Determining phase and amplitude of a digitized sinusoidal signal

As described in the main text, we used two different methods to determine the phase and amplitude of the recorded sinusoidal voltages presented in this study. For each method, 50 cycles of the recorded signal were provided for a measurement. The first method uses the nlinfit function in matlab to fit a sinusoid to the data. This fitting process is very sensitive, as very small changes in parameters can lead to very large changes in the residuals of the fit. In our experience, fitting a sinusoid in this manner is only possible because we provide the intended frequency of the sinusoid as an initial guess and this initial guess is always extremely close to the frequency of the best fitting sinusoid. Even with this accurate initial guess, on a small minority of our recordings a good fit is still not obtained from the procedure. This happened for less than 5% of our recordings.

There appears to be no obvious difference between recordings that are and are not well fit by this algorithm as demonstrated by Figure SM1 which shows two sample recordings made at 90 Hz using the same electrode. The recording made in the parallel configuration shown in panel A happens to be poorly fit with this algorithm, but the recording in the series configuration shown in panel B was fit well, though there is no obvious difference in the underlying recorded data of the two recordings. Moreover, the occurrence of recordings with successful and unsuccessful fits appeared to be random, as there were no cases in which several recordings made in immediate succession provided bad fits with this method. The recordings that were not well fit with this method always showed very large residuals resulting from a fit with a sinusoid that was much smaller than and out of phase with the actual sinusoid. As a result, a test of the residuals at a moderate value can easily automatically distinguish between recordings that were or were not well fit with this method.

For the recordings that were not well fit with the sine fit method, we used an alternate method involving the Hilbert transform of the data, with results of this method also shown for the example recordings in Figure SM1. With this method, the amplitude is defined as the average magnitude of the Hilbert transform and the phase is defined as the phase of the first sample resulting from a linear fit of the complex phase of the Hilbert transform. As can be seen from the Figure SM1, this method was much more robust. We found this provided reasonable measurements for all of our recordings and thus used the values determined from this algorithm for all recordings that were not well fit by the sine fit algorithm. For recordings in which the sine fit algorithm fit the data well, both methods gave similar answers but the residuals of the sine fit method were slightly smaller. We thus opted to use the values for the sine fit method for these cases. Amplitude ratios and phase comparisons between the recorded and actual signals were always made using the same method for each signal. So for a given recording with both the actual and recorded signals being fit by these methods, if the sine fit algorithm worked for one but not the other, the values from the Hilbert algorithm were used for both signals to make this comparison.

The code for the sine fit and Hilbert method are provided in the Supplementary Code available online. Each method is called by AddToSavedTransFxn.m, which calls the subfunctions fitsine.m and GetHilbVals.m.

We also investigated a few other methods to compare to these. We used open source software downloaded form http://www.chronux.org, which is provided from the Mitra Lab at Cold Spring Harbor Laboratory. We used the function fitlinesc which also fits a sine wave to continuous data using an alternate method to the nlinfit algorithm we previously employed. We used the code available in version 1.50 of this package, which was the newest available version of this software at the time that this article was written. We found that the values obtained from this were nowhere near the correct values, as shown in Figure SM1. In going through this code we believe that there is an oversight in this and most other programs in this package as the power and amplitude values reported from this software should be doubled from what is reported as the operations of the code are performed on values for the half-spectrum of the input data signal without the proper adjustment done later in the code to provide the values for the entire signal. Even accounting for that oversight by doubling the value reported by the chronux code, the resulting amplitudes would still be too small and the phases are incorrect, suggesting this method does not accurately fit a sine wave to sinusoidal data.

We also tried a simple zero-padded FFT measurement, which involved taking a zero-padded FFT and reporting the amplitude, phase and frequency of the maximal frequency component. The length of the zero-padded FFT was taken as the zero-padding factor times 2 to next power of 2 greater than the length of the signal. For the example data in Figure SM1, the results from this method were reasonably accurate. We also ran simulations to compare the accuracy of this with the Hilbert method, with the results shown in Figure SM2. In each iteration, we generated 50 cycles of a sinusoid with a frequency near 90 Hz with a random phase sampled at 1000 Hz. The precise frequency was set to 64 times the square root of 2, which is approximately 90.5 Hz. This was done to ensure that the frequency of the simulated sinusoid did not exactly coincide with the digital sampling to better approximate a real recorded signal and elucidate the problems of the zero-padded FFT. The amplitude of the waveform was arbitrarily set to 1, and to this noise was added which was drawn from a standard Gaussian distribution multiplied by different noise factors, 0.1 or 0.4. For the two different noise factors, 1000 iterations were run and three different phase and amplitude measurements were-compared: the Hilbert method and the zero-padded FFT method with a padding factor of 4 or 20. The number of points in the zero-padded FFT was the zero-padding factor times the next power of 2 greater than the length of the signal. Figure SM2 shows that both FFT measurements were rather intolerant to the added white noise, but this was not the case with Hilbert method. In all cases the amplitude measurements were quite accurate, with the Hilbert method showing somewhat larger errors, though the errors were still quite small. For the phase measurement however, there appears to be a consistent bias in the reported values using the FFT method. This bias seems to decrease as the zero-padding factor increases. For the high signal to noise ratio data, which most closely resembles the bulk of the data in this study, the Hilbert method is preferred for its improved accuracy in measuring the phase of the signal.

We also performed the same simulations using a sampling rate of 20000 Hz which was the sampling frequency employed for the actual data recorded in the rest of the study, and the results were similar (data not shown). However, the precise ratio between the sampling frequency and the frequency of the signal is unimportant for the purpose of relating to this study, as in this study we investigated frequencies spanning the range from very low frequencies to those approaching the Nyquist frequency.

A note on the signals in Figure 9

The signals used to obtain the data represented by the dotted lines in figure 9 had amplitude envelopes varying more slowly over time than the examples shown here for the given frequencies. We found this gave a slightly better match with the derivative measures of group delay over frequency ranges where the group delay was rapidly changing across frequencies, which we attributed to a better frequency localization of the underlying signal. The signals presented in Figure 12 were selected to more clearly visually present the delay that can be detected with this technique.

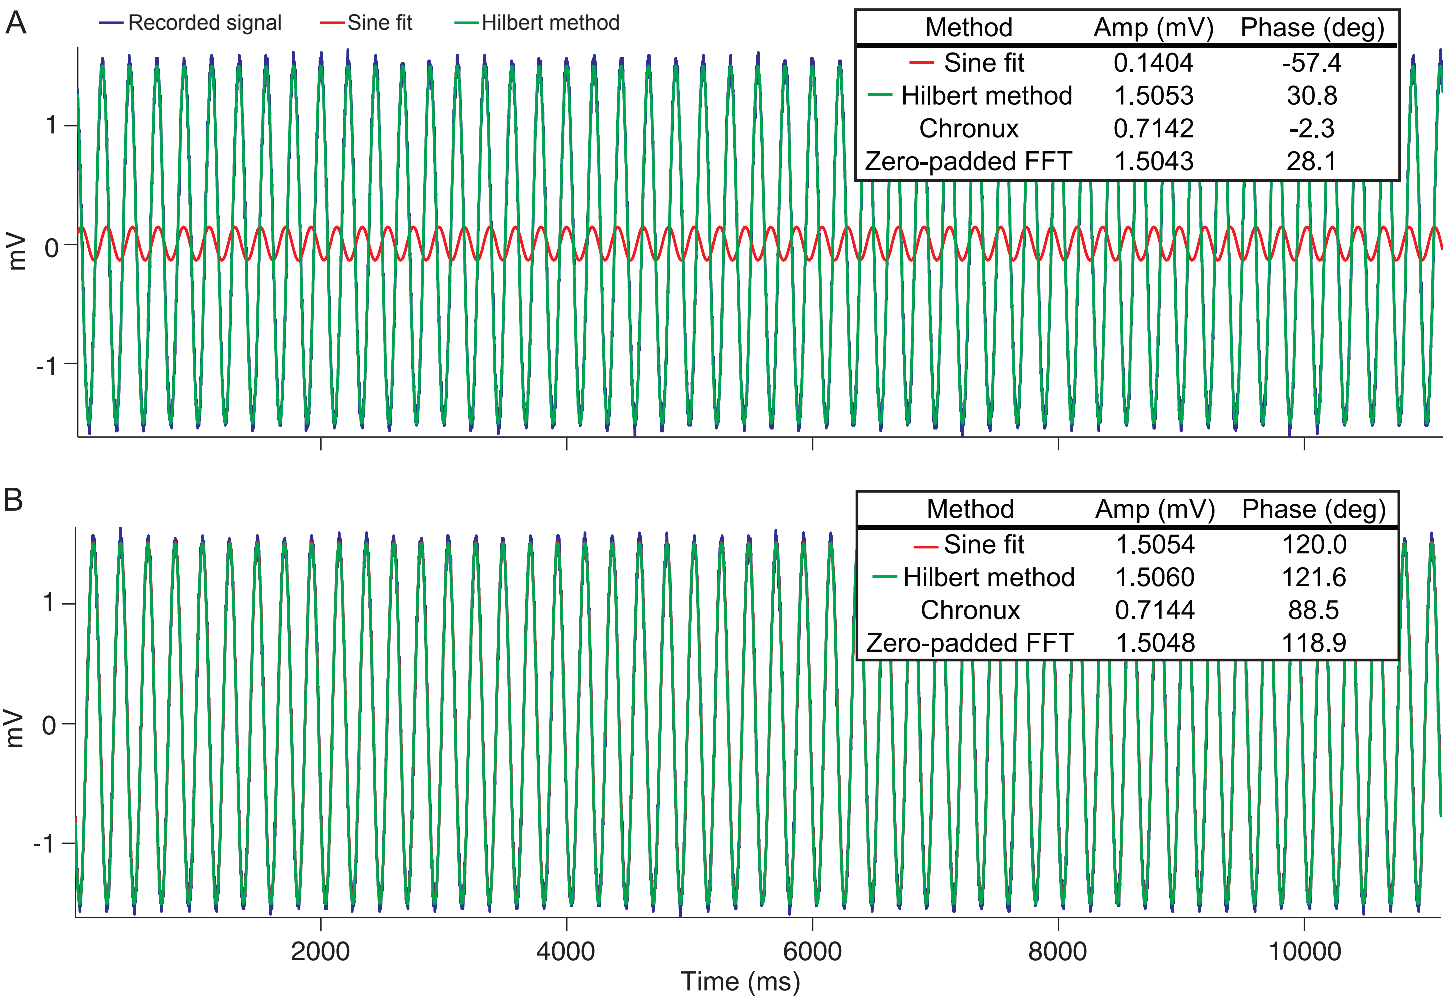

Figure SM1. Example results of raw amplitude and phase measurements. Both panels show two sample recordings done on one sample medium impedance glass electrode in dilute saline using a 90 Hz sine wave. Panel A was recorded in the parallel configuration and Panel B was recorded in the series configuration. In each panel the recorded actual signal sent to the head-stage is shown in green, the sine wave resulting from the sine fit method is shown in red, and the equivalent sine wave resulting from the Hilbert method is shown in green. The sine fit method poorly fits one recording, but fits the other well despite there being no obvious differences between the recordings. The table inset on top of each panel shows the amplitude and phase results from each of these two methods as well two methods not pictured in the plots: the result from fitlinesc.m from the Chronux open source software package, and an FFT with a zero-padding factor of 4.

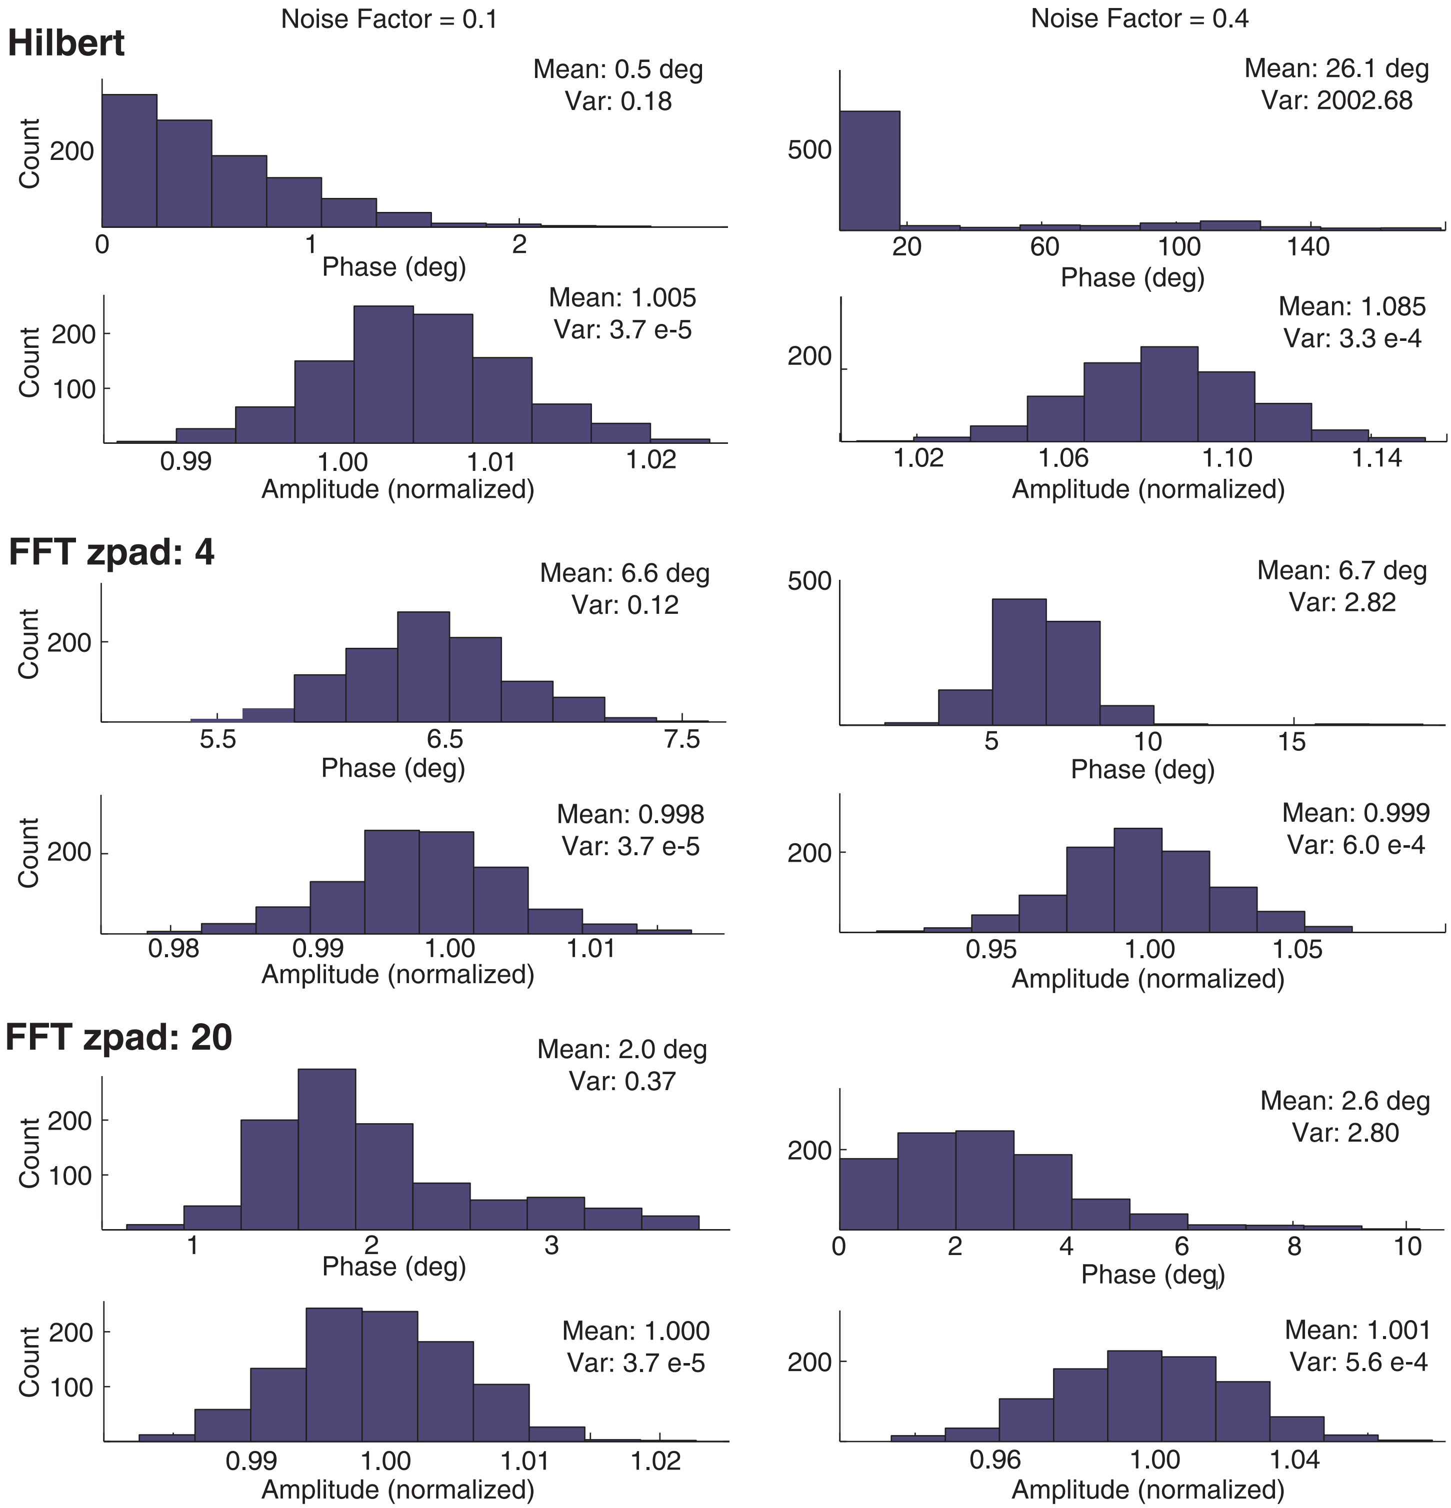

Figure SM2. Histograms of amplitude and phase measurements from simulations using the Hilbert method and zero padded FFTs. 1000 iterations were run in which 50 cycles of a sinusoid with random phase and an arbitrary amplitude of 1 were generated and added to standard Gaussian random noise multiplied by different noise factors. The phase histograms shows the angular differences between each measurement and the true value, while the amplitude histogram shows the direct measurement of amplitude. Perfect measurements in the phase and amplitude histograms would correspond to 0 degrees and 1 respectively.