Progress Report:
High-Resolution Numerical Methods for Compressible Multi-Phase
Flow in Hierarchical Porous Media

John A. Trangenstein
Duke University
Department of Mathematics
Box 90320
Durham, NC 27708-0320

March 15, 1993

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

RECEIVED
OCT 08 1993
OSTI

MASTER

EP

1 Personnel

This research program is being conducted by several people:

John Trangenstein is the principal investigator, and collaborates on all of the projects. During FY93, he has been supported jointly by this grant, and by the Defense Nuclear Agency for work in adaptive mesh refinement for wave propagation in elastic-plastic solids.

Sameer Khan is a post-doctoral fellow in Mathematics at Duke. He received his Ph.D. in Petroleum Engineering at the University of Texas in December 1992, and is supported jointly by this grant and a companion grant from NSF.

John Davies is the systems programmer in the Mathematics Department at Duke. He is supported jointly by this grant and the companion grant from NSF, by NIH grants awarded to Michael Reed and Harold Layton, and by the mathematics department.

Khaled Furati is a doctoral student in mathematics at Duke. He is supported by a fellowship from Saudi Arabia.

Richard Hornung is a doctoral student in mathematics at Duke.

2 Technical Progress

This is the first year in the proposed three-year effort to develop high-resolution numerical methods for multi-phase flow in hierarchical porous media. The issues being addressed in this research are:

Computational efficiency: Field-scale simulation of enhanced oil recovery, whether for energy production or aquifer remediation, is typically highly under-resolved. This is because rock transport properties vary on many scales, and because current numerical methods have low resolution.

Effective media properties: Since porous media are formed through complex geologic processes, they involve significant uncertainty and scale-dependence. Given this uncertainty, knowledge of ensemble averages of flow in porous media can be preferable to knowledge of flow in specific realizations of the reservoir. However, current models of effective properties do not represent the observed behavior very well. Relative permeability models present a good example of this problem. In practice, these models seldom provide realistic representations of hysteresis, interfacial tension effects or three-phase flow; there are no models that represent well all three effects simultaneously.

Wave propagation: It is common in the petroleum industry to assume that the models have the same well-posedness properties as the physical system.

An example of this fallacy is given by the three-phase relative permeability models; they were widely assumed by the petroleum community to produce hyperbolic systems for the Buckley-Leverett equations, but later the mathematics community proved that these models inherently produce local elliptic regions. Since numerical methods must use the models for computations, oscillations that develop could erroneously be attributed to numerical error rather than modeling difficulties. When wave propagation is well-understood by analysis, it can provide useful information about the expected behavior of the flows, and can assist in the development of high-resolution numerical methods.

During this year, we have made significant progress on several tasks aimed at addressing these issues:

Surfactant model: This model was chosen because of its applicability to both contaminant cleanup and enhanced oil recovery, its rich representation of physical processes, and the challenge of applying high-resolution numerical methods to it. In our formulation, this model involves seven components and from one to three phases. This makes it one of the most complicated models to which modern shock-capturing methods and adaptive mesh refinement have been applied.

Elliptic adaptive mesh refinement: This task is designed to extend adaptive mesh refinement in a new and important direction. Since multi-dimensional flow problems involve both hyperbolic behavior (through the transport of chemical species) and elliptic / parabolic behavior (through the pressure gradient involved in Darcy transport), it is reasonable to develop numerical methods that are well-suited to each type of behavior. Efficient treatment of hyperbolic conservation laws has been achieved through a combination of high-resolution schemes and adaptive mesh refinement. This task is designed to extend the benefits of adaptive mesh refinement to the elliptic / parabolic aspects of multi-dimensional flow in porous media. The approach we have chosen to handle this adaptive mesh refinement is particularly well-suited to a numerical study of effective media properties.

Analysis of wave behavior: Our third task addresses some analytical questions about a simple surfactant model, and hysteresis in relative permeabilities. Recent mathematical analysis by Johansen has raised the possibility that initial-value problems for the surfactant model have multiple solutions. Enhanced oil recovery models, such as the polymer and surfactant models, lead to the development of oil banks in enhanced oil recovery processes with continuous injections (Riemann problems). These oil banks correspond to non-monotone behavior in the saturations, and therefore trigger hysteresis. We are developing analytical solutions to these problems, both to understand the effect on the physical predictions, and to determine useful test problems for our numerical methods.

3D computation: Our fourth task involves the development of second-order Godunov methods for three-dimensional polymer flooding. Because of the intense computational requirements of three-dimensional simulation, the use of high-resolution numerical methods is particularly important to computational efficiency.

Visualization: We have developed effective visualization techniques for our simulators. These methods are useful for understanding fluid flow behavior, debugging, and collaboration at distance.

We will describe each of these tasks in detail.

2.1 Surfactant Model

The surfactant model we are studying was developed by Gary Pope and his students at the University of Texas Department of Petroleum Engineering. It models the micellar-polymer flooding process, which typically involves at least five stages:

Preflush: Since salinity can significantly lower the effectiveness of surfactants, the reservoir is first flushed with a fluid, such as water, to lower the salinity.

Slug: The injection of surfactant lowers the interfacial tension between oil and water, leading to a nearly complete sweep of oil from the pore spaces. This in turn leads to the development of an oil bank ahead of the surfactant. Often alcohol is injected with the surfactant to adjust the phase behavior.

Mobility buffer: In order to drive the slug across the reservoir, a fluid with viscosity comparable to that of the hydrocarbon is used. (Typically, this fluid is a mixture of water and polymer.) This serves to prevent the fluids in the remaining stages from developing viscous instabilities that channel through the slug and destroy the efficiency of the recovery process.

Taper: In this stage, the polymer used to thicken the water in the mobility buffer is gradually reduced. The engineering principle here is based on numerical simulation with low-order methods, which showed the possibility of recovery failure due to viscous instabilities that develop at the trailing edge of narrowly-tapered mobility buffers and punch through the slug. There is evidence from higher-resolution simulations that these tapers do not need to be as large as first thought.

Chase water: This is an inexpensive injection fluid, such as water.

This process is being used for enhanced oil recovery by Statoil, and has been examined by a number of the domestic oil companies in order to continue production from the mature domestic reservoirs. This process is also under study by Linda Abriola, of the Civil Engineering Department at the University of Michigan, for the removal of hydrocarbon contaminants from aquifers. Drs. Abriola

and Pope are currently funded jointly by the EPA to study the potential success of this process.

The surfactant model has led to the development of the UTCHEM simulator, which is distributed to essentially all of the domestic oil companies, as well as a number of foreign companies, through the Center for Petroleum and Geosystems Engineering (CPGE) at UT. This simulator is based on a fixed mesh, finite difference scheme using higher-order differencing in at most one or two variables (the relative permeabilities) along coordinate directions. The scheme is first-order in multiple dimensions, and whenever there are waves in addition to the Buckley-Leverett waves. It also requires timesteps that are significantly smaller than the Courant-Friedrich-Levy condition for numerical stability. This simulator is being modified for parallelization under a DOE HPCC grant jointly awarded to Gary Pope, and Mary Wheeler at the Department of Mathematical Sciences at Rice University.

We chose to study this model because it is one of the few enhanced oil recovery techniques that are being considered for use in aquifer remediation. (Another is steam flooding.) Thus we have the potential to influence two DOE missions at the same time, namely environment cleanup and energy production.

In FY 93 we studied both the available technical papers and the UTCHEM simulator to understand the surfactant model. We found several discontinuities in the phase behavior and transport properties, and suggested modifications to the models that removed these problems. These suggestions were presented to Gary Pope and given his approval. Further analysis also determined inexpensive and computationally stable analytic solutions for tasks that were solved iteratively in UTCHEM. We also found a number of opportunities to improve the numerical stability of the UTCHEM algorithms. Because of these difficulties in the original UTCHEM model, we completely reprogrammed the phase behavior model for our own use. The revised model was made available to Gary Pope, and may be transmitted on to the industrial subscribers of CPGE.

In order to develop a fully second-order Godunov method for this model, we continued the analysis of the model to determine a set of variables w that can determine the conserved total concentrations z and fluxes f without iteration. We call these variables the "flux variables." Although there are seven conserved quantities in z , there are only six flux variables w . This is natural because the concentrations z satisfy the constraint that they sum to one,

$$e^T z = 1,$$

where e is the vector of ones. A similar constraint is satisfied by the fluxes. If we represent the array of phase concentrations by C and the vector of phase velocities by v , then the flux f is given by

$$f = Cv.$$

The total fluid velocity

$$v_T = e^T v$$

is determined by the pressure equation and can be taken to be independent of the flux variables. Thus $e^T f$ is independent of w .

The seven conservation laws are not independent. Since the total concentrations z and the flux f are functions of the flux variables w , the conservation law can be written

$$\frac{\partial z}{\partial w} \phi \frac{\partial w}{\partial t} + \frac{\partial f}{\partial w} \frac{\partial w}{\partial x} = 0.$$

This non-conservative form of the conservation law is useful for the characteristic tracing step during the predictor phase of the second-order Godunov method. Now, $\frac{\partial z}{\partial w}$ and $\frac{\partial f}{\partial w}$ are 7×6 arrays with e in their left null-spaces. Assuming that $\frac{\partial z}{\partial w}$ has full rank (i.e., that the flux variables are independent), then the range of $\frac{\partial z}{\partial w}$ is the orthogonal complement of e , and the range of $\frac{\partial f}{\partial w}$ is contained in the range of $\frac{\partial z}{\partial w}$. As a result, it makes sense to rewrite the equation for w in the form

$$\frac{\partial w}{\partial t} + \left(\frac{\partial z}{\partial w} \right)^+ \frac{\partial f}{\partial w} \frac{1}{\phi} \frac{\partial w}{\partial x} = 0,$$

where $\left(\frac{\partial z}{\partial w} \right)^+$ is the pseudo-inverse of $\frac{\partial z}{\partial w}$. This gives us a system of six equations to consider, with six characteristic speeds given by the eigenvalues of

$$A \equiv \left(\frac{\partial z}{\partial w} \right)^+ \frac{\partial f}{\partial w} \frac{1}{\phi}.$$

Note that no information was lost in premultiplying the flux derivatives by the pseudo-inverse of the total concentration derivatives.

By working with the flux variables, we accomplish several important goals. First, we assure that the constraint that the concentrations sum to one is satisfied exactly. Secondly, we reduce the size of the system of equations. Thirdly, we achieve a better separation of the hyperbolic and elliptic aspects of the flow equations, by removing from consideration that part of the conservation law that determines the pressure equation. The selection of the flux variables is also essential to the success of the second-order Godunov method. By writing the characteristic structure in terms of the flux variables, we avoid having to call the iterative phase equilibrium calculations during the numerical solution of Riemann problems for flux computation. Furthermore, we can make simple tests on the values of the flux variables to guarantee that the intermediate stages of higher-order differencing stay within physical bounds. For example, the saturation of the non-aqueous phase is one of the flux variables for two-phase flow. If the numerical solution of the Riemann problem attempts to produce saturations that are not between zero and one at intermediate states, then we avoid the higher-order methods in favor of more diffusive flux computations that do not require Riemann problem solutions. These issues are somewhat subtle, and will be discussed in a forthcoming paper.

The hysteretic adsorption of surfactant, polymer and calcium in the model has so complicated our analysis of the model that we have not yet been able to show that the eigenvalues of A are necessarily real. At present, we have analytically determined the derivatives of z and f with respect to w , and have

used eispack routines to determine the characteristic speeds. We will test for complex characteristic speeds during the simulations, and use the numerical studies to drive future analysis of the characteristic structure.

All of our numerical computations of the phase behavior and partial derivatives have been programmed in a collection of Fortran subroutines. These routines are modular, so that they can be used with adaptive mesh refinement. We are in the process of inserting this model into the adaptive mesh refinement code at this time. This step should easily be completed before the end of the FY 93 project year.

Numerical experience with other models, such as gas dynamics and solid mechanics, has shown that 90% of simulation time is taken up by the numerical solution of the flow equations, and only 10% by the adaptive mesh strategy. Since Fortran is significantly more efficient on most machines (typically by a factor of 20 - 30 on our workstations), all of our surfactant model calculations are programmed in Fortran. On the other hand, the object-orientation features of C++ make it attractive for dealing with the complex data structures of adaptive mesh refinement. The most useful feature of C++ for our purposes has been the separation of the equation of state from the adaptive mesh refinement algorithm. The adaptive mesh refinement code stores a pointer to an abstract equation of state. At those points in the execution where some action must be taken on the data for the model, C++ uses "run-time binding" of the appropriate model procedure to perform the action. As a result, we now have users conducting adaptive mesh refinement simulations with gas dynamics, granular flow models and front-tracking, polymer flood models, and elastic-plastic solid mechanics models, in addition to the surfactant model. Each of these users can conduct their simulations independently of the other without modifications to the adaptive mesh refinement code.

The work on the surfactant model is being conducted jointly by John Trangenstein and Sameer Khan.

2.2 Elliptic Adaptive Mesh Refinement

Flow in porous media involves both hyperbolic and elliptic/parabolic behavior. While fluid fronts move slowly (on the order of a foot per day), pressure effects move rapidly (on the order of a mile per day). The pressure effects are typically modeled by summing the concentration conservation equations to obtain a divergence-free condition on the total fluid velocity, and then applying Darcy's law to obtain an elliptic pressure equation. Compressible flow is slightly more complicated, and will be treated in a later stage of this project.

While adaptive mesh refinement is well-understood for the conservation law, the use of adaptively refined meshes for equations with mixed character is more problematical. In order to provide a pressure and total fluid velocity to the conservation law, it is necessary to solve an elliptic pressure equation on the adaptively refined mesh. There are some special considerations here. First of all, the discrete gradient in the pressure equation needs to be the same as

the gradient used in the conservative difference. This is necessary so that the multi-dimensional treatment of transverse fluxes for diagonal transport in the conservation law is self-consistent. Secondly, the iterative techniques should be as efficient as possible, in order to prevent the pressure equation from dominating the computational effort. Thirdly, it would be helpful if the iterative technique were designed to assist in the numerical determination of effective media parameters.

Our approach is to use the same adaptive meshes for the pressure equation and the conservation law. This is reasonable, so that the pressure equation can provide accurate velocities to the conservation law, and so that the conservation law can provide accurate transport properties to the pressure equation. We are using adaptively refined meshes consisting of hierarchical overlying rectangular grid patches. Thus, at synchronization times, when coarse and fine grid patches have been integrated to the same time level, we must solve the pressure equation on a composite mesh, involving (possibly) both coarse and fine grid cells. Most of the time, we need only solve the pressure equation on the finest mesh at times intermediate to those on the coarser meshes.

Our approach is to use domain decomposition iterative techniques on patches at the same level of refinement, and to use multi-grid iterative techniques to communicate information between different levels of refinement. Domain decomposition is naturally suited to our use of rectangular arrays of grid cells, and allows the use of fast iterative methods on structured grids. Multi-grid iteration is a technique that communicates information globally in a very efficient way. It involves an averaging operator, to transfer information from fine grids to coarse, and an adjoint interpolation operator, to transfer information from coarse grids to fine.

Note that our use of hierarchical, overlying meshes is akin to the view of porous media as involving a hierarchy of scales. The effective properties of the porous medium, such as permeability and porosity, depend on the scale at which they are viewed. Various approaches by engineers and physicists have been developed to determine these effective media properties. We will be testing these ideas as potential averaging operators in our multi-grid iteration.

In FY 93 we have concentrated on developing the multi-grid and domain decomposition algorithms. These have been programmed and tested for one-dimensional advection-diffusion equations by Richard Hornung, as part of his thesis research under the direction of John Trangenstein. He has programmed a two-dimensional version of the algorithm for polymer flooding within the context of the adaptive mesh refinement code. In other words, he has written a Polymer equation of state C++ class, and Fortran routines for the computationally intensive work. Currently, he is debugging this code. Throughout the remainder of this year, he should be testing the code on two-dimensional cross-sections, and developing reliable error estimation procedures for the placement of refined patches.

2.3 Analysis of Wave Behavior

In FY 93, we examined two important problems: uniqueness of Riemann problem solutions for surfactant flooding, and Riemann problems for enhanced oil recovery problems involving hysteresis in relative permeabilities.

Work on the analysis of Riemann problems for polymer flooding dates back to work by Pope in 1980. In recent years, Johansen has analyzed the Riemann problem for simple surfactant models involving three components and 1-2 phases. He found cases where the solution is not uniquely determined by wave matching. This is potentially a problem for the work in this project, since we are using a surfactant model that, although more complicated, shares some of the same features as Johansen's model.

During FY 93, Khaled Furati has examined this problem as part of his dissertation work at Duke. He was able to apply an argument by Smoller and Temple to show that only one of the solutions is obtained in the limit of vanishing viscosity.

In the second part of his thesis work, Khaled Furati is also conducting research into the effect of relative permeability hysteresis on the structure of fluid fronts. Hysteresis in relative permeabilities is an important feature of multi-phase flow in porous media. Typically, reservoir simulation is conducted under the assumption that the saturations are either strictly increasing (*e.g.*, water) or strictly decreasing. The simulators are then given the relative permeability curves for the appropriate imbibition or drainage behavior. However, many enhanced oil recovery processes involve the development of oil banks, in which the water saturation first decreases and then increases. This means that hysteresis could be important to the recovery predictions.

In order to study the qualitative effects of hysteresis on enhanced oil recovery, Khaled Furati began by constructing analytic solutions to Riemann problems for two-phase Buckley-Leverett flow with hysteresis. This problem has many shared features with the work by Trangenstein and Pember on the Riemann problem for one-dimensional elastic-plastic solids. Khaled is working on extending the analysis to polymer floods and Johansen's surfactant model. The preliminary results for the polymer model indicate that the inclusion of hysteresis effects in problems resembling injection of water with polymer into an oil reservoir can cause the oil slug to be not as high in saturation and wider in spatial extent than would be predicted without hysteresis. We expect to find other, more pronounced effects on the structure of the fluid fronts when this research is complete.

2.4 3D Computations

Under contract with Statoil during the spring of 1992, we developed a 3D polymer flooding simulator. This simulator used a second-order version of Godunov's method suggested by Jeff Salzman. Under that contract we extended the Salzman algorithm to multi-phase flow by incorporating a pressure equation, and treating the velocity field as a spatially-varying quantity in the conservation

laws. The code used spatially-varying permeability tensors, but the test problem initialized the permeability field to the identity. There were two wells: a rate-specified injector in one lower corner of the cubical reservoir, and a pressure-specified producer in the opposite corner. The injector well was designed to inject at constant water saturation and polymer concentration, and the pressure equation was solved only once, at initialization. The code was developed under a short contract, principally to prove that it could be done.

During the current fiscal year, we made some minor corrections to the 3D polymer code. This work was conducted jointly by John Trangenstein and Sameer Khan. We modified the injection strategy to inject a polymer slug followed by pure water without any taper. We also modified the code to solve the pressure equation for the velocity field at each timestep. Although we could afford to run only on a relatively coarse mesh on our workstations ($17 \times 17 \times 17 = 4913$ cells for 7000 timesteps), the numerical results showed very nice resolution of the saturation fronts, and very little grid orientation. The revised code was recently transferred to Statoil; they are assisting us by making more detailed runs on the Cray at the University of Trondheim. Our work on the 3D code in this project was for the principal purpose of examining 3D visualization capabilities.

2.5 Visualization

Visualization is an important ingredient in scientific computation. Effective visualization tools can give the researcher important information about the state of the computation, and can be used effectively as debugging tools in certain situations. The combined use of color and shape is needed to give the mind maximum opportunity to obtain information quickly.

The choice and form of the visualization tools is important as well. For collaboration at a distance, portable and universal software is important, such as the XWindow system. For collaboration within a group, modular graphics software is useful to insulate students from the details of the graphics internal operations. In this regard, the XWindow system is not very advanced.

We have written some 1D and 2D graphics tools, using the XWindow system and C++, for special use with adaptive mesh refinement. For example, in 2D we draw the pictures by coloring coarse cells first and fine cells last as we examine the values on all patches on all levels of mesh refinement. Other modular aspects of the code allow the user to specify the name of variables desired for plotting on the command line, without having made any changes to the adaptive mesh refinement code at all. In this case, the plot requests are passed to the user-specified equation of state for processing. The mathematics students at Duke have found this graphics environment relatively easy to incorporate into their codes; as a result, even students working with other thesis advisors are using these tools. We also plan to make these graphics tools available to the Petroleum Engineering Department at the University of Texas.

For 3D visualization, we are using the Iris Explorer system. This is a new

product, and shares many of the same features as the AVS system marketed by Stardent. Explorer is very simple to learn, and very effective for visualizing the 3D polymer flood described in the previous section. We have used this system to prepare a short movie for presentation at the DOE project review in Bethesda. This work was conducted jointly by John Trangenstein, John Davies and Sameer Khan.

3 Future Plans

During FY 94, we plan to continue our work on the surfactant model, adaptive mesh refinement and effective media properties, and graphics.

3.1 Surfactant Model

We plan to make extensive tests of the second-order Godunov method with adaptive mesh refinement in 1D and 2D. In this effort, we will conduct numerical experiments to guarantee the reliability of the phase equilibrium calculations and to discover the structure of the fluid fronts. We will also refine our strategies for the approximate solution of Riemann problems. We plan to make extensive computational tests to examine mesh refinement and de-refinement strategies. We hope to compare our numerical methods to standard strategies in order to determine the relative efficiency of the second-order Godunov method to standard difference methods in the petroleum community. We would also like to determine the relative sizes of numerical dissipation and capillary pressure, in order to determine the conditions under which the latter is significant. Finally, we plan to continue the analysis of the characteristic structure of the flow equations, possibly under the assumption that the surfactant, polymer and calcium adsorptions are independent of fluid composition. We hope to be able to extend those results to the case in which one of the adsorptions varies with the fluid, and to choose the particular adsorption that the computations indicate is most important. This work will be conducted principally by Sameer Khan, in close cooperation with John Trangenstein.

3.2 Adaptive Mesh Refinement

The two-phase Buckley-Leverett model and three-component two-phase polymer model are ideal models with which to test adaptive mesh refinement strategies for multi-dimensional reservoir flow. Both are simple enough to permit extensive mathematical analysis, but rich enough to produce interesting wave behavior. We plan to conduct extensive tests to determine effective iterative strategies for the adaptively-refined pressure equation. We will also examine error estimation techniques in order to guarantee effective placement of the refined mesh. Afterward, we plan to use the simulator to examine the automatic computation of effective media properties through our multigrid averaging operator. We will generate random permeability fields and examine the ability of the averaging

operator to predict effective permeability tensors at coarser scales. Here, we will have to allow for the average permeability tensor to be anisotropic and possibly even non-symmetric; this may complicate the techniques for the iterative solution of the pressure equation on the adaptively-refined grid. Of course, we will depend on our connections with the University of Texas, Stanford University, Statoll and British Petroleum to give us data on interesting flow fields and test our ability to predict the average permeability. This work will be conducted principally by Richard Hornung, under the advice of John Trangenstein.

3.3 Graphics

Since the projects we plan for FY94 are principally 1D and 2D problems, we will concentrate our efforts on making graphical display as effective as possible for those problems. We plan to develop a graphical user interface for use with the surfactant model, in order to make the code as user-friendly as possible. This can easily be done with general-purpose software packages for creating these graphical user interfaces.

We also plan to investigate the incorporation of Iris Inventor into our adaptive mesh refinement algorithm. The Iris Inventor package is a collection of C++ classes that form the framework for the Iris Explorer graphics environment. We hope that this package will give us more interactive control over colormaps for 1D and 2D graphics. Looking to FY95, we also hope that we will be able to control Iris Inventor to create efficient graphical displays of 3D adaptively refined calculations. This work will be conducted by John Davies, in cooperation with John Trangenstein.

END

**DATE
FILMED
11/10/1993**