*, **59**, 294-328.
- LÉVÊQUE, J.J., L. RIVERA and G. WITTLINGER (1993): On the use of the checkerboard test to assess the resolution of tomographic inversions, *Geophys. J. Int.*, **115**, 313-318.
- MENKE, W. (1989): Geophysical Data Analysis: Discrete Inverse Theory (Academic, San Diego), revised edition.
- MINKO, S.E. (1996): A computationally feasible approximate resolution matrix for seismic inverse problems, *Geophys. J. Int.*, **126**, 345-359.
- NoLET, G. (1985): Solving or resolving inadequate and noisy tomographic systems, *J. Comput. Phys.*, **61**, 463-482.
- NOLET, G., R. MONTELLI and J. VIRIEUX (1999): Explicit, approximate expressions for the resolution and *a posteriori* covariance of massive tomographic systems, *Geophys. J. Int.*, **138**, 36-44.
- NOLET, G., R. MONTELLI and J. VIRIEUX (2001): Reply to comment by Z.S. YAO, R.G. ROBERTS and A. TRYGGVA-SON on «Explicit, approximate expressions for the resolution and *a posteriori* covariance of massive tomographic systems», *Geophys. J. Int.*, 145, 315.
- PAIGE, C.C. and M.A. SAUNDERS (1982): LSQR: an algorithm for sparse linear equations and sparse least squares, ACM Trans. Math. Soft., 8, 43-71.

- PRESS, W.H., S.A. TEUKOLSKY, W.T. VETTERLING and B.P. FLANNERY (1994): Numerical Recipes in FORTRAN (Cambridge University Press, U.K.).
- SOLDATI, G. and L. BOSCHI (2004): Whole Earth tomographic models: a resolution analysis, *Eos, Trans. Am. Geophys. Un.*, **85** (47), Fall. Meet. suppl, Abtsr. s13D-1092.
- SU, W.J., R.L. WOODWARD and A.M. DZIEWONSKI (1994): Degree12 model of shear velocity heterogeneity in the mantle, J. Geophys. Res., 99, 4945-4980.
- TARANTOLA, A. (2005): Inverse Problem Theory and Model Parameter Estimation (SIAM, Philadelphia).
- TREFETHEN, L.N. and D. BAU III (1997): Numerical Linear Algebra (Soc. Ind. and Appl. Math., Philadelphia, Penn.).
- VAN DER HILST, R.D., S. WIDIYANTORO and E.R. ENGDAHL (1997): Evidence for deep mantle circulation from global tomography, *Nature*, **386**, 578-584.
- VASCO, D.W., L.R. JOHNSON and O. MARQUES (2003): Resolution, uncertainty, and whole Earth tomography, J. Geophys. Res., 108, 2022, doi: 10.1029/2001JB000412.
- WESSEL, P. and W.H.F. SMITH (1991): Free software helps map and display data, *Eos, Trans. Am. Geophys. Un.*, 72, 445-446.
- WOODHOUSE, J.H. and A.M. DZIEWONSKI (1984): Mapping the upper mantle: three-dimensional modeling of Earth structure by inversion of seismic waveforms, J. Geophys. Res., 89, 5953-5986.
- YAO, Z.S., R.G. ROBERTS and A. TRYGGVASON (1999): Calculating resolution and covariance matrices for seismic tomography with the LSQR method, *Geophys. J. Int.*, 138, 886-894.
- YAO, Z.S., R.G. ROBERTS and A. TRYGGVASON (2001): Comment on «Explicit, approximate expressions for the resolution and *a posteriori* covariance of massive tomographic systems» by G. NOLET, R. MONTELLI and J. VIRIEUX, *Geophys. J. Int.*, **145**, 307-314.
- ZHANG, J. and G.A. MCMEHAN (1995): Estimation of resolution and covariance for large matrix inversions, *Geophys. J. Int.*, **121**, 409-426.

(received December 10, 2005; accepted June 5, 2006)