EPSRC
Engineering and Physical Sciences Research Council

InFoMM
Industrially Focused Mathematical Modelling

EPSRC Centre for Doctoral Training in Industrially Focused Mathematical Modelling

Improved Algebraic Decoupling of the Pressure Equation during Reservoir Simulation

Thomas Roy

UNIVERSITY OF OXFORD

Schlumberger

Contents

1. Introduction

Background

The use of reservoir simulation models in the oil and gas industry is essential to the development of new oil fields and in generating production forecasts for existing oil fields. Reservoir simulators solve mathematical models for the flow of fluids (typically oil, water and gas) through porous media. The complexity of such models, and the time each simulation takes motivates the need for highly efficient solvers. A major time-consuming step in reservoir simulation is solving thousands of linear systems of equations. Our aim is to find a novel transformation of these equations which reduces the number of iterations needed to solve them.

Mathematical Model

Mathematical models for oil reservoirs describe flow through porous media. For single-phase flow, the model comprises of a conservation of mass equation, where the velocity of the fluid is usually described using Darcy's law, relating the velocity to the pressure gradient. This can easily be generalized for multi-phase flow, in which mass transfer between phases is allowed. Black-oil systems are a special case of these models in which we assume that there is a water phase, an oil component which can form its own phase or get dissolved in the water, and a gas component which can form its own phase or get dissolved in the oil. This model may be written as a system of coupled partial differential equations (PDEs). In order to solve to the systems numerically, we discretise and apply an iterative nonlinear solver (e.g. the Newton-Raphson method) to the resulting nonlinear system. We obtain a very sparse linear system of the form

$$Ax = \begin{bmatrix} A_{pp} & A_{ps} \\ A_{sp} & A_{ss} \end{bmatrix} \begin{bmatrix} x_p \\ x_s \end{bmatrix} = \begin{bmatrix} b_p \\ b_s \end{bmatrix} = b, \tag{1}$$

The ellipticity of the pressure equation means that the pressure variable has an influence over the whole reservoir.

where A_{pp} is a block matrix representing the "pressure coefficients", A_{ss} is a block matrix representing the coefficients of the "secondary variables" (typically concentrations/saturations), and A_{ps} and A_{sp} represent the respective coupling coefficients. The solution of the system is the vector x, where x_p is the pressure variable and x_s represents the secondary variables. The vector b is the residual of the system, where b_p is the residual of the pressure equation and b_s represents the residuals of the secondary equations. The pressure equation is represented by the first row of this system, and the second row represents the equations for the secondary variables of the model. The pressure equation is different from the equations for the secondary variables, in that it is elliptic, and the other equations are not. The ellipticity of the pressure equation means that the effect of the pressure is global, i.e. a change in pressure in one part of the reservoir influences the flow everywhere in the reservoir. In contrast, the secondary variables are local in nature.

Decoupling the Pressure Equation

A preconditioner makes a linear system easier to solve by iterative methods.

In a decoupled system, variables do not interact with each other, and can be solved for independently.

A crucial part of the effective solution of the linear systems is the choice of preconditioners. These are transformations that one applies to a linear system in order to make it easier to solve by iterative methods. The different nature of the pressure equation and secondary equations means that a separate preconditioner is needed for each set. Indeed, the elliptic-like nature of the pressure variable requires a preconditioner that conserves the global influence of the pressure. For the other equations, a local preconditioning is sufficient.

In order to treat the pressure separately from the secondary variables, we must first reduce the coupling between them. A system is fully decoupled if the variables do not interact with each other and each variable can be solved for independently. A full decoupling ($A_{ps} = 0$) is not numerically feasible and thus we only seek to weaken the coupling, so that a large change in the secondary variables only makes a very small change in the pressure. In fact, we seek to concentrate the elliptic properties of the system inside the pressure block A_{pp}. By doing this, we can then precondition the pressure equation (represented by $A_{pp}x_p = b_p$) with a global preconditioner, and the rest of the system with a local preconditioner.

Glossary of Terms

- Decoupling operator: A matrix multiplied on both sides of the linear system in order to reduce the coupling between the equations.
- Diagonal dominance: A matrix is diagonally dominant if the coefficients on its diagonal are greater than the sums of the absolute values of the off-diagonal coefficients of their respective rows.
- M-matrix: A diagonally dominant matrix with positive coefficients on its diagonal, and non-positive off-diagonal coefficients. The M-matrix property is very important for the convergence of the algebraic multigrid (AMG) method.
- Sparse matrix: A matrix is sparse if it has relatively few non-zero coefficients. The sparsity pattern of a matrix relates to where these non-zero coefficients are located in the matrix.
- Norm: A norm is a function that assigns a positive number to a vector or a matrix. It serves as an indicator of the magnitude of vectors or matrices. The Euclidean norm of a vector x gives the ordinary distance from the origin to the point x.

2. Constrained Pressure Residual Preconditioner

The constrained pressure residual (CPR) preconditioner was introduced in the 80s and is still used today in commercial reservoir simulators. This method consists of a decoupling of the pressure equation followed by a two-stage preconditioning. The resulting linear system is solved using an iterative solution method.

We use a global preconditioner for the pressure equation, and a local preconditioner for the secondary equations.

Two-Stage Preconditioner

After applying a suitable decoupling operator to the linear system, we define a two-stage preconditioner as follows:

1. Solve the pressure system: $A_{pp}x_p = b_p$ for x_p;
2. Solve the full system and update x: $x = M^{-1}\,\hat{b} + \begin{bmatrix} x_p \\ 0 \end{bmatrix}$,

where $\hat{b} = b - A\begin{bmatrix} x_p \\ 0 \end{bmatrix}$, and the preconditioner M^{-1} is an approximation of A^{-1}. This two-stage process relies on a high quality decoupling. In fact, if the decoupling is not perfect ($A_{ps} \neq 0$), the error from this two-stage process is related to the norm of A_{ps}. Thus, it is key that the decoupling makes the A_{ps} block small.

The first stage of the preconditioner is where global preconditioning is used to solve the pressure system. An example of a global preconditioner is algebraic multigrid (AMG), which is known to preserve global properties of systems. The preconditioner in the second stage only needs to conserve local properties, so simple methods such as incomplete LU factorisation (ILU) are sufficient and very efficient.

Currently Used Decoupling Operators

Apart from weakening the couplings between equations, a decoupling operator should also improve or, at the very least, maintain desirable properties of the linear system. For example, the AMG preconditioner relies on the decoupled pressure block satisfying M-matrix properties. These properties are usually observed in elliptic-like systems such as the pressure equation. Consequently, the original pressure block usually satisfies them.

Only minor improvements have been proposed to the decoupling operators used in CPR since its creation. However, the models used in reservoir simulation are constantly evolving and the usual decoupling operators should be adapted to reflect this. Two widely used decoupling operators are true-IMPES (TI) and quasi-IMPES (QI). The QI decoupling operator reduces A_{ps} by setting its main diagonal blocks to zero. In many cases, this results in a high quality decoupling using QI. However, this technique does not consider the off-diagonal terms of A_{ps}. Hence, we will formulate methods that consider the whole of A_{ps} to reduce the coupling.

All the decoupling operators we considered are block diagonal matrices. This ensures that the decoupled system maintains a similar sparsity pattern to that of the original system. Moreover, this allows conservation of desirable properties of the pressure block such as M-matrix properties and prevent additional fill in the matrices.

Even if the TI and QI decoupling operators behave well most of the time, it is not always the case. Hence, valuable computational resources (e.g. AMG) are wasted on solving an inaccurate representation of the pressure equation. The main purpose of an improved decoupling operator is to have a more robust decoupling in the cases where QI and TI fail.

New Decoupling Operators

The LSQ decoupling operator minimises the norm of A_{ps} to reduce couplings.

Our first newly introduced decoupling operator, which we call the least squares (LSQ) decoupling operator, seeks to minimise the norm of A_{ps} to weaken the couplings between pressure and the secondary variables, whilst retaining the structure of a block diagonal matrix. The coefficients of each row in this decoupling operator are the solution of a least squares problem which minimises the Euclidean norm of the rows of A_{ps}.

In, the QI, TI and LSQ approaches, we only consider A_{ps} when choosing the coefficients inside the decoupling operator. However, the properties of the pressure block are very important to the convergence of the pressure solution. Hence, we create another decoupling operator inspired by the Dynamic Row Sum (DRS) preconditioner, which seeks to impose diagonal dominance in A_{pp}. To the minimisation problems in LSQ, we add a linear constraint on the coefficients adapted from the DRS condition. We denote this method LSQDRS.

We can generalise both the QI and LSQ approaches by using concepts for sparse approximate inverse (SPAI) preconditioners, a method for the approximation of a matrix inverse given a specific sparsity pattern. Indeed, the LSQ decoupling operator can be obtained by minimising the norm of A_{ps} while fixing the sparsity pattern of the decoupling operator to be block diagonal. By also considering only the terms in A_{ps} and A_{ss} within the block diagonal sparsity pattern, we recover QI.

Inside a reservoir, the properties of the rock and fluids can vary greatly, so different decoupling operators may be appropriate at different parts.

Reservoir simulations currently undertaken by the oil industry include very large domains where the properties of the rock and fluids can vary greatly. Consequently, different decoupling operators may be appropriate for different parts of the reservoir. This implies using a different decoupling strategy per row. In our case, we test a strategy where QI is used on some rows and LSQDRS on the others, depending on the properties of those rows. We denote this method QI-LSQDRS.

Another reason why the usual decoupling approaches fail may be that pressure is not the correct variable to decouple. For example, there are thermal cases where temperature can have a significant influence on the flow. In those cases, one option may simply be to use temperature instead of pressure as the primary variable to be decoupled. Temperature usually has a global influence in the reservoir, resulting in similar properties to those of pressure. The hope for AMG is that the newly considered block of temperature coefficients satisfies the M-matrix properties usually satisfied by the pressure block.

3. Results

In this section, we discuss some of the performance tests that were done for the reservoir simulator. We start by looking at the change of performance when using a different row scaling, and we then compare the performance of the newly introduced decoupling operators. Finally, we look at some of the properties of the system obtained from the decoupling in a case where the original QI and TI decoupling operators struggle.

Row Scaling

The new row scaling increases the performance of the solver in most test cases.

In order to prevent the decoupling operator from creating great variations in the pressure residual, we apply a row scaling. Since reservoir simulation models are constantly becoming more complex, it is important to consider adapting the previously used row scaling. To evaluate our new scaling, we have implemented 10 test simulations which are representative of various client cases. We calculate the time it takes to complete the simulation in each case, divide these by the time taken to run each test without our scaling, and then show the results in Figure 1. We see that we have improved the performance of the simulator in almost all cases; the best improvement (seen in Case 3) was 31%.

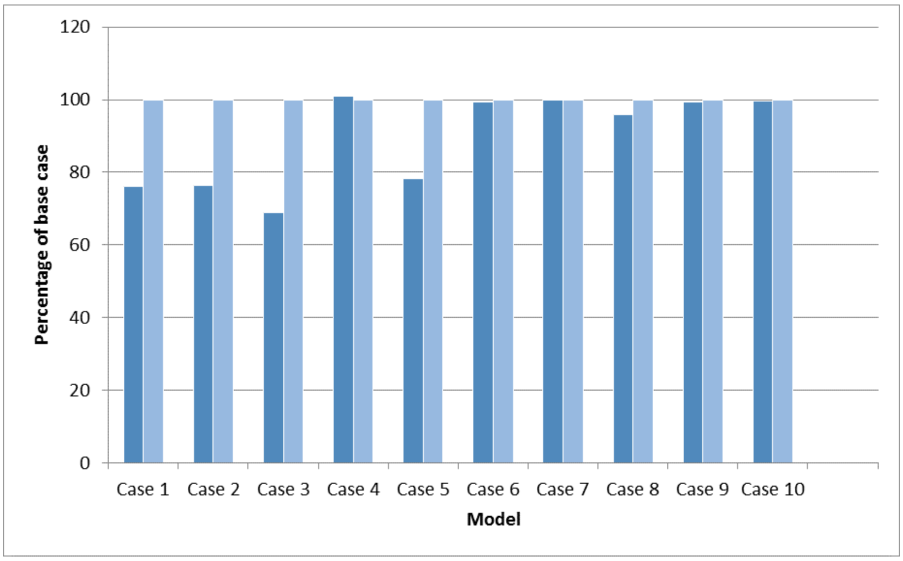

Figure 1: Relative elapsed time to solve the test cases using the new scaling (dark) and the original scaling (light).

Test Case

We now compare the performance of the reservoir simulator using some of the new decoupling operators for a single test case. The test case is a black-oil model with three phases present: oil, water and gas. During the simulation, water is injected in the reservoir to increase pressure. This part of the simulation is called the injection phase. We test the following decoupling operators: LSQ, LSQDRS, and QI-LSQDRS and compare them to the original TI and QI. In Figure 2, we illustrate the cumulative number of linear iterations needed to solve the test case at each point in the simulation. We observe that the LSQDRS decoupling operator does not perform as well as the others in the long term. We also see that QI-LSQDRS performs slightly better than QI, which is the second most efficient method. For this case, TI does not perform as well as QI, as it struggles more during the injection phase, represented by the step in the number of iterations, seen in the left-hand side of the data. The performance of the LSQ decoupling operator is not illustrated in Figure 2, because it fails to converge multiple times within the convergence criteria of the linear solver.

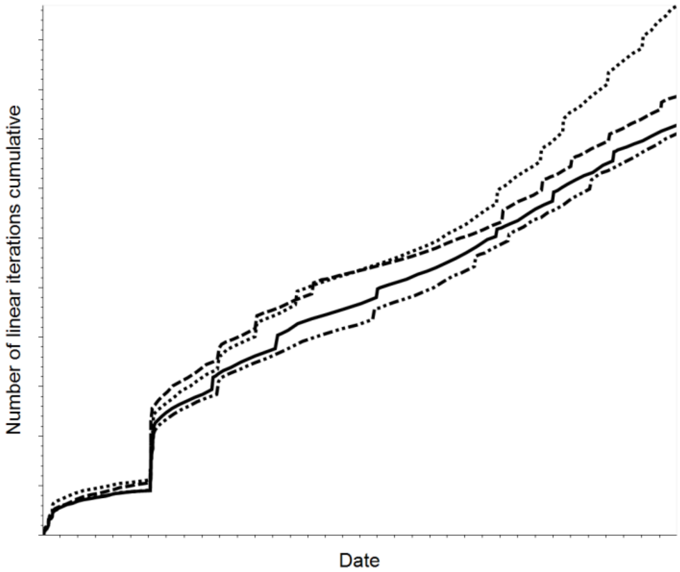

Figure 2: For a black-oil thermal model, comparison of the performance (number of linear iterations) of decoupling operators LSQDRS (dotted), TI (dashed), QI (solid) and QI-LSQDRS (dash-dot), as the reservoir simulation progresses in time. The numbers have been removed for confidentiality reasons.

Properties of the System

We now investigate the properties of a linear system from the test case. We choose a specific system where the linear solver performs better without the CPR preconditioning. In this case, using no decoupling results in convergence in a lower number of linear iterations than by using the QI or TI decoupling operators

We find that LSQ weakens couplings in the pressure system the most (small norm of A_{ps}), followed closely by QI. The parameters chosen for the DRS constraints are very stringent, which may explain why the couplings for LSQDRS and QI-LSQDRS are not weakened significantly. As for the decoupling of temperature, the reduction in the strength of the couplings is lower, but this can be explained by the fact that these couplings were not as strong to begin with.

We also observe (results not shown) that the pressure block obtained from the LSQ decoupling operator has many violations of M-matrix properties. This explains its poor performance because AMG was used to solve the pressure system. Nonetheless, these properties are not as important for other solvers. The other decoupling operators mostly satisfy M-matrix properties. Decoupling temperature instead of pressure also results in a temperature block that does not satisfy M-matrix properties for every row. This indicates that the temperature variable does not satisfy elliptic-like properties everywhere in the reservoir.

We discovered an additional desirable property for decoupling operators, which has to do with how much the pressure residual is modified

After investigating all the desirable properties a decoupling operator should have, we came to the realisation that poor performance of CPR might not be due to the violation of those properties. Indeed, we discovered a correlation between poor performance and an additional property that had not been considered before. This property is relative to how much the decoupling operator modifies the pressure residual b_p. We observe (results not shown) that the QI, TI and LSQ decoupling operators modify the pressure residual substantially. This discrepancy in the size of residuals may cause problems in the convergence of iterative solution methods. Because the LSQDRS approach allows us to put bounds on the size of the coefficients of the decoupling operator, the pressure residual retains much of its original magnitude. This was also observed for the QI-LSQDRS decoupling operator, which retains both good qualities of the QI and LSQDRS decoupling operators. Decoupling temperature also results in a decoupled system where the pressure residual is similar to the original one.

4. Discussion, Conclusions and Recommendations

In this report, we investigated the role of decoupling operators in the CPR preconditioner. We developed several decoupling operators which reduce the couplings in the pressure equation while retaining desirable properties for iterative solution methods such as AMG. During that process, we discovered an additional property related to the pressure residual, which provided an indicator of performance. Furthermore, we investigated the performance of the newly introduced decoupling operators, as well as the properties of the resulting linear systems.

Conserving the Pressure Residual

The goal of this project was to create decoupling operators that induce the desirable properties (small A_{ps}, M-Matrix) on the resulting system. However, satisfying these properties was not enough to improve the performance of the solver, which indicates that other considerations have to be taken into account. Indeed, we discovered a correlation between poor performance and how much the decoupling modifies the pressure residual. Further study will determine exactly what this property is, and how to measure and enforce it in a decoupling operator. One early finding from the LSQDRS decoupling suggests that bounds on the size of the coefficients in the decoupling operator may make it more robust.

Row-Dependent Decoupling Strategy

A key conclusion is that using different decoupling strategies for different rows is promising in order to tackle the properties of new reservoir models. The performance of the QI-LSQDRS decoupling operator in the thermal case indicates that this technique is viable. Additionally, the implementation and choice of parameters for the LSQDRS decoupling operator, as well as the choice of the condition to switch from QI to LSQDRS were not optimised by any means. Further investigation is needed into the choice of parameters for the multiple newly introduced decoupling operators. One advantage of this row-dependent decoupling approach is that its implementation does not require much change in the structure of the current implementation. Another simple alternative is to automate the use of different decoupling operators at different times in the simulation. Testing of different flavours of this method should be given priority.

Non-Pressure Decoupling

For thermal cases, temperature has a global influence in some cases. The original goal of CPR was to concentrate the global properties of the system inside the pressure block, so that global preconditioning can be done for this block and local preconditioning be done for the non-pressure system. Therefore, ignoring temperature or other variables with global influence may result in the local preconditioning being insufficient to conserve global properties. Although promising, the implementation of thermal decoupling would take considerable effort, and thus further testing should be done.

5. Potential Impact

Schlumberger is interested in increasing the performance of its reservoir simulation software and making it more robust to the various cases presented by its clients. Revisiting how the decoupling is done in CPR is essential for the software to remain cutting edge.

Tom Jönsthövel, Senior Scientific Software Engineer, Schlumberger, commented, "*The research on the algebraic decoupling has led to valuable insights in how to further optimise the performance of our simulators and to make even better use of available computational resources which is key to our clients.*"

Christopher Lemon, Software Engineer, Schlumberger, commented, *"The highly focused investigation has highlighted areas of potential development for our simulators. One of these has already shown promise for offering improved performance, and will be taken forward in a future release."*