Centrum voor Wiskunde en Informatica

REPORT*RAPPORT*

VLUGR3: A Vectorizable Adaptive Grid Solver for PDEs in 3D.
I. Algorithmic Aspects and Applications

J.G. Blom, J.G. Verwer

Department of Numerical Mathematics

VLUGR3:
A Vectorizable Adaptive Grid Solver for PDEs in 3D
I. Algorithmic Aspects and Applications

J.G. Blom, J.G. Verwer

CWI

P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

Abstract

This paper describes an adaptive-grid finite-difference solver for time-dependent three-dimensional systems of partial differential equations. The robustness and the efficiency of the solver, both for vector and scalar processors, is illustrated by the application of the code to three example problems.

1. INTRODUCTION

To solve time-dependent partial differential equations (PDEs) often a straightforward Method of Lines (MoL) approach is used. The PDE is discretized in space and the resulting system of ODEs or DAEs is solved in time with a robust 'off-the-shelf' solver. The spatial grid used can be either uniform or nonuniform but is invariable for the time interval. For problems with steep and moving fronts this can be an inefficient approach since too much grid points are needed at areas where the solution (no longer) has large spatial gradients. For this reason many attempts have been made to develop adaptive-grid methods, where the grid is adjusted to the solution either dynamically or statically (e.g., after each time step). This introduces of course always a certain amount of overhead but when the use of implicit time integrators is profitable (in the case of DAEs it is even required) or when the number of PDE components is very large such solvers can be more efficient already in two dimensions. In three dimensions the standard approach often will no longer be feasible, not only because of the inefficiency of the computations, but even more so because the physical memory of most computers will not be large enough to contain the solution values at all grid points and the work space needed for the solution of the nonlinear system.

In [16, 17, 18, 20, 19, 21, 15] an adaptive-grid finite-difference method was developed to solve time-dependent two-dimensional systems of PDEs. This method couples Local Uniform Grid Refinement (LUGR) with an implicit ODE solver and proved to be robust and efficient for the target problem class. In [3] the vectorized implementation of this method in the code VLUGR2 is described which is based on the research code MOORKOP[13] by Trompert. In this paper we describe the algorithmic aspects of VLUGR3 which is the extension of VLUGR2 to the three-dimensional case (see also [5]). This vectorizable adaptive-grid finite-difference code has been designed to solve initial-boundary-value problems for systems that fit in the following master equation

$$\mathcal{F}(t,x,y,z,\mathbf{u},\mathbf{u}_t,\mathbf{u}_x,\mathbf{u}_y,\mathbf{u}_z,\mathbf{u}_{xx},\mathbf{u}_{yy},\mathbf{u}_{zz},\mathbf{u}_{xy},\mathbf{u}_{xz},\mathbf{u}_{yz}) = 0, \quad (x,y,z) \in \Omega, \quad t > t_0, \qquad (1.1)$$

where the solution $\mathbf{u}$ may be a vector and the domain Ω an arbitrary domain that can be described by a, possibly disjunct, 'brick' structure, i.e., a domain that can be described by right-angled polyhedrons. The boundary conditions belonging to system (1.1) are formulated as

$$\mathcal{B}(t, x, y, z, \mathbf{u}, \mathbf{u}_t, \mathbf{u}_x, \mathbf{u}_y, \mathbf{u}_z) = 0, \quad (x, y, z) \in \partial\Omega, \quad t > t_0, \tag{1.2}$$

and the initial conditions satisfy

$$\mathbf{u}(t_0, x, y, z) = \mathbf{u}_0(x, y, z), \quad (x, y, z) \in \Omega \cup \partial\Omega. \tag{1.3}$$

The paper is an extension of the report[2], where some first test results were discussed. There the nonlinear systems were solved with modified Newton and the linear systems with BiCGStab[22] with ILU preconditioning. Since the test results in [2] indicated that memory demands could be a bottleneck for some problems, we added to our code a matrix-free implementation of the Newton process with as linear system solver GCRO[12]. As preconditioner we use in this case a simple (block) diagonal scaling.

To evaluate as well the robustness and the efficiency of the adaptive-grid algorithm in 3D as the efficiency of the implementation and its performance on a vector processor, we tested the code extensively on a set of three example evolutionary problems in 3D with various solution characteristics, viz.,

- Burgers' equation with as solution a wave moving skew through the unit cube,

- a rotating sphere problem described by a system of two advection-reaction equations, and

- a model for the simulation of pollution in ground-water flow.

The paper is organized as follows. In Section 2 we give an outline of the local-uniform-grid-refinement method as implemented in `VLUGR3` and the choices made in the code with respect to the refinement strategy and the time-integration strategy, including the solution of the (non)linear systems. For a more detailed discussion on the subject of the implementation of the code and how to use it we refer to the companion paper [5]. Section 3 contains the test results for the three example problems. In this section we also look into the robustness of the code on the basis of some numerical experiments. In Section 4 we discuss the (vector) performance of the solver for the example problems. Finally, Section 5 contains a summary of our findings.

2. THE ALGORITHM
2.1 Local Uniform Grid Refinement
There are two reasons to adapt a spatial grid in time to the solution. The first is to speed up the computations since the PDE is evaluated in significantly less points than what would have been the case if one uses a fixed grid. The second reason is to restrict the amount of memory needed. Naturally both goals should be reached without sacrificing overall accuracy.

The advantage of static regridding, and especially of local uniform grid refinement is that it is robust and, conceptually, very simple. The domain is covered by a, uniform, coarse base grid and nested finer uniform subgrids are recursively created in regions with high spatial activity. So all grids consist of one or more disjunct sets of interconnected grid cells, all having the same size. When a grid of a specific refinement level has been created the corresponding initial boundary value problem is solved on the current time interval.

In short, an implicit solver based on the LUGR method can be described by

0. Start with the coarse base grid, the initial solution and an initial time step

1. Solve the IBVP on the current grid with the current time step

2. If the required resolution in space is not yet reached:

   (a) Determine at forward time level a new, embedded, grid of next finer grid level
   (b) Get solution values at previous time level(s) on the new grid
   (c) Interpolate internal boundary values from old grid at forward time
   (d) Get initial values for the Newton process at forward time
   (e) Goto 1.

3. Inject fine grid solution values in coinciding coarser grid nodes

4. Estimate error in time-integration. If time error is acceptable

   • Advance time level

5. Determine new step size, goto 1 with coarse base grid as current grid.

Where interpolation is needed to obtain solution values, linear interpolation is used.

The virtue of a LUGR method lies in the fact that one reaches the accuracy of a fine mesh with considerably less computational effort and memory requirements, since the fine subgrids cover only part of the domain. A necessity for the good functioning of a LUGR method is that the refinement strategy takes care that fine subgrids are created timely and at the correct places. In [17, 18, 20, 21] the refinement strategy is based on a comprehensive error analysis of both the space discretization and the interpolation. Here we use a much more simple strategy because class (1.1)-(1.3) is too general for the error analysis to be valid. The refinement strategy is the same as in VLUGR2, i.e., based on a curvature monitor.

For each grid point (i, j, k) the space monitor is determined by

$$\texttt{SPCMON}(i, j, k) :=$$
$$\max_{ic=1,\texttt{NPDE}} \texttt{SPCTOL}(ic) \cdot \left(|\Delta x^2 \cdot u_{xx}^{ic}(i, j, k)| + |\Delta y^2 \cdot u_{yy}^{ic}(i, j, k)| + |\Delta z^2 \cdot u_{zz}^{ic}(i, j, k)| \right) \qquad (2.1)$$

where Δx, Δy, and Δz are the grid width in the x-, the y-direction, and the z-direction respectively, and

$$\texttt{SPCTOL}(ic) := \frac{\texttt{SPCWGT}(ic)}{\texttt{UMAX}(ic) \cdot \texttt{TOLS}}. \qquad (2.2)$$

The variables on the right-hand side of (2.2) are user specified quantities. Specifically, $0 \leq \texttt{SPCWGT} \leq 1$ is a weighting factor for the relative importance of a PDE component on the space monitor, UMAX the, approximate, maximum absolute value for each component, and TOLS the space tolerance. The second-order derivatives in (2.1) are approximated by second-order finite differences at the internal domain and first order at the boundaries.

A next level of refinement is created if

$$\max_{(i,j,k)} \texttt{SPCMON}(i, j, k) > \texttt{TOLWGT}. \qquad (2.3)$$

TOLWGT acts as a bar against fluctuations of the number of grid levels in subsequent time steps: if the next level of refinement existed at the previous time level $\texttt{TOLWGT} = 0.9$, otherwise $\texttt{TOLWGT} = 1.0$.

If a next level of refinement is required, then all grid points with

$$\texttt{SPCMON}(i, j, k) > 1/4 \qquad (2.4)$$

are flagged together with their 26 direct neighbors. A cell with at least one flagged corner is divided in 8 equal subcells, so we do not apply semi-refinement, i.e., refinement in only one direction parallel to an axis. The next grid level will contain all these refined cells.

2.2 Discretizing the PDE

In `VLUGR3` standard second-order finite differences are used for the spatial discretization, central on the internal domain and one-sided at the boundaries. Due to the mixed derivatives we thus deal with a 19 point coupled grid stencil.

As time integrator we use the second-order two-step implicit BDF method with variable step sizes

$$\mathbf{U}_t = a_0 \mathbf{U}^{n+1} + a_1 \mathbf{U}^n + a_2 \mathbf{U}^{n-1}, \tag{2.5a}$$

with

$$a_0 = \frac{1 + 2\alpha}{1 + \alpha} \frac{1}{\Delta t}, \ a_1 = -\frac{(1 + \alpha)^2}{1 + \alpha} \frac{1}{\Delta t}, \ a_2 = \frac{\alpha^2}{1 + \alpha} \frac{1}{\Delta t}, \ \text{and} \ \alpha = \frac{\Delta t}{\Delta t_{old}}. \tag{2.5b}$$

In the first time step we apply as usual Backward Euler ($\alpha = 0$ in (2.5)). Like in DASPK[8], we use for all error measurements in the time integration a weighted root-mean-square norm

$$\|\mathbf{v}\|_w = \|W\mathbf{v}\|_2, \tag{2.6}$$

where W is a diagonal matrix defined by

$$W = 1/\sqrt{N} \ diag(w_1, ..., w_N). \tag{2.7}$$

The time integration is controlled by the solution monitor value in time which is computed at each existing grid level

$$\texttt{TIMMON}(level) := \|\Delta t \ \mathbf{u}_t\|_w, \tag{2.8}$$

where Δt is the current time step size, $\mathbf{u}_t$ is approximated by first-order finite differences and the entries w_i of the diagonal matrix W from (2.7) are defined by

$$w_{ipt,ic} = \texttt{TIMWGT}(ic)/(\texttt{ABSTOL}(ic) + |U^{n+1}_{ipt,ic}| \cdot \texttt{RELTOL}(ic)), \quad ipt \in \Omega \ \text{and} \ ic = 1, \texttt{NPDE}, \tag{2.9}$$

with

$$\texttt{ABSTOL}(ic) = 0.01 \cdot \texttt{TOLT} \cdot \texttt{UMAX}(ic) \ \text{and} \ \texttt{RELTOL}(ic) = \texttt{TOLT}. \tag{2.10}$$

Note that the time monitor is first order although the used time-integration method is second order. The reason is that an estimator based on, for one time level, interpolated solution values, and for another computed, leads readily to unnecessarily small time steps (cf. [16]). This is also the reason to exclude the boundary points in (2.8). The variables `TIMWGT` and `TOLT` are the, user specified, analogues of the variables in (2.2).

An integration step is rejected and redone at all grid levels if

$$\max_{level} \texttt{TIMMON}(level) > 1.0. \tag{2.11}$$

A new step size is computed such that the prediction of the monitor at the next time point is 0.5, i.e.,

$$\Delta t_{new} := \frac{0.5}{\max_{level} \texttt{TIMMON}(level)} \cdot \Delta t. \tag{2.12}$$

If the step was accepted the increase in step size is restricted to a factor 2 and if the time step was rejected the decrease is restricted to a factor 4. Finally, the step size is restricted to a user specified minimum and maximum value and adjusted such that the rest of the integration interval is an integer number times Δt.

2.3 Solving the nonlinear system

Since we use an implicit time-integration method, each time step a large system of nonlinear algebraic equations has to be solved. This is done using Newton's method in combination with a preconditioned iterative linear system solver. In our experiments this combination was more robust and most of the times less expensive than nonlinear Krylov solvers like, e.g., Nonlinear Orthomin[10] or so-called hybrid Krylov methods in which it is tried to speed up the convergence of the Newton process using information from the linear Krylov solver (see [9]). `VLUGR3` offers two implementations of the Newton process: one where the Jacobian matrix

$$G = \frac{\partial \mathbf{F}}{\partial \mathbf{U}} \tag{2.13}$$

is computed and stored, once for each grid level and each time step, and a matrix-free variant. In the latter the matrix-vector product $G\,\mathbf{v}$ required in the linear system solver is approximated by a difference quotient. Note that the first option results in a *modified* Newton process, whereas the second option leads to a *true* Newton process. So the latter can have a better convergence behavior if the matrix-vector product is well enough approximated. Also note that in contrast to the common practice in MoL, it is not possible to 'freeze' the Jacobian and/or the preconditioner for a number of time steps, since the location and number of grid points of the refined grids can change every time step.

The linear system solvers are described in Section 2.4 and the approximation of the Jacobian and the matrix-vector product $G\,\mathbf{v}$ in Section 2.5. Here we discuss some of the strategies implemented in the Newton procedure. Roughly spoken we have followed the approach as used in DASSL[7, pages 123–124] and DASPK[8]. When solving the nonlinear system

$$\mathbf{F}(t, x, y, z, \mathbf{U}, \mathbf{U}_t, \mathbf{U}_x, \mathbf{U}_y, \mathbf{U}_z, \mathbf{U}_{xx}, \mathbf{U}_{yy}, \mathbf{U}_{zz}, \mathbf{U}_{xy}, \mathbf{U}_{xz}, \mathbf{U}_{yz}) = 0 \quad t > t_0, \tag{2.14}$$

which is the fully discretized form of (1.1) and (1.2), Newton generates linear systems of the form

$$G\,\mathbf{c}^{(k)} = -\mathbf{F}(\mathbf{U}^{(k)}). \tag{2.15}$$

These systems are solved with an arbitrary linear system solver and the solution is updated

$$\mathbf{U}^{(k+1)} = \mathbf{U}^{(k)} + \mathbf{c}^{(k)}. \tag{2.16}$$

The Newton process is continued until the iteration error $\|\mathbf{U} - \mathbf{U}^{(k)}\|_w$ is sufficiently small. For the Newton process and the underlying linear solvers the entries w_i of the diagonal matrix W in (2.7) are defined by

$$w_{ipt,ic} = 1.0/(\text{ATOL}(ic) + |U^{(0)}_{ipt,ic}| \cdot \text{RTOL}(ic)), \quad ipt \in \Omega \text{ and } ic = 1, \text{NPDE}, \tag{2.17}$$

with

$$\text{ATOL}(ic) = 0.01 \cdot \text{TOL} \cdot \text{UMAX}(ic) \text{ and } \text{RTOL}(ic) = \text{TOL}. \tag{2.18}$$

The user-set space and time tolerances `TOLS`, `TOLT` are buried in the tolerance

$$\text{TOL} = 0.1 \min(\text{TOLT}^2, \text{TOLS}). \tag{2.19}$$

Assuming convergence of the Newton process, the inequality

$$\|\mathbf{U} - \mathbf{U}^{(k)}\|_w \le \frac{\rho}{1-\rho} \|\mathbf{U}^{(k)} - \mathbf{U}^{(k-1)}\|_w, \tag{2.20}$$

holds, where ρ is the convergence rate which is in actual computation approximated by

$$\rho \approx \rho^{(k)} = \left(\frac{\|\mathbf{U}^{(k)} - \mathbf{U}^{(k-1)}\|_w}{\|\mathbf{U}^{(k-1)} - \mathbf{U}^{(k-2)}\|_w} \right)^{\frac{1}{2}}. \tag{2.21}$$

This leads to the stopping criterion for the Newton iteration

$$\frac{\rho^{(k)}}{1 - \rho^{(k)}} \|\mathbf{U}^{(k)} - \mathbf{U}^{(k-1)}\|_w < \texttt{TOLNEW} = 1.0. \tag{2.22}$$

Note that we use the square of the user-defined time tolerance, `TOLT`, because the BDF formula (2.5) is second order and the time monitor is first order.

If during the Newton iteration $\rho^{(k)} > 0.9$ or if the maximum number of iterations is exceeded, a new Jacobian and/or preconditioner is computed, once, and the iteration is restarted. If the Newton process does not converge with the new Jacobian / preconditioner the time step is redone with $\Delta t = \Delta t/4$.

An extra difficulty when solving nonlinear systems arising from the LUGR method is to provide an initial solution. If one employs the method of lines on a single grid fixed in time the solution of the previous time level is in general a good initial estimate. In any case, if $\Delta t \to 0$ it is *the* solution. This is, however, not the case with LUGR methods. The solution values injected from a finer grid are in general *not* a solution of the PDE system discretized on a coarser grid. So, even when $\Delta t \to 0$ it is still possible that the injected values on the previous time level are not a good initial estimate. This obviously is an exceptional situation, but for the nonlinear ground-water-flow problems we indeed experienced some serious problems with the convergence of the Newton process. So, for the sake of robustness, we also save the not-injected solution values, i.e., the solution values actually computed at the coarser grid and use these as initial estimate, with interpolated values where a grid point at the previous time did not exist.

2.4 Solving the linear system

The Conjugate Gradient method is for symmetric systems an ideal iterative solver. It has the minimal residual property, i.e., the solution vector in the current Krylov subspace is chosen such that the residual vector is orthogonal to the subspace and it can be efficiently implemented using a three-term recurrency. Krylov solvers for nonsymmetric systems can roughly be divided in two groups. In the first the minimal residual property is maintained at the cost of a more expensive computational process (GCR, ORTHOMIN, GMRES). In the second the three-term recurrency is maintained and a nonoptimal *descent* algorithm results (BiCG, CGS, BiCGStab).

In `VLUGR3` we implemented two Krylov solvers to solve the linear system (2.15). If the Jacobian matrix G and its ILU decomposition are available we use BiCGStab[22]. BiCGStab is a cheap solver (in memory and in CPU time) that appeared to be robust enough when combined with the ILU preconditioner. In the second case the matrix-vector multiply $G\,\mathbf{v}$ required in the linear system solver is approximated by a difference quotient (see (2.32)). For this matrix-free variant a comparably robust, standard preconditioner that does not make explicit use of the matrix G does not exist. Moreover, the matrix-vector multiply can be rather expensive for some PDEs. Therefore we chose for this variant GMRES to have an optimal convergence behavior and so a minimal number of matrix-vector multiplies. As preconditioner we use (block) diagonal scaling (see below in Section 2.6). However, in some cases the number of iterations and thus the number of storage required can be very large. To restrict the amount of storage one can of course restart the GMRES process after a certain number of iterations. However it is more profitable to use a recursive version of GCR in which the inner GMRES loop generates a better preconditioner and the outer GCR loop solves a well-preconditioned system. In `VLUGR2`[3] we used GMRESR[23] and we reported some stagnations. In [12] a remedy is given against these stagnations. The idea is that the new Krylov basis vectors in the innerloop should not only be orthogonal to the previous ones but also to the corrections on the approximate solution already computed in the outerloop. This leads to the GCRO algorithm which is per iteration slightly more expensive than GMRESR (extra orthogonalizations) but which in general uses less iterations and so less matrix-vector multiplies than GMRESR. A second reason not to use BiCGStab in the matrix-free variant was that in our experience BiCGStab was much more sensitive for the inaccuracies induced by

the approximation of $G\,\mathbf{v}$ by a difference quotient then GMRES-like Krylov solvers, possibly because GMRES orthogonalizes the Krylov base after the matrix-vector multiply.

It is not necessary to solve the linear system (2.15) up to machine precision (cf. [6, 19]). The stopping criterion we use in our code is

$$\|P^{-1}r^{(l)}\|_w < \texttt{TOLLSS}, \quad \texttt{TOLLSS} = \texttt{TOLNEW}/(10 \cdot 2^{\texttt{NIT}}), \tag{2.23}$$

where $\texttt{NIT}$ is the current Newton-iteration index and P^{-1} is the preconditioner in use. In the case of GCRO we take the maximum of $\|D^{-1}r^{(l)}\|_w$ and $\|P^{-1}r^{(l)}\|_w$, where D is the (block) diagonal preconditioner and P is obtained by the inner GMRES iterations. The inner GMRES iterations are stopped if the 2-norm of the (weighted diagonally scaled) residual is less than $\texttt{TOLLSS}$ and if a relative improvement of the residual of 0.01 was reached. Note that if P is a good approximation of G, then $P^{-1}r^{(l)}$ is a good approximation of the contribution to the Newton correction $\mathbf{c}^{(k)}$ of the correction $x^{(l)} - x^{(l-1)}$ on the linear solution.

2.5 The matrix-vector product $G\,\mathbf{v}$
Approximation of the Jacobian
If we use modified Newton + BiCGStab to solve the nonlinear system the Jacobian is computed by numerical differencing. The matrix G, resulting from a second-order central discretization on the internal domain and second-order one-sided differences on the boundaries consists of 19 diagonal blocks per rowblock. The storage and the matrix-vector multiply are implemented analogously to the 2D case (cf. [3]). So G is stored in block 19-diagonal storage mode, i.e., in an array $\texttt{G(NPTS,NPDE,NPDE,-9:9)}$. The placing of the 3 grid points coupled in x-direction is known. The actual placing of the grid points corresponding with the 8 lower diagonals is given in the array $\texttt{LLDG(NPTS,-9:-2)}$ and for the 8 upper diagonals in $\texttt{LUDG(NPTS,2:9)}$. If the matrix is stored in this way the matrix-vector product can be vectorized very efficiently.

The Jacobian

$$G = \frac{\partial\mathbf{F}(t,x,y,z,\mathbf{U},\mathbf{U}_t,\mathbf{U}_x,\mathbf{U}_y,\mathbf{U}_z,\mathbf{U}_{xx},\mathbf{U}_{yy},\mathbf{U}_{zz},\mathbf{U}_{xy},\mathbf{U}_{xz},\mathbf{U}_{yz})}{\partial\mathbf{U}} \tag{2.24}$$

is approximated by numerical differencing. To save residual evaluations we have made use of the identity

$$
\begin{aligned}
\frac{\partial\mathbf{F}}{\partial\mathbf{U}} =\ & \frac{\partial\mathbf{F}(.,\mathbf{U},\mathbf{U}_t,.)}{\partial\mathbf{U}} + \\
& \frac{\partial\mathbf{F}(.,\mathbf{U}_x,.)}{\partial\mathbf{U}_x}\cdot\frac{\partial\mathbf{U}_x}{\partial\mathbf{U}} + \frac{\partial\mathbf{F}(.,\mathbf{U}_y,.)}{\partial\mathbf{U}_y}\cdot\frac{\partial\mathbf{U}_y}{\partial\mathbf{U}} + \frac{\partial\mathbf{F}(.,\mathbf{U}_z,.)}{\partial\mathbf{U}_z}\cdot\frac{\partial\mathbf{U}_z}{\partial\mathbf{U}} + \\
& \frac{\partial\mathbf{F}(.,\mathbf{U}_{xx},.)}{\partial\mathbf{U}_{xx}}\cdot\frac{\partial\mathbf{U}_{xx}}{\partial\mathbf{U}} + \frac{\partial\mathbf{F}(.,\mathbf{U}_{yy},.)}{\partial\mathbf{U}_{yy}}\cdot\frac{\partial\mathbf{U}_{yy}}{\partial\mathbf{U}} + \frac{\partial\mathbf{F}(.,\mathbf{U}_{zz},.)}{\partial\mathbf{U}_{zz}}\cdot\frac{\partial\mathbf{U}_{zz}}{\partial\mathbf{U}} + \\
& \frac{\partial\mathbf{F}(.,\mathbf{U}_{xy},.)}{\partial\mathbf{U}_{xy}}\cdot\frac{\partial\mathbf{U}_{xy}}{\partial\mathbf{U}} + \frac{\partial\mathbf{F}(.,\mathbf{U}_{xz},.)}{\partial\mathbf{U}_{xz}}\cdot\frac{\partial\mathbf{U}_{xz}}{\partial\mathbf{U}} + \frac{\partial\mathbf{F}(.,\mathbf{U}_{yz},.)}{\partial\mathbf{U}_{yz}}\cdot\frac{\partial\mathbf{U}_{yz}}{\partial\mathbf{U}}.
\end{aligned}
\tag{2.25}
$$

The implementation of (2.25) is very simple and vectorizes well. If we approximate the partial derivatives of $\mathbf{F}$ in the right-hand side of (2.25) by numerical differencing, the perturbation is only local to a grid point and therefore the Jacobian can be obtained by 10 residual evaluations. A further advantage of this implementation is that different space or time discretizations only lead to different multiplying factors in (2.25). The way to compute the partial derivative of $\mathbf{F}$ with respect to $\mathbf{U}$ is copied from DASSL[7, page 124]

$$\frac{\partial\mathbf{F}(.,\mathbf{U},\mathbf{U}_t,.)}{\partial\mathbf{U}} \approx \frac{\mathbf{F}(,.\mathbf{U}+\boldsymbol{\Delta},\mathbf{U}_t+a_0\boldsymbol{\Delta},.) - \mathbf{F}(.,\mathbf{U},\mathbf{U}_t,.)}{\boldsymbol{\Delta}}, \tag{2.26}$$

where

$$\Delta_{ipt,ic} = \sqrt{uround} \cdot \text{sign}(\Delta t\, U_{tipt,ic}) \cdot \max(|U_{ipt,ic}|, |\Delta t\, U_{tipt,ic}|, \text{ATOL}(ic)). \tag{2.27}$$

For the other partial derivatives we use

$$\frac{\partial \mathbf{F}(.,\mathbf{U}_p,.)}{\partial \mathbf{U}_p} \approx \frac{\mathbf{F}(.,\mathbf{U}_p + \boldsymbol{\Delta}_p,.) - \mathbf{F}(.,\mathbf{U}_p,.)}{\boldsymbol{\Delta}_p} \quad \text{for } p = x,y,z,xx,yy,zz,xy,xz,yz. \tag{2.28}$$

The main difficulty in the numerical computation of the partial derivatives is the choice of the perturbation vector $\boldsymbol{\Delta}_p$, and especially to decide when the value to be perturbed should be considered zero. We chose the perturbation

$$\Delta_{p_{ipt,ic}} = \sqrt{uround} \cdot \max(|U_{p_{ipt,ic}}|, \text{ATOL}(ic) \cdot fac_p), \tag{2.29}$$

where

$$fac_x = \frac{1}{2\Delta x}, \; fac_{xx} = \frac{1}{\Delta x^2}, \; fac_{xy} = \frac{1}{4\Delta x \Delta y}, \quad \text{and other combinations analogously.}$$

So $\text{ATOL}/2\Delta x$ is considered to be the noise value for U_x if ATOL is the noise value for U. Finally, the 'magic tricks'

$$\Delta_{p_{ipt,ic}} = \text{sign}(U_{p_{ipt,ic}}) \cdot \Delta_{p_{ipt,ic}} \tag{2.30}$$

and

$$\boldsymbol{\Delta} = (\mathbf{U} + \boldsymbol{\Delta}) - \mathbf{U} \tag{2.31}$$

have been applied to ensure that the perturbed value has the same sign as the original one and that the perturbed value is a true machine number. Note that in (2.27) the former is not necessarily true. If the PDE is, e.g., undefined for negative values of U this can occasionally be a source of difficulties.

This way of computing a Jacobian was sufficiently accurate for the problems we solved with the code. However, if for a specific problem Newton failures would occur much more often than time step failures, it could be worthwhile to store the exact partial derivatives instead of the approximated ones.

Approximation of G $\mathbf{v}$

The matrix-vector multiply G $\mathbf{v}$ required in the linear system solvers can be approximated by a difference quotient

$$G(\mathbf{U})\, \mathbf{v} = \frac{\mathbf{F}(\mathbf{U} + \sigma\mathbf{v}) - \mathbf{F}(\mathbf{U})}{\sigma}. \tag{2.32}$$

The problem here of course is how to choose σ so that the vector $\mathbf{U}$ is appropriately perturbed. One approach is a generalization of the way the matrix elements in the Jacobian are approximated (cf. [9]). This results in

$$\sigma = \frac{\sqrt{uround}}{<\mathbf{v},\mathbf{v}>} \cdot \text{sign}(\Delta t\, \dot{\mathbf{U}}^T \mathbf{v}) \cdot \max(|\mathbf{U}^T \mathbf{v}|, |\Delta t\, \dot{\mathbf{U}}^T \mathbf{v}|, \overline{\text{ATOL}}\, |\mathbf{v}|), \tag{2.33}$$

where $\overline{\text{ATOL}} = (\text{ATOL}(1),\cdots,\text{ATOL}(\text{NPDE}), \; \cdots \; ,\text{ATOL}(1),\cdots,\text{ATOL}(\text{NPDE}))^T$ and $|\mathbf{v}| = (|v_1|,\cdots,|v_N|)$. However this is not very robust if the components of the solution differ in orders of magnitude. Therefore we adopted the approach proposed in [8], i.e., we solve

$$(WD^{-1}GW^{-1}) \cdot (W\mathbf{c}) = -(WD^{-1}F). \tag{2.34}$$

This has the advantage that the matrix-vector multiplies needed in GMRES are of the form A $\mathbf{x}$ with $A = WD^{-1}GW^{-1}$ and $\|\mathbf{x}\|_2 = 1$. So in the $G(\mathbf{U})$ $\mathbf{v}$ approximation, $\|\mathbf{v}\|_w = \|W\mathbf{v}\|_2 = 1$. Since the Newton correction is assumed small with respect to the solution and the required tolerance if $\|\mathbf{c}\|_w = 1$, this approach means that also the vector $\mathbf{v}$ is assumed small with respect to the solution $\mathbf{U}$. So we can take $\sigma = 1.0$ in the difference quotient (2.32).

In the outer GCR loop we have to scale the multiplying vector with its two-norm before the matrix-vector multiply is executed. However, the two-norm of this vector is needed for the error estimation, so the overhead is low.

2.6 Preconditioning

In our first linear solver BiCGStab is combined with a standard ILU type preconditioner, i.e., only those entries of the decomposition are computed that correspond with nonzero entries in the Jacobian. For the boundaries we use in the ILU preconditioner first-order discretization to maintain the 19-block diagonal structure. The ILU preconditioning is vectorized using a variant of the hyperplane method[1]. For the two-dimensional case this is described in [4]. The extension to 3D is straightforward. This linear solver is in our experience robust and rather efficient with respect to the number of iterations needed. Moreover, in contrast to the 2D case, the vector performance of the ILU preconditioning is satisfactorily in 3D. A disadvantage of our first linear solver is the huge amount of memory needed, especially if the number of PDE components is large.

As preconditioner in our second linear solver we offer the choice between diagonal scaling and block-diagonal scaling. The latter should be used in case the coupling between the PDE components is such that the main diagonal is not a good enough approximation for the Jacobian. The (block) diagonal preconditioner is computed analogously as the main (block) diagonal entries of the Jacobian. Since the mixed derivatives have no influence on the main diagonal, 7 (or $7 \cdot$ NPDE) residual evaluations are required to compute the preconditioner. Optionally, the contributions of the first order derivatives at the boundaries can be neglected. We then can compute the (block) diagonal of the Jacobian with only 4 (or $4 \cdot$ NPDE) residual evaluations.

3. NUMERICAL EXPERIMENTS

In this section we discuss the test results obtained with VLUGR3 for three example problems, viz., (I) Burgers' equation, a model for nonlinear convection-diffusion phenomena, (II) a system of coupled advection-reaction equations, and (III) a transport problem in ground-water flow. The main issue is to evaluate the robustness of the LUGR method. To that aim we consider for both nonlinear equation solvers the accuracy of the solution and the integration history for the example problems. The results are presented in tables and for the problem without reference solution (III) also in plots. Typical grid refinement behavior is also displayed in plots for all problems. In the description of the experiments the following notation is used:

STEPS: Number of successful time steps
JACS: Number of Jacobian and/or preconditioner evaluations
NIT: Number of Newton iterations
LSIT i: Number of linear-solver iterations in the i-th Newton iteration.

For all problems we set the space and time tolerance TOLS = TOLT = 0.1. The weight factors SPCWGT and TOLWGT are set to 1 for all components. For Problem I and II UMAX is also set to 1 for all components. For our first linear solver BiCGStab+ILU the maximum number of linear iterations was set at 100, which was never reached. In our second solver, GCRO+(block) diagonal preconditioning we imposed a maximum of 20 inner GMRES iterations and 5 outer GCR iterations. We allowed also one restart. The maximum for the number of inner iterations was often reached but the linear solver as a whole only occasionally failed using diagonal scaling and never using the block-diagonal preconditioner.

To assess the local-refinement efficiency, i.e., to compare the LUGR code with a standard MoL approach on a (non)uniform fixed grid one should take into account two points, viz., use of memory and computational efficiency. When solving the linear systems with BiCGStab and ILU preconditioning, the amount of memory is dominated by the two arrays needed to store the Jacobian and the preconditioner. Both take up 19.NPDE.NPTS.NPDE floating-point words, where NPTS is the actual number of points in the grid. The matrix-free variant requires a number of vectors of length NPTS.NPDE. Since the memory requisite is in both cases determined by the number of grid points, the LUGR method will be readily more efficient in its memory use. It is more difficult to compare the efficiency with respect to the computational effort. Assuming that the overall accuracy is the same, a first approximation would be to compare the number of grid points at the finest uniform grid with the total number

of grid points used over all grid levels. However one should keep in mind two things. First, it is in general cheaper to solve two systems of N equations than one of order $2N$. Second, the LUGR method requires extra computations for the local refinement: the determination of a new grid and the transfer of solution values between two grid levels. These overhead costs are especially significant if the number of PDE components and the number of (non)linear iterations is small. The LUGR method could be implemented computationally more efficient by adding pointers between the grid structures at different levels but that would be less efficient in memory. Another option is to 'freeze' a specific series of nested grids for a few time steps. A disadvantage of this approach is that it obscures the conceptual simplicity and probably decreases the robustness of the method.

In this section we discuss the robustness and the local-refinement efficiency of the LUGR method. The vector performance and CPU time for all three example problems will be discussed in the next section. Our experiments were done on a Cray Y-MP with a memory limit of 52 Mword.

3.1 Problem I: Burgers' equation

This model, a simple 3D analogue of the Navier-Stokes equations, has a nonlinear convection term combined with a small diffusion term

$$
\begin{aligned}
u_t + u u_x + v u_y + w u_z &= \varepsilon \Delta u, \\
v_t + u v_x + v v_y + w v_z &= \varepsilon \Delta v, \\
w_t + u w_x + v w_y + w w_z &= \varepsilon \Delta w.
\end{aligned} \tag{3.1}
$$

We determined an exact solution for (3.1) representing a wave front moving skew through a unit cube. Using the Cole-Hopf transformation (see, e.g., [24, pg. 97])

$$
u = -2\varepsilon \frac{\varphi_x}{\varphi}, \quad v = -2\varepsilon \frac{\varphi_y}{\varphi}, \text{ and, } w = -2\varepsilon \frac{\varphi_z}{\varphi} \tag{3.2}
$$

(3.1) is transformed into

$$
\varphi_t = \varepsilon \Delta \varphi. \tag{3.3}
$$

An exact solution representing a single shock is given by (cf. [24, pg. 110])

$$
\varphi = f_1 + f_2, \quad f_j = \exp\left(-\frac{c_{j1} x + c_{j2} y + c_{j3} z}{2\varepsilon} + \frac{(c_{j1}^2 + c_{j2}^2 + c_{j3}^2)t}{4\varepsilon} \right). \tag{3.4}
$$

The center of the shock is where $f_1 = f_2$, and we choose this center in the most *unfavorable* way for our grid refinement strategy, i.e., a plane skew through the cube

$$
x - y - z = -\frac{3}{4}t. \tag{3.5}
$$

This is obtained for $c_{jj} = 1$ and $c_{ij} = 1/2$ giving as exact solution for (3.1)

$$
u = \frac{f_1 + \frac{1}{2} f_2}{f_1 + f_2} = 1 - \frac{1}{2} \frac{1}{1 + \exp[(-x + y + z - \frac{3}{4}t)/(4\varepsilon)]}, \tag{3.6a}
$$

$$
v = w = \frac{3}{2} - u. \tag{3.6b}
$$

Substituting v and w from (3.6b) into (3.1) we get the scalar PDE

$$
u_t + u u_x + (\frac{3}{2} - u)(u_y + u_z) = \varepsilon \Delta u. \tag{3.7}
$$

The domain is the unit cube. The initial solution and the Dirichlet boundary conditions are given by the exact solution (3.6). We discuss the results for both the scalar PDE (3.7) and the coupled system (3.1) on the time interval [0,1].

	BiCGStab+ILU				GCRO+Diag			
	Level				Level			
	1	2	3	4	1	2	3	4
STEPS	92	92	92	92	92	92	92	92
JACS	92	92	92	92	92	92	92	92
NIT	184	184	184	184	184	184	184	184
LSIT 1	92	92	92	92	730	908	1206	1760
LSIT 2	21	84	90	91	453	885	1225	1738
# pts								
$t = 0.001$	1331	5481	22041	100637	1331	5481	22041	100637
$t = 0.6$	1331	7233	33677	164617	1331	7233	33677	164617
$t = 1.0$	1331	6789	31229	145065	1331	6789	31229	145129
Uniform	1331	9261	68921	531441	1331	9261	68921	531441

Table 1: Integration history for Problem Ia.

Figure 1: Grid for Problem Ia at $t = 1.0$. Slices at $x = 0.5$, $y = 0.5$, and $z = 0.5$.

3.1.1 Numerical results for Problem Ia   Our first test problem is the scalar Burgers' equation for a Reynolds number of 500 ($\varepsilon = 0.002$). Starting on a base grid with $\Delta x = \Delta y = \Delta z = \frac{1}{10}$ we allow a maximum of 4 grid levels resulting in a finest cell width of $\frac{1}{80}$. After 92 time steps the endpoint was reached, using at each grid level 184 Newton iterations. Since Dirichlet boundary conditions are prescribed we can compute the diagonal-scaling preconditioner in our second linear solver with only 4 residual evaluations. The accuracy and grid refinements were alike for both linear solvers. The maximum norm of the error at $t = 1.0$ was 0.07. In Fig. 1 we picture the grid at $t = 1.0$. We show plots of orthogonal slices through the middle of the cube. As expected these slices are alike since the problem is symmetric in each coordinate direction. In Table 1 we list the number of grid points used at each grid level for 3 different times. The last row contains the number of points needed to reach the same level of refinement with a uniform grid. For the lower grid levels there is not much gain in using the adaptive-grid method. This is because the LUGR method creates a certain buffer in all directions around the wave, resulting in refinement almost everywhere for the lower grid levels. However, at the finest grid level the difference in number of grid points is considerable. Predictably, the overhead of the LUGR method is for this problem rather large since it is a scalar PDE and a very small number of iterations is needed to solve the nonlinear system. Even in scalar mode for the BiCGStab-solver about 25% of the CPU time is spent in routines dealing with the refinement. For the GCRO-solver this is less than 10%, but this solver is less efficient for this problem, especially in scalar mode. As can be seen from the results in Table 1 the ILU preconditioner is optimal for this problem.

	Problem Ib				Problem Ic					
	GCRO+Diag				BiCGStab+ILU			GCRO+Diag		
	Level				Level			Level		
	1	2	3	4	1	2	3	1	2	3
STEPS	78	78	78	78	36	36	36	36	36	36
JACS	78	78	78	78	36	36	36	36	36	36
NIT	156	156	157	217	72	72	88	72	72	78
LSIT 1	659	842	1147	1790	36	36	36	377	533	894
LSIT 2	403	826	1148	1733	31	33	34	361	527	851
LSIT 3				1125			16			84
# pts										
$t = 0.001$	1331	5481	22041	100637	1331	5481	25277	1331	5481	25277
$t = 0.6$	1331	7233	35317	164593	1331	7853	37933	1331	7853	37933
$t = 1.0$	1331	6789	31229	152665	1331	7393	35141	1331	7393	35197
Uniform	1331	9261	68921	531441	1331	9261	68921	1331	9261	68921

Table 2: Integration history for Problem Ib and Ic.

3.1.2 Numerical results for Problem Ib   The second test problem is the coupled system (3.1). To solve a system of 3 PDEs using 4 refinement levels with a maximum number of points of ca. 170000 the BiCGStab solver requires more memory than available on the Cray. The transition layer of the solution is $O(\sqrt{\varepsilon})$ so a cell width of 1/80 is already rather coarse for a Reynolds number of 500 using central finite differencing, especially since the transition layer is placed skew through the cube. Since three refinement levels is too few to solve (3.1) with $\varepsilon = 0.002$, we only used the matrix-free GCRO solver. The results are in agreement with the scalar example. After 78 time steps the endpoint was reached and the maximum norm of the error at $t = 1.0$ was again 0.07. The difference in the number of time steps is caused by the fact that if the estimated time errors are slightly different this can result in a different time step size for the rest of the interval since we adjust our step size such that the rest of the integration interval is an integer number times Δt. The refinements are equivalent with the picture shown in Fig. 1. Also the number of grid points is alike (cf. Table 2). It should be mentioned that for the finest grid the number of grid points changes a great deal from one time level to the next. E.g., for the scalar Burgers' equation this number was 153133 at $t = 0.99$. It is obvious from the integration history in Table 2 that the nonlinear system is harder to solve, especially for the finer grid levels.

3.1.3 Numerical results for Problem Ic   To compare the two nonlinear solvers we decreased the Reynolds number to 200 and solved (3.1) with a maximum of 3 grid levels resulting in a finest cell width of $\frac{1}{40}$. After 36 time steps the endpoint was reached, using about two Newton iterations per time step. The maximum norm of the error at $t = 1.0$ was 0.06. Since the wave front is less sharp than in the previous example more grid points are required on the finer grid levels (cf. Table 2).

Predictably, the overhead is much smaller than for Problem Ia, ca. 1% in scalar mode and ca. 15% in vector mode for the BiCGStab-solver. The linear solver BiCGStab+ILU is again optimal with respect to the number of linear iterations needed, but not in computational time. The computation of the $19 \cdot 3$ block-diagonal Jacobian and its decomposition takes more time than the extra linear iterations needed in the second solver using the simplest possible preconditioner, viz., diagonal-scaling that requires only 4 residual evaluations (see also Table 7 in Section 4.1). One can see from the results in Table 2 that this is still an acceptable preconditioner although more linear iterations per Newton iteration are required than in the scalar Problem Ia.

3.2 Problem II: Rotating sphere problem

Our second example is the advection-reaction system

$$\mathbf{c}_t + u\mathbf{c}_x + v\mathbf{c}_y + w\mathbf{c}_z = \mathbf{f}(\mathbf{c}), \tag{3.8}$$

where u, v, and w are components of a velocity field in the x-, y-, and z-direction, respectively and the reaction term $\mathbf{f}$ consists of the components

$$
\begin{aligned}
f_1 &= -k_2 c_1 c_2 + k_1 c_2^2 \\
f_2 &= -k_1 c_2^2 + k_2 c_1 c_2.
\end{aligned}
\tag{3.9}
$$

If we set the, constant, sum of c_1 and c_2 to d, the solution of (3.8) with zero velocity is given by

$$
\begin{aligned}
c_1(t) &= \frac{d}{k_1+k_2}\frac{k_1(1-\alpha)+(k_1+k_2)\alpha e^{-dk_2 t}}{1-\alpha+\alpha e^{-dk_2 t}} \\
c_2(t) &= d - c_1(t)
\end{aligned}
\tag{3.10}
$$

with

$$\alpha = \frac{(k_1+k_2)c_1(0)-dk_1}{dk_2}. \tag{3.11}$$

For the parameters we selected $k_1 = 1000$, $k_2 = 1$. As initial values we chose $c_1(0) = 0$ and $c_2(0) = d$, so that $\alpha = -k_1/k_2$.

For the advection part we constructed an exact solution which represents a rotating sphere with the highest solution value in its center, along an ellipse on a skew plane $y = z$ through the unit cube. This rotation is obtained by choosing the velocities

$$
\begin{aligned}
u &= 2\pi\sqrt{2}\left(\frac{y+z}{2}-\frac{1}{2}\right) \\
v &= -\pi\sqrt{2}\left(x-\frac{1}{2}\right) \\
w &= v
\end{aligned}
\tag{3.12}
$$

The exact sphere solution for (3.8) with $\mathbf{f} = 0$ is given by

$$\exp(-80[(x-r(t))^2 + (y-s(t))^2 + (z-s(t))^2]) \tag{3.13}$$

with

$$r(t) = (2+\sin(2\pi t))/4, \quad s(t) = \left(2+\frac{1}{2}\sqrt{2}\cos(2\pi t)\right)/4. \tag{3.14}$$

Summarizing, our second test problem is given by (3.8) with the velocity field from (3.12) and the reaction term $\mathbf{f}$ from (3.9). At the inflow boundaries we impose the exact solution. The solution is given by (3.10) with $d = d(x,y,z,t)$ defined by (3.13) and $c_1(0) = 0$, resulting in

$$
\begin{aligned}
c_1(t) &= d(x,y,z,t)\frac{1-e^{-d(x,y,z,t)\cdot t}}{\frac{1001}{1000}-e^{-d(x,y,z,t)\cdot t}} \\
c_2(t) &= d(x,y,z,t)\frac{10^{-3}}{\frac{1001}{1000}-e^{-d(x,y,z,t)\cdot t}}
\end{aligned}
\tag{3.15}
$$

In the short time interval $[0, 0.01]$ the solution c_2 decreases over 90%. For the rest of the time interval, (3.8) is rather an advection problem with the extra difficulty that for negative solution values (3.8) becomes unstable.

	BiCGStab+ILU				GCRO+Block-Diag				GCRO+Diag			
	Level				Level				Level			
	1	2	3	4	1	2	3	4	1	2	3	4
STEPS	335	335	335	335	335	335	335	335	343	343	343	343
JACS	335	335	335	335	335	335	335	335	348	348	347	345
NIT	670	670	670	671	670	670	670	671	695	693	691	688
LSIT 1	335	335	335	335	2575	2807	3397	4148	3312	4029	6612	6231
LSIT 2	16	26	37	45	1322	1404	1693	2431	1577	2130	3647	3274
Av #pts	1331	4500	13700	35800	1331	4500	13700	35800	1331	4500	13700	35800
Uniform	1331	9261	68921	531441	1331	9261	68921	531441	1331	9261	68921	531441

Table 3: Integration history for Problem II.

Figure 2: Grid for Problem II at $t = 1.0$. Slices at $x = 0.5$, $y = 0.5$, and $z = 0.5$.

3.2.1 Numerical results for Problem II   For this test example we used a base grid with 11 points in each direction. A maximum of 4 grid levels was allowed. The initial time stepsize was 0.00001.

In the first run we encountered instabilities due to negative solution values causing a break-down of the run. We then adjusted the problem so that after every time step negative solution values were replaced by zero. With this simple adjustment, resulting in the same error values at $t = 0.5$, the endpoint was reached. The error at $t = 1.0$ was the same as obtained for the advection equation ($\mathbf{f} = 0$), viz., 0.17 for the first component (the maximum value was 0.83 instead of 1.0). The only difference in integration history between the two runs with and without reaction term was the size of the time steps used in the very first part of the time interval $[0, 0.03]$. The advection run used over the whole interval stepsizes of 0.003, whereas the run solving (3.8) used stepsizes from 0.00001 increasing to 0.003 within $[0,0.03]$. The integration performance over the rest of the interval was the same. In this case the efficiency gain in memory of the LUGR method is much larger than in the previous case. This is as expected since the refinement area is now a sphere instead of a skew plane (cf. Table 3). The plot in Fig. 2 of the grid at $t = 1.0$ shows that the sphere is approximated well by the grid refinements.

As can be seen in Table 3 the ILU preconditioner again is optimal in terms of linear iterations required. For the second linear solver, GCRO, we first used the simple diagonal scaling neglecting the first order derivatives at the outflow boundaries. For this solver the problem was a bit more difficult. Newton failures occurred 5 times, resulting in a first step size of 0.4E-7, but after the initial problems there were no more Newton failures and the time step size again increased to 0.003 within $[0,0.03]$. We then run the problem with the block-diagonal preconditioner also neglecting the first-order derivatives. Although both the computation of this preconditioner and the backsolves are more expensive then

the diagonal scaling it pays off. The integration history is as smooth as for BiCGStab+ILU and much less linear iterations are needed. For more detailed information we refer to Table 3.

3.3 Problem III: a 3D fluid-flow / salt-transport problem

We consider a model for an isothermal, single-phase, two-component saturated flow problem which consists of 2 PDEs basic to ground-water flow: the continuity equation and the transport equation. For the background of these equations we refer to [19]. We here present the model in non-conservative form. As independent variables we have the pressure p and the salt mass fraction ω. The continuity equation for the fluid and the salt transport equation are given by

$$n\rho(\beta\frac{\partial p}{\partial t} + \gamma\frac{\partial \omega}{\partial t}) + \nabla \cdot (\rho\mathbf{q}) = 0, \tag{3.16a}$$

$$n\rho\frac{\partial \omega}{\partial t} + \rho\mathbf{q}\cdot\nabla\omega + \nabla \cdot (\rho\mathbf{J}^\omega) = 0, \tag{3.16b}$$

where n is the porosity parameter of the porous medium, β a compressibility coefficient, and γ a salt coefficient. Darcy's law gives the equation for the fluid velocity $\mathbf{q} = (q_1, q_2, q_3)^T$

$$\mathbf{q} = -\frac{k}{\mu}(\nabla p - \rho\mathbf{g}), \tag{3.17}$$

with $\mathbf{g}$ the acceleration-of-gravity vector and k the permeability coefficient of the porous medium. The density ρ and the viscosity μ obey the state equations

$$\rho = \rho_0 \exp[\beta p + \gamma\omega], \tag{3.18}$$

$$\mu = \mu_0 \cdot m(\omega), \; m(\omega) = 1 + 1.85\omega - 4.1\omega^2 + 44.5\omega^3, \tag{3.19}$$

where ρ_0 is the reference density of fresh water and μ_0 a reference viscosity. The equation for the salt-dispersion flux vector is given by Fick's law

$$\mathbf{J}^\omega = -n\mathbf{D}\nabla\omega, \tag{3.20}$$

with the dispersion tensor $\mathbf{D}$ for the solute salt defined as

$$n\mathbf{D} = (nD_{mol} + \alpha_T|\mathbf{q}|)I + \frac{\alpha_L - \alpha_T}{|\mathbf{q}|}\mathbf{q}\mathbf{q}^T, \; |\mathbf{q}| = \sqrt{\mathbf{q}^T\mathbf{q}}. \tag{3.21}$$

The coefficients D_{mol}, α_T and α_L correspond, respectively, with the molecular diffusion and the transversal and longitudinal dispersion.

Our examples are connected with laboratory experiments that deal with the displacement of fresh water by a polluting fluid in a tank filled with a porous medium. Fresh water is flowing from left to right through the tank. The polluting fluid, that has a higher density than the fresh water, is injected with constant velocity through a slit at the top of the tank. We simulate here a pollution with salt water of two different concentrations, one with a salt mass fraction of 0.0935, and the other, more demanding, a pollution with brine having a salt mass fraction of 0.25. In both cases this gives rise to a fresh-salt water plume, but the steepness of the front is dependent of the salt mass fraction.

The tank is defined by a box $\Omega = \{(x, y, z)|0 \le x \le 2.5, 0 \le y \le 0.5, 0 \le z \le 1.0\}$. The acceleration of gravity vector takes the form $\mathbf{g} = (0, 0, g)^T$, where g is the gravity constant. The initial values at $t = 0$ at $\Omega \cup \partial\Omega$ are taken as

$$p(x, y, z, 0) = (0.03 - 0.012x + 1 - z)\rho_0 g, \quad \omega(x, y, z, 0) = 0. \tag{3.22}$$

For $0 < t \le t_{end}$ the following boundary conditions are imposed

n	0.35	k	$7.18_{10^{-11}}$	nD_{mol}	$1_{10^{-9}}$
ρ_0	1000	μ_0	0.001	α_T	0.001
γ	$\ln(2)$	g	10	α_L	0.01
β	0.0	t_{end}	2_{10^4}		

Table 4: Data for Problem III.

		BiCGStab+ILU				GCRO+Block-diag			
		Level				Level			
		1	2	3	4	1	2	3	4
$t = 7200$	STEPS	101	101	101	101	115	115	115	115
	JACS	103	103	103	105	126	125	125	135
	NIT	203	205	206	211	250	250	253	254
	LSIT 1	323	268	359	441	4012	4944	6634	9313
	LSIT 2	148	109	186	302	3313	2946	3682	6072
	LSIT 3							106	466
	# pts	1716	2211	7817	25647	1716	2211	7817	25647
$t = 21600$	STEPS	168	168	168	168	182	182	182	182
	JACS	170	170	170	172	193	192	193	202
	NIT	337	339	340	345	384	384	399	392
	LSIT 1	577	494	631	759	7013	8442	11806	16016
	LSIT 2	267	235	360	632	5295	4798	6682	9969
	LSIT 3							556	571
	LSIT 4							187	
	# pts	1716	4675	20419	66527	1716	4675	20419	66527
	Uniform	1716	11781	86961	667521	1716	11781	86961	667521

Table 5: Integration history for Problem IIIa.

$$
\begin{array}{llll}
x = 0, & y \in [0, 0.5], z \in [0, 1]: & p = p(x, y, z, 0), & \omega = 0, \\
x = 2.5, & y \in [0, 0.5], z \in [0, 1]: & p = p(x, y, z, 0), & \omega_x = 0, \\
y = 0, 1 & x \in [0, 2.5], z \in [0, 1]: & q_2 = 0, \omega_y = 0, \\
z = 0 & x \in [0, 2.5], y \in [0, 0.5]: & q_3 = 0, \omega_z = 0, & \quad (3.23) \\
z = 1 & x, y \notin B: & q_3 = 0, \omega_z = 0, \\
& B = \{(x, y) \mid 0.375 \le x \le 0.4,\ 0.2 \le y \le 0.3\} \\
(x, y) \in B: & & \rho q_3 = -4.95_{10^{-2}}, \omega = \omega_B.
\end{array}
$$

The last line is connected with the slit where the salt water is injected with a prescribed velocity and concentration. In the first experiment the salt mass fraction $\omega_B = 0.0935$, in the second $\omega_B = 0.25$. The other conditions are self-evident. All remaining problem data are contained in Table 4.

3.3.1 Numerical results for Problem IIIa ($\omega_B = 0.0935$) As cell width of the base grid we used 0.1 in each direction. The maximum number of grid levels allowed was 4. For this problem we set UMAX(1) $= \rho_0 g = 10000$ and UMAX(2) $= 0.1$. In Fig. 3 we give contourplots of the salt concentration in the plane $y = 0.2$ after 2 and 6 hours. The contour lines of the higher salt mass fractions (0.5 - 1.0 times ω_B) display a likely solution pattern and show no wiggles despite the steepness of the solution near the inlet slit. For the lower concentrations, however, there is an unexpected 'drip' in the solution. Since the density decreases from top to bottom, one would expect that the pollutant would be more easier taken by the flow instead of sinking to the bottom. It is not likely that this is caused by the grid refinement procedure, since the refined grids are nicely placed around the fresh-salt water plume. Moreover, a second run on a uniform grid with a cell width of 0.025 gave a solution with the same characteristics.
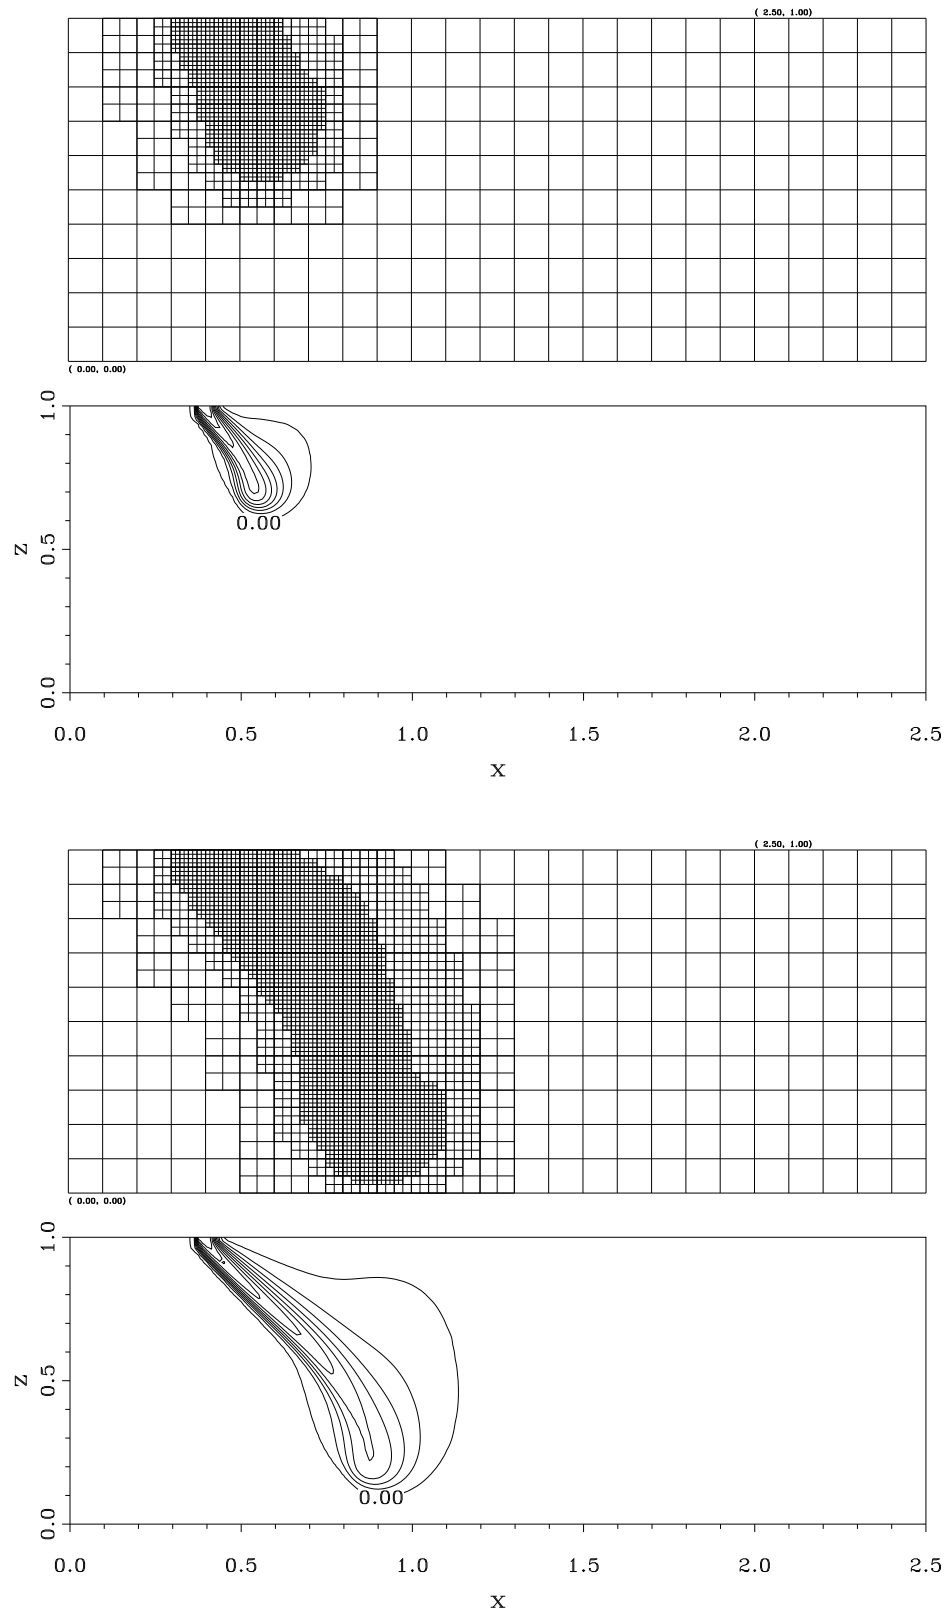
Figure 3: Problem IIIa ($\omega_B = 0.0935$), slice at $y = 0.2$.
Grid and 10% contour lines of the salt mass fraction after 2 hours (above) and after 6 hours (below).

		Level				
		1	2	3	4	5
$t = 7200$	STEPS	178	178	178	178	178
	JACS	185	184	184	184	198
	NIT	369	368	367	379	376
	LSIT 1	457	431	556	698	970
	LSIT 2	172	195	275	440	896
	LSIT 3					92
	# pts	1716	3157	10237	32541	100799
$t = 14400$	STEPS	243	243	243	243	243
	JACS	261	260	266	257	282
	NIT	521	520	526	525	558
	LSIT 1	739	689	837	991	2121
	LSIT 2	286	309	455	674	1432
	LSIT 3					223
	LSIT 4					97
	LSIT 5					63
	# pts	1716	4059	16063	52715	145394
	Uniform	1716	11781	86961	667521	5229441

Table 6: Integration history for Problem IIIb.

The integration data are given in Table 5. Using BiCGStab+ILU as linear solver the integration in time is again performed smoothly, the stepsize steadily increases from ca. 1 at the start of the problem to ca. 200 at t_{end}. Both solvers give the same solution up to plotting accuracy. In this case we could not neglect the boundary conditions in the preconditioner for GCRO. This was to be expected since the boundary conditions at the inlet do include first order derivatives. With block-diagonal preconditioning the second solver behaved reasonably smooth in time. Ten Newton failures occurred at about $t = 3300$, but in the remainder of the time interval the stepsize steadily increased to 260 at t_{end}. Restriction to diagonal scaling led to too many failures of the linear solver to be of practical use. For this problem BiCGStab+ILU is by far the more efficient in CPU time, but if memory is the bottle-neck the problem can be solved with the matrix-free solver. It should be noted that this solver is less robust. Since the quality of the stopping criterion (2.23) is dependent on the correctness of the approximation of the Jacobian G by the preconditioner P it is important to have a good preconditioner. This means for the second linear solver that the number of iterations in the inner GMRES loop should not be too small so that the system is well enough solved.

Note that the solution of the linear systems is more demanding in this case. Therefore the LUGR approach will be, from a computational point of view, more profitable than in the previous examples.

3.3.2 Numerical results for Problem IIIb ($\omega_B = 0.25$)   We started with the same base grid as in the previous case. A fifth grid level was required to solve the brine transport problem since a higher salt mass fraction results in steeper fronts. UMAX was set to (10000,0.25). The contourplots in Fig. 4, again a slice from the solution at $y = 0.2$, show that, as expected, a pollutant with higher density will sink faster to the bottom, which is reached after 4 hours. For this problem we present only the results for the BiCGStab+ILU solver. The figures for the integration history are given in Table 6. Note that, even more striking than in the previous case, the solution of the (non)linear systems at a fine grid appears to be more cumbersome than at the coarser grids. This could probably be explained by near-singularities caused by the hydrodynamic dispersion tensor in the vicinity of vortices [14].

It would not be feasible to solve problem IIIa or IIIb on a uniform fine grid. However, knowing the mass fraction of a pollutant and the pressure gradient between $x = 0$ and $x = 2.5$, one could
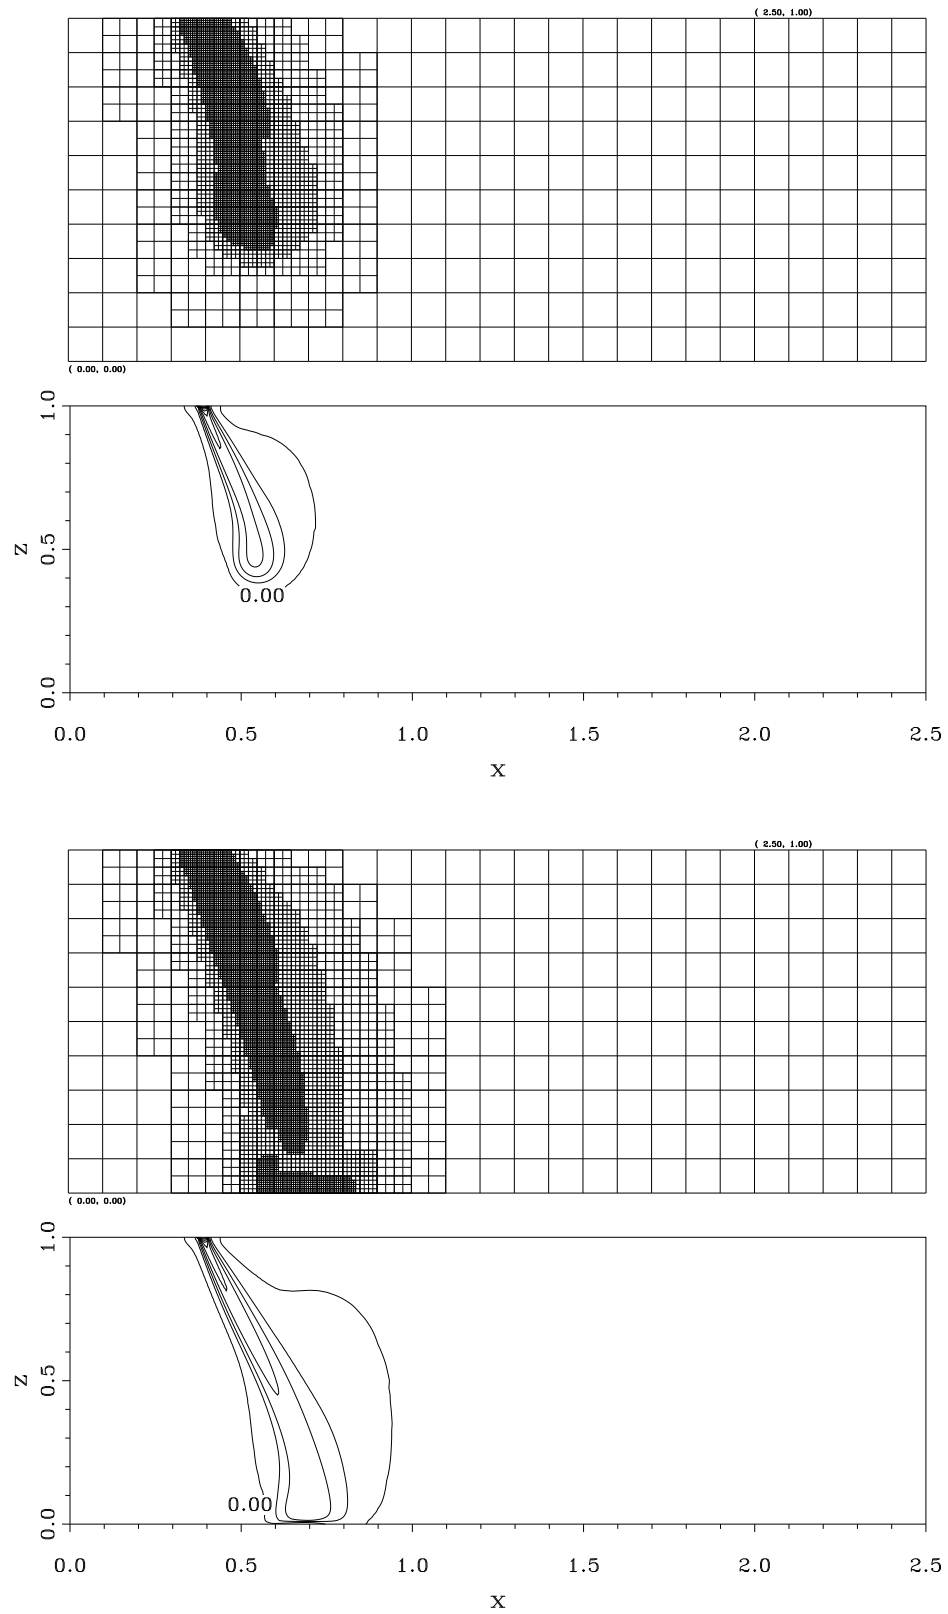
Figure 4: Problem IIIb ($\omega_B = 0.25$), slice at $y = 0.2$.
Grid and 10% contour lines of the salt mass fraction after 2 hours (above) and after 4 hours (below).

predict on forehand a region where the polluting fluid would spread and use a nonuniform grid with refinements in that specific region.

4. PERFORMANCE

Our performance evaluation was done on a Cray YMP with the CF77 compiling system. Scalar results were obtained using cf77 -Wf" -o novector", and vector results with cf77 -Zv -Wf"-o aggress". To measure the Megaflop rate and the CPU time spent in a routine we used the Cray utility Perftrace[11], that gives the hardware performance by program unit (compiler flags -F, and loader flags -F -lperf).

4.1 Global performance

		BiCGStab		GCRO		
		CP sec	Mflop	CP sec	Mflop	Prec
Problem Ia	scalar	1263	12	2446	22	Diag
	vector	430	36	509	104	Diag
Problem Ib	vector			1260	160	Diag
Problem Ic	scalar	1439	10	857	24	Diag
	vector	175	79	132	154	Diag
Problem II	vector	873	63	787	126	Diag
	vector			624	105	Block
Problem IIIa	vector	758	119	2854	195	Block
Problem IIIb	vector	4222	117			

Table 7: Global performance.

We first give a global idea of the performance of the 2 different solvers, i.e., VLUGR3 + BiCGStab, and VLUGR3 + GCRO. In Table 7 the CPU time and the Mflop rate is shown for the example problems.

BiCGStab+ILU is the more robust solver and uses the least number of linear iterations. Therefore, if memory is not a problem it is in most cases more efficient, especially on scalar processors. If the number of PDE components is large, or if (block) diagonal scaling is a good preconditioner then GCRO can be competitive or even faster. This matrix-free solver is much more memory efficient and vectorizes also much better, but since it is more sensitive to stopping criteria and restrictions on the number of iterations we feel that it should be used with care.

4.2 Vector results

In this section we discuss the vector performance of the LUGR code for the three example problems. The timings were done on 1 processor of a Cray Y-MP, which has a clock cycle time of 6ns. This gives a theoretical peak performance on 1 processor of 167 Mflops and 333 when chaining an add and a multiply. Since during one cycle time 1 store and 2 loads can be performed, indirect addressing of one of the vector operands of a triad will reduce the performance at least with a factor of 2, bank conflicts

BiCGStab+ILU					GCRO+Diag				
	# calls	Avg time	ACM %	Mflop		# calls	Avg time	ACM %	Mflop
INJON	816	7.8E-2	14.8	0	DERIVS	9641	1.4E-2	26.4	137
ILU backs	2044	2.2E-2	25.3	69	GMRESO	1393	6.0E-2	42.9	235
JACSLP	277	1.5E-1	35.0	0	INJON	816	7.8E-2	55.4	0
JACSUP	277	1.4E-1	44.3	0	MKBND	276	1.3E-1	62.5	0
MKBND	276	1.3E-1	52.6	0	SETBA	277	1.3E-1	69.5	0

Table 8: Vector performance of top 5 routines for Problem Ia.

GCRO+Diag				
	# calls	Avg time	ACM %	Mflop
DERIVS	10374	4.7E-02	38.5	145
GMRESO	1294	2.6E-01	64.9	251
PDEF	11622	1.0E-02	74.4	215
INJON	690	9.5E-02	79.6	0
MATVEC	9688	5.2E-03	83.6	183

Table 9: Vector performance of top 5 routines for Problem Ib.

BiCGStab+ILU					GCRO+Diag				
	# calls	Avg time	ACM %	Mflop		# calls	Avg time	ACM %	Mflop
ILU dec	108	5.8E-1	36.0	65	DERIVS	3849	1.3E-2	37.7	137
ILU backs	676	6.0E-2	59.4	78	GMRESO	430	8.3E-2	64.8	245
MATVEC	444	3.8E-2	69.2	143	PDEF	4281	2.7E-3	73.7	210
PDEF	3472	2.0E-3	73.2	213	INJON	208	3.3E-2	78.9	0
INJON	208	3.2E-2	77.1	0	MATVEC	3627	1.4E-3	82.7	176

Table 10: Vector performance of top 5 routines for Problem Ic.

BiCGStab+ILU					GCRO+Diag				
	# calls	Avg time	ACM %	Mflop		# calls	Avg time	ACM %	Mflop
ILU dec	1340	1.6E-1	24.6	62	DERIVS	33583	7.0E-3	30.2	137
ILU backs	5611	2.4E-2	40.2	69	GMRESO	4397	3.5E-2	50.0	239
INJON	3003	2.5E-2	48.9	0	INJON	3084	2.6E-2	60.2	0
MATVEC	2930	1.5E-2	54.0	140	PDEF	39133	1.7E-3	68.6	240
JACSLP	1006	4.3E-2	59.0	0	MKBND	1038	3.9E-2	73.7	0
					GCRO+Block-diag				
						# calls	Avg time	ACM %	Mflop
					DERIVS	22460	7.0E-03	25.5	136
					INJON	3003	2.6E-02	37.8	0
					GMRESO	4148	1.7E-02	49.4	226
					PDEF	33180	1.6E-03	58.1	240
					MKBND	1005	3.9E-02	64.4	0

Table 11: Vector performance of top 5 routines for Problem II.

BiCGStab+ILU					GCRO+Block-diag				
	# calls	Avg time	ACM %	Mflop		# calls	Avg time	ACM %	Mflop
PDEF	15006	1.6E-2	31.4	210	PDEF	83865	2.1E-2	62.9	209
ILU backs	9340	2.0E-2	56.5	72	GMRESO	4150	1.0E-1	77.7	254
MATVEC	7972	1.3E-2	70.6	142	DERIVS	72957	5.4E-3	91.7	134
ILU dec	682	1.0E-1	80.1	62	BCKBDI	72957	6.6E-4	93.4	113
INJON	1518	1.6E-2	83.3	0	MATVEC	71364	5.8E-4	94.8	175

Table 12: Vector performance of top 5 routines for Problem IIIa.

BiCGStab+ILU				
	# calls	Avg time	ACM %	Mflop
ILU backs	21071	6.6E-2	33.1	73
PDEF	29277	3.6E-2	57.9	210
MATVEC	18282	4.4E-2	77.2	142
ILU dec	1326	2.3E-1	84.2	66
INJON	3064	3.3E-2	86.7	0

Table 13: Vector performance of top 5 routines for Problem IIIb.

left out of consideration. When more vectors are indirectly addressed the (current) impossibility on the Y-MP to chain more than one gathered/scattered load/store would reduce the performance to a much larger degree.

On all problems the vector performance of the LUGR code using the matrix-free GCRO+(block) diagonal preconditioning is satisfactory or even good. For more difficult problems a vector speed of ca. 200 Mflops is reached. For Problem Ia and to a lesser extent also for Problem II the code with BiCGStab+ILU almost behaves as a scalar code. Closer inspection shows that all routines that deal with the solution of the PDE-system on *one specific grid* have a satisfactory vector performance. The definition of the PDE achieves 150-250 Mflops, the matrix-vector multiply about 150 and the preconditioner circa 70 Mflops. The disappointing overall vector performance for the first two problems is caused by the grid refinement 'overhead' routines that contain no floating-point operations and that take a considerable amount of the total CPU time. In the case of the, scalar, Problem Ia this amounts to even 70% of the total CPU time. For the two-component system of Problem II the overhead still amounts to 35%. In Tables 8-13 we can see that the most time consuming 'search' routine is INJON which transfers the solution values of previous time levels to the grid at a current time level. The time spent in this routine could be diminished by adding pointers between coinciding grid points from one grid to another during creation of the new grid. In the current implementation one has to search for coinciding grid points in both grids. It is also possible to use a one-step time-integration formula instead of the two-step BDF formula. This would approximately halve the time used to transfer solution values of previous time levels to the current grid. The less complicated way of defining the boundary in 3D compared to the structure in 2D in VLUGR2[3] has its consequences when those structures are used (as in JACSLP and JACSUP which create the dependency lists for the vectorizable version of the preconditioner) or created (MKBND). However one should keep in mind that these routines are less dominant when the problem is more nonlinear or when the number of PDE components is larger so that more work has to be done to *solve* the PDE-system on a specific grid (cf. Tables 12 and 13). For Problem IIIa the overhead costs are ca. 10% and for Problem IIIb even less than 5%.

5. SUMMARY

We have discussed the performance of a Method of Lines solver based on a Local Uniform Grid Refinement method for systems of time-dependent PDEs in three dimensions. The, vectorizable, implementation of this method is an extension of the 2D code VLUGR2[3]. The LUGR method proved to be robust and efficient with respect to the location of the refined grids, especially when a very fine grid was required in part of the domain. The experiments with VLUGR2[3] revealed that in two dimensions the overhead for the grid refinement is negligible when using an implicit time integrator, even for simple scalar problems. In three space dimensions the local refinement overhead is larger and can even dominate the CPU time for simple problems. It is possible to decrease this overhead by adding extra pointers to the datastructure or by using a one-step time-integration formula. However, if the number of components is large and/or the nonlinearity of the problem high (as for instance in Problem III) the CPU time needed for the grid structure and refinement is again negligible.

Our LUGR code offers two different iterative nonlinear solvers. In the first the Jacobian of the Newton process is approximated and stored. The linear solver used is BiCGStab[22] combined with standard ILU preconditioning. The second solver is a matrix-free Newton process and uses GCRO[12] and (block) diagonal scaling. If memory demands are no restriction we advocate the former since it is more robust and generally uses less CPU time. The matrix-free solver not only avoids the storage of Jacobian and ILU matrices, it generates also a true Newton process instead of a modified one. An additional advantage is that tailor-made space-discretization schemes, resulting in other couplings than the here used 19-point stencil, can be more easily implemented. The vector performance of this solver is good, from 100 Mflops for a simple scalar PDE to 200 Mflops for a more nonlinear system of PDEs.

The first linear solver is less efficiently vectorizable but uses also in general less CPUtime. For simple scalar problems the CPU time is dominated by the grid refinement overhead which is for this solver even larger than for the matrix-free one. It will of course still be faster then solving on a uniform fine grid. The ILU preconditioning of the linear systems vectorizes better than in the 2-dimensional case, since the number of computations that can be done independently from each other is much larger. The Mflop rate however is still hampered by the fact that only 1 indirect load/store instruction can be issued at a time, versus 2 load and 1 store instructions without gather/scatter necessity (cf. [4]).

REFERENCES
1. C. C. Ashcraft and R.G. Grimes. On vectorizing incomplete factorization and SSOR preconditioners. *SIAM J. Sci. Stat. Comput.*, 9(1):122–151, 1988.

2. J.G. Blom and J.G. Verwer. A vectorizable adaptive grid solver for PDEs in 3D. Report NM-R9319, CWI, Amsterdam, 1993.

3. J.G. Blom and J.G. Verwer. VLUGR2: A vectorized local uniform grid refinement code for PDEs in 2D. Report NM-R9306, CWI, Amsterdam, 1993.

4. J.G. Blom and J.G. Verwer. Vectorizing matrix operations arising from PDE discretization on 9-point stencils. *The Journal of Supercomputing*, 8:29–51, 1994.

5. J.G. Blom and J.G. Verwer. VLUGR3: A vectorizable adaptive grid solver for PDEs in 3D. II. Code description. Report NM-R9405, CWI, Amsterdam, 1994.

6. J.G. Blom, J.G. Verwer, and R.A. Trompert. A comparison between direct and iterative methods to solve the linear systems arising from a time-dependent 2D groundwater flow model. *Int. J. Comp. Fluid Dynamics*, 1:95–113, 1993.

7. K.E. Brenan, S.L. Campbell, and L.R. Petzold. *Numerical Solution of Initial-Value Problems in Differential-Algebraic Equations.* North-Holland, New-York, 1989.

8. P.N. Brown, A.C. Hindmarsh, and L.R. Petzold. Using Krylov methods in the solution of large-scale differential-algebraic systems. Technical Report TR 93-37, Computer Science Department, University of Minnesota, Minneapolis, 1993.

9. P.N. Brown and Y. Saad. Hybrid Krylov methods for nonlinear systems of equations. *SIAM J. Sci. Stat. Comput.*, 11(1):450–481, 1990.

10. A.T. Chronopoulos. Nonlinear CG-like iterative methods. Report UMSI 91/99, University of Minnesota Supercomputer Institute, Minneapolis, 1991.

11. Cray Research, Inc. *UNICOS Performance Utilities Reference Manual*, SR-2040 6.0 edition, 1991.

12. E. de Sturler and D.R. Fokkema. Nested Krylov methods and preserving the orthogonality. Preprint NR. 796, Department of Mathematics, University of Utrecht, the Netherlands, 1993.

13. R.A. Trompert. MOORKOP, an adaptive grid code for initial-boundary value problems in two space dimensions. Report NM-N9201, CWI, Amsterdam, 1992.

14. R.A. Trompert. A note on singularities caused by the hydrodynamic dispersion tensor. Report NM-R9302, CWI, Amsterdam, 1993.

15. R.A. Trompert. *Local Uniform Grid Refinement for Time-Dependent Partial Differential Equations*. PhD thesis, University of Amsterdam, the Netherlands, 1994.

16. R.A. Trompert and J.G. Verwer. A static-regridding method for two-dimensional parabolic partial differential equations. *Appl. Numer. Math.*, 8:65–90, 1991.

17. R.A. Trompert and J.G. Verwer. Analysis of the implicit Euler local uniform grid refinement method. *SIAM J. Sci. Comput.*, 14:259–278, 1993.

18. R.A. Trompert and J.G. Verwer. Runge-Kutta methods and local uniform grid refinement. *Math. Comp.*, 60:591–616, 1993.

19. R.A. Trompert, J.G. Verwer, and J.G. Blom. Computing brine transport in porous media with an adaptive-grid method. *Int. J. Numer. Meth. in Fluids*, 16:43–63, 1993.

20. J.G. Verwer and R.A. Trompert. An adaptive-grid finite-difference method for time-dependent partial differential equations. In D.F. Griffiths and G.A. Watson, editors, *Proc. 14-th Biennial Dundee Conference on Numerical Analysis*, pages 267–284. Pitman Research Notes in Mathematics Series 260, 1992.

21. J.G. Verwer and R.A. Trompert. Analysis of local uniform grid refinement. *Appl. Numer. Math.*, 13:251–270, 1993.

22. H.A. van der Vorst. BI-CGSTAB: A fast and smoothly converging variant of BI-CG for the solution of nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.*, 13(2):631–644, 1992.

23. H.A. van der Vorst and C. Vuik. GMRESR: A family of nested GMRES methods. Report 91-80, Faculty of Technical Mathematics and Informatics, TU Delft, the Netherlands, 1991.

24. G.B. Whitham. *Linear and Nonlinear Waves*. John Wiley & Sons, Inc., New York, London, Sydney, Toronto, 1974.