

# Seismic Explorations

## 12. Seismic Explorations

Graham Ganssle

Geophysics is an extremely diverse subject including fields of study with backgrounds in chemistry, geology, biology, and physics. This chapter attempts to show the reader a broad cross section of the geophysics of petroleum exploration. Three fundamental steps exist in the workflow of geophysical petroleum exploration. First, there is the acquisition of seismic data by using initiating acoustic energy into the Earth's subsurface and recording the reflected energy. Second, this raw acoustic data is processed using the physical and mathematical principles of wave propagation and continuum mechanics. Finally the processed seismic data is interpreted. Interpretation endeavors

12.1	<b>Seismic Data Acquisition</b> .....	420
12.2	<b>Seismic Data Processing</b> .....	423
12.3	<b>Seismic Data Interpretation</b> .....	426
12.4	<b>Summary</b> .....	431
	<b>References</b> .....	431

to create a pseudo-geological model from the geophysical data for verification or update by purely geological data. The model is then iteratively updated by both geologists and geophysicists to form a complete risk assessment of a prospective drill location.

Seismic activity is an ongoing phenomenon. Earthquakes, caused by fault slip, are a recurring reminder of seismic activity in the Earth's subsurface. In reality, seismic activity is a result of the motion of any object. An object making contact with another object creates sound waves, especially when moving. These waves, in the case of earthquakes, can have very large amplitudes. Other seismic wave sources can create sound waves with quite small amplitudes. For example, every time a person walks down the street their footsteps on the pavement are creating very small seismic disturbances. The term *seismic* is used to describe the peak frequency of the resulting sound waves. Seismic waves are generally thought of as having low frequency (in the range of 5–8 Hz) and large amplitude. It is possible to measure seismic disturbances from earthquakes to footsteps on pavement using the right equipment. There are three major types of seismic waves. P-waves, also called primary or pressure waves, are familiar to everyone in the form of sound. Nearly all seismic exploration for petroleum uses p-waves. S-waves, also called secondary or shear waves, are slower and can only travel through solids. P-waves and s-waves are collectively called body waves because they pass through the Earth. Surface waves travel along the surface, similar to an ocean wave. Surface waves are not used for petroleum exploration, but they are extensively studied because

they cause most of the damage during earthquakes.

The velocity of body waves is related to the properties of the medium they pass through by the following equations

$$V_p = \sqrt{\frac{K + 4/3\mu}{\rho}} \quad V_s = \sqrt{\frac{\mu}{\rho}}, \quad (12.1)$$

where  $K$  is compressibility,  $\mu$  is rigidity, and  $\rho$  is density. If we know both p-wave and s-wave velocities and can make some reasonable assumption about the density, we can determine the mechanical properties of the medium.

The measurement of seismic disturbances is carried out by measuring the motion of the ground for an extended period of time. This creates a record in time with each sample point being the magnitude of this motion. Recording from many locations arranged in a line or an array provides a record that approximates the local subsurface geology. When looking at a seismic data line the horizontal axis represents some spatial dimension,  $x$ , and the vertical axis represents the time dimension,  $t$ . At each  $(x, t)$  point the data is an amplitude. This amplitude is related to the reflection strength of the layer in the Earth's subsurface located at time  $t$ . In a perfect world, each reflection event on a seismic section would

represent one geologic bed, giving a perfect view of the Earth's subsurface.

Nowadays commercial seismic data is widely used in industry for the exploration of petroleum oil and gas. It is also used for the exploration of salt, certain hard-rock deposits, and engineering site investigation. The use of seismic data has proven to be so successful in petroleum exploration in the past 100 years that there are few prospects which have not had seismic shot over them at some point. Seismic data can be used in the exploration of new prospects, in the delineation of current targets, and in the classification of reservoirs. Therefore, it is extremely beneficial to use seismic data in any stage of target planning. However, the acquisition, processing, and interpretation of seismic data can be quite expensive. Because of this high cost, other techniques are employed. In large scale remote areas, gravity survey methods are often used to determine

where a seismic survey should be acquired. In areas with a large number of existing wells, the well data may be dense enough to eliminate the need for seismic data.

Inferring information about subsurface geology without going underground to take physical samples is a difficult task. Like all remote sensing problems, data is taken from a reference point far from the area of interest. There are many layers of interfering medium between the layer of interest and Earth's surface distorting the desired data. There are an incredibly large number of noise sources contaminating the data while it is being recorded. At the same time, the medium being recorded through is dispersing and absorbing the signal energy. These points demonstrate some of the reasons why a seismic survey is so expensive. They also fundamentally constrain the way seismic data must be acquired, processed, and interpreted after acquisition.

## 12.1 Seismic Data Acquisition

There are two major arenas of seismic geometry. The first is two-dimensional (2-D) seismic data. 2-D data is one line of data running in a direction of interest. The data looks like a panel of amplitudes, extending horizontally out in the direction of the line, and extending down in the direction of recording time. The second is three-dimensional (3-D) seismic data. 3-D data is a volume of data covering a lateral extent around the area of interest. The data looks like a 3-D volume of amplitudes, extending horizontally out in two directions, and extending down in the direction of recording time. 3-D data may be viewed on a computer in three dimensions or in cross section as a collection of 2-D lines. Seismic data specialists refer to increasing subsurface depth in terms of increasing recording time.

As with any specialty, the seismic industry has a wealth of nomenclature. In addition to 2-D and 3-D data there is four-dimensional (4-D) data. 4-D data does not implement a different recording geometry, like the difference between 2-D and 3-D, but is rather recorded repeatedly. 4-D data is recording a 3-D seismic survey, and then recording the same 3-D survey after some years to observe differences in the subsurface. The goal is to obtain information about the depletion of reservoirs and other time dependent processes. So, 4-D seismic is actually just two or more 3-D surveys acquired (preferentially in the same way) over the same area. There is also five-dimensional (5-D) terminology that is used in the seismic industry. 5-D is not a different recording geometry, but is a way of interpolating seismic data for special use in seismic data processing. By

means of these examples, it is realized that some of the nomenclature is overlapping and nonunique.

The initial requirement in recording seismic data is that there must be energy to record. Therefore, the acquisition of seismic data must start with a seismic source. In some cases, the seismic source is a natural one, as in the case of earthquakes. These scenarios are called passive seismic recording, as there is no operator controlled input energy. The acquisition of natural passive seismic activity is quite different than that of active seismic recording used in petroleum exploration. For example, the epicenter location of an earthquake is not known a priori, which creates a need for wide receiver spacing and continuous recording. After an earthquake has occurred, seismologists can look back in the seismic records from several stations to triangulate the signal to figure out the location of the epicenter.

Active seismic energy recording necessitates different acquisition parameters. In this case, the source position is known, so the receivers can be located according to signal requirements. The goal is to have a source that gives the best signal possible. A good seismic signal can be characterized by many features; several important ones are as follows. First, the source must transmit the bulk of its output energy into the ground. The energy that is not input into the ground is wasted for the purposes of subsurface imaging. Second, the source must input a signal of sufficient bandwidth to be used to image the strata in the exploration of petroleum. If the interpreters require a signal with a peak frequency of 40 Hz, for example, the source

may have to input a useable signal of up to 200 Hz depending on the geology. Third, the source must input enough energy into the ground that the signal is not degraded to the point of being undetectable. This parameter is a function of the absorption and dispersion characteristics of the local geology and the noise present in the area. There are many other constraints on sources, which in most cases can be controlled to an order of accuracy that is determined while planning the acquisition parameters.

There are three common ways to get the energy into the ground. These are dynamite, airgun, and vibroseis. If the survey is of an area on land, dynamite or vibroseis may be used. If the survey is over an area covered by water airguns are usually the preferred method, but in some shallow water instances dynamite may be used. Each of these three main types of sources has a huge variety of types, each meant for specific applications. Despite this large variety, dynamite, airgun, and vibroseis sources can be analyzed as groups with different signal characteristics.

Dynamite sources are used both on land and in shallow water surveys. The acquisition field crew drills holes into the ground where they place the dynamite. The holes are drilled to a depth sufficient to get the dynamite below the unconsolidated surface sediments, and into bedrock. This ensures good signal transmission from the dynamite into the ground. Unfortunately, dynamite does waste a good deal of energy in upward traveling waves going up into the air. However, dynamite is an excellent seismic source because of its desirable signal characteristics. The signal strength, or amplitude, can be very large. To get a larger signal amplitude, a larger charge or a combination of several charges can be used, or an array of charges can be laid out. Dynamite sources output an impulsive signal that has a large bandwidth. Short-duration, broadband inputs assist the seismic data processors in creating an accurate subsurface image.

Airguns are used in water surveys from shallow to ultra-deep water. The guns are towed behind a *shot boat* and are usually aligned in an array or linear pattern behind the boat. Airguns output less energy than dynamite, but can be very large in the case of deep water surveys. Because of the fluid coupling between the water and the seafloor, most of the energy the guns output is transferred into a usable signal. Some energy is lost to the atmosphere via transmission from the sea surface. Inputting more energy into the ground requires using a larger array of airguns, which can be costly. The signal from airguns is dependent on the design of the gun itself, the array design in which the guns are towed, and the timing of the individual guns. Just as dynamite, the signal is broadband impulsive.

The term Vibroseis came from an actual seismic acquisition company's brand name. Many other seismic acquisition companies now use similar designs, but the original name was so ubiquitous that it is generally used as a description of the process. The vibroseis method uses a heavy truck with a plate on the bottom attached to a vibrator driven by compressed air. The truck is moved into position, the plate is lowered to the ground and extended, lifting the truck off the ground. The plate then vibrates up and down in a specific pattern causing the ground to move as well. The vibration pattern is controlled by regulating the flow of compressed air with a servo. These trucks are usually lined up in arrays, and the vibrating pattern is decided very carefully to input a specific seismic signal. Of the three methods discussed, vibroseis inputs the smallest energy to the ground, making it ideal for use in populated areas where dynamite is not feasible. The signal obviously can be exactly specified, making it arguably the best signal shape of all of the seismic sources. The majority of vibroseis signals are called chirps. A chirp is a sinusoidal waveform that increases frequency as time increases. A simple chirp could be a sine wave with an argument of time squared; these signals sound a little like birds chirping to the human ear, hence the name. Chirps can be deconvolved with precision by the seismic processor, making vibroseis signals very useful.

Any of the three above mentioned sources inputs a seismic signal into the ground, where the signal is (among other things) reflected back to the Earth's surface. To gain any useful benefit from this system, receivers must be used to record the reflected signal. Two types of receivers are used: geophones and hydrophones. Geophones are used on land and hydrophones are used in water. Nowadays, geophones are sometimes used in the water as well. In deep water and ultra deep water plays, sometimes companies will use what are called ocean bottom nodes. Ocean bottom nodes are geophones placed on the seafloor by a remotely operated vehicle (ROV) robot. Airguns are used as a source in these instances. Ocean bottom node recording is extremely expensive, as thousands of geophones must be laid out by ROV operators and picked back up after recording is complete. The data quality however, is breathtaking.

Most geophones measure the vertical component of velocity only, although 3-component geophones are available. They record the motion at the Earth's surface as a function of time. Traditionally, only the vertical dimension of ground surface velocity is recorded to build a one-dimensional magnitude record. In some cases, all three dimensions are recorded separately to form what is called 3-component data. This data can be used to infer properties of the subsurface media by the motion of

shear waves. Geophones are laid out on the surface in groups, each of which is combined into a single trace. Each group may be an array or simply a cluster, and the groups may be laid out over an area for 3-D data or in a line for 2-D data. The location of each geophone group is recorded by a surveyor, or by a GPS unit attached to the geophone node.

Hydrophones are pressure vessels calibrated to record differences in the hydrodynamic pressure in the water in which they are towed. They are enclosed in long, flexible, oil-filled cables called streamers which are towed behind streamer boats at a depth of 5 to about 50 m. The streamers can be many miles long, so each streamer contains many hydrophones, several steering units capable of steering the streamer in depth and crosstrack (relative to the streamer boat's heading), and GPS units down the streamer to record position information. The streamer boats and the shot boats drive around in many different complicated patterns to accomplish an exact acquisition geometry decided in advance by a design team.

The geometry of the survey is of utmost importance. There is no point recording a seismic survey which does not most efficiently illuminate the subsurface area of interest so the geometry not only needs to be designed properly, but geometry also must be recorded in the field. The number one mishap in analyzing seismic data is incorrectly positioning the data relative to the original parameters. When seismic is recorded on land, the receivers and sources are laid out, their positions are recorded, and then the data recorded. In water, the shot boats and streamer boats both have GPS systems on the guns and the hydrophones to create a continuous record of position data along with the seismic data. During data processing, these positioning data are used to ensure that the final seismic section correctly indicates where to drill. The positioning data also constrains the way the seismic data must be handled during certain processing steps. If the recording locations for the seismic data are incorrect, not only will the resulting seismic section indicate incorrect locations to drill, it may also have erroneous structure and amplitude character as is demonstrated in Sect. 12.2.

When seismic data were first being recorded in the early twentieth century, the records were essentially the result of one source and one receiver being used at a time. This is called single density, or single-fold data. In modern seismic data, there are many receivers recording when a source goes off. The number of channels recording each subsurface location is the fold of the data. If the source and receiver spacings are equal, the fold will be one-half of the number of channels recorded from each shot. Today's seismic data is often recorded at up to 300-fold, an indication of why seismic

data sets are so large. In a typical land survey covering a 100 square mile area, you may have 10 000 sources with 300 records per source, for a total of three million records, also called traces. A typical sample interval is 2 ms, so if the trace lengths are (a typical) 8 s there would be 4000 samples per trace, for a total survey size of 12 billion sample points. This is hundreds of gigabytes of data even in a compressed format. The result is that it takes an incredible amount of computing power to deal with seismic data. Even still, the reason seismic data are recorded at such high fold is that data density is of paramount importance. As shown in Sect. 12.2, increasing the fold of the data input increases the quality of the data output.

In a hypothetical 3-D land survey, the recording acquisition geometry is a square. When shooting a source at the exact center of the square, which geophones should be recording? There are many answers to this question, depending on design parameters. One simple geometry, called split spread recording, would have 150 geophones recording in a line extending to the west of the source and a mirrored 150 geophones in a line to the east. This geometry would produce a fold of 300 if the shot spacing is one-half of the receiver spacing. This is very high fold and would hopefully produce a quality seismic section. The linear nature of this geometry however, means that there is only one direction, or azimuth, recorded. Azimuth is defined as a compass direction from source to receiver in the seismic survey.

The benefit of multiazimuthal acquisition geometries is to illuminate anisotropic properties of the rock. All subsurface beds are heterogeneous, meaning that they are inconsistent in lateral extent, for example, they have horizontally distributed inconsistencies. This is also true when looking in the vertical direction, as we have beds of varying rocks. Some rocks however exhibit anisotropy. Anisotropy is a change in seismic velocity as a function of direction. Nearly all sedimentary rocks have different seismic velocities in the vertical and horizontal directions. Fractured units often display azimuthal anisotropy. Multiazimuthal seismic data recording geometries can provide a characterization of the azimuthal anisotropy in subsurface beds. This happens because sound travels with different velocities through fractured and unfractured rock of the same composition. For example, imagine a shale with natural linear fracture planes oriented in the north-south direction. Sound waves would propagate with a certain speed through this shale in the north-south direction, but would be apparently slowed down in the east-west direction because of the interference created when propagating across the fractured layers. Because of this property, many modern acquisition geometries are laid out with a large azimuthal range.

Acquisition of seismic data is the first step in creating a seismic section that represents the subsurface geology. However, because of noise, waveform character, and geometrical considerations, a raw seismic record straight from the geophone does not look like a geologic section. It does look like a long record of

noise. Combining the large number of traces from the above hypothetical 3-D seismic example exhibits why the raw records of the full survey are incomprehensible. To get these noisy records to approximate anything resembling geology, seismic processing specialists must do an extremely large amount of number crunching.

## 12.2 Seismic Data Processing

Seismic data processing is more naturally accomplished in a different set of dimensions than that of the data acquisition dimensions. When acquiring seismic data, the directions to organize the data as it is streamed onto the hardware are the source, receiver, and time domains. This is because, as shown in Sect. 12.1, a source goes off while receivers are listening. The field crew records the live channels simultaneously as a function of time.

In many cases, it is more elegant mathematically to handle seismic data during processing in the offset, midpoint, and time domains. The offset domain is defined in terms of the linear distance between the source and the receiver. The midpoint domain is defined in terms of the imaginary distance halfway between source and receiver, at the surface. This domain is orthogonal to the offset domain. The time domain is the same as in the acquisition domains.

When the data comes in from the field it is often multiplexed, which means the bits from each channel are striped across one record. The seismic processor then must demultiplex the data, and sort it into midpoint, offset, time coordinates. This process is usually incorporated into the same steps as defining the geometry on the data. Each trace in a seismic data file includes a trace header which includes, among other information, the coordinates of the source and the receiver. Using the observer notes from the survey crew in the field, the data processor fills the geometry into each of these trace headers using any information available. This information may include the shotpoint, the record field ID number, or sometimes the actual GPS position of the stations themselves. Along with some other ancillary steps, these initial definitions and sorting are collectively named preprocessing.

There are many steps in the seismic data processing sequence that affect the data primarily in one of the three processing domains, midpoint, offset, and time. There are others that affect the data in more than one domain at a time, as they are intrinsically linked to the position of the data in both time and space. Each one of the dimensions can be discussed separately in many cases; however, the processing steps taken in an actual workflow can be in any order, and may be repet-

itive. A typical seismic data processing workflow can be thousands of steps long to get the data looking just right.

In the time domain among hundreds of processes that could be applied, there are three which stand out as being the categorical leaders. These are gaining, filtering, and deconvolution. All of these three processes affect the data only in the time domain, with no spatial effects in the midpoint or offset domains.

Gaining the data can be accomplished in many different ways. One of the first gain processes applied to seismic data is called geometrical spreading correction. Geometrical spreading correction gain is applied to correct for the natural absorptive effects of the Earth. As the acoustic waves propagate from the source to the receiver through rock, some of the energy is lost due to heat dissipation. The farther these waves must travel, the more energy is lost. Therefore, the late times on seismic records are always lower amplitude than the early times. Geometrical spreading correction gain is usually applied to the data as a time to the exponent gain. In many cases, a good starting guess for the exponent is two. Processors usually then scale the exponent up or down to make the records look even.

Filtering the data can be accomplished in thousands of ways. A common process uses a band-pass filter to allow a certain frequency spectrum of data to remain by eliminating everything else. For these types of filters, the data processor will apply Fourier transform to convert the data from the time domain to the temporal frequency domain, and then multiply some high and low frequency components by zero. The inverse Fourier transform is then applied, and the data is examined for content. Another filter is applied in frequency-wavenumber space, where a 2-D Fourier transform is applied and both dimensions are filtered. There are countless other 2-D filtering techniques, including methods in the time-space domain, the frequency-space domain, and the time-wavenumber domain.

Deconvolution of seismic data can be used to accomplish several things. Two major deconvolution processes are called spiking deconvolution and predictive error filtering. Spiking deconvolution aims to elimi-

nate the incoming signal waveform from the seismic source, called a wavelet, from the data. Spiking deconvolution generally tends to make the data appear to have sharper peaks and troughs on the waveform. Predictive error filtering deconvolution approximates and removes repetitive ringing from the waveform. This ringing could represent several things, but in many cases the ringing in a seismic record is produced by the propagation of multiple reflections, often simply called multiples. Multiples are caused by acoustic energy bouncing back and forth off a pair of reflectors, for example, the water bottom and the water surface in a marine seismic survey.

The offset domain is handled by many processes that affect the data only in the offset direction, but also some processes that also affect the data in the time and/or midpoint directions. The velocity of seismic wave propagation is of critical importance to the seismic data processor. The change in velocity and density between different regions in the Earth's subsurface is the cause of a seismic response. It is in the offset domain that seismic data processors estimate seismic wave propagation and correct for certain types of velocity changes.

When a source and a receiver are located at the same position, the acoustic energy measured propagates vertically. When we introduce a horizontal displacement, or offset, between the source and the receiver the acoustic energy measured has necessarily traveled not only in the vertical direction, but also in the direction of the offset. To see this, the seismic processor organizes the data traces into a group where all the energy has come from one source, and the traces increase in offset with respect to the horizontal axis. In many cases this offset is measured from the midpoint between the source and the receiver. In this case this group, or gather, of traces is called a common midpoint gather (CMP gather). CMP gathers are used to estimate the velocity with which acoustic energy propagates in the local media surrounding the midpoint in question.

Once the data is grouped into CMP gathers, the data processor begins to see a hyperbolic moveout response of the individual reflection events. Moveout is the increase in arrival time as a function of offset with each trace in a CMP (or other types of) gathers. The data processor can use the slope of the hyperbolas on the CMP gathers to estimate the velocity of the wave propagation. This is done by using a variable hyperbolic function as a function of time to sum, or stack, the traces in a CMP gather. The function's slope is varied many times and the semblance, or power average, is computed for each sample in time. Each semblance is then displayed as a function of velocity on a semblance panel. The data processor picks the highest semblance value,

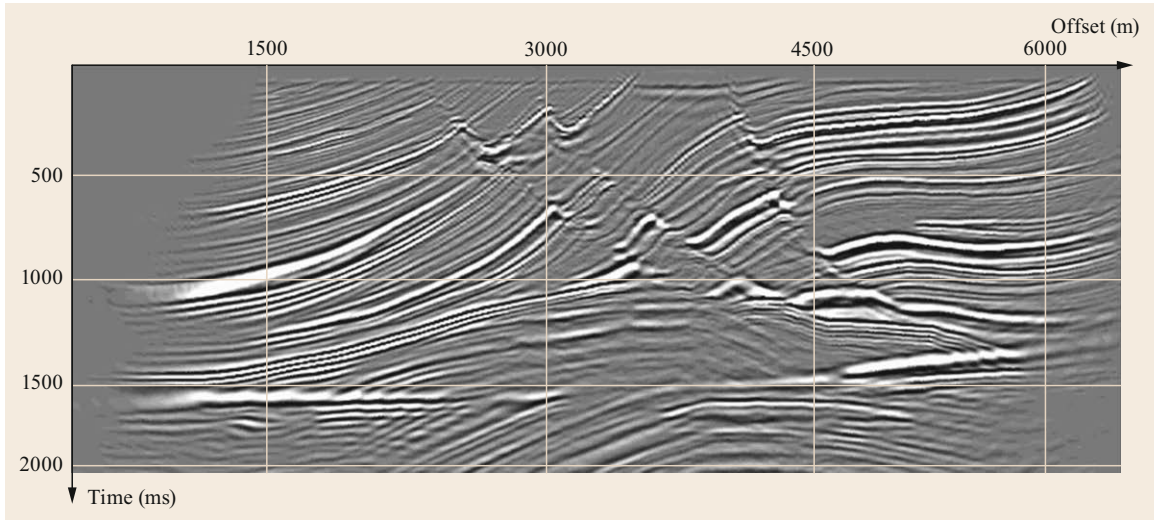
which associates the velocity value at that time sample back to the data. The CMP gathers are then corrected in time according to the slope of each hyperbola so that the seismic events appear linear rather than hyperbolic. This process is called normal moveout (NMO) correction and is an integral part of seismic data processing.

Once the data has been corrected for normal moveout, the CMP gathers appear to have flat events. The traces in the gather are then summed and multiplied by some weighting function. This process is called stacking. Stacking a CMP gather creates one trace. This is done to reduce noise and increase signal quality. Every data trace in the seismic survey before this point has been sorted into CMP gathers. Each CMP gather is stacked, creating one trace per gather, and dramatically reducing the size of the data set.

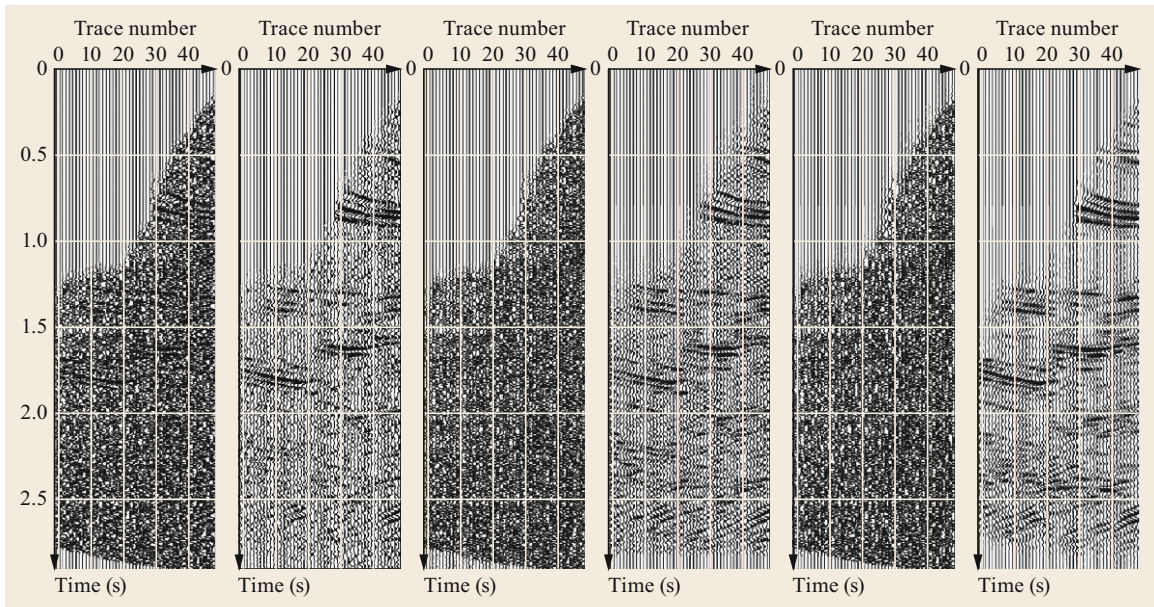
The gathers in Fig. 12.2 correspond to the stacked section in Fig. 12.1. After normal moveout correction has been applied to a set of CMP gathers they are stacked. This means that every gather is summed into one trace on the stacked seismic section. In the example of Fig. 12.2, after stacking these six CMP gathers would become six traces in the stacked section of Fig. 12.1.

Assuming several other processes have been completed accurately, and that the geologic dips are not too steep, stacking the data approximates a zero offset section. Therefore the resulting seismic section after stack represents the subsurface geology. The data processor will output the data at this point, called a structure section, for a seismic data interpreter to look at to confirm whether or not the seismic data is on par with the subsurface model built from existing well data.

The midpoint domain processes are arguably the most interesting. There are several processes that affect the midpoint domain, but the most major of these is called migration. Migration processes strive to eliminate the effects of diffraction and sometimes refraction on the acoustic energy. There are a number of different migration algorithms, but all of them require a model of the seismic velocity throughout the survey area. Migration translates energy through the time, offset, and midpoint domains, although migration techniques are traditionally studied in the category of midpoint domain processes. Because of the movement of energy on the seismic section in the midpoint domain, the data processor must be very careful to apply the correct geometry before the migration inversion processes are implemented. Because the energy moves horizontally, the survey must extend a distance, called the migration aperture, beyond the area of interest. If the geologic dips are steep or lateral velocity variations are present, migration should be applied instead of the normal moveout correction before stacking. In the case of pre-stack migration inversion



**Fig. 12.1** A 2-D seismic cross section (Reprinted with permission after [12.1])



**Fig. 12.2** A collection of CMP gathers corresponding to the seismic section in Fig. 12.1 (after [12.2])

techniques, the algorithm to complete the migration must in some cases apply geometry corrections to the trace headers.

Acoustic energy appears distorted on unmigrated seismic records. This distortion falls into several categories: first, approximate point reflectors in the subsurface cause acoustic energy to appear diffracted into hyperbolas, second, acoustic energy reflected off of dipping (sloped with respect to horizontal) reflectors appears to be translated in the down-dip direction with a smaller slope than reality, and third, all of this transla-

tion of energy creates an apparent change in amplitude of the seismic data. Migration processes work to fix all of these distortions. Migration is one of the largest subtopics of study in the seismic data processing arena, and as such there are many different approaches to implement migration. Before understanding how each of these processes works, a clear understanding of each distortion type is needed.

Diffraction of waves occurs in any wave dynamical system. The term diffraction describes the radial spreading of wave energy as it propagates away from

the spreading point. A simple analogy to the diffraction of seismic acoustic energy is a hole in a seawall off of a beach. Imagine there is a beach with a parallel seawall 100 m out from shore to stop waves. If there was a 1 m long section of seawall missing, when plane waves hit the seawall from the ocean side some of the waves would pass through the wall where the 1 m hole was. It is natural to automatically assume that beachgoers would see a flat sea with a wide wave hitting the shore. In point of fact, this is not the case. The energy from the incoming plane wave would be diffracted radially at the 1 m hole in the wall, causing a semicircular wavefront to propagate toward shore. If the beachgoers stood in a line down the beach and recorded the time at which the wavefront hit the shore in front of them, they would see a hyperbolic dependence in time with respect to their perpendicular distance from the hole in the seawall.

The same thing happens with seismic data. If there is a subsurface point diffractor the seismic receivers, which are laid out at even surface intervals extending away from the source, will record the diffracted energy with a hyperbolic dependence in time with respect to their offset from the diffractor. After stacking the data, the seismic section will show a hyperbola emanating from where the point diffractor is located in the subsurface. Assuming for the moment that the subsurface acoustic velocity is constant in the vertical and horizontal directions, the apex of the hyperbola will appear at the location of the diffractor.

The apparent shift in the location of seismic events along a dipping reflector is related to the diffraction response. The receivers only record acoustic energy as a function of time, therefore there is no directional component to their measurement. So the energy observed

after reflection from a dipping interface will appear skewed in the time and midpoint domains to be from a longer, less steep interface. This effect can also be accounted for by a process called dip moveout correction (DMO). DMO is similar to normal moveout correction with correction for dipping reflectors. Using DMO instead of normal moveout before stacking, followed by migration after stack, produces a result comparable to pre-stack migration if the lateral velocity variation is not too large. Because acoustic energy is diffracted and skewed in the subsurface, migration algorithms must also correct for an apparent loss of amplitude.

As was previously mentioned, there are many strategies for implementing migration algorithms. Some migration strategies use finite differencing methods, which can output a seismic section in depth or in time. Finite difference migration algorithms commonly use an approach to solving the wave equation by predicting the next sample point in the output section using the current sample point in the input section. Phase shift migration algorithms transform the input seismic section into frequency space to downward continue the wavefield by a specific interval for each step. After imaging has been completed, the output section is inverse transformed back into the time domain. Collapsing diffractions by a downward continuation and imaging strategy is analogous to using the focus ring on a camera. When the lenses are in the correct alignment the camera has focused the picture. When the continuation step in a migration algorithm has collapsed the diffractions to a point, the algorithm has imaged the data. While time migration corrects for diffraction, depth migration is necessary to correct for refraction effects such as the bending of seismic waves as they pass through an irregular salt body.

## 12.3 Seismic Data Interpretation

Once the data has been processed, a seismic data interpreter will analyze it to come up with recommendations on drilling target locations. The seismic data interpretation process is not as simple as point and shoot, however. There are six major subjects that fall under the umbrella of seismic data interpretation. These are: velocity profiling, structure interpretation, amplitude interpretation, well incorporation, attribute analysis, and analysis of amplitude variations with offset (AVO). Though there are many other subfields of work in the interpretation industry, it is essential to be familiar with these six.

Geophysicists have a litany of velocity definitions to deal with. In seismic data processing, there are stacking velocities, migration velocities, and many others. In

seismic data interpretation, there are sonic log velocities, checkshot velocities, and many others. Each type of velocity must be used in a particular circumstance. Stacking velocities are one measure of the acoustic velocity in the medium. Sonic logs use much higher frequencies than seismic surveys, typically thousands of hertz, and are not recorded all the way to the surface. Checkshot surveys approximate seismic acquisition in geometry and frequency and are thus the best velocity for depth conversion. If there are existing wells in the area, the first order of business when beginning interpretation of a new seismic survey is to vertically correlate the well data to the seismic data by creating a velocity function.



A velocity function in this case refers to a map from depth to time along a wellbore. The geophysicist must output a list of (time, depth) points increasing in depth down a well. This time–depth chart is discrete differentiated to come up with a set of velocity points with respect to depth in the well. These velocities are used by seismic data interpretation software to stretch or squeeze the projection of the wellbore on the seismic data to fit the geologic interpretation. When the velocity function for a well is correct the well should project on the seismic with the geologic interpretation in the same vertical position as known geophysical events.

In the simplest case, a geophysicist would create the time–depth chart by lining up geologic interpretation along the wellbore with known seismic events. For example, a geologist would provide the point at which a fault plane crosses the wellbore, so the geophysicist would find the fault plane on the seismic data. Now the depth (from the well data) and the time (from the seismic data) would be known at a single point. This would be done at several points along the wellbore to create the time–depth chart.

A sonic log is created by running an electric tool down the wellbore which creates acoustic energy at one end, and measures that energy at the other end. These tools can be very complicated, measuring different types of acoustic waves or directional acoustic energy, but overall the instrument simply measures the velocity (or slowness) of sound refracted in the borehole wall as the tool is passing through. Because of this dependence on the borehole wall, sonic logs produce questionable results in irregular boreholes as determined from the caliper log. The sonic log along with a formation density log is used to create a synthetic seismogram.

Acoustic impedance is the product of density and p-wave velocity. The reflection coefficient of a vertically incident p-wave is related to the difference in acoustic impedance across the reflecting interface by the following formula

$$RC = \frac{AI_1 - AI_2}{AI_1 + AI_2}, \quad (12.2)$$

where RC is the reflection coefficient,  $AI_1$  is the acoustic impedance of the first medium and  $AI_2$  is the acoustic impedance of the second medium. Reflection coefficients at non-normal incidence will be discussed later in the section on AVO. Because acoustic impedance contrasts are what cause variations in seismic waveforms, a simple synthetic waveform can be created in a clever way: first the density and sonic logs are multiplied to create a pseudo-acoustic impedance log. Next, an acoustic reflection coefficient series is generated

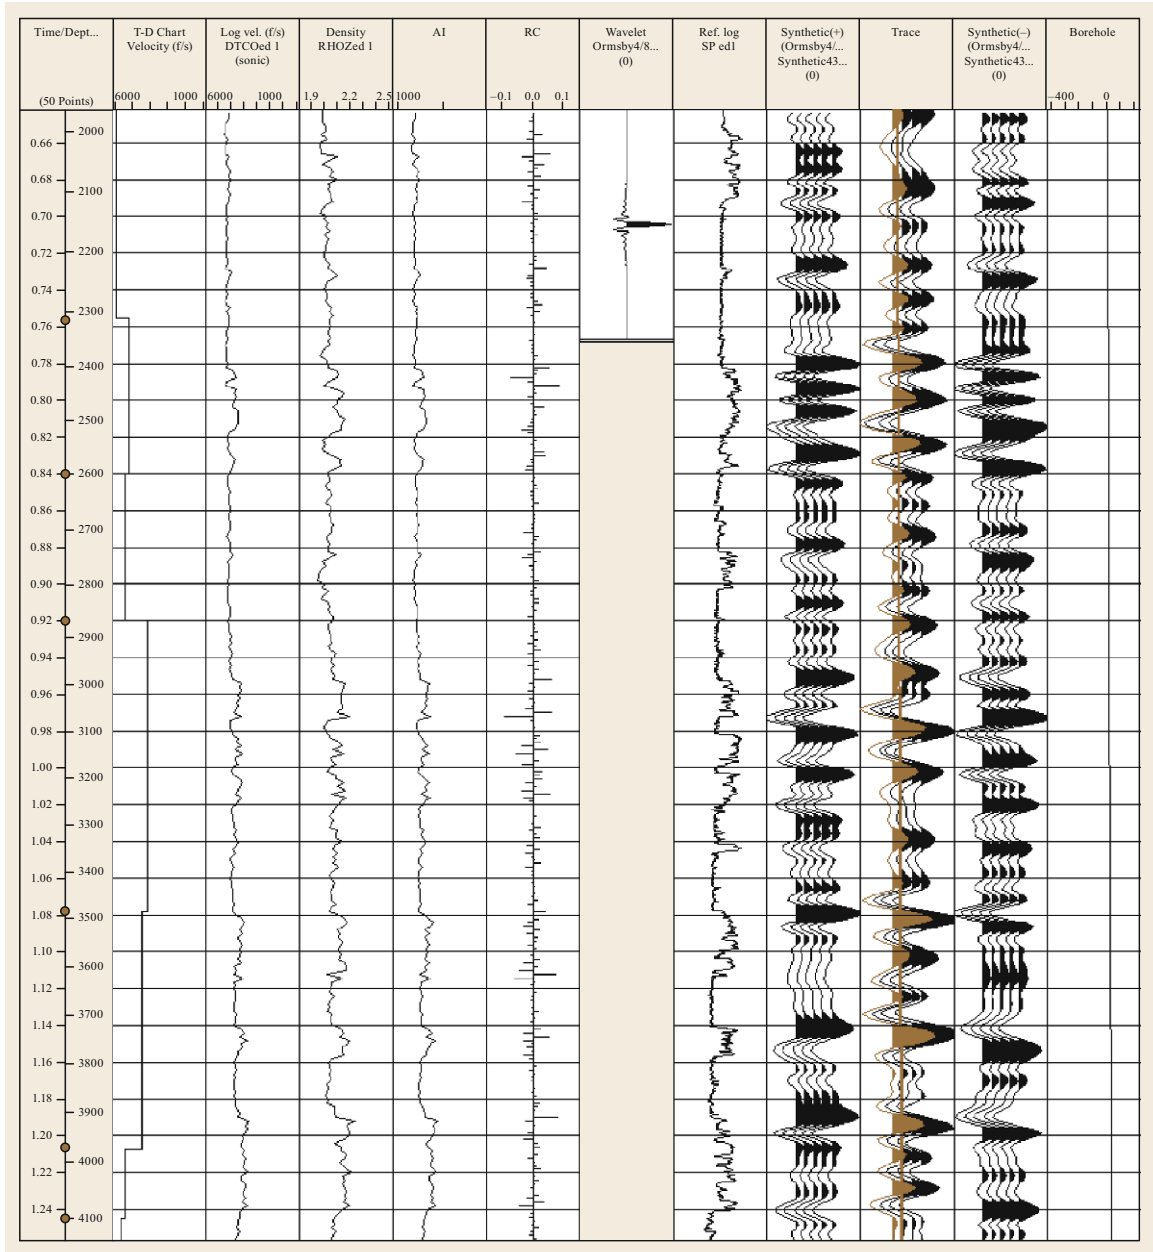
from the acoustic impedance. Then a seismic wavelet is created by either mathematical construct, or by approximation from real seismic data surrounding the wellbore. Finally, this wavelet is convolved with the reflection coefficient series to create the synthetic seismic trace (Fig. 12.3).

Once the synthetic data is created, it can be compared to the real data via the use of the time–depth chart. Recall that velocity is the derivative of the depth function with respect to time. By changing the velocity on individual intervals, the synthetic data can be stretched or squeezed vertically to match the real data. Usually a geophysicist will stretch/squeeze the synthetic data by eye next to the real data to correlate seismic events that look similar. The geophysicist will do this for many wells inside the area of the seismic survey. The benefit to all of this is that now the geologic interpretation will be qualitatively lined up to the seismic data via the projection of the wellbores on the seismic data. The next step in seismic data interpretation is to pick the seismic events that correspond to potential targets.

Picking seismic data is simply the process of selecting events in a seismic data volume by means of a colored marker so that a geophysical model can be built. There is no order in which the geophysicist must pick seismic data volumes, but there are some general guidelines to make orient the interpreter and build some structure into a large data set. When a seismic volume comes in from the processing shop, it can be overwhelmingly large. There are sometimes millions of data traces extending over many square miles of area. It is important to start picking a data set on the largest geophysical features and move toward the small scale.

First, the interpreter should pick major faults. Major faults are usually the easiest thing to see on seismic data, so beginning at one end of the data the interpreter picks a data line for the faults. Moving across the data set the interpreter picks seismic lines at some convenient interval. For a 100 square mile 3-D data set, a good gross fault picking interval may be every thirty lines. A computer program can be used to interpolate between picked lines, creating a full surface extending across the seismic data set which represents the fault scarp itself.

The next thing to pick on the seismic data should be important horizons of interest. A horizon is defined as a continuous, or discontinuous in some cases, seismic event extending in the lateral dimensions representing some geologic feature. If there is a large acoustic impedance contrast between, say, a dense limestone formation bottom and a porous sandstone bed top, the seismic data would have a large trough (or peak, depending on polarity) at the interface. This would be an



**Fig. 12.3** A synthetic seismogram generated by convolution of the Ormsby wavelet shown on *right* with the reflection coefficient series shown on *left*

easy horizon of interest to pick as a starting point to make note of large scale geology–geophysical correlations.

After picking the major faults and horizons the interpreter will begin picking finer details in the data. The subtle structural features, geologic unconformities, and potential drilling target horizons should be picked at this point. In addition to these details, geobodies can

be picked during this step. A geobody is a 3-D seismic object of any construction picked in all three dimensions. Examples of geobodies can include erosional channels, fault blocks, or salt bodies. Software can be used to eliminate geobodies from a seismic volume, as to reconstitute a pseudo-stratigraphic section in time. Similar procedures to reconstruct stratigraphic slices from seismic data in time use flattening algorithms on

the horizons, which creates a (somewhat) geologically accurate image.

It is important when interpreting seismic data to recognize the difference between a seismic image and the subsurface reality. Seismic data is sometimes misunderstood as a picture of the subsurface. This could not be further from the truth. As was shown in the previous two chapters, seismic data not only undergoes many processes with apparent structure altering characteristics, but also acquired under noisy conditions miles away from target zones. Some seismic interpreters bet millions of dollars on seismic data alone. The reality is that the subsurface geologic interpretation is much more closely related to subsurface geologic reality. It is because of this fact that good seismic interpreters use geologic interpretations as a limiting case for the seismic interpretation. When interpreting, seismic always use real geologic data to test seismic-based hypotheses.

In any commercial software for seismic data interpretation, and in any other good interpretation software, there is the ability to input well data. These well data can be used in many ways. The above mentioned time–depth chart creation is one of the things seismic interpreters do with well data. Others include projecting geologic formation top interpretations on the seismic data to verify that the events being picked are close to geologic reality. The same thing can be done with well logs; an interpreter may use gamma ray, spontaneous potential, resistivity, neutron density, or many other well logs for geologic verification of geophysical picks. Well logs in some cases can also be used for qualitative comparison of areal coverage in a seismic data set.

The traces in a seismic data survey surrounding a wellbore are acoustic responses to the rocks surrounding that well. It follows then that seismic data anywhere in the survey bearing resemblance to the data around the well may be a response to similar rock characteristics. In this way, seismic data is used to create qualitative comparisons between well locations and prospective target locations. Well logs can be semi-quantitatively incorporated into this scheme by statistically analyzing the seismic response around many wells with similar log suites. Well logs can thereby be used to determine what causes seismic amplitude response.

In many introductions to geophysics courses, students are instructed to drill locations with bright spot seismic response. Bright spots, or isolated high amplitude anomalies within a low amplitude event, can be created due to the change in acoustic impedance of a formation because the pore spaces are filled with hydrocarbons rather than water. The bright spot phenomenon is certainly a good method to locate potential targets, but they are not necessarily always direct hydrocarbon indicators.

The amplitude of seismic data at a location is a function of a huge number of factors. Part of the interpretive geophysicist's job is to evaluate within each seismic survey what drives the amplitude response. In some areas, seismic amplitude may be very closely related to the presence of oil and gas. In these cases, the drilling company just puts the wells where the seismic data lights up!

Seismic amplitudes usually behave in a more complicated manner. In addition to hydrocarbons, seismic amplitudes can also respond to gross-to-net sand ratios, bed thickness, fracture orientation, etc. There are thousands of things that can cause seismic amplitudes to be artificially inflated or deflated in a location. When interpreting seismic amplitudes, one should always take a look at the processing sequence to be sure those amplitudes have been properly preserved. Statistically incorporating well logs into the seismic data can vastly improve the geophysicists' understanding of what is the greatest geologic contributor to seismic amplitude response. It is important to note that amplitude response may not have the same driving mechanisms throughout the entire seismic survey or from one target depth to another. Because of this, interpretation must be carried out separately from zone to zone and across the entire area.

Another type of seismic information that can help geophysicists understand potential drilling targets is called seismic attributes. A seismic attribute is anything represented at a point location in a seismic survey by a pair of spatial values, a time value, and the value of the attribute ( $x$ ,  $y$ , time, attribute), for example, amplitude is a seismic attribute. Amplitude is one characteristic of a seismic waveform. Another characteristic, or attribute, of a seismic waveform would be velocity. In the seismic data processing chapter stacking velocities were defined. As was shown, every CMP gather location is analyzed and assigned a velocity for every time value in the stacked trace. Therefore, once this stacking procedure is complete the seismic data processor has a value of velocity at every point in the seismic survey. If desired the processor could output a seismic data volume of space, time, and velocity ( $x$ ,  $y$ , time, velocity), this volume would be called the velocity attribute of the seismic survey.

As seismic attributes are derived from any mathematical operation on seismic data, there are an infinite number of attributes. Fortunately, due to the literature on the subject there are a finite number of categories into which these attributes fall. The major ones are geometrical, instantaneous, wavelet, and exotic attributes. The usefulness of seismic attributes varies with the geologic setting and the information, such as well data, available. The most important aspect to consider when deciding which attributes to use is that all attributes

come from recorded data. Performing mathematical operations on any data (seismic or otherwise) can provide a different view of that data, but once the data has been recorded no new information is available to that system. Seismic attributes do not create new information; they only bring out features of the data which are not accessible by the human eye in amplitude data.

The geometrical subset of seismic attributes includes several very useful attributes. They include curvature, dip characterizers, and coherence. Curvature attributes describe the shape of individual seismic events. Seismic interpreters can use attributes such as maximum curvature, Gaussian curvature, and most positive curvature to help tease out subtle structural details in seismic data or continuity in the bedding planes associated with their target prospects. The suite of dip characteristics includes dip azimuth, dip variance, and dip of maximum similarity. These attributes in areal cross section can help to delineate reservoirs or resolve geologically complex settings. Coherence, or similarity as it is also called, is an extremely helpful attribute which creates an image of the subsurface analogous to an MRI image. The discontinuities between seismic traces are assigned a low value of coherence, so continuous zones look blank. Coherence can be used in areal cross section or vertical cross section to delineate small faults or lineaments in the subsurface, for example, fluvial channels.

The instantaneous suite of seismic attributes is created using differential type operators which take only the surrounding samples into calculation, hence the name instantaneous rather than average. The major players in this category include instantaneous frequency, instantaneous phase, and relative acoustic impedance. The frequency and phase attributes are used to assist in quantifying geology–geophysical relationships. They can also be used to delineate the lateral extent of beds in some cases. Relative acoustic impedance is often touted as an attribute that shifts the seismic data from reflection coefficient amplitude data into a true geologic image. This is not quite the case, because there has not been any new information put into the system, but either way relative acoustic impedance gives the interpreter a phase shift and some scaling on the data which in some cases can, again, help delineate bedding.

In many cases the attributes of the wavelet attribute suite are used as intermediate products to create a final attribute based on several components. However, wavelet envelope and wavelet bandwidth can be used to assist in interpretation. These attributes can be viewed in vertical cross section to give the interpreter a better feel for the gross geologic structure under certain circumstances.

The term *exotic* attribute is used here to describe any attribute that does not appear in widespread use across the literature. As previously mentioned, there is an infinite number of attributes, so the exotic attribute category can include infinitely many attributes. One example that is used with great frequency is called sweetness. Sweetness is wavelet envelope over the square root of instantaneous frequency. Sweetness can be used to pick out areas of high net to gross sand ratio, which is important in in-field development projects and near-field exploration. If an interpreter recognizes which attributes work well in an area, he/she may be able to work together with a processor to create a custom exotic attribute. Just because this custom attribute works well in this one area does not mean it will work in any other setting.

Most beneficially, many seismic attributes will be used in concert with each other to create the best interpretation of the seismic data. One way to do this is by using a multivariate statistical approach. The seismic interpreter extracts traces of several seismic attribute volumes, including amplitude if valuable. The traces should be around wells that produce and wells that do not produce, as to build up a good sample set. Each of these trace data will be culled for samples around zones of interest. The interpreter then builds up a strong multivariate equation based on the attributes which seem to have clear response to the desired geological characteristic and builds a model zone for each real zone of interest. This can be done as a map-based (2-D) or volume based (3-D) process. When complete, the model zone should display a statistical likelihood of success for drilling as a function of time and/or space.

The last major tool in the interpretive geophysicist's toolbox is the analysis of amplitude with respect to offset. Amplitude versus offset (AVO) analysis can prove to be a powerful process for finding hydrocarbon rich target zones, and also for characterizing lithology in undeveloped regions. To perform AVO analysis, the interpretive geophysicist must have access to the processed CMP gathers from the processing geophysicist. The gathers may be processed to any point that the interpreter prefers, but should represent the final seismic section that is represented. All industry standard AVO analysis is done on NMO or, better yet, DMO corrected gathers, so that the interpreter can pick and observe flat events rather than hyperbolic events. Care must be taken at each pre-processing and processing step to preserve relative amplitudes. Amplitudes on land data are often affected by the near-surface, so surface-consistent amplitude balancing may be a necessary step. Pre-stack migration or DMO improves AVO response by moving energy between traces.

The interpreter will pick the gathers with a horizon of interest across an area of interest. This is usually performed with an autopicking software initialized with the normally incident traces in each gather. The picked event is analyzed by a simple algorithm on a per CMP basis and could output an AVO class for each. An AVO class (from class 1 to class 5) is an industry standard definition of the response characteristics in the gather. The Zoeppritz equations predict the amplitude of waves propagating away from an encounter with an interface with respect to angles from a normal. There are four Zoeppritz equations to give the amplitude of the reflected p-wave, the reflected s-wave, the transmitted p-wave, and the transmitted s-wave. All four equations use both p-wave and s-wave velocities, so the reflection amplitude behavior at non-normal

incidence contains information about mechanical properties such as compressibility of the formation that is not contained in the zero-offset trace. The presence of shallower interfaces means the transmission effects at each shallower interface must be modeled as well. There are a number of commercial AVO modeling packages available. The AVO characteristics can be displayed in a crossplot with the hopes of grouping outlying data. The outliers can be culled and their CMP locations displayed on a map. By way of example, if the interpreter knows that all CMP locations exhibiting an anomalously high Shuey three-term gradient to Shuey three-term curvature response are oil bearing, then this zone would be culled for this data response, and a possible location would present itself in areal cross section.

## 12.4 Summary

The seismic data interpretation procedures outlined above are used as a starting place for building a model. It is the interpreter's job to use the seismic data for verification of an Earth model. Sometimes the model only exists in the interpreter's head, but more explicit models are recommended. Seismic modeling software can create 3-D models from seismic interpretation. The interpreter will input his/her fault surfaces and horizon surfaces into the modeling software which will create 3-D volumes from this information. The 3-D models can be used to model reservoir drainage, to plan wells, or even to steer wells in real time as they are drilled.

The key to good modeling is iteration. The first model created will never be spot on, which is why the interpreter needs to rework the model as new information arrives. If for example a well is drilled on the initial model, the interpreter should have a new set of well logs to help quantify the seismic data. The iteration process continues from the initial model through the completion of the entire drilling program, and into any other projects of similar setting. Geophysicists also must iterate new geologic and engineering data into their models. Using all of the information at hand is crucial to continued success.

The initial and updated data must drive the model. It is a waste of money to come up with a hypothesis and then use the seismic data to prove that hypothesis. The correct order should be to look at the data and then come up with a hypothesis; if any new data comes in disproving that hypothesis, the hypothesis should be thrown out. A common mistake when building implicit Earth models that only exist in the head of the interpreter is following these incorrect hypotheses all the way to drilling a dry hole. In this way an explicit Earth model updated by the geophysicist, geologist, and engineer is the most accurate and successful way to complete a drilling program. The entire seismic data workflow is a procedure in risk mitigation. Any company can throw a dart at a map and drill that location, but the chance of success is low. Seismic data acquisition, processing, and interpretation reduce the risk of drilling to some point dependent on data and interpretation quality. If this reduced risk level is below whatever arbitrary threshold risk level the drilling company has decided on, then seismic exploration has done its job.

**Acknowledgments.** The author would like to acknowledge valuable comments from Dr. Greg Kroft and Professor Kent Pederson of the University of Calgary.

## References

- 12.1 C.C. Ober: *Marmousi Model and Image* (Sandia National Laboratories, Albuquerque 1998), <http://www.cs.sandia.gov/~ccooper/seismic/marmousi.html>
- 12.2 J. Wang, M. Sacchi: Noise reduction by structure- and-amplitude-preserving multi-channel deconvolution, *CSEG Recorder* **34**(2), 24–27 (2009), Canadian Society of Exploration Geophysicists