CrossMark

Finite-Volume Discretisations for Flow in Fractured Porous Media

Ivar Stefansson[1] · Inga Berre[1,2] ·
Eirik Keilegavlen[1]

Abstract Over the last decades, finite-volume discretisations for flow in porous media have been extended to handle situations where fractures dominate the flow. Successful discretisations have been based on the discrete fracture-matrix models to yield mass conservative methods capable of explicitly incorporating the impact of fractures and their geometry. When combined with a hybrid-dimensional formulation, two central concerns are the restrictions arising from small cell sizes at fracture intersections and the coupling between fractures and matrix. Focusing on these aspects, we demonstrate how finite-volume methods can be efficiently extended to handle fractures, providing generalisations of previous work. We address the finite-volume methods applying a general hierarchical formulation, facilitating implementation with extensive code reuse and providing a natural framework for coupling of different subdomains. Furthermore, we demonstrate how a Schur complement technique may be used to obtain a robust and versatile method for fracture intersection cell elimination. We investigate the accuracy of the proposed elimination method through a series of numerical simulations in 3D and 2D. The simulations, performed on fractured domains containing permeability heterogeneity and anisotropy, also demonstrate the flexibility of the hierarchical framework.

Keywords Fractures · Hybrid-dimensional discretisation · Finite volumes

1 Introduction

When fractures are present in a porous medium, flow may be totally dominated by the fractured structure and significant complexity is added to the problem of numerical modelling. Two main concerns are complex geometries, which challenge the grid generation, and

✉ Ivar Stefansson
Ivar.Stefansson@uib.no

[1] Department of Mathematics, University of Bergen, Bergen, Norway

[2] Christian Michelsen Research, Bergen, Norway

high contrasts between permeabilities and length scales of matrix and fractures, which challenge the discretisation schemes. The most common modelling approaches may be classified according to the relative importance of the fracture and matrix contribution and to which extent the fracture geometry is honoured. In the one extreme, where large-scale fractures dominate completely, the rock matrix is irrelevant, and the model may be restricted to the fracture network. This gives the discrete fracture network models (DFN), e.g., Long et al. (1982) and Endo and Long (1984). At the other end of the spectrum, the domain may contain a high number of statistically homogeneous fine-scale fractures. In this case, it may be more reasonable to consider their effective impact on the flow instead of considering each fracture explicitly. This leads to equivalent continuum models with upscaled, possibly anisotropic, permeability fields, as described by Sahimi (2011).

Compromises between the two extremes also exist. Both the dual-porosity, the dual-permeability and more general multi-continuum models introduce computational domains of simplified geometry representing fractures and matrix connected through so-called transfer functions. In the dual-porosity models, all flow is assumed to take place through the fractures, and the matrix is only considered to have fluid storage capacity (Barenblatt et al. 1960; Warren and Root 1963). In dual-permeability and multi-continuum models, e.g. Jarvis et al. (1991) and Preuss and Narasimhan (1985), flow is also allowed through disconnected blocks of the matrix, but a connected fracture network is assumed to carry the global flow. This allows for a scale separation where the effect of the smaller fractures is upscaled to the matrix blocks.

In this work, we will use the discrete fracture-matrix model (DFM) as described by Dietrich et al. (2005). The idea is to treat the rock surrounding explicitly represented fractures as a porous medium and possibly account for smaller, non-explicit fractures by an upscaled matrix permeability, as in the equivalent continuum models. The explicitly represented fractures are accounted for individually, as in the DFN methods. In contrast to DFN models, however, flow is allowed between fractures and matrix. Reviews of fracture network models may be found in Berkowitz (2002) and Dietrich et al. (2005).

Motivated by the fractures being very thin compared to typical length scales in the reservoir, a common approach to the explicit fracture modelling introduced by Kiraly (1979) is to model them as lower-dimensional inclusions with an assigned hydraulic aperture, see also Alboin et al. (2000) and Flemisch et al. (2018). Combined with a conforming discretisation, in which fractures are located on boundaries between matrix cells, the result is a hybrid representation of the fractures. This approach has been used in combination with several different discretisation techniques, see Flemisch et al. (2018) and references therein, and is the one we will adhere to.

Finite-volume (FV) methods are among the most widely used discretisations for simulation of flow in porous media. The popularity is explained by the methods' local mass conservation as well as computational efficiency and flexibility. For the simplest version, the two-point flux approximation (TPFA), the computational cost is low and the implementation reasonably straightforward, making it the principal workhorse for a range of porous media flow problems. For more challenging problems, there are the multi-point flux approximations (MPFA).

When it comes to handling of fractures in FV discretisations, the two main approaches are distinguished by whether the grids are restricted by the fracture geometry. Recently, there has been considerable interest in the embedded DFM models (EDFM) as part of a FV discretisation, where the restriction of matrix grid alignment with the fractures is circumvented, see, e.g. Hajibeygi et al. (2011) and Moinfar et al. (2014). While beneficial for the gridding, carefully computed transfer functions are required to account for the flow between fractures and matrix. In this paper, we consider the classical FV DFM models using conforming grids. Such methods include the vertex-centred method by Reichenberger et al. (2006), the cell centre-

based one introduced by Karimi-Fard et al. (2004) and extended by Sandve et al. (2012), and the related method of Ahmed et al. (2015) using transfer functions for the fracture-matrix coupling. Other approaches include the vertex approximate gradient and hybrid FV schemes of Brenner et al. (2016) and the application of hybrid-dimensional FV methods to two-phase flow by Monteagudo and Firoozabadi (2004). We refer to Geiger and Matthäi (2014) for a review of numerical methods based on the DFM model.

Two issues for FV methods introduced by explicit and conforming representation of fractures are the fracture-matrix coupling and the handling of intersecting fractures. We propose two general approaches aimed at enhancing computational efficiency and simplifying the modelling and implementation with respect to the two challenges, at minimal loss of solution quality.

The fracture-matrix coupling is approached through a hierarchical partitioning of the domain into a number of subdomains of different dimension corresponding to matrix, fractures, and fracture intersections following Boon et al. (2017). Discretisation is then performed for each subdomain independently. This leads to a simple implementation, and allows for the application of different discretisation schemes for different domains, at the price of certain restrictions on the coupling between subdomains.

The challenge related to handling of intersecting fractures is due to the comparatively small volumes of cells at the one- and zero-dimensional fracture intersections, as was recognised by Slough et al. (1999) and Granet et al. (2001). These cells impair solution matrix condition numbers and, in the case of transport or multi-phase simulations, time step restrictions. One widely used remedy for FV schemes is the elimination of these cells following the approach introduced by Karimi-Fard et al. (2004) based on the star–delta transformation. The approach has also been extended to MPFA methods by Sandve et al. (2012). While producing satisfactory results in many cases, the procedure does not handle the intersection of fractures of highly varying permeability and some aspects of multi-phase flow, as shown by Stefansson (2016) and Walton et al. (2017).

We suggest a new approach to the elimination of fracture intersections that accounts for the permeability of the intersection cells. By discretising the domain with the intersection cells and then performing a Schur complement reduction to remove the degrees of freedom corresponding to the intersection, we obtain a reduced system possessing information on the permeability of the eliminated cells. Moreover, we show how the original elimination based on the star–delta transformation can be interpreted as the limiting case of the new scheme as the intersection permeability goes to infinity. The new technique also provides a natural way of back-calculating the pressure values at the intersections, as may be desirable e.g. for coupled problems. Finally, the technique is not restricted to FV methods and may in principle also be applied to other discretisations of the problem.

For a discretisation to be suitable for DFM modelling, it should be applicable to unstructured grids corresponding to geometrically complex fracture networks and strong parameter heterogeneity. The ability to handle strong anisotropy in permeability is also necessary in many situations. For a hybrid-dimensional FV method on a conforming grid, the geometrical complexity is mostly dealt with by the constrained meshing of the domain. The handling of parameter heterogeneity and, in particular, the anisotropy, is related to the particular FV method and flux approximation applied. These aspects will also be investigated in the current work.

The transport of heat or chemical species is oftentimes of primary interest in porous media modelling. Because of this, it is sensible to evaluate flow methods not only in terms of pressure fields but also indirectly through transport simulations on the flow fields they produce. By
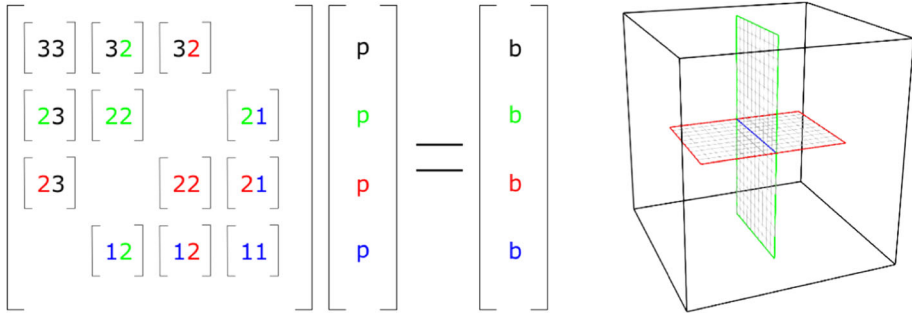
Fig. 1 To the left, a linear system matrix with off-diagonal subdomain coupling blocks and diagonal blocks for the internal subdomain discretisations. To the right, the domain with one, two and one subdomains of dimensions 3, 2 and 1, respectively. The correspondence between subdomains and matrix blocks is indicated through dimension numbers and colours

examining accumulations of tracer, we aim at revealing details and differences that may be critical to the concentration distribution, but almost indiscernible in pressure comparisons.

The implementation used in this work is available in the open-source simulation tool PorePy (Keilegavlen et al. 2017). In addition to open source code, we provide documented run scripts for the results shown in Sect. 3 at www.github.com/pmgbergen/porepy. The exception to the above is the implementation of the star–delta elimination procedure, which is used for comparison with the Schur complement procedure, and for which we rely on the Matlab Reservoir Simulation Toolbox (Lie 2016).

2 Mathematical and Numerical Modelling

As noted in the introduction, the DFM approach to fracture modelling acknowledges that fractures will strongly influence the behaviour of processes in porous media by explicitly modelling selected fractures. Any remaining fractures are considered homogenised and are represented by the matrix through average quantities. This motivates a structured framework for the domain description, as well as governing equations for both matrix and fractures. After the two have been introduced, we describe the FV discretisations used and elaborate on the general methodology for subdomain coupling and the approach to handle fracture intersections.

2.1 Mixed-Dimensional Domain Decomposition

We employ the hierarchical subdomain approach to discrete fracture modelling and associated notation introduced by Boon et al. (2017). The overarching idea is to use the constraints imposed by the fracture geometry to divide the N-dimensional (ND) domain Ω into subdomains Ω_i^d of dimension $d \leq N$. Specifically, for $N = 3$, we extract the matrix without fractures as a 3D subdomain, the fractures as 2D subdomains, fracture plane intersections as 1D subdomains and the intersection points of such lines as 0D subdomains. A similar cascade is defined for $N = 2$. As an example, consider the domain depicted in Fig. 1, which consists of one subdomain of dimension three, two fracture subdomains of dimension two and one intersection subdomain of dimension one. 0D intersections were not included to ensure visual simplicity.

Each subdomain is meshed with the immersed lower-dimensional subdomains acting as constraints. We define a mesh size indicator $h_{\min}$ to be the smallest cell diameter of the mesh of the highest dimension.

Although not required for the general hierarchical approach, we further assume the grids to match between subdomains of subsequent dimensions, so that each face in $\bar{\Omega}_i^{d+1} \cap \Gamma_{ij}^d$ corresponds to exactly one cell in Ω_j^d, where Γ_{ij}^d denotes the interface between subdomains i and j. For most realistic fracture networks this demands unstructured grids, which we limit to simplex grids in this work. For analytical purposes, however, it may also be useful to consider Cartesian grids. All descriptions in the following apply to simplex and Cartesian grids alike.

All subdomains are discretised separately before they are coupled two at a time by discretisations on the interfaces Γ_{ij}. This requires the coupling conditions to be local to the interfaces but makes for a very flexible method in several ways. Firstly, we can use different discretisations in the different subdomains, as will be shown below. Secondly, this makes the conversion from a general porous medium discretisation into a mixed-dimensional DFM method straightforward: An existing mono-dimensional method may be coupled by an appropriate coupling scheme. Thirdly, this also in principle facilitates the use of different physical models in different parts of the domain.

2.2 Governing Equations

We solve the single-phase mixed-dimensional Darcy flow problem

$$\begin{aligned} \mathbf{u} + \mathbf{K}\nabla p &= 0, \\ \nabla \cdot \mathbf{u} - [\![\hat{\mathbf{n}} \cdot \hat{\mathbf{u}}]\!] &= f \end{aligned} \quad \forall\, \Omega_i^d \subset \Omega. \tag{1}$$

Here, $\mathbf{u}$ denotes the flux, $\mathbf{K}$ the permeability tensor, p the pressure, f the sources and sinks, and $\mathbf{n}$ the outward unit normal vector of a subdomain of dimension d. Quantities related to dimension $d+1$ are indicated using the hat notation $\hat{\cdot}$. Denoting the jump operator from the neighbouring higher-dimensional grids by $[\![\cdot]\!]$, the second term in the second equation of (1) enforces coupling of the subdomains. For the highest dimension N, we define $\hat{\mathbf{u}}$ to be zero.

We divide the boundary into parts where we prescribe Dirichlet and Neumann type boundary conditions, respectively:

$$\begin{aligned} p &= p_D && \text{on } \partial\Omega_{D,f}, \\ \mathbf{u} \cdot \mathbf{n} &= u_n && \text{on } \partial\Omega_{N,f}, \end{aligned} \tag{2}$$

using subscript f to indicate the flow problem. Finally, for the inter-dimensional flux on Γ^d, Darcy's law is

$$\hat{\mathbf{n}} \cdot \hat{\mathbf{u}} = -K_n \left(p - \hat{p}|_{\Gamma^d} \right). \tag{3}$$

The notion of the effective normal permeability K_n on the fracture implies the assumption that the normal direction of the fracture is a principal direction of the local permeability.

We also model the advection of a passive tracer concentration T. For a fixed flux field and unitary porosity and density, the concentration is given according to

$$\frac{\partial T}{\partial t} + \nabla \cdot (T\mathbf{u}) - \left[\!\!\left[\hat{T}\hat{\mathbf{n}} \cdot \hat{\mathbf{u}} \right]\!\!\right] = s \quad \text{in } \Omega, \tag{4}$$

with s denoting the tracer sources and sinks. Denoting the tracer discharge over the boundary by J_N and using subscript t to denote the transport problem, the boundary conditions may be written as

$$
\begin{aligned}
T &= T_D \quad \text{on } \partial\Omega_{D,t}, \\
T\mathbf{u}\cdot\mathbf{n} &= J_N \quad \text{on } \partial\Omega_{N,t}.
\end{aligned}
\tag{5}
$$

2.3 Finite-Volume Discretisations

We explore the subdomain coupling and intersection cell elimination in a FV framework. To obtain the formulation, we start out from the integral form of the mass balance Eq. (1) over a cell c_i. Using Gauss' divergence theorem on the left-hand side, we obtain

$$
\int_{\delta c_i \backslash \Gamma} \mathbf{u}\cdot\mathbf{n}\,\mathrm{d}S - \int_{\delta c_i \cap \Gamma} [\![\hat{\mathbf{n}}\cdot\hat{\mathbf{u}}]\!]\,\mathrm{d}S = \int_{c_i} f\mathrm{d}V,
\tag{6}
$$

where $\mathbf{n}$ is the outward normal vector on the cell's boundary δc_i. Assuming the right-hand side source term to be known, and postponing the treatment of the coupling term (the second term on the left-hand side) to Sect. 2.4, we proceed to the approximation of the flux in the interior of the subdomain. Noting its dependency on the pressure as described by Eq. (1), we aim at a relationship between the flux over a face S_{ij} between cells i and j in terms of the pressures p_k at the centres of the surrounding cells, i.e.

$$
\int_{S_{ij}} \mathbf{u}\cdot\mathbf{n}dS = -\int_{S_{ij}} \nabla p\cdot\mathbf{K}\cdot\mathbf{n}dS \approx \sum_k t_k p_k.
\tag{7}
$$

The weights, or transmissibilities, t_k incorporate geometry and permeability, and may be computed in different manners leading to different FV discretisations. In this work, we limit ourselves to two versions: the two-point and multi-point flux approximations described in the following subsections and in "Appendix". Once computed, the transmissibilities are assembled into the discretisation matrix $\mathbf{A}$, leading to a global system of equations of the form

$$
\mathbf{A}\mathbf{p} = \mathbf{b},
\tag{8}
$$

where $\mathbf{b}$ accounts for boundary condition contributions and source and sink terms, and $\mathbf{p}$ is the vector of cell centre pressures. Note that even if we only solve for pressures, the fluxes may be readily back-calculated using the information in $\mathbf{A}$, which indeed is nothing else than a discretisation and summation of the fluxes in the domain.

2.3.1 Two-Point Flux Approximation

To prevent the computational cost from blowing up when the number of discretisation cells increases, the weights, t_k, should be nonzero only locally around the face. In the most extreme case, one can choose to approximate the gradient using the pressure values from the two cells immediately next to the face to compute the discharge from c_i to c_j:

$$
u_{ij} = \int_{S_{ij}} \mathbf{u}\cdot\mathbf{n}\,\mathrm{d}S = -\int_{S_{ij}} \nabla p\cdot\mathbf{K}\cdot\mathbf{n}\,\mathrm{d}S \approx -t_{ij}\left(p_j - p_i\right).
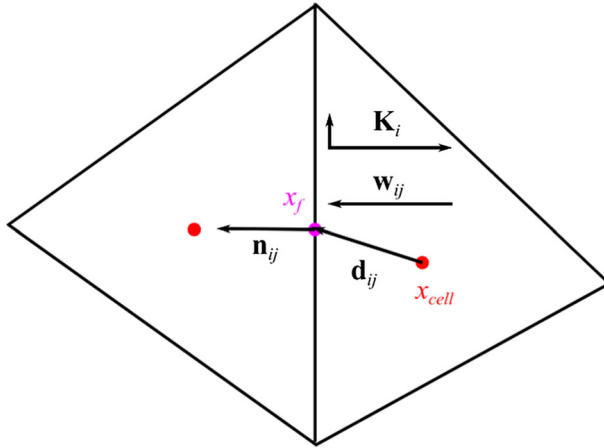\tag{9}
$$

Fig. 2 Geometrical points and entities for TPFA transmissibility computations between the left cell i and the right cell j

In this case, the face transmissibility is computed as the harmonic average of the half transmissibilities:

$$t_{ij} = \frac{\alpha_i \alpha_j}{\alpha_i + \alpha_j}, \tag{10}$$

which in turn are given as

$$\alpha_i = \frac{A_{ij} \mathbf{n}_{ij} \cdot \mathbf{K}_i}{\mathbf{d}_{ij} \cdot \mathbf{d}_{ij}} \cdot \mathbf{d}_{ij}. \tag{11}$$

Here, the aperture-weighted face area is denoted by A_{ij} and the unit normal vector $\mathbf{n}_{ij}$ on the face points outward from c_i. The permeability tensor and the distance vector between the cell centre x_{cell} and the face centre x_f, $\mathbf{d}_{ij} = x_f - x_{\text{cell}}$, also belong to c_i. In Fig. 2, the quantities needed for calculating the half transmissibility α_i of the right cell and centre face are shown. The TPFA is described in more detail in Aziz and Settari (1979) and is the standard discretisation in commercial reservoir simulation.

2.3.2 Multi-Point Flux Approximation

While straightforward to derive and implement and computationally very efficient due to matrix sparsity, the TPFA scheme is only consistent when the grid is aligned with the principal directions of the permeability tensor. Defining the normal component of the permeability as $\mathbf{w}_{ij} = \mathbf{K}_i \cdot \mathbf{n}_{ij}$, the requirement is that $\mathbf{w}_{ij}$ and $\mathbf{d}_{ij}$ be aligned for all cell-face combinations. If this is not the case, as in Fig. 2, a more sophisticated flux approximation is called for.

To construct the multi-point flux approximation, a somewhat more involved computation of face transmissibility is required. A description of the method can be found in "Appendix".

2.3.3 Transport Discretisation

The transport equation is also integrated over each cell to yield a FV formulation. The upwind term is integrated by parts and then discretised with upwind concentration evaluation. For the advection between c_i and c_j we use

$$T_{upw} = \begin{cases} T_i & \text{if} \quad u_{ij} > 0, \\ T_j & \text{if} \quad u_{ij} < 0. \end{cases} \tag{12}$$

Assuming stationary fluxes, sinks and sources, implicit Euler temporal discretisation is applied, yielding for time step n

$$V_i \frac{T_i^n - T_i^{n-1}}{\Delta t} - \sum_j A_{ij} T_{upw}^n u_{ij} = s_i \tag{13}$$

for the concentration T_i at the centre of c_i of volume V_i.

2.4 Subdomain Coupling

We now consider the coupling of the flux discretisations between different subdomains. In the discretisations of the individual subdomains, we impose Neumann conditions for all boundaries internal to the full domain. This allows us to couple the subdomains by simply discretising the flux over all faces of each internal boundary.

We consider the coupling of two subdomains Ω_i^{d+1} and Ω_j^d over the common boundary Γ_{ij}. Letting the subscript indexes i and j also indicate cells from the respective subdomains, we discretise the flux from $c_i \in \Omega_i^{d+1}$ to $c_j \in \Omega_j^d$ according to Karimi-Fard et al. (2004):

$$\hat{\mathbf{n}}_{ij} \cdot \hat{\mathbf{u}}_{ij} = -t_{ij} \left(p_j - \hat{p}_i \right), \tag{14}$$

with $t_{ij} = \frac{\hat{\alpha}_i \alpha_j}{\hat{\alpha}_i + \alpha_j}$. For the half-face transmissibility calculation, the lower-dimensional subdomain is extended by half an aperture a in the normal direction. Hence, on the fracture side of the interface we obtain

$$\alpha_j = \frac{\mathbf{n}_{ji} \cdot \mathbf{K}_j}{a/2} A_{ij}, \tag{15}$$

which equals K_n from Eq. (3). On the matrix side, $\hat{\alpha}_i$ is computed according to Eq. (11).

In the DFM approach, the fracture domains are accounted for twice both in cell volumes and distance computations, and so relies on $a \ll h_{min}$. The distance inconsistency can be corrected by using the adjusted distance vector

$$\mathbf{d}_{ij}^* = \mathbf{d}_{ij} \left(1 - \frac{a}{2 |\mathbf{d}_{ij}|} \right); \tag{16}$$

see e.g. Sandve et al. (2012).

As the flux over the face may be expressed in terms of one pressure value from each of the two subdomains only, the couplings are entirely local and can be computed independently. The coupling conditions are arranged in four blocks and assembled into the global solution matrix, as illustrated in Fig. 1. Thus, the diagonal blocks consist of the sum of the internal discretisation and the contribution from the subdomain coupling. The off-diagonal blocks will be nonzero only where the corresponding subdomains interact.

For the transport problem, upwind coupling terms are calculated just as in the internal discretisations. For the area, the face area of the higher-dimensional cell is used.
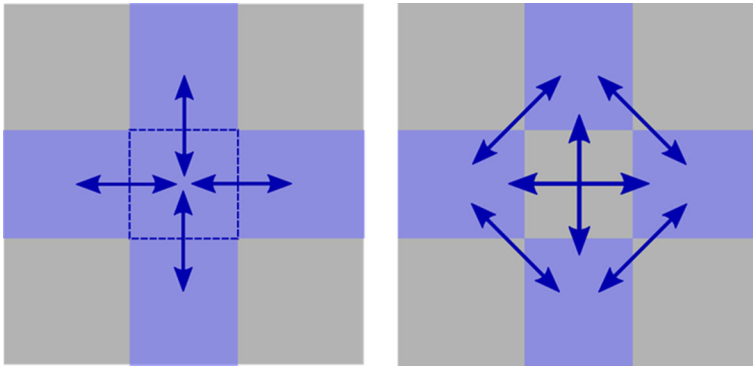
Fig. 3 Conceptual sketch of the fluxes at a fracture intersection with and without an intersection cell, left and right, respectively

2.5 Intersection Cell Removal

In realistic fractured porous domains, the fracture aperture is commonly several orders of magnitude smaller than the domain length scale. This implies high contrasts in discretisation cell size, as the smallest cell volumes in dimension d typically scale according to

$$V\left(c\right) \propto h_{\min}^d a^{N-d}. \tag{17}$$

Strong cell size variation may affect the solution quality and efficiency of an iterative linear solver through high condition number $C = cond\left(\mathbf{A}\right)$ of the matrix $\mathbf{A}$ from Eq. (8). As pointed out and demonstrated by Karimi-Fard et al. (2004) and Walton et al. (2017), removing the one- and zero-dimensional cells at fracture intersections may improve the condition number significantly, while also relaxing time step restrictions for transport and multi-phase flow simulations.

Karimi-Fard et al. (2004) proposed an elimination procedure for the TPFA, known as a star–delta transformation, where direct connections are introduced between the higher-dimensional cells. These connections are obtained by adding n_f fluxes between all the n higher-dimensional cells meeting at the particular intersection, see Fig. 3, where $n = 4$ leads to six cell pair combinations and hence $n_f = 6$ new fluxes. The transmissibility for each cell pair at the intersection is computed as

$$T_{ij} = \frac{\alpha_i \alpha_j}{\sum_k^n \alpha_k}, \tag{18}$$

where α_k is computed as in Eq. 11. This popular transformation technique was extended to the MPFA method by Sandve et al. (2012).

By ignoring the intersection cell, the star–delta elimination in effect assigns infinite permeability to the intersection. When the permeability of the intersection is in fact low, which could be the case if a blocking fracture crosses a permeable one, this may lead to significant overestimation of fracture connectivity, as demonstrated in Sect. 3.1 and further studied by Stefansson (2016).

To remedy this shortcoming, we introduce a new technique for intersection cell removal based on the Schur complement; see, for example, Zhang (2005). We first discretise the problem including intersection cells and then identify the degrees of freedom to be kept and those to be eliminated. Typically, the degrees of freedom to be eliminated are all those

corresponding to intersections of some dimension(s), i.e. all cells satisfying $c \in \Omega^m$, for $m = 0$ or $m \in \{0, 1\}$. Associating the matrix rows and columns corresponding to the two sets with subscripts k and e, respectively, we obtain the reduced system matrix and right-hand side

$$
\begin{aligned}
\mathbf{A}_r &= \mathbf{A}_{kk} - \mathbf{A}_{ke}\mathbf{A}_{ee}^{-1}\mathbf{A}_{ek}, \\
\mathbf{b}_r &= \mathbf{b}_k - \mathbf{A}_{ke}\mathbf{A}_{ee}^{-1}\mathbf{b}_e.
\end{aligned}
\tag{19}
$$

This requires the inversion of the matrix $\mathbf{A}_{ee}$. With the above-described set of cells to be eliminated, the size of this matrix is well below the size of the full system, and the computational cost of the inversion is negligible. To demonstrate that $\mathbf{A}_{ee}$ is invertible, we note that it is the internal FV discretisation on Ω^m. As described above, we first discretise a Neumann problem on these subdomains to be eliminated. Then, the TPFA coupling contributes with the addition of weights to all diagonal entries (and only these), yielding strict diagonal dominance and invertibility.

Solving the system $\mathbf{A}_r\mathbf{p}_r = \mathbf{b}_r$ for the kept or "reduced" pressures, we accurately account for intersections of fractures of arbitrary permeability ratios while avoiding the minute volumes of the removed cells. In particular, the Schur complement avoids the restriction of having to assume a certain intersection permeability in the design of the elimination scheme, leaving the modeller free to assign any permeability at the intersections. The suggested procedure accounts for the assigned intersection permeability in the solution, yielding full modelling flexibility.

There are no limitations on the number of intersecting fractures or their geometry as long as the validity of the underlying DFM model is ensured. As demonstrated in Sect. 3.1.3, the technique yields accurate pressure solutions also in cases where 1D cells are eliminated. Further, the Schur complement reduces to the above-mentioned star–delta transformation in the limiting cases of infinite normal permeability and zero tangential permeability of the eliminated cells. Note that this elimination technique is not particular to FV discretisations, but the restriction that $\mathbf{A}_{ee}$ must be invertible must be satisfied.

While this paper deals exclusively with incompressible single-phase flow, it is possible to extend the elimination technique to more complicated, possibly nonlinear, physical models such as compressible flow. The approach deals with the calculation of transmissibilities, which are only affected by the intrinsic permeabilities and geometry of the cells involved through $\mathbf{K}$, $\mathbf{n}$, $\mathbf{d}$ and A, see Eq. (11). The transmissibilities may be calculated for the intersections once at the start of the simulation, and then reused throughout linearisation or time loops, as is standard practice in reservoir simulation.

If fluxes between the grid cells are to be computed, there are two ways to proceed after the elimination. The first option is to back-compute the pressure in the intersection cells and proceeds as if the cells were never removed. While straightforward using the Schur complement, this is in many situations somewhat counterintuitive. If transport simulations are performed using the resulting flux field, the original time step restrictions must be honoured.

Not reintroducing the lower-dimensional cells calls for flux computations between the higher-dimensional neighbours directly, see Fig. 3. This may be done directly from the reduced discretisation, in the same manner as described in Sect. 2.3 for the internal subdomain fluxes. This option is not entirely unproblematic, especially if 1D cells are eliminated in a 3D domain. Flow along the intersection will be distributed more evenly between the intersecting fractures compared to a non-eliminated or back-calculated solution, leading to additional "leakage" from the main flow path to less permeable regions, as shown by Sandve et al. (2012).

The case of multi-phase flow is even more challenging due to the presence of counter-current flows and gravity effects. We offer no satisfactory solution to this problem through the Schur complement procedure.

2.6 Anisotropy

As mentioned above, anisotropic permeability calls for more sophisticated discretisations such as the MPFA. Using a DFM model, the background matrix permeability may also include the upscaled contribution from fine-scale fractures, which may enhance the matrix anisotropy.

Using TPFA for the fracture-matrix coupling might impair the solution quality in the case of matrix anisotropy as the half-face transmissibility on the matrix side, $\hat{\alpha}_i$, will be computed inconsistently, cf. Sect. 2.4. This is investigated in Sect. 3.2.

As shown by Mokhtari et al. (2015) and Tan et al. (2017), fracture planes may also exhibit intrinsic permeability anisotropy; that is, the permeability depends on the direction considered in the fracture plane itself. Given that the fracture network and its connectivity oftentimes dominate the flow through the domain, it is conceivable that fracture anisotropy may substantially influence the overall behaviour of the system. We investigate the impact of fracture anisotropy in Sect. 3.3.

3 Results

The results are presented in the form of plots of pressure and tracer distributions and plots of locally monitored tracer concentrations throughout the transport simulation time interval. We report errors of both pressure and final tracer fields as compared to various reference solutions. The errors are computed either over the entire domain or over specified subdomains, always in the discrete L^2 norm.

The numbering of the test cases corresponds to the subsections, so that Case 1.1 appears in Sect. 3.1.1 etc. We reiterate that source code is available at https://github.com/pmgbergen/porepy/tree/develop/examples/papers/arXiv_1712_08479, as is additional visualisation material intended as a supplement to the test case descriptions, figures and tables provided in this section.

3.1 Intersection Elimination

This section is devoted to an examination of the Schur complement procedure for elimination of fracture intersections introduced in Sect. 2.5. We compare the pressure solution to a solution based on a reference discretisation where the intersection cells have been kept and a solution where the star–delta transformation is used for elimination of intersection cells. In addition to assessing the solution quality through pressure and tracer field comparisons, we report the condition numbers of the respective solution matrices. An indication on the improvement offered by the eliminations is provided by the ratio between the condition number of the full solution and the eliminated, which we denote by R_C.
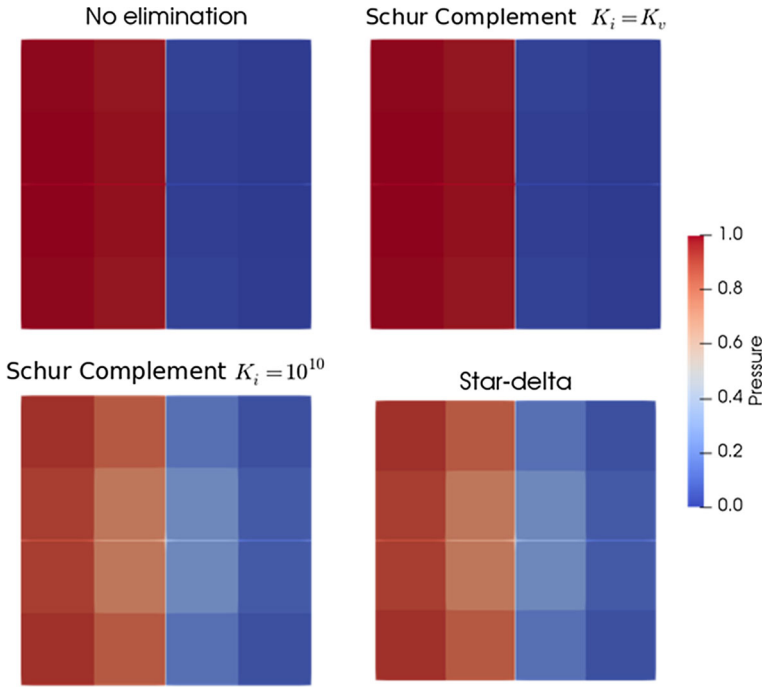
Fig. 4 Pressure fields of Case 1.1, $K_h = 10^3$ and $K_v = 10^{-3}$. The case is designed so that the intersection permeability, K_i, equals that of the vertical fracture. The bottom left Schur complement solution with $K_i = 10^{10}$ is included for comparison with the star–delta solution

3.1.1 Elimination of a 0D Intersection

We start by a systematic investigation of the elimination procedures in a synthetic two-dimensional test case containing a single intersection. We incrementally change the fracture permeabilities to reveal when the Schur complement and star–delta procedures differ.

The test case geometry is shown in Fig. 4. The domain contains two orthogonal fractures, one normal to and one parallel to the pressure gradient, both having an aperture $a = 10^{-2}$. The matrix permeability is set to unity, whereas the permeabilities of the fractures vary between 10^{-3} and 10^3 to yield a range of different ratios between the permeability of the horizontal (gradient-aligned) fracture, $\boldsymbol{K} = K_h \boldsymbol{I}$, and that of the vertical fracture, $\boldsymbol{K} = K_v \boldsymbol{I}$. The permeability of the 0D intersection cell is inherited from the vertical fracture. Homogeneous Neumann conditions are imposed on the horizontal boundaries, and the Dirichlet condition $p_D = 1 - x$ is imposed on the vertical boundaries.

In Fig. 4, the pressure plots for the permeability ratio yielding the largest difference between the elimination techniques are shown. The pressure jump across the vertical fracture is markedly higher for the reference and Schur complement solution than for the star–delta one, where the blocking effect at the intersection is lost. We include the pressure plot of a simulation where the intersection cell permeability is 10^{10}. The solution is practically identical to the results using the star–delta transformation, which illustrates how the latter is the limiting case of high intersection permeability, as discussed in Sect. 2.5.

A detailed report of the pressure errors and condition numbers is given in Fig. 5. We see machine precision agreement between the solution for which the Schur complement technique
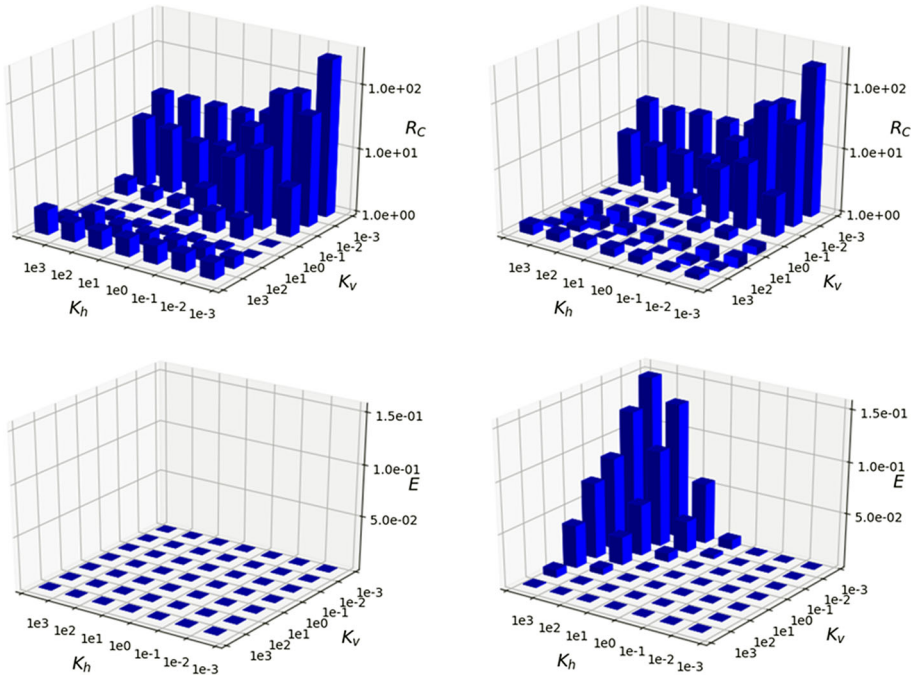
Fig. 5 Condition number ratios R_C (top) and pressure errors (bottom) for Case 1.1 for different values of the permeabilities of the horizontal and vertical fracture. Schur complement and star–delta elimination were used for the subfigures to the left and to the right, respectively

was applied and the non-reduced solution. The solution where the star–delta transformation was used, however, breaks down when the ∇p-aligned fracture is permeable and the other one is not. Under these conditions, the horizontal fracture strongly influences the solution, but the blocking effect of the vertical fracture is not captured by the star–delta transformation. The two fracture elimination schemes bring about significant and very similar condition number improvements, up to more than a factor of 10^2. Note that while the magnitude of R_c in general depends on the aperture and permeability values in question, Fig. 5 indicates that the improvement is larger when low-permeable intersections are eliminated.

3.1.2 Elimination of Multiple 0D Intersections

We now consider a 2D test case presented in the benchmark study by Flemisch et al. (2018). The case considers conducting fractures intersected by blocking fractures. The study showed that applying the star–delta intersection cell elimination decreased solution quality. We here investigate the Schur complement elimination procedure for that particular test case.

The test case contains ten fractures of two different permeabilities, 10^{-4} and 10^4, in a matrix of unit permeability. The intersection permeabilities are set to the harmonic average of the crossing fractures, making them in effect impermeable. For further details on the set-up, we refer to (Flemisch et al. 2018), where this benchmark test case appears in Sect. 4.3.2.

We display two of the solutions in Fig. 6 and in Table 1 we list the errors and the properties of the discretisation matrices, specifically the densities and condition numbers. The plots indicate that using the star–delta technique, significant flow passes through the blocking
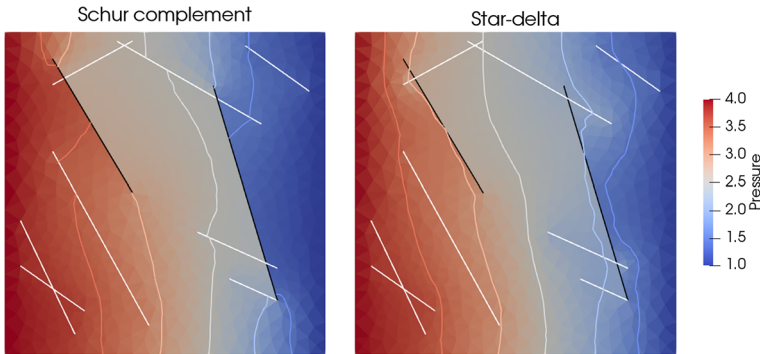
Fig. 6 The pressure distributions for Case 1.2 obtained using the two elimination techniques with MPFA discretisation. Permeable and impermeable fractures are indicated in white and black, respectively. Isolines are shown for pressure values of 1.5, 2, 2.5, 3 and 3.5

Table 1 Errors and matrix characteristics for Case 1.2

Method	Fracture error	Matrix error	Matrix density	Condition number
TPFA, no elimination	1.3×10^{-2}	1.3×10^{-2}	2.8×10^{-3}	2.0×10^{5}
TPFA, Schur complement	1.4×10^{-2}	1.4×10^{-2}	2.8×10^{-3}	2.1×10^{4}
TPFA, star–delta	5.1×10^{-2}	6.7×10^{-2}	2.8×10^{-3}	3.1×10^{4}
MPFA, Schur complement	1.2×10^{-2}	1.4×10^{-2}	8.0×10^{-3}	2.0×10^{4}
MPFA, star–delta	5.1×10^{-2}	6.7×10^{-2}	8.5×10^{-3}	3.1×10^{4}

fractures where they intersect with the conducting ones. The plotted isolines also highlight that some of the blocking effect of the low-permeable fractures is missed in the star–delta solution. The errors of the solutions obtained using the Schur complement elimination are very similar to those obtained without elimination and considerably lower than those where the star–delta elimination was applied. This indicates that the Schur complement technique does not suffer from the restrictions of the star–delta transformation related to blocking intersections. Moreover, the condition number of the Schur complement matrices is slightly lower than for the two star–delta matrices. This means that we obtain the full benefits of the star–delta technique while retaining solution quality.

3.1.3 Elimination of a 1D Intersection

We now turn to a 3D case where the 1D intersection is eliminated. In general, this is inherently different from elimination of 0D intersections, as the tangential flow along the intersection is affected. In particular, flux calculations are less straightforward, as discussed in Sect. 2.5.

The unit cube domain of unitary permeability contains two fractures lying in the xy- and yz-plane, respectively. Both have aperture 10^{-6}, and the isotropic fracture permeabilities are 10^{6} and 10^{-6}, respectively. The boundary conditions are homogeneous Neumann, except for the pressure conditions $p_D(x=0)=1$ and $p_D(x=1)=0$. Similarly, tracer concentration

Table 2 Errors of the eliminated solutions compared to the full solution (with 1D cells) for Case 1.3

Elimination	Pressure error	Tracer error	Condition number improvement
Schur complement	2.2×10^{-15}	1.5×10^{-12}	6.0×10^3
Star–delta	3.2×10^{-2}	1.2×10^{-1}	4.1×10^3
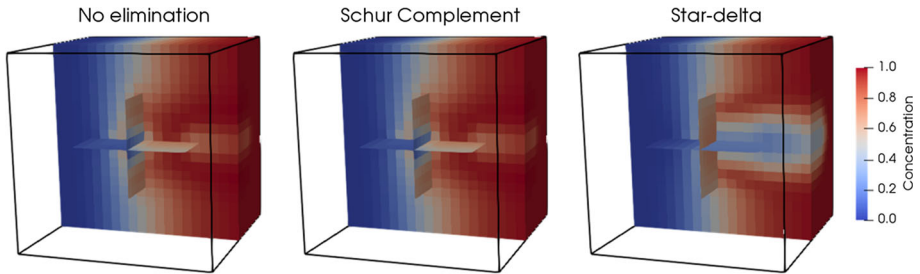
Fig. 7 Case 1.3 tracer concentrations at time $t = 0.5$
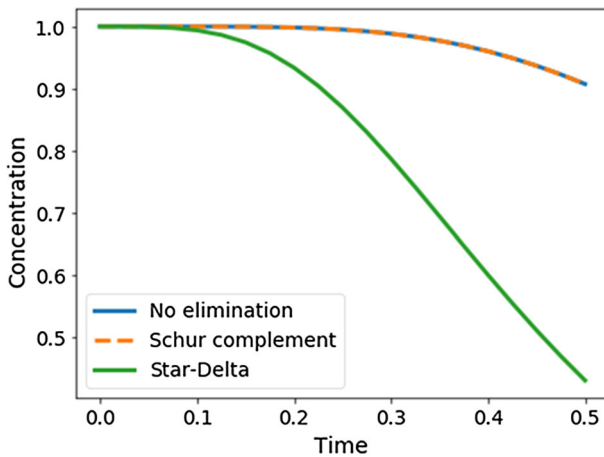
Fig. 8 Monitored tracer concentration at the centre of the outflow boundary throughout the transport simulation for Case 1.3

values of zero are prescribed at all external boundaries, whereas the initial tracer concentration is one throughout the domain. The transport simulation runs to $t = 0.5$.

We compare three solutions with TPFA in the entire domain: one where no elimination of intersection cells is performed, one where the Schur complement elimination is used and one using the star–delta elimination. The errors of the eliminated solutions compared to the full one are listed in Table 2. The tracer concentrations shown in Fig. 7 show qualitative differences: If the blocking nature of the intersection is honoured, the tracer is largely forced around the vertical fracture, whereas it shoots through in the case of the star–delta elimination. The perfect agreement between the solution without elimination and the Schur complement, and difference to the star–delta, are highlighted by the concentration time series in Fig. 8, recorded at the centre of the outflow boundary.

For the solutions where the 1D intersection cells were eliminated, the fluxes at the intersections are computed directly between the neighbouring 2D cells, as described in Sect. 2.5 and indicated to the right in Fig. 3. The very low error of the Schur complement solution illustrates that this procedure is adequate also for flux computations provided that the flow along the fracture is negligible.

3.2 Subdomain Coupling with Anisotropy in Matrix Permeability

Turning back to the unit square domain, we investigate the TPFA coupling's behaviour in face of anisotropy in the higher dimension by considering a single-fracture case. In particular, we investigate the impact of non-alignment of the distance vector $\hat{\mathbf{d}}_{ij}$ and normal permeability component $\hat{\mathbf{w}}_{ij}$ on the matrix side of the subdomain interface, with subscript indexes as in Sect. 2.4. This means that we apply the TPFA to a problem for which it is known to be inconsistent, but only in a small part of the domain.

We include one highly permeable ($\mathbf{K} = 10^4\ \mathbf{I}$) horizontal fracture of aperture 10^{-3} at $y = 0.5$. No-flow conditions are set for both pressure and tracer except at two diametrically opposing corners of the domain, where Dirichlet conditions are enforced. A series of simulations are performed with varying anisotropy ratio K_{max}/K_{min} and angle ($10°$, $30°$ and $60°$ relative to the coordinate axes) in the matrix, i.e.

$$\mathbf{K} = \mathbf{R}\,(\theta) \begin{bmatrix} K_{\max} & 0 \\ 0 & K_{\min} \end{bmatrix},$$

where $\mathbf{R}$ is the rotational matrix and θ the angle of rotation relative to the x axis. For the variation in matrix anisotropy ratio, we kept $K_{\max}$ fixed at 1 while decreasing $K_{\min}$.

We compare the reduced model using the TPFA coupling and MPFA for internal discretisations to an equi-dimensional model using MPFA in the entire domain in terms of both pressure and tracer errors. The reference solution is equi-dimensional and computed on a fine regular grid (256×256 cells), see Fig. 9. We refine the coarse solutions from 4×4 to 32×32 cells using Cartesian meshes. Because the matrix cell size $h_{\min}$ is only one order of magnitude larger than the aperture for the finest grid, we use the aperture corrected distance vector introduced in Sect. 2.4, thus avoiding that the inconsistency of the hybrid-dimensional approach obscures the convergence comparison.

In Fig. 10, we show the results for the $30°$ rotation (the two other angles yielded qualitatively equivalent results). The results indicate that the TPFA coupling performs very acceptably. There is, however, one significant limitation to the conclusion. For the anisotropy, we consider a permeability ratio range of $1 \le K_{\max}/K_{\min} \le 6$. For larger values, the results showed oscillations in the matrix pressure away from the fracture. We attribute this to monotonicity issues for the MPFA unrelated to the presence of fractures, see Nordbotten et al. (2007). We expect that the errors related to the TPFA coupling scheme are enhanced for larger anisotropy ratios, although explicitly accounting for the large fractures relieves us from the most severe anisotropies characteristic of effective permeabilities in the context of fracture upscaling.

3.3 Anisotropy in Fracture Permeability

A 3D test case investigating the effect of permeability anisotropy in a 2D fracture is now studied. We explore whether the anisotropy calls for the use of MPFA or if TPFA is admissible. We again consider a case with limited geometric complexity and assign isotropic permeability in the matrix. To further isolate the effect of fracture anisotropy, we use a Cartesian grid for
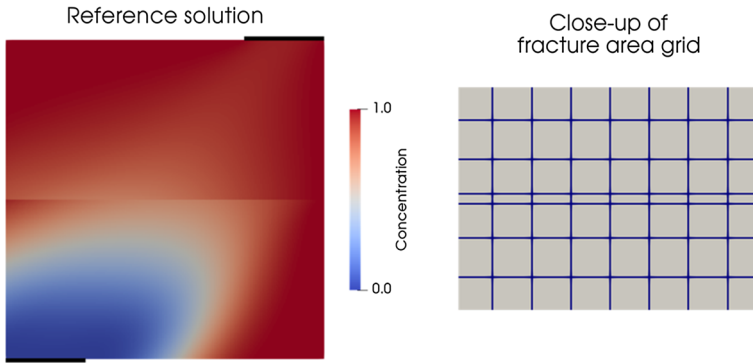
Fig. 9 To the left, reference tracer solution for $\theta = 30°$ and $K_{max}/K_{min} = 4$. To the right, a close-up of the equi-dimensional reference grid. The fracture is represented by the rectangular cells along the middle, whereas normal Cartesian cells were used for the matrix. The in- and outlet part of the boundary is indicated by the black lines at the bottom and top of the domain, respectively
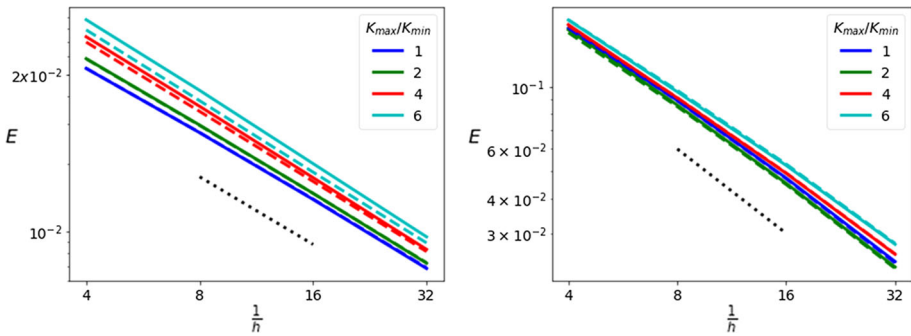

Fig. 10 Pressure and tracer error plots for the $\theta = 30°$ version of Case 2, left and right, respectively. Line colours represent different permeability ratios. Solid lines correspond to the reduced model, dashed lines to the equi-dimensional MPFA. Dotted black lines indicate a linear slope

our unit cube domain, ensuring alignment of face normals and distance vectors in the matrix, cf. Sect. 2.3.1.

The domain contains a single, horizontal fracture of aperture 10^{-3} cutting the domain at $z = 0.5$. No-flow conditions are imposed on all but a part of the top and bottom boundaries, for which we set $p_D = 0$ and $p_D = 1$, respectively. With unit permeability in the matrix and a fracture permeability tensor with principal permeabilities of 10^3, $1/3 \times 10^3$ and 10^3 rotated $45°$ in the fracture plane:

$$\mathbf{K} = \mathbf{R}\left(45°\right) \begin{bmatrix} 10^3 & 0 & 0 \\ 0 & 1/3 \cdot 10^3 & 0 \\ 0 & 0 & 10^3 \end{bmatrix} = \begin{bmatrix} 2/3 \cdot 10^3 & -1/3 \cdot 10^3 & 0 \\ -1/3 \cdot 10^3 & 2/3 \cdot 10^3 & 0 \\ 0 & 0 & 10^3 \end{bmatrix},$$

we obtain a system strongly influenced by the fracture flow. Again, the initial and boundary condition for the tracer is $T\left(t = 0\right) = 1$ in Ω and $T = 0$ on $\delta\Omega$. The transport simulation runs to $t = 30$.

In the following, we use the solution where MPFA was used for the flow problem as a reference. Comparison of the pressure solutions obtained using TPFA and MPFA yields the relative differences reported in Table 3, revealing only slight differences. The tracer solutions

Table 3 Difference to the MPFA reference in solutions and discretisation times for Case 3

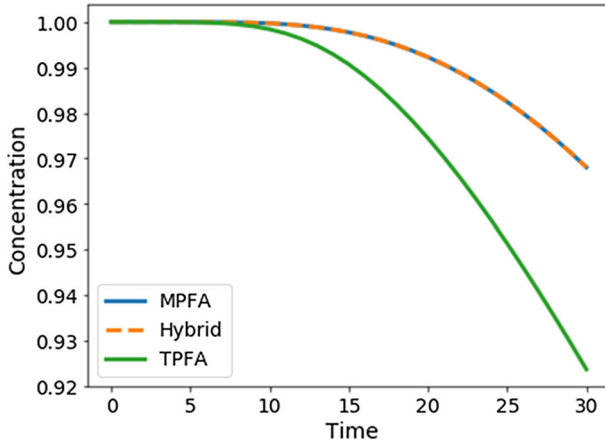| Flux discretisation	Subdomain	Difference to MPFA pressure	Difference to MPFA tracer	Discretisation time, relative to MPFA
TPFA	3D matrix	5.0×10^{-4}	5.8×10^{-2}	3.3×10^{-2}
	2D fracture	8.4×10^{-4}	1.6×10^{-1}	
Hybrid	3D matrix	1.8×10^{-10}	8.6×10^{-8}	8.3×10^{-2}
	2D fracture	2.7×10^{-10}	2.4×10^{-7}	

Fig. 11 Time series of tracer monitored in the cell of the outflow boundary, Case 3
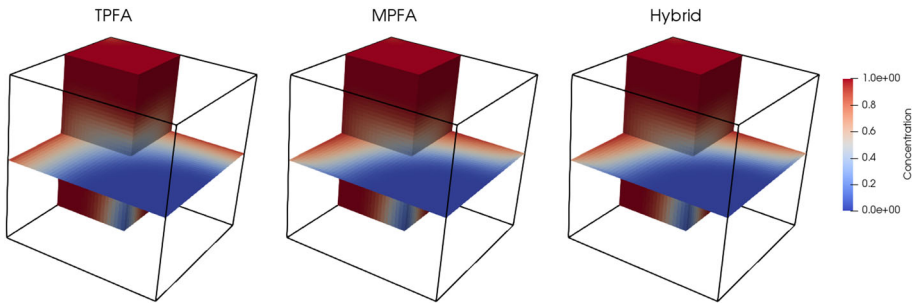

Fig. 12 Tracer distributions at $t = 30$ for Case 3. There is injection in the top, back corner and production takes place at the opposite corner. The 2D fracture plane is left uncut for visualisation purposes

are compared both in terms of L^2 errors and a time series of the concentration at the outflow boundary in Table 3 and Fig. 11. By contrast, these reveal disparities between the flow fields of significant magnitude, see also Fig. 12.

As mentioned in Sect. 1, we may also combine different discretisations on the different subdomains. Exploiting that the permeability is anisotropic only in the fracture, we construct a hybrid discretisation with TPFA in the matrix and MPFA in the fracture, coupled together with the TPFA coupling. As evidenced from the tracer results, this discretisation is sufficient to capture the anisotropy. We also report the CPU times obtained with our implementation
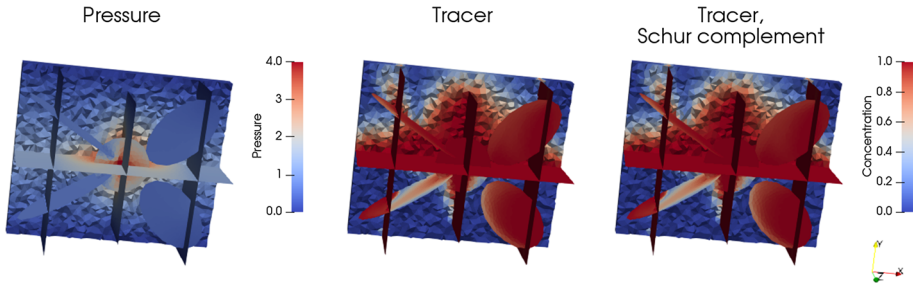
Fig. 13 Distribution of pressure (left) and tracer at $t = 2$ for a solution obtained without intersection cell removal (middle) and one where the Schur complement procedure was applied (right) for Case 4. The injection well placement may be inferred from the pressure plot

for the discretisation of the problem in Table 3. They indicate that the computational cost of the hybrid discretisation is substantially reduced compared to the MPFA, although somewhat larger than for the TPFA. Note that the times reported are those used for the assembly of the discretisation matrix $\mathbf{A}$. The time spent on inverting the linear system was left out to avoid the complexity related to different linear solvers.

3.4 Matrix Heterogeneity

In the final test case, we provide an example of how the hierarchical subdomain approach and coupling flexibility may be advantageously applied. We investigate a 3D case displaying both strong fracture influence and matrix flow with a fracture geometry containing several different intersection configurations.

We choose a mostly conductive fracture network which is symmetric about the xy-plane, shown in Fig. 13. All fractures have aperture $a = 10^{-6}$, and all but one have permeability $\mathbf{K} = 10^5 \mathbf{I}$. The circular fracture at the centre of the domain is blocking, with $\mathbf{K} = 10^{-5}\,\mathbf{I}$. Intersection permeabilities are inherited from the least permeable of the intersecting fractures. The network is surrounded by a matrix of permeability $\mathbf{K} = 10^{-2}\mathbf{I}$ in the upper half and $\mathbf{K} = 10^{-3}\mathbf{I}$ in the lower half. No-flow conditions are imposed on the vertical boundaries and $p, T = 0$ at the horizontal ones, whereas a well injection with a tracer concentration of one is located in a fracture intersection cell near the centre of the domain. Starting from an initial tracer solution of $T = 0$, the simulation runs until $t = 2$.

Since all permeabilities are isotropic, we discretise using TPFA in all dimensions. In this test case, we apply the Schur complement intersection elimination procedure and compare the solution with one where the intersection cells were kept. The solution is visualised in Fig. 13. The tracer distribution clearly shows how the symmetry of the test case is broken by the heterogeneous matrix permeability. In the upper half, a substantial part of the flow takes place in the matrix, whereas virtually all of it is conducted through the fractures in the lower half. Again, considering the solution without the elimination as reference, negligible errors and considerable condition number improvement are observed for the Schur complement procedure, see Table 4.

Table 4 Errors and condition number improvement for the elimination procedure applied in Case 4

Elimination	Pressure error	Tracer error	Condition number improvement
Schur complement	1.6×10^{-11}	4.8×10^{-11}	8.6×10^{3}

4 Conclusion

Choosing a hybrid-dimensional conforming DFM model for simulations of flow in fractured porous media may relieve some of the challenges posed by the fractures. We have considered procedures to alleviate two of the remaining restrictions, namely the coupling of fractures and rock matrix and the wide range in discretisation cell size. Firstly, we demonstrated how discretising the subdomains individually and coupling them two at a time using a hierarchical framework for mixed-dimensional porous media problems simplifies the implementation and offers considerable modelling and discretisation flexibility, also facilitating reuse of code. Secondly, we proposed a Schur complement technique to elimination of intersection cells with a considerably broader scope than existing techniques. The technique is not limited to FV discretisations and may be applied to both 0D and 1D intersections of arbitrary permeabilities without incurring additional pressure error. The technique may be extended to more general models including nonlinear single-phase flow. However, preliminary investigations suggest that the computation of fluxes introduces some issues in cases involving tangential flow along 1D intersections. Similar limitations are expected for application to multi-phase flow problems.

The numerical test cases presented demonstrate the improved accuracy of the proposed intersection cell elimination technique compared to existing solutions. Further, they show that the proposed TPFA coupling between fractures and matrix does not lead to severely reduced solution quality even in the presence of moderate anisotropy. We also demonstrate how parameter heterogeneities and anisotropies have implications on the appropriateness of different models and discretisations, and show how the inherent flexibility of our approach renders it suitable for simulations considering heterogeneous parameter regimes. The approach retains the benefits of FV discretisations and renders it feasible for simulation of various processes in fractured porous media.

Appendix

This appendix offers a description of the MPFA discretisation of a domain of a given dimension, using the following key ingredients:

- A partition of the mesh into subcells, and of cell boundaries into subfaces.
- Collection of subcells into interaction regions surrounding each node.
- Basis functions for the pressure that have support in a single interaction region and are linear on each subcell.
- Enforcement of pressure continuity at a single point on the subfaces.
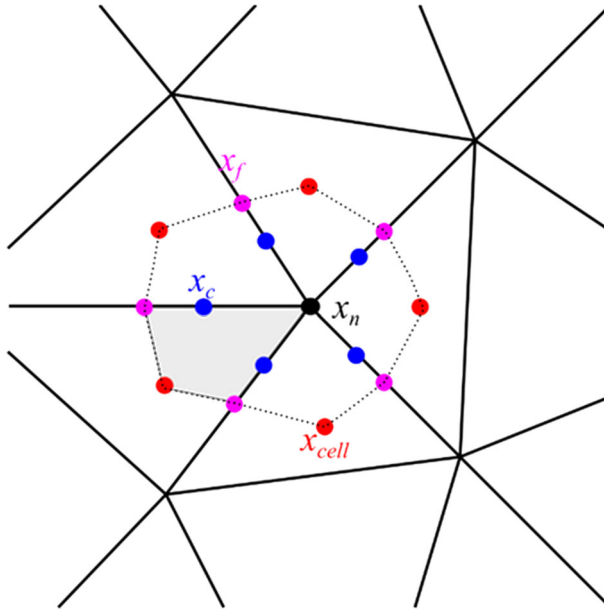- Enforcement of flux continuity at the subfaces.

Fig. 14 The O-shaped interaction region used for MPFA transmissibility calculations is marked by dotted lines connecting cell centres and face centres. Cell centres are shown in red, face centres in purple, continuity points in blue and the central vertex of the interaction region in black. A single subcell with four subfaces is indicated in grey

As for the TPFA, this ultimately yields the transmissibilities to be assembled in the discretisation matrix $\mathbf{A}$ for multiplication with the vector of cell centre pressures $\mathbf{p}$.

We commence by extending the geometrical notation. In Sect. 2.3.1, the cell and face centre locations x_{cell} and x_f are defined based on Fig. 2. Referring now to Fig. 14, we partition the mesh into subcells by splitting the faces at x_f. This results in new subcells with vertexes at cell centres, face centres and nodes x_n of the original mesh. The faces between cells are likewise split into subfaces. On the subfaces, we define the continuity points $x_c = (1 - \eta)x_f + \eta x_n$. The weight η, determining the placement along the face, will be discussed below. We define an interaction region to consist of all subcells sharing a common node x_n.

As in the case of TPFA, we discretise the pressure using cell centre values p_i. Further, the pressure is approximated using linear basis functions within each subcell. In other words, we can write the pressure at point x in the lth subcell of global cell i as

$$p\,(x) \approx p_i + \mathbf{G}_{i,l}\left(x - x_{i,\text{cell}}\right), \tag{20}$$

with $\mathbf{G}_{i,l}$ denoting the gradient of the subcell. Then, the discrete discharge over subface l in terms of the gradient of an adjacent cell i is

$$u_{i,l} = \mathbf{G}_{i,l}\left(x_{i,c} - x_{i,\text{cell}}\right) \cdot \mathbf{K}_i \cdot \boldsymbol{n}_{ij} A_l. \tag{21}$$

We now wish to write the discharge over a subface in terms of the cell centre pressures only, i.e. on the form

$$u_{i,l} = \sum_k t_k^* p_k, \tag{22}$$

where the asterisk is used to distinguish the subface weights from the full face transmissibilities introduced in Sect. 2.3. To obtain this, we proceed to eliminate the gradients from Eq. (21). This shall be done by (*i*) enforcing flux continuity over each subface and (*ii*) exploiting the additional assumption of pressure continuity at the continuity points.

Equating the discharges from either side over subface *l* shared by cells *i* and *j* then reads

$$\mathbf{K}_i \cdot \mathbf{G}_{i,l} \cdot \mathbf{n}_{ij} A_l = u_{i,l} = -u_{j,l} = -\mathbf{K}_j \cdot \mathbf{G}_{j,l} \cdot \mathbf{n}_{ji} A_l, \qquad (23)$$

where A_l denotes the subface area.

Similarly, by Eq. (20) statement (*ii*) reads

$$p_i + \mathbf{G}_{i,l} \left(x_{l,c} - x_{i,cell} \right) = p_j + \mathbf{G}_{j,l} \left(x_{l,c} - x_{j,cell} \right). \qquad (24)$$

Here, $x_{l,c}$ denotes the continuity point on subface *l*. Inspection of Eqs. (23) and (24) reveals that each relation couples a given subcell to a neighbouring subcell of the same interaction region only through the values of cell centre pressures and gradients. This means that we may collect all equations for each interaction region to a *local* system of the form

$$\begin{pmatrix} \mathbf{S} & 0 \\ \mathbf{D_G} & \mathbf{D_p} \\ 0 & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{G} \\ \mathbf{p} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ \mathbf{I} \end{pmatrix}. \qquad (25)$$

The unknowns here are the gradients $\mathbf{G}$, which are to be eliminated, and the cell centre pressures $\mathbf{p}$. The first row represents the flux balance over the subfaces, with $\mathbf{S}$ denoting products between $\mathbf{K}$ and $\mathbf{n}$, accounting for the sign of the normal vector. The second row represents the pressure continuity, with the distances between cell centres and continuity points, $\mathbf{D}_G$, and the cell centre contributions in the form of a matrix of ± 1, $\mathbf{D_p}$. Finally, the third row is a cell-wise enforcement of pressure values, with $\mathbf{I}$ denoting the identity matrix. Note that in this system, the right-hand side is a matrix rather than a vector: For each interaction region, there is one equation system corresponding to each subcell. Through the last row, the pressure value is set to unity at one of the cell centres and to zero at all others, corresponding to calculation of the basis function centred at that particular cell centre.

Solving the system (25) for $\mathbf{G}$ allows substitution into Eq. (21) to obtain a type (22) equation. In other words, we have arrived at the weights t_k^* defining the subface discharges in terms of the cell centre pressures.

Once the subface weights have been computed for all interaction regions, the global system is obtained by a two-stage assembly. First, we sum over all subfaces of a given full face *k* to obtain the face transmissibility

$$t_k = \sum_l t_l^*. \qquad (26)$$

Then, the discharges on the boundary of each cell are collected and related to the cell centre pressure of the surrounding cells through these face transmissibilities. Thus, we arrive at a linear system of the form of Eq. (8) with each row representing the local mass conservation equation of a cell.

Different choices of continuity point placement, characterised by the value of η, lead to different variants of the MPFA method. The choice affects the discretisation's symmetry properties and robustness w.r.t. grid roughness, and the optimal placement depends on the type of grid applied. In the current work, we set $\eta = 0$ on Cartesian grids, and $\eta = 1/3$ on simplex grids, as suggested in Aavatsmark (2002), Klausen et al. (2008) and Friis et al. (2009). These sources contain further information on the method and its properties, as does Edwards and Rogers (1998).

References

Aavatsmark, I.: An introduction to multipoint flux approximations for quadrilateral grids. Comput. Geosci. **6**, 405–432 (2002). https://doi.org/10.1023/A:1021291114475

Ahmed, R., Edwards, M.G., Lamine, S., Huisman, B.A., Mayr, P.: Three-dimensional control-volume distributed multi-point flux approximation coupled with a lower-dimensional surface fracture model. J. Comput. Phys. **303**, 470–497 (2015). https://doi.org/10.1016/j.jcp.2015.10.001

Alboin, C., Jaffré, J., Roberts, J.E., Wang, X., Serres, C.: Domain Decomposition for Some Transmission Problems in Flow in Porous Media, pp. 22–34. Springer, Berlin, Heidelberg (2000). https://doi.org/10.1007/3-540-45467-5_2

Aziz, K., Settari, A.: Petroleum Reservoir Simulation. Applied Science Publishers, London (1979)

Barenblatt, G.E., Zheltov, I.P., Kochina, I.N.: Basic concepts in the theory of seepage of homogeneous fluids in fissurized rocks. J. Appl. Math. Mech. **24**(5), 1286–1303 (1960). https://doi.org/10.1016/0021-8928(60)90107-6

Berkowitz, B.: Characterizing flow and transport in fractured geological media: a review. Adv. Water Resour. **25**, 861–884 (2002). https://doi.org/10.1016/S0309-1708(02)00042-8

Boon, W.M., Nordbotten, J.M., Yotov, I.: Robust discretization of flow in fractured porous media (2017). arXiv:1601.06977

Brenner, K., Groza, M., Guichard, C., Lebeau, G., Masson, R.: Gradient discretization of hybrid dimensional Darcy flows in fractured porous media. Numer. Math. **134**(3), 569–609 (2016). https://doi.org/10.1007/s00211-015-0782-x

Dietrich, P., Helmig, R., Sauter, M., Hötzl, H., Köngeter, J., Teutsch, G.: Flow and Transport in Fractured Porous Media. Springer, Berlin, Heidelberg (2005). https://doi.org/10.1007/b138453

Edwards, M.G., Rogers, C.F.: Finite volume discretization with imposed flux continuity for the general tensor pressure equation. Comput. Geosci. **2**(4), 259–290 (1998). https://doi.org/10.1023/A:1011510505406

Endo, H.K., Long, J.S.: A model for investigating mechanical transport in fracture networks. Water Resour. Res. **20**(10), 1390–1400 (1984). https://doi.org/10.1029/WR020i010p01390

Flemisch, B., Berre, I., Boon, W., Fumagalli, A., Schwenck, N., Scotti, A., Tatomir, A.: Benchmarks for single-phase flow in fractured porous media. Adv. Water Resour. **111**, 239–258 (2018). https://doi.org/10.1016/j.advwatres.2017.10.036

Friis, H.A., Edwards, M.G., Mykkeltveit, J.: Symmetric positive definite flux-continuous full-tensor finite-volume schemes on unstructured cell-centered triangular grids. SIAM J. Sci. Comput. **31**(2), 1192–1220 (2009). https://doi.org/10.1137/070692182

Geiger, S., & Matthäi, S. (2014). What can we learn from high-resolution numerical simulations of single- and multi-phase fluid flow in fractured outcrop analogues? In G. H. Spence, J. Redfern, R. Aguilera, T. G. Bevan, J. W. Cosgrove, G. D. Couples, & J.-M. Daniel, *Advances in the Study of Fractured Reservoirs* (Vol. 374, pp. 125-144). London: Geological Society of London. doi:https://doi.org/10.1144/SP374

Granet, S., Fabrie, P., Lemonnier, P., Quintard, M.: A two-phase flow simulation of a fractured reservoir using a new fissure element method. J. Petrol. Sci. Eng. **32**, 35–52 (2001). https://doi.org/10.1016/S0920-4105(01)00146-2

Hajibeygi, H., Karvounis, D., Jenny, P.: A hierarchical fracture model for the iterative multiscale finite volume method. J. Comput. Phys. **230**(24), 8729–8743 (2011). https://doi.org/10.1016/j.jcp.2011.08.021

Jarvis, N., Jansson, P.-E., Dik, P.E., Messing, I.: Modelling water and solute transport in macroporous soil. I. Model description and sensitivity analysis. J. Soil Sci. **42**, 59–70 (1991). https://doi.org/10.1111/j.1365-2389.1991.tb00091.x

Karimi-Fard, M., Durlofsky, L.J., Aziz, K.: An efficient discrete-fracture model applicable for general-purpose reservoir simulators. Soc. Petrol. Eng. **9**(2), 227–236 (2004). https://doi.org/10.2118/88812-PA

Keilegavlen, E., Fumagalli, A., Berge, R., Stefansson, I., Berre, I.: PorePy: An Open-Source Simulation Tool for Flow and Transport in Deformable Fractured Rocks (2017). arXiv:1712.00460

Kiraly, L.: Remarques sur la simulation des failles et du réseau karstique par éléments finis dans les modèles d'écoulement. Bulletin du Centre d'Hydrogéologie **3**, 155–167 (1979). https://doi.org/10.1007/s00767-006-0140-0

Klausen, R.A., Radu, F., Eigestad, G.T.: Convergence of MPFA on triangulations and for Richards' equation. Int. J. Numer. Meth. Fluids **58**(12), 1327–1351 (2008). https://doi.org/10.1002/fld.1787

Lie, K.-A.: An Introduction to Reservoir Simulation Using MATLAB. SINTEF ICT, Oslo (2016)

Long, J.S., Remer, J.S., Wilson, C.R., Witherspoon, P.A.: Porous media equivalents for networks of discontinuous fractures. Water Resour. Res. **18**(3), 645–658 (1982). https://doi.org/10.1029/WR018i003p00645

Moinfar, A., Varavei, A., Sepehrnoori, K.: Development of an efficient embedded discrete fracture model for 3D compositional reservoir simulation in fractured reservoirs. SPE J. **19**(2), 289–303 (2014). https://doi.org/10.2118/154246-PA

Mokhtari, M., Tutuncu, A.N., Boitnott, G.N.: Intrinsic anisotropy in fracture permeability. Interpretation **3**(3), 43–53 (2015). https://doi.org/10.1190/INT-2014-0230.1

Monteagudo, J.E., Firoozabadi, A.: Control-volume method for numerical simulation of two-phase immiscible flow in two- and three-dimensional discrete-fractured media. Water Resour. Res. (2004). https://doi.org/10.1029/2003WR002996

Nordbotten, J.M., Aavatsmark, I., Eigstad, G.T.: Monotonicity of control volume methods. Numer. Math. **106**(2), 255–288 (2007). https://doi.org/10.1007/s00211-006-0060-z

Preuss, K., Narasimhan, T.N.: A practical method for modeling fluid and heat flow in fractured porous media. Soc. Petrol. Eng. J. **25**(1), 14–26 (1985). https://doi.org/10.2118/10509-PA

Reichenberger, V., Jakobs, H., Bastian, P., Helmig, R.: A mixed-dimensional finite volume method for multiphase flow in fractured porous media. Adv. Water Resour. **29**(7), 1020–1036 (2006). https://doi.org/10.1016/j.advwatres.2005.09.001

Sahimi, M.: Flow and Transport in Porous Media and Fractured Rock, 2nd edn. Wiley-VCH, Weinheim (2011)

Sandve, T.H., Nordbotten, J.M., Berre, I.: An efficient multi-point flux approximation for discrete fracture-matrix simulations. J. Comput. Phys. **231**(9), 3784–3800 (2012). https://doi.org/10.1016/j.jcp.2012.01.023

Slough, K.J., Sudicky, E.A., Forsyth, P.A.: Grid refinement for modeling multiphase flow in discretely fractured porous media. Adv. Water Resour. **23**(3), 261–269 (1999). https://doi.org/10.1016/S0309-1708(99)00009-3

Stefansson, I.: A Comparison of Two Numerical Models for Flow in Fractured Porous Media and the Impact of Fracture Intersection Cell Removal. University of Bergen, Bergen (2016). Retrieved from http://hdl.handle.net/1956/15407

Tan, Y., Pan, Q., Liu, J., Wu, Y., Haque, A., Connell, L.D.: Experimental study of permeability and its anisotropy for shale fracture supported with proppant. J. Nat. Gas Sci. Eng. **44**, 250–264 (2017). https://doi.org/10.1016/j.jngse.2017.04.020

Walton, K.M., Unger, A.J., Ioannidis, M.A., Parker, B.L.: Impact of eliminating fracture intersection nodes in multiphase compositional flow simulation. Water Resour. Res. **54**(4), 2917–2939 (2017). https://doi.org/10.1002/2016WR020088

Warren, J.E., Root, P.J.: The behavior of naturally fractured reservoirs. Soc. Petrol. Eng. J. **3**(3), 245–255 (1963). https://doi.org/10.2118/426-PA

Zhang, F.: The Schur Complement and Its Applications. Springer, New York (2005)