Variability in fault size scaling due to rock strength heterogeneity: a finite element investigation

K.M. Hardacre[*], P.A. Cowie

Department of Geology and Geophysics, Edinburgh University, Edinburgh, UK

Abstract

Practical application of fault size scaling relationships usually involves extrapolation of a limited data set over a scale range beyond that observed or into adjacent unstudied areas. Here we investigate the validity of such extrapolations by considering the variability in the size–frequency distributions of fault populations that develop under identical tectonic conditions using a numerical model to generate conjugate, normal fault populations in cross-section. The deforming material is modelled using a strain-softening, Von Mises rheology with Gaussian heterogeneity in yield strength distributed randomly throughout the mesh. We present eight deformation experiments that differ only in the random spatial pattern of yield strengths. We observe power law size–frequency scaling, i.e. $N = ax^{-c}$ (where N is the cumulative number of faults and x is a measure of fault size) but with a range of values of c. The ensemble average value of c decreases with increasing percentage extension. However, for individual model runs the dependence of c on total strain shows significant variability that we can relate to small but important differences in fault growth and strain localisation. At any particular strain, the range of values of c is ~ 10 times greater than the error estimate derived from least squares regression of the cumulative frequency data. Our results suggest therefore that large uncertainties should be associated with extrapolating fault population data from one scale or region to another even if the lithology and tectonic history are similar.
© 2003 Elsevier Science Ltd. All rights reserved.

Keywords: Fault size scaling; Rock strength heterogeneity; Finite element

1. Introduction

Fault scaling relationships are used to describe the relative number, the spatial distribution and the displacement–length ratio of faults formed within the same tectonic episode. There are several practical uses of these relationships, for example calculating the total strain due to faulting (Scholz and Cowie, 1990; Marrett and Allmendinger, 1991, 1992; Walsh et al., 1991; Westaway, 1994), modelling sea floor topography (Malinverno and Cowie, 1993), and assessing the impact of faults on hydrocarbon reservoir quality (Yielding et al., 1992; Gauthier and Lake, 1993). These applications often involve extrapolating the relationships observed at one scale to another scale or region to predict the full fault population (Yielding et al., 1996). How can we assign confidence limits to such extrapolations? This

is a particularly pertinent question because fault size–frequency distributions commonly exhibit power law scaling, i.e. $N = ax^{-c}$ (where N is the cumulative number of faults and x is a measure of fault size) (see for reviews Cowie (1998) and Bonnet et al. (2001)). Thus a small change in the value of the scaling exponent, c, can lead to an order of magnitude difference in the number of faults that is predicted, especially when small-scale fault populations are predicted from larger-scale observations.

For a single data set the only available estimate of the uncertainty is the error on the best-fitting power law function, usually obtained by least squares regression of a log–log plot of cumulative number versus fault length or fault displacement. This is a statistical measure of the uncertainty that does not embody any knowledge of how the fault population developed. In this paper we present another approach to evaluating uncertainty that involves investigating the reproducibility of results in a numerical model of fault growth. Numerical models can provide insights into the natural variability in the exponent c because multiple

* Corresponding author. Correspondence address: Shell UK Exploration and Production, Aberdeen AB12 3FY, UK.

E-mail addresses: kathryn.hardacre@shell.com (K.M. Hardacre), cowie@mail.glg.ed.ac.uk (P.A. Cowie).

data sets can be generated with known initial and boundary conditions and the resulting variance in c quantified explicitly. Previous work using analogue and numerical models to simulate faulting, has provided considerable understanding of the physical controls on the value of c. However, these studies have reported values of c obtained from (1) a suite of experiments in which some fundamental control on fault population development was changed (e.g. the thickness of the deforming layer: Finch, 1998; Ackermann et al., 2001; Hardacre, 2000), or (2) ensemble averaging of results, which gives a feel for 'average' behaviour, but not the inherent variability in the system (Cox and Paterson, 1990; Spyropoulos et al., 1999). In this paper we present a detailed analysis of size–frequency distributions that arise amongst a suite of numerical experiments that have statistically identical starting points. The models only differ in the random distribution of the strength heterogeneity. However, this is sufficient to result in a different sequence of interactions between adjacent faults as they grow and thus different patterns of strain localisation. The detailed history of fault development determines the size–frequency distribution, and thus the value of c, that emerges. We compare the overall range of observed c values with the uncertainties on c estimated for individual data sets and discuss the implications for extrapolating fault populations in practical applications.

2. Modelling approach

Many published studies of the controls on fault scaling relationships are based on analogue models of faulting (e.g. Sornette et al., 1993; Davy et al., 1995; Finch, 1998; Spyropoulos et al., 1999; Ackermann et al., 2001; Mansfield and Cartwright, 2001). However, in most analogue models the fault pattern is viewed from above so that the population statistics only reflect the behaviour of faults whose tip lines cut the upper surface. It is also difficult to measure small faults during these experiments. For example, Clifton et al. (2000) note that, in their wet plaster models, the heaves on the faults that they could resolve accounted for as little as 30% of the imposed regional extension. The early numerical models used to investigate fault scaling were either limited to very small percentage strains (Cowie et al., 1993) or had periodic boundary conditions and simple homogeneous material properties (Poliakov et al., 1994). Though these problems can be addressed, computational limits on the size and number of elements mean that numerical models inevitably suffer from resolution problems as well. Despite such limitations, a numerical modelling technique was used in this study because it permits precise control of the initial and boundary conditions and the resolution limit is easily measured.

Numerical models of faulting form a spectrum between two basic approaches. At one end of the spectrum a large number of physical influences on faulting are included in a simulation but only simple geometries are considered. For example, the boundary element models of Martel (1997, 1999) give accurate solutions to the equations of linear elastic fracture mechanics but usually only a few pre-defined faults or fractures are modelled. At the opposite end of the spectrum are models that capture only the bare essentials of the physics of fracturing, but simulate large numbers of interacting faults and permit structures to growth and link over time. The cellular automaton model of Wilson et al. (1996) and the modified resistor network model of Cowie et al. (1993) represent this end member. Between these extremes of approach lie finite difference (Maillot et al., 1998; Spyropoulos, 1999; Spyropoulos et al., 2002), discrete element (Homberg et al., 1997; Finch, 1998) and finite element models (Dresden et al., 1991; Poliakov and Hermann, 1994; Schultz-Ela and Jackson, 1996; Gerbault et al., 1998; Huc et al., 1998; McKinnon and Garrido de la Barra, 1998).

The 2D finite element code used here, ADELI (described in full by Hassani (1994)), lies towards the 'complexity of space, simplicity of physics' modelling end-member. Unlike some other schemes, boundary conditions are naturally incorporated in finite element models, rheologies are defined explicitly and large strains can be simulated. ADELI uses a Lagrangian description of the geological medium and large strains are modelled under quasistatic conditions. Faulting is simulated via the mechanism of strain localisation (shear band formation) in a continuum following a strain softening, elasto-plastic constitutive law. Prior to fault initiation, the deforming material behaves as a linear elastic solid obeying Hooke's Law. The point of failure is defined by the Von Mises yield criterion, which provides a relationship between yield strength (σ_0) and deviatoric stress (mean stress is not taken into consideration in this criterion). After failure, the material loses a fixed amount of strength (σ_{drop}) whilst accumulating a certain amount of strain (the critical plastic strain, κ_c) before entering a plastic flow regime. The Von Mises constitutive model is not ideal in that it predicts failure at 45° to the principle stresses and does not incorporate yield strength increase at higher mean stress levels. However, incorporating heterogeneity in material parameters, which is the necessary for our study (see below), is easier within a simple constitutive framework. The constitutive law does not include healing: once an element has failed it remains weak throughout the experiment. Alternative yield criteria (Mohr–Coulomb, Van Eekelen and Drucker–Prager) were considered but were rejected for the purposes of this study for reasons discussed in Hardacre (2000) and Hardacre and Cowie (2003).

Elasto-plastic deformation of solids shows many phenomenological similarities with the deformation of rocks seen in nature and in rock fracturing experiments (Poliakov et al., 1994; McKinnon and Garrido de la Barra, 1998). Firstly, deformation is permanent: when the deforming stress is released, the material does not recover its

undeformed state. Secondly, the stresses required to give observed strains are reasonable tectonically. Thirdly, deformation of rocks in the upper crust does seem to be a yield point phenomenon. Finally, plasticity theory predicts the localisation of strain into narrow zones akin to faults (Rudnicki and Rice, 1975; Rudnicki, 1977).

In addition to overall rheology, heterogeneity has been shown to be a key element in the development of fault networks (Cowie et al., 1993). Variations in composition, grain size and texture constitute heterogeneities in real rocks. In analogue models, heterogeneity enters the system via microscopic variations in the sand grains, clay or gypsum used to represent the brittle crust. In some numerical models (Poliakov and Hermann, 1994; Poliakov et al., 1994; Barnichon and Charlier, 1996; Schultz-Ela and Jackson, 1996; Gerbault et al., 1998), heterogeneity emerges because of numerical 'noise' i.e. round-off errors in floating point arithmetic. This source of heterogeneity was rejected in our approach because the amplitude of the noise can neither be measured nor controlled: it is entirely dependent on machine architecture. In other models (Leroy and Ortiz, 1989; Belytschko et al., 1994), a single imperfection is introduced as a seed in an otherwise uniform mesh. Non-uniform failure can also be induced by creating geometrical instabilities in the mesh e.g. by fixing one node to create a displacement discontinuity (Dresden et al., 1991) or by cutting a notch in the mesh perimeter (Davis and Fletcher, 1990). However, heterogeneity is most commonly incorporated in numerical models by giving each mesh element a strength value drawn from a probability distribution (e.g. Cowie et al., 1993; Heimpel and Olson, 1996; Wilson et al., 1996; McKinnon and Garrido de la Barra, 1998; Spyropoulos, 1999). In this study, the yield strength was chosen as the heterogeneous material property. The strength of each model element is drawn from a Gaussian probability distribution and the spatial arrangement is random.

2.1. Experimental design

The standard model measures 10 km across and 5 km deep. The selected dimensions are large enough that the forces of gravity and applied extension are comparable, but small enough that complexities like thermally-induced rheological variation with depth need not be considered. The material is divided into quadrilateral elements composed of two overlapping pairs of triangular elements. Each of the 5000 elements in the standard model represents $100 \text{ m} \times 100 \text{ m}$. The governing equations of stress equilibrium are solved explicitly using the dynamic relaxation technique (Board, 1989), as implemented in the code FLAC (Cundall, 1989). This method damps out-of-balance forces to arrive at the quasi-static equilibrium solution. A typical run required 20,000 time steps to ensure the stability of the solution whilst guaranteeing quasi-static conditions.

The boundary conditions on the sides of the model are given in terms of velocities normal and tangential to the surface (Fig. 1). One wall remains fixed whilst the other moves with a constant normal velocity of $1.0 \times 10^{-10} \text{ m s}^{-1}$, equivalent to an extensional strain rate of 10^{-14} s^{-1}. We do not vary the strain rate in this study. Using a numerical model similar to that presented here, Lavier et al. (2000) have shown that fault development may be sensitive to the applied strain rate. The standard run time is $4.0 \times 10^{12} \text{ s}$ (approximately 127 thousand years), giving a total extensional strain of 4%. With strains in excess of 4%, grid cells near faults become so deformed that the stress equilibrium equations cannot be solved accurately without regridding, which is undesirable. The upper boundary is free in all experiments and the model lies on an inviscid fluid, resulting in an isostatic basal boundary condition. The sidewalls are not allowed to rotate, so each model can be thought of as one half of a symmetrical extensional basin with the moving boundary at the mid-point of the basin. A non-rigid, rotational boundary would apply to a model of faulting driven by gravity collapse, which is not what we consider in this paper.

The deforming material has its elastic parameters (Poisson's ratio 0.3, Young's modulus $1.36 \times 10^{10} \text{ Pa}$), mean yield strength ($275 \times 10^6 \text{ Pa}$) and post-yield, plastic parameters (strength loss, $75 \times 10^6 \text{ Pa}$; critical plastic strain 0.02) chosen to mimic those of a sample of Darley Dale sandstone deformed in uniaxial compression at a confining pressure of 50 MPa (parameter fitting by Niño (1997)). Four
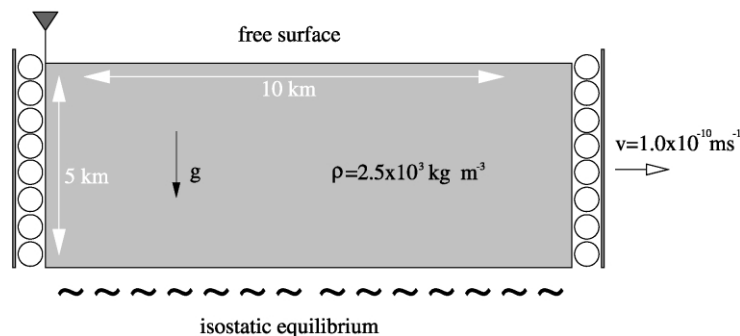

Fig. 1. The geometry and boundary conditions of the standard model: 5000 quadrilateral elements comprise the mesh, which represents a Von Mises plastic material. The left-hand wall is pinned, the right wall moves at a constant velocity. See Section 2.1 for details.

heterogeneous meshes were created by using different seeds (4067, 6007, 7177 and 7817) in the random number generator used to sample the Gaussian distribution of yield strengths. Each of these meshes (labelled n in figures) was reflected about its centre along a vertical axis to generate a further four, mirror-image meshes (labelled m in figures). As the applied boundary conditions are asymmetric, the deformation history of a mesh and its mirror image are not identical (see below). In each mesh, the mean strength was 275 MPa and the standard deviation 25 MPa. Meshes with other standard deviations were also investigated but those results are presented elsewhere (Hardacre, 2000; Hardacre and Cowie (2003)).

2.2. Map patterns

Fig. 2 illustrates the pattern of broken elements within each simulation. Fault development is *qualitatively* similar in all of the runs. Firstly, nucleation begins at the base of the model and close to the moving wall. This is because we use a Von Mises yield criterion that depends on deviatoric stress alone: the deviatoric stress increases towards the base because of the overburden and also increases towards the moving wall because of the applied load. A nucleation front then propagates upward and right-to-left across the model, away from the moving wall. The elements with the lowest strengths break first. As an element breaks it loses some of its strength and loads those elements around it. Consequently, local variations in yield strength determine whether a single broken element will develop as a left-dipping or right-dipping fault. Both synthetic (i.e. dipping towards the moving wall) and antithetic structures develop, as predicted by the modelling theory of Anderson (1951). The dip of the model faults (45°) is low compared with Andersonian faults (typically ~60°) as a consequence of modelling the deforming material as a Von Mises solid.

In finite element simulations, shear zones may 'reflect' at the boundaries of the deforming region (Davis and Fletcher,
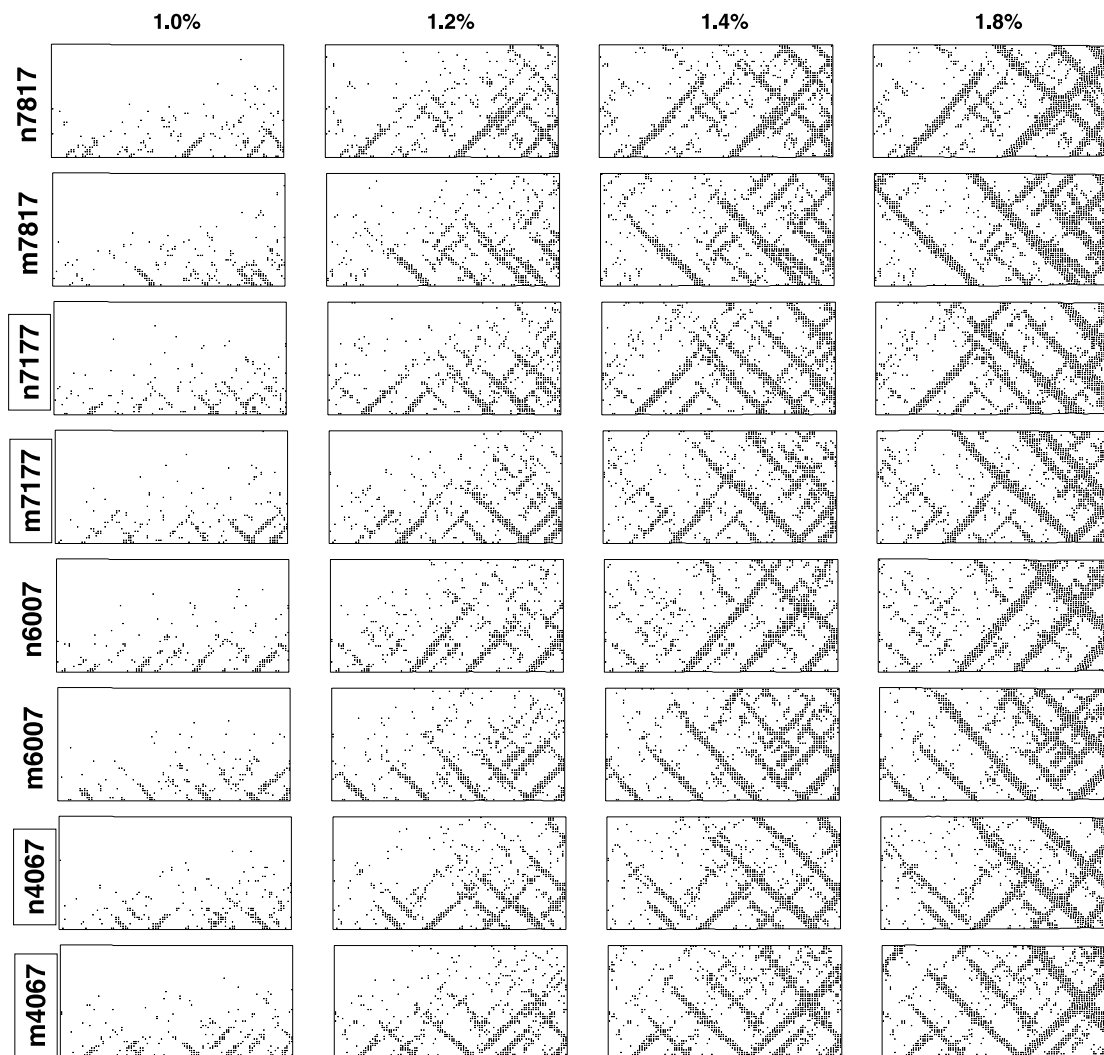

Fig. 2. Broken elements maps at 1.0, 1.2, 1.4 and 1.8% strain for model runs in which different seeds (4067, 6007, 7177, 7817) are used in the random number generator responsible for strength heterogeneity. The prefixes n and m denote, respectively, the normal and mirror image pair of the same mesh.

1990; Hobbs et al., 1990; McKinnon and Garrido de la Barra, 1998). This leads to V-shaped fault intersections and simple grabens. In such fault pairs, one structure is responsible for the formation of the other. However, they can be considered to be independent, as the growth of one structure does not necessitate the growth of its partner. Structures that intersect with the upper and lower surfaces of the model continue to accumulate displacement and the resultant thinning is compensated isostatically. Faults grow upwards by breaking the elements ahead of their tips or by coalescing with pre-existing structures. Down-dip linkage of synthetic faults of comparable size is rare but as structures lengthen and thicken during growth they commonly assimilate other smaller structures.

The largest faults have similar thicknesses and are spaced at preferred distances through the mesh. Preferred spacing of shear bands in finite element simulations of deformation in homogeneous materials is governed by the ratio of element height to model height, as described by Ord (1990). Despite the fact that the materials in the eight models used here are statistically identical, the precise location of major through-going structures varies because the pattern of yield strength heterogeneity is different in each model. Even the deformation history of a mesh and its mirror image are not identical, because the applied boundary conditions are asymmetric. In other words, the imposed yield strength heterogeneity largely suppresses the preferred spacing noted by Ord (1990) for homogeneous models. With time, strain localises into the area closest to the moving wall because structures here have been active longer and so are weaker. The eventual formation of a pair of conjugate faults forming an 'x' pattern adjacent to the moving wall is a consequence of making this a rigid, non-rotational, boundary and allowing local isostatic compensation (see Section 2.1 above).

2.3. Processing model output

The raw output of the model is data concerning the elements and nodes of the finite element mesh, not faults. Previous studies (Cowie et al., 1993; Poliakov and Hermann, 1994) used a two step algorithm for extracting fault data: first identifying failed elements and then clustering failed elements connected either at a node or along an edge to make faults. Such an algorithm would identify intersecting conjugate faults as a single structure. To identify synthetic and antithetic structures separately, we developed a three-step algorithm, illustrated by Fig. 3. First, each element is tested against a failure criterion to assess whether it is broken or unbroken (Fig. 3a). If an element has experienced finite plastic strain (measured in principal strain space as the length of the normal from the yield surface to the point representing the strain state of the element), it is considered to be broken. Then we examine the strain gradients around each element. If the UL–LR (upper left–lower right) gradients are smaller than the UR–LL gradients, the element is flagged as belonging to a right-
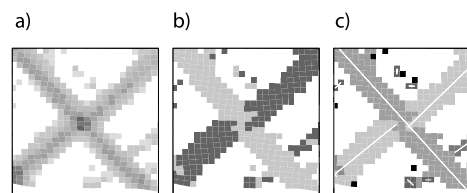

Fig. 3. The three-step clustering algorithm. (a) Elements with plastic strains above the threshold value of 0.05 (shades of grey) meet the failure criterion and are considered broken. Failed elements with low strains are shown in light grey, higher strains in dark grey. (b) Strain gradients are used to distinguish between elements with 'right-dipping' shear sense (light grey) and 'left-dipping' shear sense (dark grey). White elements are not broken. (c) Points belonging to the same cluster are shown in the same shade. Solid white lines are the best fitting straight line through each cluster.

dipping fault (Fig. 3b). If the opposite is the case, the element is assigned a left-dipping flag. Adjacent elements are then joined only if they are broken and share the same flag (Fig. 3c).

2.4. An appropriate measure of fault size

Size frequency analysis of a fault population should ideally meet the following criteria: (1) the range of size measurements should exceed one order of magnitude, and (2) the scale of measurements should be continuous to avoid artefacts caused by binning (see Pickering et al. (1995) and Main (2000) for a discussion of binning and finite range effects in cumulative size frequency statistics). Previous studies of faults in cross-section have measured their size either as displacement (throw) or length (termed fault 'height' by Finch (1998) and Ortega and Marrett (2000)). In finite element models, throw and length can only be measured between grid nodes, which makes throw measurements on structures containing just one or two elements of questionable significance and causes both finite range and binning problems for length data (discussed fully in Hardacre (2000)). The precedent set by other numerical modelling studies (Cowie et al., 1993; Poliakov and Hermann, 1994) is to use the number of elements in each structure as a measurement of size. Unlike length, the number of elements continues to increase after the structure spans the layer. However, because the number of elements in a structure will be an integer, this is not a continuous scale of measurement and the data are effectively binned.

In this study, we measure size as the sum of the plastic strains in the elements that make up a structure: a measure we term total plastic strain (TPS). Note that TPS is equivalent to the magnitude of the geometric moment of each fault, i.e. (displacement) × (fault surface area). Fig. 4a illustrates three different measures of fault size (number of elements, tip-to-tip length, TPS) applied to a single fault population. All three measures give distributions that are essentially linear in log–log space. The TPS distribution shows the same features as the length frequency distribution, for example the break in slope at cumulative frequency equals 10, but the 'number of elements'
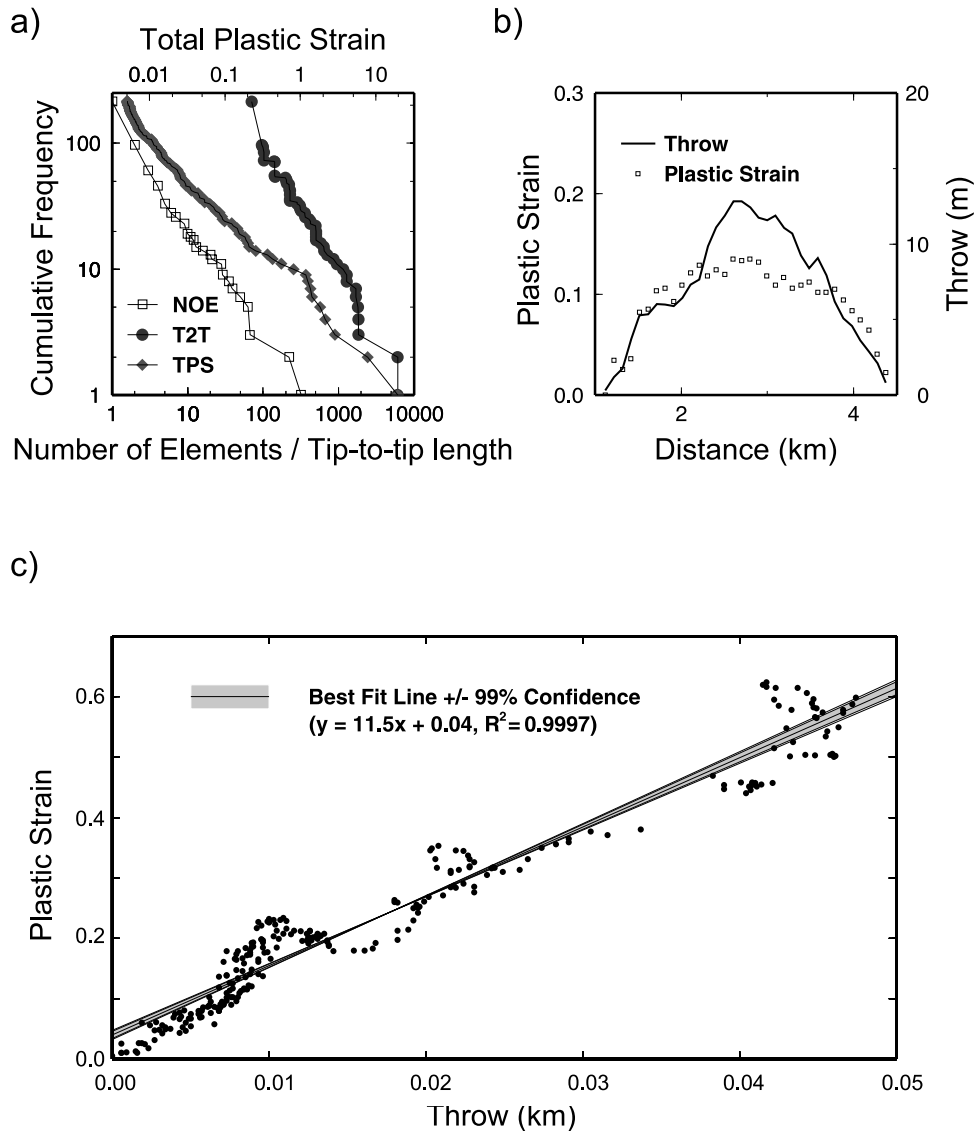
a)

b)

c)

Fig. 4. (a) Cumulative size frequency distributions calculated for a snapshot of a single model, using three different measures of fault size: number of elements (NOE), tip-to-tip length (T2T) and the total plastic strain (TPS). (b) Fault throw profile (bold line) and plastic strain profile (dotted line) across each row of elements for a single fault. (c) Relationship between plastic strain and throw measured across rows of elements for a number of faults.

distribution compares less favourably. TPS has a further advantage over the other measures of size: far more data-points define the total plastic strain distribution and so the bias towards large faults during power law fitting by linear regression is reduced. However, there is a disadvantage to using TPS as a measure of fault size. Examining Fig. 4a, the 'number of elements' and 'tip-to-tip length' distributions contain ∼250 faults, spanning three orders of magnitude. The TPS distribution contains the same number of faults but spans four orders of magnitude. This means that the 'number of elements' and 'tip-to-tip length' distributions will have similar power law exponents, but the exponent for the TPS distribution will necessarily be lower.

Having established that TPS is a good proxy for fault length, let us now examine the relationship between plastic strain and throw. For a single structure, throw profiles can be

generated by examining the vertical offset along each row of broken elements. Plastic strain can also be summed along each row, across the finite width of each fault zone, to give a plastic strain profile (Fig. 4b). The two profiles have very similar shapes: plastic strain generally varies along a fault in the same way as throw. This point is emphasised by Fig. 4c, which shows a clear linear relationship between throw and plastic strain calculated row-by-row on many faults.

2.5. Determining the exponent of the power law

In fault population studies it is most common to use the cumulative size frequency distribution to calculate the power law exponent c, i.e.

$$N = ax^{-c}$$

where x is fault size, N the number of faults of size $\geq x$, and a is a constant. An alternative method is to use the discrete frequency distribution of $\log(x)$ (hereafter referred to as the log-interval distribution). In this case the number of faults that have a size in the bin interval $\log(x) \pm \delta\log(x)$ is calculated and a histogram plotted. This latter approach is commonly used in the analysis of earthquake magnitude–frequency distributions. The relative merits of the two approaches are thoroughly described by Pickering et al. (1995), who also showed that the power law exponents calculated by either method are theoretically identical. However, in practice the exponent calculated from a log-interval distribution does depend on the choice of the bin size, $\delta\log(x)$. Consequently, the vast majority of geologists studying fault populations use the cumulative distribution to estimate the power law exponent. Cumulative frequency plots are not without their problems, in particular, the gradual roll-over in the distribution near the maximum value and a systematic increase in the apparent regression coefficient due to the smoothing effect inherent in using cumulative data (Main, 2000).

The exponent c is usually estimated by taking logarithms before fitting a straight line through the transformed data using least squares regression. This method assumes that the residuals of the dependent variable, $\log(N)$ in this case, are independent and random. However, by definition N must be an integer and on a cumulative frequency plot N is bound to increase as one goes from the largest measurement to the next largest and so on. Clarke et al. (1999) criticise least squares regression of cumulative frequency data for this very reason, proposing a maximum likelihood method instead.

However, in order to compare our results with the majority of published studies of fault size frequency distributions, we calculate c via least squares regression of cumulative frequency data. We analyse many data sets in the same way and compare the values of c that are obtained. By comparing values of c, rather than considering individual results, our conclusions are valid in spite of the criticisms raised by Clarke et al. (1999). The error on the estimate of c is calculated using Student's t-test. To avoid problems of sampling bias (i.e. censoring and truncation; see Pickering et al. (1995) for explanation), data from the extremes of the distribution are removed before fitting the power law. Although objective methods for doing this, based on maximising the goodness of fit statistic, r^2, have been suggested (Ortega and Marrett, 2000), the size distributions produced by ADELI have very obvious breaks in slope that define the limits to the regression (see Section 3.2).

The power law exponent derived from log-interval data is also calculated to show how estimates of c (and their associated errors) differ between the two methods. Although the log-interval distribution contains the same information as the cumulative distribution, Main (2000) argues that it reveals the true nature of the distribution more clearly. For

example, Main (2000) has shown that noise can cause apparent breaks in slope on the cumulative frequency distribution but not the log-interval distribution. It is also easier to justify the use of least squares regression with log-interval data as the number of faults can either increase or decrease as one goes from the largest fault size to the next largest. A log-interval width of 0.1 is used here and data were included in the regression from the smallest fault size up to the first empty bin. In the absence of an empty bin, data were included until two adjacent bins containing just one measurement had been encountered.

3. Results

3.1. Fault activity

Structures are considered active if any of the elements within them have accumulated plastic strain in the last time step. During the early nucleation phase, most faults are active and the number of faults increases (Fig. 5). Behind the nucleation front, the strain softening rheology stimulates the growth of larger, weaker faults. The localisation of strain onto some structures causes others to switch off, increasing the number of inactive structures through time (Fig. 5). Fault growth by coalescence, which is a common phenomenon by 1.2% extension in all models, leads to a decrease in the total number of faults (Fig. 5). After 1.6% model extension, faults span the layer from top to bottom and only the major structures are active. Consequently, there are fewer linkage events after this point and the number of structures decreases more slowly.

3.2. Cumulative size–frequency distributions

The cumulative size–frequency distributions and fitted power law exponents c, resulting from the eight experiments, are shown in Fig. 6. The error on the best fitting
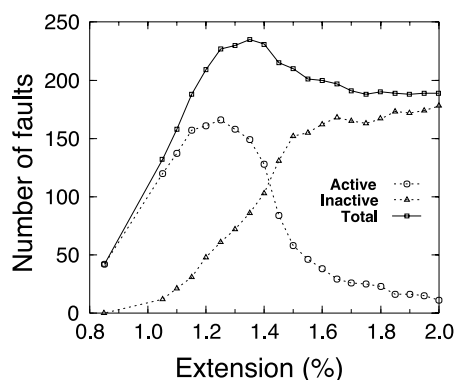

Fig. 5. Number of faults that are active and inactive in one particular experiment (mesh m4067) versus time (see Section 3.1). The pattern for the total population is consistent with that observed in analogue (Ackermann et al., 2001) and numerical (Cowie et al., 1995) studies of fault development.
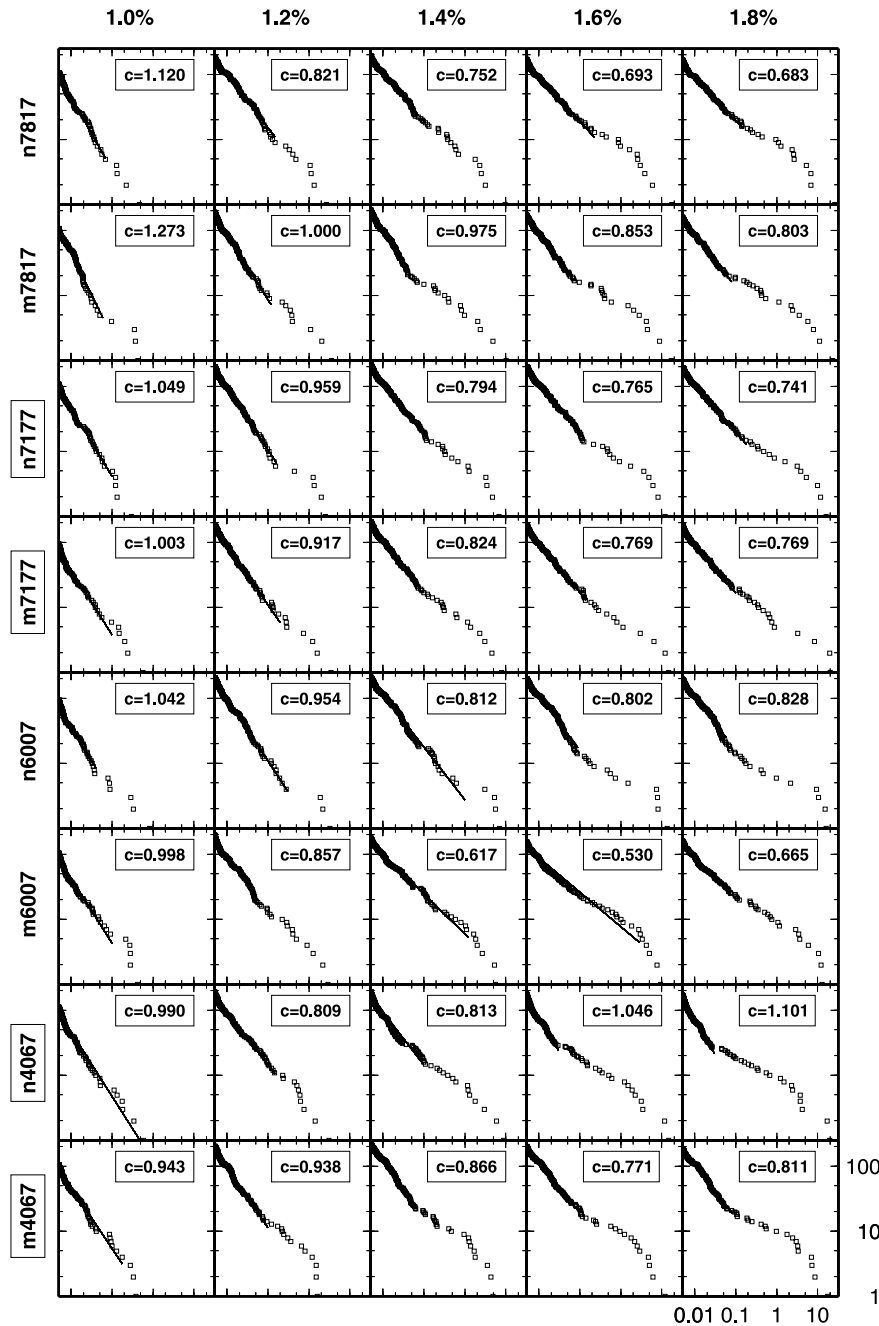
Fig. 6. Fault size (TPS) versus cumulative frequency plots at 1.0–1.8% strain for the eight experiments. The normal mesh and its mirror image for each seed are adjacent. Variation within the pairs are as large as variations between meshes with different seeds. All graphs have the same axes as the annotated plot at the bottom right of the figure. Best-fitting power laws, with exponents, c, are shown over the interval used in the least squares regression calculation. Estimated errors on c are typically less than 0.02.

value of c, estimated at the 95% confidence level, is less than 0.02 in all cases. Note that the values of c are low compared with many real fault data sets because we use TPS to represent fault size (Fig. 4). Some features are common to all runs. For example during the course of each run, the cumulative size–frequency distribution shows a change from a simple power law (i.e. a straight line on a bi-logarithmic plot) to a more complex form with one or more inflection points. The smaller faults in the population remain

power law while a breakdown in power law size scaling occurs for the largest faults. The reason for this breakdown in power law scaling is the localisation of strain onto the largest faults in the model. Hardacre (2000) shows that the large faults that penetrate the entire layer in this model are better fit by an exponential distribution (see also Ackermann et al., 2001; Spyropoulos et al., 2002; Walsh et al., 2003). In most of the experiments there is a gradual decrease in the value of c through time (Fig. 7). This is most clearly evident
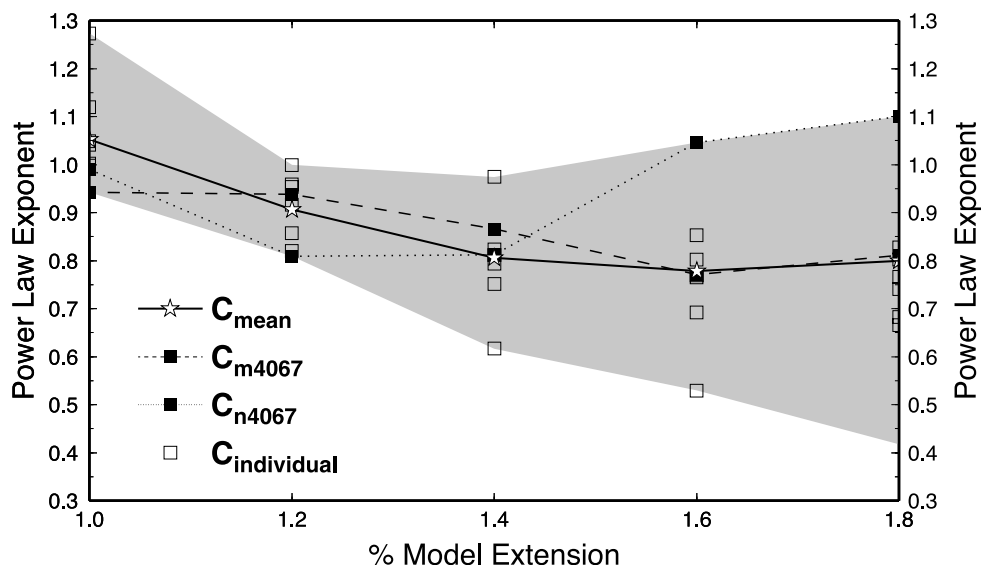
Fig. 7. Dependence of c (from Fig. 6) on total strain for the individual datasets. The mean value, c_{mean}, at each % strain is also shown. The curves for experiments n4067 and m4067 are highlighted (see text for discussion). Although c generally decreases with strain for individual runs, this signal is small compared with the variability induced by changing the seed used to generate the random heterogeneity.

for the early stages of fault development prior to the appearance of significant inflection points in the size frequency distributions, i.e. for most experiments at strains < 1.6% (see Fig. 6). In some cases however (see for example experiments m6007 and n4067) the c values actually increase with time, towards the end of the simulation.

Given the fact that all the experiments presented here are identical in terms of mechanical properties, except in the random seed used to generate the strength heterogeneity, it is reasonable to expect that the resulting fault populations should be very similar. However, in spite of the gross similarities mentioned above, the striking feature of Fig. 6 is the marked difference between experiments and in particular between mirror-image pairs (see also Fig. 2). The range in the value of c between the experiments at one value of strain is comparable with the range observed as a function of strain for any one experiment. The positions of inflection points also vary significantly. Note how the size frequency distributions from mirror images (for example n4067 and m4067) are as dissimilar as distributions from meshes produced by different random seeds. Consequently, by examining the distributions alone one could not identify those belonging to mirror pairs.

In other numerical modelling studies (Cox and Paterson, 1990; Spyropoulos et al., 1999) the fault populations from several runs are combined to investigate the 'average behaviour' of the system. To look at the average behaviour, we calculated the mean value of c for the eight experiments, c_{mean}, at each value of total strain. Fig. 7 shows c for the individual data sets and also c_{mean}, as a function of percentage extension. Note that the change in c_{mean} as total strain increases is less than the range of c values produced by varying the random distribution of the heterogeneity. The two curves

highlighted in Fig. 7 are for runs n4067 and m4067 in order to show the difference in behaviour that is possible even between mirror-image pairs. We investigate the reasons for this difference in Section 4 below.

3.3. Log-interval size frequency distributions

Pickering et al. (1995) and Main (2000) argue that log-interval distributions give a more robust indication of the true power law scaling properties of a fault population. Thus, for comparison, the log-interval distributions resulting from the eight experiments are shown in Fig. 8, along with new estimates of the power law exponent c. Plotting the log-interval data demonstrates that the central portion of the cumulative distributions, which appears to follow a power law with a very low slope (e.g. experiment n4067 at 1.8% strain), is exaggerated by plotting cumulative fault number. In reality, this portion of the cumulative distributions is simply a transition between the smaller-scale faults, which are demonstrably power law distributed, and the largest faults in the population that are better fit by an exponential distribution (Hardacre, 2000; Hardacre and Cowie (2003)).

The error estimates on c are much larger for the log-interval distributions (see Fig. 8, right hand panel); typically ± 0.5, compared with 0.02 for fits to the cumulative distribution. Although the estimates of c from the log-interval distributions are generally slightly higher than those from the cumulative distributions, they do agree within 95% confidence limits. However, there is now no statistically significant change in c with increasing strain, in contrast to that observed for the cumulative distributions (compare Fig. 8 with Figs. 6 and 7). Although the dependence of c on total strain seen in the cumulative distributions is statistically significant, this relationship is an subtle feature that can be easily masked by
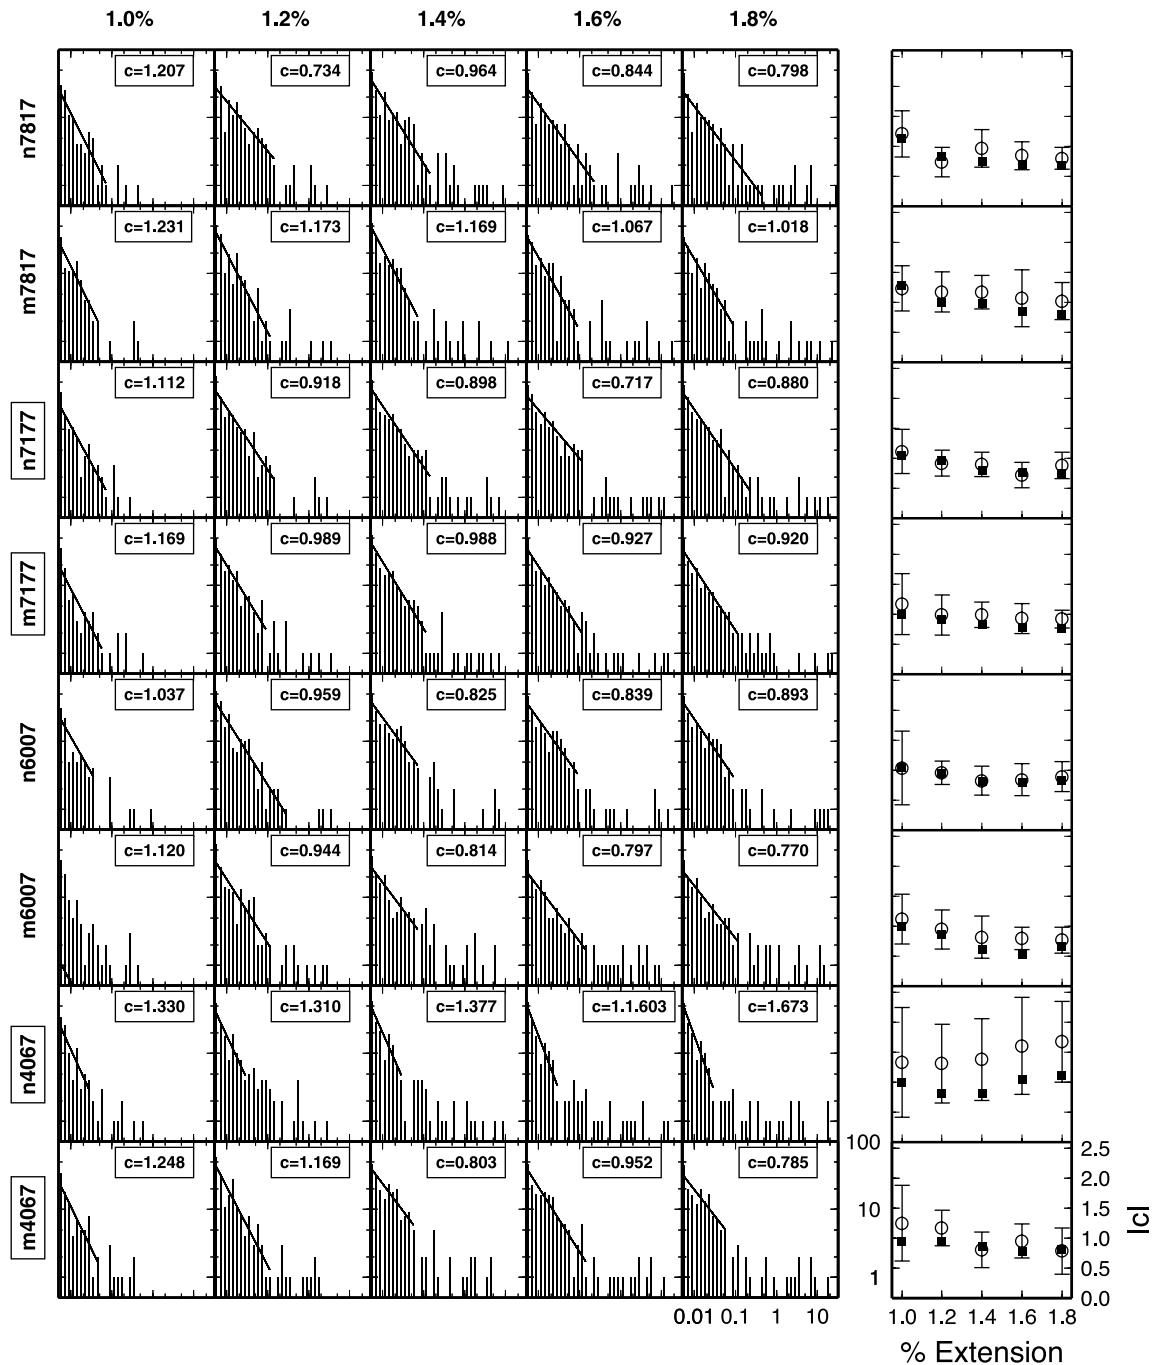
Fig. 8. Fault sized (TPS) binned in log-intervals versus discrete frequency for 1.0–1.8% strain for all eight experiments. All graphs have the same axes as the annotated plot at the bottom right of the size–frequency distributions. Power law fits are shown only over the interval used in the linear regression calculation. The right-hand column shows best fitting power law exponents, c, and 95% confidence limits determined for the log-interval (open circles) and cumulative (black squares) frequency datasets. Error bars for the cumulative frequency datasets are smaller than the size of the square symbol.

(a) intrinsic variability in the fault development even for mechanically identical systems (Fig. 7), and (b) uncertainty in the estimate of the scaling exponent (Fig. 8).

4. The source of size scaling variability

To show how the differences in fault size distributions

arise from models with statistically identical material properties, we focus on the development of the fault populations in a pair of mirror-image experiments: n4067 and m4067. Examination of the size frequency distributions at 1.8% strain shows that the small fault population in n4067 is best fit by a higher value of c than m4067 (Figs. 6 and 8). Also, the breakdown in power law scaling occurs at a smaller fault size in n4067 than in m4067. These differences

can be attributed to the spatial arrangement of large faults that are in place after only 1.2% extension in each experiment (Fig. 2). To illustrate this more clearly we show in Fig. 9 the temporal evolution of fault activity in experiments n4067 and m4067 by plotting in colour the incremental plastic strain at eight different stages of the extension history. In both experiments, in a gross sense, the deformation proceeds first by localisation of strain onto large through-going faults, followed by concentration of fault activity next to the moving wall (Fig. 9). However, in detail there are clear differences and it is these differences that affect the relative number and location of smaller faults that occur. Whereas some of the small faults that nucleate

early in the experiments may be assimilated into the large through-going 'block-bounding' faults, particularly near to the moving wall where strain becomes concentrated, many are often preserved as inactive structures within 'intra-fault' blocks. Thus, any differences in the localisation history will affect the resulting size–frequency distributions (Fig. 6). For example, the relatively high value of c seen at the end of experiment n4067 is a relict of a small fault population that formed early on and was completely abandoned by 1.4% strain within a large intra-fault block near the stationary wall (compare 1.4% and 1.6% strain for experiment n4067). In contrast, in experiment m4067 the largest intra-fault block occurs near the moving wall and localisation of strain onto the block-bounding faults is not complete until 1.8% strain. Thus a more evolved small-fault population with a wider range of sizes and a lower value of c is preserved.

The size frequency distributions of intermediate and large faults also show differences between these two experiments. A wider range of fault sizes defines the low-gradient segment at the centre of the cumulative size frequency distribution in n4067 than m4067 and fewer faults define the distribution of the very largest faults. These differences result from the different geometry of major faults in the two experiments. Because the moving wall is required to remain vertical in all experiments, the movement on the active faults must be kinematically compatible with this imposed boundary condition. Furthermore, the strain-softening rheology used in these experiments concentrates deformation onto the longer-lived, therefore softer, and therefore larger structures adjacent to the moving wall. In m4067, a pair of conjugate faults formed relatively early adjacent to the moving wall (1.4–1.6% strain in Fig. 9). This conjugate geometry is compatible with the imposed non-rotation of the side wall. In contrast in n4067, a synthetic fault array initially developed adjacent to the moving wall and this is incompatible with a non-rotating side-wall (Fig. 9). Consequently, between 1.6 and 4.0% extension in experiment n4067 all available strain energy was spent rearranging the existing large faults and developing a new antithetic fault near to the moving wall. No such rearrangement was necessary in m4067 and structures with a range of sizes continued to grow.

In Fig. 10 the difference in the evolution of faulting in experiment n4067 and m4067 is emphasised by separating the fault populations into active and inactive structures. A structure is considered to be active if its total plastic strain has increased in the last increment of extension. Thus in Fig. 10 the cumulative size frequency distributions of active and inactive fault populations are compared with the distribution for the whole population at 12 successive stages of the extension history. Overall these plots confirm the main features of Fig. 9, i.e. that strain localises onto the largest faults and that most of the small faults become inactive. However, one important point to be noted in Fig. 10 is that the size of the smallest fault that is active at any given percentage strain differs between experiments n4067 and
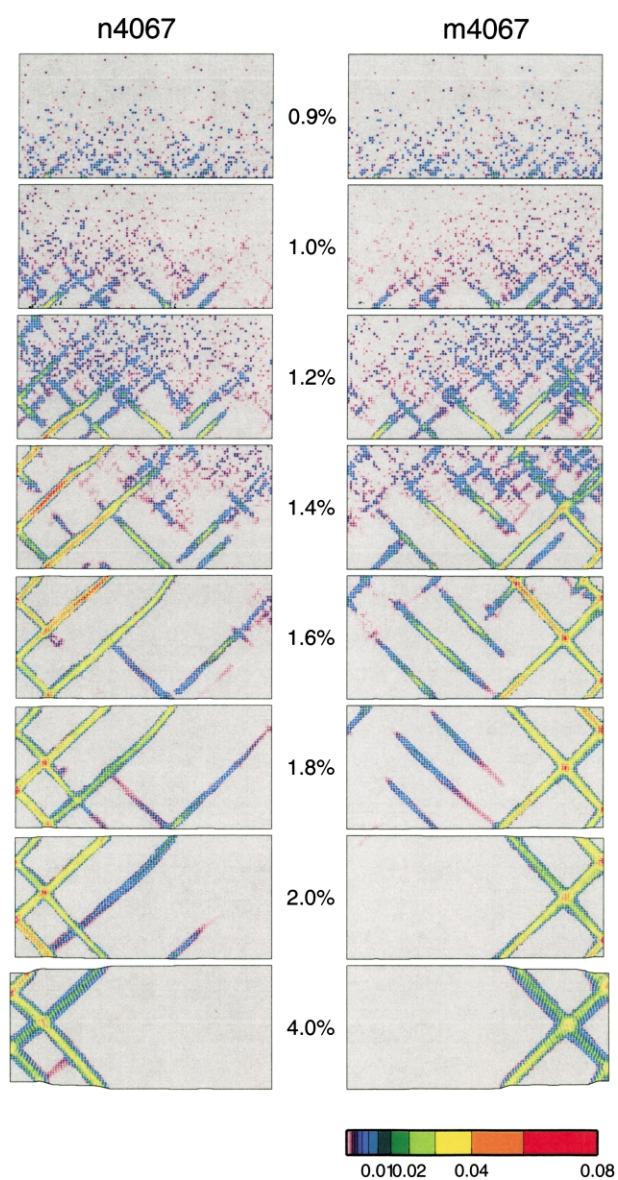

Fig. 9. Cross-sectional views of fault activity as a function of % extension for experiments n4067 and m4067. Each plot shows the increase in plastic strain in the last 0.1% (top two plots) or 0.2% (other plots) extension. Areas that experience no change in plastic strain (i.e. faults are inactive) are shown in grey.
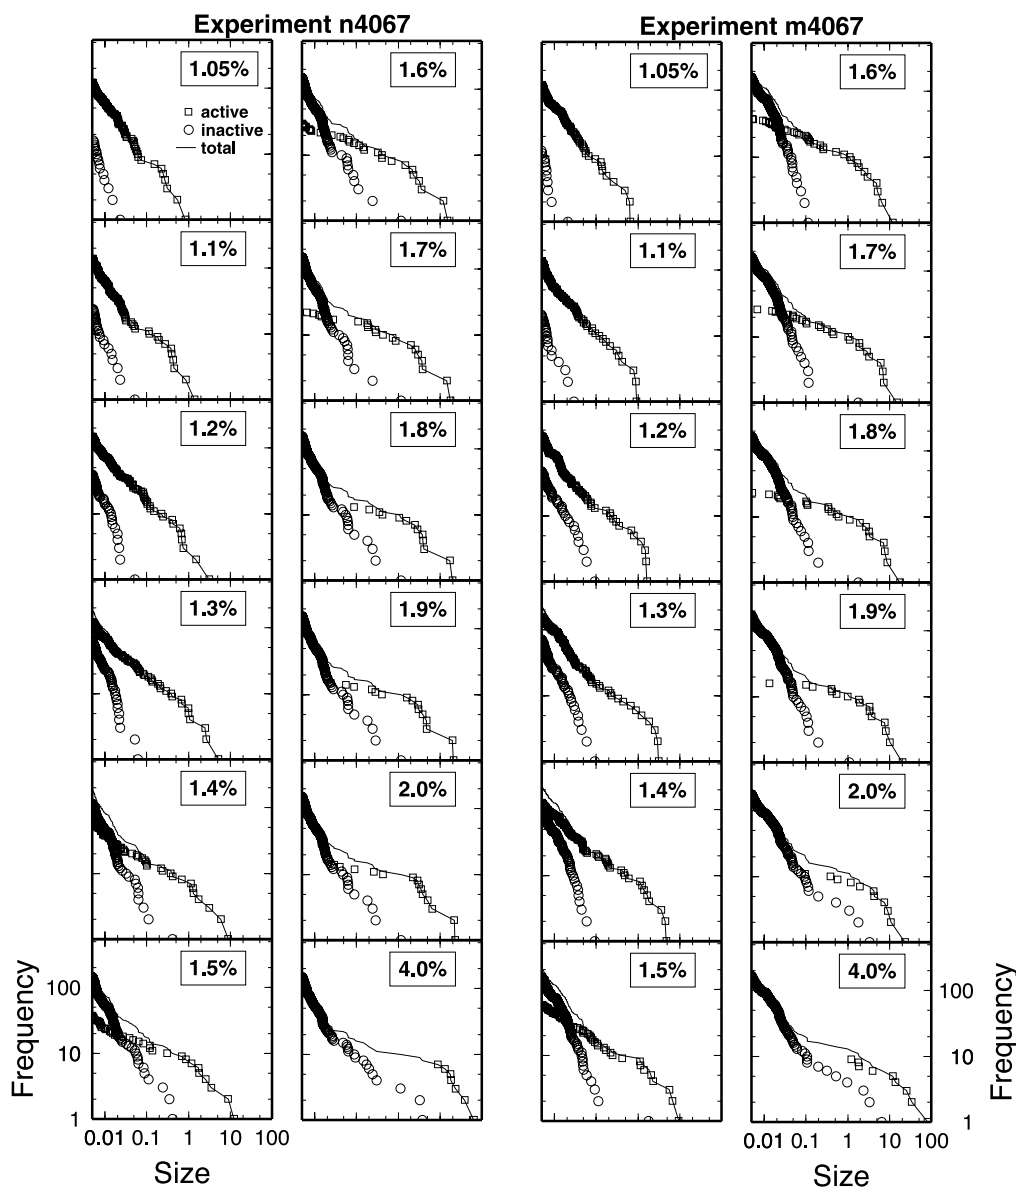
Fig. 10. Cumulative size–frequency distributions for active (squares), inactive (circles), and total (line) fault populations at 12 stages of extension from 1.05 to 4.0% strain in the mirror-image model runs: n4067 and m4067. All graphs are plotted on axes of the same scale. See Section 4 for discussion.

m4067. In experiment m4067 a significant proportion of the very smallest faults are still active at 1.8% strain, whereas in n4067 all the smallest faults are inactive after 1.7% strain. At 1.9% strain, a broad range of fault sizes in m4067 are still active while at greater strains the deformation gradually localises onto progressively larger faults. In contrast, in experiment n4067 localisation is not such a gradual or smooth process. The relatively rapid switch-off of small faults in n4067 suggests that localisation is initially proceeding rapidly (i.e. by 1.8% strain, localisation of activity onto large faults is more pronounced in n4067 than m4067; see Fig. 9). However, between 1.8 and 2.0% strain, the active and inactive fault populations in experiment n4067 do not change significantly, indicating that further localisation was inhibited during this interval of time (Fig. 10). This delay in localisation is caused by re-arrangement

of the fault structures, described above, to produce a conjugate fault geometry next to the moving wall that is compatible with the non-rotational boundary condition.

The differences in localisation behaviour between m4067 and n4067 are also reflected in the bulk stress-strain history of each system. For these two experiments, Fig. 11 presents the average extensional stress σ_{11}, on the moving wall minus the elastic component of this stress caused by the overburden. As the material properties and boundary conditions of these experiments are identical, differences in the mechanical response of the system must be due to the development of different patterns of faulting that result from simply reflecting the finite element mesh. Their behaviour is identical until 0.8% extension, at which point elements begin to fail and the material is no longer elastic. Stress continues to increase in both n4067 and m4067 as the
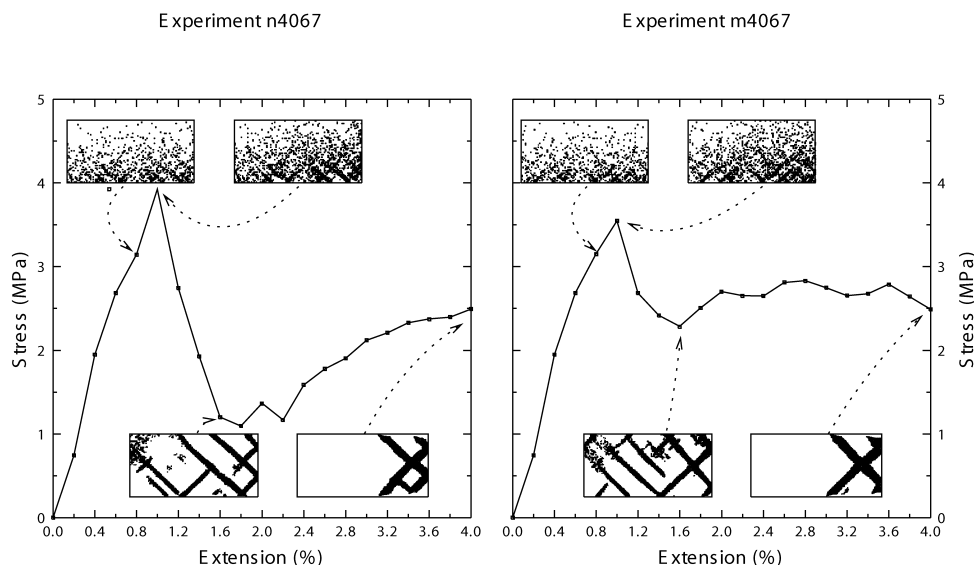
Experiment n4067   Experiment m4067

Fig. 11. Mean extensional stress on the moving side wall σ_{11} (corrected for elastic stresses induced by overburden) versus % extension for experiments n4067 and m4067. Insets show the active elements at the stages indicated. A structure is considered to be active if its total plastic strain has increased in the last 0.05% of regional extension. See Section 4 for discussion.

nucleation front spreads across the material until, at 1% extension, large faults begin to develop at the base of each model (Figs. 9 and 11). In both experiments, the bulk stress then decreases dramatically until 1.6% strain by which point many (although not all) small faults are inactive, nucleation of new faults has largely ceased and strain is being accommodated primarily by the larger structures (Fig. 9). Note that at 1.6% extension, the stress on the side wall in n4067 is actually lower than that in m4067. This indicates that it is easier to accumulate displacement on the synthetic fault set that has developed in n4067 than the crossing conjugate shear zones formed by the same stage in m4067. However, after 1.6% extension, experiment n4067 shows a marked increase in stress, i.e. overall strain hardening, even though each individual element is strain softening (Fig. 11). This strain hardening is due again to the development of the new conjugate fault near to the moving wall in experiment n4067. Thus, while localisation was initially favoured in experiment n4067 due to the early formation of a synthetic fault set, the boundary condition imposed by the vertical moving wall causes the localisation process to be inhibited at later stages in the simulation. In contrast, in experiment m4067 after 1.6% strain there is only a small increase in stress, due to shear zone thickening, followed by minor stress fluctuations.

5. Discussion

Our modelling results show that, at the same strain, materials with statistically identical material properties contain fault populations with significantly different cumulative size frequency distributions. This is because the precise arrangement of spatial heterogeneity favours certain fault growth processes over others in each of the different models. Fault nucleation, propagation and coalescence are all possible growth processes, but the process or processes that dominate can subtly vary from model to model and between different stages of the deformation in each model run. Qualitatively all the runs are similar, but in detail small changes in the relative importance of the difference growth processes at each stage of fault development controls how many faults of a given size are produced. This is not an artefact of our modelling approach—it will occur in any model (whether numerical or analogue) in which the material has heterogeneous strength properties. We might still have expected our eight experiments to generate identical statistics because the only difference between runs is a random redistribution of the strength heterogeneity. Furthermore, the boundary conditions used here (local isostasy at the base of the model) forced all the model runs to develop the same stable geometry eventually (namely two faults intersecting in an 'x' shape close to the moving wall) at the end of the experiment. This convergence towards a single active fault geometry does not occur for all basal boundary conditions (Hardacre, 2000). In other systems, which have more degrees of freedom, even greater variability in the power law exponent arises. Obviously, in nature we expect much greater complexity in rock strength heterogeneity than random variability, in addition to more complex boundary conditions that those shown here. Multiple-layered sedimentary sequences with significant contrasts in average strength between layers are common in extensional basins. Local variations in the boundary conditions are also likely, for example oblique re-activation of a pre-existing basement structure beneath an extending cover sequence.

Our results suggest therefore that extrapolating information

about fault size scaling from one area to another—for example from an area with high resolution 3D seismic imaging to a poorly-imaged adjacent area—should incorporate large uncertainty even if the areas have the same lithologies and tectonic histories. For example, the error estimate derived from least squares regression of cumulative frequency data for any one of our experiments ($\sim$0.02) is approximately 10 times smaller than the range in values of c at any given value of strain (Fig. 7). Thus if we use the results of regression through cumulative frequency data to generate other model fault populations statistically, they would not cover the full range of possible outcomes that we observe even with a very simple geological model. In contrast, the error estimate derived from least squares regression of the log-interval data is much larger ($\sim$0.5) and comparable with the range in c values. In other words, an error estimate of 0.5 (as opposed to 0.02) provides a good measure of the uncertainty involved in extrapolation for the example we present. Note that this level of uncertainty can be $>$50% of the value of c, and is sufficient to obscure any dependence of c on total strain. Clearly, real geological scenarios are likely to be associated with even greater uncertainty.

6. Summary and conclusions

Using a 2D finite element model of extensional deformation in cross-section, we simulate the development of fault populations and analyse the resulting size–frequency distributions as a function of total strain. Faulting is simulated by a process of elasto-plastic yielding and post-yield strain softening. Local isostatic compensation is included. Strength heterogeneity in the model is introduced by selecting randomly the yield strength of each element from a Gaussian probability distribution. We explain how individual faults are defined in the model and present a clustering algorithm that allows the fault populations to be extracted automatically (Section 2.3). The fault size–frequency distributions for a suite of eight deformation experiments are then compared quantitatively as a function of the amount of extension up to 4% total strain. These eight experiments are identical in a statistical sense; the only difference between each run is that the spatial arrangement of yield strengths is a different random realisation. The boundary conditions and the *average* strength properties are exactly the same in all the runs. The results of the eight experiments were used to determine both the magnitude and the cause of variability in the size–frequency scaling due to random strength heterogeneity. The purpose of this study is to provide more realistic constraints on the uncertainties involved in extrapolating fault size–frequency distributions in practical applications. The main conclusions of this study are:

1. Faults in the model appear as narrow zones along which finite plastic strain is accumulated. Fault size is defined as the total plastic strain (TPS) within each zone (Fig. 4). For this model TPS is shown to provide the most sensitive measure of fault size and represents the magnitude of the geometric moment of each structure.

2. The size–frequency distributions of all our model fault populations are well-characterised by a power law distribution, in most cases over several orders of magnitude. Power law exponents were calculated for both cumulative frequency and discrete frequency (i.e. log-interval) distributions (Section 2.5). We use least squares regression to calculate the power law exponent, c, in order to compare our results with published studies of natural fault populations.

3. We find that the value of c and the estimated error on that value depend on the method used to analyse the data (cumulative versus log-interval statistics). The standard method of fitting power laws to fault size distributions (least squares regression of log(size), log(cumulative frequency) data) was found to give error estimates ($\pm$0.02) on the value of c that were at least 10 times smaller than the variability in c observed amongst the eight experiments (typically $\pm$0.5), regardless of the strain accommodated. In contrast, using the log-interval method the estimated errors ($\pm$0.5) were comparable with the inter-model variability.

4. In the cumulative frequency data, the value of the exponent c decreases with increasing strain in most of the model runs. This observation is consistent with the results of previous physical (Sornette et al., 1993), geometrical (Cladouhos and Marrett, 1996) and numerical (Sornette and Davy, 1991; Cowie et al., 1995) models. The decrease in the value of c is partly due to fault coalescence but also arises because larger faults grow faster than smaller ones in our model. However, the dependence of c on total strain is subtle compared with the inter-model variability. In other words, the variation in c between experiments at *one* value of strain is comparable with the range observed as a function of strain for any one experiment. If natural fault population development is similarly sensitive to the spatial arrangement of rock strengths, it is unsurprising that attempts to relate the power law exponent of real fault populations to total strain have had little success (Cladouhos and Marrett, 1996).

5. By comparing two experiments in detail (Figs. 9–11), we show that the variability in size–frequency distributions between models arises as a consequence of small differences in the fault growth history. Depending on the spatial distribution of strength relative to the imposed boundary conditions, faults initiate in different locations and thus interact in subtly different ways and to differing degrees with neighbouring faults. These differences are sufficient to result in significantly different power law exponents on cumulative frequency plots even though all the fault patterns are qualitatively very similar.

6. The range in power law exponents that we observe in these experiments is clearly controlled to a significant degree by a geologically simplistic (i.e. random) variation in material properties. Furthermore, the local isostatic basal boundary condition used here tends to suppress the potential variability by limiting the range of possible fault geometries that can form. Realistic strength variations and boundary conditions in extensional provinces are likely to be much more complex and thus are expected to produce even greater variability amongst the resulting fault populations.

7. In extrapolating fault population data to smaller scales or modelling fault populations in poorly imaged areas, our results suggest that the level of uncertainty associated with such 'predictions' should be much larger than the error estimate derived from least squares regression of cumulative frequency data.

Acknowledgements

References

Ackermann, R., Withjack, M., Schlische, R., 2001. The geometric and statistical evolution of normal fault systems: an experimental study of the effects of mechanical layer thickness on scaling laws. Journal of Structural Geology 23, 1803–1819.

Anderson, E., 1951. The Dynamics of Faulting, Oliver and Boyd, London.

Barnichon, J.D., Charlier, R., 1996. Finite element modelling of the competition between shear bands in the early stages of thrusting: strain localization and constitutive law influence. In: Buchanan, P.G., Niewland, D.A. (Eds.), Modern Developments in Structural Interpretation, Validation and Modelling. Geological Society Special Publication 99, pp. 235–250.

Belytschko, T., Chiang, H.-Y., Plaskacz, E., 1994. High resolution two-dimensional shear band computations: imperfections and mesh dependence. Computer Methods in Applied Mechanics and Engineering 119, 1–15.

Board, M., 1989. FLAC (Fast Lagrangian Analysis of Continua) Version 2.20: Software Summary. Itasca Consulting Group, Inc., Minnesota.

Bonnet, E., Bour, O., Odling, N., Davy, P., Main, I., Cowie, P., Berkowitz, B., 2001. Scaling of fracture systems in geological media. Reviews of Geophysics 39, 347–383.

Cladouhos, T., Marrett, R., 1996. Are fault growth and linkage models consistent with power-law distributions of fault length? Journal of Structural Geology 22, 983–997.

Clarke, R., Cox, S., Laslett, G., 1999. Generalizations of power-law distributions applicable to sampled fault trace lengths: model choice, parameter estimation and caveats. Geophysical Journal International 136, 357–372.

Clifton, A., Schlische, R., Withjack, M., Ackermann, R., 2000. Influence of rift obliquity on fault population systematics: results of experimental clay models. Journal of Structural Geology 22, 1491–1509.

Cowie, P.A., 1998. The growth of normal faults in three dimensions in continental and oceanic crust. In: Buck, W.R. Delaney, P.T., Karson, J.A., Lagabrielle, Y. (Eds.), Faulting and Magmatism at Mid-Ocean Ridges. American Geophysical Union Monograph Series 106, pp. 325–348.

Cowie, P., Vanneste, C., Sornette, D., 1993. Statistical physics model for the spatiotemporal evolution of faults. Journal of Geophysical Research 98, 21809–21821.

Cowie, P.A., Sornette, D., Vanneste, C., 1995. Multifractal scaling properties of a growing fault population. Geophysical Journal International 122, 457–469.

Cox, S., Paterson, L., 1990. Damage development during rupture of heterogeneous brittle materials: a numerical study. In: Knipe, R.J., Rutter, E.H. (Eds.), Deformation Mechanisms, Rheology and Tectonics. Geological Society Special Publication 54, pp. 57–62.

Cundall, P., 1989. Numerical experiments on localization in frictional materials. Ingenieur Archives 58, 148–159.

Davis, R., Fletcher, R., 1990. Shear bands in a plastic layer at yield under combined shortening and shear: a model for the fault array in a duplex. In: Knipe, R.J., Rutter, E.H. (Eds.), Deformation Mechanisms, Rheology and Tectonics. Geological Society Special Publication 54, pp. 123–131.

Davy, P., Hansen, A., Bonnet, E., Zhang, S.-H., 1995. Localization and fault growth in layered brittle–ductile systems: implications for deformations of the lithosphere. Journal of Geophysical Research 100, 6281–6294.

Dresden, G., Gwildis, U., Kluegel, T., 1991. Numerical and analogue modelling of normal fault geometry. In: Roberts, A., Yielding, G., Freeman, B. (Eds.), The Geometry of Normal Faults. Geological Society Special Publication 56, pp. 207–217.

Finch, E., 1998. A crustal lattice solid model: the evolution, geometry and scaling of tectonic extension. PhD thesis, University of Ulster.

Gauthier, B., Lake, S., 1993. Probabilistic modelling of faults below the limit of seismic resolution in Pelican Field, North Sea, offshore United Kingdom. The American Association of Petroleum Geologists Bulletin 77, 761–777.

Gerbault, M., Poliakov, A., Daignieres, M., 1998. Prediction of faulting from the theories of elasticity and plasticity: what are the limits? Journal of Structural Geology 20, 301–320.

Hardacre, K., 2000. Controls on fault network evolution and population statistics—insights from field studies and numerical modelling. PhD thesis, Edinburgh University.

Hardacre, K.H., Cowie, P.A., 2003. Controls on strain localisation in a 2D elasto-plastic layer: Insights into size-frequency scaling of extensional fault populations. J. Geophys. Res., in press.

Hassani, R., 1994. Modélisation numérique de la déformation des systèmes géologiques. Thèse de Doctorat, Université de Montpellier II, France.

Heimpel, M., Olson, P., 1996. A siesmodynamical model of lithosphere deformation: development of continental and oceanic rift networks. Journal of Geophysical Research 101, 16155–16176.

Hobbs, B., Mülhaus, H., Ord, A., 1990. Instability, softening and localization of deformation. In: Knipe, R.J., Rutter, E.H. (Eds.), Deformation Mechanisms, Rheology and Tectonics. Geological Society Special Publication 54, pp. 183–192.

Homberg, C., Hu, J., Angelier, J., Bergerat, F., Lacombe, O., 1997. Characterisation of stress perturbations near major fault zones: insights from 2D distinct element numerical modelling and field studies. Journal of Structural Geology 19, 703–718.

Huc, M., Hassani, E., Chéry, J., 1998. Large earthquake nucleation associated with stress exchange between middle and upper crust. Geophysical Research Letters 25, 551–554.

Lavier, L.L., Buck, W.R., Poliakov, A.N.B., 2000. Factors controlling

normal fault offset in an ideal brittle layer. Journal of Geophysical Research 105, 23431–23442.

Leroy, Y., Ortiz, M., 1989. Finite element analysis of strain localization in frictional materials. International Journal of Numeric and Analytic Methods in Geomechanics 13, 53–74.

Maillot, B., Cowie, P., Lague, D., 1998. Simulating polyphase faulting with a tensorial 3D model of fault growth. In: Jones, G., Fisher, Q., Knipe, R. (Eds.), Faulting, Fault Sealing and Fluid Flow in Hydrocarbon Reservoirs. Geological Society Special Publication 147, pp. 209–216.

Main, I., 2000. Apparent breaks in scaling in the earthquake cumulative frequency–magnitude distribution: fact or artefact? Bulletin of the Seismological Society of America 90, 86–97.

Malinverno, A., Cowie, P.A., 1993. Normal faulting and the topographic roughness of midocean ridge flanks. Journal of Geophysical Research 98, 17921–17939.

Mansfield, C., Cartwright, J., 2001. Fault growth by linkage: observations and implications from analogue models. Journal of Structural Geology 23, 745–763.

Marrett, R., Allmendinger, R., 1991. Estimates of strain due to brittle faulting: sampling of fault populations. Journal of Structural Geology 13, 735–738.

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

Martel, S., 1997. Effects of cohesive zones on small faults and implications for secondary fracturing and fault trace geometry. Journal of Structural Geology 19, 835–847.

Martel, S., 1999. Mechanical controls on fault geometry. Journal of Structural Geology 21, 585–596.

McKinnon, S., Garrido de la Barra, I., 1998. Fracture initiation, growth and effect on stress field: a numerical investigation. Journal of Structural Geology 20, 1673–1689.

Niño, F., 1997. Modélisation numérique de la deformation localisée et de l'activité tectonique des failles. Thèse de Doctorat, Université de Montpellier II, France.

Ord, A., 1990. Mechanical controls on dilatant shear zones. In: Knipe, R.J., Rutter, E.H. (Eds.), Deformation Mechanisms, Rheology and Tectonics. Geological Society Special Publication 54, pp.183–192.

Ortega, O., Marrett, R., 2000. Prediction of macrofracture properties using microfracture information, Mesaverde Group sandstones, San Juan basin, New Mexico. Journal of Structural Geology 22, 571–588.

Pickering, G., Bull, J., Sanderson, D., 1995. Sampling power-law distributions. Tectonophysics 248, 1–20.

Poliakov, A., Hermann, H., 1994. Self-organised criticality of plastic shear bands in rocks. Geophysical Research Letters 21, 2143–2146.

Poliakov, A., Hermann, H., Podlachikov, Y., Roux, S., 1994. Fractal plastic shear bands. Fractals 2, 567–581.

Rudnicki, J., 1977. The inception of faulting in a rock mass with a weakened zone. Journal of Geophysical Research 82, 844–854.

Rudnicki, J., Rice, J., 1975. Conditions for the localisation of deformation in pressure sensitive dilatant materials. Journal of the Mechanics and Physics of Solids 23, 371–394.

Scholz, C., Cowie, P., 1990. Determination of total strain from faulting using slip measurements. Nature 346, 837–838.

Schultz-Ela, D., Jackson, M., 1996. Relation of sub-salt structures to supra-salt structures during extension. American Association of Petroleum Geologists Bulletin 80, 1896–1924.

Sornette, A., Davy, P., Sornette, D., 1993. Fault growth in brittle–ductile experiments and the mechanics of continental collisions. Journal of Geophysical Research 98, 12,111–12,139.

Sornette, D., Davy, P., 1991. Fault growth and the universal fault length distribution. Geophysical Research Letters 18, 1081–1097.

Spyropoulos, C., 1999. A study of complex systems: development of crack populations. PhD thesis, Columbia University, New York.

Spyropoulos, C., Griffith, W., Scholz, C., Shaw, B., 1999. Experimental evidence for different strain regimes of crack populations in a clay model. Geophysical Research Letters 26, 1081–1084.

Spyropoulos, C., Scholz, C., Shaw, B., 2002. Transition regimes for growing crack populations. Physical Review E 65, 056105.

Walsh, J.J., Childs, C., Imber, J., Manzocchi, T., Watterson, J., Nell, P.A.R., 2003. Strain Localisation and population changes during fault system growth within the Inner Moray Firth, Northern North Sea. J. Struct. Geol. 25, 307–315.

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

Wessel, P., Smith, W.H.F., 1998. A new improved version of Generic Mapping Tools released. EOS Trans. Am. Geophys. U. 79(47), 579.

Westaway, R., 1994. Quantitative analysis of populations of small faults. Journal of Structural Geology 16, 1259–1273.

Wilson, S., Henderson, J., Main, I., 1996. A cellular automaton model: the influence of heterogeneity in the failure process. Journal of Structural Geology 18 (2/3), 343–348.

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

Yielding, G., Neeham, T., Jones, H., 1996. Sampling of fault populations using sub-surface data: a review. Journal of Structural Geology 18, 135–146.