

Optimized Workflows with a Single Phase-Only Response Correction for Building Empirical Green's Functions for Ambient-Noise Tomography

Caio Ciardelli  * ¹, Yoweri Nseko , Albert Kabanda , Suzan van der Lee  ¹

¹Department of Earth, Environmental, and Planetary Sciences, Northwestern University, Evanston, Illinois, U.S.A.

Author contributions: *Data Curation:* Caio Ciardelli, Yoweri Nseko, Albert Kabanda. *Formal Analysis:* Caio Ciardelli, Suzan van der Lee. *Methodology:* Caio Ciardelli, Yoweri Nseko, Albert Kabanda, Suzan van der Lee. *Software:* Caio Ciardelli, Yoweri Nseko, Albert Kabanda. *Validation:* Caio Ciardelli, Yoweri Nseko. *Visualization:* Caio Ciardelli, Yoweri Nseko. *Writing – Original draft:* Caio Ciardelli. *Project Administration:* Suzan van der Lee. *Writing – Review & Editing:* Suzan van der Lee.

Abstract We present simple and optimized workflows for computing empirical Green's functions in ambient-noise tomography that enhance computational efficiency and numerical stability. A key improvement is a phase-only instrument-response correction applied only once after stacking instead of to the raw data before correlation. This prevents instability in spectral division, simplifies computations, and reduces execution time. While some of the additional optimizations we employ are already in use within the ambient-noise tomography community, we provide a detailed description along with systematic benchmarks that quantify their actual impact on runtime and stability. Key improvements include reducing redundant Fourier transforms and combining spectral equalization, cross-correlation, and stacking into a single frequency-domain step. An additional optimization reuses spectral representations of individual stations across multiple station pairs, maintaining linear complexity. We also propose a completely new optimization: applying a phase-only instrument-response correction only once *after* stacking instead of before correlation. This prevents instability in spectral division, simplifies computations, and reduces execution time. We validate the workflows using datasets from Southern California, Brazil, and Uganda. For individual station pairs, our primary optimized workflow (WF2) reduces execution time by approximately 67–75% (speed-up factors of 3.0–3.9), closely matching theoretical expectations (~ 5.1). A more scalable variant (WF3) achieves speed-up factors of 15–60 for moderate-sized networks. Furthermore, we demonstrate that a partial implementation into existing codes, requiring only minimal modifications, yields about 10% execution-time savings and improved numerical stability. The proposed workflows produce EGFs nearly indistinguishable from conventional methods and are particularly suitable for large-scale ambient-noise tomography in computationally limited environments.

Production Editor:
Gareth Funning
Handling Editor:
Paula Koelemeijer
Copy & Layout Editor:
Tara Nye

Received:
January 9, 2025
Accepted:
January 21, 2026
Published:
February 12, 2026

1 Introduction

Imaging the structure and composition of the Earth's lithosphere and asthenosphere is essential for understanding tectonics, geodynamics, and geological processes and structures, with additional implications for resource exploration and hazard assessment (e.g., Nolet, 2008; Artemieva, 2011; Routh et al., 2017; Cathles et al., 2023; Muir and Ross, 2023). Seismic tomography, using earthquake-generated seismic waves, has been paramount in constructing detailed subsurface images (e.g., Aki et al., 1977; Dziewonski, 1979; Nolet and Kennett, 1990; Thurber, 2003; Martin and Wenzel, 2006; Priestley and Tilmann, 2009; Fishwick, 2010; Plomerová and Babuška, 2010; Kumar et al., 2012; Clouzet et al., 2018; Audhkhasi and Singh, 2022; Ciardelli et al., 2022; Vaddineni and Singh, 2023). However, the uneven spatial and temporal distribution of earthquakes poses a significant limitation, leading to gaps in the resolution and coverage of tomographic models (Nolet, 2008; Rawl-

inson and Sambridge, 2003), especially in aseismic regions such as the South American Platform (Almeida et al., 2000; Petersen et al., 2018) where seismicity is sparse.

Ambient-noise tomography (ANT) has emerged as a valuable complement to traditional earthquake-based tomography. By using continuous background seismic noise, ANT offers more uniform spatial coverage controlled by the distribution of stations rather than earthquakes (Bensen et al., 2007). The methodology involves cross-correlating ambient seismic noise recorded at pairs of seismic stations to extract empirical Green's functions (EGFs) (Weaver and Lobkis, 2004; Snieder, 2004; Larose et al., 2006; Stehly et al., 2006; Yang and Ritzwoller, 2008), which are then used to infer subsurface properties such as seismic wave velocities and lithospheric units (e.g., Moschetti et al., 2007; Lin and Ritzwoller, 2011; Witek et al., 2018) and attenuation (e.g., Prieto et al., 2009). This approach has proven effective in providing high-resolution images of the Earth's lithosphere (e.g., Shapiro et al., 2005; Ritz-

*Corresponding author: caio.ciardelli@gmail.com

woller et al., 2011; Ojo, 2021), even in regions with limited seismic activity.

Removing the instrument response is a routine step in ANT to account for sensor effects before estimating EGFs (Bensen et al., 2007; Lin et al., 2008). This is typically done prior to cross-correlation and involves the removal of both amplitude and phase response (e.g., Herrmann, 2013; Lecocq et al., 2014; Jiang and Denolle, 2020; Clements and Denolle, 2020; Li et al., 2021; Magrini et al., 2022; Makus and Sens-Schönfelder, 2024). However, full deconvolution requires spectral division, which can amplify noise in frequency bands where the instrument is less sensitive, thereby introducing instability (Scherbaum, 1996; Akkar and Boore, 2009; Havskov and Alguacil, 2016).

We present a modified ANT workflow that improves numerical stability and reduces computational cost by eliminating redundant Fourier transforms, performing spectral equalization, cross-correlation, and stacking in a single frequency-domain operation, and applying phase-only instrument-response correction after stacking. This last modification is sufficient because normalization and spectral equalization largely suppress amplitude information (Bensen et al., 2007; Ritzwoller and Feng, 2019), and it avoids the instability of full spectral division. While correctly removed instrument response is secondary to more significant sources of uncertainty in ANT, this optimization is simple to implement and improves robustness at minimal cost. We validate the method on datasets from three tectonic settings, benchmark it on multiple hardware platforms, and introduce a scalable variant that reuses station-level spectra across pairs. Even partial adoption within existing codes yields measurable gains in efficiency and stability, making the approach suitable for large-scale ANT applications.

2 Methodology

Bensen et al. (2007) proposed a workflow for calculating EGFs from ambient-noise cross-correlation: (1) remove linear trend; (2) taper; (3) remove instrument response; (4) trim; (5) apply time-domain normalization; (6) taper again; (7) perform spectral equalization; (8) cross-correlate; and (9) stack. While widely used, the order of steps (5) and (7) is critical. Applying time-domain normalization first—especially one-bit—suppresses high-amplitude transient signals, yielding a more stable time series for spectral analysis and reducing bias from frequency bands where they dominate. A potential drawback is that one-bit normalization can reduce the contribution of weaker frequencies, whose influence becomes limited to the zero crossings of stronger signals. However, this effect is much less pronounced with running absolute-mean normalization, which preserves relative amplitude variations. Performing spectral equalization after normalization also enables computational optimizations discussed later.

The most computationally demanding steps are typically I/O, response removal, spectral equalization, and cross-correlation, though their relative costs depend on hardware-specific factors such as memory access

speed and cache size. Response removal and spectral equalization are applied to individual time series in the frequency domain. Cross-correlation can be performed in either the time or frequency domain, but for long time series, the frequency-domain approach is more efficient, reducing the complexity from $O(nn_c)$ to $O(n \log_2 n)$, where n is the number of samples and n_c is the cross-correlogram length. For real-valued signals, using the Real FFT (RFFT) and its inverse (IRFFT) (Cooley and Tukey, 1965) further reduces the operation count. Limiting the maximum lag time reduces n_c and allows segmentation into overlapping windows, enabling the use of multiple short FFTs in place of a single long one. FFTs are typically more efficient when the input length is a power of two. After removing linear trends and applying a taper, the instrument response is removed by deconvolving the RFFT of the signal, $S(\omega)$, by the instrument response $R(\omega)$, followed by an IRFFT to recover the corrected time series $s_c(t)$, as described in Equation 1.

$$s_c(t) = \mathcal{F}^{-1} \left[\frac{S(\omega)}{R(\omega)} \right] \quad (1)$$

To stabilize the deconvolution, we restrict the operation to the broad frequency band where the instrument response is sufficiently sensitive by applying a band-pass filter—referred to as a pre-filter—prior to deconvolution, thereby eliminating the need for a water level (Clayton and Wiggins, 1976). The signal is then transformed back to the time domain for time-domain normalization, which mitigates the influence of transient or anomalous signals such as earthquakes, instrument glitches, and local noise (Bensen et al., 2007; Ritzwoller and Feng, 2019). Afterward, the signal is returned to the frequency domain for spectral equalization—typically combined with a spectral taper—to balance frequency contributions and suppress dominance by features such as the secondary microseismic peak (Bensen et al., 2007). A final IRFFT yields the preprocessed signal for cross-correlation.

Time-domain normalization methods differ in how strongly they suppress amplitude variations. One-bit normalization converts the signal to a binary time series by retaining only the sign of each sample. This approach is highly effective at suppressing large-amplitude transients (e.g., earthquakes, instrumental glitches), ensures uniform contribution across time windows, and is computationally efficient (Bensen et al., 2007). However, it discards all amplitude information and can significantly distort phase information, especially in low signal-to-noise ratio (SNR) conditions, which reduces the quality of EGFs. Running absolute-mean normalization preserves relative amplitude variations by dividing the signal by a moving average of its absolute value over a specified window (Bensen et al., 2007; Ritzwoller and Feng, 2019). Although more computationally intensive and somewhat sensitive to window length, it is better suited for datasets with stable ambient noise and moderate transients, as its less aggressive processing reduces waveform distortion, which can yield higher-quality EGFs. We adopted running absolute-mean normalization with

a window length equal to half the maximum wavelength (Bensen et al., 2007).

Among spectral equalization strategies, one commonly used method is aggressive spectral equalization, which completely flattens the amplitude spectrum over a selected frequency band by dividing the spectrum by its own absolute value (Ritzwoller and Feng, 2019). This procedure enhances the broadband character of the cross-correlation and suppresses dominant narrow-band noise, but it also eliminates physically meaningful amplitude information and may introduce non-physical artifacts (Fichtner et al., 2020). A more conservative alternative involves dividing the spectrum by a smoothed version of its amplitude, thereby reducing large spectral peaks without fully flattening the spectrum (Bensen et al., 2007; Ritzwoller and Feng, 2019). In this scheme, the equalization acts to mitigate spectral biases while preserving some spectral variability and physical interpretability. However, this method has been found to be less successful at removing some local persistent noise sources from cross-correlations between stations in certain cases (Ritzwoller and Feng, 2019), and determining the optimal amount of smoothing in the frequency domain can be challenging. Given that both strategies produced comparable results in our tests, we selected the more aggressive equalization procedure for its simpler implementation and robust performance.

Cross-correlation in the frequency domain involves computing the RFFT of signals from two stations ($s_1(t)$ and $s_2(t)$), multiplying the complex conjugate of the transform of the first signal with that of the second, and applying an IRFFT to obtain the cross-correlogram. Because the Fourier transform is linear, the signals can be normalized and transformed once (one RFFT per signal), after which spectral equalization, cross-correlation, and stacking are all performed in the frequency domain. A single IRFFT then converts the final stack back to the time domain.

A novel optimization we introduce consists in applying the phase-only instrument-response correction only once after stacking, rather than before correlation, to improve numerical stability and computational efficiency. Since time-domain normalization and spectral equalization are nonlinear operations that suppress amplitude information while preserving phase, the amplitude component of the instrument response becomes irrelevant. Correcting only the phase avoids potentially unstable spectral division—especially under noise, sub-optimal pre-filters, or poorly chosen water levels—and is applied once to the stacked spectrum, which is a much shorter signal. This involves multiplying the stacked spectrum by a complex exponential that accounts for the combined phase responses of the two instruments. The final EGF is then obtained via an inverse transform:

$$G_e = \mathcal{F}^{-1} \left\{ \left[\sum_{j=0}^N \overline{S_1(\omega)_j} S_2(\omega)_j \right] e^{-i[\phi_1(\omega) - \phi_2(\omega)]} \right\}, \quad (2)$$

where $S_1(\omega)_j$ and $S_2(\omega)_j$ are the normalized and equalized spectra of each signal pair, $\phi_1(\omega)$ and $\phi_2(\omega)$ are the instrument phase responses, and the overline de-

notes complex conjugation. Technically, assuming the cross-correlation function is symmetric about the origin, the EGF is the derivative of the cross-correlation function for $t > 0$, taken with a minus sign (Snieder, 2004; Lin et al., 2008). Nonetheless, many studies (e.g., Shapiro et al., 2005; Bensen et al., 2007) refer to the cross-correlogram itself (or a scaled version) as the EGF, a convention we adopt here.

We define the baseline approach—using RFFT, frequency-domain cross-correlation, and traditional response removal—as Workflow 1 (WF1), and the modified version—with phase-only correction applied after stacking and fewer transformations—as Workflow 2 (WF2). Figure 1 contrasts the two workflows.

In the time domain, cross-correlation with 100% overlap requires that $s_1(t)$ be longer than $s_2(t)$. If n_1 and n_2 are their respective lengths, then the resulting cross-correlogram has $n_{cc} = n_1 - n_2 + 1$ samples. This is known as the “valid mode”, which we adopt. In the frequency domain, the same result requires zero-padding $s_2(t)$ so that $n_1 = n_2 = n$ before applying the RFFT. For simplicity, we assume $n \approx 4n_{cc}$. Since both RFFT and IRFFT involve approximately $n \log_2 n$ operations, we can estimate the computational cost of each workflow following Figure 1 and using Equations 3–4:

$$\text{WF1} \rightarrow N \times \left[10 \times n \log_2 n + \frac{n}{4} \log_2 \left(\frac{n}{4} \right) \right], \quad (3)$$

$$\text{WF2} \rightarrow N \times 2 \times n \log_2 n + \frac{n}{4} \log_2 \left(\frac{n}{4} \right), \quad (4)$$

where N is the number of stacks. Note that the expressions in Equations 3–6 represent concrete operation counts (not asymptotic big-O bounds) used to estimate the numerical speed-up factor.

The theoretical speed-up factor SPF_T for replacing WF1 with WF2 is given by:

$$SPF_T := \frac{\text{WF1}}{\text{WF2}} \approx \frac{N \times \left[10 \times n \log_2 n + \frac{n}{4} \log_2 \left(\frac{n}{4} \right) \right]}{N \times 2 \times n \log_2 n + \frac{n}{4} \log_2 \left(\frac{n}{4} \right)} \quad (5)$$

Usually, N is on the order of hundreds or even thousands. Thus, we can approximate the number of operations for WF2 as $N \times 2 \times n \log_2 n$:

$$SPF_T \approx \frac{10 \times n \log_2 n + \frac{n}{4} [\log_2 n - 2]}{2 \times n \log_2 n} = 5 + \frac{1}{8} - \frac{1}{4 \log_2 n} \quad (6)$$

A typical n ranges from a few thousands to several hundreds of thousands. Hence, it is reasonable to assume that $0.01 \leq \frac{1}{4 \log_2 n} \leq 0.02 \Rightarrow \frac{1}{4 \log_2 n} \approx 0.015$. With this final approximation, we conclude that $SPF_T \approx 5.11$.

The actual speed-up depends not only on algorithmic complexity but also on hardware-specific factors such as cache size, memory bandwidth, and I/O performance, which can significantly affect execution time.

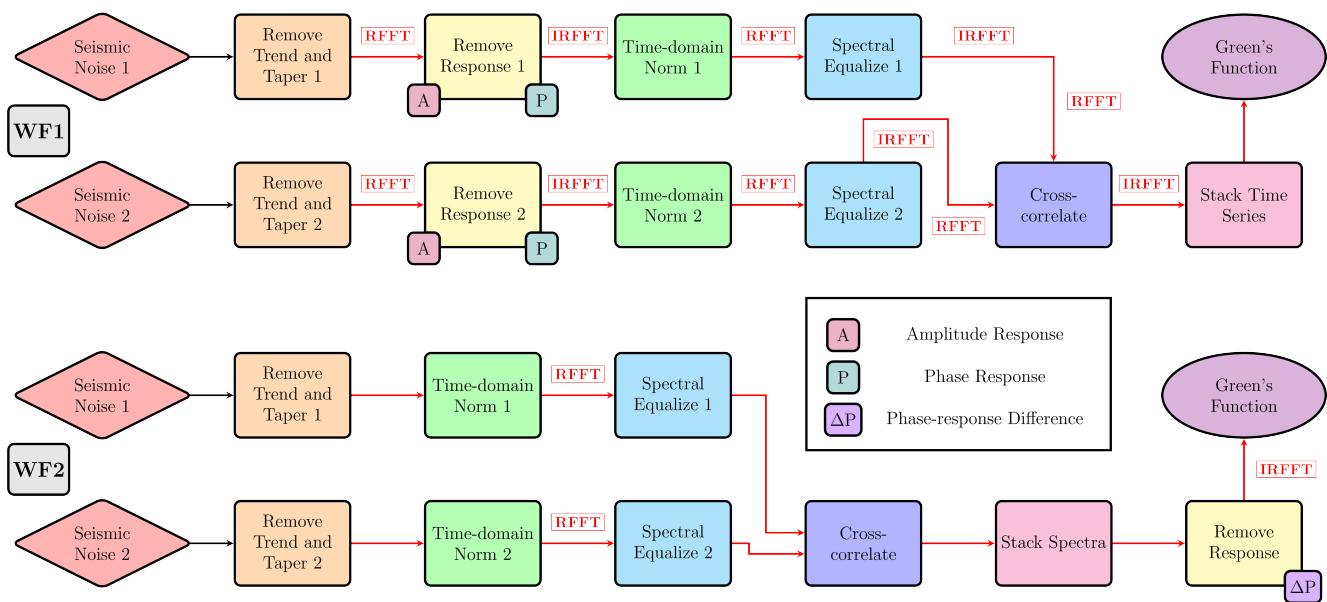


Figure 1 Flowchart comparing the two workflows. The standard workflow (WF1) removes both the amplitude (A) and phase (P) information of the instrument before time-domain normalization and spectral equalization. Our modified workflow (WF2) removes only the phase response at the end while also avoiding all redundant RFFTs and IRFFTs.

Additionally, the impact of workflow optimizations depends on the ratio between the number of stations and station pairs. While previous estimates assume comparable costs for operations performed per station and per pair, this holds primarily for small networks. In most practical cases, even when autocorrelations are excluded and symmetric pairs are not repeated, the number of station pairs grows quadratically with the number of stations:

$$n_p = \frac{n_s(n_s - 1)}{2}, \quad (7)$$

where n_s is the number of stations and n_p the number of unique station pairs. $n_p = O(n_s^2)$ in typical network configurations.

As n_s increases, execution time becomes increasingly dominated by pairwise operations such as cross-correlation and stacking, while station-level steps—such as response removal, time-domain normalization, and spectral equalization—contribute progressively less. This shift, however, depends on implementation. By reusing spectral representations of each station across its multiple pairings, the number of RFFTs can be kept proportional to n_s rather than n_p . Combined with stacking in the frequency domain, this strategy preserves linear complexity in n_s and improves scalability. The strategy of computing FFTs once per station and reusing them for all pairwise correlations is fully implemented in `SeisNoise.jl` (Clements and Denolle, 2020). Elements of this idea also appear in `Mirmex` (Fichtner et al., 2017), although the authors do not go as far as precomputing the FFTs. Instead, they perform several preprocessing operations (e.g., band-pass filtering, downsampling, detrending, and instrument-response removal) once and store the processed time-domain waveforms on disk for reuse,

thereby eliminating redundant preprocessing while still requiring a new FFT for each correlation pair.

As shown in our benchmarks, this approach reduces execution time by more than an order of magnitude even for moderate-sized networks ($n_s < 100$), without requiring parallelization. We refer to this optimized configuration as Workflow 3 (WF3), which produces EGFs identical to those of WF2 at substantially reduced computational cost. Its structure is illustrated in Figure 2.

3 Computational Experiments

For this study, we implemented the routines in Python and conducted most benchmarks on two machines. The first was a laptop with a 12th Generation Intel Core i9-12900HK processor (24 MB cache, 14 cores, 20 threads, up to 5.00 GHz), a 2 TB PCIe NVMe Gen4x4 SSD (M.2 2280), and 64 GB of DDR5 RAM (2 × 32 GB, 4800 MHz, dual-channel). The second was a Google Colab instance with an Intel Xeon processor, 2 vCPUs (virtual CPUs), and 13 GB of RAM. The code was based on a Seismo-Live tutorial on ANT (Hadzioannou and Rijal, 2024). While Python is not optimized for speed or memory usage, it is an accessible, open-source programming language that allowed us to validate the method and demonstrate its robustness. For production-scale runs, a compiled or fast just-in-time (JIT) compiled language—such as C, C++, Fortran, or Julia—would better leverage the efficiency of WF2, and especially WF3.

All tests used a sampling rate of 20 Hz and a 2-hour window length. Each continuous segment was detrended and tapered at its edges before gaps were filled with zeros, minimizing artifacts during instrument response removal and spectral equalization. Although the workflows are compatible with gap-skipping strategies,

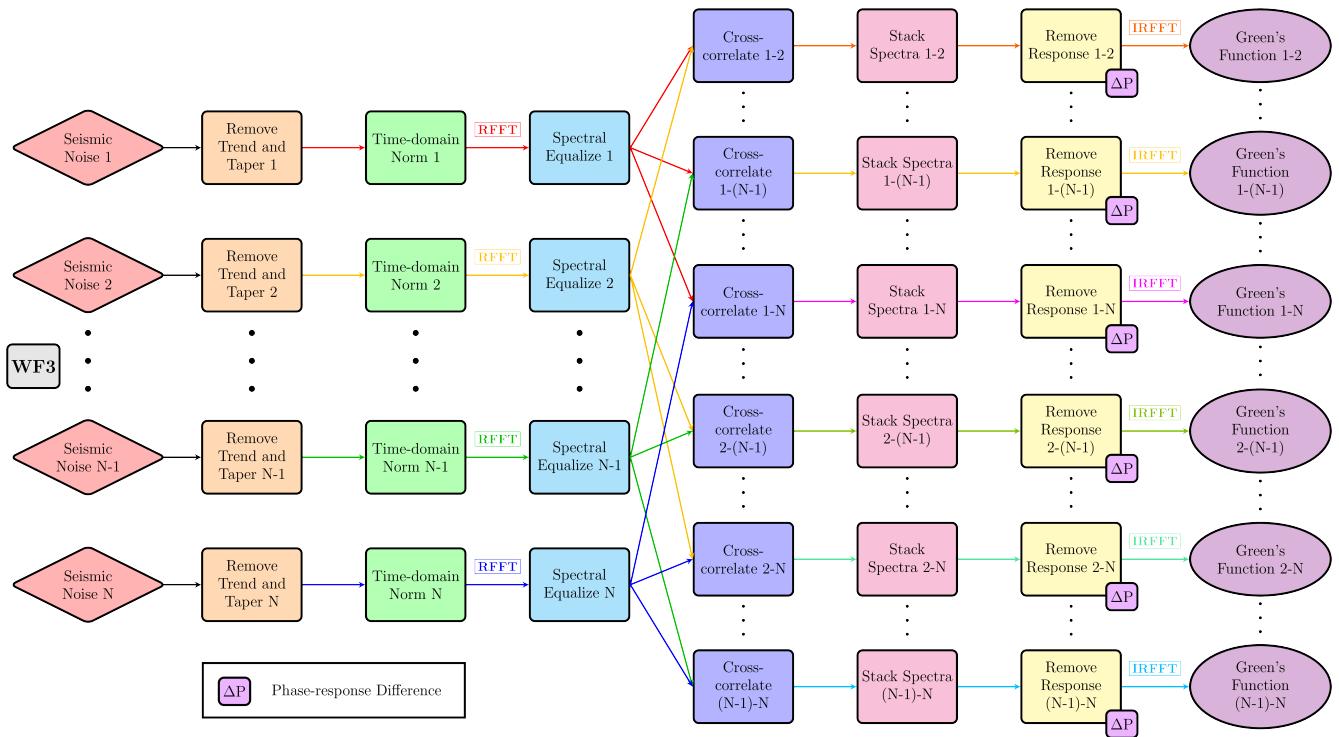


Figure 2 Schematic overview of Workflow 3 (WF3). Each RFFT is computed only once per station and reused across all cross-correlation pairs. In practice, two RFTTs per station are required to accommodate differences in time window lengths depending on whether the station appears first or second. Nevertheless, this strategy ensures that the number of RFTTs grows as $O(n_s)$ rather than $O(n_s^2)$.

we used zero-padding in this study for its simpler implementation. For actual applications, however, skipping gaps is generally a more reliable approach to avoid potential artifacts introduced by missing or corrupted data.

Figure 3 shows the EGFs computed with WF1 and WF2 using one year of ambient noise from two Southern California stations (Figure 4, right). In WF2, only the phase response is removed, resulting in stable corrections even at frequencies where instrument sensitivity is low. To allow comparison with WF1 at very low frequencies (< 0.01 Hz), we applied a pre-filter (dashed green in Figure 4, left) that is zero below 0.001 Hz (1000 s) and above 9 Hz (0.11 s), and unity between 0.0014 Hz (714 s) and 8.1 Hz (0.123 s). We compared results in two period bands: 0.11–1000 s and 10–30 s, including zoomed-in views to highlight differences. Figure 3 shows negligible differences between the results of the two workflows. WF3 plots are identical to those of WF2; the distinction lies solely in computational cost. We further quantified the differences by calculating the residuals between WF1 and WF2 (Figure S1, Supplementary Material) and the cosine similarity (CS) values for both period bands (see caption of Figure 3).

We repeated the tests for twelve station pairs in Brazil (Figure 5) using one year of continuous data. To accommodate the different instrument responses, we adjusted the retained frequency band to 0.002–2 Hz (500–0.5 s). This choice ensured stable computation of WF1 and allowed a meaningful comparison with WF2. The results again showed very close agreement between WF1

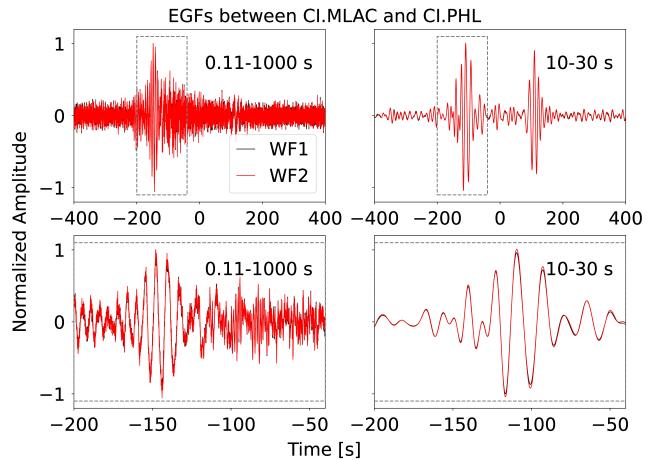


Figure 3 EGFs derived from one year of continuous data recorded at stations MLAC and PHL. WF1 is shown in black and WF2 in red; both waveforms were normalized to the maximum amplitude of WF1. Upper left: full broadband period range within instrument sensitivity; lower left: primary microseismic noise band (10–30 s); right column: zoomed-in views of the boxed regions in the left panels. Cosine similarity values measured in the 0.11–1000 s and 10–30 s period bands within the boxed regions are 0.97 and 1.00, respectively. Residuals for the same configurations are shown in Figure S1 (Supplementary Material) and exhibit markedly lower values within the primary microseismic band.

and WF2 (and thus WF3), confirming the robustness of the optimized workflows even when combining stations

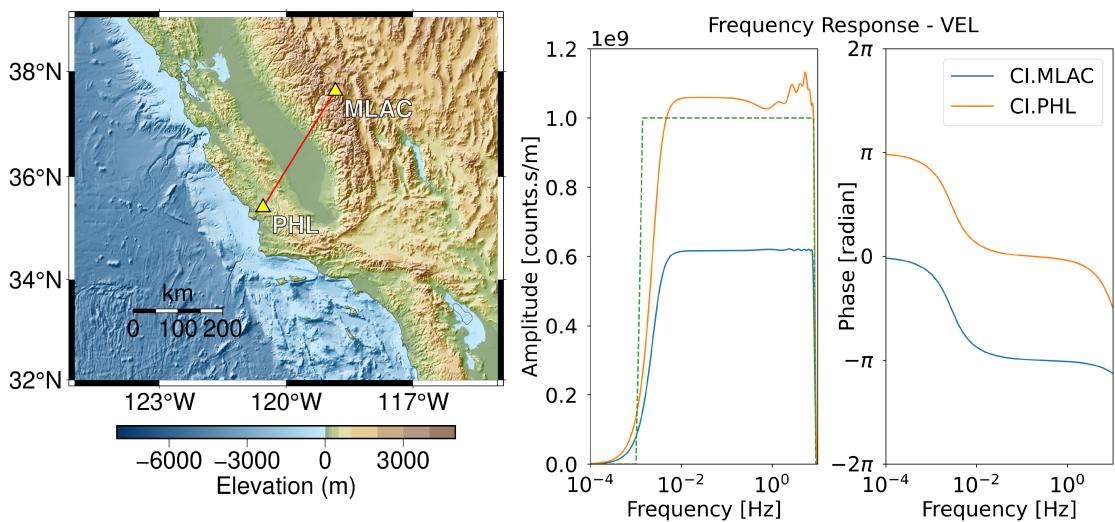


Figure 4 Left: Map of Southern California showing stations MLAC and PHL (network CI), also used by [Hadzioannou and Rijal \(2024\)](#) and [Shapiro et al. \(2005\)](#), marked as yellow triangles. The red line represents the great-circle path along which waves would propagate between the two stations. Right: Amplitude and phase response to ground velocity at stations MLAC (blue) and PHL (orange). The green dashed line indicates the broadest frequency band retained and has been scaled by 10^9 for visualization purposes.

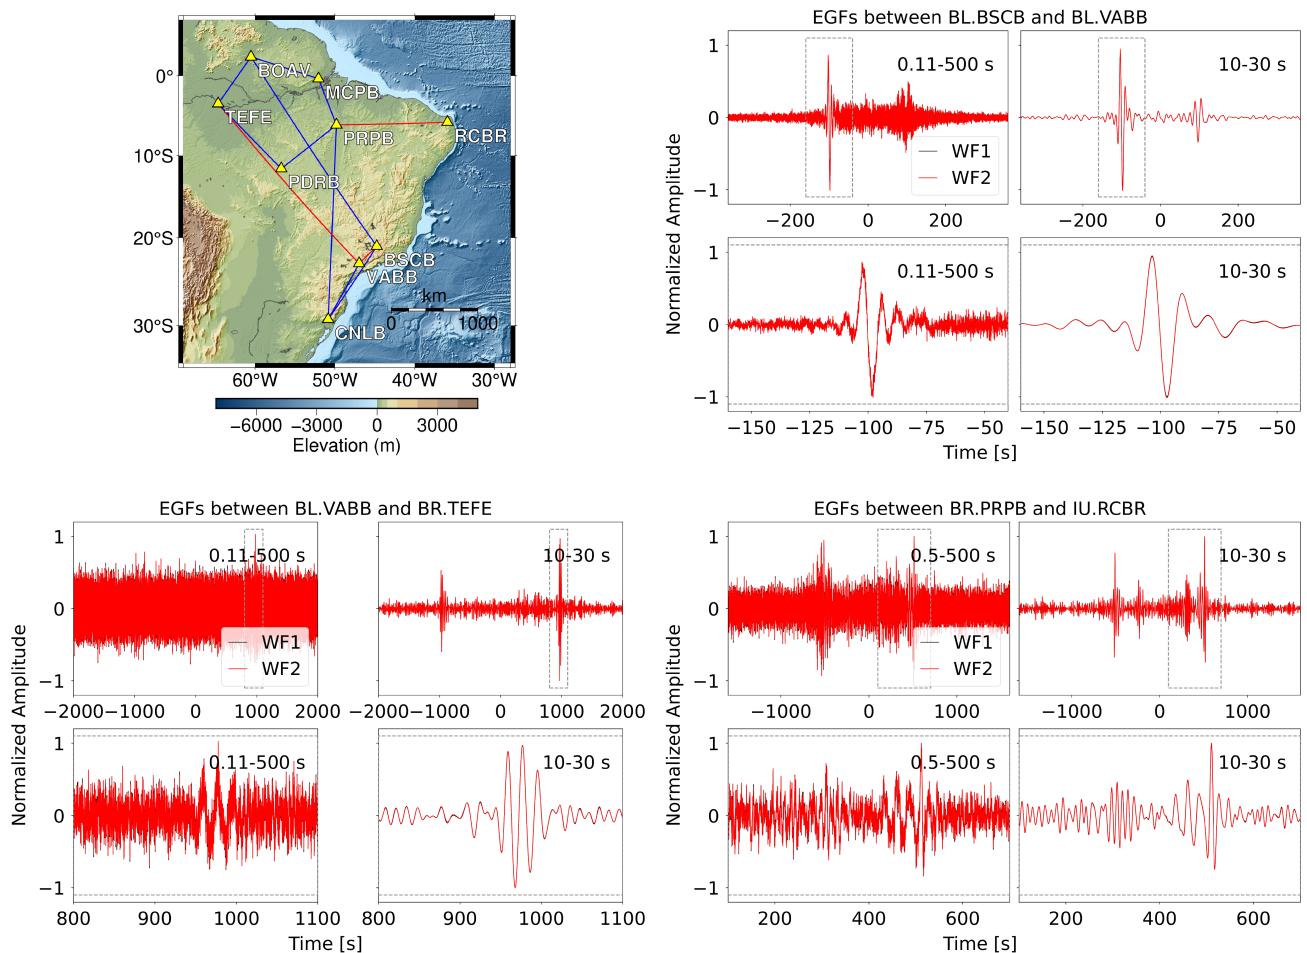


Figure 5 Upper-left panel: Analogous to Figure 4 (left), but for three station pairs in Brazil (networks BL, BR, and IU). The red lines correspond to the EGFs shown in the other three panels: BSCB–VABB (upper right), VABB–TEFE (lower left), and PRPB–RCBR (lower right); their residuals are displayed in Figure S1 (Supplementary Material). Blue lines indicate additional station pairs whose EGFs and residuals are presented in Figures S2–S5.

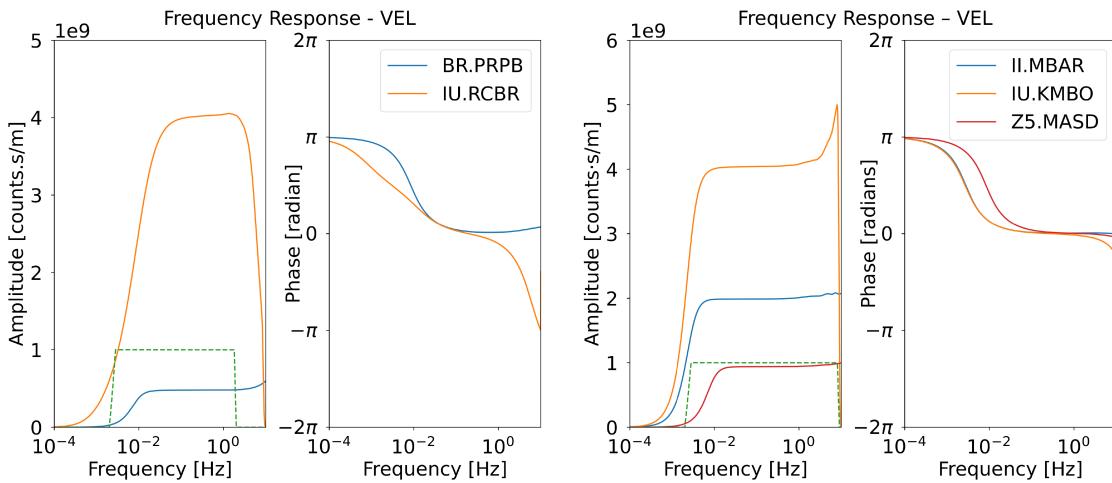


Figure 6 Left: Instrument response (amplitude and phase) to ground velocity at stations PRPB (blue) and RCBR (orange). As in Figure 4, the green dashed line marks the retained frequency band. The minimum frequency was raised to 0.002 Hz because of the low sensitivity of PRPB below this value, whereas the maximum frequency was lowered to 2 Hz owing to the sharp decline in sensitivity of RCBR above this threshold. Right: Same as the left panel, but for the three Ugandan stations MBAR (blue), KMBO (orange), and MASD (red).

Data days	Elapsed Time WF1 [s]	Elapsed Time WF2 [s]	Speed-up Factor	Execution Time Drop [%]
10	55.4	16.8	3.29	69.7
90	416.2	134.3	3.10	67.7
365	1694.5	565.0	3.00	66.7

Table 1 Elapsed times for WF1 and WF2 over different numbers of days. This benchmark was conducted on a laptop with a 12th Generation Intel Core i9-12900HK processor (24 MB cache, 14 cores, 20 threads, up to 5.00 GHz), a 2 TB PCIe NVMe Gen4x4 SSD (M.2 2280), and 64 GB of DDR5 RAM (2×32 GB, 4800 MHz, dual-channel). The reported times are averages over two runs on a single core.

Data days	Elapsed Time WF1 [s]	Elapsed Time WF2 [s]	Speed-up Factor	Execution Time Drop [%]
10	150.0	38.7	3.88	74.2
90	1147.3	292.9	3.92	74.5
365	4353.7	1158.7	3.76	73.4

Table 2 Elapsed times for WF1 and WF2 over different numbers of days. This benchmark was conducted on Google Cloud (Colab) using an Intel Xeon with 2 vCPUs (virtual CPUs) and 13 GB of RAM.

from different networks equipped with markedly different sensors (Figure 6, left). For example, station PRPB (network BR) uses a *Trillium 120P*, 120 s, 1201 V/m/s – *Trident 305*, 40 V, whereas station RCBR (network IU) is equipped with a *Geotech KS-54000 borehole seismometer*; nevertheless, the resulting EGFs exhibit no significant differences.

Quantitatively, CS values across the full retained period band remained above 0.98 for nearly all pairs. The only exception was the PRPB–RCBR pair with CS = 0.90, which is still high. Within the primary microseismic band (10–30 s), CS values were essentially 1.00 in all cases. The residuals confirm these observations: they remain small across nearly all pairs even over the 0.002–2 Hz frequency band. Only when instrument responses differ significantly, as in the PRPB–RCBR pair, residuals become noticeably larger—but only above approximately 1 Hz, which is well above the primary mi-

croseismic band's frequencies. Below 1 Hz, the EGFs are nearly identical across all frequencies, with CS reaching 0.98 between 0.002 Hz and 1 Hz.

While WF2 offers substantial speed-up (Tables 1 and 2), integrating it into existing codes may require extensive modifications. However, partial adoption is straightforward. We tested a minimal modification of an existing ANT code (Kabanda et al., 2023) by applying only the phase-only response correction at the end of the workflow. Benchmarks (Table 3) were run on Ugandan station pairs using a third machine: a laptop with a 13th Gen Intel Core i9-13900H processor (24 MB cache, 14 cores, 20 threads, up to 5.40 GHz), a 1 TB PCIe Gen4 SSD, and 32 GB of DDR5 RAM. We refer to the original and modified workflows as Workflow A (WFA) and Workflow B (WFB), respectively. WFA resembles WF1 but lacks most performance optimizations; however, it is more flexible, handling data gaps by skipping them

Data days	Elapsed Time WFA [s]	Elapsed Time WFB [s]	Speed-up Factor	Execution Time Drop [%]
10	370.8	332.0	1.12	10.5
90	3321.4	3024.6	1.10	8.9
365	13753.0	12369.4	1.11	10.1

Table 3 Elapsed times for WFA and WFB over different numbers of days. This benchmark was conducted on a laptop with a 13th Generation Intel Core i9-13900H processor (24 MB cache, 14 cores, 20 threads, up to 5.40 GHz), a 1 TB PCIe NVMe Gen4x4 SSD (M.2 2280), and 32 GB of DDR5 RAM (2x16 GB, 4800 MHz, dual-channel).

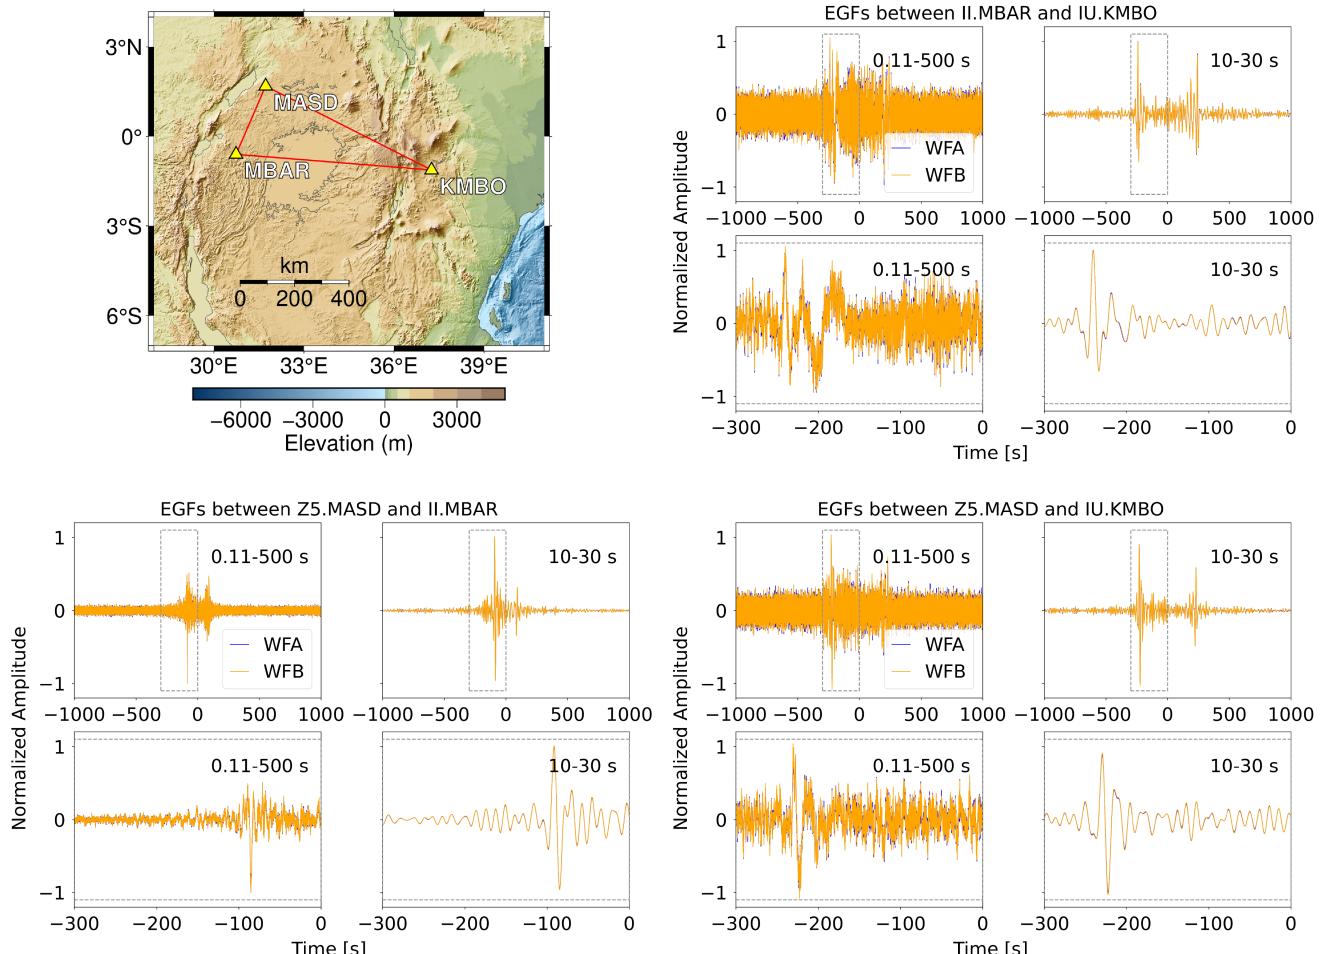


Figure 7 Analogous to Figure 5, but for three station pairs in Uganda (networks II, IU, and Z5). The upper-left panel shows the regional map, whereas the remaining three panels display the EGFs computed from one year of continuous data using the original code of Kabanda et al. (2023) (WFA in blue, WFB in yellow). The corresponding residuals between WFA and WFB are shown in Figure S6 (Supplementary Material).

rather than filling with zeros. Figure 7 shows results for three station pairs.

The three Ugandan stations used in this test belong to different networks and are equipped with different sensor types (Figure 6, right). Station MBAR (network II) uses a Streckeisen STS-6 seismometer, KMBO (network IU) a Streckeisen STS-1VBB with E300 electronics, and MASD (network Z5) a Streckeisen STS-2 Generation 3. CS values exceed 0.98 across the entire retained period band and approach 1.00 in the primary microseismic band (10–30 s). The low residual amplitudes shown in Figure S6 (Supplementary Material) further confirms the robustness of the phase-only response-correction opti-

mization.

These results confirm that the optimized workflows remain stable and effective across markedly different instruments. This stability, however, depends on the use of running absolute-mean normalization with the window length defined earlier. When one-bit normalization is applied instead, its strong nonlinearity severely distorts the phase content once the instrument response is deconvolved at the end of the workflow. For example, in the case shown in Figure 3, CS values within the boxed regions drop to 0.77 (0.11–1000 s band) and 0.93 (10–30 s band), respectively.

n_s	n_p	Time WF1 [s]	Time WF2 [s]	Time WF3 [s]	Speed-up (WF1/WF3)	Speed-up (WF2/WF3)
2	1	85.49	53.76	58.22	1.47	0.92
3	3	256.15	157.99	87.94	2.91	1.80
4	6	520.08	320.57	119.15	4.36	2.69
9	36	2932.48	1767.00	250.86	11.69	7.04
45	990	79011.34	46259.41	1286.60	61.41	35.95
Idealized scenario: station-wise preprocessing assumed to have zero cost						
2	1	3.62	3.34	5.47	0.66	0.61
3	3	11.30	10.92	9.28	1.22	1.18
4	6	22.18	20.93	12.55	1.77	1.67
9	36	150.97	144.67	34.62	4.36	4.18
45	990	4476.11	3957.23	293.01	15.28	13.51

Table 4 Execution times and speed-up factors for WF1, WF2, and WF3 across different numbers of stations (n_s) and station pairs (n_p). The top block assumes naive implementations of WF1 and WF2, where all preprocessing is repeated for each station pair. The bottom block corresponds to an idealized scenario in which all preprocessing steps that can be performed station-wise are assumed to have negligible cost.

4 Computational Benchmarks

In this section, we benchmark WF1 and WF2 on three hardware configurations—two laptops and a Google Cloud instance—using a 20 Hz sampling rate, 2-hour windows, and a maximum lag of 2000 s (Tables 1 and 2). We also benchmark WF3 on one of the laptops (Table 4), and workflows WFA and WFB on the other (Table 3). Although the observed speed-up factors for WF2 over WF1 are lower than theoretical estimates, they consistently show execution time reductions of 67–75%. Contrary to expectation, the speed-up decreases with the number of days stacked, likely due to I/O and hardware bottlenecks offsetting the computational savings. Translating the code to a faster programming language and optimizing I/O may mitigate these limitations. Runtime and speed-up are strongly hardware-dependent: absolute runtimes on the laptop are roughly three times shorter, but speed-up factors are larger on the Google Cloud instance.

Given the performance gains of WF2 over WF1, we recommend adopting all proposed optimizations. However, because fully integrating WF2 into existing ANT packages may require extensive changes, we also evaluated a partial implementation. Using the ANT package by [Kabanda et al. \(2023\)](#), we compared its default workflow (WFA) with a modified version (WFB), where only the phase response is removed after stacking, rather than performing full response removal on the raw data. This minor change yielded an $\sim 10\%$ reduction in execution time (Table 3) and simplified the workflow by removing parameter choices and reducing potential sources of numerical instability.

As discussed earlier, $n_p \gg n_s$ in most cases, which reduces the relative advantage of WF2 over WF1. However, WF3 addresses this by further optimizing WF2 for pairwise operations. Table 4 compares WF1, WF2, and WF3 using ten days of data, with earlier benchmarks addressing longer time spans. Speed-up factors depend not only on hardware but also on how the baseline work-

flows are implemented. To capture this variability, we benchmarked WF3 against two scenarios: a naive implementation where preprocessing is repeated for every pair, and a second case where all station-level preprocessing steps are disabled, simulating zero cost for those operations. These serve as endmembers, illustrating the range of expected speed-up factors.

Table 4 highlights the efficiency of WF3 over WF1 and WF2, particularly for larger networks. In the naive case, WF3 achieves speed-up factors over $60\times$ relative to WF1 and $35\times$ relative to WF2 for 45 stations. Even in the idealized scenario where per-station preprocessing is costless, WF3 still yields gains of about $15\times$ and $13\times$, respectively. These results show that shifting the computational load from pairwise to per-station operations ensures better scalability without requiring parallelization. Further improvements are possible by using high-performance FFT libraries such as FFTW ([Frigo, 1999](#)) or cuFFT ([NVIDIA Corporation, 2025](#)), which optimize FFT execution plans that can be reused across all operations due to the uniform signal length.

5 Conclusions

We developed optimized workflows for ambient-noise cross-correlation that improve computational efficiency and numerical stability. WF2 streamlines the standard method by reducing redundant Fourier transforms and applying phase-only response correction after stacking, while WF3 further enhances scalability by reusing station-level spectra across pairs. These changes simplify implementation, reduce execution time, and make large-scale ANT processing more accessible.

We demonstrated that WF2 produces EGFs that are nearly identical to those from the standard workflow (WF1), while significantly reducing execution time for a single station pair. Theoretical analysis predicts a speed-up factor of approximately 5.1, and empirical benchmarks confirm reductions in runtime between

67% and 75% (speed-up factors from 3.0 to 3.9), depending on hardware and the number of stacked days. These gains are achieved without requiring parallelization, which is particularly valuable for researchers working in regions with limited access to high-performance infrastructure. We also evaluated a partial implementation of WF2 (WFB) within an existing codebase, requiring only minor modifications. This version consistently reduced execution time by $\sim 10\%$ and improved numerical stability by avoiding unstable spectral division, showing that even partial adoption of WF2 can yield measurable benefits with minimal effort.

WF3 offers even greater efficiency, particularly in realistic scenarios where the number of station pairs far exceeds the number of stations. Benchmarks show speed-up factors exceeding $60\times$ relative to WF1 and $35\times$ relative to WF2 for networks with 45 stations. Even in idealized conditions where station-level preprocessing is assumed to have negligible cost, WF3 yields speed-ups of $15\times$ or more relative to WF1. These results demonstrate that minimizing redundant operations across station pairs by reusing station-level spectral representations is critical for maintaining performance as network size increases.

The proposed workflows increase the efficiency and robustness of ambient-noise processing. They are simple to implement, compatible with existing frameworks, and particularly well-suited for large datasets. Additional gains are possible by adopting high-performance FFT libraries and using compiled or JIT-compiled languages. These improvements can make large-scale ANT more accessible and scalable, supporting applications from local studies to global arrays.

6 Data and Code Availability

The Python routines and datasets used for Workflows 1, 2, and 3 are available at <https://doi.org/10.5281/zenodo.16785574>. Seismic data for stations BSCB, BOAV, PRPB, CNLB, MCPB, TEFE, VABB, PDRB, and RCBR were obtained from the BL, BR, and IU networks: <https://ds.iris.edu/ma/BL/>, <https://ds.iris.edu/ma/BR/>, and <https://ds.iris.edu/ma/IU/>, respectively.

Acknowledgments

We sincerely thank two anonymous reviewers for their thorough and constructive comments, which have greatly improved the clarity, rigor, and overall quality of this manuscript.

We acknowledge the networks and institutions that provided the seismic data used in this study. Data for stations MLAC and PHL were obtained from the Southern California Seismic Network (CI) (California Institute of Technology and United States Geological Survey Pasadena, 1926). Data for station RCBR and KMBO were provided by the Global Seismograph Network (IU) (Albuquerque Seismological Laboratory/USGS, 1988). Data for station MBAR were obtained from the Global Seismograph Network (II) (EarthScope Consortium, 1986).

Data for station MASD were provided by the Dry Rifting in the Albertine–Rhino Graben (DRIAR) project (Albert Kabanda et al., 2022).

Data for most Brazilian stations were provided by the Brazilian Seismographic Network (RSBR), which includes contributions from the BL and BR networks. RSBR is a collaborative initiative among the University of São Paulo (USP), University of Brasília (UnB), Federal University of Rio Grande do Norte (UFRN), and the National Observatory (ON), and is currently coordinated by the Brazilian Geological Survey (CPRM) (Bianchi et al., 2018).

References

Aki, K., Christoffersson, A., and Husebye, E. S. Determination of the three-dimensional seismic structure of the lithosphere. *Journal of Geophysical Research*, 82(2):277–296, Jan. 1977. doi: [10.1029/jb082i002p00277](https://doi.org/10.1029/jb082i002p00277).

Akkar, S. and Boore, D. M. On Baseline Corrections and Uncertainty in Response Spectra for Baseline Variations Commonly Encountered in Digital Accelerograph Records. *Bulletin of the Seismological Society of America*, 99(3):1671–1690, June 2009. doi: [10.1785/0120080206](https://doi.org/10.1785/0120080206).

Albert Kabanda, Suzan van der Lee, Lawrence Kabenge, Joseph Nyago, Fred Tugume, Geraldine Paula Babirye, and Yoweri Nseko. Dry Rifting In the Albertine–Rhino graben (DRIAR), Uganda, 2022. doi: [10.7914/SN/Z5_2022](https://doi.org/10.7914/SN/Z5_2022).

Albuquerque Seismological Laboratory/USGS. Global Seismograph Network (GSN - IRIS/USGS), 1988. doi: [10.7914/SN/IU](https://doi.org/10.7914/SN/IU).

Almeida, F. F. M. d., Brito Neves, B. B. d., and Dal Ré Carneiro, C. The origin and evolution of the South American Platform. *Earth-Science Reviews*, 50(1-2):77–111, May 2000. doi: [10.1016/s0012-8252\(99\)00072-0](https://doi.org/10.1016/s0012-8252(99)00072-0).

Artemieva, I. *The Lithosphere: An Interdisciplinary Approach*. Cambridge University Press, July 2011. doi: [10.1017/cbo9780511975417](https://doi.org/10.1017/cbo9780511975417).

Audhkhasi, P. and Singh, S. C. Discovery of distinct lithosphere–asthenosphere boundary and the Gutenberg discontinuity in the Atlantic Ocean. *Science Advances*, 8(24), June 2022. doi: [10.1126/sciadv.abn5404](https://doi.org/10.1126/sciadv.abn5404).

Bensen, G. D., Ritzwoller, M. H., Barmin, M. P., Levshin, A. L., Lin, F., Moschetti, M. P., Shapiro, N. M., and Yang, Y. Processing seismic ambient noise data to obtain reliable broad-band surface wave dispersion measurements. *Geophysical Journal International*, 169(3):1239–1260, June 2007. doi: [10.1111/j.1365-246x.2007.03374.x](https://doi.org/10.1111/j.1365-246x.2007.03374.x).

Bianchi, M. B., Assumpção, M., Rocha, M. P., Carvalho, J. M., Azevedo, P. A., Fontes, S. L., Dias, F. L., Ferreira, J. M., Nascimento, A. F., Ferreira, M. V., and Costa, I. S. L. The Brazilian Seismographic Network (RSBR): Improving Seismic Monitoring in Brazil. *Seismological Research Letters*, 89(2A):452–457, Feb. 2018. doi: [10.1785/0220170227](https://doi.org/10.1785/0220170227).

California Institute of Technology and United States Geological Survey Pasadena. Southern California Seismic Network, 1926. doi: [10.7914/SN/CI](https://doi.org/10.7914/SN/CI).

Cathles, L., Fjeldskar, W., Lenardic, A., Romanowicz, B., Seales, J., and Richards, M. Influence of the asthenosphere on earth dynamics and evolution. *Scientific Reports*, 13(1), Aug. 2023. doi: [10.1038/s41598-023-39973-y](https://doi.org/10.1038/s41598-023-39973-y).

Ciardelli, C., Assumpção, M., Bozdağ, E., and van der Lee, S. Adjoint Waveform Tomography of South America. *Journal of*

Geophysical Research: Solid Earth, 127(2), Jan. 2022. doi: 10.1029/2021jb022575.

Clayton, R. W. and Wiggins, R. A. Source shape estimation and deconvolution of teleseismic bodywaves. *Geophysical Journal International*, 47(1):151–177, Oct. 1976. doi: 10.1111/j.1365-246x.1976.tb01267.x.

Clements, T. and Denolle, M. A. SeisNoise.jl: Ambient Seismic Noise Cross Correlation on the CPU and GPU in Julia. *Seismological Research Letters*, 92(1):517–527, Sept. 2020. doi: 10.1785/0220200192.

Clouzet, P., Masson, Y., and Romanowicz, B. Box Tomography: first application to the imaging of upper-mantle shear velocity and radial anisotropy structure beneath the North American continent. *Geophysical Journal International*, 213(3):1849–1875, Feb. 2018. doi: 10.1093/gji/ggy078.

Cooley, J. W. and Tukey, J. W. An algorithm for the machine calculation of complex Fourier series. *Mathematics of Computation*, 19 (90):297–301, 1965. doi: 10.1090/s0025-5718-1965-0178586-1.

Dziewonski, A. M. Elastic and anelastic structure of the Earth. *Reviews of Geophysics*, 17(2):303–312, Apr. 1979. doi: 10.1029/rg017i002p00303.

EarthScope Consortium. Global Seismograph Network - II, 1986. doi: 10.7914/SN/II.

Fichtner, A., Ermert, L., and Gokhberg, A. Seismic Noise Correlation on Heterogeneous Supercomputers. *Seismological Research Letters*, 88(4):1141–1145, May 2017. doi: 10.1785/0220170043.

Fichtner, A., Bowden, D., and Ermert, L. Optimal processing for seismic noise correlations. *Geophysical Journal International*, 223(3):1548–1564, Aug. 2020. doi: 10.1093/gji/ggaa390.

Fishwick, S. Surface wave tomography: Imaging of the lithosphere–asthenosphere boundary beneath central and southern Africa? *Lithos*, 120(1–2):63–73, Nov. 2010. doi: 10.1016/j.lithos.2010.05.011.

Frigo, M. A fast Fourier transform compiler. In *Proceedings of the ACM SIGPLAN 1999 conference on Programming language design and implementation*, PLDI99, page 169–180. ACM, May 1999. doi: 10.1145/301618.301661.

Hadzioannou, C. and Rijal, A. Noise Correlation Wrapper: Ambient Seismic Noise Correlation, 2024. https://seismo-live.github.io/html/Ambient%20Seismic%20Noise/NoiseCorrelation_wrapper.html. Accessed: 2024-08-01.

Havskov, J. and Alguacil, G. *Correction for Instrument Response*, page 197–230. Springer International Publishing, 2016. doi: 10.1007/978-3-319-21314-9_6.

Herrmann, R. B. Computer Programs in Seismology: An Evolving Tool for Instruction and Research. *Seismological Research Letters*, 84(6):1081–1088, Oct. 2013. doi: 10.1785/0220110096.

Jiang, C. and Denolle, M. A. NoisePy: A New High-Performance Python Tool for Ambient-Noise Seismology. *Seismological Research Letters*, 91(3):1853–1866, Apr. 2020. doi: 10.1785/0220190364.

Kabanda, A., Alonso, B., van der Lee, S., Stamps, D. S., Atekwana, E. A., Nyago, J., Kabenge, L., Tugume, F., Fishwick, S., Kola-wole, F., Evans, R. L., Taylor, M. H., Katumwehe, A. B., Atekwana, E. A., and Kiberu, J. M. Interferometry of Ambient Seismic Noise Recorded by DRIAR Stations in the Northern Western Branch of the East African Rift System, Uganda, Dec. 2023. AGU Fall Meeting 2023, San Francisco, CA, 11–15 December 2023, Session: Tectonophysics / From Continental Rifts to Rifted Margins III, Poster No. T13C-0234.

Kumar, P., Yuan, X., Kind, R., and Mechie, J. The lithosphere–asthenosphere boundary observed with USArray receiver functions. *Solid Earth*, 3(1):149–159, May 2012. doi: 10.5194/se-3-149-2012.

Larose, E., Margerin, L., Derode, A., van Tiggelen, B., Campillo, M., Shapiro, N., Paul, A., Stehly, L., and Tanter, M. Correlation of random wavefields: An interdisciplinary review. *Geophysics*, 71 (4):SI11–SI21, Jan. 2006. doi: 10.1190/1.2213356.

Lecocq, T., Caudron, C., and Brenguier, F. MSNoise, a Python Package for Monitoring Seismic Velocity Changes Using Ambient Seismic Noise. *Seismological Research Letters*, 85(3):715–726, May 2014. doi: 10.1785/0220130073.

Li, Z., Zhou, J., Wu, G., Wang, J., Zhang, G., Dong, S., Pan, L., Yang, Z., Gao, L., Ma, Q., Ren, H., and Chen, X. CC-FJPy: A Python Package for seismic ambient noise cross-correlation and the frequency-Bessel transform method. Feb. 2021. doi: 10.1002/essoar.10506115.1.

Lin, F.-C. and Ritzwoller, M. H. Apparent anisotropy in inhomogeneous isotropic media. *Geophysical Journal International*, 186(3):1205–1219, July 2011. doi: 10.1111/j.1365-246x.2011.05100.x.

Lin, F.-C., Moschetti, M. P., and Ritzwoller, M. H. Surface wave tomography of the western United States from ambient seismic noise: Rayleigh and Love wave phase velocity maps. *Geophysical Journal International*, 173(1):281–298, Apr. 2008. doi: 10.1111/j.1365-246x.2008.03720.x.

Magrini, F., Lauro, S., Kästle, E., and Boschi, L. Surface-wave tomography using SeisLib: a Python package for multiscale seismic imaging. *Geophysical Journal International*, 231(2): 1011–1030, July 2022. doi: 10.1093/gji/ggac236.

Makus, P. and Sens-Schönfelder, C. SeisMIC - an Open Source Python Toolset to Compute Velocity Changes from Ambient Seismic Noise. *Seismica*, 3(1), Feb. 2024. doi: 10.26443/seismica.v3i1.1099.

Martin, M. and Wenzel, F. High-resolution teleseismic body wave tomography beneath SE-Romania - II. Imaging of a slab detachment scenario. *Geophysical Journal International*, 164(3): 579–595, Mar. 2006. doi: 10.1111/j.1365-246x.2006.02884.x.

Moschetti, M. P., Ritzwoller, M. H., and Shapiro, N. M. Surface wave tomography of the western United States from ambient seismic noise: Rayleigh wave group velocity maps. *Geochemistry, Geophysics, Geosystems*, 8(8), Aug. 2007. doi: 10.1029/2007gc001655.

Muir, J. B. and Ross, Z. E. A deep Gaussian process model for seismicity background rates. *Geophysical Journal International*, 234 (1):427–438, Feb. 2023. doi: 10.1093/gji/ggad074.

Nolet, G. *A Breviary of Seismic Tomography: Imaging the Interior of the Earth and Sun*. Cambridge University Press, Sept. 2008. doi: 10.1017/cbo9780511984709.

Nolet, G. and Kennett, B. The interaction of the S-wavefield with upper mantle heterogeneity. *Geophysical Journal International*, 101(3):751–762, June 1990. doi: 10.1111/j.1365-246x.1990.tb05581.x.

NVIDIA Corporation. *cufft Library*. NVIDIA, 2025. <https://docs.nvidia.com/cuda/cufft/>. Part of the CUDA Toolkit.

Ojo, A. ROSES 2021 Session 03: Ambient Noise Tomography, 2021. <https://www.youtube.com/watch?v=xQyV-oZju>. YouTube video, posted by ROSES.

Petersen, M. D., Harmsen, S. C., Jaiswal, K. S., Rukstales, K. S., Luco, N., Haller, K. M., Mueller, C. S., and Shumway, A. M. Seismic Hazard, Risk, and Design for South America. *Bulletin of the Seismological Society of America*, Jan. 2018. doi: 10.1785/0120170002.

Plomerová, J. and Babuška, V. Long memory of mantle lithosphere fabric — European LAB constrained from seismic anisotropy. *Lithos*, 120(1–2):131–143, Nov. 2010. doi: 10.1016/j.lithos.2010.01.008.

Priestley, K. and Tilman, F. Relationship between the upper mantle high velocity seismic lid and the continental lithosphere. *Lithos*, 109(1-2):112–124, Apr. 2009. doi: 10.1016/j.lithos.2008.10.021.

Prieto, G. A., Lawrence, J. F., and Beroza, G. C. Anelastic Earth structure from the coherency of the ambient seismic field. *Journal of Geophysical Research: Solid Earth*, 114(B7), July 2009. doi: 10.1029/2008jb006067.

Rawlinson, N. and Sambridge, M. *Seismic traveltimes tomography of the crust and lithosphere*, page 81–198. Elsevier, 2003. doi: 10.1016/s0065-2687(03)46002-0.

Ritzwoller, M. H. and Feng, L. *Overview of Pre- and Post-Processing of Ambient Noise Correlations*, page 144–187. Cambridge University Press, Apr. 2019. doi: 10.1017/9781108264808.007.

Ritzwoller, M. H., Lin, F.-C., and Shen, W. Ambient noise tomography with a large seismic array. *Comptes Rendus. Géoscience*, 343 (8–9):558–570, June 2011. doi: 10.1016/j.crte.2011.03.007.

Routh, P., Neelamani, R., Lu, R., Lazaratos, S., Braaksma, H., Hughes, S., Saltzer, R., Stewart, J., Naidu, K., Averill, H., Gotumukkula, V., Homonko, P., Reilly, J., and Leslie, D. Impact of high-resolution FWI in the Western Black Sea: Revealing overburden and reservoir complexity. *The Leading Edge*, 36(1): 60–66, Jan. 2017. doi: 10.1190/tle36010060.1.

Scherbaum, F. *Inverse and simulation filtering of digital seismograms*, page 132–160. Springer Netherlands, 1996. doi: 10.1007/978-94-010-9572-3_9.

Shapiro, N. M., Campillo, M., Stehly, L., and Ritzwoller, M. H. High-Resolution Surface-Wave Tomography from Ambient Seismic Noise. *Science*, 307(5715):1615–1618, Mar. 2005. doi: 10.1126/science.1108339.

Snieder, R. Extracting the Green's function from the correlation of coda waves: A derivation based on stationary phase. *Physical Review E*, 69(4), Apr. 2004. doi: 10.1103/physreve.69.046610.

Stehly, L., Campillo, M., and Shapiro, N. M. A study of the seismic noise from its long-range correlation properties. *Journal of Geophysical Research: Solid Earth*, 111(B10), Oct. 2006. doi: 10.1029/2005jb004237.

Thurber, C. H. *Seismic Tomography of the Lithosphere with Body Waves*, page 717–737. Birkhäuser Basel, 2003. doi: 10.1007/978-3-0348-8010-7_12.

Vaddineni, V. A. and Singh, S. C. The lithosphere–asthenosphere boundary structure at 11–21 Ma from wide-angle seismic data in the equatorial Atlantic Ocean. *Geophysical Journal International*, 235(3):2743–2757, Sept. 2023. doi: 10.1093/gji/ggad392.

Weaver, R. L. and Lobkis, O. I. Diffuse fields in open systems and the emergence of the Green's function (L). *The Journal of the Acoustical Society of America*, 116(5):2731–2734, Nov. 2004. doi: 10.1121/1.1810232.

Witek, M., van der Lee, S., Kang, T., Chang, S., Ning, J., and Ning, S. S Velocity Model of East Asia From a Cluster Analysis of Localized Dispersion. *Journal of Geophysical Research: Solid Earth*, 123 (11):9712–9732, Nov. 2018. doi: 10.1029/2018jb016060.

Yang, Y. and Ritzwoller, M. H. Characteristics of ambient seismic noise as a source for surface wave tomography. *Geochemistry, Geophysics, Geosystems*, 9(2), Feb. 2008. doi: 10.1029/2007gc001814.

The article *Optimized Workflows with a Single Phase-Only Response Correction for Building Empirical Green's Functions for Ambient-Noise Tomography* © 2026 by Caio Ciardelli is licensed under CC BY 4.0.