PRACTICAL PETROPHYSICS

Martin Kennedy

Developments in Petroleum Science

Volume 62

Practical Petrophysics

Developments in Petroleum Science

Volume 62

Practical Petrophysics

Martin Kennedy
MSK Scientific Consulting, pty ltd. Perth, Australia

For information on all Elsevier publications
visit our website at http://store.elsevier.com/

Contents

Series Editor's Preface

This is the third book in the *Developments in Petroleum Science* series since it incorporated *The Handbook of Petroleum Exploration and Production* in 2013. After books on geophysics and stratigraphic reservoir characterization, we now look at the equally important field of petrophysics.

Petrophysics is described as the study of the physical properties of rocks, their pore systems and the fluids they contain. As such it plays a key role in the geosciences and reservoir engineering and is a cornerstone to petroleum exploration and production. In J.H. Schon's 2011 workbook in the *Handbook of Petroleum Exploration and Production* entitled *Physical Properties of Rocks*, he discusses and defines the fundamental parameters we can measure by petrophysics including fluid types and volume, porosity, and lithological rock types as well as some of the properties needed for reservoir characterization and simulation.

Martin Kennedy builds on these fundamentals in this book and in his words 'show how to achieve a balance between the rigorous principles that ultimately determine the petrophysical properties and how our measuring instruments respond to them on the one hand and what is realistically achievable with limited time and resources on the other'. In other words he is moving away from purely idealized or theoretical petrophysics to the more complicated and multi-disciplinary nature of the real petrophysical world.

The idealized or theoretical vision of petrophysics has been established over the last 90+ years, and is essential to make sense of our rock, fluid and log measurements. However, as our application of petrophysics has grown ever more sophisticated so has the risk that we employ techniques or equipment that exceed their theoretical limits. Here Martin Kennedy has emphasized how important it is for all practicing geoscientists or petroleum engineers to understand how a particular equation or technique is derived and what its limits are. He describes these techniques, equipment and limitations in detail and provides a book that will as a result be essential reading for those active in petroleum exploration and production.

John Cubitt
Holt, Wales

Page left intentionally blank

Preface

The number of specialist textbooks dealing with Petrophysics or Log Analysis is relatively small and some of them are getting quite 'long in the tooth' now. This book was written to fulfil three objectives:

1. To be up to date both in terms of general measurements and techniques.
2. To present the basic principles of petrophysics in a straightforward and – I hope – readable form.
3. To show how these principles can be applied in a pragmatic way to estimate petrophysical properties.

It is not intended to be a comprehensive description of every equation and algorithm that is available to petrophysicists. Neither does it attempt to describe every logging tool and core analysis measurement currently available. The latter is a hopelessly ambitious task and to my certain knowledge half a dozen new logging tools have appeared on the market between finishing the text and writing this preface. More importantly, it is not intended to be a 'How to' manual. There is undoubtedly a place for such books but blindly following recipes without some understanding of the underlying principles is asking for trouble in any technical discipline (including cooking). It is particularly dangerous in petrophysics where most of the equations are either purely empirical or are based on a number of approximations and assumptions. This book is intended to provide that essential background.

In terms of column inches much of this book is devoted to log analysis. This is quite deliberate and is a reflection of the relative amounts of log and core data that are typically available in the field. No matter how highly an individual or organization values core the fact of the matter is that there will inevitably be large gaps in the core record that only logs can fill. The best interpretations use the two types of data to complement each other and although not often explicitly stated, I hope the book leaves the reader with that impression.

The emphasis on principles means that some of the day-to-day tasks of a petrophysicist are given scant coverage in this book. Topics which I have at best mentioned in passing and generally ignored include getting logging tools into places where they do not want to go and the limitations created by high temperature and pressure, corrosive fluids, general HSE and money. These are all important but at the end of the day they have no influence on how well Archie's equation works and so have no place in this book.

Martin Kennedy

Page left intentionally blank

Chapter 1

Introduction

Chapter Outline

1.1 WHAT IS PETROPHYSICS?

Petrophysics is the study of the physical properties of rocks. As a pure science its objective would probably be to explain why rocks have the properties they do. In particular how the relative amounts and arrangements of the minerals that comprise them determine their physical properties. In practice, most of the time we are concerned with the reverse problem of using physical properties to try and find out what the rock is made of. This is valuable information for anyone who works with rocks whether as a resource, a substrate or a storage medium. But, as will be seen below, petrophysics has its origins in the oil industry and is still most widely used for describing the rocks that make up hydrocarbon traps. For this reason most of the tools and techniques that are described in this book were originally developed to deal with porous, sedimentary rocks in the sub-surface. In particular the problem of determining what the rock is made of often reduces to finding how much of the rock is fluid, how much of that fluid is water and how that fluid is distributed (as that will give some indication of how easily it can be extracted).

More succinctly petrophysics in the oil industry is used to find the following:

1. Porosity – How much fluid can the rock store?
2. Saturation – How much of it is water?
3. Permeability – How quickly can it be extracted?

These are often referred to as 'petrophysical properties' or even just 'properties'. The tools and techniques that were developed to estimate them can often be used to find other information of practical importance, for example identifying special minerals or modelling the seismic response of a sand/shale interface

(we will look at this later in the book). Moreover in order to estimate these three properties we often have to go through intermediate steps so that a full petrophysical analysis may well end up producing a lot more information.

Since this book is ultimately concerned with the properties of rocks we should explain what we mean by a 'rock' and also how big it is. For our purposes rocks are physical mixtures of minerals. Minerals are for the most part chemically pure substances that may be solid, liquid or gas (so in this book at least water, oil and gas are considered minerals). For convenience we will also include mixtures of similar compounds as minerals. An obvious example is crude oil, which is invariably a mixture of hydrocarbon molecules as well as some more complicated organic compounds. Examples of solid mixtures are some of the clay minerals, which can have a range of compositions and a single grain may show a variation in composition from one side to the other.

The size of the rocks we are interested in is largely determined by our measurements. In the laboratory, samples may be minute, in fact some techniques can be applied to single mineral grains. But in this book we will frequently deal with borehole logging measurements, which typically cover volumes from tens of cubic centimetres to several cubic metres. Even small core plugs have volumes of several cubic centimetres. So to put it simply, the volumes we deal with vary in size from hand specimens to boulders.

1.2   EARLY HISTORY

No doubt scientists have been measuring and exploiting certain physical properties of rocks for centuries but most petrophysicists would date their profession to the 1940s. Fittingly the noun 'Petrophysics' was coined by G.E. ('Gus') Archie in the late 1940s to satisfy what he felt was the need for a word to describe the study of the physics of rocks. Even if someone else had invented the name, Archie would almost certainly still be regarded as the founder of the profession. In 1941 he developed the empirical equation, that bears his name, which relates the electrical resistivity of a porous rock (R_0) to its porosity (Ø) and the resistivity of the fluid – invariably salt water – contained within its pores (R_w). In general

$$R_0 = f(R_w, \O) \qquad (1.1)$$

This is a classic case of petrophysics in action. The equation describes how the resistivity of the rock depends on the relative amount of one of the minerals in the rock (water); and as we will see later, how that water is distributed within the rock. Historically, it is considered to be the first attempt to explain why a physical property has the value it does.

Of course being able to predict how resistivity depends on porosity or vice versa, might be interesting but if it was limited to laboratory measurements on core plugs it would have few practical applications. Fortunately, resistivity had been measured in boreholes since 1929, when the Schlumberger brothers ran an experimental tool in a well in Alsace. This was the first wireline log (in this book

we will henceforth simply refer to wireline logs as 'logs'). The technique rapidly caught-on and by the time Archie published his results, resistivity logs were routinely run in many parts of the world. Their principle application was however, correlation and qualitative interpretation such as identifying sands and sometimes distinguishing water and oil in the pore space. Archie's work allowed the logs to be used to estimate porosity along the well bore. Log analysis is now the standard way to determine the petrophysical properties in the sub-surface.

Almost from the start, logs were an oil industry tool and it is hardly surprising that Archie too came from that industry (specifically Shell Oil). To this day the major developments in petrophysics hardware and interpretation tend to be driven by the needs of the hydrocarbon industry. Nevertheless, it can, and is applied to all industries that deal with rocks.

We will look at Archie's equation in a bit more detail in subsequent sections and a lot more detail in a later chapter. Before doing that it is only fair to point out that Archie himself had much wider interests than electrical resistivity. He studied almost any rock property that could be expressed numerically and in the 1950 AAPG paper in which he introduced the word 'petrophysics' he was already describing applications of porosity–permeability cross-plots, capillary pressure curves, the SP log, neutron logging and of course resistivity. Significantly he also showed how these properties depend on the geometry of the pore system. In short he did not leave much for his successors to work on.

1.3   PETROPHYSICAL DATA

Almost all the petrophysical data discussed in this book comes from wells, this imposes some important constraints on the accuracy of our estimates. There are two, fundamentally different, sources of data:

1. Instrumental methods that measure physical properties.
2. Actual samples of rocks, which can be analysed in a laboratory.

(For completeness we should add the various types of well test to this short list but we will defer any further discussion of these until much later in the book.)

The former obviously refers to the various types of geophysical log (which we will simply call 'logs'). These provide a continuous record of one or more physical properties along the path of the well. Log analysis converts physical properties to petrophysical properties and much of this book concerns it. This is entirely appropriate: log analysis is not a synonym for petrophysics but it is an indispensable part of it.

'Samples' include cuttings and various types of core. Cores are normally quite localised, either in relatively short intervals, in the case of a whole core, or widely spaced depth points, in the case of sidewall cores. Cuttings do give a continuous record, but the drilling process always results in a certain amount of mixing, possible loss of some minerals and sometimes they are so finely ground that it is impossible to tell their original lithology. Even so, all these different

types of information should complement each other, and if properly integrated their individual shortcomings can be overcome to an extent.

Well-bores are difficult places to make measurements and so there are relatively few instrumental techniques we can adapt for that environment. Of these, few can read more than a few centimetres into the formation. Unfortunately, drilling inevitably alters the formation near to the well bore so even when the logging tool is working perfectly, it will make an accurate measurement of rock that has been changed in some way.

On the other hand, core material can be studied 'at leisure' using almost any technique one desires. Unfortunately, whole cores are expensive and sometimes nearly impossible to acquire. Sidewall cores are a cheaper alternative but there is a limit to how many can be taken from one well and so they are generally quite widely spaced. Also, depending on the type of tool that was used to obtain them, they may not be suitable for all types of analysis. In any case, regardless of what type of core was taken, the rock goes through some drastic changes between coring, being brought to surface and then being cleaned and prepared for analysis. Cuttings give the greatest coverage for any sample type and they are always present (although for the top-holes of some offshore wells they never get beyond the sea bed). On the other hand they also suffer the greatest alteration on their journey to the surface.

Even if we can obtain a complete set of well logs, cores and cuttings we can only really be certain that we have characterised the reservoir in the near well bore region. Because most logs can only read at most a few metres we have no direct knowledge of what happens beyond. The net effect of this is that petrophysics can often provide a very accurate and precise description of the sub-surface but only at a few points across the reservoir (i.e. the wells). The greatest uncertainty is often associated with what is going on between the wells.

1.4 QUANTITATIVE DESCRIPTION OF MIXTURES

As noted earlier a lot of applied petrophysics involves finding the relative proportions of the minerals – including water and hydrocarbons – that make up a rock. When we describe a mixture we have a choice in how to express the relative amounts of each of the components. For describing rocks the simplest choices are:

1. by volume fraction.
2. by mass fraction.

By convention, but also for convenience, in petrophysics and log analysis the proportions are invariably expressed as volume fractions. Porosity, for example is the volume fraction of fluids in the rock and 'shale volume' is self-explanatory. Many of the laboratory techniques that are applied to cores and cuttings, however give the results as mass fractions. It is obviously important to know which system is being used and since we often wish to integrate

data from the two sources we need to know how to convert from one system to the other. To do this requires knowledge of the density of each component. This is generally easier said than done, but for now we will assume we do know the densities of all the minerals making up the sample. The calculation is best illustrated by an example.

Consider the analysis of a sandstone sample (this is based on a real sample but it has been simplified to four minerals by excluding some clay and mica that made up about 5% of the total). The numbers give the mass percentage of the solid minerals. In other words any porosity is excluded. The densities of the four minerals are written in the row below the mass fractions and it can be seen that the two iron-containing minerals on the right, are significantly denser than the silicates on the left.

Mineral	Quartz	Kaolinite	Siderite	Pyrite	Sum
Per cent by mass	31.5	15.3	48.6	4.6	100
Density (g/cm^3)	2.65	2.64	3.96	5.00	

To convert to equivalent volumes divide the mass fraction by the density.

Mineral	Quartz	Kaolinite	Siderite	Pyrite	Sum
Mass frac./density	11.9	5.8	12.3	0.9	30.9

Finally to convert to volume fractions, divide by 30.9.

Mineral	Quartz	Kaolinite	Siderite	Pyrite	Sum
Per cent by volume	38.5	18.8	39.8	2.9	100

The differences in this case are large enough to be significant and if we wished to integrate a log analysis with the sample analysis we would be well advised to go through these steps.

Sometimes it may be necessary to convert volume fractions to mass fractions. For example, one of the most widely used quantities to describe a source rock is the total organic carbon (TOC) content. This is the ratio of the mass of carbon present in organic molecules, to the total mass of the rock. It is so familiar and so widely used that it is futile to protest that it would be better to use the volume fraction of organic matter (Φ_o). So, although the latter can be estimated using log analysis, sooner or later it needs to be converted to TOC. Again density is the key to the conversion. If the density of the source rock is ρ and the average density of the organic matter is ρ_o, then the mass fraction of the organic component is:

$$\text{Mass fraction organic matter} = \frac{\rho_o . \Phi_o}{\rho}$$

(Try deriving this for yourself.) So if a source rock with a density of 2.2 g/cm^3 contains 10% organic matter by volume with a density of 1 g/cm^3, the mass fraction of organic matter is 4.5%. The TOC is actually less than this because it

refers to the mass fraction of carbon alone. To find this we need to know the ratio of the mass of carbon to the total mass of the organic matter. This depends on precisely what the organic matter is made of but varies from 0.75 for methane, up to 0.95 for hydrocarbons and as low as 0.4 for carbohydrates (which do not survive long in the sub-surface). A value of 0.8 is often assumed giving a TOC of 3.6% in the above example.

To conclude this section we will look at some of the ways to describe concentrations of specific chemical compounds and elements. It generally does not matter whether these are expressed as mass, volume or some other fraction but it is important that the appropriate system is used for the particular application. For example, to apply Archie's equation we need to know the resistivity of the formation water. This depends on the concentrations of the salts dissolved in the water and temperature. For most formation waters the majority of the salt is sodium chloride and its concentration is specifically called 'salinity'. In oilfield applications salinity is most often expressed as mass of sodium chloride per unit mass of solution. It is normally expressed in parts per million (ppm). Present day sea water, for example has a salinity of 35,000 ppm (meaning 35 g of NaCl per kg of solution it is sometimes written 35 kppm). Providing we have a chart or formula to find the resistivity from salinity there is no need to convert to volume fractions but it is important to confirm the concentration really has been given as a mass fraction. To be completely sure some reports quote the salinity and then write 'w/w' in brackets (spoken as 'weight–weight'). Other ways of expressing salinity are mass per unit volume ('w/v') and molarity (number of moles per unit volume of solution).

Concentrations of exotic elements, that are only present at trace levels, are normally expressed in parts per million but it is generally difficult or impossible to find exactly what these ratios refer to. In petrophysics, this is not a serious problem because it is often the trends in concentration with depth that are of more interest than the absolute numbers. Nevertheless, when comparing data from different sources we ought to be confident that we are comparing like with like, but all too often this is not possible.

The commonest system refers to numbers of atoms, so, for example a uranium concentration of 5 ppm means that five out of every million atoms in the sample are uranium. Another system commonly used by geochemists quotes the number of atoms of the element of interest for every million silicon atoms. This obviously acknowledges the high abundance of silicon in the Earth's crust and mantle, but is of limited use in carbonates, for example. In this book the two most important examples of trace elements are uranium and thorium, as they are responsible for much of the background radioactivity in the Earth's crust. We will assume that concentration refers to the total numbers of atoms.

Finally, mention needs to be given to hydrogen because some logging measurements are designed to be particularly sensitive to it (neutron porosity and NMR). These are often discussed in terms of a property known as the 'hydrogen

index' which is a measure of the number of hydrogen atoms per unit volume (sometimes written HI). By definition, pure water at 75°F (23.9°C) has a HI of one, this can easily be shown to be equivalent to 0.11 mole hydrogen atoms per cubic centimetre (or 6.6×10^{25} H atoms/litre). HI is also an important characteristic of source rocks but unfortunately for these applications it is defined in a different way, although it is still written HI. Specifically, when applied to source rocks the HI is the mass of hydrocarbon per unit mass of organic carbon. A good source rock will typically have a HI of 0.5 (i.e. 0.5 g hydrocarbon per gram of organic matter). Unfortunately, HI is by no means the only property that petrophysicists have their own definition for, some more examples are given in Section 1.9. In this book HI will always have the first meaning, that is hydrogen atoms per unit volume relative to water.

1.5 THE PRACTICE OF PETROPHYSICS AND PETROPHYSICS IN PRACTICE

Petrophysics is a quantitative discipline and as noted earlier Archie was interested in almost any property of a rock that could be measured. It exploits theoretical calculations made on simplified models of rocks, numerical modelling of more realistic models, measurements on very well characterised physical models of porous solids and measurements on samples from well characterised real rocks that are chosen for their simplicity and uniformity. Much of this fundamental work was and still is carried out in academia and industrial research laboratories.

Ultimately, however this book is concerned with applying this knowledge to rocks of economic importance. We therefore have to deal with what we are given by nature and this mostly means rocks that are complicated mixtures with properties that may show a lot of variation over a short distance. A theoretical model to explain how their properties vary with composition is likely to be impossibly difficult to develop and not particularly useful. From the earliest days progress has been made by combining the general results of research with specific measurements made on the rocks of interest.

There are some practical issues that arise from time to time:

1. The theoretical models are often overly simplistic and get pushed too far.
2. The experimental data inevitably only applies to a limited range of compositions and conditions and we have no knowledge of what happens outside these limits.
3. Equations that are really only empirical fits to a specific set of measurements, acquire the status of a law of nature.

It is important to realise that most equations in petrophysics, indeed most of petroleum engineering, are actually empirical. Often they are only intended to work in a limited range of circumstances or they rely on certain assumptions. Problems typically occur when they get applied outside these limits.

1.5.1 The Archie Equation: A Case Study

The Archie equation epitomises the approach of combining models and experimental data and we will use it here to illustrate how petrophysics is applied to real-world problems and some of the pitfalls and misunderstandings that can occur.

The general equation given as Eq. 1.1 can be written more specifically:

$$R_0 = a.R_{\mathrm{w}} / \emptyset^m \qquad (1.2)$$

Where 'a' and 'm' are constants that are derived by simply fitting a curve to measurements of rock resistivity against porosity. The equation quantifies what you might intuitively expect.

1. The resistivity of the rock is proportional to the resistivity of the salt water – 'brine' – in its pores.
2. All other things being equal, as porosity increases, resistivity decreases (note that porosity has to be expressed as a fraction in Eq. 1.2).

The first point tacitly acknowledges that the only part of the rock capable of conducting electricity is the brine. So, if a particular brine results in the rock having a resistivity R_0 say, then replacing it with a different brine that has double the resistivity of the original, causes the rock to have a resistivity of $2R_0$. Archie's experiments confirmed this.

The second point basically says as you increase the porosity, you are putting more brine into the rock and it will conduct better. In other words its resistivity will drop. Again Archie's experiments confirmed this. What is more difficult to predict is precisely how the resistivity will fall with increasing porosity. In fact, this can only be done analytically for very simple systems such as porosity consisting of parallel plane-sided fractures. But at this point experiment takes over and Archie measured the resistivities of hundreds of core plugs saturated with brine of known resistivity to find an equation linking resistivity and porosity.

This combination of a descriptive explanation of the physics combined with empirical observation is still the way most progress is made in petrophysics. The former can give some confidence in the empirical equations and they in turn allow the physics to be applied.

In the case of Archie's equation there is one porosity where we can be absolutely sure of the resistivity without making a measurement or a calculation: at 100% porosity – pure brine – the resistivity must be R_{w}.

But now consider the real – and rather unexceptional – data that is shown in Fig. 1.1. This consists of porosities and resistivities measured on approximately 20 core plugs that have been saturated with brine with a resistivity of $1.00\,\Omega$m. (This data is the same type as generated by Archie, although, does not include any of his original measurements).

The solid light grey line in Fig. 1.1, labelled as the 'Archie fit', is the best fit to the data using Eq. 1.2. It actually passes through a resistivity of $0.56\,\Omega$m at 100% porosity. In terms of Eq. 1.2 this simply means the constant 'a' takes
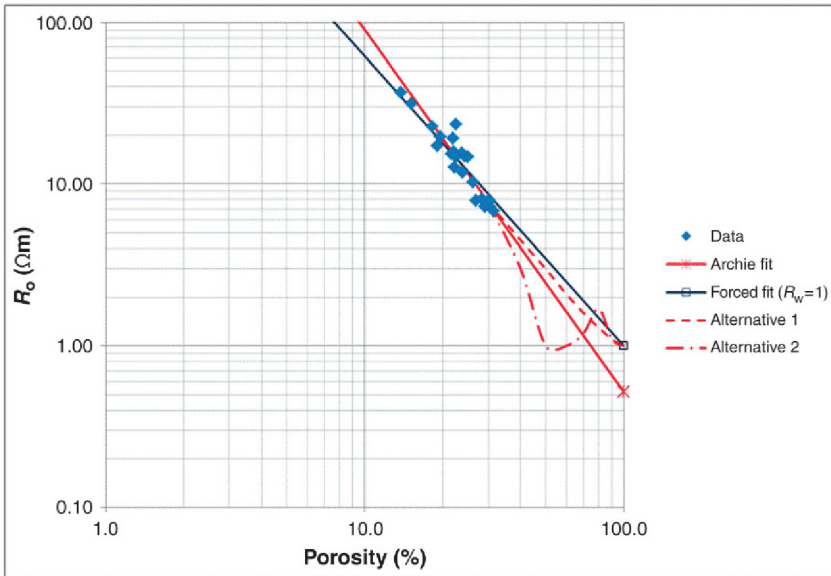
FIGURE 1.1   Some porosity and resistivity data measured on Cretaceous sandstone core plugs from the North West Shelf of Australia. See the text for an explanation of the various curves.

the value 0.56 but it is often argued that a value of anything other than unity is non-sensical because we must measure $1.00\,\Omega$m at that porosity.

   This misses the point that Archie's equation was developed using rocks with a limited range of porosity. Most of Archie's published data lay in the range from 10% to 40% and the data shown in Fig. 1.1 only spans 13–32%. The equation is not intended for use outside this range and although it is sometimes extrapolated to a 100% porosity to find R_w, the extrapolation may well give a value that is higher or, as in this case, lower than, the true value. We have no idea how resistivity varies at higher porosities, most likely it actually follows a curve like the dashed light grey curve labelled 'alternative 1' but, for all we know, it could follow the more complicated behaviour of alternative 2' (light grey dash-dot). For the specific case of the Cretaceous Sandstones from NW Australia it is academic because porosities never exceed 35%.

   We could and often do avoid the contradiction by forcing the line to go through a resistivity of $1.00\,\Omega$m at 100% porosity (the solid dark grey line). This is equivalent to forcing 'a' to unity or, if you prefer, fitting the simpler equation:

$$R_0 = R_w\,/\,\text{\O}^m \tag{1.3}$$

   This may lead to a quieter life and certainly involves one less constant, which was a major advantage in the days when Archie's equation had to be applied by hand. But these advantages come at the price of a poorer fit to the data, we will consider whether that price is worth paying later.

1.6 THE PETROPHYSICAL MODEL

Sandstone will typically comprise at least 15 different minerals, but if we are using log analysis to solve the problem we are often limited to four or five components to describe the rock. This is because we rarely have more than a handful of truly independent log measurements available. In effect, to accurately describe most sandstones we have to solve for more unknowns than we have equations. Even if we did have a sufficient number of logs it is unlikely we could solve for every mineral because the logs have limited accuracy and some minerals have quite similar log responses. In any case, some of the minerals will only be present at trace levels.

To make matters worse the relative amounts of the different minerals probably vary from place to place, even in a sand body that looks superficially uniform. Fortunately, more often than not we do not need to know the relative amounts of every mineral that makes up the rock. So, to make progress, the rock is simplified to four or five key components. Two of these will almost certainly be the volume fractions of water and hydrocarbon (which together give the porosity). The solid minerals are then grouped together as, for example 'matrix' and 'shale'.

The commonest models consist of four components: shale, matrix, water and hydrocarbon. The two fluids are self-explanatory (although we will define them more carefully later in the book). The 'matrix' could probably be better described as the 'reservoir lithology' but, in any case, it has a different meaning to the one typically used by petrographers. 'Shale' may refer to an argillaceous lithology such as claystone or it may actually be referring to the volume fraction of clay minerals present or some intermediate definition. Sometimes, the shale and/or the matrix has to be further sub-divided. For example, if a dense mineral is present in varying amounts we may wish to try and resolve it because it can have quite a significant effect on the interpretation.

The process of building a petrophysical model is illustrated in Fig. 1.2. The real rock, in this case a sandstone, is shown by the thin section in Fig. 1.2a. The relative amounts of all the different minerals and the porosity is shown in Fig. 1.2b. This is ideally what our petrophysical interpretation would show. In this particular case there are only five minerals plus whatever is in the pore space. Three of the minerals only comprise a few volume per cent of the total and in practice log analysis is unlikely to be able to determine them separately. They are therefore combined with the quartz as 'matrix'. Because the matrix is mainly quartz it would have similar properties to pure quartz and in fact could conceivably be given the properties of quartz. The end product is thus a three-component system shown in Fig. 1.2c. Notice that although in this case clay is only a minor component it is still being treated as a separate component. We will see in a later chapter why clay or 'shale' is singled out for special treatment.

At the end of the day any interpretation has to be a compromise between the maximum number of components that can be resolved with the available data
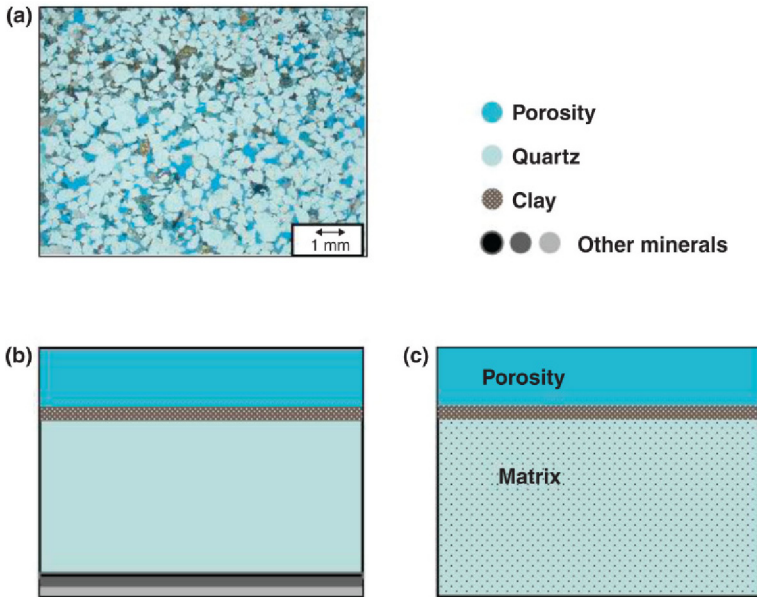
FIGURE 1.2 The relationship between the real rock and the petrophysical model (see text for explanation).

and the minimum number of components to adequately describe the rock. In most oilfield applications it is the porosity that needs to be most accurate and the model should be chosen to ensure that occurs. In reality we are unlikely to be able to faithfully reproduce the volume fraction of each component at every point in the well. But the model should at least underestimate a volume fraction as often as it overestimates it and at any depth the departure from the true value should not be excessive. More rigorously stated, it should accurately reproduce the mean and distribution in a particular interval.

Archie's equation can be used as an illustration. By re-arranging Eq. 1.2 we can write an equation to estimate the porosity of a water-bearing sandstone from its resistivity. To do this we need to find the constants a and m (we will assume we know R_w). For the core plug data for Cretaceous sands shown in Fig. 1.1 these turned out to be $a=0.56$ and $m=2.21$. But close inspection of the plot shows that there is scatter in the data. It is tempting to dismiss this as experimental error, but in fact, most if not all, of the scatter is real. Every plug has a different pore geometry and we will see later that leads to a unique value of 'm'. But unless we have some independent way of finding 'm' at every point in the reservoir we have to assign it a constant value to make progress. In other words we have to simplify things and so we assume an acceptable approximation is to make 'm' the same everywhere. The difference between the measured porosity and the prediction from the line is an indication of how good this

approximation is. The fact that the line passes through the cloud of points shows that we would underestimate porosity as often as we overestimate it.

Finding composition from one or more physical properties is the most common application of petrophysics, but the reverse problem of predicting one or more of the physical properties knowing the composition is also important. This is required for properties that cannot be directly measured with logging tools. A particular application is to calculate the density and acoustic properties of porous rocks filled with different mixtures of water, oil and gas. This allows geophysicists to predict how reliable seismic is for identifying hydrocarbon pools. It is also routinely used in log analysis to estimate the physical properties of the solid part of a rock – the matrix, which are an input to porosity calculations.

1.7 PHYSICAL PROPERTIES OF ROCKS

The rocks that concern us typically consist of a large number of minerals plus some fluid in the pore space. The latter invariably includes water but there may also be oil or gas in its natural state and/or an artificial fluid that has been introduced by the drilling process (e.g. mud filtrate). In the laboratory the original fluids will almost certainly be replaced by other purer fluids that are used for 'cleaning' the pore space and for making measurements under controlled conditions. These include the gases – air, nitrogen and helium and liquids such as chloroform, decane and mercury as well as water of course.

Carbonates are normally simpler than clastics but still consist of at least one solid mineral plus fluids. In some basins thick evaporate beds are encountered, in which a single mineral with very well-defined properties is present (most commonly halite or anhydrite but occasionally something more exotic). But in general the rocks that form reservoirs, seals and/or sources are complex and contain at least a dozen components at more than trace level. The log analyst must deduce what these components are and in what proportion. In Table 1.1 some of the commonly occurring minerals are listed, together with the properties that can be measured by, or are exploited by logging tools. More extensive tables can be found in text books on mineralogy, the chart books published by logging companies and possibly the help files associated with log analysis software.

Some of the properties shown in Table 1.1 are familiar (density and velocity) but more than half are properties that come from atomic and nuclear physics. These are the 'cross-sections' shown in the right-hand columns. Rightly or wrongly the cross-sections are converted into quantities, which are more easily related to petrophysical properties. The best example is probably the so-called neutron porosity, which is the standard output from most modern neutron tools. The neutron porosity is computed from count-rates that are themselves determined by the nuclear properties of the formation, borehole and the tool itself.

One property that is almost always measured during logging but is not included in Table 1.1 is formation resistivity. This will be discussed at length

TABLE 1.1 Some Physical Properties of Commonly Occurring Minerals

		Density (g/cm^3)	V_p (m/s)	V_s (m/s)	U (b/cm^3)	Sigma (cu)	Capture x_c (mb)	Scatter x_c (mb)
Silicates								
Quartz		2.65	5443	3464	4.79	4.3	40	
Orthoclase	K-Keldspar	2.56	4417		7.21	15.5		
Plagioclase	Na-Feldspar	2.59	6220	3586	4.35	7.5		
Muscovite	Mica	2.83	6220	2046	6.74	16.9		
Biotite	Mica	3.01	6000	1361	18.75	29.8		
Clays								
Kaolinite		2.59			4.44	14.1		
Illite		2.64			8.73	17.6		
Montmorillonite		2.06			4.04	14.1		
Carbonates								
Calcite		2.71	6220	3448	13.77	7.08	84	145
Dolomite		2.87	6927	4233	9	4.7	47	135
Siderite	FeCO$_3$	3.94	6485		57.14	52.31	650	
Sulphides								
Pyrite		5.01	7776	4908	84.68	90.1	1200	

(Continued)

TABLE 1.1 Some Physical Properties of Commonly Occurring Minerals *(cont.)*

		Density (g/cm^3)	V_p (m/s)	V_s (m/s)	U (b/cm^3)	Sigma (cu)	Capture x_c (mb)	Scatter x_c (mb)
Evaporites								
Halite		2.16	4549	2540	9.45	754		
Anhydrite		2.96	6096		14.93	12.45	150	
Coal								
Anthracite		1.47	2903		0.23	8.7	120	
Lignite		1.19	1905		0.24	12.8		
Fluids								
Water		1.00	1550		0.36	22		
Light oil		0.65	850			20		
Heavy oil		0.85	1300			22		
Gas		0.1	440			3		

Notes: Properties of the fluids depend to a varying degree on pressure, temperature and composition. Typical values are given in the table.

when saturation is considered for now it will just be noted that all the minerals in Table 1.1 except pyrite and water have resistivities that are too high to be measured by logging tools.

It is worth noting that there is some disagreement on the value a particular property has for a particular mineral. Often, the differences are too small to bother about but for clay minerals there can be substantial disagreements between different sources. The reasons for this will be discussed in the sections on clay and shale volume. But as an example different contractor chart books give the density of dolomite as anything from 2.84 to 2.87 g/cm^3 but for illite, a commonly occurring clay, the range is from 2.52 to 2.64 g/cm^3.

Since we invariably deal with mixtures, the properties we measure are some form of weighted average of the properties of the components. The equations that describe how the property of a mixture depends on the properties and proportions of its components are called 'mixing laws'. Since in log analysis we always deal with volume fractions, the mixing laws we use are written in terms of these.

The simplest conceivable mixing law is the volume weighted average of the individual component's properties. Density strictly follows this law and so the density of a mixture is given by:

$$\rho = \Sigma \cdot V_i \cdot \rho_i \qquad (1.4)$$

Where V_i is the volume fraction of component 'i' and ρ_i is its density. This is subject to the additional constraint that the volume fractions sum to unity.

$$\Sigma \cdot V_i = 1 \qquad (1.5)$$

In the special case of a two-component mixture consisting of 'fluid' and 'matrix', the volume fraction of fluid is the porosity $\emptyset$. So Eqs 1.4 and 1.5 can be combined to give:

$$\rho = \emptyset \cdot \rho_{ma} + (1 - \emptyset) \cdot \rho_{fl} \qquad (1.6)$$

In which ρ_{ma} and ρ_{fl} are the density of the matrix (solid) and fluid components. This can be re-arranged to give porosity as a function of density. This is a rare example of an equation in petrophysics that comes from first principles rather than being empirical.

In general the mixing laws are more complicated than Eq. 1.4. There are several reasons for this:

1. The physics that determines the way the components interact (e.g. velocity of sound).
2. Exotic minerals characterised by extreme values of a property. Even at trace levels these will strongly influence the overall value (e.g. neutron porosity).
3. The tool does not directly measure the property of interest, but rather responds to something else that correlates with it (e.g. the density log actually measures gamma-ray absorption and scattering).

Some of the other mixing laws that are found in log analysis are:

1. Power law

$$X = \Pi \cdot X_i^{V_i} \tag{1.7}$$

(The symbol Π means multiply all the terms together.)

The equation gives the 'geometric mean' of a population but in this case it is subject to the additional requirement that the V_i sum to unity. An example of its use is in the calculation of the thermal conductivity of a mixture.

2. Wood's law

$$X = \frac{1}{\sum} (V_i / X_i) \tag{1.8}$$

This is the equation of the harmonic mean of a population although again there is the additional constraint that the V_i's sum to unity. An example of its use is to estimate the bulk modulus of a mixture of fluids, which is used in predicting seismic velocities (e.g. oil and water). Note that if X is replaced by its reciprocal $1/X$, the equation takes exactly the same form as Eq. 1.4. This is a common trick to simplify a response equation.

Other equations are more specific. For example, one equation for the electrical conductivity (C) of a shaly-sand is:

$$C = C_w \cdot V_w^2 + C_{sh} \cdot V_{sh}^2 \tag{1.9}$$

Where V_w is the volume of water (the porosity) and V_{sh} is the volume fraction of shale. In this case V_w and V_{sh} do not sum to unity. This is because the remaining component – the matrix (given by $1 - V_w - V_{sh}$) – is assumed to have a conductivity of zero. Incidentally, conductivity is just the reciprocal of resistivity, we could have written 1.9 in terms of resistivity but it would be a lot more complicated and less obvious what it was saying. When we come to look at saturation equations later in the book we will frequently use this trick of working with conductivity.

1.8 FUNDAMENTALS OF LOG ANALYSIS

Log analysis is a specific case of using physical properties to gain information on the composition of rocks. Normally we are interested in making a quantitative analysis but sometimes all we can achieve or need is a qualitative or semi-quantitative interpretation. Examples of these various levels of interpretation are:

Level	Examples
Qualitative	Coal, facies, net reservoir, fractured interval
Semi-quantitative	Source rock maturity, fracture index
Quantitative	Porosity, pyrite concentration, water saturation

This is analogous to analytical chemistry where sometimes all that is required is an indication of presence or absence of a particular chemical (e.g. routine screening at airports for explosive residues). On other occasions the analyst needs to quantify the amount of a substance (e.g. measuring the H_2S content of a gas sample to determine whether special alloys are needed to construct a processing plant). Analytical chemists can use 'wet chemistry' or instrumental methods to perform these tests, but either way they seek a method, which is very sensitive to the chemical of interest and insensitive to everything else that might be present. The same is true in log analysis but here we are limited to a few instrumental methods and may have to accept that they respond to an extent to all the components.

We have already seen that to make progress we have to simplify the description of the rock to a handful of components and at its simplest log analysis attempts to determine the amount of shale, matrix and water in the system. By convention these are expressed as volume fractions. The hydrocarbon volume is normally found from the difference between porosity and water volume fraction. The familiar petrophysical properties are all either equal to these volume fractions or are defined in terms of them. For example, porosity is the sum of the volume fractions of water and hydrocarbon. Water saturation is the ratio of the volume fraction of water to porosity.

Life would be simple if we could find three different measurements that individually respond strongly to one of shale, matrix or water and not at all to the other two. Then measurement one would determine the shale volume, measurement two water volume and measurement three the volume of matrix. Unfortunately, no such combination exists, although we may come close with the gamma ray giving the shale volume more or less independently of anything else and electrical resistivity giving the volume of water. Unfortunately, most logs respond quite strongly to clay and hence shale so the analysis must somehow correct for this. Normally we rely on different logs responding slightly differently to the different components to unravel the mixture. The details of how this is accomplished are what a lot of this book is about.

Assuming we can find a set of logs that allow us to determine the volume fractions of each component there are two basic approaches to log analysis.

1. Deterministic Analysis. The petrophysical properties are written as explicit equations of log and possibly other petrophysical properties. Re-arranging Archie's equation to give porosity as a function of resistivity is a case in point.
2. Probabilistic Analysis. In which the computer answers the question 'What combination of minerals gives the observed set of log responses?' This involves a lot of computation and in fact these methods would be better described as matrix inversions. In principle, they allow as many components to be solved for, as there are independent input logs. Consequently, these methods often claim to be able to not just find the shale volume, but to actually resolve volume fractions of individual clay minerals (illite, kaolinite

etc.). This very detailed description relies on (a) accurate input logs and (b) mixing equations that reproduce what happens in nature.

Although probabilistic – or matrix inversion – methods are popular, particularly with service companies; we will not discuss them in depth in this book. This is because it is easier to appreciate what is happening with a deterministic equation and ultimately the same underlying principles apply to both.

1.9   A WORD ON NOMENCLATURE

Petrophysicists and log analysts often use the same words as geologists, but have a different definition in mind. We have already come across a couple of examples. Normally the geologists got in first and most Earth sciences dictionaries give their definitions, but the terms are so deeply embedded in petrophysics that we are unlikely to stop using them.

Anyway here are a few of the worst offenders:

1. Matrix: to a petrophysicist this is the solid part of a porous rock excluding any shale or clay. There are numerous parameters in log analysis with 'matrix' in them, for example matrix density, neutron matrix.
2. Secondary porosity: in log analysis this is porosity other than inter-granular porosity. Fractures are an obvious example but it also includes vugs, moulds etc. The reason for distinguishing it is that secondary porosity can affect logging measurements differently to inter-granular porosity. In fact it is claimed that some measurements are not affected by secondary porosity at all.
3. Shale: some petrophysicists actually mean 'clay' when they talk about 'shale', so when they discuss 'shale volume' they are really referring to the volume fraction of the clay minerals in the formation. We will return to this at length later in the book but for now we note that shale is a type of rock that normally contains a lot of clay – 60–80% by volume – but is not composed exclusively of clay. Clays are a class of silicate minerals and it is quite possible for a bed to contain clay but no shale. To make matters worse many petrophysicists do not even stick to either clay or shale, but use the two interchangeably.

1.10   THE FUTURE OF THE PROFESSION

For many years petrophysics seems to have been regarded as quite a specialist, stand-alone discipline (despite occasional indignant articles in the trade press vehemently claiming the petrophysicist was central to understanding the sub-surface). In practice log and core data from wells was converted to petrophysical properties either as a continuous function of depth or a single average value that characterised an interval (geological formation, reservoir unit or a hydrocarbon leg). The results were passed on to the other disciplines – often as an attachment on an e-mail – to incorporate in volumetric calculations or a

coarse dynamic reservoir model. The petrophysicist, their work done, moved to the next problem. There appeared to be no real reason for the other disciplines to understand what exactly the petrophysical interpretation involved or for the petrophysicist to know how their results were being used.

The increasing use of software to build detailed 3D geological models of reservoirs has meant that petrophysics has to be properly integrated with the other sub-surface disciplines. The model builder needs to know what assumptions have gone into the creation of the petrophysical property curves and the petrophysicist needs to know that their results are being used appropriately. Consequently a working knowledge of practical petrophysics is no longer just a 'nice to have'. This book is designed to provide that knowledge.

Page left intentionally blank

Chapter 2

Petrophysical Properties

Chapter Outline

2.1   INTRODUCTION

In the previous chapter, it was noted that much of practical petrophysics boils down to finding porosity, water saturation and permeability. For sure, some or all of these petrophysical properties may require some intermediate property to be determined first (e.g. shale volume). Furthermore, the so-called unconventional reservoirs require several additional properties for an adequate description, but even here the 'Big Three' are still key inputs to their evaluation. In this chapter, we will look at the petrophysical properties that are the end product of an interpretation. This includes apparently easily defined properties such as porosity, as well as the more subjective measures that are used to describe reservoirs ('Net' for example). Many of these properties are easy to define on paper but when they are applied to real rocks with water occupying at least some of the pore space, the definitions can become quite ambiguous.

2.2 POROSITY

Porosity is the ratio of pore volume to bulk volume or, if you prefer, the volume fraction of fluids in the rock. That seems simple enough and for many artificial systems such as a porous solid made from plastic spheres it really is that simple. Unfortunately, for real rocks with water in the pore space it is not so obvious what constitutes the pore volume. As we shall see this is the result of water interacting strongly with many of the commonly occurring minerals. This ultimately leads to two rival ways of describing porosity: total and effective. Both are equally valid descriptions but one needs to be consistent and clear which is being used. We will return to this issue later but to begin with we will assume there is no ambiguity about what is solid and fluid.

As a ratio, porosity can have a value between zero and one (the latter is pure fluid of course). It is also commonly expressed as a percentage or in porosity units (pu), which are actually the same thing. Porosities for real rocks can vary from 0 to at least 50% although porosities in excess of 35% are unusual. Where porosities in excess of 35% do occur, they typically – but not exclusively – apply to carbonates. Figure 2.1 shows some porosity ranges for some real examples of clastic and carbonate reservoir rocks.

To make progress in petrophysics it is helpful to start with simple models (either physical or theoretical). Suitable models to understand porosity are porous solids made from spherical grains. The simplest examples use uniform spheres arranged in an array. In these cases the porosity can be calculated exactly. Some examples of small parts of such arrays are shown in Fig. 2.2a, it shows how four
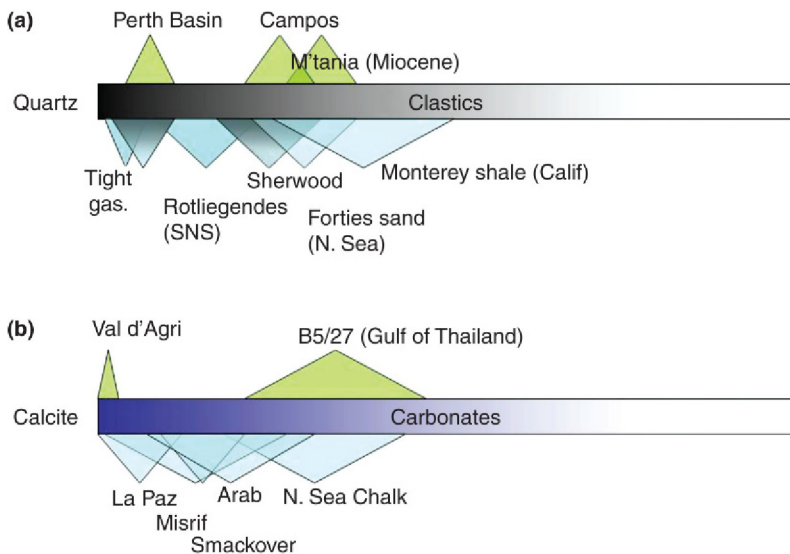
FIGURE 2.1 Porosity ranges for clastics (a) and carbonates (b).
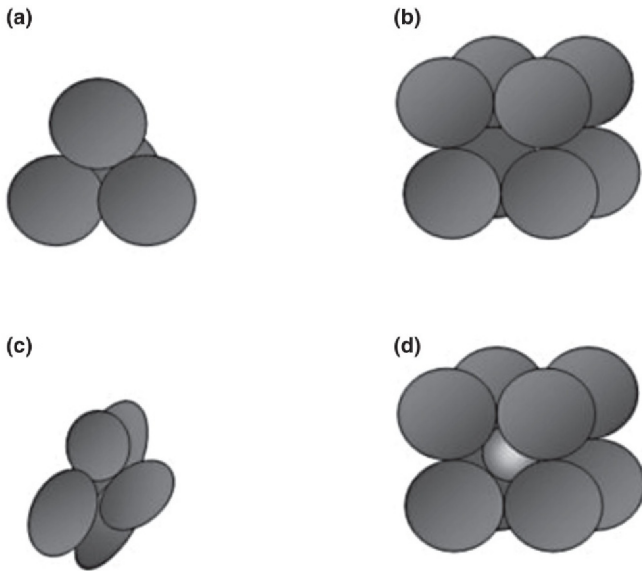
FIGURE 2.2 Simple models of porous solids. (a) Close packed array of identical spherical grains. (b) Simple cubic array of cubic grains (it is particularly easy to calculate the porosity of this). (c) A porous solid made from elongated grains. (d) The simple cubic array with smaller spherical grains occupying some of the original pore space.

spheres pack together to occupy the smallest volume. If this is extended in all directions we get a 'close packed' array so-called because it is the most efficient way of packing uniform spheres. It has a porosity of 25.9% and this is therefore the lowest porosity possible with uniform spherical grains (the calculation of the porosity is simple but tedious). Any other arrangement, including one in which the spheres are not uniformly arranged will have a higher porosity.

The porosity of the appropriately named simple cubic array is much easier to calculate and is left as an exercise (Fig. 2.2b). Using these simple models we can qualitatively predict what happens as we move to something more complicated and more realistic. In Fig. 2.2d smaller spheres are introduced into the spaces between the larger spheres. This simple model is therefore saying that porosity falls as sorting gets poorer. Similarly if we start to elongate the grains we expect the porosity to be reduced.

The simple models using spheres are a good way to understand clastic rocks but a similar approach could be applied to fracture porosity. In that case we would probably start by putting plane parallel-sided cracks with a well-defined aperture into a solid block. This approach illustrates a point made in Chapter 1 that practical petrophysics relies on integrating experimental results from real rocks with highly idealised models such as those in Fig. 2.2. The latter helps us to explain the former and in the process to give us confidence in the results.

It is now time to consider real rocks in contact with water because this is what leads to the two alternative descriptions of porosity. Water has the ability to form weak chemical bonds to the surfaces of some mineral grains, this causes the grains to be surrounded by a strongly adhering film of water. Substances with free oxygen atoms at their surfaces, such as the silicates, are particularly prone to this. Clays have a particularly strong affinity for water and furthermore their morphology results in a high surface area so, they can bind a relatively large volume. Because this water is not going to move under natural conditions it could be considered part of the matrix. Nearly all the silicates will bind to water and other examples that immobilise relatively large volumes of water include the micas – including glauconite – and chlorite. To keep the word count down, for the remainder of this chapter we will refer to 'clays' when we really mean any silicate that binds significant volumes of water. Actually, log analysts often include glauconite and chlorite as clays precisely because they can fix large amounts of water.

The way bound water is accounted for determines which of the two different ways of defining porosity to use. In the effective porosity system, bound water is no longer counted as part of the pore space. It's proponents argue that, as it cannot move, it might just as well be solid rock. To supporters of the total porosity model, however, it is still water – i.e. fluid – and therefore it contributes to the pore volume. If no clay is present the two descriptions give the same or almost the same answer. Conversely, if all the water in the system is associated with clay, as will occur in a claystone or shale, the effective porosity will be zero but the total porosity will be finite. The relationship between total and effective porosity and the volume of clay is shown graphically in Fig. 2.3. The intelligent reader will see that there is merit to both arguments and will keep an open mind.

As already noted, clay is not the only type of mineral capable of binding and immobilising water so that strictly speaking the above descriptions should have been made a bit more general. In fact, it is not even just silicates that have an affinity for water. Nevertheless in practice the difference between the two descriptions tends to be most significant in 'shaly sands' in other words sands with a high shale or clay content. We will look at these in more detail later.

Total porosity is often justified as the better description because that is what is measured in core analysis. This is because most companies thoroughly dry the core plugs before measuring porosity and even the water bound to the clays is driven off. So, if a porosity calculated from logs is compared to measurements on core plugs we are implicitly comparing to a total porosity. Furthermore, to find the total porosity from logs we only need to find the total volume of water in the system (although with conventional logs that may be easier said than done).

To find effective porosity we need to determine the proportion of the water that is tied up with the clays. To do this with core plugs we have to stop the drying process after the water in the pores has been removed but before the clay bound water starts to evaporate. In reality that is going to be almost impossible to judge. Nevertheless, effective porosity is arguably what we are really interested
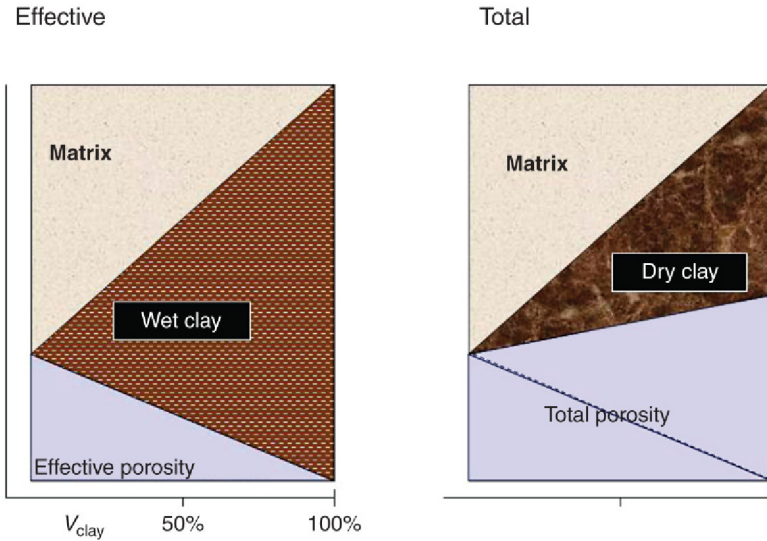
Effective

Total

FIGURE 2.3 Relationship of total and effective porosity to the volume of clay in the system. In the effective porosity model water associated with the clay is excluded from porosity, so that at 100% clay the porosity is zero. In the total porosity model the clay water is still counted as water and in this particular case the porosity as at its highest at 100% clay. When there is no clay both systems agree.

in because it is related to the volume of fluid that can be extracted from the rock. Furthermore, as effective porosity is the volume fraction of moveable fluid, it controls permeability and is likely to influence the seismic response. For many people the most compelling reason for using effective porosity is, however, that it avoids the counter-intuitive implication of shales having non-zero porosity.

Total and effective porosity are end member descriptions of porosity which are related by

$$\emptyset_T = \emptyset_E + V_{clay} \cdot \emptyset_{clay} \tag{2.1}$$

where V_{clay} is the volume fraction of clay ('clay volume') and $\emptyset_{clay}$ is the porosity of the clay, in other words the volume fraction of water in pure clay. Suffixes 'T' and 'E' refer to total and effective porosity, respectively. This equation is just an analytical description of the diagrams in Fig. 2.3.

For completeness, it should be noted that the concept of effective porosity is sometimes introduced by considering a porous rock in which a proportion of the porosity is present as isolated pores. This porosity is non-effective because it's contents cannot ever move. In reality, however, truly isolated pores are unusual if only because fluid is generally needed to create the pore in the first place. Examples of truly isolated pores are vesicles in igneous rocks or fluid inclusions within secondary over-growths.

Finally, although all of the above discussion has concerned water bound to mineral grains it should be noted that the difference between total and effective porosity can be due to organic matter. The North Sea Piper sandstones, for example often include bitumen in their make-up, that the usual cleaning and drying procedures applied prior to core analysis, do not remove. Special solvents can dissolve it, however, and as a result the measured porosity increases by typically 1 pu. Is the bitumen part of the oil or part of the rock?

There are two properties that are directly related to porosity that should be mentioned for completeness. The void ratio is used in engineering geology and soil science and is defined as the ratio of pore volume to solid volume. In symbols

$$\text{Void ratio} = \frac{\varnothing}{1 - \varnothing}$$

where porosity is given as a fraction. Void ratio is always greater than porosity and approaches infinity at the fluid point. It is obviously easily converted to porosity and is only likely to be encountered in research papers. The other quantity that can be encountered in the petroleum literature is the 'solidity' this literally compliments porosity and is given by

$$S = 1 - \varnothing$$

It is occasionally used in compaction curves (i.e. general relationships describing how porosity – or solidity – changes with depth). Introducing a new symbol just to avoid writing '1' seems a bit extravagant to this author, but there is no escaping the fact that some important papers have been written in terms of it. Neither property will be discussed again in this book.

2.3 SATURATION

Saturation is the proportion of the pore space occupied by water. To be safe it is better to explicitly talk of 'water saturation' because if unspecified it is natural to assume it refers to what we are really interested in: the hydrocarbons. Like porosity it is a ratio so can either be expressed as a fraction or a percentage. The term 'saturation units' is also sometimes used, this is simply saturation expressed as a percentage (SU).

Water saturation is described in two ways:

1. Evaluated level by level at the well(s).
2. Modelled at every point in a structure.

The former invariably involves calculating water saturation as a continuous curve along the well track using log analysis. More often than not Archie's equation or one of it's numerous modifications forms the basis of the analysis. The modelling approach seeks an equation that calculates water saturation at any point in the structure (including away from the wells). The equation is known as a saturation-height function (SHF) because as a minimum it is a func-

tion of the height above the contact (or more properly the 'free water level'). Normally, there will be other inputs such as permeability and geological zone or facies. The SHF is needed to model the distribution of hydrocarbons throughout a structure, which in turn is needed to accurately calculate in-place volumes and as a starting point for dynamic modelling. We will consider both these approaches in this book.

Log analysis generally works by calculating porosity by some means and the volume of water in the system using a different tool (often, although not necessarily, resistivity). Water saturation is then given by:

$$S_w = \frac{V_w}{\varnothing} \tag{2.2}$$

where V_w is the volume fraction of water in the rock. Notice that any error in the porosity carries through to saturation and since the volume of water will be subject to some uncertainty as well, saturation is always less certain than porosity.

Water saturation is a ratio that takes values between one and zero. The latter is hardly if ever achieved in the subsurface except possibly above the water table in arid locations (in the laboratory plugs are routinely dried completely). Instead the water saturation reaches a lower limit known as the 'irreducible water saturation'. This is such an important property of the rock that it will be given its own symbol (S_{wir}). It can vary from a few per cent in high-quality reservoir rocks to a hundred per cent in shales. It is mainly a result of the strong affinity of commonly occurring minerals for water that was discussed in the porosity section.

Most rocks start out saturated with water and if oil or gas enters the pores, this occurs later. But a reservoir rock that has contained oil or gas cannot normally be returned to a state of complete water saturation (apart from in a laboratory). A low saturation of oil or gas typically remains trapped in the rock. This is referred to as residual oil or gas and the saturations are given the symbols S_{or} (S_{gr}). Values typically range from a few per cent to 25%.

There is not much more to say other than to ask 'What Porosity?' In an effective porosity model the volume of water (V_w) should not include the water associated with the clays. In a total porosity model on the other hand the clay water is included in V_w.

Log analysis software often computes effective and total saturations where

$$S_{we} = \frac{(V_{wt} - V_{w,clay})}{\varnothing_e} \tag{2.2a}$$

$$S_{wt} = \frac{V_{wt}}{\varnothing_t} \tag{2.2b}$$

where V_{wt} is the total volume of water in the system. Notice that the volume of hydrocarbon calculated should be the same in either model. The volume of hydrocarbon per unit volume of formation is given by:

$$V_{hc} = \emptyset(1 - S_w) \qquad (2.3a)$$

So that

$$\emptyset_e(1 - S_{we}) = \emptyset_t(1 - S_{wt}) \qquad (2.3b)$$

2.4 PERMEABILITY

Permeability differs from the previous two properties in that it is not dimensionless and also, generally, it is anisotropic (see below). Permeability is the constant that links flow-rate to the pressure applied to a fluid. For a simple system such as a cylindrical core plug with flow parallel to the axis of the plug, the flow-rate (Q) is given by:

$$Q = \frac{kA\Delta P}{l\eta} \qquad (2.4)$$

where A is the cross-sectional area, l is the length of the plug and η is the viscosity of the fluid. ΔP is the pressure drop across the plug. The constant of proportionality k is the permeability. Equation 2.4 can be generalised to deal with more complicated geometries but here we will deal with the simple linear flow case. A dimensional analysis shows k has units of length squared so the SI unit is m². The petroleum industry of course has its own unit: the Darcy (D) which, by coincidence, is very close to $1\,\mu m^2$ which is technically an SI unit ($1\,D = 0.987\,\mu m^2$).

Most permeabilities of reservoir rocks are actually given in milliDarcy (mD).

Permeabilities of rocks that concern us range from zero to tens of Darcy. The highest values are typically found in tar sands. Commercial laboratories using the apparatus described in Chapter 3 normally do not quote values below 0.01 mD and many people believe 0.1 mD is the lower limit of reliable data. Tight gas reservoirs are often defined as those with average permeabilities of less than 0.1 mD and samples with permeabilities of the order of $1\,\mu D$ are not unusual. Special apparatus and/or a lot of patience is needed to produce trustworthy data at such low values.

As with porosity it is often helpful to consider highly idealised systems when thinking about what controls permeability. This allows permeability to be calculated analytically using the principles of fluid mechanics. Here we consider a block of rock with one or more cylindrical pores running through it (Fig. 2.4). It can then be shown that a plug of cross-sectional area A with a cylindrical capillary of radius 'a' running the length of the plug has a permeability of:

$$k = \frac{\pi a^4}{8A} \qquad (2.5)$$

$$Q = \frac{\Delta p.\pi a^4}{8\Delta l.\mu} \quad \text{Flow through a cylindrical pipe}$$

$$= \frac{kA\Delta p}{\Delta l.\mu} \quad \text{permeability equation}$$

$$k = \frac{\pi a^4}{8A}$$

$$\varnothing = \frac{n.\pi a^2}{A} \quad \text{porosity}$$

$$\Delta l$$

$$k = \frac{\varnothing.a^2}{8n}$$
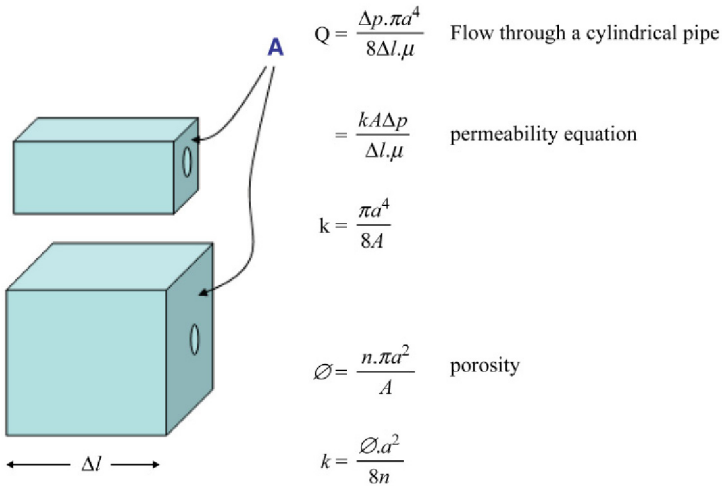
FIGURE 2.4 Simple capillary tube models (LHS) allow permeability to be computed from first principles (first three equations on RHS). Both models have a single cylindrical pore of radius 'a'. The upper model has a lower cross-sectional area and therefore a higher porosity (fourth equation on RHS).

Where K will be in milliDarcy if a is in micrometers (or very close). The fourth power term suggests permeability increases rapidly with pore diameter. This has two consequences:

1. The largest diameter conduits tend to dominate the permeability (open fractures are a case in point).
2. Within a 'conduit' the restrictions dominate the permeability. In other words pore throats exert far more control than the pore bodies.

Consider two plugs: one with a single 10-μm capillary through it and one with 100 1-μm capillaries. The total cross-sectional area of the capillaries is the same in both cases but the single large capillary gives a permeability that is 100 times higher than the alternative arrangement.

It is often stated that permeability has no relationship to porosity but actually for a particular rock-type the two are closely linked. Increasing porosity puts more pathways through the rock and so permeability increases. Again the simple capillary model shows this well. The porosity for this system is actually given by:

$$\varnothing = \frac{n\pi a^2}{A} \tag{2.6}$$

Where n is the number of capillaries within the cross-sectional area of the plug. Substituting this into 2.5 gives:

$$k = \frac{\varnothing a^2}{8n} \tag{2.7}$$
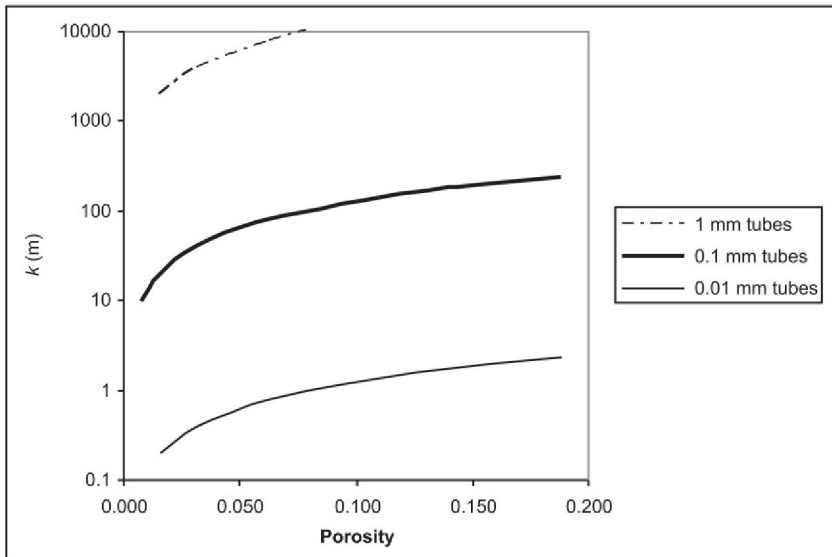
FIGURE 2.5 Permeability as a function of porosity for capillary tube systems like those illustrated in Fig. 2.4. The different trends refer to different capillary sizes. Increasing the capillary diameter by a factor of 10 gives a hundred fold increase in permeability.

So K increases linearly with porosity and the constant of proportionality includes the term capillary radius squared (The relationship is plotted on a conventional "semi-log" grid in Fig. 2.5).

Of course the conduits in real rocks are far more complicated than groups of parallel-sided capillaries embedded in an impermeable matrix. But we have already seen that petrophysics loses much of its mystique if one starts with simple conceptual models. In this case the 'bundle of straws' model has suggested:

1. permeability is controlled by the largest diameter conduits and
2. permeability tends to increase with porosity.

Moving from the 'bundle of straws' model to a real inter-granular porosity system the two most obvious complications are as follows:

3. The conduits no longer have a uniform diameter but rather have numerous constrictions produced by the pore throats.
4. The different conduits are interconnected.

For inter-granular porosity, it is then the pore throats that generally determine permeability. If these become cemented or if clay minerals start to form in them the porosity is only marginally reduced but permeability may be drastically reduced. This is one reason that nearly all petrophysical models single out clay as one of the key components.

2.4.1 The Klinkenberg Effect

Permeability is normally measured in the laboratory on cylindrical core plugs using a relatively low-pressure gas as the fluid. In a petroleum reservoir the moving fluids will almost certainly be at much higher pressures and may well be liquids. The change to reservoir fluids causes a reduction in permeability, this is known as the Klinkenberg effect after the physicist who first produced a quantitative model to explain it. The size of the change depends on the absolute permeability and tends to be relatively small for conventional rocks with permeabilities of more than 10 mD but in tight rocks it can amount to a 50% reduction or more.

The Klinkenberg effect is a result of the distance a molecule can travel between collisions – the mean-free path. This is inversely proportional to pressure. In a capillary, molecules that move more or less parallel to the axis are the ones that contribute to the flow and the further these can travel before being deflected by a collision, the higher the permeability. As the pressure on a gas increases, the mean-free path falls and the permeability appears to fall. In effect all the molecules ricochet around in the capillary and take longer to negotiate it.

Klinkenberg found that:

$$k = k_{inf}(1 + 4C\lambda/r) \tag{2.8a}$$

where C is a constant, λ is the mean-free path and r is the radius of the capillary. k_{inf} is the hypothetical permeability at infinite pressure but in practice it is assumed to be closer to the permeability at reservoir pressure. In practice the inverse relationship between mean-free path and pressure is used to re-write second term in brackets, giving:

$$k = k_{inf}(1 + b/P) \tag{2.8b}$$

where b is a constant. Although the equation was developed for a bundle of capillaries it is still assumed to apply for real porous rocks and b will be specific to a particular plug.

For the gases that are of interest in petroleum engineering the mean-free path at ambient conditions is about 0.04 μm (or a hundred molecular diameters). At reservoir conditions, it will be at least an order of magnitude less. Although these are small relative to the pore throats in a conventional reservoir rock they are large enough to have noticeable effect and typically permeabilities reduce by 10% on going from ambient to reservoir pressures. For tight gas reservoirs with permeabilities of the order of microDarcy the pore throats are much smaller and the Klinkenberg corrections can be substantial (50% or more).

The Klinkenberg correction is typically measured by taking a subset of plugs and repeating the permeability measurement at a series of increasing gas pressures. A plot is then produced of permeability against inverse pressure and a line is extrapolated to an inverse pressure of zero to give Klinkenberg permeability. It is variously given the symbol k_{inf}, k_{∞} or k_l (in this book we will use

the first of these). k_l is a reference to 'liquid permeability' which in turn is an acknowledgement that at high enough pressure the gas will liquefy.

It is important to realise that the Klinkenberg correction is designed to account for the increase in pressure on the gas in the reservoir. It is nothing to do with the compaction of the rock, which causes a further reduction in permeability.

2.4.2 Effective and Relative Permeability

All the discussion so far has assumed a single fluid is present in the pore space and in fact it is tacitly assumed that the fluid does not interact strongly with the rock. Under these conditions the permeability is sometimes called the 'absolute permeability'. In the natural state the pores may contain water and oil or gas. When the two immiscible fluids are present they are treated separately and each is assigned an 'effective permeability'. There are two of these of course: effective permeability to water and effective permeability to oil (or gas). In this book, we will give them the symbols K_w and K_o (or K_g in the case of gas). In general these depend on the water saturation but will always be less than the absolute permeability. There is a certain finite value of water saturation where the effective permeability *to water* will fall to zero and the effective permeability to oil (or gas) will reach its maximum value (this is the irreducible water saturation introduced in Section 2.3). There is also a maximum water saturation at which the rock becomes impermeable to oil and the effective permeability to water reaches its maximum value.

Effective permeabilities are related to absolute permeabilities by a special factor known as the 'relative permeability'.

$$k_w = kr_w (S_w) k \qquad (2.9a)$$

$$k_o = kr_o (S_w) k \qquad (2.9b)$$

Despite its name, relative permeability is dimensionless and always lies between zero and one.

The reason effective permeability is lower than absolute permeability can again be explained using the 'bundle of straw' model. In the section on porosity, we noted that many minerals have a strong affinity for water, that means, amongst other things, that the water saturation will never fall to zero. In the case of the cylindrical pores water forms a film on the surface of the pore, which effectively reduces its radius. Oil or gas is forced to flow in what is effectively a smaller pore and so the effective permeability is reduced. This simple explanation predicts that the smaller the pores, the higher is the irreducible water saturation and the lower is the effective permeability. At higher water saturations both water and oil are mobile and they both have finite effective permeabilities but the flows normally interfere so that the sum of their flow-rates is still less than could be achieved using a single fluid. At higher water saturations the oil(gas)
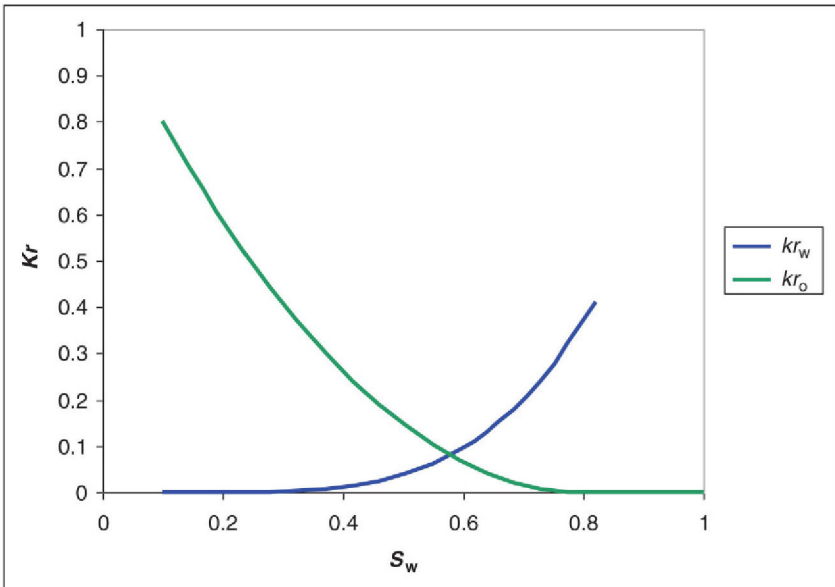
FIGURE 2.6 An example of a pair of relative permeability curves for a good quality clastic. The curves are colour coded but in fact could be distinguished by the direction they increase (kr_w in dashed line [blue in web version] and kr_o in dotted line [green in web version]). The curves within the range of water saturations defined by irreducible water (the lowest value of S_w) and residual oil (the lowest value for $1 - S_w$).

is immobilised because a small quantity remains trapped in dead end pores or is unable to squeeze through the small pore throats (this is the residual oil or gas).

An example of a pair of relative permeability curves for a clastic rock are shown in Fig. 2.6. The end point saturations are irreducible water saturation and the residual oil saturation. The curves do not continue beyond these limits. The highest effective permeability will be to oil when the water is at irreducible water saturation. The highest effective permeability to water occurs at the residual saturation but in this case it is only about 40% of the absolute permeability. At an intermediate water saturation of about 50% the effective permeabilities for both oil and water are low and even their sum is a lot less than the absolute permeability. This behaviour is fairly typical of inter-granular pore systems.

In Fig. 2.6 the relative permeability curves are limited to the saturation range defined by irreducible water and residual oil. This is the range that could exist in the reservoir today, but at some point in the past it would have been completely saturated with water so that a high-end effective permeability at an S_w of 100% could occur naturally. The argument that was used to show that effective permeabilities in the presence of water are generally less than absolute values still applies, however. In other words water adhering to the grain surfaces effectively reduces the diameter of the pore throats, even if it is the only fluid present. The net effect is that the permeability measured with pure
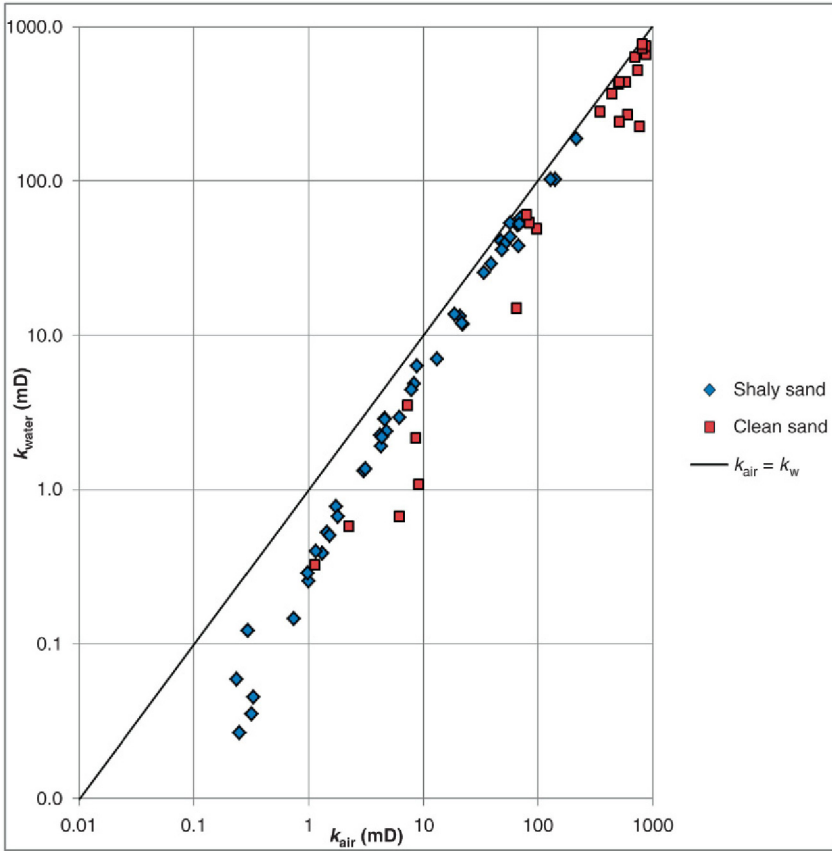
FIGURE 2.7 A cross-plot of permeability measured with pure water against permeability measured with dry air for approximately 100 plugs from two different reservoirs.

water should be less than that measured with an inert fluid such as air, helium or even oil. Figure 2.7 shows some real measurements made with air and water and indeed the permeability in water is lower. In order to distinguish this permeability we will give it the symbol 'k_w' (this is to distinguish it from the true effective permeability to water – in the presence of oil or gas – for which we decided to use an upper case K). Notice that the difference tends to get smaller as the permeability increases.

2.5 SHALE AND CLAY VOLUME (V_{shale} AND V_{clay})

In Chapter 1 we noted that real rocks are commonly modelled as a mixture of matrix, shale and fluids (or porosity). The shale volume is the volume fraction of the shale and obviously lies between zero and unity, the latter corresponding to a shale bed (or 0–100%). At any depth it is constrained by the volume of the

matrix, the porosity and any other minerals that are included in the model, since all the components of the rock have to sum to unity.

The shale volume in a typical conventional reservoir rock normally only amounts to a few per cent, so it is reasonable to ask why it is singled out for special treatment. There are several answers but probably the most important is that it gives some idea of how much clay is in the reservoir and how it is distributed. In fact some models actually use 'clay volume' instead of shale volume and the petrophysical model becomes matrix, clay and porosity. It must be realised that shale volume and clay volume are not the same thing and in fact with enough logs we could calculate both.

There are pros and cons to using either shale or clay volume, but regardless of which is ultimately used we should be clear what exactly we are trying to estimate. Regrettably, this is often not the case and the terms shale volume and clay volume are used interchangeably. As the name suggests the shale volume is the volume fraction of shale in the formation. Shale is a rock type consisting largely, but not exclusively of clay minerals. In fact they typically contain 60–80% clay by volume, the remainder being made up of silt or sand-sized grains of quartz, other silicates, carbonates and possibly organic matter. Clay volume on the other hand is the volume fraction of clay minerals in the formation and so in a pure shale bed it is likely to be about 70%. Sandstones often contain at least a small amount of clay and so clay volume will be more than 0% even if no shale is present.

To summarise what has been said above.

In a pure shale the following inequality holds true:

$$V_{clay} < V_{shale}$$

In a clean sandstone the following holds true:

$$V_{clay} > V_{shale} = 0$$

In the intermediate case of some shale being present in a sandstone bed it is not possible to generalise. But clay that is associated with the sand may just about compensate for the non-clay component of the shale. This is probably the main reason we have got away with using shale and clay interchangeably.

Knowing how much clay is in the system is important because it tends to have a disproportionate influence on the petrophysical properties.

Clays are of particular interest because of the following:

1. They introduce most of the bound water, which represents the difference between total and effective porosity.
2. The physical properties of clays can be significantly different to other silicates and so they need to be quantified in order to accurately estimate porosity.
3. Some clays introduce an excess electrical conductivity which if not accounted for results in water saturation being overestimated.

4. When present in a reservoir rock, clay minerals often reduce permeability dramatically.
5. They contain hydrogen as part of their structure which can be wrongly assigned to water and hence porosity.

To appreciate why clays do all these things requires a bit more understanding of their structure and chemical make-up and the next sub-section addresses this. Before embarking on this, however remember that in Section 2.2 we extended the name 'clay' to include some other silicate minerals that petrophysically behave in a similar way. We will continue to include them in the discussion.

2.5.1   Clay Minerals

Discussion of clays is complicated by the fact that the word describes a class of silicate minerals, a texture and a rock type. Texturally clay refers to grains that are smaller than 3.9 μm (1/256 mm). As a rock, clay consists of clay minerals and 'accessory' non-clay minerals. The clay minerals are generally quite plastic, which can have important consequences for clastic reservoir rocks, which include clay grains in their make-up. The clay minerals nearly all consist of stacked silicate sheets and they are therefore highly anisotropic at the microscopic level. The ease with which the sheets can slide past each other is the origin of the plasticity of clays.

Chemically they can be thought of as being derived from quartz by replacing some of the silicon with aluminium and other metal atoms. Quartz has a highly symmetrical crystal structure in which every silicon atom is connected to four oxygen atoms, which form the vertices of a tetrahedron. This arrangement allows the oxygen atoms to get as far apart as possible. Each oxygen atom is in turn bonded to two silicon atoms so that the silicon atoms also form a tetrahedral arrangement with an oxygen atom lying between every neighbouring pair. This highly symmetrical 3D lattice is very rigid and hence quartz forms hard grains that are difficult to deform. So replacing some silicon with aluminium completely changes the mechanical properties.

Kaolinite

The simplest of the true clays is kaolinite, this can be thought of as quartz with half the silicon atoms replaced by aluminium. The stoichiometric composition is $Al_4[Si_4O_{10}](OH)_8$ (cf. Si_8O_{16} for quartz). The replacement is not random, however and each sheet essentially comprises a layer of silicon atoms and a layer of aluminium atoms. A highly simplified 2D representation of the structure is shown in Fig. 2.8. The sheets are arranged so that the aluminium layer of one sheet is always next to the silicon layer of the next sheet. Within the layers oxygen atoms lie between neighbouring silicon atoms and neighbouring aluminium atoms and they also bond the layers together to form a single sheet. In order to satisfy aluminium's bonding requirements approximately half the oxygen atoms in the aluminium layer also bond to hydrogen atoms.

Kaolinite structure
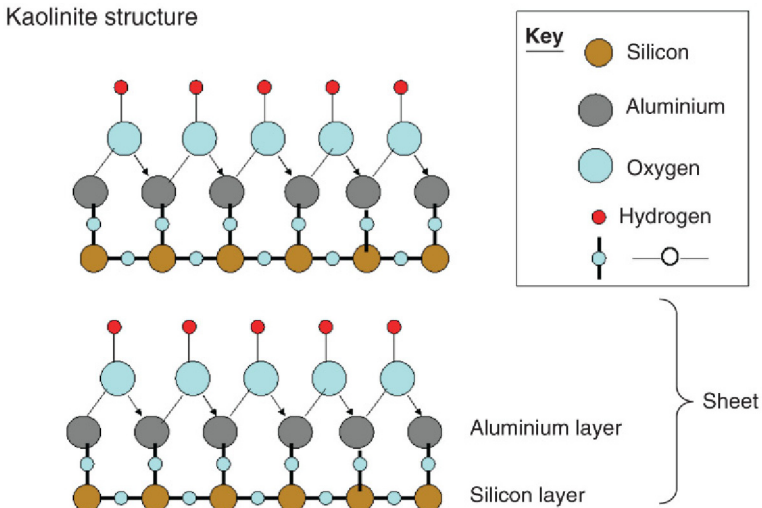
FIGURE 2.8 A simplified representation of the kaolinite structure. Two sheets are shown.

The structure of the individual sheets and their arrangement explains at least two of kaolinite's physical properties. The density is similar to quartz – it is variously reported as 2.60–2.68 g/cm^3 – and it has a high hydrogen index and therefore a high neutron porosity (of nearly 50 pu depending on the tool type). The high neutron porosity reflects the hydrogen atoms that are an intimate part of the molecular structure. In effect there is the equivalent of a water molecule for every silicon or aluminium atom in the structure. The density reflects the fact that the structure is similar to quartz with aluminium replacing some of the silicon. The atomic masses of silicon and aluminium are similar:

Si = 28.1 g
Al = 27.0 g

The gap between the sheets would produce a lower density relative to quartz but the two additional oxygen atoms increase the density. So, all in all a similar density is to be expected. These observations will be important when we use logs to quantify the amount of clay in shaly sand.

Illite

In terms of complexity the simplest clay mineral after kaolinite is arguably illite. This is also generally the most abundant clay mineral in the sub-surface. Illite basically has the following chemical formula:

Illite	$K\{Al_5Si_7O_{20}\}(OH)_4$	compare this with
Kaolinite	$\{Al_6Si_6O_{15}\}(OH)_{12}$	
Quartz	$Si_{12}O_{24}$	

Note: to aid comparison the formulas have been written so that there are always 12 silicon and aluminium atoms in each 'molecule'.

In other words there is a slight surplus of silicon compared to kaolinite. The precise formula varies, however and different samples can have slightly more potassium and slight variations in the silicon/aluminium ratio (this type of variation is typical of clays). The change in silicon/aluminium ratio causes a charge imbalance, which is compensated for by having less hydrogen and the addition of a potassium ion.

Illite is therefore a potassium compound that makes it radioactive. This is one of the reasons shales tend to be more radioactive than reservoir rocks. The individual sheets are thicker than in kaolinite and basically consist of a sandwich of an aluminium layer between two silicon layers. The potassium ions lie between the sheets and hold the structure together by electrostatic attractions to the negatively charged silicon layers. As with kaolinite the similarity to quartz in chemical composition and structure suggests that the density should be similar, it is in fact variously reported as 2.6–2.9 g/cm^3.

The Smectites

The smectites are typically found in young sediments, including soils and freshly deposited muds as well as tertiary deep-water deposits that form the reservoirs in the Gulf of Mexico, offshore Brazil and West Africa. When present they have a significant effect on log responses and it is these clays above all others that can produce quite misleading results if they are not accounted for.

The smectites exhibit a huge range of compositions. In particular, the aluminium/silicon ratio varies, some of the aluminium can and normally is replaced by iron and/or magnesium and the role of the potassium ion in illite can be taken by sodium or calcium. This variation can occur at the level of individual layers so that different samples from the same rock may show large changes in composition and properties. A consequence of this is that there is little agreement on their physical properties. Density, for example is quoted as anywhere from 2 to 3 g/cm^3 depending on the source of the clay. They also have a strong affinity for water, which is so intimately associated with them that, arguably, the dry clay is no longer smectite. This further contributes to uncertainty in defining their physical properties and is particularly significant to their electrical properties.

Hydrogen enters the smectite structure both as an integral part of the layers and as water that is tightly bound to the layer surface. So in general neutron porosities are high but it is mute point whether the bound water should be considered part of the clay or whether it is actually total porosity (this is important for the total vs. effective porosity description of reservoir rocks). Regardless of how that is resolved the presence of iron will increase the neutron porosity and also probably density. Like kaolinite, the smectites are not potassium compounds and so in their purest state they will not contribute to gamma activity.

Glauconite

Glauconite is responsible for some of the largest 'clay effects' on logs, this is ironic as it is not actually a clay. Most mineralogists consider it part of the mica group of minerals, it is indisputably a sheet silicate, however and it affects logs and contributes to the water content in the same fundamental way as the true clays. In particular, it has a strong affinity for water and it includes some hydrogen in its chemical make-up. It is particularly significant in older rocks where the smectites have generally disappeared as a result of diagenesis. Like the smectites its chemical formula is quite variable. It has a higher silicon–aluminium ratio than illite and like the smectites some of the aluminium is replaced by magnesium and/or iron. The presence of potentially large amounts of iron leads to high densities and high neutron porosities. The sheets from which it is formed are separated by potassium ions, so glauconite often contributes to high gamma activity. In some cases the potassium is replaced by sodium or calcium however. One further complication is that glauconite layers can be substituted by smectite layers, this contributes even more variability.

Glauconite can make up a significant part of a sandstone in which case it's high density, neutron porosity and gamma activity can lead to the sand being misidentified as a shale!

Chlorite

Chlorite is another 'non-clay', in fact the 'chlorites' are often given the status of a separate group. From our point of view however, their effects on logs and the ways they are handled are essentially the same as clays and so we will not bother with the distinction. Chlorite is probably best known for its ability to preserve permeability and to trap large volumes of irreducible water. This gives the apparently contradictory situation of a sand with high water saturations that produces dry hydrocarbons at prodigious rates.

Once again the chlorites have quite a variable chemical compositions and like glauconite a lot of the aluminium is normally replaced by magnesium and iron (other metals including zinc, manganese and nickel replace aluminium can also do this. Some commercial nickel ores are chlorites). There are two types of sheet that alternate in chlorite. Both types of sheet include hydrogen in their make-up and so even the dry mineral has a high hydrogen index. Unlike all of the other layer silicates – except kaolinite – discussed above, chlorite does not host cations such as potassium between the layers. This means it generally contributes little gamma activity and is unlikely to significantly reduce electrical resistivity.

2.5.2 Physical Properties of the Clays

Published values for the density and neutron porosity for the commonly encountered clays described above are given in Table 2.1. The salient features of Table 2.1 are the disagreements between different sources and the large

TABLE 2.1 Summary of Published Density and Neutron Porosity Values for Commonly Occurring Clay Minerals

Mineral	Schlumberger				Baker				DHZ		Fert. and Frost	Hill, Soo and Thil	
	Density	PEF	NPHI	NPHIf (epi)	Density	PEF	NPHI	NPHI (epi)	Density	Hardness	Density	Density	NPHI
Quartz	2.64	1.8	−0.02	−0.01	2.64	1.8	−0.02	−0.01	2.65	7			
Kaolinite	2.41	1.8	0.37	0.34	2.59	1.5	0.45	0.48	2.61–2.68	2–2.5	2.61	2.60	0.37
Illite	2.52	3.5	0.3	0.2	2.64	3.5	0.16	0.13	2.6–2.9	1–2	2.64–2.69	2.53–2.65	0.3
Smectites									2–3	1–2	2.2–2.7		
Montmoril-lonite	2.12	2	0.44	0.4	2.06	2	0.13	0.12				2.12–2.53	0.44
Glauconite	2.54	6.4	0.36	0.23	2.58	6.4			2.4–2.95	2			
Chlorite	2.76	6.3	0.52	0.37	2.88	6.3			2.6–3.3	2–3	2.6–2.96	2.7–3.3	0.52

Sources: (a) DHZ: Deer, Howie and Zussmann 'An Introduction to the Rock Forming Minerals'; (b) Fert. and Frost: W.H. Fertl and E. Frost 'Evaluation of Shaly Clastic Reservoir Rocks' in JPT Sept. 1980; (c) Hill, Soo and Thil.: J.A. Hill, D.K.Y. Soo, Thilagavathi Verriah. Geol. Soc. Malaysia Bull. 32 (1992) 15–43.

ranges quoted in some publications. There are two main reasons for the disagreements:

1. Natural variability in the composition of the clays.
2. Definition of what is pure clay.

There is not much that can be done about the natural variability. Aluminium and silicon are sufficiently alike that for some of the clays, almost any ratio gives a similar structure and the resulting charge imbalance can be taken care of by adding or subtracting cations. Moreover, the individual layers from which the pure clays are built can often be mixed to create new compositions. Add to this the fact that aluminium can be replaced by a range of other metals and it is clear why the 'pure' clays show so much variation in their properties.

More can be said about what precisely constitutes a clay. All the clays described above except the smectites, do not include water in their basic composition. All do contain hydrogen, but the hydrogen is part of a hydroxyl group bound to an aluminium atom (Fig. 2.7). As far as most logging tools are concerned this hydrogen is almost indistinguishable from hydrogen in a water molecule. If it is removed from the clay however, a new compound has been formed. It should therefore be possible to define the densities and hydrogen indices quite precisely (the presence of heavy metals in some clays makes the conversion to neutron porosity less straight-forward).

In the case of the smectites, water molecules are generally considered to form part of the clay structure. These are trapped between the individual layers of the clay and so truly are intimately mixed with the clay. On the other hand, they can be driven off by heating: most is removed by heating at 250°C but some is retained even at 500°C when the hydroxyl groups start to break down. In other words it is difficult or impossible to remove all the water without starting to alter the basic clay structure. The difference between the Schlumberger and Baker chart book neutron porosity values probably reflects very different treatments before the measurement (although interestingly they almost agree on density).

Although in principle the other clays should not suffer from this compositional ambiguity the fact is that different sources even disagree about kaolinite's properties. This is probably caused by the same issue that dogs the smectites, that although pure kaolinite does not have water in its structure it is capable of adsorbing a lot of water on its surface. If this is not removed the density will be lower and the neutron porosity will be higher (although note that this does not explain the trends between the two logging companies!).

2.5.3 Petrophysics of Clay and Shale

As we have seen shale and clay volume are different properties and in this book we will be careful to distinguish them.

Shale volume is easier to find from logs because, in principle, the values that characterise a pure shale can simply be read in a suitable bed. In practice
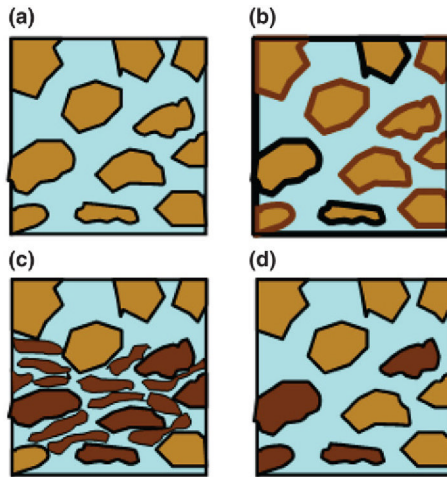
FIGURE 2.9 Cartoons showing the different ways shale (or clay) is present in a shaly-sand (the dimension of the sketches are in the range 0.5–5 mm). (a) Clean sand with little or no clay. (b) Dispersed 'shale', in which clay coats the sand grains. (c) Laminated shale in which clay is located in thin but continuous beds. (d) Structural shale in which some of the grains are composed of shale.

that means one that is thick enough to be resolved and is not suffering from bad-hole. It also should be genetically related to the reservoir rock. In other words, it should have been deposited as part of the same general system. Clay volume will be more difficult to obtain from logs as it is unlikely that a bed of pure clay will be found in the sub-surface. It will therefore be based on published log readings for the pure clay minerals, which as we have seen vary widely depending on the source. Some programmes simply calculate clay volume as a certain fraction of the shale volume (70% say).

The issue of the way in which clay enters a reservoir becomes important when log resolution is taken into account. For example, in some environments shale clasts can be incorporated into the sands, and in others alternating thin beds of sand and shale occur. Log analysts describe these cases as 'structural shale' and 'laminated shale', respectively. In either case the shale is removing effective storage capacity but is not harming the permeability of the sand component. On the other hand, if the clay is intimately associated with the sand it can drastically reduce permeability. This is referred to as 'dispersed shale' (although it ought to be called 'dispersed clay'). In the Morecambe gas fields in NW England the presence of a few per cent by volume of illite is sufficient to reduce permeability by a factor of a thousand.

Three cases are illustrated by sketches in Fig. 2.9. Most logs lack the resolution to distinguish these different cases and will simply show log responses that are intermediate between clean sand and pure shale. In any case the three cases are really end members and all three may be encountered at the same location.

The core shown in Fig. 2.10 is a case in point. It is from a lower shore-face environment and a quick look suggests the sand is very dirty. Closer inspection
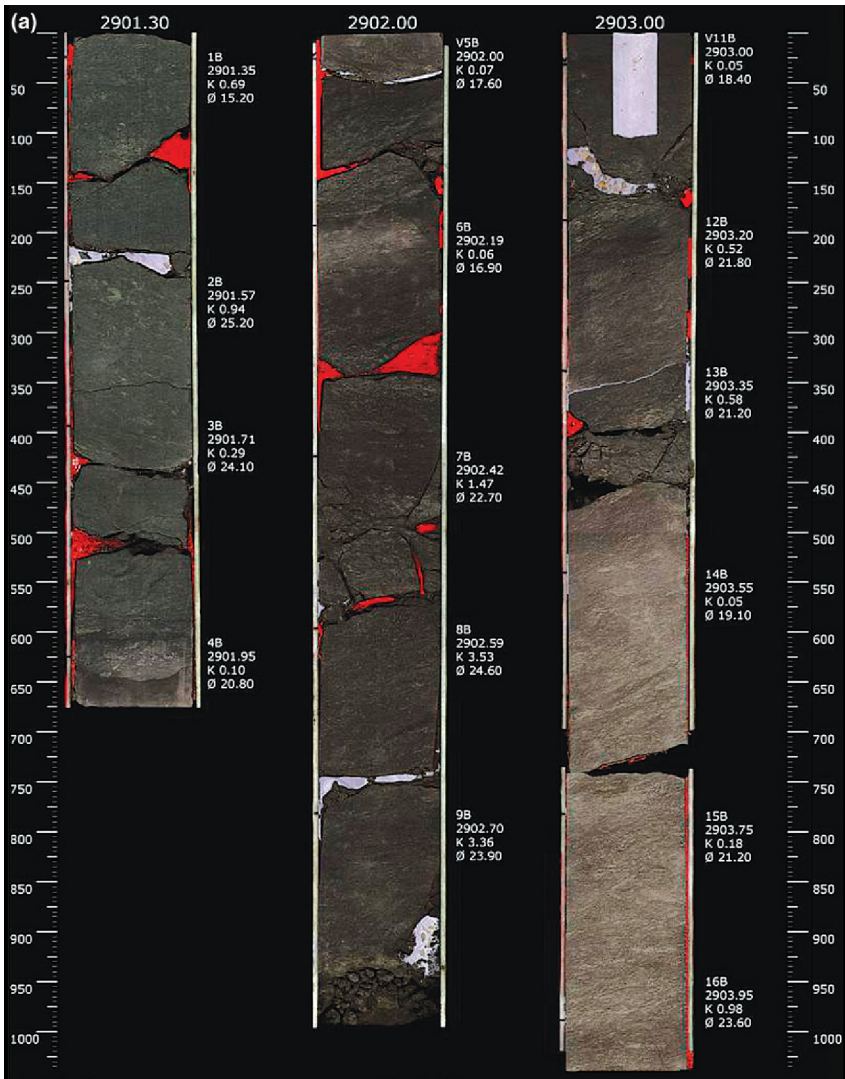
FIGURE 2.10   (a) Photographs of a core from a lower shore-face setting. Note the distribution of the clay in the bottom 0.5 of the core.

shows it has been intensely bioturbated and much of the clay appears to be concentrated in clay-rich laminae (see in particular depths around 2903.5 m). A thin section photograph shows, however that the clay actually fills much of the inter-granular pore space and so is also structural. At the microscopic level kaolinite can be seen to have re-crystallised in the pore space to form a dispersed component. This completely blocks the few remaining pores. The net effect is that most of the porosity is actually present as water associated with clays.
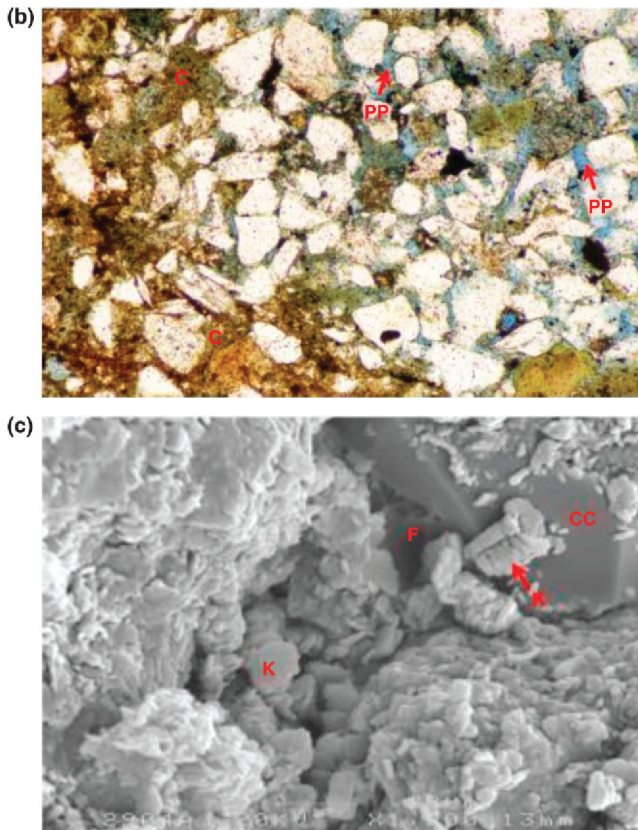
FIGURE 2.10 *(Cont.)* (b) Thin section photomicrograph of part of the core shown in sub-part a. Much of the inter-granular space is filled with detrital clay ('dispersed shale'). The better pore space in the upper RH quadrant is part of a sand-filled burrow. The horizontal dimension of the photograph is approximately 1 mm. (c) SE micrograph of the pore space in Fig. 2.10b. The kaolinite (K) has re-crystallised from the detrital clay and has blocked much of the remaining pore space.

2.5.4 Shale Volume and Clay Volume from Log Analysis

Regardless of what name is being used the end result is a curve which is supposed to represent the variation in shale or clay volume along the well track. Since this is almost certainly a function of one or more logs it follows that the shale volume refers to the volume of investigation of those logs. As we will see in a later chapter these volumes are typically of the order of several hundred litres.

As we have seen it is the clays that cause the most problems for log analysis and so it is fair to say that it is clay volume that we would like to measure. Unfortunately, beds of pure clay are very uncommon in the sub-surface (or at least in conjunction with petroleum-bearing rocks). So we will not find a bed in which we can determine the log readings in pure clay. On the other hand,

shale – or claystone – beds are commonly present and so we can then easily find the log readings corresponding to pure shale. In other words computing the shale volume is relatively easy and in practice this is mostly what is calculated regardless of what it has been named. As a result shale volume was and often still is, used as a proxy for clay volume. This is one reason the two names are so often used interchangeably. Another reason is that many commercial log analysis packages use 'V_{CLAY}' – or something similar – as the default curve name, even if it is actually a shale volume that is being calculated.

Notwithstanding these inconvenient truths, the first step in deterministic log analysis normally is the calculation of shale and/or clay volume. The way this is subsequently used depends on a variety of factors including the nature of the formation, company procedures and whether an effective or total porosity system is being used. At its simplest, the shale volume is simply used as a cut-off to define reservoir quality rock. For example, if it exceeds 50% the porosity is set to zero or at least is excluded from any reservoir averages. In more sophisticated models the shale volume is actually included in the porosity and saturation calculations. For advanced geophysical work such as fluid substitution the shale volume is an essential input into the calculation of the acoustic properties of the rock.

The shale volume is found by using a log, or combination of logs, that gives a good contrast between shale (or claystone) and the reservoir lithology. The gamma ray often fulfils this condition because one or more of the gamma-emitting elements is concentrated in the shales. Regardless of what specific method is used the computation of shale volume involves finding the log reading(s) in the clean reservoir lithology and the pure shale and interpolating between them.

2.6   RELATIONSHIPS BETWEEN PROPERTIES

Petrophysical properties are often correlated and in practice these correlations are regularly exploited to estimate the value of a property at a particular depth or to show that groups of data come from the same population (or not as the case me be). Either way you could say that looking for relationships between properties is largely what this book is about.

Examples of the two reasons for looking for relationships are:

1. To use one property to predict another, e.g. predicting permeability from porosity in an uncored interval.
2. To demonstrate that groups of data come from the same population, e.g. demonstrating that two sands encountered in two different locations share a common source.

There are rigorous statistical measures to establish these relationships and to test their reliability but scientific in-sight, experience and common sense are at least as important. The best place to start looking for a relationship is to construct a cross-plot and we have already met two examples (Figs 1.1 and 2.6). As

a general rule if you cannot see a relationship in a cross-plot, it does not exist. Assuming the cross-plot shows there is an obvious relationship between the two properties we can progress to the next step of deriving an equation, which explains how one property depends on the other.

Before that it is worth listing the reasons that relationships exist in the first place.

Relationships exist between properties because:

1. The properties are controlled by some deeper underlying relationship (e.g. density and sonic slowness of evaporites are both controlled, amongst other things, by the atomic masses of the elements from which they are formed).
2. The properties are both dependent on a third variable (e.g. density and neutron porosity both depend on porosity).
3. 'Self induced'.
4. 'Closure'.

These will be discussed in more detail below, but for now note that the first two causes are what we normally seek, they are a result of the underlying physics in the problem. The other two are created artificially by the way the properties are defined. A trivial case of 'self induction' would be a relationship between a property X and some combination of X and another property (X/Y say).

'Closure' refers to properties which sum to a constant value. This is commonly encountered in quantitative petrophysics where we deal with volume fractions that sum to unity. A trivial example would be water and hydrocarbon saturations. Trivial examples should be easy to spot, but it is not always obvious that closure or self-induced correlations exist.

Most people looking at a cross-plot have a good sense of how good the relationship between the variables is. For the two cross-plots of porosity against permeability shown in Fig. 2.11 there seems to be a closer relationship between porosity and permeability in formation B than A. In both cases, on average, permeability increases with porosity but in formation B the grouping is tighter. The question is can we quantify the 'closeness' of the relationship and if so at which point should we dismiss the relationship as too loose to reliably predict permeability from porosity.

There is no right answer to the latter, it depends on how much potential error we can live with, but there certainly are ways to quantify the closeness of the relationship. The most familiar but often misunderstood way of quantifying 'closeness' is the correlation coefficient (r). (Often the value of r^2 is actually given.)

The correlation coefficient is a property of the data pairs that have been plotted. In fact, there are several different correlation coefficients but they all have the following characteristics:

1. They have values between -1 and $+1$ (or 0 and 1 if it is squared).
2. A value of one indicates that the variables increase in step (in other words the lowest value of x corresponds to the lowest value of y, the second highest value of x corresponds to the second highest y value and so on).
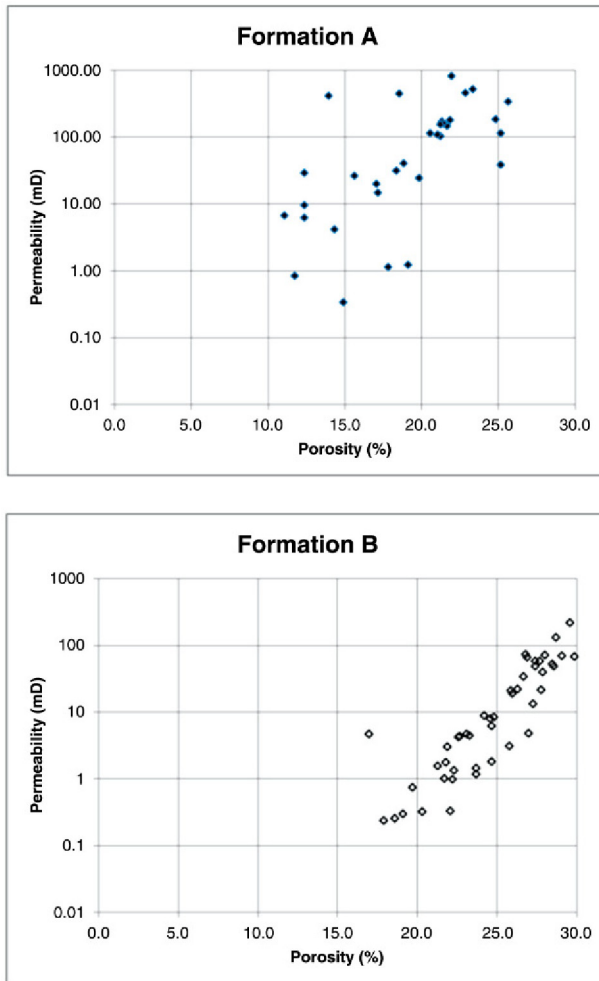
FIGURE 2.11 Porosity–permeability cross-plots for core data from two different sandstone formations. Note the common practice of plotting permeability on a logarithmic scale has been used.

3. A value of −1 indicates the minimum value of x corresponds to the maximum value of y, the second highest x corresponds to the second highest y etc.
4. A value of zero indicates there is no monotonic trend between the pairs of measurements.

This is important because if the relationship between x and y passes through a peak – or a trough – we will get a low correlation coefficient even if there is actually quite a good relationship between the variables. In petrophysics peaks and troughs are quite unusual but it is still advisable not to rely completely on the correlation coefficient.

In this book what we refer to as the 'correlation coefficient' is strictly speaking the 'Pearson product moment coefficient'. This has a specific formulation that imposes some tighter requirements on the limiting values. In particular to get a value of one it is not enough for the variables to increase in step, they must also be linearly related. So a low value of r can arise simply because the variables are not linearly related. If we consider the data sets in Fig. 2.11 the correlation coefficients are 0.33 and 0.70 for formation A and B, respectively. Although these confirm formation B shows the closer relationship between porosity and permeability the correlation coefficients are generally quite low. The reason is that the correlation coefficient quantifies how good a linear fit is, but the cross-plot actually shows permeability on a logarithmic scale. The correlation coefficients for logarithm permeability against porosity are 0.67 and 0.89, respectively. In other words there is a much better relationship between logarithm of permeability and porosity than – un-adulterated – permeability and porosity. In summary if we are going to appeal to the correlation coefficient as proof of a close relationship between two variables we may need to transform one or both of them first. Conversely, a low correlation coefficient does not necessarily imply a lack of a relationship.

The correlation coefficient should not be used blindly to support the existence of a relationship. An example of it's potential misuse is given in Fig. 2.12 the data suggests there is a linear relationship between x and y. The correlation
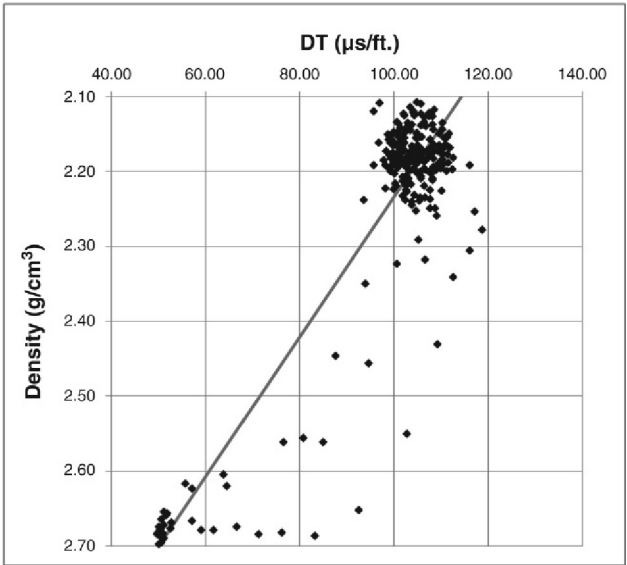
FIGURE 2.12   A cross-plot of density against sonic slowness for a deepwater sandstone. In this example the correlation coefficient is −0.89 suggesting a good relationship between the two properties but this is driven by six points in the lower left quadrant. If these are removed the correlation coefficient falls to −0.4.

coefficient is −0.9, a high value which is often stated as good evidence for there being a relationship (the negative value shows density decreases with increasing sonic slowness). In this case, however the good correlation is due to a few points with unusually high density and low sonic readings. If these are removed the correlation becomes far less convincing. This is easily seen by looking at the cross-plot but will be unknown to any end user of the relationship. The opposite case can also occur where one or two outliers far from the main trend drastically reduce the correlation coefficient (an example is described Section 2.6.2 on 'Closure').

Figure 2.13 shows a less convincing relationship between x and y but the data can still be fitted by an equation, which suggests y increases with x. Once again inspection of the cross-plot would lead to caution in using it however. In this case the relationship is suspect because it produces far less variation in y than is actually observed in the data. This is enough in itself to cause worry to the potential user, but if more evidence is needed – because it is being 'sold'
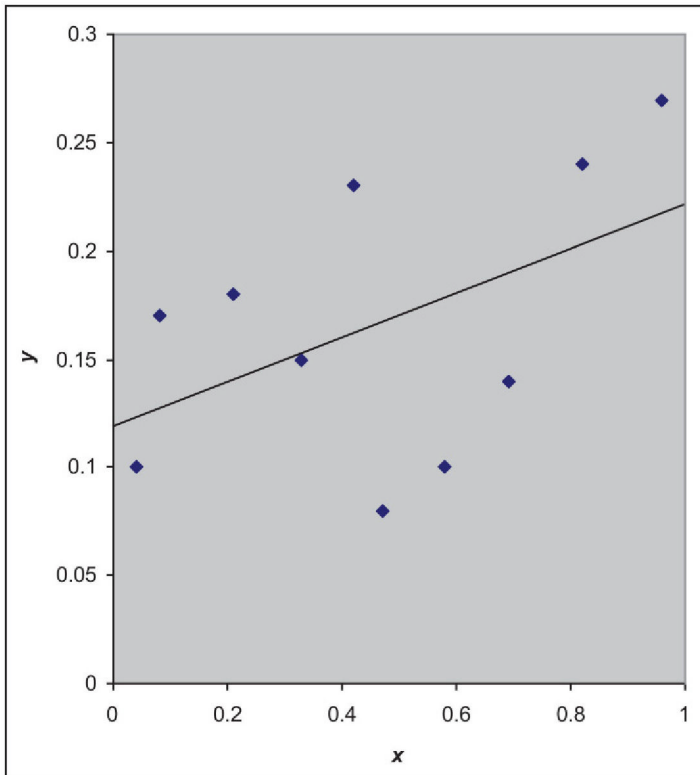
FIGURE 2.13   Regression suggests a relationship between x and y but this is probably purely due to chance ($r = 0.25$). There are only 10 points and the variation in y about it's mean is as large or larger than the range suggested by the regression line.

TABLE 2.2 The Probability that the Correlation Coefficient *R* for a Sample of Size *n* has Occurred by Chance

R	3	5	7	10	20	30
0.1	0.94	0.87	0.83	0.78	0.67	0.6
0.2	0.87	0.75	0.67	0.58	0.40	0.29
0.3	0.81	0.62	0.51	0.40	0.20	0.11
0.4	0.74	0.50	0.37	0.25	0.08	0.03
0.5	0.67	0.39	0.25	0.14	0.02	0.00
0.6	0.59	0.28	0.15	0.07	0.01	0.00
0.7	0.51	0.19	0.08	0.02	0.00	0.00
0.8	0.41	0.10	0.03	0.01	0.00	0.00
0.9	0.29	0.04	0.01	0.00	0.00	0.00

to you – there are statistical tests. Actually the data in Fig. 2.13 are random numbers so there really is no relationship between them!

To properly use correlation coefficients one needs the sample size as well. In the case of Fig. 2.13 the sample size is 10. Table 2.2 shows how the correlation coefficient can be combined with the sample size (*n*) to find the probability that the relationship has arisen entirely by chance. In this case there is an approximately 50% chance that the relationship is entirely due to chance.

The correlation coefficient also allows uncertainty bands to be put on regression lines.

If a relationship

$$Y = mX + c$$

Is found to fit data pairs with a correlation coefficient of *r*, the standard error in the gradient is given by

$$S_m = m\sqrt{[(1 - r^2)/n]} \qquad (2.10)$$

For example, the fit to the data in Fig. 2.11 is:

$$Y = 0.101X + 0.119$$

The correlation coefficient is 0.23 and so S_m is ±0.028. In other words there is 67% chance the gradient lies between 0.073 and 0.129. Which together with the fact that there is only a 50% chance that the relationship arose by chance suggests great caution needs to be exercised using it.

There are a number of ways of fitting a curve to experimental data, the process of doing this is known as regression. The particular method used depends

on the end objective of the fit and the nature of the data that the fit is being made to. Commonly available fits are:

1. *Y* on *X* (or *X* on *Y*). The objective is to reduce the sum of the squared differences between the measured value of *Y* and the value predicted by the curve fit. The method tacitly assumes that the independent variable *X* can be measured precisely. This type of regression is appropriate if *Y* is to be estimated from *X*. This is because the estimate of *Y* is unbiased – all the data points on which the relationship is based contribute equally. But just as the mean is not necessarily the best average to describe a population the *Y* on *X* regression line is not necessarily the best function to map *X* to *Y*.

   The gradient of the regression line is related to the correlation coefficient. It is given by the ratio of the standard deviations in the two variables multiplied by the correlation coefficient (cf. the RMA regression). A corollary of this is that if the correlation coefficient is low the regression line will have a low gradient and so will suggest *Y* varies little.

   (*Y* on *X* uses the same principle but *X* is the dependent variable. This is not the conventional way of plotting data but would be appropriate for finding how a property varies with depth, for example).

2. Weighted Least Squares. This is another method suitable for finding a relationship to predict one property from another. It is used in cases where there is some uncertainty in the independent variable (hence the weighting). It is not used very much in petrophysics but would be expected to be a more robust way of finding *Y(X)*.

3. Reduced Major Axis (RMA). This method is best used for establishing trends in data sets rather than finding a way to predict one property from another. For example, if one wanted to compare two sets of porosity–permeability data from different wells. The RMA fit often gives the 'best looking' fit to the data, this is a result of the way it is defined. Although not strictly correct it may be a better correlation to use to develop a prediction if there is significant uncertainty in the dependent variable. It also has the convenient quality that the gradient of the regression line is given by the ratio of the standard deviations in the two variables. (This is *not* how it is defined, however and actually in *Y* on *X* regression the gradient of the regression line is still proportional to the ratio of the standard deviations.)

Regression should ideally only be attempted on data sets that have been re-scaled to give an approximately normal distribution. The best-known example of this is when permeability is involved it is frequently plotted on a logarithmic scale. Regressions are thus between logarithm of permeability and the other variable.

2.6.1 Self-Induced Correlation

We briefly looked at some reasons for correlation above and noted that ideally we wanted correlations that were the result of some underlying relationship or

for the two properties to be functions of a third property (arguably these amount to the same thing). We also noted that correlations can be caused by the way we choose our variables. These were referred to as:

Closure (the variables sum to a constant amount, often unity or 100%) (Case 4 above)
Self-induced (the variables are both functions of the same property) (Case 3 above)

A special case of self-induced correlation is the so-called ratio effect. It is best introduced using an example. Log data from a 40 m zone of siltstones and mudstones is shown in Fig. 2.14. The properties of interest are density, compressional slowness and thorium concentration and the plot shows the ratio of density/Th as a function of the ratio slowness/Th. There is a very strong correlation between these ratios ($r = 0.91$). What does this mean?

On the face of it we have a way of accurately predicting density from sonic and thorium values (something that one is commonly required to do). The straightforward cross-plot of density against sonic gives a less satisfactory relationship ($r = 0.31$). Figure 2.15 actually shows that both density and sonic logs have a very limited range in the interval of interest.

Sonic	117 (±3)	A coefficient of variation of 0.026
Density	2.41 (±0.04)	A coefficient of variation of 0.017

Thorium on the other hand varies from 16 to 26 ppm (mean 21(±1.5)). So, most of the variations seen in Fig. 2.13 are actually a result of variation in the denominator of the ratios (i.e. thorium). What this means is that we can predict the ratio density/Th very accurately but the value of density at any depth is not going to be predicted with much better accuracy than from the simple density(sonic) regression. In fact, if you were trying to produce a density curve it would probably be better to simply use a constant – average – density reading.

2.6.2 Closure

Much of log analysis relies on the fact that volume fractions sum to unity. All saturation methods, for example rely on independently measuring porosity and the volume fraction of water in the system. The difference is assumed to be made up of hydrocarbon. Another example is the effective porosity equation:

$$\emptyset_e = \emptyset_t - V_{sh}\emptyset_{sh}$$

which shows effective porosity will most likely decrease with increasing shale volume. Amongst other things this shows that: (a) applying porosity and shale cut-offs is often tantamount to applying the same cut-off twice and (b) the common observation that 'shale exerts a strong control on porosity' is simply a result of choosing to use an effective porosity model.

Depth 2660–2700 metres
Filter:
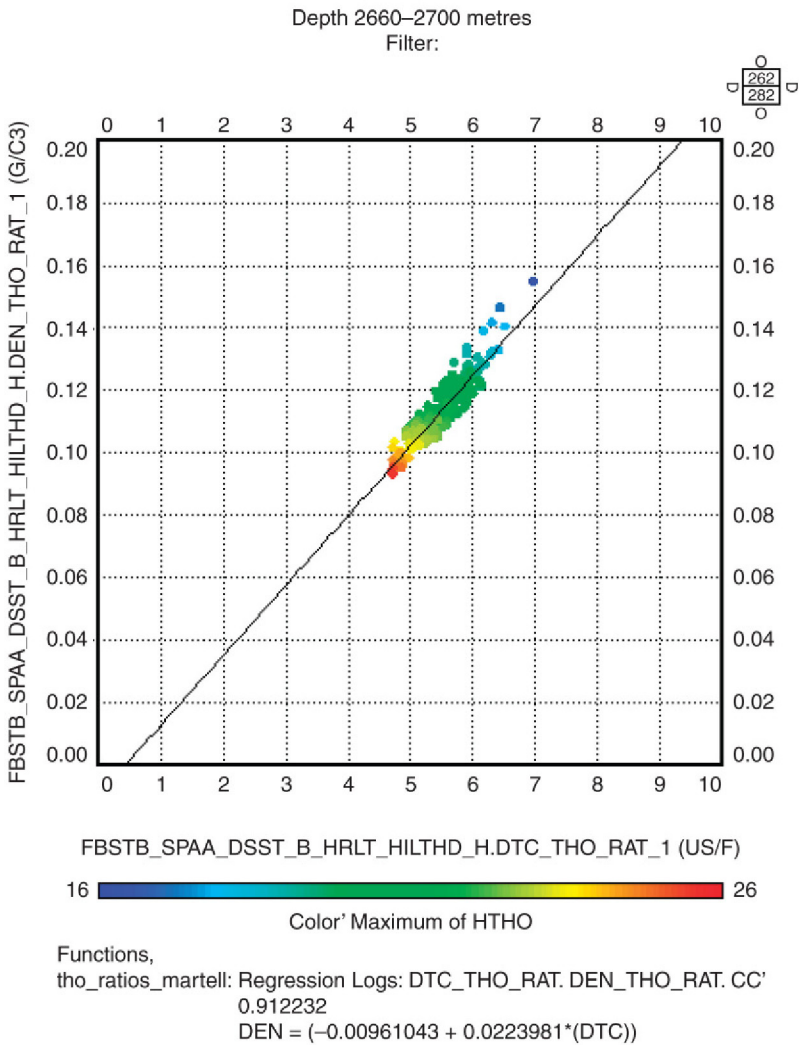
FIGURE 2.14 A cross-plot of the ratio of density to thorium concentration against slowness to thorium concentration. The points are shaded by thorium concentration the highest concentrations occur at the lower end of the plot. The line is a *Y* on *X* regression.

Problems caused by closure were first recognized in geochemistry and at the risk of sounding complacent it is probably more of a problem there than in log analysis.

Quantities expressed as fractions are known as 'compositional data' and geochemists can in principal avoid the issue altogether by presenting their data as dimensional data (masses or concentrations, for example). In practice this is
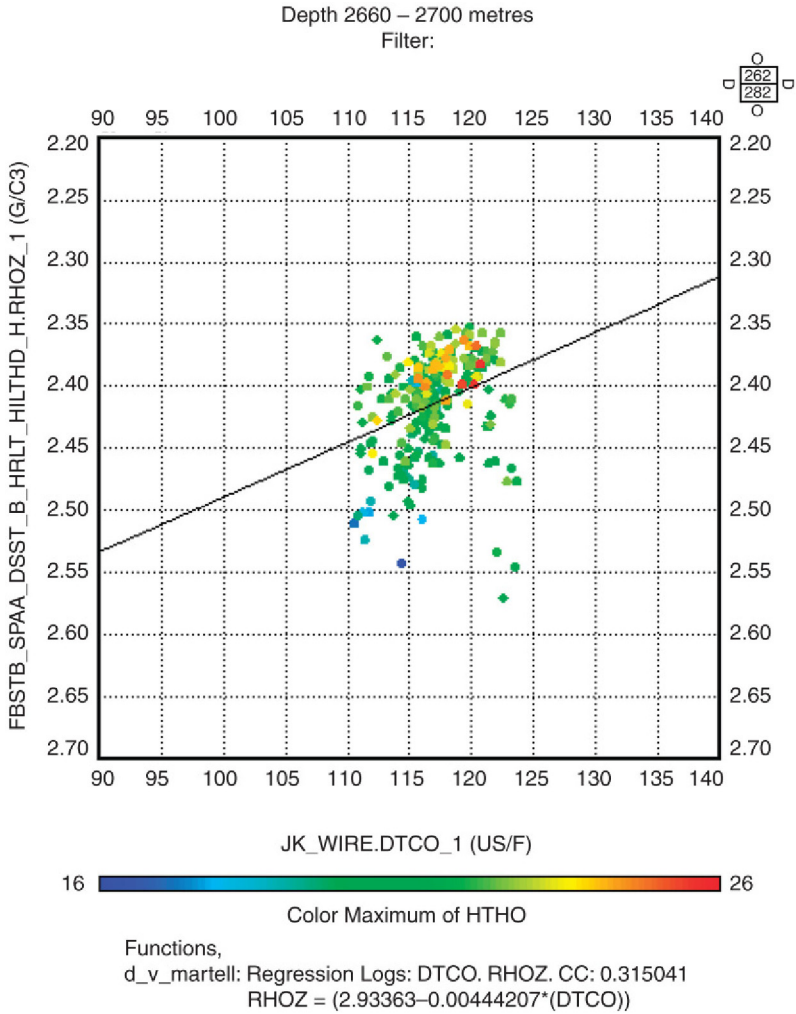
FIGURE 2.15 The same data shown in Fig. 2.14 but plotted simply as density against slowness.

often impractical and for petrophysicists is simply unavoidable because nearly all petrophysical properties are examples of compositional data.

The fundamental problem is the way fractions are defined:

$$X_i = x_i/(x_1 + x_2 + x_3 \ldots x_n) \tag{2.11}$$

where X_i is the fraction of component i and lower case x_1, x_2 etc. are the measured amounts of component 1, 2 etc. Any increase in X_i must be at the cost of a reduction in one or more of the other components. Good correlations are almost

inevitable and they are real, the issue is whether they reveal any underlying association between the components.

Problems normally occur when correlations are produced between the volumetrically most significant component and something else. In Table 2.3 XRD data for some sandstone core plugs is shown quartz is the most significant component for nearly every sample. Kaolinite is in most cases the next most significant component and inspection of the data – or better constructing a cross-plot – shows this tends to increase with decreasing quartz. The correlation coefficient is −0.2 and for a sample size of 24 there is a chance of approximately one-third that the inverse relationship between quartz and kaolinite has arisen by chance.

As an aside, note that the complete data set, with a correlation coefficient of −0.2, is an example of a low correlation coefficient produced by a few outliers. Inspection of the data shows two outliers with unusually low quartz fractions of 25% and 29% (by weight). If these are removed the correlation coefficient increases to 0.7 (the probability that this is produced by chance is very low).

There is no question that there is a good relationship between the fraction of kaolinite and quartz, the question is whether it is useful or whether it simply arises for the same reason that there is a perfect correlation between the functions 'x' and '$1 - x$'.

In classical log analysis only a few properties are calculated, the most significant component of which is the so-called matrix (e.g. 'sand' or 'limestone'). But even the less significant components: shale volume, porosity, hydrocarbon volume etc. are likely to be greater than 5–10%. Correlations between them will thus always have a significant closure contribution. With measurements like NMR and neutron-based geochemical techniques a much larger number of components are provided and there is a greater danger of closure effects being mistaken for something of deeper significance.

As already noted geochemists have been aware of the problem for many years. They normally deal with a much larger number of components than petrophysicists and so closure-based correlations are often difficult to spot (although as noted above some of the more recent logging measurements may produce a similar number of components). Some robust techniques have been developed to try and unravel the relationships between components. These techniques involve a lot of computation and even with the help of a computer can be quite labour-intensive, they are not discussed further here.

2.6.3 How the Correlation Coefficient is Calculated

The correlation coefficient is a function of a more fundamental statistic: the covariance. This is similar to the variance of a single measurement and is – approximately – the mean of the product of the variations of the two measurements from the mean. It is simpler to write this as a formula:

$$COV = \frac{\sum (x - X)(y - Y)}{(n - 1)} \tag{2.12}$$

TABLE 2.3 Petrogrpahic Data for the Sandstone Discussed in the Text

Depth	Qtz.	KF	I/M	I/S	Kaolinite	Calcite	Sid.	Ank.	Ap.	Ana.	Pyrite	Marcasite
3700.8	57.2	14.0	9.7	0.8	16.0		0.8			1.1	0.4	
3705.97	81.4	7.5	5.2		5.3					0.6		
3712.3	79.5	14.2	3.7		2.0	0.1	0.3					
3714.83	87.6	8.2	2.3		1.5	0.1	0.3					
3724.2	25.0		1.1	20.6	2.0		50.0				1.3	
3725.9	89.0				3.6		3.0			0.6	3.2	
3726.8	80.2		2.9	3.6	8.8		1.3			0.9	2.3	
3736.01	93.9				4.1		0.7			0.5	0.2	
3740.23	92.9				5.8					0.4	0.4	
3742.57	89.4				6.8			2.6		0.6	0.2	
3744.35	88.1				7.3		0.5	2.1		0.6	0.8	
3746.6	81.6				6.2		4.6	2.1	3.7	1.0	0.2	

3750.87	86.7				11.2		0.1	0.6		1.0		
3753.05	85.5				12.4		0.1	0.2		0.7	0.8	
3773.1	30.1		3.6	0.7	14.7		46.5				4.4	
3774.48	94.4				3.5		1.9					
3775.77	29.1				1.1		35.9				12.3	20.8
3782.6	96.4		0.3		1.1						1.9	
3789.3	95.9		0.5		3.1						0.3	
3792.15	94.8		0.1		1.9						3.2	
3796.9	88.7		0.2		4.1						6.7	
3804.4	87.1		0.8		10.0		0.7				0.9	
3808.1	93.6				5.9							
3812.5	96.4		0.5		2.8							
Mean	80.19	10.98	2.38	6.43	5.88	0.10	9.78	1.52		0.73	2.32	
SD	21.75	3.62	2.75	9.55	4.30	0.00	18.04	1.05		0.23	3.14	

where X, Y are the means of the samples of x and y and n is the sample size. Note the similarity to the variance of a single property.

$$S = \sigma_x^2 = \frac{\sum (x - X)^2}{(n-1)} \tag{2.13}$$

It is easy to see how the covariance indicates the quality of a correlation. If x and y generally increase together then so will $x - X$ and $y - Y$ and the covariance will be large and positive. If there is not much of an association then $x - X$ will often be large when $y - Y$ is small or even negative and vice versa. The covariance will then be relatively small.

The correlation coefficient is given by

$$r = \text{COV}/\sigma_x \sigma_y \tag{2.14}$$

A magnitude of one indicates 'perfect' correlation and a value of zero indicates no correlation. Notice that because r is defined in terms of standard deviation it is susceptible to outliers. Statisticians would say it is not robust.

2.7 HETEROGENEITY AND ANISOTROPY

So far all the properties we have discussed have had quite specific definitions. For example, porosity is the volume fraction of the rock occupied by fluids. Granted we have to be careful in explaining exactly what constitutes 'fluid' but porosity is always going to be a ratio of some fluid volume to the bulk volume of the rock. In this section we discuss two attributes that can be applied to anything that can be measured, calculated or even just described. If you like, they are properties of properties.

Put simply, heterogeneity describes how much a property varies from place to place and anisotropy refers to properties that depend on direction. Both these or strongly dependent on the scale we are working at and so this section is also a good place to introduce the important concept of 'volume of investigation'.

2.7.1 Anisotropy

Mathematicians and physicists classify properties as scalar, vector or tensor. Scalar properties have no directional dependence, they include mass, volume, temperature, heat, charge as well combinations of these such as density or specific heat capacity. These are all properties that can be described by a single number. In petrophysics porosity, saturation and shale volume are scalar properties. Vector and tensor properties on the other hand depend on the direction they are being measured, examples are electrical resistivity and the speed of sound. Permeability is the most significant directional petrophysical property. To properly describe any of these, at the very least we need to specify how big it is in the particular direction we are interested in (which makes it a vector). To

give a complete description we need a matrix that allows the value to be calculated in any direction (and that is the tensor).

Although three out of the four properties introduced in the first half of this chapter are scalars we often calculate them from measurements that are actually vectors. So for example, Archie's equation, that was discussed in some detail in Chapter 1, uses resistivity (a vector) to estimate saturation (a scalar). Clearly if we are using a measurement that depends on the direction it was made, we will end up with the non-sensical conclusion that saturation too depends on the direction. The true implications of this particular case will be discussed when we look at log analysis later in the book.

So far we have discussed measurements and properties that are inherently scalar or directional. But in fact as long as the rock has no preferred direction or fabric any measurement can be treated as a scalar and the rock is said to be isotropic. Thick massive sandstones and carbonates are good examples and since these often form reservoirs it is possible to ignore anisotropy (Fig. 2.16a). There are plenty of cases where it cannot or should not be ignored, however. Rocks that are strongly layered are a particularly common case of anisotropy complicating reservoir description (the so-called thin bed pays are a case in point). This type of anisotropy is intermediate between the general case where different values will be obtained in each of three orthogonal directions and the isotropic – same in any direction – case. Essentially the measurements will differ parallel and perpendicular to bedding: the particular direction chosen in the bedding plane makes no difference. Since bedding plains are often horizontal measurements in the two directions are distinguished by suffixes 'h' and 'v' (Fig. 2.16b). In many cases this issue was only recognised when high angle wells became more commonplace.

As an aside by convention cylindrical core plugs are normally cut parallel to bedding. That is to say the axis of the cylinder is parallel to the bedding plains and as permeability is measured along that axis it is invariably given the symbol

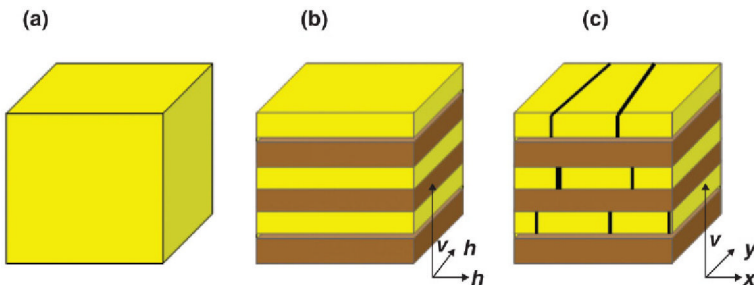**FIGURE 2.16** The various levels of anisotropy. (a) Isotropic rock: any property has the same value regardless of the direction it is measured. (b) Layered fabric in which directional properties depend on whether they are measured perpendicular (v) or parallel (h) to bedding. (c) Layered fabric with superimposed joints: directional properties now differ in vertical (v) and both horizontal directions (x or y).

'*kh*'. A few plugs are normally cut at 90° to the bedding and the permeabilities measured on these 'vertical' plugs are given the symbol '*kv*'.

For the specific case of permeability the ratio *kv/kh* (pronounced '*k-v-k-h*') is an important parameter for characterising a reservoir unit. It determines the preferred direction for flow and pressure support. In general it varies between zero and one, in other words the vertical permeability can be equal to but rarely exceeds the horizontal value and more often than not, it is a lot less. This however is a scale-dependent property and if it is being used to model reservoir performance it is important to be clear at what scale we are describing the reservoir. Core plugs are often quite homogenous and horizontal and vertical plugs cut at more or less the same depth give similar values. So at the core plug scale *kv/kh* is normally close to unity and may be greater than one. At the scale of individual sand packages, however it is normally significantly lower than unity because thin shales inhibit vertical flow (Fig. 2.16).

2.7.2   Heterogeneity

Heterogeneity describes how much and how quickly properties vary from place to place. In fact we are specifically interested in variation in three dimensions because, after all, layered fabrics like those discussed in the previous section produce a lot of variation, but only in the direction perpendicular to bedding. When we say a formation is heterogeneous we are implicitly assuming that we see a lot of variation in any direction and the fabric of the rock is really quite chaotic.

Heterogeneity is of particular interest to petrophysics because it causes the value of a measurement to be strongly dependent on precisely where it is made. As a corollary we would expect a large amount of variation between measurements and therefore a large number of measurements are needed to confidently characterise the formation. In the specific case of logs, the different volumes of investigation of each tool means they typically respond to a different mixture of minerals and fluids and therefore appear to be inconsistent. Ultimately, this causes disagreements between different ways of calculating the same property. It is this which is largely responsible for a widespread belief that petrophysical interpretation of carbonates is particularly difficult. The problem is not so much due to the carbonates but rather the porosity systems they sometimes host, which can be highly heterogeneous.

All rocks are heterogeneous at some scale, but we are most interested when they vary significantly at the scale we are making our measurements. In the previous chapter, it was stated that petrophysics typically deals with volumes that vary from cubic centimetres to cubic metres and it is variation at this type of scale that concerns us in this book. Specific examples are conglomerates, breccias and vuggy carbonates. It is helpful to consider some examples. Figure 2.17 shows two cores from wells in the Canning Basin. On the left the core consists almost entirely of fine-grained sand, the only visible variation is the oil staining in the middle box (Fig. 2.17a). This is a very homogenous formation (it is also probably isotropic).
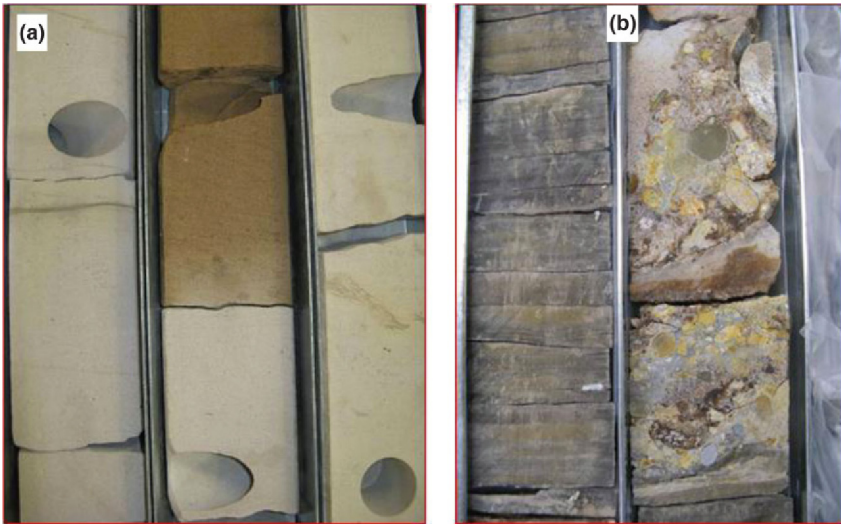
FIGURE 2.17 Slabbed ores from the Canning Basin, Western Australia. The width of the slabbed surfaces is approximately 9 cm. (a) Massive homogeneous fine-grained sand. (b) Claystone (left) and overlying conglomerate (right).

The right hand example includes conglomerate with individual clasts clearly visible in the photograph – they have dimensions of the order of a centimetre.

Recognising the conglomerate core is more heterogeneous than the sandstone implicitly assumes that the scale of the measurements (or samples) is small compared to the width of the core. Samples of the order of 1 cm could variously be all pebble, all matrix or some mixture of the two. But if the samples were approximately $10 \, cm^2$ they would mostly consist of several clasts with their surrounding matrix. There would be less variation between different samples. Think of this as replacing Fig. 2.17b with a very low resolution photograph.

Because the majority of petrophysical data comes from wells, we can generally say more about heterogeneity in the 'along hole' direction than in the plane at 90° to the hole direction. But as discussed above variation in the along hole direction may simply be due to bedding (see the previous section). Relying on wells to provide information also constrains the number of measurements we can realistically make so that we may never be able to confidently characterise a very heterogeneous formation. In fact apart from cores and image logs there are few direct ways to look at heterogeneity in wells but it is sometimes possible to argue that a formation is heterogeneous simply because different measurements are apparently incompatible.

The conglomerate shown above is obviously more heterogeneous than the sandstone but it would be nice to devise an objective way of quantifying how much more. One approach is to take several samples at random locations, measure some property on each one and use statistics to quantify the amount of variation. A commonly used statistic is the 'coefficient of variation', which is

basically the ratio of some measure of variation to the average value of the sample. Most commonly it is defined as:

$$Cv = \frac{Standard\ deviation}{Mean} \tag{2.15}$$

Any quantitative property could be used but in practice it is often permeability that is chosen because variation in that has the most drastic influence on reservoir performance.

2.7.3 The Lorenz Coefficient

Some of the best measures look at relationships between two different properties measured at various places. One of the best known is the Lorenz coefficient, which is calculated from continuous measurements or estimates of porosity and permeability. The Lorenz coefficient essentially quantifies how well porosity and permeability are correlated. It varies between one for a perfectly heterogeneous rock and zero for a perfectly homogeneous one. The easiest way to appreciate how it works is through a worked example. Figure 2.18 shows a porosity–permeability cross-plot of some core data from a sandstone reservoir. In this case porosity covers quite a large range (5–17%) so the coefficient of variation is quite high and on that basis this is a heterogeneous reservoir.

The first 10 measurements are reproduced in Table 2.4a the cross-plot. There are actually 89 individual measurements and in reality all would be used

TABLE 2.4A The First – Shallowest – 10 Core Measurements from the Data Set Shown in Fig. 2.18

| Sample number | Sample depth (m) | Confining stress 4370 psi | | | Grain density (g/cm³) |
| | | Permeability | | Porosity (%) | |
		K_{inf} (md)	K_{air} (md)		
1	2750.08	<0.001	0.002	3.9	2.645
2	2750.37	0.005	0.013	5.0	2.646
3	2750.68	0.181	0.269	7.6	2.647
4	2750.87	0.112	0.171	9.0	2.642
5	2751.17	0.060	0.072	6.2	2.648
6	2751.47	0.002	0.005	4.0	2.645
7	2751.76	0.162	0.219	8.0	2.644
8	2752.06	4.85	5.99	9.0	2.646
9	2752.36	0.026	0.031	5.7	2.652
10	2752.66	9.89	11.5	10.3	2.670
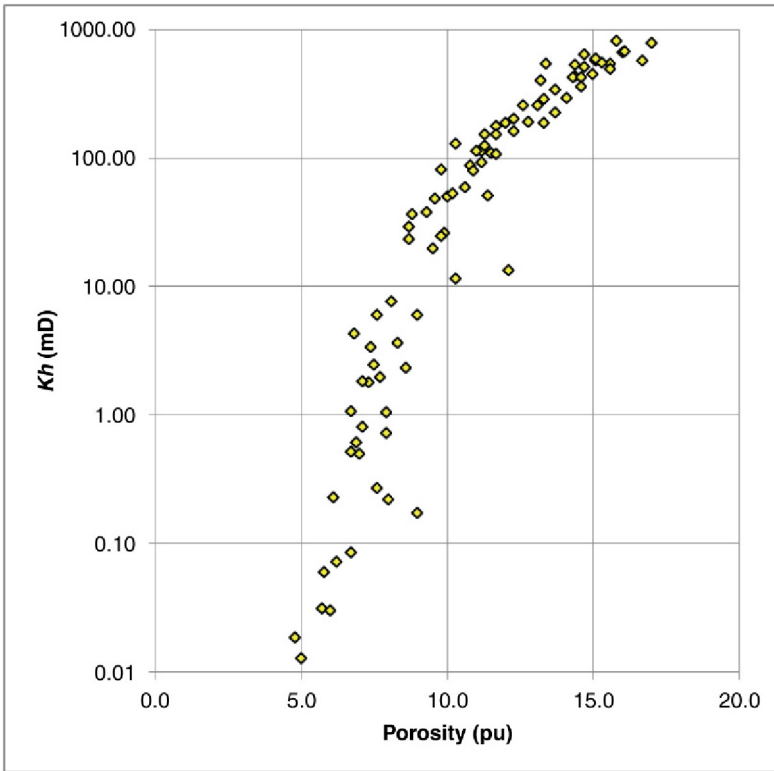
FIGURE 2.18 Porosity–permeability measurements from a sandstone reservoir used in the example of the Lorenz coefficient calculation.

to find the Lorenz coefficient, but to save space we will show the working with just the first 10. The same measurements are shown in Table 2.4b but now they have been re-arranged to be in order of ascending air permeability (the grain densities have been discarded since they are not used). Two additional columns have been added which give the cumulative porosity and permeability.

So the highest permeability is 11.5 mD and this sets the bottom row. The next highest permeability is 5.99 mD so the cumulative permeability is 17.5 mD. The third highest is 0.27 mD and so the cumulative permeability for the third row up is 17.8 mD and so on. The cumulative permeability is conventionally named the 'flow capacity'. At the same time the cumulative porosity is calculated in each row. Conventionally this is known as the 'storage capacity'.

The final step is to plot each flow capacity value against the corresponding storage capacity value. This has been done for the complete set of measurements in Fig. 2.19 (i.e. all 89).

The Lorenz coefficient is then defined by the ratio of the area between the smooth curve and the diagonal to the area under the diagonal. In this case

TABLE 2.4B The Data Shown in Table 2.4a, Re-arranged in Order of Increasing Permeability

Sample number	Sample depth (m)	Confining stress 4370 psi			Cumulative	
		Permeability		Porosity	Porosity	Permeability
		K_{inf} (md)	K_{air} (md)	(%)	Storage	Flow capacity
1	2750.08	<0.001	0.002	3.9	68.7	18.3
6	2751.47	0.002	0.005	4.0	64.8	18.3
2	2750.37	0.005	0.013	5.0	60.8	18.3
9	2752.36	0.026	0.031	5.7	55.8	18.3
5	2751.17	0.060	0.072	6.2	50.1	18.2
4	2750.87	0.112	0.171	9.0	43.9	18.1
7	2751.76	0.162	0.219	8.0	34.9	18.0
3	2750.68	0.181	0.269	7.6	26.9	17.8
8	2752.06	4.85	5.99	9.0	19.3	17.5
10	2752.66	9.89	11.5	10.3	10.3	11.5

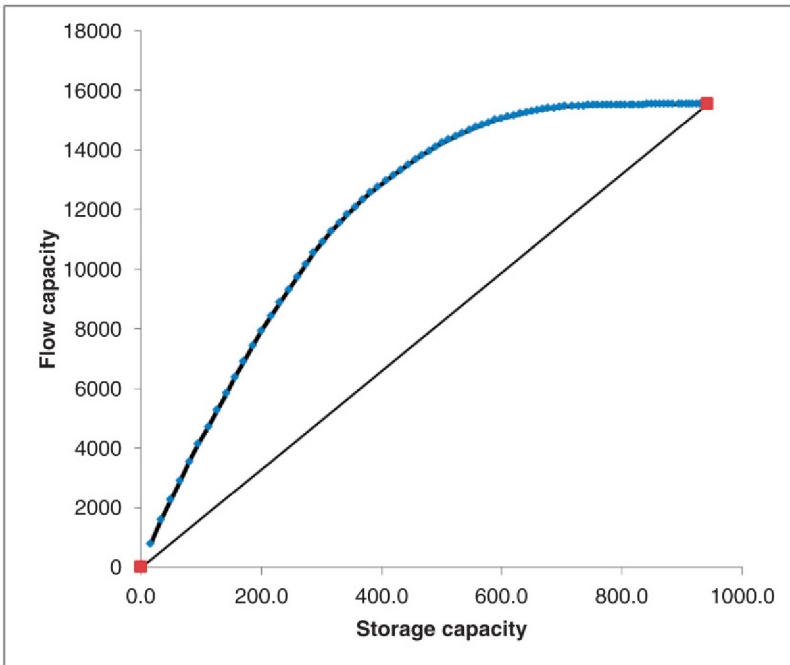The two right-hand columns show porosity and permeability accumulating upwards.

FIGURE 2.19 Lorenz plot for the core data shown in Fig. 2.18 (89 points in total).

the answer turns out to be 3.9 million/7.3 million or 0.53. This suggests a reservoir of intermediate homogeneity.

There is no fundamental reason for this definition, it is just the way the coefficient was defined when it was invented. By making it a ratio of areas it becomes insensitive to the number of points used to construct it. For example, if we had doubled the number of core plugs the total storage capacity would have doubled to approximately 2,000 but the flow capacity would also have doubled to about 31,000 and the shape of the curve would remain the same. A completely heterogeneous reservoir would be the one in which all the flow capacity was located at one point. In that case the Lorenz curve would rise almost vertically from the origin and then track horizontally to the maximum storage capacity. The area between the Lorenz curve and the diagonal would be the same as the area under the diagonal and the Lorenz coefficient would be one. The ideal homogeneous reservoir is one in which the permeability is directly proportional to porosity.

In reality homogeneous reservoirs tend to have a limited porosity range and permeability range and so the two cumulative properties often do increase approximately linearly. The Lorenz coefficient and other similar properties are most useful when used to compare reservoirs rather than as an absolute number.

The Lorenz coefficient has been around for a long time, it was originally designed to help reservoir engineers predict water flood efficiency. It was reasoned that the ideal piston like displacement was favoured by homogeneous reservoirs. Conversely, water was more likely to find high permeability conduits in a heterogeneous reservoir. With the advent of 3D simulators such screening tools fell out of favour but they are now experiencing a resurgence as they can help explain scale effects (when comparing well test and core permeabilities, for example).

2.8   NET, PAY AND AVERAGING

The most useful output of a petrophysical interpretation is a series of curves that show how porosity, water saturation, permeability and other properties vary along the well path. These property curves will nearly always be derived from the logs and so they will consist of values given at regular depth intervals (e.g. every 15.24 cm or 6 in. in imperial units). This can amount to a lot of information and for various reasons there is often a requirement to distil these down to a single number that characterises the reservoir, zone or well of interest.

This is always a contentious issue, not least because there is no agreed nomenclature let alone methodology. In terms of arithmetic it is easy enough to reduce a column of numbers to a single value. The arithmetic mean will give an unbiased estimate but there may be good reasons for using some other average. The real problem is, however that we want numbers that in some sense convey the essence of the zone but at the same time can be applied consistently to different reservoirs, possibly from very different geological environments and parts of the world. For example, in one prospect we may have discovered several thin, high porosity sandstones in a background of shale and in another a thick, low porosity

sandstone with a few thin shales embedded in it. The average porosity may be the same in each case, but they are clearly completely different reservoirs. We need a way to compare them and moreover a way that conveys their relative value.

The way this is achieved in practice is more often than not, by means of cut-offs. That is to say we restrict attention to rock that passes certain simple criteria. An example is 'porosity greater than 10% and shale volume less than 50%'. Of course, the skill is picking the cut-offs in the first place: that is to select the property(s) to use and its (their) limiting values.

The output from this process is, for the interval of interest:

1. A ratio that describe how much of the interval is judged to be reservoir quality (or equivalently the thickness represented by the reservoir quality material).
2. The average value of each property of interest for the reservoir quality rock.

In this book we will use the following definitions:

Gross   Total thickness of the interval.
Net   Thickness of reservoir quality rock. That is rock that is capable of contributing to reservoir energy. It need not contain hydrocarbons.
Pay   Thickness of reservoir quality rock that contains moveable hydrocarbons. Rock that is Pay is necessarily also Net by our definition.
Net-to-Gross (NTG)   The ratio of Net to Gross.
Pay-to-Gross (PTG)   The ratio of Pay to Gross.
Net Average Porosity   The average value of porosity for the Net. (An equivalent Pay average porosity can also be defined.) Normally the average used is the arithmetic mean.
Net Average Water Saturation   The volumetric average water saturation for the net. (As with porosity a Pay averaged S_w can be defined.)

It is important to note these definitions are by no means universally accepted and can mean different things in different organisations. The only way to avoid confusion is to provide a clear statement of what the terms mean every time they are used.

Typical cut-offs that are used are a lower limit on porosity and/or an upper limit on shale volume for Net and these together with an upper limit on water saturation for Pay. Variations on the theme include using a continuous permeability curve either on its own or as an additional cut-off and other curves to remove artefacts caused by coals or bad-hole (which could be mistaken for high porosity gas-bearing sands).

In terms of arithmetic the process is simple and is best illustrated by an example. This is done in Fig. 2.20. But again it should be stressed that the difficult problem is exactly what we mean by reservoir quality rock or equivalently 'net' and having decided that, how we select cut-offs to identify it. Here we have defined reservoir quality rock as rock, which is capable of contributing to production that means it must be permeable. This is so important that we will write it out again.
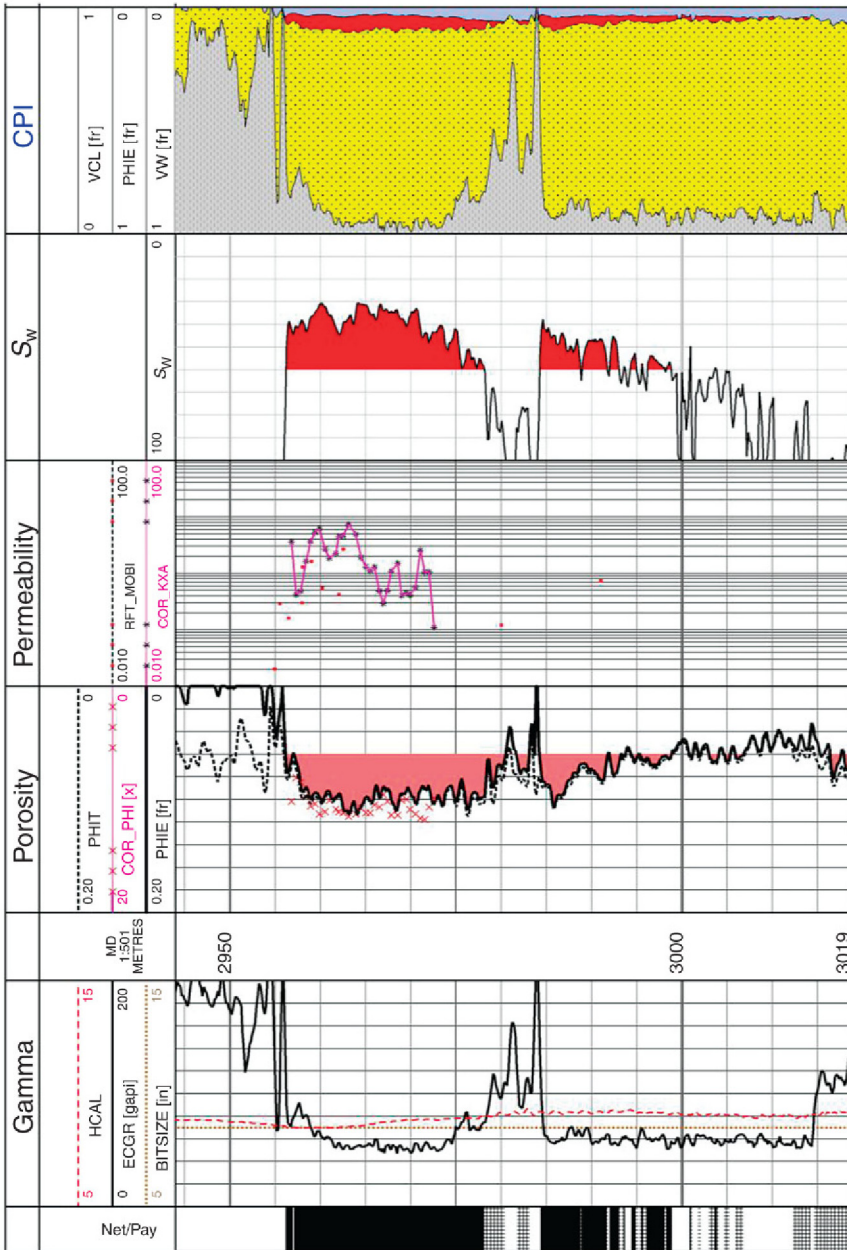
FIGURE 2.20 A log interpretation of an oil-bearing sand showing how porosity, shale volume and saturation cut-offs are used to define Net and Pay. Pay (and Net) rock are shown by the shading in the LH track.

Definition: Net rock is rock that is capable of contributing to production.

It is worth repeating that this is just the definition that we will use in this book, other sources will use different words in place of 'net' or use 'net' to mean something different.

It is important to realise that the fluids it contains need never be produced. They can always contribute to production through a 'knock-on effect' in which they move and in the process push other fluids towards the wells. In the limit they may contribute to production simply by expanding into the surrounding pore space as the reservoir pressure drops. In other words the net contributes to reservoir energy.

Unfortunately, it is not just the words that change between sources. Words after all are just labels and we can always produce a table showing how our definitions relate to those in use by company X. There is a more fundamental problem that company X may have a different understanding of reservoir quality rock, this goes to the heart of picking cut-offs and ensuing arguments. Over the years there have been arguments for 'Net' having to contribute flow into a well and even to contributing an economic flow-rate (life is too short to discuss the latter). The problem with relying on in-flow is that it only tells you what is happening very near the well bore where conditions are extreme compared to most of the reservoir (high flow-rates, changes in the nature of the reservoir fluids and possibly formation damage). It also requires that an in-flow profile has been recorded, more often than not this is not the case.

We will conclude this section with an example. Figure 2.20 shows an interpretation of some logs that were recorded in a sandstone reservoir. There is some core data from the top of the sand this, has also been plotted and is shown as discrete points. The calculated porosity is a bit lower than the core values this is partly because the core has not been compaction corrected. Porosities vary from under 5% to just over 10% and water saturation decreases from 100% at the base of the sand to about 35% at the top of the sand. Most of the interval of interest is sand but there is a shaly interval in the middle of the reservoir. Net and Pay are shown by the shading in the LH track. Net is defined by porosity and shale volume cut-offs, specifically:

$$\text{Porosity} > 6\%, \text{shale volume} < 50\%$$

Pay has to pass the additional cut-off that $S_w < 60\%$ and is shown by the solid black coding. Rock that is Net but not Pay has been shown with cross-hatch coding. To emphasise these cut-offs the porosity curve has been shaded whenever it exceeds the cut-off value and similarly, the S_w curve is shaded whenever it falls below 60% (the shale volume cut-off is not explicitly highlighted but the reservoir has so little shale that most of it passes this cut-off).

That is all quite straightforward, the problem, repeatedly stated, is how the cut-offs were chosen. They probably are arbitrary and are values used by the operator in this type of play, that is quite acceptable for comparing one well with another.

TABLE 2.5 Reservoir Average Properties Calculated with a Different Porosity Cut-Off

Top	mMDRT			2956		
Base	mMDRT			3013.9		
Gross	m			57.9		
Net	m	55.0	54.4	41.9	25.0	7.5
NTG		0.9	0.9	0.7	0.4	0.1
Net av. poro.	v/v	0.077	0.077	0.085	0.096	0.105
Net av. S_w	v/v	0.54	0.53	0.48	0.42	0.39
Pay	m	34.1	34.1	33.8	24.4	7.5
Pay av. poro.	v/v	0.089	0.089	0.089	0.096	0.105
Pay av. S_w	v/v	0.44	0.44	0.43	0.41	0.39
HCPV	m	1.7	1.7	1.7	1.4	0.5
Poro. cut-off		0.02	0.04	0.06	0.08	0.1

A table of reservoir average properties using different porosity cut-offs is shown in Table 2.5. Obviously, as the porosity cut-off is reduced the Net thickness increases but the reservoir average porosity decreases. The same applies to Pay but as this is also defined by an S_w cut-off a point is reached where no matter how low the porosity cut-off the Pay thickness remains the same.

The increase in Net – or Pay – thickness as the cut-offs are relaxed is to an extent offset by the increase in average porosity and decrease in average S_w. This can be seen with the property HCPV in the lowermost row of Table 2.5. HCPV stands for 'hydrocarbon pay volume' although it is actually a thickness. It is calculated from:

$$\text{HCPV} = \text{Pay.Pay average porosity}(1 - \text{Pay average } S_w) \qquad (2.16)$$

Physically this is the interval in metres that would be occupied by gas if the sandstone was compacted so that all the fluid was expelled into the resulting head space. Inspecting Table 2.5 shows that in this case the HCPV remains constant until the porosity cut-off exceeded 6% and that may be why that value was chosen. Of course this does not really say anything about the rocks ability to produce or at least provide reservoir energy but it is a useful device for showing how sensitive the averages to the choice of cut-offs.

Since in this book we have defined Net and Pay in terms of ability to produce, permeability is actually going to be an important input although we have said nothing about that so far. Figure 2.21 shows a porosity–permeability cross-plot for the core data from the reservoir. A semi-log relationship between them has been drawn onto the plot. A common way to define Net is actually in
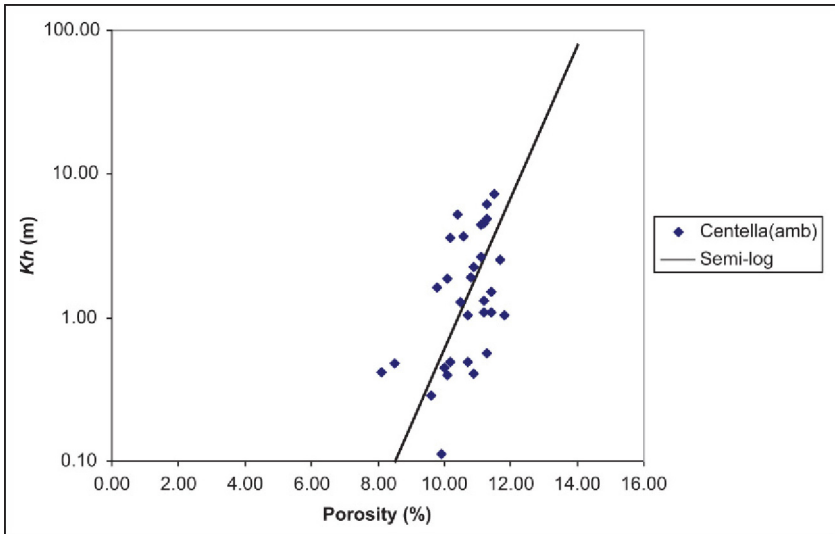
FIGURE 2.21 Porosity–permeability cross-plot for the core data obtained from the reservoir shown in Fig. 2.18. The data was measured at laboratory conditions.

terms of a permeability cut-off. Common values are 0.1 mD for gas and 1.0 mD for oil. The porosity–permeability cross-plot can then be used to convert these to porosity cut-offs. In this case the gas cut-off would be approximately 8.5%. Note however that these numbers are still arbitrary all that has been achieved is that a 'recipe' has been written for deriving a porosity cut-off.

There is plenty of scope for argument using such a recipe. Common objections that are raised include the following:

1. The core data refers to laboratory conditions but the production occurs at reservoir conditions.
2. Permeability was measured using an inert gas on dried cores, in reality production is of hydrocarbons from rock that contains a certain amount of water.
3. The origin of the function relating permeability to porosity.

The first two points can be accounted for by more elaborate – and expensive – measurements, but actually they take attention away from the perennial issue that 1 mD – or whatever number was chosen – is still an arbitrary number.

If permeability can be measured, fluid can flow out of the rock and so perhaps we should not have cut-offs at all. Some operators do indeed use very relaxed cut-offs that do nothing more than remove the obvious non-reservoir rock (e.g. shale beds). This is fine so long as the reservoir has a binary nature: that is to say it is nearly all reservoir rock with a relatively high permeability or non-reservoir with a permeability several orders of magnitude lower. When reservoirs have a continuous range of permeabilities, where a large amount of

material has intermediate values, there is a real danger of creating significant reserves that have no hope of being recovered in the life of a project.

All of this argument comes about because Net and Pay are artificial constructs rather than properties of reservoir rock that can be defined in terms of their chemical make-up. Nevertheless, Net and Pay are essential for describing reservoirs and so the arguments will go on as long as hydrocarbons are produced.

2.9   UNCONVENTIONAL RESERVOIRS

So-called unconventional reservoirs include oil and gas shales and coal bed – or seam – methane (CBM or CSM). Whether these are fundamentally different types of reservoir or simply represent the end members of a continuous spectrum of reservoir types is a mute point. But whether you are a 'splitter' and prefer the former or a 'lumper' and tend towards the latter, the fact remains that in practice they do require some additional properties to be adequately described.

It is worth discussing the place that coals and organic rich shales occupy in the petroleum system in a bit more detail. These lithologies are the source of nearly all-commercial petroleum and so unconventional reservoirs could be defined as those where the hydrocarbons hardly move any distance from their point of origin. CBM and oil and gas shales are often lumped together as 'self-sourcing reservoirs' for this reason. From a practical point of view many unconventional plays were originally studied as the source rocks for conventional sandstone and limestone reservoirs. One of the reasons that shale gas – and more recently oil shale – reservoirs have been so successfully exploited in the United States is that hundreds of wells had been drilled through them on the way to deeper conventional reservoirs, long before they were considered as reservoirs in their own right. As source rocks the minimum knowledge that is needed is how much hydrocarbon – if any – have they produced, what type(s) and how much has been expelled. The same is really true when they are treated as reservoirs but the emphasis and scale of the problem changes. For example, when considered as source rocks it is obviously desirable for a large proportion of the petroleum to have been expelled (in order to have any chance of ending up in a conventional trap). As a reservoir on the other hand it is better that little if any petroleum is expelled. Similarly, when considered as a source rock, studies typically want to know how much hydrocarbon could be generated across an entire basin or at least a large part of one. As a self-sourcing reservoir on the other hand the interest is in how much hydrocarbon is stored in a particular structure.

Because coals and organic-rich shales were recognised as source rocks from the earliest days of the petroleum industry, a wide range of techniques have been developed to quantify their capacity to generate and retain hydrocarbons. Many of these techniques have been routinely used for decades and are generally considered the preserve of organic geochemists. Many of them do however

involve determining petrophysical properties such as porosity or other important quantities such as total organic content (see Section 1.4).

One of the main ways that coal and organic-rich shales differ from conventional reservoirs is that a significant amount of petroleum is stored by adsorption. Adsorption basically involves some of the molecules in the fluid forming weak bonds to solid surfaces. We have already met the important case of water adsorbing to mineral grains. This gives rise to irreducible water and is the principal reason for the difference between total and effective porosity, organic molecules are more likely to bind to certain clays and organic solids – including coal. For this reason in conventional reservoirs the amount of adsorbed hydrocarbon is insignificant but in coals and black shales a large proportion of the stored hydrocarbons may actually be adsorbed. Some workers consider adsorbed hydrocarbons as contributing to the porosity and others prefer to treat them separately (the interested reader may notice the analogy with the treatment of clay-bound water in the total and effective porosity descriptions).

To further complicate matters the hydrocarbons are ultimately derived from the organic component of the reservoir. As the temperature increases the organic matter starts to 'crack' and form light organic molecules (i.e. petroleum). One of the most widely used techniques to characterise organic rich shales – whether considered as a source or a reservoir – is pyrolysis, which basically involves simulating the cracking process in the laboratory. Pyrolysis involves increasing the temperature of a sample at a well-defined rate and measuring the volume of hydrocarbons that are evolved as a function of time.

Chapter 3

Core and Other Real Rock Measurements

Chapter Outline

3.1 INTRODUCTION

Core is highly valued by most petrophysicists because it offers the opportunity to make measurements on a real piece of the reservoir and moreover they can measure the properties of interest directly. Core suffers none of the limitations caused by the large and generally irregular volumes of investigation of logging tools. Consequently core is sometimes investigated at a microscopic scale. In fact some core measurements suffer from the opposite problem to the relatively poor resolution of logs: they sample such tiny volumes that it becomes difficult to be sure they are representative of the reservoir. In this course we are primarily interested in quantitative analysis, most obviously measurements of porosity, permeability and other petrophysical properties and also the mineralogy of the formation, which enables us to refine the matrix properties used in log analysis.

3.2 TYPES OF CORE

When we talk of core we normally think of continuous cylinder of rock, cut during drilling using a special bottom hole assembly. These are often referred to as 'conventional cores'. Obtaining such cores is expensive because drilling has to be interrupted in order to run the coring assembly. Coring is also often, although by no means always, slower than normal drilling so more time is lost. Finally, having cut the core yet more time is spent recovering the core and running back into the hole to commence drilling again (to reduce the risk of damaging the core the operator should specify a maximum rate that the core can be pulled out of the hole). Each core typically adds a day or two of

FIGURE 3.1 Schlumberger mechanical coring tool from the 1970s. The two diamond saws that cut a triangular core from the borehole wall can be seen in the middle of the tool.

drilling time. The alternative is to use wireline logging tools that take small cylindrical core plugs at pre-determined depths. These are known as 'sidewall cores' and have the advantage of being relatively quick and cheap to acquire but, since the samples are small, most of the geological context is lost. In the 1970s Schlumberger and Shell together developed a wireline tool that was a compromise between sidewall and conventional coring. This took a short but continuous core from the borehole wall. The core was cut by two inclined diamond saw blades and so had a triangular cross-section (like a piece of 'Tobleron' chocolate). The tool named the mechanical coring tool has long since disappeared (Fig. 3.1).

No matter how the cores are acquired the process of taking them and then returning them to surface exposes them to some drastic changes. Amongst other things the formation fluids will be at least partially replaced by borehole fluid. Furthermore as the core is taken to surface the pressure falls so that oil may start to evolve gas and gas expands expelling part of the fluid or even fracturing the core (this type of damage can be avoided by pulling out of hole slowly). The reduction in pressure also leads to the rock expanding and possibly disaggregating. The accompanying fall in temperature may also produce changes such as precipitation of salts from saline formation water. All these factors have led some petrophysicists to challenge the validity of core data. The fact remains however that some measurements are not affected by these changes and even when they do occur, for many rocks, the changes are small and can be quantified. When large changes occur they are generally sufficiently obvious to warn the end user to be cautious. Furthermore, there are techniques available – at a price – that allow cores to be kept under pressure as they are brought to surface and/or to capture all the fluid that is expelled.

The core diameter is typically about 50% of the bit size so that in an 8.5 in. hole a 4 in. OD core is taken and in 6 in. hole the core OD is 2.5 in. (Fig. 3.2). The maximum length of a single core has increased over the years and individual cores of 30 m length are now routinely taken. Once coring starts it should continue until the barrel is full. Most problems start when coring is interrupted to see if it is still engaged.

Sidewall cores can be acquired by two different types of tool. The majority are taken with a Percussion side-wall gun. This literally fires hollow bullets into the formation. The bullets are attached to the tool with steel cables and should pull free when the tool is pulled up. In the worst case the cables will snap well before anything else in the tool string so that although the bullet
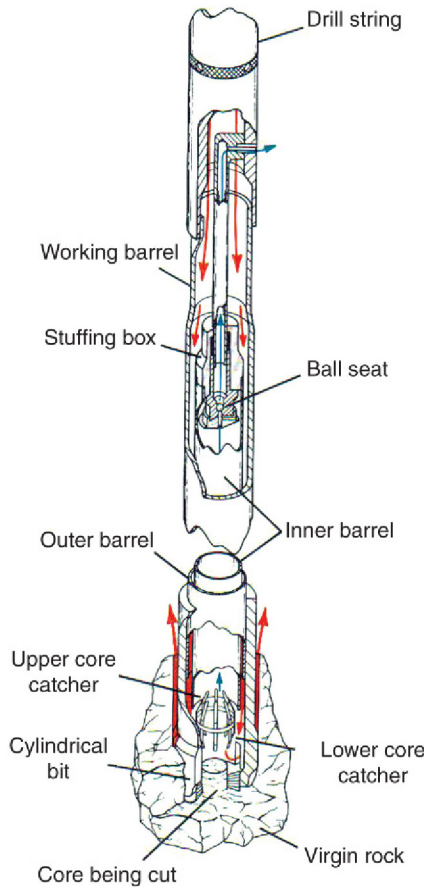
FIGURE 3.2 Conventional coring bottom hole assembly. The core is cut by a rapidly rotating hollow bit and passes into a sleeve around which the drill string rotates. When the barrel is full the core is snapped free off the formation and held in place by the core catcher.

is lost the tool can still be recovered. A typical gun holds 25–30 bullets and normally at least two guns can be run on one trip in the hole. The main objection to this type of tool is that the rock is likely to be compacted and possibly fractured as a result of the bullet being fired into it. Therefore quantitative measurements of porosity and permeability are considered unreliable. The plugs are perfectly acceptable for defining mineralogy and lithology, however and offer a cheap way of finding the cause of unusual log responses. Their main application is in bio-stratigraphy however, where the ability to place fossils at a precise depth is invaluable.

An alternative to percussion sidewall coring is the rotary or mechanical sidewall coring tool. This uses a miniature diamond-coring bit to cut the core plugs. These plugs are normally limited to 1 in. diameter and most tools can only

take a maximum of 25 plugs on one trip in the hole (tools that can cut 1.5 in. diameter plugs are now starting to enter service). The plugs can be used for quantitative measurements and the tool is often used to get plugs for core analysis from potential reservoirs that were not conventionally cored.

Historically, coring has been targeted exclusively at reservoir rocks (apart from percussion side-walls). Recently however the heavy reliance of marginal developments on a few high-angle wells has emphasised the need for good understanding of borehole stability in the overburden. This in turn requires estimates of rock strength whose reliability is enhanced by at least a few direct measurements on core samples. Hence the recovery of at least one core in the overburden is often a key objective of appraisal wells. Cores are also essential for understanding oil and gas shales. So increasingly cores are targeted at lithologies that a few years ago were deliberately avoided.

3.3   CORE MEASUREMENTS

Having acquired core it is possible to make as many measurements as are desired. Compared to the cost of acquiring core, measurements and studies are generally cheap and there is little point in stinting here. Any study should however look at a representative range of samples and not simply 'cherry-pick' what appears to be the best quality reservoir or most unusual mineralogy. It should also be noted that some measurements are very time consuming and as commercial laboratories only have finite resources, an overly ambitious study may take literally years to complete (and therefore arrive too late to have much impact on a typical modern development). Some measurements are also destructive and there will be a limit to how many samples can be sacrificed to gain that information.

Quantitative petrophysical measurements are normally made on 1.5 in. or 1 in. diameter core plugs that are cut from the conventional core using a similar bit to that employed by the rotary sidewall coring tool. The size of the plug depends on the diameter of the core and possibly the particular measurement to be made. Plugs are ideally cut parallel to bedding and modern practice tends to emphasise taking plugs at regular intervals (e.g. 25 cm, 50 cm or 1 ft.). This is supposed to avoid biasing the data set to the best-quality reservoir. It is however also desirable to take plugs in as homogeneous a piece of rock as possible and this will inevitably result in some variation from an even spacing. In thinly bedded formations the plugs will inevitably be taken in the sands, which will be unevenly spaced.

A real example of the coring and core analysis of an appraisal well follows. The well appraised a large discovery in which the reservoir consists of massive, well consolidated but high-quality sandstones (porosities of the order of 20% and permeabilities of the order of 1000 mD). Four cores with a total length of 207 m were cut, between them they covered nearly all of the reservoir. A total of 688 1.5 in. diameter plugs were taken for porosity and

permeability measurements. About 85% of these were horizontal plugs – that is parallel to bedding – with an average spacing of 30(±10) cm. The remainder were vertical plugs cut close to the location of a horizontal plug. On the other hand just 23 samples were taken for petrology and petrography studies. The latter included quantitative mineralogy, grain size analysis and light and electron microscopy. The core plugs represent about 10% of the total length of the core but less than 4% of its volume. The petrology samples represent less than 0.1% of the volume of the core they therefore have to be carefully chosen to be representative of the core.

It is possible to make measurements on sections of un-plugged core. The core is either cut into suitable lengths or pre-existing breaks are used to define the samples. Sample lengths can be quite variable but rarely exceed 30 cm. Whole core analysis is particularly useful for vuggy carbonates because the vugs can have similar dimensions to a core plug.

3.4 PREPARATION FOR ANALYSIS

When the core arrives in a laboratory it is likely to contain a mixture of the original reservoir fluids, mud filtrate and air. The proportions of these will vary continuously along the core and depend on the nature of the rock, the drilling fluid, the coring assembly, how it has been handled etc. The first step in core analysis is normally to clean and dry the core plugs (or the whole core sections if that is being used for analysis). This involves firstly removing the original fluids using two or more solvents and then drying the core plugs in an oven. The cleaned and dried core is of course in a very different state to its natural condition in the reservoir. The intention is at this stage is to get all the plugs into a consistent state not a natural one.

The cleaning and drying steps can take several days and will lead to misleading results if they are not carried through to completion. By convention cleaning normally involves flushing the core plugs with a non-polar and a polar solvent. The former that is invariably toluene, is intended to remove any oil and the latter, which is generally methanol, removes water (formation water and/or mud filtrate). There are various ways to introduce the solvents. The two most common are to flow cold solvent through the plug using an apparatus similar to the steady-state permeability equipment shown in Fig. 3.3 or to bathe the plug in warm solvent. The latter uses special laboratory glassware so that the contaminated solvent is physically separated from the core plug. Often several plugs are cleaned at the same time but if the plugs are cleaned separately it is possible to work out how much oil and water has been removed by the solvent at the end of the process.

A slightly more complicated process known as Dean–Stark extraction is designed to quantify the volume of oil and water present in the core plug. The plug is bathed in a continuous flow of toluene vapour from a heated flask. Both oil and water in the core evaporate and combine with the toluene. The
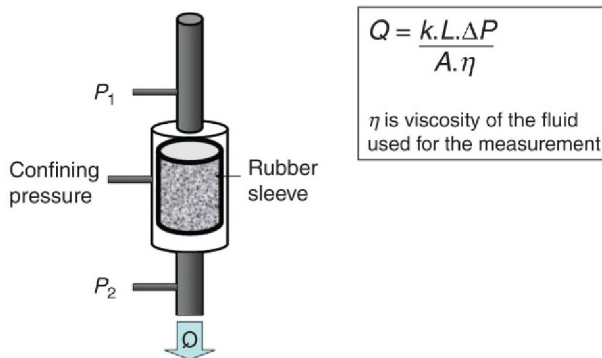
$$Q = \frac{k.L.\Delta P}{A.\eta}$$

η is viscosity of the fluid used for the measurement

FIGURE 3.3   Schematic of a permeameter.

vapour is then condensed away from the plug and the water, which is immiscible in liquid toluene, is collected in a graduated vessel. The toluene and any oil that has dissolved in it flows back to the flask to continue the process. At the end of the process, which typically takes a few days, the water volume is measured directly and the oil volume is inferred from the weight loss of the plug.

Once all the original fluids have been removed the plug is dried in an oven to remove the solvents and any remaining traces of water. Drying conditions vary depending on the nature of the rock and/or company procedures. Some operators prefer a very aggressive drying process with a temperature in excess of 100°C and a dry atmosphere; others prefer a lower temperature and a humid atmosphere. There are pros and cons to both: aggressive drying more or less ensures all core plugs end up in the same condition and results in a porosity that is as close to true total porosity as possible. Humidity drying is less likely to damage the clays and therefore produce core permeabilities that are more representative of *in situ* conditions, but different plugs will lose different proportions of the water.

3.5   CORE POROSITY

Porosity is the volume fraction of fluids in the sample and for a core plug can be defined as the ratio of fluid volume (V_f) to bulk volume (V_b).

$$\varnothing = \frac{V_f}{V_b} \tag{3.1}$$

A porous rock consists of fluid and grains (solid) and in general the bulk volume of a sample is given by:

$$V_b = V_f + V_g \tag{3.2}$$

$$V_b \qquad V_g \qquad \varnothing = (V_b - V_g)/V_b$$

Direct measurement
Mercury displacement

$$\varnothing = V_f/V_b$$

Retort
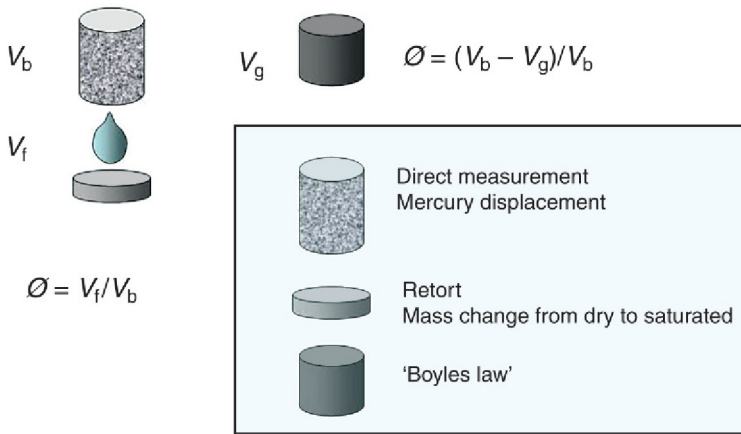Mass change from dry to saturated

'Boyles law'

FIGURE 3.4 Principles of porosity measurement. The total volume of a porous solid is the sum of the volumes of fluid and grains. If any two of these volumes can be determined the porosity can be found.

Provided we can find two of the three quantities we can calculate porosity (Fig. 3.4). For example, knowing bulk and grain volume porosity is calculated using:

$$\varnothing = \frac{V_b - V_g}{V_b} \qquad (3.3)$$

Most modern laboratories measure bulk volume and grain volume and use the formula given in Fig. 3.4 to find porosity. The bulk volume for a cylindrical plug is given by:

$$V_b = \pi r^2 l \qquad (3.4)$$

Where r is the radius of the plug and l is its length (diameter $-2r-$ is normally 1 in. or 1.5 in.: 2.54 cm or 3.81 cm). The grain volume is found by exploiting one of the ideal gas equations using the apparatus shown schematically in Fig. 3.5.
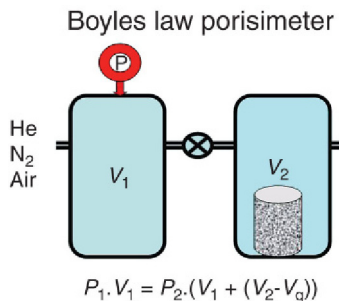
Boyles law porisimeter

He
N_2
Air

V_1

V_2

$$P_1 \cdot V_1 = P_2 \cdot (V_1 + (V_2 - V_g))$$

FIGURE 3.5 Boyles law porisimeter.

The core plug is placed in chamber 2 and all the air is evacuated so the pressure is zero. At the same time gas is introduced into chamber 1 at a known pressure P_1. The valve connecting the two chambers is then opened so that the gas can flow into chamber 2. The pressure falls to P_2 which can be calculated by re-arranging ideal gas equation:

$$P_1 V_1 = P_2 (V_1 + (V_2 - V_g)) \tag{3.5}$$

In practice we measure the pressures and know the chamber volumes V_1 and V_2 so that Eq. 3.4 can be re-arranged to give V_g. The apparatus is constantly checked by replacing the core plug with a stainless steel cylinder of known volume. The preferred gas is helium, which has the smallest molecular diameter of any substance. It can therefore enter even the smallest pores and will very quickly occupy all the pores in the sample.

As noted above providing the plug is perfectly cylindrical the volume can be accurately calculated. The plug may not be perfectly cylindrical however: grains may be lost from the surface, the ends may not be planar and corners may get chipped. Furthermore there is a limit to how accurately the dimensions can be measured. An alternative to calculating V_b is to measure it by displacement. A fluid needs to be chosen which will not enter the pore system and mercury is the best candidate for this.

There are several alternative ways of measuring porosity. In the Soviet Union it was common practice to find the pore volume directly. To do this the weights of the plug dried and then saturated with kerosene were measured. The difference would give the weight and hence volume of the kerosene in the pore space. Old core reports often quote a 'porosity by fluid summation'. This involved collecting the fluids extracted from the plug during cleaning and drying and measuring their volume. In both these techniques the pore volume – V_f – is measured directly and so porosity would be calculated from Eq. 3.1.

3.6   GRAIN DENSITY

If the grain volume and mass are known the grain density can be found. This is one of the most important parameters in log analysis and was briefly mentioned in Chapter 1 (Eq. 1.6). The grain volume is actually what is measured by the Boyles law porisimeter and providing the apparatus is used properly it can be accurately measured regardless of the sample shape. The mass of the plug can be measured on a chemical balance to four or five significant figures (the humble chemical balance is actually one of the most accurate and precise scientific instruments). The mass of the plug is very close to the mass of the grains alone because the pore space is filled with air with a density of $0.0006\,\text{g/cm}^3$. So the grain density is simply given by:

$$\rho_g = \frac{\text{Mass of plug}}{V_g} \tag{3.6}$$

Even if one questions the validity of core porosity measurements – because of the changes the rock undergoes between being drilled and measured – the grain density is the true density of the solid part of the rock regardless of whether it is in the laboratory or *in situ*. That alone should be sufficient reason to cut core.

3.7 PERMEABILITY

To measure permeability fluid has to flow through the plug. Permeability can be measured in two ways. The simplest to interpret, known as the 'steady state' method, exploits the simplest form of Darcy's equation. This was the standard method until the 1980s and is still used in many commercial laboratories.

$$Q = \frac{kA\Delta P}{l\eta} \tag{3.7}$$

For cylindrical plugs the cross-sectional area (A) and length (l) are determined by the plug dimensions. Normally the fluid is a well-characterised gas such as helium or nitrogen. Both these have very well understood relationships between viscosity and pressure. It is then possible to fix the flow-rate and measure pressure difference across the plug or vice versa or simply measure both when the system has come to equilibrium. As with porosity it is normal practice to frequently check the accuracy of the apparatus by replacing the samples with a well-characterised standard.

The alternative method is essentially a miniature well test in which the response to a step change in applied pressure is recorded as a function of time. The shape and duration of the transient can be analysed to find permeability. All other things being equal, the slower the response the lower the permeability. These are the standard methods for very low permeability rocks such as oil and gas shales because they can produce reasonably accurate matrix permeabilities in a time scale of hours if not minutes. The technique is also exploited in mini-permeameters, which can measure permeability from anything with a flat surface. They are used to measure permeability on outcrops and also to produce lines or maps of closely spaced readings from slabbed core surfaces. In fact, the transient method is so quick that in many commercial laboratories it is now the standard method for determining permeability of any rock. The results are normally in good agreement with steady-state measurements although for permeabilities in excess of 1000 mD the steady-state method is still preferred.

For low-permeability samples there is a danger that most of the gas will flow around the outside of the sample by exploiting the roughness of the plug surface. To try and avoid this, the plug is placed in an elastic sleeve, which is forced to conform to the plug surface by applying external pressure. Unfortunately, permeability is dependent on the confining stress applied to the plug and so will depend on the magnitude of this sleeve pressure. Some data for four plugs from the North Sea Forties sands are shown in Fig. 3.6. It can be seen that
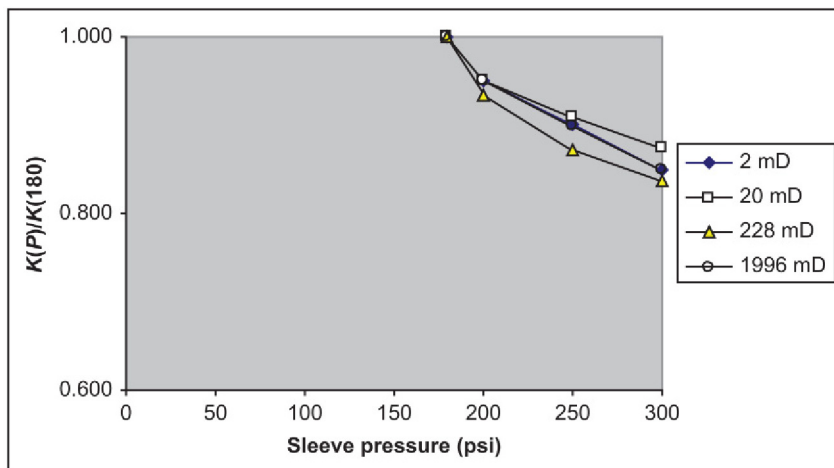
FIGURE 3.6 Decline in permeability with sleeve pressure for four plugs from the North Sea Palaeocene.

increasing confining pressure from 180 to 300 psi causes permeability to fall by over 10%. The preferred confining pressures have been increasing steadily over the years. In the 1990s applied pressures of 200 psi were typically used but 800 psi is not unusual for contemporary measurements. Care must obviously be exercised when comparing dataset of different vintages.

3.8   SPECIAL CORE ANALYSIS

Special core analysis includes a large range of measurements that are more complicated than those described earlier. They typically involve the use of two immiscible fluids, at least one of which is found in the reservoir (e.g. air–brine, oil–brine). The measurements normally take a long time to produce satisfactory results and so realistically they can only be applied to a small sub-set of the core plugs. These plugs need to be carefully selected to ensure they are representative of the rock types that are of greatest significance to the development (i.e. hold most of the reserves or perhaps represent most of the aquifer). They also need to be quality checked to make sure there are no fractures or heterogeneity that might invalidate the measurement (Fig. 3.7).

Special core measurements that are of particular interest to this book are electrical properties, capillary pressure curves and compressibility. The first two will be discussed in more detail later and the former is basically the measurement that Archie made that was discussed in Chapter 1. The objective is to obtain accurate values for the parameters so that resistivity logs can be converted to water volume and hence saturation.

Capillary pressure curves involve finding how much non-wetting fluid can be forced into a plug when a particular external pressure is applied. By increasing
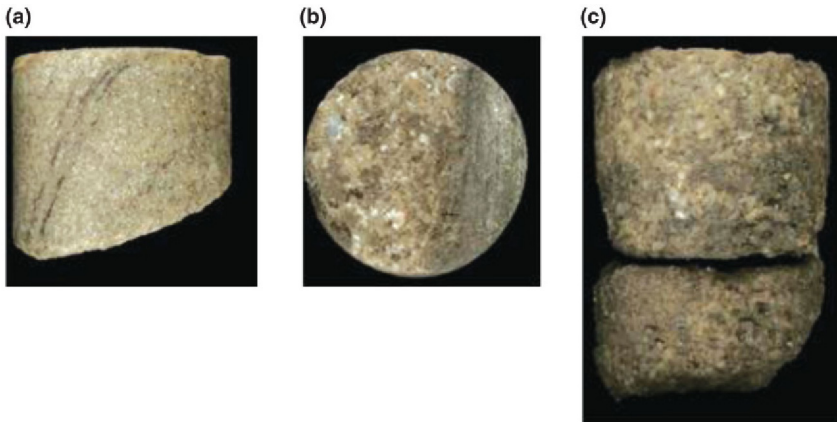
FIGURE 3.7 Examples of plugs that are unsuitable for SCAL work. (a) Too short. (b) Heterogeneous. (c) Fractured.

the pressure in steps a curve of saturation against applied pressure is produced. This can be used to derive the saturation-height function (SHF) that was briefly introduced in Chapter 2. The SHF in turn determines how hydrocarbon distributes itself within the reservoir.

3.8.1 Compressibility

Compressibility measurements are used to determine how porosity and permeability are reduced on going from laboratory conditions to reservoir stress. In the case of porosity this is important if the core porosity is being used to calibrate logs or even to just check the match between log derived and core porosity. Rock compressibility can also be a significant source of reservoir energy. As a rule the greatest reductions in porosity occur in high porosity, unconsolidated sandstones where reductions by a factor of 0.9 are not unusual.

Permeability reduction under stress is obviously important for predicting well performance. As a rule the greatest reductions occur at the lowest permeabilities. Tight gas sands in particular can show large reductions in permeability when the measurements are made at simulated over-burden conditions. The data for the highly cemented sandstone shown in Fig. 3.8, suggests *in situ* permeabilities are three times lower than the values measured in the laboratory. By contrast unconsolidated sands with permeabilities of the order of 1000 mD often show reductions of, at most, 20% on going to reservoir conditions.

3.8.2 Klinkenberg Effect

Compressibility measurements account for the higher stresses in the reservoir caused by hundreds or thousands of metres of over-burden. The increased pressure on the gas in the reservoir further reduces permeability through the Klinkenberg
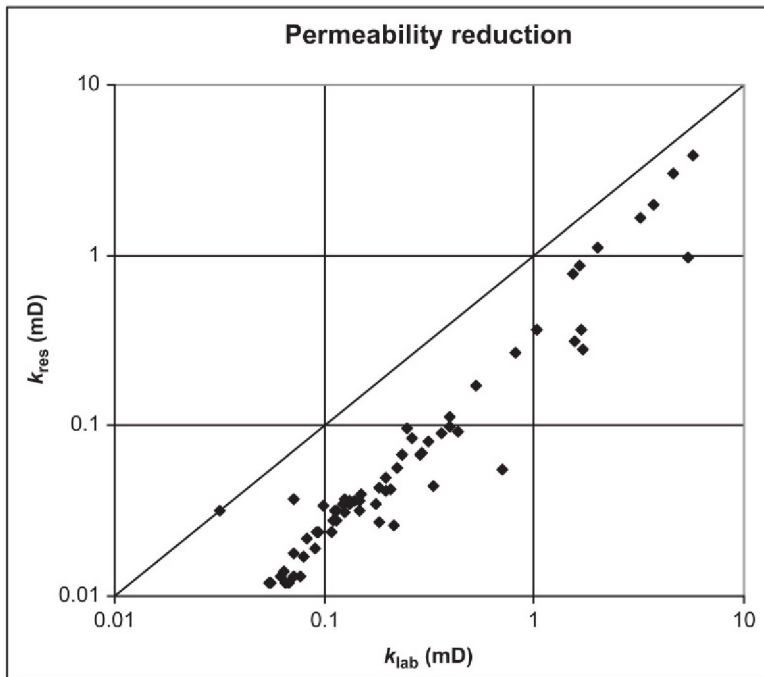
FIGURE 3.8 Routine permeability measured at laboratory stress (400 psi) and simulated reservoir stress (4000 psi) for plugs from a tight gas reservoir. A permeability of 0.1 mD measured at laboratory conditions will be one third this at reservoir stress.

effect (discussed in Section 2.4.1). The measurement of the Klinkenberg effect is normally considered part of routine core analysis but it is discussed here because it involves a correction from laboratory to reservoir conditions. For steady-state permeability measurements the Klinkenberg permeability – k_{inf} – is found by measuring the permeability at several different gas pressures. A graph is then constructed of permeability against reciprocal pressure and the intercept with the permeability axis gives the Klinkenberg permeability. An example is shown in Fig. 3.9. Because this is quite a time-consuming experiment it is normally only applied to a small sub-set of plugs. These are then used to find a general relationship between the k_{inf} and ambient pressure permeabilities (k_{air}). In the case of transient measurements however, the Klinkenberg permeability, can be found by analysing the pressure transient and so a value for k_{inf} is produced for every plug.

3.9 OIL AND GAS SHALES

All the techniques described in the preceding sections were developed to deal with conventional reservoirs. That is reservoirs that store most of their producible fluids in macro-porosity (inter-granular pores or larger systems such as fractures or vugs). But in fact petroleum has been produced from shales
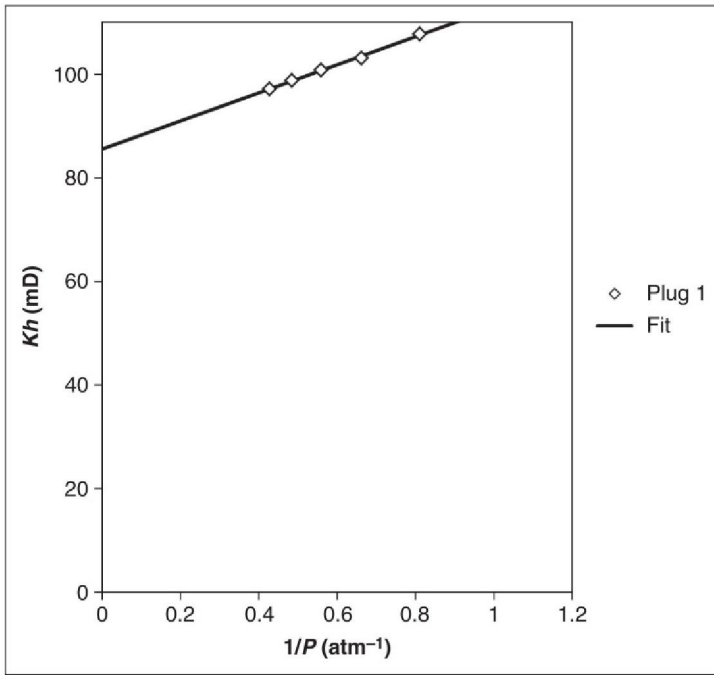
FIGURE 3.9   Determination of Klinkenberg permeability for a sandstone core plug with a permeability to air at ambient conditions of approximately 100 mD. The Klinkenberg permeability is 85 mD.

for as long as the oil industry has been in existence and oil shales were mined in quarries long before conventional sources were tapped with wells. Since the 1990s shale gas and oil has become such a significant resource that objective ways of assessing the size of these resources has become essential. As with conventional reservoirs cores are used to ground truth properties estimated from logs but if anything they are even more important for evaluating oil and gas shales. There are a number of reasons for this:

1. Hydrocarbons in shales are not simply stored in pores, they are also adsorbed on mineral surfaces. The only way of reliably quantifying the amount of adsorbed matter and estimating how much can be recovered is by measurements on the core.
2. Log analysis of sandstones and limestones can often give reasonable estimates of porosity by just using textbook matrix parameters. Shales have more complex mineralogies and so their matrix parameters are far more variable and need to be measured.
3. Like it or not, effective porosity becomes a key characteristic of shales when they are reservoir rocks (in conventional reservoirs the shales are effectively assumed to have zero-effective porosity).

In principle all the techniques described earlier could be applied to shales. In practice, however the extremely low permeability of shales means that prohibitively long times are required to both remove the original fluids and re-saturate with the measuring fluid (even if that is helium). Together these mean that both preparation and the measurements take far longer and may require more sensitive equipment. Consequently, a set of procedures was developed in the early 1990s to allow operators to routinely measure petrophysical properties of shales on time scales comparable to those applied to conventional reservoir lithologies.

The distinguishing step of these procedures is that the core plugs are deliberately crushed prior to the measurements. This may seem to represent such a drastic change that the measurements are worthless. But by crushing the sample the cleaning solvents can access the original fluids quickly and the analyst can be sure they have all been removed. The volumes of the extracted fluids are invariably measured in oil and gas shales because in such low-permeability rocks they are unlikely to change appreciably as a result of drilling.

The grain volume of any sample, including a crushed one, can be found using the Boyles law method described earlier. Having found the grain volume it is simple to find the grain density and if the volume of the original plug was known total porosity can be calculated as well. So in fact crushing the sample simply speeds up the time it takes for the gas to occupy all the free space.

The permeability of a crushed sample can be estimated using a transient technique. The apparatus to do this is shown schematically in Fig. 3.10a. The experiment involves charging a chamber filled with the crushed sample with air or some other suitable gas at a pressure of about 1000 psi (1200 psi in Fig. 3.10a). The system is allowed to come to equilibrium, which is indicated by the pressure attaining a constant value, in Fig. 3.10b this has occurred at time 10. A second chamber is charged with gas to a slightly higher pressure: greater than 1300 psi in Fig. 3.10b. The two chambers are then instantaneously connected by opening a valve. Initially the pressure in the system will rapidly rise to an intermediate value as the higher-pressure gas enters the gaps between the fragments. But then the pressure will slowly fall as the gas enters the pore spaces within the fragments. By analysing this slow decay in pressure with time the permeability of the rock can be calculated. The dimensions of the apparatus are designed so that the reservoir volume is similar to the pore volume of the crushed sample, as this gives the greatest change in pressure when the connecting valve is opened. In practice this means the reservoir chamber has a volume of a few cubic centimetres (as the reader can confirm by calculation). The experiment can also be performed by charging the reservoir chamber to a lower pressure than the sample chamber and in practice two reservoirs may be connected to the sample chamber to allow the sample to be subjected to both a sudden increase and a sudden decrease in pressure. The same technique can be applied to complete core plugs but obviously it takes much longer for the high-pressure gas to reach the pores in the centre of the plug.
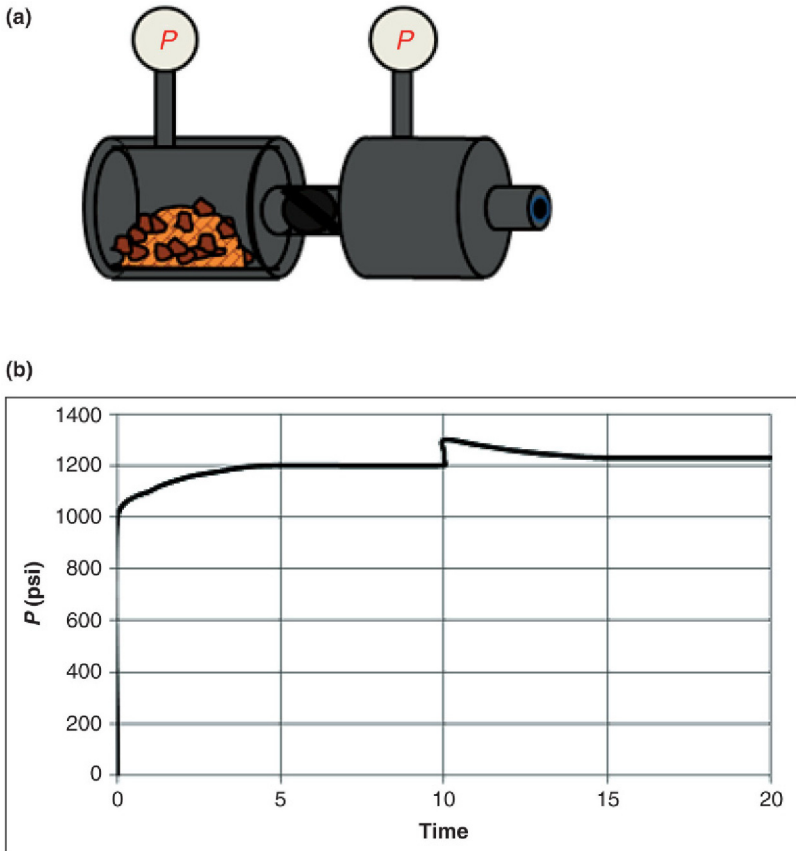
(a)

(b)

FIGURE 3.10 (a) Schematic of the apparatus used to measure permeability of crushed shales. (b) Change of pressure with time that is interpreted to give the matrix permeability. The time units are arbitrary (see text for explanation).

3.10 CUTTINGS

Most of this section has been concerned with cores or core plugs that have been deliberately targeted at points or intervals of particular interest. Drilling produces a continuous record of rock samples in the form of cuttings, however and for large parts of well these will be only rock samples available. In fact, often they will be the only samples. The main problem with cuttings is that it is not clear exactly where they originate. For a typical well near TD it takes several hours for the cuttings to reach the surface and in that time they will become mixed with cuttings from above and below their point of origin.

Most of the information routinely obtained from cuttings is qualitative or at best semi-quantitative. This includes lithology and mineralogy of the major components as well as features such as colour and hardness. For clastic rocks

information on grain size and shape, sorting and degree of cementation is also generally given. Nevertheless, this is all very useful information that complements the physical properties measured by logs. Taken together the two types of information complement each other well and the logs can often be used to 'unravel' the mixing process described in preceding section.

Unfortunately, modern insert bits tend to grind the cuttings to a paste known as 'rock flour' and as a result a lot of information is lost. Where good cuttings are produced it is possible to make some quantitative measurements on them. Any of the techniques described earlier that do not rely on samples having a precise geometrical shape can be applied to cuttings.

Chapter 4

Logs Part I: General Characteristics and Passive Measurements

Chapter Outline

4.1 INTRODUCTION

It is true that petrophysics involves more than just log analysis but the fact is logs normally represent the bulk and possibly all of the data used to estimate the formation properties. It follows that log analysis is an essential tool in petrophysics and much of the effort in a typical interpretation goes into converting logs to property curves. In other words petrophysicists may do a lot more than just log analysis but they still have to be good log analysts. There are many publications that describe the different types of logs and how they work in detail. Some of the best examples are produced by the logging companies and are generally available free. It is not the intention here to duplicate these accounts rather we will look at how the measurements can be used and what exactly is being measured.

Logs are graphs of a physical property against depth. In practice modern computer logging units produce a series of evenly spaced measurements (computer logging units first appeared in the late 1970s and by the mid-1980s had largely replaced the older 'camera' systems). This is true for both wireline and logging while drilling (LWD) measurements, although, the latter normally require some additional re-sampling to produce the even spacing. The commonest depth increments are 6 in. or some fraction of this even if the depth is ultimately measured in metres (although a few companies do offer increments that are fractions of a metre, e.g. 10 cm). Although measurements are made at fixed

depth points they are actually the average value over a relatively large and often irregularly shaped volume centred at the depth point. These volumes normally extend a much larger distance along the hole than the depth increment. This means that logs are actually some form of rolling average of the property and may not be able to resolve the true variation of a property.

Logs are completely objective in the sense that they produce a measurement every time the sensor moves a fixed distance (or at least that is the intention). This occurs regardless of the lithology, hole conditions or anything else. Cores on the other hand are normally cut at carefully chosen locations to deliberately target a particular lithology. As a corollary logs can give almost complete coverage of a well bore whereas, with a few exceptions, cores are never intended to cover more than a small part of the hole.

4.2   WIRELINE AND LOGGING WHILE DRILLING

There are two ways to acquire logs:

1. Wireline
2. LWD

By and large these are just different ways of getting the sensors into and out of a well. With one or two notable exceptions a particular type of tool uses the same principles whether it is in LWD or wireline form. In fact, increasingly, wireline and LWD tools offered by the same company share a lot of components. For much of this book we will not bother to make a distinction between the two types of tool but it should be noted that each type of tool does have a few advantages and disadvantages. Some of these are marginal but in some applications they may be significant enough to determine what type of tool to use.

Wireline tools were of course the first to be developed and for the first 50 years or so had the 'field' to themselves. During that time nearly all the logs currently available were developed. As the name suggests wireline logging involves lowering the tool down the well on the end of a cable – the 'wireline' – and then pulling it back out. More often than not, the wireline includes electrical conductors to provide power to the tool and to transmit the signals back to the recording instruments at surface. Increasingly, control signals are also sent from surface to the tool. Since the 1980s it has been possible to dispense with an electrical connection to the tool and simply run the tool on 'slick-line'. The tool is then powered by a battery and the raw data, as a function of time, is stored in down-hole memory. Providing one knows the time the tool was at a particular depth, the time-based raw data can be converted to a conventional depth-based log. Whether run on an electrical cable or slick-line, depth is ultimately measured by the length of cable that passes over an accurately calibrated depth wheel at surface. If used properly, such a system can measure depth to an accuracy of a fraction of a per cent.

In LWD form the sensors are mounted in a drill collar located a few metres behind the bit. LWD first appeared in the late 1970s and at first only the gamma ray and a few simple resistivity tools were available. Within 20 years however all the basic wireline tools were available in LWD form and although there is still debate as to the relative quality of the measurements compared to wireline, in the author's opinion modern LWD measurements are as good as their wireline equivalents. The most obvious advantage of this method of logging is that the logs are continuously produced as the well gets deeper; in particular the log is recorded soon after the formation is drilled rather than several days later when the well has reached some key depth like a casing point.

At the time of writing the standard way of communicating with the surface is through pulses in the mud. Compared to electrical signals this is very slow (as a rough guide it is possible to transmit several hundred thousand bits a second over an electrical cable versus a few bits per second using mud pulses). Consequently, the transmitted – or 'real time' – logs are often quite sparsely sampled and restricted to a handful of key curves. The majority of the data has to be stored in down-hole memory and is not available until the drill string is brought back to surface. It is this 'memory data' that generally produces the wireline quality logs. There is an alternative to mud pulses known as 'wired pipe' which goes someway to giving wireline quality logs in real time. The key component is special drill pipe with an electrical conductor running through it. Although not giving quite the transmission rates of a logging cable it still allows transmission rates several thousand times higher than the mud-pulse system. The main disadvantage is the cost of manufacturing the pipe.

Memory data can offer significant advantages over wireline logs even if one has to wait until the LWD tools are back at surface. The two most general advantages are:

1. LWD avoids the time to rig-up, run and rig-down wireline services. This may amount to several days, which on a deepwater well can equate to millions of dollars worth of rig time. If nothing else it reduces the time the well is open and in danger of collapse.
2. No special equipment is required to log high-angle wells; if you can drill it you can log it. By contrast the equipment needed to run wireline logs in a high-angle well is complicated, failure prone and requires experienced logging and drilling crews to operate it.

There are other advantages to LWD, which sometimes apply.

3. The LWD sensors normally pass over a bed sooner after it was drilled than a wireline tool so there *may* be less formation alteration.
4. LWD sensors normally rotate with the drill string so that for some measurements an image can be generated and the effects of heterogeneity on directional measurements can be reduced.
5. LWD allows wells to be geo-steered.

It is fair to ask what are the disadvantages of LWD? Arguably, the biggest disadvantage is depth control since LWD is ultimately dependent on the drilling depth curve to produce a conventional depth-based log. This actually is not a great limitation, contrary to popular belief the driller's depth is normally as good an estimate of depth as the wireline logger's. Where wireline scores however, is in measuring depth intervals. So for example, if a wireline log suggests two tops are 100 m apart they almost certainly are separated by that distance to within a few centimetres. The drillers depths on the other hand may produce an uncertainty of several metres. Both LWD and wireline loggers have the technology to be much more accurate in their depth estimates and are quite happy to oblige. Logging cable and drill pipe are very well-characterised components and if a continuous record of the loads, temperatures and pressures exerted during logging is available then their lengths can be corrected. The reason this correction is not done routinely is a bit of a mystery but can probably be blamed on the operators' desire for results as soon as possible.

The other major disadvantage is that whilst most wireline tools will run in a wide range of different hole sizes, a different LWD tool is needed for each of the main bit sizes. This means contractors have to hold more tools and that, ultimately, will be passed on to the customer as higher rentals.

4.3 CHARACTERISTICS OF LOGS

As noted earlier a log is a graph of a physical property against depth (or strictly speaking distance along the well track). Although modern logs consist of a series of closely spaced measurements, for all intents and purposes they are a continuous record (prior to computerised logging units they were truly continuous). In this section, we will look at what exactly the measurement refers to. Most logs give the average reading of a volume of formation whose dimensions are large relative to the depth increment. This is why we could assert that logs are essentially continuous records. The volume of investigation is characterised by two key metrics:

1. Vertical Resolution. What distance along the hole contributes to a measurement.
2. Depth of Investigation. How far into the formation does the tool sense.

These typically vary from one type of log to another and more often than not, they actually depend on the value of the property that the tool is trying to measure. In other words they vary continuously along the well. The depth of investigation of the density tool, for example decreases with increasing formation density. For some tools, notably the resistivity devices, tool designers have a lot of control over the resolution and depth of investigation, but mostly these are constrained by the size and power limitations imposed by working in a borehole. Even with resistivity tools vertical resolution and depth of investigation tend to be linked, so that, to see deeper one has to accept lower vertical
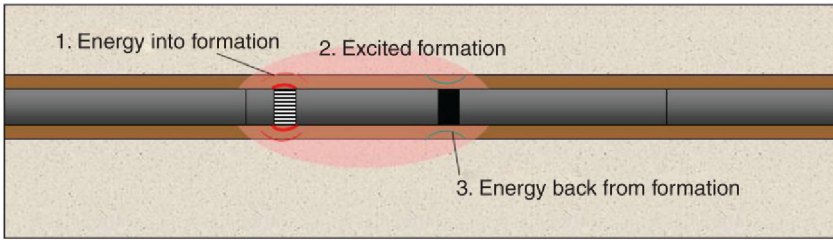
FIGURE 4.1 A generalised logging tool.

resolution. It is certainly the case that very few tools have volumes of investigation, which have simple geometrical shapes like a cylinder or a sphere. Furthermore, within the complex volumes of investigation which do result, the contribution to the measurement is not everywhere same. In short, most log measurements are some sort of weighted average property over a complicated 3D shape. The shapes vary from tool to tool and normally their size changes in response to changes in the formation.

4.3.1 Vertical Resolution

Most logging tools work by putting energy into the formation at one point and measuring the response some distance away (Fig. 4.1). The distance between the source of the energy and the receiver puts a lower limit on the vertical resolution (in practice it is often larger than the source–receiver spacing). So it might seem that simply reducing the spacing will improve vertical resolution. For some tools this is true, but unfortunately the depth of investigation also falls as the source and receiver are moved closer together. In other words there is a trade-off between vertical resolution and depth of investigation. As we have already seen the drilling process alters the formation in the near well bore region so that as the depth of investigation reduces, the altered formation becomes larger and larger component of the measurement. Worse still, measurements with very shallow depths of investigation will be sensitive to even small imperfections in the borehole wall, which allow mud to get between the tool and the formation. The shallower the depth of investigation the smaller the imperfections that can be accommodated, so tools are normally designed to read at least 10 cm beyond the borehole wall. Other practical constraints may also limit source–receiver spacing so that for many tool types designers actually have very little latitude in that distance.

We can be a bit more precise defining both vertical resolution and depth of investigation. In this book vertical resolution is defined as the minimum bed thickness for which the tool will return a true reading.

Definition: Vertical resolution is the minimum bed thickness required for the logging tool to measure the true value of the formation (within the depth of investigation of the tool).
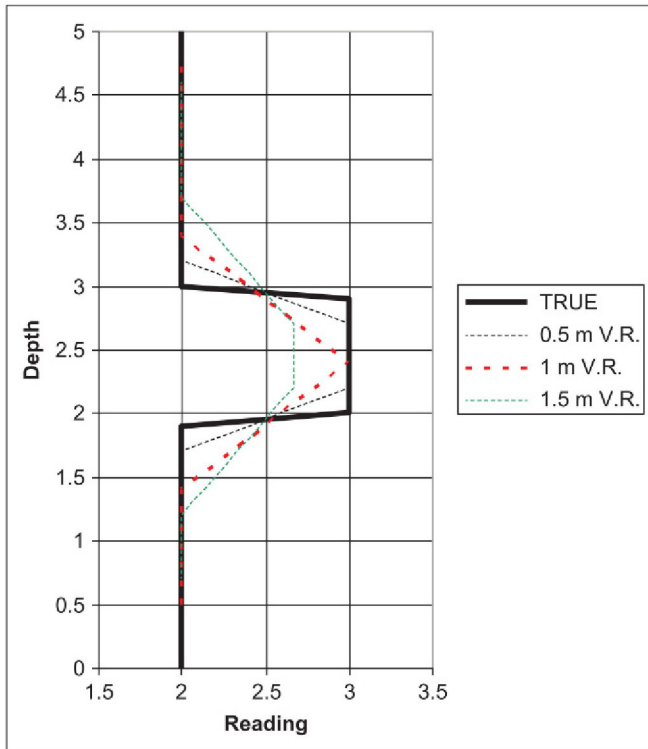
FIGURE 4.2 Response of three different tools to a 1-m thick bed with true property reading of 3. The tools have vertical resolutions of 0.5, 1.0 and 1.5 m. The tool with the 1.5 m vertical resolution cannot resolve the bed although it certainly responds to it.

If the bed is thinner than the vertical resolution, the measurement will be influenced by the formation above and below the bed. This is not to say the tool will not respond to the bed – it almost certainly will – it is just that it will not be able to tell us what the true reading of the bed is. This is illustrated by the hypothetical general tool responses shown in Fig. 4.2. Here, a 1-m thick bed with some physical property with a value of 3 units lies in a background of rock with a uniform value of 2 units.

The tool with the best resolution (0.5 m) returns a true reading for much of the bed thickness and it therefore does properly resolve it. The tool with the 1 m vertical resolution does read the true property value in the centre of the bed and therefore also properly resolves it – just. The lowest vertical resolution device however, never reads more than 2.67 units and therefore does not properly resolve the bed. Often logging contractors report the minimum bed thickness that a tool can detect as 'the vertical resolution'. Few tools can properly resolve beds less than a metre thick but a tool can detect something that is much thinner, possibly just a few centimetres.

In practice the vertical resolution is further reduced by 'filtering' or 'smoothing'. This is done to improve the appearance of the log and also to make it more repeatable. The actual value plotted at a particular depth point is made up of the measurement at that point together with a contribution from the measurements either side of it. A commonly used algorithm is:

$$X(i) = 0.25x(i-1) + 0.5x(i) + 0.25x(i+1) \tag{4.1}$$

Where $X(i)$ is the value plotted on the log at depth point i and $x(i)$ is the raw measurement at depth point i. The unfiltered measurements are still recorded so it is possible to play back an unfiltered log or for that matter a more heavily filtered log.

4.3.2 Depth of Investigation

Very few tools have a definite depth of investigation beyond which there is no influence on the measurement. Rather, there is a diffuse boundary over which the contribution to the measurement falls to insignificant levels. In order to compare different tools a definition is required that accounts for this diffuse boundary. Here we will adopt a common definition that the depth of investigation is the distance from the tool inside which 90% of the signal originates.

Definition: Depth of investigation is the distance measured from the outer surface of the tool, from which 90% of the measurement originates.

As we have already noted the depth of investigation typically varies as the formation properties change. It may also depend on external factors such as temperature and mud type. The depth of investigation is normally a more clear-cut specification than vertical resolution and tool designers put a lot of effort into determining it either by experiments or detailed computer models. But for certain tool types it is actually quite difficult to define exactly what it is – notably the sonic log.

The twin metrics of depth of investigation and vertical resolution are illustrated for the hypothetical tool in Fig. 4.3 (in this case the depth of investigation is quite straight-forward).

4.4 VOLUME OF INVESTIGATION OF LOGS

In order to completely specify the 3D volume of investigation of the tool, the radial or azimuthal distribution of the measurement also needs to be specified. This is basically the extent of the arc, looking down the hole, which contributes to the measurement (Fig. 4.4). Together with the vertical and horizontal distributions this defines the 3D volume that the tool 'sees'. For many tools there is no preferred direction and the signal is gathered from all around the borehole: the measurement is therefore averaged over 360°. Other tools, particularly those in which the sensors are deliberately pushed against the borehole wall, only investigate a limited arc of perhaps 60°–90°. Some other
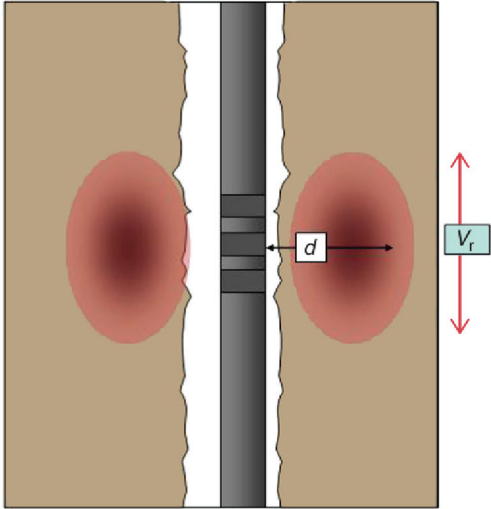
FIGURE 4.3   A hypothetical logging tool illustrating the source of the signal contributing to the measurement (oval shape [pink coloration in the web version]). The resulting vertical resolution (V_r) and depth of investigation (d) have been annotated. Assuming the tool has a typical 3 5/8 in. outside diameter, the depth of investigation is approximately 30 cm (12 in.) and the vertical resolution is approximately 50 cm.
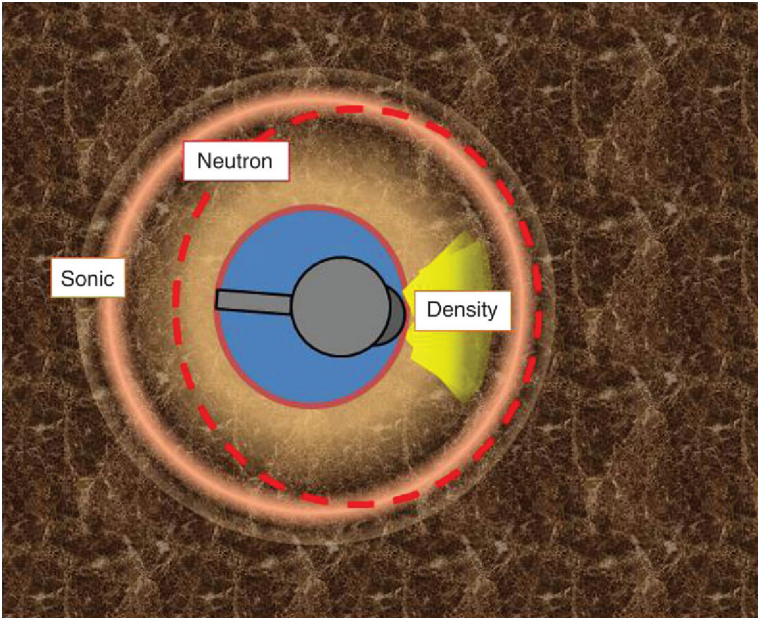

FIGURE 4.4   Azimuthal distributions for three conventional logs. The density is a shallow reading tool with a narrow azimuthal distribution. The sonic has a 360° volume of investigation and a deeper depth of investigation.

tools are intermediate between these cases and gather information from all directions but not equally.

The azimuthal distribution determines how the tool responds to heterogeneity. In a homogeneous formation it does not really matter whether the measurement extends right around the borehole or is confined to a narrow arc. In a heterogeneous formation however, a directional measurement will depend on exactly which way the tool is orientated whereas a 360° measurement will give the same averaged reading regardless of where the tool is located. For example, in a conglomerate an 'all-round' measurement will average the values for the clasts and the matrix, whereas a directional measurement may be orientated so that all that appears in the volume of investigation is a clast or pure matrix.

LWD tools can average readings around the borehole because the sensors are rotating. For a typical rotation rate of 100 rpm the tool will rotate approximately 10 times for every 10 cm of log, even at a high rate of penetration.

4.5 PASSIVE LOG MEASUREMENTS

Although most logging tools work by putting some form of energy into the formation, a few just passively measure a property. Examples include temperature, calliper, spontaneous potential and natural gamma ray logs.

4.5.1 Temperature Logs

A temperature log shows how the temperature of the borehole fluid varies with depth. They have been available as a stand-alone measurement, particularly for cased holes, since the 1930s but are now often recorded routinely with open hole logs. A knowledge of temperature as a function of depth is important for a number of disciplines but in this book we are mainly interested in temperature because it determines the resistivity of formation water (R_w). Some measurements are also temperature dependent and so temperature is needed to correct them prior to log analysis. Unfortunately, a temperature log is not generally the same as the true geothermal profile of the well because it takes days our even weeks for the borehole fluid to come to equilibrium with the formation. Normally, equilibrium is reached relatively quickly at the bottom of the hole however, so that the temperature measured there is a reasonable estimate of the true temperature at that depth. Estimating temperature will be discussed later in the book.

4.5.2 Calliper Logs

The calliper measurement is another 'passive' measurement that is only used indirectly in log analysis. Calliper logs measure the radius or more commonly the diameter of the borehole, like the temperature log, stand-alone tools have been available for many years but now it is more commonly measured as part of a complicated logging string. Depending on the particular tool string anything from one to over a hundred separate measurements can be made. The latter,

which are measured using a rotating ultra-sonic transducer, allow a very detailed picture of the borehole geometry to be built. The main application of the calliper log in this book is for log quality control. As already noted many logs have shallow depths of investigation and rely on a smooth borehole to produce a reliable reading. The calliper shows at a glance where the borehole is smooth and where it is too rough to guarantee a reliable reading. In those circumstances it is not possible to 'repair' the log but at least we have a 'health warning' about using anything calculated from it. Roughness is quantified as a property known as 'rugosity', which is simply the rate of change of the calliper reading. Many log analysis packages calculate a rugosity curve automatically to aid with quality control.

The remaining 'passive' logs that we will look at here, tell us something about the rocks themselves rather than the condition of the borehole.

4.5.3   Spontaneous Potential

Spontaneous potential – or 'SP' – is a continuous record of the voltage between an electrode in the well and an electrode at surface. It is one of the oldest logging measurements and more often than not it is acquired with resistivity tools. To measure it the hole must be filled with water-based mud. There are two reasons for this.

1. To measure a voltage we need an electrical circuit, which in turn demands a conductive medium between the electrode in the borehole and the formation.
2. The measurement relies on an exchange of ions between the formation water and the drilling fluid.

The requirement for water-based mud means that in some regions and environments it is hardly used anymore and even where water-based muds have made a come-back the SP has been unavailable for so long that it has fallen into disuse. But there are plenty of other places where it is regarded as a primary measurement of great importance.

Because the SP is a voltage or more precisely a potential difference, the absolute value is not significant. What is of interest is the deflections in the curve at bed boundaries corresponding to sudden changes in potential difference. The magnitude of these varies from a few to over a hundred millivolts (mV). Providing the well has been drilled with water-based mud, deflections in the SP can be used to:

1. Identify permeable beds in a background of impermeable shales.
2. Estimate formation water salinity and hence R_w.
3. Estimate shale volume.

The magnitude of the deflection depends on the contrast in salinity between the mud filtrate and the formation water and it may be positive or negative

depending on which is larger. If the mud filtrate is fresher than the formation water the SP gives a negative deflection. Conversely, a positive deflection is a sign of relatively fresh formation water. Formation water salinity – or equivalently resistivity (R_w) – can be found from the deflection providing the salinity – or equivalently resistivity – of the mud filtrate is known. In practice this involves applying two or three special charts, which are normally found towards the front of a chart book.

The special case of no deflection in front of the permeable beds leads to the convenient result that the formation water has the same salinity and hence resistivity as the mud filtrate. The mud resistivity is routinely measured by the logging engineer and should be reported on the log header.

A good example of an SP response in a sand containing relatively fresh water is shown in Fig. 4.5. The SP deflects by +60 mV in the sand relative to the shale. The sign shows the formation water is fresher than the mud filtrate, the deflection will tell by how much. In fact, a deflection of 60 mV is on the
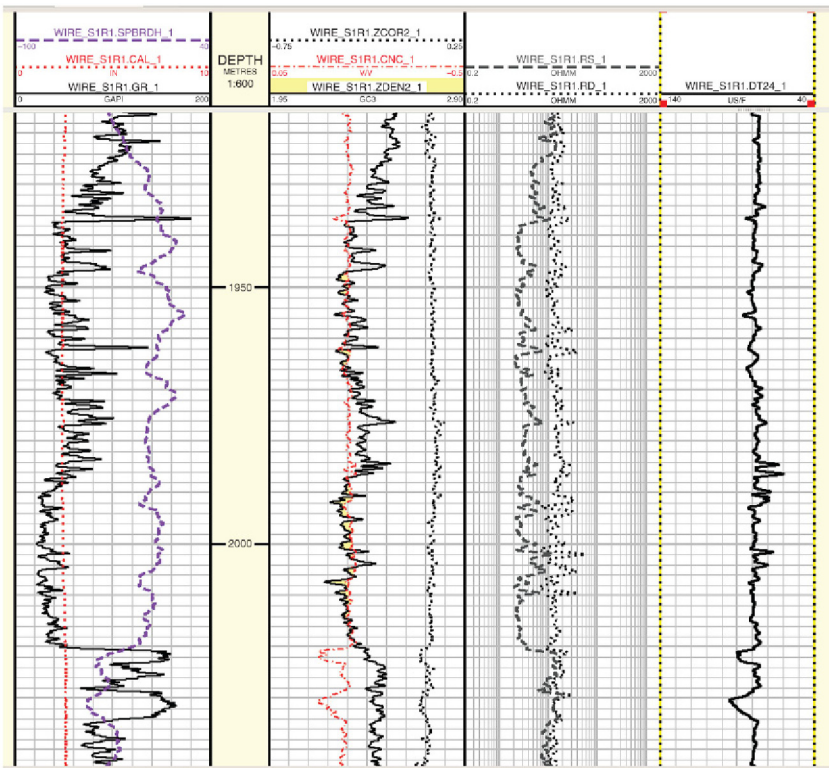

FIGURE 4.5 An example of an SP log and other measurements from a water-bearing sand (the SP is the dashed curve [purple dashed curve in the web version] in track 1). The SP deflects by more than 60 mV to more positive potentials in front of the sand. This shows the formation water is fresher than the mud filtrate.

large side, suggesting there is a big contrast in salinity between the two solutions. In other words the formation water is much fresher than the filtrate.

Unfortunately, there can be a few complications to interpreting SP deflections. Firstly, in order to develop the full deflection the sand must be clean – that is largely shale and/or clay free. If this is not the case the magnitude of the SP deflection is reduced. In fact this can be turned to advantage as a way of estimating shale volume. Secondly, the presence of hydrocarbons in the permeable formation can reduce the magnitude of the SP deflection. This is by no means universal behaviour, however and it is often found that the magnitude of the deflection is the same in water and oil or gas-bearing sands. Finally, the use of potassium chloride (KCl) muds complicates the electrochemical reactions that lead to the SP deflection. The models that are used to explain the origin of the SP assume the formation water and the mud filtrate are both sodium chloride solutions albeit of different concentrations. The potential difference is a result of the sodium and chloride ion concentration gradients between the permeable bed and the borehole. But if the borehole actually contains KCl there are three concentration gradients contributing to the potential difference.

The potential difference is measured by a circuit that is several kilometres long and comprises several different components ranging from simple conductors such as the logging cable to complex and poorly understood media in the sub-surface. The deflections that we are interested in are the result of concentration gradients set up between the mud and the formation water in permeable bed so the depth of investigation is closely related to the invasion depth. Of more interest is the vertical resolution, which is poor compared to almost any other log. Even at a sharp sand-shale boundary the deflection is quite 'lazy' requiring 2–3 m to develop. This means that although the SP can be used as a shale indicator it gives a curve with a low vertical resolution and for sands that are less than a few metres thick it will not be clear whether they are a shaly-sand or just a clean sand that has not been properly resolved (see the base of Fig. 4.5 for an example).

4.5.4   Gamma Ray

The SP may be the simplest measurement to make but the gamma-ray log is probably the simplest to understand. Nevertheless it is still one of the most important measurements. Gamma rays are very high-energy electromagnetic photons that are created when unstable nuclei decay. Traditionally, photon energies are expressed in 'electron volts' and gamma rays typically have energies of hundreds of thousands to millions of electron volts (a large fraction to several mega electron volts). These are very high energies and as a result they are very penetrating and can do a lot of damage when they pass through matter (the photons that make up visible light have energies of a few electron volts). For the types of materials we are interested in, with densities of a few grams

per cubic centimetres, gamma rays can, on average, penetrate between 0.1 and 0.5 m before being completely absorbed. This more or less determines the depth of investigation of the measurement. The penetration distance has to be defined, statistically, as the thickness of material that will absorb a certain fraction of the incident gamma rays. It decreases with decreasing gamma-ray energy. It also decreases with increasing density in a predictable manner. (This explains why nuclear facilities and X-ray equipment are shielded by lead and other dense materials. The use of dense materials means a relatively thin barrier is sufficient to absorb a large proportion of the radiation.)

Natural gamma activity comes almost exclusively from the radioactive decay of three elements: potassium, uranium and thorium. In the case of potassium it is actually only the relatively rare ^{40}K isotope that contributes, this makes up roughly 0.001% of the total amount of potassium in the universe. There are in fact numerous other isotopes of other elements that contribute to the background gamma activity but these are either at such low concentrations or they decay so slowly, that their contribution is insignificant. Uranium, thorium and potassium-40 (K-40) have half-lives that are of the same order as the age of the Earth. This means they decay fast enough to produce a detectable number of gamma rays, but not so fast that they have all but disappeared in the present day. The crustal abundance for these key elements is given in Table 4.1, together with their half-lives.

Gamma activity depends both on the abundance of the gamma emitter and the half-life. In Table 4.1 these have been combined in the column named 'activity' and it can be seen that although K-40 is relatively scarce, it has the highest activity. This is because of its relatively fast decay (if 1.3 billion years can ever be described as fast). So, on average, its activity is actually higher than the more abundant, but slower decaying, heavy elements. The change in the

TABLE 4.1 Crustal Abundances and Half Lives of the Main Contributors to the Natural Gamma Activity

	Abundance weight (ppb)	Abundance molar (ppb)	Half-life (Gya)	Activity
Uranium	1,800	150	4.51	33
Thorium	6,000	540	14.1	38
Potassium (all)	1,500,000	7,800,000		
K-40	180	94	1.28	73
Silicon*	270,000,000	200,000,000		
Oxygen*	460,000,000	600,000,000		

*Silicon and oxygen have been included because they are the most abundant elements in the Earth's crust. Between them they make up about 80% of the total.
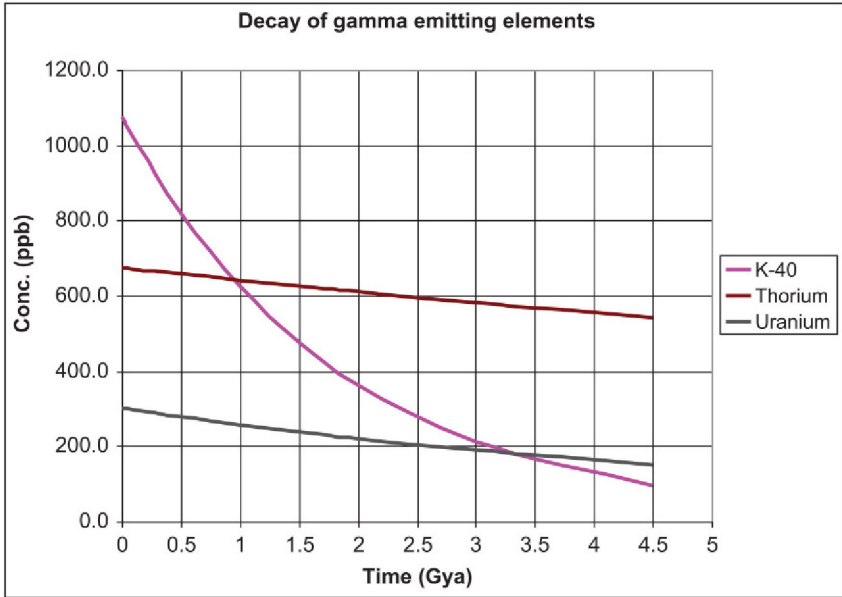
FIGURE 4.6   Decay of K-40, uranium and thorium in the Earth's crust. Note the significantly higher decay rate of K-40. In reality volcanic processes tend to concentrate these elements in the crust so that in the past their levels – in the crust – would have been lower than suggested.

concentrations of the gamma emitters with time is shown in Fig. 4.6. This emphasises the high decay rate of K-40 and shows that in the past it was far more abundant than today. So, early in the Earth's history, heat production would have been an order of magnitude higher than the present day and potassium would have been the major contributor of that additional heat.

The logging tool consists of a gamma-ray detector – very often a scintillator and photo-multiplier combination – together with its associated power supplies and electronics. Because it is so simple, the tool can be made very rugged and, if necessary, it can be packaged in a very slim housing. The tool will work in any type of borehole fluid and can still record a useable curve through several strings of casing (although the absolute activity is reduced by each string). This fact is exploited in one of its most important roles, which is to put cased-hole logs and production equipment on depth with the open-hole logs.

The output of the tool is a count-rate recorded at a particular depth. The count rate depends on the particular tool being used, but it is almost always converted to another unit known as the API or GAPI (API – American Petroleum Institute). The API in turn refers to what the tool should read in a special test well located at the University of Houston. Although two tools may record a different count-rate at a particular depth in a well, they should give the same reading in API. In practice most formations will give readings between a few and a few hundred API and this actually tends to be quite close to the true count

rate measured by the tool in counts per second (in fact you can calculate that on average there are about 200 decays per second in the volume of investigation of a gamma-ray tool although many of the gamma-rays produced by these, will never be detected). In reality the reading is quite strongly dependent on environmental factors such as hole size, mud weight, where precisely the tool is located in the hole (central or off to one side) and possibly mud type. In principle these factors can be corrected for.

The Houston test pit is a good illustration of how logging tools are calibrated. Although ideally every tool would be placed in the pit from time to time to find its unique conversion to API, in practice this is impractical and secondary standards are used to calibrate tools in the field. The normal standard is a weak source that is placed at a precise distance from the detector when the tool is in air. This arrangement then gives a known API reading and the conversion can be found. The use of secondary calibrators for the field is standard practice with all types of logging tool.

Although the gamma-ray tool is very simple, the measurement itself is based on some subtle physical effects, which need to be understood to properly interpret it. The tool normally responds to gamma-ray activity resulting from the decay of naturally occurring radioisotopes. Although it is possible for a gamma ray to travel directly from a decaying nucleus to the detector, they typically follow a more tortuous path. If they collide with an electron they will lose energy and change direction. This happens repeatedly so there is a chance that:

1. The gamma ray will be absorbed completely before it reaches the detector.
2. The gamma ray will end up travelling away from the detector.

Even if it does eventually arrive at the detector, it may pass through undetected. The measurement thus relies on the detector being sensitive enough to give a count-rate that is high enough to be statistically meaningful.

The detector has no directional capability so it cannot distinguish between gamma rays that originate directly in front of it from those that originate some distance above or below it (Fig. 4.7). Obviously, the further from the detector a gamma ray originates, the more likely it is to be absorbed before reaching the tool. In practice, this means the count-rate at any depth reflects the number of gamma rays that originate within a roughly spherical volume surrounding the detector. The radius of this volume is roughly 25 cm for most tools but this depends on many factors including the density of the formation, the borehole diameter, the borehole fluid, whether it is cased or not and where the tool is located in the hole. A consequence of this is that the same reservoir unit in two different wells can give significantly different gamma activity even if it has the same mineralogy. As noted earlier the different environmental factors can be 'corrected' but in practice it is the shape of the gamma-ray curve rather than the absolute value that is used to identify different lithologies or reservoir units.

The decay of K-40, uranium and thorium is responsible for most of the Earth's heat production and so gamma activity also tells you something about
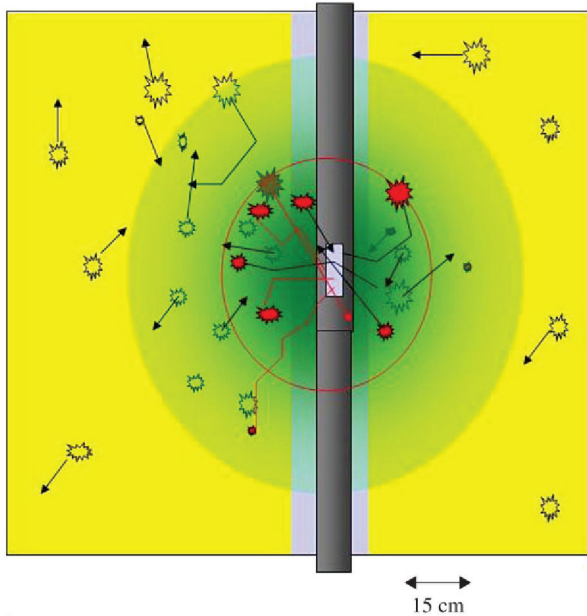
= Radio-isotope decay that is detected. (90% are within the volume of investigation of the tool.)

= Decay that is not detected.

FIGURE 4.7 The 'depth of investigation' of the gamma-ray tool. Unstable isotopes decay all the time but only those that give rise to gamma rays that reach the detector contribute to the response. These events mostly occur within the DOI (by definition 90% of the detected counts originate inside this). Most decays within the DOI go undetected because the gamma rays head off in the wrong direction or pass through the detector unscathed.

the heat budget! In fact, heat production is about the only physical quantity one can calculate from a gamma-ray log in isolation. The power of the tool for more conventional petrophysics relies on demonstrating that the naturally oc-curring radiogenic elements are associated with particular lithologies. Ideally, one would like the gamma activity to be predictably related to the mineralogy at any depth. This is generally not the case but there are general rules that hold true most of the time.

High gamma activity is often assumed to indicate argillaceous rocks such as shales but there is no fundamental reason why this should be so. Of the three elements, only potassium forms part of the chemical structure of any of the commonly occurring sedimentary minerals. It is found in illite, glauconite, mica and orthoclase.

The other elements – U and Th – can be found as trace impurities in some commonly occurring minerals or as part of the chemical structure of minerals that occur in trace quantities. Often these too are associated with fine-grained rocks. One reason for this is that both elements either form or

become concentrated in stable chemically inert minerals like zircon. These often formed billions of years ago and have had plenty of time to be ground to clay-grade material, which is deposited in shales and other argillaceous rocks. Uranium is also often associated with organic matter and can occur at trace levels within some carbonates.

4.5.4.1   Artificial Gamma Activity

Although in most circumstances the gamma ray effectively only responds to the gamma-ray activity produced by naturally occurring radio-isotopes, there are a few circumstances where it is affected by artificial gamma activity. Firstly, in production wells the scale, which precipitates from produced water can be associated with relatively high concentrations of radium (a daughter product of uranium and thorium decay). Radium decays very fast and so small amounts create a lot of activity that can completely swamp the natural activity. Scale build-up only occurs in production wells however, so it is only really a problem for reservoir monitoring logs. The other source of artificial gamma activity in boreholes is neutron activation, which will be discussed in the section on the neutron tool.

KCl muds also give rise to artificial gamma activity. In this case the K-40 in the mud adds to the background activity. The actual shape of the gamma-ray curve remains the same but the overall activity is boosted by the contribution from the mud. The actual contribution from the mud depends on the KCl concentration, the size of the hole and the position of the tool in the hole but it can more than double the activity measured by a logging tool.

4.5.4.2   LWD and Geo-steering

Although the gamma-ray detector that is at the heart of any tool has no inherent directional dependence it is possible to build a tool with an azimuthal dependence. The trick is to shield most of the detector with a dense material like tungsten, but to leave a window so that gamma rays from one direction can still be detected. Obviously only a fraction of the gamma rays that a conventional tool would contribute. But this is a price worth paying if one needs to know in what direction most of the gamma rays are coming from. This information is most useful if the detector rotates and so most LWD contractors have directional gamma-ray tools for a range of hole sizes. In fact most of these tools have several detectors located around the tool body. The information allows the orientation of bed boundaries and fault planes to be inferred relative to the well.

4.5.5   Spectral Gamma Ray

The gamma-ray tools described in the previous sections measure the total gamma-ray count-rate regardless of the gamma-ray energies. But in fact the energy of the individual gamma rays can be measured using the scintillator-photo-multiplier detector at the heart of nearly all modern gamma-ray tools. This is useful information because the energy of the gamma ray produced

by the decay of an unstable nucleus is a characteristic property of the nucleus. For example, when K-40 decays it emits a gamma ray with an energy of 1.46 MeV. The number of 1.46 MeV gamma rays detected at a particular depth can thus be related to the number of potassium atoms in the volume of investigation of the tool.

The scintillator is a cylinder of dense transparent material that emits a flash of light when a gamma ray is absorbed. The photo-multiplier is an electronic device that converts light to an electrical pulse. Dense materials are preferred for the scintillator because they are most likely to absorb the incoming gamma ray completely and they have to be transparent in order for the flash of light to be detected. For many years the standard scintillator was a single crystal of sodium iodide doped with thallium but in the 1990s alternative materials began to appear that were more efficient gamma-ray absorbers (the density of NaI is $3.67 \, g/cm^3$, BGO one of materials that is replacing it, has a density of $7.13 \, g/cm^3$!). Regardless of what the scintillator is made from, it turns out that the intensity of the light is proportional to the energy of the gamma ray. The photo-multiplier in turn converts the light flash to a voltage pulse whose magnitude is proportional to the light intensity so the net effect is that the electrical signal gives the energy of the incoming gamma ray.

In practice the measurement is complicated by the fact that the energy recorded by the tool is not necessarily the energy of the original primary gamma ray. In fact the majority of gamma rays arriving at the tool will have much lower energy because of scattering either in the formation, mud or the tool itself. Even if the gamma ray arrives at the detector without being scattered it may not be completely absorbed, so in the case of a K-40 decay the tool actually records an energy of less than 1.46 MeV. The net effect is that the primary gamma rays that can be related to the concentrations of uranium, potassium and thorium make up only a small fraction of the total count-rate recorded by the tool (one in a thousand or less). In order to get statistically meaningful count-rates spectral gamma-ray tools use scintillators that are many times larger than those used in a standard gamma-ray tool. Even then some sophisticated signal processing is used to try and extract more information from the scattered gamma rays.

Furthermore, the measurement becomes less accurate at high temperatures. It is also more difficult to build a spectral tool in LWD form because the large amount of steel in the tool absorbs and scatters most of the primary gamma rays. The recent improvements in detector technology however, have allowed at least one contractor to offer an LWD spectral gamma-ray service.

Chapter 5

Logs Part II: Porosity, Resistivity and Other Tools

Chapter Outline

5.1 INTRODUCTION

In this chapter we continue to look at well logs but the tools discussed in this chapter work by putting some form of energy into the formation and measuring the formation's response. Although more complicated than the simple 'listening devices' that were discussed in Chapter 4, these tools tend to have a more predictable response to a particular lithology and are often less sensitive to environmental factors such as hole size. Amongst other things this makes these tools better suited to quantitative log analysis and nearly all of them can be used to estimate porosity and/or saturation. This is not to say that they render tools such as the gamma ray unnecessary but typically if one or more is available the interpretation of the gamma ray response becomes a lot more reliable.

As with Chapter 4 the intention is not to repeat descriptions of the tool physics that have been covered well in other publications. Unfortunately however, for at least some measurements, standard accounts of the tool physics are oversimplified and lead to misconceptions, which ultimately result in inaccurate and even incorrect interpretations. Where appropriate the accounts in later sections will address these. Nevertheless the main intent here is to describe exactly what

the measurement refers to in terms of the size and shape of the volume of investigation and the limitations on the measurement.

In Chapter 1 brief mention was made of the similarity between a petrophysical interpretation and analytical chemistry. In both cases the aim is to unravel a complicated mixture. Analytical chemists however have the advantage of being able to work in a laboratory and in principle at least can use the most sensitive technique for a particular component. By contrast, in a borehole there are only a limited number of measurements that can realistically be made and the log analyst is forced to work with those. The principle limitation is the type(s) of energy that can be used. These need to interact strongly enough with the formation that they produce a useful measurement but not so strongly that they are almost completely absorbed before they have any chance of being detected. In practice this means:

1. certain parts of the electromagnetic spectrum (either much longer or much shorter wavelengths than visible light)
2. neutrons
3. sound (at roughly audio frequencies).

Significantly, these are all electrically neutral. There are a large number of very specific analytical techniques that use high-energy charged particles, but it is virtually impossible to generate these in a down-hole tool. Similarly, analytical chemists regularly exploit infra-red, visible and ultra-violet light to make simple but very reliable determinations. Unfortunately, this part of the spectrum is absorbed so strongly that it would barely penetrate the film of mud that inevitably gets between the tool and the formation. The net effect of all this is that tool designers do not have much to 'play with' and nearly all logs in use today had appeared in some form by 1960.

5.2 DENSITY TOOLS

Density can be accurately measured by logging tools. Although the tool exploits atomic physics to make the measurement, it does measure something that is very close to the true formation density (i.e. the ratio of mass to volume). Density contributes to the seismic response of the formation and so it is important for relating wells to seismic sections. Its most important role however is to estimate porosity and this is basically accomplished by re-arranging Eq. 1.6. Before considering its role in log analysis however it is helpful to understand how the tool works and what its characteristics and limitations are.

The density log uses the fact that the scattering and absorption of gamma rays increases with density. They therefore exploit the same principles as a medical X-ray machine. Note that X-rays and gamma rays are fundamentally the same, they are both high-energy forms of electromagnetic radiation and the distinction is mainly to do with how they are generated. X-rays being produced on demand, using a special generator and gamma rays are produced by the decay of radioactive isotopes. On average gamma -rays have higher energies than X-rays

but there is considerable overlap and for our purposes they behave in the same way when they interact with matter.

When X-rays/gamma rays interact with matter it is specifically with the electrons and the more electrons there are, the more likely it is that the gamma ray will be scattered. It turns out that to a good approximation the number of electrons per unit volume is proportional to the density of the material. So by doubling the density, you double the electron density and the number of scattering events. The net effect is that all other things being equal there is a very good correlation between the count rate of gamma ray/X-rays and the density of the material under-investigation. This is how a medical X-ray works: a generator produces a source of X-rays of known energy and intensity which pass through the patient and on to a film that darkens in proportion to the number of X-rays hitting it. If dense bone lies between the source and the film relatively few X-rays avoid being scattered or absorbed but if soft tissue is in the path then more get through.

Several changes have to be made to the equipment to create a viable logging tool. Firstly, the X-ray generator used in medicine is far too bulky to put in a logging tool so a 'chemical' source is used instead. The source consists of a sealed metal vessel containing a small quantity of Cs-137, which is an unstable isotope produced as a by-product by nuclear reactors. It emits 'gamma rays' which actually have a higher energy than the output from a typical X-ray generator although lower than the natural gamma rays discussed in Chapter 4.

The primary gamma rays produced by the decay of Cs-137 have an energy of 662 keV whereas most medical applications use energies in the range 100–200 keV. Cs-137 decays with a half-life of 37 years in other words about a billion times faster than the naturally occurring gamma emitters. Very little natural Cs-137 is left in the crust.

A more fundamental distinction between the logging tool and a medical X-ray, is the way the source and detector are arranged. In a medical system the X-rays travel through the sample and on to a photographic plate ('the detector,' see Fig. 5.1a). But this is not possible in a borehole so there the gamma rays are scattered back to a detector that sits adjacent to the source (Fig. 5.1b). The rule that high-count rates occur with low densities and vice versa still applies, although it may not be as easy to appreciate why this is the case. A plot of count rate versus density is given in Fig. 5.2. This plot is unique to the particular source and detector used and their relative positions. If any of these factors change a new relationship between counts and density has to be determined. In practice logging tools are calibrated at least monthly and after any major repairs, so that the count rate to density transform should always be up to date.

From the point of view of measurement accuracy, the stronger the source the better, as this will reduce measurement statistics and completely overwhelm the natural background activity. Unfortunately, HSE considerations make it impractical to use a strong source and in practice the weakest sources that are still consistent with an accurate measurement are used. The sources are actually
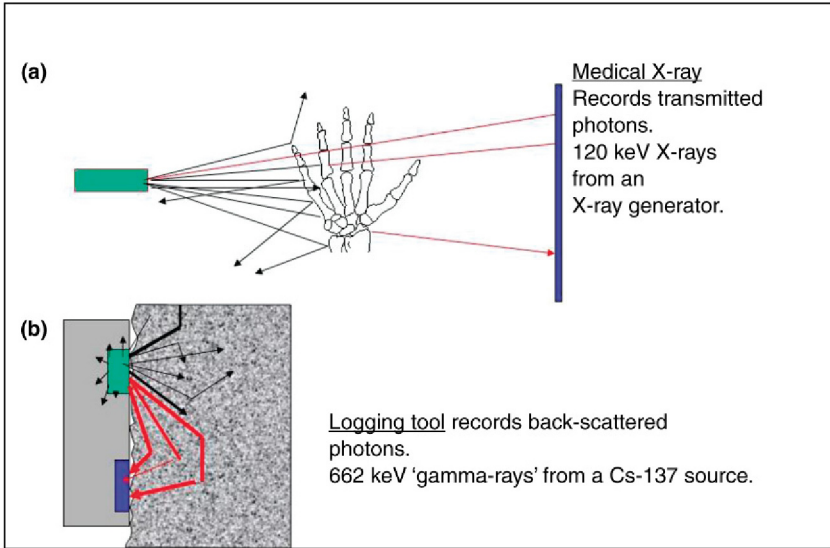
FIGURE 5.1 (a) The density log uses the same basic principle as a medical X-ray photograph but making the measurement in a borehole imposes some changes to the basic equipment. (b) The X-ray generator is replaced by a shielded gamma-ray source and a different configuration of source (light grey rectangle [green in the web version]), target and detector (dark grey rectangle [blue in the web version]) is required. Nevertheless the basic rule that the count rate falls with increasing density still applies.
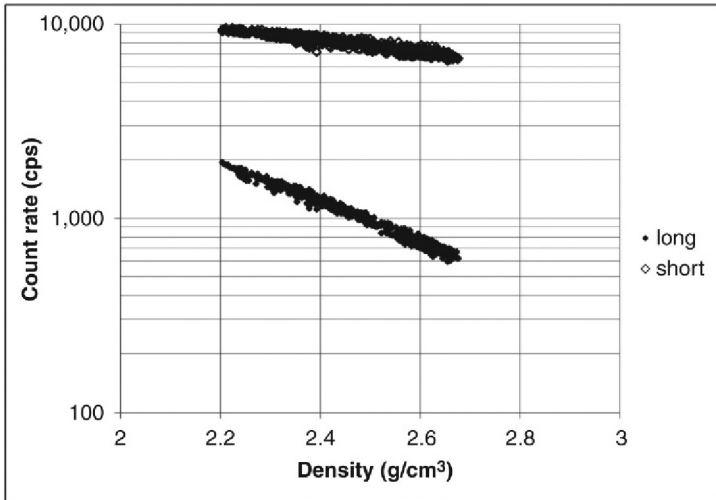
FIGURE 5.2 A plot of count rate versus density for a wireline logging tool. Note the count rate falls exponentially with increasing density.

quite feeble but still produce count rates that are orders of magnitude higher than the natural background activity.

Figure 5.2 shows the count rate falls logarithmically with increasing density. This means that for high-density rocks the count rate can be a factor of 10 lower than for low-density ones. This has some consequences for the accuracy of the measurement; it will be less accurate in denser rocks!

In practice density tools actually use two detectors, this gives some redundancy in the event of a detector failure but the main reason is to give a more accurate reading in cases where the density is not uniform. The reason can be seen in Fig. 5.3, which is a sketch of the business end of density tool. This consists of a heavy skid that is pushed against the borehole wall. The skid contains the gamma-ray source and the two detectors located at different distances from it. The weight of the skid is due to extensive use of tungsten shielding which is used to absorb all gamma rays except those coming out of the front of the tool. The detectors are also shielded so that only gamma rays returning from the formation are actually detected (the same idea is exploited in directional gamma-ray tools). The distance between the source and the detector roughly determines the average distance into the formation a gamma ray travels before it is scattered back to the detector. The larger the separation, the greater the depth of investigation.

In a homogeneous formation the two detectors should give the same density but if the pad does not lie flush against the borehole wall, mud will get between the formation and the tool. The inclusion of mud in the volume of investigation
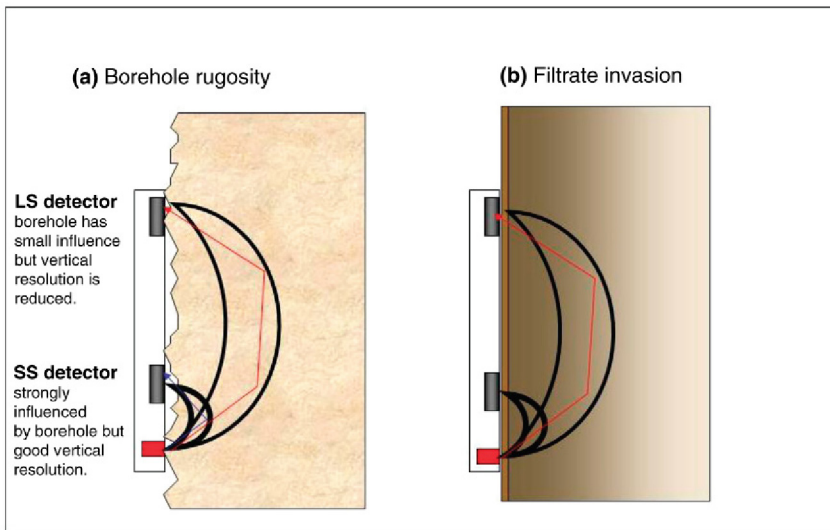

FIGURE 5.3   A sketch of the pad of a typical density tool. The pad contains the gamma-ray source and two detectors located at two different offsets from it. The different depths mean that on average the detectors have different depths of investigation.

of the tool means the density measured by the tool will under-estimate the true density. Note however that there will be proportionately more mud seen by the short-spaced detector than the long-spaced one and the former will therefore record a lower density. This situation is shown in Fig. 5.3a. The difference in density is output as curve on the log that is generally referred to as the 'density correction'. This is actually used to correct the measurement for the mud so the output from the tool is actually derived from both detectors. Providing the correction is not too large the effect of the mud can be accounted for and the tool gives a reading, which is the true density of the formation. There is a limit to how much mud the tool can correct for however and if it becomes excessive the measurement is basically useless and cannot be recovered. The correction is therefore a very useful quality control (QC) check on the measurement. As a rule of thumb a log with a correction greater than $0.075\,g/cm^3$ should be treated with caution.

A similar situation occurs in Fig. 5.3b where although the skid is now pushed flush against the borehole, it is separated from it by the mud cake. Furthermore, invasion by filtrate has replaced some of the formation fluid close to the borehole wall. Again the density seen by the near detector will be different to the far one (it may be higher or lower depending on the relative densities of mud cake, mud filtrate and formation fluids). The tool again attempts to correct for the density difference.

The effect of borehole roughness – or rugosity – and the use of the density correction as a QC tool are well illustrated for the short section of logs shown in Fig. 5.4. In this case the calliper shows the borehole wall is becoming quite rough below 3100 m. The density suddenly reduces from 2.6 to $2.7\,g/cm^3$ to less than $1.95\,g/cm^3$. At the same time the density correction (DENC) increases from about zero to $0.1–0.2\,g/cm^3$. This suggests the measurement is useless most likely because in the enlarged and rugose borehole the density pad can no longer make good contact with the borehole wall.

5.2.1 Vertical Resolution and Depth of Investigation

Recall that the depth of investigation typically depends on the source-detector spacing. So to see deeper one needs to move the energy source and the detector further apart. In the density tool this comes at a price:

1. The count rate reduces with increasing separation.
2. The vertical resolution of the tool is determined by the source-detector spacing.

The count-rate reduction could be overcome by using a stronger source, but as noted earlier there are practical limits to this imposed by HSE considerations. The vertical resolution is however a genuine trade-off with depth of investigation. In the case of the density tool the long-spaced detector is typically placed 35–50 cm from the gamma-ray source. A quick look at Fig. 5.3 shows that this
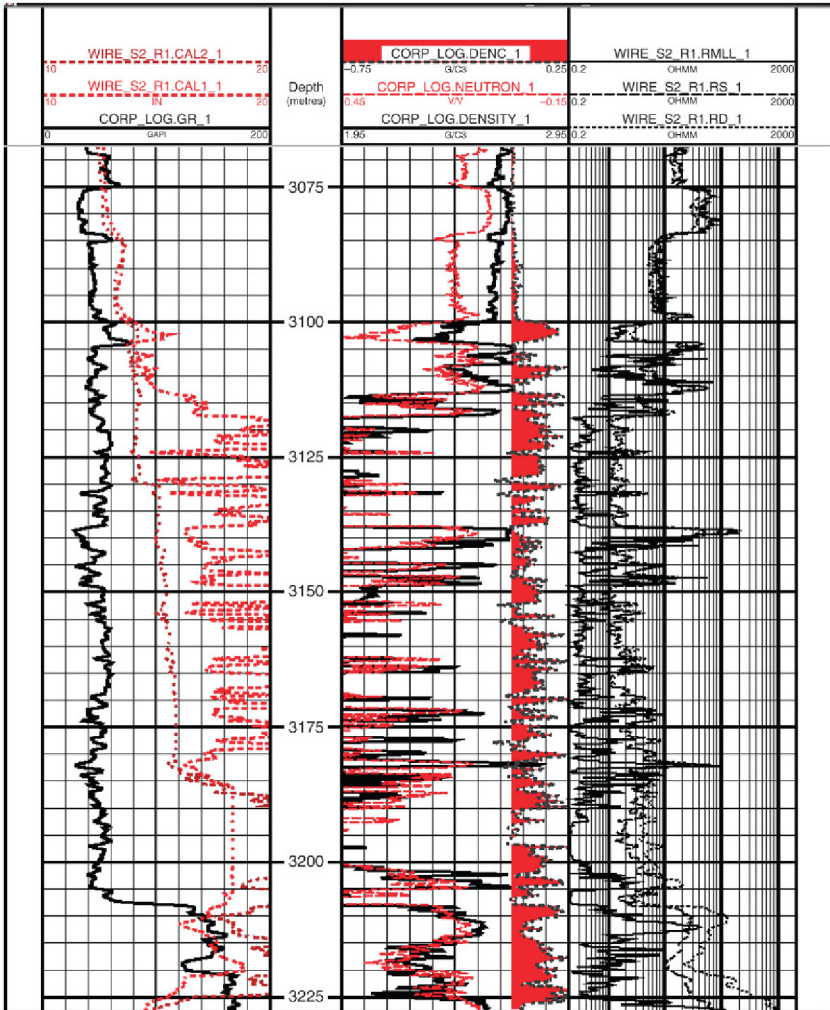
FIGURE 5.4   An example of the effect of 'bad hole' on the density log. Below 3100 m the hole becomes large and rugose (rough). The density pad can no longer lie flush against the borehole wall and for much of the remaining interval the measurement is useless. The density correction curve shows where the density is unreliable (shaded dark grey [shaded red in the web version]).

determines the thinnest bed that the tool can properly resolve. In fact the vertical resolution will be worse than this because the tool accumulates data over a particular interval – 3 in. or 7.62 cm say – and so this is added to the fundamental limit created by the source-detector spacing. Finally, recall that the log is normally filtered to improve its appearance and this further worsens the vertical resolution. The net effect is that the density log struggles to properly resolve anything thinner than 1 m.

For the source-detector spacing given earlier the density tool can read 10–15 cm into the formation. This is largely determined by the long-spaced detector spacing. The short-spaced detector on its own measures density within about 5 cm of the borehole wall. It would give improved vertical resolution but as discussed earlier, the shallow depth of investigation makes it very vulnerable to borehole imperfections (a typical source to near detector spacing is about 15 cm). Single detector density measurements are routinely recorded by modern tools but these are normally only presented if specifically asked for. If one needs a high-resolution density measurement however, the use of the short-spaced detector alone may be an option.

In Chapter 4 a statistical measure of depth of investigation was introduced: the distance within which 90% of the signal originates. This type of information is often expressed as a plot of the cumulative signal versus distance into the formation. Examples for a density and a neutron tool are shown in Fig. 5.5.
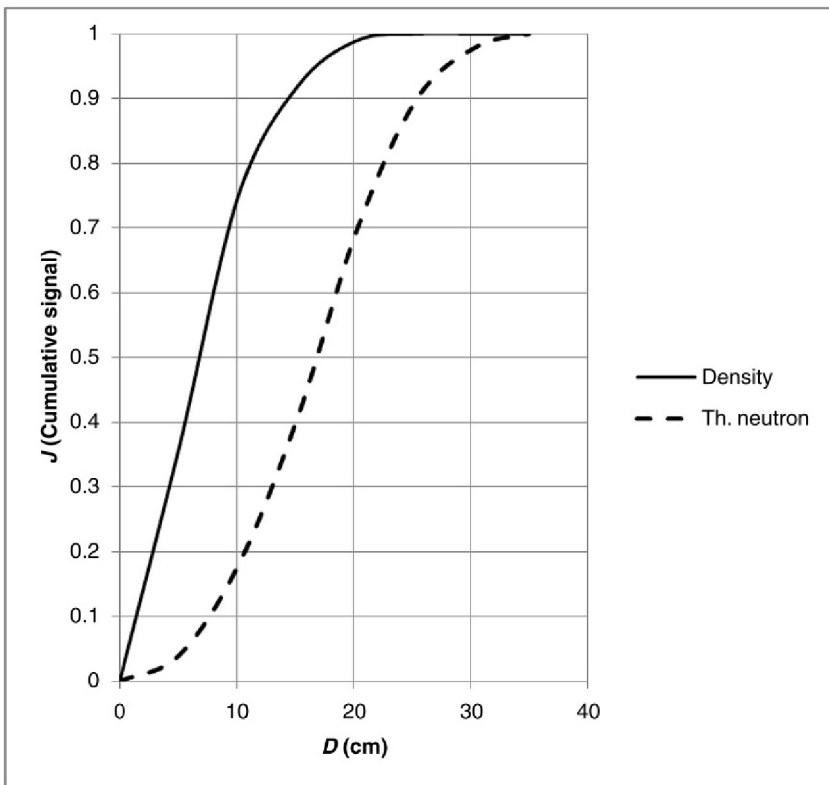
FIGURE 5.5 Depths of investigation for a wireline density tool (solid) in a 30% porosity sandstone. The dashed curve gives the same information for a typical two detector neutron porosity tool.

It should be noted that the depth of investigation for the density tool actually depends on the density of the medium. It increases as density decreases because on average gamma rays can penetrate deeper before being scattered or absorbed (the same reason the simple gamma-ray tool reads deeper in low-density formations).

Knowing how the density tool works we can now draw its volume of investigation. Remembering the shielding is designed to absorb all gamma rays except those that come out of the front of the pad and go straight into the formation, it is apparent that the tool only 'sees' formation in front of the skid. The depth of investigation was discussed earlier and very little signal is received from outside the interval by the source-detector spacing (Fig. 5.3). Taking all this together shows the volume of investigation of the density tool is roughly shaped like a giant orange segment with a length of 35–50 cm and a maximum depth of 15 cm. Radially, the segment roughly covers a 45° arc.

5.3   NEUTRON LOGS

Neutron logging tools include a wide variety of devices that exploit a source of high-energy neutrons. These represent the energy that is used to excite the formation. Like gamma rays, the energies of neutrons are normally expressed in electron volts and high energy means several mega electron volts (in terms of speed this is equivalent to a substantial fraction of the speed of light so they are also sometimes called fast neutrons). Neutrons interact with matter in a number of different ways and different tools exploit different types of interaction. Broadly speaking tools can be divided into two types:

1. those that detect neutrons that have been scattered back towards the tool.
2. those that detect secondary gamma rays that are produced when neutrons interact with the formation.

In both cases the neutrons interact with atomic nuclei, this is in contrast to gamma rays that interact with electrons.

Tools that detect scattered neutrons are analogous to the density tool in as much as they detect the same type of particle that they emit (Fig. 5.6). The raw outputs of the tool are the count rates from the detectors and like the density tool these are converted to something more user friendly, in this case the so-called 'neutron porosity'. That really is where any similarity with the density measurement ends as although neutron logs are often classified together with density tools as 'nuclear logs', the two exploit some very different physics.

As well as interacting with matter in different ways to gamma rays, neutrons typically suffer many more collisions between source and detector and lose proportionately far more energy. In fact the loss of energy is essential to the measurement because the detectors will only count low energy – slow – neutrons (for neutrons 'slow' means about 2 km/s). Furthermore, having lost most of their energy they are not simply absorbed by the nearest atom, but
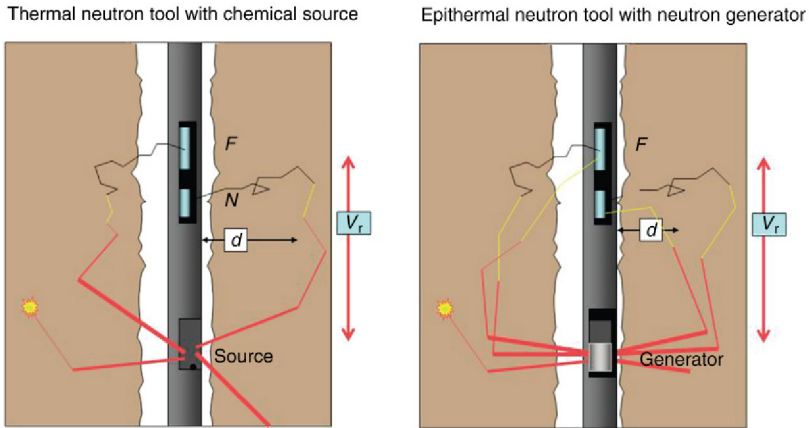
Thermal neutron tool with chemical source   Epithermal neutron tool with neutron generator

FIGURE 5.6 Typical two detector neutron tools consisting of a source and two neutron detectors. The whole tool is pushed against the borehole wall by a 'bow-spring'. High-energy neutrons are emitted from the source and are slowed down by collisions with nuclei in the formation and borehole. After several collisions they are slowed to energies where they can be detected by the two neutron detectors. The tool on the left uses a chemical source that constantly emits neutrons in all directions. The tool on the right uses a generator and produces pulses of high-energy neutrons that initially travel at right angles to the tool. The colours and thicknesses of the lines depict the energy of the neutrons.

continue to diffuse away from the source. The measurement physics is typically modelled as a two-stage process in which:

1. The high-energy neutrons are slowed to thermal energies. It is in this phase that their speed drops from a significant fraction of the speed of light to about 2 km/s.
2. The thermal neutrons diffuse away from the source and form a 'cloud' around the tool. Eventually they will be absorbed.

Thermal neutrons that happen to reach the detectors will be absorbed there and are counted. The slowing down and subsequent diffusion makes the neutron porosity a far more complicated measurement than the density.

A typical tool of this type consists of a source of high-energy neutrons and two neutron detectors located at two points further along the tool. Ultimately, the neutron porosity is calculated from the count rates but the name is unfortunate, as the tool is certainly not measuring porosity (in fact the neutron porosity comes from the ratio of the counts from the two detectors). It might be better if – like the density – the output was a physical property, which could then be further transformed to one or more petrophysical properties. Like the density however, every tool has a unique transform from count rates to neutron porosity and like the density the transform needs to be checked and updated at regular intervals.

The source can either be a radio-active – chemical – source which constantly emits neutrons or a neutron generator which can be switched on and off and

emits short pulses of very high-energy neutrons. The chemical source emits neutrons in all directions equally. Neutron generators emit neutrons in a plane perpendicular to the tool axis, but within the plane there is no preferred direction (they are also on average more energetic and have a narrower range of energies). Neutron generators avoid the HSE issues associated with a radioactive source and actually have some additional benefits (early chemical sources used plutonium making neutron-logging tools particularly unpopular with the CIA!).

The detectors are superficially similar to gamma-ray detectors but they can only detect neutrons with low energies (a few electron volts at most). Since both types of source emit neutrons in all directions the basic neutron measurement investigates a volume with a 360° azimuth. The tool is normally pushed against one side of the hole however, so the volume of investigation is not completely symmetrical with respect to the borehole axis. The depth of investigation is deeper than the density log and for the type of tool described earlier is shown in Fig. 5.5.

The source-detector spacings are chosen to try and make the tool reasonably insensitive to borehole effects (the borehole contains a lot of hydrogen). This means that for most tools the count rates actually fall with increasing porosity (see Fig. 5.7). So unlike the density tool it is actually more accurate at low porosity.

Like the density and the gamma-ray log it is possible to built tools which measure a narrower arc by shielding parts of the detectors but these tools require special detectors that will only detect neutrons of a certain energy (so-called epithermal neutrons). These will be discussed later.

To properly understand, the neutron log requires a good understanding of nuclear physics because the tool exploits interactions between neutrons and the nuclei of the atoms that make up the formation. Simple explanations for its function abound however and most of these emphasise the fact that hydrogen is more effective at scattering neutrons than any other element. This is because, once they have slowed down, neutrons have almost the same mass as a hydrogen nucleus (i.e. a proton). Snooker, pool and billiard players will know that having particles of the same mass is the most efficient way of losing energy in a collision A good snooker player can cause the cue ball to lose all its energy in one collision. But replace the cue ball with either a ping-pong ball or a bowling ball and even the best player will fail to make it lose much energy.

The snooker analogy is good but it is often pushed too far. For a start a bowling ball may not lose much energy in a collision with a snooker ball, but it will lose some. In other words hydrogen is not necessary to slow neutrons down it is just particularly good at it. Secondly, neutrons are more massive when they are emitted from the source because they are travelling so fast (relativity enters into the tool physics). So early on, collisions with heavier nuclei such as carbon or oxygen are likely to be more effective than hydrogen at slowing it down. (If you like the snooker analogy you will have to modify the game to one in which the cue ball gets lighter as the game goes on!)
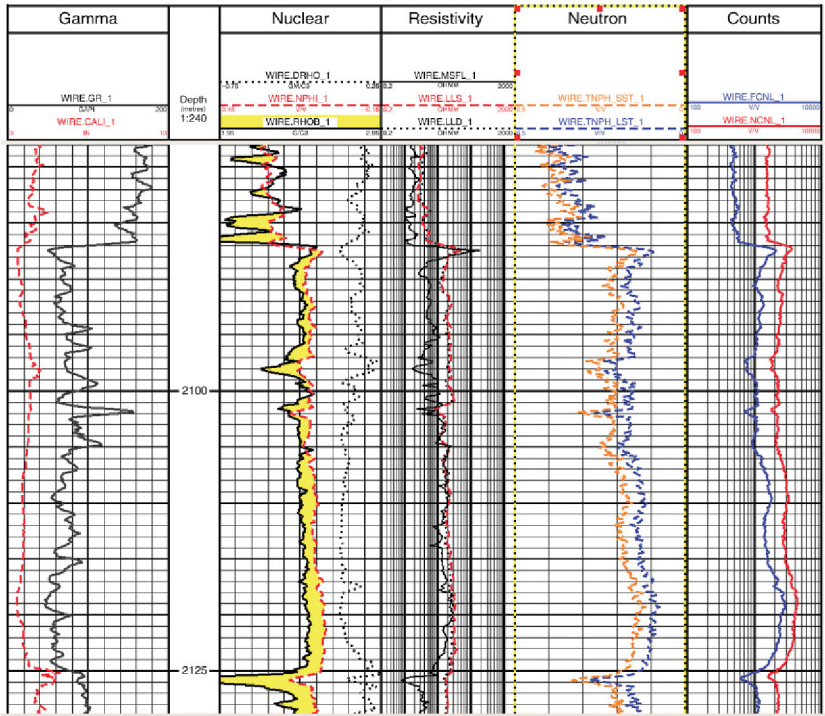
FIGURE 5.7   Neutron log recorded in a clastic sequence, showing how the choice of matrix changes the measurement. The raw count rates from which the neutron porosity is calculated are shown in the right-hand track (note the logarithmic scale). The track to the left (number 5) shows the neutron porosity curves for lime (dark grey [blue in the web version]) and sand (light grey [orange in the web version]) matrix. Track 3 shows the neutron porosity for lime matrix plotted together with the density curve using conventional 'limestone compatible' scales.

The net effect of all this is that the tool does respond strongly to hydrogen, but it also responds to just about everything else. So, for example the tool will read differently in a 20% porosity limestone and a 20% porosity quartz sandstone. In other words neutron porosity is not porosity and if they are equal, it is really just a coincidence (see Fig. 5.7). Having said that, the tool is set up so that in a particular lithology saturated with water, the neutron porosity should equal the true porosity. More often than not the 'special' lithology is limestone but the engineer can select sandstone or even dolomite. We will return to this in later sections because it is an important issue that is often not properly appreciated.

The tool is often said to measure the property known as the hydrogen index (HI) that was introduced in Chapter 1 (recall this is the amount of hydrogen per unit volume of formation, by definition the HI of freshwater is 1). The reasoning behind calling the primary output of the tool porosity is as follows:

1. The tool is designed to be particularly sensitive to hydrogen.

2. Hydrogen mostly occurs in the sub-surface as part of a water molecule.
3. The volume fraction of water is the total porosity (in a water-bearing formation).

Conveniently, the HI of many oils is also close to one so the logic can sometimes be extended to oil-bearing formations as well. Unfortunately, it should be clear by now that each stage of the argument is actually flawed. Firstly, the tool is sensitive to more than just hydrogen and is very sensitive to some elements. Secondly, hydrogen is attached to more than just water molecules (notably clays, gas and coal).

5.3.1   The Neutron Matrix

As noted already the primary output of the tool – the neutron porosity – will equal the true porosity under certain special circumstances.

1. The formation is water bearing.
2. The formation has a particular lithology and is free of clays (shale volume is zero).
3. It has been properly corrected for environmental factors such as hole size and temperature.

The lithology in which the tool reads the correct porosity is known as the 'neutron matrix' or simply the matrix. It is set by the logging engineer and fundamentally it is the algorithm that converts the counts to neutron porosity. In many parts of the world the matrix is set to limestone even if the reservoir is sandstone, dolomite or something more exotic. This means that unless the formation is a clean limestone the neutron porosity will not equal the true porosity (in fact it will under-estimate the true porosity by several porosity unit if the formation is actually a clean sandstone). But this is by no means universal and in the Americas it is quite common to find neutron logs that have been recorded with a sandstone matrix. In that case the neutron porosity will be the same as the true porosity in a water-bearing clean sand. In a limestone the porosity will be over-estimated.

The effect of choosing a different matrix is shown in Fig. 5.7. In track 5 neutron porosities that have been computed using lime and sand matrices are both shown. It is important to note that they are calculated from the same raw data – the count rates – and they are both correct. The example shown in Fig. 5.7 is a water-bearing sandstone with porosities varying from 9% to 15%. The neutron porosity recorded using the sand matrix will be closest to the true porosity although as it happens in this case the neutron porosity tends to under-estimate the true porosity by about 2 pu. This is probably mainly caused by environmental factors.

Summary: The 'neutron matrix'
The output for the tool is a property known as the 'neutron' porosity. This may or may not have the same value as the true total porosity. The measurement is calibrated so that the neutron porosity is equal to the total porosity in a standard

lithology – the 'neutron matrix' – when saturated with freshwater. In Europe, Asia and Australasia, the matrix tends to be limestone, but in many parts of the world it is sandstone and practices do vary between operators. It is essential that the matrix the tool is set up for is known, it should be clearly indicated on the log header and preferably on the curve (unfortunately it is often 'buried' within rows and rows of other logging parameters that are normally only of interest to the logging company).

To give an indication of the potential errors that arise by assuming the wrong matrix, consider a 25% porosity water-bearing sandstone. If this is recorded using a sandstone matrix the tool will read a neutron porosity of 25% but if it is recorded with a limestone matrix it will read a neutron porosity of 21%! Both are correct readings and when they are converted to true porosity they will both give 25%. But if one was not aware of the importance of the matrix one might assume that the value of 21% was the true porosity.

5.3.2 Neutron-Absorbing Elements

There is one other complication that the log analyst needs to be aware of which is a result of the fact that neutrons are strongly absorbed by some elements. If they are present they reduce the overall neutron population and consequently the count rates drop. The tool interprets this as an increase in neutron porosity and the log analyst interprets that as an increase in porosity. Many of the best neutron absorbers are quite exotic and are only likely to occur at trace levels (e.g., boron and cadmium are very effective neutron absorbers). Unfortunately, trace levels can be quite sufficient to significantly reduce the neutron flux reaching the detectors. More abundant elements, such as iron, are often quite effective neutron absorbers as well and what they lack in neutron-absorbing power they make up for in abundance. The neutron porosity of shales is generally high and can reach values of 60–70% or more. This is partly because clays contain some hydrogen in their structure but traces of strong neutron absorbers concentrated in the shales also contribute.

The worst effect of strong neutron absorbers is accounted for by using two detectors with different spacing. The neutron counts at both detectors will be reduced in the presence of neutron-absorbing elements and so by basing neutron porosity on the ratio of the counts rather than a single count rate the effects can be mitigated. Nevertheless, neutron absorbers still tend to increase the neutron porosity. A case in point is that iron-containing minerals are characterised by high neutron porosities.

5.3.3 Neutron Activation

The neutron flux from the tool actually 'activates' the formation. That is to say certain nuclei are exited as a result of being bombarded by neutrons and they decay back to the ground state by emitting one or more gamma rays. This is a

rapid process and can 'swamp' the natural background activity that the gamma-ray tool is trying to measure. Fortunately, the worst effects can be avoided by ensuring that the gamma-ray tool passes over the formation before the neutron does. In practice, this means placing the gamma-ray tool above the neutron for a typical wireline run and below it for an LWD tool. If for some reason the tools are not run in a conventional configuration one should not be surprised to see elevated gamma activities.

So-called geochemical tools actually exploit the effect because the gamma-ray energies are characteristic of the element that emitted them. We will look at these in Section 5.7.

5.3.4   Epithermal Neutrons

In the introduction it was noted that in order to detect the neutrons they have to be slowed down considerably. Most of the neutrons detected by the tools described earlier have been slowed down to 'thermal energies'. This means that they are in 'thermal equilibrium' with their surroundings so that on average they neither lose nor gain energy from their surroundings. The average energy of a thermal neutron depends on the temperature of its surroundings and for typical borehole conditions varies from 0.04 to 0.06 eV (it is proportional to the absolute temperature and this is the reason the neutron porosity depends on temperature). As already noted these energies correspond to speeds of 2–3 km/s.

Some neutron tools are designed to count neutrons that have not slowed all the way to thermal energies. These are known as epithermal neutrons and have energies of a few electron volts (corresponding to speeds of 10–20 km/s). This is achieved by covering the detectors in a layer of cadmium, this absorbs thermal neutrons very strongly but leaves epithermal neutrons largely unscathed. By doing this the tool becomes much less sensitive to the effects of traces of neutron absorbers in the formation which as discussed in Section 5.3.3 can lead to neutron porosities that are far higher than the true porosity (Fig. 5.8).

The price that is paid for using epithermal neutrons is that the count rates will inevitably be lower. This is because epithermal neutrons represent a transitional state on the way to becoming thermal neutrons. An epithermal neutron does not survive long before losing energy to become a thermal neutron whereas thermal neutrons survive until they are absorbed. On the other hand, the epithermal neutron population is constantly being replenished by high-energy neutrons losing energy and their short lifetime does mean that they have less opportunity to diffuse away from the tool. Epithermal tools have been available commercially since the mid-1960s and the latest tools demonstrably avoid most of the problems that complicate the conventional – thermal – neutron logs. Furthermore, they can be designed to have similar volumes of investigation to the density tool, which can be a major advantage in heterogeneous formations.
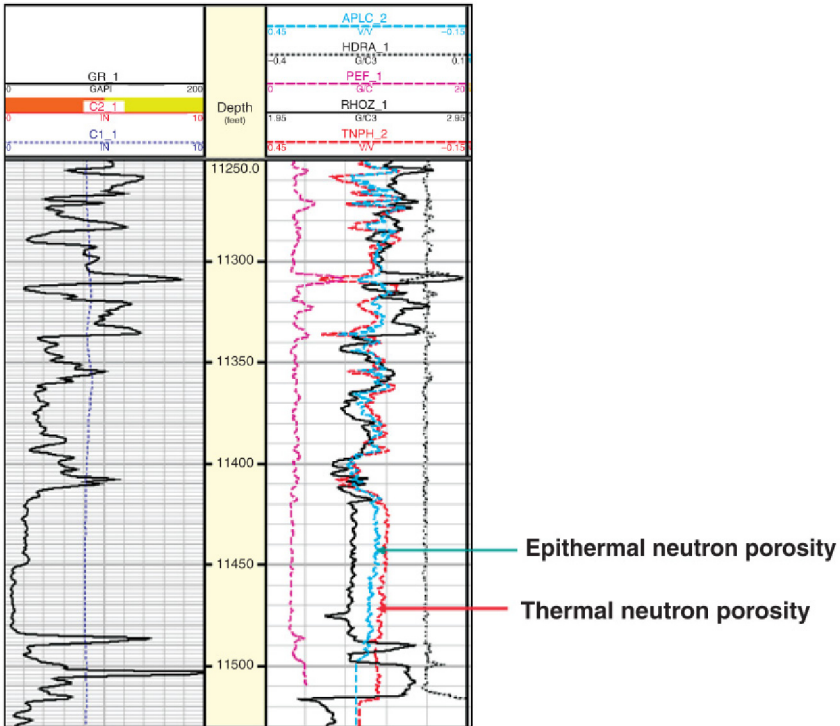
FIGURE 5.8 A comparison between epithermal neutron (middle curve (light grey)) and thermal neutron (dark grey [red in the web version]) porosities. The peaks in the neutron porosity between 11,300 and 11,350 ft. are caused by the presence of iron-containing minerals. The iron reduces the number of thermal neutrons by absorption and this leads to an increase in neutron porosity. The epithermal neutron count rate is not affected.

5.3.5   Neutron Logs: Conclusion

Despite its complications the neutron tool has some important virtues:

1. It will work through casing, this is because a significant proportion of the low energy neutrons pass through the steel casing.
2. It is deeper reading and therefore less sensitive to bad-hole than the density log, although it still needs a good hole to work.
3. Some tools do not use a radioactive source and can therefore be run in wells where one cannot afford to lose a source.
4. Tools are normally designed so that the highest count rates correspond to the lowest porosities so that they are most accurate in the low porosity range (cf. the density tool).

Density and neutron logs were one of the earliest combinations of tools to be run together. Nowadays, combining tools is routine and is typically done to

reduce rig-time. But the density–neutron combination was always recognised as particularly useful for distinguishing different lithologies or the presence of gas and the curves are still plotted together in the same track. It should be noted however that the volumes of investigation of the two tools are different.

1. Neutron reads 20–30 cm into the formation against 12.5 cm for the density.
2. Neutron is basically a 360° measurement against 45–90°.
3. Vertical resolution is worse because the detectors are positioned further from the source.

This means in heterogeneous and anisotropic formations the two tools may be looking at significantly different average properties (see Chapter 2).

5.4 SONIC

The sonic is one of the most versatile logs although, arguably, its most important use is to provide the link between the surface seismic and the borehole. So it is of particular importance to geophysicists (processing and interpreters). It is also used by petrophysicists as a porosity tool and to determine mineralogy (when used with other logs). In recent years acoustic measurements have become increasingly important for determining rock strength *in situ*. This is vital for predicting borehole stability and sanding potential and the latest tools to enter service have been designed to address these areas. Other applications include pore pressure prediction and fracture characterisation.

In order to perform all these functions the sonic tool has evolved further than any other log since its introduction in the late 1950s. Nevertheless, the basic principle of exciting the formation with a short pulse of audio frequency sound and recording the resulting wavetrain produced by the borehole and near wellbore region has not changed. The main advances in technology are as follows:

1. A steady increase in the number of transmitters and especially receivers.
2. Digital recording of the wavetrains from each receiver.
3. Introduction of different types of transmitter (of increasing complexity).
4. Introduction of directional transmitters and receivers so that anisotropy can be detected.

The sonic log has been available in LWD form since the 1990s. The latest wireline tools are really more of a kit of parts, which can be assembled to give a tool that is tailor-made for a particular application.

A modern sonic tool is shown in Fig. 5.9 – such a tool whether LWD or wireline consists of:

1. One or more transmitters that emit pulses of audio frequency sound. The pulses are emitted at regular intervals of the order of a second.
2. An array of receivers located at least a metre from the transmitters.
3. A 'low velocity housing' to prevent the sound energy that passes through the tool body from arriving before the formation signals (and therefore overwhelming them).
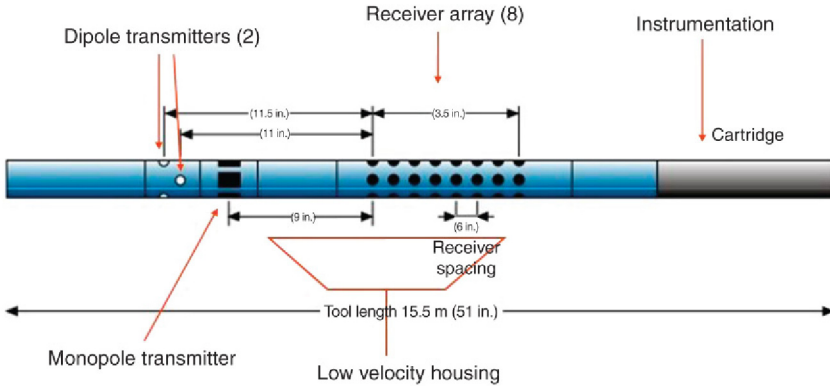
FIGURE 5.9   A modern sonic tool. The tool consists of a transmitter producing audio frequency pulses, an array of receivers and the necessary instrumentation to transmit data to and receive instructions from surface. Wireline and LWD tools are essentially the same but the latter use fewer receivers (4 or 5 vs. 8+).

4. The electronics to control the tool and to convert the outputs from the receivers to a form suitable for transmission. (In the case of LWD the electronics in the form of a micro-processor also calculates the velocities.)

The primary output of the tool at the well site is the compressional velocity of the formation in the near well-bore region. Most tools will also provide a shear velocity but the reliability of this depends on the tool type and the nature of the formation it is run in. As a general rule wireline tools are better than LWD tools in this respect and tools work better in relatively hard-consolidated rocks. Providing the wavetrains have been recorded the data can be re-processed to improve the log and produce additional information (e.g. the so-called Stoneley wave velocity).

By convention sonic log curves are actually given as a slowness, which is the reciprocal of velocity. The symbol for slowness is Δt – 'delta-T' in words. The normal units for slowness are either μs/ft. (micro-seconds per feet) or μs/m. Typical slowness values for sedimentary rocks vary from 140 to 600 μs/m. Note that the lowest values correspond to the highest velocities and would only be encountered in something such as an anhydrite bed. The highest values might be encountered in a marine claystone, which has undergone a small amount of compaction but little else or a gas-bearing, unconsolidated sand.

To convert from slowness to velocity use the following formulae:

$$\text{Velocity (m/s)} = 1,000,000 / \text{slowness (μs/m)}$$
$$= 304,800 / \text{slowness (μs/ft.)}$$

Sonic logs are normally plotted on a linear scale, which increases from right to left (in other words the fastest formations plot towards the right).
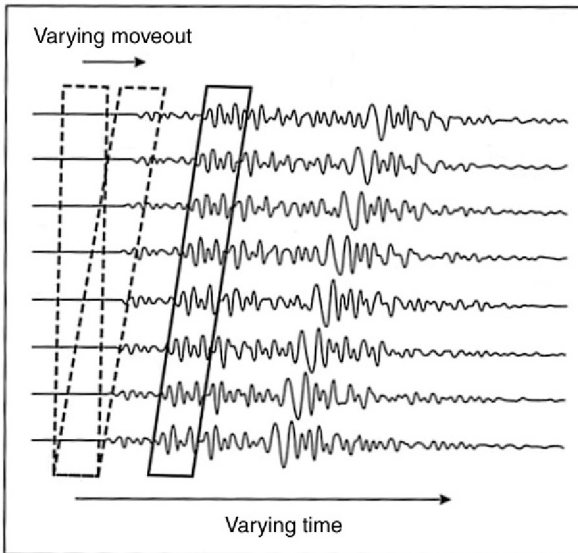
FIGURE 5.10   Primary output from the sonic tool is a series of wavetrains. One complete wavetrain is recorded from every receiver each time the transmitter fires. Each wavetrain consists of several milliseconds worth of data and is made up of compressional waves in the early part of the train, followed by – larger amplitude – shear waves and finally tube waves and direct arrivals from the mud.

The raw output from the tool consists of a series of wavetrains that individually resemble a single seismic trace (Fig. 5.10). Each receiver produces a wavetrain every time a transmitter fires. In a wireline tool the wavetrain is digitised and transmitted to surface. In an LWD tool the wavetrain is stored in a downhole memory. Typically, several milliseconds (ms) of data are recorded in each wavetrain. This means an enormous amount of data is generated by these tools.

The wavetrain may resemble a seismic trace but the likeness is superficial. Whereas a seismic trace is the result of sound waves interacting with a number of reflectors, a sonic tool normally only deals with one reflector (the borehole wall). When the transmitter fires a spherical wavefront is produced in the mud. This moves away from the tool at the speed of sound of the mud (normally 1.5–1.8 km/s). As soon as the front hits the borehole wall it sets up secondary waves within the formation. This also propagates in all directions including parallel to the borehole wall. Because the formation is solid it can support compressional and shear waves so both types of wave spread into the formation but at different speeds. In addition more exotic waves are also set-up at the mud-formation boundary. As all these wavefronts move along the borehole wall they generate more wavefronts in the mud which propagate towards the borehole axis and ultimately reach the tool's receivers. At some point the original disturbance created directly in the mud will reach the receivers as will the signal that has been transmitted along the tool body (Fig. 5.11). By this time there is not a lot of point recording any more of the wavetrain.
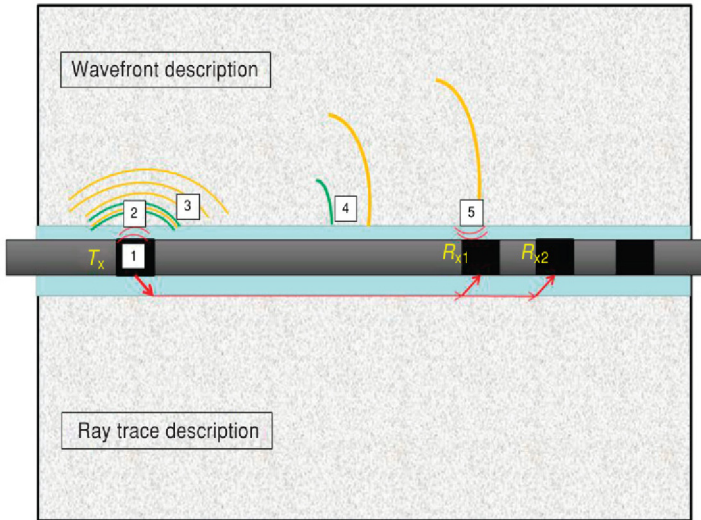
FIGURE 5.11 Operation of the sonic tool. A transmitter (1) generates audio frequency pulses which cause spherical wavefronts (2) in the mud. When these strike the borehole wall they set up secondary wavefronts (3) in the formation. Both compressional (light grey [yellow in the web version]) and shear (dark grey [green in the web version]) fronts are generated. These spread away from and along the borehole axis (4). At the borehole wall the secondary wavefronts set up further disturbances in the mud (5), which ultimately are detected by the receivers (R_{x1}, R_{x2} etc.).

Notice that the process of generating formation compressional and shear waves continues all along the borehole wall as the original disturbance reaches further. The net effect is that instead of receiving three pulses, corresponding to the formation compressional and shear waves and the surface wave, the receivers pick up a complicated and continuous waveform (Fig. 5.10).

Until the 1980s most of the information in the wavetrain was discarded. The tools used a system known as 'first arrival detection', which only allowed the compressional velocity to be measured.

Modern tools compare the complete wavetrains from each receiver (or at least 'windows' of data that include several peaks). This effectively cancels out the noise that occurs at different points within the individual wavetrains. Various algorithms exist to do this processing, but all work by 'sliding' one wavetrain relative to the next until the best match is obtained. These techniques are normally referred to as 'semblance processing' and involve an enormous amount of repetitive computation.

The time difference needed to achieve a good match is directly proportional to the velocity. In fact the velocity is given by the time difference divided into the separation between the receivers. A good match will be obtained at a time corresponding to the compressional velocity and another good match will be obtained at a larger shift corresponding to the shear velocity. Because modern

tools have an array of at least four receivers, there is plenty of redundancy in the measurement, which assists in reducing the effects of noise so, in practice, the whole receiver array is used to calculate the slowness values (LWD tools have a relatively small number of receivers, wireline tools have eight or more).

The quality of the picks can be checked by quantifying the closeness of the match between the wavetrains from the different receivers. The better the match, the more reliable the velocity that is calculated. This is routinely displayed on logs as a 'coherence' curve who's value quantifies the closeness of the match and/or using colour coding in a dedicated track. For the latter, the scale refers to the shift that is applied between the waveforms from different receivers. When no shift is applied the match should be very poor but as it increases the match should get better and better until the value matches the compressional slowness of the formation. Further increases in the shift should then lead to a poorer match. The 'goodness' should therefore pass through a sharp peak with a high value. By colour coding the 'goodness' value both these features can be checked by the end user (Fig. 5.12) Parts of the log where:

1. the goodness of fit is low and/or
2. there is a wide-range of shifts over which a good match is obtained should be regarded as unreliable.

Compressional and shear slownesses can be measured through casing provided there is good coupling between the casing and the formation. The first arrival will be from the casing but this has a very predictable slowness so can easily be identified. The next part of the wavetrain will be from the formation compressional arrival and this can be characterised using normal semblance processing. In practice good coupling means a good cement bond. If this is not satisfied, little energy can reach the formation and of course even less arrives back at the tool, it is therefore best to log in open-hole if at all possible.

The way the sonic tool works makes defining its depth of investigation quite difficult. Various 'rules of thumb' exist such as one wavelength or the more memorable but probably less accurate:

'depth of investigation in inches is roughly the transmitter-receiver spacing in feet' (for a typical tool that amounts to about 3 in. or 7.5 cm).

The fundamental problem is that when the mud arrival is detected the formation disturbance will have travelled quite a long way into the formation. At that point the part of the wavefront moving away from the borehole may be a long way from the borehole but it still influences the part of the wavefront travelling along the borehole wall.

One thing that is certainly not true is that the sonic barely investigates beyond the borehole wall. This belief, which is still stated, comes from applying the 'ray-trace' description, which is also shown in Fig. 5.11. This is a good example of a model that may help explain some of the tool principles, that has been pushed too far.
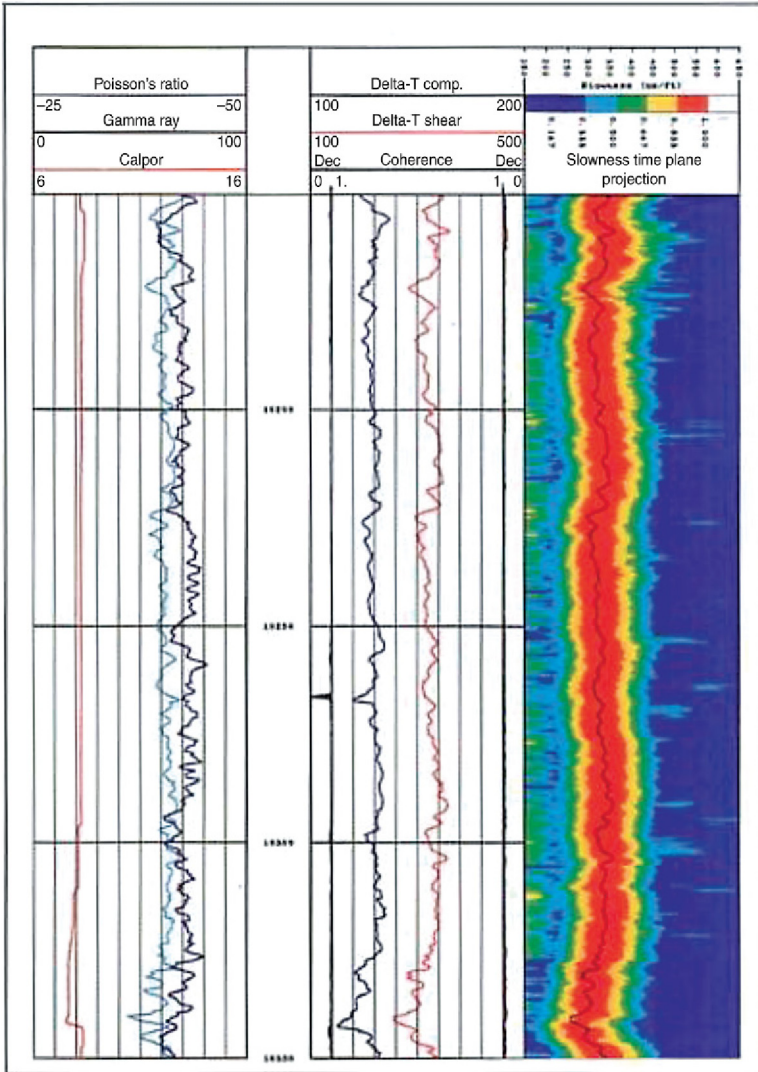
FIGURE 5.12 A modern sonic log plot. The compressional and shear slowness values are shown in track 2 (note the different scales and the un-conventional presentation of slowness increasing L–R). The very colourful track on the right-hand side is a QC track showing how confident the computer is in the shear arrival pick. Central medium grey (red in the web version) indicates high confidence and dark grey (blue in the web version) very low. In this case the pick is reliable over the whole log interval.

More certain is the azimuthal distribution of the measurement, which is determined by the design of the transmitter and receivers. Until the 1980s these were isotropic devices that emitted and detected sound from all directions (known as 'monopole sources'). Most modern tools have the option of highly

directional transmitters and receivers to allow anisotropy to be detected and characterised. This is particularly important for modelling borehole stability, which in turn is important for designing high-angle wells and completions.

The vertical resolution of the sonic tool is often assumed to be equal to the spacing of the receivers, which is typically about 10 cm. Although at first sight this seems to make sense, a look at any modern sonic log suggests this is optimistic. There are various reasons for this including the way the signal(s) are processed but the fundamental problem is again trying to apply the ray-trace description to something that is more complicated.

5.5   NUCLEAR MAGNETIC RESONANCE

Nuclear magnetic resonance (NMR) has been available in commercial logging tools since the early 1980s but only became a reliable service with the advent of NUMAR's MRIL tool 10 years later (NUMAR was a company set up specifically to develop and market an NMR logging tool, they were eventually taken over by Halliburton). Now all the large logging companies offer at least one type of NMR tool. There are also LWD versions of the tool and laboratory instruments that make the measurement on core plugs. The tool responds almost exclusively to hydrogen and since hydrogen is normally encountered as either part of a water molecule or a hydrocarbon molecule the tool responds almost exclusively to porosity. The composition of the matrix is almost irrelevant so the tool basically measures porosity independent of lithology (contrast this with the neutron).

The tool excites the formation with a pulse of electromagnetic radiation similar to that used by LWD resistivity tools, the principle difference is that the measurement takes place in a strong magnetic field generated by the tool. The physics of the measurement is difficult, but all the analyst needs to know is that the tool not only detects hydrogen but it can provide information on the chemical environment of the hydrogen as well. This means, for example the tool can distinguish hydrogen that is part of a water molecule bound to clay from hydrogen in an oil molecule that is free to move around in a large pore. This is in contrast to the neutron log which also responds to hydrogen but is un-affected by the environment the hydrogen is in. In fact the flawed logic that was applied to the neutron measurement to demonstrate that it measured porosity in Section 5.3 actually comes close to working for a modern NMR tool. To paraphrase:

1. The tool is designed to be particularly sensitive to hydrogen.
2. Hydrogen mostly occurs in the sub-surface as part of a water molecule.
3. The volume fraction of water is the total porosity (in a water-bearing formation).

In the case of NMR the tool really is only sensitive to hydrogen, so if only water is present it will measure the total porosity. As will be seen later hydrogen that is part of the clay structure (or coal) is not detected so the relationship between hydrogen and total porosity only starts to fail when hydrocarbons are present (and even then only gas and certain types of oil). Furthermore the ability

to distinguish different microscopic environments means that NMR can actually determine clay-bound water and hence effective porosity.

In summary the physics of the measurement produces three important differences to the neutron log:

1. The tool only responds to hydrogen. This is in contrast to the neutron tool, which to a greater or lesser extent responds to all elements.
2. The tool has a very precisely defined volume of investigation. Different tools have different shapes but they all have non-diffuse boundaries and any atom outside that barrier cannot contribute to the signal.
3. The tool does not respond to the hydrogen that forms part of a clay molecule (Chapter 2). This is again in contrast to the neutron log, which cannot distinguish the environment the hydrogen is in. In particular hydrogen associated with clay is interpreted as water-filled porosity.

As mentioned above the tool contains a permanent magnet that creates a strong magnetic field. This magnetises the formation in the vicinity of the tool and in particular, causes hydrogen nuclei to align themselves with it. When the magnetic field is removed hydrogen nuclei lose this fixed orientation and the magnetisation decays. The hydrogen atoms that are relatively free to move, lose their alignment relatively slowly (up to several seconds). Whereas hydrogen atoms that are in environments where movement is restricted lose their magnetisation much faster (milliseconds). Examples of the former are hydrogen atoms in water molecules that occupy large pores and the latter are represented by water molecules that are bound to clays.

In practice it is impossible to suddenly switch-off the tool's magnetic field as it is created by a permanent magnet. Instead an additional component is added to the field by the pulse of electromagnetic radiation, which can be switched off suddenly. The tool measures the decay of the magnetisation caused by the pulse and analyses it as a sum of individual decays corresponding to different microscopic environments. This information is summarised as the 'T2 distribution', which is the defining feature of the NMR log.

The T2 is the characteristic decay time for a particular microscopic environment and so the complete distribution gives the total amounts of hydrogen in each state. By convention the T2 distribution, which is continuous, is binned into a few specific environments: e.g. clay-bound water, capillary-bound water and free fluid.

A highly idealised example is shown in Fig. 5.13a. In this case the hydrogen is found as part of a water molecule in one of the three possible environments. The first is characterised by tightly bound water and the magnetisation decays very rapidly with a T2 of 10 ms. The second environment consists of small pores in which the water is free to move. The decay is considerably slower with a T2 of 100 ms. The remaining water is held in very large pores and is free to move virtually un-impeded, the T2 is 500 ms. The decay of the magnetisation in each of the environments is shown in Fig. 5.13a. The tool measures the sum of these
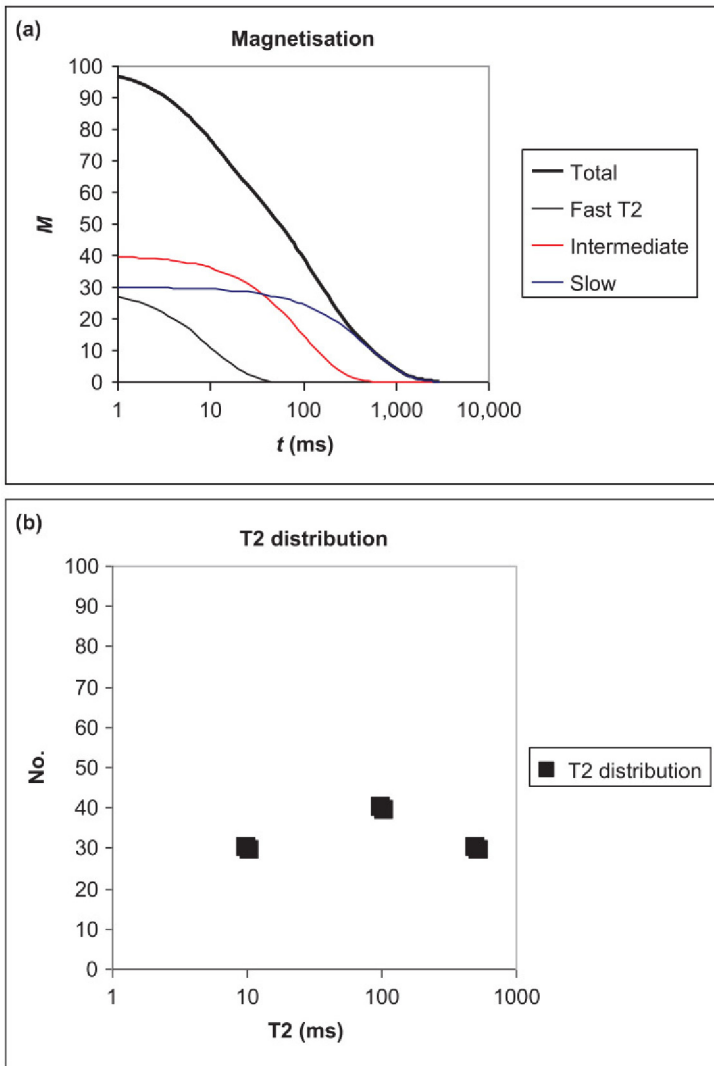
FIGURE 5.13   (a) The decay of magnetisation of a formation containing hydrogen in three differ-
ent environments. Thirty per cent of the hydrogen atoms are in an environment where movement is
restricted and the decay is fast, a further 30% are in an environment where movement is easy and the
decay is slow. The remaining 40% occupy an intermediate environment. The NMR tool measures
the sum of the three decays. (b) The T2 distribution resulting from the analysis of the decay shown
in Fig. 5.12. In this idealised example 30% of the hydrogen magnetisation decays with a T2 of
10 ms, 40% with a T2 of 100 ms and 30% with a T2 of 500 ms.

decays, which is the thick black line in Fig. 5.13a. The information contained in
the decay is presented as a spectrum or distribution of T2 values which shows
the total number of hydrogen atoms in each environment (Fig. 5.13b). In this
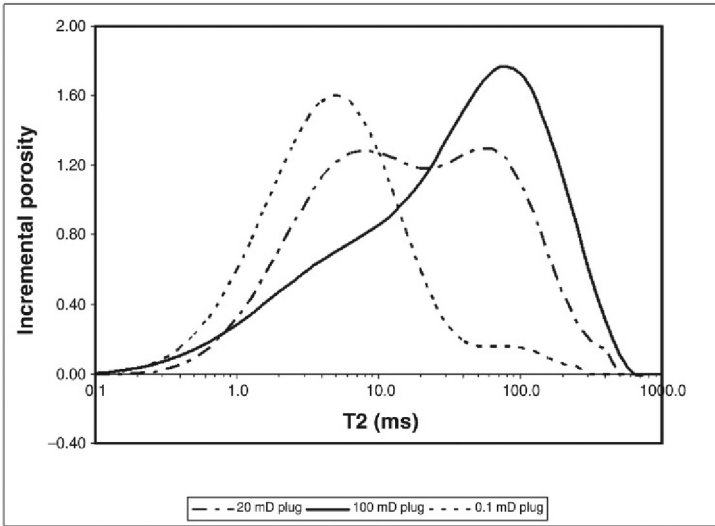
FIGURE 5.14 T2 distributions measured on three sandstones. These were measured using a laboratory NMR tool that uses very similar instrumentation to a logging tool.

case 30% of the water is bound to clay 40% is in small pores and 30% in large pores. In reality the decay will be generated by hydrogen in a continuum of states and the T2 distribution will be a smooth curve (three examples are shown in Fig. 5.14).

The total area under the curve is proportional to the total amount of magnetised hydrogen in the system. If the hydrogen is present as water the area under the curve is proportional to the total porosity. A total porosity curve is one of the standard outputs of the tool.

If some of the water is replaced by hydrocarbon with a HI of less than one, the NMR total porosity will be lower than the true total porosity. The difference is particularly marked with gas because of its low HI and in this the NMR response is analogous to the neutron tool (see Section 5.3). Actually there is a second more subtle reason why the NMR porosity under-estimates the true porosity in gas. Because gas molecules are so mobile the T2 is very long and in the presence of gas the magnetisation may still be decaying when recording stops. So not only is there less hydrogen in gas but some of it may never actually be recorded. (A related problem is that some of the hydrogen in the gas may never have been magnetised in the first place.)

So far it has been implicitly assumed that hydrogen only occurs in the subsurface in the liquid or gas state (as part of a water, oil or gas molecule). In fact, as explained in Chapter 2, clays include some hydrogen in their chemical make-up (this is part of the alumino-silicate structure and is quite distinct from the hydrogen in the water bound to the clay's surface). The hydrogen in the clay is

in the solid state and is actually not detected by the type of NMR measurement exploited in logging tools (and core analysis instruments). In a shale then, the NMR tool only responds to water, this will prove to be very useful for converting between total and effective porosity. It is worth emphasising again that this is in contrast to the neutron log, which responds to all the hydrogen in the formation regardless of its environment. The hydrogen in clay is one reason the neutron tool gives such a high porosity in shales.

For the NMR log T2 distributions are generated at each depth increment and depending on the tool, will have a vertical resolution varying from 0.5 to 2 m (Fig. 5.15). In addition to the T2 distribution the log normally shows a number of curves that are supposed to represent the amount of clay-bound water, irreducible water, effective porosity etc. These are produced by splitting the distribution into a number of bins defined by an upper and a lower limit of T2 values. The precise location of the bins therefore determines how much water is assigned to the clays, irreducible water etc.

The total porosity is obtained from the total area under the curve or alternatively the sum of all the porosity bins.

The log recorded at the well-site uses default values to define the bins but the log can be re-processed using field-specific values if required. The amount of clay-bound water depends on the amount and type of clay present in the formation and so could be used as the basis of a clay indicator.

NMR tools are shallow reading devices and so are adversely affected by bad hole.

5.6 RESISTIVITY

5.6.1 Introduction

Resistivity was the first property to be measured with a wireline tool. Resistivity tools can be designed to read several metres into the formation, often well beyond the effects of invasion and certainly much deeper than any of the porosity tools. Log analysis implicitly uses them to measure the properties of the virgin reservoir.

The danger in interpreting resistivity measurements comes from ignoring the fact that the current is more or less forced to follow a particular path through the formation. In an isotropic medium this does not matter but real rocks are seldom truly isotropic and so the resistivity depends on the direction that it is measured. As an example, consider a thinly bedded formation in which resistive gas sands alternate with conductive shales (Fig. 5.16).

If the current is forced to flow across the beds, the tools response will be dominated by the gas sands and it will read a high value. If it flows parallel to bedding, the current will tend to flow in the shales and the tool will read a low value. Which is the correct reading?
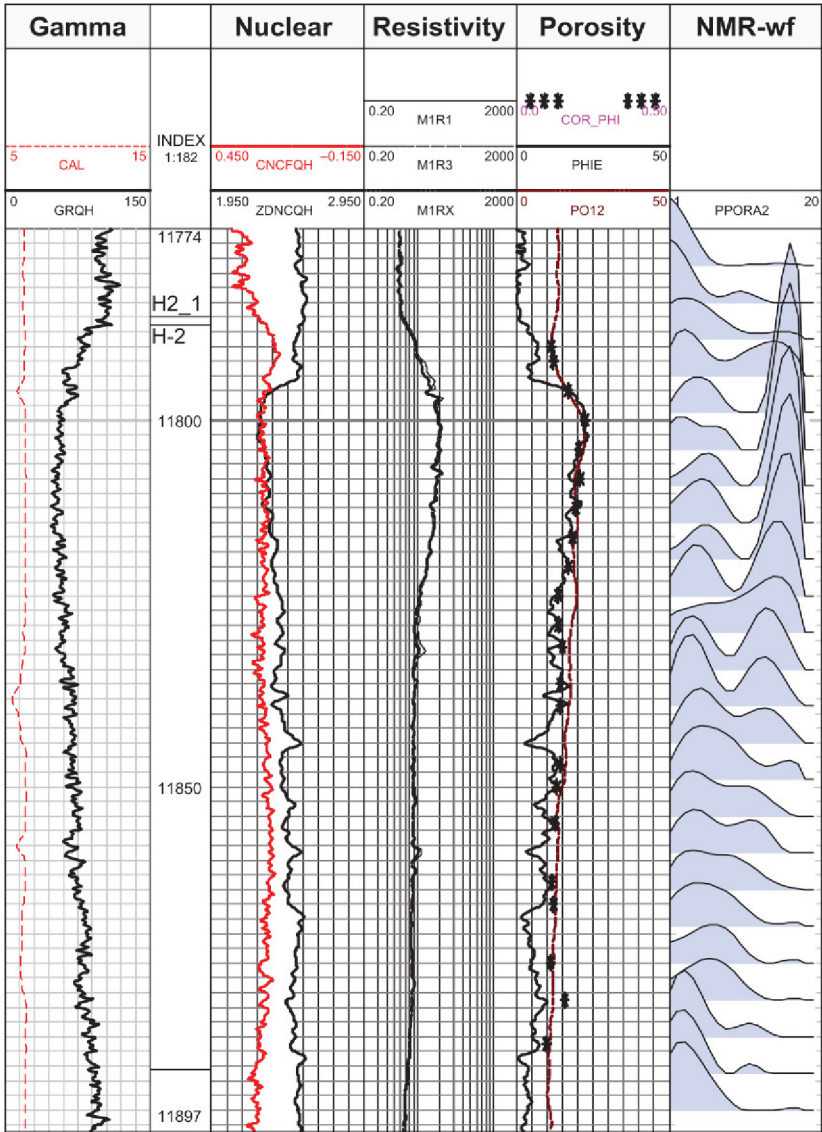
FIGURE 5.15 An NMR log in a shaly sand reservoir. Conventional logs are shown in the left-hand tracks. The right-hand track shows the T2 distribution on a logarithmic scale from 2 to 200 ms. The top of the reservoir is at 11,790 ft. and can be seen as an increase in GR and fall in resistivity. The T2 distribution above that point has a single peak with a short T2 showing all the water is bound to clays. The reservoir has two components in the T2 distribution: the slow T2 peak represents inter-granular porosity and the short T2 is caused by clay and silt. Notice the total porosity (solid dark grey line [blue line in the web version]) falls at the top of the reservoir to the value of the shale porosity.
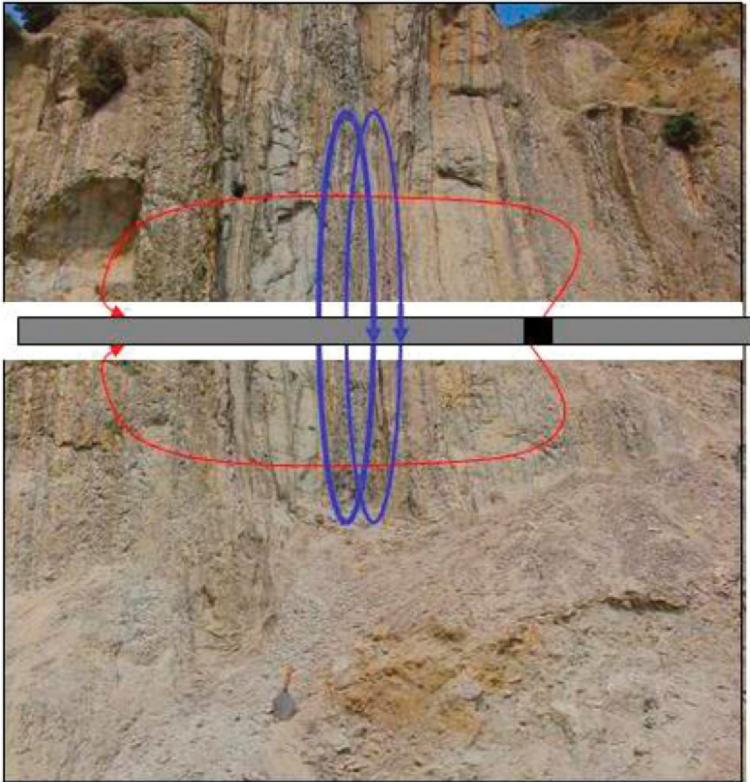
FIGURE 5.16 Current flow for two different types of resistivity tool in a thinly bedded sand-shale sequence. The light grey (red in the web version) current flows perpendicular to bedding and is forced to pass through the more resistive beds. The dark grey (blue in the web version) current flows parallel to bedding and most will flow in the less resistive beds.

Answer: both, because resistivity is a directional property and you need to define the direction it is being measured in. Fortunately, this problem has received for more attention with the advent of high-angle wells.

The unit of resistivity is the Ωm: a unique example of the oil industry using an SI unit where the rest of the world doesn't! Roughly speaking a low value is anything less than $1\,\Omega$m and a high value is anything above $10\,\Omega$m. Resistivity varies more than any other physical property (from zero for a superconductor to $10^{18}\,\Omega$m for some plastics used as insulators). In petrophysics the values we are interested in vary from $0.01\,\Omega$m for saline formation waters at a high temperature to a few thousand Ohm metres in tight limestones. Most modern tools will give readings from $0.1\,\Omega$m to at least a $1000\,\Omega$m, but the results may be suspect above about $100\,\Omega$m. This is not normally a problem for log analysis, as for quantitative work we are normally working at values of $100\,\Omega$m or less.

Because resistivity can vary so much, it is not generally plotted on a simple linear scale. Two commonly encountered ways of coping with the high dynamic range are:

1. To plot on a logarithmic scale (by convention wireline logs use a scale of $0.2–20 \Omega$m or $0.2–2000 \Omega$m depending on the range expected. LWD logs are often plotted on a scale of $0.2–200 \Omega$m).
2. To use a 'times 10 back-up' (or $\times 10$) scale. In this, the resistivity is plotted on a linear scale, from 0 to 20Ωm say, but if it exceeds 20Ωm it goes onto a $20–200 \Omega$m scale.

An example using both systems is shown in Fig. 5.17.

The latter can be quite confusing, particularly if the log has been badly copied. It is well worth checking and re-checking the original log if you find a long oil column that was 'missed' in the original interpretation! By the same token, beware of logarithmic curves being plotted on linear grids or vice-versa.
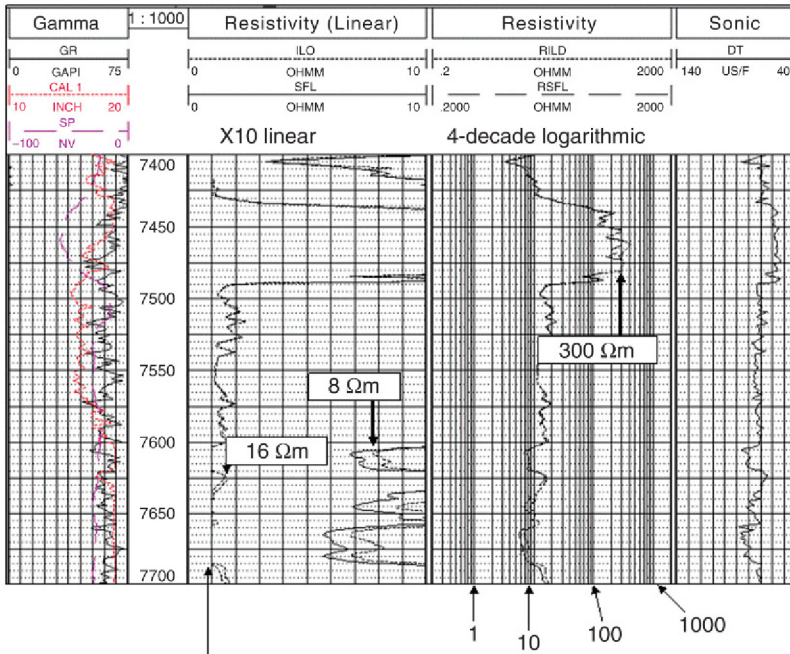
FIGURE 5.17 Resistivity log scales. In this case the scale must cope with a range of values from 6 to 300Ωm (compare this with the sonic in the right-hand track).

5.6.2   Unfocussed Resistivity Tools

The simplest and earliest tools were un-focussed devices in which current was passed between two electrodes and the voltage drop was measured between two others. These devices are extremely simple to build and have no down-hole electronics. The problem with them is that the current path will vary depending on the relative resistivities of the borehole (mud) and formation. In a resistive formation – tight limestone say – most of the current will flow in the borehole and the measured resistivity will be dominated by that. In a highly conductive formation – a porous sand saturated with saline formation water say – most of the current will flow in the formation and the formation will dominate the tools response.

The proportion of current flowing in the formation can be increased by increasing the spacing between the electrodes. With a large enough spacing, one might ensure that between 95% and 99.9% of the current flows in the formation. The response will then always be dominated by the formation, but this is at the cost of very poor vertical resolution. To get good vertical resolution one needs closely spaced electrodes but that, of course, means the signal is generally dominated by the borehole. Un-focussed tools have been almost completely replaced by more sophisticated tools in which the current is forced to follow the same path regardless of the resistivity of the formation.

5.6.3   Focussed Resistivity Tools

The problem with un-focussed devices is that the current path depends on the formation resistivity, mud resistivity, hole size and probably other factors as well. Focussed devices were developed to overcome this major shortcoming by forcing the current to always follow the same path through the formation. They do this by splitting the total current into two parts; the measure current (light grey inside the dashed line [shaded yellow in the web version] in Fig. 5.18) and the 'bucking' current (light grey [blue in the web version]), which are constantly varied to keep the measure current on the same path. Since current paths cannot cross the measure current is effectively 'squeezed' into a particular volume by the bucking current. The currents are constantly monitored, together with voltages at various points on the tool axis and the bucking current is adjusted to ensure the measure current always follows the same path. Current focussing brings with it the cost of complexity, although in it's simplest incarnations it is still possible to avoid any down-hole electronics (tools such as the 'laterolog-3' and the 'laterolog-7').

There are two basic ways of applying the bucking current:

1. It can flow between the tool housing and the measure current, thereby forcing the measure current to flow deeper into the formation.
2. It can flow outside the measure current, thereby limiting the distance away from the tool it flows.
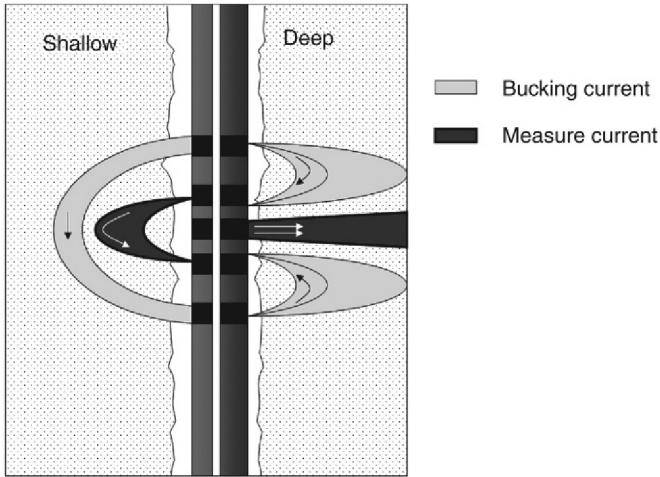
FIGURE 5.18 Current paths for schematic focussed resistivity devices. The bucking currents are constantly adjusted to keep the measure current flowing along the same path. In the shallow device (left) the bucking current flows outside the measure current forcing it to keep close to the tool. In the deep measurement (right) the bucking currents force the measure current to flow deep into the formation before returning to the tool or – in this case – the surface. (Although the diagram shows the shallow and deep measurements on different sides of the hole, this is simply to produce a clearer diagram; in reality these are both 360° measurements.)

The former would typically be used to produce a deep reading device and the latter a shallower reading one. In practice numerous different configurations have been employed over the years, which give depths of investigation from about 25 cm to several metres.

The most widely used focussed resistivity tool is the so-called 'laterolog,' which is a relatively deep reading device. The 'measure' and 'bucking' currents flow out of a single electrode in the middle of the tool and back to different electrodes on the tool body. Nearly all tools currently in service make two measurements with different depths of investigation at the same time. These are invariably named 'deep' and 'shallow' although these terms are relative: the shallow reading is still much deeper than any of the porosity measurements described earlier. It is possible to use the same tool to make the different measurements at the same time by using different frequencies for the two depths. The shallow current passes out of the centre of the tool, turns to flow parallel to the borehole axis before returning to the tool. The deep current is forced radically out from the tool to a distance of several meters before returning to an electrode at the surface.

Using the criteria the depth of investigation is the distance from the tool inside which 90% of the signal originates; a typical shallow laterolog curve reads about 1.5 m into the formation whilst the deep curve reads more than 4 m into the formation. The latest tools, generally referred to as array laterologs, make

measurements at five or six different depths of investigation although they rely on some quite sophisticated signal processing to do this.

Whatever the precise design of the tool, they all require a conductive path to the formation and so in practice will only work in holes filled with water-based mud. Even then, if the mud is very fresh, the tool may over-estimate the resistivity because the current is forced to flow through the resistive mud to get to the formation.

5.6.4 Induction Tools

The induction log was developed in order to measure resistivity in holes filled with air, oil or oil-based mud. The induction log uses a high frequency current flowing in a coil to induce a current in the formation that flows circumferentially around the tool axis. The dark grey current in Fig. 5.16 is an example of this. The formation current induces a secondary current in a receiver coil. The lower the resistivity of the formation, the higher the current that flows and the larger the signal induced in the receiver. Commercial tools actually use an array of subsidiary coils to sharpen up the tool response so that the current path varies little over quite a wide range of formation resistivities, hole sizes and mud resistivities. These tools have evolved considerably since their introduction and now use a combination of hardware and signal processing to produce multiple depths of investigation and good vertical resolution. The deepest reading measurements read 2–3 m into the formation.

5.6.5 Micro-resistivity Tools

Micro-resistivity tools are focussed devices that are designed so that the measure current only penetrates a few centimetres into the formation. If invasion has occurred, their response tends to be dominated by the saturation in the invaded zone (S_{x0}). They can therefore be used to detect invasion and to determine what type of fluid the shallow reading porosity tools are being influenced by. Because they have such a shallow depth of investigation they are very sensitive to borehole roughness and in large washouts may simply read the mud resistivity. An example is shown in Fig. 5.4 (which also shows the effect of bad-hole on the density another shallow reading tool).

5.6.6 Propagation Tools (LWD)

Resistivity is one of the few measurements where LWD tools use a different principle to wireline ones. The reason for this is that because LWD tools are part of the bottom-hole assembly they have to be as strong as any other part of the drill string. In practice this means they are made of metal and so would short circuit either an induction or a laterolog type measurement. LWD resistivity tools actually use a principle that is more akin to RADAR than a conventional resistivity measurement. A typical tool is shown in Fig. 5.19. That particular tool has five transmitters that emit, in turn, short pulses of high-frequency
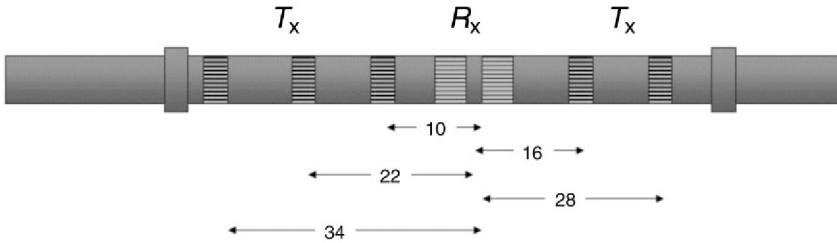
FIGURE 5.19   A schematic of a modern LWD resistivity tool (this is based loosely on Anadrill's ARC series but all modern tools have a similar layout). The numbers refer to the spacing in inches between the transmitter (T_x) and middle of the receiver pair (R_x).

electromagnetic radiation and a pair of receivers. Nearly all tools use a frequency of 2 MHz and most modern tools also use 400 kHz radiation as well. These are much higher frequencies than used by any of the wireline tools (20–40 kHz for the induction, 10–100 Hz for the laterolog).

Different tools have different numbers of transmitters but for our purposes they all operate on the same principle. The use of a number of transmitters allows measurements to be made at different depths of investigations. The usual rule that the wider the T_x–R_x spacing, the deeper the tool reads applies.

The principle of the measurement is shown in Fig. 5.20. Each time a transmitter fires a more or less spherical electromagnetic wavefront moves into the borehole and formation. As the wavefront spreads out it is attenuated in part
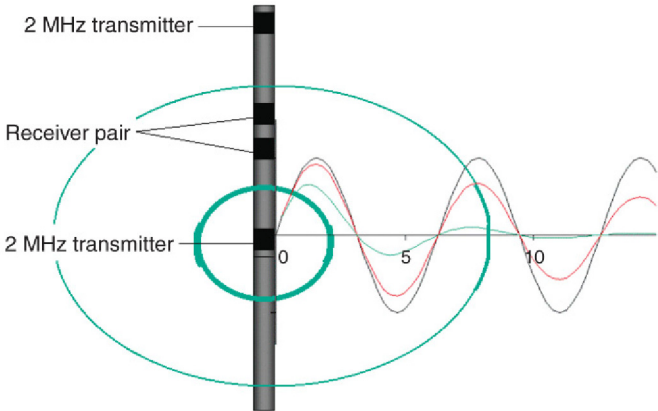
FIGURE 5.20   Principle of the LWD resistivity tool. A pulse of high-frequency electromagnetic radiation is emitted from a transmitter. The wavefronts – ellipses (green ellipses in the web version) – propagate away from the transmitter into the formation. As the wavefront moves away from the transmitter the amplitude is reduced as a result of spreading and losses to the formation (sine waves to the right-hand side of the tool sketch). *In vacuo* geometric spreading is the only factor reducing amplitude and it drops relatively slowly with distance (black curve). In a borehole, energy is lost to the formation and the amplitude drops more quickly with distance. For a resistive formation the losses are smaller and the amplitude drops more slowly with distance (dark grey [red in the web version] curve) than in a low resistivity formation (light grey [green in the web version] curve).

this is simply because it is being spread more thinly but – more importantly – it is also losing energy to the formation. The rate it loses energy depends on the resistivity of the formation, the higher the resistivity the lower the losses. So by measuring the amplitude at a known distance from the transmitter it is possible to deduce the resistivity of the formation.

In addition to the amplitude of the electromagnetic wave changing with distance, the phase shifts in a predictable way with resistivity as well so this provides a second way of measuring resistivity. In general the amplitude technique gives a deeper reading measurement than the phase and some of the simplest tools only provide two measurements (phase and amplitude). The type of tool shown in Fig. 5.19 can provide 5 amplitude and 5 phase measurements and so in principle 10 different depths of investigation. Furthermore, as noted earlier, most modern tools actually emit two different frequencies, which again give different depths of investigation (the lower the frequency the deeper the tool reads).

So we have potentially 20 different resistivity measurements. These are a bit like the gears on a modern bicycle however, where for example there are 3 cogs at the pedals and 6 or 7 on the rear axle giving about 20 different combinations. For most cyclists many of those combinations are almost redundant because they give very similar end results. In the log some of the redundancy is used to determine and then remove borehole effects and furthermore some of the measurements will simply not function in certain environments (e.g. the deepest attenuation measurements are often incapable of measuring resistivities over $100 \, \Omega$m).

LWD tools generally give excellent results but they do have one or two drawbacks compared to the traditional wireline tools. Most importantly the depth of investigation depends on the resistivity. It increases in more resistive formations. As an example, for the tool shown in Fig. 5.19 the depth of investigation of the 34 in. spaced, high-frequency phase resistivity increases from 17in. to 28 in. as the resistivity changes from 1 to $10 \, \Omega$m. The vertical resolution changes slightly as well. The second undesirable characteristic is the phenomenon of 'polarisation', which is normally limited to high-angle wells and is discussed later. These drawbacks have led to attempts to build LWD induction logs some of which have been successfully field-tested.

5.6.7   Horizontal Wells

It has already been pointed out that the resistivity measurement can depend on the relative orientation of the tool with respect to the beds. This issue has received much more attention with the advent of high-angle wells, where it was often found that the resistivity measured in a horizontal well was different to that found in a vertical pilot. There are various explanations for this:

1. Different current paths through laminated formations (discussed earlier).
2. The tools depth of investigation extends into the cap rock or across a contact in a high-angle well.
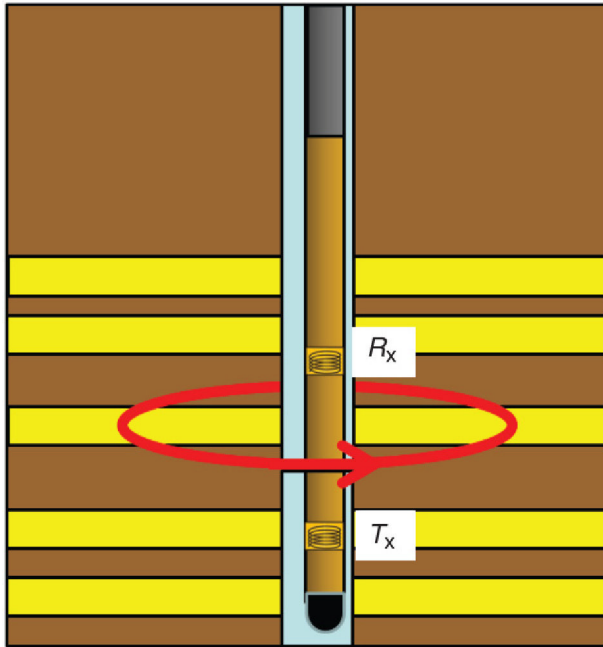3. For propagation tools: polarisation.

FIGURE 5.21   Induction resistivity tool schematic. The current (dark grey [red in the web version]) flowing circumferentially around the borehole is induced by a transmitter coil in the tool (T_x). The current itself induces a signal in a receiver coil (R_x) which is processed to determine the formation resistivity.

Polarisation occurs when propagation tools run parallel or close to parallel to an interface. The electromagnetic pulse emitted by the tool creates a sheet of charge at the interface and this drastically alters the receiver response. The net effect is that the resistivity curves show large spikes near the interface. These are known as 'polarisation horns' and they give no quantitative information. They can be useful however as they provide a very clear indication that the tool is passing close to an interface and so they are often exploited in geo-steering (Fig. 5.21).

5.7   MORE USES OF NEUTRONS: GEOCHEMICAL LOGS

In Section 4.5.4 (gamma ray) it was noted that neutrons from a neutron-logging tool could temporarily increase the gamma activity of the formation. This is a result of interactions between the neutrons and the atoms making up the formation and is known generally as 'neutron activation'. Specifically, neutron activation involves collisions between neutrons and the atomic nuclei. There are several different types of collision.

1. Inelastic Scattering. The neutron collides with the target nucleus and loses some energy. It continues to exist as an independent particle.
2. Fast Reaction. A neutron reacts with the nucleus to produce a different element and one or more fundamental particles.
3. Thermal Absorption. The neutron can be absorbed to create a different isotope (this gives the same element before and after absorption, but the atomic number is increased by one).

Inelastic scattering and fast reactions involve high-energy neutrons (for the tools discussed in Section 5.3 inelastic scattering was actually exploited in slowing the neutrons down to energies where they could be detected). Thermal absorption involves slow neutrons. In all three interactions, immediately after the collision the nucleus is in a higher-energy state but it can decay to its ground state by emitting a gamma ray. These gamma rays have characteristic energies and so by measuring the gamma-ray energy it is possible to infer which element emitted it. The technology to measure gamma-ray energies had already been developed in spectral gamma-ray tools (Section 4.5.5). In Section 5.3 neutron logging tools were classified according to whether they detect neutrons or secondary gamma rays. These tools are clearly examples of the latter, in fact they can be thought of as a spectral gamma-ray tool with its own neutron source.

Several different logging tools now use the gamma–neutron method to deduce the elemental make-up of the formation. Because these tools can determine the 'chemical formula' of the formation they are generally referred to as geochemical logging tools.

The fundamental output of the tool is a gamma-ray spectrum with energies ranging from a bit less than $1\,MeV$ to about $10\,MeV$, one is produced each time the tool moves one depth increment and an example is shown in subsequent section (Fig. 5.22). It can be seen that the spectrum consists of a series of peaks super-imposed on a continuous background of counts, which falls with increasing gamma-ray energy. The latter is made up of all the gamma rays that have undergone several scattering events before reaching the tool and therefore have lost the information as to what emitted them.

In detail the spectrum consists of the sum of the individual elemental spectra plus the degraded background. The individual elemental proportions are found by effectively forward modelling the measured spectrum (Fig. 5.22).

The log that is presented consists of the relative quantities of the different elements and an estimate of the uncertainty on each of these. The relative amounts are normally presented as dry weight fractions (this is what was referred to as 'weight–weight' in Section 1.4). As an example pure limestone ($CaCO_3$) would appear on a log with the following dry weight fractions: Ca (40%), C (12%) and O (48%). These can be converted to relative numbers of atoms if the atomic masses are known (i.e. the chemical formula). Contractors often present additional curves that are derived from the dry weight fractions that are designed to give a quick impression of lithology and mineralogy. So, for example a clay or carbonate curve is often produced. More detailed processing
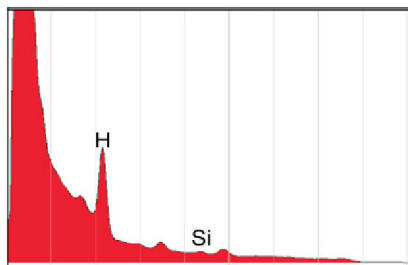
FIGURE 5.22   The gamma-ray spectrum from a geochemical tool. The prominent peaks in this case are caused by hydrogen and silicon.

allows mineral proportions to be estimated using either general or field-specific relationships. An example for a fairly simple sand-shale sequence is shown in Fig. 5.23. Weight fractions of seven key elements are shown in the right-hand tracks and the interpreted lithology is shown in track 2. Notice that aluminium and iron proportions correlate well with the classic shale indicators of gamma ray and neutron–density separation (recall that aluminium is always found in clays and iron often is). Notice also that the tool suggests that even the sand contains nearly 10% clay by weight.

From the relative proportions of the different elements it is possible to make at least an educated guess as to which minerals are present and in what proportions. As a trivial example if at a particular depth one found gamma rays characteristic of calcium, carbon and oxygen with the dry weight ratios given above it could be deduced that the formation consists of limestone. The elements that can be determined by logging tools are, in order of increasing atomic number:

H, C, O, Al, Si, S, Cl, Ca, Ti, Fe, Sm, Gd.

Fortunately, this list comprises most of the elements that make up the common minerals found in the crust. Not all tools can determine all these elements however; the precise list depends mainly on the type of neutron source (chemical or generator). The maximum neutron energy available from a chemical source is not high enough to react with aluminium and so only tools that use a generator can determine it directly. Because aluminium (Al) is a key component of the clay minerals and glauconite this could be seen as major weakness. On the other hand, the rare earth gadolinium (Gd, atomic number 64) can be readily determined by these logging tools at parts per million levels. Most people would consider it an exotic element that is normally only present at trace levels but it is often associated with clay minerals so goes some way to make up for the inability to detect aluminium.

Magnesium is also a 'nice to have' that is absent from the list, as it is part of the atomic structure of dolomite. But in fact some tools can actually determine magnesium.

Geochemical logging tools first entered service in the 1980s as quite bulky devices that were primarily intended to be run in cased holes. They had to be run
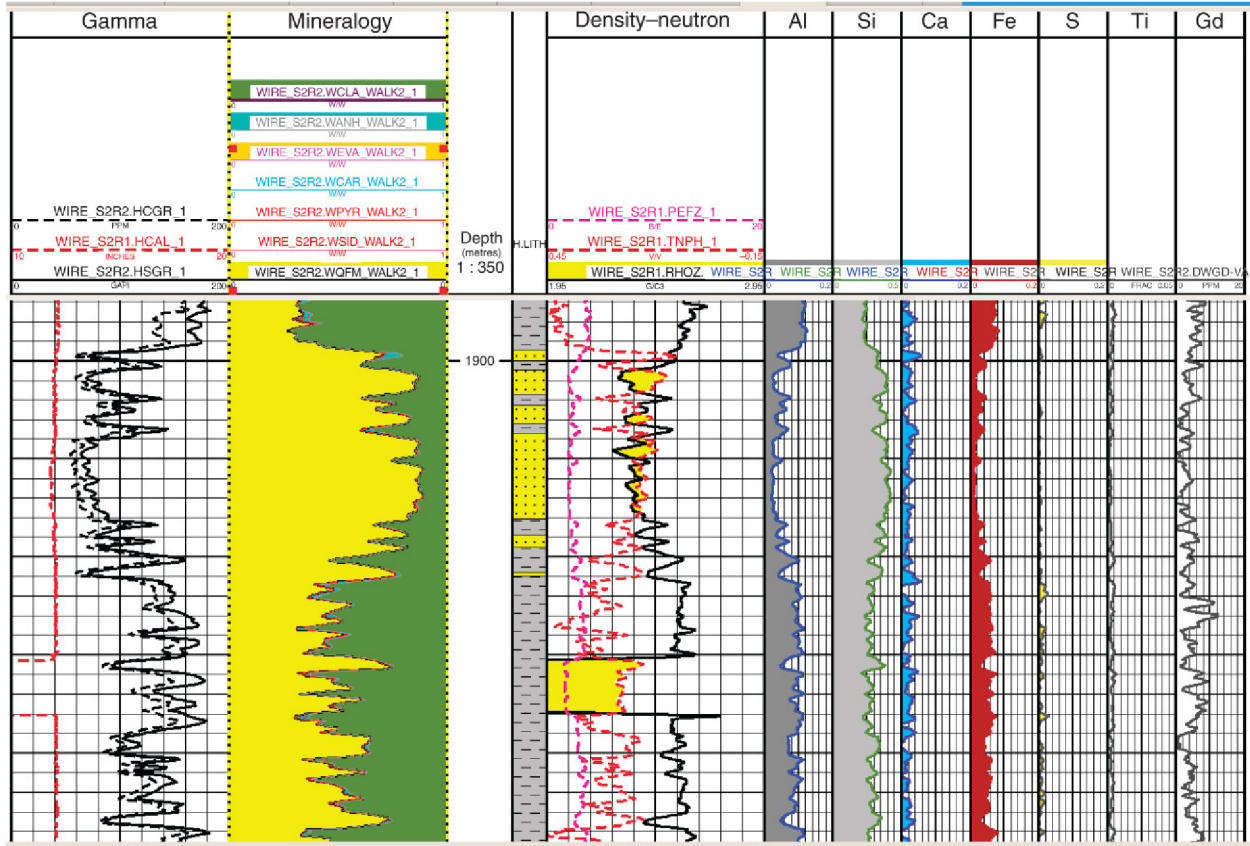
FIGURE 5.23 An example of a geochemical log. Seven of the dry weight fractions are shown in the right-hand tracks. For most elements the scales run from 0% to 20%, for silicon the scale is from 0% to 50% and for titanium it is from 0% to 5%. The rare earth element gadolinium is scaled from 0 to 20 ppm. Lithology indicators are shown in track two: quartz–feldspar–mica (QFM) is coloured light grey (yellow in the web version) and the dark grey (green in the web version) indicates clay.

at very low logging speeds and in some cases were only capable of stationary measurements. They have since evolved into more compact tools that are available in wireline and LWD form and can be combined with many other tools (amongst other things they have benefited from the new scintillator materials). The different tools offered by different contractors do have some significant differences including the type of source they use – chemical or a neutron generator – and the details of how they measure the induced gamma activity.

5.8   ENVIRONMENTAL CORRECTIONS

As a general rule logging tools are designed to give an accurate measurement of the formation within their volume of investigation and to be relatively insensitive to general environmental factors such as hole size, mud type and temperature. In addition specific property measurements are generally designed to be reasonably insensitive to variations in the value of that property within the borehole. So, for example resistivity tools should be reasonably insensitive to the actual value of the mud resistivity. In practice, it is impossible to make measurements completely insensitive to the affects of the borehole and some measurements are affected to a small degree by temperature and/or pressure as well. In order to compare like with like most tools have a set of associated 'environmental corrections', which allow the measurement to be corrected for these environmental factors. Traditionally, these corrections were provided by the logging contractors as charts but now they are normally applied by computer. This can take place either during log analysis or as the measurement is being made.

Most of the corrections have been developed for the tools described in this chapter but the gamma ray can be corrected for hole size, location within the hole and mud weight and potassium level.

There are many misunderstandings associated with environmental corrections, which at best result in unnecessary effort being expended and at worst result in interpretations being based on bad measurements. Here are a few general statements about environmental corrections:

1. They are normally small, often insignificant. If you ignore them, you will not make a bad interpretation.
2. They are based on specific assumptions, which may or may not apply in a particular case.
3. Increasingly they are applied by the contractor as the log is recorded. If the same corrections are applied during analysis the measurement will be less accurate (although see point 1).
4. Tools that are operating outside their working limits cannot be corrected.

The intention of the corrections is to get all logs back to a consistent basis. So, for example if the discovery well was drilled with an 8½ in. bit using a heavy, oil-based mud but the appraisal well was drilled with using a 12¼ in. bit and a much lighter, water-based mud, the environmental corrections should remove the differences due to the hole conditions. The corrected measurements are what would be measured if the tool were run in a 'standard hole'.

As will be seen later the most commonly used type of neutron porosity log is also affected by natural factors such as temperature, formation pressure and formation water salinity. If these are going to be corrected for, a set of standard conditions needs to be defined. For the neutron log these are a pressure of zero, freshwater and a temperature of 10°C. None of these conditions are ever likely to apply in a hydrocarbon reservoir of course, but all the algorithms for converting neutron logs to porosity assume these conditions. The intention is simply to get to a consistent measurement. Having said all of that, remember again the corrections are small and at the end of the day there are often more significant factors affecting the accuracy of our estimates.

As a specific example consider the wireline density log – normally the preferred measurement for estimating porosity. The tool is designed to give a true reading in an 8 in. hole filled with a drilling fluid with a density of $1.05\,g/cm^3$. If the log is actually acquired in a 17½ in. hole filled with a $2.1\,g/cm^3$ mud the measurement will actually be about $0.02\,g/cm^3$ higher than the true value. Although these conditions amount to quite an extreme departure from what the tool was designed to work in, the correction is still less than the inherent uncertainty in the measurement!

The corrections for conventional neutron porosity logs can be relatively large (as corrections go). Neutron porosity is sensitive to a number of factors including borehole size, mud weight, mud salinity, stand-off (the distance of the tool from the borehole wall), temperature, pressure and formation water composition. The magnitude of each individual correction depends on the neutron porosity itself and as a rule they increase with increasing neutron porosity. But in fact the individual corrections tend to cancel each other out and the tools are designed so that in a typical well the overall correction is small. For the example given earlier for the density tool, the corrections for hole size and mud weight add up to about −4 pu for a neutron porosity of 20 pu (in other words a 20% reduction). But the temperature correction alone for typical borehole conditions is +5 pu and the other corrections will add and subtract from this so that – normally – the overall correction will be within 10%.

Resistivity tools can also require relatively large corrections. For wireline tools the largest corrections occur in the largest hole sizes. For induction tools the worst situation is a large hole filled with a very conductive fluid. For laterologs it is large holes filled with resistive fluids that lead to the largest corrections.

If you are going to apply environmental corrections using a commercial log analysis package it is strongly recommended that a few spot checks are made against hand calculated corrections using the appropriate chart book. In our experience little effort goes into implementing or quality checking the corrections in log analysis packages.

5.9   CONCLUSIONS

All logs have limitations on where they can be used, Tables 5.1 and 5.2 give a quick summary of where the basic open-hole logs can and cannot be run. Some of these limitations are universal. For example, a dual laterolog will not work in hole filled with oil-based mud and in fact there is a good chance that the tool

TABLE 5.1 Limitations on Where Different Tools can be Run

Tool type	OBM	Through casing	Air-filled hole	Rugose hole	Large hole
GR	Yes	Yes	Yes	Yes	Yes
SP	No	No	No	Yes	Yes
Sonic	Yes	Yes	No	Yes	Yes
Density	Yes	No	Yes	No	No
Neutron	Yes	Yes	No	Just	Just
Laterolog	No	No	No	Yes	Yes
Induction	Yes	No	Yes	Yes	Yes
Micro-res.	No	No	No	No	No
Dielectric	Yes	No	Yes	No	No
NMR	Yes	No	Yes	Depends	No
Neutron activation	Yes	Yes	No	Just	Just

TABLE 5.2 Summary of Volumes of Investigation for Different Types of Logging Tools

Type	Specific tool	Measure-ment			DOI (cm)	Azimuth (°)	Vertical resolution (m)*	Vertical resolution (m)**
Gamma		Gamma ray			12	360	1	0.5
Acoustic		Sonic			20	360	2	0.75
Resistivity		R_{x0}	▮		8	25	0.5	0.3
	Laterolog	LLS	▮		150	360	1.5	
		LLD	▮		750	360	1.5	
	Induction	Deep induction	▮		600	360	2	
	Propaga-tion	Phase resistivity	▮	3	75	360	1	0.6
		Attenuation"	▮	3	100	360	1	1.2
Density		Density	▮		10	45	0.8	0.3
		PEF	▮		8	25	0.5	0.3
Neutron		Thermal neutron			25	360	1.5	0.3
	APS	Epithermal neutron			17	45	1	0.3
Misc.	CMR	NMR pad	▮		4	20	1	0.3
	MRIL	NMR centred	▮		20	360	1.5	1.5
	EPT	Dielectric	▮		4	20	0.5	0.3

*Vertical resolution defined as the thinnest bed the tool can return a true measurement from (whilst stationary)
**Vertical resolution typically specified by service company
3 Greatest T_x–R_x spacing at $10\,\Omega m$

and the surface equipment will be destroyed if it is attempted. Other limitations are not so strict, for example a good density log can be obtained 0from a washed-out hole providing the borehole wall remains smooth (see Section 5.8). Although making measurements in boreholes is far from ideal and there are many potential pitfalls with logs, the fact is that in most circumstances they do a good job. This is just as well because they provide the bulk of the data that a petrophysicist uses.

Page left intentionally blank

Chapter 6

Introduction to Log Analysis: Shale Volume and Parameter Picking

Chapter Outline

6.1 INTRODUCTION

As pointed out in the introduction to Chapter 4, petrophysical models rely heavily on logs. This not only reflects the fact that logs contribute far more 'bits and bytes' to a typical petrophysical dataset than anything else but is also a result of their 'objectivity' and high – ideally 100% – coverage of the reservoir. Log analysis provides the tools and techniques to convert the physical properties measured by logging tools to petrophysical properties. In this chapter some generalities will be discussed before looking at the specific problem of estimating shale volume. In the next two chapters the general tools and techniques introduced here will be employed calculating porosity and saturation.

6.2 FUNDAMENTALS: EQUATIONS AND PARAMETERS

Log analysis involves converting physical to petrophysical properties. To achieve this, the log analyst has to do two things:

1. Select an equation (or more generally algorithm).
2. Choose the constants that define the particular transform. By convention the constants are referred to as 'parameters'.

As an example if we decide to calculate shale volume from the gamma ray we have the choice of several general equations. The simplest is a linear relationship in which gamma activity is proportional to the shale volume but many workers have found that if gamma activity is plotted against shale volume a curved relationship fits the data better. Some of these relationships have been published and usually several are available in commercial log analysis packages. Five examples are shown in Fig. 6.1.

Once the general equation has been decided it is made specific to the particular well and zone of interest by specifying the parameters. For the particular example of shale volume from gamma ray, the parameters are normally the end
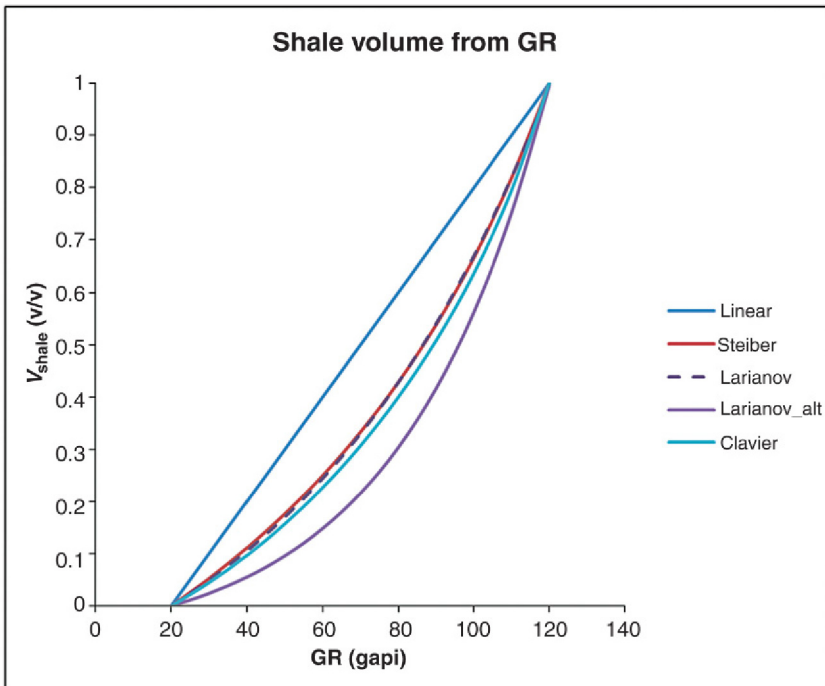
FIGURE 6.1   Some published relationships for finding shale volume from gamma ray. In this case the shale-free gamma activity is 20 api and pure shale is characterised by a gamma activity of 120 api (n.b., these values are simply examples and have no fundamental significance).

points of the relationship (although some relationships require a third parameter that defines the amount of curvature). In other words the gamma-ray activity of the 'clean' – that is shale free – rock and the pure shale. These normally have to be found by noting the gamma activity in a suitable bed (in the case of Fig. 6.1 they are 20 api and 120 api, respectively).

In general, parameters can be found from a variety of sources and apart from log readings in known lithologies, they may be found from laboratory measurements, offset wells, experience or perhaps just read from text books. Normally, only a few parameters can be found directly from logs, this is because most of the components of a petrophysical model do not exist in isolation at the scales the logs work at. But in the case of shale volume it is often possible to find a bed of pure shale so that the log responses can be read directly. The same may be true of the clean – that is shale free – lithology but this should be carefully checked. In the earlier example, it is not enough to simply note that 20 api is the lowest gamma reading in the zone of interest we should confirm that this really does correspond to a bed of clean formation that is thick enough to give a true reading.

If the gamma-ray log had been environmentally corrected the parameters should obviously refer to the corrected curve. In the case of shale volume from gamma ray the same answer will be computed whether the uncorrected or corrected curve is used, providing the parameters match the curve. This is always the case when parameters are estimated from the logs themselves. But if the parameters are 'text book' values or are obtained from laboratory measurements it is implicitly assumed the logs have been corrected.

6.2.1 Deterministic and Matrix Inversion Methods

In this book we will mainly consider deterministic log analysis. This tends to use equations that can be written explicitly. For example, shale volume could be found from gamma ray using an equation of the form:

$$V_{shale} = \frac{GR - GR_{clean}}{GR_{shale} - GR_{clean}} \tag{6.1}$$

Where GR_{clean} and GR_{shale} are the parameters and GR is the gamma-ray reading. This is the linear relationship shown in Fig. 6.1. The other relationships are a bit more complicated but can still be written in the general form:

$$V_{shale} = f(GR) \tag{6.2}$$

Archie's equation is another example, whether it is used to find porosity from resistivity or, as more usually, to find water saturation from porosity and resistivity. There are some equations, which cannot be re-arranged into the simple explicit form but these are relatively unusual and they can generally still be solved satisfactorily by numerical analysis (preferably with a computer).

Regardless of the particular equations used, the basic properties are calculated in the order shale volume, porosity and saturation. In the case of porosity and saturation there is really no alternative as saturation is defined in terms of porosity. In all but the simplest cases where shale volume is simply used as a criteria for net reservoir, shale volume needs to be calculated first as it is an input to effective porosity and/or shaly-sand saturation equations. The final interpretation may use the saturation to estimate the hydrocarbon effects on the input logs in which case the whole process will have to run through at least two iterations at each level. Nevertheless, the basic sequence: shale volume–porosity–saturation still applies. In the book we will calculate the basic properties in that order.

Most log analysis packages offer an alternative way of performing log analysis, which is variously known as probabilistic, matrix inversion or by various brand names. In this book we will refer to them as matrix inversion techniques. The details of the algorithms often vary from vendor to vendor but the basic approach is to answer the question:

'What combination of minerals and fluids gives this set of log responses'? There is a clear link back to Chapter 1 where it was stated that a lot of practical petrophysics involves using a rock's physical properties to infer its composition. Although deterministic and matrix inversion techniques are often presented as fundamentally different approaches, the reality is that they have a lot in common. In particular the fundamental equations linking physical and petrophysical properties are generally the same, the difference is the way they are combined. Deterministic analysis normally applies the equations one after the other (like a production line). Matrix inversion finds shale volume, porosity and other properties at the same time.

6.2.2   Computer Log Analysis

Log analysis is ideal for computers – it mostly involves simple but repetitive calculations – and not surprisingly, it was one of the first applications for computers in the industry. With the advent of computer-based acquisition systems in the late 1970s the digital data really became the primary product produced by the logging contractor. That is to say files consisting of blocks of data at regular depth increments. Almost from the outset standard formats were developed for digital data and by and large these have been adopted by all contractors and operators. The two most common formats currently in service are LAS and DLIS. The former is an ASCII format, which can be viewed using a text editor and loaded into a spreadsheet. Its main disadvantage is that file sizes are large relative to the amount of data stored. DLIS is a binary format that makes much more efficient use of the storage medium but it cannot be read using a text editor. Modern log analysis packages can normally read LAS, DLIS and some older formats that are now obsolete and convert them to a form that the package uses.

Government regulations and unitisation agreements were written in terms of optical logs for many years, even after the advent of computers, but increasingly

even these formal documents specify digital files as the primary product. The optical records – originally prints or films – are increasingly actually digital files (e.g. pdfs). Recording the logs in digital form allows far more data to be acquired and it is easier to manage and store. Prior to computers, a typical log would consist of six or seven primary curves (e.g. gamma ray, calliper, density). The cable had always been capable of transmitting far more information but the total number of curves recorded was limited by the optical recording equipment – the camera – or even just to avoid putting so much data on the print that it became cluttered and difficult to read. Recording digital data removes these limitations and so computer log units not only record the primary curves but also the raw data that they are computed from, tool diagnostics and variations on the primary curves (e.g. neutron porosities using sandstone, limestone and dolomite matrices). Amongst other things this allows logs to be re-processed if there is a problem with calibration, for example.

At their simplest log analysis programmes work on a level-by-level basis and treat one block of data at a time. This means for a typical log there are over 650 blocks of data for every 100 m of hole. Even if the interpretation algorithms are simple explicit equations, this represents a prohibitive number of manual calculations and prior to computers log analysis was typically performed by estimating average properties over intervals of several metres. Computers can now very quickly convert the basic primary curves to continuous property curves with individual points at the same depths as the input logs.

In order to do this, all the input data need to exist at the same depth point. Sometimes it is necessary to re-sample some or all of the curves to ensure that this is the case (e.g. when combining logs with different depth increments). Many applications will actually implicitly do this but it is better for the analyst to exert some control.

6.3 PREPARATION

Before starting to calculate properties such as shale volume from the gamma-ray curve some basic preparation needs to be carried out. This includes editing the logs, environmental corrections (if they are to be applied), identifying any intervals of bad data and identifying any special minerals. Depending on the company and the software package that is being used, temperature and pressure depth curves may also need to be calculated.

Editing includes various changes that are made to the curves provided by the logging contractor. This may involve any or all of the following: depth matching, re-sampling to change the depth increment, re-naming, filtering and de-spiking. Often no changes are required and the curves can be used as provided. On other occasions, particularly if the hole has proved difficult to log, a lot of editing may be needed to avoid artefacts in the log analysis. In general, the nature and severity of edits are down to company and personal preferences, but if bad data does exist for any reason, some process needs to be applied to account for it.

6.3.1 Environmental Corrections

Environmental corrections have been discussed in Chapter 5. There it was noted they are generally small changes that are designed to put logs from different wells on a consistent basis. As noted earlier logging parameters need to be consistent with the curves that are input to the interpretation equations. It follows that an interpretation based on uncorrected curves can be as accurate as the one that uses environmentally corrected curves; the analyst just needs to make sure the parameters are appropriate for the input curves.

It is advisable to perform environmental corrections if a matrix inversion method is being employed. This is because the model is based on a specific collection of minerals that are mixed together at a scale that is much smaller than log resolution. The log responses of the individual minerals therefore have to come from external sources such as laboratory measurements or theoretical models. These values implicitly assume that all the environmental factors have been removed. Curves that are particularly in need of correction before being used in a matrix inversion programme are the PEF, gamma ray and spectral gamma-ray curves and – for certain tool types – the neutron porosity.

6.3.2 Re-sampling

As noted in Chapter 4 modern logs are provided as digital data with a regular depth increment. Often this increment is 6 in. or some fraction of that (15.24 cm) but different tools and contractors employ different increments and to get them all consistent the logs may have to be re-sampled. Because the inherent resolution of the measurements is normally a lot lower than 6 in. it does not really matter whether the increment is increased or decreased although it is important to realise that re-sampling to a smaller increment is in no way improving the vertical resolution. The intention is simply to get all the curves to give values at the same depth points.

A particular problem that sometimes occurs with LAS files is that the logging engineer fails to specify enough decimal places for the depth reading. Recall that the fundamental depth increment is often 6in. or some fraction of that. In metric units that is 0.1524 m and so depth needs to be given to at least four decimal places. Unfortunately, for whatever reason, the depth may actually only be given to three or even two decimal places with the exact depth being rounded up or down. For example, consider the series of three depth points with an increment of 0.1524 m in Table 6.1.

The depth inaccuracy is trivially small – about half a millimetre – but as far as the computer is concerned the data no longer has a constant depth increment and for some commercial packages that can cause a major headache! All this can be avoided if the logging engineer is 'on the ball' and generates data files to the required accuracy (or better the software they use ensures this). Note that the file shown in Fig. 6.2 has depth that is given to four decimal points.

TABLE 6.1

True depth	Rounded depth	Depth difference
xx.1524	xx.152	–
xx.3048	xx.305	0.153
xx.4572	xx.457	0.152

```
RB_PPC2.DEG                                          :PPC2 Relative Bearing {F11.4}
RB_PPC1.DEG                                          :PPC1 Relative Bearing {F11.4}
SCD   ..1IN                                          :Speed Corrected Depth {F11.4}
SCDV  .F/HR                                          :Speed Corrected Depth Velocity {F11.4}
SDEV  .DEG                                           :Sonde Deviation {F11.4}
SP    .MV                                            :Spontaneous Potential {F11.4}
SPAR  .MV                                            :SP Armor Return {F11.4}
SPHI  .V/V                                           :Sonic Porosity {F11.4}
STIA  .M                                             :Stuck Tool Indicator, Adjusted {F11.4}
STIT  .M                                             :Stuck Tool Indicator, Total {F11.4}
TDEP  ..1IN                                          :6-Inch Frame Depth {F11.4}
TENS  .LBF                                           :Cable Tension {F11.4}
TIME  .MS                                            :6-Inch Frame Time {F11.4}
#-------------------------------------------------------------------------------
#
#     DEPT      AF10     AF20     AF30     AF60     AF90    AFCO[5]   AFCO10    AFCO20    AFCO30
#
~A
   2885.8464   1.1956   1.1151   1.0608   1.0270   1.0437  836.3788  836.3788  896.7791  942.7033
   2885.9988   1.1956   1.1151   1.0608   1.0270   1.0437  836.3788  836.3788  896.7791  942.7033
   2886.1512   1.1956   1.1151   1.0608   1.0270   1.0437  836.3788  836.3788  896.7791  942.7033
   2886.3036   1.1956   1.1151   1.0608   1.0270   1.0437  836.3788  836.3788  896.7791  942.7033
   2886.4560   1.1956   1.1151   1.0608   1.0270   1.0437  836.3788  836.3788  896.7791  942.7033
   2886.6084   1.1956   1.1151   1.0608   1.0270   1.0437  836.3788  836.3788  896.7791  942.7033
   2886.7608   1.1956   1.1151   1.0608   1.0270   1.0437  836.3788  836.3788  896.7791  942.7033
   2886.9132   1.1956   1.1151   1.0608   1.0270   1.0437  836.3788  836.3788  896.7791  942.7033
   2887.0656   1.1956   1.1151   1.0608   1.0270   1.0437  836.3788  836.3788  896.7791  942.7033
   2887.2180   1.1956   1.1151   1.0608   1.0270   1.0437  836.3788  836.3788  896.7791  942.7033
   2887.3704   1.1956   1.1151   1.0608   1.0270   1.0437  836.3788  836.3788  896.7791  942.7033
   2887.5228   1.1956   1.1151   1.0608   1.0270   1.0437  836.3788  836.3788  896.7791  942.7033
   2887.6752   1.1956   1.1151   1.0608   1.0270   1.0437  836.3788  836.3788  896.7791  942.7033
   2887.8276   1.1956   1.1151   1.0608   1.0270   1.0437  836.3788  836.3788  896.7791  942.7033
   2887.8276   1.1956   1.1151   1.0608   1.0270   1.0437  836.3788  836.3788  896.7791  942.7033
```

FIGURE 6.2 A small part of an LAS file, showing part of the header and a few lines of data. The depth is in the left-hand column.

6.3.3   Depth Shifting

Depth shifting is normally required when tools do not move smoothly up or down the well. This is obviously a potential issue when tools are run separately but, surprisingly perhaps, it can also occur when tools are combined in a single run. Combining tools together is now standard practice for wireline and for logging while drilling (LWD), all the services are run in the same string as a matter of course. If tools are combined the different positions of all the sensors on the tool string have to be accounted for to get all the different measurements on depth.

As an example, consider the wireline tool string shown in Fig. 6.3 (the same arguments apply to LWD). If the log is being recorded upwards, the upper sensor A encounters a particular depth first and its response is stored in memory until the tool has moved the distance to the lowermost sensor. At that point the
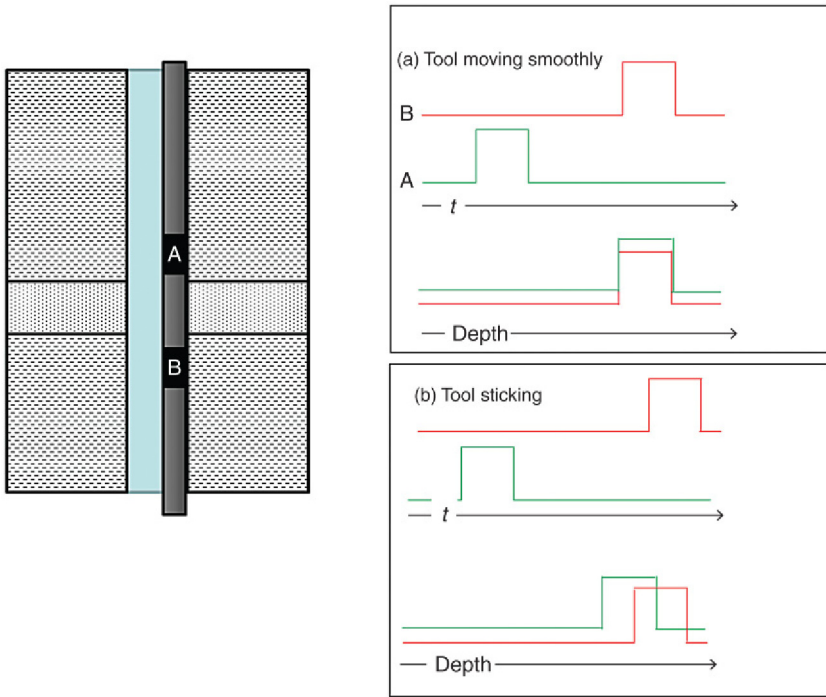
FIGURE 6.3 A cartoon showing how measurements made with the same tool string can appear off depth. The tool string has two sensors – A and B – at different points on the tool body. Sensor A logs the sand first and the log is stored in memory until sensor B has moved the distance between them. (a) If the tools move smoothly – upper traces – the different measurements are put on depth. (b) If the tool string sticks – lower traces – the measurement from B will actually be from a deeper point than from A.

outputs from both sensors are recorded (in fact the tool string may travel even further before the measurements are recorded what matters is that all the sensors have passed the depth of interest). If the sensors are 3 m apart the measurements will be recorded when the tool has moved at least 3 m.

The problem with depth mismatches originates in the fact that depth is actually measured by noting how much cable or drill pipe has passed the drill floor. The fundamental assumption is that if 3 m of cable moves past the drill floor then the tool string moves by 3 m down hole. In reality cable and pipe is sufficiently elastic that it can move at surface with no tool movement at depth. The recording system does not know this of course and assumes the tool has moved by the same amount.

If the tool sticks for a short time, the lower sensor will actually be more than 3 m below the upper sensor when the individual measurements are combined and they will be off depth.

The same problem will occur if the measurements are recorded on different runs and one or both tools stick. Most modern log analysis packages have a module that allows different curves to be depth matched interactively. The log analyst just needs to decide which curve to use as the reference curve to match the other curves to. Depth matching is invariably required when wireline logs are being combined with LWD or core measurements since the latter will have used the driller's depth measuring system.

There is also a possibility of systematic depth offsets between different runs or even within a single run. The former occurs because the two runs have not been correctly put on depth during acquisition and the latter happens because the computer is given incorrect tool lengths (often some additional modules are put in the tool string to centralise it or make it more flexible and if the engineer forgets to account for these the memorisation will be wrong). Digital recording means that the logging contractor should be able to correct systematic errors before the final product is delivered.

6.3.4   Filtering and De-spiking

Filtering is basically a smoothing process that is designed to avoid artefacts caused by the different vertical resolutions of the tools. In fact the data provided by logging contractors is normally filtered already. This is partly to improve its appearance, partly to remove noise from 'nuclear' measurements such as density and neutron porosity and partly to remove the dependence on the precise depth that logging commences. Unfiltered data can always be provided either when the log is recorded or as a play-back, but this normally has to be explicitly requested.

De-spiking is a more localised edit, probably to only one or two curves, in which thin features with relatively extreme values are arbitrarily excised. Of all the preparation stages, this and filtering are the most controversial because we are deliberately altering the readings.

Filtering implicitly reduces the vertical resolution of the different curves and so is probably best avoided unless there is good reason to believe the highest resolution curves are imposing some unwanted fine structure on the interpretation. If a filter is going to be applied we still have to decide between several different filtering algorithms. The simplest of these take an average of all the readings in an interval centred on the depth point. The window should be equivalent to an odd number of depth increments. Obviously, the larger the window the lower the resolution of the output curve. There is also a choice over what type of average to take. An obvious choice is the mean but this will always be influenced by a small number of unusually high values within the window. In other words a high amplitude spike will shift the whole curve to higher values over an interval equal to twice the window width. A better choice may be to use the median, which is unlikely to be influenced by a few extreme values. Alternatively more complicated weighted averages could be used.

6.4   PARAMETER PICKING AND DISPLAYING LOGS

As noted on more than one occasion log analysis involves the following:

1. Selecting the form of the equation that relates one or more logs to a petrophysical property.
2. Selecting the constants that make the equation specific to the particular well and zone that is being analysed.

The particular equations that are used will be discussed later for shale volume and in later chapters for the other properties. Here we will look at some of the general tools for picking the parameters. As a general rule, deterministic analysis models the formation as a mixture of end members whose log responses can be readily determined from the log readings themselves. 'Shale' for example, is invariably one of the components although as we have seen, in reality, shale is quite a complicated mixture of minerals. The matrix inversion methods generally model the formation as mixtures of specific minerals whose properties ultimately have to be obtained from laboratory measurements or even theoretical models (a few examples were given in Tables 1.1 and 2.1).

The simplest parameters to pick are those of a component that exists somewhere in the section as a pure, resolvable bed. The component most likely to fulfil this requirement is the shale, but other examples include special minerals and occasionally limestone or dolomite in a carbonate reservoir. Assuming a suitable bed exists there are two ways to find the parameters. A suitable bed means one that is:

1. Thick enough to be resolved by the log(s) of interest.
2. In good hole.

The simplest way of finding the parameters, which is perfectly adequate for 'quick-look' interpretations, is to simply read the log in the bed of interest. For more sensitive work it is better to generate a histogram of the log readings and use some statistical definition to find the parameter. The latter is actually not likely to produce an interpretation that is any more reliable than just 'eyeballing' the logs but it does allow the procedure for picking the parameter to be specified, so that in principle anyone can duplicate the result.

6.4.1   Histograms

Histograms are a useful tool for illustrating variation in a log reading and properties that are derived from it. Even if they are not used to pick parameters, histograms are used to show how properties are distributed, which is an important consideration for building static geological models. So any log analysis package worth the name should include a module for generating histograms. Moreover, it is not too much to expect an option to annotate the plot with key statistics and to be able to filter out data points when a different curve – such as a bad hole indicator – passes some threshold.

In general there are an infinite number of histograms that can be produced from a particular data set. These are distinguished by the number and size of the classes – or 'bins' – into which the data is separated. The histogram module in a log analysis package should allow the analyst to specify these as well as the left- and right-limiting values (which are obviously chosen to include the majority, if not all, of the data points). For most petrophysical properties, the classes are normally made to have equal sizes. This makes the shape of the histogram and hence the distribution easier to appreciate. The main exception to having equally sized bins occurs for data types, which are naturally plotted on logarithmic scales, notably permeability and resistivity, where the class size expands for higher values.

Although any class size can be used there are clearly upper and lower limits to what is sensible. The extreme case of a class that takes in all the data points is clearly of no value. A histogram based on a large number of classes that are so small, that the majority have one or no data points in them are almost as bad. As a rough rule of thumb a good number of classes is equal to the square root of the sample size. This assumes the high and low limits of the histogram are more or less the same as those of the data set. In practice, to compare histograms from different zones and/or wells we may wish to use a larger range and the number of classes needs to be expanded accordingly (e.g. the lower limit for a histogram of gamma-ray readings may be better set to zero than the lowest reading of the data set). In any case the 'square-root' is only a rough guide and one may need to experiment to get the most informative plot. The default number of classes for most log analysis packages is about 50 and more often than not this gives a satisfactory plot.

The shape and location of a histogram can be defined by various statistics. These really fall into three categories depending on whether they define:

1. The location.
2. The variation.
3. The shape of the distribution.

The location is basically another word for the average of the data points. There are a number of different averages most of which are independent of the choice of classes. The exception is the mode – the centre value of the most frequently occurring class. Although often dismissed as 'useless' this may actually represent the best choice for a parameter. It does however depend on how the classes are defined (both width and limiting values).

Variation quantifies the spread of the readings: how high and low they can typically be. There are various measures but the most frequently used are: standard deviation, minimum and maximum, and particular percentiles. The standard deviation can only really be used to define the distribution if it approximates a 'bell curve'. For parameter picking particular percentiles are generally better.

The shape specifies any skew to high or low values and whether, for example the points cluster around a single or several peaks. Plots are often annotated with two specific statistics known as 'skewness' and 'kurtosis.' 'Skewness' as the name suggests is a measure of the asymmetry of the distribution: roughly

speaking how far the peak is from the centre point of the data. The higher the magnitude, the more asymmetrical the distribution. A specific example of a skewness measure is the Pearson skewness given by:

$$\frac{\text{Mean} - \text{Mode}}{\text{Standard deviation}}$$

So, if the distribution has a long 'tail' of high values, the Pearson skewness will be positive and the distribution is said to be 'skewed to the right' (this obviously eludes to the convention for plots that values increase from left to right).

Kurtosis basically quantifies the 'sharpness' of the peak with high values characteristic of a sharp peak.

6.4.2 Cross-plots

We have already met cross-plots in earlier chapters, at their simplest they just involve plotting the values of two different curves at each depth point against each other. There are really two applications for a cross-plot:

1. To help determine a quantitative relationship between two properties.
2. To help show that two sets of data come from the same population (or conversely that they come from different populations).

This is similar to the role of regression and correlation discussed in Section 2.6 and indeed cross-plots and correlation are closely related. For the specific case of parameter picking cross-plots are used to estimate values for components which do not exist in isolation and so cannot simply be read from logs or found from a histogram. An example of this was given in Section 1.5.1 where a cross-plot of resistivity against porosity was used to find the resistivity of the formation water. In order to do this the best line through the data points had to be extrapolated to 100% porosity – pure water. Another common example, applied extensively during the main phase of North Sea developments, involved cross-plotting density against porosity in order to find the densities of the end point components: fluid and matrix.

For the remainder of this section we will look at some of the specific tools that can be applied to a basic cross-plot in order to make it more informative. As with histograms it is desirable to be able to exclude points from a plot when the value of a third curve exceeds certain limits. Most obviously this can be used to exclude bad data caused by bad hole for example, but it may also be used to concentrate on one particular lithology or fluid type. Most packages allow points to be colour coded or to use different symbols according to the value of a third curve. The third 'curve' may be a different log or it may actually code for different zones or even different wells. In fact, most packages allow colour to be controlled by a third curve and the symbol to be determined by a fourth curve. So in principle it is possible to construct a cross-plot summarising relationships between four different curves. In practice most people find it quite difficult to

make sense of so much information, except when the relations are so clear that they could be seen in the original curves. It is worth stating at this stage that all a cross-plot is doing is removing depth as a variable from the logs.

The simplest – normally default – cross-plot consists of one point for every depth point but it is also possible to plot the data in classes defined by upper and lower values for the *x* and *y* variables. Each class is then coded according to the number of data points that lie within it. In effect the cross-plot is then a two-dimensional histogram.

Finally, several packages allow a 3-D cross-plot to be produced. These simply use a suitable projection to give the impression of a 3-D plot. These are really only useful for illustrating particular features but on the basis that a 'picture can be worth a thousand words' they are certainly more than just a gimmick.

6.5   SHALE VOLUME

As we have seen 'shale volume' is often used as a proxy for and often, incorrectly, as a synonym for the volume fraction of clay in the formation (implicitly within the volume of investigation of the logs). The amount of clay needs to be quantified because clay minerals can have a strong influence on log readings. The general approach to finding shale volume is to find one or more logs which show a strong contrast between the shales and the reservoir lithology. In principle any log or combination of logs can be used providing a reasonable correlation exists with shale volume. All software packages offer a range of standard techniques that may or may not fulfil this criteria (normally, at least one or two of them will).

As a rule the two most universally applicable and commonly used methods are the gamma ray – including the spectral gamma ray – and the density–neutron combination. These shale indicators actually include a number of different specific algorithms (the gamma ray as a shale indicator was used as an example in the introduction to the chapter). Other techniques that have more limited application but still work well in the right circumstances are the density–sonic combination, sonic–neutron and the SP. Geochemical and NMR logs which have more recently moved into the mainstream actually offer some of the most accurate ways of finding shale volume and even a true clay volume. In this section, we will look at the more common methods and use them to illustrate the important topic of parameter picking.

6.5.1   Shale Volume from Gamma Ray

Although there is no fundamental reason why the gamma ray should make a good shale indicator, the fact remains that more often than not it does and in practice it is probably the most popular method. The gamma ray is an indication of how much potassium (strictly ^{40}K), uranium and thorium are in the formation. Providing one or more of these elements is closely associated with the main components of the shale, the gamma ray will show a good contrast between shales and other lithologies.

Potassium is part of the composition of some clay minerals and in practice it is often present as an impurity in the non-potassium containing clays. Thorium is also commonly associated with shales as it tends to occur in hard, insoluble chemically inert oxides that ultimately get ground down to shale or silt sized particles (e.g., zircon). Uranium is less straightforward as although, like thorium, it is often incorporated into small fragments of chemically inert oxides it can also form water-soluble compounds that allow it to be leached out of the rock. It is well known for its association with organic matter, which tends to create conditions that encourage it to precipitate from solution.

It should be noted however, that there are plenty of cases where one or more of these elements is more closely associated with the reservoir quality rock. For example, sands that contain potassium feldspar will have high gamma-activities and 'hot' limestones with uranium impurities are also widely known. Figure 6.4,
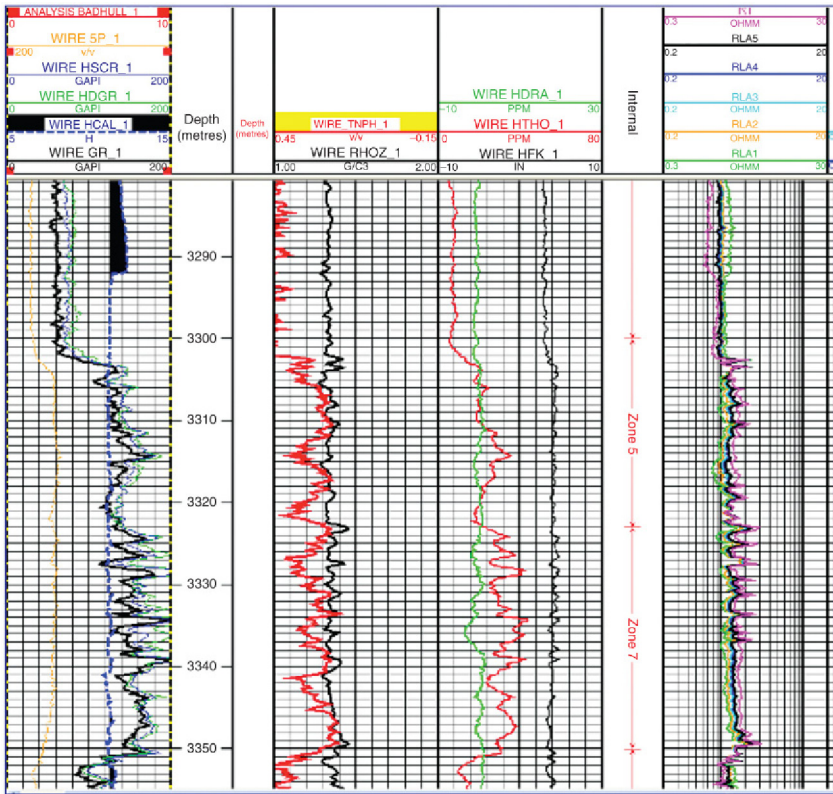
FIGURE 6.4 An unusual set of logs in which the shales (above 3302 m) are characterised by relatively low gamma activities. The shales can be reliably identified by density–neutron separation (track 3). The spectral gamma ray (track 4) shows the high gamma in the activity in the sand is due to increases in thorium (grey) and a lesser extent potassium (black).

shows some logs where the gamma ray is at its lowest in the shales although these are clearly identifiable as argillaceous with the density–neutron combination.

Spectral gamma-ray tools determine the separate contributions from potassium, thorium and uranium and it may be that whilst the combined gamma-ray activity is not a very sensitive shale indicator, one or two of the individual elements are. In particular removing the uranium contribution often results in a more sensitive shale indicator. On the other hand as noted in Chapter 2 there are limitations to where spectral gamma-ray tools can be used and the curves representing the individual contributions are not as accurate as the total gamma activity. Older tools may simply lack the accuracy to give a robust shale indicator.

Assuming the gamma ray or one of its components, does make a good shale indicator the next task is to define a relationship between gamma-ray activity and shale volume. As with all deterministic log analysis, this involves the following two steps:

1. Define the form of the relationship (the shape of the curve).
2. Define the parameters (the gamma activity in clean formation and shale).

The simplest relationship is a linear one in which the gamma-ray activity increases in direct proportion to the amount of shale, but over the years a number of non-linear relations have also been proposed. Graphs for some of these were shown in Fig. 6.1, they always produce a lower shale volume for a particular gamma reading than the linear model (except at 0% and 100% shale volume). They are designed to account for the changing volume of investigation of the gamma-ray tool with formation density and some were quite rigorously derived. Depth of investigation falls with increasing density and in a typical sand-shale system density is highest in the shales. This in turn means a smaller volume of formation contributes to the gamma-ray activity in a shale than a sand. The greatest discrepancy occurs at shale volumes of about 50%.

The linear equation has the virtue of simplicity and is actually the correct model to apply when the shale density is close to the sandstone density. The non-linear models are expected to work better when the shales are relatively dense, although in practice the differences will only be significant in intermediate compositions. Unless there are good reasons for selecting a non-linear method it is generally best to stick to the linear equation.

The linear relationship is defined by two parameters that are almost always the gamma-ray reading in the shale and the clean formation. For quick look work they can normally be satisfactorily estimated by simply reading from the log. For more precise work, they are normally defined statistically. Generally the best approach is to produce a histogram of the gamma-ray readings in the interval of interest and depending on the distribution of sand and shale the parameters can variously be defined in terms of the mode(s), mean or percentiles. It should be noted that the purpose of using statistics is simply to introduce precision into the choice of parameters and there is no fundamental reason why,

for example the 90th percentile on a histogram should represent the gamma reading in pure shale.

As noted in Section 6.2 here are a few issues to be wary of when defining parameters. Specific issues to consider are:

1. The parameters apply to a particular set of curves. For example, if the gamma ray has been environmentally corrected then the parameters must be derived from the corrected curves.
2. The clean reservoir lithology need not give the lowest gamma-ray activity. If the formation contains evaporite beds or tight limestone stringers these will probably be characterised by the lowest gamma-ray activities. The clean reservoir lithology is likely to show higher activity (albeit still low compared to the shale). But if the lowest readings are used to define the 'clean reservoir' value, shale volume will be over-estimated everywhere except the pure shales.
3. Does the reservoir interval contain any resolvable, completely clean formation? For that matter does it include a representative shale that can be resolved by the gamma-ray log? If this is not the case the gamma ray will never reach the true clean (or shale) readings. This is a particular problem in thinly bedded reservoirs.
4. If a suitable shale can be found is it genetically related to the reservoir interval?
5. Is it shale volume or clay volume that is actually required? If it is clay volume, the lowest gamma reading observed need not necessarily correspond to 0% clay.

As an example consider the short interval of sands and shales in Fig. 6.5. These give the 'conventional' behaviour in which the shale is associated with relatively high gamma activity. A simple inspection of the logs suggests the sandstone is characterised by a gamma activity of approximately 50 api and the shale by just over 200 api.

A histogram of the gamma-ray readings is shown in Fig. 6.6, the programme that generated it produces a lot of statistical information automatically. Some pertinent values are as follows:

Minimum = 15 api
Maximum = 229 api
5th percentile = 42 api
95th percentile = 205 api
Mean = 91.8 api standard deviation = 42.2 api
Mode = 77 api

The extreme values are not recommended as parameters to define the clean and shale point. This is partly for the reasons discussed earlier although in this particular case there is reason to believe the lowest value really does correspond to the cleanest sand. Ultimately however the gamma ray is a statistical measurement and the extreme values are unlikely to be reproduced at exactly the same depths. Many procedures specify the 5th and 95th percentile readings as the
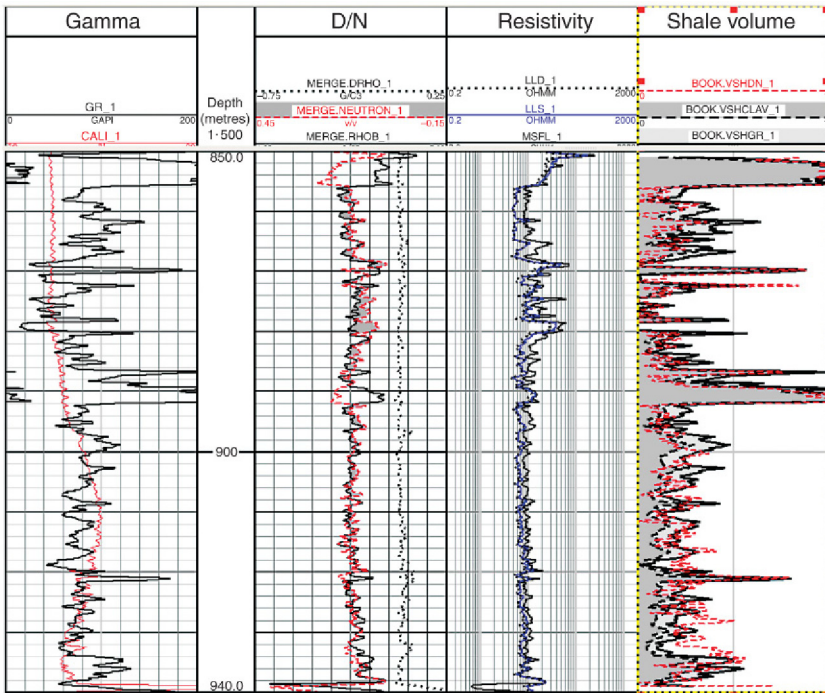
FIGURE 6.5 A 90-m section of logs over a sandstone reservoir sand. The logs have not been environmentally corrected. The computed shale curves are shown in the right-hand track. Three different shale volume estimates are shown: two using the gamma ray (black curves) and one using the density–neutron separation (dashed [dashed red in the web version]).

clean and shale points respectively and these have been used here to calculate the shale volume curves in the right-hand track of Fig. 6.5, but as already noted there is no fundamental reason for using these values and other procedures might specify the 10th and 90th percentiles, for example.

It was noted earlier that the shale parameter should come from a shale bed that is genetically related to the sand. In other words was deposited by the same general system. It is almost impossible to be completely sure that this condition is fulfilled using logs alone. But there are ways to increase confidence:

1. A gradational boundary between sand and shale.
2. If a resolvable shale can be found within a sand package.
3. If the shale readings above and below a sand are the same.

Conversely a sudden change from sand to a thick shale may indicate an unconformity with the two lithologies separated by large time spans.

In this case the high gamma-ray values that are assumed to characterise the shale, come mainly from two beds: one at the top of the interval between 852–856 m and a 2 m bed in the middle of the sand (890–892 m). There are other
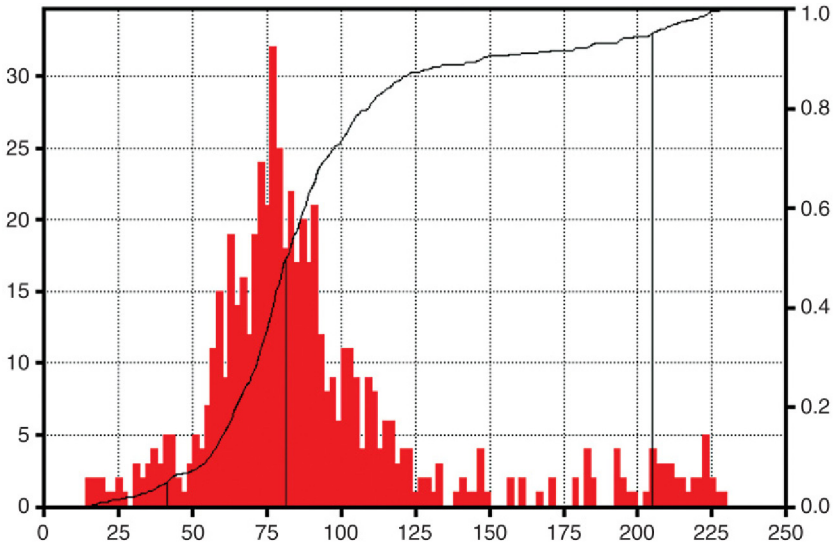
FIGURE 6.6   A histogram of the gamma-ray readings for the log shown in Fig. 6.5. Some pertinent statistics are given in the text. Each bin on the histogram represents 2 api units.

candidates, particularly above 890 m but these are too thin to be properly resolved and the gamma-ray reading will under-estimate the true value for a shale. The 2-m thick bed in particular provides confidence that it is representative of the shale in the interval precisely because it does lie in the middle of the sand.

6.5.2   Density–Neutron Cross-plot

In good hole conditions the density–neutron combination is arguably the most generally applicable method for finding shale volume. It is not however infallible and is not suitable for sandstone reservoirs which include varying amounts of heavy minerals. Furthermore, if the reservoir consists of dolomite the method is unlikely to work at all. The method basically relies on the fact that clays and therefore shales have a lot of hydrogen associated with them and therefore give a relatively high neutron porosity for a particular density reading.

In fact the density–neutron is a particularly powerful pairing that can be used to quickly identify lithology and to distinguish gas from oil. It is therefore worth spending time to discuss the combination in general. In Chapter 5 it was noted that density and neutron tools have been run in combination for decades (in fact several LWD and wireline tools now combine the two measurements in a single tool). This is partly because both measurements employ weak radioactive sources, which attract a lot of onerous rules and regulations regarding their handling and storage. Consequently, it is best to get both measurements out of the way at the same time. More importantly, however the two measurements complement

each other which is exploited by plotting the curves in the same track using standard scales. Several examples have already been given (e.g. Figs 6.4 and 6.5). Before discussing the cross-plot we will briefly look at these scales.

If density and neutron porosity curves are available from the same section of hole they will almost certainly be plotted on the same grid. This is just another aspect of the combination being worth more than the sum of its parts. By default the density is plotted increasing from $1.95 \, \text{g/cm}^3$ on the left-hand side to $2.95 \, \text{g/cm}^3$ on the right-hand side. The neutron porosity on the other hand is plotted increasing from -0.15 on the right to 0.45 at the left. These default scales are not very convenient for reading values. Why not plot the density from 2 to $3 \, \text{g/cm}^3$ and the neutron porosity from say 0 to 0.5? (In high porosity formations the scales may be shifted to account for the lower average densities and higher neutron porosities but they are still not the most obvious choices.) The answer is that the scales are chosen to give a clear visual indication of lithology. In particular if the neutron log has been recorded using a lime matrix (see Section 5.3.2), the density and neutron curves will overlay in water-bearing limestone. In a typical water-bearing, quartz-rich sandstone the neutron will plot just to the right of the density curve and in water-bearing dolomite, the neutron will plot well to the left of the density curve. (If a sand matrix is used the curves will overlay in a water-bearing sandstone.) Strictly speaking the logs should be environmentally corrected for this to apply and obviously the tools should be properly calibrated etc.

So just by glancing at the log it is possible to make a quick inference of the lithology. The cross-plot reveals the same relationships but removes the depth as a variable and in fact density–neutron is such a useful combination that default grids have been produced for just about every combination of density and neutron tool types that a particular contractor has produced. Nowadays that includes charts for wireline and LWD tools, charts for neutron logs that have and have not been environmentally corrected and charts for each type of neutron tool, if a contractor has more than one tool in service. Most of these will also be available in digital form in a log analysis package.

An example of a typical density–neutron cross-plot grid is shown in Fig. 6.7. This has been drafted specifically for this book and does not apply to any real tools. However, all density–neutron cross-plots are very similar and any real example from a chart book will look much the same. It will be used to illustrate how the density–neutron combination can be used to find shale volume and other useful information.

By convention neutron porosity is plotted on the *x*-axis and density plots on the *y*-axis. Neutron increases from left to right and the lowest value is normally a small negative value. This is because the neutron porosity of quartz is less than zero – the exact figure depends on the tool type. Density increases from top to bottom and whilst this seems perverse it is more or less equivalent to porosity increasing upwards. So, all things being equal, points plot from the lower left to upper right as porosity increases. The blank cross-plot invariably has three
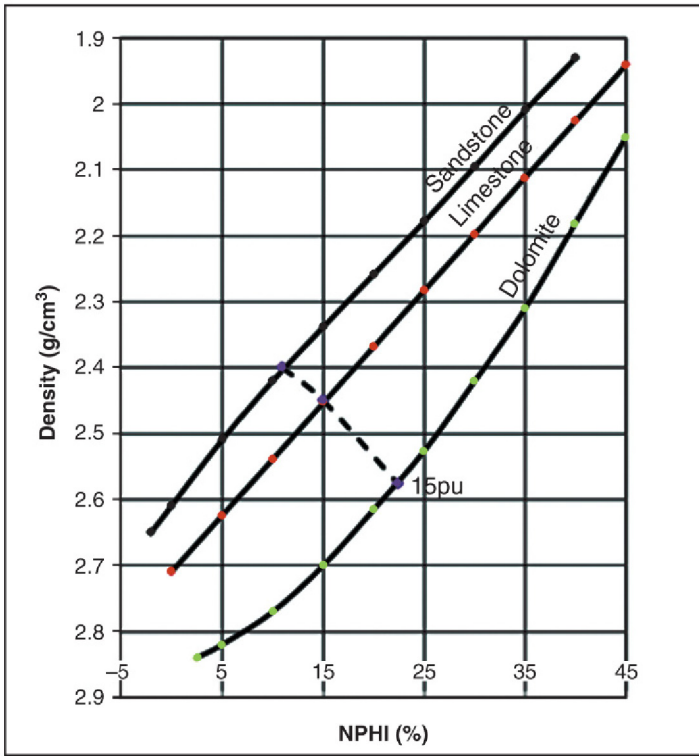
FIGURE 6.7 A blank density–neutron chart showing the lithology lines marked off in 5% porosity increments. An iso-porosity line has been constructed for 15% porosity.

'lithology lines' marked on it which are approximately parallel and lie in the direction of changing porosity (lower left to upper right). These are from top to bottom the 'sandstone', 'limestone' and 'dolomite' lines. For clean water–bearing limestone and dolomite the points should fall close to these lines (reasons why this might not occur will be discussed later). The sandstone line is a bit less clear-cut because sandstones normally have a more complicated mineralogy. The 'sandstone' line refers to a sand with a matrix comprising a lot of quartz (greater than 70% with a matrix density close to $2.65\,g/cm^3$). If some heavier minerals are present, the points will fall between the sandstone and dolomite lines or even on the limestone line. If there is a relatively large amount of feldspar, which has a relatively low density, the points may actually plot just above the sandstone line.

The lithology lines are normally marked off with ticks corresponding to total porosity. In most chart books the ticks are at intervals of 1 pu and they are normally annotated every 5 pu. For water-bearing formations they allow total porosity to be estimated and for clean formations the porosity can simply be read off

the chart even if the point falls between lithology lines. In log analysis software the ticks are often replaced with lines of equal total porosity. An example is given in Fig. 6.7, with the 15 pu line being added (normally a line is added every 5 pu so that a grid is produced). Some cross-plots also include single points corresponding to minerals that can form beds that are thick enough to resolve (salt and anhydrite are the commonest examples).

Although the lithology lines are approximately parallel, a closer look shows that they actually start to converge towards the high porosity end of the plot. This is to be expected because if the axes were extended to a density of – or close to – $1.0\,g/cm^3$ and a neutron porosity of 100% all three lines should meet. This is the fluid point where the tools are effectively suspended in a swimming pool. Chart books typically produce charts for both fresh and very saline water with densities of 1.0 and $1.1\,g/cm^3$ respectively. Some software packages allow the grid to be adjusted to correspond to any fluid type the analyst desires (e.g. oil-based mud filtrate).

If something other than water occupies the pore space the lines will converge to a different point. The most drastic change occurs in good quality gas reservoirs where the pore space is largely filled with a low-density fluid (90% gas and 10% water, for example). Invasion may displace some gas in the near well bore region but in high-porosity reservoirs invasion tends to be very shallow and so makes little difference. The net effect is that the fluid point where the lithology lines converge is shifted to lower densities and in the case of gases and high GOR oils it will also be shifted to lower neutron porosities. In other words the lithology lines will rotate upwards about the matrix points. The opposite situation where the pore fluid is significantly denser than $1.0\,g/cm^3$ occasionally occurs and then the lithology lines rotate downwards. This occurs naturally in some carbonate fields when the formation waters contain a high concentration of calcium chloride. It can also occur with modern heavy muds that are based on formate solutions rather than chlorides.

We will now return to the specific case of finding shale volume from the density–neutron combination. Figure 6.8 shows a short interval of logs which includes a claystone overlying an interval of water-bearing sandstones and thin shales. Figure 6.9 shows the density and neutron porosity points plotted on the same chart as in Fig. 6.7. Some additional porosity lines have also been added although they are not necessary for the shale volume calculation.

The log data shown in Fig. 6.8 and cross-plotted in Fig. 6.9 falls into two distinct groups corresponding to a 25-m thick claystone and a series of thin sands and shales below it. Consider the cross-plot divided into two halves by the limestone line (the limestone line corresponds to points on the log where the density and neutron porosity curves overlay). The points from the claystone lie well within the lower-right half of the plot, this is typical of any 'shale' (which here includes claystones, mudstones as well as shales). The exact location of the shale's points can vary a lot depending on its composition and how much compaction it has undergone, but they will always lie in the lower right-hand
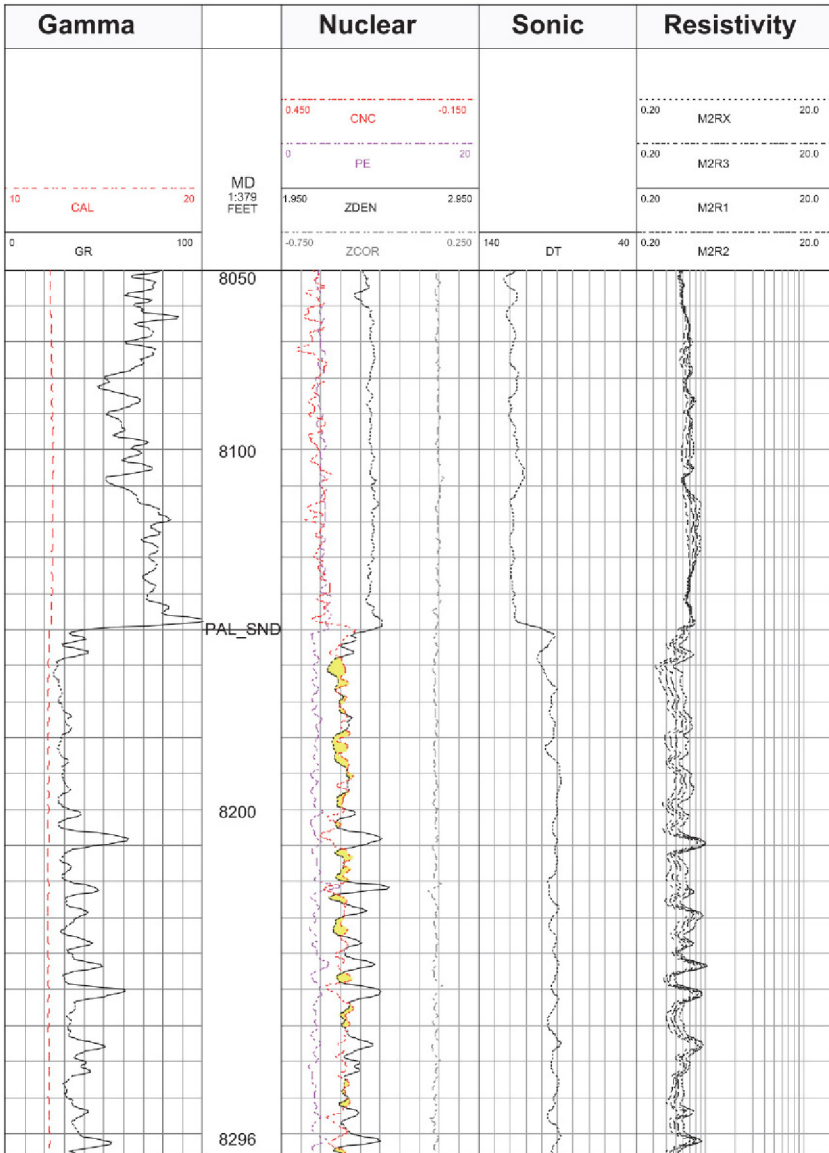
FIGURE 6.8 Conventional open hole logs over the interval of sand and claystone discussed in the text.

side of the plot. Often they lie on or close to the dolomite line, which explains why this is not a good technique to use with dolomite reservoirs. Even if it lies off the dolomite the contrast between the shale and dolomite responses will be small and so the method will never be very accurate.
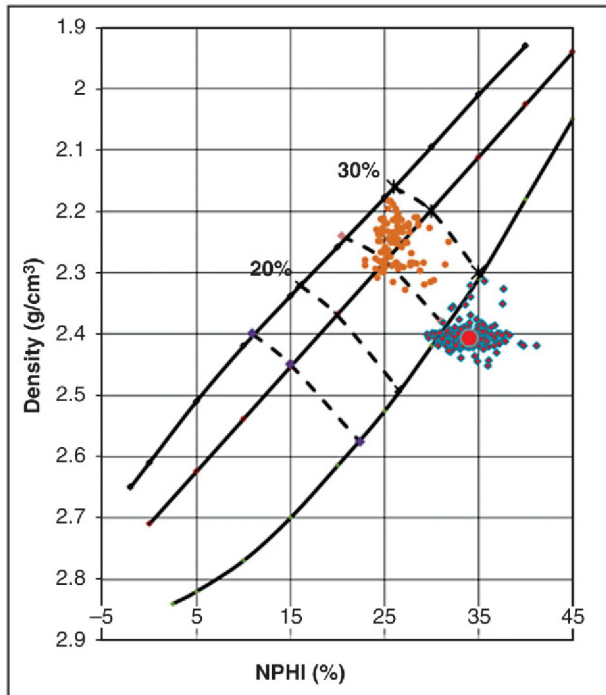
FIGURE 6.9 The cross-plot shown in Fig. 6.7 with log data added. The dark grey dots (red dots in the web version) are taken from the 25-m thick claystone bed. The other points come from the underlying interval of sands and thin shales.

On a log plot such as Fig. 6.8 using conventional scales for the density and neutron curves, shales always appear with the neutron plotting several divisions to the left of the density. The main reason that the shales plot where they do on a cross-plot is the large amount of clay they contain. Recall that clays contain a significant amount of hydrogen associated with their crystal structure. For a neutron tool this is indistinguishable from hydrogen associated with water (i.e. total porosity). So the net effect is to pull the log points to the right relative to a rock with similar water content but no clay. Furthermore, clays often have high grain densities as a result of including heavy elements such as iron in their layers. This has the effect of moving the shale point downwards. Finally traces of neutron-absorbing elements, which tend to associate with shales, reduce the neutron count rate leading to an apparent increase in neutron porosity (at least for thermal neutron tools). In summary, for a number of reasons, shales will always shift log readings into the lower right half of the plot. This is exactly what is needed since it is producing a good contrast in responses between sandstone or limestone and the shales.

The calculation of a shale volume from the cross-plot is really just a geometrical construction. The parameters are the density and neutron readings of the shale point and the end points of the clean lithology line (six in total). The

shale volume of any point is found by constructing a line from the shale point through the log point to the lithology line. The shale volume is simply the ratio of the distance from the clean line to the log point to the distance from the clean line to the shale point. This is shown in Fig. 6.10, which is basically Fig. 6.9 with most of the log data removed. If the various points are given the symbols cl, log, SH for the clean point, log point and shale point, respectively and if the distance between any two points is given by x–y then the shale volume is given by:

$$V_{shale} = \frac{cl - log}{cl - SH} \tag{6.3}$$

Which is similar in form to Eq. 6.1, the equation for shale volume from gamma ray. The difference is that the clean point actually changes with the log reading. To be precise the clean point now moves up and down the clean lithology line as the log point changes.
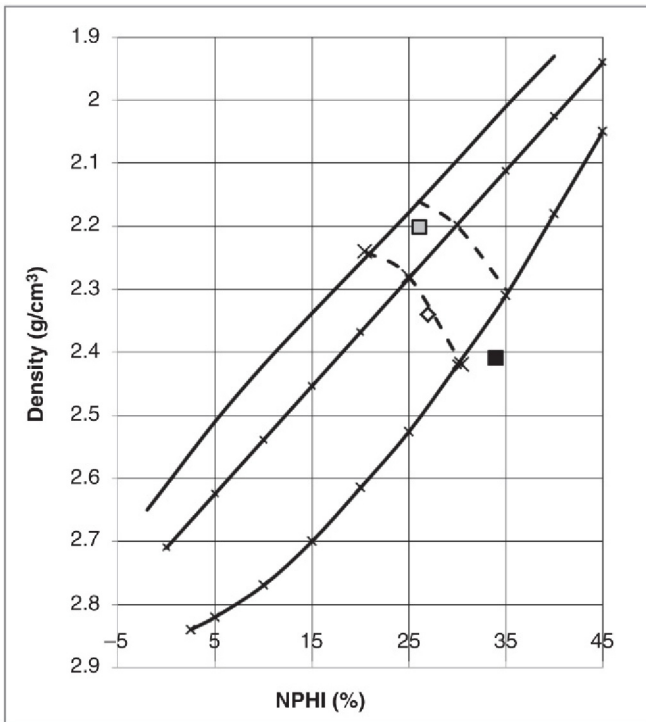
FIGURE 6.10 The cross-plot shown in Fig. 6.9 with most of the data points removed for clarity. The shale point is the black square. Two log points have been left. One with a significant shale volume is shown by the open diamond and a point. This has a shale volume of 50%. The point represented by the grey square has a shale volume of 9%.

As always the parameters that define the clean line and shale point must correspond to the particular curves that are being used. Most obviously if the logs have been environmentally corrected the parameters must be picked from the corrected curves. More fundamentally however the clean line must account for whatever fluid is in the pore space. The lithology lines added to blank charts generally refer to pore volume that is saturated with water and as explained earlier, if the pore space contains a significant amount of another fluid with a density different to $1.0\,g/cm^3$ the whole line will rotate upwards – or occasionally downwards – around the matrix point. If this is not accounted for the shale volume will be under-estimated (or over-estimated for the dense brines).

The simplest way to account for hydrocarbons is to simply choose a different fluid point but for accurate work the fluid point will continuously vary in response to changes in oil or gas saturation and/or invasion. To account for this a more elaborate method is needed and this will be discussed later when we consider hydrocarbon effects. It is also worth noting that hydrocarbons, especially gas, move points characteristic of a water-bearing formation in the opposite direction to shale. When a point is influenced by both gas and shale it is impossible to unravel the two opposing changes using the density–neutron cross-plot alone. To make any progress we have to find either shale or hydrocarbon saturation by some independent means. Shale volume, for example could be found from the gamma ray and its influence on the density–neutron response could then be estimated and accounted for.

6.5.3 Other Cross-plots

Chart books normally include blank grids for constructing density–sonic and sonic–neutron cross-plots. These can be used in a similar fashion to the density–neutron cross-plot but they do not have such general applicability. The fundamental problem is that the sonic slowness is sensitive to rock texture and fabric as well as mineralogy and porosity. One result of this is that a large number of different relations between sonic slowness and porosity have been proposed over the years (discussed in Chapter 7). The shape and position of the lithology lines depend on which, if any, of these apply to the formation under study. In other words the lithology lines can never be as universal as for the density–neutron cross-plot, where both tools are insensitive to fabric. In practice, this means that in some cases there is a good contrast between clean and shale points, but in others they overlap.

Where a good contrast does occur the same basic methodology developed for the density–neutron log can be applied. Again six parameters are needed to define the clean line and the shale point and the shale point can be found by noting the readings in a suitable shale bed. The density–sonic cross-plot does make a good shale indicator in young, unconsolidated formations deposited in deepwater environments (e.g. the deepwater reservoirs of the Gulf of Mexico and West Africa).

6.5.4 Nuclear Magnetic Resonance (NMR)

Estimating shale or clay volume from nuclear magnetic resonance (NMR) log is a recent application for NMR data, which is itself a relatively new addition to the log analyst's armoury. Nevertheless, NMR offers the promise of a more objective method than any of the 'classical' techniques discussed earlier. The reason is the NMR tool's ability to distinguish water that is bound to clay from water in other environments. If one knows how much clay-bound water is associated with pure shale then it is possible to use the clay-bound water volume to quantify the amount of shale everywhere. This assumes that the shale does not change its composition of course, but actually we have tacitly made this assumption in all the previous methods as well.

The use of NMR to estimate shale volume is best illustrated by an example. Figure 6.11 shows conventional logs and some shale volume curves calculated from them in an argillaceous gas-bearing sand. The gamma ray and density–neutron combination suggest the clay content reduces upwards to give the cleanest sand at the top of the reservoir. The question is how clean? The density–neutron combination gives a separation, which is typical of clean sand for 1 or 2 m at the top. This coincides with the lowest gamma-ray reading. The simplest interpretation is that the top of the reservoir is quite clean and the gamma reading there can be used as the clean point for the transform.

The shale volume curves based on the assumption that the top of the reservoir is clean are shown in the track named 'shale 1' (track 6). They include shale volumes based on gamma ray and one based on the density–neutron combination. Shale volumes from the gamma ray include a simple linear model and the non-linear Clavier. The parameters used for each model are summarised in Table 6.2.

The two basic methods of calculating shale volume: gamma ray and density–neutron both involve making an assumption. In the case of the gamma ray it is assumed that the reservoir is clean at the top and the gamma-ray reading there is the value characteristic of clean sand (80 api in this case). In the case of the density–neutron it is assumed that the fluid in the near well bore region is water (with density $1.0\,g/cm^3$ and a neutron porosity of 1 v/v). The negative separation of 0.06–0.09 at the top of the sand is then what is expected of a clean sand with a grain density close to $2.65\,g/cm^3$. (The actual grain density of $2.67\,g/cm^3$ comes from routine core analysis in this example.)

Based on these assumptions, the shale volume is 0 at the top of the reservoir but that is only because we decided that the reservoir is clean there.

The next track to the right uses the bound water fraction curve from the NMR to guide parameter selection (this is track 7 named 'shale final'). The NMR curve has been plotted from 0 on the left to 0.28 on the right. The value of 0.28 is the value in the shale immediately below the reservoir. We are assuming this value characterises 100% shale (note it is characterising shale and not clay). The advantage of the NMR bound water curve is that we know that its value in clean sand should be zero (or possibly a few per cent to account for small amounts of dispersed clay and silt). The fact that it never actually falls to 0 or even close, suggests that even the top of the reservoir has a high shale volume.
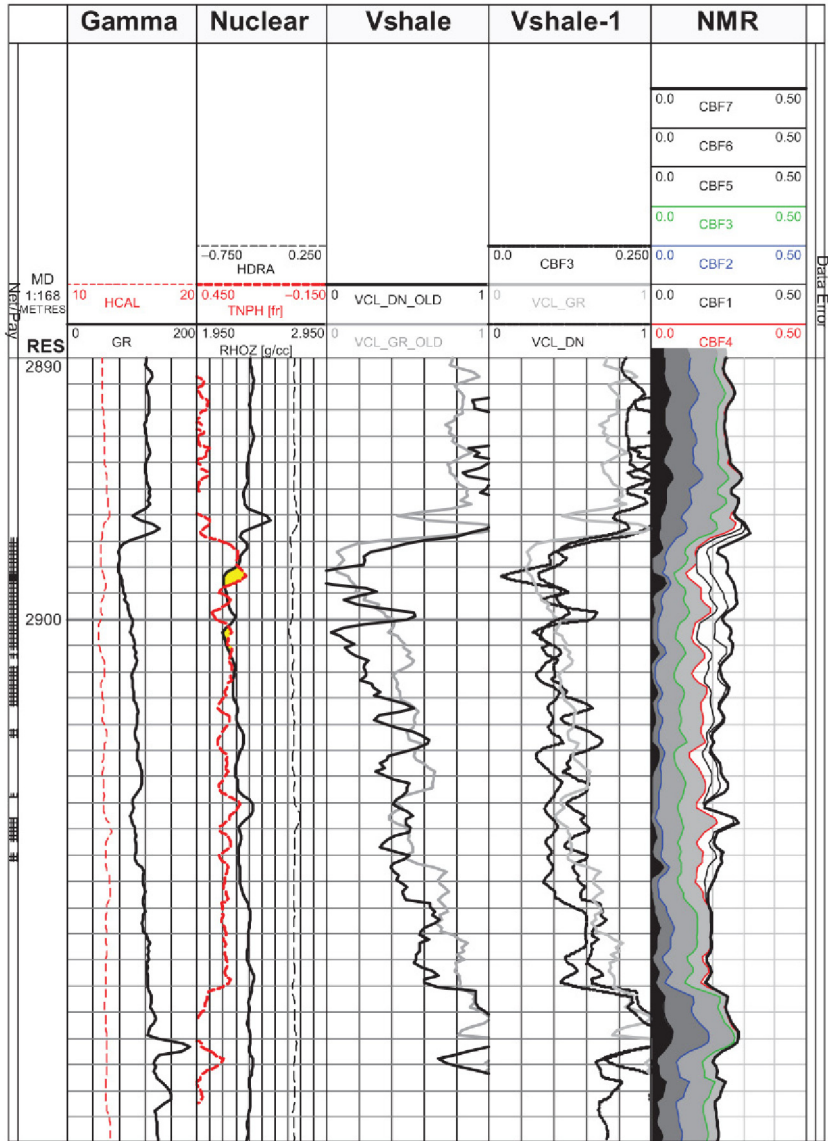
FIGURE 6.11   Logs and shale volume curves for a shaly-sand gas reservoir. The original shale curves calculated from gamma-ray (grey) and density-neutron (black) are shown in track 4 (V_{shale}). Revised shale volumes calculated with reference to the NMR log are shown in track 5 ($V_{shale-1}$). The NMR 'bins' are shown in the RH track on a scale of 0 - 0.5.

Contrast this with the gamma ray where we require a bed of clean sand in order to be confident of the clean point.

The right-hand track shows some X-ray diffraction measurements of clay content, made on core chips, through the reservoir. Although these are mass fractions and not volume fractions they nevertheless show that the reservoir

TABLE 6.2 Shale Parameters Obtained from Figure 6.11 Assuming the Top of the Sand is Clean

	Clean	Fluid	Shale
Gamma ray	80	–	140
Density	2.67	1.00	2.35
Neutron	−0.03	1.00	0.48

has a clay content of typically 50% and therefore an even higher shale volume (glauconite is also shown).

Even in the cleanest part of the reservoir the bulk fluid volume reads 12% corresponding to a shale volume of 43% (i.e. 12/28). The NMR bound fluid curve was used to guide the selection of the clean gamma-ray value and the fluid point for the density–neutron shale volume. A trial and error approach was used to get the new shale volume curves to roughly track the curve from the NMR. The new parameters are given in Table 6.3.

The resulting curves show shale volumes that are 20–30% higher than originally calculated and which are more consistent with the X-ray diffraction data. Even at the top of the reservoir shale volumes of 30–40% are calculated.

Although, in the example the bound water volume was used to guide picking of parameters for more conventional shale indicators, it could have been used as the basis of a shale volume curve in its own right. The key assumption is that all the bound fluid is associated with shale so that that zero bound fluid volume corresponds to zero shale volume. The 100% shale value would then be 28% bound fluid volume. Note that the NMR data was simply the well-site generated curves using default parameters, there is no reason why the raw NMR data should not be re-processed to come up with an improved bound fluid volume.

6.5.5 Geochemical Logs

Geochemical logs were introduced in Section 5.7. The basic outputs are dry weight fractions for a large number of elements. These in turn are used to estimate the mineralogy at a particular depth; this is basically an optimisation

TABLE 6.3 Shale Volume Interpretation Parameters Found Using the NMR

	Clean	Fluid	Shale
Gamma ray	55	–	140
Density	2.67	0.5	2.35
Neutron	−0.03	0.5	0.48
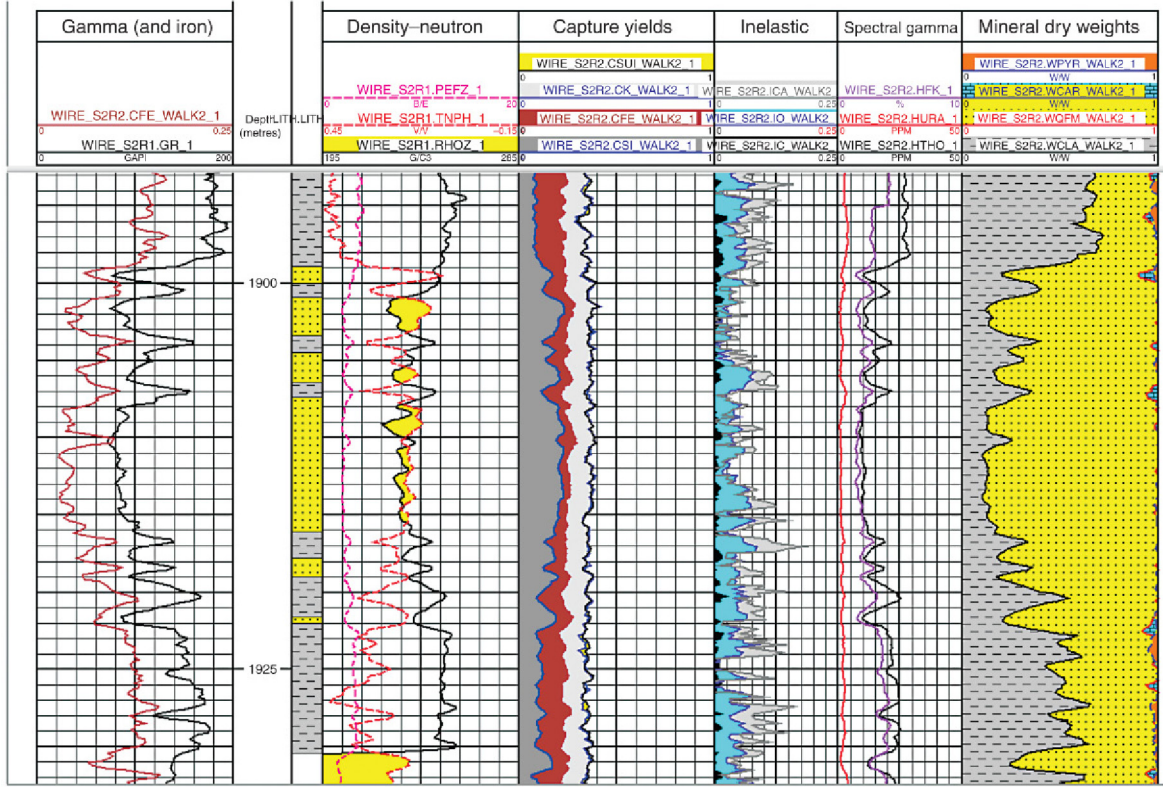
FIGURE 6.12 A geochemical log recorded over a sand-shale interval. The right hand track shows the dry weight of clays (grey dashed left), dry weight of QFM (dotted, middle) and dry weight of carbonate (minor component on right) estimated from the elemental make-up.

process that seeks to find the best combination of minerals to give the observed proportions of elements. Various models have been proposed to accomplish this and field specific ones may be available.

In addition to the dry weight fractions of the different elements, most tools also produce curves that give estimates of the relative amounts of the significant types of minerals. This includes a clay fraction, a carbonate fraction and a QFM fraction (quartz–feldspar–mica). An example is shown for a simple sand-shale sequence and it can be seen that the clay fraction shows the same trends as the more conventional shale indicators of gamma ray and density–neutron separation (Fig. 6.12). Note that the tool is calculating 30–50% 'sand' in the shales. This is actually quite reasonable as the tool is responding to the chemical make-up of the formation and not its texture. Shales normally do include a significant amount of quartz in their make-up.

6.6   COMBINING SHALE VOLUME CURVES

There is a lot to be said for calculating shale volumes by more than one method. At the very least if several methods agree this improves confidence in the calculations. Most packages allow curves from different methods to be combined into a single, composite curve. Various ways of combining curves exist: the simplest involve simply averaging the curves (e.g. mean or median).

Probably the commonest method is to take the lowest value at each depth and some simpler packages actually impose this method. The rationale for this is that individual methods typically err on the high side. For example the linear gamma-ray method over-estimates shale volume when the shales are significantly denser than the sands and any gamma-ray method over-estimates shale volume if the clean point is actually taken in a thin limestone bed or a coal. The density–neutron method on the other hand over-estimates shale volume in the presence of heavy minerals as these have the affect of moving points towards the lower right where the shales invariably plot.

It is however also possible to err on the low side. The density–neutron method for example will give zero shale volume in bad hole because the density reads too low. Gas will also move density and neutron points to lower values which, as we saw earlier, can result in a shaly-sand looking clean. All these artefacts can be satisfactorily accounted for by a more elaborate programme. For example density–neutron shale volumes could be excluded from inclusion in bad hole. But, at the end of the day, software is just a tool and to get the best out of it the user must be prepared to sometimes manually over-ride it. It is largely a matter of personal choice whether to try and write a 'programme' that uses the optimum curve or combination at every depth or whether it is easier to just specify where a particular method is to be applied. Just because we can calculate a large number of shale volume curves does not mean we have to use them all.

Chapter 7

Log Analysis Part I: Porosity

Chapter Outline

7.1   INTRODUCTION TO POROSITY

Porosity is the most important output from log analysis, if for no other reason than it is a necessary condition to form a reservoir. In regional work porosity may in fact be the only output required as it is typically used for trend work, to predict how it varies with depth, for example. In conventional volumetric work it is arguably the most fundamental property that is calculated as saturation is defined in terms of porosity, permeability relations invariably include a porosity term and net (pay) normally involves a porosity cut-off.

As discussed in Chapter 2 although porosity appears to be a straight-forward concept, the interaction of water with silicate and carbonate minerals complicates matters. The water that adheres more or less permanently to the surface of the mineral grains can either be considered as part of the pore space or not. Clays in particular bind large volumes of water because of their high surface areas. Indeed, some clays are so intimately associated with water that it is arguably part of their chemical structure.

The way the bound water is treated leads to the two end member definitions of porosity: *total* and *effective*. These were discussed in detail in Chapter 2 but

in summary there are pros and cons to both descriptions. Total porosity has the least ambiguity: the pore space is by definition equal to the volume of water and hydrocarbon regardless of its chemical environment but effective porosity is closer to what is of interest economically and can be more reliably calculated in the presence of shale.

Modern core measurements are generally closer to a total porosity because they are made with inert fluids on cleaned and dried plugs. The quantitative relationship between them was given in Eq. 2.1 which for convenience is repeated below.

$$\varnothing_T = \varnothing_E + V_{clay} \cdot \varnothing_{clay} \qquad (2.1)$$

Where V_{clay} is the volume fraction of clay (clay volume) and $\varnothing_{clay}$ is the porosity of the clay, in other words the volume fraction of water in pure clay. Suffixes 'T' and 'E' refer to total and effective porosity, respectively.

7.2   POROSITY CALCULATION FUNDAMENTALS

A porous rock can be considered as a mixture of fluid and matrix (i.e., the solid minerals). In order to calculate porosity a physical property is needed that shows a good contrast between the two components. The most familiar examples are acoustic velocity, density and neutron scattering cross-section. Other examples are neutron absorption cross-section(s) and magnetic resonance response.

In general the physical property of the mixture is given by:

$$X = X(v_f X_f, v_m X_m) \qquad (7.1a)$$

Where v_i is the volume fraction and X_i is the physical property of pure i. The volume fraction of fluid v_f is simply porosity and since the volumes sum to unity we have

$$X = X(\varnothing X_f, (1-\varnothing)X_m) \qquad (7.1b)$$

In principle this can be re-arranged to give porosity in terms of X.

$$\varnothing = f(X, X_f, X_m) \qquad (7.2)$$

As the properties of the pure components are known, this is 'one equation in one unknown'. The properties of the pure fluid and pure matrix – X_f and X_m – are the 'parameters'.

Density has the desirable property that the density of a mixture is the weighted average density of the components so that Eq. 7.1 takes the simple form:

$$\rho = \varnothing\rho_f + (1-\varnothing)\rho_m \qquad (7.3a)$$

which can be re-arranged to:

$$\varnothing = \frac{(\rho - \rho_m)}{(\rho_f - \rho_m)} \qquad (7.3b)$$
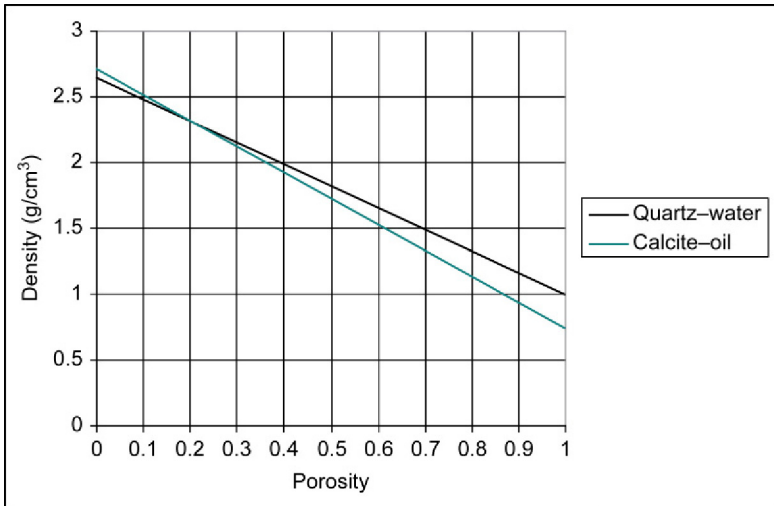
FIGURE 7.1   Relationship between density and porosity for a mixture of quartz and water (water-bearing sand) and calcite and oil (oil-bearing limestone).

This equation is quite general and therefore universally applicable, this is one of the reasons for the density often being the preferred porosity log. Graphically Eq. 7.3a is the straight line crossing from the matrix density at 0% porosity to the fluid density at 100% (Fig. 7.1).

Unfortunately, other physical properties generally do not give such simple, predictable relationships with porosity, although for certain formations and/or limited porosity ranges similar equations may be applicable (e.g. the Wyllie equation relating porosity to the sonic log discussed in the subsequent section).

Many techniques make use of a combination of logs, which is best visualised on a cross-plot. Cross-plot methods are discussed in more detail later but for now we note that the objective is still to find a clear distinction between matrix and fluid.

7.3   SINGLE LOG POROSITY METHODS

Single log methods are based on a one-to-one relationship between a physical property and porosity. The most widely used properties are the so-called porosity logs: density, neutron and sonic but NMR is becoming increasingly accepted as a primary porosity measurement as well. As always the log analyst has to decide two things when using a single log method:

1. The form of the relationship between the log and porosity.
2. The choice of the parameters for the particular problem under consideration. These determine the precise form of the relationship.

7.3.1 Density Porosity

For the density log Eq. 7.3b applies, this is repeated below.

$$\emptyset = (\rho - \rho_m) / (\rho_f - \rho_m) \tag{7.3b}$$

Graphically, this is equivalent to reading-off a porosity from a straight line drawn between the end points at 0 and 100% porosity (Fig. 7.1). As in this case there is only one equation to choose from, the log analyst is left with the task of setting the parameters; that is the densities of the matrix and the fluid.

The matrix density depends on the mineralogy: for sandstones it typically varies from 2.64 g/cm^3 to 2.68 g/cm^3. Limestone (calcite) normally has a matrix density that is very close to 2.71 g/cm^3 but dolomite can have a range of values from 2.84 g/cm^3 to 2.9 g/cm^3 depending on its precise composition and mode of formation.

The fluid density depends on the composition of the fluid, which will certainly include water but may also contain hydrocarbon. In general the composition will change with depth because of variations in reservoir quality and saturation-height effects. Even if the formation is known to be water bearing, the fluid density can vary from 0.95 g/cm^3 to 1.25 g/cm^3 depending on salinity, pressure and temperature. If these are known the water density can be calculated quite accurately (see Fig. 7.2).

Oil and gas densities can also be precisely calculated, but reservoir rocks will never be completely saturated with hydrocarbons and so the tool will always respond to a mixture of hydrocarbon and water. In general, as mentioned earlier, the relative amounts will vary with depth. A further complication is that the shallow depth of investigation of the density tool means that it is actually responding to a mixture of reservoir fluids and mud filtrate. In the case of a gas well drilled with oil-based mud there could be three different fluids affecting the density log!

By using combinations of logs it is often possible to have a reasonable idea of which fluids are within the volume of investigation of the density log and then to account for their effects. This is discussed later.

7.3.2 Parameters and Uncertainty

It is worth noting that because porosities are normally much less than 50%, they will be more sensitive to a small change in the matrix density than to a similar change in the fluid density. As a general rule, the lower the porosity the more sensitive it becomes to changes in matrix density and the less one has to worry about fluid density (see Fig. 7.1 again). For this reason a constant fluid density value will often produce an acceptable porosity curve in the hydrocarbon leg, even when in reality the fluid consists of variable mixture of water, hydrocarbon and mud filtrate. Fortunately, the matrix density is often constrained to lie within a narrow range: for most sandstones the matrix density lies between 2.64 g/cm^3
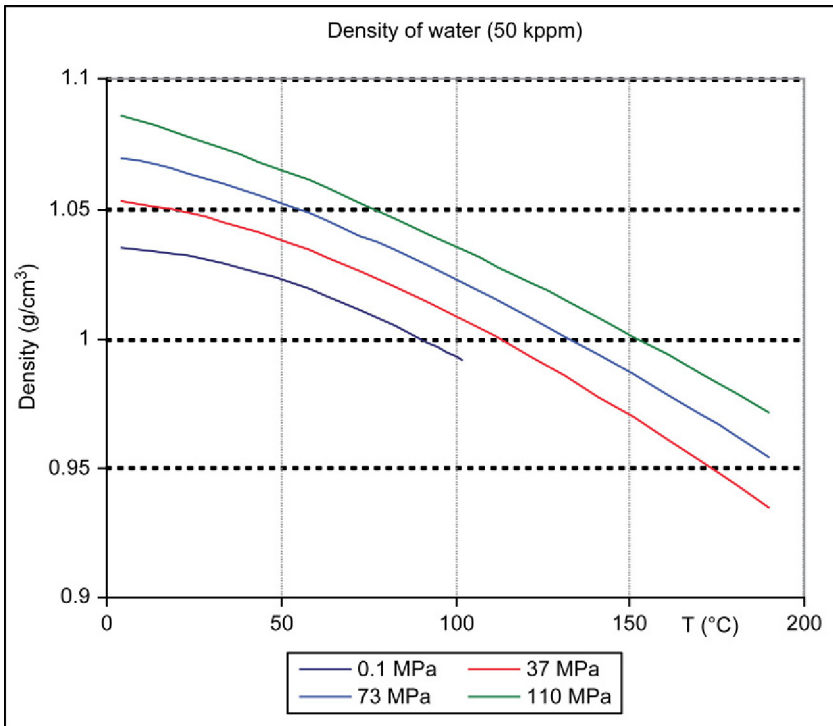
FIGURE 7.2 Density of water as a function of temperature and pressure. Salinity also affects density; these graphs are for a salinity of 50 kppm.

and 2.68 g/cm^3 and for limestones it is likely to be very close to 2.71 g/cm^3. Although a small uncertainty in matrix density produces a larger uncertainty in porosity than the same uncertainty in fluid density, it should be remembered that in an oil or gas leg the fluid density can potentially vary more than the matrix. For a gas-bearing reservoir fluid density could vary from close to 1 g/cm^3 for pure water to less than 0.1 g/cm^3 at high gas saturations.

There are a number of ways the parameters can be picked. The methodology chosen will depend on the time available, the purpose of the analysis, additional data and prior knowledge of the reservoir. Thus for a quick look analysis of a water-bearing sandstone, 'text-book' values for matrix and fluid densities of 2.65 and 1.0 g/cm^3 can be quite acceptable. For a field study more care will be needed and if the field is a tight gas reservoir great care will be needed picking the matrix density.

If core data is available, the grain density can be found from the individual grain densities of the core plugs (even the most ardent supporter of log analysis for porosity determination is forced to admit that core data is useful for providing matrix density). This may involve constructing a histogram or simply calculating a suitable average.

7.3.3 Shale Volume and Porosity

The matrix density is often assumed to be constant throughout the reservoir, even if it is known to vary as a result of changing mineralogy. How reasonable this assumption is depends on the minerals present and how their proportions vary from place to place, as well as the purpose of the interpretation. One class of minerals that normally are accounted for are clays. There are two reasons for this:

1. The volume of clay – or at least 'shale' – is normally calculated as part of a log analysis.
2. Clay minerals have high grain densities and therefore could shift the grain density significantly.

In addition, it is the clays that give rise to most, if not all, of the difference between total and effective porosity and so clay volume will certainly be needed for an effective porosity model.

In order to account for clay, the density porosity equation is modified to include a V_{shale} term (or if possible a V_{clay} term). This means that a third parameter the density of the shale – or clay – is required. As will be seen in the subsequent section the exact value of this depends on whether a total or an effective porosity is being calculated.

The methodology is the same but as the mixture now consists of three components the equation giving the density becomes:

$$\rho = \varnothing \rho_f + (1 - \varnothing - V_{sh})\rho_{ma} + V_{sh}\,\rho_{sh} \tag{7.4}$$

The volumes of the three components sum to unity so that the volume of the matrix is now given by $(1 - \varnothing - V_{sh})$.

Equation 7.4 can be re-arranged to give porosity in terms of density and shale volume:

$$\varnothing = \frac{(\rho_{ma} - \rho)}{(\rho_{ma} - \rho_f)} - \frac{V_{sh}(\rho_{ma} - \rho_{sh})}{(\rho_{ma} - \rho_f)} \tag{7.5}$$

Equation 7.5 can calculate either a total or an effective porosity depending on the value of the shale density (ρ_{sh}). If the shale density is the wet shale value, in other words the density of the shale including all its associated water, then the porosity is effective. If the shale density is the dry shale value, that is after all the water has been eliminated then a total porosity has been calculated. It is actually quite difficult to determine the dry shale density whereas the wet shale value can simply be read from the logs (for accurate work one might use a histogram to find it). Consequently, Eq. 7.5 will most often be used to find effective porosity. For this reason it is actually easier to compute effective porosity and convert to total than the other way (even though total porosity is often thought to be more 'fundamental'). Many log analysis packages, for example calculate total porosity from effective.

Another way to look at the influence of the shale is to combine the matrix and shale terms in Eq. 7.5 into a single 'shaly-matrix' term

$$\rho_{ma} = \rho_{ma} + V_{sh}(\rho_{sh} - \rho_{ma}) \qquad (7.6)$$

which then allows the density porosity to be written in the same two component form as Eq. 7.3b. But now the matrix density depends on shale volume. If the effective porosity model is being used the shale density can be taken as the log reading in a suitable shale (see Chapter 6). For the total porosity model the dry shale or dry clay density will have to be estimated using either published dry clay densities or by considering other logs.

Notice that the multiplier of V_{sh} in the second term of Eq. 7.5 is similar to Eq. 7.3b but with the shale density replacing the log reading. The term is sometimes referred to as the 'shale porosity' or better, the 'shale total porosity'. It always refers to wet shale (dry shale has a porosity of zero).

$$\varnothing_{sh} = (\rho_{ma} - \rho_{sh})/(\rho_{ma} - \rho_f) \qquad (7.7)$$

It should be noted that all the earlier discussion has been framed in terms of 'shale volume', in principle it could just as easily have been stated in terms of clay volume but, obviously, this assumes there is a reliable way to quantify the amount of clay. As always one needs to be clear which is being used: sands can be free of shale but contain a significant amount of clay.

There is one further complication in the choice of porosity models. This is that strictly speaking V_{sh} will be smaller for the dry value than the wet. The earlier discussion shows that whilst mathematically the inclusion of a shale term is easily accomplished, from a practical point of view it brings a lot of additional complexity. For this reason it is often desirable to work with the simple two-component model and accept some lack of rigour and accuracy. If the reservoir consists mainly of clean reservoir rock or shale adopting the simple approach will not materially affect the volumetrics.

The density porosity is often said to be a total porosity (or even the total porosity). It is certainly true that the tool cannot distinguish unconnected porosity from connected and so for situations where the non-effective porosity is represented by enclosed vugs, for example the density porosity is a total value. Where the difference between total and effective porosity is due to porosity associated with shale however, the density log can give either porosity depending on the choice of parameters. As effective porosity is always less than or equal to total porosity, the matrix density for the effective porosity model should be less than for the total. Intuitively this makes sense, in the effective porosity model the shale, consisting largely of wet clay, is included as part of the matrix. This means some water with a low density is being included in the matrix. In the total porosity model, the water associated with the shale is specifically excluded from the matrix.

7.3.4 Porosity from the Sonic Log

The lower vertical resolution and lack of a universal relationship between compressional velocity (or slowness) and porosity means that sonic is not normally the first choice porosity tool. It does however have some advantages:

1. It can produce a valid log in bad-hole.
2. It can read through casing.
3. It does not lose accuracy at low porosities.

The sonic log reading certainly does depend on porosity but not necessarily in a linear manner and there is normally some dependence on 'fabric' as well. So for example a sonic tool responds differently to fracture porosity and inter-granular porosity. The density tool simply responds to the relative amounts of matrix and fluid regardless of how they are distributed.

The simplest of the sonic–porosity relationships is a linear one: Wyllie's (time average) equation.

$$\emptyset = \frac{(\Delta t - \Delta t_{\mathrm{m}})}{(\Delta t_{\mathrm{f}} - \Delta t_{\mathrm{m}})} \tag{7.8}$$

Wyllie's equation is precisely analogous with the density porosity equation with slowness/transit times (Δt) substituted for densities. Unlike the density equation however, there is no fundamental basis for the linear relationship. Of course the same issues arise with respect to parameter picking.

Like the density equation it can be extended to a three-component (matrix, fluid and shale) system. The parameters are the slowness' for the matrix, fluid and possibly shale. Unlike matrix density, matrix slowness is not a common output of routine core analysis (and is not likely to be very reliable anyway) so one is more reliant on 'text book' values. Values for pure minerals may not be suitable however. The transit time for pure calcite, for example refers to a single crystal, whereas the porous solid is more likely to consist of grains and the 'matrix' is better thought of as a solid mass of these.

Fluid acoustic properties are better defined however (Δt_{f}). Values of velocity and thus slowness for water in particular are well established and can be found at any temperature, pressure and salinity. Gases such as methane are also well characterised but oils can be more problematical. As with density, one is often interested in a mixture of fluids and moreover the mixture can be very complicated (e.g. formation water, gas and OBM filtrate). The acoustic properties of mixtures are not simple volume averages and so cannot be found as easily as for the density. In general a process of trial and error and educated guesses needs to be applied although a Wood's law relation is often assumed (Eq. 1.8).

Cross-plotting core porosity against sonic is still a good way to find the matrix and fluid transit times however. The matrix and fluid transit times are found by extending the best fit to the data to the 0% and 100% porosity points, respectively. The values obtained need not be the same as those expected for the

pure matrix (e.g. calcite for limestone) or fluid (e.g. water), they simply provide the best fit to the data. One common use of sonic porosity is to substitute for density porosity in bad-hole. In that case one wants the best match between the sonic and the density porosities in good-hole and the parameters are chosen to achieve that.

Despite the comments earlier, Wyllie's equation can work remarkably well, even if the parameters are simply text-book values. Well-consolidated sand-stones often give a good match between Wyllie and core data (e.g. the Rotli-egendes of the Southern North Sea). It also normally works well for any low porosity rock, this is not really surprising and is really just another example of the choice of fluid parameter having little influence at low porosities. Con-versely, unconsolidated high porosity sands such as those found in deepwater environments normally produce very poor agreement, with the Wyllie equa-tion, invariably over-estimating porosity and sometimes giving impossibly high values.

In order to cope with situations such as un-consolidated sands the equation is sometimes modified with a compaction factor (C_f). The equation is:

$$\emptyset = \frac{1}{C_f} \times \frac{(\Delta t - \Delta t_m)}{(\Delta t_f - \Delta t_m)} \tag{7.9}$$

Values of C_f vary from 1 in well-consolidated sands to 1.6 or more in un-consolidated sands. 'Compaction factor' really is just a fudge factor and to find a suitable value one needs something to match to such as a density-based porosity curve or core data. An alternative way of achieving the same end is to select the parameters Δt_m, Δt_f so that they give a good match.

Wyllie is the simplest of a large number of published equations with poros-ity as a function of sonic slowness. Two other equations that are commonly found in Chart Books and Commercial software are:

Wyllie–Rose:

$$\emptyset = 1 - (\Delta t_m / \Delta t)^{1-x} \tag{7.10}$$

Where x depends on the lithology: for sandstone it is 1.6, limestone 1.76 and dolomite 2.0.

Raymer–Gardner–Hunt (RGH):

$$\emptyset = -a - \sqrt{[a_2 + (\Delta t_m / \Delta t) - 1]} \tag{7.11a}$$

Where $a = (\Delta t_m / 2\Delta t_f) - 1$ \qquad (7.11b)

The RGH equation often does a better job of matching core data in carbon-ates. Some of these equations are shown graphically in Fig. 7.3.

Shale can be accounted for in the Wyllie equation using the same reasoning as applied to the density equation. Three parameters are then needed: transit
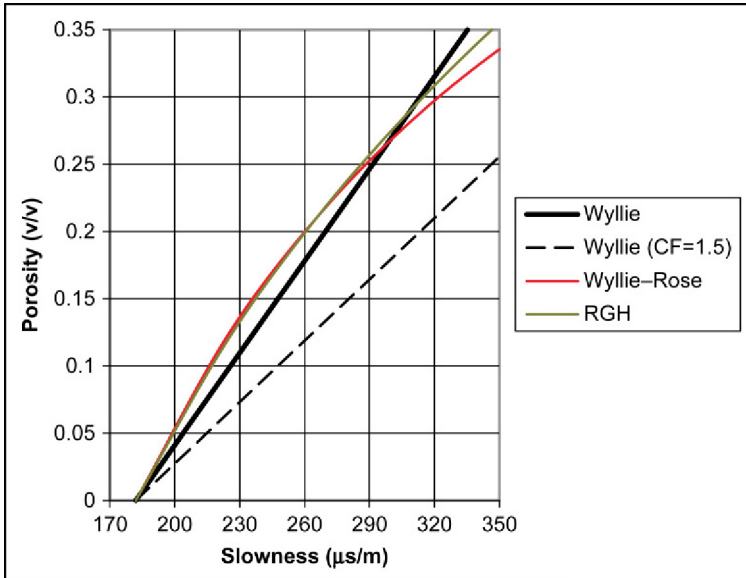
FIGURE 7.3 Some published sonic–porosity equations for water-bearing sandstones. (Matrix transit time 182 μs/m, fluid 620 μs/m.)

times for matrix fluid and shale. As always the latter is normally read from the logs in a suitable bed.

Secondary Porosity Index

Some log analysts advocate exploiting the differences in the sonic and density measurements to help distinguish connected and un-connected porosity. The latter is confusingly referred to as 'secondary porosity' although this is a different use to the more common definition of porosity created by diagenesis. The reasoning is that nuclear tools respond to all forms of porosity, regardless of how it is distributed whereas the sonic transit time is the fastest path through the formation, which means skirting around any un-connected pores. In other words the un-connected porosity does not affect the sonic ray and the sonic porosity will only include connected pores. The sonic porosity should thus be less than or equal to the density porosity and the difference shows how much of the porosity is un-connected. Unfortunately, there are a number of drawbacks with the technique, the principle one being that it does not generally work! Often one finds the sonic porosity is higher than the density porosity. There are cases where it does appear to work, however although there are few documented cases where an independent check has been made.

7.3.5 Neutron Log

An approximately linear relationship can often be found between the neutron porosity and the true porosity. Indeed in the right lithology the tool is set-up so that neutron porosity is numerically equal to the total porosity in a

water-bearing formation (providing the environmental corrections have been applied). Normally the 'right lithology' is clean limestone but the engineer can arrange for it to be sandstone or dolomite. But, in general, the precise relationship depends on the nature of the matrix, the fluid and the tool type. Most neutron tools are very sensitive to environmental conditions including temperature, pressure, salinity and the presence of even traces of neutron-absorbing elements. The net effect is that the measurement has to be corrected for the environmental effects or a different relationship needs to be established for each new well (The epithermal tools avoid many of these problems.)

A chart for converting a particular neutron porosity measurement to porosity is shown in Fig. 7.4. The chart assumes that the neutron tool has been set up for a lime matrix which can be seen from the fact that in limestone the output of the tool is identical to the porosity. In sandstone at intermediate to high porosities the true porosity is approximately 4 pu greater than the tool reads. At lower porosities the difference is smaller.

As with sonic and density tools the reading is also effected by fluid type and clay. Gas in particular can drastically reduce the neutron porosity. Clay minerals also tend to have disproportionate effect on the neutron log (a fact which is exploited when it is used as a shale volume indicator). As has been explained in the previous chapter clays always increase the neutron porosity sometimes to values that are un-realistically high for a true porosity (Fig. 7.5).
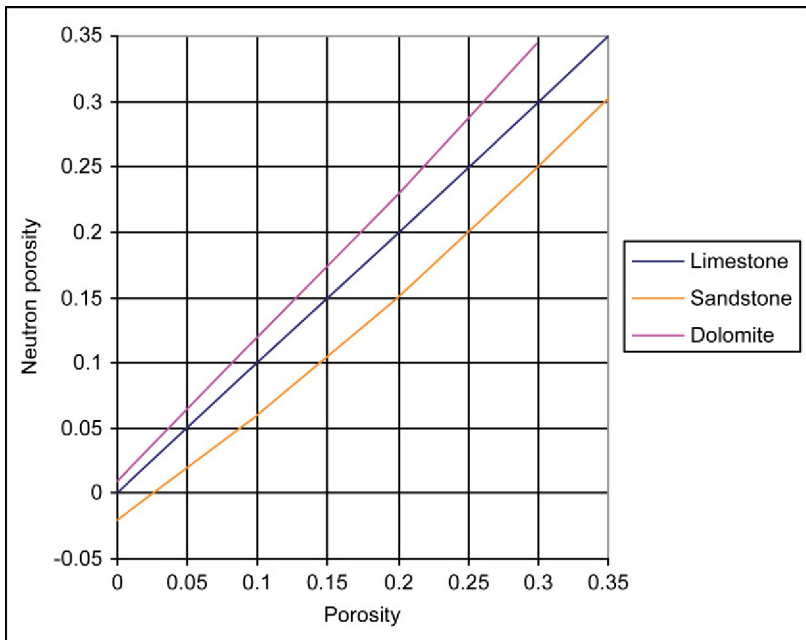
FIGURE 7.4   Chart for converting neutron porosity to true porosity. The chart only applies to Schlumberger CNL that has been environmentally corrected (for temperature, pressure, borehole size etc.). A different tool or a CNL that has not been corrected would require a different chart.
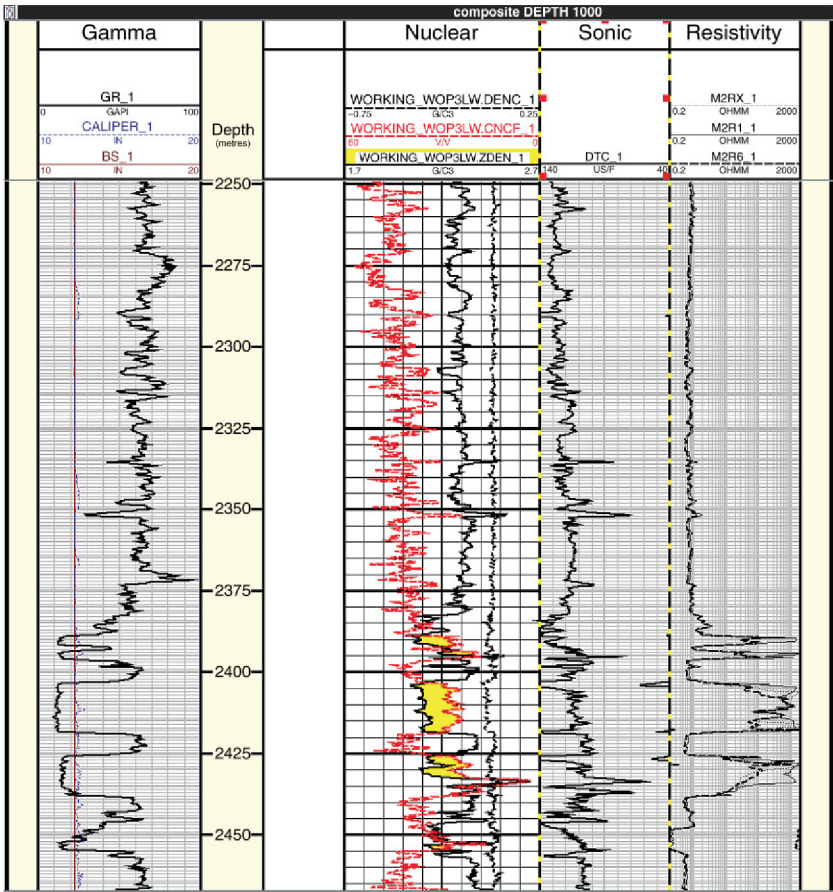
FIGURE 7.5 A set of logs from a deep-water play, offshore West Africa. Note the very high neutron reading in the shales (these are easily identified by their high gamma activity). At 2275 m, for example the neutron porosity is over 50% this cannot be explained by a high-water content alone.

7.4 METHODS INVOLVING MORE THAN ONE INPUT CURVE

There are a large number of methods that exploit two or more input logs. These have the advantage of being less sensitive to a problem with a single curve but conversely will introduce errors if there is a problem with one of the inputs that would be avoided if that curve were never used. Moreover the resolution is limited to that of the worst input curve and given that different tools generally have different volumes of investigation, problems may arise if the formation is heterogeneous.

Some of the simplest ways of combining the two measurements involve averaging the porosities estimated from the two logs. Three commonly used equations which combine density and neutron data are given here:

$$\emptyset_{nd} = (\emptyset_n + \emptyset_d)/2 \qquad \text{simple average} \qquad (7.12a)$$

$$\emptyset_{nd} = 0.7\emptyset_n + 0.3\emptyset_d \qquad \text{weighted average} \qquad (7.12b)$$

$$\emptyset_{nd} = \sqrt{(\emptyset_n^2 + \emptyset_d^2)} \qquad \text{RMS average} \qquad (7.12c)$$

Here $\emptyset_n$ is the neutron porosity and $\emptyset_d$ the density porosity calculated with the appropriate matrix but assuming the fluid is water. The density porosity is simply given by Eq. 7.3b (or Eq. 7.5b).

In some mature basins in North America where the matrix is known, the neutron log is corrected for the appropriate matrix and the density log is presented as a density porosity so the input values can simply be read from the logs. The porosities given by Eq. 7.12 may be total or effective depending on whether the input porosities have been corrected for shale (clay) or not.

7.4.1   Density–Neutron Cross-plot Methods

Cross-plots were introduced in Chapter 6 as a way of estimating shale volume. In fact they allow two properties to be estimated simultaneously. Cross-plot methods effectively solve two equations for two unknowns, for example porosity and shale volume or porosity and matrix density. This is a useful feature, especially in reservoirs where the mineralogy is quite variable (e.g. limestone–dolomite mixtures).

In Chapter 6 most of the discussion concerned the density–neutron cross-plot because it has the most predictable response to shale. It also has the most predictable response to porosity and as noted in Chapter 6 the lithology lines on published charts are marked with porosity values (normally every porosity unit). These can be used to construct 'iso-porosity' lines by connecting points with the same porosity on the three lithology lines (iso-porosity lines mean lines of equal porosity). These are orientated in a roughly lower right to upper left direction, more or less normal to the lithology lines (SE to NW if you prefer). Examples were shown in Figs 6.7 and 6.9. Depending on the accuracy required these lines could either be curved or constructed of two straight segments running from sandstone to limestone and limestone to dolomite (the latter construction is used in Figs 6.7 and 6.9). The simplest construction would use single straight lines.

As an example of how the density–neutron cross-plot is used to find porosity consider the short section of log shown in Fig. 6.8. Two points from the sand have been plotted on the cross-plot shown in Fig. 7.6. This is exactly the same grid that was used for finding shale volume in the previous chapter, the shale point has also been plotted. The readings of the two points are summarised in Table 7.1.

The clean sand point plots close to but just below the 30% iso-porosity line (the grey square). The porosity is close to 28%. The point appears to be so clean that this is both the total and the effective porosity. The porosity value
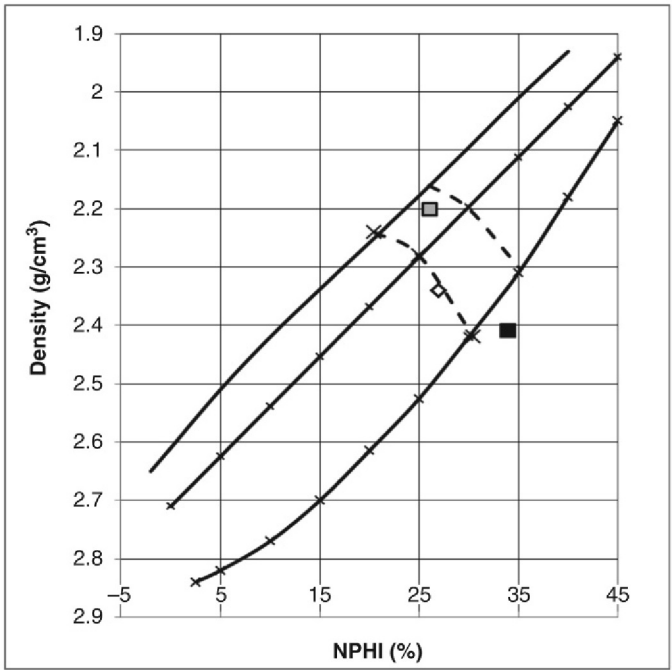
FIGURE 7.6   Density–neutron cross-plot used to find porosity.

TABLE 7.1 Density and Neutron Values Discussed in Text

Depth (ft.)	Lithology	Density (g/cm³)	Neutron porosity (%)
8160	Clean sand	2.20	26
8200	Shaly sand	2.34	27

could be fine-tuned by ensuring that the fluid point is the correct one for the water in the near well bore region. Inspection of the resistivity log shows very low readings ($0.3\,\Omega$m) showing that R_w is very low and the salinity is high. The water density may well be higher than the $1.0\,g/cm^3$, which the plot has been produced for. This would actually pull the sandstone line closer to the point but overall would have a very small affect on the porosity.

The second point (black diamond) is more difficult to interpret because it may be quite shaly. In fact using the approach suggested in Chapter 6 the shale volume from the density–neutron is about 50%, although the gamma ray suggests it is a lot less. Simply reading off the chart gives a porosity of 25%. This is related to a total porosity because the tools do not distinguish clay-bound water. But it may

not be the true total porosity because, as we know, the neutron tool is responding to the extra hydrogen provided by the clay. If we ignore the extra hydrogen the total porosity is 25%. The effective porosity is given by Eq. 2.1 (this time written in terms of shale):

$$\text{Ø}_T = \text{Ø}_E + V_{shale}.\text{Ø}_{shale} \tag{2.1}$$

The cross-plot gives the porosity at the shale point as close to 27%. Because this includes clay it will actually be an over-estimate of the shale porosity but ignoring that for the time being gives the effective porosity as:

$$\text{Effective porosity} = 25 - 27 \times 0.5 = 11.5\%$$

Since the shale porosity is an over-estimate this is a conservative value but even if the true shale porosity is only 22%, which is a large reduction, the effective porosity only increases by 2.5%.

7.4.2 Grain Density from the Density–Neutron Cross-plot

Although the lithology lines are labelled with different – well – lithologies, they could equally be labelled with grain density.

Sandstone = 2.65 g/cm^3
Limestone = 2.71 g/cm^3
Dolomite = 2.85 g/cm^3

The 'calcareous sand' points then become points with grain densities between 2.65 g/cm^3 and 2.71 g/cm^3 (depending exactly where they fall). The plot therefore enables us to estimate grain density, a key parameter for calculating a density porosity. This suggests another way of using the density–neutron cross-plot to find porosity. This is basically a two-stage process:

1. Use the cross-plot to estimate grain density.
2. Use the grain density as an input to the density–porosity equation.

The cross-plot can either be used to estimate an average grain density for the interval of interest or it could be used to estimate a single value of grain density at each level. The latter is very useful in so-called complex lithologies, where the mineralogy and hence grain density varies a lot within a reservoir. This approach was first developed by Bateman and Konen as an algorithm for computer log analysis.

To illustrate how this technique works we will again refer to Fig. 7.6. The clean sand actually plots just below the sandstone line which suggests a grain density of slightly more than 2.65 g/cm^3. If its position were accurately measured, the grain density is 2.67 g/cm^3. Calculating a density–porosity with that value and assuming for the sake of argument that the chart is compatible with the water density, the porosity comes out to 28.1%.

If we ignore the shale volume, then the Bateman–Konen approach should give the same porosity as was simply read off the cross-plot. The technique becomes useful when shale is accounted for, but because we are now using the cross-plot to find grain density we have to find some other way of estimating shale volume. In this case gamma ray is probably the best choice. At 8200 ft. the gamma ray gives a shale volume of about 25% (this is left to the reader to confirm). To find the clean lithology the log point needs to be shifted away from the shale point by an additional 25% (this is effectively running the shale volume calculation in reverse).

The shift to the density is given by:

$$\Delta\rho = (\rho_b - \rho_{sh})V_{sh} / (1 - V_{sh}) \qquad (7.13a)$$

Which in this particular case equals 0.025 g/cm^3. The shifted density is lower than the log value because, in this case, the shale density is higher than the log reading. The shift in the neutron is given by the same form of equation and equals about 2.5 pu again the shift is to a lower value. The shale corrected log point is therefore:

	Log point	Correction	Corrected point
Density	2.33	−0.025	2.30 g/cm^3
Neutron	34	−2.5	31.5%

This point plots just below the limestone line and has a grain density of 2.72 g/cm^3. Remember this is the matrix density of the non-shale component and in order to calculate porosity we have to use the three-component density Eq. 7.5b. Furthermore, since we are using wet shale readings we will compute an effective porosity (like it or not). For the values above, this comes out as just under 18%. To find the total porosity we have to add back the contribution from the clay-bound water, which – again – raises the issue of what is the shale porosity? If we take the value of 27% from the cross-plot the contribution is about 7 pu and the total porosity is 25%. But we know that the shale porosity read from the cross-plot is an upper limit and so the true value for the total porosity could be 23% or less.

Although the main objective in this chapter is to calculate porosity it is worth closing this section by asking what the significance of the high grain density – 2.72 g/cm^3 – at 8200 ft. is. Once again remember this is the matrix density of the shale-free component. One possibility is that it is a thin shaly-limestone bed but overall the formation seems to consist largely of quartz-rich sand, so that even at the point of interest there is likely to be a lot of quartz. To get a grain density of 2.72 g/cm^3 it requires the presence of some heavy minerals. Qualitatively, this could have been deduced from the fact that the density is plotting to the right of the neutron curve despite the low gamma activity suggesting relatively little shale in the formation. Of course another possibility is that the point actually contains the higher amount of shale suggested by the density–neutron combination and the gamma ray is actually under-estimating the shale volume. Without more information either explanation could apply.

7.4.3 Hydrocarbon Effects

Hydrocarbons will change the density, neutron porosity and sonic readings from the values expected in a water-bearing formation. The density for example is given by Eq. 7.3a, which can be written explicitly in terms of the densities of water, hydrocarbon and the water (or hydrocarbon) saturation:

$$\rho = \text{\O} S_w \rho_w + \text{\O}(1 - S_w)\rho_{hc} + (1 - \text{\O})\rho_m \qquad (7.3c)$$

How this is solved is described in Chapter 9 but note that in order to do this the saturation is needed and that relies on knowing the porosity! The density–neutron cross-plot to an extent avoids this chicken and egg problem because the hydrocarbons have opposite effects on the two input logs. In the presence of hydrocarbons density is reduced which results in porosity being over-estimated but the neutron porosity is underestimated (this is also why an average of the density and neutron porosity works even when the saturation is varying). On a conventional cross-plot this results in the points being moved upwards and to the left (or to the NW if you prefer a geographical description). For small shifts at least, the porosity can be accurately estimated without needing to do any of the number crunching described in Chapter 9. For large gas effects the points move so far from the lithology lines that it is difficult to estimate the porosity accurately, but even then a reasonable quick-look estimate is often still possible.

7.4.4 Other Cross-plots

Although the most widely used cross-plot uses the density and neutron logs because these curves give reasonably consistent responses for a particular lithology, charts are also normally produced for density–sonic and sonic–neutron pairs. The problem with these combinations is the lack of a universal relationship between the sonic log and porosity. The charts are often constructed assuming a Wyllie relationship but as discussed earlier there are many occasions when this does not apply.

Figure 7.7 shows a blank density–sonic cross-plot. The lithology lines correspond to limestone, sandstone and dolomite and if the relation between slowness and porosity is given by the Wyllie time-averaged equation these are perfectly straight and converge at the fluid point (density of $1.0\,\text{g/cm}^3$ and a slowness of $189\,\mu\text{s/ft.}$ for water). This is because the equations linking porosity to density and porosity to sonic are both linear. As was explained earlier however, the Wyllie equation is just the simplest of a large number of empirical equations linking porosity and slowness. There is no fundamental reason for it to apply in practice and if it does not apply, the points will plot well away from the lithology lines. A different sonic porosity equation will produce curved lithology lines and post-1980 chart books frequently include lithology lines based on both the Wyllie equation and the RGH equation.
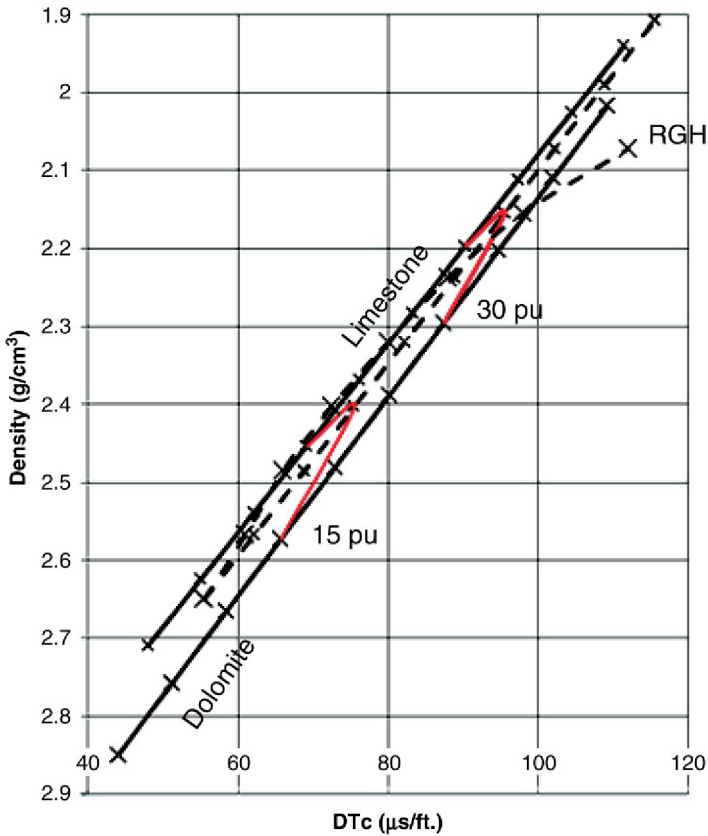
FIGURE 7.7 A density–sonic cross-plot showing lithology lines and two iso-porosity lines. The lithology lines for limestone and dolomite are marked and the sandstone lines are dashed. Two sandstone lines are shown corresponding to the Wyllie time-averaged equation (straight) and the Raymer–Gardner–Hunt (curves and marked 'RGH'). Porosities are marked every 5 pu. Two iso-porosity lines for 15% and 30% are shown (light grey line [red line in the web version]).

The relative values of the matrix densities and slownesses mean that the iso-porosity lines are not even approximately linear and actually form a series of 'hair-pins'. This means that even if the appropriate sonic–porosity relation is known, a slight change in density and/or sonic reading can lead to a large change in interpreted lithology and porosity. These cross-plots are therefore best limited to shale volume estimation.

7.5 NUCLEAR MAGNETIC RESONANCE

Nuclear magnetic resonance (NMR) is unique among logging tools in that it responds exclusively to one element: hydrogen. Furthermore it will only respond to hydrogen that is part of a fluid, which in the sub-surface means water, oil or gas. Hydrogen that is part of a clay layer is not detected. Not only can it

determine how much hydrogen, in the form of water or hydrocarbon, is in the formation but it can identify through the T2 value its microscopic environment and in particular how mobile the hydrogen is. For example, it can distinguish water that is tightly bound to clay from water that is free to move around the centre of a large pore.

The neutron tool is also supposed to respond mainly to hydrogen but it cannot distinguish hydrogen that is part of the chemical make-up of clay, from water that is bound to clay from water is part of the effective porosity. One consequence of this is the high neutron porosities measured in argillaceous rocks.

As discussed in Chapter 5 the defining feature of an NMR log are the T2 distributions (plotted at every depth increment). The more tightly bound the hydrogen the shorter is the T2 value. Hydrogen in a water molecule that is tightly bound in a clay has a T2 value of a few milliseconds. Whereas hydrogen in the gas phase within the inter-granular pore space of a coarse-grained sand, has a T2 value of over 1000 ms. At any depth the T2 distribution reflects the complete range of fluid hydrogen mobility within the depth of investigation of the tool.

The NMR tool responds to the quantity of hydrogen in a fluid state regardless of whether it is part of a water, oil or gas molecule. The conversion of the T2 distribution to porosity implicitly assumes the pore space is filled with water (with a hydrogen index of 1). If water is replaced by oil or particularly gas with a lower hydrogen index, the total porosity measured by the tool will be lower than the true value. Gas creates additional problems for the measurement because it is so mobile; this essentially means the T2 for the gas phase is actually longer than the tool normally measures. In other words gas not only reduces the amount of hydrogen in the formation, but not all of that hydrogen is even recognised.

Nevertheless, if the formation is water bearing the NMR tool will provide an accurate measurement of total porosity regardless of the mineralogy. No other porosity tool is insensitive to the reservoir lithology. Furthermore, effective porosity can also be measured directly. In principle the tool can measure the volume of clay-bound water and effective porosity is then simply:

$$\emptyset_E = \emptyset_T - \text{clay-bound water volume} \qquad (7.13)$$

An example of the use of a NMR porosity log is shown in Fig. 7.8. This shows a 200 ft. interval of shales and shaly sands that have been logged with conventional density, neutron and induction resistivity logs and a mandrel-type NMR tool. The T2 distributions have been separated into approximately 10 different bins, which are displayed in track 5 using different colours. The shorter T2s have been shaded blue and the longer T2s have been shaded yellow/red. Track 6 shows total and effective porosities calculated from the density–neutron and NMR tools. The NMR total porosity corresponds to the sum of all the bins shown in track 5, the effective porosity uses the higher T2 bins only. The black curves have been calculated from the density–neutron combination and suggest that the NMR effective porosity is too high (it includes some of the dark grey bins). In fact even in the shales an effective porosity is calculated. In other
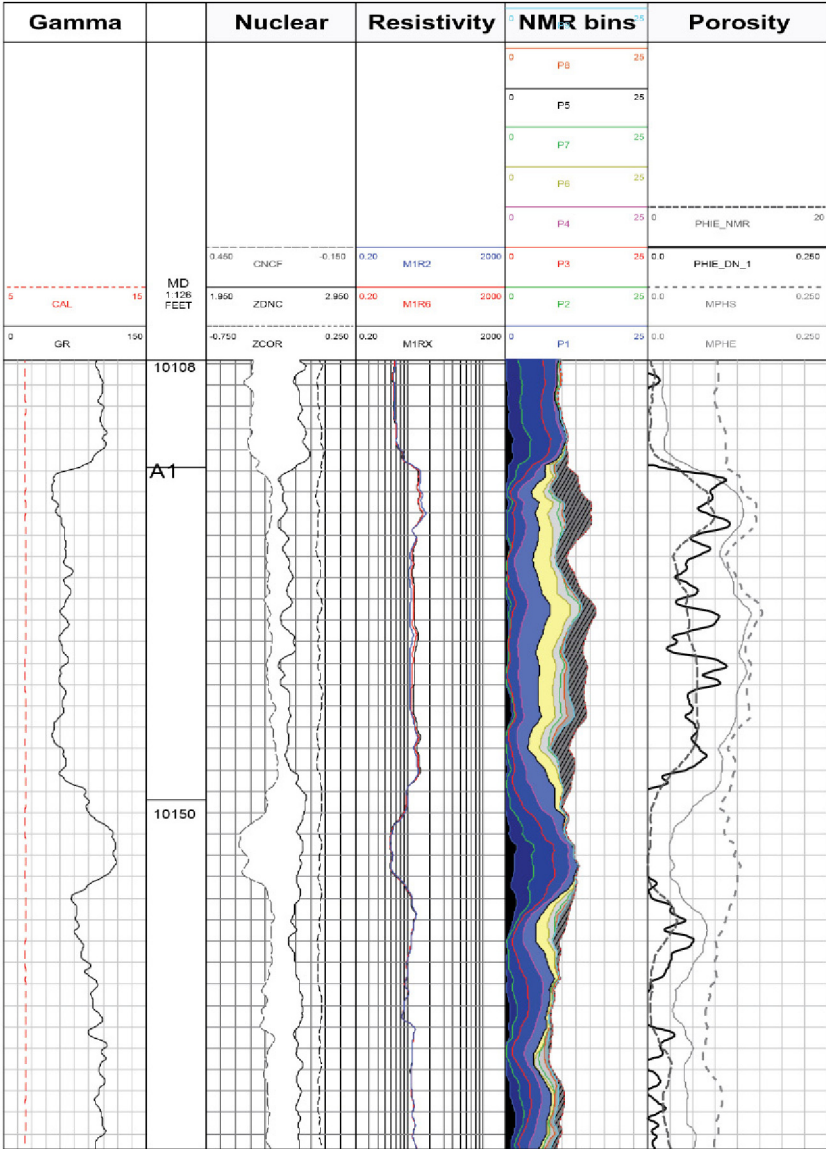
FIGURE 7.8 Conventional and NMR porosity data for an interval of shales and water-bearing shaly sands. The right-hand track shows effective porosities calculated using the density-neutron combination (black) and from the NMR (dark grey). The NMR effective porosity (light grey) is provided by the logging contractor using default T2 cut-offs and is consistently higher than the effective porosity from the density-neutron combination.

words the T2 cut-off needs to be increased. A new effective porosity has been calculated using only the highest T2 bins and produces a better agreement to the conventional effective porosity.

Because NMR can distinguish the different types of porosity it offers the possibility of objectively measuring shale porosity, which is essential for converting between total and effective porosity. Furthermore, the dry shale density can be found from the shale porosity and its density.

In the case of the example the porosity in the shales varies from 7.5% to 10% and the shale density is 2.6(±0.05) g/cm^3. By re-arranging the density–porosity equation the dry shale density is given by:

$$\rho_{dsh} = \frac{(\rho_{sh} - \emptyset_{sh}\rho_w)}{1 - \emptyset_{sh}} \tag{7.14}$$

Using the above values this gives a dry shale density of between 2.72 g/cm^3 and 2.78 g/cm^3.

7.6 INTEGRATION WITH CORE DATA

If core is available, porosity can be measured independently of the logs (Section 3.5). In addition core allows grain density to be measured, which is a key parameter to any porosity model that uses the density log. In principle other matrix parameters can also be measured on core plugs, for example the acoustic parameter Δt_m although acoustic measurements are difficult to make on plugs and may produce misleading results.

Direct measurements of porosity can be used in one of two ways:

1. To calibrate the logs.
2. As an independent check on log analysis.

In the former case it is tacitly assumed that the core porosity is correct and the analyst then seeks an equation and parameters that transform the log readings to the core value. Potential problems with this approach are as follows:

1. The changes the core undergoes between reservoir and laboratory.
2. The volume of the core plugs compared to the volumes of investigation of the logging tools.

In addition there is normally a depth offset between the core depths and the log depths (the core depths are ultimately determined by the driller). This can amount to several metres and is normally accounted for by shifting the core to match the log. The appropriate shift can either be found by overlaying the porosity measurements on the density curve (or any other curve that shows a strong relationship to porosity) or by matching the core gamma ray to the gamma-ray log. For single cores a simple block shift is normally sufficient but if multiple cores have been cut, different shifts may be required even when the cores are supposedly contiguous.

The changes the core undergoes are reasonably well understood and at least some can be quantified in the laboratory. Two important ones are the lower confining stress, which are used in laboratory measurements and the different fluids that are used in the laboratory. In routine porosity measurements the plugs are normally dry and filled entirely with an inert gas at relatively low pressure (e.g. helium, nitrogen, air). Other fluids that may be used are kerosene or mercury. These fluids are selected because they do not react with the minerals that make up the matrix. This is in stark contrast to water that is however, always present in the sub-surface. Routine measurements are also normally made at low confining stresses whereas in a typical reservoir the rock is compacted by several kilometres of over-burden.

7.6.1 Confining Stress

Historically, the laboratory measurements were made at confining stresses of a few hundred pounds per square inch (see Chapter 3). In the reservoir, the rock is normally under a much higher stress due to the weight of the hundreds or even thousands of metres of material above it. A commonly used rule of thumb is that the overburden produces a stress of 1 psi/ft. of vertical depth (i.e. 3.28 psi/m). This stress is known as the lithostatic stress (LS). The actual stress that the reservoir rock is under, is less than this because the pore pressure (PP) of the fluids opposes it. The so-called net over-burden (NOB) is given by the difference between LS and PP. As an example consider a normally pressured reservoir at 2500 m true vertical depth. The PP is given by:

$$PP = 1.42 \times depth = 3550 \, psi$$

(for now just accept the formula is true, we will see where it comes from when we discuss fluid distribution later in the book).

The LS is given by:

$$LS = 3.28 \times depth = 8200 \, psi$$

And the NOB stress is given by the difference, which at 2500 m is 4650 psi. This is an order of magnitude higher than the typical confining stress applied during laboratory measurements.

Before going any further it is important to note that the way stress is typically applied in the laboratory and in the reservoir are different. In the laboratory it is normal to apply a 'triaxial' stress that is to say the same pressure is applied everywhere to the core plug surface. The plug can then respond to the pressure by contracting in all three dimensions. In the reservoir the rock is unable to deform in the horizontal directions and shortens in the vertical direction only. This is uniaxial stress and in order to simulate it using a normal core holder, the confining stress to the plug has to be lower than the calculated NOB stress. A good rule of thumb is to reduce it by a factor of 2/3 although the exact amount

is a function of Poisson's ratio. So, for the example of the 2500 m reservoir, the NOB to apply in the laboratory, in order to simulate reservoir conditions, is 3100 psi. Given that we calculated the LS using a rule of thumb (1 psi/ft.), there is not a lot of point calculating the triaxial to uniaxial correction exactly (see inset).

Often, measurements of the PP are available so that at least they can be found directly. Furthermore, in a deep-water well the LS is given by the sum of the stress due to the water and the stress due to the rock. The former can be found quite precisely. So, if the example well had actually been drilled in 1000 m of water the LS would be given by:

$$LS = 1000 \times 1.42 + 1500 \times 3.28 = 6340 \text{ psi}$$

In the past the normal way to account for the higher stress in the reservoir was to measure how porosity changed between laboratory and reservoir conditions on a small sub-set of plugs (this was often part of the special core analysis programme). The information was then used to develop a general transform to convert all the laboratory measurements to reservoir conditions. Nowadays, it is common practice to actually measure the porosity of all the plugs at the simulated reservoir stress. Note that measurements are still also made at – relatively – low confining stresses if only to allow comparison with historical data sets.

Determining Simulated Reservoir Stress

The LS is given by the weight of all the rock above the reservoir (or the depth of interest). In general this is given by:

$$LS = g \int \rho(h)\,dh \qquad (7.15)$$

Where g is the acceleration due to gravity, $\rho(h)$ is the density at depth h. (If you do not like integrals then think of 9.13 as adding up the weight of a large number of thin slithers of the overlying rock). LS is actually given in Pascals (Pa) if the density is in kilogram per cubic metres, depth is in metres. Acceleration due to gravity, g, is more or less constant everywhere at $9.8\,\text{ms}^{-2}$. If we want to work in pounds per square inch and gram per cubic metres the equation becomes:

$$LS = 1.42 \int \rho(h)\,dh \qquad (7.15a)$$

(the factor 1.42 takes care of the units conversion and the constant g). In a good hole, the density as a function of depth is the density log. But unfortunately it is rare to have a density log all the way to surface (or seabed) or even most of the way there, so some assumption has to be made. Earlier we suggested 1 psi/ft. is a good rule of thumb and it turns out this corresponds to an average density of $2.31\,\text{g/cm}^3$. This sits nicely inside the $2\text{–}2.7\,\text{g/cm}^3$ range occupied by most sedimentary rocks.

7.6.2 Other Core-Log Integration Issues

Accounting for over-burden is a reasonably well-understood process that, moreover, often only results in small corrections. The other problems reconciling core and log data cannot be dealt with using a simple workflow and may actually be insurmountable. Most of these problems are to do with scale: typically core measurements are made on volumes that are far smaller than the volumes of investigation of logs. Even whole core measurements are made on volumes that are small compared to those used by logs.

A related issue is that the core itself is never investigated by the logs. It is a statement of the obvious, but the part of the reservoir that was the core has gone by the time the logs are run. It is tacitly assumed that nothing significantly changes on the scale of the order of a metre in the vicinity of the borehole. In heterogeneous formations this will not be the case and not only will different logs end up looking at different average properties but they will all look at different average properties to the core.

In some situations operators try to improve the chances of coring the intervals of interest by 'by-pass coring'. This involves drilling the well to the target depth and running logs. The well is then side tracked a few metres from the original hole and coring commences at a depth determined by the logs from the original hole. In homogeneous formations the core can generally be satisfactorily matched to the logs but again in heterogeneous formations the distance between the original and the core hole may be enough to give some large changes. Ideally, the core hole should also be logged but often, in order to reduce costs, this is not done.

Obviously if the individual sands are thinner than log resolution, log analysis will give different porosities to core analysis even in homogeneous formations.

7.6.3 Log Calibration

Most of the large North Sea Oil and Gas fields were developed in the 1970s. The reservoirs generally consist of clean sandstones with porosities ranging from 15% to over 30%. With a few exceptions coring and core analysis worked well in these and consequently most operators acquired a lot of core during the exploration and appraisal phases. Furthermore, these sands were thick enough and homogenous enough that the problems discussed earlier were generally not an issue. The core was fully utilised for all aspects of reservoir description, not least for determining the petrophysical properties.

The most common approach was to use the core to actually calibrate logs. The workflow would typically involve the following:

1. Measuring porosity on plugs taken at more or less regular intervals of 1 ft. (or half or a quarter of a metre).
2. Converting the porosities from laboratory to reservoir conditions (using the methodology described earlier).

3. Depth matching the core to the logs.
4. Cross-plotting one or more logs against the core porosity and using regression to find the 'best' relationship between them.
5. Applying the relationship to the logs to generate a continuous porosity curve over the entire logged interval.

The most common calibration involved predicting porosity from the density log alone because of its good vertical resolution and the predictable relation of density with porosity. Normally, different relations were developed for water-, oil- and gas-bearing sands but these could be further sub-divided into relations for different formations and relations for wells drilled with oil- and water-based muds. The largest fields were appraised by 10 or more wells drilled over a period of several years so it was not unusual for the wells to use a variety of different mud systems and types of logging tools.

An example of such a calibration is shown in Fig. 7.9. Figure 7.9 shows a cross-plot of compaction corrected porosity against the density log for a 30-m interval of a high-porosity gas-bearing sand. The core has been depth matched to the logs. Although, ultimately porosity is the dependent variable – we seek to predict it from density – we have used a common convention of plotting porosity on the x-axis. The lines are then X on Y regressions. Two have been shown, the light grey line is the free regression, in other words the unbiased estimate of porosity from density. The black line is forced to go through a density of 2.69 g/cm^3 at zero porosity, that is the mean – and in this case modal – matrix density of the core plugs.

The parameters given by the two different regressions are summarised in Table 7.2.

The free regression will give the better estimate of porosity despite having an apparent grain density that is much higher than the true value and a fluid density that is undoubtedly lower than the true value. Remember, the parameters are not intended to correspond to real physical properties, all they do is produce the best match between the log and the core. Nevertheless, many petrophysicists prefer the forced regression so that the matrix point, at least, corresponds to a real measurement. In reality both relationships do a good job of estimating porosity for the majority of the interval.

The example is as simple as core calibration gets, but nevertheless it has worked very successfully in areas such as the North Sea. More complicated models use grain densities that are functions of shale volume and multiple regressions on more than one log (see Eq. 7.6, for example).

7.6.4 Using Core to Guide Log Analysis

Despite its successes, using core to calibrate logs has declined in popularity since its heyday in the North Sea. There are undoubtedly many reasons for this including that it has simply gone out of fashion. Two important objections to calibration however, are the very different scales of core and log
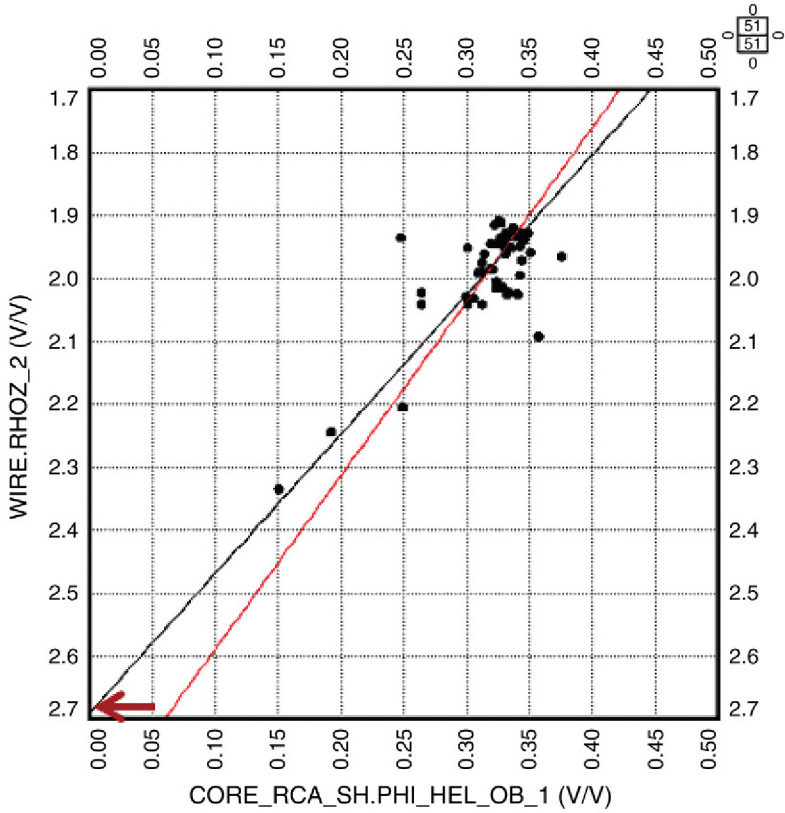
FIGURE 7.9 A cross-plot of core porosity against density for a high-porosity gas sand, showing regression lines for predicting porosity from density ('*X* on *Y*' in this case). The black line is forced through a density of 2.69 g/cm^3 at zero porosity (arrow). The light grey line (red line in the web version) is a free regression.

TABLE 7.2 Parameters Given by the Two Different Regressions

	Matrix density	Fluid density
Free regression	2.86	0.101
Forced regression	2.69	0.466

measurements and the requirement for different relations for every combination of reservoir fluid, mud and formation. Core porosity is now more commonly used as an independent check on log analysis and if necessary to guide parameter selection (this can be particularly important when invasion effects have to be accounted for).

The two most-effective ways of using core in this role are simply to plot the core porosity in the same track with the porosity(s) from log analysis and to cross-plot core porosity against the log-derived porosity.

7.7 OIL AND GAS SHALES

In theory most of the techniques described earlier can be applied to oil and gas shales, in practice there are two issues that mean estimating porosity is more difficult than most conventional reservoirs.

1. Matrix densities are not as predictable as with conventional sandstones and limestones (the other matrix properties are even more difficult to define).
2. Effective porosity is only a small fraction of the total porosity and so is a lot more uncertain.

There are various ways to tackle these problems but some of the so-called modern logs, in particular, NMR and geochemical tools, are particularly important. In addition core measurements, discussed in Chapter 3, are invaluable for ground-truthing the interpretation.

Page left intentionally blank

Chapter 8

Log Analysis Part II: Water Saturation

Chapter Outline

8.1 INTRODUCTION

Water saturation is the volume fraction of the pore space occupied by water.

$$S_w = \frac{V_w}{\emptyset} \tag{8.1}$$

where V_w is the volume fraction of water in the rock. Providing there are at most two fluids in the pore space it is related to the hydrocarbon saturation very simply by:

$$S_{hc} = 1 - S_w \tag{8.2}$$

We nearly always quote water saturation because log analysis relies on measurements that respond strongly to water. Sometimes, we talk about

'saturation' and take it for granted that this refers to water. This may not be obvious to everyone however and since ultimately it is oil or gas that will be sold, it is not unreasonable that 'saturation' refers to these. As always ambiguity and possible embarrassment can be avoided by being completely explicit.

If the reservoir has been interfered with by some artificial process (e.g. production or drilling) a third fluid may be introduced. This situation can complicate log analysis or at least increase uncertainty. Examples are invasion by mud filtrate and – on a larger scale – injection of gas, CO_2 or polymer.

Water saturation is computed from logs by applying Eq. 8.1. Porosity is estimated using one of the methods discussed in Chapter 7 and the volume of water V_w is estimated using a log or combination of logs that are particularly sensitive to water. Any log measurement that responds strongly to water and weakly to everything else is suitable. Because water is quite an unusual substance there are a number of possibilities. For example:

1. Dielectric constant.
2. Electrical resistivity.
3. Neutron scattering.
4. Neutron capture.
5. Nuclear magnetic resonance.

Although all these measurements are particularly sensitive to water, none of them respond exclusively to it and so applying Eq. 8.1 can be quite complicated in practice. It is also worth pointing out that most of these logs have shallow depths of investigation and therefore may not be capable of providing an S_w estimate in unaltered formation.

As water saturation is defined in terms of porosity it will, in general, differ depending whether a total or effective porosity model is being used. The two possibilities are illustrated in Fig. 8.1 and are sometimes distinguished as total or effective water saturation (S_{wt} or S_{we}). As with porosity itself, either possibility is acceptable and both have pros and cons. Inspection of Fig. 8.1 shows that in a normal inter-granular pore system, where the hydrocarbons are all located in the effective porosity, the total water saturation should be greater than the effective water saturation.

Regardless of which system is used it is obviously essential to be consistent, in other words to use total water saturation with total porosity and effective water saturation with effective porosity.

Since the volume of hydrocarbon should be the same in either model the following should hold true:

$$\varnothing_t (1 - S_{wt}) = \varnothing_e (1 - S_{we}) \tag{8.3}$$

This can be quite a useful check if a number of different analysis methods have been tried.

$$S_{we} = \frac{\text{free water volume}}{\text{effective porosity}} \qquad S_{wt} = \frac{\text{total water volume}}{\text{total porosity}}$$
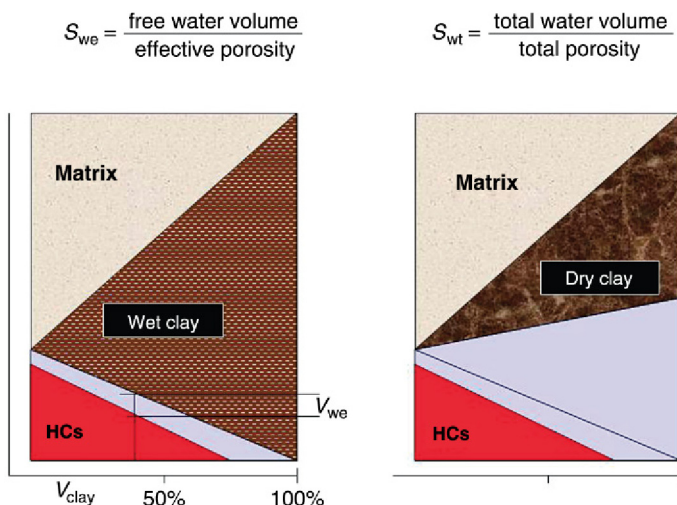
FIGURE 8.1 Diagrams illustrating the difference between effective (LH) and total water saturation and how they depend on clay volume. In the former clay-bound water is not counted as porosity and so the hydrocarbons occupy a greater fraction of the pore space. The two models agree in completely clean formations.

8.2 BASIC PRINCIPLES

8.2.1 Determining Water Volume

As already mentioned, in order to calculate water saturation a log is needed which responds particularly strongly to water. Fortunately, for once we are spoilt for choice. Water is actually a very unusual substance, although it is easy to take it for granted. For example, consider the melting and boiling points of water and the other 'Group VI hydrides' H_2S, H_2Se etc.

	MP (°C)	BP (°C)	MW (g)
H_2O	0	100	18
H_2S	−82.3	−60.3	34.1
H_2Se	−68.7	−41.3	81
H_2Te	−49	−2.2	129.6

All other things being equal, one would expect water, being the lightest compound, to have the lowest melting and boiling point. In fact they are far higher than any of the other hydrides, although the heavier compounds do follow the expected trends. These anomalies are a result of water molecules being able to form strong bonds to neighbouring molecules (this will be discussed in more detail in chapter 10). To break these bonds needs a lot of heat and so the melting and boiling points are high. The presence of these strong bonds results in water having unusually high values for a number of physical properties. These can be exploited in log analysis.

8.2.2   Dielectric Constant

One of the best properties to use is the dielectric constant – ε_r – also known as the relative permittivity. This is a fundamental physical property that defines the effect matter has on an electric field (roughly speaking it describes how rapidly an electric field decays in matter, compared to a vacuum). It is dimensionless as it compares the effect of matter to that of a perfect vacuum. For pure water it varies from 80 at room temperature to 55 at 100°C, at the frequencies used by logging tools. Oil, gas and most minerals that make up sedimentary rocks have dielectric constants of less than 10. (This basically means an electric field decays far more rapidly in water than most other substances found in sedimentary environments.)

The dielectric constant of a mixture is, to a reasonable approximation, the weighted average of the component dielectric constants (this is the same mixing law as density). So for a quartz sand with oil and water in the pore space the dielectric constant is given by

$$E = \emptyset S_w E_w + \emptyset(1 - S_w)E_{oil} + (1 - \emptyset)E_{qtz} \tag{8.4}$$

where E_x is the dielectric constant of substance x and the suffixes w, oil and qtz correspond to water, oil and quartz. E_{oil} is typically about 2, E_{qtz} is 4.7 and as noted earlier E_w depends on temperature. If these are known, the equation can be re-arranged to find $\emptyset S_w$ and ultimately S_w.

As with the density porosity equation, finding the volume of water is equivalent to reading off a graph. The dielectric constant is exploited in very high frequency electromagnetic tools, which first appeared in the 1980s. These do not measure dielectric constant directly but measure the attenuation and phase shift of a pulse of microwave radiation. These in turn depend on the dielectric constant and conductivity of the medium.

As we will see in subsequent sections – dielectric constant has some advantages over electrical conductivity as a way of measuring water volume. Despite this, the original high frequency propagation tools were never very popular in practice and were hardly used at all after 1995. Probably their biggest disadvantage was that they were inherently shallow reading devices and consequently required an almost perfect borehole (the depth of investigation is considerably less than even the density tool). Even if the hole was good they generally only measured water saturation in the invaded zone.

A more fundamental problem is that the dielectric constant is actually a complex quantity (i.e., it has real and imaginary parts). Providing the medium is non-conductive the imaginary part is insignificant but formation waters invariably are electrical conductors because of the salts dissolved in them. This in effect means that water-filled porosity not only attenuates the field, it also phase shifts it. This is actually a plus, in that it means the tool response can be solved for both water volume and water salinity, but the analysis is complicated and the necessary computing power has only recently been available at the well site.

When the tools first appeared the raw data had to be re-processed in a computer centre so that the measurement could only be exploited sometime after the well was logged.

Recently, improved versions of the dielectric tools have appeared and consequently there is renewed interest in them. Modern well-site computers are easily powerful enough to carry out the analysis in real time and the tools use lower frequencies and can read deeper into the formation (although still shallow enough to be affected by invasion).

8.2.3 Neutron Lifetime

Another tool type, routinely run in cased holes for monitoring water movement, is the so-called neutron lifetime log. These actually respond to chlorine, but if the salinity of the water is known the amount of chlorine can be converted to water volume and hence saturation.

The tool works by firing pulses of very high-energy neutrons into the formation. As with any neutron tool these are slowed down by collisions with atoms in the borehole, formation and, in cased holes, the completion. Once they have reached low energies the neutrons are very susceptible to capture by chlorine. When this happens a gamma ray is produced which is detected by the tool. If there is a lot of chlorine present the neutron population is rapidly absorbed and the gamma activity rapidly decays. For low chlorine concentrations it takes longer for the neutrons to be absorbed and the gamma activity decays more slowly. The tool actually measures the decay rate of these secondary gamma rays and this is quantified as a property known as 'sigma' (Σ) which has dimensions of length^{-1}. Sigma is conventionally expressed in the prosaically named 'capture unit' or CU (1 CU = $0.1\,m^{-1}$). A high value of sigma denotes a rapid decay and a high concentration of chlorine. This could be due to a large volume of water with a moderate salinity or a smaller volume of water with a high salinity. The sigma value as a function of salinity is shown in Fig. 8.2. For freshwater sigma has been measured as 22.3 CU. Most minerals have sigma values of a few capture unit (halite is an obvious exception at 748 CU!).

It should be obvious that the technique works best with very saline waters. Conversely, the technique will not work for freshwater and in fact is not really suitable for formation waters fresher than marine compositions. Furthermore, these tools still have relatively shallow depths of investigation and so may be strongly affected by invasion (typically 25 cm). They are primarily used for monitoring saturation in producing wells, where the drilling fluids should have been flushed from the near well-bore region when the well was put on production.

8.2.4 Nuclear Magnetic Resonance

Nuclear magnetic resonance (NMR) quantifies the amount of hydrogen in the fluid state within the formation, as well as providing information on its chemical
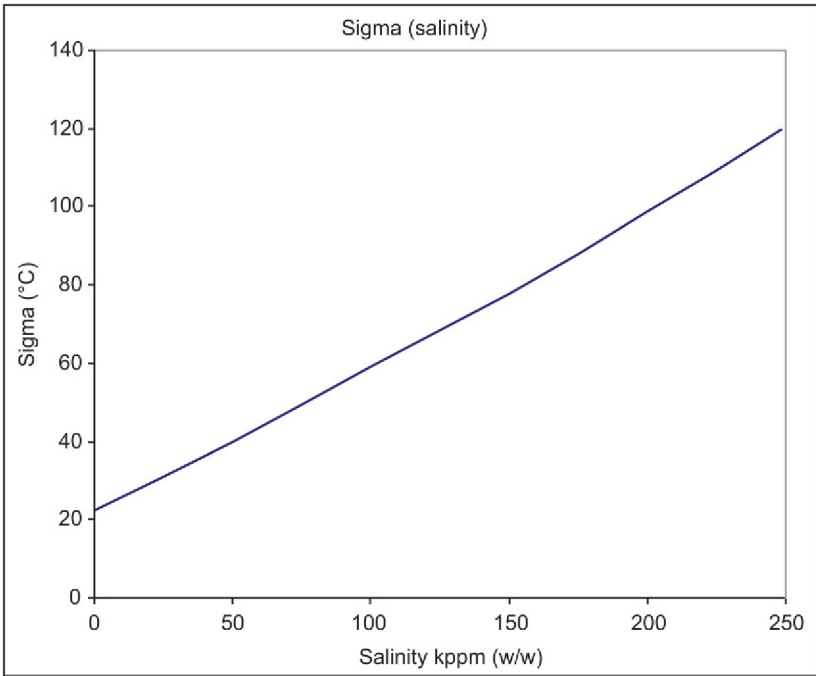
FIGURE 8.2 Sigma for sodium chloride solutions as a function of salinity.

environment. In the subsurface, hydrogen is found as either part of a clay layer – which NMR does not see – a water molecule, a hydrocarbon molecule or if unlucky part of an H_2S molecule. With gases – including H_2S – the concentration of hydrogen atoms is low and to a first approximation can be ignored (the concentration is quantified as the hydrogen index introduced in Section 1.4). So any hydrogen detected in a gas well can be related to a water volume. For oil wells however this is not the case as the hydrogen index of oils is similar – sometimes identical – to water. The basic NMR measurement is therefore not suitable for measuring water volume in the presence of oil.

The latest tools can however be set up to distinguish oil from water and so can be used as a saturation device in any situation (this exploits the fact the NMR measurement is affected, amongst other properties, by the diffusion coefficients of the molecules that contain the hydrogen). Unfortunately, these measurements are relatively slow and in any case NMR is still a shallow reading technique and therefore may only respond to the invaded zone.

8.2.5 Electrical Resistivity

Electrical resistivity is by far and away the most common way of estimating water volume and hence calculating saturation from logs. It is used so frequently that it is often forgotten that fundamentally it is still just a variation on Eq. 8.1.

But it is so important that it deserves a section to itself later. For now we note that its popularity stems from the fact that:

1. Resistivity tools have by far the greatest depths of investigation and so have the best chance of avoiding the invaded zone.
2. In most circumstances resistivity logs really do respond to water alone (inevitably there are exceptions to this which will be discussed later).
3. The parameters that are needed to compute saturation can be measured on core plugs.

8.3   WATER SATURATION FROM RESISTIVITY

8.3.1   Introduction

The most widely used tool for finding water volume from resistivity is the empirical equation, introduced at the beginning of the book, known as Archie's equation. Before discussing this we will briefly revise what resistivity is.

For any sample the resistance (R) is given by the ratio of the current to the voltage drop across the sample. In general resistance depends on the size and shape of the sample. If the sample is homogeneous and has a simple geometry, the resistance can be simply expressed in terms of an intensive property known as resistivity and the dimensions of the sample. For a cylindrical sample in which the voltage drop is parallel to the axis

$$R = \frac{\rho l}{A} \qquad (8.5)$$

where ρ is the resistivity, A is the cross-sectional area and l is the length the current flows along. This is the first and last time we use the physicist's and electrical engineer's convention that R is resistance and ρ is the resistivity.

Here, we are only interested in resistivity because it is a property of the medium and is independent of the shape of the sample (it is an 'intensive property'). This means measurements from different types of logging tools and from core samples are directly comparable. Because we never deal with resistance we can get away with using the symbol 'R' for resistivity without danger of confusion (just as well as we are already using ρ for density). The SI unit for resistivity is the Ωm – 'Ohm-metre' – and for once, the oil industry does use the SI unit (unlike electrical engineers, who tend to use Ωcm).

8.3.2   Archie's Equation

Archie's work on the resistivity of water saturated porous rocks was described in Chapter 1. The experiments basically involved measuring the resistivity of a large number of core plugs that were saturated with water of known composition. The porosity and permeability of the plugs were known from conventional core analysis and so Archie was able to cross-plot resistivity against these.

Qualitatively Archie's work suggested:

1. The resistivity of a saturated plug (R_0) is proportional to the resistivity of the water (R_w) in it.
2. The resistivity typically falls with increasing porosity.

As already noted in Section 1.5.1 both of these observations are what one would expect. Consider first the effect of the water composition. For all intents and purposes, the minerals that make up clastic and carbonate rocks are nearly all electrical insulators and so the current in a porous rock is carried entirely by the water. If water with a resistivity of $1\,\Omega m$ is replaced by water with a resistivity of $0.1\,\Omega m$, one would expect the resistivity of the rock to fall by a factor of 10. This is what was observed most of the time (in fact Archie used waters with salinities between 20,000 and 100,000 ppm so would never get a 10fold change in water resistivity).

This allowed Archie to define a dimensionless property known as 'formation resistivity factor' or simply 'formation factor.' This is variously given the symbol *F*, *FF* or *FRF*. Providing condition 1 holds, the formation factor is a property of the rock and is independent of the water resistivity (which depends on composition and temperature). The resistivity of the water saturated rock R_0 is then given by:

$$R_0 = FR_w \qquad (8.6)$$

In other words, *F* is the resistivity of a rock saturated with $1\,\Omega m$ water.

Now consider the effect of the porosity. An increase in porosity is equivalent to an increase in the volume of water and so more current can be carried for a particular voltage drop. In short the resistivity falls. This can be exploited to calculate porosity from resistivity.

One of Archie's original figures showing some of his results is reproduced in Fig. 8.3. In fact there are two separate graphs on the same figure, which was how he presented resistivity data in his early papers. On the left resistivity is plotted against permeability and on the right, taking up less than 20% of the figure, is resistivity against porosity. It is the latter that most people would regard as the more significant plot. (In fact Archie has plotted porosity and permeability against the 'formation factor', but as explained earlier this is closely related to resistivity and to begin with we will discuss the plots as if they plot resistivity against porosity.)

The particular plot shows measurements for approximately 25 plugs, covering a porosity range from 12% to about 25%. Note that both resistivity and porosity are plotted on logarithmic grids, which means that the straight line fit can be described by:

$$\mathrm{Log}(R_0) = A - m\mathrm{Log}(\varnothing) \qquad (8.6a)$$

Or, re-writing:

$$R_0 = \frac{aR_w}{\varnothing^m} \qquad (8.6b)$$

FIG. 4.—Relation of porosity and permeability to formation resistivity
factor, Eocene Wilcox sandstone, Mercry, Texas.
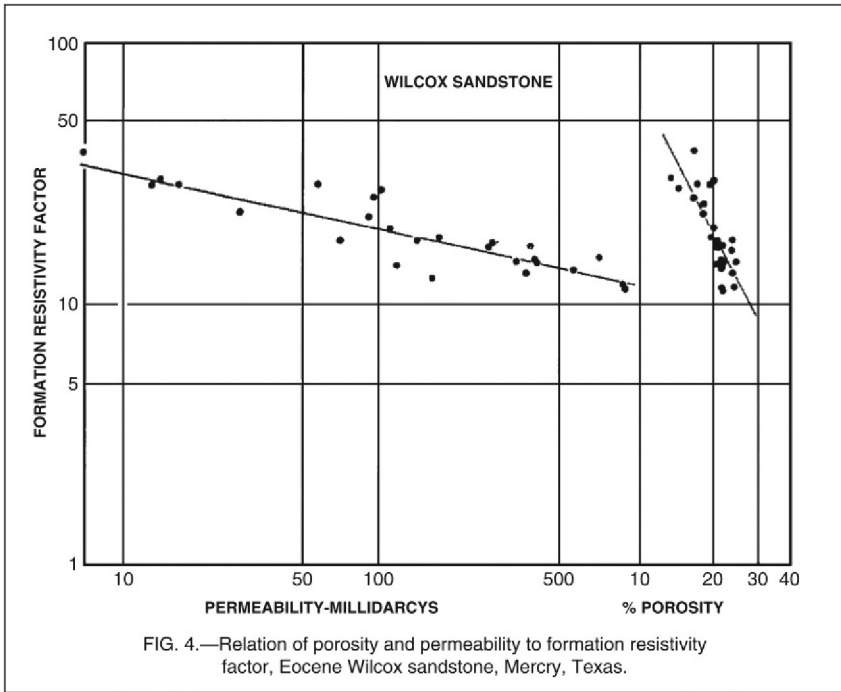
FIGURE 8.3 One of Archie's original figures. There are in fact two plots: permeability against resistivity on the left and porosity against resistivity on the right. The relationship with porosity is the one that concerns us here.

where R_0 is the resistivity of the water saturated plug, R_w the resistivity of the water, $\emptyset$ the porosity and a and m are constants. This is Archie's equation (although the 'a' value was actually added by researchers at Humble Oil about 10 years after Archie first published). The porosity has to be input as a fraction and providing $m > 1$ it can be seen that it is completely consistent with points 1 and 2 discussed earlier. By re-arranging the equation it is possible to estimate porosity from resistivity providing the pore space is saturated with water. The constants a and m can be obtained from the left hand plot. The constant m is sometimes called the 'cementation exponent' and it is the gradient of the line on a log–log plot (it is also often referred to as the 'Archie-m'). Using Archie's data in the above plot, gives an m value of 2, which is typical for clastics.

The form of Eq. 8.6b is illustrated in Fig. 8.4. With porosity and resistivity plotted on linear scales it can be seen that resistivity increases steeply, with falling porosity at low porosity (so as a porosity tool it will be most accurate at low porosity). The same calculation is shown plotted on a log–log grid in Fig. 8.5 this is the normal way of presenting the data because it produces a straight line. Either plot confirms that for a particular porosity, resistivity increases with increasing R_w and/or m.
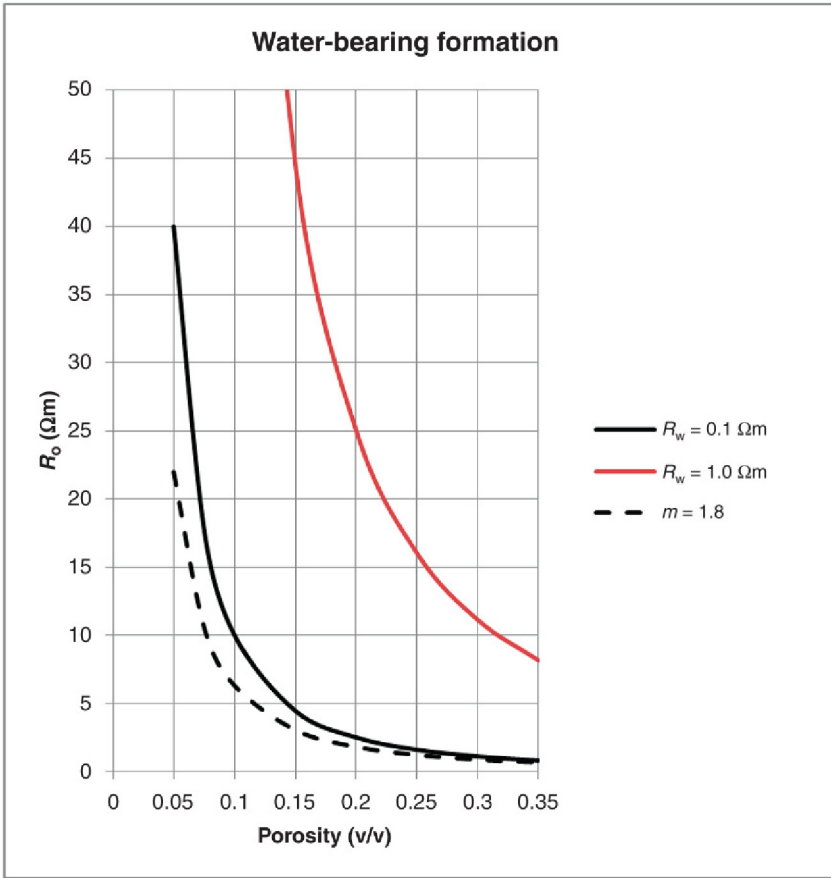
FIGURE 8.4   Resistivity as a function of porosity for a water-bearing formation, plotted on a linear grid. Resistivities have been calculated using Archie's equation with two values of R_w and two values of m. (m is set to 2 for the different values of R_w and R_w is set to 0.1 Ωm for the two values of m.)

There are numerous extra symbols and properties that have been introduced to help or hinder the interpretation of resistivity. We have already met the 'formation factor', which is the ratio of formation resistivity to the resistivity of the fluid. In terms of porosity it is given by:

$$F = \frac{R_0}{R_w} = \frac{a}{\varnothing^m} \qquad (8.7)$$

In Archie's data F varied from 'a few' to about a hundred. A high value indicates a low porosity and/or a high m value.
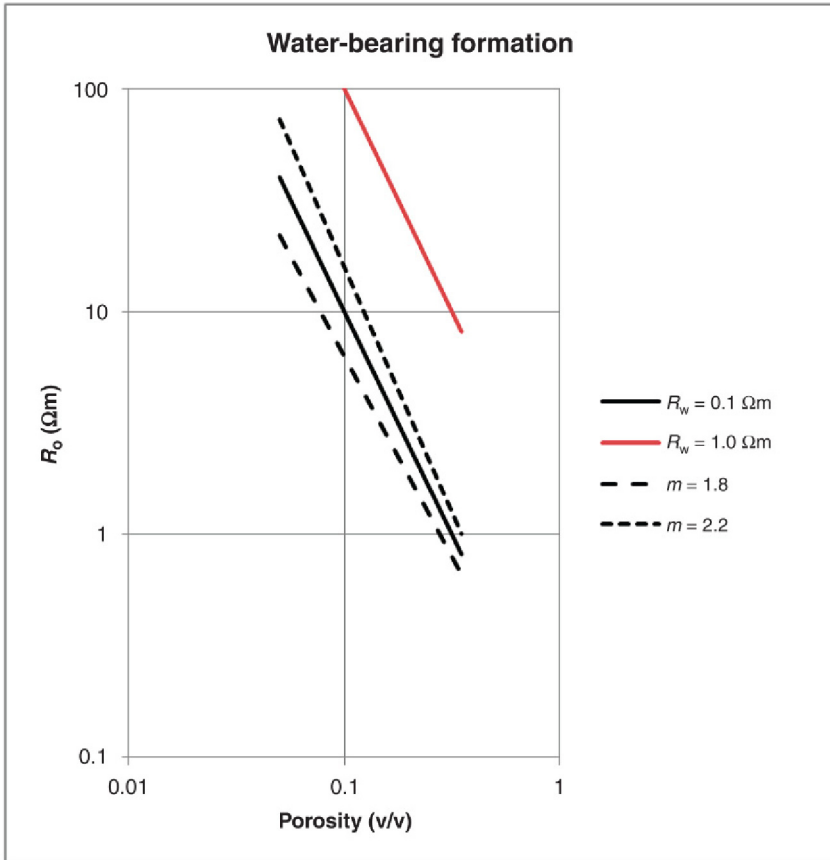
FIGURE 8.5 Resistivity as a function of porosity for a water-bearing formation, plotted on a log–log grid. Resistivities have been calculated using Archie's equation with two values of R_w and three values of m. (m is set to 2 for the different values of R_w and R_w is set to $0.1\,\Omega$m for the three values of m.)

8.3.3 Water Saturation and the Archie Equation

So far we have only discussed water-saturated rocks as Archie's experiments involved these. But even in his original 1942 paper Archie had suggested incorporating some even earlier studies that looked at the change in resistivity as a result of reducing the water saturation. At least four papers published in the late 1930s had all shown that to a good approximation:

$$\frac{R_0}{R_t} = S_w^{-n} \tag{8.8}$$

where R_t is the resistivity of the plug at a water saturation of S_w. The value of the constant n is approximately equal to 2 for the sandstones and sand packs that

were used in the experiments. Curiously, none of these studies were concerned primarily with resistivity. They were actually looking at two-phase permeability and the researchers needed a way of measuring water saturation *in situ*. Resistivity was an obvious quantity and Eq. 8.8 was simply the calibration curve used to convert resistivity to saturation.

Archie combined Eq. 8.8 with Eq. 8.6b to give:

$$R_t = \frac{aR_w}{(\emptyset^m S_w^n)} \qquad (8.9)$$

This is also called 'Archie's equation' and Eq. 8.6b is then just a special case.

The effect of the S_w term is to increase R_t (as with porosity, S_w has to be given as a fraction). This makes sense as the water volume has been further reduced. Figure 8.6 shows resistivity as a function of porosity for three different water saturations. Notice that the symbol for formation resistivity is now R_t – the true formation resistivity.

The ratio of R_t to R_0 is called the 'resistivity index' (RI) and the constant n is known as the 'saturation exponent' (or often just the 'Archie-n'). Summarising:

$$RI = \frac{R_t}{R_0} = \frac{1}{S_w^n} \qquad (8.10)$$
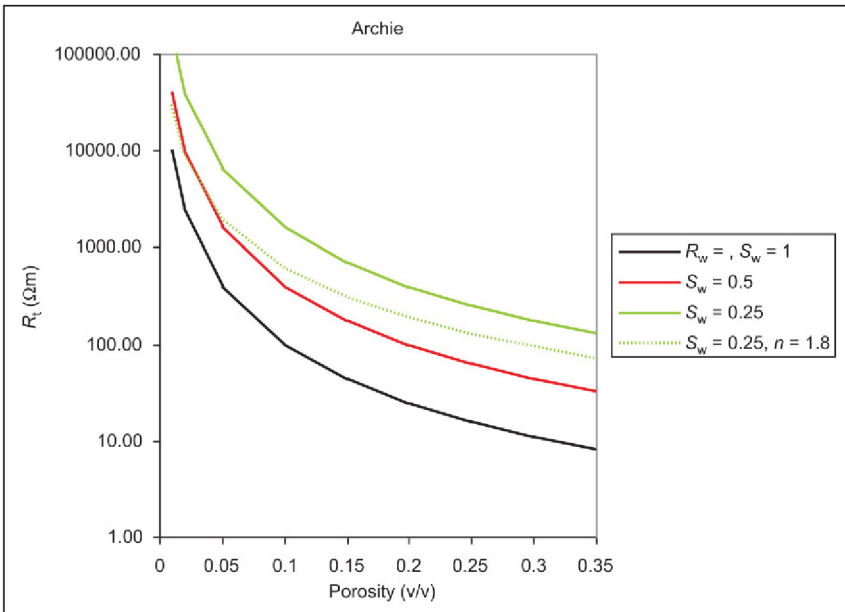
FIGURE 8.6   Resistivity as a function of porosity for various water saturations between 0.25 and 1.0. R_w is set to $1.00\,\Omega m$ and $m=2.0$ throughout. The dashed line shows the effect of changing the saturation exponent from 2.0 to 1.8.

Or in logarithms:

$$Log(RI) = -nLog(S_w) \tag{8.11}$$

As already mentioned, n tends to take values that are close to 2 for clastics, in other words similar to m. At least that is the case for water-wet rocks, it can have much higher values in oil-wet rocks (wettability will be discussed in a later chapter). For any combination of porosity and water saturation, reducing n results in a reduction in resistivity. This is illustrated in Fig. 8.6 and particularly in Fig. 8.7 where the effect of reducing n from 2.2 to 1.8 is shown. So when using the Archie equation to calculate saturation, a lower n value results in a lower water saturation for a particular porosity and resistivity.
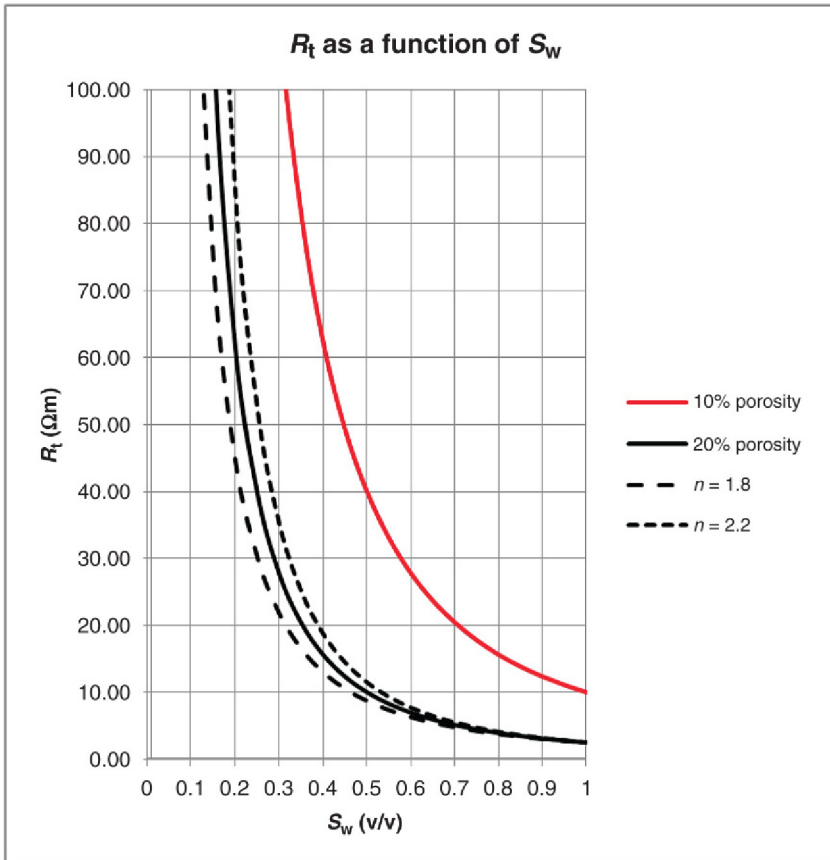
FIGURE 8.7   Resistivity as a function of water saturation for two porosities and three values of n; n is held constant at 2 for the two porosity cases and porosity is constant at 20% for the three n values.

As a large number of symbols and properties has been introduced in quick succession it is worth summarising them in one place:

R_0 The resistivity of water saturated rock.
R_{x0} The resistivity of the near well-bore region.
R_t The formation resistivity (for water-saturated rock this is equal to R_0).
R_w The formation water resistivity.
R_{mf} The mud filtrate resistivity.
F Formation factor. The ratio of the resistivity of the rock to the resistivity of the fluid in the pore space.
R_I Resistivity index. The ratio of the actual resistivity of the rock to the resistivity if it were saturated with water.
S_w The true water saturation of the rock.
S_{x0} The water saturation in the near well-bore region.

8.3.4   Calculating Saturation and Saturation Parameters

To calculate saturation the Archie equation is re-arranged to give:

$$S_w = \left| \frac{aR_w}{R_t} \frac{1}{\emptyset^m} \right|^{1/n} \tag{8.12}$$

if $a=1$ and $n=m=2$ this takes the simpler, more convenient form:

$$S_w = \frac{\sqrt{(R_w / R_t)}}{\emptyset} \tag{8.12b}$$

By replacing the subscript 'w' with 'x0' we can write equivalent equations for the near well-bore zone (where invasion is likely to change the saturation). As a result of invasion, the formation water will be replaced by mud-filtrate, so that R_w is replaced by R_{mf} when calculating invaded zone saturations.

In order to use Eq. 8.12 one needs at least three – generally four – parameters (a, m, n and R_w). As with any other equation in log analysis these introduce some uncertainty. In this section we will look at some of the ways they can be estimated.

8.3.4.1   Formation Water and Water Resistivity R_w

R_w is the only saturation parameter that is a property of the fluid. Pure water is actually a poor conductor with a resistivity of approximately 250,000 Ωm at 25°C. No commercial logging tool is capable of reading such a high value even if it was simply immersed in a swimming pool. But dissolving as little as a few parts per million of a salt in it is sufficient to reduce the resistivity to a measurable value. The variation of resistivity with concentration for sodium chloride solutions is shown in Fig. 8.8. It can be seen that resistivity falls by roughly
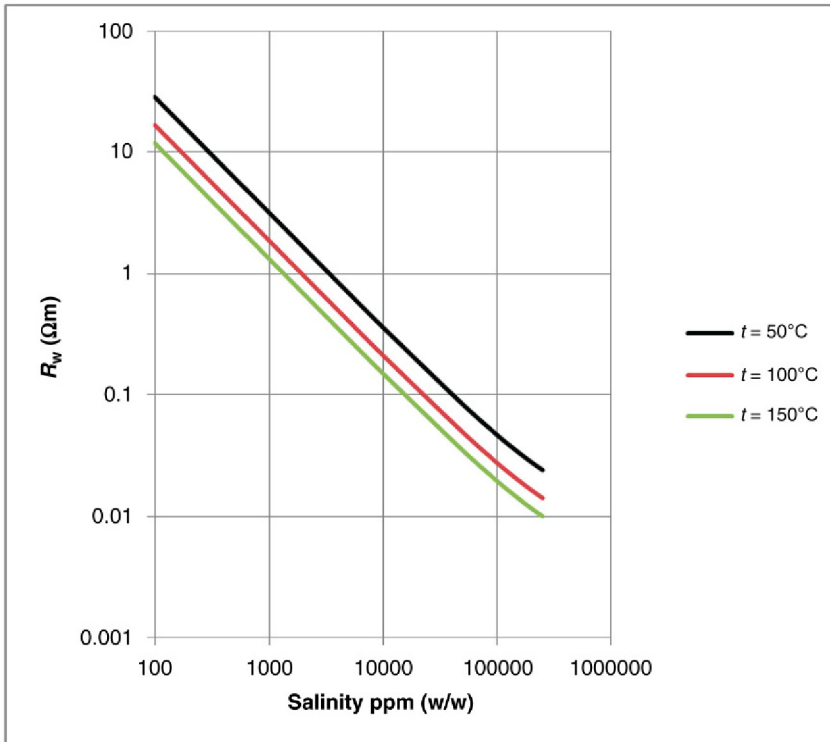
FIGURE 8.8   Water resistivity as a function of salinity for three different temperatures.

a factor of 10 for every 10-fold increase in salinity. Figure 8.8 only refers to solutions of sodium chloride in water but in fact, for our purposes, that is what most formation waters are. Although other salts are present, they are normally a very minor component. Exceptions are sometimes found in carbonate reservoirs where large amounts of calcium chloride may also be present (calcium chloride not only affects R_w it can also result in formation waters with densities that are much higher than $1 \, g/cm^3$).

Figure 8.8 also shows that resistivity falls with increasing temperature. Chart books include a chart or nomogram to estimate R_w given temperature and salinity (log analysis software should also include a module that does the same thing). A commonly used equation to find R_w at a temperature t_1 given its value at a reference temperature t_0 is:

$$R_w(t_1) = R_w(t_0) \frac{(t_0 + 21.5)}{(t_1 + 21.5)} \tag{8.13}$$

where t_1, t_0 are given in degrees Celsius.

There are a large number of ways of finding R_w. These vary from laboratory measurements on samples of formation water to educated guesswork. Some of these are very accurate and reliable and others need to be treated with caution. Broadly speaking they fall into three categories:

1. Laboratory measurements on samples.
2. Log-based techniques.
3. Regional and geological knowledge.

Laboratory measurements are made on samples of water recovered from the well. The main cause for concern is the quality of the samples rather than the quality of the measurements, which for the major ions are well established and very accurate. Most laboratories will measure the concentrations of all the commonly occurring ions and also measure the water resistivity directly.

Samples come from a variety of sources including well tests, wireline formation testers and water recovered from cores. Although it is unlikely that a well test would be performed with the main aim of getting a water sample, this is still the best way of getting a reliable sample. Tests can produce thousands of litres of water so that any mud filtrate is flushed clear. More likely water samples will be taken using a wireline formation tester. Since the 1990s these have been capable of:

1. Pumping fluid into the borehole until the engineer decides to divert the flow to the sample chamber(s).
2. Continuously measuring resistivity and other properties as the sample flows through the tool.
3. Taking 12 or more samples.

So although it may take a long time for the tool to pump out all the mud filtrate, the engineer can be reasonably sure that pure formation water is flowing into the sample chambers. Older tools had neither pump-out nor down-hole measurement capability and could take at most two samples. Sampling was then a matter of opening the chamber and hoping for the best!

Water will only flow from an aquifer or the transition zone, so samples from tests or wireline formation testers can only come from these parts of the reservoir. If samples are used to provide R_w values it is tacitly assumed that the water does not change in composition between the aquifer and the hydrocarbon legs. As will be seen later this need not be the case.

The smallest and arguably least reliable samples are obtained from core. Various methodologies have been tried, but the two extremes are to extract fluids from the centre of the core, where hopefully the drilling fluids have not penetrated, or to drain all the free fluid from the core plugs and then 'spin' the tightly bound water out of the plugs using a centrifuge. The samples are always tiny but this is the only way to obtain samples of formation water from high within an oil or gas column. In fact arguably the best opportunity to get reliable water samples from core plugs is from cores that have been cut in the

hydrocarbon leg using oil-based mud as then the only water in the core should be formation water.

If a water sample is not available, methods based on log responses have to be used. Even when samples are available, log-based methods add a lot of value by, for example acting as an independent check or revealing changes in water composition within a well. There are three basic log-based techniques two of which will only work in wells drilled with water-based mud.

As with samples, most of these techniques rely on a water-bearing formation being available and assume that the water in the aquifer has the same composition as the hydrocarbon leg. Probably the commonest log-based technique is the Pickett plot, which is very similar to the type of plot that Archie published. In a Pickett Plot resistivity, rather than formation factor, is plotted against porosity on a log–log grid. In principle, such a plot will give the m value from the gradient of the straight line and R_w from the intercept with the 100% porosity axis. More often than not, however there is simply not enough variation in porosity to find both properties. Rather than forming a neat straight line the data points then form a tight 'blob'. To find R_w a line with gradient m is forced through the 'blob' and continued to 100% porosity. The R_w value obtained is at reservoir conditions, so no temperature corrections are needed.

The main problem with the Pickett Plot is the value 'a' in Archie's equation because the intercept is actually the product aR_w. If a is known, or assumed to be close to unity, the technique can work very well.

In practice, resistivity readings from a deep-reading logging tool are cross-plotted against porosity computed in a water-bearing interval. The best straight line is forced through the points and it's intercept with the resistivity axis is read-off. This is R_w or possibly aR_w. Figure 8.9 shows an example where there
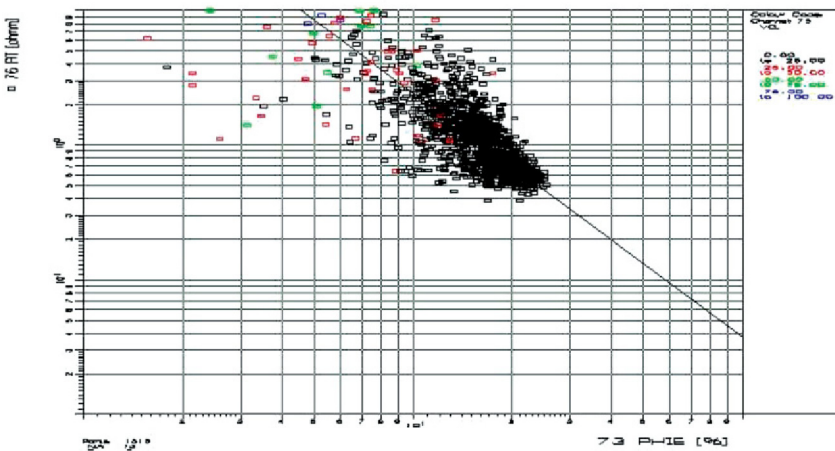
FIGURE 8.9   An example of a Pickett plot. In this case $R_w = 0.037\,\Omega$m at reservoir conditions. In this case the value of m can also be estimated to be (1.83).

is a large enough variation in porosity that m and R_w can both be estimated. In that case, which comes from a North Sea chalk reservoir, porosity varies from about 5% to over 25%. There is a lot of scatter in the data, which leads to some uncertainty in both m and R_w. Although R_w has been given as $0.037\,\Omega\text{m}$ on the plot it could be as high as 0.05 or as low as $0.025\,\Omega\text{m}$ (similarly m could lie between 1.7 and 2 and still fit the data.)

It is worth recalling some of the points made in Chapter 1 at this point. It is tempting to dismiss the scatter in Fig. 8.9 as 'experimental error'. Examples might be measurement uncertainty or small errors in the porosity calculation. But most of the scatter is probably real and reminds us that Archie's equation is just the best empirical fit to the data. In reality every point in the reservoir requires a unique value of m (and possibly a) to make the equation work. At the end of the day, we seek methods and parameters that produce estimates that are not too far from reality and over-estimate as often as they under-estimate.

A related plot to the Pickett plot is the 'Hingle plot'. These were developed before calculators and computers were in widespread use and were primarily intended for well-site work. They are really obsolete now but the blank grids can still be found towards the back of most chart books. These charts plot porosity on a linear scale against $(aR_t)^{1/m}$. The interested reader can see this is just another way of getting a straight line out of the Archie equation. Unlike the Pickett plot, a different chart is needed for each combination of a and m and normally three alternatives are given in a typical chart book ($a=1$, $m=2$, $a=1$, $m=1.8$ and $a=0.62$, $m=2.15$). If porosity is on a linear scale there is actually no need to convert density, neutron porosity or, if the Wyllie equation applied, sonic readings to porosity. In other words the log analyst could simply plot log readings straight onto the grid without needing to calculate porosity. This was a major plus in the days before calculators but now it is almost as easy to calculate porosity and plot that on a Pickett plot.

Obviously, Pickett plots work equally well for water or oil-based muds. The other two log-based methods can only be used in water-based muds.

The SP can be used to find R_w because the deflection depends on the contrast between formation water and mud-filtrate salinities. As the latter is known, the formation water salinity and hence resistivity can be found. In practice this involves applying several charts. The SP can obviously only be used in water-based muds and needs a sand of at least 5-m thick to develop its maximum deflection. In principle the SP should allow R_w to be found in oil- and gas-bearing formations but there is a lot of evidence to suggest the deflection can be reduced in the presence of hydrocarbons. This is particularly true at high hydrocarbon saturations. Furthermore the sand must be clean to develop the full deflection.

If a micro-resistivity device is available and deep invasion is known to have occurred then the following relation holds true:

$$\frac{R_{x0}}{R_t} = \frac{R_{mf}}{R_w} \qquad\qquad (8.14)$$
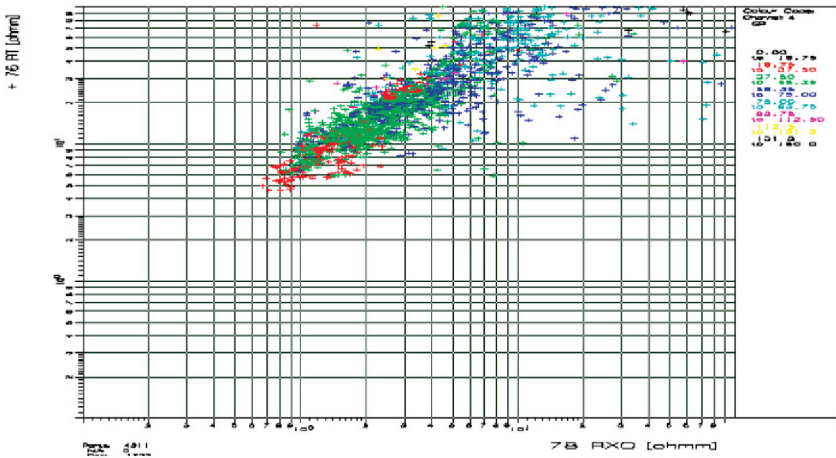
FIGURE 8.10   An example of a R_t/R_{x0} cross-plot. In this case the deep resistivity reads 6–7 times the micro-resistivity reading. This implies R_w is 6–7 times the mud filtrate resistivity ($R_{mf} = 0.07\,\Omega m$ at 32°C). Note the gradient is 1:1.

which can be re-arranged to give R_w in terms of R_{mf} (at reservoir temperature). An example of how it is used in practice is shown in Fig. 8.10 where the deep resistivity (R_0) has been cross-plotted against the micro-resistivity (R_{x0}). In that case when R_{x0} is $1\,\Omega m$ the deep resistivity tool reads $7\,\Omega m$, in other words $R_0 = 7R_{x0}$ or $R_w = 7R_{mf}$.

The resistivities are normally plotted on a log–log grid although in principle linear scales would be just as good. It is important to note however, that the gradient has to be 1:1. If it is any different to that, Eq. 8.14 does not fit the data and the method should not be used.

It was implied earlier that the R_t/R_{x0} method will not work in oil-based muds. This is simply because all micro-resistivity tools need a conductive mud to function. But two relatively modern tools designed for other applications, the oil-based image log and the latest dielectric tool, can in fact provide a near well-bore resistivity measurement in oil-based mud.

Other ways of finding R_w include local knowledge, which for some basins has been made available in the form of catalogues of values by formation and area. It should also be noted that R_w may be constrained by geological process-es. Reservoirs that are capped by salt, for example will have formation waters that are salt saturated (e.g. the Rotleigendes sands of the Southern Gas Basin).

If possible, it is best to use more than one method to find R_w (even if a sam-ple is available).

It should be noted that no matter how the formation water resistivity is de-termined, compositions can vary within reservoirs. In particular it is possible for the water composition in the aquifer to be different to that in the oil leg! A relatively common situation is to have the original formation water preserved

in the hydrocarbon column but to be replaced in the aquifer. This is found in many basins and is particularly common in onshore or near-shore settings. In these cases meteoric water flowing down dip displaces the original water from the aquifer. The incoming water bypasses the structural highs that trap the hydrocarbons however, leaving the original water in place there.

Near salt diapers the reverse process can occur. Water flowing past the diaper dissolves salt and becomes a lot more saline than the original water retained in hydrocarbon columns.

8.3.4.2 Cementation Exponent (Archie 'm') and 'a'

The cementation exponent is found by cross-plotting the resistivity of water-saturated rock against porosity. This is another example of the use of the Pickett plot but this time it is the gradient that is of interest rather than the intercept. The data may either come from logs or from measurements on core plugs. In Fig. 8.9 log data was plotted and consequently there are a lot of points. Core-based measurements are considered part of special core analysis (SCAL) but they are nevertheless a standard laboratory service. Normally, the laboratory estimate will be based on relatively few plugs, but it does have the advantage that conditions can be tightly controlled, in particular R_w is known (typically about 10 separate measurements are made).

In order to estimate 'm' we ideally need a good range of porosities, this is certainly the case in Fig. 8.9 where there is fivefold variation in porosity. Unfortunately, as discussed earlier, porosity often only covers a limited range. In these circumstances 'm' can still be estimated providing R_w is known. This effectively provides a point at 100% porosity. (This is the reverse of the previous section where we had to assume m in order to find R_w.)

With core measurements R_w is certainly known so this is not a limitation. Nevertheless plugs are selected to give as wide a range of porosity as possible. A plot for some typical core data from a single well is shown later.

Core reports normally present plots of formation factor rather than resistivity against porosity. They also normally give several alternative interpretations of the raw data.

1. The data is forced through the point FRF=1.00, porosity=100%. This is equivalent to fixing $a=1$.
2. The data is fitted with a free regression. The best-fit line will then generally pass through a value different from 1.00 at 100% porosity. This is equivalent to having an 'a' value different from 1.
3. Individual m values are calculated for each plug assuming $a=1$. (Each line is constructed using only two points.)

This actually shows that a and m are not independent. In the case of the data shown in Fig. 8.11 a free regression has been used in which a turns out to be less than one. The data could also have been fitted with a lower m value and an a value of unity. The pairs will give slightly different saturations of course.
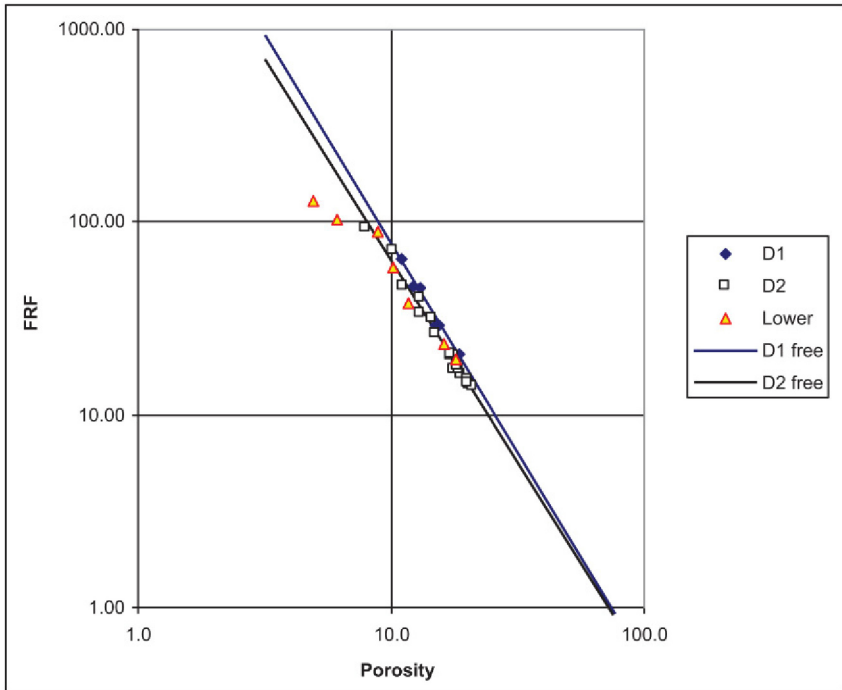
FIGURE 8.11   Special core analysis determination of cementation exponent '*m*'. The resistivity is presented as formation factor (FRF). The data has been split into reservoir units D1, D2 and lower so that individual *m* values can be defined for each unit. The lines are free regressions for D1 and D2 and do not pass through FRF=1 at 100% porosity. This is equivalent to having the '*a*' parameter different from 1 (about 0.5 in this case).

8.3.4.3   Saturation Exponent (Archie 'n')

The saturation exponent can realistically only be obtained from measurements on core plugs. The experiment consists of saturating a plug with brine, measuring its resistivity and then draining it in stages and re-measuring resistivity at each stage. A plot of resistivity against water saturation on a log–log grid then gives '*n*' from the gradient of the best straight line through the data (Eq. 8.10b). Since the plug starts out saturated with water the line has to pass through the value of R_0 at 100% water (contrast this with the derivation of *m* described earlier).

In practice core laboratories normally plot resistivity index (RI) against water saturation, this simply means the resistivity has been normalised to the value at 100% water. Notice that one value of '*n*' is measured for each plug in contrast to '*m*' which is based on measurements from several plugs. Sometimes core laboratories present all the individual measurements, from all plugs on the same plot and then put the best straight line through the complete data set. This is simply a way of averaging all the individual plug measurements.

In principal 'n' could be derived from logs alone providing one has an independent measure of S_w that does not depend on resistivity. Any of the logs described in Section 8.2 could do this, but they are all shallow reading tools that are more likely to give invaded zone saturations. This could still be combined with a measurement of invaded zone resistivity (R_{x0}) but in practice all the tools have different depths of investigation and the resulting value of 'n' would have a large error bar. If there is no core available it is probably better to assume $n = 2$.

8.4 BACK TO THE ROCKS. WHAT CONTROLS THE SATURATION PARAMETERS?

The Archie equation is so familiar and so commonly used that it can be easy to forget that it is purely empirical and each formation or rock type requires a different set of parameters to give a good fit. Even in the original paper Archie had pointed out that the constant m seemed to be related to the degree of cementation with well-consolidated sands having values in the range '1.8–2.0' and artificial sand-packs having lower values of 1.3. He speculated that the unconsolidated sands typical of the Gulf Coast would have values between 1.3 and 1.8. In a later paper he produced his characteristic two part graphs for four different formations – two sandstones and two limestones – to again show that each one was characterised by a different m value (one of these figures is shown earlier). These observations have stood the test of time and indeed it is found that the value of m is dependent to a large extent on various aspects of the pore system.

8.4.1 A Simple Model for 'm'

Earlier, we stated that it is helpful to consider highly idealised systems when trying to explain what controls a petrophysical property. In the case of m, it is ultimately resistivity that we are trying to understand and to make progress we need to calculate it for a simple porosity model. The easiest place to start is the capillary tube model that was introduced in Chapter 2 to help understand permeability. An example has been reproduced in Fig. 8.12. Because water is assumed to carry all the current each parallel-sided pore is an example of the simple resistor discussed in Section 8.3.1 The resistivity of the pores is R_w and their total resistance is given by:

$$\text{Resistance} = R_w \frac{l}{\mathbf{a}} \tag{8.15}$$

where $\mathbf{a}$ is the cross-sectional area of the pores, l is the length of the sample and we have written 'resistance' to avoid any confusion with resistivity. The resistivity of the sample is given by:

$$R_0 = \frac{\text{Resistance} \cdot \mathbf{A}}{l} \tag{8.16}$$
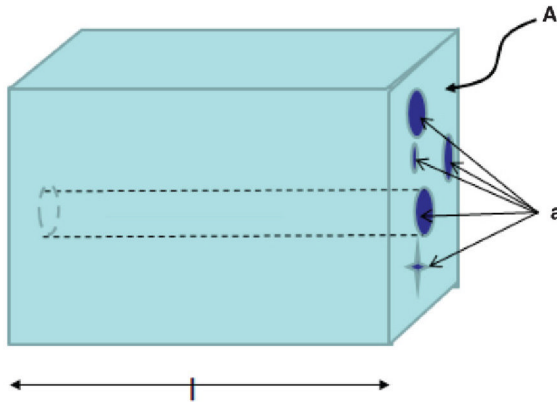
FIGURE 8.12   Capillary pore model for formation factor. The sample has cross-sectional area and length *l*. The porosity is represented by parallel-sided capillaries with a total cross-sectional area of *a*, that run parallel to the axis of the sample.

where **A** is the cross-sectional area of the sample (e.g. a plug).

$$= R_w \frac{\mathbf{A}}{\mathbf{a}} = \frac{R_w}{\emptyset} \tag{8.17}$$

where we have used the fact that the ratio of the cross-sectional area of the pores to the cross-sectional area of the sample – **a/A** – is in this case the porosity. So for this type of system the *m* value has been found from first principles to be 1.00. The model can be made more realistic by

1. putting restrictions in the capillary tubes
2. putting bends and kinks in the capillary tubes so that they are longer than the sample length *l*
3. blocking off some capillaries completely.

Any or all of these will increase R_0 above the value given by Eq. 8.17. Numerically, an increase in R_0 can be achieved by having an *m* that is greater than one. As the restrictions get smaller or more numerous and as more dead-ends enter the pore system *m* has to be increased to correctly predict resistivity.

In a real clastic rock these changes could be caused by poor sorting, cements or clays forming in the pore throats or some combination. So there is a good basis for Archie's assertion that as the degree of consolidation increases, *m* rises. The opposite change could occur if a clastic rock is fractured, as then a type of porosity with few restrictions is introduced. The most extreme *m* values are expected for vuggy systems where much of the porosity is located in isolated packages with only limited connectivity to the outside world. In these systems much of the water is inaccessible and R_0 is expected to be high. Vuggy carbonates for example often have *m* values of 2.5 or more.

Note that increasing the length of the capillary tubes by a factor '*a*' allows the resistivity to be calculated exactly as:

$$R_0 = a \frac{R_w}{\emptyset} \qquad (8.16b)$$

(Note this *a* is not the same as bold **a** which referred to the pore area, although confusing we have used '*a*' to keep it consistent with the Archie equation.)

This is the Archie equation, with an *m* of one, after it had been modified by the Humble Oil group. This sort of reasoning explains why the constant *a* is sometimes known as 'tortuosity constant'. Unfortunately, this model predicts a must always be greater than one but in fact with real data it is quite often less than one (yet another case of pushing a simple model too far).

8.4.2   The *m* Value for Real Rocks

So there is good reason to think that *m* (and possibly *a*) is related to the structure of the pore system with more open pore systems leading to lower *m* values than those involving fewer more tortuous paths through the rock. In practice these features are related to petrographic properties such as sorting. Petrographers describe sedimentary rocks using a large number of parameters but most of these are at best semi-quantitative and they are generally quite interpretive as well. So calculating the *m* value to two decimal places from a description or a table of petrographic parameters is hopelessly ambitious. Such information can, however, certainly help narrow the likely range that is expected and conversely the approximate *m* value can tell us something about the nature of the rock.

Figure 8.13 shows thin-section micrographs of three different sandstones, together with a highly abridged petrographic description, porosity, permeability and the *m* value of the plug (so *a* is fixed at 1). On the left (sample *a*) is an unconsolidated sand, the pores are all highly connected and there are numerous pathways through the rock with only minor restrictions. The *m* value of 1.66 is the lowest of the three and most log analysts would regard it as at the lower end of possible values for *a* sandstone. At the other extreme on the right (sample *c*) is a well-consolidated sandstone in which the porosity appears to be concentrated in a relatively few, large pores that are connected together by narrow 'slots'. The *m* value is at the high end of values expected for clastics (*m* = 2.14). Between these extremes with an *m* value of close to 2 (*m* = 1.99) is a a very fine-grained sand with a high clay content. Although it has a lot of clay in it, this sand still has a more open pore system than the right-hand sample and consequently *m* is lower.

Obviously, many more than three specimens are needed to come up with any relationships between fabric and cementation exponent. Specific studies using scores of samples as well as everyday experience have identified parameters that seem to exert most control on *m*. As we have seen, anything that leads to an open pore system with numerous connections between pores and only minor

(a)

Unconsolidated.
Lower medium grained,
Moderate sorted
Porosity = 30%, kh = 8000 mD
$m = 1.66$

(b)

Shaly.
Very fine grained,
Well sorted
Porosity = 26%, kh = 18 mD
$m = 1.99$

(c)

Well consolidated.
Bimodal sorting
Porosity = 12%, kh = 12 mD
$m = 2.14$
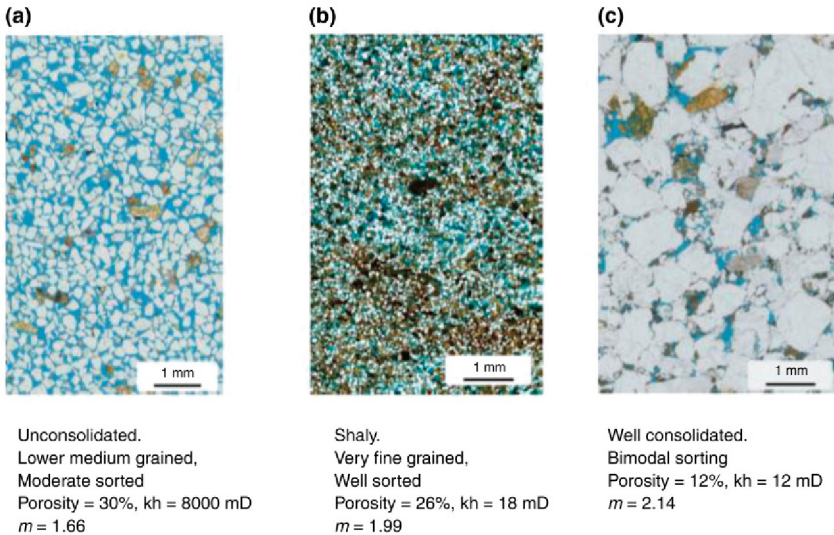
FIGURE 8.13 Thin section photographs of three different sandstones showing how m is related to the nature of the pore system (see text for explanation).

constrictions gives lower m values. In practice that means unconsolidated sands with reasonably uniform grain sizes. At the other extreme, heavily cemented sands, possibly together with poor sorting, where the pores are connected by narrow restrictions, are characterised by high m values. Of all the different factors to consider, degree of cementation is probably the single best predictor of m value (which explains why m is called the 'cementation factor'). A very rough rule of thumb is:

Heavily cemented	2.0–2.2
Moderately cemented	1.8–1.9
Slightly cemented	1.6–1.7
Unconsolidated	1.3–1.5
Beach sand	1.2–1.3

Carbonates can give some of the highest m values because they are prone to developing vugs and moulds where most of the porosity is concentrated in relatively large, poorly connected clumps. But carbonates can also produce some of the 'best behaved' porosity systems with m values a bit lower than two. Examples are chalks, oolitic limestones and sucrosic dolomites. Although they are formed in very different environments these all have in common, very uniform grain sizes and shapes. In the case of chalks it is the skeletal remains of a single species of foram that form the grains. Ooids are spherical particles of remarkably uniform size that form in high energy, shallow water environments and the grains forming sucrosic dolomite are single crystals with the rhombohedral shape characteristic of pure crystalline dolomite (often simply known as 'rhombs').

8.4.3 Relationship of *m* to Porosity and Permeability

As resistivity and permeability both describe flowing phenomena in porous solids some workers have speculated that there should be some relationship between them. In this book we have even used the same, idealised pore system model to try and understand what controls them. But, in fact the capillary pore model suggests that in general the two properties are not likely to be related. Permeability is dominated by the largest diameter pores, whereas in the resistivity model the individual pore diameters have no control on resistivity. Ultimately, at very small diameters, pore size will affect resistivity because the ions that carry the current have a finite size, but for most reservoir rocks this will not be an issue.

The analysis above suggests that both the cementation and saturation exponents should change as the rock undergoes compaction. In fact the analysis suggests that *m* is most likely to increase as the net overburden stress increases. This is because connections between pores will be squeezed and even closed-off completely as the stress increases. Nowadays it is normal to measure the saturation parameters at simulated reservoir stress but in the past this was not always done and many laboratories simply did not have the capability to make resistivity measurements under stress. Figure 8.14 compares *m* values measured at ambient
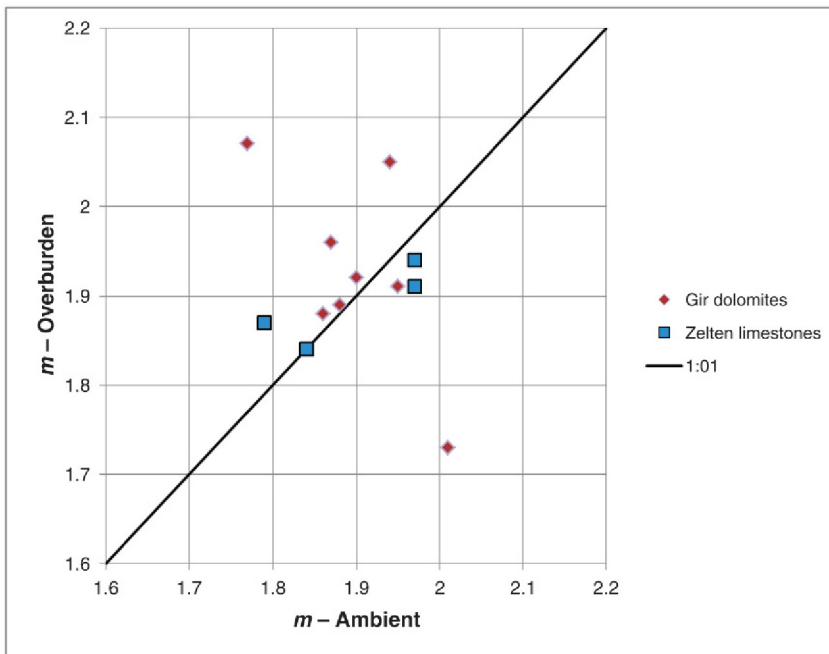
FIGURE 8.14   Saturation exponents measured at ambient conditions and under a simulated reservoir stress of 2100 psi compared. The measurements refer to individual 1½ in. plugs. On average the *m* value increases at overburden conditions.

and reservoir stress for some Libyan dolomites and limestones. In that case the average m value does increase for the dolomites but it hardly changes at all for the limestones and some individual plugs from both lithologies actually show a reduction. Some of the changes are large enough to suggest caution needs to be exercised when using any data that was obtained at ambient conditions.

8.4.4   Saturation Exponent

So far all the discussion has concerned m (and possibly a). The saturation exponent could be analysed in a similar way, but we then have to consider exactly how the water is distributed in the pore system when oil or gas is introduced. The fact that n often has similar values to m suggests that the hydrocarbons are tending to replace water more or less uniformly in a typical clastic rock. The interested reader can satisfy themselves that the capillary tube model would give an n value of 1 providing the oil or gas distributes itself uniformly along the pores.

   It is interesting to ask if there is any general trend between n and m. The data shown in Fig. 8.15 for, admittedly, a very limited number of cases suggests the only obvious trend is that n tends to cover a wider range of values than m and that for the sandstones at least n tends to be higher than m.

8.5   UNCERTAINTY AND ERROR ANALYSIS

Water saturation will always be more uncertain than porosity if only because it involves both porosity and V_w in its definition. In fact an analysis of Eq. 8.1 shows that the uncertainty in S_w increases as the inverse square of porosity regardless of what particular method is used. In other words uncertainty increases rapidly with decreasing porosity.

   In the case of the Archie equation with its three or four parameters, there are plenty of sources of uncertainty. It is impossible to generalise on how significant these different uncertainties are because the sensitivity of S_w to changes in parameters, depends on the magnitude of porosity and saturation itself. What is certain however is that changes in saturation parameters cause the following changes in calculated S_w.

   Calculated S_w increases with:

1. increase in m;
2. increase in n;
3. increase in a;
4. increase in R_w.

   These are at least reasonably memorable, as an increase in parameter always leads to an increase in water saturation. Conversely, if one feels the need to 'enhance' the quality of a reservoir, decreasing any or all of the saturation parameters will do the trick (as will be seen later the worse the quality of the
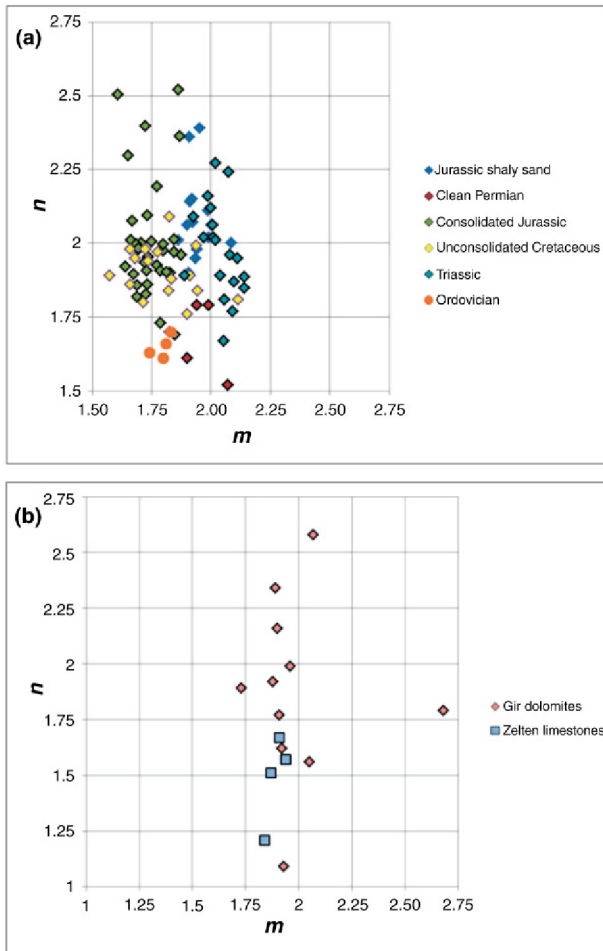
FIGURE 8.15   Cross-plots of saturation exponent n against cementation exponent m for (a) several sandstone formations from NW Australia and (b) two carbonate formations from the Sirte Basin, Libya. The m and n scales are the same in each plot but note the enlarged scale used for the carbonates.

reservoir the larger the impact of such changes). Fortunately, there are limits to how much tinkering can be got away with and the more information one has, the lower the room for manoeuvre. If a water leg is present the analyst is normally quite constrained and even if one parameter is decreased, another generally has to be increased to give 100% water (e.g. consider a Pickett plot where reducing the gradient m results in an increase in the intercept). For inter-granular pore systems at least, it is difficult to reconcile a log response in which resistivity falls with increasing porosity with anything other than a water leg or a low saturation of residual oil or gas.
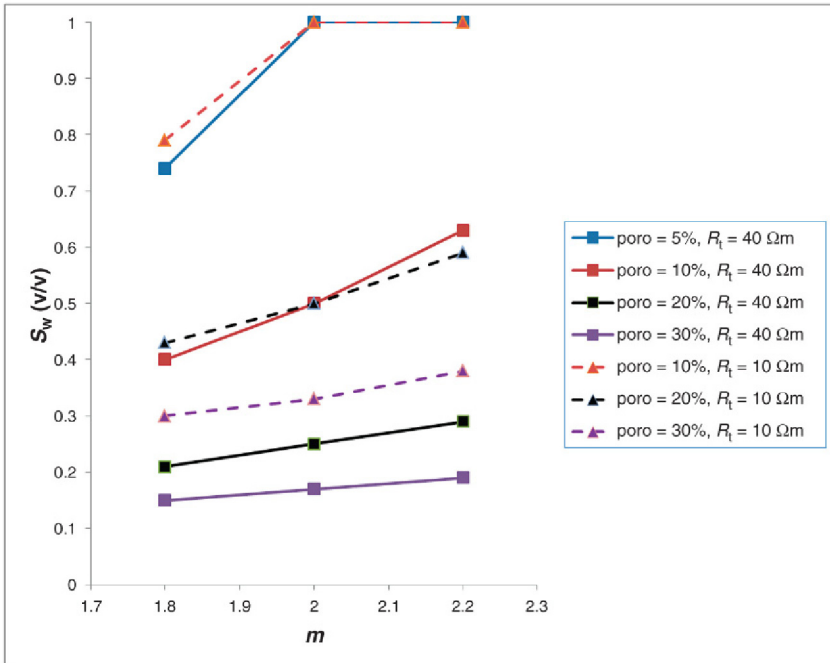
FIGURE 8.16   A diagram showing how sensitive an S_w calculation is to a change in m value. Porosities of 5, 10, 20 and 30% have been investigated with resistivities of 40 Ωm (solid lines and square symbols) and 10 Ωm (dashed lines and triangles). The other parameters R_w, n and a are fixed at 0.1 Ωm, 2 and 1, respectively.

Exactly how significant a change in parameter is, could be found by a careful analysis of the Archie equation but it is easier and probably more informative to experiment either by repeating a computer log analysis with different m, n, R_w etc. or by doing a few Archie calculations on a typical pair of porosity and resistivity values. An example of the latter, used to investigate a change in m of ±0.2 is shown in Fig. 8.16. Obviously this only applies to the particular cases described in Fig. 8.16, but it does illustrate a couple more generalities:

1. The absolute and relative change in S_w for a particular change in m is greatest at low porosities.
2. The absolute change in S_w is greatest at high water saturations (or low resistivities).

In the introduction to the book and in the previous sections it was pointed out that Archie's original equation was modified by including an extra – fourth – parameter named 'a' or the tortousity constant. Parameters a and m are not independent and as a increases, m decreases and vice versa. So in Figure 1.1 the best fit to the data used an a of 0.56 and an m of 2.21 (although they were not given there). The simpler fit with a fixed at unity gives an m value of 1.78.

TABLE 8.1 Calculated Formation Factor, R_0 and Water Saturation for Alternative Archie Parameters

Porosity	A = 0.56, m = 2.21			a = 1, m = 1.78		
	FF	R_0	S_w	FF	R_0	S_w
0.15	37.34	3.73	0.31	29.18	2.92	0.27
0.27	10.16	1.02	0.16	10.26	1.03	0.16
0.30	8.05	0.80	0.14	8.51	0.85	0.15

It is interesting to ask what difference the two alternatives make to a saturation calculation. This is easily seen by substituting in a few plausible values of R_w and R_t for the minimum and maximum porosities (15–30% in this case). An example is given in Table 8.1 for R_t of 40 Ωm and an R_w of 0.1 Ωm.

The differences in S_w are never very large and at one porosity (26.5%) there is no difference at all (this is where the two, alternative fits, cross). Obviously the more different a is from one the greater will be the differences in S_w at the extremes, but mostly the choice of forced or free regression makes no material difference (again remember we are just trying to find a single set of parameters that give an acceptable approximation to the saturation everywhere). For this reason many log analysts prefer to always use $a = 1$ and reduce the number of parameters they have to deal with.

As a general rule then, Archie will give a good indication of what is in the reservoir even if the only source of saturation parameters is an educated guess. That is to say the Archie equation will reliably distinguish water bearing from hydrocarbon-bearing rocks with conventional inter-granular pore systems, even if the computed saturations are actually rather inaccurate. Uncertainty becomes a concern at low porosities and high-water saturations. In practice this means the situations that are most likely to lead to gross errors are tight gas sands, wells that only penetrate the base of the transition zone and oil and gas shales. Actual accuracy becomes important in unitisation, particular if the reservoir quality differs significantly on either side of the median line.

One type of reservoir where care with saturation parameters is needed are vuggy carbonates which, as explained, earlier tend to be characterised by high m values. If however a default value of 2 is assumed then relatively low water saturations will be calculated even if the formation is actually water bearing. A useful exercise when dealing with carbonates is to calculate R_0 assuming an m of say 2.5 and then input that back into the Archie equation with an m of 2.

The effect of changing saturation exponent n can also be investigated but it is worth remembering that n only has any influence when S_w is less than one. This only happens when R_0 is less than R_t and as R_0 ultimately depends on m, it is arguably more important to get m right.

8.6 CONDUCTIVE MINERALS AND SHALY-SAND EQUATIONS

All the analysis in the previous sections is based on the assumption that the only conductor in the rock is the formation water. In that case Eq. 8.6 holds true and the formation factor is independent of R_w. This means, amongst other things that the resistivity of a porous rock saturated with water will be much lower than the same rock containing hydrocarbons (Fig. 8.7). However, if a second conductor is present the contrast will be lower, possibly much lower. In order to go on using Archie to find saturations, the conductivity of this second conductor has to be accounted for. In effect it has to be removed so that Archie can be applied to the water alone. The conductivity introduced by the second conductor is known as 'excess conductivity'.

Excess conductivity is an enormous topic and several general techniques as well as a huge number of specific equations have been developed since the 1950s to deal with it. These equations are generally called 'shaly-sand models' because it is shales – or strictly speaking clays – that are the most common cause of excess conductivity. It is not the intention here to cover every equation, or even most of them, rather this section will look at the general principles for dealing with excess conductivity and the physical models that have been developed to explain it.

Before getting started it is important to give a health warning. Ultimately, shaly-sand equations are trying to explain why the resistivity in a hydrocarbon-bearing formation is 'so low'. An entirely plausible reason is that there is simply a lot of water in the formation and that Archie is doing a good job at quantifying that. Applying a shaly-sand model may well calculate lower water saturations, but that is simply a result of the arithmetic and the choice of parameters. The real problem is justifying their use in the first place. In the wrong hands these equations can lead to some very expensive mistakes!

Although the equations for dealing with excess conductivity are collectively known as 'shaly-sand' models, excess conductivity can be caused by other minerals. Pyrite for example is an excellent electrical conductor that commonly occurs in sedimentary rocks. In fact it is such a good conductor that at concentrations of a few per cent by volume it can reduce the resistivity to values that are below the lowest resistivities possible with water alone. Furthermore, at these concentrations pyrite forms a 'short circuit' and the resistivity has the same low value regardless of how much gas or oil is present in the pore space. An example of this is shown in Fig. 8.17. Pyrite at a concentration of several per cent by volume is present between 3045 m and 3060 m and has reduced the resistivity to about $2\,\Omega$m despite the fact that the water saturation is only about 10%. In the pyrite-free sand above 3060 m the resistivity can exceed $100\,\Omega$m. In this particular case the effect of the pyrite is so strong that it cannot be corrected and some other method would be needed to find S_w.

At lower concentrations (between 0.5% and 5% by volume say) pyrite will still suppress resistivity but not to the extent that it can no longer be used to
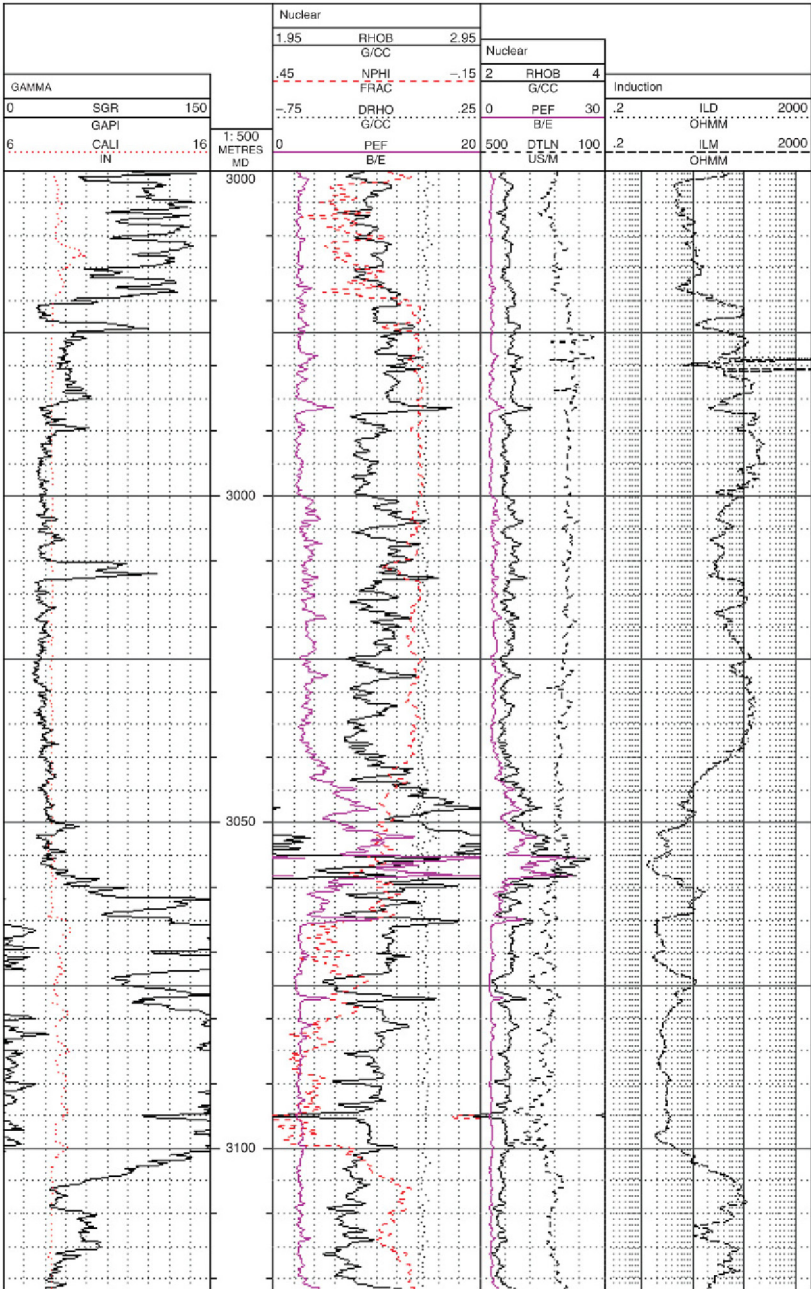
FIGURE 8.17 Conventional logs recorded in high-quality gas-bearing sands responding to pyrite. Pyrite is present at levels of several per cent by volume between 3045 m and 3060 m and causes the resistivity (right-hand track) to fall from 100 Ωm to < 5 Ωm (note also the high-density readings in this interval).

calculate S_w. In these cases if the pyrite is not accounted for the Archie equation will over-estimate S_w.

In order to account for the pyrite a quantitative model of the resistivity is needed.

It turns out to be easier to work in terms of conductivity (C), which is simply the reciprocal of resistivity. This makes the algebra simpler and clearer at the relatively cheap price of introducing a new set of symbols.

$$C = \frac{1}{R}$$

The SI unit for conductivity is the mho/m. By and large we will use the same suffixes that we used with resistivity so, for example 'C_w' is the conductivity of formation water. In terms of conductivity, Archie's equation is:

$$C_t = C_w \cdot \Phi^m \cdot S_w^n \qquad (8.9a)$$

A plausible general equation to quantify the conductivity in the presence of pyrite is:

$$C_t = C_o = C_w \Phi^m S_w^n + \text{Cond of pyrite} \qquad (8.17)$$

Equation 8.17 is saying that the conductivity of the rock is given by the sum of the conductivity due to the water and the conductivity due to the pyrite and the two components are independent. It has been assumed that the conductivity of the water is given by Archie's equation. The pyrite term is more difficult to determine, but the simplest conceivable equation is:

$$C_t = C_o = C_w \Phi^m S_w^n + C_{pyr} V_{pyr} \qquad (8.17a)$$

where V_{pyr} is the volume fraction of pyrite and C_{pyr} is its conductivity. In fact pyrite's conductivity can vary from 1 to 10,000 mho/m depending on its precise composition and impurity level. It is important to emphasise that at this stage (Eq. 8.17a) is purely a guess – it is just the simplest possible equation – if pyrite was dispersed in a similar fashion to the pore space – and therefore the water – the equation might actually be closer to:

$$C_t = C_o = C_w \Phi^m S_w^n + C_{pyr} V_{pyr}^2 \qquad (8.17b)$$

Without doing some careful laboratory experiments it is impossible to say. Nevertheless, it is instructive to do a few simple calculations to look at how the resistivity is suppressed by small volumes of pyrite.

Inserting some typical values for porosity and C_w (R_w) of 20% and 10 mho/m (0.1 Ωm) respectively, the Archie term is 0.4 mho/m (taking $m = 2$). If V_{pyr} is taken as 1% then a pyrite conductivity of 40 mho/m would result in the pyrite contributing as much to conductivity as the formation water. A conductivity of 40 mho/m (R of 0.025 Ωm) is well within the range of measured pyrite values.
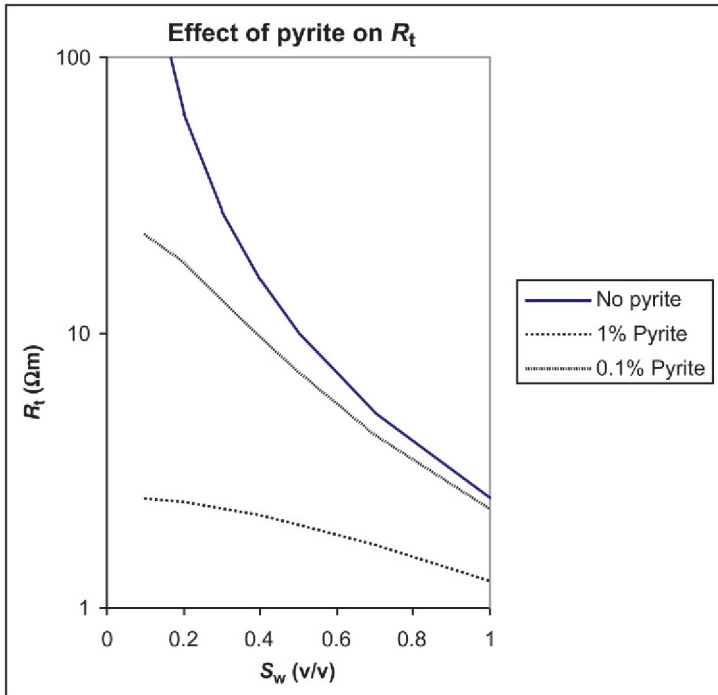
FIGURE 8.18 Resistivity as a function of water saturation for a 20% porosity sand with varying amounts of pyrite (expressed as volume per cent). R_w is 0.1 Ωm and the pyrite is assumed to have a resistivity of 0.025 Ωm.

If now hydrocarbon is introduced so that S_w is 10%, the pyrite term contributes about 99% of the total conductivity. Obviously, if this is not accounted for the S_w will be grossly overestimated. The results of varying S_w and pyrite levels are illustrated in Fig. 8.18. A pyrite content of 1% is sufficient to reduce the resistivity of *a* sand with an S_w of only 10% to the same value as a water saturated one. In other words, if the pyrite is not accounted for the hydrocarbon-bearing sand will be interpreted as water saturated!

Note that Eqs 8.17 say that the rock will conduct even if all the water is removed (by completely drying a core plug, for example). In other words the pyrite conductivity is constant and independent of the water conductivity.

8.6.1 Shaly-sand Equations

As noted in the preceding sections pyrite is a good way to introduce the idea of excess conductivity because its contribution is independent of the water composition. When clays contribute to excess conductivity this is no longer the case. It is found that the excess conductivity falls with salinity and when the rock is filled with a non-conducting fluid it no longer conducts at all.

'Archie rock'   'Classic' shaly sand   Waxman–Smits model
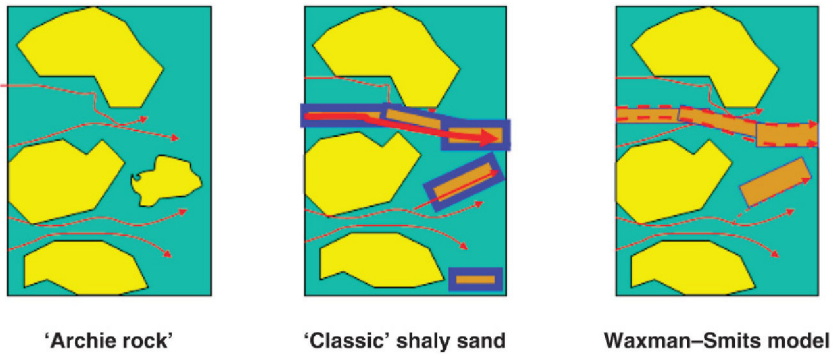
FIGURE 8.19 Alternative descriptions of excess conductivity due to clays. On the left is a clean sandstone in which all the current is carried by formation water (thin arrows [red arrows in the web version]). In the middle the clays – or shales – add extra conductivity (thick arrow [thick red arrow in the web version]). The precise mechanism is not specified, but at least some current is carried by the clay-bound water that is included in the clay component. On the right is the Waxman–Smits model where the two conductive components are all the formation water and the surface of the clay particles (dashed line [dashed red line in the web version]).

From a practical point of view, pyrite rarely affects more than a small part of the reservoir, as it tends to be quite localised. Clays on the other hand can be pervasive so that the proper treatment of excess conductivity is essential if the water saturation is not to be over-estimated everywhere. Several different types of model have been proposed to explain and deal with this. They are distinguished by exactly how the excess conductivity is generated.

Arguably the simplest models consider the whole 'shale' component as a conductor (this is very similar to the situation with pyrite). Normally, the shale includes the clay-bound water so that the two components of the conductivity are:

1. The water in the effective porosity.
2. The shale.

The alternative model considers the excess conductivity to be located at the surface of the clay particles. The two components of the conductivity then consist of:

1. All the water.
2. The surface of the clay.

The former description is thus particularly suitable for effective porosity models and the latter is better suited to total porosity models. The two rival descriptions are summarised in the cartoons in Fig. 8.19.

8.6.2 Shale Volume Models

The equations to deal with whole shale conductivity normally have a form similar to Eq. 8.17:

$$C_t = C_w \varnothing^m S_w^n + C(V_{shale}) \qquad (8.18)$$

The different equations are, roughly speaking, distinguished by the precise form of $C(V_{shale})$ (many actually are a bit more complicated than Eq. 8.18 and often they cannot be re-arranged to give an explicit equation for S_w.) The shale term invariably involves shale conductivity, which is therefore one of the parameters of the equation (C_{shale}, although in practice this is normally entered as the shale resistivity). One of the reasons for the popularity of these models is that this key parameter can simply be read from the logs.

The simplest example is the Simandoux equation where the excess conductivity term due to the shale has the form

$$C(V_{shale}) = kC_{shale}V_{shale} \qquad (8.19a)$$

where k is a constant that is equal to one in water-bearing sands and is less than one in the presence of oil or gas (how much less is a matter of trial and error although see in a later section for one common interpretation). Almost as simple is the Hossin equation in which the shale term is given by:

$$C(V_{shale}) = C_{shale}V_{shale}^2 \qquad (8.19b)$$

Note the similarity to our pyrite Eqs 8.17a and b. The conductivity of the formation is the sum of contributions from:

1. Formation water in the conventional pore system.
2. The entire shale component.

The porosity is implicitly the effective porosity and these models require that to be calculated as part of the overall analysis. The conductivity of the shale – C_{shale} – is found from the resistivity reading in a suitable shale bed. The criteria for a suitable shale are exactly the same as for any shale parameter (thick enough to be properly resolved etc.).

In order to apply Eq. 8.18 it is re-arranged to give:

$$C_w \varnothing^m S_w^n = C_t - C(V_{shale}) \qquad (8.18a)$$

which can then be further manipulated to give an explicit equation for S_w. The equation only makes sense if the term on the right-hand side is positive, in other words the total conductivity is greater than the conductivity due to the shale.

$$C_0 > C_t > C(V_{shale})$$

Whether this applies or not, depends on the parameters including C_{shale} (or R_{shale}). For Simandoux and Hossin the shale term is proportional to C_{shale} so the higher C_{shale} (lower R_{shale}) the more likely the inequality breaks down, even though it is at low R_{shale} values that a shaly-sand equation is most likely to make a material difference. Most software includes 'traps' which simply set S_w to 100% if the in-equality is not obeyed. This is a mixed blessing it avoids computational 'busts' but it does not alert the log analyst to a bad choice of parameters.

Most of the other shaly-sand equations are a bit more complicated than the Simandoux-type equations. In fact even the Simandoux equation – as normally implemented – is more complicated than Eq. 8.19a. In general these changes make it impossible to solve for S_w explicitly (in other words, to re-arrange the equation to write $S_w = \dots$). On the other hand, at least some of the equations reduce the possibility of the excess conductivity exceeding the total conductivity.

The Simandoux equation, for example is often implemented in log analysis software as the 'modified Simandoux equation'. The modification consists of making the constant k equal to S_w, so that the equation we are trying to solve is:

$$C_t = C_w \varnothing^m S_w^n + C_{sh} V_{sh} S_w \qquad (8.18b)$$

In common with most shaly-sand equations this cannot generally be solved analytically (although for the special cases of $n = 1$ or 2 it could be). The normal approach is to use numerical techniques to find S_w.

As already noted there are numerous alternatives to the Simandoux equation (the Baker Atlas Chart book lists no fewer than 16, in addition to the core-based methods discussed in subsequent sections). These involve at least one more parameter than the Archie equation and hence they introduce more uncertainty (the resistivity of the shale). The shale volume curve also introduces some uncertainty of course.

In much the same way that Archie's equation can best be understood by drawing graphs of resistivity against porosity and/or saturation so, an equation like Simandoux can be better appreciated by constructing graphs. Now however, shale volume is also an input so there are several possible parameter combinations to consider.

As noted at the start of this section, this type of model in which the excess conductivity is a function of the shale volume, is designed to work with effective porosity and implicitly give an effective water saturation (S_{we}). This is an important consideration when deciding how much difference a shaly-sand model makes because as pointed out in the introduction S_{we} is always less than or equal to S_{wt} (at least for conventional reservoirs). In fact re-arranging Eq. 8.3 gives:

$$\frac{\varnothing_e}{\varnothing_t} = \frac{(1 - S_{wt})}{(1 - S_{we})} \leq 1 \qquad (8.3a)$$

So for example if the ratio of effective to total porosity is 0.8 and S_{wt} is 50%, S_{we} will be 37.5%. This is simply the result of switching from a total at an effective porosity model. Unfortunately, not all packages overtly distinguish S_{we} and S_{wt} and so there is a danger that when comparing, say, an Archie and a shaly-sand S_w, that like is not being compared with like. Worse still, is if the effective water saturations computed by models such as Simandoux are combined with a total porosity. To avoid any chance of mixing total porosity with effective saturations it is probably best to do a few spot calculations using a

spreadsheet in order to get a feel for what typical saturations are. If necessary, effective water saturations can always be converted to total water saturations using Eq. 8.3.

Before moving on, mention should be made of the most commonly encountered of these equations: The Indonesia equation:

$$C_{t} = (C_{w}^{1/2} \emptyset^{m/2} + C_{shale}^{1/2} f(V_{shale}))^{2} S_{w}^{n} \tag{8.19}$$

Mathematically, the important feature is the summing of the Archie and shale terms within the brackets before multiplying by the S_w term. Although complicated the equation has the benefit that when re-arranged it will never result in S_w being the nth root of a negative number. This probably explains why an equation with a strong clue in its name for where it was developed, is used all over the world! In fact, the only mathematical 'bust' occurs if the term in brackets is smaller than C_t (when S_w will be calculated as greater than one). Exactly the same thing will happen with Archie, if C_o is less than C_t and in both cases the issue can be avoided by appropriate choice of parameters.

8.6.3 Total Porosity Models

Although the shale volume models described earlier can always be converted to a total porosity basis, they are fundamentally effective porosity models. The electrical current is split between the water in the effective porosity and the shale. In general no attempt is made to explain which part(s) of the shale is contributing to conduction. The alternative approach is to assume that the two conductive components consist of all the water and an excess conductivity due to the solid minerals. Specifically, it is assumed that the excess conductivity occurs at the surface of the mineral grains in contact with water. All mineral surfaces can conduct to an extent, but it is certain types of clays that are capable of adding significant excess conductivity. The excess conductivity is quantified by a property known as the cation exchange capacity (or 'CEC,' this is defined in later section). A high-excess conductivity is favoured by a high CEC and or a high surface area.

The best known of these models is the Waxman–Smits equation which was developed by scientists at Shell Oil and will be described in detail later. Like the Archie equation it is firmly based on resistivity measurements on core plugs.

As was explained in Section 2.5.1 clays are sheet silicates. They can be thought of as being derived from quartz by replacing some of the silicon with a metallic element (generally aluminium, magnesium or iron). In quartz, every silicon atom bonds to four oxygen atoms, which are arranged tetrahedrally around it. Every oxygen atom bridges two silicon atoms so that a highly symmetrical lattice is built. This structure is chemically very stable and mechanically very strong. In clays, the silicon still bonds to four oxygen atoms but the oxygen bridges from silicon to a metallic element. This creates layers of silicates, which may be electrically charged.
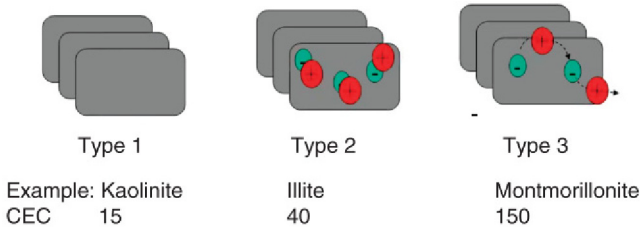
FIGURE 8.20 Schematic diagram illustrating the differences between the three clay types. Only Type 3 clays have large numbers of mobile cations at the surface of the clay particles.

From a petrophysical perspective clays can be classified into three types:

Type 1. The layers are not charged.

Type 2. The layers are negatively charged but the charge is balanced by fixed cations (positively charged ions).

Type 3. The layers are negatively charged and the charge is balanced by mobile cations.

The three types are illustrated schematically in Fig. 8.20.

It is the mobile cations in the Type 3 clays that cause the significant excess conductivity in shaly sands.

The best-known example of a Type 1 clay is kaolinite, and Type 2 is exemplified by illite, where potassium ions balance the charge on the sheets. The micas and chlorite also exhibit this type of charge balancing. Because these charges are fixed they cannot carry a current.

The Type 3 clays are the immature minerals such as montmorillonite and smectite that tend to be associated with young sediments. This provides an important clue as to where classical shaly-sand analysis is most likely to be required. The CEC is basically the number of mobile cations per unit mass of rock, typical values for some commonly occurring clay minerals are given in Table 8.2.

The unit for CEC is 'milli-equivalents per gram', written meq/g (or more commonly per 100 g). A 'milli-equivalent' is equal to 6.10^{20}, this is Avagadro's number divided by 1000. So the CEC really does give the number of mobile

TABLE 8.2 Typical Cation Exchange Capacities for Four Clay Minerals

Clay	Type	CEC (meq/100 g)
Illite	2	10–40
Kaolinite	1	3–15
Smectite	3	80–150
Chlorite	2	10–40

cations per unit mass and it is basically a 'concentration'. It is measured by exposing the rock to an excess of ammonium ions (NH_4+), which replace the naturally occurring moveable ions, normally sodium (they are said to 'exchange' with mobile cations, which explains the name). The ammonium ions are then themselves removed by exposure to boric acid. Their concentration in the acid is then measured using conventional analytical techniques. Needless to say the measurement is destructive.

Clearly the more mobile ions there are, the higher the excess conductivity. So excess conductivity depends on both the type of clay and its volume fraction. A small amount of montmorillonite can suppress resistivity far more than a large amount of kaolinite. For a shaly sand – as opposed to the pure clay minerals – a high value of CEC is something in excess of 10 meq/100 g.

8.6.3.1 Waxman–Smits Equation

The Waxman–Smits equation provides the link between CEC and excess conductivity and is based on experimental work carried out in the 1950s by Waxman, Smits, Thomas, Hill, Milburn and many others (appropriately these all worked in the Shell Oil research laboratories where Archie had done his original work). The derivation of the equation is elegant but quite involved. Some readers may prefer to skip this section until they really need to use the equation in anger.

The first step is to convert CEC to milli-equivalents per unit volume rather than mass. Specifically, it is unit pore volume that is of interest, the resulting property is Q_v with units in milli-equivalent per cubic centimetre (it is simply called '$Q-V$'). Conversion from mass to volume can be fraught with pitfalls but the basic steps were explained in Section 1.4. The CEC normally refers to dry rock and if core is available, routine core analysis provides the necessary inputs: porosity and grain density (see box).

It is easiest to show the conversion with an example:

Sample	Porosity	Grain density	CEC
1	10.7%	2.68 g/cm^3	0.36 meq/100 g

Firstly, find the dry rock volume equivalent to 100 g, this is simply 100 divided by the grain density which in this case turns out to be 37.13 cm^3. This is the volume of 100 g of sandstone grains, that is it does not include the pore volume. The total volume of the porous rock is therefore given by:

$$V_b = \frac{V_g}{1-\varnothing} = \frac{37.13}{(1-0.107)} = 41.78 \, \text{cm}^3$$

The pore volume is 10.7% of this or 4.47 cm^3. Finally Q_v, which is the number of milli-equivalents per unit pore volume, is 0.139 meq/cm^3.
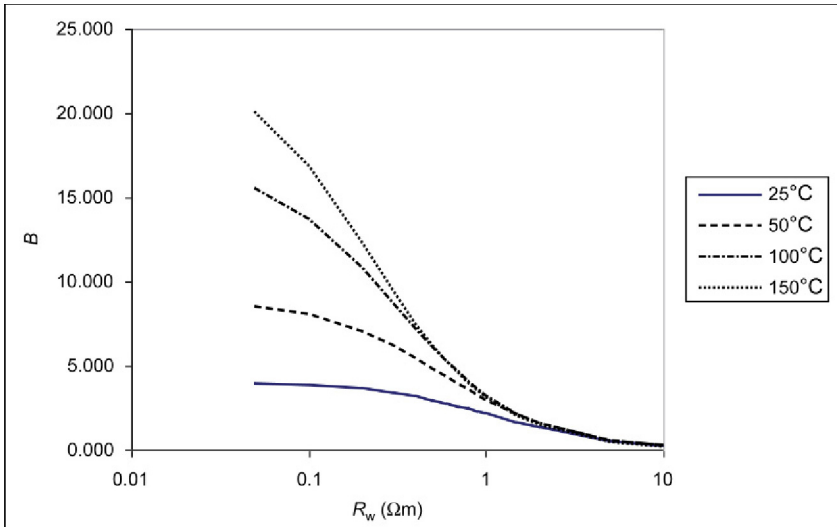
FIGURE 8.21   Waxman and Thomas' B function showing dependence on R_w and temperature.

The excess conductivity is directly proportional to Q_v and the constant of proportionality B is a function of temperature and R_w.

$$C_{excess} = B(R_w,T)Q_v \qquad (8.21)$$

Although simple this is quite a profound equation as it says the excess conductivity is a product of a property of the water (B) and the rock (Q_v). The function B was determined experimentally by Waxman and Thomas (1974). Its form is shown in Fig. 8.21 and it is worth noting that although this function is used every time the Waxman–Smits equation is used it is based on only a handful of core plugs.

The most important feature of Fig. 8.21 is the dependence on R_w, this is tacitly implying that the excess conductivity varies with R_w. It actually gets smaller with increasing R_w (decreasing salinity). As we will see however, the conventional Archie conductivity falls even more rapidly so that the excess term becomes more significant at low salinities.

The best way of showing the effect of the excess conductivity, is to plot the conductivity of the plug (C_o) against the conductivity of the water used to saturate it (C_w). For rocks that obey the Archie equation this results in a straight line that passes through the origin. Some real examples with non-zero excess conductivity are shown in Fig. 8.22.

For the clean rocks that Archie studied the two are proportional (the constant of proportionality is just $1/F$ (see Section 8.3.2).

$$C_o = \frac{C_w}{F} \qquad (8.22)$$
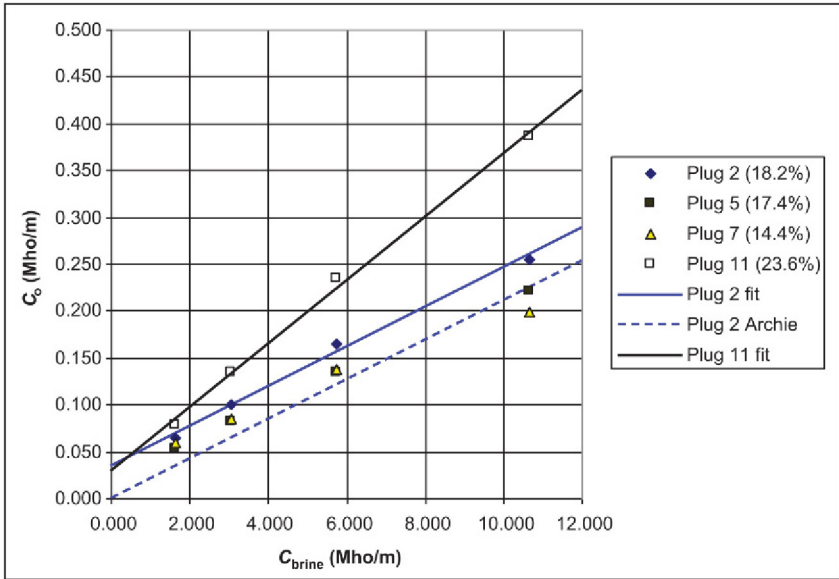
FIGURE 8.22   A plot of plug conductivity against brine conductivity for four shaly-sand plugs (porosity in brackets). The intercept with the C_o axis gives the excess conductivity (the hypothetical conductivity of the dry plug). The gradients are $1/F$ and so depend on porosity and 'm'.

For shaly plugs however the relationship has the extra term so that

$$C_o = \frac{C_w}{F} + C_{excess} \qquad (8.23)$$

This is of course almost the same as Eq. 8.18 but with F being written in place of the explicit porosity term and without a specific dependence on shale volume. In terms of a plot like Fig. 8.22, C_{excess} is simply the intercept on the C_o axis. In principle it is the conductivity of a plug filled with a non-conductive fluid. But because excess conductivity caused by clay depends on the composition of the formation water, at low water conductivities – freshwater – the conductivity of the plug is lower than would be predicted by Eq. 8.23. At the limit of zero water saturation, the plug will not conduct ($C_o(C_w=0)=0$). To summarise C_{excess} is a function of C_w (R_w) and is zero when C_w is zero.

So in the Waxman–Smits equation, C_{excess} is the product of the two terms B and Q_v (Eq. 8.21). To repeat, B is a property of the formation water and Q_v is a property of the rock. Ideally Q_v is measured on core plugs as part of the special core analysis but B is just a function of salinity and temperature. The Waxman–Smits equation is actually written as:

$$C_o = \frac{(C_w + B(C_w)Q_v)}{F^*} \qquad (8.22)$$
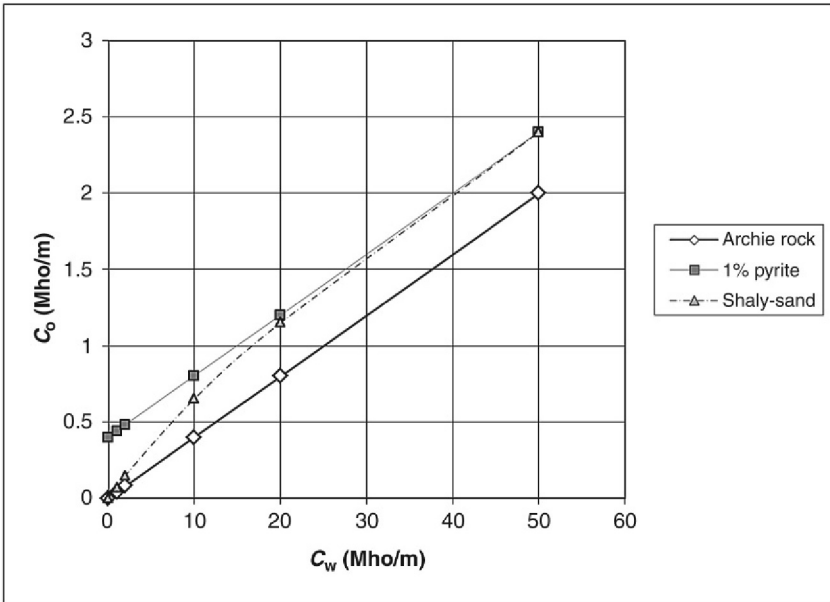
FIGURE 8.23 Plots of rock conductivity (C_o) against water conductivity (C_w) for various types of sand. The Archie rock has no conductive minerals in it and all the current is carried by the water. In this case the plug conductivity (C_o) is proportional to the conductivity of the water (C_w). If pyrite is present it will add a constant additional conductivity and the proportionality will be lost. Clays also add excess conductivity but the magnitude of this falls as the water gets fresher.

So the formation factor F^* is slightly different than that of a simple Archie rock. It is the gradient of the straight line of the C_0/C_w plot and can be written as a function of porosity and a special cementation exponent m^*. Equation 8.22 has been explicitly written with B as a function of C_w to remind us that the excess conductivity varies with formation water salinity. The intercept on the C_0 axis is equal to BQ_v/F^* but the B is the value at high C_w. At low salinities the function curves down towards the origin. Inspection of Fig. 8.22 shows evidence of this for C_{ws} less than 6.0 mho/m.

This curvature is in contrast to pyrite where C_{excess} is independent of C_w (Eq. 8.17). So a plug in which excess conductivity is caused by pyrite will conduct even when it is dry. Three hypothetical C_0/C_w plots are given in Fig. 8.23 for the case of no excess conductivity, constant excess conductivity and excess conductivity that falls with C_w.

In principle Q_v could be calculated from a knowledge of the clay component in the rock (which could be determined by X-ray diffraction, for example). In practice it is better to measure it by constructing a plot like Fig. 8.22. This is a standard special core analysis technique known as C_0/C_w (pronounced 'C-nought-C-W'). As explained in the introduction to this section, it is also possible to measure the CEC using 'wet chemistry' and ultimately that can be

converted to Q_v and excess conductivity. Most petrophysicists prefer the C_0/C_w method because it measures BQ_v directly.

8.6.3.2 Hydrocarbons and Waxman–Smits

Although we have covered a lot of ground we have yet to put any hydrocarbon into the system. We will cover this next but first let us summarise the key ideas of the model. The excess conductivity occurs at the surface of the clay layers and will only occur if they are in contact with water. That water however is counted as porosity and is not distinguishable from water in the normal inter-granular pore space.

When hydrocarbons enter the pore space it is assumed they lie in the centres of the pores away from the mineral surfaces (the reason for this will be discussed in a later chapter). The reduction in water makes the sand appear more shaly and Q_v increases to:

$$Q_{vhc} = \frac{Q_v}{S_w} \tag{8.23}$$

The full Waxman–Smits equation is then written:

$$C_t = \frac{S_w^{n*}(C_w + BQ_v / S_w)}{F*} \tag{8.24}$$

This reduces to Archie's equation if $Q_v = 0$. Like most shaly-sand equation it is not generally soluble explicitly. (In other words we cannot rearrange it so that $S_w = \ldots$).

For certain values of $n*$ an explicit equation can be written however and given the assumptions involved in getting to 8.24 one could be justified in making that simplification ($n* = 2$).

8.6.3.3 Applying the Waxman–Smits Equation

The Waxman–Smits equation was the first attempt to deal with the excess conductivity introduced by clays in a total porosity model. It is a remarkably successful model that is still routinely used but it does suffer from a number of disadvantages, which are summarised here:

1. The method relies on special core analysis (to measure the excess conductivity) and so in practice is dependent on measurements that are made on a handful of core plugs (which may not even have come from the reservoir of interest).
2. The equation(s) cannot generally be solved exactly (not a great disadvantage in an age of cheap computing power but originally this would have been a deterrent to its use).
3. The method only applies to total porosity models.
4. The whole model is quite complicated and requires care and patience to understand what is going on. There are some unusual properties with equally unusual units and dimensions that go into the equations (e.g. meq/100 g).

Point 1 is of practical importance because in order to apply the equation Q_v is needed at every point along the well path, although in practice it is unlikely to be measured on more than a few plugs. The normal approach is to look for a relationship between Q_v and another property. Often a relationship with porosity is used, but shale volume is another obvious candidate. The relationship is then used to compute a continuous Q_v curve, which is combined with resistivity and porosity in Eq. 8.24.

Q_v itself can be obtained from $C_0 - C_w$ plots or CEC measurements made as part of the SCAL programme. The former is normally preferred as it is a direct measurement of the excess conductivity. The latest geochemical tools can also produce a CEC curve but this is simply based on general empirical relationships.

One thing the Waxman–Smits equation does deserve credit for is emphasising that the saturation parameters m^* and n^* are different to the values that would be used with the Archie equation (the asterisk emphasises that they are different parameters). In fact m^* is greater than m and n^* is greater than n, this means that the effect of the excess conductivity is to an extent compensated by the higher saturation parameters. The reason can be seen by considering Fig. 8.24 which shows C_0/C_w plots for a hypothetical 20% porosity sand
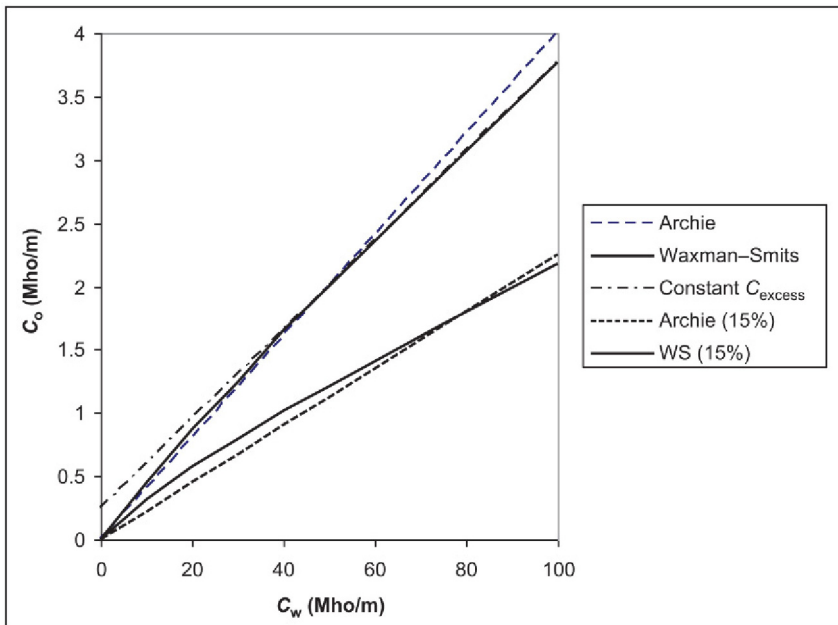
FIGURE 8.24 C_0/C_w plots for a hypothetical 20% porosity sandstone. The Archie line (dashed line [blue dashed line in the web version]) is constructed by drawing a straight line from the origin to a single measurement on a plug saturated with 50 mho/m. The curved solid line is constructed from a series of measurements using different C_w values. The two lines will cross at a C_w of 50 mho/m. The plot also shows lines for a 15% porosity sand with the same m and m^* values as the 20% sand.

saturated with water. If a single conductivity measurement is made using water with a conductivity of 50 mho/m, C_o turns out to be 2 mho/m (n.b., these are just the values chosen for this example). In the absence of any other information this rock conductivity could be explained using an Archie model with an m value of 2.00 ($F = 25$). Now assume a C_o/C_w test is conducted on a core plug with a range of C_w values. Both tests will agree on the conductivity of the rock when saturated with 50 mho/m water but for other values they will disagree. The linear portion of the full C_o/C_w line has a lower gradient, which is equivalent to a higher m value of 2.08 (or an F^* of 28.4). If we wish to use the Waxman–Smits model then we are obliged to use the higher m value. Using the Archie m in a Waxman–Smits model will under-estimate S_w, at least at low salinities where shaly-sand models are most likely to be applied.

8.7   CONCLUSIONS

This is the longest chapter in the book and it has introduced over 20 different symbols (R_t, F, m etc.) and about two-dozen equations (most of which are far from memorable). Furthermore, although long, this is far from an exhaustive account of saturation calculations. It can be easy to lose sight of what we are trying to do: finding a relationship between a physical property that can be measured with a logging tool and the volume of water in the system.

There are a large number of measurements that respond strongly to water thanks to its molecular structure. But in practice only resistivity has the depth of investigation to have any chance of determining saturation in un-altered formation. The way to do this was discovered by Archie in the early 1940s and his equation has been in continuous use ever since.

The Archie equation has subsequently been modified to deal with excess conductivity and 'briefly' describing those modifications has taken up nearly a third of the chapter. Even then, several important methods have not even been mentioned. I will close by noting again that it is very tempting to invoke excess conductivity when disappointingly low resistivities are logged across a reservoir interval. But more often than not, a low resistivity is simply a reflection of a large amount of water in the system. In other words the Archie equation is doing a good job at estimating S_w. True excess conductivity requires something other than water to conduct the current, such as pyrite or the mobile cations that occur on the surface of a special group of clays. It is important to remember that even a large amount of clay need not produce much excess conductivity but it will bind a lot of water.

Chapter 9

Hydrocarbon Corrections

Chapter Outline

9.1 INTRODUCTION

The previous two chapters explained how logs can be used to estimate porosity and water saturation but we have avoided the inconvenient fact that in hydrocarbon-bearing formations, they are not independent. Most obviously water saturation is defined in terms of porosity but calculating porosity requires a fluid parameter that is a function of saturation. So there is a classic 'chicken and egg' problem, in which we need to know the saturation in order to calculate the porosity and the porosity to calculate the saturation. The general way of approaching problems like this is by iteration. Specifically:

1. Assume a saturation and calculate the fluid parameter for that fluid mixture.
2. Make a first estimate of porosity using this fluid parameter.
3. Use this initial porosity to re-calculate saturation.
4. Use this new saturation to re-calculate the fluid parameter.

Repeat steps 1–4 until subsequent estimates of porosity and saturation are hardly changing.

More often than not the opening assumption is that the fluid is entirely water but in principle any composition could be used. This type of simple, repetitive calculation is ideal for computers and any log analysis package will offer at least some combinations of porosity and saturation models to do this. Many packages allow the use to mix and match any porosity model with any saturation model. This is a short chapter that looks at the way iterative methods work as well as some refinements and potential pitfalls.

9.2   INTEGRATING DENSITY POROSITY WITH ARCHIE SATURATION

As a specific example of the iterative method consider calculating porosity from density and saturation using the Archie equation. This is one of the simplest conceivable combinations.

The general steps outlined above are implemented by:

A1 Calculate porosity assuming the fluid is all water with density RHO_{water}.
A2 Input this porosity to the Archie equation to produce a first estimate of water saturation S_{w1}.
A3 Calculate the fluid density of the mixture:

$$RHO_{fluid1} = S_{w1}RHO_{water} + (1 - S_{w1})RHO_{hc}$$

B1 Re-calculate porosity using RHO_{fluid1}.
B2 Input this porosity to the Archie equation to produce an improved estimate of water saturation S_{w2}.
B3 Re-calculate the fluid density of the mixture:

$$RHO_{fluid2} = S_{w2}RHO_{water} + (1 - S_{w2})RHO_{hc}$$

C1 Re-calculate porosity using RHO_{fluid2}.
Etc., etc…

The calculation is normally stopped if the changes in porosity and/or S_w between iterations, fall below a particular threshold, or a particular number of iterations is reached. For a simple method, like the one above, this normally occurs after one or two iterations but obviously without such limits the programme will continue in an endless loop.

Porosity is calculated using Eq. 7.3b but now there are three parameters that are needed: the density of the matrix, the density of water (RHO_{water} in the above) and density of the hydrocarbon (RHO_{hc}). Saturation is calculated using the Archie equation and no additional parameters are needed for that (yet). To see how this specific method works in detail we will look at one example. The values of porosity and water saturation at each iteration are shown in Fig. 9.1. The parameters are as follows:

Density g/cm³	Matrix	Water	Hydrocarbon
	2.65	1.05	0.20

Saturation	R_w	a	m	n
	0.1	1	2.0	2.0

The initial fluid density assumed was that of the formation of water 1.05 g/cm³ but after three iterations it had settled down as 0.39 g/cm³, the value for a mixture of 78% gas and 22% water. Inspection of the graphs showed that the porosity and saturation hardly changed after two iterations and quite accurate
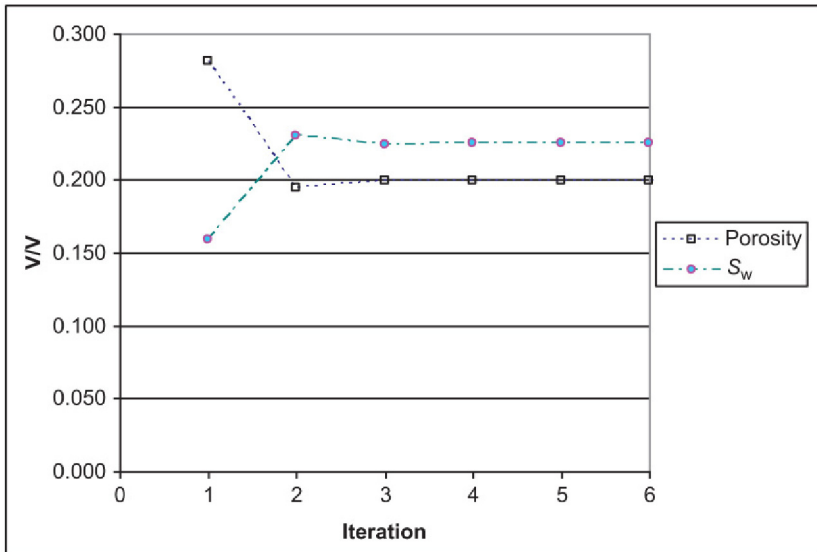
FIGURE 9.1   Iterative calculation of density porosity and water saturation using the methodology described in Section 9.2. The formation density is 2.2 g/cm³ and the resistivity is 50 Ωm (parameters are given in the text).

results had been calculated after just one iteration. This is quite typical for a simple scheme like this one.

9.3   COMPLICATIONS AND REFINEMENTS

The basic principles applied in the simple model of Section 9.2 apply to any integrated porosity and saturation model. The great virtue of these methods is that the hydrocarbon correction to the porosity varies continuously in response to changes in saturation. There are some problems with the basic scheme, however which are addressed by various refinements. Broadly speaking the problems can be classified as:

1. Physics of the measurements.
2. Invasion.
3. Extending the method to include shale volume and shale effects.

9.3.1   The *Z/A* Correction

Of these the simplest to deal is the physics of the measurement and to illustrate the problem we will stick with the simple density porosity and Archie saturation model described in Section 9.2. In Section 5.2 it was noted that the density log exploits the scattering of gamma rays by electrons and that there is an almost linear relationship between the true density – mass per unit volume – and the

number of electrons per unit volume. 'Almost' is the operative word. Most of the lighter elements that dominate the composition of the commonly occurring minerals in sedimentary rocks have almost equal numbers of protons and neutrons in their nuclei. For example, nearly all oxygen is the ^{16}O isotope with 8 protons and 8 neutrons, similarly most silicon nuclei comprise 14 neutrons and 14 protons. The relative numbers of protons and neutrons is summarised in a ratio known as Z/A where Z is the atomic number (equal to the number of protons) and A is the atomic mass number (protons plus neutrons). The rule for light elements can be written as:

$$\frac{Z}{A} = \frac{1}{2}$$

Since Z is also the number of electrons, a measurement of electrons per unit volume is directly proportional to the mass per unit volume (i.e. the density).

Unfortunately, hydrogen represents an exception to the rule as the normal isotope has no neutrons and so Z/A is one. Therefore, as hydrogen is introduced to the mix Z/A gets larger and the electron density will overestimate the true density. Most hydrogen enters the system as water and because this is a well-understood substance the relationship between electron density and true density can be corrected. But if some of the hydrogen enters as hydrocarbon the difference between the two densities is no longer constant. If however the porosity and the true saturation are known, the amount of hydrogen can be determined and the electron density can then again be transformed to the true density.

For completeness it should be noted that heavy elements have proportionately more neutrons, so that Z/A is less than half (e.g. the most abundant isotope of uranium has 146 neutrons and 'only' 92 protons). But the heavy elements with a large excess of neutrons are rare in most of the minerals of interest to us.

Returning to the iterative hydrocarbon correction the simple scheme described in Section 9.2 needs to be modified to account for the Z/A change. There are various ways this could be accomplished but the normal way is to use the saturation to make a correction to the density. The actual corrections are empirical and can normally be found in the help files of the software. A commonly used density correction, originally published by Schlumberger, and for use in oil-bearing reservoirs is:

$$\Delta\rho' = 1.07\emptyset(1 - S_{X0})\{1.11 - 0.15P\rho_w - 1.11\rho_{oil} - 0.03\} \tag{9.1}$$

Where the ρ_w and ρ_{oil} are the densities of water and oil, respectively in gram per cubic centimetre and P is the salinity of the water expressed as a fraction by weight. Porosity and invaded zone saturation have to be given as fractions. This correction is intended for use with oil, a slightly different correction was developed for gas-bearing formations.

Note that this density correction – $\Delta\rho'$ – is nothing to do with the 'delta–rho' curve presented on a density log, which was described in Section 5.2. The
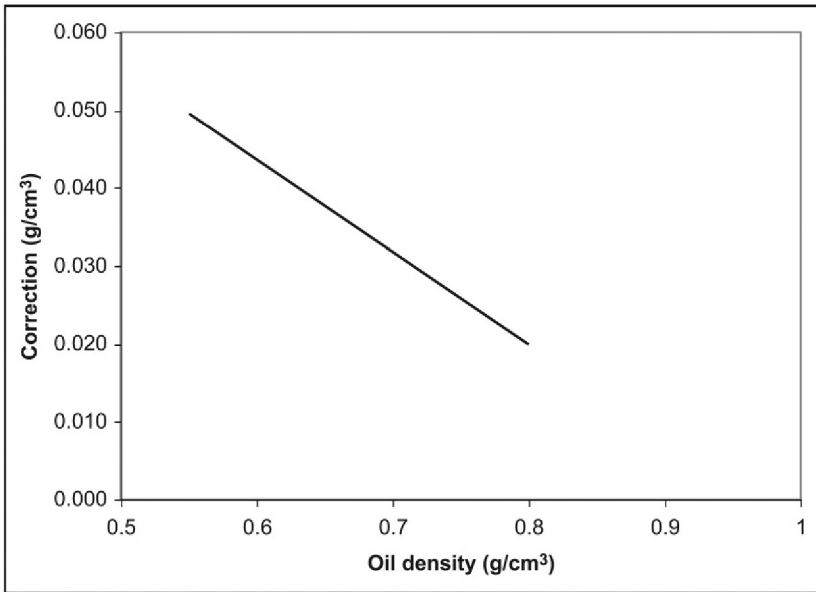
FIGURE 9.2   The magnitude of the *Z/A* correction to the density reading in an oil-bearing formation with a porosity of 20% and an invaded zone saturation of 50%. Salinity is 0.03 (30 kppm) and the water density is $1.00 \, \text{g/cm}^3$.

correction normally amounts to just a few per cent of the measured density (as can be seen by inserting some typical values into Eq. 9.1). The actual correction for a formation with 20% porosity and 50% invaded zone saturation as a function of oil density, are shown in (Fig. 9.2). If the lithology is clean sandstone, the measured density would be of the order of 2.3 g/cm^3 so the maximum correction is of the order of 2%.

The *Z/A* correction is applied iteratively along with steps A1 to A3, B1 to B3 described in Section 9.2. Having made an initial estimate of porosity and saturation, the programme would not only calculate the density of the mixture of fluids – step A3, B3 etc. – but also the correction to the formation density before making an improved estimate of porosity – step B1.

The *Z/A* correction is positive and is expected to increase with porosity and decrease with S_{x0} and oil density. This is exactly what Eq. 9.1 does, less obvious are the effects of water density and salinity. Note that here 'water refers to whatever is in the invaded zone, that could be formation water, mud filtrate or some mixture of the two'. Fortunately, the correction is not very sensitive to either water density or salinity so for this type of correction it is not really necessary to know exactly what is in the invaded zone (as explained in Section 9.3.2 it can matter for calculating the invaded zone saturation, however).

Similar corrections have been developed for neutron porosities to account for the change in hydrogen index caused by hydrocarbons.

9.3.2 Accounting for Invasion

As noted on several occasions in this book all porosity tools have limited depths of investigation and so will tend to measure the properties of altered formation. In the simple density porosity–Archie saturation scheme described earlier the density of the fluid mixture should actually refer to the near well-bore region rather than the unaltered formation implied by step A3. In fact the discussion on Z/A corrections explicitly used invaded zone saturation as an input to the density correction.

There are various ways to estimate S_{x0} but none of them are fool-proof and some can only be applied to wells drilled with water-based mud. The problem is that invasion profiles can be quite complicated and different tools have different depths – and volumes – of investigation. Two possible invasion profiles for a gas-bearing formation are shown schematically in Fig. 9.3. The cartoons show how the distributions of the three fluids – gas, formation water and filtrate – vary
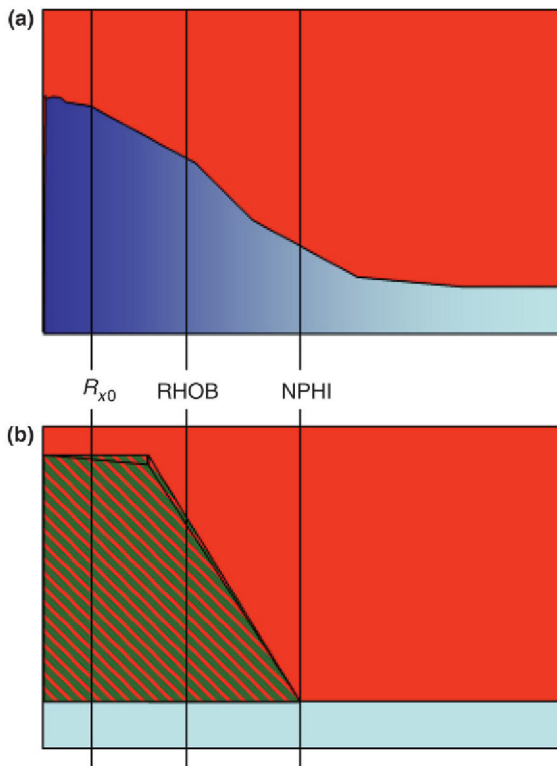
FIGURE 9.3 Cartoons showing how the fluid composition varies away from the borehole wall in a permeable bed drilled (a) with a water-based mud and (b) an oil-based mud. The distance into the formation increases left to right. The natural fluids are gas (grey, upper part of plot) and formation water (light grey at base of plots). The filtrate is coded dark grey for water-based mud and cross-hatched for oil-based mud.

with distance into the formation and the depths of investigation of three different shallow reading tools.

When a microresisivity tool is available S_{x0} can be calculated from Archie or if appropriate a more complicated saturation equation.

$$S_{x0} = \left| \frac{aR_{mf}}{R_{x0}\varnothing^m} \right|^{1/n} \tag{9.2}$$

Equation 9.2 is basically Eq. 8.12 with R_{mf} in place of R_w and R_{x0} in place of R_t. The filtrate resistivity is measured by the logging engineer and should be written on the log header. Using Eq. 9.2 as the basis of a hydrocarbon correction involves two assumptions:

1. The microresistivity (R_{x0}) is responding entirely to mud filtrate. In other words the mud filtrate has completely replaced formation water to the depth of investigation of the microresistivity tool.
2. S_{x0} is a good estimate of the saturation to the depth of investigation of the density tool (or whatever tool is being used to estimate porosity).

The first assumption is reasonably safe because the composition of the water near to the borehole wall is altered by the rapid diffusion of ions in solution, rather than bulk displacement of formation water by filtrate. Assumption 2 is more problematical because microresistivity tools have shallower depths of investigation than even the density tool (let alone neutron and sonic tools). This is certainly the case in Fig. 9.3a where there is a gradual change in gas saturation with distance and all three tools respond to different gas/water ratios. In this case S_{x0} calculated with Eq. 9.2 will overestimate the amount of water that determines the density and especially neutron readings. This will ultimately result in the density porosity being overestimated. This problem can often be seen with high-quality gas sands drilled with water-based mud where the microresistivity tool gives a low reading, characteristic of a sand almost saturated with mud filtrate, but the density corresponds to a sand with a very low water saturation (an example is shown in Fig. 9.4).

A commonly used alternative to finding S_{x0} is to use S_w taken to a fractional power (P).

$$S_{x0} = S_w^P \tag{9.3}$$

where S_w and S_{x0} are given as fractions and P is less than one.

Since S_w is always less than or equal to unity, S_{x0} will always lie between S_w and unity. A common default is to use a value of 5/8 for P this gives a value of 37% for S_{x0} when S_w is 20%. Smaller values of P give higher values of S_{x0} at a particular S_w. In other words if deep invasion is suspected a lower value of P should be used.

Both the previous ways of estimating S_{x0} apply to water-based muds. If a well is drilled with oil-based mud the shallow reading porosity tools will generally respond to a mixture of three fluids: formation water, natural hydrocarbons
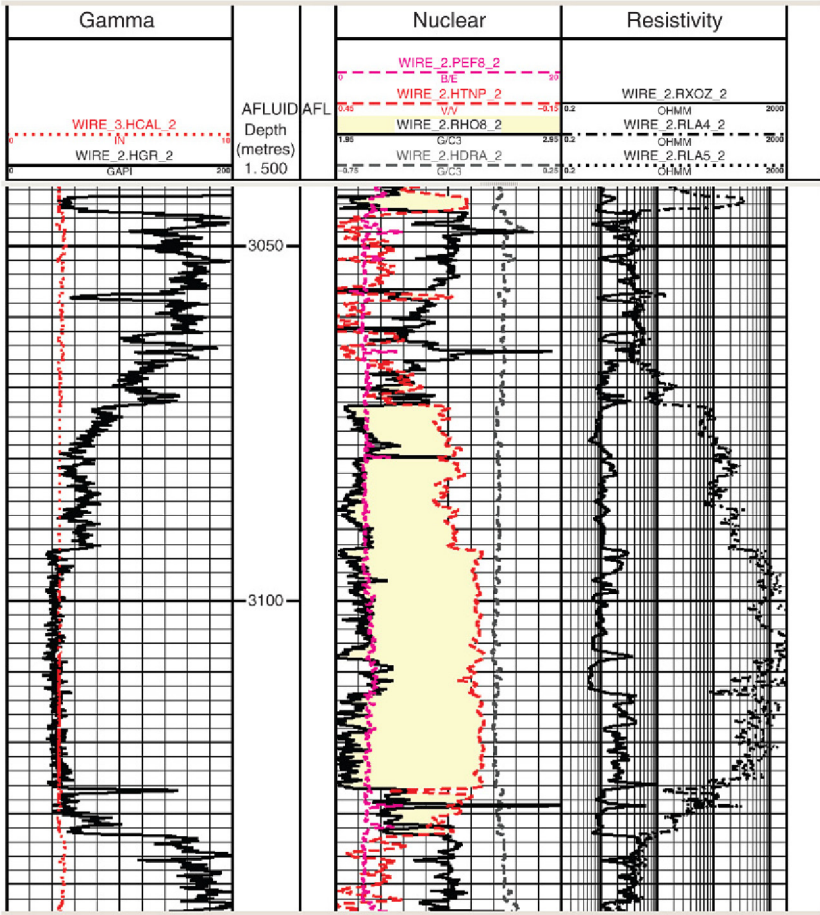
FIGURE 9.4 An example of shallow invasion by water-based mud into a permeable gas-bearing sand. The microresistivity tool reads 1 Ωm or less suggesting a very high filtrate saturation at it's depth of investigation. The density–neutron combination on the other hand has a huge gas effect. The low density reading of 2 g/cm³ or less suggests it is responding to an average water saturation of 30% or less.

and oil-based mud filtrate. In this case because the only water present is formation water, the Archie equation can find the water saturation quite accurately. But clearly, resistivity cannot help find the relative amounts of the filtrate and the gas. In Fig. 9.3b the invaded zone is coded with cross-hatching to emphasise that gas and filtrate are present but in an unknown ratio.

There are two approaches to this problem. Firstly, one can arbitrarily fix the saturation of the filtrate saturation at a particular value, S_{mf} say. The density of the mixture is then given by:

$$\text{RHO}_{\text{fluid}} = S_w \text{RHO}_{\text{water}} + (1 - S_w - S_{mf}) \text{RHO}_{\text{hc}} + S_{mf} \text{RHO}_{mf} \qquad (9.4)$$

The density of oil-based filtrates – RHO$_{mf}$ – is typically about $0.85\,g/cm^3$ and the exact value can be found from the mud engineer's reports. Notice that the water saturation is S_w even though we are working in the near well bore region. This assumes the water is at irreducible saturation and so is immobile (see Chapter 10). In that case the filtrate can only displace hydrocarbon (this is what is shown in Fig. 9.3b).

In a transition zone and the water leg, water is mobile and can be displaced by filtrate so that the water saturation seen by the porosity tool may actually be lower than the unaltered formation, S_w. One could possibly model this by using Eq. 9.3 with a value of P greater than one. But it is worth remembering the assumption of a fixed filtrate saturation is an approximation and so it is questionable whether it is worth trying to find the invaded zone water saturation to a high degree of accuracy.

The second way of dealing with invasion by oil-based mud filtrate is to actually attempt to measure the oil saturation. This is just about possible with one of the new NMR or sonic tools or possibly by combining density and neutron readings but it is not often done and it is unlikely to give a unique solution.

9.3.3   Shale Volume

If shale volume is calculated using a cross-plot method then it too will be influenced by the presence of hydrocarbons. A gas effect on the density–neutron combination, for example will move the log point away from the shale so that a shaly gas-bearing sand can give similar response to a clean water-bearing sand. The net result is that the density–neutron cross-plot will underestimate the shale volume in the presence of gas. The problem can be avoided by correcting the density and neutron logs for hydrocarbons. For example, the corrected density can be found from the porosity by substituting the formation water density into the density porosity equation and a similar reasoning gives the neutron porosity.

In a simple scheme like the one discussed in Section 9.2 this is easy to apply, problems will occur when the shale volume becomes an input to porosity and/or saturation as the iteration scheme is becoming quite complicated. At best it can take quite a lot of iterations to get everything – shale volume, porosity and saturation – to converge and at worst the algorithms can become quite unstable and end up giving extreme – and obviously erroneous – porosities.

9.4   THE NEUTRON LOG RE-VISITED

Everything discussed earlier for the density log can be generalised to any porosity tool. Each tool has one or two peculiarities, however. Moreover the neutron and sonic tools tend to read deeper into the formation and the influence of invasion will on average be less.

It is worth spending a bit of time on the neutron log, if only to cover some of the terms that the reader may come across. More importantly, methods based on the neutron–density cross-plot that account for hydrocarbons are so widely used

that it is worth looking at the specific corrections to the neutron log. We have already seen that the neutron porosity is reduced when gas replaces water. The simplistic explanation is that gas has a lower hydrogen index than water (i.e. it has less hydrogen atoms per unit volume). This is not the whole story, however. Consider two identical sandstones one saturated with water and one containing 90% gas and 10% water. To a first approximation at ambient condition the gas-bearing sand could be treated as a water-bearing sand with a lower porosity of:

$$\text{Apparent porosity} = \emptyset S_\text{w}$$

(9.5)

But this is actually quite a poor approximation for estimating the neutron response. At reservoir conditions methane densities are typically in the range $0.15–0.2\,\text{g/cm}^3$ which is equivalent to 0.01 mol of methane per cubic centimetre. Since each methane atom has four hydrogen atoms, the concentration of hydrogen is $0.04\,\text{mol/cm}^3$. For water with density $1.0\,\text{g/cm}^3$ and only two hydrogen atoms per molecule the equivalent number is $0.11\,\text{mol/cm}^3$ (Chapter 1). So gas has a hydrogen index that is just under half the value for water. An example of how the hydrogen index varies for water- and gas-bearing sands with a porosity of 25% is shown in following sections.

Porosity	S_w	S_gas	HI
0.25	1.00	0.00	0.25
0.25	0.20	0.80	0.12

These imply the gas-bearing sand should give the same neutron porosity as a water-bearing sand with approximately half the porosity. In reality neutron porosities of gas-bearing sands can be a lot lower than half the true porosity. For example, in Fig. 9.4 the neutron porosity is approximately 6 pu although the true porosity is about 25%.

The neutron porosity is actually suppressed more than would be expected from the reduction in hydrogen index. The reason is that a true 12% porosity, water-bearing sand consists of 88% solid minerals, mainly silicates if it is a typical sandstone, whereas our example only has 75% solids. This change in the volume fraction of the solid component causes an additional change in the neutron porosity. In Section 5.3 it was noted that although the tool is designed to respond strongly to hydrogen it does respond to everything else as well so it is not surprising that the 25% porosity gas-bearing sand behaves differently to a 12% water-bearing sand. The difference is quantified as the 'excavation effect'. The name comes from the fact that in the gas-bearing case some of the solid has been 'excavated' to make way for the gas. The excavation effect results in a further reduction in neutron porosity over and above that caused by the reduction in hydrogen index. Log analysis packages typically account for it as a function of porosity, hydrocarbon saturation and hydrocarbon HI. One equation that is often used is:

$$E = 2\emptyset^2(1 - S_\text{x0})(1 - \text{HI})(1 - (1 - S_\text{x0})(1 - \text{HI}))$$

(9.6)

Where $\emptyset$ is the true porosity and S_{x0} is the water saturation seen by the neutron tool (invasion will probably make this higher than the value calculated in the reservoir). Equation 9.6 predicts excavation factor increases with porosity, hydrocarbon saturation and decreases with the hydrogen index of the hydrocarbon. In other words it will be at its maximum in very porous, gas-bearing reservoirs.

A number of accounts – especially on the internet – incorrectly claim the excavation effect is simply caused by the lower amount of hydrogen in the formation when gas replaces water. It should be clear from the above that it is only indirectly related to the gas. In a hydrocarbon-bearing formation the neutron log needs to account for both the reduction in the hydrogen index and the excavation effect. The former is quantified by the hydrogen index of the hydrocarbon (or possibly it's neutron porosity). The excavation factor is the additional term. These are the third and fourth terms in Eq. 9.7.

In a clean hydrocarbon-bearing formation the neutron porosity is given by:

$$\Phi_n = \emptyset + (1-\emptyset)\Phi_{n,ma} - B\emptyset(1-S_{x0}) - E(\emptyset) \tag{9.7}$$

Where $\emptyset$ is the true porosity and Φ_n is the neutron porosity as read by the tool. $\Phi_{n,ma}$ is the neutron porosity of the matrix, E is the excavation factor and B is a constant that gives the hydrogen index of the hydrocarbon. If shale is present an additional term for the neutron porosity of the shale is needed, but for simplicity we will just deal with clean formations here. Writing the equation for neutron porosity in this way emphasises that the excavation factor reduces the neutron porosity over and above the reduction caused by the lower hydrogen index of the hydrocarbons.

The constant B is related to the hydrogen index of the hydrocarbons. The precise way it is related to density and molecular weight was explained in Chapter 1. For gas consisting mainly of methane the hydrogen index is given by

$$HI = 2.2\rho_{gas} \tag{9.8}$$

For oil it is often assumed to be $\rho_{oil} + 0.3$ but it could be calculated if the density and composition of the oil is known. In either case the density has to be given in gram per cubic centimetre.

If the formation water is pure water HI is the value of B. But in most practical realisations B is also used to account for the change in the HI of the water resulting from the dissolved salts that are invariably present in formation waters and mud filtrates. A typical expression for B for use in a log analysis programme is:

$$B = \frac{\rho_w(1-P) - 2.2\rho_{gas}}{\rho_w(1-P)} \tag{9.9}$$

where P is the salinity written as a fraction.

As an example of the size of the two hydrocarbon corrections consider the neutron log shown in Fig. 9.4. At 3100 m the porosity is about 25%, the water saturation is 5% and the fluid properties are:

Density water	$1.0 \, g/cm^3$
Density gas	$0.2 \, g/cm^3$
Salinity	20 kppm (0.02)

The well was drilled with water-based mud and the microresistivity curve shows some shallow invasion has taken place so we will assume S_{x0} is 10%. Substituting these into Eq. 9.9 gives $B = 0.53$. Then substituting this into Eq. 9.7 and assuming the lithology is sandstone with a matrix neutron porosity of -0.03 the neutron porosity becomes:

$$\Phi_n = 0.10 - E$$

Calculating the excavation factor using Eq. 9.6 and the above parameters gives a value of 0.03 so that the neutron porosity is predicted to be 7%. The actual value read by the tool at 3100 m is 7–8%.

In practice of course we normally wish to find porosity from the neutron log and so need to re-arrange Eq. 9.7 to get an explicit equation for porosity. In practice this will not be possible and the equation will have to be solved using the same type of iterative approach as was introduced at the start of the chapter.

Chapter 10

Fluid Distribution

Chapter Outline

10.1 INTRODUCTION

Water saturation at a point in the reservoir was defined in Section 2.3. It is simply the fraction of the pore volume occupied by water. In this chapter, we describe the physical processes that determine the water saturation and consequently how it varies with location in the reservoir. This can be used as a predictive tool – the saturation-height function – to determine the saturation away from well control.

Hydrocarbon column height is always a factor in determining saturation, as it is directly related to the buoyancy forces trying to force hydrocarbon into the pores. Saturation is also determined by the fluids, the minerals making up the rock and the nature of the pore system. A knowledge of all of these allows, in principle, water saturation to be calculated at any point in the reservoir. This has many practical applications of which the most important is to accurately calculate the hydrocarbon in place.

In practice, the saturation has to be written as a reasonably simple function of column height, rock properties and fluid properties. The equation

that does this is conventionally named the 'saturation-height function' even though saturation is usually determined by more than just column height. This is an acknowledgement that even in a perfectly homogeneous reservoir, saturation will vary with height. Practical saturation-height functions involve at least one rock property, for example permeability. So we could write

$$S_w = S_w(h, k) \qquad (10.1)$$

It is worth emphasising that a saturation-height function is just an empirical fit to some experimental data and/or log analysis. A lot of time can be wasted arguing over the precise form of the equation(s) to be used when in reality anything that captures the right shape will do.

Fundamentally, saturation is determined by the interplay of gravity and the microscopic forces that exist between the water, hydrocarbon and mineral grains. The latter make up the so-called capillary forces and they depend on pressure, temperature and the chemistry of the rock and the fluids.

One other factor needs to be introduced before discussing the fundamental forces and that is the order in which fluids are introduced to the pore space. Nearly all the rocks we are interested in were originally deposited in aquatic environments and start out saturated with water (the exception is Aeolian sandstones and even they are likely to be saturated with water early in their lives). The hydrocarbon charge arrives later and displaces some of the water from the pore space. This process is referred to as 'drainage' and in the sub-surface it occurs:

1. When the reservoir is charged.
2. In gas storage schemes when gas is injected into water-bearing rocks.
3. In geosequestration when CO_2 is injected into an aquifer (most likely as a supercritical fluid).

It also occurs in the laboratory when water is removed from core plugs for whatever reason.

It is also possible for water to displace hydrocarbons. This is imbibition and in the sub-surface this occurs when:

1. Hydrocarbon leaks from a structure because of a seal failure.
2. Hydrocarbons are produced by water drive.
3. Aquifer pressure increases and forces water into the hydrocarbon column.

It occurs in the laboratory whenever a core plug is saturated with water.

Drainage and imbibition are not normally reversible so, for example a rock initially saturated with water that is first drained and then allowed to imbibe water, seldom if ever returns to a completely saturated state. This means that the saturation at a particular point in the reservoir will be different depending on whether the rock there has most recently undergone drainage or imbibition (Fig. 10.1).
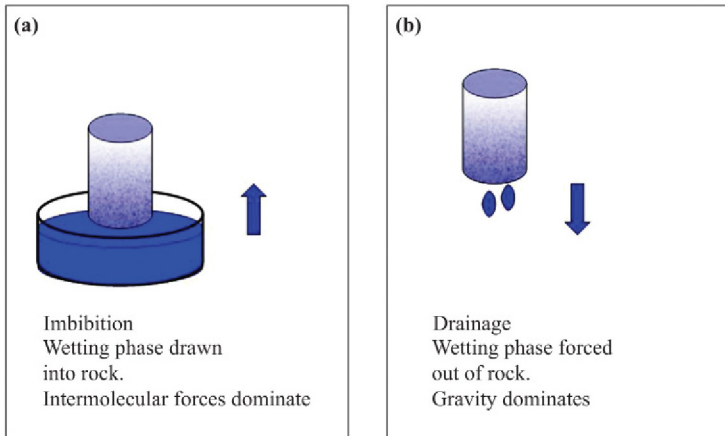
FIGURE 10.1 The concepts of (a) drainage and (b) imbibition.

10.2   GRAVITATIONAL FORCES AND BUOYANCY

In any fluid column pressure increases with depth: the deeper one goes the greater is the mass of fluid pushing down on a unit area. The rate the pressure increases is directly proportional to the density of the fluid. This can be confirmed by measuring pressure at different depths and constructing a graph. For the typical fluids, conditions and column lengths of interest here, fluid densities are more or less constant in a particular reservoir. This means the pressure increases linearly with depth and at a contact between two fluids of different densities, the rate will change suddenly. Some typical densities of interest are given in Table 10.1. The pressure at any depth within a fluid column is given by:

$$P = gh\rho \tag{10.2}$$

where ρ is the density, h is the depth below the datum and g is the acceleration due to gravity. (If SI units are used pressure is given in Pascals.)

Gravity acts to segregate the water and the hydrocarbon. Since water is almost always the densest fluid the gravitational force, pulling towards the centre of the earth, is stronger on the water than the hydrocarbon and so the former sinks. If gravity were the only force acting, all the water would sink to the bottom of the structure and the water saturation would make a step change to zero at the contact. This is what happens with a very simple system like a glass of water where the fluids of interest are water and air with densities 1.0 and 0.0006 g/cm³, respectively (Fig. 10.2a). The saturation-height function is simply given by:

$$S_w = 1 (d > \text{FWL})$$
$$S_w = 0 (d < \text{FWL}) \tag{10.3}$$

where d is the depth and FWL is the depth of the water surface. This is illustrated in Fig. 10.2c. If pressure was measured at various points within the glass

TABLE 10.1 Densities of Some Fluids Encountered in Hydrocarbon Reservoirs and the Laboratory

Fluid	Pressure (psia)	Temperature (°C)	Density (g/cm³)	Pressure gradient (psi/m)
Water (fresh)	14.7	25	1.00	1.421
Water (fresh)	7000	200	0.985	1.40
Light oil	2000	100	0.70	0.99
Heavy oil	2000	100	0.85	1.20
Methane	14.7	25	0.0007	0.0003
Methane	7000	200	0.17	0.24
CO_2	14.7	25	0.020	0.028
CO_2	5000	120	0.34	0.48
Mercury	14.7	25	13.53	19.2

the gradient will suddenly increase at the water surface ($d = $ FWL). In practice the density of the air is so low that the pressure will appear not to change above the water surface – zero pressure gradient – but in the water column it will start to increase at a measurable rate of 1.42 psi for every metre below the surface. The point where the gradient changes, in other words where the water and air lines intersect, is known as the free water level (FWL). In the case of the glass of water it coincides with the gas water contact (GWC) but this need not be the case when the fluids are held in a porous solid.

In real reservoir rocks *in situ*, it is very unusual to reach a water saturation of zero and the pore space contains at least a few per cent by volume of water.
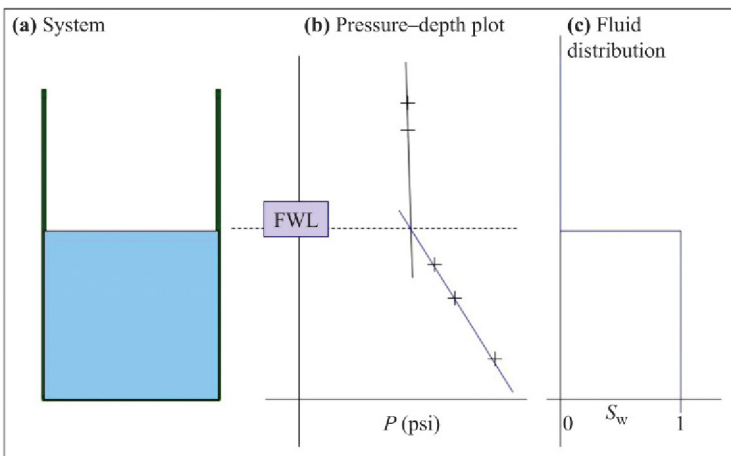
FIGURE 10.2 Fluid distribution in a glass of water. (a) Sketch of the system. (b) Pressure–depth plot through the contact showing the location of the FWL. (c) A graph of the saturation-height function that describes how water is distributed.

This observation tells us straight away that in real porous rocks gravity is not the only force influencing the distribution of the fluids. (Note: a zero water saturation might be encountered in porous rocks that are above the water table on land and possibly at greater depths as a result of some highly unusual recovery mechanisms. An S_w of zero is also deliberately achieved in core analysis during drying.)

Before discussing the capillary forces we will briefly introduce the concept of buoyancy. This is really just a different way of expressing the effects of gravity. Instead of looking at fluid segregation as the sinking of the denser fluid, buoyancy emphasises the lighter fluid rising. This is a more intuitive way of describing the charging of a structure: buoyancy forces drive the hydrocarbon upwards to the crest, but in reality it is still gravity that drives the whole process. The pressure P driving the lighter fluid upwards is given by:

$$P = gh\Delta\rho \qquad (10.4)$$

where g is the acceleration due to gravity, h is the height above the contact (or strictly FWL) and $\Delta\rho$ is the density difference between water and hydrocarbon. If density is measured in gram per cubic centimetre and the pressure in pounds per square inch. The equation can be written:

$$P = 1.42h\Delta\rho \qquad h \text{ in metres}$$
$$P = 0.433h\Delta\rho \qquad h \text{ in feet}$$

Note: The latter equations only apply near the Earth's surface, if you ever have to derive a saturation-height function for use on another planet or in the mantle use 10.4!

Equation 10.4 says that, all things being equal, the buoyancy pressure increases with density contrast and height above the contact (Fig. 10.3). This means:
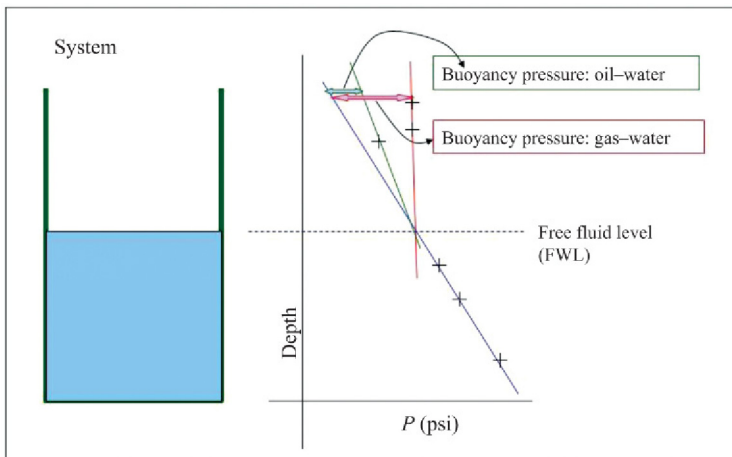
FIGURE 10.3 Diagram illustrating the dependence of buoyancy pressure on the contrast in density between the two fluids and the height above free water level (HAFWL).

1. At a particular height above the contact there will be more pressure pushing gas into the pore space than oil.
2. The higher up a column one goes, the greater is the force pushing hydrocarbon into the pore space. This is the reason why, on average, water saturation tends to fall with height above a contact.

10.3   CAPILLARY FORCES

To understand capillary forces we need to consider what is happening at the molecular level. In particular, it is the concepts and models of surface chemistry that explain what is happening. We are typically interested in three substances: water, hydrocarbon and mineral grains (in reality the latter two consist of a number of different chemical compounds but for simplicity we will ignore this complication). The molecules making up these substances interact with each other and so there are at least six different interactions to consider:

Water–water	Water–HC	Water–mineral
	HC–HC	HC–mineral
		Mineral–mineral

These interactions are called intermolecular forces and they are similar to, but much weaker than, the bonds that hold individual molecules together. Water molecules form particularly strong intermolecular bonds to other water molecules and certain mineral grains. One consequence of this is that the boiling point of water is surprisingly high for a light molecule (essentially it takes a lot of heat to break the intermolecular bonds).

The intermolecular forces in water result from the fact that the hydrogen–oxygen bonds are polarised so that the hydrogen atom has a slight positive charge and the oxygen a negative one. This in turn means that hydrogen from one water molecule is attracted to the oxygen atom of a neighbouring water molecule and overall the water molecules tend to clump together. In a hydrocarbon molecule the carbon–hydrogen bonds are not strongly polarised and the forces between molecules are far weaker than in water. Amongst other things this means methane boils at $-162°C$ whilst water boils at $+100°C$ (note that both molecules have almost the same mass). If water is mixed with methane the water molecules attract each other far more strongly than they attract methane. So even if the two fluids are thoroughly mixed initially, the water molecules quickly tend to clump together and separate out (see Fig. 10.4). Notice it is the mutual attraction of the water molecules that drives the segregation and not a repulsion between water and methane. Neither is mutual attraction between methane molecules a factor in the segregation. Methane molecules are in fact no more strongly attracted to each other than they are attracted to water. They end up separating out together because there is nowhere else for them to go.
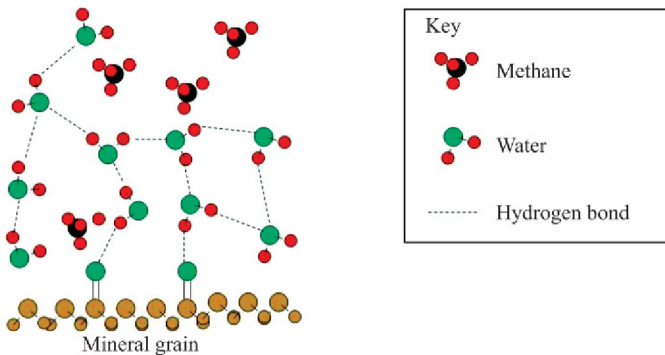
FIGURE 10.4   Grain surface in the presence of water and methane at the molecular level. Water molecules form relatively strong bonds to the mineral surface and other water molecules. Methane only forms weak bonds to the surface and thus tends to get displaced into the centre of the pores.

10.3.1   Solid–Fluid Interactions

The three fluid interactions have been covered (water–water, water–hydrocarbon and hydrocarbon–hydrocarbon). We now need to consider the two fluid to mineral interactions. The minerals we typically deal with are silicates and carbonates. As is often the case in petrophysics however, it is best to start with simple idealised model systems before attempting to understand real reservoir fluids in real porous rocks. Capillary forces are often introduced by discussing glass tubes of varying diameters (capillary tubes). Glass is a silicate that, unless specially treated, has a high affinity for water (it is said to be 'strongly water wet'). If a small diameter tube touches the surface of a large volume of water, some of the water will be spontaneously drawn into the tube and will rise to a level above the surface. Note that this is an example of imbibition, as the wetting fluid is being drawn into a system where it was at a low saturation.

The rise is caused by the strong attraction of the water to the glass, it stops when the force of attraction balances the weight of the water in the tube. Another way of explaining the rise is to assume that the water within the tube is at a different pressure to the bulk of the water. The difference is the capillary pressure which is an important input to saturation models. In order for the water to rise the pressure must be lower in other words the capillary pressure is negative. To bring the water in the tube level with the bulk of the water – in the beaker – an excess positive pressure would have to be applied to the tube. This is numerically equal to the capillary pressure.

The mass of the water in the tube is given by:

$$M = \pi r^2 H \rho \tag{10.5}$$

Where the symbols have their usual meanings (H and r are shown in Fig. 10.5). The weight is M multiplied by g, acceleration due to gravity and

r: radius of the tube
H: 'capillary rise'
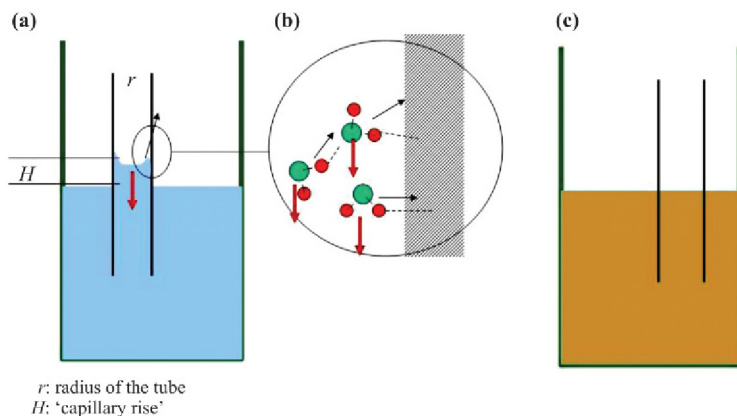
FIGURE 10.5 Capillary rise. (a) The balance of forces causing water to rise within the glass tube. The intermolecular attraction between the water molecules and the glass (black arrow) balance the weight of the water that has risen (dark grey [red in the web version] arrow). (b) Intermolecular forces cause water molecules to adhere to the glass and gravitational forces attempt to pull them back into the container. (c) A fluid composed of molecules with identical forces with similar molecules or the surface of the tube. In this case there is no rise.

this is the force that the intermolecular attraction between water and glass must balance. The total force of attraction on the water is directly proportional to the circumference of the tube ($= 2\pi r$). So as the tube gets bigger, the downward force due to the weight increases more rapidly than the upward force due to intermolecular attraction. The net effect is that the rise (h) is inversely proportional to the radius of the tube.

The pressure the column of water would exert if it were not held up by the capillary forces is given by:

$$P = Mg/A \tag{10.6a}$$

where A is the cross-sectional area of the tube. So,

$$P = H\rho g \tag{10.6b}$$

Note the similarity to the equation for the buoyancy pressure (Eq. 10.4). There is an important difference, however in Eq. 10.4; h, the height above the free water level, is a variable in Eq. 10.6b; H, the capillary rise, is a constant that is determined by the balance of forces.

If the water is replaced by a light oil, the capillary might be expected to be higher because the hydrocarbon is less dense. But in fact the rise is likely to be small. This is because, as discussed earlier the hydrocarbon is not particularly strongly attracted to the glass (if at all).

A careful examination of the water in the tube will show that its surface is curved downward (Fig. 10.5a). At the microscopic level this is a result of the

water molecules nearest the glass adhering more strongly to the glass than to neighbouring water molecules (Fig. 10.5b). (Careful examination of a glass of water shows the same 'pull-up of the water surface at the sides of the glass). It follows that when oil is the liquid it will not only rise less but the curvature of its surface will be lower too. If the strength of the interaction between liquid molecules is the same as that between the molecules and the glass there will be no rise and the surface of the liquid will be perfectly flat (Fig. 10.5c).

Mercury actually gets depressed in the glass tube, that is to say the level within the tube will be below the bulk of the liquid. This is a result of the forces between the mercury atoms being far stronger than the forces between the mercury and the glass. Consequently, pressure actually has to be applied to the bulk of the mercury to push it into the tube. This pressure however, is still the capillary pressure. (If a metal tube were used the forces between the walls of the tube and the mercury atoms are similar to those between the mercury atoms and the rise/depression will be small.)

10.3.2   Interactions Between Water and Real Rocks

The simplest of the silicates is quartz (SiO_2) in which ideally every silicon atom is bonded to four oxygen atoms and every oxygen atom is bonded to two silicon atoms. This arrangement is easily accommodated within a single crystal but at it's surface, the ideal arrangement cannot be satisfied. An oxygen atom, for example may end up bonded to a single silicon atom and its other bond 'dangles'. This 'dangling' oxygen atom, with its associated negative charge' will attract the hydrogen of a water molecule very strongly (it may even actually satisfy its bonding requirements by 'stealing' a hydrogen atom). The net effect is that a layer of water molecules adheres strongly to the surface of the mineral grains (Fig. 10.6). By contrast, hydrocarbon molecules are not strongly attracted to the mineral grain and so tend to get displaced to the centres of the pores (again note
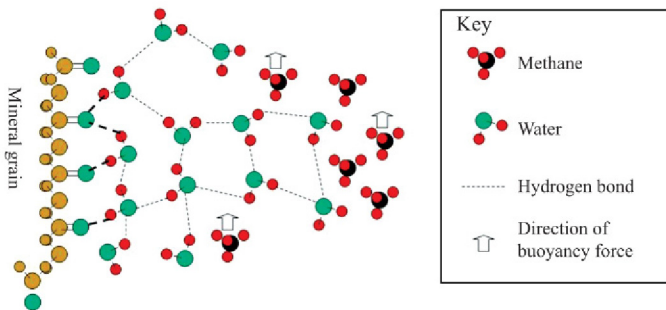
FIGURE 10.6   The interactions between water and the surface of a silicate grain at the molecular scale. A layer of water molecules adheres tightly to the mineral grain. These in turn adhere to other water molecules resulting in a layer several molecules thick. The water molecules furthest from the mineral surface are more easily displaced by methane molecules responding to the buoyancy force.

they are not repelled by the mineral grains they are just not attracted as strongly as the water).

The intermolecular forces between the water molecules and the mineral grains are actually stronger than the gravitational forces pulling the water molecule towards the centre of the Earth. Some water will therefore remain adhering to the mineral grain and the water saturation will never drop to zero. Furthermore, the layer of water molecules bound to the mineral surface will itself bind other water molecules so that a layer, several molecules thick, will be present. This water is referred to as irreducible water and its saturation is the irreducible water saturation (S_{wir}). This water cannot be displaced no matter how great the buoyancy force of the hydrocarbon or equivalently how much pressure is applied to the non-wetting phase.

The irreducible water saturation depends on the total surface area of the mineral grains and consequently tends to be higher in fine-grained rocks than coarse-grained rocks of the same porosity. As an extreme example if all the porosity is represented by a single large vug, the irreducible water saturation may be very low because it has a relatively low surface area. In a very fine-grained siltstone, the total surface area of the pores is large and a lot of water can be held in place.

Real clastic rocks consist of more minerals than quartz but the same principles apply. This is also true of carbonates where again oxygen can end up at surface with its normal bonding arrangement unfulfilled. In clay minerals water is often considered to be part of the chemical structure and so if they are present, water is too. In any case water adheres strongly to clays and irreducible water saturations are normally high in clay-rich rocks.

10.4 WATER IN POROUS ROCKS

In Sections 10.2 and 10.3, respectively, the two competing processes of buoyancy forces (gravity) and capillary forces (intermolecular forces) were introduced. Although they were discussed separately in reality they are both present and it is the competition between them that controls the distribution of water and hydrocarbon in reservoir rocks.

Put simply buoyancy forces act to drive hydrocarbon towards the crest of a structure and water towards the base. The downward movement of water – and therefore the upward movement of hydrocarbon – is opposed by capillary forces, which tend to bind water to the mineral grains. The buoyancy forces increase with height above the FWL and so are capable of overcoming stronger and stronger intermolecular forces. This means that as the height above the FWL increases, hydrocarbon can enter smaller pores and overall the water saturation falls. As noted above however, there is a limit to how much water can be removed and high in the column the water saturation falls to a limiting value: S_{wir}.

If hydrocarbon is pushed, by buoyancy forces, into a structure it will tend to be confined to the centres of the pores. Even so it will form a continuous column all the way down to the contact. Therefore at any depth the pressure will reflect

the mass of the hydrocarbon pressing down from above and the gradient will be determined entirely by the density of the hydrocarbon. The water on the other hand is ultimately bound to the mineral grains and does not contribute to the pressure. Below the FWL, some water is still bound to the mineral grains but in the centre of the pores it connects up to form a continuous phase that determines the pressure. The pressure gradient is then given by the density of the water. The water and hydrocarbon pressure–depth lines intersect at the FWL (which has some important practical applications).

In a porous rock the FWL is not necessarily the same as the gas (or oil) water contact which is defined as the depth at which the water saturation first falls below 100%.

To summarize:

FWL is defined by pressures.
Gas/oil water contact (GWC/OWC) is defined by saturation.

In the example of the glass of water they coincided and it is a philosophical point as to whether they coincide or not in porous rocks. Normally the GWC is believed to occur above the FWL, in other words the first gas is not encountered until a certain distance above the intersection of the water and gas lines. The reason for this is that the buoyancy forces are not large enough to overcome the capillary forces until a certain distance above the FWL. In practice it is worth remembering that saturation can only be measured with limited accuracy and so the depth at which the water saturation falls to 99% is going to be subject to an uncertainty of at least tens of centimetres. In any case it is frequently found that a well intersects the contact in a non-reservoir facies, so the contact cannot be picked.

10.5   WETTABILITY

The relative strength of the intermolecular forces between water and the mineral grains and oil and the mineral grains is termed the wettability. The reservoir described earlier, in which water adheres more strongly to mineral grains than hydrocarbons, is said to be a 'water wet rock'. The qualitative model described earlier, suggests that all clastic rocks composed of silicate minerals (i.e. most of them) and all carbonate rocks are water wet. It is possible however to come across oil wet rocks in which the hydrocarbon molecules apparently form stronger bonds to the mineral grains than water.

Oil wet rocks typically occur where the hydrocarbon is a heavy, biodegraded oil. These oils contain a relatively high concentration of fatty acids and other natural surfactants and it is these that allow hydrocarbons to bind to the mineral grain surfaces. These molecules contain two parts:

A hydrocarbon 'tail' similar to a typical light oil molecule.
A polar group, often containing oxygen, that can bond as or more, strongly than water to the mineral surface.
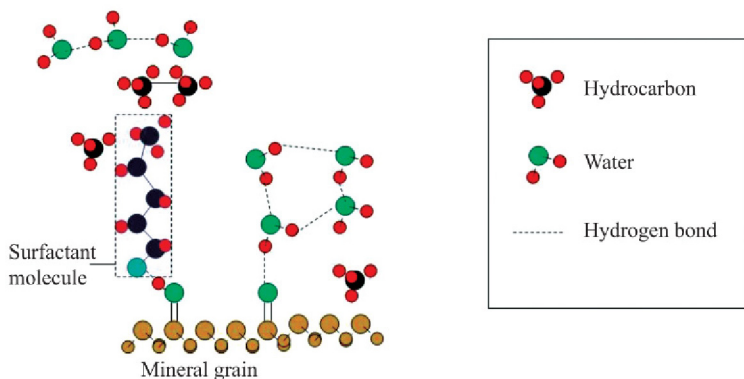
FIGURE 10.7   The origin of oil wet rocks. Natural surfactants can allow hydrocarbons to bond to mineral grains that are fundamentally water wet. The polar end of the surfactant adheres to the mineral grain for the same reason that water does. The hydrocarbon tail is directed into the pore space and prevents water from approaching.

The result is that the surfactant molecule is orientated such that the polar head faces the grain surface and the non-polar tail is directed into the pore. The grain thus appears to be coated in oil and water can no longer access the mineral surface. The water will then concentrate in the centre of the pore (Fig. 10.7).

There are a few generalities that can be drawn from this:

1. Sedimentary rocks start out as water wet systems.
2. Wettability is generally the result of the fluid the rock was in contact with rather than a primary property of the rock.
3. Oil wetness is most likely to be found with heavy, biodegraded oils.
4. If the rock has only ever been in contact with light hydrocarbons it will remain water wet (the most obvious case being gas).

Although it is components in the oil that cause a reservoir rock to become oil wet it is a fact that some minerals are more prone to oil wetness than others. For example, carbonates tend to be more prone to oil wetness than clastics and amongst carbonates dolomites seem to be most prone to becoming oil wet. Of the silicate minerals it is typically clay minerals that promote oil wetness (chamosite, a member of the chlorite family, is often mentioned as promoting oil wetness).

Wettability is of great practical importance because it determines fluid distribution and the residual oil saturation (and hence recovery factor). Unfortunately it is very difficult to measure precisely.

Surface chemists measure a related property known as the contact angle (θ_w or 'theta'). This is defined as the angle a liquid droplet makes when it is in contact with a polished, planar solid surface. In order to be of any use the two fluids and the solid have to be specified. For example, 'the contact angle of water in air with glass is 30°' is a valid statement. The 'contact angle of mercury is 170°' is not.

If the edge of the droplet forms an angle of less than 90° with the mineral surface the liquid is said to wet the surface and is called the 'wetting phase'. The other fluid is by definition the non-wetting phase and it will have a contact angle of

$$\theta_{nw} = 180 - \theta_w \qquad (10.7)$$

A familiar example is water in air on glass: a window pane that has been exposed to rain for example. Water is normally the wetting phase and air is the non-wetting phase. It might be possible to replace air with a second fluid that adheres even more strongly to glass than water in which case water becomes the non-wetting phase. In air, the water will form droplets that make contact angles of less than 90° but in the presence of our – hypothetical – replacement for air, the contact angle for the water will increase to greater than 90°. Wetness is thus a relative concept.

The fundamental forces that determine the contact angle will be discussed in subsequent sections, when we consider interfacial tension (IFT). For now however, it is worth discussing the practical problems of trying to use contact angle as a measure of wettability. We noted earlier that the contact angle is defined for a pair of fluids in contact with a polished, planar solid surface. We are interested in naturally occurring porous rocks that are typically made up of:

1. Very small, irregularly shaped grains possibly with non-planar surfaces.
2. More than one mineral.

These make it impractical to extend the contact angle measurement to real rocks and hence to use it as a measure of wettability. Some laboratories will measure contact angles on samples of reservoir rocks by preparing a specimen with a highly polished surface. It is debatable how useful these measurements are, since the surface is a composite of different minerals. More informative is the environmental scanning electron microscope (ESEM) which allows wet specimens to be studied. This allows individual grains with associated water and oil droplets to be examined at high magnifications. Contact angles for a particular grain can then be measured (see Fig. 10.8 for an example).

Even the ESEM only allows the contact angle of individual grains to be measured and there is no obvious way to combine all these separate determinations into an objective measure of wettability for the whole specimen. To overcome these problems alternative ways of measuring wettability on complete core plugs were devised. The two most widely quoted are the AMOTT and USBM methods, both of which can be made by most commercial laboratories. Both methods subject a core plug initially saturated with water to a cycle of drainage – imbibition – drainage. In the Amott method, water saturation is measured at various points in the cycle and these are converted to a 'wettability index'. The USBM method is more complicated but again the output is a single number. In both methods a positive wettability index implies a water wet rock and a negative value an oil wet rock.
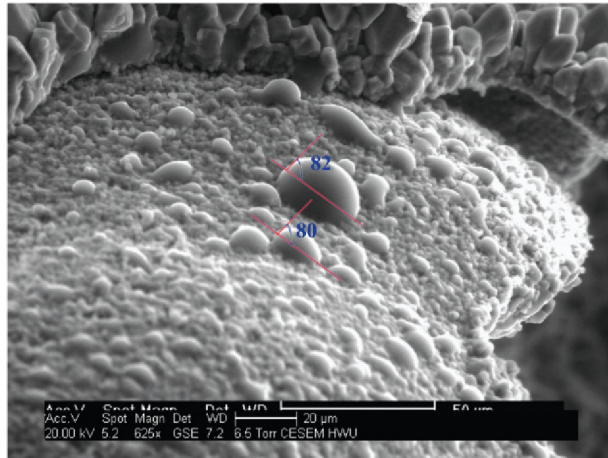
FIGURE 10.8 ESEM micrograph of the surface of an oil wet limestone. The droplets are oil droplets and the contact angle is less than 90° so the limestone is oil wet.

10.6   INTERFACIAL TENSION AND CAPILLARY PRESSURE

IFT is the property that quantifies the relative strengths of the intermolecular forces between the same and different types of molecule. It can be thought of as a measure of the 'strength' of the interface between two immiscible substances whether two liquids, a liquid and a solid, a liquid and a gas or a solid and a gas. The surface of water in air, for example is strong enough to allow certain insects to walk on it (and with skill it is possible to 'float' a steel pin on the water surface). If the water is replaced by a light oil, the unfortunate insects will sink and drown. In the reservoir we are concerned with three IFTs (rock–water, rock–hydrocarbon and water–hydrocarbon). In the laboratory, we often use air–brine or air–mercury as the fluids so different IFTs come into play.

IFTs are determined by the contrast in strength of the intermolecular forces between the same type and different types of molecules. They can be accurately measured by a variety of laboratory techniques and have dimensions of force per unit length. The SI unit is Newton per metre but, as is often the case, the oil industry tends to use a non-SI unit: dynes/cm (1 N/m = 1000 dyne/cm). Some representative values are given in next section. The conventional symbols for IFT are lower case 'lambda' λ used by chemists or 'sigma' σ more commonly used by petroleum engineers.

Water has a high IFT when in contact with air and simple hydrocarbons (light oil or gas) because, as discussed earlier, the water molecules form strong intermolecular forces with each other but hardly interact with the non-polar molecules in the air or hydrocarbon phase. Mercury has an extremely high IFT with air because, like all metals, the atoms that form it are very strongly bonded to each other. Replacing air with water does not change the IFT significantly

because the forces between the mercury atoms are so strong. Light hydrocarbons that make up light oils form relatively weak intermolecular forces so that even in contact with air they have a low IFT.

IFTs fall with increasing temperature and pressure and at the critical point they drop to zero. For miscible fluids, water and ethanol, for example the IFT is again zero this is equivalent to saying an interface cannot form. As soon as the two liquids come into contact they start to diffuse across the boundary and given sufficient time will completely mix. IFTs also depend on chemical composition. So, for example all other things being equal the IFT between water and gas increases with salinity. Conversely there is some evidence that increasing the CO_2 content and/or CGR of a gas decreases the IFT with water. Impurities can have a drastic effect on IFT generally reducing it, this means great care must be taken when making a measurement to exclude contaminants from samples and laboratory apparatus. It is generally better to make measurements on synthetic fluids than real reservoir fluids, which may be contaminated with mud filtrate.

10.6.1 Glass Tube Re-visited

IFT is important because it determines how tightly water bonds to the mineral grains or equivalently the capillary pressure. For a tube of radius r the capillary pressure is given by:

$$P_c = \frac{2\sigma\cos\theta}{r} \qquad (10.8a)$$

where we use σ for IFT –because this a petroleum engineering book – and θ is the contact angle. Remember the contact angle refers to the two fluids we are dealing with and the material that the tube is made out of. Change any of these and the contact angle will change.

Remember also from Section 10.3 that P_c is the pressure that has to be exerted to move the contact back to the water level in the glass or if you prefer the pressure needed to support the column of wetting liquid in the tube. This column's mass is given by Eq. 10.5 and it would exert a pressure given by Eq. 10.6b, this has to be balanced by the capillary forces so we can write:

$$P_c = \frac{2\sigma\cos\theta}{r} = H\rho g \qquad (10.8b)$$

At this stage we should note that real porous rocks do not consist of bundles of vertical cylindrical pores, standing in a more or less infinite reservoir of fluid. Nevertheless Eq. 10.8b gives some useful insights into how the fluid properties influence saturation. For a start we now know that in order to push the contact down in the tube we have to exceed a threshold pressure given by Eq. 10.8b. This is commonly observed in laboratory measurements on real rocks saturated with water. In this case it is found that no water is expelled until the applied pressure exceeds a minimum value: often named the entry pressure.

TABLE 10.2 Properties of Some Pure Fluids Discussed in the Text

Fluid	IFT (dynes/cm)	Density (g/cm^3)	IFT/density
Water	72.8	1.00	72.8
Ethanol	22.3	0.79	28.2
Hexane	18.4	0.65	28.3

Values at 20°C and atmospheric pressure.

Also note that the equation can be re-arranged to give IFT in terms of capillary rise, this is one of the ways that IFT is measured in the laboratory.

Capillary pressure increases with IFT and fluid density and decreases with contact angle. In practice these properties are not independent so that, for example replacing water with ethanol changes IFT, density and contact angle (assume the other fluid is air). If we assume the contact angle is the same in each case the capillary rise or capillary pressure is proportional to the ratio of IFT to density.

Inspection of Table 10.2 shows that the greatest rise occurs with water despite this being the densest liquid. It's very high IFT more than compensates for the higher density. Although we do not know the contact angle is unchanged, we do know that for all three fluids in contact with glass it is low. The contact angle enters as a cosine and so the effect of small changes tends to be masked. For mercury on the other hand the contact angle in air with glass is very large – nearly 180° – and cannot be ignored. The cosine is actually negative and the interface gets depressed inside a glass tube.

Explicit equations for capillary pressure analogous to Eq. 10.8b, can be derived for more complicated pore geometries but they always have the same general form and do not provide many more insights (the term $1/r$ is replaced by a more general curvature term). In any case, real rocks introduce so many complications that even these more complicated models do not get us much closer to reality.

10.7 CAPILLARY PRESSURE CURVES

The complications eluded to above include:

1. Different minerals with different contact angles.
2. Pores of different sizes and shapes.
3. Pores that are connected to each other.
4. Mineral grains with rough surfaces.

On the other hand there is at least one simplification which is that one of the fluids is always water and the other is either oil, gas or in the case of geo-sequestration CO_2.
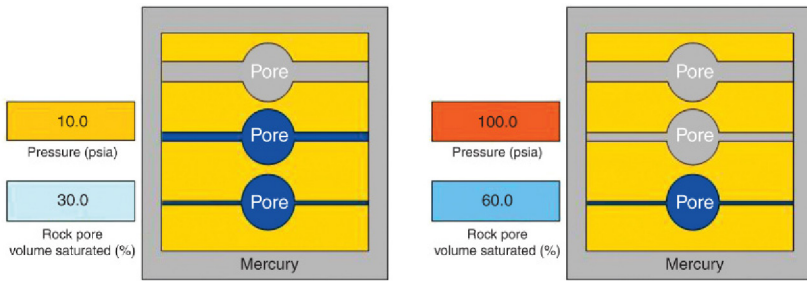
FIGURE 10.9 Schematic showing the general measurement of a mercury–air capillary pressure curve.

The effect of some of the complications can be determined at least qualitatively, but for most we just accept they will cause deviations from the ideal capillary tube model. In order to obtain quantitative information we must resort to measurements on samples from the reservoir. These are routinely made as part of a special core analysis program (SCAL) and although a variety of methods are available they all work by saturating the core with a wetting fluid and then forcing a non-wetting fluid into the pore space by applying external pressure. The pressure is increased in increments so that a curve of wetting phase saturation versus pressure can be constructed. This is the capillary pressure curve.

The simplest method is simply known as 'air–mercury'. Air is the wetting fluid by virtue of the fact that mercury bonds to itself far more strongly than it does to any mineral that we are likely to encounter. The curve is constructed by immersing a cleaned and dried sample of rock in mercury and then simply increasing the pressure on the mercury reservoir (Fig. 10.9). It has the advantage that it can be used with small and/or irregularly shaped samples (some operators routinely run the technique on cuttings). The volume of the sample is found by the volume of mercury displaced when the sample is first immersed.

The technique also allows very high capillary pressures to be reached and because it is simple and quick, a lot of data can be obtained in a relatively short time.

The disadvantages are the technique that is destructive and mercury has completely different physical and chemical properties to any reservoir fluid. In particular as it is forced into the pores it can 'flatten' the delicate clays, which often play a major role in controlling water saturation. Moreover air is the wetting phase but, unlike water, it has no particular affinity for commonly occurring minerals.

The Capillary Pressure Curve produced by the idealised system shown in Fig. 10.9 is shown in Fig. 10.10. Note that for the system shown in Fig. 10.10 the saturation would change quite suddenly when the threshold pressure for each tube is reached but in practice the data would probably be interpreted as giving a smooth change. In any case real laboratory measurements use more pressure steps.
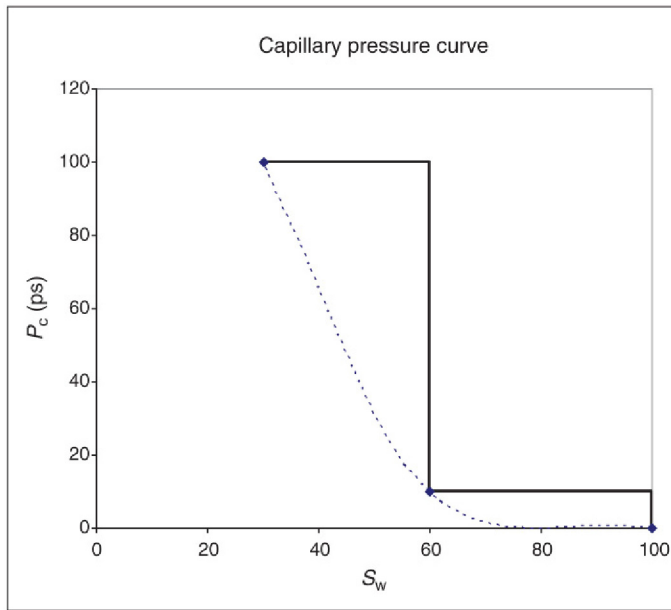
FIGURE 10.10   The capillary pressure curve that would result from the experiment illustrated in Fig. 10.9. The true curve for the model is in black, the dotted curve shows the curve that would probably be presented by a laboratory.

The other techniques make use of fluids that are closer to those found in the reservoir if not actual reservoir fluids. In particular they use water – brine – as the wetting fluid. Non-wetting fluids include air, kerosene, decane, dead crude and live crude (decane is a straight chain hydrocarbon that can be used as a model for light crudes). Two basic methods are used to generate the capillary pressure curve: porous plate and centrifuge. These both use cylindrical core plugs and are generally more involved than air–mercury so in practice the number of samples that can be processed is limited. Consequently a lot of effort should go into plug selection and QC to ensure they are as representative as possible. Furthermore, far fewer steps are used to generate the curve (five or six).

In the porous plate technique, the plug is initially saturated in brine and one end of the plug is covered by a special membrane – the porous plate. The membrane only allows the wetting phase, which is normally water, to pass through. The whole assembly is then immersed in the non-wetting fluid within a special holder. As pressure is applied to the non-wetting fluid it enters the pores and expels the wetting fluid through the membrane. The volume expelled allows the saturation to be calculated at any pressure.

The pressure is increased in increments and should be held steady until no further change is observed. Most of the change occurs soon after the pressure is increased but it can take months for equilibrium to be attained even for plugs

with permeabilities of 10 s of milliDarcy. This makes it very difficult to know when the equilibrium saturation has been attained and so there is a suspicion that curves generated using a porous plate overestimate the water saturation at any particular pressure.

In recent years an alternative method using a centrifuge has become popular, this allows curves to be constructed more quickly but there is a limit to how high the pressure can be taken and the physics of the measurement mean the pressure actually varies along the length of the plug. An example of some typical, real data and the curve they define are shown in Fig. 10.11. As can be seen there are seven pressure increments that roughly double at each step.

(a)

Determined at 1160 psi NOBP

Sample no.	Depth (m)	K_{inf} (md)	K_{air} (md)	Porosity (%)	P_c (psi)	S_w (%)
RSCT 5	1144.00	1.15	1.83	25.5	0	100.0
					0.5	100.0
					1	100.0
					2	95.4
					4	88.9
					8	82.8
					14	78.2
					22	74.6

(b)
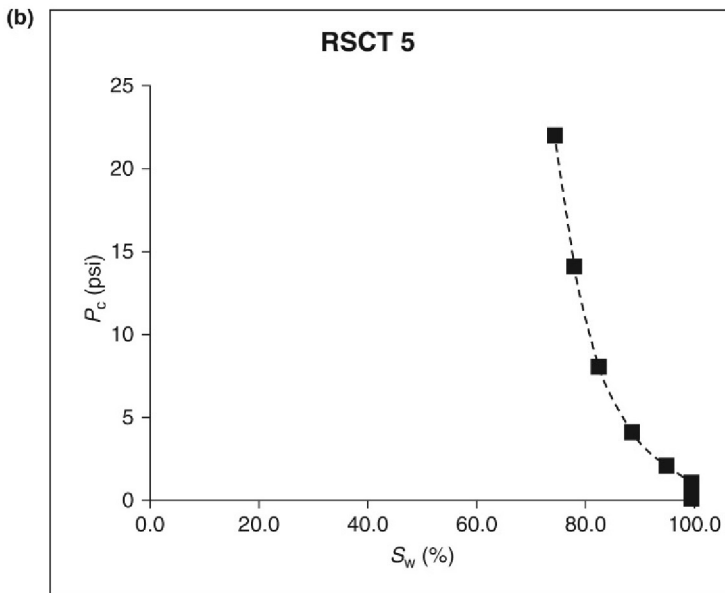
FIGURE 10.11 A real example of an oil–brine capillary pressure curve measured by a commercial core laboratory. (a) Table of the data as provided in the final report. (b) The resulting curve. Note the generally high water saturations and modest permeability.

10.7.1 Converting from Laboratory to Reservoir

Regardless of which technique is used the fluids are unlikely to be exactly the same as those in the reservoir. Laboratory experiments are normally made at ambient temperatures and pressures with fairly bland fluids (water, air/nitrogen, dead-oil or kerosene), which will have different IFTs and contact angles to reservoir fluids at reservoir conditions. Equation 10.8b allows experiments made in a laboratory to be converted to reservoir conditions. If we wish to convert from a laboratory measurement to reservoir conditions the capillary pressure changes by the factor:

$$\frac{P_c(\text{RES})}{P_c(\text{LAB})} = \frac{\sigma(\text{RES})\cos\theta(\text{RES})}{\sigma(\text{LAB})\cos\theta(\text{LAB})} \tag{10.9}$$

For example, if the IFT of water and air at ambient conditions is 72 dynes/cm and the IFT for water and a wet gas at 5000 psi and 100°C is 25 dynes/cm and assuming the contact angle does not change significantly, the capillary pressure in the reservoir is about one third the value in the laboratory. This means:

1. The capillary rise would only be one third that seen in the laboratory.
2. Gas will enter a particular pore size at one third the height above FWL than it would in the laboratory.

(in short, the IFT changes help us squeeze more gas into the structure).

10.8 PUTTING IT ALL TOGETHER: REAL ROCKS AND REAL FLUIDS

Real rocks have a continuous range of pore and pore throat sizes so the saturation varies continuously with increasing capillary pressure. As will be seen later a lot of information on reservoir quality can be deduced from the shape of a capillary pressure curve. Figure 10.12 shows some hypothetical curves. These are drainage curves in other words the core plugs have started out saturated with the wetting fluid and the capillary pressure has gradually been increased. In this case the wetting phase is water (the axis is labelled S_w) but the same conclusions could be drawn from air–mercury curves (or any other pair of fluids). There are three features of the curves that can be used to gain quantitative and qualitative information about the rock.

1. The entry pressure (the minimum pressure that has to be applied to start displacing wetting fluid).
2. The lowest wetting phase saturation that can be achieved at high capillary pressures.
3. The shape of the curve between these extremes.

If the wetting phase is water the minimum saturation that can be achieved is the irreducible water saturation. These components are marked on the saturation-height function as shown in Fig. 10.13.
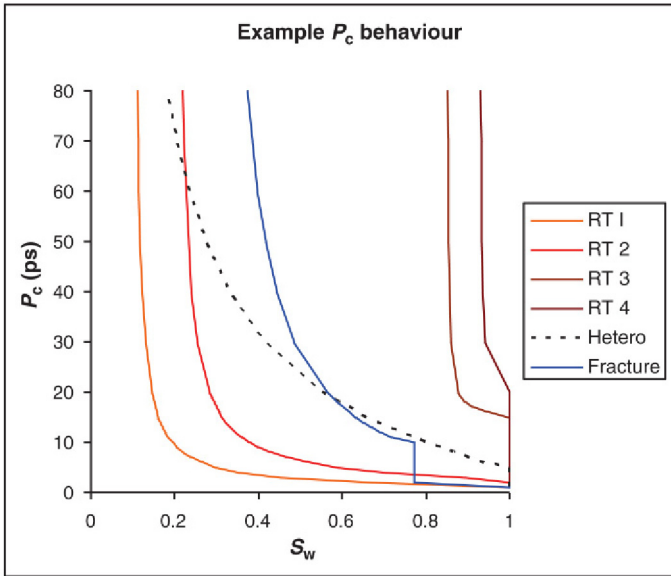
FIGURE 10.12   Some hypothetical capillary pressure curves (see text for explanation).
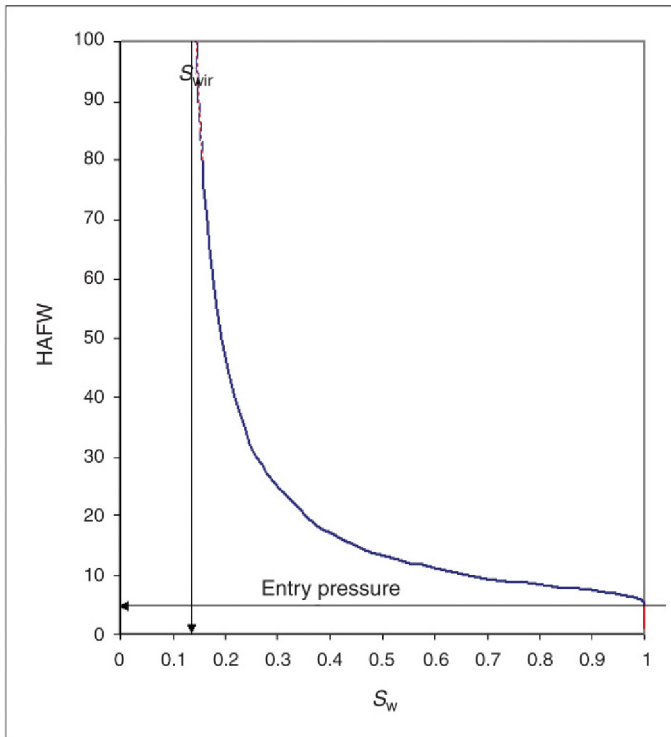
FIGURE 10.13   A general saturation height function showing how the terms that are used to describe it are defined.

The curves labelled RT1 and RT2 represent sandstones or possibly carbonate rocks with a fairly narrow range of pore and pore throat sizes. In other words they are well sorted. More specifically:

RT1 is likely to have a high permeability (>100 mD) and to be relatively coarse grained. The latter can be deduced from the low entry pressure: it is relatively easy to force the non-wetting phase into the pore throats.

RT2 is still likely to have good permeability (tens of milliDarcy) but is significantly finer grained than RT1. This can be deduced from the higher entry pressure and the higher irreducible water saturation.

RT3 and RT4 represent further declines in reservoir quality and RT4 is so poor that in some circumstances it would be considered a seal rather than a reservoir rock. Nevertheless they still seem to have relatively uniform pore size distributions. This can be seen from the short range of pressures between the entry pressure and the pressure at which S_{wir} is reached. RT3 is probably either a siltstone or a heavily cemented sandstone with permeability below 0.1 mD. RT4 is probably a mudrock with a permeability of the order of a microDarcy or less.

The remaining curves belong to more complicated rocks. The curve labelled 'hetero' belongs to a rock with a wide range of pore and pore throat sizes. It has a fairly high entry pressure suggesting the largest pore throats are a lot smaller than those seen in RT2, but they are certainly larger than those in RT3. The slow decline in S_w with increasing pressure shows there is a wide range of pore sizes and the rock is therefore likely to be poorly sorted or generally quite heterogeneous.

Finally, the curve labelled 'fracture' has a dual porosity system in which approximately one third of the porosity has a low entry pressure and the remainder consists of a wide range of small pores/pore throats. Most likely the low entry pressure is caused by microfractures, which can be almost completely drained as soon as any pressure is applied. Once they are drained the pressure has to be increased until the pore throats in the conventional 'matrix' porosity can be accessed. These have a fairly high entry pressure and so are probably formed in a very fine-grained background.

So far we have looked at qualitative information that can be obtained from capillary pressure curves but in fact it is possible to obtain quantitative information on pore throat sizes from the entry pressure. The formula used is Eq. 10.8:

$$P_c = \frac{2\sigma\cos\theta}{r} \tag{10.8a}$$

which can be re-arranged to give

$$r = \frac{2\sigma\cos\theta}{P_c} \tag{10.8b}$$

where all the measurements are in SI units. In oilfield units – dyne/cm and psi – the equation becomes

$$r = 0.29 \frac{\sigma \cos \theta}{P_c} \qquad (10.8c)$$

Where r is in micrometres. An entry pressure of 1 psi in an air–brine capillary pressure measurement therefore corresponds to a pore throat radius of 20 μm. Equation 10.8 allows any capillary pressure to be converted to an equivalent pore throat size and so a capillary pressure curve can be converted to a pore throat size distribution. This is routinely done with air–mercury data where there are a large number of individual measurements. An example is shown in Fig. 10.14 where air–mercury capillary pressure curves and the derived pore size distribution for two samples from a sandstone reservoir are shown.
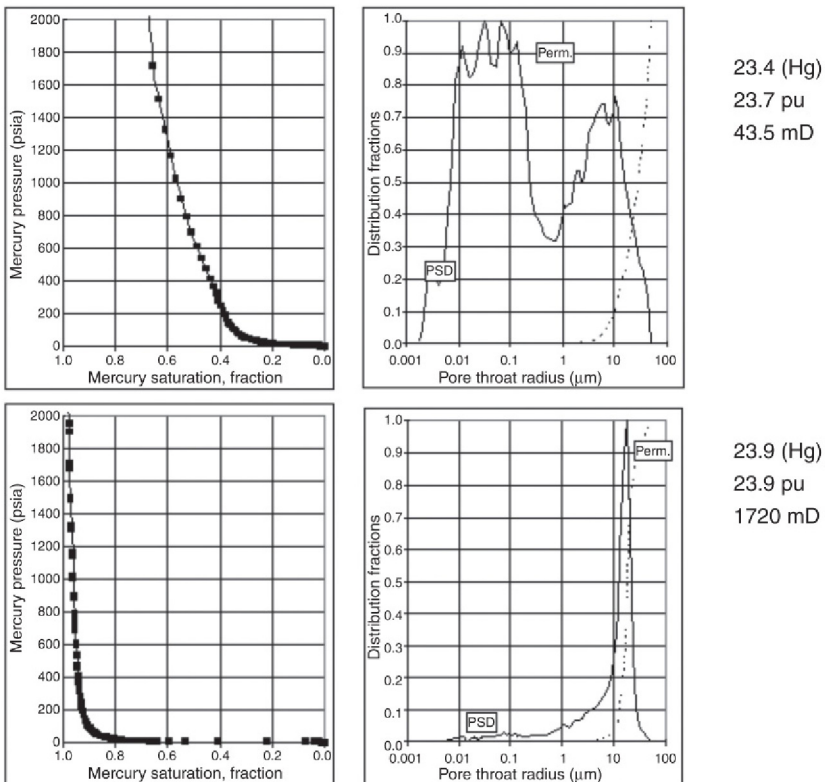
FIGURE 10.14   Air–mercury capillary pressure curves for two contrasting samples from the same reservoir. The plots show the relationship between the shape of the P_c curve and the pore throat size distribution. Note the agreement between the measured porosity of the nearby plug used for conventional core analysis and the porosity inferred from the amount of mercury injected into the sample.

The upper sample is poorly sorted with a wide range of pore throat sizes whilst the lower one has a narrow range of pore throat sizes centred on 20 μm. Notice that although porosities are similar the poorly sorted sample has a measured permeability that is approximately 40 times lower than the well-sorted sand.

10.9 DEVELOPING A SATURATION-HEIGHT FUNCTION

10.9.1 Introduction

All the proceeding theory is needed to predict the saturation at any point in the reservoir. In practice this involves writing an equation that relates saturation to height above FWL, rock properties and possibly fluid properties (although the particular fluids have a profound effect on saturation this may be taken care of by writing an equation which is specific to a particular reservoir). This is the saturation-height function that may be derived from capillary pressure curves and/or the results of log analysis. In either case the saturation-height function is derived by fitting a curve to the data. Any equation will do providing it reproduces the way water saturation drops with height above the FWL. This means the curve should:

1. Give a water saturation of 100% below the FWL.
2. Give a continuous decrease in water saturation with height above the FWL.
3. Asymptotically converge on the irreducible water saturation at large heights above FWL.
4. Possibly give water saturations of 100% for a distance above the FWL corresponding to the entry pressure.

Numerous functions have been proposed that have these general properties. The best one to use is really a matter of how well they match the field data and possibly how convenient they are to implement. The simplest will typically involve negative exponentials and/or reciprocals both of which rise asymptotically with falling S_w.

No matter what form is actually employed the aim is simply to fit a curve(s). This will give a curve of general form:

$$S_w = S_w (P_c, \text{rock properties}) \tag{10.10}$$

To apply this to the reservoir P_c needs to be converted to height above FWL.

$$\text{HAFWL} = P_c /(1.42\Delta\rho) \tag{10.11}$$

Where HAFWL is in metres, P_c in pounds per square inch and the fluid density term is in gram per cubic centimetres. P_c must be converted from laboratory to reservoir conditions using Eq. 10.9. The density difference $\Delta\rho$ refers to fluids at reservoir conditions (if other units are used the constant 1.42 will change).

10.9.2 Saturation-Height Functions Based on Capillary Pressure Curves

Capillary pressure curves are generally the best starting point for deriving a saturation-height function because they cover the full range of saturations ($S_w = 1$ to S_{wir}). To get the same information from logs requires a well or wells that have intersected a contact, a substantial part of the transition zone and part of the column above the transition zone. Even with large fields intersecting a contact can be largely a matter of luck.

So the pragmatic way of dealing with the complications introduced by real rocks is to measure capillary pressure curves on core plugs and then fit a function to the results. This is saturation-height modelling. The individual steps are:

1. Fit a function to the laboratory capillary pressure curve(s).
2. Convert the capillary pressure measured in the laboratory to reservoir conditions using Eq. 10.9.
3. Convert the reservoir capillary pressure to HAFWL using Eq. 10.11.

Here these steps will be implemented in the order they are given earlier but one could equally well convert laboratory capillary pressure to HAFWL first and then fit a curve (i.e. perform the steps in order 2, 3, 1).

As already noted capillary pressure curves tend to have the following characteristics:

1. The pressure has to reach a finite value before the non-wetting phase enters the pores (this is the entry pressure).
2. Above a certain pressure the wetting phase saturation cannot be reduced by further increases in pressure.
3. Between these two regions the wetting phase saturation falls continuously with increases in pressure.

The saturation-height function needs to reproduce this behaviour. There are an infinite number of functions that will achieve this and moreover provided one is content to accept a lot of complexity, an experimental capillary pressure curve can be matched exactly. In practice, however it is better to use relatively simple functions and accept a small amount of mismatch. If nothing else it must be remembered that ultimately we are trying to model saturation throughout a large structure using the experimental curves measured on a handful of plugs with volumes of a few cubic centimetres. No matter how carefully the plugs are selected they will never represent every point in the reservoir.

Examples of general functions that will capture the general shape of the simpler capillary pressure curves include:

$$S_w = g(\exp(f(P_c))) \tag{10.12}$$

Or

$$S_w = A/f(P_c) \tag{10.13}$$

where the functions g and $f(P_c)$ are simple functions whose precise form can be varied to achieve a match with experimental data. Some commonly used functions that are specific examples of Eq. 10.13 are:

$$S_w = AP_c^{-\lambda} + B \qquad\qquad \text{The Lambda function} \qquad (10.14)$$

$$S_w = 1 - A\exp(-\lambda/(\log(P_c) + B)) \qquad \text{Thomeer function} \qquad (10.15)$$

And a more recent function that is an example of Eq. 10.12 is

$$S_w = 1 - A\exp(B/(P_c + D)) \qquad \text{Harrison–Skelt function} \qquad (10.16)$$

An even simpler function which is an example of Eq. 10.12 can work well.

$$\text{Log}(S_w) = A + BP_c \qquad\qquad (10.17)$$

(as far as I know this does not have a special name). The parameters A, B, D and λ (lambda) are what are varied to match the experimental data. The three functions above have been fitted to some real data in Fig. 10.15. It is left to the reader to decide which gives the best match (hint: they are all equally good).

Because P_c is a function of HAFWL any of these can and often are, written in terms of height. Furthermore most of them can be written in ways that look quite different but actually are just rearrangements, although it is also true that slightly modified versions are sometimes given in the literature (e.g. Harrison–Skelt, as published, employs four constants).

It is worth pointing out that these equations although fairly simple are also quite clumsy and it is very easy for typographic errors to creep into reports! It is also important to know whether the equation is intended to give saturation as a percentage or a fraction. The parameters A, B etc. will be quite different in each case and in general they will not simply be a factor of 100 different!

In general any of the equations above will only model one capillary pressure curve. Different plugs will produce different curves and ideally there will be a systematic change with a property that can be estimated from logs. This is what is shown with curves RT1 to RT4 in Fig. 10.13. So we will most likely want to develop a family of curves that can take into account variations in reservoir quality (Eq. 10.10a).

More often than not the quantity that is used to describe reservoir quality is permeability but there is no fundamental reason for this and porosity or shale volume or some combination, may work just as well. Given that permeability will normally have to be predicted from porosity there is a lot to be said for just using porosity.

In any case the problem is actually starting to become quite complicated. If for example we decide to use Eq. 10.18 and follow the common practice of using permeability as our other variable. The steps are as follows:
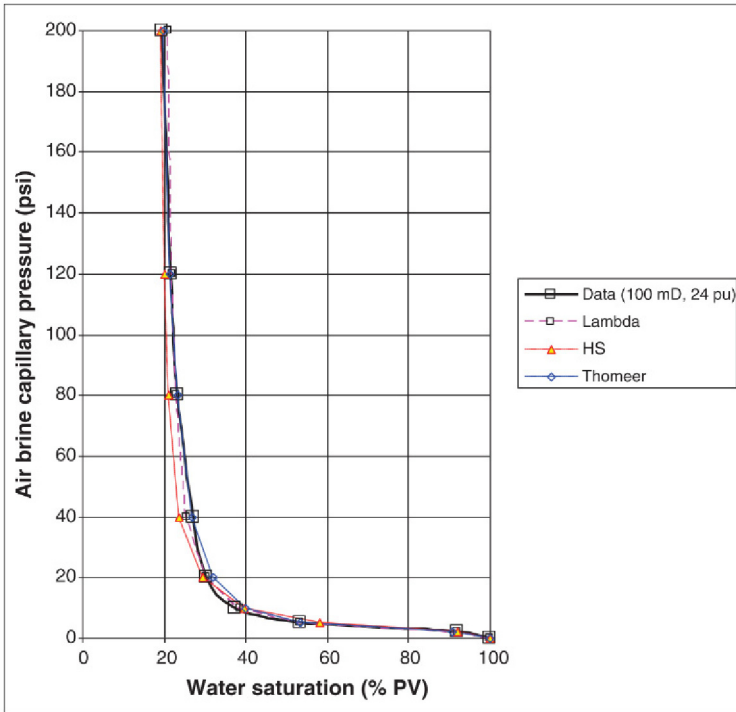
FIGURE 10.15   Curve fits for the Thomeer, Lambda and Harrison–Skelt function to some real air–brine capillary pressure data. Realistically, all three functions do an equally good job! The precise equations are:

Lambda $S_w = 19.48 + 128.4 P_c{}^{-0.83}$,

HS $S_w = 100 - 81.89 \exp(-2.82/(P_c - 0.78))$ and

Thomeer $S_w = 100 - 910.68 \exp((-0.43/(\log(P_c) - 0.12)))$.

1. Fit each experimental capillary pressure curve with Eq. 10.18 by varying A, B.
2. Plot A and B against the plug permeability and attempt to find a simple relationship between them.

   For example, $A(k) = A1 + A2.\log(k)$ $B(k) = B1 + B2.\log(k)$
   Our full equation to predict S_w from P_c and k is then:

$$\text{Log}(S_w) = A1 + A2.\log(k) + \{B1 + B2.\log(k)\} P_c \qquad (10.18)$$

And we still need to convert P_c to height!

Unfortunately, Eq. 10.18 is only likely to apply everywhere in a small structure with little variation in the nature of the reservoir rock. We may well have to derive separate functions for different formations, facies, fault blocks etc.

10.9.3 Other Approaches to Saturation-Height Functions

If core data is not available the only alternative is to try to match the saturation curve calculated from logs. To do this requires a more or less complete saturation curve to be available. This might require obtaining different parts of the curve from different wells but we certainly need an estimate of the irreducible water saturation and the variation of saturation above the contact to proceed.

If this information is available Cuddy proposed plotting saturation against height above contact on a log–log grid. The relationship is thus:

$$\text{Log}(S_w) = A\text{log}(\text{HAFWL}) \qquad (10.19)$$

In fact Cuddy went further and used the plot to estimate the FWL. Locating the FWL is discussed further in subsequent sections.

10.9.4 Leverett *J*-Function

Although this section is not intended to be an exhaustive treatment of every saturation-function that is available, one specific function should be mentioned if only because it is still the first choice for many petroleum engineers. This is the Leverett *J*-function, named after it's inventor M.C. Leverett who introduced it in a paper published in 1941 (a year before Archie published his results). It is not actually a saturation-height function at all, but is a transform that is applied to capillary pressure.

$$J = CP_c.\sqrt{(k/\Phi)} \qquad (10.20)$$

where k and Φ are the permeability and porosity of the plug and C is a constant that depends on the units in use. All the transform appears to do is to stretch or squeeze the capillary pressure curve by the term in brackets. Note that because permeability has the dimensions of area, the term in brackets has the dimensions of length (porosity is dimensionless of course). For those who are interested, a dimensional analysis shows that J/C has the dimensions $[\text{M}][\text{T}]^{-2}$ (in SI units, for example this would be kg/s^2). In fact, C is chosen to make J dimensionless, so it has dimensions $[\text{M}]^{-1}[\text{T}]^2$. In practice nobody really bothers to calculate C and it is often not even written in the equation, i.e. it is implicitly set to 1. This is permissible because the real point of the exercise is to apply Eq. 10.20 to all the capillary pressure curves and then plot the *J* curves on the same graph. In an ideal world all the different P_c curves lie on a single curve when transformed to *J*. (I have to confess I have never come across a single case where this happens, but I have seen plenty of cases where a large number of P_c curves end up 'collapsing' to three or four *J* curves). A real example is shown in Fig. 10.16. In that particular case the transformation to *J* does not produce any consolidation despite the fact that all the data comes from within a single 50-m thick sand.

It is important to realise that Leverett chose the particular form of his transform to make *J* dimensionless. In particular, the square root term gives a length
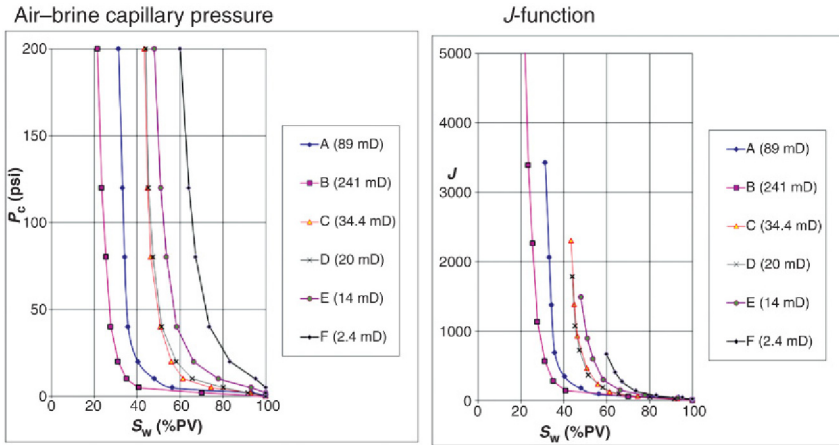
FIGURE 10.16 Air–brine capillary pressure curves from a shaly-sand reservoir. The measured curves are shown on the left and the same curves with P_c transformed to J are shown at right.

dimension when applied to permeability. Some engineers try to get a single curve by using some other power of permeability. They may succeed, but at the cost of losing the rigour imposed by Leverett.

Since P_c depends on the fluids used to make the measurements Eq. 10.20 is often written with an IFT and contact angle term as well. This essentially allows Eq. 10.9 to be included within the expression for J so that it applies at reservoir conditions without any further manipulation.

Equation 10.20 is not a saturation-height equation of course and in order to find saturation we have to fit a curve in just the same way that previously curves were fitted to $S_w(P_c)$. In fact, exactly the same form of equations are often used. Assuming this has been done the work-flow would be:

1. Convert height to P_c using Eq. 10.11 (and 10.9).
2. Find permeability and porosity at each point and convert P_c to J.
3. Find the saturation from $S_w = S_w(J)$.

10.10   THE FREE WATER LEVEL AND FORMATION TESTERS

In Section 10.2 the FWL was defined as the depth that the buoyancy force dropped to zero. In the simple 'glass of water' system discussed there this was the surface of the water. In that particular case there are in effect two different ways of defining the FWL:

1. A hydrostatic definition based on pressure gradients.
2. An analytical chemistry definition based on the presence or absence of water.

With real reservoir rocks these two definitions each lead to a different depth. The analytical chemistry definition is actually referred to as the GWC or OWC
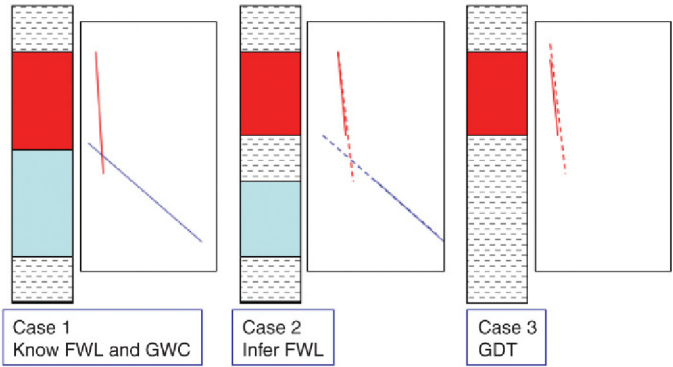
FIGURE 10.17 Diagram illustrating the conditions for determining the GWC and the FWL. The FWL can be inferred from pressure gradients in the LH and middle scenarios. The contact can only be determined in the LH case where it lies within a sandstone. In the RH case it is not even possible to determine the FWL because no pressure measurements are available from a water-bearing formation.

if there is an oil leg. The hydrostatic definition gives the FWL. For water wet rocks the FWL is normally deeper than the GWC/OWC although the difference is often too small to detect.

The GWC has to be located using log analysis and it is by definition the depth below which the water saturation is constant at 100%. It can only be located if it is penetrated by a well. If a reservoir interval is completely gas bearing the best that can be said is that the contact is below the base of the sand (which then defines the 'gas down to' or GDT).

The FWL on the other hand can often be located quite accurately in a well even, if it lies in a non-reservoir lithology (e.g. shale, Fig. 10.17). The only requirement is that there is gas- and water-bearing rock of sufficient quality to measure accurate formation pressures. The technology to do this first appeared in the early 1970s with the advent of Schlumberger's repeat formation tester (RFT). It is no exaggeration to say this revolutionised reservoir engineering and it did not take long for Schlumberger's competitors to introduce their own versions. It should be noted that formation testers had been in existence long before the RFT appeared, but these older tools could only make one or two pressure measurements per trip in the hole and often required a specially trained crew to operate them.

The heart of any formation tester is a pressure gauge, the purpose of the tool is quite simply to put this in pressure communication with the formation at a well-defined depth and to isolate it from the mud. The RFT as its name implied, could in theory do this an unlimited number of times and so detailed plots of formation pressure against depth could be made.

A diagram of the essential components of a formation tester tool is shown in Fig. 10.18. Most modern formation testers are more complicated
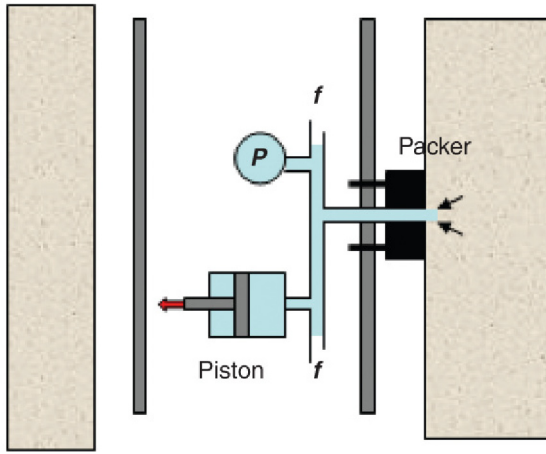
FIGURE 10.18   Wireline formation tester principles. The packer is pushed firmly against the borehole wall thereby isolating the flow line (*f*) from the mud. The piston then moves back, drawing a few cubic centimetres of fluid from the formation into the flow-line. The pressure is then allowed to build-up to formation pressure. The pressure in the flow-line is measured continuously by the pressure gauge (*P*).

than this because they are also designed to take samples of formation fluid. Formation pressure is measured by pushing a packer against the borehole wall, this isolates the formation from the mud. The centre of the packer has a hole that is connected to a flow-line within the tool, this in turn is connected to the pressure gauge and one or more pistons. The sequence of events is as follows:

1. The gauge is initially connected to the well bore and measures the pressure of the mud column at the tool depth. This is known as the 'hydrostatic pressure'.
2. The packer is pushed against the borehole wall, trapping fluid in the flow-line at hydrostatic pressure (if anything the pressure increases slightly as it is compressed by the packer).
3. A small piston moves backwards withdrawing a few cubic centimetres of fluid from the formation. The pressure drops.
4. The piston stops and the pressure is allowed to build-up. Given sufficient time it will build up to the formation pressure.
5. The packer is pulled off the borehole wall and the pressure returns to hydrostatic.

The formation tester therefore performs a miniature well test. The duration is determined by the time it takes for the pressure to build up to formation pressure in step 4. This in turn is mainly determined by the permeability of the formation and typically takes a few minutes.

By taking measurements at a number of depths in a well, a plot of pressure against true vertical depth can be produced. Under ideal conditions these plots can provide the following information:

1. The pressure regime.
2. The *in situ* density of all the different reservoir fluids.
3. FWL.
4. *Location of the gas oil contact if one exists.*

The pressure regime simply means the reservoir pressure at a datum depth, it is extremely important in reservoir engineering but we will not discuss it here. The *in situ* density of the fluids is given by the gradient of the plot. This is just applying Eq. 10.2. In oil-field units it becomes:

$$\rho = \frac{\Delta P}{1.42 \Delta h} \qquad (10.21)$$

where density is in gram per cubic centimetre, depth (h) is in metres and pressure (P) is in pounds per square inch (the constant 1.42 takes care of the units conversion and includes g the acceleration due to gravity). In other words the pressure gradient is directly proportional to the *in situ* fluid density.

If pressures can be obtained in the water and hydrocarbon legs they will have different gradients and the depth where these intersect is the FWL. In high permeability formations with a sufficient depth range formation tester data is the most accurate and precise data to come from wells. *In situ* densities can be calculated to three significant figures and the FWL can be measured to a precision of 10 cm (the accuracy also depends on how accurate the depth measurements are of course).

In practice life may not be so convenient. For a start a water-bearing formation is required (see Fig. 10.17). This may only be available in a bed that is much deeper than the hydrocarbon-bearing formation (or of a completely different age). There will then be some doubt as to whether the water-bearing formation is actually in communication with the hydrocarbon. Ultimately the only way to solve this is to drill an appraisal well that is designed to penetrate the down dip water leg of the reservoir.

There are also 'experimental' uncertainties to deal with, particularly in low permeability formations (<10 mD). In these formations the build-up may be so slow that formation pressure is not reached in a sensible timeframe. The measured pressure then underestimates the true formation pressure. It is also possible for the measured pressures to overestimate the true value. This occurs because of a phenomenon known as 'super-charging' in which the formation in the near well bore region is 'pressured-up' by invasion. The formation tester is in communication with this over-pressured region and thus measures it's pressure. The typical symptom of super-charging and/or slow build-ups is a lot of scatter in the pressure depth plots. If this is found one has to accept that the FWL has a

lot of uncertainty. It is not sufficient just to put the best straight line through the pressure points and assume this is close to the true value. As has been seen above, in low permeabilities some or all the pressure measurements may be too low, or too high or a mixture of the two. Whatever the case the offset from the true pressure will not be constant and a lot of manual 'validation' of points will be needed.

The latest tools have a few features that may help reduce the instance of these problems or at least allow them to be recognised. But there will always be formations where pressure data is subject to significant uncertainty and in the tightest formations (<1 mD) formation testers are unlikely to work at all.

10.11   CONCLUSIONS

We now have everything needed to derive and use a saturation height function. The work-flow suggested here is:

1. Fit a suitable equation to each capillary pressure curve.
2. Look for a systematic change in the constants with a petrophysical property such as permeability.
3. Implement the resulting function $S_w = S_w(P_c, \text{property})$.
4. Locate the FWL
5. Convert height above FWL to a continuous P_c curve (using Eq. 10.12, fluid densities and IFTs).
6. Substitute P_c and the appropriate property curve into $S_w(P_c, \text{property})$ to obtain a prediction of S_w along the well track.
7. Compare the predicted S_w with the value obtained from conventional log analysis.

As noted previously, it is equally acceptable to convert from P_c to height above free water in the reservoir at a much earlier stage so that S_w is given as a function of height and – say – permeability.

Compare this with the Leverett J-function work-flow given in Section 10.9.4.

1. Convert all the P_c curves to J-curves using Eq. 10.20.
2. Plot all the J-curves together and if they all, more or less lie on the same line, fit a suitable equation to the line.
3. Locate the FWL.
4. Find porosity and permeability at each height above FWL and calculate the J value.
5. Convert J to S_w using the equation found in step 2.

If – and this is a big if – a single function $S_w(J)$ can be defined then this does involve less computation. Nowadays, this is at best a marginal advantage but in the 1940s when the function was introduced to the world it would have saved a lot of time.

Page left intentionally blank

Chapter 11

Permeability Re-visited

Chapter Outline

11.1 INTRODUCTION

In Chapter 1 a throw away remark was made that petrophysicists are primarily employed to find out how much fluid a rock can hold, how much of that is water and how quickly it can be extracted (basically porosity, water saturation and permeability, respectively).

Chapters 7, 8 and 9 have discussed how curves that show how porosity and water saturation vary along a well, can be obtained from well logs. This chapter looks at how permeability curves can be generated and in particular how logs can be transformed to permeability curves. There are fundamental and practical reasons why this is a much more ambitious task.

Permeability curves have a number of applications including the following:

1. Populating static and dynamic reservoir models.
2. An input to saturation height equations.
3. Defining net (pay) through a cut-off.
4. Quantitatively defining heterogeneity.
5. Predicting well performance.
6. Operational work such as defining perforation programmes.

Logs do a good job of estimating porosity because the physical properties they measure are determined by the relative amounts of fluids and solid minerals that make up the rock. Similarly, certain measurements are very sensitive to the amount of water in the formation and so can be exploited to accurately determine water saturation. So, providing there is a reasonably comprehensive

set of logs, which are working properly, log analysis can normally be relied on to give accurate estimates of porosity and saturation.

Permeability on the other hand is determined by the structure of the pore system and this is not something that most logs respond strongly to, if at all. The consequence of this is that there are no general tools and equations that can be relied on to generate an accurate permeability curve ('relied-on' is the operative phrase, there are several general equations that have been proposed over the years). It also means that to have any confidence in a permeability curve, some form of raw permeability data has to be available to ground truth it. In fact the normal way of constructing a continuous permeability curve is to base it on routine core measurements and typically one would look for some correlation between one or more raw or computed logs and the core permeability. The correlation will never be as strong as between logs and static properties like porosity, but unless there is complete core coverage of the reservoir, we rely on logs to at least fill the un-cored intervals.

11.2   CHARACTERISTICS OF PERMEABILITY

As explained in Chapter 2 permeability has a number of characteristics that mean its relationship to physical properties – logs – is normally a lot looser than is the case for porosity.

These characteristics are as follows:

1. Permeability is not dimensionless (it has the dimensions of an area). The static properties porosity, saturation and V_{shale} are dimensionless.
2. Permeability has an unlimited range (in practice permeability ranges from 0 to tens of Darcy). Saturation and shale volume are limited to values between 0 and 1 and porosity rarely falls outside the range 0–0.35.
3. As discussed earlier, permeability is controlled by the morphology of the pore system (the sizes, shapes of the individual pores and their interconnections). By contrast most logs are only weakly influenced if at all, by the pore system (the most notable exception is NMR but even then it responds primarily to the pores rather than the complete system).
4. Permeability can be highly scale dependent. So a volumetrically insignificant part of the reservoir may dominate the movement of the fluids.
5. Permeability depends on the fluid(s) in the pore system.
6. Permeability can change proportionately more than porosity with the pressure of the fluid and the stress on the rock.
7. In general permeability is a directional property: it varies depending on the direction it is measured in.

Broadly speaking the first three items mean that permeability is rarely well correlated with the physical properties that are measured by logs. Points four to six refer to exactly what the permeability curve is trying to model. As the starting point for a permeability model is often routine core data, the permeability

curve implicitly refers to absolute gas permeabilities with the gas at relatively low pressure. In order to simulate an effective permeability in the reservoir, the curve will need corrections to account for water (relative permeability), the higher pressure of the fluid (Klinkenberg factor) and possibly the lithostatic stress on the rock (overburden correction). It is important that the user of the curve understands what exactly it refers to and therefore what further corrections need to be made for it to describe what goes on in the reservoir. In this book, the curve will describe a single-phase gas permeability, in other words the absolute permeability measured by routine core analysis. This is simply to avoid endlessly re-stating what the curve is referring to, in reality the end user may well expect the petrophysicist to produce an effective permeability curve.

11.3   PERMEABILITY DATA

True permeability data is obtained from core and well tests (including wireline formation testers). Both types of measurement involve moving fluids through the pore system and so directly respond to permeability. Continuous log measurements such as NMR and 'Stonely Wave' measure static properties of the formation, which may or may not correlate well with permeability. Inferring invasion by mud filtrate from resistivity profiles and/or comparing LWD and wireline logs tells us that the formation is permeable but gives no quantitative information. The same is true of an SP deflection. Furthermore a lack of invasion or an SP deflection is no guarantee that the formation is impermeable.

   Core data is particularly good for generating a permeability curve because the data density is similar to logs (approximately 10 measurements per metre). It is worth re-iterating however that core measurements are made on far smaller volumes of rock than logging tools respond to, so care needs to be exercised when comparing the two. Core permeabilities are also normally measured with a single fluid, more often than not an inert gas at relatively low pressure. A whole series of corrections are needed to convert the measured absolute permeabilities at ambient conditions to effective permeabilities for reservoir fluids at reservoir conditions. These corrections are normally quite reliable and are typically based on SCAL measurements, made on a sub-set of plugs that should be representative of the whole data set. The challenge is to decide at what stage to apply the corrections and to QC the entire workflow to ensure they are applied once and only once. As stated in the previous section, for most of this chapter we will describe the generation of curves that reproduce absolute gas phase permeabilities (normally at over-burden conditions).

   Well testing avoids the need for corrections because it involves real reservoir fluids flowing at more or less reservoir conditions. The real reason for testing is that it measures the 'deliverability' of the formation directly. The latter is unquestionably crucial information for field development planning but it has led to a naïve belief, in some quarters, that well testing is sufficient to provide all other permeability information. This is nonsense, the dynamic behaviour of a

well involves a complex interplay of formation properties, varying fluid properties and well hydraulics. Additional complications arise from the possibility of formation damage and limitations of the instrumentation. A well test interpretation relies on a number of assumptions and simplifications and if any of these is unjustified the permeability coming out of the analysis will be wrong.

In fact no test produces an estimate of the permeability alone. All interpretation models actually give an estimate of permeability–thickness product KH (or 'K–H' as it is frequently referred, in this book we will use capital letters to distinguish it from horizontal permeability kh). Of course if one knows the thickness (H), the average permeability (K) follows immediately but this can be a big 'if'. Frequently used values for H include the perforation interval and the formation thickness. Both suffer from the fact that there is no guarantee that the whole interval is contributing to the flow (i.e. the true H is less than has been assumed). On the other hand if the perforation interval is narrower than the formation thickness it is quite likely H will exceed the former (true H is greater than the assumed value). Narrower perforation intervals are used where the well has penetrated a contact or when there are limitations on gun lengths. Particularly thorny issues arise when the test interval is close to a contact as relative permeability effects then come into play.

Wireline formation testers produce a permeability estimate as a by-product of the pressure measurements. They, in effect, conduct a miniature well test and results from the latest modular tools often show good agreement with core plugs. There are two features of the measurement, which mean it should be treated with caution, however.

1. The measurement is made very close to the borehole wall and may therefore be taking place in formation that has been altered by drilling.
2. Related to the above the fluid that is moving may be reservoir fluid, filtrate or worse an intermediate mixture. These generally have different viscosities.

From the brief discussion earlier it should be apparent that cores and tests actually compliment each other. Core-based measurements provide a large number of data points, made under highly controlled conditions but they are made on small, well-defined volumes of rock with simple, inert and well characterised fluids. Tests provide one or two data points made under *in situ* conditions. Tests typically involve a large volume of rock whose thickness may be poorly defined.

It is actually quite unusual for core and test permeabilities to agree. Reconciling the two can provide new insights however.

As noted earlier much of this section will deal with core data and so it is worth spending a bit more time looking at this. The way permeability is measured on an individual core plug was discussed in Section 3.7 here we are more concerned with the complete set of core measurements. Part of a typical modern core report is reproduced in Fig. 11.1.

Sample number	Depth	Confining stress 800 psi.				Ermeabili Steady STA	Porosity HG bulk	Grain density		Confining stress 2840 psi.				Commen
		Permeability		Orchheime Beta	Porosity					Permeability		Forchheime Beta	Porosity	
		K_{inf}	K_{air}				volume	(g/cm^3)		K_{inf}	K_{air}			
	(m)	(mD)	(mD)	$(ft.^{-1})$	(%)	Kair (md)	(%)			(mD)	(mD)	$(ft.^{-1})$	(%)	
						Core 1								
54	3786.80	20.1	22.3	4.11E+08	12.7			2.644		18.5	20.4	4.77E+08	12.3	
55	3787.10	178	189	2.75E+07	16.5	16.4		2.645		167	177	2.91E+07	16.1	Healed f
56	3787.41	99.6	106	3.52E+07	16.8			2.643		94.4	101	3.80E+07	16.4	
57	3787.70	64.7	69.8	5.28E+07	15.9			2.643		60.9	65.8	5.73E+07	15.5	
15V	3787.97	35.5	38.8	–	14.0			2.643		32.6	35.6	–	13.6	
58	3788.03	22.9	25.6	1.71E+08	15.1			2.644		21.4	23.9	2.79E+08	14.6	
59	3788.30	105	111	3.57E+07	16.0			2.647		99.6	106	3.78E+07	15.6	
60	3788.58	77.0	81.7	7.28E+07	15.3			2.705		72.8	77.4	7.90E+07	14.8	Pyrite
61	3788.90	84.9	89.7	5.29E+07	15.3			2.644		79.9	84.6	5.74E+07	14.8	
3DS	3789.08	117	124	–	16.2	16.1		2.641		–	–	–	–	
62	3789.50	96.3	102	4.06E+07	16.6			2.642		90.9	96.3	4.44E+07	16.1	
63	3789.80	75.9	80.9	5.25E+07	16.0	15.7		2.643		71.5	76.2	5.75E+07	15.4	
16V	3789.97	84.6	89.5	–	16.0			2.640		80.1	84.9	–	15.6	
64	3790.11	50.3	54.1	9.64E+07	15.0			2.644		47.1	50.7	1.09E+08	14.5	
65	3790.41	44.5	49.9	1.36E+09	17.2			2.643		42.5	47.2	1.57E+09	16.7	
66	3790.72	8.41	9.74	1.50E+09	13.2			2.647		7.18	8.28	1.75E+09	12.6	
17V	3790.97	0.011	0.017	–	6.4			2.659		0.002	0.005	–	5.7	
67	3791.04	0.173	0.256	1.74E+11	7.3			2.689		0.037	0.073	1.40E+13	7.0	Pyrite

FIGURE 11.1   Part of a routine core analysis report for a high net to gross sand (ca. 2005).

The example covers approximately 4 m of a massive sandstone reservoir (the full set of porosity–permeability data is shown in Fig. 11.2). Plugs have been cut approximately 30 cm apart, but few if any consecutive plugs are separated by exactly that difference. Modern practice is to take equally spaced plugs if possible, to remove sampling bias. But the core analyst also needs to avoid fractures, damaged sections, thin shales, pebbles etc. and so is unlikely to achieve truly uniform sampling. Notice also that the measurements have been made at two confining stresses: 800 and 2840 psi (right-hand columns). The latter has been calculated to simulate reservoir stress and would probably form the
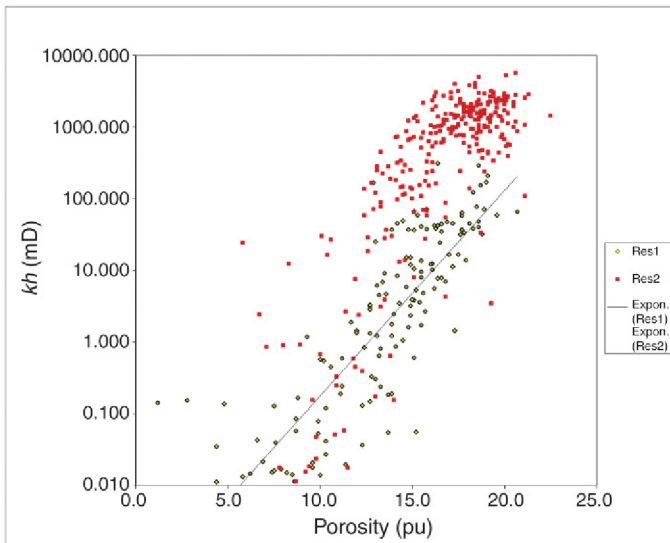
FIGURE 11.2   Example of compaction-corrected core data from a sandstone reservoir.

basis of the permeability curve. Complete sets of curves measured at overburden stress are a relatively new feature and so measurements are generally also made at the lower 'ambient' stress so that they can be compared with older data sets. At each stress two permeabilities are given: K_{inf} and K_{air}. The latter is the uncorrected measured permeability using air at low pressure. The former is the Klinkenberg corrected permeability, which accounts for the behaviour of the gas at high pressure. The permeabilities have been measured using a transient method so that there is a Klinkenberg permeability for every plug.

The measurement that is closest to reservoir conditions is the Klinkenberg value at simulated reservoir stress. This still requires the relative permeability to be accounted for to convert it to a true effective permeability but the effects of overburden and reservoir pressure is taken care of.

The simplest and most accurate way to generate a permeability curve is simply to use the core data directly. It needs to be re-sampled to give it the same regular increment as the logs and properties derived from them but otherwise it is a good indication of how permeability varies through the cored part of the reservoir. The problem is of course that the core may not cover the whole reservoir and even if it does it probably has not been cut in every well. Nevertheless, using core as far as possible has a lot to commend it.

11.4 PERMEABILITY PREDICTION

Figure 11.2 is a cross-plot of permeability at overburden against porosity at overburden for the complete set of data of which a small part was listed in Fig. 11.1. This cross-plot is the starting point for finding a relationship between porosity and permeability.

The points have been coded according to which reservoir unit they belong (named Res1 and Res2). The reservoir units themselves were picked on log character and correspond to different depositional environments. The Res1 points define a reasonable log-linear trend for which a Y on X regression gives the following:

$$\text{Log}(kh) = -3.699 + 0.2865\Phi \tag{11.1a}$$

Or equivalently,

$$kh = 0.002 \times 10^{0.2865\Phi} \tag{11.1b}$$

where, in this case, porosity is given as a percentage. Equations 11.1a and 11.1b are known as 'semi-log' relations and they are commonly applied to porosity–permeability data. In fact most models implicitly model logarithm permeability, this is a reflection of the fact that permeability typically varies over several orders of magnitude.

There are two things to note about Eqs 11.1a, 11.1b and 11.1c. Firstly, they are a regression for logarithm of permeability against porosity and so give an unbiased estimate of logarithm permeability. Secondly, ignoring a few very
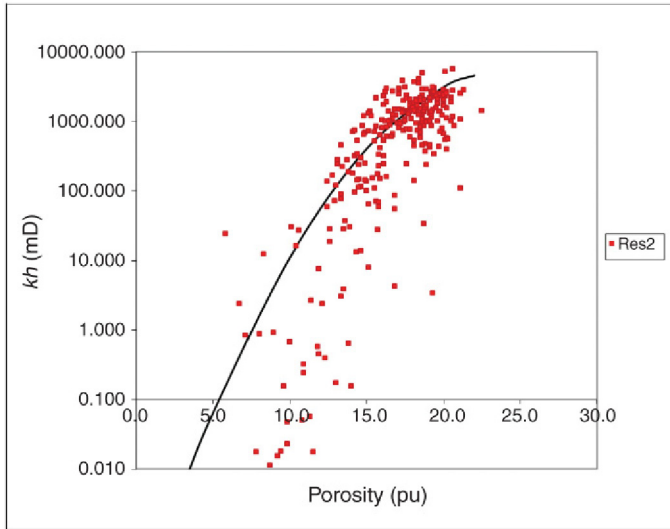
FIGURE 11.3   Core data from the Res2 zone and a relationship to model it.

low porosity plugs, the porosities in the Res2 unit vary from 7.5% to 19%, which equates to a calculated permeability range of 0.03–56 mD. Inspection of Fig. 11.2 shows permeabilities actually range up to 300 mD (in this case the lower end is about right). This reduction in the range is typical of predictions and it is particularly severe when logarithms are involved. The equation may do a good job of finding the average behaviour but at the cost of making the reservoir appear more homogenous than is really the case.

A similar approach could be taken to the Res2 data but inspection shows a common feature of permeability data from very high quality reservoirs. On average the permeability increases with increasing porosity but the rate of increase seems to fall as porosity – and permeability – increases. In short a convex curve would appear to fit the data better than the simple straight line used for Res1. An example is shown in Fig. 11.3 where a relationship of the form:

$$\mathrm{Log}(kh) = A + B.\Phi + C.\Phi^2 \tag{11.1c}$$

has been used. The criticism that the semi-log relationship reduces variation in the permeability applies equally well to this relationship. At any particular porosity, the core permeability can vary by about an order of magnitude. More complicated functions of porosity that technically fit the data better can certainly be found but they can do nothing about the scatter and so it is difficult to justify the additional complexity.

At this stage it is important to remember the purpose of Eqs 11.1a, 11.1b and 11.1c is to generate a permeability curve. This will be applied to log porosity, which implicitly, refers to volumes that are much larger than a core plug.
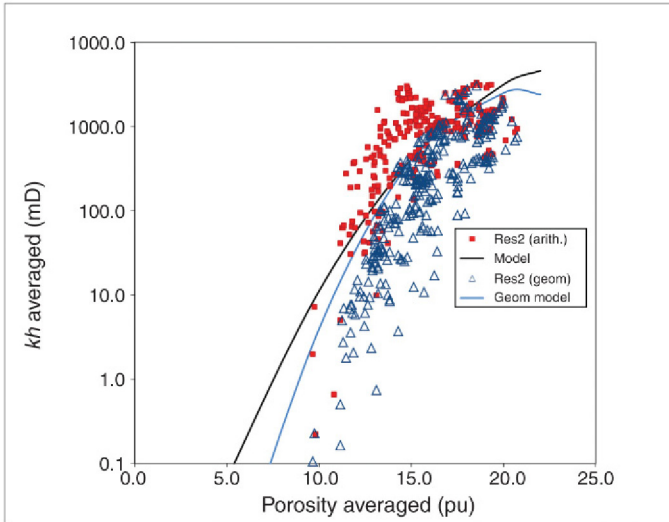
FIGURE 11.4   Cross-plot of moving averaged core data for the Res2 interval (the raw data is shown in Fig. 11.3). The averages are over five consecutive core plugs, the squares (red squares in the web version) use arithmetic averaging of the permeabilities (the mean of the five plugs) and the open triangles (blue open triangles in the web version) use a geometric average.

For the core data in Fig. 11.3 the plugs are an average 0.3 m apart so about five plugs cover an interval sampled by the logging tool. So, although rarely done it may actually be better to find a relationship that describes the moving averaged core data. This is easier said than done because of the enormous range covered by permeability data. The porosities can be averaged using an arithmetic mean, but is this the appropriate way to average the permeability? The answer in part depends on what sort of relationship is going to be fitted to the data, an equation such as Eq. 11.1c involves a logarithm or permeability so it really looks as if the appropriate average is a geometric average. In Fig. 11.4 a cross-plot of the averaged core data is shown with both arithmetic and geometric averaging applied to the permeability.

   The issue of how to average the core permeability is part of the more general problem of 'up-scaling'. With a core we can describe a reservoir rock at microscopic scales if we wish but it is impractical to model a whole field – or even a small part of one – at anything less than a metre scale.

11.5   KOZENY–CARMEN EQUATION

Special mention should be made of the Kozeny–Carmen equation if only because it is frequently invoked. The equation is an extension of the Kozeny equation which itself is derived from a similar model pore system to the capillary bundle discussed in Chapter 2. In the Kozeny equation the capillaries have a

range of radii. The Kozeny–Carmen equation modelled a system consisting of spherical grains of identical size. The resulting equation for permeability as a function of porosity is:

$$k = \frac{\Phi^3}{5S^2(1-\Phi)^2} \qquad (11.2)$$

where S is the 'specific surface area of the grains' (that is the surface area per unit volume so it has dimensions of inverse length. Squaring S makes the equation dimensionally correct). To apply Eq. 11.2 to real rocks the 'five' is replaced by a function of the ratio of 'hydraulic radius' to grain size (the hydraulic radius is related to the size of the pore throats). The Kozeny–Carmen equation has the attractive feature of being derived from first principles. The disadvantage is, of course, that it requires a knowledge of the size and shape of the grains. With the possible exception of NMR, this information will not be available from logs alone.

11.6 PERMEABILITY AS A FUNCTION OF POROSITY AND IRREDUCIBLE WATER SATURATION

The discussion at the end of Section 11.4 is completely academic if, as is commonly the case there is no core data. Although reservoir rocks are simply too variable to make general porosity–permeability relationships – analogous to log–porosity relationships – a useful proposition, functions of porosity and saturation are available and are worth trying in the absence of anything better. The rationale behind these equations is that permeability and irreducible water saturation are both controlled, to an extent, by grain size. As explained in Chapter 2, permeability is mainly controlled by pore throat size, but in an inter-granular rock that is itself strongly dependent on grain size.

As grain size decreases, the surface area increases, even if the porosity remains the same. So the irreducible water volume, that is largely made up of water adhering to the grain surfaces also increases. Therefore, the irreducible water saturation can be used as a measure of the surface area of the grains. The inability to estimate surface area directly with logs was what limited the application of the Kozeny–Carmen equation, but providing we are above the transition zone we have, in S_{wir}, a related property that can be readily found from logs.

In general then, permeability is expected to increase with porosity and decrease with irreducible water saturation. Several empirical equations, that quantify these trends, have been proposed over the years. Because they are functions of irreducible water saturation however, they originally could only be applied above the transition zone (thereby excluding water-bearing formations from further consideration).

These equations mostly have the form:

$$kh = \frac{A.\Phi^B}{S_{\text{wir}}{}^C} \qquad (11.3)$$
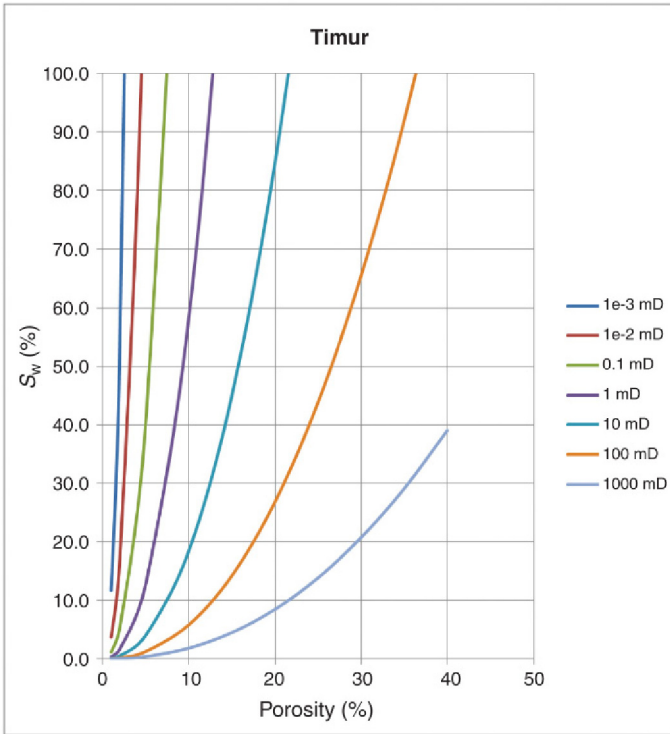
FIGURE 11.5   Graphs of constant permeability calculated from irreducible water saturation and porosity using the Timur–Coates equation.

with *A*, *B* and *C* constant. Most of the equations specify the values of constants but obviously there is nothing stopping the user modifying them if they have a good reason to (matching a well test for example). For example, the Timur–Coates equation of 1968 uses $A=0.136$, $B=4.4$ and $C=2$ and the similar Coates–Dumanoir equation (1973) uses $A=1,000,000$, $B=4.5$ and $C=2$. The large difference in *A* is due to the fact that in the former porosity and saturation are entered as a percentage and in the latter as fractions (Fig. 11.5).

Another equation that is sometimes used is due to Morris and Biggs (1967). The constants are $B=6$ and $C=2$ with *A* of the order of 10,000 (the inputs are entered as fractions). The Morris–Biggs equation gives significantly lower permeabilities at a particular S_{wir} or porosity than the other equations (Fig. 11.6). It is particularly good in clean but heavily cemented sandstones, which often have low irreducible water saturations and very low permeabilities.

These equations are used to generate the permeability curve that is a standard output of NMR logs. The NMR log provides both the inputs to equations such as Eq. 11.2. The total area under the T2 distribution gives the total porosity and because NMR tools can actually allocate water among different
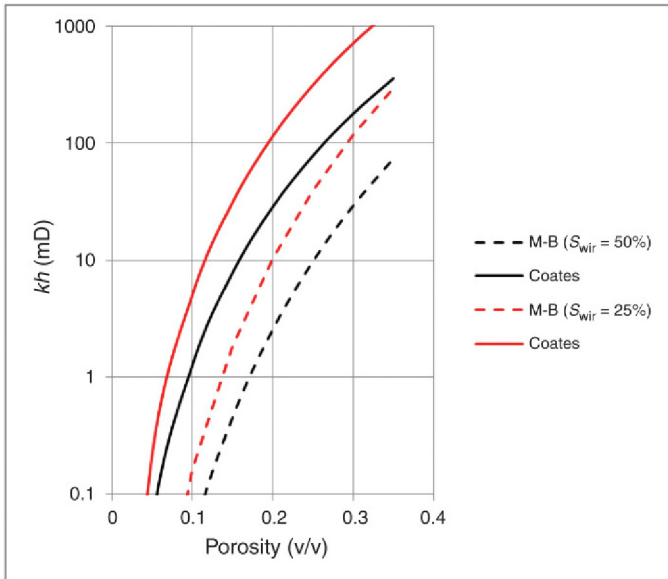
FIGURE 11.6 Permeability as a function of porosity calculated using the Morris–Biggs equation (dashed) and Coates equation (solid). The S_{wir} is fixed at 50% (black) and 25% (light grey [red in the web version]).

environments – clay-bound, capillary-bound, effective porosity etc. – it can also estimate the irreducible water saturation. This is the case even in the water leg. The Coates–Timur equation is commonly used to generate the curve, but any standard relationship could be used and if core data is available in one or two key wells a field-specific relation could be developed.

A variation on the theme was developed by Schlumberger researchers for use with their NMR tools. This equation has the form:

$$k_{sdr} = A.\Phi^B (T2gm)^C \tag{11.4}$$

The 'sdr' stands for 'Schlumberger-doll research', the company's centre for fundamental research that was located in Ridgefield, Connecticut. The term T2gm is the geometric mean of the T2 distribution. So the higher the value of the average, the greater the contribution of the effective porosity and the smaller the proportion of clay-bound and irreducible water. So the equation can be seen to again fundamentally depend on the ratio of pore volume and grain surface area. The logging engineer can change the values of the constants but typical values are $A = 4$, $B = 4$ and $C = 2$ (with porosity given as a fraction).

Care needs to be taken using NMR permeabilities in gas-bearing formations because there the total porosity will be under-estimated (as will be the case in any fluid which has a hydrogen index significantly lower than one). NMR will

still estimate S_{wir} however, so Eq. 11.2 could still be used with a porosity from conventional log analysis.

These equations are unlikely to be applicable to vuggy or fracture porosity. This is because the irreducible water saturation is determined by the pores rather than the connections between them, but in reality permeability depends mostly on the radii of the connections. With these porosity types the assumption that the two or quite closely related does not apply. In fact even with inter-granular system there is no universal relationship between them. For example, clays and cements may preferentially form in the pore throats drastically reducing the permeability but having a minor effect on S_{wir}. So these relationships tend to work best in clean sandstones that have undergone little diagenesis. In short these relationships may work well with no modifications, they may work well after some modification or they may not work at all.

11.7 ANALOGUES AND ROCK TYPES

As mentioned already, the problem with permeability prediction from logs is that the principal controls on it are not properties that logs are particularly sensitive to. Consequently, hard permeability measurements are needed to derive or at least 'ground-truth' log-based estimates. Ideally these come from the formation/well under consideration but there is a good chance that no such information exists. This is particularly true for wells that were drilled early or late in the history of a play. The alternative is to use analogue data; in other words permeability data from a similar formation in another well. The most obvious source is a nearby well drilled in the same play but a number of companies have developed the idea of rock typing where the analogue is selected because of its similar petrography and mineralogy, rather than geographical location or age. Rock-typing systems vary from a stand-alone exercise to develop a permeability transform for one particular field, to sophisticated databases that are designed to be universally applicable. All rely on having a good description of the reservoir rock at the pore scale. Their application relies on the assumption that petrophysical properties are controlled by properties such as grain size, clay content, degree of cementation rather than age, basin and depositional environment.

As an example the 'World wide rock catalogue' (or WWRC), which is a commercially available product, classifies sandstones using five characteristics:

1. Mean grain size
2. Sorting
3. Degree of consolidation
4. Argillaceous content
5. Measured porosity

The value of each of these is expressed as a single digit between 1 and 5. For example, grain size varies from pebble (1) to very fine (5). Any sandstone is fully described by a five-digit code between 11111 and 55555 (there are some

combinations that simply do not exist and others which for whatever reason are not in the database). Core plugs from an example of each one of these has been subjected to the same series of routine and special core analysis tests (including permeability). The user can use a mixture of cuttings, logs and possibly sidewall core descriptions to assign values to as many characteristics as possible. The tool will then search for all the matches, obviously the more characteristics that can be assigned a value, the shorter the list of candidates. If the only information that is available is porosity, the list will be long and we will be no further forward, so a reasonable knowledge of at least some petrographic properties is essential. But if this exists the permeability range can be substantially narrowed. There is also a carbonate catalogue, which uses a different set of characteristics.

One problem with the WWRC for permeability modelling, is that it is based on short sections of core so there is only a limited range of porosities and permeabilities for any particular sandstone (or carbonate). Nevertheless, a single porosity–permeability estimate is a lot better than nothing and the catalogue could be used as a first step towards identifying full sets of analogue core data.

Many of the larger operators have proprietary tools for predicting porosity–permeability relations but these invariably rely on a knowledge of at least some of the characteristics listed earlier.

11.8 MORE LOG-BASED METHODS

NMR is the most reliable and widely used log for directly estimating permeability but there are one or two other modern logs that are sometimes tried. In the 1990s the Stonely wave generated by sonic logging tools was investigated as a possible permeability measurement. The Stonely wave is a particular mode generated in mud-filled boreholes. It normally forms the late part of the wavetrain generated by a sonic log. Its use as a permeability measure relied on the fact that it would be more strongly attenuated in permeable formations. The reason being that in permeable formations energy would be lost squeezing fluid away from the borehole. This is quite an attractive model because it means the measurement does depend on moving fluid and so in a sense really is measuring permeability. Unfortunately, in practice results have been mixed and although there are published examples of it doing a reasonable job matching core permeability there are plenty of un-published examples where it has proved to be an expensive failure.

A permeability curve can also be generated by geochemical logs. This is yet another implementation of the Kozeny–Carmen equation. Specifically, the relative amounts of clay minerals determined by the tool are combined to estimate the specific surface area (S) introduced in Section 11.5. This could be used in conjunction with porosity estimated by conventional log analysis to find permeability using Eq. 11.2 but in fact, like with NMR tools, a permeability curve is produced using a semi-empirical function of the outputs of the tool. Given that this is yet another permeability that is based on the ratio of pore surface area to volume, it will have similar limitations to the NMR permeability.

11.9 A CASE STUDY

We will conclude this chapter with a case study illustrating how a permeability curve can be built using limited core measurements. This is not a real example but it is based on one. It represents a common situation where a curve has to be built using limited or non-existent core data (generating a curve when a comprehensive set of core data is available was discussed in Section 11.4).

11.9.1 Using *K–H* from a Well Test

A common situation is to have a carefully designed and executed well test but no core data. This means that an accurate value for the permeability–thickness product – *KH* – is the only permeability measurement available. As noted earlier if *H* is known with reasonable certainty, this is tantamount to knowing the average permeability (as also noted, this may be a big 'if'). Depending how much is known about the reservoir the equation for the permeability curve could be based on analogues, NMR or geochemical logs or one of the practical ways of implementing the Kozeny–Carmen equation described in Section 11.6. The parameters in these equations are then chosen so that the average for the curve equals the average permeability calculated from the well test.

This seems simple enough but we have to decide which average is the appropriate one to use. The two best candidates are arithmetic and geometric. The geometric mean will always be less than the arithmetic but if the function gives a maximum variation in permeability over the tested interval of, say, one order of magnitude the difference will not be great and it does not really matter which average is used for the match. Remember we are basically trying to reproduce a function that ideally would have been developed from core plug measurements of porosity and permeability made every 30 cm or so. These measurements would invariably show quite a lot of scatter and neither the geometric nor the arithmetic average of them is going to match the well test exactly (even if all the corrections to reservoir conditions have been correctly applied). In other words the difference between the two averages is much smaller than the inherent uncertainty in permeability.

As the range of computed permeability increases however, the difference between the two averages starts to become quite significant. It is impossible to generalise but if there is a variation of two orders of magnitude the two averages could differ by a factor of 10 and then we need to think carefully about which to use. Modelling work performed by researchers at Heriot-Watt University in Edinburgh suggests that the arithmetic average is the appropriate one to use in a layered system. This means layers with more or less constant permeability that are laterally extensive and more or less normal to the well bore. The geometric average is the one to use when the reservoir is actually quite heterogeneous and permeability varies in all directions. In the absence of any other information I would go for the former.

There is really no point trying to calculate the constants in, say, Eq. 11.3 to several significant figures. Once again remember we are trying to model
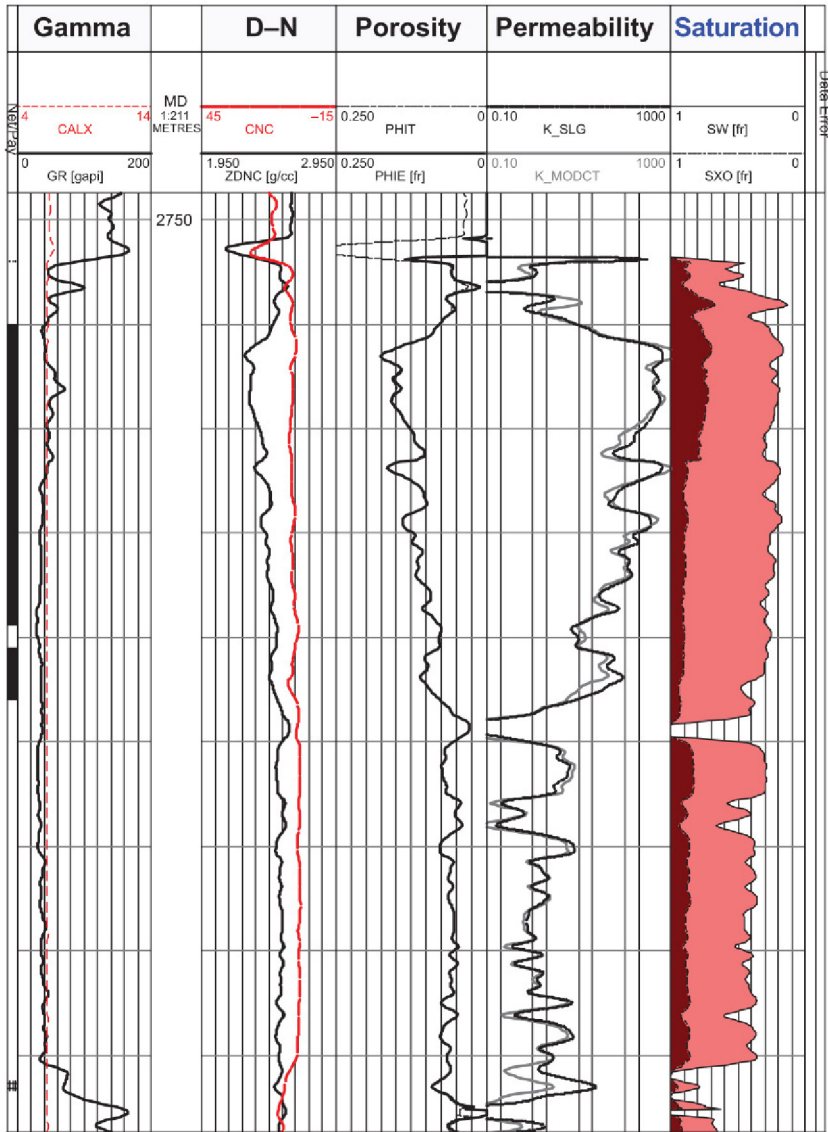
FIGURE 11.7 A 40-m thick, gas-bearing sand showing computed porosities from log analysis (track 4), and permeability curves calculated using two different models on a logarithmic scale from 0.01–1000 mD. The RH Track is a summary of the static properties. (track 5).

something with a lot of inherent scatter so it is actually more informative to use parameters with one or two significant figures that give a close but not exact match to the test *KH*. As an example consider the gas-bearing sand shown in Fig. 11.7. The whole sand was tested and gave a *KH* of 5080 mD.m, the total sand thickness is 40 m. Two alternative permeability models are shown in
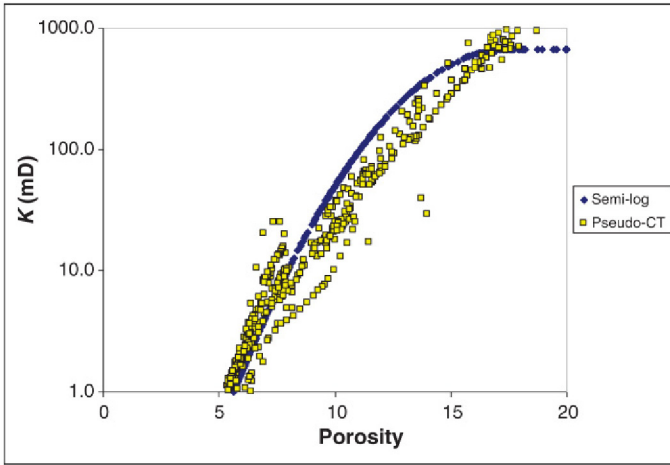
FIGURE 11.8   Calculated permeabilities against porosity using a semi-log porosity equation (dia-monds [blue diamonds in the web version]) and a modified version of the Timur–Coates equation (squares [yellow squares in the web version]).

the figure. One is a semi-log porosity-permeability relationship based on ana-logues and the second is a modified form of the Timur–Coates equation 11.3 (Section 11.6).

The specific equations used were:

(black dashed curve) $\text{Log}(kh) = -3.3 + 0.7\Phi - 0.02\Phi^2$ (11.1c)

The permeability is fixed at 668 mD for porosities in excess of 17.4%.

$$(\text{solid grey curve}) \qquad kh = \frac{0.3\Phi^5}{S_{\text{wir}}^2} \qquad (11.3a)$$

In both equations porosity is input as a percentage. First pass permeabil-ity curves were calculated using the parameters for the analogue or the default Coates–Timur parameters (see Section 11.6) and then these were gradually changed until a satisfactory match was obtained. The match was made between a permeability thickness product calculated for the curve(s) over the entire 40-m thickness and the well test value. For a typical log increment of 0.1524 m there are 263 separate permeability points, the permeability–thickness product is ob-tained by summing these and multiplying by 0.1524 m.

Cross-plots of permeability calculated from both equations at each poros-ity calculated from the logs are shown in Fig. 11.8. The permeability thickness product calculated for the two models are 4634 and 4584 mD.m for Eqs 11.1c and 11.3a respectively. These are approximately 10% lower than the value ob-tained from the well test. A closer match to the test could be obtained by using different parameters in the equations, but as noted several times already the

inherent scatter in permeability data means that we have not improved accuracy of the model.

Note that because we are matching a test, the equations implicitly model effective permeability (to gas) at reservoir conditions. This is different to the curve fits to the core data in Fig. 11.2 where the equations model absolute permeability at reservoir conditions.

Page left intentionally blank

Chapter 12

Complex Lithology

Chapter Outline

12.1 INTRODUCTION

The log analysis methods described so far have generally assumed the reservoir has a uniform mineralogy. Most of the examples have dealt with clean, quartz-rich sandstones or in a few cases limestone. In either case the physical properties of the matrix show little variation and can satisfactorily be modelled by constant values. For example, the matrix density of sandstone is fixed at $2.65\,g/cm^3$. In reality the mineralogy is unlikely to be truly uniform and grain densities measured on core plugs from even the cleanest and most massive sandstones typically show a variation of $\pm 0.01\,g/cm^3$. This creates relatively little uncertainty in all but the tightest sands; however and at any point the true porosity is unlikely to be more than about 0.5 pu different from the calculated value. Furthermore, providing the matrix density is close to the true mean grain density the average porosity will be even closer to the true value. To put it another way the assumption of a grain density fixed at the mean value results in porosity being over-estimated as often as it is under-estimated.

Some formations are not even approximately uniform however, in clastics variable amounts of clays and heavy minerals lead to grain density and other matrix parameters varying quite significantly along the well and in carbonates variable proportions of limestone, dolomite and anhydrite can produce similar variations. In these cases it should still be possible to calculate an accurate average porosity, but at any point the calculated porosity may be several porosity units different to the true value. To improve on this and calculate a porosity curve

319

that does a reasonable job in matching the true porosity at every point the matrix parameter(s) have to vary and therefore some way of estimating them at each point is needed (at least if we are to apply the equations given in Chapter 7). In fact this chapter is largely about ways of estimating grain density and other matrix parameters from logs.

One type of reservoir where the mineralogy is nearly always more complicated than 'predominantly' quartz or limestone are the so-called oil and gas shales. Although called 'shales' these generally are made of between one-third and one-half carbonate or quartz by volume (recall that most shales consist of 70% clay by volume). The relatively low clay content is necessary to make them brittle enough to fracture. Furthermore the mineralogy will often vary significantly through the reservoir so not only are the matrix parameters very different to the 'text-book' values used previously but they change continuously as well. Fortunately, there are a number of tools and combinations of tools that allow the mineralogy to be deduced to at least the accuracy needed to estimate the logging parameters. Some of these tools have not been described before in this book so we will take the opportunity to describe them here.

12.2 PHOTO-ELECTRIC FACTOR

The photo-electric factor or PEF is a by-product of the density measurement that first appeared in the 1980s when improvements to detectors and instrumentation resulted in more accurate density tools entering service. The PEF is related to the average atomic number (Z) of the formation and typically has values ranging from a bit less than two for a high-porosity quartz sand to 5 or 6 for a low porosity limestone. The relationship:

$$PEF = \alpha\, Z^{3.6} \qquad\qquad (12.1)$$

is sometimes given in publications, where Z is the average atomic number and α is constant. It is important to realise that this is only an empirical equation and a lot of real substances have PEF values very different from its predictions. Equation 12.1 does at least emphasise that PEF is disproportionately influenced by heavy elements (i.e. those with a high atomic number). In practice this means that PEF can be dominated by traces of heavy elements such as iron. As is often the case in practical petrophysics, this can be viewed as either a limitation or a very useful feature.

As noted earlier the PEF measurement is a by-product of the density measurement, it uses the same hardware and like the density measurement exploits the interaction of gamma rays with electrons. The PEF measurement concentrates on the count-rates of low energy gamma rays, which are identified using the same principles that allow spectral gamma-ray tools to measure the energy of gamma rays. Gamma rays with energies less than 100 keV are particularly susceptible to absorption by the electrons associated with heavy elements. The process is known as photo-electric absorption and as the name implies it results

TABLE 12.1 Densities and PEF Related Properties for Some Commonly
Occurring Minerals

Mineral	Formula	Density (g/cm³)	U (b/cm³)	PEF (b/g)	Mean Z
Quartz	SiO_2	2.65	4.79	1.8	10.0
Calcite	$CaCO_3$	2.71	13.77	5.1	10.0
Dolomite	$Ca/MgCO_3$	2.84	9.0	3.5	9.2
Pyrite	FeS_2	5.00	84.9	17.0	19.3
Siderite	$Fe(CO_3)$	3.89	57.2	14.7	11.2
Anhydrite	$CaSO_4$	2.96	14.93	5.0	11.3
Water	H_2O	1.00	0.36	0.36	3.3

in the gamma-ray disappearing (by contrast Compton scattering which is the basis of the density measurement results in the gamma-ray losing energy and changing direction but not disappearing completely). As a rule of thumb the higher the atomic number the stronger the absorption and so a low count-rate of low energy gamma rays implies a high atomic number.

The PEF curve is actually related to a more fundamental property known as the photo-electric absorption cross-section. This is given the symbol U and is conventionally measured in 'barns/cubic centimetres', this is another example of an unholy mixture of units which occur throughout this book. The barn, symbol b, is equal to 10^{-28} m² and so U has the dimensions of reciprocal length with SI unit m⁻¹ but barns/cm³ gives more manageable numbers and providing everyone is consistent in its use, no problems will arise. Some U values and densities of the commonly occurring minerals – including fluids – were given in Table 1.1 in the first chapter and the ones that are particularly relevant to this section are reproduced in Table 12.1 (the mean atomic number is included so that the reader can see how well they fit Eq. 12.1).

The mixing rule for U is the arithmetic mean (like density). So at any point in the reservoir the following holds true:

$$U = \sum v_i . U_i \qquad (12.2)$$

Where U_i is the photo-electric absorption cross-section of component i and v_i is its volume fraction. The PEF is defined as U divided by density (in gram cubic centimetre) and since the tool measures both these properties, PEF can be output in real time. The unit of PEF can be seen to be b/g although for some reason it is always quoted as 'barns per electron (b/e)'.

It is worth asking why the derived quantity PEF is plotted when the tool could output the more fundamental property U. The reason is that PEF has the convenient property of being almost independent of porosity. This can be appreciated by actually calculating some densities and U values for some clean quartz

TABLE 12.2 Calculated Porosity, U and PEF Values for Water Bearing Quartz Sands with Various Porosities

Porosity	0%	10%	20%	30%	100%
Density (g/cm^3)	2.65	2.49	2.32	2.16	1.00
U (b/cm^3)	4.79	4.35	3.90	3.46	0.36
PEF (b/g)	1.81	1.75	1.68	1.61	0.36

sands. This has been done for model sandstones consisting simply of quartz and water in Table 12.2.

The densities and U values for quartz and water are in the left- and right-hand columns, respectively and the values for the mixtures are found by applying Eqs 12.2 (U), 7.3a (density) and the definition of PEF. Although PEF does fall with porosity, the change over the range of realistic porosities is small (about 10%). (The reader may like to confirm that if the pore space is filled with a very low-density fluid like air – which also has a very low U value – PEF hardly changes at all.)

So the PEF curve has the convenient property that its value at a particular depth can immediately identify the mineralogy. For example, limestone has a value between 5 and 4.7 and dolomite varies between 3.2 and 2.8. Clays vary from 1.8 for kaolinite to over 6 for chlorite (the latter contains iron). Normally the PEF curve is plotted in the same track as the density – and neutron – using a scale of 0–10 or 20.

However, the PEF does have a few draw-backs. Firstly, it is based on just a fraction of the gamma rays that contribute to the density measurement and so it is not very accurate. In practice this means one should be wary of reading too much into small variations. Secondly, as already mentioned it is disproportionately sensitive to high atomic number elements so that even traces of some heavy elements can dominate the measurement. A good example is barite (barium sulphate) with a U value of 1070 b/cm^3 and a density of 4.1 g/cm^3 (the high U-value is largely attributable to the barium with atomic number 56). It can occur naturally in petroleum reservoirs, particularly in mineralised fault zones, but it is most likely to be encountered as a drilling mud additive. Because of its high density it is used to increase the mud weight, but with such a high U-value it does not take much to dramatically increase the PEF of a quartz sand. Using the same methodology that was used to construct Table 12.2 it is possible to show that 0.1% barite by volume is enough to increase the PEF of the 20% porosity sand from about 1.7 b/g to 2.1 b/g. Although the density tool is designed to exclude mud from between the tool and the borehole wall some mud inevitably gets trapped and furthermore, if a mudcake forms it will certainly include barite, which is part of the solid component of the mud. In other words, even in a perfect hole, enough barite is likely to be present to increase the PEF.

Although barite does occur naturally a more likely cause of a high PEF within the formation itself, are various iron-containing minerals such as siderite and pyrite. The PEF and related properties of these are given in Table 12.1. The U values are nothing like as high as barite but sufficiently high that a few per cent by volume is enough to significantly increase the PEF.

12.2.1 Using PEF to Estimate Matrix Density and the Density/PEF Cross-plot

As noted earlier, this sensitivity to heavy elements could be considered a draw-back but actually it is a powerful way of quantifying 'heavy' – in other words dense – minerals. The commonly occurring heavy minerals tend to be iron compounds and they are dense enough that a few per cent by volume can significantly increase the grain density. Fortunately because iron has a high atomic number – 26 – they also have a high U and so are easily identified using the PEF curve.

As an example, consider a sandstone that includes variable amounts of pyrite up to several per cent by volume. This is by no means unusual and a specific example was looked at in Section 1.4. The increase in sandstone grain density with pyrite content is given in Table 12.3.

Failure to account for the pyrite will result in porosity being under-estimated. For example, with 5% pyrite by volume, a 20% porosity water-bearing sand has the same density as a 14% porosity sand with no pyrite. The U and PEF values for these quartz–pyrite mixtures can be found from Eq. 12.2 (Table 12.4).

So the PEF value could be used to find the pyrite volume fraction, which in turn allows the appropriate matrix density to be picked. In fact this reasoning can be extended to calculate the density and PEF value for any mixture of known

TABLE 12.3 Calculated Grain Densities for Quartz Sand Containing Minor Amounts of Pyrite

Pyrite fraction (% by volume)	0	1	2	3	4	5
Grain density (g/cm^3)	2.65	2.67	2.70	2.72	2.74	2.77

TABLE 12.4 Calculated U and PEF Values for Quartz with Minor Amounts of Pyrite

Pyrite fraction (% by volume)	0	1	2	3	4	5
U (b/cm^3)	4.79	5.59	6.39	7.19	7.99	8.80
PEF (b/g)	1.81	2.09	2.37	2.64	2.91	3.18
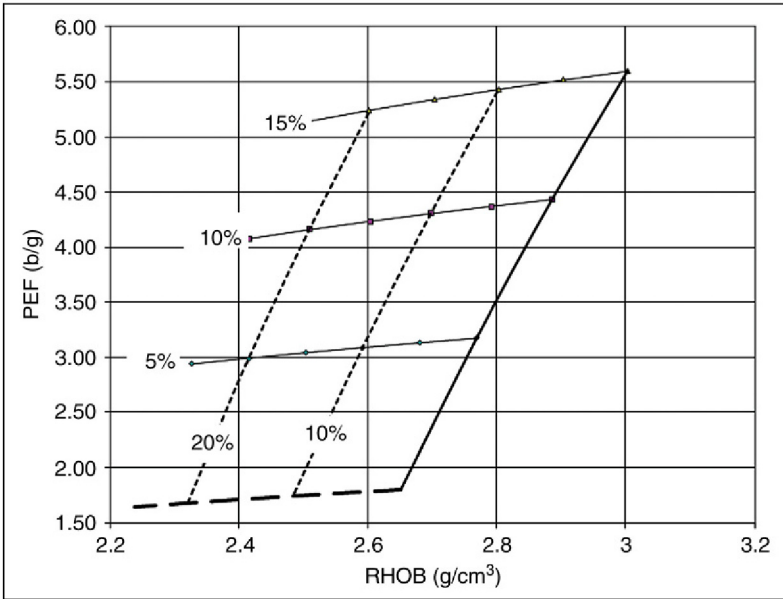
FIGURE 12.1 A density–PEF cross-plot. The overlay has been computed for a quartz sand with up to 15% pyrite by volume. The porosity is filled with water of density $1.00 \, g/cm^3$. The steep lines are lines of equal porosity (the porosity values are marked on two lines). The almost horizontal lines are lines of equal pyrite volume fraction, the value is shown at the left hand end of the line. The unmarked lines are for a clean, pyrite-free sand and a quartz–pyrite mixture with no porosity.

composition. In particular it is possible to calculate the density and PEF value for a pyritic sand of any porosity and this can be used to construct an overlay for a cross-plot of PEF against density. The overlay gives the porosity and the pyrite level for each point. An example is given in Fig. 12.1.

Obviously, to do this it is necessary to have some independent evidence that the high PEF is due to the presence of pyrite and not siderite or some other heavy mineral. A PEF value of 3.2 could be an indication of 5% pyrite or it could imply a dolomite bed, as always some independent information on lithology and mineralogy is needed (dolomite with a porosity of 5% also has a density close to $2.77 \, g/cm^3$).

12.3 DENSITY–NEUTRON CROSS-PLOT

The density–neutron combination was discussed at some length in Section 6.5.2 particularly as a shale indicator. There it was noted that, as a shale indicator, it was not suitable for sands of variable composition, this is because the sandstone line would no longer be fixed. Once again this limitation for one application can become an opportunity in a different application. In fact how this is possible has already been covered in Chapter 7 (Section 7.4.2). We will revise it here and illustrate its use by applying it to a particularly complicated interval.

In Section 7.4.2 it was noted that the three 'lithology lines' that are invariably super-imposed on a neutron–density cross-plot can be thought of as lines of equal matrix density. The values are 2.65, 2.71 and 2.85 g/cm^3 for sandstone, limestone and dolomite, respectively (the exact value for the dolomite line varies depending on the publisher of the chart, this reflects the range of densities that are quoted for dolomite). An example of a chart is given in Fig. 12.3, like the charts shown in Chapter 6 this is not for a real tool. If, for the time being, we assume the formation is clean and that there are no hydrocarbon effects to complicate matters then every point on the plot corresponds to a unique combination of porosity and matrix density. The lines of equal matrix density run parallel to and include the three lithology lines and as discussed in Chapter 6 converge to the fluid point which is beyond the top right corner of the plot. Lines of equal porosity are roughly orthogonal to these.

To illustrate how the density–neutron combination can be used to infer something about the matrix consider the short section of logs shown in Fig. 12.2. This was recorded in what is normally a dolomite formation. The most obvious feature of this particular section is the high-density beds that occur throughout. In the conventional density–neutron presentation to the right of the depth track they repeatedly go off scale (i.e. exceed 2.95 g/cm^3). They have been plotted on a reduced scale, together with the PEF, in the next track and it can be seen there that they reach 3.5 g/cm^3 at some points. This is of course far higher than the density of pure dolomite. Note that the PEF generally exceeds 20 b/g in these high-density beds suggesting they contain a significant quantity of heavy elements.

The same information is shown in Fig. 12.3 as a density–neutron cross-plot (although the cross-plot covers the whole formation). An additional 'lithology line' has been added that is parallel to the dolomite line, this is equivalent to a grain density of 2.97 g/cm^3. In addition two iso-porosity lines have been added corresponding to 0% and 15%, these have been extended to meet the new, high-matrix density lithology line. In order to do this involves some educated guesswork because the author of the original chart only provides information up to the matrix density of dolomite. It is easy to imagine a whole series of 'lithology lines' or 'iso-grain density' lines more or less parallel to the ones shown in Fig. 12.3.

The formation was cored and it is known the high-density, high-PEF beds which contain metal sulphides including pyrite (iron), chalcopyrite (copper/iron), galena (lead) and the zinc sulphides. These are all very dense minerals and the metals are all in the 'heavy' part of the periodic table. Pyrite has a density of 5 g/cm^3 and Galena is 50% denser again. It does not take much of either of these minerals to give the observed densities. Furthermore, the heavy elements will also tend to produce high-neutron porosities as they are generally strong neutron absorbers. So the points affected by these metal sulphides will be pulled towards the lower right corner of a conventional neutron–density cross-plot. The densest points, with densities in excess of 3.2 g/cm^3 and neutron porosities of 10–15%
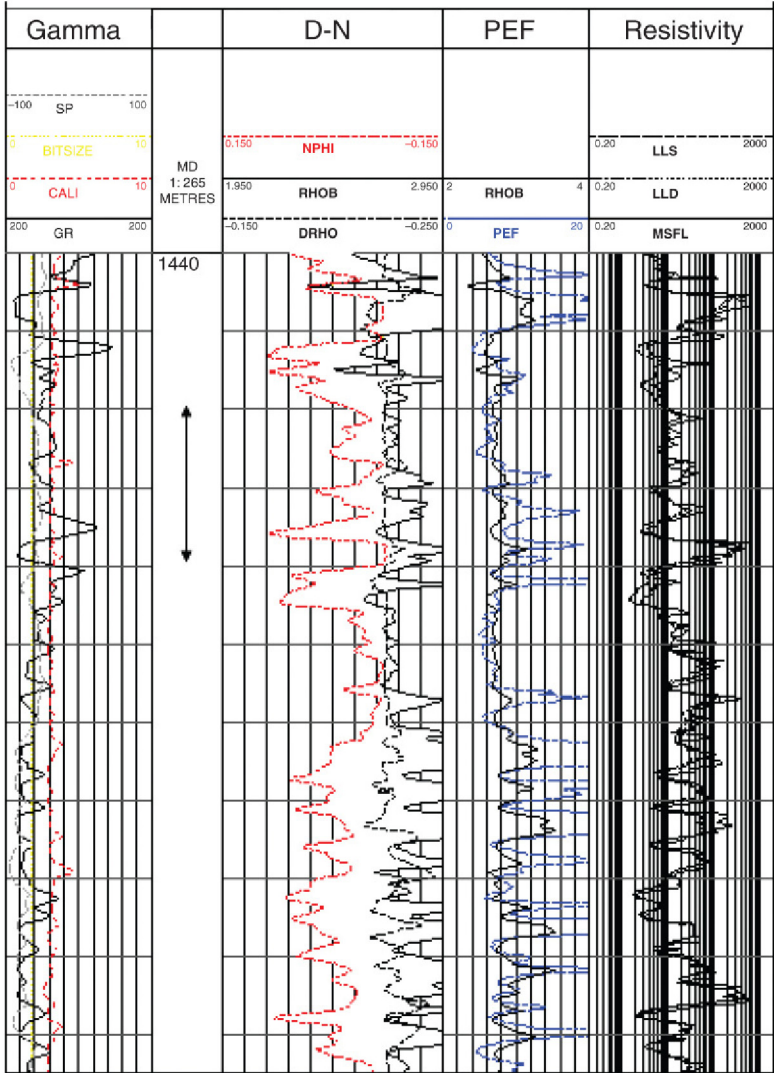
FIGURE 12.2 A 50-m section of logs from part of a mineralised fault zone encountered in an exploration well drilled in the Canning Basin of NW Australia. The log responses are discussed in the text. The arrow in the depth track is 10-m long.

are probably heavily mineralised with little or no remaining porosity. But the points with lower densities may well have some primary porosity remaining.

As an example consider a point with a density of 2.8 g/cm³ and a neutron porosity of 15 pu. If the reservoir is considered to be dolomite the calculated porosity would be between 3% and 4% depending on the value assumed for the matrix density of dolomite. But of course the matrix is not pure dolomite,
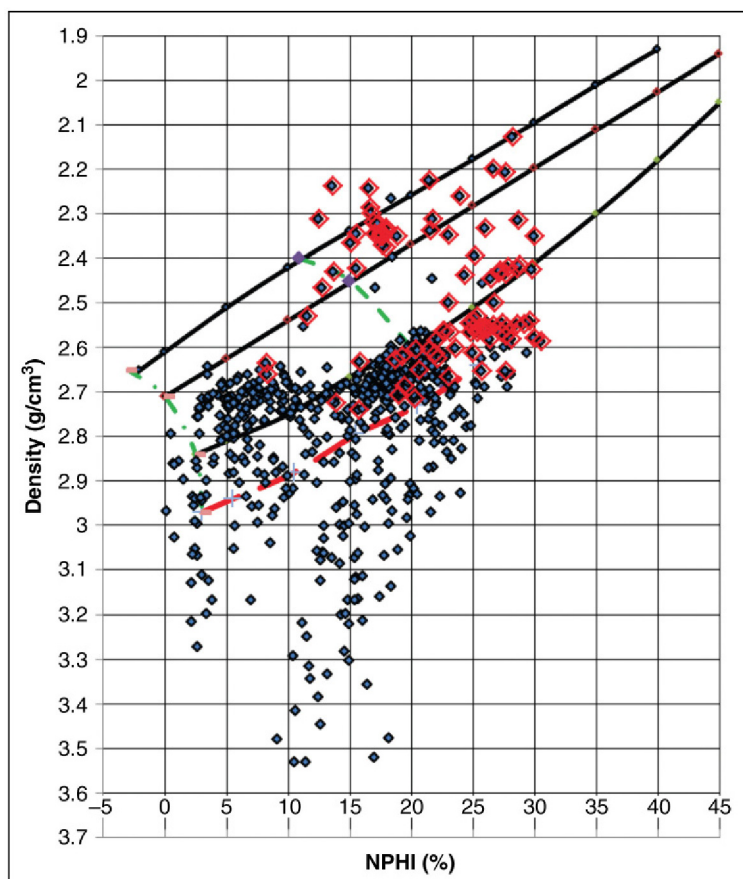
FIGURE 12.3 Density–neutron cross-plot constructed from the logs shown in Fig. 12.2. The points outlined light grey (coloured red in the web version) are shales identified by their high gamma activity. The dashed line (dashed red line in the web version) is an additional lithology line corresponding to a matrix density of 2.97 g/cm³.

we know this because the point lies well below the dolomite line. It may not be dolomite at-all but the plot suggests that whatever it is has a matrix density of 2.97 g/cm³. Using this in the density equation gives a porosity of 9%. Incidentally, using the neutron porosity and assuming pure dolomite would give a porosity of 7–9% depending on the tool type. Notice also that the average of the porosity calculated from the density assuming dolomite and the neutron porosity as read is 9%. This method for calculating porosity was mentioned in Section 7.4 and this result suggests that providing the matrix has not departed too far from pure dolomite, the method is reasonably robust. Alternatively, one could argue that the agreement between the cross-plot porosity and the simple average of density and neutron porosity shows that, despite where the point

plots, there is still a substantial amount of dolomite in the matrix. The reader can confirm that only 5–6% pyrite or even less galena is needed to produce the high-matrix density.

The further the points move from the standard lithology lines the less certain is the matrix density and porosity. For example, the point with a density of 3 g/cm^3 and a neutron porosity of 15 pu has a matrix density somewhere between 3.1 g/cm^3 and 3.15 g/cm^3 which leads to quite a lot of uncertainty in the density porosity equation. The simple average method seems to break down completely as well.

12.3.1 Complex Lithology in the Presence of Hydrocarbons

Hydrocarbons, particularly gas, complicate the use of the density–neutron cross-plot because they tend to shift the points in the direction of a lower matrix density. A gas-bearing dolomite, for example may end up giving the same density–neutron response as a water-bearing sandstone. It may still be possible to deduce matrix density and mineralogy but the effects of the hydrocarbon have to be removed. In principle, this is possible if the porosity and saturations are known. The normal approach would be to use the equations introduced in Chapter 9 to find the magnitude of the hydrocarbon effects (Eq. 9.8 for the neutron 9.4 for the density). These can then be removed to find the density and neutron readings in water-bearing formations and that in turn gives the matrix density. In practice this can involve many iterations and under certain conditions the whole calculation may be quite poorly conditioned. In particular if shale volume also has to be accounted for there is a real danger of the problem settling on a nonsensical result.

12.4   CASE STUDY: LIMESTONE–DOLOMITE SYSTEMS

12.4.1   Chemistry and Physics of Dolomite: Properties and Occurrence

A special case of complex lithology occurs in carbonates where the matrix may be limestone or dolomite or some intermediate mixture. We will use this to illustrate the use of some of the tools discussed earlier and also as a good way to revise some of the other topics in this book.

The densities of dolomite and limestone are so different that if the wrong mineralogy is assumed a gross error in porosity will result. Fortunately, in principle the two can be easily distinguished using logs and in particular the density and neutron tools. Incidentally, the term 'limestone' is a good example of a subtle difference in meaning depending on who is using it and why. Log analysts typically use limestone to mean the pure mineral calcite. The limestone line on a cross-plot implicitly refers to a mixture of calcite and water. In geological usage, on the other hand, a limestone is basically a carbonate rock with a high proportion of calcite in it. More often than not 'limestones' are sufficiently close

TABLE 12.5 Physical Properties Relevant to Logging for Calcium and Magnesium Carbonates

Mineral	Density (g/cm³)	PEF (b/g)	U (barns/cm³)	Sigma	Dt$_{comp}$ (μs/m)	Dt$_{shear}$ (μs/m)
Calcite	2.71	5.1	13.7	7.1	151	292
Dolomite	2.85	3.1	8.8	4.7	138	236
Magnesite	3			1.5	144	246

to the log analyst's ideal that it is not worth making the distinction. But there are plenty of examples where 'limestones' contain significant amounts of other minerals and calcite may even be present at less than 50% by volume. In this section at least we will be referring to the minerals and will endeavour to use 'calcite' whenever $CaCO_3$ is involved. The same comments apply to dolomite, although in its case the nomenclature is unavoidably blurred because dolomite is certainly a stable chemical compound with formula $Ca_{0.5}Mg_{0.5}CO_3$ but the name also describes rocks with more than just dolomite in them. In this section, as far as possible, 'dolomite' will mean the pure mineral. Some key physical properties of the two minerals were given in Table 1.1, at beginning of the book. For convenience they have been given again in Table 12.5 with the acoustic properties converted to slowness.

Dolomite is related to calcite by replacing half the calcium atoms with magnesium. The replacement is not a random process and the calcium and magnesium atoms occupy well-defined sites in the crystal lattice. Both minerals are stable chemical compounds and a particular mixture of calcium, magnesium and carbonate ions will react to give calcite and dolomite and not some compound with an intermediate composition. The laws of chemical thermodynamics cause most carbonate reservoir rocks to be primarily calcite or dolomite. As an example the 'compound' $Ca_{0.7}Mg_{0.3}CO_3$ will spontaneously decompose according to the following reaction:

$$10.Ca_{0.7}Mg_{0.3}CO_3 \rightarrow 6.Ca_{0.5}Mg_{0.5}CO_3 + 4.CaCO_3$$
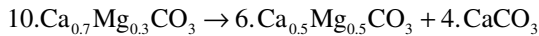
It must be remembered that logs have large volumes of investigations so if the reaction above results in a mixture of calcite and dolomite crystals that are barely visible to the naked eye then, as far as any log is concerned, the matrix is still intermediate between the two end members. In fact concentrating the individual minerals in volumes that are anything smaller than a boulder will be difficult to resolve. On the other hand, the example chemical reaction above only applies in a 'closed' system (one where nothing can be added or removed). Real rocks are exposed to water that can bring in additional ions from outside or remove surplus ones. The net effect is that in general one or other mineral tends
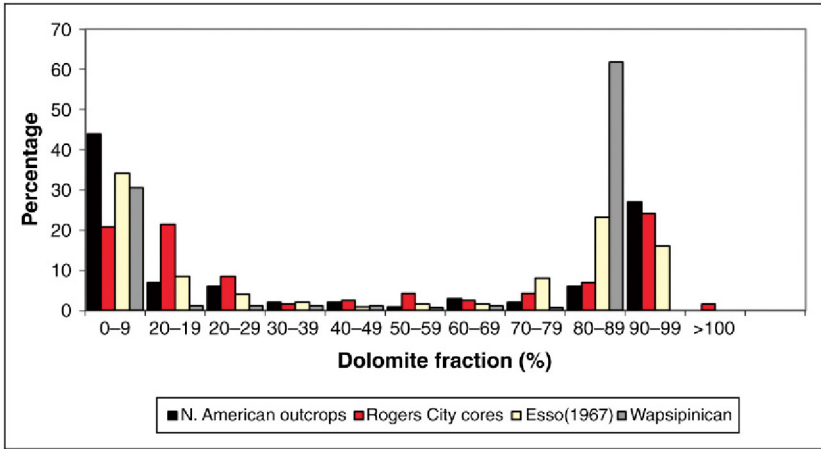
FIGURE 12.4   Limestone–dolomite ratios that have been reported in the literature. The 'dolomite fraction' is in reality based on the Mg/Ca ratio with a value of 1 referring to more than 50% Mg.

to form the reservoir. This is not to say that both minerals cannot be found in the same field, they commonly are but typically only one of them provides most of the pore space. Figure 12.4 reinforces this 'limestone or dolomite' tendency. This shows histograms of carbonate composition from several publications. The source of this information varies from hand-specimens collected from outcrops to cores cut in a single field but in each case it turns out that the mineralogy is far more likely to be calcite or dolomite than something in between.

The reader may reasonably ask why the reaction stops at dolomite and does not go all the way to a mixture of magnesium carbonate and calcite. Magnesium carbonate exists in nature as the mineral magnesite. But it is relatively rare compared to calcite and dolomite and is typically found in metamorphic rocks. It therefore does not often trouble petrophysicists. It is however a stable compound with a similar crystal structure to calcite and dolomite and it is even denser thanks to the small magnesium ion allowing a more compact structure (approximately $3\,g/cm^3$). But, of the two magnesium-containing carbonates, dolomite is thermodynamically the most stable. For those who know anything about chemical thermodynamics the heat of formation for dolomite is about – 133 kJ/mol compared to – 68 kJ/mol for calcite and – 64 kJ/mol for magnesite (crudely speaking the more negative the heat of formation the more stable the compound). These figures tell us that any mixture of calcium, magnesium and carbonate will form as much dolomite as possible. If the original system started with more magnesium than calcium the end result would be dolomite plus magnesite but this situation rarely if ever happens in sedimentary rocks. In the Universe in general and the Earth in particular magnesium is far more abundant than calcium (it is a lighter element) but if attention is restricted to the crust the relative abundances are reversed and for every magnesium atom there are two

calcium atoms. On balance then chemistry produces as much dolomite as it can and the left over calcium forms calcite.

Before getting too carried away with global geochemistry it is time to get back to logs. A representative collection of physical properties is given in Table 12.5. As was discussed in Section 1.7 different sources do quote different values, but for carbonates they normally agree to within 5%. Firstly, notice the density difference: if porosity is going to be calculated from density we better be sure which carbonate mineral we are dealing with. For example, a density of 2.6 g/cm^3 corresponds to a porosity between 13 pu and 14.5 pu in dolomite but only 6.5 pu in limestone! Moreover, the consideration of thermodynamic stability discussed earlier suggests it is much more likely to be one or other, than something in between. This is different to typical clastic rocks where there might be quite a range of compositions but the resulting grain densities typically form a mono-modal distribution (in other words they cluster around a single peak). In the worst case using a fixed grain density based on an average value in a sandstone may not give the right answer at every depth but the average porosity for an interval will be reliable. With carbonates if we assume calcite when the matrix is actually made up of dolomite we will systematically under-estimate the porosity.

12.4.2 Density–Neutron Cross-plot in Limestone–Dolomites Systems

Fortunately, some combinations of logs can normally distinguish calcite from dolomite. Remember that the standard form of the density–neutron cross-plot found in a chart-book invariably includes limestone and dolomite 'lithology lines'. These show that for a particular density, pure water-bearing dolomite has a neutron porosity 10 pu or more higher than a calcitic rock. Conversely, a calcitic rock will have a density that is about 0.25 g/cm^3 lower than the dolomite at a constant neutron porosity. These differences – which do vary slightly between different tools – are easily large enough to be resolved (their exact magnitudes may also depend on whether the neutron porosity has been environmentally corrected).

The differences are driven partly by the density difference between calcite and dolomite but this is enhanced by the difference in the response of the neutron porosity. Since this is partly a revision section we will consider the reasons why in a bit more detail. Firstly, consider a water-saturated limestone consisting of pure calcite and water. Assume that the hole is good, the neutron has been recorded with a lime matrix and that it has been environmentally corrected. The density and neutron values should then plot on the limestone line on the cross-plot and on a conventional log plot the curves will overlay.

Now consider what happens if the calcite is replaced by dolomite with the porosity remaining the same. Clearly, the density increases because calcite has been replaced by a denser mineral. The actual increase is easy to calculate and drops from 0.14 g/cm^3 for tight rocks to 0.09 g/cm^3 at a porosity of 35%. In terms of the cross-plot this means the point moves vertically downwards.

Slightly more difficult to predict is the change to the neutron porosity. In terms of the chemistry of the system two changes occur on replacing calcite with dolomite.

1. Half the calcium atoms are replaced by magnesium a lighter element.
2. The higher density of the dolomite means there are more atoms in the volume of investigation of the neutron tool compared to the 'limestone'.

Calculating exactly what these changes will do to the neutron porosity is difficult but it is possible to qualitatively predict what happens. Recall, that the tool works by slowing neutrons to thermal velocities and then allowing them to diffuse away from the source. Hydrogen is particularly good at slowing neutrons but other elements can still do this and the lighter the element the better. Dolomite is related to calcite by replacing half the calcium (atomic mass 40) with magnesium (atomic mass 24). At the same time because dolomite has a higher density, there will more dolomite 'molecules' in a particular volume than calcite. The volume that we are interested in is the volume of investigation of the tool. So the net effect is not only half the heaviest element replaced by a lighter one but more of the carbonate atoms, which are lighter still, are packed into the tool's volume of investigation. Together, these mean dolomite is more effective at slowing down neutrons and that will ultimately result in a higher neutron porosity being measured. Inspection of any density–neutron chart shows that the pure dolomite point is indeed offset to a positive value by a few porosity units (1–3 pu is typical).

Now consider a constant density ($2.6 \, \text{g/cm}^3$ say). This corresponds to porosity in a calcitic rock of 6.5% and since this is water-bearing limestone the neutron porosity is also 6.5%. The porosity for dolomite is 13%, but the neutron porosity should be this plus the extra resulting from replacing calcite. This gives a final neutron porosity of 15–16%. Is this what is observed in practice? Actually the two-detector thermal neutron tool, that is still the most common neutron porosity tool, typically reads considerably higher than this. The chart in Fig. 12.3, for example suggests the neutron porosity corresponding to $2.6 \, \text{g/cm}^3$ dolomite is actually about 20 pu (Schlumberger's chart CP-1c gives an even higher value of 22 pu). Clearly, the neutron porosity is responding to more than just the introduction of magnesium and some extra carbon and oxygen atoms. Discussion of this will be deferred to the end of the section.

To summarise, replacing calcite by dolomite results in both the density and the neutron porosity increasing. In terms of the chart dolomite plots below and to the right of limestone and the equivalent situation on the log is that the neutron curve moves to the left and the density to the right. This is best shown with an example. Figure 12.5 shows the basic open-hole logs recorded over a 200-m interval in an oil well drilled in the Kirkuk area of Iraq. This is a classic carbonate province. The density and neutron curves are plotted on standard – limestone compatible scales – and the neutron porosity was set up for a limestone matrix. In water-bearing limestone the two curves are expected
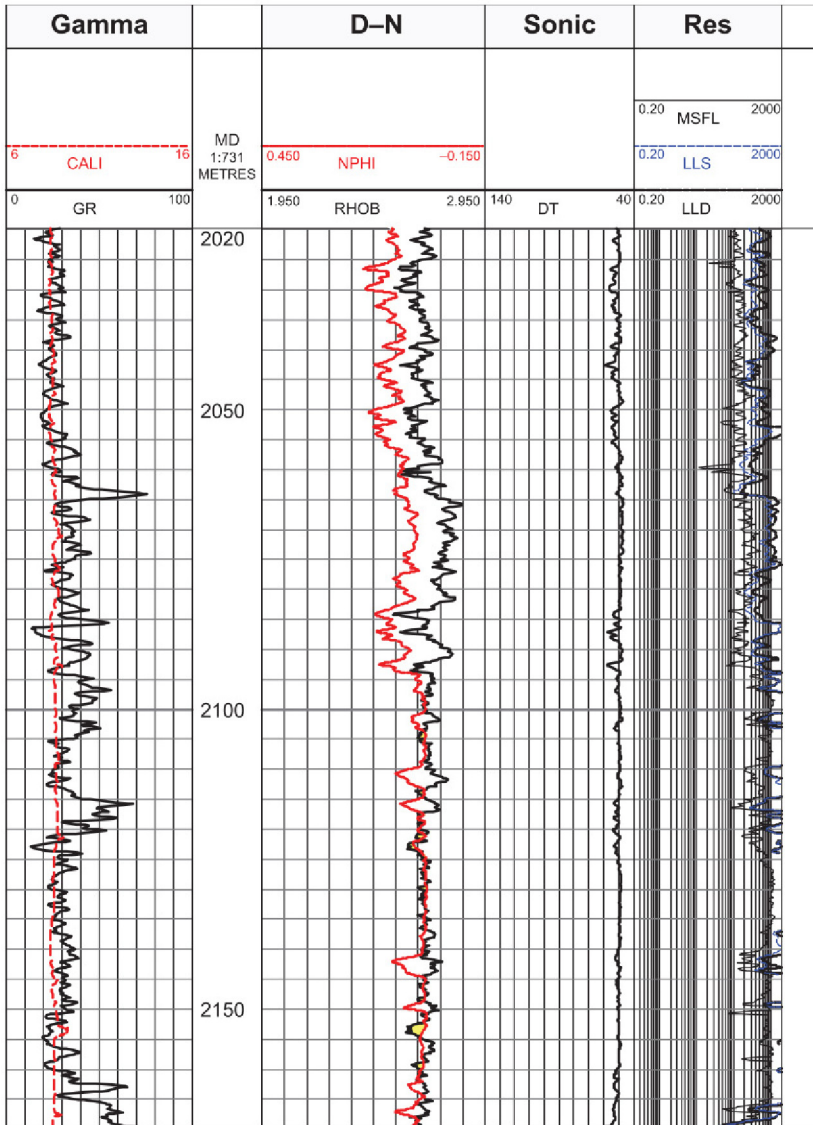
FIGURE 12.5 A 150-m interval of conventional logs from a carbonate reservoir in the Kirkuk region of Iraq. The mineralogy changes from dolomite above 2095 m to dominantly calcite below. The change is most obvious from the separation between the density and neutron curves.

to overlay and this is exactly what is observed below 2095 m. Furthermore, in the limestone the density is more or less constant at about 2.69 g/cm³ so the limestone has almost no porosity associated with it. Above 2095 m the curves show a positive, shale like separation equivalent to 12 pu. This is the signature of dolomite.
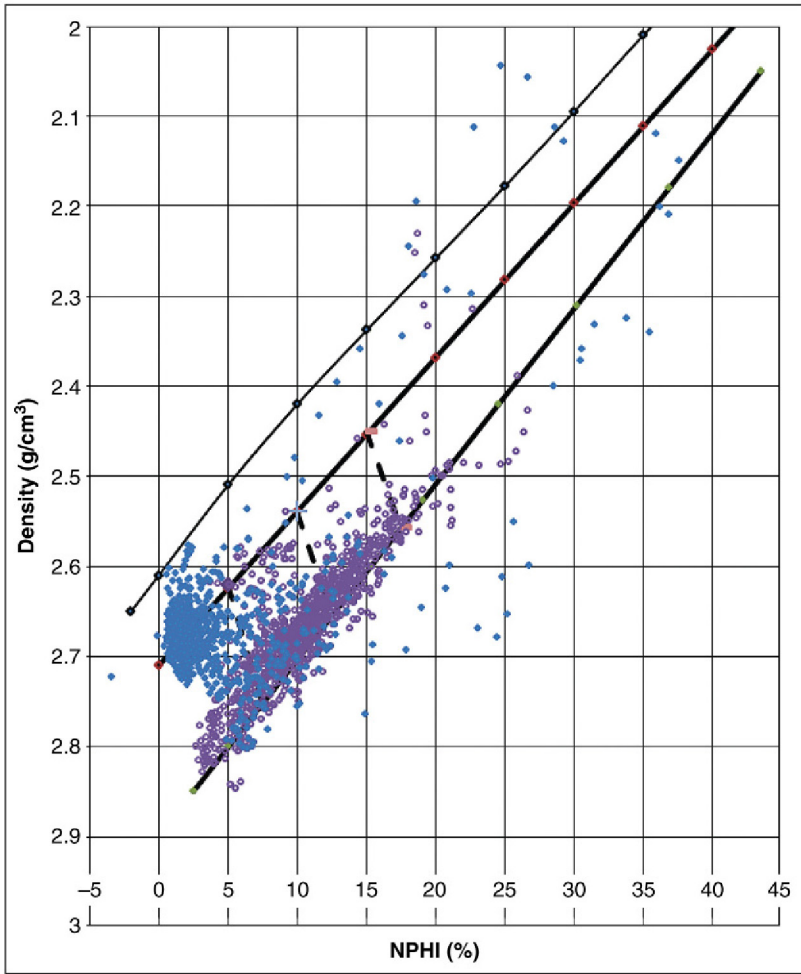
FIGURE 12.6 A density–neutron cross-plot of the logs shown in Fig. 12.5. The points have been colour coded by mineralogy: calcite in dark grey (blue in the web version) and dolomite in light grey (purple in the web version).

Figure 12.6 shows the density–neutron curves shown in Fig. 12.5 plotted on a standard density–neutron cross-plot. This really just reinforces what had been observed on the logs: the interval shown tends to be either tight calcite or dolomite with not much in between. Furthermore, the porosity is almost entirely associated with dolomite. Careful inspection of Fig. 12.5 however shows the dolomite line is shifted less to the right than was the case in Fig. 12.3 say. In fact the dolomite line is close to where the simple model discussed earlier predicted it should be. Why is this? The logs in Fig. 12.5 were actually recorded using an earlier type of neutron tool known as the SNP, which used epithermal neutrons.

These were discussed in Section 5.3.5 where it was stated that they had the advantage of not being strongly absorbed by some trace elements. Conventional thermal neutron tools give higher neutron porosities in dolomites than epithermal tools largely because they are responding to impurities which are almost always present in dolomite rocks.

12.4.3   Other Cross-plots

Chart books include blank cross-plots for every combination or sonic, density and neutron curves. They all include lithology lines so the implication at least, is that any pair of these tools could be used to distinguish calcite from dolomite. In practice however that is not generally the case. The sonic–density and neutron–sonic cross-plots for the same interval discussed earlier are shown in Fig. 12.7. The sonic density cross-plot does not allow the two minerals to be distinguished because for all but the lowest porosities limestone and dolomite have similar sonic and density responses. This could have been anticipated by simply inspecting Fig. 12.5 where it can be seen that what little variation there is in the sonic is more or less mirrored by the density curve.

The neutron–sonic cross-plot is more promising as in this case it does a good job separating the minerals. Inspection of the logs shows this separation is driven largely by the neutron porosity. Although in this case the cross-plot could be used to distinguish the limestones and dolomites, as it stands, it could

Neutron–sonic
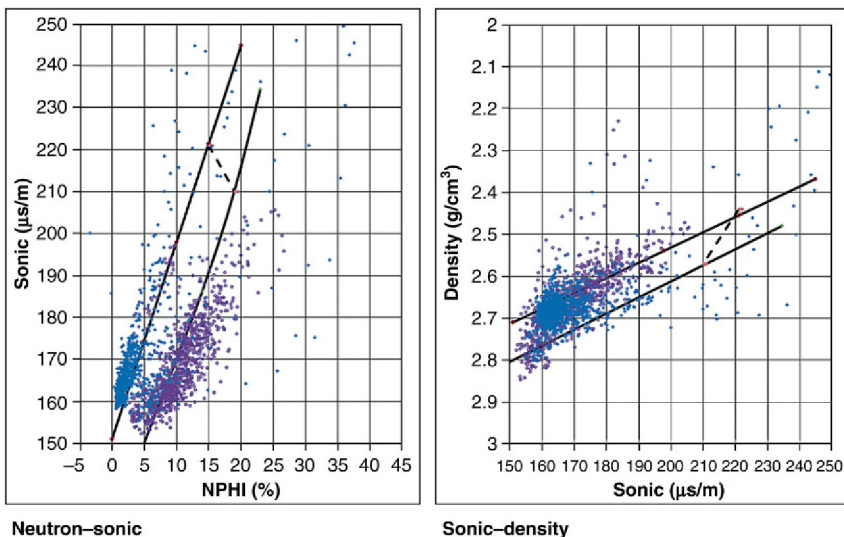
Sonic–density

FIGURE 12.7   Neutron–sonic and sonic–density cross-plots for the logged interval shown in Fig. 12.5. The lithology lines are for limestone and dolomite (dolomite is always the lower line). The dashed line corresponds to 15% porosity.
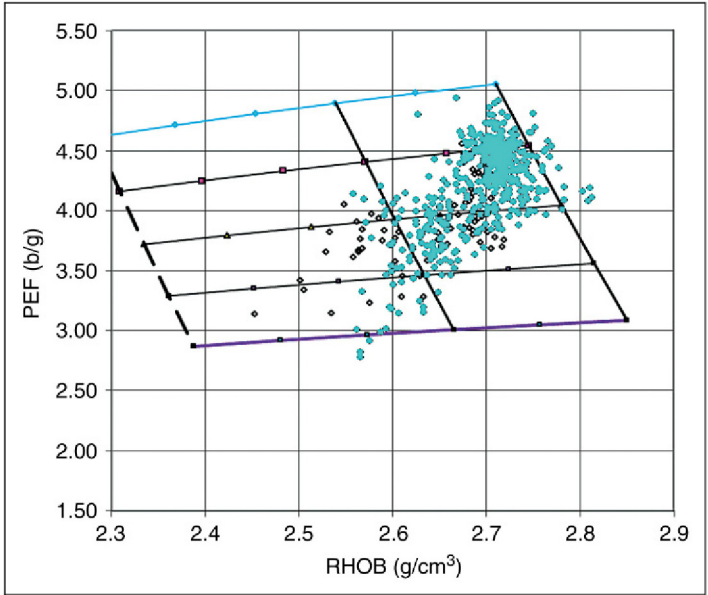
FIGURE 12.8   Density–PEF cross-plot for the logs shown in Fig. 12.9. The light grey (blue in the web version) points correspond to lower gamma activities. The light grey (blue in the web version) line at the top of the cross-plot is for pure calcite (or limestone). The dark grey (purple in the web version) line at the bottom is for pure dolomite, the intervening lithology lines are for limestone dolomite ratios of 25, 50 and 75%. The porosity lines are for 0, 10 (solid) and 25% (dashed).

not be used to estimate porosity. The 15% porosity line is well beyond most of the points suggesting that porosity never gets close to this value. The density–neutron cross-plot however shows that in the dolomite porosity reaches at least 15%. The problem in both cases is the sonic log which as discussed in Section 7.3.3 does not respond to porosity alone but is also influenced by the fabric of the rock. The overlays for the cross-plots below use the Wyllie time average equation and that is clearly not the appropriate equation to use in this case.

The density–neutron cross-plot is the preferred one of the three possibilities because both the input logs respond to the chemistry of the formation and are largely unaffected by the rock fabric. Earlier in the chapter the density–PEF cross-plot was introduced and as PEF is another measurement responding only to chemistry that too looks promising. The way the overlay is constructed was described for a three-component mixture consisting of quartz, pyrite and water but by substituting densities and U-values for calcite and dolomite an overlay can be generated for limestone/dolomite mixtures (Fig. 12.8). In fact of course the PEF curve should be able to distinguish limestones and dolomites on its own.

Unfortunately, the logs shown in Fig. 12.5 pre-date the PEF curve but a more recent set from a formation consisting of complicated mixture of carbonates and shales is shown in Fig. 12.9. The shales have the classic signature of a high gamma activity and they also have relatively low resistivity, so they are easy to
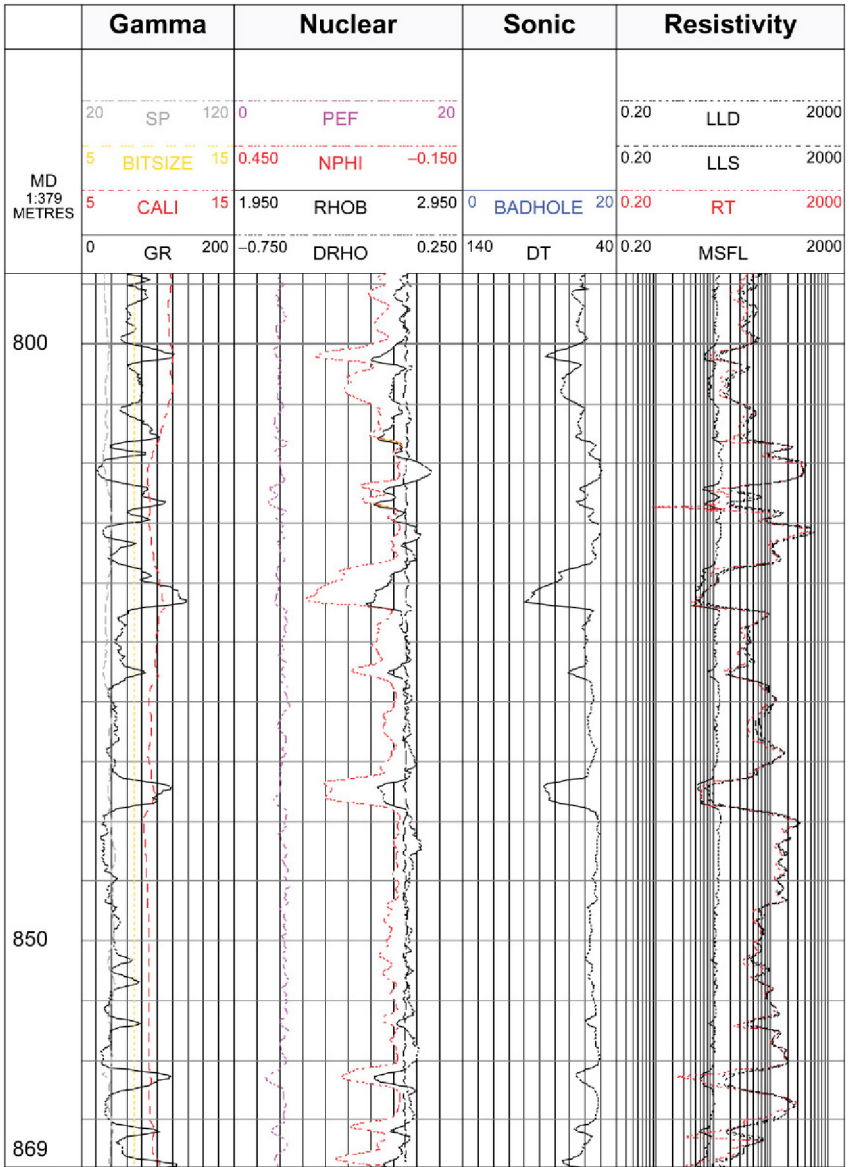
FIGURE 12.9 Conventional logs from a Devonian carbonate.

identify. The carbonate mineralogy was described as 'limestone' but the density–neutron response shows a positive, dolomite-like separation.

The density–neutron cross-plot is shown in Fig. 12.10 and as a thermal neutron tool was used, the dolomite line is shifted further to the right than in the previous example. There appears to be a continuous range of compositions from limestone to dolomite and to make matters more confusing the trend continues
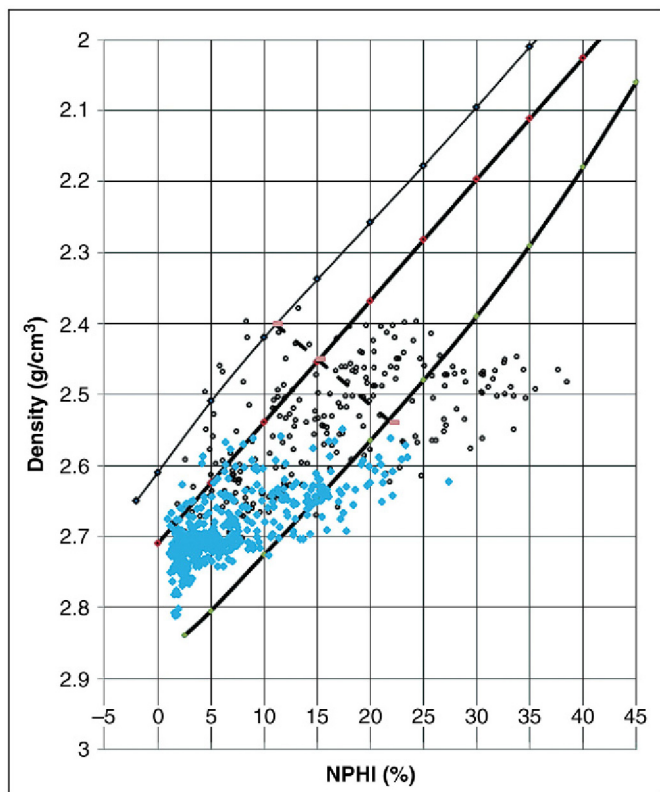
FIGURE 12.10 Density–neutron cross-plot for the logs shown in Fig. 12.9. The light grey (blue in the web version) points correspond to lower gamma activities.

into the shales (that have been identified by high gamma activity). Is there really a continuous range of calcite–dolomite or is it really pure limestone as suggested by the cuttings description? The density–PEF cross-plot confirms the continuous range of compositions suggested by the density–neutron response. In reality the formation probably consists of thin beds or even inter-bedded nodules of limestone and dolomite that are simply too thin to be resolved.

12.5 GEOCHEMICAL TOOLS

Geochemical tools are an obvious solution to complex lithology because they can identify and to an extent quantify many of the major components of a sedimentary rock (Section 5.7). They are not a universal panacea, however and some care needs to be taken when using them for interpretation. An example of a geochemical log showing some of the outputs of particular relevance to this chapter is shown in Fig. 12.11. The basic outputs are, the relative amounts by *weight* of the oxides of the elements that can be detected by the tool. In Fig. 12.11
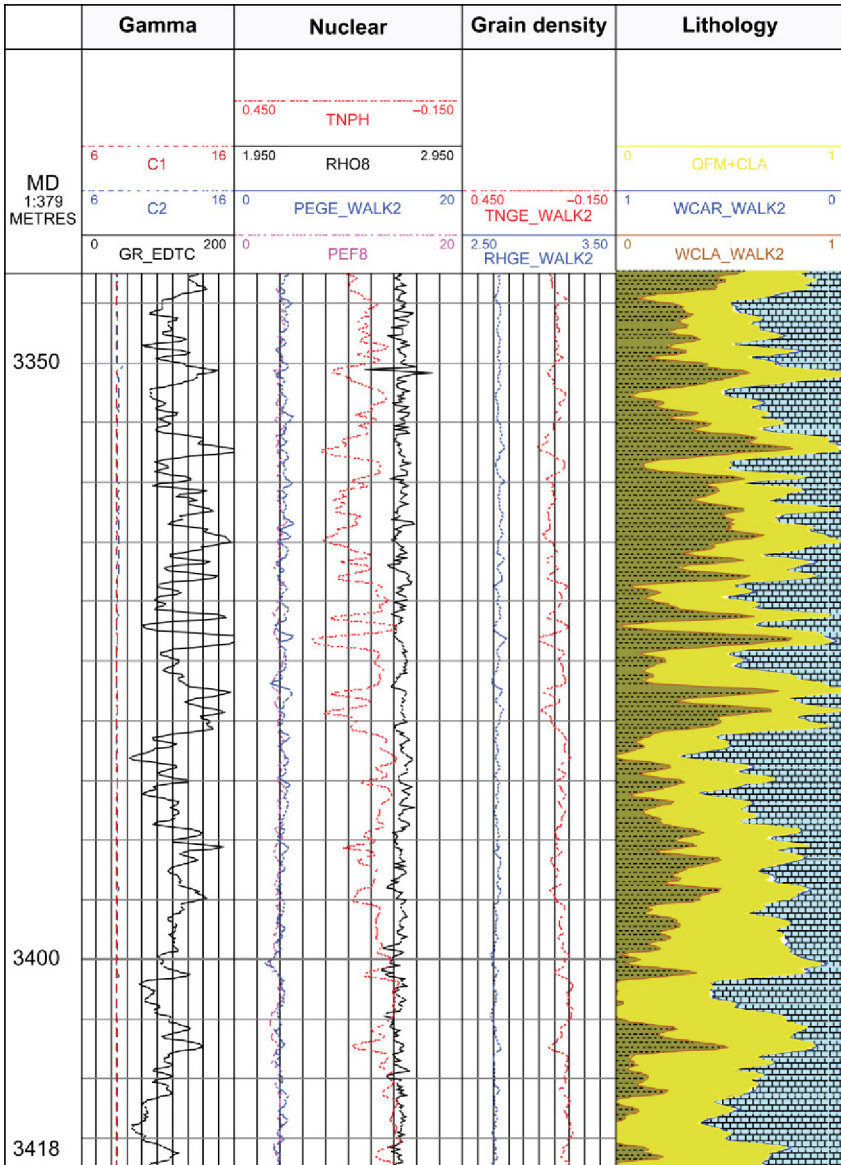
FIGURE 12.11   Some of the outputs from a geochemical tool from an 80-m interval of mixed carbonates and shales. The right-hand track shows the principal rock types: clay (left), quartz and feldspar (middle) and carbonate (right) inferred from the elemental ratios. Track 4 shows the grain density – on a scale of 2.5–3.5 g/cm^3 – inferred from this makeup and the equivalent neutron matrix porosity. Also shown are the density–neutron curves in track 3 ('nuclear').

(and also Figure 5.22) these have been used to estimate the relative amounts of some key mineral groups. In the case of Fig. 12.11 these are carbonates, clays and 'QFM' (i.e. quartz, feldspar and mica). The model used to find these also estimates the amount of anhydrite, pyrite, coal and other evaporites, but in the particular example these other components are insignificant.

The dry weight fractions can also be used to estimate grain density and some other matrix properties such as the U-value (and therefore PEF) and the neutron porosity of the matrix. The curves for matrix density and the neutron porosity are also included in Fig. 12.11. In principle this is a solution to the problem stated in the introduction to the chapter: to find the matrix parameters at each level in the well. In practice, this depends on how reliable these curves are. They are based on three steps, each of which becomes successively more interpretive.

1. The raw spectra are converted to elemental ratios.
2. The elemental ratios are converted to dry weight fractions of the oxides.
3. The dry weights of the oxides are combined to give the matrix properties.

The raw spectra are subject to some noise and the processing steps are based on assumptions and approximations. Providing these are reasonable there is no reason to doubt the accuracy of the parameters, but ideally one would use core data to produce a field- or formation-specific algorithm to generate the parameters.

Chapter 13

Thin Bed Pays: Dealing with the Limitations of Log Resolution

Chapter Outline

13.1 INTRODUCTION

At various points in this book it has been noted that few logs can properly resolve features thinner than about half a metre and most do worse than that. Remember that in this book 'vertical resolution' refers to the thinnest bed that the tool will measure the property of. A log may – and probably will – reveal the presence of something much thinner but the value it measures will be an average for the bed and whatever is either side of it. Unfortunately, it is the properties of the bed alone that are needed to calculate petrophysical properties.

Attempting to do quantitative analysis when log resolution becomes an issue is fraught with danger. In most cases simply ignoring the issue and just applying the tools and techniques described in Chapters 6–9 to the raw logs produces an interpretation that is pessimistic. This is because the log measurements have effectively diluted the reservoir lithology with non-reservoir material so that, for example, an interval that actually consists of 50% of high-porosity sandstone and 50% non-reservoir claystone is interpreted as a homogeneous but very poor quality shaly sand. This issue has been recognised for many years and interpretation methods have been developed to recognise and deal with it. There are two dangers with these:

1. The methods implicitly assume that log resolution is a problem and they will produce a result, which accounts for that. If the logs really are just responding to a poor-quality reservoir, the interpretation will be hopelessly optimistic.
2. These methods should not only estimate the porosity and water saturation of the reservoir lithology, they should also quantify what it's volume fraction is

at any depth. The end user must understand that although the interpretation is indicating good quality sandstone this only represents a fraction of the total volume at any depth.

Point 1 is a warning similar to the last paragraph of Chapter 8, it is very tempting to invoke log resolution as the cause of an apparently poor-quality reservoir. Unfortunately, the logs may be responding to just that: a poor-quality reservoir. Point 2 simply says that we should make sure that the results of such an analysis are being used appropriately. Up to this point in the book net-to-gross has been considered a single number that applies to an interval of at least several meters, but in these reservoirs it is a continuously varying curve with a different value at each depth increment (compare with Section 2.8).

In the last 20 years or so a large number of logging tools have appeared that can help identify situations where log resolution is an issue. The most obvious of these are image logs which with their twin attributes of high resolution and the ability to produce a 'picture' of the borehole wall can provide a lot more confidence that conventional log resolution is an issue. Even some logs, with relatively low vertical resolution such as NMR, can distinguish between a homogenous shaly sand and a heterogeneous mixture of sandstone and shale.

13.2   THE PROBLEM OF LOG RESOLUTION

When different lithologies are concentrated in volumes with dimensions that are comparable or smaller than the volume of investigation of the logging tool, the measurement will give a value intermediate between the highest and lowest values of the pure lithologies. The precise value will depend on the mixing law for the property of interest, as well as the way the tool works and probably the way the different lithologies are distributed (clasts, thin beds, lens etc.). It is difficult enough to calculate what the log would read knowing all of this but in log analysis it is the inverse problem that needs to be solved and in general; there is no unique solution to that. It is worth stating explicitly what we are trying to do here:

1. Determine how much of the tools volume of investigation is occupied by the reservoir lithology.
2. Determine what different logs would read in a thick bed of that lithology.
3. Convert those log readings to petrophysical properties.

Some well-known examples are the so-called 'thin bed pays' in which reservoir quality sand and shale are present in alternating sheets which are laterally extensive but much thinner than the resolution of conventional logs. A cartoon illustrating this type of reservoir is shown in Fig. 13.1.

The diagram shows the distribution of the lithologies, but gives no information on their properties. In general these will vary from bed to bed and even within individual beds. Fig. 13.2 shows a short interval (16 m) of alternating
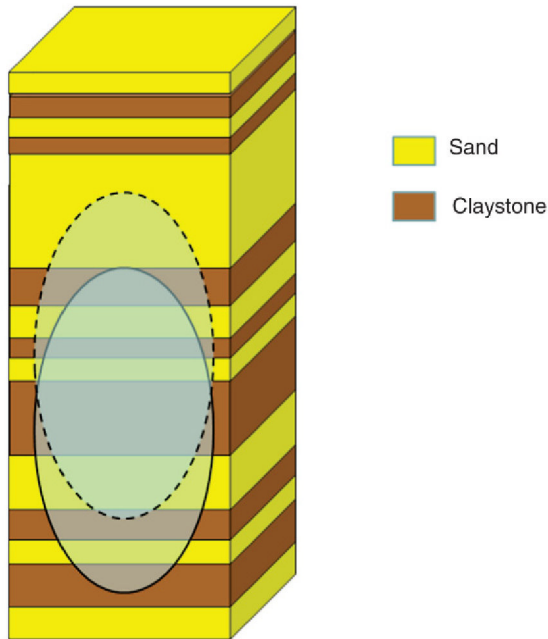
FIGURE 13.1   A block diagram showing a series of inter-bedded sands and shales covering about 1 m. The ovals represent the volume of investigation of a logging tool, they are separated by the depth increment of the log.

beds. For the sake of argument let us assume the log is a gamma ray and the beds consist of sand, with a relatively low gamma activity and shale.

The log, which has a 0.1 m depth increment, has been generated by simply taking a moving average over a 0.6-m interval. In reality the response will be more complicated than that as it involves a trade-off, with increasing distance from the tool, between a larger volume of formation against a lower number of gamma rays surviving to reach the detector. The influence of the borehole further complicates the response. Here we are more interested in the consequences of the tool's vertical resolution and the moving average model is as good as any for showing that.

Some points to note from Fig. 13.2 are:

1. Any bed over 0.3-cm thick causes a response in the tool (in fact even 0.2-cm features are picked up on the log, but in reality these would be difficult to see above the natural variation in gamma activity).
2. A bed needs to be about 1-m thick to be properly resolved (in this simple model a thinner bed can be resolved but in reality the movement of the tool and the necessity to account for natural variation will come into play).
3. A system with thin – 10 cm – beds of sand and shale gives a uniform gamma activity of 60 api and is indistinguishable from a uniform bed with that activity.
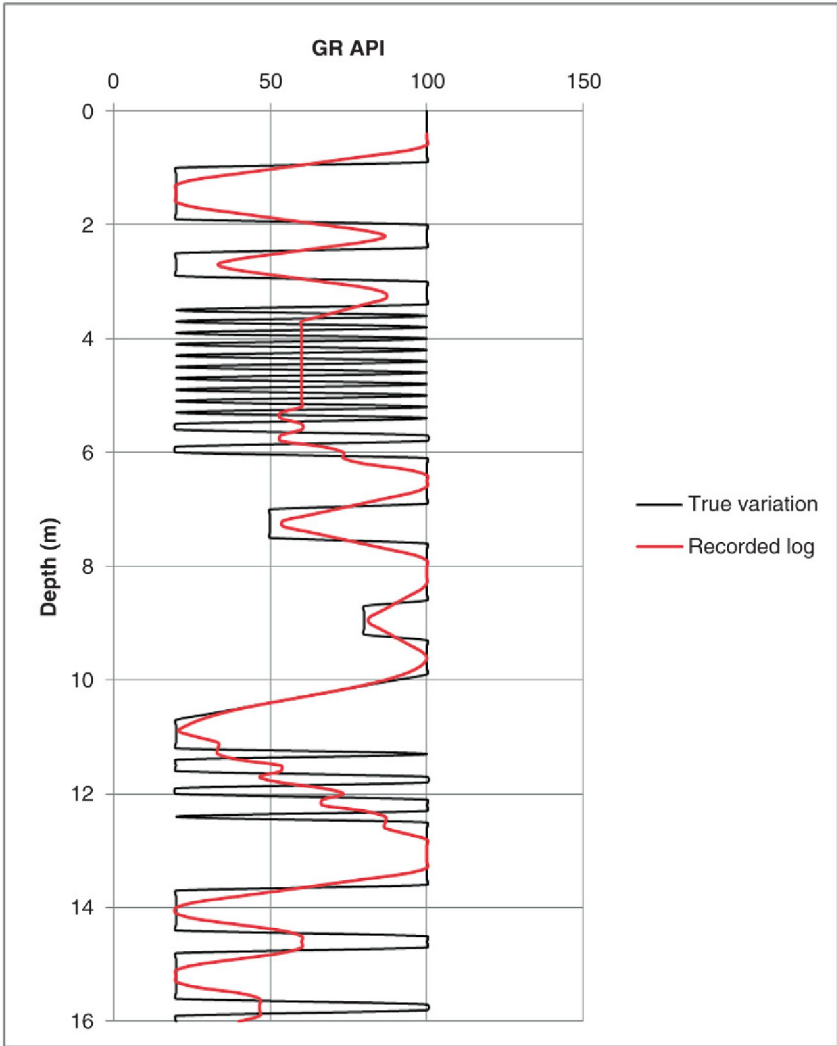
FIGURE 13.2 The response of a hypothetical gamma-ray tool to a system of thin sands and shales. The true variation of the gamma activity is shown in black, the response of the tool is in grey (red in the web version).

4. The gradational gamma reading at the top of the bed at 10 m really is caused by a gradational change. The gradational base is the result of increasing numbers and thicknesses of thin shales.

All this serves to illustrate that a particular log shape does not have a unique explanation and there is no unique solution in terms of sand response, shale response and sand–shale ratio. In order to make any progress assumptions have to be made and the reliability of the interpretation is largely determined by how

realistic these assumptions are. This is a very important consideration that directly pertains to the 'first danger' highlighted in the introduction.

13.3 THOMAS-STIEBER METHOD

One of the first and still most widely used techniques for dealing with thin beds was described in a paper published in 1975 by E.C.Thomas and S.J.Stieber. This was a time when all the conventional logs were available but neither image logs nor NMR were in commercial service. Thomas and Stieber were particularly concerned with tertiary sand-shale systems in and around the Gulf of Mexico. The specific example they used to illustrate the problem and how the method works was a Miocene reservoir from South Louisiana. They reproduced a photograph of approximately 15 ft. of core (4.6 m), which showed individual sand thicknesses between 1 in. and 1 ft. (2.5–30 cm) and shale thicknesses up to 3 cm. In other words well below log resolution.

The idea is to re-allocate the homogeneous mixture of sand and shale seen by conventional logs into individual sand and shale components, which are assumed to be present as thin beds. The method makes no claim to be able to quantify the thickness of individual beds let alone their precise location.

Figure 13.3 shows the model on which the Thomas-Stieber method is based. The key assumption that underlies this technique is that the clean sand and
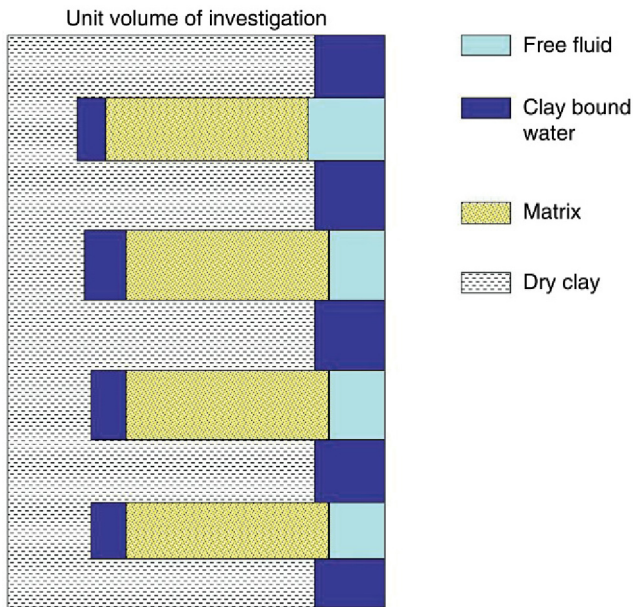
FIGURE 13.3   Basic model of a thinly bedded sand/shale system that was used to develop the Thomas-Stieber method. Shale or clay is assumed to be present in laminations and dispersed in the sand. Logging tools cannot resolve these components.

shale have constant petrophysical and physical properties. In particular they have known, constant total porosities. In their worked example, the total porosity of the sand was 33% and the shale 15%. Note that an individual sand bed can itself be a shaly sand: that is a homogeneous mixture of sand and shale. But the shale has the same properties whether it is dispersed in a sand or present as a thin bed. These correspond to the dispersed and laminated shales described in Section 2.5.3. So the model is designed to distinguish and quantify the dispersed 'shale' and the laminated shale present as distinct beds. The basic outputs from the process are:

1. Sand-shale ratio. (or 'net-to-gross')
2. The shale volume of the sand component.

The heart of the method is a ternary diagram whose end points are pure sand, pure shale and a hypothetical point where the inter-granular pore space of the sand is completely filled with shale. The individual points are defined by total porosity and shale volume, which come from the logs (although as formulated in the original paper it was actually sand volume that was used). The analysis simply involves converting the log point to volume fractions of sand, laminated shale and dispersed shale. The way these are defined is shown in Fig. 13.4. As originally presented sand volume was a linear function of gamma ray and porosity was calculated from density. But there is no reason why the method should not be extended to other shale and porosity methods and nowadays it frequently is (e.g. NMR).
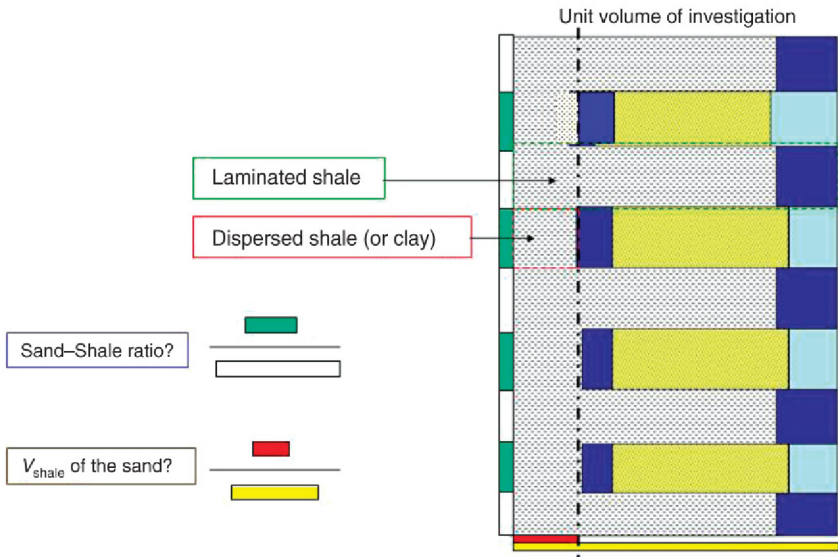
FIGURE 13.4 The Thomas-Stieber model showing how laminated and dispersed shale are defined as well as the sand fraction.

The two assumptions that underpin the method are:

1. There are only two rock types: clean sand and pure shale. These have fixed total porosities, which are generally different.
2. There is no change in shale character within the interpretation interval.

(In the original paper there are three other assumptions, which relate to the use of the gamma ray as a shale indicator).

An example of a Thomas-Stieber plot is shown in Fig. 13.5. The sand fraction, confusingly named 'Gamma', is on x-axis and total porosity is on y-axis. The pure sand point is on the RHS (Gamma=1) and in this case corresponds to a total porosity of 35%. The pure shale point is at the LHS (Gamma=0) and in this case the shale has a total porosity of 20%. The line connecting these points corresponds to laminated sand and shale and the distance along the line gives the sand–shale ratio (the black line in Fig. 13.5). Any points lying below this have some dispersed shale present as well.

The dispersed shale point has a porosity equal to the product of the pure sand porosity and the pure shale porosity (as all the inter-granular pore space is filled with shale with its own characteristic porosity). In the case of Fig. 13.5 this is $0.35 \times 0.2 = 0.07$. The gamma value for pure dispersed shale cannot be found from a simple formula and is normally chosen so that the three end points cover all the – valid –log points in the interval. The line connecting the pure sand point to the dispersed shale point gives the porosity of a sand with dispersed shale (the dark grey line in Fig. 13.5). Lines connecting these points to the pure shale point are lines of equal sand porosity – iso –porosity lines (dot-dashed) – and will vary from the pure sand value to the pure dispersed shale value (35–7% in this case).
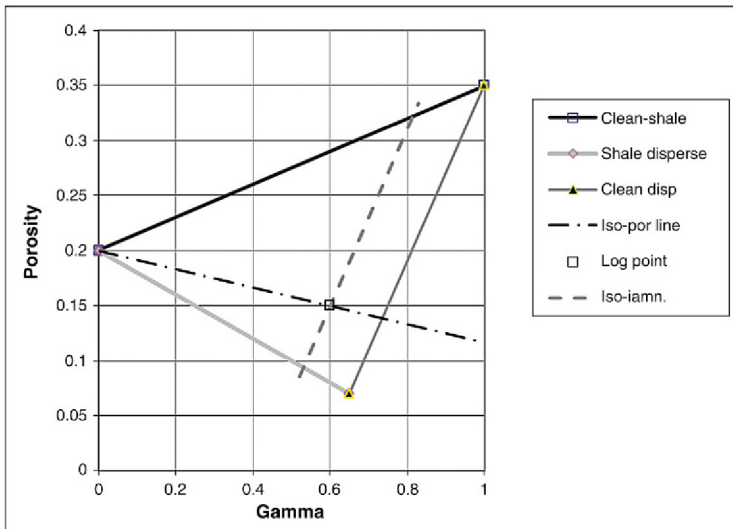
FIGURE 13.5   An example of a Thomas-Stieber plot (see text for a full explanation).

Any point within the triangle defined by these end points consists of laminated shale and shaly sand. The relative amounts are found by a geometric construction. Lines running parallel to the sand-dispersed shale point give the amount of laminated sand and shale (dashed). The total porosity of the shaly sand component is given by the iso-porosity line, on which the point lies. The example point on Fig. 13.5, with a porosity of 15% and a gamma value of 0.6 have a sand–shale ratio of close to 80%. The sand is very shaly however and has a total porosity of only 14%. The reader is recommended to take a ruler to Fig. 13.5 and confirm these values.

As a check on these figures the total porosity of the log point (0.15) should agree with the sum of the porosity contributions from the sand and shale components.

$$\begin{aligned} \text{PHIT} &= (\text{Porosity of sand}) \times (\text{Fraction of sand}) + (\text{Porosity of Shale}) \\ &\quad \times (\text{Fraction of shale}) \\ &= 0.14 \times 0.8 + 0.2 \times 0.2 = 0.15 \quad \text{as required.} \end{aligned}$$

Assuming the basic assumptions are valid the accuracy of the results obviously depends on the choice of the pure sand and shale points. Although the technique was developed to work in environments where thin beds dominate, there often are one or two beds that are thick enough for conventional log analysis to work. In this case these beds are used to define the sand and shale end points. If there are no thick beds then the end points are necessarily educated guesses. The triangle joining the end points should enclose all the valid log points, but beyond that there is a lot of latitude. Core data might help define the end points, particularly the sand point but the shale volume of the plugs needs to be carefully determined and since this is not something that core laboratories measure, one requires quantitative measurements of mineralogy.

Not shown in Fig. 13.5 is another part of the plot that deals with structural shale, in other words grains that are composed of shale. These will actually add to the conventional sand porosity and therefore pull points above the pure sand–pure shale line. For simplicity these have been ignored here and the interested reader is referred to the original paper. More importantly we have completely ignored hydrocarbon effects in the previous discussion, which, given that one of the primary inputs is a density log, is a major omission. In reality this will involve estimating saturation, probably using one of the shaly sand equations discussed earlier in the book (Section 8.6). The saturation can then be used to correct the density for hydrocarbon effects, which will shift the point on the plot thereby producing new estimates of porosity and sand–shale ratio. These will result in a revised estimate of S_w and so on. In principle this is easily accomplished but in practice it involves a lot of computation and complexity and there is a real danger of the whole problem becoming poorly conditioned. There are some other practical issues
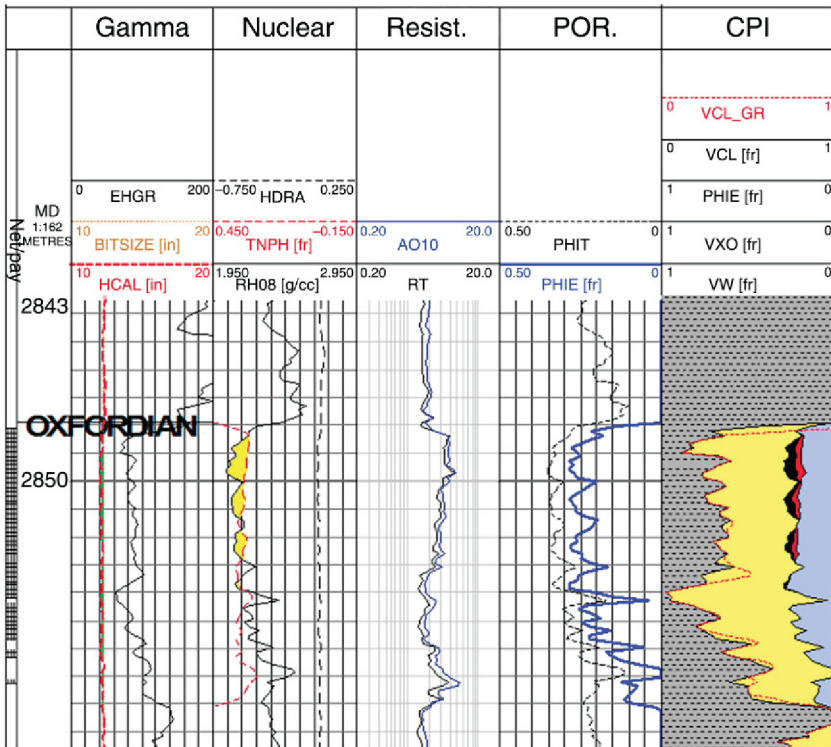
FIGURE 13.6 Raw and interpreted logs across what appears to be a very poor-quality, gas-bearing sand. Note the low resistivities (<5 Ωm) and the apparent lack of a gas effect on the density–neutron logs.

relating to the measurement of resistivity that will be discussed in the subsequent section.

To conclude this section we will look at an example. The logs and conventional interpretation for an apparently poor-quality gas sand is shown in Fig. 13.6. Although gas was sampled with a wireline formation tester, resistivities are low and the density-neutron logs show no gas effect. The conventional interpretation suggests the log responses are caused by a reservoir with a lot of dispersed clay (shale) that result in high water saturation. There is of course an alternative explanation that the log responses are caused by the sand being present as thin beds. The sand–shale ratio increases upwards but the individual sands never get thick enough to be properly resolved.

Figure 13.7 shows the Thomas-Stieber cross-plot produced from the logs in the 10-m interval highlighted by the shaded bar on the LHS of the log. Note there are no clean sands that are thick enough to be resolved and so the clean sand point is chosen to enclose the log points. The cross-plot suggests that the
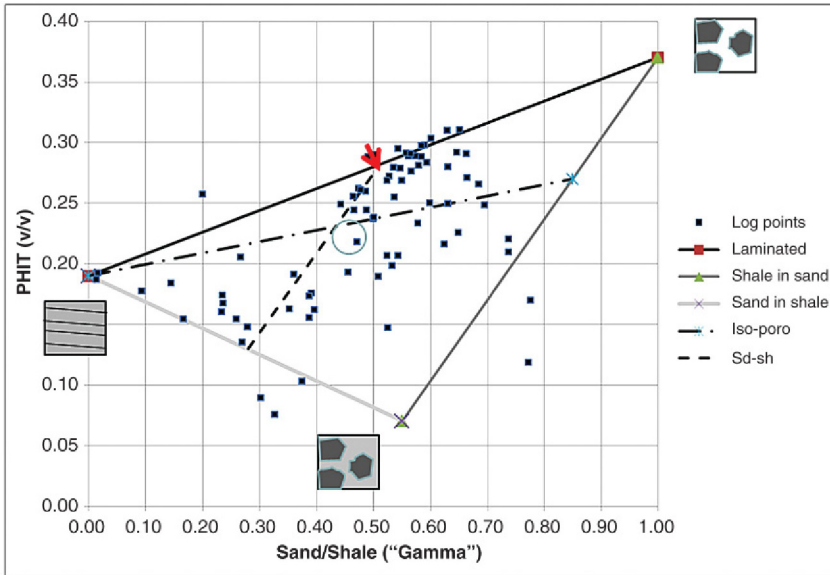
FIGURE 13.7 Thomas-Stieber plot generated from the logs in Fig. 13.6. The black (red in the web version) arrow represents equal thicknesses of thin shales and clean sands. The dot-dashed line (blue line in the web version) leading away from the arrow, generalises this to equal thicknesses of shale and shaly sand. The dashed line represents a total porosity for the sand of 17%.

reservoir is always laminated and that at log resolution the sand laminations make up at most 60% of the thickness but that there are intervals where the laminations consist of nearly clean sand. These are the points that lie along the sand shale line.

The Thomas-Stieber interpretation is shown in Fig. 13.8. The volume of lamination in track 3 refers to the shale component so that a VLAM of 67% means two-thirds of the reservoir consist of shale beds at that depth. VDIS is a measure of the amount of shale dispersed in the sand but it is not the same as the shale volume for the sand. A value of zero means the sand is clean but a value of 1 means all the inter-granular space is occupied by shale. In other words the sand component looks like the dispersed shale point in the Thomas-Stieber diagram. Together these curves can be used to find the shale volume of the complete system of sand and shale beds (this is the red curve in the right-hand track) and this can be used to find the effective porosity of the system. Of more interest is the effective porosity of the sand since this tells us how good the reservoir quality is. This is the solid curve in the porosity track with the higher values and it looks very respectable (exceeding 35% in the upper 3–4 m). Because some points lie above the shale–clean sand line the calculated effective porosity
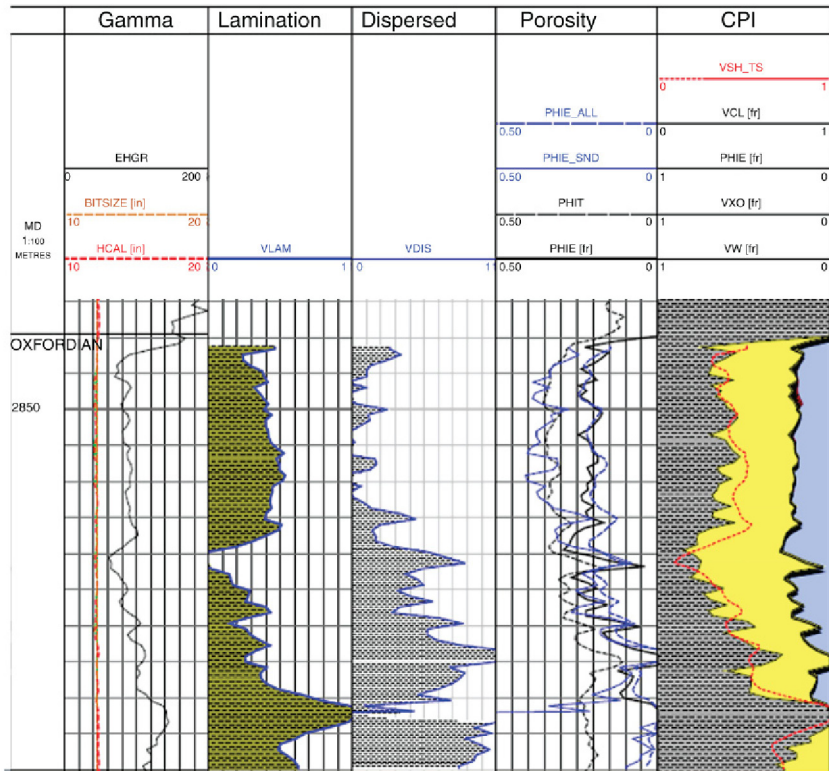
FIGURE 13.8 Interpretation of the shaly sand using the Thomas-Stieber approach. The fraction of shale layers is shown by the curve VLAM in track 3 and the proportion of dispersed clay is given by the curve VDIS in track 4. The effective porosity of the sand is shown by the grey curve in the porosity track which is 10 pu or more higher than the effective porosity calculated by the conventional method (black). The total effective pore volume is almost the same for both interpretations, however because even in the best reservoir sand only makes up about 60% of the total thickness.

slightly exceeds 37% in places, this is not really possible and is an indication of the limits imposed by the fundamental assumptions of the method. The total porosity of the sand is the value on the iso-porosity line.

It is now worth repeating the warning to proceed with caution. The Thomas-Stieber method is based on a very specific model (alternating sands and shales bounded by parallel planar surfaces). As with any log analysis technique, it will only work if the assumptions and approximations on which it is based apply. The converse does not apply: the fact that you can perform a Thomas-Stieber analysis does not mean the formation is thinly bedded. This may seem a trivial point but the method is specifically designed to deal with beds that are too thin to be resolved by conventional logs. To apply the method one needs independent

evidence that there is a chance that a thinly bedded system is present. This could come from:

1. Cores.
2. Image logs.
3. 'Modern' logs such as NMR and 3D-resistivity.
4. Shows and gas readings in an apparently shaly interval.
5. Some form of flow test (including wireline formation test data).
6. Offset well experience and general knowledge of the reservoir.

With the exception of a core, all of these techniques can give misleading results as well.

As it happens the reservoir shown in Fig. 13.6 really is as bad as the interpretation suggests. In other words all the analysis embodied in Fig. 13.7 was based on erroneous assumptions! In fact core photographs and thin sections from the same formation in a different well were shown in Fig. 2.10. Core plugs showed the clay was all present as dispersed shale and the conventional log analysis actually is a good representation of the reservoir!

13.4   RESISTIVITY AND SATURATION

The Thomas-Stieber method basically assumes that logs respond linearly to mixtures of sand and shale: that is the mixing law is a simple linear average. This is true for density and a good approximation for the gamma ray but is unlikely to hold for the resistivity measurement, which may be strongly biased towards the shale value (in fact deviations of the gamma ray from a linear response were considered in the original paper). This is particularly true of induction tools and unfortunately these may be the only type of tool that is available (because thin bed pays are often associated with deep-water un-consolidated sediments which require(d) oil-based mud to drill). Thomas and Stieber's worked example from South Louisiana was logged with a deep induction tool.

Induction tools cause current to flow in a plane that is perpendicular to the borehole. For low angle wells this is normally parallel – or close to parallel – with the bedding.

In a thinly bedded sand-shale system more of the current will flow in the more conductive component and the output of the tool will reflect this bias. When the sands are water bearing the contrast in resistivity is normally quite small and the tool gives a reasonable approximation of the true average but when the sand is hydrocarbon bearing the resistivity is strongly biased towards the shale value.

In order to get a valid Rt measurement to compute saturation in the first place may require some additional processing. For conventional induction tools forward modelling is probably the best approach. Resistivity tools are so well understood that it is possible to compute the tool's response to a particular arrangement of beds of varying thicknesses and resistivities. This is known as

forward modelling. The latest forward modelling programmes can cope with invasion, dipping beds and different sized boreholes filled with different fluids. To get the true resistivity of the beds involves constructing a model of the thin sands and shales and modelling the tool response. The resistivity of the sands is then adjusted until the tool response matches the actual measurement (Fig. 13.8). This can be quite computationally and manually intensive, however and until recently it was impractical to do this routinely.

If several different resistivity curves are available it may be possible to 'inverse model' the formation. In this case the resistivities of the beds are adjusted until a satisfactory match between the computed and the measured curves is obtained. Note that the location and thickness of the beds still have to be provided by the analyst.

When the Thomas-Stieber method was published, resistivity modelling would not normally have been a practical proposition (most logs were not even recorded digitally at that time). Hagiwara developed a simple model to allow the resistivity of a laminated system measured by an induction tool to be predicted. The inputs are sand–shale ratio, sand resistivity, shale resistivity and bed dip (relative to the tool).

The algorithm combines resistivities computed parallel and normal to bedding. These are given by

$$C_{\text{parallel}} = C_{\text{sh}}(1-L) + C_{\text{t}}L \qquad (13.1a)$$

$$R_{\text{normal}} = R_{\text{sh}}(1-L) + R_{\text{t}}L \qquad (13.1b)$$

Where L is the sand–shale ratio, R_{sh} is the resistivity (conductivity) of a shale bed and R_{t} (C_{t}) is the resistivity (conductivity) of the sand beds. (If you remember simple circuit theory these will be recognisable as the equations for parallel and series resistors). Although it may appear a bit clumsy to mix equations for resistivity and conductivity together this has the virtue of making the underlying reasoning clearer.

According to Hagiwara the conductivity measured by an induction tool is then given by:

$$C_{\text{ild}} = C_{\text{parallel}}(\cos^2(\theta) + C_{\text{parallel}} \times R_{\text{series}} \times \sin^2(\theta)) \qquad (13.2)$$

Where θ is the dip relative to the well. So, for the 'normal' situation of a vertical well drilling through horizontal beds, θ is zero and Eq. 13.2 becomes:

$$C_{\text{ild}} = C_{\text{parallel}} = C_{\text{sh}}(1-L) + C_{\text{t}}L$$

Since in a hydrocarbon zone $C_{\text{sh}} >> C_{\text{t}}$, unless the sand–shale ratio is very high, the measured conductivity – and therefore resistivity – will be close to the shale value, which is of course the problem we are trying to solve. In practice the equations would be used in a forward modelling mode. That is by adjusting C_{t} until a reasonable match with the log reading is obtained.

Modern '3D' induction tools such as the Baker 3DEX can actually produce a good estimate of R_t even in thinly bedded pays. These tools were specifically developed to deal with thin beds and first appeared in the late 1990s. They use the same principle as a traditional induction tool but have three sets of coils that are mutually orthogonal. One of these sets of coils induces currents to flow circumferentially around the borehole in the same way as a normal induction tool. The resistivity measured by this set of coils is known as the 'horizontal resistivity'. The other sets of coils induce currents that flow parallel to the borehole (there are two such currents at 90° to each other). These measure the vertical resistivity(s). In the 'normal' situation of a vertical well intersecting horizontal beds the latter currents are forced to flow through all the beds even if some are relatively resistive. All three measurements will agree in a homogeneous formation but in a thinly bedded formation with alternating high and low resistivity beds the vertical resistivity will be significantly higher than the horizontal.

An example is shown in Fig. 13.9 in which a homogeneous gas bearing sand transitions to the claystone seal through approximately 10 m of thinly bedded sands and claystones.

The upper few metres of Fig. 13.9 represent the claystone seal for the structure. Even here the vertical resistivity is higher than the horizontal, although only by a factor of 2–3. This is commonly found in claystones and shales and is a result of their laminated nature. They are anisotropic and resistivity differs depending whether it is measured parallel or perpendicular to laminations (this almost certainly explains why micro-resistivity logs often read consistently higher than deeper reading tools in shales).

The laminations could also be seen on an image log, although the colour scheme used previously does not have enough contrast in the low resistivities to show this. The response of these 'modern' logs therefore needs to be treated with care. Alternating high and low resistivity layers on an image log and a separation in horizontal and vertical resistivities is not a sufficient condition to prove thin beds much less thin bed pay. As pointed out in the introduction, invoking thin bed pays needs evidence from other logs or cores.

13.5   IMAGE LOGS

Image logs have been mentioned on several occasions in this chapter and appear in the last two figures, so it is high time they were discussed in a bit more detail. They are a class of measurements, which now exploit a number of different physical measurements to build a map of the borehole wall. Regardless of the physical property they exploit they all share the following characteristics.

1. The ability to make discrete measurements at several points around the borehole wall.
2. Some way of locating the measurements in space so that the image can be related to North and/or the high or low side of the hole.
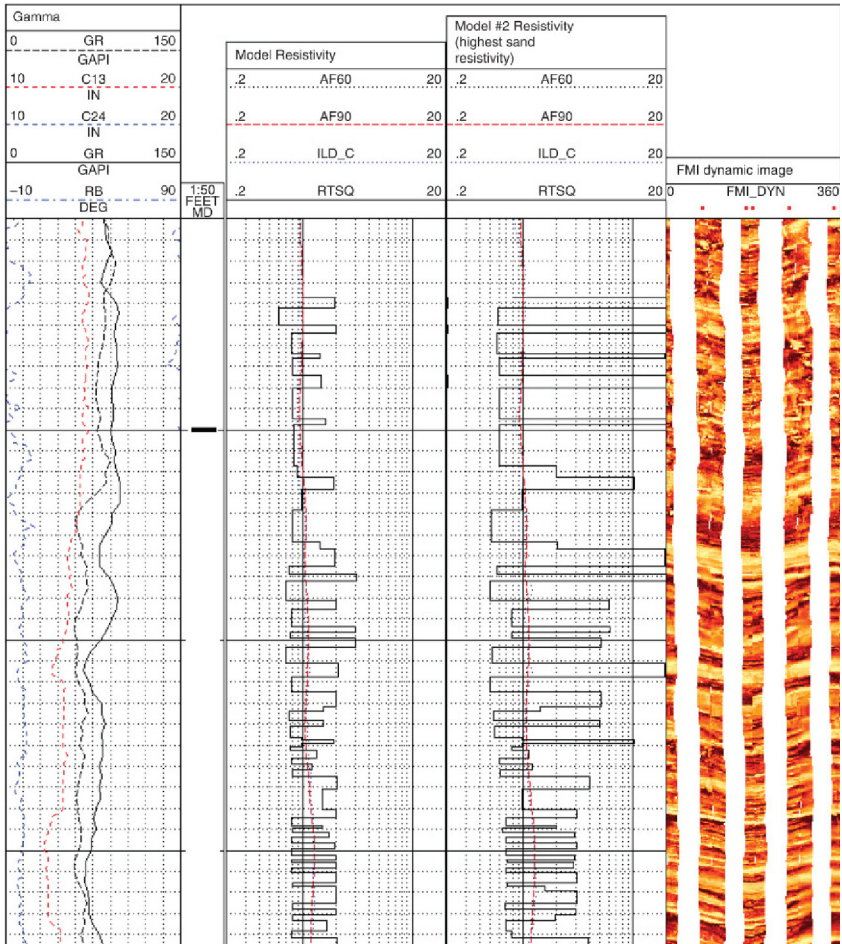
FIGURE 13.9   An example of thinly bedded pay from the Gulf of Mexico. The measured resistivity (grey dashed [red dashed in the web version]) is low and shows little variation and the density–neutron curves (not shown) showed a generally shaly character. The image log suggests the interval actually consists of thin sands and shales. The image allowed a model of the formation to be built (black squared curves) which could be used to model the induction tool (dark grey doted curve [blue doted curve in the web version]). The sand and shale resistivities in the model were adjusted until the computed induction curve (dark grey dots [blue dots in the web version)]) overlaid the original log. This showed resistivities in the thin sands could actually be as high as $10\,\Omega\,m$, which is characteristic of a gas sand.

Images can be generated using wireline or LWD logs. The former tend to have the additional property of very high vertical resolution and under good conditions can easily resolve features with a dimension of a centimetre or less. LWD measurements exploit logs that have an inherent directional volume of investigation such as the density. Unlike a wireline tool the rotation of the BHA allows the

measuring volume to sweep around the borehole circumference. Since most LWD images are created from conventional logs they do not have especially good vertical resolution. Nevertheless, they can still produce images that are quite capable of distinguishing thinly bedded systems form more homogeneous ones.

13.5.1 Electrical Image Logs

Electrical image logs have evolved from an earlier generation of tools known as 'dipmeters'. As the name suggests the property they exploit is resistivity and the image is a map of resistivity variation in a small volume between the borehole wall to a few centimetres beyond. The absolute value of the resistivity is not important as it is the variation that is exploited (although some tools can produce a high-resolution resistivity measurement). This is analogous to the difference between a medical X-ray and a density measurement discussed in Section 5.2. Although both techniques exploit the same physics the medical application is more interested in the spatial variation in count-rates than the absolute numbers.

Image logs evolved from well-established tools known as the dipmeters that were designed to measure the dips of bed boundaries and structural features such as fault planes. Any dipmeter has several caliper arms, ending in pads that are pushed against the borehole wall. The earliest tools had three, these were quickly supplanted by four-arm tools and by the 1980s several contractors were offering six-arm devices. The pad contained a small electrode array, which in most tools resembled a miniaturized laterolog type tool. So each pad produced a high-resolution curve: a typical depth increment was 2.5 mm (0.1 in.) although the inherent resolution of the measurement was lower than that. A typical pad with a single electrode system is shown in Fig. 13.11a and consists of a central electrode with a diameter of about 5 mm embedded in plastic. The latter ensures the electrode is electrically isolated from the rest of the pad (and the sonde). A pencil of current is emitted from the electrode and flows several metres into the formation before returning to the upper part of the tool body (Fig. 13.11c). Current is also emitted from the rest of the pad and it also flows back to the upper part of the tool, which is electrically isolated from the sonde by an insulated section. This current serves to push the measure current into the formation in much the same way that bucking currents force the measure current to follow the same path in a conventional resistivity tool (Section 5.6.3). Although the measure current flows deep into the formation it is mostly controlled by the resistivity immediately in front of the pad so, in practice, dipmeters are shallow reading tools with all that implies for log quality. Bed boundaries were detected by changes in resistivity and so there was no need to produce a calibrated resistivity measurement and in fact often the raw dipmeter log simply consisted of an un-calibrated number, which had no resemblance to resistivity or any other property.

An electrical image-logging tool simply added more electrodes to each pad so that instead of a single trace being produced at say four points around the borehole, multiple traces are produced from each pad. An example of an image

log pad is shown in Fig. 13.11b where there are 20 electrodes. In order to get more electrodes on a pad they are actually smaller than a typical dipmeter's. They are shown arranged in two rows that are slightly offset horizontally. This is common practice and ensures 100% coverage by the electrodes within the array although, obviously, quite significant gaps remain between the pads.

'Evolution' is the right word to describe the change from dipmeter to image tool as by the 1980s dipmeters were already appearing with two or three electrodes per pad and crude images were being produced. The increase to a dozen or more electrodes per pad that image logs use was just an increase in complexity. Normally the raw numbers are converted to a colour/shade for individual pixels from which the image is built. Figure 13.9 is an example of an image generated using a four-arm tool (each strip is the output from one pad). Most often an 'earth like' colour scheme based on yellows and browns is used for the image. Normally the darker colours represent lower resistivities, so that in the thinly bedded system shown in Fig. 13.9 the shales have the darker shades and the more resistive sands are pale yellow or light grey. This convention does mean that coals appear white however! In reality any colour code will do and more lurid colours are sometimes used to enhance particular features.

The type of tool described previously that is related to a laterolog, does have the major drawback that it needs a conductive mud to function. In recent years several tools have appeared that will work in oil-based mud. These use a variety of techniques to measure resistivity but all have very shallow depths of investigation both in terms of current path and what affects the measurement. The measuring systems are normally slightly larger than the embedded buttons described previously so that the vertical resolution is poorer and the number of electrodes per pad is generally lower (six or seven per pad is typical). The image log shown in the right-hand track of Fig. 13.10 was recorded using one of these tools.

13.5.2 Acoustic Image Logs

The first image logs to be developed exploited the contrast in acoustic impedance between the borehole fluid and the formation. The earliest devices could only be used in relatively shallow wells filled with simple fluids but they have been continuously improved over the years and now they can be run in quite hostile environments and in strongly attenuating mud, including oil-based mud (before oil-based electrical tools were developed, these were the only image logs available in OBM). The heart of any tool is a rotating transducer that emits pulses of ultra-sound. These bounce off the borehole wall and are detected by the same transducer. Modern tools make somewhere between 150 and 250 measurements for each rotation. Two properties are measured:

1. The 'two-way time' for the sound to leave the transducer, bounce off the borehole wall and travel back to the transducer.
2. The amplitude of the returning pulse.
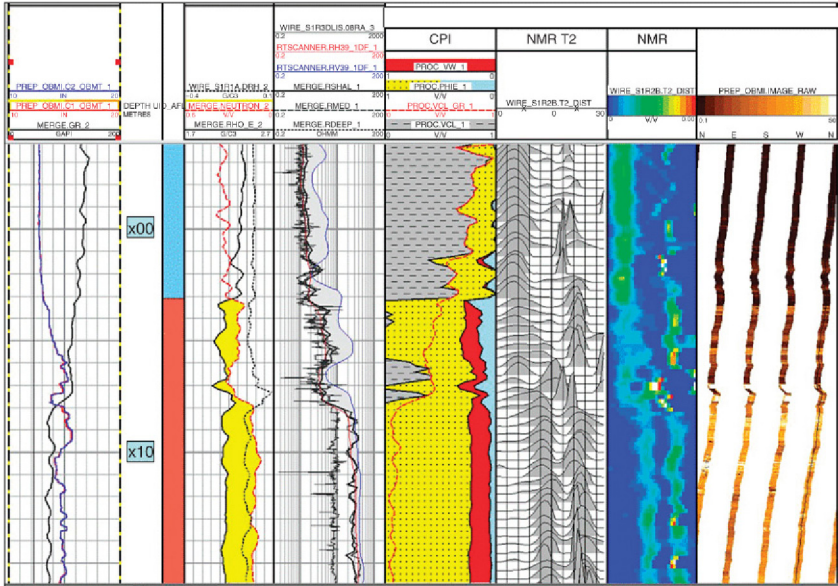
FIGURE 13.10   Example of logs in a thinly bedded interval (x00–x09 m). The horizontal induction measurement falls to a few Ωm in this interval but the vertical resistivity (dark grey curve [blue curve in the web version)]) exceeds 10 Ωm. NMR and the image log (right-hand tracks) reveal the presence of a thinly bedded component.

The two-way time depends on the speed of sound in the mud and the distance from the transducer to the borehole wall. If the speed of sound is known the two-way time can be converted to borehole radius and the tool has effectively become a very high-resolution borehole geometry tool. Most tools include a second, fixed transducer that fires pulses of sound at a plate a fixed distance away. This allows the speed of sound to be measured *in situ* and so the time of flight can be confidently converted to a radius measurement. With over 150 separate measurements of the borehole radius a very detailed picture of the shape of the hole can be built.

The amplitude depends on the contrast in acoustic impedance between formation and mud, borehole roughness and how attenuating the mud is. The latter should be more or less constant, at least over an interval of tens of metres, so the image is basically controlled by the acoustic impedance of the borehole wall and its roughness. A weak signal can be the result of a rough wall, which scatters a lot of the incoming energy and/or a low acoustic impedance contrast so that a lot of energy is transmitted into the formation. Conversely strong signals are favoured by smooth boreholes and hard formations giving a large contrast in acoustic impedance. For our purposes one or both these properties need to be related to 'geology' in order to produce a useful image. This is a major limitation of acoustic imagers. The tool can be
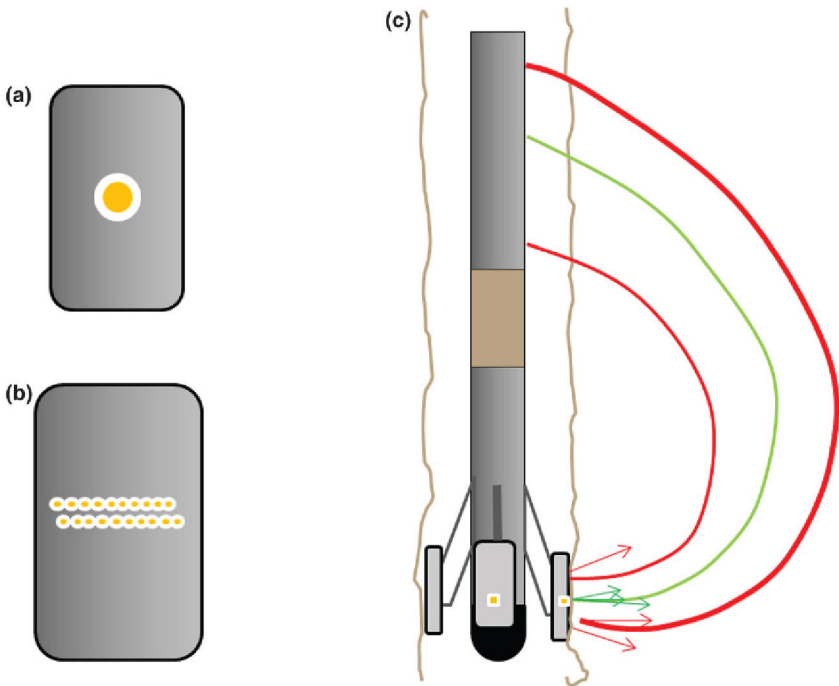
FIGURE 13.11 A typical electrical image log. (a) A pad from a dipmeter with a single electrode. (b) A pad from an image log with twenty small electrodes arranged in two rows. (c) A four-arm dipmeter/image log showing the current path from the pads to the tool body.

working perfectly well but if the borehole wall does not reflect geological features the image will not help.

13.5.3 The Inclinometer

All image logs – and dipmeters – include a device known as an inclinometer, which measures the way the tool is orientated in space. The original devices used a compass mounted in gimbals to determine the deviation of the tool relative to the vertical and the direction of one of the pads relative to North. In the 1980s these high-precision mechanical devices were replaced with a solid-state system consisting of three orthogonal accelerometers and three orthogonal magnetometers. The former determined the angle of the tool relative to the vertical and the latter its direction relative to North. This information is essential for correctly orientating bed boundaries, fault planes etc. with respect to the Earth. But even for our purposes, where we are really just interested in identifying and characterising, thin beds the images need to be orientated so that all features are given there correct vertical relationships. The role of the inclinometer can clearly be seen in Fig. 13.10 where the stripes corresponding to the four pads

can be seen to 'wander' towards the right as the depth decreases. This is a result of the tool rotating through 90° over the 20-m interval shown. So a pad that was facing north at the base of the log is facing East at the top.

13.5.4 Health Warning

An image log provides a lot of confidence that an apparently poor-quality shaly sand actually consists of thin sands and shales. Had an image log been available for the reservoir shown in Fig. 13.6 we would have known that a Thomas-Stieber model was not appropriate. However, it is possible for image logs to mislead. Many formations, particularly shales have a laminated fabric that produces a 'stripy' image log. The stripes are not alternating sands and shales, however, they are simply subtle changes in resistivity produced by the anisotropic fabric.

The other problem with image logs is that the apparent size of the features on the log is not necessarily a reflection of the true size. Conductive features tend to draw current towards themselves and therefore appear larger than they should and conversely current is drawn away from resistive features that appear smaller. For thin-bed systems this may result in the sand thickness being under- or over-estimated. Either way the thinner the sand, the larger the error. For vuggy carbonates it is tempting to the use image logs to quantify the vuggy porosity. There is a real danger of over-estimating this because the vugs exposed at the borehole wall are filled with mud and will be highly conductive.

13.6 NMR LOGS

Image logs reveal thin beds directly because they can resolve much smaller features than conventional logs. NMR logs are particularly good at identifying thin beds precisely because their vertical resolution is limited. Recall that the NMR log not only measures the amount of water in the formation but the T2 distribution gives information on it's microscopic environment (Section 5.5). In a thinly bedded formation the water bound to clays in the shale beds gives a different response to the water contained in the inter-granular pores of the sand beds. Furthermore, the relative amounts of water in the two environments can be quantified so that providing the total porosity of the sand and shale are known the sand–shale ratio can be found. Hydrocarbons and especially gas complicate the interpretation but they can be accounted for using more elaborate processing.

Figure 13.10 includes two representations of the T2 distributions from an NMR log (these are in the two tracks to the left of the image log). The 'waveform' display shows an individual T2 distribution approximately every 25 cm. Above x00 m the T2 distribution consists almost entirely of a single fast event with a peak at about 20 ms. This is caused by the clay bound water in the shale and therefore above x00 m the formation consists almost entirely of shale (this is consistent with all the other logs including the image log). Below x08 m there are two peaks with slower T2 times. The fastest of these is due to capillary

bound water adhering to the sand grains (in other words the irreducible water). The second, slower peak is most likely cause by filtrate that has displaced some of the gas. The gas itself would be expected to give a broad peak with a T2 in excess of 1000 ms, which is outside the range that has been plotted. The conventional logs are consistent with a clean, gas-bearing sand with the gas saturation reduced by invasion (although there is a clear gas effect on the density–neutron combination, the separation is not as great as expected for the low-calculated water saturation). The interesting part of the log occurs between these two depths where the conventional logs show a gradual deterioration in reservoir quality upwards (e.g. the gamma ray gradually increases). The NMR log shows that the peak due to the filtrate continues up to at least x03 m with little or no reduction in amplitude. In the faster T2 range the capillary water peak decreases in amplitude and a shorter T2 peak starts to appear as well and in fact these two peaks tend to coalesce into a single broad event. The interpretation of this is that the tool is responding to beds of sand and shale that are too thin to be resolved. The proportion of sand decreases markedly above x03 m and the sand appears to be of poorer quality because the T2 peak corresponding to filtrate shifts to slightly shorter T2 values.

A system of thinly bedded sands and shales is not the only explanation for the presence of two or three peaks in the T2 distribution. In fact multiple peaks only show that sand and shale are present as discrete units that are much smaller than the volume of investigation of the measurement. Shale clasts incorporated in a sand would give the same signature. In the case of the example discussed previously the NMR and image logs taken together provide strong evidence that there is a thinly bedded system in the upper 5–10 m of the reservoir.

Page left intentionally blank

Chapter 14

Geophysical Applications

Chapter Outline

14.1 INTRODUCTION

In Section 5.4 the sonic log was introduced and there it was stated that one of its most important functions is to provide the link between the borehole and surface seismic surveys. This is because the sonic log is a continuous, high-resolution record of compressional velocity along the well path. This information is sufficient to calculate the time for a sound wave to travel between any two points in the well. In particular, it should be possible to calculate the time for sound energy to travel from surface to any depth penetrated by the well. In other words the relationship between time and depth can be established. In practice there are a number of reasons why this may be more difficult than this simple recipe suggests. Nevertheless, it is worth persevering because the time–depth relationship allows logs to be converted from a depth-based graph to a time-based one, so that they can be superimposed on a seismic section. In this way prominent features on the logs can be related to reflectors seen on the seismic. Conversely, seismic sections can be converted to a depth basis.

There are other ways that the sonic – and other logs – can complement seismic surveys. Not only does the sonic give the compressional velocity at any depth it tells us the contrast that occurs at a bed boundary. This together with the density contrast determines the reflectivity of a bed boundary. By computing reflectivity as a function of depth along the well a synthetic seismic trace can be generated for comparison with an actual trace obtained at or near the well. Furthermore, by calculating the velocity and density that would be expected if a different fluid were present in a porous bed it is possible to predict the seismic

response for different scenarios. For example, if an exploration well actually intersected a water-bearing sand, the response of a gas-bearing sand could be predicted. This could then be used to high grade the un-drilled structures. Even more sophisticated modelling can be carried out if a shear slowness curve is available.

In the petroleum industry the term 'geophysics' is often synonymous with all things seismic but strictly speaking it covers far more than acoustic properties. The gravitational field, for example depends on the 3D distribution of density. In practice this offers an alternative to gamma-ray scattering to measure density along a borehole. Furthermore, a conventional density log can help refine a gravity survey. Similar comments apply to resistivity surveys and logs.

14.2   INTEGRATED TRANSIT TIME AND THE TIME–DEPTH CURVE

As noted in the preceding sections a basic sonic log shows the continuous variation of sonic slowness along the well path. The time – T – taken by sound to travel a short distance δh along the borehole is given by:

$$T = \delta h \times DT \tag{14.1}$$

where DT is the sonic slowness read by the tool. If the slowness is in microseconds per metre (μs/m) and the depth is measured in metres, T is given in microseconds (μs). This is one reason sonic logs record slowness rather than, the intuitively more familiar, velocity. By summing all the Ts between two depth points, the time for the sound wave to travel between them can be found. In symbols:

$$TT = \sum DT(h)\delta h \tag{14.2}$$

In the limit of a truly continuous sonic curve this becomes an integral, but as all modern digital logs are really a series of discrete measurements there does not seem much point going this far. Equation 14.2 applies to any two points in the well but in practice to get a time–depth curve one of those points is going to be some reference datum (such as mean sea level, ground level or a seismic reference datum). In this case TT is the time for sound to travel from, say, sea level to the depth of interest. This is a one-way time (OWT). A seismic section is conventionally given in two-way time (TWT), the time taken for sound to travel vertically down to the reflector and back. So in order to compare seismic and well data one needs to double the OWT.

The calculation of the transit time is normally done routinely whilst logging and the result is a curve named TT or possibly ITT (integrated transit time). Because the integration is automatic, the value of TT will increase in the direction of logging which for, wireline at least, is generally upwards. This is the opposite of the time–depth curve, but nevertheless the ITT still contains all the information
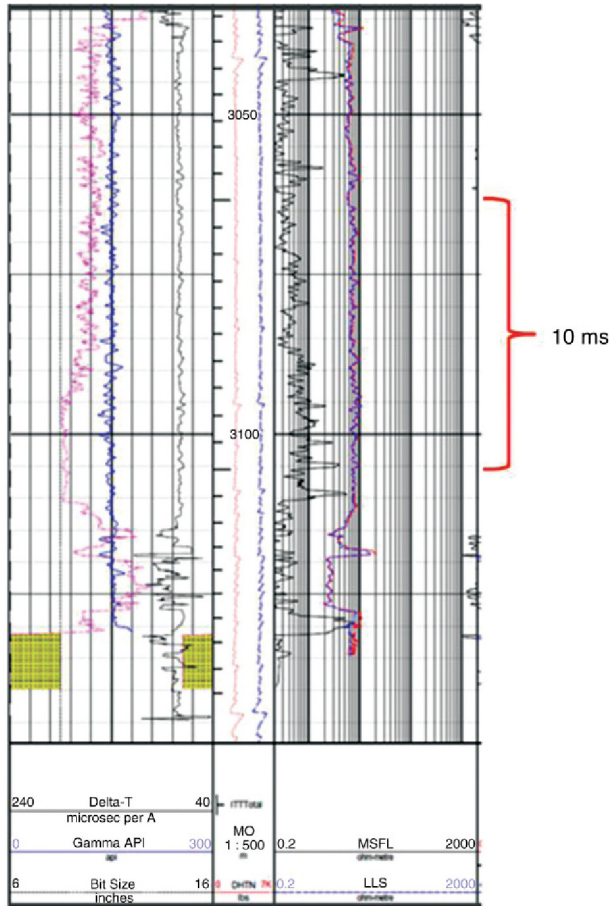
FIGURE 14.1 A short section of log showing the ITT as a series of ticks on the left-hand side of the depth track. Each small tick represents a transit time of 1 ms, large ticks are produced every 10 ms. The sonic log is the black curve in the left-hand track.

needed to generate the time–depth curve. The *ITT* is often presented as a series of ticks in the depth track: the time between adjacent ticks is 1 ms and by counting the number from the datum depth to each depth of interest, a time–depth curve can be constructed. An example of the integrated transit time on a log is shown in Fig. 14.1. Consider the interval 3050–3100 m. The total transit time between these depths is 12 ms and so the average slowness is given by:

$$DT = 12{,}000/50 \ \mu\text{s/m} \ (240 \ \mu\text{s/m or } 73.2 \ \mu\text{s/ft.})$$

The sonic log from which the *ITT* is calculated, is shown in track 1 on a scale of 240–40 μs/ft. It can be seen that the slowness between 3050 m and 3100 m is fairly constant with, as expected, a slowness of the order of 75 μs/ft.

Some sort of time–depth curve can be constructed from any sonic log but in practice there are a number of reasons why it may not be an accurate representation of the time–depth relation for the seismic section. These can broadly be classified as problems of:

1. Formation alteration by drilling.
2. Differences in frequency employed by seismic and sonic logging tools.
3. Problems caused by well and bed geometry.

A fourth possibility is that the sonic tool is simply not reading correctly. It is fair to say this is uncommon – at least for compressional slowness – but older tools could occasionally be affected by noise or weak signals. More often not these produced very localised spikes, which would have only a small effect on the integration. Long intervals of erroneous data could – and can – occur however and these will lead to large errors.

14.2.1   Formation Alteration

Sonic logging tools are shallow-reading devices and they will to some extent measure the slowness of formation that has been altered as a result of drilling. In permeable beds they will respond to formation that has been invaded by borehole fluids. For reasons that will be explained later, invasion rarely changes slowness significantly. The greatest problems tend to be associated with shales and other clay-rich rocks where there is a danger that they will react with the mud (at least if it is water based). In this case their velocities can be significantly reduced. The advent of LWD sonic tools has given the opportunity to compare sonic logs acquired a few hours after the formation was drilled, with a wireline log acquired several days later. A change of 10% is not unusual and changes of 20–25% have been observed over a period of 3–4 weeks. As a general rule the greatest changes occur in young un-compacted claystones (i.e. the ones that were slow even before alteration).

14.2.2   Dispersion

Even if the formation is unaltered there will be differences in velocity between sound waves generated by seismic and sonic tool sources. Sonic tools use sources that produce high frequencies (tens to thousands of Hertz, basically the same as human hearing) whereas seismic energy is generally much lower in frequency (a few to tens of Hertz). Actually, most seismic sources produce plenty of high frequencies but these are strongly attenuated in the earth so that at any depth of interest to the petroleum industry, the low frequencies dominate. This might not matter if all sound moved at the same speed, but in general the speed of sound depends on its frequency. This is known as 'dispersion' and is a general feature of any type of wave propagating in an attenuating medium (e.g. the electromagnetic waves exploited by LWD resistivity tools).

Dispersion is particularly marked for porous solids filled with naturally oc-curring fluids: precisely the type of materials we are interested in! As a general rule the speed of sound at high frequencies is higher than at low frequencies and for sedimentary rocks the most significant changes tend to occur in the frequen-cy range of a few Hertz to a few thousand Hertz (i.e. from seismic frequencies at the low end to sonic tool frequencies at the high end!). This change in velocity can amount to 25%. For non-porous rocks dispersion is almost insignificant.

Together, dispersion and formation alteration mean that the sonic tool is re-sponding to the velocity of high-frequency sound in altered formation whereas the seismic experiment responds to low-frequency sound in unaltered forma-tion. To an extent they compensate each other. Velocities tend to slow when shales are altered and so move closer to the low-frequency velocity of the un-altered formation. But the two components are unlikely to perfectly compensate each other.

14.2.3   Hole Geometry

A seismic section is normally processed so that the reflectors are located cor-rectly in space. This means the section is based on the assumption that every-where sound has travelled vertically down to a reflector and is reflected back along the same path. In reality relatively little if any of the seismic energy does this but the processing is designed to 'correct' for this. The raw sonic log can be thought of as being generated by sound that follows the path of the well. Provid-ing the well is vertical this path is parallel to the processed seismic signals but if the well is deviated this is no longer the case. The transit time between any two points will increase but providing the deviation is not too great this can be accounted for by a simple TVD correction. Either the original sonic log can be converted to TVD and the transit time re-calculated or the integrated transit time can be converted to TVD.

Problems start to occur at high deviations when the well covers significant horizontal distances and each reflector is intersected at a different position on the horizontal plane. More fundamental problems are caused by anisotropy where the horizontal velocity is significantly different to the vertical. The seis-mic response is dominated by the latter whereas the sonic log is being affected by both components.

14.3   SONIC CALIBRATION

The sonic tool is one of the few logging measurements that is not calibrated (it is implicitly assumed that the tool's timing chips are always accurate). Neverthe-less it is possible to 'calibrate' a sonic log. In this case however the calibration involves adjusting the sonic log so that it agrees with a seismic experiment. In this way the effects of formation alteration and dispersion are removed. The procedure relies on making a direct measurement of the time taken for sound to

travel from a seismic source at surface to a particular point in the well. This is known as a 'check-shot'. The sound will have travelled almost entirely in unaltered formation and furthermore because it was generated by a seismic source and has had to pass through a significant thickness of rock, its frequency content will be close to that of the surface seismic.

As an example of how the calibration is applied consider the example shown in Fig. 14.1. Assume checkshots were made at 3050 and 3100 m and the times taken for the sound to travel from an air gun at surface were:

Depth	TT (ms)
3050	1006.5
3100	1018.0

The time taken for the seismic signal to travel from 3050 m to 3100 m is therefore 11.5 ms. This is slightly faster than the sonic log, which gave an ITT of 12.0 ms for the same interval. This suggests the sonic signal is travelling slightly slower than the seismic velocity and in order to agree, the sonic velocity needs to be increased by the factor 12/11.5 = 1.043. Since we are adjusting a sonic log the slowness needs to be reduced and it in fact the raw sonic readings would be divided by 1.043 over the whole interval 3050–3100 m.

In practice, additional corrections are applied to account for geometrical factors such as the horizontal offset of the seismic source from the wellhead. Furthermore the actual block shifts to the sonic log would be made at formation boundaries to avoid introducing artificial reflectors but in essence the above describes how a sonic log is calibrated.

14.4 FLUID SUBSTITUTION

Broadly speaking logs provide two links to seismic:

1. The time-to-depth transform that puts a seismic section on a depth basis or conversely puts logs onto a time basis (discussed earlier).
2. Modelling the amplitude and phase of a reflector.

In this section, we are interested in the latter and in particular we wish to calculate a synthetic seismic trace at the well to compare with the surface seismic.

In principle this is a simple process providing sonic and density curves are available. If this is the case a continuous curve of acoustic impedance (Z) can be calculated, this is given by

$$Z = \rho v = \rho / \Delta t \tag{14.3}$$

Because logs use a variety of units, which are often inconsistent, some units conversion is normally desirable. But what we are really interested in is the reflection coefficient given by:

$$R = \frac{Z1 - Z2}{Z1 + Z2} \tag{14.4}$$

So the units cancel out (in other words it does not really matter if the acoustic impedance is given in g.ft./cm^3.μs rather than the SI unit kg/m^2.s). Of more concern here is that both sonic and density are shallow-reading tools that may actually measure properties of the altered formation. Of particular concern are the effects of invasion in permeable beds.

It is worth investigating whether this concern is justified. The density curve can easily be corrected for invasion by applying the general density equation. In a clean sand:

$$\rho_b = \rho_{ma}(1-\varnothing) + \rho_w \times \varnothing \times S_w + \rho_{hc} \times \varnothing \times (1-S_w) \qquad (14.5)$$

The log actually measures the density found by replacing S_w with S_{xo}. The largest discrepancies will occur in high porosity gas sands with deep invasion. For water-based mud with a filtrate density close to that of the formation water the change in density is given by:

$$\Delta \rho_b = \varnothing(S_{xo} - S_w)(\rho_w - \rho_{hc}) \qquad (14.6)$$

For a worst-case scenario of a 30% gas sand with S_w of 20% and S_{xo} of 80%, the change is 0.14 g/cm^3 to a true formation density of about 2 g/cm^3. This is a worst case and it amounts to a correction of less than 10%, in an oil reservoir and/or less invasion and/or lower porosity the correction will be much smaller. Furthermore, deep invasion is actually quite uncommon in high porosity gas sands and so even in the worst-case scenario density changes are likely to be small.

Changes to the sonic caused by invasion are more difficult to calculate, but a theoretical model is available, this is the Gassman equation. Before discussing this further, it is worth pointing out that compressional velocity or slowness can be very sensitive to the presence of gas but tends to vary little with changes in the gas saturation. In other words the sonic can be a very strong indicator of the presence of gas but is poor at quantifying gas saturation. One of the consequences of this is that providing some residual gas remains in the invaded zone the measured slowness is likely to be very close to the slowness in the unaltered formation. In other words, invasion is not expected to have a large effect on reflection coefficients and using the raw logs is likely to produce a satisfactory synthetic seismic trace. This is frequently observed in practice and a bigger issue is likely to be whether a continuous density curve is even available (either because it was never recorded or because parts of it have been rendered useless by poor hole conditions).

The calculation of synthetic seismic traces can be extended to modelling situations, which may never have been encountered in the sub-surface. For example, using the properties measured in a water-bearing sand to model what it would look like if it had been filled with gas, or vice versa. This can then be used to predict the fluid fill from reflection amplitude in an undrilled structure in the same play. In this case the changes in compressional velocity can be large.

14.4.1   The Gassman Equation

For any material, slowness is related to density by

$$\Delta tc = \sqrt{(\rho / (K + 4/3\mu))} \qquad (14.7)$$

assuming a self-consistent set of units. For typical log units of slowness in microseconds per metre, density in gram per cubic centimetres and K and μ in GigaPascal, this becomes:

$$\Delta tc = 1000\sqrt{(\rho / (K + 4/3\mu))} \qquad (14.7a)$$

(A rare example of an equation in petrophysics that is rigorously true.) K is the bulk modulus, which describes how easy it is to compress the material. It is defined as the ratio of pressure to volume strain (the fractional change in volume). A high bulk modulus implies the material is relatively incompressible. A slow formation – high slowness – is thus favoured by a high density or low bulk modulus. The shear modulus μ describes how easily the rock deforms when subjected to a shear stress. The higher the value the more it resists deformation (for fluids it is zero).

The Gassman equation gives K as a function of porosity and the bulk moduli of the matrix, fluid mixture and something referred to as the 'dry-frame matrix'. It is quite a complicated equation but the real challenge is to be able to apply it.

$$K = K_{df} + \frac{(1 - K_{df}/K_{ma})^2}{\left(\dfrac{\O}{K_f} + \dfrac{1-\O}{K_{ma}} - \dfrac{K_{df}}{K_{ma}^2}\right)} \qquad (14.8)$$

where subscripts ma, f and df refer to the matrix, fluid and dry frame, respectively. Arguably the main problem with using Eq. 14.8 is defining these parameters. The fluid is normally a mixture of water and hydrocarbon and conventionally K_f is assumed to be given by the following mixing law:

$$\frac{1}{K_f} = \frac{S_w}{K_w} + \frac{(1 - S_w)}{K_{hc}} \qquad (14.9)$$

The bulk modulus of water is about 2.2 GPa and gas is typically less than 0.05 GPa (oil depends on the GOR and gravity). Because the bulk modulus of gas is so low compared to that of water, K_f will be small even at high water saturations. This is the reason why sonic slowness is such a good indicator of the presence of gas.

K_{ma} is more difficult to define but the values for most of the reservoir-forming minerals are known. Some values are given in Table 14.1.

More complicated mineralogies are taken care of by various mixing laws, often a simple weighted average is sufficient given how little is known. The

TABLE 14.1 Mechanical Properties for Some Commonly Occurring Minerals

	K (GPa)	ρ (g/cm^3)
Quartz	38	2.65
Calcite	77	2.71
Dolomite	95	2.87
Siderite	124	3.96
Sandstone	37.5	
Limestone	70	
Dolomite	83	

biggest problem is the dry-frame bulk modulus as this refers to something which does not exist (in this respect it is not unlike finding the shale porosity in order to convert between total and effective porosity). As the name suggests it is the bulk modulus of the porous solid minus any fluid. Unfortunately, it is not the same as the bulk modulus of a dried core plug. The action of drying the plug produces something that is acoustically different to the original rock.

K_{df} is normally found from the density and compressional slowness logs assuming that the porosity and fluid mixture within the depth of investigation of the tools is known. In other words the Gassman equation is re-arranged to give K_{df} in terms of K (which can be found from Eq. 14.7. Having found K_{df} the Gassman equation can be used to find K and hence slowness for any mixture of fluids in the pore space. The conventional workflow is as follows.

1. Compute S_{xo} and hence K_f in the invaded zone.
2. Use the log readings to find K_{df}.
3. Use the new value of K_{df} to calculate K and hence compressional slowness for water saturated formation.
4. Calculate K_f and hence K for the desired mixture of fluids one wishes to model.

Formation density would be modelled at the same time using Eq. 14.5. The shear slowness is also normally modelled although the changes between water- and gas-bearing sand are always small for reasons that are explained later.

An example of a fluid substitution is shown in Fig. 14.2. In this case it is being used as a 'post mortem' on an exploration well that was drilled to investigate a seismic bright spot. The sand was essentially water bearing but there were good gas shows as it was drilled. The fluid substitution suggests the sand contains gas at a low saturation (ca. 10%). If the gas–water mixture is replaced by pure water the slowness falls by about 10%. This is sufficient for the 'bright spot' to dim considerably.
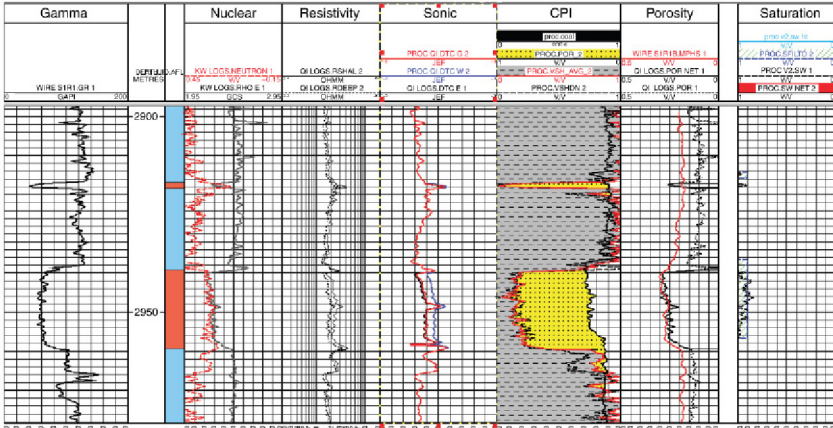
FIGURE 14.2 An example of fluid substitution into a sand suspected of being filled with residual gas. The compressional slowness curves are shown in track 6 ('sonic'). The black curve is the recorded sonic log and the grey curve (1 division to the right in the sand) curve is the modelled sonic assuming 100% water in the pore space. Modelling the sonic response to the low gas saturation shown in the RH track almost perfectly reproduces the measured curve.

14.4.2 Shear Slowness

Acoustic energy is transmitted by two modes in solids.

1. As a compressional wave (discussed earlier).
2. Shear wave.

These are distinguished by the direction of the displacement with respect to the direction of the wave. Up until now we have mainly been interested in compressional waves but for completeness we should also discuss shear waves. Compressional waves consist of alternating compressions and rarefactions and so the individual particle displacements occur in the direction of the wave. With shear waves the displacement is transverse to the direction of the wave. One consequence of this is that shear waves cannot propagate in fluids including the mud (in fluids there are no restoring forces to a shear displacement). So the primary disturbance created by the tool is a compressional wave in the mud. Any disturbance at a fluid–solid interface will however set up a compressional and a shear wavefront in the solid. The shear wavefront is always slower than the compressional. Array sonic tools were introduced to measure formation compressional and shear slowness. This comes at the cost of considerably more complicated hardware and a huge increase in computation than used in the earlier, 'first arrival' tools. It is worth asking why there is a demand for the shear slowness.

The two principle applications of shear data are for rock mechanics and for advanced geophysical applications. We have already seen how compressional

slowness is used to model seismic traces for well ties and to predict how reflectivity changes with fluid type. It turns out that reflectivity changes with the angle of incidence at the reflector. This is the basis of amplitude variation with offset (AVO). The way that reflectivity changes with angle of incidence can be predicted but only if the shear velocity above and below the reflector is known. If compressional and shear velocity and density are known, reflectivity can be calculated as a function of the angle of incidence by solving a set of non-linear simultaneous equations known as the Zoepritz equations. This is well beyond the scope of this book.

It is worth noting here that the shear velocity hardly changes with fluid type in the pore space. This is because shear waves are not supported by fluids and so it is basically the grains that allow the shear wave to pass through a porous solid. The shear velocity – or in this case slowness – is related to density by an equation that is similar to Eq. 14.7.

$$\Delta ts = \sqrt{(\rho / \mu)} \tag{14.10}$$

where μ is the shear modulus. We know that density is unlikely to change by more than 10% when a sandstone originally saturated with water is filled with gas and so DTS is not going to change by more than 5%. Shear slowness increases with increasing density and so all other things being equal will be slightly higher for a water-bearing than a gas-bearing sand (i.e. it is slower when water bearing). This is the opposite to the trend expected of compressional slowness but that is mainly caused by the dramatic reduction in K when gas is introduced.

An example of the contrasting effects of gas on compressional and shear slowness is shown in Fig. 14.3. The target of the well were the high porosity unconsolidated sands shown in the short section of log. Although gas was seen when drilling through the sands, log analysis showed at best residual gas saturations (note the low resistivities). A Gassman fluid substitution was performed to model the effect of replacing a gas saturation of 10% with pure water. The result is a dramatic reduction in slowness from measured values of 670 μs/m to a value of 560 μs/m for the water-saturated sand. On the other hand density hardly changes at all because the fluid has simply changed from a mixture of 90% water and 10% gas with a density of 0.92 g/cm^3 to pure water with a density of 1.00 g/cm^3. This causes a density change to the sand of about 1.5% and the percentage change in shear slowness will be even less.

The almost immeasurable change in DTS and the dramatic increase in DTC when gas is introduced to the pore space, forms the basis of several qualitative gas indicators. For example, producing a V_p/V_s (DTS/DTC) curve shows the presence of gas by a significant fall compared to its value in a water-saturated sand. Another technique involves selecting scales for DTS and DTC so that they overlay in water-bearing sands and diverge in the presence of gas.
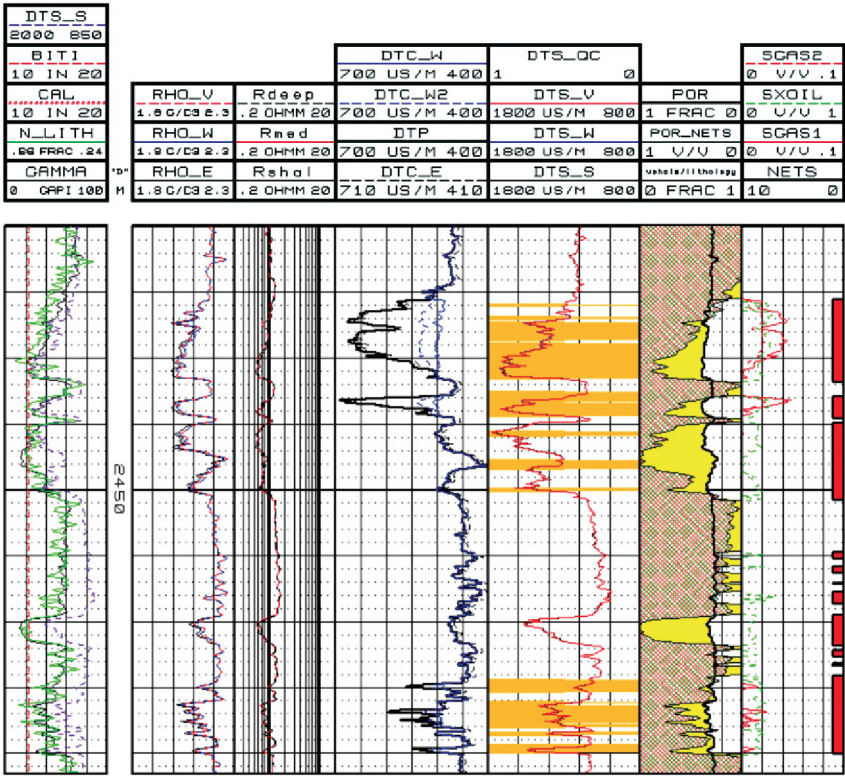
FIGURE 14.3 Another example of a large gas effect on the compressional sonic curve caused by a low gas saturation. The compressional slowness' is shown in track 4 and the shear slowness' in track 5. The measured compressional slowness (black) shows a very slow formation. The modelled response (grey) is significantly faster.

14.4.3 V_p/V_s Ratios: Empirical Relationships and Fundamental Limits

Direct measurements of shear slowness only became available when array sonic tools with full wavetrain recording appeared in the mid-1980s. Even then shear slowness could not, at first, be measured in 'slow formations'. That is rocks in which the shear velocity is slower than the mud velocity (which typically means shear slowness greater than $600\,\mu s/m$). The advent of dipole sources in the 1990s removed that limitation but the fact remains there are many wells drilled before this which could help understand seismic responses if only they had shear data. Fortunately, shear and compressional velocity – or slowness – are often closely correlated and various empirical relations have been published to allow shear velocity to be estimated from a compressional log. Probably the best known are those of Castagna and Han which were published in the 1980s
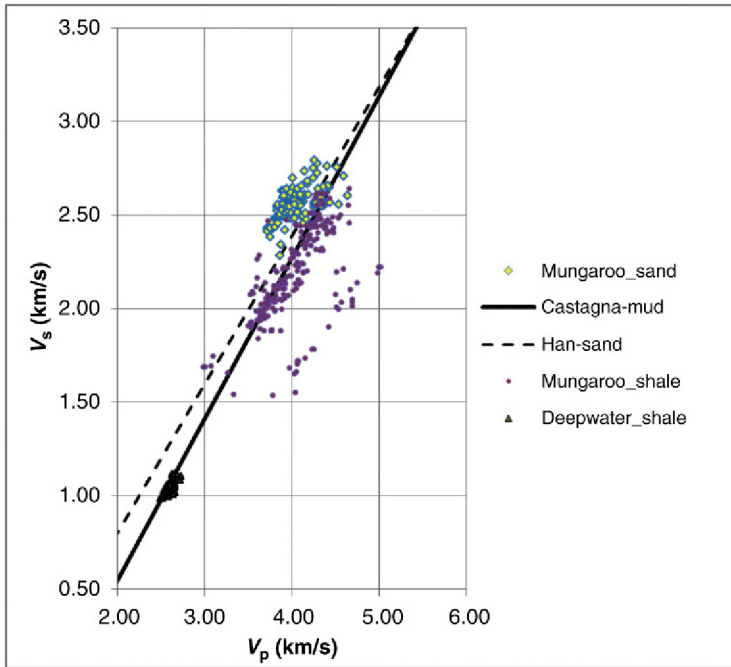
FIGURE 14.4 Log data for some Triassic sands and shales of the Mungaroo formation sediments from Western Australia and some deepwater Tertiary sediments from East Africa. Also shown are two standard V_s (V_p) relationships that are used to predict V_s in sands (dashed) and mudstones (solid).

but these type of relationships have a long history and some date back to the 1960s. Typically they are linear and specific to a particular lithology.

The Castagna 'mud-line', for example has the form:

$$V_s = 0.8621V_p - 1.1724 \qquad (14.11)$$

where V is the velocity in kilometres per second. Similarly Han's relationship for sandstones is:

$$V_s = 0.7936V_p - 0.7868 \qquad (14.12)$$

These particular relationships have been plotted together with some real log data for water sands and shales in Fig. 14.4. In this particular case the Castagna mudstone relationship is doing a reasonable job for both the old and young shales although they have very different properties and histories. The Han relationship seems to under-estimate the true shear velocity in the Triassic sands however. This behaviour is by no means typical and as a general rule these type of relationships are more reliable in sands than shales.

Many log analysis programmes have a module, which includes these equations. Apart from being essential for estimating shear velocity in wells where

only compressional data is available, they can also be used as QC checks on shear logs. These equations predict that the ratio V_p/V_s increases as compressional velocity slows.

The ratio V_p/V_s is related to Poisson's ratio, which basically describes how a stressed material responds in a direction normal to the one in which it is being stretched or squeezed. The relationship between V_p, V_s and Poisson's ratio is:

$$P.R. = \frac{V_p^2 - 2V_s^2}{2\left(V_p^2 - V_s^2\right)} \qquad (14.13)$$

Poisson's ratio is constrained to lie in the range 0–0.5 at least for rocks (strictly speaking P.R. can be negative but this only occurs with special materials that have been engineered at the molecular level to expand when they are stretched). A value of zero applies to a material that does not change in thickness when it is stretched or squeezed (cork is the most familiar example, this gives it the useful property that it does not expand when it is pushed into the neck of a wine bottle). It can be shown that a material with a Poisson's ratio of zero has V_p/V_s equal to the square root of two and so this represents a fundamental lower limit (i.e. V_p/V_s should never be less than 1.414). So this represents a simple QC check on shear velocity. At the high end, a Poisson's ratio of 0.5 is equivalent to compressional velocity being much greater than shear velocity (rubber is the example that is most often quoted). Most rocks have Poisson's ratios between 0.1 and 0.4.

14.5   BOREHOLE GRAVITY SURVEYS

Borehole gravity surveys provide a high-resolution measurement of 'g' – the acceleration due to gravity – along the well path. They are to gravity surveys, what sonic logs are to seismic. In practice, however the measurements are made at a series of stations that are typically between 2.5 m and 10 m apart. So the depth increment is considerably larger than a sonic log (which is for all intents and purposes a continuous measurement of velocity). Borehole gravity surveys have two main applications:

1. To relate the surface gravity survey to the vertical variations in g.
2. To provide a deep-reading density log.

From a petrophysical point of view the latter is the most important application but before looking at that application we will briefly look at gravity surveys in general.

A gravity survey, whether conducted at the surface or along a borehole, involves measuring acceleration due to gravity at different points (this is invariably given the symbol g). It is a vector quantity and the latest surface surveys

measure its magnitude and direction but here we will just consider the former. At the Earth's surface it has a value of $9.81\,\mathrm{ms}^{-2}$ to two decimal places. But if more precise measurements are made it is found to vary slightly. There are several reasons for these small variations including:

1. Large scale changes in the radius of the Earth.
2. Latitude.
3. Elevation.
4. Local topography.

These are all predictable changes and in a gravity survey they are calculated and removed from the measured value of g. Whatever variation remains is basically due to variations in the composition of the rocks making up the crust in the vicinity of the survey. For example, a basalt intrusion into a sedimentary basin represents a local increase in density which will give in an increase in g. Surface surveys of g allow these features to be mapped out. Unfortunately, there is no unique explanation for a particular anomaly so that, for example a local increase in gravity might be due to a small, very dense intrusion near the surface or it may be due to some larger, deeper feature. One way to distinguish the two possibilities is to drill a well.

The typical variations in a gravity survey amount to $0.00001\,\mathrm{ms}^{-2}$, that is to say one part per million of the value of g at the Earth's surface. Most instruments are designed to detect changes of 1 part in 10^8 or 10^9. Gravity surveys normally use a unit known as the Galileo or -'Gal' which is equal to $0.01\,\mathrm{ms}^{-2}$. Because this is still much larger than a typical anomaly in a survey the milliGal (mGal) is actually the usual unit used.

Gravity is measured in surface and downhole instruments using a very accurate and precise balance. A schematic of a typical downhole instrument is shown in Fig. 14.5. The tension in the spring (T) is adjusted so that the hinged bar that connects the mass m to the instrument chassis is horizontal. In a downhole tool the adjustment to the tension is made using some electro-mechanical device but with surface tools the adjustment can be made manually. At this point:

$$Tc = md \times g \tag{14.14}$$

Since the mass m and the lengths d and c are known g can be calculated. Because of the high precision required the instrument has to be isolated from any mechanical disturbance and needs to be kept at a fixed temperature, for this reason each measurement takes several minutes and the tool has to be stationary. This explains why the log is recorded in stations.

The geometry of the measuring instrument and the requirement for high precision imposes a limit on how compact the tool can be. Downhole tools in particular have relatively large outside diameters of 4 or 5 in. The balance can be mounted on gimbals to allow it to remain vertical in deviated wells but size constraints limit the deviations to low values ($15°$ say).
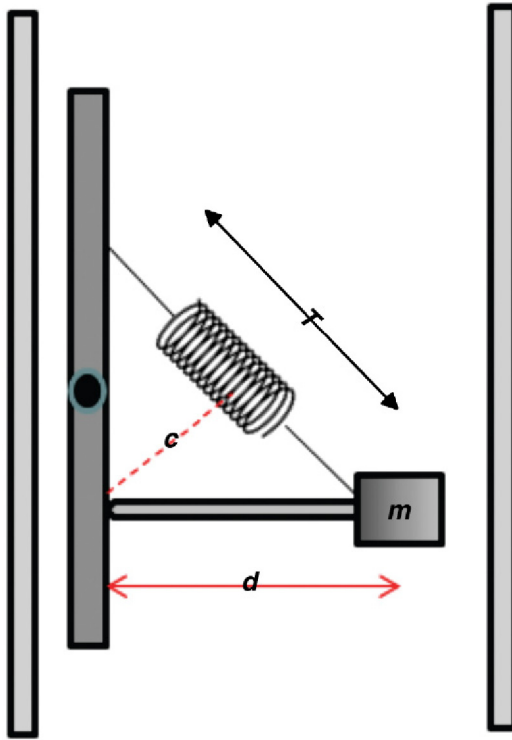
FIGURE 14.5   Schematic of a borehole gravity instrument.

The gravity log will consist of a series of measurements of g against depth. Acceleration due to gravity – g – falls in a predictable way with height above the Earth's surface. The relationship is:

$$g = \frac{G.M}{R^2} \tag{14.15}$$

$$\frac{dg}{dR} = \frac{-G.M}{R} = F \tag{14.15a}$$

where R is the distance to the centre of the Earth, M is the Earth's mass and G is the universal gravitation constant (one of the fundamental constants of nature). 'F' is known as the 'free air gradient of gravity'. Substituting for G and M (7.10^{24} kg) gives, at the Earth's surface ($R = 6880$ km) a value of approximately.

$$F = 0.308 \text{ mGal/m}$$

The precise value depends on the exact location on the Earth's surface that the survey was made. The situation is more complicated below the Earth's

surface because the mass (M) contributing to g falls as one goes deeper (i.e. R decreases). So Eq. 14.15a becomes:

$$\frac{\mathrm{dg}}{\mathrm{d}R} = \frac{-G.M(R)}{R} \qquad (14.15\mathrm{b})$$

Gravity data is typically modelled by treating the Earth as a series of concentric shells, each with a constant density ρ. Then the rate of change of g with depth depends on the density (ρ):

$$\frac{\mathrm{dg}}{\mathrm{d}R} = F - 400,000\pi G\rho \qquad (14.16)$$

where the gradient is again given in mGal/m providing G and density are given in SI units (i.e. density in kilograms per cubic metres). In practice because g is measured at stations, the formula is actually given in terms of the change in g divided by the depth difference. The final density log will thus consist of a series of straight lines with step changes at the station depths.

In practice, like any logging tool, the borehole gravity meter (BHGM) has a limited depth of investigation. As a rule of thumb 90% of the signal originates from within a radius of approximately five times the station interval. For a typical log this means somewhere between 10 m and 50 m. That means the density measurement is almost unaffected by invasion, badhole and can be made inside casing. Although the default processing assumes horizontal slabs of uniform density and infinite extent (or at least world-wide extent) it is possible to forward model more complicated situations such as a volcanic intrusion or salt diaper that is close to the well. The model can then be compared with the actual survey and if necessary modified to produce a better match.

14.6 DEEP READING RESISTIVITY SURVEYS

Resistivity logs are by far the deepest reading conventional logs and most of the time their depths of investigation are far greater than the altered zone. Nevertheless in the late 1960s resistivity tools were developed that could read a kilometre or more away from the borehole. The purpose of these tools was not to avoid the effects of borehole alteration but rather to detect and characterise large resistive bodies in the vicinity of the borehole. More often than not these were salt diapirs but the technique would be equally good at detecting volcanics or carbonate reefs. The tools were actually very simple devices and consisted of three electrodes located on a long length of insulated cable known as a bridle. Spacings were typically several hundred metres but tools with electrode spacings of 1200 m have been built. The usual rule that the larger the spacing the deeper the tool reads applies (Fig. 14.6).

Like the borehole gravity survey the log is actually measured at stations with each measurement taking about 1 min to make. Current is passed between

FIGURE 14.6 Schematic of a deep-reading resistivity tool. Current is passed from electrode A to a fourth electrode at or near the surface. The resulting voltage drop is measured between electrodes M and N. The ratio AM/AN varies from 0.1 to 0.5.

the electrode at the base of the tool (A) to an electrode at or near the surface (conventionally referred to as the B electrode). This typically produces a voltage drop of the order of a micro-volt between the M and N electrodes. The stations are located between 10 m and several 10s of metre apart. As a general rule the longer the electrode spacing the larger the distance between stations.

The raw data from the tool is interpreted by comparing it to the calculated response of an 'Earth model' consisting of laterally extensive, parallel sided layers of uniform resistivity. The model is generally built from a conventional resistivity log. If the actual resistivity measured by the tool is higher than the calculated response, then a resistive body is near-by. The larger the difference the closer the body.

14.7 CONCLUSIONS

This chapter moves beyond the traditional realms of petrophysics which, although literally the 'physics of rocks', tends in practice to concern itself with volumes of 'boulder' size or less. In Chapter 1 it was pointed out that the largest source of uncertainty in resource estimates is more often than not the result of not knowing what happens beyond the volume of investigation of logging tools. The techniques described earlier, either allow the depth of investigation of borehole measurements to be extended by orders of magnitude or allow surface measurements such as seismic surveys to be calibrated to physical and ultimately petrophysical properties. Many of these techniques have been around for decades but were largely the preserve of specialist companies. It is hard to imagine that these and related deep-reading techniques will not become more common as the oil and gas developments become less able to afford wells that miss their intended targets. These deep-reading techniques offer the best chance of avoiding under achieving development wells.

Chapter 15

Epilogue: High-Angle Wells

Chapter Outline

15.1 INTRODUCTION

This is the end of the book, but rather than just summarise the last 14 chapters I have chosen to finish by describing how the tools, techniques and overall philosophy can be adapted and applied to the particular case of high-angle wells. This is why I have given the chapter the title 'Epilogue' when purists will rightly argue that the epilogue of a book comes after all the chapters. Anyway, as far as this book is concerned the important feature of high-angle wells is that bed boundaries are more or less parallel to the well path. More often than not the bed boundaries are close to horizontal, so the wells of interest here are typically deviated at – say – 90°(±10) but the arguments apply equally well to vertical wells penetrating steeply dipping strata. For sure, truly horizontal wells introduce all sorts of practical problems for coring and logging but those are not the focus of this chapter or, for that matter, this book, rather it is what happens to logs when the bed boundaries are parallel to the well that is of interest.

Horizontal wells date back a surprisingly long way. In fact the first horizontal well was drilled in the late 1920s (at about the same time as the first log was run). But it was really only in the 1980s that high-angle wells became 'mainstream'. This was a result of necessity and technological advances. It is fair to say that they are essential for getting adequate production rates in plays as diverse as small–medium sized offshore oil fields, tight gas and unconventional oil and gas. In short these days it is difficult for a practicing petrophysicist to avoid interpreting data that has come from a high-angle well.

In Section 1.5 it was asserted that the best way to make progress is to combine the general results of research with specific measurements made on the rocks of interest. This approach has been implicitly used throughout the book. In this context 'research' does not necessarily mean laboratory work in a dedicated facility, rather it refers to any study of an ideal(ised) system, by anyone, in order to make sense of measurements on real rocks in real boreholes. It may be a long-term study, involving several full-time researchers using the latest laboratory

tools but it could equally involve half an hour and a few sketches on the back of an envelope. The latter is the approach we will adopt here to explain the responses of conventional logs at a bed boundary that cuts a well at a low angle. Suffice it to say the major service companies have spent a lot of time and money on detailed mathematical modelling and to truly understand the log responses there is no alternative. Certainly the fruits of that work are now used on a daily basis.

15.2 LOGGING HIGH-ANGLE WELLS

As more and more horizontal wells were drilled it became apparent that the logs recorded in them were visibly different to those from nearby vertical wells. Various explanations were proposed most of which were quite valid but rarely provided the whole explanation. These included the following:

1. Invasion by filtrate is deeper on the low side of the hole.
2. Differences between wireline tools that were predominantly used in vertical wells and LWD that are preferred for the horizontal wells.
3. Formation anisotropy.
4. The presence of two different lithologies within the volume of investigation of the tool.

The importance of each of these can be appreciated by simple sketches showing the relative locations of well and bed boundary. Figure 15.1 shows a
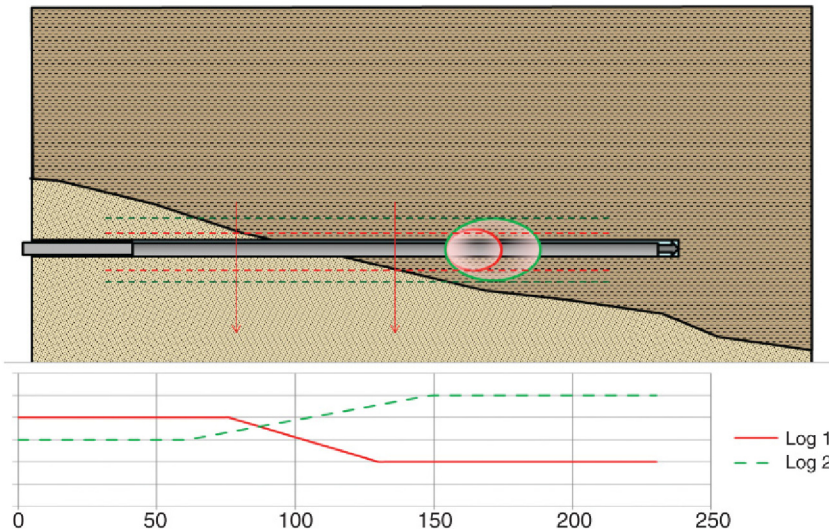
FIGURE 15.1 Cartoon illustrating how log responses in a high-angle well are influenced by the depth of investigation. The depths of investigation are shown by the dotted lines running parallel to the well. The idealised logs that would result are shown immediately below the sketch.

sketch of a horizontal well cutting through a shallow dipping bed boundary that separates a sand (below) from a shale. This is a typical situation in a development well where the intention is to place the well just below the reservoir seal. Subtle changes in the well path or the bed boundary have in this case caused the well to stray into the shale. The diagram shows two measurements with different depths of investigations. Both have 360° volumes of investigation. The key point to note is the distance along the hole that the measurements are influenced by both sand and shale. In contrast to a vertical well, there will be a gradual change from the sand response to the shale response over this interval and the deeper the depth of investigation the longer the transition continues. The raw measurements in a vertical well may show some sort of a transition at a bed boundary but this is a result of the finite vertical resolution rather than the depth of investigation (and at the scales logs are typically plotted at it may not be apparent at all).

An important issue is whether a transitional log response is an artefact of the finite volume of investigation or a reflection of a gradual change from one lithology to another. In the case of the horizontal well a systematic increase in the length of the transition with depth of investigation is a clear sign that the boundary is actually quite sharp. But in the real world such trends may not be as clear-cut as this simple model suggests. Complicating factors include:

1. Depth of investigation is not a hard boundary and in any case it generally varies with the physical properties of the formation.
2. The natural variability of the formations on either side of the boundary is superimposed on the ramp predicted by the simple model.
3. Horizontal wells may be drilled with the tools 'sliding' so that tools that only look in a limited range of directions will respond differently to '360°' measurements.

15.2.1   Some Real Data

Some real data is shown in Fig. 15.2. This shows a short section of 'triple-combo' memory data acquired in a horizontal production well drilled in an oil field. The bit-size is 8½ in. or 0.22 m. The interval shown includes a shale at the 'top' and 'base' of the section and a very good quality oil-bearing sand in between. The logging tools were a gamma ray with two detectors that 'look-at' opposite sides of the hole, a modern resistivity tool with multiple depths of investigation and density and neutron tools. Also shown, in track 3, are density readings calculated from the individual near and far detectors. The gamma ray shows a continuous change over an interval of approximately 6 m from the shale reading of 130 api to the sand reading of just over 25 api. By contrast the density shows a more complicated and generally more rapid change between the shale value of 2.5 g/cm^3 and the sand value of $2.15–2.2 \text{ g/cm}^3$. The neutron also shows
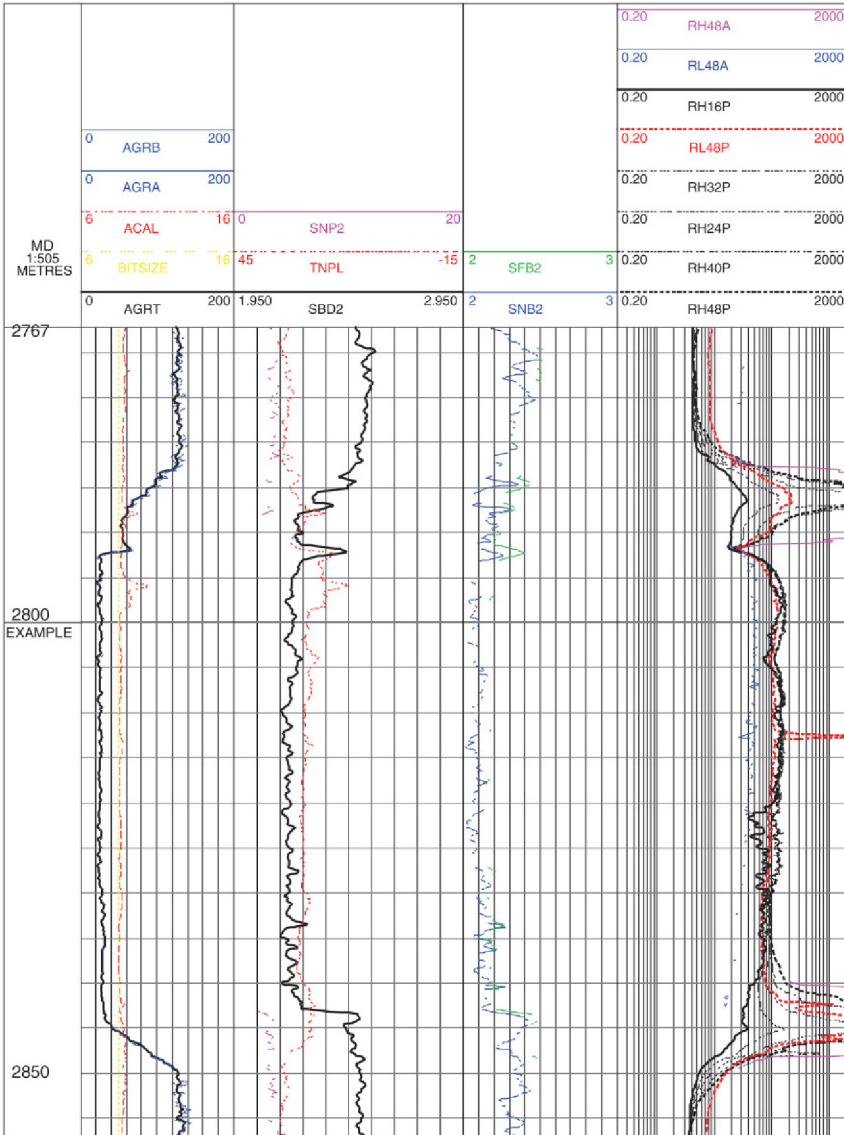
FIGURE 15.2   Memory LWD data from a 100-m interval of a horizontal development well showing a transition from shale- to oil-bearing sand and then back into shale. Depth lines are a spacing of 5 m.

a more complicated response but like the gamma ray shows a gradual change from the shale reading of 33 pu to the sand reading of 27 pu.

The gamma-ray response is very reminiscent of the idealised model of logs in a horizontal well shown in Fig. 15.1. The depth of investigation of a typical gamma-ray tool in formations with densities of the magnitude seen here is about

0.25 m so the distance equivalent to the separation between the two dashed lines in Fig. 15.1 is:

$$O = 0.25 + 0.2 + 0.25 = 0.7 \text{ m}$$

This interval represents the distance across the well bore where the tool is influenced by both sand and shale. The condition occurs over a distance of 6 m and so a bit of elementary trigonometry allows us to estimate – θ – the dip of the sand-shale interface relative to the well:

$$\tan(\theta) = O/A = 0.7/6 \ \text{ or } \ \theta \approx 6°$$

The actual deviation at 2800 m is 86° so the true bed dip could be almost horizontal.

The density does not really follow the predictions and even the neutron shows a more complicated behaviour than the simple model would suggest. For this particular case this can mainly be explained by the fact that both density and neutron show more variation than the gamma ray, particularly in the sand, so that the change as the well cuts through the interface is modified by the natural variation in these properties. In some cases the density response might also be affected by the fact that the density tool's volume of investigation is quite narrow. This will not apply here however because the tool was rotating at over 150 rpm when drilling the interval shown. Combined with a rate of penetration of 25 m/h this means the tool rotates approximately 40 times for each data point.

Arguably, the resistivity logs show an even greater departure from the predictions of our simple model. Most of the curves show very high-resistivity spikes in the vicinity of the point where the bed boundary cuts the well. These are the manifestation of a well-known phenomenon, peculiar to propagation resistivity logs, known as 'polarisation horns'. They are a natural result of the physics of the measurement and basically are caused by an electrical charged layer that is induced at the bed boundary. This creates a secondary field, which adds to the signal the tool uses to determine resistivity. The resulting reading is completely erroneous and attempting to use it to find saturation will produce the wrong answer. This illustrates a key message of the book, that simple conceptual models are helpful for understanding properties but they are only models and do not include all the physical and geological complexity.

If one ignores the horns, the multiple resistivity curves show the influence of the shallow dipping boundary nicely. At the upper boundary the shallowest reading curve – a short spaced, high-frequency phase curve named RH16P (heavy black coding in Fig. 15.2) starts to respond to the sand some 3 m further into the hole than the deeper reading RH48P (heavy black dotted coding). The even deeper reading low frequency curve RL48P (coded red) actually reads several Ohm metres higher in the shale than any of the high-frequency curves. This

may be a result of the well path being close enough to the bed boundary for the deepest resistivities to be influenced by both lithologies.

15.3   FORMATION ANISOTROPY AND THIN BEDS

Traditionally, courses and books on logging tools introduce resistivity tools early on and leave the density and neutron tools until much later. In this book however the order was reversed with density being the first tool to be discussed in any detail and resistivity being left to the end of Chapter 5. The reason was that whereas density is a scalar property that gives the same reading no matter how the tool is orientated with respect to bedding resistivity depends on the path the current follows through the formation. This means that in anisotropic formations the measurement depends on the orientation of the tool even in a thick bed. This was discussed in Sections 5.6 and 13.4 a model for predicting how the induction resistivity depends on the angle between the tool – or borehole – and bedding was presented. This only applies to the induction tool but the key equation was:

$$C_{\text{ild}} = C_{\text{parallel}}(\cos^2(\theta) + (C_{\text{normal}}/C_{\text{parallel}})\sin^2(\theta)) \tag{13.2}$$

where θ is the dip relative to the well. In this section, we are only really interested in the cases where θ is either zero or 90°. This means that the resistivity (conductivity) is equal to the first term or the second term, respectively.

For the vertical well:

$$C_{\text{ild}} = C_{\text{parallel}}$$

In a truly horizontal well in the same formation Eq. 13.2 becomes:

$$C_{\text{ild}} = C_{\text{normal}} = 1/R_{\text{normal}} \tag{15.1}$$

The formulae for C_{parallel} and R_{normal} are functions of the resistivity (conductivity) of the shale and sand components and the sand shale–shale ratio (they were given in Section 13.4). Substituting, gives for the horizontal well:

$$C_{\text{ild}} = 1/((L \times R_{\text{t}} + (1 - L)R_{\text{sh}}) \tag{15.2}$$

For the vertical well the equivalent equation was:

$$C_{\text{ild}} = C_{\text{sh}}(1 - L) + C_{\text{t}}L$$

To make further progress we need to substitute in some actual numbers, we will also revert to resistivity as that is what the log uses. For the sake of argument assume a sand–shale ratio of 50%, $R_{\text{sh}}=1\,\Omega\text{m}$, $R_{\text{t}}=10\,\Omega\text{m}$ these values give R_{ild} of 5.5 Ωm for the horizontal well and 1.8 Ωm for the vertical. There should be no surprise here but it does show that for directional measurements in general different values will be obtained in low- and high-angle wells. For much the same reason different tool types can give different readings in the same well. A simplistic model of the shallow laterolog suggests that in a vertical

well the current is forced to flow through both the sand and shale layers and the resistivity is given essentially by Eq. 15.2. In a horizontal well the current has the opportunity to flow preferentially in the more conductive shales and the measured resistivity will be lower. In reality life is more complicated than that and to properly model a laterolog, advanced numerical methods are needed. But as a tool to explain or predict the changes between vertical and horizontal wells this model is quite adequate.

Resistivity is the most obvious measurement that depends on the orientation of the well but any anisotropic property will show the same general behaviour. An important example is the sonic log.

15.4 CONCLUSIONS

This really is the end of the book. The reader will undoubtedly be able to name whole topics that have been barely mentioned: fracture characterization, for example. Certainly, many readers will feel strongly that not nearly enough emphasis has been put on topics that they feel are of particular importance. I cannot argue with this but to paraphrase a popular maxim 'If you cannot please all the people all the time then at least try to disappoint everyone the same'. The purpose of the book is to show how to achieve a balance between the rigorous principles that ultimately determine the petrophysical properties and how our measuring instruments respond to them on the one hand and what is realistically achievable with limited time and resources on the other. Real wells do not drill through perfectly homogenous beds with boundaries that are mathematical planes normal to the well bore. Real pore systems do not consist of bundles of parallel-sided capillaries. I strongly believe these idealized systems are essential to make sense of our measurements and observations. But there will be a point where the model has been pushed too far. This typically occurs when the assumptions on which it is based become invalid, so it is incumbent on all of us to understand how a particular equation or technique was derived. Consider the Thomas–Steiber method for interpreting thin beds. This is based on a series of strict assumptions that the authors have carefully listed in their original paper, if one or more of these is invalid the method will not produce valid results and will probably produce an overly optimistic view of the reservoir.

Page left intentionally blank

Bibliography

The bibliography is not intended to be an exhaustive list of references rather I have included a few of the most important papers that heralded a new tool or technique (and that have stood the test of time). I have had to consult most of these papers at various times during my career and I consulted many of them again whilst writing this book. There are one or two however which I admit I have never read but are so important that for completeness they have to be included (Gassman's paper on the acoustic properties of porous solids is a case in point). Most of these papers were written decades ago but they are still completely relevant today. They are also of a far higher standard than most contemporary papers.

There are also some textbooks that go into specific topics in more depth and may prove useful for anyone wanting to tackle a specific problem. A lot of 'further reading' material, not specifically referenced here, is available for free on the websites of the logging companies. Most of these reports are essentially case studies and some would argue they are thinly disguised sales brochures but they do normally offer clear explanations of tools and techniques and rarely stint on high-quality graphics. They are also available for free although some companies insist you register to access the material.

I have accompanied most of the references with a sentence or two explaining its significance.

Chapter 1

G.E. Archie "Introduction to Physics of Reservoir Rocks" AAPG Bull. *34*(1950) p. 943.

This is the paper that introduced the word 'Petrophysics' in the public domain. As with all Archie's papers it is very readable and was in many ways ahead of its time.

L.A. Allaud and M.H. Martin "Schlumberger: The History of a Technique" (1978).

Published by John Wiley.

This book is long since out of print but is an interesting account of the history of the Schlumberger wireline business from the early days of surface resistivity surveys to the post-war period during which most of the conventional logs were developed. Part history, part biography and part technical it is not a

389

great read but it does capture the excitement of the early days of a key technology. A few copies are available on Amazon.

ISBN-13: 978-0471016670

D. Taub and E.C. Donaldson "Petrophysics: Theory and Practice of Measuring Reservoir Rock and Fluid Transport Properties" 3rd edition (2011) (first published 1996).

ISBN-13: 978-0123838483

Although nearly 1000 pages long and often criticized for being too 'mathematical' this is arguably the definitive textbook on the fundamentals of petrophysics.

Chapter 2

J.W. Neasham "The Morphology of Dispersed Clay in Sandstone Reservoirs and it's Effect on Sandstone Shaliness, Pore Space and Fluid Flow Properties" SPE 6658 SPE In the 1977 Technical Conference and Exhibition Proceedings.

M.D. Wilson and S.D. Pittman "Authigenic Clays in Sandstones: Recognition and Influence on Reservoir Properties and Palaeoenvironmental Analysis" J. Sed. Pet. *47* (1977) p. 3.

Two papers with figures that have subsequently appeared in countless papers, reports and training manuals that summarize the subdivision of 'shales' into 'structural', 'dispersed' and 'laminated' forms.

Chapter 4

M. Rider and M. Kennedy "The Geological Interpretation of Well Logs" 3rd edition(2011).

Published by Rider-French Consulting Ltd.

ISBN-13: 978-0-9541906-8-2

As the title suggests this book is ultimately about using logs to make inferences about geological processes and events such as environments of deposition and flooding surfaces. Nevertheless much of the book is devoted to describing the function and characteristics of the different types of log. There is more detail on specific tool types than in this book.

Chapter 8

G.E. Archie "The Electrical Resistivity Log as an Aid in Determining Some Reservoir Characteristics" Trans. AIME *46* (1942) p. 54.

This is the classic paper that introduces the Archie equation. It is well worth making the effort to read (at the time of writing it was available on-line for free).

H.J. Hill and J.D. Milburn "Effect of Clay and Water Salinity on Electrochemical Behaviour of Reservoir Rocks" Trans. AIME *207* (1956) p. 65.

This empirical study extended Archie's work to 'shaly sands'. These are the measurements that showed the need for shaly-sand equations and in particular led to the Waxman–Smits equation.

A. Poupon and J. Leveaux "Evaluation of Water Saturation in Shaly Formations" The Log Analyst (1971) p. 3.

This became the Indonesia equation. The Waxman–Smits equation appeared before it but in this book equations like the Indonesia equation are discussed first.

M.H. Waxman and L.J.M. Smits "Electrical Conductivities in Oil-Bearing Shaly Sands" SPE Paper 1863-A Trans. AIME *243* (1968) pp. 107–122.

The paper introducing the Waxman–Smits equation.

M.H. Waxman and E.C. Thomas "Electrical Conductivities in Shaly Sands II. The Temperature Coefficient of Electrical Conductivity" SPE Paper 4094 J. Pet. Tech. (1974) p. 213.

This paper describes measurements that quantify the 'B' in B.Qv. It is worth reading for the subject matter and as an example of how an experimental paper should be written.

C. Clavier, G. Coates and J. Dumanoir "Theoretical and Experimental Bases for the Dual Water Model for the Interpretation of Shaly Sands" SPE Paper 6859. Presented at the 1977 SPE Conference.

This is the Dual Water Model which is not discussed in the book but is so widely used that a reference is justified.

Chapter 10

M.C. Leverett "Capillary Behaviour in Porous Solids" Trans. AIME *42* (1941) p. 159.

This introduces the Leverett J-function and more generally introduces the whole concept of saturation-height modelling.

Chapter 11

P.C. Carman "Fluid Flow Through Granular Beds" Trans. Inst. Chem. Eng. London, *15* (1937) pp. 150–166.

This is the paper that basically introduces the Kozeny–Carmen equation. Although Carman is the sole author he actually modified an earlier equation developed by J. Kozeny (1927).

R.L. Morris and W.P. Biggs "Using Log-Derived Values of Water Saturation and Porosity" SPWLA Logging Symposium Transactions, 1968.

A. Timur "An Investigation of Permeability, Porosity and Residual Water Saturation Relationships" SPWLA Logging Symposium Transactions, 1968.

A. Timur "Pulsed NMR Studies of Porosity, Moveable Fluid and Permeability" J. Pet. Tech. (1969) p. 775.

Three papers describing general functions that relate permeability to porosity and irreducible water saturation. They are specific cases of Kozeny–Carman.

Chapter 12

R.M. Bateman and C.E. Konen "The Log Analyst and the Programmable Pocket Calculator" The Log Analyst (1977).

Despite the title this paper describes one of the best ways to deal with complex lithology: a density–neutron method that now bears the author's names. The basic principles are explained in this book but this paper includes the algorithm for implementing the technique.

M.C. Kennedy "Solutions to Some Problems in the Analysis of Well Logs in Carbonate Rocks" Chapter 6 in AAPG methods in Exploration 13 "Geological Applications of Well Logs" (2002).

Published by the AAPG and edited by Mike Lovell and Niel Parkinson.

A general account of carbonate log analysis including ways to distinguish limestone and dolomite.

Chapter 13

E.C. Thomas and S.J. Stieber "The Distribution of Shale in Sandstones and its Effect upon Porosity" SPWLA Logging Symposium Transactions, 1975 (paper T).

This is the paper that explains the principles and practical implementation of the Thomas–Stieber method. Note that Thomas also made important contributions in the field of shaly-sand saturation equations (Chapter 8).

C. Skelt "The Influence of Shale Distribution on the Sensitivity of Compressional Slowness to Reservoir Fluid Changes" SPWLA Annual Logging Symposium Transactions, 2004 (paper V).

This is an interesting paper that discusses the acoustic properties of shaly-sand systems, including thin bed pays.

Chapter 14

F. Gassman "Uber die Elistizitat Porozer Medien" Quarterly Reports of Research from the Zurich Research Centre **96** (1951) p. 1.

Obviously this is an old paper written in German but as it is still applied every day it must be included here.

R.J. Runge, A.E. Worthington, D.R. Lucas "Ultra-Long Spaced Electric Log (ULSEL)" Paper H in the Transactions of the 10th Annual SPWLA Symposium (1969).

Paper introducing the – then – new ultra deep reading resistivity ULSEL tool. It is quite technical but does include a couple of case studies. An easy to read article on the ULSEL appeared in January 1991 'Oilfield Review' and can be found on the Schlumberger website.

Index