EPSRC Centre for Doctoral Training in Industrially Focused Mathematical Modelling

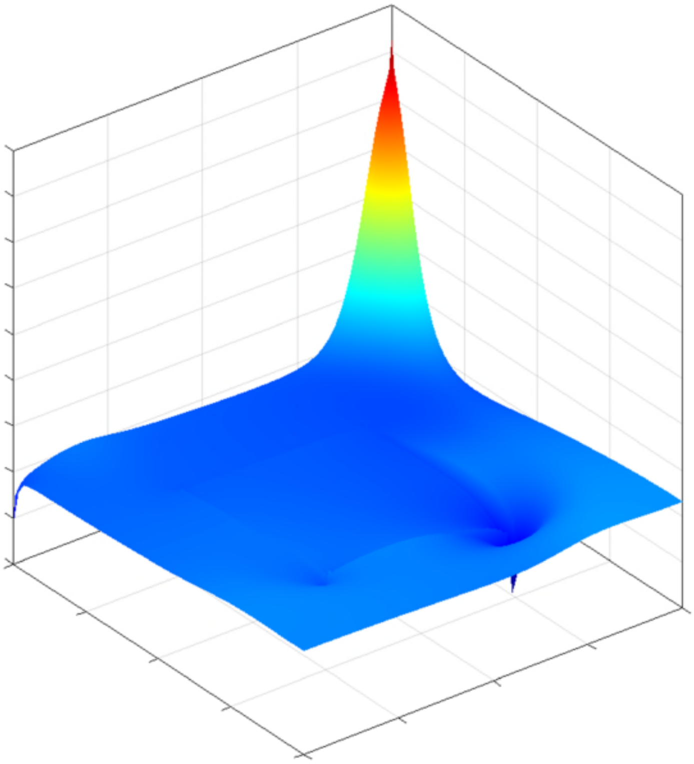

Block Preconditioning for Incompressible Two-Phase Flow

Niall Bootland

Table of Contents

1. Introduction

Background

Two-phase flows arise in many coastal and hydraulic engineering applications such as the study of coastal waves (see Figure 1) and the designing of channels and coastal structures. However, computational solutions of the incompressible Navier-Stokes models used in such applications are frequently limited by the computational cost of solving the two-phase flow models. One particular step of the numerical model dominates the computation time: the solution of linear systems derived from the variable density and variable viscosity Navier-Stokes equations. Preconditioners are crucial for the practical and effective solution of these linear systems.

Figure 1: A real air and water two-phase flow in action.

Mesh independence refers to the computational effort required being proportional to the mesh size.

A preconditioner can be thought of as a transformation of the linear system of equations to one which is more amenable to the iterative solution methods required for these large systems, while still allowing the solution of the original problem to be easily constructed. For one-phase flow problems, certain block preconditioners have been shown to be efficient solvers and can exhibit nearly mesh independent behaviour. The primary challenge for constructing such block preconditioners is to incorporate a faithful approximation of the Schur complement, a matrix appearing in the process, which can be more readily inverted. Here we consider if such preconditioners can be adapted for the case of two-phase flow.

The Reynolds number determines the relative importance between viscous and inertial effects.

Problems of interest at the US Army Coastal and Hydraulics Laboratory are typically high Reynolds number (10^3–10^7) flows computed on very fine meshes (10^6–10^8 vertices). Thus, of particular importance will be to consider how the preconditioners we develop depend on the Reynolds number and mesh size. We will also consider the effect of different viscosity and density ratios of the two fluids (air and water being the most common fluids in nature) and, in the time-dependent case, the length of the time-steps. We will see that newly adapted Schur complement preconditioners appear to offer significant benefits over the currently used “self-p” method (a simple and computationally inexpensive approximation to the Schur complement).

In the remainder of the introduction we will briefly describe the details of the problem and where our focus lies. We then explain the role of the preconditioner and describe the newly adapted methods developed. Our numerical results are summarised before we discuss conclusions and recommendations. Finally we outline the potential impact of our work.

We consider the simplest level set method along with the Navier-Stokes equations. We use a backwards Euler time-stepping scheme, the Picard nonlinear iteration method, and a finite element discretisation in space.

Focusing the problem

The process of solving numerically the full set of equations for incompressible two-phase flow using a level set method (which tracks where each of the fluids are) involves a number of different steps. We use a backwards Euler scheme to advance the solution in time, a finite element methodology to discretise the problem in space, and the Picard iteration method to treat the nonlinearity in the equations. A key approach to tackling the nonlinearity is to solve an approximating set of linear equations iteratively until the solution converges. Since such a method requires these linear equations to be solved many times, their efficient solution is vital to obtaining practical solution methods.

We consider a general form of the linear system required to be solved at one step of the overall solution process. This system depends on the parameters in the problem; namely the mesh size, the Reynolds number, the viscosity and density ratios and the time-step. Our results will illustrate how these parameters affect the performance of the preconditioners we test.

Our focus is in developing and testing preconditioners based on block factorisation of the linear system, these use the structure in the equations and an understanding of the iterative methods used to gain efficiency in the solution process. The ideas used to define the preconditioners are fairly general in their approach and thus can be used for different treatments of the nonlinearity, time-stepping methods and discretisations. We will see that different preconditioners also have differing benefits and drawbacks in their use, implementation, and performance.

2. Preconditioning

Iterative methods for linear systems

Very large linear systems are typically solved on parallel, high-performance computers using iterative algorithms. These algorithms are efficient and reliable but their efficiency varies with certain characteristics of the linear system, particularly a quantity known as the condition number which measures how much the solution changes for a small change of the input arguments. Unfortunately, linear systems arising from numerical approximations of two-phase flow have the undesired property that the condition number grows rapidly as the mesh is refined. We seek a strategy for counteracting this growth in condition number of the linear systems which utilises the power of preconditioners.

The preconditioning strategy

When using a preconditioner to transform the system of linear equations there is a balance between the preconditioner remaining a faithful representation to the problem and yet still allowing for the solution to be easily constructed. Our aim is to identify a computationally inexpensive transformation that limits the growth in the condition number of the resulting linear system, thus preserving the efficiency of the overall iterative method.

We consider two preconditioners known as PCD and LSC which we have adapted for use in the two-phase case. We also compare with a method known as self-p.

One effective approach for the linear systems arising from Stokes and Navier-Stokes equations is to construct a preconditioner from a 2-by-2 block factorisation of the linear system into a lower-triangular and an upper-triangular system. These blocks correspond to the velocity and pressure components of the system. In this approach a new block, known as the Schur complement, arises in the pressure space. In the one-phase case, efficient preconditioners have been successfully proposed based on computationally inexpensive approximations of the Schur complement. We seek to adapt these preconditioners for the case of two-phase flow problems. We consider two such preconditioners, called pressure convection diffusion (PCD) and least square commutator (LSC), which we will compare to the self-p method.

In adapting the PCD and LSC methods for two-phase flow we find that appropriate scaling is necessary to incorporate the two-phase nature of our problem. In particular, our numerical tests found that a viscosity scaling within the preconditioners greatly improved their performance and that, for one-phase flow, this modification reduces to the original PCD and LSC methods.

Comparison of the preconditioners

We now describe the various known benefits and drawbacks of the three preconditioners in terms of their use and implementation. The PCD method requires the construction of additional matrices and the correct choice of boundary conditions to apply is not fully understood; these problems are not encountered explicitly by the LSC method. On the other hand, PCD requires only one sub-problem solution per iteration while correspondingly LSC requires two. We note that the computational effort required by the self-p method is similar to that of PCD but does not require the construction of additional matrices. The PCD method, as well as self-p, is also readily applicable in the case where stabilised finite elements are used; the extension of the LSC methodology in the stabilised case is more involved.

Comments

We have adopted a practical approach: future theoretical work is needed to extend the current theory to multi-phase flow. Our approach will provide grounding on which to base our choice of preconditioners.

Our experience with these newly adapted preconditioners for incompressible two-phase flow comes primarily from our numerical experiments. We will see in our results that these preconditioners can exhibit much improved performance over the self-p method which is currently used in the US Army Coastal and Hydraulics Research Laboratory's computational facility. Our focus has been on the practical experimentation to see if such preconditioners are worth pursuing. To back up our numerical results, current theory will need to be extended to incorporate two-phase flow, primarily the variable density and viscosity. Future research providing such analysis may well yield a better understanding of the behaviour and trends that we see and stimulate new improvements in their performance and application. In particular, theoretical consideration of the choice of appropriate scaling could either confirm our choice or suggest more fruitful approaches; this may be especially beneficial for the PCD method.

3. Numerical Results

In our numerical experiments we consider a classic test problem, the driven cavity problem, adapted for two-phase flow. We focus on solving one instance of the linear system arising from the full problem; our experience suggests that this is representative in displaying the trends and behaviour of the preconditioners seen throughout the full solution and when other set-ups are used. We first consider a typical convergence profile for the preconditioners before presenting results on how the performance of each method depends upon the parameters in the problem.

Convergence profiles

A convergence profile shows how the error in the approximate solution reduces throughout the iteration procedure used to solve the linear system. To obtain an approximate solution satisfying a given tolerance criterion, we would ideally like the number of iterations required to be small so that less computational work is needed. We recall that the LSC method requires approximately twice as much work per iteration as the PCD method while the self-p method requires a similar amount of work to PCD. In Figure 2, typical convergence profiles are given for these three preconditioners; we now discuss what this tells us about their performance.

The PCD and LSC methods show much more rapid convergence than the self-p method. In particular, the best performance is seen in the LSC method which requires a tenth of the iterations that self-p needs.

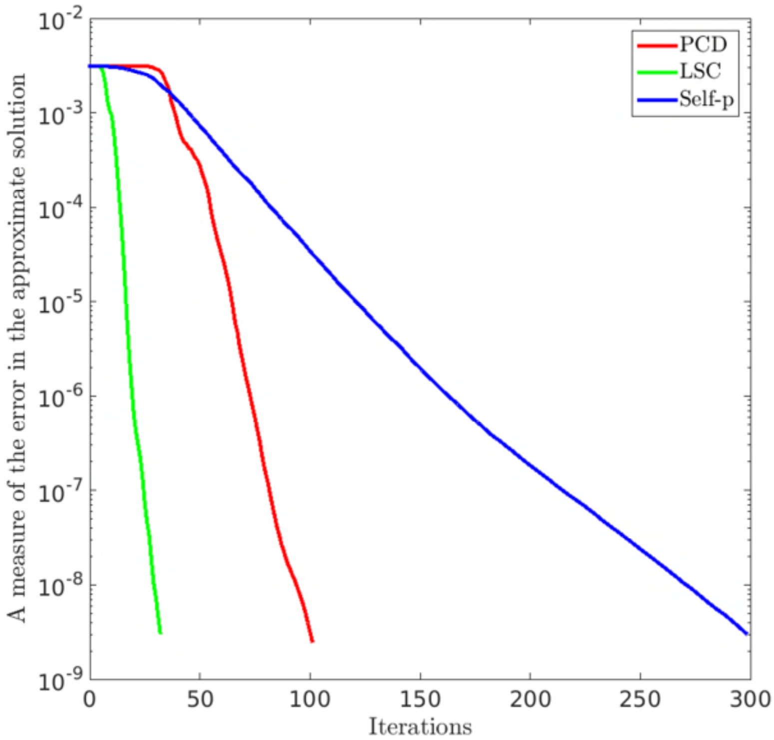

Figure 2: Comparison of typical convergence profiles for the newly adapted PCD and LSC methods as well as the self-p method. The plot shows a measure of the error over the course of the iterative procedure to solve the linear system until a required tolerance is satisfied. The ideal profile is one in which the error reduces rapidly to within the tolerance over a small number of iterations. The example shown here is for the steady problem with Reynolds number 1000, a viscosity and density ratio between the two fluids of 100, and a mesh of 4096 square elements.

What we learn from the convergence profiles is that the PCD and LSC methods display much more rapid convergence rates, since the corresponding gradients in Figure 2 are much steeper. We note however that the PCD method initially plateaus so that the fast rate of convergence does not kick in straight away. Despite LSC requiring roughly twice as much work per iteration it still offers the best performance and this preconditioner's superiority is a generic feature from our numerical results; indeed in all our tests it required the fewest number of iterations. On the other hand, the self-p method exhibits a much slower convergence rate, with a shallower gradient as seen in Figure 2. This provides an explanation for the poorer performance of self-p and suggests that it is not able to represent the Schur complement as successfully as the other methods.

Overall, these typical convergence profiles display characteristic features of the preconditioners. The LSC method performs very well and is seen to be robust. The PCD method can also perform well but is less robust; in particular the initial plateau appears to be obstructing better convergence and is seen to lengthen in certain cases giving poorer results. It may be that a different scaling or adaptation is needed within PCD to give the improved performance seen by LSC. Finally, self-p is hampered by its slow rate of convergence; this limits its effectiveness in applications, especially for the larger and more difficult problems which exacerbates the issue.

Dependence upon parameters

During our numerical tests we considered a range of mesh sizes, Reynolds numbers, viscosity and density ratios and length of time-steps. We found that the performance of the preconditioners were effectively independent of both the viscosity and density ratios. In the case of PCD and LSC we found that, for small Reynolds numbers, the number of iterations required increased slightly as the mesh became finer, while for larger Reynolds numbers, the number of iterations levels out, suggesting nearly mesh independent behaviour. However, for self-p, as the mesh becomes finer the amount of iterations required continues to increase, suggesting that this method is not as well suited for fine meshes. The dependence upon Reynolds number and time-step is interrelated, thus we split into two cases, the steady flow problem and the time-dependent problem.

In the case of the steady flow problem we find that, as the Reynolds number increases, the number of iterations required also increases. For coarse meshes this increase is larger, but this may be because the flow is not well resolved for such meshes. As the mesh becomes finer, the effect of increasing the Reynolds number reduces so that there is only a mild increase in the iteration count for the PCD and LSC methods. This is not true in the case of self-p; however, here the increase remains considerable. It is worth noting that, in real applications, meshes would typically be much finer than in our numerical experiments so the trend in mesh size is particularly important and it is only with PDC and LSC that we see nearly mesh independent behaviour.

The LSC preconditioner proves the most robust method and shows positive trends in its dependence upon the parameters. For problems of practical interest on fine meshes it is likely that LSC will beat self-p by an even greater factor than 10.

For the time-dependent problem we see differing behaviours in the PCD method when compared with the LSC and self-p methods. In the latter methods, we observe that as the time-step decreases the number of iterations required also decreases and that, once the time-step becomes small enough, the dependence on the Reynolds number changes so that now as the Reynolds number increases the iteration count decreases slightly. In contrast, for the PCD method the number of iterations increases with either a decrease in time-step or an increase in Reynolds number. Thus we see a qualitative difference between the PCD and LSC methods in the case of the time-dependent problem where PCD shows poorer performance. Again we see that LSC appears most robust and the method of choice.

4. Discussion, Conclusions and Recommendations

In our numerical study, we have considered two newly adapted preconditioners for equations describing incompressible two-phase flow, compared their performance, and discussed how this is affected by the parameters in the problem. We also made a comparison with the self-p method. Our results demonstrate that the most robust and efficient method is the least square commutator (LSC) preconditioner which, for the finer meshes and larger Reynolds numbers used in our tests, showed a ten-fold reduction in the number of iterations required to solve the arising linear system over the self-p method. The trends and behaviours observed suggests that, for problems of interest in real applications where much finer meshes and larger Reynolds numbers are encountered, this factor of improvement should only increase, thus providing tangible benefits by significantly reducing the computational cost of solving these problems.

Further work will include considering modifications to allow LSC to work with stabilised methods and performing a theoretical analysis of the methods to understand their behaviour better; this could further lead to improving the PCD method. A theoretical approach is likely to give an understanding of the predictability and scalability of these preconditioners.

We also note that the only method which shows positive trends for fine meshes, large Reynolds numbers and small time-steps is the LSC preconditioner. Looking at the implementation comparison we see that this would require using stable finite element methods; however this appears to be the only main drawback. Nevertheless the much improved performance seen by LSC appears very favourable and further work should be done to consider if modified versions of LSC, which allow for stabilised elements but are more complicated in their formulation, can also be adapted to work well for two-phase problems. The PCD method naturally extends to work with stabilised elements so further study into the theory of these methods might also yield improvements that see PCD become a more robust and favourable method.

A more theoretical approach will provide additional understanding of the properties of these preconditioners and such analysis will clarify the appropriate scaling for two-phase problems, in particular for the PCD method. This is likely to give an understanding of the predictability and scalability of these preconditioners.

Overall we have seen promising results from our study which suggest that significant benefits can be gained by a more careful choice of preconditioner and that the block preconditioners featured here may well prove fruitful in applications and successfully offer substantial increases in computational efficiency.

5. Potential Impact

The applications of incompressible two-phase flow are many and varied. Practical experimentation of such flows can be lengthy and expensive, thus an important tool is numerical simulation. However, this is limited by the computational efficiency of the solution methods used. For solving the large linear systems arising in the problem, the use of effective preconditioners, such as those developed here, could see substantial improvements in the computational runtime required for simulations and see significant savings in overall resource utilisation.

Dr. Chris Kees, Research Hydraulic Engineer, US Army Coastal and Hydraulics Laboratory, commented "*The dramatic reductions in iteration counts for LSC and PCD are impressive. If we see comparable reductions for realistic two-phase flow simulations, then this work could have a big impact through reductions in our computation time. Extension of the single-phase theory to multi-phase systems would be very helpful to guide our choice of preconditioner based on theory rather than experimentation.*"