Response to Editor and Reviewers

Dear Dr. Koelemeijer,

Thank you very much for handling our manuscript. We sincerely appreciate the time and effort that both you and the reviewers have dedicated to evaluating our work.

We found the reviewers' comments constructive for improving the manuscript. In particular, Reviewer 1's observation that, in moderate to large-scale networks, the main computational bottleneck often lies in the cross-correlation step—rather than in preprocessing—was especially important. This prompted us to include an additional optimization (Workflow 3) that specifically addresses this issue by reusing spectral representations of each station across all pairs. As a result, the workflow achieves linear complexity in the number of stations, eliminating the quadratic scaling that characterizes traditional implementations. This new addition significantly improves scalability and strengthens the practical value of our approach.

We understand the reviewers' concern that the previously proposed optimizations may not be impactful in all use cases. We agree and do not intend to suggest otherwise. Most algorithmic improvements involve trade-offs—such as reduced flexibility or modularity. Rather than prescribing a replacement for existing practices, we offer simple, low-cost alternatives designed to improve computational efficiency and numerical robustness. We also emphasize that some of the proposed optimizations not only enhance performance but can also improve numerical stability. Due to their simplicity, these improvements are worth adopting even when the resulting computational gains are modest. In regions where high-performance computing (HPC) resources are readily available, it is often reasonable to prioritize ease of use, modularity, or broader applicability. However, in many parts of the world, limited access to HPC infrastructure makes such optimizations not only beneficial but sometimes essential to ensure the feasibility of large-scale ambient-noise tomography studies.

Following Reviewer 2's suggestion, we substantially shortened the manuscript by tightening the prose, condensing the background, and reducing the number of figures. Where appropriate, we combined maps and cross-correlation plots to improve clarity and avoid redundancy. Although the length of the main text could not be reduced—due to the addition of new references, figures, results, and a new workflow—the supplementary material has been completely eliminated, reducing the overall number of pages. As a result, the revised manuscript conveys substantially more useful information within fewer pages. In the original version, we inadvertently overemphasized the response-correction modification and underrepresented other optimizations; this balance has now been corrected. Fully explaining the rationale and implementation of all optimizations required additional space, particularly to ensure accessibility for researchers new to ambient-noise tomography or signal processing.

Finally, we revised the bibliography and data acknowledgments to ensure that all seismic networks used in our analysis are properly cited, including persistent identifiers (DOIs) for each dataset. We agree that proper citation is essential for recognizing and supporting the work of data providers.

We are grateful for the opportunity to revise our work and believe the updated manuscript has been significantly improved thanks to the reviewers' input. Please let us know if any further modifications are required.

With kind regards,

Caio Ciardelli
on behalf of all co-authors

Responses to Reviewers

Reviewer 1

A Phase-only Response-correction Method for Improved Efficiency and Stability in Building Empirical Green's Functions for Ambient-Noise Tomography An important first step in ambient noise tomography experiments that employ heterogeneous seismic station networks composed by a variety of seismic sensors is to remove the instrument response. The main idea of this manuscript is to remove the phase instead of the full response. This approach has sense considering the spectrum of the seismograms are typically whitened afterwards, when extracting the empirical Green's function (EGF). Besides, it avoids eventual numerical instability caused by the amplitude spectrum division operation of conventional instrumental response removal approaches. However, I have to disagree on the claimed improved computational efficiency of the proposed modified processing flow. Basically, because the instrument response correction, the spectral whitening and the amplitude normalization is computed per station, and the cross-correlation is computed per station pairs. Considering that modern ambient noise tomography experiments easily have more than 100 stations, the cross-correlation is by far the most time consuming operation, and when using the standard normalized cross-correlation reading and writing operations takes much more time than the calculations themselves.

We thank the reviewer for highlighting a key limitation of the original manuscript. We agree that for moderate to large seismic networks, the cross-correlation step typically dominates the computational cost, particularly when the number of station pairs far exceeds the number of stations. This important point prompted us to introduce a new optimization (Workflow 3) that keeps in memory and reuses the spectral representations of each station across all pairs, thereby reducing the complexity of the workflow from quadratic to linear with respect to the number of stations. This improvement directly targets the bottleneck identified by the reviewer and significantly enhances the scalability of the method. We also acknowledge that, in the initial version, we overemphasized the benefits of postponing instrument response correction, without sufficiently highlighting the broader set of optimizations included in the workflow. We have now clarified that this step is only one component of the overall method and is particularly useful because of its simplicity, ease of implementation in existing codes, and its positive impact on numerical stability. However, the most substantial performance gains come from the full combination of optimizations, especially when applied to large networks. These clarifications have been incorporated into the revised manuscript, along with new benchmarks that reflect the improved scalability of the method.

Pg. 4, lines 90–91 and from Pg. 5, line 115 to Pg. 6, line 133. In the vast majority of cases, the spectral whitening has to be applied before the amplitude normalization to improve convergence. This is clear when using one-bit amplitude normalization (the most widely used method and if I understood well the one you employ). Otherwise, the weakest frequency components of the seismograms would only contribute to the EGF at the zero crossing of the strongest ones.

We thank the reviewer for this insightful comment. We would like to clarify that we do not use one-bit normalization, but rather apply running absolute-mean normalization, which is less aggressive and better preserves amplitude modulation. An advantage of applying time-domain normalization first is that it more effectively suppresses large transients, yielding a more stable signal for spectral analysis and reducing bias from transient-dominated frequency bands. At the same time, as the reviewer pointed out, this can suppress weaker signals by limiting their contribution to the zero crossings of stronger components. But, since we use running absolute-mean normalization, this issue is largely mitigated in our workflow. We have now clarified this point in the revised manuscript and included a discussion of the advantages and disadvantages of both one-bit and running absolute-mean normalization, as well as the trade-offs associated with the ordering of normalization and spectral equalization.

Pg. 4, line 96 to Pg. 5, line 107. This is an extremely basic and well-known topic, I personally would not spend so much space on this. In case you still want, you should mention that the maximum lag time plays an important role in the numerical complexity of the cross-correlation operation because it allows you to break the convolution of two long sequences in pieces and compute many FFTs of short sequences instead of a single one of a long sequence, and that the FFT is optimal for sequences whose length is equal to a power of 2.

We thank the reviewer for this suggestion. We agree that the material on frequency-domain cross-correlation is standard and have condensed the corresponding section to focus only on the most relevant aspects. In doing so, we incorporated the reviewer's point regarding the role of maximum lag time in controlling the computational cost of the FFT. Specifically, we now mention that restricting the lag enables the use of shorter time series segments and that numerical efficiency is maximized when the sequence length is a power of two. We appreciate this clarification, which enhances both the precision and relevance of the text.

Pg. 5, line 113: "prefilter" is quite vague. Do you mean that we typically apply the deconvolution in the frequency band of interest? Aside, I would say that fixing a proper water level is even more important in practice.

Thank you for this comment. We have revised the paragraph to clarify that the prefilter is applied solely to stabilize the deconvolution by restricting it to frequency bands where the instrument has non-vanishing sensitivity, thereby removing the need to choose a water level. We do not use it to define the band of interest; instead, we apply band-pass filters only after obtaining the final broadband empirical Green's function.

Pg. 6, paragraph 2. When I compute one-bit amplitude normalization and the conventional normalized cross-correlation, the time spent reading the data (even from a fast SSD drive) is much higher than the time spent doing the calculations (essentially

3 FFT operations).

We completely agree that hardware performance and I/O operations can have a significant impact on runtime—sometimes exceeding the cost of floating-point operations themselves. In the original manuscript, we benchmarked the workflows across different hardware configurations and discussed how I/O and memory access can affect performance. In the revised version, we clarify these aspects further—especially in the context of Workflow 3—by benchmarking both endmembers: one with all preprocessing steps enabled, and another with preprocessing bypassed to isolate the cost of cross-correlation. While we did not bypass I/O explicitly, we note that, in our implementation and hardware setup, I/O time was negligible compared to other stages of the workflow.

Pg. 6, line 141. When you mention stacking, I guess you mean linear stacking. I agree that combining the correlation and the stacking operation does contribute to speed up calculations because you only have to save one correlation sequence instead of many. Having a modular processing flow is however convenient to simplify the application of more aggressive stacking methods such as weighted or non-linear stacking, or to discard anomalous correlations. This is in my opinion the reason to save the individual correlations in the first place.

We thank the reviewer for this observation. Indeed, we refer to linear stacking as the default method, but our implementation also supports weighted stacking, as noted in the revised manuscript. We agree that modular workflows offer advantages, particularly for applying non-linear stacking or performing outlier rejection, and we now explicitly acknowledge this trade-off. While our frequency-domain strategy improves performance by avoiding repeated transformations and I/O operations, we clarify that more advanced stacking methods can still be adapted to this approach with appropriate care. We hope this strikes a useful balance between efficiency and flexibility, and we thank the reviewer for prompting us to expand this discussion.

Pg. 7, lines 153–157. Avoiding spectral division is nice. In experiments where the sensors used are the same for all the stations, you could even avoid instrument response removal when you apply any sort of spectra whitening.

Thank you for this comment. Indeed, the main advantage of postponing response correction is avoiding spectral division, which improves numerical stability. While the associated performance gains are certainly welcome, our primary motivation for this adjustment is the simplification and increased robustness it brings to the workflow.

Pg. 8, line 180 to Pg. 9, line 195. When you compute the speed up from the standard to the modified workflow, it is important to take into account that some operations are applied per station and others per station pairs. As we have much less stations than station pairs, the operations that in practice matter in the overall execution time are the cross-correlation and the stacking. Moving the instrument response correction to the end, after the stacking where we have a single sequence, makes sense from the point of view of execution time.

We thank the reviewer for this observation. We fully agree that the distinction between operations applied per station and those applied per station pair is critical when assessing overall performance. In our revised manuscript, we clarify that the theoretical speed-up for Workflow 2 refers to a single station pair. As previously mentioned, the scaling limitations introduced by pairwise operations are addressed through an additional optimization introduced in Workflow 3, which shifts most computations to the station level and achieves linear complexity in the number of stations. We now *emphasize this improvement more clearly in the manuscript.*

Pg. 9, lines 196–199. It is nice you mention that execution time is not only about computational time. The time spent in reading from/writing to disk or RAM are often important.

We appreciate the reviewer's comment and fully agree. As previously mentioned, we further emphasize the impact of hardware on performance in the revised manuscript.

In the three examples you show, the pair of stations of the first two are from the same network. Do they use the same sensor? In this case, the phase response would be very close (if not the same), and removing the phase response difference at the end, as you propose in the modified workflow, would be equivalent to not removing any. I think it would be much more interesting to use stations with different sensors, to proof that moving instrument response removal to the end of the workflow is possible, even when applying aggressive non-linear operations such as the one-bit amplitude normalization. Aside, could you mention the sensors used in the three examples.

We thank the reviewer for this helpful observation. For the Southern California test, although both stations belong to the same network (CI), we explicitly included the instrument responses for MLAC and PHL in a figure to show that they are significantly different. For Brazil, while the twelve examples were already present in the supplementary material of the original submission, we have now grouped all results into a single figure and emphasized that some pairs include stations from different networks and with different sensors. For instance, PRPB (BR) uses a Trillium 120P, 120 s, 1201 V/m/s–Trident 305, 40 V, and RCBR (IU) is equipped with a Geotech KS-54000 Borehole Seismometer. In the Uganda test, the three stations span three distinct networks and sensors: MBAR (II) uses a Streckeisen STS-6, KMBO (IU) a Streckeisen STS-1VBB with E300 electronics, and MASD (Z5) a Streckeisen STS-2 Gen 3. Also, as previously mentioned, we chose running absolute-mean normalization rather than one-bit because it is less aggressive and better preserves amplitude modulation.

Minor comments

- **In the title: Bulding → Building**

  We thank the reviewer for catching the typo. It has been corrected in the revised manuscript.

- **Fig. 1. I think that drawing the two workflows separately would be clearer.**

  We agree and have redrawn the workflows as two separate diagrams to improve clarity. Workflow 3, which was introduced in the revised manuscript, is in a dedicated figure.

- **Pg. 9, equations between lines 188 and 189. The number of the equation and the numerical complexity symbol are missing.**

  We appreciate the reviewer's attention to detail. To maintain consistency, we opted to remove the complexity symbol from the preceding paragraph rather than add it to the equations, since we were referring to the approximate number of operations rather than formal asymptotic complexity. Equation numbers have now been added.

- **Pg. 10, line 212. Filling gaps with zeros may introduce spurious signals in the instrument response removal and the whitening operations due to border effects.**

  This is an important point. In our implementation, we first detrend and taper each continuous segment before zero-filling, which helps to minimize spurious artifacts caused by border effects.

We also note that a more robust strategy—such as skipping gaps entirely—could be adopted without altering the overall workflow. However, such implementations would be more complex, and our priority in this study was to introduce and benchmark the new algorithms in a clear and accessible manner.

Reviewer 2

Most importantly, I am not sure that such an easy modification requires such a long manuscript. For "Seismica" readers, who mostly have a seismology background, I really see no need for a lengthy review on noise correlation processing, etc. I am sure that the idea, as beautifully simple as it is, can be presented in just a few pages. (The summary I wrote above is really all there is to it.) Hence, my major suggestion is to radically shorten the manuscript.

We thank the reviewer for their constructive feedback. In response to this suggestion, we made substantial efforts to tighten the manuscript. We streamlined the prose, condensed the background sections, and removed or merged several figures to reduce redundancy. In particular, we combined maps and cross-correlation plots where appropriate to improve focus and clarity.

While we agree that the core concept of postponing phase response correction is simple and can be summarized succinctly, both the original and revised manuscripts present a broader set of improvements. However, we acknowledge that we inadvertently overemphasized this particular step in the initial version, which may have led to the impression that it was the sole or primary contribution of the work. This was our oversight, and we appreciate the reviewer's comment for highlighting it.

In reality, the response-correction step is only one of several workflow optimizations we propose. Other enhancements—such as restructuring the frequency-domain stacking procedure and introducing Workflow 3 to reduce the overall computational complexity from quadratic to linear in the number of stations—are less straightforward and required more careful explanation. These additions further expand the scope of the revised manuscript. Additionally, because one of our goals is to make the paper accessible to researchers with limited background in ambient-noise tomography or digital signal processing, we chose to retain a level of exposition that may be unnecessary for expert readers but helpful to new practitioners. We hope that the more concise and focused version now submitted strikes a better balance between clarity and completeness.

It really surprises me that the authors portray the computation of noise correlations as a high-precision science. In fact, there are a few open secrets in the community: (i) The subjective choices made in the preprocessing strongly affect the noise correlations and can lead to largely different results. Caring about instrument responses is, honestly, a second-order effect compared to all the extreme nonlinear procedures that people apply to noise recordings to obtain a correlation that "looks nice". (ii) Real noise sources fail all requirements needed to approximate a Green's function. It is well known that standard ambient noise tomographies do not produce reasonable seismograms when used for numerical wave propagation. (Recent work of Artie Rodgers for the Western US contains some prominent examples.) Hence, there does not seem to be a need to worry too much about the instrument response, when a few first-order issues are still open.

We appreciate this perspective and agree that instrument response correction is a secondary concern compared to more fundamental sources of uncertainty in ambient-noise tomography, such as prepro-

cessing choices and the distribution of noise sources. It was not our intention to suggest otherwise, and we have revised the manuscript to avoid giving that impression. However, response correction can become a more important issue when applied inconsistently; simplifying it removes subjectivity.

We do not present our modifications as essential corrections but as optional, low-cost enhancements aimed at improving robustness and numerical stability with minimal implementation effort. While their impact may vary across applications, they are particularly useful in resource-constrained environments or when integrating into larger processing pipelines. Postponing response correction, for instance, reduces numerical instability from spectral division and enables modest performance gains without sacrificing generality. As with any algorithmic adjustment, our approach involves trade-offs—such as reduced modularity—and is not meant to replace existing methods. We hope the revised manuscript better reflects this intent and thank the reviewer for the opportunity to clarify these points.

Response to Editor and Reviewers

Dear Dr. Koelemeijer,

Thank you very much for handling our manuscript for the second time and for giving us the opportunity to further improve it. We sincerely appreciate the considerable time and effort that you and the two reviewers have invested in evaluating and strengthening our work.

Once again, the reviewers' comments were insightful, constructive, and highly relevant. We fully understand the concerns regarding the perceived lack of novelty in some of the optimizations we discuss. Following Reviewer 1's excellent suggestion, we have revised both the title and the text throughout the manuscript to place greater emphasis on our main original contribution: the phase-only instrument-response correction applied after stacking. Following Reviewer 2's equally valuable recommendations, we now clearly state that several of the other optimizations are already known within the community, even if they are not very often explained in detail in the literature. As stated in the previous letter, our intention in presenting them in a simple, step-by-step manner—accompanied by systematic benchmarks and easy-to-follow Python code—is to make these techniques more widely accessible, especially to early-career researchers and those working without access to high-performance computing facilities. While we believe that making existing knowledge more accessible can, in certain contexts, have an impact comparable to developing entirely new methods, we completely understand that the current version may not fully meet the traditional criteria for a standard research article. We therefore gratefully accept your suggestion to publish the manuscript as a technical report instead.

Below we address each reviewer comment point-by-point, indicating the changes made or explaining why certain aspects were retained. We have also added an acknowledgment of the two anonymous reviewers in the Acknowledgments section.

We are deeply grateful for the careful and constructive review process. Thanks to the reviewers' input, we believe the revised manuscript is now significantly clearer and considerably more useful to the broader seismology community. Please do not hesitate to let us know if any further modifications are required.

With kind regards,

Caio Ciardelli
on behalf of all co-authors

Responses to Reviewers

Reviewer 1

Thanks a lot for the detailed answers to my comments and all the changes made throughout the manuscript. As I mentioned previously, I think that correcting solely for phase instrument response is the most interesting contribution of the manuscript, and I liked that this idea was reflected in the previous title of the article. However, the manuscript now has shifted its focus toward workflow optimizations, such as reducing the number of FFT operations. While these are very important in practice, they are widely known by practitioners, and I feel they somewhat distract from the core contribution. I fully agree that they are important in scenarios when computational resources are limited, and they are often overlooked when HPC clusters are available, in favor of modularity or non-linear correlation and stacking methods improving convergence speed, such as phase cross-correlation or phase weighted stacking. That said, moving the instrument response correction to the very end of the workflow is an original approach that may require reviewing some of these workflow optimizations. For this reason, I think it might be sufficient to go directly to workflow 3 (FW3) and compared it to the equivalent non-optimized workflow.

Thank you very much for this helpful comment. Following your suggestion, we have modified both the title and the abstract to better highlight the novelty and practical importance of the phase-only instrument-response correction applied after stacking. We have also added a clear sentence at the beginning of the paragraph preceding Equation 3 to explicitly emphasize the originality of this optimization.

We also greatly appreciate your suggestion to skip WF2 and move directly to WF3. After careful consideration, however, we decided to retain WF2 in the presentation because it serves an important pedagogical and analytical purpose: it isolates and clearly demonstrates the cumulative benefit of all per-station (i.e., $O(N)$) optimizations. In contrast, WF3 introduces a fundamentally different pairwise ($O(N^2)$) optimization that becomes increasingly dominant as the number of stations grows. For sufficiently large networks, this pairwise step overshadows the contributions of the earlier optimizations, making their individual impact difficult to discern. Conversely, when applied to a single station pair, WF3 actually performs slightly worse than WF2 (by approximately 8%, as shown in our benchmarks), which could obscure the clear advantages of the per-station optimizations. Since the only difference between WF2 and WF3 is this pairwise reuse of spectral representations, presenting WF2 first allows us to rigorously quantify the gains from all preceding optimizations in a clean and transparent way. Transitioning from WF2 to WF3 then becomes conceptually straightforward and instructive.

Minor comments

- **1. - Pg. 2, line 34: "constitution" is awkward. I would write "composition"**

  Thank you. We updated this sentence and replaced "constitution" by "composition". It sounds better, indeed.

- **2.- Pg. 8, line 177: n log2 n should be O(n log2 n) . Similarly in equations (2-6).**

  We thank the reviewer for the suggestion. However, in Equations (4)–(6) we are not providing an asymptotic complexity analysis but a concrete operation count to derive the expected numerical speed-up factor between the two workflows. The factor of 10 vs. 2 RFFT operations *per stack is the core of the theoretical 5.1× speed-up we claim, and big-O notation would hide this crucial constant as it only considers the fastest growing term without any multipliying constants. We therefore prefer to keep the current notation, but we have added a clarifying sentence on page 8: "Note that the expressions in Equations (4)–(6) represent concrete operation counts (not asymptotic big-O bounds) used to estimate the numerical speed-up factor." We hope this small addition eliminates any possible confusion.*

- **3.- Fig. 4, and similarly in Figs 6 and 7: Having the differences, maybe in a supplement, can be useful to better appreciate the impact of the changes introduced in WF2 on the waveforms.**

  Thank you very much for this excellent suggestion. Following your advice, we have included the residuals in the Supplementary Material, and they are indeed highly informative. In fact, examining them revealed that we had not properly adapted the retained period band for the newer instruments in Brazil and Uganda. The original band of 0.11–1000 s was well suited to the California stations but proved inappropriate for some receivers in the other regions. Although both WF2 and WF3 remain unconditionally stable regardless of instrument sensitivity, WF1 becomes unstable whenever the response approaches zero. In such cases, the comparison between WF1 and WF2 loses all meaning.

  We have now adjusted the retained period band individually for each region (Southern California, Brazil, and Uganda) to ensure meaningful and fair comparisons across all workflows. For completeness and transparency, we have also added much more cosine-similarity (CS) values in the main text. For each region we now report both the minimum and maximum CS values obtained across all station pairs, separately for (i) the full retained period band and (ii) the primary microseismic band (typically 10–30 s).

- **4.- Fig. 6: These subfigures are too small to see any detail. I would show a selection here, and move the rest to a supplement.**

  That is totally true. We kept only three representative pairs in the main text and send all the others to the Supplementary Material. Besides, we increased label sizes for all figures. Thank you for suggesting that.

Reviewer 2

This is a substantially revised version of the authors' original manuscript about strategies to accelerate and stabilise large-scale ambient noise correlations. Without losing content, the authors managed to produce a more precise manuscript that describes their contribution more clearly. The text is very well written and a pleasure to read.

Thank you very much for this encouraging comment. We are truly happy to know that our revisions have paid off.

Although the authors added new content in the form of an additional workflow that further reduces computational cost, its novelty is not obvious. Hence, the main issue from the original manuscript may persist: scientific relevance.

We understand this concern and agree that our current manuscript may not have enough inovations to be published as a research paper, despite the originality of our phase-only response removal method. That is why we agreed with the editor's suggestion to publish it as a technical report instead.

[1] The newly added workflow 3 certainly increases the scientific value of this manuscript. However, its novelty is somewhat unclear. The authors mention that they borrowed the concept from Clements & Denolle, 2020 ("Seisnoise.jl: Ambient seismic noise cross correlation on the CPU and GPU in Julia"), and there also seem to be elements from Fichtner et al., 2017 ('Seismic noise correlations on heterogeneous supercomputers"). The Whisper Suite (Briand et al., 2013 ("Processing of terabytes of data for seismic noise analysis with the Python codes of the Whisper Suite")) also applies various strategies for the reduction of computational requirements, although its details do not seem to be well documented in the literature. Still, the important point is that the authors should better explain if and how workflow 3 is novel.

Thank you for raising this concern. When we originally developed WF3, we did so indepentenly, without using any external references. However, when revising the text, we realized that Clements & Denolle (2020) briefly mention this optimization in their manuscript. That is why we cite it, despite having indepentenly reinvented the method. That, however, is the only source we could find where this optimization is clearly stated. Fichtner et al. (2017) does have elements of it, as you pointed out, but only claims to reuse the preprocessed data, not the FFT calculation istself. Briand et al. (2013) also does not mention doing that in their abstract.

To provide proper credit and improve clarity, we have added the following sentences in the Methodology section: "The strategy of computing FFTs once per station and reusing them for all pairwise correlations is fully implemented in the Julia package `SeisNoise.jl` (Clements & Denolle, 2020). Elements of this idea also appear in the `Mirmex` toolbox (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, down-sampling, 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 already explained in our response to Reviewer 1, we have also clarified in the abstract (and throughout the manuscript) that several of the optimizations we describe are already known in the ambient-noise community. Our primary original contribution remains the phase-only instrument-response correction applied after stacking—an optimization we could not find documented anywhere else in the literature.

[2] Almost all the figure labels are too small. (Figure labels should approximately have the same size as the main text.) Fig. 6 in particular is barely readable, also because the authors may have gone a bit too far with combining multiple figures into one. My suggestion would be to modify the figure(s) such that they really help to convey the authors' important messages. Less may be more in this case.

Thank you very much for highlighting this issue. As already explained in our response to Reviewer 1, we have significantly increased the font size in almost all figures throughout the manuscript. The only exception is the flowcharts, where the fixed cell dimensions inherently constrain the maximum font size.

We have also reduced the number of station-pair examples shown in the main text for the Brazilian network: only three representative pairs are now displayed in the main manuscript, while all remaining figures have been moved to the Supplementary Material. This greatly improves readability while preserving full access to the complete set of results.

We hope these changes fully address your concern.