

A Fully Implicit WENO Scheme on Stratigraphic and Unstructured Polyhedral Grids

Knut-Andreas Lie · Trine S. Mykkeltvedt · Olav Møyner

Received: date / Accepted: date

Abstract Many authors have used higher-order spatial discretizations to reduce numerical diffusion, which can be particularly pronounced when simulating EOR processes involving active chemical substances that are transported by linear or weakly nonlinear waves. Most high-resolution methods reported in the literature are based on explicit temporal discretizations. This imposes severe time-step restrictions when applied to the type of grids seen in industry-standard simulation models of real assets, which usually have orders-of-magnitude variations in porosities and Darcy velocities that necessitate the use of implicit discretization. Herein, we propose a second-order WENO discretization suitable for complex grids with polyhedral cell geometries, unstructured topologies, large aspect ratios, and large variations in interface areas. The WENO scheme is developed as part of a standard, fully implicit formulation that solves for pressure and transported quantities simultaneously. We investigate the accuracy and utility of the WENO scheme for a series of test cases that involve corner-point and 2D/3D Voronoi grids and black-oil and compositional flow models

Keywords Reservoir simulation · WENO scheme · Fully implicit discretization · Unstructured grids · Corner-point grids

Knut-Andreas Lie
SINTEF Digital, Mathematics and Cybernetics, P.O. Box 124 Blindern, N-0314 Oslo, Norway
E-mail: knut-andreas.lie@sintef.no

Trine S. Mykkeltvedt
NORCE Energy, Nygårdssporten 112, N-5008 Bergen, Norway
E-mail: tmyk@norceenergy.no

Olav Møyner
SINTEF Digital, Mathematics and Cybernetics, P.O. Box 124 Blindern, N-0314 Oslo, Norway
E-mail: olav.moyner@sintef.no

1 Introduction

The workhorse in reservoir simulation is a first-order finite-volume method with implicit temporal discretization and intercell fluxes computed by a single-point, upstream-mobility weighting scheme. The method is flexible and robust, but can suffer from (severe) grid-orientation effects and excessive numerical diffusion. Numerical diffusion can be particularly detrimental in compositional simulations, which tend to contain many displacement fronts that can be difficult to distinguish when smeared out, and in simulation of chemical EOR, in which the active chemical components propagate as linear or weakly linear waves that are particularly susceptible to numerical diffusion.

Many authors have proposed the use of high-resolution spatial discretizations reduce numerical smearing, see, e.g., [4, 5, 3, 20, 12, 8, 28, 35, 13, 21, 27, 40, 42, 24, 9, 55]. TVD methods based on flux- and slope-limiter approaches dating back to the early work of [56] and (W)ENO reconstructions both rely on local polynomial reconstructions computed from the cell average of each grid cells and its adjacent cell neighbors and are designed to maintain high-order accuracy on smooth parts of the solution and at the same time minimize the creation of spurious oscillations around discontinuities. Such methods are readily applicable to Cartesian grids and similar grids with structured topology, but have also been extended to unstructured grids consisting of simplices (triangles and tetrahedrons) or prismatic elements, e.g., as discussed in [26, 61] and references therein. Another approach is to use discontinuous Galerkin [51, 22, 23, 47] and similar methods in which the higher-order approximation relies entirely on unknowns localized inside each cell.

Except for [52], little work has been devoted to develop and analyze high-resolution methods for the type of grids found in contemporary reservoir models, which are often characterized by large aspect ratios, large differences in cell

volumes, nonmatching cells, and various types of degenerate cell geometries. In particular, their extension to fully unstructured grids with general polyhedral cell geometries is unclear.

The widely used corner-point format [50] was introduced to represent the stratigraphy and structure of reservoirs with high accuracy. The stratigraphy of the rock is represented through a rectangular ijk topology, in which cells with the same k index typically represent rock units that have been deposited in the same period. Each grid cell is by default a hexahedral volume delimited by its eight corner points. The eight corner points are placed in pairs of two on four vertical/inclined coordinate lines that extend downward from a rectilinear (or curvilinear) areal mesh. Each pair of points specifies the depth of the top and bottom surface of the grid along a given coordinate line and can collapse to a single point to model erosion of the deposited sediments, effectively leading to various forms of degenerate cell geometries. Corner points of lateral cell neighbors (i.e., same i and k or j and k index) are specified on the same coordinate lines, but need not coincide. This is used to model fault displacement and generally gives non-matching cell faces and an unstructured topology in which each cell may have (significantly) more than four lateral neighbors. This, combined with large aspect ratios and orders-of-magnitude variations in the area of the cell interface shared by neighboring cells, can cause severe difficulties for spatial reconstruction methods (slope-limiter, WENO, etc) that have been developed for simpler geometries and topologies.

Corner-point grids constitute just one example of the many grid formats used in reservoir simulation. Another popular format is the perpendicular bisector (PEBI) grids [19, 49, 57], which in other fields of science are known as Voronoi grids. The stratigraphic or 2.5D form of these grids is constructed along the same lines as corner-point grids by extrusion along vertical/inclined coordinate lines, except that these coordinate lines now extend downward from the vertices of an areal Voronoi mesh and not from a rectilinear mesh. Use of polygonal tessellations offers improved resolution control in the lateral direction, but 2.5D PEBI grids have more lateral neighbors that must be accounted for in a spatial reconstruction and pose the same challenges as corner-point grids in the vertical direction. Several methods have also been proposed to construct truly 3D PEBI grids adapting to various types of geological objects and curvilinear well paths [43, 29]. Such grids will have a fully unstructured topology and general polyhedral cell geometries.

Another example of unstructured formats are the so-called cut-cell grids [39], which consist of hexahedral, highly orthogonal cells that are arranged in a structured manner away from faults. Near faults, the hexahedral cells have been clipped against the triangulated fault surfaces and thus are converted into general polyhedra. If the clipping surfaces represent

sealing faults, one can argue that having an accurate higher-order reconstruction may not be very important since flow will mostly be stagnant. In other cases, however, faults or fractures may be the main flow conducts and having accurate spatial reconstructions in the polyhedral cells may be essential to accurately resolve the flow.

A second challenge is the temporal discretization. With a few exceptions [7, 6, 36, 48, 14, 47], high-resolution schemes applied in reservoir simulation have primarily relied on explicit temporal discretizations. This requires that the multiphase flow equations are solved using a sequential solution procedure that computes the flow (pressure and fluxes) and transport of saturations and component concentrations in separate steps. Coupling of between fluid pressure and transport of phases and components can be strong, e.g., in systems with significant compressibility, and sequential solution procedures are generally not as robust as methods that seek to solve for all primary variables simultaneously. More important, sector and field-scale models tend to have large variations in time constants arising because of high local flow rates in the near-well zone, cells with small pore volumes, etc. This means that the well-known CFL condition will impose severe time-step restrictions that quickly can render an explicit high-resolution scheme computationally infeasible. [18] and [62] showed that implicit five-point TVD schemes are conservative and unconditionally stable for a scalar equation in 1D if the discrete nonlinear equations are solved exactly, whereas [16] showed that implicit time-integration schemes of order higher than one are only conditionally TVD. To be efficient, the cost of solving the nonlinear equations, e.g., by a Newton method, must be offset with the ability to take larger time steps.

Herein, we develop weighted essentially nonoscillatory (WENO) high-resolution schemes for all the classes of unstructured grids discussed above, formulated so that they can be applied in a fully implicit or a sequentially implicit setting. WENO schemes compute a set of local polygonal interpolations with accompanying nonlinear smoothness indicators that are used to compute a local reconstruction in each cell that introduces as few spurious oscillations as possible. The development and application of WENO schemes on unstructured grids is still ongoing, see, e.g., [15, 54, 33, 58, 59, 37]. In previous research [46], we have developed fully implicit WENO schemes and slope-limiter schemes for rectilinear grids and showed how one can easily overcome the cumbersome task of linearizing the discrete flow equations and computing the Jacobian matrix necessary in a Newton-type nonlinear solver by use of automatic differentiation. To apply these schemes to realistic reservoir models, we can reuse most of the same ideas, but need to develop compact and effective sets of local polynomials and nonlinear interpolation weights that are sufficiently robust to large aspect ratios, significant variations in cell sizes and the number of

cell neighbors, and various forms of geometrical degeneracies.

In the following, we restrict ourselves to demonstrating a proof-of-concept in MRST [34, 45] and validating the new WENO schemes by applying them to various types of black-oil and compositional problems posed on a variety of grid types that are representative of contemporary models of real petroleum assets. Local WENO reconstructions can be computed in many different ways. Herein, we only consider a simple pragmatic choice that gives a relatively small stencil and seems to work well on representative grids. A deeper analysis is needed to assess the level of grid-orientation effects (for adverse mobility displacements), investigate and compare various strategies for reducing the local stencil (e.g., by setting the weights of the least smooth polynomials to zero), and compare the efficacy of WENO schemes to slope-limiter schemes. Likewise, we only consider cases with a single rock type (i.e., a single set of relative permeability and capillary pressure curves), so that we avoid the problem of interpolating saturations between regions with different capillary pressure curves.

2 Governing equations

In reservoir simulation, one is generally interested in flow systems consisting of N fluid phases that may contain M different components. Each component can either refer to a single chemical species or be a pseudo-component that consists of a collection of different chemical species that are lumped together and assumed to have a distinct behavior. For simplicity, we will disregard diffusion. The mass conservation of component $\ell = 1, \dots, M$ then reads

$$\frac{\partial}{\partial t} \left(\phi \sum_{\alpha} c_{\alpha}^{\ell} \rho_{\alpha} S_{\alpha} \right) + \nabla \cdot \left(\sum_{\alpha} c_{\alpha}^{\ell} \rho_{\alpha} \mathbf{v}_{\alpha}^{\ell} \right) = \sum_{\alpha} c_{\alpha}^{\ell} \rho_{\alpha} q_{\alpha}. \quad (1)$$

Here, ϕ is rock porosity; S_{α} , ρ_{α} , and q_{α} denote the saturation, density, and source term of fluid phase α ; and c_{α}^{ℓ} is the mass fraction and $\mathbf{v}_{\alpha}^{\ell}$ the superficial velocity of component ℓ in phase α . The velocities are given by Darcy's law,

$$\mathbf{v}_{\alpha}^{\ell} = \frac{K k_{r\alpha}}{\mu_{\alpha}^{\ell}} (\nabla p_{\alpha} - g \rho_{\alpha} \nabla z), \quad (2)$$

where K is the absolute permeability; p_{α} is pressure and $k_{r\alpha}$ the relative permeability of phase α ; μ_{α}^{ℓ} is the effective viscosity of component ℓ in phase α ; g is gravity acceleration; and z the vertical coordinate. The fluid phases are assumed to fill the void space completely, so that $\sum_{\alpha} S_{\alpha} = 1$. In addition, we need closure relationships for the phase densities, mass fractions, and phase pressures, as well as models for the relative permeabilities and effective viscosities.

The standard black-oil equations describe a system consisting of three phases (an aqueous, an oleic, and a gas phase) and three pseudo-components (water, oil, and gas). At surface conditions, oil is only found in the oleic phase and gas only in the gaseous phase. At reservoir conditions, however, oil can be vaporized in the gaseous phase, and gas can be dissolved in the oleic phase. To describe the fluid behavior, one uses a relatively simple PVT model that consists of pressure-dependent shrinkage/expansion factors $b^{\ell} = V_s^{\ell}/V^{\ell}$ that relate the volume V_s^{ℓ} of component ℓ at surface condition to the volume V^{ℓ} at reservoir conditions. Solubility of gas in oil is modelled through the solution gas-oil ratio, $R_s = V_{gs}/V_{os}$ defined as the volume of gas, measured at standard conditions, that at reservoir conditions is dissolved in a unit of stock-tank oil. The solubility of oil in gas is modeled similarly by a factor R_v , defined as the amount of surface oil that can be vaporized in a unit volume of surface gas at reservoir conditions. Phase pressures are related through saturation-dependent capillary pressure functions, $p_o - p_w = P_{cow}(S_w, S_o)$ and $p_g - p_o = P_{cgo}(S_o, S_g)$. The effective viscosities are the same for all components within each phase and are uniquely given by pressure. There are several possible choices for primary unknowns; herein, we use pressure of the oleic phase p_o , water saturation S_w , and gas saturation S_g when all three phases are present, and R_s if all gas is dissolved or R_v if all oil is vaporized.

As an example of enhanced oil recovery, we also consider a basic model for polymer flooding, which is a model consisting of a single-component oleic phase and an aqueous phase that contains a mixture of water and dissolved polymer. Here, the primary unknowns are (oil) pressure p , water saturation S_w , and polymer concentration c . Effective viscosities for water and polymer are given by a Todd-Longstaff mixture rule. The model also contains an additional accumulation term that accounts for adsorption of polymer onto the reservoir rock, and reduced effective permeability of long-chained polymer molecules, which makes K depend on c . Full details are given in [46].

Compositional models describe the same general three-phase system as the standard black-oil equations, but allow the oleic and gaseous phases to consist of any number of hydrocarbon components. Mass exchange between the two hydrocarbon phases is governed by the isofugacity relation for each component ($f_g^l = f_o^l$) if both phases are present. MRST uses a generalized cubic equation-of-state [41], and in the following we use Peng-Robinson for the oleic and gaseous phases with viscosities calculated from the correlation in [38]. The aqueous phase is described as in the black-oil case and only consists of the water component; see [44] for a detailed description of the implementation of natural variables in MRST.

3 Discretizations and solution method

We start by subdividing the spatial domain Ω into a set of finite volumes (cells) Ω_i . In a stratigraphic grid, the basic shape of these cells will typically be a hexahedron or a triangular or hexahedral prism, and each cell may potentially have nonmatching faces with its neighbors. Herein, we subdivide nonmatching faces and assume that all grids are fully unstructured and consist of cells with general polyhedral geometry. Let $\mathcal{N}(i)$ denote the indices of all neighbors of cell i , and let the interface Γ_{ij} between two neighboring cells i and j have normal vector \mathbf{n}_{ij} pointing from Ω_i to Ω_j .

We use a finite-volume discretization so that each unknown quantity u is represented as discrete cell averages,

$$u_i(t) = \frac{1}{|\Omega_i|} \iint_{\Omega_i} u(\mathbf{x}, t) d\mathbf{x}. \quad (3)$$

For brevity, we only present the discretization for a simple two-phase, two-component system ($\alpha = \{w, n\}$, $c_w^1 = c_n^2 = 1$, and $c_w^2 = c_n^1 = 0$) in the absence of gravity and capillary forces. Picking the wetting phase, integrating (1) over cell Ω_i from time t_n to t_{n+1} , and inserting (2) gives the flow equation on residual form,

$$\begin{aligned} \mathcal{R}_w = & [\rho_w \phi S_w]^{n+1} - [\rho_w \phi S_w]^n \\ & + \frac{\Delta t}{|\Omega_i|} \sum_{j \in \mathcal{N}(i)} \int_{\Gamma_{ij}} (\rho_w \lambda_w K \nabla p \cdot \mathbf{n})_{ij}^m ds = 0. \end{aligned} \quad (4)$$

Setting $m = n$ gives an explicit scheme whereas $m = n + 1$ gives a fully implicit scheme. The two accumulation terms can be computed directly from cell-averaged quantities if we approximate each term by a product of cell averages. The flux integral is more difficult, since any numerical quadrature rule will require *point values* of the unknown quantities along Γ_{ij} . Herein, we only consider schemes of order two or less and it is hence sufficient to apply the midpoint rule. For the density at the interface, we simply use the arithmetic average of the cell averages $\rho_{ij} = \frac{1}{2}(\rho_i + \rho_j)$ in the case of immiscible flow. For compositional models, phase properties use the saturation-weighted average to account for the possibility of an absent phase: $(\rho_{ij})_\alpha = \left(\frac{s_i \rho_i + s_j \rho_j}{s_i + s_j} \right)_\alpha$. For the gradient term, we use a standard two-point flux approximation:

$$(K \nabla p \cdot \mathbf{n})_{ij} = \frac{p_i - p_j}{T_{i,j}^{-1} + T_{j,i}^{-1}}, \quad (5)$$

$$T_{i,j} = \frac{K_i(\mathbf{x}_{ij} - \mathbf{x}_i) \cdot \mathbf{n}_{ij}}{|\mathbf{x}_{ij} - \mathbf{x}_i|^2}, \quad T_{j,i} = \frac{K_j(\mathbf{x}_{ij} - \mathbf{x}_j) \cdot \mathbf{n}_{ji}}{|\mathbf{x}_{ij} - \mathbf{x}_j|^2},$$

where \mathbf{x}_i , \mathbf{x}_j , and \mathbf{x}_{ij} denote the centroids of Ω_i , Ω_j , and Γ_{ij} , respectively, and $\mathbf{n}_{ji} = -\mathbf{n}_{ij}$. The difference between first and second-order schemes lies in how we compute the mobility term, $\lambda_w = k_{rw}/\mu_w$, which governs how the flux depends upon saturations (and component concentrations).

For a first-order scheme, we *reconstruct* point values by assuming that the mobility is constant inside each cell and can be computed from the cell-averaged saturation (and component concentration) values. This gives two values at the midpoint, a value λ^- reconstructed inside the cell the normal vector \mathbf{n} is pointing *from*, and a value λ^+ reconstructed inside the cell the normal vector is pointing *to*. Given these one-sided point-values, we use the standard upstream method to evaluate the integrand at each integration point,

$$\lambda_{ij} = \begin{cases} \lambda^-, & \text{if } (K \nabla p \cdot \mathbf{n})_{ij} \geq 0, \\ \lambda^+, & \text{otherwise.} \end{cases} \quad (6)$$

Higher accuracy is achieved if we use a higher-order reconstruction of the point values. To this end, we can either reconstruct point-values for the primary variables and evaluate one-sided mobilities at each interface, or we can first compute mobilities from cell averages and use these ‘‘cell-averaged’’ values to reconstruct point values. To not distinguish between the two, the following discussion considers the reconstruction of point values \hat{u} from a set of cell averages u_i . To keep the presentation as simple as possible, we start by outlining the basic concepts in 2D.

3.1 Local polynomial reconstruction

To obtain a second-order reconstruction, we start by forming linear planes that interpolate the cell averages in \mathbf{x}_i and at the centroids for any two neighbors from $\mathcal{N}(i)$. For a cell with N neighbors, it is possible to construct $\binom{N}{2}$ different planes. Figure 1 shows all stencils that can be used to compute local interpolation planes for a cell with five faces.




Fig. 1 Example of a cell with five neighbors with a common face and the ten corresponding stencils. Red dots indicate the cell centers and blue dots indicate faces centers (used in the two-point flux approximation). Primary stencils are formed from cells that share a common vertex, whereas ancillary stencils involve three cells that do not share a single vertex.

In practice, one only chooses a subset $k = \{1, 2, \dots, N_j\}$ of the possible stencils to form the basis of the reconstruction inside each cell Ω_i . As an example, we could use the five primary stencils in Figure 1 and disregard the five ancillary ones. We will return to this discussion later for 3D stratigraphic grids.

To compute the local interpolation corresponding to stencil k , we let $\mathbf{x}_\ell^k = (x_\ell^k, y_\ell^k)$ and u_ℓ^k denote the centroids and the cell averages of the corresponding three cells. If we introduce barycentric coordinates, the local plane is constructed as follows:

$$\hat{u}^k(x, y) = \underbrace{[u_1^k \ u_2^k \ u_3^k] \begin{bmatrix} x_1^k & x_2^k & x_3^k \\ y_1^k & y_2^k & y_3^k \\ 1 & 1 & 1 \end{bmatrix}^{-1}}_C \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}. \quad (7)$$

One possibility to get a multidimensional slope-limiter method would now be to use the minimum-angle-plane reconstruction [11], which picks the local plane having the minimum magnitude gradient. See [30] for an up-to-date discussion of alternative slope-limiter reconstructions.

3.2 General WENO reconstruction in 2D

The next step is to write the reconstruction as a convex combination of the local interpolation planes

$$\hat{u}_i(\mathbf{x}) = \sum_{k=1}^{N_i} w_i^k \hat{u}_i^k(\mathbf{x}). \quad (8)$$

To get optimal order of reconstruction, the weights w_i^k should be chosen so that the overall polynomial has the same formal order on smooth data as we would get if we had used the $N_i + 1$ data points to directly form a single polynomial. For a fully unstructured grid with general polyhedral cells, these weights will have to be computed uniquely for each unique cell geometry.

In the WENO reconstruction, the linear weights are replaced by nonlinear weights that try to put less emphasis on nonsmooth parts of the solution. These weights are defined as follows,

$$w_i^k = \beta_i^k / \sum_{k=1}^{N_i} \beta_i^k, \quad \beta_i^k = \gamma_i^k / (\varepsilon + IS_i^k)^2. \quad (9)$$

Here, γ_i^k are linear weights that sum to unity and ε is a small positive parameter to avoid division by zero (herein: $\varepsilon = 10^{-7}$ or $\varepsilon = 10^{-12}$). The smoothness indicator IS_i^k measures how smooth the k th local polynomial $\hat{u}_i^k(\mathbf{x})$ is on the cell [25]; the smaller the smoothness indicator, the smoother the function \hat{u}_i^k is on Ω_i . The general expression for an m th order polynomial on an unstructured grid is given by [37] as (here $\eta = (\eta_1, \eta_2)$ is a multi index)

$$IS_i^k = \sum_{1 \leq |\eta| \leq m} \int_{\Omega_i} |\Omega_i|^{|\eta|-1} (D^\eta \hat{u}_i^k(\mathbf{x}))^2 d\mathbf{x}, \quad (10)$$

$$D^\eta \hat{u}_i^k(\mathbf{x}) = \frac{\partial^{|\eta|} \hat{u}_i^k(x, y)}{\partial x^{\eta_1} \partial y^{\eta_2}}.$$

This formula applies to meshes with uniform cell sizes. In 3D, the scaling factor is $|\Omega_i|^{2|\eta|/3-1}$. The purpose of the scaling factor is to make the smoothness indicator invariant under spatial scaling [54].

For our linear polynomials ($m = 1$) and the smoothness indicator for cell Ω_i simplifies to:

$$IS_i^k = \int_{\Omega_i} \left((D^{(1,0)} \hat{u}_i^k(\mathbf{x}))^2 + (D^{(0,1)} \hat{u}_i^k(\mathbf{x}))^2 \right) d\mathbf{x} \quad (11)$$

$$= \int_{\Omega_i} |\nabla \hat{u}_i^k(\mathbf{x})|^2 d\mathbf{x}.$$

This gradient is quick to compute. Referring back to (7) and let \tilde{C} denote the first two columns of the inverse coordinate matrix C , the local gradient is given as a simple vector-matrix product

$$\boldsymbol{\sigma}_i^k = (\nabla \hat{u}_i^k(\mathbf{x}))^T = [u_1^k \ u_2^k \ u_3^k] \tilde{C}. \quad (12)$$

With this, the weights are

$$w_i^k = \beta_i^k / \sum_{k=1}^{N_i} \beta_i^k, \quad \beta_i^k = \gamma_i^k / (\varepsilon + |\boldsymbol{\sigma}_i^k|^2 |\Omega_i|)^2. \quad (13)$$

Using that the weights sum to unity, the reconstruction of the one-sided point value u_{ij}^- at the centroid \mathbf{x}_{ij} of Γ_{ij} is another simple vector-vector product

$$u_{ij}^- = \sum_{k=1}^{N_i} w_i^k (u_i + \boldsymbol{\sigma}_i^k (\mathbf{x}_{ij} - \mathbf{x}_i)) \quad (14)$$

$$= u_i + \left(\sum_{k=1}^{N_i} w_i^k \boldsymbol{\sigma}_i^k \right) (\mathbf{x}_{ij} - \mathbf{x}_i).$$

3.3 Changes required by complex stratigraphic grids

The construction is analogous in 3D with the obvious modifications necessary for four points and coordinate/gradient vectors with three components and works robustly as long as the polygonal grid is sufficiently regular. In the following we will discuss in detail some of the extra precautions and modifications we have found necessary to ensure robust reconstructions for the type of complex grids seen in industry-standard models of petroleum reservoirs. To illustrate a few of the typical issues, we use the grid geometry from the simulation model of Norne as an example, see github.com/OPM/opm-data. There are several other public data sets that contain similar complexities, but we chose Norne since it is the only one that has been made specifically to represent a real asset.

By default, each cell in a corner-point grid is hexahedral and thus has six logical neighbors. Exceptions occur for cells that are adjacent to external boundaries, adjacent




Fig. 2 Illustration of cell geometry and local topology for a stratigraphic corner-point grid, here shown for two cells from the Norne simulation model.




Fig. 3 Illustration of the complex geometry and topology arising near faults in real reservoir models, here represented by two cells from the Norne simulation model. The left figure shows a cell at the top of the reservoir that lies adjacent to three fault faces. Altogether, the cell shares faces with twenty other cells, which together with the external top boundary means that the cell has 21 unique faces. The right figure shows another cell along with the neighboring cells above, below, and to the south, where the vertical pillar of cells contains an inactive cell; this inaccessible rock volume is shown as void space in the figure. No-flow boundary conditions must be imposed on faces marked in white color.

to faults, or have degenerate faces that have collapsed entirely. To ensure that the complex layering of the reservoir is represented with as few cells as possible, each cell may be tilted axially and have large aspect ratios. Because the eight corner-points can be shifted independently up and down the four coordinate lines that delimitate each pillar, the cell faces will generally be bilinear and can deviate far from being planar. Pairs of corner-points can also collapse to a single point so that the cell faces reduce to a triangle or disappear entirely. These effects are illustrated in Figure 2.

Each cell in a stratigraphic grid will only have a single neighbor above and below but may have multiple neighbors in each of the four lateral directions if any of the corresponding hexahedral faces are adjacent to a fault. This can complicate the geometry and topology significantly, as seen in Figure 3. Even if we restrict the number of local polynomials to primary quadruples defined analogously to the primary triples in Figure 1, the number of local polynomials can be very large, which results in a dense local stencil. There is also the risk of putting too much emphasis on polynomials that interpolate across subfaces with small areas. Herein, we simplify the reconstruction by limiting the local stencils so that they only include a single neighboring cell in any of the six logical directions (up, down, north, south, east, west). That is, if a face of the original hexahedral cell

has been subdivided to create a matching grid, we pick the neighboring cell with the largest subface and disregard the others when forming the local polynomials. This is done as part of a preprocessing step. Our approach is obviously a significant simplification and it is not difficult to come up with special cases where this choice is not optimal. However, we are more concerned with robustness for complex grids, which has been verified in a number of challenging test cases.

Cell faces defined by four or more points will generally not be planar, and this introduces a certain ambiguity in how to define geometrical quantities such as face areas, cell volumes, and face/cell centroids. Herein, these quantities are computed by use of a tetrahedral subdivision described in more detail in [34]. The resulting cells are not necessarily convex so that cell centroids may lie *outside* of the cell itself, as illustrated in Figure 4.

3.4 Robust scaling of the smoothness indicator

In our experience, nonconvex cells with high aspect ratios, such as the one shown in Figure 4, do not seem to adversely affect the construction of the local polynomials. However, the smoothness indicator (10) is not invariant to aspect ratios and cannot be used directly, since interpolation using a small gradient can give a large overshoot when multiplied by a large coordinate distance. To make the indicator more robust and invariant under spatial scaling and stretching, we make an affine transformation into the local coordinate space defined by a singular-value decomposition of the vectors from the cell centroid to the face centroids of the cell under consideration, computed as part of the preprocessing.

The transformed cell, shown to the upper-right in Figure 4, is scaled, rotated, and shifted so that the centroids fit into a unit cube, where smoothness is naturally defined. Interpolating in this space only affects the smoothness indicators and not the values of each local polynomial. The local coordinate space is defined by an affine transformation $g(\mathbf{x}) = S(\mathbf{x} - \mathbf{z})$, where S is linear map onto the new coordinate axes and \mathbf{z} is the origin of the new coordinate system. If we let X be the matrix of all vectors from cell centroid to face centroids for a given cell in d dimensions, the singular-value decomposition $UDV^T = X$ has d nonzero singular values in D , with corresponding right-singular vectors in V . These vectors are the major axes of variation in the set of interpolation points, giving us the affine transformation $S = \tilde{D}V^{-1}$, where \tilde{D} corresponds to the top square part of D . After rotation, translation and scaling, the transformed set of face points local to each cell are contained in a cube with approximately equal axes, with the cell centroid as origin. As a simple demonstration of the effect large aspect ratios can have on the general smoothness indicator (10), we consider interpolation inside a parallelogram with sides L




Fig. 4 A single grid cell from the Norne model shown in three different coordinate systems: original coordinates, z -direction scaled by a factor 10, and a coordinate system defined from a singular-value decomposition of the vectors from cell to face centroids.




Fig. 5 Illustration of reconstruction of point values for cells with large aspect ratios. Here, cell averages from four cells are used to reconstruct a point value at the interface between the red and green cells. The left and middle plots show how the quadrilateral interpolation region, which is not aligned with the major axes (left), is transformed to square in the local coordinate system by an affine transformation. The right plot shows the error in the interpolated value when using unscaled variables as a function of height and length. The error increases significantly even for modest deviations from unit aspect ratio.

and H that connects the midpoints of four cells, as in the left plot of Figure 5. The right plot shows the error in the interpolated value at the midpoint of the interface between two of the cells, compared to the (constant) interpolated value computed in SVD-coordinates, for a span of L and H values. This error increases rapidly as the aspect ratio deviates from unity. (We note in passing that the original smoothness indicator for Cartesian grid in [32] is scale-invariant under axis-aligned changes in aspect ratio.)

For simplicity, we define the linear weights γ_i^k in (13) to be proportional to the volume of the triangle/tetrahedron spanned by the centroids that define the corresponding polynomial. This ensures that very small, degenerate triplets or quadruples do not adversely impact the interpolation quality. Another option would be to select linear weights to obtain a third-order reconstruction in regions where the interpolated quantity is smooth, but requires additional treatment for negative linear weights and we thus leave this for future work.

3.5 Linearization and solution of the nonlinear discrete problems

Inserting the reconstructions discussed above into the residual equations of the type shown in (4) for each conserved component gives a nonlinear system of discrete equations, $\mathbf{F}(\mathbf{y}) = \mathbf{0}$, where the vector \mathbf{y} collects all the unknown cell averages in all cells of the model. For compressible rock and fluids, the system will generally be nonlinear also if we use an explicit discretization (i.e., set $m = n$ in (4)). To solve the nonlinear system, we use a standard Newton method iterative solver,

$$\mathbf{0} \approx \mathbf{F}(\mathbf{y}^n) + \mathbf{J}\delta\mathbf{y}, \quad \mathbf{y}^{n+1} = \mathbf{y}^n + \delta\mathbf{y}, \quad (15)$$

where \mathbf{J} is the Jacobian matrix evaluated at \mathbf{y}_0 . The linearized system does not explicitly enforce saturations to be in the unit interval and to mitigate this and to ensure sufficient convergence rates for nonlinear fluxes, we only allow

for a maximum saturation update of 0.2 in each Newton update for the examples with nonlinear flux functions. If convergence issues are detected by oscillating or stagnant maximum residuals, the solver employs a global relaxation for the remainder of the current Newton loop. Another option is to use a line-search, but this may require a large number of costly WENO evaluations. Likewise, we may have to halve the time step to obtain convergence when the prescribed number of nonlinear iterations is insufficient (in the explicit case, the time step is limited by a standard CFL condition).

A main difficulty in developing a fully implicit, high-resolution simulator is to derive and compute the Jacobian matrix. The combination of complex constitutive relationships and the high-order reconstruction can lead to very intricate nonlinearities, and deriving the required linearizations analytically and then coding them is a time-consuming and error-prone process. To avoid this problem, we use automatic differentiation (AD) as explained in [46]. AD relies on the premise that the evaluation of any residual equation can be broken down to a nested sequence of elementary binary operations (addition, multiplication, etc.) and unary operations (logarithm, exponential, etc). These operations satisfy known derivative rules, and we can use these rules in combination with the chain rule to analytically evaluate the derivatives of any function evaluation. In the AD-OO framework of MRST [31], this is done by operator overloading, and as a user, all you have to do is code the residual evaluation, and then the Jacobian matrix is computed simultaneously with the residual by the software.

4 Numerical examples

This section reports a series of numerical examples that compare the performance of the fully implicit WENO method to a first-order, single-point upwind (SPU) method. The study is a continuation of [46], and if not stated explicitly otherwise, all two-phase cases with polymer or waterflooding use the same fluid model. In all simulations reported herein we only reconstruct primary variables; we have run a number of test cases to verify that the discrepancy between reconstructing mobilities and primary variables is negligible. Likewise, we have already presented several comparisons between explicit and implicit schemes in [46], and explained in detail why implicit schemes should generally be preferred. In the following, we therefore only present results from the implicit versions of the WENO and first-order schemes.

4.1 Example 1: Convergence study

We start by studying the numerical error and the order of convergence for the method. In many EOR scenarios, the

chemical fronts will propagate as linear or weakly nonlinear waves. Such waves contain no or very little self-sharpening and are thus more susceptible to numerical smearing than leading nonlinear displacement fronts. To illustrate typical behavior, it is sufficient to study a 1D, single-phase, two-component displacement process. That is, we consider a $L = 100$ m wide reservoir initially filled with a 'blue' fluid except for a 'chemical bank' containing 'red' fluid to the left in the reservoir. We displace the bank of 'red' fluid by injecting 'blue' fluid from the left boundary, assuming a constant pressure drop Δp over the reservoir. The governing equations reduce to a linear advection equation

$$u_t + au_x = 0, \quad u(x, 0) = u_0, \quad a = \mu K \Delta p / (L \phi),$$

which has exact solution $u(x, t) = u_0(x - at)$. Here, however, we solve the problem as a two-phase model with equal fluid properties for the two phases. This setup represents a worst-case scenario since the linear wave lacks the self-sharpening mechanisms that tend to counteract numerical smearing for nonlinear waves.

It is simple to show that the numerical smearing for the SPU scheme is proportional to $\Delta x + a \Delta t$. Introducing the CFL number $\nu = a \Delta t / \Delta x$ that relates the time step to the spatial discretization, we get a smearing $(1 + \nu) \Delta x$, which for a fixed spatial discretization decreases with the time step. (This is in contrast to *explicit* schemes, for which the smearing *increases* with reduced CFL number.) For the WENO scheme, we expect a formal L^1 error to have a $\mathcal{O}(\Delta x^2)$ contribution from the spatial discretization and a $\mathcal{O}(\Delta t)$ contribution from the temporal discretization.

Figure 6 reports the L^1 error on a sequence of refined, uniform Cartesian meshes for two different initial data. For a smooth Gauss-pulse, the SPU scheme exhibits the expected linear convergence. The order of convergence for the WENO scheme is dictated by the choice of the time step, and we can only expect to observe quadratic convergence if we choose $\Delta t \propto \Delta x^2$ or introduce a second-order temporal discretization. It is more interesting to observe the improvement in accuracy for $\nu \sim 1$. For many EOR models, linear waves are typically trailing waves that move slower than leading displacement fronts. Thus, if the CFL number of the leading nonlinear waves are chosen to be moderately above unity, the effective CFL number of the linear wave would be in a range where the second-order WENO discretization would give significantly better resolution than the standard first-order SPU scheme. For the discontinuous double-step profile, the convergence order is one half for both schemes, as expected. For both initial conditions, the nonlinear solver converged in one iteration for all time step for SPU and WENO.

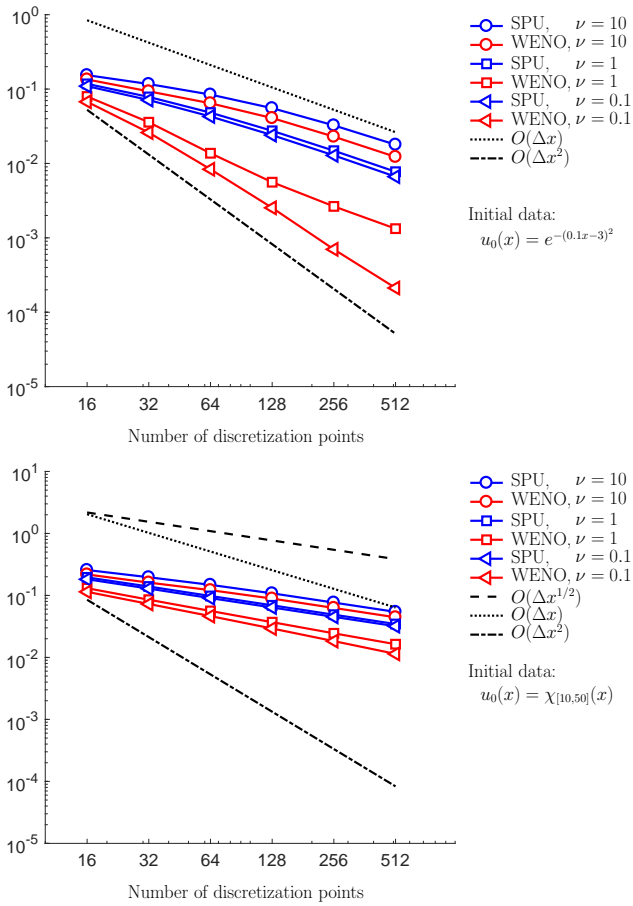


Fig. 6 Convergence study of the L^1 -error of the SPU (blue) and WENO (red) schemes for smooth and discontinuous initial data.

4.2 Example 2: Quarter five-spot

To compare how the SPU and WENO schemes perform for a full displacement profile, we consider the well-known quarter five-spot setup for a displacement of oil by water containing polymer. Instead of using a full well model, we model the injector placed in the south-west corner as a source term with constant injection rate and represent the producer placed in the north-east corner as a boundary condition with fixed pressure. We compare how the two schemes perform on a simple Cartesian grid, a Voronoi grid, and a triangle grid, which all have a comparable number of cells.

Figure 7 reports saturation profiles at a fixed time, chosen so that water has just broken through at the producer in all three grids. The leading water front and the trailing chemical front are both resolved more sharply by the second-order WENO scheme. The computed saturations are more patchy on the two unstructured grids than on the Cartesian grid. The effect is somewhat exaggerated for WENO since we plot cell averages and not the piecewise bilinear reconstruction. A simple grid-refinement study on the Voronoi grid shows that the first-order scheme needs approximately

four times as many grid cells as WENO to achieve the same accuracy. This is consistent with observations made for Cartesian grids in [46].

4.3 Example 3: Layer from SPE 10

To see how the difference in resolution between SPU and WENO translates to a highly heterogeneous setting, we consider a fluvial formation consisting of high-permeable sandstone channels embedded in a low-permeable background of shale and coal, modelled by a 60×220 Cartesian grid with permeability sampled from Layer 45 of the SPE 10 benchmark [10]. We inject water with polymer from a point source placed in grid cell (36, 1) and produce fluids at constant bottom-hole pressure from a producer located in cell (1, 217). From the contours of cell-averaged saturations in Figure 8, sampled 3000 days after injection started, it is evident that WENO scheme captures the displacement fronts with approximately the same accuracy as SPU on a 2×2 refined grid. We observed the same trend for similar simulations on other layers of the SPE 10 model, including the more smoothly varying, shallow-marine, Tarbert formation.

For this example we have used a time-step control built into MRST, which, in its simplest form, takes a set of time step targets (control steps) and a desired number of nonlinear iterations as input. Here, we have used five iterations as our target and a step target that consists of 30 equally spaced time steps of 100 days. To avoid a large initial error, the first step is replaced by a $(\frac{1}{32}, \frac{1}{32}, \frac{1}{16}, \frac{1}{16}, \frac{1}{4}, \frac{1}{2})$ subdivision to form a gradual ramp-up. For some control steps, the time-step controller will reduce the actual time step to achieve convergence within the specified number of nonlinear iterations. The bar graphs in Figure 8 report the number of nonlinear iterations for the two schemes. For SPU, most steps require three iterations and all controls steps complete without chopping. This gives an average time step of 81 days and a minimum step size of 3 days (CFL number 50 and 1.87, respectively). For WENO, the Newton solver struggles more and requires three times as many iterations as SPU in total, giving an average step size of 68 days and a minimum step size of 2 days (CFL number 42.5 and 1.25, respectively).

4.4 Example 4: Norne field model

This example seeks to demonstrate how the WENO scheme can significantly improve the resolution of linear or weakly linear waves at the typical grid resolution seen in field-scale models by considering the conceptual fluid model from Example 2 posed on the grid from the Norne simulation model. We impose a constant injection rate at the end of the reservoir, with no-flow conditions elsewhere, and place two producer wells operating at fixed bottom-hole pressure on the

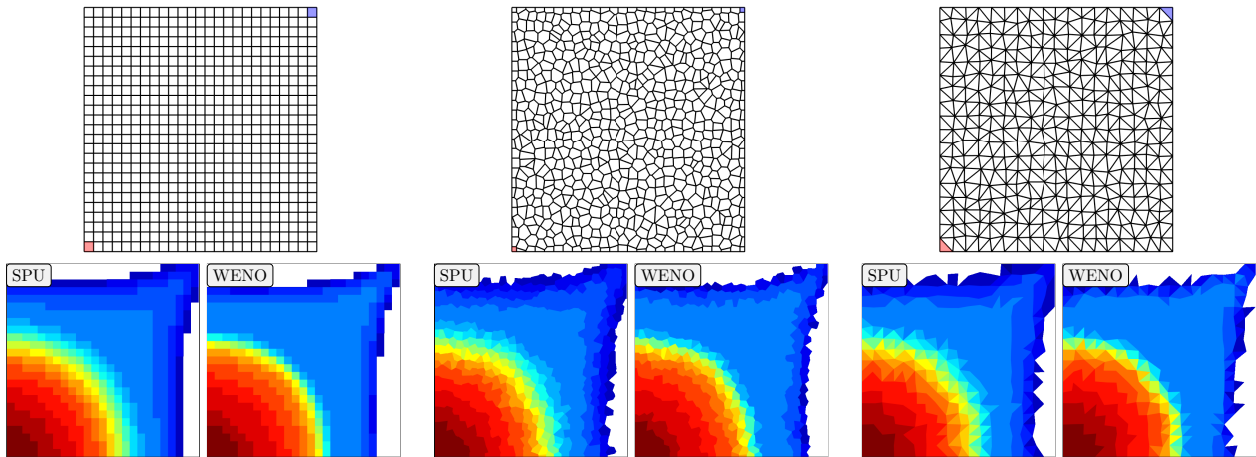


Fig. 7 Quarter five-spot with polymer flooding computed on three different grids with constant (first-order) and WENO (second-order) reconstruction. The color plots show cell-averaged saturation values.

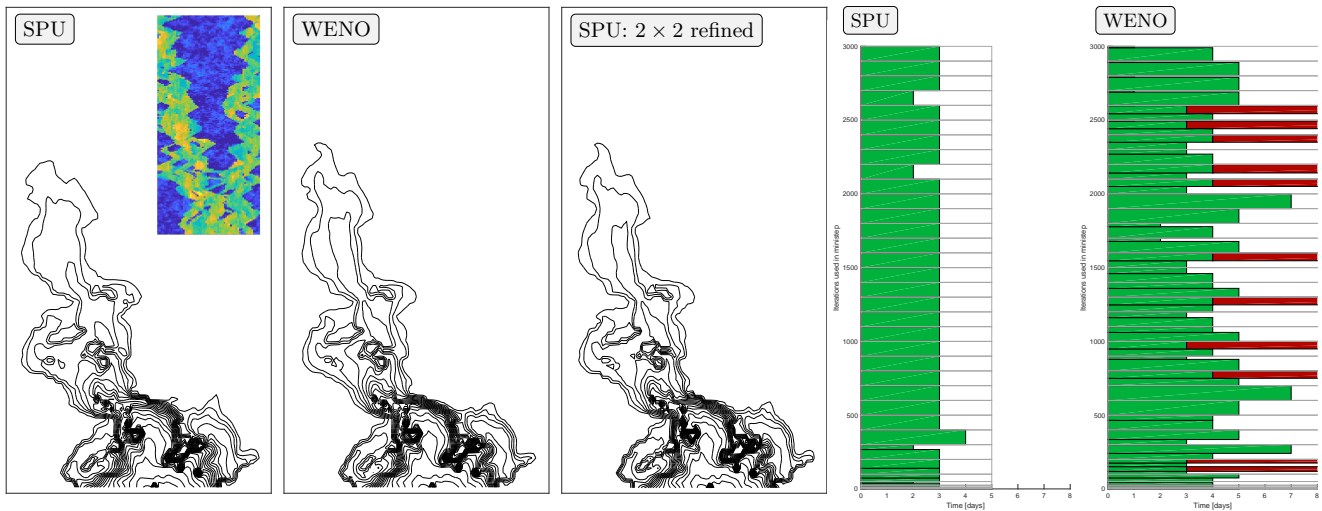


Fig. 8 Polymer injection into a fluvial reservoir (Layer 45 from SPE 10). The contours show water saturation at values 0.22:0.030:0.78. The bar graphs show the number of nonlinear iterations per targeted time step (control step) for SPU (left) and WENO (right) schemes. Some steps have been divided to ensure that the number of iterations per step (green bars) do not exceed five. Steps that have not converged within 10 iterations are considered to fail and illustrated by a red bar.

opposite end of the domain. This simple piston-like displacement of a 'blue' fluid by an identical 'red' fluid, see Figure 9, clearly exhibits how

WENO reduces the numerical smearing significantly compared with the SPU scheme for the same time step. In fact, WENO reduces the smearing more than if we let SPU use ten times as many time steps. Comparing water-cuts from the three simulations, we see how the smearing effect causes SPU to predict almost 80 days earlier breakthrough than WENO. We now repeat the same experiment with permeability and porosity from the original field model. In this case, we observe less discrepancy between the water cut for SPU and the two other simulations. This is to be expected, as the layering of the real model leads to very different flow paths through the model, which reduces the impact of numerical diffusion on the aggregate water production in a

long well, as the first arrival time varies between layers. We use 300 time-steps to simulate 2000 days of displacement with maximum CFL-numbers of 185.5 and 161.8 for the homogeneous and heterogeneous cases, respectively. The cell-average CFL numbers are 1.32 and 0.88. When we examine the nonlinear iterations used by WENO in Figure 10, we note that WENO uses more iterations (3.58 on average) or the heterogeneous permeability than for homogeneous permeability (1.56 on average), largely because of a number of time-step cuts. As this is a linear problem, SPU uses a single iteration per time-step.

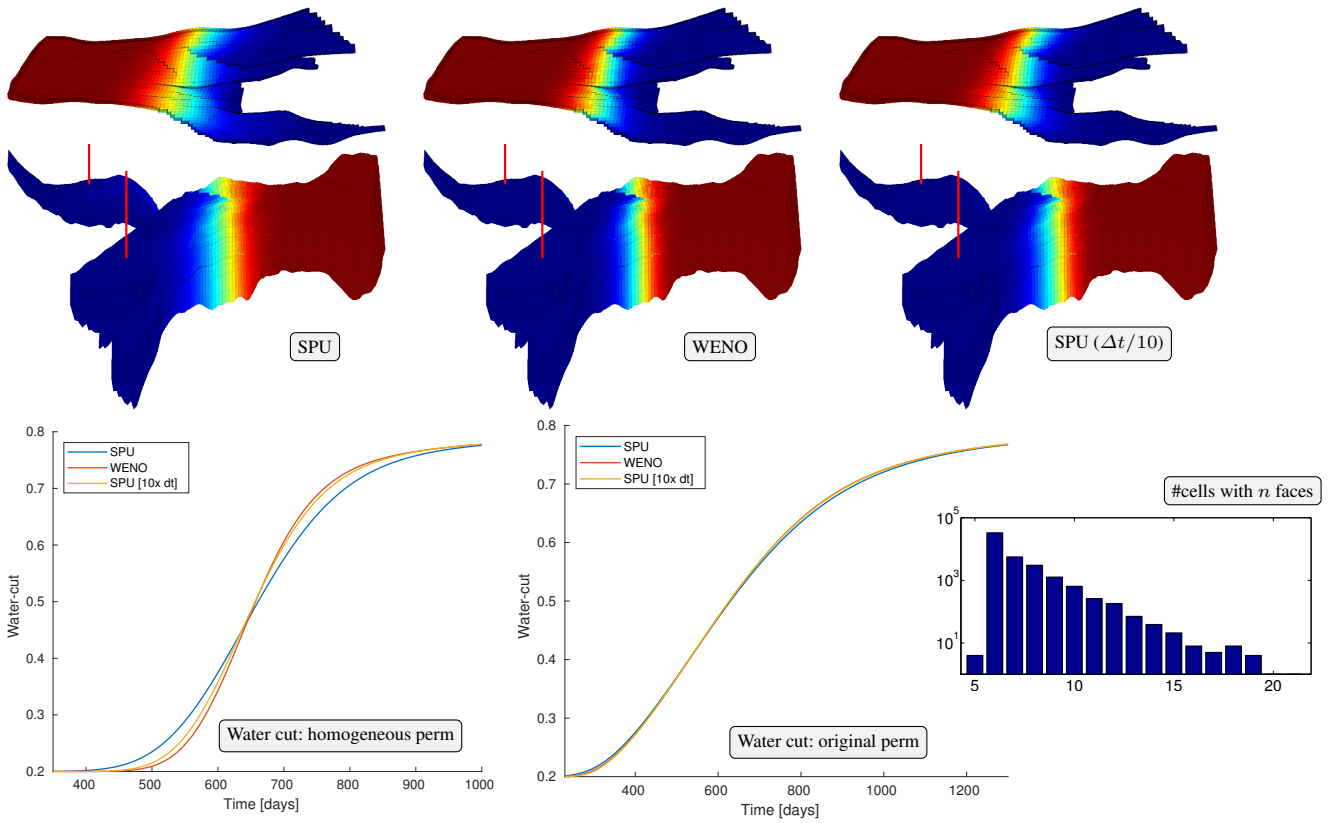


Fig. 9 The two upper rows report water saturations for a piston displacement on the Norne field model with homogeneous permeability simulated by SPU and WENO with the same time step and by SPU with ten times as many time steps; locations of the producer wells are plotted in red. The lower row shows water cut for homogeneous and original petrophysical data, as well as a histogram of the number of faces per cell.




Fig. 10 Number of nonlinear iterations per time step used by WENO for the Norne model with uniform permeability and porosity (left) and original petrophysical data (right).

4.5 Example 5: Unstructured grids

The examples discussed so far have focused on illustrating how the WENO scheme improves the resolution of linear and nonlinear waves. To illustrate the utility of WENO on fully unstructured grids, we consider a case from [2] that describes a vertical cross-section with a producer represented by a single point and an injector whose curved trajectory spans a relatively large part of the domain. The setup consists of four grids:

G1: a coarse Cartesian grid with 231 cells;

G2: a composite grid in which the coarse Cartesian grid is refined locally around the wells by adding Voronoi cells, giving in total 939 cells;

G3: an unstructured Voronoi grid with 1926 cells, which is adaptively refined near the wells; the grid is constructed so that the centroids of the perforated cells are placed exactly on the well trajectory;

G4: a fine Cartesian grid with 20 000 cells.

Figure 11 verifies that WENO gives improved resolution of the leading saturation front as well as the weakly nonlinear chemical front for the two unstructured polyhedral grids.

On grid G3, in particular, WENO seems able to capture the shape of the chemical front with much higher accuracy. Unlike the first-order scheme, WENO captures small undulations in the trailing rarefaction waves caused by the tabulated relative permeabilities on G3 and G4, see [46].

Looking at the oil and water production, we see that the finest grid (G4) has so high resolution that both methods seem to capture the fluid production with almost similar accuracy. On all the other grids, WENO gives much sharper resolution of the water breakthrough and the corresponding decay in oil production. This effect would be even more pronounced if the simulations were continued past the time the chemical displacement front breaks through. Trailing chemical fronts generally have weaker self-sharpening mechanisms than the primary displacement front and are thus more difficult to capture accurately. Use of a higher-order reconstruction can therefore be a feasible alternative to increasing the grid resolution.

4.6 Example 6: Grid effects

The two previous examples demonstrated that the WENO scheme is applicable to grids with the complexity seen in contemporary simulation models. The method nevertheless has two features that may create numerical artifacts. First of all, like in most commercial simulators, we compute inter-cell fluxes $K \nabla p \cdot \mathbf{n}$ using a two-point discretization, which is only consistent on K-orthogonal grids [1, 60]. To illustrate the resulting *single-phase* errors, we use the same setup as in Example 1, except that we now simulate the 1D advection on a 2D quadrilateral mesh in which the nodes have been perturbed so that none of the cells satisfy the condition for K-orthogonality.

The 2D plots in Figure 12 report cell-averaged values for both schemes, whereas the scatter plots report reconstructed point values at the cell centroids for SPU and at face centroids for WENO, plotted as function of the corresponding x -coordinates. In the absence of grid effects, the dots should all fall on a single curve. Here, they do not, and both schemes obviously suffer from errors induced by the inconsistent pressure discretization. On the positive side, WENO resolves the width and the height of the 'chemical bank' significantly more accurately than SPU. This is very important in EOR studies in which enhanced local displacements tend to depend strongly and nonlinearly on the concentration of the active chemical substances. The nonlinear solver uses two iterations in the first time step and one iteration for the remaining time steps for the SPU scheme. For WENO, the nonlinear solver needs two iterations in all time step.

The type of error illustrated in Figure 12 is a *single-phase* phenomenon that stems from an incorrect discretization of the linear operator $K \nabla$ and should not be confused by grid-orientation errors that arise because of insufficient

multidimensional approximation of the (nonlinear) *multi-phase* phase flux, which can give particularly severe grid-orientation errors for adverse mobility displacements also on K-orthogonal grids; see e.g., [17] and references therein. Experiments reported in [46] indicate that second-order TVD and WENO discretizations are somewhat less susceptible to such errors compared with the SPU scheme for a (rotated) quarter-five spot problem posed on Cartesian grids. On the other hand, using a higher-order reconstruction will reduce the stabilizing role of numerical diffusion, and more numerical experiments are needed to assess whether WENO mitigates or enhances this nonlinear error mechanism in general.

Secondly, one could imagine that the stencil reduction discussed in Section 3.3 will introduce artifacts. To investigate this, we consider a nonmatching interface between two rectilinear meshes of different resolution, see Figure 13. At the interface between the two submeshes, the eastern faces of some of the fine cells to the left are split in two. Likewise, the western face of all the coarser cells to the right are subdivided into three or four subfaces with largely different areas. This introduces grid effects both for the SPU scheme and for the WENO scheme with full stencil. On the other hand, reducing the stencil to only involve cell pairs on opposite sides of the subfaces with the largest area does not introduce any notable adverse effect on the scatter plot. For SPU, the nonlinear solver uses two iterations for the first time step and one iteration for the remaining time steps, whereas it requires one or two iterations for all time steps when using WENO with full or compact stencil.

4.7 Example 7: Compositional flow

We consider a gas injection case described using a two-phase, three-component compositional model with Peng–Robinson equation-of-state. The reservoir is a rectangular block with homogeneous properties and no-flow boundaries and produced by a pair of vertical injector/producer wells placed in the southwest/northeast corners. We consider two different variations of the same injection scenario. In both cases, the compositional fluid model is comprised of methane, CO₂, and n-decane. Initially, the reservoir contains a mixture of C1 and C10 in a pure liquid phase. One pore-volume of a gas with a mixture of 10% C1 and 90% CO₂ by moles is injected. The first scenario is at a high pressure of 400 bar at datum depth. The resulting injection is completely miscible with no free gas and gives a displacement profile consisting of a single shock. Gravity forces have limited effect on this single-phase flow scenario. For the second scenario, the initial reservoir pressure is 120 bar at datum depth. Now, the injection gives rise to free gas upon injection and significant gravity segregation and multiphase behavior. High viscosity of the resident liquid results in a weak and unstable displacement profile.

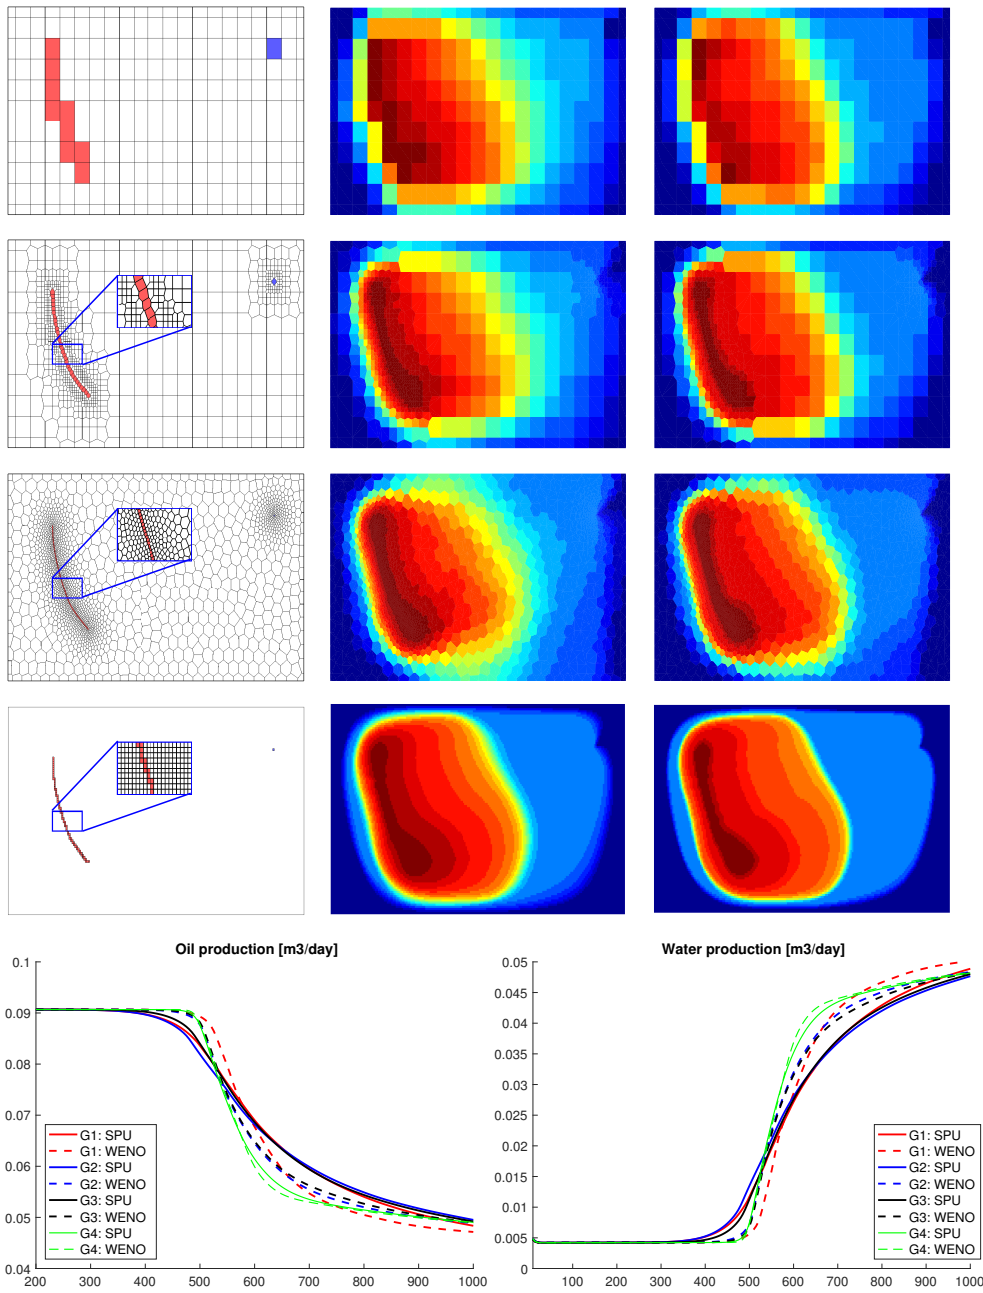


Fig. 11 Example taken from [2] with a point well and a curved well path with the perforated cells marked in blue and red color, respectively. The lower plots show oil and water production rates. In the production plots, solid lines are the first-order method and dashed lines the WENO scheme.

The basic setup is a 2.5D Voronoi grid with 2470 cells, i.e., a grid that is laterally unstructured, but vertically structured. For comparison, we consider three refined grids with twice the vertical resolution and/or approximately 2×2 the lateral resolution. To keep the CFL number approximately the same, we also use twice the number of time steps. To assess the spatial contribution to the smearing for SPU, we also run a simulation with the time step reduced by a factor one hundred. Figure 14 shows 3D snapshots of the CO₂ fraction for some of the simulations.

For the high-pressure scenario, we observe the same behavior as in the other cases above: WENO gives much sharper resolution than SPU on the same grid, and slightly better resolution than SPU on all the refined grids. This is particularly evident when looking at the production rate of CO₂ shown in Figure 15. Because the CO₂ front travels faster the higher you are in the reservoir, the breakthrough will be sharper and occur earlier when the vertical resolution is increased for WENO. Figure 16 shows that the WENO scheme only requires slightly more iterations than SPU on the same grid. At low pressure, we see a much bigger change for both dis-

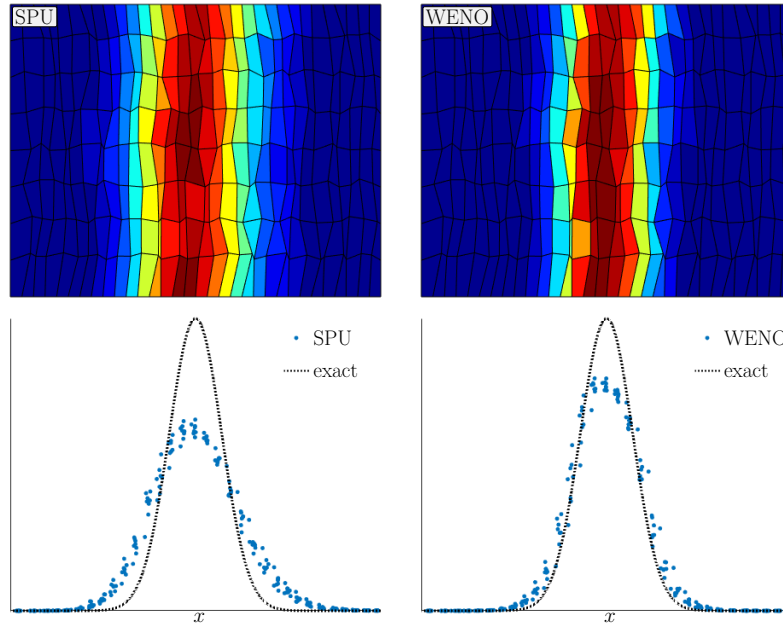


Fig. 12 Simulation of a 1D displacement on a 2D quadrilateral grid that does not satisfy the condition for K-orthogonality.

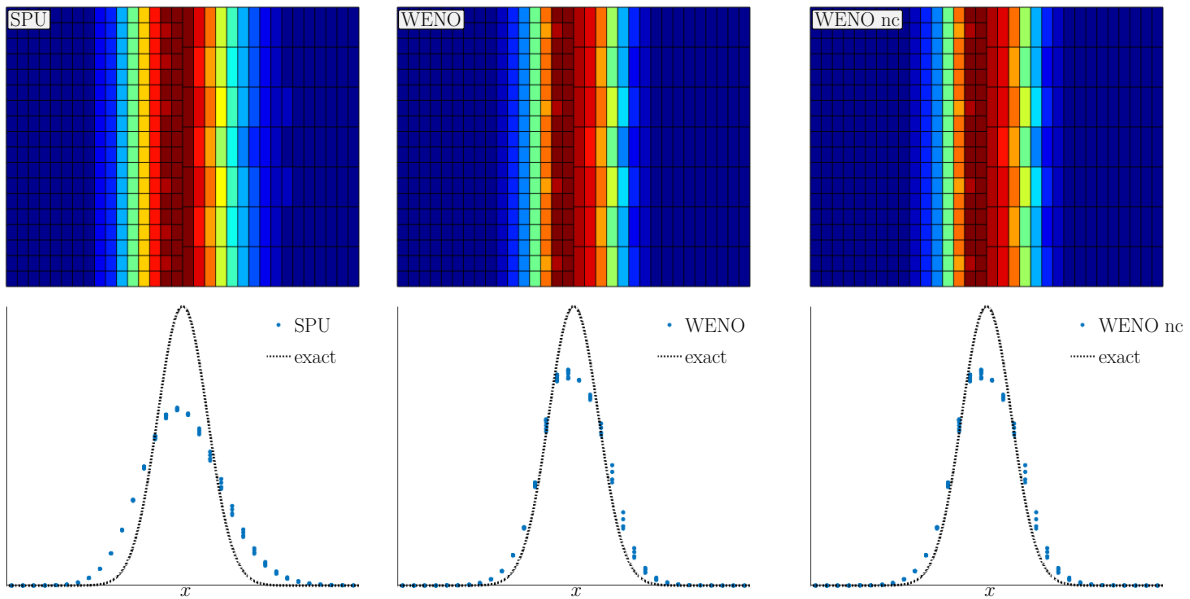


Fig. 13 Simulation of a 1D displacement across a nonmatching grid interface.

cretizations when the vertical resolution is increased. The stair-stepping artifact is likely a result of mobility contrasts between different layers. Using a high-resolution reconstruction does not significantly reduce this artifact. For this, and comparable scenarios, the most efficient method would be to use a high-resolution method on a grid with relatively coarse lateral resolution and high vertical resolution.

5 Concluding remarks

We have presented the formulation of a fully implicit WENO scheme, which is applicable to black-oil type and compositional simulations, and discussed some special adaptations necessary to obtain an efficient and robust scheme on the types of grids found in contemporary reservoir models. In particular, we promote the use of a local coordinate transformation to robustly handle cells with high aspect ratios and nonconvex geometries and the use of automatic differenti-

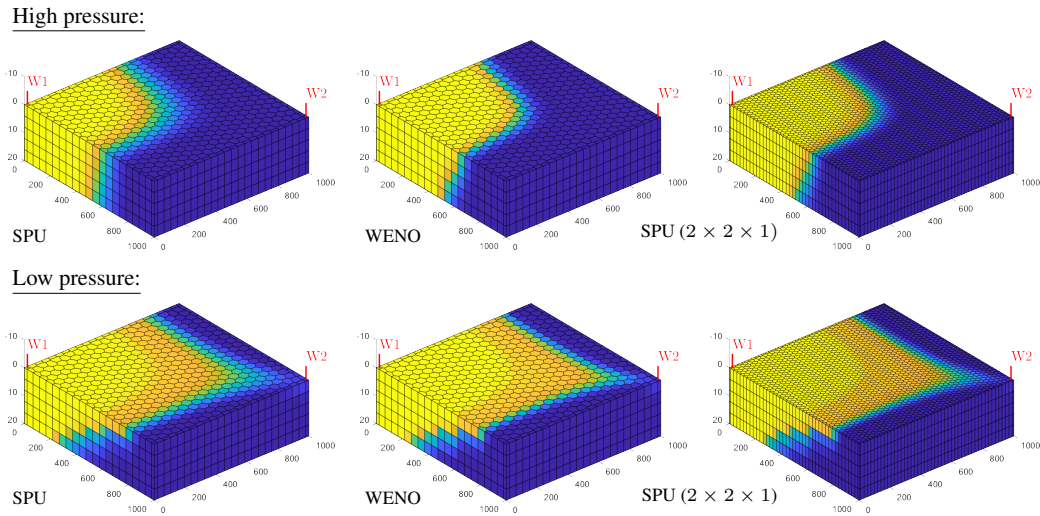


Fig. 14 Snapshots of the CO₂ front part-way during the simulation of the compositional case for two different pressure regimes.




Fig. 15 CO₂ production in unit tonnes per day for the compositional test cases.

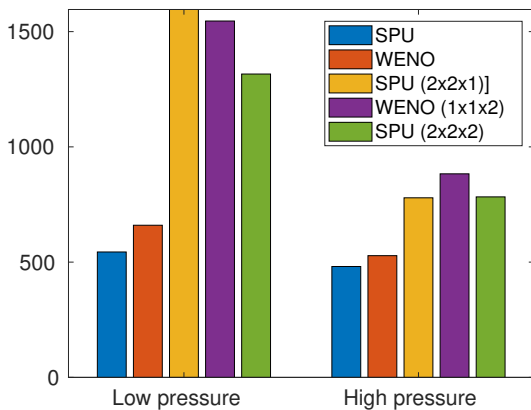


Fig. 16 Number of nonlinear iterations for the compositional gas injection test cases.

ation to overcome the challenge of linearizing the resulting system of discrete nonlinear equations.

A series of numerical tests, some of which are reported herein, show that the WENO scheme improves the resolution of both linear and nonlinear waves significantly, typically giving the same resolution as the standard SPU scheme

with twice as many grid cells in each spatial direction or ten times as many time steps. The results presented herein should nonetheless only be considered as preliminary proof-of-concept. We believe that one can obtain even better results by optimizing the choice of local polynomials and linear weights for each cells. More research is required to this end.

The computational cost of WENO is obviously higher than SPU, not only because of the reconstruction procedure and the denser local stencil, which incurs more evaluations of partial derivatives, but also because the scheme requires more iterations. One can reduce the cost of the reconstruction somewhat by precomputing all geometric parts of the stencil (the local SVD coordinate transformation and the inverse matrix \tilde{C} in (12)). In [46], we also discussed lagged evaluation of the nonlinear weights β_i^k to reduce the nonlinearity of the discrete stencil. Lagging the evaluation over the whole time step seems to work well for imbibition or drainage processes with monotone displacement profiles, but breaks down almost immediately in water-alternating-gas (WAG) type scenarios. Lagging the evaluation in the nonlinear iteration process does not seem to cause similar break-

down, but has little effect on the computational efficiency. We believe a better approach would be to localize reconstruction to regions with significant fluid movement (see e.g., [53]) and try to reuse previous nonlinear weights for stencils where changes in cell averages are below a prescribed value. This requires more research.

Acknowledgements Lie was funded by the Research Council of Norway through grant no. 244361. Mykkeltvedt was funded by the Research Council of Norway and the industry partners of The National IOR Centre of Norway (ConocoPhillips Skandinavia AS, Aker BP ASA, Eni Norge AS, Total E&P Norge AS, Equinor ASA, Neptune Energy Norge AS, Lundin Norway AS, Halliburton AS, Schlumberger Norge AS, Wintershall Norge AS, and DEA Norge AS). Møyner was funded by VISTA, a basic research program funded by Statoil and conducted in close collaboration with The Norwegian Academy of Science and Letters.

References

1. Aavatsmark, I.: Interpretation of a two-point flux stencil for skew parallelogram grids. *Comput. Geosci.* **11**(3), 199–206 (2007). DOI 10.1007/s10596-007-9042-1
2. Bao, K., Lie, K.A., Møyner, O., Liu, M.: Fully implicit simulation of polymer flooding with MRST. *Comput. Geosci.* **21**(5–6), 1219–1244 (2017). DOI 10.1007/s10596-017-9624-5
3. Bell, J.B., Colella, P., Trangenstein, J.A.: Higher order Godunov methods for general systems of hyperbolic conservation laws. *J. Comput. Phys.* **82**(2), 362–397 (1989). DOI 10.1016/0021-9991(89)90054-5
4. Bell, J.B., Shubin, G.R.: Higher-order Godunov methods for reducing numerical dispersion in reservoir simulation. In: *SPE Reservoir Simulation Symposium*. Society of Petroleum Engineers (1985). DOI 10.2118/13514-MS
5. Bell, J.B., Shubin, G.R., Trangenstein, J.A.: A method for reducing numerical dispersion in two-phase black-oil reservoir simulation. *J. Comput. Phys.* **65**(1), 71–106 (1986). DOI 10.1016/0021-9991(86)90005-7
6. Blunt, M., Rubin, B.: Implicit flux limiting schemes for petroleum reservoir simulation. *J. Comput. Phys.* **102**(1), 194–210 (1992). DOI 10.1016/S0021-9991(05)80015-4
7. Chang, Y.B., Pope, G.A., Sepehrnoori, K.: A higher-order finite-difference compositional simulator. *J. Petrol. Sci. Eng.* **5**(1), 35–50 (1990). DOI 10.1016/0920-4105(90)90004-M
8. Chen, W.H., Durlafsky, L.J., Engquist, B., Osher, S.: Minimization of grid orientation effects through use of higher order finite difference methods. *SPE Advanced Technology Series* **1**(02), 43–52 (1993). DOI 10.2118/22887-PA
9. Chertock, A., Kurganov, A., Petrova, G.: Fast explicit operator splitting method for convection–diffusion equations. *Int. J. Num. Meth. Fluids* **59**(3), 309–332 (2009)
10. Christie, M.A., Blunt, M.J.: Tenth SPE comparative solution project: a comparison of upscaling techniques. *SPE Reserv. Eval. Eng.* **4**, 308–317 (2001). DOI 10.2118/72469-PA. URL <http://www.spe.org/csp/>
11. Christov, I., Popov, B.: New non-oscillatory central schemes on unstructured triangulations for hyperbolic systems of conservation laws. *J. Comput. Phys.* **227**(11), 5736–5757 (2008). DOI 10.1016/j.jcp.2008.02.007
12. Edwards, M., Christie, M.A.: Dynamically adaptive Godunov schemes with renormalization in reservoir simulation. In: *SPE Symposium on Reservoir Simulation*. Society of Petroleum Engineers (1993). DOI 10.2118/25268-MS
13. Edwards, M.G.: A higher-order Godunov scheme coupled with dynamic local grid refinement for flow in a porous medium. *Comp. Meth. Appl. Mech. Eng.* **131**(3), 287–308 (1996). DOI 10.1016/0045-7825(95)00935-3
14. Epshteyn, Y., Rivi re, B.: Fully implicit discontinuous finite element methods for two-phase flow. *Appl. Numer. Math.* **57**(4), 383–401 (2007). DOI 10.1016/j.apnum.2006.04.004
15. Friedrich, O.: Weighted essentially non-oscillatory schemes for the interpolation of mean values on unstructured grids. *J. Comput. Phys.* **144**(1), 194–212 (1998). DOI 10.1006/jcph.1998.5988
16. Gottlieb, S., Shu, C.W., Tadmor, E.: Strong stability-preserving high-order time discretization methods. *SIAM Review* **43**(1), 89–112 (2001). DOI 10.1137/S003614450036757X
17. Hamon, F.P., Mallison, B.T.: Fully implicit multidimensional hybrid upwind scheme for coupled flow and transport. arXiv e-prints p. arXiv:1811.04830 (2018)
18. Harten, A.: On a class of high resolution total-variation-stable finite-difference schemes. *SIAM J. Numer. Anal.* **21**(1), 1–23 (1984). DOI 10.1137/0721001
19. Heinemann, Z.E., Brand, C.W., Munka, M., Chen, Y.M.: Modeling reservoir geometry with irregular grids. *SPE Reserv. Eng.* **6**(2), 225–232 (1991). DOI 10.2118/18412-PA
20. Holing, K., Alvestad, J., Trangenstein, J.A.: The use of second-order Godunov-type methods for simulating EOR processes in realistic reservoir models. In: *ECMOR II - 2nd European Conference on the Mathematics of Oil Recovery* (1990). DOI 10.3997/2214-4609.201411104
21. Hornung, R.D., Trangenstein, J.A.: Adaptive mesh refinement and multilevel iteration for flow in porous media. *J. Comput. Phys.* **136**(2), 522–545 (1997). DOI 10.1006/jcph.1997.5779
22. Hoteit, H., Firoozabadi, A.: Multicomponent fluid flow by discontinuous Galerkin and mixed methods in unfractured and fractured media. *Water Resources Research* **41**(11) (2005). DOI 10.1029/2005WR004339
23. Hoteit, H., Firoozabadi, A.: Compositional modeling by the combined discontinuous Galerkin and mixed methods. *SPE J.* **11**(01), 19–34 (2006). DOI 10.2118/90276-PA
24. Jessen, K., Gerritsen, M.G., Mallison, B.T.: High-resolution prediction of enhanced condensate recovery processes. *SPE J.* **13**(2), 257–266 (2008). DOI 10.2118/99619-PA
25. Jiang, G.S., Shu, C.W.: Efficient implementation of weighted ENO schemes. *J. Comput. Phys.* **126**(1), 202–228 (1996). DOI 10.1006/jcph.1996.0130
26. Jiang, J., Younis, R.M.: An efficient fully-implicit multislope MUSCL method for multiphase flow with gravity in discrete fractured media. *Adv. Water Resour.* **104**, 210–222 (2017). DOI 10.1016/j.advwatres.2017.04.009
27. Kaibara, M.K., Gomes, S.M.: A fully adaptive multiresolution scheme for shock computations. In: E.F. Toro (ed.) *Godunov Methods: Theory and Applications*, pp. 497–503. Springer (2001). DOI 10.1007/978-1-4615-0663-8_49
28. Khan, S.A., Trangenstein, J.A., Horning, R.D., Holing, K., Schilling, B.E.R.: Application of adaptive mesh-refinement with a new higher-order method in simulation of a North Sea micellar/polymer flood. In: *SPE Reservoir Simulation Symposium*, 12–15 February, San Antonio, Texas. Society of Petroleum Engineers (1995). DOI 10.2118/29145-MS
29. Klemetsdal,  .S., Berge, R.L., Lie, K.A., Nilsen, H.M., Møyner, O.: Unstructured gridding and consistent discretizations for reservoirs with faults and complex wells. In: *Proceedings of the SPE Reservoir Simulation Symposium* (2017). DOI 10.2118/182666-MS
30. Kl fkor, R., Kvashchuk, A., Nolte, M.: Comparison of linear reconstructions for second-order finite volume schemes on polyhedral grids. *Comput. Geosci.* **21**(5), 909–919 (2017). DOI 10.1007/s10596-017-9658-8

31. Krogstad, S., Lie, K.A., Møyner, O., Nilsen, H.M., Raynaud, X., Skaflestad, B.: MRST-AD – an open-source framework for rapid prototyping and evaluation of reservoir simulation problems. In: SPE Reservoir Simulation Symposium, 23–25 February, Houston, Texas, pp. 1–25 (2015). DOI 10.2118/173317-MS
32. Levy, D., Puppo, G., Russo, G.: Compact central WENO schemes for multidimensional conservation laws. *SIAM J. Sci. Comp.* **22**(2), 656–672 (2000). DOI 10.1137/S1064827599359461
33. Levy, D., Puppo, G., Russo, G.: A fourth-order central WENO scheme for multidimensional hyperbolic systems of conservation laws. *SIAM J. Sci. Comp.* **24**(2), 480–506 (2002). DOI 10.1137/S1064827501385852
34. Lie, K.A.: An Introduction to Reservoir Simulation Using MATLAB: User guide for the Matlab Reservoir Simulation Toolbox (MRST). Cambridge University Press, Cambridge (2019)
35. Liu, J., Delshad, M., Pope, G.A., Sepehrnoori, K.: Application of higher-order flux-limited methods in compositional simulation. *Transp. Porous Media* **16**(1), 1–29 (1994)
36. Liu, J., Pope, G.A., Sepehrnoori, K.: A high-resolution, fully implicit method for enhanced oil recovery simulation. In: Proceedings of SPE Reservoir Simulation Symposium, San Antonio, Texas, USA, 12 - 15 February, pp. 35–50 (1995). DOI 10.2118/29098-MS
37. Liu, Y., Zhang, Y.T.: A robust reconstruction for unstructured WENO schemes. *J. Sci. Comp.* **54**(2-3), 603–621 (2013). DOI 10.1007/s10915-012-9598-3
38. Lohrenz, J., Bray, B.G., Clark, C.R.: Calculating viscosities of reservoir fluids from their compositions. *J. Petrol. Tech.* **16**(10), 1–171 (1964)
39. Mallison, B., Sword, C., Viard, T., Milliken, W., Cheng, A.: Unstructured cut-cell grids for modeling complex reservoirs. *SPE J.* **19**(02), 340–352 (2014). DOI 10.2118/163642-PA
40. Mallison, B.T., Gerritsen, M.G., Jessen, K., Orr, F.M.: High order upwind schemes for two-phase, multicomponent flow. *SPE J.* **10**(03), 297–311 (2005). DOI 10.2118/79691-PA
41. Martin, J.J.: Cubic equations of state - which? *Indust. & Engng. Chem. Fund.* **18**(2), 81–97 (1979). DOI 10.1021/i160070a001
42. Matthäi, S.K., Mezentsev, A.A., Pain, C.C., Eaton, M.D.: A high-order TVD transport method for hybrid meshes on complex geological geometry. *Int. J. Numer. Meth. Fluids* **47**(10-11), 1181–1187 (2005). DOI 10.1002/fld.901
43. Merland, R., Caumon, G., Lévy, B., Collon-Drouaillet, P.: Voronoi grids conforming to 3D structural features. *Comput. Geosci.* **18**(3-4), 373–383 (2014). DOI 10.1007/s10596-014-9408-0
44. Møyner, O., Tchelepi, H.A.: A mass-conservative sequential implicit multiscale method for isothermal equation of state compositional problems. *SPE J.* (2018). DOI 10.2118/915-PA
45. MRST: The MATLAB Reservoir Simulation Toolbox. www.sintef.no/MRST (2017b)
46. Mykkeltvedt, T.S., Raynaud, X., Lie, K.A.: Fully implicit higher-order schemes applied to polymer flooding. *Comput. Geosci.* **21**(5), 1245–1266 (2017). DOI 10.1007/s10596-017-9676-6
47. Natvig, J.R., Lie, K.A.: Fast computation of multiphase flow in porous media by implicit discontinuous galerkin schemes with optimal ordering of elements. *J. Comput. Phys.* **227**(24), 10108–10124 (2008). DOI 10.1016/j.jcp.2008.08.024
48. Oldenburg, C.M., Pruess, K.: Simulation of propagating fronts in geothermal reservoirs with the implicit Leonard total variation diminishing scheme. *Geothermics* **29**(1), 1–25 (2000). DOI 10.1016/S0375-6505(99)00048-6
49. Palagi, C., Aziz, K.: Use of Voronoi grid in reservoir simulation. *SPE Adv. Tech. Ser.* **2**(2), 69–77 (1994). DOI 10.2118/22889-PA
50. Ponting, D.K.: Corner point geometry in reservoir simulation. In: ECMOR I – 1st European Conference on the Mathematics of Oil Recovery (1989). DOI 10.3997/2214-4609.201411305
51. Rivière, B., Wheeler, M.F.: Discontinuous Galerkin methods for flow and transport problems in porous media. *Comm. Numer. Meth. Eng.* **18**(1), 63–68 (2002). DOI 10.1002/cnm.464
52. Sammon, P.H., Kurihara, M., Jialing, L.: Applying high-resolution numerical schemes in reservoirs described by complex corner-point grids. In: SPE Reservoir Simulation Symposium, 11–14 February, Houston, Texas. Society of Petroleum Engineers (2001). DOI 10.2118/66344-MS
53. Sheth, S.M., Younis, R.M.: Localized solvers for general full-resolution implicit reservoir simulation. In: SPE Reservoir Simulation Conference, 20–22 February, Montgomery, Texas (2017)
54. Shi, J., Hu, C., Shu, C.W.: A technique of treating negative weights in WENO schemes. *J. Comput. Phys.* **175**(1), 108–127 (2002). DOI 10.1006/jcph.2001.6892
55. Trangenstein, J.A.: Numerical solution of hyperbolic partial differential equations. Cambridge University Press, Cambridge (2009)
56. van Leer, B.: Towards the ultimate conservative difference scheme. II. monotonicity and conservation combined in a second-order scheme. *J. Comput. Phys.* **14**(4), 361–370 (1974). DOI 10.1016/0021-9991(74)90019-9
57. Verma, S., Aziz, K.: A control volume scheme for flexible grids in reservoir simulation. In: SPE Reservoir Simulation Symposium, 8-11 June, Dallas, Texas. Society of Petroleum Engineers (1997). DOI 10.2118/37999-MS
58. Wieteska, R., Majewski, J., Rokicki, J.: WENO schemes for finite volume simulation of compressible flows. *Math. Model. Anal.* pp. 135–140 (2005)
59. Wolf, W.R., Azevedo, J.L.F.: High-order ENO and WENO schemes for unstructured grids. *Int. J. Num. Meth. Fluids* **55**(10), 917–943 (2007). DOI 10.1002/fld.1469
60. Wu, X.H., Parashkevov, R.: Effect of grid deviation on flow solutions. *SPE J.* **14**(01), 67–77 (2009). DOI 10.2118/92868-PA
61. Xie, Y.w., Edwards, M.G.: Higher resolution total velocity vt and va finite-volume formulations on cell-centred structured and unstructured grids. *Comput. Geosci.* **21**(5-6), 921–936 (2017). DOI 10.1007/s10596-017-9669-5
62. Yee, H., Warming, R., Harten, A.: Implicit total variation diminishing (tvd) schemes for steady-state calculations. *J. Comput. Phys.* **57**(3), 327–360 (1985). DOI 10.1016/0021-9991(85)90183-4