Electronic Appendices: to appear as supplementary material on journal archive

Appendix 1: Typical computation of effective rock properties

Here we illustrate how the effective rock properties for a single layer might be computed from a full suite of fundamental model parameters (equation (1)) for the layer.

- Compute the brine+matrix density

$$\rho_{\rm sat,b} = \phi_s \rho_{\rm b} + (1-\phi_s)\rho_{\rm g}. \tag{43}$$

- Compute the associated moduli

$$\mu_{\rm sat,b} = \rho_{\rm sat,b} v_{s,s}^2 \tag{44}$$
$$M_{\rm sat,b} = \rho_{\rm sat,b} v_{p,s}^2 \tag{45}$$
$$K_{\rm sat,b} = M_{\rm sat,b} - \frac{4}{3}\mu_{\rm sat,b} \tag{46}$$
$$\mu_{\rm sat,fl} = \mu_{\rm sat,b} \tag{47}$$

- Use the fluid properties to compute $K_{\rm b} = \rho_b v_{p,b}^2$ and $K_h = \rho_h v_{p,h}^2$, and then

$$K_{\rm fl} = (\frac{S_h}{K_h} + \frac{1-S_h}{K_{\rm b}})^{-1}$$

- Using literature values of $K_{\rm g}$, compute $K_{\rm sat,fl}$ using equation (13). This comes out to be

$$K_{\rm sat,fl} = K_{\rm g} / \left\{ 1 + \left[\frac{1}{\phi_s} \left(\frac{1}{K_{\rm g}/K_{\rm fl} - 1} - \frac{1}{K_{\rm g}/K_{\rm b} - 1} \right) + \frac{1}{K_{\rm g}/K_{\rm sat,b} - 1} \right]^{-1} \right\} \tag{48}$$

  This equation has regions of parameter space with no positive solution (often associated with, say, small ϕ). The forward model must flag any set of parameters that yield a negative $K_{\rm sat,fl}$ and furnish a means of coping sanely with such exceptions. The same remark applies to equations (46) and (51) below.
- Compute the new effective density

$$\rho_{\rm sat,fl} = (1-\phi_s)\rho_{\rm g} + \phi_s(S_h\rho_h + (1-S_h)\rho_{\rm b}),$$

- Compute also the new effective p–wave modulus from

$$M_{\rm sat,fl} = K_{\rm sat,fl} + \frac{4}{3}\mu_{\rm sat,fl}$$

- Compute the impermeable–rock moduli

$$\mu_{\rm m} = \rho_{\rm m} v_{s,m}^2 \tag{49}$$
$$M_{\rm m} = \rho_{\rm m} v_{p,m}^2 \tag{50}$$
$$K_{\rm m} = M_{\rm m} - \frac{4}{3}\mu_{\rm m} \tag{51}$$

- Mix the *fl*–substituted permeable rock with the shale properties using the rock–mixing Backus averaging formula (9). Compute also the mixed density from (10).
- From the mixed rock, back out the velocities $v_{\rm p,eff} = (M_{\rm eff}/\rho_{\rm eff})^{1/2}$, and $v_{\rm s,eff} = (\mu_{\rm eff}/\rho_{\rm eff})^{1/2}$.

Appendix 2: Error sampling rates

When comparing synthetic seismic traces to co–located actual seismic, it is a question of some nicety as to what rate of sampling of the difference trace must be applied to quantify the error between the two. It is trivial to show that if a random process has, eg. a Ricker-2 power spectrum (i.e. $\sim f^2 \exp(-(f/f_{\rm peak})^2)$), then 95% of the spectral energy in the process can be captured by sampling at times

$$\Delta T_s = 0.253/f_{\rm peak}, \tag{52}$$

where $f_{\rm peak}$ is the peak energy in the spectrum (and often about half of the bandwidth). Most practical seismic spectra will yield similar results.

The sampling rate in equation (52) is about 2 times the Nyquist rate associated with the peak frequency. If the error is acquired with a faster sampling rate than this, adjacent error samples will be strongly correlated and the error term will become large. Since the error covariance is unknown, and bearing in mind that we seek an error measure corresponding to a maximum number of "independent" data points, we use the rate given by equation (52) and model the error points as i.i.d. values. The peak frequency $f_{\rm peak}$ is estimated at initialisation from the FFT spectrum of the wavelet.

Appendix 3: Mode–location starting points

The location, characterisation, and enumeration of local modes in the posterior distribution is performed by a loop over a set of strategically chosen starting

points in the model parameter space. Since one of the major functions of the prior distribution is to exclude improbable posterior modes, it makes sense to base the choice of starting points on the prior distribution.

The fundamental variables associated with multi–modality are the layer times, in that the posterior likelihood surface is frequently oscillatory in the layer times. Convergence to a secondary or minor solution is usually called loop–skipping, and may or may not be desirable. In situations where loop skipping may be permissible or even desirable, the prior uncertainty on the layer times is usually set to be quite broad, as these secondary solutions may be feasible. Convergence to these solutions can then be prompted by setting the initial layer times fed to the minimisation routine at values likely to be within the basin of attraction of the desired solution. The hueristic method used for defining these starting points is as follows:

Form a set of starting configurations, for each fluid state, by:

(1) Forming an approximate multi–Gaussian posterior distribution for the properties by updating the prior to account for any isopach constraints and the depth consistency requirement. This posterior is characterised by the mean $\tilde{\mathbf{m}}$ and covariance $\tilde{C}$. Recall that the layer times are the first $i = 1 \ldots (n+1)$ elements in the vector $\tilde{\mathbf{m}}$. Fix all non-time parameters $i = n+2 \ldots$ at the values in $\tilde{\mathbf{m}}$.

(2) From this posterior, compute the expected layer time thickness $\overline{\Delta t_i}$, layer time uncertainties σ_{t_i} and layer time thickness uncertainty $\sigma_{\Delta t_i}$, from

$$\overline{\Delta t_i} = \tilde{\mathbf{m}}_{i+1} - \tilde{\mathbf{m}}_i \tag{53}$$

$$\sigma^2_{t_i} = \tilde{C}_{i,i} \tag{54}$$

$$\sigma^2_{t_{i+1}} = \tilde{C}_{i+1,i+1} \tag{55}$$

$$\sigma^2_{\Delta t_i} = \tilde{C}_{i,i} + \tilde{C}_{i+1,i+1} - 2\tilde{C}_{i,i+1}. \tag{56}$$

(3) Now, if $\sigma_{\Delta t_i}/\overline{\Delta t_i} > 0.5$, this amounts to a $\approx 5\%$ chance of the layer pinching out ($\Delta t_i < 0$) before taking into account the seismic, so this is likely enough to warrant investigation. The starting points

$$t'_i = \frac{\sigma_{t_{i+1}}\tilde{\mathbf{m}}_i + \sigma_{t_i}\tilde{\mathbf{m}}_{i+1}}{\sigma_{t_{i+1}} + \sigma_{t_i}} \tag{57}$$

$$t'_{i+1} = t'_i - 2\epsilon_{FD} \tag{58}$$

are then used as suitable starting times for the times t_i, t_{i+1}. The remaining times t_j, $j \neq i, i+1$ are set at the posterior means $t_j = \tilde{\mathbf{m}}_j$. If more than one layer is potentially pinched out, we form a set of starting points by looping over the set of potentially–pinched–out layers, and setting only one of these layers at a time to the pinched–out starting configuration just described, with remaining times at the values in $\tilde{\mathbf{m}}$.

Appendix 4: An Independence Sampler

For each fixed fluid combination k, a successful location of a mode in the full parameter space will produce a "maximum likelihood" vector $\tilde{\mathbf{m}}_{k_j}$ and local covariance $\tilde{C}_{k_j}$, by successive Newton iterates (equations (33)–(37)). Here the index j runs over the set of distinct modes we have found for each fluid combination by starting from separated starting points, as per Appendix 3 (in many cases there will be only one). To save notational baggage, we will use k henceforth to mean k_j - the context makes the meaning clear.

The Laplace approximation to the mode posterior marginal likelihood is formed by assuming the posterior surface is fully Gaussian, centred at the maximum–likelihood point, and characterised by the covariance $\tilde{C}_k$. Integrating over all the model parameters then gives the marginal probability to be proportional to

$$q_k \sim (2\pi)^{d_k} |\tilde{C}_k|^{1/2} \Pi(\tilde{\mathbf{m}}_k | \mathbf{S}), \tag{59}$$

where the posterior Π is evaluated as per (27,28), and the determinant and dimension–varying terms in (28) are carefully retained.

A full independence sampler can then be implemented using a Metropolis Hastings method. With the chain in mode k, with model vector $\mathbf{m}_k$, we propose a new mode k' with probability $\sim p'_k$, whose new model vector $\mathbf{m}'_k$ is drawn from $q(\mathbf{m}'_{k'}) = N(\mathbf{m}'_{k'} | \tilde{\mathbf{m}}'_{k'}, \tilde{C}_{k'})$. We accept this candidate with the usual probability

$$\alpha = \min(1, \frac{\Pi(\mathbf{m}'_{k'} | \mathbf{S}) p_k q(\mathbf{m}_k)}{\Pi(\mathbf{m}_k | \mathbf{S}) p_{k'} q(\mathbf{m}'_{k'})}, \tag{60}$$

or leave the state of the chain unchanged. This method has have somewhat mixed success - it is perfect when drawing from the prior (since the prior is multi–Gaussian), but is susceptible of being trapped in particularly favourable states for long periods when the posterior departs sharply from Gaussianity. Further, any thick tails in the posterior density are not likely to be explored by the proposal density, since it has compact Gaussian tails. The proposal overdispersion recommended by Gelman (Gelman et al., 1995, Ch. 11) does not work well, since the dimensionality is so high that even mildly overdispered t–distributions dramatically reduce the acceptance rates, even when sampling from the prior.

It has proven useful to implement an approximate sampler where the modes are proposed from the posterior marginal ($\sim q_{k'}$) with normal approximations at each mode ($q(\mathbf{m}'_{k'}) = N(\mathbf{m}'_{k'} | \tilde{\mathbf{m}}'_{k'}, \tilde{C}_{k'})$, and mandatory acceptance ($\alpha = 1$).

Such candidates appear to very rarely produce unlikely χ^2 values for the seismic and isopach likelihoods.

Appendix 5: Modified SU trace formats for properties

5.1 Model–prior trace formats

The local prior information for the inversion at each trace is communicated to the inversion code by a SU data set with a set of traces on a pattern and ordering that matches the SU seismic data set exactly. The xml *Modelparameters.xml* file which sets up the inversion will specify which prior quantities vary on a trace-by–trace basis, and the block indices k at which they are located. Suppose there are N_p properties which vary, and there are N_l layers. The value for property k in layer l is then the $(l + (k-1)N_l)$th float value in the current trace of prior data. There are some special rules as well

- The layer times must always be present in the prior trace, and they must always be the last block ($k = N_k$). Further, one additional float value is appended for the base time of the last layer t_{base}).
- All values are expected to be positive, and are converted to positive numbers if not. The value -999.0 signifies a "bad value", and will be ignored (the default prior will then apply)
- All times are in milliseconds.

5.2 Stochastic trace outputs

Delivery will write stochastic samples from the posterior to the requested output file in a very similar format to the model–prior traces. A sequence of blocks is written, each of size N_l, and the last block of $N_l + 1$ floats is either time or depth, as specified by the obvious entry in the *ModelDescription.xml* file. The names of the properties can be written to a single–line ascii file by supplying **-BHPcommand filename** to the *delivery* command line.

Similarly, the *deliveryAnalyser program* can collate the large statistical ensemble produced by the *delivery* output into a compact set of traces representing either means, standard deviations or quantiles of all the salient quantities of interest. The set of property names of the blocks in these summary traces is obtained by adding **-BHPcommand filename** to the *deliveryAnalyser* command line.

Appendix 6: Wavelet format

Wavelets are expected to be single trace SU files, with the fields ns, dt and f1 set to reflect the number of samples, sampling interval, and starting time of the first sample (beginning of precursor). The ns and f1 fields are set so as to make the wavelet as compact as possible with respect to the tapering, so excessive samples are discouraged, but the tapering must also be smooth to the first and last sample.

Appendix 7: Usage of the code

7.1 Inputs

Typical usage of the **delivery** inversion code demands 3 input files:

- An XML file **ModelDescription.xml** specifying the disposition of the layers, the rock and fluid physics, and so meta–information about the su traces containing the variable prior information.
- One or two SU files for the seismic data. These will represent near (and possibly far) offset data.
- An SU file containing the variable prior information.

In distributed systems, the latter two items may refer to named pipes on a unix/Linux system.

A typical inversion would be run with the command
% delivery ModelDescription.xml -m prior_traces.su -v 3 -PD -o realisations.su -N 1000
which will generate (**-N**) 1000 realisations per trace and write them (**-o**) to the file **realisations.su**, using the pinchout detection (**-PD**) BFGS methods described in section 4.1. The prior model (**-m**) is read from the SU file **prior_traces.su**. Many other commandline options are available, and can be perused in the code's SU–style self documentaion.

7.1.1 The **ModelDescription.xml** *file*

The inversion is chiefly driven by an XML file that specifies the necessary rock–physics, layer descriptions, and information about the seismic. A well formed XML file is crucial to this end, so the file can be created using a schema–driven GUI editor expressly built for the purpose, which is also available at the website. "Help" documentation for the inverter is available through standard

menus. The editor is also capable of creating, running and controlling complex scripts on distributed systems, using a schema for common unix, SU, BHP_SU and **delivery** commands.

The structure of the XML document is controlled by the schema, and the meaning of the various entries is readily inferred by reference to the examples in the distribution.

Some particular entries in the XML file are worth further explanation

- The attribute "varies_areally" may be set to true or false for a certain *property*. If true, the quantity in question will vary on a trace–by–trace basis, and is expected to be read in from the SU model file. The "<property_indices>" block then will contain an entry "*property*_link" which defines the block of the SU model file containing the property (see Appendix 5).
- Various properties may contain the attribute "same_as_layer_above". This is set to true when the property in question should be a replica of the same property in the layer above. For example, the low–frequency–interval–velocity (LFIV) is usually a coarse scale measure representing an average across many layers, so it is best fixed at the same value (or mapped to the same stochastic variable) for all of the relevant layers. Likewise, the "depth" input is an approximation used to sample from the trend curves, and may safely be set to the same value for many adjacent thin layers. Prudent use of this attribute helps to reduce the dimensionality of the final model.
- Each "<layer>" entry may have the attribute "synchronise_rock_matrix" set to true. This is a flag to indicate that any stochastic rock properties of a given rock type in this layer must be matched to those of the matching rock type in the layer above (for example, the v_p, v_s and ϕ_s of the permeable member in the two layers is mapped to the same underlying stochastic variable By this means, two adjacent layers with matching end–member rocks are made acoustically indistinguishable but for fluid property differences. Any internal reflections are then due to fluid contrasts, and this feature can then be used to make the layer boundaries track a fluid interface.
- Linear trend curves are of form $v_s = \text{slope} \times v_p + \text{intercept} \pm \text{sigma}$. The Gardner-Gardner-Gregory regression curves for impermeable rocks are of form $\rho = \text{factor} \times v_P^{\text{exponent}} \pm \text{sigma}$.
- The "<output>" block has the following useful flags.
  - "<density_ordering>" is set to one of full/partial/none, with the meanings as described in section 2.2.2.
  - "<realisation_number_header_field>" denoted the SU header word used to store the realisation number.
  - "<master_depth_layer_number>" denotes the layer number from which depths of the other layers will be computed using summations of layer thicknesses ($v_{\text{p,eff}}\times$ one–way time thickness). The actual depth of the master layer is the depth variable d in the model vector $\mathbf{m}$.

The inversion code dumps stochastic realisation of the posterior model to a series of SU traces, so the fastest scan direction in the resulting (large) SU file is that of the realisation number, which defaults to the **mark** header word. Posterior analysis of the simulations is performed with a second program **deliveryAnalyser**, which can form various statistical quantities of interest, synthetic seismic traces, and blocked layer maps of the simulations. Some of the salient quantities of interest include

- All the basic model variables in **m** of equation (3).
- Fluid type.
- Effective layer properties such as $v_{\text{p,eff}}$, ρ_{eff}, Z_{eff} (impedance).
- Effective reflection coefficients for each stack at the layer top.
- A normalised χ^2 variable indicating the degree of match with the seismic ($-\log(L_{\text{seis}})/N_{\text{error-sampling-points}}$ from equation (25)).
- Thickness (one–way time–thickness $\times$ $v_{\text{p,eff}}$.
- Net sand (net sand = thickness $\times N_G$).
- Net hydrocarbon (Net sand $\times$ porosity $\times$ saturation).

Like **delivery** and SU programs, the program has self documentation, but some typical analysis of a suite of realisations **realisation.su** might be

(1) **deliveryAnalyser -i realisations.su -p-ascii t 6 | ascii–graphing–program**
Print the time of layer 6 from succesive realisations out as an ascii stream and graph it.

(2) **cat realisations.su |deliveryAnalyser --mean NG 6**
Print a series of realisation averages of N_G in layer 6, one for each trace location.

(3) **cat realisations.su |deliveryAnalyser --filter fluid_type = 1 4 --trace--filter 'fldr>1234' --mean NG 6**
Print a series of realisation averages of N_G in layer 6, one for each trace location, but only for realisations with oil (fluid_type=1) in layer 4, and for the traces with fldr>1234 (note the shell protection).

(4) **deliveryAnalyser -i realisations.su --histogram net-sand 6 --trace--filter 'fldr=34,tracf=12' | ascii–graphing–program**
Print a histogram of net-sand ($N_G\times$thickness) in layer 6 at the specified location.

(5) **deliveryAnalyser -i realisations.su -s 4 5 R_near wavelet.su | suswapbytes | suximage**
Produce a synthetic near–stack seismic from the model realisations over the time interval 4-5 seconds from the wavelet **wavelet.su** and display this on a little endian machine (all the java codes wrote big–endian SU

files).

End of electronic appendices