

## Hypocentric Relocations Aided by Virtual Receivers Constructed Via Seismic Interferometry

Franklin G. Horowitz and Larry D. Brown

Dept. EAS, Snee Hall, Cornell University, Ithaca, NY 14853, U.S.A.

frank.horowitz@cornell.edu

**Keywords:** microseismic, hypocenter relocation, virtual receivers, seismic interferometry

### ABSTRACT

The 3D elastic wave propagation program WPP (Petersson & Sjogreen, 2011) has been used to investigate whether the technique of Curtis et al. (2009) can be used to improve hypocentral relocations by employing virtual receivers near a cloud of microearthquakes. The virtual receiver technique can be loosely described as the "dual" of the ambient noise technique from seismic interferometry -- replacing noise sources on the boundary of a region of interest with physical receivers. Seismograms from events in the interior of the region of interest can be cross-correlated and integrated over all boundary receivers to estimate a seismogram from one of the interior events as if it were recorded at the location of another interior event.

Unlike ambient noise interferometry, where raypaths from all directions impinge on the region of interest, Virtual Receivers' raypath directions are constrained by the location of the physical receiver array. Hence, approximating the surface integral plays a large role in the practical success of the technique. Fortunately, stationary-phase arguments suggest that only a few physical receivers nearby the interior-source to virtual-receiver ray direction suffice to reconstruct the seismogram (as described in Curtis et al., 2009). Arrival time error statistics supporting this conclusion from WPP simulations are shown. Additionally, relocations of perturbed synthetic hypocenters using virtual receiver arrivals and hypoccc (Foulger and Julian, 2013) are shown. Improvements to the numerical least squares solver have been implemented, and the results indicate that the choice of solver strongly affects the quality of the relocations.

### 1. INTRODUCTION

A remarkable recent result from seismic interferometry allows one to create *virtual seismic receivers* (Curtis et al., 2009). By recording an original earthquake on an array of seismic receivers, one can estimate seismograms from other events as if they were recorded at the original event's location. This estimation occurs by signal processing the other events' seismic records from the array in conjunction with the original event's records.

Here, we present some *preliminary* efforts attempting to use Curtis et al.'s result to improve estimates of hypocentral locations in microseismic clouds.

### 2. THEORY

#### 2.1 Seismic Interferometry


For our purposes, the key expression from Curtis et al. (2009) is their equation (SI-16) from the Supplementary Information, which we reproduce here, as the top line of Eq. 1:

$$\begin{aligned} iK\omega \int_S u_n(x' | x_2) u_n^*(x' | x_1) dx' &= M_{ip}^2 M_{mq}^1 \partial_p \partial'_q G_{im}^h(x_2 | x_1) \\ &= \delta_{ip} \partial_p \delta_{mq} G_{im,q}^h(x_2 | x_1) \\ &= \partial_i u_i^h(x_2 | x_1) \end{aligned} \quad (1)$$

This (complex valued) expression is shown in the Fourier domain – hence the multiplications are convolutions – and the explicit dependence on frequency has been suppressed in the expression. The geometry is as displayed in Figure 1 with the boundary integral due to the use of the Green/Gauss theorem – in practice the integral is approximated by a sum over discrete locations. The vectors  $x_1$  and  $x_2$  are the locations of two distinct earthquakes, while the  $x'$  are the locations of the array receivers. The homogenous Green's function  $G_{im}^h$  has a source at  $x_1$  and is received at location  $x_2$  – it is a linear combination of the causal and the conjugated a-causal Green's functions, and it is symmetric with respect to source and receiver by reciprocity. The repeated index summation convention is used. The constant  $K$  absorbs all of the  $C_{ijkl}$  elasticity constants of the full expression, and is different for different geometries of earthquakes (Curtis et al., 2009) – we use it here as a mere constant of proportionality. Spatial gradients in the  $p$  direction are denoted both by  $\partial_p$  and by the subscripted comma convention.  $M_{mq}^1$  and  $M_{ip}^2$  denote the moment tensors for events number 1 and 2 respectively. The  $i\omega$  in the left-hand side is a temporal derivative in the Fourier Domain (transform sign convention dependent) probably being used here as an approximation of spatial gradients in this wave equation based theory. The expression  $u_n(x' | x_1)$  denotes the  $n$ -th component of a displacement seismogram from a source at  $x_1$  to the array receiver at one of the  $x'$  locations. The superscripted star denotes complex conjugation, such that the expression  $u_n u_n^*$  in the LHS is the sum of the cross-correlations between each of the 3 components of the displacement seismograms.

The full expression in Curtis et al. (2009) – for which the top line of Equation 1 is an approximation – requires unrealistic (or at the very least, rare in the real world) strain-gauge "dipole" seismometers. Equation 1 is the more practical displacement "monopole" seismometer expression, which allows arbitrary moment-tensor events to be used as either a source or a virtual receiver event. To keep the numerical experiments simple and appropriate for our first-arrival estimates, we choose to restrict our attention to

explosive sources and virtual receivers with simple (spherical) moment tensor representations. Hence, substituting Kronecker deltas for the moment tensors appropriately (i.e.  $M^1_{mq} = \delta_{mq}$  and  $M^2_{ip} = \delta_{ip}$ ) we find the second line of Equation 1. Finally, using the definition of moment tensors ( $u_i = M_{mq} G_{imq}$ ; Aki and Richards, 1980, eqn. 3.22), we find the bottom line.



**Figure 1: (From Curtis, et al. 2009).** a) The standard interferometry geometry. Seismic events are depicted as red stars, receivers as blue triangles. b) The virtual receiver geometry. An array of receivers (blue triangles) receives seismograms from 2 distinct seismic events (red stars). The Green's function between the events can be estimated via Equation 1.

In words, the bottom line of Equation 1 states that the homogenous volumetric strain seismogram (i.e. proportional to the seismogram recorded by a hydrophone) propagating from  $x_1$  to  $x_2$  can be estimated from observing the displacement seismograms from both earthquakes using a summation over the entire set of receiver stations as an approximation of the full surface integral.

In essence, we have created a virtual hydrophone receiver at position  $x_2$ !

An interesting feature of Equation 1, is that the math places no restrictions on whether  $u_n(x'|x_2)$  occurs before or after  $u_n(x'|x_1)$  in time. As long as we have both sets of seismograms recorded, we can employ the equation. For our present purposes, this means that the choice of the best located (or “nicest” moment tensor) virtual receiver(s) event(s) can be made or adjusted at any time after the activity of a microseismic swarm. There are other (related) important consequences of this order-blind feature (e.g. Curtis et al., 2012) that we hope to flesh out with future work.

## 2.2 Improving the Microseismic Locations

We focus on the Double Difference (DD) location algorithm of Waldhauser and Ellsworth (2000), and its variants (e.g. Zhang and Thurber, 2003; or the hypoccc code of Foulger and Julian, 2013). The basic technique minimizes traveltime residuals of the form:

$$dr_k^{ij} = (t_k^i - t_k^j)^{obs} - (t_k^i - t_k^j)^{cal} \quad (2)$$

In this expression,  $dr$  is a DD residual, the  $t$ 's are travel times, the superscripted indices refer to the earthquake events, the subscripted index refers to the recording station, and the superscripted “*obs*” and “*cal*” refer to the measured *observation* and the velocity model derived *calculated* values respectively. The actual residual being minimized by an inversion procedure is the sum over all stations of Equation 2.

Notice that the DD residual is composed of two conceptually different parts. The “*obs*” component relates to the difference in travel time between two events observed from the same instrument. This is the primary observation of the method, and contains all of the information about the relative hypocentral locations, source times, and the velocity field encountered by the actual source-receiver ray-paths. The “*cal*” component is composed of differences in travel times due to the assumed or estimated velocity models.


Virtual receivers can help improve inversions based on such residuals in several ways. By using Equation 1 to increase the effective number of arrival time observations, the statistics in the DD inversion are improved. Virtual receivers also improve the condition number of the linear system to be solved – discussed below. By placing virtual receivers closer to the swarm of microseismicity than possible with real receivers, the effects of inaccurate velocity models can be ameliorated. The virtual receivers’ close-in geometry also allows for better ray-path coverage for tomographic estimates – improving velocity model estimates for the rockmass surrounding the swarm, hence allowing the possibility of iterative improvement in the absolute locations. There are probably other reasons that will become apparent as this new technique matures.

## 3. SIMULATIONS

### 3.1 The Wavefield Simulator

We simulate a microseismic cluster (Figure 2) as being located on a planar fault with attitude N45W 45NE, passing through the point (6000, 5000, 1000) meters. The (right handed) coordinate system is shown and has its x component increasing due north, its y component increasing due east, and its z component increasing downwards. The receiver array is arranged at 30-degree increments around a circle on the surface with radius 2500 meters, and its center located at (5000, 5000, 0). Note that the cluster is offset with respect to the center of the receiver circle. There are 25 (blue) microseismic cluster events distributed on the fault plane, with the center event (red) designated as the “Virtual Receiver” (VR) – which plays the role of  $x_2$  in Equation 1. The numerical model extends 10x10x2 km as shown. There is a free surface at  $z=0$ , and all other boundary surfaces use the non-reflecting boundary conditions implemented in the modeling code.

We use the parallel fully elastic 3D finite difference Wave Propagation Program from Lawrence Livermore (WPP; e.g. Petersson and Sjögren, 2011) for our wave field simulator. The grid spacing is 20 meters in all three directions. We use a P speed of 2500 m/s, an S speed of 1500 m/s, and a mass density of 2600 kg/m<sup>3</sup> throughout the entire volume. After trial and error, we use their predefined source-time-function called “Smoothwave” – a 7th order polynomial in time with reasonable spectral behavior – as producing solutions meeting the stability criteria of their simulator. (Triangular source-time-functions have too broad a spectrum for stable solutions.) WPP allows one to record seismograms from any position in the model including in the interior, so we can compare the seismograms estimated via Equation 1 with those directly reported by WPP.



**Figure 2: Model Geometry.** The event cloud is dipping away from you towards the NE, and the entire model domain is displayed, with distances in meters. The microseismic “swarm” is shown in blue dots, while one event (colored red) has been chosen to form the “virtual receiver” estimates via Equation 1. The physical receivers are shown at the top (free) boundary.


Model run-times for the 3-second seismograms we generated were on the order of 5 minutes per simulation on a 24 core (hyperthreaded) machine, and there were 25 simulations required (one per aftershock) to generate the full suite of seismograms.


### 3.2 The Virtual Receiver Process

Because Equation 1 represents the required computations in a very compact notation, we display the actual workflow graphically in Figure 3.

Figure 3 indicates that estimating the P arrival time from interferometry requires some subtlety. Thinking about computing  $u u^*$  in the Fourier domain allows us to see that phase relations are strongly perturbed by this operation (to the point of extinguishing all phase information in the limit of autocorrelation). Thus we should not expect the “first break” in any interferometric seismogram to correlate with the P arrival. In fact, we find heuristically that the maximum of *abs* (amplitude) is the best “pick” for arrival time – in the spirit of zero-phase arrivals. We do not currently have a detailed explanation for the observation that the picks are all minima rather than maxima (as you would expect for a cross correlation), but suspect that the imaginary unit *i* in Equation 1 has a role here.

From experience using other source-time-functions, we infer that the high frequencies evident in some of the red seismograms are an indicator of incipient numerical instability in the WPP solution.






**Figure 3:** Top right panel (previous page): the geometry of Figure 2, with the source (“new”) event indicated by the gray arrow, and the VR event in red. Top left two panels: 3 component seismograms recorded on the (simulated) surface array from both the VR event and the new event. Mid-right panel: cross correlations of the individual seismogram components. Note how the arrival times are decreased by the cross-correlation “re-datuming” the traces to the position of the VR. Bottom panel (this page), top: the resulting interferometric pressure VR seismogram is shown in blue, while the pressure seismogram computed in by WPP is shown in red. The vertical red bar is the known P arrival time, while the green bar is the S arrival. Arrival time “picking” for the interferometric seismogram is described in the text. The corresponding fault image on the right indicates the origin event in red. The mismatch in amplitude between the red and blue curves is attributed to our incorrect assumption of  $K=1$  in Equation 1. Bottom panel, middle and bottom: results from the equivalent operations performed on events originating at two other locations (indicated in the corresponding right fault images).

### 3.3 Relocations

In order to compare the relocation results with and without VR seismograms, we use the double difference code “hypocc” (Foulger and Julian, 2013) – modified with a better performing damped least squares solver as discussed below – to relocate randomly perturbed event locations. We start with the true locations of events used in the wavefield simulations, randomly perturb their depths, and feed the result into hypocc to perform the relocation (Figure 4). It is noteworthy that the RMS error of the relocations does not improve very much via the inclusion of the VR seismogram. This is clearly discouraging for the utility of this technique. However we note that the more VRs, perhaps located more optimally with respect to the stationary phase arguments in the discussion below, might well improve the relocations further, but we have not yet tested this.



**Figure 4: A)** Along-strike view of DD relocations. The upper left shows a perspective view of the true locations. The upper right shows those locations with their depths randomly perturbed from a uniform distribution of length  $\pm 50$ m. The lower left shows the relocation from hypocc using the surface (Figure 2) receivers only. The lower right shows the relocations after including the VR seismograms in the double difference calculation. Note the slight decrease in RMS residual with the inclusion of the virtual seismogram results. **B)** A view with the same layout, but adding a less stable part of the singular spectrum into the solution. See section 4.2 for a full discussion.


## 4. DISCUSSION

### 4.1 Error Sources in the Interferometry


#### Full Surface Integral Approximation

The integral appearing in Equation 1 is meant to be a fully continuous surface integral over the entire bounding surface of the volume of interest. However, in real-world applications of the theory we are only able to approximate that integral via a discrete sum over our physical receivers. Such receiver geometries, which occur because of obvious operational constraints for seismic surveys, are unlikely to be reasonable geometric approximations of the full surface integral.

To investigate how much of an improvement is available from a better surface integral approximation, we recorded seismograms via WPP from locations arrayed on a 1km grid over the bounding surface (Figure 5). The resulting interferometric estimates of the seismograms from the ground surface array (Figure 2) and from the better surface integral approximation (Figure 5) are shown in Figure 6. Encouragingly, in this example the quality of the virtual receiver estimates formed via the ground surface array is quite comparable to the quality from those from the bounding surface array. Once again we see that the (zero-phase style) deepest minimum of the estimate is a reasonable “pick” for the first arrival, and indeed this is visually confirmed across all of the seismograms from all of the events on the modeled fault plane.




**Figure 5:** Summing over these locations forms a better approximation to the Equation 1 surface integral. Each receiver is on a 1km grid on all of the boundaries of the original volume displayed in Figure 2.



**Figure 6:** Top: similar to the results found in the bottom panel of Figure 3 – using only seismograms recorded from the ground surface array – but from a different source location than those displayed in Error! Reference source not found. Bottom: everything is kept the same as the top, except that the blue seismogram is estimated from a better approximation of the full surface integral of Equation 1. Recording seismograms at discrete gridded locations shown in Figure 5 and summing the results form this better approximation.

#### Stationary Phase Stacking


The surface integral in Equation 1 stacks point-wise seismograms with a wide variety of phase angles (e.g. Schuster, 2009). From purely geometric considerations, those physical receivers a small angle away from the source-to-virtual-receiver ray should have phases that interfere constructively, while those physical receivers from a larger angle away from that ray should interfere destructively. Obviously, this interference is wavelength dependent. **Error! Reference source not found.** illustrates the geometry.



**Figure 7:** The red starburst is the microseismic source event. The blue triangle is the virtual receiver event connected to the source event by a blue ray-path. The thin black line is the leading edge of the propagating wave. The other blue lines shown below are ray-paths to physical receivers, with their corresponding leading edges under the action of Equation 1 shown by black lines. The dotted lines from those leading edges back to the source-VR ray-path show the projections of the physical receiver arrivals. During the stacking operation implicit in Equation 1, small angles away from the source-VR path interfere constructively, while larger angles interfere destructively. Such constructively interfering wave phenomena are known as “Stationary Phases” in asymptotics.

It is apparent from examination of the geometry in Figure 7 that the full surface integral in Equation 1 is not necessary to construct an asymptotic approximation of the waveform. Only those ray-paths “close” in some (wavelength dependent) angular sense to the source-to-VR ray-path interfere constructively in the stacking process. Accordingly, we now numerically investigate first-arrival-time error statistics as a function of angle from the source-to-VR ray-path.

Because we have available all of the seismograms recorded at the locations in Figure 5, we can compute Equation 1 stacks including seismograms only within a specified angle from the source-VR ray-path. Figure 8 summarizes the resulting arrival time “picking error” statistics. The “picking error” is defined to be the arrival time of the deepest minimum in the Equation 1 stack minus the theoretical arrival time computed from the known distances and velocity for each of the blue microseismic events.



**Figure 8:** P arrival time error statistics for VR seismograms estimated according to Equation 1. The y axis is arrival time error in seconds. The x axis is the angle in degrees from the source-VR rays up to which we included seismograms from the physical receiver locations of Figure 5. Also shown on the x axis are the special cases corresponding to the “Top Circle” locations of Figure 2 and the full surface integral respectively. Violin diagrams summarize the arrival time errors for each case.

There are several things noteworthy in Figure 8: Firstly, the “Top Circular” physical array arrival time errors have a median of about 20 msec., which corresponds to a median raw location error of about 50 meters from that geometry using this technique with a single virtual receiver. Obviously, incorporating multiple virtual receivers into a relocation effort could improve that raw error by averaging out the arrival time error fluctuations. Secondly, one explanation for the trend observed in the median results from 20 to 60 degrees is that more physical receivers give better counting statistics in our Equation 1 stack. That is consistent with the tight, low-median errors found for the “All Receivers” result. Thirdly, the 20 to 60 degree median errors are all positive, implying that the stack and this “deepest minimum” picking strategy appears to produce events that arrive later than they should. We currently have no explanation for this observation. Fourthly, there are some very bad outliers evident in our method, with arrival time errors on the order of 150 msec. or so. We have no explanation for this observation either. We speculate that some numerical instabilities from

the wavefield simulator are still found in some of the solutions which contaminate the constructive interference — and hence our statistics.

#### 4.2 Error Sources in the Double Difference Relocations

A surprising finding in performing this work was the delicate nature of the relocation solution strategy required by the stock hypoccc code (Foulger and Julian, 2013). Since hypoccc's ancestor HypoDD (Waldhauser and Ellsworth, 2000) uses the same sparse numerical solver for large relocation problems (LSQR; Paige and Saunders, 1982), we infer that HypoDD has a similar delicate nature — something we expect to be able to test before the meeting. We think that the underlying reason for the finicky numerical performance — at least in the case of our relatively limited numerical study — is that the so-called condition equations (the linear algebraic system to be solved) are quite ill-conditioned.


##### 4.2.1 Using the LSMR Solver

We have attempted to improve the performance of hypoccc via two different routes. Firstly, we replaced LSQR with its successor LSMR (Fong and Saunders, 2011). LSMR improves on LSQR's performance by possessing the following property (quoted directly from Saunders' website <http://www.stanford.edu/group/SOL/software/lsmr.html>): “*Special feature*: Both  $\|r\|$  and  $\|A^T r\|$  decrease monotonically, where  $r = b - Ax$  is the current residual. For LSQR, only  $\|r\|$  is monotonic. With loose stopping tolerances, LSMR may be able to terminate significantly sooner than LSQR.” LSMR has essentially the same calling signature as LSQR, also works with a very similar sparse matrix interface, and the Fortran90 implementation was easily adapted to be used by hypoccc. (HypoDD, being implemented in Fortran, should in principle also be an easy “port” to use LSMR.) We found that the LSMR solver indeed was less finicky than LSQR. Because there is a wide audience for double difference relocations, we recommend that other users investigate LSMR in their use cases. (A git repository, containing modifications to hypoccc including the port of LSMR is available upon request to Frank Horowitz. Be warned, though, it is research code.)

##### 4.2.2 Solving with Singular Value Decomposition

Secondly — following the suggestion found in Waldhauser and Ellsworth (2000) — we performed a preliminary investigation of the ill-conditioning of the condition equations' matrix for our test-case by using Singular Value Decomposition (SVD). A strategy that enables us to pick among a library of many solvers — including Python implementations of both a well-tested sparse SVD routine as well as LSMR — uses the SciPy interface (Jones, *et al.*, ongoing). We chose to port hypoccc's solver to a Python/SciPy interface using the interface language Cython (Behnel, *et al.*, 2011).

The singular spectra corresponding to the damped solutions shown above in Figure 4A and B (and their corresponding undamped solutions) are shown in Figure 9. Note that damping had very little effect on the character of the singular spectra — at least for the very low value of the damping factor that we used — indicating that our damped least-squares solutions were not effective in helping stabilize our DD problem. Under the usual definition of the condition number for the problem — the ratio of the highest to the lowest singular value — the addition of virtual receiver information does indeed render our least squares solutions to be formally less ill-conditioned.



**Figure 9: Singular Spectra.** Shown are the singular values for the SVD decomposition of 4 different relocation problems built from the geometry of Figure 2. The solid blue and dashed cyan curves are for a damped and undamped (respectively) least squares DD problem including the virtual receiver seismogram estimates in the relocation. The solid red and dashed black curves are for the corresponding surface-receiver-only DD problems. All solutions rescaled the time parameter appropriately using the mechanism built-in to hypoccc. The damping parameter was set to a very low value of 0.0003 by trial and error. Note that all 4 spectra display essentially similar characteristics for singular values higher than about 2.E-3 (singular vector number 68). We defined the “noise floor” for this problem (i.e. the singular values/vectors that do not contribute very much resolving power to the solution) to be below 2.E-3. For both the damped and undamped solutions in the “noise floor”, the singular values of the solutions that include seismograms from the single virtual receiver are greater than those corresponding to the solutions using the surface stations only. This means that the inclusion of virtual receiver seismograms does indeed lead to a DD problem that is formally better conditioned.

Our practical problem is that this (formal) better conditioning of the solution does not help. To demonstrate this we include into the solution the highest set of 4 singular values in the noise floor – 4 to allow the corresponding singular vectors to span the x, y, z, and t components of the solution. The RMS error results shown in Figure 4B indicate that our DD relocations begin to diverge. This divergence worsens by including more singular values/vectors from the noise floor – which essentially form a basis for the (effective) nullspace of the condition equations.

These observations lead us to infer that the condition equations from our DD problem are simultaneously over-constrained (have more equations than unknowns) and under-constrained (ambiguous combinations of our relocation vectors exist – e.g. see the discussion in section 15.4 of Press *et al.*, 1992). A corollary is that double differences do not provide adequate resolving power for the unknowns in our problem.

We suspect that such ambiguous combinations of relocation vectors might play a role in other deployments of double difference relocations.

Following on from a conversation with Bill Ellsworth (pers. comm., 2014) to the effect that receiver “geometry is king” for relocation problems – not to mention that our circular array of physical receivers is a particularly weak geometry – we have begun to investigate replacing that circular physical array of receivers with a subset of the top surface receivers from Figure 5. While preliminary results are not encouraging, we expect to have performed a more systematic investigation by the time of the meeting, and expect to be able to report on those results at that time.

## 5. CONCLUSION

We have succeeded in constructing a virtual receiver in the interior of the microseismic source cloud, but it appears not to be as helpful in improving microseismic relocations as we initially expected. However, the virtual receiver approach clearly improves the condition number of the solutions. This leads us to speculate that incorporating more than one virtual receiver into relocation problems might be beneficial. However, clearly some augmentation of our strategy that needs to happen before we can successfully deploy this technique

## ACKNOWLEDGEMENTS

FGH thanks Klaus Regenauer-Lieb of the University of Western Australia for allowing access to his parallel cluster for the WPP computations performed during this work. He also gratefully acknowledges helpful conversations with Bruce Julian and Bill Ellsworth on the relocation aspects of this work.

## REFERENCES

- Aki, K., & Richards, P. (1980). *Quantitative Seismology*, vol. 1. W.H. Freeman and Co., San Francisco, 1st ed.
- Behnel, S., Bradshaw, R., Citro, C., Dal Cin, L., Seljebotn, D. S., & Smith, K. (2011). Cython: The best of both worlds. *Computing in Science Engineering*, 13 (2), 31-39. <http://dx.doi.org/10.1109/MCSE.2010.118>
- Curtis, A., Nicolson, H., Halliday, D., Trampert, J., & Baptie, B. (2009). Virtual seismometers in the subsurface of the Earth from seismic interferometry. *Nature Geoscience*, 2 (10), 700-704. <http://dx.doi.org/10.1038/ngeo615>
- Curtis, A., Behr, Y., Entwistle, E., Galetti, E., Townend, J., & Bannister, S. (2012). The benefit of hindsight in observational science: Retrospective seismological observations. *Earth and Planetary Science Letters*, 345-348, 212-220. <http://dx.doi.org/10.1016/j.epsl.2012.06.008>
- Fong, D. C., & Saunders, M. (2011). LSMR: An iterative algorithm for sparse Least-Squares problems. *SIAM Journal on Scientific Computing*, 33 (5), 2950-2971. <http://dx.doi.org/10.1137/10079687X>
- Foulger, G. R., & Julian, B. R. (2013). Seismological software for geothermal monitoring. In *PROCEEDINGS, Thirty-Eighth Workshop on Geothermal Reservoir Engineering*, SGP-TR-198. Stanford University. <https://pangea.stanford.edu/ERE/pdf/IGAstandard/SGW/2013/Foulger.pdf>
- Jones, E., Oliphant, T., Peterson, P., & contributors, o. (2001-ongoing). Scipy: Open source scientific tools for Python. <http://www.scipy.org/>
- Paige, C. C., & Saunders, M. A. (1982). LSQR: An algorithm for sparse linear equations and sparse least squares. *ACM Trans. Math. Softw.*, 8 (1), 43-71. <http://dx.doi.org/10.1145/355984.355989>
- Petersson, N. A., & Sjögreen, B. (2011). *User's guide to WPP version 2.1.5*. Lawrence Livermore National Laboratory. <https://computation.llnl.gov/casc/serpentine/download/WPP-2.1-UsersGuide.pdf>
- Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. (1992). *Numerical Recipes in C: The Art of Scientific Computing, Second Edition*. Cambridge University Press.
- Schuster, G. T. (2009). *Seismic Interferometry*. Cambridge University Press, first ed.
- Waldhauser, F., & Ellsworth, W. L. (2000). A double-difference earthquake location algorithm: Method and application to the Northern Hayward Fault, California. *Bulletin of the Seismological Society of America*, 90 (6), 1353-1368. <http://dx.doi.org/10.1785/0120000006>
- Zhang, H., & Thurber, C. H. (2003). Double-difference tomography: The method and its application to the Hayward Fault, California. *Bulletin of the Seismological Society of America*, 93 (5), 1875-1889. <http://dx.doi.org/10.1785/0120020190>