Pergamon

Journal of Structural Geology, Vol. 18, Nos 2/3, pp. 135 to 146, 1996

0191-8141/96 $15.00 + 0.00

0191-8141(95)00108-5

Sampling of fault populations using sub-surface data: a review

GRAHAM YIELDING, TIM NEEDHAM and HELEN JONES

Badley Earth Sciences Ltd, North Beck House, North Beck Lane, Hundleby, Spilsby, Lincolnshire PE23 5NB, U.K.

(*Received* 3 *January* 1995; *accepted in revised form* 22 *August* 1995)

Abstract—In favourable circumstances, seismic reflection data can give an unrivalled view of faulted rocks in the sub-surface, imaging features down to the seismic resolution (typically 20–30 m). The lack of finer detail can, in part, be addressed by analysing well cores through the same rock volume. Samples of fault populations from such data often exhibit power-law size distributions where 'fault size' can be trace-length or fault-displacement. Analysis of a synthetic fractal model (the 'fragmentation model') demonstrates that changing the dimension of the sampling domain (e.g. volume to plane, plane to line) changes the power-law exponent of the sample's size distribution. The synthetic model also suggests how best to treat faults that extend out of the sample area, and illustrates potential problems in comparing samples from very different scales (e.g. regional and detailed mapping).

Analysis of a variety of interpreted seismic-reflection data sets has provided a range of power-law exponents for different sample types:

(i) fault-trace lengths (two-dimensional samples): −1.1 to −2.0;
(ii) fault-trace maximum displacements (two-dimensional sample): −1.0 to −1.5;
(iii) 'arbitrary' displacements (one-dimensional sample): −0.5 to −1.0.

Fault-trace lengths are very sensitive to truncation (resolution) effects, and tip regions should be re-assessed using displacement gradients. Maximum displacements, and displacements obtained by line-sampling, are much more robust attributes. Well data are useful in constraining the extrapolation of populations to smaller scales. Fault populations scale differently than earthquake populations, because the latter represent only the instantaneous deformation, whereas fault populations represent the deformation accrued over geological time. A valuable data set to clarify these relationships would be a true three-dimensional sample of faults in an actively-deforming area.

INTRODUCTION

In recent years there has been considerable interest in the scaling properties of faults and fault populations. A commonly-observed characteristic of many sampled fault populations is that the size–frequency distribution is described by a power-law. Denoting the 'size' of a fault by S (e.g. its length or maximum displacement), the number N of faults having a size greater than or equal to S is given by:

$$N = aS^{-D}. \quad (1)$$

The variable a is a measure of the size of the sample, and the power-law exponent D is often referred to as the fractal dimension of the population. Larger values of D imply a greater proportion of smaller faults in the population.

The results of this work have been applied in several distinct but related areas. First, the scaling properties have been used to constrain and test models of fault growth (e.g. Walsh & Watterson 1987, 1992, Cowie & Scholz 1992a,b). Second, the properties of the fault population have been used to make predictions about the overall brittle strain in a region (e.g. Scholz & Cowie 1990, Walsh *et al.* 1991, Jackson & Sanderson 1992, Marrett & Allmendinger 1992). Third, observations of a restricted part of a fault population have been extrapolated to predict the numbers of unobserved faults in the same area, in order to constrain their likely impact on hydrocarbon production (e.g. Heffer & Bevan 1990, Yielding *et al.* 1992, Gauthier & Lake 1993). A critical step in the latter two areas (i.e. strain summation and prediction of small-scale faulting) is the deduction of the fault population in the three-dimensional rock volume from observations made on cross-sections or maps. Marrett & Allmendinger (1991) showed that this 'sampling domain' effect can be accounted for, provided that the fault population is governed by power laws and the displacement–length relationship is also known.

A significant proportion of recent fault-population observations has been derived from interpretation of oil industry sub-surface data, i.e. seismic reflection surveys and wells. Traditionally seismic data has been 'two-dimensional', comprising a grid of shot-lines each of which is essentially a vertical cross-section. The last decade has seen an explosive development in 'three-dimensional' seismic acquisition, where shot-lines are sufficiently close (e.g. 10–20 m) that the data can be regarded as being a data volume rather than a cross-section. Such continuous sampling of many cubic kilometres of rock generates significantly improved imaging of sub-surface structure (e.g. Dalley *et al.* 1989), and indeed generally provides a better overall structural picture than could be obtained by onshore surface map-

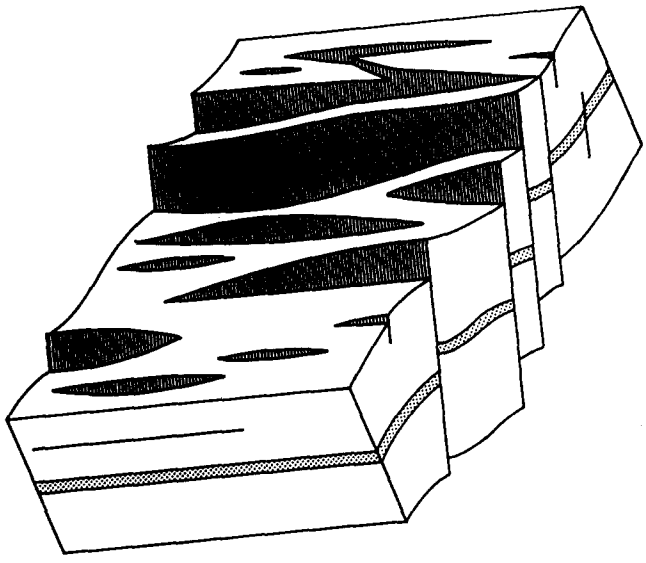

Fig. 1. Block diagram of a faulted rock volume. A three-dimensional sample of this fault population would require counting of all the faults in the volume. In practice, interpretation of sub-surface data typically provides a suite of cross-sections and a suite of 'horizon maps'. Counting fault offsets on one horizon on a cross-section constitutes a one-dimensional sample of the fault population. Counting fault traces on a one horizon map (e.g. upper surface of block) constitutes a two-dimensional sample of the fault population.

ping. The unavoidable drawback, however, is the resolution of the seismic acquisition technique. At depths of 3 km the seismic resolution is typically 30 m (although in favourable circumstances may be as low as 10 m): fault offsets less than this are essentially invisible on seismic reflection data. Wells provide the opportunity to sample faults in the sub-surface at a much finer scale, particularly if cores are recovered. However, the width of the core barrel (typically *ca* 10 cm) imposes an upper limit on the scale of information that can be obtained. Interpolation of the fault population across this 'data gap' must be performed carefully, since wells and seismic reflection data may represent different sampling domains.

The principal objective of this contribution is to review a number of examples of fault-population samples obtained from sub-surface data. Interpretation of a seismic data set typically involves the 'picking' or mapping of a number of stratigraphic reflections (horizons) at different levels within the subsurface. The 'results' of interpretation in an area of faulting are therefore:

(i) a suite of cross-sections (seismic sections) on which fault offsets (apparent throw) can be measured;

(ii) a suite of structure maps, showing fault traces at each horizon: for each trace a trace-length and a maximum displacement can be measured.

These two aspects of the interpretation correspond, respectively, to one- and two-dimensional samples of the fault population (see Fig. 1). To be characterized completely (a 'three-dimensional sample'), the fault population would have to be interpreted and measured throughout the three-dimensional rock volume covered by the data set. To our knowledge, such analysis has not been performed, since it would be much more time-consuming than the degree of subsurface mapping usually required commercially. In practice then, we use fault maps (the intersection of the faults with a sample plane) or line samples (the fault offsets seen at one horizon on a cross-section).

It is important to understand the effects of these sampling procedures on real data distributions. A good way to do this is to use an idealized distribution to explore sampling strategies relevant to the kind of fault-population samples that are available in practice. The next section of this contribution uses a fractal geometrical model (which has a number of properties in common with real fault systems) to illustrate:

(i) the difference between sampling domains (i.e. volume vs maps vs line);

(ii) the effect of incompletely-sampled faults; and

(iii) a possible difficulty in comparing samples at different scales (e.g. regional and detailed mapping).

These illustrations provide a framework within which the real fault samples of the subsequent sections can be considered.

SAMPLING OF A THREE-DIMENSIONAL FRACTAL MODEL

The model that we use to examine sampling strategies is the 'fragmentation model' of Sammis *et al.* (1987) (Fig. 2). This comprises a self-similar array of cubes that has a fractal dimension of 2.58, i.e. the cumulative frequency distribution of the cube sizes (and bounding surfaces) is controlled by a power law with exponent −2.58:

$$N = aL^{-2.58}, \qquad (2)$$

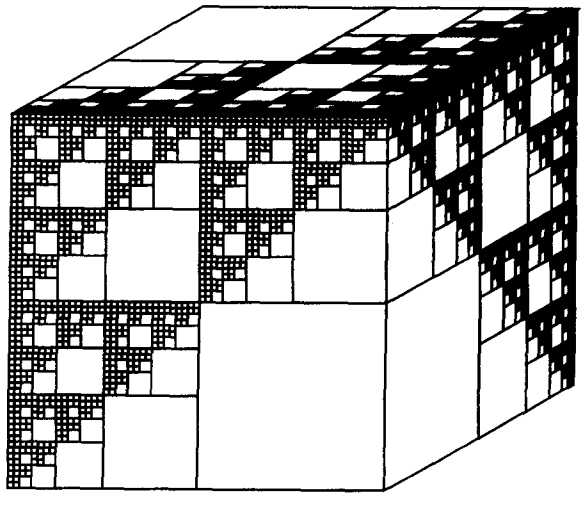

Fig. 2. The fragmentation model of Sammis *et al.* (1987). The cubes comprise a population with a fractal dimension of 2.58, i.e. $N = aL^{-2.58}$, where N is the number of cubes having a size greater than or equal to L. The planes bounding the cubes have the same distribution.

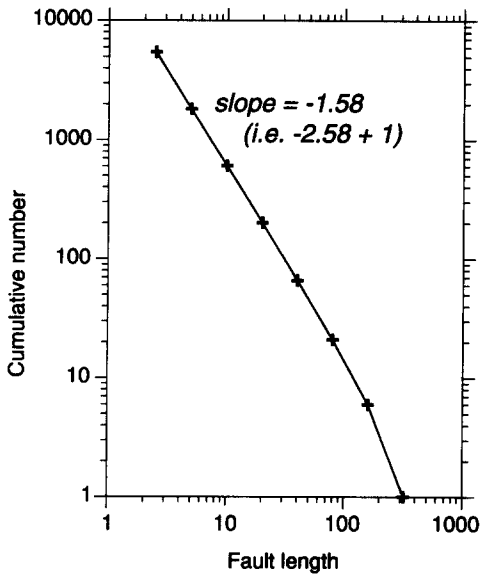

Fig. 3. Cumulative frequency distribution for line lengths (arbitrary units) sampled on one face of the model in Fig. 2. This is analogous to a map sample of fault planes. The data follow the power law $N \sim L^{-1.58}$. Note that this exponent differs by 1 from the exponent for the population of planes in the volume (Fig. 2).

where N is the number of objects having a size greater than or equal to L. Alternatively,

$$\log N = (\log a) - 2.58(\log L). \quad (3)$$

Sammis *et al.* (1987) showed that some fault gouges are statistically identical to this model, in that the distribution of particle sizes is controlled by the same power law. They proposed a mechanism whereby such a distribution can arise from repeated tensile splitting of grains. Whilst we would not necessarily expect the same mechanism to apply to the growth of a fault population, the model is a useful analogue for fault–length populations because:

(a) it is three-dimensional and yet simple;

(b) it contains a population of planes whose frequency distribution, like faults, is power-law controlled;

(c) the spacing distribution of the planes is controlled by the same power law, and this is possibly true of faults (e.g. Gillespie *et al.* 1993).

These similarities between the fragmentation model and real faults are apparent to some extent by visual comparison of one face of the model with a fault map: both exhibit significant clustering of lines/faults and corresponding areas where lines/faults are absent. The following examples explore the properties of the population of lines on one face of the model. This is a two-dimensional sample of the population of planes in the volume, and is analogous to the population of fault traces on a map.

Figure 3 is a cumulative frequency diagram of the 'fault' lengths on one face. Using log–log axes, the sample plots as a straight line of slope −1.58, indicating that the cumulative frequency is controlled by a power law with exponent −1.58. Note that this is exactly 1 greater than the exponent for the population of planes in the volume. This is a general property of isotropic fractal sets (Mandelbrot 1982), i.e. changing the sampling domain by 1 (e.g. from three to two dimensions) changes the measured fractal dimension by 1. There is a smaller proportion of small 'faults' represented on the sample plane than in the volume because the probability of a 'fault' plane being cut by the sample plane is directly proportional to the 'fault' size. If we can measure the population on a sample plane, and the population is isotropic, then the population distribution in the volume can be obtained simply by decreasing the exponent by 1; this relationship was described by Marrett & Allmendinger (1991).

The data distribution in Fig. 3 is slightly curved for the longest 'faults' because the sample area is physically bounded by the edges of the model. This 'censoring' or 'finite range' effect is explored further in Fig. 4, where a square sample area is placed randomly within the confines of the model. Figure 4 shows three cumulative frequency curves. The central curve comprises lengths for all of the faults within the sample area: for those faults that extend beyond the sample area boundaries the measured length is that part which lies within the sample area. Although a little irregular, this curve gives a good approximation to the 'correct' slope of −1.58. The left-hand curve omits all length measurements where only part of the fault is within the sample area, and has a rounded form (long faults are under-represented). Conversely, for the right-hand curve the complete length of the part-faults has been measured (i.e. including the extension out of the sample square), and here the longest faults are over-represented. These results suggest that the most reliable and simplest way of dealing with faults that extend out of a sample area is to include that part which can be measured, rather than discarding them completely.

Real observations of faults are often restricted, for practical reasons of sampling, to a limited range of sizes (typically one or two orders of magnitude). In order to characterize the population over a wide scale range, it is therefore common to compare a number of samples from different scales on one cumulative frequency graph (e.g. Heffer & Bevan 1990, Yielding *et al.* 1992, Scholz *et al.* 1993). This process is simulated in Fig. 5. On the fractal model, successively smaller squares indicate a sequence of sample areas; these might be thought of as regional mapping and detailed exposure mapping, for example. On the graph, cumulative frequencies are normalized per unit area (i.e. cumulative fault density) in order to combine the samples into one population. Each individual sample accurately reflects the true population slope of −1.58, but the overall trend of the combined data has a slope of −2.00. This latter slope has nothing to do with the fault population but is simply the dimension of the sampling domain. Successive sample lines on the graph are offset from one another because smaller sample squares are never positioned in areas that have no faults at all: therefore the normalized fault density must be higher in the smaller sample areas. The slope of −2 arises in the following way. In moving from one sample to the next, the length of the sample area

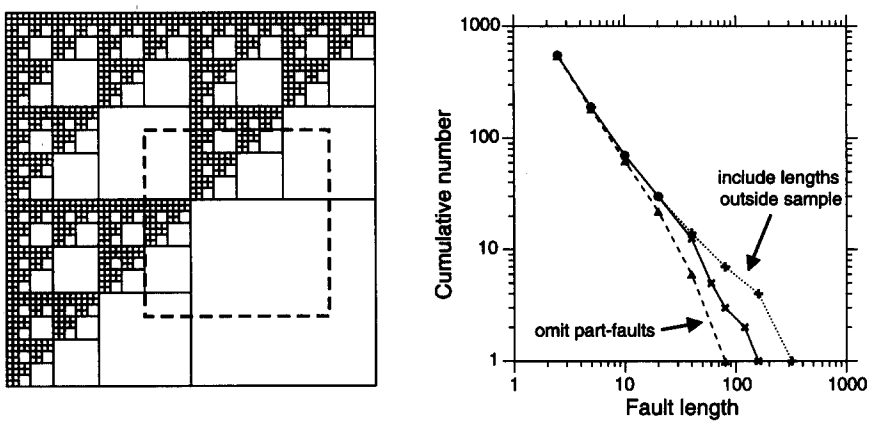

Fig. 4. The effect of sample censoring on 'fault' length distributions. The sample area is shown as a dashed outline on the fractal figure: many of the faults extend out of the sample area. Three different samples are plotted on the graph. The central line represents all faults in the sample area; for those that extend beyond the sample area the length measured is that part that lies within the area. For the left-hand curve, faults that extend out of the area are omitted from the sample. For the right-hand curve, the true length of those faults was measured, i.e. including the part outside the sample square. Note that the 'best' sample (i.e. the line which best approximates a slope of −1.58) is the one that includes part-faults.

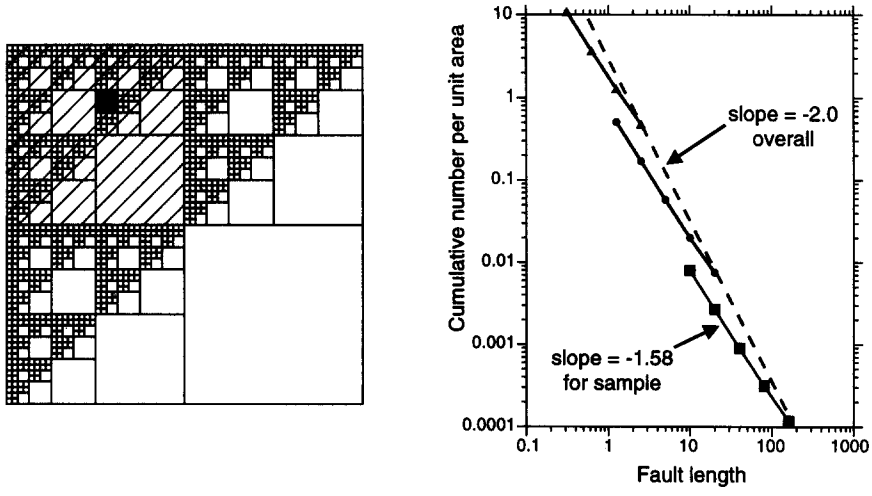

Fig. 5. Comparison of fault sampling at different scales. The diagonally-ruled and stippled squares and black dot indicate successively smaller sample areas on the fractal model, giving rise to successive samples on the graph. Each individual sample accurately reflects the population exponent of −1.58. However the samples are offset from one another on the plot because the average fault density is higher in the smaller sample areas (large unfaulted areas are not present at smaller scales). The overall slope of the combined sample points is −2.0, which simply reflects the dimension of the sample plane.

might change by a factor s and so does the length of the longest fault in the sample. However the area of the sample changes by s^2 and so the fault density at the largest sampled fault size ($n = 1$) will change by s^2. Hence between samples the change in fault density varies as the square of the fault length. This biasing of the densities in successive sub-areas was also recognized by Sammis *et al.* (1987) in their study of fault gouge (see their appendix 1). A similar effect is likely with fault maps at different scales, because a detailed fault map will generally be at a locality with more faulting rather than a locality where fault density is low.

Figure 6 shows the effect of another change in the sampling domain, this time to line (one-dimensional) sampling. Five arbitrarily-placed sample lines have been drawn on the fractal model, and 'fault-trace lengths' have been measured for all lines intersecting each sample line. This procedure is analogous to measuring the lengths of faults intersected by a well (though in practice this would not be possible). The resultant cumulative frequency plots show significant scatter depending upon the exact position of the sample line through the model, but they cluster around a slope of −0.58. Once again a change of 1 in the sample slope reflects a change of 1 in the dimension of the sample domain. There is a smaller proportion of small faults represented on the sample line than on the plane because the probability of a fault trace being cut by the sample line is directly proportional to its length.

The preceding examples illustrate a number of points relevant to sampling of faults:

(i) a cumulative frequency distribution for fault-trace lengths on a map gives an apparent fractal dimension that differs by 1 from the fractal dimension of planes in the volume (for an isotropic population);

(ii) an equivalent sample for a line (a well) gives an apparent fractal dimension that differs by 2 from that of the fault planes;

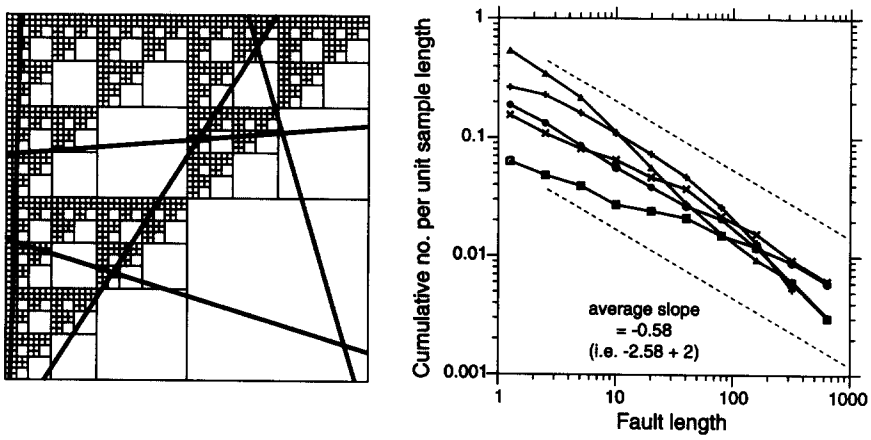

Fig. 6. Line sampling of the fractal model. Five lines are arbitrarily located on the model, and the data curves on the graph show the length distribution for all faults cut by each line. Sample slopes show considerable scatter depending on precise sample-line location, but they are clustered about a slope of −0.58 (denoted by dashed lines). Note that this slope differs by 1 from that obtained by two-dimensional sampling (−1.58) and by 2 from that of the three-dimensional model (−2.58).

(iii) where fault traces extend out of the sample area, it is better to include these part-lengths rather than discard them;

(iv) care should be taken when comparing samples from different scales of interpretation, because detailed maps are often from those areas where faulting is most intense.

The next section examines examples of real fault populations and discusses them in the light of the model sampling described above.

FAULT POPULATIONS FROM SUB-SURFACE DATA

The samples described in this section include new and previously published examples derived from interpretation of seismic data sets, together with supporting well cores. Two-dimensional samples (from maps) comprise measurements of fault-trace lengths and fault-trace maximum displacements; cross plots of these data are also useful in defining fault scaling properties. One-dimensional (line) samples comprise measurements of fault displacement on seismic lines.

Fault-trace lengths

Figure 7 shows three examples of fault-length samples derived from interpretations of North Sea three-dimensional seismic reflection data sets. Fault-trace lengths have been extracted from horizon maps and therefore these are two-dimensional samples. The longest fault trace represented varies from *ca* 4 km in the smallest data set (*ca* 50 faults) to *ca* 25 km in the largest data set (nearly 300 faults). The smallest and largest data sets are from Jurassic horizons affected by end-Jurassic extension in the northern North Sea; the third is from a Palaeozoic interval affected by two or more episodes of Mesozoic extension. The smallest data set barely covers one order of magnitude of fault length, and its cumulative frequency curve has an irregular form; this sample is

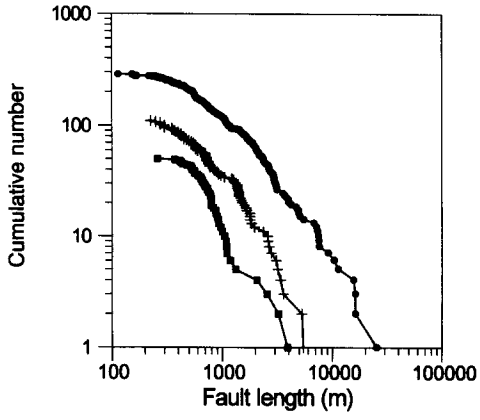

Fig. 7. Examples of fault-trace length samples derived from seismic interpretations of North Sea fields. The left-hand data set is small, has restricted range, and has an irregular form. The two larger data sets (A1 and F1 of Table 1) both exhibit a gently curved form because of resolution effects (see Figs. 8 and 9).

too small to be a reliable guide to the actual fault population. The larger two samples display a more regular data distribution, but both are curved rather than straight as would be expected for a power law. Heffer & Bevan (1990) explained this as a degradation of an underlying power law because of the limited resolution of the seismic data. The tip regions of faults are not detected by the seismic interpreter because the fault throw has dropped below the seismic resolution (typically 30 m), and therefore the mapped fault length is always an underestimate. This effect is proportionately greater for smaller faults because a larger fraction of their length has a throw below resolution. Hence the real population curve becomes shifted progressively to the left for smaller fault lengths.

A corrected population curve can be retrieved from such a sample if the tip regions can be reconstructed. This can be done if the lateral displacement gradient is simple. Recent models of fault growth (e.g. Walsh & Watterson 1987, Scholz *et al.* 1993) suggest a 'bell-shaped' displacement profile on isolated fault surfaces; to a first approximation the displacement gradient in the tip region is linear. Figure 8 shows an example of a

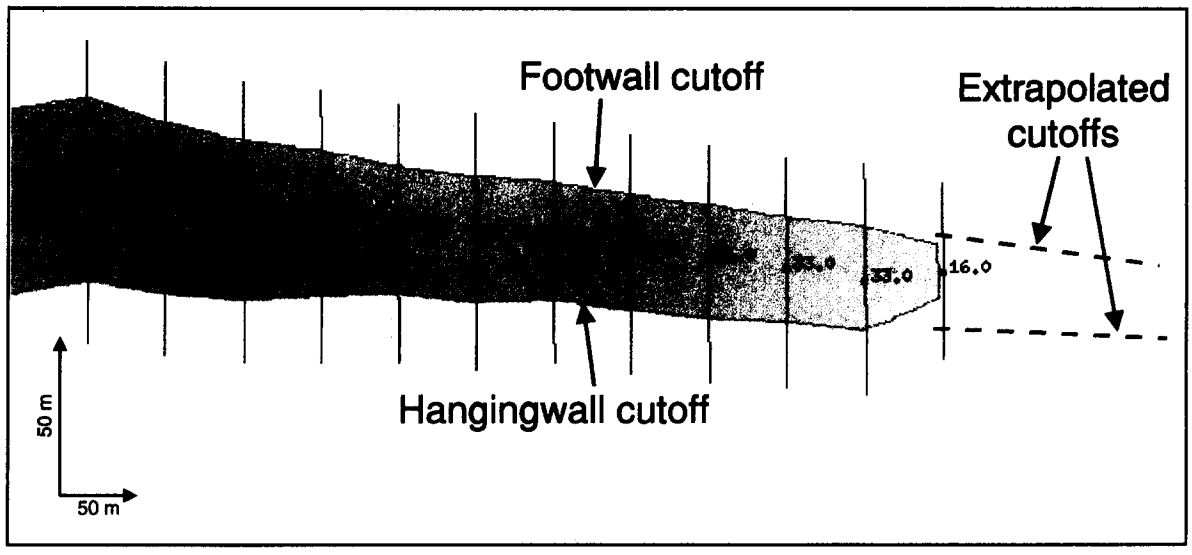

Fig. 8. Example of length truncation for a small fault. Display shows a vertical strike projection of interpreted horizon cutoffs on part of one fault. The fault throw is shown at 50 m intervals along the fault: at this horizon the maximum throw is at the left end of the diagram. At the interpreted end of the fault the throw appears to drop suddenly from 33 to 16 m and then to zero. The reading of 16 m is below seismic resolution (30 m) and is unreliable. The displacement gradient suggests that this fault may extend another 350 m beyond the mapped end; this end section was not interpreted because the throw is below the seismic resolution.

vertical strike projection of part of a small fault from the largest data set of Fig. 7. The footwall and hangingwall cutoffs of one mapped horizon are shown, together with interpreted throw values at 50 m intervals along the fault; the throw maximum at this horizon is at the left end of the diagram. At the mapped fault tip the interpreted fault throw drops abruptly from 33 to 16 m, before no longer being resolved. The last throw value (16 m) is unreliable because it is less than the seismic resolution in this data set (20–30 m). Extrapolation of the observed throw gradient along the length of the fault, however, would suggest that this fault probably extends for a few hundred metres beyond the interpreted tip. This unresolved tip region of the fault is likely to be between 250 and 500 m in length (depending upon how much of the observed fault is used in the extrapolation); 350 m would be a fair estimate. Re-interpreting all the faults and extrapolating their tips in this way allows a corrected data curve to be plotted, as in Fig. 9. This curve is much straighter than the raw data curve, and gives a power-law exponent of −1.37.

If all the faults have the same displacement gradient then the unresolved tip regions will all be the same length (for a given resolution), and this calculated length can potentially be used as a short-cut to a corrected length sample. This is illustrated by the right-hand curve in Fig. 9, where 700 m has simply been added to all the raw fault lengths (350 m at each end). Further steepening of the data curve is now apparent, with a best-fit slope of −1.59. However, this is likely to be an over-correction because a proportion of the faults link onto other faults rather than ending at tips.

For a total of 10 data sets for which we have extracted trace-length samples, the best-fit slopes (power-law exponents) lie in the range −1.18 to −2.04 (see Table 1 and Fig. 10a). Six of these lie in the centre of the range, at

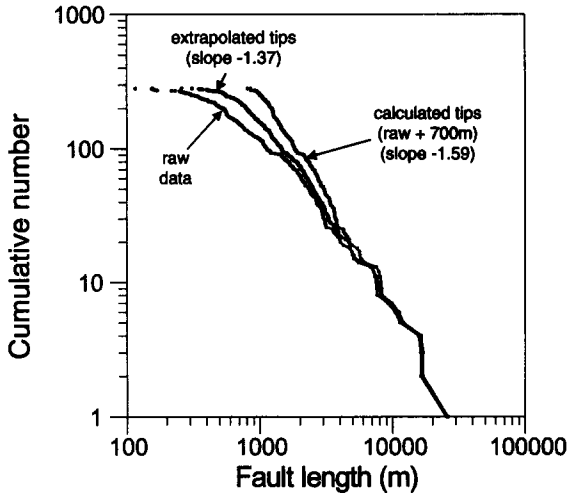

Fig. 9. Correction of raw trace-length data for resolution effects. The left-hand curve is the raw data (largest data set of Fig. 7, A1 in Table 1). The central curve represents lengths that have been individually corrected by analysis of displacement gradients, as in Fig. 8. The right-hand curve was obtained by simply adding a nominal tip length (350 m at each end) to all faults in the raw sample. Note the straightening and steepening of the curves when these corrections are applied.

−1.35 to −1.6. The two largest values (1.87 and 2.04 from data sets F1 and F2) are from an area of the U.K. Continental Shelf that has experienced at least two episodes of extension, whereas all other samples reflect a single episode of faulting (late Jurassic faulting affecting Jurassic horizons). A larger exponent implies a greater number of small (short) faults for every large (long) fault. It is tempting to suggest that further strain (reactivation) in an area of existing faults has required the growth of large numbers of shorter faults, to accommodate space problems around existing structures that might not be of optimum orientation for the new exten-

Table 1. Cumulative frequency exponents for fault-trace lengths and fault-trace maximum displacements, derived from mapped horizons in seismic reflection data sets. Data sets A, B, C, etc., indicate different seismic surveys; numbers (e.g. C1, C2, C3) indicate different stratigraphic horizons interpreted on the same survey. Samples F1 and F2 are from Palaeozoic intervals on part of the U.K. Continental Shelf that has undergone two or more episodes of extension. All other samples are from Jurassic horizons affected by end-Jurassic extension. Survey D is located within survey E

Data set	Trace length exponent	Maximum displacement exponent	Ratio *L/D*
A1	1.37	1.34	1.02
A2	1.42	1.25	1.14
B	1.5	1.45	1.03
C1	1.43	1.03	1.39
C2	1.49	1.29	1.16
C3	1.35	1.1	1.23
D1	1.18	—	—
D2	—	1.42	—
D3	—	1.25	—
E	1.75	1.48	1.18
F1	2.04	0.96	2.13
F2	1.87	1.04	1.8

sion. Further data from a variety of settings are required to test this suggestion.

The overall range of values for trace-length exponents in Fig. 10(a) is similar to published values obtained by field mapping as well as seismic-reflection mapping. For example, Okubo & Aki (1987) found a fractal dimension of about 1.3 for mapped fault traces in the San Andreas Fault System. Watterson *et al.* (1996) found values of 1.36–1.87 for fault traces in a compilation of U.K. coalfield maps, and Villemin *et al.* (1995) reported an average value of 1.4 for fault traces in the Lorraine Coal Basin. Scholz *et al.* (1993) reported a value of 1.3 for mapped fault traces in the Volcanic Tableland of eastern California. Heffer & Bevan (1990) suggested that a cumulative frequency exponent of *ca* 2.0 is consistent with many seismic reflection and map data sets. Gauthier & Lake (1993) found values of 1.10–1.69 from analysis of a North Sea data set. Barton (1995) analysed a variety of rock pavements and found fractal dimensions in the range 1.32–1.70.

In order to increase the size range of a sample, and thereby better constrain the population parameters, we can attempt to combine mapping from different scales of seismic survey. This is done in Fig. 11, using an extensive two-dimensional seismic survey (covering 1500 km^2) and a three-dimensional survey that covers a 220 km^2 sub-area within the two-dimensional survey. The samples in Fig. 11 come from the same stratigraphic level, a Jurassic horizon cut by end-Jurassic faulting. However, the two data sets do not obviously form a continuous curve; their slopes appear different and the fault density is higher in the sub-area. This is perhaps not surprising since the three-dimensional survey was shot

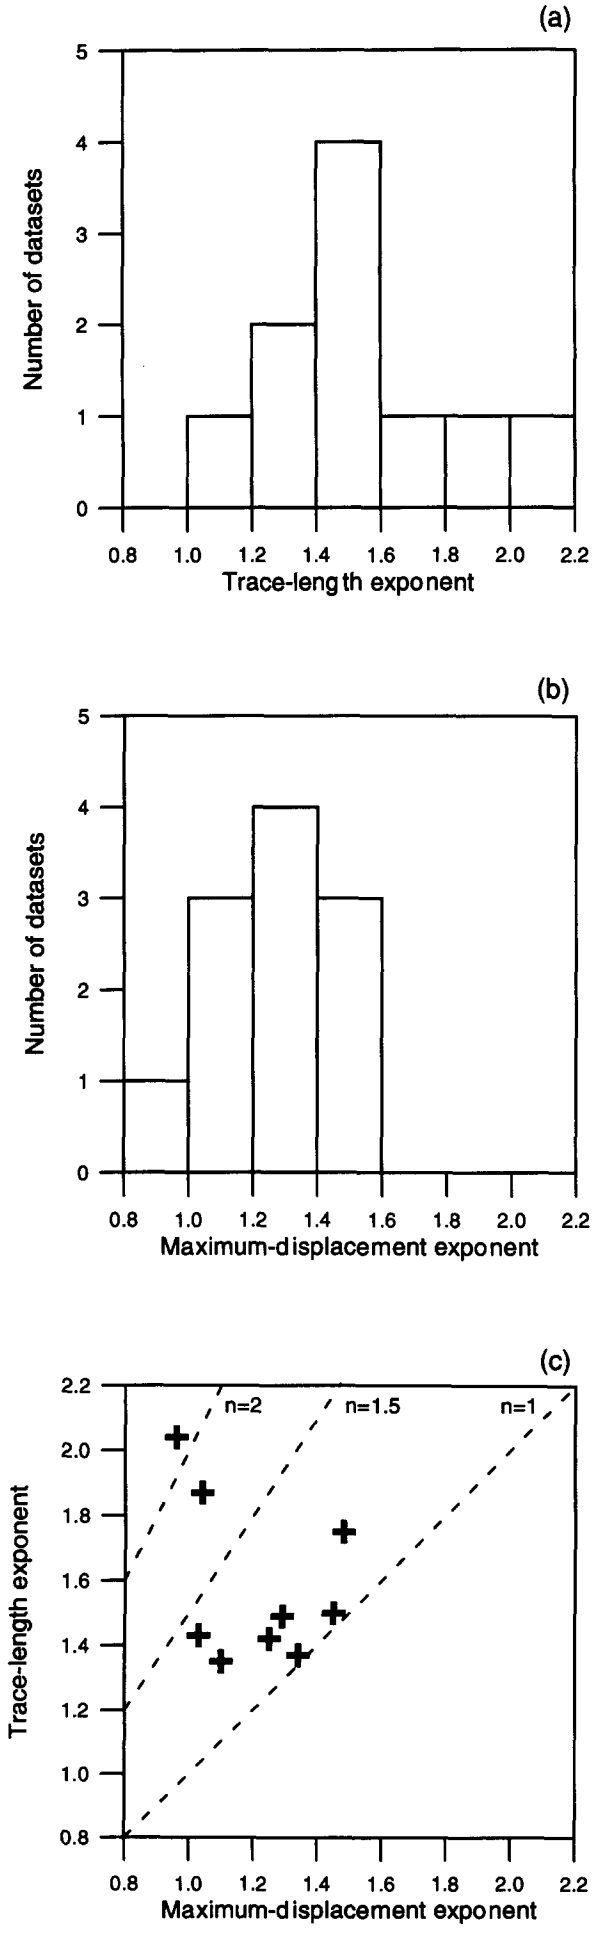

Fig. 10. (a) Histogram of trace-length exponents listed in Table 1. (b) Histogram of maximum-displacement exponents listed in Table 1. (c) Cross-plot of exponents displayed in (a) and (b). The ratio of the two exponents is a guide to the value of n in $D = cL^n$. (Ratios are listed in Table 1.)

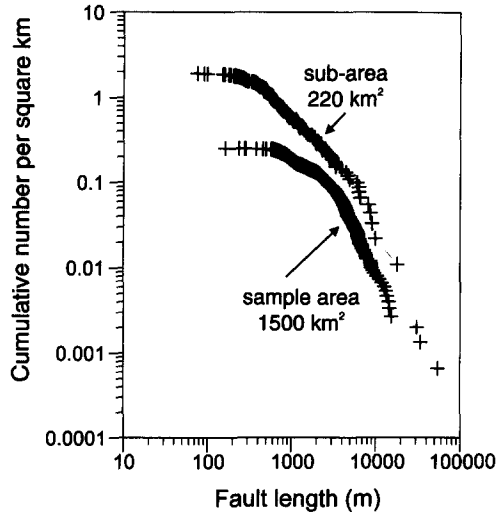

Fig. 11. Combination of fault-length samples at different scales of seismic mapping in the North Sea. The lower curve represents data from one horizon map from a large two-dimensional survey (1500 km^2) (data set E in Table 1, from Yielding *et al.* 1992). The upper curve represents data from the same horizon, mapped on a three-dimensional survey covering a 220 km^2 sub-area within the larger survey (data set D1 in Table 1). Fault density is higher in the sub-area and it is not clear how to relate the two samples, cf. Fig. 5.

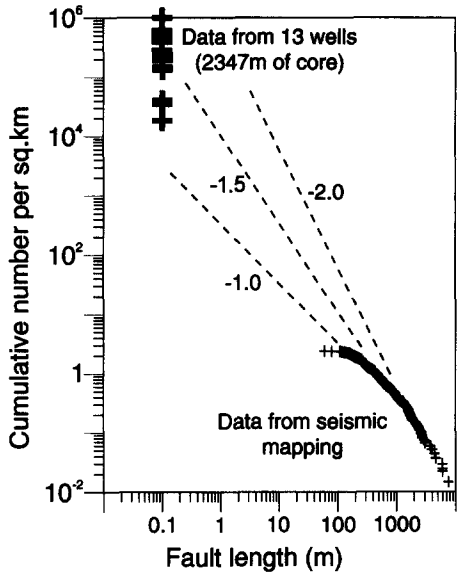

Fig. 12. Combination of fault-length data from the seismic and well scales in the same field in the North Sea. The seismic-scale measurements represent a reservoir horizon mapped on three-dimensional seismic data (data set B in Table 1). The core measurements are from the same reservoir layer, in a total of 13 wells scattered through the mapped area. The fault densities in core are derived from linear fault densities using the mean fault lengths calculated from the proportion of fracture terminations seen in the core slab. The seismically-mapped faults and core fractures are believed to be coeval because they are both restricted to pre-Cretaceous rocks and have similar azimuth distributions (see Needham *et al.* 1996). Although the well data show almost 2 orders of magnitude range in fault density, they help to constrain extrapolation of the population observed on seismic data.

to resolve the details in an area of more intense faulting. (The seismic resolution in the two-dimensional data set is about 20 ms (30 m) whereas that in the three-dimensional data set is about 10 ms (15 m).) The form of the data curves is similar to those in the model sampling of Fig. 5 and illustrates the difficulty of combining data sets even when from the same area. Sub-areas mapped in detail are rarely representative of the whole: they are usually mapped in detail because they are more complex, i.e. they have higher fault densities. A possible additional complexity in the samples in Fig. 11 concerns the longest fault-traces; faults with length > 10 km are likely to cut the entire seismogenic layer and therefore the population is no longer isotropic in three dimensions.

One potential way of obtaining a more representative view of smaller-scale faulting is to use well core data. In finely-laminated sedimentary rocks, small faults and fractures can often be readily identified, though the size of the core (typically *ca* 10 cm) places a limit on the maximum size of shear offset or fracture length that can be measured. If the fractures in a well core can be reasonably shown to be coeval with faults seen on seismic reflection data, then fault densities measured in core should be a useful constraint on the extrapolation of cumulative frequency relationships inferred from the seismic data. If a number of wells are available in a particular survey area, then sampling bias towards areas of locally-high fault density is likely to be avoided. Fault lengths are not directly measurable in core, since the majority of them extend beyond the limits of the core slab. However, the proportion of fault terminations can be used to estimate the mean fault length sampled by the core. A large proportion of terminations implies that the mean fault length is not much greater than the core width; conversely, a small proportion of fault terminations implies that the mean fault length is significantly greater than the core width. Dividing the mean fault length into the linear fault density gives an estimate of the areal fault density, for fault lengths greater than or equal to the core diameter (see also Sassi *et al.* 1992).

Well data collected in this way are displayed on Fig. 12. Over 2 km of core from a Jurassic reservoir in 13 wells in one North Sea field have been analysed to provide estimates of areal fault density at a scale of 10 cm fault length. These estimates vary by almost 2 orders of magnitude. At the scale of a well, the fault network is clearly very heterogeneous, with some wells being highly faulted and others having very few fractures. Nevertheless, the well data do constrain extrapolation of the seismically-mapped fault lengths (from a three-dimensional survey of the same field), which are 3–5 orders of magnitude longer than the core width. For the data set shown, a population exponent of *ca* −1.5 would be consistent with both the well and seismic data; this estimate is more reliable than if seismic data alone had been used. Preliminary estimates of fault-length populations using core and seismic data in this way (Yielding *et al.* 1992) had suggested an exponent of −2, but this has not been borne out by more extensive data shown in Fig. 12.

Plots such as Fig. 12 can be used as a guide to the likely numbers of faults at the intermediate size range, by using the estimated power-law trend (see Needham *et al.* 1996). It is possible, of course, that the mapped and core-scale faults are not related, and that extrapolation

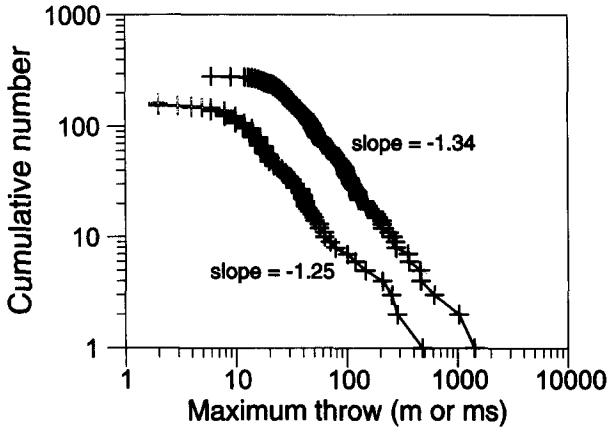

Fig. 13. Examples of maximum-displacement samples for fault traces mapped on seismic data. Both data sets (A1 and A2 in Table 1) show a good power-law form above the limit of seismic resolution (*ca* 10 and 25 m respectively).

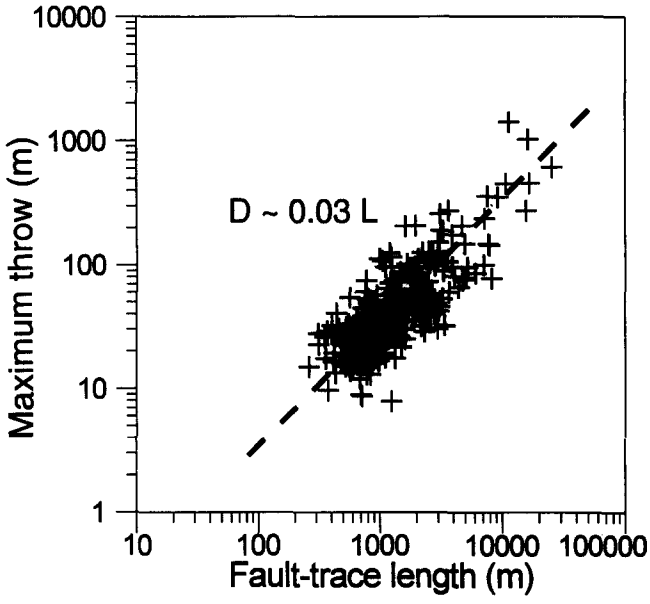

Fig. 14. Length–displacement cross-plot for a set of fault traces (data set A1 of Table 1, shown in Figs. 8 and 11). Note the significant variation in displacement at any given trace length. A linear relationship is shown ($D \sim 0.03L$) but the power law $D \sim L^{1.10}$ fits the data equally well.

from the mapped sample to the core sample is fortuitous. However, examples of detailed mapping at the intermediate scale in other areas (e.g. Scholz *et al.* 1993, Watterson *et al.* 1996) also show power-law behaviour. The simplest interpretation of these observations is that power-law behaviour is applicable across the range of scales.

Fault-trace maximum displacements

Examples of maximum vertical displacements measured on fault traces in North Sea fields are shown in Fig. 13. (In terms of sampling domain, these are two-dimensional samples of the fault displacement.) The maximum displacement is a relatively robust attribute and does not suffer from resolution problems in the manner discussed above for lengths. The samples shown in Fig. 13 are typical in showing good power-law form. The abrupt changes of slope at 10 and 25 m indicate the limit of resolution of the respective seismic data sets: below these values very few faults are detected.

Constraint from well data is not possible as maximum displacements cannot be determined from cores. A displacement of 10 cm seen on a fracture in a well core may be near the centre of a small fault or near the tip of a much larger structure.

The exponents shown in Fig. 13 are typical of seismic reflection data sets we have analysed (see Table 1 and Fig. 10b). The range over 11 data sets is from -0.96 to -1.48. This is a tighter distribution than that for fault-trace lengths (Fig. 10a), possibly because the latter have greater errors caused by the resolution problems discussed above.

Published exponent values from field and seismic mapping show a similar range of -1.2 to -1.5. Scholz & Cowie (1990) give a cumulative frequency exponent of 1.2 for maximum displacements on small Neogene faults in Japan. Marrett & Allmendinger (1992) infer values between 1.3 and 1.5 for the same Japanese faults and for fault traces mapped on Gulf of Mexico seismic reflection data. Villemin *et al.* (1995) report an average exponent of 1.35 for maximum offsets on fault traces in the Lorraine Coal Basin.

Length–displacement relationship

The relationship between fault length and fault displacement has been a matter of some controversy in recent years, in particular as to whether the relationship is linear ($D = cL$) or power-law ($D = cL^n$) (e.g. Cowie & Scholz 1992c, Gillespie *et al.* 1992, Hatton *et al.* 1994). Individual data sets usually exhibit considerable scatter and have a restricted range. As an example, in Fig. 14 we show a displacement–length cross-plot from a large data set (cumulative frequencies of lengths and maximum displacements are shown in Figs. 9 and 13, respectively). At any given length, there is almost an order-of-magnitude range of maximum displacement, and vice-versa. The fitted relationship shown on the figure is linear, with a constant of proportionality of 0.03; however, a power law with exponent 1.10 fits the data equally well. Extrapolation of these relationships to smaller sizes is only justified over perhaps 1 order of magnitude: beyond that the uncertainties would be unacceptably large.

An alternative approach to the length–displacement relationship is to compare the cumulative frequency exponents for each attribute. If length and maximum displacement are linearly related, the cumulative frequency exponents should be the same. If they are power-law related, the cumulative frequency exponents should be consistently unequal, and their ratio is a guide to the value of n in $D = cL^n$. Table 1 lists this ratio for the data sets discussed already, and the trace-length and displacement cumulative frequency exponents are cross-plotted in Fig. 10(c). For most data sets the ratio lies

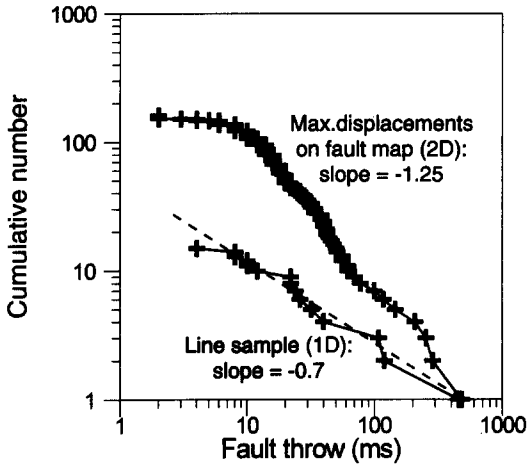

Fig. 15. Comparison of two-dimensional (map) and one-dimensional (line) samples of displacement from the same mapped horizon in a North Sea field (data set D3 in Table 1, from Yielding *et al.* 1992). Note that the map sample shows a steeper slope, as it should for the higher dimension of the sampling domain; however the difference is not 1 as would be expected if maximum displacement scaled with fault length (cf. Fig. 6).

between 1.0 and 1.4; for none of them is it less than 1.0, and for two it is about 2.0. This observation suggests that a power-law length–displacement relationship (i.e. $n >$ 1.0) might be more appropriate than a linear relationship. However, such a relationship might not be 'universal' over a wide range of scales (see discussion by Hatton *et al.* 1994).

It is perhaps noteworthy that the two samples (F1, F2) with highest trace-length exponent also have low displacement exponents, and therefore plot separately from the other data sets on Fig. 10(c). It is not clear whether this is a result of the multiple extension episodes suffered by this area.

Line (one-dimensional) sampling of displacements

A line, or one-dimensional, sample can be obtained by making measurements along a traverse drawn on an horizon map or along a single horizon on a seismic section. In the latter case the measurements are usually restricted to displacements, since the length of the fault is not discernible from the cross-section alone. Displacements sampled in this way are relatively objective since they do not require the fault pattern to be interpreted in three dimensions.

As shown with the fractal model, line samples of a set of faults must exhibit a lower fractal dimension than a map (two-dimensional) sample of the same faults. Figure 15 shows line and map samples of displacements on faults cutting a Jurassic horizon in a North Sea field mapped with a detailed seismic survey. The sample of maximum displacements on the fault traces on the horizon map displays a power-law form, with slope −1.25. The line sample across the same horizon map gives a much lower slope, −0.7; the line sample obviously crosses only a small subset of the total number of mapped faults, and will generally not encounter the maximum displacement on any fault trace that it crosses. The slope estimate is quite stable (±0.1) across the several hundred seismic lines that comprise the survey.

It is noteworthy that the one- and two-dimensional samples shown in Fig. 15 have slopes that differ by an amount that is not exactly 1.0, in contrast to the fault lengths examined in the fractal model (Fig. 6). The same is true of the ranges of published exponents for many different data sets: −1.0 to −1.5 for maximum displacements on traces (discussed above) and −0.5 to −1.0 for displacements on line samples (Childs *et al.* 1990, Yielding *et al.* 1992, Walsh *et al.* 1994). If maximum displacement were linearly proportional to fault length, then changing the sampling domain ought to change the slope by 1. A difference of slopes less than 1 implies $D \sim L^n$ with $n > 1$, in agreement with the analysis above of two-dimensional samples.

Line sampling of displacements is particularly useful in that measurements at the seismic scale and at the well-core scale are relatively straightforward to compare. Average linear fault density (and population slope) can be obtained by combining many parallel seismic lines into one sample. Similarly, if fine-scale bedding is present, linear fault density and population slope can be readily measured from slabbed core. The seismic line sample is effectively horizontal and the well sample is (usually) vertical, and so a simple trigonometric correction (relative to fault dip) must be applied before they can be compared. More detailed examples of line-sampling are given by Marrett & Allmendinger (1992), Walsh *et al.* (1992, 1994) and Needham *et al.* (1996).

DISCUSSION

The fault-population samples described above, and all others published in the literature, represent fault cuts seen on a slice or line through the faulted rock volume. Modern three-dimensional seismic reflection data offers the possibility of obtaining a true three-dimensional sample of a fault population, but to date no comprehensive analysis of this type has been attempted. In the meantime, inferences about the true population must be made from the lower dimensional samples. It is to be hoped that true three-dimensional samples will be available in the future.

The samples of fault-trace lengths on maps exhibit power-law exponents between −1.1 and −2.0, with the better-constrained data sets lying in the central part of this range. If these data represent isotropic fractal populations, then the exponents for the population of faults in the volume would lie in the range −2.1 to −3.0. Faults that are large relative to the thickness of the seismogenic layer (*ca* 10 km) would not meet this criterion, since these structures are essentially cuts in a brittle sheet (Scholz 1990). Smaller faults, however, will in many cases be a true population of planes in a volume and therefore the above argument should be valid. The faults plotted in Fig. 7 are 'small' faults except for the few longest traces; all the data sets represent areas of Mesozoic crustal extension. Whilst we cannot define the

population distribution for the largest faults (which define the main tilted fault blocks), we can infer that the length populations for the 'small' faults appear to be in the range −2.1 to −3.0. A fractal dimension between 2 and 3 seems intuitively reasonable in that faulting is a process that deforms a rock volume (i.e. three dimensions) by a set of planes (two-dimensional features), cf. King (1983).

A power-law control on earthquake frequencies has been recognized for many decades (e.g. Gutenberg & Richter 1954). Cumulative frequency plots of earthquake magnitudes are typically log-linear with a slope (b-value) of *ca* −1 for 'small' earthquakes and *ca* −1.5 for 'large' earthquakes (Pacheco *et al.* 1992). Magnitude is a logarithmic measure of the size of an earthquake, and can be related to the logarithm of the seismic moment, which is the product of seismic slip, slip surface area and rigidity modulus. Aki (1981) showed that for 'small' earthquakes the seismic moment scales with the cube of the length (L^3), and therefore a b-value of −1 implies a power-law exponent of −2 for the length distribution of the slip surfaces. For 'large' earthquakes the slip surfaces cut the seismogenic layer and seismic moment is proportional to L^2; the b-value of −1.5 then also corresponds to a fault-length exponent of −2. Thus active slip surfaces show a self-similar length distribution over all scales: the cumulative frequency exponent is −2. This contrasts with the fault-length distributions discussed above, where the exponent is −2 to −3 (typically *ca* −2.5).

Why should active slip surfaces and inactive faults have different frequency distributions? There are two possible explanations. First, the recurrence interval for earthquake occurrence on a fault may be a function of fault size (Turcotte 1992). If the recurrence interval is longer for smaller faults, then small slip surfaces will be proportionally less well represented in the seismicity. Fractal dimensions of 2.5 for the fault population and 2.0 for the seismicity, for example, would imply that recurrence interval is proportional to $1/L^{0.5}$. This kind of behaviour is clearly *not* observed in plate boundary areas (e.g. San Andreas) where the strain rate is relatively uniform and small segments slip more frequently than large ones (Scholz 1990). Intra-plate areas such as the Wasatch Fault Zone do show longer recurrence intervals on shorter distal segments (Cowie & Scholz 1992b, using data from Machette *et al.* 1991). However this is probably just a reflection of strain-rate variations along the strike of the structure, and is not strong evidence that small faults in a deforming region have consistently longer recurrence intervals.

The second explanation is that faults do not remain active throughout the deformation episode. The earthquake population therefore represents the instantaneous deformation on the fault network, which evolves over time. It is to be expected that the instantaneous deformation would become progressively localized on through-going faults and many small faults become inactive (King 1983, Westaway 1992, Cowie *et al.* 1993); the 'geological' fault population however must obviously include all faults that have ever been active. This behaviour can be replicated by models of fault-population growth (Walsh & Watterson 1992). Analogous behaviour is observed in diffusion-limited aggregation of 'viscous finger' structures in fluids, where the active growth zone has a lower fractal dimension than the entire structure (Feder 1988, pp. 96–99). A valuable way to test such growth models for faulting would be to examine fault populations in areas of active faulting, so that 'seismic' and 'geological' samples could be compared directly.

CONCLUSIONS

(i) As stated in previous studies, seismic reflection data and well cores are capable of providing a wealth of information about fault networks in the subsurface. High-quality seismic reflection data is volumetrically continuous and has a resolution of 20–30 m. Well cores can provide samples of fault displacements in the range 1–100 mm. Core data place valuable constraints on the extrapolation of populations from the scale of seismic mapping.

(ii) 'Fault' samples derived from an isotropic fractal fragmentation model can be used to illustrate many of the properties of real fault systems. They show that:

—the exponent of cumulative frequency distributions is affected by the dimension of the sampling domain (the exponent for fault lengths changes by 1 when moving from volume to plane or from plane to line);

—censoring limits the maximum size of fault that can be sampled, but faults that extend beyond the sample area should still be included;

—it is difficult to compare fault samples at different scales because those areas mapped in more detail are usually the most highly faulted.

(iii) Fault-trace length distributions derived from seismic reflection and other data exhibit a power-law exponent between −1.1 and −2.0, with the better constrained data sets lying at the centre of this range. Seismic resolution effects degrade the sample but can be corrected for if displacement gradients are used to extrapolate mapped fault tips to realistic lengths.

(iv) Maximum displacements on fault traces usually provide a robust sample, controlled by a power-law exponent of about −1.2 to −1.5.

(v) Conclusive data on the relationship between trace-length and maximum displacement are difficult to collect, but the available analyses suggest that the relationship is not linear in the size range covered by seismic reflection data.

(vi) Displacements on line samples have power-law exponents in the range −0.5 to −1.0.

(vii) The observed exponents for trace-length distributions (i.e. typically *ca* −1.5) imply that 'small' faults in the rock volume have a population exponent of *ca* −2.5. This is significantly different from the length distribution of active slip surfaces in earthquakes (exponent −2.0), which suggests that the instantaneous deformation on a

fault population does not affect all the faults. The geologically-observed fault population is the time-integrated result of many episodes of fault slip and fault growth.

Acknowledgements—We are grateful to BP, Norsk Hydro, Schlumberger and Statoil for permission to use the data and interpretations described in this paper. Ian Davison and Rob Knipe are thanked for constructive reviews of a first draft.

REFERENCES

Aki, K. 1981. A probabilistic synthesis of precursory phenomena. In: *Earthquake Prediction: An International Review* (edited by Simpson, D. W. & Richards, P. G.). *Am. Geophys. Un. Geophys. Monogr., Maurice Ewing Series* **4**, 566–574.

Barton, C. C. 1995. Fractal analysis of sealing and spatial clustering in fractures. In: *Fractals in the Earth Sciences* (edited by Barton, C. C. & La Pointe, P. R.). Plenum Press, New York, 141–178.

Childs, C., Walsh, J. J. & Watterson, J. 1990. A method for estimation of the density of fault displacements below the limits of seismic resolution in reservoir formations. In: *North Sea Oil and Gas Reservoirs—II* (edited by Buller, A. T., Berg, E., Hjelmeland O., Kleppe, J., Torsæter, O. & Aasen, J. O.), Graham & Trotman, London, 309–318.

Cowie, P. A. & Scholz, C. H. 1992a. Physical explanation for the displacement-length relationship of faults using a post-yield fracture mechanics model. *J. Struct. Geol.* **14**, 1133–1148.

Cowie, P. A. & Scholz, C. H. 1992b. Growth of faults by accumulation of seismic slip. *J. geophys. Res.* **97**, 11,085–11,095.

Cowie, P. A. & Scholz, C. H. 1992c. Displacement–length scaling relationship for faults: data synthesis and discussion. *J. Struct. Geol.* **14**, 1149–1156.

Cowie, P. A., Vanneste, C. & Sornette, D. 1993. Statistical physics model for the spatiotemporal evolution of faults. *J. geophys. Res.* **98**, 21,809–21,821.

Dalley, R. M., Gevers, E. C. A., Stampfli, G. M., Davies, D. J., Gastaldi, C. N., Ruijtenberg, P. A. & Vermeer, G. J. O. 1989. Dip and azimuth displays for 3D seismic interpretation. *First Break* **7**, 86–95.

Feder, J. 1988. *Fractals*. Plenum Press, New York.

Gauthier, B. D. M. & Lake, S. D. 1993. Probabilistic modelling of faults below the limit of seismic resolution in Pelican field, North Sea, offshore United Kingdom. *Bull. Am. Ass. Petrol. Geol.* **77**, 761–777.

Gillespie, P. A., Walsh, J. J. & Watterson, J. 1992. Limitations of dimension and displacement data from single faults and the consequences for data analysis and interpretation. *J. Struct. Geol.* **14**, 1157–1172.

Gillespie, P. A., Howard, C. B., Walsh, J. J. & Watterson, J. 1993. Measurement and characterisation of spatial distributions of fractures. *Tectonophysics* **22**, 113–141.

Gutenberg, B. & Richter, C. F. 1954. *Seismicity of the Earth and Associated Phenomena*. Princeton University Press, Princeton, New Jersey.

Hatton, C. G., Main, I. G. & Meredith, P. G. 1994. Non-universal scaling of fracture length and opening displacement in the Krafla fissure swarm, NE Iceland. *Nature* **367**, 160–162.

Heffer, K. J. & Bevan, T. G. 1990. Scaling relationships in natural fractures—data, theory and applications. *Proc. European Petrol. Conf.* **2**, 367–376 (SPE paper No. 20981).

Jackson, P. & Sanderson, D. J. 1992. Scaling of fault displacements from the Badajoz-Cordoba shear zone, SW Spain. *Tectonophysics* **210**, 179–190.

King, G. 1983. The accommodation of large strains in the upper lithosphere of the Earth and other solids by self-similar fault systems: the geometrical origin of *b*-value. *Pure & Appl. Geophys.* **121**, 761–815.

Machette, M. N., Personius, S. F., Nelson, A. R., Schwartz, D. P. & Lund, W. R. 1991. The Wasatch fault zone, Utah—segmentation and history of Holocene earthquakes. *J. Struct. Geol.* **13**, 137–149.

Mandelbrot, B. B. 1982. *The Fractal Geometry of Nature*. W. H. Freeman, New York.

Marrett, R. & Allmendinger, R. W. 1991. Estimates of strain due to brittle faulting: sampling of fault populations. *J. Struct. Geol.* **13**, 735–738.

Marrett, R. & Allmendinger, R. W. 1992. Amount of extension on "small" faults: an example from the Viking graben, *Geology* **20**, 47–50.

Needham, D. T., Yielding, G. & Fox R. J. 1996. Fault population description and prediction using examples from the offshore U.K. *J. Struct. Geol.* **18**, 155–167.

Okubo, P. G. & Aki, K. 1987. Fractal geometry in the San Andreas fault system. *J. geophys. Res.* **92**, 345–355.

Pacheco, J. F., Scholz, C. H. & Sykes, L. R. 1992. Changes in frequency-size relationship from small to large earthquakes. *Nature* **355**, 71–73.

Sammis, C., King, G. & Biegel, R. 1987. The kinematics of gouge deformation. *Pure & Appl. Geophys.* **125**, 777–812.

Sassi, W., Livera, S. E. & Caline, B. P. R. 1992. Reservoir compartmentalisation by faults in Cormorant Block IV, U.K. northern North Sea. In: *Structural and Tectonic Modelling and its Application to Petroleum Geology* (edited by Larsen, R. M., Brekke, H., Larsen, B. T. & Talleraas, E.), *Spec. Publ. Norweg. Petrol. Soc.* Amsterdam. **1**, 355–364.

Scholz, C. H. 1990. *The Mechanics of Earthquakes and Faulting*. Cambridge University Press, Cambridge.

Scholz, C. H. & Cowie, P. A. 1990. Determination of total strain from faulting using slip measurements. *Nature* **346**, 837–839.

Scholz, C. H., Dawers, N. H., Yu, J.-Z. & Anders, M. H. 1993. Fault growth and fault scaling laws: preliminary results. *J. geophys. Res.* **98**, 21,951–21,961.

Turcotte, D. L. 1992. *Fractals and Chaos in Geology and Geophysics*. Cambridge University Press, Cambridge.

Villemin, T., Angelier, J. & Sunwoo, C. 1995. Fractal distribution of fault length and offsets: implications of brittle deformation evaluation—the Lorraine Coal Basin. In: *Fractals in the Earth Sciences* (edited by Barton, C. C. & La Pointe, P. R.). Plenum Press, New York, 205–226.

Walsh, J. J. & Watterson, J. 1987. Distributions of cumulative displacement and seismic slip on a single normal fault surface. *J. Struct. Geol.* **9**, 1039–1046.

Walsh, J. J. & Watterson, J. 1992. Populations of faults and fault displacements and their effects on estimates of fault-related regional extension. *J. Struct. Geol.* **14**, 701–712.

Walsh, J., Watterson, J. & Yielding, G. 1991. The importance of small-scale faulting in regional extension. *Nature* **351**, 391–393.

Walsh, J. J., Watterson, J. & Yielding, G. 1994. Determination and interpretation of fault size populations: procedures and problems. In: *North Sea Oil and Gas Reservoirs—III* (edited by Aasen, J. O., Berg, E., Buller, A. T., Hjelmeland, O., Holt, R. M., Kleppe, J. & Torsæter, O.). Kluwer, Dordrecht, 141–155.

Watterson, J., Walsh, J. J., Gillespie, P. A. & Easton, S. 1996. Scaling systematics of fault sizes on a large-scale range fault map. *J. Struct. Geol.* **18**, 199–214.

Westaway, R. 1992. Evidence for anomalous earthquake size distributions in regions of minimal strain. *Geophys. Res. Lett.* **19**, 1499–1502.

Yielding, G., Walsh, J. & Watterson, J. 1992. The prediction of small-scale faulting in reservoirs. *First Break* **10**, 449–460.