Processing seismic ambient noise data to obtain reliable broad-band surface wave dispersion measurements

A P Bensen, M M Ritzwoller, M M Barmin, A P Levshin, F P Lin, M M Moschetti, N M Shapiro, Yanyan Yang

Geophys. J. Int. (2007) **169,** 1239–1260 doi: 10.1111/j.1365-246X.2007.03374.x

Processing seismic ambient noise data to obtain reliable broad-band surface wave dispersion measurements

G. D. Bensen,[1] M. H. Ritzwoller,[1] M. P. Barmin,[1] A. L. Levshin,[1] F. Lin,[1] M. P. Moschetti,[1] N. M. Shapiro[2] and Y. Yang[1]

[1] *Center for Imaging the Earth's Interior, Department of Physics, University of Colorado at Boulder, Campus Box* 390, *Boulder,* CO 80309, *USA. E-mail: gbensen@ciei.colorado.edu*

[2] *Laboratoire de Sismologie, CNRS, IPGP,* 4 *place Jussieu,* 75252 *Paris Cedex* 05, *France*

Accepted 2007 January 31. Received 2007 January 24; in original form 2006 July 28

SUMMARY

Ambient noise tomography is a rapidly emerging field of seismological research. This paper presents the current status of ambient noise data processing as it has developed over the past several years and is intended to explain and justify this development through salient examples. The ambient noise data processing procedure divides into four principal phases: (1) single station data preparation, (2) cross-correlation and temporal stacking, (3) measurement of dispersion curves (performed with frequency–time analysis for both group and phase speeds) and (4) quality control, including error analysis and selection of the acceptable measurements. The procedures that are described herein have been designed not only to deliver reliable measurements, but to be flexible, applicable to a wide variety of observational settings, as well as being fully automated. For an automated data processing procedure, data quality control measures are particularly important to identify and reject bad measurements and compute quality assurance statistics for the accepted measurements. The principal metric on which to base a judgment of quality is stability, the robustness of the measurement to perturbations in the conditions under which it is obtained. Temporal repeatability, in particular, is a significant indicator of reliability and is elevated to a high position in our assessment, as we equate seasonal repeatability with measurement uncertainty. Proxy curves relating observed signal-to-noise ratios to average measurement uncertainties show promise to provide useful expected measurement error estimates in the absence of the long time-series needed for temporal subsetting.

Key words: seismic noise, seismic processing, surface waves, tomography.

1 INTRODUCTION

Theoretical studies have shown that the cross-correlation of diffuse wavefields (e.g. ambient noise, scattered coda waves) can provide an estimate of the Green function between the stations (e.g. Weaver & Lobkis 2001a,b, 2004; Derode *et al.* 2003; Snieder 2004; Wapenaar 2004; Larose *et al.* 2005). Seismic observations based on cross-correlations between pairs of stations have confirmed the theory for surface waves using both coda waves (Campillo & Paul 2003; Paul *et al.* 2005) and long ambient noise sequences (Shapiro & Campillo 2004; Sabra *et al.* 2005a) and for crustal body waves using ambient noise (Roux *et al.* 2005). Oceanic applications are also feasible (Lin *et al.* 2006). An example of a year-long cross-correlation between a pair of Global Seismic Network (GSN) stations in the US filtered into several subbands is shown in Fig. 1.

The first attempts to use ambient noise for surface wave tomography, called ambient noise surface wave tomography, were applied to stations in Southern California (Shapiro *et al.* 2005; Sabra *et al.* 2005b). These studies resulted in group speed maps at short periods (7.5–15 s) that displayed a striking correlation with the principal geological units in California with low-speed anomalies corresponding to the major sedimentary basins and high-speed anomalies corresponding to the igneous cores of the main mountain ranges.

Ambient noise tomography is now expanding rapidly. Recent applications have arisen across all of California and the Pacific Northwest (Moschetti *et al.* 2007), in South Korea (Cho *et al.* 2006), in Tibet (Yao *et al.* 2006), in Europe (Yang *et al.* 2007), across New Zealand (Lin *et al.* 2007), as well as elsewhere in the world. Most of the studies, to date, like the earlier work of Shapiro *et al.* (2005), have been performed in the microseism band below 20 s period. Broad-band applications extending to considerably longer periods are now emerging (e.g. Bensen *et al.* 2005; Yao *et al.* 2006; Yang *et al.* 2007) and the method is also being applied to increasingly large areas such as Europe (Yang *et al.* 2007). In spite of these developments, the data processing procedures that underlie ambient noise tomography remain poorly documented, even as they have become increasingly refined. The purpose of this paper is to summarize the state of data processing as it has developed since the first

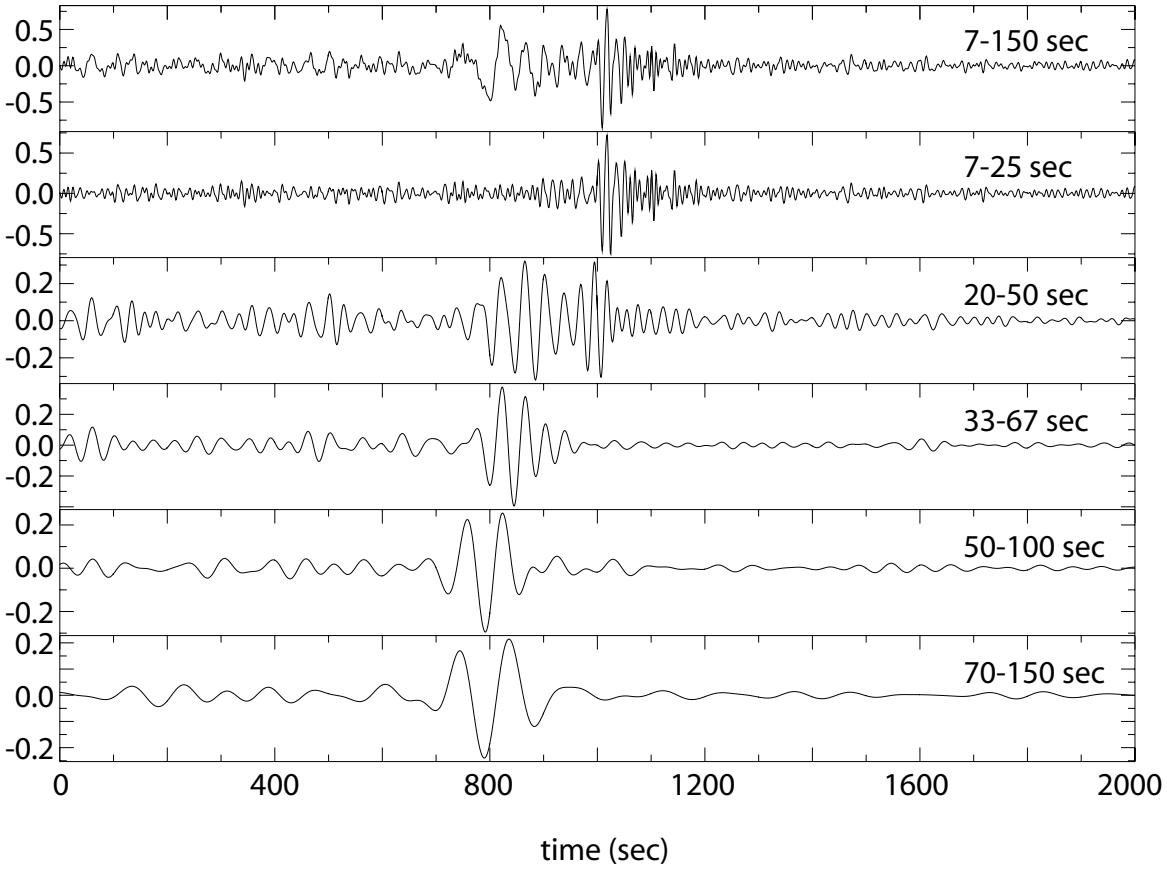

Figure 1. Example of a broad-band symmetric-component cross-correlation using 12-months of data from stations ANMO (Albuquerque, NM, USA) and HRV (Harvard, MA, USA). The broad-band signal (7–150 s passband) is shown at top and successively longer period passbands are presented lower in the figure. (The symmetric component is the average of the cross-correlation at positive and negative lags.)

papers on the use of ambient noise to obtain surface wave dispersion measurements (Shapiro & Campillo 2004).

In its current state, the ambient noise data processing procedure divides into four principal phases that are applied roughly in order: (1) single station data preparation, (2) cross-correlation and temporal stacking, (3) measurement of dispersion curves and (4) quality control, including error analysis and selection of the acceptable measurements. These steps are presented schematically in Fig. 2. After data processing is complete, tomography for group or phase speed maps (e.g. Yang *et al.* 2007) and inversion for a V_s model (e.g. Cho *et al.* 2006; Lin *et al.* 2007) may follow, but discussion of these steps is beyond the scope of the present paper. The procedures in this paper are exclusively applied to Rayleigh waves, but Love wave studies have also begun to emerge (e.g. Cho *et al.* 2006).

In judging between candidate components of the data processing procedure, we have assigned significant weight to flexibility and the applicability to a wide variety of observational situations. The procedures described here, therefore, are designed to be applied over a broad range of periods, interstation distances and geographical scales. Examples are shown in this paper from regional to continental scales, from very short to long periods, and are drawn from Europe, North America and New Zealand. Applications are, however, taken exclusively from continental or ocean island stations. Most are, in fact, taken from GSN stations within the US. As discussed by Lin *et al.* (2006), broad-band cross-correlations of ambient noise obtained at ocean bottom or subbottom seismometers (OBS) are contaminated at long periods (above ∼25 s) by tilting under fluid flow and seafloor deformation under gravity waves. Crawford *et al.* (2006) argue that these effects can be mitigated on the vertical component using horizontal component data and a co-located differential seafloor pressure gauge. The success of this process will be needed for broad-band ambient noise measurements to be obtained from OBS data. We are unaware of research that has tested this idea in the context of ambient noise measurements, however.

Our principal purpose, therefore, is to summarize the status of the ambient noise data processing procedure that we have developed over the past several years. The paper is intended to explain, justify, and present salient examples of this development. It is also intended to act as a primer to help provide guidance and act as a basis for future efforts in surface wave studies based on ambient seismic noise. Each of the four following sections presents a discussion of one phase of the data processing procedure, which ranges from processing data from a single station (Section 2), cross-correlating and stacking data from station-pairs (Section 3), measuring surface wave dispersion (Section 4) and applying data quality control measures, particularly estimating uncertainties and selecting reliable measurements (Section 5).

2 SINGLE STATION DATA PREPARATION

The first phase of data processing consists of preparing waveform data from each station individually. The purpose of this phase is to accentuate broad-band ambient noise by attempting to remove

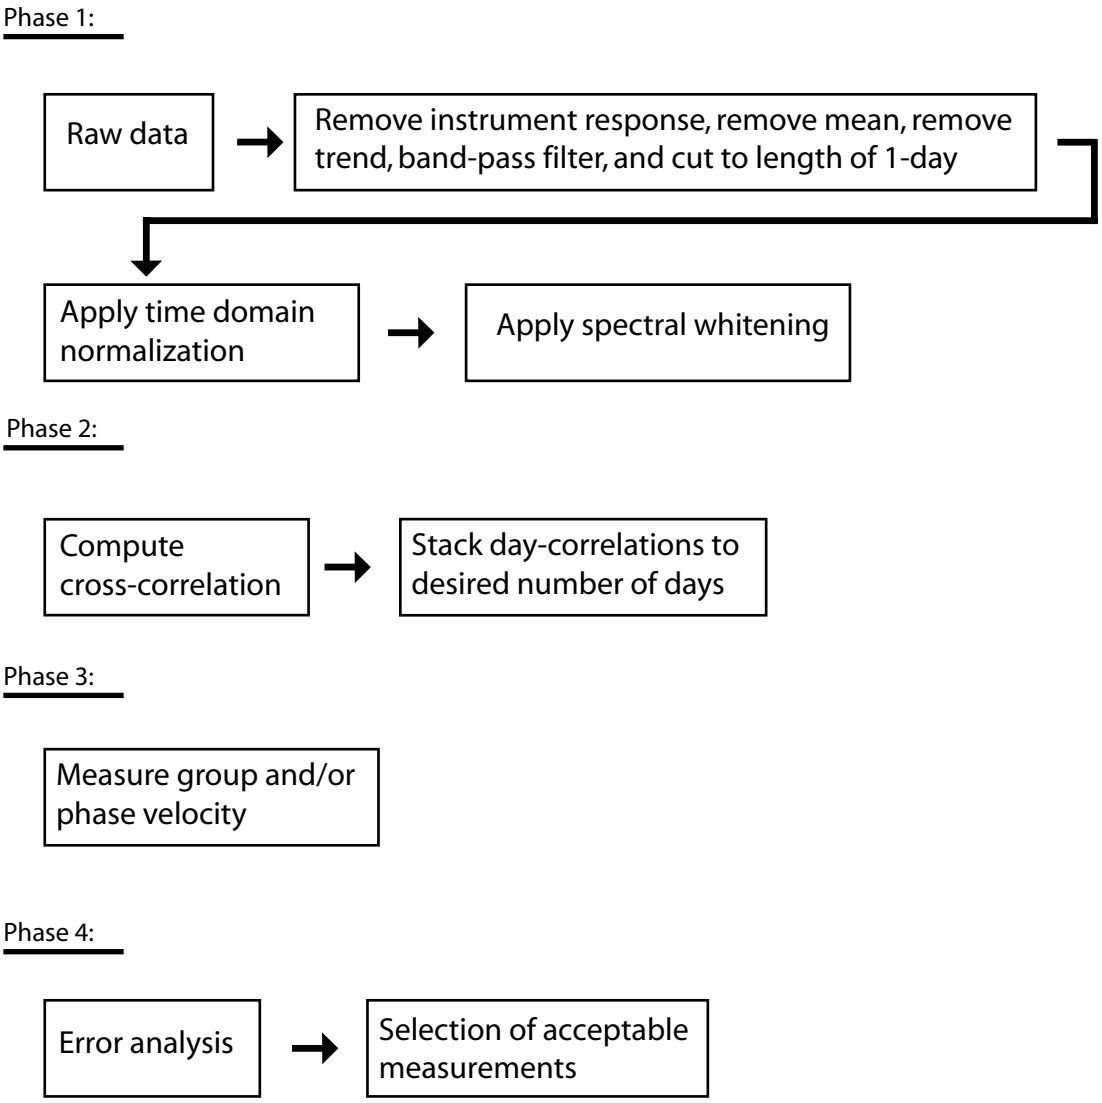

Figure 2. Schematic representation of the data processing scheme. Phase 1 (described in Section 2 of the paper) shows the steps involved in preparing single-station data prior to cross-correlation. Phase 2 (Section 3) outlines the cross-correlation procedure and stacking, Phase 3 (Section 4) includes dispersion measurement and Phase 4 (Section 5) is the error analysis and data selection process.

earthquake signals and instrumental irregularities that tend to obscure ambient noise. Obscuration by earthquakes is most severe above about 15 s period, so this step of the data processing is most important at periods longer than the microseism band (∼5 to ∼17 s period). In addition, because the spectral amplitude of ambient noise peaks in the microseism band, methods have to be devised to extract the longer period ambient noise from seismic records. Fig. 2 shows the steps that compose Phase 1 of data processing: removal of the instrument response, de-meaning, de-trending and bandpass filtering the seismogram, time-domain normalization and spectral whitening. This procedure is typically applied to a single day of data. Day data with less than 80 per cent 'on-time' are currently rejected, but this may be modified at the user's discretion. Some of the steps, such as the temporal normalization and spectral whitening, impose non-linear modifications to the waveforms, so the order of operations is significant. Because this phase of data processing is applied to single stations, rather than to station-pairs, it is much less time consuming and computationally intensive than subsequent cross-correlation, stacking and measurement phases that are discussed in later sections of the paper. Our current applications involve from several dozen (e.g. 41 stations across New Zealand) to several hundred (e.g. 110 stations across Europe, ∼250 stations across North America) stations.

2.1 Temporal normalization

The most important step in single-station data preparation is what we call 'time-domain' or 'temporal normalization'. Time-domain normalization is a procedure for reducing the effect on the cross-correlations of earthquakes, instrumental irregularities and non-stationary noise sources near to stations. Earthquakes are among the most significant impediments to automated data processing. They occur irregularly and, although the approximate times and locations of large earthquakes can be found in earthquake catalogs, small earthquakes over much of the globe are missing from global catalogs. In addition, the time of arrival of surface wave phases at short periods is not well known. Thus, removal of earthquake signals must be data-adaptive, rather than prescribed from a catalogue.

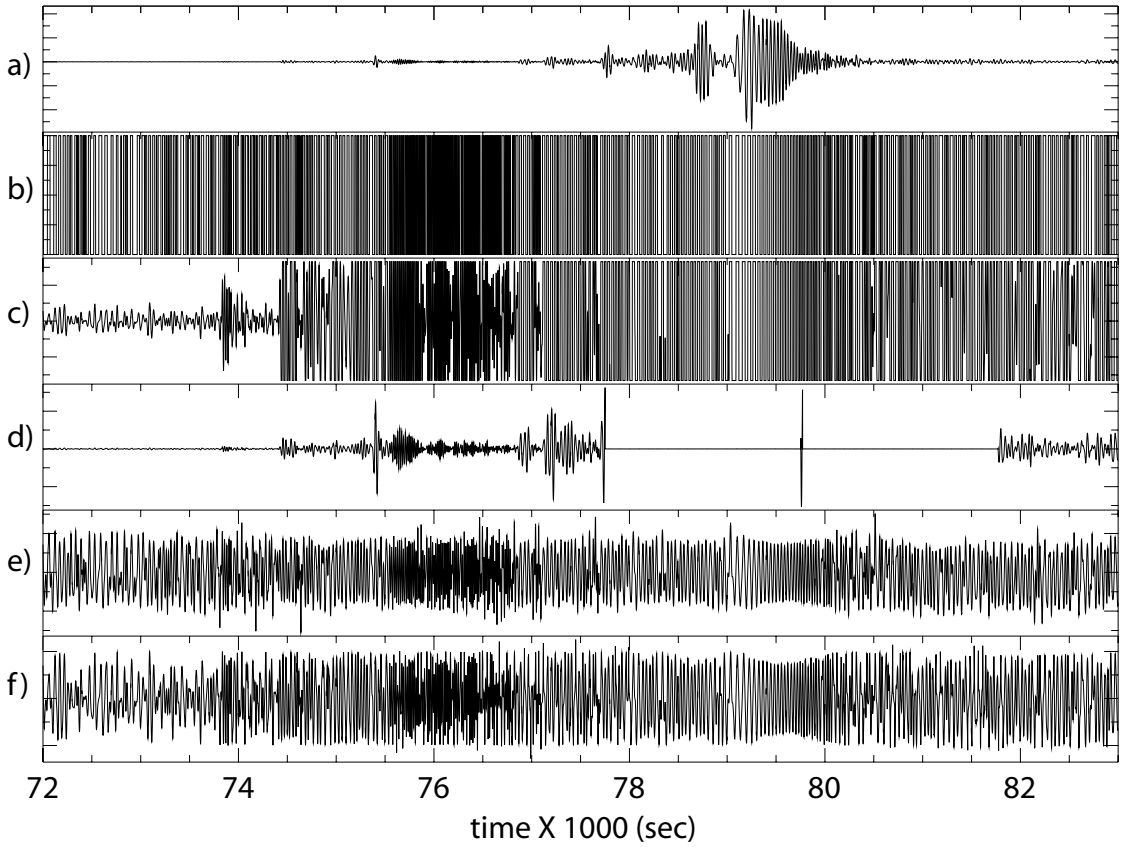

Figure 3. Waveforms displaying examples of the five types of time-domain normalization tested. The examples are bandpass filtered between 20 and 100 s period to clarify the contamination by the earthquake signal. (a) Raw data showing ∼3 h of data windowed around a large earthquake ($M_s = 7.2$, Afghanistan-Tajikistan border region) recorded at station ANMO. (b) One-bit normalized waveform, whereby the signal is set to ±1 depending on the sign of the original waveform. (c) Clipped waveform, where the clipping threshold is equal to the rms amplitude of the signal for the given day. (d) Automated event detection and removal. If the amplitude of the waveform is above a certain threshold, the next 30 min of it are set to zero. (e) Running absolute mean normalization whereby the waveform is normalized by a running average of its absolute value. (f) 'Water level normalization' whereby any amplitude above a certain multiple of the daily rms amplitude is down-weighted. It is run iteratively until the entire waveform is nearly homogeneous in amplitude.

We have considered five different methods to identify and remove earthquakes and other contaminants automatically from seismic waveform data. An illustrative example is shown in Fig. 3. The first and most aggressive method is called 'one-bit' normalization (Fig. 3b), which retains only the sign of the raw signal by replacing all positive amplitudes with a 1 and all negative amplitudes with a −1. This method has been shown to increase signal-to-noise ratio (SNR) when employed in acoustic experiments in the laboratory (Larose *et al.* 2004) and has been used in a number of early seismic studies of coda waves and ambient noise (Campillo & Paul 2003; Shapiro & Campillo 2004; Shapiro *et al.* 2005; Yao *et al.* 2006). The second method, employed for example by Sabra *et al.* (2005a), involves the application of a clipping threshold equal to the rms amplitude of the signal for the given day. An example is shown in Fig. 3(c). The third method, illustrated by Fig. 3(d), involves automated event detection and removal in which 30 min of the waveform are set to zero if the amplitude of the waveform is above a critical threshold. This threshold is arbitrary and its choice is made difficult by varying amplitudes at different stations. The fourth method is running-absolute-mean normalization, which is the method of time normalization that we promote here. This method computes the running average of the absolute value of the waveform in a normalization time window of fixed length and weights the waveform at the centre of the window by the inverse of this average. That is, given a discrete time-series d_j, we compute the normalization weight for time point n as:

$$w_n = \frac{1}{2N+1} \sum_{j=n-N}^{n+N} |dj|, \tag{1}$$

so that the normalized datum becomes $\tilde{d}_n = d_n / w_n$. The width of the normalization window $(2N + 1)$ determines how much amplitude information is retained. A one-sample window ($N = 0$) is equivalent to one-bit normalization, while a very long window will approach a re-scaled original signal as $N \to \infty$. After testing various time window widths, we find that about half the maximum period of the passband filter works well and that this length can be varied considerably and still produce similar results. An example result of the application of this method is shown in Fig. 3(e). This method is not without its faults, however. For example, it does not surgically remove narrow data glitches, as it will inevitably down-weight a broad time interval around the glitch. One-bit normalization does not suffer from this shortcoming. Finally, there is a method that we call iterative 'water-level' normalization in which any amplitude above a specified multiple of the daily rms amplitude is down-weighted. The method is run repeatedly until the entire waveform is below the water-level, which is six times the daily rms level in the example shown in Fig. 3f. This method of time-domain normalization is the most time intensive of the candidates considered here.

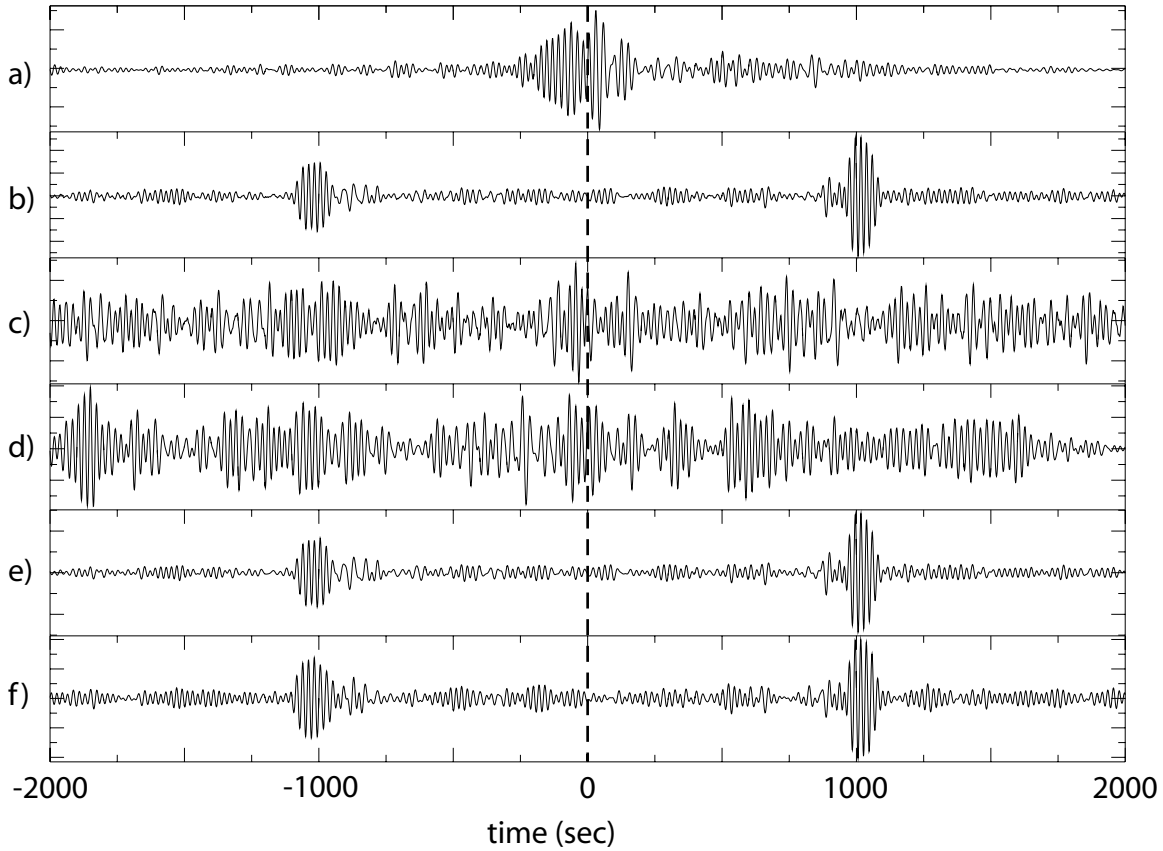

Figure 4. Twelve-month cross-correlations between the station-pair ANMO and HRV for the time-domain normalization methods shown in Fig. 3. The passband is 20–100 s period. The panels of the Figs (a)–(f) correspond to those in Fig. 3.

Fig. 4 presents examples of year-long cross-correlations, bandpass filtered between 20 s and 100 s period, using each of these methods of time-domain normalization. The raw data (Fig. 4a), the clipped waveform method (Fig. 4c), and the automated event detection method (Fig. 4d) produce noisy cross-correlations in this period band. The one-bit normalization (Fig. 4b), the running-absolute-mean normalization (Fig. 4e), and the water-level normalization (Fig. 4f) methods produce relatively high SNR waveforms displaying signals that arrive at nearly the same time. In this example, the one-bit and the running-absolute-mean normalizations are nearly identical. A systematic test has been performed using 15 GSN stations in North America using the observed spectral SNR (defined in Section 3) at 20 s period to compare the methods at five periods. The resulting SNR values are similar for one-bit normalization and the running-absolute-mean normalization. The water-level normalization method also allows meaningful results to be recovered. The running-absolute-mean method provides a small enhancement to SNR values above one-bit normalization at all periods and a more significant improvement over the water-level normalization.

The principal reason we prefer running-absolute-mean normalization over the water-level or one-bit normalization methods is its greater flexibility and adaptability to the data. For example, in areas with high regional seismicity it is desirable to tune the time-domain normalization to the frequency content of the seismicity. Fig. 5 shows that if the temporal weights of the running-absolute-mean normalization are computed on the raw waveform data, small earthquakes can get through the procedure because they exist in the raw waveform near the background noise level. Earthquakes are revealed by a low-pass filter both in the raw waveform (Fig. 5b) and the temporally normalized waveform (Fig. 5d). Alternately, the temporal weights of the running-absolute-mean normalization can be computed on the waveform filtered in the earthquake band (Fig. 5b). In this case, if d_j is the raw seismogram and $\hat{d}_j$ is the seismogram bandpass filtered in the earthquake band, we define new temporal weights calibrated to the regional seismicity

$$\hat{w}_n = \frac{1}{2N+1} \sum_{j=n-N}^{n+N} |\hat{d}j|. \tag{2}$$

These weights are then applied to the raw data as before ($\tilde{d}_n = d_n/\hat{w}_n$). This procedure severely down-weights time-series during earthquakes (Fig. 5e), which more effectively removes them from low-pass filtered seismograms (Fig. 5f). Contamination by earthquakes of the cross-correlations, therefore, should be ameliorated.

Earthquake signals that pass through the temporal normalization tend to appear on cross-correlations as spurious precursory arrivals, such as the high amplitude arrivals appearing between 0 and 100 s in the 12-month cross-correlation shown in Fig. 6(a). Defining the temporal normalization weights in the earthquake band, however, reduces the amplitude of the precursors, as Fig. 6(b) illustrates. This process will be most important in regions with significant regional seismicity. The example shown in Fig. 6 is from New Zealand where, because of high levels of seismicity in the Fiji and Tonga–Kermadec regions, the process is recommended strongly (Lin *et al.* 2007).

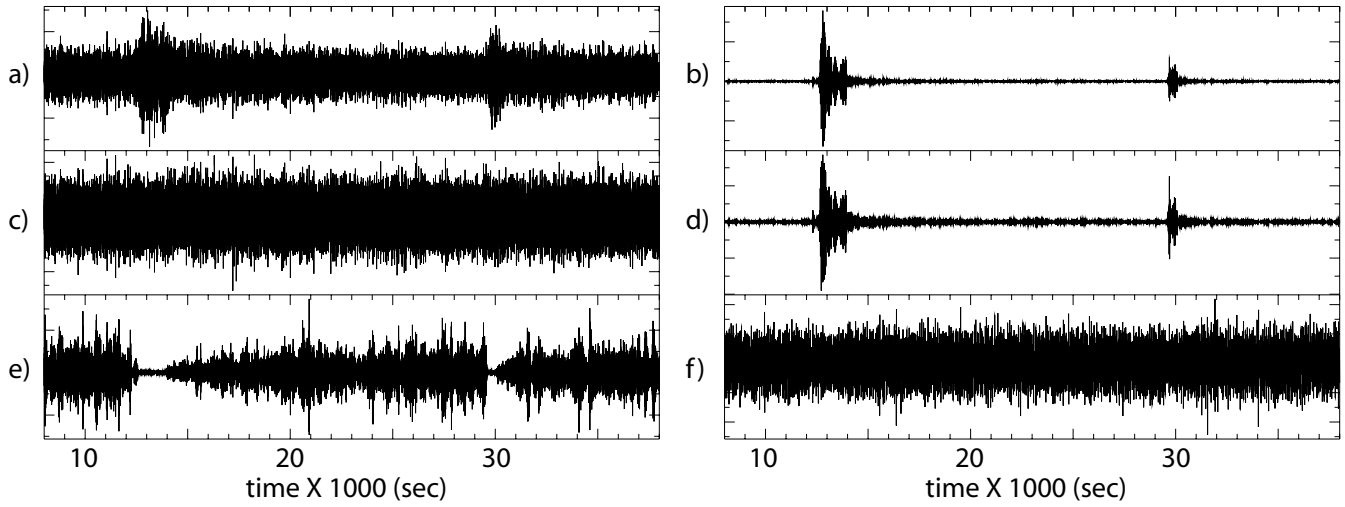

Figure 5. Example of the effect of tuning time-domain normalization to earthquake signals for data from GeoNet station CRLZ in New Zealand. (a) Raw broad-band data from Oct. 14, 2005 showing two earthquakes barely emerging above background noise. (b) Data from (a) bandpass filtered between 15 and 50 s period, more clearly showing the two earthquake signals (first: S. Fiji, $m_b = 5.4$; second: S. of Kermadec, $m_b = 5.1$). (c) Data after temporal normalization using the running-absolute-mean method in which the weights are defined on the raw (unfiltered) data in (a). (d) Data from (c) bandpass filtered between 15 and 50 s period, showing that the earthquake signals have not been removed by temporal normalization defined on the raw data. (e) Data after temporal normalization using the running-absolute-mean method in which the weights are defined on the bandpass filtered data in (b). (f) Data from (e) bandpass filtered between 15 and 50 s period, showing that the earthquake signals have been removed by temporal normalization defined on the bandpass filtered data.

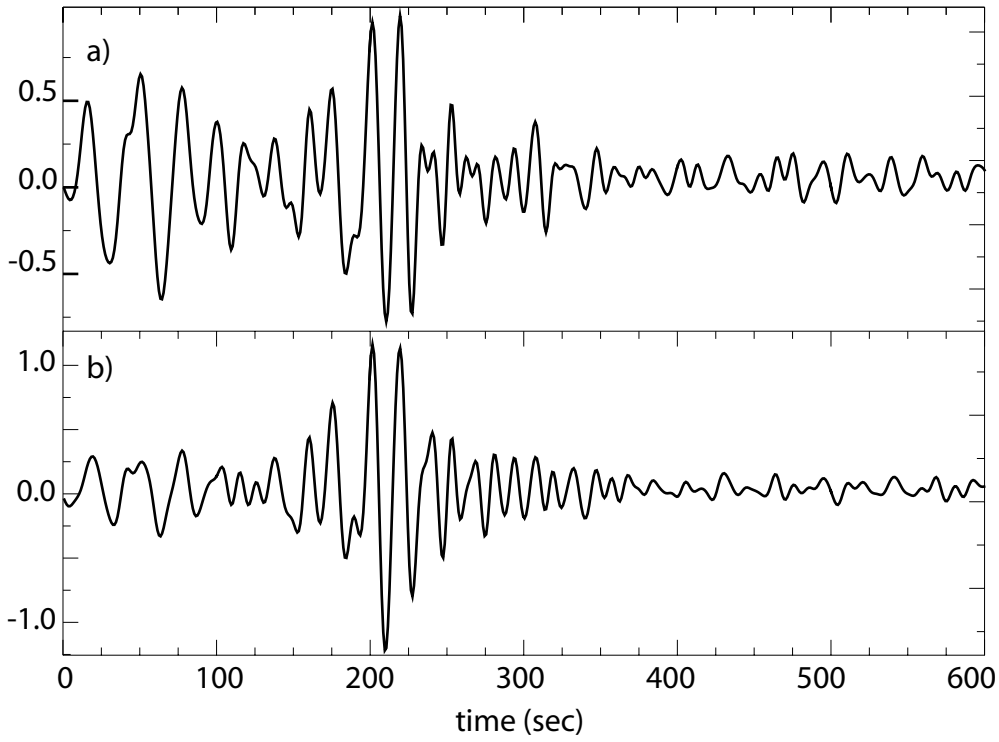

Figure 6. Example of the effect of tuning time-domain normalization to earthquake signals on cross-correlations computed between GeoNet stations CRLZ and HIZ in New Zealand. (a) Year-long cross-correlation in which the temporal normalization is defined on the raw data. (b) Year-long cross-correlation in which the temporal normalization is defined on data bandpass filtered between 15 and 50 s period. Spurious precursory arrivals are substantially reduced in (b) relative to (a). Waveforms are bandpass filtered between 5 and 50 s period.

2.2 Spectral normalization or whitening

Ambient noise is not flat in the frequency domain (i.e. is not spectrally white), but is peaked near the primary (around 15 s period) and secondary (around 7.5 s period) microseisms and rises at very long periods above 50 s to form a signal now referred to as Earth 'hum' (e.g. Rhie & Romanowicz 2004). Fig. 7(a) presents an example of an amplitude spectrum for a day long time-series obtained after temporal normalization. Primary and secondary microseisms

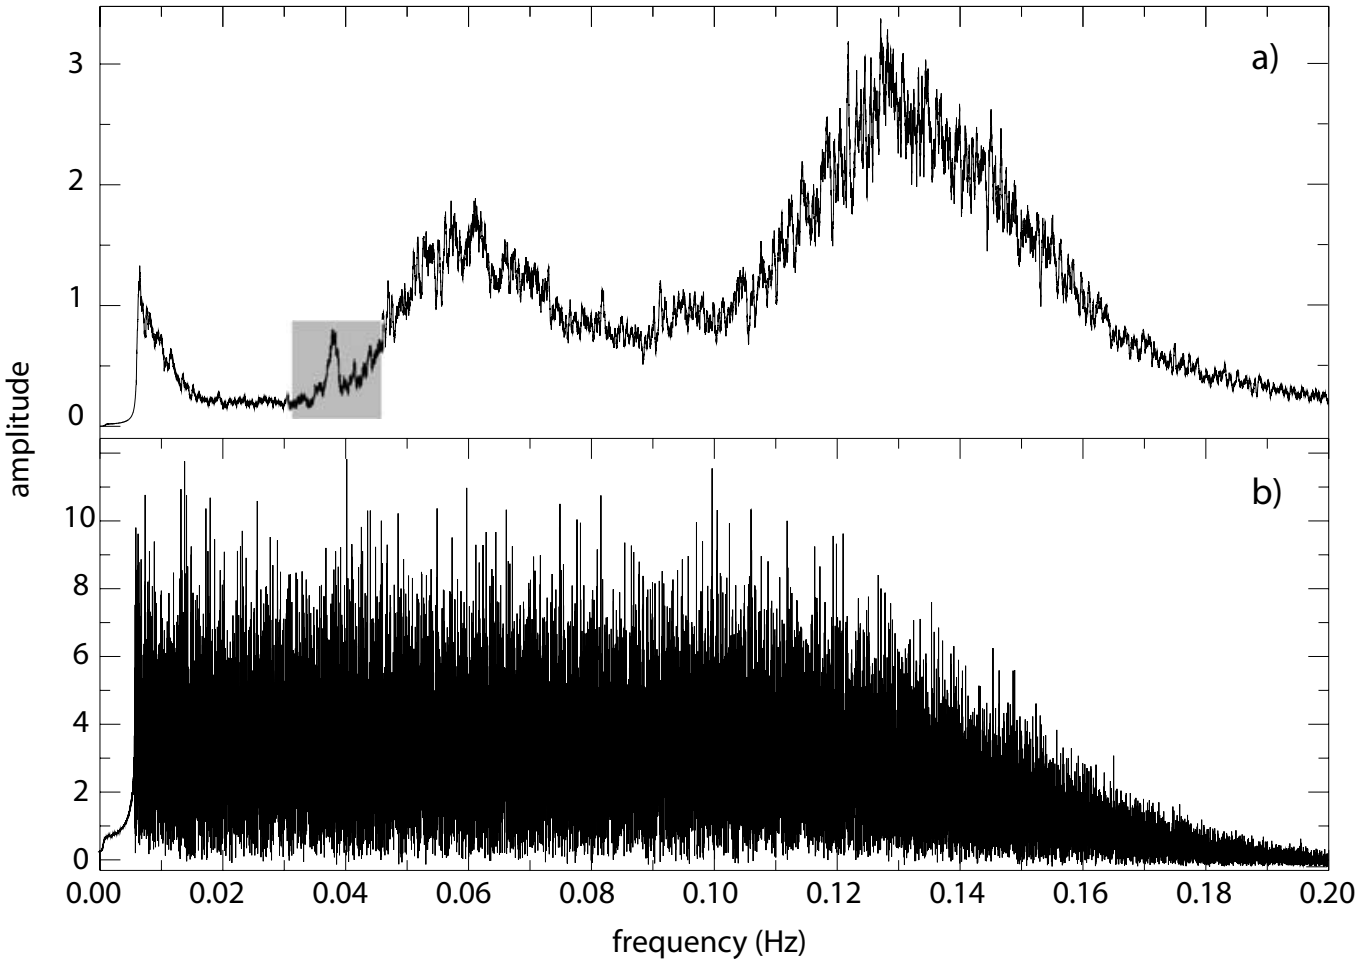

Figure 7. (a) Raw and (b) spectrally whitened amplitude spectra for 1 sample per second vertical component data at station HRV for July 5, 2004. The shaded box indicates the location of the 26 s period signal originating from the Gulf of Guinea. The taper seen at both ends of the spectra is largely attributable to a 7–150 s bandpass filter.

as well as Earth hum signatures can be seen clearly on this record which was bandpass filtered between 7 and 150 s period. In addition to these signals, there is a smaller peak near 26 s that is caused by a persistent narrow-band noise source in the Gulf of Guinea (Shapiro *et al.* 2006). Without the temporal normalization, which reduces the effect of earthquakes, the 26 s resonance typically is not seen. Ambient noise is minimum in the period range from about 30 to 70 s. Inversely weighting the complex spectrum by a smoothed version of the amplitude spectrum produces the normalized or whitened spectrum shown in Fig. 7(b). Spectral normalization acts to broaden the band of the ambient noise signal in cross-correlations and also combats degradation caused by persistent monochromatic sources such as the Gulf of Guinea source.

First, regarding the problem of an isolated, persistent nearly monochromatic noise source, the grey box in Fig. 7(a) highlights the noise peak at 26 s period as observed at the station HRV on a northern summer day. As documented by Holcomb (1998), this signal is seasonal, being much stronger in the northern summer than in the winter. Fig. 8(a) shows a 12-month cross-correlation between GSN stations ANMO and CCM in which spectral normalization has not been applied. The 26 s resonance appears as a broad envelope in the time domain and corrupts the cross-correlation at positive correlation lag. Shapiro *et al.* (2006) used the apparent arrival time of the 26 s signal observed at stations in North America, Europe, Africa and Asia to locate the source in the Gulf of Guinea. The amplitude spectrum of this cross-correlation displays the prominent peak at ∼26 s period (∼0.038 Hz) as seen in Fig. 8(b). In contrast, Figs 8(c) and (d) show the cross-correlation and its amplitude spectrum where spectral normalization has been applied. The effect of the 26 s resonance is greatly reduced. Shapiro *et al.* (2006) recommend eradicating this problem by applying a narrow band reject filter centred around 26 s period. Figs 8(e) and (f) show the effect of this filter. The cross-correlation is largely unchanged compared to spectral whitening. In many cases, therefore, the more gentle approach of spectral whitening is sufficient to eliminate the 26 s problem from the cross-correlations. The band-reject filter also creates problems for automated dispersion measurement in a later stage of processing, so spectral whitening is preferable if it suffices to ameliorate the effect of the 26 s microseism.

Second, spectral normalization seeks to reduce broad imbalances in single-station spectra to aid in the production of a broad-band dispersion measurement. Figs 9(a) and (b) show a one-month broad-band cross-correlation between stations CCM (Cathedral Cave, MO, USA) and SSPA (Standing Stone, PA, USA) for spectrally unwhitened and whitened data taken during the northern spring (when the 26 s resonance is weak). Figs 9(c) and (d) display the amplitude spectra of the unwhitened and whitened cross-correlations, respectively. Without the whitening, Fig. 9(c) shows that the resulting

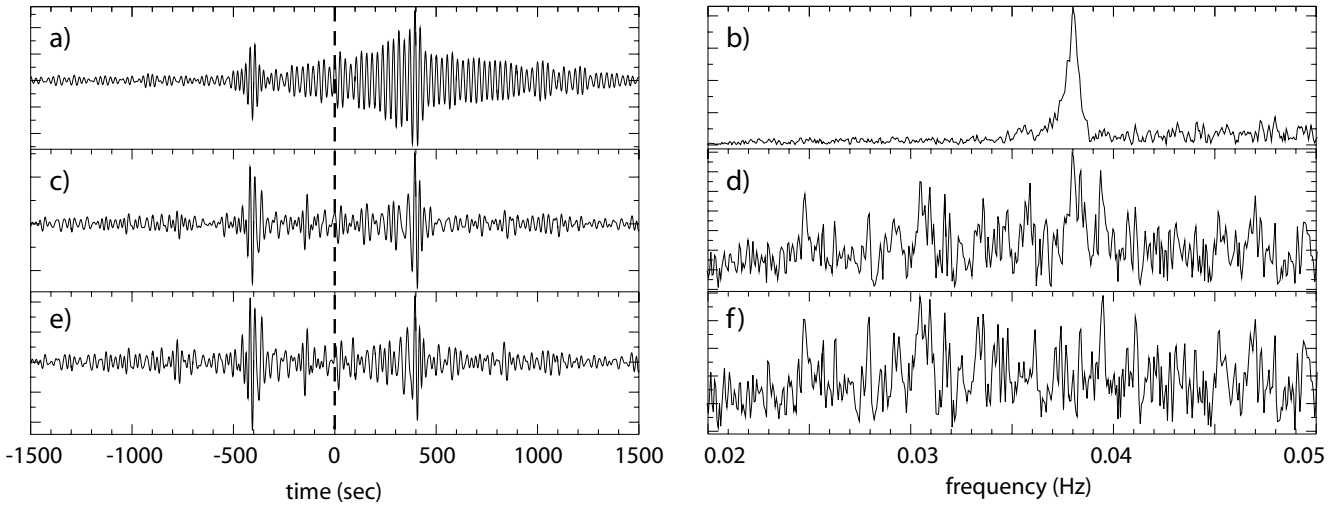

Figure 8. Effect of the 26 s microseism on cross-correlations and attempts to remove it. (a) Twelve-month cross-correlation between data from stations ANMO and CCM (Cathedral Cave, MO, USA). The broad, nearly monochromatic 26 signal at positive lag dominates the waveform. (b) Amplitude spectrum of the cross-correlation in (a) showing the spectral peak at about 26 s period. (c) Cross-correlation between data from the same two stations that have been spectrally whitened prior to cross-correlation. (d) Amplitude spectrum of the cross-correlation in (c) showing that the 26 s spectral peak is largely missing. (e) Cross-correlation between the data that have been spectrally whitened prior to cross-correlation with a notch filter applied around 26 s period. (f) Amplitude spectrum of the cross-correlation in (e). Application of the notch filter changes the cross-correlation only minimally.

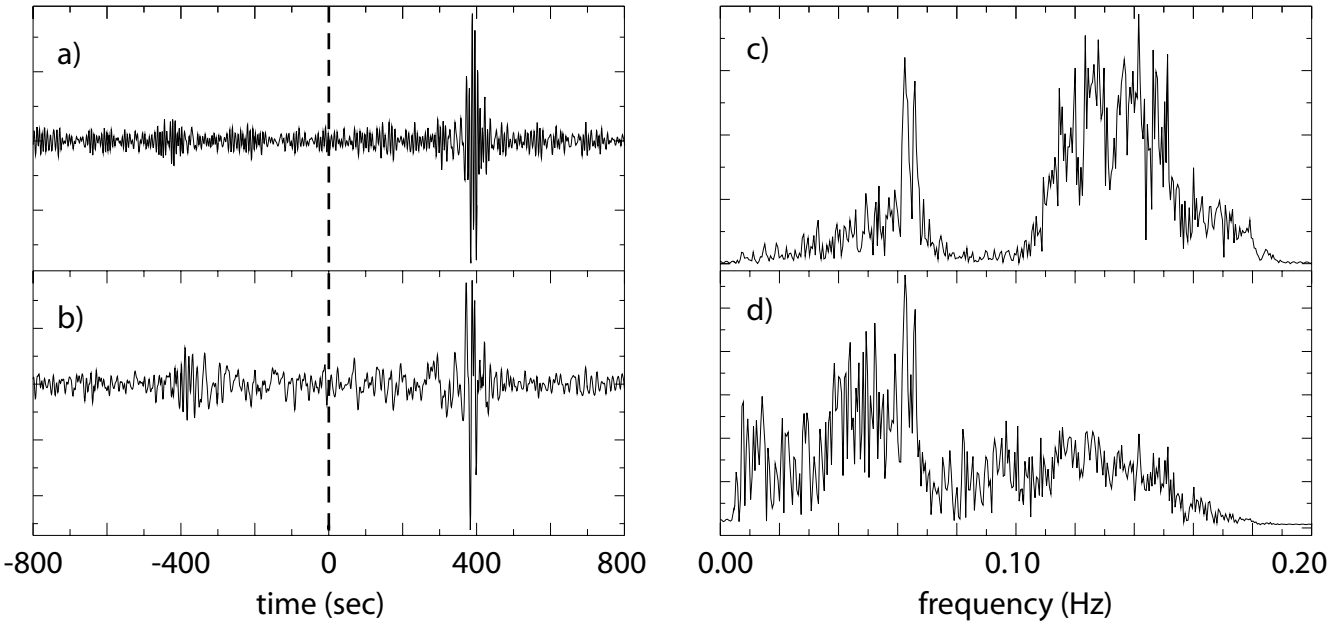

Figure 9. Comparison of cross-correlations with and without spectral whitening. Cross-correlation are for the month April, 2004 for data from stations CCM and SSPA (Standing Stone, PA, USA) bandpass filtered from 7 to 150 s period. (a) Cross-correlation without spectral whitening. (b) Cross-correlation with spectral whitening. (c) Amplitude spectrum of the unwhitened waveform in (a). The primary and secondary microseisms dominate the spectrum. (d) Amplitude spectrum of the pre-whitened waveform in (b).

cross-correlation is dominated by signals in the microseism band, predominantly from 15 to 17 s and from the 6 to 9 s period. Not surprisingly, spectral whitening produces a broader-band signal. In many cases, the cross-correlation amplitude spectrum is shaped with the longer periods having higher amplitudes than the shorter periods, as in Fig. 9(d). This is apparently because the longer period ambient noise, although naturally lower in amplitude than microseismic noise, propagates more coherently over long distances. Additional whitening of the cross-correlation prior to dispersion measurement is an added option.

3 CROSS-CORRELATION, STACKING AND SIGNAL EMERGENCE

After the preparation of the daily time-series described in Section 2, the next step in the data processing scheme (Phase 2) is cross-correlation and stacking. Although some interstation distances may be either too short or too long to obtain reliable measurements, we perform cross-correlations between all possible station pairs and perform data selection later. This yields a total of $n(n-1)/2$ possible station pairs, where n is the number of stations. Obtaining tens of thousands of cross-correlations is common when ambient noise

data processing is performed over large spatial scales such as Europe (e.g. Yang *et al.* 2007) or North America (e.g. Bensen *et al.* 2005).

Cross-correlation is performed daily in the frequency domain. After the daily cross-correlations are returned to the time domain they are added to one another, or 'stacked', to correspond to longer time-series. Alternately, stacking can be done in the frequency domain which would save the inverse transform. We prefer the organization that emerges from having daily raw time-series and cross-correlations that are then stacked further into weekly, monthly, yearly, etc. time-series. In any event, the linearity of the cross-correlation procedure guarantees that this method will produce the same result as cross-correlation applied to the longer time-series. The resulting cross-correlations are two-sided time functions with both positive and negative time coordinates, i.e. both positive and negative correlation lags. We typically store the correlations from −5000 to 5000 s, but the length of the time-series needed will depend on the group speeds of the waves and the longest interstation distance.

The positive lag part of the cross-correlation is sometimes called the 'causal' signal and the negative lag part the 'acausal' signal. These waveforms represent waves travelling in opposite directions between the stations. Several examples of cross-correlations have been shown earlier in the paper. Figs 4, 8 and 9 display some two-sided cross-correlations for different time-series lengths. Fig. 1 clearly shows the broad-band content of ambient noise. If sources of ambient noise are distributed homogeneously in azimuth, the causal and acausal signals would be identical. However, considerable asymmetry in amplitude and spectral content is typically observed, which indicates differences in both the source process and distance to the source in the directions radially away from the stations. We often compress the two-sided signal into a one-sided signal by averaging the causal and acausal parts. We call this the 'symmetric' signal or component. An example was shown in Figs 1 and 6.

Stacking over increasingly long time-series, on average, improves SNR ratio. An example is shown in Fig. 10, which displays cross-correlations of different length time-series from the stations ANMO and DWPF (Disney Wilderness Preserve, FL, USA). The causal and acausal signals are seen to emerge as the time-series length increases in both of the period bands that are displayed in Figs 10(a) and (b).

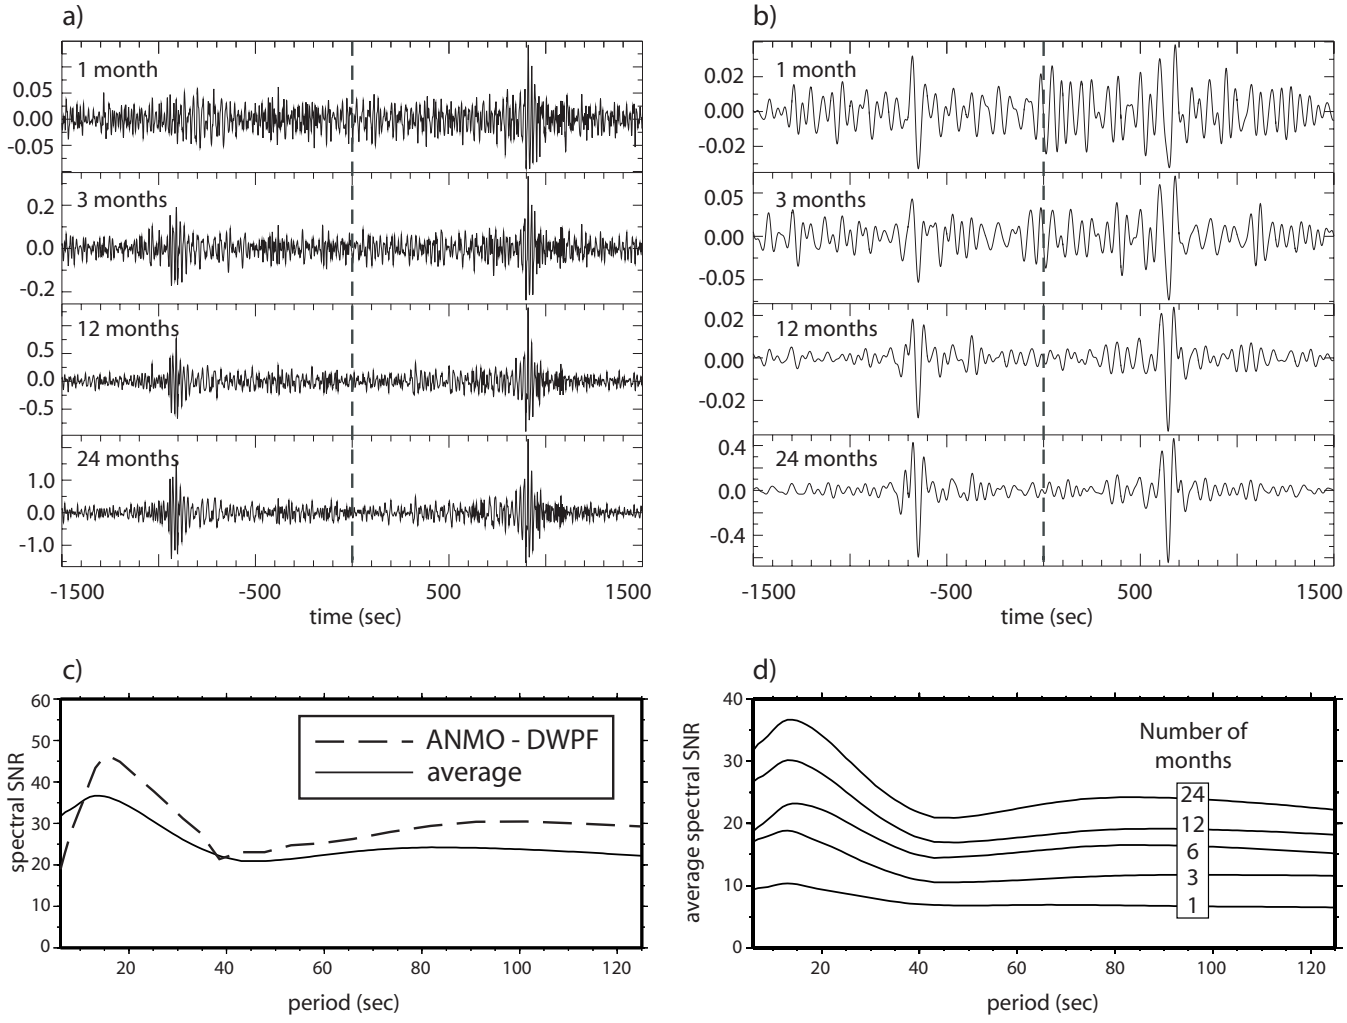

Figure 10. Example of the emergence of the Rayleigh waves for increasingly long time-series. (a) Cross-correlations at the specified time-series lengths for the station pair ANMO and DWPF (Disney Wilderness Preserve, FL, USA) bandpassed between 5 and 40 s period. (b) Same as (a), but for a passband between 40 and 100 s period. (c) Spectral SNR for the 24-month ANMO-DWPF cross-correlation shown with a dashed line, and the spectral SNR averaged over all cross-correlations between GSN stations in the US shown with a solid line. (d) Spectral SNR averaged over all cross-correlations between GSN stations in the US for different time-series lengths of 1, 3, 6, 12 and 24 months.

Measurements of the frequency dependence of the SNR are useful to quantify observations of the emergence of the signals with increasing time-series length. We also use it as part of data selection in Phase 4 of the data processing procedure. Fig. 11 illustrates one way in which the frequency dependence of SNR may be measured. From the 3-D model of Shapiro & Ritzwoller (2002), we predict the maximum and minimum group arrival times ($t_{\min}$, $t_{\max}$) expected for the path between the station-pair over the period band of interest ($\tau_{\min}$, $\tau_{\max}$). We perform a series of narrow bandpass filters centred on a discrete grid of frequencies and measure the peak in the time domain in a signal window ($t_{\min} - \tau_{\max}, t_{\max} + 2\,\tau_{\max}$) shown with solid vertical lines in Fig. 11. We also measure the rms noise level in a 500 s noise window (vertical dashed lines) that trails the end of the signal window by 500 s. This rms level is shown with dotted lines in Fig. 11 in the noise window. The resulting ratio of peak signal in the signal window to rms noise in the trailing noise window on the grid of centre frequencies is the 'spectral' SNR measurement. Centre periods and SNR are identified in each panel of Fig. 11. Note that although we call this a spectral SNR measurement, it is, in fact, a measurement of SNR in the time domain. It is 'spectral' only in the sense that the measurements are a function of frequency.

This spectral SNR, which takes the ratio of signal to trailing noise, mostly is a measure of the signal level, as the trailing noise does not strongly depend on signal-generated noise. Alternately, one could define the ratio of signal to leading noise, which is strongly dependent on signal-generated noise, as discussed earlier. Although we do not use signal-to-precursory noise here, it has the advantage of quantifying precursory noise which interferes with dispersion measurements more than trailing noise. Further research is needed to determine if it is a better predictor of the quality of dispersion measurements than the spectral SNR that we use.

A spectral SNR curve for the 24-month cross-correlations between stations ANMO and DWPF, shown in Figs 10(a) and (b), is presented as the dashed line in Fig. 10(c). It is contrasted with the average SNR over all GSN station pairs within the US. For this example, spectral SNR, on average, peaks in the primary microseism band around 15 s period, minimizes near 40 s period and then is fairly flat to much longer periods, although it rises slightly. The details of the curve, however, will vary geographically, with path length and season. Fig. 10(d) shows how spectral SNR increases with time-series length. The shapes of the SNR spectra also change subtly with time-series length.

In general, therefore, as time-series length increases so does SNR, so the longer the time-series the better. The details of how the signal emerges from noise depends on frequency, and also on the location and interstation spacing. Fig. 12 presents an example of how SNR depends on time-series length computed for the 15 GSN stations in the US. The emergence of the signal can be fit well with a power law, and Fig. 12 shows the fit power law rather than the raw data: SNR = $At^{1/n}$, where A and n are period dependent. For the periods shown in Fig. 12, n varies from about 2.55 at 10 s period to 2.88 at 25 s. It attains a maximum of about 3.4 at 50 s and then diminishes again so that at 100 s period n is about equal to 2.66. Inspection of Fig. 7, which is a typical daily amplitude spectrum for temporally normalized data, reveals that n maximizes at intermediate periods between about 25 and 50 s where ambient noise is generally weakest. In this period band, the emergence of the signal is slowest. At shorter and longer periods, in the microseismic and 'Earth-hum' bands, n ranges from about 2.5 to 2.9, and the signal emerges at a faster rate than at the intermediate periods. As discussed in Section 6 below, the curves in Fig. 12 are useful in designing experiments based on ambient noise tomography. Further work, however, is needed to understand the frequency dependence of the power law behaviour of the emergence of the signal from ambient noise, as well as its geographic variability.

4 DISPERSION MEASUREMENT

After the daily cross-correlations have been computed and stacked, the resulting waveform is an estimated Green function. Using the estimated Green function, the group and phase speeds as a function of period can be measured by using traditional frequency–time analysis (FTAN) (e.g. Dziewonski *et al.* 1969; Levshin *et al.* 1972, 1992; Herrin & Goforth 1977; Russell *et al.* 1988; Levshin *et al.* 1989; Ritzwoller & Levshin 1998; Levshin & Ritzwoller 2001). This is Phase 3 of the data processing procedure. As with Phases 1 and 2, because the number of interstation pairs can be very large, the dispersion measurement process needs to be automated. The method that we promote is based on a version of FTAN described in detail by Levshin *et al.* (1989), which obtains measurements on single waveforms and involves significant analyst interaction. However, the computational structure of FTAN allows automation and this is what we describe here. Although FTAN has been applied dominantly to measure group speeds, phase speed curves are also measured naturally in the process.

We roughly follow the notation and terminology of Bracewell (1978), but if $s(t)$ is the waveform of interest its Fourier transform is defined with a positive exponent as $S(\omega) = \int_{-\infty}^{\infty} s(t) \exp(i\omega t) dt$. Dispersion measurements are obtained by considering the 'analytic signal', which is defined simply in the frequency domain as

$$S_a(\omega) = S(\omega)(1 + \mathrm{sgn}(\omega)), \tag{3}$$

and upon inverse Fourier transforming is expressed in the time domain as follows:

$$S_a(t) = s(t) + i\,H(t) = |A(t)| \exp(i\phi(t)). \tag{4}$$

$H(t)$ is the Hilbert transform of $s(t)$. To construct a frequency—time function, the analytic signal is subjected to a set of narrow bandpass Gaussian filters with centre frequencies ω_0:

$$S_a(\omega, \omega_0) = S(\omega)(1 + \mathrm{sgn}(\omega))G(\omega - \omega_0), \tag{5}$$

$$G(\omega - \omega_0) = e^{-\alpha\left(\frac{\omega-\omega_0}{\omega_0}\right)^2}. \tag{6}$$

Inverse transforming each bandpassed function back to the time domain yields the smooth 2-D envelope function, $|A(t, \omega_0)|$, and phase function, $\phi(t, \omega_0)$. α is a tunable parameter that defines the complementary resolutions in the frequency and time domains and is commonly made distance dependent (Levshin *et al.* 1989). Group speed is measured using $|A(t, \omega_0)|$ and phase speed using $\phi(t, \omega_0)$. In particular, the group arrival time, $\tau(\omega_0)$, as a function of the centre frequency of the Gaussian filter is determined from the peak of the envelope function so that the group speed is $U(\omega_0) = r/\tau(\omega_0)$, where r is the interstation distance. We follow Bracewell and replace ω_0 with the 'instantaneous frequency', defined as the time rate of change of the phase of the analytic signal at time τ. We, therefore, replace the centre frequency of the narrow-band filter, ω_0, with the instantaneous frequency, $\omega = |d\phi(t, \omega_0)/dt|_{t=\tau(\omega_0)}$. This correction is most significant when the spectrum of the input waveform is not flat, in which case, due to spectral leakage, the central frequencies of the narrow-band filters will not accurately represent the frequency content of the output of the filters.

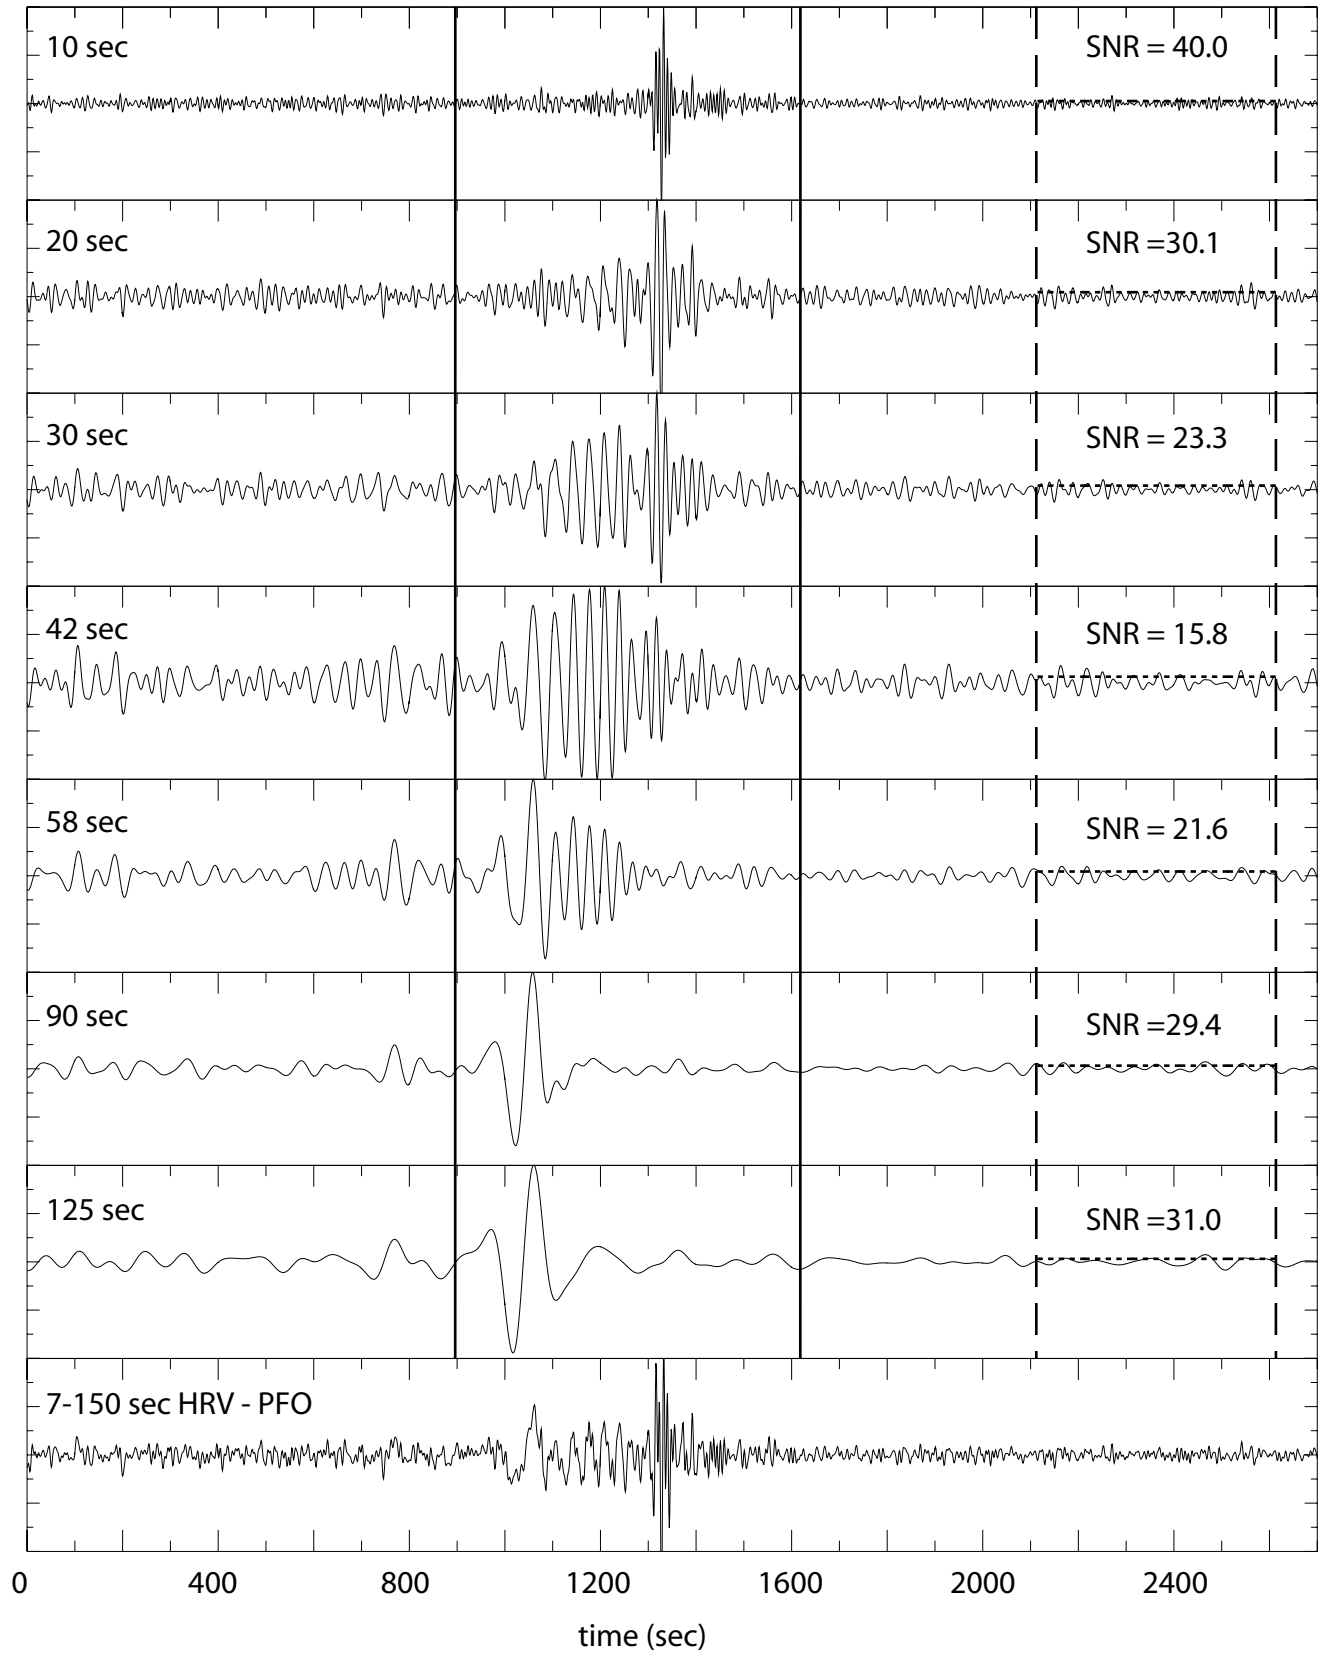

Figure 11. Example of how spectral SNR measurements are obtained on a 12-month cross-correlation between data from stations HRV and PFO (Pinyon Flat, CA, USA). Vertical solid lines indicate the signal windows and vertical dashed lines the noise windows. Waveforms are centred on the period indicated at the left-hand side in each panel, and SNR is defined as the ratio of the peak within the signal window and rms noise in the noise window. The noise level is presented as the horizontal dotted lines in the noise windows. SNR in each band is indicated at right-hand side in each panel.

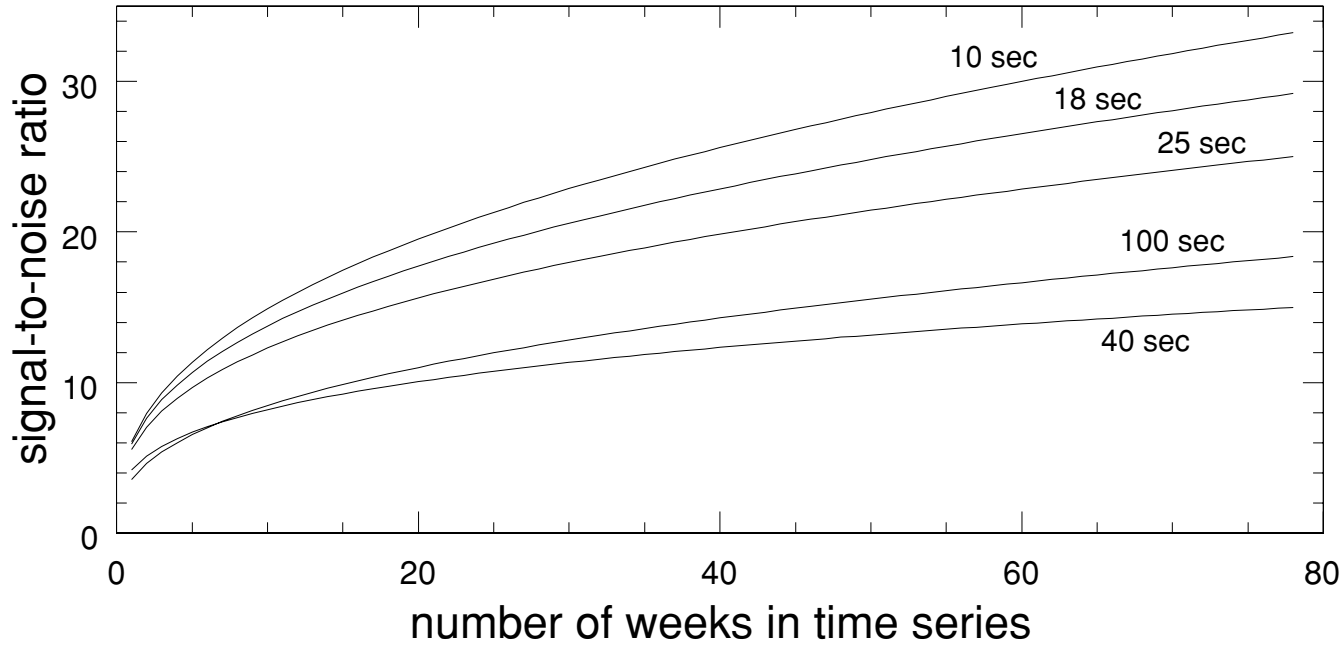

Figure 12. Emergence of the signal with time-series length. The power law fit to the average of the measured SNR from cross-correlations between the GSN stations within the US at each of the five indicated periods is plotted versus variable time-series length (in weeks).

The measurement of dispersion curves divides into eight steps. We discuss each step and then indicate how the analyst-driven and the automated FTAN processes differ. This will be done in the context of group velocity measurements in Section 4.1 and then we will follow with a discussion of how FTAN measures phase speed curves in Section 4.2. Fig. 13 graphically illustrates the process. In this figure, all results are for the automated FTAN process.

4.1 Group speed measurements

Fig. 13(a) shows a broad-band signal obtained from a one-year cross-correlation between stations ANMO and COR in the US. In Step 1 of FTAN, a frequency (period)—time (group speed) or FTAN image is produced by displaying the logarithm of the square of the envelope of the analytic signal, $\log |A(t, \omega_0)|^2$. Fig. 13(b) shows the FTAN image of the waveform in Fig. 13(a). The envelope functions $\log |A(t, \omega_0)|^2$ are arrayed vertically on a grid of different values of ω_0 to produce a matrix that can be displayed as a 2-D image. There is a similar phase matrix not displayed here. Typically, group speed replaces time and period replaces filter centre frequency. In Step 2, the dispersion ridge is tracked as a function of period to obtain a raw group speed curve. Fig. 13(b) shows this curve and the prediction from the 3-D model of Shapiro & Ritzwoller (2002). This raw group speed measurement may be sufficient for many applications.

Steps 3-8 of FTAN involve phase-matched filtering to clean the waveform of potential contamination and generate an alternative group speed curve. This measurement may be preferable in some applications. In Step 3, an anti-dispersion or phase-matched filter is defined on a chosen period-band. Levshin & Ritzwoller (2001) discuss the phase-matched filtering method in detail. In Step 4, this anti-dispersion filter is applied to the waveform in the period band chosen to produce the undispersed signal. Fig. 13(c) shows the undispersed or 'collapsed' signal. In Step 5, contaminating noise is identified and removed from the undispersed signal. Typically, for earthquakes this noise is signal generated, being composed of multipathed signals, seismic coda, body waves, and so forth. An example cut is shown with the red line in Fig. 13(c). In Step 6, the cleaned collapsed waveform is redispersed. It is shown as the red line in Fig. 13(a). In Step 7, the FTAN image of the cleaned waveform is computed using the same process applied to the raw waveform in Step 1. Fig. 13(d) shows the FTAN image of the cleaned waveform. To improve frequency resolution, the Gaussian filters that are applied during phase-matched filtering are broader than those that are applied to the raw waveform. For this reason, the time-width of the FTAN image is broader in Fig. 13(d) than in Fig. 13(b), but this does not reflect a lower intrinsic temporal resolution because interfering signals have been removed. Finally, in Step 8, the dispersion ridge is tracked as a function of period on the cleaned FTAN image to obtain the cleaned group speed curve. Fig. 13(d) shows this curve and the predicted curve from the 3-D model.

The traditional analyst-driven FTAN procedure has been applied to earthquake data by analysts for more than 200 000 individual paths globally (e.g. Shapiro & Ritzwoller 2002). The analyst, however, only enters the process in Steps 3 and 5. In Step 3, the analyst defines the phase-matched filter and the frequency band of interest, which usually depends on the bandwidth of the signal that is observed. The analyst either can use the group speed curve that is automatically produced on the raw FTAN image in Step 2 or can define a curve interactively. The latter approach is usually chosen, as FTAN images of earthquake data commonly display spectral holes that vitiate the automated group speed measurement. The automated group speed measurements are also often tricked by scattered or multipathed arrivals and, therefore, do not track the dispersion branch of interest accurately. Multipathing and scattering is a problem mostly for large epicentral distances. In Step 5, the analyst interacts with the collapsed signal to remove noise. It is, therefore, only Steps 3 and 5 that require automation beyond the existing method.

To automate Step 3, the group speed measurements that result in Step 2 must be used to define the phase-matched filter. Therefore, these measurements must be robust to spectral holes and scattered or multipathed arrivals. Fortunately, high SNR FTAN images that result

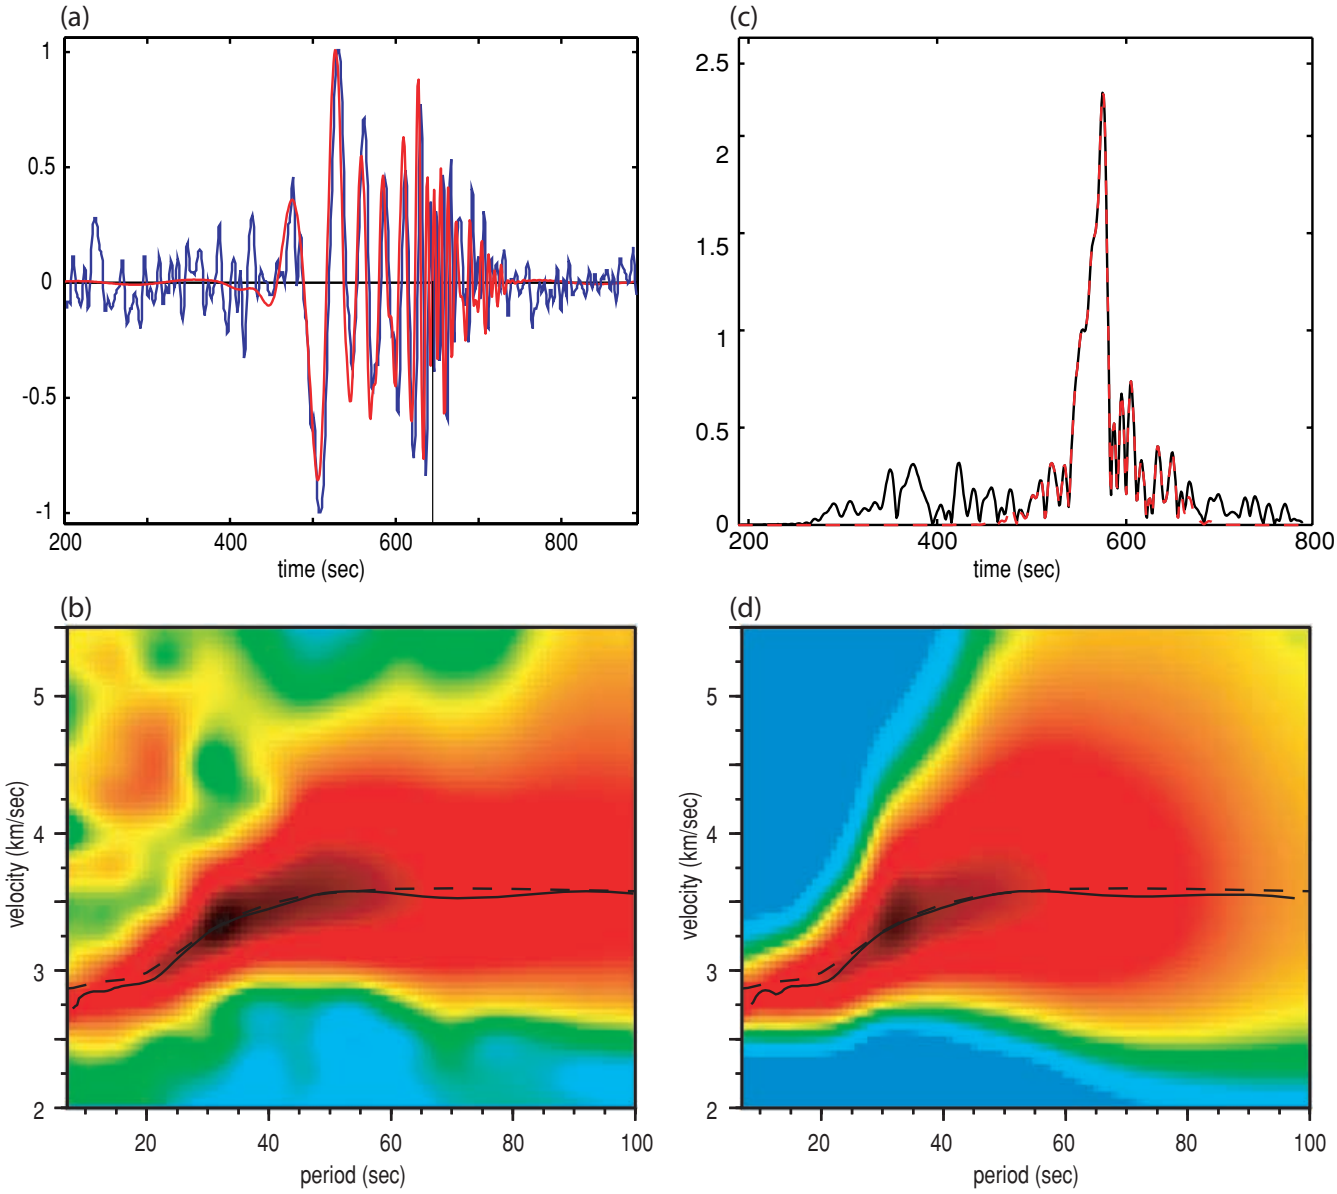

Figure 13. Graphical representation of FTAN. (a) Raw (blue) and cleaned (red) waveforms for the 12-month stacked cross-correlation between stations ANMO and COR (Corvallis, OR, USA). (b) Raw FTAN diagram, measured group speed curve as the solid line and prediction from the 3-D model of Shapiro and Ritzwoller (2002) as the dashed line. (c) Undispersed or collapsed signal (black) and cleaned signal (red dashed). (d) Cleaned FTAN diagram, measured group speed curve and prediction from the 3-D model of Shapiro and Ritzwoller (2002).

from cross-correlations of ambient noise tend to be much simpler than those from earthquakes, and spectral holes are rare. Interstation spacings for ambient noise measurements are also typically shorter than epicentral distances, so multipathing is not as severe of a problem. The automated procedure, therefore, only differs from the raw group velocity procedure applied during interactive FTAN in that in Step 2 added measures are taken to ensure the continuity of the dispersion curve by rejecting spurious glitches or jumps in group times. Formal criteria are set to reject curves with distinctly irregular behaviour or to interpolate through small glitches by selecting realistic local instead of absolute maxima. When gaps or jumps are too large in amplitude or persistent in period, the dispersion curve is rejected. Spectral whitening (Section 2.2) helps to minimize jumps in the measured curve as well as the incompleteness of measurements at the long period end of the spectrum. The raw group speed curve that emerges from Step 2 is one of two alternative curves that emerge from the automated process.

To automate Step 5, the undispersed signal is selected from the surrounding noise automatically. Fig. 13(c) illustrates this procedure graphically using the waveform from Fig. 13(a). In an ideal case, the anti-dispersed signal will collapse into a single narrow spike. The collapsed waveform, given by the red line in Fig. 13(c), is then cut from the surrounding time-series and re-dispersed to give the clean waveform shown with the red line in Fig. 13(a). In this example the collapsed waveform is more complicated than a single spike. The principal advantage of this phase-matched filtering method arises when there exists strong neighbouring noise that can be removed from the undispersed signal. In the case of ambient noise cross-correlations, spurious precursory arrivals exist in many cases, particularly at long periods. A good example can be seen in

Fig. 6(a), and the example in Fig. 13(a) also contains precursory arrivals. Such arrivals tend to interfere with the primary signals and can make the resulting group velocity curves undulatory. Phase-matched filtering helps to reduce the effect of precursory arrivals and smooths the measured group speed curve in general. The FTAN image after phase-matched filtering is broader in time (Fig. 13d) because a larger α-value (eq. 6) is applied compared to the raw image (Fig. 13b) In the example in Fig. 13, however, there is little difference between the group speed curves that emerge from the raw and phase-match filtered FTAN images.

A problem occurs with phase-matched filtering, however, when the waveform of interest is narrow-band. In this case, the undispersed signal will possess prominent side lobes that will need to be included in the cleaned collapsed signal cut from surrounding noise. If these side lobes extend broadly enough in time, the cutting procedure may not effectively eliminate contaminating noise. Alternately, if the side lobes are not included in the selected waveform, the redispersed signal will be biased and the dispersion curve will often be undulatory at the long period end of the measurement. For these reasons, phase-matched filtering (i.e. FTAN Steps 3-8) is only recommended for application to broad-band signals.

4.2 Phase speed measurements

By analyzing the envelope function, $|A(t, \omega)|$, the group speed curve, $U(\omega)$, is measured. Phase speed cannot be derived directly from group speed, but the group speed can be computed from phase speed. To see this, let $U = \partial\omega/\partial k$ and $c = \omega/k$ be group and phase speed, respectively, $s_u = U^{-1}$ and $s_c = c^{-1}$ be group and phase slowness, respectively, and k be wavenumber. Then $s_u = \partial k/\partial\omega = \partial(\omega s_c)/\partial\omega$, which gives the following first-order differential equation relating the group and phase slownesses at frequency ω:

$$\frac{\partial s_c}{\partial \omega} + \omega^{-1} s_c = \omega^{-1} s_u. \tag{7}$$

If the phase speed curve $c(\omega)$ is known, the group speed curve $U(\omega)$ can be found directly from this equation. If the group speed curve is known, this differential equation must be solved to find $c(\omega)$, which involves an integration constant that is generally unknown. The solution is

$$s_c(\omega) = \omega^{-1}\left(\int_{\omega_n}^{\omega} s_u(\omega) d\omega + \omega_n s_c^n\right), \tag{8}$$

where the constant of integration has been written in terms of a boundary condition that the phase speed curve is known at some frequency $\omega_n : s_c(\omega_n) = s_c^n$. This is a condition that will generally not apply. Nevertheless, knowledge of the group speed can help to find the phase speed, as we now show.

Measurement of the phase speed curve requires information in addition to the envelope function on which the group speed has been measured. This information derives from the phase of the analytic signal which is approximately composed of a propagation term, an initial source phase and a phase ambiguity term that will be discussed further below. At instantaneous frequency ω, this can be written:

$$\phi(t, \omega) = k\Delta - \omega t - \phi_s - \phi_a, \tag{9}$$

where t is the traveltime, Δ is distance (interstation or epicentral), k is wavenumber, ϕ_s is source phase and ϕ_a is the phase ambiguity term. To proceed, we evaluate the observed phase at the observed group arrival time, $t_u = \Delta/U$, and let $k = \omega s_c$ to find the expression for phase slowness:

$$s_c = s_u + (\omega\Delta)^{-1}(\phi(t_u) + \phi_s + \phi_a), \tag{10}$$

where we now suppress the ω notation for simplicity. The group speed curves, therefore, enter this process by defining the point in time at which the observed phase is evaluated.

Eq. (10) prescribes the phase slowness (and hence the phase speed) curve. Its use, however, depends on knowledge of the initial source phase and the extra phase ambiguity term. In earthquake seismology, ϕ_s is typically computed from Centroid Moment Tensor (CMT) solutions. One of the traditional advantages of studies of group speed over phase speed is that source phase plays a secondary role in group speed (Levshin *et al.* 1999), particularly at short periods. Group speeds, therefore, can be measured at short periods unambiguously using small earthquakes without prior knowledge of the CMT solution. For cross-correlations of ambient noise, however, the situation is considerably easier, as the source phase should be zero: $\phi_s = 0$.

For both earthquake and ambient noise studies, the phase ambiguity term contains a part derived from the 2π ambiguity inherent to any phase spectrum: $\phi_a = 2\pi N$, where $N = 0, \pm1, \pm2, \ldots$. Typically, this ambiguity can be resolved by using a global 3-D model (e.g. Shapiro & Ritzwoller 2002) or phase velocity maps (e.g. Trampert & Woodhouse 1995; Ekstrom *et al.* 1997) to predict phase speed at long periods. The value of N then is chosen to give the closest relation between these predictions and observation. If observations extend to long periods (e.g. greater than 40 s, preferably longer), a global model or observed phase velocity maps may predict phase speed well enough to get N right in most cases. As discussed in Section 5, we recommend making dispersion measurements only up to a period (in s) equal to $\Delta/12$, where Δ is in km. To obtain a 40 s measurement, therefore, requires an interstation spacing of about 500 km. If resolution of the phase ambiguity requires 100 s observations, then an interstation spacing of at least 1200 km is recommended. For ambient noise cross-correlations, if observations are limited to short periods or short interstation distances, the phase ambiguity may not resolve in a straightforward way.

For ambient noise cross-correlations, the phase ambiguity appears to be exacerbated by another factor. Eq. (23) of Snieder (2004) shows that the phase of the cross-correlation between displacement waveforms possesses a $\pi/4$ term that arises from the stationary phase integration (effectively over sources) in the direction transverse to the two stations. The sign of the term depends on the component of the seismometer, positive for the vertical component and negative for the radial component for a Rayleigh wave. The assumption, however, is that sources are homogeneously distributed with azimuth. An inhomogeneous distribution may produce a different phase shift and, because this distribution may vary with frequency, the shift could be frequency dependent. More theoretical work is needed on this problem, but an empirical argument made by Lin, in preparation, 2007 demonstrates that for velocity waveforms the value appears to be $-\pi/4$ for the vertical component. Thus, following Lin, in preparation, 2007 for vertical component ambient noise cross-correlations between velocity waveforms the phase ambiguity term is $\phi_a = 2\pi N - \pi/4$.

In summary, phase-slowness derived from a vertical component ambient noise cross-correlation can be written

$$s_c = s_u + (\omega\Delta)^{-1}(\phi(t_u) + 2\pi N - \pi/4), \tag{11}$$

where $N = 0, \pm1, \pm2, \ldots$. More theoretical work and simulations are needed to determine the uncertainty in the $-\pi/4$ phase shift as well as the possible dependence on frequency and geographical location.

5 QUALITY CONTROL

Because the number of interstation paths grows as the square of the number of stations, the data processing procedure that is applied to ambient noise cross-correlations must be designed to require minimal human interaction. Erroneous dispersion measurements are more likely to arise than if analysts were providing guidance at strategic intervals along the process. Data quality control measures, therefore, must be devised to identify and reject bad measurements and compute quality assurance statistics for the accepted measurements.

First, we have found that a reliable dispersion measurement at period τ requires an interstation spacing (Δ in km) of at least 3 wavelengths (λ): $\Delta > 3\lambda = 3c\tau$ or $\tau < \Delta/3c$. Because phase speed $c \sim 4$ km s^{-1}, for measurements obtained at an interstation spacing of Δ, there is a maximum cut-off period of about $\tau_{\max} = \Delta/12$. We clearly observe the degradation of dispersion measurements at periods greater than about $\tau_{\max}$, at least for group speeds. This imposes a severe constraint on measurements obtained from small regional arrays such as PASSCAL experiments. A broad-band network 500 km in extent, for example, can only produce measurements up to about 40 s period, and that only for the stations across the entire array which is a small subset of the interstation paths. Intermediate and long period measurements most likely will be obtained from the array to surrounding stations, which indicates the importance of permanent (back-bone) stations in the context of regional deployments. At present, we have less experience with phase speed measurements obtained on cross-correlations of ambient noise, so it is possible that the period cut-off may be able to be relaxed for phase speeds.

Second, we need the means to determine the reliability of dispersion measurements that satisfy the period cut-off criterion. One way to estimate reliability is comparison with ground truth. The best case would be when an earthquake has occurred beneath one of the stations. Figs 14 and 15 present an example comparison, using an earthquake that occurred near station PFO (Pinyon Flat, CA, USA). (Date = Oct 31, 2001, $m_b = 5.2$, lat = 33.508, lon = -116.514, depth = 15.2 km). The five paths that are selected are shown in Fig. 14 and comparison between the cross-correlation and the earthquake signals is presented in Fig. 15. To limit the comparison to the period range where both signals are strong, we multiplied the earthquake amplitude spectrum by the cross-correlation amplitude spectrum. This was then used as the amplitude spectrum for both signals. In general, the arrival times of the fundamental Rayleigh waves (the largest amplitude arrivals in each panel) are similar, particularly in light of the source phase shift that affects the earthquake. To compensate for the earthquake radiation pattern, we flip the sign of the earthquake records for stations CMB (Columbia, CA, USA) and LLLB (Lillooet, BC, Canada) which are to the north of station PFO. Also, because the earthquake is south-west of station PFO, the epicentral distances to the stations are greater than the distance between PFO and the other stations. To compensate for this,

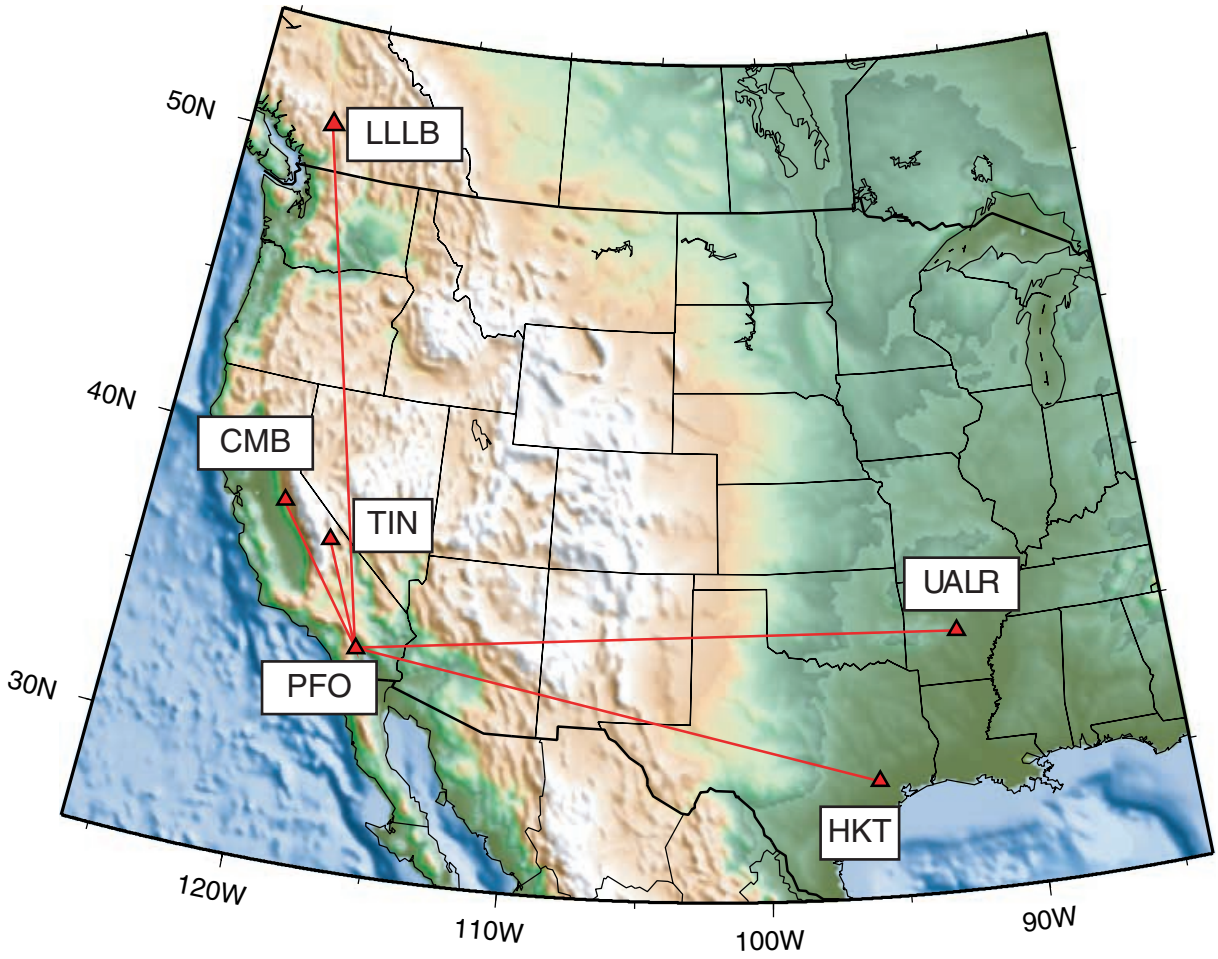

Figure 14. Location map for Fig. 15 showing the ray paths between station PFO and the five other stations.

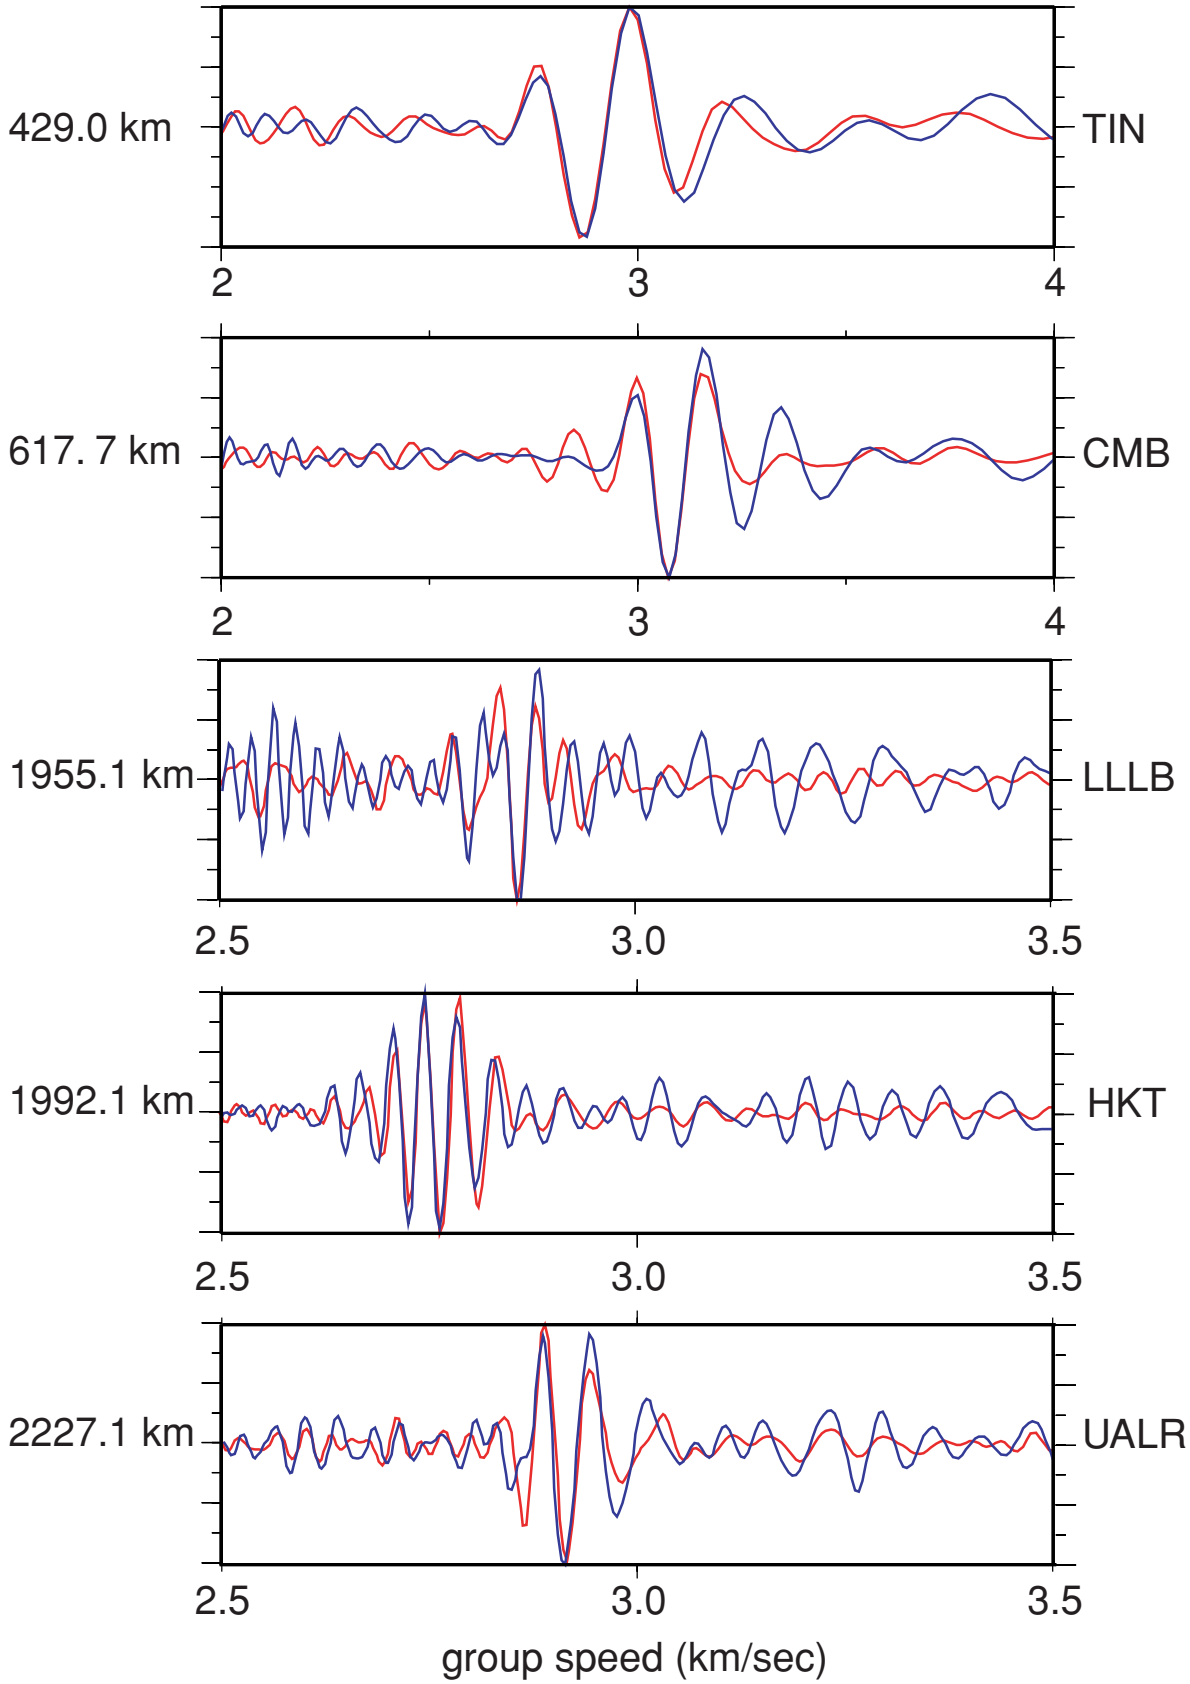

Figure 15. Comparison of symmetric component cross-correlations and earthquake records. Blue lines are 12-month cross-correlations between station PFO in southern California with five other stations around North America. Red lines are earthquake waveforms recorded at the indicated stations following an earthquake near station PFO. The time-series are plotted against group velocity to account for slightly differing path lengths. Station names are indicated at the right-hand side and interstation distances are at the left-hand side. Earthquake records for stations CMB and LLLB are sign flipped.

we plot the recovered signal versus group speed rather than time thus accounting for the different path length. Relocating the earthquake by roughly 6 km to the south-east relative to the PDE location improved the match between the earthquake and cross-correlation signals. Examples such as this give confidence to the ability to interpret cross-correlations of ambient noise in terms of earth structure, similar to the way earthquake signals are interpreted. Coincidences between earthquake and station locations are, however, too rare to be of general use for data selection or uncertainty estimation.

The principal metric on which to base a judgment of the quality of the measurements is stability, the robustness of the measurement to perturbations in the conditions under which it is obtained. The stability of spatially clustered and temporally repeated measurements is particularly useful to identify erroneous measurements and to quantify uncertainties.

Clustering measurements obtained at a particular station from a set of earthquakes located near to one another is commonly used to assess uncertainties in earthquake dispersion measurements (e.g. Ritzwoller & Levshin 1998). A similar cluster analysis can be applied to ambient noise data. For example, Fig. 16 presents a spatial cluster analysis that exploits the high station density in southern California. Numerous measurements between southern California and distant stations were obtained with similar paths (see Fig. 16a). Cross-correlations between the southern California stations and the GSN station ANMO provide one estimate of uncertainty. In this example that there is substantial difference in velocity compared to the CU-Boulder global model (Shapiro & Ritzwoller 2002) at periods below about 35 s for group speeds and 30 s for phase speeds. Measurements between southern California and more distant stations typically are closer to the model prediction. Spatial cluster analyses such as this can be performed when a tight cluster of stations subtends a small angle to a relatively distant station (located many interstation spacings away from the cluster). These conditions typically will not hold for most measurements, although the growth of regional arrays like the Transportable Array component of USArray/EarthScope will help to make this method increasingly applicable. At present, however, cluster analysis provides only an assessment of average uncertainty for long path measurements or a data rejection criterion for a subset of the measurements.

A more useful method to estimate reliability is to assess temporal repeatability. The physical basis for this method is that sources of ambient noise change seasonally and provide different conditions for the measurements. Given the changing conditions, therefore, the repeatability of a measurement is a significant indicator of reliability. This standard is elevated to a high position in our assessment, as we equate seasonal repeatability with measurement uncertainty. It is one of the salutary features of ambient noise dispersion measurements generally that uncertainties can be measured in this way, unlike earthquake derived measurements.

Fig. 17 presents an example of seasonal variability for the GSN station-pair CCM and DWPF. Four disjoint 3-month broad-band cross-correlations are shown in Fig. 17(a) comprising winter, spring, summer and fall months. The long period part of the cross-correlations displays a strong seasonal variability. Group velocity curves from individual 3-month stacks are plotted in Fig. 17(b). Using a year of data, in principle there are twelve 3-month stacks; i.e. January-February-March, February-March-April, . . . , December-January-February. Only curves from the 3-month stacks in which spectral SNR >10 at all periods are shown. 10 of the 12 stacks satisfied this criterion. Overplotted with the red line is the group speed curve measured for the 12-month stack. It appears in the middle of the shorter measurements and is smoother than most of the 3-month stacks. This indicates that the use of variability among the 3-month stacks to estimate the uncertainty in the dispersion measurement for the 12-month time-series is conservative. The predicted curve from the 3-D model of Shapiro & Ritzwoller (2002) is also overplotted in green.

In earlier applications of the data processing procedure described herein (Yang *et al.* 2007; Lin *et al.* 2007a), dispersion measurements are obtained on 12-months of data. To estimate uncertainties in these measurements, we also measure dispersion on all sequential 3-month stacks if SNR exceeds some threshold. The standard deviation is computed if a sufficient number of the 3-month stacks exceeds the SNR criterion. In the high ambient noise environment of New Zealand, Lin *et al.* (2007a) required seven of the 3-month stacks to have SNR >10. Yang *et al.* (2007), working with the lower ambient noise conditions that prevail across most of Europe yielded lower SNR values of the resulting cross-correlations. They were forced to loosen this criterion (four 3-month stacks with SNR >7). Both studies rejected any measurement for which an uncertainty measurement could not be determined. Yang *et al.* (2007), in particular, rejected many measurements because uncertainty could not be determined even with the loosened criteria. They argued, therefore, that at least across much of Europe, two years of data would be preferable to one in order to estimate uncertainties and reject far fewer measurements. Presumably this would be true for most other continents around the world.

If a seismic station is operated or a pair of stations are run simultaneously only for a short period of time, however, acquisition of two years of data may be out of the question and temporal subsetting to estimate uncertainties may not be feasible. Temporal overlap between neighbouring deployments of stations also may not be long enough to estimate uncertainties based on temporal variability. In this case, SNR measurements can provide a useful proxy for uncertainties. An example is shown in Fig. 18. In these figures, the average standard deviation measured from the temporal variability of cross-correlations of ambient noise observed over one-year is plotted as a function of spectral SNR. The cross-correlations are obtained on more than 200 stations across the US and southern Canada from the year 2004. Results at 10 s period (green circles) and 20 s period (red triangles) are shown, and are segregated into two interstation distance ranges, 1000–2000 km and distances greater than 2000 km. At both periods there is a clear linear relation between standard deviation and spectral SNR for $10 < \mathrm{SNR} < 40$. For $\mathrm{SNR} < 10$, the standard deviation increases rapidly and non-linearly. These curves illustrate that SNR may provide a useful proxy for measurement error if $\mathrm{SNR} > 10$. In addition, because, as Fig. 12 shows, SNR is also related to the number of weeks in the stack, the number of weeks stacked is related to the expected measurement uncertainty. For example, inspection of Fig. 12 shows that at 10 s period, a SNR of 10 is expected after 4-weeks of observation. Fig. 18(a) reveals that a SNR of 10 relates to a measurement error of about 55 m s^{-1}. Thus, four-weeks of observation (in North America) is expected to give a measurement error of about 55 m s^{-1}. Increasing the observing length to 20 weeks at 10 s period is expected to increase the SNR to about 20 and the measurement error is expected to reduce to about 45 m s^{-1}. Observing still longer to 60 weeks is expected, on average, to increase SNR to about 30 and reduce error to about 35 m s^{-1}. Thus, although it is tempting to stack data indefinitely, the power-law dependence of SNR on time-series length generates diminishing returns in reducing measurement errors after SNR is high enough to provide a reasonable dispersion measurement. Continued observation past this point may best be used to measure temporal variability directly, which is preferable over the use of proxy curves.

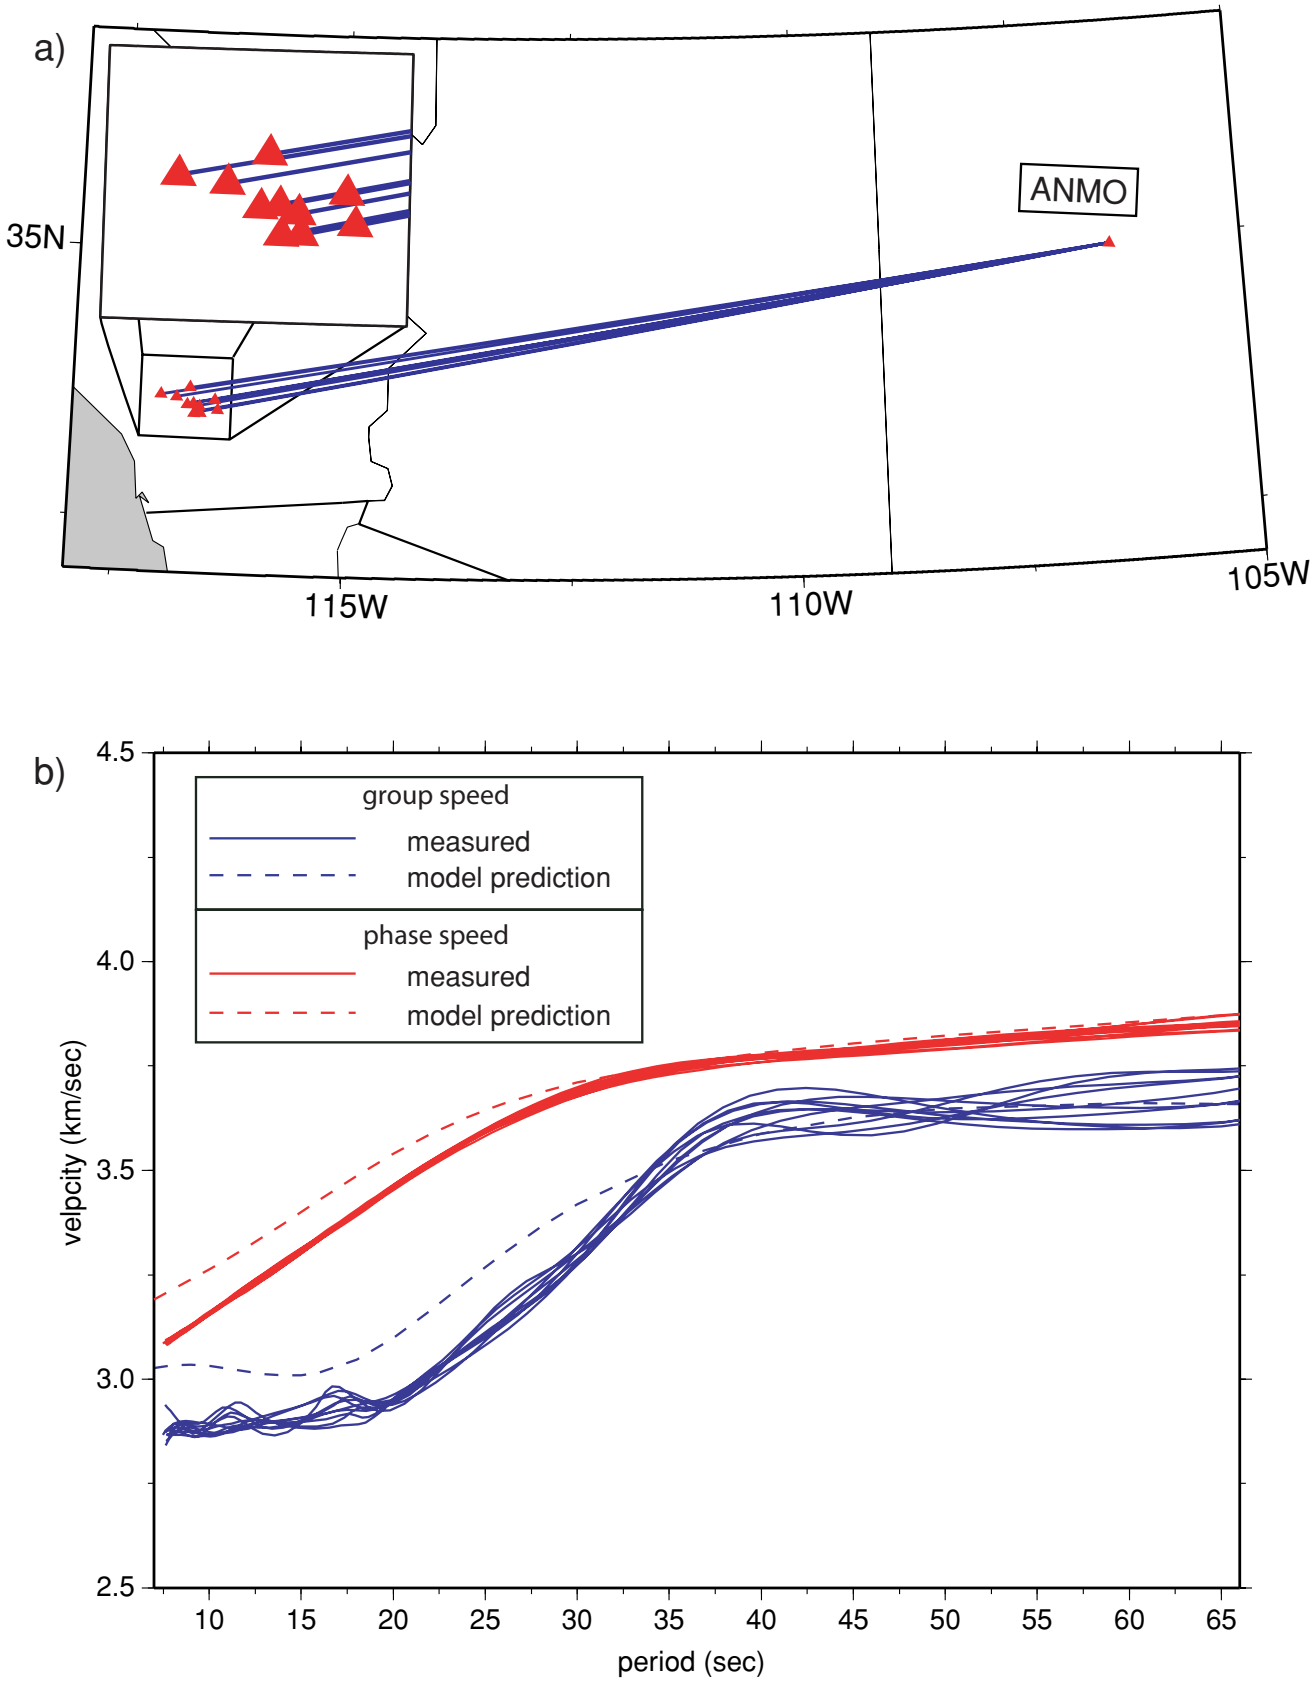

Figure 16. Example assessment of the spatial variability of dispersion measurements using a cluster of 10 stations in southern California. (a) The cluster of 10 paths used in this analysis with a detail plot of the stations used in southern California. (b) Measurements shown with solid lines are from 12-month stacks observed between station ANMO and the southern California cluster of 10 stations and the dashed lines are the predictions from the 3-D model of Shapiro & Ritzwoller (2002).

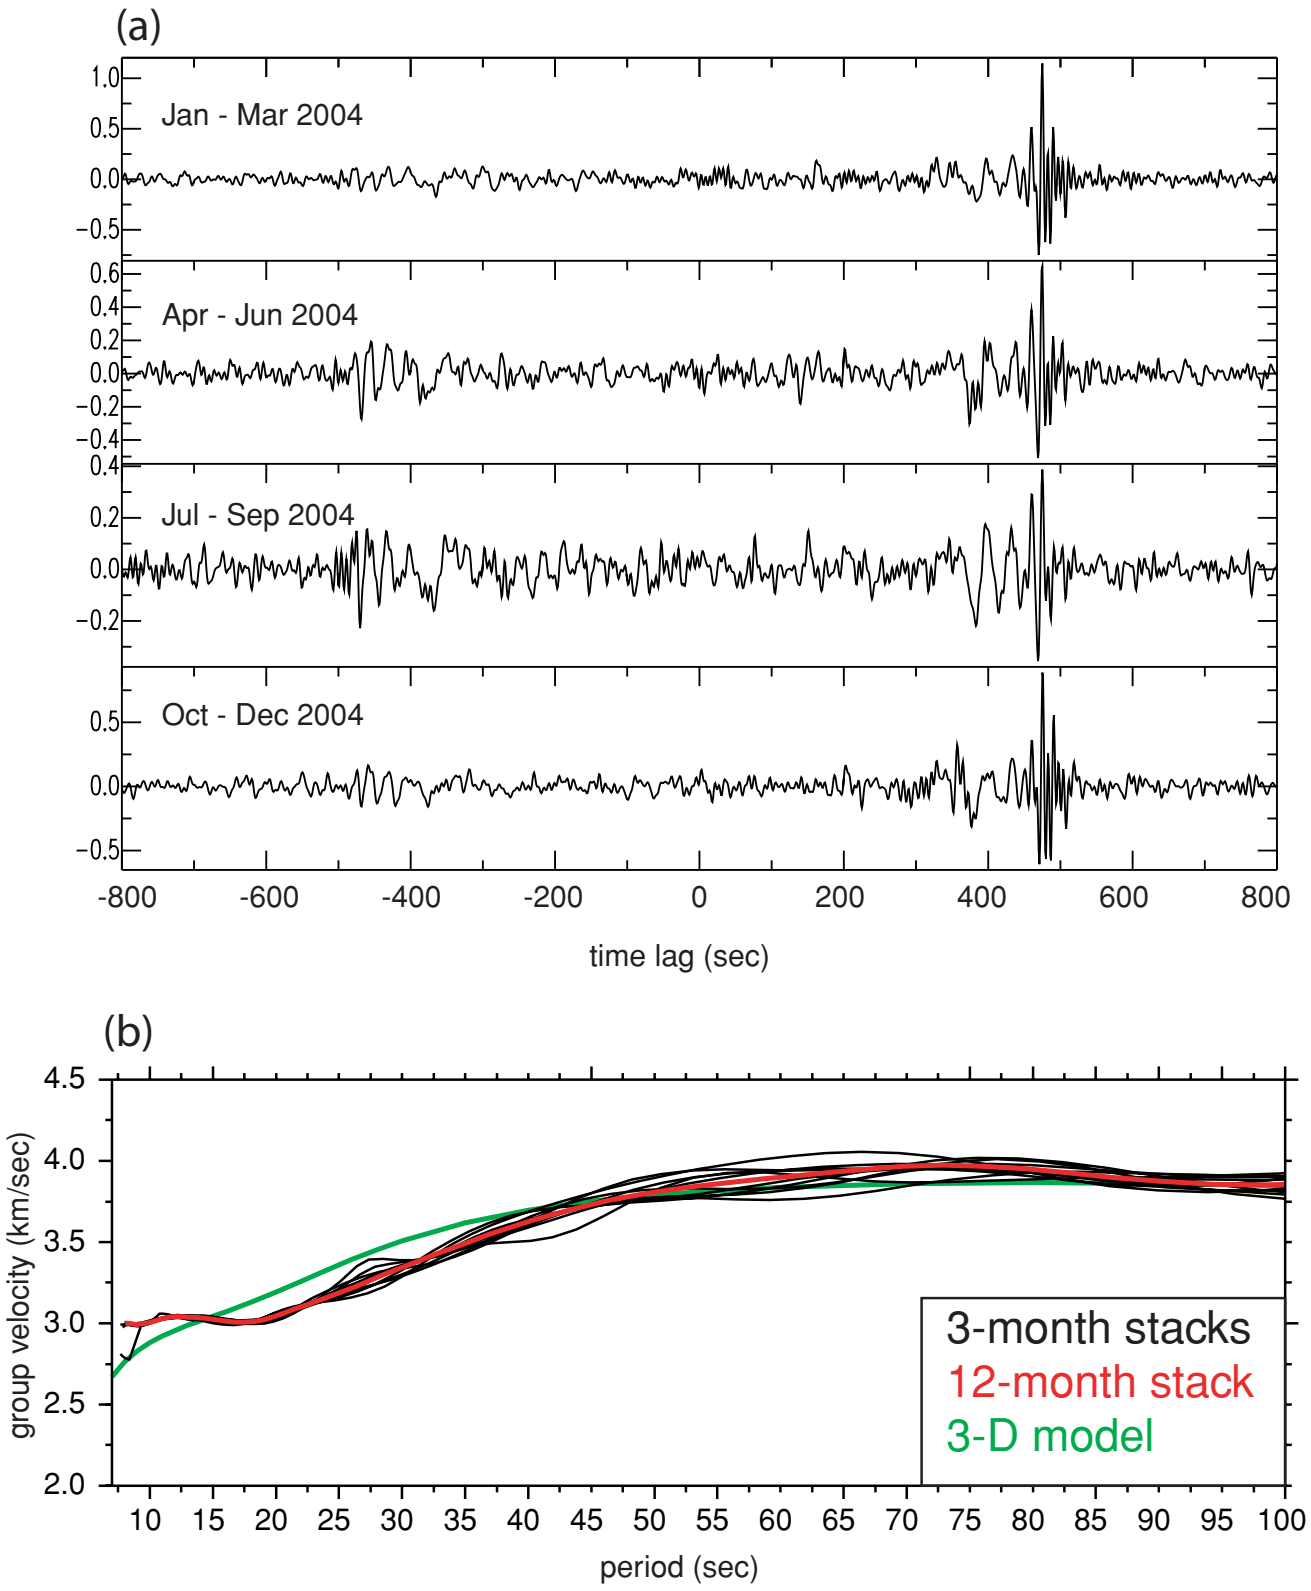

Figure 17. Example assessment of temporal variability of dispersion measurements. (a) Four disjoint broad-band (5–150 s) three-month stacks are shown for the station-pair CCM-DWPF. Arrivals at positive correlation lag are for waves travelling from the northwest, from CCM to DWPF and negative correlation lag corresponds to waves arriving from the southeast, from DWPF to CCM. (b) Group speed measurements obtained on the symmetric-component from ten 3 months are presented versus period as the black curves. The measurement for the 12-month stack is indicated by the red line and the green line is the prediction from the 3-D model of Shapiro & Ritzwoller (2002).

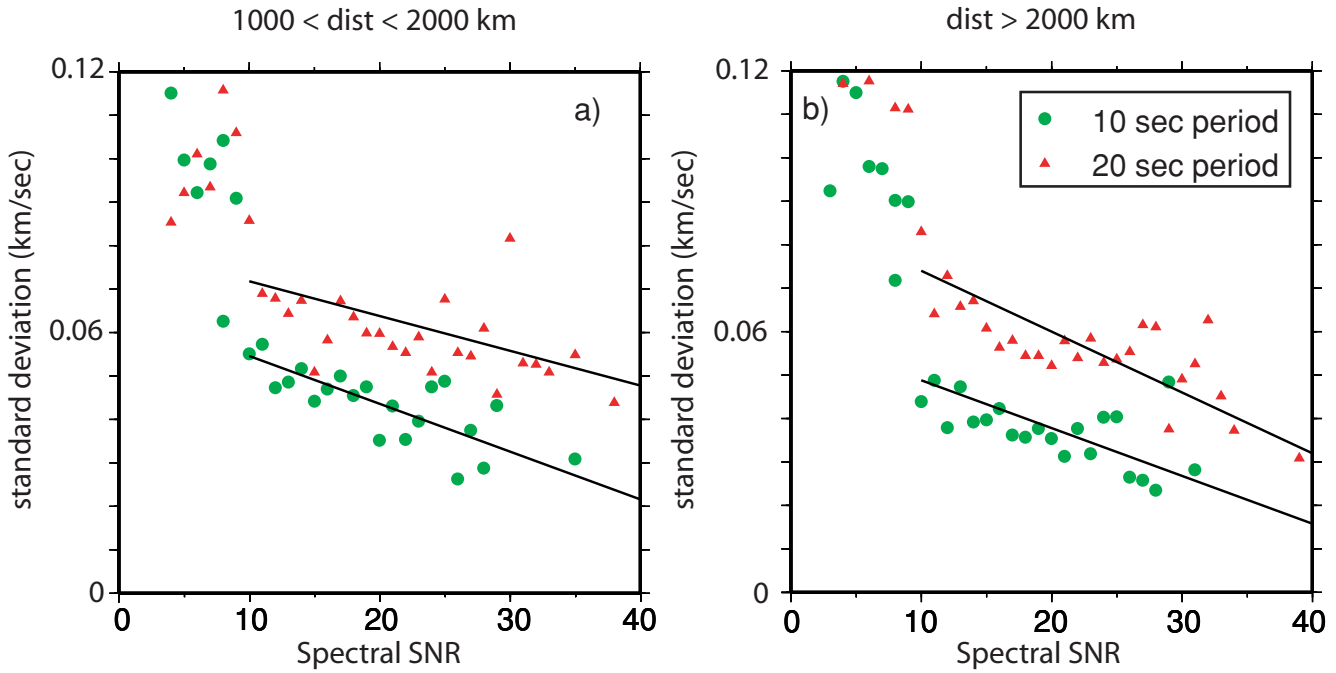

Figure 18. SNR as proxy for measurement uncertainty. Average measurement error, estimated from seasonal variability for a year of data using more than 200 stations across North America, is plotted versus spectral SNR. Two distance ranges are shown and two periods (red triangles—20 s; green circles—10 s). Below SNR $\sim$ 10, measurements become unreliable. Black lines show the best-fit linear trends for 10 < SNR < 40.

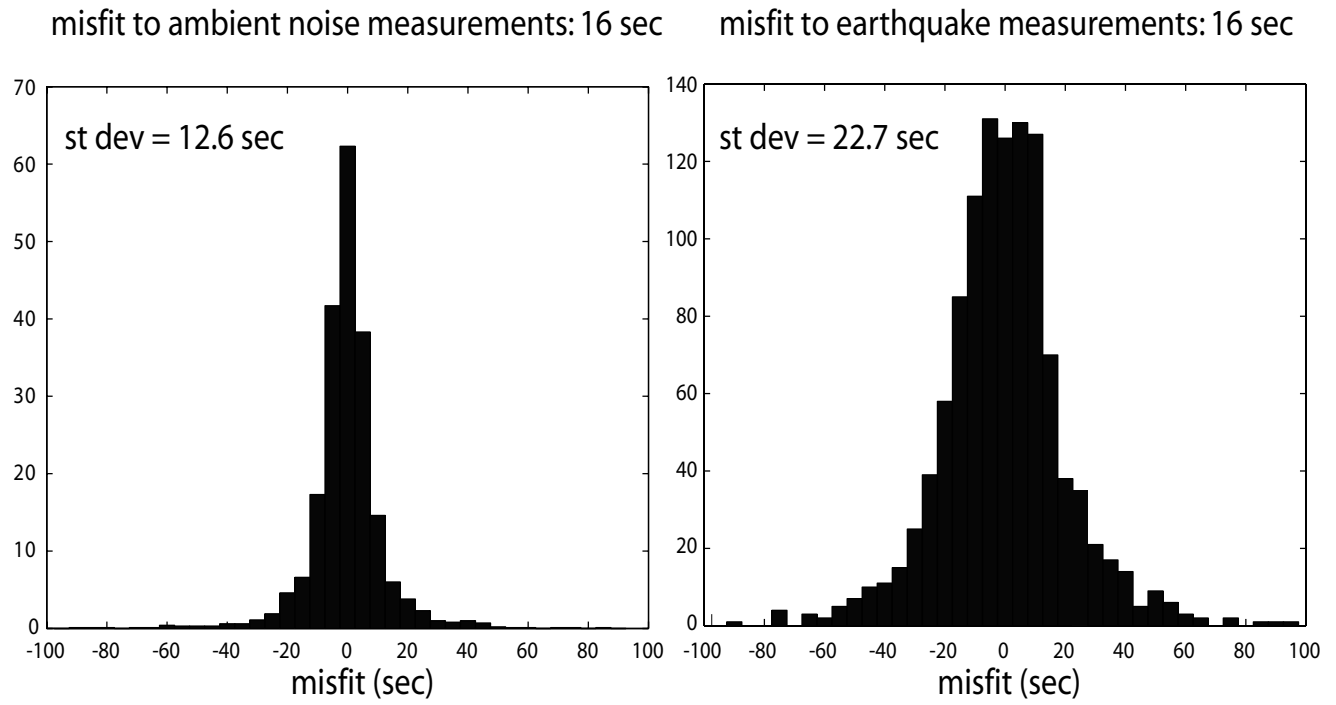

Figure 19. Measurement coherence. Histograms of misfit for ambient noise group time measurements (left-hand side) and earthquake group time measurements (right-hand side) at 16 s period across Europe. Misfit is calculated from the corresponding smooth group speed maps inverted from the ambient noise and earthquake data, respectively. The standard deviation is indicated at the top left-hand side in each panel.

The discussion in the previous paragraph is relevant to the design of seismic experiments to use ambient noise cross-correlation. It is, however, intended to be more illustrative than definitive, and more work is needed to understand the distributions of the quantities in order to produce better proxy curves and guide their use.

Third, we seek measurements that cohere as a whole; that is, that agree with other accepted measurements. This condition can be tested tomographically. Measurements that can be fit with a smooth tomographic map are said to agree with one another. Yang *et al.* (2007) presents a detailed discussion of the application of this

criterion across Europe. He finds that, on average, dispersion measurements that derive from ambient noise tomography can be fit better than those that derive from earthquake data. Moreover, the distribution of misfit is tight. While erroneous measurements do pass the previous selection criteria, they are small in number. An example comparison between the misfit histograms of ambient noise and earthquake derived group speed measurements across Europe is shown in Fig. 19.

6 SUMMARY AND CONCLUSIONS

The data processing procedure for extracting estimated Rayleigh wave Green functions from ambient seismic noise has now largely stabilized and is sufficiently well evolved to describe without fear of radical modification in the near future. The procedures that are described herein have been designed not only to deliver reliable measurements, but to be flexible and applicable to a wide variety of observational settings. The procedure divides into four phases: (1) single-station data preparation, (2) cross-correlation and stacking to a desired time-series length, (3) dispersion measurement and (4) quality control. Because the number of cross-correlations grows as the square of the number of stations, the procedure we describe is entirely automated.

The principal step in single-station data preparation is 'temporal normalization' which is designed to ameliorate the contamination of the ambient noise signals by earthquakes, instrument irregularities, and non-stationary noise sources near to stations (such as passing storms and high local sea heights). We advocate the use of 'running-absolute-mean' normalization, which is an effective method that allows for tuning to regional earthquake conditions. In addition, spectral whitening is advisable prior to cross-correlation in order to minimize contamination by the 26 s Gulf of Guinea resonance and to broaden the measurement band.

The use of long time-series helps to optimize the SNR, which is anti-correlated with measurement error. SNR displays a power law dependence on time-series length, with the most rapid emergence of signals from noise in the microseismic ($<$20 s period) and 'Earth hum' ($>$50 s) bands. The greatest challenge for ambient noise tomography, therefore, lies between about 30 and 60 s period.

Automated dispersion measurement is performed with a modification of traditional FTAN (e.g. Levshin *et al.* 1992). We have described methods that measure group velocity curves reliably with and without phase-matched filtering. Phase-matched filters are helpful to extract the estimated Green function from adjacent contaminating signals, if they exist. However, phase-matching filtering works best on broad-band waveforms. For signals that are significantly band-limited, it would be best to forego phase-matched filtering. Reliable group velocity measurements, on average, require a SNR > 10.

Experience reveals that reliable group velocity measurements require an interstation spacing, Δ in km, of at least three wavelengths. This creates a period cut-off $\tau_{\max} \sim \Delta/12$. At periods longer than this, the interstation spacing will be less than three wavelengths, and the measurements are significantly more likely to be unreliable. This condition can impose a stringent constraint on the use of ambient noise tomography with data from local or regional arrays.

Phase velocity curves also emerge naturally from the automated FTAN and preliminary results indicate that the curves are stable and smooth as long as SNR is above a threshold value of about 10. For small interstation spacings where only short period phase velocities can be measured, the 2π phase-ambiguity may be difficult to resolve unless station density is high enough to exploit the observed phase 'move-out'. Further work is also needed to determine if the period cut-off can be relaxed for phase velocity measurements.

Within the context of an automated data processing procedure, data quality control measures are particularly important to identify and reject bad measurements and compute quality assurance statistics for the accepted measurements. The principal metric on which to base a judgment of quality is stability, the robustness of the measurement to perturbations in the conditions under which it is obtained. Temporal repeatability, in particular, is a significant indicator of reliability. The physical basis for this method is that sources of ambient noise change seasonally and provide different conditions for the measurements. This standard is elevated to a high position in our assessment, as we equate seasonal repeatability with measurement uncertainty. It is one of the commendable features of ambient noise dispersion measurements generally that uncertainties can be measured, unlike earthquake derived measurements. Although one year of data is sufficient to estimate uncertainties through temporal repeatability in some cases, two years of data are preferable.

Acquisition of one to two years of data may be out of the question in many circumstances, so that temporal subsetting to estimate uncertainties may not be feasible. In this case, SNR measurements can provide a meaningful proxy for uncertainties. Such proxy curves relating measurement uncertainty to SNR (e.g. Fig. 18) can be used with information about the emergence rate of the signal (e.g. Fig. 10) to help design experiments that seek to perform ambient noise tomography. The results in Fig. 18 are are not yet definitive, and more work is needed to understand the statistical distributions of the quantities in order to produce better proxy curves and guide their use. Nevertheless, we believe that this approach promises to provide approximate uncertainty estimates for experiments that are too short to exploit temporal repeatability.

The data processing procedures that have been developed for ambient noise tomography currently have a history shorter than three years and need to continue to develop. In particular, efforts are needed to tune the method further for phase velocities (e.g. understanding potential phase ambiguities related to source distribution) and Love waves. We also believe that work on proxy curves in which SNR (or time-series length) is used to infer an expected measurement uncertainty is a fertile area for future research.

ACKNOWLEDGMENTS

The authors thank Roel Snieder for a discussion of the $\pi/4$ phase term that appears in eq. (11). All of the data used in this research were downloaded either from the IRIS Data Management Center, the European Orfeus Data Center, or the New Zealand GeoNet Data Center. This research was supported by a contract from the US Department of Energy, DE-FC52-2005NA26607 and two grants from the US National Science Foundation, EAR-0408228 and EAR-0450082.

REFERENCES

Bensen, G.D., Ritzwoller, M.H., Shapiro, N.M. & Levshin, A.L., 2005. Extending ambient noise surface wave tomography to continental scales: application across the United States, *EOS, Trans. Am. geophys. Un.,* **86**(52), Fall Meeting Suppl., Abstract S31A-0274.

Bracewell, R.N., 1978. *The Fourier Transform and its Applications*, McGraw-Hill, 2nd edition, New York.

Campillo, M. & Paul, A., 2003. Long-range correlations in the diffuse seismic coda, *Science,* **299,** 547–549.

Cho, K.H., Hermann, R.B., Ammon, C.J. & Lee, K., 2006. Imaging the crust of the Korean peninsula by surface wave tomography, *Bull. seism. Soc. Am.,* in press.

Crawford, W.C., Stephen, R.A. & Bolmer, S.T., 2006. A second look at low-frequency marine vertical seismometer data quality at the OSN-1 stie off of Hawaii for seafloor, buried and borehole emplacements, *Bull. seism. Soc. Am.,* **96,** 1952–1960.

Derode, A., Larose, E., Tanter, M., de Rosny, J., Tourim, A., Campillo, M. & Fink, M., 2003. Recovering the Green's function from field-field correlations in an open scattering medium, *J. acoust. Soc. Am.,* **113,** 2973–2976.

Dziewonski, A.M., Bloch, S. & Landisman, M., 1969. A technique for the analysis of transient seismic signals, *Bull. seism. Soc. Am.,* **59,** 427–444.

Ekstrom, G., Tromp, J. & Larson, E.W.F., 1997. Measurements and global models of surface wave propagation, *J. geophys. Res.,* **102,** 8137–8157.

Herrin, E.E. & Goforth, T.T., 1977. Phase-matched filters: Application to the study of Rayleigh Waves, *Bull. seism. Soc. Am.,* **67,** 1259–1275.

Holcomb, L.G., 1998. Spectral structure in the Earth's microseismic background between 20 and 40 s, *Bull. seism. Soc. Am.,* **88,** 744–757.

Larose, E., Derode, A., Campillo, M. & Fink, M., 2004. Imaging from one-bit correlations of wideband diffuse wavefields, *J. Appl. Phys.,* **95,** 8393–8399.

Larose, E., Derode, A., Corennec, D., Margerin, L. & Campillo, M., 2005. Passive retrieval of Rayleigh waves in disoredered elastic media, *Phys. Rev. E.,* **72,** 046607, doi:10.113/PhysRevE.72.046607.

Levshin, A.L. & Ritzwoller, M.H., 2001. Automated detection, extraction, and measurement of regional surface waves, *Pure appl. Geophys.,* **158**(8), 1531–1545.

Levshin, A.L., Pisarenko, V.F. & Pogrebinsky, G.A., 1972. On a frequency-time analysis of oscillations, *Ann. Geophys.,* **28,** 211–218.

Levshin, A.L., Yanovskaya, T.B., Lander, A.V., Bukchin, B.G., Barmin, M.P., Ratnikova, L.I. & Its, E.N., 1989. *Seismic Surface Waves in a Laterally Inhomogeneous Earth*, ed. Keilis-Borok, V.I., Kluwer, Norwell, Mass.

Levshin, A.L., Ratnikova, L. & Berger, J., 1992. Peculiarities of surface wave propagation across central Eurasia, *Bull. seism. Soc. Am.,* **82,** 2464–2493.

Levshin, A.L., Ritzwoller, M.H. & Resovsky, J.S., 1999. Source effects on surface wave group travel times and group velocity maps, *Phys. Earth planet. Int.,* **115,** 293–312.

Lin, F.-C., Ritzwoller, M.H. & Shapiro, N.M., 2006. Is ambient noise tomography across ocean basins possible?, *Geophys. Res. Lett.,* **33,** L14304, doi:10.1029/2006GL026610.

Lin, F.-C., Ritzwoller, M.H., Townend, J., Savage, M. & Bannister, S., 2007. Ambient noise Rayleigh wave tomography of New Zealand, *Geophys. J. Int.,* in press.

Moschetti, M.P., Ritzwoller, M.H. & Shapiro, N.M., 2007. Ambient noise tomography from the first two years of the USArray Transportable Array: group speeds in the western US, *Geophys. J. Int.,* submitted.

Paul, A., Campillo, M., Margerin, L., Larose, E. & Derode, A., 2005. Empirical synthesis of time-asymmetrical Green function from the correlation of coda waves, *J. geophys. Res.,* **110,** doi:10.1039/2004JB003521.

Rhie, J. & Romanowicz, B., 2004. Excitation of earth's incessant free oscillations by Atmosphere-Ocean-Seafloor coupling, *Nature,* **431,** 552–556.

Ritzwoller, M.H. & Levshin, A.L., 1998. Surface wave tomography of Eurasia: group velocities, *J. geophys. Res.,* **103,** 4839–4878.

Roux, P., Sabra, K.G., Gerstoft, P., Kuperman, W.A. & Fehler, M.C., 2005. *P*-waves from cross-correlation of seismic noise, *Geophys. Res. Lett.,* **32,** L19393, doi:10.1029/2005GL023803.

Russell, D.W., Herrman, R.B. & Hwang, H., 1988. Application of frequency-variable filters to surface wave amplitude analysis, *Bull. seism. Soc. Am.,* **78,** 339–354.

Sabra, K.G., Gerstoft, P., Roux, P., Kuperman, W.A. & Fehler, M.C., 2005a. Extracting time-domain Green's function estimates from ambient seismic noise, *Geophys. Res. Lett.,* **32,** L03310, doi:10.1029/2004GL021862.

Sabra, K.G., Gerstoft, P., Roux, P., Kuperman, W.A. & Fehler, M.C., 2005b. Surface wave tomography from microseism in southern California, *Geophys. Res. Lett.,* **32,** L14311, doi:10.1029/2005GL023155.

Shapiro, N.M. & Campillo, M., 2004. Emergence of broadband Rayleigh waves from correlations of the ambient seismic noise, *Geophys. Res. Lett.,* **31,** L07614, doi:10.1029/2004GL019491.

Shapiro, N.M. & Ritzwoller, M.H., 2002. Monte-Carlo inversion for a global shear velocity model of the crust and upper mantle, *Geophys. J. Int.,* **151,** 88–105.

Shapiro, N.M., Campillo, M., Stehly, L. & Ritzwoller, M.H., 2005. High resolution surface wave tomography from ambient seismic noise, *Science,* **307,** 1615–1618.

Shapiro, N.M., Ritzwoller, M.H. & Bensen, G.D., 2006. Source location of the 26 sec microseism from cross correlations of ambient seismic noise, *Geophys. Res. Lett.,* **33,** L18310, doi:10.1029/2006GL027010.

Snieder, R., 2004. Extracting the Green's function from the correlation of coda waves: a derivation based on stationary phase, *Phys. Rev. E,* **69,** 046610.

Trampert, J. & Woodhouse, J.H., 1995. Global phase velocity maps of Love and Rayleigh waves between 40 and 150 s, *Geophys. J. Int.,* **122,** 675–690.

Wapenaar, K., 2004. Retrieving the elastodynamic Green's function of an arbitrary inhomogeneous medium by cross correlation, *Phys. Rev. Lett.,* **93,** 254301, doi:10.1103/PhysRevLett.93.254301.

Weaver, R.L. & Lobkis, O.I., 2001a, Ultrasonics without a source: Thermal fluctuation correlation at MHz frequencies, *Phys. Rev. Lett.,* **87,** doi:10.1103/PhysRevLett.87.134301.

Weaver, R.L. & Lobkis, O.I., 2001b. On the emergence of the Green's function in the correlations of a diffuse field, *J. acoust. Soc. Am.,* **110,** 3011–3017.

Weaver, R.L. & Lobkis, O.I., 2004. Diffuse fields in open systems and the emergence of the Green's function, *J. acoust. Soc. Am.,* **116,** 2731–2734.

Yang, Y., Ritzwoller, M.H., Levshin, A.L. & Shapiro, N.M., 2007. Ambient noise Rayleigh wave tomography across Europe, *Geophys. J. Int.,* **168,** 259–274.

Yao, H., van der Hilst, R.D. & de Hoop, M.V., 2006. Surface-wave tomography in SE Tibet from ambient seismic noise and two-station analysis: I.—Phase velocity maps, *Geophys. J. Int.,* **166,** 732–744, doi: 10.1111/j.1365–246X.2006.03028.x.