

THE INSTITUTE OF MATHEMATICS & ITS
APPLICATIONS CONFERENCE SERIES

New Series Number 3

MULTIGRID METHODS FOR INTEGRAL AND DIFFERENTIAL EQUATIONS

Edited by D. J. PADDON
and H. HOLSTEIN



Multigrid methods for integral and differential equations

Based on lectures at a Summer School/Workshop
held at Burwalls Conference Centre, University of Bristol, England,
September 1983

Edited by

D. J. PADDON
University of Bristol

and

H. HOLSTEIN
University College of Wales, Aberystwyth

CLARENDON PRESS

Oxford University Press, Walton Street, Oxford OX2 6DP
Oxford New York Toronto
Delhi Bombay Calcutta Madras Karachi
Kuala Lumpur Singapore Hong Kong Tokyo
Nairobi Dar es Salaam Cape Town
Melbourne Auckland
and associated companies in
Beirut Berlin Ibadan Nicosia

Oxford is a trade mark of Oxford University Press

Published in the United States
by Oxford University Press, New York

© *The Institute of Mathematics and its Applications, 1985*

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior permission of Oxford University Press

British Library Cataloguing in Publication Data
Multigrid methods for integral and differential equations : based on lectures at a summer school/ workshop held at Burwalls Conference Centre, University of Bristol, England, September 1983.— (The Institute of Mathematics and its Applications Conference series. New series; 3)
1. Equations—Numerical solutions
I. Paddon, D.J. II. Holstein, H. III. Institute of Mathematics and its Applications IV. Series
512.9'4 QA218
ISBN 0-19-853606-2

Library of Congress Cataloging in Publication Data
Main entry under title:
Multigrid methods for integral and differential equations. (The Institute of Mathematics and Its Applications conference series : new ser., 3)
Includes indexes.
1. Integral equations—Numerical solutions—Addresses, essays, lectures. 2. Differential equations—Numerical solutions—Addresses, essays, lectures. 3. Numerical grid generation (Numerical analysis)—Addresses, essays, lectures. I. Paddon, D. J. II. Holstein, H. (Horst) III. Institute of Mathematics and Its Applications. IV. Series.
QA431.M75 1985 515.3'5 85-11567
ISBN 0-19-853606-2

Printed in Great Britain by
St Edmundsbury Press Ltd.
Bury St Edmunds, Suffolk

PREFACE

These proceedings consist of papers whose authors addressed a small but enthusiastic group of numerical analysts on the Multigrid Method, during a Summer School/Workshop held at the Burwalls Conference Centre, University of Bristol, England, in September 1983. Most of the papers were substantially revised after the meeting, and bring together the authors' views and research up to the effective closing date of the proceedings (July 1984).

During the first part of the meeting, participants had the privilege of hearing the guest speakers, all of whom are leading authorities in this relatively new and exciting branch of numerical analysis: Achi Brandt, Wolfgang Hackbusch, Pieter Hemker, Ulrich Trottenberg and Pieter Wesseling. Each presented two papers on the theory and practice of the Multigrid Method. The second part of the meeting consisted of the presentation of contributed papers. Throughout the meeting, participants responded warmly to the excellent quality of the papers, and to the open and friendly discussions led by the guest speakers.

Multigrid Methods came into prominence in the late 1970's. The guest speakers have made major contributions to its development and its acceptance as a significant computational tool. Even though computers have shown orders of magnitude performance increases during the last decade, the demand for efficient, robust methods for solving field equations remains open. It is in this direction that the Multigrid Method has made notable advances, and has made accessible new areas of computation, e.g. three-dimensional simulation. In many cases, the Multigrid Method allows a realisation of optimal algorithms, whose computational work grows linearly with the number of unknowns in a system of equations. The Multigrid Method may itself be regarded as a model numerical process, capable of influencing future research into optimal algorithms.

These proceedings fall into two sections.

Section 1. This starts with a witty introduction by A. Brandt, from which we may gain insight into the deep-felt philosophical convictions held by a principal founder of the Multigrid Method. It is followed by the paper of W. Hackbusch, which gives a substantial review of the Multigrid Method for integral equations, together with a considerable innovative content. A very wide range of problems is shown to be amenable to the integral equation approach. P. Hemker's contribution gives a state-of-the-art guide to the design and implementation of programs for the Multigrid Method, applied to the solution of difference equations derived from general elliptic equations of the convection-diffusion type. Scalar and vector codes are considered. The paper of P. Sonneveld, P. Wesseling and P.M. de Zeeuw gives a most lucid treatment of conjugate gradient and Multigrid as methods for the acceleration of iterative convergence. The efficiency and robustness of the methods are compared for a variety of difficult test problems derived from the convection-diffusion equation.

Section 2. This section contains the contributed papers, and represents a spectrum of research ranging from new theory (J.F. Maitre and F. Musy; S. McCormick) to applications of Multigrid in engineering design (K.E. Barrett et al.) and in particular, to three-dimensional simulation (T. Scott). Hybrid techniques, in which Multigrid plays a key role, are described in papers on Spectral Methods (T.N. Phillips, T.A. Zang and M.Y. Hussaini) and the Hierarchical Finite Element Method (A.W. Craig and O.C. Zienkiewicz).

Both finite difference (B. Favini and G. Guj) and finite element approaches are exemplified in this section. An outstanding contribution is made by J. Ruge and K. Stüben to the understanding of the Algebraic Multigrid Method. Here the Multigrid approach is extended to apply to a fairly wide class of sparse matrix systems, irrespective of their origin from a finite difference or finite element discretisation, or indeed from any geometric grid structure at all. While there are computational overheads compared to the most efficient geometric multigrid solvers when applicable, the authors exhibit remarkably successful results on a variety of problems, using a truly "pitch black box" Algebraic Multigrid routine.

The Summer School/Workshop and subsequent communication with the authors has given the editors a great deal of pleasure. We warmly express our gratitude to the authors for their cooperation and friendship while producing these proceedings, and for the opportunity to present for publication this collection of stimulating papers.

We would like to thank Roger Moses and the staff of the Burwalls Conference Centre for creating an atmosphere so conducive to the meeting and discussions that took place. We would also like to thank Catherine Richards and her staff at the Institute of Mathematics and its Applications for the careful typing and efficient yet cheerful way in which they coordinated the production of these proceedings. We are grateful to the staff of the OUP for publishing this volume, and in particular to Anthony Watkinson, for the help, encouragement and patience he has shown throughout its preparation.

Derek Paddon

Horst Holstein

ACKNOWLEDGEMENTS

The Institute thanks the authors of the papers, the editors, Dr. D. Paddon, AFIMA (University of Bristol) and Dr. H. Holstein, AFIMA (University College, Wales) and also Mrs. Janet Parsons, Miss Denise Wright, Miss Pamela Irving and Miss Karen Jenkins for typing the papers.

LIST OF CONTRIBUTORS

- K.E. Barrett, Department of Mathematics, Coventry (Lanchester)
Polytechnic, Coventry CV1 5FB, UK
- A. Brandt, Department of Applied Mathematics, The Weizmann Institute
of Science, Rehovot, Israel
- D.M. Butterfield, Department of Mathematics, Coventry (Lanchester)
Polytechnic, Coventry CV1 5FB, UK
- A.W. Craig, Civil Engineering Department, University College of
Swansea, Singleton Park, Swansea SA2 8PP, UK
- S.E. Ellis, Department of Mathematics, Coventry (Lanchester)
Polytechnic, Coventry CV1 5FB, UK
- B. Favini, Dipartimento di Meccanica e Aeronautica, Università di
Roma, 'La Sapienza', Italy
- G. Guj, Dipartimento di Meccanica e Aeronautica, Università di Roma,
'La Sapienza', Italy
- W. Hackbusch, Institute für Informatik und Praktische Mathematik,
Christian-Albrechts-Universität Kiel, Olshausenstr. 40, D-2300,
Kiel 1, Germany
- P.W. Hemker, Centre for Mathematics and Computer Science, Department
of Numerical Mathematics, Kruislaan 413, 1098 SJ Amsterdam,
The Netherlands
- M.Y. Hussaini, Institute for Computer Applications in Science and
Engineering, Mail Stop 132C, NASA Langley Research Centre,
Hampton, Virginia 23665, USA
- C.J. Judd, Department of Mathematics, Coventry (Lanchester)
Polytechnic, Coventry CV1 5FB, UK
- J.F. Maitre, Département de Mathématiques Informatique Systèmes,
Ecole Centrale de Lyon, B.P. 163-69131, Ecully, Cedex, France
- S. McCormick, Department of Mathematics, Colorado State University,
Fort Collins, Colorado 80523, USA
- F. Musy, Département de Mathématiques Informatique Systèmes, Ecole
Centrale de Lyon, B.P. 163-69131, Ecully, Cedex, France
- T.N. Phillips, Institute for Computer Applications in Science and
Engineering, Mail Stop 132C, NASA Langley Research Center,
Hampton, Virginia 23665, USA

- J. Ruge, Gesellschaft für Mathematik und Datenverarbeitung, Postfach 1240, D-5205 St. Augustin 1, West Germany*
- T. Scott, UKAEA, Winfrith, Dorchester, Dorset DT2 8DH, UK*
- P. Sonneveld, Department of Mathematics and Informatics, Delft University of Technology, Julianalaan 132, 2628 BL Delft, The Netherlands*
- K. Stüben, Gesellschaft für Mathematik und Datenverarbeitung, Postfach 1240, D-5205 St. Augustin 1, West Germany*
- J.H. Tabor, Department of Mathematics, Coventry (Lanchester) Polytechnic, Coventry CV1 5FB, UK*
- P. Wesseling, Department of Mathematics and Informatics, Delft University of Technology, Julianalaan 132, 2628 BL Delft, The Netherlands*
- T.A. Zang, NASA Langley Research Center, Mail Stop 132C, Hampton, Virginia 23665, USA*
- P.M. de Zeeuw, Centre for Mathematics and Computer Science, Department of Numerical Mathematics, Kruislaan 413, 1098 SJ. Amsterdam, The Netherlands*
- O.C. Zienkiewicz, Civil Engineering Department, University College of Swansea, Singleton Park, Swansea SA2 8PP, UK*

CONTENTS

List of contributors	xi
1. Introduction - Levels and Scales by A. Brandt	1
2. Multigrid Methods of the Second Kind by W. Hackbusch	11
3. Some Implementations of Multigrid Linear System Solvers by P.W. Hemker and P.M. de Zeeuw	85
4. Multigrid and Conjugate Gradient Methods as Convergence Acceleration Techniques by P. Sonneveld and P. Wesseling	117
5. Efficient Solution of Finite Difference and Finite Element Equations by J. Ruge and K. Stüben	169
6. Algebraic Formalisation of the Multigrid Method in the Symmetric and Positive Definite Case - A Convergence Estimation for the V-Cycle by J.F. Maitre and F. Musy	213
7. A Variational Theory for Multi-Level Adaptive Techniques (MLAT) by S. McCormick	225
8. Spectral Multigrid Methods for Dirichlet Problems by T.N. Phillips, T.A. Zang and M.Y. Hussaini	231
9. MG Techniques for Staggered Differences by B. Favini and G. Guj	253
10. Multigrid Analysis of Linear Elastic Stress Problems by K.E. Barrett, D.M. Butterfield, S.E. Ellis, C.J. Judd and J.H. Tabor	263
11. Multigrid Methods for Oil Reservoir Simulation in Three Dimensions by T. Scott	283
12. A Multigrid Algorithm using a Hierarchical Finite Element Basis by A.W. Craig and O.C. Zienkiewicz	301
Subject Index	313
Author Index	321

INTRODUCTION - LEVELS AND SCALES

A. Brandt

*(Department of Applied Mathematics,
The Weizmann Institute of Science, Israel)*

Open your eyes and you see multi-level processes all around. They have always been here. The organization and operation of military forces is an obvious example: soldiers are grouped in squads, which in turn are grouped in sections, grouped in platoons, then in companies, battalions, regiments, brigades, divisions, corps and armies. Civil society, less strict and more complicated, still operates in a variety of hierarchical structures: geographical, economical, political, judicial, educational, and so forth.

Such hierarchies are necessary because there are very many, sometimes millions, interdependent decisions to be made; they cannot be made by one governor, because of their multitude and their complicated interdependence, nor can they be decided by many independent administrators without being coordinated with each other. The hierarchical structure effectively deals with this situation by exploiting the fact that each decision has a certain "scale". The location of a new hospital, for example, is a decision which has the scale of a district. It strongly affects the district served by it, marginally affects neighbouring districts, and very weakly affects others. The decision can therefore be made by a certain administration assigned to the district, in some coordination with neighbouring administrations. The decision cannot effectively be negotiated at a too low administrative level, e.g., at the level of individual families living in the district; the relevant information (concerning needs, constraints, etc.) must be gathered into one point of decision. Similarly, the general outline of principal thoroughways should be decided at the inter-state level, while local back roads should be regulated at the district or town levels.

The two kinds of roads should be connected, of course. The effective way to manage the inter-dependence between the different levels is based on the assumption that global decisions should only marginally be affected by local ones (otherwise the latter would not be local). Thus, ideally, a two-level hierarchical structure should operate in the following way. The global government first gathers some general figures summarized at the local level, representing sum totals of local needs, important overall constraints, etc. Based on these it prepares preliminary global plans. These global plans give the local governments the framework for devising their own, more detailed, plans. In course of doing that, the local government may realize that some, usually marginal, aspects of the global plans do not quite fit the local situation and, therefore, need some adjustments or corrections. So, at a second round, the global government again gathers information summarized at the local level, now representing sum totals of needed corrections ("defect corrections", in the language of numerical analysts). Since in practice this process is seldom fully recognized, let alone fully effectively organized, more such rounds may be needed. When more levels of government are involved, the process is applied recursively, in a variety of manners.

For very much similar reasons, iterative hierarchical procedures, similar to the process we have just described, are very efficient in solving large and complicated problems on computers. Such procedures were naturally introduced to solve problems where the hierarchical structure was already explicit in the problem itself. A good example is the field of production planning, notably in the Soviet Union, where hierarchic divisions into sectors and pyramidal management naturally led to the introduction of iterative "aggregation/disaggregation" (a/d) algorithms, starting in the mid sixties (Dudkin and Yershov (1965)) and growing in the seventies into extensive Russian literature on iterative a/d procedures for large linear programming problems (see Vakhutinsky, Dudkin and Ryvkin (1979)). Multi-level approaches have in fact quite naturally emerged in all branches of computer technology, as in the structured organization of computer hardware (see for example Stone (1972)), the top-down structured design of software (Yourdon, Swann (1978)), the pyramidal data structures (trees, heaps, etc.) and many of the most efficient algorithms in computer science, such as fast sorting (sorting n numbers in $O(n \log n)$ operations) the "divide and conquer" class of algorithms, etc. (See, e.g., Aho, Hopcroft and Ullmann (1974)). Most of these procedures are not iterative, though.)

Also, for very much the same reasons, multi-level algorithms have come forth as the most efficient algorithms in solving the very large algebraic systems arising in discretising partial differential boundary value problems, especially those describing steady-state physical systems. Here the fully hierarchical structure is not at all generally explicit in the problem itself, so it took some effort, and interesting historical development, to realise it. Fully effective multi-level algorithms were first developed as *direct*, not iterative, solvers, treating very special situations where it was algebraically clear enough how to construct recursively hierarchical solvers. I refer here to the fast solvers based on fast Fourier transforms (FFT) and/or reduction methods, especially the cyclic odd-even reduction, both of which are clearly recursive, but non-iterative, multi-level processes (see Buzbee, Golub and Nielson (1970), Hockney (1970) and Temperton (1979)). The total reduction method also belongs to this class (see Trottenberg (1977)). The solution of n equations by this kind of solvers require at most $O(n \log n)$ computer time and storage, but the class of problems for which this full efficiency is attained is quite limited: separable problems, essentially meaning constant-coefficient elliptic equations with constant-coefficient boundary conditions on rectangular domains. This class has been substantially enlarged by using these fast solvers, and various combinations of them, *iteratively*. Thus, for example, if the coefficients are not constant but sufficiently uniform, the iterative application of a constant-coefficient fast solver can be very effective (the number of iterations depending on the uniformity of the coefficients, but not on the meshsize). Extensions to domains of arbitrary shapes have been obtained by "capacitance matrix methods", using the fast direct solvers iteratively, with conjugate-gradient acceleration, each solution typically costing the equivalent of some 15 applications of the fast solver (see e.g. Proskurowski (1981)). The nested-dissection approach to elimination ordering (George and Lin (1981)) is another powerful multi-level approach, more general but less efficient than the FFT and reduction solvers.

Meanwhile, steady-state PDE problems and the solution of their discretised equations were examined from two other points of view, which jointly led to the realisation that each such problem contains a natural hierarchy of levels, not immediately explicit, but very powerful and much more general than the mechanical hierarchy exploited by the above methods. First, studying reasons for slow convergence of various relaxation solvers, it became clear that relaxation is a "local" process which cannot efficiently treat "global" or "smooth" solution components. A smooth error component shows relatively small "residuals", i.e., small errors in the individual difference equations. The smooth error can be much larger than shown by the residuals because, throughout a large region, these residuals have the same sign, so they reinforce each other. Now, since the relaxation corrections are based on the individual residuals, they are necessarily small compared with the actual error, if the latter is smooth. Thus, in order to get a correction comparable to a smooth error, information concerning the residuals throughout sufficiently large regions must be summed up to one point of decision, very much as in the case of global social decisions mentioned above. The size of the regions over which residuals should be summed up - so that a correct picture about the error magnitude is obtained - must be comparable to the scale of the error, i.e., to the typical distance over which the error substantially changes. As long as each relaxation step (each step of correction) works on a much smaller scale, convergence must be slow.

A second, complementary viewpoint evolved from examining the nature of the discretisation error, i.e., the difference between the true solution of the differential problem and the exact solution of the discretised equations. The relative magnitude of this error is clearly determined by the relative magnitudes of the discretisation meshsize and the solution scale. A smooth solution, which is a large-scale solution, can thus be approximated on a coarse grid. The same is obviously also true for a smooth error. Thus, it became clear, exactly those errors that are slow to converge by relaxation processes on some fine grid can be approximated on a coarser grid, where the meshsize is comparable to their scale and hence their convergence need not be slow.

A natural hierarchy of levels emerges, based on viewing the solution to each boundary value problem as a linear combination of components with different scales. Each component is most effectively controlled by grids with meshsize comparable to its scale, and efficient multi-level control can thus be realised as a multi-grid processing.

A two-grid process, for example, is fully analog to the ideal operation of the two-level hierarchical government described above. The problem is first represented on the coarser grid, e.g., by averaging its equations to the scale of that grid. The (approximate) solution to the resulting coarse-grid problem, once computed, is then interpolated to the fine grid, serving there as a first approximation, a framework, to be next improved by fine-grid processes, such as relaxation. This fine-grid processing finds the fine features of the solution which were invisible to the coarser grid, and also, as a result, encounters some residuals of global (smooth) errors, which it cannot efficiently reduce. (These are smooth errors caused by aliasing, i.e., by the previous coarse-grid processing having misinterpreted coarse-grid traces of the fine features. Now that those fine features have been removed from the error by the fine processing, that aliasing error becomes the dominant one.) So, in the next round, the residual problem is approximately transferred, by some

averaging, to the coarse grid, where it can efficiently be solved and its solution is then interpolated back to the fine grid and added as a correction to the previous fine-grid solution.

The process just described is the two-level "full multigrid" (FMG) algorithm. It can be used recursively in a variety of manners in case more levels are involved.

The number of levels that should be used depends on the ratio between the size of the domain and the finest scale one wants to see in the solution. Between these two scales as many should be introduced as practical. Namely, the ratio between successive scales (successive meshsizes) should be as small as possible, as long as this does not substantially increase the total number of gridpoints. The ratio 1:2 between successive meshsizes is very convenient: the total number of gridpoints is still dominated by their given number on the finest grid, but successive meshsizes are close enough to effectively treat any solution component. In fact, with such a ratio, the experience so far showed that suitable FMG algorithms could solve all test problems "to the level of truncation error" (i.e., to the point where the error in approximating the differential solution is dominated by the discretisation error, not by the error in solving the discrete system) in just few (less than 10) "work units", where the work unit is the amount of operations involved in *expressing* the given (finest-grid) system of discrete equations. The 1:2 ratio is also most convenient in programming the inter-grid transfers. Note that in this respect the multigrid processing is different from multi-level social structures. Its levels are chosen much more tightly, to achieve maximum efficiency.

Note also that the decomposition of the PDE solution into components of different scales is only implicit; it is used above to motivate and explain the validity and strength of the multigrid process; but the actual multigrid algorithm does not use any such decomposition. It only transfers equations (or residual equations) from fine grids to coarse grids, and solutions (or corrections) from coarse to fine. Decomposition in terms of Fourier components, in particular, can be used as a powerful tool to *analyse*, and even exactly predict, the performance of multigrid algorithms, but the algorithms themselves do not employ such decompositions, and their efficiency extends far beyond the cases where the Fourier analysis is rigorously valid. (Incidentally, it is important to realize that in some cases Fourier analysis is the wrong tool to separate local effects from global ones. For example: a local discontinuity gives rise to global high-frequency Fourier components. Other tools should then be used to understand quantitatively, and optimise, the multigrid performance.)

The FMG solvers (solving discretised PDEs to the level of truncation errors) have been developed to the point that they are today even faster than the FFT and reduction solvers mentioned before. (The multigrid Poisson solver described in Barkai and Brandt (1983) is the fastest we know.) More importantly, of course, these FMG solvers are much more general. They solve with the same efficiency (i.e., in few work units) complicated nonlinear systems on general domains. Moreover, it is possible to integrate into each application of an FMG solver, for small extra computer work, various processes of local mesh refinement, mesh optimisation and local coordinate transformations, making it very effectual in treating singularities, unbounded domains, curved

boundaries, boundary layers, discontinuities, etc.

Furthermore, multigrid solvers can be directly applied to "higher" problems - such as optimisation, optimal design and optimal control problems, or system identification problems - whose solution would normally be accomplished through solving a sequence of boundary value problems. An important principle indeed is always to try to multigrid the given, *original* problem, instead of merely using fast multigrid solvers to a sequence of intermediate subproblems. The *original* problem (e.g, the optimisation problem itself) should first be solved on a coarser grid, then relaxed on the finer grid (including for example, local optimisation of parameters, in case some of the functions to be optimised do have local scales), then brought back to the coarser grid, etc.. The entire solution of the *original* problem may thus cost only few work units. Multigridding the original problem is especially advantageous in case that problem is autonomous (i.e., having solution-dependent but not directly space-dependent coefficients, as for example most fluid dynamic problems), while the subproblems are nonautonomous (having coefficients spatially depending on the solution of a previous subproblem).

Sometimes it is still required to solve a sequence of subproblems. A typical example is the interactive design of a certain structure, where between two solutions of the system the structure is changed in some specific parts and/or in some of its global parameters. The new solution can then be obtained from the old one by a remarkably short multigrid processing, in which the finer grids are relaxed only around the changed parts. This technique may allow the design of a large structure to be done mostly in core memory, since for several design steps only the currently designed parts, and some neighbouring parts, should fully be kept in memory while the rest of the structure may be represented by coarser grids (using the FAS version of multigrid). Re-solving by such techniques should be so efficient as to allow the designer to introduce some changes to a large structure and immediately view the new solution on the screen.

Similarly, in evolution problems with implicit time differencing, the solution of a new system of equations seems to be required at each time step. Each of these systems could be solved very efficiently by an FMG algorithm, costing the equivalent of just few explicit time steps. But, here again, multigrid techniques may be even more effective if they are designed in terms of the original evolution problem itself. For example, one can often drastically reduce the overall work by exploiting the fact that the solution is a superposition of pure convection and smooth changes, or changes which are smooth throughout most of the domain, hence requiring fine-grid interactions only at some small subdomains.

The full efficiency of multigrid solvers, as described here, is not easy to obtain, though. It sometimes depends on the correct treatment of each feature of the problem at each multigrid stage. Many things can easily go wrong, such as relaxing a certain boundary condition in a way which conflicts with the interior smoothing; or improper fine-to-coarse transfers of boundary conditions and their residuals; or treating at relaxation conditions which seem local but are not; or using a relaxation scheme which is not as powerful a smoother as it should, and can, be; or wrong order or type of interpolations; or any inadequate treatment of any

difficulty, from structural singularities (e.g., reentrant corners) to discontinuities in the solution or in the equations, anisotropies, non-ellipticity, etc., etc. A single mistake at any of these may substantially degrade the whole performance, not to mention plain programming bugs, which, due to the corrective nature of the algorithm, may well disguise themselves in the innocent form of slow convergence. To obtain full efficiency it is therefore necessary to construct the algorithmic concepts and the actual programs in a gradual, systematic way, using available knowhow (see Hackbusch and Trottenberg (1982) , Brandt (1984)).

The class of partial differential equations that can be solved by multigrid solvers has been ever extending, from second order equations to arbitrary orders, from linear to nonlinear, from smooth coefficients to strongly discontinuous ones (Alcouffe, Brandt, Dendy and Painter (1981)), from definite to indefinite problems, from scalar equations to general systems (Brandt (1984) §3.8), and from elliptic type to other types.

Unlike evolution problems, where properties like hyperbolicity and parabolicity are all important, the only feature that matters concerning the type of the differential operator in (stationary) boundary value problems is whether it is nicely (isotropically) elliptic or not. If it does not have a good ellipticity measure, it does not make any difference whether it is anisotropic elliptic (like $\epsilon \partial^2/\partial x^2 + \partial^2/\partial y^2$) or semi-elliptic, weakly elliptic (e.g., having elliptic singular perturbation), hyperbolic or any other nonelliptic type. All these types have the same basic difficulty and can be treated essentially by the same approach. The difficulty is that the solution scales discussed above are *not isotropic*. (Incidentally, one dimensional problems, such as $du/dx = f$, are elliptic, unlike their two-dimensional counterparts. Their use as models for treating non-ellipticity is thus erroneous.) In other words, at each point in the domain of the problem there passes a "characteristic line" (sometimes just a characteristic surface; and in case of non-scalar PDEs there may pass several such lines or surfaces at each point, corresponding to different solution components) such that a "global" error component (in the sense that its residuals are small compared with its size) can change rapidly in directions perpendicular to the characteristic lines. Such global but rapidly changing components are called *characteristic components*. As global components, their convergence by relaxation is inefficient, but as rapidly changing components they cannot be well approximated by coarser grids.

The key for efficient multigrid treatment of such anisotropic problems is to distinguish clearly between two very different situations, depending on whether or not one wishes to approximate characteristic components far from boundaries, where "far" is meant relative to the component's smaller scale, i.e., the scale of its rapid change perpendicular to the characteristic. To so approximate characteristic components by the discrete system, the grid must (roughly) be aligned with the characteristic lines throughout the domain. Fast multigrid solvers can then be based on these aligned gridlines, either by relaxing these gridlines simultaneously, or by coarsening the grid only along these lines and not in the perpendicular directions. (The latter approach is preferred when the characteristics are surfaces rather than lines.) If, on the other hand, the grid is not consistently aligned with the characteristics, then characteristic components cannot be

approximated far from boundaries (except accidentally, in some regions of accidental alignment), and it is then unwise to attempt to have fast algebraic convergence for such components. This convergence, which is meaningless in terms of approximating the *differential* solution, is harder to obtain: costlier and more complicated relaxation schemes and/or intergrid transfers must be developed, and even then fast convergence is not always guaranteed. It is absurd, we believe, to invest most of your computer resources and programming effort to get fast algebraic convergence exactly for those components whose algebraic solution is generally (out-side accidental regions of alignment) no approximation to the differential solution.

This fashion in treating anisotropy, as well as some other developments, has led to the recognition that fast algebraic convergence should not be the main objective of multigrid solvers. The objective is of course to get the desired accuracy, in terms of solving the *differential* problem, for minimal computer (and also human) resources. This is obtained by FMG solvers which, especially in cases of anisotropic equations without corresponding grid alignment, do not necessarily employ uniformly good smoothers, hence nor do they attain uniformly good algebraic convergence rates. Working with such solvers requires of course certain modifications in the traditional "smoothing factor" approach for measuring the effectiveness of relaxation (see Brandt (1984), §20.3.1), as well as new approaches for a priori predicting, and a posteriori judging, the overall success of the FMG solver, (see Brandt (1984), §7.4, §7.5 and §1.6). An important advantage of these approaches for performance evaluation is that the performance becomes less sensitive to the precise treatment of all problem features at all algorithmic stages.

These approaches also allow the evaluation of important schemes which deliberately avoid any algebraic convergence. Such for example is the "double discretization" scheme (Brandt (1984), §10.2), which employs less accurate, stable discretizations in relaxation processes, and others, more accurate but not necessarily stable ones, in the fine-to-coarse transfers of residuals, thus combining easier local stability with higher global accuracy.

Another interesting development has recently led to the extension of multigrid-like techniques to cases where no grid is actually present. It started, in a way, with the development of usual multigrid algorithms for diffusion problems which were isotropically elliptic but in which the diffusion coefficients were strongly discontinuous. It turned out (Alcouffe, Brandt, Dendy and Painter (1981)) that to obtain the usual multigrid efficiency in such cases, the coarse-to-fine interpolation of corrections should be based on the difference equations themselves, rather than being the standard polynomial interpolation. This later led to the recognition that interpolation can be based solely on the algebraic equations, without even using the geometry of the grid. In a similar manner, the fine-to-coarse transfer of residual equations can completely be based on the *adjoint* (transposed) algebraic system. Since these intergrid transfers are what give the coarse grid a definite meaning, it was further realized that even the choice of the coarse-grid *variables* can be freed from its traditional geometric context, and purely "algebraic multigrid" (AMG) schemes were introduced (Brandt (1983), Brandt (1984), §13.1, Brandt, McCormick and Ruge (1984), Ruge and Stüben (1984)). In these schemes the selection of coarser levels is based on the principle that each variable of any level should have a sufficiently strong "total

algebraic connection" to variables of the next-coarser level. The entire processing is thus made in terms of the given (sparse) algebraic system of equations, with no reference to their geometric origin. AMG algorithms can thus be used as very efficient "black box" solvers for important classes of matrix equations. (For some other classes the current AMG solvers are not suitable.) The typical multigrid efficiency is obtained by AMG even for cases where it would be very difficult to construct conventional (geometric) multigrid algorithms, such as cases of finite element discretisation on arbitrary, irregular triangulations, or even cases where topologically rectangular grids are used, but with highly and non-uniformly stretched coordinates (Lagrangian discretizations in particular) or with peculiarly distributed physical coefficients. In addition, AMG solvers can be applied to many large algebraic systems which are not at all derived from continuous problems, such as the geodetic problem treated in Brandt, McCormick and Ruge (1983) .

The scope of multi-level computations has thus been extended very much. To state it most broadly, consider any matrix equation $Ax = b$. (That the system is linear is not really essential; but it simplifies the following statements.) Denote by \tilde{x} the evolving computed approximation to x , and define the error vector $e = x - \tilde{x}$ and the *normalized* residual vector $r = (r_i) = (a_i e / \|a_i\|)$, where a_i is the i -th row of A and $\|\cdot\|$ is the ℓ_2 norm. For a suitable relaxation scheme it can be shown (Brandt (1983), Theorem 3.4) that the decrease in $\|e\|$ per sweep can be slow only when $\|r\| \ll \|e\|$. Since r is properly normalized, for most error components $\|r\|$ is comparable to $\|e\|$. Hence, convergence can be slow only for *special* types of error components. Slowly converging errors can therefore be approximated by far fewer parameters, that is, by a much smaller algebraic system — a coarser level. To exploit this fact one of course needs some characterisation for those vectors e for which $\|r\| \ll \|e\|$. In case the matrix A approximates a differential operator L , those vectors e approximate functions v for which $\|Lv\| \ll \|L\| \|v\|$. This usually implies smoothness of v , or, when L is anisotropic, at least smoothness in characteristic directions. In other cases other characterisations can be found, so the general rule which emerges is that slow convergence should always be avoidable.

This corresponds, more or less, to the "golden rule of computational physics", which states that the amount of computational work should be proportional to the amount of real physical changes in the computed system. Stalling numerical processes must be wrong. Indeed, multi-level processing is a general vehicle to effect this rule. So, whenever you have stalling computations — either in the form of slowly converging iterative procedures, or in the form of computational grids, in space and/or in time, which almost everywhere tend to excessively over-resolve the scales of real physical changes — try to think in terms of multi-level techniques.

Multi-level methods are now in the process of being introduced into a variety of new fields, including various systems of tomography, image processing and pattern recognition; statistical physics; queueing theory; network simulation and design; geodesy; multivariate interpolation; large transportation problems and linear programming. The aggregation/disaggregation methods developed earlier for linear programming can now be improved by an AMG-type approach, because the latter provides a more mathematical basis for defining the levels, hence "tighter" hierarchies, exploiting implicit levels not necessarily recognized by the real-life