Bulletin of the Seismological Society of America, Vol. 97, No. 1B, pp. 91–99, February 2007, doi: 10.1785/0120060088

Designing Compact Causal Digital Filters for Low-Frequency Strainmeter Data

by Duncan Carr Agnew and Kathleen Hodgkinson

Abstract For the strainmeter component of the Plate Boundary Observatory, filters are needed to produce low-frequency series (5-minute samples) from the higher-frequency (1 Hz) data generated by the instruments. We present design methods for finding filters that are efficient, causal, and compact. We use standard methods for generating symmetric finite impulse response filters, followed by root finding, selection of roots, and reconstruction of the weights, using procedures that make these processes numerically stable. The final filters show appropriate performance even in the presence of large teleseismic signals, but introduce unavoidable artifacts for strain data from large local earthquakes.

Introduction

Most seismic data are collected from sensors that (often by design) do not respond to signals at periods much longer than those of seismic radiation: though a suitably equipped vertical seismometer can measure changes in gravity, or a horizontal one tilts, they are not usually designed to do so with low noise levels. Most earth strain measurements have, in contrast, focused on much longer period variations, though some have provided seismic data (Benioff, 1935; Fix and Sherwin, 1970; Sacks *et al.*, 1971; Johnston *et al.*, 1987; Borcherdt *et al.*, 1989).

The Plate Boundary Observatory (PBO) component of the EarthScope project will include a large number of strainmeters capable of recording strain from periods of years to frequencies of 1 Hz or more. Recognizing that many users of these data will wish to study variations over very long times, the PBO is providing data that are sampled at a rate low enough (300 sec) to keep data files manageable even if they cover long spans of time. Digital filtering is used to create these low-frequency data from the original 1-Hz data. For data with seismic signals, the filters have to meet slightly unusual goals.

Because the resulting filters may be applicable to other situations, we have therefore written this article both to document the PBO strainmeter filters and to provide a detailed guide to the procedures used to design them. The digital signal-processing literature offers a somewhat bewildering panoply of methods. In the hope of assisting others in choosing from these, we have sought to use methods that are the most conceptually obvious; because these often entail numerical instabilities, we provide ways to minimize these.

Goals for Filters

We start with the properties we wish the filter to have, namely:

1. Efficiency (that is, reasonably small amounts of computation) for large decimation. This is measured by the number of arithmetic operations (multiplications and additions) needed by the filter. In what follows we assume an initial sample interval of one second, so that we will be decimating by 300, to a final Nyquist frequency, f_N, of 0.001666 Hz.
2. Adequate reduction of energy above f_N to minimize aliasing of the final series. In particular, since many of the signals being studied are at much lower frequencies than f_N, the filter should especially reject frequencies at aliases of zero frequency.
3. Causality. If the data have sudden changes (from earthquakes or from possible instrument problems), it is desirable (Scherbaum and Bouin, 1997) that there be no leakage of energy into earlier times.
4. Finiteness. Again, given the potential for sudden large changes, we want the effects of these to extend over only a finite, and known, time.
5. Compactness. Further, we would like the filter to spread any impulses or steps over the minimum possible time span. As we will see, this is equivalent to minimizing the filter phase lag and group delay.

Our filter design is based on a multistage procedure with modest decimation at each stage. The filtering at each stage is done with a finite impulse response (FIR) filter; given an input series x_k, we convolve it with a set of N weights w_k to form

$$y_m = \sum_{k=0}^{N-1} w_k x_{m-k}$$

and then decimate (downsample) the series y by some factor

to form a new series x for the next stage. The sum as written enforces causality. For decimation by 300, we use stages of filtering and decimation by factors of 2 (twice), 3, and 5 (twice), at each stage combining the downsampling and convolution to avoid excess computation. This multistage method greatly simplifies the filter design, since for each filter N can be small.

Filter Theory: Minimum-Delay Filters

To describe the design methods, we first review, some aspects of FIR filters (Oppenheim and Schafer, 1975, 1989). The convolution filter of equation (1) has a complex-valued frequency response:

$$W(f) = \sum_{k=0}^{N-1} w_k e^{-2\pi i f k} \tag{2}$$

over the range of frequencies $-0.5 \le f \le 0.5$. (We have assumed a sample interval of one unit.) This can be generalized to the entire complex plane by forming the polynomial

$$W(z) = \sum_{k=0}^{N-1} w_k z^{-k} \tag{3}$$

where z is a complex number; this is called the z-transform. $W(f)$ is then $W(z)$ evaluated on the unit circle, $z = e^{-2i\pi f}$.

The complex polynomial of equation (3) is degree $N - 1$, and so must have $N - 1$ roots, which we term $r_1, r_2, \ldots, r_{N-1}$; these are the zeros of $W(z)$. Given these roots, the z-transform can be written as a product of degree-one polynomials:

$$W(z) = w_0 \prod_{k=1}^{N-1} (1 - r_k z^{-k}) \tag{4}$$

with the w_0 providing appropriate scaling.

For common choices of w_n there are several "symmetry" relationships that apply to the roots r_k. First, for real weights w_k, the roots r_k must either be real, or occur as complex conjugate pairs. Second, most FIR filter-design methods are for filters with a phase response that is linear with frequency, equivalent to a pure time delay. Such filters usually have an odd number $N = 2M + 1$ of weights that are symmetric around the midpoint of the filter:

$$w_k = w_{2M-k} \quad \text{for } k = 0, 1, \ldots, M.$$

For symmetric weights, if r_k is a root, so is $r_l = 1/r_k$: the roots occur as reciprocal pairs unless they are on the unit circle. To see this, write the z-transform as

$$\begin{aligned} W(z) &= \sum_{k=0}^{2M} w_k z^{-k} \\ &= z^{-M}\left[w_M + \sum_{k=0}^{M-1} w_k (z^{M-k} + z^{k-M}) \right] \\ &= z^{-2M} W(z^{-1}) \end{aligned}$$

so that if $W(r_k) = 0$, then $W(1/r_k) = 0$ as well. So, for real and symmetric weights, the roots of $W(z)$ occur in reciprocal pairs plus single roots on the unit circle, and all pairs (or singlets) not purely real occur as complex conjugates. In this case the response on the unit circle is:

$$\begin{aligned} W(f) &= e^{-2\pi f M}\left[w_M + 2\sum_{k=0}^{M-1} w_k \cos 2\pi k f \right] \\ &= e^{-2\pi f M} W_a(f), \end{aligned} \tag{5}$$

where $W_a(f)$ is the amplitude response, which is purely real. If there are single roots of $W(z)$ on the unit circle, there are values of f for which $W_a(f) < 0$. If $W_a(f) > 0$ for all f, there are, obviously, no roots on the unit circle, but only reciprocal pairs. And if $W_a(f) \ge 0$ for all f, the values of f at which $W_a = 0$ correspond to double roots of $W(z)$, for which the root and its reciprocal are equal.

We describe the z-transform polynomial in terms of its roots because the locations of these affect the compactness of the filter. Suppose some sequence w_n, with z-transform $W(z)$, has a root within the unit circle; without loss of generality we can take this to be the "last" root r_{N-1}. Now consider another z-transform polynomial

$$W'(z) = W(z) \frac{z^{-1} - r_{N-1}^*}{1 - r_{N-1} z^{-1}}. \tag{6}$$

In the polynomial W' the root at r_{N-1} has been divided out and replaced by its reciprocal complex conjugate, which is outside the unit circle. It is not difficult to show that, for $z = e^{-2i\pi f}$,

$$\left| z^{-1} \frac{1 - r_{N-1}^* z}{1 - r_n z^{-1}} \right|^2 = 1, \tag{7}$$

so that the amplitude responses $|W(f)|^2$ and $|W'(f)|^2$ are identical.

For our purposes the important result is a relationship for the corresponding sequences w_k and w'_k. It can be shown (Appendix A) that

$$\sum_{k=0}^{N-2} |w_k|^2 - |w'_k|^2 = (1 - |r_{N-1}|^2)\,|v_{N-2}|^2, \tag{8}$$

where v_{N-2} is the last term of the sequence corresponding to the z-transform of w with the last root factored out (deflation of the polynomial). But because $|r_{N-1}| < 1$, the sum (8) has

to be positive; this means that the sequence w_k has more energy (measured by the sum of squared amplitudes) concentrated before the last point than does the sequence w'_k. It is also true (Appendix A) that, if we express $W(f)$ in terms of amplitude and phase ($W(f) = A(f)\, e^{-\phi(f)}$, then $\phi(f) > \phi'(f)$; that is, the phase lag of $W(f)$ (lags being negative) is less than that of $W'(f)$. And, finally, for the group delay D, defined by $\partial_f \phi(f)$, $D(f) < D'(f)$.

This result can be extended to show that, among all FIR filters with the same amplitude response $|W(f)|^2$, the one whose z-transform has all its roots on or inside the unit circle will have the most concentration of energy toward the early terms of the sequence; that is,

$$\sum_{k=0}^{m} |\, w_k\,|^2$$

will be maximized for all $m < N - 1$. (For $m = N - 1$ the sum is the same, being the filter response at $f = 0$.) This makes the filter as compact as possible, satisfying goal (5). The additional results show that the sequence with this property also minimizes the group delay and makes the phase lag as small as possible for a given amplitude response. Such sequences are therefore usually referred to as providing minimum-phase FIR filters.

Filter-Design Methods

The results just given show that to approach our design goals, we need to find weights for FIR filters that, at each stage,

- Minimize aliasing into frequencies below the new Nyquist frequency f_N, especially aliasing to very low frequencies;
- Have an amplitude response very close to 1 over most of the range from 0 to f_N;
- Have all the roots inside or on the unit circle; and
- Have as few weights as possible.

There are trade-offs between these different goals, notably between the first two and the last one: fewer weights give a poorer approximation to the ideal frequency response. Figure 1 shows the ideal response for the three different filters, with (1) a passband with unit amplitude response; (2) a stopband with ideal response close to zero throughout, but constrained to be closer to zero in the bands that will alias to the frequencies near zero; and (3) a "don't care" band between the passband and stopband, which is needed to allow the response to vary gradually. Because the high-frequency data remain available for those who wish to study signals near the Nyquist of the low-frequency series, we have not been concerned by aliasing at these frequencies, and so have allowed this last band to span the Nyquist in each case.

Usually, the error introduced by filtering and aliasing data with a given power spectrum defines how poor an approximation to the ideal we can accept. In the seismological

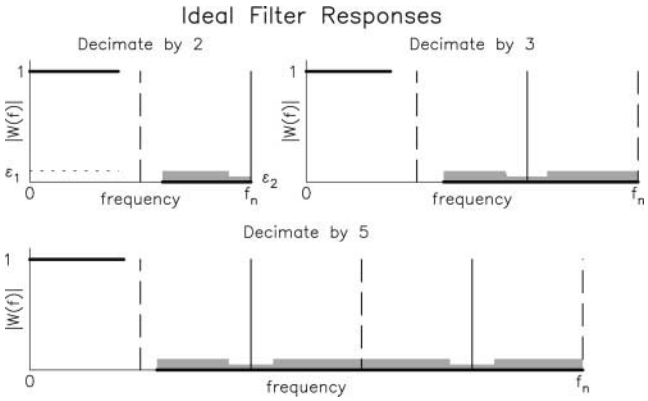

Figure 1. Ideal responses for the decimation filters for the three stages, shown schematically. The black line shows the ideal response, and the gray band shows the acceptable deviation from it, which is ε_1 over most of the stopband, and a smaller value ε_2 around the aliases of zero frequency. There is a "don't-care" transition band between the passband (ideal response 1) and stopband (ideal response 0). The frequency is always dimensionless, with assumed interval 1 unit. Vertical dashed lines are the new Nyquist frequency and its aliases, and vertical solid lines the aliases of zero frequency.

case two issues complicate this approach: first, that the largest signals, being the transients from earthquakes, do not possess a power spectrum, and second that there is no standard level for acceptable contamination of low-frequency data. We have therefore proceeded somewhat intuitively, trying to produce short filters that will only rarely cause significant aliasing. We demonstrate below that, even for large-earthquake signals, the low-passed data are not affected in ways that are likely to alter their interpretation.

To obtain each filter we need to:

1. Find symmetric real weights that approximate the ideal response.
2. Find the roots of the corresponding z-transform polynomial, and modify them so that all roots are on or within the unit circle.
3. Construct the sequence corresponding to the modified set of roots.

There are two different ways to accomplish step 2; which one we choose will set requirements on step 1. We have used two approaches for step 1, giving four choices in all.

One way, suggested by equation (7), is to replace all roots outside the unit circle with their reciprocal complex conjugates, leaving roots on the unit circle in place. This procedure is called the allpass decomposition, because the factors of the from (7) describe allpass filters, which have amplitude response of one. Equation (6), generalized to contain as many such factors as there are roots outside the unit circle, becomes the result that any FIR filter can be expressed as the convolution of a minimum-phase FIR filter and an

allpass filter. Note that the allpass filter itself is not a FIR filter, but must be computed recursively; although this is acceptable for correcting for the application of a noncausal FIR filter (Scherbaum, 2001), it is not for our application.

A second approach, called spectral factorization, is much more common in the signal-processing literature, because it (in general) performs better (less delay) for a given number of weights. This method depends on the result that, for symmetric filters, if $W_a(f) \geq 0$ on the unit circle, all the roots can be paired into reciprocals, with double roots on the unit circle. The factorization consists of taking all the roots inside the unit circle and half of the roots on it, and constructing the corresponding sequence. The resulting filter will have an amplitude response $|W(f)| = \sqrt{W_a(f)}$.

To use spectral factorization we first need to construct a filter with nonnegative amplitude response, $|W(f)| \geq 0$. One way to do this is to use a filter-design method that allows such a constraint; for example, the METEOR program of Steiglitz *et al.* (1992), which uses linear programming to solve the filter-design problem. An advantage of this approach is that it allows constraints on dW_a/df as well as on W_a; we may, for example, make the amplitude response monotonically decrease within the passband.

However, there appear to be limits to the number of weights for which this approach can be made to work. An alternative procedure (and so far as we know a novel one) is to use the well-known Parks–McClellan algorithm (McClellan *et al.*, 1973) in a different way than usual. This algorithm finds the filter that, over J specified frequency bands, satisfies $|W_a(f) - W_j| \leq \varepsilon/s_j$, where W_j is the desired response in the jth band, s_j is a relative weighting, and ε is the overall error level. The level ε, which depends on N, is found in the course of determining the weights. Usually W_j is set to zero in the stopband, which gives a filter with $W_a(f) < 0$ for some frequencies. Although the entire response can be shifted to make $W_a(f)$ nonnegative (Herrmann and Schuessler, 1970), this is suboptimal if the constraints s_j vary, as they do in our case. To allow such variation, we first design the filter with $W_j = 0$ to find ε for a given N. Then we rerun the design with W_j set to ε/s_j in the stopbands, which gives $W_a(f)$ a minimum value of 0 and a maximum of $2\varepsilon/s_j$. This procedure allows the design of factorizable filters with variable stopband response and large N.

For a filter constructed using the allpass decomposition, we can have $W_a(f) < 0$, so the most straightforward approach is to use the Parks–McClellan algorithm.

Once the starting filter is designed, we need to find the roots of the associated z-transform polynomial. This procedure is nonlinear and subject to numerical instability, especially for large numbers of roots. Alternative methods to perform the factorization have therefore been suggested, such as cepstral analysis (Damera-Venkata *et al.*, 2000), or direct solution for the coefficients (Orchard and Willson, 2003). (We attempted the second of these but were not successful.) Working with the roots is the most conceptually obvious approach, and we found no difficulty in using it for filters of the size we were developing. However, some care was needed. We found it necessary to use an optimized root-finding algorithm; both the *rootsl* program of Lang and Frenzel (1994) and the MATLAB root-finding routine *roots* were satisfactory. The most important step was to follow a suggestion of Orchard and Willson (2003) and make $W_a(f) \geq \delta$, where δ is a value smaller than the allowed stopband rejection. Making $W_a(f)$ positive (rather than nonnegative) separates any double roots on the unit circle (of which there are, usually, many) into reciprocal pairs. Because double roots are intrinsically determined with much less precision than single roots, this stabilizes the root finding considerably. In the event that double roots did occur, we found the roots of the polynomial corresponding to $dW(z)/dz$; any double roots on the unit circle for the original polynomial are replaced by single roots for the derivative.

Having found the roots and either changed them (in the allpass method) or discarded half of them (in spectral factorization) we need to construct the coefficients of the associated polynomial. As with the factorization, this procedure can be numerically unstable; but this instability can be avoided if the multiplication of roots is done in the right order, which is called Leja ordering (Calvetti and Reichel, 2003). An algorithm for sorting $N - 1$ roots $r_1, r_2, \ldots, r_{N-1}$ into this order is as follows.

First find the root (call it r'_1) with maximum magnitude $|r|$. Then to find the kth Leja-ordered root r'_k, out of $r'_2, r'_3, \ldots, r'_{N-1}$, select from all the roots r_i, the one that maximizes

$$\prod_{j=1}^{k-1} |r_i - r'_j|.$$

(Obviously those roots already selected will give a product of zero, which will not be a maximum; so this method, while simple, has some redundancy.) Double roots off the unit circle, which will occur if the allpass decomposition is used, should be put in together. Once the roots have been Leja ordered, a recursion can be used to find the coefficients. Start by finding $a_0^1 = -r'_1$. Then repeat the following steps $N - 2$ times, for $k = 1, \ldots, N - 1$:

$$a_k^{(k+1)} = a_{k-1}^{(k)} - r'_k$$

$$a_j^{(k+1)} = a_{j-1}^{(k)} - r'_k a_j^{(k)} \quad \text{for } j = 1, \ldots, k-1$$

$$a_0^{(k+1)} = -r'_k a_0^{(k)}$$

where the superscripts refer to the order of the recursion (and of the polynomial). Note that the as are complex; a useful check on the whole procedure is that the final coefficients should have imaginary parts that are zero to within machine precision.

Design Results

We now describe the specific designs we developed; inevitably, these are compromises among the several goals. As a check on our work, we used (aside from METEOR) two sets of programs: D.C.A. used various stand-alone routines, and K.H. used MATLAB; both of these were satisfactory.

We designed the filters for decimation by 2 and by 3 using spectral factorization, with the initial filter designed using METEOR, using the specifications given in Appendix B. We constrained the passband to have an amplitude response that monotonically decreased with increasing frequency, and was 1 at zero frequency. This minimizes variations in the passband. This program can also specify the response at a particular frequency, so we set the response to zero at exact aliases of zero frequency. Otherwise the response in the stopband was constrained to be positive, usually by an amount that was one-tenth or less of the maximum amplitude response allowed. Figure 2 (top) shows the amplitude response for both filters.

We developed two designs for the decimate-by-5 filter, perhaps the most important part of the system, since its amplitude response dominates the overall response. The first filter used spectral factorization of a design found from METEOR. Although this gives the smallest delay, it proved difficult to get adequate rejection in the stopband. We therefore developed an alternative filter using the allpass decomposition applied to an initial filter design from the Parks–McClellan algorithm. Although even better rejection could be obtained using spectral factorization applied to positive-response filters designed with this algorithm, the resulting delays were larger, so we used only the combination of allpass decomposition with this algorithm. The lower panel of Figure 2 shows the amplitude response for both versions.

The final response and the filtering efficiency depend on the order in which the stages of filtering and decimation are combined. Because all the filters have about the same number of weights, minimizing the total length of the filter means applying the greatest decimation (by 5) to the final stages, since this minimizes the sample interval for the final filter. For the filters as designed, the shortest effective length and the minimum delay occurs for the stages in order 2-2-3-5-5. If the last two stages use the spectral factorization filter, the group delay at zero frequency is 305 sec; if they use the allpass-decomposition filter, it is 356 sec. The effective lengths of the two cases are 2552 sec and 2624 sec. We can construct a filter equivalent to the multistage processing by convolving together the individual filters, each one interpolated with zeros to place its weights at the correct interval. Figure 3 shows the resulting time-domain weights for the two cases. If these weights were applied to the data with a decimation of 300, there would be about 8.7 multiplications and additions per output point; if the filters are applied sequentially, this number is 25.2, about three times as large. Given the data rates involved, these differences are not significant in practice.

Figure 4 shows the passband responses of the combined filter. Because of the monotonicity constraint the design using only spectral factorization is more nearly constant than is the equiripple design produced by the Parks–McClellan algorithm. The latter, however, gives a more nearly constant group delay. Over the band usually used for long-term strain studies (hourly sampling) the differences are small; at the periods of the tides they are negligible.

Filter Performance: Noise, Teleseisms, and Local Earthquakes

Whether the filters described perform adequately depends on the spectrum of the input, especially on the relative

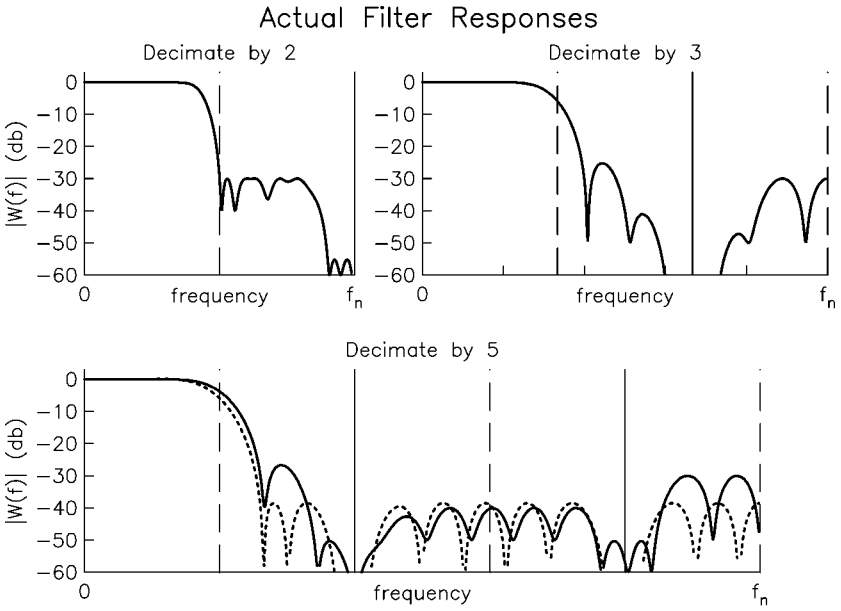

Figure 2. Actual responses for the decimation filters for the three stages. For the final (decimation by 5) filter two responses are shown: solid for the spectral factorization filter derived from METEOR, and dashed for the allpass decomposition filter derived from the Parks–McClellan algorithm. Vertical lines are as in Figure 1.

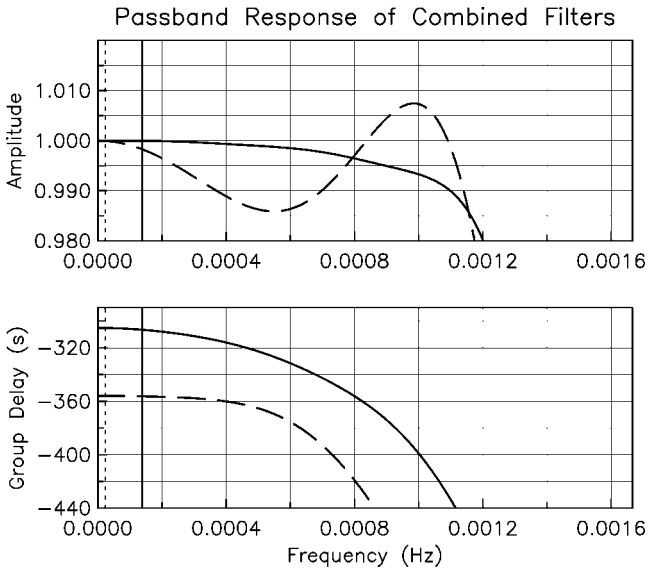

Figure 3. Time-domain response for the combined filters; as in Figure 2, the solid line is for the case of using the spectral factorization filter for the last two stages (decimation by 5); the dashed curve is if the allpass decomposition filters are used for these stages.

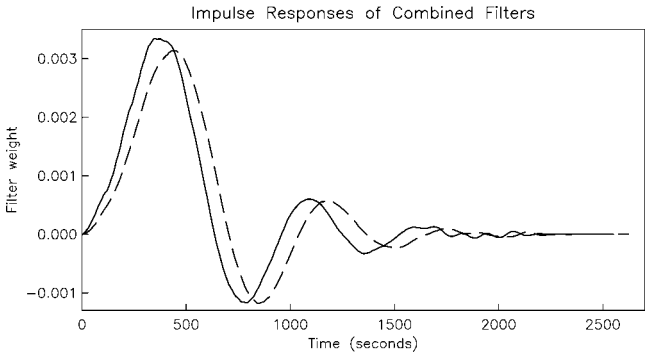

Figure 4. Passband response for the combined filters; dashed and solid as in Figure 3. The vertical solid line shows what the Nyquist frequency would be for hourly sampling; the vertical dashed line is at the frequency of the semidiurnal tides.

amplitudes of the spectrum in the final band of interest, and at frequencies that alias into this. Between earthquakes, the spectrum of earth strain up to 10 sec (Fix and Sherwin, 1972; Berger and Levine, 1974; Beavan and Goulty, 1977) diminishes smoothly with frequency. There will thus be little aliasing from frequencies just above the final Nyquist frequency. The largest levels at higher frequency are in the microseism band. A maximum plausible value comes from taking the high-noise model of Peterson (1993), and converting the levels to strain using a phase velocity of 3 km/sec. The resulting power spectral density is -175 dB (ε^2/Hz), comparable to the level shown for storm microseisms by Beavan and Goulty (1977). For periods from 3 to 10 sec the combined filter has at least 96 dB of rejection, making the aliased signal no larger than -270 dB, some 70 dB below the spectral level found by Berger and Levine (1974) at the 10-min period.

A much more severe test comes from large transient signals from earthquakes. Although ranges for these signals could be found from models, we have preferred to work with actual examples to capture the full complexity of real signals. Because the PBO strainmeters have not been operating very long, our examples come from earlier records from long-base laser strainmeters, which have broadband response and high-dynamic range. Our examples are for two large earthquakes, one at regional distances, for which the filter performs well, and one local, which shows the unavoidable limitations of any filtering scheme.

Our first example is for a magnitude 7.2 event off Cape Mendocino in 2005. Figure 5 shows data from the laser strainmeter installed at Yucca Mountain, Nevada, the closest laser strainmeter (960 km away) and a seismically quiet site. The top panel of the figure shows 0.6 hours (36 minutes) of 1-Hz data as recorded, with peak strains approaching 0.5×10^{-6}. The bottom panel shows the low-passed data that is produced if we convolve these data with the 1-sec weights of Figure 3, and the decimated 5-min data (crosses). The energy above the Nyquist frequency is reduced by a factor of 10,000, to well below the level of the variation between samples (the noise level at long periods).

Our second example is the 1999 Hector Mine earthquake, a magnitude 7.1 shock 110 km from the strainmeters at Piñon Flat Observatory. Figure 6 shows the data from the northwest–southeast laser strainmeter there, plotted as in Figure 5. In this case the original data are clipped by the dynamic range limitations of the recorder (not the instrument). Again, the filter has completely removed the high-frequency energy. However, at this distance there is a step in strain created by the near-field source terms. In the filtered version this step shows several features (a delay, overshoot, and subsequent ringing) that result from the form of the filter, not the actual strain. This "distortion" of the step is in fact unavoidable. The step has a rise time comparable to the rupture time, which is about 20 sec (Ji *et al.*, 2002); to represent it correctly without aliasing would require data sampled at a much higher rate than 300 sec. Getting unaliased data with 300-sec sampling requires removing energy actually present in the strain step, with effects (overshooting and ringing) that are as expected from the Gibbs phenomenon in Fourier theory. Any study of a local earthquake with strain data should not be based on low-frequency data.

Conclusions

The two main results from this article are the actual filters, and the methods used to design them. Our primary conclusion is that this type of filter can be designed most easily, and perfectly adequately, by finding roots of the

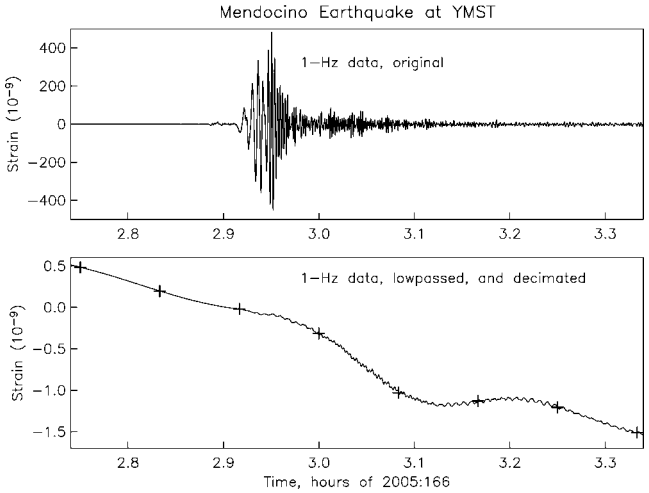

Figure 5. Data from the earthquake off Cape Mendocino, 15 June (day 166) 2005, as recorded at the long-base laser strainmeter at Yucca Mountain, Nevada (36.828° N −116.45° W, oriented 91° east of north). (Top) Raw data recorded at 1 Hz. (Bottom) The data after filtering with the minimum-phase filters without decimation (line) and decimated to 5-min samples (pluses).

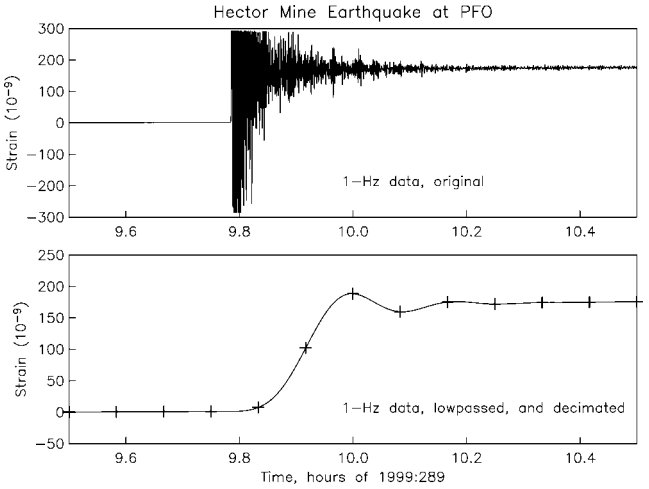

Figure 6. Data from the Hector Mine earthquake, 16 October (day 289) 1999, as recorded at one of the long-base laser strainmeters at Piñon Flat Observatory, California (33.609° N −116.455° W, oriented 135° east of north). (Top) Raw data recorded at 1 Hz. (Bottom) The data after filtering with the minimum-phase filters without decimation (line) and decimated to 5-min samples (pluses).

z-transform, and manipulating these to create a minimum-phase system. Although there can be problems with numerical stability, these can be avoided by using well-designed algorithms, aided by some specific methods we described previously—most notably by requiring a positive (nonzero) amplitude response for filters that will be factored. We believe that the methods we have described are relatively easy to understand—advantages for users who, like ourselves, need to design only a few filters.

Acknowledgments

This research was supported by National Science Foundation Grants EAR 0454475, EAR 0350028-01, and EAR 0323700.

References

Beavan, J., and N. R. Goulty (1977). Earth-strain observations made with the Cambridge laser strainmeter, *Geophys. J. R. Astr. Soc.* **48,** 293–305.

Benioff, H. (1935). A linear strain seismograph, *Bull. Seism. Soc. Am.* **25,** 283–309.

Berger, J., and J. Levine (1974). The spectrum of earth strain from 10^{-8} to 10^{2} Hz, *J. Geophys. Res.* **79,** 1210–1214.

Borcherdt, R. D., M. J. S. Johnston, and G. Glassmoyer (1989). On the use of volumetric strain meters to infer additional characteristics of short-period seismic radiation, *Bull. Seism. Soc. Am.* **79,** 1006–1023.

Calvetti, D., and L. Reichel (2003). On the evaluation of polynomial coefficients, *Numer. Algorithms* **33,** 153–161.

Damera-Venkata, N., B. L. Evans, and S. R. McCaslin (2000). Design of optimal minimum-phase digital FIR filters using discrete Hilbert transforms, *IEEE Trans. Signal Proc.* **48,** 1491–1495.

Fix, J. E., and J. R. Sherwin (1970). A high-sensitivity strain-inertial seismograph installation, *Bull. Seism. Soc. Am.* **60,** 1803–1822.

Fix, J. E., and J. R. Sherwin (1972). Development of LP wave discrimination capability using LP strain instruments, *Rep. TR-72-3,* Teledyne-Geotech, Garland, Texas, NTIS AD-748-232.

Herrmann, O., and G. W. Schuessler (1970). Design of nonrecursive digital filters with minimum-phase, *Electron. Lett.* **6,** 329–330.

Ji, C., D. J. Wald, and D. V. Helmberger (2002). Source description of the 1999 Hector Mine, California, earthquake, part II: Complexity of slip history, *Bull. Seism. Soc. Am.* **92,** 1208–1226.

Johnston, M. J. S., A. T. Linde, M. T. Gladwin, and R. D. Borcherdt (1987). Fault failure with moderate earthquakes, *Tectonophysics* **144,** 189–206.

Lang, M., and B. C. Frenzel (1994). Polynomial root finding, *IEEE Signal Proc. Lett.* **1,** 141–143.

McClellan, J. H., T. W. Parks, and L. R. Rabiner (1973). A computer program for designing optimum linear phase FIR filters, *IEEE Trans. Audio Electroacoust.* **21,** 506–526.

Oppenheim, A. V., and R. W. Schafer (1975). *Digital Signal Processing*, Prentice Hall, Englewood Cliffs, New Jersey.

Oppenheim, A. V., and R. W. Schafer (1989). *Discrete-Time Signal Processing*, Prentice Hall, Englewood Cliffs, New Jersey.

Orchard, H. J., and A. N. Willson (2003). On the computation of a minimum-phase spectral factor, *IEEE Trans. Circuits Systems I* **50,** 365–375.

Peterson, J. (1993). Observations and modeling of seismic background noise, *U.S. Geol. Surv. Open-File Rept. 93-322*.

Sacks, I. S., S. Suyehiro, D. W. Evertson, and Y. Yamagishi (1971). Sacks-Evertson strainmeter, its installation in Japan and some preliminary results concerning strain steps, *Pap. Meteor. Geophys.* **22,** 195–207.

Scherbaum, F. (2001). *Of Poles and Zeros: Fundamentals of Digital Seismology*, Kluwer Academic Publishers, Boston, Massachusetts.

Scherbaum, F., and M.-P. Bouin (1997). FIR filter effects and nucleation phases, *Geophys. J. Int.* **130,** 661–668.

Steiglitz, K., T. W. Parks, and J. F. Kaiser (1992). METEOR: a constraint-based FIR filter design program, *IEEE Trans. Signal Proc.* **40,** 1901–1909.

Appendix A: Root Locations and Minimum Filter Responses

We begin with equation (6). Rearranging this, we get

$$W(z) = V(z)(1 - r_{N-1}z^{-1}) \quad \text{and}$$
$$W'(z) = V(z)(z^{-1} - r^*_{N-1}) \tag{A1}$$

where $V(z)$ is a z-transform (polynomial) of degree $N - 2$. From these polynomial relationships the expressions for the coefficients of W and W' are

$$w_k = v_k - r_{N-1}v_{k-1} \quad \text{and}$$
$$w'_k = -r_{N-1}v_k + v_{k-1} \tag{A2}$$

from which it is easy to find the result (8) for the difference in sums of squares up to the penultimate term.

For the phase and group delay, we can use the results (Oppenheim and Schafer, 1989) that the phase of the allpass term in equation (6) is, for $r_{N-1} = ae^{2\pi i\theta}$,

$$-2\pi f - 2 \arctan\left[\frac{a \sin 2\pi(f - \theta)}{1 - a \cos 2\pi(f - \theta)}\right]$$

and the group delay is

$$\frac{1 - a^2}{|1 - ae^{2\pi i\theta}e^{-2\pi if}|^2},$$

which is positive for $a < 0$ (the root inside the unit circle) so that $W'(f)$ has a greater delay, and a larger (more negative) phase than $W(f)$.

Appendix B: Specifications for Initial Filters

In this section we give the detailed specifications used for the design of the initial FIR filters and the final results for the minimum-phase filters.

METEOR requires the responses in different frequency bands, each band extending from f_l to f_h. For each band, the specification needs a minimum and maximum amplitude response. W_l and W_h); these may in general vary across the band, but in all cases we held them constant. The program also requires that we specify if the limit is "hugged" or not; the design moves the response as far as possible from limits that are not hugged, allowing the response to be arbitrarily close to those that are. So, for example, in a passband, we would hug the upper limit of 1, and not the lower limit; the latter will, in practice, set the response at the edge of the frequency band. The program also allows us to specify if the amplitude response is concave up or down in a band; we always specified this, for the passband only, as concave down.

For the decimate-by-two filter we specified a stopband closer to zero in the region that would alias into the passband, with the response set to zero (but still nonnegative) at the original Nyquist. This last feature, while desirable, could

Table B1
Specification for Decimate-by-Two Filter

Band	f_l	f_h	W_l	W_h	Hugged
1	0.0	0.2	0.81	1.00	W_h
2	0.25	0.40	0.0001	0.001	W_l
3	0.40	0.45	0.0	—	
4	0.45	0.49	0.000001	0.0000003	W_l
5	0.49	0.50	0.0	—	
6	0.50	0.50	—	0.0	

Table B2
Specification for Decimate-by-Three Filter

Band	f_l	f_h	W_l	W_h	Hugged
1	0.0	0.137	0.81	1.0	W_h
2	0.14	0.20	0.0	—	W_l
3	0.197	0.280	0.00001	0.003	W_l
4	0.303	0.363	0.0000001	0.000001	W_l
5	0.38	0.50	0.00001	0.001	W_l

Table B3
Specification for Decimate-by-Five Filter

Band	f_l	f_h	W_l	W_h	Hugged
1	0.0	0.085	0.81	1.0	W_h
2	0.09	0.12	0.0	—	
3	0.125	0.165	0.0001	0.01	W_l
4	0.170	0.195	0.000001	0.00001	W_l
5	0.198	0.202	0.0000003	0.000001	W_l
6	0.205	0.220	0.000001	0.00001	W_l
7	0.225	0.375	0.00001	0.0001	W_l
8	0.380	0.420	0.0000003	0.00001	W_l
9	0.440	0.500	0.00001	0.001	W_l

Table B4
Specification for Decimate-by-Five Filter

Band	f_l	f_h	R	Weight
1	0.0	0.01	1.0	10
2	0.02	0.07	1.0	1
3	0.13	0.18	0.0	1
4	0.19	0.21	0.0	30
5	0.22	0.38	0.0	1
6	0.39	0.41	0.0	30
7	0.42	0.50	0.0	1

Table B5
Weights for Minimum-Phase Filters

k	w_k^{II}	w_k^{III}	w_k^{Va}	w_k^{Vb}
1	0.0983262	0.0373766	0.0218528	0.0137003
2	0.2977611	0.1165151	0.0458359	0.0323895
3	0.4086973	0.2385729	0.0908603	0.0627611
4	0.3138961	0.3083302	0.1359777	0.1018581
5	0.0494246	0.2887327	0.1830881	0.1454735
6	−0.1507778	0.1597948	0.1993418	0.1742102
7	−0.1123764	0.0058244	0.1957624	0.1888950
8	0.0376576	−0.0973639	0.1561194	0.1786665
9	0.0996838	−0.1051034	0.0994146	0.1440888
10	0.0154992	−0.0358455	0.0346412	0.0931153
11	−0.0666489	0.0359044	−0.0236544	0.0316546
12	−0.0346632	0.0632477	−0.0580081	−0.0217322
13	0.0322767	0.0302351	−0.0703257	−0.0580216
14	0.0399294	−0.0168856	−0.0555546	−0.0707892
15	−0.0097461	−0.0356758	−0.0287709	−0.0601628
16	−0.0341585	−0.0190635	0.0032613	−0.0350987
17	−0.0039241	0.0126188	0.0267938	−0.0047034
18	0.0246776	0.0159705	0.0358952	0.0202246
19	0.0099725	0.0082144	0.0311186	0.0332569
20	−0.0157879	−0.0087978	0.0134283	0.0322796
21	−0.0099098	−0.0037289	−0.0028524	0.0222429
22	0.0078510	−0.0017068	−0.0170042	0.0075065
23	0.0081126	0.0028335	−0.0176765	−0.0051609
24	−0.0026986		−0.0123123	−0.0117785
25	−0.0061424		−0.0036798	−0.0142579
26	0.0007108		0.0057730	−0.0085583
27	0.0039659		0.0059817	−0.0023946
28	−0.0006209		0.0083501	0.0030231
29	−0.0017117		0.0000581	0.0062511
30	0.0007240		0.0005724	0.0017000
31			−0.0033127	0.0019444
32			0.0004411	0.0000312
33			−0.0030766	−0.0014347
34			0.0016604	−0.0030002
35				0.0018199

not be included in the filters for more decimation without rendering the design problem insoluble.

These specifications (given in Table B1) could be met by a FIR filter with 51 weights; spectral factorization produced a minimum-phase filter with 30 weights. The specifications for the decimate-by-three filter are given in Table B2; these could be satisfied by a filter with 45 weights; spectral factorization produced a minimum-phase nta filter with 23 weights. The specifications for the decimate-by-five filter using METEOR are in Table B3; these could be satisfied by a filter with 67 weights; spectral factorization produced a minimum-phase filter with 34 weights. As can be seen in Figure 2, the resulting response had relatively large sidebands that could alias into the upper part of the passband. We therefore designed a filter using the Parks–McClellan method, for which the specifications are given in Table B4; for this design we also need to specify the number of weights (35).

Finally, Table B5 gives the filter weights: w^{II} for the decimate-by-two; w^{III} for the decimate-by-three; and w^{Va} and w^{Vb} for the decimate-by-five.

Institute of Geophysics and Planetary Physics
Scripps Institution of Oceanography
University of California, San Diego
La Jolla, California 92093-0225
(D.C.A.)

UNAVCO, Inc.
New Mexico Institute of Technology
Socorro, New Mexico 87801
(K.H.)

Manuscript received 10 April 2006.