

CPR-Based Multilevel Preconditioning for Efficient Solution of Two-Phase Flow in Highly Heterogeneous Petroleum Reservoirs

José C. A. dos Santos¹, Darlan K. E. de Carvalho², Ramiro B. Willmersdorf², Alessandro R. E. Antunes³, José R.P. Rodrigues⁴, Paulo R. M. Lyra²

¹*Civil Engineering Department, Universidade Federal de Pernambuco
Avenida da Arquitetura, 50.740-550, Recife, Pernambuco, Brasil.*

²*Mechanical Engineering Department, Universidade Federal de Pernambuco
Avenida da Arquitetura Recife, Pernambuco, Brasil.*

³*Technological Center, Agreste Campus, Universidade Federal de Pernambuco
Avenida Marielle Franco, 55.014-900, Caruaru, Pernambuco, Brasil*

⁴*PETROBRAS R&D Center (CENPES)*

Av. Horácio Macedo, 21.941-915, Rio de Janeiro, Rio de Janeiro, Brasil

*josecicero.santos@ufpe.br, Darlan.ecarvalho@ufpe.br, ramiro.willmersdorf@ufpe.br,
alessandro.antunes@ufpe.br, jrprodriques@petrobras.com.br, paulo.lyra@ufpe.br.*

Abstract. Simulating multiphase flow in heterogeneous petroleum reservoirs presents significant computational challenges due to ill-conditioned discrete systems of equations, with a huge number of unknowns, that result from the numerical discretization of these problems. Industry applications generally require the analysis of multiple scenarios, such as during management optimization and uncertainty quantification, requiring efficient numerical methods for the decision-making process. Multiscale methods emerge as a reduced-order model or alternatively as a preconditioner to accelerate the solution at the fine scale. In this work, we present the Non-Uniform Algebraic Dynamic Multilevel (NU-ADM) method, combining the Algebraic Multi-Scale (AMS) and the Two-Point Flux Approximation (TPFA) finite volume formulations for fully implicit two-phase flow simulation. The multiscale operators serve as preconditioners for an iterative solver for the pressure subsystem, which is obtained via the Constrained Pressure Residual (CPR) method. We compare our NU-ADM against standard industrial preconditioners (ILU, AMG) for the CPR pressure system. Our preliminary results show the advantages of the NU-ADM in heterogeneous cases, where traditional methods struggle with slow convergence, and present themselves as a promising alternative in terms of computational cost. The method maintains accuracy while improving computational efficiency through dynamic resolution adaptation. Current investigations aim to explore the use of multi-stage preconditioning (MSP) to further improve performance. This approach combines multiple preconditioning steps, potentially improving robustness for complex cases. Future work will extend our approach to more complex physical processes and optimize preconditioning strategies.

Keywords: Constrained Pressure Residual (CPR), Fully implicit Non-Uniform Algebraic Dynamic Multilevel (FINU-ADM) Method, Reservoir Simulation, Finite Volume, Heterogeneous Porous Media.

1 Introduction

Numerical simulation of fluid flow in petroleum reservoirs is a key tool for acquiring insights and helping decision-making processes supporting oil recovery maximization. While static geocellular models in the industry currently reach up to a billion grid blocks, dynamic simulation models typically operate one hundredth of this amount, being feasible only using supercomputer systems [1]. To reduce computational demands, homogenization techniques such as upscaling are often employed to approximate solutions on a coarser grid, which causes fine-scale information loss and reduces accuracy. To address this, multiscale finite volume (MsFV) methods have

emerged, offering more accurate results than traditional upscaling by transferring information between fine and coarse scales using transfer operators (restriction and prolongation) at a lower computational cost than full fine-scale simulations [2].

In this work we apply Algebraic Multiscale/Multilevel procedures as preconditioners to accelerate the computation of high-fidelity solutions directly on the fine mesh with Fully Implicit Method (FIM). The computational structure includes introducing a level adaptation mechanism based on normalized pressure gradients. The system of equations associated with the pressure field is derived using the Constrained Pressure Residual (CPR) method, which extracts the elliptic block from the Jacobian matrix of the system of equations. The full system of equations is solved using the Generalized Minimum Residual Method (GMRES) with ILU-0 (Incomplete LU factorization) preconditioning. The developed methodology is compared with the use of Algebraic Multigrid (AMG) as preconditioner.

2 Mathematical model

In this work, we consider the flow of immiscible, isothermal, and incompressible fluids, oil and water, through an incompressible porous medium, neglecting the effects of capillarity and adsorption. A fully implicit (FIM) approach is employed to solve the mass conservation equation. Under these assumptions, the mass conservation equation to a phase f ($f = o$, for oil, $f = w$ water) can be expressed as follows [3]:

$$\phi \rho_f \frac{\partial S_f}{\partial t} + \vec{\nabla} \cdot (\rho_f \vec{v}_f) - q_f = 0, \quad (1)$$

where \vec{v}_f and ρ_f denote the Darcy velocity and the density of phase f , respectively; ϕ is the effective porosity of the rock i.e., the fraction of interconnected void space available for fluid flow; and q_f represents the specific source/sink term, such as injection or production wells. S_f is the saturation of phase f , that is, the fraction of the pore volume occupied by that phase. Assuming a fully saturated porous medium, the saturation constraint $\sum S_f = 1$ is imposed, where S_f , with $f = o, w$, corresponds to the saturations of oil and water, respectively.

Neglecting capillary effects, the Darcy velocity can be expressed as:

$$\vec{v}_f = -\lambda_f K (\vec{\nabla} p - \rho_f g \vec{\nabla} z), \quad (2)$$

where p , $\lambda_f = k_{rf}/\mu_f$, k_{rf} e μ_f represent, respectively, the pressure, the phase mobility, the relative permeability and the viscosity of the phase f , K is the absolute permeability tensor of the rock, $\vec{\nabla}$ is the gradient operator and g gravity acceleration module, and z is the an unitary vector on the direction of the gravity force. Moreover, the simplifying assumptions inherent to Darcy's law are considered valid in the studies conducted [3].

3 Numerical approximation

Substituting Darcy's law (2) into the mass conservation Eq. (1) and integrating over the domain Ω , we obtain:

$$\int_{\Omega} \{ \phi \rho_f \frac{\partial S_f}{\partial t} - \vec{\nabla} \cdot [\rho_f \lambda_f K (\vec{\nabla} p - \rho_f g \vec{\nabla} z)] - q_f \} d\Omega = 0, \quad (3)$$

Considering a computational mesh with N control volumes, that discretizes the computational domain Ω , the discrete form of Eq. (3) can be written as:

$$\sum_1^N \int_{\Omega_i} \{ \phi \rho_f \frac{\partial S_{fi}}{\partial t} - \vec{\nabla} \cdot [\rho_f \lambda_{fi} K (\vec{\nabla} p_i - \rho_f g \vec{\nabla} z)] - q_{fi} \} d\Omega_i = 0. \quad (4)$$

3.1 Non-Uniform Algebraic Dynamic Multilevel (NU-ADM)

The Algebraic Dynamic Multilevel (ADM) method combines an algebraic multilevel approach (inspired by multigrid methods) with Algebraic Multiscale Solver (AMS) to efficiently represent solutions across different scales. This method builds a hierarchy of approximation spaces, where AMS basis functions capture local features of the problem, such as heterogeneities, without explicitly depending on the physical mesh. This enables a compact and accurate representation of the solutions, reducing computational cost. The adaptive dynamics adjusts the basis functions as needed, maintaining critical regions on the fine scale and mapping less relevant areas to the coarse scale. In its classical implementation [5], this adaptation considers the saturation front and well locations as critical regions to retain the fine-scale resolution. The NU-ADM method operates similarly to ADM. However, it defines non-uniform levels for scale transfers, offering greater flexibility in defining coarse-scale spaces and, consequently, a lower percentage of volumes retained at the fine scale compared to the ADM method. Additionally,

the NU-ADM strategy adapts the mesh to pressure and saturation fields, well locations, and Multiscale basis functions, thus providing tighter control over error norms [6].

Apart from the ADM and NU-ADM alternatives, we consider, for comparison purpose, the ADM* method that uses the same level definition parameters as NU-ADM but with uniform levels, like ADM, resulting in the highest number of volumes retained at the fine scale among the three methods and lower error norms, but at a higher computational cost.

3.2 Constrained Pressure Residual (CPR)

In general, the linear system solution step accounts for a large portion of the total computational time. It is well known that the performance of iterative linear solvers depends heavily on the choice of robust and efficient preconditioners that are suitable for massive parallelization. Additionally, for realistically scaled problems, memory requirements also become a significant concern.

The Constrained Pressure Residual (CPR) method [4] was developed based on treating the parabolic/elliptic part of the system as a separate internal step (predictor stage), combined with a residual correction stage. The combination of this predictor–corrector strategy resulted in improved convergence properties compared to the strategy adopted in the FIM. In the case of the FIM, the CPR pressure equation is extracted from the fully implicit approach.

Initially, a decoupling matrix \mathbf{M} is applied to the system such that $J_{ps}^* \rightarrow 0$.

$$\underbrace{\mathbf{M}\mathbf{J}}_{\mathbf{J}^*} \Delta x = \underbrace{\mathbf{M}\mathbf{r}}_{\mathbf{r}^*} \rightarrow \begin{bmatrix} \mathbf{J}_{pp}^* & \mathbf{J}_{ps}^* \\ \mathbf{J}_{sp} & \mathbf{J}_{ss} \end{bmatrix} \begin{bmatrix} \Delta x_p \\ \Delta x_s \end{bmatrix} = \mathbf{r} = \begin{bmatrix} r_p^* \\ r_s \end{bmatrix} \quad (5)$$

The choices for \mathbf{M} can be made such that the off-diagonal term J_{ps}^* becomes small or, in the case of the Schur complement, becomes zero. The elimination of J_{ps}^* can be expressed algebraically as the multiplication of a matrix $\mathbf{C} = [\mathbf{I} \ \mathbf{0}]^T$ to the system. The pressure equation system is then formed by solving for an approximate pressure increment $\Delta x'_p$:

$$\underbrace{(\mathbf{C}^T \mathbf{M} \mathbf{J} \mathbf{C})}_{\mathbf{A}_{CPR}} \Delta x'_p = \mathbf{C}^T \mathbf{r}^* \quad (6)$$

In this way, the coefficient matrix of the CPR pressure system is given by $\mathbf{A}_{CPR} = \mathbf{C}^T \mathbf{M} \mathbf{J} \mathbf{C}$. In this work, we used the so called True-IMPES reduction [5]:

$$\mathbf{M} = \begin{bmatrix} \mathbf{I} & -\text{colsum}(\mathbf{J}_{ps})\text{colsum}(\mathbf{J}_{ss})^{-1} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \quad (7)$$

In general, the linear system (8) is solved using the GMRES method with two-stage preconditioning. In the first stage, an AMG method with V-cycle is used to obtain an approximate solution of the CPR pressure subproblem, Eq. (9). In the second stage, a classical preconditioner, such as ILU(0), is applied to the global system.

3.3 Multilevel CPR Algorithm

The multilevel CPR-based algorithm also employs a two-stage preconditioning strategy. However, in this work, a multilevel method is used to precondition the pressure system instead of the commonly adopted AMG. By splitting the fully implicit Jacobian matrix \mathbf{J} , the iterative solution procedure can be summarized as follows:

(1) Initialization:

- Using the current FIM solution x^{*v} , compute the FIM residual r^{*v} .

(2) Stage 1: Pressure Stage:

- Define the CPR system (9);
- Apply the pressure solver (FINU-ADM, ADM, or ADM*);
- Update the pressure solution $x_p^{v+1} = x_p^v + \Delta x_p^{*v}$ and the corresponding residual r_p^{v+1} .

(3) Stage 2: Global Stage:

- Recompute the FIM residual based on the updated pressure correction $r^{*v+1} = r^v - \mathbf{J}(\mathbf{C}\Delta x_p^{v+1})$;
- Apply the preconditioner to the global system to obtain the updated correction $\Delta x^{*v+1} = \mathbf{L}^{-1}r^{*v+1}$.

(4) Solution Update:

- Update the full solution using the global and pressure stages $x^{*v+1} = x^{*v} + \Delta x^{*v+1} + \mathbf{C}\Delta x_p^{v+1}$.

4 Results

In this section, we present the results, $x^{FIM-NU-ADM}$, obtained by applying multilevel preconditioners on the NU-ADM framework to a set of test cases and compare to those obtained solution on the fine-scale mesh, x^f ($x = p$, for the pressure and $x = s$, for the saturation). The error norm is defined as,

$$\|e_x\| = \left\| \frac{x^{FIM-NU-ADM} - x^{fs}}{x^{fs}} \right\|. \quad (8)$$

In the example we have used a fixed coarsening ratio $CR = (5, 5)$. The permeability field and the wells positions are shown in Figure 1. As initial condition we set a constant water saturation of $Sw = 0.0$ equal to the irreducible, the residual oil saturation was set as 0.0. The viscosity is 0.3 cP for the water and 3cP for the oil, in addition null flux was set in all external boundaries. The permeability field is shown In Figure 2, we used the following NU-ADM level control parameters, $\delta_s^{lim} = 0.1$, $\delta_p^{lim} = 0.1$ and $\beta^{lim} = 10$. The definition of these parameters is available on [6].

For the preliminary analysis of how the number of fine-scale control volumes affect computational time, the reservoir depicted in Figure 1 was discretized using three different grid resolutions: $60 \times 60 = 3.600$; $120 \times 120 = 14.400$; and $240 \times 240 = 57.600$ control volumes.




Figure 1. Logarithmic-scale permeability field of the 2D reservoir with semi-circular barriers and the location of injection and production wells.

The following nomenclature has been adopted to describe how the approximate solutions were obtained using the different multiscale-multilevel method variants analyzed:

FI-ADM (Fully Implicit Algebraic Dynamic Multilevel): The ADM method employs a multilevel algebraic strategy inspired by multigrid techniques, combined with Algebraic Multiscale Basis Functions (AMS) to represent solutions across different scales. It builds a hierarchy of approximation spaces by solving local flow problems and capturing local heterogeneities. The dynamic adaptation retains critical regions, such as saturation fronts and control volumes containing wells, at the fine scale, while less relevant areas are mapped to the coarse scale. This approach reduces computational cost while maintaining accuracy.

FINU-ADM (Fully Implicit Non-Uniform ADM): The FINU-ADM follows similar principles as ADM but introduces non-uniform levels for scale transfer, enhancing flexibility in defining coarse-scale approximation spaces. As a result, fewer volumes are retained at the fine scale compared to ADM [6]. Additionally, FINU-ADM adapts the mesh not only based on saturation and well locations, but also incorporates pressure fields and multiscale basis functions, improving control over error norms. This strategy provides greater computational efficiency without sacrificing accuracy and is particularly advantageous in problems with high spatial variability.

FI-ADM* (Extended FI-ADM): FI-ADM* retains the level-definition parameters of FINU-ADM but applies uniform levels, like the traditional FI-ADM. Consequently, more volumes are kept at the fine scale compared to other methods, resulting in lower error norms but increased computational time.

For all the multiscale-multilevel methods described above, when preceded by the CPR prefix, it indicates

that the Constrained Pressure Residual (CPR) strategy is used to extract the pressure submatrix and apply the corresponding multiscale method as a preconditioner, accelerating the convergence of fine-scale solutions. Finally, CPR-AMG refers to the industry-standard approach, which uses Algebraic Multigrid (AMG) as the preconditioner.

The simulation results for the different benchmark cases presented below were obtained using the methods listed in Table 1. The results are grouped into two categories: approximate solutions and fine-scale solutions. The approximate solutions are derived directly from the multiscale formulations, while the fine-scale solutions are obtained using either multiscale-multilevel or classical preconditioners. In all cases, the CPR-AMG strategy is used as the reference benchmark, representing the industry-standard approach for comparison purposes.

Figure 2 presents the average computational time required to solve a single time step, considering the simulation up to 60% of the Porous Volume Injected (PVI). In addition to the computational times obtained from the evaluated numerical models, the figure includes the best-fit power-law regression curve, $t(n^f) = a.(n^f)^b$, which was derived using the least squares method.

For this example, it is observed that methods using multiscale formulation as an approximation technique exhibit a sublinear relationship (with exponent values below or close to 1) between computational time and the size of the fine-scale model. Methods that employ multiscale formulation as a preconditioner produce exponents of approximately 1.16 for CPR-FINU-ADM, 1.27 for CPR-FI-ADM, and 1.24 for FI-ADM*, compared to the reference method, CPR-AMG, which has an exponent of about 1.13.

Table 1: Level control parameters adopted.

Método	β^{lim}	δ_p^{lim}	δ_s^{lim}	Level criteria
FI-ADM	∞	∞	0.1	Uniform
FI-ADM*	10	0.1	0.1	Uniform
FINU-ADM	10	0.1	0.1	Non-uniform
CPR-FI-ADM	∞	∞	0.1	Uniform
CPR-FI-ADM*	10	0.1	0.1	Uniform
CPR-FINU-ADM	10	0.1	0.1	Non-uniform
CPR-AMG	-	-	-	Finescale

These results indicate that all tested variants show good scalability for the problem under consideration. Furthermore, the performance of multiscale methodologies using the standard ADM or the modified ADM* is slightly lower than that of FINU-ADM, which exhibits performance nearly equivalent to that of AMG.




Figure 2: Average computational time per time step, considering the simulation up to the time corresponding to 60% PVI, for the reservoir simulation case with semi-circular barriers.

Figure 3 presents the error norms for pressure (a) L_2 and (b) L_∞ and for saturation (c) L_1 along with the percentage of active volumes (d), simulation time (e), and water cut (f), obtained from the simulation of the model with a resolution of $60 \times 60 = 3.600$ control volumes.

It is observed that, in the context of approximate solutions, the classical multilevel method (FI-ADM), which retains only the saturation front at the fine scale, is not suitable for maintaining pressure errors under control. This approach leads to L_2 norm errors of approximately 9% and L_∞ norm errors of up to 300% during most of the

simulation.

In contrast, methods that incorporate adaptation parameters based on the pressure field significantly reduce pressure error norms, achieving L_2 norm errors around 1% and L_∞ norm errors near 3%.

The saturation field errors were generally higher for the methodology using non-uniform levels. This behavior is attributed to the lower percentage of volumes retained at the fine scale by this approach. To improve this metric, one may reduce the adaptation threshold parameter for the saturation field or adopt non-constant basis functions for saturation.

Regarding simulation time, Figure 3 (e) shows that applying multiscale methodologies in conjunction with CPR leads to a reduction in simulation time, for the problem sizes analyzed, when compared to the combination of CPR with AMG. This improvement is attributed to the substantial reduction in the number of degrees of freedom in the pressure system, as shown in Figure 3 (d).

Finally, in terms of water cut, the methods using the multiscale formulation as a preconditioner accurately reproduce the reference solution obtained with CPR-AMG, as expected. Meanwhile, the approximate multiscale methods reproduce the reference solution only approximately.

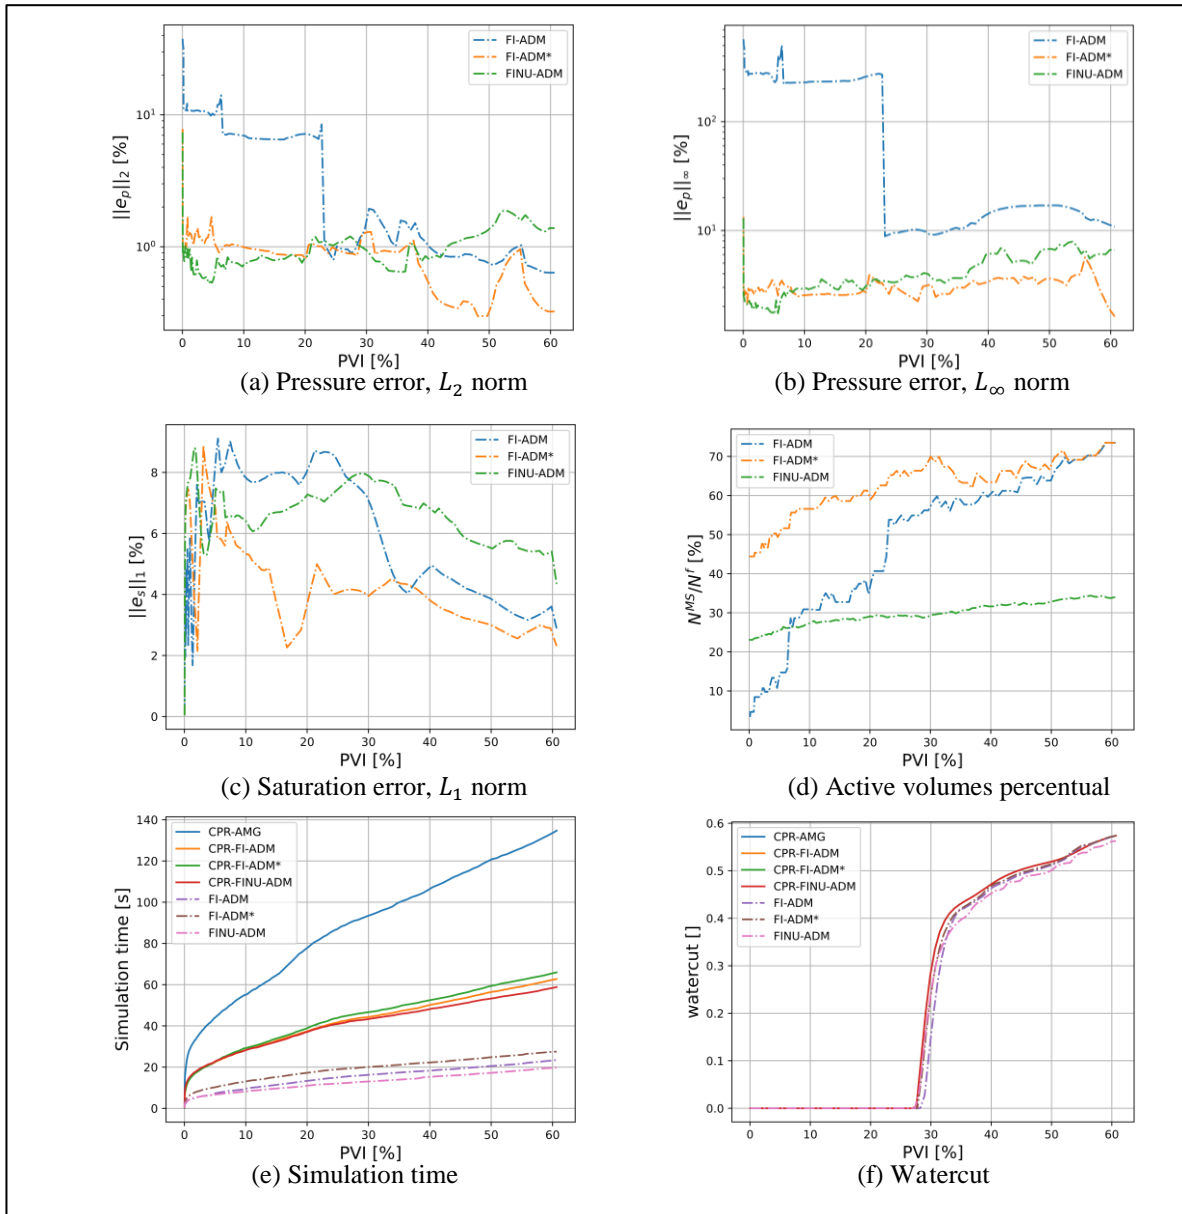


Figure 3: Average computational time per time step, considering the simulation up to the time corresponding to 60% of the injected pore volume (VPI), for the reservoir.

5 Conclusions

For the two-phase flow example tested, the FINU-ADM methodology provided solutions that were at least six times faster when used as an approximate solver, while maintaining pressure field errors below 3% (in the L_2 norm) and saturation errors below 7% (in the L_1 norm). When used as a preconditioner, FINU-ADM also achieved speedups ranging from 20% to 100% compared to the reference CPR-AMG solution.

Results demonstrated that CPR-FINU-ADM offers better scalability than traditional approaches such as CPR-AMG and FI-ADM, with almost linear growth in computational time relative to fine-scale model size. These findings highlight the method's potential for large-scale simulations, with future work focused on extending its robustness through multi-stage preconditioning and application to more complex flow scenarios.

Acknowledgements. We would like to thank PETROBRAS (Petróleo Brasileiro S.A.) and CNPq (Conselho Nacional de Desenvolvimento Científico e Tecnológico) for the financial support for this research.

Authorship statement. The authors hereby confirm that they are the sole liable persons responsible for the authorship of this work, and that all material that has been herein included as part of the present paper is either the property (and authorship) of the authors, or has the permission of the owners to be included here.

References

- [1] A. C. R. Souza, L. M. C. Barbosa, F. R. L. Contreras, P. R. M. Lyra and D. K. E. Carvalho, "A Multiscale Control Volume framework using the Multiscale Restriction Smooth Basis and a non-orthodox Multi-Point Flux Approximation for the simulation of two-phase flows on truly unstructured grids". *Journal of Petroleum Science and Engineering*, vol. 1, n. 188, pp. 1–21, 2020, doi: 10.1016/j.apm.2021.11.017.
- [2] S.B. Bosma, S. Klevtsov, O. Møyner and N. Castelletto, "Enhanced multiscale restriction-smoothed basis (MsRSB) preconditioning with applications to porous media flow and geomechanics,". *Journal of Computational Physics*, Vol. 1, n. 109934, 2020, doi:10.1016/j.jcp.2020.109934.
- [3] Ertekin, T., Abou-Kassem, J. H., & King, G. R. (2001). *Basic Applied Reservoir Simulation*. Richardson, TX: Society of Petroleum Engineers (SPE).
- [4] Wallis, J. R., Kendall, R. P., Little, T. E., & Nolen, J. S. (1985) "Constrained residual acceleration of conjugate residual methods. SPE Reservoir Simulation Symposium" *SPE Reservoir Simulation Symposium*. DOI: <http://dx.doi.org/10.2118/13536-MS>.
- [5] M. Cusini, C. Van Kruijdsdijk e H. Hajibeygi, "Algebraic Dynamic Multilevel (ADM) method for fully implicit simulations of multiphase flow in porous media." *Journal of Computational Physics*, pp. 60-79, 2016. DOI: <https://doi.org/10.1016/j.jcp.2016.03.007>
- [6] Santos, J. C., Lyra, P. R., Andrade, J. P., Souza, A. C., Lira, R. J., & Carvalho, D. K. (2022). An algebraic dynamic multilevel and multiscale method with non-uniform mesh resolution and adaptive algebraic multiscale solver operator for the simulation of two-phase flows in highly heterogeneous petroleum reservoirs." *Journal of Computational Physics*, 1-24. DOI: <http://dx.doi.org/10.1016/j.jcp.2022.111174>