Michael John Welch
Mikael Lüthje
Simon John Oldfield

Modelling the Evolution of Natural Fracture Networks

Methods for Simulating the Nucleation, Propagation and Interaction of Layer-Bound Fractures

Springer

Modelling the Evolution of Natural Fracture Networks

Michael John Welch · Mikael Lüthje ·
Simon John Oldfield

Modelling the Evolution of Natural Fracture Networks

Methods for Simulating the Nucleation, Propagation and Interaction of Layer-Bound Fractures

Springer

Michael John Welch
Danish Hydrocarbon Research and Technology Centre
Technical University of Denmark
Kongens Lyngby, Denmark

Mikael Lüthje
Danish Hydrocarbon Research and Technology Centre
Technical University of Denmark
Kongens Lyngby, Denmark

Simon John Oldfield
Danish Hydrocarbon Research and Technology Centre
Technical University of Denmark
Kongens Lyngby, Denmark

ISBN 978-3-030-52413-5 ISBN 978-3-030-52414-2 (eBook)
https://doi.org/10.1007/978-3-030-52414-2

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface

Natural fracture networks are important in controlling the mechanical behaviour and flow of fluids through geological formations. It is, therefore, essential to include fracture networks in the static geomodels that are used to model such behaviour, in applications ranging from tunnel excavation and mining to groundwater management, hydrocarbon exploration and production, geothermal energy extraction and CO_2 sequestration. However, fractures cannot normally be mapped directly in the subsurface, as they are below the resolution of geophysical data, and boreholes provide only very limited data coverage. Traditional solutions to this problem include building stochastic fracture models, in which fractures of arbitrary size are placed at random locations in an attempt to match the fracture densities measured in the boreholes, or using numerical methods such as the finite element method to simulate the nucleation and growth of the fractures. However, the former method produces inaccurate and geologically unrealistic fracture models, while the latter is too computationally expensive to be practical for any but the simplest fracture systems.

In this book, we present a new method of simulating the growth of layer-bound fracture networks, which is based on fundamental geomechanical principles, but is simple enough to model large networks containing hundreds of thousands of fractures across major geological structures such as anticlines and diapirs. This is achieved by combining the established theories of subcritical fracture propagation, linear elastic fracture mechanics and fracture distribution to derive quantitative expressions describing the evolution of the fracture network, based on the mechanical properties of the host rock and the deformation history.

We start by modelling the growth of small circular fractures within a homogeneous layer. We then focus on layer-bound fractures that are confined within brittle competent layers (which often act as aquifers or reservoirs) sandwiched between more ductile layers (which often act as seals). By applying the expressions for fracture growth to cumulative distribution functions describing the fracture populations as a whole, we can model the evolution of these populations through time without needing to model the growth of each fracture individually. To do this accurately, we must also model the effects of fracture interaction; this includes both intersection with perpendicular or oblique fractures and stress shadow interaction with parallel fractures. The resulting expressions can predict the density of fractures of different

sizes, and in different orientations, at different locations across the geological structure we are modelling, as well as predicting important properties of the fracture network such as fracture porosity and connectivity. Finally, we can use these results to generate a geomechanically consistent explicit Discrete Fracture Network (DFN) model.

This method can be used to explore the controls on different aspects of the fracture networks that develop under different conditions. We show that the mean linear density (P_{32}) of layer-bound fractures is mainly controlled by the presence of stress shadows and the layer thickness, while the volumetric density (P_{30}) and mean length of layer-bound fractures are controlled by many factors, including the duration of deformation, the subcritical fracture propagation index (which controls the fracture propagation rate), the initial microfracture population and the interactions between the fractures (which are in turn dependent on anisotropy in the applied strain). We show that fracture networks may not start to develop immediately when horizontal strain is applied, but that once they do start to develop they grow very rapidly, often reaching "saturation" (when no further fractures can nucleate or propagate) within tens of thousands of years or less. On a large scale, the fracture network often appears first at the location of greatest strain on the large-scale geological structure (e.g. on the crest of an anticline or diapir), and then sweeps outwards across the structure as a sharp "deformation front". If we look at the large-scale structure at any specific point in time, therefore, we can divide it into unfractured and fully fractured (saturated) zones separated by a clearly defined boundary, rather than seeing a gradual variation in fracture density across the structure. We also show that mode of the fractures (i.e. Mode 1 dilatant fractures verse Mode 2 shear fractures) is dependent not only on the mechanical properties and fluid overpressure (which we might expect), but also the propagation rate (Mode 1 fractures are more likely to form if the subcritical fracture propagation index is high). Finally, we investigate ways of characterising the anisotropy and connectivity of the fracture network and show how these properties evolve as the fracture network develops.

A method of this complexity requires calibration against real fracture networks before we can apply it with confidence to the subsurface. We have selected three fractured outcrops from the UK for this purpose. The Nash Point outcrop in south Wales exposes a uniaxial fracture network (i.e. comprising a single set of parallel fractures) in thin brittle limestone beds sandwiched between ductile shales (Maerten et al. 2016). We show that the thickness of the limestone beds controls the observed fracture spacing, and we show how to replicate the observed fracture lengths by adjusting the subcritical fracture propagation index and the initial microfracture density. The Robin Hood's Bay outcrop in northeast England exposes an orthogonal fracture network, comprising two perpendicular sets of layer-bound fractures, again confined within thin brittle limestone beds sandwiched between ductile shales (Rawnsley et al. 1993). The anisotropy of this network varies laterally across the outcrop; we show that this variation in fracture anisotropy is the result of lateral variation in the local strain around the km-scale Peak Fault, which lies to the east of the outcrop. Finally, the Pegwell Bay outcrop in southeast England exposes a series of fracture corridors propagating outwards from strike-slip faults in chalk (Souque et al. 2019). We show

that it is possible to reproduce the growth of these fracture corridors, but only in conditions of high fluid overpressure, high subcritical fracture propagation index and without the development of stress shadows around the fractures.

Finally we use this new method to predict fracture networks in two subsurface examples. The Kraka field is a hydrocarbon field offshore Denmark, producing from a fractured chalk reservoir overlying a salt pillow. We show that the fracture network generated by a combination of growth of the salt pillow and local strain around a set of seismically-mapped faults gives a good match for the fractures observed on borehole images and in core (described by Aabø et al. 2019). The Anloo salt diapir near Drenthe in the Netherlands is considered a potential prospect for geothermal energy extraction, utilising several fractured reservoir layers overlying the diapir. Although the large-scale structure has been mapped out from seismic data, there is very little well data available with which to constrain the fracture population. We show how the new method can be used to identify the areas most likely to be fractured and predict the most likely fracture geometries at different points around the structure. We then show how to quickly build multiple, geologically realistic fracture models from this limited dataset for use in uncertainty and risk analysis.

Kongens Lyngby, Denmark
July 2020

Michael John Welch
Mikael Lüthje
Simon John Oldfield

Acknowledgements

The authors kindly acknowledge the Danish Underground Consortium (Total E&P Denmark, Noreco & Nordsøfonden) for providing data for the Kraka field and granting the permission to publish this work. This research has received funding from the Danish Hydrocarbon Research and Technology Centre (DHRTC) under the Advanced Water Flooding programme.

We would also like to thank Total E&P Denmark, and especially Amit Singh, Alain Lejay, Maurits de Heer, Klaas Kostro and Alan Cunningham for their help in testing the model output in their dynamic simulations, providing feedback and good discussions.

We would like to thank everyone at DHRTC and partner institutions who have contributed to this research project. In particular we would like to thank Florian Smit providing us with input and help building the Drenthe model. Thanks also goes to Tala Aabø, Jesper Dramsch and Solomon Seymun for providing us with results from their analysis of borehole images, seismic data and core for the Kraka field. We would like to acknowledge Frederic Amour letting us use results from his extensive work characterising the elastic moduli of the chalk in Kraka. We would like to give a special thanks to Aslaug Clemmensen Glad for her work on the sensitivity analyses.

We would like to thank Ole Rønø Clausen and Kenni Petersen at Aarhus University for providing us with strain data for the Kraka field, and for allowing us to use the strain history modelling software that they developed to generate strain data for Drenthe.

We would like to thank Bertrand Gauthier and his colleagues in the Naturally Fractured Reservoir Team at Total for helpful feedback and discussion.

We would like to acknowledge Ordnance Survey data and style sheets shared under the Open Government Licence.

We would like to acknowledge data shared by the Geologische Dienst Nederland, part of TNO, through dinoloket.nl.

Finally, we would also like to acknowledge our colleagues, students, families and friends who have often contributed to our work through discussion and questioning of the processes and concepts behind this book.

Contents

Constants and Symbols

The following fundamental geometric and material properties are user-specified input parameters in the calculations. Where appropriate, typical values are given for consolidated sedimentary rocks at c.2km burial depth.

Table 1 List of fundamental geometric and material properties used in these calculations

Symbol	Property	Typical value (SI Unit)
A	Rupture velocity in the rock matrix; similar to the sonic velocity	2000 m/s
b	Subcritical fracture propagation index	10–50
G_c	Crack surface energy of the rockmass	1000 J/m^2
E_r	Young's Modulus of the intact rock	10^{10} Pa (=10 GPa)
ν_r	Poisson's ratio of the intact rock	0.25
υ	Initial compaction factor	0–1
μ_{fr}	Friction coefficient of fracture	0.5
B	Initial microfracture density coefficient	0.001 fractures/m$^{3\text{-}c}$
c	Initial microfracture distribution coefficient	2
h	Layer thickness	0.1–100 m
z	True Vertical Depth (TVD)	0–5000 m
g	Gravitational constant	9.81 m/s^2
ρ_b	Bulk rock density	2200–2500 kg/m^3
ρ_f	Pore fluid density	1000 kg/m^3
t_r	Strain relaxation time constant	0.1 ma $\approx 3.15\times10^{11}$ s

The following symbols represent derived variables that are parameters of or output from the calculations described in this document. In some cases, modifiers (subscripts or superscripts) are applied to constrain the scope of the variables (for example, so they apply only to a specific fracture type or set, or a specific timestep). Where these modifiers are absent, the variables can be assumed to apply generally or in situations in which the scope is implied or not defined; e.g. if we are considering only one fracture type or set, if the timestep is undefined.

Table 2 List of variables used in these calculations

Symbol	Property	SI Unit
N	Current timestep	none
M	Timestep in which specified macrofracture nucleated	none
J,K	Counters used for iteration through timesteps	none
I,J	Represent two perpendicular-striking fracture sets	none
t_0	Initial time	s
t_n	A time during timestep **N**	s
t_N	Time at the end of timestep **N**	s
Δt	Time increment (e.g. a timestep duration)	s
ω	Fracture dip	rad
r	Microfracture radius	m
L	Macrofracture length (along strike)	m
ℓ	Half-macrofracture length (along strike)	m
K_c	Critical stress intensity (also known as fracture toughness)	$Nm^{3/2}$
K_f	Stress intensity at the tip of a fracture	$Nm^{3/2}$
σ_v'	Vertical effective stress (+ve for compressive stress)	Pa
σ_n'	Effective normal stress acting on a fracture	Pa
τ	Shear stress acting on a fracture	Pa
σ_{dN}	Driving stress acting on a fracture during timestep **N**	Pa
ε_{hmax}	Horizontal strain in the maximum (most compressional) strain direction (+ve for compressional strain)	ratio
ε_{hmin}	Horizontal strain in the minimum (most extensional) strain direction	ratio
${}_{MF}\varepsilon_h$	Horizontal strain accommodated by elastic displacement on a set of macrofractures	ratio
α_{MF}	Macrofracture propagation rate factor	variable
β	Factor related to subcritical fracture propagation index	none
γ	Instantaneous microfracture growth factor	variable
Γ_N	Cumulative microfracture growth factor at the end of timestep **N**	variable
$\Lambda_{N,M}$	Cumulative macrofracture growth factor between the end of timestep **M** and the end of timestep **N**	variable

(continued)

Table 2 (continued)

Symbol	Property	SI Unit
$_{\mu F}P_{30}$	Volumetric microfracture density (number of microfractures per unit volume)	fracs/m^3
$a_{\mu F}P_{30}$	Volumetric density of active microfractures	fracs/m^3
$s_{\mu F}P_{30}$	Volumetric density of static microfractures	fracs/m^3
$_{\mu F}P_{30}(r,t)$	Volumetric density of microfractures with radius $\geq$ r at time t	fracs/m^3
$_{\mu F}P_{32}$	Mean linear microfracture density (microfracture area per unit volume)	m^2/m^3 = fracs/m
$_{MF}P_{30}$	Volumetric macrofracture density	fracs/m^3
$_{MF}P_{30}(L, t)$	Volumetric density of macrofractures with length $\geq$ L at time t	fracs/m^3
$_{MF}P_{32}$	Mean linear macrofracture density	fracs/m
$_{\underline{MF}}P_{30}$	Volumetric half-macrofracture density	fracs/m^3
$_{\underline{MF}}P_{30}(\ell, t)$	Volumetric density of half-macrofractures with half-length $\geq \ell$ at time t	fracs/m^3
$_{\underline{MF}}P_{32}$	Mean linear macrofracture density	fracs/m
$\mathbb{I}_{30}$	Anisotropy of the fracture network, calculated from volumetric macrofracture density	ratio
$\mathbb{I}_{32}$	Anisotropy of the fracture network, calculated from mean linear macrofracture density	ratio
$\mathbb{V}_{MF}(\mathbf{L})$	Internal volume of a Mode 1 dilatant macrofracture with length L	m^3
ϕ_{MF}	Total macrofracture porosity (for Mode 1 macrofractures)	ratio
η_{MF}	Total macrofracture volumetric heave (for Mode 2 macrofractures)	ratio
W	Macrofracture stress shadow width	m
ψ_{MF}	Total macrofracture stress shadow volume	ratio
$\mathcal{X}_{MF}$	Total macrofracture exclusion zone volume	ratio
q	Instantaneous probability of microfracture deactivation	s^{-1}
$\overline{q}_{\mathbf{N}}$	Mean probability of microfracture deactivation during timestep **N**	s^{-1}
$\theta_{\mathbf{N}}$	Cumulative microfracture activation probability to the end of timestep **N**	ratio
$\theta_{\mathbf{N}}{}'$	Effective cumulative microfracture activation probability to the end of timestep **N**	ratio
$F_{I,I,\mathbf{N}}$	Instantaneous probability of deactivation for a macrofracture of set I due to stress shadow interaction during timestep **N**	s^{-1}
$\Phi_{I,I,\mathbf{N}}$	Cumulative macrofracture activation probability for a macrofracture of set I due to stress shadow interaction during timestep **N**	ratio
$\bar{F}_{I,J,\mathbf{N}}$	Mean probability of deactivation for a macrofracture of set I due to intersection with a macrofracture of set J during timestep **N**	s^{-1}
$\Phi_{I,J,\mathbf{N}}$	Cumulative macrofracture activation probability for a macrofracture of set I due to intersection with a macrofracture of set J during timestep **N**	ratio

(continued)

Table 2 (continued)

Symbol	Property	SI Unit
$\bar{F_N}$	Mean probability of macrofracture deactivation for any reason during timestep **N**	s^{-1}
$\Phi_{\mathbf{N,M}}$	Cumulative macrofracture activation probability from the end of timestep **M** to the end of timestep **N**	ratio
$\xi(t)$	Stress shadow interaction box growth rate at time t	$m^{-2}s^{-1}$
${}^{I,I}\varsigma$	Stress shadow obstruction modifier for static half-macrofracture tip density	ratio
${}^{I,J}\varsigma$	Intersection obstruction modifier for static half-macrofracture tip density	ratio
${}_{IB}P_{32}(t + dt)$	Total macrofracture area per unit volume at time t plus the cross-sectional area of the stress shadow interaction boxes for the interval dt	ratio
$\mathcal{X}_{IB}(t + dt)$	Total macrofracture exclusion zone volume at time t plus the volume of the interaction boxes for the interval dt	ratio
S(x)	Cumulative macrofracture spacing distribution function, for macrofracture spacing x	m^{-1}
$\mathbb{A}$	Coefficient for cumulative macrofracture spacing distribution function	m^{-1}
$\mathbb{B}$	Exponent for cumulative macrofracture spacing distribution function	m^{-1}
M_{hh}	Fracture mode factor relating horizontal stress to horizontal strain	none
M_{vh}	Fracture mode factor relating vertical stress to horizontal strain	none
M_{hv}	Fracture mode factor relating horizontal stress to vertical strain	none
M_{vv}	Fracture mode factor relating vertical stress to vertical strain	none

Chapter 1
Introduction

Abstract In this chapter, we review previous approaches to modelling natural fracture networks in the subsurface and explain the drawbacks of each. The most common method is implicit stochastic modelling, where the fracture population is represented statistically by one or more spatially variable properties. The main drawback is that there is generally very limited data available to populate such models, so they are usually incomplete and unreliable. Another approach is to build an explicit Discrete Fracture Network (DFN) model in which fractures are represented as geometric objects; however, because fractures cannot be imaged directly in the subsurface, these models also suffer from the drawback of uncertainty due to limited data. A solution to this problem is to simulate the growth of the fracture network dynamically. This has the advantage that it will generate geologically realistic models based only on the geometry of the larger geological structure and the mechanical properties, which are generally known. However, conventional numerical methods such as the finite element method are far too computationally expensive to simulate a network of hundreds of thousands of fractures across a large geological structure. We therefore lay out the requirement and specifications for a new geomechanically-based method of simulating the growth of large fracture networks. We explain that this could be achieved by combining the theories of subcritical fracture propagation, linear elastic fracture mechanics and fracture size distributions. Such a method would be able to generate both implicit fracture models and explicit DFNs.

Keywords Fracture modelling · Discrete fracture networks · Naturally fractured reservoir · Geomechanics · Reservoir modelling

The purpose of this book is to present a new method for simulating the growth of a network of layer-bound fractures during an episode of geological deformation. Our method provides detailed statistical information on the development of the fracture network as a whole, as well as a means for visualising this development in the form of an explicit geometric model of the fractures. It is hoped that this method will help in understanding the controls on the development and resulting geometry of natural fracture networks in real geological formations, as well as provide improved

M. J. Welch et al., *Modelling the Evolution of Natural Fracture Networks*,
https://doi.org/10.1007/978-3-030-52414-2_1

tools for modelling geomechanical behaviour and fluid flow in fractured layers in the subsurface.

1.1 Background and Previous Work

Many geological formations contain extensive networks of natural fractures that developed in response to an externally-applied horizontal strain. This strain can be generated by various mechanisms: regional tectonic deformation, flexure over large folds or diapirs, local deformation around larger faults, or high pressure fluids injected into the formation. The fractures that develop in response to the strain may be vertical Mode 1 dilatant fractures or inclined Mode 2 shear fractures, depending on the in situ stress state and fluid pressure at the time of deformation (Anderson 1951). If the formation comprises interbedded layers with contrasting mechanical properties, the fractures may often be confined within brittle layers sandwiched between more ductile layers, forming "layer-bound fractures" (Gross and Engelder 1995). In such cases, the spacing of the layer-bound fractures may be controlled by the "stress shadows" that develop around them. These stress shadows are zones, surrounding each fracture, in which the applied horizontal strain is accommodated by displacement on the fracture rather than by elastic strain in the host rock. The in situ stress that drives fracture propagation is much reduced within these stress shadows, and therefore other fractures cannot propagate into them (Pollard and Segall 1987; Bai and Pollard 2000; Bai et al. 2000a, b).

The overall fracture population within each brittle layer can typically be subdivided into sets of approximately parallel-striking fractures. If the applied horizontal strain was biaxial, we may see two orthogonal fracture sets, one striking perpendicular to the maximum horizontal strain and the other striking perpendicular to the minimum horizontal strain (Pollard and Aydin 1988; Lorenz et al. 1991); outcrop examples include those described by Rawnsley et al. (1993), Engelder and Peacock (2000), and Gillespie et al. (2001).

Such fracture networks will have a major impact on fluid flow through a fractured layer, especially when the host rock has low permeability, as they will provide important long-distance fluid conduits. They will also have a significant impact on the strength and elastic properties of the fractured layer as a whole (the "bulk rock" properties), potentially introducing planes of weakness and elastic anisotropy to the layer. It is, therefore, important to incorporate the fracture network into any 3D models of the subsurface geology ("static geomodels") that may be used for fluid flow modelling, geomechanical modelling or assessing rock stability.

Direct evidence of fractures in the subsurface is, however, often elusive. Their impact on seismic anisotropy is known, but may be hard to differentiate from stress effects, and they are typically both too small and lack sufficient contrast in elastic properties to be resolved or detected by seismic imaging. There may be some limited data from boreholes: fractures can be identified and logged in core, where this is

available, and those fractures which have acoustic or electrical properties with sufficiently contrast to the host rock may be visible on borehole image logs. It may also be possible to identify fractures indirectly through their impact on bulk rock properties (e.g. on bulk density or sonic velocity logs) or on fluid flow (e.g. from production logging tools or temperature logs).

In outcrop, fractures are usually visible, but may be subject to preferential erosion. Outcrops often give us a good understanding of the overall geometry of a fracture network, although this is normally limited to two-dimensions. At a smaller scale, some laboratory analyses (e.g. CT scanning) provide three-dimensional analysis of microfracture networks. In both cases, the fracture geometries and properties may differ from those expected at depth since any outcrop will, by definition, have experienced a different geological history to a layer in the subsurface.

1.1.1 Impact of Fracture Networks on Bulk Rock Properties: The "TM" Approach

The conventional solution to the problem of modelling subsurface fracture networks, described comprehensively in Nelson (2001) and Mäkel (2007), is to try to characterise the subsurface fracture network as best we can from available data, and then extrapolate this to build a stochastic fracture model. In the simplest stochastic models, the fractures are not represented individually, but rather the fracture network is described statistically by one or more spatially variable properties, such as fracture density, fracture orientation or fracture connectivity (which describes proportion of fractures that are connected to other fractures in the network, versus the proportion of isolated fractures). These properties are then used to directly modify key bulk rock properties, such as porosity and permeability (for flow modelling) or failure strength and elastic moduli (for geomechanical modelling). In the case of permeability, we can define a multiplier for the transmissibility between each gridblock in the static geomodel based on the estimated fracture density in that gridblock, known as a Transmissibility Multiplier or TM. Models of this type are often called "implicit" fracture models; examples are given in Conway et al. (1996) and Gauthier et al. (2000).

The properties that can be measured and extrapolated in these models are usually very limited. From borehole data, we can only measure the density of fractures along a one-dimensional transect through the fractured layer, although we can correct this for borehole orientation (as described in Terzaghi 1965) to obtain the density of fractures along a linear transect perpendicular to the fracture planes (the "mean linear fracture density" or P_{32}). This is equivalent to the total fracture area per unit volume and can thus be used as a proxy for fracture porosity ϕ, if the fracture aperture is known or can be estimated. The boreholes also give us data on fracture orientation. The mean linear fracture density and orientation of each fracture set can then be interpolated between boreholes, with guidance from surface attributes or seismic anisotropy data if available (see e.g. Hart 2006; Singh et al. 2008; Chopra et al.

2009; Donati et al. 2016; Williams et al. 2017). However, the boreholes provide no data on the fracture sizes or connectivity, so we cannot determine the fracture size distribution, or determine how many of the fractures are connected to other fractures and how many are isolated. Nor can we measure the total number of fractures per unit volume (the "volumetric fracture density" or P_{30}). This lack of data makes it difficult to calculate the fracture permeability.

1.1.2 Modelling Fractures Explicitly: The "DFN" Approach

A more sophisticated approach is, therefore, to build an "explicit" stochastic model, comprising a series of geometric objects representing individual fractures, also known as a Discrete Fracture Network model or DFN (Dershowitz et al. 2000; also described by Bowen et al. 2013). Examples of this approach are given in Fonta et al. (2005), Barr et al. (2007), Casabianca et al. (2007), Rawnsley et al. (2007), Rogers et al. (2007), and Souche et al. (2012).

The fractures in the DFN do not correspond directly with actual fractures in the subsurface; instead, we aim to generate a DFN which has the same overall statistical and topological properties as the real fracture network. Key properties that we would like to replicate are the fracture density (both volumetric fracture density and mean linear fracture density), the fracture size distribution and the fracture connectivity.

The fracture density and size distribution can both be characterised using the cumulative fracture density distribution functions $P_{30}(S)$ or $P_{32}(S)$, which give the volumetric or mean linear density of fractures of size S or greater, as a continuous function of S. The fracture size S can be represented by a variety of parameters, but typically we use the fracture radius for circular fractures and the fracture length for layer-bound fractures. Many studies have suggested that natural fracture populations often follow fractal (power-law) cumulative fracture density distributions (e.g. Marrett and Allmendinger 1991, 1992; Westaway 1994; Ouillon et al. 1996; Marrett 1996), although this is by no means universal and other types of cumulative density distributions have also been recorded (e.g. Nicol et al. 1996; Bonnet et al. 2001).

The fracture connectivity is partly dependent on the fracture density and size distribution (Odling 1997; Manzocchi 2002), but also reflects other geometric attributes of the fracture network such as the fracture spacing distribution and fracture anisotropy. Fracture connectivity has a significant (and non-linear) impact on the bulk rock permeability, since fluid can flow much more easily through a network of connected fractures than it can through isolated fractures (Balberg et al. 1991; Berkowitz 1995; Odling et al. 1999).

Unfortunately neither the fracture size distribution nor the fracture connectivity can be measured directly in the subsurface, even from seismic anisotropy data. We must, therefore, rely on measurements from outcrop analogues to condition our stochastic DFN (e.g. De Kejzer et al. 2007; Wennberg et al. 2007; Bisdom et al. 2014). However, both are emergent properties that derive from the growth of the fracture network. An alternative approach is, therefore, to generate the DFN using

a simple kinematic growth model. In this approach, fractures nucleate and grow at arbitrarily specified rates. New fractures are placed at random locations, with orientations corresponding to the orientations of observed fractures, and they interact following simple rules. A DFN generated in this way should exhibit a more realistic fracture size distribution and connectivity than a DFN generated by simply placing fractures of arbitrary size in arbitrary locations. This technique is explored in detail by Bonneau et al. (2013). However, we should bear in mind that these are not dynamic models, i.e. they do not take into account the in situ stress, applied strain or mechanical properties of the fractured layer. Furthermore they do not give direct expressions for the fracture size distribution or connectivity; rather these properties must be calculated empirically from the explicit DFN.

1.1.3 Simulating Fracture Growth: The Mechanical Approach

A third approach is to generate a fracture model by simulating the growth of the fracture network dynamically, based on fundamental geomechanical principles. The main advantage of this approach is that we do not require a priori knowledge of the fracture density, size distribution, connectivity or any of the other statistical parameters that are so difficult to measure in the subsurface, and nor do we need to specify arbitrary fracture nucleation and propagation rates; all we need is a knowledge of the geomechanical properties and deformation history of the fractured layer. This information is usually readily available: the geomechanical properties can be taken from mechanical tests on core samples, wireline log geophysical data or published data for similar lithologies, while the deformation history is often described in studies of the regional geology, or can be inferred from the geometry of larger-scale geological structures (e.g. major faults, folds or diapirs). DFNs generated in this way are more realistic than stochastic DFNs as they are by definition consistent with the geomechanical processes of fracture development, and hence honour the geology of the fractured layer and the larger-scale geological structure. The statistical properties of the fracture network emerge from the simulation, unlike the stochastic modelling approaches where they are specified in advance. These simulations, therefore, generate new information, unlike the stochastic DFNs which simply reflect the (often unreliable) data used to create them.

In most cases, this approach involves using some numerical method to simulate the propagation and interaction of individual fractures dynamically, and thereby generate an explicit DFN. The finite element method is most commonly used, but other methods such as the boundary element method or the discrete element method have also been used for this purpose. This approach has been used to model both natural and artificially induced fracture networks; examples include Advani et al. (1987), Sousa et al. (1993), Huang and Desroches (2004), Bunger et al. (2007),

Adachi et al. (2007), Welch et al. (2015), Guo et al. (2017), Huang et al. (2017), Paluszny et al. (2018), Salimzadeh et al. (2018), Vik et al. (2018).

Dynamic numerical methods allow detailed modelling of local stress anomalies that build up around fractures and faults, as well as local mechanical heterogeneities (e.g. cemented nodules, stiffer or more ductile layers or previous fractures). They can also incorporate thermal or chemical controls on fracturing. They can thus predict the effect that such heterogeneities will have on the ultimate geometry of the fracture network. However, numerical methods are computationally expensive, as the models must have sufficient resolution to capture the local stress anomalies around all of the fractures. It is, therefore, rarely practical to use them to model more than a few fractures in a restricted area. Certainly it is not possible to generate DFNs across large geological structures, which may contain hundreds of thousands of individual fractures, using dynamic numerical methods.

Some studies have, therefore, tried to use simplified (but still geomechanically consistent) rules governing fracture nucleation, propagation and interaction to generate larger DFNs. The best example of this is a series of papers by Olson and co-workers (Olson 1993, 2004; Olson et al. 2001, 2007, 2009), in which they generate explicit DFNs comprising a single set of layer-bound fractures by applying a horizontal strain to a layer containing an initial population of "seed" microfractures, in order to demonstrate the impact of fracture propagation rate on fracture geometry. Their method is based on subcritical fracture propagation theory, which broadly speaking distinguishes slow "subcritical" fracture propagation at low stress, and fast "critical" fracture propagation at high stress (see Sect. 1.2 and Sect. 3.1 for a more detailed review of subcritical fracture propagation theory).

The Olson studies show that, for a single set of parallel fractures, slow subcritical fracture propagation results in a random fracture length and spacing distribution, while intermediate rates of fracture propagation result in a series of regularly-spaced fractures of similar lengths, and fast critical fracture propagation leads to a fractal population of fractures following a power-law length and spacing distribution. It should be noted, however, that their method does not give the cumulative fracture density distribution functions directly; instead, these must be calculated empirically from the explicit DFNs. Therefore, in order to build an implicit fracture model using the Olson method (or other similar methods), it is necessary to first build an explicit DFN which includes all fractures within the size range we are interested in.

If we want to avoid this step and generate the implicit fracture model directly, we must apply the mechanical laws governing fracture propagation to functions describing the fracture population as a whole (i.e. the cumulative fracture density distribution functions), rather than to individual fractures. This is much more complicated as it requires combining statistical and mechanical laws, but it would allow large fracture populations to be modelled very quickly, since the computation time for such calculations will not be proportional to the number of fractures. Such a method could be applied to problems that are very difficult to solve by conventional numerical methods, for example, predicting the variation in size distribution and connectivity of small, centimetre-scale fractures across a large geological structure such as a thrust sheet or a salt dome.

However, the only significant developments in this technique to date are those of Davy and co-workers (Davy et al. 2010, 2013; Maillot et al. 2014), who use it to model the evolution and size distribution of a population of randomly oriented circular fractures in the subsurface. No previous studies have applied this technique to interacting sets of layer-bound fractures, or to examine the effect of fracture stress shadows, and none have derived expressions for the cumulative fracture density distribution functions that take into account fracture deactivation. These steps are necessary if we want to generate realistic implicit fracture models of subsurface layers acting as reservoirs.

1.1.4 Implementation and Practical Application of Fracture Models

As we have previously noted, fracture models (implicit or explicit) are mostly used to enable more accurate simulation of subsurface fluid flow, geomechanical response or both (for example, to predict subsidence or fracture reactivation in response to changes in subsurface fluid pressure). The approach taken will depend on the input data available, the time and resources available for modelling and the accuracy required.

Often, when an approximate solution is all that is required or when time and input data are limited, an implicit fracture model may be sufficient. The simplest implicit models may comprise only one parameter, for example, total mean linear density or total porosity of all fractures, extrapolated laterally across the fractured layer. This property is then scaled to generate a transmissibility multiplier for each gridblock in the static geomodel. Where greater accuracy is required, dual permeability or dual porosity models may be used, in which the flow through the fractures and/or storage of fluid in the fractures is calculated separately from the flow of fluid through the pores of the host rock. Such models require an estimate of the rate of fluid transfer between the pores and the fractures, which is dependent on the fracture spacing and geometry, as well as an estimate of the fracture orientation and connectivity (Cottereau et al. 2010). We must, therefore, build a more sophisticated implicit fracture model that includes these parameters.

To obtain a full statistical description of the fracture population in a specific layer, we must subdivide all the fractures in the layer into sets of approximately parallel fractures. Each set is then described statistically by two cumulative density distribution functions:

- the volumetric fracture density distribution function, $P_{30}(S)$, is the number of fractures of any specified size S or larger, per unit volume of rock and
- the mean linear fracture density distribution function, $P_{32}(S)$, is the total area of all fractures of any specified size or larger within a unit volume of rock,

where the fracture size S is expressed as fracture radius for circular fractures and length for layer-bound fractures. These functions can be defined as continuous analytical functions or as piecewise or discretised functions. The total volumetric fracture density and total mean linear fracture density, P_{30} and P_{32} respectively, are scalar parameters that represent the total number and total area of all fractures within a unit volume of rock, regardless of fracture size. Other statistical measurements, such as the fracture spacing distribution, the total fracture porosity ϕ (i.e. the combined volume of all fractures per unit volume of rock) or the fracture connectivity (i.e. the proportion of fractures that are connected to other fractures) can also be used to enhance the statistical description of the fracture population.

However it is very difficult to obtain all this data for a subsurface fractured layer. The required data can be obtained more easily from an explicit DFN model, so a standard approach is to build a DFN based on the limited statistical data available (e.g. total mean linear fracture density from a wellbore, plus an approximate breakdown into fracture sets), and then extract the more detailed statistical data from the DFN. This is in fact the most common use for explicit DFN models. The drawback with this approach is that a stochastic DFN, generated with no geomechanical constraint, cannot really give out any more information than was used to generate it; any other statistical measurements derived from it must be regarded as extremely unreliable.

More sophisticated flow models are able to simulate fluid flow through individual fractures, as well as fluid exchange between the fractures and the host rock pores (e.g. Sarda et al. 2002; Matthäi et al. 2007; Paluszny et al. 2007). This approach may be particularly useful when fluid flow is concentrated in a few large fractures, which provide pressure communication over long distances. Obviously, an explicit DFN is required to do this.

A major issue when building these more detailed fracture models is the enormous complexity of natural fracture networks in the subsurface. Such networks may contain hundreds of thousands of fractures at a whole range of scales, from mm- or cm-sized microfractures, up to fractures or small faults that are hundreds of metres long and are on the limits of seismic resolution. We must be aware, when comparing (for example) fractures observed in core or borehole image data with fractures mapped in outcrop or inferred from seismic anisotropy analysis, that we may be comparing fractures of very different sizes. Fractures on all these scales may have a significant impact on fluid flow or geomechanical properties in the subsurface: the smaller microfractures may provide additional porosity, or significant local permeability enhancement (or permeability reduction, if they are cemented or filled), or they may reduce the bulk rock strength and elastic moduli; while the largest fractures may provide long-distance fluid flow pathways, or fluid baffles (if they are cemented or filled), and may be reactivated in response to in situ stress or fluid pressure changes. Ideally, therefore, a comprehensive fracture model should include fractures on all scales.

It is impractical, however, to build a DFN across an entire geological structure that includes fractures on all these scales, as it will contain too many fractures to calculate within a reasonable timeframe. One solution is, therefore, to generate a combined implicit and explicit fracture model, in which only the larger fractures