



Unstructured grid optimization for improved monotonicity of discrete solutions of elliptic equations with highly anisotropic coefficients

Martin J. Mlacnik^{*}, Louis J. Durlofsky

Department of Petroleum Engineering, Stanford University, Green Earth Sciences Building, 367 Panama Street, Room 65, Stanford, CA 94305-2220, United States

Received 23 July 2005; received in revised form 7 December 2005; accepted 8 December 2005

Available online 25 January 2006

Abstract

Multipoint flux approximation (MPFA) techniques are commonly applied for discretizing the porous media flow equations within the context of finite volume numerical procedures. Although these methods can be applied to heterogeneous, anisotropic systems on generally unstructured grids, the inverse of the resulting linear operator can suffer from a loss of monotonicity at high permeability anisotropy ratios, resulting in spurious oscillations of the pressure solution. The purpose of this paper is to develop a method for optimizing unstructured grids in two and three dimensions such that the monotonicity behavior of the MPFA technique is significantly improved. The method employs anisotropic triangulation and can be readily combined with permeability upscaling procedures. Results are presented for a variety of examples and the technique is shown to perform well on problems involving complex grid point distributions and heterogeneous permeability fields with strong anisotropy ratios (of $O(100)$). Oscillation free pressure solutions are achieved in all cases considered, even for examples in which the original grid shows large oscillations.

© 2005 Elsevier Inc. All rights reserved.

Keywords: Unstructured gridding; Anisotropic triangulation; MPFA; Upscaling; Reservoir simulation; Permeability; Maximum principle; Anisotropy; Elliptic equations

1. Introduction

Subsurface formations such as petroleum reservoirs and aquifers often display significant geological and geometric complexity. These systems consist of anisotropic porous rock characterized by significant variation in the magnitude of permeability. Accurate simulation of flow in these formations requires grids that are capable of resolving geological complexity coupled with discretization schemes suitable for unstructured grids with

^{*} Corresponding author. Present address: ConocoPhillips Company, 282 Geoscience Building, Bartlesville Technology Center, Bartlesville, OK 74004, United States. Tel.: +1 918 6611575.

E-mail addresses: mmlacnik@gmail.com, martin.mlacnik@conocophillips.com (M.J. Mlacnik), lou@stanford.edu (L.J. Durlofsky).

highly variable, anisotropic coefficients. A number of discretization techniques have been presented in the literature, e.g., multipoint flux approximation (MPFA) for logically structured grids [2,15,29,33], MPFA for unstructured grids [3,14], k -orthogonal grids with two-point flux approximation [21], mimetic finite differences [38,23,24] and mixed finite element approaches [5,12,43]. There are close relationships between many of these techniques; for example, links between MPFA, mimetic finite differences and mixed finite elements are analyzed in [27].

Within the context of petroleum reservoir simulation, MPFA methods are commonly considered, as they are control-volume based, locally conservative, introduce only one unknown per block, are applicable for multiphase flow problems, and allow the use of grids with general geometry. This paper focuses on MPFA for unstructured grids, as introduced (independently) by Aavatsmark et al. [3,4] and Verma and Aziz [41]. It is known, however, that MPFA may lead to severe oscillations of the pressure solution, caused by the non-monotonicity of the inverse of the coefficient matrix, in cases with high degrees of permeability anisotropy.

In practical applications, permeability fields for flow simulation models are usually obtained by upscaling (or coarsening) an underlying highly detailed geological description (this geological model is often generated geostatistically). This upscaling procedure can lead to high degrees of anisotropy, as is readily apparent by considering layered systems, which qualitatively characterize many geological settings. For such a system, the effective permeability in the direction along the layers is given by the arithmetic average of the individual layer permeabilities, while the effective permeability normal to the layers is given by the harmonic average. For highly variable layer permeabilities, these two averages can differ by orders of magnitude. In the context of a two-point discretization scheme, where the grid is aligned with the principal directions of permeability, this anisotropy may not pose a problem. However, with unstructured grids that are not generally aligned with the permeability tensor, MPFA discretizations can lead to oscillatory pressure solutions. In fact, this behavior was observed in the context of upscaling and gridding for unstructured models in recent work by Prevost [36]. His observations provide a key motivation for this study.

The work presented in this paper addresses the problems caused by non-monotonicity from a grid optimization perspective that focuses on the link between MPFA discretization and grid geometry. Currently existing approaches within an MPFA context modify the polygonal/polyhedral grid to enhance the diagonal dominance [15] or symmetry [16] of the coefficient matrix. Here, we identify the inner angles of the triangulation underlying the polygonal/polyhedral grid as the primary target for optimization, which leads to the development of an anisotropic triangulation optimization algorithm. Though our emphasis in this study is on subsurface flow modeling, the methods developed here are also applicable for solutions of elliptic equations in other applications.

Many of the applications of triangular mesh optimization arise in the context of dynamic adaptive gridding for the finite element technique, where mesh adaptation is performed to control the solution error. The grid optimization metrics strongly depend on the specific problem requirements, e.g., satisfying the Delaunay criteria or directional metrics recovered from interpolating over the solution obtained from a mesh used for a previous time step. In this context, work with a focus on directional metrics has been published by [10,22,31,34]. From a strongly geometric perspective, anisotropic triangulations have been discussed by [8,20,25]. The grid optimization technique developed in this paper exploits an approach based on topological operators, similar to the edge swapping used by Freitag and Ollivier-Gooch [18]. Our approach will be shown to be well suited for improving grids for control-volume discretization techniques, as it has the ability to handle directional metrics which may exhibit strong variation and local randomness. A complication that arises within the context of upscaling is that, as the grid is changed (based on the directional metrics), the permeability tensor itself changes because it is computed from an underlying fine grid description. This effect is treated within our grid optimization by ensuring that, upon convergence, consistency between the grid and the upscaled permeability tensor is achieved.

This paper is organized as follows. The governing equations and the MPFA formulation are presented in Section 2, where the criteria guaranteeing an M -matrix for the three-dimensional case are developed. Section 3 discusses the details of grid optimization techniques for unstructured grids in two and three dimensions. Numerical examples demonstrating the convergence of MPFA on optimized grids and the elimination of spurious oscillations in heterogeneous problems are presented in Section 4. Conclusions and future directions are discussed in Section 5.

2. Governing equations and multipoint flux approximation

In this section we will first formulate the equations governing the flow in porous media in a control-volume framework. Next, we introduce the MPFA discretization technique for unstructured grids in two and three dimensions. Finally, the monotonicity behavior of linear operators deriving from MPFA discretization on unstructured grids will be discussed by focusing on the link between discretization and grid generation and its impact on the M -matrix property of the resulting linear operator.

2.1. Control-volume formulation of flow in porous media

Modeling flow in subsurface formations requires the solution of conservation equations coupled with Darcy's law and descriptions of phase behavior. In the case of multiple components flowing in multiple phases, the flow equations can be formulated in terms of a so-called pressure equation coupled with a sequence of hyperbolic conservation laws. The form of this pressure equation and the complications introduced by strong permeability anisotropy are similar to those of the pressure equation describing flow of only a single phase. Thus, in this work we will study the single-phase pressure equation. In the limit of incompressible fluid and rock, this equation is elliptic and is given by

$$\nabla \cdot (\mathbf{k} \nabla p) = Q, \quad (1)$$

where p denotes the pressure, Q represents sources and sinks and \mathbf{k} is the permeability tensor of the form

$$\mathbf{k}(\mathbf{x}) = \begin{pmatrix} k_{xx} & k_{xy} & k_{xz} \\ k_{yx} & k_{yy} & k_{yz} \\ k_{zx} & k_{zy} & k_{zz} \end{pmatrix}. \quad (2)$$

The permeability tensor is in general a function of location \mathbf{x} and is assumed to be symmetric and positive-definite. In a control-volume formulation, the reservoir domain is subdivided into a number of finite subvolumes Ω_j and an approximate solution for the pressure is sought at the center of each of these control volumes. Assuming that \mathbf{k} is constant within Ω_j , the flux q across the control-volume boundary $\partial\Omega$ can be expressed by integrating Eq. (1):

$$q_j = - \oint_{\partial\Omega_j} (\mathbf{k} \nabla p) \cdot \hat{\mathbf{n}} \, dS, \quad (3)$$

where $\hat{\mathbf{n}}$ denotes the outward-pointing unit normal vector on the surface S of $\partial\Omega$. This expression is valid because the Darcy velocity \mathbf{u} is given by $\mathbf{u} = -\mathbf{k} \nabla p$.

2.2. Principles of multipoint flux approximation in 2D and 3D

In this section we will briefly review the derivation of MPFA transmissibility coefficients in physical space. A detailed description for two-dimensional polygonal grids was presented by Aavatsmark et al. [3]. An approach for three-dimensional polyhedral grids was given by Verma and Aziz [41]. Our descriptions below follow those given in these references.

The MPFA technique under discussion derives transmissibility coefficients for unstructured grids formed from triangular meshes. These meshes are often Delaunay triangulations in two dimensions and Delaunay tetrahedrizations in three dimensions (note that we will also refer to *tetrahedrization* using the more common term *triangulation*). The triangular meshes are also termed *primal* grids as they are the basis for the generation of the *dual* polygonal or polyhedral control-volume grids that are actually used for flow simulation. Individual control volumes of the dual grid are commonly referred to as *blocks*. Block parameters such as permeability are assumed to be constant within each control volume but may vary in an arbitrary manner from block to block. In the standard case in two dimensions, the control-volume boundaries are formed from the sub-interfaces generated by joining the barycenter \mathbf{b} of each triangle with the mid-points \mathbf{m} of the three triangle edges (Fig. 1). Hence, a full block interface in two dimensions is composed of a maximum of two sub-interfaces.




Fig. 1. Two-dimensional control volume and interaction region showing vectors used for the MPFA discretization.

In three dimensions, control-volume boundaries are formed from the barycenter of the tetrahedral mesh elements \mathbf{b}_t , the barycenters of their faces \mathbf{b}_f and the mid-points \mathbf{m} of their edges, as illustrated in Fig. 2. This construction ensures that all sub-interfaces are planar [9]. As an edge in a three-dimensional triangulation can be shared by an arbitrary number of tetrahedra, there is no restriction on the number of sub-interfaces contributing to a full block interface. The dual grids resulting from these procedures are always non-overlapping.

This construction decomposes the boundary surface of a control volume into a set of discrete interfaces between neighboring control volumes and hence the boundary flux given by Eq. (3) is decomposed into a series of interface fluxes q_i . MPFA seeks to approximate the fluxes across each interface i by a weighted combination of multiple pressures

$$q_i \approx \sum_{j \in \Gamma} t_{ij} p_j, \tag{4}$$

where Γ denotes the set of points contributing to the flux, p_j is the pressure at the center of grid block j , and the weights t_{ij} are termed transmissibility coefficients, which account for the contribution from the pressure in block j to the interface flux q_i . To ensure $q_i = 0$ when the pressure is constant, the coefficients must satisfy

$$\sum_{j \in \Gamma} t_{ij} = 0. \tag{5}$$

The computation of the transmissibility coefficients for each sub-interface is based on the geometrical simplices of the primal grid. Each simplex forms a so-called interaction region, comprising three sub-cells and three sub-interfaces in two dimensions and four sub-cells and six sub-interfaces in three dimensions. We first describe the determination of t_{ij} for the two-dimensional case.

Integration of Eq. (3) over a sub-interface yields

$$q_{ik} = -\mathbf{n}_i^T \mathbf{k}_k \nabla p_k, \tag{6}$$

where ∇p_k is the pressure gradient in sub-cell k and \mathbf{n} is the interface normal vector of magnitude equal to the interface area. For the derivation of t_{ij} , a linear variation of the pressure in each sub-cell of the interaction




Fig. 2. 3D interaction region with the points involved in the construction of the sub-cell around vertex 1 (left) and tetrahedron face normals with numbering as used for the MPFA discretization.

region is assumed; i.e., ∇p_k is a constant that can be directly expressed as a linear combination of the pressures at the interaction region vertices. For the case where \mathbf{k} is homogeneous over the global domain, analytical expressions for the transmissibility coefficients can be derived. Specifically, as given for the two-dimensional case in [3],

$$t_{ij} = \frac{1}{2F} \mathbf{n}_i^T \mathbf{k} \mathbf{v}_j, \tag{7}$$

where F is the area of the triangular interaction region and \mathbf{v}_j are the vectors perpendicular to the triangle edge opposite to corner j , with their magnitudes equal to the edge length (see Fig. 1).

For heterogeneous media and general grid geometry, no simple analytical expressions for the transmissibility coefficients can be derived. If \mathbf{k} is heterogeneous, the pressure gradient in each sub-cell must be defined separately by assuming pressure continuity between neighboring sub-cells at the shared edge mid-points, as illustrated in Fig. 3 for the two-dimensional case. This introduces one additional pressure unknown \bar{p} per interaction region edge. Imposing flux continuity over each of the sub-interfaces in the interaction region provides sufficient conditions to eliminate the additional unknowns. Using Eq. (6) and the numbering introduced in Fig. 1, flux continuity for each of the three sub-interfaces can be written as:

$$\begin{aligned} \mathbf{n}_1^T \mathbf{k}_1 \nabla p_1 &= \mathbf{n}_1^T \mathbf{k}_2 \nabla p_2, \\ \mathbf{n}_2^T \mathbf{k}_2 \nabla p_2 &= \mathbf{n}_2^T \mathbf{k}_3 \nabla p_3, \\ \mathbf{n}_3^T \mathbf{k}_3 \nabla p_3 &= \mathbf{n}_3^T \mathbf{k}_1 \nabla p_1. \end{aligned} \tag{8}$$

After substituting the expressions for the pressure gradient into Eq. (8) and introducing the vectors $\mathbf{p} = [p_1, p_2, p_3]^T$ and $\bar{\mathbf{p}} = [\bar{p}_1, \bar{p}_2, \bar{p}_3]^T$, the system of equations specifying flux continuity can be expressed in matrix notation as $\mathbf{E}\bar{\mathbf{p}} + \mathbf{F}\mathbf{p} = \mathbf{G}\bar{\mathbf{p}} + \mathbf{H}\mathbf{p}$, where \mathbf{E} , \mathbf{F} , \mathbf{G} and \mathbf{H} are 3×3 matrices depending on permeability and geometry only. Now, defining the transmissibility matrix \mathbf{T} (with coefficients t_{ij}) as the 3×3 matrix that relates fluxes to pressures, we have

$$\mathbf{E}\bar{\mathbf{p}} + \mathbf{F}\mathbf{p} = \mathbf{G}\bar{\mathbf{p}} + \mathbf{H}\mathbf{p} = \mathbf{T}\mathbf{p}, \tag{9}$$

from which \mathbf{T} can be determined (see below).

The derivation of transmissibility coefficients in three dimensions is analogous to the two-dimensional case. In the case of a homogeneous system, using the principles applied above, we obtain, in analogy to Eq. (7), the following expression for the transmissibility coefficients:

$$t_{ij} = \frac{1}{3V} \mathbf{n}_i^T \mathbf{k} \mathbf{v}_{\Delta_j}, \tag{10}$$

where V is the tetrahedron volume and \mathbf{v}_{Δ_j} is the outwards normal vector on the tetrahedron face Δ_j (this face lies opposite to vertex j). The vector \mathbf{v}_{Δ_j} is of magnitude equal to the area of Δ_j .

In the more general case of heterogeneous \mathbf{k} , a linear pressure variation for each of the four sub-cells is introduced, leading to six additional pressure unknowns, one for each interaction region edge. The linear pressure variation in a sub-cell, for example that around vertex 1 in Fig. 2, can be expressed in terms of the pressure at vertex 1 and the pressures at the mid-points ($\mathbf{m}_{12}, \mathbf{m}_{13}, \mathbf{m}_{14}$) of the three tetrahedron edges where pressure continuity is enforced. As in the two-dimensional case, the additional pressure unknowns can be eliminated by enforcing flux continuity across the sub-interfaces,




Fig. 3. Points involved in the definition of the pressure gradients in the sub-cells of a 2D interaction region.

$$\begin{aligned}
\mathbf{n}_{12}^T \mathbf{k}_1 \nabla p_1 &= \mathbf{n}_{12}^T \mathbf{k}_2 \nabla p_2, \\
\mathbf{n}_{13}^T \mathbf{k}_1 \nabla p_1 &= \mathbf{n}_{13}^T \mathbf{k}_3 \nabla p_3, \\
\mathbf{n}_{14}^T \mathbf{k}_1 \nabla p_1 &= \mathbf{n}_{14}^T \mathbf{k}_4 \nabla p_4, \\
\mathbf{n}_{23}^T \mathbf{k}_2 \nabla p_2 &= \mathbf{n}_{23}^T \mathbf{k}_3 \nabla p_3, \\
\mathbf{n}_{24}^T \mathbf{k}_2 \nabla p_2 &= \mathbf{n}_{24}^T \mathbf{k}_4 \nabla p_4, \\
\mathbf{n}_{34}^T \mathbf{k}_3 \nabla p_3 &= \mathbf{n}_{34}^T \mathbf{k}_4 \nabla p_4,
\end{aligned} \tag{11}$$

where \mathbf{n}_{jk} designates the normal on the sub-interface between vertices j and k pointing out of sub-cell j . For example, referring to Fig. 2, \mathbf{n}_{12} is the normal on the sub-interface between the sub-cells around vertices 1 and 2. This sub-interface is defined by the points \mathbf{m}_{12} , \mathbf{b}_{f123} , \mathbf{b}_t and \mathbf{b}_{f124} .

Inserting the expressions for the pressure gradient in each sub-cell into Eq. (11) leads to an equation system of the same general form as Eq. (9). The vectors holding the pressure unknowns are now written as $\mathbf{p} = [p_1, p_2, p_3, p_4]^T$ and $\bar{\mathbf{p}} = [\bar{p}_1, \bar{p}_2, \bar{p}_3, \bar{p}_4, \bar{p}_5, \bar{p}_6]^T$, the matrices \mathbf{E} and \mathbf{G} are of size 6×6 , and \mathbf{F} , \mathbf{H} and \mathbf{T} are of size 6×4 .

After manipulating Eq. (9) to eliminate the sub-interface pressure unknowns $\bar{\mathbf{p}}$, we have

$$\mathbf{T} = \mathbf{E} \cdot (\mathbf{E} - \mathbf{G})^{-1} \cdot (\mathbf{H} - \mathbf{F}) + \mathbf{F}. \tag{12}$$

This equation can be solved for all transmissibility coefficients in the interaction region and applies for both the two- and three-dimensional cases. Each row of \mathbf{T} consists of the transmissibility coefficients for one sub-interface, so the sub-interface fluxes in the interaction region can be expressed by

$$\mathbf{q} = \mathbf{T}\mathbf{p}, \tag{13}$$

where \mathbf{q} is a 3×1 vector in two dimensions and a 4×1 vector in three dimensions holding the sub-interface fluxes. Application of this procedure to the other interaction regions and assembly of all of the contributions gives the full discretization of Eq. (1).

2.3. Monotonicity and the M -matrix property

Numerical schemes for elliptic problems of the form of Eq. (1) should ideally lead to solutions that are free of nonphysical oscillations, i.e., the solutions should obey a discrete maximum principle. As discussed in, for example, [1], this can be guaranteed if the inverse of the matrix \mathbf{A} of the linear system $\mathbf{A}\mathbf{p} = \mathbf{Q}$ that approximates Eq. (1) satisfies the monotonicity inequality

$$\mathbf{A}^{-1} \geq \mathbf{0}, \tag{14}$$

where $\mathbf{0}$ is the zero matrix; i.e., all entries in \mathbf{A}^{-1} are non-negative.

A sufficient condition for Eq. (14) to hold is that \mathbf{A} is an M -matrix, i.e., \mathbf{A} is irreducible and the matrix coefficients $a_{i,j}$ satisfy

$$\begin{aligned}
a_{i,i} &> 0 \quad \forall i, \\
a_{i,j} &\leq 0, \quad \forall i, j, \quad i \neq j, \\
\sum_j a_{i,j} &\geq 0 \quad \forall i
\end{aligned} \tag{15}$$

with strict inequality for at least one row [40]. As Eq. (14) is only a sufficient condition for monotonicity one may still obtain a monotone inverse for a matrix that does not satisfy Eq. (15). For MPFA on structured uniform parallelogram grids in homogeneous media, criteria for the monotonicity property for non M -matrices were recently developed in [32].

It has been observed by several authors, e.g., [3,15,16], that the transmissibility coefficients computed by an MPFA discretization in two dimensions are not necessarily positive and will therefore not generally lead to an M -matrix. For unstructured grids in the 2D homogeneous case the sign of the transmissibility coefficients can be directly related to the internal angles of the triangle that forms the according interaction region. Eigestad

et al. [16] analyzed the contribution of a single interaction region to the full matrix \mathbf{A} by investigating the 3×3 matrix \mathbf{L} arising locally for a single interaction region. Using the numbering of Fig. 1, off-diagonal element $l_{2,1}$ can be written as

$$l_{2,1} = t_{21} - t_{11} = \frac{1}{2F} \mathbf{n}_2^T \mathbf{k} \mathbf{v}_2 - \frac{1}{2F} \mathbf{n}_1^T \mathbf{k} \mathbf{v}_2 = \frac{1}{4F} \mathbf{v}_3^T \mathbf{k} \mathbf{v}_2. \tag{16}$$

If \mathbf{k} is isotropic, it follows that the sign of $l_{2,1}$ will depend on the inner product $\mathbf{v}_3^T \mathbf{v}_2$. This product will only be non-positive if the angle between \mathbf{v}_3 and \mathbf{v}_2 is greater than or equal to $\pi/2$, which is equivalent to the inner angle around vertex 1 being less than or equal to $\pi/2$. This reasoning can be extended to the general homogeneous case by investigating the inner product in a computational space that reduces the permeability tensor to the identity matrix. If any inner angle of the triangle in computational space exceeds $\pi/2$ there will be a positive off-diagonal element in the local matrix. On the level of the full matrix \mathbf{A} , coefficients include contributions of the two sub-interfaces forming a full block interface. This relaxes the angle criteria and, as a consequence, the sum of the two angles opposite the shared edge being less than or equal to π becomes a sufficient condition for a non-positive off-diagonal element [4]. This mesh property is satisfied by a two-dimensional Delaunay triangulation. The same results have been reported for the control-volume finite element technique by Forsyth [17].

In three dimensions it can be shown via the counter-example presented in [30] that even in the homogeneous, isotropic case a Delaunay triangulation does not guarantee an M -matrix. The conditions that guarantee an M -matrix for the general homogeneous case will be derived in the following. Again, we proceed by analyzing the contributions of a single interaction region to the full matrix \mathbf{A} . We consider the 4×4 matrix \mathbf{L} arising for the local interaction region. Using the vertex numbering and the sub-interface normal indexing convention introduced in Fig. 4, the following expression for off-diagonal element $l_{1,2}$ is obtained for isotropic \mathbf{k} :

$$l_{1,2} = t_{12} + t_{22} + t_{32} = \frac{1}{3V} \mathbf{k} \mathbf{v}_{\Delta_2} (\mathbf{n}_1 + \mathbf{n}_2 + \mathbf{n}_3). \tag{17}$$

As a consequence of the divergence theorem, the sum over all interface normals of a polyhedron yields zero if each normal has a magnitude equal to the area of its associated interface. Furthermore, the barycenter based grid construction ensures that the magnitude of an interface normal on a sub-cell face that is part of a tetrahedron face is exactly one third of the magnitude of the normal on the full tetrahedron face (see Fig. 4). Hence, the following relations apply to the tetrahedron interface normals

$$\mathbf{v}_{\Delta_1} + \mathbf{v}_{\Delta_2} + \mathbf{v}_{\Delta_3} + \mathbf{v}_{\Delta_4} = 0 \tag{18}$$

and to the sub-cell around interaction region (vertex) 1

$$\mathbf{n}_1 + \mathbf{n}_2 + \mathbf{n}_3 + \frac{1}{3} (\mathbf{v}_{\Delta_2} + \mathbf{v}_{\Delta_3} + \mathbf{v}_{\Delta_4}) = 0. \tag{19}$$

Combining Eqs. (17)–(19) leads to the following expression for the off-diagonal element $l_{1,2}$




Fig. 4. Interface normals of the sub-cell around vertex 1 in a 3D interaction region.

$$l_{1,2} = \frac{1}{9V} \mathbf{k} \mathbf{v}_{\Delta_2} \mathbf{v}_{\Delta_1}. \quad (20)$$

The sign of $l_{1,2}$ depends on the inner product $\mathbf{v}_{\Delta_2} \mathbf{v}_{\Delta_1}$, which will only be negative if the dihedral angle between the tetrahedron faces Δ_2 and Δ_1 is smaller than $\pi/2$. This extends the two-dimensional angle criterion to three dimensions and indicates that each interaction region will only lead to negative contributions to the off-diagonal elements of the full matrix if all of its dihedral angles are smaller than $\pi/2$. This criterion can also be applied to a general homogeneous \mathbf{k} if the vector product is evaluated in a computational space that maps the permeability tensor to the identity matrix. In general, a Delaunay triangulation of an arbitrary set of vertices does not satisfy this criterion [7] and therefore is not a sufficient condition for an M -matrix.

Assuming that \mathbf{k} is isotropic and using the following expression for the tetrahedron volume

$$V = \frac{2}{3e_{ij}} A_{\Delta_i} A_{\Delta_j} \sin(\theta_{ij}), \quad (21)$$

where A_{Δ_i} and A_{Δ_j} are the areas of the tetrahedron faces Δ_i and Δ_j , e_{ij} is the length of the tetrahedron edge shared by faces Δ_i and Δ_j and θ_{ij} is the dihedral angle between tetrahedron faces Δ_i and Δ_j , and expressing the dot product between the two outwards normals as

$$\mathbf{v}_{\Delta_i} \mathbf{v}_{\Delta_j} = -A_{\Delta_i} A_{\Delta_j} \cos(\theta_{ij}), \quad (22)$$

Eq. (20) can be rewritten as follows

$$l_{1,2} = -\frac{1}{6} e_{12} \cot(\theta_{ij}). \quad (23)$$

Because a full block interface may be composed of arbitrarily many sub-interfaces, a given off-diagonal element of the full matrix combines contributions from multiple interaction regions. Clearly, the full matrix can still be an M -matrix in cases where the dihedral angle criterion is locally violated, as one interaction region may compensate for a positive contribution from another. Any off-diagonal element $a_{i,j}$ of the full flow matrix can be expressed by summing over the contributions of all interaction regions written in the form of Eq. (23). The summation leads to the following generalized expression for an off-diagonal element in the full flow matrix

$$a_{i,j} = -\frac{1}{6} \sum_{T \supset \bar{i}\bar{j}} (e_{ij} \cot(\theta_{ij}))_T, \quad (24)$$

where $\sum_{T \supset \bar{i}\bar{j}}$ indicates a summation over all tetrahedra T that share edge $\bar{i}\bar{j}$ and $(\cdot)_T$ means that the expression is evaluated for tetrahedron T .

Eq. (24) effectively relaxes the dihedral angle criterion, as it is now sufficient that the sum of the contributions to an off-diagonal element remains negative, rather than each contribution itself. Eqs. (23) and (24) also hold for homogeneous cases with anisotropic permeability if they are evaluated in a computational space which reduces the permeability tensor to the identity matrix.

It is important to note that, in previous work, expressions identical to Eqs. (23) and (24) were derived by Xu and Zikatanov [44] for the Galerkin method with linear basis functions. Our derivation above shows that the discrete equations are identical for the control-volume scheme considered here. This agreement will also hold for anisotropic cases if \mathbf{k} is constant. This correspondence would be expected, since the two methods are known to give identical discretizations on two-dimensional triangular grids (e.g., [19]), but we are not aware of the previous demonstration of this for three-dimensional tetrahedral grids. For heterogeneous cases, with \mathbf{k} defined on the control volumes (in either two or three dimensions), these correspondences no longer hold.

Another method related to our control-volume scheme is the orthogonal subdomain collocation scheme analyzed in [37]. With this method, the control-volume grid is constructed using the circumcenters of the tetrahedra rather than their barycenters. For a Delaunay grid and isotropic permeability, this approach guarantees an M -matrix. In the context of petroleum reservoir simulation, this technique is similar to the use of Voronoi (or PEBI; i.e., perpendicular bisection) grids for isotropic \mathbf{k} and k -orthogonal Voronoi (or k -PEBI) grids for anisotropic \mathbf{k} , as suggested by Heinemann et al. [21]. Voronoi or k -orthogonal Voronoi grids (as

appropriate) yield two-point flux approximations and hence guarantee an M -matrix. For anisotropic permeability, however, k -orthogonal Voronoi grids can only be formed for systems with relatively low degrees of anisotropy. Limits for the anisotropy ratio were investigated in [4].

3. Grid optimization

In this section we introduce grid optimization techniques in two and three dimensions with the objective of improving the quality of the primal grid to better accommodate the multipoint discretization stencil.

The criteria described in Section 2.3 identify the internal angles of the geometric simplices that form the primal grid as critical parameters that impact the M -matrix property. Hence, the focus of this work is on the underlying triangulation rather than the associated dual polygonal/polyhedral grid. In the absence of a triangulation algorithm that is able to honor discrete directional metrics and at the same time account for angular restrictions, we propose a posteriori grid optimization techniques that operate on initial two- and three-dimensional triangular meshes. The resulting techniques can be classified as anisotropic triangulation optimization approaches.

Anisotropic triangular mesh optimization is governed by the objective of constructing a triangular mesh that locally conforms to size and direction constraints. These constraints are specified as metrics at each point in space through symmetric, positive-definite tensors that depend on the problem requirements and on the particular notion of mesh quality. For the purpose of mesh optimization, this work relies on the application of topological operators, e.g., [20], as powerful but inexpensive basic building blocks of grid optimization algorithms. These act only on a subset of an initial triangulation by locally modifying the mesh connectivity. Generally, the use of topological operators arises from the observation that there exist repeated topological patterns in a grid that are potential candidates for optimization.

The algorithms described in the following sections do not modify the original set of vertices; i.e., there is no shifting, adding or deleting of nodes. This is not a limitation of the technique itself, but rather an algorithm design decision based on the assumption that the grid node distribution has been obtained from some preprocessing step, e.g., flow-based gridding. Adding nodes locally will improve the quality of the grid and, in the three-dimensional case, will act to eliminate non-acute tetrahedra. The addition of points will best be accomplished by considering both grid issues and other data (approximate flow information, geological features, permeability connectivity) in combination. The algorithms described here can be used in conjunction with meshes that resolve internal features such as faults (see [41] for an illustration of fault resolution with control-volume grids). This is accomplished by restricting the optimization to edges and faces of the triangular elements that are not part of these internal boundaries.

3.1. Permeability and geometry

To account for tensorial permeability in grid optimization, a direct link between geometry and permeability must be established. Following Bear [6], it can be shown that the geometric interpretation of a permeability tensor is an ellipse in \mathbb{R}^2 and an ellipsoid in \mathbb{R}^3 . The following discussion will only refer to the three-dimensional case, with the two-dimensional case as a subset.

A general permeability ellipsoid is given by

$$\frac{\eta^2}{k_1} + \frac{v^2}{k_2} + \frac{\xi^2}{k_3} = 1, \quad (25)$$

where η , v and ξ are the independent coordinate variables in the axis directions of a vector space spanned by the eigenvectors \mathbf{e}_1 , \mathbf{e}_2 and \mathbf{e}_3 of the permeability tensor, and k_1 , k_2 , k_3 are the respective eigenvalues. The notions of permeability and spatial distance become equivalent in a computational space which reduces the permeability ellipsoid to a sphere. The relationship between physical and computational space is given by

$$\mathbf{x}' = \mathcal{M}\mathbf{x}, \quad (26)$$

where \mathbf{x}' and \mathbf{x} denote the coordinate vector in computational and physical space, respectively. The transformation matrix \mathcal{M} maps the permeability ellipsoid to a sphere through a combined rotation and stretching operation and is defined as

$$\mathcal{M} = (\mathcal{R}^{-1} \sqrt{\mathcal{S}} \mathcal{R})^{-1}, \quad (27)$$

where the columns of the matrix \mathcal{R} are composed of the eigenvectors of the permeability tensor in x - y - z space, accounting for the rotational part of the transformation

$$\mathcal{R} = (\mathbf{e}_1(x, y, z), \mathbf{e}_2(x, y, z), \mathbf{e}_3(x, y, z)). \quad (28)$$

\mathcal{S} is a diagonal matrix with the eigenvalues as elements

$$\mathcal{S} = \text{diag}(k_1, k_2, k_3), \quad (29)$$

accounting for the stretching operation.

3.2. Grid optimization in two dimensions

In two dimensions, Delaunay triangulations are commonly used for generating the initial primal grid, as efficient algorithms for their construction are readily available and Delaunay triangulations have several favorable properties [13]. For example, Delaunay triangulations maximize the minimum inner angle of the triangular elements (MaxMin angle property).

To render a triangulation more suitable for the MPFA discretization, it should account for directional metrics defined by the local permeability tensor. In the simplest case of globally homogeneous \mathbf{k} , the directional metric is constant throughout the domain. As proposed in [4], an optimal grid can be constructed in this case directly from a grid node distribution that has been mapped into a computational space. The primal grid resulting from this triangulation in computational space is optimal in a Delaunay sense and therefore satisfies the angle criterion that guarantees an M -matrix.

Rather than mapping grid points into computational space, an initial grid generated in physical space can be optimized through purely local manipulations of its topology. This local optimization is based on an edge swapping technique, introduced by Lawson [28], which exploits the fact that in two dimensions there are exactly two ways in which four non-degenerate vertices can be triangulated. Without accounting for directional metrics, any arbitrary triangulation of a given set of n points can be transformed into a Delaunay triangulation by looping over all edges that are shared by two triangles which form a convex quadrilateral. Each of these configurations can be optimized with respect to the MaxMin angle property by swapping the shared edge if the minimum internal angle of the two triangles is thus maximized. Regardless of the initial mesh, the algorithm converges to a global optimum, i.e., the Delaunay triangulation of the point set, with complexity $O(n^2)$.

This local grid optimization operation can be achieved by employing a single topological operator. The operator used for optimizing the MaxMin angle property of triangulations in \mathbb{R}^2 acts on a convex quadrilateral formed by two triangles sharing an edge and returns a topology representing two triangles. This operation will therefore be denoted by $Tr_{2 \rightarrow 2}$, following the nomenclature used by, e.g., [20]. To account for an additional directional metric such as full-tensor permeability, the MaxMin angle property must be evaluated in computational space. The required coordinate transformation follows Eq. (26). Fig. 5 illustrates the mechanics of the $Tr_{2 \rightarrow 2}$ operator.

The application of the $Tr_{2 \rightarrow 2}$ operator can be extended to heterogeneous media. Although applied here for general cases, this extension is most suitable for systems in which the eigenvectors (principal directions) of permeability vary relatively slowly (the magnitudes of k_1 and k_2 can, however, vary abruptly). We note that relatively slow variation in the eigenvectors is frequently observed in upscaled models of geological systems. This is because, although the principal values of upscaled permeability often display large local variations in magnitude, the principal directions are impacted by large-scale geologic phenomena (such as flooding events) that display consistent patterns and trends.

The local orientation of the permeability tensor can be captured by an average $\langle \mathbf{k} \rangle$ of the tensors associated with the four vertices. This tensor can then be used as the basis of a coordinate transformation. In this work we apply a simple arithmetic average, i.e.,




Fig. 5. Illustration of the mechanics of the Tr_{2-2} operator. (a) shows the original configuration in physical space, (b) the two alternative triangulations in computational space, and (c) the updated configuration in physical space.

$$\langle \mathbf{k} \rangle = \frac{1}{4} \sum_{i=1}^4 \mathbf{k}_i, \tag{30}$$

as it maintains symmetry and positive-definiteness of the resulting tensor. In regions with local randomness in permeability orientation, little or no grid optimization will be introduced, as illustrated in Fig. 6. We note that other types of local permeability averages could also be applied, though we have found that the simple arithmetic average performs well for the cases considered here.

The resulting grid optimization algorithm starts from an initial triangulation of the grid node distribution. Each iteration cycle consists of a loop over all internal edges (i.e., edges shared by two triangles) and repeatedly applies the topological operator Tr_{2-2} to each of these configurations. The operator evaluates the Max-Min angle property in computational space and, if necessary, updates the topology of the shell in physical space by flipping the central edge. The process is repeated until convergence is achieved. For homogeneous media the method will converge to the grid obtained from directly constructing the Delaunay triangulation in computational space. For cases with variable anisotropy, there is no guarantee of convergence to a global optimum, but the improvement in grid quality was almost always sufficient to overcome monotonicity problems for the anisotropy ratios considered in this paper. It is expected that the addition and/or movement of points (which were not investigated here) will act to further improve the grid. Note that the initial triangulation does not need to be Delaunay, but for problems with moderate anisotropy ratios Delaunay meshes are a good starting point and reduce the number of optimization iterations.

3.3. Grid optimization in three dimensions

In three dimensions the relaxed angle criterion given by Eq. (24), which guarantees an M -matrix, appears to be the obvious choice for an optimization metric. This criterion does not however measure the quality of a




Fig. 6. Arithmetic averaging of four tensors. The tensors are represented by solid ellipses, ellipses with dashed lines show the arithmetic average. If tensors are well aligned their average reproduces the trend in orientation and axis ratio (left figure), but with high degrees of variability the average ellipse relaxes towards a circle (right figure).

tetrahedron itself, but rather the quality associated with a tetrahedron edge. This makes it difficult to directly apply Eq. (24) in the commonly used triangulation optimization algorithms which aim at improving tetrahedron quality.

One alternative would be to revert to the more restrictive criterion that requires tetrahedrizations with no dihedral angle larger than $\pi/2$. This criterion is straightforward to implement in the context of topological operators, but is limited by the severe restrictions it imposes on the primal grid. In general, a Delaunay tetrahedrization of an arbitrary set of vertices does not satisfy this criterion and, moreover, a tetrahedrization in \mathbb{R}^3 satisfying the non-obtuse dihedral angle property may not necessarily be Delaunay [7]. Improving the grid quality measured in terms of dihedral angles is an active research area in the field of computational geometry. Along these lines, however, Bern et al. [7] showed that the dihedral angle criterion can only be strictly enforced by adding Steiner points. As indicated above, however, in our algorithm we do not add points to the original set of vertices. Thus we cannot expect to satisfy this angle criterion.

From the above discussion, it is clear that we cannot satisfy Eq. (24) directly in our optimization or guarantee reduction of the maximum dihedral angle to a value less than $\pi/2$. However, from Eq. (24) it is apparent that angles that approach π will lead to $a_{i,j}$ that are large and positive. The existence of such $a_{i,j}$ means that the resulting \mathbf{A} matrix is far from satisfying the M -matrix criteria (see Eq. (15)). The degree of violation of the M -matrix criteria will be reduced if we apply an optimization scheme to minimize the largest dihedral angles appearing in the model (with θ_{ij} evaluated in computational space). Although this will not in general lead to an M -matrix, the results obtained from our optimized grids do offer clear improvement in terms of the quality of the solution. This is consistent with the findings of Petrovskaya [35], who showed that the performance of his finite volume scheme could be understood to some extent in terms of simple measures of non-positivity (non-negativity in our case) of the matrix coefficients.

We now discuss the reconnection operations (referred to here as topological operators) employed to minimize the maximum dihedral angles in each shell. Lawson [28] showed that the grid optimization through edge swapping can be extended from \mathbb{R}^2 to \mathbb{R}^3 , although this extension is not straightforward. In three dimensions additional complications arise through the correspondence of edges and faces in swapping operations. This is illustrated in Fig. 7, which shows the two topologically valid tetrahedrizations of five non-degenerate vertices, i.e., either two tetrahedra sharing one face, or three tetrahedra sharing one edge, can be formed. Swapping the common face of the two-tetrahedra configuration will not generate a new face, but rather leads to a new edge that is shared by three tetrahedra. This operation can be encapsulated by the topological operator $Tr_{2 \rightarrow 3}$, which takes a two tetrahedra configuration (left triangulation in Fig. 7) as input and returns a three tetrahedra configuration (right triangulation in Fig. 7). $Tr_{2 \rightarrow 3}$ cannot handle the reverse operation, as it only acts on two tetrahedra configurations. Therefore, an additional operator needs to be defined that transforms a three tetrahedra configuration with a common edge into a two tetrahedra configuration with a common face and which will hence be denoted by $Tr_{3 \rightarrow 2}$.

In this work we use a topological operator technique similar to the edge swapping initially presented by de l'Isle and George [25] and adapted by Freitag and Ollivier-Gooch [18]. However, we extend these ideas to account for strongly directional metrics. In our procedure, a topological operator stores all alternative topol-




Fig. 7. The two topologically valid triangulations of five points in 3D. The left figure shows two tetrahedra sharing one face, while the right figure shows three tetrahedra sharing one edge in an exploded view. This illustrates $Tr_{2 \rightarrow 3}$ (left to right transform) and $Tr_{3 \rightarrow 2}$ (right to left transform).

ogies. Among these, the algorithm chooses the best solution according to the optimization metric(s) evaluated in computational space. This avoids the potentially expensive re-triangulation of a shell (i.e., the polyhedron formed by the tetrahedra sharing an edge) and the problem of developing a triangulation algorithm that can account for criteria other than the Delaunay property. Other mesh optimization techniques have been reported in the literature, e.g., [26] or [39], and could also be applied. The topological operator approach employed here proved to be very effective for the cases investigated while remaining simple and robust.

Following the discussion presented in Section 2.3, the quality of a mesh in an anisotropic medium should be evaluated in a computational space that maps the permeability ellipsoid to a sphere. The according coordinate transformation was also discussed in Section 3.1. If the medium is heterogeneous, as in the two-dimensional case the coordinate transformation is based on a local average of the permeability tensors

$$\langle \mathbf{k} \rangle = \frac{1}{n_{\text{shell}}} \sum_{i=1}^{n_{\text{shell}}} \mathbf{k}_i, \quad (31)$$

where n_{shell} is the number of vertices involved in the shell.

A rigorous application of topological operators for grid optimization, where each operator can only account for a specific number of tetrahedra sharing an edge, requires the definition of one operator for each possible configuration. The number of alternative tetrahedrizations increases rapidly with the number of vertices involved in the shell following the Catalan number [20], but de l'Isle and George found that for practical purposes it is sufficient to account for a maximum of nine tetrahedra sharing an edge. Freitag and Gooch [18] further decreased the number of topological operators, reporting good results for an approach based on operators that account for a maximum of seven tetrahedra sharing an edge. As the algorithm described in this paper is especially concerned with optimization in the presence of anisotropy affected by a degree of randomness, it is critical to preserve the representative character of the local average tensor, which is increasingly lost with increasing numbers of vertices.

For our application, it proved to be sufficient to introduce only two operators in addition to $Tr_{2 \rightarrow 3}$ and $Tr_{3 \rightarrow 2}$ defined above. These additional operators are designated $Tr_{4 \rightarrow 4}$ and $Tr_{5 \rightarrow 6}$. $Tr_{4 \rightarrow 4}$ is locally applied to entities in the full triangulation where four tetrahedra share an edge. It stores the two alternative triangulations (again consisting of four tetrahedra), evaluates the original and alternative configurations in a computational space defined by the optimization metric, and returns the optimum configuration, i.e., it either leaves the original configuration unchanged or updates the mesh locally, using one of the stored alternatives. $Tr_{5 \rightarrow 6}$ acts on configurations where five tetrahedra share an edge. The mechanics of this operator are similar to $Tr_{4 \rightarrow 4}$, but $Tr_{5 \rightarrow 6}$ needs to store and evaluate five alternative configurations consisting of six tetrahedra.

The alternative tetrahedrizations of a shell are sufficiently characterized through the possible triangulations of the equatorial polygon, which is the polygon that is formed by the vertices that are not part of the shared edge. The triangulation of the equatorial polygon, together with the two points of the central edge in the original configuration, define the tetrahedra that form the alternative configurations. To visualize the alternative tetrahedrizations of the operators $Tr_{4 \rightarrow 4}$ and $Tr_{5 \rightarrow 6}$, the triangulations of the according equatorial polygons are shown in Figs. 8 and 9. See [20] for further details on these topological operators.

Using topological operators, several different mesh optimization strategies can be envisioned [20]. Here we minimize the largest dihedral angle in all tetrahedra by applying topological operators to appropriate target entities (faces and edges) in the mesh. For this purpose a two step algorithm proved to be most effective. In the first step, the algorithm optimizes the grid with respect to the Delaunay property. That is, the topological operators choose the optimal configuration based on the Delaunay property evaluated in computational space. In the second step the grid is optimized with respect to the dihedral angles, again in computational space. In this step each topological operator returns the optimum configuration with respect to the dihedral angles, i.e., the configuration that minimizes the largest dihedral angle in all tetrahedra involved in a shell.

Both steps of the optimization are based on the same iterative scheme. Regardless of the optimization metric, a loop of the scheme can be summarized as follows. Note that the optimization defined here is for specified permeability tensors. Modifications required to incorporate an upscaling step (in which the tensors depend on the mesh) will be described below.




Fig. 8. Configuration in which four tetrahedra share an edge and the two alternative tetrahedrizations. The alternative triangulations of the equatorial polygon $\overline{1234}$ are shown below the corresponding tetrahedrizations. This represents the tetrahedrizations involved in $Tr_{4 \rightarrow 4}$.




Fig. 9. Alternative triangulations of the equatorial polygon for a configuration where five tetrahedra share an edge. Application of $Tr_{5 \rightarrow 6}$ eliminates the shared edge (original configuration with shared edge not shown).

- Generate a list of all faces in the grid.
- Loop over the face list and apply the topological operator $Tr_{2 \rightarrow 3}$ to each of the tetrahedra faces. $Tr_{2 \rightarrow 3}$ returns the optimal configuration with respect to the optimization metric and updates the mesh topology. If a face is part of a boundary (internal or external), it will not be optimized.
- Generate a list of all edges in the grid.
- Loop over all edges in the edge list and apply the topological operators $Tr_{3 \rightarrow 2}$, $Tr_{4 \rightarrow 4}$ and $Tr_{5 \rightarrow 6}$ to each edge that is shared by either 3, 4 or 5 tetrahedra. The operators return the optimal configuration with respect to the optimization metric and update the mesh topology. If an edge is part of a boundary, it will not be optimized.
- If the number of topological changes drops below a prescribed minimum, stop the loop; otherwise start a new iteration cycle.

4. Numerical results

The examples presented in this section involve solutions of the incompressible, single-phase flow problem on original and optimized grids. Sections 4.1 and 4.2 investigate the convergence of MPFA on optimized grids for homogeneous \mathbf{k} with strong anisotropy ratios. Section 4.3 shows the application of the technique to a heterogeneous permeability field with a complex global structure but a high degree of local randomization.

Section 4.4 illustrates the coupling of grid generation and upscaling. In Section 4.5 the effect of grid optimization on monotonicity is investigated for a three-dimensional example.

For the convergence studies, convergence rates are estimated for the maximum norm

$$L_{max} = \|P_i - p_i\|_{max} \tag{32}$$

and the mean-square norm

$$L_2 = \left[\sum_{i=1}^n (P_i - p_i)^2 A_i \right]^{1/2} \tag{33}$$

where P_i is the exact solution, p_i is the numerical solution at the location of vertex i , and A_i is the area of the control volume associated with vertex i . In the examples the errors are computed for successively refined grids and the convergence rates q_{max} and q_2 are estimated from the last two refinement steps (last three points).

To investigate the performance of the technique with respect to monotonicity, the inverse coefficient matrices are checked for negative entries. For the three-dimensional case, where monotonicity problems cannot be resolved through pure grid optimization, the solution is checked for the appearance of local minima or maxima. The results obtained from optimized grids are compared to results obtained from the original grids formed by Delaunay triangulation in physical space.

The following examples are based on strong permeability anisotropy ratios which will be quantified by the relation k_1/k_2 , where k_1 and k_2 are the permeability values in the vector space defined by the eigenvectors of the tensor (i.e., the principal values of \mathbf{k}). In the unrotated orientation, k_1 aligns with the x -axis direction and k_2 with the y -axis.

4.1. Homogeneous \mathbf{k} with Dirichlet boundary conditions

This example is a generalization of a problem presented by Crumpton et al. [11]. A constant, symmetric permeability tensor

$$\mathbf{k} = \begin{pmatrix} k_{xx} & k_{xy} \\ k_{xy} & k_{yy} \end{pmatrix}, \tag{34}$$

is defined on a unit square and the problem is solved subject to Dirichlet boundary conditions consistent with the analytical solution given by

$$p = e^{xy}. \tag{35}$$

For a general symmetric tensor the corresponding right-hand side (source term) of the pressure equation is given by

$$Q(x, y) = e^{xy} [k_{yy}x^2 + 2k_{xy}(1 + xy) + k_{xx}y^2]. \tag{36}$$

The numerical tests use a full permeability tensor with an anisotropy ratio of $k_1/k_2 = 100/1$ rotated 60° in the clockwise direction. The grids are based on regular $N \times N$ Cartesian grid point distributions. Results for grids based on Delaunay and optimized triangulations are presented in Table 1. The table provides the number of points in one coordinate direction N , the square root of the mean control-volume area $\langle A \rangle$, results for the maximum and least-square norms, corresponding convergence rates based on the last two refinements, and results for the checks of $\mathbf{A}^{-1} \geq \mathbf{0}$.

For both grid types the maximum and least-square norms indicate approximately second order convergence. Although the results for the original grids are obtained from non-monotone matrices, the corresponding pressure solutions do not show local minima or maxima. Note that the optimized grids in all cases lead to monotone \mathbf{A}^{-1} . For a given N , both errors are larger using optimized grids. Investigation of the error map shows that the maximum errors in the optimized grid occur at the transition from boundary blocks (unaffected by the grid optimization) to the stretched blocks inside the domain. These larger errors may be caused by the abrupt transitions or by the handling of boundary conditions. Although a more detailed analysis would be required to clarify this issue, we expect that the errors with the optimized grids would be significantly reduced in this region if we forced smoother transitions.

Table 1
Convergence rates for original and optimized grids (see Section 4.1)

Grid	N	$\sqrt{\langle A \rangle}$	L_{\max}	L_2	q_{\max}	q_2	\mathbf{A}^{-1} monotone
Original	20	0.050	6.35E – 04	3.21E – 04			No
	30	0.033	2.78E – 04	1.39E – 04			No
	40	0.025	1.55E – 04	7.74E – 05			No
	50	0.020	0.99E – 04	4.92E – 05			No
	60	0.017	0.68E – 04	3.40E – 05	2.01	2.02	No
Optimized	20	0.050	1.80E – 03	4.09E – 04			Yes
	30	0.033	8.02E – 04	1.72E – 04			Yes
	40	0.025	4.56E – 04	9.40E – 05			Yes
	50	0.020	2.94E – 04	5.92E – 05			Yes
	60	0.017	2.05E – 04	4.07E – 05	1.97	2.06	Yes

4.2. Homogeneous \mathbf{k} with localized source term

The problem domain is formed by a regular dodecagon inscribed into the unit square. A source injecting at a constant rate is located at the center of the domain and zero Dirichlet boundary conditions are imposed on the domain boundary. Due to the lack of a simple analytical solution that accounts for permeability anisotropy, a surrogate analytical solution is generated through a fine grid solution on a 401×401 grid defined over the unit square (with blocks outside the dodecagon set to be inactive). The value of the “true” solution on the coarse grid is obtained by bilinearly interpolating the values at the fine grid points that surround a particular coarse grid vertex.

The advantage of this model set-up lies in its radial symmetry. The surrogate solution on the fine grid can be generated using a diagonal tensor aligned with the x – y coordinate axes, which allows the use of a two-point discretization. Any rotation of this tensor in 30° increments, such that the full-tensor eigenvectors align with diagonals of the dodecagon, can be used for the coarse grid coefficients. The coarse grid solution can then be brought into alignment with the surrogate analytical solution through a simple rotation operation.

For the numerical testing, a series of subsequently refined grids were generated with the number of grid vertices n ranging from 317 to 8615. The grid node distributions at subsequent refinement levels include the grid node distributions of previous levels and were generated such that each triangle of the Delaunay triangulation covered approximately the same area, leading to generally unstructured grids. The problem involves a permeability tensor with an anisotropy ratio of $k_1/k_2 = 100/1$ rotated 30° in the counterclockwise direction. An optimized grid with 317 vertices and the reference fine grid solution (solved on a 401×401 grid oriented with the principal directions of \mathbf{k}) are shown in Fig. 10. In order to avoid spurious effects due to the source term (which leads to the solution being unbounded in L_{\max}), the convergence rates were estimated using only vertices located outside a circular area with radius 0.1 around the source. Convergence results are summarized in Table 2 and Fig. 11. Note that the convergence is affected by a certain degree of randomness, which can be attributed to the point insertion process at subsequent refinement levels.

The results on the original grids are obtained from non-monotone matrices, and oscillations appear in the areas where the pressure surface flattens out towards 0. The L_2 errors are significantly lower and the convergence rate higher for the optimized grids than for the original grids. This may be a consequence of non-monotonicity in the solutions on the original grids. The L_{\max} error for both cases shows a low rate of convergence, but the rate for the optimized grids is higher. The maximum error consistently occurs just outside the circular cutout around the source, so these results are likely sensitive to the gridding in that region.

4.3. Heterogeneous permeability field

The domain of this example is the unit square with pressure conditions imposed at the lower left and upper right corner and no-flow conditions imposed elsewhere along the domain boundary. The grid is based on a 16×16 Cartesian point distribution with an initial node spacing of $\Delta x = \Delta y = 1/15$. The location of each node i away from the boundaries is randomly perturbed by $-\Delta x \leq r_{x,i} \leq \Delta x$ in the x -direction and by $-\Delta y \leq r_{y,i} \leq \Delta y$




Fig. 10. Optimized dual grid for Section 4.2 with 317 grid points and pressure contours from surrogate analytical solution.

Table 2
Convergence rates for original and optimized grids (see Section 4.2)

Grid	n	$\sqrt{\langle A \rangle}$	L_{\max}	L_2	q_{\max}	q_2	A^{-1} monotone
Original	317	0.049	3.55	0.35			No
	627	0.035	2.38	0.24			No
	1997	0.019	2.43	0.18			No
	4077	0.014	1.71	0.11			No
	8615	0.0093	1.82	0.094	0.39	0.89	No
Optimized	317	0.049	0.32	0.037			Yes
	627	0.035	0.21	0.017			Yes
	1997	0.019	0.14	0.0089			Yes
	4077	0.014	0.093	0.0054			Yes
	8615	0.0093	0.074	0.0043	0.82	0.99	Yes

in the y -direction. The permeability field is defined in any control volume by a tensor with an anisotropy ratio of $(k_1 + r_k)/k_2$ where $k_1 = 1$, $k_2 = 80$ and r_k is a randomization term such that $-20 \leq r_{k,i} \leq 20$. Each tensor is rotated in the counterclockwise direction by rotation angle α_i which is given by the bilinear function

$$\alpha_i = 60 + 100x_i - 10y_i + 30x_iy_i + r_{\alpha,i}, \tag{37}$$

where $r_{\alpha,i}$ is an angle randomization term such that $-30 \leq r_{\alpha,i} \leq 30$. The resulting tensor field is shown in Fig. 12. A global trend in the orientation can be identified, but a high degree of local randomization also exists.

The inverse of the matrix resulting from the original grid is non-monotone, leading to the severe oscillations in the pressure solution shown in Fig. 13 (lower left). The inverse matrix for the optimized problem is monotone, and the resulting pressure field is free of spurious oscillations (see Fig. 13, lower right). This example vividly illustrates both the problems that can result from the use of simple gridding procedures and the improvement offered by our grid optimization technique.

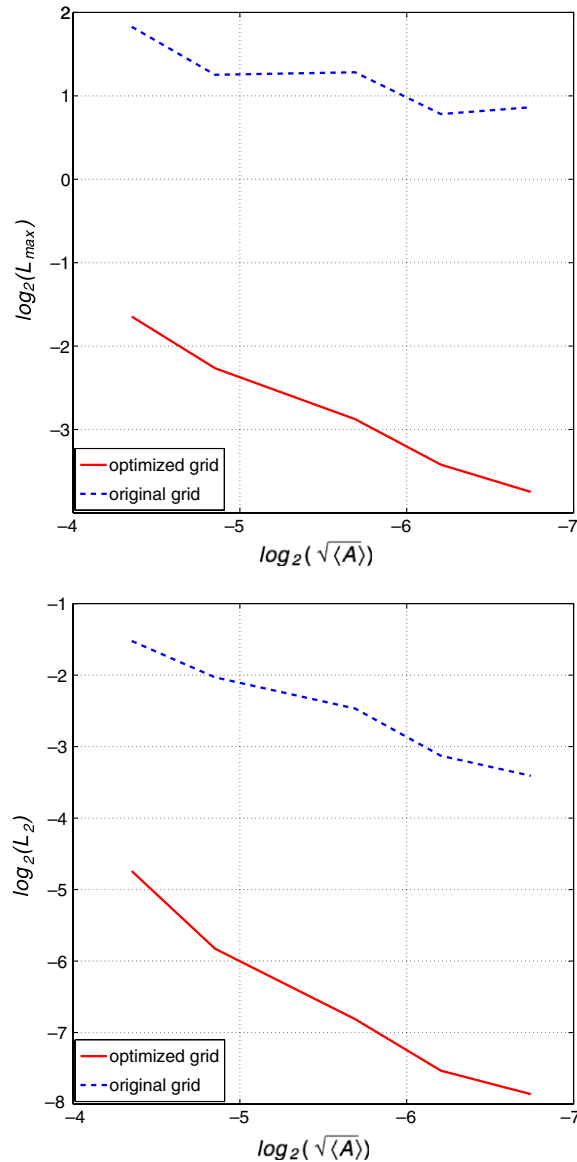


Fig. 11. Convergence of the maximum (top) and L_2 (bottom) errors for Section 4.2.

4.4. Coupling gridding and upscaling

In subsurface flow applications, permeability tensors are not defined by a continuous function. Rather, they result from a numerical procedure that computes each tensor from a high resolution geological model, which is too fine to be used directly for flow simulation. Numerical upscaling techniques proceed by solving local problems over fine scale regions corresponding to control volumes of the flow simulation grid (in some cases bordering regions are also included in these calculations) in order to determine the local permeability tensor. A variety of techniques are available in the literature; here we apply an approach using periodic boundary conditions which was adapted for use with unstructured grids. The particular challenge for the optimization arises from the fact that the tensor is a function of the control-volume geometry, which changes during the grid optimization procedure, requiring an iterative process between upscaling and grid optimization. To keep computational costs low, we compute the upscaled \mathbf{k} based on the initial grid, run the grid optimization until




Fig. 12. Permeability tensor field for Section 4.3. The ellipses indicate the orientation of the permeability tensor at each point. The average anisotropy ratio is 80.

convergence, and then recompute the permeability tensors using the “optimized” grid. The procedure is repeated until consistency between the grid and the permeability tensors is achieved.

We illustrate this procedure with an example modified after Wen et al. [42]. This case involves a two-dimensional permeability field with oriented layers. The fine grid is of dimensions 100×100 and displays layering and orientation of the principal axes of permeability at an angle of 30° relative to the x -axis. The permeability in the x - y coordinate system is therefore a full tensor quantity, with an anisotropy ratio of $k_1/k_2 = 30$ (this level of anisotropy is a factor of three greater than that used in [42]; the permeability fields are otherwise identical). The coarse grid is based on an a priori estimate of the flow field (see [42] for a discussion of flow-based gridding) and consists of approximately 230 control volumes. The fine scale permeability field and the optimized (flow-based) dual grid are shown in Fig. 14. The optimization procedure starts from an initial point distribution (determined from the a priori flow estimate) and repeatedly performs grid optimization and upscaling. Five full optimization-upscaling loops were required for convergence, which was achieved when the number of edge swaps in the optimization procedure dropped below a prescribed minimum (20 for this example). To minimize the computational cost, the algorithm does not necessarily need to run until full convergence; for practical purposes it will often suffice to perform only a single optimization cycle.

For this problem, Dirichlet boundary conditions are specified along the lower and upper domain boundaries and no-flow boundary conditions are imposed elsewhere. The inverse of the matrix for the original unstructured grid is non-monotone, but the pressure solution shows only minor oscillatory behavior. The inverse coefficient matrix for the optimized grid is monotone and the computed pressure is free of spurious oscillations. The grid obtained using only a single optimization cycle was also tested and was found to provide a monotone inverse, demonstrating that the procedure need not be run to full convergence for this case.

The quality of the upscaling-grid optimization procedure is assessed by comparing the pressure solutions obtained from the fine and coarse grids. The pressure surfaces in Fig. 15 show that the overall structure of the fine scale pressure solution is preserved with reasonable accuracy on the coarse grid. The coarse model does however lose some of the fine scale detail, as would be expected. Discrepancies between the two solutions are most evident in regions of high ∇p . This is perhaps not surprising since the grid was formed to resolve high flow regions rather than variations in pressure. In any event, the results in Fig. 15 demonstrate the viability of the combined upscaling-grid optimization procedure.




Fig. 13. Original and optimized primal and dual grids (top row) and pressure surfaces and contours (bottom row) for Section 4.3.

We note that the same model set-up was tested for anisotropy ratios up to 100/1. The inverse of the coefficient matrix for all of the upscaled problems was monotone, while that obtained from the original grids was non-monotone (though the resulting pressure fields showed relatively minor oscillations).

4.5. Three-dimensional corner to corner flow

This example illustrates the impact of grid optimization in \mathbb{R}^3 . Several examples were tested, and in general we observed that the oscillatory behavior of the pressure solution in three dimensions using the original grid was not as strong as that observed in two dimensions. This may be because multiple interaction regions contribute to a given connection and therefore, as discussed in Section 2.3, positive contributions in the coefficient matrix are more likely to be offset by negative contributions from other connections. In the following, we present an example for which the initial grid did lead to relatively strong non-physical effects.

The domain is the unit cube with an initial Cartesian grid point distribution of $n_x \times n_y \times n_z = 12 \times 4 \times 5$. Vertices away from the boundaries are perturbed by a maximum of \pm one third of the grid point spacing




Fig. 14. $\log(k_1)$ of fine scale permeability field and optimized dual grid for Section 4.4.

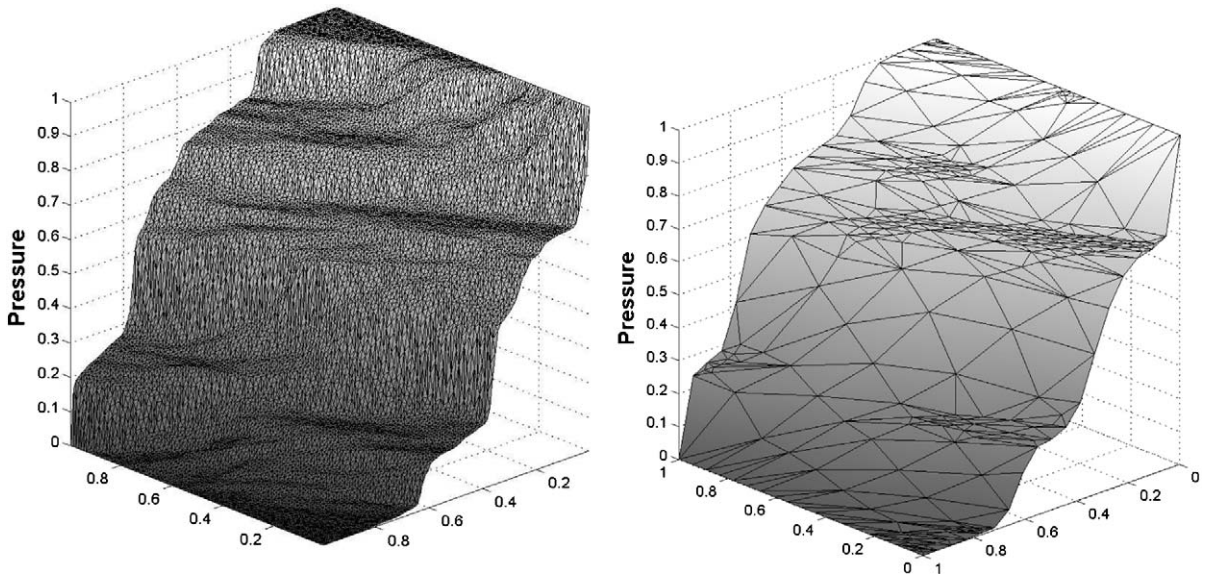


Fig. 15. Pressure surfaces for Section 4.4. The fine grid reference result is shown on the left and the result for the coarse grid is shown on the right.

Table 3
Grid statistics for 3D example (see Section 4.5)

Dihedral angle range θ	Tetrahedra fraction	
	Original grid	Optimized grid
$0 < \theta \leq 90$	0.00	0.08
$90 < \theta \leq 120$	0.11	0.42
$120 < \theta \leq 140$	0.16	0.19
$140 < \theta \leq 160$	0.34	0.19
$160 < \theta \leq 180$	0.40	0.12

in the respective coordinate direction. The permeability field is based on a tensor with $k_1:k_2:k_3 = 100:1:1$ and the tensor rotation angle defined by Eq. (37) after neglecting the randomization term. Table 3 summarizes the statistics of the original and optimized primal grids. The table presents results for the largest dihedral angle for each tetrahedron, computed in a computational space based on the average of the permeability tensors at the four vertices. These results show that the algorithm is effective in reducing very high angles (recall that the optimization seeks to minimize the maximum dihedral angle), although more grid optimization would be desirable. The effect of the grid optimization on a single control volume and its underlying tetrahedrization is shown in Fig. 16. The configurations are plotted in computational space and it is evident that the optimized control volume is much less elongated than the original block. This occurs as a result of the reduction in the maximum dihedral angles of the tetrahedrization.

The problem is solved subject to Dirichlet boundary conditions in the corners $(0,0,z)$ and $(1,1,z)$ applied over the full model thickness and no-flow conditions specified elsewhere along the boundary. Because the grid

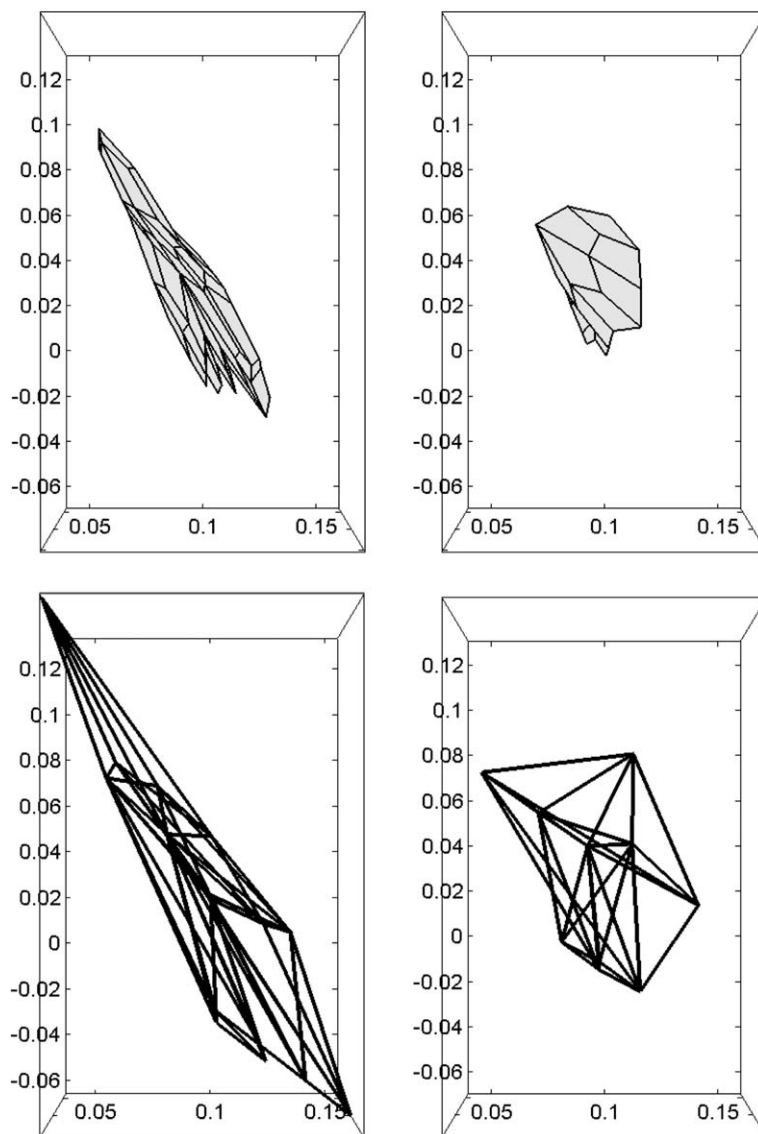


Fig. 16. A single control volume (top row) and its underlying tetrahedrization (bottom row) for Section 4.5 plotted in computational space, viewed along the z -axis. Plots on the left show the configurations before and on the right after optimization.

optimization does not guarantee satisfaction of Eq. (24), the inverse matrices for both the original and optimized problem are non-monotone. Nonetheless, the benefit of the optimized grid is clearly evident in the pressure solution. Specifically, using the original grid, 60 out of a total of 240 blocks show oscillations or exceed the imposed boundary conditions. The solution on the optimized grid, by contrast, does not show any oscillations or block pressures that exceed the boundary conditions.

As an alternative to the dihedral angle based optimization, an algorithm using a solid angle criterion was also investigated for this problem. Although the algorithm led to an improved grid, measured in terms of dihedral angles, it was not sufficient to eliminate the non-physical oscillations in the pressure solution.

5. Conclusions and possible extensions

The purpose of this paper was to develop and apply a grid optimization technique that eliminates the oscillatory behavior of solutions computed using multipoint flux approximations caused by strong anisotropy ratios of the permeability tensors. In Section 2, the relation between monotonicity and grid geometry was investigated and a criterion guaranteeing an M -matrix for three-dimensional problems was derived for the homogeneous case. In Section 3, grid optimization techniques based on the application of topological operators and anisotropic triangulation were introduced for two- and three-dimensional grids. The key feature of the underlying algorithm is its ability to account for variable anisotropy in the permeability field, as results from permeability upscaling procedures. The grid generation techniques introduced here do not require any modification of the underlying MPFA schemes.

The application of the method for improving the monotonicity performance of MPFA when solving elliptic problems was demonstrated in Section 4. Convergence for the MPFA discretization on unstructured grids was established numerically and examples illustrating the applicability of the technique for heterogeneous cases were presented. A methodology that couples grid optimization and upscaling was also developed and applied.

A number of extensions of the grid optimization techniques presented in this paper would be of use. In the current algorithms, efficient data structures for meshing were not of primary concern, so computational efficiency issues must be addressed and appropriate data structures implemented. The extension of the grid optimization through the introduction of geometric operators that add vertices and the application of local grid smoothing to further improve the grid quality, both accounting for heterogeneity and anisotropy, would also be useful. For three-dimensional grids, where our algorithm can only alleviate oscillations in the pressure solutions, a more advanced triangulation or triangulation optimization technique accounting for a dihedral angle criterion and variable anisotropy would be a major step forward. An algorithm that links upscaling and grid optimization was presented in \mathbb{R}^2 . A fast and accurate three-dimensional implementation of this coupled procedure is currently under development.

Angular restrictions are common to many discretization schemes that are based on triangular meshes. Therefore, the grid optimization technique introduced in this paper may prove useful for schemes other than the MPFA method considered here. It is possible that a hybrid solution incorporating both a modified MPFA stencil and grid optimization will provide a very general and robust overall methodology.

Acknowledgments

We are grateful to Chevron Energy Technology Company for partial funding of this work and to W.H. Chen for facilitating this support. We also thank G.T. Eigestad, R. Juanes, B. Mallison and H.A. Tchelepi for many stimulating discussions and helpful suggestions.

References

- [1] I. Aavatsmark, An introduction to multipoint flux approximations for quadrilateral grids, *Comput. Geosci.* 6 (2002) 405–432.
- [2] I. Aavatsmark, T. Barkve, Ø. Bøe, T. Mannseth, Discretization on non-orthogonal, quadrilateral grids for inhomogeneous, anisotropic media, *J. Comput. Phys.* 127 (1996) 2–14.
- [3] I. Aavatsmark, T. Barkve, Ø. Bøe, T. Mannseth, Discretization on unstructured grids for inhomogeneous anisotropic media. Part I: derivation of the methods, *SIAM J. Sci. Comput.* 18 (1998) 1700–1716.

- [4] I. Aavatsmark, T. Barkve, Ø. Bøe, T. Mannseth, Discretization on unstructured grids for inhomogeneous, anisotropic media. Part II: discussion and numerical results, *SIAM J. Sci. Comput.* 18 (1998) 1717–1736.
- [5] T. Arbogast, C.N. Dawson, P.T. Keenan, M.F. Wheeler, I. Yotov, Enhanced cell-centered finite differences for elliptic equations on general geometry, *SIAM J. Sci. Comput.* 19 (1998) 404–428.
- [6] J. Bear, *Dynamics of Fluids in Porous Media*, Elsevier, New York, 1972.
- [7] M. Bern, L. Chew, D. Eppstein, J. Ruppert, Dihedral bounds for mesh generation in high dimensions, in: *Proceedings of the 6th Symposium on Discrete Algorithms*, ACM and SIAM, 1995, pp. 189–196.
- [8] H. Borouchaki, P.L. George, P. Hecht, P. Laug, E. Saletl, Delaunay mesh generation governed by metric specification: Part I. *Algorithms, Finite Elem. Anal. Des.* 25 (1997) 6183.
- [9] G.F. Carey, Hexing the tet, *Commun. Numer. Meth. Eng.* 18 (2002) 223–227.
- [10] M.J. Castro-Diaz, F. Hecht, B. Mohammadi, O. Pironneau, Anisotropic unstructured mesh adaption for flow simulations, *Int. J. Numer. Meth. Fluids* 25 (1997) 475–491.
- [11] P.I. Crumpton, G.J. Shaw, A.F. Ware, Discretization and multigrid solution of elliptic equations with mixed derivative terms and strongly discontinuous coefficients, *J. Comput. Phys.* 116 (1995) 343–358.
- [12] L.J. Durlofsky, A triangle based mixed-finite element-finite volume technique for modeling two phase flow through porous media, *J. Comput. Phys.* 105 (1993) 252–266.
- [13] H. Edelsbrunner, *Geometry and Topology for Mesh Generation*, Cambridge University Press, Cambridge, 2001.
- [14] M.G. Edwards, Unstructured, control-volume distributed, full-tensor finite-volume schemes with flow based grids, *Comput. Geosci.* 6 (2002) 433–452.
- [15] M.G. Edwards, C.F. Rogers, Finite volume discretization with imposed flux continuity for the general tensor pressure equation, *Comput. Geosci.* 2 (1998) 259–290.
- [16] G.T. Eigestad, I. Aavatsmark, M. Espedal, Symmetry and M -matrix issues for the O -method on an unstructured grid, *Comput. Geosci.* 6 (2002) 381–404.
- [17] P.A. Forsyth, A control volume finite element approach to NAPL groundwater contamination, *SIAM J. Sci. Stat. Comput.* 12 (1991) 514–539.
- [18] L.A. Freitag, C. Ollivier-Gooch, Tetrahedral mesh improvement using swapping and smoothing, *Int. J. Numer. Methods Eng.* 40 (1997) 3979–4002.
- [19] L.S.-K. Fung, A.D. Hiebert, L.X. Nghiem, Reservoir simulation with a control-volume finite-element method, *SPE Res. Eng.* 7 (1992) 349–357.
- [20] P.L. George, H. Borouchaki, *Delaunay Triangulation and Meshing*, HERMES, Paris, 1998.
- [21] Z.E. Heinemann, C.W. Brand, M. Munka, Y.M. Chen, Modeling reservoir geometry with irregular grids, *SPE Res. Eng.* 6 (1991) 225–232.
- [22] W. Huang, Metric tensors for anisotropic mesh generation, *J. Comput. Phys.* 204 (2005) 633–665.
- [23] J. Hyman, J. Morel, M. Shashkov, S. Steinberg, Mimetic finite difference methods for diffusion equations, *Comput. Geosci.* 6 (2002) 333–352.
- [24] J. Hyman, M. Shashkov, S. Steinberg, The numerical solution of diffusion problems in strongly heterogeneous non-isotropic materials, *J. Comput. Phys.* 132 (1997) 1029–1057.
- [25] E.B. de l’Isle, P.L. George, Optimization of tetrahedral meshes, in: I. Babushka, W.D. Henshaw, J.E. Olinger, J.E. Flaherty, J.E. Hopcroft, T. Tezduyar (Eds.), *Modeling, Mesh Generation, and Adaptive Numerical Methods for Partial Differential Equations*, Springer, Berlin, 1995.
- [26] B. Joe, Construction of three-dimensional improved-quality triangulations using local transformations, *SIAM J. Sci. Comput.* 16 (1995) 1292–1307.
- [27] R.A. Klausen, T.F. Russell, Relationships among some locally conservative discretization methods which handle discontinuous coefficients, *Comput. Geosci.* 8 (2005) 341–377.
- [28] C.L. Lawson, Software for C^1 surface interpolation, in: John D. Rice (Ed.), *Mathematical Software III*, Academic Press, New York, 1977.
- [29] S.H. Lee, H.A. Tchelep, P. Jenny, L.J. DeChant, Implementation of a flux-continuous finite-difference method for stratigraphic hexahedron grids, *SPE J.* 7 (2002) 267–277.
- [30] W. Letniowski, Three-dimensional Delaunay triangulations for finite element approximations to a second-order diffusion operator, *SIAM J. Sci. Stat. Comput.* 13 (1992) 514–539.
- [31] D.J. Mavriplis, Adaptive mesh generation for viscous flows using Delaunay triangulation, *J. Comput. Phys.* 90 (1990) 271–291.
- [32] J.M. Norbotten, I. Aavatsmark, Monotonicity conditions for control volume methods in uniform parallelogram grids in homogeneous media, *Comput. Geosci.* 9 (2005) 61–72.
- [33] J.M. Norbotten, G.T. Eigestad, Discretization on quadrilateral grids with improved monotonicity properties, *J. Comput. Phys.* 203 (2005) 744–760.
- [34] J. Peraire, J. Peiró, K. Morgan, Adaptive remeshing for three-dimensional compressible flow computations, *J. Comput. Phys.* 103 (1992) 269–285.
- [35] N.B. Petrovskaya, Modification of a finite volume scheme for Laplace’s equation, *SIAM J. Sci. Comput.* 23 (2001) 891–909.
- [36] M. Prevost, *Accurate coarse reservoir modeling using unstructured grids, flow-based upscaling and streamline simulation*, Ph.D. Thesis, Stanford University, 2003.
- [37] M. Putti, C. Cordes, Finite element approximation of the diffusion operator on tetrahedra, *SIAM J. Sci. Comput.* 19 (1998) 1154–1168.

- [38] M. Shashkov, S. Steinberg, Solving diffusion equations with rough coefficients in rough grids, *J. Comput. Phys.* 129 (1996) 383–405.
- [39] J. Shewchuk, Two discrete algorithms for the topological improvement of tetrahedral meshes, manuscript, 2002. Available from: <http://www-2.cs.cmu.edu/jrs/jrspapers.html>.
- [40] R. Varga, *Matrix Iterative Analysis*, Prentice-Hall, New York, 1963.
- [41] S. Verma, K. Aziz, A control volume scheme for flexible grids in reservoir simulation, SPE Paper 37999, in: *Proceedings of 14th SPE Symposium on Reservoir Simulation* held in Dallas, TX, 1997, pp. 215–227.
- [42] X.H. Wen, L.J. Durlofsky, M.G. Edwards, Upscaling of channel systems in two dimensions using flow-based grids, *Transport Porous Media* 51 (2003) 343–366.
- [43] J.A. Wheeler, M.F. Wheeler, I. Yotov, Enhanced velocity mixed finite element methods for flow in multiblock domains, *Comput. Geosci.* 6 (2002) 381–404.
- [44] J. Xu, L. Zikatanov, A monotone finite element scheme for convection–diffusion equations, *Math. Comp.* 68 (1999) 1429–1446.