A Scalable Algorithm for Cross-correlations of Compressed Ambient Seismic Noise

Eileen R. Martin, *Dept. of Mathematics, Computational Modeling and Data Analytics, Virginia Tech*

Last updated: April 14, 2019

Abstract

Companies and academic geophysicists are increasingly collecting continuous seismic data on denser arrays, and are looking to a variety of lossy compression methods to store and quickly access this data. Some researchers turn to ambient noise interferometry for low-cost near-surface imaging to avoid to cost and permitting required for active source experiments, but the computation can be very expensive. For each window of time, typical ambient noise interferometry scales as the product of the number of time samples per window and the number of sensors squared. This paper proposes a new algorithm for data stored in a low-rank matrix factorized form, performing interferometry in compressed form, and separating scalability in sensors from time samples. Application to real data shows nearly identical results at orders of magnitude lower cost. The algorithm can be extended to tensor compressions, averaging cross-correlations over many time windows.

1   Introduction and Prior Work

Consider an array of N_s continuously recording sensors. Each sensor collects N_t regularly spaced time samples per window of time. We wish to calculate the cross-correlation between each pair of sensors at N_τ time lags. The typical serial time-domain algorithm walks through each pair of sensors (that is, $N_s(N_s+1)/2$ pairs), calculates the inner product between the pairs' time-lagged data (each inner product takes $O(N_t - N_\tau)$ calculations in the time domain) for each time lag (we consider N_τ time lags). This scales as $O(N_s^2(N_t - N_\tau)N_\tau)$. There have been efforts to speed up these calculations through parallelism and GPUs in Mirmex, which breaks up the number of sensors into small groups or tiles (Fichtner et al., 2017).

As technologies such as MEMS accelerometers (Evans et al., 2014) and Distributed Acoustic Sensing (DAS) (Martin et al., 2018) make long-term continuous recording easier than ever, the number of sensors has grown by orders of magnitude (particularly for DAS, which enables meter-scale channel spacing over multiple kilometers of fiber). The number of windows of data averaged in ambient noise interferometry is typically at least hundreds, although modern scalable algorithms must be able to work with streaming data. For ambient noise interferometry on modern seismic acquisition systems, scales of interest are roughly: Number of sensors, $N_s \approx 100$s to 10,000s; Number of time samples per window or data, $N_t \approx 10,000$s to 1,000,000s; Number of time lags of interest in the cross-correlation, $N_\tau \approx 100$s to 1,000s.

Further, with new acquisition technologies becoming increasingly prevalent, energy companies and academic seismic data repositories are increasingly considering lossy compression methods to increase their access speed for large quantities of passive seismic data. Particularly when compressing passive seismic data, collected in large quantities with a variety of uses, a general purpose compression such as a low-rank matrix factorization of the data is appealing. Say that the true data in one time window is represented by the matrix $D \in \mathbb{R}^{N_s \times N_t}$, and is approximately equal to $D \approx UV^T$ where U has N_s rows, and V has N_t rows, and U and V are fast-to-multiply (can be expressed in terms of a small number of low-rank or sparse matrices). Some examples of fast methods to obtain low-rank factorizations include randomized SVD (combining the singular value matrix into either the left or right singular values) and randomized QR (Halko et al., 2009), (Xiao et al., 2017).

This paper proposes a new algorithm that performs ambient noise cross-correlations on data compressed through a low-rank matrix factorization such that the data never need to be decompressed. By leaving the data in their low-rank factorized form, the algorithm splits the scalability between number of passive time samples and number of sensors $O(k^2(N_t - N_\tau)N_\tau) + O(N_s^2 k)$ where $k << N_s, N_t$ and $N_\tau << N_t$. Further, the algorithm reduces to level-3 BLAS operations, rather than sliding inner products, so it can easily take advantage of highly optimized parallel linear algebra libraries in a way that prior ambient noise interferometry algorithms could not. I show the algorithm applied to DAS data acquired at the Stanford Fiber Optic Seismic Observatory, and that the results of the new algorithm are identical (up to expected numerical precision) to those of cross-correlations of reconstructed data following the traditional algorithm. However, the results with the new algorithm are achieved more than 100 times faster per time window compared to the old algorithm on this particular dataset.

2 Theory

Consider a low-rank matrix approximation of the data in one time window $D \approx UV^T$, where $U \in \mathbb{R}^{N_s \times k}$, $V \in \mathbb{R}^{N_t \times k}$, and $k << N_s, N_t$. Denote the data subset at receiver vs:

$$
\begin{aligned}
\tilde{D}^{(vs,0)} &= D_{vs,N_\tau/2:N_t-N_\tau/2-1} &(1)\\
&\approx U_{vs,0:k-1}V^T_{0:k-1,N_\tau/2:N_t-N_\tau/2-1} &(2)\\
&= \tilde{U}^{(vs)}\tilde{V}^{(0)T} &(3)
\end{aligned}
$$

and the time-lagged data at receiver r

$$
\begin{aligned}
\tilde{D}^{(r,i_\tau)} &= D_{r,i_\tau+N_\tau/2:i_\tau+N_t-N_\tau/2-1} &(4)\\
&\approx U_{r,0:k-1}V^T_{0:k-1,i_\tau+N_\tau/2:i_\tau+N_t-N_\tau/2-1} &(5)\\
&= \tilde{U}^{(r)}\tilde{V}^{(i_\tau)T} &(6)
\end{aligned}
$$

then the cross-correlation between receivers vs and r at time lag index i_τ

$$
\begin{aligned}
C_{vs,r,i_\tau+N_\tau/2} &= \tilde{D}^{(vs,0)} \cdot \tilde{D}^{(r,i_\tau)} &(7)\\
&\approx \tilde{U}^{(vs)}\tilde{V}^{(0)T}\left(\tilde{U}^{(r)}\tilde{V}^{(i_\tau)T}\right) &(8)\\
&= \tilde{U}^{(vs)}\left(\tilde{V}^{(0)T}\tilde{V}^{(i_\tau)}\right)\tilde{U}^{(r)T} &(9)
\end{aligned}
$$

Note that the center factor in parenthesis is the same no matter which sensor pair is being cross-correlated. Denote

$$W^{(i_\tau)} = \tilde{V}^{(0)T}\tilde{V}^{(i_\tau)} \tag{10}$$

which is a $k \times k$ matrix that costs $O(k^2(N_t - N_\tau))$ to calculate. Notice that $W^{(i_\tau)}$ is shared for all sensor pairs (vs, r). This process can be carried out for all sensor pairs organized in a larger matrix multiplication for any given time-lag:

$$C_{0:N_s-1,0:N_2-1,i_\tau+N_\tau/2} \approx \tilde{U}^{(vs)}W^{(i_\tau)}\tilde{U}^{(r)T} \tag{11}$$

Thus, we could calculate all $W^{(i_\tau)}$, then perform a multiplication of a tall-skinny matrix in $\mathbb{R}^{N_s \times k}$, by a small matrix in $\mathbb{R}^{k \times k}$, by a short-fat matrix in $\mathbb{R}^{k \times N_s}$. These multiplications can take advantage of pre-optimized routines for dense matrix-matrix multiplications. The complexity based on number of sensors can thus be split apart from the complexity based on the number of time samples in the window of data by this rearrangement of the cross-correlation.

3 Proposed algorithm

The theory above suggests the following algorithm:

Algorithm 1 Scalable cross-correlations of matrix-factorized data $D \approx UV^T$
---
procedure SCALABLE_XCORRS($U \in \mathbb{R}^{N_s \times k}, V \in \mathbb{R}^{N_t \times k}$)
    Let $C \in \mathbb{R}^{N_s \times N_s \times N_\tau}$
    Let $\tilde{V}^T = V^T_{0:k-1,N_\tau/2:N_t-N_\tau/2-1}$
    for $i_\tau = -N_\tau/2, \ldots, N_\tau/2$ **do**
        Let $\tilde{V}^{(0)} = V_{N_\tau/2:N_t-N_\tau/2-1,0:k-1}$
        Let $\tilde{V}^{(i_\tau)} = V_{i_\tau+N_\tau/2:i_\tau+N_t-N_\tau/2-1,0:k-1}$
        $W^{(i_\tau)} = \tilde{V}^{(0)T}\tilde{V}^{(i_\tau)}$
        $\tilde{C} = U_{0:N_s-1,0:k-1}W^{(i_\tau)}U^T_{0:k-1,0:N_s-1}$
        $C_{0:N_s-1,0:N_s-1,i_\tau+N_\tau/2} = \tilde{C}$
    Return C

This algorithm loops over all time lags, and for each time lag, it performs $O(k^2(N_t - N_\tau))$ calculations to obtain the small matrix W, then multiplies a tall-skinny by a small square by a short fat matrix, adding another $O(N_s k^2) + O(N_s^2 k)$ calculations. We assumed that $k << N_s$. Thus, the total operation complexity is $O(k^2(N_t - N_\tau)N_\tau) + O(N_s^2 k)$ for each time window. The scalability in time and number of sensors (both large quantities) is completely separated.

4 Application to Passive DAS Data

Since September 2016, the Stanford Fiber Optic Seismic Observatory has been continuously collecting passive seismic data at a rate of 50 samples per second at 4 meter channel spacing on 2.5 km of fiber optic cables in existing telecommunications conduits underneath the Stanford

University campus (Martin et al., 2018). It is the longest running, ultra-dense urban seismic experiment, and has been used to ambient noise interferometry, earthquake detection (Lindsey et al., 2017), and active seismic acquisition (Martin et al., 2017).


Figure 1: The number of singular values at least 5% the largest is plotted for each 5 minute file (620 x 15000 samples) over two days. The rank increases sharply as morning traffic begins.

An array with similar settings that covered 1000 km (a reasonable length for an array covering a major metropolitan area in an earthquake-prone region) would collect 0.5 petabytes of data annually, so some form of lossy compression will need to be part of any scalable plan for future DAS data archiving.


Figure 2: The the compressed and reconstructed data (bottom) during a small earthquake shows the same moveouts as the true data (top). The vertical axis is sensor number, and the horizontal axis is time sample with 0.02 seconds between samples.

The data are broken into 5 minute files with all 620 sensors recording data at 0.02 second sampling (15,000 time samples). As an initial test of the compressibility of these data, I calculated the singular value decomposition (SVD) of each file, and observed relatively fast decay of singular values. For each file in order, the number of singular values at least 5% the size of the largest

is plotted in Figure 1 as a proxy for that file's compressibility. There is a clear jump up in the number of larger singular values starting around 5 am local time (when commuting starts). This makes sense, as it would indicate a wider variability in the vibrations being recorded, caused by a wider variability in the sources of ambient noise (both anthropogenic and natural during the day, but mostly natural late at night).

One may rightly wonder whether the truncated SVD is an appropriate data compression scheme, and it is likely that it is not the optimal one, but for ambient noise interferometry, the important question is whether moveouts between sensors over a wide variety of random sources are preserved. A plot of data during an earthquake that was compressed by truncating the SVD at 1/20th the maximum singular value for the file starting at UTC 2016-12-12 02:22:21.330.

To compare the speed of the typical cross-correlation method versus the new algorithm, I calculated the truncated SVD of the first five minute file from Figure 1, which starts at 2016-12-12 00:02:21.330, and has 38 singular values at least 5% the largest singular value. I applied both the new and the traditional algorithms to this five-minute window of data with $N_\tau = 100$ (one second of time lag in either direction, since the array has a small aperture).

I implemented both the old and new algorithms in python using matrix operations in the numpy package. For one 5-minute window on a laptop with a 3.1 GHz Intel Core i5 processor, the typical pair-wise cross-correlation algorithm took 175 seconds to calculate all cross-correlations at all time lags, with < 1 second initially decompressing data. The new algorithm took a total of 1.33 seconds to calculate the same cross-correlations, which broke down into 0.63 seconds calculating W matrices plus 0.70 seconds calculating UWU^T products. The Frobenius norm of the output tensor C in both cases was 0.34134182, and the Frobenius norm of the difference in the typical and new algorithms' C tensors was 3.73×10^{-8}. Neither the old algorithm nor the new algorithm had any parallelism or additional optimizations (for example, to take advantage of time-reversed symmetry between receiver pairs), and the for-loops over sensors are a major inefficiency, particularly in python. The two algorithms yielded the same results up to expected floating point precision, but the new algorithm performed the same task more than 100 times faster than the old algorithm.

In a complex noise environment such as the Stanford campus, being able to calculate cross-correlations faster means that computational resources can be freed up to test a wider variety of pre-processing schemes, or to test multiple weighting strategies between windows of data. These effect of types of changes on cross-correlations can be difficult to predict, especially in the presence of localized or transient noise sources often present in urban areas or around infrastructure, so it is important that geophysicists test a variety of processing schemes to ensure robust results. Further, many field sites have limited compute power or bandwidth to send data into centralized high-performance computing facilities. This algorithm would allow scientists to analyze denser/larger seismic arrays in the field.

5   Extensions: Other Factorizations, Tensors

Other algorithms focused on hierarchical relationships due to spatio-temporal locality, or factorizations of data transforms such as singular spectrum analysis (Sacchi, 2009) may yield more reliable compression of seismic data. Such a factorization might allow researchers to better

preserve the characteristic moveouts of any given window of data, thus enabling extraction of reliable virtual source response functions from cross-correlations of fewer windows of data. This algorithm should be adapted to consider those factorizations, which should still allow separation of the scalability in number of sensors and number of time samples, although the scalability of the matrix multiplications may vary depending on the sparse/low-rank structure.

Additionally, if N_w windows of data are represented by the tensor $T \in \mathbb{R}^{N_s \times N_t \times N_w}$, and an approximate low-rank tensor decomposition is stored, one can adapt the matrix factorization presented here to calculate all pairs of approximate cross-correlations and average them over N_w windows while separating out the scalability with N_w. Given that N_w is often 100s to 10,000s or more, there is significant savings in splitting its scalability from N_s and N_t. Consider an approximate CP tensor decomposition where for each entry: $T_{s,t,w} \approx \sum_{\kappa=1}^{k} \lambda_k a_s^{(\kappa)} b_t^{(\kappa)} c_w^{(\kappa)}$ (Kolda and Bader, 2009). Consider the average cross-correlation of receivers vs and r at time lag index i_τ:

$$\sum_{w=0}^{N_w-1} C_{vs,r,i_\tau+N_\tau/2,w} \quad = \quad \sum_{w=0}^{N_w-1} \sum_{j=0}^{N_t-N_\tau-1} T_{vs,N_\tau/2+j,w} T_{r,i_\tau+N_\tau/2+j,w} \tag{12}$$

$$= \quad \sum_{w=0}^{N_w-1} \sum_{j=0}^{N_t-N_\tau-1} \sum_{\kappa,\gamma=1}^{k} \lambda_\kappa a_{vs}^{(\kappa)} b_{J-i_\tau}^{(\kappa)} c_w^{(\kappa)} \lambda_\gamma a_r^{(\gamma)} b_J^{(\gamma)} c_w^{(\gamma)} \tag{13}$$

where $J := i_\tau + N_\tau/2 + j$ for more concise notation. Denote $A \in \mathbb{R}^{N_s \times N_s \times k \times k}$ have its (vs, r, κ, γ) entry as $\lambda_\kappa \lambda_\gamma a_{vs}^{(\kappa)} a_r^{(\gamma)}$, so it takes $O(k^2 N_s^2)$ operations to fill. Let $B \in \mathbb{R}^{N_\tau \times k \times k}$ have its (i_τ, κ, γ) entry as $\sum_{j=0}^{N_t-N_\tau-1} b_{\frac{N_\tau}{2}+j}^{(\kappa)} b_{i_\tau+\frac{N_\tau}{2}+j}^{(\gamma)}$, so it takes $O(k^2(N_t-N_\tau)N_\tau)$ operations to calculate. Let $\mathcal{C} \in \mathbb{R}^{k \times k}$ have its (κ, γ) entry as $\sum_{w=0}^{N_w-1} c_w^{(\kappa)} c_w^{(\gamma)}$, so it takes $O(N_w k^2)$ operations to fill. Then we can rearrange the calculations above so that:

$$\sum_{w=0}^{N_w-1} C_{vs,r,i_\tau+N_\tau/2,w} = \sum_{\kappa,\gamma=1}^{k} A_{vs,r,\kappa,\gamma} B_{i_\tau+N_\tau/2,\kappa,\gamma} \mathcal{C}_{\kappa,\gamma} \tag{14}$$

By calculating A, B, and $\mathcal{C}$, then multiplying their (κ, γ) entries for all pairs of (vs, r) and i_τ this algorithm adaptation requires just $O(k^2(N_t - N_\tau)N_\tau) + O(k^2 N_w) + O(k^2 N_s^2 N_\tau)$ operations to calculate the cross-correlation between all sensor pairs averaged over N_w windows. This separates the scalability in sensors from number of windows and from time samples per window.

6 Conclusions

In this paper I show that if ambient seismic noise data has been stored in an approximate low-rank matrix-factorized form, it should be cross-correlated in its compressed form. I propose a new algorithm that separated the scalability in number of sensors from scalability in number of time samples. The algorithm applied to a real DAS data set shows that the same results are achieved at much lower computational cost. This algorithm can be generalized to averaging cross-correlations over many time windows, separating the scalability in windows from sensors and time samples per window. Further work is needed to parallelize these algorithms, to efficiently incorporate preprocessing workflows, and to adapt these algorithms to other data compression schemes.

7   Acknowledgements

References

[1] Fichtner, A., L. Ermert, and A. Gokhberg. *Seismic noise correlation on heterogeneous supercomputers.* Seismological Research Letters 88, no. 4 (2017): 1141-1145.

[2] Martin, E.R., F. Huot, Y. Ma, R. Cieplicki, S. Cole, M. Karrenbach, and B.L. Biondi. *A seismic shift in scalable acquisition demands new processing: Fiber-optic seismic signal retrieval in urban areas with unsupervised learning for coherent noise removal.* IEEE Signal Processing Magazine 35, no. 2 (2018): 31-40.

[3] Evans, J. R., R.M. Allen, A. I. Chung, E. S. Cochran, R. Guy, M. Hellweg, and J. F. Lawrence. *Performance of several lowcost accelerometers.* Seismological Research Letters 85, no. 1 (2014): 147-158.

[4] Halko, N., P.-G. Martinsson, and J.A. Tropp. *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions.* SIAM review 53, no. 2 (2011): 217-288.

[5] Xiao, J., M. Gu, and J. Langou. *Fast parallel randomized QR with column pivoting algorithms for reliable low-rank matrix approximations.* In 2017 IEEE 24th International Conference on High Performance Computing (HiPC), pp. 233-242. IEEE, 2017.

[6] Lindsey, N.J., E.R. Martin, D.S. Dreger, B. Freifeld, S. Cole, S.R. James, B.L. Biondi, and J.B. AjoFranklin. *Fiberoptic network observations of earthquake wavefields.* Geophysical Research Letters 44, no. 23 (2017): 11-792.

[7] Martin, E.R., C.M. Castillo, S. Cole, P.S. Sawasdee, S. Yuan, R.G. Clapp, M. Karrenbach, and B.L. Biondi. *Seismic monitoring leveraging existing telecom infrastructure at the SDASA: Active, passive, and ambient-noise analysis.* The Leading Edge 36, no. 12 (2017): 1025-1031.

[8] Sacchi, M.D. *FX singular spectrum analysis.* In CSPG CSEG CWLS Convention, pp. 392-395. 2009.

[9] Kolda, T.G., and B.W. Bader. *Tensor decompositions and applications.* SIAM review 51, no. 3 (2009): 455-500.