# A GLOBAL 3D P-VELOCITY MODEL OF THE EARTH'S CRUST AND MANTLE FOR IMPROVED EVENT LOCATION

Michael L. Begnaud[1], Sanford Ballard[2], Christopher J. Young[2], James R. Hipp[2], Marcus C. Chang[2], Andre V. Encarnacao[2], Charlotte A. Rowe[1], W. Scott Phillips[1], and Lee K. Steck[1]

[1]Los Alamos National Laboratory and [2]Sandia National Laboratories

Sponsored by the National Nuclear Security Administration

Award Nos. DE-AC52-06NA25396/LA09-IRP-NDD02[1] and DE-AC04-94AL8500/SL09-3D_Earth-NDD02[2]

## ABSTRACT

To test the hypothesis that high quality 3D Earth models will produce seismic event locations that are more accurate and more precise than currently used 1D and 2/2.5D models, we are developing a global 3D P wave velocity model of the Earth's crust and mantle using seismic tomography. In this paper, we present the most recent version of our model, SALSA3D (SAndia LoS Alamos 3D) version 1.7, and demonstrate its ability to reduce mislocations for a large set of realizations derived from a carefully chosen set of globally-distributed ground truth events, compared to existing models and/or systems. Our model is derived from the latest version of the Ground Truth (GT) catalog of P and Pn travel time picks assembled by Los Alamos National Laboratory. To prevent over-weighting due to ray path redundancy and to reduce the computational burden, we cluster rays to produce representative rays. Reduction in the total number of ray paths is ~50%. The model is represented using the triangular tessellation system described by Ballard et al. (2009), which incorporates variable resolution in both the geographic and radial dimensions. For our starting model, we use a simplified layer crustal model derived from the NNSA Unified model in Eurasia and Crust 2.0 model everywhere else, over a uniform *ak135* mantle. Sufficient damping is used to reduce velocity adjustments so that ray path changes between iterations are small.

We obtain proper model smoothness by using progressive grid refinement, refining the grid only in areas where the data warrant it. In previous versions of SALSA3D, we based this refinement on velocity changes from previous model iterations. For version 1.7, we utilize the diagonal of the model resolution matrix to control where grid refinement occurs, resulting in more consistent and continuous areas of refinement than before. In addition to the changes in grid refinement, we also employ a more robust convergence criterion between successive grid refinements, allowing a better fit to first broader model features, then progressively to finer ones. Our approach produces a smooth, multi-resolution model with node density appropriate to both ray coverage and the velocity gradients required by the data. This scheme is computationally expensive, so we use a distributed computing framework based on the Java Parallel Processing Framework, providing us with ~400 processors.

We compare the travel-time prediction and location capabilities of SALSA3D to standard 1D and 2/2.5D models via location tests on a global event set with GT of 5 km or better. These events generally possess hundreds of Pn and P picks from which we generate different realizations of station distributions, yielding a range of azimuthal coverage and ratios of teleseismic to regional arrivals, with which we test the robustness and quality of relocation. The SALSA3D model reduces mislocation over the standard 1D *ak135* model regardless of Pn to P ratio, with the improvement being most pronounced at higher azimuthal gaps. SALSA3D also reduces mislocation compared to the combined RSTT/*ak135* model (2.5D – RSTT for regional phases), with SALSA3D and RSTT performing about the same when using only Pn arrivals in location tests. We currently are testing the use of the full model covariance matrix in order to produce realistic path-dependent travel time uncertainty during location tests, replacing the standard distance-dependent, path-independent uncertainty typically used in location algorithms.

## OBJECTIVES

The evolving U.S. monitoring needs require accurate location of ever smaller events. Given the limited station coverage of a typical global monitoring network, such as the International Monitoring System (IMS), it is likely that such events will be detected by very few stations, likely with poor network geometry. For locating such events, it is essential to have extremely high fidelity travel-time predictions, particularly at regional distances where lateral heterogeneity can be significant. Accurately accounting for lateral heterogeneity implies using 3D models of the Earth to calculate travel times, but relatively few of the available 3D Earth models are appropriate for high fidelity travel-time prediction, and it is unclear whether any of those that are appropriate actually improve event locations. Rowe et al. (2009) demonstrated that when 3D models are used to predict travel times in a manner inconsistent with how the model was developed, the resulting values tend to have higher residuals than those produced when using standard 1D velocity models. This is not surprising in itself, but does indicate problems utilizing contributed models that have vastly different assumptions and/or ray tracing methodologies. It is crucial that a seismic location algorithm used with a 3D model also use the same ray tracer as during model development. Thus, we are developing our own global 3D P-wave velocity model of the Earth's crust and mantle—SALSA3D using seismic tomography based on a carefully assembled data set of P-phase travel times collected by LANL over the past decade. An important difference between our effort and previous efforts is that our model was produced specifically for improving event location; the model provides potentially valuable information about the structure of the Earth, but this is not our focus. Hence, all decisions about data processing and tomography were made with this goal in mind.

In this paper, we present the most recent version of our model, SALSA3D version 1.7, and demonstrate its ability to reduce mislocations for a large set of realizations derived from a carefully chosen set of distributed GT events.

## RESEARCH ACCOMPLISHED

### Data Set

The data used for the tomographic inversion was collected by LANL over the past decade and includes events categorized with a GT level of 25 km or better (Bondár et al., 2004). Arrivals for these events were merged from various sources to produce a single set of arrivals for each event with redundancies removed, maximizing the available arrivals (Begnaud, 2005). Regional Pn and teleseismic P arrivals were extracted for use in the tomography and resulted in ~118K events, ~12K stations, and almost 12 million individual ray paths (Figure 1). Unique data sets include USArray data and the Deep Seismic Soundings in the former Soviet Union (e.g., Li and Mooney, 1998).

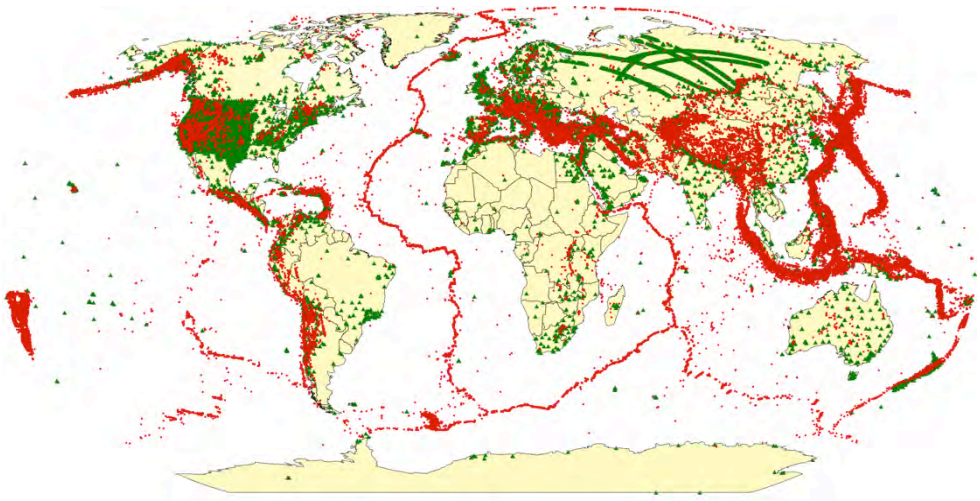

**Figure 1. Stations (green triangles) and events (red points) used as the raw data for global tomography inversion. These data were collected by LANL over the past decade and are all categorized with a ground truth level of 25 km or better. Recent additions to this data set include the mid-western portion of USArray as well as several PASSCAL deployments in Asia.**

Summary rays are produced from the ~12 million rays to remove redundant ray paths, search for data outliers, and reduce the computational load of ray tracing required during tomographic inversion. Each unique station/GT level

combination is clustered separately, preserving GT levels and their relative weighting (lower GT levels are weighted more heavily), giving ~7 million summary rays.

**Tomographic Inversion**

We use the LSQR algorithm of Paige and Saunders (1982) to invert the data to find the P wave velocity distribution in the mantle that optimally reduces the misfit between observed and predicted travel times for P and Pn phase arrivals. Our starting model for the inversion consists of the NNSA Unified Model for Eurasia (Pasyanos et al., 2004; Steck et al., 2004) with the Crust 2.0 model elsewhere (Laske et al., 2000; Laske and Masters, 1997) modified by combining their six layers into four by merging the three sedimentary layers into one. The crust is damped heavily in the inversion so it is essentially fixed.

The full iterative inversion is run three separate times, relocating the GT>5 events with the SALSA3D model after each full inversion (including grid refinement below). After each relocation step (which we call outer tomography iterations), the tomographic inversion is restarted using the original starting model and the updated locations. This has the effect of improving the locations for events with higher GT levels (e.g., GT25) whose original locations were based on the *ak135* model (Kennett et al., 1995). This shifting of locations tends to occur mostly along subduction zones, where the locations migrate up-slab as the use of *ak135* model along with the commonly one-sided station distribution has pushed locations farther off-shore.

For the inner tomographic iterations, we retrace rays in each iteration using the pseudo-bending algorithm (Ballard et al., 2009; Um and Thurber, 1987; Zhao and Lei, 2004). During early iterations, the mantle is moderately damped to discourage the inversion from progressing too rapidly but then damping is relaxed in later iterations. Convergence of the inversion is achieved when successive iterations result in overall residual changes of $\leq 0.25\%$, refining to ~0.20% at finer scales of grid refinement (see below). Regularization is not required during the inversion since it is accomplished by refining the grid several times during the inversion (Simmons et al., 2009).

**Adaptive Grid Refinement**

P wave velocity in our model is stored on a 3D grid of nodes that has variable resolution in both geographic and radial dimensions. Our basic starting point is a 2 dimensional, multi-level triangular tessellation of a unit sphere (Ballard et al., 2009; Wang and Dahlen, 1995). We assign a tessellation with 8° triangles to the lower mantle, a tessellation with 4° triangles to the transition zone and upper mantle, and a third tessellation with variable resolution to all crustal layers. The crustal tessellation (not shown) has 2° triangles in oceanic regions, 1° triangles in most continental regions, and 0.5° triangles in the United States where the additional resolution is warranted due to the dense station coverage afforded by the USArray. The two tessellations in the mantle are refined, in both the geographical and radial dimensions, during the progress of the tomographic inversion.

Grid refinement was originally accomplished based on a minimum change in the slowness values during tomographic iterations. We have recently been able to solve for the full 3D resolution matrix (Hipp et al., 2011, these Proceedings) and applied the resolution matrix for a coarse grid to drive the grid refinement steps. By using the diagonal of the resolution matrix, we are able to take advantage of data coverage. Figure 2 demonstrates adaptive grid refinements steps for a cross section of the mantle.

**Velocity Model**

The final SALSA3D 1.7 model is shown in Figures 3 and 4. Although it is not our intent to resolve detailed tectonic features, many structural features present themselves. Clearly visible in the figures are subduction zones, mid-ocean ridges, slower velocities in western North America, including the Yellowstone hot spot, faster velocities in eastern North America associated with the craton, the Red Sea rift, detailed variations in Europe and the Middle East, and faster velocities in northern Russia and into Tibet and India. These features are consistent with previous 3D models (e.g., Li et al., 2006) in most areas. The full 3D model resolution matrix indicates that the best laterally-resolved areas correspond to natural areas of higher seismicity, of course, typically in the western United States, Europe, and into most of Asia. Resolution at depth is best in the upper portion of the lower mantle, due to the large number of teleseismic raypaths in our dataset that bottom near this depth.

**Reduction of Travel-Time Residuals**

When validating a new velocity model, the first step is usually to test the reduction in travel-time residuals, typically with distance. We plotted the median residual with distance (bias) along with the L1 norm spread or the median absolute difference from the median (MAD1) (Figure 5).

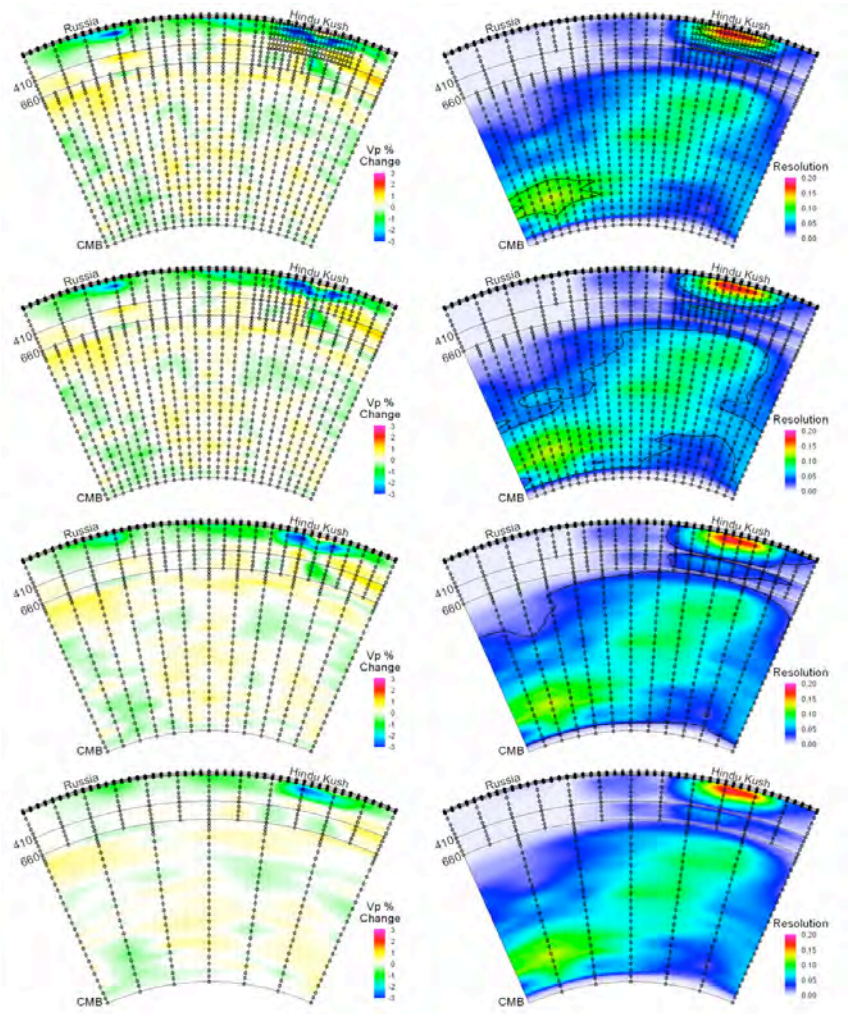


**Figure 2. Evolution of the model during grid refinement. Images on the left show the % change in P wave speed relative to *ak135* (Kennett et al., 1995) while those on the right show the diagonal elements of the resolution matrix obtained in an earlier tomography run. Symbols show the positions of grid nodes in the plane of the image which extends along the meridian at 72°E from 74.2°N to 26.7°N. Note the change in the distribution of the grid nodes as the model is refined from coarse (bottom) to fine (top).**

For the *ak135* model (Kennett et al., 1995), the bias is significant, especially for regional distances (~<25°), as well as the spread. For teleseismic distances, *ak135* still shows an observable bias, although greatly reduced over regional distances. When using the SALSA3D model, the bias is significantly reduced over *ak135*, most noticeably for regional distances, although some bias is still observed for triplication distances (~18-25°). The bias is almost zero for teleseismic distances when using SALSA3D. The spread of the residuals for SALSA3D is displaying a significant reduction over *ak135*. We tend to associate this reduction in global spread with distance actually to travel time bias reductions within tectonic regions, while real residual spread is more likely due to pick error which can't be corrected using a different model. The SALSA3D model reduces bias in tectonic regions and, thus, shows up in global residuals with distance as a reduction in residual spread.

**Reduction in Event Mislocation**

For monitoring purposes, the requirement for any new velocity model is the degree to which it reduces mislocation of seismic events with well-constrained locations. We are comparing location results to the standard *ak135* model, as well as the recently developed Regional Seismic Travel Time (RSTT) model (Myers et al., 2010). The RSTT model is designed for regional phases (e.g., Pn) and is combined with *ak135* when using teleseismic arrivals. If the SALSA3D model doesn't display location improvement over *ak135* and RSTT, we cannot justify the level of effort required to develop and implement 3D models for location in the context of nuclear explosion monitoring.

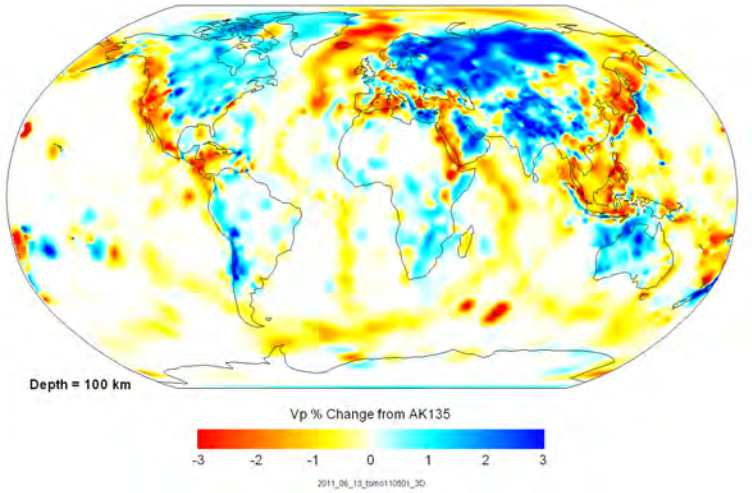


**Figure 3. Depth slice through the SALSA3D 1.7 model at 100 km demonstrating the Vp velocity structures produced. Colors are percent change from the *ak135* model. Clearly visible are tectonic features such as mid-ocean ridges, subduction zones, and large cratons.**

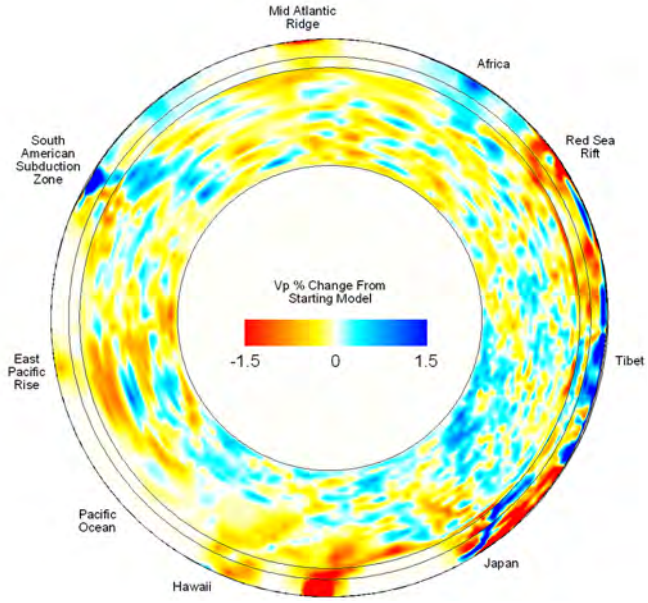


**Figure 4. Cross section through the SALSA3D 1.7 model in a pseudo-great circle pattern to highlight structures. Colors are percent change from the starting model (at this view, mostly the *ak135* model). Of particular notice is the delineation of the Japanese subduction zone, the thicker, faster Tibetan area, and the other subduction zones, rifts, and mid-ocean ridges.**

*Model/Prediction Uncertainty*

To run relocation tests, an uncertainty model must first be determined for SALSA3D. For standard location procedures, a simple residual spread with distance is used to estimate the prediction or model uncertainty for any arrival. The problem with this simple assumption is that it does not account for aseismic areas or areas that do not predict travel times well. To perform location tests with SALSA3D, however, we needed a standard prediction uncertainty model as a function of distance to use with the location code. Comparison of the spread with distance is a validation test to compare SALSA3D to other models and should show a reduction in spread with distance, as illustrated above.

When comparing these distance-dependent or prediction uncertainty curves, we follow these guidelines to ensure a fair comparison:

1. The distance-dependent uncertainty values should be created **using the same data for each model**. Comparing curves developed with different validation data is inconsistent.
2. Since most travel times (i.e., origin times) were determined using a 1D model such as *ak135*, **all travel times should be baselined to the models that are being compared**. The baselining is accomplished by adjusting origin times using average residuals per event. This is done so travel times are determined relative to the specific model being compared.
3. If possible, the data used for the uncertainty curves should not be used in the tomography models. For best results, the uncertainty curve with distance should be based on similar data that will eventually be used in practice with the model and uncertainty.

To follow guideline (3), a new SALSA3D model was determined but having ~5% of the raw arrival data removed (within the boundaries of the RSTT model for consistency). Summary ray production and relocation steps continued normally after this data removal. After tomography, the original origins were relocated through the alternative SALSA3D model and the removed arrival data were grouped into 1° distance bins to evaluate uncertainty (Figure 6). We use the same removed arrivals to evaluate both *ak135* and RSTT uncertainty; however, the removed data were not necessarily removed from the RSTT model determination, violating guideline (1) above. This could not be rectified for this analysis as the RSTT model was developed with a different data set base. The origin times were baselined to each model used. One disadvantage of using residual spread with distance as a proxy for uncertainty is that the residual spread consists of both pick errors and model errors. To isolate only the model errors, pick error must be removed. In a mini-study using a method similar to Flanagan et al. (2007), we extracted the median prediction uncertainty values from IMS stations and GT25 or better data to produce a prediction uncertainty estimate as a function of distance.

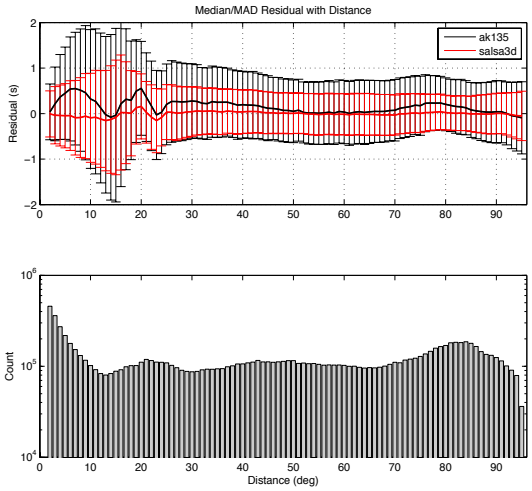


**Figure 5. Travel-time residuals with distance (top) showing the median residual (solid line) and L1 spread (error bars) for residuals binned in 1° bins. Data are those used in the tomography model. The *ak135* model displays significant bias, especially for regional data, while the SALSA3D reduces this bias to almost zero, with a corresponding reduction in residual spread. Bottom: counts of arrivals in each bin demonstrating about equal and adequate coverage over all distance ranges.**

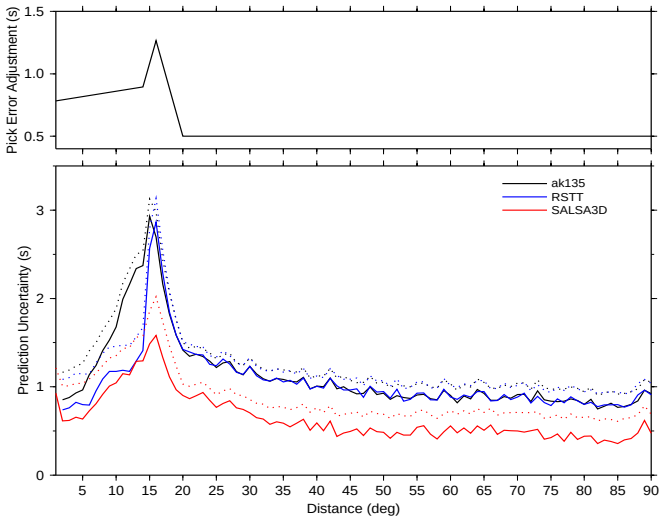


**Figure 6. Distance-dependent uncertainty using ~5% of the original SALSA3D data set left out of the 3D model. These data were not necessarily left out of the RSTT model, however, due to a different data set base. Dashed lines are raw uncertainty with solid lines being uncertainty after an estimate of pick error was removed.**

The adjusted uncertainty values for regional Pn distances (< 15°) show a significant reduction in uncertainty for SALSA3D over *ak135* and to a lesser degree over RSTT. The uncertainty associated with triplication distances (~15-20°) is reduced for the SALSA3D model. For teleseismic distances (> 20°), the SALSA3D uncertainty has decreased by almost 0.5 s.

*Location Validation Tests*

The paramount model validation is testing for improved location accuracy and precision or reduced mislocation. In order to compare location results between models, **it is best to use the same validation events among models**, if possible, that were left out of the tomographic inversions. By using similar validation events, any concerns over inconsistency of the events themselves are eliminated, as well as circularity issues. Comparisons are then strictly affected only by model differences.

To be consistent with RSTT, we created another SALSA3D model and left out the 52 Eurasian validation events used in the validation of the RSTT model (Myers et al., 2010) in the "Unified" area (lat 0-85°, lon -20-150°) (Figure 7). These events were chosen from GT0-5 events to maximize both coverage as well as the number of available P and Pn phases for location tests.

To test improvements in location using the SALSA3D model and compare with other models, a set of origins and associated arrivals has to be generated for location testing. Since we are interested in the combination of P and Pn arrivals and how different combinations affect the resulting location, we generate a grid of random combinations of P and Pn arrivals from **existing arrivals** of the validation events, using the LANL Location Database. The various combinations simulate different network configurations arising from a variety of possible event magnitudes.

A target set of P and Pn arrivals is determined from a total of 3 to 20 phases, and all of the available arrivals for an event are assigned a random number, then sorted by that random number. For the target number N phases, the first N phases are chosen from the sorted set. The process is repeated 100 times for each event in each target number of phases for a total of 5200 possible realizations per selected target phases (Figure 8). **We use 100 realizations per event as it provides a large number of origins and a valid statistical base. Using too few realizations per event can result in erratic looking mislocation grids that also show significant variations with different realization sets.** For our set of 100 realizations per event, we obtained 1,121,570 unique event realizations.

The random realization "events" were relocated using *ak135*, RSTT/*ak135*, and SALSA3D. No arrivals were eliminated during relocations unless the model was not defined for that particular ray path. The uncertainty-with-distance values from Figure 6 were used to estimate model uncertainty during relocation. The median values for each grid of target phases were determined along with the L1 spread for each grid position (Figure 9). For origin

realizations with very few arrivals (<~10), the SALSA3D model exhibits significant mislocation reduction over the other models.

Perhaps an easier method of comparing the mislocation results is to difference the plots and look at relative values from one model over another (Figure 10). When very few arrivals are used, the differences between SALSA3D and *ak135* are on the order of tens of kilometers improvement (numbers listed on grid positions in (Figure 10)).

Another validation test is to plot mislocation values as a function of azimuthal gap (Figure 11). Higher-dimensional models are generally more effective for reducing mislocation values in the presence of larger azimuthal gaps. For teleseismic P, SALSA3D reduces mislocation by over 5 km compared to *ak135*/RSTT for a 200° gap. Using only Pn, SALSA3D and RSTT again give about the same mislocation values.

Coverage ellipse area and percent are also used as measures of model improvement. During location, 95% coverage ellipses were calculated which directly depend on the *a priori* model uncertainty assumptions. For the median coverage ellipse area over the grid, SALSA3D shows improvement (2297 ± 912 $km^2$) over *ak135* (2706 ± 1050 $km^2$) and RSTT (2391 ± 977 $km^2$), with similar patterns as those for the mislocation by number of arrivals and by azimuthal gap. The median coverage ellipse areas are higher than the 1000 $km^2$ target for on-site inspection; however, these numbers include results for small numbers of phases. The point for this measure is to reduce the uncertainty ellipse size, which SALSA3D is doing, sometimes by several 100 square kilometers.

**Figure 7. RSTT and SALSA3D validation events in Eurasia.**

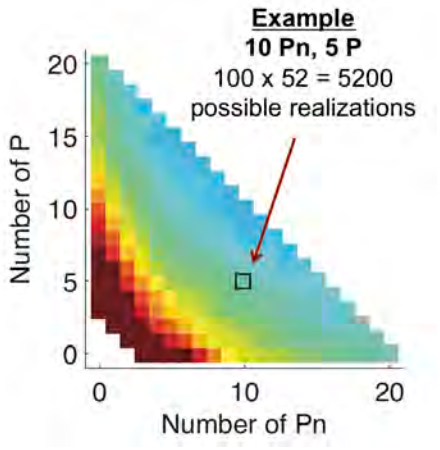


**Figure 8. Example of random realizations of phases, generating selections of origins for different network configurations and event magnitudes. Each grid position is a target number of P and Pn phases, with a total possible of 5,200 origin/realizations.**

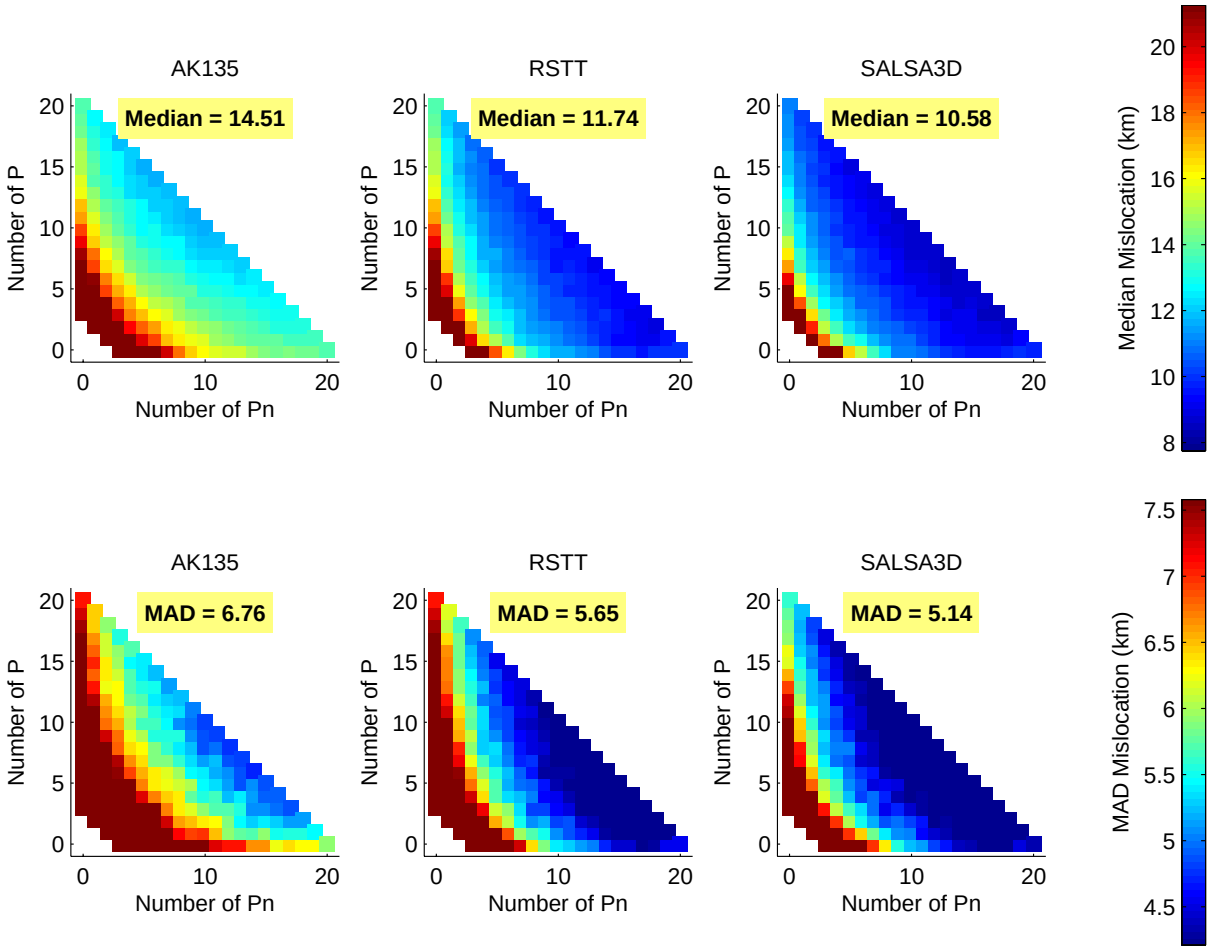


**Figure 9. Median and MAD1 mislocation values for each grid position of target phases and for each model tested (*ak135*, RSTT/*ak135*, SALSA3D). The SALSA3D model displays an overall reduction in mislocation by 26% over ak135 and 9% over RSTT/ak135 with significant reduction in mislocation spread.**

The coverage ellipse percentage reflects the fraction of the time the original GT location falls within the uncertainty ellipse of the test location. For each of the models, an uncertainty ellipse percentage grid was produced (Figure 12), similar to the mislocation grid structures shown above. The *ak135* results approach ~90%, largely for teleseismic arrival grid positions. For positions with more Pn arrivals, this drops to ~87% thus suggesting the regional uncertainty values are too low. RSTT percentages are much closer to the target percentage than *ak135*, with higher values for fewer teleseismic P and mid-range numbers of Pn. The SALSA3D model has a slightly higher average (94.1%) than RSTT (93.3%), but with a better spread, suggesting more consistent percentage results.

## CONCLUSIONS AND RECOMMENDATIONS

In this paper, we have presented the latest version (1.7) of SALSA3D, the Sandia-Los Alamos global tomography 3D P-velocity model of the crust and mantle. Our model is multi-resolution, with the resolution determined as part of the tomographic inversion. Additional nodes are added where ray coverage is high and the full 3D resolution matrix suggests sufficient resolution; other areas maintain the starting grid sampling. SALSA3D successfully images many well-established features within the Earth, such as subducting slabs and hot spots, and is generally consistent with other recent global and regional models.

More important than the structural features in the model is the fact that SALSA3D provides unambiguous improvement in travel-time prediction and in event location compared to *ak135* and, to a lesser extent, RSTT/*ak135*.

Event location improvement is evaluated by re-locating a set of 52 well-characterized events in Europe and Asia. By forming many different random realizations of these events, we can evaluate relocation improvement for a large number of different scenarios. SALSA3D improves locations to a lesser extent when the number of stations is large, but as the number of available phases becomes smaller, the 3D model performs significantly better. Further, because our model is developed to fit both P and Pn data, it performs well with regional data, teleseismic data, or a combination of the two, eliminating the requirement of using multiple models based on path length.

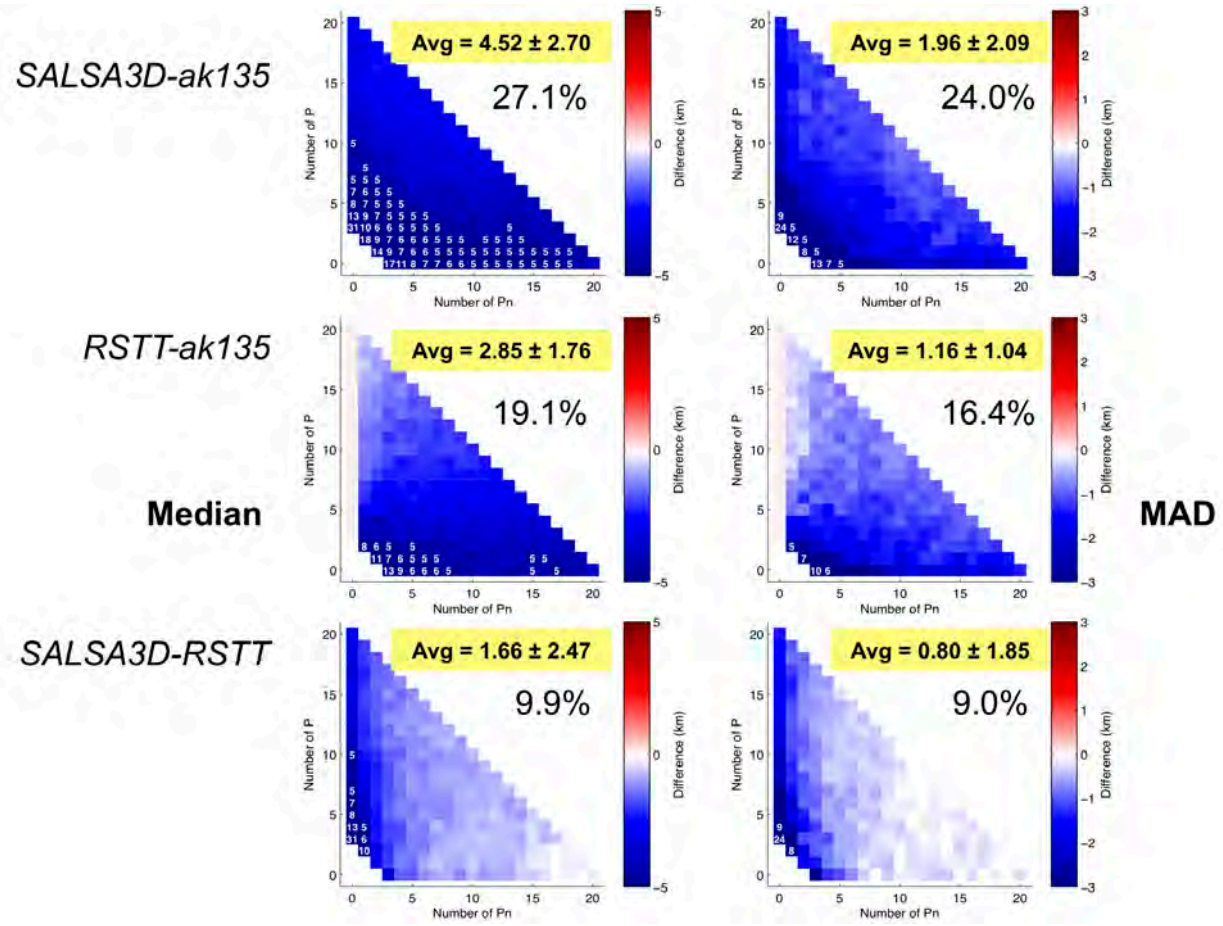


**Figure 10. Mislocation differences between the three models tested. The median and MAD1 values in each of the grid positions in Figure 9 have been differenced. The percentage shown is the overall percentage improvement of the first model compared to the second. Any grid position with a difference of 5 km or more has the integer value plotted on the grid position. At very small numbers of phases (<~10), SALSA3D shows a reduced mislocation by tens of kilometers compared to *ak135*.**

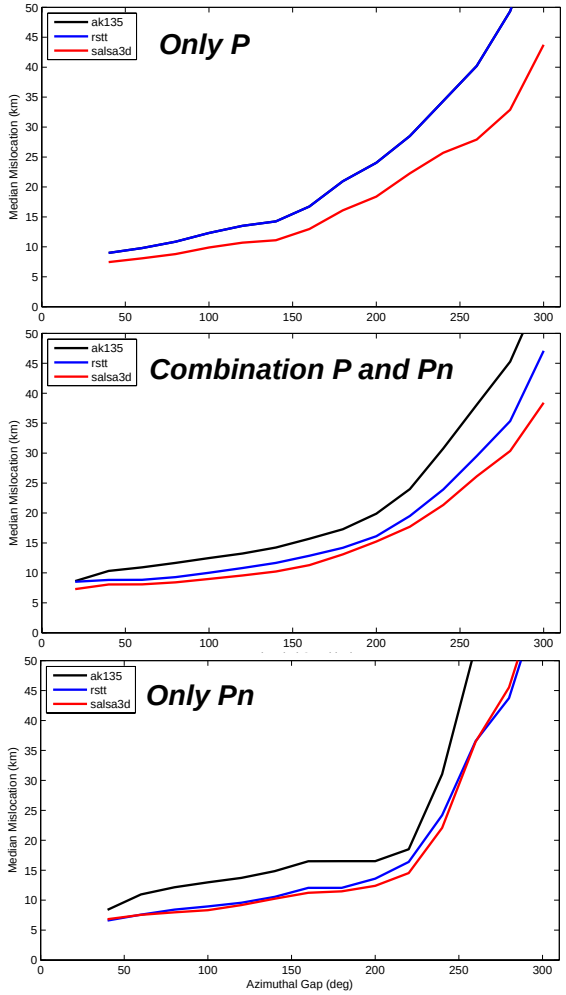


**Figure 11. Median mislocation relative to azimuthal gap and different combinations of P and Pn phases. SALSA3D shows improvement over other models for only P and combinations of P and Pn. For locations with only Pn phases, SALSA3D and RSTT/*ak135* exhibit about the same levels of mislocation, both improvements over strictly *ak135*.**

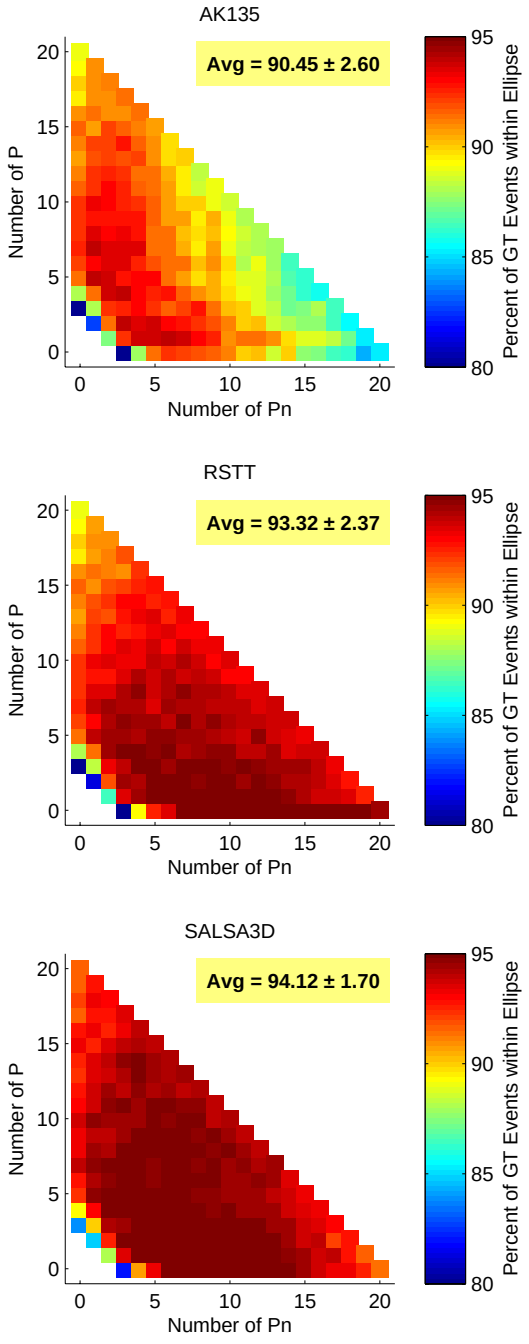


**Figure 12. 95% coverage ellipse percentages for ak135 (top), RSTT (middle), and SALSA3D (bottom). Each grid position displays the median ellipse percentage (percent of the time the original GT location falls within error ellipse of the test location).**

## REFERENCES


Ballard, S., J. R. Hipp, and C. J. Young (2009) . Efficient and accurate calculation of ray theory seismic travel time through variable resolution 3D earth models, *Seis. Res. Lett.* 80, 6: 989–998, doi:10.1785/gssrl.80.6.989.

Begnaud, M. L. (2005). Using a dedicated location database to enhance the gathering of ground truth information, *Los Alamos National Laboratory document* LA-UR-04-5992, 22 pp.

Bondár, I., S. C. Myers, E. R. Engdahl, and E. A. Bergman (2004). Epicentre accuracy based on seismic network criteria, *Geophys. J. Int.* 156: 483–496.

Flanagan, M. P., S. C. Myers, and K. D. Koper (2007). Regional travel-time uncertainty and seismic location improvement using a three-dimensional a priori velocity model, *Bull. Seismol. Soc. Am.* 97, 3: 804–825.

Kennett, B. L. N., E. R. Engdahl, and R. Buland (1995). Constraints on seismic velocities in the Earth from travel times, *Geophys. J. Int.* 122: 108–124.

Laske, G., G. Masters, and C. Reif (2000). CRUST 2.0: A New Global Crustal Model at 2x2 Degrees, http://igppweb.ucsd.edu/~gabi/crust2.html.

Laske, G. and T. G. Masters (1997). A global digital map of sediment thickness, Abstract F483 presented at Fall Meeting, American Geophysical Union, San Francisco, California,

Li, C., R. D. van der Hilst, and M. N. Toksöz (2006). Constraining P-wave velocity variations in the upper mantle beneath Southeast Asia, *Phys. Earth Planet. Inter.* 154: 180–195.

Li, S. and W. D. Mooney (1998). Crustal structure of China from deep seismic sounding profiles, in *Proc. of the 7th international symposium on deep seismic profiling of the continents*, Asilomar, CA, United States, Tectonophysics 288, 105–113.

Myers, S. C., M. L. Begnaud, S. Ballard, M. E. Pasyanos, W. S. Phillips, A. L. Ramirez, M. S. Antolik, K. D. Hutchenson, G. S. Wagner, J. J. Dwyer, C. A. Rowe, and D. R. Russell (2010). A crust and upper mantle model of Eurasia and North Africa for Pn travel time calculation, *Bull. Seismol. Soc. Am.* 100: 2, 640–656.

Paige, C. C., and M. A. Saunders (1982). LSQR: An algorithm for sparse linear equations and sparse least squares, *ACM Trans. Math. Softw.* 8, 1: 43–71.

Pasyanos, M. E., W. Walter, R., M. P. Flanagan, P. Goldstein, and J. Bhattacharyya (2004). Building and testing an a priori geophysical model for western Eurasia and North Africa, *Pageoph* 161: 235–281.

Rowe, C., S. Ballard, M. Begnaud, C. Young, L. Steck, and J. Hipp (2009). Validating 3D geophysical models for use in global travel-time calculation for improved event locations, in *Proceedings of the 2009 Monitorning Research Review: Ground-Based Nuclear Explosion Monitoring Technologies*, LA-UR-09-05276, Vol. 1, pp. 408–415.

Simmons, N. A., S. C. Myers, and A. L. Ramirez (2009). Multi-resolution seismic tomography based on a recursive tessellation hierarchy, in *Proceedings of the 2009 Monitoring Research Review: Ground-Based Nuclear Explosion Monitoring Technologies*, LA-UR-09-05276, Vol. 1, pp. 211–220.

Steck, L. K., C. A. Rowe, M. L. Begnaud, W. S. Phillips, V. L. Gee, and A. A. Velasco (2004). Advancing seismic event location through difference constraints and three-dimensional models, in *Proceedings of the 26th Seismic Research Review - Trends in Nuclear Explosion Monitoring*, LA-UR-04-5801, Vol. 1, pp. 346–355.

Um, J. and C. Thurber (1987). A fast algorithm for 2 point ray tracing, *Bull. Seismol. Soc. Am.* 77: 972–986.

Wang, Z. and F. A. Dahlen (1995). Spherical-spline parameterization of three-dimensional Earth models, *Geophys. Res. Lett.* 22: 3099–3102.

Zhao, D. and J. Lei (2004). Seismic ray path variations in a 3D global velocity model, *Phys. Earth Planet. Inter.* 141: 153–166.