

In search of the vibroseis first arrival

Stewart Trickett*

Juniper Bay Software, Calgary, Canada

Received October 2021, revision accepted March 2022

ABSTRACT

The first step in correcting for time delays of land seismic data due to low-velocity weathered layers is to pick the first-arrival times of the refracting energy. But doing so for vibroseis data can be difficult, as the seismic wavelet is often ringy and uncompact, resulting in cycle-skipped picks. Even when we manage to pick a waveform feature consistently, it is not clear where the first-arrival time is in relation to it. I present a novel method that shapes the seismic wavelet to a Ricker wavelet whose peak is located at the true arrival time, so the time of the first arrival is unambiguous. Further, the arrivals are less ringy and their energy more focused, so that they are less likely to cycle skip or be overwhelmed by random noise. The result is more accurate and consistent first-arrival picks.

Key words: Attenuation, Noise, Seismics, Signal processing.

INTRODUCTION

Correcting for time delays caused by low-velocity near-surface weathered layers is one of the oldest steps in land seismic processing (e.g. Gardner, 1939). The standard approach is this (Cox, 1999, Chapter 5):

- Pick the times of the first arrivals (also called first breaks).
- Interpret the depth and velocity of the weathered layers from these picks.
- Apply statics to correct for the weathered layers, in effect turning the near surface into a constant-velocity layer.

Weathering statics can correct for severe short-spatial-wavelength statics that otherwise might cause surface-consistent residual statics to “bust” or cycle skip. They can also correct for longer spatial-wavelength statics than residual statics correction is able to resolve (Wiggins et al., 1976) and provide a velocity model of the near surface which can be exploited by normal moveout and migration (Ellison et al., 2017).

Here we are concerned with the first step – picking the first-arrival time of each trace. The most fundamental definition of the first-arrival time of a given trace is:

Def 1: The shortest time it takes for seismic energy to travel from source to receiver.

For impulsive sources, this is often interpreted in practice as:

Def 2: The time of the initial onset of source energy.

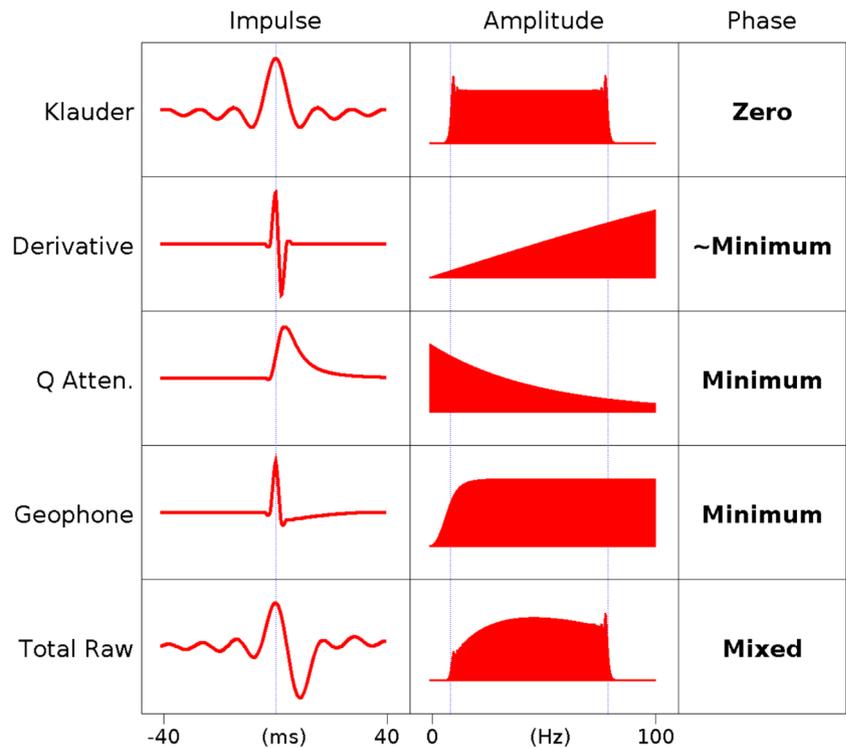
For correlated vibroseis data, however, this interpretation is unworkable. Taking a model-based approach similar to Hart et al. (2001), the raw vibroseis seismic wavelet is composed of, at a minimum, a convolution of the following:

- Klauder wavelet
- Far-field temporal derivative
- Q attenuation response
- Geophone response

The Klauder wavelet is the autocorrelation of the vibroseis sweep and thus is zero phase. We assume we have a reasonable approximation to it, although that is not always the case (Baeten and Ziolkowski, 1990). The temporal derivative is a consequence of recording the far-field particle velocity generated from the applied ground force (Aki and Richards, 2002, §4.2.1). It can be closely approximated by a minimum-phase wavelet, especially within a limited sweep band. Q attenuation is a broadening of the seismic wavelet caused by anelastic propagation and other effects, and tends to be severe

*E-mail: stewart.trickett@gmail.com

Figure 1 The impulse responses, amplitude spectra and phase spectra of various components making up a raw vibroseis seismic wavelet with an 8–80 Hz linear sweep. Starting from the top is the Klauder wavelet, a temporal derivative, the Q attenuation response and the geophone response. The total response is at the bottom. Time zero (that is, the first-arrival time) and the sweep band are indicated with light vertical lines. Note the ringy precursors of the total response before time zero.



in the near surface (Aki and Richards, 2002, §5.5). Its response is minimum phase, but the amount of attenuation is typically unknown. The geophone response depends on the sensor type. The traditional moving-coil velocimeter has the impulse response of a damped spring, with a natural frequency typically around 10 Hz (Hons et al., 2008). The more recent MEMS accelerometer has, within the sweep band, the response to particle velocity of a temporal derivative. Both responses are minimum phase. Other effects like recording instrument responses are typically so mild in modern acquisition that they can be ignored. We also assume that the SEG polarity standard for vibroseis data (Landrum et al., 1994) has been followed. These components, and the resulting seismic wavelet (the convolution of the components), are shown in Fig. 1.

We now see why the vibroseis first-arrival time cannot be defined as the initial onset of source energy. The seismic wavelet contains the zero-phase Klauder wavelet, whose onset is many seconds before time zero. Instead, we might try to pick a consistent waveform feature near a strong energy surge. But ringy precursors can often cause cycle skipping, and even when we can avoid this, it is not clear where the true first-arrival time is in relation to the picked feature.

Another problem is that the first arrival suffers from increasing amounts of anelastic attenuation as the arrival time increases. Figure 2 shows how the first-arrival wavelet changes

with time assuming a constant Q value of 30. As the arrival time increases, the wavelet becomes less ringy but increasingly delayed. This makes it difficult to know what feature to pick. If we chose, say, a strong positive peak as the first arrival, the arrival time is too early at early times and too late at late times.

To overcome these problems, we might define the vibroseis first-arrival time as:

Def 3: The time of the initial onset of source energy when the Klauder wavelet is replaced with a minimum-phase wavelet.

Motivated by the above definition, the seismic wavelet is often first shaped using an all-pass filter that converts the Klauder wavelet to minimum phase (Ristow and Jurczyk, 1975; Gibson and Larner, 1984). This can help to pick a consistent feature. But the difficulty of creating a causal wavelet out of one whose amplitude spectrum is both sharp edged and band limited means that precursors are not fully removed, so that cycle skipping remains a risk. Even if you can avoid ringy precursors, it is still uncertain where the true first arrival is in relation to the picked feature, as discrete minimum-phase band-limited wavelets tend to have a gradual onset of energy which is easily overwhelmed by noise. These problems produce at best a bulk shift from the true weathering static, and at worst static inaccuracies at all spatial wavelengths.

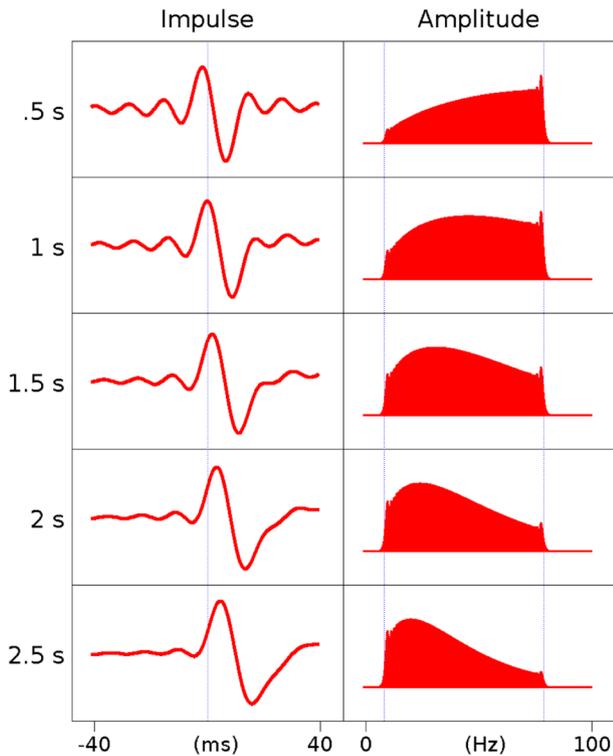


Figure 2 The first-arrival wavelet changes with the amount of Q attenuation it has suffered. Here we show the first-arrival wavelet at various arrival times assuming a constant Q of 30. An arrival time of 0.5 s has a very ringy wavelet, with its maximum peak just before the true arrival time. The wavelet at 2.5 s is less ringy, but its maximum peak is now later than the true arrival time.

Instead, I propose exploiting a fourth definition for the first-arrival time:

Def 4: The time of the earliest signal peak when the seismic wavelet’s initial peak is at its time zero.

Below I present a novel method which modifies not just the phase but the amplitude spectrum of the seismic wavelet, producing a simply shaped first arrival whose first peak is at the true first-arrival time. The first arrival is also less ringy and has its energy more focused, so it is less likely to cycle skip or be overwhelmed by random noise. The result is more accurate and consistent first-arrival picks.

METHOD

Shaping to a zero-phase Ricker wavelet

Instead of trying to make the seismic wavelet minimum phase, I propose replacing it with a wavelet which is:

- Zero phase

- Compact and simply shaped. Notably, it is not ringy.
- Has almost all of its energy contained within the sweep frequency band.
- Has a single strong positive peak at time zero, but no other peaks.

Here I use a zero-phase Ricker wavelet (Ricker, 1940; Hosken, 1988), whose impulse response as a function of time t is

$$(1 - 2a) \exp(-a), \text{ where } a = (\pi f_p t)^2 \tag{1}$$

with amplitude spectrum¹ as a function of frequency f :

$$\frac{2 f^2}{\sqrt{\pi} f_p^3} \exp\left(-\frac{f^2}{f_p^2}\right). \tag{2}$$

A zero-phase Ricker wavelet and its amplitude spectrum are shown in Figure 3. It has three extrema (or ‘loops’ in Hosken’s parlance) comprised of one peak (that is, a local maximum) and two troughs (local minima). This is critical. If a first-arrival picker is keying on peaks, there will be no signal precursors to confuse it, although of course noise might still do so.

The sole parameter for this wavelet is its peak-amplitude frequency f_p . I recommend a value of 45% of the highest sweep frequency to ensure that most of its energy is contained within the sweep frequency band.

Given the modelled input wavelet (the total response described in the introduction) and desired output wavelet (a zero-phase Ricker), applying the shaping is straightforward: divide the discrete Fourier transform (DFT) of a seismic trace by the DFT of the modelled input and multiply by the DFT of the desired output. Prewhitening is needed for the DFT of the modelled input to avoid division by small values, but the results are not sensitive to its precise level, as almost all of the energy of the desired output wavelet is within the sweep band.

Once a peak is chosen as the first arrival, there is no need to adjust the pick to the actual onset of energy by moving it to a previous trough, inflection point or zero crossing – an exercise which is prone to error due to noise, and which rarely gives the true first-arrival time even without noise. Because the peak is at time zero of the seismic wavelet, it is at the first-arrival time, and because it is a strong peak, its exact location is not much affected by noise.

¹ This is the amplitude spectrum for continuous signals. For discrete signals (that is, when using Discrete Fourier Transforms), different constant scalings are commonly used.

Figure 3 The impulse responses and amplitude spectra of various output seismic wavelets. Starting from the top is the total raw seismic wavelet shown in Figure 1 with an 8–80 Hz linear sweep, the raw seismic wavelet with its Klauder wavelet converted to minimum phase, a zero-phase Ricker wavelet and a four-loop Ricker wavelet and a four-loop Ricker wavelet shifted so that its first peak is at time zero. All Ricker wavelets have a peak frequency of 36 Hz. Time zero (that is, the first-arrival time) and the sweep band are indicated with light vertical lines.

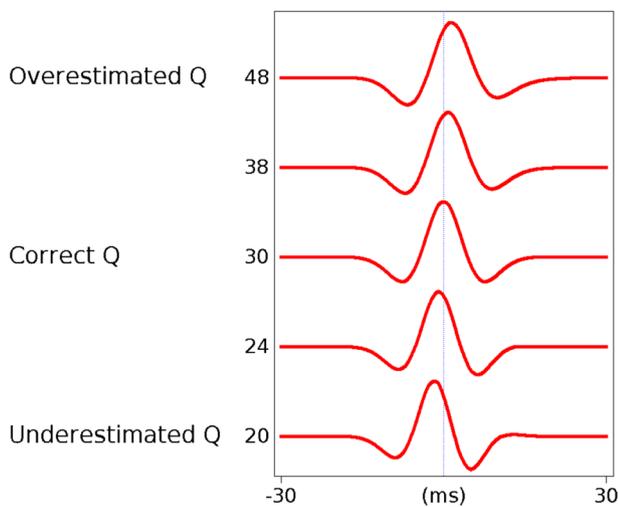
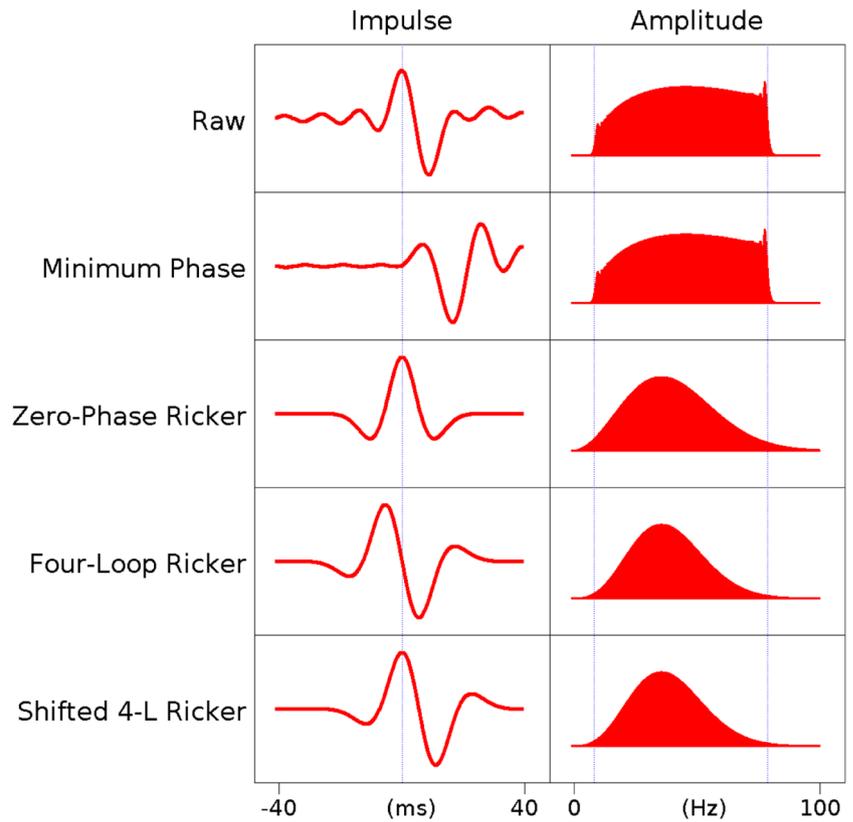


Figure 4 Results of shaping a seismic wavelet to a 50-Hz zero-phase Ricker wavelet using a variety of Q attenuation values. The raw seismic wavelet has a Q value of 30 at the arrival time of 1 s. Using the wrong amount of Q attenuation shifts the peak of the output wavelet away from the true arrival time (that is, from time zero) by a few milliseconds. Although the wavelet becomes lopsided, it maintains its simple shape without peak precursors.

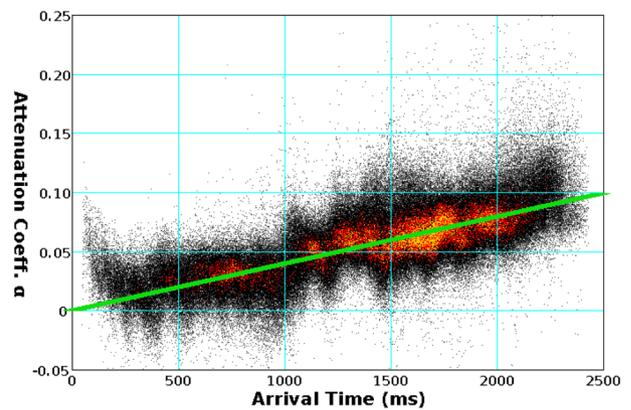


Figure 5 A plot of estimated attenuation coefficients α versus expected arrival time for the real dataset using the method described in Appendix B. The green line is a manually selected $\tilde{\alpha}(\tau)$ function. It corresponds to a constant Q value of about 78 through the formula $Q = \pi\tau/\alpha$, where τ is the arrival time.

Shaping to a four-loop Ricker wavelet

Hosken (1988) notes that the Ricker wavelet is based on an incorrect model of seismic propagation and has only one

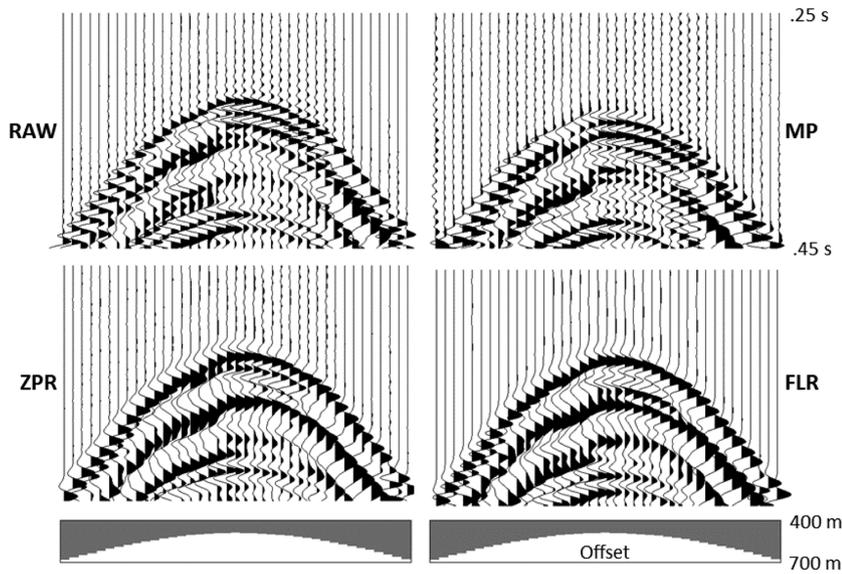


Figure 6 Shallow first arrivals with various shaping filters applied, including RAW (no filter), MP (Klauder wavelet converted to minimum phase), ZPR (wavelet shaped to a zero-phase Ricker), and FLR (wavelet shaped to a four-loop Ricker). The last two are far less ringy, decreasing the chances of the picker cycle skipping.

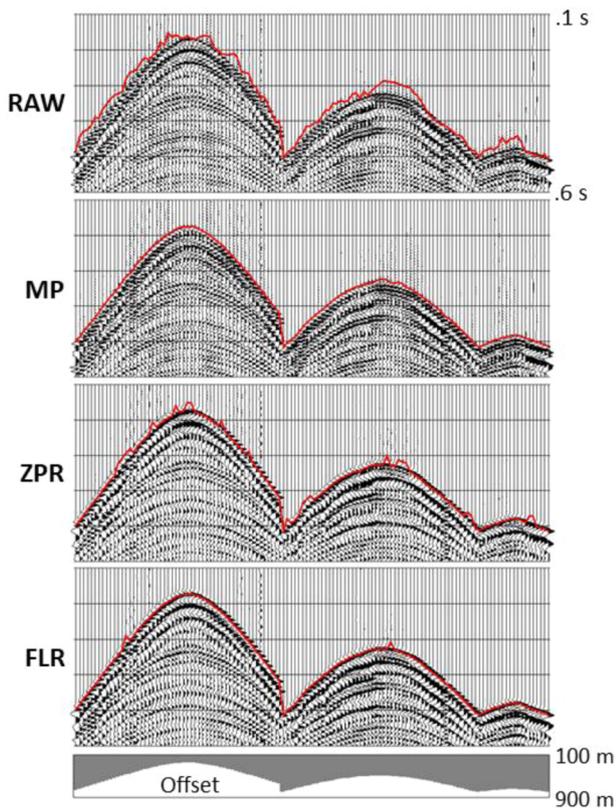


Figure 7 Picking the four data sets with a trace-by-trace first-arrival picker. The RAW dataset suffers from considerable cycle skipping due to the ringiness of the data, while the other three are reasonable. The MP picks are surprisingly good given how ringy the data look.

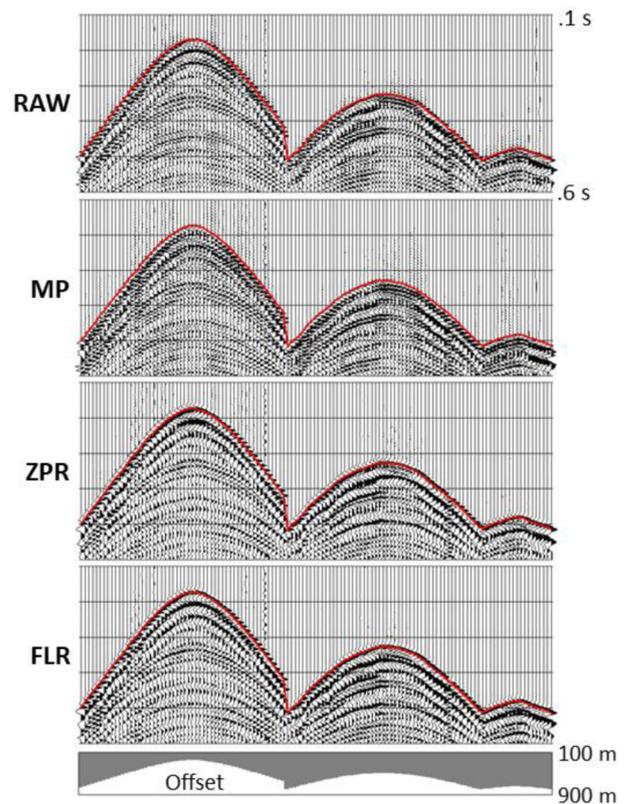


Figure 8 Picking the four datasets with a first-arrival picker that encourages continuity between adjacent picks. A consistent feature on each dataset is now picked, despite the ringiness of the top two.

parameter (the peak frequency) to control its shape. For these and other reasons, he recommended that the Ricker wavelet never be used. Its simple compact shape and limited bandwidth seem, however, well suited to this application. If it has one flaw, it is that it does not fall off quickly enough at the low and high frequencies. The Klauder wavelet is assumed to be the autocorrelation of the applied ground force, but it is not always accurate. These inaccuracies tend to be most pronounced near the limits of the sweep band. It might be advantageous, therefore, to shape to a wavelet that has a narrower frequency band, so that the amplitudes at the sweep-band limits are smaller. In doing so, I am going to relax the constraints that the output wavelet is zero phase and has only one peak.

Taking the temporal derivative of the Ricker wavelet results in a (again following Hosken's parlance) four-loop Ricker wavelet with peak frequency g_p (Appendix A):

$$s_t t (2b - 3) \exp(-b), \text{ where } b = 2(\pi g_p t)^2 / 3 \quad (3)$$

with amplitude spectrum

$$s_f f^3 \exp\left(-\frac{3f^2}{2g_p^2}\right). \quad (4)$$

The wavelet is shown in Figure 3 and has a constant phase spectrum of 90° . The scalars s_t and s_f , defined by Equations (A9) and (A10), are selected so that the maximum value of the wavelet is one. To place the leading peak at time zero (that is, at the first-arrival time), we must shift the wavelet by about $0.2045 / g_p$ s (Equation (A8)), shown at the bottom of Figure 3. Although the four-loop Ricker has two peaks, the second peak is small in magnitude and trails the larger peak, so that it is not likely to be mistaken for the first arrival. Note how the four-loop Ricker does indeed have a narrower frequency band than the zero-phase Ricker, allowing it to better fit within the sweep band.

Estimating Q attenuation

All of this assumes we know the amount of Q attenuation that the seismic wavelet has suffered, and normally we do not. There are at least two approaches to overcome this. The first is to assume a fixed, reasonable rate of attenuation for all first arrivals. Although this will rarely be correct for any trace, it works surprisingly well. Due to the simple shape of the attenuation response, the main effect of incorrect Q is to cause a time shift in the position of the wavelet peak of a few milliseconds (Fig. 4). Importantly, peak precursors are not generated. Although this shift is undesirable, it also arises when the seis-

mic wavelet is not shaped, or when the Klauder wavelet is converted to minimum phase (Kobayashi, 2001).

The second approach is to estimate the amount of Q attenuation, such as in Hatherly (1986). This is a difficult problem, as the first-arrival event is not isolated, but often closely followed by subsequent events. Even if we restrict our analysis to a small time window surrounding the first arrival, the totality of these events has an amplitude response which can distort the estimation. Noise also distorts Q estimation, usually resulting in too large of a value. We do not, however, require precision in order to produce an improved estimate of the first arrival. A single 'attenuation versus expected arrival time' function for the entire survey should normally suffice. Appendix B describes a novel method for determining this that is adapted from the Centroid Frequency Shift algorithm of Quan and Harris (1997).

Adapting the automatic picking algorithm

Some changes to the automatic picking algorithm are needed for first arrivals shaped to a Ricker wavelet. Specifically:

- Peaks should be picked.
- Tests such as Coppens' (1985) energy-ratio test should be centred about $0.75 / f_p$ s *before* each candidate peak.
- Once a peak has been selected as indicating the first arrival, no adjustment should be made to the time.

REAL DATA EXAMPLE

Here I compare the methods on a real seismic survey. I generated four datasets filtered as follows:

RAW	No filter
MP	Klauder wavelet shaped to minimum phase
ZPR	Seismic wavelet shaped to a zero-phase Ricker wavelet (first proposed method)
FLR	Seismic wavelet shaped to a shifted four-loop Ricker wavelet (second proposed method)

The picking algorithms key on specified features like peaks and troughs. For RAW, ZPR and FLR data, I keyed on peaks (the first because the modelled first-arrival time is close to a strong peak, the last two because they are designed so that the first-arrival time is on a strong peak). The MP data are more difficult as there is no strong consistent feature near the true first-arrival time, as Figure 3 indicates. Keying on a trough, whose time was then decreased to the nearest zero crossing, seemed to give the best result.

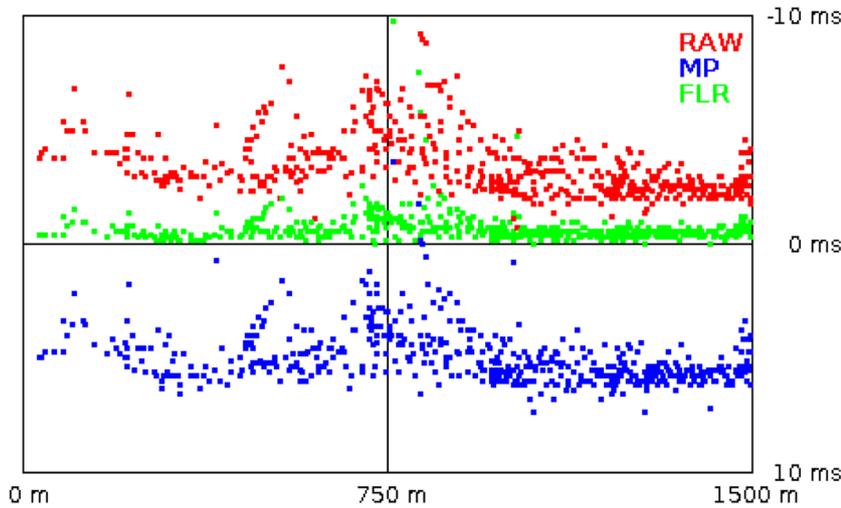


Figure 9 The time difference between the ZPR picks and the other three datasets for a single clean shot record, plotted as a function of offset. The ZPR and FLR picks are nearly identical. The RAW picks are roughly 3 ms early, while the MP picks are about 5 ms late.

Our example seismic survey has as its source a single vibrator with an 8–115 Hz, 20 s linear sweep. Some parts of the survey had first arrivals contaminated with strong random noise.

Figure 5 shows the results of estimating anelastic attenuation as a function of expected arrival time using the method described in Appendix B. Based on a small part of the survey, a coherent pattern of attenuation coefficients α emerges. This allows us to draw a simple linear function through it corresponding to a constant Q value of around 78 for all times, which was then used in the wavelet shaping.

Figure 6 shows a set of first arrivals with the four filters applied. Like many vibroseis datasets, there are ringy precursors prior to the first arrivals in both the RAW and MP datasets. The ZPR and (especially) FLR datasets, however, removed most ringiness prior to the first arrival, leaving the first arrivals easier to pick. A close examination shows that their first arrivals look remarkably like our modelled wavelets in Figure 3.

Many first-arrival pickers use a trace-by-trace strategy, picking each trace independently without reference to surrounding traces. Figure 7 shows what happens when a popular commercial trace-by-trace picker is applied. The ringiness in the RAW dataset throws off the picker badly, while the two proposed shaping methods give reasonable results. Surprisingly, the MP dataset picks are also more spatially continuous than the RAW dataset. It seems to have focused the energy of the first arrival enough to allow the picker to do its job.

It is better, however, to use an automatic picker which encourages spatial continuity between adjacent picks (Appendix

C). This is what I will use from hereon in. Figure 8 shows that the results are now much improved for all datasets.

The picks between the datasets are not identical, however. Figure 9 compares the difference between the ZPR picks to the other three datasets for a single clean shot record, plotted as a function of offset. The picks for FLR are very similar to ZPR, with about a 0.5 ms bulk difference. The RAW picks are about 3 ms earlier on average, although this difference decreases at later offsets. This is likely due to the fact that the RAW data suffer increasing delays due to anelastic attenuation, while the ZPR data compensate for this effect. The MP picks are about 5 ms later than the ZPR picks.

Where the two proposed shaping methods shine, however, is when the data are noisy. Figure 10 shows parts of a shot record at around 0.8 s, with Figure 11 comparing the difference between the ZPR picks with the other three datasets. Figure 12 shows part of a shot record at around 2 s, flattened for ease of display. The RAW and MP first arrivals are difficult to discern for many traces. This is a common problem at deeper times – the amplitudes of the first arrivals are often low compared to the ambient noise, resulting in cycle-skipped picks. The arrivals for the ZPR and FLR datasets stand out better, as these shaping filters reduce low-frequency noise and better focus the energy of the first arrivals so that they are not easily overwhelmed by noise, resulting in more consistent picking.

The impact of Q is mild for this dataset. At offsets of 3000 m, the difference between compensating and not compensating for Q in the shaping filter is about -4 ms in the picked arrival times. Other datasets have lower average Q values (e.g., 30 rather than 78), resulting in differences of over

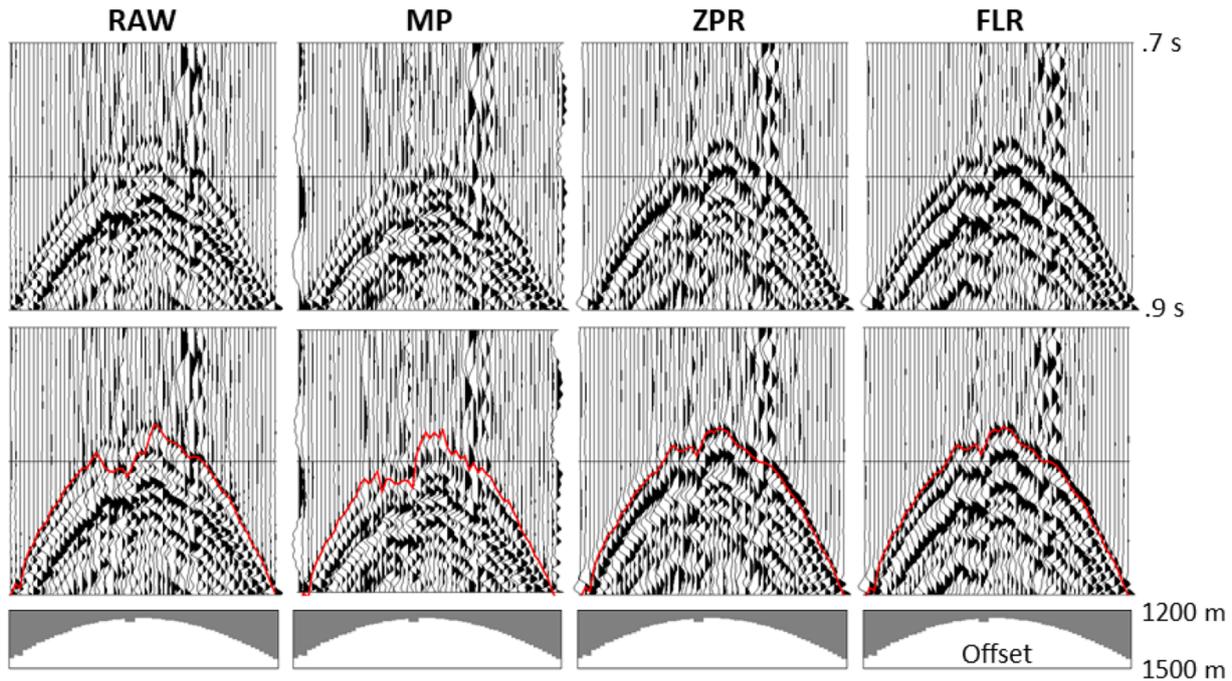


Figure 10 First arrivals from the four datasets at moderately deep times of around 0.8 s. The bottom plots are the same as the top, but with the picks shown. The arrivals are contaminated by noise as the signal is lower amplitude than at earlier times. The arrivals are clearer on the ZPR and FLR datasets, resulting in better picks.

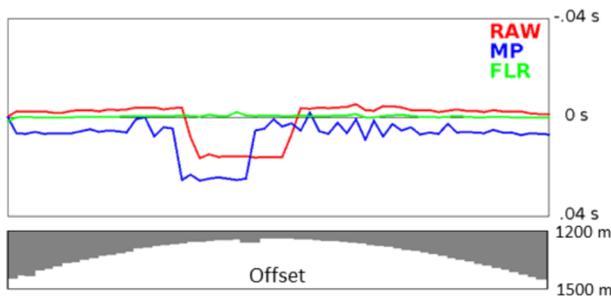


Figure 11 The difference between ZPR pick times and the pick times from the three other datasets from the previous figure. The RAW and MP datasets cycle skip in the noisy traces.

–10 ms at the far offsets. The difference a bulk shift of δt to the first-arrival times at far offsets makes to weathering statics in a two-layer weathering model is roughly $\delta t (v_1/v_r - 1)$, where v_1 is the second-layer velocity and v_r is the replacement velocity. Thus, compensating for Q can produce a significant bulk in statics when Q is small and $v_1 \ll v_r$. Neither is true, however, for this data set.

Figure 13 displays pick times for a single shot record, plotted as a function of offset for all four datasets. For display purposes, the picks have been partially flattened using a

single linear function. The picks for the RAW and MP datasets show clear signs of cycle skipping, particularly after 3000 m offset. The ZPR picks are the most consistent.

The obvious question is whether the ZPR and FLR picks are closer to the true first-arrival times than the others. It would seem so based on theory, but proving this on a real dataset is difficult. We can test, however, if the picks are more surface consistent. I derived weathering statics for each dataset using a well-known weathering interpretation package based on work by Hampson and Russell (1984). A feature of this package is that it uses the derived weathering model to predict back the first-arrival times, and then measures the difference between these and the input first-arrival times. High-quality picks should give a close fit. A table of the fitting errors is given in Table 1. Both the ZPR and FLR picks have about two-thirds the fit errors as the RAW and MP picks, suggesting they have done a better job at picking a consistent feature.

Another way of determining the quality of the picks is to derive weathering statics from them and see how well they remove the short-wavelet statics between adjacent first arrivals. Figure 14 shows flattened raw first arrivals for a single receiver, with weathering statics applied using the RAW and ZPR picks. On the cleaner traces (each trace representing a

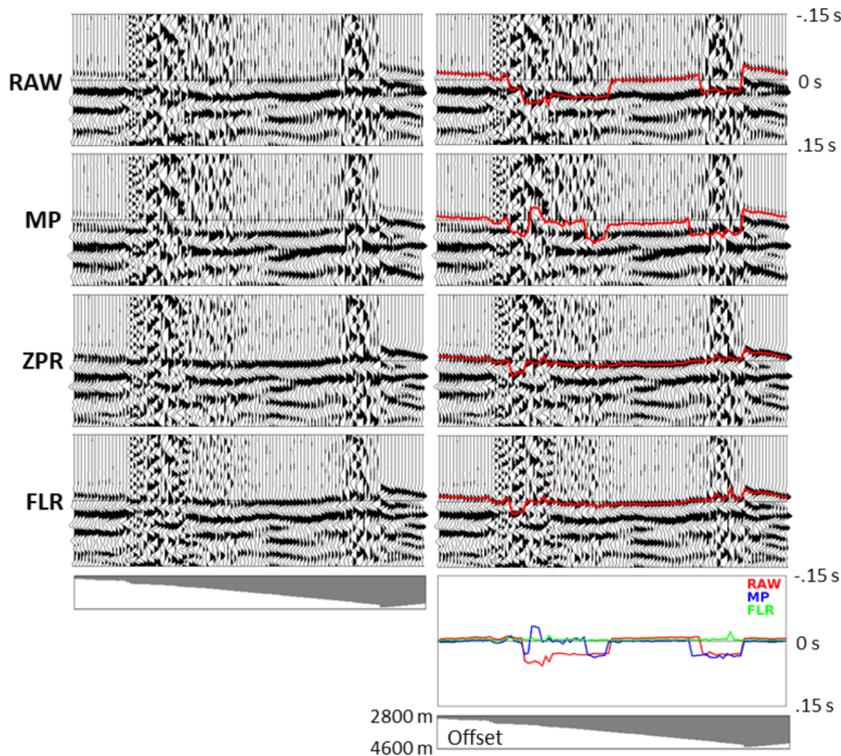


Figure 12 First arrivals from the four datasets at times between 1.6 and 2.4 s, flattened for display purposes using a static shift of $-0.00043 \times \text{offset} - 0.45$ s. The right displays are the same as the left, but with the picks shown. The bottom-right graph shows the difference between ZPR pick times and the other three. The RAW and MP datasets cycle skip badly in the noisy traces. The ZPR and FLR datasets have far less cycle skipping as the picked feature is more energetic, so it is less overwhelmed by noise.

Table 1 Mean absolute fit error between actual and reversed-modelled first-arrival times for each dataset. The two datasets shaped to a Ricker wavelet have lower fit errors, suggesting higher quality picks

Fit Errors of the Four Dataset Picks	
Dataset	Fit error (ms)
RAW	12.1
MP	13.7
ZPR	8.8
FLR	8.5

different shot), there is little difference between the two except for a small bulk shift. On the noisier traces, the ZPR weathering statics do a significantly better job of flattening the arrivals.

DISCUSSION

Picking vibroseis first arrivals is a long-standing problem that has been poorly addressed in the literature. Shaping the first arrival to a simple wavelet such as a zero-phase or four-loop Ricker can improve picking by removing ringy precursors and

by focusing the energy of the first arrival better. In theory, the picks should also be closer to the true first-arrival time, although that is difficult to verify in practice.

Shaping to a Ricker wavelet is not the only possible choice. Low-order Butterworth filters (Oppenheim and Schaffer, 2010, §7.3), for example, might also serve. The key is that the wavelet be simply shaped and have a narrow frequency band.

It may seem counterintuitive that we improved the pickability of the first arrivals by narrowing rather than broadening the amplitude spectrum of the seismic wavelet. But as Berkhout (1988) points out, the meaning of resolution depends entirely on what we are trying to resolve. Here we are trying to locate the time of a single strong event, with no events preceding it. This is quite different than trying to distinguish between many closely clustered events. Even when using a common definition of resolution such as short-wavelet length, Berkhout shows that it is not so much the bandwidth but the smoothness of the amplitude spectrum that dictates resolution. By shaping the amplitude spectrum so that the sharp cutoffs at the sweep limits are removed, I have shortened and simplified the seismic wavelet to reduce the chances of cycle skipping.

Figure 13 Pick times plotted as a function of offset for a single shot record for the four datasets. The pick times have been partially flattened using a linear function of $-0.5 \times \text{offset}$ s. The RAW and MP datasets show clear signs of cycle skipping, particularly beyond 3000 m offset. The ZPR dataset has the most consistent picks.

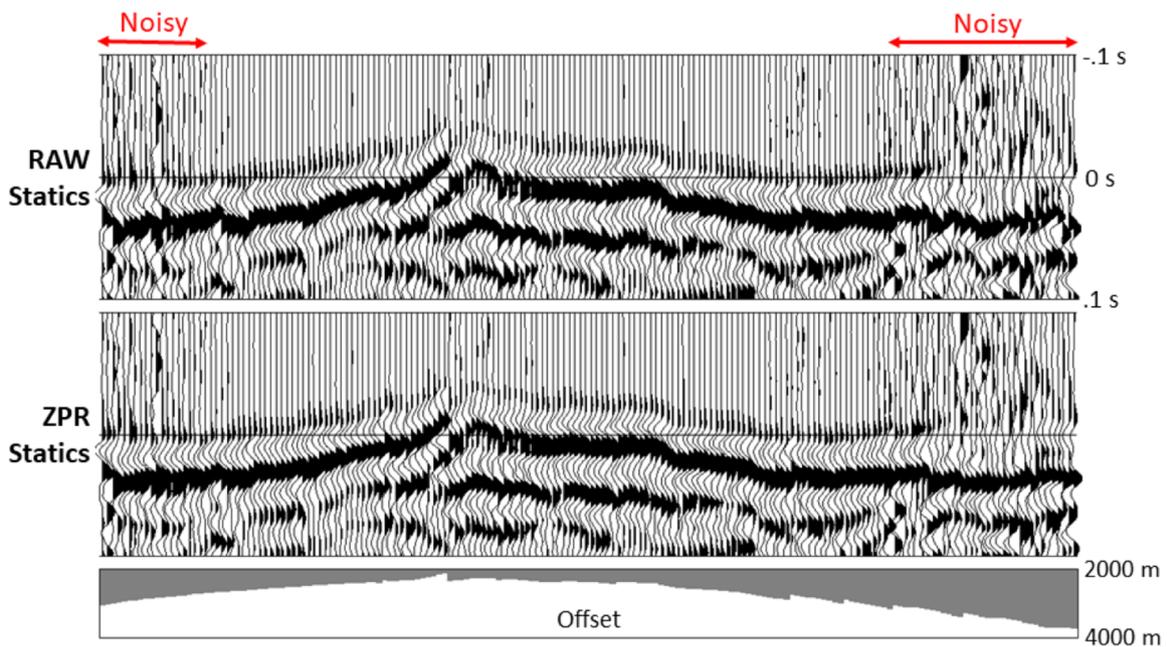
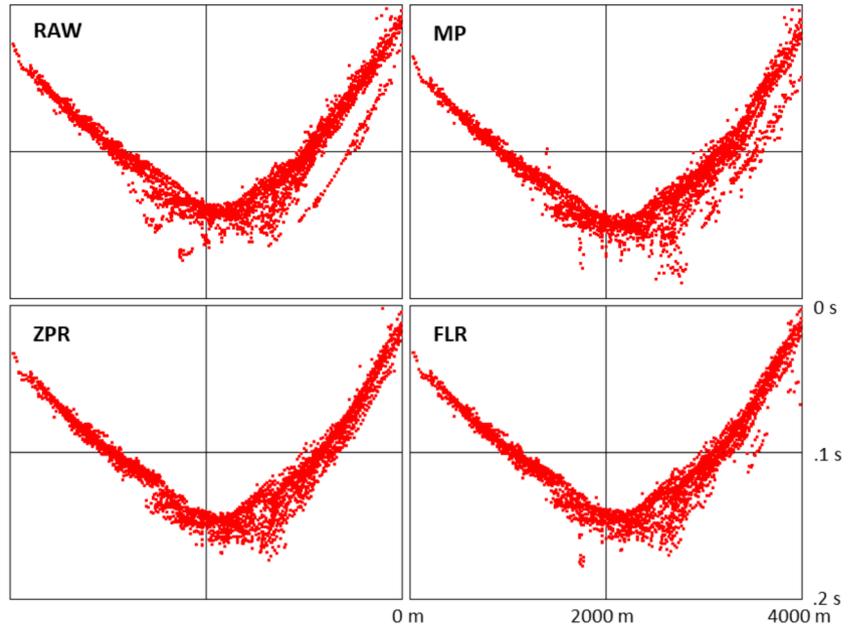


Figure 14 Raw first arrivals for a portion of a single receiver gather (each trace representing a different shot) with weathering statics applied based on RAW (top) and ZPR (bottom) picks. The picks have been partially flattened with a linear function of $-0.47 \times \text{offset} + 0.2$ s. The ZPR statics do a better job of flattening the first arrivals in the noisy shots, suggesting that it has produced a better weathering solution there.

Most industrial automatic pickers try to pick a consistent waveform feature without a serious attempt to find the true first-arrival time relative to it. The proposed method attempts to find the true arrival time, but requires knowing the amount

of Q attenuation undergone by each first arrival to be highly accurate. I have presented one simple method for estimating the attenuation, but improvements on the algorithm might be worthwhile. Such information is not just needed for the

methods described here; it would seem to be needed for any attempt to find the true first-arrival time.

One fly in the ointment is the issue of trace polarity. It is critical to the proposed method that the acquired data follow the SEG polarity standard (Landrum et al., 1994), but in practice this is not always so (Sakallioğlu, 2011). One might guess at the true polarity by examining the shape of the first arrivals, but this seems self-confirming in regard to our method. A more independent means of determining polarity from within the processing centre would be worthwhile.

We should not expect this method to work in every case. There are many effects that are not accounted for, such as estimated Klauder wavelets that poorly reflect the true applied ground force (Baeten and Ziolkowski, 1990), array effects (Vermeer, 1990), geophone ground coupling (Krohn, 1984) and shingling (Cassinis and Borgonovi, 1966). In most cases, however, it is surprising how closely the data match the modelled seismic wavelet, suggesting that the seismic wavelet is more knowable than is generally thought.

ACKNOWLEDGEMENTS

Thanks to Absolute Imaging Inc. for the use of their facilities, and to Jack Bouska, Samo Cilensek, Dave Bloxham and Dennis Quinn for their help and feedback.

DATA AVAILABILITY STATEMENT

Seismic data associated with this research are confidential and cannot be released.

REFERENCES

- Aki, K. and Richards, P.G. (2002) *Quantitative Seismology*, 2nd edition. Sausalito, CA: University Science Books.
- Baeten, G. and Ziolkowski, A. (1990) *The vibroseis Source*. New York: Elsevier.
- Boschetti, F., Dentith, M.D. and List, R.D. (1996) A fractal-based algorithm for detecting first arrivals on seismic traces. *Geophysics*, 61, 1095–1102.
- Berkhout, A.J. (1988) *Seismic Resolution: A Quantitative Analysis of Resolving Power of Acoustical Echo Techniques*. London: Geophysical Press.
- Cassinis, R. and Borgonovi, L. (1966) Significance and implications of shingling in refraction records. *Geophysical Prospecting*, 14(4), 547–565.
- Coppens, F. (1985) First arrival picking on common-offset trace collections for automatic estimation of static corrections. *Geophysical Prospecting*, 33(8), 1212–1231.
- Cox, M. (1999) *Static Corrections for Seismic Reflection Surveys*. Tulsa, OK: SEG.
- Ellison, D., Innanen, K. and Cameron, G. (2017) Improved resolution in depth imaging through reflection static corrections derived from model-based moveout. 87th Annual International Meeting, SEG, Expanded Abstracts.
- Futterman, W.I. (1962) Dispersive body waves. *Journal of Geophysical Research*, 67(13), 5279–5291.
- Gardner, L.W. (1939) An areal plan of mapping subsurface structures by refraction shooting. *Geophysics*, 4(4), 247–259.
- Gibson, B. and Larner, K. (1984) Predictive deconvolution and the zero-phase source. *Geophysics*, 49(4), 379–397.
- Hampson, D. and Russell, B. (1984) First-break interpretation using generalized linear inversion. 39th Annual International Meeting, SEG, Expanded Abstracts.
- Hart, D.I., Hootman, B.W. and Jackson, A.R. (2001) Modeling the seismic wavelet using model-based wavelet processing. 56th Annual International Meeting, SEG, Expanded Abstracts.
- Hatherly, P.J. (1986) Attenuation measurements on shallow seismic refraction data. *Geophysics*, 51(3), 250–254.
- Hons, M., Stewart, R., Lawton, D., Bertram, M. and Hauer, G. (2008) Accelerometer vs. geophone response, a field case history. CSPG/CSEG/CWLS GeoConvention.
- Hosken, J.W.J. (1988) Ricker wavelets in their various guises. *First Break*, 6(1), 24–33.
- Kobayashi, Y. (2001) Delay of first arrival caused by dispersion of seismic body waves. 56th Annual International Meeting, SEG, Expanded Abstracts.
- Krohn, C.E. (1984) Geophone ground coupling. *Geophysics*, 49(6), 722–731.
- Landrum, R.A., Brook, R.A. and Sallas, J.J. (1994) Polarity convention for vibratory source/recording systems. *Geophysics*, 59(2), 315–322.
- Ma, Y., Cao, S., Rector, J.W. and Zhang, Z. (2020) Automated arrival-time picking using a pixel-level network. *Geophysics*, 85, V415–V423.
- Oppenheim, A.V. and Schaffer, R. W. (2010) *Discrete-Time Signal Processing*, 3rd edition. Upper Saddle River, NJ: Prentice Hall.
- Quan, Y. and Harris, J.M. (1997) Seismic attenuation tomography using the frequency shift method. *Geophysics*, 62(3), 895–905.
- Ricker, N. (1940) The form and nature of seismic waves and the structure of seismograms. *Geophysics*, 5(4), 348–366.
- Ristow, D. and Jurczyk, D. (1975) Vibroseis deconvolution. *Geophysical Prospecting*, 23(2), 363–379.
- Sakallioğlu, Y. (2011) Field Testing Seismic Phase and Polarity – Vibratory Source Signature and Recording Systems. 6th Congress of Balkan Geophysical Society, Budapest, Hungary.
- Vermeer, G.J.O. (1990) *Seismic Wavefield Sampling*. Tulsa, OK: Society of Exploration Geophysicists.
- Wiggins, R.A., Larner, K.L. and Wisecup, R.D. (1976) Residuals statistics analysis as a general linear inverse problem. *Geophysics*, 41, 922–938.
- Yang, D., Liu, J. and Li, D. (2020) Q-factor estimation using bisection algorithm with power spectrum. *Geophysics*, 85(3), V233–V248.

APPENDIX A

THE FOUR-LOOP RICKER WAVELET

Here I derive the four-loop Ricker wavelet briefly referred to by Hosken (1988). From Equations (1) and (2), the standard three-loop zero-phase Ricker wavelet has an impulse response of

$$(1 - 2a) \exp(-a), \text{ where } a = (\pi f_p t)^2, \quad (\text{A1})$$

where f_p is the peak frequency, and a continuous amplitude spectrum of

$$\frac{2 f^2}{\sqrt{\pi} f_p^3} \exp\left(-\frac{f^2}{f_p^2}\right). \quad (\text{A2})$$

To convert to discrete amplitude spectra (that is, for DFTs), an additional scaling of twice the Nyquist frequency is required, assuming the most common scaling convention for DFTs. As well, time and frequency must have consistent units. If time is in seconds, for example, frequency must be in cycles per second (Hertz).

For now I will drop constant-scaling terms, as the proper scaling will be addressed at the end. The four-loop Ricker wavelet is proportional to the temporal derivative of Equation (A1):

$$t (2a - 3) \exp(-a). \quad (\text{A3})$$

Its phase spectrum is 90° . Its amplitude spectrum is proportional to Equation (A2) multiplied by f :

$$f^3 \exp\left(-\frac{f^2}{f_p^2}\right). \quad (\text{A4})$$

The peak frequency f_p of the zero-phase Ricker wavelet will not be the peak frequency of its temporal derivative. Find the peak frequency g_p of Equation (A4) by setting its spectral derivative to zero, giving:

$$\frac{2}{3} g_p^2 = f_p^2. \quad (\text{A5})$$

Substituting Equation (A5) into (A3) results in the four-loop Ricker wavelet parameterized by its peak frequency:

$$t (2b - 3) \exp(-b), \text{ where } b = \frac{2}{3} (\pi g_p t)^2 \quad (\text{A6})$$

The amplitude spectrum is now proportional to

$$f^3 \exp\left(-\frac{3 f^2}{2 g_p^2}\right). \quad (\text{A7})$$

The time of the maximum value of this wavelet is determined by setting the temporal derivative of Equation (A6) to

zero, resulting in four roots. The negative root nearest to zero is the time t_p of the global maximum:

$$t_p = -\frac{\sqrt{3(3-\sqrt{6})}}{2\pi g_p} \approx -0.2045 / g_p. \quad (\text{A8})$$

By dividing (A6) by its value at t_p , we have the final form of the four-loop Ricker wavelet scaled so that its maximum value is one:

$$s_t t (2b - 3) \exp(-b), \text{ where } s_t = \frac{\pi e^\beta g_p}{3\sqrt{\beta}} \text{ and}$$

$$\beta = \frac{3 - \sqrt{6}}{2}. \quad (\text{A9})$$

Its correctly scaled amplitude spectrum is

$$s_f f^3 \exp\left(-\frac{3 f^2}{2 g_p^2}\right), \text{ where } s_f = \frac{3^{3/2} e^\beta}{\sqrt{8\pi\beta} g_p^4}. \quad (\text{A10})$$

APPENDIX B

ESTIMATING Q ATTENUATION

The wavelet shaping method described in this paper requires an estimate of the amount of anelastic attenuation (Aki and Richards, 2002, §5.5) that each trace's first arrival has suffered. Assuming a single physically reasonable value for Q (30, for example) for all traces works quite well, but measuring Q as a function of the expected first-arrival time should result in more accurate first-arrival picks, and thus a better near-surface model and better weathering statics. Yang et al. (2020) provide an extensive list of Q estimation methods, but I know of only one paper by Hatherly (1986) which is specifically for first arrivals. Here I describe a novel method for first arrivals inspired by the reflection Q estimation technique of Quan and Harris (1997).

Consider a short time window (400 ms, for example) of a trace centred at τ , a rough estimate of the trace's first-arrival time. Assuming the constant Q (that is, frequency-independent Q) attenuation model of Futterman (1962), its amplitude spectrum can be modelled as

$$w(f) = s p(f) r(f) \exp\left(-\frac{\pi \tau f}{Q}\right), \quad (\text{B1})$$

where s is a scaler accounting for the magnitude of the energy source, geometrical spreading, and other effects, $p(f)$ is the product of the Klauer, differential and geophone amplitude responses, $r(f)$ is the amplitude response of the reflection coefficients within the window, and Q is the 'rock quality factor' describing the average rate of anelastic attenuation. This last value is what we wish to estimate.

Table B1 Attenuation coefficient α versus centroid frequency f_c in Hertz for a seismic wavelet with a linear sweep and 10 Hz critically damped velocity geophones. Such tables can be used to estimate the amount of Q attenuation of first arrivals

Attenuation coefficients versus centroid frequency					
α	f_c	α	f_c	α	f_c
-0.05	75.6	0.05	39.2	0.15	20.8
-0.04	73.6	0.06	35.7	0.16	20.2
-0.03	71.2	0.07	32.7	0.17	19.6
-0.02	68.4	0.08	30.1	0.18	19.1
-0.01	65.0	0.09	28.0	0.19	18.6
0.00	61.1	0.10	26.2	0.20	18.2
0.01	56.8	0.11	24.7	0.21	17.9
0.02	52.2	0.12	23.5	0.22	17.6
0.03	47.6	0.13	22.5	0.23	17.3
0.04	43.2	0.14	21.6	0.24	17.0

Assuming that the reflection-coefficient amplitude response is white (that is, $r(f) = 1$) gives

$$w(f) \propto p(f) \exp\left(-\frac{\pi \tau f}{Q}\right). \tag{B2}$$

The centroid frequency f_c (Quan and Harris, 1997) of a discrete amplitude spectrum $a(f)$ is defined as

$$\frac{\sum_{f=f_1}^{f_2} f a(f)}{\sum_{f=f_1}^{f_2} a(f)}, \tag{B3}$$

where f_1 and f_2 are the start and end frequencies of the signal band. Values just within the interior of the sweep band are usually good. The centroid frequency is a measure of the dominant frequency; generally the higher the anelastic attenuation (that is, the lower the value of Q), the lower the centroid frequency.

We begin by creating a table that lists the centroid frequency values of the amplitude spectra

$$p(f) \exp(-\alpha f) \tag{B4}$$

for a series of attenuation coefficients $\alpha = \pi \tau / Q$; values of -0.05 to 0.25 by 0.01 (for time units in seconds) are reasonable. Unless the vibroseis sweep is highly nonlinear, the centroid frequency will decrease monotonically with α . A table using typical acquisition parameters is shown in Table B1.

For a given trace, find its centroid frequency within a small window (e.g. 400 ms wide) centred roughly about the first-arrival time. This window should be heavily tapered so as to avoid windowing effects. From this centroid frequency, look up its α value from the table, using linear interpolation for values that fall between tabulated frequency values. This value will not usually be accurate due to noise and the fact

that the reflection-coefficient response is rarely white. But if we plot these values for a large number of traces over a wide range of expected arrival times then a pattern should emerge, and we can fit a simple and physically reasonable function $\tilde{\alpha}(\tau)$ to them (Fig. 4).

Finally, for a given trace with an expected first-arrival time τ , shape it to a Ricker wavelet using

$$Q = \frac{\pi \tau}{\tilde{\alpha}(\tau)}. \tag{B5}$$

Although for most surveys a single simple $\tilde{\alpha}(\tau)$ function should do, it is possible that some projects may require a function that varies with geographical location as well as expected arrival time.

APPENDIX C

SPATIALLY CONTINUOUS PICKING

First-arrival pickers play a central role in the examples, so a brief discussion of them seems warranted.

Many automatic first-arrival pickers determine picks on a trace-by-trace basis, meaning each trace is picked independently of the others. Such pickers generally try to find a time point where some trace attribute or group of attributes changes abruptly in character. The most common and successful of these attributes is the energy level of the samples (Coppens, 1985), but more exotic ones like fractal dimension (Boschetti et al., 1996) can be used as well.

Trace-by-trace pickers can work fairly well for dynamite data (more accurately, data with an explosive source), where the first arrivals tend to be cleaner, and where there is an abrupt onset of source energy on each trace. But on vibroseis

data these pickers often work poorly. This is because vibroseis first arrivals tend to be noisier and have a broader and non-causal seismic wavelet, resulting in a less abrupt change in trace attributes near the first-arrival time. Indeed, as pointed out in the introduction, there is no true onset of source energy on a vibroseis trace.

To pick vibroseis data well, I suggest that a picker must consider more than one trace simultaneously. One strategy is to encourage spatially continuity of the picks – that is, after linear-moveout correction, there should be few or no large time differences between nearby traces. There are surprisingly few papers that describe such strategies; Ma et al. (2020) is one example.

Describing in detail the spatially continuous picker used in this paper's examples would be lengthy and beyond our scope. A brief description is as follows. For a single shot record:

1. Determine the times of all peaks or troughs near the expected first-arrival time for each trace.
2. From these set of times, choose one or more 'candidate picks' for each trace, where these candidates display a significant change in one or more attributes such as energy.
3. From these candidate picks, determine the final picks by finding a shortest path in the trace direction through the candidate picks.

The shortest path will depend on how the distance between two candidate picks of adjacent traces is calculated. It should mostly be based on their absolute time difference (after linear-moveout correction), which favours spatially continuous picks. But countless games can be played with the distance, such as taking into consideration the energy ratios of the two candidates.