

Chapter 1

Processing seismic monitor data

1.1 Overview

In this chapter, I discuss a seismic data processing strategy to obtain high-resolution subsurface images which are consistent from time-lapse survey to survey. Processing issues will be illustrated with a field data example from the Duri Field, Indonesia, in which six 3-D surveys have been recorded at a shallow steam injection site over a total period of 20 months. I describe the data acquisition geometry and parameters, and the preprocessing steps I used to enhance weak reflection energy in the data. Finally, I discuss my approach to optimizing coherent signal in the NMO stack, migration, and difference sections, and suggest future work to further improve seismic images of the steamflood process. Chapter 5 gives an interpretation of the Duri seismic data integrated with fluid flow, rock physics, seismic modeling and velocity analysis.

1.2 Introduction

The most important feature of time-lapse seismic monitoring data is the opportunity to compare seismic images as a function of elapsed time. Careful attention to data processing issues is needed to ensure that images obtained at one time are validly comparable to subsequent images. This is especially true for seismic difference sections, in which one seismic image is subtracted from a second seismic image acquired at a different calendar time. Small artifacts in amplitude, phase, time and depth of imaged seismic events can lead to disastrously noisy difference sections which may completely obscure seismic fluid-flow anomalies. The following is a list of processing issues which I consider important for seismic monitoring data:

- experimental repeatability
- survey position accuracy
- wavelet shape

- spectral content
- amplitude preservation
- velocity accuracy
- depth accuracy

1.2.1 Repeatability

Experimental repeatability is the degree to which the same experiment can be repeated without incurring differences in raw seismic observations or processed images. At the Duri Field site, two baseline surveys were acquired before steam injection to test repeatability (?). Each survey was acquired in one full day, at three days separation. Analysis of both data sets revealed that the seismic monitoring experiment was repeatable within 0.6 ms in rms time error, and 6% in rms amplitude error. This ensures that seismic time-lapse anomalies greater than these values can be attributed to changes in reservoir fluid flow rather than ambient noise and experimental design. To some extent, seismic data processing can be used to enhance repeatability by wavelet deconvolution, spectral shaping and surface-consistent amplitude balancing.

1.2.2 Position accuracy

Survey position errors of source and receiver locations can obscure seismic monitor anomalies. Johnstad et al. (?) estimated serious errors in survey coordinates between successive time-lapse data sets acquired over the Oseberg Field in the North Sea, and showed how these errors almost completely masked a fluid-flow anomaly. Raw prestack traces or partially processed images could be correlated between two surveys to estimate any relative acquisition geometry coordinate errors. At the Duri Field, fixed shot and receiver holes were pre-drilled and reoccupied for each survey, in part to ensure that position errors would not contaminate seismic time-lapse results.

1.2.3 Wavelet consistency

It is desirable to maintain wavelet shape between surveys to enhance difference-section anomalies. A highly repeatable source waveform is advantageous. At Duri, dynamite sources placed in shot holes were of a small enough charge size so that the hole was not damaged between successive surveys. Hence, little change in source waveform is expected since the shot cavity is fairly stable with time. Wavelet deconvolution, spectral shaping, and match filtering can be applied to time-lapse seismic data to help regularize wavelet characteristics between successive surveys. It is possible that different wavelet processing parameters may be required for each survey to regularize wavelet shape. For example, if two surveys have different spectral content, it may be useful to reduce both data sets to only that portion of the spectrum which overlaps.

The original region of overlap may be extended with deconvolution if S/N conditions are favorable. I believe that model-based processing is preferable to statistical data-dependent processing, to ensure that wavelet shaping is based on a consistent model of the earth physics, as opposed to a potentially inconsistent signal character among multiple seismic data sets. This underscores the importance of surface-consistent deconvolution.

1.2.4 Amplitude preservation

Amplitude variations can make it difficult to interpret seismic difference sections, even if time difference errors are negligible. Many factors can cause amplitude variations between time-lapse data sets. A common cause of amplitude errors, especially for land data, is the variation in source and receiver coupling at the surface. In this case, surface-consistent amplitude balancing can be applied to the raw data, e.g., Brzostowski and McMechan (?); Lumley et al. (?). Other factors affecting geometric spreading loss, Q -attenuation, and offset-dependent AVO quality control (?) may need to be carefully addressed *simultaneously* for each survey. At the Duri site, hydrophones were used as receivers in water-filled receiver holes. This ensured that the receiver coupling between surveys was consistent, and the depth of burial reduced potential amplitude variations caused by changing near-surface conditions.

1.2.5 Velocity accuracy

Seismic stacking and migration velocities for each survey need to be carefully considered. Changes in velocity with time lead to changes in the arrival time (pseudodepth) of imaged seismic events, and changes in the hyperbolic moveout focusing of prestack data. In general, for time-domain processing such as NMO stack and time migration, separate rms velocity functions should be estimated for each survey. This ensures the best focusing of prestack reflection events. In some cases, seismic velocity changes are small enough that time-domain processing and interpretation is sufficient. This is the case for the Troll study in Chapter 3, since the seismic difference anomaly is mainly due to a reflection amplitude change, rather than a change in structural pseudodepth. However, for situations such as the Duri Field that exhibit strong lateral changes in velocity, depth-domain velocity analysis is preferable.

1.2.6 Depth accuracy

If velocity changes are strong enough to cause significant structural changes in imaged events between surveys, the best method for comparison is in the depth domain. Pseudodepth image times are affected by velocity, whereas in reality, the physical depth of the reservoir does not change between surveys. The advantage would be that one could compare seismic results in the physical depth domain where difference

sections are more meaningful. However, depth domain comparisons require time-to-depth conversion or depth migration. Images can be calibrated in depth using well log information and depth migration velocity analysis. These are extremely computational and time-intensive processes, especially for several 3-D data sets, and depend critically on the ability to accurately estimate depth interval velocities from seismic data. Prestack 3-D depth migration and velocity analysis in the presence of rapid lateral velocity variations is still a challenging problem in seismic data processing research.

In the next sections, I review the data acquisition and processing strategy I employed for the Duri 4-D seismic data set.

1.3 Seismic data acquisition

Six repeated 3-D surveys have been collected in time-lapse mode over a shallow heavy oil reservoir undergoing steamflooding at the Duri Field, Sumatra, Indonesia. The steam injection pattern, well locations and seismic data coverage area are shown in Figure 1.1.

1.3.1 Shot array

Each 3-D survey consists of 301 shots fired into a fixed array of 480 receivers. A single shotpoint consists of a small dynamite charge (50 g) fired in a shot hole at 15 m depth. The shots are laid out in a staggered grid of 22 shot lines aligned in the Y direction, as shown in Figure 1.2. Each shot line contains 14 shotpoints at an inline spacing of 30 m. The 22 shot lines are separated by 20 m in the crossline direction.

1.3.2 Receiver array

The receivers are laid out in a regular grid of 16 receiver lines along the X direction, as shown in Figure 1.3. Each receiver is a single hydrophone immersed in water at 7 m depth. Each receiver line contains 30 receivers spaced at a 10 m interval. Receiver lines are separated by a 20 m interval. The receiver array is stationary such that each shot is recorded in the same complete receiver array.

1.3.3 CMP coverage

Figure 1.4 shows an overlay of the shot and receiver positions. The total dimensions of the pattern are about 350 m by 350 m, and shot lines are perpendicular to receiver lines. Frequencies in the data exceed 200 Hz at a 1 ms sample interval. The maximum offset in the data is 480 m, with a good range of offset and azimuthal coverage. Figure 1.5 shows the locations of the CMP bin centers, and Figure 1.6 shows the

CMP fold of each 3-D survey, which varies spatially from a minimum of zero at the edges of the pattern, to 24 fold at the 7-spot hexagon, to a maximum fold of 80 at the center steam injector location.

Six repeated seismic monitor surveys have been obtained with identical acquisition parameters. Shot and receiver holes were predrilled and reoccupied for each monitor survey. Both sets of holes are cased down to just above the source or shot depth. Dynamite charges were lowered and tamped with sand for each survey, and the small charge size was chosen for good bandwidth and negligible hole damage. Hydrophones were lowered into each receiver hole and filled with water. Repeat surveys were acquired at intervals ranging from two months after steam injection, to 19 months after steam injection, with an average 4 month repeat interval.




Figure 1.1: Duri 4-D seismic coverage over a single 7-spot production pattern. Steam injection well is at the center, flanked by two temperature observation wells.
proc-7-spot [NR]

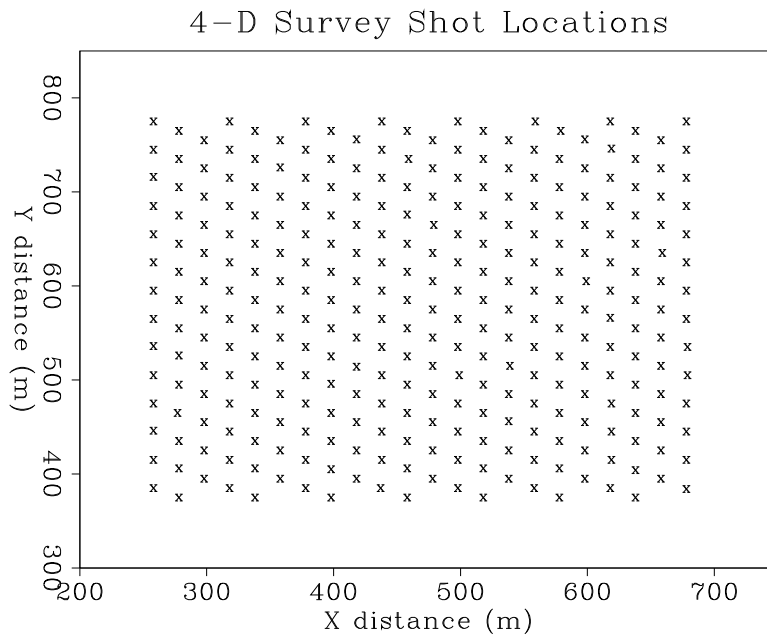


Figure 1.2: Shot locations of the 4-D monitor survey. `proc-shots` [CR]

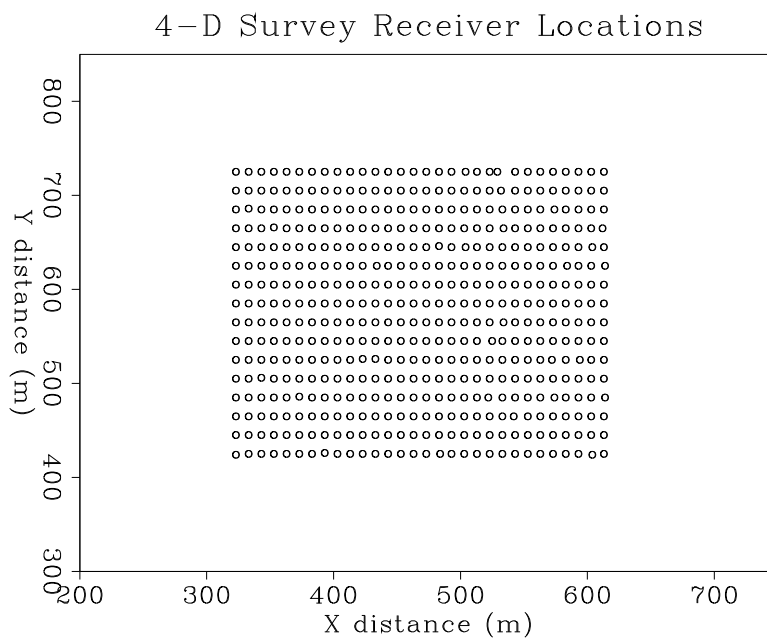


Figure 1.3: Receiver locations of the 4-D monitor survey. `proc-rcvrs` [CR]




Figure 1.4: Overlain shot and receiver locations of the 4-D monitor survey. proc-survey [CR]




Figure 1.5: CMP bin-center locations of the 4-D monitor survey. proc-bins [CR]

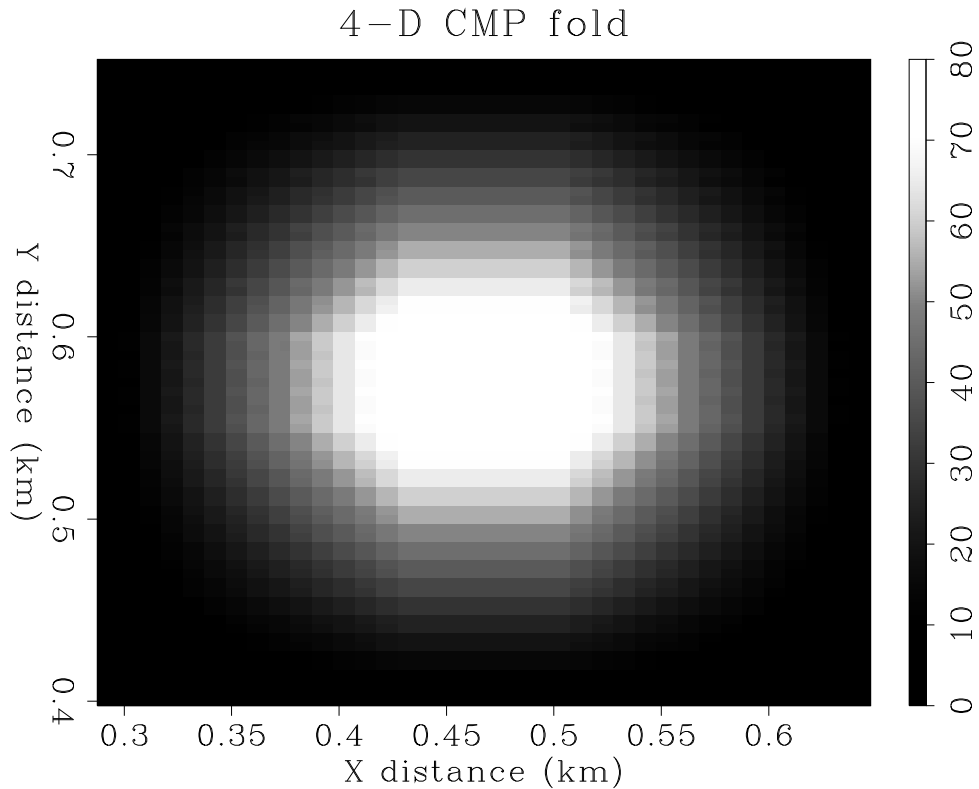


Figure 1.6: CMP spatially-variable fold of the 4-D monitor survey. `proc-fold` [CR]

1.4 Preprocessing

1.4.1 Raw data

All six 3-D seismic data sets were partially preprocessed by CalTex Petroleum Indonesia (CPI) before I received them. I will call this the “raw” data. These gathers have been partially preprocessed by CalTex as follows:

- Demultiplex, Geometry
- Datum corrections
- Spreading correction (4 dB/s)
- First break mute
- Spiking deconvolution
- Residual statics

Deconvolution was done in order to increase the frequency content of the data, and to help maintain a consistent wavelet shape and spectral content from survey to survey. Two loops of residual statics were applied independently for each survey to

help reduce static time shifts between surveys caused by variations in the near surface velocities (water table level, soil moisture, etc.). Residual static estimates were on the order of 1–2 ms, probably kept small by burying shot and receivers below the weathering layer and water table in the predrilled holes.

Figure 1.7 shows two raw CMP gathers near the center of the acquisition pattern where the fold attains a maximum of 80. Noise levels are high enough that only a few reflections are faintly visible. Several traces are badly gained, and strong groundroll at the near offsets, and refractions at the far offsets dominate the data. The left panel of Figure 1.11 shows a brute stack of ten raw CMP gathers near the center of the pattern. The result is not very good compared to the right panel, which is my best stack after optimizing the remaining preprocessing steps, as described next.

1.4.2 Geometric spreading

The first step I took was to remove the 4 dB/s spreading correction applied by CPI. This type of gain correction is unphysical and often overcompensates data at late arrival times. I then applied a physical gain correction based on travelttime t raised to the power p , such that

$$g(t) = t^p . \quad (1.1)$$

I found that a value of $p = 1.5$ gave a good geometric spreading correction to the trace amplitudes.

1.4.3 Trace balance

Several of the raw traces appeared over or under gained, as shown in Figure 1.7. Perhaps this is due to some variation in the amount of mud around hydrophones in the receiver holes, or shot energy variations. These amplitude imbalances can deteriorate reflection coherency within a single data set, and can lead to noisy difference sections when comparing a suite of images from several time-lapse surveys. Ideally, a surface-consistent amplitude balancing process should be applied in 3-D (e.g., Brzostowski and McMechan (?); Lumley et al. (?)). I used a quick and approximate alternative to surface-consistent amplitude balancing by using automatic gain control (AGC) on each trace with a window length equal to the total trace length. This ensures that each trace in every survey has the same energy, measured as the sum of squared trace sample values.

1.4.4 Bandpass and deconvolution

I found that the spiking deconvolution applied by CPI created a lot of high-frequency noise above 200 Hz in the original data, as shown in the left panel of Figure 1.9. Additionally, some low frequency noise below 20 Hz contaminated the reflection data, perhaps also related to the deconvolution operator applied, or ground roll energy. In

the preprocessing phase, I applied a 10–350 Hz bandpass filter to suppress the worst of the noise at the extreme ends of the spectrum. Then I made a second pass of wavelet deconvolution with soft parameters (short operator) that enhanced the waveform similarity across time-lapse data subsets, followed by a second 10–350 Hz bandpass filter. After stack and migration, I applied a final bandpass filter of 20–250 Hz to remove spectral noise not suppressed by imaging.

1.4.5 Mute

The original data had been processed with an outer mute by CPI. This left too much residual refraction energy and possibly postcritical reflection energy at the far offsets, in my opinion, as seen in Figure 1.7. I applied an extra outer linear mute to suppress this noise from the ensuing stacks and migrations, using a mute velocity of 900 m/s and a time intercept of -25 ms. Additionally, none of the ground roll had been removed by CPI from the near offsets at late times. Although the groundroll was not as strong as it often can be, I found that the stacks and migrations improved with the application of an inner mute, using a velocity of 250 m/s and time intercept at the origin.

1.4.6 Final preprocessing flow

I converged on an optimal preprocessing flow for these raw data as follows:

- Remove original dB/s spreading correction
- Trace balance (AGC, full window)
- Apply new spreading correction (tpow=1.5)
- Bandpass filter (10–350 Hz)
- Wavelet deconvolution
- Bandpass filter (10–350 Hz)
- Apply new mute (inner and outer)

Figure 1.8 shows the same CMP gathers as Figure 1.7 after geometric spreading corrections, trace amplitude balancing, bandpass filtering, wavelet deconvolution and refraction/groundroll mutes. Note that reflection energy is more coherent than in the raw CMP gathers, but is still contaminated by large noise levels. Figure 1.9 shows the amplitude spectrum of the raw data (left panel) and the final preprocessed data (right panel). The raw amplitude spectrum corresponds to the raw CMP gathers of Figure 1.7, and the preprocessed spectrum corresponds to the preprocessed CMP gathers of Figure 1.8. Coherent signal is now visible in the preprocessed data above 200 Hz.

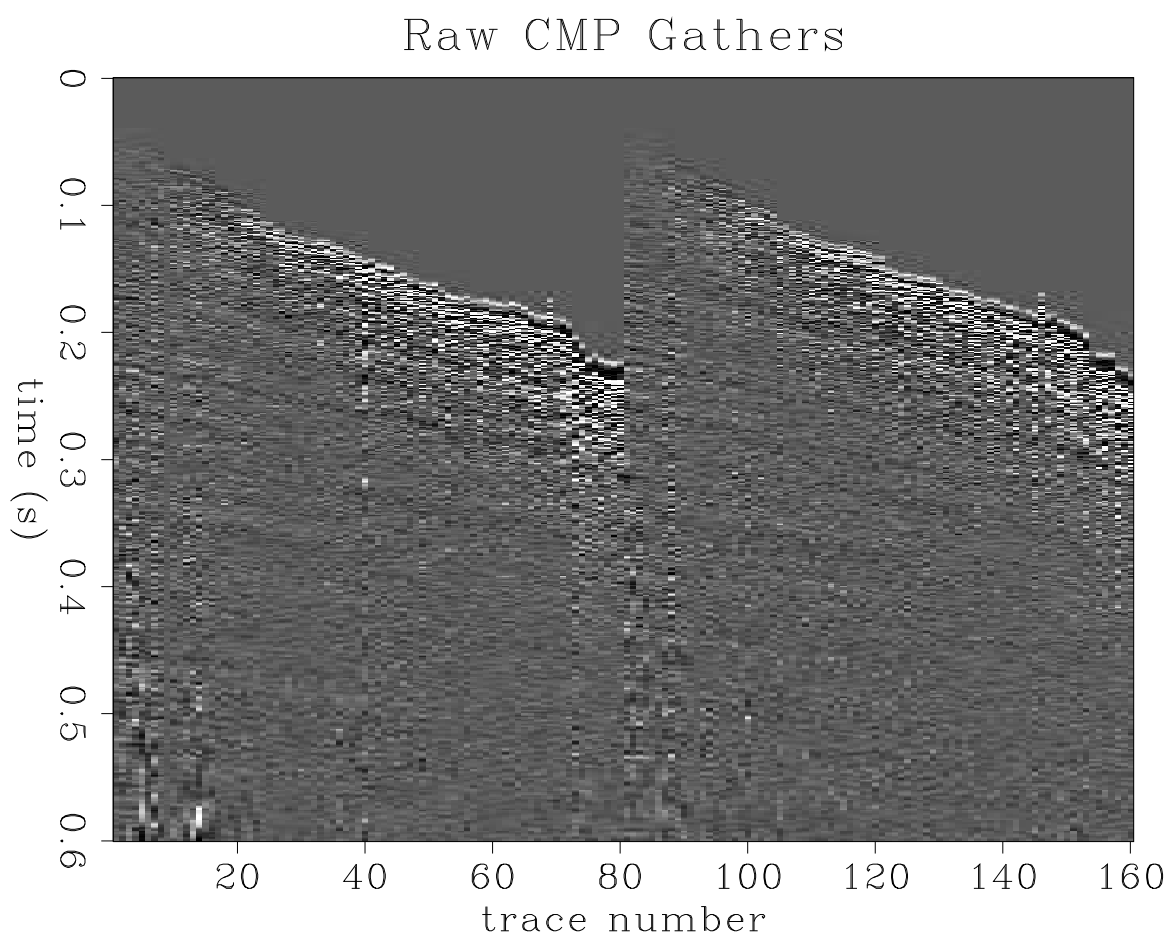


Figure 1.7: Raw CMP gathers near the center of the 3-D survey pattern.

proc-raw

[CR]

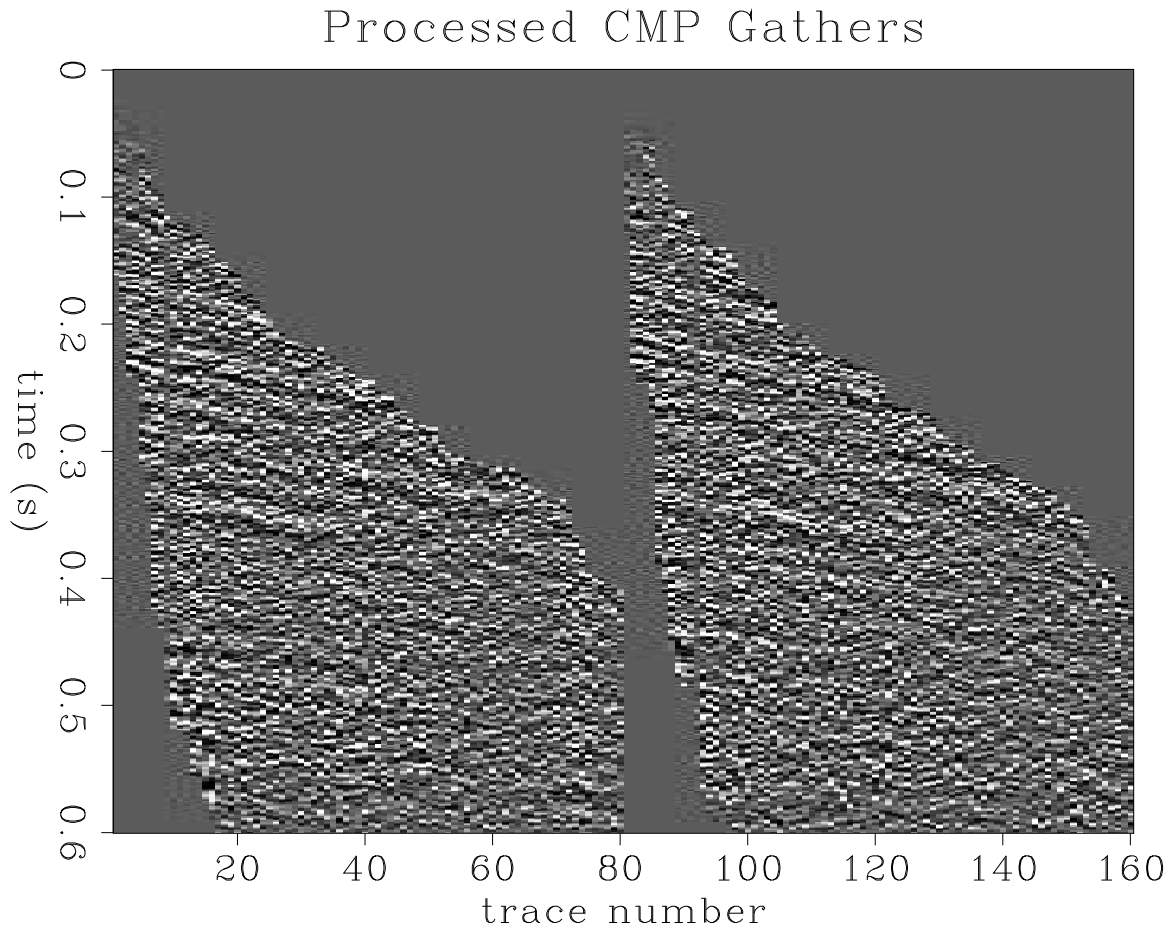


Figure 1.8: Processed CMP gathers near the center of the 3-D survey pattern. `proc-proc` [CR]

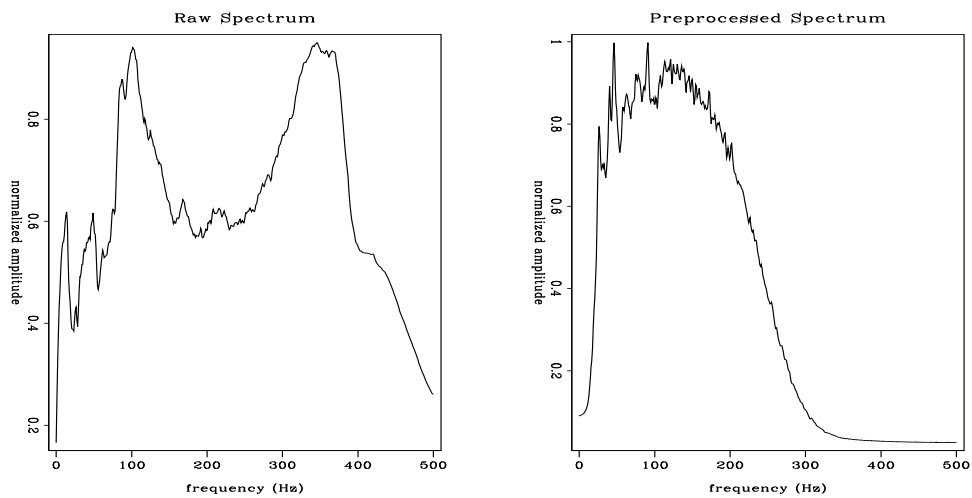


Figure 1.9: Amplitude spectra before and after preprocessing. `proc-spectral` [CR]

1.5 NMO stack

Normal moveout (NMO) stack is the process of summing seismic reflection data over hyperbolic trajectories within CMP gathers. Based on the assumption of horizontal reflectors in the earth, seismic data at a non-zero offset are time-shifted to earlier traveltimes along stacking velocity hyperbolas, and summed (stacked) into the zero-offset location. The net result is to emulate a zero-offset seismic experiment with coincident source and receivers, and to increase the S/N ratio through the summing process (?). If the subsurface geology is not too complex, the resulting stacked section is an approximate image of the geologic structure.

The S/N improvement in the data generated by the NMO stack process is approximately proportional to \sqrt{N} , where N is the fold of the data. Since the fold of the Duri data ranges from 24 at the edges of the 7-spot hexagon, to 80 at the center of the survey, the anticipated S/N improvement should be on the range of 5–9 in the stacked images. As the prestack data are very noisy, the NMO stack should greatly improve the quality of the data.

1.5.1 Optimal stack parameters

I estimated a stacking velocity function “interactively” until the events were flat (by eye) after NMO. This velocity is constant at 1.75 km/s until 200 ms, whereafter it increases gradually with time with a gradient of about 0.5 Hz. I used the same $v(t)$ stacking velocity function for each survey. In reality, there is some slight lateral velocity variation above the reservoir, and strong lateral variations within the reservoir, as a function of time. A better approach might be to find a separate 3-D stacking velocity function $v(x, y, t)$ for each survey. However, to demonstrate the first order effects, I chose a single velocity function that gave the best results for all six data sets.

I paid careful attention to make sure that the frequency content in the data was preserved as best as possible after stack for each survey. In particular, I tried to preserve the shallow high-frequency events with less than 100 ms arrival time in the stacks. Additionally, I sought to optimize the stack energy at the target zone around 150–200 ms travelttime. This was achieved by a combination of the preprocessing parameters discussed in the previous section, and choosing a simple stacking velocity function that gave the best result *simultaneously* for subset of all six data sets.

Figure 1.10 shows the CMP gathers after NMO correction. Many flattened reflections can now be seen from about 50 ms to as deep as 450 ms. Figure 1.11 shows 10 stacked traces near the center of the pattern. Note the dramatic improvement in frequency content and coherency compared to the brute stack (right versus left panel). The increase in S/N ratio gained by the NMO stacking process is evident in that clear and coherent reflections are visible above the noise level from 50–350 ms (right panel of Figure 1.11), which were not so obvious in the prestack data of Figure 1.8.

Figures 1.12 and 1.13 show inline stacks from all six 3-D time-lapse surveys, arranged in chronological order. Figures 1.14 and 1.15 show inline stack difference sections from all six 3-D surveys, obtained by doing an unweighted subtraction of each monitor survey from the baseline survey image. These figures illustrate that the preprocessing has been designed well enough to ensure that most seismic differences above the reservoir zone (shallower than 140 ms) are negligible, and that most significant seismic changes occur within and below the reservoir where steam injection is altering the seismic response as a function of elapsed time. The first three monitors after the baseline survey show excellent repeatability above the reservoir, while monitors 4 and 5 are becoming increasingly noisy. However, the overall results look encouraging, and will be the subject of a detailed integrated interpretation in Chapter 5.

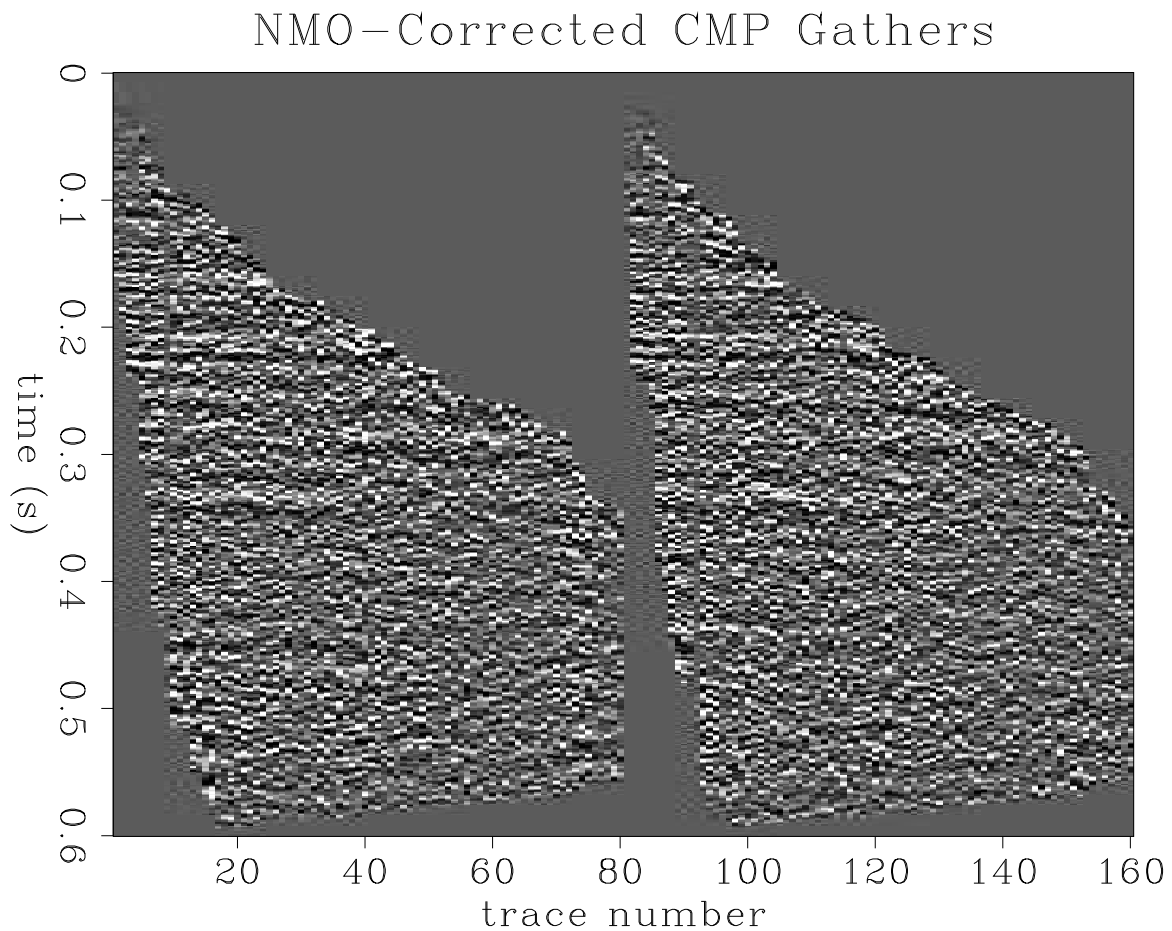


Figure 1.10: NMO-corrected CMP gathers near the center of the 3-D survey pattern.
proc-nmo [CR]

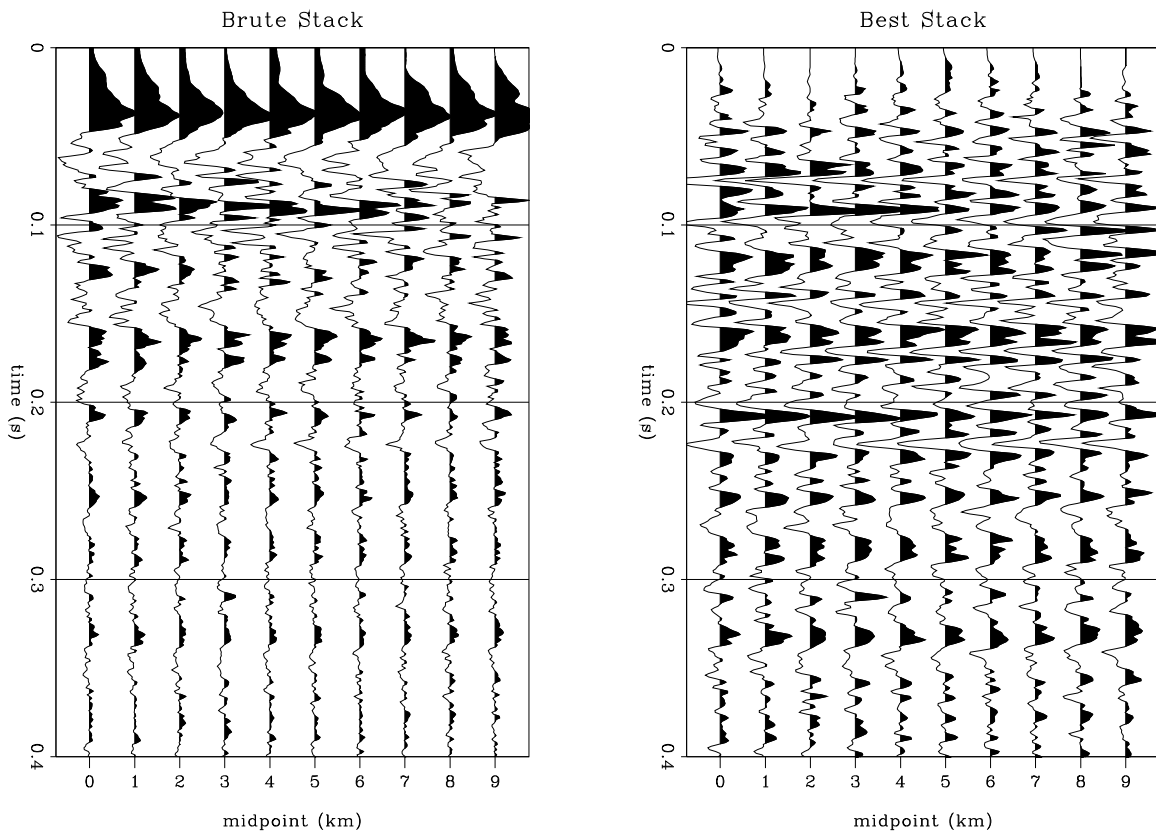


Figure 1.11: Stacked traces near the center of the 3-D survey pattern. Brute stack (left) and best stack (right). `proc-stks` [CR]




Figure 1.12: Stacked inline sections from the first three 3-D surveys. Baseline survey (left), first monitor after steam injection (center), and second monitor survey (right).

`proc-stk05a-ann` [CR]




Figure 1.13: Stacked inline sections from the last three 3-D surveys. Third monitor survey after steam injection (left), fourth monitor (center), and fifth monitor survey (right). `proc-stk05b-ann` [CR]




Figure 1.14: Stacked difference sections from the first three 3-D surveys. Baseline survey (left), first monitor difference (center), and second monitor difference (right).

`proc-stk05difa-ann` [CR]

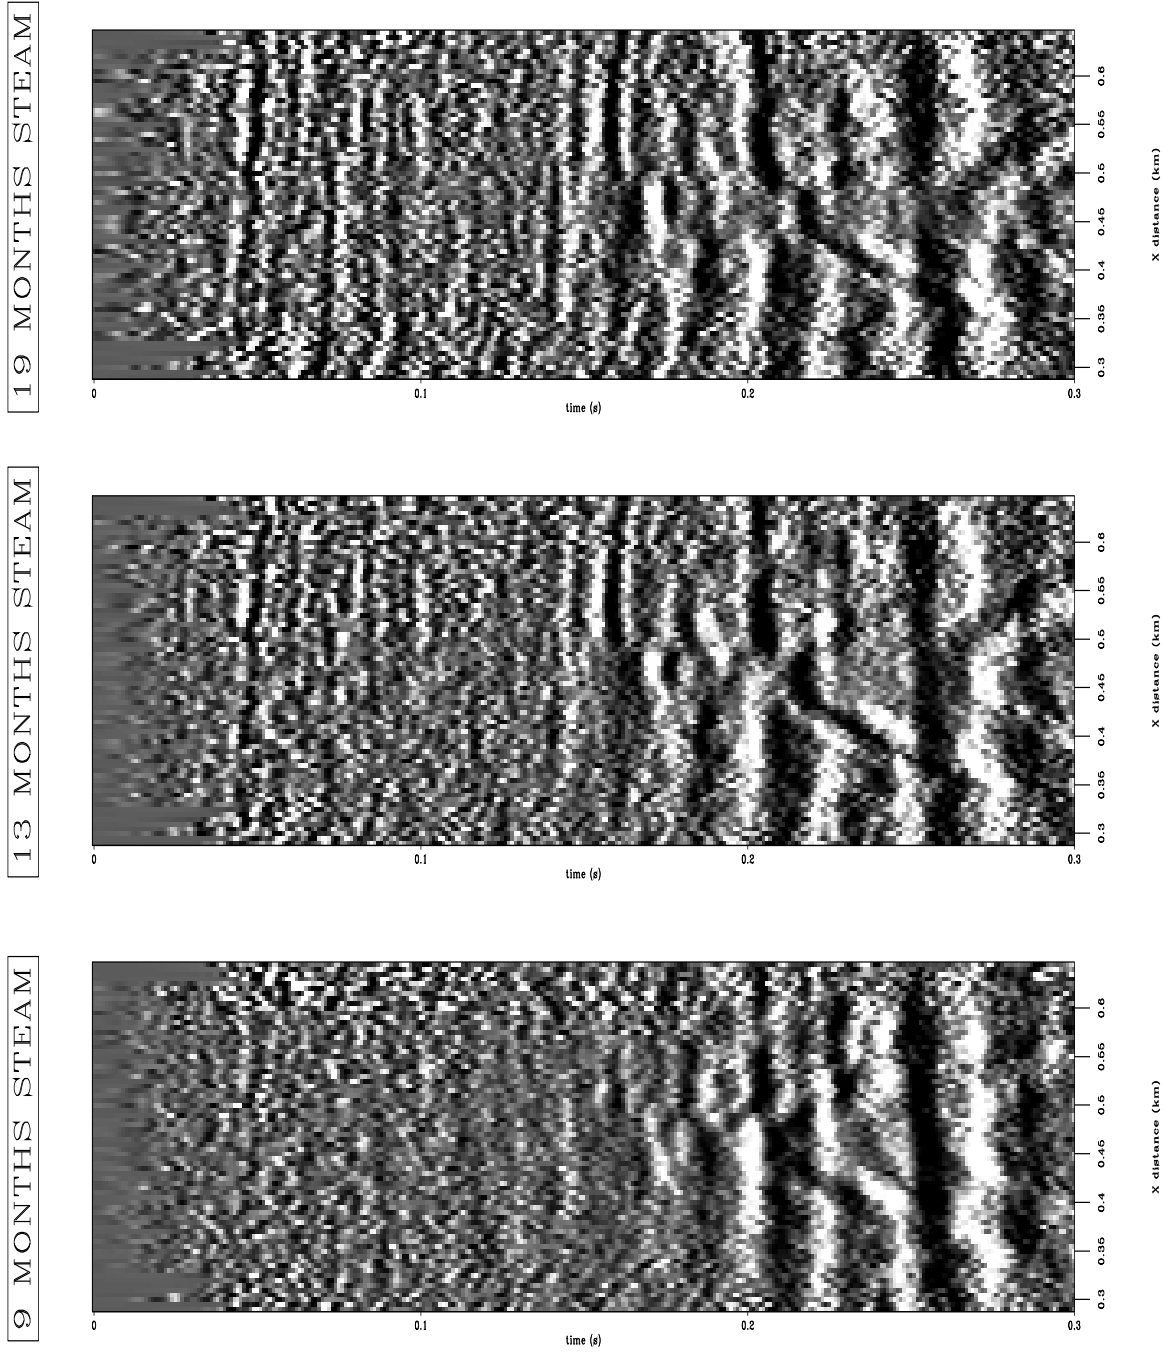


Figure 1.15: Stacked difference sections from the last three 3-D surveys. Third monitor difference (left), fourth monitor difference (center), and fifth monitor difference (right). `proc-stk05difb-ann` [CR]

1.6 Migration

Migration is a seismic imaging method that treats recorded seismic reflection data as a boundary condition to a wave equation (?). Migration theory is discussed in some detail for the Kirchhoff boundary integral method in Chapter 2. Reflectivity images obtained by seismic migration are useful in the following sense:

- Increased S/N ratio vs. stack.
- Increased resolution vs. stack.
- Improved geologic image vs. stack.
- Improved estimates of rock properties.

Migration can increase the S/N content of seismic images by at least one order of magnitude, compared to stacked sections. The resolution of features in a stacked section is on the order of a Fresnel radius, which is typically a few hundred meters in conventional surface seismic data. However, the resolution of features in a migrated image is on the order of a fraction of a seismic wavelength, which is typically tens of meters in conventional surface seismic data. Hence migration resolution is at least one order of magnitude better than stack resolution. Migration improves the geologic image compared to stack, because reflection events are more correctly located in space, and diffractions are collapsed to points. Finally, migration makes a better estimate of subsurface reflectivity, which is proportional, to first order, to relative P-wave impedance contrasts of velocity and density in rocks.

1.6.1 Signal-to-noise enhancement

Migration increases the S/N ratio over NMO stack by a factor F of about

$$F_{2D} = \sqrt{W/\Delta} \quad (1.2)$$

in 2-D, and

$$F_{3D} = \sqrt{A/\Delta^2} \quad (1.3)$$

in 3-D, where W is a 2-D width measure, A is a 3-D area measure, and Δ is the average midpoint spacing in the 2-D or 3-D data. For migration of diffractions, the width measure is the width of the diffraction hyperbola that contains most of the diffraction energy, denoted $W = \Delta x_1$ in Figure 1.16. For 3-D diffraction hyperboloids, $A \approx W^2$. For migration of reflections, the width measure is the first Fresnel zone width, denoted $W = \Delta x_2$ in Figure 1.16. A Fresnel radius is given approximately as

$$R \approx \frac{v}{2} \sqrt{\frac{t}{f}}, \quad (1.4)$$

where v is the rms velocity of the reflection event, t is its arrival time, and f is the dominant frequency of the event (?). In 2-D, the Fresnel width $W = R$, and in 3-D the Fresnel area $A \approx R^2$.

In the Duri data, the Fresnel radius at the reservoir zone is about

$$R \approx \frac{1.750}{2} \sqrt{\frac{0.200}{100}} \approx 40\text{m} . \quad (1.5)$$

Since the midpoint spacing Δ is 5 m, we can expect a S/N increase of about 2.8 for 2-D migration of reflections, and an increase in S/N of about 8.0 for 3-D migration of reflections, compared to 2-D and 3-D stacked images at Duri. The maximum width of diffraction hyperbolas emanating from steam zone is about 300 m. Given the midpoint spacing of 5 m, the increase in S/N ratio would be about 7.7 for a 2-D migrated diffraction, and a S/N increase of about 60.0 for a 3-D migrated diffraction, compared to 2-D and 3-D stacked images at Duri. Due to the significant noise levels in the prestack and stacked data, it is highly worthwhile to attempt 3-D migration of the data in order to gain the anticipated S/N improvements of 10–60 on steam zone features in the reservoir.

1.6.2 Spatial resolution

The spatial resolution of features in a stacked section is on the order of a Fresnel radius. For the Duri data, seismic features in stacked images at the reservoir cannot be resolved below a length scale of about 40 meters. The spatial resolution of features in a migrated image is on the order of 1/4th a seismic wavelength $\lambda = v/f$. This length scale is denoted Δx_3 in Figure 1.16. At the Duri reservoir depth the rms velocity is about 1.75 km/s, and the dominant frequency content is at least 100 Hz. Therefore, I anticipate resolving steamflood features down to a length scale of about 5 m, which is coincidentally the CMP trace spacing, and is approximately one order of magnitude improvement over the stack resolution of 40 meters.

1.6.3 Computational times

3-D seismic migration can be extremely computational-intensive. Before starting a 3-D migration project, it is well worthwhile to estimate the amount of computer work required. Because of irregular trace spacings and the flexibility to image subsections of the earth, Kirchhoff migration is the dominant method for 3-D seismic data sets. I will now give estimates for Kirchhoff poststack and prestack migration of 2-D and 3-D data.

To migrate a single trace into a single 2-D section, $N_x \cdot N_z \cdot N_{op}$ floating point operations (flops) are required, where $N_x \cdot N_z$ is the number of image pixels in the x - z image plane, and N_{op} is the amount of operations required to image a single pixel with a single trace. To perform a single 2-D migration therefore requires $N_t(N_x \cdot N_z \cdot N_{op})$ flops, where N_t is the number of input traces migrated. To perform a full




Figure 1.16: Width measures for migration resolution and S/N enhancement. (after Claerbout, 1985). proc-fwidth [NR]

3-D migration requires $N_t(N_x \cdot N_y \cdot N_z \cdot N_{op})$ flops, where N_y is the number of image samples in the y direction.

3-D poststack migration

A two-pass 3-D poststack migration is an efficient approximation to full 3-D poststack migration when the velocity does not vary laterally too much. This is obtained by performing 2-D migrations individually on all inline sections in the first pass, followed by 2-D migrations of the first-pass results in the crossline direction. A two-pass 3-D poststack migration requires $(N_x + N_y)[N_x \cdot N_y \cdot N_z \cdot N_{op}]$ flops, if there are N_x and N_y poststack traces in the inline and crossline directions respectively. A full one-pass 3-D poststack migration requires $(N_x \cdot N_y)[N_x \cdot N_y \cdot N_z \cdot N_{op}]$ flops under the same assumption.

The increase in computational time of one-pass versus two-pass migration is given by the factor $(N_x \cdot N_y)/(N_x + N_y)$, where N_x and N_y are the number of input poststack traces in the inline and crossline directions. Therefore, for 3-D poststack migration, a two-pass migration can be about one to three orders of magnitude faster than a full one-pass migration. In the Duri case, the number of traces in the inline and crossline directions are 72 and 71, which means that a two-pass 3-D poststack migration will be about 36 times faster than a full one-pass migration. A two-pass 3-D poststack *time* migration takes me about 1 CPU hour on an IBM RS/6000 550E workstation, therefore a full one-pass 3-D poststack *time* migration would take about 36 CPU

hours on the same computer. However, I can complete a full one-pass 3-D poststack time migration on our 32-node Connection Machine CM5 in 1.5 CPU hours, which runs my Kirchhoff code at about 500 Mflops/s.

For Kirchhoff 3-D poststack *depth* migration, the N_{op} value can be as much as 10 times the value for time migration, which includes raytracing Green's function traveltimes, and accessing and interpolating the traveltimes tables on the fly during the migration. In this case it would take about 2 CPU weeks to complete a full 3-D one-pass poststack depth migration on an RS/6000 workstation. The same result should take about 15 CPU hours on a CM5. As an aside, Kirchhoff one-pass 3-D poststack depth migration is very inefficient and would be better replaced by an $f-x$ finite-difference or PSPI phase-shift migration. Kirchhoff two-pass time/depth and one-pass time migration are competitive however.

3-D prestack migration

Kirchhoff migration currently offers the most efficient method for 3-D prestack migration of seismic data. To perform a full 3-D Kirchhoff prestack migration requires $N_t(N_x \cdot N_y \cdot N_z \cdot N_{op})$ flops, where N_t is the total number of input prestack seismic traces to be migrated, and $(N_x \cdot N_y \cdot N_z)$ is the total number of voxels in the 3-D output image. The increase in computational time of full prestack versus one-pass poststack migration is given by the factor $N_t/(N_x \cdot N_y) \approx$ fold, if there are approximately $(N_x \cdot N_y)$ poststack seismic traces. This ratio is approximately equal to the average CMP fold of the input data.

For the Duri example, there are about 150,000 prestack traces per survey, and the average fold is about 30 over the survey area. These values imply that a single 3-D prestack time migration of the Duri data would require 6 CPU weeks on a workstation, which is clearly out of the question, or about 2 CPU days on a CM5, which is feasible, but still costly. A full volume 3-D prestack depth migration of the Duri data could take as long as 30 CPU days on a CM5, which is not practical in my current computing environment.

Practical choices

The results of computational times for 3-D migration are summarized in Table 1.1. These times would need to be multiplied by at least a factor of six to complete the 3-D prestack migrations required for all six time-lapse surveys. Because of the computational times involved for all six surveys, I decided to do two-pass 3-D poststack time migrations on an RS/6000 workstation. I checked the 2-pass results by comparing with a 1-pass poststack time migration on the CM5, and the differences were so subtle that they did not justify showing here. Therefore, I am currently getting the quality of a full 1-pass 3-D time migration for the much cheaper cost of a 2-pass algorithm.

	RS/6000	CM5
POSTSTACK		
2-pass time	1 hr	—
1-pass time	36 hrs	1.5 hrs
1-pass depth	2 weeks	15 hrs
PRESTACK		
1-pass time	6 weeks	2 days
1-pass depth	1 year	3 weeks

Table 1.1: Computation times for Kirchhoff 3-D migration of a single seismic survey from the Duri Field.

To do anything more elaborate would cost much more time. A 1-pass poststack depth migration for all six surveys could be reasonably done in one weeks worth of computer time on the CM5. However, the quality of the results would depend on the accuracy of the depth interval velocity model (depth migration is less robust to velocity error than time migration). To be worthwhile, a good estimate of the lateral velocity variation within the reservoir is needed. That is not an easy task since the steamflood process causes large velocity contrasts, as will be shown in Chapter 5. Also, one would have to hope that valuable information was not destroyed during the stacking process before migration.

A 3-D prestack time migration of all six survey would take about 2 weeks worth of CM5 time. Since the overburden is nearly constant velocity, a prestack time migration would not improve much in the image above the reservoir, and would be inaccurate within and below the reservoir where depth migration is needed. A full volume prestack depth migration for all six surveys would require at least 5 months worth of computer time on our CM5, which is not practical in my current computing environment. Even if feasible, its value depends critically on the accuracy of the estimated depth interval velocity model. A 3-D prestack depth migration into a few 2-D image planes, or a few constant-offset target-oriented migrations might be a practical alternative to full volume imaging, however, and that is something I hope to explore in future work.

1.6.4 Migration operator design

To obtain optimal Kirchhoff 3-D poststack time migration images, I carefully designed the migration operators to reduce aliasing artifacts and preserve frequency content. Kirchhoff migration employs space-time migration operators which are subject to spatial aliasing artifacts (?). Lumley et al. (?) showed that while migration operator aliasing might be tolerable in some 2-D images, operator aliasing could severely degrade Kirchhoff 3-D migration images. I use the triangle filters presented by Lumley et al. to anti-alias the Kirchhoff 3-D migration operators on the fly during migration.

To design an optimal migration operator, I conducted several migration response tests. This included single-trace migration responses for aperture design, 10-trace migration responses for anti-alias filter design, 2-D inline migration tests, and finally full 3-D migration tests. The longest tests involved computing several full 3-D migrations on multiple 3-D data sets with various operator designs, and selecting the best results. For this reason, it was very helpful to have the fast two-pass algorithm.

Aperture design

I first tried to find an optimal migration operator for standard (non-anti-aliased) Kirchhoff migration. Figure 1.17 shows the standard migration operator response to a single trace. An aperture of 60 degrees was used to capture most diffraction energy without incurring too much spatial aliasing. An aperture dip-filter weight of $\cos^2 \theta$ was used to further dampen the aliasing artifacts, and to avoid truncation artifacts at the aperture limit. The characteristic “speckle” along steep portions of the operator response in Figure 1.17 is due to image-space aliasing.

In comparison, Figure 1.18 shows the anti-aliased migration operator response to a single trace. A nearly full aperture of 80 degrees was used to capture all of the available diffraction energy. An operator aperture weight of $\cos \theta$ was used to retain as much steep-dip energy as possible, while tapering smoothly to vertical dip at 90 degrees. The triangle anti-alias filters were scaled by 1.5 in both the inline and crossline directions (?). Notice that most of the speckle due to image-space aliasing is suppressed in Figure 1.18 compared to Figure 1.17.

Anti-alias filter design

I next computed the migration operator response to 10 adjacent traces. Figure 1.19 shows the standard operator response. Notice that the aliasing speckle at steep dips has not been removed by destructive interference of adjacent trace migrations. This interference speckle from migrated adjacent traces is due to operator aliasing. Figure 1.20 shows the identical response from the anti-aliased migration operator. Notice that the operator aliasing speckle has been largely suppressed, at little or no apparent loss in frequency content.

Inline migrations

Figure 1.21 shows a 2-D inline migration using the best standard migration operator parameters. The 2-D image contains a lot of noise that is related to the operator aliasing speckle. In comparison, Figure 1.22 shows the 2-D inline migration using the best anti-aliased migration operator parameters. The anti-aliased migration image is much cleaner than the standard migration, again at no apparent loss of frequency content.

Optimal design parameters

The best set of standard migration parameters I obtained were:

- standard operator: no anti-aliasing
- aperture: 60 degrees
- aperture weight: $\cos^2 \theta$

However, using an anti-aliased migration operator gave superior results to the standard migrations. Therefore, I used anti-aliased 3-D migrations for all of the final migration images shown in the remainder of this thesis. The optimal migration parameters I obtained were:

- anti-aliased operator
- anti-alias scale factors: $aafx=1.5$, $aafy=1.5$
- aperture: 80 degrees
- aperture weight: $\cos \theta$

1.6.5 Final 3-D migration images

Using the optimal migration operator design parameters listed above, including the triangle anti-alias filters, I made full 3-D poststack time migrations of each of the six time-lapse surveys. The final migrations were filtered and scaled for best display with a bandpass filter of 20–250 Hz. Figures 1.23 and 1.24 show inline sections from the full 3-D migrations of all six 3-D surveys. Figures 1.25 and 1.26 show inline difference sections obtained by subtracting a full 3-D migration of a monitor survey from the 3-D migrated cube of the baseline survey.

Notice in all the images that the S/N ratio is dramatically enhanced over the comparable stacked images of Figures 1.12–1.15, as predicted by the section on migration noise reduction. A dramatic increase in resolution is visible in the central region at about 75 ms pseudodepth, where an apparent steam leak at the steam injector borehole has caused seismic time lags due to heating. This region should have caused diffractions in the stacked sections, but they are not visible due to poor S/N ratio. In the migrated section they give a two-pixel resolution (10 m) of the lateral location of the borehole, which is an amazing sight rarely seen in seismic data.

1.7 Future work

As discussed in the section on computational times, further migration work would be interesting to pursue. The next step I would take would be to estimate a 3-D

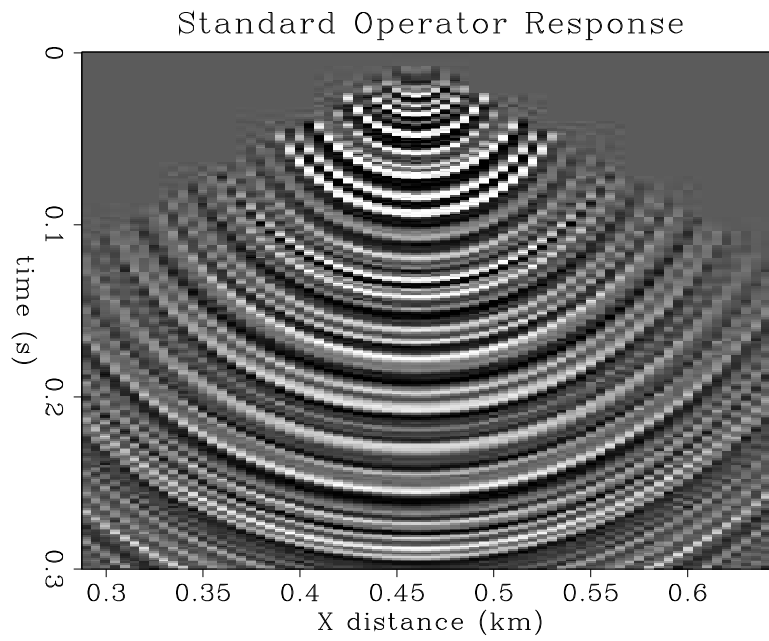


Figure 1.17: Standard migration operator response to a single trace. `proc-migst1` [CR]

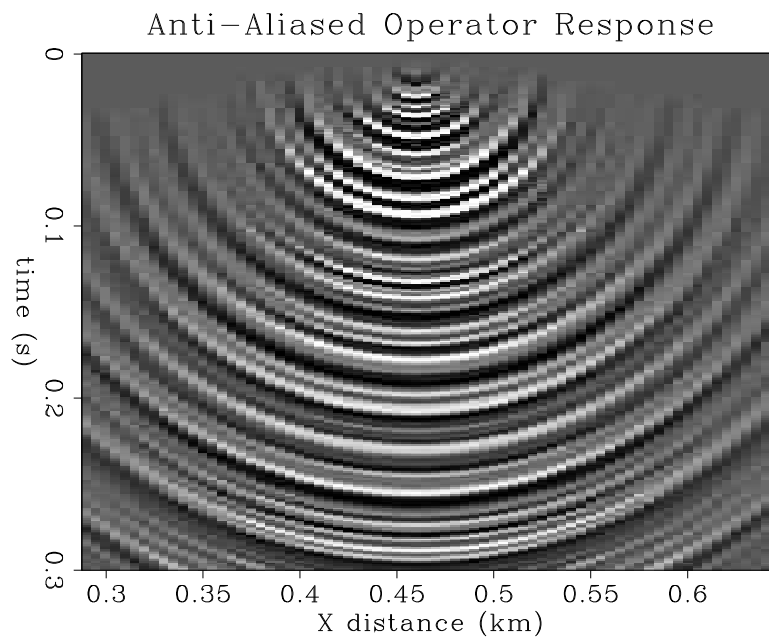


Figure 1.18: Anti-aliased migration operator response to a single trace. `proc-migaa1` [CR]




Figure 1.19: Standard migration operator response to 10 adjacent traces. proc-migst10 [CR]

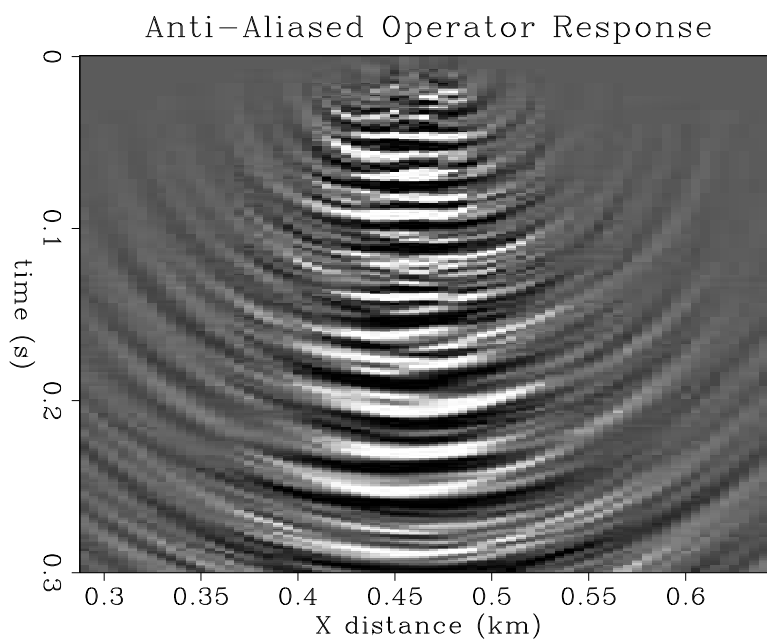


Figure 1.20: Anti-aliased migration operator response to 10 adjacent traces. proc-migaa10 [CR]

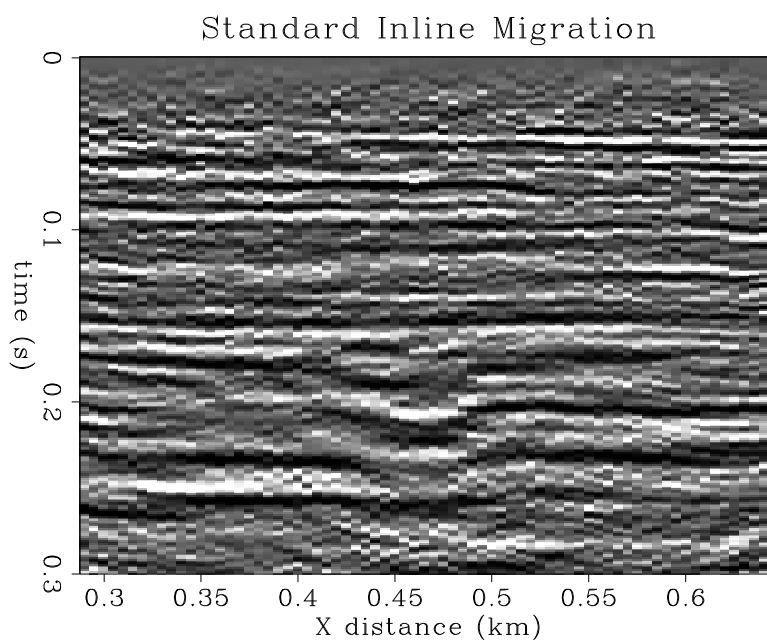


Figure 1.21: Standard 2-D inline migration. `proc-migst35` [CR]

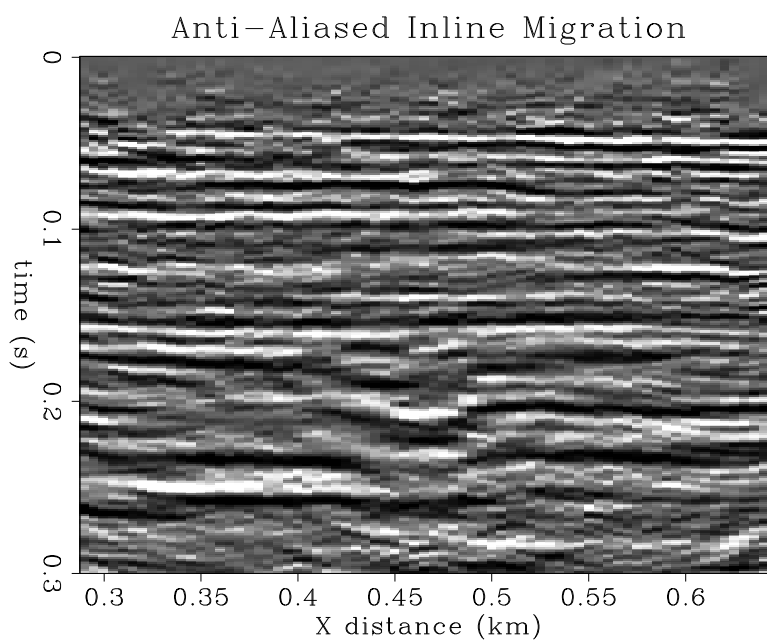


Figure 1.22: Anti-aliased 2-D inline migration. `proc-migaa35` [CR]




Figure 1.23: 3-D migrated inlines from the first three 3-D surveys. Baseline survey (left), first monitor after steam injection (center), and second monitor survey (right).

`proc-mig05a-ann` [CR]




Figure 1.24: 3-D migrated inlines from the last three 3-D surveys. Third monitor survey after steam injection (left), fourth monitor (center), and fifth monitor survey (right). `proc-mig05b-ann` [CR]




Figure 1.25: 3-D migration difference sections from the first three 3-D surveys. Baseline survey (left), first monitor difference (center), and second monitor difference (right). `proc-mig05difa-ann` [CR]




Figure 1.26: 3-D migration difference sections from the last three 3-D surveys. Third monitor difference (left), fourth monitor difference (center), and fifth monitor difference (right). `proc-mig05difb-ann` [CR]

stacking velocity function $v(x, y, t)$ for each survey independently, and then repeat the 3-D poststack time migrations. This might help better focus features within the reservoir, and remove some of the time shifts in the overburden of surveys 5 and 6 which add noise to the difference sections. A 3-D prestack time migration could then be performed on one of the monitor surveys, Monitor 2 for example, to determine if prestack time migration adds any clarity to the images compared to poststack time migration. If so, all six surveys could be prestack time-migrated.

Following that, I would convert each $v(x, y, t)$ rms velocity field to interval velocity $v(x, y, z)$ and try a 3-D poststack depth migration. Furthermore, this could give very interesting results if combined with a reasonable estimate of the depth velocity variations within the reservoir (see Chapter 5). At the very least, this would allow a practical time-to-depth conversion of stacked data to the depth domain, wherein migrated depth slices could be examined and interpreted in physical depth.

Of particular interest is to attempt a 3-D prestack depth migration into a small target volume, as opposed to a full volume image which is prohibitive. This could be done in an iterative manner to try and estimate the depth interval velocity field directly within the reservoir using prestack depth migration focusing analysis.

Finally, the layout of the sources and receivers makes this survey ideal for attempting 3-D plane-wave migration. As described by Schultz and Claerbout (?), plane-wave sources are simulated at the surface by summing individual time-delayed shot gathers appropriately. The resulting plane-wave section has about the same S/N increase as an NMO stacked section, but makes no assumption (and therefore no error) with respect to the subsurface velocity. Each plane-wave section could then be 3-D migrated by a finite-difference of phase-shift method for about the cost of two poststack depth migrations. This procedure could be iterated for each trial depth interval velocity model, and so constitute an efficient method for performing 3-D depth migration velocity analysis.

1.8 Conclusions

In this chapter I have discussed the desirable features of data processing for time-lapse seismic monitor surveys in terms of experimental repeatability, survey positioning accuracy, wavelet consistency, amplitude preservation, velocity accuracy and depth accuracy. I have discussed 4-D acquisition design using the example of the Duri Field experiment, and preprocessing issues related to maintaining consistency between processed time-lapse images and difference sections. Finally, I have discussed advantages in spatial resolution and signal-to-noise enhancement of 3-D migration compared to NMO stacked images, migration computational time issues, and optimal design of anti-aliased 3-D migration operators.

1.9 Acknowledgments

I would like to thank Steve Jenkins of CalTex Petroleum Indonesia for discussions related to the Duri 4-D data set, and for providing a copy of their IPA paper on the subject (?). I also thank Fred Herkenhoff of Chevron Overseas for getting me involved in the Duri project, and helping to release the data to me at Stanford. Ray Ergas, Steve Cole and Jeff Wright of Chevron Research were also helpful in providing data and ancillary information.