


**ma
the
ma
tisch**

**cen
trum**



AFDELING NUMERIEKE WISKUNDE
(DEPARTMENT OF NUMERICAL MATHEMATICS)

NW 154/83

JUNI

P.W. HEMKER, P. WESSELING & P.M. DE ZEEUW

A PORTABLE VECTOR-CODE FOR AUTONOMOUS MULTIGRID MODULES

amsterdam

1983

**stichting
mathematisch
centrum**



AFDELING NUMERIEKE WISKUNDE
(DEPARTMENT OF NUMERICAL MATHEMATICS)

NW 154/83

JUNI

P.W. HEMKER, P. WESSELING & P.M. DE ZEEUW

A PORTABLE VECTOR-CODE FOR AUTONOMOUS MULTIGRID MODULES

kruislaan 413 1098 SJ amsterdam

BIBLIOTHEEK MATHEMATISCH CENTRUM
AMSTERDAM

Printed at the Mathematical Centre, Kruislaan 413, Amsterdam, The Netherlands.

The Mathematical Centre, founded 11 February 1946, is a non-profit institution for the promotion of pure and applied mathematics and computer science. It is sponsored by the Netherlands Government through the Netherlands Organization for the Advancement of Pure Research (Z.W.O.).

1980 Mathematics subject classification: 65F10, 65N20

Copyright © 1983, Mathematisch Centrum, Amsterdam

A portable vector-code for autonomous multigrid modules

by

P.W. Hemker, P. Wesseling & P.M. de Zeeuw

ABSTRACT

The implementation is described of two multigrid algorithms for use as standard subroutines for the solution of linear systems that arise from 7-point discretizations of elliptic PDEs on a rectangle. For both algorithms a tuned scalar-version and a tuned vector-version have been constructed and run on a CYBER 170, a CRAY 1 and a CYBER 205. The CPU-times are given and compared. The implementation is available in portable ANSI-FORTRAN.

KEY WORDS & PHRASES: *Multigrid methods, vector computers, elliptic PDES.*

1. INTRODUCTION

In this paper we describe software for the solution of discretized 2^{nd} order linear elliptic PDEs in two space dimensions. The domain of definition is assumed to be a rectangle and the discretization is assumed to result in a regular 7-diagonal matrix. The algorithms, based on multigrid cycling, are selected for efficiency. The aim was to obtain software that is perceived and can be used just like any standard subroutine for solving systems of linear equations. The user has to specify only the matrix and the right-hand-side, and remains unaware of the underlying multigrid method. Such a subroutine, that operates without outside interference, will be called autonomous. We find that a large class of equations can be solved efficiently by use of our autonomous multigrid subroutines. The equation may be non-self-adjoint, and its coefficients are arbitrary.

The two algorithms use saw-tooth multigrid cycles [8, 9]. One algorithm is based on ILU-relaxation, the other on ZEBRA-relaxation [7]. The discretization on coarse grids is provided automatically by means of a built-in Galerkin approximation.

Various scalar- and vector-versions of the code have been constructed and run on a CYBER 170, a CRAY 1 and a CYBER 205. It appeared that a code written for automatic vectorization in portable ANSI FORTRAN runs efficiently in all cases. Specially tuned versions are only a small fraction more efficient.

In section 2 we describe the class of problems that can be solved. In section 3 we describe the general algorithm for multigrid cycling. In the sections 4, 5, 6 we specialize the general algorithm and come to the various versions of the codes. In sections 7, 8 we compare the various programs. In the last section we formulate some conclusions.

2. THE PROBLEM


We consider the linear 2^{nd} order elliptic PDE in two dimensions

$$(2.1.a) \quad \sum_{i,j=1,2} a_{ij} \left(\frac{\partial}{\partial x_i} \right) \left(\frac{\partial}{\partial x_j} \right) u + \sum_{i=1,2} a_i \left(\frac{\partial}{\partial x_i} \right) u + a_0 u = f \text{ on } \Omega \subset \mathbb{R}^2,$$

with variable coefficients and with boundary conditions on $\delta\Omega = \Gamma_N \cup \Gamma_D$

$$(2.1.b) \quad \begin{aligned} \left(\frac{\partial}{\partial n}\right)u + \alpha\left(\frac{\partial}{\partial s}\right)u + \beta u &= \gamma \quad \text{on } \Gamma_N, \\ u &= g \quad \text{on } \Gamma_D. \end{aligned}$$


The coefficients are arbitrary but should satisfy the ellipticity condition. If this equation on a rectangle Ω is discretized on a regular triangularization of the form



then the resulting discretization

$$(2.2) \quad A_h u_h = f_h$$

can be a linear system with a regular 7-diagonal structure. We consider



The shape of
a matrix A_h

codes for the solution of linear systems with a structure corresponding to this kind of 7-point discretization. On the rectangle Ω equidistant computational grids Ω^k , $k = 1, 2, \dots, \ell$, are defined

$$\Omega^k = \{(x_1, x_2) \mid x_i = m_i 2^{1-k}, m_i = O(1)N_i C \cdot 2^{k-1}\}.$$

The user has to provide the discrete operator A_h and the data f_h only on the finest grid Ω^ℓ .

To solve the linear system efficiently, multigrid methods are used [2, 7, 8]. These methods make also use of Ω^k , $k = 1, 2, \dots, \ell-1$, but when using the autonomous subroutines the user remains unaware of this fact. Much effort has been spent in the search for efficient variants of the MG-method [3, 4, 5, 6, 7, 8]. In this paper we consider only two variants that are found to belong to the more promising ones.

3. THE GENERAL MULTIGRID ALGORITHM

The general multigrid cycling algorithm for the solution of (2.2) is an iterative process, which makes use of a sequence of discretizations on the grids Ω^k , $k = 1, 2, \dots, \ell$.

Each iteration cycle consists of

- (1) p relaxation sweeps, followed by
- (2) a coarse grid correction, followed by
- (3) q more relaxation sweeps.

The coarse grid correction consists of

- a) the computation of the current residual $r_h = f_h - A_h \tilde{u}_h$;
- b) the restriction of the residual to the next coarser grid $r_{Hh} = R_{Hh} r_h$;
- c) the computation of \tilde{c}_H , an approximation to the solution of the correction equation $A_H c_H = r_H$. This approximation is obtained by application of s multigrid iteration cycles to this equation, and
- d) updating the current solution \tilde{u}_h by addition of the prolonged correction

$$\tilde{u}_h := \tilde{u}_h + P_{hH} \tilde{c}_H.$$

On the coarsest grid the correction equation (3.1) has to be (approximately) solved by another method (at choice). The coarse grid discrete operators A_H can be constructed by analogy to (2.2) or by Galerkin approximation:

$$(3.2) \quad A_H = R_{Hh} A_h P_{hH}$$

(cf. [8]).

4. THE ALGORITHMS

In this paper we consider the implementation of two particular instances of the multigrid algorithms: MGD1 and MGE0Z. Based on various comparisons [4, 6, 9], the parameters p , q and s are chosen to be 0, 1 and 1 respectively. The resulting strategy is called a saw-tooth cycle [9]. For the prolongation and the restriction 7-point operators are chosen, that correspond to linear interpolation on coarse-grid triangles in the triangulation of Ω (for P_{hH}) and to its adjoint operator (for R_{Hh}). The Galerkin approximation (3.2) is chosen for the construction of the coarse grid operators A_H . For an approximate solution on the coarsest grid a single relaxation sweep is used. The two algorithms differ only with respect to the relaxation method.

In MGD1 the Incomplete LU (ILU-) relaxation is used [9]. For this relaxation the 7-diagonal matrix A_h is decomposed as

$$A_h = LU - C$$

where L is a lower-triangular matrix (with 1 for all main-diagonal elements) and U is an upper triangular matrix. The requirement that L and U have non-zero diagonals only where A_h has, determines L and U . The rest-matrix C has only two non-zero diagonals, of which the elements are easily derived from L and U . One relaxation sweep of the Incomplete LU-relaxation is now the solution of the system

$$LU u^{(i+1)} = f + C u^{(i)}.$$

After such a relaxation sweep a residual is efficiently computed by

$$r^{(i+1)} := f_h - A_h u_h^{(i+1)} = C(u_h^{(i+1)} - u_h^{(i)})$$

In MGE0Z the ZEBRA-relaxation is used [7]. This is a line-Gauss-Seidel

relaxation in which first all points on even lines (lines that appear in the coarser grid) are simultaneously relaxed and secondly all points on the odd lines. An important advantage of this relaxation is the fact that many points can be relaxed simultaneously and that the residual computation simplifies, because the residual vanishes at all odd lines after a relaxation sweep. For ZEBRA relaxation tridiagonal systems have to be solved. The solution of these systems can be accelerated by storage of the decomposition of the tridiagonal matrices. We have chosen to do this at an extra storage cost of 2 reals per grid-point.

5. THE STRUCTURE OF MGD1 AND MGEOZ

The general structure of both MGD1 and MGEOZ is the same. First, in the preparational phase, the sequence of coarse grid discrete operators is constructed by a subroutine RAP, using (3.2). Then the decomposition is performed (in DECOMP or DECOMPZ) and the initial estimate of the solution is set to zero. Finally, in the cycling phase, at most MAXIT iterations of the cycling process are performed. On the basis of intermediate results the iteration can be stopped earlier; this necessitates the computation of a vectornorm (in VL2NOR).

In the following the structure of the cycling process of MGD1 is described in quasi-FORTRAN.

```

DO 10 k=l-1(-1)1
CALL RESTRICTION (f,f,k).           fk = rk+1fk+1
10 CONTINUE
C   START OF maxit MULTIGRID ITERATIONS
DO 50 n=1, maxit
IF (n.EQ.1) GO TO 30
CALL CTUMV (C,u,v)                 vl = Cl(ul-vl)
C   vl IS THE NEW RESIDUE fl-Alul
CALL RESTRICTION (f,v,l-1)         fl-1 = rlvl
DO 20 k=l-2(-1)1
CALL RESTRICTION (f,f,k)           fk = rk+1fk+1
20 CONTINUE

```

30 CALL <u>SOLVE</u> (u,f,1)	$u^1 = (L^1 U^1)^{-1} f^1$
DO 40 k=2 (1)l-1	
CALL <u>PROLONGATION</u> (u,u,k)	$u^k = p^k u^{k-1}$
CALL <u>CTUPF</u> (v,u,f,k)	$v^k = C^k u^k + f^k$
CALL <u>SOLVE</u> (u,v,k)	$u^k = (L^k U^k)^{-1} v^k$
40 CONTINUE	
CALL <u>PROLONGATION</u> (v,u,l)	$v^l = p^l u^{l-1}$
$v^l = v^l + u^l$	
CALL <u>CTUPF</u> (u,v,f,l)	$u^l = C^l v^l + f^l$
CALL <u>SOLVE</u> (u,u,l)	$u^l = (L^l U^l)^{-1} u^l$
CONTINUE	

In the actual implementation of MGD1, the matrix A_h is not kept in storage, but is overwritten by L and U. At minimal costs, the rest-matrix $C = LU - A_h$ is recomputed each time from L and U (in the subroutines CTUMV and CTUPF). All subroutines mentioned have their own particular features that make them more or less feasible for vectorization. This can be seen in table (8.2).

The structure of MGEOZ is more straightforward and follows directly from the MG-algorithm in section 3. Here the original matrix is not overwritten by the decomposition.

Two main alternatives exist for the implementation of ZEBRA relaxation. The lines in the grid can be relaxed successively and vectorization can be applied to speed up the solution of each tridiagonal system. However, because the solution of tridiagonal systems is not very suitable for vectorization, and all tridiagonal systems have the same size, we have chosen the other possibility to exploit vectorization, namely we solve all linear systems in each half relaxation sweep simultaneously by treating the even (odd) systems in parallel. Because of the rectangular shape of Ω the data-structure both in MGD1 and MGEOZ is simple. The grid-values of u_h and f_h are stored sequentially in one-dimensional arrays. They are ordered by grid and in each grid they are ordered by meshline. The diagonals of A_h are stored similarly in a two-dimensional array; the columns of the array corresponding to the diagonals of the matrix.

6. VERSIONS OF THE PROGRAM

To investigate the advantages of vectorization, different versions of the programs have been constructed. One version is written in portable FORTRAN and is tuned for execution on sequential hardware (MGD1S and MGEOZS). Another version, also written in portable FORTRAN was tuned for vectorization (MGD1V and MGEOZV). To keep the FORTRAN portable, we had to rely on the automatic vectorization capabilities of the compilers at hand. All loops for which vectorization was required could indeed be expressed in standard ANSI FORTRAN.

Other versions of the programs were considered as well. An interesting variant was (not portable) MGD1D. This version is the same as MGD1V except for two statements, containing a call to a STACKLIB routine for the recursive parts of the routine SOLVE in MGD1. The STACKLIB library, supplied for the CYBER 200 series, contains particularly efficient routines for vector operations that are not vectorizable because of recursion. For the comparison of MGD1D and MGD1V see table (8.2) and (8.3). For details about the implementation of the rectorized versions cf. [10].

7. THE TEST PROBLEM

In this study we are not interested in the *numerical* behaviour of the algorithms for different problems. We consider here only the efficiency of their implementation. Therefore, we may restrict ourselves to a single testproblem: the solution of Poisson's equation on the unit square with Dirichlet boundary conditions. Of course, for this problem other alternatives exist for its efficient solution. A program implementing one cycle in a Full Multigrid (FMG) algorithm, specially tuned for this problem on a CYBER 205, is described in [1]. Such a program may run much faster than our general purpose code. They report the solution on a 129×129 grid (with the usual 5-point discretization in 0.006 sec. However, we did *not* adapt our codes in any way to this particular problem. For our codes the cost of one iteration is the same for any 7-point discretization of a problem (2.1) on grids a given size.

In all cases reported here, the problem was solved on a mesh with

$$(2^{\ell+1}+1)^2$$

meshpoints. The length of most vectors in the program is $(2^{n+1}+1)^j$, $n = 1(1)\ell$. $j = 1$ and $j = 2$ (i.e. they represent lines or complete grids on level n). We performed experiments for $\ell = 4,5,6,7$. For problems with $\ell > 5$ the size of the problem was too large to run on the available CYBER 170; for $\ell > 6$ the problem was too large for the available CRAY 1 (Daresbury 1983).

8. THE EFFECT OF VECTORIZATION

In the tables (8.1)-(8.5) we give CPU-times for the programs mentioned in section 6. On CYBER 205 and CRAY 1 the vector-tuned versions ran with the vector option, the CPU-times mentioned for the scalar-tuned versions ran without. If we run the portable vector-code in scalar mode we sacrifice about 5% CPU-time on the CYBER 205 (CRAY 1: about 9%) when compared with the tuned scalar-code in scalar mode.

In the tables we give the total CPU-time in seconds spent in runs with 10 iteration cycles, including the preparational work. Also the time spent in the various subroutines is presented. From these numbers we derive the time spent in a single cycle. Additionally, we give the average convergence factor in the iterative cycling (CONV).

MGD1S	CYBER 170	CRAY 1		CYBER 205		
LEVELS	5	5	6	5	6	7
GRID	65 × 65	65 × 65	129 × 129	65 × 65	129 × 129	257 × 257
CONV	4.7E-2	4.7E-2	4.3E-2	4.7E-2	4.3E-2	4.3E-2
RAP	0.186	0.088	0.313	0.084	0.317	1.232
DECOMP	0.061	0.031	0.120	0.050	0.195	0.764
SOLVE	0.311	0.150	0.582	0.153	0.594	2.265
CTUMV	0.098	0.066	0.261	0.079	0.315	1.134
CTUPF	0.143	0.088	0.347	0.099	0.390	1.498
PROLON	0.041	0.030	0.113	0.024	0.093	0.360
RESTRI	0.066	0.017	0.062	0.014	0.054	0.202
VL2NOR	0.030	0.022	0.087	0.015	0.059	0.233
TOTAL	1.030	0.522	1.994	0.565	2.141	8.101
CYCLE	0.066	0.035	0.137	0.037	0.145	0.546

Table 8.1. CPU-times (in seconds) of the program MGD1S for problems with different sizes, run in scalar mode on different machines.

MGD1V	CRAY 1		CYBER 205		
	5	6	5	6	7
LEVEL	5	6	5	6	7
GRID	65 × 65	129 × 129	65 × 65	129 × 129	257 × 257
CONV	4.7E-2	4.3E-2	4.7E-2	4.3E-2	4.3E-2
RAP	0.033 (2.7)	0.085 (3.7)	0.022 (3.8)	0.053 (6.0)	0.151 (8.2)
DECOMP	0.010 (3.1)	0.037 (3.2)	0.012 (4.2)	0.043 (4.5)	0.162 (4.7)
SOLVE	0.086 (1.7)	0.324 (1.8)	0.091 (1.7)	0.325 (1.8)	1.251 (1.8)
CTUMV	0.008 (8.2)	0.032 (8.2)	0.003 (26!)	0.010 (31!)	0.042 (27!)
CTUPF	0.011 (8.0)	0.043 (8.1)	0.004 (25!)	0.014 (28!)	0.059 (25!)
PROLON	0.007 (4.3)	0.018 (6.3)	0.009 (2.7)	0.022 (4.2)	0.061 (5.9)
RESTRI	0.004 (4.2)	0.011 (5.6)	0.012 (1.2)	0.031 (1.7)	0.092 (2.2)
VL2NOR	0.003 (7.3)	0.010 (8.7)	0.001 (15)	0.004 (15)	0.015 (16)
TOTAL	0.169 (3.1)	0.575 (3.5)	0.177 (3.2)	0.533 (4.0)	1.882 (4.3)
CYCLE	0.012 (2.9)	0.043 (3.2)	0.012 (3.1)	0.040 (3.6)	0.151 (3.6)

Table 8.2. CPU-times (in seconds) of the program MGD1V run in vector-mode.

Between brackets: the acceleration factor by vectorization (compared with the tuned scalar version).

MGD1D	CYBER 205		
	65 × 65	129 × 129	257 × 257
SOLVE	0.094 (1.6)	0.263 (2.3)	0.831 (2.7)
TOTAL	0.181 (3.1)	0.469 (4.6)	1.442 (5.6)
CYCLE	0.012 (3.1)	0.034 (4.3)	0.108 (5.1)

Table 8.3. CPU-times in seconds of the program MGD1D run in vector-mode

MGEOZS	CYBER 170	CRAY 1		CYBER 205		
LEVEL	6	6	7	6	7	8
GRID	65 × 65	65 × 65	129 × 129	65 × 65	129 × 129	257 × 257
CONV	2.3E-1	2.3E-1	2.2E-1	2.3E-1	2.2E-1	2.02E-1
RAP	0.188	0.089	0.315	0.085	0.319	1.240
DECOMPZ	0.012	0.006	0.023	0.011	0.044	0.175
ZEBRA	0.333	0.135	0.511	0.148	0.585	2.309
RESIDU	0.106	0.049	0.191	0.045	0.180	0.733
PROLON	0.058	0.035	0.131	0.027	0.100	0.390
RESTRI	0.030	0.013	0.049	0.010	0.037	0.142
VL2NOR	0.018	0.013	0.048	0.008	0.033	0.128
TOTAL	0.797	0.348	1.286	0.353	1.333	5.216
CYCLE	0.053	0.023	0.088	0.023	0.090	0.357

Table 8.4. CPU-times (in seconds) of the program MGEOZS for problems with different sizes, run in scalar-mode on different machines.

MGEOZV	CRAY 1		CYBER 205		
LEVEL	6	7	6	7	8
GRID	65 × 65	129 × 129	65 × 65	129 × 129	257 × 257
CONV	2.3E-1	2.2E-1	2.3E-1	2.2E-1	2.0E-1
RAP	0.033 (2.7)	0.085 (3.7)	0.022 (3.9)	0.054 (5.9)	0.151 (8.2)
DECOMPZ	0.006 (1.0)	0.023 (1.0)	0.010 (1.1)	0.040 (1.1)	0.158 (1.1)
ZEBRA	0.034 (4.0)	0.103 (5.0)	0.084 (1.8)	0.210 (2.8)	0.623 (3.7)
RESIDU	0.010 (4.9)	0.034 (5.6)	0.007 (6.4)	0.020 (9.0)	0.067 (10.9)
PROLON	0.008 (4.4)	0.022 (6.0)	0.013 (2.1)	0.032 (3.1)	0.093 (4.2)
RESTRI	0.004 (3.2)	0.009 (5.4)	0.009 (1.1)	0.022 (1.7)	0.059 (2.4)
VL2NOR	0.003 (4.3)	0.009 (5.3)	0.002 (4.0)	0.004 (8.3)	0.012 (10.7)
TOTAL	0.102 (3.4)	0.293 (4.4)	0.162 (2.2)	0.400 (3.3)	1.190 (4.4)
CYCLE	0.006 (3.8)	0.017 (5.2)	0.011 (2.1)	0.028 (3.2)	0.084 (4.3)

Table 8.5. CPU-times (in seconds) of the program MGEOZV run in vector-mode.
Between brackets the acceleration by vectorization.

We see that certain parts of the algorithms benefit greatly from vectorization viz CTUMV and CTUPF (a factor 25-30 on CYBER 205, a factor 8-9 on CRAY 1). Other parts vectorize also well: VL2NOR, RAP, DECOMP, PROLON, RESIDU (on CRAY 1 also RESTRI and ZEBRA, in which vectoroperations with stride 2 occur). Other parts hardly benefit because of their recursive structure (DECOMPZ and SOLVE). If we give up portability, SOLVE can be speeded up on the CYBER 205 by use of the STACKLIB library.

9. CONCLUSIONS

Implementation of general-purpose multigrid solvers on vectorcomputers is feasible. Efficient programs in portable FORTRAN are now available for variable coefficient elliptic problems in a rectangular domain, discretized by a 7-point difference molecule. For the implementation implicit vectorization (auto-vectorization) can be used. The effect of vectorization (the

factor by which the program accelerates) depends strongly on the size of the problem and, of course, on the algorithm used.

Our implementation with ILU-relaxation vectorized well on the CYBER 205 (factor 3.2-4.3) but slightly worse on a CRAY 1 (factor 3.1-3.5). The implementation with ZEBRA relaxation vectorized better on a CRAY 1 (3.4-4.4) and less well on a CYBER 205 (factor 2.2-4.4). The reduced vectorizability of MGEOZ on the CYBER 205 is due to the frequent occurrence of vectors with stride unequal 1. For MGEOZV the CRAY1 is faster than the CYBER 205; for MGDIV the CYBER 205 is faster for large problems.

We notice that for the determination of the efficiency of an algorithm on a vectormachine the usual measure of complexity - the operations count - appears to be completely irrelevant. Many more aspects have to be taken into account such as: are the computations arranged in small or large do-loops, are they recursive, vectorizable, how are the data stored etc.

The relative efficiency of the various algorithms depends - of course - on the complexity of the algorithms and on the rate of convergence. The complexity of MGDEOZ is less, but generally MGD1 has a better convergence rate. Based on the present experiments we see that roughly one iteration with MGD1 taken twice the CPU-time of a MGEOZ iteration. On the other hand, in our Poisson testproblem, the empirical convergence rate of MGD1 is twice the rate of MGEOZ. (Which ratio is also found from theory [6]) In general, the relative efficiency of MGD1 and MGEOZ depends on the difference problem to solve, the size of the system of equations and on the machine used.

ACKNOWLEDGEMENT

We are indebted to mr. W. Lioen who constructed the different versions of MGEOZ.

NOTE

The codes discussed in the paper can be obtained by sending a tape to Mr. A. van Deursen, Dept. of Mathematics and Informations, Delft University of Technology, Julianalaan 132, 2628 Delft, The Netherlands.

REFERENCES

- [1] BARKAI, D. & A. BRANDT, *Vectorized Multigrid Poisson Solver for the CDC CYBER 205*, In: Procs. Int. MG. Conference, Copper Mountain, Colorado, April 6-8, 1983.
- [2] BRANDT, A., *Multi-level adaptive solutions to boundary-value problems*, Math. Comp. 31, 333-390, 1977.
- [3] HEMKER, P.W., *On the comparison of line-Gauss-Seidel and ILU relaxation in multigrid algorithms*, In: J.J.H. Miller (ed.), Computational and asymptotic methods for boundary and interior layers, pp. 269-277. Boole Press, Dublin, 1982.
- [4] HEMKER, P.W., *Multigrid methods for problems with a small parameter*, To appear in: Procs. Dundee. Conf. 1983, LNM, Springer-Verlag.
- [5] HEMKER, P.W., P. WESSELING & P.M. DE ZEEUW, *Multigrid methods: development of fast solvers*. In: Procs. Int. MG-Conference, Copper Mountain, Colorado, April 6-8, 1983.
- [6] KETTLER, R., *Analysis and comparison of relaxation schemes in robust multigrid and preconditioned conjugate gradient methods*. In: W. Hackbusch and U. Trottenberg (eds.), Multigrid methods. Proceedings, Köln-Porz, 1981. Lect. Notes in Math. 960, pp. 502-534, Springer-Verlag, Berlin etc., 1982.
- [7] STÜBEN, K. & U. TROTTEBERG, *Multigrid methods: fundamental algorithms, model problem analysis and applications*. In: W. Hackbusch and U. Trottenberg (eds.), Multigrid methods. Proceedings, Köln-Porz, 1981. Lect. Notes in Math. 960, pp. 1-176, Springer-Verlag, Berlin etc., 1982.
- [8] WESSELING, P., *Theoretical and practical aspects of a multigrid method*. SIAM J. Sci. Stat. Comp. 3, 387-407, 1982.
- [9] WESSELING, P., *A robust and efficient multigrid method*. In: W. Hackbusch and U. Trottenburg (eds.), Multigrid methods. Proceedings, Köln-Porz, 1981. Lect. Notes in Math. 960, pp. 614-630, Springer-Verlag, Berlin etc., 1982.

- [10] ZEEUW, DE P.M., W. LIOEN & P.W. HEMKER, *Vectorized Multigrid Codes*,
To appear as Mathematical Centre report, Amsterdam, 1983.