IMPROVED REPRESENTATION AND CALCULATION OF BASE MODEL TRAVEL TIMES USING THE PARAMETRIC GRID LIBRARY

Paul C. Reeves, Sanford Ballard, James R. Hipp, Christopher J. Young,
Daniel S. Myers, and Breanna Ammons

Sandia National Laboratories

Sponsored by National Nuclear Security Administration
Office of Nonproliferation Research and Engineering
Office of Defense Nuclear Nonproliferation

Contract No. DE-AC04-94AL85000

ABSTRACT

Efficient and accurate event location is critical to nuclear event monitoring, during both the event association processing and event characterization stages. Computation of predicted (i.e., base model) travel times, slowness values, and their spatial derivatives is a core component of these calculations. Common base model algorithms represent travel time using a standard, rectilinear, two-dimensional (distance, depth) lookup table for each seismic phase. The values used are typically computed as a preprocess using ray theory, with dummy values inserted into the table at distance/depth points beyond where the seismic phase is observed.

Linearized least-squares event location algorithms require computing the expected travel times for each observed phase at each iteration of the hypocenter calculation. Using the lookup tables, travel time values are computed using a high-order interpolation scheme so spatial derivatives of travel time can be computed. These are needed by the location algorithm to formulate the gradient of the predicted travel time with respect to event location, which is used to iteratively move towards a minimum-residual solution. Second derivatives of travel time with respect to event location are needed whenever slowness observations are used. When a hypocenter lies outside the valid distance/depth range of a given phase, then extrapolation of information from valid regions is required. Together, these calculations comprise the vast majority of the computational effort involved in seismic event location (See Hipp et al., 2005).

We present results demonstrating the potential weakness of this method if spatial discretization is too coarse. We also show the danger of computing slowness from the numerical derivative of travel time with respect to distance. These problems arise due to discontinuities in the travel time surface where different branches cross and/or where different phases arise.

We present two alternate representations to standard, rectilinear lookup tables, that overcome the problems that arise from discontinuities. They represent the two-dimensional lookup table on an optimized, irregular, triangular grid whose density is proportional to the curvature of the travel-time field. One uses a single gridded region to represent a particular seismic phase, while the other uses sub-regions representing particular branches or phases in order to eliminate discontinuities. This representation should provide fast numerical performance while achieving high accuracy.

We also advocate pre-computing all relevant values (travel time, travel time error, slowness, *dtt/dz*, *dsh/dz*, *dsh/d*Δ, and model error) as a preprocess, including extrapolation beyond the valid region for the phase. This will allow simple bilinear interpolation and minimize significantly the amount of on-the-fly calculations, both of which will increase speed. The preprocessing can be done at high resolution to provide improved accuracy.

OBJECTIVES

Fast earthquake location algorithms are required for nuclear explosion monitoring. These are typically supported with simple one-dimensional earth velocity models such as IASPEI91 (Kennett and Engdahl, 1991) or ak135 (Kennett et al., 1995). Calculations of predicted travel time values are facilitated with two-dimensional (distance-depth), rectilinear lookup tables for each phase of interest (e.g., LocSAT [Bratt and Bache, 1986; Nagy, 1996]). The alternative, ray tracing or ray theory methods (e.g., TauP [Crotwell et al., 1999]), are currently too slow for operational purposes.

Lookup tables offer both representational and numerical advantages. They are easy to use and can require negligible memory if discretization is coarse. For gradient descent location algorithms, a four-by-four patch of neighboring points may be used to support bicubic interpolation which yields travel time and its first and second spatial derivatives. If slowness observations are used in the solution formulation, then slowness predictions are often supplied as the numerical derivative of travel time with respect to distance.

Despite these advantages, this representation poses significant risk if spatial discretization is too coarse. Fine discretization can address accuracy problems, but degrades computational speed and inefficiently utilizes memory. We discuss these issues below and present an alternative representation that could provide both accuracy, efficiency, and speed through optimal tessellation and natural neighbor interpolation using the Parametric Grid Library (Hipp et al., 1999, 2005)

RESEARCH ACCOMPLISHED

Figure 1 shows computed travel-times for the “first” P phase throughout its defined source depth and distance range based on the IASPEI91 velocity model (Kennett and Engdahl, 1991). The data for the plot was generated using TauP Toolkit (Crotwell et al., 1999), which is commonly used to populate travel time lookup tables. The surface appears to be smooth and well behaved except at the jump between Pdiff and PKIKP at 114º; however, there are subtle discontinuities along any given depth where different branches intersect. These can be seen in Figure 2, which presents a complete set of P-phase travel time curves for an event at 28 km depth. There are triplications arising from the velocity discontinuities at 410 and 660 km depth that show up near 20º, but in practice they are not distinguished since they are too difficult to pick on a seismogram. The discontinuity at the Pdiff – PKIKP transition is apparent, and there is also a small range near 145º where the PKPab branch arrives before PKIKP (a.k.a. PKPdf). The “first” P table (Figure 1) uses the earliest arriving phase. Pdiff is an exception due to its weak signal, so PKIKP and PKPab are substituted for Pdiff where they exist.

Representing “first” P phase travel times in a lookup table would appear benign at first glance. However significant errors can result if spatial resolution is too coarse. Figure 3 and 4 show the difference between the look up table-predicted and TauP travel times using a common spatial discretization for the lookup table (black points) and libloc’s bicubic interpolation algorithm (Nagy, 1996; Nagy, 2001). The algorithm does an excellent job at the grid points, as expected, with errors increasing in proportion to the distance from grid support. This yields a “quilted” appearance. Discrepancies between the interpolated and TauP travel times become substantial near triplications caused by velocity discontinuities in the IASPEI91 velocity model at depths of 210, 410 and 660 km. These discrepancies are evident in the distance range of 10° to 30° in Figures 3 and 4. These anomalies occur because the supporting four-by-four grid straddles the triplication and is unable to accurately represent the discontinuous derivative of travel time with respect to position at the triplication.

While the problems at depth may be of lesser concern to the monitoring community, there are still problematic areas near the surface. Figure 4 is a detail of Figure 3 at shallower depth that shows how travel time prediction errors persist even within the crust, exceeding ±1 second in some areas. The largest errors occur near zero distance where the upgoing p is the first arrival. The discretization fails to capture its concave shape and the discontinuity with downgoing P. The Pdiff – PKIKP transition at distances of 114° is also problematic, as expected, as is the area where PKPab arrives first (around 148°). These problems are all symptomatic of the coarse discretization, which is, in fact, much finer at the shallow depths shown here.

Figures 5 and 6 show details of the travel time curve (Figure 2) together with predicted values where the p phase arrives first (Figure 5) and where the P phase triplications arrive (Figure 6). The 1º table discretization completely misses the p phase (Figure 5), so the predictions fail to capture its concave shape at distances less than about 0.7°. Figure 6 shows the how the discretization and cubic spline undershoot the triplications within the P phase.

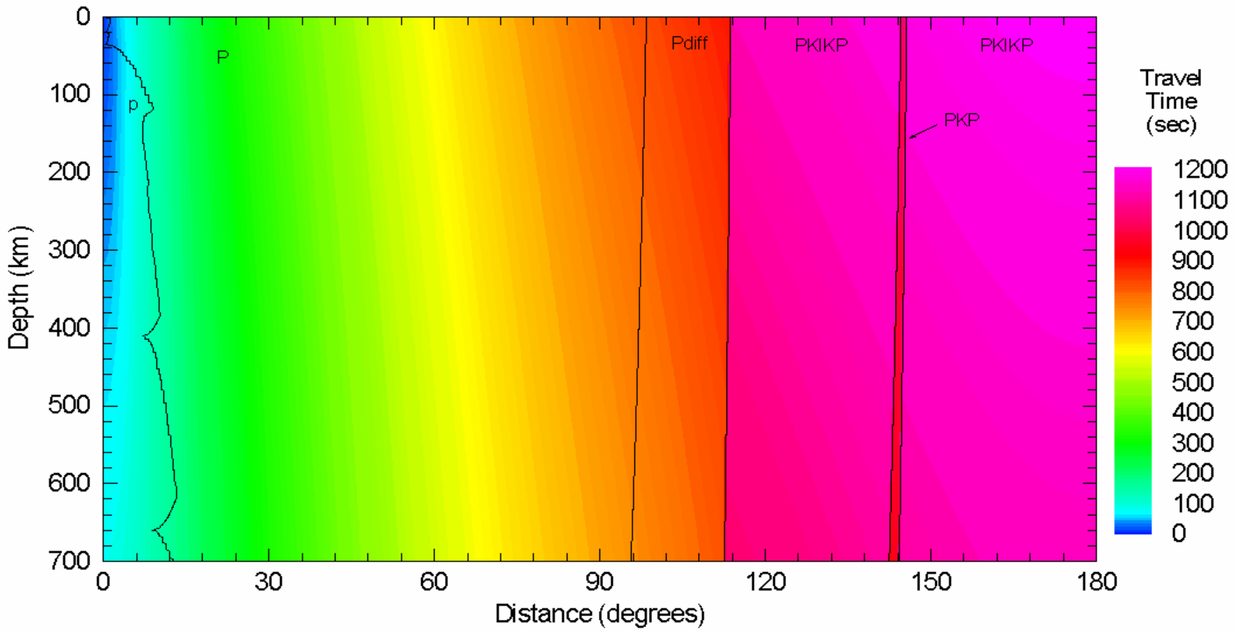

Figure 1. "First" P-phase travel times computed with TauP. Spatial resolution is 0.36º in distance and 5 km in depth.

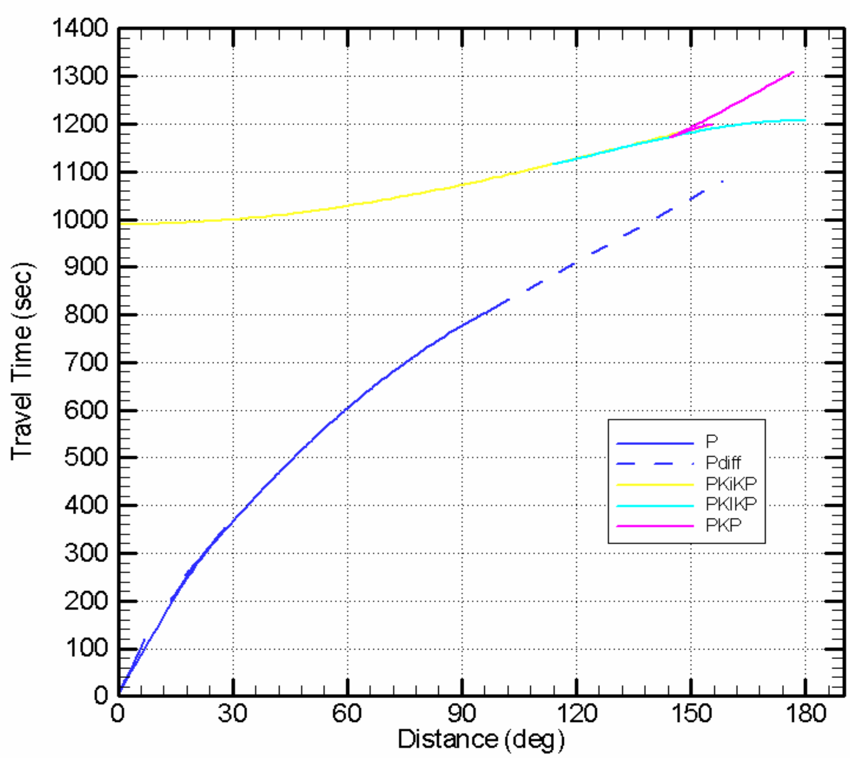

Figure 2. "First" P-phase travel time curve at a depth of 28 km. All phases and triplications are shown.

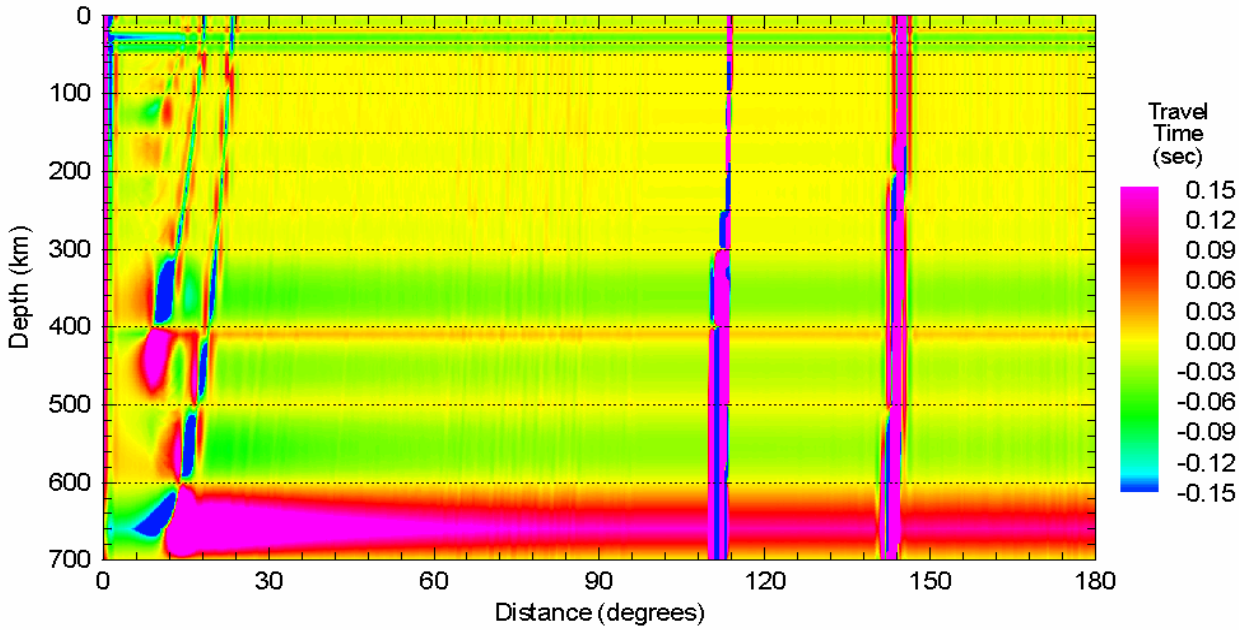

Figure 3. Difference between the TauP and libloc computed travel time values using a standard travel time table discretization (shown as black points).

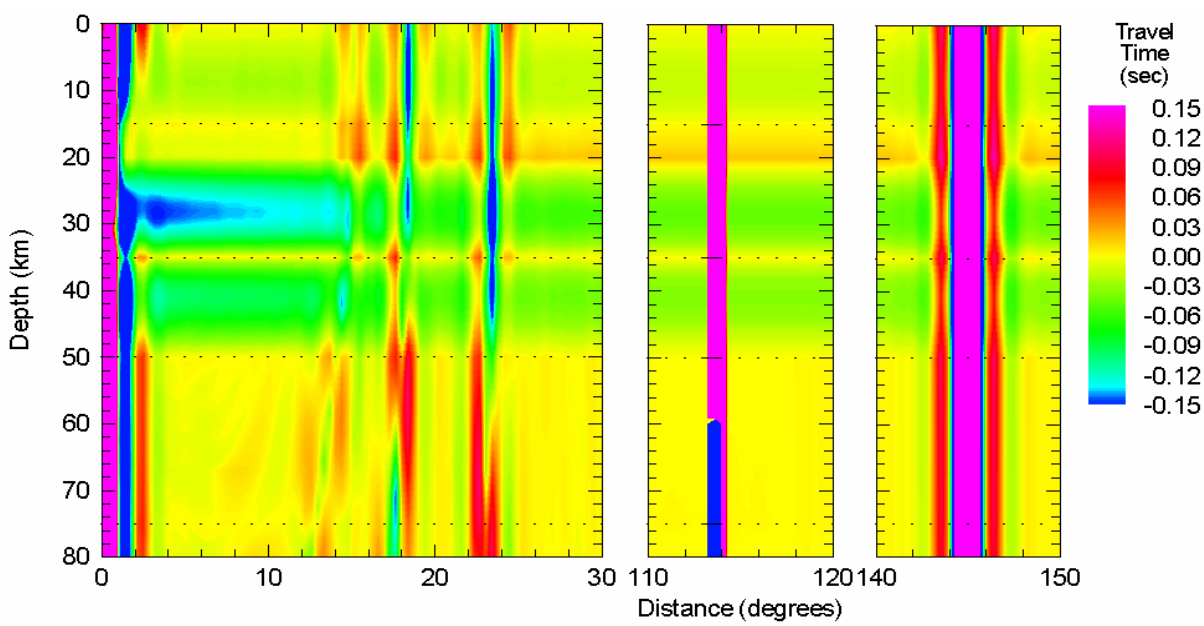

Figure 4. Detail of the difference between the TauP and libloc computed travel-time values. Data Between 30°-110° and 120°-140°have been omitted to better highlight problematic areas.

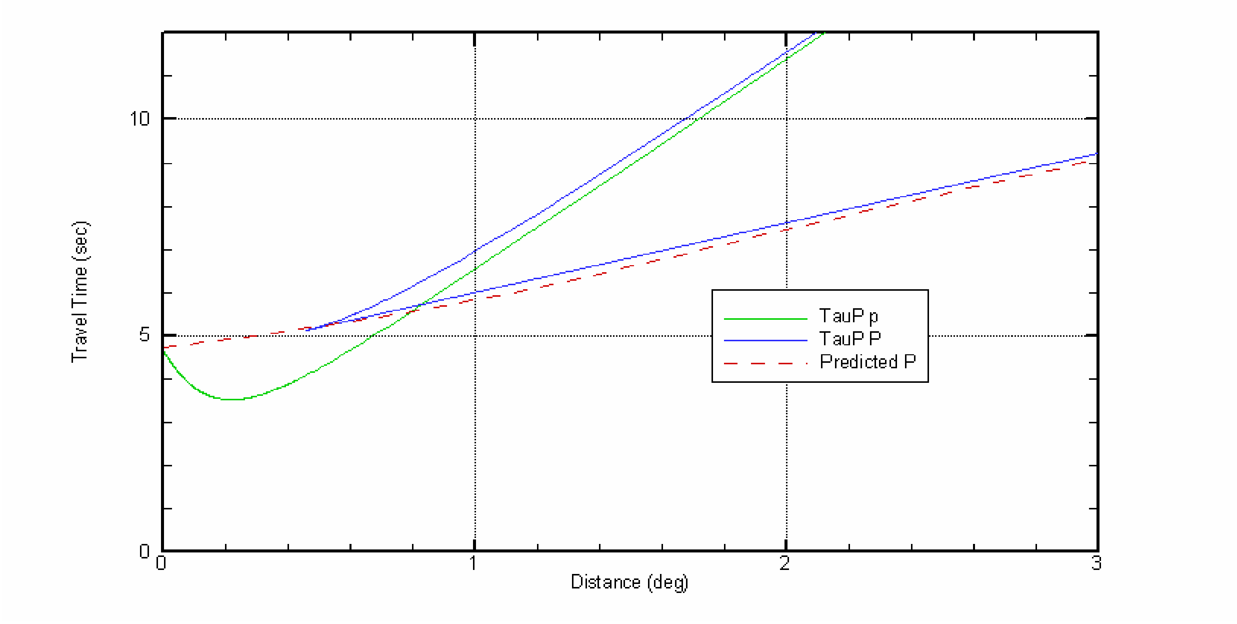

Figure 5. Reduced travel time as a function of distance at a depth of 28 km. With Predicted P discretization of 1°, the p phase is missed entirely, resulting in a discrepancy of more than 1 second at very close distances.

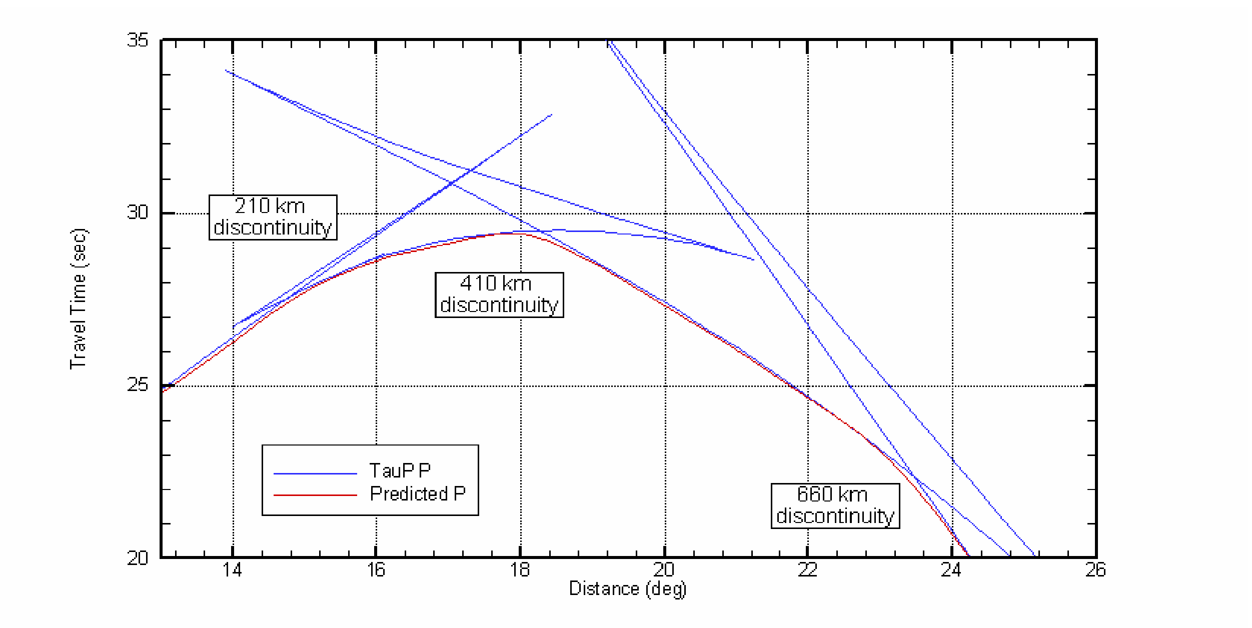

Figure 6. Reduced travel time as a function of distance at a depth of 28 km, emphasizing the triplications arising from the 210 km, 410 km, and 660 km velocity discontinuities. Note the discrepancies between the predicted first arrivals very close to the triplications.

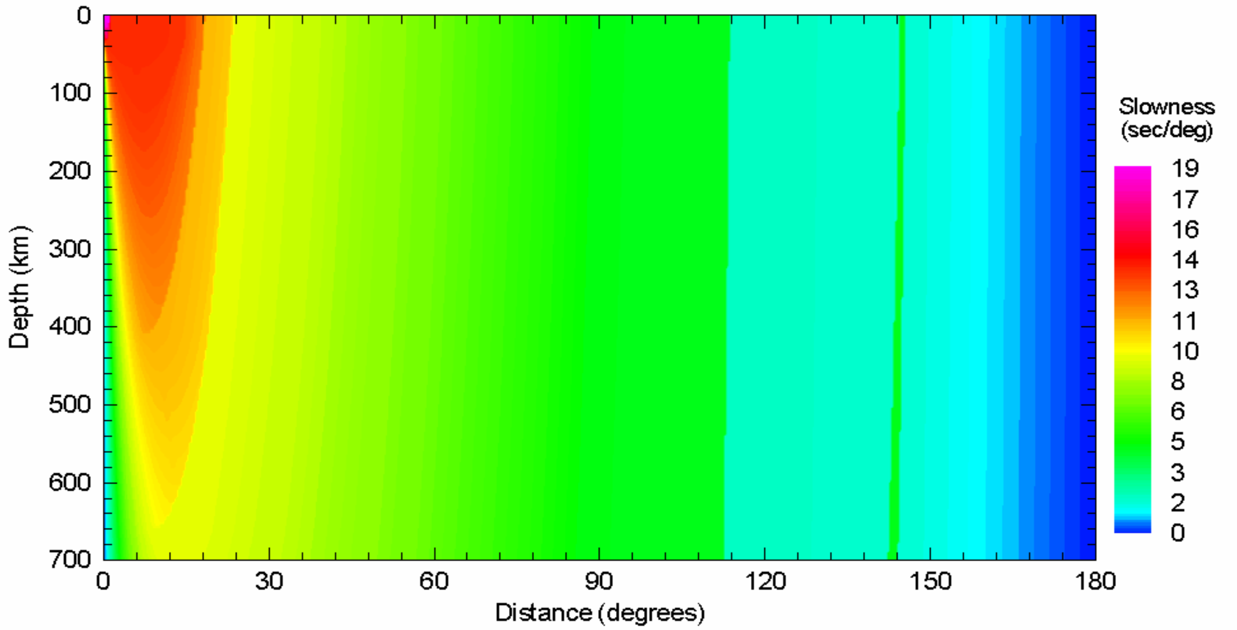

Figure 7. "First" P-phase slowness values computed with TauP. Spatial resolution is 0.36º in distance and 5 km in depth.

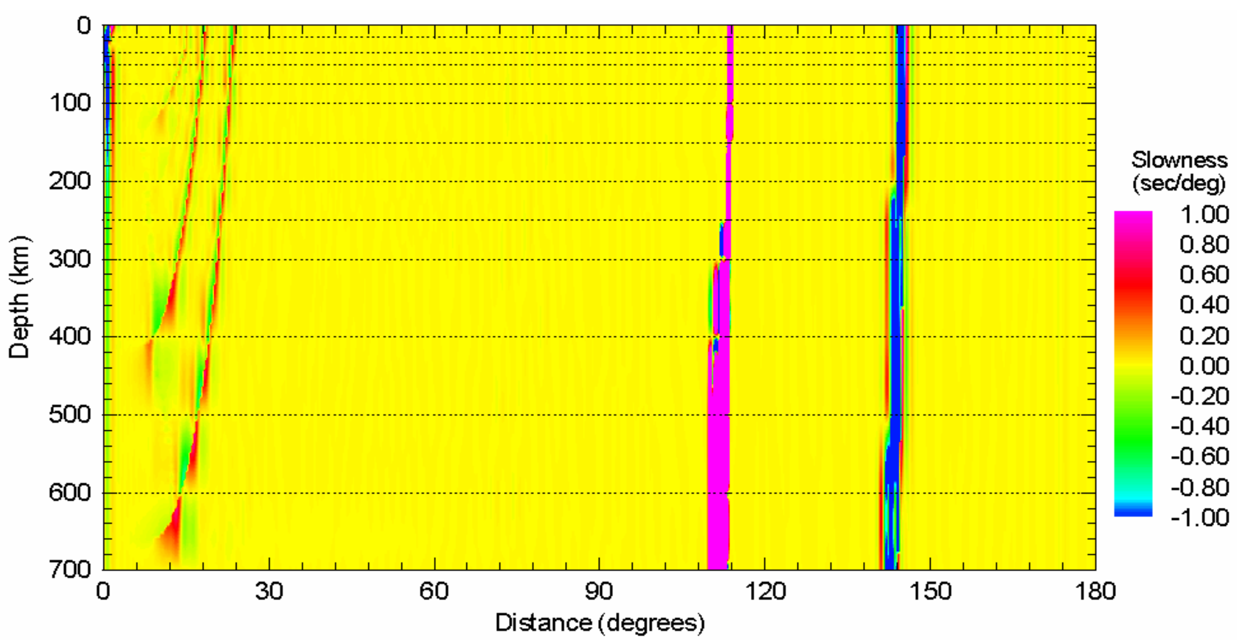

Figure 8. Difference between the TauP and libloc computed slowness values using a standard travel time table discretization (shown as black points).

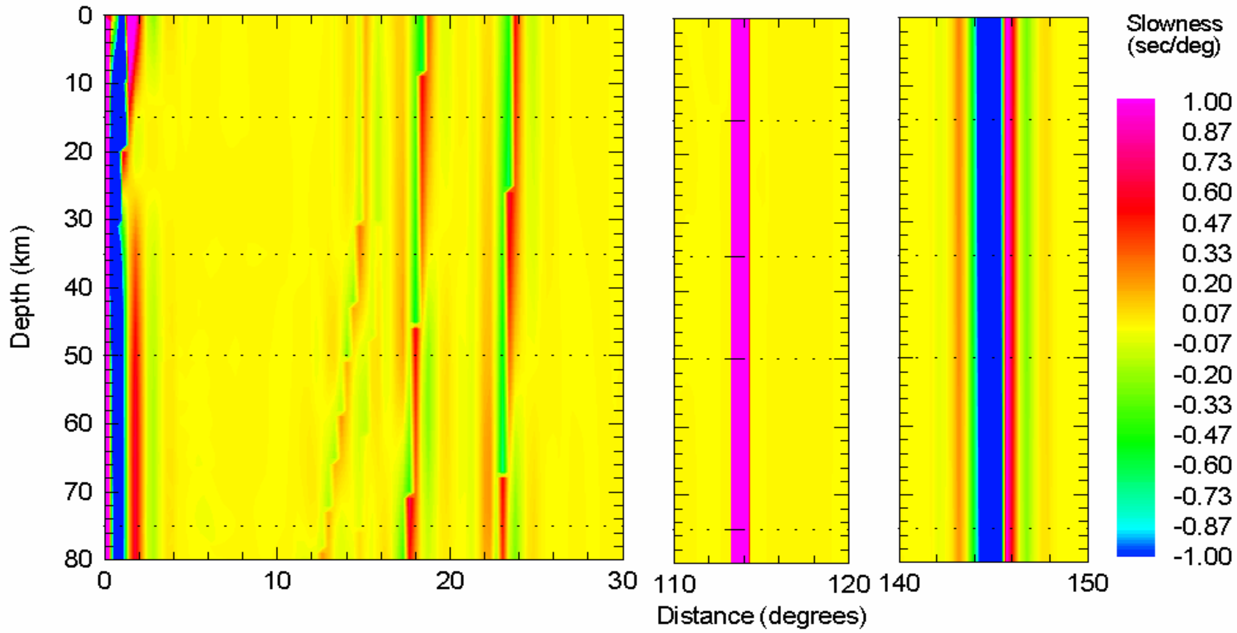

Figure 9. Detail of the difference between the TauP and libloc computed slowness values. Data between 30º-110º and 120º -140º have been omitted to better highlight problematic areas.

Results for slowness, when computed as the numerical derivative of travel time with respect to distance, are even more striking. Figure 7 displays the TauP computed slowness, while Figures 8 and 9 show the difference between look up table predictions and TauP using the same discretization as above.

Again we see problems along triplications, but the pattern is now asymmetric with under-predictions to the left and over-predictions to the right of the triplications. Unlike travel time predictions, slowness predictions are degraded throughout the depth profile, because the error is proportional to the grid spacing in the distance direction and does not vary with depth. Consequently, the magnitude of the slowness prediction errors exceeds 10% in places, even at the surface. This has significant implications with respect to phase identification and event locations involving slowness observations.

Our studies indicate that the use of coarsely gridded lookup tables can lead to significant errors in predicted travel time and slowness values, and undoubtedly affects computed spatial derivatives as well. Likewise, a coarse discretization is unable to represent the lateral boundaries of the valid phase ranges. We can expect these problems to be manifested as (a) poor event locations and (b) larger location uncertainties, whenever an observation falls into the poor prediction region. It is not clear how these problems impact model error. In addition, if the numerical derivative of travel time with respect to distance is used to support slowness predictions, then the current formulation can be expected to yield slowness prediction errors of several percent. This problem should also be expected to impact predicted event locations and confound phase identification during association.

We intend to explore two categories of modification to the current implementation that could significantly improve both accuracy and computational performance:

A) slowness values should also be represented as a table lookup with values derived directly from TauP. High-accuracy derivatives (*dtt/dz*, *dsh/dΔ*, and *dsh/dz*) should also be precomputed and represented on the lookup table. This would then allow us to use a lower order interpolation scheme, which would be faster. Values beyond the valid phase region should be pre-extrapolated to eliminate costly on-the-fly calculations. We anticipate that all values (travel time, slowness, model error and derivatives) could be represented on the same support grid, thereby amortizing lookup costs. These modifications would reduce significantly the numerical cost of supporting location calculations, while improving accuracy.

B) Ellipticity corrections are also phase-specific and require the calculation of three coefficients that vary with distance and depth. These are normally computed via a lookup table. We recommend using the same grid as (A) to further amortize lookup costs.

C) An alternative lookup-table discretization should be developed. We intend to investigate three possible approaches, which we present here in order of increasing complexity.

Approach 1: Rectilinear Grid with Fine Discretization

The simplest approach would be to refine the lookup table while maintaining a simple, rectilinear grid. In particular, the grid would need to be refined about the triplications. However, because the triplications are not always aligned with the grid, this would lead to over sampling in some areas. While this approach would require no algorithm modifications to perform interpolations, the memory requirements would probably increase by an order of magnitude. Likewise, the numerical cost of table lookup (via bisection) would also increase. As noted above, the values of interest are smooth and well-behaved except near the discontinuities, so a fine discretization would be inefficient in regions where the bicubic spline performs well.

Approach 2: Single-Region Model with Optimal Grid Sampling

An alternative approach would be to abandon rectilinear sampling. Instead, an optimal tessellation could be developed using algorithms already implemented in the Parametric Grid Library for the support of path correction calculations (see Hipp et al., 2005). Recent upgrades allow it to tessellate a region with a uniform triangular grid and then refine and decimate the grid until a desired level of accuracy is obtained. This yields an irregular triangular grid suitable for natural neighbor interpolation using the points of the triangle containing the interpolation point [e.g., Lawson, 1977; Sambridge et al., 1995]. For the problem at hand, the grid should be fine in and around discontinuities with coarser support in between. Interpolation supported by an irregular triangular grid requires first determining which triangle contains the point in question, which we propose to accomplish by a walking triangle search. This operation is $O(\text{Sqrt } N_t)$, where N_t is the number triangles in the table (Lawson, 1977). This compares favorably with two passes of $O(\log_2 N_n)$ for bisection of a rectilinear grid, where N_n is the number of nodes in a given direction, so long as the grid doesn't become too refined.

This method should provide accuracy to any desired level while also minimizing memory usage, but it is difficult to know in advance how computationally expensive the triangle lookup will be since we do not know how many triangles the representation will require. If the discretization becomes too fine near the discontinuities, then this could hinder a single-region approach.

Approach 3: Multi-Region Model with Optimal Grid Sampling

A more complex approach would involve a multi-region representation similar to how travel time path corrections are represented by the Parametric Grid Library (Hipp et al., 2005). Regions defining a given branch or phase would be individually tessellated and assembled to form a master lookup table. This would yield a series of overlapping meshes, each of which would be optimally sampled. Because there would no longer be discontinuities within any region, each could be accurately modeled with a relatively coarse mesh, thereby minimizing memory usage. Table lookup would proceed as follows: first the region containing the point of interest would be determined, then triangle lookup would be performed in that region's mesh as in the single-region model. Because of the coarseness of the regional discretization and the two-step lookup process, lookup should be much faster than a single-region optimal sampling. Interpolation would proceed via natural neighbor interpolation (e.g., Lawson, 1977; Sambridge et al., 1995). In summary, this representation should yield (1) arbitrary accuracy, (2) optimal memory usage, and (3) fast lookup.

Figure 10 shows this method schematically. The bottom portion shows three regions of support for the lookup table. The regions overlap and represent individual branches and/or regions where different phases arrive first. Crossovers are depicted in red. The upper portion shows the corresponding travel time curve at zero depth. When an interpolation is required, the first arrival's lookup table data is used, in this case, that of region B.

This representation could be used to construct a "global" first P table that combines compressional phases p, P, Pdiff, etc. This would be done by combining phase-specific regions within which a given phase is expected to be the first arrival. These phase-specific regions could be further sub-divided into sub-regions as necessary to capture finer branch regions.

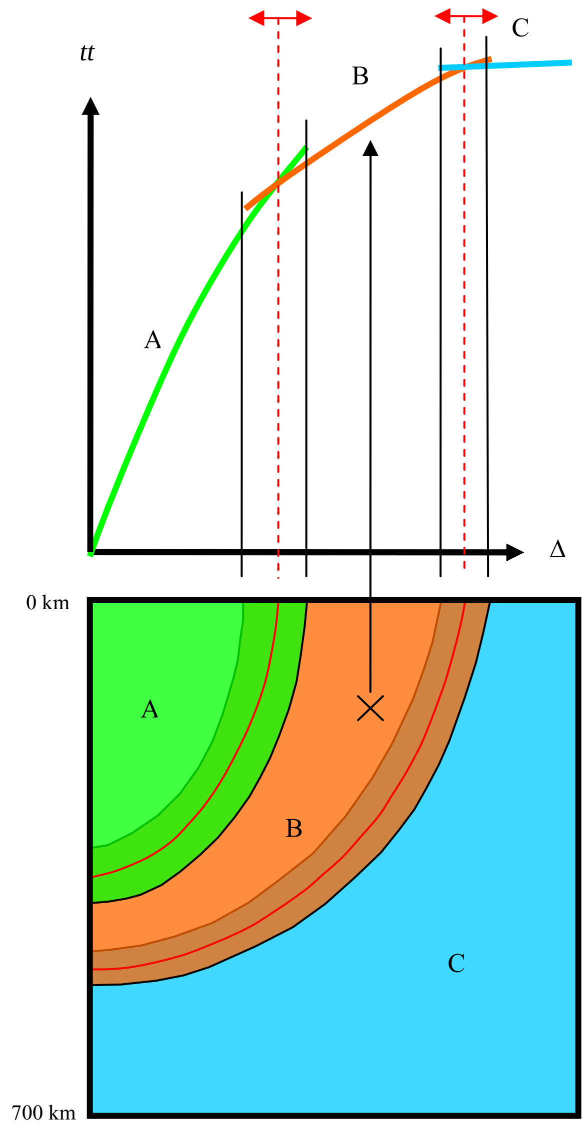

Figure 10. **Depiction of a multi-region model with optimal grid sampling. Top: Travel time curve at zero depth showing how different branches or phase regions of the travel time curve support calculations. Bottom: Regions supporting calculations for various branches.**

CONCLUSIONS AND RECOMMENDATIONS

In summary, the use of a rectilinear grid for representing travel time tables fails to capture the discontinuities of the surface and can lead to significant travel time and slowness prediction errors. An optimal tessellation scheme, either single or multi-region, could overcome these problems while maintaining numerical efficiency. Further efficiencies could be gained by pre-computing spatial derivatives and ellipticity correction coefficients and by pre-extrapolating values for lookup tables for phases having a limited valid range. As a final comment, it should be noted that Approaches 2 and 3 introduce discontinuous spatial derivatives, thereby impacting any gradient descent method. Consequently, a location algorithm designed to handle nonlinearities (e.g., Ballard, 2003) will be required if it is to be supported by such high-fidelity base models.

ACKNOWLEDGEMENTS

We would like to thank Mark Harris and David Gallegos for their thoughtful reviews. Their critiques and suggestions have improved the manuscript.

REFERENCES

Ballard, S. (2003). Seismic event location: Dealing with multi-dimensional uncertainty, model non-linearity and local minima, in *Proceedings from the 25th Seismic Research Review—Nuclear Explosion Monitoring: Building the Knowledge Base,* LA-UR-03-6029, Vol. 1, pp. 193-202.

Bratt, S.R. and T.C. Bache (1988). Locating events with a sparse network of regional arrays. *Bull Seism. Soc. Am.* 78: (2), 2019-2035.

Crotwell., H. P., T.J. Owens, and J. Ritsema (1999). The TauP Toolkit: Flexible seismic travel-time and ray-path utilities. *Seismological Research Letters* 70: (2), 154-160.

Hipp, J.R., C.J. Young, S.G. Moore, E.R. Shepherd, C.A. Schultz, and S.C. Myers (1999). Parametric Grid Information in the DOE Knowledge Base: Data Preparation, Storage, and Access, Sandia National Laboratories, SAND99-2277.

Hipp, J.R., S. Ballard, M.C. Chang, and L.A. Jensen (2005). Representing kriged travel time correction surfaces with optimal tessellations, this issue.

Kennett, B.L.N., and E.R. Engdahl (1991). Traveltimes for global earthquake location and phase identification. *Geophys. J. Int*. 105: (2), 429-465.

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

Lawson, C.L. (1977). Software for C1 Surface Interpolation, in *Mathematical Software III*, J. Rice (ed). Academic Press, New York.

Nagy, W. (1996). New travel-time handling facilities at the IDC: Functionality, testing and implementation details. Science Applications International Corporation, Technical Report, SAIC-96/1179.

Nagy, W. (2001). International Data Centre Documentation: Event Location Software. Science Applications International Corporation, Technical Report, SAIC-01/3010.

Sambridge, M., J. Braun, and H. McQueen (1995). Geophysical parameterization and interpolation of irregular data using natural neighbors, *Geophys. J. Int*. 122: (3), 837-857.