Tutorial: Good practice in well ties

Roy White[1] and Rob Simm[2]

Well ties are a very important part of the interpreter's trade. They provide a means of 1. correctly identifying horizons to pick, and 2. estimating the wavelet for inverting seismic data to impedance. Just as well ties are paramount in the calibration of a seismic interpretation, so too they are the cornerstone of using seismic amplitudes in impedance and AVO inversion, and ultimately of inferences fed into the risking process. Figure 1 shows an example where the geologically relevant amplitude information is highly focussed in a particular seismic loop (the blue loop in this case). The wavelet is ~30 degrees rotated from symmetry and its main loop is 24 ms deeper in time than time zero. This is nominally minimum phase data displayed with European polarity, i.e. with negative values corresponding to compressions and plotted as troughs. Given this convention, together with the fact that the red loop is closest to the checkshot time, some interpreters might expect that the red pick is the one to make. If you are after the amplitudes that map out the sedimentology, then it clearly isn't. Once you know the wavelet, it is evident that the wavelet energy is concentrated in the blue loop and this loop is delayed relative to the checkshot time. Readers are referred to our previous tutorial article (Simm and White 2002) for a discussion of phase, timing, polarity and the interpreter's wavelet and their impact on seismic interpretation.

What is important in this example is that both the wavelet shape and timing were estimated without making assumptions about the wavelet (i.e. what the wavelet should look like) or the timing (i.e. which loop represents the top of the reservoir).

The subject of this paper then is the process by which wavelets are estimated through a well tie procedure that results in quantitative measures of synthetic to seismic goodness-of-fit and likely wavelet accuracy. In the authors' view this is a good practice approach, which should form at the very least the initial stages of a well tie study. Activities such as stretch and squeezing a synthetic to fit a seismic trace are unscientific and definitely not good practice! Regrettably, while 'stretch and squeeze' is a common feature of well tie software, q.c. features are at best rudimentary.

The well tie process

Making a well tie (Figure 2) is deceptively simple. A synthetic seismogram is matched to a real seismic trace and features from the well are correlated to the seismic data. The prime concept in constructing the synthetic is the convolutional model. This represents a seismic reflection signal as a sequence of interfering reflection pulses of different amplitudes and polarity but all of the same shape. This pulse shape is the seismic wavelet, formally the reflection waveform returned by an isolated reflector of unit strength at the target depth. Because reflecting boundaries are spaced much more finely in depth than the length of the reflected pulses, the degree of interference is generally severe and only the strongest reflectors or reflection complexes stand out in the reflection signal.

Well tie procedure

In outline the procedure for tying a well-log synthetic seismogram to seismic data comprises the following steps (Figure 3).

1. Edit and calibrate the sonic and density logs.
2. Construct the synthetic seismogram from the calibrated well-logs:
   a. choose the appropriate reflection series (usually primaries only),
   b. construct the reflection series in two-way time.
3. Perform the match, comprising
   a. determine the best match location,
   b. estimate the wavelet and its accuracy.

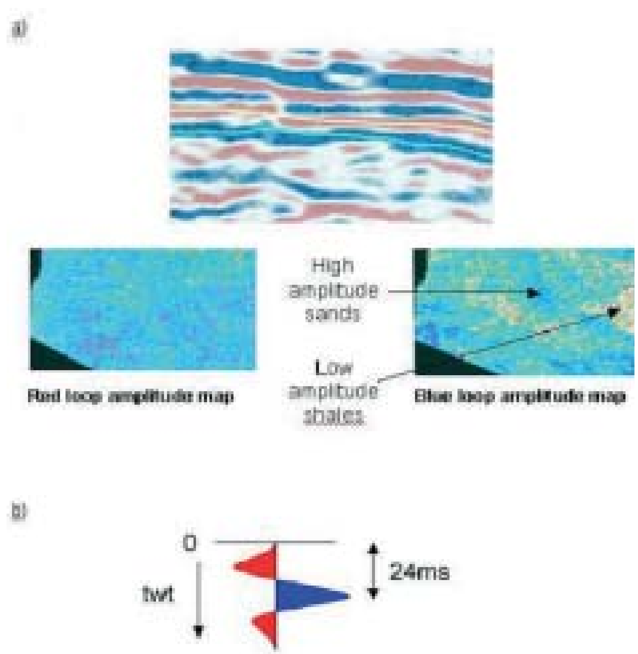

Figure 1 *Making the wrong pick can make a difference to geological interpretation: a) horizon picks and amplitude maps, b) the wavelet*

[1] *School of Earth Sciences, Birkbeck, University of London (roy.white@ucl.ac.uk)*
[2] *Rock Physics Associates Ltd. (rob.simm@rock-physics.com)*

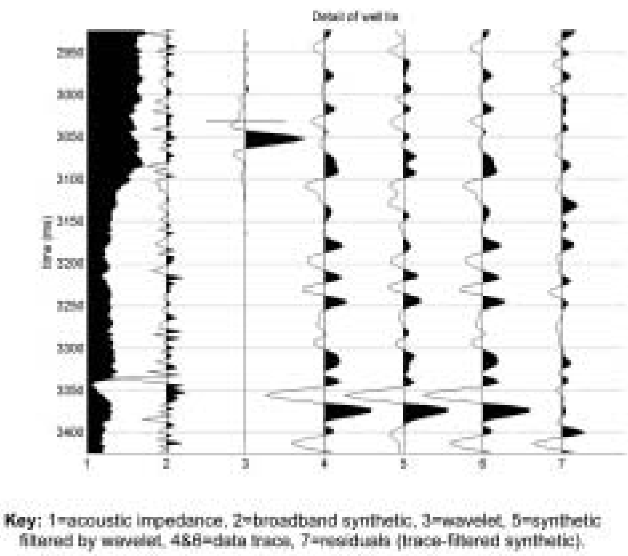

Figure 2 *The components of a well tie. Time zero of the wavelet is indicated by the horizontal line.*

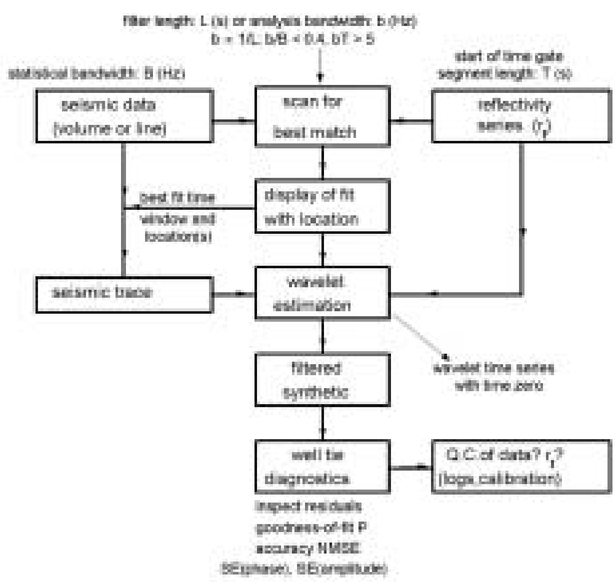

Figure 3 *Well tie procedure, indicating the key parameters required by the process.*

In some situations there is also the issue of comparing wavelets estimated from different wells in the survey area and combining them if appropriate.

Log calibration

In the interests of brevity we shall omit any discussion of editing and conditioning the logs. This is not to deny its importance but it is too big a topic to be satisfactorily reviewed here.

Well log calibration brings the timing of the sonic log into agreement with seismic times from a checkshot survey or Vertical Seismic Profile (VSP). Practical and theoretical analysis of the factors that influence the accuracy of well ties (White 1997) shows that timing is paramount. Timing errors in the synthetic seismogram are much more detrimental than amplitude errors. They are especially damaging to estimating correctly the higher frequency components of the seismic wavelet. Well-log calibration establishes the timing of the synthetic seismogram. Without it, the tie can only be fudged.

Calibration obviously starts with picking checkshot times. Current practice with regard to picking checkshot times seems to favour trough-to-trough picking of VSP times in which times are picked from the first trough of the signal on the source hydrophone to the first trough of the downhole signal. Some geophysicists have advocated picking times from the onsets of these signals. A correction for dispersion-attenuation may then be necessary as the onsets tend to travel at the speed of the highest frequency components in the seismic waveform. The first-order effect of dispersion-attenuation in the earth at common target depths is a slight stretching of travel times and trough picking can be regarded as a way of compensating for this effect (White 1998).

Nowadays the interpretation geophysicist is likely to receive either calibrated sonic and density logs or times picked and corrected by the VSP contractor. In the latter case, a calibration (or drift) curve is fitted to the difference between the integrated sonic log times and the corrected VSP times, as illustrated in Figure 4. In this case the calibration uses straight line segments between knee points. Knee points have to be chosen at major jumps in the sonic log so as to avoid introducing artificial or abnormal reflection coefficients. A popular alternative is a smooth spline fit which provides smoothly varying corrections to the log. Spline calibration requires a choice of spline knots and this can be automated. Nonetheless calibration points often show sudden bends at unconformities and changes in character of the logs and knee points allow direct control of the calibration. In practice the density of VSP recordings usually offers very little scope to vary the calibration curve significantly and any reasonable fit to this many points generates a satisfactory timing.

The drift curve is applied to the sonic log which is then integrated to give a time-depth relationship that fits the checkshots to within 1-2 ms. The picking accuracy of VSP times is of the order of 1 ms and the calibration should constrain the timing of the synthetic seismogram to a similar accuracy. Stewart et al (1984) discuss the causes of different types of drift and their physical explanation.

Constructing the synthetic seismogram

Constructing the synthetic involves creating a reflection series in time and convolving this with a wavelet. There is also the question as to whether to use a primaries-only reflectivity or whether to include internal multiples or even to allow for other multiples. In most cases it is the correlation of the data with the primary reflectivity that interpretation requires and this usually produces a satisfactory well tie.

a) Creating a reflection series in time

There is no single approach to creating a broadband reflection series in time. Some modern software gives the option of either using a velocity log with the effect of drift applied (i.e. the sonic log values are effectively changed) or using the original sonic values together with the time-depth relationship constructed from applying the drift curves to the sonic log (i.e. the depths are relabelled in time). The conversion to time obviously involves a re-sampling step. In order to mitigate any aliasing effects the initial sampling gives a sample rate finer than the initial depth sampling. These days there is no reason why the synthetic trace cannot be computed from this finely sampled log and a wavelet with the same sample rate. However the synthetic seismogram generally benefits from some blocking of the logs. In the past, limited computing power may have dictated further sample coarsening, commonly to match the seismic sample rate.

This coarsening of the sample rate is effectively a time averaging process. Theoretically, in order to account for the effects of velocity dispersion between sonic and seismic frequencies the sonic logs need to be averaged using the Backus average rather than the time-average approach (Marion et al 1994, Rio et al.1996). In this approach velocities and densities are used to derive the bulk and shear moduli, which are then averaged using the harmonic average. Velocities are re-constituted from the averaged moduli and the arithmetically averaged density log. Experience has shown that in sand/shale environments there is very little difference between well ties made using the time-average and the Backus average approach. It is usually safe to block to a fine sample interval (0.5 ms say) and then resample the synthetic to the required sample interval. If there are large contrasts in velocity then the Backus average should be used and in this case it may pay to optimise the blocking length.

For well ties on angle stacks, reflection series need to be computed for a range of angles of incidence or recording offsets. Converting the logs to elastic impedance (Connolly 1999) offers a quick and convenient way of computing angular reflectivities using standard normal incidence code. This procedure becomes seriously inaccurate on going from siliciclastic to carbonate lithologies or at any sudden change in the background V_s/V_p ratio. Alternatively computing angle or offset-dependent primary reflection series is a straightforward task.

b) Estimating the wavelet to be convolved with the reflection series

The broadband reflectivity has to be filtered by the seismic wavelet in order to tie the well to the data. A problem in estimating the wavelet is that the transmission response of the earth has to be included if the filtered synthetic seismogram is to match the seismic data. A matching process is therefore required. There are a number of approaches to estimating the wavelet:

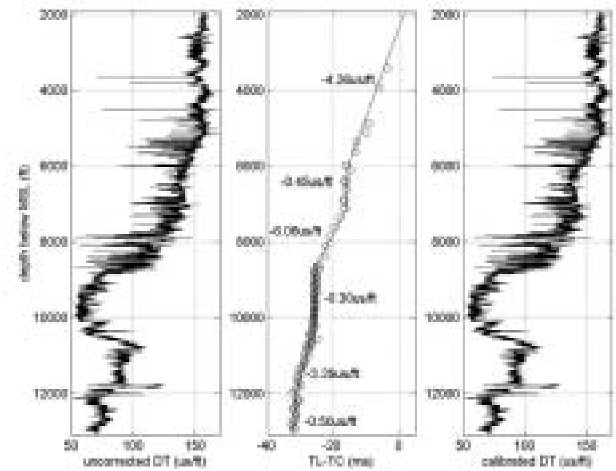

Figure 4 *An example of well-log calibration from the central North Sea. Left: uncalibrated sonic log; centre: integrated log time minus check-shot time fitted using knee points; right: calibrated sonic log. The slopes of the calibration curve (centre panel) provide the correction factors subtracted from the uncalibrated log.*

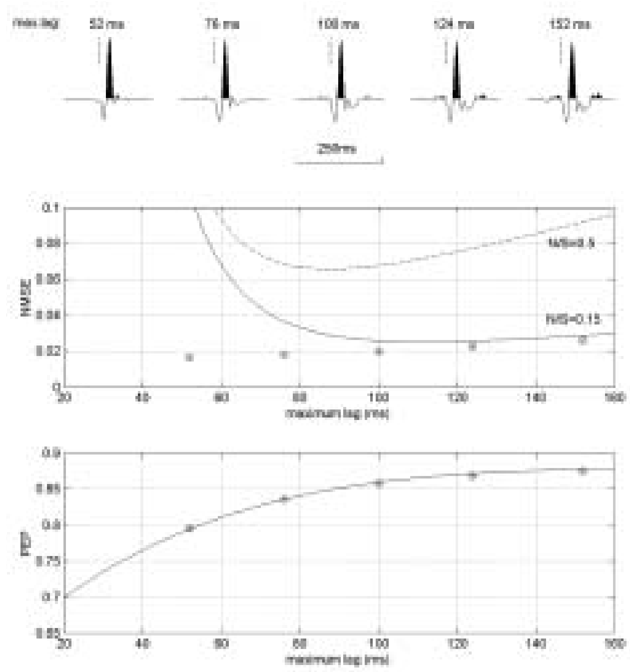

Figure 5 *The trade off between bias (distortion) and noise-induced errors in matching a synthetic seismogram and seismic trace. The top panel shows wavelets extracted from the data shown in Figure 6 using various maximum correlation lags; the dashed vertical lines mark zero time. The centre panel plots the computed NMSE of these wavelets against maximum lag; the solid curve is a theoretical fit to these values using a N/S ratio of 0.15; the dashed curve is for a N/S ratio of 0.5. The bottom panel plots the predictabilities (P) fitted by a smooth curve.*

1. Forward modelling approach

One approach is to build a model of the seismic wavelet from the response of the acquisition and processing system, adding perhaps a constant-Q filter to represent attenuation in the earth. Constant-Q attenuation is at best a simple approximation to the characteristic amplitude decay with frequency from transmission through the earth and it may or may not approximate the concomitant phase and timing response. Ghosting, the effect of processing and an 'off-the-shelf' source signature inject further significant uncertainties into this approach. After stacking and migration the wavelet will no longer be a causal or even a physical wavelet. Without some further adjustment, the well tie is likely to be crude and sub-optimal.

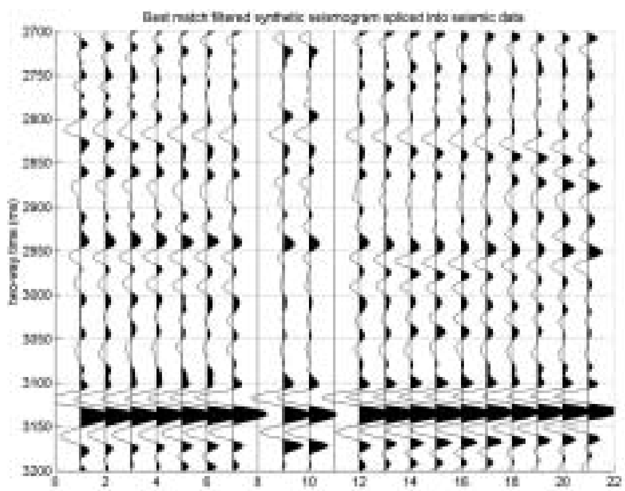

Figure 6 *The well tie using the 124ms wavelet in Figure 5. Two copies of the synthetic seismogram are spliced into the seismic data at traces 9 and 10. The best matching trace is repeated on each side of the synthetic seismogram at traces 7 and 12. The best matching trace was found by scanning a cube of data around the well (Figure 7).*

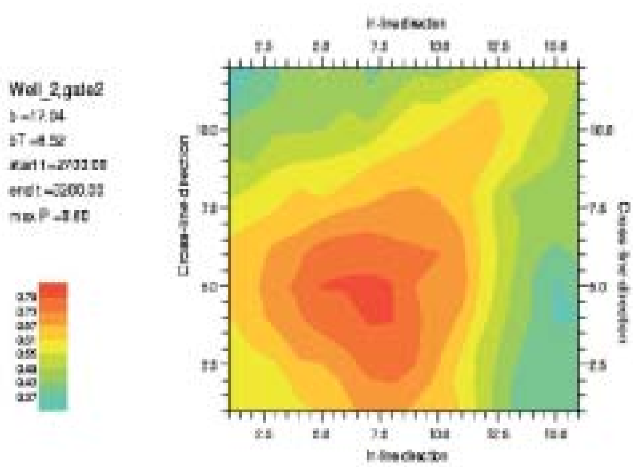

Figure 7 *Contours of goodness-of-fit from matching a 2.7-3.2s segment of well-log synthetic seismogram with a cube of data centred on the well. The goodness-of-fit is measured by P, the proportion of trace energy in the seismic trace predicted by the match, which is approximately the square of the cross-correlation coefficient between the synthetic seismogram and data.*

2. Phase rotation of zero phase synthetic

A more straightforward alternative is to construct a zero-phase wavelet that matches the energy spectrum of the data window and phase rotate and time shift it to optimise the match. This can be a very effective way of making a character match of synthetic to seismic. The reason it works is that, after processing, the phase of the seismic wavelet is usually fairly linear across the seismic bandwidth. Deviations from linearity at the edges of the bandwidth are likely to matter only if the wavelet is to be used for inversion. Owing to differences in spectral character between Ricker wavelets and real seismic data, Ricker wavelets should not be used for the initial zero phase synthetic (Hosken 1988).

3. Wavelet estimation through matching

The method that is recommended here and that is described in the following sections is to estimate the wavelet through a coherency matching technique (White 1980,Walden and White 1998). This procedure gives a number of outputs that effectively define whether the tie is good or not.

a. the shape of the wavelet
b. the phase characteristics of the wavelet
c. quantitative measures related to the tie
   i. an estimate of the goodness-of-fit of the synthetic to seismic
   ii. an estimate of the likely phase error (or accuracy) of the wavelet

Measuring the tie: goodness of fit and accuracy

The distinction between goodness-of-fit and accuracy is a very important one in assessing the reliability of the tie. In order to define goodness-of-fit and accuracy, we introduce two terms:

1. the *energy* of a trace is the sum of squares of a segment of a time series
2. the *residuals* are the difference between a seismic trace and its matched or filtered synthetic seismogram.

A simple measure of goodness-of-fit is the proportion of trace energy predicted by the synthetic seismogram, which we can call the predictability P:

$P = 1 -$ (energy in the residuals/trace energy).

The correlation coefficient R is another goodness-of-fit measure but implicit in it is the assumption of an error-free synthetic seismogram. If the assumption holds, then $P = R^2$. Many statisticians consider R^2 a less deceptive measure of correlation than R.

To measure accuracy we start with the normalised mean square error in the synthetic seismogram:

NMSE = (energy in the errors in the synthetic)/(energy in the synthetic)

Since the 'true' synthetic is an abstraction, we cannot know what the errors are but estimation theory tells us what the expected *NMSE* is, namely:

$$E\{NMSE\} \approx \frac{1}{bT}\frac{1-P}{P}$$

T is the duration of the matched data segment and b is called the analysis bandwidth (White 1980). bT is dimensionless. In frequency domain terms, b is the bandwidth of a moving spectral window within which the frequency response is analysed. In time domain terms, b is a constant divided by the maximum correlation lag employed in the analysis. The maximum lag is a proxy for wavelet length. Thus the *NMSE* is proportional to the ratio of wavelet length to data segment length. $(1-P)/P$ approximates the ratio of misfit energy to fitted energy; and it depends on the S/N ratio of the seismic data and the accuracy of the synthetic model built from the calibrated well logs.

Simulations show that this approximation measures the *NMSE* in the synthetic seismogram rather well. If the reflectivity spectrum is fairly flat, it also estimates the *NMSE* in the wavelet. However it measures only the misfit from noise in the seismic data and errors in the synthetic. If the analysis bandwidth b is too broad, or equivalently the maximum lag is made too short to accommodate the wavelet, the estimated wavelet will be distorted. On the other hand, if the maximum lag is too long, the tie will be overfitted, matching seismic noise as well as reflection signal.

The *NMSE* is a somewhat crude measure of wavelet accuracy. Error bars on the amplitude and phase spectra of the wavelet are more informative. These are illustrated later in the section on well ties and seismic data quality. A rough idea of the phase error within the seismic bandwidth can be found from

$$S.E\{phase\} \approx \sqrt{NMSE/2} \ \text{radians}$$

This equation also approximates the relative standard error in the amplitude spectrum.

A practical example

Figure 5 illustrates a range of wavelets which have been extracted from the data shown in Figure 6. At a wavelet length of 52 ms the wavelet is compressed and has limited side lobe energy whereas at 152 ms the wavelet has well developed side lobes with secondary oscillations. As wavelet length increases so does the goodness-of-fit (bottom panel, Figure 5). The accuracy of the wavelet as measured by the NMSE is shown by the points in the middle panel of Figure 5. The solid curve in this panel is a theoretical estimate of the NMSE when wavelet distortion is included, for a noise-to-signal power ratio of 0.15, corresponding to a predictability P of 0.87 (R≈0.93). The minimum in the curve is very flat and any choice of maximum lag (or analysis bandwidth) that avoids distortion yields a satisfactory wavelet estimate. The synthetic seismogram in Figure 6 used the wavelet from the 124 ms maximum lag.

NMSE increases for wavelets shorter than the optimum maximum lag owing to the effects of distortion or bias. It increases for wavelets longer than the optimum through the inclusion of noise into the wavelet estimation. In practice since distortion rises rapidly at low maximum lags, it pays to err towards larger not smaller values. Lengthening the wavelet may make the side lobes look noisier but bias distorts the main lobe which is more serious since it carries most of the wavelet energy. Unfortunately simply inspecting a wavelet estimate does not necessarily reveal that it is seriously distorted, although there may be indications from inspecting its spectrum. If wavelet length is a problem, there is no substitute for simulation. Automated criteria which take no account of spectral distortion, such as the Akaike Information Criterion, do not optimise wavelet length (Walden and White 1984). The dashed curve shows how the minimum in the total *NMSE* would move to a smaller maximum lag if the N/S ratio increased to 0.5.

The mechanics of matching: finding the best match location

Matching is usually performed on time migrated seismic data. A volume of data around a well is scanned for the best match location. For various reasons, the best match often does not occur at the well location. Since velocity typically increases with depth, time migration commonly moves the best match location up-dip from the well. The effect of imperfect migration velocities is less predictable and it is even possible that the well location is suspect. During the scanning

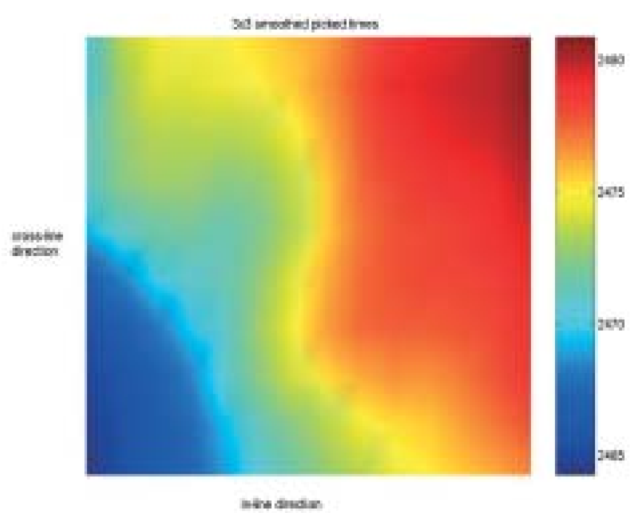

Figure 8 *Contour plot of two-way time to a prominent reflection at ~2.5 s from the cube of data scanned in Figure 7. The up-dip direction is from NE to SW.*

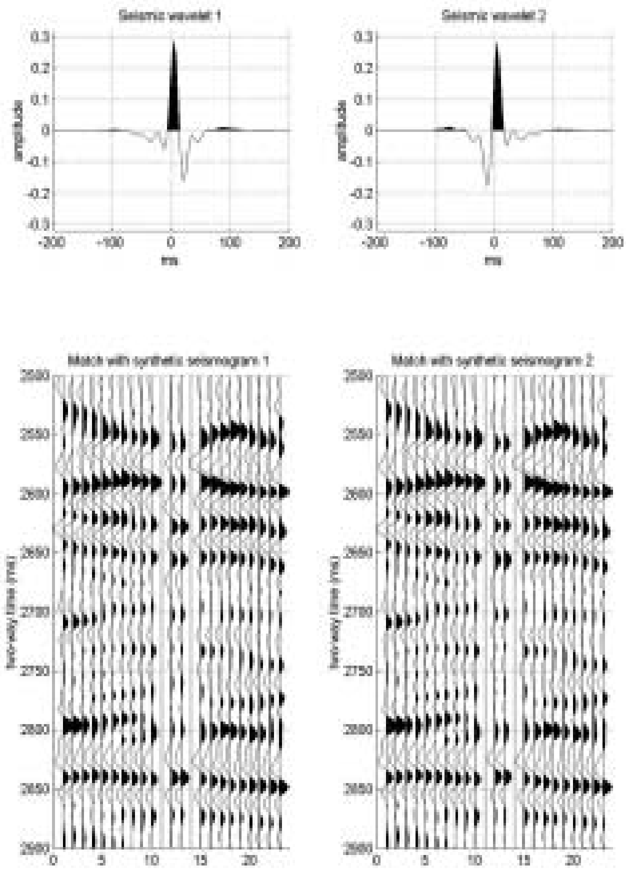

Figure 9 *Which is the better well tie? Two equally good well ties from a time interval showing poor seismic bandwidth. The well ties were constructed from a calibrated impedance log using the two wavelets shown.*

procedure the choice of analysis parameters in estimating the wavelet can be investigated.

The best match time location is found by scanning a number of traces around the well for goodness-of-fit using, say, a 500 ms segment of synthetic seismogram. Over such a segment it can be generally assumed that phase is not varying significantly. Reliability of the procedure deteriorates rapidly if shorter segments are used. When more than 500 ms of calibrated sonic is available the best time gate in which to establish the tie should also be investigated. Figure 7 shows contours of goodness-of-fit from matching a 2.7-3.2 s segment of well-log synthetic seismogram with a cube of data centred on the well. The best match location is about 100 m SW of the well. A contour plot of two-way time to a prominent reflection above the matched time gate (Figure 8) confirms that the overlying strata dip upwards from NE to SW. Good well ties reveal this up-dip pattern. When tying wells to poor seismic data the best match location may be pulled around by variations in S/N ratio. The well tie at the best match location of Figure 7 is shown in Figure 6.

Well ties and seismic data quality

The quality of the seismic data has an impact on our confidence in well ties and wavelet extractions. This is illustrated with regard to a well which has sections of good and poor ties. Starting with the zone of poor data, Figure 9 shows two possible ties in a time section from 2.5-2.9 s. Visually there is nothing to choose between them nor can any measurement of goodness-of-fit help. For example, the correlation coefficients between the nearest data trace and the synthetic seismograms are 0.796 and 0.790. Yet seismic wavelet 1 has a 36 degree phase advance and wavelet 2 a 48 degree phase

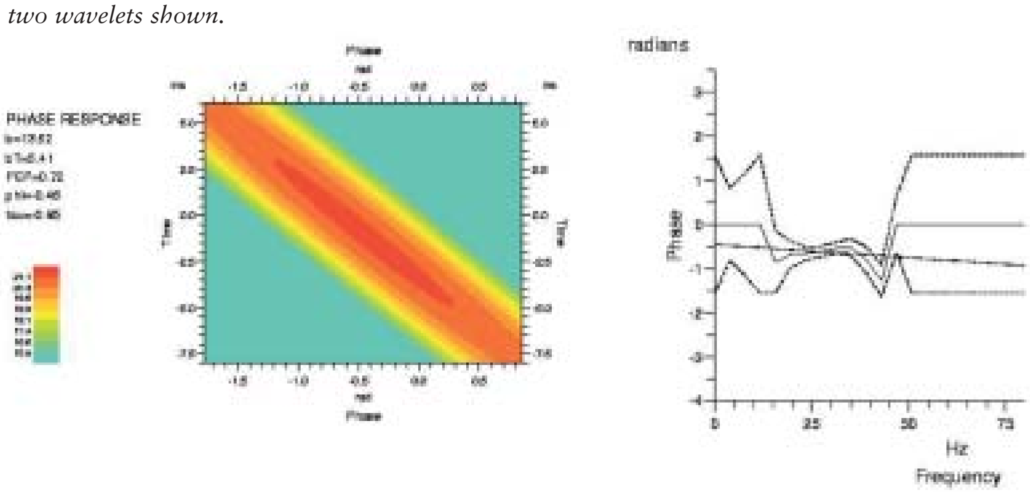

Figure 10 *Left: Log likelihood function, measuring goodness-of-fit, from estimating timing and phase for the well tie of Figure 9. The axes are time and phase advance in ms and radians respectively; i.e. negative values are lags. The maximum likelihood estimates are a time lag of 1.0+/-3.5 ms and a phase lag of -0.46+/-0.66 radians. The third contour from the centre marks the 90% confidence zone. The correlation coefficient between timing and phase is –0.98.*
Right: phase spectrum (solid curve) with standard errors (dashed curves) of the seismic wavelet estimated from the well tie for the 2.5-2.9 s time window shown in Figure 9. The intercept and slope of the sloping straight line give the maximum likelihood best fit of time and phase.

lag, a difference of almost 90 degrees. Their timing differs by 8 ms. In seismic inversion an error of 90 degrees in phase will reveal a step in impedance as a thin bed and vice versa.

The problem stems from the 400 ms segment of data. Its spectrum and that of the well-log reflectivity segment both show a spectral peak at 28 Hz. The data bandwidth is just 19 Hz. This bandwidth is inadequate for determining timing and phase accurately. Figure 10 (left) illustrates the ambiguity. It shows the log likelihood function, a weighted measure of goodness of fit, from estimating the phase and timing of this data segment. Timing and phase estimates are negatively correlated, and strongly so: the correlation coefficient between them is always in the range –0.866 to –1.0. Timing and phase combinations represented by the central red zone will give a reasonable tie. The lower the data bandwidth the more elongated is this zone.

Figure 10 (right) illustrates another way of looking at the problem. Estimated phase spectra and their error bounds from which the phase and timing are computed are shown. In the frequency domain the slope of the phase spectrum determines the timing and the intercept at zero frequency the phase. There is clearly a number of different fits (slope and intercept) to the data that can be made. This ambiguity of timing and phase is similar to the anti-correlation of intercept and slope also seen in AVO intercept-gradient crossplots (Simm, White & Uden 2001) and velocity-depth analysis (Al-Chalabi 1997).

The data with good bandwidth (2.9-3.4 s) are shown in Figure 11. The tie is clearly a good one and the effect of the

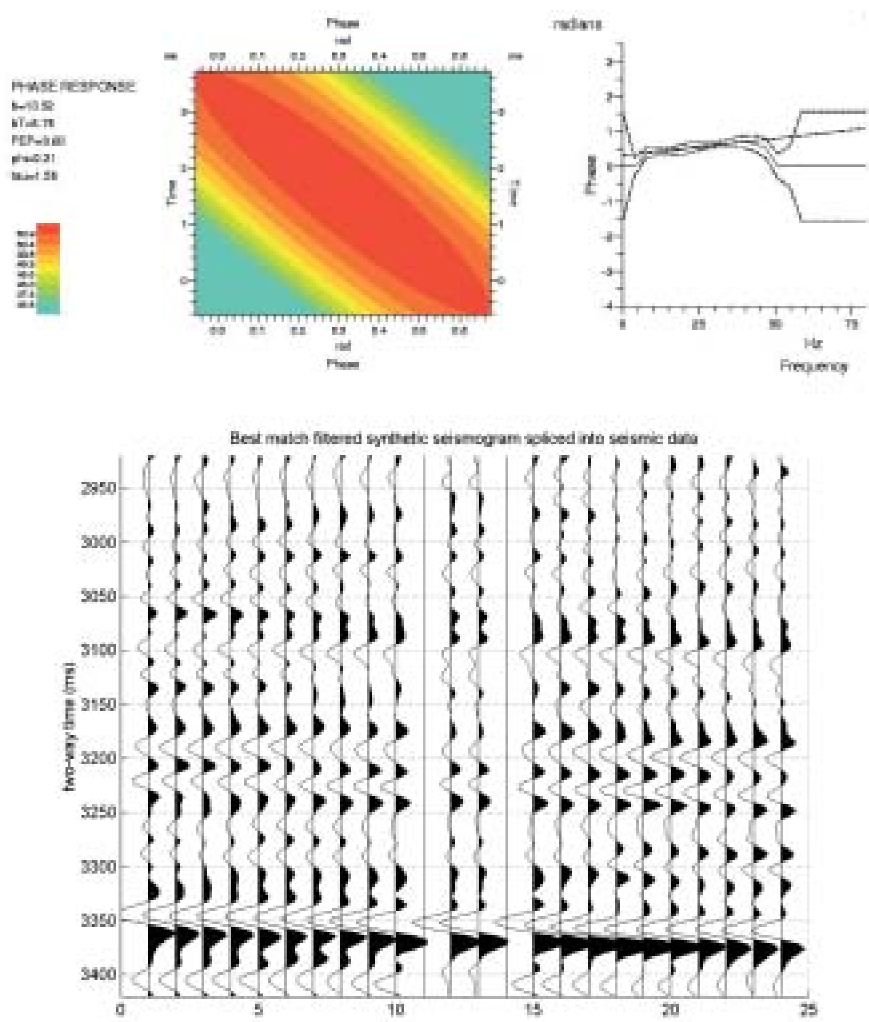

Figure 11 *Left: log likelihood function from estimating timing and phase for the well tie over the 2.9-3.4 s time window shown below. The axes are time and phase advance in ms and radians respectively; i.e. negative values are lags. The maximum likelihood estimates are a time lag of 1.6+/-1.0 ms and a phase lag of 0.31+/-0.18 radians. The third contour from the centre marks the 90% confidence zone. The correlation coefficient between timing and phase is -0.92.*
Right: phase spectrum (solid curve) with standard errors (dashed curves) of the seismic wavelet estimated from the well tie in the 2.9-3.4 s time window. The intercept and slope of the sloping straight line give the maximum likelihood best fit of time and phase.

broader bandwidth in resolving the intercept and slope is evident.

This example illustrates that the quality of a well tie depends primarily on the quality of the data – seismic and log. One cannot expect to make good well ties from poor data. The key seismic indicators are first of all bandwidth, and then the S/N ratio and duration of the data available for making a tie. The example also illustrates a distinction between goodness-of-fit and accuracy: the ties of Figure 9 look reasonable but their accuracy is poor. The eye can judge goodness-of-fit fairly well but accuracy needs to be measured.

Comparing wavelets

Is the seismic wavelet from well A in Figure 12 different from that from well B? At a superficial level they may appear to be but the question cannot be answered reliably without knowing how accurate the wavelets are. One simple test would be to compare the normalised mean square difference between the two wavelets with their estimated *NMSE*. Comparing their amplitude and phase spectra is more informative. The ratio of their amplitude spectra should hover around 1.0 within expected error limits (and it is in this case). Similarly the difference of their phase spectra should be zero, again within error bounds. There are two qualifications to this: first, the phase difference must be unwrapped to remove spurious 2π jumps and second, it may show a linear trend due a time difference. The error bounds for a ratio and a difference can be readily calculated from the individual standard errors.

The lower panel in Figure 12 shows the phase difference between wavelets A and B. The dashed lines show the standard errors. Standard errors correspond to 68% confidence bounds; i.e. roughly two thirds of the estimates are expected to fall within these bounds. The light line in this panel is the best fit straight line through the origin of the phase difference within the 8-50 Hz seismic bandwidth. It is within the error bounds for 65% of that bandwidth. A regression of the phase difference against frequency within the 8-50 Hz seismic bandwidth has a phase intercept is 22° ± 22°, making it doubtful whether the wavelets really differ in phase.

Overfitting always makes wavelets appear different since noise has been propagated into them. Measuring their accuracy puts the differences into perspective.

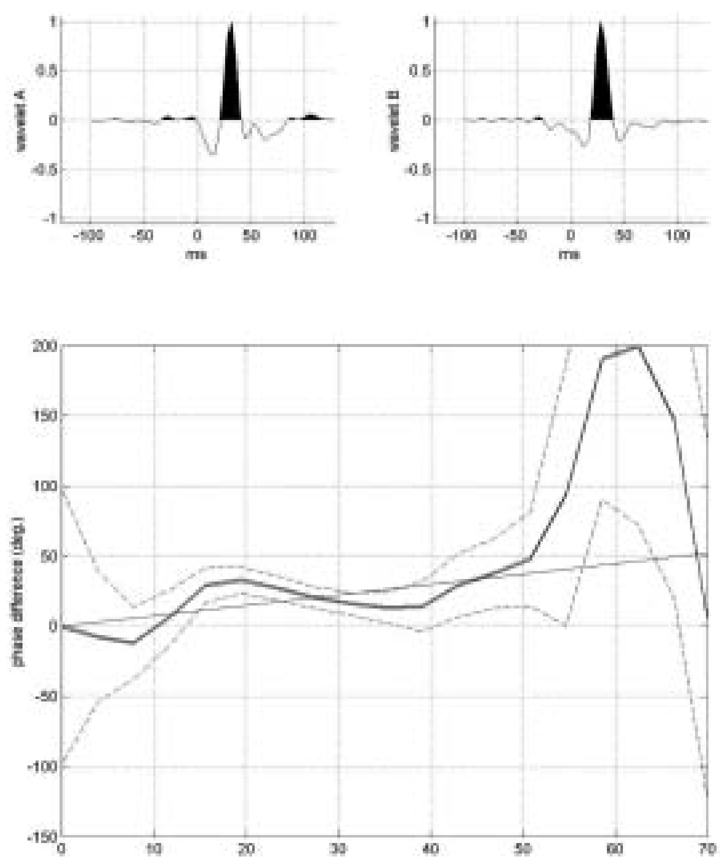

Figure 12 *Comparing the phase response of two seismic wavelets. The top panel shows (left) wavelet A extracted from the data in Figure 6 and (right) wavelet B extracted from the data shown in Figure 3 from the same 3D survey. The phase spectrum of wavelet B is shown in Figure 11. The bottom panel plots the unwrapped phase difference between the wavelet phase spectra: the dashed lines indicate the standard error bounds; the light solid line is a best fit through the origin of the phase difference within the 8-50 Hz seismic bandwidth.*

When comparison shows wavelet differences are within error bounds, the estimates can be combined, possibly through a weighted average, into a survey-wide wavelet. This is especially beneficial where no one wavelet stands out as the most accurate. It is very important that the combination does not blur the wavelet through small misalignments during averaging and this is most conveniently achieved in the frequency domain.

Comments on stretch and squeeze

One common bad practice is the stretching and squeezing of the synthetic seismogram to beautify the tie without any confirmatory evidence for the time shifts applied. The time shifts involved are often implausible on grounds of both geophysical reasoning and timing accuracy (White 1998). For example, any geophysical explanation of a time shift of just 2 ms would exhibit other measurable effects. The structural dip, lateral velocity variation or attenuating layer that produced this time shift would be unmissable in the seismic data. Stretch and squeeze may indicate a problem but it is not a valid remedy.

What do you do when all the ties are bad?

When a well tie is poor, it is difficult to offer any prescription for improving it beyond the general one of going back to q.c. the seismic data and calibrated logs. In general effort spent on practical details of the logs and seismic data are more likely to improve a well tie than niceties such as the time variance of the seismic wavelet. Time variance can be built into the synthetic model or, more dangerously for overfitting, into the matching procedure. Any improvement in the fit can be measured and tested. Our experience is that many perceived conceptual limitations of well ties, such as time variance of the seismic wavelet, are generally overridden by practical issues. A more productive line of action is to look for improvements in the log conditioning and calibration and, if feasible, some reprocessing of the seismic data. White, Simm and Xu (1998) documented an example where log conditioning and reprocessing combined to improve a well tie significantly. The outstanding well tie reported by White and Hu (1998) came about after a 2D line was reprocessed.

Conclusions

Well ties are a key part of the interpreter's art. Owing to the potential pitfalls in making assumptions about phase and timing, good practice in well ties demands a quantitative approach to the determination of wavelet shape and timing. We have shown how an appreciation of wavelet accuracy is required as well as the goodness-of-fit of the tie (as described by parameters such as the cross-correlation coefficient). Approximate measures of accuracy can be computed using the equations in this paper. A quantitative approach helps build experience on what matters in making a good well tie in a particular area. It relies on measurements from the data, not preconceptions of what seismic wavelets should look like.

There are good reasons for adopting a rigorous approach to well ties. Making assumptions about phase and timing can be misleading. A quantitative approach to estimating wavelet shape, timing and likely accuracy is recommended. Approximate measures of accuracy can be computed from the equations in this paper, starting from the cross-correlation coefficient between the filtered synthetic seismogram and the matched data trace. Measuring accuracy as well as goodness-of-fit is important in the quality control of a well tie.

A quantitative approach helps build experience on what matters in making a good well tie. It relies on measurements from the data, not preconceptions of what seismic wavelets should look like.

Acknowledgements

We thank Mike Bacon for practical suggestions that improved the paper.

References

Al-Chalabi, M. [1997] Parameter nonuniqueness in velocity versus depth functions. *Geophysics* **62**, 970-979.

Connolly, P. [1999] Elastic impedance. *The Leading Edge*, **18**(**4**), 438-452.

Hosken, J.W.J. [1988] Ricker wavelets in their various guises. *First Break* **6**(**1**), 24-33.

Marion, D., Mukerji, T. and Mavko, G. [1994] Scale effects on velocity dispersion: from ray to effective medium theories in stratified media. *Geophysics* **59**, 1613-1619.

Rio, R., Mukerji, T., Mavko, G. and Marion, D. [1996] Velocity dispersion and upscaling in a laboratory-simulated VSP. *Geophysics* **61**, 584-593.

Simm, R., White, R. and Uden, R. [2000] The anatomy of AVO crossplots. *The Leading Edge* **19**(**2**), 150-155.

Simm, R. and White, R.E. [2002] Phase, polarity and the interpreter's wavelet. *First Break* **20**(**5**), 277-281.

Stewart, R.R., Huddleston, P.D. and Kan, T.K. [1984] Seismic versus sonic velocities: A vertical seismic profiling study, *Geophysics* **49**, 1153-1168.

Walden, A.T. and White, R.E. [1998] Seismic wavelet estimation: a frequency domain solution to a geophysical noisy input-output problem, *IEEE Transactions on Geoscience and Remote Sensing* **36**, 287-297.

White R.E. [1980] Partial coherence matching of synthetic seismograms with seismic traces. *Geophysical Prospecting* **28**, 333-358

White, R.E. [1997] The accuracy of well ties: practical procedures and examples, *Expanded Abstract RC1.5*, 67th SEG Meeting, Dallas.

White, R.E. [1998] Stretch and squeeze – just keeping up appearances? EAGE 60th Conference and Technical Exhibition, Leipzig, *Extended Abstract P138*.

White, R.E. and Hu, T. [1998] How accurate can a well tie be? *The Leading Edge* **18**(**8**),1065-1071.

White, R.E., Simm, R. and Xu, S. [1998] Well tie, fluid substitution and AVO modelling - a North Sea example. *Geophysical Prospecting* **46**, 323-346.